Chapter 3

Inner tracking system

3.1 Introduction

The inner tracking system of CMS is designed to provide a precise and efficient measurement of the trajectories of charged particles emerging from the LHC collisions, as well as a precise reconstruction of secondary vertices. It surrounds the interaction point and has a length of 5.8 m and a diameter of 2.5 m. The CMS solenoid provides a homogeneous magnetic field of 4 T over the full volume of the tracker. At the LHC design luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ there will be on average about 1000 particles from more than 20 overlapping proton-proton interactions traversing the tracker for each bunch crossing, i.e. every 25 ns. Therefore a detector technology featuring high granularity and fast response is required, such that the trajectories can be identified reliably and attributed to the correct bunch crossing. However, these features imply a high power density of the on-detector electronics which in turn requires efficient cooling. This is in direct conflict with the aim of keeping to the minimum the amount of material in order to limit multiple scattering, bremsstrahlung, photon conversion and nuclear interactions. A compromise had to be found in this respect. The intense particle flux will also cause severe radiation damage to the tracking system. The main challenge in the design of the tracking system was to develop detector components able to operate in this harsh environment for an expected lifetime of 10 years. These requirements on granularity, speed and radiation hardness lead to a tracker design entirely based on silicon detector technology. The CMS tracker is composed of a pixel detector with three barrel layers at radii between 4.4 cm and 10.2 cm and a silicon strip tracker with 10 barrel detection layers extending outwards to a radius of 1.1 m. Each system is completed by endcaps which consist of 2 disks in the pixel detector and 3 plus 9 disks in the strip tracker on each side of the barrel, extending the acceptance of the tracker up to a pseudorapidity of $|\eta| < 2.5$. With about $200\,\mathrm{m}^2$ of active silicon area the CMS tracker is the largest silicon tracker ever built [15, 16].

The construction of the CMS tracker, composed of 1440 pixel and 15 148 strip detector modules, required the development of production methods and quality control procedures that are new to the field of particle physics detectors. A strong collaboration of 51 institutes with almost 500 physicists and engineers succeeded over a period of 12 to 15 years to design, develop and build this unique device.

3.1.1 Requirements and operating conditions

The expected LHC physics program [17] requires a robust, efficient and precise reconstruction of the trajectories of charged particles with transverse momentum above 1 GeV in the pseudorapidity range $|\eta| < 2.5$. A precise measurement of secondary vertices and impact parameters is necessary for the efficient identification of heavy flavours which are produced in many of the interesting physics channels. Together with the electromagnetic calorimeter and the muon system the tracker has to identify electrons and muons, respectively. Tau leptons are a signature in several discovery channels and need to be reconstructed in one-prong and three-prong decay topologies. In order to reduce the event rate from the LHC bunch crossing rate of 40 MHz to about 100 Hz which can be permanently stored, tracking information is heavily used in the high level trigger of CMS.

The operating conditions for a tracking system at the LHC are very challenging. As already mentioned, each LHC bunch crossing at design luminosity creates on average about 1000 particles hitting the tracker. This leads to a hit rate density of $1\,\mathrm{MHz}/\,\mathrm{mm}^2$ at a radius of 4 cm, falling to 60 kHz/mm^2 at a radius of 22 cm and 3 kHz/mm^2 at a radius of 115 cm. In order to keep the occupancy at or below 1% pixelated detectors have to be used at radii below 10 cm. For a pixel size of $100 \times 150\,\mu\mathrm{m}^2$ in r-ϕ and z, respectively, which is driven by the desired impact parameter resolution, the occupancy is of the order 10^{-4} per pixel and LHC bunch crossing. At intermediate radii ($20\,\mathrm{cm} < r < 55\,\mathrm{cm}$) the reduced particle flux allows the use of silicon micro-strip detectors with a typical cell size of $10\,\mathrm{cm} \times 80\,\mu\mathrm{m}$, leading to an occupancy of up to 2–3% per strip and LHC bunch crossing. In the outer region ($55\,\mathrm{cm} < r < 110\,\mathrm{cm}$) the strip pitch can be further increased. Given the large areas that have to be instrumented in this region, also the strip length has to be increased in order to limit the number of read-out channels. However, the strip capacitance scales with its length and therefore the electronics noise is a linear function of the strip length as well. In order to maintain a good signal to noise ratio of well above 10, CMS uses thicker silicon sensors for the outer tracker region (500 μm thickness as opposed to the 320 μm in the inner tracker) with correspondingly higher signal. These thicker sensors would in principle have a higher depletion voltage. But since the radiation levels in the outer tracker are smaller, a higher initial resistivity can be chosen such that the initial depletion voltages of thick and thin sensors are in the same range of 100 V to 300 V. In this way cell sizes up to about $25\,\mathrm{cm} \times 180\,\mu\mathrm{m}$ can be used in the outer region of the tracker, with an occupancy of about 1%. These occupancy-driven design choices for the strip tracker also satisfy the requirements on position resolution.

CMS is the first experiment using silicon detectors in this outer tracker region. This novel approach was made possible by three key developments:

- sensor fabrication on 6 inch instead of 4 inch wafers reduced the sensor cost to 5–10 CHF/cm^2 and allowed the coverage of the large required surfaces with silicon sensors,
- implementation of the front-end read-out chip in industry-standard deep sub-micron technology led to large cost savings and to an improved signal-to-noise performance,
- automation of module assembly and use of high throughput wire bonding machines.

The radiation damage introduced by the high particle fluxes at the LHC interaction regions is a severe design constraint. Table 3.1 shows the expected fast hadron fluence and radiation dose

2008 JINST 3 S08004

Table 3.1: Expected hadron fluence and radiation dose in different radial layers of the CMS tracker (barrel part) for an integrated luminosity of $500\,\mathrm{fb}^{-1}$ (≈ 10 years). The fast hadron fluence is a good approximation to the 1 MeV neutron equivalent fluence [17].

Radius (cm)	Fluence of fast hadrons (10^{14} cm^{-2})	Dose (kGy)	Charged particle flux (cm^{-2}s^{-1})
4	32	840	10^8
11	4.6	190	
22	1.6	70	6×10^6
75	0.3	7	
115	0.2	1.8	3×10^5

in the CMS barrel tracker for an integrated luminosity of $500\,\mathrm{fb}^{-1}$ corresponding to about 10 years of LHC operation [15, 17]. Neutrons generated by hadronic interactions in the ECAL crystals make up a substantial contribution to the fast hadron fluence, which actually dominates in the outer tracker close to the ECAL surface. The uncertainties on these estimates due to the extrapolation error of the inelastic proton proton cross-section, momentum distributions and multiplicities to $\sqrt{s} = 14\,\mathrm{TeV}$ and in the Monte Carlo description of the cascade development lead to a safety factor of 1.5 (2 in regions where the neutron contribution dominates) which was applied to these estimates in order to define the design requirements for the tracker.

Three different effects had to be considered in the design of a radiation tolerant silicon tracker. Surface damage is created when the positively charged holes, generated by the passage of an ionizing particle, get trapped in a silicon oxide layer. This is mostly a concern for the front-end chips where this additional space charge changes for instance the characteristics of MOS structures. Surface damage simply scales with the absorbed dose. The silicon sensors are mainly affected by bulk damage, i.e. modifications to the silicon crystal lattice which are caused by non-ionizing energy loss (NIEL) and lead to additional energy levels in the band gap. NIEL is a complicated process, depending on particle type and energy, but is found to scale approximately with the fast hadron fluence. The consequences are an increase of the leakage current (linear in fluence), a change in the doping from n- to p-type with a corresponding change in depletion voltage by a few hundred volts over the lifetime of the tracker, and the creation of additional trapping centers which will reduce the signal by roughly 10% after 10 years of LHC running [18]. The design of the silicon sensors and the read-out electronics has to take this into account and assure a signal-to-noise ratio of 10:1 or better over the full lifetime of the detector, in order to guarantee a robust hit recognition at an acceptable fake hit rate. Finally, transient phenomena due to the generation of charge by ionizing particles in the electronic circuitry can change for instance the state of memory cells and therefore disturb or even stop the correct functioning of the read-out (single event upset, SEU).

The increased detector leakage current can lead to a dangerous positive feedback of the self heating of the silicon sensor and the exponential dependence of the leakage current on temperature, called thermal runaway. This has to be avoided by efficient coupling of the silicon sensors to the cooling system and by a low operating temperature. For this reason it is foreseen that the whole tracker volume will be operated at or slightly below $-10°$C. After 10 years of operation it is

2008 JINST 3 S08004

expected that this will require a cooling fluid temperature of about $-27°C$ which in turn means that all structures in the tracker have to survive temperature cycles between room temperature and about $-30°C$. A second effect, called reverse annealing, requires to keep the silicon sensors permanently well below 0°C except for short maintenance periods. This effect is caused by the interaction of radiation induced defects in the silicon sensors which can lead to more serious damage and to an even stronger change in depletion voltage with fluence. Experimentally it is found that reverse annealing becomes insignificant for temperatures roughly below 0°C [18].

The read-out chips employed in the CMS tracker are fabricated in standard 0.25 μm CMOS technology which is inherently radiation hard due to the thin gate oxide (and special design rules). The lifetime of the silicon strip tracker is therefore limited by the radiation damage to the silicon sensors. For efficient charge collection they always need to be over-depleted, requiring bias voltages up to 500 V after 10 years of LHC operation. This reaches the limit of the typical high voltage stability of current sensor layouts. Furthermore, the increased leakage currents of the sensors will at some point lead to thermal runaway. All tests have shown that the silicon strip tracker will remain fully operational for 10 years of LHC running. For the pixel detector on the other hand, which has to survive even higher radiation doses, under-depleted operation is possible due to a different sensor layout. Its lifetime reaches from at least 2 years at full LHC luminosity for the innermost layer to more than 10 years for the third layer.

The ultimate position resolution of the pixel and strip sensors is degraded by multiple scattering in the material that is necessary to precisely hold the sensors, to supply the electrical power (in total about 60 kW for the CMS tracker) and to cool the electronics and the silicon sensors. Nuclear interactions of pions and other hadrons in this material reduce significantly the tracking efficiency for these particles. In addition, this material leads to photon conversion and bremsstrahlung which adversely affect the measurement accuracy of the electromagnetic calorimeter. It was therefore a requirement to keep the amount of this material to a minimum.

3.1.2 Overview of the tracker layout

A schematic drawing of the CMS tracker is shown in figure 3.1. At radii of $4.4, 7.3$ and $10.2\,$cm, three cylindrical layers of hybrid pixel detector modules surround the interaction point. They are complemented by two disks of pixel modules on each side. The pixel detector delivers three high precision space points on each charged particle trajectory. It is described in detail in section 3.2. In total the pixel detector covers an area of about $1\,m^2$ and has 66 million pixels.

The radial region between $20\,$cm and $116\,$cm is occupied by the silicon strip tracker, which is described in detail in section 3.3. It is composed of three different subsystems. The Tracker Inner Barrel and Disks (TIB/TID) extend in radius towards $55\,$cm and are composed of 4 barrel layers, supplemented by 3 disks at each end. TIB/TID delivers up to 4 r-ϕ measurements on a trajectory using $320\,\mu$m thick silicon micro-strip sensors with their strips parallel to the beam axis in the barrel and radial on the disks. The strip pitch is $80\,\mu$m on layers 1 and 2 and $120\,\mu$m on layers 3 and 4 in the TIB, leading to a single point resolution of $23\,\mu$m and $35\,\mu$m, respectively. In the TID the mean pitch varies between $100\,\mu$m and $141\,\mu$m. The TIB/TID is surrounded by the Tracker Outer Barrel (TOB). It has an outer radius of $116\,$cm and consists of 6 barrel layers of $500\,\mu$m thick micro-strip sensors with strip pitches of $183\,\mu$m on the first 4 layers and $122\,\mu$m on

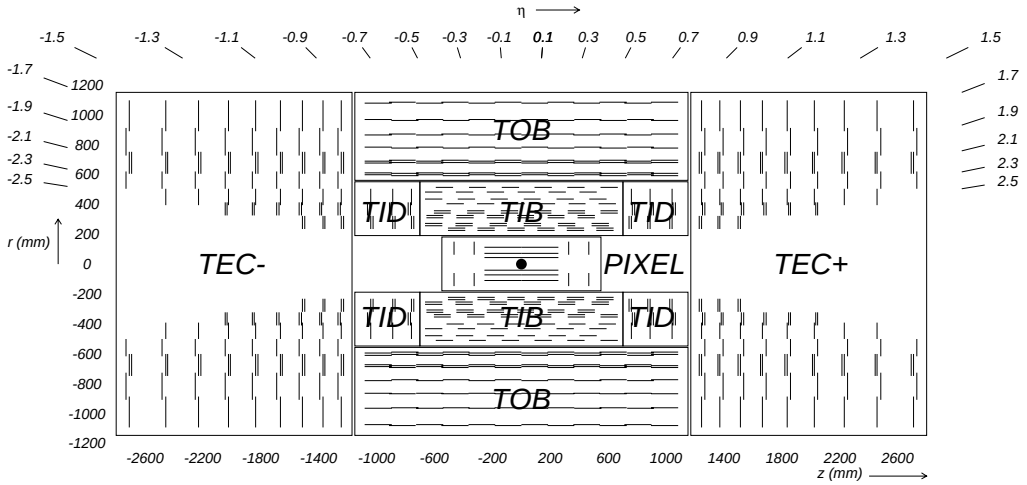

Figure 3.1: Schematic cross section through the CMS tracker. Each line represents a detector module. Double lines indicate back-to-back modules which deliver stereo hits.

layers 5 and 6. It provides another 6 r-ϕ measurements with single point resolution of 53 μm and 35 μm, respectively. The TOB extends in z between $\pm$118 cm. Beyond this z range the Tracker EndCaps (TEC+ and TEC- where the sign indicates the location along the z axis) cover the region 124 cm $< |z| <$ 282 cm and 22.5 cm $< |r| <$ 113.5 cm. Each TEC is composed of 9 disks, carrying up to 7 rings of silicon micro-strip detectors (320 μm thick on the inner 4 rings, 500 μm thick on rings 5-7) with radial strips of 97 μm to 184 μm average pitch. Thus, they provide up to 9 ϕ measurements per trajectory.

In addition, the modules in the first two layers and rings, respectively, of TIB, TID, and TOB as well as rings 1, 2, and 5 of the TECs carry a second micro-strip detector module which is mounted back-to-back with a stereo angle of 100 mrad in order to provide a measurement of the second co-ordinate (z in the barrel and r on the disks). The achieved single point resolution of this measurement is 230 μm and 530 μm in TIB and TOB, respectively, and varies with pitch in TID and TEC. This tracker layout ensures at least ≈ 9 hits in the silicon strip tracker in the full range of $|\eta| < 2.4$ with at least ≈ 4 of them being two-dimensional measurements (figure 3.2). The ultimate acceptance of the tracker ends at $|\eta| \approx 2.5$. The CMS silicon strip tracker has a total of 9.3 million strips and 198 m^2 of active silicon area.

Figure 3.3 shows the material budget of the CMS tracker in units of radiation length. It increases from 0.4 X_0 at $\eta \approx 0$ to about 1.8 X_0 at $|\eta| \approx 1.4$, beyond which it falls to about 1 X_0 at $|\eta| \approx 2.5$.

3.1.3 Expected performance of the CMS tracker

For single muons of transverse momenta of 1, 10 and 100 GeV figure 3.4 shows the expected resolution of transverse momentum, transverse impact parameter and longitudinal impact parameter, as a function of pseudorapidity [17]. For high momentum tracks (100 GeV) the transverse momentum resolution is around $1-2\%$ up to $|\eta| \approx 1.6$, beyond which it degrades due to the reduced lever arm. At a transverse momentum of 100 GeV multiple scattering in the tracker material accounts for 20 to

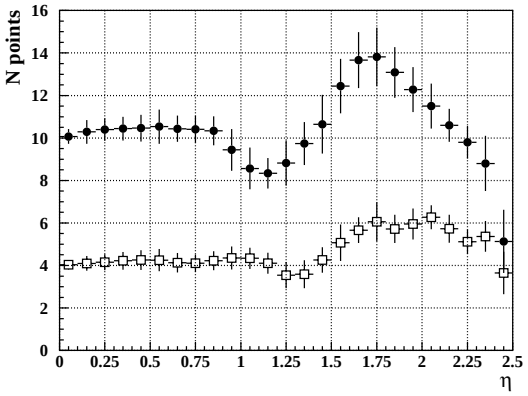

Figure 3.2: Number of measurement points in the strip tracker as a function of pseudorapidity η. Filled circles show the total number (back-to-back modules count as one) while open squares show the number of stereo layers.

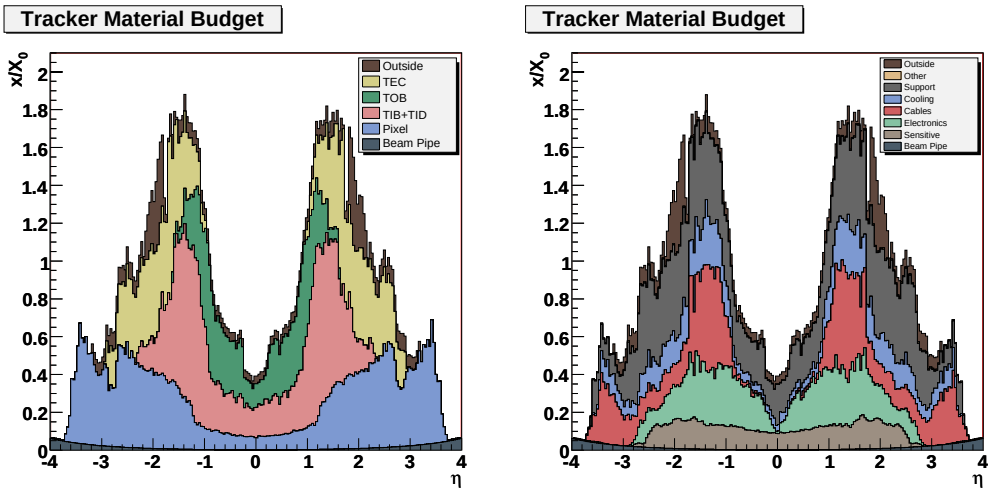

Figure 3.3: Material budget in units of radiation length as a function of pseudorapidity η for the different sub-detectors (left panel) and broken down into the functional contributions (right panel).

30% of the transverse momentum resolution while at lower momentum it is dominated by multiple scattering. The transverse impact parameter resolution reaches $10\,\mu$m for high p_T tracks, dominated by the resolution of the first pixel hit, while at lower momentum it is degraded by multiple scattering (similarly for the longitudinal impact parameter). Figure 3.5 shows the expected track reconstruction efficiency of the CMS tracker for single muons and pions as a function of pseudorapidity. For muons, the efficiency is about 99% over most of the acceptance. For $|\eta| \approx 0$ the efficiency decreases slightly due to gaps between the ladders of the pixel detector at $z \approx 0$. At high η the efficiency drop is mainly due to the reduced coverage by the pixel forward disks. For pions and hadrons in general the efficiency is lower because of interactions with the material in the tracker.

2008 JINST 3 S08004

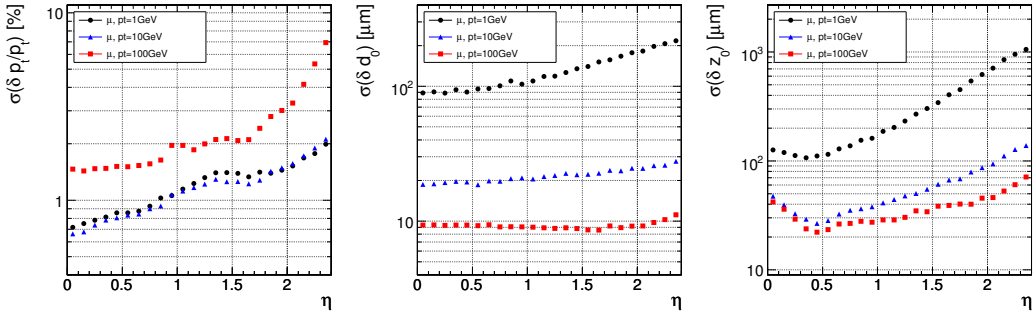

Figure 3.4: Resolution of several track parameters for single muons with transverse momenta of 1, 10 and 100 GeV: transverse momentum (left panel), transverse impact parameter (middle panel), and longitudinal impact parameter (right panel).

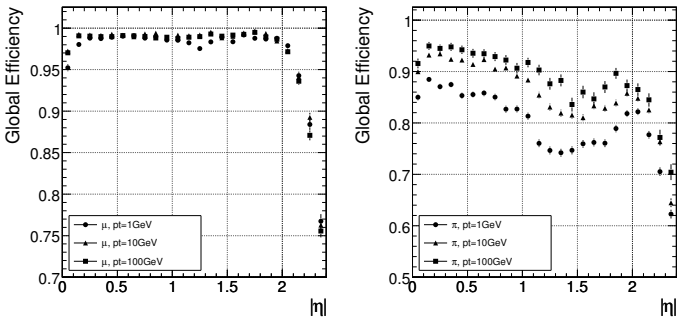

Figure 3.5: Global track reconstruction efficiency for muons (left panel) and pions (right panel) of transverse momenta of 1, 10 and 100 GeV.

3.1.4 Tracker system aspects

All elements of the CMS tracker are housed in the tracker support tube, which is suspended on the HCAL barrel. The tracker support tube is a large cylinder 5.30 m long with an inner diameter of 2.38 m. The 30-mm-thick wall of the cylinder is made by two 950-1/T300 carbon fiber composite skins, 2 mm in thickness, sandwiching a 26-mm-high Nomex core. Over the entire length of the tube's inner surface, two carbon fiber rails are attached on the horizontal plane. The tracker outer barrel (TOB) and both endcaps (TEC+ and TEC-) rest on these rails by means of adjustable sliding pads. The tracker inner barrel and disks (TIB/TID) are in turn supported by the TOB. The angle between the guiding elements of these rails is controlled to better than 0.183 mrad, corresponding to a parallelism between the guides better than ± 0.5 mm in all directions over the full length.

An independent support and insertion system for the pixel detectors, the central section of the beam pipe and the inner elements of the radiation monitor system spans the full length of the tracker at its inner radius. This is composed of three long carbon fiber structures, joined together during tracker assembly to form two continuous parallel planes, on which precision tracks for the installation, support and positioning of each element are machined. The central element is a 2266.5-mm-long and 436-mm-wide cylinder which is connected with flanges to the TIB/TID detector. This element provides support and accurate positioning to the pixel detectors. Two 2420-

2008 JINST 3 S08004

mm-long side elements are coupled to it only by very precise pinned connections, bridging the gap between the faces of the TIB/TID and the closing flanges of the tracker without direct contact to the TEC detectors. These side elements are therefore structurally decoupled from the silicon strip detectors and can be installed and removed at any time with no impact on the strip detectors. They serve several purposes: they provide support and alignment features for the central section of the beam pipe, they allow the installation of the inner elements of the radiation monitor system, and they are used for installation and removal of all the components permanently or temporarily housed in the inner region of the tracker: beam pipe, bake-out equipment, pixel barrel, pixel disks and radiation monitor. This system of permanent tracks, light but very stiff and stable, installed in the core of the tracker will allow for the quickest possible intervention in this region during maintenance, inducing no disturbance to the volume occupied by the silicon strip detectors. This feature will be extremely valuable after some years of operation, when activation of components and radiation damage on sensors will start becoming an issue.

The outer surface of the tracker tube faces the electromagnetic calorimeter, which is operated at room temperature and requires good temperature stability. The surface of the electromagnetic calorimeter must be kept at (18 ± 4)°C while the tracker volume needs to be cooled to below -10°C. In order to achieve this thermal gradient over a very limited radial thickness, the inside surface of the tracker support tube is lined with an active thermal screen. It ensures a temperature below -10°C inside the tracker volume even when the sub-detectors and their cooling are switched off, while maintaining a temperature above $+12$°C on the outer surface of the support tube in order to avoid condensation. It also reduces the thermal stress across the support tube structure. The thermal screen consists of 32 panels. On the inside, cold fluid is circulated in a thin aluminium plate whilst, separated by 8 mm of Rohacell foam, several polyimide-insulated resistive circuits are powered to heat up the outer surface to the required temperature. The system is feed-back controlled, based on 64 temperature sensors.

The total power dissipation inside the tracker volume is expected to be close to 60 kW. Mainly for robustness in operation, the CMS tracker is equipped with a mono-phase liquid cooling system. The liquid used for refrigeration of the silicon strip and pixel detector as well as the thermal screen is C_6F_{14}. It has a sufficiently low viscosity even at the lowest required temperature, excellent behaviour under irradiation and is extremely volatile (with practically no residues) thus minimizing eventual damages from accidental leaks. The cooling system provides up to 77 m^3/hour of C_6F_{14} liquid to the tracker, at a temperature of down to -35°C and with a pressure drop of up to 8 bar. This corresponds to a cooling capacity of up to 128 kW.

The full tracker volume (about 25 m^3) is flushed with pre-chilled dry nitrogen gas at a rate of up to one volume exchange per hour.

3.2 Pixel detector

3.2.1 Pixel system general

The pixel system is the part of the tracking system that is closest to the interaction region. It contributes precise tracking points in r-ϕ and z and therefore is responsible for a small impact parameter resolution that is important for good secondary vertex reconstruction. With a pixel cell

2008 JINST 3 S08004

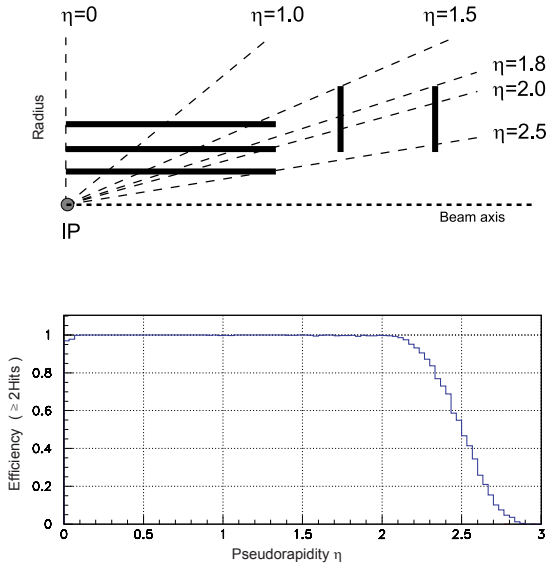

Figure 3.6: Geometrical layout of the pixel detector and hit coverage as a function of pseudorapidity.

size of $100 \times 150\ \mu\text{m}^2$ emphasis has been put on achieving similar track resolution in both r-ϕ and z directions. Through this a 3D vertex reconstruction in space is possible, which will be important for secondary vertices with low track multiplicity. The pixel system has a zero-suppressed read out scheme with analog pulse height read-out. This improves the position resolution due to charge sharing and helps to separate signal and noise hits as well as to identify large hit clusters from overlapping tracks.

The pixel detector covers a pseudorapidity range $-2.5 < \eta < 2.5$, matching the acceptance of the central tracker. The pixel detector is essential for the reconstruction of secondary vertices from b and tau decays, and forming seed tracks for the outer track reconstruction and high level triggering. It consists of three barrel layers (BPix) with two endcap disks (FPix). The 53-cm-long BPix layers will be located at mean radii of 4.4, 7.3 and 10.2 cm. The FPix disks extending from $\approx$6 to 15 cm in radius, will be placed on each side at z=$\pm$34.5 and z=$\pm$46.5 cm. BPix (FPix) contain 48 million (18 million) pixels covering a total area of 0.78 (0.28) m^2. The arrangement of the 3 barrel layers and the forward pixel disks on each side gives 3 tracking points over almost the full η-range. Figure 3.6 shows the geometric arrangement and the hit coverage as a function of pseudorapidity η. In the high η region the 2 disk points are combined with the lowest possible radius point from the 4.4 cm barrel layer.

The vicinity to the interaction region also implies a very high track rate and particle fluences that require a radiation tolerant design. For the sensor this led to an n+ pixel on n-substrate detector design that allows partial depleted operation even at very high particle fluences. For the barrel layers the drift of the electrons to the collecting pixel implant is perpendicular to the 4 T magnetic field of CMS. The resulting Lorentz drift leads to charge spreading of the collected signal charge over more than one pixel. With the analog pulse height being read out a charge interpolation allows

to achieve a spatial resolution in the range of 15–20 μm. The forward detectors are tilted at 20° in a turbine-like geometry to induce charge-sharing. The charge-sharing is mainly due to the geometric effect of particles entering the detector at an average angle of 20° away from normal incidence [19]; charge-sharing is also enhanced by the $\vec{E} \times \vec{B}$ drift. A position resolution of approximately 15 μm in both directions can be achieved with charge-sharing between neighbouring pixels. The reduction in the depletion depth or the increase in bias voltage will lead to a reduction of charge-sharing and therefore a degradation of the spatial resolution with radiation damage.

In order to allow a replacement of the innermost layers the mechanics and the cabling of the pixel system has been designed to allow a yearly access if needed. At full LHC luminosity we expect the innermost layer to stay operational for at least 2 years. The 3 layer barrel mechanics as well as the forward disks are divided into a left and a right half. This is required to allow installation along the beam pipe and to pass beyond the beam pipe support wires at z=$\pm$ 1632 mm. The 6 individual mechanical pieces are referenced to each other through precisely machined rails inside the TIB cylinder. Power, cooling, the optical controls as well as the optical read-out lines are brought to the detector through supply tube shells. In case of the barrel pixel system the supply tubes have a flexible connection that needs to bend by a few degrees during insertion following the slightly curved rails around the beam pipe support ring.

The pixel system is inserted as the last sub-detector of CMS after the silicon strip tracker has been installed and after the central section of the beam pipe has been installed and baked out.

3.2.2 Sensor description

Technological choices

The sensors for the CMS-pixel detector adopt the so called *n*-on-*n* concept. The pixels consist of high dose n-implants introduced into a high resistance n-substrate. The rectifying pn-junction is placed on the back side of the sensor surrounded by a multi guard ring structure. Despite the higher costs due to the double sided processing this concept was chosen as the collection of electrons ensures a high signal charge at moderate bias voltages ($< 600\,\mathrm{V}$) after high hadron fluences. Furthermore the double sided processing allows a guard ring scheme keeping all sensor edges at ground potential.

The isolation technique applied for the regions between the pixel electrodes was developed in close collaboration with the sensor vendors. Open p-stops [20] were chosen for the disks and moderated p-spray [21] for the barrel. Both types of sensors showed sufficient radiation hardness during an extensive qualification procedure including several test beams [22, 23].

Disk sensors

The disk sensors use the p-stop technique for interpixel isolation. To maximize the charge collection efficiency and minimize the pixel capacitance within the design rules of the vendor a width of 8 μm for the p-stop rings and a distance of 12 μm between implants was chosen. Figure 3.7 shows a photograph of 4 pixel cells. The open ring p-stops, the bump-bonding pad and the contact between the aluminium and the implanted collecting electrode are highlighted.

2008 JINST 3 S08004

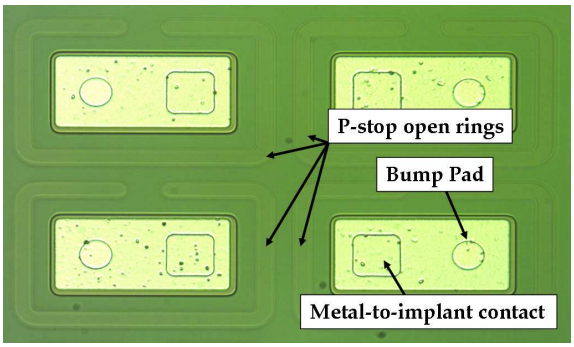

Figure 3.7: Picture of four pixels in the same double column for a pixel disk sensor.

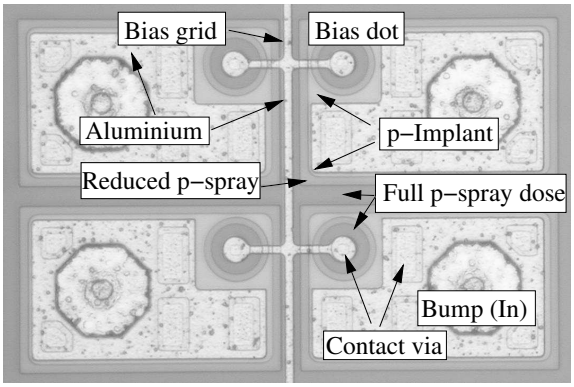

Figure 3.8: Photograph of four pixel cells. The Indium bumps are already deposited but not yet reflown.

The opening on the p-stop rings provides a low resistance path until full depletion is reached to allow IV (current-voltage) characterization of the sensor on wafer and a high resistance path when the sensor is over-depleted (10–20 V over-depletion) to assure interpixel isolation.

The process is completely symmetric with five photolithographic steps on each side to minimize the mechanical stress on the silicon substrate and the potential bowing of the diced sensors.

The sensors were all fabricated in 2005 on 4 inch wafers. The depletion voltage is 45–50 V and the leakage current is less than 10 nA per cm^2. The 7 different sensor tiles needed to populate a disk blade, ranging from 1×2 read-out chips (ROCs) to 2×5 ROCs, are implemented on a single wafer.

A production yield higher than 90% has been achieved and 150 good sensors for each of the seven flavours are available to the project for module assembly.

Barrel sensors

The sensors for the pixel barrel use the moderated p-spray technique for interpixel isolation. A photograph of four pixels in a barrel sensor is shown in figure 3.8. Most area of a pixel is covered with the collecting electrode formed by the n-implant. The gap between the n-implants is kept small ($20\,\mu$m) to provide a homogeneous drift field which leads to a relatively high capacitance of the order of 80-100 fF per pixel.

In one corner of each pixel the so called *bias dot* is visible. They provide a high resistance punch-through connection to all pixels which allows on-wafer IV measurements which are important to exclude faulty sensors from the module production.

The dark *frame* around the pixel implants visible in figure 3.8 indicates the opening in the nitride covering the thermal oxide. In this region the p-spray dose reaches the full level. Close to the lateral pn-junction between the pixel implant and the p-sprayed inter-pixel region the boron dose is reduced.

The sensor shown in figure 3.8 has undergone the bump deposition process. The Indium bumps are visible as roughly 50 μm wide octagons.

The sensors are processed on n-doped DOFZ-silicon [24] with $\langle 111 \rangle$ orientation and a resistivity of about 3.7 kΩcm (after processing). This leads to a full depletion voltage of 50-60 V of the 285 μm thick sensors. All wafers for the production of the barrel sensors come from the same silicon ingot to provide the best possible homogeneity of all material parameters.

The pixel barrel requires two different sensor geometries, 708 full (2×8 ROCs) and 96 half modules (1×8 ROCs). They were processed in 2005 and 2006 using two different mask sets.

3.2.3 Pixel detector read-out

System overview

The pixel read-out and control system [25] consists of three parts: a read-out data link from the modules/blades to the pixel front end driver (pxFED), a fast control link from the pixel front end controller (pFEC) to the modules/blades and a slow control link from a standard FEC to the supply tube/service cylinder. The latter is used to configure the ASICs on the supply tube/service cylinder through a I^2C protocol. Figure 3.9 shows a sketch of the system.

The front end consists of a Token Bit Manager (TBM) chip which controls several read-out chips (ROCs). The pFEC sends the 40MHz clock and fast control signals (e.g. trigger, reset) to the front end and programs all front end devices. The pxFED receives data from the front end, digitizes it, formats it and sends it to the CMS-DAQ event builder. The pFEC, FEC and pxFED are VME modules located in the electronics room and are connected to the front end through 40 MHz optical links. The various components are described in the following sections.

Read-out chip

Sensor signals are read out by ROCs bump bonded to the sensors. A ROC is a full custom ASIC fabricated in a commercial 0.25-μm 5-metal-layer CMOS process and contains 52×80 pixels [26]. Its main purposes are:

- Amplification and buffering of the charge signal from the sensor.
- Zero suppression in the pixel unit cell. Only signals above a certain threshold will be read out. This threshold can be adjusted individually for each pixel by means of four trim bits. The trim bits have a capacitive protection against single event upset (SEU), which has shown to reduce SEUs by 2 orders of magnitude [26]. The mean threshold dispersion after trimming at T$= -10$°C is 90 electrons equivalent with a noise of 170 electrons.

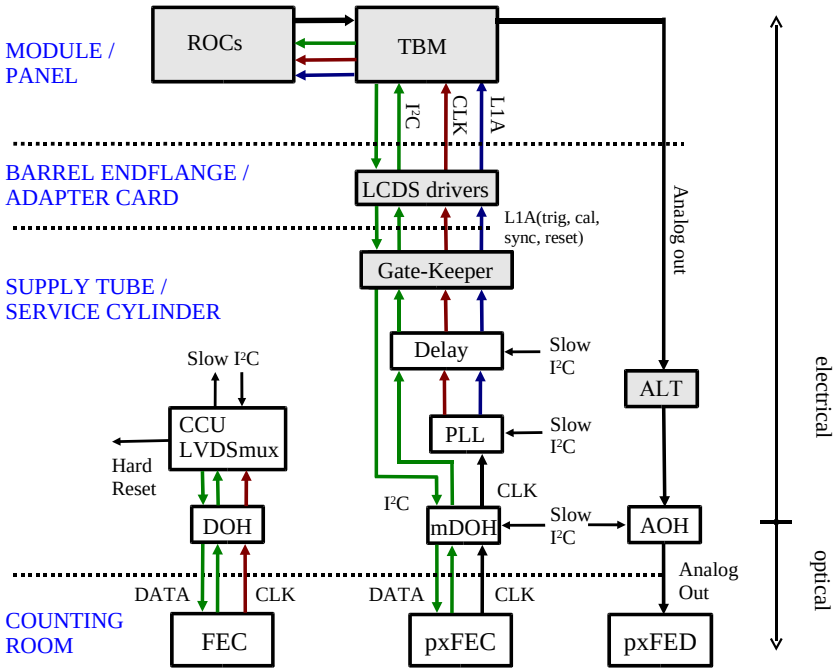

Figure 3.9: Block diagram of the pixel control and read-out system.

- Level 1 trigger verification. Hit information without a corresponding L1 trigger is abandoned.
- Sending hit information and some limited configuration data (analog value of last addressed DAC) to the TBM chip. Pixel addresses are transferred as 6 level analog encoded digital values within 5 clock cycles (125ns) while the pulse height information is truly analog.
- Adjusting various voltage levels, currents and offsets in order to compensate for chip-to-chip variations in the CMOS device parameters. There are a total of 29 DACs on the chip.

The ROC needs two supply voltages of 1.5 V and 2.5 V. There are 6 on chip voltage regulators. They compensate for differences in supply voltage due to voltage drops in module cables of different lengths, improve AC power noise rejection and strongly reduce intermodule cross-talk. An on-chip temperature sensor allows the monitoring of the module temperature online. The ROC is controlled through a modified I^2C interface running at 40 MHz. The configuration data can be downloaded without stopping data acquisition.

There are a few architecture inherent data loss mechanisms. The particle detection inefficiency has been measured in a high-rate pion beam. It is in fairly good agreement with expectations and reaches 0.8%, 1.2% and 3.8% respectively for the three layers at a luminosity of 10^{34} $s^{-1}cm^{-2}$ and 100 kHz L1 trigger rate.

The power consumption depends on the pixel hit rate. At the LHC design luminosity, the ROC contributes with 34 μW per pixel about 88% (62%) to the total pixel detector front end power budget before (after) the detector has received a total fluence of $6\times10^{14}/cm^2$.

Token Bit Manager chip

The TBM [27] controls the read-out of a group of pixel ROCs. The TBM is designed to be located on the detector near to the pixel ROCs. In the case of the barrel, they will be mounted on the detector modules and will control the read-out of 8 or 16 ROCs depending upon the layer radius. In the case of the forward disks, they will be mounted on the disk blades and will control the read-out of 21 or 24 ROCs depending on blade side. A TBM and the group of ROCs that it controls will be connected to a single analog optical link over which the data will be sent to the front end driver, a flash ADC module located in the electronics house. The principal functions of the TBM include the following:

- It will control the read-out of the ROCs by initiating a token pass for each incoming Level-1 trigger.
- On each token pass, it will write a header and a trailer word to the data stream.
- The header will contain an 8 bit event number and the trailer will contain 8 bits of error status. These will be transferred as 2 bit analog encoded digital.
- It will distribute the Level-1 triggers, and clock to the ROCs.

Each arriving Level-1 trigger will be placed on a 32-deep stack awaiting its associated token pass. Normally the stack will be empty but is needed to accommodate high burst rates due to noise, high track density events, or trigger bursts. Since there will be two analog data links per module for the inner two layers of the barrel, the TBMs will be configured as pairs in a Dual TBM Chip. In addition to two TBMs, this chip also contains a Control Network. The Hub serves as a port addressing switch for control commands that are sent from the DAQ to the front end TBMs and ROCs. These control commands will be sent over a digital optical link from a front end controller in the electronics house to the front end Hubs. The commands will be sent using a serial protocol, running at a speed of 40 MHz. This high speed is mandated by the need to rapidly cycle through a refreshing of the pixel threshold trim bits that can become corrupted due to single event upsets. There are four external, write only ports on each Hub for communicating with the ROCs and there is one internal read/write port for communicating with the TBMs within the chip. The first byte of each command will contain a 5-bit Hub address and a 3-bit port address. When a Hub is addressed, it selects the addressed port, strips off the byte containing the Hub/port address and passes the remainder of the command stream unmodified onto the addressed port. The outputs of the external ports consist of two low voltage differential lines for sending clock and data.

Analog chain

The hit information is read out serially through analog links in data packets containing all hits belonging to a single trigger. Within such packets a new analog value is transmitted every 25 ns and digitized in the Front End Driver (pxFED) at the same rate. Each pixel hit uses 6 values, or 150 ns. Five values are used to encode the address of a pixel inside a ROC and the sixth value represents the signal charge. Only the charge signals are truly analog while headers and addresses are discrete levels generated by DACs. No ROC IDs are sent but every ROC adds a header, whether

2008 JINST 3 S08004

it has hits or not in order to make the association of hits to ROCs possible. The sequential read-out is controlled by a token bit which is emitted by the TBM, passed from ROC to ROC and back to the TBM. The differential electrical outputs of the ROCs are multiplexed by the TBM onto either one or two output lines. On the same lines the TBM transmits a header before starting the ROC read-out. After receiving the token back from the last ROC in the chain the TBM sends a trailer containing status information. From the TBM to the end ring of the pixel barrel the read-out uses the module Kapton cable. The Kapton cable has a ground mesh on the back side and the differential analog lines are separated by quiet lines from the fast digital signals. Nevertheless, cross-talk from LVDS signals was found to be unacceptable and a low swing digital protocol is being used instead. On the end ring the analog signals are separated from the digital and all analog signals of the sector are sent on a separate Kapton cable to a printed circuit board that houses the Analog Optical Hybrids (AOH). The signal path between TBM and AOH is designed with a constant impedance of 40 Ω and terminated on the AOH. The optical links of the pixel system are identical to those used in the silicon strip tracker. An ASIC that adapts the output levels of the pixel modules to those expected by the laser driver has been added to the AOH of the pixel system. A clean identification of the six levels used for encoding pixel addresses is crucial for the reconstruction of hits. The ratio of RMS width to separation of the digitized levels after the full read-out chain is 1:30. The rise-time at the digitizer input is 3 ns which makes corrections from previously transmitted levels negligible. The full read-out chain adds a noise equivalent to 300 electrons to the analog pulse height, dominated by baseline variations of the laser drivers.

Front End Driver

Optical signals from the pixel front end electronics (ROCs and TBMs) are digitized using the pixel Front End Digitizer (pxFED). A pxFED is a 9U VME module. It has 36 optical inputs each equipped with an optical receiver and an ADC. The ADC converts at LHC frequency supplied by the TTC system which can be adjusted by an individually programmed phase shift (16 steps within 25 ns) for precise timing. A programmable offset voltage to compensate bias shifts can also be set. The output of the 10 bit-ADC is processed by a state machine to deliver pixel event fragments consisting of header, trailer, input channel number, ROC numbers, double column numbers and addresses and amplitudes of hit pixels all at a subject-dependent resolution of 5 to 8 bits. Event fragments are strobed into FIFO-1 (1k deep $\times$ 36 bit wide) which can be held on demand to enable read-out via VME. In normal processing mode FIFO-1 is open and data of 4 (5) combined input channels are transmitted to 8 FIFO-2 memories (8k $\times$ 72 bits). In order to determine thresholds and levels required for the state machine, FIFO-1 can alternatively be operated in a transparent mode making unprocessed ADC output data available. The output from FIFO-2 is clocked into two FIFO-3 memories (8k $\times$ 72 bits) whose outputs are combined to provide the data now at a frequency of 80 MHz (twice the common operating pxFED-frequency) to the S-Link interface acting as a point-to-point link with the CMS-DAQ system. Parallel to the data flow spy FIFOs are implemented (restricted in size) to hold selected event fragments and make them available for checking data integrity. Error detection takes place in the data stream from FIFO-1 to FIFO-2 and corresponding flags are embedded in the event trailer and also accessible from VME. A selected DAC output from each ROC (on default representing the ROC's temperature) is available as well.

2008 JINST 3 S08004

In addition, errors are directly transmitted to the CMS-TTS system using a dedicated connector on the S-Link supplementary card. A histogramming feature has been implemented to monitor the rate of double column hits. This histogram is intended to be read out via VME periodically to check for dead or overloaded columns. The pxFED houses an internal test system which, when enabled, replaces the normal ADC input by a pattern of 256 clocked analog levels simulating a normal pixel event. There are three test DACs (10-bit) available to generate such a pattern meaning that every third input channel receives the same simulated event. This test system allows to test most of the features of the pxFED without the need of external optical input signals. All FIFOs, the state machine with its adjustable parameters, the VME protocol, error detection and histogramming features are integrated into several FPGAs mounted on daughter cards making the pxFED flexible to changes and improvements. The corresponding firmware can be downloaded via VME or using a JTAG bus connector mounted on the mother board. The whole pixel read-out system will consist of 40 pxFED modules (32 for the barrel and 8 for the forward) set up in three 9U VME crates located in the electronics room. Individual modules can be accessed by VME geographical addressing.

Front End Controller

The Pixel Front End Controller (pFEC) supplies clock and trigger information to the front end, and provides a data path to the front end for configuration settings over a fiber optic connection. The pFEC uses the same hardware as the standard CMS FEC-CCS [28]. The firmware which defines the behaviour of the mezzanine FEC (mFEC) module has been replaced by a pixel specific version, converting the FEC into a pFEC. Each mFEC board becomes two command links to the front end. The Trigger Encoder performs all trigger transmission functions, encoding TTC triggers to match the pixel standard, block triggers to a given channel, generate internal triggers, either singly, or continuously, for testing purposes. Within each command link are a one kilobyte output buffer for data transmission, and a two kilobyte input buffer for data receiving. All data, whether write or read operations, are retransmitted back from the front end for possible verification. To minimize the VME data transfer time, the pFEC uses several data transfer modes. When transferring pixel trim values to the front end, the pFEC calculates the row number information for a given column of pixels on the fly. This results in nearly a 50% reduction in the time required to transfer trim values over VME to a given command link buffer. In this way, the entire pixel front end trims can be reloaded in 12 s. Another 2 s are used to load the other configuration registers, for a total of 14 s to reload the front end completely. This column mode is also the reason that the return buffer is twice as big as the transmit buffer. The return buffer receives the row number as well as the trim value for each pixel. Once data is loaded into an output buffer, the transfer may be initiated either by computer control, or by a signal from the TTC system. Since single event upsets are expected to occur in the front end registers, it is anticipated that periodic updates will be necessary. Since updating the front end may disrupt data taking, it is preferable to perform small updates synchronized to orbit gaps or private orbits. This is done through the TTC initiated downloads. For transmission verification purposes, the number of bytes transmitted is compared to the number of bytes returned from the front end. Also, the returning Hub/port address is compared to the transmitted address. Status bits are set with the results of this comparison, and these values are stored, for possible review, should an error condition occur.

2008 JINST 3 S08004

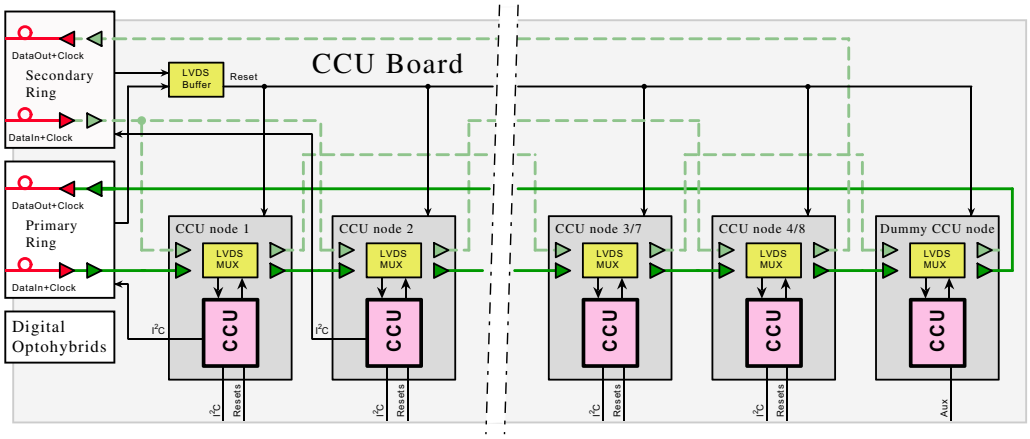

Figure 3.10: Block diagram of the Pixel front end control system. Note that the total number of CCU nodes is 9 for the BPix and 5 for the FPix.

The detector front end control system

The CMS Pixel detector front end control system for both the barrel (BPix) and the forward (FPix) detectors consists of four communication and control unit boards (CCU Boards). Each CCU board controls a quarter of the detector with eight Barrel read-out sectors or twelve Forward port cards. Figure 3.10 shows the block diagram of a CCU Board. The same ring topology configured as a local area network as in the silicon strip tracker is used. The front end controller (FEC) module is the master of the network and uses two optical fibers to send the timing and data signals to a number of slave CCU nodes, and another two fibers to receive return communication traffic. The two receiver channels on the digital optohybrid (DOH) transmit the 40 MHz clock and control data at 40 Mbit/s in the direction from the FEC to the ring of communication and control units (CCUs). The two transmitter channels send clock and data back to the FEC from the ring of CCUs. The CCU is the core component developed for the slow control, monitoring and timing distribution in the tracking system [29]. To improve system reliability against a single component failure a redundant interconnection scheme based on doubling signal paths and bypassing of faulty CCUs is implemented. An additional "dummy" CCU node allows to mitigate a single DOH failure preserving complete functionality. A CCU node failure leads to a loss of communication to all electronics attached to that CCU. The first two CCU nodes in the ring provide also the I^2C data channels necessary to control the digital optohybrids on the CCU boards.

In the BPix each read-out sector is controlled by a separate CCU node. Eight active and one dummy CCU node build a single control ring. One I^2C data channel is used to access and control the front end read-out electronics and three output channels generate the necessary signals to reset the digital and the analog optohybrids as well as the read-out chips (ROCs) in one read-out sector. The FPix control ring consists of four active and one dummy CCU node. Each of the active CCU nodes control 3 port cards, which constitute a 45° sector in the detector coverage at one end. A connection between a CCU and a port card includes a bi-directional 100 KHz I^2C communication channel and two reset signals. One reset signal is for the port card electronics, and the other one goes to the read-out chips on the detector panels.

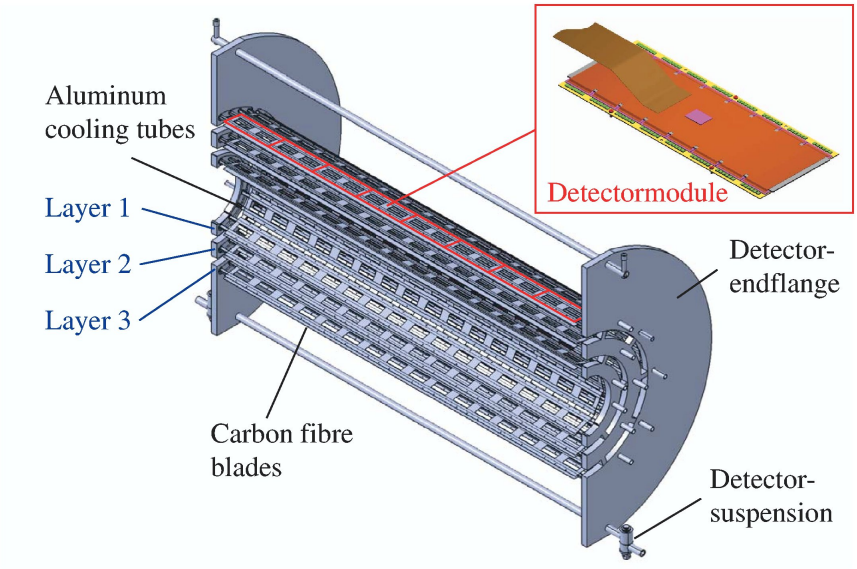

Figure 3.11: Complete support structure half shell with the three detector layers.

3.2.4 The pixel barrel system

The pixel barrel system as installed inside CMS comprises the barrel itself, i.e. detector modules mounted on their cylindrical support structure, as well as supply tubes on both sides. The barrel with its length of 570 mm is much shorter than the Silicon Strip Tracker inside which it is installed. Supply tubes carry services along the beam pipe from patch panels located outside of the tracker volume to the barrel. The supply tubes also house electronics for read-out and control. The length of the full system is 5.60 m. Support structure and supply tubes are split vertically to allow installation in the presence of the beam-pipe and its supports. Electrically the $+z$ and $-z$ sides of the barrel are separated. Each side is divided in 16 sectors which operate almost independently, sharing only the slow control system.

Pixel barrel support structure

The detector support structure for the three layers at the radii of four, seven and eleven centimeters equipped with silicon pixel modules has a length of 570 mm ranging from -285 mm to $+285$ mm closest to the CMS interaction region. Figure 3.11 shows a sketch of a complete support structure half shell.

Aluminium cooling tubes with a wall thickness of 0.3 mm are the backbones of the support structure. Carbon fiber blades with a thickness of 0.24 mm are glued onto the top or bottom of two adjacent cooling tubes in such a way that their normal directions alternate pointing either to the beam or away from it. The tubes have trapezoidal cross sections defined by the azimuthal angles of the ladders they hold.

Four to five of these tubes are laser welded to an aluminium container which distributes the cooling fluid. The resulting manifold provides the necessary cooling of the detector modules to

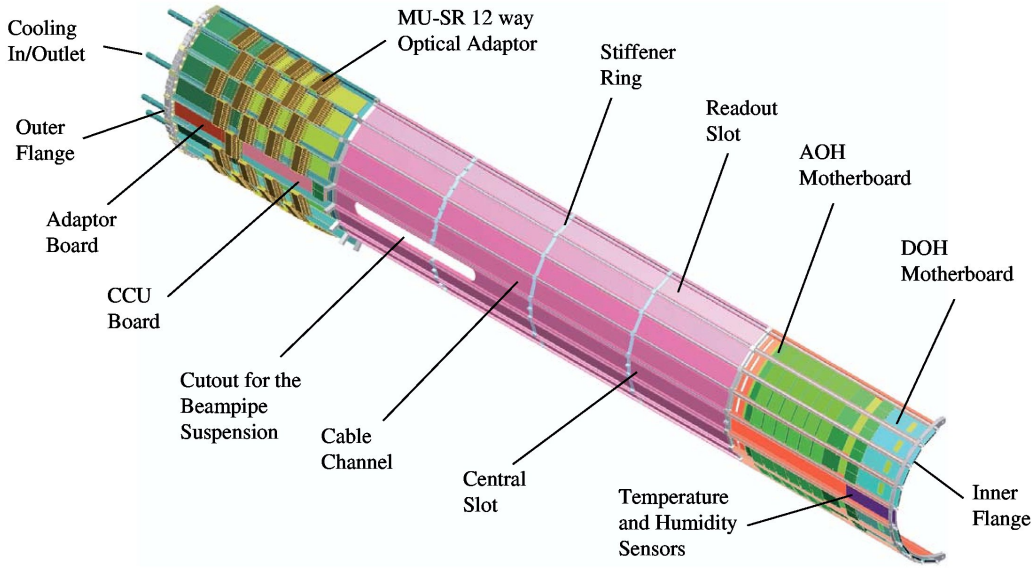

Figure 3.12: Overview of a supply tube half shell.

about $-10\,^{\circ}\mathrm{C}$ with C_6F_{14}. Support frames on both ends, which connect the single segments, build a complete detector layer half shell. These flanges consist of thin fibreglass frames (FR4) that are filled with foam and covered by carbon fibre blades.

Printed circuit boards mounted on the the flanges hold the connectors for the module cables and provide control signal fan-out and power distribution to the individual modules of a sector.

Pixel detector supply tube. The electrical power lines, the electrical control signal and the optical signals as well as the cooling fluid are transferred across the supply tubes to the pixel barrel. The two supply tube parts of a half shell in $+z$ and $-z$ direction have a length of 2204 mm (figure 3.12).

The supporting elements of the basic structure are the stainless steel tubes with a wall thickness of 0.1 mm running along the z-direction connected to the stiffener rings (FR4) and the inner and outer flanges made out of aluminium. The tubes supply the detector with the cooling fluid. The gaps in between are filled with foamed material with the corresponding shape to guarantee the necessary rigidity. All power and slow control leads are embedded in the supply tube body. This allows a clear layout of the wiring and also makes the system more reliable.
The motherboards, which hold the optical hybrids for the analog and digital control links, are installed in the eight read-out slots near the detector on the integrated supply boards. The corresponding boards at the outer ends carry the power adapter boards, which provide the detector power and the bias voltage for this sector. In the central slot the digital communication and control board (CCU Board) is installed. From here the digital control signals are distributed to the individual read-out boards in each of the eight read-out sectors. Here also all slow control signals like temperatures, pressures and the humidity are brought together and connected by the dedicated slow control adapter board to the cables. The optical fibres are installed in the cable channels. The 36 single fibres for the analog read-out and the eight fibres for the digital control of the detector

2008 JINST 3 S08004

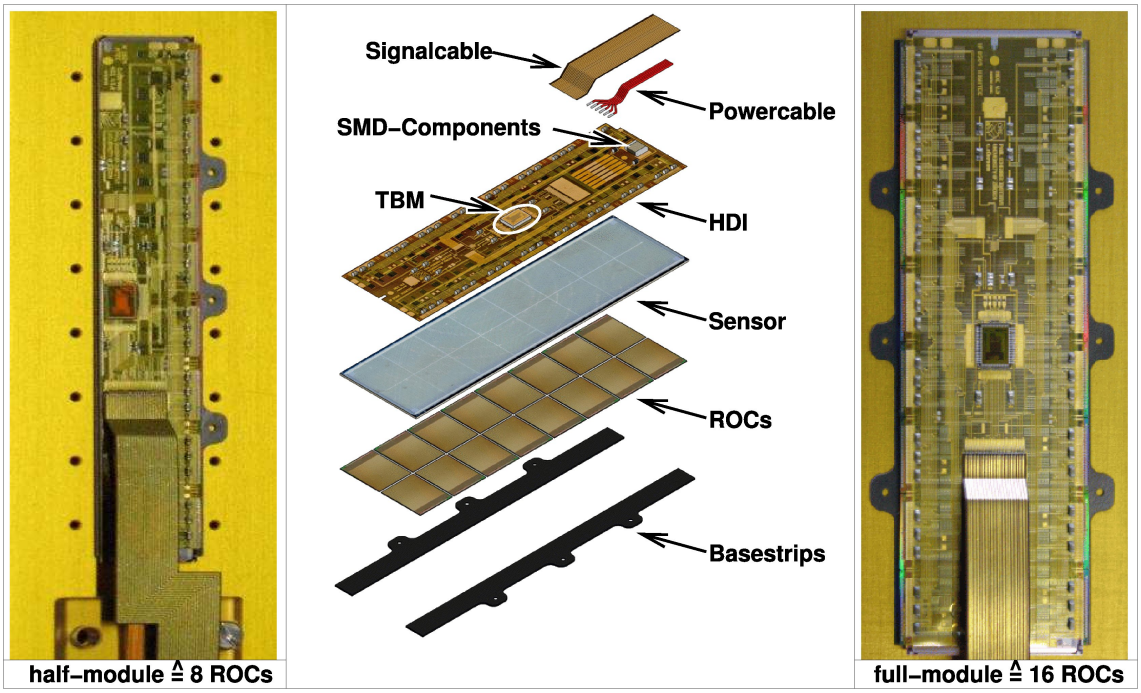

Figure 3.13: Exploded view (middle panel) of a barrel pixel detector full module (right panel) and picture of an assembled half module (left panel).

modules will then be connected through the MUSR connector to the optical ribbon cable. These adapters are mounted at the circumference in the first part of the supply tube. The length of each supply tube is 2204 mm. Only a flexible mechanical connection is made between the barrel and the supply tube.

Pixel barrel detector modules

The barrel part of the CMS pixel detector consists of about 800 detector modules. While the majority of the modules (672) are full modules as seen in figure 3.13 on the right, the edges of the six half-shells are equipped with 16 half-modules each (96 in total, see figure 3.13 on the left).

Geometry and components. A module is composed of the following components (figure 3.13). One or two basestrips made from 250 μm thick silicon nitride provide the support of the module. The front end electronics consists of 8 to 16 read-out chips with 52×80 pixels of size 100×150 μm^2 each, which are bumpbonded to the sensor. The chips are thinned down to 180 μm . The High Density Interconnect, a flexible low mass 3 layer PCB with a trace thickness of 6 μm equipped with a Token Bit Manager chip that controls the read-out of the ROCs, forms the upper layer of a module and distributes signals and power to the chips. The signals are transferred over an impedance matched 2 layer Kapton/copper compound cable with 21 traces and 300 μm pitch. The module is powered via 6 copper coated aluminium wires of 250 μm diameter.

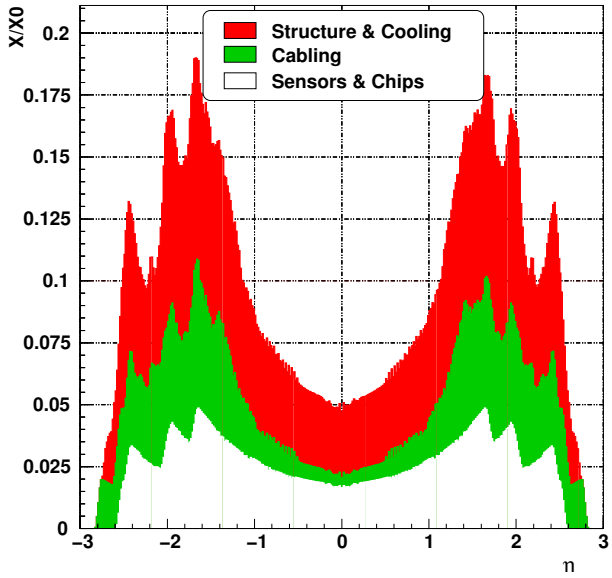

Figure 3.14: Material budget of the pixel barrel in units of radiation length versus rapidity. The plot does not contain contributions from the pixel support cylinder, the supply tube and cabling from the detector end flange to the supply tube.

A completed full-module has the dimensions 66.6×26.0 mm^2, weights 2.2 g plus up to 1.3 g for cables, and consumes 2 W of power. The material of the pixel barrel amounts to 5 percent of a radiation length in the central region. Sensors and read-out chips contribute one third of the material while support structure and cooling fluid contribute about 50 percent. The distribution of material as a function of pseudorapidity is shown in figure 3.14.

3.2.5 The forward pixel detector

The FPix detector consists of two completely separate sections, one on each side of the interaction region. They are located inside the BPix supply tube but are mounted on separate insertion rails. Each section is split vertically down the middle so the detector can be installed around the beam-pipe and its vertical support wire and so it can also be removed for servicing during major maintenance periods without disturbing the beam-pipe. Each of these four sections is called a *half-cylinder*.

Mechanics of a half-cylinder

Each half-cylinder consists of a carbon fiber shell with two half-disks located at its front end, one at 34.5 cm from the IP and the other at 46.5 cm. The half-disks support the actual pixel detectors that extend from 6 cm to 15 cm in radius from the beam.

The half-disk has 12 cooling channels (each in the shape of a “U”) assembled between a half ring shown in figure 3.15. The assembly requires three slightly different types of cooling channels. Each channel is made by Al-brazing two blocks of Al with the channel for the cooling fluid already

Figure 3.15: The FPix half-disk cooling channels mounted in the outer half-ring structure. The turbine-like geometry is apparent. Panels are mounted on both sides of the cooling channels.

machined in the two parts. The brazed parts are then machined to their final shape. The walls of the channels are 0.5 mm thick. The average weight of the channels is 8.21 g.

All channels passed a Helium leak test at 1.33×10^{-8} mbar-litre/s. The pressure drop of the individual cooling channels for a flow of 2600 sccm of dry N_2 is 0.49 $\pm$ 0.02 mbar. Six daisy-chained cooling channels form a *cooling loop*. The pressure drop over a loop (for a flow rate of 1230 sccm) of dry N_2 is 0.96 ± 0.13 mbar. For C_6F_{14} at -20°C with a rate of 12cc-s the pressure drop is 294 mbar.

Each of the twelve cooling channels of a half-disk has trapezoidal beryllium *panels* attached to each side. The panels support the sensors and read-out chips that constitute the actual particle detectors. As explained above, the cooling channels are rotated to form a turbine-like geometry to enhance charge-sharing. The panels are made of 0.5mm beryllium. The beryllium provides a strong, stable and relatively low-mass support for the actual pixel detectors. The cooling channels are supplied with C_6F_{14} at about -15°C. A single cooling channel with panels mounted on both sides forms a subassembly called a *blade*. There are 24 panels, forming 12 blades, in each half-disk.

Powering up the electronics on one blade increases the temperature by $\approx$2°C. The temperature of each ROC is part of the information available for each event. Each panel also has a resistance temperature detection sensor. The pixel sensors have fiducial marks visible with a coordinate measuring machine (CMM). Their position is then related to reference marks mounted on the half-disk units.

After installing the half-disks in the half-cylinder, the disk position is measured relative to the half-cylinder using a CMM and also by photogrammetry. This permits relating the position of the sensors to the CMS detector. The detector is surveyed at room temperature but operated at

2008 JINST 3 S08004

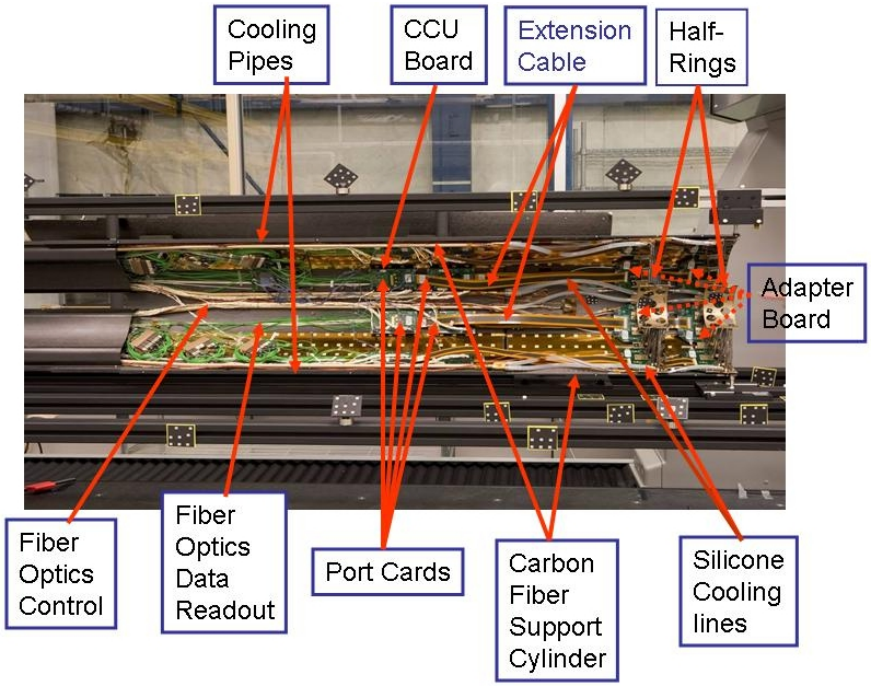

Figure 3.16: Overview of the Forward Pixel half-cylinder. A photograph of the portion of the first production half-cylinder facing the interaction region (IR). The aluminium flange, the filter boards (see below), and the CCU board are not shown. The half-cylinder is mounted in a survey fixture. The carbon fiber cover at the end away from the IR protects the downstream components during insertion of the beam pipe suspension wires that run through a slot in the half-cylinder towards the left end of the picture.

about $-10°$C. The deformation (magnitude and direction) of the panels on a half-disk, when its temperature changes from 20°C to $-20°$C has been measured to be 150 μm. This result has been reproduced by a finite element analysis of the half-disk and it will be used in the final alignment of the pixels. We anticipate knowing the pixel geometry to a few tens of microns before the final alignment with tracks.

The service half-cylinder also contains all the mechanical and electrical infrastructure needed to support, position, cool, power, control and read out the detector. In particular, it contains electronics for providing bias voltage to the sensors, power to the read-out chips, signals for controlling the read-out chip via optical fibers linking it to the control room, and laser drivers for sending the signals (address and energy deposition) off the detector to the data acquisition system. The service half-cylinder also provides the path for cooling fluid necessary to remove the heat generated by the sensors and read-out chips.

At the end of each service half-cylinder there is an annular aluminium flange that contains holes to pass the power cables, cooling tubes, control and monitoring cables, and fiber optic read-out from intermediate patch panels to the FPix detector. The electronics cards needed for the operation of the detector are mounted on the inner surface of the half-cylinder. A picture of a half-cylinder is shown in figure 3.16.

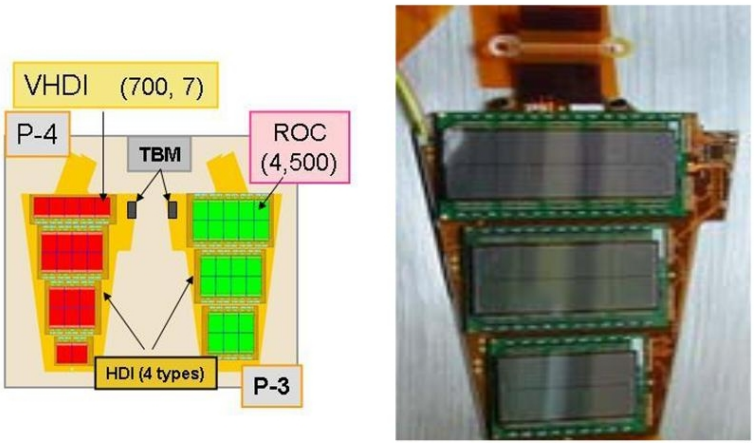

Figure 3.17: Sketches of the two types of FPix panels showing the different sizes and numbers of the plaquettes on each (left side). A photograph of an actual 3-plaquette panel (right side).

Forward pixel detection elements - the plaquettes

The basic unit of construction for the forward pixel detector is the *plaquette*. A plaquette consists of a single pixel sensor bump-bonded to an appropriate number of Read-Out Chips (ROCs) and wire-bonded to a very-high-density-interconnect (VHDI) that provides power, control, and data connections.

In order to cover the trapezoidal or *pie-shaped* panels without leaving cracks, five different sizes of plaquettes are needed. These are respectively 1×2, 2×3, 2×4, 1×5, 2×5, where the first digit refers to the number of rows and the second to the number of columns of read-out chips that are attached to a given sensor. The largest plaquette, the 2×5, has dimensions of 16 mm × 35 mm. The panels on the side of the cooling channel closest to the IP contain 1×2, 2×3, 2×4, and 1×5 plaquettes or a total of 21 ROCs. The panels on the side of the cooling channels farthest from the IP contain 2×3, 2×4 and 2×5 type plaquettes with a total of 24 ROCs. The sensors are offset on the upstream and downstream panels so that there are no cracks in the coverage due to the ROC read-out periphery. The two types of panels are shown in figure 3.17. A total of 672 plaquettes are needed.

The joining, or *hybridization*, of the pixel sensors and the pixel unit cells of the ROC is achieved by fine-pitch bumping using Pb/Sn solder and then flip-chip mating. The bumping is done on the 8" ROC wafers and the 4" sensor wafers. After bumping, the ROC wafers are thinned by backside grinding to 150 μm and then diced. Finally, each of the 5 different types of sensors are mated to the appropriate number of ROCs. The sensor with its ROCs bump-bonded to it is called a module. For FPix, the hybridization was done in industry. The fraction of broken, bridged, or missing bumps is at the level of a few 10^{-3}.

After delivery from the vendor, the bump-bonded pixel detector module is then installed on a Very High Density Interconnect (VHDI). The VHDI is a two-layer flexible printed circuit, laminated to a 300 μm thick silicon substrate, whose trace geometry and characteristics (impedance, low intrinsic capacitance, and low cross-talk) have been optimized for the intended use of conveying digital control and analog output signals to and from the sensors and ROCs.

The VHDI is made as follows. A bulk 6" silicon wafer is laminated to a flexible sheet containing several VHDI circuits. Passive components are also attached using surface-mount solder techniques. The wafer of populated circuits is then diced into individual circuits using a diamond saw. The circuits are then electrically tested.

The hybridized pixel module is attached and wire bonded to a populated VHDI to become a *plaquette*. The joining is made using parallel plate fixtures aligned on linear rails. The alignment of components is inspected using a coordinate measuring machine. A flexible plate is used for fine adjustments on the fixtures resulting in alignments between joined components within 100 μm. The adhesive bond between plaquette components is made in a vacuum at 60°C, to soften the adhesive and prevent air entrapment. An air cylinder applies and controls the mating pressure, which is limited by the compression allowed on the bump-bonds.

The effects of thermal cycling and radiation on the assembled plaquettes have been extensively tested. The tests demonstrate that the adhesive and the application method mitigate warping due to temperature changes, and provide reliable strength and thermal conductivity.

Once plaquettes are mechanically joined, they are clamped in cassettes that accommodate all processing steps such as wirebonding which provides electrical connections between the ROCs and the VHDI. After wirebonding we encapsulate the feet of wirebonds. This encapsulation is necessary due to periodic $I\vec{d}l \times \vec{B}$ forces expected to occur during actual CMS operation. The encapsulant acts as a damping force on the wire, preventing large resonant oscillations to work harden the wire and cause eventual breakage [30]. Finally the plaquettes undergo quick testing at room temperature. During this test the quality of the plaquettes is evaluated in terms of the characteristics of the sensor, the read-out chip, the number of bad pixels and missing bonds. The assembly and testing rate is optimized for a rate of six plaquettes per day.

The completed plaquettes are subjected to a quick plug-in test. Then they are loaded into a *Burn-In Box* where they undergo 10 temperature cycles between 20°C and -15°C. These cycles can take up to 2 days to complete, depending upon the thermal load. During these cycles, the plaquettes are monitored for electrical operation. We have seen no failures during the cycling. After the burn-in process is completed, the plaquettes are subjected to a series of electrical tests to ensure their suitability for their eventual mounting on a panel. These tests, at the operating temperature of -15°C, include the functionality of the ROC, the integrity of the bump-bond, and the I-V characteristics of the sensor. Other tests measure the thresholds and noise characteristics of each pixel on the entire plaquette assembly, and the individual pixel thresholds are trimmed via the ROC capability. We have found that the pixel trim values from the plaquette test on each pixel remain valid even after subsequent steps of the assembly process. After testing the plaquette data is loaded into the Pixel Construction Database and the plaquettes are graded. We have three main categories of grades:

- A - the plaquette is available for immediate mounting on a panel;
- B - potential issues have been found during testing and need further analysis;
- C - the plaquettes are unsuitable for mounting.

The data on each B-grade plaquette are examined carefully. In many cases, the plaquettes are found to have missed being classified as A-grade due to very minor deficiencies (e.g. slightly

2008 JINST 3 S08004

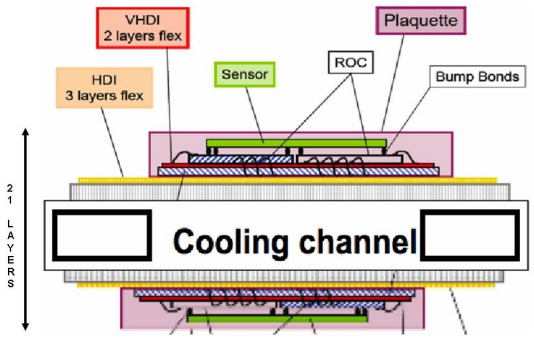

Figure 3.18: Sketch of a plaquette mounted on a panel showing its several layers.

too many noisy pixels) which will not be significant when an entire panel's quality is assessed. These are "promoted" to A-grade and declared usable on panels. Current plaquette yields, based on an original grade of A or a promotion to A-grade from B, are in the 80% range, varying slightly according to plaquette size.

Panel assembly

A panel is formed from three or four plaquettes attached to an assembly of a High Density Interconnect (HDI) laminated to a beryllium plate. The HDI is a three-layer flexible printed circuit whose trace geometry and characteristics (impedance, low cross-talk) have been optimized for the intended use of transferring digital control and analog output signals.

The process by which a panel is assembled is as follows. A single HDI circuit is laminated to a trapezoidal-shaped 0.5 mm thick beryllium plate. Passive components are attached using surface-mount solder techniques. The Token Bit Manager (TBM) is attached to the corner tab of the HDI using a die attach method and wire-bonding. After functional and burn-in tests with only the TBM, the individual plaquettes are attached to the HDI using adhesive for mechanical attachment and wire-bonds for electrical connection.

There are four types of panels, a right and left 3-plaquette version, and a right and left 4-plaquette version. The right and left handed versions have their TBMs on opposite sides of the panel centerline. Both types are required so that no panel part projects past a line in the vertical plane. The reason for the "3" and "4" type panels is that they are eventually mounted on opposite side of a blade, and the gaps between plaquettes on one type are covered by the active area of the other.

A panel is built up out of several layers of components. These are shown in figure 3.18. The total number of panels in all eight half-disks is 192.

Final detector assembly validation

The panels are attached onto the front and back of the half-disk cooling channels. The 4-plaquette panels are mounted on the side closest to the interaction region (IR), and the 3-plaquette versions on the opposite side. The half-disk assembly is mounted onto the half-service cylinder and is again tested.

Electronics chain

Each HDI is connected to another flexible printed circuit board, the *adapter board*. Each adapter board serves three blades (or 6 panels). One important purpose of the adapter board is to send and receive signals from the panels, which are mounted perpendicular to the axis of the service cylinder to and from the electronics mounted on the inner surface of the service cylinder. This is done by a *pigtail* at the end of each panel that plugs into connectors on the fingers of the adapter board.

The adapter board has three types of ASICs mounted on it. These are used to pass the clock, trigger and control data signals to each panel and return the received control signals back to the pFEC.

The adapter board is connected to another printed circuit board, the *port card*, by a light-weight extension cable. These cables are of two types, a power cable which distributes the power to the ROCs and TBMs, and the HV bias to the sensors. The other cable is to transmit the pixel data and control signals to the panel from the port card. The port card is a low-mass printed circuit board. It houses the electronics needed to interface the front-end chips with the VME electronics (the pFEC and pxFED) and power supplies located in the counting room. The port card transmits the clock signal, L1 trigger and slow control signals to the front end electronics. It distributes the power and bias voltages to the chips and sensors. It also monitors the currents and voltages as well as the temperature on some panels. These functions are done by various ASICs that are common to the CMS tracker. These ASICs include the DCU for monitoring, the TPLL for regenerating the trigger and timing signal, the gatekeeper for keeping the optical up and down links open as needed.

To control and monitor the various ancillary chips and optohybrids, there is a CCU board for each half service cylinder, as described above.

The port card contains the Analog Optohybrid (AOH). Each of the 6 laser diodes of the AOH chip receives data from one panel via its TBM and sends it over its own optical fiber to the Front end Driver (FED).

The control of the ROCs is achieved through the Pixel Front End Controller. Optical signals are sent from it to the Digital Optohybrids on the port card, through the extension cables to the adapter board, then to the TBM on the panel, through the HDI and the VHDI to the ROCs.

Power and monitoring

Power connections are made from CAEN power supplies via cables that run through the flange at the end of the half-cylinder away from the IR into a set of power/filter boards. From these boards, it is sent along wires to the port card, in the case of low voltages, and directly to the adapter board in the case of the sensor bias voltage.

Monitoring points for temperature are distributed throughout the service cylinder. There are also humidity sensors. Additional temperature sensors are mounted on the panels. High and low voltage and detector monitoring are connected to the DCS system described below.

Testing

Testing is a key element of quality assurance in the assembly process. While rework is possible, it is difficult and error prone. At every step, we confirm that we are using only “known good parts”.

2008 JINST 3 S08004

Testing must keep up with the driving assembly step, plaquette production. Full characterization of a plaquette requires hundreds of thousands of measurements. To accomplish this, we have developed special read-out hardware and software that can carry out these measurements quickly and efficiently. A software using a USB-based data acquisition scheme is employed when flexibility is needed to develop measurement programs of modest complexity and duration, such as the burn-in procedure. For the most extensive measurements, including plaquette testing and characterization, we use a PCI-based system and a software program called Renaissance [31].

Final testing is performed using the real data-acquisition and control hardware and prototype data-acquisition software and constitutes an end-to-end system test. Detailed testing also establishes an initial set of parameters for the many DACs and thresholds in the system.

3.2.6 Power supply

All needed high and low DC voltages are generated by means of a commercial modular system of the type CAEN-EASY4000 . This system is also employed by the CMS silicon strip tracker for which the main regulating cards (A4601H) were custom designed [32]. Only small changes in hard- and firmware were necessary for adaptation to the pixel project.

The core of this system, accessed through LAN, is located in the detector control room (USC55) and consists of one main controller (SY1527) containing 3 branch controllers (A1676A). The actual power supply cards are placed in two racks of 5 crates in close proximity to the detector thanks to their radiation tolerance and magnetic field resistance. This has been chosen in order to reduce power loss in the cables. The power supply crates are connected by flat cables ($\approx 100\,\mathrm{m}$) to the branch controllers. They are fed by local 3-phase $230\,\mathrm{V_{AC}}$ to $48\,\mathrm{V_{DC}}$ master converters of each 2 kW (A3486H) also suited for operation at hostile environments.

The crates house two types of electronic cards, one of 4 channels of 2.5 V/7 A (A4602) feeding the service electronics on the supply tubes (auxiliary power), while the other (A4603H) deliver 2 complex channels of each 2 low (1.75 V/7 A and 2.5 V/15 A) and 2 high voltage lines ($-600\,\mathrm{V}/20\,\mathrm{mA}$) for ROC and sensor biasing respectively. Each of these channels contains floating pulse-width-modulated DC/DC switching transformers with a common ground return for the 1.75 and 2.5 V lines. The isolation resistance (ground return versus earth on the racks) is typically $100\,\Omega$ at 5 MHz. Every card is controlled by an optically decoupled microprocessor for setting and measuring voltages, currents, ramp times, trip parameters, interlocks and others.

The DC levels are regulated over sense lines. The reaction time of the sensing circuit (typically 200 μs) is subject to fine tuning to comply with capacitive load, cable impedance and length (typically 50 m). The line drop in the cables amounts to roughly 2 V, while the regulators would allow for a maximum of 6 V. Fourteen A4602 cards, yielding 40 independent channels of auxiliary power, feed the 32 slots of the barrel service tubes with each 2 DOHs and 6 AOHs as well as 4 groups of each 12 port cards of the forward half disks. The main supplies of 112 complex LV and HV channels (56 A4603H cards) feed the 64 barrel groups (192 ROCs, this contains groups with half size modules) of each 12 detector modules, and 48 forward groups (135 ROCs) of each 3 disk blades. Each of these groups draws a typical current of 4.6 A on the analogue (1.75 V) and 9 A on the digital (2.5 V) line respectively. The large current reserve of the supplies is needed to comply with conditions during bootstrapping where the ROCs remain briefly in an undefined state. It was

verified that the regulators undergo a smooth transition from the constant-voltage to the constant-current regime if the programmed current limits are approached. Beside microprocessor controlled actions (1 s) fast over-current security is guaranteed by various solid state fuses (10 ms) as well as crowbars (100 μs) for over-voltage.

Noise levels are typically 5 mV_{pp} on the LV and 50 mV_{pp} on the HV outputs which can easily be accepted thanks to the LV regulators in the ROCs and the intrinsically small sensor capacitances respectively. Of major concern in the overall design were fast drops in the digital current consumption (2.5 V line) in case of low ROC activity like in orbit gaps. Due to the cable inductance a typical current drop of 2 A per group generates over-voltage spikes at the module level in the order of some Volts depending on local buffer capacitors. The integrity of the cable-module-ROC circuit was therefore checked by a full simulation in SPICE together with measurements on pulsed current loads. This served for the designs of the cables and the electronic layout, e.g. grounding or HV distribution. (In one sector layer-1 modules are fed by one line while layer-2,3 modules are commonly fed by the other.) Finally a $6\times4\,mm^2$ shielded copper cable was chosen for the 40 m from the power supply cards to the patch panel (PP1) located in the HCAL with alternating current directions between adjacent lines. Two twisted pair lines for the senses and a bunch of 10 commonly shielded lines for HV are contained in the same cable complex ($0.1\,mm^2$).

Inductance, capacitance and characteristic impedance between two of the main lines were measured to be 6 μH/m, 0.13 nF/m and 24 Ω respectively. The 4 m connection between PP1 and PP0 (tracker bulkhead) uses Al conductors in the cable. The auxiliary power cable is also shielded and contains $26\times0.75\,mm^2$ and 4 twisted pair copper lines with $0.1\,mm^2$ for the sense wires.

3.2.7 Cooling

The power consumption per pixel amounts to around 55 μW, including about 13 μW from the sensor leakage current at final fluences of $6\times10^{14}/cm^2$. For the total of $\approx$ 66 million pixels this adds up to 3.6 kW. The power load on the aluminium cooling tubes is therefore expected to be about 50 W/m. The sensor temperature will be maintained at around -10°C. As for the strip detectors, liquid phase cooling with C_6F_{14} is used. To keep the temperature increase of the coolant below 2°C, a total flow rate of 1 litre/s is required.

The pixel system is cooled by a total of 18 cooling loops: 10 for the barrel and 4 for each of the two end disk systems. For the barrel, the coolant enters at $+z$ and exits at $-z$, or vice versa. The coolant for the two disk sets on each side of the interaction region is supplied and reclaimed from the same z side. One barrel loop feeds in parallel 9 thin-walled aluminium pipes, each cooling 8 modules in series. One disk loop cools in parallel one quarter of each of the 2 disks; inside the quarter disks the 6 blade loops are connected serially. The coolant flow at the pixel modules is turbulent. The total lengths of the cooling loops starting from and returning to the pixel cooling rack amount to about 80 m, resulting in pressure drops of below 2 bar.

3.2.8 Slow controls

The safe operation of the barrel and forward pixel detectors is guaranteed by the CMS Pixel slow controls system (DCS). Its tasks are to monitor temperatures and humidities at different locations

of the detector and to monitor and control the high and low voltages necessary for operation of the on-detector electronics.

The monitoring of temperatures and humidities is based on a commercial Siemens S7–300 modular mini Programmable Logic Controller (PLC) system. The Siemens S7–300 system monitors a total of 192 temperature and 8 humidity sensors installed in the Pixel barrel and forward endcap disks. For the temperature sensors, platinum resistance temperature detection sensors with a nominal resistance of 1 kΩ (Pt1000 RTD) have been chosen. The measurement of humidity is based on detecting the water vapor induced shear stress in a small polymer element that is connected to a Wheatstone Bridge piezoresistor circuit [33]. This circuit provides a small (mV) output signal that is linearly proportional to relative humidity (RH) between the full range of 0% to 100% RH and is amplified by the same kind of conditioning electronics that is used by the silicon strip tracker. The PLC of the Siemens S7–300 system is programmed in the Statement List (STL) language [34] to convert the currents and voltages of the temperature and humidity sensors into calibrated physical units (i.e. degrees Celsius for temperatures and percentages for humidities). For the purpose of avoiding damage to the detector in case the cooling system (dry air supply) fails, routines are programmed within the PLC to interlock the CAEN power supplies (shut-off the cooling) in that case.

An additional 96 Pt1000 temperature sensors are read out via the data-acquisition (DAQ) system, together with the temperature dependent voltage sources integrated into each one of the pixel read-out chips. The temperatures recorded by the DAQ system are passed to the slow controls system by means of a dedicated software interface [35].

The Barrel and Forward Pixel slow controls system is integrated into the PVSS graphical user interface (chapter 9) of the main CMS DCS.

3.3 Silicon strip tracker

3.3.1 Silicon sensors

The sensor elements in the strip tracker are single sided p-on-n type silicon micro-strip sensors [36, 37]. They have been manufactured on 6 inch wafers in a standard planar process, leading to significant cost reduction per unit area when compared to the more traditional 4 inch wafers. The base material is n doped float zone silicon with $\langle 100 \rangle$ crystal orientation. This crystal orientation was preferred over the more common $\langle 111 \rangle$ orientation because measurements [38] have shown that the built-up of surface charge on $\langle 100 \rangle$ wafers due to irradiation is much smaller and consequently irradiation causes less inter-strip capacitance increase on this material.

In TIB/TID and on the inner 4 rings of the TECs (figure 3.1), thin sensors of $(320 \pm 20)\,\mu$m wafer thickness are used, with substrate resistivity of $\rho = 1.55 - 3.25\,$kΩcm. TOB and the outer 3 rings of the TECs are equipped with thicker sensors of $(500 \pm 20)\,\mu$m thickness, with substrate resistivity of $\rho = 4 - 8\,$kΩcm. Due to the single sided processing, these sensors show a significant bow, which is required to be less than $100\,\mu$m.

A uniform n^+ implantation on the back side of the wafers, covered by aluminium, forms an ohmic contact which is connected to positive voltage up to about 500 V. Those sensors which are penetrated by the beams of the laser alignment system (section 3.3.7) feature a 10 mm hole in the

2008 JINST 3 S08004

back side metalization, as well as anti-reflective coating in order to achieve transmission through up to four sensors with a sufficient signal on a fifth sensor.

On the front side, strip shaped diodes are formed by p^+ implantation into the n type bulk. Due to the radiation damage to the crystal lattice, the bulk material will undergo type inversion and change to p type. At this point, the pn junction moves from the strip side of the wafer to the rear side contact. Each implanted strip is covered by an aluminium strip from which it is electrically insulated by means of a silicon oxide and nitride multilayer. This integrated capacitor allows for AC coupling of the signals from the strips to the read-out electronics, which is thus protected from the high leakage currents after irradiation. Each metal strip has two bond pads on each end, which are used to make a wire bond connection to the read-out chip and in case of the daisy chained sensors to make a wire bond connection between the two sensors in one detector module. For testing purposes there is also a DC pad connected to the p^+ implant. Each strip implant is connected via a (1.5 ± 0.5) MΩ polysilicon bias resistor to a p^+ bias ring which encloses the strip region and also defines the active area of the sensor.

For all sensors in the CMS strip tracker the ratio of p^+ implant width over strip pitch is $w/p = 0.25$, leading to a uniform total strip capacitance per unit length of about 1.2 pF/cm across all sensor geometries [38]. The actual w/p value was chosen in order to minimize the strip capacitance while still maintaining a good high voltage behaviour of the sensor. The aluminium strips feature a metal overhang of 4 to 8 μm on each side of the strip which pushes the high field region into the silicon oxide where the breakdown voltage is much higher, leading to stable high voltage operation. For the same reason, the bias ring is surrounded by a floating guard ring p^+ implant. It gradually degrades the electric field between the n^+ implant at the cut edge of the sensor and the bias ring, which are at backplane potential (high voltage) and ground, respectively. Figure 3.19 shows the layout of a corner of the active region of a sensor.

In order to equip all regions in the CMS tracker, 15 different sensor geometries are needed [36] (figure 3.19): two rectangular sensor types each for TIB and TOB, and 11 wedge-shaped sensor types for TEC and TID. They have either 512 or 768 strips, reflecting the read-out modularity of 256 channels (two 128-channel front-end chips multiplexed to one read-out channel). Since the sensors are fabricated on 6 inch wafers, they can be made rather large. Typical dimensions are for instance about $6 \times 12\,\mathrm{cm}^2$ and $10 \times 9\,\mathrm{cm}^2$ in the inner and outer barrel, respectively. The total number of silicon sensors in the strip tracker is 24 244, making up a total active area of $198\,\mathrm{m}^2$, with about 9.3 million of strips [36].

3.3.2 Read-out system

The signals from the silicon sensors are amplified, shaped, and stored by a custom integrated circuit, the APV25 [39]. Upon a positive first level trigger decision the analogue signals of all channels are multiplexed and transmitted via optical fibers to Front End Driver (FED) boards in the service cavern where the analogue to digital conversion takes place. This read-out scheme brings the full analogue information to a place where it can be used for accurate pedestal and common mode subtraction as well as data sparsification. Clock, trigger, and control signals are transmitted by optical links as well. A schematic view of the silicon strip tracker read-out scheme is given in figure 3.20. This analogue read-out scheme was chosen for several reasons: optimal spatial reso-

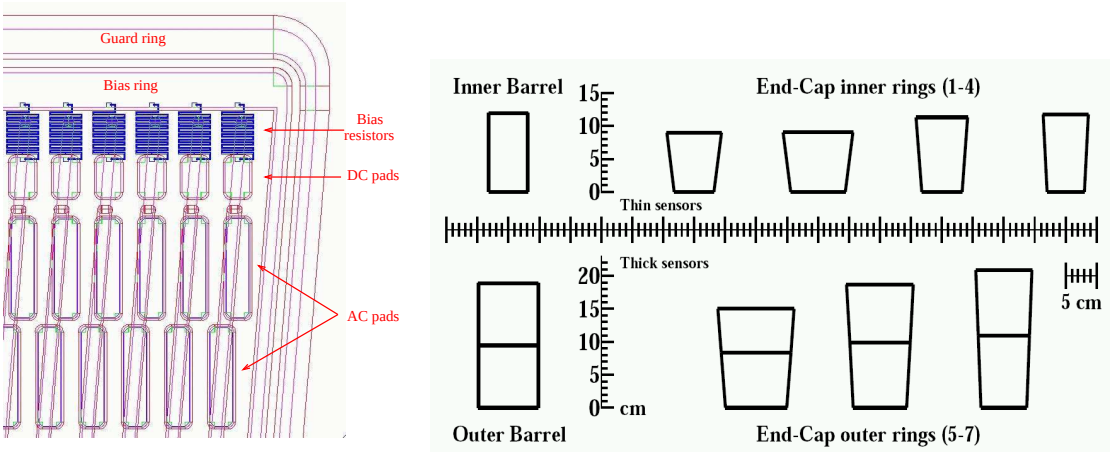

Figure 3.19: Left panel: drawing of one corner of the active region of a wedge-shaped silicon strip sensor for the tracker endcaps. Right panel: silicon sensor geometries utilized in the CMS tracker. In the outer layers the sensors are paired to form a single module, as shown in the figure. The Inner Barrel and Outer Barrel sensors exist in two types, of same area and different pitch. The sensors utilized for the first inner ring exist in two different versions, one for TID and one for TEC, respectively. (Only the TEC version is shown.)

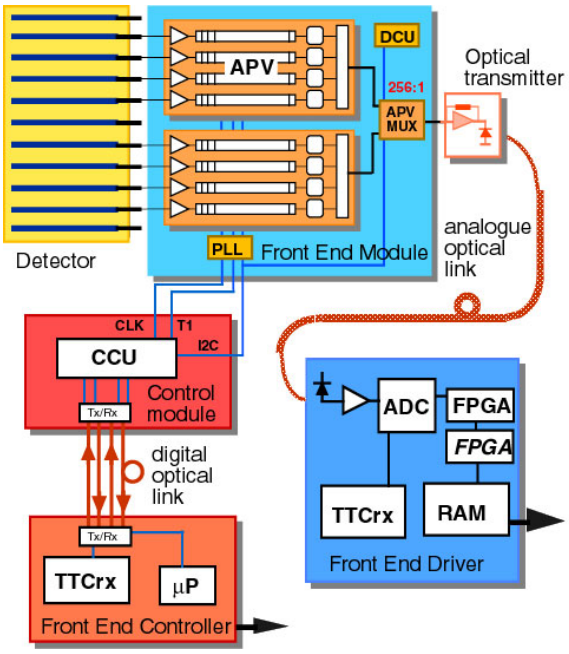

Figure 3.20: Read-out scheme of the CMS tracker.

lution from charge sharing, operational robustness and ease of monitoring due to the availability of the full analogue signal, robustness against possible common mode noise, less custom radiation hard electronics and reduced material budget as the analogue to digital conversion and its power needs are shifted out of the tracker volume.

2008 JINST 3 S08004

Front-end ASICs

The APV25 has been designed in an IBM 0.25 μm bulk CMOS process. Compared to processes with bigger feature sizes, the thin gate oxide inherent to this deep sub-micron process is much less affected by radiation induced charge-up and thereby, in conjunction with special design techniques, ensures radiation tolerance [40]. The APV25 has 128 read-out channels, each consisting of a low noise and power charge sensitive pre-amplifier, a 50 ns CR-RC type shaper and a 192 element deep analogue pipeline which samples the shaped signals at the LHC frequency of 40 MHz. This pipeline is used to store the data for a trigger latency of up to 4 μs and to buffer it. A subsequent stage can either pass the signal as sampled at the maximum of the 50 ns pulse (peak mode) or form a weighted sum of three consecutive samples which effectively reduces the shaping time to 25 ns (deconvolution mode). The latter is needed at high luminosity in order to confine the signals to the correct LHC bunch crossing. The pulse shape depends linearly (linearity better than 5%) on the signal up to a charge corresponding to 5 minimum ionizing particles (MIPs, one MIP is equivalent to 25 000 electrons in this case), with a gradual fall off beyond. When a trigger is received, the analogue data from all 128 channels of the appropriate time slice in the pipeline are multiplexed and output at a rate of 20 MS/s (mega-samples per second) as a differential bi-directional current signal, together with a digital header. Due to the tree structure of the analogue multiplexer the order in which the channels are output is non-consecutive and therefore re-ordering is necessary prior to the actual data processing. An internal calibration circuit allows to inject charge with programmable amplitude and delay into the amplifier inputs in order to be able to monitor the pulse shape.

The APV25 needs supply voltages of 1.25 V and 2.5 V with a typical current consumption of about 65 mA and 90 mA respectively, leading to a total power consumption of typically around 300 mW for one APV25 or 2.3 mW per channel. The noise of the analogue read-out chain is dominated by the front end MOSFET transistor in the APV25. Measurements have shown that the total noise for an APV25 channel depends linearly on the connected detector capacitance C_{det}. The equivalent noise charge is found to be $ENC_{\mathrm{peak}} = 270e + 38e/\mathrm{pF} \cdot C_{\mathrm{det}}$ in peak mode and $ENC_{\mathrm{deconv}} = 430e + 61e/\mathrm{pF} \cdot C_{\mathrm{det}}$ in deconvolution mode, both measured at room temperature [39]. Mainly due to the MOSFET characteristics, the noise reduces with temperature approximately as $ENC \sim \sqrt{T}$. Therefore, the noise at operating temperature is about 10% lower.

More than 100 APV25 chips from all production lots have been irradiated with X-rays to 10 Mrad ionizing dose, in excess of the expectation for 10 years of LHC operation. No significant degradation in pulse shape or noise level has been observed.

The APV25 is fabricated on 8 inch wafers with 360 chips per wafer. More than 600 wafers corresponding to 216 000 chips have been manufactured and probe-tested. After initial yield problems were solved, an average yield of 88% was achieved.

Another custom ASIC, the APVMUX, is used to multiplex the data streams from two APV25 chips onto one optical channel by interleaving the two 20 MS/s streams into one 40 MS/s stream, which is then sent to a laser driver of the optical links. One APVMUX chip contains 4 such multiplexers.

Optical links

Analogue optical links are used to transmit the data streams from the tracker to the service cavern over a distance of about 100 m at 40 MS/s. Likewise, the digital timing and control signals (see below) are transmitted by digital optical links running at 40 Mb/s [41]. Optical links are superior to an electrical distribution scheme mainly since they have minimal impact on the material budget and are immune to electrical interference. The transmitters are commercially available multi-quantum-well InGaAsP edge-emitting devices, selected for their good linearity, low threshold current and proven reliability. Epitaxially grown planar InGaAs photo-diodes are used as receivers. The optical fiber itself is a standard, single-mode, non dispersion shifted telecommunication fiber. The fibers are grouped in ribbons of 12 fibers which in turn are packaged in a stack of 8 inside a 96-way ribbon cable, which features a small diameter (< 10 mm) and a low bending radius (8 cm). For the analogue data link up to three transmitters are connected to a laser driver ASIC on an Analogue Opto-Hybrid (AOH), one of which sits close to each detector module. The electrical signals from the APVMUX are transmitted differentially over a distance of a few centimeters to the laser driver, which modulates the laser diode current accordingly and provides a programmable bias current to the diode. For the bi-directional digital optical link a set of two receivers and two transmitters is mounted on a Digital Opto-Hybrid (DOH), converting the optical signals to electrical LVDS [42] and vice versa.

Front End Drivers

The strip tracker Front End Driver (FED) is a 9U VME module which receives data from 96 optical fibres, each corresponding to 2 APV25 or 256 detector channels [45]. All 96 fibres are processed in parallel. The optical signals are converted to electrical levels by opto-receivers [43] and then digitized by a 40 MHz, 10 bit ADC. The ADC sampling point for each fibre can be programmed independently in 1 ns steps. After auto-synchronization to the APV data stream, pedestal corrections are applied and the common mode subtracted. The common mode correction is calculated for each trigger and each APV separately. The samples are then re-ordered to restore the physical sequence of detector channels which is essential for the following step of cluster finding. Pedestal values for each detector channel and thresholds for cluster finding are stored in look up tables. The digital functionality of the FED is implemented in FPGAs and can therefore be adjusted with considerable flexibility. In zero suppression mode, which is the standard for normal data taking, the output of the FED is a list of clusters with address information and signal height (8-bit resolution) for each strip in the cluster, thus passing to the central DAQ only those objects which are relevant for track reconstruction and physics analysis. In this way an input data rate per FED of about 3.4 GB/s, at LHC design luminosity, is reduced to roughly 50 MB/s per percent strip occupancy. Other modes are, however, available which suppress one or more steps in the processing chain and therefore transmit additional data to the central DAQ to be used mainly for debugging and system analysis. There are a total of 450 FEDs in the final system.

2008 JINST 3 S08004

Control and monitoring

Clock, trigger and control data are transmitted to the tracker by Front End Controller (FEC) cards [44]. These are VME modules, located in the service cavern, as close as possible to the tracker in order to reduce trigger latency. They receive clock and trigger signals from the global Timing Trigger and Command (TTC) system and distribute those as well as control signals via digital optical links and the digital opto-hybrids to LVDS token ring networks (control rings) inside the tracker volume. Several Communication and Control Units (CCU) [46] participate in one token ring. These are custom ASICs which interface the ring network to the front-end chips. One CCU is mounted on a Communication and Control Unit Module (CCUM) and is dedicated to a set of detector modules. A combined clock and trigger signal is distributed to Phase Locked Loop (PLL) chips [47] on each detector module while the industry standard I^2C protocol [48] is used to send control signals to the APV chips as well as to the other ancillary chips. One CCU can control up to 16 units so that one FEC ring typically controls a set of several tens of detector modules. The PLL chips decode the trigger signals and provide a very low jitter, phase adjustable clock signal to the local electronics.

Detector Control Unit (DCU) ASICs [49] on the detector modules are used to monitor the low voltages on the hybrid, the silicon sensor leakage current, and the temperatures of the silicon sensors, the hybrid and the DCU itself. For this purpose, each DCU contains eight 12 bit ADCs. The DCUs are read out through the control rings and digital links so that these readings are only available when the control rings and the detector modules are powered.

Hybrids

The front-end read-out electronics for a detector module is mounted onto a multi chip module called hybrid [50]. Due to the different detector module geometries 12 different types of hybrids are needed in the CMS silicon strip tracker. Each hybrid carries 4 or 6 APV25 read-out chips which are mounted as bare dies, and one APVMUX chip, one PLL chip and one DCU chip which are packaged components. The main features of the hybrid are to distribute and filter the supply voltages to the chips, to route clock, control and data lines between the chips and to remove the heat from the chips into the cooling system. No high voltage is present on the CMS tracker hybrids. The hybrid substrate is fabricated as a four layer polyimide copper multilayer flex circuit (figure 3.21). It is laminated onto a ceramic (Al_2O_3) carrier plate using double sided acrylic adhesive. A polyimide cable is integrated into the layout of the hybrid. The minimal feature sizes are $120\,\mu$m for via diameter and line width. Large metalized through holes under the chips transfer the heat to the underlying ceramic plate, from where it is removed through the frame of the module into the cooling system. Three different flex circuit types (one each for TIB/TID, TOB and TEC) combined with different geometries of the ceramic plates, different connector orientations and different number of APV25 chips (4 or 6) make up the total of 12 different hybrid flavours.

Power supplies

Silicon strip modules are grouped into 1944 detector power groups in order to share the power services. Each group is supplied by a power supply unit (PSU) [32], featuring two low-voltage

2008 JINST 3 S08004

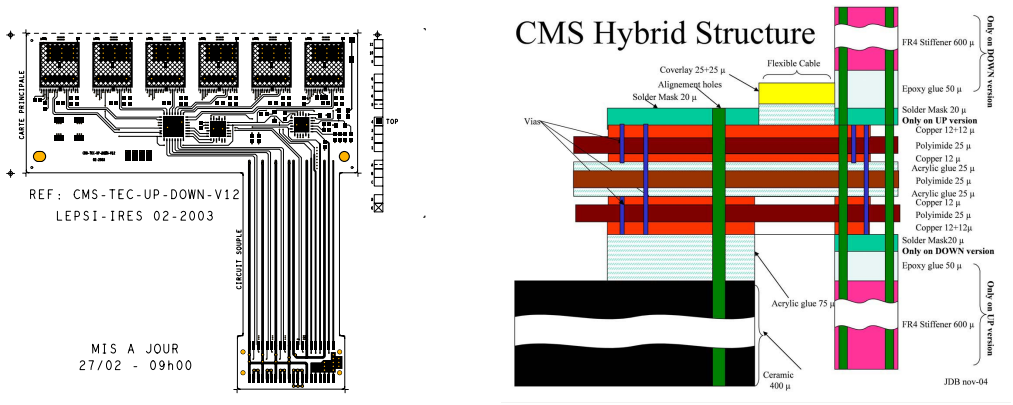

Figure 3.21: Front-end hybrid layout (example for TEC shown on the left) and arrangement of layers.

regulators, respectively 1.25 V (up to 6 A) and 2.5 V (up to 13 A), and two high-voltage regulators (0-600 V, up to 12 mA). All regulators are "floating" (return line isolated from the local earth). The two low-voltage channels share the same return line and use the sensing wire technique to compensate, up to 4 V, the voltage drop along the cables. The two high-voltage regulators are fanned out at the PSU exit into 8 lines; each silicon strip sensor is connected to one of these lines. Two PSU are combined into one power supply module (PSM, A4601H model). In total 984 A4601H boards are needed to power the detector groups; they are located on 129 EASY 4000 crates, disposed on 29 racks, around 10 m away from the beam crossing region, and operate in a "hostile" radiation and magnetic field environment, powering the detector through $\approx$ 50-m-long low impedance cables [32]. The 356 control rings require a separate power at 2.5 V. This is provided by a different set of 110 control power supply modules (A4602, four 2.5 V channels per module), fully integrated in the same system of the A4601H units and located on the same crates. Both A4601H and A4602 units require two distinct 48V power sources, one source (48Vp) for the regulators, the other (48Vs) for the service electronics. They are both provided by AC-DC converters, CAENs A3486 ("MAO"), disposed on the same racks. Each EASY 4000 crate hosts up to 9 boards (A4601H mixed to A4602) and provides 48Vp and 48Vs rails, interlock and general reset bus lines. The first slot in the crate (*slot 0*) hosts one interlock-card, which interfaces the interlock and reset lines to the control and safety systems (section 3.3.8). The average power consumption of each silicon strip module with 6 (4) APV25 chips is about 2662 mW (1845 mW). The total power supplied by A4601H and A4602 boards is approximately 68 kW, of which nearly 50% is dissipated on power cables. The power consumption is foreseen to increase with the aging of the detector; the power supply system is dimensioned to cope with up to 60% increase of the low-voltage currents, corresponding to a total consumption of nearly 150 kW.

2008 JINST 3 S08004

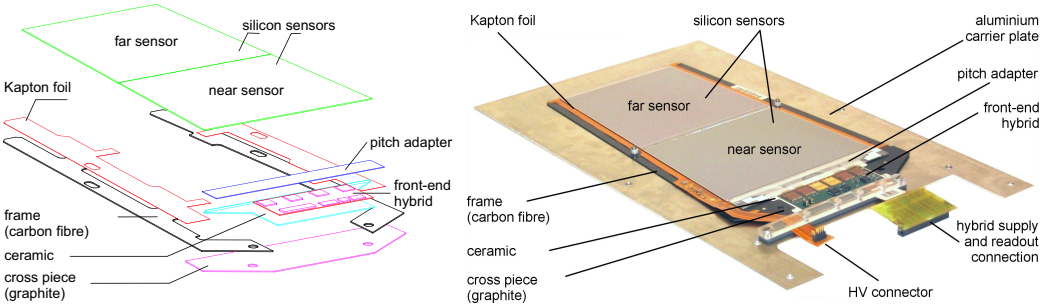

Figure 3.22: Left panel: exploded view of a module housing two sensors. Right panel: photograph of a TEC ring 6 module, mounted on a carrier plate.

3.3.3 Silicon modules

Module design

The silicon strip tracker is composed of 15 148 detector modules distributed among the four different subsystems (TIB, TID, TOB, TEC). Each module carries either one thin (320 μm) or two thick (500 μm) silicon sensors from a total of 24 244 sensors. All modules are supported by a frame made of carbon fiber or graphite, depending on the position in the tracker. A Kapton circuit layer is used to insulate the silicon from the module frame and to provide the electrical connection to the sensor back plane, i.e. bias voltage supply and temperature probe read-out. In addition the module frame carries the front-end hybrid and the pitch adapter. Figure 3.22 shows an exploded view and a photograph of a TEC module.

Modules for the inner barrel, the inner disks and rings 1 to 4 in the endcaps are equipped with one sensor, modules in the outer barrel and rings 5 to 7 in the endcaps have two sensors. In the case of two sensors, their corresponding strips are connected electrically via wire bonds. Depending on the geometry and number of sensors the active area of a module varies between 6243.1 mm^2 (TEC, ring 1) and 17202.4 mm^2 (TOB module). In total 29 different module designs, 15 different sensor designs and twelve different hybrid designs are used in TIB, TOB, TID and TEC. For alignment purposes special modules are prepared with etched holes in the aluminium back plane to allow a laser ray to traverse up to five modules.

The module frame provides the stability, safety and heat removal capability needed in the sensor support and carries the read-out electronics. In addition it has to remove the heat generated in the electronics and the silicon sensor(s) into the cooling points. In the endcaps the frame for the one-sensor modules is U-shaped and made of (780$\pm$5) μm thick graphite (FE779 carbon). For the two-sensor modules a similar U-shaped support structure is obtained by gluing two (640$\pm$40) μm thick carbon fiber legs (K13D2U CFC, 5 $\times$ 125 μm fabric, cyanate ester resin (CE3)) on a 800 μm thick graphite cross-piece (FE779 carbon) which holds the front end electronics. In the inner barrel a 550 μm thick carbon fiber frame that surrounds the silicon sensor on all sides is used . For the TOB, U-shaped module frames are obtained by gluing two carbon fiber legs (K13D2U CFC, 5 $\times$ 125 μm fabric, cyanate ester resin (CE3)) on a carbon fiber cross piece made of the same material.

Both graphite and carbon fiber fulfil the requirements of high stiffness, low mass, efficient heat removal from the sensors, and radiation hardness. Differences in the expansion coefficients need to be compensated by the glue joint between the frames and the silicon. Several types of glues are used in module construction which all comply with the requirements of radiation hardness, good thermal conductivity and thermal stability. Among them are e. g. Epoxy AW 106 (Araldit, Novartis), silicone glue RTV 3140 (Dow Corning) to compensate for different thermal expansion coefficients and the conductive glue EE 129-4 (Polytec) between the silicon sensor back plane and the HV lines on the Kapton bias strips (see below).

Different types of aluminium inserts and precision bushings in the module frames are used to position and attach the modules to the larger support structures with high precision. TIB/TID and TEC modules are mounted using four points, two being high precision bushings that allow for a mounting precision of better than 20 μm while all four provide thermal contact between the module and the cooling pipes. For TOB modules two Cu-Be springs give the precision positioning and four screws ensure thermal contact.

The high voltage supply to the silicon back plane is provided by Kapton bias circuits running along the legs of the modules between the silicon sensor and the carbon fiber support frame. The connection of the bias voltage to the back plane is done via wire bonds. Thermal probes are placed on the Kapton foil to measure the temperature of the silicon. The glue joint between the temperature sensor and the back plane is done with the silicone glue RTV 3140.

The pitch adapter between the front end hybrid and the silicon sensor adjusts the strip pitch of the sensor (80 μm–205 μm depending on sensor type) to the APV pitch of 44 μm. It also allows placing the heat producing front end electronics farther away from the silicon sensors. A pitch adapter for TOB and TEC consists of a 550 μm thick glass substrate (Schott D263 glass), cut to the correct dimensions, with a pattern of low resistivity aluminium strips. For TIB 300 μm thick glass (Corning 1737F or G glass) is used. The 10 μm narrow lines are etched on a (1.0–1.5) μm thick aluminium layer deposited on a chromium base, resulting in less than 25 m$\Omega/\Box$.

Module assembly and testing

Sensors and front end hybrids are glued to the frames by high precision gantry robots. The components are aligned by cameras surveying special fiducial marks with a pattern recognition algorithm. In total seven institutes shared the responsibility for the assembly of all modules. The assembly rate was about 20 modules per day per gantry robot. A positioning precision of approximately 10 μm (RMS) has been achieved and one example from the quality control can be seen in figure 3.23.

Thin wire wedge bonding is used in several places on the modules to make electrical connections: APV chip to front-end hybrid, APV chip to pitch adapter, pitch adapter to sensor, sensor to sensor (in case of two-sensor-modules), bias voltage connection to the sensor back plane. In total 15 institutes *(bonding centers)* shared the responsibility for wire bonding all modules. The bonding rate was approximately 1 Hz. Bonding wire (99% aluminium, 1% silicon) with a diameter of 25 μm was used for all connections.

For the TEC and TOB modules the line of bonding wires connecting the hybrid pitch adapter to the silicon strips, and in the case of two sensor modules the strips of the two sensors, can be damaged by vibration during transport. As a protection for the TEC modules the silicon is glued

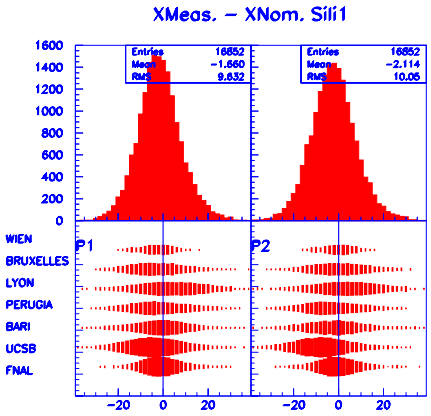

Figure 3.23: A typical residual distribution (in μm) for a reference point on the modules is shown for the different module assembly centers, indicating a precision of 10 μm (RMS) in the module production.

to a supporting strip (400 μm thin ceramic Al_2O_3) which in the case of the pitch adapter-sensor connection is also glued to the graphite cross piece. The reinforcement for the TOB modules was done by dispensing Sylgard 186 glue on the backside of the modules, between the two sensors and between the near sensor and the edge of the hybrid. For the TOB modules the sensor-sensor bonds and the backside APV bondings are encapsulated by Sylgard 186 glue across the bonding wires. For TIB modules no reinforcement was done.

After wire bonding each module was tested and graded, using the ARC system [51]. A detailed description of all tests performed and the acceptance criteria for good channels is given in the reference. Modules were graded A if fewer than 1% of the channels were failing the quality acceptance criteria (due to high noise, open bondings, oxide defects) and B if the failure rate was less than 2%. The remaining modules were graded C and were not used in the experiment. Other reasons to reject modules were imperfect mechanical precision or poor high voltage behaviour. All relevant test results are stored in the central CMS tracker data base. The yield of module production was greater than 97%.

3.3.4 Tracker Inner Barrel and Disks (TIB/TID)

Introduction and mechanics

The Tracker Inner Barrel (TIB) consists of four concentric cylinders placed at radii of 255.0 mm, 339.0 mm, 418.5 mm, and 498.0 mm respectively from the beam axis that extend from -700 mm to $+700$ mm along the z axis. The two innermost layers host double sided modules with a strip pitch of 80 μm, while the outer two layers host single sided modules with a strip pitch of 120 μm. Each cylinder is subdivided into four sub-assemblies ($\pm z$, up/down) for ease of handling and integration. Each of these sub-assemblies (half-shells) hosts an independent array of services from cooling to electronics and thus can be fully equipped and tested before being mechanically coupled to each other during final assembly.

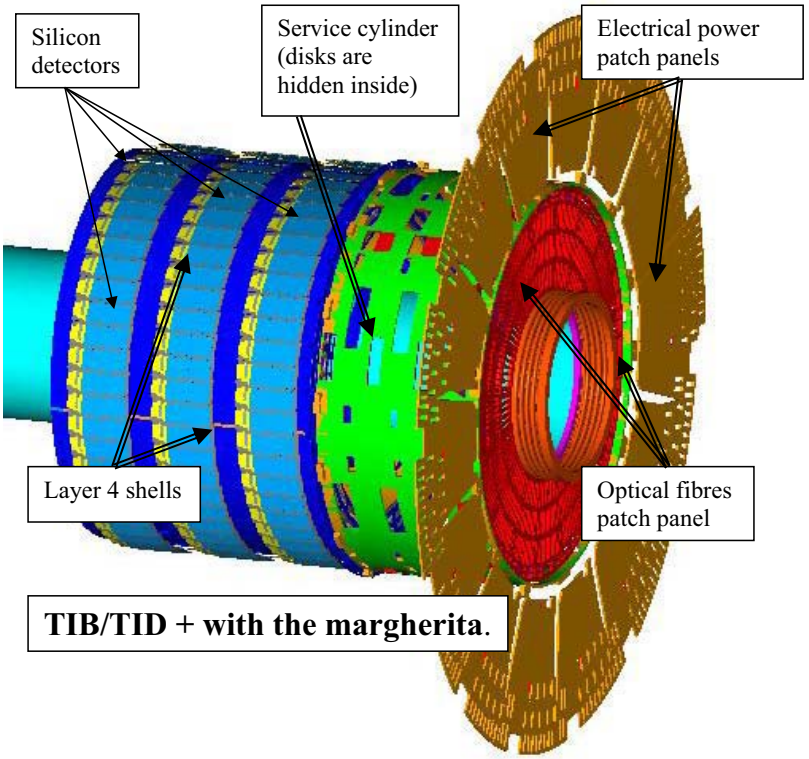

Figure 3.24: Schematic drawing of the TIB/TID+ subassembly. This structure and its twin (TIB/TID-) nest inside the Tracker Outer Barrel (TOB), one for each end. Services routed out from the margherita consist of copper cables for powering and slow controls, optical fibers for signals and controls and also cooling fluid supply lines made of aluminium tubing.

Two service cylinders are coupled to the ends of TIB$\pm$ (referring to $+z$ or $-z$) which end in a service distribution disk called the *margherita* (see below). These service cylinders play a dual role: one is to route out services from the shells to the margherita, the other is to support the Tracker Inner Disks (TID) which sit inside them. Figure 3.24 shows a schematic drawing of one half TIB/TID structure together with its corresponding margherita.

The TID$\pm$ are assemblies of three disks placed in z between $\pm$800 mm and $\pm$900 mm. The disks are identical and each one consists of three rings which span the radius from roughly 200 mm to 500 mm. The two innermost rings host back-to-back modules while the outer one hosts single sided ones. Just like the TIB shells each individual ring can be fully equipped and tested independently of the others before final assembly. Together the full TIB/TID guarantee hermetical coverage up to pseudorapidity $\eta = 2.5$.

All mechanical parts like shells, disks and service cylinders are made of high strength low deformation carbon fiber chosen both for its lightness and its low material budget. The margherita is instead made of conventional G-10 fiber epoxy with 30 μm copper on both sides.

The silicon detector modules are mounted directly on the structure's shells and rings. Thus, while a large number of modules has to be integrated and tested at any one time, the approach chosen allows for far greater precision of assembly. The individual components of a TIB shell,

some of which not only service the silicon detector needs but also define its geometric position in space, will be described in some detail in next paragraphs.

Cooling

The cooling circuits must be able to efficiently cool the detectors with a cooling liquid temperature down to about $-25°C$, while keeping the material budget as low as possible. For the TIB/TID the decision was made to use aluminium piping of 6 mm cross section and 0.3 mm wall thickness. These pipes are bent into loops and soldered to inlet/outlet manifolds which connect several loops in parallel. The thermal connection between pipes and silicon modules is made with aluminium ledges which are glued to the pipes. On each ledge there are two threaded M1 holes onto which the modules are tightened. For the TIB each loop hosts three modules placed in straight row (figure 3.25), while in the TID arrangements are more varied even though the number of modules per cooling loop is similar.

Since the position of the ledges defines the position in space of the modules, after the glue has hardened the whole half cylinder is surveyed with a precision measuring machine. Before gluing, the circuits are tested individually for leaks both at cold temperatures ($-30°C$) and at high pressure (20 bars). It is only after the survey that the TIB cylinders (or TID disks) are available for the integration of the electrical parts including the detector modules. The dimensions of the cooling circuit vary from layer to layer and depend on the amount of power dissipated by the modules used for that specific layer. The cooling circuits vary from a minimum of four loops (12 modules equivalent) for the double sided layers to a maximum of 15 loops for the outer single sided ones where individual module heat dissipation is much lower. The TIB/TID uses a total of 70 independent cooling circuits so that in case of an accidental break in one of the circuits only a small part of the tracker is affected. The TIB thus is organized in three module ladders (the cooling loop) which cover the outer and inner surface of the four layers. The same concept applies to the TID with the only difference that the number of modules per cooling loop varies with the ring radius. The electrical grouping which we now describe takes this mechanical distribution into account.

Electrical grouping

The modules have been grouped together electrically. The basic group consists of three modules which sit on any given cooling loop (figure 3.25). The three modules are interconnected through a Kapton circuit (mother cable) through which powering, detector biasing and controls are distributed. At the top of a mother cable sits a CCUM which takes care of clock, trigger and I^2C distribution. These mother cables are then electrically joined in a more complex group called the control ring which distributes trigger, clock and slow control signals to the CCUMs. Control ring groups never straddle two different cooling loops and are dimensioned so that a reasonable compromise between granularity and complexity is achieved. Control rings in the TIB/TID make use of a unit called the DOHM (Digital opto-hybrid module) which receives all the signals from the optical fibers coming from the front end controllers (FEC) and converts them to electrical LVDS signals that are then distributed to up to 45 detector modules (15 mother cables) via CCUs. Given the high number of modules belonging to a Control Ring, TIB/TID has implemented redundancy in its DOHM hardware.

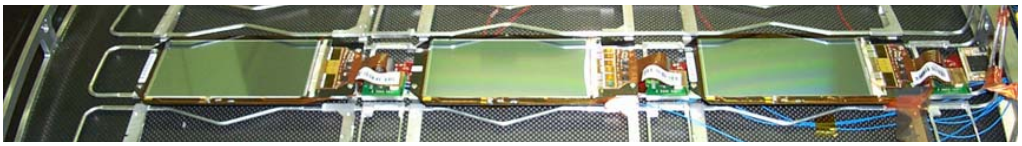

Figure 3.25: Three TIB modules mounted on a layer 3 shell. The Kapton mother cable runs underneath. A CCUM module at the end of the string interfaces the modules to the control ring. Also visible are the three analog opto-hybrids (see text) and fibers.

Modules have been grouped together to keep the number of power supplies down to a manageable level. The smallest power group consists of three modules (one mother cable) while the largest comprises up to 12 modules (four mother cables). Power groups are contained within a control ring (i.e. there is no straddling across control ring boundaries) and are fed by a specific power supply unit (PSU) developed for the tracker which also supplies HV biasing for the detectors.

Analog signals from the detector front end are converted to optical by analog opto-hybrids which sit next to the silicon modules and are connected directly to the front end hybrids. Thus the system is completely optically decoupled from the DAQ which helps preserve signal integrity while avoiding ground loops.

Grounding of the TIB/TID relies on the cooling circuits which are made of aluminium. The return current wires are connected to the cooling manifolds for all mother cables and DOHMs. The cooling inlet and outlet pipes run along the service cylinder across the margherita, making electrical contact with it. Outside the tracker volume these pipes are then connected to the CMS detector ground. Power cable shields are connected to the margherita which hosts all of the connectors. All detector modules have their own carbon fiber frame directly connected to the front end hybrid local ground. The shells are grounded through the cooling manifolds.

3.3.5 Tracker Outer Barrel (TOB)

Mechanical structure and layout

The Tracker Outer Barrel consists of a single mechanical structure (*wheel*) supporting 688 self-contained sub-assemblies, called *rods*.

The wheel is composed by four identical disks joined by three outer and three inner cylinders (figure 3.26). Disks and cylinders are made of carbon fiber epoxy laminate. The cylinders have a core of aramid-fiber honeycomb. The joints between disks and cylinders are realized with aluminium elements glued to the carbon fiber parts on precision fixtures, and then bolted together. Each of the disks contains 344 openings, in which the rods are inserted. Each rod is supported by two disks, and two rods cover the whole length of the TOB along the z axis. The wheel has a length of 2180 mm, and inner and outer radii of 555 mm and 1160 mm, respectively. With cabling at its two ends the TOB has a total length of 2360 mm. The openings in the disks form six detection layers with average radii of 608, 692, 780, 868, 965, 1080 mm. Within each layer, the centers of gravity of the rods are displaced by ± 16 mm with respect to the average radius of the layer, thus allowing for overlap in ϕ and therefore full coverage within each layer. The rod mechanics are designed in such a way to implement overlap of the silicon sensors at $z = 0$. In each layer, the

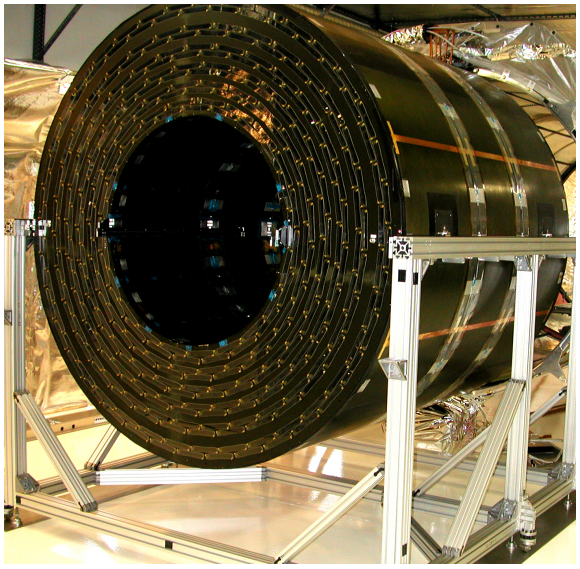

Figure 3.26: Picture of the TOB wheel.

overlap in the r-ϕ view between neighboring rods is always larger than 1.5 mm or 12 strips, while the overlap around $z = 0$ is precisely 1.5 mm. Inside the disk openings, the rod support spheres are held by precision elements made of polyetherimide plastic that are glued to the carbon fiber structure. The four disks have all been assembled in a temperature-controlled room on one single precision table, ensuring a precision on the relative positions of the rod holding elements and the aluminium elements joining disks and cylinder of 100 μm, and a reproducibility between different disks at the 10 μm level.

The wheel is equipped with *targets* for measurements of the geometrical precision of the assembled structure. Photogrammetry, theodolites, and 3D coordinate measurement systems have been used for survey and alignment of the wheel structure. Some of these targets remain visible after insertion of the TOB in the *tracker support tube*, for a precise measurement of the TOB positioning in the tracker reference frame, and even after integration of TIB, to monitor possible movements due to deformations of the loaded structure. The wheel mechanics has been thoroughly measured before starting rod integration, and the relative positioning of the precision elements has been found to be typically within 100 μm of nominal values over the whole TOB dimensions, with maximum deviations observed around 200 μm.

The rod mechanics

The rods are self-contained assemblies providing support and cooling for 6 or 12 silicon detector modules, together with their interconnection and read-out electronics.

The mechanical structure consists of two 1130 mm long carbon fiber C-shaped profiles, joined by several transverse carbon fiber ribs and plates. All rod components are contained in an envelope of $159 \times 1130 \times 22$ mm^3, except the four supporting spheres that stick out laterally in correspondence of the two disks of the wheel, and the *z-stops* that block the rod against the outer disk surface after insertion in the wheel.

2008 JINST 3 S08004

A U-shaped cooling pipe runs around the rod, inside the C-profiles; 24 aluminium inserts are glued through openings along the profiles to the carbon fiber and around the cooling pipe; these inserts provide support and cooling to the detector modules, that are mounted in six positions along the rod, three per side. Each detector is supported by four inserts, two close to the read-out hybrid, and two close to the sensor-to-sensor bonds. The two inserts close to the hybrid implement pins on which the Cu-Be springs on the module frame are clamped, determining the precision of the module positioning; all four inserts have a threaded hole for the fixation of the module to the rod: cup-shaped washers together with a calibrated torque used in tightening the screw ensure efficient cooling contact between the aluminium heat spreader on the module frame and the rod support inserts. On the cooling pipe side, the shape and the size of the inserts is optimized to minimize the thermal impedance of the contact, which in turn allows to minimize the cross section of the cooling pipe.

In single-sided rods, which populate layers 3–6, one detector module is mounted in each of the six positions, with the strips facing the central plane of the rod. In double-sided rods, which populate layers 1 and 2, two detectors are mounted in each position, the inner one as in single-sided rods and the outer one with the backplane facing the backplane of the first module. The distance between the sensor and the middle plane of the rod is $\pm$3.3 mm in single-sided rods, $\pm$3.3 mm and $\pm$7.6 mm in double-sided rods.

The rod cooling pipes, and the manifolds housed on the outer disks of the wheel, are realized in CuNi 70/30 alloy. This material is chosen for its corrosion resistance, and as it allows reliable solder joints to be made relatively easily, avoiding the use of o-rings or ferrules in the pipe connections; the reliability of the cooling circuits is a crucial issue for the tracker, and particularly so for the TOB, which is the most inaccessible subsystem once the detector is fully integrated. The rather high density of the material (its radiation length of about 1.4 cm is 6 times shorter than that of aluminium) is compensated by the reduced thickness of the walls that this technology allows: rod pipes and manifolds have 100 μm and 200 μm wall thickness, respectively. In addition the design of the cooling circuit has been optimized (as already mentioned above), to minimize the cross section of the pipes (the cooling fluid also gives a non-negligible contribution to the material budget), and to maximize the number of rod pipes served by a single manifold (within the constraints of the desired cooling performance). An outer diameter of 2.2 mm is chosen for single-sided rod pipes (providing cooling to 6 detectors), 2.5 mm for double-sided rod pipes (providing cooling to 12 detectors), and 6 mm for the manifolds; one manifold serves on average more than 15 rod pipes, the actual number varying between 8 and 22 depending on the region of the TOB. Overall, the whole TOB is served by 44 cooling lines, giving an average of 118 detectors, or 550 read-out chips, per line.

Rod electrical design

The 6 or 12 modules housed in a rod form a power group, i.e. they are supplied by a single power supply unit. The low voltage lines supplying the front-end hybrids and the Analogue Opto-Hybrids (AOHs) run in the Inter-Connect-Bus (ICB), a 700 mm long printed circuit board sitting in the middle plane of the rod (figure 3.27). The Communication and Control Unit Module (CCUM) is plugged to one end of the ICB. The clock and the control signals issued by the CCUM are also routed to the final destinations through the ICB. The distribution of power, clock and signals to front-end hybrids and AOHs proceeds through four other PCBs, the Inter-Connect-Cards

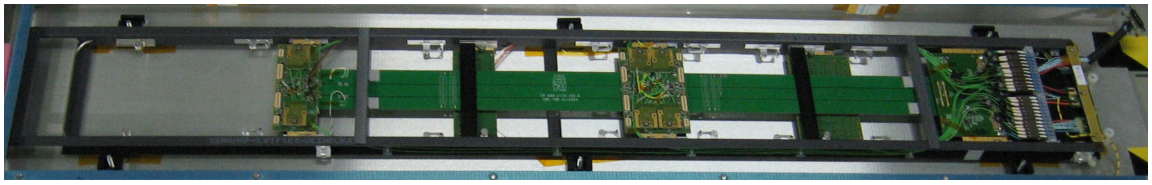

Figure 3.27: Photo of a rod frame equipped with electronics components, ready to receive silicon detector modules.

(ICCs). Two ICCs serve one module position and two other ICCs serve two module positions. ICCs have different design in single-sided rods and double-sided rods, which have one and two modules mounted in each module position, respectively; therefore there are in total four different ICC flavours in the TOB.

The ICB is held in place by small transverse carbon fiber plates; the ICCs and the CCUM are plugged to the ICB and screwed to the aluminium module support inserts (on the opposite side of the module), which also provide a good cooling contact to those boards. The AOHs are supported and cooled only by the connector that plugs to the ICCs. In addition to distributing LV power and CTRL signals, the ICCs receive the data lines from the read-out hybrid and route them to the AOHs (a few cm away) where they are converted to optical signals. The ICCs also receive lines carrying temperature information from the module frame Kapton circuit and route them to the ICB. The optical fibers leaving the AOHs travel inside the carbon fiber profiles, guided by dedicated plastic holders. The only electrical lines not integrated in the ICB/ICCs distribution system are the bias lines for the sensors. These run in dedicated wires (size AWG 26) housed in the carbon fiber profiles, while the line with the return current is integrated in the ICB. There are six lines in single-sided rods (one per module), and 8 lines in double-sided rods (four serving one module each, and four serving two modules each). The LV lines and the HV lines go in separate connectors in the rod end-panel, each of which also hosts some temperature lines, and then run all together to the back-end in one multi-service cable plus low-impedance cable. At the power supply backplane the six or eight bias lines are connected to the two independent high-voltage supply lines in such a way that each line powers one side of the rod. The clock and control lines as well as the LV lines powering the CCUM leave the rod through a short cable which plugs into the next rod of the control ring. The first and the last rod of a control ring are connected to the Digital Opto-Hybrid Module (DOHM). This board houses the digital opto-hybrids optically connected to the remote control system and distributes the clock and the control signals through a token-ring 40 MHz LVDS-based protocol to the connected rods (up to 10). The length of the optical fibers coming from the AOHs is chosen so that all fibers end at the same location near the CCUM, where the connectors of the 12-fiber ribbons are integrated (figure 3.28). The choice of including the optical patch panel inside the rod volume was made to reduce the thickness of the TOB services on the TOB end-flanges, so minimizing the inactive volume between TOB and TEC.

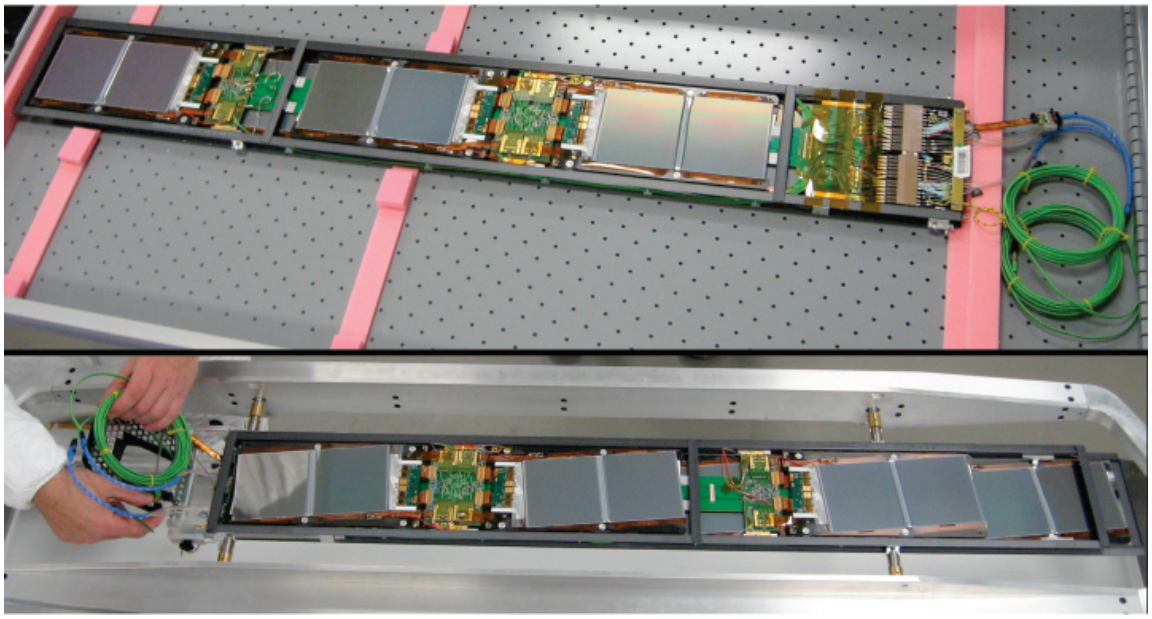

Figure 3.28: Top panel: photo of an assembled double-sided rod, showing the CCUM side, with the 12-way optical ribbons connected to the AOH fibers. Bottom panel: double-sided rod being prepared for insertion in the TOB mechanics; the side opposite to the CCUM is shown.

Electrical and read-out grouping

The grouping of the rods into control rings is designed primarily to avoid having control rings spanning across two different cooling segments, while maximizing the size of a control ring (to reduce cost and material budget) within the recommended limit of 10 CCUMs per ring. This logic results in two or three control rings per cooling segment, with a single exception of a cooling segment containing one control ring only. The average number of CCUMs (i.e. of rods) per ring in the TOB is 7.5. Within a control ring, rods are clustered in groups that are read out by the same FED. Again, a read-out group never spans over two control rings, and the grouping is optimized to minimize the number of unused channels in the FEDs (to reduce cost). The average FED occupancy in the TOB is 94%. In summary, the TOB is made of 688 rods read out by 134 FEDs, controlled by 92 DOHMs, and cooled by 44 independent lines.

Grounding

In each rod the return line of LV and bias is connected inside the CCUM to the return line of the LV power of DOHM and CCUMs, and connected through a short multi-wire cable to the cooling manifold serving the rod: this is the main ground connection of the rod. The grounding is improved by additional ground connections in each ICC, implemented through metalization around the mounting holes.

The DOHMs, mounted on the TOB end-flange (figure 3.29), are protected by alodyned aluminium plates of 0.5 mm thickness, which are locally connected to the power return line.

2008 JINST 3 S08004

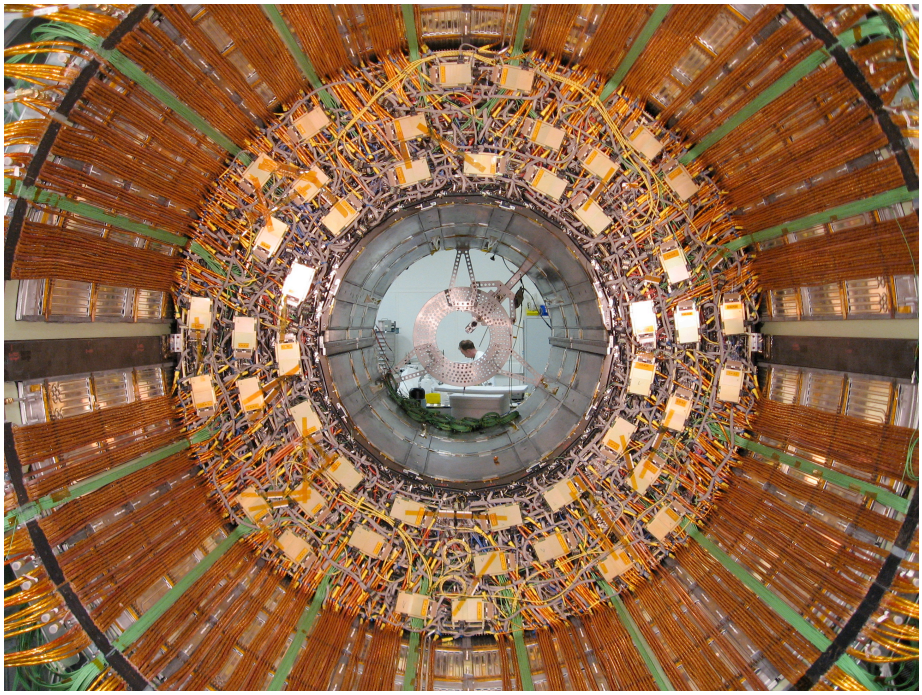

Figure 3.29: Photo of the completed z+ side of the TOB. The DOHMs form the outer layer of the services on the TOB end flange. Optical ribbons (green) run out, grouped in 16 channels. Power cables and feeding pipes run parallel to each other on the thermal screen panels.

The cooling circuits of the different segments are then connected electrically through short multi-wire cables soldered to the radial pipes feeding the manifolds (or to the manifolds themselves, for the outer layer) and screwed to the *ground rings*: an alodyned aluminium bar of 10×10 mm^2 square section bent to round shape and equipped all along with threaded holes, which is installed at the outer radius of the TOB, on both sides. Gold-coated copper strips of 30 mm width and 0.2 mm thickness connect the ground ring to the carbon fiber structure of the outer cylinder, in eight locations in ϕ. The connection to the carbon fiber is realized with conductive araldite. The same strip material is used to realize the electrical connections between outer cylinders and disks, and inner cylinders and disks, again in eight locations in ϕ. In addition, copper strips as long as the whole TOB are added on the outer surface of the outer cylinder (visible in figure 3.26); for the inner cylinder instead, which is inside the tracking volume, it was decided to rely on the conductivity of the carbon fiber.

Such design of the grounding scheme ensures good electrical connection of mechanical structures and power return lines making efficient use of the existing conductive materials (cooling pipes and carbon fiber parts), with minimal amount of added metallic elements.

2008 JINST 3 S08004

3.3.6 Tracker EndCaps (TEC)

Mechanical structure

The endcaps extend radially from 220 mm to 1135 mm and from $\pm$1240 mm to $\pm$2800 mm along the z-direction. The two endcaps are called TEC+ and TEC- (according to their location in z in the CMS coordinate system). Each endcap consists of nine disks that carry substructures on which the individual detector modules are mounted plus an additional two disks serving as front/back termination. A sketch of one endcap and a photograph of the completed TEC+ is shown in figures 3.30 and 3.31. Eight U-profiles, referred to as service channels because all services are grouped in their vicinity, join the disks together along their outer periphery, while at its inner diameter each disk is attached at four points to an inner support tube. To preserve the envelope necessary for the insertion of the pixel detector, the last six disks have a larger inner radius (309 mm) as compared to the first three (229 mm).

The disks are Carbon Fiber Composite (CFC) / honeycomb structures. The honeycomb core is 16 mm thick NOMEX, 3.2-92 with a border of epoxy potting. On either side of the core there is a symmetric layup of CFC skins (0.4 mm thickness). The skin material is CF-fabric THENAX HTA 5131,3K (T300) impregnated with EP121 epoxy resin. The same material is used for the service channels and the inner support tube. The latter has a thickness of 3 mm and is azimuthally segmented into four 90° segments. Each of these segments is attached to the disks and the gaps at the joints between segments are filled with epoxy glue so that they are gas tight. A thin cylindrical skin made of 0.5 mm thick CFC panels surrounds the endcaps on the outside and serves as a gas envelope for the atmosphere of dry nitrogen. The front plate has the same function and consists of a 5 mm NOMEX core with 0.2 mm CFC skins on each side. The back plate provides an additional thermal shielding for the cold silicon volume and is considerably thicker. The NOMEX core is 45 mm with each CFC skin 1.5 mm thick. The back plate also serves to make the overall structure rigid in the z-direction. The back plate is covered by another carbon fibre disk, the bulkhead, which is, however, mechanically detached from the TEC and supported by the tracker support tube. The bulkhead carries the outer connectors of all TEC cables, thereby forming a patch panel for the electrical connection of the TEC to the external power cables. It is covered by panels with heating foils which close the thermal screen at the end face of the tracker support tube.

Ten different module types are arranged in rings around the beam pipe. For reasons of modularity they are mounted on substructures called *petals*, which in turn are mounted on the disks. Disks 1 to 3 carry seven rings of modules, ring 1 is missing on disks 4 to 6, rings 1 and 2 are missing on disks 7 and 8, and disk 9 carries rings 4 to 7 only. Rings 1, 2 and 5 are built up of so-called double sided modules: two modules are mounted back-to-back with a stereo angle of 100 mrad. This provides space information perpendicular and parallel to the strip orientation.

Petals

To allow easy access to the detector modules they are mounted on modular elements, the petals (figures 3.32 and 3.33). Petals can be individually removed from the endcaps without uncabling and/or disassembling the entire structure. A total of 16 petals are mounted on each of the nine disks of one endcap, eight on the front face of the disk — as seen from the interaction point —

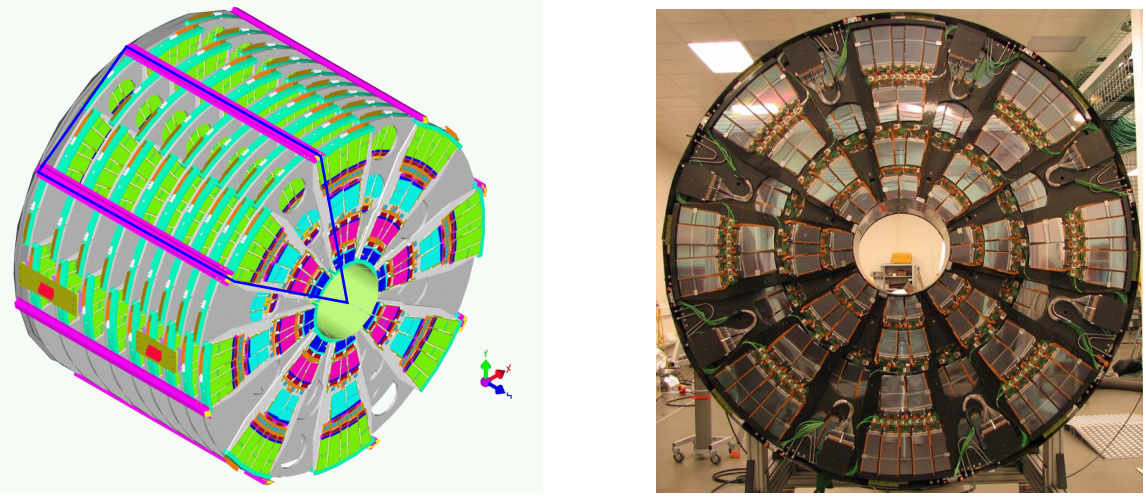

Figure 3.30: Left panel: Sketch of one tracker endcap. Modules are arranged in rings around the beam axis. They are mounted on trapezoidal sub-structures called *petals*. One sector, indicated with a line, consists of nine front petals mounted on the disk sides facing the interaction point (3 FD13, 3 FD46, 2 FD78, 1 FD9) and nine back petals mounted in the opposite side of a disk (3 BD13, 3 BD46, 2 BD78, 1 BD9). Right panel: Photograph of a TEC as seen from the interaction point. The diameter of the TECs is 2.3 m.

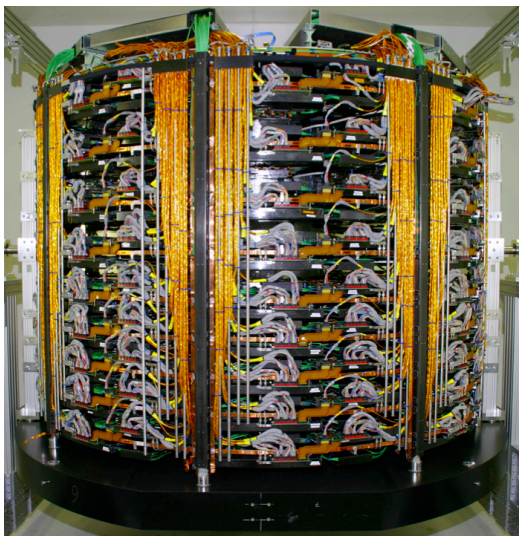

Figure 3.31: Side view of a TEC.

(front petals) and eight on the back face (back petals). Mechanically there are two types each of front and back petals, long petals for disks 1–3 and short ones for disks 4–9. As described above, the front and back petals on disks 1–3 carry all seven rings of modules and are labelled FD13 and BD13, respectively. Petals on disks 4–6 carry rings 2 to 7 (FD46/BD46), those on disks 7 and 8 carry rings 3 to 7 (FD78/BD78), and on disk 9 the petals carry rings 4 to 7 (FD9/BD9). The petals have a structure similar to the disks, consisting of a 10 mm NOMEX core sandwiched between

0.4 mm CFC skins. As viewed from the interaction point the modules belonging to rings 1, 3, 5, 7 are mounted on the petal front side (A-side and C-side for the front and back petals, respectively), while modules in rings 2, 4, 6 are mounted on the back side of each petal (B-side and D-side for the front and back petals, respectively). On a given disk the front petals overlap azimuthally with the back petals, as do, for a given petal, detector modules belonging to the same ring. Detectors in adjacent rings are arranged to overlap radially, thus providing full coverage. Each petal is mounted on inserts in the main disks using a three point fixation: one point fixed in x, y and z, one fixed only in phi, and one fixed only in z.

Cooling

The heat generated by all electronic components on a petal must be removed efficiently. In addition the silicon sensors must be operated at a temperature of about -10°C to reduce the effects of radiation damage. The silicon sensors and front end hybrids are cooled via the CFC frames of the detector modules, for which carbon fiber of high thermal conductivity is used (800 W/(m K)). The aluminium inserts for positioning the modules serve at the same time for the coupling to the cooling pipe. The two inserts along the legs of the module frame provide primarily for the cooling of the sensors, while the inserts on the frame base are heat sinks for the front end hybrid. Each petal contains two cooling circuits traversing the petal longitudinally and meandering from one cooling point to the next. The cooling pipes are made of titanium with an outer diameter of 3.9 mm and a wall thickness of 0.25 mm. They are embedded in the petal and serve to cool the components on both back and front side. The tubing is pre-bent into the proper shape. The input/output manifolds are laser welded onto the cooling pipes. After having milled the corresponding grooves and holes into the petals, the tubing is inserted. Gluing jigs are used to position the cooling inserts and to glue them to the pipes and to the petal. To close the grooves and re-establish the integrity of the petal a CFC skin with holes at the location of the inserts is glued onto the petal face. The inserts are then machined to the precision required for module positioning. The maximum heat load from the electronics on a petal is about 87 W, including the heating of the sensors after ten years of LHC operation. In these conditions a mass flow of 2.3 kg/min of the C_6F_{14} coolant gives a temperature difference of 2°C between petal inlet and outlet. The connection of the petal circuits to the piping running along z is done at the outer periphery of the petal. These connections can be undone easily in case the petal needs to be removed. A pair of longitudinal pipes serves either 4 or 5 petals, which are connected in parallel. A total of 64 longitudinal stainless steel pipes with 11 mm inner diameter are used per endcap.

Electrical system design

The silicon modules, AOHs and CCUMs on the petals are connected to motherboards, called InterConnect Boards or ICBs, which are mounted on both sides of the petal. In figure 3.32 photos of a bare front petal equipped with ICBs only are shown. There are five individual boards: the main board ICB_46 on side B/D, which carries all the connectors for the cables and two CCUM boards and transmits power and signals to the modules of rings 4 and 6, and four smaller boards, which provide the power and signals for the other rings (ICB_2 on side B/D and ICB_1, ICB_3 and ICB_57 on side A/C, where the numbers correspond to the number of the ring to which the

2008 JINST 3 S08004

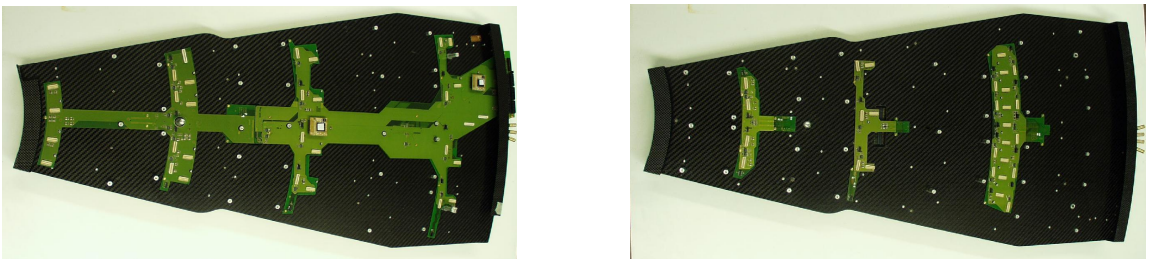

Figure 3.32: The different ICBs on the two sides of a front petal: ICB_2 and ICB_46 on side B, and ICB_1, ICB_3 and ICB_57 on side A (from left to right). On ICB_46, the two CCUMs are plugged.

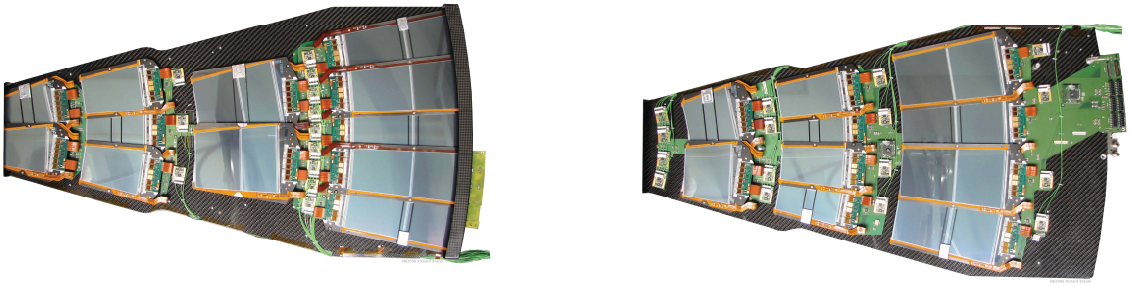

Figure 3.33: Left photograph: front side of a TEC Petal. Right photograph: back side.

connected modules belong). These four boards are connected to the main board. The ICB brings the ground, the various supply voltages and the bias voltage to the electrical devices on the petal, and transmits LVDS and I^2C signals. In addition analogue data from the FE hybrids are transmitted differentially to the AOHs over distances of a few centimetres.

To keep the number of low voltage power supplies and connections relatively small while limiting the current that must be provided by one power supply, the modules are organized in three low voltage (LV) groups, which are served by individual power supplies. The LV group 1 consists of rings 1 and 2, group 2 contains rings 3, 4 and 6 and finally rings 5 and 7 belong to group 3. This corresponds to 8/11/9 (4/8/11) modules or 48/44/44 (24/32/56) APVs on front (back) petals in LV group 1/2/3. In total there are eleven power rails on ICB_46, which must carry a current of up to 12 A. Sensing is implemented for the low voltage connections. The sense resistors are located in the electrical centre of each power group. Capacitances are implemented on the ICB near the power input connectors as well as near the front-end connectors to suppress ripples and minimize a possible voltage overshoot caused by switching off the FE-hybrids.

For each low voltage group, two high voltage channels are provided. For each HV channel there are up to four single HV lines, which bias one or two silicon modules.

The ICB_46 and ICB_57 have six copper layers, while the smaller boards have only four layers. To limit the contribution to the material budget, the copper layers are rather narrow and thin.

2008 JINST 3 S08004

The layer thickness amounts to 17 and 25 μm for the inner four and outer two layers, respectively, except for the innermost layer of boards ICB_1, ICB_3 and ICB_57 on front petals and ICB_1, ICB_2 and ICB_3 on back petals, which has a thickness of 35 μm. Digital and data traces are shielded by power and ground layers.

Two petals, one back and one front petal, are connected in a control ring. The front petal is the first in the control loop. Both on back and front petals, rings 1–4 and 5–7 are connected to one CCU, respectively. The Digital Opto-Hybrid (DOH) converts the optical signals to electrical LVDS signals and vice versa. Two DOHs are located on a separate PCB, the Digital Opto-Hybrid Module (DOHM), which is mounted on the back petal. From the DOHM, which also distributes the power for the DOHs, electrical signals are transmitted to the CCUMs on the petal. For the control ring, a redundancy scheme is implemented on the ICB. Each CCU can be bypassed electrically in case of a problem, so that the functionality of the control ring is maintained. The second DOH is needed for redundancy purposes only. To allow also the last CCU on the ring to be bypassed, a fifth CCU is located on the DOHM. It is used only in this special case. However, if two consecutive CCUs are faulty, the complete control ring is lost.

Low-pass filters are implemented for the traces of the temperature signals that are brought out via power cables, to ensure that noise is not coupled in via these lines. In addition to the thermistors located on the Kapton of the silicon modules, several temperature and humidity probes are located on or connected to the ICB. Two 10 kΩ NTC thermistors are located on ICB_46 on front petals and read out via the power cable of low voltage group 2. Both on front and back petals, four 10 kΩ NTC thermistors are glued to the cooling inserts of the ring 6 modules. They are read out via the DCU that is present on each CCUM. On both petal types, a humidity sensor can be connected to ICB_46. For back petals, this sensor is read out via the power cable of LV group 2. On each z-side in total 12 hardwired humidity sensors are distributed over the TEC volume. For front petals, the humidity sensor is read out via the DCU on the CCUM. Front petals of all disks of the top and bottom sectors carry these additional humidity sensors, providing detailed information on the relative humidity along the z-direction.

Kapton cables of about 15 cm length are used to link the petals inside one control ring with each other and with the DOHM, providing the electrical digital signals and the power for the CCUMs. These cables consist of two copper layers with a thickness of 35 μm each, separated by a 100 μm thick polyimide layer.

Each TEC LV group is supplied by one so-called multiservice cable, which transmits the analogue power and the bias voltage and brings out signals from temperature or humidity sensors. Inside the tracker support tube, power cables are arranged around the main TEC cooling pipes that run along the z direction, and end at the bulkhead. These cables implement silver-plated aluminium conductors to minimize the impact on the material budget. Typical currents per cable range from about 5 A to 11 A, depending on the number of APVs connected. Therefore three cable types exist, with wire cross-sections tailored to the differing needs.

The connection from the bulkhead to the so-called patch panel 1, located outside of the tracker volume, is provided by power cables implementing tinned copper conductors. The control power is transmitted via separate cables, which also break at the bulkhead. In this case tinned copper conductors are used both inside and outside the tracker volume.

The grounding scheme

The so-called TEC common ground is located at the back end of each TEC. It is realized by means of a 5 cm wide and 150 μm thick copper ring, which is glued to the outer radius of each back disk and tied to the brackets that connect the tracker support tube to the hadron calorimeter. The material of the hadron calorimeter represents a very solid ground. The shields of all cables, the reference points of all power groups, the cooling manifolds that are used to connect the cooling pipes of the petals to the main tubes that are mounted on the TEC, the CF skins of the disks and petals and the outer aluminium shields of the TEC are connected to this TEC common ground. On the petal side, one common analogue ground is implemented per petal. This so-called local petal ground is distributed via a 2 cm wide and 20 μm thick copper path along the ICBs as a reference rail. The LV and HV supplies of all power groups are referenced to this local petal ground at the geometrical/electrical centre of each group. The digital ground of a control group is referenced once to the local petal ground. The local petal ground of each petal is connected to the TEC common ground. Copper strips glued to the outer radii of the disks and along the service channels that connect all disks with the back disk provide the electrical connection to the TEC common ground. These copper strips are connected via short copper braids to the ICBs on the petals. The carbon frames of the silicon detectors are connected via a conductive glue spot to the bias Kapton and finally via the ICB to the FE hybrid ground. To avoid ground loops, the frames are electrically insulated from the cooling pipes by an anodized layer between the cooling inserts and the pipe.

3.3.7 Geometry and alignment

The deviation of true position and orientation of tracker modules from their nominal values as specified in the engineering drawings depends on many factors with different origin, some of them time-dependent: the achieved assembly precision, deformation due to tracker cooling, stress from access and magnetic field, out-gassing of components in dry nitrogen. This leads to a degradation of the track parameter resolution (figure 3.4), which needs to be recovered by determining true module position and orientation, called *alignment*.

Alignment of the tracker relies on three key components: the various data about assembly gathered during the integration process, the Laser Alignment System and the alignment with tracks, ordered by increasing precision and availability with time.

For alignment purposes, modules with two sensors are treated as they would have one large sensor with identical active area coverage. This is justified by sensor mask design [36] and achieved sensor placement accuracy (figure 3.23).

The CMS tracker alignment task thus consists of the determination of three translational and three rotational parameters for each of the 15 148 tracker modules. To achieve ultimate precision, it might be necessary to consider additional parameters, e.g. the sensor bow due to single-sided processing.

Geometry

Two methods are mainly used for measuring tracker component assembly precision: survey with coordinate measurement machines with a typical accuracy of a few μm to a few tens of μm, and

Table 3.2: Estimated assembly precision (RMS, in μm) of tracker components. Values are given with respect to the next level in the hierarchy, e.g. the position accuracy of sensors in modules is 10 μm.

TIB		TID		TOB		TEC	
Sensor	10	Sensor	10	Sensor	10	Sensor	10
Module	180	Module	54	Module	30	Module	20
Shell	450	Ring	185	Rod	100	Petal	70
Cylinder	750	Disc	350	Wheel	140 ($r\phi$), 500 (z)	Disc	150
Tube		Cylinder	450	Tube	1000	TEC	600
		Tube		CMS		Tube	

photogrammetry with an accuracy of 150 μm (80 μm) under good (optimal) conditions for relative measurements. The measured and expected mounting precision from those data are summarized in table 3.2. It should be noted that structure deformations due to loading as well as temperature and humidity variations have not been taken into account.

The software description of the position and orientation of the active detector volumes has been validated with survey data and reconstructed tracks from test beams and cosmic muons recorded in various test and integration setups.

Laser Alignment System

The Laser Alignment System (LAS, figure 3.34) uses infrared laser beams with a wavelength $\lambda = 1075$ nm to monitor the position of selected tracker modules. It operates globally on tracker substructures (TIB, TOB and TEC discs) and cannot determine the position of individual modules. The goal of the system is to generate alignment information on a continuous basis, providing geometry reconstruction of the tracker substructures at the level of 100 μm, which is mandatory for track pattern recognition and for the High Level Trigger. In addition, possible tracker structure movements can be monitored at the level of 10 μm, providing additional input for the track based alignment.

In each TEC, laser beams cross all nine TEC discs in ring 6 (ray 2) and ring 4 (ray 3) on back petals, equally distributed in ϕ. Here special silicon sensors with a 10 mm hole in the backside metalization and an anti-reflective coating are mounted. The beams are used for the internal alignment of the TEC discs. The other eight beams (ray 4), distributed in ϕ, are foreseen to align TIB, TOB, and both TECs with respect to each other. Finally, there is a link to the Muon system (ray 1), which is established by 12 laser beams (six on each side) with precise position and orientation in the tracker coordinate system.

The signal induced by the laser beams in the silicon sensors decreases in height as the beams penetrate through subsequent silicon layers in the TECs and through beam splitters in the align-

2008 JINST 3 S08004

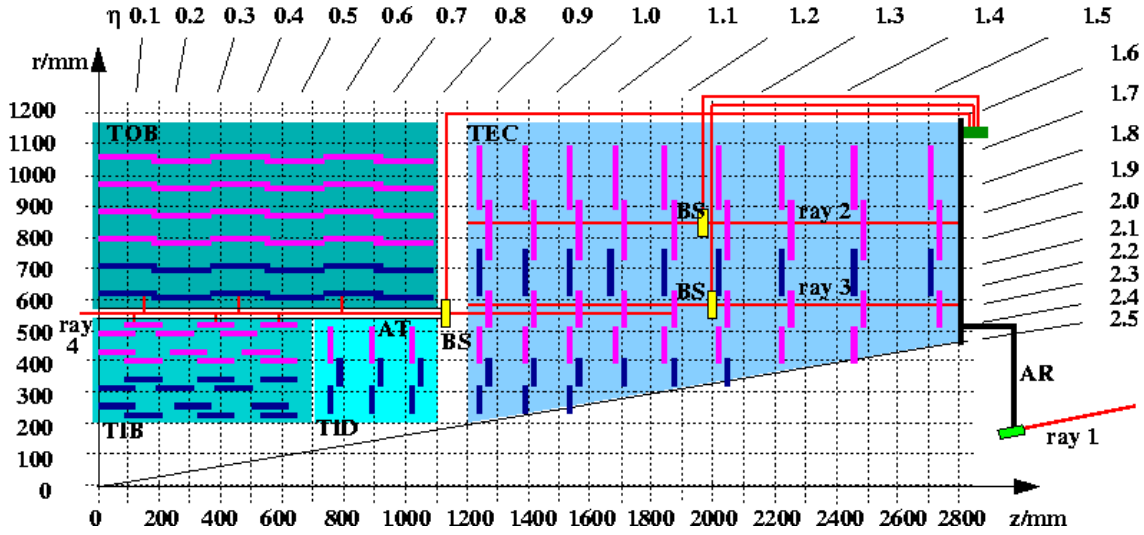

Figure 3.34: Overview of the CMS Laser Alignment System.

ment tubes that partly deflect the beams on TIB and TOB sensors. To obtain optimal signals on all sensors, a sequence of laser pulses with increasing intensities, optimized for each position, is generated. Several triggers per intensity are taken and the signals are averaged. In total, a few hundred triggers are needed to get a full picture of the alignment of the tracker structure. Since the trigger rate for the alignment system is around 100 Hz, this will take only a few seconds. These data will be taken at regular intervals, both in dedicated runs and during physics data taking.

Alignment with tracks

CMS pursues the development of two novel track-based alignment algorithms that allow to quickly solve the system of linear equations of order $\mathcal{O}(100\,000)$. The first is an extension to the well-known global Millepede algorithm [52], that takes all correlations into account and has been shown to successfully align the most sensitive 50 000 parameters. The second is a novel approach using a Kalman Filter [53], which bridges the gap between global and local algorithms by taking into account the most important correlations. In addition the HIP [54] algorithm, which is local in the sense that it takes into account only correlations of parameters within a module, is developed in parallel. In this algorithm, correlations between modules are dealt with implicitly by iterating the alignment many times. All three methods are expected to be able to provide alignment constants for the full silicon pixel and strip tracker.

Experience from other experiments has shown that collision data are not sufficient to constrain certain correlated module movements well enough to obtain a unique set of alignment constants. Therefore complementary data and constraints need to be exploited. Examples are tracks from cosmic muons (with and without magnetic field) that constrain the tracker barrel modules, or beam halo muons for the endcap. Beam gas and minimum bias events are also under consideration. Typical examples of constraints are a vertex constraint for decay particles e.g. from $Z \rightarrow \mu^+\mu^-$ or jets, mass constraints, measurements from the Laser Alignment System, and survey constraints. First studies indicate that those data will provide a unique alignment parameter set.

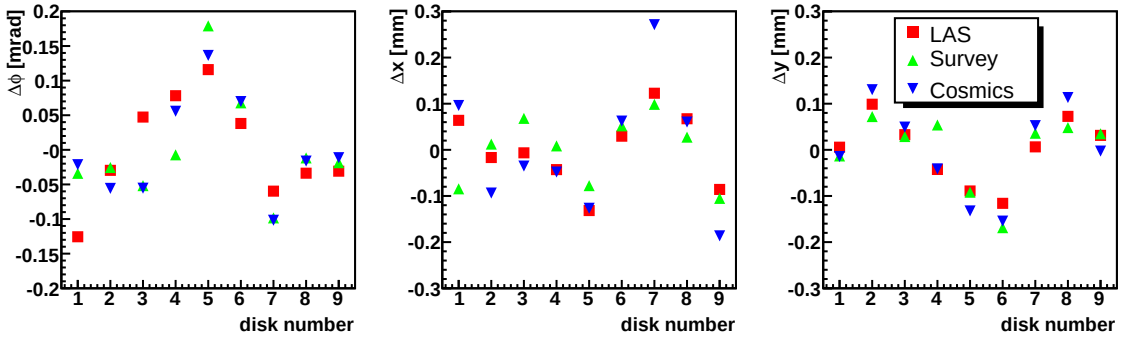

Figure 3.35: TEC+ disk rotation $\Delta\phi$ (around the beam axis) and displacements $\Delta x, \Delta y$ (in the disk plane) as determined from survey, LAS and cosmic muon tracks.

During integration of the TEC+, deviation of disk positions and rotations from nominal values have been determined from survey with photogrammetry, the LAS, and tracks from cosmic muons. Figure 3.35 shows the results from the three complementary methods. The global degrees of freedom (absolute position and orientation, torsion and shear around the symmetry axis) have been fixed by requiring the average displacement and rotation as well as torsion and shear to be zero. The values agree within 60 μm and 80 μrad with each other, which can be taken as an upper value on the precision of each method.

3.3.8 Detector control and safety system

The Tracker Detector Safety System (TDSS) and tracker Detector Control System (tracker DCS) is a two pillar system. The TDSS ensures independently the safety, with a large PLC (Programmable Logical Controller) system, occupying 6 LHC racks. A limited set of around 1000 hardwired temperature and humidity sensors are evaluated and out of limit states interlock power supplies. The tracker DCS, as a complementary partner, controls, monitors and archives all important parameters. The heart of the DCS is composed out of an industrial SCADA program (Supervisory Control And Data Acquisition) PVSS (Prozessvisualisierungs- und Steuerungssystem from ETM Austria, chapter 9) together with a Finite State Machine written in SMI++, a derivative of the former DELPHI control software; thus using the standard control software framework for all LHC experiments. The main task of the DCS is to control about 2000 power supplies for silicon module low and high voltage power and about 100 low voltage control power supplies via the OPC (OLE for Process Automation) protocol. Detector interdependencies of control, low and high voltages are handled, as well as fast ramp downs in case of higher than allowed temperatures or currents in the detector, experimental cavern problems, etc. All this is ensured by evaluating 10^4 power supply parameters, 10^3 data points from DSS via a Siemens S7 driver and 10^5 readings from the DCUs situated on all front end hybrids and control units CCUs. Several passive alarms and warning levels are defined for temperature, relative humidity, voltages, currents, etc. and are reported in a global warning panel as well as limits that, if surpassed, would result in automatic shutdown. Information from the tracker cooling plant, the thermal screen, beam conditions and the dry gas system are crucial for safe running and are accessible from the tracker DCS and TDSS.

2008 JINST 3 S08004

All parameters are archived to ORACLE. The TDSS (tracker DCS) system is fully implemented in the global CMS DSS (DCS) and Run Control system.

3.3.9 Operating experience and test results

Performance in test beam experiments

The system performance of integrated structures of the silicon strip tracker and its data acquisition chain as well as the performance of the silicon strip modules themselves has been studied in various test beam experiments at CERN and the Paul Scherrer Institut (PSI), Villigen (CH). In test beam campaigns, performed in May and October 2004 at the X5 test beam complex in the CERN west area, large substructures of TIB, TOB and TEC were exposed to a secondary pion beam with an energy of 120 GeV and a tertiary muon beam with muon energies ranging from 70 to 120 GeV. The TIB setup comprised a prototype half-shell structure of layer 3, equipped with eight single-sided strings, plus four double-sided strings, mounted on a custom support structure. For the TOB, the so-called cosmic rack, a precise mechanical telescope-like structure equipped with four single-sided and two double-sided rods, was used in the beam tests. The TEC setup consisted of one back and one front petal [55]. These setups corresponded to about 1% of the complete TIB, TOB and TEC detectors, respectively. The TOB and TEC setups were operated at a temperature below -10°C, while the TIB setup was operated at room temperature. Typical primary trigger rates for the pion beam were 600 000 pions per spill (a 2.2 s long period within a 12 s long SPS cycle during which particles are delivered) corresponding to a mean occupancy of 15 Hz/cm^2.

In the strip-cluster finding the cuts for the signal-to-noise ratio, S/N, of the cluster seed / neighbour strips / total cluster are 4/3/5 for TIB, 5/2/5 for TOB and 3/2/5 for TEC, respectively. The cluster noise is calculated by adding the single strip noise values in quadrature (TIB, TEC) or by taking the seed noise as the cluster noise (TOB). To determine the most probable value for the S/N of a module, a Landau distribution convoluted with a Gaussian is fitted to the signal-to-noise distribution, and the most probable value of the fitted function is quoted as the S/N.

The mean most probable S/N values for all module types, together with their strip length, pitch and abbreviations used in the following, are summarized in table 3.3. For thin (thick) TEC sensors, most probable S/N values of 29–33 (36–42) in peak mode and 19–22 (20–24) in deconvolution mode have been observed [55]. For the thick TOB OB1 (OB2) modules a S/N of typically 36 (38) and 25 (27) was found in peak and deconvolution mode, respectively [17], while the thin TIB IB1 (IB2) modules exhibited a S/N of 26 (30) in peak mode and 18 (20) in deconvolution mode.

Assuming that a MIP creates 24 000 electrons in a 300 μm thick layer of silicon [16], and assuming that the beam particles can be treated as MIPs, the S/N can be used to calculate the equivalent noise charge, ENC. The common mode subtracted noise depends on the capacitance of the sensor, which depends linearly on the strip length and the ratio between strip width and pitch, w/p [16]. Since $w/p = 0.25$ for all sensor types, the ENC varies between different module types according to the strip length. Results for all module types except W1TID are summarized in table 3.3. Measurements performed at low temperature (for the TEC, typically hybrid temperatures of $+10$°C and 0°C were reached for hybrids with six and four APVs, respectively) are plotted versus the strip length in figure 3.36. A linear fit to these data yields the following dependence of

the ENC on the strip length L:

$$\begin{aligned} \mathrm{ENC_{peak}} &= (36.6 \pm 1.9)\, e^-/\mathrm{cm} \cdot L + (405 \pm 27)\, e^-, \\ \mathrm{ENC_{dec}} &= (49.9 \pm 3.2)\, e^-/\mathrm{cm} \cdot L + (590 \pm 47)\, e^-. \end{aligned}$$

The common mode noise is the standard deviation of the common mode, calculated per APV from a certain number of events. The mean common mode noise has been evaluated and amounts to (173 ± 38) and (299 ± 76) electrons for TEC (mean from all APVs in the setup) and (265 ± 36) and (300 ± 19) electrons for TIB (mean from all APVs of TIB2 modules) in peak and deconvolution mode, respectively.

Although no dedicated beam telescope was available, efficiency studies have been performed both with the TOB and TEC setups, exploiting the fact that in both cases the beam penetrated several layers of modules. Efficiencies of above 99% have been observed in all such studies.

The uniformity of the module performance along and perpendicular to the strip direction has been studied in 2003 with several TIB modules in a test beam experiment at the X5 complex. Two single-sided strings equipped with IB2 modules were mounted on a structure corresponding to a portion of a layer 3 half-shell, and operated at room temperature. To study the uniformity across the strips, the strips read out by three APVs (on two different modules) were exposed to a pion beam, and between 1000 and 8000 events were collected per strip. A cluster was associated to a strip if the centre of gravity x of the cluster was reconstructed within $(n-0.5) \cdot p < x < (n+0.5) \cdot p$ for strip n and pitch p. The uniformity, defined as the ratio between the RMS and the mean of the respective distribution, was 1.3% for the cluster noise, with an increase close to the APV chip edges. The cluster charge uniformity was of the order of 1.4%, but dropped to 0.5% if calculated separately for groups of 32 adjacent strips. A uniformity of the S/N of 1.6% on average and of 1.0% for groups of 32 strips was measured. To investigate the uniformity along the strips, a muon beam was used for its uniform particle density. The cluster position along the strip could be obtained from the TOB setup that was operated in the same test beam, since the strip direction of the TOB modules was perpendicular to that of the TIB modules. The clusters were binned in 24 intervals according to their centre of gravity, corresponding to length intervals of 5 mm, and about 1500 events were accumulated per bin. Both the uniformity of cluster charge and S/N were found to be 1.4%.

Performance during integration

Testing during integration consisted typically of checks of the control ring functionality, tests of the I^2C communication of all chips, tests of the gain of the optical connections, commissioning (i.e. tuning of chip operation parameters), pedestal runs in peak and deconvolution mode, bias voltage ramping up to 450 V, read-out of currents and module and hybrid temperatures through the DCUs, and a functionality check of the temperature and humidity sensors.

In the following sections the performance of TEC, TIB/TID and TOB during integration is described. Two comments apply to all three sub-detectors:

- Numbers of dead and noisy strips are given below. While dead strips can be identified reliably, the noisiness of strips depends on external conditions such as grounding and the APV read-out mode and the figures given should be regarded as estimates only. APV edge strips

2008 JINST 3 S08004

Table 3.3: Pitch, strip length, signal-to-noise ratio and equivalent noise charge after common mode subtraction for different module types. The TEC and TOB measurements are for hybrid temperatures of below 0 °C, the TIB measurements were performed at room temperature. Sensors of type IB1 and IB2 are used in TIB, layers 1 and 2 and layers 3 and 4, respectively. In the TOB, layers 1–4 are equipped with OB2 sensors, layers 5 and 6 with OB1 sensors. The sensor geometries abbreviated with W are wedge-shaped sensors used in TEC and TID, with the number corresponding to the ring. W1 sensors have a slightly different geometry in TID and TEC.

Module type	Pitch [μm]	Strip length [mm]	S/N Peak mode	S/N Dec. mode	ENC [e^-] Peak mode	ENC [e^-] Dec. mode
IB1	80	116.9	25.8 ± 1.3	18.3 ± 0.5	931 ± 48	1315 ± 37
IB2	120	116.9	29.5 ± 1.4	20.3 ± 0.6	815 ± 37	1182 ± 31
OB1	122	183.2	36	25	1110 ± 47	1581 ± 75
OB2	183	183.2	38	27	1057 ± 17	1488 ± 22
W1TEC	81–112	85.2	33.1 ± 0.7	21.9 ± 0.6	714 ± 23	1019 ± 37
W2	113–143	88.2	31.7 ± 0.5	20.7 ± 0.4	741 ± 25	1068 ± 51
W3	123–158	110.7	29.2 ± 0.6	20.0 ± 0.4	802 ± 16	1153 ± 48
W4	113–139	115.2	28.6 ± 0.5	19.2 ± 0.3	819 ± 21	1140 ± 26
W5	126–156	144.4	42.2 ± 1.1	24.1 ± 1.1	971 ± 29	1354 ± 57
W6	163–205	181.0	37.8 ± 0.6	23.0 ± 0.4	1081 ± 26	1517 ± 47
W7	140–172	201.8	35.5 ± 1.0	20.3 ± 1.1	1155 ± 40	1681 ± 107

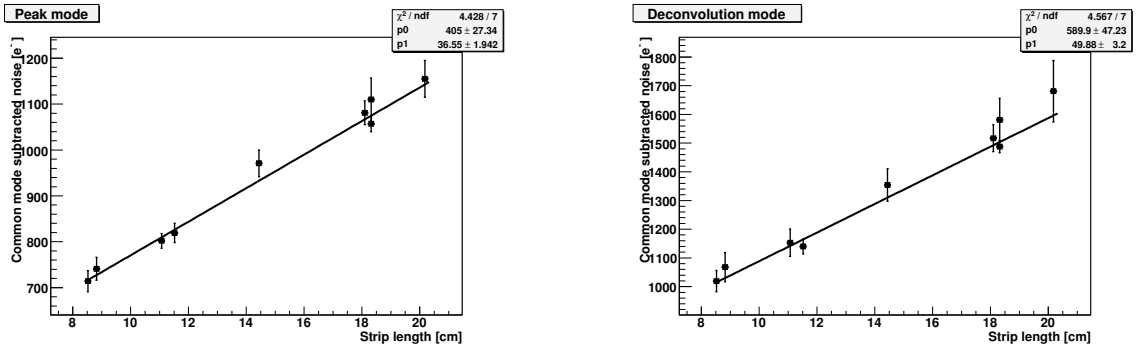

Figure 3.36: Equivalent noise charge after common mode subtraction versus strip length for all TOB and TEC module types, in peak (left panel) and deconvolution mode (right panel).

show typically an increased noise and are frequently flagged as noisy, especially when a fixed noise cut is used for all strips. These edge strips are included in the numbers of flagged strips, although they are usually fully efficient.

- Although all components (petals, rods, single modules in case of the TIB/TID) were tested before insertion and components not fulfilling strict quality criteria were rejected, several defects have been observed during integration. Typical defects are broken optical fibers, bad APVs (i.e. with many noisy or dead strips), and missing or unreliable I^2C communication

2008 JINST 3 S08004

of complete modules or single chips. Most of the problems are assumed to be caused by mishandling during insertion or cabling. Since the exchange of components bears a considerable risk, not all defective components have been exchanged. Additional defects could be introduced by any following handling step, such as cabling of the tracker. The numbers given below should thus be regarded as a snapshot reflecting the situation right after integration of the single sub-detectors.

TEC Performance during integration

The TEC petals were integrated sector-wise, where one sector corresponds to one eighth of a TEC in ϕ, and comprises 18 petals that share nine control rings and four cooling circuits. After integration of one sector, a read-out test of the full sector was performed at room temperature, with the coolant at $+15$°C and mean silicon and hybrid temperatures of about $+23$°C and $+33$°C, respectively.

During integration a flaw in the crimping of the connectors of the multiservice power cables was found. After all such connectors had been replaced on both TECs, the system performance observed during integration was very robust. In figure 3.37, left side, the common mode subtracted noise of all strips of both TECs is shown for deconvolution mode. Since the measured noise depends on the gain of the optical chain, the noise was normalized to the digital output of the APV (scale on upper x-axis in figure 3.37). In addition, the number of ADC counts in the FED was converted to ENC according to the following method: with a nominal digital APV output of $\pm\,4$ mA and a nominal APV gain of 1 MIP/mA for thin sensors, the height of the digital output corresponds to 8 MIPs or 200 000 electrons. This method allows a direct comparison of the measurements from different optical channels and delivers an approximate absolute calibration of the equivalent noise charge. Cross-checks with cosmic muon data performed during TIB/TID integration indicate that this scaling agrees with the real ENC within 10–20%. Furthermore, the noise depends on the strip capacitance and thus on the strip length, i.e. on the module type. For this reason the noise of all strips was normalized to the strip length of modules of ring 1 (8.52 cm). In addition a correction was applied to TEC- data to account for the fact that they were taken with other chip parameter settings than TEC+ data. The common mode subtraction was performed assuming a constant common mode per APV. To extract the mean noise, a gaussian was fitted to the distribution. The resulting mean common mode subtracted noise amounts to 1693 ± 75 electrons in this normalization.

The mean common mode noise, calculated per APV, amounts to (22 ± 4)% and (21 ± 3)% of the mean intrinsic noise in peak and deconvolution mode, respectively (figure 3.38, left, for all non-defective APVs of TEC+).

The flatness of the noise across the APV is a good indicator for the quality of the grounding. The relative spread of the total noise (before common mode subtraction), i.e. the RMS of the noise divided by the mean noise, both calculated per APV, can be used to quantify the flatness. The relative spread is (2.5 ± 0.2)% in both read-out modes, as shown in figure 3.38, right, indicating that the grounding scheme implemented by the TEC works well.

Strips are counted as noisy or dead if their noise is more than five times the RMS of the noise above or below the mean noise of the respective APV. Edge strips are counted as noisy, if

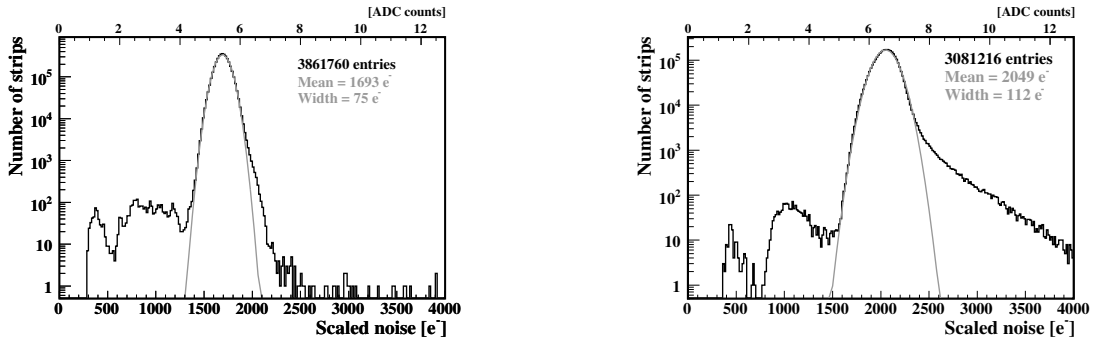

Figure 3.37: Normalized common mode subtracted noise of all strips (scaled to the strip length of ring 1 sensors) of both TECs (left panel) and the TOB (right panel), in deconvolution mode. Details are described in the text.

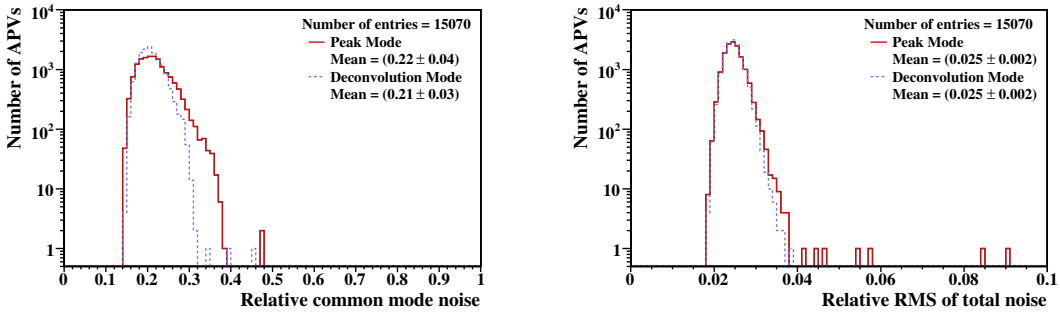

Figure 3.38: Ratio between common mode noise and mean intrinsic noise (left panel) and ratio between the RMS of the total noise and the mean total noise (right panel), calculated per APV, in peak and deconvolution mode for all non-defective APVs of TEC+.

their noise is more than seven sigma above the mean noise. In total, there are 3.0 per mille of bad channels in TEC+, while TEC- has 2.7 per mille of bad channels.

TIB and TID performance during integration

During TIB/TID integration [56], modules and AOHs were assembled onto half layers and disks and tested extensively for functionality, including pedestals, once a mother cable was completed (corresponding to a string in the TIB and three single-sided or five double-sided modules in the TID). Completed disks and half layers were then subjected to a burn-in in a climatic chamber, during which the structures were operated at a silicon sensor temperature of about $-15°$C. The complete half layers and disks were read out during these tests. Typically, the structures underwent 2–3 cooling cycles during a five day measurement period. After-wards disks and half layers were assembled into the complete TIB/TID+ and TIB/TID- structures and shipped to CERN, where the last integration operations were performed, such as connection of fibers to the final ribbons and cabling of the margherita.

After optimization of the grounding scheme, the noise performance observed in the TIB and TID structures was very good. For the TIB and TID structures the scaled common mode subtracted noise of all strips, except for two TIB half-shells for which data were taken under non-final running conditions and three half-shells for which the proper grounding scheme was not yet implemented, is shown in figure 3.39 for deconvolution mode. Scaling and common mode subtraction have been implemented as previously described in this section. These data have been taken under nominal CMS conditions, with a mean silicon sensor temperature of about -15°C, hybrid temperatures ranging from -4°C (TID double-sided modules) to -14°C (TIB single-sided modules) and APV parameters set as intended for this temperature range. The mean noise, taken from a gaussian fit, amounts to (1233 ± 87) electrons in the TIB and (1246 ± 76) electrons in the TID. Measurements with a silicon sensor temperature of about $+10$°C and hybrid temperatures of $+10$°C to $+30$°C show a mean noise about 20% larger. In contrast to the TEC, a strip is flagged as dead if its noise is below 75% of the average noise of the APV, and APV edge strips are not treated differently. The total number of bad channels is 4.4 per mille in TIB/TID+ and 3.4 per mille in TIB/TID-.

TOB performance during integration

Fully equipped and tested rods were integrated cooling segment-wise. After a first functional test, the cooling connection was soldered and a leak test was performed. Then the cooling segment was cabled, and a full read-out test, including pedestals, was performed at room temperature. During these measurements, the silicon sensor temperature was about $+24$°C and the hybrid temperature about $+30$°C.

During integration, a sensitivity to pick-up noise has been observed, which leads to non-flat, wing-like common mode subtracted noise distributions. This sensitivity is especially pronounced for layers 3 and 4, which are equipped with single-sided 4 APV modules, and within these layers the effect is worst for modules mounted closest to the CCUM. Defining as a figure of merit the ratio of the highest noise amplitude (taken from a parabola fit to the noise distribution) to the flat noise baseline, and counting all APVs with a ratio above 1.25 as "in the wings", the fraction of APVs in the wings is about 30% in layers 3 and 4 and about 7% and 1% in layers 1/2 and 5/6, respectively. In total, 11.4% of all APVs are found to be in the wings according to this criterion. It has been verified that either with adjusted cluster cuts or with a linear online common mode subtraction the increase in the cluster width and occupancy is negligible.

The normalized noise of all TOB strips is shown in figure 3.37, right. The tail to high noise values comes from the non-flat noise distributions. The mean noise from a gaussian fit amounts to (2049 ± 112) electrons.

Due to this wing-like noise structure, a special algorithm has been adopted to evaluate the number of dead and noisy strips. A parabola is fitted to the noise distribution of each APV in an iterative procedure, and strips are flagged as bad if their noise deviates more than ten times the RMS of the distribution of fit residuals from the fitted function.

Only very few permanent defects, corresponding to 0.6 per mille of lost channels, have been introduced during TOB integration. Including the number of noisy and dead strips, the number of bad channels amounts to 0.6 per mille in TOB+ and 1.9 per mille in TOB-.

2008 JINST 3 S08004

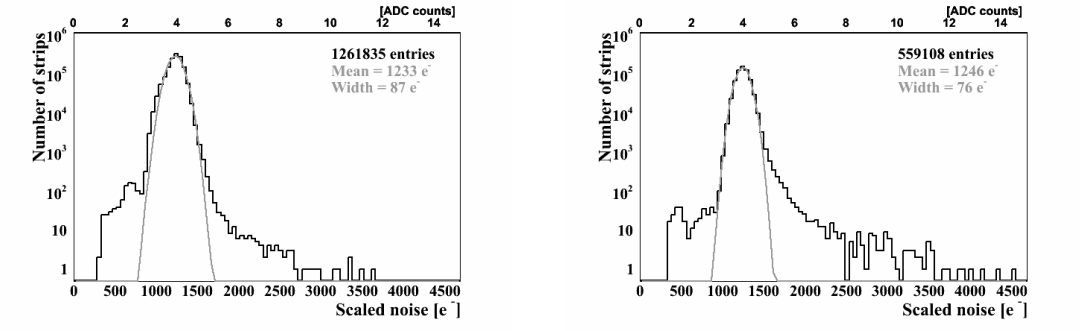

Figure 3.39: Normalized common mode subtracted single strip noise for TIB (left panel) and TID (right panel), in deconvolution mode. Details are described in the text.

Irradiation studies

As already discussed in detail in section 3.1.1, the silicon strip tracker will suffer from a severe level of radiation during its 10 year long lifetime: up to $1.8 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$ for TIB/TID and TEC and up to $0.5 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$ for TOB, assuming an integrated luminosity of $500\,\mathrm{fb^{-1}}$. The radial and z dependence of the fluence both for fast hadrons and neutrons is described in detail in [15, 16]. Hadrons are expected to dominate in the inner part of the tracker, up to a radius of about 0.5 m, while neutrons backscattered off the electromagnetic calorimeter dominate further outside. Safety factors of 1.5 and 2.0 on the fluence are typically applied for TIB/TID and TOB/TEC, respectively.

To ensure that both the FE electronics and the silicon sensors can be operated safely and with satisfactory performance after such an irradiation, several irradiation tests with neutrons and protons have been carried out. Neutron irradiation was usually performed at the isochronous cyclotron of the Centre de Recherches du Cyclotron, Louvain-la-Neuve, which delivers neutrons with a mean energy of 20 MeV (hardness factor 1.95 relative to 1 MeV neutrons [58]). Proton irradiation has been carried out e.g. at the compact cyclotron of the Forschungszentrum Karlsruhe, where a 26 MeV proton beam (hardness factor 1.85 relative to 1 MeV neutrons) with a current of $100\,\mu\mathrm{A}$ and a beam spot diameter of 1 cm is available.

To study the performance of complete irradiated modules, several OB1 and OB2 modules (table 3.3 for explanation) were irradiated with a proton fluence ranging from $0.1 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$ to $0.7 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$, and one OB2 module was subjected to a neutron fluence of about $1.2 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$ [57]. Two TEC W5 modules were irradiated with a proton fluence of about $1.1 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$, and three TIB IB1 modules were subjected to a proton fluence of $0.5 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$ to $2.1 \times 10^{14}\,\mathrm{n_{eq}cm^{-2}}$. The effect of annealing was simulated by heating the modules for 80 minutes at 60°C and afterwards storing them at room temperature for at least two hours. To prevent uncontrolled annealing, the modules were stored at −20°C between the irradiation or annealing steps. Measurements were performed at −15°C.

As expected from inversion from n- to p-type doping, the full depletion voltage increased with the fluence, as shown in figure 3.40 (left). However, the required depletion voltage stays below 500 V, which is the maximum depletion voltage for which the sensors are specified. The dependence of the depletion voltage on annealing time was studied as well and found to be in

2008 JINST 3 S08004

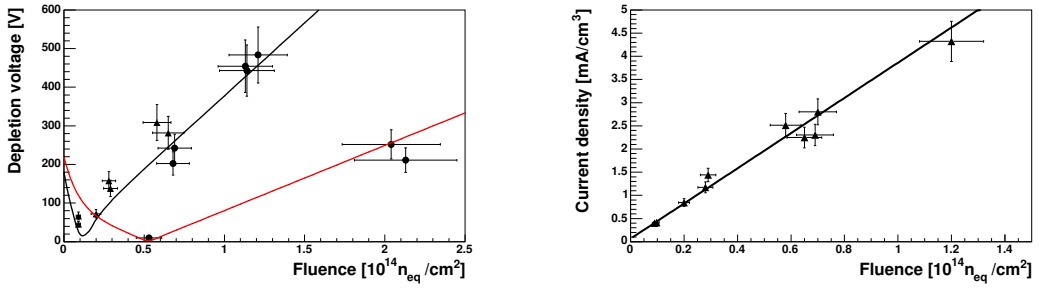

Figure 3.40: Left panel: Variation of depletion voltage with fluence for OB1 (triangles), OB2 and W5 (dots on upper curve) and IB2 (dots on lower curve) modules, after an annealing time of 80 minutes after each irradiation step. The curves correspond to calculations for 500 μm (upper curve) and 320 μm (lower curve) sensors for an annealing time of 80 minutes at 60°C. Right panel: Current density, scaled to 20°C, versus fluence after annealing for 80 minutes at 60°C.

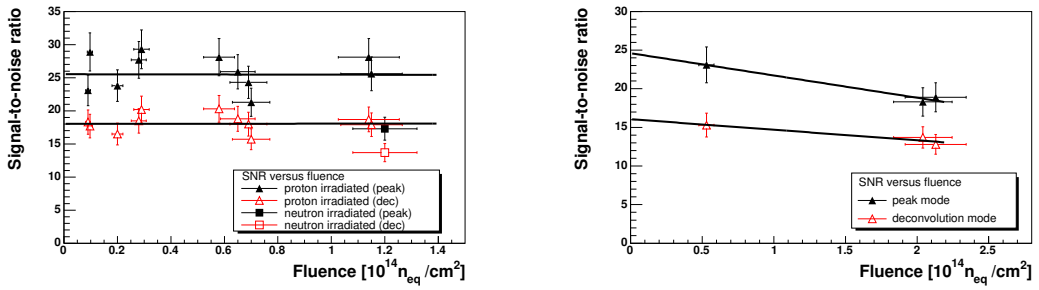

Figure 3.41: Signal-to-noise ratio versus fluence for modules with 500 μm (left panel) and 320 μm thick sensors (right panel) in peak (filled symbols) and deconvolution mode (open symbols).

excellent agreement with the Hamburg model [59], with a minimum at 80 minutes annealing time, corresponding to a 10 day shut down period at room temperature.

The leakage current is expected to increase with fluence, leading to a larger heat dissipation and increased noise. In figure 3.40 (right) the dependence of the current density on the fluence is shown. The current related damage rate, defined as the current increase, scaled to 20°C, per sensor volume and equivalent neutron fluence, amounts to $(3.79 \pm 0.27) \times 10^{-17}$ A/cm, which is in good agreement with literature and measurements from test structures.

Measurements of the signal-to-noise ratio, S/N, of irradiated modules have been performed with a ^{90}Sr source. Due to an increase of the noise and a decrease of the charge collection efficiency, the S/N is expected to decrease with fluence. The dependence of the S/N on the accumulated fluence for thick and thin sensors in both read-out modes is shown in figure 3.41. For thick sensors, the S/N decreased from 23 (35) to 15 (21) in deconvolution (peak) mode, while for thin sensors a decrease from 18 (24) to 13 (18) was observed. These figures ensure a hit finding efficiency of above 95% even after 10 years of operation at the LHC [60, 61].

2008 JINST 3 S08004