

## Chapter 3


# Central detectors

### 3.1 Inner Tracking System (ITS)

The main tasks of the Inner Tracking System (ITS) are to localize the primary vertex with a resolution better than  $100 \mu\text{m}$ , to reconstruct the secondary vertices from the decays of hyperons and D and B mesons, to track and identify particles with momentum below  $200 \text{ MeV}/c$ , to improve the momentum and angle resolution for particles reconstructed by the Time-Projection Chamber (TPC) and to reconstruct particles traversing dead regions of the TPC. The ITS therefore contributes to practically all physics topics addressed by the ALICE experiment, as discussed in detail in [20].

The ITS surrounds the beam pipe, for which it provides the mechanical support so that no relative movement will take place during operation. The beam pipe is a  $800 \mu\text{m}$ -thick beryllium cylinder of 6 cm outer diameter, coaxial with the ITS detector layers. As shown schematically in figure 3.1 the ITS consists of six cylindrical layers of silicon detectors, located at radii between 4 and 43 cm. It covers the rapidity range of  $|\eta| < 0.9$  for all vertices located within the length of the interaction diamond ( $\pm 1\sigma$ , i.e.  $\pm 5.3 \text{ cm}$  along the beam direction). The number, position and segmentation of the layers were optimized for efficient track finding and high impact-parameter resolution. In particular, the outer radius is determined by the necessity to match tracks with those from the TPC, and the inner radius is the minimum allowed by the radius of the beam pipe. The first layer has a more extended pseudo-rapidity coverage ( $|\eta| < 1.98$ ) to provide, together with the Forward Multiplicity Detectors (FMD), continuous coverage for the measurement of charged-particles multiplicity. The detectors and front-end electronics are held by lightweight carbon-fibre structures. The geometrical dimensions and the technology used in the various layers of the ITS are summarized in table 3.1.

Because of the high particle density expected in heavy-ion collisions at LHC, (as many as 50 particles per  $\text{cm}^2$  have been predicted for the inner layer), and in order to achieve the required impact parameter resolution, Silicon Pixel Detectors (SPD) have been chosen for the innermost two layers, and Silicon Drift Detectors (SDD) for the following two layers. The two outer layers, where the track density is expected to be below one particle per  $\text{cm}^2$ , are equipped with double-sided Silicon micro-Strip Detectors (SSD). The four outer layers have analogue readout and therefore can be used for particle identification via  $dE/dx$  measurement in the non-relativistic ( $1/\beta^2$ ) region. The analogue readout has a dynamic range large enough to provide the  $dE/dx$  measurement for low-



**Figure 3.1:** Layout of the ITS.

**Table 3.1:** Dimensions of the ITS detectors (active areas).

Layer	Type	$r$ (cm)	$\pm z$ (cm)	Area (m <sup>2</sup> )	Channels
1	pixel	3.9	14.1	0.07	3 276 800
2	pixel	7.6	14.1	0.14	6 553 600
3	drift	15.0	22.2	0.42	43 008
4	drift	23.9	29.7	0.89	90 112
5	strip	38.0	43.1	2.20	1 148 928
6	strip	43.0	48.9	2.80	1 459 200
Total area				6.28	

momentum, highly ionising particles, down to the lowest momentum at which tracks can still be reconstructed. This feature gives the ITS stand-alone capability as a low- $p_t$  particle spectrometer. The main parameters for each of the three detector types are summarized in table 3.2. The overall cooling system of the ITS outer layers has been designed [28–30] to satisfy the stringent requirements of the TPC in terms of temperature stability and uniformity.

The momentum and impact parameter resolution for low-momentum particles are dominated by multiple scattering effects in the material of the detector; therefore the amount of material in the active volume has been kept to a minimum. The silicon detectors used to measure ionisation densities (drift and strips) must have a minimum thickness of approximately  $300 \mu\text{m}$  to provide acceptable signal-to-noise ratio. In addition, detectors must overlap to cover the solid angle entirely. Therefore, the detectors effective thickness amounts to 0.4% of  $X_0$ . The additional material in the active volume, i.e. electronics, cabling, support structure, and cooling system, has been kept at a comparable effective thickness (table 3.3). The values in the table are for the detector as built, and they are remarkably close to the goals set at the time of the design [11]. Figure 3.2 shows the integral of the material traversed by a particle crossing perpendicularly the ITS as a function

**Table 3.2:** Parameters of the various detector types. A module represents a single sensor element.

Parameter		Silicon Pixel	Silicon Drift	Silicon Strip
Spatial precision $r\phi$	( $\mu\text{m}$ )	12	35	20
Spatial precision $z$	( $\mu\text{m}$ )	100	25	830
Two track resolution $r\phi$	( $\mu\text{m}$ )	100	200	300
Two track resolution $z$	( $\mu\text{m}$ )	850	600	2400
Cell size	( $\mu\text{m}^2$ )	$50 \times 425$	$202 \times 294$	$95 \times 40000$
Active area per module	( $\text{mm}^2$ )	$12.8 \times 69.6$	$72.5 \times 75.3$	$73 \times 40$
Readout channels per module		40 960	$2 \times 256$	$2 \times 768$
Total number of modules		240	260	1698
Total number of readout channels	(k)	9 835	133	2608
Total number of cells	(M)	9.84	23	2.6
Max. occupancy for central Pb-Pb (inner layer)	(%)	2.1	2.5	4
Max. occupancy for central Pb-Pb (outer layer)	(%)	0.6	1.0	3.3
Power dissipation in barrel	(W)	1350	1060	850
Power dissipation end-cap	(W)	30	1750	1150


**Table 3.3:** ITS material budget traversed by straight tracks perpendicularly to the detector surface. Units are percentages of radiation length.

Detector	Pixel		Drift		Strip	
	Inner	Outer	Inner	Outer	Inner	Outer
Layer	1.14	1.14	1.13	1.26	0.83	0.86
Thermal shield/Support		0.52	0.25		0.53	
Total	7.18 (7.26 including Air)					


of radius. In figure 3.3 the total material encountered by a perpendicular track crossing the ITS is shown versus the azimuthal angle. The inevitable variations due to the structure of the ladders and their overlaps remain within an acceptable range.

As discussed in [20], the relative momentum resolution achievable with the ITS is better than 2% for pions with transverse momentum between 100 MeV/ $c$  and 3 GeV/ $c$ .

The granularity of the detectors was optimized to cope with a track density of 8000 tracks per unit of rapidity at mid-rapidity, the upper limit of theoretical predictions. Under these conditions, the ITS would detect simultaneously more than 15 000 tracks. Keeping the system occupancy low, at the level of a few per cent, requires several million effective cells in each layer of the ITS. As shown in table 3.2, the occupancy is roughly constant in the different ITS layers. The ITS detectors have a spatial resolution of the order of a few tens of  $\mu\text{m}$ , with the best precision (12  $\mu\text{m}$ ) for the detectors closest to the primary vertex. They thus provide a resolution on the impact-parameter measurement adequate for heavy-flavoured particle detection (better than 60  $\mu\text{m}$  in the  $r\phi$  plane for  $p_t > 1$  GeV/ $c$ ). In addition, for momenta larger than 3 GeV/ $c$  the spatial precision of the ITS is an essential element of the momentum resolution.



**Figure 3.2:** Integral of material thickness traversed by a perpendicular track originating at the primary vertex versus radius.



**Figure 3.3:** Material thickness traversed by a perpendicular track originating at the primary vertex versus azimuthal angle.

For the ITS, the total dose expected during the lifetime of the experiment (see section 2.2) varies from a few krad (tens of Gy) for the outer parts of the ITS to about 270 krad (2.7 kGy) for the inner parts [20]. Although these levels are not very severe, compared, for instance, with the levels in ATLAS and CMS, all components used in the ITS design were tested for their radiation hardness to levels exceeding significantly the expected doses, and full radiation tolerant technologies have been used throughout the system for critical components, such as the front end electronics.


### 3.1.1 Silicon Pixel Detector (SPD)

#### 3.1.1.1 SPD system overview

The Silicon Pixel Detector (SPD) constitutes the two innermost layers of the ITS. It is a fundamental element for the determination of the position of the primary vertex as well as for the measurement of the impact parameter of secondary tracks originating from the weak decays of strange, charm, and beauty particles [20]. The SPD will operate in a region where the track density could be as high as 50 tracks/cm<sup>2</sup>, and in relatively high radiation levels: in the case of the inner layer, the integrated levels (10 years, standard running scenario) of total dose and fluence are estimated to be  $\approx 2.7$  kGy and  $\approx 3.5 \times 10^{12}$  n/cm<sup>2</sup> (1 MeV neutron equivalent), respectively [31]. The SPD design implements several specific solutions to minimize the material budget. The average material traversed by a straight track perpendicular to the detector surface is  $\approx 1\%$   $X_0$  per layer.

The SPD is based on hybrid silicon pixels, consisting of a two-dimensional matrix (sensor ladder) of reverse-biased silicon detector diodes bump-bonded to readout chips. Each diode is connected through a conductive solder bump to a contact on the readout chip corresponding to the input of an electronics readout cell. The readout is binary: in each cell, a threshold is applied to the pre-amplified and shaped signal and the digital output level changes when the signal is above a set threshold. This technique had already been successfully applied in the WA97 and NA57 experiments at CERN [32]. The basic detector module is the half-stave. Each half-stave consists of two ladders, one Multi-Chip Module (MCM) and one high density aluminium/polyimide multi-layer interconnect. The ladder consists of a silicon sensor matrix bump bonded to 5 front-end chips. The sensor matrix includes  $256 \times 160$  cells measuring  $50 \mu\text{m}$  ( $r\phi$ ) by  $425 \mu\text{m}$  ( $z$ ). Longer sensor cells are used in the boundary region to ensure coverage between readout chips. The sensor matrix has an active area of  $12.8 \text{ mm}$  ( $r\phi$ )  $\times$   $70.7 \text{ mm}$  ( $z$ ). The front-end chip reads out a sub-matrix of  $256$  ( $r\phi$ )  $\times$   $32$  ( $z$ ) detector cells. The thickness of the sensor is  $200 \mu\text{m}$ , the smallest that can be achieved with an affordable yield in standard processes. The thickness of the readout chip is  $150 \mu\text{m}$ ; the readout wafers are thinned after bump deposition, before bump bonding. The two ladders are attached and wire bonded to the high density aluminium/polyimide interconnect (pixel bus). A  $200 \mu\text{m}$  clearance between the short edges allows for dicing tolerances and ease of assembly. The pixel bus carries data/control bus lines and power/ground planes. The multi-chip-module, wire bonded to the pixel bus and located at the end of the half-stave, controls the front-end electronics and is connected to the off-detector readout system via optical fibre links.

Two half-staves are attached head-to-head along the  $z$  direction to a carbon-fibre support sector, with the MCMs at the two ends, to form a stave. Each sector supports six staves: two on the




**Figure 3.4:** Carbon-fibre support of the Si-pixel staves.

inner layer and four on the outer layer. Ten sectors are then mounted together around the beam pipe to close the full barrel. In total, the SPD (60 staves) includes 240 ladders with 1200 chips for a total of  $9.8 \times 10^6$  cells. The inner (outer) SPD layer is located at an average distance of 3.9 cm (7.6 cm) from the beam axis; the clearances between the detector modules are very small (figure 3.4). The power dissipated in the front-end electronics is  $\approx 1.35$  kW. The cooling system is of the evaporative type and is based on  $C_4F_{10}$ . The sectors are equipped with cooling capillaries embedded in the sector support and running underneath the staves (one per stave). The heat transfer from the front-end chips is assured with high thermal conductivity grease. The SPD barrel is surrounded by an Al-coated carbon-fibre external shield to prevent radiation of heat towards the SDD layers.


Figure 3.5 shows one half-barrel assembled and ready to be integrated in the pixel mechanics. Figure 3.6 shows the SPD installed around the beryllium beam pipe in the experiment.

### 3.1.1.2 Pixel chip and sensor, ladder, bus and front-end electronics

The ALICE pixel readout chip is a mixed signal ASIC developed in an IBM  $0.25 \mu\text{m}$  CMOS process (6 metal layers) with radiation-tolerant layout design [33]. Each chip contains 8192 readout cells of  $50 \mu\text{m} \times 425 \mu\text{m}$  arranged in 32 columns and 256 rows. The size of the chip is  $13.5 \text{ mm} \times 15.8 \text{ mm}$  including internal DACs, JTAG controller, chip controls and wire bonding pads. The chip clock frequency is 10 MHz. A detailed description of the chip architecture can be found in [34, 35]. Each pixel cell contains a preamplifier-shaper with leakage current compensation, followed by a



**Figure 3.5:** Half barrel assembled on reference table.



**Figure 3.6:** The SPD installed around the beryllium beam pipe.

discriminator. A signal above threshold results in a logical 1 which is propagated through a delay line during the L1 trigger latency ( $\approx 6 \mu\text{s}$ ). A four-hit-deep multi-event buffer in each cell allows derandomization of the event arrival times. Upon arrival of the L1 trigger, the logical level present at the end of the delay line is stored in the first available buffer location. The power consumption per chip is  $\approx 990 \text{ mW}$ . Electrical tests indicated a mean minimum threshold of  $\approx 1000$  electrons and a spread of  $\approx 200$  electrons r.m.s. [36].

**Table 3.4:** Main specifications of the ALICE SPD front-end chip.

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

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


Each pixel chip generates a pulse (Fast-OR) whenever a pixel cell (or a group of them) detects a particle signal above threshold. The Fast-OR is used to self-trigger the front-end in test mode and to implement a prompt physics trigger as described in section 3.1.1.9.

The pixel chip has proven to be insensitive to a total ionization dose (TID) exceeding 10 Mrad. The measured cross-section for Single-Event Effects (SEU) is  $\approx 3 \times 10^{-16} \text{ cm}^2$ , corresponding to an upset rate of 0.1 bit/hour in the full barrel for Pb-Pb collisions at nominal luminosity. The main specifications of the ALICE SPD front-end chip are summarized in table 3.4.


A detailed protocol was developed for testing and classifying the readout chips into three classes. Only class I chips are used for flip-chip bonding. These chips show full electrical functionality, normal current consumption and less than 1% faulty pixels. Class II chips are fully functional, but have more than 1% pixel defects. Class III chips show a serious defect, such as an excess in current consumption or a JTAG malfunction. Each wafer (200 mm diameter) contains 86 ALICE pixel chips. A total of 58 wafers were tested. An average yield of 45 class I chips per wafer was found. The pixel chip wafers are thinned down from  $725 \mu\text{m}$  native thickness to  $150 \mu\text{m}$  thickness to reduce the overall material budget in the SPD. Wafer thinning through backside grinding is done by the bump bonding supplier (VTT Center for Microelectronics, Espoo, Finland) after deposition of the solder bumps, which are  $\approx 20 \mu\text{m}$  in diameter.

The pixel sensors have an active size of  $12.8 \text{ mm} \times 70.7 \text{ mm}$ . They are produced on  $200 \mu\text{m}$  thick, high resistivity n-type silicon wafers supplied by CANBERRA (CANBERRA Semiconductor, B-2250 Olen, Belgium) and ITCirst (ITCirst, 38100 Povo di Trento, Italy). Five readout chips are flip chip bonded to one sensor ( $200 \mu\text{m}$  thick) using Pb-Sn solder bump bonds. The performance of each ladder is evaluated using the on-chip test pulse generator as well as a radioactive source. The Fast-OR functionality is verified. According to the test results the ladders are classified into three classes [37]. Only class I ladders are used for half-stave mounting. They have less than 1% defects per chip and a total leakage current of less than  $2 \mu\text{A}$  at nominal operating voltage of 50 V. A rework procedure was developed by VTT in order to replace chips that show an unaccept-





**Figure 3.7:** Number of pixel defects per chip in class 1 ladders. The arrow corresponds to the 1% cut acceptance.

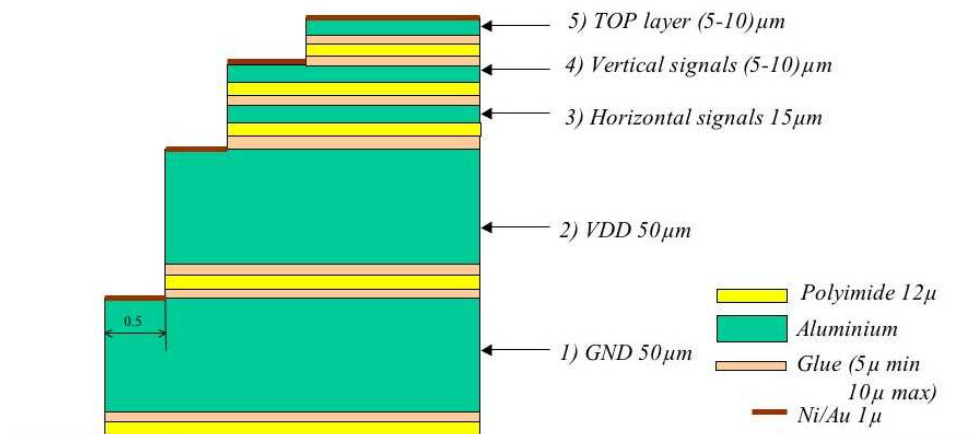


**Figure 3.8:** Block diagram of the electronic system.


able number of defects after bump bonding. A total of 445 ladders were tested, with a class I yield of 77% (including reworked units). Figure 3.7 shows the number of defects per chip of all class I ladders.

A schematic diagram of the SPD electronics system is shown in figure 3.8. The pixel bus (figure 3.9) is a 5-metal layer sequential build-up (SBU) substrate (aluminium/polyimide).

The use of aluminium in place of copper is dictated by the low-mass requirements; it is not an industrial standard and has required a custom development (CERN TS-DEM Workshop) [38]. Two  $50 \mu\text{m}$  thick layers are used as LV power/ground planes. The other three layers carry data/control lines; two of these layers are made by vacuum deposition and are  $\approx 7 \mu\text{m}$  thick (see figure 3.9).




**Figure 3.9:** Pixel bus layers structure.




**Figure 3.10:** Wire bonding of ladders to pixel bus.

The edges of adjacent layers are step-wise receded to make room for wire bonding pads and to allow connections to the power and ground planes without using vias. The overall thickness of the pixel bus is  $\approx 280 \mu\text{m}$ . A picture of the wire bonding of ladders to the bus is shown in figure 3.10. The MCM (figure 3.11) performs the readout of a half-stave. It includes the digital PILOT, the analogue PILOT and the gigabit optical link driver GOL, all developed in the same IBM  $0.25 \mu\text{m}$  CMOS process used for the pixel chip, see [39–41].

The digital PILOT receives serial trigger, configuration data and clock from the off-detector electronics via two PIN diodes in the optical package and the receiver chip RX40 (see figure 3.8 on the preceding page). The digital PILOT initiates the pixel chip read-out, performs data multiplexing



**Figure 3.11:** Multi-Chip Module (MCM). Left to right: wire bonds connecting the MCM ASICs via the pixel bus to the read-out chips, MCM ASICs, optical package with three optical fibres.



**Figure 3.12:** Half-stave assembly.

and sends the data to the G-link compatible 800 Mbit/s serializer GOL, which drives the laser in the optical package. The analogue PILOT contains DACs that generate precision reference levels for the pixel chips and an ADC that reads a Pt1000 temperature sensor chain on the half-stave.

The MCM is based on a 5-metal-layer sequential build-up substrate (polyimide-copper). The footprint is 110 mm×12 mm. The thickness is less than 1.5 mm; this was achieved by mounting bare die ASICs and by the development of a custom optical package that is 1.2 mm thick (STMicroelectronics, Milano, Italy). As the MCM processes and transmits digital data at 800 Mbit/s and also provides reference voltages to the pixel read-out chips with a precision of 10 mV, the layout has required special care. The MCM data/signal lines are wired bonded to the pixel bus. Power is supplied to the pixel bus and the MCM via two independent flexible copper/polyimide laminates (power extenders).

### 3.1.1.3 Half-stave assembly

The half-stave assembly requires the attachment and wire bonding of the two ladders and the multi-chip module onto the high density aluminium-polyimide multi-layer interconnect (pixel bus) (figure 3.12). A grounding foil, consisting of an aluminium-polyimide laminate (25  $\mu\text{m}$  and 50  $\mu\text{m}$  thick, respectively), is included to provide electrical isolation with respect to the carbon fibre. The grounding foil has openings for the thermal coupling of the chip substrates to the cooling ducts via thermal grease pads. The cooling ducts run in grooves embedded in the carbon-fibre support.

The micrometric alignment of the half-stave components is performed with a Mitutoyo coordinate measuring machine (Crysta Apex 9166) equipped with custom developed tools and jigs

for vacuum-suction handling of the components. The glue is the two-component epoxy adhesive Eccobond 45. It is electrically insulating and has adequate thermal conductivity. It can be cured at room temperature; the flexibility can be controlled by the amount of Catalyst 15 used. In our case a mixture of glue and catalyst in the 1 : 1 ratio was used.

The half-stave assembly is a three step procedure based on micrometric components alignment and gluing. The tight dimensional tolerances of the SPD require that the thickness and the planarity of the half-stave be kept under control very carefully during every production step. A few  $\mu\text{m}$  precision in the alignment of components and planarity better than  $100\ \mu\text{m}$  were obtained.


In the first step the two ladders and the MCM are aligned with respect to one another and are glued to the grounding foil. In the second step the pixel bus is aligned and glued on top of them. The third and final step is the ultrasonic wire-bonding connection of the pixel bus to the front-end chips and the MCM. There are more than 1,100 bonds on each half-stave. The quality of the component alignment and gluing is crucial for the reliability and robustness of the wire-bonds. Details on the half-stave assembly procedure can be found in [42–44].

After assembly, each half-stave is extensively tested at nominal working temperature (Peltier coolers). The protocol includes functionality and performance tests such as minimum threshold, threshold uniformity, noisy pixels, dead pixels and Fast-OR. A thorough test is essential since the replacement of a malfunctioning half-stave on a sector is difficult and risky, as it could result in damage to a neighbouring half-stave. The half-staves are then shipped to the facility where the mounting on sector is performed. They are tested again on reception to check for any damage due to transport and handling.

#### 3.1.1.4 Sector assembly

In view of the stringent requirements on the relative positioning of the components, material budget and mechanical handling, specific tools and techniques were developed for mounting the half-staves on sectors. A detailed description of the SPD barrel sector assembly procedure can be found in [45]. The main steps of the procedure can be summarized as follows:

- the half-staves are tested on reception at the clean room facility where the system for mounting them onto the  $200\ \mu\text{m}$  thick Carbon-Fibre Support Sectors (CFSS, see next section) is installed;
- the planarity of the surface on which the half-staves are to be attached is measured for compliance with specifications ( $100\ \mu\text{m}$ ). A thin layer ( $\approx 150\ \mu\text{m}$ ) of thermal grease is dispensed in rectangular pads, each pad corresponds to the back side of a pixel chip. The choice of the thermal grease (AOS 52029, AOS Thermal Compounds, NJ, USA) was made based on its mechanical, thermal and radiation hardness characteristics;
- two half-staves are then aligned to form a stave and positioned on the CFSS using a vacuum holding tool. The average uncertainty in the relative positioning of the two half-staves was found to be better than  $20\ \mu\text{m}$  with respect to the nominal position;



**Figure 3.13:** In the upper plane of the CFSS, thermal grease pads on the cooling duct and UV dots are visible. Such material is dispensed before the assembly of the half-staves. In the lower CFSS planes, three already assembled half-staves are shown. The half-staves are attached to the CFSS with carbon fibre clips and UV glue dots.

- the aligned stave is finally glued to the CFSS using a UV curable glue (Norland NEA 123, Norland Products Inc., NJ, USA). The upper part of the stave is fixed using carbon fibre clips, to protect the wire-bonding and to ensure good mechanical stability, as shown in figure 3.13.

The assembly of the whole SPD barrel was completed with a yield of 100% (all accepted half-staves were assembled on the CFSS without any loss or damage). For two half-staves, additional thermal grease had to be injected during the test phase in order to improve the thermal contact.

### 3.1.1.5 Carbon-Fibre Support Sector (CFSS)

The CFSS is the mechanical support structure of the barrel with an embedded cooling system. The design provides a direct thermal contact between the cooling duct and the silicon chip backplane. This allows to maximise the thermal coupling between the stave and the cooling system and to avoid problems of anisotropic heat conductivity typical of carbon fibre structures. The shape of the groove for the cooling duct together with the shape of the CFSS also provide a very high local stiffness. The global and local deformations of the CFSS are expected to be of the order of  $1\ \mu\text{m}$  or less in the service conditions. The CFSS are made by winding two layers of unidirectional high-modulus,  $100\ \mu\text{m}$ -thick carbon fibre tapes, with fibres parallel and perpendicular to the beam axis, around a metallic mandrel. In order to obtain a good surface finishing quality metallic counter-mandrels are also used. The wall thickness of the CFSS is about  $200\ \mu\text{m}$  in the sensitive region increasing to  $600\ \mu\text{m}$  at the ends in order to allow the positioning of the mechanical references and of the components needed for the global assembly of the sectors. The global thickness of


the support end-sections is locally increased to 1.8 mm in the connections area, in order to allow for an adequate pin length. All the coupling pins and pin seats are made in carbon fibre in order to minimise the use of high-Z material near the sensitive volume. The CFSS have a parilene deposition (about 10  $\mu\text{m}$  thick) in order to increase the DC electrical insulation and to also avoid the accidental release of debris from the CFSS surface. The CFSS qualification is performed by means of a 3D survey machine, with software that allows direct comparison with the 3D CAD model of the components. The planarity, the angular distortion and the difference between the nominal and the measured dimensions are precisely determined. Since the half-stave structure is a combination of materials with a non-negligible mismatch of thermal expansion coefficients, the global  $\Delta T$  from the assembly reference temperature was minimized and the half-staves are firmly attached to the stiff carbon-fibre structure by means of UV glue tags and carbon-fibre clips. Possible deformations do not exceed 10  $\mu\text{m}$  (in any direction) with stresses well within acceptable limits. Further details can be found in [46].

### 3.1.1.6 Cooling

The major contribution to the on-detector power dissipation is due to front-end chips; they generate a heat load of  $\approx 23$  W (nominal) in each stave. The design of the cooling system was driven by several constraints such as low material budget, long-term stability against corrosion, chemical compatibility, minimal temperature gradients, cooling duct temperature above the dew point. Several possible solutions based on different coolants were considered [47]. The final choice of an evaporative system with  $\text{C}_4\text{F}_{10}$  as coolant was found to fulfil all the requirements. The  $\text{C}_4\text{F}_{10}$  follows a Joule-Thomson cycle (rapid expansion at constant enthalpy and subsequent evaporation). The liquid, overcooled and compressed by a pump, is brought to the coexistence phase inside the cooling duct by a pressure drop inside the capillaries (0.5 mm internal diameter, 550 mm long). Heat abduction through phase transition takes place inside the cooling tube at 15–18°C (1.9–2.0 bar); a compressor raises then the pressure pushing the gas towards a condenser, where the liquid phase is re-established by heat transfer to cold water ( $\approx 6^\circ\text{C}$ ). The evaporation temperature can be controlled by regulating the pressure in the return line, setting then the coexistence conditions of the mixed phase. Each stave is put in thermal contact with the cooling duct mounted in a groove on the CFSS by a thermal grease layer (see section 3.1.1.5).

The cooling duct is obtained using Phynox tubes with a wall thickness of 40  $\mu\text{m}$  and an initial diameter of 2.6 mm, squeezed down to flat profile with an overall thickness of 600  $\mu\text{m}$  in the thin dimension. Each sector is equipped with cooling collectors at the two ends, one functioning as an inlet and the other as an outlet for the whole sector. Extensive corrosion tests were performed on tubes, together with tests to optimize the choice of surface treatment and of fitting materials.

In the final test, detectors are turned on at nominal power. An infrared (IR) camera is used to measure the temperature distribution on the external part of the pixel matrix (i.e. the bus), on the MCM and on the power extenders (copper-polyimide laminates). Residual gaps due to mechanical tolerances of the surfaces in contact can locally affect the quality of thermal coupling. The average operating temperature of the half-staves is in the range of 25 °C to 30 °C. The temperature profile measured on a half-stave with the IR camera is shown in figure 3.14. The temperature does not exceed 35 °C and the local variations are limited to a few °C.



**Figure 3.14:** Left: infrared image of an SPD half-stave during operation. Right: temperature profiles - transversal across the ladders (top two plots) and longitudinal along the full half-stave (bottom plot).


In normal operation, if a sudden failure of the cooling were to occur, the temperature at the half-stave would increase at a rate of  $1\text{ }^{\circ}\text{C/s}$ . Continuous monitoring and a fast, reliable safety interlock on each half-stave are therefore mandatory. They are based on Pt1000 temperature transducers mounted on the pixel bus, next to the pixel chips. Two daisy chains of 5 transducers each (interleaved positions) provide redundant measurements of the average temperature. One chain is read out in the MCM, the other is hard-wired to the remote interlock system, based on a programmable logic controller (PLC) that is part of the detector control and safety system. Temperature values are logged. If the temperature reaches a preset threshold, the low-voltage power supply is promptly switched off by the safety interlock and an alarm is generated.

### 3.1.1.7 Power supply and grounding

The low-voltage power supply system is based on CAEN A3009, (CAEN, Viareggio, Italy) dc-dc converter modules, installed in racks in the experimental area at a distance of  $\approx 40\text{ m}$  from the detector. Each half-sector (6 half-staves) is powered by one module, which has 12 independent floating outputs; 2 outputs are used for each half-stave, one for the pixel chips ( $1.85\text{V}/0.5\text{A}$ ), the other one for the MCM ( $2.6\text{V}/0.5\text{A}$ ). Remote sensing is used throughout. In each half-stave, the pixel chip/MCM supply return lines are shorted and define the half-stave ground. In the CAEN A3009, all return lines are connected via  $10\text{ k}\Omega$  resistors to a power supply reference ground that is connected to the ALICE ground (on the absorber and the space frame).

The detector bias voltage ( $50\text{ V}$  typically at start of detector operation) is provided by CAEN A1519 modules located in the control room at a distance of  $\approx 120\text{ m}$  from the detector.

One module (12 channels) is used for each sector. The two sensor ladders in one half-stave share one high voltage module output but are connected by one coaxial cable each to the HV-



**Figure 3.15:** Power supply and grounding scheme.

module in the control room. This allows the individual connection of a sensor ladder to the bias voltage. The return line of the high voltage is connected to the half-stave ground via a  $100\text{ k}\Omega$  resistor. The half-stave ground is isolated from the carbon-fibre support by means of the grounding foil. The carbon-fibre support itself is connected to the ALICE ground.

Figure 3.15 illustrates the power supply and grounding scheme.

### 3.1.1.8 Readout and Calibration

After the reception of the L1 trigger signal, pixel hit data are extracted from the digital delay line and stored in one of the four multi-event buffers in the pixel chip. Upon arrival of the second level trigger (L2), the data contained in the multi-event buffer locations corresponding to the first (oldest) L1 trigger are loaded onto the output shift registers. Then, for each chip, the data from the 256 rows of cells are shifted out during 256 cycles of a 10 MHz clock. At each cycle, a 32-bit word containing the hit pattern from one chip row is output on the 32-bit data bus where it is processed



by the MCM and sent optically to the read-out electronics located in the control room. One pixel chip is read out in  $25.6 \mu\text{s}$ . The 10 chips on each half-stave are read out sequentially. The 120 half-staves are read out in parallel. In ion beams operation, the dead time introduced by the readout of the SPD is estimated to be less than 10% in the worst case, corresponding to ALICE running with Ar-Ar beams at high luminosity, with an L1 rate of 2.5 kHz.


The SPD off-detector read-out electronics is located in the control room [41]. Twenty 9U VME- and FPGA-based electronics cards (routers) contain each 3 plug-in type daughter cards (link receiver cards). The link receiver cards are connected via optical fibres to two half-staves. The link receivers zero-suppress and re-format the hit data to an ALICE off-line analysis compatible format. The routers multiplex data from the six half-staves into one ALICE optical Detector Data Link (DDL) and attach trigger and status information. Each of the 20 router cards provides the interface to the trigger, DAQ and DCS/ECS for 6 half-staves. The trigger information is delivered by the Trigger, Timing and Control (TTC) system via optical fibres. Each router is interfaced to the DAQ via one DDL. The 20 DDLs are connected to Local Data Collectors (LDCs) housed in 4 PCs. The data access from the DCS/ECS to the routers is established via the router VME ports. The same port also allows monitoring and copying the data flow during data taking.

Detector calibration is performed using the on-chip test pulse generator [34, 35]. The test pulse amplitude is controlled by the analogue PILOT chip and is programmable. The pulse can be sent to individual pixel cells or to clusters of them. The SPD DCS is required for executing the calibration procedures [48].

### 3.1.1.9 Fast-OR pixel trigger

Each pixel chip provides a Fast-OR digital pulse when one or more of the pixels in the matrix are hit. The Fast-OR signals of the 10 chips on each of the 120 half-staves are read by the PILOT chip ([39–41]) and transmitted every 100 ns on the 120 optical links that are also used for the data readout. The Fast-OR allows the implementation of a unique prompt trigger capability. The pre-processed Fast-OR data can be used to contribute to the Level 0 trigger decision in the ALICE Central Trigger Processor (CTP). This feature is very useful in particular in the case of events with very low multiplicities in pp runs. Various trigger algorithms including the Fast-OR are being investigated [49]. These studies have shown that background rejection in pp interactions and event selection in heavy-ions runs is significantly improved by using the Fast-OR information in the L0-trigger decision.

The pixel trigger signal generated by the Fast-OR processor must reach the CTP within  $\approx 800$  ns of the interaction in order to meet the latency requirements of the L0-trigger electronics [50]. Short connections and minimal processing overhead are therefore essential. However, some constraints deriving from the pixel chip design have to be taken into account. The Fast-OR output is synchronized with the 10 MHz pixel system clock, hence the signal is effectively integrated over 100 ns corresponding to 1 bunch-crossing in the case of operation with heavy ions, or 4 consecutive bunch-crossings in pp operation. All 120 half-staves are synchronized to the same integration period covering the same set of bunch-crossings in order to limit this ambiguity when all Fast-OR signals are simultaneously processed and to minimize the latency. In the ALICE trigger system, the bunch-crossing ambiguity can be resolved by considering the coincidence between the pixel trigger signal and the ALICE V0 detector signal.



**Figure 3.16:** Pixel trigger system architecture.

The pixel trigger system architecture is shown in figure 3.16. A full description of the pixel trigger system is given in [51]. The development of the system was challenging due to the low latency and extreme compactness required. Two-way optical splitters equipping the optical fibres outgoing from the detector allow the simultaneous transmission of the readout and trigger data to the readout system and to the pixel trigger system respectively. The pixel trigger system consists of a set of ten receiver boards that deserialize the optical data streams and extract the 1200 Fast-OR bits received every 100 ns from the detector. Each receiver board includes a customized 12 channels optical parallel receiver module [52], 12 G-Link compatible deserializer ASICs and a FPGA to tap the Fast-OR signals from the twelve 800 Mb/s input channels. The receiver boards are all connected as mezzanines onto a motherboard and provide the 1200 Fast-OR bits as simultaneous inputs to the processing unit on the motherboard, a fast FPGA with a large number of pins and large logic space. Various algorithms, based on global multiplicity or on predefined topologies, can be implemented as Boolean logic functions of the Fast-OR bits. The algorithm to process the Fast-OR bits is implemented in programmable hardware to allow fast execution, upgrading and remote reconfiguration. The result of the pixel trigger algorithm implemented on the FPGA is directly fed to the ALICE CTP.

The implementation on a large FPGA of some of the proposed trigger algorithms was completed. The behavioural simulations showed that even the most complex function among the proposed ones could be evaluated in less than 15 ns in a Xilinx Virtex 4 device. The critical delays in the system are associated with the Fast-OR data deserialization, extraction and transfer between the peripheral FPGAs and the processing unit. The overall latency budget of 800 ns can be subdivided among each of the processes along the data flow. The on-detector electronics takes 400 ns from the time of the collision to transmit the Fast-OR bits. The shortest path for the optical fibres from the detector to the location of the pixel trigger system is  $\approx 37$  m. This implies a signal propagation delay of  $\approx 185$  ns. Therefore only  $\approx 215$  ns are left for the deserialization, extraction and processing phases. The system is fully developed, and its compliance with these constraints was verified.

### 3.1.1.10 Tests of the sector, half-barrel and Fast-OR with cosmics

A dedicated test and integration facility was set up at CERN within the Departmental Silicon Facility (DSF) clean room area (class 100 000). The facility includes the complete readout chain and trigger/DAQ system, the final LV and HV power supplies and cables, the cooling plant and the detector control system and safety interlock to allow testing one complete half-barrel. The facility was used to develop and test major parts of the system. Each sector is tested individually before its integration into a half-barrel. The complete test sequence is repeated for each of the two half-barrels.

A thermal survey of each sector was performed using a thermal imaging camera together with the measurement of the Pt1000 transducer chains mounted on each half-stave. The average temperature on the half-stave in normal operating conditions was  $\approx 28^\circ\text{C}$ .

On each sector, all the half-staves were tested following a well defined procedure for a full functionality check. Power supply and leakage currents were measured. The on-chip test pulse was used to determine the minimum threshold, the homogeneity of the matrix response and the number of noisy pixels. No degradation of the minimum threshold compared to the half-stave tests was found. The total number of noisy pixels in all 10 sectors is  $\approx 50$ . Sector tests were also carried out using a  $^{90}\text{Sr}$  radioactive source to identify non-working pixels. The bump bonding yield was in agreement with the corresponding measurements on half-staves.


On a subset of half-staves mounted on a sector the Fast-OR trigger setting was adjusted to carry out dedicated runs with cosmic rays. For this purpose the sector was oriented for maximum vertical acceptance. A similar test was later carried out on part of one half-barrel. In a 6 hour continuous run, in which the trigger was based on the coincidence of the Fast-OR of the inner and outer layer, approximately 7000 events were collected. The observed distribution of clusters in the two layers and the clusters correlation along the sector axis were compatible with the expectations, considering that the limited test time available did not allow the optimization of uniformity response and Fast-OR DACs settings required for maximum efficiency.

### 3.1.1.11 Results from beam tests

SPD prototype assemblies were tested in proton- and pion- secondary-beams at the SPS [53, 54] in the years 2002 and 2003. The experimental setup consisted of a movable assembly under test placed between two fixed doublets of pixel planes of the same type. The doublets were used as the reference and tracking system. Measurements were done at different threshold settings and orientations of the test device with respect to the beam line. The configuration of the tracking doublets, together with a detailed cluster analysis of the hits, allowed a study of the intrinsic spatial precision of the planes under test and their detection efficiency [55, 56].

At normal track incidence and standard threshold (200 DAC value, corresponding to about 3000 electrons [57]), mainly single and double pixel clusters occurred. After subtracting the uncertainty on the telescope prediction, the global detector precision (mixing all cluster patterns) along the short pixel side ( $y$ -axis) was  $\sigma_{\text{pixel}}(y) = (11.1 \pm 0.2) \mu\text{m}$ . In a wide plateau around this threshold setting the detector efficiency was above 99%.

Data runs were taken in 2002 with the detector under test tilted by 0, 5, 10, 15, 20, 25 and 30 degrees and with thresholds set at 185, 200 and 210 DAC units (increasing DAC values correspond



**Figure 3.17:** Test-beam results on the intrinsic precision in the  $y$  coordinate as a function of the track incidence angle on the detector: at different threshold settings (left) and at 200 DAC units threshold for the main cluster patterns separately (right).


to decreasing effective threshold). Results of the study on the spatial precision as a function of the track incidence angles are shown in figure 3.17. The plot on the left side shows that the highest spatial precision is obtained at an incidence angle in the range  $5\text{--}10^\circ$ ; for each threshold this occurs at the angle for which the fractions of single and double pixel clusters are roughly equal. The precision also degrades for all the threshold settings with the increasing track angle, as expected when higher cluster topologies are involved. The plot on the right side shows the intrinsic precision, together with the contributions due to the main pixel cluster patterns, in the case of 200 DAC threshold setting. These results are particularly useful in the tuning of the tracking errors to be associated to the SPD clusters, in simulations as well as in the analysis of experimental data.

### 3.1.2 Silicon Drift Detector (SDD)

The Silicon Drift Detectors (SDD) [11] equip the two intermediate layers of the ITS, where the charged particle density is expected to reach up to  $7\text{ cm}^{-2}$ . They have very good multitrack capability and provide two out of the four  $dE/dx$  samples needed for the ITS particle identification.

#### 3.1.2.1 Sensor layout

The ALICE SDDs were produced from very homogeneous high-resistivity ( $3\text{ k}\Omega\text{cm}$ )  $300\text{ }\mu\text{m}$  thick Neutron Transmutation Doped (NTD) silicon [58]. As shown in figure 3.18, they have a sensitive area of  $70.17(r\phi)\times 75.26(z)\text{ mm}^2$  and a total area of  $72.50\times 87.59\text{ mm}^2$ . The sensitive area is split into two drift regions by the central cathode strip to which a HV bias of  $-2.4\text{ kV}$  is applied. In each drift region, and on both detector surfaces, 291  $p^+$  cathode strips, with  $120\text{ }\mu\text{m}$  pitch, fully deplete the detector volume and generate a drift field parallel to the wafer surface. To keep the biasing of the collection region independent on the drift voltage, a second bias supply of  $-40\text{ V}$  is added. The overall detector performance, when averaged over its entire area, does not depend significantly on the applied bias voltage in a fairly large range from  $-1.65\text{ kV}$  to  $-2.4\text{ kV}$ , so the precise value of bias voltage to be applied during long term operation is to be adapted to the specific running conditions. The degrading of the high voltage to the zero potential of the detector boundary



**Figure 3.18:** Layout of the ALICE SDD. The sensitive area is split into two drift regions by the central, highest voltage, cathode. Each drift region has one row of 256 collection anodes and three rows of 33 point-like MOS charge injectors for monitoring the drift velocity. Drift and guard regions have independent built-in voltage dividers.

is implemented by two insensitive guard regions biased by 145 cathode strips with  $32 \mu\text{m}$  pitch. To improve the detector reliability, all the drift and guard regions have their own built-in voltage dividers. Their total power dissipation is 1 W per detector and is removed by an appropriate air circulation system. Each drift region has 256 collection anodes with  $294 \mu\text{m}$  pitch and three rows of 33 MOS charge injectors ( $20 \times 100 \mu\text{m}^2$  each) to monitor the drift velocity which depends on temperature:  $v_{\text{drift}} \propto T^{-2.4}$ , [59], which gives a  $0.8\%/K$  variation at room temperature. They will be triggered at regular intervals during the gaps in the LHC orbit.

At the bias voltage of  $-2.4(-1.65)$  kV the drift velocity is  $8.1(5.6) \mu\text{m}/\text{ns}$ . Since the front-end electronics samples the signal of each anode at a frequency of 40.08 MHz, the size of the sensitive element (cell) is  $294 \times 202(130) \mu\text{m}^2$ , corresponding to  $89.1(138.4) \times 10^3$  cells per detector, which are readout by 512 electronic channels (256 at each end).

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

The double-track resolution was calculated using the detailed simulation program, included in AliRoot, see section 6.4 after its preliminary tuning based on the single-track results from beam tests [61]. Due to charge diffusion during the drift process, the double-track resolution is a function of the drift time for a given separation efficiency. The relative distance at which two clusters are disentangled with a 70% efficiency grows almost linearly from  $600 \mu\text{m}$  near the anodes to  $800 \mu\text{m}$  at the maximum drift distance [62]. The average double-track resolution measured in the beam test for tracks perpendicular to the SDD plane is  $570 \mu\text{m}$  for 70% resolution efficiency [63]. The main parameters of the ALICE SDD are summarised in table 3.5.


**Table 3.5:** The main characteristics of the ALICE silicon drift detectors.

Sensitive area	$70.17 \times 75.26 \text{ mm}^2$
Total area	$72.50 \times 87.59 \text{ mm}^2$
Collection anodes (readout channels)	$2 \times 256$
Anode pitch	$294 \mu\text{m}$
Operating voltage	$-1.65 \text{ to } -2.4 \text{ kV}$
Nominal bias of the collection region	$-40 \text{ V}$
Drift velocity	$5.6 \text{ to } 8.1 \mu\text{m/ns}$
Maximum drift time	$4.3 \text{ to } 6.3 \mu\text{s}$
Cell size at drift velocity $8.1 \mu\text{m/ns}$	$294 \times 202 \mu\text{m}^2$
Cells per detector at drift velocity $8.1 \mu\text{m/ns}$	$2 \times 256 \times 174$
Total number of cells (260 SDDs)	$23 \cdot 10^6$
Average resolution along the drift ( $r\phi$ )	$35 \mu\text{m}$
Average resolution along the anode ( $z$ )	$25 \mu\text{m}$
Detection efficiency	$99.5\%$
Average double-track resolution at 70% efficiency at max. field	$700 \mu\text{m}$


### 3.1.2.2 Front-end electronics and readout

The SDD front-end electronics is based on three types of ASICs. The first two, PASCAL and AMBRA, are assembled on the front-end hybrid, shown in figure 3.19 on its assembly jig. PASCAL contains three functional blocks: preamplifier, analogue storage and Analogue-to-Digital Converter (ADC) [64–67]. AMBRA, which receives data from PASCAL, is a digital four-event buffer which performs data derandomisation, baseline equalization on an anode-by-anode basis and 10-bit to 8-bit non-linear data compression. AMBRA also sends the data to the third ASIC, CARLOS, which is a zero-suppressor and data-compressor mounted on one of the end-ladder boards [68, 69]. The three ASICs have been designed using radiation-tolerant layout technique (enclosed gate geometry) based on a commercial CMOS deep-submicron process ( $0.25 \mu\text{m}$ ). The PASCAL chip designed with this technology has proven to be insensitive to total ionisation dose up to  $300 \text{ kGy}$ . The average power dissipation of each PASCAL-AMBRA front-end channel is about  $6 \text{ mW}$ .

The signal generated by an SDD anode feeds the PASCAL amplifier-shaper which has a peaking time of  $40 \text{ ns}$  and a dynamic range of  $32 \text{ fC}$  (the charge released by an 8-MIP particle hitting near the anode). The amplifier is DC coupled to the sensor and is able to cope with a leakage current of  $500 \text{ nA}$  per anode. The amplifier output is sampled at  $40.08 \text{ MHz}$  by a ring analogue memory with 256 cells per anode. This is the mode of operation in the idle state of the front-end. When an L0 trigger is received, the SDD BUSY is immediately set and, after a programmable delay which accounts for the L0 latency ( $1.2 \mu\text{s}$ ) and the maximum detector drift time ( $\sim 5 \mu\text{s}$ ), the analogue memories are frozen. The BUSY being still set, their contents are then digitised by a set of 10-bit linear successive-approximation ADCs which write the data into one of the free AMBRA buffers. On a run-by-run basis, PASCAL can be programmed to use half of this nominal frequency for the sampling, thus reducing the amount of data and, therefore, the sub-system readout dead-time. Analysis of beam test data and simulation have shown that the cost in terms of both spatial resolution and double-track resolution is negligible. The advantages of an A/D



**Figure 3.19:** Left: a Silicon Drift Detector on the assembly jig with the two front-end hybrids connected, via their test extensions, to test interface circuits. Right: a single hybrid on its assembly jig: the four chips on the top row are the PASCALS, the ones on the bottom are the AMBRAs.

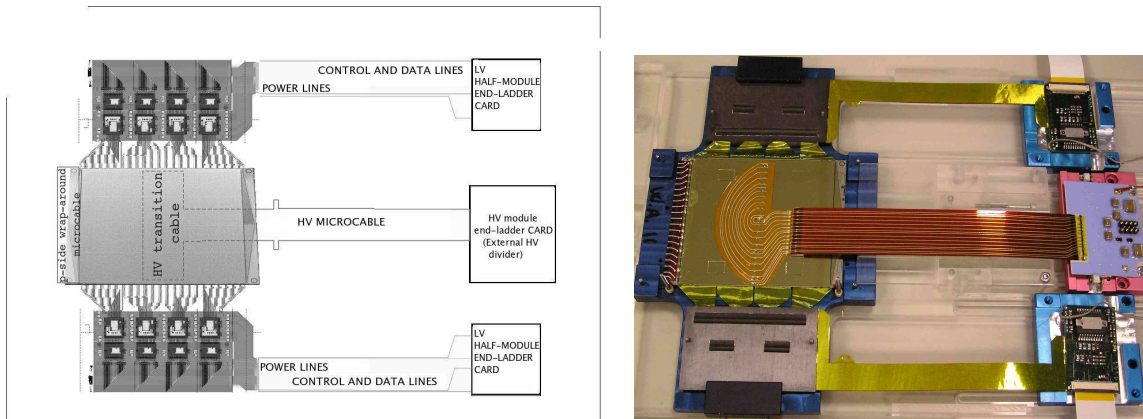


**Figure 3.20:** Left: photograph of a CARLOS end-ladder board, with the GOL and CARLOS chips visible on the left. Right: photograph of a CARLOS-rx board.

conversion on board of the front-end chip are the noise immunity during signal transmission, and the possibility of inserting a multiple-event buffer to derandomise the data and, therefore, to slow down the transfer rate to the DAQ system. This greatly reduces the material budget of the cabling. The digitisation lasts for about  $230 \mu\text{s}$  ( $120 \mu\text{s}$  in half-frequency mode) and can be aborted by the absence of the L1 trigger or by the arrival of an L2-reject signal; in both cases, the front-end electronics resets the SDD BUSY and returns to the idle state within 100 ns. On the successful completion of the analogue-to-digital conversion the SDD BUSY is reset if at least one buffer is still available in the AMBRAs. As soon as the conversion is completed, all the AMBRAs transmit the data in parallel to the CARLOS chips on the end-ladders, an operation which takes 1.23 ms (0.61 ms in half-frequency mode) if 192 samples per anode are read out. By means of a 2-dimensional 2-thresholds algorithm and with no additional dead time, the CARLOS chips reduce the SDD event size from the raw 24.4 MB by more than one order of magnitude. They also format data and feed the GOL ASICs [70] which in turn drive the optical links. The board on which the CARLOS and GOL chips are mounted is shown in figure 3.20 (the chips are visible on the left). In the counting room, 24 VME boards, CARLOS-rx [71], concentrate the data coming from the 260 SDDs into 24 DDL (Detector Data Link) [72] channels and embody the trigger information in the data flow.

**Table 3.6:** Main parameters of the SDD front-end electronics.

Shaping time	40 ns
Sampling frequency	40 MHz or 20 MHz
Noise	350 electrons rms
Gain	50 mV/fC
ADC resolution	10 bit
ADC linearity	better than 1%
Minimum dead-time	129 $\mu$ s
Power consumption	6mW/channel



**Figure 3.21:** Scheme (left) and photograph (right) of a SDD module.

CARLOS-rx, shown in figure 3.20, also uploads in parallel the configuration parameters received over the DDL to the ladder electronics it controls, and monitors the error-flag words embedded in the data flow by the CARLOS chips in order to signal potential Single-Event Upsets (SEU) on the ladder electronics.

To allow the full testability of the readout electronics at the board and system levels, the three ASICs embed a JTAG standard interface. In this way it was possible to test each chip after the various assembly stages. The same interface is used to download control information into the chips before the data taking. A programmable test pulse generator on PASCAL allows for a fast, yet detailed, test of the whole chain for calibration purposes.

The average noise of the front-end electronics is 2 ADC counts rms, which corresponds to an Equivalent Noise Charge (ENC) of 350 electrons. This value measured at the system level is close to that obtained with a single chip (320 electrons) [73].

Table 3.6 summarizes the main parameters of the SDD front-end electronics.

### 3.1.2.3 SDD modules

A SDD module (see figure 3.21) consists of one silicon drift detector and two front-end hybrids, each connected to the corresponding end-ladder LV board. A micro-cable, specially designed to carry high voltage (up to over 2.4 kV), connects the detector to the HV end-ladder board. All the assembly and test steps were performed in-house [74]. A photograph of the final assembly is shown in figure 3.21.



The detector bias voltage is provided by specially designed flexible printed circuit called "micro-cables". The connections to the central bias cathode and to the injector lines is provided by a micro-cable called 'transition cable', glued on the detector p-side. The bias lines are wire bonded to the corresponding bonding pads. The high voltage is then brought to the n-side using the so called 'wrap-around' cable. The transition cable is Tape Automatic Bonded (TAB) to the long HV cable connected to the corresponding HV end-ladder board. Each sensor was carefully tested and after being accepted for module production. It was equipped with transition and wrap-around cables and sent to module assembly. The transition cable and the wrap-around cable are clearly visible in figure 3.21, on the centre of the sensor and on its left-hand side respectively.

After wafer-testing, with a yield of 78% for AMBRA and 89% for PASCAL [73], thinning to 150  $\mu\text{m}$  and dicing, the front end chips were TAB bonded to aluminium multilayer printed circuits called 'chip-cables', one chip-cable per PASCAL-AMBRA pair. This circuit provides connections of data lines between AMBRA and PASCAL as well as the connections to the sub-hybrid circuit. After testing, the bondings on the chip-cables were protected by a thin layer of glue.

After this encapsulation process, chip-cables were glued four by four on a sub-hybrid circuit, in turn glued onto a rigid carbon-fibre heat-dissipater called heat-bridge, which clips to the cooling tubes running along both sides of a ladder to ensure good thermal contact. Each hybrid was then tested before and after a series of thermal cycles from 20 to 65°C and the bonds encapsulated with protective glue.


Testing at the different assembly stages was facilitated by the presence of extensions to the micro-cables interfacing to the test equipment, which were later removed before further assembly. Such test extensions are visible in figure 3.19 on page 40, which shows a sensor being connected to its two front-end hybrids.

Finally, the module was completed by the assembly of the long cables which connect it to the LV and HV boards, both sitting at the end of the ladders. These micro-cables have typical polyimide and aluminium thicknesses varying from 12 to 20  $\mu\text{m}$ , leading to total thicknesses of the multilayer assembly of 220  $\mu\text{m}$  for the digital cable and 150  $\mu\text{m}$  for the analogue one. The high voltage cable is 340  $\mu\text{m}$  thick.

The LV boards ensure the signal interfaces and Low Voltage distribution, carry the radiation-hard low voltage regulators, the LVDS signal-receivers, the interface with the DCS and the ancillary circuitry to drive the MOS charge injectors, while the HV boards carry a support High-Voltage divider. The development of the long micro-cable was particularly delicate since it was designed to minimize material while keeping excellent High-Voltage insulation, signal quality and power dissipation.

The complexity of the assembly and test procedures (over 20 process steps were necessary) led to an average time of 3 days to have the complete module ready for final test with laser.

This test was developed to get a complete characterization of each individual module prior to assembly in the ladders. By mounting the fully assembled module on a moving stage and generating charge at known positions with a laser, each sensor was mapped with over 100 000 points, providing a detailed description of its response. In particular, this method allows for the generation of correction tables compensating for doping fluctuations, which affect a subset of the detectors [75, 76], defects in the integrated voltage divider and noisy or dead channels.



**Figure 3.22:** The SDDs are mounted at different radii in both  $r_z$  and  $r_\phi$  planes to obtain the full coverage in the acceptance region. Units are millimetres.


**Table 3.7:** Main parameters of the ALICE SDD layers and ladders.

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


### 3.1.2.4 SDD ladders

The SDDs modules are mounted on linear structures called ladders. There are 14 ladders with six modules each on layer 3, and 22 ladders with eight modules each on layer 4. Modules and ladders are assembled (figure 3.22) to have an overlap of the sensitive areas larger than  $580 \mu\text{m}$  in both  $r_\phi$  and  $z$  directions. This ensures full angular coverage for vertices located within the interaction diamond,  $\pm\sigma = 10.6 \text{ cm}$ , and for  $p_t > 35 \text{ MeV}/c$ .

The main geometrical parameters of the SDD layers and ladders are summarised in table 3.7. The ladder space frame is a lightweight truss made of Carbon-Fibre Reinforced Plastic (CFRP) and has a protective coating against humidity absorption. The sagging of the space frame under the load of the ladder components applied to the mid point of the truss was evaluated with a Finite Element Analysis model and verified by measurement. When the load acts perpendicularly to the detector plane the sagging is  $21 \mu\text{m}$ ; when it acts parallel to the detector plane the sagging is  $7 \mu\text{m}$ . All ladders were measured individually, and only those with sags smaller than  $40 \mu\text{m}$  were accepted for mounting.




**Figure 3.23:** Left: support structure of the SDD ladders; units are millimetres. Right: ladder foot with its V-shaped nut and the reference ruby sphere which is pressed against the nut by a spring. The ruby sphere glued onto the ladder foot is used to verify the ladder position once the assembly of a layer is completed.



**Figure 3.24:** Left: the ladder assembly jig in operation. Right: cut-view of the ladder cross section showing the front-end micro-cables and the cooling tube.

The ladders were assembled on a CFRP structure made of a cylinder, two cones and four support rings (figure 3.23). The cones provide the links to the outer Silicon Strip Detector barrel and have windows for the passage of the SDD services. The support rings are mechanically fixed to the cones and bear reference ruby spheres for the ladder positioning. The accuracy of the re-positioning of the ladders after a dismount-mount sequence was measured to be better than  $10 \mu\text{m}$  full width.

The detectors are attached to the support structure using ryton pins and have their anode rows parallel to the ladder axis ( $z$ ). The front-end electronics, assembled on two hybrid circuits, one per anode row, is glued with a thermo-conductive compound on rigid carbon fibre heat-exchangers, which in turn are clipped to the cooling pipes running along the ladder structure (figure 3.24). The pipes are made of MP35N alloy and have  $80 \mu\text{m}$  wall thickness and 2 mm inner diameter. The coolant is demineralised water and the system is designed for underpressure operation. The connections between the detectors, the front-end hybrids and the end-ladder boards, and the connection for the detector biasing are all provided with flexible aluminium-polyimide micro-cables that are



**Figure 3.25:** Left: the SDD detector completely assembled, ready to be integrated with the Silicon Strips. Right: the CAD design of the SDD layers, showing the support cones and the ladders of the two layers.

TAB bonded. Before the hybrids were clipped to the cooling pipes, the positions of the detectors were measured with respect to the reference ruby spheres glued to the ladder feet. The whole procedure of assembly of detector modules to the ladder frame was performed using a precision measurement machine and a set of specially-designed tools, which allowed the assembly of one ladder in approximately one day. Figure 3.24 shows a ladder during assembly and the assembly tool. Each ladder, once assembled and measured, was completely tested. The few modules with problems which could not be fixed on the ladder were removed and replaced. Although the MOS injectors have proven to be an excellent monitoring instrument for the drift velocity, the cooling system of the SDD layers has been designed with extreme care to provide a temperature stability of 0.1 K, undergoing not only detailed simulations and prototype tests, but also the construction of a full scale model of the entire system, which allowed the verification and optimization of the proposed solutions [28, 30]. The final system is a combination of two independent underpressure water circuits, one coupled to the front-end electronics via the pipes running along the ladders and the other coupled to the readout electronics, LV and HV boards via pipes embedded in the end-ladder structure. The water cooling is complemented by a moderate air flow preventing the creation of hot spots while providing additional cooling to the sensors [28]. To ensure proper operation and the persistence of the underpressure conditions, each of the 13 SDD cooling circuits is controlled by four pressure regulators and monitored by 4 flow-meters and 6 pressure sensors. Depending on its position, a cooling circuit cools two or four ladders (and the related end-ladders). This kind of grouping has been imposed by the limited available space for the piping.

In figure 3.25 the completely assembled SDD detectors are shown, prior to insertion in the SSD detectors, together with a CAD drawing of the SDD layers, showing ladders from both layers and the support cones.

**Table 3.8:** SSD system parameters (see text).

component layer	layer 5	layer 6
Radius	378 and 384 mm	428 and 434 mm
Sensitive area	2.2 m <sup>2</sup>	2.8 m <sup>2</sup>
Number of ladders	34	38
Modules/ladder	22	25
Total modules	748	950
Total weight	1.24 kN	

### 3.1.3 Silicon Strip Detectors (SSD)

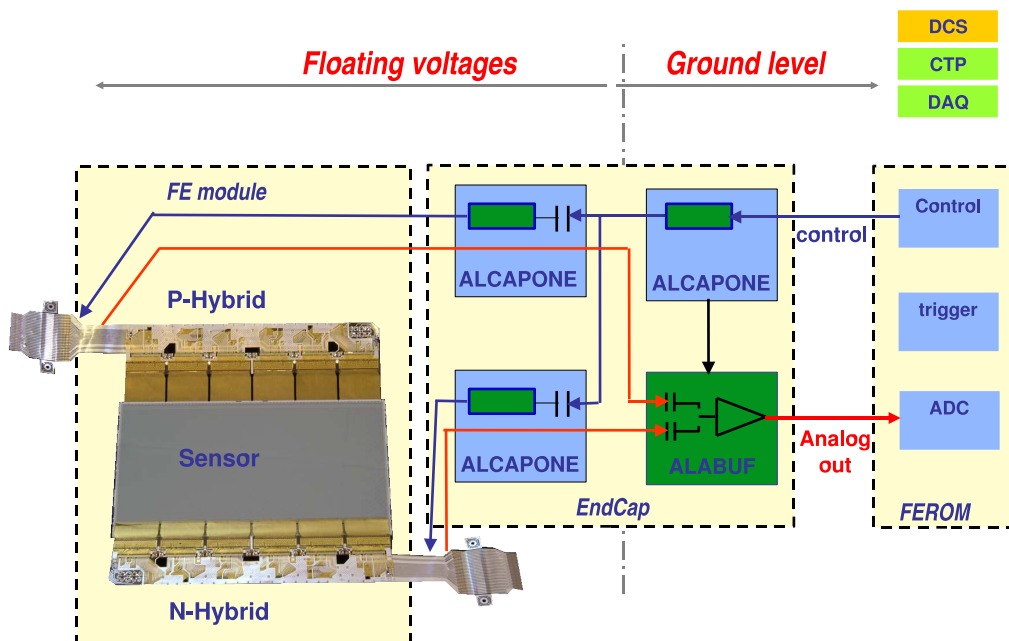
#### 3.1.3.1 Overview

The outer layers of the ITS are crucial for the matching of tracks from the TPC to the ITS. They provide a two dimensional measurement of the track position. In addition they provide  $dE/dx$  information to assist particle identification for low-momentum particles. The system is optimized for low mass in order to minimise multiple scattering.

Both outer layers use double sided Silicon Strip Detectors (SSD) [77]. The detection modules [78] consist of one sensor each, connected to two hybrids with six HAL25 [79] chips each. All interconnections between the sensor and the electronics in the detection module are made using aluminium on polyimide cables (micro-cables) [80]. Minimization of the material budget of the mechanical support for the detection modules is achieved by using linear Carbon Fibre Composite (CFC) material [81] for all support structures in the active volume. The cooling system is also optimized [29, 30] in line with the zero heat balance required for all ALICE detectors. The application of 300  $\mu\text{m}$  thick CFC with high thermal conductivity as miniature motherboards [30] (stiffeners) significantly reduces the material needed for the support and cooling of the front-end chips [30]. The modules are assembled on ladders [81] of the same design as those supporting the SDD. These ladders are one module wide and up to 25 modules long along the beam direction. The 72 ladders, carrying a total of 1698 modules, are mounted on CFC support cones in two concentric cylinders. For each layer the ladders are mounted in two slightly different radii such that full azimuthal coverage is obtained. The modules are cooled by water running through two thin (40  $\mu\text{m}$  wall thickness) phynox tubes along each ladder [30, 82].

The electronic signals from the modules are AC-coupled and buffered in custom made electronics, the EndCap Modules (ECM) [83] located at each end of each ladder. The analogue to digital conversion of the signals performed in the Front-End ReadOut Modules (FEROM) [84] located outside the ALICE magnet, see figure 3.26.

The technology of each component was chosen corresponding to the radiation environment and magnetic field in which they need to operate. The segmentation of the power, cooling and readout system avoids failure of significant parts of the system due to a single point failure. Down to the level of the front-end chips malfunctioning components can be isolated to maintain system integrity. The main geometrical parameters of this system are listed in table 3.8. A detailed description of the SSD system design can be found in [11].



**Figure 3.26:** Schematic overview of the SSD readout system. The detection module is located in the active volume. The endcap is located at the end of the ladder, just outside the active volume and the FEROM is located outside the ALICE magnet (see text).

**Table 3.9:** SSD sensor parameters. The sensors were produced by three different suppliers (Cannerra, ITC and Sintef) according to the same specification.

Sensor area	$75 \times 42 \text{ mm}^2$	Active area	$73 \times 40 \text{ mm}^2$
Sensor thickness	$300 \mu\text{m}$	Strip pitch	$95 \mu\text{m}$
Strips per side	768	Stereo angle	35 mrad
Strip tilt p-side	+7.5 mrad	Strip tilt n-side	-27.5 mrad
Operating voltage	20–90 V	Bias+guard current	$\leq 7 \mu\text{A}$

### 3.1.3.2 Detection module

The sensors are  $300 \mu\text{m}$  thick and they have 768 strips on each side with a pitch of  $95 \mu\text{m}$ . The stereo angle is 35 mrad which is a compromise between stereo view and reduction of ambiguities resulting from high particle densities. The stereo angle is obtained by defining strips with an angle of 7.5 mrad with respect to the sensor short edge on the p-side and with an angle of 27.5 mrad on the n-side. Sensors are mounted with the strips nearly parallel to the magnetic field in order to optimize the resolution in the bending direction, see table 3.9. The sensor p-side (n-side) of layer 5 (layer 6) faces the interaction region. This results in four almost equally spaced strip orientations in the two layers, which significantly reduces the number of ambiguities seen by the tracking software.

The hybrids are obtained by connecting six HAL25 front-end chips [79] to an aluminium on polyimide hybrid circuit (flex) which is supported by the stiffener. The chips are glued directly

**Table 3.10:** HAL25 front-end chip parameters.

Non-linearity	<5%
Readout	sequential analogue samples
Readout speed	10 MHz at 20 pF/200 $\Omega$
Supply	+2.5 V single supply
Dissipation	340 $\mu$ W /channel
Internal parameter control	JTAG
Size	3.8x11 mm <sup>2</sup>
Thickness	150 $\mu$ m
Technology	0.25 $\mu$ m IBM CMOS, 3 metal layers
Design	mixed analogue and digital readout ASIC
Analogue channels	128
Input connection pitch	80 $\mu$ m
Output connection pitch	125 $\mu$ m
ENC with sensor connected	< 400 $e^-$
Shaping time	1.2–2.4 $\mu$ s adjustable
Output	250 $\mu$ A bipolar differential current
Range	14 MIP

on the stiffener which provides a thermal conductivity about 1.3 times better than copper while keeping the multiple scattering to a minimum (about 0.03%  $X_0$ ). Each stiffener is connected to the cooling tubes by small aluminium clamps. The flex consists of two layers of 30  $\mu$ m thick aluminium traces on a 20  $\mu$ m thick polyimide foil. These layers are glued together and the electrical connections between them are made using single point tape automated bonding (spTAB) [85]. The flex hosts power lines as well as digital and analogue lines for driving and reading the chips. The decoupling capacitors and transmission line termination resistors are soldered onto the flex.

The HAL25 is a mixed analogue and digital ASIC realized in IBM 0.25  $\mu$ m CMOS technology, see table 3.10.

It is designed to be radiation tolerant up to several Mrad. It includes 128 analogue channels. Each channel has a preamplifier to convert the positive or negative charge input from the sensor into an analogue voltage step whose magnitude is proportional to the charge. Its input range is about 300 000 electrons, corresponding to about 14 MIPS in this set-up. The preamplifier is followed by a shaping circuit whose shaping time is remotely adjustable between 1.4  $\mu$ s and 2.4  $\mu$ s. The analogue peak value is stored in a sample and hold circuit controlled by an external HOLD signal generated in the FEROM, see section 3.1.3.4. The analogue samples are shifted out serially through an integrated analogue multiplexer at a maximum speed of 10 MHz under control of a clock and token generated in the FEROM. All internal parameters of the HAL25 can be set through a JTAG bus. For testing purposes the HAL25 contains an internal pulser and boundary scan registers.

All connections between the chips and the flex and between the chips and the sensor are made by chip-cables connected via spTAB. Chip-cables consist of 14  $\mu$ m thick aluminium traces supported by a 10  $\mu$ m thick polyimide foil [80]. They allow for an easy testing of the chip before



**Figure 3.27:** Photograph of a hybrid. Each hybrid was tested individually, using the connection area at the left side of the flex. This part was cut off before connecting to the sensor. The lower part of the photograph shows the six chip-cables which will connect the HAL25 chips to the sensor. The chips themselves are covered by the chip-cables and not visible in this picture. The aluminium clamps connecting to the cooling tubes are on the reverse side of the hybrid.

further use. The chip-cables also function as a pitch adapter between the sensor pads ( $95\ \mu\text{m}$  pitch) and the chip input pads ( $80\ \mu\text{m}$  pitch). In addition the use of a flexible connection between sensor and electronics makes it possible to fold the hybrids on top of the sensor, creating a compact object. The detection modules are mounted on 1 m long carbon fibre supports [81] such that the active areas of the sensors overlap 0.9 mm in the  $z$ -direction. Each module is connected to the support by three pins made of G10. A photograph of a hybrid is shown in figure 3.27.

### 3.1.3.3 Endcaps

The detection modules are connected to the Front-End ReadOut Module (FEROM) by means of the endcap modules [83]. These modules are placed at both ends of each ladder.

Each endcap module controls the detector modules on a half ladder, the number of modules depending on which layer and which side. Due to the limited available space ( $7 \times 7 \times 5\ \text{cm}^3$ ) the electronics is miniaturized and active cooling is necessary. The electronics is able to survive Single Event Effects (SEE) and about 500 Gy of total ionizing dose. It is built around two ASIC's in the same IBM  $0.25\ \mu\text{m}$  CMOS technology as the HAL25. The control functions for the front-end and the endcap, like power regulation and latch-up protection (over current), signal coupling, readout control and error detection are integrated in the ALCAPONE chip. The ALABUF includes drivers for the analogue signals and switches between the analogue signals from the p-side and n-side of the detection modules.

The electronics that reads out each side of the sensor, operate at a different potential, defined by the sensor bias. This sensor bias voltage is provided symmetrically with respect to ground potential. Therefore, all signals are connected via AC coupling to each detector side, see also figure 3.26 on page 47. Each hybrid has its own ALCAPONE with a power regulator that has a programmable voltage, but a fixed current limit. The HAL25 front-end chips need to be configured via the JTAG serial bus. This bus is distributed by the ALCAPONE chips. The endcap itself is also controllable via the same JTAG connection.

At the start of the readout, a token is sent by the FEROM to all the modules, first to the



p-hybrids and then to the n-hybrids. Analogue switches in the ALABUF multiplex the analogue signals from the p-side and n-side of the detection module to create a single differential signal per module which is transmitted to the ADC's in the FEROM. The endcap is connected to the outside via a custom made cable which provides the LVDS control signals for the readout, the JTAG communication, the detector bias, and the low power supply voltages from the floating power supplies. The analogue outputs are also transmitted via this cable bundle. All signals on this cable are differential to minimize interference and are at safe ground potential, because the AC coupling is done completely inside the endcap.

### 3.1.3.4 Readout

The main task of the Front-End ReadOut Module (FEROM) [84] is to digitize the 2.6 million analogue samples from the front-end modules, keeping up with the trigger rate in ALICE. This is achieved by digitizing the signals from each of the 1698 detection modules in parallel. The detection modules themselves can store only one analogue sample per input. They serialize the 1536 analogue samples before sending them to the FEROM.

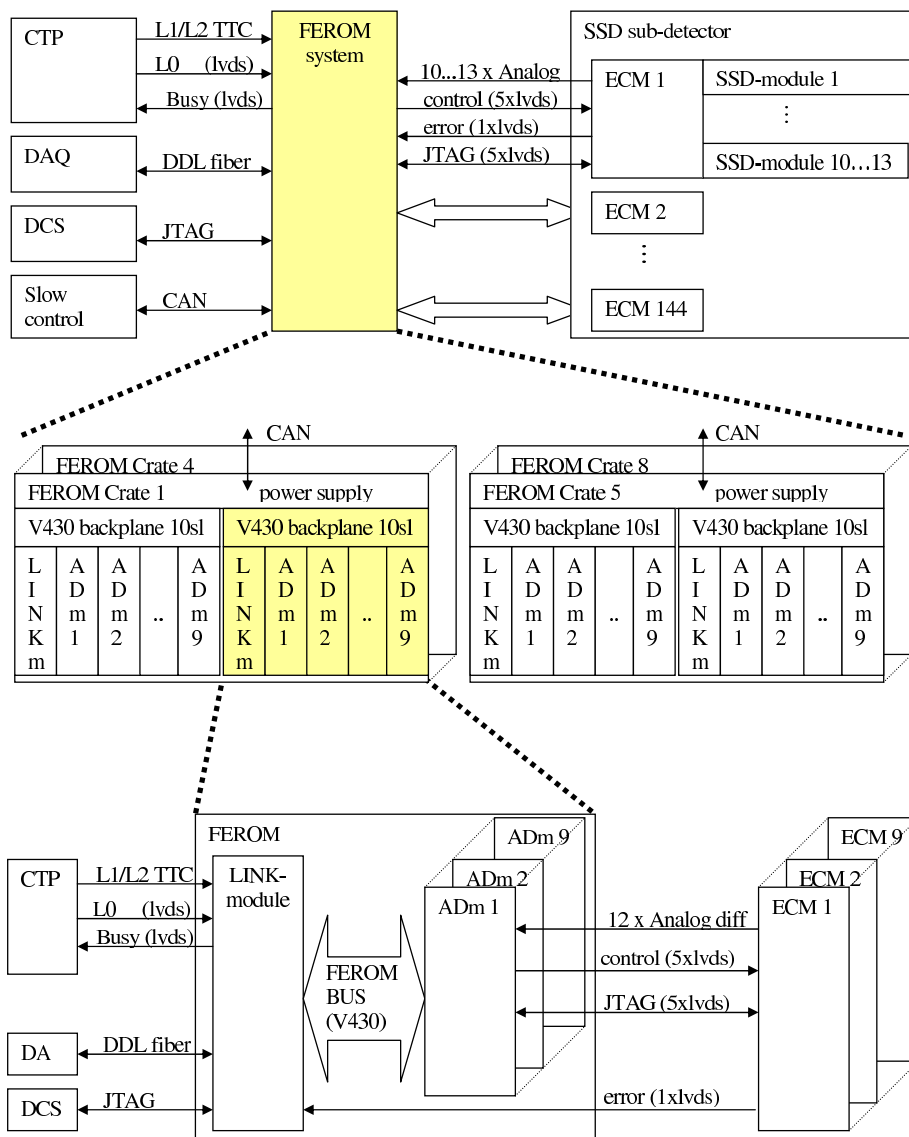
The FEROM system, see figure 3.28, consists of 8 crates each containing 216 analogue-to-digital converters (AD modules) and two interface cards (LINK modules) connected to the data acquisition system, the detector control system and the central trigger processor. These crates are located in the cavern just outside the ALICE magnet. They connect to the endcaps via approximately 40 m long twisted pair cables.

Upon an L0 trigger the FEROM activates the HOLD signal with an adjustable delay, matching the shaping time of the amplifiers in the HAL25 on the detection module, see section 3.1.3.2.

If the event is accepted by the L1 trigger decision the FEROMs generate the control signals needed for the readout of the detection modules. The analogue samples of the two hybrids of each module are shifted out. With the chosen readout rate of 10 MHz the total conversion time for one event is less than  $160 \mu\text{s}$ . The design of the front-end electronics does not allow overlapping L0 triggers before an L2 trigger decision is made. In addition the BUSY of the readout system remains set for  $40 \mu\text{s}$  after completion of a readout cycle in order to allow the ladder electronics to quiet down. Therefore the readout rate is limited to about 5 kHz. In case the zero suppression in the FEROM is switched off or the occupancy increases above 6% the readout rate is limited by the bandwidth of the optical link to the data acquisition system at 25 MB/s per FEROM crate. The expected occupancy for the most central events in Pb-Pb collisions is below 3%.

In case the event is rejected by a L2 trigger decision a fast-clear signal is sent to the endcaps and to the front-end chips and the obtained digital and analogue information is discarded. In order not to overload the data acquisition system the digitized data are corrected for pedestals and zero-suppressed for each sample individually. The remaining data are stored in an event buffer and sent to the data acquisition system asynchronously. The FPGA based design allows easy upgrading of the algorithm, when needed. Currently an additional algorithm to suppress baseline distortion and its implementation is being studied so that it can be installed should the experimental environment prove different from the laboratory environment.

The detector control system communicates with the FEROM using the JTAG protocol. The JTAG chain can be used to check interconnections inside the FEROM and the connections between



**Figure 3.28:** The embedding of the FEROM between the ALICE data acquisition and control system (left) and the SSD detector system (right) is shown in the top part of this figure. The central part shows the configuration of 8 FEROM crates, each equipped with AD modules and LINK modules. The bottom section shows the internal communication between the 9 AD modules and the single LINK module in each crate.

the FEROM and the endcap. Each endcap can include the corresponding HAL25 chips in the JTAG chain when needed.

As the FEROMs are installed in the cavern and therefore inaccessible during beam-time, much emphasis was put on high reliability. Each card in a FEROM crate has a multidrop addressable JTAG switch. In this way the scope of a single point failure is reduced to the branch behind the switch. Each card serves up to 12 front-end modules, corresponding to about 0.7% of the SSD system. Special precautions were taken to reduce SEU problems due to the 0.1 Gy irradiation

dose expected. In case the neutron dose outside the ALICE magnet proves to be much higher than expected, an upgrade based on the much more radiation tolerant ACTEL flash FPGA instead of the currently used Xilinx SRAM FPGA is already prepared. A dedicated addressing protocol with TTL comparators, a parity check over the measured data and poly switch fuses on each module limit the consequence of single event effects. Furthermore, each backplane in a FEROM crate is split in two and each half is connected with a separate data link to the DAQ system.

### 3.1.3.5 Power and cooling

All the sensors mounted on the same half-ladder are biased in parallel by means of two power supply channels belonging respectively to a positive and a negative polarity power supply. These are CAEN power supplies (A3501P and A3501N) which are connected in series to provide the sensors with a bias voltage which is symmetric with respect to a reference ground connected to the central point. The central point is connected to ground in each endcap. The signal ground is connected to the safety ground at the input of the ADC modules in the FEROM. The ladder structure is connected to the endcap signal ground but otherwise floating with respect to the rest of the mechanical structure.

Each hybrid floats upon the potential level of the sensor side to which the hybrid is connected. The voltage needed for operating the hybrids (2.5 V) is supplied by regulators in the endcap which are powered at 4 V by low-voltage cards (CAEN A3602). Two channels of the low-voltage cards are used to power respectively the p-side and n-side hybrids of a half-ladder, providing a typical current of 2 A. The third low-voltage channel is used to power the endcap electronics, which is referenced to ground. The powering of both hybrid and endcaps electronics makes use of sense wires to correct for the voltage drop in the connecting cables.

For homogeneity purposes all three low-voltage channels are identical, even though the power needed by the endcap electronics is lower. The choice of using floating power supplies helps in reducing possible common mode phenomena and allows a unique reference ground thereby preventing the occurrence of ground loops.

Both the bias and the low-voltage cards are part of the CAEN EASY system, which is designed to operate in high radiation environment, being able to tolerate moderate radiation doses and to operate inside intense magnetic fields. This choice was necessary since the power supplies will be sitting just outside the L3 magnet in an area where the magnetic field will be as high as 300 Gauss and where the integrated dose of order 0.1 Gy. The power supply modules will be located in the experimental area while the monitoring and control will be made using an interface card (CAEN-A1676A) sitting in a mainframe (CAEN-SY1527) located in the counting room.

The segmentation of the power supplies allows us to switch on each half-ladder individually. In addition each hybrid on each detection module can be switched on individually by the endcaps. Therefore the malfunction of one module will not influence the operation of the others. A short in the power cabling outside the ITS will also affect at most one half-ladder.

The average power dissipated inside the barrel is 2.2 kW of which one third in the detection modules. The actual power depends mainly on the readout rate. Outside the barrel up to 1 kW is dissipated in the cables for which no special cooling is provided.

**Table 3.11:** SSD performance.

Spatial precision $r\phi$	20 $\mu\text{m}$	Spatial precision $z$	820 $\mu\text{m}$
Material budget L5	0.84%	Material budget L6	0.84%
Support cone	0.53%	readout duration	180 $\mu\text{s}$
Dead channels	< 3%	S/N for 1 MIP	> 30

Each ladder has two cooling pipes operated in counter flow with the endcap in series with the return on its side. The ladders are grouped into 8 segments corresponding to the segmentation of the FEROMs. The power supplies are interlocked with temperature sensors on each endcap. The water system is designed as a leak proof system; it operates below atmospheric pressure. A moderate airflow is provided to even out moisture and temperature gradients inside the SSD volume. Tests and simulations [29, 30] have shown that the SSD can operate in a thermal neutral way with the water temperature about 5 K below ambient temperature.

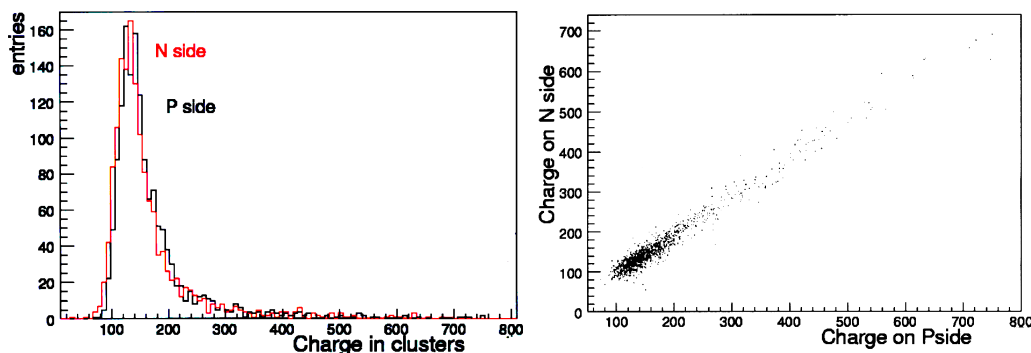
### 3.1.3.6 Performance

The spatial resolution of the SSD system is determined by the 95  $\mu\text{m}$  pitch of the sensor readout strips and the charge-sharing between those strips (see table 3.11). Without making use of the analogue information the r.m.s spatial resolution is 27  $\mu\text{m}$ . Beam tests have shown that a spatial resolution of better than 20  $\mu\text{m}$  in the  $r\phi$  direction can be obtained by analysing the charge distribution within each cluster. In the direction along the beam the spatial resolution is 820  $\mu\text{m}$ .

Almost one half of the material budget of the SSD layers consists of the sensors. The sensors amount to 0.36% of radiation length; the remainder is distributed between mechanics (0.20%), cabling (0.16%), and electronics components (0.11%) for a total thickness in each layer of only 0.84%  $X_0$ . The hybrids and cabling on the ladders were reduced to the minimum amount of material required for an efficient heat transfer between the electronics and the cooling tubes and for an acceptable electrical resistance in the cabling. The cylinder connecting the two cones which support the ladders adds 0.53%  $X_0$ . Thus, the average amount of material seen by particles traversing the two layer SSD at  $\theta=90^\circ$  is 2.2%  $X_0$ .

The positions of the detection modules referred to the ladder ends were measured for every ladder before mounting. The sag of a ladder due to gravity was modelled and checked by measurements [86] on selected ladders to be less than 50  $\mu\text{m}$ . The position of the ladder mounting points was measured relative to the ITS reference points before mounting the ladders. Using all these data, the module positions on the completed SSD are expected to be known within 60  $\mu\text{m}$ . However, the procedure of precisely checking the module's positions of the completed SSD is based on real tracks.

The efficiency of the SSD system is mainly determined by the number of defective channels and the signal to noise ratio (SNR) obtained in the final system. Beam tests on single modules and laboratory tests with the finally produced ladders have shown that a SNR of better than 30 is achieved in all modules in the final system. In figure 3.29 a typical response of a module to cosmics is shown. For details of the cluster reconstruction method see [87]. The number of non-working channels, i.e. high noise or no digital communication is about 3%.



**Figure 3.29:** Typical response of a module to cosmics. The left panel shows a histogram of the total signal measured at the ADC for the reconstructed clusters. The right panel shows the correlation between the signals measured at the P and N sides of the sensor.

## 3.2 Time-Projection Chamber (TPC)

### 3.2.1 Design considerations

The Time-Projection Chamber (TPC) [14] is the main tracking detector of the central barrel and is optimised to provide, together with the other central barrel detectors, charged-particle momentum measurements with good two-track separation, particle identification, and vertex determination, see [14] and chapter 8 of this document. In addition, data from the central barrel detectors are used to generate a fast online High-Level Trigger (HLT) for the selection of low cross section signals, see section 6.3.

The phase space covered by the TPC in pseudo-rapidity is  $|\eta| < 0.9$  for tracks with full radial track length (matches in ITS, TRD, and TOF detectors); for reduced track length (at reduced momentum resolution), an acceptance up to about  $|\eta| = 1.5$  is accessible. The TPC covers the full azimuth (with the exception of the dead zones). A large  $p_t$  range is covered from low  $p_t$  of about 0.1 GeV/c up to 100 GeV/c with good momentum resolution, see chapter 8.

At the Pb-Pb design luminosity of the LHC, an interaction rate of 8 kHz is expected, of which 10% are to be considered as central collisions. In the TPC design phase, before RHIC data allowed a better extrapolation to LHC energy, an extreme charge particle multiplicity density of  $dN_{ch}/d\eta = 8000$  was assumed, which would result in 20 000 charged primary and secondary tracks in the TPC acceptance, an unprecedented track density for a TPC. These extreme multiplicities set new demands on the design which were addressed by extensive R&D activities. Careful optimisation of the TPC design finally resulted in maximum occupancies (defined as the ratio of the number of readout pads and time bins above threshold to all pads and time bins) of about 40% at the innermost radius and 15% at the outermost radius for the quoted extreme multiplicity.

It is expected that in Pb-Pb runs the TPC can be operated at central collision rates of up to 200 Hz. While there is not much operational experience with large TPCs at these rates, the current load on the readout chambers is not excessive. Simulations have shown that, at this rate, the space charge due to the ion feed-back during gate-open time starts to be comparable to the space charge due to the ionisation in the TPC drift volume itself, resulting in tracking distortions of order a few

mm. The true space-charge limit, however, might eventually depend on the background conditions due to the beam quality. Offline corrections for the space charge are expected to recover part of the resolution loss and should thus extend this limit.

For proton-proton runs, the memory time of the TPC is the limiting factor for the luminosity due to the  $\sim 90 \mu\text{s}$  drift time. At a pp luminosity of about  $5 \times 10^{30} \text{ cm}^{-2}\text{s}^{-1}$ , with a corresponding interaction rate of about 350 kHz, ‘past’ and ‘future’ tracks from an average of 60 pp interactions are detected together with the triggered event; the detected multiplicity corresponds to about 30 minimum-bias pp events. The total occupancy, however, is lower by more than an order of magnitude than in Pb-Pb collisions, since the average pp multiplicity is about a factor  $10^3$  lower than the Pb-Pb multiplicity for central collisions [4]. Tracks from pile-up events can be eliminated because they point to the wrong vertex.

The space charge, due to both ionisation in the drift volume and to the ion feed-back during gate-open time, is about one order of magnitude lower than for Pb-Pb, and thus not seen as a problem. Trigger rates of up to 1 kHz seem to be realistic.

Before installation of the TPC in the ALICE cavern in January 2007, beam tests were performed with an inner readout chamber in a special setup [88] and the completed detector was commissioned sector-by-sector with laser and cosmic muon tracks triggered by ACORDE (see chapter 3.8), so that essentially all the design features underwent a serious test. As a result, the main specifications as laid down in the TPC TDR [14] were verified by actual measurements. In the following, we briefly review design features and components of the TPC.

### 3.2.2 Detector layout

The TPC design is ‘conventional’ in overall structure but innovative in many aspects. A summary of its main parameters is presented in table 3.12. The TPC is cylindrical in shape; 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.


The detector is made of a large cylindrical field cage, filled with  $90 \text{ m}^3$  of Ne/CO<sub>2</sub>/N<sub>2</sub> (90/10/5), in which the primary electrons are transported over a distance of up to 2.5 m on either side of the central electrode to the end plates. Multi-wire proportional chambers with cathode pad readout are mounted into 18 trapezoidal sectors at each end plate.

The TPC layout is shown in several figures: figure 3.30 shows a 3D cut-open view of the field cage. Details of the support rods and mylar strips for the 3 bottom sectors of one half of the TPC are shown in figure 3.31. figure 3.32 is a cut along the beam ( $yz$ ), figure 3.33 is a cut perpendicular to the beam ( $xy$ ) at one of the two, identical end plates.


**Field cage.** The field cage is based on a design with a central high-voltage electrode and two opposite axial potential dividers which create a highly uniform electrostatic field in the common gas volume, see figures 3.30 to 3.33. The central electrode is a stretched aluminized Mylar foil of  $22 \mu\text{m}$  thickness to satisfy the requirement of minimal material near  $90^\circ$  relative to the beam direction. The electrical potential in the drift region is defined by aluminised Mylar strips wound around 18 inner and outer support rods which also contain two resistive potential dividers (inner and outer), see figure 3.31. The rods are aligned with the dead zones in-between the readout chambers.

**Table 3.12:** Synopsis of TPC parameters.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$ for full radial track length $-1.5 < \eta < 1.5$ for 1/3 radial track length
Azimuthal coverage	$360^\circ$
Radial position (active volume)	$848 < r < 2466$ mm
Radial size of vessel (outer dimensions)	$606.5 < r < 2780$ mm
Radial size of vessel (drift gas volume)	$788 < r < 2580$ mm
Length (active volume)	$2 \times 2500$ mm
Segmentation in $\varphi$	18 sectors
Segmentation in r	2 chambers per sector
Total number of readout chambers	$2 \times 2 \times 18 = 72$
Inner readout chamber geometry	trapezoidal, $848 < r < 1321$ mm active area
pad size	$4 \times 7.5$ mm ( $\varphi \times r$ )
pad rows	63
total pads	5504
Outer readout chamber geometry	trapezoidal, $1346 < r < 2466$ mm active area
pad size	$6 \times 10$ and $6 \times 15$ mm ( $\varphi \times r$ )
pad rows	$64 + 32 = 96$ (small and large pads)
total pads	$5952 + 4032 = 9984$ (small and large pads)
Detector gas	Ne/CO <sub>2</sub> /N <sub>2</sub> 90/10/5
Gas volume	90 m <sup>3</sup>
Drift length	$2 \times 2500$ mm
Drift field	400 V/cm
Drift velocity	2.7 cm/ $\mu$ s
Maximum drift time	92 $\mu$ s
Total HV	100 kV
Diffusion	$D_L = D_T = 220 \mu\text{m}/\sqrt{\text{cm}}$
Material budget	$X/X_0 = 3.5\%$ near $\eta = 0$
Front-End Cards (FEC)	121 per sector $\times 36 = 4356$
Readout Control Unit (RCU) scheme	6 per sector, 18 to 25 FEC per RCU
Total RCUs	216
Total pads — readout channels	557 568
Pad occupancy (for $dN/dy = 8000$ )	40 to 15% inner / outer radius
Pad occupancy (for pp)	$5$ to $2 \times 10^{-4}$ inner / outer radius
Event size (for $dN/dy = 8000$ )	$\sim 90$ MB (conditions see text)
Event size (for pp)	$\sim 1\text{--}4$ MB (depending on pile-up and noise threshold)
Total bandwidth	$\sim 30$ GB/s
Trigger rate limits	300 Hz Pb-Pb central events 1000 Hz proton-proton events
ADC	10 bit
sampling frequency	5–10 MHz
time samples	500–1000
Conversion gain	6 ADC counts / fC
Position resolution ( $\sigma$ ) in $r\varphi$	1100 to 800 $\mu\text{m}$ inner / outer radii
in $z$	1250 to 1100 $\mu\text{m}$
$dE/dx$ resolution, isolated tracks	5.0%
$dN/dy = 8000$	6.8%




**Figure 3.30:** 3D view of the TPC field cage and service support wheels, without readout chambers. The containment vessels are cut open. The support rods and mylar strips are not shown, see figure 3.31 for these details.




**Figure 3.31:** Detail of the field cage structure for three sectors at the bottom, showing the support rods and the mylar strips (near the end plate, the first few strips are all shown, further on for clarity only a few. There are 166 strips in total on each side). One inner and one outer rod contain a resistive divider.





**Figure 3.32:** Cross section along the beam axis of the TPC field cage with some dimensions.



**Figure 3.33:** Layout of the TPC end plate.


Because of the Ne/CO<sub>2</sub>/N<sub>2</sub> gas mixture used in the TPC, the field cage is operated at high voltage gradients, of about 400 V/cm, with a high voltage of 100 kV at the central electrode, which results in a maximum drift time of about 90  $\mu$ s.

An insulating gas envelope of CO<sub>2</sub> in containment vessels surrounds the field cage. The field cage and containment volumes are each constructed from two concentric cylinders, sealed by the end plate on either side. To provide high structural integrity against gravitational and thermal loads while keeping the material budget low, composite materials were used. Hence the mechanical stability and precision is guaranteed to be about 250  $\mu$ m. The coplanarity of the central electrode and of the entrance grids of the readout chambers have been carefully surveyed and adjusted to this precision by individual shimming of the readout chamber mounts.


**Drift gas system.** The drift gas Ne/CO<sub>2</sub>/N<sub>2</sub> (90/10/5) is optimised for drift speed, low diffusion, low radiation length and hence low multiple scattering, small space-charge effect, and ageing and stability properties. Mixtures containing CH<sub>4</sub> and CF<sub>4</sub> were rejected due to their ageing properties. The N<sub>2</sub> admixture improves the quenching and allows higher maximum gas gains [89]. The drawback of Ne/CO<sub>2</sub> is that this mixture is a ‘cold’ gas, with a steep dependence of drift velocity on temperature [90]. For this reason, the TPC is aiming for a thermal stability with  $\Delta T \leq 0.1$  K in the drift volume, see below.

The gas system circulates and purifies the gas mixture, with very low fresh gas injection. The pressure follows the ambient pressure. The CO<sub>2</sub> and N<sub>2</sub> fractions are kept stable to 0.1%, necessary to ensure stable drift velocity and gas gain of the readout chambers. The O<sub>2</sub> impurity as achieved with a reduced flow during commissioning is extrapolated to be about 1 ppm or lower in the final installation in the cavern, limiting the signal reduction due to attachment for the maximum drift length of 2.5 m to < 5%.

**Readout chambers.** The readout chambers instrument the two end plates of the TPC cylinder with an overall active area of 32.5 m<sup>2</sup> [91–93]. The chambers are multi-wire proportional chambers with cathode pad readout. Because of the radial dependence of the track density, the readout is segmented radially into two readout chambers with slightly different wire geometry adapted to the varying pad sizes mentioned below. The radial range of the active area is from 84.8 cm to 132 cm (and from 134.6 cm to 246.6 cm) for the inner (and outer) chamber, respectively. Note that the active area of the readout chambers are larger than the openings of the end plate in figure 3.33, as seen, e.g., in figure 3.36. This design optimizes active area; it required, however, the mounting of the readout chamber from the inside. This was accomplished with a mounting tool that allowed insertion of a chamber through the end plates into the field cage, rotating and positioning of the chamber once inside, and finally engaging the chamber in its proper orientation with the mount points on the end plate from inside. The inactive areas between neighbouring inner chambers are aligned with those between neighbouring outer chambers. Such an arrangement optimises the momentum precision for detected high-momentum tracks but has the drawback of creating cracks in the acceptance — in about 10% of the azimuthal angle the detector is non-sensitive. The readout chambers are made of standard wire planes, i.e. they consist of a grid of anode wires above the pad plane, a cathode wire plane, and a gating grid, see figures 3.34 and 3.35 showing cuts along and perpendicular to the wires for the inner readout chamber. They combine, however, a very



**Figure 3.34:** Cross section of inner readout chamber along wires, with wire diameters and wire tensions.




**Figure 3.35:** Cross section of inner readout chamber perpendicular to wires.

tight electrode configuration (anode-to-cathode distance 2 and 3 mm for inner and outer readout chamber, respectively) with unprecedented operating gain (up to 20 000).

To keep the occupancy as low as possible and to ensure the necessary  $dE/dx$ , position and two-track resolution, there are about 560 000 readout pads of three different sizes:  $4 \times 7.5 \text{ mm}^2$  in the inner chambers,  $6 \times 10 \text{ mm}^2$  and  $6 \times 15 \text{ mm}^2$  in the outer chambers.

The mechanical mount of the readout chambers is a 3-point kinematical mount. The seal to the end plate is a flexible foil pressed against O-rings both on the chambers and the end plate, see figure 3.36. Coplanarity was checked by survey and adjusted (using individually machined spacers for the 3-point mounts) with an accuracy of  $\leq 250 \mu\text{m}$ .



**Figure 3.36:** Cut of end plate with details of readout chamber mount points and seal with a foil (cross-hatched). The pad plane side of the chambers faces downwards in the drawing, it is larger than the end plate opening. Mounting is thus from the inside (see text).

**Gating.** The readout chambers are normally closed by a gating grid for electrons coming from the drift volume and for ions from the amplification region. The gate is opened only by the L1 trigger ( $6.5 \mu\text{s}$  after the collision) for the duration of one drift-time interval, i.e. of about  $90 \mu\text{s}$  (see section 6.1). This helps to prevent space charge due to positive ions from drifting back from the multiplication region for non-triggered interactions and background. It has been verified that the closed gate reduces the drift of electrons by a factor better than  $10^5$ .

**Laser calibration.** The laser system [94] provides straight tracks distributed over all regions of the drift space. On each end plate, a laser beam of about 2 cm diameter is distributed to 6 of the 18 outer support rods of the field cage. At  $2 \times 4$   $z$ -positions, these beams are deflected by micro-mirror bundles, producing 7 fine laser rays of about 1 mm diameter traversing the drift volume perpendicular to the beam axis at strategic angles. This system is used for precise position inter-calibration for the readout chambers and will allow online monitoring of temperature and space-charge distortions, both of the order of a few mm. The laser can be triggered with a rate of about 10 Hz.


**Cooling and temperature control.** The requirement of 0.1 K thermal uniformity in the TPC drift volume, in the presence of thermal gradients inside the ALICE solenoid of about 5 K according to simulations, requires an elaborate system of heat screens and cooling circuits: (i) a heat screen



**Figure 3.37:** The TPC as assembled above ground on its way to the ALICE cavern.

at the outer radius toward the TRD, (ii) heat screens at the inner radius shielding from the ITS services, (iii) heat screens in the readout chamber bodies shielding from the FEE heat dissipation, (iv) cooling of the front-end electronics itself, (v) cooling of the resistive potential dividers, (vi) a heat screen at the FEE towards the outside. All these circuits are leakless i.e. operated at pressure below ambient. In the case of the resistive potential divider of the 100 kV drift voltage very high resistivity water is used. For material budget reasons, no extra heat shield is used in the central region around the ITS barrel; the ITS controls its surface temperature there itself.

**Integration into ALICE.** The TPC field cage, which weighs 8 t including the end plates and readout chambers, occupies the central opening of the spaceframe. The photograph in figure 3.37 shows the TPC assembled above ground on its transport to the ALICE cavern. The TPC is mounted inside the spaceframe on a rail system, with 4 feet gliding on Teflon. In the installation phase of the ITS and beam line, and possibly in future service periods, the TPC is partially retracted from the spaceframe by about 5 m to a ‘parking position’, allowing access to the area around the interaction point (see figure 3.38).




**Figure 3.38:** The TPC in the spaceframe, seen from the muon absorber side, at the ‘parking position’ which allows access to the ITS. The ITS is supported by temporary rails that are finally removed. Note the thermal screen towards the TRD.

The front-end electronics and the services are not directly mounted on the TPC end plates but in two separate structures called service support wheels which cover the end plates and are independently supported by the TPC rails. This avoids loading the readout chambers with the weight of the electronics and services (1 t for the readout cards alone on each side) and reduces thermal coupling to the drift volume. The ITS occupies the inner free radius of the TPC field cage, see chapter 3.1. The weight of the ITS is supported at two points by the innermost shell of the field cage.

### 3.2.3 Front-end electronics and readout

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

A single readout channel is comprised of three basic functional units, see block diagram of figure 3.39: i) a charge sensitive amplifier/shaper (PASA); ii) a 10-bit 25 MHz low-power ADC; and iii) a digital circuit that contains a shortening filter for the tail cancellation, baseline subtraction



**Figure 3.39:** Block diagram of the TPC front-end electronics.

and zero-suppression circuits, and a multiple-event buffer. A front-end card (FEC) contains 128 channels. For the subsequent readout via a front-end bus and RCU see below.

The charge collected on the TPC pads is amplified and integrated by a low-input impedance amplifier. It is based on a Charge Sensitive Amplifier (CSA) followed by a semi-Gaussian pulse shaper of the 4<sup>th</sup> order. These analogue functions are realized by a custom integrated circuit (PASA), implemented in 0.35  $\mu\text{m}$  CMOS technology, which contains 16 channels with a power consumption of 12 mW per channel. The circuit has a conversion gain of 12 mV/fC and a dual output dynamic range of  $2 \times 1$  V with a linearity of 0.2%. It produces a pulse with a rise time of 150 ns and a shaping time (FWHM) of 190 ns. The single channel has a noise value (r.m.s.) of  $\sim 550 e$  at the typical input capacity of 12 pF and a channel-to-channel cross-talk below  $-60$  dB. Immediately after the PASA, a 10-bit pipelined ADC (one per channel) samples the signal at a rate of 5–10 MHz. The digitized signal is then processed by a set of circuits that perform the baseline subtraction, tail cancellation, zero-suppression, formatting and buffering. The ADC and the digital circuits are contained in a single chip named ALTRO [95, 96] (ALice Tpc ReadOut). The ALTRO chip integrates 16 channels, each of them consisting of a 10-bit, 25 MHz ADC, a pipelined Digital Processor and a multi-event buffer memory. When an L1 trigger is received, 500 or 1k time samples are temporarily stored in the multi-event buffer. Upon L2 trigger arrival the event is frozen, otherwise it is overwritten by the next event. The Digital Processor, running at the sampling frequency, implements several algorithms that are used to condition and shape the signal. After digitization, the Baseline Correction Unit I is able to perform channel-to-channel gain equalization and to correct for possible non-linearity and baseline drift of the input signal. It is also able to adjust DC levels and to remove systematic spurious signals by subtracting a pattern stored in a dedicated memory. The next processing block is an 18-bit, fixed-point arithmetic, 3<sup>rd</sup> order Tail Cancellation Filter [97]. The latter is able to suppress the signal tail, within 1  $\mu\text{s}$  after the pulse peak, with an accuracy of 1 LSB. Since the coefficients of this filter are fully programmable, the circuit is able to cancel a wide range of signal-tail shapes. Moreover, these coefficients can be set independently for each channel and are re-configurable. This feature allows a constant quality of the output signal

regardless of ageing effects on the detector and channel-to-channel fluctuations. The subsequent processing block, Baseline Correction Unit II, applies a baseline correction scheme based on a moving average filter. This scheme removes non-systematic perturbations of the baseline that are superimposed on the signal. At the output of this block, the signal baseline is constant with an accuracy of 1 LSB. Such accuracy allows an efficient zero-suppression procedure, which discards all data below a programmable threshold, except for a specified number of pre- and post-samples around each pulse. This produces a limited number of non-zero data packets, thus reducing the overall data volume. Each data packet is formatted with its time stamp and size information in such a way that reconstruction is possible afterwards. The output of the Data Processor is sent to a Data Memory of 5 kB, able to store up to 8 full events. The data is read out from the chip at 40 MHz through a 40-bit wide bus, yielding a total bandwidth of 200 MB/s. Moreover, the readout speed and the ADC sampling frequency are independent. Therefore, the readout frequency does not depend on the bandwidth of the input signal being acquired. The ALTRO chip is implemented in the ST 0.25  $\mu\text{m}$  HCMOS-7 process.

The readout chain up to the ALTRO chip is contained in the Front-End Cards [93] (FEC), which are plugged in support frames mounted to the service support wheel. Each FEC contains 128 channels (8 PASA and 8 ALTRO chips), and is connected to the cathode plane by means of 6 flexible cables. One row of FECs (up to 25) is controlled by a Readout Control Unit (RCU) [98], which interfaces the FECs to the DAQ, the trigger, and the Detector Control System (DCS) [99]. The RCU broadcasts the trigger information to the individual FEC modules and controls the readout procedure. Both functions are implemented via a custom bus, based on low-voltage signalling technology (GTL), which provides a data bandwidth of 200 MB/s. The interfacing of the RCU modules to the Trigger and to the DAQ follows the standard data acquisition architecture of the experiment. In summary, for each of the 36 TPC sectors, the front-end electronics consists of 121 FECs in 6 rows, 6 RCUs, and 6 Detector Data Links (DDL).

The event sizes and rates quoted in table 3.12 are for the 1000 time-sample option. After zero-suppression and data encoding, the event size from the TPC for a central Pb-Pb collision will be about 90 MB at a charged multiplicity of  $dN_{\text{ch}}/d\eta = 8000$ , or 30 MB at  $dN_{\text{ch}}/d\eta = 2500$ . For the 500-time-sample option, the event sizes reduce to 60% (less than a factor 2 due to the constant coding overhead). For pp-events, the data volume is very sensitive to the zero-suppression threshold. For a threshold at  $3\sigma$ , 1.4 MB of the data volume arise from the noise alone, the pp track data itself are less than 100 kB (both including coding). The data rate capabilities of the TPC readout, with a total bandwidth of 27 GB/s measured in commissioning, allow transfer of 300 central Pb-Pb events/s for  $dN_{\text{ch}}/d\eta = 8000$ . This exceeds the recording capabilities of the DAQ (see chapter 6.2) by far. The recorded data volume can be reduced using the ‘region-of-interest’ option of the trigger, reading out only a few sectors of the TPC.

In order to increase the physics potential of ALICE especially for jet and electron physics, rare signals like dielectron pair candidates have to be enriched in the trigger mix. Therefore an ‘intelligent’ readout is available via a HLT processor farm, which will operate on the raw data shipped through optical links. The HLT allows lossless data compression, selective readout of electron candidates identified by the TRD (the ‘region-of-interest’ option of the trigger), as well as online track finding and tracking of the whole TPC. Details of the intelligent readout including strategies for its implementation are discussed in section 6.3.



### 3.3 Transition Radiation Detector (TRD)

#### 3.3.1 Design considerations

The main purpose of the ALICE Transition Radiation Detector (TRD) [16] is to provide electron identification in the central barrel for momenta above 1 GeV/c. Below this momentum electrons can be identified via specific energy loss measurement in the TPC. Above 1 GeV/c transition radiation (TR) from electrons passing a radiator can be exploited in concert with the specific energy loss in a suitable gas mixture to obtain the necessary pion rejection capability. In conjunction with data from the ITS and the TPC it is possible to study the production of light and heavy vector-meson resonances and the dilepton continuum both in pp as well as in Pb-Pb collisions. Exploiting the excellent impact parameter resolution of the ITS it is furthermore possible to reconstruct open charm and open beauty in semi-leptonic decays.

The TRD was designed to derive a fast trigger for charged particles with high momentum. It is part of the Level 1 trigger and can significantly enhance the recorded  $\Upsilon$ -yields, high- $p_t$   $J/\psi$ , the high-mass part of the dilepton continuum as well as jets.

The design parameters of the TRD are the following:

**Pion rejection capability** This is governed by the signal-to-background ratio in the measurement of  $J/\psi$  production and its  $p_t$  dependence. This led to the design goal for the pion rejection capability of a factor 100 for momenta above 1 GeV in central Pb-Pb collisions/c [6]. The measurement of the lighter vector-mesons and the determination of the continuum between the  $J/\psi$  and the  $\Upsilon$  will only be feasible when this level of rejection can be reached.

**Position and momentum resolution** It is crucial to be able to match to the TPC in order to exploit the combined momentum resolution leading to an overall mass resolution of about 100 MeV/c<sup>2</sup> at the  $\Upsilon$ -mass (for B = 0.4 T). The required pointing accuracy for electrons needs to be on the level of a fraction of a TPC pad. The anticipated momentum resolution of the TRD of 3.5 (4.7)% at 5 GeV/c (depending on multiplicity) see table 3.13 will crucially determine the sharpness of the trigger threshold in  $p_t$  as well as the capability to reject fake tracks.


**Radiation length** It has to be minimized in order to reduce Bremsstrahlung leading to incorrect momentum determination or loss of electrons and to reduce photon conversions resulting in increased occupancy as well as incorrect matching.

**Detector granularity** In the bending direction this is governed by the desired momentum resolution and in longitudinal direction by the need to correctly identify and track electrons through all layers of the detector even at the largest anticipated multiplicities. For the desired quality of the reconstructed dilepton pair signal (mostly dictated by the tracking efficiency) this leads to pads with an average area of about 6 cm<sup>2</sup>, a tracking efficiency of 90% can thus be achieved for single tracks.

Including secondary particles this leads to a maximum occupancy of 34% at the highest simulated multiplicity density of  $dN_{ch}/d\eta = 8000$ .

**Table 3.13:** Summary of TRD parameters.

Pseudo-rapidity coverage	$-0.84 < \eta < 0.84$
Azimuthal coverage	$360^\circ$
Radial position	$2.90 < r < 3.68$ m
Total number of detector modules	540 arranged in 18 super modules
Largest/smallest module	$113 \times 145$ cm <sup>2</sup> ( $91 \times 122$ cm <sup>2</sup> )
Active detector area	683 m <sup>2</sup>
Radiator	fibre/foam sandwich, 4.8 cm per layer
Radial detector thickness	$X/X_0 = 23.4\%$
Module segmentation	144 in $\phi$ 12-16 in $z$
Typical pad size	$0.7 \times 8.8$ cm <sup>2</sup>
Total number of pads	$1.18 \times 10^6$
Detector gas	Xe/CO <sub>2</sub> (85%/15%)
Gas volume	27.2 m <sup>3</sup>
Depth of drift region	3 cm
Depth of amplification region	0.7 cm
Drift field	0.7 kV/cm
Drift velocity	1.5 cm/ $\mu$ s
Longitudinal/transverse diffusion	$D_L = 250 \mu\text{m}/\sqrt{\text{cm}}$ $D_T = 180 \mu\text{m}/\sqrt{\text{cm}}$
Lorentz angle	$8^\circ$ at magnetic field 0.4 T
Number of readout channels	$1.18 \times 10^6$
Time samples in $r$ (drift)	20
ADC	10 bit, 10 MHz
Number of multi-chip modules	70848
Number of readout boards	4104
Pad occupancy for $dN_{ch}/d\eta = 8000$	34%
Pad occupancy in pp	$2 \times 10^{-4}$
Space-point resolution at 1 GeV/ $c$	in $r\phi$ $400(600)\mu\text{m}$ for $dN_{ch}/d\eta = 2000$ ( $dN_{ch}/d\eta = 8000$ )
	in $z$ 2 mm (offline)
Momentum resolution	$\delta p/p = 2.5\% \oplus 0.5\%(0.8\%)p/(\text{GeV}/c)$ for $dN_{ch}/d\eta = 2000$ ( $dN_{ch}/d\eta = 8000$ )
Pion suppression at 90% electron efficiency	better than 100 for $p = 1-6$ GeV/ $c$
Event size for $dN_{ch}/d\eta = 8000$	11 MB
Event size for pp	6 kB
Trigger rate limit	100 kHz



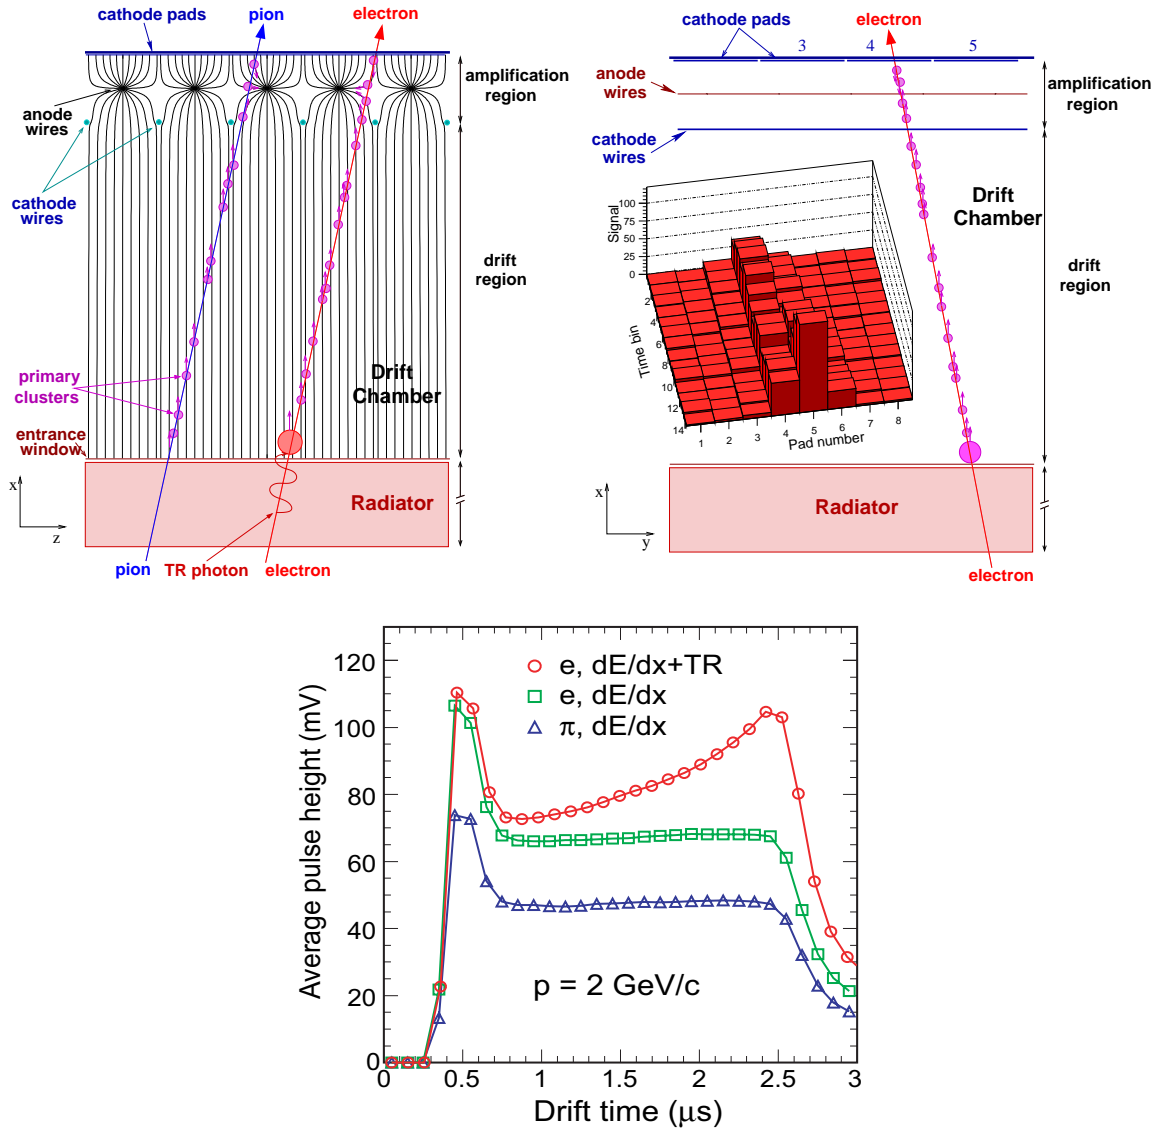
**Figure 3.40:** Left panel: schematic drawing of the TRD layout in the ALICE space frame. Shown are 18 super modules each containing 30 readout chambers (red) arranged in five stacks of six layers. One chamber has been displaced for clarity. On the outside the TRD is surrounded by the Time-Of-Flight (TOF) system (dark blue). On the inside the heat shield (yellow) towards the TPC is shown. Right panel: super module during assembly with the first three layers installed.

### 3.3.2 Detector layout

The final design of the TRD is depicted in figure 3.40. The TRD consists of 540 individual readout detector modules. They are arranged into 18 super modules (right panel of figure 3.40) each containing 30 modules arranged in five stacks along  $z$  and six layers in radius. In longitudinal ( $z$ ) direction the active length is 7 m, the overall length of the entire super module is 7.8 m, its total weight is 1650 kg.

Each detector element consists of a carbon fibre laminated Rohacell/polypropylene fibre sandwich radiator of 48 mm thickness, a drift section of 30 mm thickness, and a multi-wire proportional chamber section (7 mm) with pad readout. The pad planes are supported by a honeycomb carbon-fibre sandwich back panel (22 mm). While very light, the panel and the radiator provide enough mechanical rigidity of the chamber to cope with overpressure up to 1 mbar to ensure a deformation of less than 1 mm. The entire readout electronics is directly mounted on the back panel of the detector. Including the water cooling system the total thickness of a single detector layer is 125 mm. In the bending plane ( $r\phi$ ) each pad row consists of 144 pads. The central chambers consist of 12, all others of 16 pad rows. This leads to an overall channel count of  $1.18 \times 10^6$ . The total active area subtended by the pads is  $716 \text{ m}^2$ . The operating conditions of the detector are summarized in table 3.13.

Cross-sectional views of one TRD chamber together with average signals are shown in figure 3.41. Ionizing radiation produces electrons in the counting gas ( $\text{Xe}/\text{CO}_2$  (85:15)). Particles exceeding the threshold for transition radiation production ( $\gamma \approx 1000$ ) will in addition produce about 1.45 X-ray photons in the energy range of 1 to 30 keV. X-rays in this energy regime are efficiently converted by the high-Z counting gas with the largest conversion probability at the very beginning of the drift region. All electrons from ionization energy loss and X-ray conversions will drift towards the anode wires. After gas amplification in the vicinity of the anode wires the signal is induced on the readout pads. Signals of a typical track are shown in the inset of the central



**Figure 3.41:** Schematic cross-sectional view of a detector module in  $rz$ -direction (top left panel) and  $r\phi$ -direction (top right panel). The inset shows the charge deposit from an inclined track which is used for momentum reconstruction. The bottom panel shows the average pulse height as a function of drift time for pions (triangles), electrons without a radiator (squares) and electrons with a radiator (circles) for  $2 \text{ GeV}/c$  particles.

panel of figure 3.41. The inclination of the track in the bending direction is a direct measure of its transverse momentum. For particles with a momentum of  $2 \text{ GeV}/c$  the average amplitude of the cathode pad signal versus drift time is shown in the right panel. Two effects lead to an efficient discrimination between electrons and pions: i) the increased specific energy loss of electrons compared to pions at this momentum (difference between triangles and squares) and ii) the absorption of transition radiation generated by electrons (circles) predominantly at the beginning of the drift section corresponding to large drift times.

### 3.3.3 TRD performance

As outlined above the performance of the TRD is characterized by its particle identification capability and its tracking performance.


#### 3.3.3.1 Electron identification

The figure of merit for the TRD in terms of electron identification is its power to reject pions at a given electron efficiency (we choose 90%). This is usually expressed in terms of the pion efficiency (we aim at 1%). A parameterization of the measured amplitude spectra as a function of drift time and momentum both for electrons and pions provides the necessary likelihood distributions [100] that allow to evaluate the electron identification performance.


**Specific energy loss.** Measurements of the specific energy loss of electrons and pions of 2 GeV/ $c$  momentum in the Xe/CO<sub>2</sub> gas mixture are shown in figure 3.42 along with simulations. A correct understanding of the particle separation capability relies on a precise understanding of the details in the specific energy deposit of electrons and pions. Figure 3.42 demonstrates the level of agreement that has been achieved in the simulations. In that context it has been shown that a correct description of the escape probability of energetic  $\delta$ -electrons is needed to describe the tails of the energy distributions as well as their momentum dependence [101].

**Transition radiation.** The radiator is a composite structure using different inhomogeneous materials. The front and back face consist of 8 mm Rohacell foam covered with 0.1 mm carbon fibre laminate and 25  $\mu\text{m}$  aluminized mylar. It is filled with irregular polypropylene fibre mats (average fibre diameter 20  $\mu\text{m}$ ). The transition radiation production of this structure has been evaluated in prototype tests following [103]. The production of transition radiation sets in at an electron momentum of 0.5 GeV/ $c$  and levels off at about 2 GeV/ $c$  where on average 1.45 transition radiation photons are produced of which 1.25 are detected per incident electron [102]. The measured transition radiation energy spectra along with simulations are shown in figure 3.43.

**Pion rejection.** The resulting pion rejection capability was determined as a function of particle momentum. The results are shown in figure 3.44 for two different sets of data. One was obtained using six full-size TRD chambers (“big stack”), the other are extrapolations of results obtained with smaller prototype chambers (“small DCs”). Plotted is the percentage of pions erroneously identified as electrons versus momentum at an electron efficiency of 90%. The analysis was based on the likelihood of total deposited charge (L-Q). Along with the measurements, results from simulations are shown applying two-dimensional likelihoods on total deposited charge and position (L-Q1,Q2) as well as analysis based on neural networks (NN) [100]. These methods provide a factor 2–5 increase of the pion rejection capability compared to the simple likelihood analysis on total deposited charge.



**Figure 3.42:** Measurement of the specific energy deposit of 2 GeV/c pions (top) and electrons (bottom) in the nominal Xe,CO<sub>2</sub> gas mixture along with simulations [101].




**Figure 3.43:** Measurement and simulation of the transition radiation spectrum from 2 GeV/c electrons with the ALICE TRD radiator sandwich. The top panel shows the energy distribution of TR photons, the bottom panel the response per incident electron [102].

### 3.3.3.2 Tracking

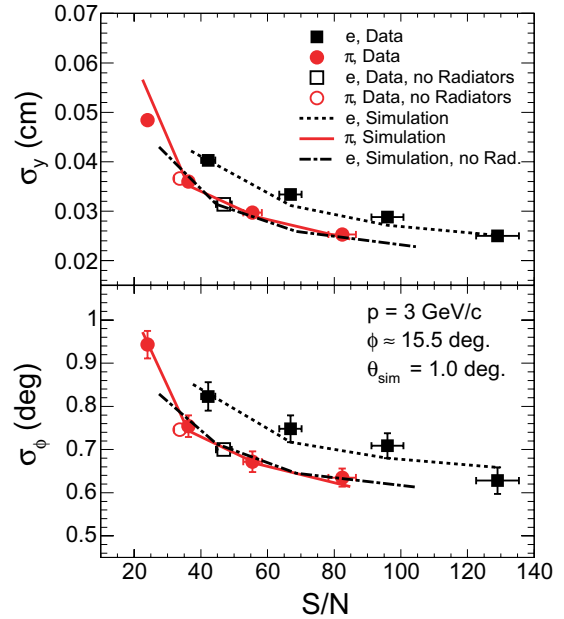
The chief tracking device in ALICE is the TPC. The needed tracking capability of the TRD is derived from the requirement of a reasonably sharp trigger threshold for individual particles in the range of up to  $p_t \approx 10 \text{ GeV}/c$  ( $\delta p_t/p_t \approx 5\%$  at 5 GeV/c). For particles with large momentum the tracking through the six layers of the TRD improves the overall momentum resolution of the tracking through the central barrel, see chapter 8.

**Position resolution.** The aforementioned tracking requirements in the nominal magnetic field of 0.5 T, require a position resolution for each time bin of  $\sigma_y \lesssim 400 \mu\text{m}$  and a resulting angular resolution per layer of  $\sigma_\phi \lesssim 1^\circ$ . The achievable respective resolutions were measured and are a function of the signal-to-noise ratio [104]. They are shown in figure 3.45. At a signal-to-noise ratio of about 40 the detector meets the requirements.

**Momentum resolution.** Using the above position and angular resolutions the stand-alone tracking resolution of the TRD can be derived for different momenta as a function of multiplicity density. For momenta below 2 GeV/c the stand-alone momentum resolution of the TRD is around  $\delta p_t/p_t \approx 2.5 - 3\%$  with little multiplicity dependence. Through the inclusion of the TRD into the



**Figure 3.44:** Measured pion efficiency as a function of beam momentum applying likelihood on total deposited charge (L-Q) (full symbols) measured with a stack of six chambers and smaller test chambers. Results are compared to simulations (open symbols) for 90% electron efficiency and six layers. These simulations were extended to two-dimensional likelihood on deposited charge and position (L-Q1,Q2) and neural networks (NN) [100].




**Figure 3.45:** Measurements and simulations of the position (top) and angular resolution as a function of the signal-to-noise ratio. Open (full) symbols refer to electrons and pions measured with (without) radiator in front of the drift chamber along with the respective simulations [104].

tracking in the central barrel an overall momentum resolution of better than 5% can be obtained up to momenta of about 100 GeV, see chapter 8.

### 3.3.4 Readout electronics

#### 3.3.4.1 General layout

An overview of the TRD readout electronics is shown in figure 3.46. The electronics, including the optical serializers ORI (two per chamber), is directly mounted on the backside of the detector modules. The data are transmitted to the Global Tracking Unit through 60 optical fibres per super module. The GTU either passes the data directly to the DAQ via the Detector Data Link (DDL) or it further analyzes the data in order to derive a fast Level-1 trigger decision. In that case individual tracklets from different layers of a stack are combined to determine the multiplicity of high- $p_t$  particles or to detect high-momentum  $e^+e^-$ -pairs. At Level-1 this trigger is transmitted to the Central Trigger Processor (CTP).



**Figure 3.46:** Schematic overview of the readout electronics of the TRD.

### 3.3.4.2 Detector electronics

Most of the on-detector readout electronics are realized as ASICs. Two such chips, an 18-channel full-custom preamplifier shaper (PASA), and a so-called Tracklet Processor (TRAP) were integrated into a multi-chip module (MCM). The PASA [105] is a folded cascode with differential output (120 ns shaping time,  $850e$  ENC for an input capacitance of 25 pF, 12.4 mV/fC gain, and 12 mW/channel power consumption). It was realized using the AMS 0.35  $\mu\text{m}$  CMOS process. Each PASA has 18 input and 21 output channels. For reconstruction of tracklets the extra output channels are fed into the analogue inputs of the TRAPs on neighbouring MCMs to allow for continuous charge sharing across MCM boundaries.

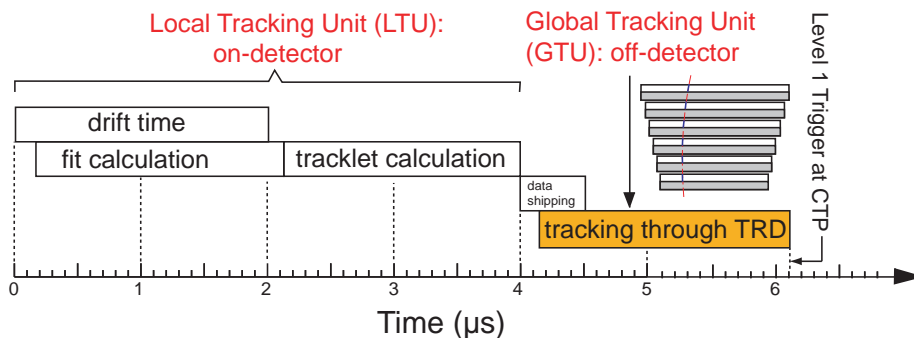
The Tracklet Processor is a mixed signal design realized in the UMC 0.18  $\mu\text{m}$  process. It comprises 21 ADCs, digital filters [106], event buffers, and processing units [107] that allow to calculate the inclination of a track in the bending direction as well as the total charge deposited along the track (Local Tracking Unit - LTU). This feature allows to identify high- $p_t$  particles at the trigger level. Evaluation of the deposited energy will furthermore allow to tag possible electron candidates at the trigger level.

### 3.3.4.3 Global Tracking Unit (GTU)

After preprocessing of the analogue signals as described above, the resulting track segments from the different detector layers have to be matched in three dimensions for transverse momentum reconstruction. Based on the data of all 1.2 million analogue channels, the reconstruction has to be performed within 6.1  $\mu\text{s}$  to derive the Level-1 trigger decision. The entire trigger timing sequence involving the LTU and the GTU is depicted in figure 3.47.

The massively parallel hardware architecture of the GTU is capable of processing up to 20 000 track segments within 2  $\mu\text{s}$ . The core of the GTU, along with a custom bus system, is the so-called Track Matching Unit (TMU). It is an FPGA-based system utilizing PCI and 12 fibre-optical SFP transceiver interfaces gathering the data from a stack of six chambers. It is realized as a CompactPCI plug-in card. The main FPGA is a Xilinx Virtex-4 FX chip which includes an integrated multi-gigabit serializer/ deserializer and PowerPC processor blocks.





**Figure 3.47:** Trigger timing diagram for the generation of a high- $p_t$  trigger.

### 3.3.4.4 Detector Control System (DCS)

The entire configuration, monitoring, and control of the detector electronics is realized on a single DCS board with an autonomous power supply. There is one DCS board per chamber. The boards comprise an Altera FPGA with ARM9 core, flash memory, Ethernet-, JTAG-, and I<sup>2</sup>C-interfaces, 8 analogue and 8 digital inputs as well as 8 digital outputs.

The board reshapes and distributes the central clock signal received via an on-board Trigger Timing Controller (TTC) to all MCMs on the detector. All internal clocks of the MCMs are derived from this central clock. The configuration of different filter and gain parameters as well as zero suppression conditions used for the reconstruction of the tracklets within the Local Tracking Unit (LTU) are handled by the DCS board. It also monitors the temperatures and currents of all individual MCMs. Through its external control lines it is able to shutdown voltage regulators in case of over-temperature or over-current. Moreover, the DCS board can reroute the so-called Slow Control Serial Network in the configuration chain which allows us to bypass faulty MCMs without loss of data from entire readout boards.

The states of all services (low voltage, high voltage, gas, cooling) as well as all states generated or monitored by the DCS board are implemented as Finite-State Machines (FSM) which are operated and controlled via PVSS.

## 3.4 Time-Of-Flight (TOF) detector

### 3.4.1 Design considerations

The Time-Of-Flight (TOF) detector [15] is a large area array that covers the central pseudo-rapidity region ( $|\eta| \lesssim 0.9$ ) for Particle IDentification (PID) in the intermediate momentum range, below about 2.5 GeV/ $c$  for pions and kaons, up to 4 GeV/ $c$  for protons, with a  $\pi/K$  and  $K/p$  separation better than  $3\sigma$  [21, 108]. The TOF, coupled with the ITS and TPC for track and vertex reconstruction and for  $dE/dx$  measurements in the low-momentum range (up to about 1 GeV/ $c$ ), will provide event-by-event identification of large samples of pions, kaons, and protons. In addition, at the inclusive level, identified kaons will allow invariant mass studies, in particular the detection of open heavy-flavoured states and vector-meson resonances such as the  $\phi$  meson. Details on the physics

observables that can be addressed and the expected performance are detailed in [4, 21, 108, 109] and in section 8 of this document.

A large-coverage TOF detector, operating efficiently, with an excellent intrinsic response and an overall occupancy not exceeding the 10–15% level at the highest predicted charged-particle density of  $dN/d\eta = 8000$  was required; see [21]. This led to the current design with more than  $10^5$  independent TOF channels. Since a large area had to be covered, a gaseous detector was chosen. In the framework of the LAA project at CERN an intensive R&D programme has shown that the best solution for the TOF detector was the Multi-gap Resistive-Plate Chamber (MRPC) [110, 111].

The key aspect of these chambers is that the electric field is high and uniform over the full sensitive gaseous volume of the detector. Any ionisation produced by a traversing charged particle immediately starts a gas avalanche process which generates the observed signals on the pick-up electrodes. Unlike other types of gaseous detectors, there is no drift time associated with the movement of the electrons to a region of high electric field. Thus the time jitter of these devices is caused by the fluctuations in the growth of the avalanche.

The main advantages of the MRPC technology [112] with respect to other parallel-plate chamber designs are that: i) the chamber operates at atmospheric pressure; ii) the signal is the analogue sum of signals from many gaps, so there is no late tail and the charge spectrum is not of an exponential shape — it has a peak well separated from zero; iii) the resistive plates quench the streamers so there are no sparks, thus high-gain operation becomes possible; iv) the construction technique is in general rather simple and makes use of commercially available materials.

The final tests of several MRPC multicell strips from mass production confirmed that these devices indeed reach an intrinsic time resolution better than about 40 ps and an efficiency close to 100%.

### 3.4.2 Detector layout

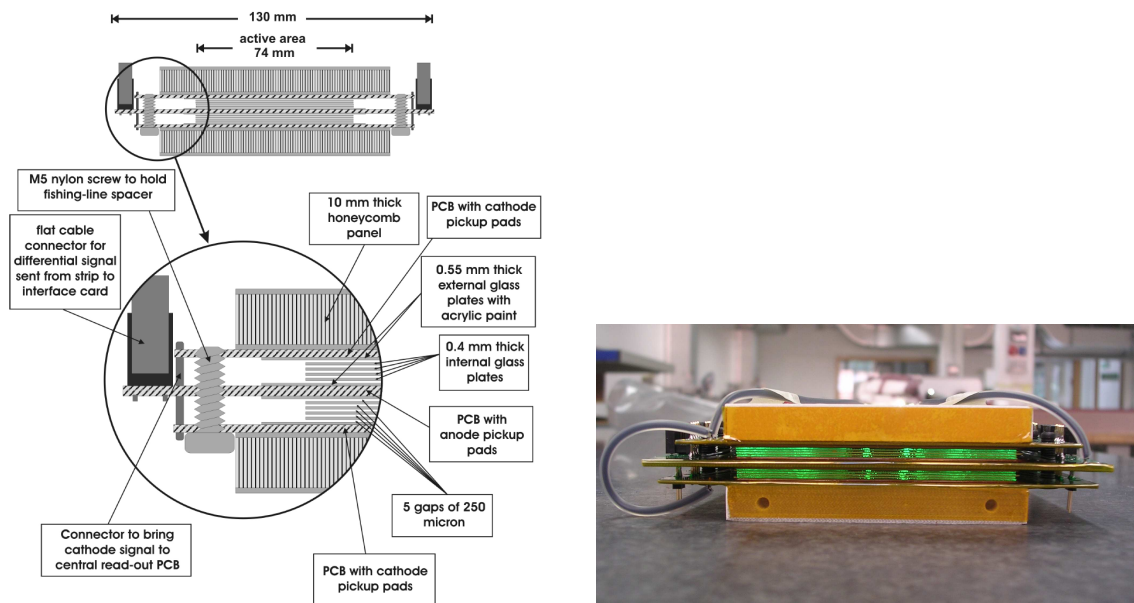
The detector covers a cylindrical surface of polar acceptance  $|\theta - 90^\circ| < 45^\circ$ . It has a modular structure corresponding to 18 sectors in  $\varphi$  and to 5 segments in  $z$  direction.

The whole device is inscribed in a cylindrical shell with an internal radius of 370 cm and an external one of 399 cm. The whole device thickness corresponds to 30% of a radiation length.


The basic unit of the TOF system is a 10-gap double-stack MRPC strip (figure 3.48) 122 cm long and 13 cm wide, with an active area of  $120 \times 7.4 \text{ cm}^2$  subdivided into two rows of 48 pads of  $3.5 \times 2.5 \text{ cm}^2$ . An overall view of a strip is shown in figure 3.49. The strips are placed inside gas-tight modules (which also act as Faraday cages) and are positioned transversely to the beam direction. Five modules of three different types are needed to cover the full cylinder along the  $z$  direction.

They all have the same structure and width (128 cm) but differ in length. The actual dimensions and overall design are defined in such a way that the joining areas of the modules are aligned with the dead areas of the other detectors projected from the interaction point, thus creating a configuration of minimal disturbance for the external detectors. The length of the central module is 117 cm, the intermediate ones 137 cm, and the external ones 177 cm. The overall TOF barrel length is 741 cm (active region).

The granularity of the TOF detector is dictated by the requirement of identifying, on an event-by-event basis, as many charged particles as possible, even at the maximum predicted charged-particle multiplicity density of  $dN_{\text{ch}}/d\eta = 8000$ . Detailed simulation studies have shown that, with



**Figure 3.48:** Schematic cross section (left) and photograph (right) of a 10-gap double-stack MRPC strip.




**Figure 3.49:** Photograph of an MRPC strip during (left and centre) and after (right) construction.

the chosen pad size of  $3.5 \times 2.5 \text{ cm}^2$  and the tilted-strip geometry, the occupancy of the detector is  $\sim 14\%$ , at the highest charged-particle density (including secondary particles) with a magnetic field of 0.5 T. Lower values of particle density result in linearly lower occupancy.


In order to minimise the transversal path of the incident particles through the chamber strips, a special tilted positioning of the strips was devised. This reduces the number of very oblique transversal paths that can create a sharing effect of the signal among adjacent pads, thereby increasing the occupancy and the time jitter of the detected signals. The angle of the strips with respect to the axis of the cylinder is progressively increasing from  $0^\circ$  in the central part ( $\theta = 90^\circ$ ) of the detector to  $45^\circ$  in the extreme part of the external module ( $\theta = 45^\circ$ ). This arrangement makes the median zone of a strip perpendicular to a radius coming from the interaction point, thus minimising the angle of the incoming particles with respect to the normal direction. In order to avoid dead areas, adjacent strips were overlapped inside the modules so that the edge of the active area of one pad is aligned with the edge of the next one. This gives the possibility of creating a full active area with no geometrical dead zones. The modules were designed in such a way as to avoid any loss of the sensitive area, along the  $z$ -axis. The only dead area is due to the unavoidable presence of the supporting spaceframe structure.

Every module of the TOF detector consists of a group of MRPC strips (15 in the central, 19 in the intermediate and in the external modules) closed inside a box that defines and seals the gas volume and supports the external front-end electronics and services. An aluminium-honeycomb plate is the backbone of the module and gives the necessary mechanical stiffness to the system. The honeycomb plate is 1.5 cm thick including two aluminium skins, each 0.1 cm thick. A cover connected to the backbone by means of screws and standard sealing O-Ring closes the total gas volume; it has to withstand the overpressure required for gas circulation inside the chamber (less than 3 mbar). The cover is made of fibreglass 0.3 cm thick, reinforced with moulded ribs, protruded by 1.8 cm; there are 6 ribs in the central module, 9 in the intermediate and 8 in the external ones. A 0.04 cm fibreglass web, covered by a  $20\ \mu\text{m}$  aluminium layer, is glued to the inside surface of the cover in order to create an electromagnetic shielding. Inside the gas volume, fixed perpendicularly to the honeycomb plate, there are two macrolon plates, 0.5 cm thick, to which the MRPC strips are attached. The system turns out to be very simple and allows for a fast insertion of the strips at whatever angle is needed. Holes that accommodate feed-through for the signal cables coming from the readout pads, the HV connectors and the gas inlet and outlet, are machined into the honeycomb plate. The signal feed-through is made of a PCB having on one side connectors receiving the cables coming from the strips inside the gas volume and, on the other side, connectors that accommodate the front-end electronic cards, see figure 3.50. The PCB is glued to the honeycomb plate via special fibreglass gaskets which are inserted into the holes. The volume containing the electronic cards, input and output cables, gas pipes, water cooling pipes and radiators is closed by a 0.15 cm aluminium cover. This allows access to the electronics with no disturbance to the active part of a module.

The complete TOF system consists of 90 modules. Three central modules in front of PHOS have not been installed in order to reduce the amount of material for this high resolution electromagnetic calorimeter in three of its five modules. The five modules in a row are located inside a supermodule framework of longitudinal and transverse aluminium beams for each of the 18 sectors. Pairs of horizontal and vertical rollers are attached to the supermodule body permitting the insertion of the supermodule into the spaceframe sector where two rails are fixed. Each supermodule has an overall length of  $\sim 930$  cm and a weight of  $\sim 1.4$  t. A schematic layout of one supermodule inside the ALICE spaceframe and a photo of an assembled supermodule are shown in figure 3.51. An overview of the TOF detector parameters is shown in table 3.14.




**Figure 3.50:** Photographs of a TOF module during assembly. The strips are installed inside the gas volume (upper) and the FEA cards are plugged onto the PCB connectors (lower).



**Figure 3.51:** Schematic drawing of one TOF supermodule, consisting of 5 modules, in the ALICE spaceframe (left panel). A supermodule after assembly; in the foreground, the two custom crates with readout electronics (right panel).

**Table 3.14:** Overview of TOF parameters.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$
Azimuthal coverage	$360^\circ$
Radial position	$370 < r < 399 \text{ cm}$
Length (active region)	741 cm
Segmentation in $\varphi$	18-fold
Segmentation in $z$	5-fold
Total number of modules	90
Central module (A)	$117 \times 128 \text{ cm}^2$
Intermediate module (B)	$137 \times 128 \text{ cm}^2$
External module (C)	$177 \times 128 \text{ cm}^2$
Detector active area (total area)	$141 \text{ m}^2$ ( $171 \text{ m}^2$ )
Detector radial thickness	$X/X_0 = 29.5\%$
Number of MRPC strips per module	15 (A), 19 (B), 19 (C)
Number of readout pads per MRPC strip	96
Module segmentation in $\varphi$	48 pads
Module segmentation in $z$	30 (A), 38 (B), 38 (C) pads
Readout pad geometry	$3.5 \times 2.5 \text{ cm}^2$
Total number of MRPC strips	1638
Total number of readout pads	157 248
Detector gas	$\text{C}_2\text{H}_2\text{F}_4(90\%), \text{i-C}_4\text{H}_{10}(5\%), \text{SF}_6(5\%)$
Gas volume	$17.5 \text{ m}^3$
Total flow rate	$1 \text{ m}^3/\text{h}$
Working overpressure	$< 3 \text{ mbar}$
Fresh gas flow rate	$0.02 \text{ m}^3/\text{h}$
Number of readout channels	157 248
Number of front-end analogue chips (8-ch)	19 656
Number of front-end boards	6 552
Number of HPTDC chips (8-ch, 24.4 ps bin width)	20 520
Number of HPTDC readout boards (TRM)	684
Number of readout boards (DRM) and crates	72
Occupancy for $dN_{\text{ch}}/d\eta = 8000$	14% ( $B = 0.5 \text{ T}$ )
Occupancy for pp	$10^{-4}$ ( $B = 0.5 \text{ T}$ )
$\pi, \text{K}$ identification range (with contamination $< 20\%$ )	$0.5\text{--}3.0 \text{ GeV}/c$
p identification range (with contamination $< 15\%$ )	$0.5\text{--}6.0 \text{ GeV}/c$
e identification range in pp (with contamination $< 10\%$ )	$0.3\text{--}0.5 \text{ GeV}/c$
Event size for $dN_{\text{ch}}/d\eta = 8000$	110 kB
Event size for pp	24 kB




**Figure 3.52:** Efficiency and time resolution as a function of high voltage for a sample of mass-production MRPC strips. The operating voltage is 13 kV.

Results obtained at the CERN/PS T10 test-beam with ten MRPC strips, chosen randomly from two years of mass-production, and with the full electronic chain are shown in figure 3.52 and figure 3.53 for the efficiency and the time resolution. Similar results were obtained over a sample of several tens of pre-production strips. This demonstrates a detector of full efficiency and an intrinsic time resolution better than 40 ps.

The contribution of the full electronic chain (FEA, TRM, CPDM: see next section) to the overall time resolution of the detector has been measured to be about 38 ps.

Extensive tests [113] were performed at the CERN PS and Gamma Irradiation Facility (GIF) to study different gas mixtures, ageing effects and rate capability. The gas mixture was composed of  $C_2H_2F_4$  (90%),  $i-C_4H_{10}$  (5%),  $SF_6$  (5%). No ageing effect was observed after 3.5 times the dose foreseen for the ‘standard running scenario’ in the first 10 years of LHC operation, and a rate capability far in excess of the  $50 \text{ Hz/cm}^2$  maximum expected rate at the ALICE experiment was demonstrated.

Details on the TOF detector mass-production procedures and results are available in [114].



**Figure 3.53:** Efficiency and time resolution for a sample of mass-production MRPC strips measured in several positions (pads) along and across each strip, at fixed high voltage.

### 3.4.3 Front-end electronics and readout

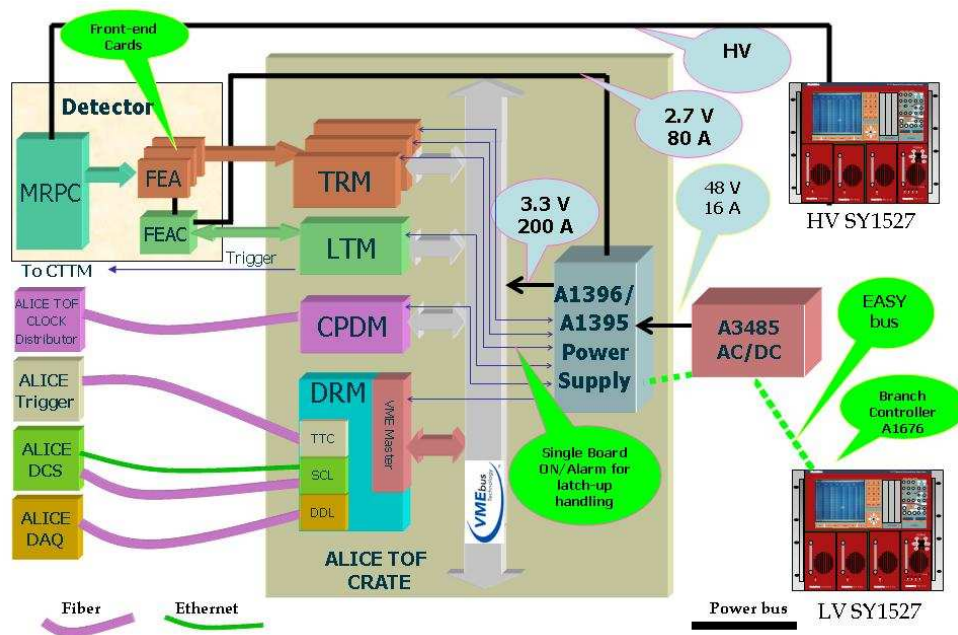
An overview of the TOF electronics is shown in figure 3.54. The front-end electronics for the TOF were designed to comply with the basic characteristics of the MRPC detector, i.e. very fast differential signals from the anode and cathode readout pads and intrinsic time resolution better than 40 ps.

During the R&D phase of the detector a solution based on commercial chips (a very fast non-differential amplifier and a comparator) was adopted. In 2002 an ASIC chip (‘NINO’) in CMOS 0.25  $\mu\text{m}$  technology was developed and tested at PS in 2003 with excellent results [115], thus allowing a substantial reduction of the power dissipation (now at about 40 mW/channel). All the stages (input amplifier and comparator stages) of the 8-channel ASIC are fully differential, its input impedance and capacitance are matched to the transmission line and the peaking time is smaller than 1 ns.

The basic Front-End Analogue card (FEA) contains 3 ASIC chips (24 channels) with a common threshold regulation; it is connected to the PCB interface cards on top of the honeycomb support plate of a module. The LVDS output signals, routed along the two sides of the supermodule, carry the information of the hit time (leading edge) and of the Time-Over-Threshold (TOT), related to the input charge, that is needed for the time-slewing correction.

The readout electronics, located in custom crates at both ends of a supermodule, consists of 9 (right side crate) or 10 (left side crate) TRM (TDC Readout Module) cards and one DRM (Data Readout Module) card in each crate. The TRM card houses the HPTDC [116] (High Performance TDC) 8-channel chips that are used in the very high resolution mode (24.4 ps bin width). Each






**Figure 3.54:** Schematic overview of the TOF electronics.

TRM card contains 30 HPTDC chips, i.e. 240 channels, corresponding to the readout pads of 2.5 MRPC strips. The DRM card is the TOF interface to the ALICE DAQ system, it reads and encodes the data from the TRM cards and sends them to the DAQ via the DDL optical link. The DRM card receives the trigger information (Level 1, Level 2-accept, Level 2-reject) from the CTP (Central Trigger Processor) via the TTCrx (Timing Trigger and Control receiver) chip and performs a slow-control function with a dedicated FPGA.

Two other electronic cards are located in the custom crates. The CPDM (Clock and Pulser Distribution Module) card distributes to the TRM, DRM and LTM (see below) cards a clock signal with a very low jitter (17 ps) and to the MRPC strips a pulser signal for calibration and monitoring purposes.

The LTM (Local Trigger Module) card, besides setting and monitoring the thresholds of the FEA ASICs, monitoring the FEA low voltage and the temperature inside the electronics volume of a module, collects the 48 OR signals from each pair of FEAs. The purpose of these signals is twofold: it sends a fast pretrigger to the Transition Radiation Detector (TRD) and provides a Level 0 (L0) trigger to select Ultra Peripheral Collisions (UPC) (see section 6.10.3 of [21]), Minimum Bias (MB) events in pp collisions, and Cosmic muons both for central detectors calibration and for cosmic-ray physics (see section 6.11.5 of [21]).

A large electronics board, the CTM (Cosmic and Topology Trigger Module), located under the muon spectrometer platform, receives the LTM signals, asserts the L0 and L1 triggers and sends them to the CTP.



**Figure 3.55:** View of the seven modules of the HMPID mounted on the cradle.

## 3.5 High-Momentum Particle Identification Detector (HMPID)


### 3.5.1 Design considerations

The High-Momentum Particle Identification Detector (HMPID) [8], is dedicated to inclusive measurements of identified hadrons at  $p_t > 1 \text{ GeV}/c$ . The aim is to enhance the PID capability of ALICE by enabling identification of charged hadrons beyond the momentum interval attainable through energy-loss (in ITS and TPC) and time-of-flight measurements (in TOF). The detector was optimised to extend the useful range for  $\pi/K$  and  $K/p$  discrimination, on a track-by-track basis, up to  $3 \text{ GeV}/c$  and  $5 \text{ GeV}/c$ , respectively, see chapter 8. The HMPID was designed as a single-arm array with an acceptance of 5% of the central barrel phase space. The geometry of the detector was optimized with respect to the particle yields in pp and heavy-ion collisions at LHC energies, and with respect to the large opening angle required for two-particle correlation measurements, see section 6.3 of [21]. In addition the identification of light nuclei and anti-nuclei (d, t,  $^3\text{He}$ ,  $\alpha$ ) at high transverse momenta in the central rapidity region can also be performed with the HMPID [117].

### 3.5.2 Detector layout

The HMPID is based on proximity-focusing Ring Imaging Cherenkov (RICH) counters and consists of seven modules of about  $1.5 \times 1.5 \text{ m}^2$  each, mounted in an independent support cradle (figure 3.55) [118]. The cradle is fixed to the spaceframe at the two o'clock position.

The radiator, which defines the momentum range covered by the HMPID, is a 15 mm thick layer of low chromaticity  $\text{C}_6\text{F}_{14}$  (perfluorohexane) liquid with an index of refraction of  $n = 1.2989$  at  $\lambda = 175 \text{ nm}$  corresponding to  $\beta_{\min} = 0.77$  (i.e. a momentum threshold  $p_{\text{th}} = 1.21 m$ , where  $m$  is the particle mass). Cherenkov photons, emitted by a fast charged particle traversing the radiator,



**Figure 3.56:** Working principle of a RICH detector employing CsI thin films deposited onto the cathode plane of a MWPC. The Cherenkov cone refracts out of the liquid radiator of  $C_6F_{14}$  and expands in the proximity volume of  $CH_4$  before reaching the MWPC photon detector. Electrons released by ionising particles in the proximity gap are prevented to enter the MWPC volume by a positive polarisation of the collection electrode close to the radiator.

are detected by a photon counter (figure 3.56) which exploits the novel technology of a thin layer of CsI deposited onto the pad cathode of a Multi-Wire Pad Chamber (MWPC). The HMPID, with a surface of about  $11 \text{ m}^2$ , is the largest scale application of this technique [119].

A relevant feature of the adopted design for the HMPID photo-detector is the open geometry, which suppresses the electrode elements that in TMAE photo-detectors were specifically implemented to prevent spurious avalanches from feedback photons. Each HMPID module consists of a stack of five independent frames made out of a stabilised aluminium alloy (peralumin) which minimises out-gassing into the chamber active volume. The photo-detector, filled with pure methane and operated at ambient temperature and pressure, is closed on one side by an end-flange which supports six independent CsI photocathode panels,  $64 \times 40 \text{ cm}^2$  large, segmented into pads of  $8 \times 8.4 \text{ mm}^2$ . On the opposite side, a honeycomb panel supports three  $C_6F_{14}$  radiator vessels placed at a distance of 80 mm from the anode wire plane. The gas tightness is ensured by soft O-rings placed in grooves in the chamber frames.

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

Cathode and collection wire planes are made out of  $100\ \mu\text{m}$  diameter gold-plated Cu/Be wires stretched at 200 g per wire by using crimping pins. Stiffness and flatness of pad cathodes is obtained by gluing together two multilayer printed circuit boards on an aluminium frame. The photo-converter is a 300 nm thick layer of CsI, evaporated by the Joule effect. During the photo-cathode production the quality cut was set such to allow for the detection of at least 16 resolved clusters per ring (for  $\beta = 1$ ). The performance of the selected 42 photocathodes allows, on average, for the detection of more than 20 resolved clusters. The production technique and quality evaluation of the 42 photo-cathodes equipping the seven HMPID modules are fully described in [120].

The liquid radiator containers consist of trays of  $1330 \times 413\ \text{mm}^2$  made out of glass-ceramic material (NEOCERAM), thermally compatible (thermal coefficient  $0.5 \times 10^{-6}\ \text{°C}^{-1}$ ) with the fused silica plates used as UV-transparent windows. The thickness and size of the tray's elements were carefully optimised by investigating the best compromise between the detector total radiation length and the perfluorohexane hydrostatic pressure: the quartz window is 5 mm thick, while the NEOCERAM base plate is 4 mm thick. To withstand the hydrostatic pressure, thirty cylindrical spacers are glued on the NEOCERAM bottom plate on one side and the quartz window on the other side. They consist of fused silica rods, 10 mm diameter, arranged in three rows of 10 equally spaced elements.

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

The radiator is 15 mm thick, a choice based on detailed MC simulations and test beam results [121]. Thus it is possible to operate the detector at a low gain (less than  $10^5$ ) under stable conditions of operation while achieving good performance (see refs. [122, 123]). The high photon yield coupled with the improved quantum efficiency of the photo-cathodes will be used to tune the detector gain adapting it to the specific running conditions. A low gain can be used in ion-ion collisions, minimising the photon feedbacks and MIPs contribution to the occupancy, while in proton-proton runs, where the much lower track density eases the pattern recognition, higher gain can be used to improve the efficiency. A positive voltage of 2050 V, applied to the anodes, while cathodes are grounded, provides a total 'nominal' gas gain of  $4 \times 10^4$ .

A liquid circulation system was implemented to purify  $\text{C}_6\text{F}_{14}$ , fill and empty the twenty-one radiator trays at a constant flow, independently, remotely and safely. Due to the inaccessibility of the detector during the run and the fragility of the radiator trays, a system based on a gravity flow principle was chosen owing to its intrinsically safe nature. Since  $\text{C}_6\text{F}_{14}$  is not available in a high-purity grade form, filters are implemented in the circulation system in order to remove contaminants (mainly water and oxygen) and achieve the best transparency in the UV region where the detector operates. The full description of the liquid system can be found in [122]. The total HMPID gas volume is  $1.4\ \text{m}^3$  of pure  $\text{CH}_4$ . It is split equally into the seven independent RICH modules by supplying each one of them individually with a gas flow rate such that five volume changes are obtained per day while maintaining an operating pressure of a few mbar above ambient pressure. When not operational, for instance during machine shutdown, the detector is flushed with argon. During the detector pre-commissioning prior to installation in the experimental cavern, a leak was detected in one (out of 21) of the radiator trays. Due to the little impact on the physics potential of

**Table 3.15:** Synopsis of HMPID parameters.

Pseudo-rapidity coverage	$-0.6 < \eta < 0.6$
Azimuthal coverage	$1.2^\circ < \phi < 58.8^\circ$
Radial position	5 m
Total number of modules	7
Detector active area	10.7 m <sup>2</sup>
Detector thickness radially	$X/X_0 = 18\%$
Radiator thickness	15 mm
Radiator medium	liquid C <sub>6</sub> F <sub>14</sub>
Refractive index	1.2989 at 175 nm
Independent radiator trays	3 per module (21 total)
$\beta$ threshold	0.77
Detector gas	CH <sub>4</sub>
Gas volume	1.4 m <sup>3</sup>
MWPC anode-cathode gap	2 mm
Operating voltage	2050 V
Photon converter	Caesium Iodide (CsI)
CsI thickness	300 nm
Quantum Efficiency (QE)	> 25% at 175 nm
Number of Photo-Cathodes (PC)	42
Number of pads/PC	3840 (161 280 total)
Pad size	$8.0 \times 8.4 = 67.2 \text{ mm}^2$
Front-end chip (GASSIPLEX-07-3)	10 080
Peaking Time	1.2 $\mu$ s
Readout chip (DILOGIC-3)	3360
Front-end cards	3360
Readout cards	672
ADC (12-bit)	3360
Power consumption/module	450 W (total 3.15 kW)
Number of DDL	14
Multiplexing frequency	10 MHz (maximum)
Multiplexing time	5 $\mu$ s at 10 MHz
Readout time (for 12% occupancy)	< 300 $\mu$ s
Event size (for 12% occupancy)	< 0.1 MB
Number of readout channels	161 280
Occupancy ( $dN_{\text{ch}}/d\eta = 8000$ )	12%

the detector (5% loss in the acceptance) and the non-negligible risks, including delays, involved in the reparation, it was decided to install the detector inside the ALICE solenoid with 20 operational radiators while flushing with methane the leaking one.

The parameters of the detector are summarised in table 3.15.

### 3.5.3 Front-end electronics and readout

The front-end electronics is based on two dedicated ASIC chips, GASSIPLEX [124] and DILOGIC [125], successfully developed in the framework of the HMPID project in the ALCATEL-MIETEC 0.7  $\mu\text{m}$  technology. The GASSIPLEX-07 chip is a 16-channel analogue multiplexed low-noise signal processor working in TRACK&HOLD mode. It features a dedicated filter to compensate for the long ion drift tail, a semi-Gaussian shaper and internal protection against discharges. The noise on detector is found to be  $1000e$  r.m.s. During the commissioning before installation the number of dead or noisy pads was found to be less than 200 out of 161 280 pads. In the HMPID application the GASSIPLEX analogue output is presented to the input of a commercial 12-bit ADC (AD9220ARS). The multiplexing level is 3 chips (48 channels) per ADC at a maximum frequency of 10 MHz. The digitization and zero-suppression time does not depend on the occupancy, being determined only by the multiplexing rate. The highest foreseen luminosity ( $5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ , interaction rate of about 200 kHz) in proton-proton interactions is acceptable for the HMPID since the front-end electronics has a baseline recovery better than 0.5% after 5  $\mu\text{s}$ . Moreover, the low multiplicity per event makes event overlap negligible.

The DILOGIC chip is a sparse data scan readout processor providing zero suppression and pedestal subtraction with individual threshold and pedestal values for up to 64 channels. Several chips are daisy-chained on the same 18-bit output bus. Asynchronous read/write operations are allowed. Data are read out via the standard ALICE Detector Data Link (DDL). The ALICE DDL Source Interface Unit (SIU) is housed, together with the TTCrq mezzanine board, in the Readout and Control Board (RCB) which interfaces the detector readout electronics with the DAQ and Trigger subsystems. Two RCBs boards are required to operate one HMPID module. The readout time after L2 arrival (for 12% occupancy) is of the order of 300  $\mu\text{s}$ . Since momentum information is vital to exploit the HMPID detector, only events for which the TPC information is available are of interest and the HMPID can perfectly cope with the readout rates foreseen for the TPC.

## 3.6 PHOton Spectrometer (PHOS)

The PHOton Spectrometer (PHOS, [9, 20]) is a high-resolution electromagnetic spectrometer covering a limited acceptance domain at central rapidity. The main physics objectives are the test of thermal and dynamical properties of the initial phase of the collision extracted from low  $p_t$  direct photon measurements and the study of jet quenching through the measurement of high- $p_t$   $\pi^0$  and  $\gamma$ -jet correlations.

### 3.6.1 Design considerations

The high particle multiplicity in nuclear collisions requires a dense, highly segmented calorimeter with small Molière radius at a large distance from the interaction point in order to keep the cell occupancy at a manageable level of about 10 – 20%. A good energy and position resolution improves the signal to background ratio for meson identification, in particular at low  $p_t$  where the combinatorial background is very large. The identification of photons requires high discrimination power against charged hadrons, neutrons and anti-neutrons. Topological analysis of the shower

development in the electromagnetic calorimeter, time of flight measurement and charged particle identification provide the discriminating criteria (see [21], section 5 and [126]). The required performance is met through a high granularity electromagnetic calorimeter, timing resolution of the individual detector cells of the order of a few ns and a charged particle detector in front of the calorimeter. Direct photons are discriminated against decay photons either through shower shape analysis at high  $p_t$  or through invariant mass analysis at low  $p_{Tmt}$ .

The high-energy resolution and granularity is provided by using dense scintillator material (lead-tungstate,  $\text{PbWO}_4$ ) of  $20 X_0$  with high photo-electron yield. The energy resolution of PHOS is measured to be [127]:

$$\frac{\sigma_E}{E[\text{GeV}]} = \sqrt{\left(\frac{0.0130 \pm 0.0007}{E[\text{GeV}]}\right)^2 + \left(\frac{0.0130 \pm 0.0007}{\sqrt{E[\text{GeV}]}}\right)^2 + (1.12 \pm 0.3)^2}. \quad (3.1)$$

The high spatial resolution is achieved by choosing the dimension of the individual detection cells of the order of the Molière radius of the scintillator. The position resolution of PHOS is measured to be [9]:

$$\sigma_{x,y}[\text{mm}] = \sqrt{\left(\frac{3.26}{\sqrt{E[\text{GeV}]}}\right)^2 + 0.44^2}. \quad (3.2)$$

The resulting two-photon invariant mass resolution at the  $\pi^0$  peak is 3.5%.

The required timing resolution is achieved by use of fast scintillator and preamplifier. The timing resolution which can be reached with PHOS is about 2 ns at energies above 1.5 GeV.

The large dynamic range (0.005–80 GeV) is achieved by selecting an appropriate detector length that minimizes shower leakage for the highest particle-energies without deteriorating the energy resolution for the lowest particle energies due to light attenuation.

### 3.6.2 Detector layout

PHOS is designed as a single-arm high-resolution high-granularity electromagnetic spectrometer consisting of a highly segmented electromagnetic calorimeter (PHOS) and a Charged-Particle Veto (CPV) detector. A synopsis of the parameters of PHOS is given in table 3.16. PHOS is subdivided into five independent PHOS+CPV units, called PHOS modules. It is positioned on the bottom of the ALICE setup at a distance of 460 cm from the interaction point. After its final installation it will cover approximately a quarter of a unit in pseudo-rapidity,  $0.12 \leq \eta \leq 0.12$ , and  $100^\circ$  in azimuthal angle.

#### 3.6.2.1 Electromagnetic calorimeter (PHOS)

Each PHOS module is segmented into 3584 detection cells arranged in 56 rows of 64 cells. The detection cell consists of a  $22 \times 22 \times 180 \text{ mm}^3$  lead-tungstate crystal,  $\text{PbWO}_4$  (PWO), coupled to a  $5 \times 5 \text{ mm}^2$  Avalanche Photo-Diode (APD) followed by a low-noise preamplifier [128]. The APD and the preamplifier are integrated in a common body glued onto the end face of the crystal (figure 3.57). The crystals are assembled in strip units of two rows of eight detection cells (figure 3.57). The 16 analog signals from one strip unit are fed into a T-shaped connector, connected to the shaper/digitizer/trigger electronics.

**Table 3.16:** Summary of PHOS parameters.

Coverage in pseudo-rapidity	$-0.12 \leq \eta \leq 0.12$
Coverage in azimuthal angle	$220^\circ < \phi < 320^\circ$
Distance to interaction point	4600 mm
Modularity	Five modules of $1.2 \times 1.4 \text{ m}^2$
<b>PHOS</b>	
Material	Lead-tungstate crystals (PWO) Crystal
Crystal dimensions	$22 \times 22 \times 180 \text{ mm}^3$
Depth in radiation length	20
Number of crystals	17 920, 3584 crystals per module
Total area, crystal volume	$8.6 \text{ m}^2, 1.5 \text{ m}^3$
Total crystal weight	12.5 t
Operating temperature	$-25^\circ \text{ C}$
<b>CPV</b>	
Gas	80% Ar/20% CO <sub>2</sub>
Thickness	$0.05 X_0$
Active area	$8.6 \text{ m}^2$
Wire diameter	$30 \mu\text{m}$
Number of wires per module	256
Wire pitch	5.65 mm
Pad size, pad inter-distance	$22 \times 10.5 \text{ mm}^2, 0.6 \text{ mm}$
Number of pads per module	7168

To increase the light yield of the PWO crystals (temperature coefficient  $-2\%$  per  $^\circ\text{C}$ ), the PHOS modules are operated at a temperature of  $-25^\circ \text{ C}$  stabilized with a precision of  $0.3^\circ \text{ C}$ . The crystals are located inside a cold enclosure, whereas the readout electronics are located outside this enclosure. The temperature is monitored by resistive temperature sensors of thickness  $30\text{--}50 \mu\text{m}$ , inserted in the gap between crystals.

A calibration system using Light Emitting Diodes (LED) and stable current generators monitors every PHOS detection cell.





### 3.6.2.2 Charged-Particle Veto detector

The Charged-Particle Veto (CPV) detector is a Multi-Wire Proportional Chamber (MWPC) with cathode-pad readout [129, 130]. Its charged-particle detection efficiency is better than 99%. The spatial resolution of the reconstructed impact point is about 1.54 mm along the beam direction and 1.38 mm across the beam.

The CPV is placed on top of the PHOS modules at a distance of about 5 mm. The material budget is less than 5% of  $X_0$ . The active volume of 14 mm thickness is filled with a gas mixture 80% Ar/20% CO<sub>2</sub> at a pressure slightly (1 mbar) above atmospheric pressure.

The cathode plane is segmented into 7168 pads of size  $22 \times 10.5 \text{ mm}^2$  with an inter-pad distance of 0.6 mm. The largest dimension is aligned along the wires and the total active area of the CPV module is about  $1.8 \text{ m}^2$ .





**Figure 3.57:** Left top: Crystal detector unit with glued photo-detector; photo-detector, APD mounted on the preamplifier substrate. Right top: Strip unit containing  $8 \times 2$  crystal detector units. Left bottom: PHOS module with strip units installed onto cooling plates. Right bottom: 5 PHOS modules.

**Table 3.17:** Summary of PHOS front-end electronic parameters.

Least count energy single channel	5–10 MeV
Dynamic range	0.005-80 GeV
Energy channels	‘high’ and ‘low’ gains
Timing resolution	around 2 ns at 2 GeV
Trigger	L0, L1
max channel counting rate	
in Pb-Pb	1 kHz
in pp	10 Hz
APD gain control	individual bias setting

### 3.6.3 Front-end electronics and readout

A substantial effort was made for PHOS to investigate alternative shaper designs and shaping times. In the case of PHOS, a primary consideration is optimum performance for measurement of low energy photons, in the region around 1 GeV, of interest for observation of thermal photon radiation from the Quark Gluon Plasma. With the good intrinsic resolution of  $\text{PbWO}_4$ , care must be taken to minimize the electronic contribution to the noise. This can be done by a judicious choice of the shaping time of the amplifier. For a detector capacitance of  $C_D=100$  pF as for the Hamamatsu S8664-55 APD, the results indicated an electronic noise minimum of about  $300e$  for a shaping time of about  $2\mu\text{s}$ . The competing consideration of a reasonably good timing measurement motivated a final choice of  $1\mu\text{s}$  shaping time for PHOS.

#### 3.6.3.1 PHOS front-end electronics

The PHOS electronic chain includes energy digitization and programmable trigger logic generating L0 and L1 triggers to ALICE. Each input channel supplies two shaper channels, one with ‘low’ and one with ‘high’ amplification, digitized in separate ADCs. The gains of the APDs are tuned to equalize the channel response by means of a control system where the bias is set individually for each APD. The timing information is derived offline from the shape of the energy pulse. The main physics requirements for the front-end electronics are summarized in table 3.17. There are 4 electronic channels per PHOS crystal: high and low energy gain, fast trigger logic and APD bias control. The total number of electronic channels for the five PHOS modules is 89 600. The main parameters of the preamplifier are [127]: sensitivity 1 V/pC, max. input charge 8 pC, ENC around 400 e for  $C_{\text{APD}} = 100$  pF, and power dissipation 62 mW. The shaping amplifier has a time constant of  $1.0\mu\text{s}$ . The electronic noise in the output signal corresponds to around 10 MeV. The front-end electronics of one PHOS module is segmented (see figure 3.58) into 112 FEE boards each processing the signal from 32 crystals, eight trigger boards (TRU) each covering a region of  $28 \times 16$  crystals and 4 readout controller units (RCUs) serving two branches of 14 FEE boards through a GTL bus. The trigger algorithm calculates in parallel all  $4 \times 4$  partial energy sums of the digitized fast energy signals from the trigger region. This algorithm is implemented in a FPGA, programmable via the JTAG input. To control the gain of the energy channels, the APD bias control system regulates the bias voltage of each of the APDs with an precision of 0.2 V.




Figure 3.58: Layout of the PHOS electronics.

### 3.6.3.2 CPV front-end electronics

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

### 3.6.3.3 Readout

The readout is based on the concept and modules developed for the TPC detector see section 3.2. A readout controller unit (RCU) transfers formatted event data over the ALICE Digital Data Link (DDL) to the ALICE DAQ system. The RCU operations are executed by means of firmware in the on-board FPGA. The firmware includes a Detector Control System (DCS) card with processor core for handling the Ethernet connection to the ALICE detector control system. The data and control interfaces between the RCU and the TRUs of the front-end electronics are implemented by means of a GTL bus.

## 3.7 ElectroMagnetic CALorimeter (EMCal)

### 3.7.1 Design considerations

The construction of a large ElectroMagnetic Calorimeter (EMCal) [7] began in 2008 with the aim to enable ALICE to explore in detail the physics of jet quenching (interaction of energetic partons with dense matter) over the large kinematic range accessible in heavy-ion collisions at the LHC [20].

The scope and basic design parameters of the calorimeter were chosen to match the physics performance requirements of the high- $p_t$  physics goals [21]. The EMCal is a large Pb-scintillator

sampling calorimeter with cylindrical geometry, located adjacent to the ALICE magnet coil at a radius of  $\sim 4.5$  metres from the beam line. It covers  $|\eta| \leq 0.7$  and  $\Delta\phi = 107^\circ$ , and is positioned approximately opposite in azimuth to the high-precision ALICE Photon-Spectrometer (PHOS) calorimeter. The size of the EMCal is constraint by the available free space and the maximum weight which can be supported by the L3 magnet.

The choice of a large-acceptance, moderate-resolution electromagnetic calorimeter provides a cost-effective pathway into jet physics in ALICE. The EMCal increases the electromagnetic calorimeter coverage of ALICE by nearly an order of magnitude. It provides a fast and efficient trigger (L0, L1) for hard jets, photons and electrons, allowing ALICE to exploit fully the luminosity of the LHC. The EMCal also measures the neutral energy component of jets, enabling full jet reconstruction in all collision systems, from pp to Pb–Pb. The combination of the EMCal, the excellent ALICE charged particle tracking capabilities, and the modest ALICE magnetic-field strength, is a preferred configuration for jet reconstruction in the high-background environment of heavy-ion collisions, allowing the optimization of background rejection while preserving the crucial jet-quenching signals at moderate transverse momentum. The EMCal in ALICE in conjunction with the TPC has good jet energy resolution in Pb–Pb collisions, and excellent sensitivity to the full range of jet-quenching effects expected at the LHC.


The emphasis on high  $p_t$  means that the intrinsic energy resolution of the EMCal can be modest and the detector granularity can be coarser, with moderately high occupancy. These design criteria point toward the choice of a sampling-calorimeter design, similar to those that were built for STAR and PHENIX to study high- $p_t$  phenomena in heavy-ion collisions at RHIC and then in the LHC-B experiment at the LHC.

### 3.7.2 Mechanical issues

The overall design of the EMCal is heavily influenced by its location within the ALICE L3 magnet. Figure 3.59 shows a schematic integration drawing of the end view of the ALICE central barrel. The EMCal is to be located inside the large room-temperature L3 magnet within a cylindrical integration volume 112 cm deep sandwiched between the ALICE central-detectors space frame (housing the TPC, TRD and TOF) and the magnet coils. Due to the installation of the PHOS carriage below the ALICE TPC and the HMPID above the ALICE TPC, the EMCal is limited to a region of about  $107^\circ$  in azimuth above the TPC adjacent to the HMPID.

The EMCal support structure bridges the azimuthal gap between two dedicated rails mounted to the ends of the L3 magnet yoke. These rails support the full, 100 t weight of the detector. The EMCal is installed into the support structure as twelve super module units, with the lower two super modules being of one-third size.

Small azimuthal gaps ( $\sim 3$  cm) between super modules facilitate installation and alignment. These gaps are positioned in line with the TPC sector boundaries. Along these sector boundaries there is substantial additional structural material required for the support of the TPC and other detectors that would significantly degrade any electromagnetic measurements made in these gaps. Thus, the gaps create no additional loss of electromagnetic acceptance. Detailed simulations show that these gaps have no significant influence on the measured jet energy or jet resolution.




**Figure 3.59:** Schematic integration drawing of the end view of the ALICE central barrel.

The chosen technology is a layered Pb-scintillator sampling calorimeter with a longitudinal pitch of 1.44 mm Pb and 1.76 mm scintillator with longitudinal wavelength-shifting fibre light collection (Shashlik). The detector is segmented into 12 288 towers, each of which is approximately projective in  $\eta$  and  $\phi$  to the interaction vertex. The towers are grouped into super modules of two types: ‘full size’ which span  $\Delta\eta = 0.7$  and  $\Delta\phi = 20^\circ$ , and ‘one-third size’ which span  $\Delta\eta = 0.7$  and  $\Delta\phi = 7^\circ$ . There are 10 full-size and 2 one-third-size super modules in the full-detector acceptance. The super module is the basic structural unit of the calorimeter. These are the units handled as the detector is moved below ground and rigged during installation into their final resting place in the L3 magnet.


Each full-size super module (see figure 3.60) is assembled from  $12 \times 24 = 288$  modules arranged in 24 strip modules of 12 modules each (see below). Each module has a rectangular cross section in the  $\phi$  direction and a trapezoidal cross section in the  $\eta$  direction with a full taper of 1.5 degrees. The resultant assembly of stacked strip modules is approximately projective with an average angle of incidence at the front face of a module of less than 2 degrees in  $\eta$  and less than 10 degrees in  $\phi$ .

The smallest building block of the calorimeter is the individual module. A cross section of a single module is shown in figure 3.61. Each individual module contains  $2 \times 2 = 4$  towers built up from 76 alternating layers of 1.44 mm Pb (Pure Pb-standard mill spec.) and 77 layers of 1.76 mm polystyrene base, injection molded scintillator (BASF143E + 1.5%pTP + 0.04%POPOP) with an intrinsic light output of 50% Anthracene [131]. To best account for inactive materials in the space immediately before the calorimeter, the first layer of the detector is scintillator.

White, acid free, bond paper serves as a diffuse reflector on the scintillator surfaces and, in addition, provides the required degree of static friction between the calorimeter layers. The scintillator edges are treated with a  $\text{TiO}_2$  loaded reflector to improve the transverse optical uniformity



**Figure 3.60:** EMCAL super module.



**Figure 3.61:** Cross sectional view of a single EMCAL module. This view shows the 1.5 degree tapered cross section of the module. The principal components of the module are indicated. The size of module is  $\sim 12 \text{ cm} \times 12 \text{ cm} \times 24.6 \text{ cm}$ .

within a single tower. Each scintillator edge also receives an over-coating of black paint to provide tower to tower optical isolation better than 99%. The total thickness of the reflective and optical isolating material on the scintillator edges is  $\sim 50 \mu\text{m}$ . Before completing the final mechanical closure of the module, it is wrapped in a thin layer of black adhesive paper to improve the light tightness.

All modules in the calorimeter are mechanically and dimensionally identical. The front-face dimensions of the towers are  $6 \times 6 \text{cm}^2$  resulting in individual tower acceptance of  $\Delta\eta \times \Delta\phi \sim 0.014 \times 0.014$ .

The Pb-scintillator stack in a module is secured in place by the static friction between individual layers under the load of an internal pressure of  $2 \text{kg/cm}^2$ . The module is closed by a skin of  $150 \mu\text{m}$  thick stainless steel on all four transverse surfaces. Flanges are welded to the ends of the stainless steel skin (straps) to permit secure attachment to the front and rear plates of the module. This thin stainless skin plus the optical treatment of the scintillator edges is the only inert material between the active tower volumes. The internal pressure in the module is stabilized against thermal effects, mechanical relaxation and long term flow of the Pb and/or polystyrene by a customized array of five non-linear spring sets (Belleville washers) per module. In this way, each module is a self supporting unit with a stable mechanical lifetime of more than 20 years with a large safety factor on the final compressing force. This compressed module design allows the detector to be completely supported from its rear surface.

### 3.7.3 Strip modules and super modules

An array of 12 modules in the  $\phi$  direction and a strong back (which links them) form an EMCal ‘strip module’. Like the individual module, the strip module is a self-supporting unit. The super module, discussed above, is made from 24 strip modules mounted in  $\eta$  direction. The strong back, an aluminium beam, which integrates the modules into the strip module, functions as the mechanical support for modules and a stiff structural element of the super module structure. It provides protection for the optical fibres, structural mount for the light guide, APD and charge-sensitive preamplifier housing, and a light-tight enclosure of these elements.

The super module shown in figure 3.60 is the largest building block of the calorimeter. Starting with 288 individual EMCal modules which are rather compact and heavy, a stiff super module structure was designed that performs with very small deflections in any  $\phi$  orientation yet does not require extensive, heavy external stiffening components that would reduce the available volume for the active detector.

The concept adopted for the ALICE EMCAL is a super module ‘crate’ which functions not as a box for the individual modules, but rather as an integrated structure in which the individual elements (strip modules) contribute to the overall stiffness. The super module crate thus functions as a large I-beam in which the aluminium flanges of the beam are the longitudinal sides of the crate and the 24 transverse rows of strip modules together form the ‘core’. With 1444 mm spaced flanges, this configuration gives to the super module a good stiffness for any  $\phi$  orientation.



**Figure 3.62:** EMCAL support structure and the associated rail supports for super modules is shown on its assembly fixture, above ground at the ALICE site (Point 2).

The interface between the super module crate and the EMCAL support structure (figure 3.62) is achieved with a set of rails mounted on the inner surface of the EMCAL support structure and rollers fitted to carriages mounted on super module crates (see figure 3.60). Each super module slides into its resting place inside the ALICE L3 magnet on two U-shaped aluminium rails. This system allows installation of super modules as needed during the annual LHC shutdown period until the full EMCAL is in place.

### 3.7.4 Module physical parameters

The average density of the detector's active volume is  $5.68 \text{ g/cm}^3$  which results from a 1:1.22 Pb to scintillator ratio by volume. This results in a compact detector, consistent with the EMCAL integration volume at the chosen detector thickness of 20.1 radiation lengths ( $20.1 X_0$ ). In simulations, this number of radiation lengths gives a maximum deviation from linearity (due mainly to shower leakage) of 1.5% for the most probable energy response in the range up to 100 GeV photons.

The impact of detector energy resolution on the proposed physics programme was studied extensively. Given the main focus on jet physics and the significant non-detector contributions to the energy resolution in the heavy-ion environment, the calorimeter resolution is not very critical. Simulations and test-beam measurements indicate that the final detector resolution [7] will be better than  $10\%/\sqrt{E}$  which is completely adequate to accomplish the proposed scientific programme. The physical characteristics of the EMCAL are summarized in table 3.18.



**Table 3.18:** The EMCal physical parameters.

Quantity	Value
Tower size (at $\eta = 0$ )	$6.0 \times 6.0 \times 24.6 \text{ cm}^3$ active
Tower size	$\Delta\phi \times \Delta\eta = 0.0143 \times 0.0143$
Sampling ratio	1.44 mm Pb / 1.76 mm scintillator
Number of layers	77
Effective radiation length $X_0$	12.3 mm
Effective moliere radius $R_M$	3.2 cm
Effective density	$5.68 \text{ g/cm}^3$
Sampling fraction	10.5
Number of radiation length	20.1
Number of towers	12,672
Number of modules	3,168
Number of super modules	10 full size, 2 half size
Weight of super modules	7.7 metric tons (full size)
Total coverage	$\Delta\phi = 107^\circ, -0.7 \leq \eta \leq 0.7$

### 3.7.5 Module optical system

Scintillation photons produced in each tower are captured by an array of 36 Kuraray Y-11 (200 M), double-clad, wavelength-shifting (WLS) fibres that run longitudinally through the Pb/scintillator stack. Each fibre terminates in an aluminized mirror at the front face of the module and is integrated into a polished, circular group of 36 at the photo-sensor end at the back of the module. Because the tower transverse shape deviates slightly from square as a function of longitudinal depth, a fibre pattern was chosen which has exactly the same aspect ratio as the mechanical tower shape at a depth close to the electromagnetic shower maximum. This has the effect of making the fibre pattern uniform across tower boundaries when weighted by the shower energy deposition.

The mirroring on the distant end of the fibre creates an approximate compensation for the effects of the finite attenuation lengths within the fibre making the longitudinal response of the detector quite uniform. This is most significant for the shorter attenuation lengths associated with cladding light that is quite important in this rather short detector. The properties of the selected fibres are given in table 3.19.

The fibre bundle in a single tower terminates in a 6.8 mm diameter disk and connects to the APD photo sensor through a short light guide/diffuser with a square cross section of  $7 \text{ mm} \times 7 \text{ mm}$  that tapers slowly down to  $4.5 \text{ mm} \times 4.5 \text{ mm}$  as it mates (glued) to the  $5 \text{ mm} \times 5 \text{ mm}$  active area of the APD.

Figure 3.63 shows 4 pre-fabricated fibre bundles inserted into the towers of a single prototype module. In this picture all of the module rear enclosing and structural elements are omitted so that the wavelength shifting fibres may be seen as they converge to the light guide (inside the black plastic tube) and finally to mate with the APD and charge sensitive preamplifier.

The selected photo sensor is the Hamamatsu S8664-55 Avalanche Photo Diode. This photodiode has a peak spectral response at a wavelength of 585 nm compared to an emission peak of

**Table 3.19:** Characteristics of the EMCal module wavelength shifting fibres.

Quantity	Value
WLS fibre	Y-11 (200) M-DC
Manufacturer	Kuraray
WLS Fluor	K27 200 mg
Absorption peak	430 nm
Emission peak	476 nm
Decay time	7 ns
Core material	PS
Refractive index of core	1.59
Inner cladding	PMMA
Refractive index of cladding	1.49
Outer cladding	FP
Refractive index of outer cladding	1.42
Long fibre attenuation length	3.5m
Fibre diameter	1.0 mm



**Figure 3.63:** Prototype EMCal module. The back enclosing structure is removed so that the converging, wavelength shifting fibres are visible from each of the 4 towers. The charge sensitive preamps are shown integrated directly to the APD (not visible).

476 nm for the Y-11 fibres. However, both the spectral response and the quantum efficiency of the APD are quite broad with the latter dropping from the maximum by only 5% at the WLS fibre emission peak. At this wavelength, the manufacturer's specification gives a quantum efficiency of 80%.

### 3.7.6 Readout and trigger

Both the selected APD and the charge sensitive preamp are identical to those used for the PHOS calorimeter, described in section 3.6. The light yield, per unit of energy, deposit in the EMCal is similar to that of the PHOS [9, 127]. Since the electronic noise performance requirements of the EMCal are less stringent than those of PHOS (due to the larger intrinsic energy resolution of the EMCal), the PHOS readout electronics were adopted for the EMCal readout, with only minor modification. The only significant difference with the PHOS readout is the difference in the FEE amplifier due to the chosen dynamic range, the EMCal light yield, and the amplifier shaping time.

#### Dynamic range

Based on the expected annual yield of photons and  $\pi^0$  at high  $p_t$ , the full scale energy range for an EMCAL tower is chosen to be 250 GeV as compared to 80 GeV for PHOS, well beyond the expected maximum photon or  $\pi^0$  energy. Setting the EMCAL full-scale energy range to 250 GeV sets the Least Significant Bit (LSB) on the lower of the two overlapping gain ranges to 250 MeV (10-bits) with the corresponding maximum energy on the high-gain range at 16 GeV and least-significant bit at 16 MeV.

#### Light yield

An important parameter of the EMCal readout is the light output to the APD per MeV of energy deposited in a tower. The test-beam measurements with the EMCal prototype modules gave a measured light yield of about 5 photo-electrons/MeV from the APD at gain  $M=1$ , similar to the PHOS which produces 4.4 photo-electrons/MeV.

#### 3.7.6.1 Shaper time constant optimization

Several considerations motivate a much shorter shaping time for the EMCal than the one for PHOS. First, simulations of central Pb–Pb collisions within the AliRoot offline framework and using HIJING event generator (see chapter 6.4) indicate that the EMCal, with its hydrogen-rich scintillator, will be affected by a large slow-neutron contribution that has a tail extending for hundreds of ns after the collision. Second, the number of ALTRO chip digitizer samples recorded is dictated by the total shaped pulse width. With the approximately ten times coarser EMCal granularity, the occupancies will be correspondingly higher in EMCal than PHOS. With the goal to keep the total data volume per Readout Control Unit (RCU) similar to that of PHOS, the number of EMCal samples will be reduced to keep the product of occupancy  $\times$  no. samples similar. This can be achieved with a shaping time of about 100 ns. With 100 ns shaping time the voltage noise will dominate such that one expects a total electronics noise contribution of about 1500  $e$  for the APD+preamplifier+shaper chain, corresponding to an electronics noise contribution of about 12 MeV per EMCal tower.

The significance of this electronic noise is best illustrated relative to the estimated energy fluctuations due to late-neutron background. Simulations of the EMCal response for central Pb–Pb collisions have shown that there is a large background energy deposit predominantly from late neutrons produced in secondary interactions in the surrounding materials of the ALICE experiment [132]. The time distribution of the background has a long tail with arrival times at the

calorimeter extending for several hundreds of ns after the collision. (Note that this background is generated by the primary central Pb–Pb collision, and hence it is a centrality dependent effect, expected to scale with the multiplicity of produced particles). The result indicates that as short a shaping time as feasible should be used in order to minimize this background contribution to the energy measurement. Without any additional timing cut, the average background energy deposit is 36 MeV per tower, i.e. several times greater than the expected electronics noise contribution per tower with 100 ns shaping time.

### 3.7.6.2 EMCal trigger


The PHOS readout, which was adopted for the EMCal, includes a Trigger Region Unit card (TRU) [133] on each GTL readout bus that provides sliding  $4 \times 4$  tower sums for generation of a Level-0 and three Level-1 triggers with programmable thresholds for simultaneous low-, mid-, and high-energy shower trigger outputs. All of the trigger outputs generated locally from each TRU are logically ‘OR’ed and input to the ALICE Central Trigger Processor (CTP). The PHOS trigger was designed to trigger efficiently on photons and electrons and will serve that same purpose equally well for the EMCal. The PHOS TRU will be used for the EMCal with only minor modifications. First, because the EMCal input connections to the FEE cards do not have exactly the same geometrical mapping as for PHOS, the TRU FPGA code will be modified to take into account the different tower mapping of PHOS and the EMCal.

The PHOS TRU single shower trigger provides a good leading-photon, leading- $\pi^0$  trigger that also triggers efficiently on high-energy jets. However, because a single-shower trigger will have a bias toward jets including a leading  $\pi^0$  it is of interest to have the capability to trigger on larger regions of EMCAL towers.

For this purpose a Summary Trigger Unit (STU) will take input from the EMCal TRU modules and perform energy sums over larger regions of the EMCal. The PHOS TRU design was modified for this purpose to provide raw trigger sum output that can input to the STU to form trigger sums over larger regions that cross TRU boundaries. Also, the STU will use the available raw data in order to improve the single shower-trigger by TRU boundaries for  $4 \times 4$  tower sums. The STU will also receive particle multiplicity information from the V0 trigger detector that will allow to apply centrality dependent single shower and jet triggers. At the end of the trigger sequence (Level 2) the STU will provide, via sums DDL link, all raw trigger sums to the DAQ and HLT.

### 3.7.7 EMCal calibration

The initial pre-calibration of the EMCal will be performed at the time of super module assembly using the ‘mip’ peak obtained with cosmic-ray particles triggered by scintillator arrays above and below the super module. The individual APD voltages will be adjusted to match the cosmic mip peak in each tower to bring all towers to the same relative gain. Tests of the planned cosmic calibration procedure using prototype EMCal modules following beam-test measurements demonstrated that the cosmic mip peak can provide relative gain matching of each tower with an rms variation of less than 10%. This level of initial gain matching is required to provide a uniform single-shower trigger response. Ultimately the absolute energy scale for each tower will be



**Figure 3.64:** Photograph of the ACORDE scintillator module array on the upper faces of the magnet yoke in the ALICE pit. An ACORDE module unit includes two overlapping scintillator counters, each with  $190 \times 20 \text{ cm}^2$  effective area, and 10 mm thickness, plus two PMTs placed at the end of each scintillator, with optic guide and iron magnetic shielding. All components are attached to an aluminium structure, assembling a 40 kg robust box. There are twenty modules on each one of the three upper faces.

obtained using physics data from proton–proton collisions. This will be extracted using a combination of measurements of the mip peak, electron  $E/p$  ratio, and  $\pi^0$  mass for showers centered on each tower.

## 3.8 ALICE COsmic Ray DEtector (ACORDE)

### 3.8.1 Design considerations

ACORDE, the ALICE cosmic ray detector, is an array of plastic scintillator counters placed on the upper surface of the L3 magnet. It plays a two-fold role in ALICE: i) the first task is to provide a fast (Level-0) trigger signal, for the commissioning, calibration and alignment procedures of some of the ALICE tracking detectors; ii) it will also detect, in combination with the TPC, TRD and TOF, single atmospheric muons and multi-muon events (so-called muon bundles) thus allowing us to study high-energy cosmic rays in the energy region of the knee in the cosmic ray spectrum. A feasibility and performance study is detailed in [20, 21, 134, 135].

### 3.8.2 Detector layout

An ACORDE module consists of two scintillator counters, each with  $190 \times 20 \text{ cm}^2$  effective area, placed on top of each other and read out in coincidence (figure 3.64). With this setup a uniform efficiency higher than 90% along the whole length of a test module is achieved. Details of the construction and performance of this counters are explained in [134]. The ACORDE scintillator

**Table 3.20:** Synopsis of ACORDE parameters.


Pseudo-rapidity coverage	$-1.3 < \eta < 1.3$
Azimuthal coverage	$-60^\circ < \phi < 60^\circ$
Radial position	$r = 8.5$ m
Total number of detector modules	60
Module size	$300 \times 26 \times 10$ cm <sup>3</sup>
Active detector area	$190 \times 20$ cm <sup>2</sup>
Radial detector thickness	$X/X_0 = 4.7\%$
Module segmentation	30 in $\phi$ 2 in $z$
Number of readout channels	120
Number of readout boards	60
Trigger time delay electronics for	single muon
	multi-coincidence
	TRD wake-up
	220 ns
	270 ns
	170 ns
Rate for cosmic single muon events	4.5 Hz/m <sup>2</sup>
Rate for cosmic 5-fold coincidence module events	$5 \times 10^{-3}$ Hz

module array, which includes 60 scintillator counter modules placed on top of the ALICE magnet, is shown in figure 3.64. ACORDE provides a fast Level-0 trigger signal to the Central Trigger Processor, when atmospheric muons impinge upon the ALICE detector. The signal is used for the calibration, alignment and performance of several ALICE tracking detectors, mainly the TPC, TOF, HMPID and ITS. The operational Cosmic Ray Trigger is delivering trigger signals independent of the LHC beam. Table 3.20 summarize the operation conditions of this detector.

The typical rate for single atmospheric muons reaching the ALICE detector is relatively low (4.5 Hz/m<sup>2</sup>, on top of the magnet), the rate for multi-muon events is expected to be much lower (less than  $10^{-3}$  Hz/m<sup>2</sup>). This is statistically sufficient for the study of these type of events, provided we can trigger and store tracking information from cosmic muons parallel to the ALICE normal data taking with colliding beams. Atmospheric muons need an energy of at least 17 GeV to reach the ALICE hall, while the upper energy limit for reconstructed muons in the TPC will be near 2 TeV, at a magnetic field intensity of 0.5 T. This allows us to measure and analyze the atmospheric muon momentum spectra in a wide range (0.1–2 TeV) with a very high precision.

### 3.8.3 Readout electronics.

The ACORDE electronics was designed to fulfil the following requirements: i) produce the single muon trigger signal and provide a rough spatial position when atmospheric muons hit the ALICE magnet. This information will be used to calibrate the ALICE TPC detector and other tracking detectors; ii) provide a multi-coincidence trigger signal and store kinematics information from cosmic muons. Given the low rate of this trigger (less than  $10^{-3}$  Hz), it would greatly benefit from running in parallel with ALICE normal data taking; iii) provide the wake-up signal to the ALICE Transition Radiation Detector (TRD), and iv) supervise the performance of the ACORDE scintillator counter array.



**Figure 3.65:** Block diagram of the ACORDE electronics system.

The ACORDE electronics (figure 3.65) is formed by 60 FEE Cards, one per ACORDE module; and the ‘ACORDE OR’ Card to generate the TRD wake up signal. This card receives the 60 coincidence LVDS signals coming from the FEE cards. The final component is the Main Card, which contains the electronics to receive the 120 LVDS signals coming from 120 scintillator. This card produces the single and the multi-coincidence trigger signals and provides the connectivity to the ALICE trigger and DAQ systems.