



ELSEVIER

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

Technical Notes

The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events

J. Alme^a, Y. Andres^d, H. Appelshäuser^g, S. Bablok^a, N. Bialas^a, R. Bolgen^a, U. Bonnes^f, R. Bramm^h, P. Braun-Munzinger^{f,h,m,n}, R. Campagnolo^d, P. Christiansen^j, A. Dobrin^j, C. Engster^d, D. Fehlker^a, Y. Foka^{h,*}, U. Frankenfeld^h, J.J. Gaardhøje^e, C. Garabatos^h, P. Glässelⁱ, C. Gonzalez Gutierrez^d, P. Gros^j, H.-A. Gustafsson^j, H. Helstrup^b, M. Hoch^d, M. Ivanov^h, R. Janik^c, A. Junique^d, A. Kalweit^f, R. Keidel^o, S. Kniege^g, M. Kowalski^k, D.T. Larsen^a, Y. Lesenechal^d, P. Lenoir^d, N. Lindegaard^e, C. Lippmann^d, M. Mager^d, M. Mast^d, A. Matyja^k, M. Munkejord^a, L. Musa^d, B.S. Nielsen^e, V. Nikolic^l, H. Oeschler^f, E.K. Olsen^e, A. Oskarsson^j, L. Osterman^j, M. Pikna^c, A. Rehman^d, G. Renault^e, R. Renfordt^g, S. Rossegger^d, D. Röhrich^a, K. Røed^b, M. Richter^a, G. Rueshmann^g, A. Rybicki^k, H. Sann^{h,1}, H.-R. Schmidt^h, M. Siska^c, B. Sitár^c, C. Soegaard^e, H.-K. Soltveitⁱ, D. Soyk^h, J. Stachelⁱ, H. Stelzer^h, E. Stenlund^j, R. Stock^g, P. Strmeň^c, I. Szarka^c, K. Ullaland^a, D. Vranic^h, R. Veenhof^d, J. Westergaard^e, J. Wiechulaⁱ, B. Windelbandⁱ

^a Department of Physics, University of Bergen, Bergen, Norway

^b Faculty of Engineering, Bergen University College, Bergen, Norway

^c Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovakia

^d European Organization for Nuclear Research (CERN), Geneva

^e Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark

^f Institut für Kernphysik, Technische Universität Darmstadt, Darmstadt, Germany

^g Institut für Kernphysik, Johann-Wolfgang-Goethe Universität Frankfurt, Frankfurt, Germany

^h GSI Helmholtzzentrum für Schwerionenforschung GmbH, Darmstadt, Germany

ⁱ Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany

^j Division of Experimental High Energy Physics, University of Lund, Lund, Sweden

^k The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, Cracow, Poland

^l Rudjer Bošković Institute, Zagreb, Croatia

^m ExtreMe Matter Institute, EMMI, GSI, Darmstadt, Germany

ⁿ Frankfurt Institute for Advanced Studies, J.W. Goethe University, Frankfurt, Germany

^o Zentrum für Technologietransfer und Telekommunikation (ZTT), Fachhochschule Worms, Worms, Germany

Acronyms: AC, alternating current; ACORDE, ALICE COsmic ray DETector; ADC, analog to digital converter; ALEPH, apparatus for LEP PPhysics; ALICE, a large ion collider experiment; ALTRO, ALICE TPC ReadOut chip; BC, board controller; CAD, computer aided design; CAN, controller area network; CCD, charge-coupled device; CERN, Conseil Européen pour la Recherche Nucléaire (European Organization for Nuclear Research); CFD, computational fluid dynamics; CMOS, complementary metal-oxide-semiconductor; CSA, charge sensitive amplifier; CTP, central trigger processor; CU, control unit; DAC, digital to analog converter; DAQ, data acquisition system; DC, direct current; DCS, detector control system; DDL, detector data link; DIM, distributed information management system; DNL, differential non-linearity; D-RORC, DAQ RORC data ReadOut receiver card; DU, device unit; ECS, experiment control system; EEPROM, electrically erasable programmable read only memory; ELMB, embedded local monitor board; EMCAL, ElectroMagnetic CALorimeter; ENC, equivalent noise charge; ENOB, equivalent number of bits; FEC, front-end card; FEE, front-end electronics; FEM, finite element method; FET, field-effect transistor; FMD, forward multiplicity detector; FPGA, field programmable gate array; FSM, finite state machine; FWHM, full width half maximum; GEM, gas electron multipliers; GTL, gunning transistor logic (FEE-bus technology); HCMOS, high-speed CMOS; HF, high frequency; HLT, high level trigger; HMPID, high momentum particle identification detector; HV, high voltage; INL, integral non-linearity; IROC, inner ReadOut chamber; ITS, inner tracking system; L0, level 0 trigger; L1, level 1 trigger; L2(a,r), level 2 trigger (accept, reject); L3, magnet used by LEP-L3 experiment; LAN, local area network; LDC, local data concentrator; LEP, large electron positron collider; LHC, large hadron collider; LSB, least significant bit; LV, low voltage; LVCMOS, low-voltage CMOS; MEB, multiple event buffer; MSPS, mega-samples per second; MTBF, mean time between failures; MWPC, multiwire proportional chamber; NMOS, N-type metal-oxide-semiconductor field effect transistors; NTP, normal temperature and pressure; OCDB, offline conditions data base; OLE, object linking and embedding; OPC, OLE for process control; OROC, outer ReadOut chamber; PASA, PreAmplifier ShAper; PCB, printed circuit board; PEEK, polyarylether-etherketone; PHOS, PHOTon spectrometer; PID, proportional-integral-derivative; PLC, programmable logic controller; PMD, photon multiplicity detector; PMEM, pedestal MEMORY; PMOS, P-type metal-oxide-semiconductor field effect transistors; PS, power supply; PVSS, prozessvisualisierungs- und steuerungs- system; RCC, ring cathode chamber; RCU, readout control unit; RHIC, relativistic heavy ion collider; RICH, ring imaging Cherenkov; RMS, root mean square; ROC, ReadOut chamber; SCADA, supervisory controls and data acquisition; SEL, single event latchup; SEU, single event upset; SFDR, spurious-free dynamic range; SIU, system interface unit; S/N, signal-to-noise ratio; SRAM, static random access memory; SSW, service support wheel; STAR, solenoidal tracker At RHIC; TCF, tail cancellation filter; TCP/IP, transmission control protocol/internet protocol; TDR, technical design report; TID, total ionization dose; TOF, time-of-flight detector; TP, technical proposal; TPC, time projection chamber; TQFP, thin quad flat pack (chip package); TRD, transition radiation detector; TRG, TRIGger system; TTC, timing, trigger and control; UPS, uninterruptible power supply; UV, ultra violet; ZDC, zero degree calorimeter.

* Corresponding author. Tel.: +41 76 487 4387; fax: +41 22 767 7910.

E-mail address: Yiota.Foka@cern.ch (Y. Foka).

¹ Deceased.

ARTICLE INFO

Article history:

Received 5 February 2010

Received in revised form

22 March 2010

Accepted 12 April 2010

Available online 7 July 2010

Keywords:

ALICE

Time Projection Chamber

ABSTRACT

The design, construction, and commissioning of the ALICE Time-Projection Chamber (TPC) is described. It is the main device for pattern recognition, tracking, and identification of charged particles in the ALICE experiment at the CERN LHC. The TPC is cylindrical in shape with a volume close to 90 m³ and is operated in a 0.5 T solenoidal magnetic field parallel to its axis.

In this paper we describe in detail the design considerations for this detector for operation in the extreme multiplicity environment of central Pb–Pb collisions at LHC energy. The implementation of the resulting requirements into hardware (field cage, read-out chambers, electronics), infrastructure (gas and cooling system, laser-calibration system), and software led to many technical innovations which are described along with a presentation of all the major components of the detector, as currently realized. We also report on the performance achieved after completion of the first round of stand-alone calibration runs and demonstrate results close to those specified in the TPC Technical Design Report.

© 2010 CERN for the benefit of the ALICE collaboration. Published by Elsevier B.V. All rights reserved.

1. Introduction

The ALICE [1,2] Time-Projection Chamber (TPC) [3] is the main device, in the ALICE ‘central barrel’, for tracking of charged particles and particle identification.

The main goal of the ALICE experiment at the CERN Large Hadron Collider (LHC) is the investigation of Pb–Pb collisions at a center-of-mass energy of 5.5 TeV per nucleon pair. Tracking of charged particles in such an environment can only be performed with a detector which can cope with unprecedented densities of charged particles: the maximum expected rapidity density in Pb–Pb collisions at LHC energy is about 3000 [4]. Furthermore, a comprehensive experiment needs to cover full azimuth and provide a significant acceptance in pseudo-rapidity $\eta = -\ln \tan \theta/2$ with θ the polar angle. In addition, the detector should provide excellent momentum and energy-loss resolution and run at extremely high rates (> 300 Hz for Pb–Pb central collisions, > 1.4 kHz for proton–proton collisions).

The resulting detector choice was a large-volume TPC with overall ‘conventional’ lay-out but with nearly all other design parameters beyond the state of the art. This manuscript describes in detail the resulting detector and outlines the path from design considerations to construction and commissioning.

In outline the ALICE TPC consists of a hollow cylinder whose axis is aligned with the beams from the LHC and is parallel to the ALICE detector’s solenoidal magnetic field. The active volume has an inner radius of about 85 cm, an outer radius of about 250 cm, and an overall length along the beam direction of 500 cm. A conducting electrode at the center of the cylinder, charged to 100 kV, provides, together with a voltage dividing network at the surface of the outer and inner cylinder, a precise axial electric field of 400 V/cm. The detector is filled with a counting gas consisting of a Ne–CO₂–N₂ mixture at atmospheric pressure. Charged particles traversing the detector ionize the gas. The ionization electron drift, under the influence of the electric field, to the endplates of the cylinder, where their arrival point in the cylinder plane is precisely measured. Together with an accurate measurement of the arrival time (relative to some external reference such as the collision time of the beams from the LHC) the complete trajectory in space of all charged particles traversing the TPC can be determined with precision.

The ALICE set-up is shown in Fig. 1. The TPC surrounds the Inner Tracking System (ITS) which is optimized for the determination of the primary and secondary vertices and precision tracking of low-momentum particles. On the outside the Transition Radiation Detector (TRD) is designed for electron identification. The outermost Time-Of-Flight (TOF) array provides pion, kaon, and proton identification. In addition, there are three single-arm detectors: the Photon Spectrometer (PHOS), the Electro-Magnetic CALorimeter (EMCAL) and an array of RICH

counters optimized for High-Momentum Particle Identification (HMPID).

The 0.5 T magnetic field in the central barrel is provided by the L3 solenoidal magnet previously used by the L3 experiment.

The ALICE TPC was designed to cope with the highest conceivable charged particle multiplicities predicted, at the time of the Technical Proposal (TP), for central Pb–Pb collisions at LHC energy [1,5,6], i.e. rapidity densities approaching $dN_{ch}/dy = 8000$ at center-of-mass energy of 5.5 TeV.² Its acceptance covers 2π in azimuthal angle and a pseudo-rapidity interval $|\eta| < 0.9$. Including secondaries, the above charged particle rapidity density could amount to 20 000 tracks in one interaction in the TPC acceptance.

Furthermore, the design of the readout chambers, electronics, and data handling allows inspection of up to several hundred such events per second with a maximum interaction rate of 8 kHz for Pb–Pb collisions, implying special precautions to minimize the effects of space-charge built-up in the drift volume of the TPC on the track reconstruction.

To realize a detector which performs efficiently in such an environment required the development of many new components and procedures. A summary of the design parameters is presented in Tables 1–3. A summary and system overview can be found in Ref. [2].

In this paper we describe the major components of the detector as currently realized and report on the performance achieved after completion of the first round of calibration runs.

The first major challenge was the design and construction of the field cage, whose overall thickness should not exceed 5% of a radiation length while providing, over a volume of nearly 90 m³, an axial electric field of 400 V/cm with distortions in the 10^{−4} range. The realization of this device is described in Section 2.

The readout chambers are installed at the two endplates of the cylinder. Their design is based on the Multi-Wire Proportional Chamber (MWPC) technique with pad readout. To ensure low diffusion of the drifting electrons and a large ion mobility, Ne was chosen as the main component of the counting gas. Furthermore, the size of the readout pads had to be adapted to the expected large multiplicities, implying pad sizes as small as 4×7.5 mm² in the innermost region. As a consequence, the readout chambers have to be operated safely at gains near 10⁴. In Section 3 we describe the technical implementation and report on the first operating experience of these detectors.

In Section 4 we discuss the design and implementation of the electronics chain. Because of the high granularity (557 568 readout channels) special emphasis was placed on very low power consumption. To cope with the large dynamic range needed to track particles from very low to high momenta, and to

² More recent estimates [4] put this number at $dN_{ch}/dy < 3000$.




Fig. 1. ALICE schematic layout [2].

Table 1
General parameters of the ALICE TPC.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$ for full radial track length $-1.5 < \eta < 1.5$ for 1/3 radial track length
Azimuthal coverage	360°
Radial position (active volume)	$848 < r < 2466$ mm
Radial size of vessel (outer dimensions)	$610 < r < 2780$ mm
Radial size of vessel (gas volume)	$788 < r < 2580$ mm
Length (active volume)	2×2497 mm
Segmentation in φ	20°
Chambers per sector ^a	2
Total number of readout chambers	$2 \times 2 \times 18 = 72$
Inner readout chamber geometry	Trapezoidal, $848 < r < 1321$ mm active area
Pad size	4×7.5 mm ² ($r\varphi \times r$)
Pad rows	63
Total pads	5504
Outer readout chamber geometry	Trapezoidal, $1346 < r < 2466$ mm active area
Pad size	6×10 and 6×15 mm ² ($r\varphi \times r$)
Pad rows	$64 + 32 = 96$ (small and large pads)
Total pads	$5952 + 4032 = 9984$ (small and large pads)
Detector gas	Ne–CO ₂ –N ₂ [85.7–9.5–4.8]
Gas volume	90 m ³
Drift voltage	100 kV
Anode voltage (nominal)	1350 V (IROC) 1570 V (OROC)
Gain (nominal)	7000–8000
Drift field	400 V/cm
Drift velocity (NTP)	2.65 cm/ μ s
Drift time (NTP)	94 μ s
Diffusion (longitudinal and transversal)	220 μ m/ $\sqrt{\text{cm}}$
Material budget (including counting gas) ^b	$X/X_0 = 4.1\%$ near $\eta = 0$

^a See Fig. 1.

^b See Table 4.

Table 2
ALICE TPC electronics parameters.

Front-End cards (FECs)	121 per sector $\times 36 = 4356$
Readout partitions	6 per sector, 18 to 25 FECs each
Total readout control units	216
Total pads—readout channels	557 568
Pad occupancy (for $dN/dy = 8000$)	40–15% inner/outer radius
Pad occupancy (for pp)	$5\text{--}2 \times 10^{-4}$ inner/outer radius
Event size (for $dN/dy = 8000$)	≈ 70 MByte
Event size (for pp)	0.1–0.2 MByte
Total bandwidth	35 GByte/s
Maximum trigger rate	300 Hz Pb–Pb central events 1.4 kHz proton–proton events
ADC	10 bit
Sampling frequency	5–10 MHz
Time samples	500–1000
Conversion gain	6 ADC counts/fC

Table 3
Expected resolution parameters.

Position resolution (σ)	
in $r\varphi$	1100–800 μ m inner/outer radii
in z	1250–1100 μ m
dE/dx resolution, isolated tracks	5.0%
$dN/dy = 8000$	6.8%

provide low noise performance combined with efficient baseline restoration and zero-suppression, the signals from the preamplifier/shaper chip were fed into a 10-bit, 10 MSPS ADC integrated into a digital chip. We report on the implementation and running experience of the electronics chain as realized in the ALICE TPC.

Successful operation of a very large detector like the ALICE TPC depends on a considerable amount of infrastructure and services,

along with sophisticated gas and cooling systems. One of the major challenges in this context was to provide a temperature stability of $<0.1\text{K}$ across the full volume of the TPC. This requirement originates from the strong temperature dependence of the drift velocity in the Ne–CO₂–N₂ mixture at realistically accessible electric fields. Furthermore, it is essential to control the O₂ content of the counting gas below a level of 5 ppm to keep to a minimum the absorption of electrons over the long drift length. The approach to solve these and many other technological challenges is described in the sections on cooling, gas system, infrastructure and services, and Detector-Control System (DCS).

Calibration and commissioning of the ALICE TPC relied, before the availability of any collisions from the LHC, on three different methods: a set of external UV laser beams was used to characterize field distortions and to determine the magnitude of the correction from $E \times B$ effects on the drifting electrons originating from the residual non-parallelism of the electric and magnetic field inside the drift volume. Furthermore, radioactive krypton was inserted through the gas system into the detector to provide efficient and precise amplitude calibration of all 557 568 readout channels. Finally, extensive measurements with cosmic rays were performed to determine tracking efficiencies, energy loss, and momentum resolution of the detector. The methods used and results obtained during these calibrations are described in detail in Sections 10 and 11. They demonstrate that detector performance is close to that specified in the original technical design report [3].

2. Field cage

The purpose of the field cage is to define a uniform electrostatic field in the gas volume in order to transport ionization electrons from their point of creation to the readout chambers on the endplates without significant distortions. The field cage provides a stable mechanical structure for precise positioning of the chambers and other detector elements while being as thin as possible in terms of radiation lengths presented to the tracks entering the TPC (see Fig. 2). In addition, the walls of the field cage provide a gas-tight envelope and ensure appropriate electrical isolation of the field cage from the rest of the experiment.

It is a classical TPC field cage with the high voltage electrode in the middle of the detector. Electrons drift to both end plates in a uniform electric field that runs parallel to the axis of the cylinder.




Fig. 2. 3D view of the TPC field cage. The high voltage electrode is located at the center of the drift volume. The endplates with 18 sectors and 36 readout chambers on each end are shown.

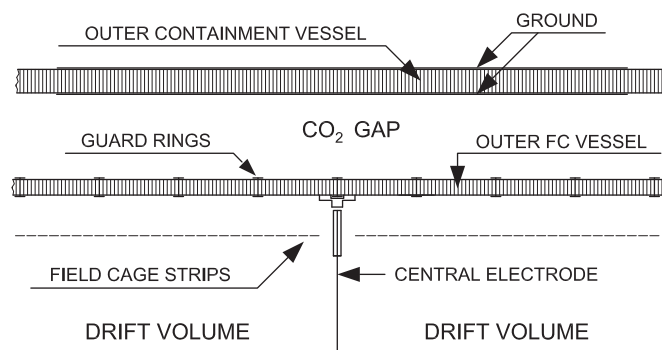


Fig. 3. Detail view of the outer field cage near the central electrode.

The TPC is filled with a mixture of neon, carbon dioxide, and nitrogen because the multiple coulomb scattering in this gas mixture is relatively low, it has good diffusion characteristics, and it has a high positive ion mobility that helps to clear positive ions out of the drift volume in a short amount of time (see Section 6). However, to also have fast electron drift velocities requires putting 100 kV on the central electrode. The isolation of the high voltage field cage from the rest of the experiment is ensured by using CO₂ filled gas gaps between the containment vessels and the field cage vessels; see Fig. 3.

The design of the ALICE field cage is similar to the design of the field cage used in the NA49 experiment [7]. An important part of the design is the requirement to prevent charge build-up, and possible breakdown, on solid insulator surfaces between the field-defining strips and so the use of these insulators is minimized or completely avoided.

The ALICE field cage consists of two parts; a field cage vessel with a set of coarsely segmented guard rings and a finely segmented field cage which is located inside the field cage vessel. The guard rings on the field cage vessel help to avoid large electric fields due to charge build-up on the surface of the vessel. The rings have a 92 mm gap between them and this corresponds to a relatively low field gradient of 46.7 V/mm on the insulating surface between the rings. The guard rings are made of 13 mm wide strips of aluminum tape and they are placed on both sides of the containment vessel with a pitch of 105 mm. Small holes were drilled through the walls of the vessel to allow for electrical contact between corresponding rings and filled with Al foil feed-throughs and sealed with epoxy. The potentials for the guard rings are defined by an independent chain of $24 \times 500\text{M}\Omega$ resistors (per end). The first of these resistors is connected to the rim of the high-voltage electrode. The last one is connected to ground through a 100 k Ω resistor, across which the voltage drop is measured for monitoring purposes. The field gradient between the guard rings matches the field gradient on the finely segmented field cage which lies inside the guard ring vessel.

The finely segmented field cage is made of 165 free standing mylar strips. In principle, there is space for 166 strips but mechanical considerations near the central electrode prevents the installation of the first strip and so it is left out. (The resistor chain for the field cage includes 166 resistors and so in this way the missing strip is included, see below.) The mylar strips do not come into contact with the field cage vessel or the gas containment vessel but, instead, are wrapped around a set of 18 Makrolon[®] rods that are regularly spaced around the circumference of the TPC. The Makrolon[®] rods are located at a radius of 2542.5 mm on the outer field cage and 815 mm on the inner field cage. The field cage strips are made of 13 mm wide aluminized mylar. They are stretched over the Makrolon[®] rods with a pitch of 15 mm. This leaves a 2 mm insulation gap between each pair of strips and




Fig. 4. Cross-sectional side view of the TPC with relevant dimensions (in mm). The service support wheels and one of the I-bars are also shown.

creates a voltage gradient of 300 V/mm across the gap. The neon gas mixture is the only insulator that separates the strips of the field cage except for the region where the strips touch the rods. The resistor chains for the field cages are located inside one of the Makrolon[®] rods for the outer field cage, one on each end, and also inside one rod for the inner field cage, one on each end.

The combination of the fine segmentation of the field cage strips and the coarsely segmented guard rings is a robust and stable electrostatic design when 100 kV is applied to the central electrode. This design minimizes the electric field distortions that would occur inside the drift volume of the TPC if the electric field lines were to go directly to ground from the field cage. Electrostatic calculations demonstrate that the field shape distortions inside the drift volume are below 10^{-4} at a distance of 15 mm from the strips [8].

2.1. Vessels

Four cylinders are required to make the complete field cage structure; two field-cage vessels (one inner and one outer) and two containment vessels (inner and outer). The cylinders are composite structures made with a Nomex[®] honeycomb core sandwiched between prepreg sheets (epoxy fiberglass) and Tedlar[®] foils, to provide a light, rigid, and gas-tight structure. A cross-sectional view of the TPC with some relevant dimensions is shown in Fig. 4.

The inner and outer field-cage vessels (see Fig. 4) define the gas volume of the TPC. They have radii of 788 and 2580 mm, respectively. The thickness of the Nomex[®] honeycomb core is 20 mm for both vessels. Clamps are glued to the inside walls of the cylinders to provide support for the outer rim of the high voltage electrode and all the rods.

The containment vessels surround the field cage vessels and they provide gas tight and grounded enclosures at the inner and

Table 4

The thickness of the inner and outer field cage components are listed in radiation lengths.

Part	X/X_0 (%)
Central drum	0.540
Inner CO ₂ gap	0.085
Inner field cage vessel	0.730
Inner field cage strip	0.012
Inner field cage total	1.367
Drift gas	0.607
Outer field cage strip	0.012
Outer field cage vessel	0.730
Outer CO ₂ gap	0.081
Outer containment vessel	1.330
Outer field cage total	2.153

The total thickness presented to a particle entering the TPC at $\eta=0$ is about 1.4% X_0 .

outer diameters of the TPC. To maintain a good ground, both walls of the cylindrical composite structures are covered with 50 μ m thick aluminum foil. The containment vessels are separated from the field cage vessels by an insulating gap and these gaps are continuously flushed with CO₂ to isolate the field cage voltage from the grounded containment walls. The distance between the outer field cage vessel and the outer containment vessel is 147 mm. The Nomex[®] core for the containment vessel is 30 mm thick. The inner containment vessel is made of three parts: a central drum which surrounds the inner tracking system (ITS) and two cones that support the drum; see Fig. 4. The central drum is cylindrical in shape (1420 mm long and 610 mm in radius), and has a Nomex[®] core that is 5 mm thick. The support cones are made of 3 mm thick aluminum and they span the distance from the central drum to each endplate. They provide support for the

ITS while leaving room for a muon absorber on one end and services for the inner tracking system on the other. The attachment between the central drum and the cones is sealed with a 2 mm thick flat neoprene rubber ring. The CO₂ gap between the inner containment vessel and the inner field cage vessel is 156 mm thick at the centerline of the detector and decreases to 80 mm near the endplates.

The thickness of the critical components of the field cage vessels are listed in Table 4 in units of radiation length.

2.2. Central electrode

The central electrode is made of a stretched 23 μm thick mylar foil which is aluminized on both sides and held flat by circular inner and outer aluminum rims. Three foils were glued together by laying them side by side and gluing 50 mm wide aluminized mylar bands over the junctions. The resulting 6 × 6 m² foil was stretched with pneumatic jacks and glued onto a set of inner and outer rims. After curing, a second set of rims was lowered into position and glued to the foil.

2.3. Rods

A total of 72 rods are positioned axially on the internal walls of the inner and outer field-cage vessels and in the corners between the readout chambers. Their main role is to hold the field cage strips for the inner and outer field cages. The rods are made of several 178 and 209 mm long pieces of Makrolon[®] tube (a special sort of Plexiglas) which have been glued together. The final rod assemblies have an outer (inner) diameter of 44 mm (36 mm). Their outer surfaces were machined with 2 mm wide and 2.5 mm

deep grooves, at a pitch of 15 mm, to increase the distance along the insulator surface between the strips. An aluminum ring at each glue junction helps to minimize and redistribute the accumulation of charge along the rods. The gluing operation was performed on a precision jig in order to achieve a uniform spacing of the strips to within 100 μm. The rods for the outer field cage, except for their grooves, are coated with copper to avoid charge accumulation on their exposed surface. The rods are held in position with holding clamps which are glued to the walls of the field cages with a 500 mm spacing between the clamps.

2.3.1. Resistor rods

The voltage dividers are integrated with the so-called resistor rods, and they are inserted into four of the rods of the field cage: inner, outer, and on both sides. The resistor rods contain a chain of resistors which define the potential on each strip of the field cage. The innovative design of these rods allows for water cooling and serviceability. The power dissipated by the resistors is removed by a water-cooling circuit that runs back and forth through the rods. A set of contacts ensures a good connection between the resistors in the chain and to each strip. Provisions are made for insertion, contacting, locking, and removal of the rods for service. Details of both ends of a resistor rod can be seen in Figs. 5 and 6.

For each resistor rod, a set of 165 copper plates, 0.5 mm thick, are held together by short sections of PEEK[®] (polyacryletheretherketone, a thermoplastic) tubing which are glued together, thus defining the 15 mm pitch for the strips. The resulting tube is 2.5 m long and the central hole is used to flush the system with drift gas since the Makrolon[®] rod is not necessarily gas-tight. In addition to the central hole, two more holes are drilled into the copper plates through which two ceramic pipes, 3 mm inner and 9 mm outer diameter, are inserted. The copper plates are connected to these ceramic pipes with thermally conductive glue in an alternating pattern. The pipes are bridged together at the high voltage end by a stainless steel tube so that de-ionized water flows into one of the ceramic pipes and returns through the other pipe. In this manner, the power dissipated by the voltage divider is coupled to the copper plates and is removed by the cooling water.

A 7.5 MΩ resistor is connected between each of the Cu plates. The first resistor, from the central membrane to the first Cu plate, has a value of 15 MΩ to compensate for the missing first strips in the field cage. In all cases, the resistors are soldered to washers which are then screwed to the plates. This results in a resistor chain with a total resistance of 1245 MΩ inside the TPC drift volume plus an additional 4.286 MΩ at the end of the chain to allow for precisely tuning the voltage on the last strip of the field cage; see Fig. 6. A small piece of PEEK[®] material is screwed to




Fig. 5. Detailed view of the high-voltage end of the resistor rod, showing the cooling pipes, the central PEEK pipe, the heat-dissipating copper plates, the contacts to the strips, and the high-voltage contact, which matches the contact at the housing rod.




Fig. 6. Schematic of the mechanical and electrical arrangement of the ground side of the resistor rod. The currents through the resistor chain and through the cooling water are measured independently.

each copper plate, onto which a flexible, gold-plated stainless steel electrical contact is again screwed. A 50 μm Au–W wire is used to make the connection between the plate and the contact, thus minimizing the amount of heat transmitted to the contacts and into the drift volume. At the high-voltage end, a connector consisting of a crown of flexible contacts provides the electrical connection to the corresponding part in the field-cage rod; see Fig. 5.

The ground end of the resistor rod, shown in Fig. 6, is equipped with various resistors for properly terminating the assembly to ground, tuning the potential of the last strip, measuring the current through the resistor chain and measuring the current through the cooling water. The contacts on the field-cage rods are made of gold-plated brass and are glued into holes in the rod wall. The hooks to which the strips are attached are screwed onto these contacts.

With 100 kV on the central electrode, the total current flowing through one of the resistor chains is 91 μA ; this is a sum of 80 μA flowing through the field cage resistor chain, 8.4 μA through the guard ring chain and 2.5 μA flowing through the cooling water for the rod.

2.3.2. High-voltage cable rod

The cable that provides high voltage to the central electrode is inserted into one of the outer Makrolon[®] rods. The ground shield for the cable has been removed over the entire length of the rod (250 cm) and replaced by a semi-conductive carbon loaded polyethylene sleeve that provides a smooth voltage gradient inside the rod. The contact for the cable is similar to the resistor-rod contact, and again a special cable connects the rod's contact to the rim of the central electrode. A rod with a spare contact is installed on the other side of the TPC. A flange in the endplate ensures gas tightness of the rod and mechanical support for the cable.

2.3.3. Laser rods

Six outer rods per side are devoted to the laser calibration system. The laser rods are spaced uniformly around the perimeter of the TPC. The corresponding flanges for the rods include a quartz window for introducing a laser beam into the rod. The rod itself holds, in its interior, a set of mirrors which deflect the light into the drift volume of the TPC through openings machined in the rods for this purpose. The laser calibration system is described in detail in Section 7.

2.3.4. Gas rods

Ten rods from the outer field cage, and 17 rods for the inner field cage, are empty and so these rods plus the partially obscured laser rods are used to circulate gas through the TPC. The rods are machined with an array of 1 mm holes which have a 15 mm pitch. The inner rods are used for the gas inlet, and the outer rods are used for the gas outlet and this is the only way that gas goes in and out of the TPC. In this manner, the gas flows radially through the system thus minimizing the forces exerted on the central electrode.

2.4. Strips

The field-defining strips are made from aluminized mylar, 25 μm thick and 13 mm wide. Under a tension of 3.5 N, they are cut to the right length (5.246 and 16.018 m for the inner and outer strips, respectively). A custom-made tool was then used to fold Cu–Be foil around the end of each strip to produce a hook. The strips were then strung around the rods and connected to similar




Fig. 7. The field cage strips connect to the resistor rods with Cu–Be hooks. Metallic screws hold the hooks onto the rod and provide the electrical connection to contacts inside the rod.




Fig. 8. A view inside the field cage where the strips and supporting rods are visible. The central electrode reflects a view of the field cage and the readout chambers. The subdivision of the pad planes of the OROCs into four boards can be seen. The skirt electrodes around the OROCs are also visible.

hooks on the resistor rods, as shown in Fig. 7. A photograph showing the interior of the finalized field cage is shown in Fig. 8.

2.5. Skirts

The strips of the inner field cage run close to the inner edge of the readout chambers thus enabling a good match of the drift field with the potentials on the cover electrodes of the inner readout chambers (see Section 3). The voltages on the cover electrodes are tunable and this helps to ensure a good match. However, there is a gap between the outer readout chambers and the strips of the outer field cage which is too large to be left unfilled. The electric field would be distorted if it were left exposed and so a 38 mm wide skirt is inserted into the gap. The skirts are parallel to the endplate and are electrically interconnected so they can be set to an appropriate potential to minimize the distortions of the field. A temperature sensor (PT1000) is glued on the back side of each skirt sector, thus allowing for temperature measurements inside the volume of the TPC.

2.6. Endplates

The function of the endplates is to align the cylinders for the field cage vessels and to hold the readout chambers in position. The four cylinders are screwed to the flanges that connect the field cage vessels and the containment vessels, and are made gas-tight with O-rings. The aluminum structure of the endplate is 60 mm thick and the spokes are 30 mm wide. The cut-outs for the

readout chambers are equipped with provisions for the alignment of the chambers relative to the central electrode and are independent of the endplate itself (see Section 3). Gas tightness is achieved by a sealing foil and a double O-ring; one on the chamber and one on the endplate. The endplates also provide feed-throughs and flanges for gas, laser and electrical connections.

2.7. I-bars

The TPC is installed at an angle of 0.79° with respect to the horizontal due to the inclination of the LHC accelerator at the ALICE collision hall. This puts a gravity load on the endplates and leads to a displacement of the inner field cage with respect to the outer field cage. The elastic deformation of the endplates is removed by pulling on the inner field cage with a pair of I-bars. In Fig. 4, the I-bars are shown attached on the right hand side of the TPC and were designed so that they do not obstruct the area around the beam-pipe. The I bars are attached to the outer ring of the endplate and can push or pull on the inner field cage ring in order to re-align the field cages. During assembly in the ALICE detector, it was necessary to pull on the inner field cage with a force of 3 kN and an alignment of about $150\mu\text{m}$ was actually achieved.

3. Readout chambers

3.1. Design considerations

Large-scale TPCs have been employed and proven to work in collider experiments before [9], but none of them had to cope with the particle densities and rates anticipated for the ALICE experiment [5,6].

For the design of the Read-Out Chambers (ROCs), this leads to requirements that go beyond an optimization in terms of momentum and dE/dx resolution. In particular, the optimization of rate capability in a high-track density environment has been the key input for the design considerations.

The ALICE TPC has adopted MWPCs with cathode pad readout. In preparation of the TPC TDR [3] alternative readout concepts had also been considered, such as Ring Cathode Chambers (RCCs) [10] or Gas Electron Multipliers (GEMs) [11] as amplification structures. However, those concepts seemed, though conceptually convincing, not yet in an R&D state to be readily adopted for a large detector project, which had to be realized within a relatively short time span.

3.2. Mechanical structure

The azimuthal segmentation of the readout plane is common with the subsequent ALICE detectors TRD and TOF, i.e. 18 trapezoidal sectors, each covering 20° in azimuth. The radial dependence of the track density leads to different requirements for the readout-chamber design as a function of radius. Consequently, there are two different types of readout chambers, leading to a radial segmentation of the readout plane into Inner and Outer ReadOut Chamber (IROC and OROC, respectively). In addition, this segmentation eases the assembly and handling of the chambers as compared to a single large one, covering the full radial extension of the TPC.

The dead space between neighboring readout chambers is minimized by a special mounting technique (described in Section 3.4) by which the readout chambers are attached to the endplate from the inside of the drift volume. The dead space between two adjacent chambers in the azimuthal direction is 27 mm.




Fig. 9. Cross-section through a readout chamber showing the pad plane, the wire planes and the cover electrode.

This includes the width of the wire frames of 12 mm on each chamber (see Fig. 9) and a gap of 3 mm between two chambers. The total active area of the ALICE TPC readout chambers is 32.5 m^2 . The inner and outer chambers are radially aligned, again matching the acceptance of the external detectors. The effective active radial length (taking edge effects into account) varies from 84.1 to 132.1 cm (134.6–246.6 cm) for the inner (outer) readout chambers. The mechanical structure of the readout chamber itself consists of four main components: the wire planes, the pad plane, made of a multilayer Printed Circuit Board (PCB), an additional 3 mm Stesalit insulation plate, and a trapezoidal aluminum frame.

3.2.1. Wires

The wire length is given by the overall detector layout and varies from 27 to 44 cm in the inner chambers, and from 45 to 84 cm in the outer chambers.

At constant potential, the gas gain increases with decreasing anode-wire diameter. Thus, a small anode-wire diameter is preferred. Owing to their superior strength, gold-plated tungsten is preferable to copper–beryllium (an alloy of 98% Cu and 2% Be) for the thin anode wires. However, for the thicker cathode and gating grid wires this dense material would require unaffordable tensions on the thin wire ledges. Therefore, copper–beryllium is used.

However, electrostatic and gravitational forces cause the anode wires to sag, leading to gas-gain variations along the wire. The electrostatic sag is approximately proportional to the square of the length of the wire, and inversely proportional to the stretching force, while the gravitational sag depends on the density of the wire material. Therefore, the wires need to be mechanically strong enough to withstand the required stretching forces. We have chosen for the anode wires a diameter of $20\mu\text{m}$ and a stretching force of 0.45 N. The cathode and gating grid wires have a diameter of $75\mu\text{m}$ and a stretching force of 0.6 and 1.2 N for inner and outer chamber, respectively. The wire tension has been measured during production for all wires (see Section 3.3.1). The measured values ensure a wire sag around $50\mu\text{m}$ and thus are below the specified limit of $70\mu\text{m}$ [3].

3.2.2. Wire planes

The ALICE-TPC readout chambers employ a commonly used scheme of wire planes, i.e. 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 (see below), their wire geometry is different, as shown in Fig. 10. The gap between the anode-wire grid and the pad plane is 3 mm for the outer chambers, and only




Fig. 10. Wire geometries of the outer (left) and inner (right) readout chambers.

2 mm for the inner chambers. The same is true for the distance between the anode-wire grid and the cathode-wire grid. The gating grid is located 3 mm above the cathode-wire grid in both types of chamber. The anode-wire grid and the gating grid are staggered with respect to the cathode-wire grid. Henceforth we abbreviate the wire geometry of the inner chamber by (2–2–3), and that of the outer chamber by (3–3–3).

3.2.3. Anode-wire grid

Because of the expected high particle multiplicity and the relatively large gas gains required for the readout chambers (see below) a small anode-wire pitch was chosen for the ALICE TPC to minimize the accumulated charge per unit length of the anode wire and hence the risk of rate-induced gas-gain variations. This led to the choice of a 2.5 mm pitch for the anode wires. There are no field wires since they would reduce the signal coupling to the pads, as they pick up a significant fraction of the signal. The absence of field wires also considerably reduces the mechanical forces on the wire frames. However, a chamber without field wires requires a somewhat higher voltage to achieve the required gas gain and a higher geometrical precision in the positioning of the wires.

3.2.4. Cathode-wire grid

The cathode-wire grid separates the drift volume from the amplification region. A large number of the ions produced in the amplification avalanche are collected at the cathode wires without causing a noticeable reduction in electron transmission. The cathode wire pitch is 2.5 mm. Electrostatic calculations substantiating the above layout numerically are described in detail in Ref. [3].

3.2.5. Gating-wire grid

The gating grid is located above the cathode-wire grid, with alternating wires connected together electrically. In the open gate mode, all the gating grid wires are held at the same potential U_G , 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 bipolar field $U_G \pm \Delta U$ (see Section 8.5), which prevents electrons from the drift volume to get to the amplification region. This considerably reduces the integral charge deposit on the anode wires. In addition, the closed gate stops ions created in the avalanche processes of previous events from drifting back into the drift volume. This is important because escaping ions accumulate in the drift volume and can cause severe distortions of the drift field [12]. The goal is therefore to avoid increasing the ion charge density above that created by primary ionization. The resulting requirement is that the ion leakage from the amplification region has to be $< 10^{-4}$. To achieve an electron transparency close to 100% in the open mode while trapping ions and electrons in the closed mode, the offset and bias potentials of the gating grid are carefully adjusted. On the other hand, any ionization produced by particles traversing the gap between the gating grid and pad

plane will unavoidably be amplified at the anode wires and thus contribute to the integral charge accumulation. To minimize this effect, 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 μ s. To keep the alternating bias voltages low, the pitch between the gating grid wires is 1.25 mm.

3.2.6. Cover and edge geometry

The standard wire configuration (see Fig. 10) has a discontinuity at the transition to the next chamber in the radial direction. Electrostatic simulations, as shown in Fig. 11a for the standard wire configuration, revealed a substantial inefficiency of the ion gate.

The drift lines of positive ions originating from the amplification zone around the anode wire are shown. A sizable number of positive ions could leak back into the drift zone for this particular configuration. In order to improve the electrostatic configuration additional electrodes, i.e. ground and cover strips, were introduced (see Fig. 9). The voltage of the cover strip, which frames the whole chamber, can be tuned to maximize the homogeneity of the drift field in the amplification zone. The ground strip, together with the HV of the cover strip forces all drift lines to end on either the cover or ground strip. In addition, two thicker edge anode wires (75 μ m) were introduced. Their HV can be set independently thus providing a lower gain in the edge region. The corresponding field lines from electrostatic simulations are shown in Fig. 11b. Measurements of the ion-back flow for this configuration are given in Section 3.3.

3.2.7. Pad plane, connectors and flexible cables

The readout pad structure has been optimized for signal-to-noise ratio and position resolution at the desired gas gain. A detailed account for the considerations leading to the chosen pad layout is given in Ref. [3]. The adopted pad sizes are given in Table 1.

The pad size increases with radius in two steps following the radially decreasing track density. The pad plane itself is a 3 mm thick halogen-free FR4 printed circuit board. The signal from the pad is routed in three layers of traces and vias to the connector side. The routing of the traces from the pads to their connector pad was realized employing an auto-router and was optimized for minimum trace length and maximum trace-to-trace distance. The boundary conditions for electrical design of the inner (outer) readout chamber pad plane were the line width of four (8) mil³ and the minimum distance between lines of 13 (31) mil.

The pad plane connectors are standard for vertical connection of flat flexible cables. They have 23 pins each with a pitch of 1 mm. Six connectors in the radial direction are grouped to connect to the 128 readout channels of one FEC; four of them use two ground lines and two use one ground line to connect the ground on the pad plane with the ground of the FEC. The cables themselves are flexible KaptontTM cables, 8.2 cm long.

3.2.8. Pad plane capacitance measurements

An important optimization parameter of the pad plane is the minimization of the pad and traces-to-board capacitances. One way to reduce the pad-to-board (ground) capacitances is to make the traces as short as possible. Typically, traces from the border pads to their connectors are the longest ones. After optimization, i.e. basically overwriting auto-routers choice 'by hand', the capacitances were generally below 9 pF and as low as 6 pF for the shortest connections.

³ 1 mil = 25.4 μ m.

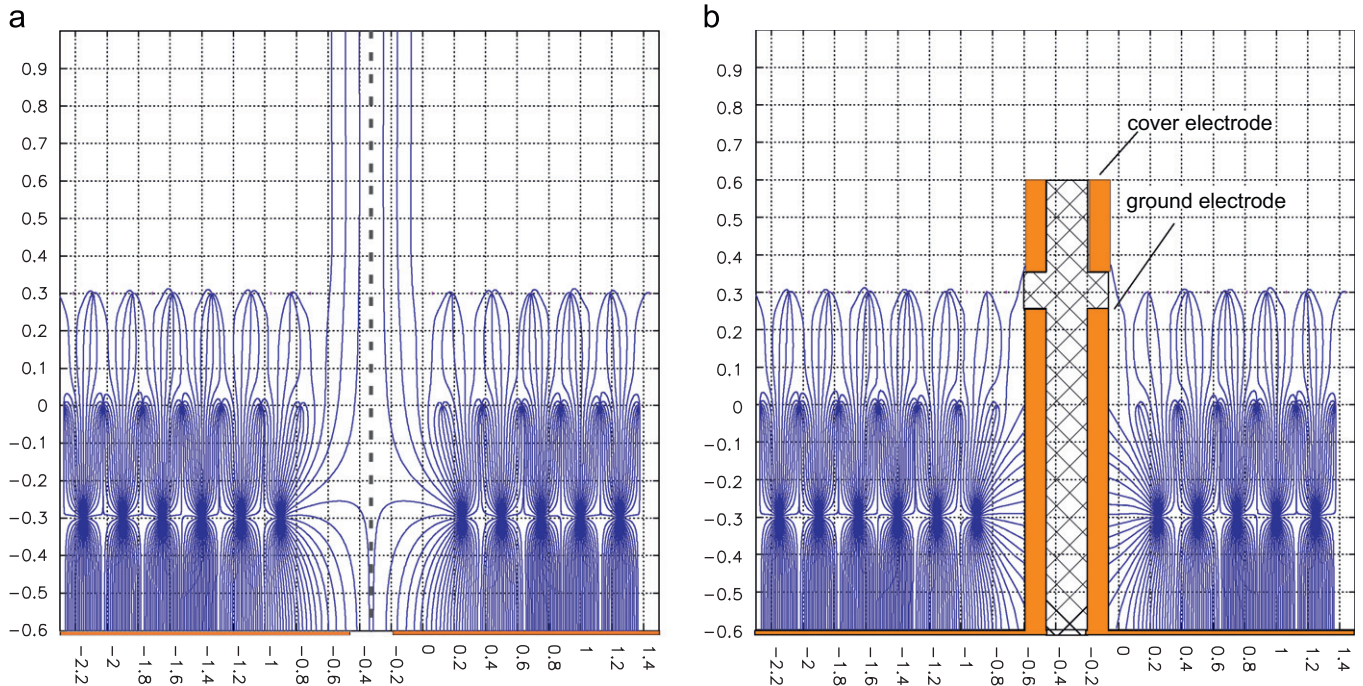


Fig. 11. Drift lines for positive ions at the border of two readout chambers with gate closed and standard wire configuration before (a) and after (b) the optimization of the electrostatic configuration at the borders of the chambers.




Fig. 12. Drawing of the Al-body of an outer readout chamber. Shown is the FEC side with the cut-outs for the flexible cables.

3.2.9. Al-body

Fig. 12 shows the aluminum body for an IROC, which holds the pad and wire planes. Its stability against deformation from the forces of the wire planes was optimized employing Finite-Element (FEM) calculations.

The Al-body has cut-outs to allow for the connection of the FECs to the connectors on the backside of the pad-plane. A cooling pipe was introduced into the Al-body to remove residual heat not taken away by the main cooling of the FECs or heat transmitted to the Al-body via the flat Kapton cables (see Section 5). The mechanical deformation of the readout chambers under the forces of the stretched wires was estimated via FEM calculation. The input to the FEM calculations is based on the technical specification of the TPC readout chambers and on the material parameters specified by the producer. The calculations yield the mechanical stress, the stiffness as well as the deformation of the overall structure: the maximal overall deformation of an inner module is 10 and 25 μm for an outer module. These values, as well as the corresponding values for the stress, are below the values, which are considered to be critical, i.e. would influence noticeably the performance of the chamber.

3.3. Tests with prototype chambers

Tests were performed both with several small custom-built chambers to investigate specific properties of TPC components and with real-size prototypes to verify the design before mass production of the chambers. The tests, with both small and real-size prototypes, are described in detail in Refs. [13–15]. In summary it was verified that the:

- *gating efficiency* measured both with a radioactive source and the laser is better than $\approx 0.7 \times 10^{-4}$, i.e. of the same order than the inverse of the envisioned gain of 2×10^4 ;
- *cross-talk* in the flat cables is of the order of 0.5–1.0%, i.e. it is thus not expected that the tracking performance is deteriorated in a significant way;
- *gas gain* estimated as a function of anode voltage, is of the order of 3×10^4 at 1280 V (for the original mixture without N_2) and thus sufficiently high to achieve a signal/noise ratio $S/N=20$;
- average current is stable (*long term stability*) during the irradiation of a chamber with a source corresponding to one year of Pb–Pb ($dN_{\text{ch}}/dy=8000$) running at 400 Hz trigger rate;
- chamber performance does not suffer from *aging* or *electron attachment* induced by out-gassing of construction materials; and
- chambers are stable at *high beam rate* and perform according to the design values for position and energy resolution.

3.3.1. Description of production steps

The work to produce readout chambers is split into work packages defined such that they can be carried out in parallel. One work package was the preparation of the module body, which included the insertion and gluing of the cooling loop, the insulation plate and the pad plane onto the Al-body.

After geometrical tests and cleaning, the module bodies were ready to receive the wire planes. The work package with the

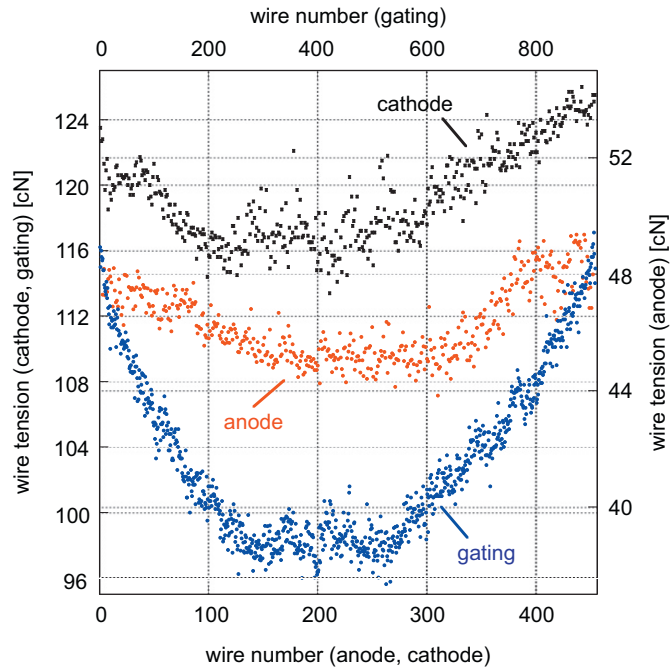


Fig. 13. Measured wire tension for all types of wires (anode, cathode, gating) for one of the outer readout chambers.

longest irreducible time span was the winding and gluing of the three wire planes (5 days); which defined the maximum chamber production rate, i.e. 1 chamber/week. The total production time per chamber was 30 days. A third work package was the testing of the final chamber which included a measurement of the wire tension (see Fig. 13) and the connectivity of the wire planes, as well as leak and performance tests, which are described below. The wire tension, e.g. for the anode wires, varies by $\pm 5\%$. The wire sag at the nominal voltage (1500 V) is for the longest wires (90 cm) about $50\ \mu\text{m}$, which changes the gain by about 4%. Such gain variations are fully equilibrated by the krypton calibration (see Section 10.4.1).

The production time for all 80 chambers was nominally 400 days, which, adding 25% contingency, amounts to an total effective production time of ≈ 2.5 years. In fact, the production of the multi-wire proportional readout chambers started in March 2001 and finished in May 2004, i.e. took a little more than three years.

3.3.2. Quality assurance and tests

All chambers were tested during and after production in order to validate them for the final assembly into the TPC. These tests included:

- gas tightness;
- pad plane deformation;
- performance: measurement of the gain as a function of high voltage;
- continuous operation under irradiation; and
- uniformity of response: irradiation scans over the active area addressing the gain homogeneity.

For these tests a dedicated setup was used, consisting of a gas box with a short field cage and auxiliary sensors, into which a readout chamber could be mounted. The installed sensors allowed us to monitor the O_2 and H_2O content of the gas, temperature and

pressure, as well as the currents of anode, cathode and gating grid wires.

Leak tests: The leak rate of each chamber was estimated from the O_2 contamination at the chamber outlet when flushing the chamber with a certain flow of fresh gas as described in Ref. [14]. Typically, the leak rate is 0.2 ml/h (5.5×10^{-5} mbar l/s) at a flow rate $f \approx 0.023\ \text{m}^3/\text{h}$. The acceptable O_2 contamination for the whole TPC is < 5 ppm [3], for a gas regeneration rate of $15\ \text{m}^3/\text{h}$. This translates into a leak rate of 0.5 l/h. If this leak was entirely due to the 2×36 readout chambers, each of them would be allowed to contribute with a maximum leak rate of 5 ml/h.

Long term stability tests: Each chamber is subjected to a long term stability test. For this, the anode voltage for each of the chambers is set to a value corresponding to a gain of 3×10^4 . A collimated iron source is placed at a fixed position for a full two-days irradiation test. The currents, X-ray fluxes, pulse height spectra and ambient pressure and temperature are continuously (every 15 min) recorded. The chamber is validated if no visible deterioration of its performance is observed.

Gain homogeneity tests: After the long term test, a scan over the active area of the chamber is performed. Keeping the same voltages, the ^{55}Fe source is consecutively placed in various predefined positions. Currents were recorded for each position to map the gain uniformity of the chamber. Misalignment of wire planes or sags due to insufficient wire tension would result in observable patterns on such a gain scan. Fig. 14 shows the scan performed on one of the chambers. The spots on the corners fall partially outside the active area of the chamber and present therefore a lower efficiency. Owing to the tight electrode geometry, high gains and the absence of field wires, a gain uniformity of the order of 10–20% was measured. However, no evidence of loose wires was observed.

From the 20 IROCs tested this way, 17 showed a stable and uniform ($\Delta G/G < \pm 20\%$) performance. Three of them did not pass the validation tests. Two chambers showed large (order of μA) dark or leakage currents at voltages below the operational ones. It is suspected that the reason for this behavior was a bad pad plane: either dirty or with a rough surface. These two chambers were therefore discarded for installation into the TPC. In one chamber an anode wire broke after several minutes under nominal voltage. The anode was burned at some 5 mm from the holding ledge. This was traced to faulty wire material. After this incident it was decided to inspect the wire quality employing electron microscopy before winding any anode wire plane.

Pad plane deformation: The spatial homogeneity of the chamber gain depends on the distance between anode wire and the cathode (wire) planes, i.e. on the wire sag due to gravitational and




Fig. 14. Scan performed on OROC after the long term irradiation test, which was done at the central position.




Fig. 15. Top: RMS of the deviation of the pad plane references points from the mean for an IROC chamber. Bottom: maximum distance between any two measured reference points of an IROC chamber.

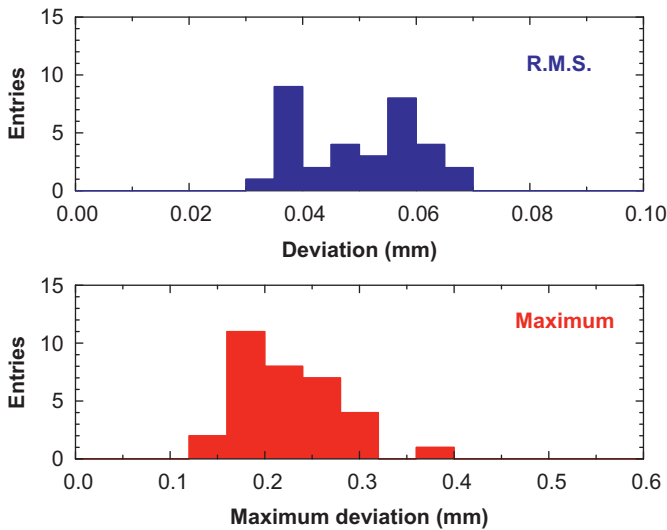


Fig. 16. Top: RMS of the deviation of the pad plane references points from the mean for an OROC chamber. Bottom: maximum distance between any two measured reference points of an OROC chamber.

electrostatic forces and on the planarity of the cathode pad plane. To ensure an acceptable contribution of the pad plane inhomogeneity to the gain variation, the pad plane deformation should be comparable to the average wire sag. After gluing the pad planes onto the Al-body the homogeneity of the pad plane has been measured for each chamber on a *xy*-table. 24 and 28 reference points were surveyed for IROC and OROC, respectively. The results are depicted in Figs. 15 and 16. The RMS value of the pad plane deformation is of the same magnitude as the average wire sag ($\approx 50\ \mu\text{m}$) and thus contributes with a value of $< 5\%$ to the gain variation. The maximum deviations are significantly higher and can contribute with values of up to 15% to the gain inhomogeneity.

3.4. Chamber mounting and pre-commissioning

The chambers are attached to the endplate from the inside of the TPC to minimize the dead space between neighboring chambers and



Fig. 17. Insertion of an OROC through the endplate. The tilt, polar and azimuthal angles of the chambers can be adjusted via handles and a transmission system.




Fig. 18. Rotation of an IROC inside the field cage. To prevent dirt falling into the field cage the FEC side of the chambers is closed with a cover.

to allow for accurate adjustment of the position of each chamber. The installation of the chambers required a special mounting tool that has a long manipulator arm and which allows the rotation and tilting of the chambers as they go into the TPC. The chambers were attached to the tip of the manipulator arm and inserted at an angle so they could be moved through the endplate (see Fig. 17). Thereafter the angle was re-adjusted so the chambers could be retracted (see Fig. 18) and fit into their final positions. The mounting technology was previously used by the ALEPH collaboration, from whom we inherited the manipulator device.

After installation, the position of each chamber relative to the endplate was determined by a photogrammetric survey. The adjustment of the chambers was achieved by shimming them at three points; two on the outer radius of each chamber and one at the inner radius. For each chamber, the shims were then individually re-machined based on the photogrammetric measurements. After putting the final shims in place, the relative accuracy of the alignment of the chambers was approximately $250\ \mu\text{m}$.

A flexible gasket was used to achieve a gas seal around the edges of the chambers. The gasket covered the gap between the

chamber and the endplate and was pressed over two O-ring seals; one on the chamber, and the other on the endplate.

Once the chambers were installed, four types of electrical measurements were done on each sector:

- pedestal and noise measurements;
- calibration pulser measurements (study the shaping properties of the electronics);
- measurements with the TPC laser system (for alignment purposes); and
- gain and drift-time measurements (using a cosmics trigger).

These measurements were performed before and after the TPC was installed in the L3 magnet and the results are reported in Section 10.

4. Front-end electronics and readout

4.1. General specifications

Charged particles traversing the TPC volume ionize the gas along their path, liberating electrons that drift towards the endplate of the chamber. The signal amplification is provided through avalanche effect in the vicinity of the anode wires of the readout chambers. The electrons and positive ions created in the avalanche, which move, respectively, towards the anode wire and the surrounding electrodes, induce a positive current signal on the pad plane. The current signal of a single avalanche, which is characterized by a fast rise time (< 1 ns) and a long tail (of the order of $50 \mu\text{s}$), carries a charge that, in the ALICE TPC, can be as low as a few fC. It is then delivered on the detector impedance which, to a very good approximation, is a pure capacitance of the order of a few pF. The shape of the signal tail, which is due to the motion of the positive ions, is rather complex and depends on the details of the chamber and pad geometry [16,17]. This tail, causing pile-up effects, sets the main limitation to the maximum track density at which a MWPC can be operated.

The readout of the signal is done by 557 568 pads that form the cathode pad plane of the readout chambers. The signals from the pads are passed to four 356 Front-End Cards (FECs), located 7 cm away from the pad plane, via flexible Kapton cables. In the FECs a custom-made charge-sensitive shaping amplifier, named PASA (PreAmplifier ShAper), transforms the charge signal induced in the pads into a differential semi-Gaussian voltage signal that is fed to the input of the ALTRO (ALice Tpc Read Out) chip. Each ALTRO contains 16 channels operating concurrently that digitize and process the input signals. Upon arrival of a first-level trigger, the data stream corresponding to the detector drift time ($\leq 100 \mu\text{s}$) is stored in a memory. When a second-level trigger (accept or reject) is received, the latest event data stream is either frozen in the data memory, until its complete readout takes place, or discarded. The readout can take place at any time at a speed of up to 200 MByte/s through a 40-bit-wide backplane bus linking the FECs to the Readout Control Unit (RCU), which interfaces them to the Data Acquisition (DAQ), the Trigger and the Detector Control System (DCS).

The main requirements for the readout electronics and the way they are derived from the detector performance requirements are discussed in Ref. [3, Section 5.1.1] and listed in Table 5.

One of the tightest requirements is defined by the extremely high pulse rate with which the ALICE TPC Front-End Electronics (FEE) has to cope. Indeed, the FEE has been designed to cope with a signal occupancy as high as 50%. Furthermore the extremely large raw data volume (750 MByte/event) asks for zero suppres-

Table 5
Readout electronics requirements.

Number of channels	557 568
Dynamic range	900:1
Noise (ENC)	$< 1000 e$ (rms)
Conversion gain	12 mV/fC
Cross-talk	$< 1\%$
Shaping time	≈ 200 ns
Sampling rate	5–10 MHz
Tail correction after $1 \mu\text{s}$	0.1%
Maximum readout rate (Pb–Pb)	300 Hz
Maximum readout rate (pp)	1.4 kHz
Power consumption	< 100 mW/channel

sion already in the FEE in order to fit events at the foreseen event rate into the DAQ bandwidth (216 links at 160 MByte/s). For example, for a trigger rate of 1 kHz as planned for pp collisions, this leads to a raw data throughput of 750 GByte/s, which is beyond the present data handling capabilities. Zero suppression, however, reduces the event sizes to about 0.1–0.2 MByte for pp and up to about 70 MByte for Pb–Pb, respectively (see Table 2). It should be noticed that in a high occupancy environment, in order to preserve the full resolution on the signal features (charge and arrival time), a very accurate cancellation of the signal tail and correction of the baseline have to be performed before the zero suppression.

Besides the optimization for the maximum rate the detector can be operated at, its dead time has to be minimized. This is done by introducing a derandomizing Multiple Event Buffer (MEB) to eliminate the contribution due to the random nature of the trigger arrival times. Simulations showed that four entries were a good trade-off between cost/size and effect (refer to Ref. [3, Section 5.1.7]).

The front-end electronics system has to satisfy many other constraints, while meeting the required performance specifications. Mainly, the readout 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 radiation load on the TPC is rather low (< 1 krad over 10 years). Thus standard radiation-soft technologies are suitable for the implementation of this electronics. However, some special care has to be taken to protect the system against Single Event Upset (SEU), see Section 4.7.

4.1.1. System overview

A single readout channel is comprised of four basic functional units (Fig. 19): (1) a charge sensitive amplifier/shaper (PASA); (2) a 10-bit 25 MSPS (mega samples per second) low power ADC; (3) a digital circuit that contains a shortening filter for the tail cancellation, the baseline subtraction and zero suppression circuits; and (4) the MEB.

The charge collected at the TPC pads is amplified and integrated by a low input impedance amplifier. It is based on a charge sensitive amplifier followed by a semi-Gaussian pulse shaper. These analog functions are realized by a custom integrated circuit, PASA (Section 4.2), which contains 16 channels. The circuit has a conversion gain of 12 mV/fC, an output dynamic range of 2 V, and produces a differential semi-Gaussian pulse with a shaping time (FWHM) of 190 ns.

The output signals of the PASA chip are digitized by a 10-bit pipelined 25 MSPS ADC (one per channel) operated at a sampling rate in the range of 5–10 MHz. The digitized signal is then processed by a set of circuits that perform the baseline subtraction, tail cancellation, zero-suppression, formatting and

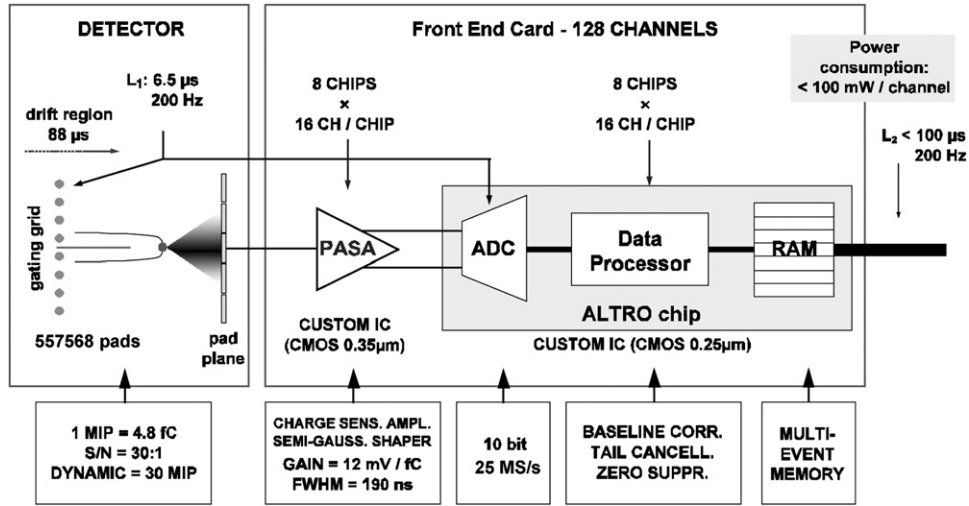


Fig. 19. An overview of the ALICE TPC front end electronics (© 2003 IEEE [19]).

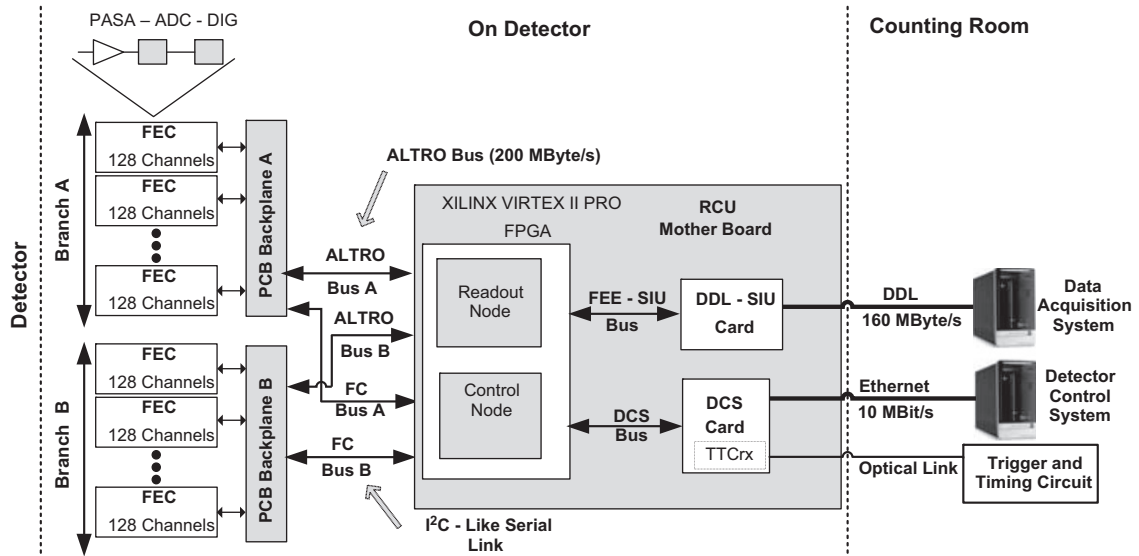


Fig. 20. Block diagram of the basic readout partition. The overall TPC readout consists of 216 readout partitions.

buffering. The ADC and the digital circuits are contained in a single, custom-made chip named ALTRO (Section 4.3).

The complete readout chain is contained in FECs (Section 4.4), which are plugged in crates that are supported by the service support wheel. They are mechanically decoupled with respect to the detector by Kapton cables (see Section 8.2). Each FEC houses eight PASAs and eight ALTRO chips, 128 channels in total. Another important component of the FEC is the Board Controller (BC), which implements a number of key functions for the readout and system monitoring. The FECs are connected to the cathode plane by means of six 8.2 cm long flexible cables. As illustrated in Fig. 20, FECs are grouped in readout partitions controlled by RCUs (Section 4.5). The number of FECs per partition varies according to its radial position on the detector, due to the trapezoidal shape of the sectors and the different pad sizes. Within a readout partition, the FECs are organized in two branches, connected to the RCU via separate backplanes.

The FECs communicate with the RCU by means of two independent buses: the ALTRO bus and the front-end control bus, both based on low-voltage signaling technology (Gunning Transistor Logic, GTL). The configuration of the FECs, the

distribution of the clock and trigger signals and the readout of trigger related data are performed via the ALTRO bus. The RCU uses the front-end control bus for all operations related to the safety and monitoring of the readout partition. Each of the 2×18 TPC sectors is equipped with six readout partitions with, respectively (from the innermost to the outermost) 18, 25, 18, 20, 20, and 20 FECs, accounting for a total of four 356 FECs and 216 RCUs. From the readout and control point of view, each partition represents an independent system.

4.2. PASA

The PASA [18] integrates 16 identical Charge Sensitive Amplifiers (CSAs) followed by a pole-zero cancellation network and a shaping amplifier. A simplified block diagram of the signal processing chain is shown in Fig. 21. The positive polarity CSA, with a capacitive and resistive feedback connected in parallel, is followed by a pole-zero cancellation network with a self-adaptive bias network, a CR-filter, two $(RC)^2$ -bridged-T filters, a common-mode feed-back network and two quasi-differential gain-2

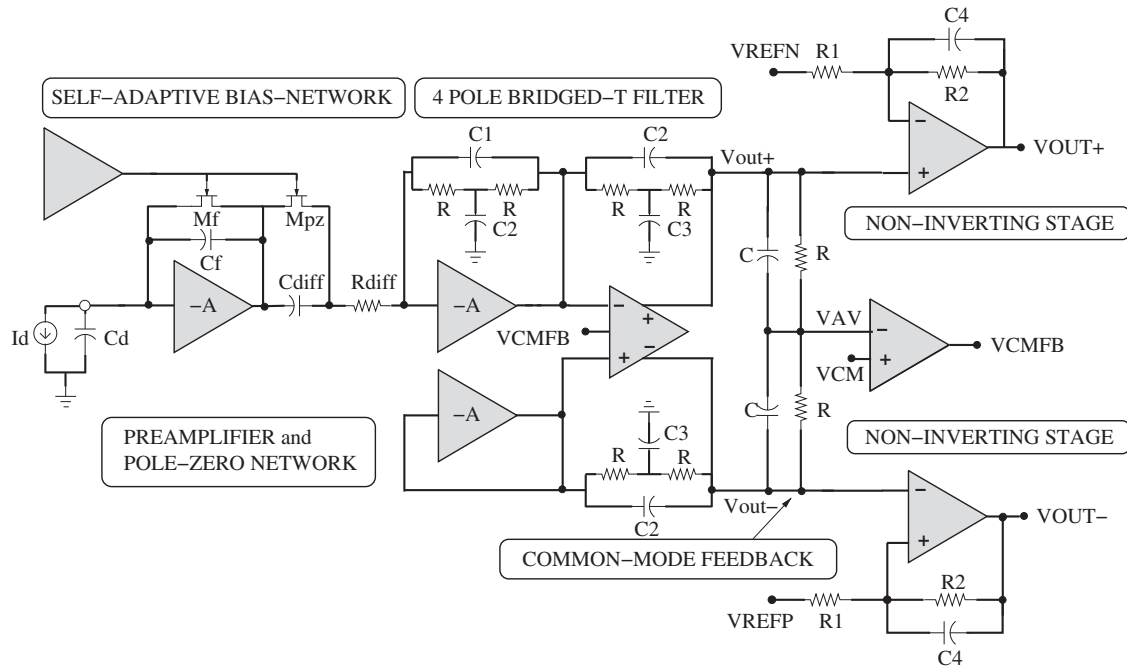


Fig. 21. A simplified block diagram of the PreAmplifier-ShAper (PASA) signal processing chain.

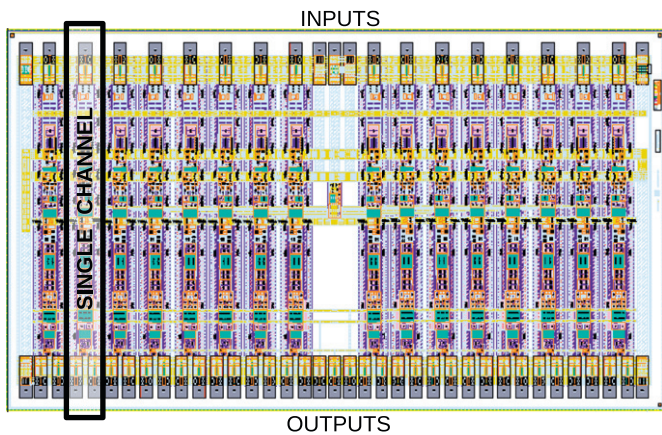


Fig. 22. Layout of the PASA chip.

amplifiers. The circuit is optimized for a detector capacitance of 18–25 pF.

The amplifier topology is based on a single-ended folded cascode amplifier followed by a source follower. As seen in Fig. 21, an NMOS transistor, which is operated in the subthreshold region to implement a large resistor, is connected in parallel to the feedback capacitor C_f . The purpose of this resistor is to avoid saturation of the CSA by continuously discharging the feedback capacitance. This resistor contributes to the parallel noise at the CSA input. A value of 10 M Ω is chosen as trade off between noise and count rate. Still, the relatively long discharge time constant of the CSA may cause signal pile-up. For this reason, the CSA is followed by a pole-zero cancellation network (M_{pz} and C_{diff}), which is combined with the $C_{diff}R_{diff}$ -filter stage. The signal is then amplified and further shaped by two second order bridged-T filters to optimize the signal-to-noise ratio and to limit the signal bandwidth. In the last PASA stage the signal levels are adjusted to match the input of the ALTRO chip.

The chip was manufactured in the C35B3C1 0.35 μ m CMOS technology featured by Austriamicrosystems. It has a width of 5.3 mm and a length of 3.4 mm and gives in total an area of

Table 6
PASA key performance figures.

Parameter	Specs	Simulation	Test
Noise	< 1000 e	385 e	\approx 385 e
Shaping time (ns)	190	212	\approx 190
Non-linearity	< 1%	0.19%	0.2%
Cross-talk	< 0.3%	–	< 0.1%
Baseline variation (mV)	–	–	\pm 80
Conv. gain (mV/fC)	12	12.74	\approx 12.8
Power (mW/ch)	< 20	11	11.67

All values are given for a detector capacitance of 12 pF.

18 mm². A picture of the produced PASA is shown in Fig. 22. The input and output pins are distributed along the length of the chip. The 16 channels, divided in two groups of eight channels each, are placed on each side of the bias network, located at the center of the chip. The de-coupling capacitors together with pads connected to ground are placed adjacent to each channel, creating a physical distance between the channels in order to reduce the cross-talk. In addition each channel is surrounded by a guard ring connected to the substrate which isolates them from each other and further reduces the cross-talk.

As detailed in Section 4.8, the performance of about 48 000 PASA was tested; 98% of the chips was fully functional. The general performance of the chip is listed in Table 6.

4.3. ALTRO

4.3.1. Circuit description

The ALTRO [19] is a mixed-signal custom integrated circuit containing 16 channels operating concurrently and continuously on the analog signals coming from 16 independent inputs. It is designed to process a train of pulses sitting on a common baseline. Fig. 23 shows a simplified block diagram of the chip. When a first-level trigger is received, a predefined number of samples is processed and temporarily stored in a data memory

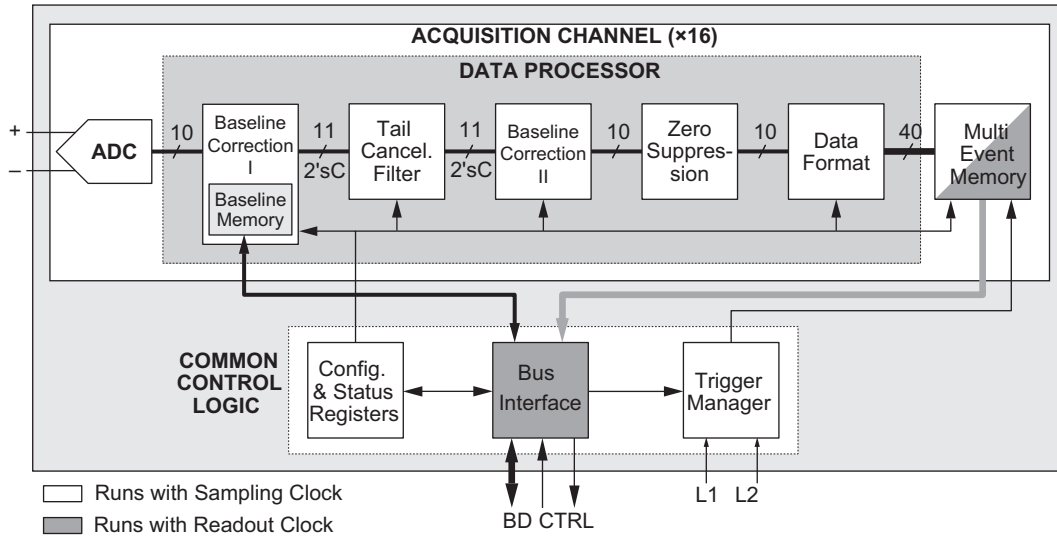


Fig. 23. ALTRO chip block diagram (© 2003 IEEE [19]).

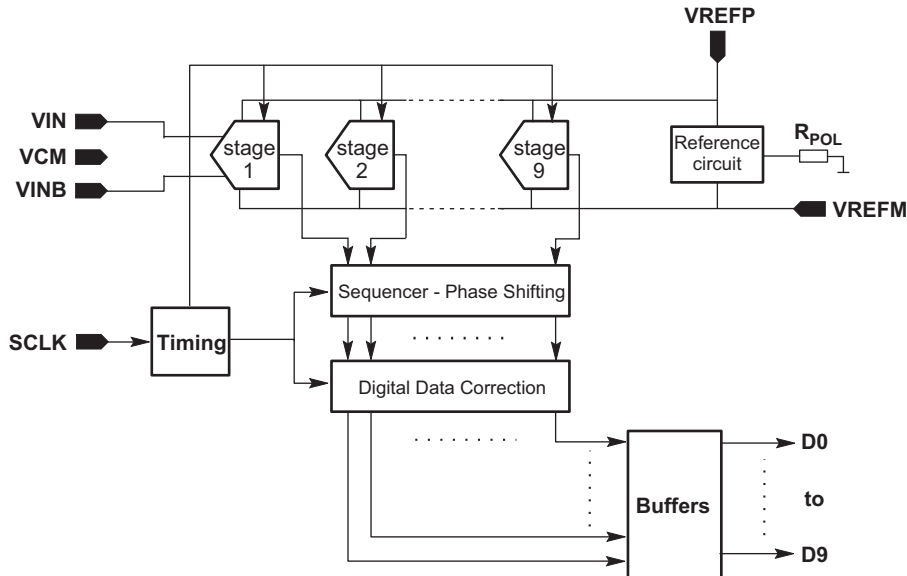


Fig. 24. Block diagram of the ALTRO ADC (© 2003 IEEE [19]).

(acquisition). The acquisition is frozen if a positive second-level trigger is received; otherwise it is overwritten by the next acquisition.

With reference to Fig. 23, a short description of the main building blocks is given following the signal processing path.

ADC. The Analog to Digital conversion is based on the ST Microelectronics TSA1001 [20], a CMOS 10-bit pipelined ADC. The block diagram of this ADC is presented in Fig. 24. The conversion pipeline consists of nine stages with an overall latency of 5.5 clock cycles.

The internal construction of the ADC is fully differential, and allows up to 2V differential swing. A polarization current, provided for each channel by an internal resistor, defines the ADC bandwidth and power consumption. The polarization resistor is divided in multiple taps such that only one metal layer has to be changed in order to optimize the power consumption for the required bandwidth. Two versions of the ALTRO chip have been produced. They were optimized for maximum sampling rates of 25 and 40 MHz and power consump-

tions of 12.5 and 43 mW/channel, respectively. The ALICE TPC uses the 25 MSPS version.

Data processor: The Data Processor conditions the signal in several processing stages. The first stage is the Baseline Correction I. Its main task is to prepare the signal for the tail cancellation by removing low frequency perturbations and trigger correlated effects. While the first is implemented by a self-calibration circuit based on a first order infinite impulse response filter, the latter is achieved by subtracting a pattern stored in a 10-bit wide 1k deep Pedestal Memory (PMEM). The next processing block, the Tail Cancellation Filter, which is based on a third order infinite impulse response filter, is able to suppress the tail of the pulses within 1 μ s after the peak, with an accuracy of 1 LSB [16]. Since the filter coefficients for each channel are fully programmable, the circuit is able to cancel a wide range of signal tail shapes. This also allows maintaining a constant quality of the output signal regardless of the actual detector operation parameters (gas and anode voltage), aging effects on the detector, and channel-to-channel fluctuations. The subsequent processing

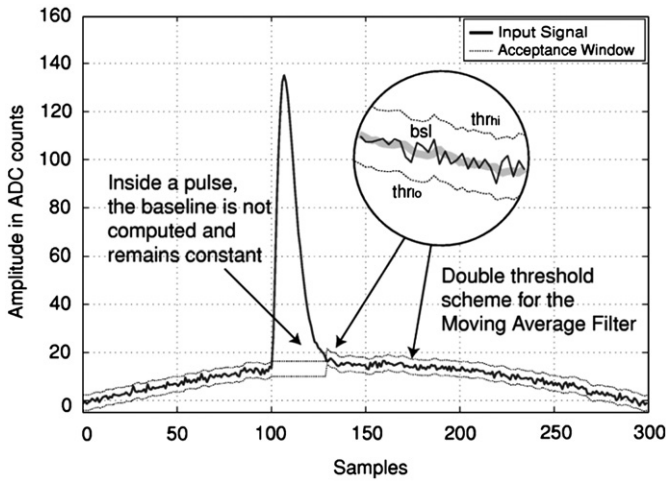


Fig. 25. Baseline correction II block operation principle (© 2003 IEEE [19]).




Fig. 26. Zero suppression scheme (© 2003 IEEE [19]).

block, the Baseline Correction II, applies a baseline correction scheme based on a 8-tap moving average filter. This scheme removes non-systematic perturbations of the baseline that are superimposed to the signal (see Fig. 25). At the output of this block, the signal baseline is constant with an accuracy of 1LSB. Such accuracy allows an efficient signal compression implemented in the Zero Suppression unit, which discards all data below a programmable threshold (see Fig. 26). In addition this unit features the following three functions. (1) A glitch filter checks for a consecutive number of samples above the threshold, confirming the existence of a real pulse, and thus reducing the impulsive noise sensitivity. (2) In order to keep enough information for further extraction, the complete pulse shape must be recorded. Therefore, the possibility to record pre- and post-samples is provided. (3) Finally, the merging of two consecutive data sets that are closer than three samples is performed. In the Data Format unit, each data packet is formatted with its time stamp and size information such that reconstruction is possible afterward. The output of the Data Processor is sent to a 5kByte data memory able to store up to eight acquisitions.

Multi-event memory: The ALTRO channel data memory (1024×40 bit) is partitioned in either four or eight buffers. The size of the memory allows storing four complete 1000-sample acquisitions with non-zero-suppressed data. If the Data Processor is configured to process < 500 samples, the 8-buffer partitioning can be used. In order to reduce the noise, the basic principle of operation is that all bus activity should be stopped during the

Table 7
ALTRO physical characteristics.

Process	ST HCMOS-7 (0.25 μm)
Area	64 mm^2
Dimensions	7.70 \times 8.35 mm^2
Transistors	6 million
Embedded memory	800 kbit
Supply voltage	2.5 V
Package	TQFP-176




Fig. 27. ALTRO chip layout (© 2003 IEEE [19]).

acquisition time. For this reason, the data memory manager interrupts the readout when a trigger is received and resumes only once the acquisition has finished.

Data can be read out from the chip, as standalone circuit, at a maximum speed of 60 MHz through a 40-bit wide bus, yielding a total bandwidth of 300 MByte/s. The readout speed and the ADC sampling frequency are independent. In the FEC the ALTRO chips are readout at a frequency of 40 MHz.

4.3.2. Physical implementation

The ALTRO chip is manufactured in the ST Microelectronics CMOS 0.25 μm (HCMOS-7) technology. The main physical characteristics of the circuit are summarized in Table 7.

The integration of the ADC imposes certain restrictions to the layout and the pin-out of the chip in order to guarantee a good performance in terms of noise and conversion reliability. The 16 ADCs are arranged in two octal ADC macros. The pedestal memories are placed close to the macros on the left and right side as shown in Fig. 27. The data memories are placed towards the center of the chip, distant from the ADCs macros. The placement of the memories reflects the regular structure of the 16 concurrent processing channels. The processing logic is distributed in the remaining space. To reduce the effect of digital noise on the ADC, the following strategy was applied during the layout phase. Since 95% of the logic works on the sampling clock, the phase of the clock signal distributed to all the flip-flops can be adjusted

such that the switching of all digital nodes occurs outside the aperture time of the ADC. Each ADC block contains a passive clock tree balanced with an accuracy of 1 ps.

As detailed in Section 4.8, 48 000 chips have been fabricated with a production yield of 84%. A number of tests were implemented to assess the performance of the chip (see Table 8).

4.4. Front-end card (FEC)

The Front-End Card (FEC) contains the complete readout chain for the amplification, shaping, digitization, processing and buffering of the TPC signals; it must handle the signal dynamic range of about 10 bits with minimal degradation of resolution.

4.4.1. Circuit description

With reference to Fig. 28, hereafter the FEC layout following the signals flow is described. The FEC receives 128 analog signals

through 6 flexible Kapton cables and the corresponding connectors. The input signals are very fast, with a rise time of < 1 ns. Therefore, to minimize the channel-to-channel cross-talk, the eight PASA circuits have to be very close to the input connectors. The analog to digital conversion and the digital processing are done by the ALTROs, which are connected to the corresponding PASAs with differential signals. It should be noticed that the PASA and ALTRO can also be interconnected in a single-ended mode. However, according to measurements, the noise increases by a factor two with respect to the differential-mode interconnection adopted in the FEC.

The data lines of the ALTROs are multiplexed, at the board level, through an LVCMOS bus. It features an asynchronous protocol, which is enhanced by a clocked block-transfer that provides a bandwidth of up to 200 MByte/s. The FEC is interfaced to the RCU through a bus based on the GTL technology, named ALTRO bus. At the board output the bus signals are translated from LVCMOS level to GTL level by bi-directional transceivers. The configuration, readout and test of the board are done via the GTL bus. Moreover, the FEC contains the BC, implemented in an FPGA, which provides the RCU with an independent access to the FEC via an I²C link. This secondary access is used to control the state of the voltage regulators and monitor the board activity, power supplies and temperature.

The board offers a number of test facilities. As an example a data pattern can be written in the ALTRO chip and read out back exercising the complete readout chain. The BC allows verifying the bus activities, the presence of the clock and the number of triggers received.

The ALTRO chips and the BC work synchronously under the master clock with a frequency up to 40 MHz. The ALTRO circuits usually perform the same operations concurrently, under the control of the RCU. However, the latter can also control a single channel at a time. This is performed in the configuration phase and for test purposes. The RCU broadcasts the trigger information to the individual FECs and controls the readout procedure. Both functions are implemented via the GTL bus.

Table 8
ALTRO key performance figures.

Power consumption	320 mW ^a
Max. readout bandwidth	300 MByte/s ^b
ADC resolution	10 bit
ENOB	9.7 bit ^{a,c}
N _{rms} (RMS noise)	0.35 LSB rms
DNL	< 0.2 LSB rms ^{a,c}
INL	< 0.8 LSB abs. ^{a,c}
SFDR	78 dB ^{a,c}
Cross-talk	0.05 LSB rms ^{a,d}

The three and four letter acronyms stand for: effective number of bits (ENOB), differential non-linearity (DNL) integral non-linearity (INL), and spurious-free dynamic range (SFDR).

^a $f_s = 10$ MHz, $R_{po1} = 90$ k Ω , internal.

^b $f_{RDO} = 60$ MHz.

^c $f_{in} = 960$ kHz, 1 V_{pp}.

^d Aggressor: $f_{in} = 960$ kHz, 1 V_{pp}, victim closed to 100 Ω .




Fig. 28. FEC layout. The components are mounted on both sides of the board. The figure shows the board topside with four PASAs, four ALTROs, one FPGA, the voltage regulators and some other minor components. On the bottom side are placed the other four PASAs and four ALTROs and, close to the readout bus connectors, the GTL transceivers.

4.4.2. Physical implementation

In order to match the position of the connectors on the chamber pad plane, the FEC has a width of 190 mm. Moreover, in order to fit into the available space its height and thickness are 170 mm and 14 mm, respectively.

The FEC Printed Circuit Board (PCB) contains four signal layers and four power layers (two supply layers with the corresponding ground layers). The power layers have essentially the same geometry. The duplication of the power and ground layers provides the following advantages: (1) it eases the implementation of controlled-impedance lines; (2) it reduces the voltage drop over the power layers; and (3) it reduces the noise produced by ground bouncing. From the power supply point of view the board is divided in three main sections: the PASA section, the ALTRO/ADC section and the digital section. Each power layer consists of three different power planes. The ALTRO/ADC and the digital planes are supplied with the same input voltage (+2.5 V), and are closed together at the input of the voltage regulators. The PASA plane is supplied at +3.3 V. The three ground planes (PASA ground, ALTRO/ADC ground and digital ground) are closed together with a pad, which is located upstream of the voltage regulators.

The FEC has a maximum power consumption of about 6 W. In order to minimize the heat transfer to the detector sensitive volume, the FECs are embedded in two copper plates cooled by water (see Section 5.4.1).

4.5. RCU

The readout control unit (RCU) is the central node in the readout and control networks. It acts as bridge between the different interfaces to the TPC (DAQ, Trigger and DCS) and its underlying electronics (FECs). In addition, the RCU provides core functionality to configure, trigger, readout, monitor and debug the FEE.

The RCU consists of a motherboard, which contains the main control circuit, and two daughter cards: the SIU (Source Interface Unit) and DCS cards. The RCU functions, which are mostly implemented on programmable logic devices, can be re-programmed remotely. This has proven to be an essential feature while commissioning and debugging the TPC FEE and its interplay with the backend systems. Moreover, the programmability of the RCU allows us to incorporate new features that might become relevant when operating the detector with high luminosity heavy ion beams. The SIU, which is the common interface card to the ALICE DAQ, is described in detail in Ref. [21]. The motherboard and DCS cards will be discussed in the following sub-sections.

4.5.1. RCU motherboard

The motherboard hosts the main FPGA, a Xilinx Virtex-II Pro (XC2VP7 [22,23]), the interface to the daughter cards, as well as the line drivers for the buses that connect to the FECs.

The architecture of the firmware for the FPGA is coarsely sketched in Fig. 20. The readout node provides a sequencer to readout the ALTRO channels in a programmed order and allows the program to skip empty channels to speed up the readout process (called 'sparse readout'). Special emphasis was given to the handling of erroneous bus transactions and their reporting, which is essential to operate a system with half a million channels in a stable fashion. In the course of implementing a fault-proof system, the RCU wraps the ALTRO event data (10-bit words) into a simple 32-bit format, adding marker bits to re-align to single channels in case of data corruption. This allows the system to decode the event data packet even when part of the system fails. The control node uses the slow control bus to read and monitor the BC values. It also implements an interrupt handler that can

turn off FECs within 100 μ s upon a severe error, e.g. a temporary short circuit produced by a highly ionizing particle.

The RCU's main FPGA is also responsible for distributing the clock signals (readout and sampling clock) to the FECs. They are derived from the global LHC clock (≈ 40 MHz) that is distributed via the LHC Timing Trigger and Control (TTC) system [24] to the FEE. The sampling clock can be selected to be either 2.5, 5, 10 or 20 MHz, and the RCU provides the necessary synchronization logic to keep the sampling clock phases of all RCUs equal. In this context, it should be mentioned that the phase of the trigger signal with respect to the sampling clock is measured for each event by the RCU, in LHC bunch crossing periods (25 ns), and is included in the trailer of the event data packet.

Besides the main FPGA, the RCU motherboard also hosts a flash FPGA (ACTEL ProAsic+ [25]) and some flash memory. This set of circuits is used to verify the configuration of the main FPGA and to possibly reconfigure it while it is operating, making the RCU a radiation tolerant circuit as further discussed in Section 4.7.

4.5.2. DCS board

The DCS board provides the interfaces to the ALICE Trigger and DCS systems [26,28]. It is based on an Altera Excalibur FPGA (EPXA1) with an embedded processor core (ARM 922T) [29] running a tailored version of Linux. The connection to the DCS system is established via a 10 Mbit/s Ethernet network interface, which is electrically adapted to run in a magnetic field. This architecture has proven to be extremely flexible and easy to maintain as it is built on widely supported hard- and software platforms.

The main application running on the DCS boards is a server application, called FeeServer, that provides communication channels to the DCS system for configuration and monitoring. It is described in more detail in Section 9.2.

4.6. Trigger subsystem

The ALICE trigger system (CTP) is based on three trigger levels [30]:

- L0:** The 'level zero' trigger pulse has a fixed latency of about 1.2 μ s with respect to the interaction.
- L1a/L1r:** Each L0 can be followed by a 'level one accept' pulse after a fixed latency of about 7.7 μ s with respect to the interaction time. If this is not the case, it is referred to as 'level one reject' and the trigger sequence has finished. If an L1a was issued, an asynchronous message containing basic event information (containing the event ID) will follow.
- L2a/L2r:** A third level trigger ('level two accept' or 'level two reject') completes the trigger sequence by deciding if the triggered event should be transferred from the FEE data buffers to the DAQ. This trigger level is dispatched as an asynchronous message after a minimum time of about 100 μ s, which corresponds to the TPC drift time, in order to ensure the completion of the TPC readout. It is relevant to mention that the rejection of events that are superimposed within the TPC drift time, can be implemented at this trigger level.

The synchronous trigger pulses (L0 and L1a) and the trigger messages (L1 message, L2a/L2r) are transmitted together with the LHC clock signal via the TTC optical fibers.

The TPC data acquisition is started either upon an L0 or upon an L1a, according to the configuration of the trigger detectors participating in the run. The readout process always starts after an L2a.

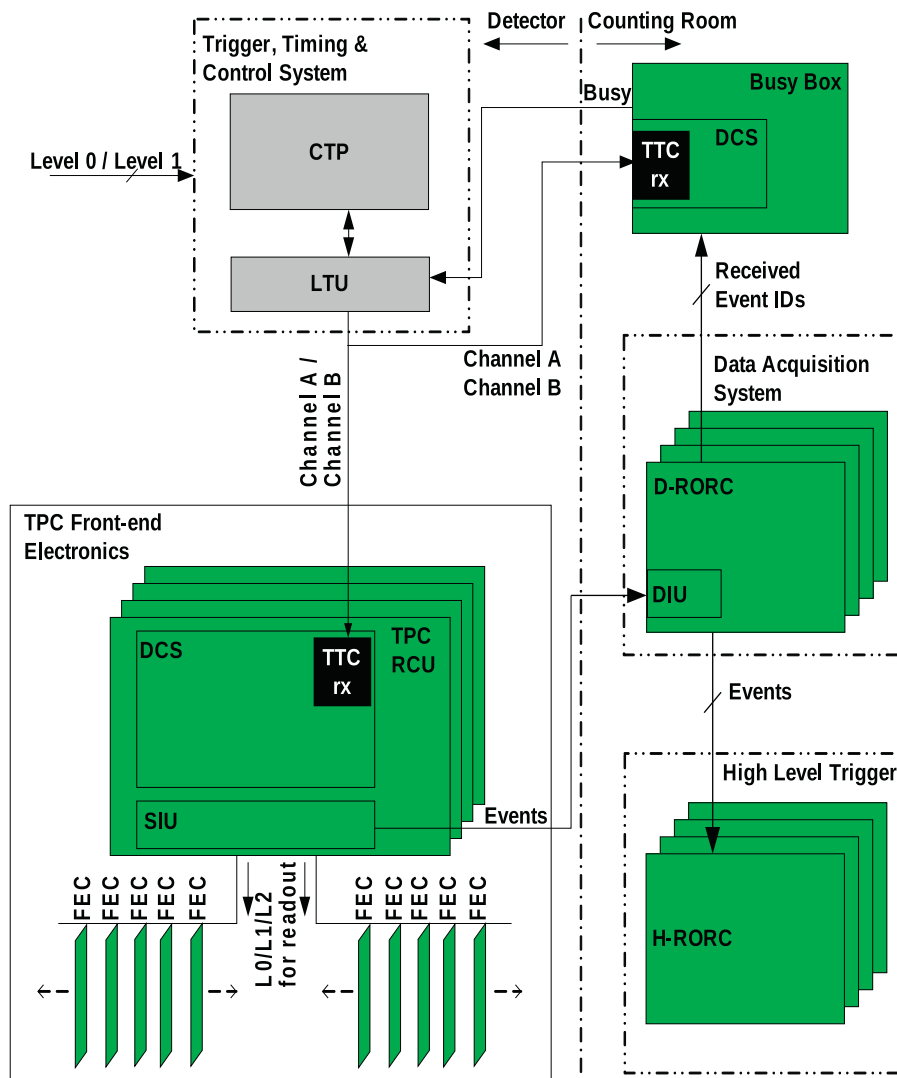


Fig. 29. Overview of the trigger subsystem.

The dead-time generated by the TPC has two contributions: detector dead time, i.e. the drift time, and the FEE dead time (event readout time). Whenever the TPC cannot process any further events, a signal (busy signal) is asserted to prevent the CTP from issuing subsequent triggers. The busy signal is generated by the BusyBox, which keeps track of the triggers issued and events received by the DAQ machines (via separate links to each of the 216 D-RORCs), as well as of the number of free MEB entries in the ALTROs (refer to Fig. 29).

4.7. Radiation tolerance

Exposure to energetic particles can produce instantaneous failures and, over time, the degrading of electronic components. The requirements for the TPC electronics in terms of radiation tolerance are twofold. On the one hand the estimated radiation is rather low both in terms of flux and dose (see below), which allows for the use of standard, radiation-soft technologies. On the other hand it has to be assured that the impact of rare (but existing) effects due to radiation are not affecting the overall performance or even causing irreversible damages to the system.

Simulations suggest that the total radiation load on the TPC will be < 13 Gy over 10 years, and the flux through the electronics less than about 800 particles/cm²s [31]. Nevertheless, for the design and implementation of the TPC FEE, the following three effects were carefully evaluated:

- **Total Ionization Dose (TID) effects:** TID is a cumulative effect, related to the damage of the semiconductor lattice, which causes slow gradual degradation of the device's performance.
- **Single Event Upset (SEU):** When an ionizing particle traverses the sensitive region of a memory cell it can change its logical state (0–1 or vice versa).
- **Single Event Latch-up (SEL):** The ionized track of a particle can form a conducting path through the substrate of the semiconductor creating a short between the supply rails.

Concerning the TID effects, all components were qualified to withstand a dose corresponding to 150 years of LHC operation. The measures to protect the FEE components against SEUs and SELs were chosen according to their potential impact on the system. Two examples, the protection of the ALTRO chips and the RCU FPGA, are hereunder discussed in more detail.

4.7.1. SEU

ALTRO: In the ALTRO all control state machines are protected against SEU by Hamming coding that implements an algorithm to recover from single bit and to detect double bit errors. In particular this prevents electrical conflicts of two ALTRO chips on the I/O bus. As this virtually makes them fail-safe, the most severe problem that can still occur is a corruption of the configuration registers content. Moreover, data could also be affected, but the estimated bit error rate does not justify any protection on this level. Combining the simulated particle flux with the SEU cross-sections measured for the ALTRO chip, the corresponding mean time between failures (MTBF) have been determined as listed in Table 9.

RCU: In the RCU there is a single device to be protected against SEUs, the main FPGA, which is based on SRAM technology. In this context, it should be mentioned that the hardware resources needed for the implementation of the RCU functions are featured only in large SRAM programmable devices. However, these devices store their configuration, which defines their logical function, inside a radiation sensitive SRAM. If an SEU occurs in this configuration memory, it might cause a circuit malfunction, which, however, can be corrected by rewriting the originally stored value to the memory cell. Partial reconfiguration is a feature offered by several Xilinx FPGAs, which allows to reconfigure a subset of the configuration memory without interrupting the operation of the device. This has been implemented in the RCU system in order to detect and correct SEUs

Table 9
MTBF of the internal elements of the ALTRO chip.

Part	MTBF
Registers	36 h
PMEM	168 min
MEB	42 min
FSM ^a	58 days

The MTBF values are quoted for a readout partition with 25 FECs.

^a It refers to the occurrence of SEUs in a finite state machine (FSM), which is recovered by the Hamming protection.

[32]. As shown in Fig. 30, the main parts of the verification and reconfiguration circuit are an Actel flash-based FPGA (auxiliary FPGA), and a flash-memory device.

4.7.2. SEL

The radiation tests carried out on the FEE components have never shown the occurrence of SELs. However, the FECs were designed before the radiation tests were completed and, therefore, were protected against the occurrence of short circuits induced by SEL as a precautionary measure. The BC continuously monitor the FEC's power consumption and asserts an interrupt if a programmable threshold is exceeded. Whenever the RCU detects an interrupt it turns off the corresponding FEC, removes it from the readout process, and notifies the event to the DCS.

4.8. Testing procedure

The testing and qualification procedure for the electronics was implemented in two consecutive stages: first the individual custom made chips (ALTROs and PASAs) and then the assembled FECs.

After having been packaged, all chips had to go through a thorough test that included burn-in testing and functional testing based on a semiautomatic set-up using a robot [33]. The digital and analog parts were subject to different acceptance tests. The digital parts were required to be 100% functional. The analog parts were allowed a small variation in the most critical parameters (conversion gain, peaking time and output DC offset). The requirements and sources of failures are displayed in Tables 10 and 11, which also include the numbers as obtained for a later production of the circuits mainly for STAR in 2006. The yields for ALTROs and PASAs are 84% and 98%, respectively. The PASA yields were so good, that we were able to apply very stringent selection criteria as listed in Table 11.

Once the parts had been qualified, the assembled FEC was tested to reject clear electrical faults (e.g. shorts) and to characterize the interplay of the components. The corresponding yields are shown in Table 12.

During all tests the components were marked and their characteristics filled into a database. For the FECs the ID is also stored in an EEPROM on the card itself.

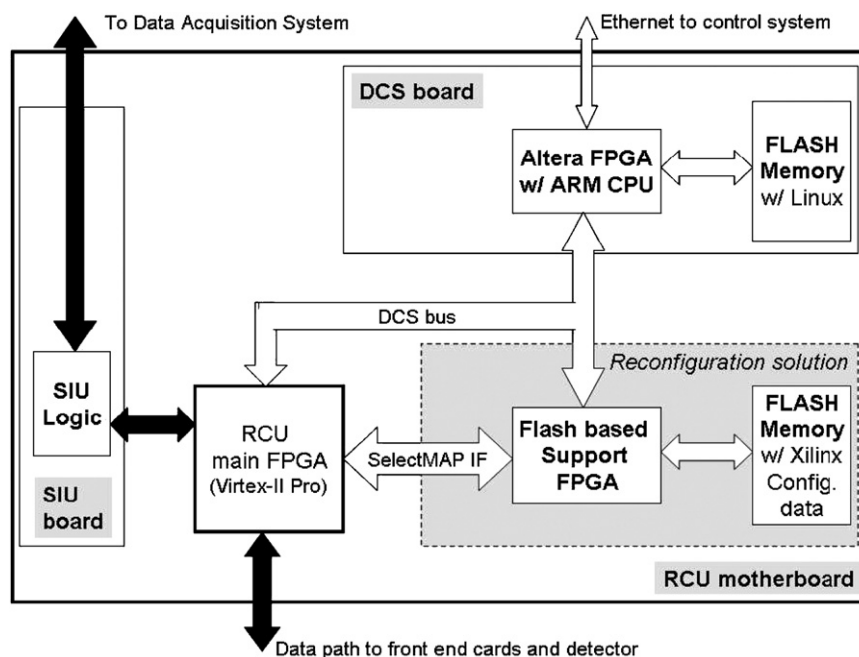


Fig. 30. RCU block diagram, emphasizing the circuit for active partial reconfiguration at the bottom right. The data path is given by black arrows.

Table 10
Failures in the ALTROs.

Failure	Frequency			
	ALICE		STAR	
None	41 297	84.1%	14 273	
Power	2307	4.7%	674	4.1%
Register	1032	2.1%	190	1.2%
PMEM	621	1.3%	199	1.2%
MEB	2203	4.5%	633	3.8%
DSP	753	1.5%	241	1.5%
ADC	712	1.4%	228	1.4%
Misc.	153	0.3%	49	0.3%
Not tested	48	0.1%	2	0.01%
Total	49 127	100.0%	16 489	100.0%

Tests were performed in the listed order and the procedure was stopped upon the first detected failure.

Table 11
Acceptance levels for the PASAs as chosen for the ALICE TPC.

Not accepted	Frequency			
	ALICE		STAR	
None	40 938	85.9%	20954	91.9%
Power	776	1.6%	72	0.3%
Conv. gain (5%)	1009	2.1%	179	0.8%
Peaking time (6%)	1408	3.0%	153	0.7%
Outp. offset (50 mV)	5 28	11.4%	2122	9.3%
Total	47 637	100.0%	22 795	100.0%

Multiple counting occurs but in the case of too high power consumption where the testing procedure was aborted.

Table 12
Failures in the FECs.

Failure	Frequency	
None	4320	90.0%
Any	380	7.9%
Damaged PASAs	80	1.7%
Damaged ALTROs	40	0.8%
Improper placement or soldering of connectors	35	0.7%
PCB traces	5	0.1%
Improper soldering of IC pins	107	2.2%
Defective, misplaced or missing passive components	73	1.5%
Unknown	140	2.9%
Total	4800	100.0%

Multiple counting occurs. The numbers are biased by the ‘unknown’ entry, which refers to FECs that fail the test but still have to be debugged/repared, as the FECs with most obvious errors were repaired first.

5. Cooling and temperature stabilization system

5.1. Overview

In this section we describe the water based cooling and temperature stabilization system of the ALICE TPC. We substantiate the need for a temperature stabilization system in addition to the cooling (heat removal) of the front-end electronics. The principle of the leakless operation of the cooling system is briefly introduced and the cooling strategy is outlined. In Section 5.6 the experience gained during the commissioning of the cooling

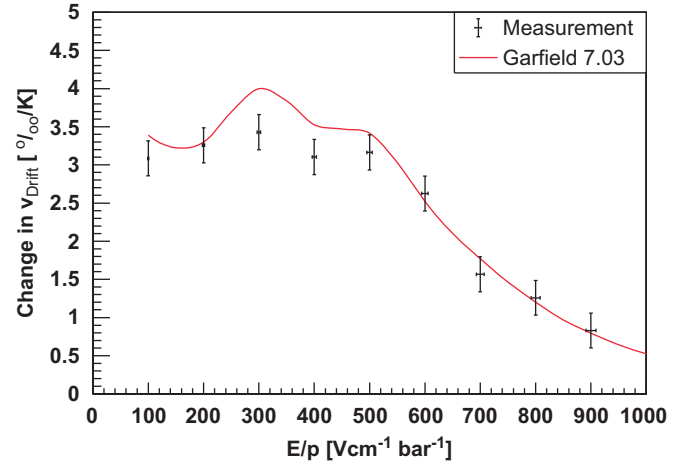


Fig. 31. Relative change of the drift velocity with temperature as a function of the drift field for the binary mixture Ne–CO₂ in proportion 90–10 [34,35].

systems (2006–2008) is reported. The achieved temperature homogeneity during these first runs is also discussed.

5.2. The necessity for uniform temperatures

The design goal for the temperature stability and homogeneity within the TPC drift volume is $\Delta T < 0.1^\circ\text{C}$ [3]. This value is a consequence of our particular gas choice, a mixture of Ne–CO₂–N₂.

Fig. 31 shows the relative change in drift velocity with temperature as a function of the drift field [34,35]. As can be seen from the graph, the change in drift velocity is about 0.35%/K at 400 V/cm. This means that a temperature difference larger than 0.1 K over the full drift length of 250 cm results in longitudinal position variations Δz of the order of 1 mm at $v_{\text{drift}} = 2.65 \text{ cm}/\mu\text{s}$. This exceeds the internal resolution of the readout chambers.

5.2.1. Heat load and computational fluid dynamics calculations

The main heat contribution stems from the FECs, which are connected via short (8.2 cm) flexible cables to the cathode pad plane of the readout chambers. The 2×18 sectors, each being equipped with 121 FEC cards, dissipate a total of 28 kW. This heat load has to be removed by the FEC-cooling circuits. The bus bars, providing the low voltage power to the FEC, are integrated into the SSW spokes and dissipate a total of about 0.54 kW.

Another important heat source affecting the TPC gas is the power produced by the four field-cage resistor rods. While the power is relatively small (8 W/rod), it would, without counter-measures, be dissipated directly into the gas volume.

Other heat sources are neighboring detectors, namely the Inner Tracking System (ITS) inside and the Transition Radiation Detector (TRD) outside of the TPC.

In the context of the optimization of the ventilation scheme inside the ALICE L3 magnet, a Computational Fluid Dynamics (CFD) study has been carried out to estimate the residual heat distribution within the L3 magnet [36]. The resulting temperature profile within the L3 magnet shows a variation within 20 and 30 °C which leads to a temperature inhomogeneity of up to 5 K inside the TPC volume. A 90% cooling efficiency of the various detector cooling systems has been assumed for this study leading to a dissipated heat of 17 kW. It was shown that, irrespective of the detailed layout of the ventilation scheme, a temperature gradient of about 5 K develops across the vertical dimension of the TPC. This study clearly demonstrates that




Fig. 32. Principle of underpressure operation.

the TPC fiducial volume needs additional thermal shielding and compensation of temperature gradients.

5.3. Principle of underpressure cooling

The principle of underpressure cooling [37] is depicted in Fig. 32. The cooling-liquid circuit is a closed circuit, which allows to operate all or part of the cooling lines below atmospheric pressure. The cooling-liquid tank is kept at underpressure, which, by the proper choice of length and diameter of the return pipes and of the circulation-pump output pressure, ensures that the water pressure inside the detector is below atmospheric pressure. This has the obvious advantage of an active protection against the occurrence of leaks.

In our case, the space constraints due to the extremely dense front-end readout does not allow space-consuming high-pressure certified fittings. Therefore, we chose simple silicon hoses without any special lock mechanism to couple to the copper tube of the card cooling envelope (see Section 5.4.1). Though the connection between the silicon hose and the copper tube has been tested to hold overpressure of 2.5 bar over an extended period (24 h), it is mechanically fragile, e.g., against tears or cuts. These considerations led us to the choice of the sub-atmospheric ‘leakless’ technology for our cooling circuits. In the following paragraphs, we will give a short description of this, a more detailed description can be found in Ref. [37].

An apparent disadvantage of the sub-atmospheric ‘leakless’ technology is the limited range of operation ($p < 1$ bar). This implies that the allowed pressure loss, Δp_{loss} , in the detector is rather limited. The situation is further aggravated, in our case, due to the height difference of about 8 m between the highest inlet

and the cooling plant. In addition the height difference between the inlets of sector 4 and 13 located at the top and bottom of the TPC, respectively, is 5 m. While the input pressure at each inlet can be adjusted independently via balancing valves, this has not been foreseen for the return lines. Thus, all detectors ‘see’ a combination of the reservoir pressure, the hydrostatic pressure and the pressure loss in the return pipes. This can result in very low pressure values in some of the return lines, which might cause cavitation phenomena (see Section 5.6.3).

5.4. TPC cooling plants

The TPC main plant [38], schematically depicted in Fig. 32, consists of a 1200 l reservoir, large enough to buffer most of the water in the installation, a pump, a heat exchanger connected to the CERN mixed water network, a supply manifold and a return manifold. A total of 48 circuits are connected on these manifolds. All 48 circuits can be temperature and flow adjusted. Temperature regulation is done via individual heaters on the supply manifold. These heaters are controlled via a Proportional–Integral–Derivative (PID) controller using a temperature sensor immediately downstream of the heater as a feedback signal. The precision of the regulation is 0.1 K. The effective power of the heaters is ≈ 3 kW, which allows a temperature swing of about 2 K at a flow of 20 l/min. To allow a larger flow and/or temperature swing for the chamber body, the heaters for these loops were upgraded to 10 kW (see Sections 5.5 and 5.6).

Flow adjustment was done manually by adjusting valves during the commissioning of the system, in order to ensure that all circuits got the right flow despite of the varying pipe lengths and heights between the plant and the detector. The pressure at the inlet of the reservoir is controlled by a so-called back-pressure valve. This valve had not been foreseen in the original design of the plant. However, during commissioning (see Section 5.6) the necessity of being able to vary the return line pressure became evident. A PID controller adjusts the pump speed via a frequency inverter in order to keep a constant pressure in the supply manifold independent of the number of circuits in service. Another PID controller adjusts the flow of mixed water via a 3-way valve in order to maintain a constant temperature at the heat-exchanger outlet. All regulation loops and the plant control is done by a Programmable Logical Controller (PLC). A comprehensive description of the TPC cooling plant, its operational parameters as well as the alarm handling is presented in Ref. [38].

A conceptually similar, but much smaller plant ($V_{\text{reservoir}} = 80$ l) is used for the resistor-rod cooling. It has, in addition, an ion-exchanger filter to purify the circulation water. The conductivity of both the supply and return water is measured. An alarm is raised if its value exceeds a threshold value. For this plant the pressure of the four return lines was adjusted individually by introducing fixed restrictions. They compensate the different pressure drops in the return lines due to the hydrostatic pressure given by the different height of the resistor rods.

5.4.1. Cooling circuits

A schematic overview of the different cooling units of the TPC, consisting of circuits to remove the heat produced by (i) the front-end electronics, bus bars and resistor rod and (ii) screens to define iso-thermal surfaces, is shown in Fig. 33.

Cooling and temperature stabilization of the TPC is provided via 60 individual loops which are supplied by three different cooling plants. The main TPC cooling plant supplies:

- 2×18 loops for the front-end electronics cooling at the A- and C-side, respectively;

- 2×2 loops for the bus bar and cover cooling. The two loops at each side furnish the top and bottom half of the bus bars and covers, respectively;
- 2×2 loops for the chamber body cooling. The two loops at each side supply the top and bottom half of the chamber bodies, respectively;
- 2×1 loops for the inner thermal screen, which separates the TPC from the ITS services. Each of the loops is split after the balancing valve and supplies the upper and lower manifold of the screen panel; and
- 1 loop supplies the resistor-rod heat exchanger.

Another plant, the resistor rod-cooling plant, supplies the 2×2 loops for the resistor-rod cooling. On each side of the TPC, the inner and the outer resistor rods have their own cooling circuits. A separate plant for the resistor-rod cooling is needed because of the special demand on the purity of the cooling water.

The outer thermal screen, decoupling the TPC and the TRD thermally, is supplied by nine independent cooling circuits. The Al-panels of the screen require deionized water, which is provided by the TRD cooling plant.

All loops of the main TPC plant are independent from each other in the sense that the flow and the temperature (within limits) can be regulated independently. The resistor-rod lines have a common temperature set point and individual flow regulation via balancing valves. The outer thermal-screen panels are supplied by a common temperature water flow. Regulation, e.g. between top and bottom panels, is possible only via different water flow settings.

FEC and Al-body cooling: Cooling to the FECs is provided sector by sector. As shown in Fig. 34 (top), six FECs are grouped together and are connected to the sector manifold. It has been estimated [38] that a flow of about 0.5 l/min per group of six FECs, i.e. 10 l/min in total per sector, is sufficient to cool the electronics. A photograph of the FEC and its cooling envelope is shown in Fig. 34 (bottom); the copper envelope is flipped open and exhibits the FEC inside. As can be seen, a 5 mm copper tube for the cooling water is soldered to one of the copper plates. The two plates are held together via six screws, which at

the same time serve as heat bridges between the two plates. The main heat load of the FECs is taken up by the cooled copper envelopes. However, to stabilize the temperature and to absorb residual heat transfer, e.g., by the Kapton cables, auxiliary cooling circuits have been integrated into the Al-body of the chambers (see Section 3).

The busbars, supplying the low voltage to the FEC, run along the spokes of the service support wheel. To remove the resistive heat of the busbars, additional, water-cooled copper bars, which are in thermal contact with the busbars, are integrated into the service support wheel spokes. The busbars and the TPC covers are supplied by the same manifold.

Thermal screens. The TPC vessel is, to a certain extent, shielded against outside heat sources or temperature variations by thermal screens. Fig. 35 (right) shows an example of one of the outer heat-screen panels, which is located between TPC and TRD. The panels are installed into the space frame structure, which also defines the size of an individual panel. The outer heat screen design follows the space frame structure, i.e. 18 ‘super-panels’ in φ are subdivided into five basic panels in z-direction. A design restriction for the panel was the required short radiation length, which led to the choice of 0.5 mm thick Al-sheets and Al-tubes of 10 mm outer and 6 mm inner diameter, respectively. The corresponding average radiation length X/X_0 is $\approx 1.2\%$. The upper nine super-panels are supplied by six loops (two loops per three sectors), the lower panels are supplied by three loops (one loop per three sectors).

The inner thermal screen shields against the heat of the ITS services. It does not cover the central drum of the TPC and is hence outside of the acceptance of the TRD. On each side of the TPC the screen consists of 12 double-wall stainless steel panels. An example of a panel is shown in Fig. 35 (left).

A third thermal screen is defined by the IROC and OROC covers: each readout chamber is covered by a 1 mm thick copper sheet, onto which 3×1 mm cooling coils are soldered. The covers are fixed to the service support wheel.

Resistor rod cooling: Though not large in quantity (4×8 W), the heat produced by the field-cage resistor chains is dissipated directly into the fiducial volume of the TPC. To avoid this direct




Fig. 33. Schematic view of the various TPC cooling elements.




Fig. 34. Photograph of a FEC with its cooling envelope. The top part of the figure shows a sketch of the routing of the cooling pipes connecting six FECs.




Fig. 35. Inner (left) and outer (right) heat screen panel.




Fig. 36. Schematic overview of a resistor-rod cooling loop.

heating, a cooling scheme for the resistor rods has been introduced. A schematic drawing is shown in Fig. 36. The cooling of the 166 pairs of $15\text{ M}\Omega$ resistors is provided by water flowing through ceramic tubes of 9 mm outer and 3 mm inner diameter, respectively. Copper plates serve both as heat bridges to transfer the heat from the resistors to the ceramic tubes and as terminal for the field strips. To avoid electrical breakthrough between the high-voltage terminal connecting to the central membrane and the cooling water, the water column is at 100 kV and ground potential at the two ends of the ceramic tube, respectively. This requires ultra-pure water with a resistivity close to $18\text{ M}\Omega\text{m}$ in order to keep the current through the water around $3\text{ }\mu\text{A}$. This value is sufficiently small compared to the resistor rod current of nominally $80.05\text{ }\mu\text{A}$ at 100 kV.

The water flow through each of the four resistor-rod cooling tubes is only about 0.5 l/min, which, however, corresponds to an almost 15-fold volume exchange per minute of the ceramic cooling tubes. This relatively small flow requires that the temperature of the cooling water is defined precisely at the TPC: owing to the long distance from the cooling plant to the TPC the water with this small amount of flow is likely to pick up the ambient temperature on its way to the TPC, which might be variable and different from the desired gas operating temperature. We have therefore installed four heat exchangers at the input of each of the resistor rod cooling tubes. The primary heat-exchanger circuit is supplied with a high flow ($> 101\text{ l/min}$) and is thus largely insensitive to ambient temperature changes.

5.5. Cooling strategy

The stringent requirement on the TPC temperature stability and homogeneity necessitate an elaborate cooling approach. The strategy to stabilize the TPC temperature was validated experimentally in a small test setup and is described in Refs. [39,40]. Basically, it was demonstrated that thermal neutrality of the FECs can be achieved by ‘undercooling’, i.e. the cooling water is injected into the FEC cooling loops several degrees below the desired TPC operating gas temperature. Furthermore, it was shown that the body cooling loops are required to establish stable equilibrium values of the gas operating temperature.

5.6. Commissioning of the cooling system

5.6.1. Test with mock-up sectors

The reaction of the leakless plant to the sudden appearance of a major leak, e.g. the inadvertent removal of a silicon hose with water circulation on, has been tested employing specially made mock-up sectors. The two mock-up sectors used have approximately the same properties as a real sector in terms of silicon-hose tubing, i.e. the pressure distribution of the water is similar to a real sector. The most important result of these tests with mock-up sectors is summarized in Fig. 37: on incidence of a large leak (e.g., a silicon hose open) at $t=15\text{ s}$ the system is able to keep the detector cooling sub-atmospheric only for a limited time ($< 15\text{ s}$). Thereafter the pressure, measured at the inlet of the detector, rises above atmospheric pressure and water spills out. The reason the system goes beyond 1000 mbar is due to the fact that air bubbles sucked into the system clog the return line, increasing the resistance for the water flow. Therefore, we have installed pressure sensors at all inlets. The data from the pressure sensors are fed to a dedicated PLC, which, on detecting a pressure value higher than a preset value, sends a signal to the PLC of the cooling plant upon which the plant circulations is stopped, water is collected in the tank and underpressure is re-established in all loops.

5.6.2. Startup procedures and operation

At the very beginning of the cooling-system commissioning it was found that water circulation in many loops could not be established without violating the paradigm of the sub-atmospheric cooling, i.e. that the pressure in the detector should not




Fig. 37. Pressure as function of time after a forced leak at $t=15$ s in one sector. The occurrence of the leak results in a short, immediate pressure spike at which, however, no water is spilled. This appears only after about 10 s, when the pressure raises permanently above atmospheric.

exceed the atmospheric value. The reason for this has been traced back to the routing of both the supply and return lines which were laid out with siphons. In a tube with an air-water mixture, as it exists during the startup of the plant when the tubes are initially air-filled, any siphon reduces the over- or underpressure relative to atmospheric. Hence, in the presence of many or large siphons the limited pressure range of an under-pressure system is not sufficient to start the circulation, i.e. to overcome the initial hydrostatic pressure imbalance. Only a careful rerouting of the supply and return lines avoiding siphons as much as possible finally allowed startup of the circulation in all loops. In addition, many of the loops reach circulation only if the pressure in the tank is lowered to at least 350 mbar.

While the startup requires a tank pressure as low as possible, once a steady circulation has been established the pressure in the tank can be raised to 550 mbar or higher. The necessity to raise the tank pressure to a level of at least 550 mbar is discussed below in Section 5.6.3. During the commissioning runs in 2008/9 the cooling systems for resistor rods and FECs have been in operation over an extended period (several month) without problems.

5.6.3. Cavitation problem

As already mentioned above, a low value of the reservoir pressure (350 mbar), together with a large negative hydrostatic pressure difference (850 mbar for the highest sectors) in the return line, gives, by calculation, negative pressure values which produce instabilities. Physically, this leads to turbulences in the return line and to a phenomenon called cavitation. Cavitation is defined as the phenomenon of formation of vapor bubbles of a flowing liquid in a region where the pressure of the liquid falls below its vapor pressure (23 mbar at 22 °C). The collapsing bubbles produce shock waves, which might be strong enough to entail significant damage. In fact, at low-pressure operation a ‘knocking’ noise had been observed in the return lines. To avoid possible damage to the tubing the plant has to be operated at tank pressures above 550 mbar, a value at which the dynamic effects

seem to be reduced and audible signs of cavitation cease. To be able to switch quickly from low- to high-pressure values of the tank a ‘back-pressure’ valve has been installed at the inlet of the reservoir which regulates the pressure in the return lines.

For the resistor-rod cooling plant, cavitation constitutes a particularly severe problem owing to the ceramic cooling tubes which are inside the TPC gas volume. Individually adjusted restrictions, introduced in the return lines, raise the pressure to a level above the cavitation threat. To control the pressure distribution in the ceramic tubes, pressure sensors were installed both at the inlet and outlet of the resistor-rod cooling tubes.

5.7. Temperature monitoring system

To monitor the temperature distribution of the TPC, 496 PT1000 sensors are mounted both inside and outside of the gas volume of the TPC. In addition to sensors covering the outside of the Inner and Outer Field Cage containment vessels, several sensors are mounted onto each IROC and OROC. For each sector, sensors measure the cooling water inlet and outlet temperature. Several sensors (2×18) are attached to a circular skirt inside the gas volume. Additional temperature sensors on the front-end electronic cards (one sensor for each of the 4356 FECs) complete the monitoring system. A comprehensive description of the temperature monitoring system including the sensors’ positions, its calibration and its readout system is given in Ref. [41].

5.7.1. Temperature profile and homogenization

Fig. 38 shows the temperatures as measured with the skirt sensors, which are located inside the gas volume, as a function of time. The temperature data were taken during the detector commissioning run in fall 2008. The plot demonstrates the sensitivity of the TPC gas volume to the heat load from other detectors. During the beginning of the measurement period the TOF detector, surrounding the TPC, had been continuously running. Later it was switched off each night for 8 h. The influence on the temperature inside the TPC is clearly visible by temperature excursions of the order of 0.2 K. Even though the TOF detector has been on for $\frac{2}{3}$ of the day no stable temperature (i.e. equilibrium) is reached. Measurements have shown that the temperature relaxation times are of the order of 16 h due to the large mass involved.

The histogram of the skirt temperature sensor distribution is shown in **Fig. 39**. The temperatures were sampled over a period of 24 h with stable environmental conditions. The RMS value of the histogram is below 50 mK showing that the desired temperature homogenization of the TPC gas volume below 0.1 K is within reach.




Fig. 38. Temperatures measured with the skirt PT1000 sensors as a function of time. Each of the curves represents one of the skirt sensors.

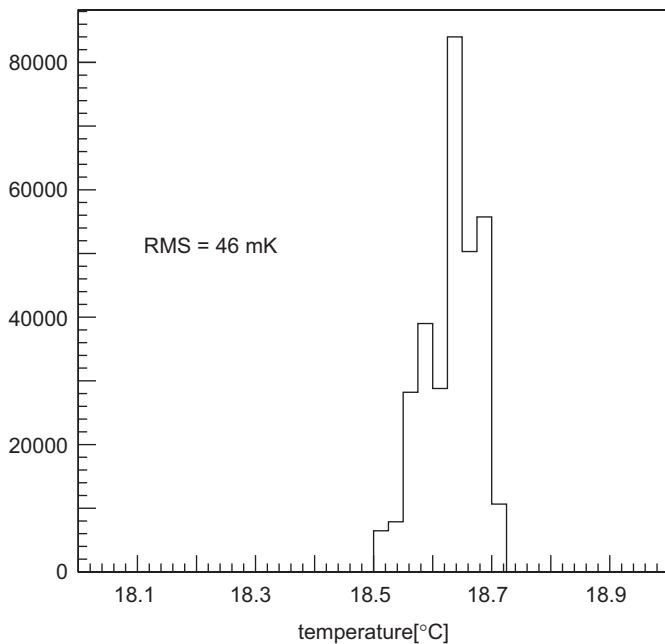


Fig. 39. Temperatures distribution measured with the skirt PT1000 sensors.

It should, however, be noted that the skirt temperature probes represent the x – y gradient close to the readout plane. The optimization of the distribution close to the sectors is comparatively straightforward: it involves the proper ‘overcooling’ of the readout cards and a slight top–bottom difference in the chamber body cooling. The feedback information from the sensors, both on the readout chamber and on the skirts, is relatively fast, i.e. within several hours an equilibrium value is established. The temperature homogenization over the full TPC volume via flow and temperature adjustment of the thermal screens is considerably more involved: since the sensors are outside of the TPC vessel and are, in addition, insulated from the gas volume by the CO_2 volume they thus reveal limited information about the temperature inside the TPC. A better understanding of the gradients over the full TPC volume requires an analysis of straight, laser induced tracks (see Section 7).

Above we have described a complex cooling and temperature homogenization system, which involved altogether 60 individual, adjustable cooling circuits. Conceptually, the circuits have the tasks, to remove heat, i.e. cool the front-end electronics, and to define iso-thermal surfaces around the TPC in order to provide an as much as possible gradient-free gas volume. Overall, we have been able to reduce the temperature gradients inside the TPC to a $\text{RMS} < 0.05^\circ\text{C}$, as it is required for a full exploitation of the TPC internal position resolution.

6. Gas and gas system

The TPC is essentially a 90 m^3 volume filled with gas, where the gas is the detecting medium. The detector performance depends crucially on the gas choice, stability and quality, since these influence the charge transport in the drift volume and the amplification processes in the readout chambers. The choice of the gas composition is constrained by a set of performance requirements and boundary conditions, and in turn the selected gas mixture determines the performance of the detector and conditions various aspects of its design, from the shaping time of the front-end electronics to the temperature uniformity of the gas in the detector. In particular, the gas system is designed to fulfill the requirements derived from the gas choice and the expected

performance of the detector. In the next section we discuss the selection procedure that led to the choice of a $\text{Ne}\text{--}\text{CO}_2\text{--}\text{N}_2$ gas mixture and the implications of this choice. We then describe the gas system that injects, circulates and cleans this gas.

6.1. Gas choice

The selection of both the noble gas and the quencher was made by a process of elimination rather than choosing the gas by its merits. Between Ar and Ne, the former, although providing larger primary statistics, was discarded because of material-budget considerations and a slow ion mobility. In particular, Argon would substantially enhance space-charge effects in the drift volume of the TPC due to its relatively slow ion drift velocity. As far as the quencher is concerned, hydrocarbons were excluded due to aging considerations (Malter currents would set in after about 1 year of operation with heavy-ion beams). CF_4 , on the other hand, presented many concerns with material compatibility at the time of the system design. Therefore, CO_2 was chosen. Since the maximum drift field in the field cage was designed to be 400 V/cm , and the maximum drift time $94\ \mu\text{s}$, the composition results in 10% CO_2 in Ne [42]. During the prototyping phase, a further 5% of N_2 was added to the mixture [43]. This addition reduces the drift velocity at the nominal field by about 5%, but it provides a more stable operation of the readout chambers at high gain. This also eliminates the problem that N_2 could build up in the gas due to small leaks. The changing nitrogen content would then affect the detector performance since it cannot be removed by the cleaning agents of the gas system. Excited states of neon have energies around 17 eV, for which the quenching capabilities of CO_2 are poor. More CO_2 in the mixture would rapidly decrease the drift velocity. N_2 , on the other hand, presents a slightly higher ionization cross-section at these energies and so helps to quench the Ne, and affects the drift velocity modestly. The resulting mixture is also less sensitive to the exact composition. The slow proton production due to neutron bombardment of N_2 molecules has been shown in simulations to be reasonably small. The main implication of the choice of a $\text{Ne}\text{--}\text{CO}_2\text{--}\text{N}_2$ gas mixture for the gas system is the necessity of monitoring and controlling a ternary mixture. Fig. 40 shows the drift velocity and the longitudinal and transverse diffusion coefficients calculated with Magboltz [44,45], as a function of the electric field for the mixtures with and without N_2 . In both cases, the drift velocity is not saturated at the nominal field, thus making it very sensitive to gas density fluctuations and to the exact electric field.

6.1.1. Implications of the gas choice

Fig. 41 shows the drift velocity dependence on temperature at 400 V/cm . To achieve a position resolution of order $200\ \mu\text{m}$ [3], the necessary temperature uniformity in the drift volume is 0.1 K , a stringent requirement which drives the cooling strategy (see Section 5). The fluctuations due to ambient pressure variations, which are followed by the pressure regulation of the gas system, are corrected for.

In addition, the drift velocity changes by -6.4% per % change in CO_2 concentration and by -1% per % in N_2 , while the gain dependence is 15% per % change in CO_2 and 6% per % change in N_2 . To keep the drift velocity constant at the 10^{-4} level it would be necessary to control the CO_2 concentration to better than 0.01% , which is beyond the precision of current mass-flow controllers. Therefore, in addition to the laser system (see Section 7), the gas system is equipped with diagnostic tools to measure the gas composition: a gas chromatograph with a thermal conductivity detector and a high precision drift velocity monitor.




Fig. 40. Drift velocity (top) and longitudinal and transverse diffusion coefficients (bottom) as a function of the electric field for the Ne-CO₂ (dashed lines) and the Ne-CO₂-N₂ (solid lines) mixtures calculated with Magboltz at 750 Torr and 20 °C. While the diffusion coefficients do not change with the addition of N₂, the drift velocity decreases by 5% at the nominal field of 400 V/cm.




Fig. 41. Dependence of the drift velocity of the ternary mixture on the temperature as calculated with the Magboltz simulation package at 400 V/cm and 750 Torr.

The attachment coefficient of electrons to O₂ is greatly enhanced in the presence of CO₂, reaching 400 bar⁻² μs^{-1} . This is because negatively ionized excited oxygen molecules rapidly decay into the ground ionic state by energy-transfer collisions

with CO₂ molecules. An electron drifting over 2.5 m in this gas, contaminated with 5 ppm O₂, has a 25% chance to get attached. Therefore, besides the tightness of the detector itself, the gas system and pipe work must be certified to be leak free, and provisions to remove oxygen from the gas are necessary. This issue, combined with the large volume of the system, also determines the total gas flow through the detector. A reasonable choice is to flush the detector volume at a rate close to 5 times a day. The oxygen contamination is monitored by constantly sampling the gas with an oxygen analyzer. The water content and the CO₂ concentration are also analyzed with appropriate sensors in the same chain.

Finally, since neon is a high cost gas, a CO₂ absorber system is implemented in the gas system, so that the filling of the detector with the mixture can be done at no waste of the main gas, as explained below.

In addition, the gas is humidified, in a controlled way, with 50 ppm of water vapor, in order to avoid charging up of insulators near the wires of the readout chambers. Other than decreasing the drift velocity by 0.16%, the effect on the gas properties is negligible.

6.2. Description of the gas system

Most of the gas systems and control programs of the LHC experiments are designed and built under a common modular scheme coordinated by CERN infrastructure groups. Each function of the gas system, such as the gas mixing, the circulation, the cleaning, the analysis, etc., is integrated into a logical module which usually corresponds to one or more racks. A Programmable Logic Controller (PLC) runs the system by executing actions, like opening a valve, or a sequence of actions, regulating devices according to set points, like flows or pressures, reacting to alarms, reading analog and digital values, including Profibus networks, and publishing information to a user interface. The user interface [46] allows the operator to change the state of the gas system, for example from fill to run states, and to act on individual modules or on single components of the system. It also logs the various data points into a data base for inspection of trends. A set of recipes, or configuration files, specific for each module, is loaded into the PLC for determination of set points, limits, regulation parameters, timers, etc., which the operator can edit and reload at any time from the user interface.

The racks are distributed in a surface building, in the plug of the shaft (a shielded platform just above the cavern), and in the cavern, according to their functions and specifications.

6.2.1. Configuration

The gas system is a gas circulation loop with injection, distribution, regulation and other tasks distributed in modules located at different elevations in the ALICE hall (see Fig. 42). The functionality of the loop is schematically depicted in Fig. 43. Gas is circulated through the detector by a compressor module, which extracts the gas from the TPC and fills a high-pressure buffer volume for gas storage. A regulated bypass proportional valve (B in the figure) re-injects part of the compressed gas back into the compressor inlet, as part of a feedback loop to regulate the pressure inside the TPC. This valve is driven by a pressure sensor installed at the detector, which determines the operating pressure set point with respect to atmosphere. The TPC overpressure is thus regulated to 0.4 mbar. The high-pressure buffer, a 1 m³ tank at 2–4 bar overpressure, stores gas that can be delivered to the detector in case of an increase of the atmospheric pressure, or can accept gas from the detector when the ambient pressure decreases.




Fig. 42. Schematic view of the distribution of the various modules of the TPC gas system on the surface, on the plug (a platform in the shaft of the experiment, just above the cavern) and in the cavern.




Fig. 43. Simplified diagram of the gas system loop, showing its main modules and some of the main elements used for operating and regulating the system in various modes. Bold lines indicate gas pipes that operate at high pressure (about 3 bar).

Once at a few bars overpressure, the gas, depending on the system mode, flows through several modules, located on the surface, or is directly vented out through a back-pressure regulator (V) if the detector is being flushed with CO₂ (purge mode). The purifier, CO₂ absorber, the exhaust, the analysis and the mixer modules are installed on the surface. The mixer can be programmed to inject a fixed amount of fresh gas into the loop. Part of the excess gas in the system is exhausted through the analysis lines (A). The remainder of the excess gas in the loop is exhausted through a mass-flow controller (E) regulated against the pressure set-point in the high-pressure buffer. In run mode, the Cu-catalyzer purifier is activated, so that one of its two cartridges continuously removes O₂ and H₂O from the gas. These cartridges can be regenerated on the spot by a programmed sequence of purging them with hydrogenated gas and heating them.

After the gas injection from the mixer into the gas loop, the pressure is decreased to a fraction of a bar by a high-precision needle valve (D). The resulting pressure will determine the total gas flow (from surface to cavern) through the detector, whose pressure is fixed by the compressor module. This configuration allow for independent settings of the detector pressure and the gas flow, a feature of closed-loop gas systems.

6.2.2. On-detector distribution

A final module, the distribution module, is located in the cavern. Here the detector inlet and outlet are distributed and collected. Since the gas is distributed into the drift volume through the strip-holding rods, the inlet and outlet gas pipes have to reach both endplates. For mechanical and access reasons, each gas pipe is split into two in order to service the top and bottom halves of each endplate. Finally, a set of eight half-circular manifolds is installed at the inner and outer rims of the service support wheels to service the rods through flexible bellows. The diameter and length of these bellows are adjusted to tune the fraction of gas flow for each rod, such that special rods like the HV cable rod, the laser mirror rods and the resistor rods receive a limited flow. In this manner, the gas flows radially from the inner to the outer field cage vessels, and the flow is uniform in the z-direction. Therefore, there is no force on the central electrode.

The distribution module also holds a single-pass gas system to supply CO₂ to the two insulating volumes surrounding the field cage. The flow in this case is in the z-direction, and the gas is exhausted through an extraction system coupled to bubblers installed in the distribution rack.

6.2.3. Filling

Before injecting neon, the detector is flushed with CO₂ and the return gas is exhausted through the V back-pressure regulator in the high-pressure area (purge mode). Once the air contamination is at the trace level, filling with Ne starts. In this fill mode, the gas is made to recirculate in the loop by switching the three-way valve T (see Fig. 43). Pure neon is injected from the mixer as cartridges filled with molecular sieve alternatively trap CO₂ and remove it from the system. The CO₂ cartridges are regenerated after saturation, cooled down and put back into service when needed. A thermal conductivity CO₂ analyzer and a gas chromatograph are used to measure the composition. The last step is to inject a fixed amount of N₂ in order to establish the final composition.

6.2.4. Running

Once the final mixture has been blended, the system is switched to run mode. The high pressure is regulated within precise limits according to the ambient pressure fluctuations. A

small amount of fresh gas (40 l/h) is continuously injected while the excess gas is vented out.

The purifier is activated during normal operation in order to clean residual water and oxygen from the gas. After the gas is decompressed at the surface, a bypass in the gas stream may be used to introduce a Rubidium source, which releases gaseous ^{83}Kr isotopes into the gas for detector calibration purposes.

6.2.5. Back-up system

A set of alarms ensures that the system does not produce conditions which are dangerous for the detector integrity, in which case the system stops, i.e. the compressor stops and the high pressures are isolated from the detector via pneumatic valves. The ultimate safety of the detector for excessive over- or under-pressures is a paraffine oil-bubbler directly connected at the endplate of the TPC. This bubbler has been designed to produce small bubbles and leaves ample space for gas to be exhausted or air to be injected, depending on the pressure conditions. If the system was stopped for a long time, inevitably air would enter the detector through the safety bubbler when the atmospheric pressure increases. In order to avoid this situation, a back-up supply of gas is put in place as shown in Fig. 43. Premixed gas flows through a line connected to the TPC as soon as the power fails or the system goes to stop mode. This line is equipped with a pressure regulator tuned such that if the pressure in the detector decreases, the differential pressure that the regulator sees increases and therefore opens to the gas stream. The back-up system is located on the platform just above the cavern (i.e. on the plug).

6.2.6. Analysis

An analysis module at the surface is connected to various high-pressure points of the system (points A in Fig. 43) through a pressure reducer. The gas to be sampled is extracted from the high-pressure buffer, clean gas from a point after the purifier module, and fresh gas from the mixer outlet. The analysis module has two single-pass lines (i.e., the outlet gas is vented out). One of them contains the CO_2 , the H_2O , and the O_2 analyzers, which operate at a flow rate of a few l/h. The user interface allows the operator to define which lines to sample and for how long, in a periodic way. The second line contains a drift velocity monitor which also provides a means of calculating the gas composition by measuring, in addition, a gas amplification factor.

Another line from the gas system is connected to a gas chromatograph through a manual valve. The instrument is equipped with a so-called tandem capillary column which allows for separation of the three gases of the gas mixture as well as the rest of gases that compose the air. The detection of the effluents is carried out by a thermal conductivity cell. In this way, the gas composition and the air contamination can be measured. Fig. 44 shows the CO_2 and N_2 concentrations as a function of time as measured with the chromatograph over one year.

7. Laser system

Precise reconstruction of particle tracks in the TPC requires a thorough understanding of the drift velocity and any inhomogeneities in the drift field. A non-uniform electron drift can be caused by mechanical or electrical imperfections in the field cage and readout chambers, whereas deviations of the electron drift from the ideal paths inside the gas volume are caused by temperature variations, relative misalignment of the electrical and magnetic fields ($E \times B$ effects) and local variations of the electric field from moving charges (space-charge effects).

To calibrate the drift field parameters against a known standard, a laser calibration system was built, using a large

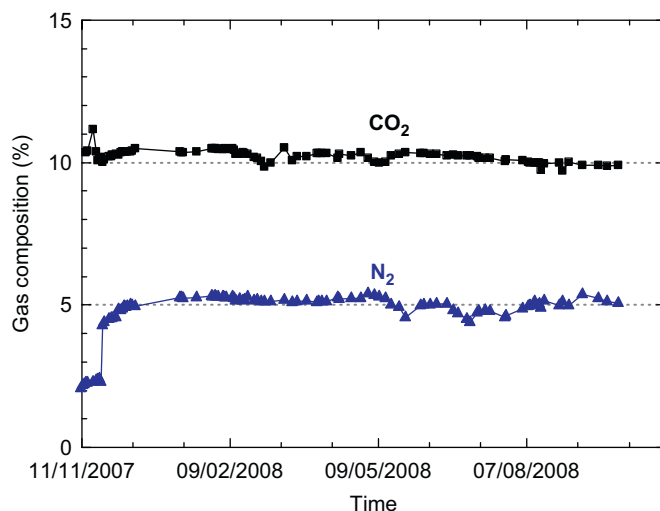


Fig. 44. CO_2 and N_2 contents in the TPC during 2008, as measured with a gas chromatograph.

number of narrow ultraviolet rays at predefined positions inside the drift volume to generate tracks. The system was designed to make fast and accurate measurements of time varying drift velocities. It will run every half hour interspersed between physics events to measure the drift velocity and assess space charge effects. The laser system was used extensively during the detector commissioning for testing of the electronics and the alignment of the readout chambers and central electrode.

Many features of the ALICE laser calibration system follow the system built and operated by the STAR experiment at RHIC [47,48].

7.1. Requirements

The goal is to measure distortions in the TPC drift field with a relative error of 5×10^{-4} . Narrow beams of pulsed UV laser light can be used to generate tracks in the active volume of the TPC. If the track positions are well known by construction, they can be used to calibrate the electron drift velocity. For a comprehensive review of the use of lasers in gaseous detectors, see Ref. [49].

To measure the drift velocity on a single event basis to the required precision, the position of the tracks must be known to a spatial resolution of $\sigma_{r\phi} \approx 800 \mu\text{m}$ and $\sigma_z \approx 1000 \mu\text{m}$ and the individual laser tracks have comparable transverse dimensions. The stability of their position must be assured at the same level. The nature of the laser tracks assures that the tracks are always straight lines.

We use pulsed monochromatic laser beams of 266 nm wavelength ($E = h\nu = 4.66 \text{ eV}$) and $\approx 5 \text{ ns}$ pulse duration with approximately Gaussian cross-section with $\sigma \approx 400 \mu\text{m}$. The ionization in the gas volume along the laser path occurs via two photon absorption by organic impurities with ionization potentials in the range 5–8 eV. The molecules of the pure Ne- CO_2 - N_2 drift gas have ionization potentials above 10 eV and are not ionized by the laser.

Because the ionization process is mostly a result of gas impurities, it is difficult to determine the necessary beam intensity a priori. Experience from this and other experiments show that energy densities of approximately $20 \mu\text{J}/\text{mm}^2$ for a 5 ns pulse at 266 nm wavelength are sufficient to obtain an ionization corresponding to several minimum ionizing particles. We designed our system to have up to $40 \mu\text{J}/\text{mm}^2$ per pulse.

The aim is to measure the response of the TPC to several hundred laser tracks generated simultaneously throughout the

TPC drift volume at predefined positions. The laser events can be generated in special calibration runs or interspersed between physics events. To obtain the best precision of the measured tracks, the preferred geometry is one where the tracks have constant drift times and are perpendicular to the wires. For this configuration, clusters are smallest and the electronics and reconstruction programs give the best possible single point resolution. Simultaneously, an extensive coverage of the full drift volume is desired. This led us to provide tracks in planes at constant z , of which some radiate with approximately constant φ . Tracks generated at different z throughout the drift volume allow easy determination of drift velocities from single laser events.

Most metallic surfaces have work functions below 4.66 eV and emit electrons by photoelectric effect when hit by UV light above this energy. Being a first order effect in the light intensity, a considerable amount of low energy electrons are seen from the diffusely scattered, time correlated UV light produced by reflections. The signal from the aluminum surface of the central electrode is used to give a precise picture at the maximum drift time across the electrode.

7.2. System overview

The idea of generating hundreds of narrow laser beams simultaneously was developed for the STAR experiment and was modified appropriately for ALICE. The basic principle is that the narrow beams are generated very close to the drift volume by optics in a mechanically very stable configuration. Fig. 45 shows a sketch of the principle.

A commercial laser outside the TPC generates an energetic pulsed beam of UV light with 25 mm diameter and very low divergence. Through an optical system of semitransparent beam splitters, mirrors and bending prisms, this wide beam is split in several lower intensity beams and guided into the TPC at different entry points through quartz windows. The wide beams travel along the inside of the hollow outer rods of the field cage, used for holding the mylar strips that define the electric field. Inside the rods, the wide beams are intersected by a number of very small mirrors (1 mm diameter) that each deflect a small part of the wide laser beam into the TPC drift volume. The dimensions, points of origin and directions of the narrow beams are given by the size, positions and angles of the micromirrors and only to a very minor degree by the parameters of the wide beam. The micromirrors are grouped in small bundles and placed along the length of the rod so that they do not shadow each other. The undeflected part of the wide beam is used for position and intensity monitoring by cameras placed at the far end of the rod. All elements of the optical guidance and splitting system are static, except for a few remotely controllable mirrors used to fine tune the beam path.

Six rods in each half of the TPC were equipped with four micromirror bundles each. Each mirror bundle contains seven small mirrors. The wide beam originates from one laser for each TPC half and is split and guided into the six rods. The two lasers are synchronized to provide simultaneous laser pulses in the full TPC, thus resulting in a total of 336 simultaneous narrow laser rays in the TPC volume. It is also possible to operate the system with just one laser for the full TPC using an additional beam-splitter near the laser.

7.3. Optical system

7.3.1. UV lasers

Energetic pulsed laser light in the UV region is obtained from a Nd:YAG laser ($\lambda = 1064 \text{ nm}$) equipped with two frequency




Fig. 45. Schematic 3D view of the TPC and the laser system. Two wide pulsed laser beams enter horizontally at the bottom of the TPC and are guided around the two end-caps by mirrors, prisms and beam splitters before entering the TPC. Bundles of micromirrors in the hollow laser rods intersect the beams and generate a large number of thin rays in the TPC drift volume. The undeflected part of the beams continue through the monitor rods to cameras at the far end. All elements are fixed mechanically, except for the remote controllable entrance mirrors at the bottom.

doublers, generating pulses of UV light of 266 nm wavelength. The same kind of laser was used for the STAR experiment at RHIC [48], and also in NA49 [7] and CERES/NA45 [50] at the CERN SPS and ALEPH [51] at LEP. The typical beam diameter from this kind of laser is 9–10 mm, but our lasers were fitted with telescopes to expand the beam diameter to about 25 mm. The power density of the narrow beams is given by that of the wide beam inside each rod. A $40 \mu\text{J}/\text{mm}^2$ density in each of the beams translates into a requirement of the total energy out of the laser of 100 mJ per pulse.

The laser from one side of the TPC was provided by Spectron Laser Systems Ltd, model SL805-UPG. Operated in Q-switched mode, it provides 130 mJ/pulse of $\approx 5 \text{ ns}$ duration at 266 nm wavelength and a repetition rate of 10 Hz. A computer controlled tracking system continuously optimizes the orientation of the second frequency doubling crystal to compensate for temperature drifts. Built into the laser is a beam expanding telescope to enlarge the beam diameter to 25 mm and reduce the beam divergence to $\approx 0.3 \text{ mrad}$. Close to the laser, the beam has a flat intensity profile across the beam spot which develops smoothly into a Gaussian profile after 20–30 m.

A second laser for the other end of the TPC is a similar Q-switched Nd:YAG laser from EKSPLA uab, model NL313, similarly fitted with a computer controlled frequency quadrupling system and beam expanding telescope. It provides up to 150 mJ pulses at $\lambda = 266 \text{ nm}$ and 3–5 ns duration at 10 Hz repetition rate. After the expander telescope, the 25 mm diameter beam also has a flat top profile and a divergence of $< 0.5 \text{ mrad}$.

The lasers are triggered by a fixed rate 10 Hz external clock, such that their pulses are synchronized to each other and to the readout clock of the TPC. Both lasers are placed in optically stable conditions in a hut outside the L3 magnet at $z \approx -10 \text{ m}$, 2.5 m under the LHC beam line. Together with the actual laser heads and their power supplies, the hut contains remote adjustable mirrors to point the wide beams in the correct direction toward the TPC and remote control electronics for the lasers, monitor cameras and adjustable mirrors. Each laser beam is deflected through a ‘knee’ of one fixed and one adjustable mirror before it exits the

hut. The hut ensures personnel safety against UV light in the underground hall.

7.3.2. Laser beam transport system

From the laser hut, the two laser beams are guided to the entrance windows at the outer radius of the TPC field cage by a system of mirrors, beam splitters and bending prisms, all enclosed in pipes to ensure personnel safety and stable optical conditions. Both beams pass through a vertical slit in the L3 magnet. One beam hits the nearest A-side endplate close to its outer radius, where a mirror reflects it by 90° into a vertical plane parallel to the TPC endplate. The other beam passes slightly lower, and after another knee of two 90° reflections it enters a tube mounted below the TPC field cage. It continues in a straight line to the far C-side endplate where another 90° mirror bends it into the vertical plane parallel to this plate.

Fig. 46 shows an overview of the optical elements in the guidance system. High quality fused silica is used for all optical elements and all surfaces are antireflection coated for UV light. Dielectric coatings of the mirrors and beam splitters are designed to divide the beam intensity evenly between the six rods. First, a 50% beam splitter directs half of the beam in each direction around the periphery. Prisms deflect the beams by 30° such that each half of the beam passes over the prolongation of three of the outer TPC rods. At these points, beam splitters at 45° angles direct equal intensity beams into each rod along the z axis by deflecting 33%, 50% of the remaining and ≈ 99% of the then remaining beam through a 90° angle. A small remaining beam is monitored by a camera and dumped after the last splitter. The beam paths on the two endplates are virtually identical. By mirror symmetry they are arranged such that the prolongation of each of the six laser rods at one TPC end beyond the central electrode corresponds to a hollow rod (monitor rod) in the opposite end of the TPC.

All optical elements on the endplates are placed in small boxes. Each box is firmly attached to the endplate and the angles of the optics were fine adjusted manually inside the box after installation. Fig. 47 shows examples of the mechanics in two such boxes.

A few of the mirrors both in the laser hut and at the entrance to the endplates are remotely adjustable. Together they define the beam vector at the entry point on the two vertical planes parallel to the endplates. The rest of the beam guidance system is based on fixed optics, carefully aligned during the construction and using the endplates as stable mechanical support.

7.3.3. Micromirrors and laser rods

After entering the six laser rods at each end of the TPC through sealed quartz windows, the wide laser beams travel along the inside of the rods as illustrated in Fig. 45. They are intersected by four, roughly equally spaced, micromirror bundles before arriving at the TPC central electrode. Here, the beam passes through another sealed quartz window to the hollow rod in the other half of the TPC and exits at the far end through a third window. At the far end, the beam position is monitored by a camera before being dumped.

The generation of narrow beams happens inside the laser rods by reflecting the wide laser beam off micromirrors at a 45° incidence angle. The mirrors were made from short 1 mm diameter quartz fibers, cut at a 45° angle at one end. The resulting elliptical surface was polished and coated for total reflectivity for 266 nm light. To increase the number of laser tracks, bundles of micromirrors were assembled with 7 fibers in a unit which generate 7 narrow beams when hit by the wide laser beam at the cut fiber ends. The rays spread out from the bundle roughly in a plane perpendicular to the wide laser beam. The fibers were rotated along their axis to give predefined azimuthal reflection angles: 2.5°, ±9.2°, ±16.0° and ±31.8° relative to the direction towards the TPC axis. The bundles were constructed with a

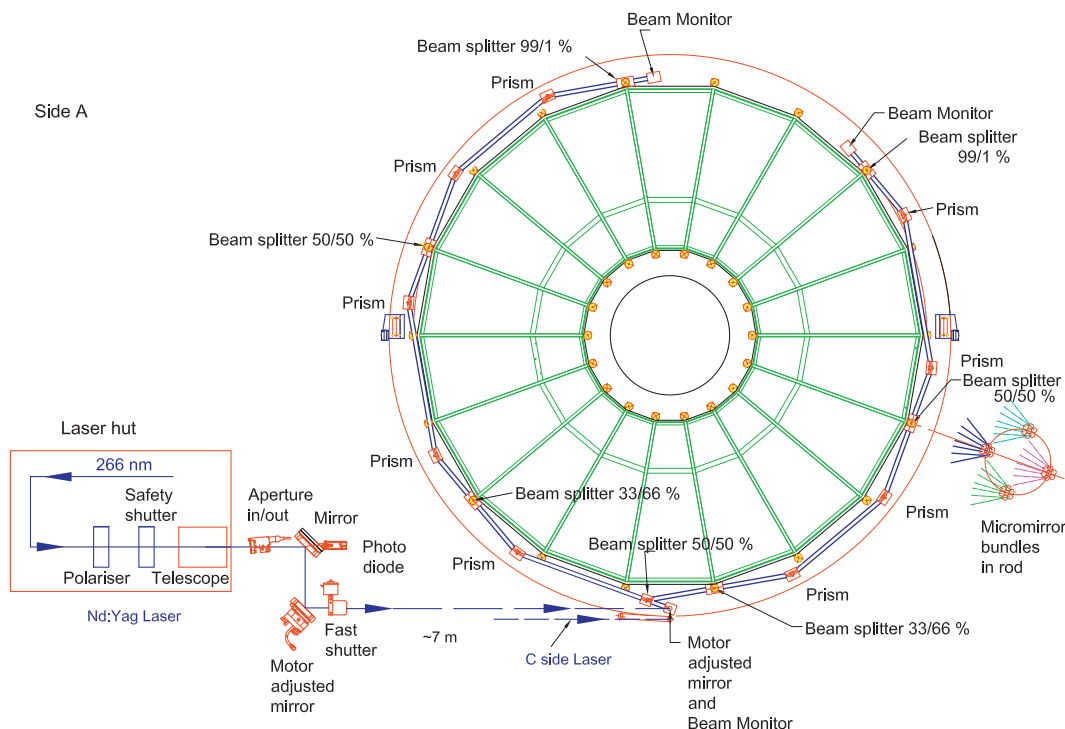


Fig. 46. Overview of the optical elements to guide the laser beam from the laser to the entrance windows in the TPC field cage. The A side system is shown; the C side system is obtained by mirror symmetry in a vertical plane along the TPC axis.




Fig. 47. Example of the design of the interior of optics boxes installed on the TPC endplate. The shown boxes contain a 30° bending prism and a beam splitter, respectively.




Fig. 48. Section of a laser rod with a micromirror bundle and its support. An end view of the tube with the position of the four mirror bundles is also shown.

tolerance of 1° in both azimuth and dip angle, and the angles of each bundle were measured to a precision of 0.05 mrad.

Fig. 48 shows the principle of mounting the micromirrors in the TPC rods. The holders are integrated into aluminum rings and glued between the polycarbonate tube pieces that build up the 2.5 m long rod. Holes are drilled in the tube to allow the narrow beams to exit the rod and enter the drift volume.

Because the position and in particular the mounting angle of the micromirrors define the narrow beam positions inside the drift volume, care was taken to assure the mechanical stability of the mirror holders. The gluing procedure for these rings was specially adapted to control the position of the mirror holders. The angles of the reflected beams were measured after the assembly of the 2.5 m long rod. Also, special care was taken to place the mirror holders close to the rod supports to the outer field cage to minimize movements due to mechanical stresses on the rods.

7.4. Laser beam characteristics and alignment

7.4.1. Narrow beam characteristics

A narrow ray generated by reflection from a circular surface (as seen along the beam direction) is equivalent to the beam progressing beyond a screen with a similarly shaped hole. Thus, the narrow beams in the TPC are approximated by pure Fresnel

diffraction of an infinite planar wave through a circular aperture of 1 mm diameter.

The profile and total energy of the narrow beams, generated by reflection from each of the micromirror bundles, were measured in the lab. The reflected beams were measured by a calibrated energy meter and imaged with a CCD camera as a function of the distance, z , from the mirror bundle. The energies did not vary substantially for different micromirrors and were stable as a function of time, reflecting the quality of the coated surfaces and the laser. The patterns matched qualitatively what one expects from Fresnel diffraction and the measured FWHM remains at or below 1 mm up to $z=200$ cm. For further details, see Ref. [33].

7.4.2. Narrow beam layout

The transverse pattern of the narrow beams in the TPC volume follows from the micromirror angles given in Section 7.3.3. In the z -direction, the planes of laser tracks are situated at $z \approx \pm 115, 820, 1660, 2440$ mm. When defining the angles and z positions, we have aimed at generating beams radiating at constant z that cross-sector boundaries strategically, i.e. at points where alignment between sectors would benefit the most. We have also avoided having too many tracks with small angles relative to the wires of the readout chambers.

Fig. 49 shows the resulting pattern for the beams at a single position in z . Beams from neighboring laser rods in φ are offset by a few cm in z relative to each other to avoid most of the apparent

beam crossings. The central ray (of the seven) from each bundle radiates in (r,φ) in a way very similar to tracks from the interaction point while the other angles were chosen in order to illuminate all sectors and assure beams that cross-over all sector boundaries. In this way, we defined a pattern optimized for testing and sector alignment as well as for drift distortion measurements. One should note that, although the position of the beams was measured to a high precision, the production tolerance of 1° on the angles results in deviations of the shown paths of up to 40 mm near the inner cylinder.

7.4.3. Spatial precision and stability

The TPC calibration would ideally require absolute knowledge with infinite precision of the spatial position of all laser tracks in absolute ALICE coordinates. Given the mechanical tolerance, the best absolute coordinate frame for each half TPC is defined by the plane of the endplate. All readout chambers and the plane of the central electrode defining the high voltage surface were aligned and adjusted relative to the endplates by optical survey. These surfaces were defined relative to each other during construction to a precision of approximately $100\mu\text{m}$ with the aim of obtaining a relative electric drift field error below 5×10^{-4} .

A final precision goal of $800\text{--}1000\mu\text{m}$ for space points translate into matching requirements on spatial coordinates and angles of the laser systems:

$$(\Delta x, \Delta y, \Delta z) \leq (800\text{--}1000)\mu\text{m}$$

$$(\Delta\theta, \Delta\varphi) \leq (0.4\text{--}0.5)\text{ mrad.}$$

By far the most important issue in the definition of the laser track positions is the placement of the micromirrors, both in (x, y, z) and in particular in the angles (θ, φ) . The only other deviation from the ideal rays that matters is the incidence angle of the wide laser beam on the micromirrors and this is relatively easy to measure and keep constant because the optical system has long lever arms.

Even if small movements of the rods and external beam optics cannot be excluded during and after the assembly of the TPC,

the stability of the finished and installed TPC is well below the requirements.

7.4.4. Construction and surveys

The construction errors of the micromirror bundles were specified to be $< 100\mu\text{m}$ in the spatial measures and $< 1^\circ$ in all reflection angles. The critical surfaces were, however, measured to remain within a $50\mu\text{m}$ tolerance and all the angles of the reflected beams from the mirror bundles were subsequently measured to a precision of 0.05 mrad. The mechanics of assembly of the rods assured very precise z position of all the mirror bundles and, after assembly of the rods, a second measurement of the angles was performed to a precision of 0.1 mrad, using a green laser.

Relative (x, y, z) shifts of micromirrors within a rod during the installation in the TPC are unlikely. However, the absolute position of the rods relative to the endplates was not guaranteed to a precision good enough for the laser system, and the mounting of mylar strips was seen to cause a small bending of the rods between the support points and result in small rotations of the micromirrors, especially in the dip angles. In order to monitor such shifts, the (θ, φ) angles of the central micromirror in each bundle was remeasured by surveying the intersection of the narrow laser beam with the inner cylinder using the green laser (see Fig. 49).

The construction of the TPC required the assembly of the field cage while it was standing on first one, and then the other end, before it was finally turned into its final horizontal position. Changes in the mechanical stresses may have influenced the absolute position of the rods. However, the stiffness of the rods guarantees a continued good relative alignment of the micromirrors in the same rod. In this way, the absolute (x, y, z) of the micromirrors is known to $100\text{--}200\mu\text{m}$ and the mirror angles given by the measurements of the rods prior to installation. Furthermore, the relative angles of the micromirrors, which are glued together in a bundle, are determined to a precision of 0.1 mrad by the lab measurements.

The initial setup of the optics on the endplates needed careful manual adjustment of all the mirrors and prisms in the system.

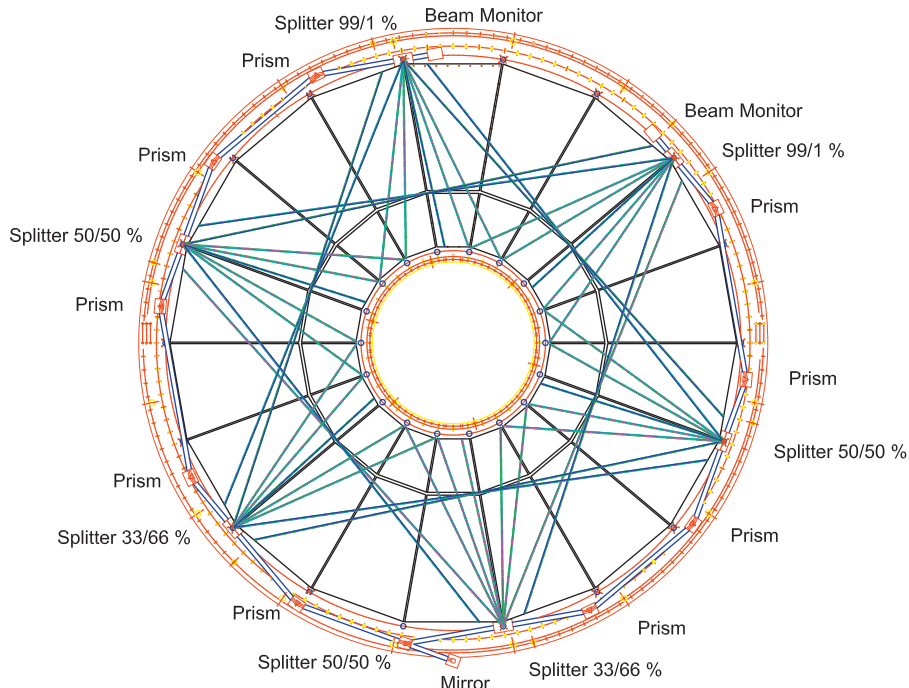


Fig. 49. Ideal laser tracks in (r,φ) in the TPC drift volume, projected to the endcap. The pattern repeats eight times through the full length of the TPC.

Special measurement tools were installed temporarily in the beam path. Adjustments of the beam paths between the lasers and the endplates of the TPC must be realigned after interventions, using the remotely adjustable mirrors and cameras.

7.4.5. Online and offline alignment

Given that the readout chambers and the central electrode constitute the best aligned surfaces during the TPC construction, it is an advantage to use these surfaces as references also for the absolute position of the laser tracks. As mentioned, all metallic surfaces inside the TPC, which are hit by stray laser light, emit electrons by the photoelectric effect, synchronously with the laser pulse. This is in particular the case for the central electrode made of a stretched aluminized mylar foil. We use these electrons to image the whole plane at the end of the drift volume. The central electrode signal is used to calculate the full drift time and its (x, y) variations and to obtain an online measurement of drift velocity variations in time. The full 3D map of the drift field is obtained in an offline alignment procedure where the reconstructed tracks are used together with all the survey information of relative and absolute laser beam positions.

Even without the ultimate, absolute precision of all laser tracks, much can be learned from a few well determined tracks, in particular if they span a large lever arm in z . Close to the outer TPC radius, the angular uncertainties of the rays play only a minor role, and it is possible to obtain a very good drift velocity measurement near the six rod positions in each half of the TPC.

Furthermore, variations over time are tracked to a very good accuracy throughout the TPC. An uncertainty in the time variation of the laser beam positions could come from a possible torsion in the field cage due to variations in the mechanical loads, magnetic field and the external temperature. In stable running conditions, these effects are minimal.

7.5. Operational aspects

Under the operational conditions of the LHC, it is essential that all controls and monitoring of the system are done remotely and the software interfaces are integrated in the Detector Control System (DCS). The two lasers and the angle adjustable mirrors are controlled this way. Along the beam path, a number of cameras are used to look at the laser beam from the control room. Furthermore, several electronics modules are introduced in the system. One is used to synchronize the laser pulses to the ALICE trigger and the TPC readout clock. Another controls various shutters and apertures that can be inserted into the beam paths by remote control.

7.5.1. Beam monitoring and steering

The laser heads with their power supplies are equipped with remote control facilities through RS-232 communication. After conversion to a network protocol using a Digi terminal server [52], all controls are transferred to the online computers. Both lasers use an active feedback system to optimize the frequency conversion from 532 to 266 nm by very slight adjustment of the angle of the conversion crystal inside the laser head. The algorithms delivered by the laser manufacturers were improved and ported to run on the online computers.

Eight mirrors were inserted in the beam paths where they make 90° bends. They define the correct alignment of the beams entering the field cage volume. Five of the mirrors are adjustable by remote control in two angular dimensions. For the A-side laser beam, one adjustable mirror is placed in the laser hut and another at the beam intersection point on the TPC endplate. The C-side laser beam is guided through the knee at the bottom of the TPC by

a third adjustable mirror. The mechanical movement is driven by the New Focus Picomotors [53] which allow angular steering with 0.7 μ rad resolution, even in the magnetic field. The direction and quality of the wide laser beams are monitored online by simple CMOS cameras [54], focused on mat glass screens, on which the UV pulses are converted to visible (blue) light. The cameras are interfaced to the computer by coaxial cables and a frame grabber card [55]. In total, there are 18 cameras in the full system, placed at the beam entrance points on the endplates, at the end of the two paths on the periphery of each endplate and finally at the far end of the rods, downstream from each laser rod. Fig. 50 shows the graphical user interface with a display of 16 (out of 18) cameras and the options for beam steering by the five adjustable mirrors. A trained operator can manipulate the few parameters of the mirrors while observing the images on the cameras, and the system then remains stable over months. To facilitate the adjustment of the mirrors, a 3 mm aperture can be inserted by remote control in the middle of each of the wide beams just after the exit from the laser heads, thus reducing the beam to a thin one useful as a 'pointer'.

Finally, two shutters were installed on each laser beam. One set is internal to the laser head, used as a safety device to block all light from exiting the head. The other set at the exit from the laser hut is used in conjunction with the laser synchronization module (see Section 7.5.2). They operate rapidly enough to decide on a pulse-by-pulse level whether to block the pulse or let it pass.

7.5.2. Trigger and synchronization

Optimal thermal stability of the lasers requires that they are operated at a constant 10 Hz pulse rate. Equally important for good calibration is a good synchronization of the laser pulses to the TPC readout clock. Furthermore, the system is designed to take data with laser tracks in various run configurations, either as dedicated calibration runs with one or both lasers running at a fixed 10 Hz trigger rate or in a mode where the laser events are interspersed between physics triggers. A dedicated laser synchronization module was built to handle all trigger and timing conditions of the laser operation. It is based on the common RCU module where the optical communication daughter board was replaced by a dedicated signal driver board. The module provides programmable timing outputs to control the operation of the lasers and the shutters, and to interface to the ALICE trigger system.

Based on the LHC 40 MHz clock, the module generates a 10 Hz clock in phase with the TPC readout clock. Each laser is controlled by signals from the module to trigger their flash lamps and Q-switches in order to generate laser pulses synchronous to the 10 Hz base clock. In case of stand-alone calibration runs, the trigger module provides an ALICE trigger at the 10 Hz base clock rate. For laser events interspersed between physics triggers, the module also generates signals at the 10 Hz rate to fire the laser flash lamps and generate a calibration event trigger. If this trigger is vetoed by the central ALICE trigger system, the laser pulse can be suppressed by vetoing the Q-switch signal to the lasers for this event. Otherwise, the generation of a Q-switch signal assures a timely laser pulse from one or both lasers. During LHC collisions, we foresee to run in a mode where the lasers are first warmed up at 10 Hz with the shutters closed, for as much as 1 min until they are thermally stable. The warm-up is followed by a short burst (< 1 min) of laser events interleaved between physics triggers and a period of about one half hour where the lasers are put into standby mode, without firing the flash lamps. This burst mode of operation is handled automatically by the detector control system and is designed to ensure a laser flashlamp

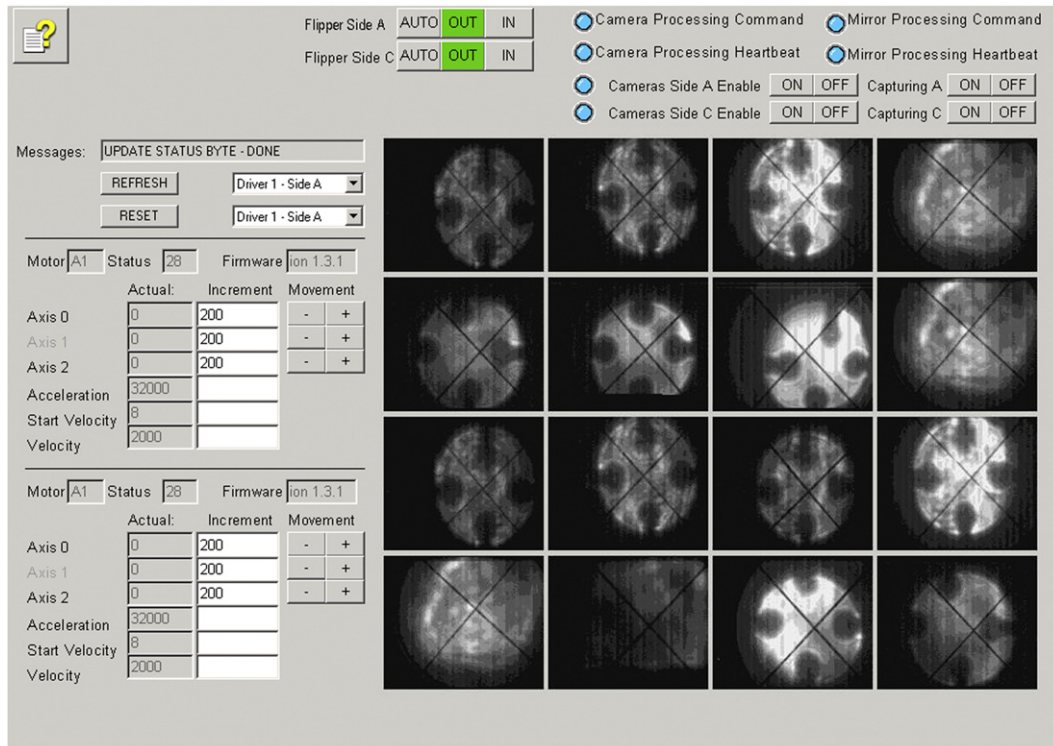


Fig. 50. User interface panel for the steering of the laser beams and monitoring by cameras.

lifetime commensurate with the length of typical LHC runs (of order one year).

8. Infrastructure and services

The TPC together with the ITS, TRD and TOF detectors of the 'central barrel', (see Fig. 1), are mounted in the so-called space frame (see Ref. [33, Section 8]).

The space frame is supported with four adjustable feet on two beams traversing the full length of the L3 magnet. The TPC rests on two rails, located in the median plane, in the central opening of the space frame. The ITS is suspended on two points in the inner opening of the TPC. An overall view of the TPC together with the Service Support Wheels (SSWs), the rails and I-bars (but without the support for the ITS) is shown in Fig. 51.

The TPC services are organized as follows: the FEE is housed in the SSWs, which are located on the A and C side at a distance of about 20 cm from the TPC endplates. Connection to the readout chambers is via flexible Kapton cables. Outgoing services are routed via the so-called baby frame and back frame on the A and C side, respectively. These frames are short, about 2 m long extensions of the space frame with nearly the same geometry as the space frame. They are decoupled mechanically from the space frame except for the flexible services.

8.1. Moving the TPC

The TPC can be moved on the rail system. This was necessary during installation, and is foreseen in case of future servicing of TPC or ITS. When the TPC needs to be moved, the rails are extended on the A side by transfer rails that are connected to a support structure with rails outside the L3 magnet. For initial installation and later servicing of the ITS, the TPC is moved to the so-called parking position, 4.8 m toward the A side. All services,




Fig. 51. Overall view of the TPC with the SSW, rails and I-bars.

except the 100 kV HV cable and temporary gas connections have to be removed during movement.

The TPC sits on its rail with 4 ft Teflon-coated gliders which are adjustable in x and y . In Fig. 52 the detailed design of a pair of feet (one for the TPC one for the SSW, see below) with the Teflon padded gliders is shown.

The gliders on the I side have no lateral play, the ones on the O side have ± 4 mm play. During movement, two of the feet, either both A or both C side, are supported vertically by coupled hydraulic jacks, providing effectively a 3-point support and completely eliminating torsional stress on the field cage due to imperfect parallelism of the rails. The rails were initially, before loading the space frame with detectors, parallel with a tolerance of about 0.3 mm vertically and about ± 1 mm laterally. At rest, the feet are fixed with adjustable screws. The procedure of letting two feet hydraulically float is also applied during movement of other heavy detector components that might lead to a deformation of the space frame at a level of more than a few tenths of a mm.

The two SSWs are equipped with the same type of gliders. During movement, the TPC and the two SSWs are coupled by four

so-called track rods between the feet. The ensemble, TPC and SSWs, is always moved by (stepwise) *pulling* with two hydraulic jacks attached to the feet of a SSW and the rail. In order to limit the tensional stress of the field cage, two steel cords between the A and C side feet on each rail are pre-tensioned to about 8000 N. The friction coefficient of the system Teflon-glider and rail is about 10%, resulting in a pulling force ≤ 8000 N on either side, controlled by the hydraulic pressure. For safety, the hydraulic pressure is limited.

When moving the TPC, the ITS is disengaged from its fixtures at the TPC and supported by a temporary second rail system, the ITS




Fig. 52. Design of the feet to support the TPC and the SSW. The surfaces in contact with the rails are covered with Teflon pads.

rails. These rails are each attached at the hadron absorber cone (see Fig. 1) on one end and at a support outside the L3 magnet on the other end. In between, they have sliding supports at the two TPC endplates which allows the TPC to move with respect to the ITS. The ITS rails are removed when the TPC is at the working position.

8.2. Service support wheel

The SSW houses and supports the front electronics and its services (LV, DCS, DAQ), the manifolds for the various cooling circuits, and the drift gas manifolds. In Fig. 53 details of the SSW are shown in a CAD model. In Fig. 54 more details of the connection between one I-bar and the endplate are shown.

Forces on the endplates and readout chambers are kept to a minimum. The flexible Kapton cable connections from the front-end cards to the readout chambers have a maximum play of ± 5 mm. The position of the SSW relative to the respective endplate has thus to be controlled to much better than this. While moving, the constant distance is assured by four so-called tie rods between each SSW and the corresponding endplate, at the angles of about 45° , 135° , 225° and 315° at the outer circumference. The x - y position is tuned by adjusting the glider feet vertically and laterally. In the final position, the support of the SSW is transferred from the tie rods to fixation points on the space frame and close to the location of the tie rods, minimizing the forces on the TPC endplates.

The SSW is constructed from 18 essentially identical trapezoidal frames matching the sectors of the endplate and the readout chambers. In Fig. 55, one frame of the SSW with the holder for the front-end cards is shown together with an inner and an outer readout chamber.

In Fig. 56, a close up of the mounting support for the FECs is given. On the side facing the endplate, front-end-card holders are attached, adjusted with a precision of about one mm in x - y . The LV services, 12 per sector, are bundled with a cooling circuit and fixed to one long side of the trapeze. The outer face of each




Fig. 53. Closeup of the SSW together with the endplate of the TPC cylinder showing some components of the services: cooling lines, gas distribution pipe, a patch box and low voltage bus bars.

sector is closed with a cooled cover. A number of distribution boxes is mounted near the outer circumference of each SSW (see below).

In addition to TPC services, the A side SSW has to serve as the support for ITS services and beam pipe parts. The carriers of the ITS services, two half cones jutting from the SSW to the ITS, are each fixed in three points on the inner circumference of the SSW. Sector 13, pointing downward, carries a framework which supports a pump and valve of the beam pipe.

In Table 13, an overview of the weights of all components described in this section is given.




Fig. 54. Details of the connection between I-bar and endplate.




Fig. 55. One frame of the SSW together with the holders for the FECs and the wire chambers. For clarity only a few front-end cards are shown.

8.3. Low-voltage distribution

The FEE Low-Voltage (LV) system is based on modular, low-noise, high-efficiency switching power supplies (W-IE-NE-R model PL512 [56]). In each of the 19 water-cooled crates the modules to supply 2 complete sectors are housed (rated: 2–7V, 2 × 100A, 2 × 200A). Remote control by DCS is provided via a network connection (see Section 9.1.4).

The power supplies are not connected to ground internally. The ground potential is defined at the level of the front-end cards. The connection between power supplies and TPC is based on large cross-section unshielded and uncooled copper cables. Inside each sector the power is distributed by bus bars running along the spokes of the SSW. Voltage drops along these bus bars are about 20 mV. For the distribution of the LV to the individual FECs 12 small distribution boards are connected by screws to the busbars. Up to 15 FECs are supplied by one distribution board. To buffer




Fig. 56. Construction detail of the holder mechanics for the FECs. The six Kapton cables connecting to the wire chambers are also shown.

Table 13
Weights of all components described in this section.

Item	Weight (kg)	Units	Full TPC (kg)
<i>Field cage</i>			
Field cage w/o ROC	6000	1	6000
IROC	15	36	540
OROC	40	36	1440
ROC total			1980
Inner cooling panels			100
Gliders	80	4	320
I-bars C side	100	2	200
Total field cage			8600
<i>SSW</i>			
SSW frame alone	570	2	1140
Frontend cards	0.56	4356	2440
RCU	0.5	216	110
Backplanes per sector	4	36	144
Frontend card supports per sector	18	36	648
Services per sector	14	36	504
Cooling covers	9	36	324
Gliders	80	4	320
Total per SSW	2800	2	5600
TPC total			14200

Table 14
Characteristics of the low-voltage supply system.

Parameter	Analog supply	Digital supply
Supply voltage (V)	4.9 (5.2)	4.1 (4.4)
Cable cross-section (mm ²)	150	300
Current (A)	83	133
ΔU in cables (V)	0.65 (0.92)	0.9 (1.2)
Total power per sector (W)	407 (432)	545 (585)

The voltages are measured at the power supply. The currents and the power dissipation refer to one sector connected to the TPC by 40 m long cables. Due to the longer routing path and additional patch panels the voltage drops and corresponding supply voltages are higher on the A side (numbers given in brackets).

small voltage variations and also prevent voltage surges which could occur when the supply lines are interrupted, large capacitors of 10 mF each are mounted on the distribution boards. In addition smaller capacitors of 2.2 μ F for the suppression of HF noise are also mounted. The sense lines of the power supplies are connected to these local bus bars and a dynamic regulation always ensures the correct voltages at the front-end cards. An overview of the parameters of the system is presented in Table 14.

8.4. Chamber HV system

The high-voltage system for the chamber anode-wire voltages is based on the models EDS 20025p_204-K1 from ISEG, see Ref. [57]. These modules have 32 output channels grouped on two independent 16 channel boards. Each group of 16 channels is supplied by a common HV source with independent control of the individual output channels. For the auxiliary voltages (edge anode wires, cover electrodes and skirts, see Sections 3.2 and 2.5) models EHQ 8005n_156_SHV and EHQ 8010p_805_SHV with eight independent output channels are used. All power supplies reside in one crate and are remotely controlled via CAN bus from DCS (see Section 9.1.4). Since the grouping of channels in the power supplies does not match the grouping of wire chambers in the TPC (16 compared to 18) two types of patch boxes were built. The first type is located directly below the HV crate and redistributes the outputs of the 16 channel modules to multi-wire HV cables combining 18 channels plus 2 spares. They connect to the second type of patch box located directly on the SSWs of the TPC. From there individual HV cables connect to the wire chambers.

For the supply of the edge anode wires and cover electrodes the single outputs of the HV units are fanned out in splitter boxes into groups of 18 (plus 2 spares) and are also connected via multi-wire HV cables to patch boxes on the SSWs and from there via individual cables to the chambers. In case of shorts (possible in particular in the last anode-wire circuits) individual bridges are installed in the splitter boxes allowing easy disconnection of the affected circuit.

8.5. Gate pulser

A gating grid is installed between the cathode grid and the drift region, to prevent positive ions, generated in the gas-amplification process, to drift back into the drift volume and create distortions of the drift field (see Section 3.2). The gate is 'closed' when a voltage of $U_G \pm \Delta U$ is applied to alternating wires. The necessary value of ΔU , given by the magnetic field, the wire spacing, and drift field [3], is ± 90 V. A gate pulser system has been devised to enable the rapid transition of the gating grid from the 'closed' to the 'open' state upon the receipt of a trigger.

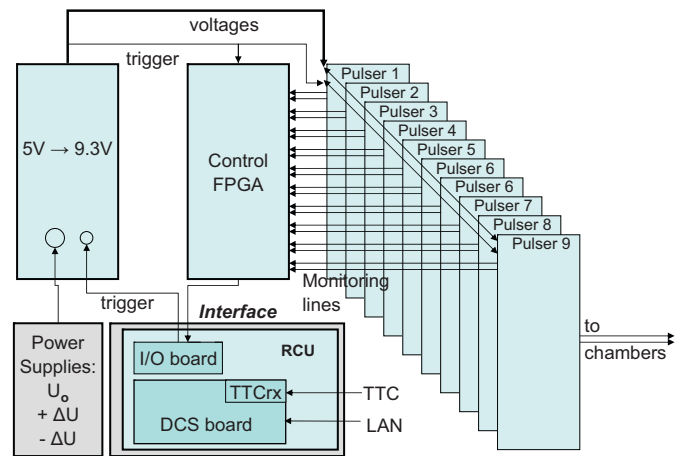


Fig. 57. Schematic setup of the gate pulser system.

Its general layout is shown in Fig. 57. The three main components of the system are:

- *The interface to DCS and trigger system* (see Section 9.2): This is an RCU board as used in the readout chain with its add-on DCS board connecting to the network and the trigger system. The firmware of the RCU is modified to control the state of the gate-pulser system (on/off), select the trigger (L0 or L1) and adjust the duration of the 'open' state of the gating grid via the DCS system. In addition the actual states of the pulsers are controlled by monitoring modules and in case of a malfunction an alarm is sent to DCS (see below).
- *The individual pulser modules*: A detailed description of the pulser module design is given in Ref. [3]. Groups of nine pulsers are housed in 3U Euro crates (six in total) together with a monitoring module. This monitoring module is connected via opto-couplers to the output lines of all pulsers and generates an error flag in case any of the pulsers does not change its state as a function of the trigger. In addition, each crate houses a 5V power supply and a module to receive the trigger signal and distribute it via the backplane. This module also receives the voltages U_G and $\pm \Delta U$ and distributes them via the backplane and, in addition, generates 9.3V via DC-DC conversion which is needed by the gate-pulser modules.
- *Three dual-channel power supplies* (Zentro model LD 2 \times 150/1 [58]): The supplies provide U_G and $\pm \Delta U$. IROCs and OROCs are supplied by separate channels. The power supplies are remotely controlled via RS232 connections from the DCS (see Section 9.1.4).

The gate-pulser system consists of two subsystems, one per TPC side. The gating circuit has to put large voltage swings ΔU on the gating grid as fast as possible with minimal pick-up on the readout electronics resulting from the transients. To achieve this a three step procedure is followed. First, the FET switches (the same FET type is used for both polarities) are grouped according to their individual switching time. Then, the FETs with similar characteristics are mounted together on the same board. In the last step, the relative timing of the positive and negative side of a pulser is tuned by a potentiometer such that a minimum in residual signal is reached when summing the two outputs and monitoring them on an oscilloscope.

For $\Delta U = \pm 90$ V the amplitude of the pickup signal on the TPC reaches 80–100 ADC counts for the IROCs, 70–90 ADC counts for the inner and 110–130 ADC counts for the outer part of the OROCs, respectively. The duration of the pickup signal is about 10

time bins and 21 time bins for IROCs and OROCs, respectively. Since these pickup signals are constant in time they are subtracted out in the pedestal subtraction procedure. The time between trigger and full transparency of the gating grid was measured to be $1.5\ \mu\text{s}$ [13]. The measured gating efficiency against positive ion feedback is described in Section 3.3.

8.6. Calibration pulser

A calibration pulser system is used to measure the gain and timing calibrations of all readout channels. In addition, the calibration pulser can be used to make general tests of the readout chain such as to identify dead channels. To generate a signal in the readout electronics, charge is injected onto the pads by pulsing the cathode wire plane. For homogeneous gas amplification and good position resolution, the mechanical tolerances of the chambers (pad sizes and distances between wire and pad-planes) are quite narrow. Therefore, the variation of charge injected onto different pads is quite small and should, in the worst case, reflect only large-scale variations across the chamber surface. In addition, by controlling the cable length between calibration pulser and wire chambers, the timing of the readout can be calibrated.

The working principle of the calibration pulser system is described in Ref. [3]. A step function is generated by an arbitrary waveform generator. Due to the capacitive coupling between the cathode-wire plane and the pad-plane the signal is differentiated and a narrow spike is detected by the readout electronics. The falling edge of the signal is always kept outside the readout time window of $100\ \mu\text{s}$ to avoid unnecessary data flow.

The schematic layout of the calibration-pulser system is shown in Fig. 58. It comprises three main components:

- *The interface to DCS and the trigger system* (see Section 9.2): This is an RCU as described before for the gate-pulser system. The firmware of the RCU is modified and includes a communication module with the control FPGA of the pulser.
- *The controller module*: Its main components is a Field Programmable Gate Array (FPGA) acting as the controller for the pulser system. Different pulse shapes can be stored as a sequence of amplitudes in its memory. On command, the FPGA is set to an active pulse generating state taking into account the desired amplitude, the pulse delay relative to the trigger and the pulse shape. In addition, the FPGA controls which driver channels are activated. Besides a simple step function, generating one single signal per readout cycle, a sequence of signals with identical

amplitudes or a ramp (signals with increasing amplitudes) can be generated. A special fine delay unit shifts the position of the output signal in 1 ns steps. The standard step size is 50 ns given by the 20 MHz internal clock of the pulser.

- *The output drivers*: They are connected via a high-impedance input to a DAC which is controlled by the FPGA. Their function is to drive the $50\ \Omega$ cables connected to the individual TPC readout chambers. The 18 sectors of one side of the TPC are supplied by four modules with nine channels each. An additional control circuit on the driver cards allows the activation/deactivation of individual channels.

There is one such system (6U Euro crate) per TPC side.

The pulser cables are terminated by $50\ \Omega$ resistors on the readout chamber side. Due to the large capacitance of the wire plane this termination is far from ideal. To avoid reflections back to the chambers the output stage of the calibration pulser has a $50\ \Omega$ impedance. This has the disadvantage that the effective output amplitude of the driver stage has to be twice the amplitude desired at the chambers ($7.8\ \text{V}$, see below). Since the pulse width is about $110\ \mu\text{s}$ in order to keep the falling edge outside the readout time window, a considerable power is dissipated. As a consequence, each of the four driver modules has its own 18 V power supply.

In Fig. 59 the shapes of the rising edges at various stages of the signal generation in the calibration pulser are shown. The simulation of an OROC chamber, with their rather large capacitive load, shows that the rise time and the width of the induced signal is given by the charging process and much less by the shaping time of the preamplifier.

The output amplitudes of the drivers are equalized to $< 1\%$. The measured non-linearity is $< 1\%$. The remote control of the amplitudes allows to check the linearity of individual readout channels. The maximum possible amplitude at the chambers is $7.8\ \text{V}$. This corresponds roughly to the dynamic range of the readout chain. The rather high output voltage is necessary since a considerable part of the signal strength arriving at the cathode-wire plane is absorbed by the anode wire plane: the buffer capacities of $4.7\ \text{nF}$ installed to stabilize the anode wire voltages represent an effective AC connection to ground.

9. Detector control system (DCS)

The primary task of the Detector Control System (DCS) is to ensure safe and reliable operation of the TPC. It provides remote

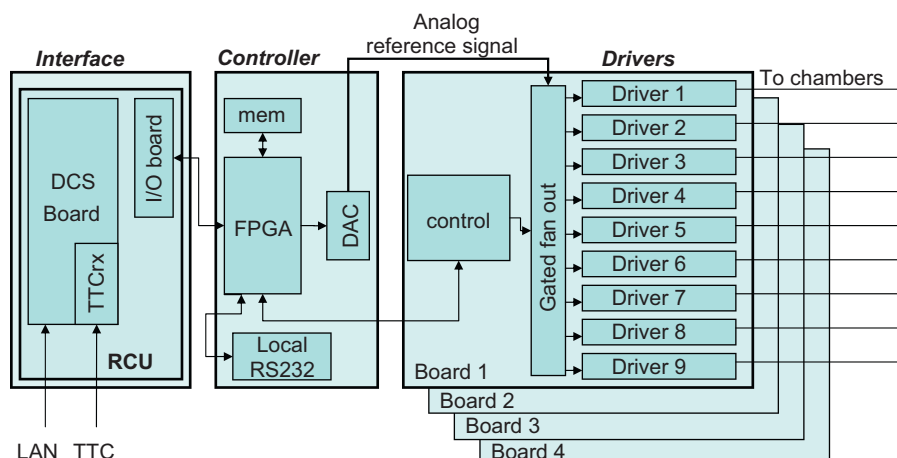


Fig. 58. Schematic setup of the calibration pulser system. Each side of the TPC is equipped with one such system.

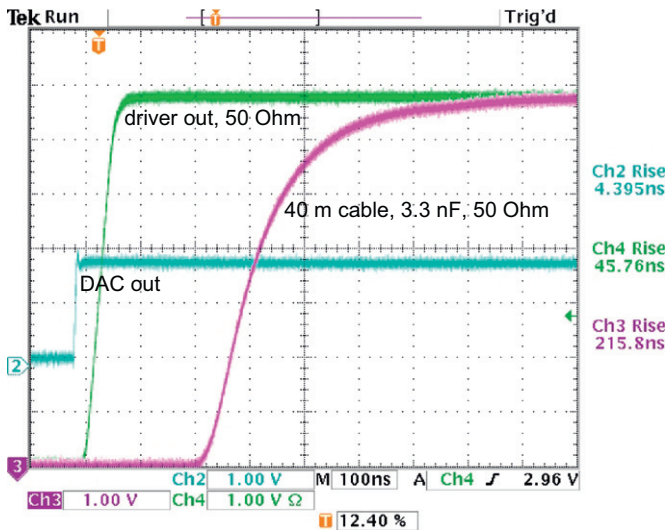


Fig. 59. Signal shapes in the calibration-pulsar system: output of the DAC (channel 2); output of the driver, terminated by 50Ω (channel 4); test setup simulating a chamber using a 40m cable with 3.3nF and 50Ω termination (channel 3). On the right side the measured rise times are indicated.

control and monitoring of all detector equipment in such a way that the TPC can be operated from a single workplace (the ALICE experimental control room at LHC Point 2) through a set of operator interface panels. The system is intended to provide optimal operational conditions so that the data taken by the TPC is of the highest quality. More information about the TPC DCS can be found in Ref. [33].

9.1. Overview

The TPC control system is part of the ALICE DCS [59]. Like the other three ALICE online systems [60] (Data Acquisition system (DAQ), Trigger system (TRG) and High Level Trigger system (HLT)), the ALICE DCS is controlled by the Experiment Control System (ECS), where the ECS is responsible for the synchronization between the four systems. This is schematically shown in Fig. 60.

9.1.1. Hardware architecture

The hardware architecture of the TPC DCS can be divided in three functional layers. The *field layer* contains the actual hardware to be controlled (power supplies, FEE...). The *control layer* consists of devices for collecting and processing information from the field layer and making it available to the *supervisory layer*. At the same time the devices of the control layer receive commands from the supervisory layer to be processed and distributed to the field layer. The equipment in the *supervisory layer* consists of personal computers, providing the user interfaces and connecting to disk servers holding databases for archiving data, etc. The three layers interface mainly through a Local Area Network (LAN).

9.1.2. Software architecture

The software architecture is a tree structure that represents the structure of the TPC, its sub-systems and devices. The structure, as shown in Fig. 61, is composed of nodes, each having a single 'parent', except for the top node called the 'root node'. Nodes may have zero, one or more children. There are two types of nodes, the parent nodes are called Control Units (CU) and the leaf nodes are called Device Units (DU). The control unit controls the sub-tree below it, and the device unit 'drives' a




Fig. 60. Overview of the ALICE online systems. The DCS interfaces to the other systems through the ECS.

device. The behavior and functionality of each control unit is implemented as a finite state machine.

The control system is built using a 'controls framework' that is flexible and allows for easy integration of separately developed components [61,62]. This framework includes drivers for different types of hardware, communication protocols, and configurable components for commonly used applications such as high or low voltage power supplies. The framework also contains many other utilities such as interfaces to the various databases (configuration, archiving), visualization tools, access control, alarm configuration and reporting, etc.

9.1.3. System implementation

The core software of the control system is the commercial SCADA (Supervisory Controls And Data Acquisition) system PVSSII (Prozess Visualisierung und Steuerungs System) from the company ETM [63], which is also used by the other LHC experiments. PVSSII is an object-oriented process visualization and control system that is used in industry and research as well as by the four LHC experiments. PVSSII is event-driven and has a highly distributed architecture. The SCADA System for the TPC is distributed over 12 computers.

9.1.4. Interfaces to devices

Where possible, commercial servers using the OPC standard of process control are used to interface the SCADA system to devices like power supplies. OPC servers interface the ROC high voltage, the field cage high voltage, the front-end electronics low voltage and the temperature monitoring system [33]. The power supplies for the gating grid of the readout chambers and the lasers are controlled via a TCP/IP to RS232 bridges and RS232 interfaces. For non-commercial hardware the communication has been developed based on the communication framework Distributed Information Management (DIM [64]); it is used in the laser system, in the drift velocity monitor, the electronics control and the pulser control.

9.1.5. Interlock

The safety of the equipment and the detector is based on three layers of interlocks:

- **Internal interlock:** The internal mechanism of devices (e.g. power supply trip) are used wherever applicable to guarantee the highest level of reliability and security. The threshold and status of these interlocks are controlled by the SCADA system, but their function is independent of the communication between hardware and software.
- **Cross-system interlock:** The interlocks between different sub-systems are realized by open loop contacts. Programmable Logical Device (PLC) systems are used to delay the signals or to give the possibility to enable or disable these interlocks.
- **Software interlock:** The software interlocks are realized in the supervisory layer. They rely on the communication between

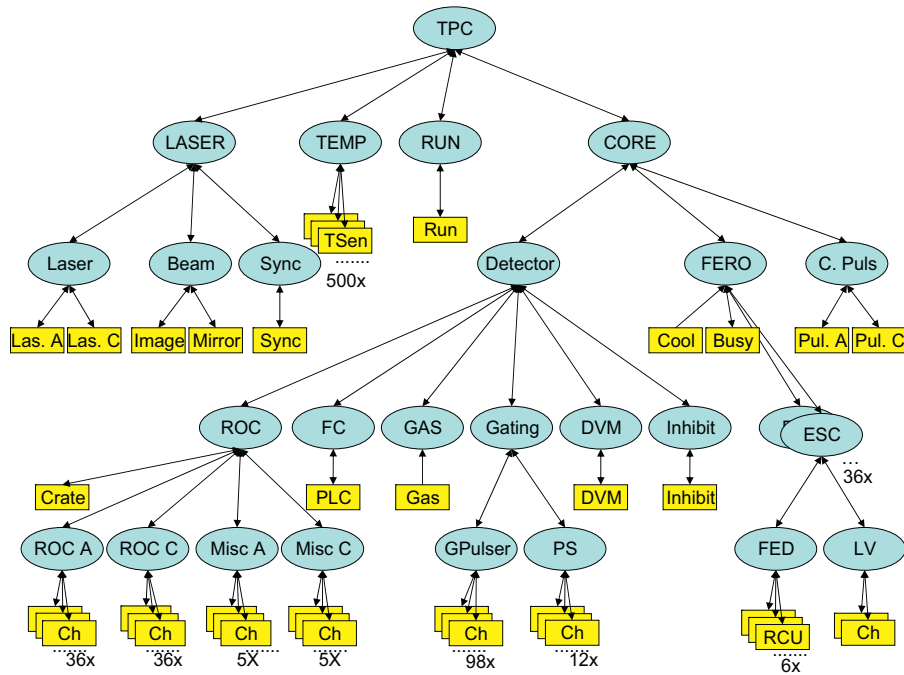


Fig. 61. Overview of the software architecture of the DCS. The tree structure is build out of Device Units (boxes) and Control Units (ellipses).

the hardware and the SCADA system. Therefore, they are only used to prevent the system from unwanted but not harmful events like switching off the power supplies under full load. The safety of the equipment does not rely on the software interlocks.

Internal interlocks are used for the ROC high voltage, the field cage high voltage, the front-end electronics low voltage, the cooling and the gas system. External interlocks (cross-system or software interlocks) are implemented for the field cage high voltage, the front-end electronics low voltage and the cooling system. Software interlocks are used for the ROC high voltage, the front-end electronics low voltage and the front-end electronics [33].

In addition to the interlocks the alert system of the SCADA system is set-up to inform the shifter of unusual or potentially dangerous situations.

9.2. Electronics control

The three functional levels (as described in Section 9.1) for the DCS for the FEE are shown in Fig. 62. The 216 RCUs and the FECs with the ALTRO chips (Section 4) are part of the field layer. The interface between the different layers relies on communication using DIM. Independent, but identical setups are used for the two sides of the TPC.

The FEE communication software for the TPC is described in Ref. [65] and is used for many detectors in the ALICE experiment. A front-end electronics server (FeeServer) interfaces the hardware, publishes monitoring data and receives commands and configuration data. The software can be used for different hardware devices, since it has been divided into a device-independent core and an interface to the hardware dependent functionality (*ControlEngine*).

The TPC FeeServer runs on each of the 216 DCS boards (Section 4) that control the corresponding RCU and FECs. It reads temperatures, voltages and currents on the FECs and a number of

status and configuration registers on the RCU and publishes the acknowledge and message channels. For one readout partition with 25 FECs (for the full TPC) about 165 (30 000) services are published.

The InterComLayer gathers, buffers and bundles monitoring data from the 108 FeeServers and redistributes them to the SCADA system. The services of each FeeServer are grouped in dedicated service channels, reducing the number of services to which the SCADA system has to subscribe to 216 per TPC side. Since the performance of the DIM client in the SCADA system is limited to < 1 kHz, an algorithm is implemented in the InterComLayer that effectively reduces the network traffic to below that limit, while still ensuring that the latest monitoring information is made available. The implementation of the hardware dependent functionality is moved to a configuration database and a separate software component that retrieves the configuration data from the database (*CommandCoder*).

9.2.1. Front-end monitoring

The SCADA system implements the graphical user interfaces to display monitoring data. For the TPC FEE this data is mainly the temperatures, voltages and currents from the FECs and status information of the FEE. In the commissioning phase this functionality allowed the system to identify unwanted voltage drops and insufficiently cooled FECs.

9.2.2. Front-end configuration and control

There are about 5 million configurable parameters for each TPC configuration. The structure of the TPC configuration database follows the structure of the hardware. The parameters for the ALTROS, FECs and RCUs are collected in dedicated tables which are linked via relational tables.

The configurations for the FEE are subdivided in 216 blocks (one block per readout partition). The configurations change over time due to dysfunctional or replaced hardware or due to changing hardware behavior. One complete configuration needs ≈ 30 MByte of space in the configuration database, not including the ALTRO

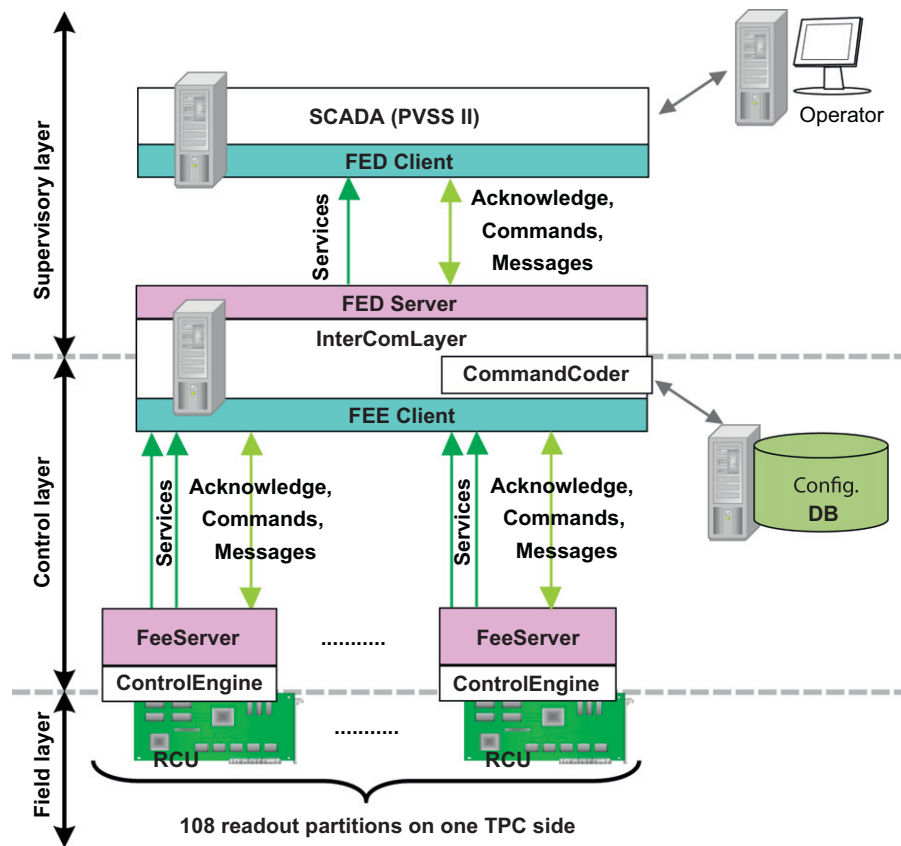


Fig. 62. Schematic setup of the DCS for the TPC FEE. The operator in the ALICE control room can open graphical user interfaces to display the monitoring values as they are provided by two PVSS systems, corresponding to the two sides of the TPC. The operator can also send commands to the control and field layer.

pedestal memories. It is expected that the TPC will need 20–100 different FEE configurations to be stored in the database.

Only simple configuration commands containing a parameter which describes the configuration type are passed from the supervisory layer to the control layer. The actual complex configuration process takes place mainly in the control layer. The CommandCoder assembles queries to the configuration database, which are executed serially and retrieve the configuration data for each readout partition. The InterComLayer sends the data to the FeeServers, which execute the commands that are contained in the configuration data block. The time to compile, send and execute the configurations is of the order of 20 s for 108 readout partitions. A new version of the InterComLayer will make use of parallel invocation of the CommandCoder, which should effectively speed up the configuration process.

The FeeServer stores the latest retrieved configuration data locally on the DCS board making possible a very fast reconfiguration of the FEE, without the overhead of querying the configuration database.

The TPC FEE can also be configured via the optical Detector Data Link (DDL), bypassing the DCS and the FeeServer. This method is used for configuring the ALTRO pedestal memories, where the volume of the configuration data is about 700 MByte for the full TPC. Freshly calculated pedestal and noise values can be used to configure the FEE directly after pedestal data has been acquired.

9.3. Interfaces to experiment control and offline

The DCS interfaces to ECS, which steers the whole experiment. It also interfaces to the offline data analysis framework, since for a proper interpretation of the recorded data the configuration and

status of the FEE and the environmental conditions in the cavern and on/inside the detector have to be known.

A run type defines how the FEE will be configured and which subsystems are to be activated at the start of the run. For the TPC, four run types are of relevance. A *Physics* run is the general run type for recording data with beam–beam collisions. The remaining run types are used for the calibration of the detector. For a *Pedestal* run (typically performed once a day) the FEE is configured to read out black events (no zero suppression) in order to analyze the pedestal and noise values in all channels. For *Laser (Pulsar)* runs the TPC laser system (the calibration pulser system) is activated. The run type is propagated from ECS to DCS, together with the run number and the list of readout partitions which will be read out in the upcoming data taking. This is needed in order to properly configure the BusyBox.

10. Commissioning and calibration

10.1. Calibration requirements

The main goal for the calibration procedures is to provide the information needed for the offline software to reconstruct the position and energy of clusters with sufficient precision so that the design performance can be achieved (see Section 11).

To cope with the huge amount of raw data (about 750 MByte/event), zero suppression is performed on the level of the FEE. The first step in the calibration chain is to obtain the parameters that are uploaded to the FEE and used to process the raw data online. Because the online zero suppression uses a threshold for removing noise, the noise design value is included in the

calibration requirements. The requirements are:

- **Noise:** The TPC and FEE were designed to have an overall Equivalent Noise Charge (ENC) better than $1000e$ in the whole TPC, corresponding to ≈ 1 ADC channel. With the nominal gain (2×10^4) this should give a signal⁴ to noise ratio of about 20 in the IROCs and 30 in the OROCs for minimum ionizing particles. This still leaves a large (order of 30) dynamic range for the larger ionization of low-energy tracks.
- **Gain homogeneity:** The gain has to be calibrated to better than 1.5% over all pads. The time dependence of the residual gain variations can then be obtained from precision measurements of temperature and pressure variations.
- **Space-point resolution:** The systematic contribution of each of the following effects to the space-point resolution has to be kept below $200 \mu\text{m}$:
 - **Drift velocity:** The drift velocity has to be known to a precision better than 10^{-4} . This results from considering a resolution of $200 \mu\text{m}$ over the full drift length (250 cm).
 - **Alignment:** The residual rotation (translation) after alignment has to be kept below 0.1 mrad ($200 \mu\text{m}$).
 - **$E \times B$ effects:** The $E \times B$ effect was estimated from simulations using the measured magnetic field to be of the order of 8 mm for the full drift length. The precision of the correction therefore has to be on the level of 2%.

The alignment and $E \times B$ calibration will be described in detail in a future publication dedicated to the calibration and performance of the TPC.

10.2. Commissioning

10.2.1. Commissioning phases

After the TPC was assembled, the commissioning activities began. The main activities in the commissioning procedures were:

- **Phase 1:** First tests at the surface (2006). For these tests limited services were available so that only 2 sectors could be powered at a time. No zero suppression was applied to the data (so called 'black events').
- **Phase 2:** First commissioning in the ALICE experimental area underground (December 2007). A full side (side C) was successfully operated with final services attached and online zero suppression.
- **Phase 3:** Commissioning under final running conditions (March to September 2008). The TPC was operated stably over several months. Different run types (see Section 10.2.2) were successfully implemented, and extensive calibration data was taken.

The bulk of the calibration data was taken in phase 3, and most of the results shown in the following sections are from this phase.

10.2.2. Data sets

- **Pedestal runs:** This run type, where the zero suppression is switched off, is used to determine the pedestal and noise for all readout channels. The extracted values are used to perform zero suppression with the FEE.
- **Calibration pulser runs:** The cathode wire grids of the ROCs can be pulsed to determine the response of the electronics chain (see Section 8.6). The extracted data is used for drift-time calibration as well as for the detection of dead channels and floating wires.

- **Laser runs:** The TPC laser system (see Section 7) provides well defined straight tracks within the TPC volume. In addition, scattered laser light creates photo-electrons on the central electrode. Both sets of data are used to determine the drift velocity and study the inter-chamber alignment.
- **Cosmic runs:** Cosmic runs were taken primarily with the ACORDE scintillator array [2] as trigger.
- **Krypton runs:** To determine the gain of each individual readout channel with high precision, krypton gas was released into the TPC gas system. Krypton data were accumulated with cosmic-triggered events for three different anode wire voltage settings (gains).

10.3. Electronics calibration

10.3.1. Pedestal and noise determination

In pedestal runs the electronics baseline (pedestal) and its width (noise) are determined. The typical baseline of one electronic channel is displayed in Fig. 63. The insert shows its distribution. The Gaussian mean defines the pedestal value, while the sigma corresponds to the noise. The measured pedestal and noise values are stored in the OCDB [2] for offline reconstruction usage, as well as on the LDC machines to be uploaded into the FEE. The zero-suppression threshold at this stage is set at three sigma of the baseline distribution.

Results of noise measurements: In Fig. 64 the histogrammed noise distributions obtained during the final stage of commissioning are shown for all pads and for the different pad sizes separately. Truncated mean values in the range 0–2 ADC channels and the corresponding RMS of the distributions are summarized in Table 15.

Only 1.7 (0.14)% of the channels show noise values above 1 (1.5) ADC channels. The largest pads with 5.7 (0.4)% contribute most to this value. The mean ENC in the TPC is about $730e$.

A systematic variation in the noise level is observed increasing from the center of each readout partition towards its edges [66]. This variation is directly related to the variation in length of the traces on the pad-plane PCB board, which connects the pads with the connectors on its back side. With the trace length the capacitance at the input of the charge-sensitive preamplifier/shaping chips (PASAs) rises and hence the noise level rises. Fig. 65 shows the dependence of the noise on the trace length for the medium-sized pads.

A detailed discussion on noise measurements in the TPC can be found in Ref. [66].

Improvements to decrease the noise level: Measurements during the first commissioning in 2006 showed that a large fraction of pads ($\approx 10\%$) had noise values above 1 ADC channel. In the largest pads this fraction was $\approx 24\%$. Two modifications reduced the noise to the desired level. First, the start of readout for groups of channels was desynchronized to minimize the peak current drawn by the FEE (reduce ground-bounce effects). Secondly, the grounding scheme for the FECs was revised and optimized in terms of the noise behavior.

10.3.2. Tail-cancellation filter parameter extraction

The signal from a gas detector with a MWPC readout is often characterized by a long tail with a rather complex shape. Detailed simulations can be found in Refs. [17,67].

A Tail-Cancellation Filter (TCF) (see Section 4.3) is implemented in the ALTRO chip for filtering the digital signal after the initial baseline subtraction so that zero suppression can be applied in an efficient way. The TCF is based on the approximation of the tail by the sum of exponential functions. The parameters for the TCF were extracted from non-zero suppressed cosmic data. The method is described in detail in Ref. [68, Chapter 3].

Fig. 66 shows an example of mean pulses in the IROCs at two different gains before and after the TCF was applied. Different

⁴ Here signal means the maximum value of the charge for a pad-timebin cell in a cluster (often denoted Q_{max}).




Fig. 63. Typical electronic baseline of one channel. The insert shows its distribution.




Fig. 65. Correlation of noise and the trace length on the pad-plane PCB board. A straight line was fit to the data.




Fig. 64. Noise distribution; for all chambers and separately for the different pad sizes (© 2009 Elsevier [72]).

Table 15

Truncated mean and RMS of the noise distributions in Fig. 64 (see text).

Pad size (mm ²)	Mean	RMS	> 1(1.5) ADC (%)
4 × 7.5	0.686	0.068	0.419 (0.072)
6 × 10	0.719	0.064	0.244 (0.015)
6 × 15	0.792	0.127	5.692 (0.420)

Values are given separately for the different pad sizes. In addition, the fraction of pads above 1 (1.5) ADC channels is shown.

gain (anode voltage) settings were applied to achieve a detailed ion-tail characterization of the MWPCs.⁵ At an anode wire voltage of 1350 V the ion tail reveals an immediate undershoot after the main peak of the signal (Fig. 66 (bottom)).

The pulse-by-pulse fluctuations were found to have a bigger impact than the pad-per-pad systematics. Differences between medium and long pad sizes in the OROC were not significant either. The geometrical differences between IROC and OROC as well as the different gain settings were found to have the biggest impact on the ion-tail shape. Therefore, shaping parameters for the different sectors were found to be sufficient. Since small gas composition changes do not have a major influence on the shape of the tail either, it is foreseen to redo the shaping-parameter-finding-procedure just once a year in order to study the influence of long-term variations like aging effects on the chambers.

The TCF was not applied for any of the results presented in the following sections. When it is employed, in the future, we anticipate that we will make the time signal more symmetric.

10.4. Gain calibration

10.4.1. Krypton calibration

The krypton calibration method was developed by the ALEPH [51,69] and DELPHI [70] collaborations. It was also successfully applied by the NA49 Collaboration [7]. The advantage of this method is that it provides an absolute calibration of the total gain (electronics and gas amplification) for each pad.

The pad-gain factors are measured from decay clusters of radio-active krypton (⁸³Kr) which is released into the TPC gas. Dedicated krypton data taking is planned once a year.

Analysis: Data samples were collected for three different gains (anode voltage settings) over one week of data taking. In the following paragraphs only results for the case of 1350 V and 1550 V applied to the IROCs and OROCs anode wires, respectively

⁵ With increasing gain, the avalanche size around the multiplication wire increases. Therefore, the ratio between the number of ions, which go to the pad, and the number of ions, which go to the cathode, changes [17].

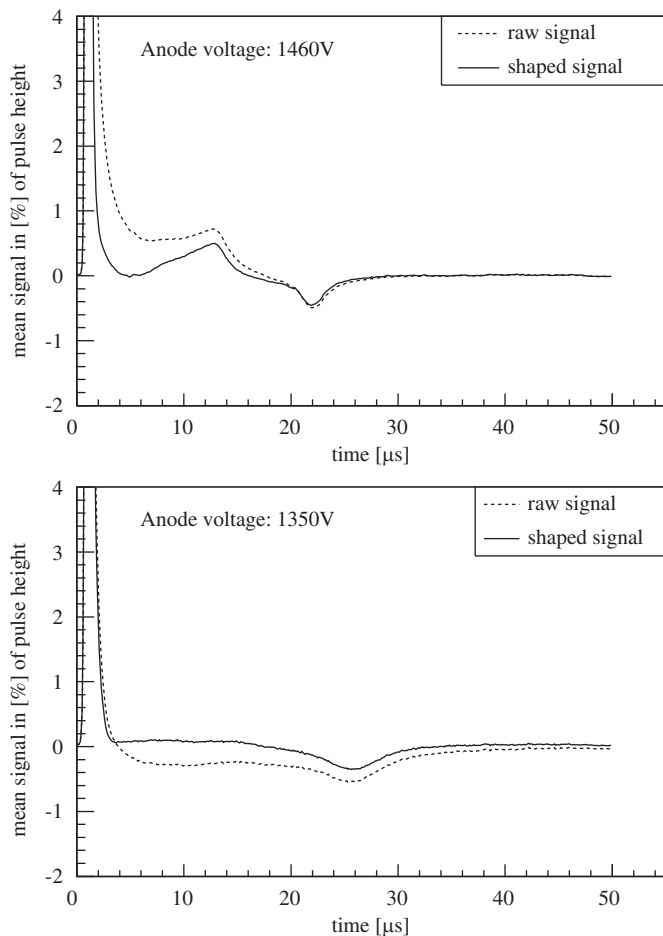


Fig. 66. Mean pulses in the IROC at two different gains, before (raw) and after (shaped) the tail-cancellation filter was applied.

(quoted as ‘nominal voltage setting’) are presented. The gain for the nominal voltage setting is about 6500–7500. A gain curve can be found in Ref. [33]. Events of 11.3 million were collected where each event had roughly 80 krypton clusters.

A dedicated krypton-cluster finder is used to reconstruct the krypton decays. The reconstructed cluster charge is associated with the pad with the maximum amplitude. For each pad the charge spectrum is accumulated. The pad gain is then defined as the mean of a Gaussian fit to the main peak (41.6 keV) of the charge spectrum. The error on the mean obtained from the fit is of the order 0.2% on the single pad level, which is well below the requirements of 1.5% specified in Section 10.1.

Fig. 67 shows an example of the accumulated charge spectrum of all OROCs, corrected for the pad-by-pad variations. The resolution of the main peak for inner and outer ROCs is of the same magnitude, 4.0% for IROCs and 4.3% for OROCs.

Results: The results of the krypton calibration are used to produce pad-by-pad calibration constants which reflect the gain topology.

Gain variations within a chamber reflect mechanical deformations and imperfections (see Refs. [7,71]). The geometrical characteristics of a sector are visible in the radial and azimuthal projections. Fig. 68 shows the average gain variations in radial (top) and azimuthal (bottom) direction over all chambers.

Within a single sector, sizable systematics are observed in the radial direction, typically reaching up to 18%, 23% and 11% for IROCs, OROCs short pads and OROCs long pads, respectively. The maximal variations in the azimuthal direction are 27% for IROCs and 22% for OROCs. A decrease of gain on the edges is visible




Fig. 67. Krypton spectrum of all OROCs at nominal gain (© 2009 Elsevier [72]).

especially in the azimuthal direction. It is related to the fact that on the edges the full krypton cluster cannot be reconstructed. For this reason a parabolic extrapolation of the gain is used for gain correction in these regions.

10.5. Drift-time calibration

10.5.1. Shaping variations in the FEE

To determine the shaping characteristics of the front-end electronics, a pulse is injected on the cathode-wire plane of the readout chambers (see Section 8.6). This induces a signal on the pads, without gas amplification. Due to manufacturing tolerances of the PASA chips the shaping of the signal is expected to vary, resulting in the detection of different arrival times and integrated charges.

In order to correct for these effects, calibration pulser runs (see Section 10.2.2) will be taken on a regular basis to monitor the chip characteristics. The resulting correction values are stored in the OCDB and used in the offline reconstruction.

Fig. 69 shows a typical pulser signal of one channel. A calibration algorithm accumulates a number of pulser events and calculates the position (center of gravity), width (RMS) and area (integral) for each pad signal. The signal analysis is done in a window of minus two to plus two time bins around the maximum bin, as used in the offline cluster finder.

Fig. 70 shows the timing differences within one IROC. Clear patterns can be seen: groups of 16 pads are found, showing nearly the same values. Differences between the groups can be larger. Each of the groups corresponds to one PASA chip. The variations result from manufacturing tolerances.

In Fig. 71 the distribution of the timing variations in the complete TPC is shown. The RMS of the distribution is 0.052, corresponding to 5.2 ns. Considering a drift velocity of about 2.65 cm/μs, this would yield a systematic error in the cluster-position resolution of about 140 μm. Compared to the intrinsic cluster-position resolution of 300–800 μm, given by the diffusion and therefore depending on the z-position, this is a second order effect.




Fig. 68. Radial (top) and azimuthal (bottom) systematics of gain variations within one sector. The step observed in the radial variations is due to the IROC-OROC transition.

Details on the pad-by-pad shaping variations are discussed in Ref. [66].

10.5.2. Drift velocity

The drift velocity is a function of the field (electric, magnetic) and the mobility [12]. The mobility depends on the gas density which is a function of the environment variables as well as the gas composition. The drift velocity is therefore a function of many parameters:

$$v_d = v_d(E, B, N(P, T), C_{CO_2}, C_{N_2}) \quad (1)$$

where E and B are the field values (electric, magnetic), N is the gas density, P is the atmospheric pressure, T is the temperature inside the TPC and C_{CO_2} and C_{N_2} are two concentrations out of three components of the drift gas Ne–CO₂–N₂ (85.7–9.5–4.8). We assume that these parameters, especially the environment variables, will vary in time within a reasonable range. According to *Magboltz-2* [44,45] simulations, a first order Taylor expansion of the dependencies around the nominal values, $\Delta v_d = v_d - v_{d0}$, is sufficient.

Within the TPC volume, the parameters in the expansion are changing with different time constants. A significant change of the drift velocity due to changes in the gas composition as well as E and B field variations has a time constant of several hours, while the changes due to pressure and temperature variations have to be corrected on the level of minutes. In the following we will therefore disentangle the two time scales and summarize the long




Fig. 69. Typical calibration pulser signal in one readout channel.




Fig. 70. Topology of pulser-timing variations in one IROC.

term variations under the term $k_0(t)$:

$$x = \frac{\Delta v_d}{v_{d0}} = k_0(t) + k_N \frac{\Delta N(P, T)}{N_0(P, T)} = k_0(t) + k_{P/T} \frac{\Delta(P/T)}{(P/T)_0} \quad (2)$$

The correction factor x can be measured using different methods:

- matching laser tracks with the surveyed mirror positions;
- matching with tracks from the Inner Tracking System (ITS);
- matching of the TPC primary vertices from the two halves of the TPC; and
- matching tracks from two halves of the TPC using cosmic tracks.

The unknown parameters $k_0(t)$ and $k_{P/T}$ can be determined using a Kalman filter approach.

Precision of the correction: The precision of the drift velocity correction is proportional to the precision of the pressure and temperature measurement (first term) and to the length of the

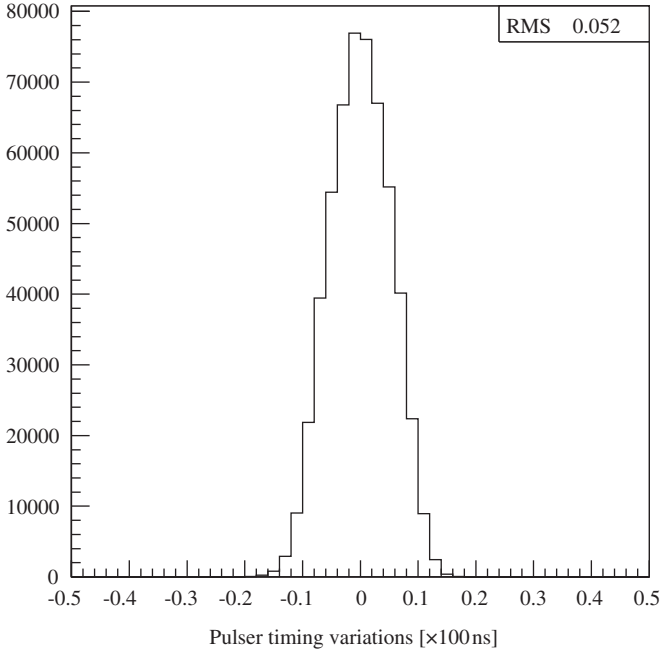


Fig. 71. Distribution of pulser-timing variations of all pads.

time interval between the laser measurements (second term):

$$\sigma_x^2 = \frac{k_{P/T}^2}{(P/T)_0^2} \sigma_{P/T}^2 + \frac{\partial \sigma_{k_0(t)}^2}{\partial t} \Delta t = \dot{\sigma}_{k_0(t)}^2 \Delta t + k'_{P/T} \sigma_{P/T}^2. \quad (3)$$

The typical relative resolution of the pressure and temperature measurement is on the level of 6×10^{-5} and 1×10^{-5} , respectively.⁶ For a cool gas the coefficient k_N is close to unity. The contribution of the P/T correction to the drift velocity uncertainty is approximately 6.1×10^{-5} (150 μm for the full drift length of 250 cm).

Fig. 72 shows the input to the Kalman filter and Fig. 73 shows the results after drift velocity correction.

A similar method will be applied to correct for the time variation of the gain.

The $\dot{\sigma}_{k_0(t)}$ from Eq. (3) was estimated from Fig. 73 and is on the level of 0.001 in a four day period. This estimate was obtained for the period of largest change in the present data sample. Further studies will be performed for extended time periods.

For the TPC drift velocity determination, the required relative resolution is on the level of 6×10^{-5} . Entering the observed sigmas into Eq. (3) the minimal frequency of the drift velocity updates were estimated to be about 1 h.

11. Performance

The ALICE TPC is the main tracking device of the experiment, therefore its performance is a crucial issue. In this section we discuss the space-point resolution, which mainly determines the tracking performance (momentum and angular resolutions), the track matching performance and the particle identification performance.

⁶ This is the resolution of the pressure sensor used in 2008. As this resolution was close to the required drift velocity precision (see Section 10.1) it has now been replaced with a sensor with higher precision.




Fig. 72. Drift velocity as function of time (top) and as a function of $\Delta(T/P)$ (bottom).

11.1. Space-point resolution

In general, the space-point resolution (σ_{COG}) depends on many parameters; namely, the readout geometry, the gas composition and the track characteristics. Here we discuss those, which are dominant for the present detector, namely:

- the drift length (L_{Drift});
- the track inclination angle (α); and
- the charge deposited on the anode wire (Q).

The space-point resolutions presented in this section were determined as a function of these values.

For further studies, it is convenient to parametrize the space-point resolution as a function of the parameters mentioned above. We obtained the parametrization of the space-point resolution by fitting parameters p_0 , p_L and p_A in the formula below to data with cosmic ray measurements.

$$\begin{aligned} \sigma_{\text{COG}}^2 &\propto p_0^2 + p_L^2 L_{\text{Drift}} + p_A^2 \tan^2 \alpha \\ p_L^2 &\propto \frac{\sigma_D^2 G_g}{N_{\text{ch}}} \\ p_A^2 &\propto \frac{L_{\text{pad}}^2 G_{\text{Lfactor}}}{N_{\text{eprim}}} \end{aligned} \quad (4)$$

where N_{ch} is the number of electrons created during the amplification process, L_{pad} is the pad length, and N_{eprim} is the number of primary electrons per pad. There are three main factors which degrade the space-point resolution, namely the gas gain




Fig. 73. Drift velocity corrected for P/T variations as function of time (top). In the lower plot the correction for time-dependent offset is also applied.

fluctuations (factor $G_g \approx 2$), the Landau fluctuations of the ionization energy loss (factor $G_{L\text{factor}}$) and the diffusion σ_D (here, normalized to the drift length). One should note that:

$$N_{\text{ch}} \propto L_{\text{pad}}$$

$$N_{\text{eprim}} \propto L_{\text{pad}} \quad (5)$$

and therefore

$$p_L \propto \frac{1}{\sqrt{L_{\text{pad}}}}$$

$$p_A \propto \sqrt{L_{\text{pad}}} \quad (6)$$

Analysis of cosmic ray data determines the resolution parameters p_0 , p_L and p_A . Results from a fit which are scaled according to Eq. (6) are shown in Table 16. The fit was done separately in the z (drift) and the $r\phi$ directions, denoted further as z and y accordingly.

In the ALICE TPC three different pad geometries are used, thus the space-point resolution was obtained for each of them separately. One should note that the scaled values of $p_{Ly} \sqrt{L_{\text{pad}}}$, $p_{Lz} \sqrt{L_{\text{pad}}}$ and, separately, $p_{Ay} / \sqrt{L_{\text{pad}}}$, $p_{Az} / \sqrt{L_{\text{pad}}}$ are equal to within 10%, as expected.

In the previous formula we assumed that all electrons created in the ionization process contribute to the measured signal. In a real experiment, because of the applied zero suppression, part of the signal is lost. The fraction of the signal below the threshold is proportional to the width of the response function and increases

Table 16

Values of parameters describing the space-point resolution.

Pad size	$0.75 \times 0.4 \text{ cm}^2$	$1.0 \times 0.6 \text{ cm}^2$	$1.5 \times 0.6 \text{ cm}^2$
p_{0y} (cm)	0.026	0.031	0.023
p_{0z} (cm)	0.032	0.032	0.028
$p_{Ly} \sqrt{L_{\text{pad}}}$ (cm)	0.0051	0.0060	0.0059
$p_{Lz} \sqrt{L_{\text{pad}}}$ (cm)	0.0056	0.0056	0.0059
$p_{Ay} / \sqrt{L_{\text{pad}}}$ ($\text{cm}^{1/2}$)	0.13	0.15	0.15
$p_{Az} / \sqrt{L_{\text{pad}}}$ ($\text{cm}^{1/2}$)	0.15	0.16	0.17

with the drift length and the track inclination angle. We have corrected for these effects, replacing p_L and p_A from Eq. (4):

$$p'_L \propto p_L \cdot p_{LC} = p_L \cdot (1 + p_{L1} \cdot L_{\text{Drift}} + p_{L2} \tan^2 \alpha)$$

$$p'_A \propto p_A \cdot p_{AC} = p_A \cdot (1 + p_{A1} \cdot L_{\text{Drift}} + p_{A2} \tan^2 \alpha) \quad (7)$$

where p_{L1} , p_{L2} , p_{A1} and p_{A2} are parameters to be fitted. We also added to the Eq. (7) terms proportional to $1/Q$, where Q is a total charge of the cluster, to account for the number of electrons which contribute to the signal. However, the space-point resolution improves only slightly, and after a certain value of Q deteriorates (see Fig. 74). This is due to the production of δ -electrons. The influence of δ -electrons is much smaller in presence of the magnetic field because of their smaller effective range.

The measured space-point resolution in y ($r\phi$) and z (drift) directions are shown in Fig. 75. The parametrization (see Eq. (7)) describes the data within 2%.

The observed dependence of the space-point resolution on drift length and inclination angle is, for these cosmic ray tracks, mostly determined by geometrical factors. The values observed for small inclination angles are close to those specified in the TPC TDR [3]. Note that small inclination angles are common for tracks originating from the collision vertex of an LHC pp or Pb–Pb event.

11.2. Momentum resolution

The momentum resolution achievable with the ALICE TPC used as a stand-alone detector can be determined by using cosmic ray tracks passing through the center of the TPC to simulate tracks coming from an interaction vertex. Comparison of the momenta for the first and second half of each track yields the momentum resolution curve depicted in Fig. 76. At the current stage of calibrations a momentum resolution of better than 7% is reached at 10 GeV close to the value listed in the TDR [3]. We are currently continuing to improve the correction of various (small) distortions and further improvements are expected.

11.3. Particle identification performance

The simultaneous measurement of the momentum p of a particle and its specific ionization loss in the TPC gas provides particle identification over a wide momentum range. In practice, only relative values of the ionization need to be known to distinguish between different particle species. The dE/dx information for a given track must be extracted from the n_{cl} clusters ($50 < n_{cl} < 160$), which are assigned to the track. For each cluster its maximal charge Q_{max} (the highest ADC value) and its total charge Q_{tot} can be obtained. The question of whether the dE/dx information should be extracted from Q_{max} or Q_{tot} depends on the

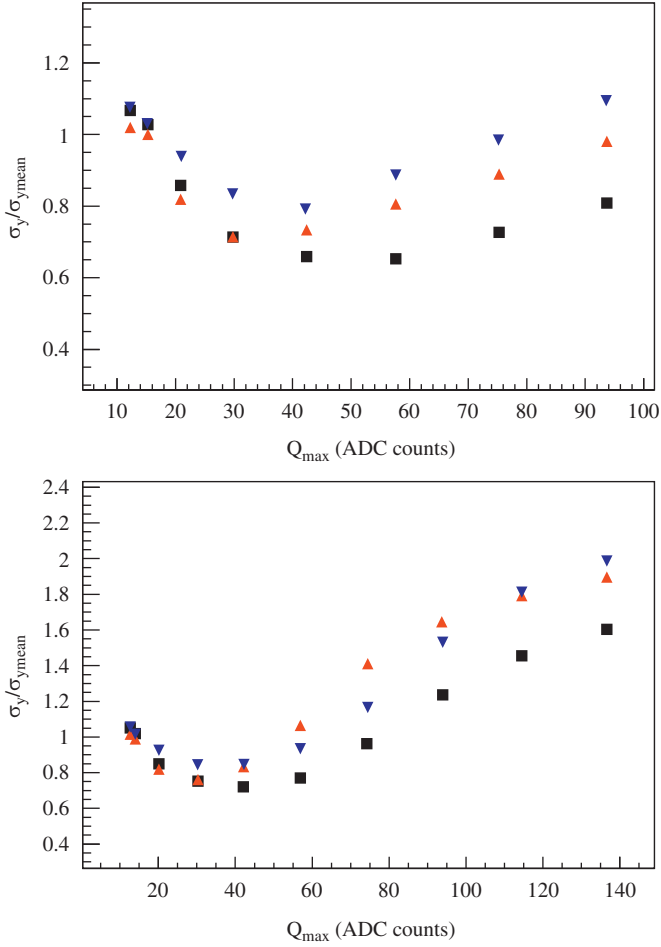


Fig. 74. Space-point resolution in y ($r\phi$) direction as a function of the maximum of deposited charge within a cluster Q_{\max} , with (upper panel) and without (lower panel) magnetic field. The different curves correspond to the three pad regions short (squares), medium (triangles), and long (inverted triangles).

signal to noise ratio. Results shown here are based on evaluations from Q_{tot} .

Because of the long tail towards higher energy losses in the straggling function, the average energy loss is not a good estimator as it would be for a Gaussian distribution. Therefore, the so-called truncated mean $\langle Q \rangle_\eta$, also called TPC signal, is used [12, Section 10.3]. It is defined as the average over m lowest values, which correspond to the η -fraction of the whole sample,

$$\langle Q \rangle_\eta = \frac{1}{m} \sum_{i=1}^m Q_i \quad (8)$$

with $Q_i \leq Q_{i+1}$ for $i = 1, \dots, n-1$ and $m = \lceil \eta n \rceil$. Values of $\langle Q \rangle_\eta$ follow an almost perfect Gaussian distribution. At present, the value of η is set to 0.7, as the result of an optimization process, but will be a subject of further investigation. The dE/dx or the energy loss resolution $\sigma_{dE/dx}$ is given by the variance of the Gaussian distribution of $\langle Q \rangle_\eta$.

Fig. 77 shows the TPC signal of cosmic tracks versus their momentum, from 8.3 million events. For these data, the maximal inclination angle of $\tan(\alpha) < 1$ and at least 120 out of 160 possible TPC points per track were required. Characteristic bands for various particles (electrons, muons, protons, deuterons) are clearly visible. The energy loss is described by the Bethe–Bloch function:

$$\left\langle \frac{dE}{dx} \right\rangle = \frac{4\pi N e^4 Z^2}{m c^2 \beta^2} \left(\ln \frac{2 m c^2 \beta^2 \gamma^2}{I} - \beta^2 - \frac{\delta(\beta)}{2} \right) \quad (9)$$

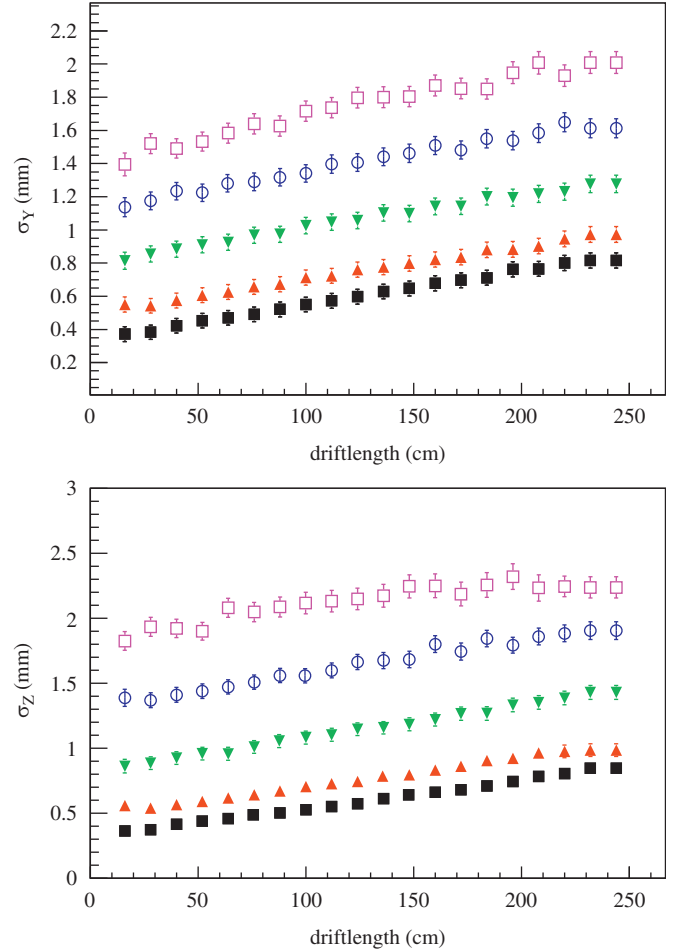


Fig. 75. Space-point resolution in y ($r\phi$) and z (drift) direction as a function of the drift length. The different symbols correspond to inclinations from $\tan(\alpha) = 0$ (full squares) to $\tan(\alpha) = 0.92$ (open squares) in steps of 0.23.

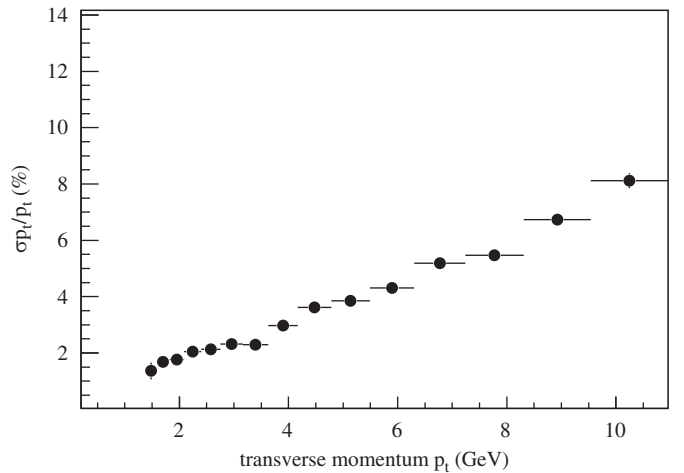


Fig. 76. Transverse momentum resolution measured with cosmic rays.

where $m c^2$ is the rest energy of the electron, Z the charge of the projectile, N the number density of electrons in the traversed matter, e the elementary charge, β the velocity of the projectile and I is the mean excitation energy of the atom. In the analysis of experimental data, other parametrizations than the Bethe–Bloch function are often used. Here we use the parametrization proposed by the ALEPH




Fig. 77. dE/dx spectrum of cosmic rays (© 2009 Elsevier [72]).




Fig. 78. Dependence of the dE/dx -resolution on the number of TPC track points measured with cosmic tracks.

experiment of the form [12]:

$$f(\beta\gamma) = \frac{P_1}{\beta^{P_4}} \left(P_2 - \beta^{P_4} - \ln \left(P_3 + \frac{1}{(\beta\gamma)^{P_5}} \right) \right). \quad (10)$$

They are shown as lines in Fig. 77.

The decisive quantity for particle identification is the resolution $\sigma_{dE/dx}$ of the dE/dx -measurement. Assuming a perfect gain calibration, it depends on the number of samples n , the pad size x and the gas pressure p . In a given gas cell, the energy loss distribution depends only on the cluster-size distribution and on the number of primary interactions in the gas. This implies that the ionization distribution varies with p in the same way as it does with x and therefore the width of the distribution scales inversely proportional with the product xp .

For the remaining dependence on n we expect a statistical scaling according to the law $\sigma_{dE/dx} \propto 1/\sqrt{n}$. In addition to this, the measurement of the energy loss is influenced by systematic uncertainties σ_{syst} . Therefore, the overall resolution is assumed to be of the form

$$\sigma_{dE/dx}^2(n) = \sigma_{\text{syst}}^2 + \frac{\sigma_{\text{stat}}^2}{n}. \quad (11)$$

A measurement of this dependence with cosmic tracks is shown in Fig. 78.

The results demonstrate that the energy loss resolution reaches 5% for cosmic tracks with 160 clusters, (corresponding to about 1.5 times the minimum ionizing energy loss) which is close to and actually slightly better than the design value. In summary, space-point and dE/dx resolutions as specified in the TPC TDR [3] have been reached with the ALICE TPC.

12. Conclusions

The ALICE TPC has been constructed over a period of five years between 2002 and 2006 with most of the assembly taking place in a large clean room located above ground near the ALICE experiment at CERN. In January 2007 it was transported into the ALICE underground cavern and installed into the ALICE experiment. In late 2007 and 2008 there were extensive campaigns to test all components and the full system, using pulsers, laser beams, and cosmic rays. In September 2008 the TPC was ready for first collisions. After the LHC incident [73] a major effort started in late 2008 to improve accessibility of the TPC electronics with all other ALICE detectors installed. Since August 2009 the TPC is in full operating mode. As described in this paper, all systems perform close to or even exceed specifications and the calibration scheme is sufficiently advanced that momentum resolution of better than 6% at 10 GeV and dE/dx resolution of better than 5% are reached with the TPC alone. Calibration is further continuing with cosmic rays and the TPC team very much looks forward to taking first data with proton–proton and Pb–Pb collisions at high rate (around 1400 and 300 Hz, respectively) in the coming runs.

Acknowledgments

The ALICE TPC Collaboration wishes to acknowledge significant contributions to the TPC project given by: G. Augustinski, L. Bozyk, B. Cantin, H.-C. Christensen, H.W. Daus, A. Ferrand, F. Formenti, D. Fraissard, S. Gärtner, J. Hehner, S. Lang, J. Lien, A. Martinsen, T. Morhardt, A. Przybyła, A. Rosseboe, A. Stangeland, R. Toft-Petersen, T. Weber.

In particular we wish to thank J.H. Thomas for his very helpful suggestions and careful proofreading.

We acknowledge the support of the following funding agencies:

- CERN;
- Danish National Science Research Council and the Carlsberg Foundation;
- German BMBF and the Helmholtz Association;
- Research Council of Norway;
- Polish Ministry of Science and Higher Education;
- Ministry of Education of the Slovak Republic; and
- Swedish Research Council (VR) and Knut & Alice Wallenberg Foundation (KAW).

Open Access

This article is published Open Access at sciencedirect.com. It is distributed under the terms of the Creative. This article is distributed under the terms of the Creative Commons Attribution License 3.0, which permits unrestricted use, distribution, and reproduction in any medium, provided the original authors and source are credited.

References

- [1] ALICE Collaboration, Technical proposal for a large ion collider experiment at the CERN LHC, CERN/LHCC 95-71, 1995.
- [2] ALICE Collaboration, *J. Instr.* 3 (2008) S08002. doi:10.1088/1748-0221/3/08/S08002.
- [3] ALICE Collaboration, Technical Design Report of the Time Projection Chamber, CERN/LHCC 2000-001, 2000.
- [4] N. Armesto, et al., *J. Phys. G* 35 (2008) 054001. doi:10.1088/0954-3899/35/5/054001 (arXiv:0711.0974).
- [5] ALICE Collaboration, *J. Phys. G* 1517 (2004) 30. doi:10.1088/0954-3899/30/11/001.
- [6] ALICE Collaboration, *J. Phys. G* 32 (2006) 1295. doi:10.1088/0954-3899/32/10/001.
- [7] S. Afanasev, et al., NA49 Collaboration, *Nucl. Instr. and Meth. A* 430 (1999) 210. doi:10.1016/S0168-9002(99)00239-9.
- [8] D. Vranić, Drift distortions in Alice TPC field cage, ALICE-INT-1997-22, CERN-ALICE-INT-1997-22, 1997.
- [9] H. Wieman, et al., *Nucl. Phys. A* 638 (1998) 559. doi:10.1016/S0375-9474(98)00385-6.
- [10] R. Janik, M. Pikna, B. Sitar, P. Strmen, I. Szarka, *Nucl. Instr. and Meth. A* 598 (2009) 681.
- [11] F. Sauli, *Nucl. Instr. and Meth. A* 522 (2004) 93.
- [12] W. Blum, W. Riegler, L. Rolandi, Particle Detection with Drift Chambers, second ed., Springer-Verlag, 2008.
- [13] U. Frankenfeld, et al., The ALICE TPC inner readout chamber: results of beam and laser tests, ALICE-INT-2002-030, 2002.
- [14] H. Stelzer, et al., The ALICE TPC readout chamber: from prototypes to series production, ALICE-INT-2003-017, 2003.
- [15] P. Christiansen, et al., Performance test of the ALICE TPC IROC module at the PS, ALICE-INT-2005-029, 2005.
- [16] B. Mota, et al., *Nucl. Instr. and Meth. A* 535 (2004) 500. doi:10.1016/j.nima.2004.07.179 in: Proceedings of the 10th International Vienna Conference on Instrumentation.
- [17] S. Rossegger, W. Riegler, Signal shapes in a TPC Wire Chamber, ALICE-INT-2009-038, 2009.
- [18] H.K. Soltveit, The preamplifier-shaper for the ALICE TPC-detector, ALICE-INT-2009-039, 2009.
- [19] R. Esteve Bosch, A. Jimenez de Parga, B. Mota, L. Musa, *IEEE Trans. Nucl. Sci.* NS-50 (2003) 2460. doi:10.1109/TNS.2003.820629.
- [20] STMicroelectronics, TSA1001 product information and data sheet, 2001 <http://www.st.com/stonline/books/pdf/docs/7333.pdf>.
- [21] G. Rybun, C. Soós, ALICE DDL—hardware guide for the front-end designers, ALICE-INT-1998-21, 2007.
- [22] Xilinx Inc., Virtex-II Pro and Virtex-II Pro X Platform FPGAs: Complete Data Sheet, DS083 (v4.7), 2007, Available from: <http://www.xilinx.com>.
- [23] Xilinx Inc., Virtex-II Pro and Virtex-II Pro X FPGA User Guide, UG012 (v4.2), 2007, Available from: <http://www.xilinx.com>.
- [24] B. Taylor, Timing distribution at the LHC, in: Eighth Workshop on Electronics for LHC, Colmar, 2002, pp. 63–74.
- [25] Actel Corporation, Actel ProASICPLUS Flash Family FPGAs, v5.5, 2007, Available from: <http://www.actel.com>.
- [26] T. Krawutschke, Reliability and Redundancy of an Embedded System used in the Detector Control System of the ALICE Experiment, Ph.D. Thesis, University of Applied Sciences Cologne, University of Heidelberg, Germany, 2008.
- [27] T. Krawutschke, International Conference on Field Programmable Logic and Applications 155 (2008) ISBN: 978-1-4244-1960-9.
- [28] Altera, Excalibur devices, Hardware Reference Manual, Version 3.1, 2002, Available from: <http://www.altera.com>.
- [29] A. Bhasin, et al., Implementation of the ALICE Trigger System, in: Real-Time Conference, 2007 15th IEEE-NPSS, 2007, pp. 1–8, doi:10.1109/RTC.2007.4382861.
- [30] A. Morsch, B. Pastirčák, Radiation in ALICE detectors and electronic racks, ALICE-INT-2002-028, 2002.
- [31] K. Roed, Single event upsets in SRAM FPGA based readout electronics for the Time Projection Chamber in the ALICE experiment, Ph.D. Thesis, University of Bergen, 2009.
- [32] ALICE TPC Collaboration, The ALICE TPC, a large 3-dimensional tracking device with fast read-out for ultra-high multiplicity events, ALICE-INT-2009-034, 2009.
- [33] J. Wiechula, et al., *Nucl. Instr. and Meth. A* 548 (2005) 582. doi:10.1016/j.nima.2005.05.031.
- [34] J. Wiechula, Präzisionsmessung der Elektronen-Driftgeschwindigkeit in NeCO₂, Master's Thesis, Universität Frankfurt, 2004.
- [35] A. Müller, The CFD study of the Alice L3 volume ventilation, CFD-2004-05, 2004.
- [36] P. Bonneau, M. Bosteels, Liquid cooling systems (LCS2) for LHC detectors, Prepared for Fifth Workshop on Electronics for the LHC Experiments (LEB 99), Snowmass, Colorado, 20–24 September 1999.
- [37] M. Pimenta dos Santos, ALICE TPC Readout Chambers Cooling System, CERN-ST/ CV-2003-490540, 2003.
- [38] S. Popescu, U. Frankenfeld, H.R. Schmidt, *IEEE Trans. Nucl. Sci.* NS-52 (2005) 2879. doi:10.1109/TNS.2005.862794.
- [39] U. Frankenfeld, S. Popescu, H. Schmidt, Experimental evaluation of the ALICE TPC front-end electronics cooling strategy, ALICE-INT-2005-001, 2005.
- [40] U. Frankenfeld, S. Popescu, H. Schmidt, Temperature monitoring system for the ALICE TPC, ALICE-EN-2005-001, 2005.
- [41] R. Veenhof, Choosing a gas mixture for the ALICE TPC, ALICE-INT-2003-29, 2003.
- [42] C. Garabatos, *Nucl. Instr. and Meth. A* 535 (2004) 197. doi:10.1016/j.nima.2004.07.127.
- [43] S.F. Biagi, *Nucl. Instr. and Meth. A* 421 (1999) 234. doi:10.1016/S0168-9002(98)01233-9.
- [44] S. Biagi, Magboltz-2, Transport of electrons in gas mixtures <http://consult.cern.ch/writeup/magboltz>.
- [45] G. Thomas, et al., LHC-GCS: a framework for the production of 23 homogeneous control systems, CERN-OPEN-2005-031, 2005.
- [46] A. Lebedev, *Nucl. Instr. and Meth. A* 478 (2002) 163. doi:10.1016/S0168-9002(01)01747-8.
- [47] J. Abele, et al., STAR Collaboration, *Nucl. Instr. and Meth. A* 499 (2003) 692. doi:10.1016/S0168-9002(02)01966-6.
- [48] H.J. Hilke, *Nucl. Instr. and Meth. A* 252 (1986) 169. doi:10.1016/0168-9002(86)91177-0.
- [49] D. Miskowiec, P. Braun-Munzinger, *Nucl. Instr. and Meth. A* 593 (2008) 188. doi:10.1016/j.nima.2008.02.034 arXiv:0801.4920.
- [50] D. Decamp, et al., *Nucl. Instr. and Meth. A* 294 (1990) 121. doi:10.1016/0168-9002(90)91831-U.
- [51] Digi PortServer TS 16, Digi International Inc., 11001 Bren Road East, Minnetonka, MN 55343, USA.
- [52] Intelligent Picomotor (iPico), New Focus, 2584 Junction Avenue, San Jose, CA 95134, USA <www.newfocus.com>.
- [53] M3185A B/W Camera Module, Quasar Electronics Ltd, PO Box 6935, Bishops Cleeve, CM23 4WP, UK <www.quasarelectronics.com>.
- [54] FALCONquattro Express, Imaging Development Systems GmbH, Dimbacher Strasse 6, 74182 Obersulm, Germany <www.ids-imaging.com>.
- [55] W-IE-NE-R, Plein & Baus GmbH, Müllersbaum 20, D - 51399 Burscheid, Germany.
- [56] iseg Spezialelektronik GmbH, Bautzner Landstrae 23, D-01454 Radeberg, Germany.
- [57] Zentro-Elektrik GmbH KG, Sandweg 20, D-75179 Pforzheim, Germany.
- [58] A. Augustinus, et al., *ICFA Beam Dyn. Newslett.* 47 (2008) 90.
- [59] ALICE Collaboration, Technical Design Report of the Trigger, Data Acquisition, High Level Trigger and Control System, CERN-LHCC-2003-062, 2004.
- [60] S. Schmeling, *IEEE Trans. Nucl. Sci.* NS-53 (2006) 970. doi:10.1109/TNS.2006.873706.
- [61] B. Franek, C. Gaspar, *IEEE Trans. Nucl. Sci.* NS-45 (1998) 1946. doi:10.1109/23.710969.
- [62] ETM professional control GmbH, A Siemens Company, Kasernenstrae 29 A-7000, Eisenstadt, Austria.
- [63] C. Gaspar, M. Dönszelmann, P. Charpentier, *Comput. Phys. Commun.* 140 (2001) 102.
- [64] S. Bablok, et al., *Nucl. Instr. and Meth. A* 557 (2006) 631. doi:10.1016/j.nima.2005.11.208.
- [65] J. Wiechula, Commissioning and calibration of the ALICE-TPC, Ph.D. Thesis, Goethe-Universität Frankfurt am Main, 2009.
- [66] B. Mota, et al., *Nucl. Instr. and Meth. A* 535 (2004) 500. doi:10.1016/j.nima.2004.07.179.
- [67] S. Rossegger, Simulation and calibration of the ALICE TPC including innovative space charge calculations, Ph.D. Thesis, University of Technology, Graz, CERN-THESIS-2009-124, 2009.
- [68] W. Blum, The ALEPH Handbook, CERN, Geneva, CERN-ALEPH-89-077, 1989.
- [69] A. De Min, et al., *IEEE Trans. Nucl. Sci.* NS-42 (1995) 491. doi:10.1109/23.467923.
- [70] A. Rybicki, Charged hadron production in elementary and nuclear collisions at 158 GeV/c, Ph.D. Thesis, H. Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, Krakow, 2002.
- [71] J. Wiechula, *Nucl. Phys. A* 830 (2009) 531c. doi:10.1016/j.nuclphysa.2009.10.046 (Quark Matter 2009—The 21st International Conference on Ultrarelativistic Nucleus–Nucleus Collisions).
- [72] Summary of the analysis of the 19 September 2008 incident at the LHC <http://cdsweb.cern.ch/record/1135729/>, 2008.