

The CMS experiment at the CERN LHC

CMS Collaboration

Abstract

The Compact Muon Solenoid (CMS) detector is described. The detector operates at the Large Hadron Collider (LHC) at CERN. It was conceived to study proton-proton (and lead-lead) collisions at a centre-of-mass energy of 14 TeV (5.5 TeV) and at luminosities up to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ ($10^{27} \text{ cm}^{-2}\text{s}^{-1}$). The detector is characterised by high hermeticity. A high-field superconducting solenoid surrounds a silicon tracker, an electromagnetic calorimeter, and a hadron calorimeter. The return yoke is instrumented with muon detectors covering most of the 4π solid angle. Forward sampling calorimeters extend the CMS pseudorapidity coverage at high $|\eta|$ values.

Submitted to the Journal of Instrumentation (JINST)

Contents

1	Introduction	1
2	General concept	2
3	Superconducting magnet	5
3.1	Overview	5
3.2	Main features of the magnet components	5
3.2.1	Superconducting solenoid	5
3.2.2	Yoke	6
3.2.3	Electrical scheme	7
3.2.4	Vacuum system	8
3.2.5	Cryogenic plant	8
3.2.6	Other ancillaries	9
3.3	Operating test	10
3.3.1	Cool-down	10
3.3.2	Charge and discharge cycles	10
3.3.3	Cold mass misalignment	11
3.3.4	Electrical measurements	12
3.3.5	Yoke mechanical measurements	13
3.3.6	Coil stability characteristics	14
3.3.7	Coil warm-up	15
4	Inner tracking system	23
4.1	Introduction	23
4.1.1	Requirements and operating conditions	23
4.1.2	Overview of the tracker layout	25
4.1.3	Expected performance of the CMS tracker	27
4.1.4	Tracker system aspects	27
4.2	Pixel detector	28
4.2.1	Pixel system general	28
4.2.2	Sensor description	29
4.2.3	Pixel detector read-out	31
4.2.4	The pixel barrel system	34
4.2.5	The forward pixel detector	37
4.2.6	Power supply	43
4.2.7	Cooling	44
4.2.8	Slow controls	44
4.3	Silicon strip tracker	44

4.3.1	Silicon sensors	44
4.3.2	Read-out system	46
4.3.3	Silicon modules	49
4.3.4	Tracker Inner Barrel and Disks (TIB/TID)	51
4.3.5	Tracker Outer Barrel (TOB)	53
4.3.6	Tracker EndCaps (TEC)	56
4.3.7	Geometry and alignment	61
4.3.8	Detector control and safety system	63
4.3.9	Operating experience and test results	63
5	Electromagnetic calorimeter	70
5.1	Lead tungstate crystals	70
5.2	The ECAL layout and mechanics	71
5.3	Photodetectors	73
5.3.1	Barrel: avalanche photodiodes	73
5.3.2	Endcap: vacuum phototriodes	74
5.4	On-detector electronics	76
5.5	Off-detector electronics	77
5.5.1	Global architecture	77
5.5.2	The trigger and read-out paths	79
5.5.3	Algorithms performed by the trigger primitive generation	79
5.5.4	Classification performed by the selective read-out	80
5.6	Preshower detector	80
5.6.1	Geometry	80
5.6.2	Preshower electronics	81
5.7	ECAL detector control system	82
5.7.1	Safety system	82
5.7.2	Temperature	82
5.7.3	Dark current	82
5.7.4	HV and LV	83
5.8	Detector calibration	83
5.9	Laser monitor system	84
5.9.1	Laser-monitoring system overview	86
5.10	Energy resolution	87
6	Hadron calorimeter	93
6.1	Barrel design (HB)	93
6.2	Endcap design (HE)	100
6.3	Outer calorimeter design (HO)	106

6.4	Forward calorimeter design (HF)	112
6.5	Read-out electronics and slow control	113
6.6	HF luminosity monitor	120
7	Forward detectors	121
7.1	CASTOR	121
7.2	Zero degree calorimeter (ZDC)	122
8	The muon system	125
8.1	Drift tube system	127
8.1.1	General description	127
8.1.2	Technical design	129
8.1.3	Electronics	134
8.1.4	Chamber assembly, dressing, and installation	138
8.1.5	Chamber performance	142
8.2	Cathode strip chambers	153
8.2.1	Chamber mechanical design	155
8.2.2	Electronics design	157
8.2.3	Performance	161
8.3	Resistive Plate Chamber system	169
8.3.1	Detector layout	169
8.3.2	Readout electronics	172
8.3.3	Low voltage and high voltage systems	175
8.3.4	Temperature control system	177
8.3.5	Gas system	177
8.3.6	Chamber construction and testing	180
8.4	Optical alignment system	186
8.4.1	System layout and calibration procedures	187
8.4.2	Geometry reconstruction	192
8.4.3	System commissioning and operating performance	192
9	Trigger	195
9.1	Calorimeter trigger	196
9.2	Muon trigger	198
9.3	Global Trigger	202
9.4	Trigger Control System	203
10	Data Acquisition	205
10.1	Sub-detector read-out interface	207
10.2	The Trigger Throttling System and sub-detector fast-control interface	208

10.3	Testing	209
10.4	The Event Builder	209
10.5	The Event Filter	212
10.6	Networking and Computing Infrastructure	214
10.7	DAQ software, control and monitor	215
10.8	Detector Control System	218
11	Detector infrastructures and safety systems	222
11.1	Detector powering	222
11.2	Detector cooling	222
11.2.1	Front-end electronics cooling	222
11.2.2	Cryogenics	223
11.3	Detector cabling	223
11.4	Detector moving system	224
11.4.1	Sliding system	224
11.4.2	Pulling system	224
11.5	Beam and Radiation Monitoring systems	224
11.5.1	Introduction	224
11.5.2	Protection systems	226
11.5.3	Monitoring systems	227
12	Computing	230
12.1	Overview	230
12.2	Application Framework	230
12.3	Data formats and processing	231
12.4	Computing centres	232
12.5	Computing services	234
12.6	System commissioning and tests	236
13	Conclusions	237

1 Introduction

The Compact Muon Solenoid (CMS) detector is a multi-purpose apparatus due to operate at the Large Hadron Collider (LHC) at CERN. The LHC is presently being constructed in the already existing 27-km LEP tunnel in the Geneva region. It will yield head-on collisions of two proton (ion) beams of 7 TeV (2.75 TeV per nucleon) each, with a design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ ($10^{27} \text{ cm}^{-2}\text{s}^{-1}$). This paper provides a description of the design and construction of the CMS detector. CMS is installed about 100 metres underground close to the French village of Cessy, between Lake Geneva and the Jura mountains.

The prime motivation of the LHC is to elucidate the nature of electroweak symmetry breaking for which the Higgs mechanism is presumed to be responsible. The experimental study of the Higgs mechanism can also shed light on the mathematical consistency of the Standard Model at energy scales above about 1 TeV. Various alternatives to the Standard Model invoke new symmetries, new forces or constituents. Furthermore, there are high hopes for discoveries that could pave the way toward a unified theory. These discoveries could take the form of supersymmetry or extra dimensions, the latter often requiring modification of gravity at the TeV scale. Hence there are many compelling reasons to investigate the TeV energy scale.

The LHC will also provide high-energy heavy-ion beams at energies over 30 times higher than at the previous accelerators, allowing us to further extend the study of QCD matter under extreme conditions of temperature, density, and parton momentum fraction (low- x).

Hadron colliders are well suited to the task of exploring new energy domains, and the region of 1 TeV constituent centre-of-mass energy can be explored if the proton energy and the luminosity are high enough. The beam energy and the design luminosity of the LHC have been chosen in order to study physics at the TeV energy scale. A wide range of physics is potentially possible with the seven-fold increase in energy and a hundred-fold increase in integrated luminosity over the previous hadron collider experiments. These conditions also require a very careful design of the detectors.

The total proton-proton cross-section at $\sqrt{s} = 14 \text{ TeV}$ is expected to be roughly 100 mb. At design luminosity the general-purpose detectors will therefore observe an event rate of approximately 10^9 inelastic events/s. This leads to a number of formidable experimental challenges. The online event selection process (*trigger*) must reduce the huge rate to about 100 events/s for storage and subsequent analysis. The short time between bunch crossings, 25 ns, has major implications for the design of the read-out and trigger systems.

At the design luminosity, a mean of about 20 inelastic collisions will be superimposed on the event of interest. This implies that around 1000 charged particles will emerge from the interaction region every 25 ns. The products of an interaction under study may be confused with those from other interactions in the same bunch crossing. This problem clearly becomes more severe when the response time of a detector element and its electronic signal is longer than 25 ns. The effect of this pile-up can be reduced by using high-granularity detectors with good time resolution, resulting in low occupancy. This requires a large number of detector channels. The resulting millions of detector electronic channels require very good synchronization.

The large flux of particles coming from the interaction region leads to high radiation levels, requiring radiation-hard detectors and front-end electronics.

The detector requirements for CMS to meet the goals of the LHC physics programme can be summarised as follows:

- Good muon identification and momentum resolution over a wide range of momenta and angles, good dimuon mass resolution ($\approx 1\%$ at 100 GeV), and the ability to determine unambiguously the charge of muons with $p < 1 \text{ TeV}$;
- Good charged particle momentum resolution and reconstruction efficiency in the inner tracker. Efficient triggering and offline tagging of τ 's and b -jets, requiring pixel detectors close to the interaction region;
- Good electromagnetic energy resolution, good diphoton and dielectron mass resolution ($\approx 1\%$ at 100 GeV), wide geometric coverage, π^0 rejection, and efficient photon and lepton isolation at high luminosities;
- Good missing-transverse-energy and dijet-mass resolution, requiring hadron calorimeters with a large hermetic geometric coverage and with fine lateral segmentation.

The design of CMS, detailed in the next Section, meets these requirements. The main distinguishing features

of CMS are a high-field solenoid, a full-silicon-based inner tracking system, and a homogeneous scintillating-crystals-based electromagnetic calorimeter.

The coordinate system adopted by CMS has the origin centered at the nominal collision point inside the experiment, the y -axis pointing vertically upward, and the x -axis pointing radially inward toward the center of the LHC. Thus, the z -axis points along the beam direction toward the Jura mountains from LHC Point 5. The azimuthal angle ϕ is measured from the x -axis in the x - y plane and the radial coordinate in this plane is denoted by r . The polar angle θ is measured from the z -axis. Pseudorapidity is defined as $\eta = -\ln \tan(\theta/2)$. Thus, the momentum and energy transverse to the beam direction, denoted by p_T and E_T , respectively, are computed from the x and y components. The imbalance of energy measured in the transverse plane is denoted by E_T^{miss} .

2 General concept

An important aspect driving the detector design and layout is the choice of the magnetic field configuration for the measurement of the momentum of muons. Large bending power is needed to measure precisely the momentum of high-energy charged particles. This forces a choice of superconducting technology for the magnets.

The overall layout of CMS [1] is shown in Fig. 1. At the heart of CMS sits a 13-m-long, 6-m-inner-diameter, 4-T superconducting solenoid providing a large bending power (12 Tm) before the muon bending angle is measured by the muon system. The return field is large enough to saturate 1.5 m of iron, allowing 4 muon *stations* to be integrated to ensure robustness and full geometric coverage. Each muon station consists of several layers of aluminium drift tubes (DT) in the barrel region and cathode strip chambers (CSC) in the endcap region, complemented by resistive plate chambers (RPC).

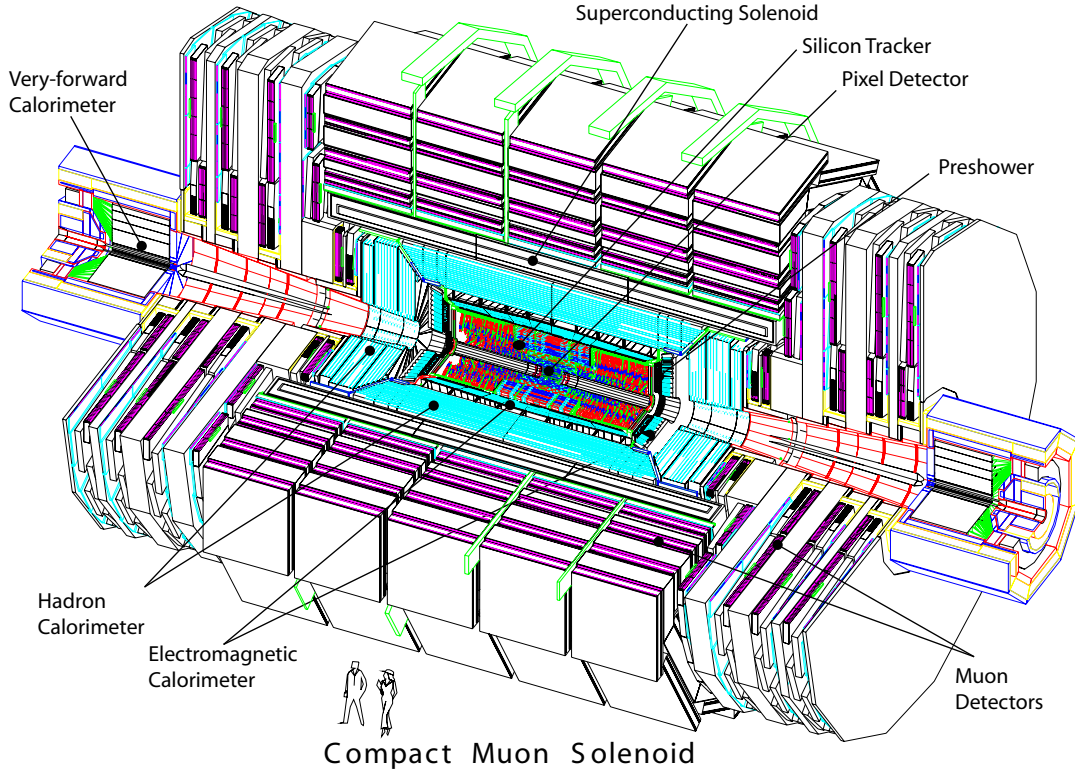


Figure 1: A perspective view of the CMS detector.

The bore of the magnet coil is large enough to accommodate the inner tracker and the calorimetry inside. The tracking volume is given by a cylinder of 5.8-m length and 2.6-m diameter. In order to deal with high track multiplicities, CMS employs 10 layers of silicon microstrip detectors, which provide the required granularity and precision. In addition, 3 layers of silicon pixel detectors are placed close to the interaction region to improve the measurement of the impact parameter of charged-particle tracks, as well as the position of secondary vertices. The expected muon momentum resolution using only the muon system, using only the inner tracker, and using both sub-detectors is shown in Fig. 2.

The electromagnetic calorimeter (ECAL) uses lead tungstate (PbWO_4) crystals with coverage in pseudorapidity up to $|\eta| < 3.0$. The scintillation light is detected by silicon avalanche photodiodes (APDs) in the barrel region and vacuum phototriodes (VPTs) in the endcap region. A preshower system is installed in front of the endcap ECAL for π^0 rejection. The energy resolution of the ECAL, for incident electrons as measured in a beam test, is shown in Fig. 3; the stochastic (S), noise (N), and constant (C) terms given in the figure are determined by fitting the measured points to the function

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2. \quad (1)$$

The ECAL is surrounded by a brass/scintillator sampling hadron calorimeter (HCAL) with coverage up to $|\eta| < 3.0$. The scintillation light is converted by wavelength-shifting (WLS) fibres embedded in the scintillator tiles and channeled to photodetectors via clear fibres. This light is detected by photodetectors (hybrid photodiodes, or HPDs) that can provide gain and operate in high axial magnetic fields. This central calorimetry is complemented by a *tail-catcher* in the barrel region (HO) ensuring that hadronic showers are sampled with nearly 11 hadronic interaction lengths. Coverage up to a pseudorapidity of 5.0 is provided by an iron/quartz-fibre calorimeter. The Cerenkov light emitted in the quartz fibres is detected by photomultipliers. The forward calorimeters ensure full geometric coverage for the measurement of the transverse energy in the event. An even higher forward coverage is obtained with additional dedicated calorimeters (CASTOR, ZDC, not shown in Fig. 1) and with the TOTEM [2] tracking detectors. The expected jet transverse-energy resolution in various pseudorapidity regions is shown in Fig. 4.

The CMS detector is 21.6-m long and has a diameter of 14.6 m. It has a total weight of 12 500 t. The ECAL thickness, in radiation lengths, is larger than $25 X_0$, while the HCAL thickness, in interaction lengths, varies in the range 7–11 λ_I (10–15 λ_I with the HO included), depending on η .

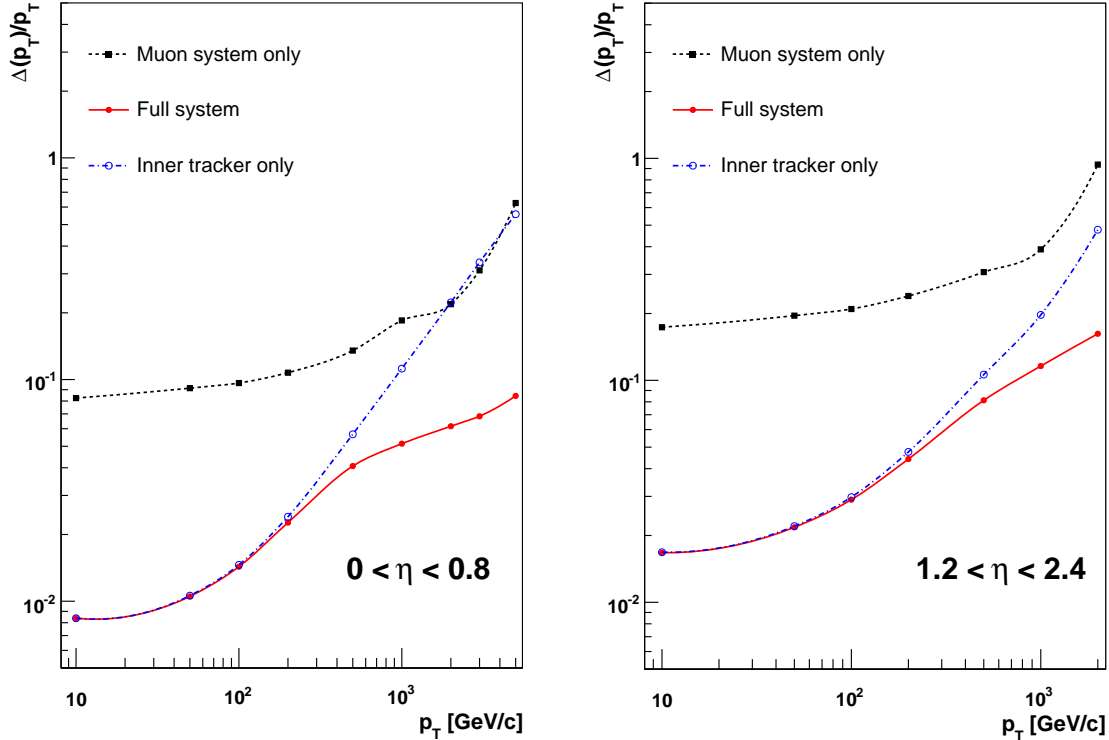


Figure 2: The muon momentum resolution as a function of the momentum (p_T) using the muon system only, the inner tracking only, and both. Left panel: $|\eta| < 0.8$, right panel: $1.2 < |\eta| < 2.4$.

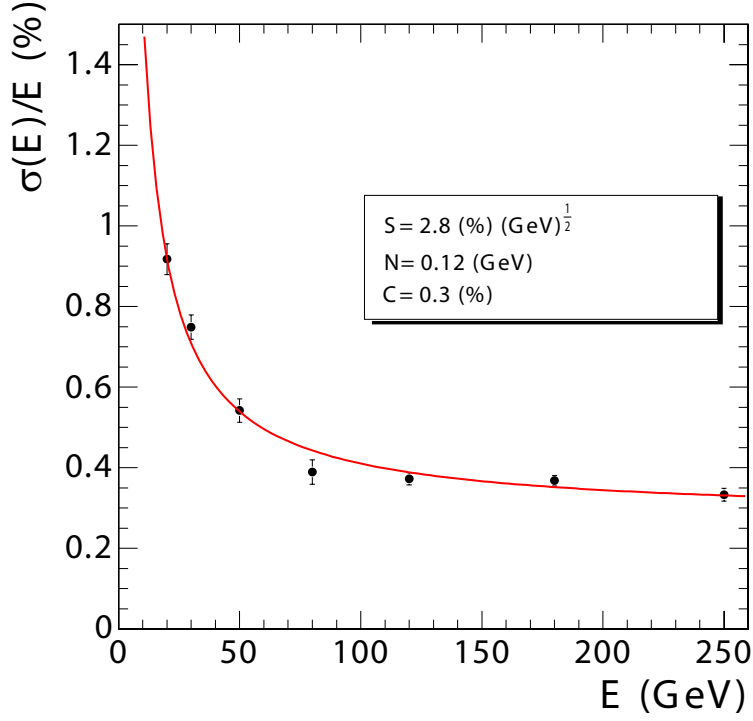


Figure 3: ECAL energy resolution, $\sigma(E)/E$, as a function of electron energy as measured from a beam test. The energy was measured in an array of 3×3 crystals with an electron impacting the central crystal. The points correspond to events taken restricting the incident beam to a narrow ($4 \times 4 \text{ mm}^2$) region. The stochastic (S), noise (N), and constant (C) terms are given.

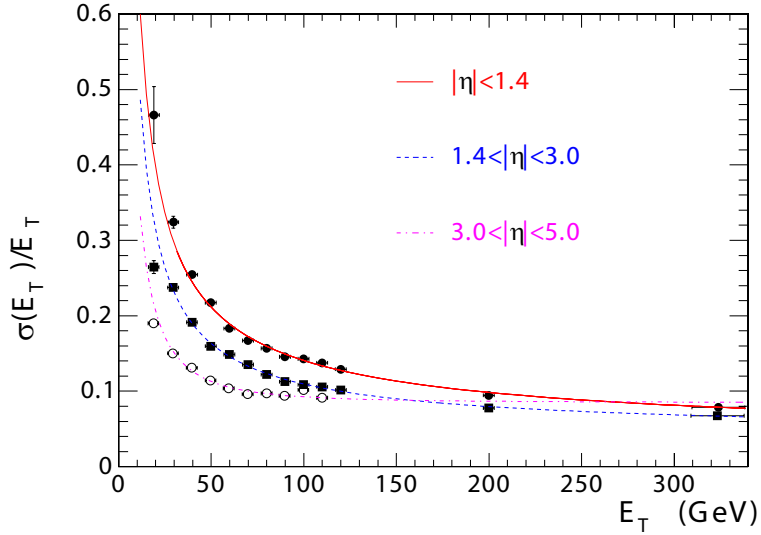


Figure 4: The jet transverse-energy resolution as a function of the simulated jet transverse energy for barrel jets ($|\eta| < 1.4$), endcap jets ($1.4 < |\eta| < 3.0$), and very forward jets ($3.0 < |\eta| < 5.0$). The jets are reconstructed with an iterative cone algorithm (cone radius = 0.5).

3 Superconducting magnet

3.1 Overview

The superconducting magnet for CMS [3, 4, 5, 6] has been designed to reach a 4 T field in a free bore of 6 m diameter and 12.5 m length with a stored energy of 2.6 GJ at full current. The flux is returned through a 10 000 t yoke comprising 5 wheels and 2 endcaps, composed of three disks each (Fig. 1). The distinctive feature of the 220-t cold mass is the 4-layer winding made from a stabilised reinforced NbTi conductor. The ratio between stored energy and cold mass is high (11.6 KJ/kg), causing a large mechanical deformation (0.15%) during energising, well beyond the values of previous solenoidal detector magnets. The parameters of the CMS magnet are summarised in Table 1. The magnet was designed to be assembled and tested in a surface hall (SX5), prior to being lowered 90 m below ground to its final position in the experimental cavern. After provisional connection to its ancillaries, the CMS Magnet has been fully and successfully tested and commissioned in SX5 during autumn 2006.

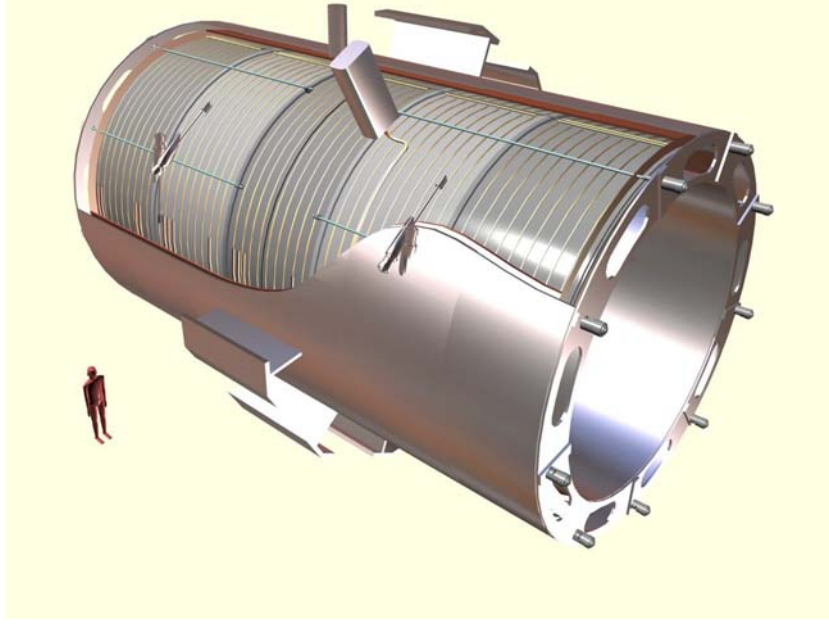


Figure 5: General artistic view of the 5 modules composing the cold mass inside the cryostat, with details of the supporting system (vertical, radial and longitudinal tie rods).

3.2 Main features of the magnet components

3.2.1 Superconducting solenoid

The superconducting solenoid (see an artistic view in Fig. 5 and a picture taken during assembly in the vertical position in SX5 in Fig. 6) presents three new features with respect to previous detector magnets:

- Due to the number of ampere-turns required for generating a field of 4 T (41.7 MA-turn), the winding is composed of 4 layers, instead of the usual 1 (as in the Aleph [7] and Delphi [8] coils) or maximum 2 layers (as in the ZEUS [9] and BaBar [10] coils);
- The conductor, made from a Rutherford-type cable co-extruded with pure aluminium (the so-called insert), is mechanically reinforced with an aluminium alloy;
- The dimensions of the solenoid are very large (6.3-m cold bore, 12.5-m length, 220-t mass).

For physics reasons, the radial extent of the coil (ΔR) had to be kept small, and thus the CMS coil is in effect a “thin coil” ($\Delta R/R \sim 0.1$). The hoop strain (ϵ) is then determined by the magnetic pressure ($P = \frac{B_0^2}{2\mu_0} = 6.4$ MPa), the elastic modulus of the material (mainly aluminium with $Y = 80$ GPa) and the structural thickness ($\Delta R_s = 170$ mm i.e., about half of the total cold mass thickness), according to $\frac{PR}{\Delta R_s} = Y\epsilon$, giving $\epsilon = 1.5 \times 10^{-3}$. This value is high compared to the strain of previous existing detector magnets. This can be better viewed looking at a more

Table 1: Main parameters of the CMS magnet.

General parameters	
Magnetic length	12.5 m
Free bore diameter	6.3 m
Central magnetic induction	4 T
Total Ampere-turns	41.7 MA-turns
Nominal current	19.14 kA
Inductance	14.2 H
Stored energy	2.6 GJ
Cold mass	
Layout	Five modules mechanically and electrically coupled
Radial thickness of cold mass	312 mm
Radiation thickness of cold mass	$3.9 X_0$
Weight of cold mass	220 t
Maximum induction on conductor	4.6 T
Temperature margin wrt operating temperature	1.8 K
Stored energy/unit cold mass	11.6 kJ/kg
Iron yoke	
Outer diameter of the iron flats	14 m
Length of barrel	13 m
Thickness of the iron layers in barrel	300, 630 and 630 mm
Mass of iron in barrel	6000 t
Thickness of iron disks in endcaps	250, 600 and 600 mm
Mass of iron in each endcap	2000 t
Total mass of iron in return yoke	10 000 t

significant figure of merit, i.e. the E/M ratio directly proportional to the mechanical hoop strain according to $\frac{E}{M} = \frac{PR}{2\Delta R_s \delta} \frac{\Delta R_s}{\Delta R} = \frac{\Delta R_s}{\Delta R} \frac{Y\epsilon}{2\delta}$, where δ is the mass density. Figure 7 shows the values of E/M as function of stored energy for several detector magnets. The CMS coil is distinguishably far from other detector magnets when combining stored energy and E/M ratio (i.e. mechanical deformation). In order to provide the necessary hoop strength, a large fraction of the CMS coil must have a structural function. To limit the shear stress level inside the winding and prevent cracking the insulation, especially at the border defined by the winding and the external mandrel, the structural material cannot be too far from the current-carrying elements (the turns). On the basis of these considerations, the innovative design of the CMS magnet uses a self-supporting conductor, by including in it the structural material. The magnetic hoop stress (130 MPa) is shared between the layers (70%) and the support cylindrical mandrel (30%) rather than being taken by the outer mandrel only, as was the case in the previous generation of thin detector solenoids. A cross section of the cold mass is shown in Fig. 8.

The construction of a winding using a reinforced conductor required technological developments for both the conductor [11] and the winding. In particular, for the winding many problems had to be faced mainly related to the mandrel construction [12], the winding method [13], and the module-to-module mechanical coupling. The modular concept of the cold mass had to face the problem of the module-to-module mechanical connection. These interfaces (Fig. 9) are critical because they have to transmit the large magnetic axial force corresponding to 14 700 t, without allowing local displacements due to possible gaps. These displacements can be partially converted into heat, causing a premature quench. A construction method which involved the machining of the upper surface of the modules and a local resin impregnation during the mechanical mounting allowed us to get an excellent mechanical coupling between the modules.

3.2.2 Yoke

The yoke (Fig. 10) is composed of 11 large elements, 6 endcap disks, and 5 barrel wheels, whose weight goes from 400 t for the lightest up to 1920 t for the central wheel, which includes the coil and its cryostat. The easy relative movement of these elements facilitates the assembly of the sub-detectors. To displace each element a combination of heavy-duty air pads plus grease pads has been chosen. This choice makes the system insensitive to metallic dust on the floor and allows transverse displacements. Two kinds of heavy-duty high-pressure air pads with a capacity of either 250 t (40 bars) or 385 t (60 bars) are used. This is not favourable for the final approach when closing



Figure 6: The cold mass mounted vertically before integration with thermal shields and insertion in the vacuum chamber.

the detector, especially for the YE1 endcap that is protruding into the vacuum tank. A special solution has been adopted: for the last 100 mm of approach, flat grease-pads (working pressure 100 bar) have been developed in order to facilitate the final closing of the detector. Once they touch the axially-installed z-stops, each element is pre-stressed with 100 t to the adjacent element. This assures good contact before switching on the magnet. In the cavern the elements will be moved on the 1.23% inclined floor by a strand jacking hydraulic system that ensures safe operation for uphill pulling as well as for downhill pushing by keeping a retaining force. The maximum movements possible in the cavern are of the order of 11 meters; this will take one hour.

To easily align the yoke elements, a precise reference system of about 70 points was installed in the surface assembly hall. The origin of the reference system is the geometrical center of the coil. The points were made after loading the coil cryostat with the inner detectors, the hadronic barrel in particular which weights 1000 t. A mark on the floor was made showing the position of each foot in order to pre-position each element within a ± 5 mm tolerance. Finally, all the elements were aligned with an accuracy of 2 mm with respect to the ideal axis of the coil.

3.2.3 Electrical scheme

The CMS solenoid can be represented as a 14 H inductance mutually coupled with its external mandrel. This inductive coupling allows for the so-called *quench back* effect, as the eddy currents, induced in the external mandrel at the trigger of a current fast discharge, heat up the whole coil above the superconducting critical temperature. This is the fundamental basis of the protection system, which, in case of a superconducting to resistive transition of the coil, aims at keeping the lowest possible thermal gradients and temperature increase in the superconducting windings, and prevents the occurrence of local overheating, hence reducing the thermal stresses inside the winding. A diagram of the powering circuit with protection is shown in Fig. 11.

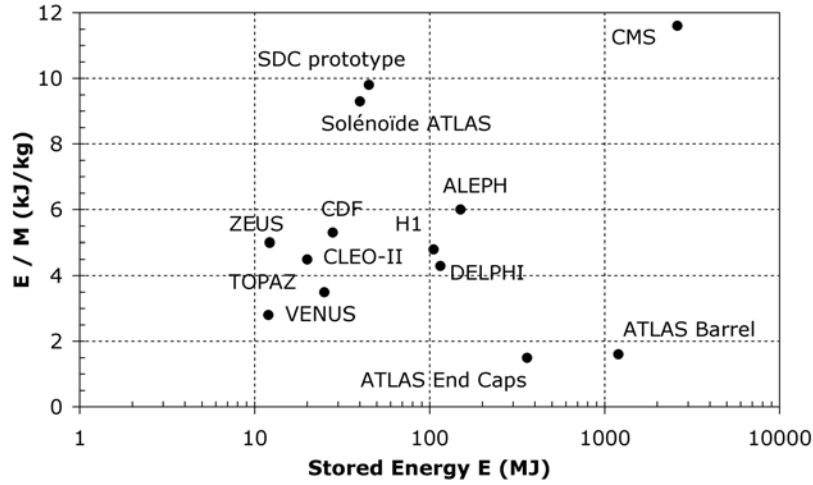


Figure 7: The energy-over-mass ratio E/M , for several detector magnets.

A bipolar thyristor power converter rated at 520 kW with passive L-C filters is used to power the CMS solenoid. It covers a range of voltages from +26 V to -23 V, with a nominal DC current of 19.1 kA. In case of a sudden switch off of the power converter, the current decays naturally in the bus-bar resistance and through the free-wheel thyristors until the opening of the main breakers. Inside the power converter, an assembly of free-wheel thyristors, mounted on naturally air-cooled heat sinks, is installed. In case of non-opening of the main switch breakers, the thyristors are rated to support 20 kA DC for 4 minutes. The current discharge is achieved by disconnecting the electrical power source by the use of two redundant 20 kA DC normally-open switch breakers, leaving the solenoid in series with a resistor, in a L-R circuit configuration. The stored magnetic energy is therefore extracted by thermal dissipation in the so-called dump resistor. This resistor is external to the solenoid cryostat and is designed to work without any active device. It is positioned outdoors taking advantage of natural air convection cooling. The fast discharge (FD) is automatically triggered by hardwired electronics only in case of a superconductive-to-resistive transition, a so-called quench, and for unrecoverable faults which require fast current dumping. The FD time constant is about 200 s. An emergency FD button is also available to the operator in case of need. As the coil becomes resistive during the FD, energy is dissipated inside the coil, which heats up. As a consequence, this necessitates a post-FD cool-down of the coil. The FD is performed on a 30 m Ω dump resistor, as a compromise to keep the dump voltage lower than 600 V, and to limit the coil warm-up and subsequent cool-down time. For faults involving the 20 kA power source, a slow discharge (SD) is triggered through hardwired electronics on a 2 m Ω dump resistor. The SD current evolution is typically exponential, and its time constant is 7025 s, but the coil stays in the superconducting state as the heat load, about 525 W, is fully absorbed by the cooling refrigerator. For current lower than 4 kA, a FD is performed in any case, as the heat load is small enough for the refrigerator. The same resistor is used in both cases for the FD and the SD, using normally open contactors, leaving the dump resistor modules either in series (FD) or in parallel (SD). For other cases, and depending on the alarms, the coil current can be adjusted by the operator, or ramped down to zero, taking advantage of the two-quadrant converter.

3.2.4 Vacuum system

The vacuum system has been designed to provide a good insulation inside the 40 m³ vacuum volume of the coil cryostat. It consists of 2 double-primary pumping stations, equipped with 2 rotary pumps and 2 Root's pumps, that provide the fore vacuum to the two oil diffusion pumps located at the top of CMS and connected to the coil cryostat via the current leads chimney and the helium phase separator. The rotary pumps have a capacity of 280 m³/h while the two Root's pumps have a flow of 1000 m³/h. The biggest oil diffusion pump, installed via a DN 400 flange on the current leads chimney, has a nominal flow of 8000 l/s at 10⁻⁴ mbar of fore vacuum. The smallest one delivers 3000 l/s at the phase separator.

3.2.5 Cryogenic plant

The helium refrigeration plant for CMS is specified for a cooling capacity of 800 W at 4.45 K, 4500 W between 60 and 80 K, and simultaneously 4 g/s liquefaction capacity. The primary compressors of the plant have been installed, in their final position, while the cold box, as well as the intermediate cryostat which interfaces the phase separator

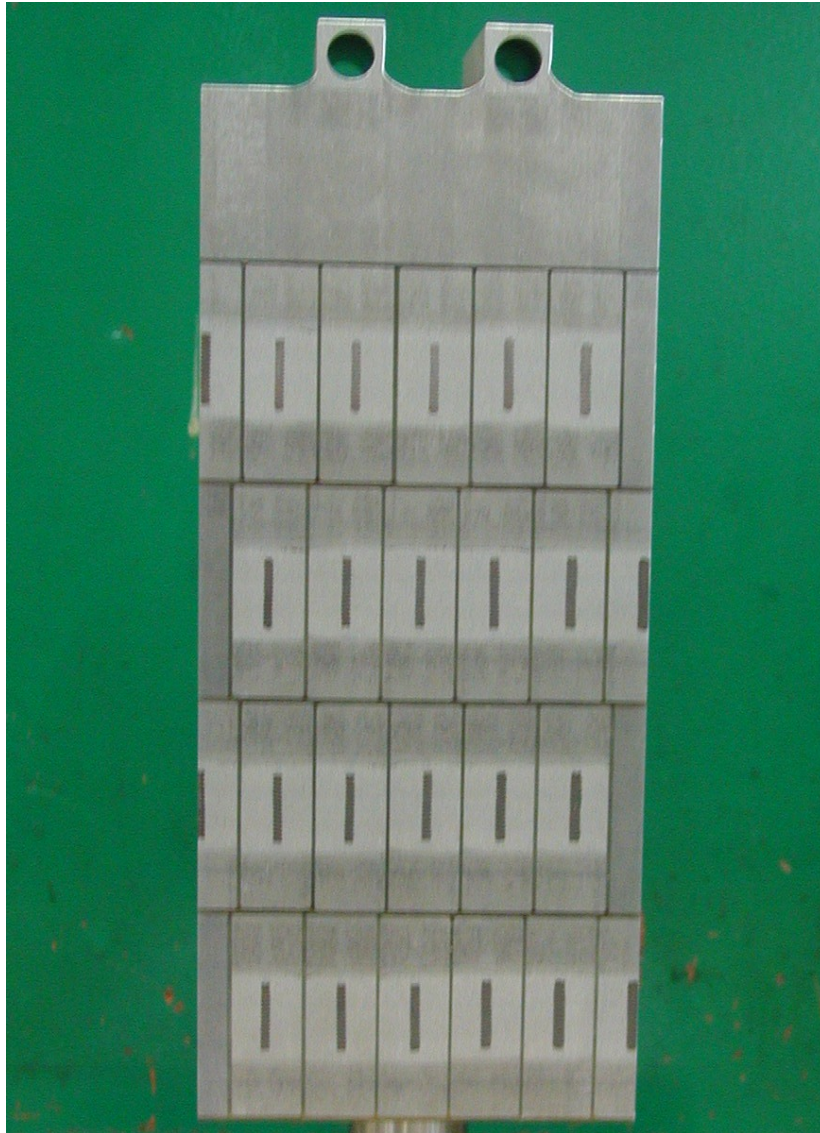


Figure 8: Cross section of the cold mass with the details of the 4-layer winding with reinforced conductor.

and the thermo-syphon, were moved underground after the completion of the magnet test. These components were commissioned with the help of a temporary heat load of 6.5 kW that simulated the coil cryostat which was not yet available. The performance of the cold box has been measured in cool-down mode and in nominal and operation mode.

3.2.6 Other ancillaries

- *Current leads.* The two 20-kA current leads are made of a high purity copper braid, having a cross section of 1800 m^2 and RRR (Residual Resistivity Ratio) of 130, placed inside a conduit and cooled by circulating helium gas. Without cooling, the current leads are able to hold a current of 20 kA for 5 minutes, followed by a FD without any damage, as the temperature at the hot spot stays below 400 K [14].
- *Grounding circuit.* The grounding circuit is connected across the solenoid terminals. It fixes the coil circuit potential, through a $1 \text{ k}\Omega$ resistor, dividing by two the potential to ground. The winding insulation quality is monitored by continuously measuring the leakage current through a $10 \text{ }\Omega$ grounding resistor.
- *Quench detection system.* The quench detection system is a key element of the Magnet Safety System (MSS). The role of the quench detection system is to detect a resistive voltage between two points of the coil, whose value and duration are compared to adjustable thresholds. The voltage taps are protected by $4.7 \text{ k}\Omega$, 6 W resistors. There are 2 redundant systems, with resistor bridge detectors and differential detectors. For each

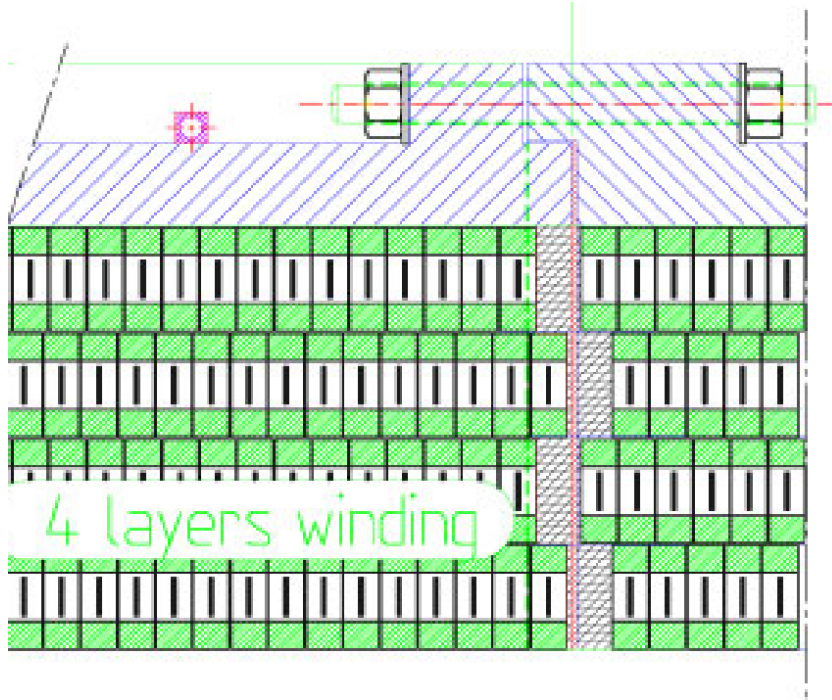


Figure 9: Detail of the interface region between 2 modules. In order to guarantee mechanical continuity, false turns are involved. The modules are connected through bolts and pins fixed through the outer mandrels.

system, there are 5 detectors. Each resistor bridge detector spans two modules and one detector spans the whole solenoid. Each coil module is compared with two other modules through two differential detectors.

3.3 Operating test

The magnet and all its ancillaries were assembled for testing in SX5 and ready for cool-down in January 2006. Figure 12 shows the test layout.

3.3.1 Cool-down

The cool-down of the solenoid started on February, the 2nd, 2006 and in a smooth way brought the cold mass to 4.6 K in 24 days. Figure 13 shows the cool-down curve. The only glitch was due to an overpressure on a safety release valve that stopped cooling for one night before the system was restarted.

One important aspect monitored during the cool-down was the amount of coil shrinkage. In order to explain this point, we refer to the coil suspension system inside the cryostat (Fig. 5), made of longitudinal, vertical, and axial tie-rods in Ti alloy. The magnet is supported by 2×9 longitudinal tie rods, 4 vertical tie rods, and 8 radial tie rods. The tie rods are equipped with compensated strain gauges to measure the forces on 2×3 longitudinal, plus the vertical and radial tie rods. The tie rods are loaded in tension and flexion. To measure the tension and flexion strain, 3 strain gauges are placed on the tie rods at 0° , 90° , and 180° .

The measured stresses in the tie bars due to the cool-down, causing a shrinkage of the cold mass and putting the tie-bars in tension, are shown in Table 2. A comparison with the expected values is provided as well. The measured axial and radial shrinkage of the cold mass is shown in Fig. 14.

3.3.2 Charge and discharge cycles

The magnetic tests took place during August 2006, with additional tests during the magnet field mapping campaign in October 2006. The current ramps for the field mapping are detailed in Fig. 15. The tests were carried out through magnet charges to progressively higher currents, setting increasing dI/dt , followed by slow or fast discharges. During these current cycles all the relevant parameters related to electrical, magnetic, thermal, and mechanical behaviours have been recorded. Depending on the level of the current at the trigger of a fast discharge, the time needed for re-cooling the coil can be up to 3 days.



Figure 10: A view of the yoke at an early stage of magnet assembly at SX5. The central barrel supports the vacuum chamber of the superconducting coil. At the rear, one of the closing end cap disks is visible.

Table 2: Calculated and measured cold mass displacements and related stresses on tie-rods due to the cool-down to 4.5 K.

	Expected value	Measured value
Cold Mass Shrinkage		
Longitudinal	26 mm	27 mm
Radial	14 mm	15 mm
Tie rod stress due to cool-down		
Vertical	315 MPa	310±45 MPa
Radial	167 MPa	153±20 MPa
Longitudinal	277 MPa	260±20 MPa

3.3.3 Cold mass misalignment

The support system is designed to withstand the forces created by a 10 mm magnetic misalignment, in any direction of the cold mass with respect to the iron yoke. Geometrical surveys were performed at each step of the magnet assembly to ensure a good positioning. Nevertheless, the monitoring of the coil magnetic misalignment is of prime importance during magnet power test. The misalignment can be calculated either by analysing the displacement of the cold mass or the stresses of the tie rods when the coil is energised. The displacement is measured at several locations and directions at both ends of the coil with respect to the external vacuum tank wall, by the use of rectilinear potentiometers. Results are displayed in Figs. 16 and 17. The displacement of the coil's geometric centre is found to be 0.4 mm in z , in the $+z$ direction. According to the computations, such a displacement indicates that the coil centre should be less than 2 mm off the magnetic centre in $+z$. As the coil supporting system is hyper-static, the tie rods are not all initially identically loaded. But the force increase during energising is well distributed, as shown in Fig. 18 and Fig. 19, giving the force measurements on several tie rods. These figures also indicate the forces computed in the case of a 10-mm magnetic misalignment, together with forces calculated for the ideally-centred model, showing there is no noticeable effect of misalignment on the forces.

Using the strain gauges glued on the cold mass (outer mandrel of the central module, CB0), one can determine the Von Mises stress. The cold mass Von Mises stress versus the coil current is given in Fig. 20. The measured value of Von Mises stress at 4.5 K and zero current is 23 MPa. The value at 19.1 kA is 138 MPa. These values are in agreement with computations done during design [3, 6].

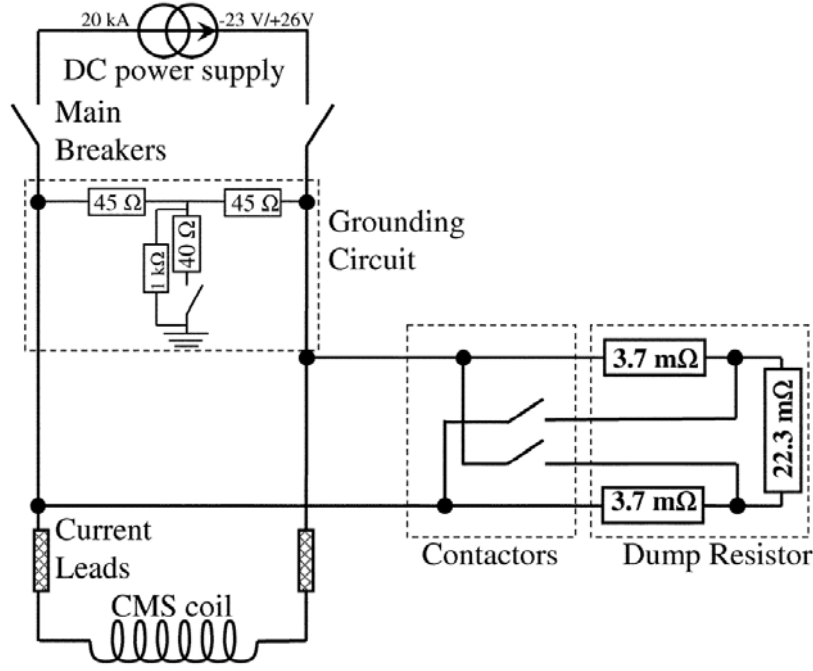


Figure 11: The electrical scheme of the magnet with the protection circuit. One of the main components of the protection is the dump resistor, made of three elements.

3.3.4 Electrical measurements

The apparent coil inductance measured through the inductive voltage $V = LdI/dt$ is decreasing while increasing the current, as the iron yoke reaches the saturation region. From voltage measurements at the coil ends in the cryostat, while ramping up the coil current at a regulated dI/dt , the inductance is calculated and results are given in Fig. 21. Initially the apparent inductance of the coil is 14.7 H at zero current, and then it decreases to 13.3 H at 18 kA. The 21 resistive electrical joints, which connect the 5 modules together and, for each module, the 4 layers, are positioned externally to the coil, on the outer radius of the external mandrel, in low magnetic field regions. The resistance measurements of the joints indicate values ranging from 0.7 nΩ to 1.6 nΩ at 19.1 kA, corresponding to a maximum dissipation in the joint of 0.6 W. The specific joint cooling system is fully efficient to remove this local heat deposit in order to avoid that the resistive joints generate a local quench of the conductor. As mentioned above, the fast discharge causes a quench of the coil, through the quench-back process. The typical current decay at the nominal current of 19.14 kA is given in Fig. 22.

The effect of the mutual coupling of the coil with the external mandrel is clearly visible at the beginning of the current fast discharge as shown in the zoomed detail of Fig. 22. It appears clearly that a high dI/dt of about 500 A/s occurs at the very beginning of the discharge. The minimum and maximum temperatures of the coil are displayed in Fig. 23 for a fast discharge at 19.14 kA. A maximum temperature difference of 32 K is measured on the coil between the warmest part, located on the coil central module internal radius, and the coldest part, located on the external radius of the mandrel. It should be noted that the thermal gradient is mainly radial. The temperature tends to equilibrate over the whole coil 2 hours after the trigger of the fast discharge. The average cold mass temperature after a fast discharge at 19 kA is 70 K.

During a magnet discharge, the dump resistor warms up, with a maximum measured temperature increase of 240 °C, resulting in an increase of the total dump resistance value by up to 19%. Also the coil internal electrical resistance is increased by up to 0.1 Ω at the end of a FD at 19.14 kA. The effect of both the dump resistor and the magnet electrical resistance increasing was revealed through the measurement of the discharge time constant, which was equal to 177 s, 203 s, 263 s, 348 s and 498 s for fast discharges respectively at 19 kA, 17.5 kA, 15 kA, 12.5 kA and 7.5 kA. This is visible in Fig. 24. The temperature recovery of the dump resistor is achieved in less than 2 hours after the trigger of a fast dump. It is 5 hours after the trigger of a slow dump.

In the case of a fast dump at 19.14 kA, typically half of the total energy (1250 MJ) is dissipated as heat in the external dump resistor. The energy dissipated in the dump resistor as a function of the magnet current at the trigger of a FD was measured for each FD performed during the magnet tests and is given in Fig. 25. The magnet current

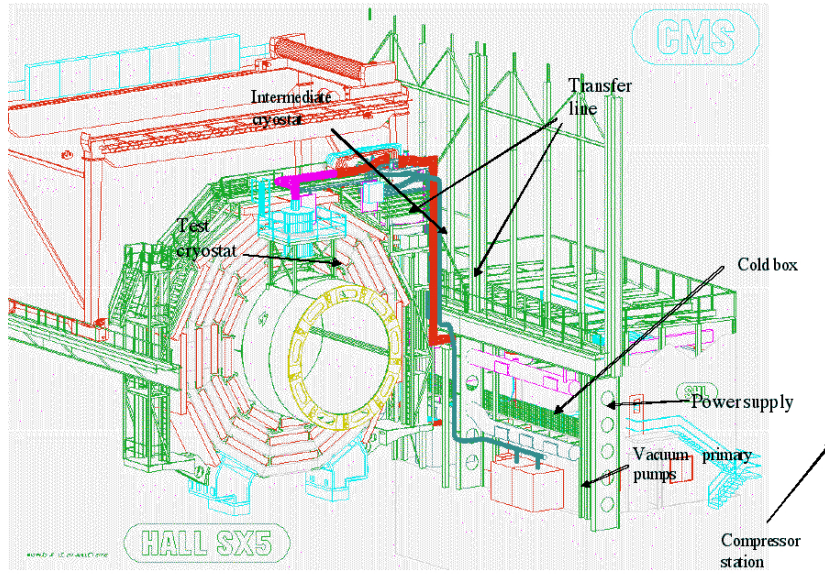


Figure 12: The layout for the surface test at SX5, showing only the central barrel. The magnet is connected to the cryoplant (through the proximity cryogenics), the vacuum and the power systems.

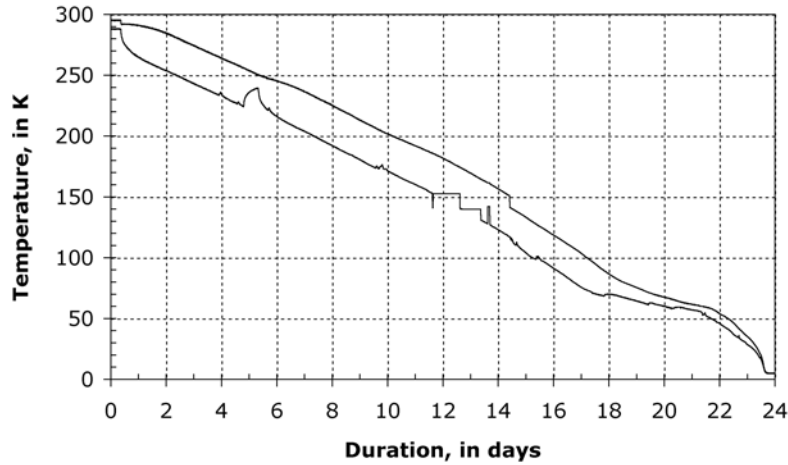


Figure 13: Graph of the coil minimum and maximum temperatures during the cool-down from room temperature to 4.5 K.

is precisely measured by the use of two redundant DCCTs (DC current transformer). The peak-to-peak stability of the current is 7 ppm with a voltage ripple of 2.5% (0.65 V). In order to gain on the operation time, an acceleration of the slow dump has been tested and validated by switching to the fast dump configuration at 4 kA. It has been checked that the cryogenic refrigerator can take the full heat load, and the magnet stays in the superconducting state. This Slow Dump Accelerated (SDA) mode was tested in semi-automatic mode through the cryogenics supervisory system and the magnet control system, and it will be fully automatic for the final installation in the cavern.

3.3.5 Yoke mechanical measurements

The elements of the return yoke, barrels and endcaps, are attached with several hydraulic locking jacks, which are fixed on each barrel and endcap. They are pre-stressed in order to bring the barrels and endcaps into contact at specific areas using the aluminium-alloy Z-stop blocks. There are 24 Z-stops between each barrel and endcap. A computation of the total axial compressive force gives 8900 tons. The stresses are measured on some Z-stops; the forces on these Z-stops are given in Fig. 26 and compared to the case of a uniformly distributed load on all

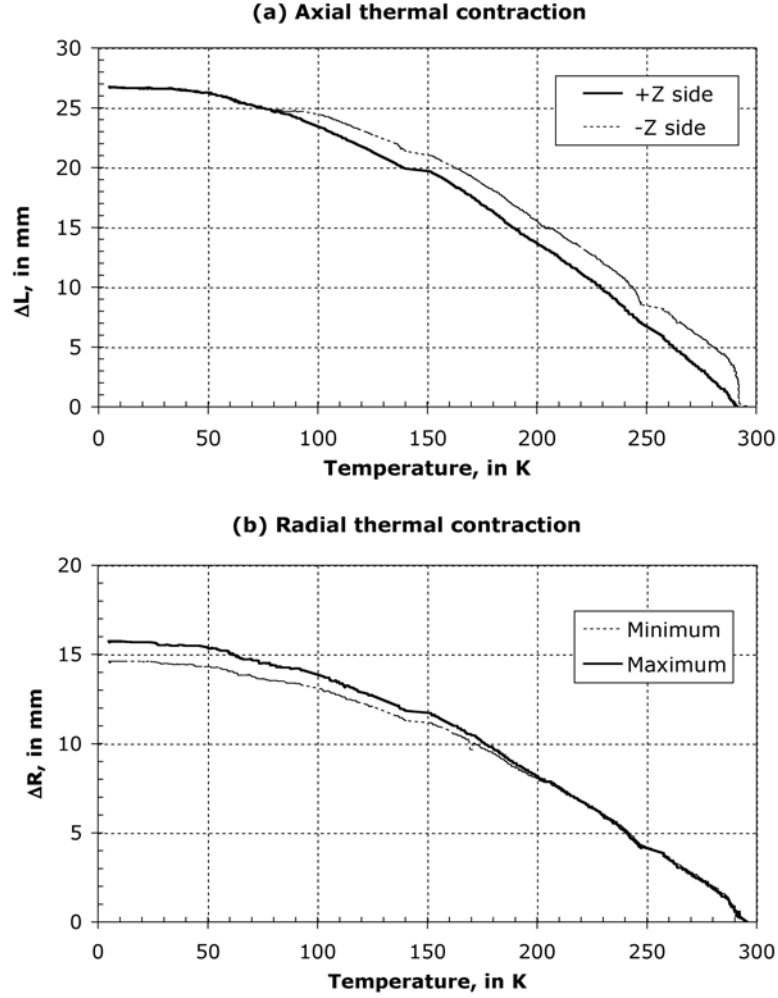


Figure 14: Axial (a) and radial (b) shrinkage of the cold mass from 300 K to 4.5 K.

the Z-stops. To allow for uniform load distribution and distortion during magnet energising, the yoke elements are positioned on grease pads. During magnet energising, the displacement of the barrel yoke elements under the compressive axial force is very limited, while the displacement of the yoke end cap disk YE+1 is clearly noticeable on the outer radius of the disk, due to the axial attraction of the first yoke endcaps towards the interaction point. The measurement of the distance between the barrel elements parallel to the axial axis of the detector is given in Fig. 27. The endcap YE+1 disk is equipped with rosette strain gauges on its inner face, under the muon chambers and near the bolts at the interface between two adjacent segments. The main stresses measured in these regions do not exceed 88 MPa.

3.3.6 Coil stability characteristics

The NbTi superconductor critical temperature is $T_c = 9.25$ K at zero field. At $B = 4.6$ T (peak field on the conductor), $T_c = 7.30$ K. The current-sharing temperature T_g is defined as the maximum temperature for which the current can flow, with no dissipation, in the superconducting part. For CMS the operating current is 19 143 A, while the critical current, according to the measurements done on a short sample extracted from the length used in the inner layer of the central module (the one exposed to the higher field), is $I_c (T = 4.5$ K, $B = 4.6$ T) = 62 kA leading to $T_g = 6.44$ K, i.e., the temperature margin is 1.94 K. This margin is a little higher than the designed one (1.83 K) because the nominal current is less than the one used in this kind of computation (19.5 kA) and the expected conductor critical current was from 7% to 10% lower than the real one obtained through advanced and qualified processes. The T_c value was confirmed at 9.3 K during cryogenic recovery tests (Fig. 28) at zero field. The conductor pure-aluminium stabilizer RRR, deduced from electrical measurements during cool-down, is found to be above 1800.

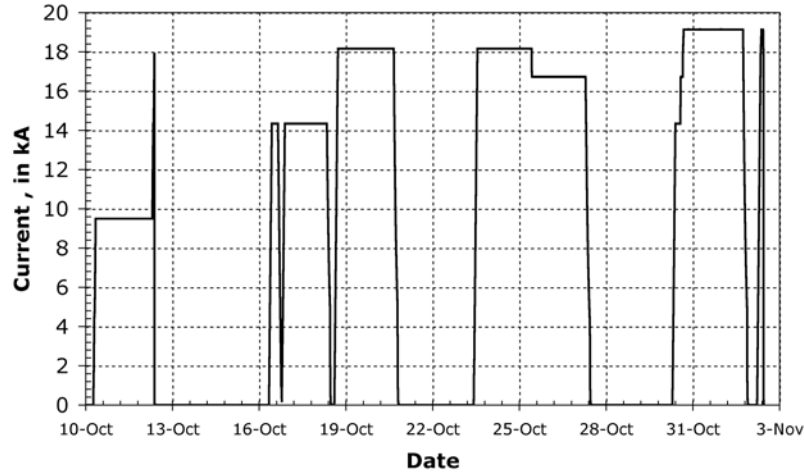


Figure 15: Magnet cycles during the CMS magnet tests in October 2006.

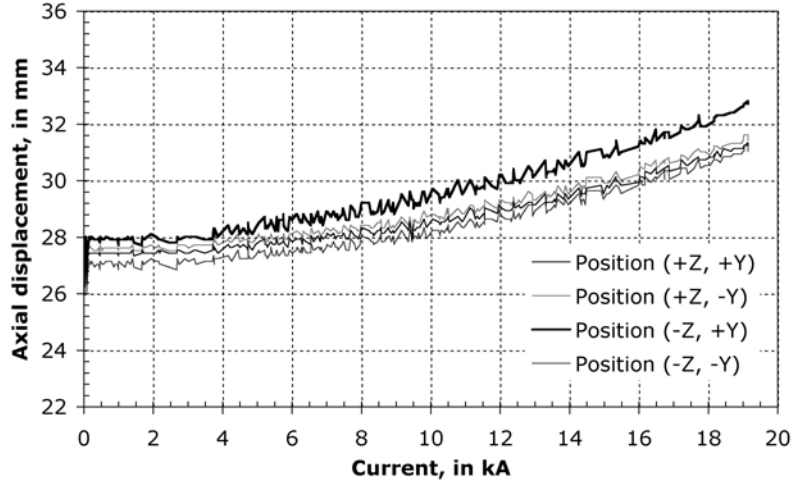


Figure 16: Axial displacement in Z at both ends of the coil in different positions during energising.

3.3.7 Coil warm-up

Following the test of the magnet on the surface, the cold mass had to be warmed up to room temperature before lowering. The coil, inside its cryostat, was attached to the central barrel YB0 to avoid any risk due to vacuum degradation during the transport operations. The warm-up was performed using a dedicated power supply (200 V-300 A DC) to maintain integrity of the coil/mandrel interface. Knowing the temperature dependence of both the electrical resistivity and the specific heat of the coil materials, the temperature increase for a given electrical power is calculated. Taking into account the capacity of the warm-up supply, and limiting the temperature increase to 1 K/hour, the warm-up was performed as shown in Fig. 29. As the warm-up was done after a fast discharge, the coil temperature was already at 70 K. Nevertheless, the warm-up took place only at night as the yoke was opened to continue integration activities inside the detector. Ultimately, the warm-up lasted only 3 weeks. The maximum temperature gradient across the coil during the warm-up exercise was less than 9 K.

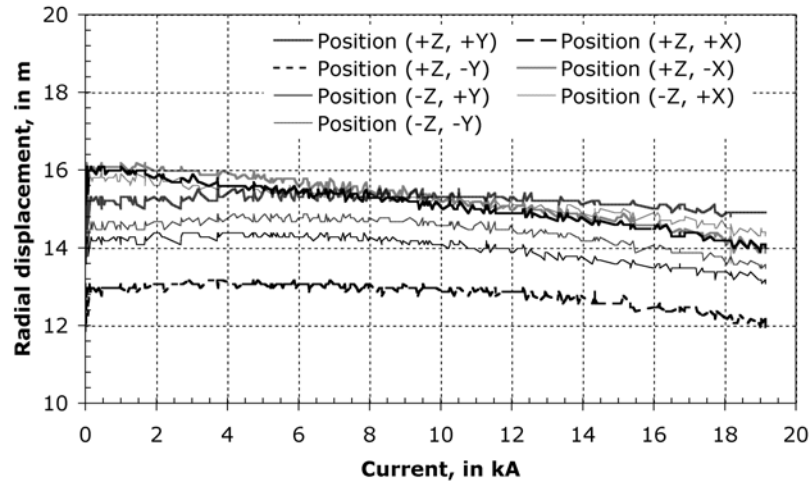


Figure 17: Radial displacement at both ends of the coil in different positions during energising.

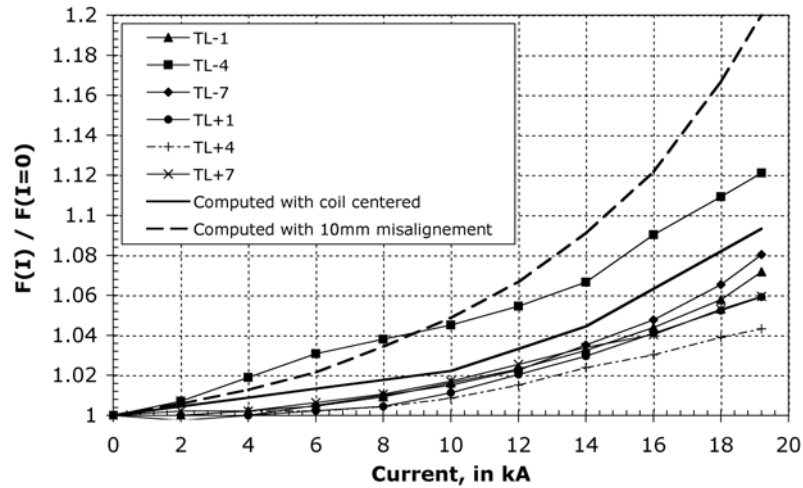


Figure 18: Force increase on several axial tie rods; the average force at zero current is 45 t.

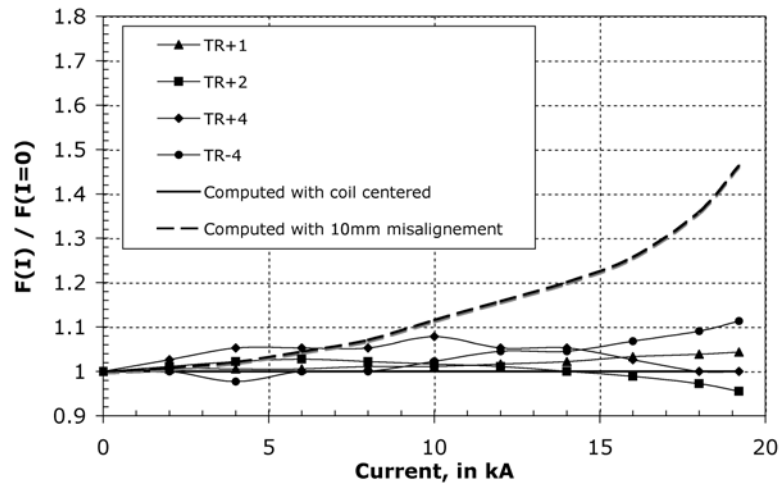


Figure 19: Force increase on several radial tie rods; the average force at zero current is 15 tons.

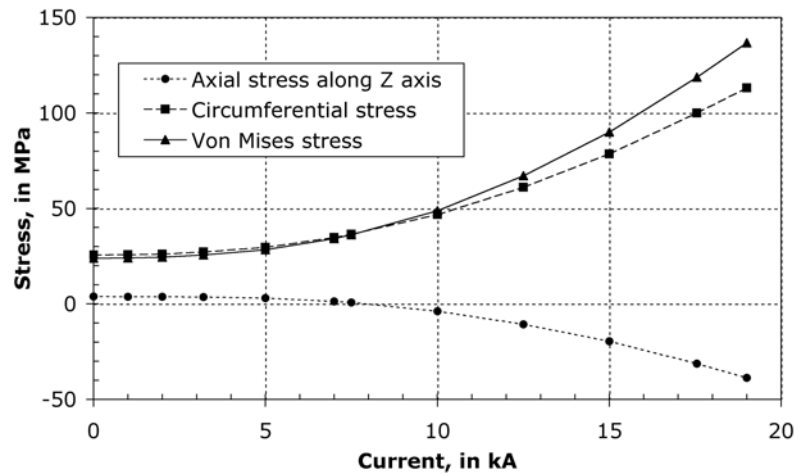


Figure 20: Stresses measured on the CB0 module as a function of the current.

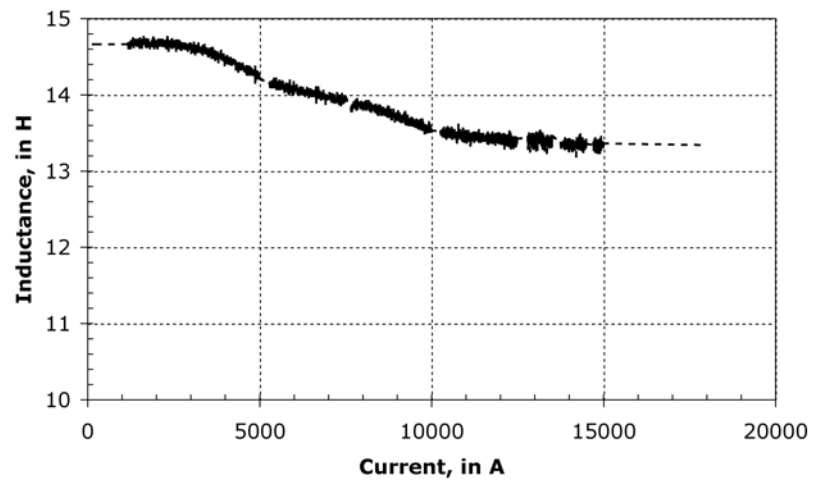


Figure 21: Coil inductance as a function of the magnet current.

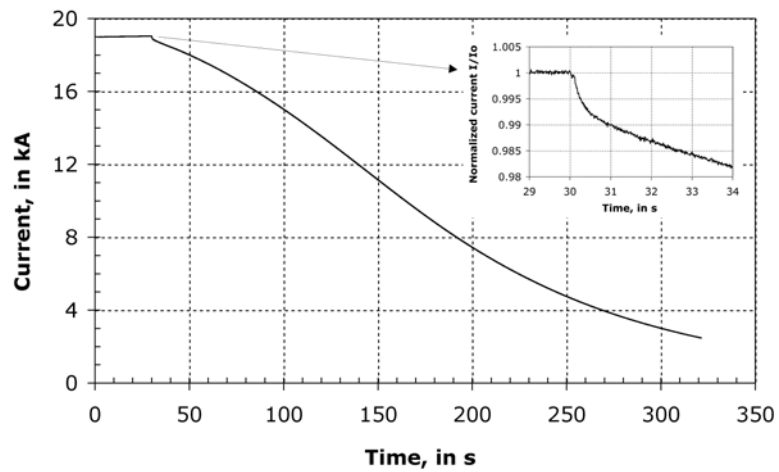


Figure 22: Magnet current during fast discharge at the nominal field of 4 T. The insert shows the details at the beginning of the discharge.

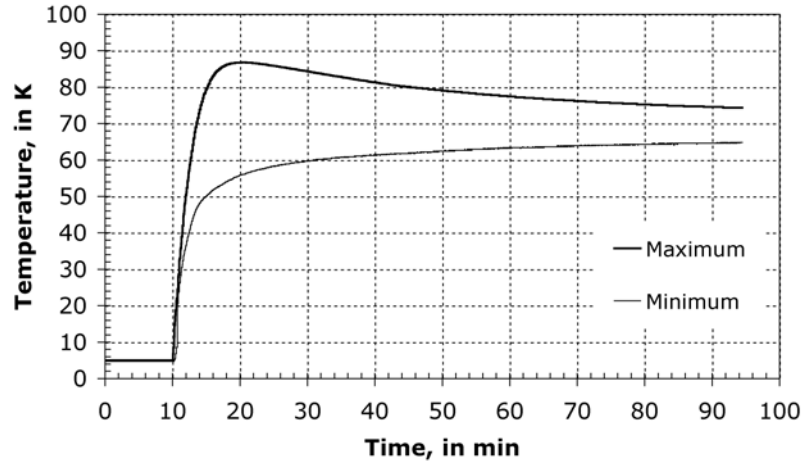


Figure 23: Minimum and maximum temperatures detected on the cold mass during the fast discharge from 19.1 kA.

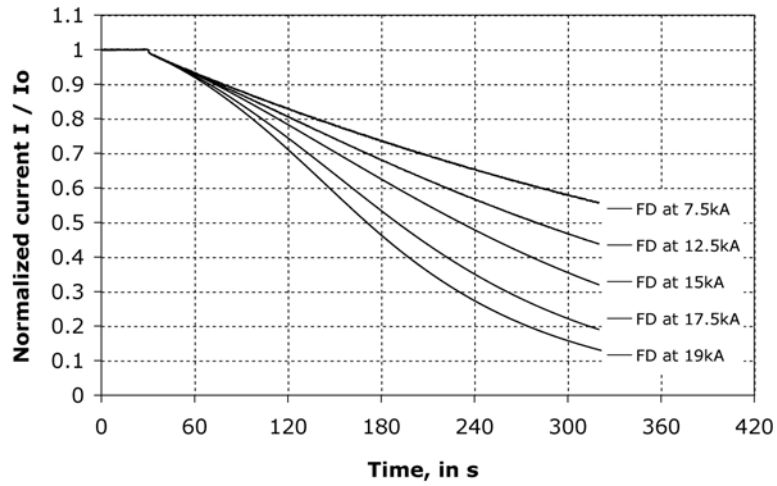


Figure 24: The normalised discharge current as a function of time for different initial currents, showing the effect of the increase in magnet and external dump resistance with current.

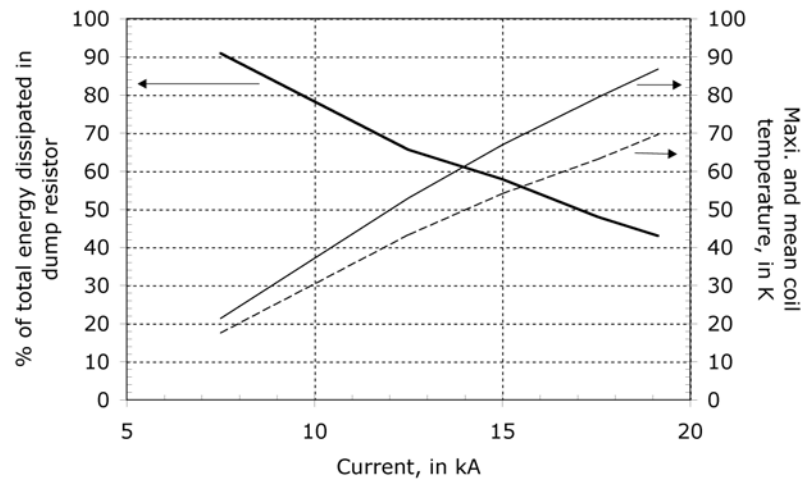


Figure 25: Energy dissipated in the external dump resistor and the mean and maximum temperatures of the coil during FD.

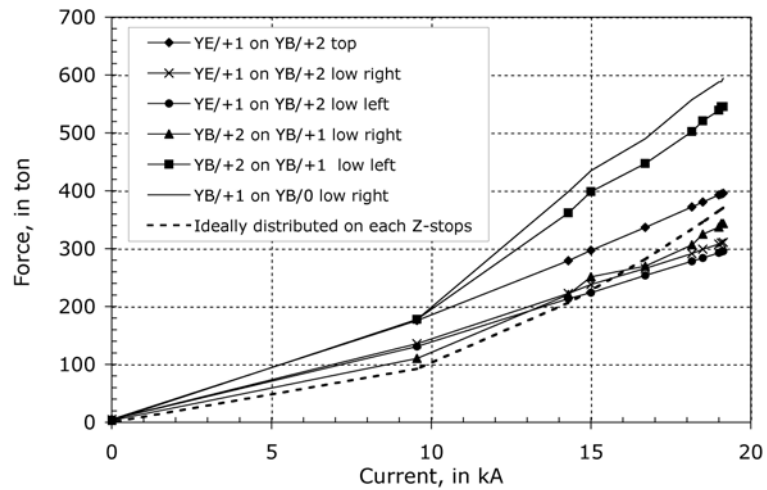


Figure 26: Axial forces acting on the yoke Z-stops during the coil energizing.

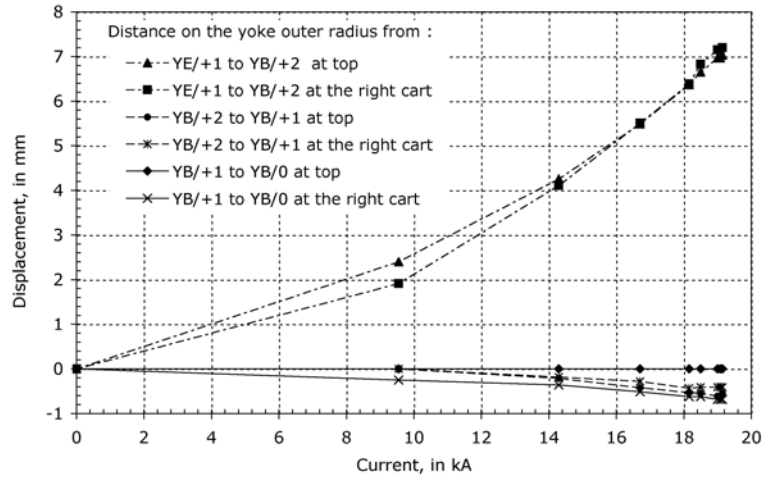


Figure 27: Measured displacement of the yoke during the coil energising.

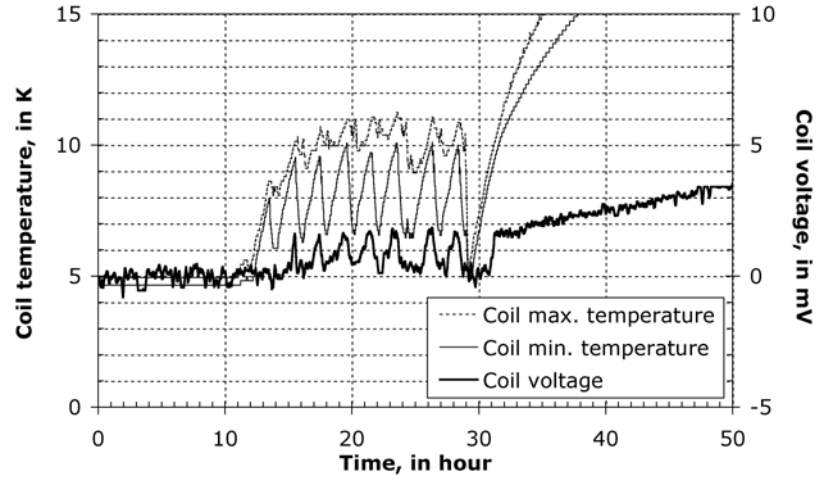


Figure 28: The minimum and maximum temperatures and voltage of the coil as a function of time, with only a few amperes of current, showing the superconducting-to-resistive-state transition at around 9.3 K.

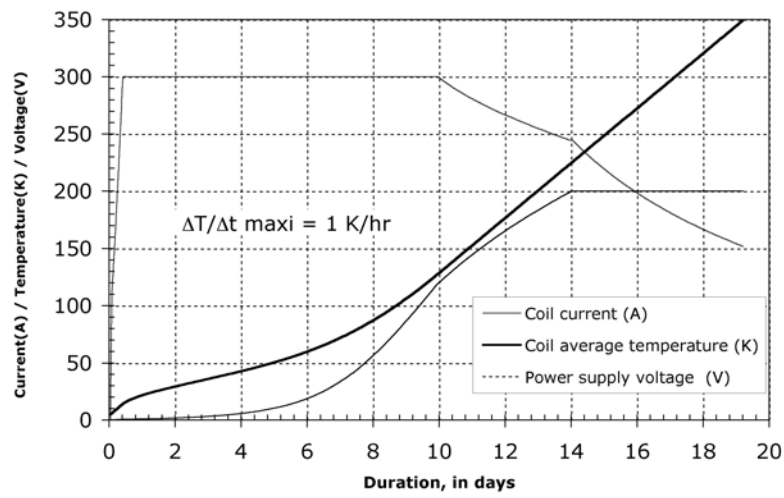


Figure 29: Measurements of the coil warm-up behaviour as a function of time; the Y-axis scale is common for all the three curves.

4 Inner tracking system

4.1 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 $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ 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 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.

4.1.1 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 MHz/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. The deviation from the $1/r^2$ law results from the strong magnetic field of 4 T. 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 $100 \times 150 \mu\text{m}^2$ 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 ($20 \text{ cm} < r < 55 \text{ cm}$) the reduced particle flux allows the use of silicon micro-strip detectors with a typical cell size of $10 \text{ cm} \times 80 \mu\text{m}$, leading to an occupancy of up to 2–3% per strip and LHC bunch crossing. In the outer region ($55 \text{ cm} < r < 110 \text{ cm}$) 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 ($500 \mu\text{m}$ thickness as opposed to the $320 \mu\text{m}$ 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 $25 \text{ cm} \times 180 \mu\text{m}$ 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/cm² 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 shows the expected fast hadron fluence and radiation dose in the CMS barrel tracker for an integrated luminosity of 500 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 $\sqrt{s} = 14$ 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.

Table 3: Expected hadron fluence and radiation dose in different radial layers of the CMS tracker (barrel part) for an integrated luminosity of 500 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

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 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.

4.1.2 Overview of the tracker layout

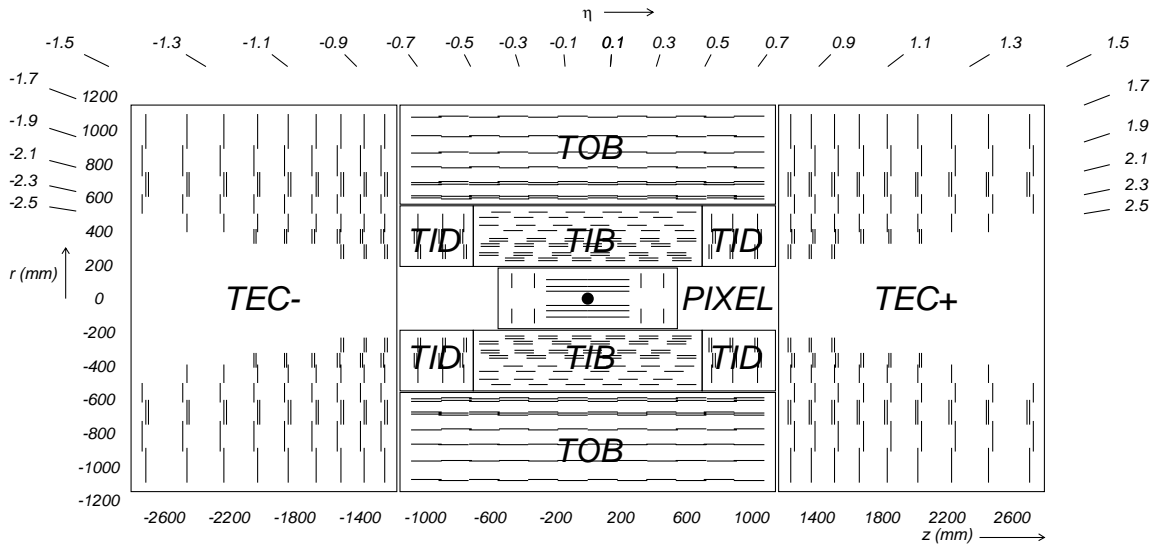


Figure 30: Schematic cross section through the CMS tracker. Each line represents a detector module. Double lines indicate back-to-back modules which deliver stereo hits.

A schematic drawing of the CMS tracker is shown in Fig. 30. 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 Sect. 4.2. In total the pixel detector covers an area of about 1 m² 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 Sect. 4.3. 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 μm 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 μm on layers 1 and 2 and 120 μm on layers 3 and 4 in the TIB, leading to a single point resolution of 23 μm and 35 μm , respectively. In the TID the mean pitch varies between 100 μm and 141 μm . 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 $500\text{ }\mu\text{m}$ thick micro-strip sensors with strip pitches of $183\text{ }\mu\text{m}$ on the first 4 layers and $122\text{ }\mu\text{m}$ on layers 5 and 6. It provides another 6 r - ϕ measurements with single point resolution of $53\text{ }\mu\text{m}$ and $35\text{ }\mu\text{m}$, respectively. The TOB extends in z between $\pm 118\text{ cm}$. Beyond this z range the Tracker EndCaps (TEC+ and TEC- where the sign indicates the location along the z axis) cover the region $124\text{ cm} < |z| < 282\text{ cm}$ and $22.5\text{ cm} < |r| < 113.5\text{ cm}$. Each TEC is composed of 9 disks, carrying up to 7 rings of silicon micro-strip detectors ($320\text{ }\mu\text{m}$ thick on the inner 4 rings, $500\text{ }\mu\text{m}$ thick on rings 5-7) with radial strips of $97\text{ }\mu\text{m}$ to $184\text{ }\mu\text{m}$ 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 $230\text{ }\mu\text{m}$ and $530\text{ }\mu\text{m}$ 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 (Fig. 31). 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 32 shows the material budget of the CMS tracker in units of radiation length. It increases from $0.4\text{ }X_0$ at $\eta \approx 0$ to about $1.8\text{ }X_0$ at $|\eta| \approx 1.4$, beyond which it falls to about $1\text{ }X_0$ at $|\eta| \approx 2.5$.

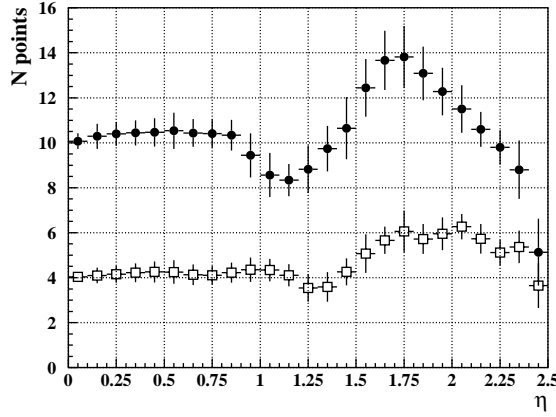


Figure 31: Number of measurement points in the strip tracker as a function of pseudorapidity η . Filled circles show the total number (back-to-back modules count as one) while open squares show the number of stereo layers.

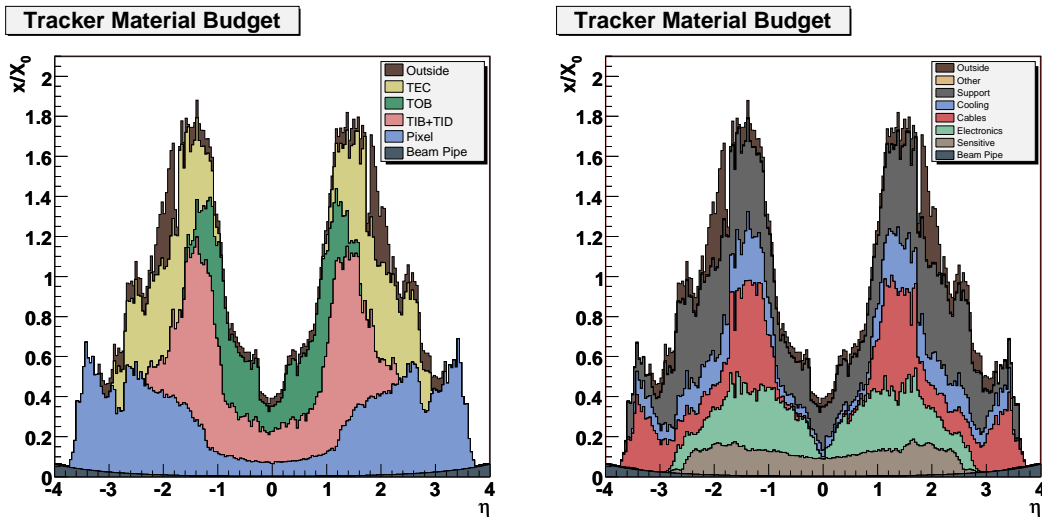


Figure 32: Material budget in units of radiation length as a function of pseudorapidity η for the different sub-detectors (left panel) and broken down into the functional contributions (right panel).

4.1.3 Expected performance of the CMS tracker

For single muons of transverse momenta of 1, 10 and 100 GeV Fig. 33 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 30% of the transverse momentum resolution while at lower momentum it is dominated by multiple scattering. The transverse impact parameter resolution reaches 10 μm 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 34 shows the expected track reconstruction efficiency of

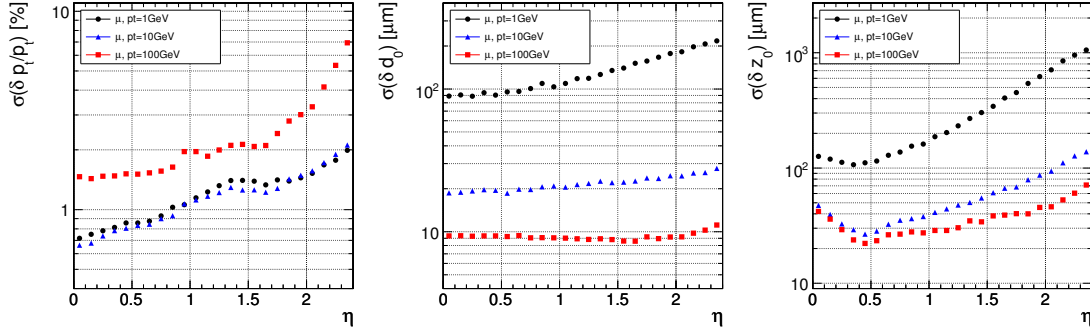


Figure 33: Resolution of several track parameters for single muons with transverse momenta of 1, 10 and 100 GeV: transverse momentum (left panel), transverse impact parameter (middle panel), and longitudinal impact parameter (right panel).

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 $|\eta|$ 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.

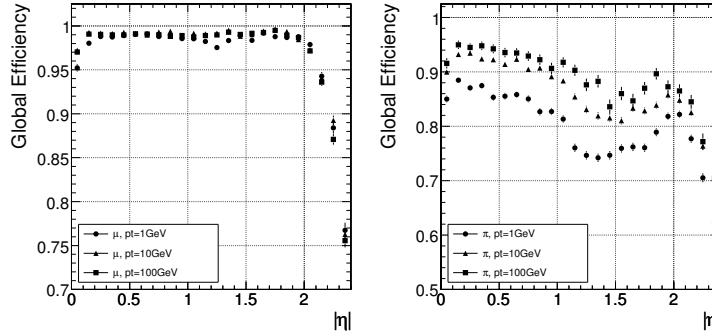


Figure 34: Global track reconstruction efficiency for muons (left panel) and pions (right panel) of transverse momenta of 1, 10 and 100 GeV.

4.1.4 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-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 m³/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 m³) is flushed with pre-chilled dry nitrogen gas at a rate of up to one volume exchange per hour.

4.2 Pixel detector

4.2.1 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 size of 100×150 μm^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 ≈ 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) m². 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 35 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.

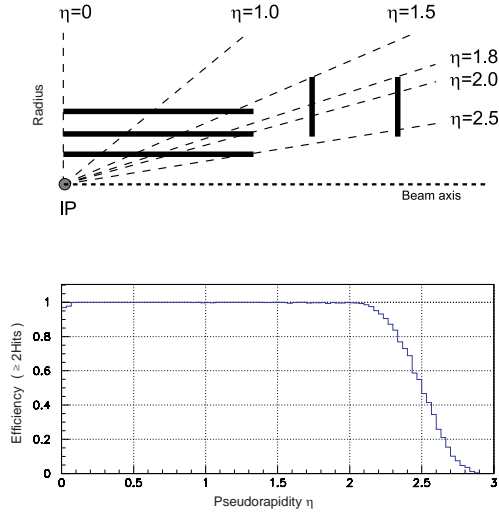


Figure 35: Geometrical layout of the pixel detector and hit coverage as a function of pseudorapidity.

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 so that the drift direction is not parallel to the magnetic field. 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.

4.2.2 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 V) after high hadron fluences. Furthermore the double sided processing allows a guard ring scheme keeping all sensor edges at ground potential.

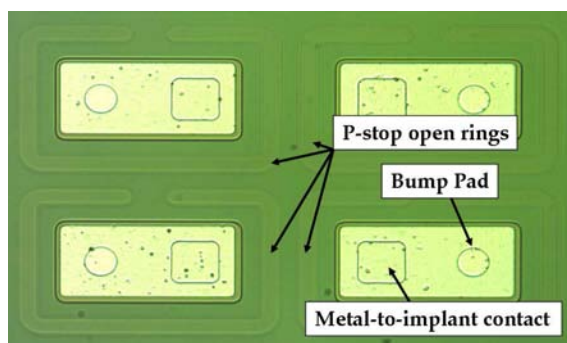


Figure 36: Picture of four pixels in the same double column for a pixel disk sensor.

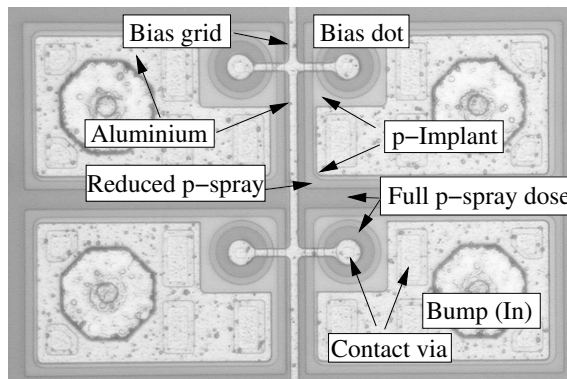


Figure 37: Photograph of four pixel cells. The Indium bumps are already deposited but not yet reflowed.

The isolation technique applied for the regions between the pixel electrodes was developed in close collaboration with the sensor vendors. Open p-stops [19] were chosen for the disks and moderated p-spray [20] for the barrel. Both types of sensors showed sufficient radiation hardness during an extensive qualification procedure including several test beams [21, 22].

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 $8\ \mu\text{m}$ for the p-stop rings and a distance of $12\ \mu\text{m}$ between implants was chosen. Figure 36 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.

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 $10\ \text{nA per cm}^2$. 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 Fig. 37. 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 ($20\ \mu\text{m}$) 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 connec-

tion 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 Fig. 37 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 Fig. 37 has undergone the bump deposition process. The Indium bumps are visible as roughly $50\text{ }\mu\text{m}$ wide octagons.

The sensors are processed on n-doped DOFZ-silicon [23] with $\langle 111 \rangle$ orientation and a resistivity of about $3.7\text{ k}\Omega\text{cm}$ (after processing). This leads to a full depletion voltage of 50-60 V of the $285\text{ }\mu\text{m}$ 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.

4.2.3 Pixel detector read-out

System overview The pixel read-out and control system [24] 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 I²C protocol. Figure 38 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.

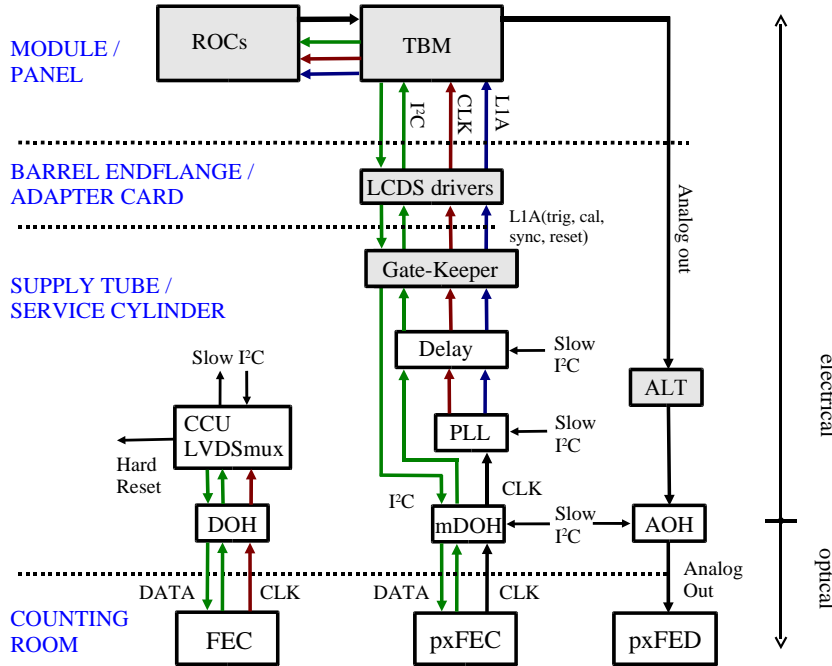


Figure 38: Block diagram of the pixel control and read-out system.

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 $0.25\text{-}\mu\text{m}$ 5-metal-layer CMOS process and contains 52×80 pixels [25]. 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 [25]. The mean threshold dispersion after trimming at $T = -10\text{ }^{\circ}\text{C}$ is 90 electrons equivalent with a noise of 170 electrons.
- 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 I²C 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 $10^{34}\text{ s}^{-1}\text{cm}^{-2}$ and 100 kHz L1 trigger rate.

The power consumption depends on the pixel hit rate. At the LHC design luminosity, the ROC contributes with $34\text{ }\mu\text{W}$ per pixel about 88% (62%) to the total pixel detector front end power budget before (after) the detector has received a total fluence of $6 \times 10^{14}/\text{cm}^2$.

Token Bit Manager chip The TBM [26] 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 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 $40\ \Omega$ 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. 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 [27]. 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.

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 39 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 [28]. 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 I²C 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 I²C 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 I²C 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.

4.2.4 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

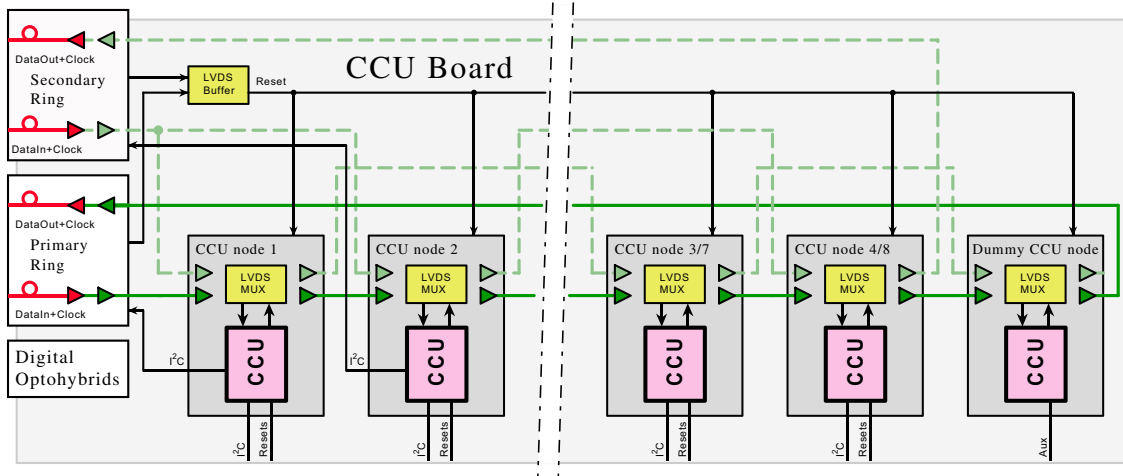


Figure 39: Block diagram of the Pixel front end control system. Note that the total number of CCU nodes is 9 for the BPix and 5 for the FPix.

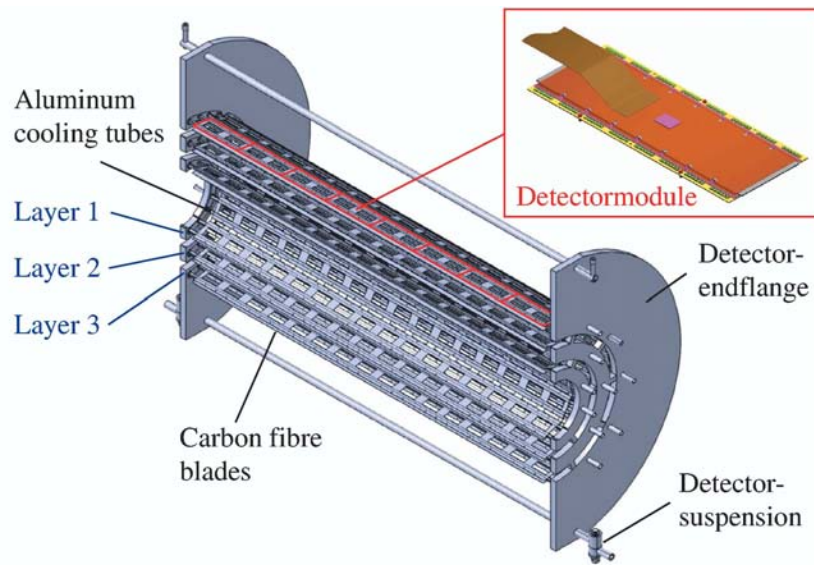


Figure 40: Complete support structure half shell with the three detector layers.

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 40 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 about -10 °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

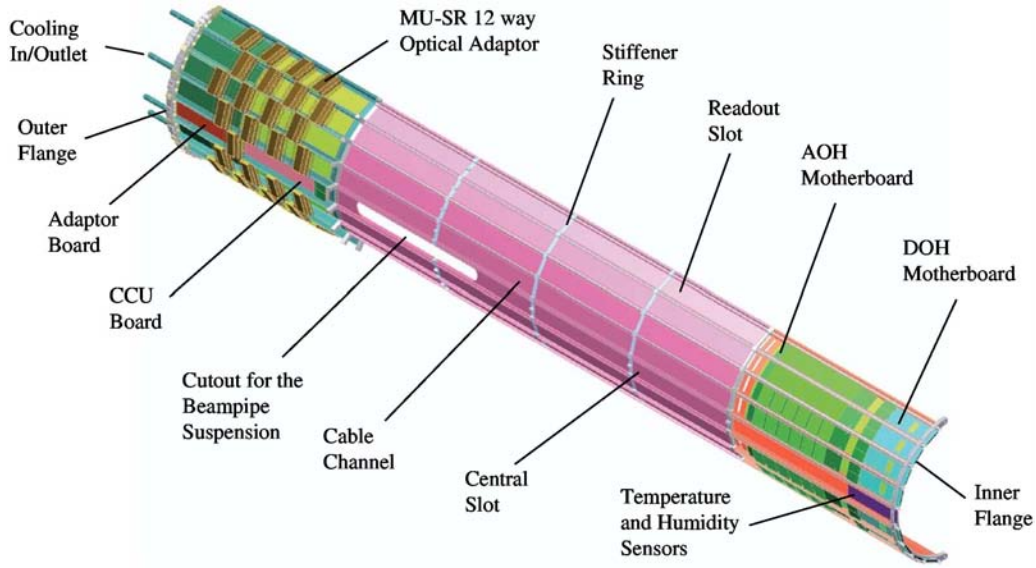


Figure 41: Overview of a supply tube half shell.

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 (Fig. 41).

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 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 Fig. 42 on the right, the edges of the six half-shells are equipped with 16 half-modules each (96 in total, see Fig. 42 on the left).

Geometry and components A module is composed of the following components (Fig. 42). One or two base-strips 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 $100 \times 150 \mu\text{m}^2$ 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

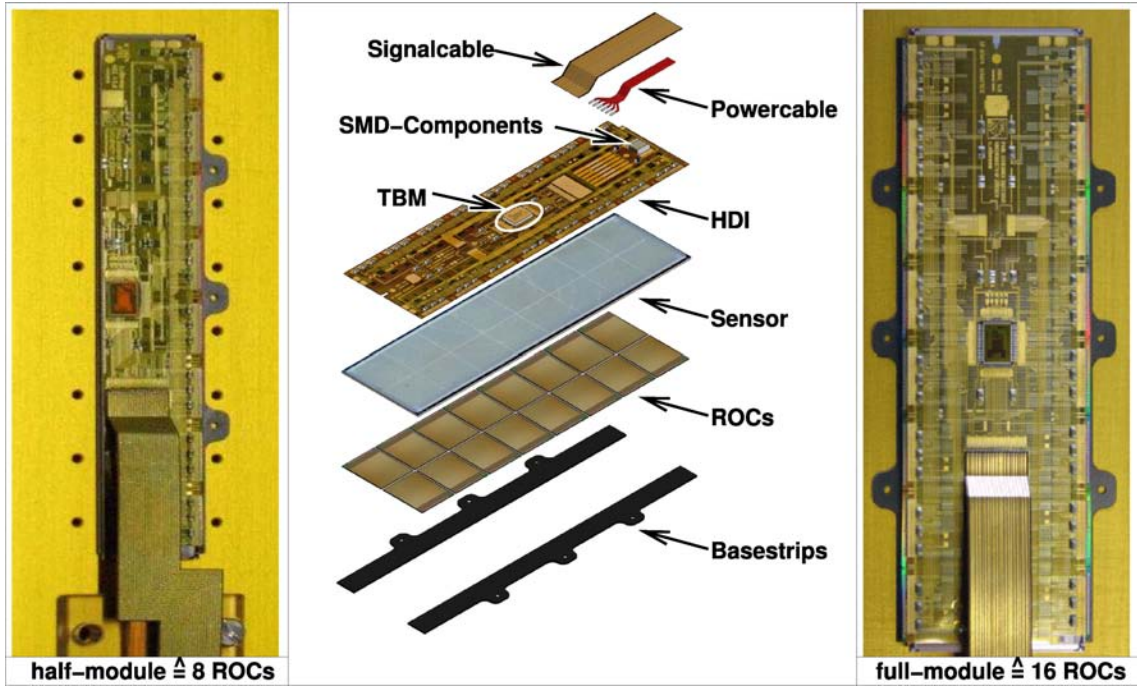


Figure 42: Exploded view (middle panel) of a barrel pixel detector full module (right panel) and picture of an assembled half module (left panel).

copper coated aluminium wires of $250\ \mu\text{m}$ diameter.

A completed full-module has the dimensions $66.6 \times 26.0\ \text{mm}^2$, weighs 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 Fig. 43.

4.2.5 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 Fig. 44. 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 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 ± 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 $-20\ ^\circ\text{C}$ 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 induced by the $\vec{E} \times \vec{B}$ drift.

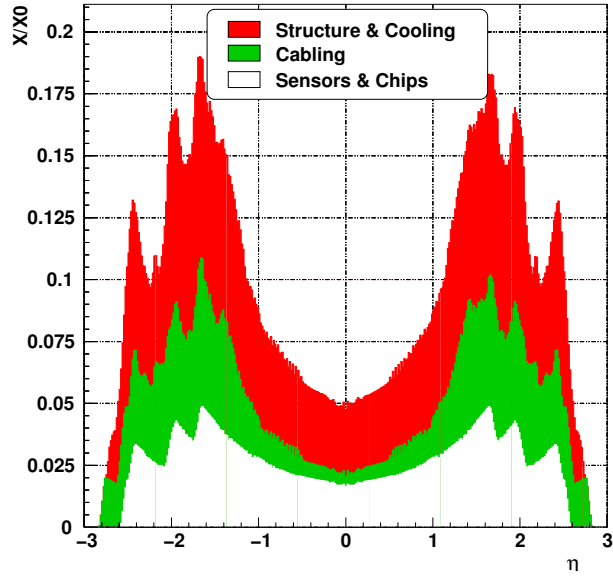


Figure 43: Material budget of the pixel barrel in units of radiation length versus rapidity. The plot does not contain contributions from the pixel support cylinder, the supply tube and cabling from the detector end flange to the supply tube.



Figure 44: 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.

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 $-15^\circ C$. 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 $\approx 2^\circ C$. 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 about $-10^\circ C$. The deformation (magnitude and direction) of the panels on a half-disk, when its temperature changes from $20^\circ C$ to $-20^\circ C$ has been measured to be $150\ \mu m$. 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 Fig. 45.

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 $16\text{ mm} \times 35\text{ mm}$. 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 Fig. 46. 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 $150\ \mu m$ 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 $300\ \mu m$ 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.

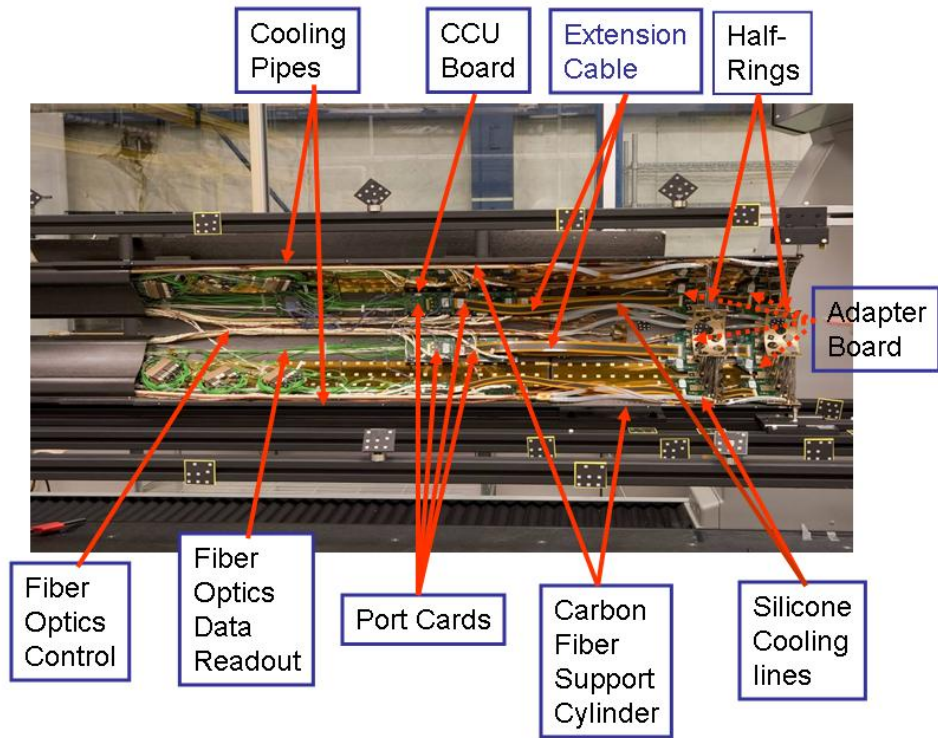


Figure 45: Overview of the Forward Pixel half-cylinder. A photograph of the portion of the first production half-cylinder facing the interaction region (IR). The aluminium flange, the filter boards (see below), and the CCU board are not shown. The half-cylinder is mounted in a survey fixture. The carbon 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.

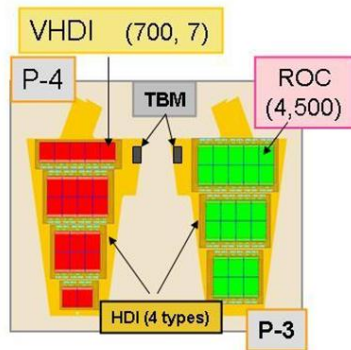


Figure 46: Sketches of the two types of FPix panels showing the different sizes and numbers of the plaquettes on each (left side). A photograph of an actual 3-plaquette panel (right side).

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 $I d\vec{l} \times \vec{B}$ 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 [29]. 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 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.

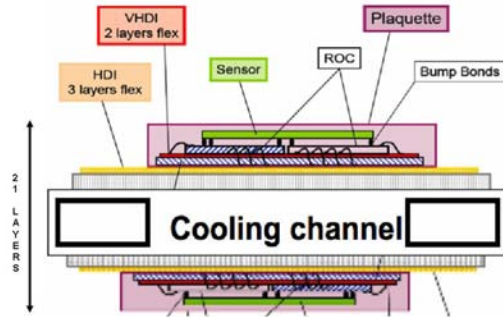


Figure 47: Sketch of a plaquette mounted on a panel showing its several layers

A panel is built up out of several layers of components. These are shown in Fig. 47. 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”. 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 [30].

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.

4.2.6 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 [31]. 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 (≈ 100 m) to the branch controllers. They are fed by local 3-phase 230 V_{AC} to 48 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 V/20 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 Ω 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 μ s) 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 μ s) 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 $6 \times 4 \text{ mm}^2$ 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 $6 \mu\text{H/m}$, 0.13 nF/m and 24Ω 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 $26 \times 0.75 \text{ mm}^2$ and 4 twisted pair copper lines with 0.1 mm^2 for the sense wires.

4.2.7 Cooling

The power consumption per pixel amounts to around $55 \mu\text{W}$, including about $13 \mu\text{W}$ from the sensor leakage current at final fluences of $6 \times 10^{14}/\text{cm}^2$. For the total of ≈ 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.

4.2.8 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 $1 \text{ k}\Omega$ (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 [32]. 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 [33] 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 [34].

The Barrel and Forward Pixel slow controls system is integrated into the PVSS graphical user interface (Sect. 10) of the main CMS DCS.

4.3 Silicon strip tracker

4.3.1 Silicon sensors

The sensor elements in the strip tracker are single sided p -on- n type silicon micro-strip sensors [35, 36]. 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 [37] 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 (Fig. 30), thin sensors of $(320 \pm 20) \mu\text{m}$ wafer thickness are used, with substrate resistivity of $\rho = 1.55 - 3.25 \text{ k}\Omega \text{ cm}$. TOB and the outer 3 rings of the TECs are equipped with thicker sensors of $(500 \pm 20) \mu\text{m}$ thickness, with substrate resistivity of $\rho = 4 - 8 \text{ k}\Omega \text{ cm}$. Due to the single sided processing, these sensors show a significant bow, which is required to be less than $100 \mu\text{m}$.

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 (Sect. 4.3.7) feature a 10 mm hole in the 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 $(1.5 \pm 0.5) \text{ M}\Omega$ 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 [37]. 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 $8 \mu\text{m}$ 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 48 shows the layout of a corner of the active region of a sensor.

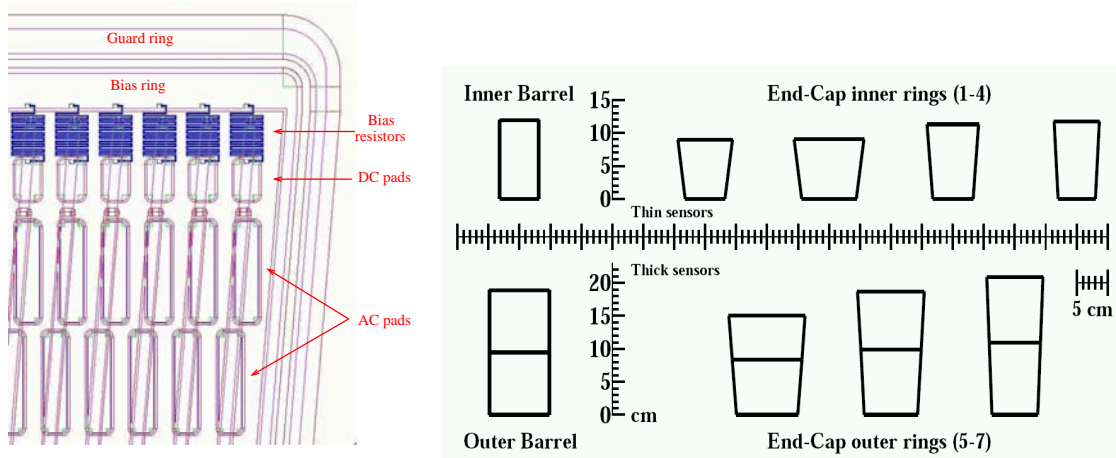


Figure 48: Left panel: Drawing of one corner of the active region of a wedge-shaped silicon strip sensor for the tracker endcaps. Right panel: Silicon sensor geometries in the CMS tracker.

In order to equip all regions in the CMS tracker, 15 different sensor geometries are needed [35] (Fig. 48): 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 $6 \times 12 \text{ cm}^2$ and $10 \times 9 \text{ cm}^2$ 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 m^2 [35].

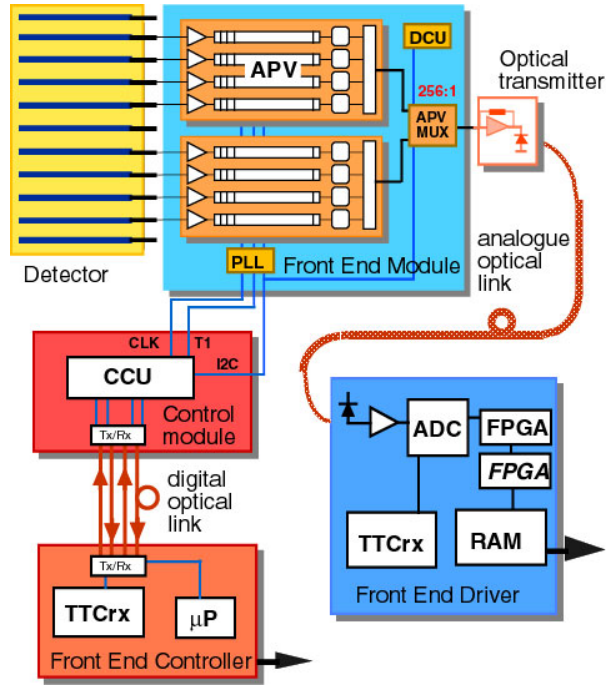


Figure 49: Read-out scheme of the CMS tracker.

4.3.2 Read-out system

The signals from the silicon sensors are amplified, shaped, and stored by a custom integrated circuit, the APV25 [38]. 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 Fig. 49. This analogue read-out scheme was chosen for several reasons: optimal spatial resolution 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.

Front-end ASICs The APV25 has been designed in an IBM $0.25\ \mu\text{m}$ 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 [39]. 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 $4\ \mu\text{s}$ 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 $ENC_{peak} = 270e + 38e/pF \cdot C_{det}$ in peak mode and $ENC_{deconv} = 430e + 61e/pF \cdot C_{det}$ in deconvolution mode, both measured at room temperature [38]. 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 [40]. 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 [41] 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 [44]. All 96 fibres are processed in parallel. The optical signals are converted to electrical levels by opto-receivers [42] 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.

Control and monitoring Clock, trigger and control data are transmitted to the tracker by Front End Controller (FEC) cards [43]. 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) [45] 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 [46] on each detector module while the industry standard I²C protocol [47] 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 [48] 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.

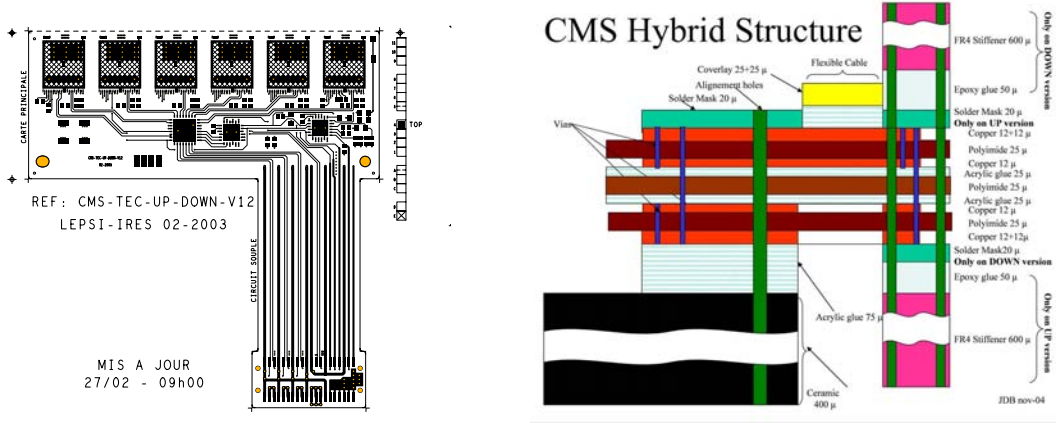


Figure 50: Front-end hybrid layout (example for TEC shown on the left) and arrangement of layers.

Hybrids The front-end read-out electronics for a detector module is mounted onto a multi chip module called hybrid [49]. 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 (Fig. 50). 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 $120 \mu\text{m}$ 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) [31], featuring two low-voltage 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 $\approx 50\text{-m-long}$ *low impedance cables* [31]. 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 (Sect. 4.3.8). 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.

4.3.3 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 ($320\ \mu\text{m}$) or two thick ($500\ \mu\text{m}$) 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 51 shows an exploded view and a photograph of a TEC module.

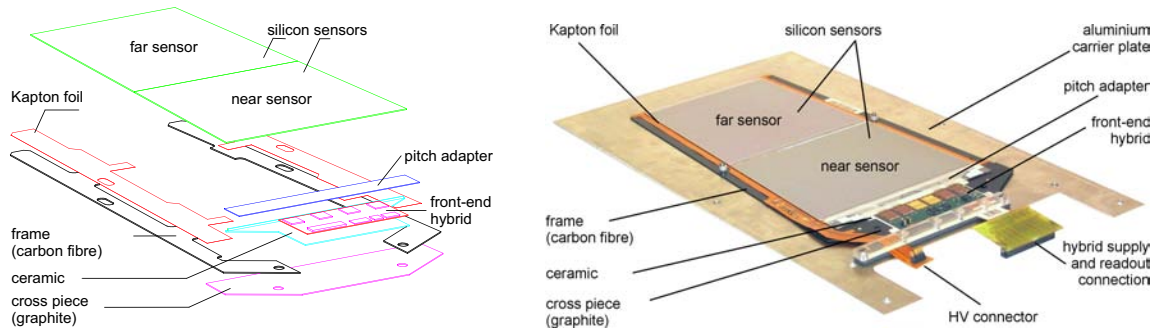


Figure 51: Left panel: Exploded view of a module housing two sensors. Right panel: Photograph of a TEC ring 6 module, mounted on a carrier plate.

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 $6243.1\ \text{mm}^2$ (TEC, ring 1) and $17202.4\ \text{mm}^2$ (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 $(780\pm5)\ \mu\text{m}$ thick graphite (FE779 carbon). For the two-sensor modules a similar U-shaped support structure is obtained by gluing two $(640\pm40)\ \mu\text{m}$ thick carbon fiber legs (K13D2U CFC, $5 \times 125\ \mu\text{m}$ fabric, cyanate ester resin (CE3)) on a $800\ \mu\text{m}$ thick graphite cross-piece (FE779 carbon) which holds the front end electronics. In the inner barrel a $550\ \mu\text{m}$ 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, $5 \times 125\ \mu\text{m}$ 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

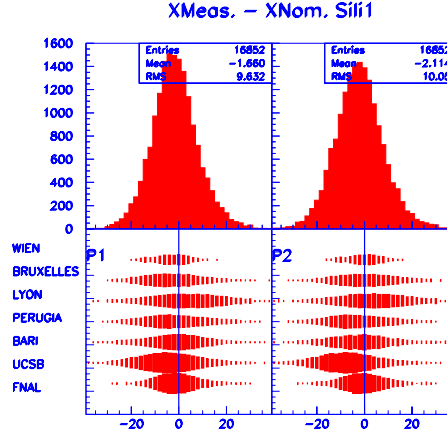


Figure 52: A typical residual distribution (in μm) for a reference point on the modules is shown for the different module assembly centers, indicating a precision of $10 \mu\text{m}$ (RMS) in the module production.

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 $20 \mu\text{m}$ 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 ($80 \mu\text{m}$ – $205 \mu\text{m}$ depending on sensor type) to the APV pitch of $44 \mu\text{m}$. 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 $550 \mu\text{m}$ thick glass substrate (Schott D263 glass), cut to the correct dimensions, with a pattern of low resistivity aluminium strips. For TIB $300 \mu\text{m}$ thick glass (Corning 1737F or G glass) is used. The $10 \mu\text{m}$ narrow lines are etched on a $(1.0\text{--}1.5) \mu\text{m}$ thick aluminium layer deposited on a chromium base, resulting in less than $25 \text{ m}\Omega/\square$.

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 $10 \mu\text{m}$ (RMS) has been achieved and one example from the quality control can be seen in Fig. 52.

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 $25 \mu\text{m}$ 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 to a supporting strip ($400 \mu\text{m}$ 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 [50]. 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%.

4.3.4 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 $80\text{ }\mu\text{m}$, while the outer two layers host single sided modules with a strip pitch of $120\text{ }\mu\text{m}$. 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.

Two service cylinders are coupled to the ends of $\text{TIB}\pm$ (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 53 shows a schematic drawing of one half TIB/TID structure together with its corresponding *margherita*.

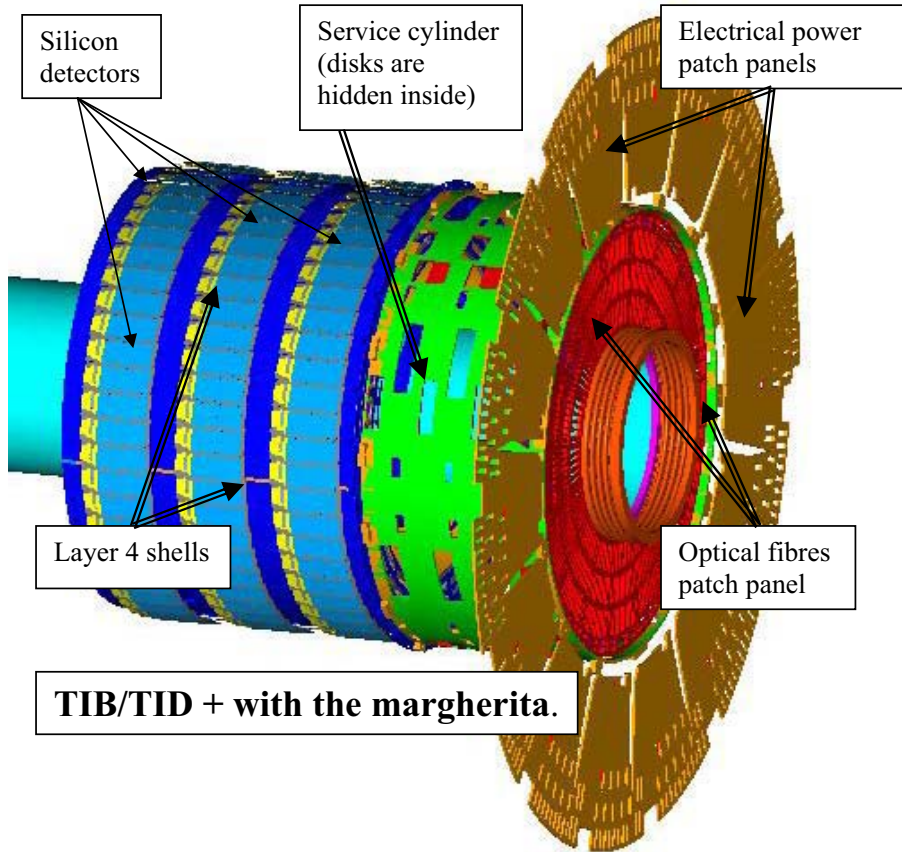


Figure 53: Schematic drawing of the TIB/TID+ subassembly. This structure and its twin (TIB/TID-) nest inside the Tracker Outer Barrel (TOB), one for each end. Services routed out from the *margherita* consist of copper cables for powering and slow controls, optical fibers for signals and controls and also cooling fluid supply lines made of aluminium tubing.

The $\text{TID}\pm$ are assemblies of three disks placed in z between ± 800 mm and ± 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

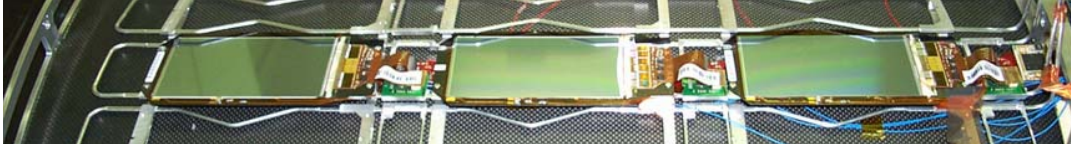


Figure 54: Three TIB modules mounted on a layer 3 shell. The Kapton mother cable runs underneath. A CCUM module at the end of the string interfaces the modules to the control ring. Also visible are the three analog opto-hybrids (see text) and fibers.

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 (Fig. 54), 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 (Fig. 54). 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.

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.

4.3.5 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 (Fig. 55). 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 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 $100\ \mu\text{m}$, and a reproducibility between different disks at the $10\ \mu\text{m}$ 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 $100\ \mu\text{m}$ of nominal values over the whole TOB dimensions, with maximum deviations observed around $200\ \mu\text{m}$.

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 $159 \times 1130 \times 22\ \text{mm}^3$, 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.

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

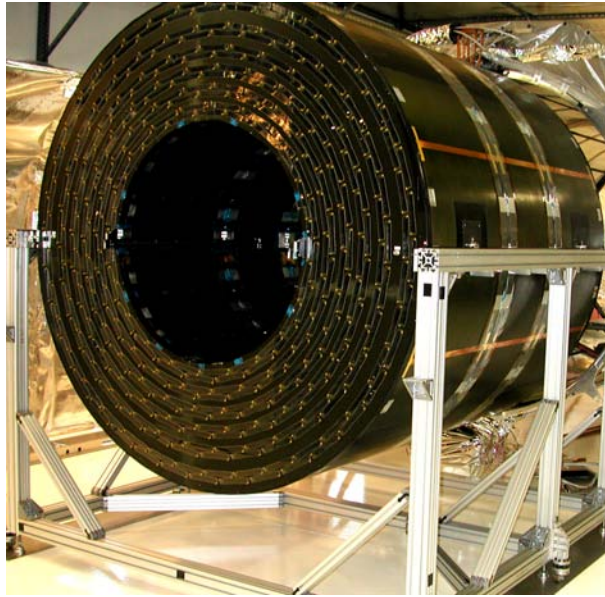


Figure 55: Picture of the TOB wheel.

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 ± 3.3 mm in single-sided rods, ± 3.3 mm and ± 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 $100\text{ }\mu\text{m}$ and $200\text{ }\mu\text{m}$ 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 (Fig.56). 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 (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



Figure 56: Photo of a rod frame equipped with electronics components, ready to receive silicon detector modules

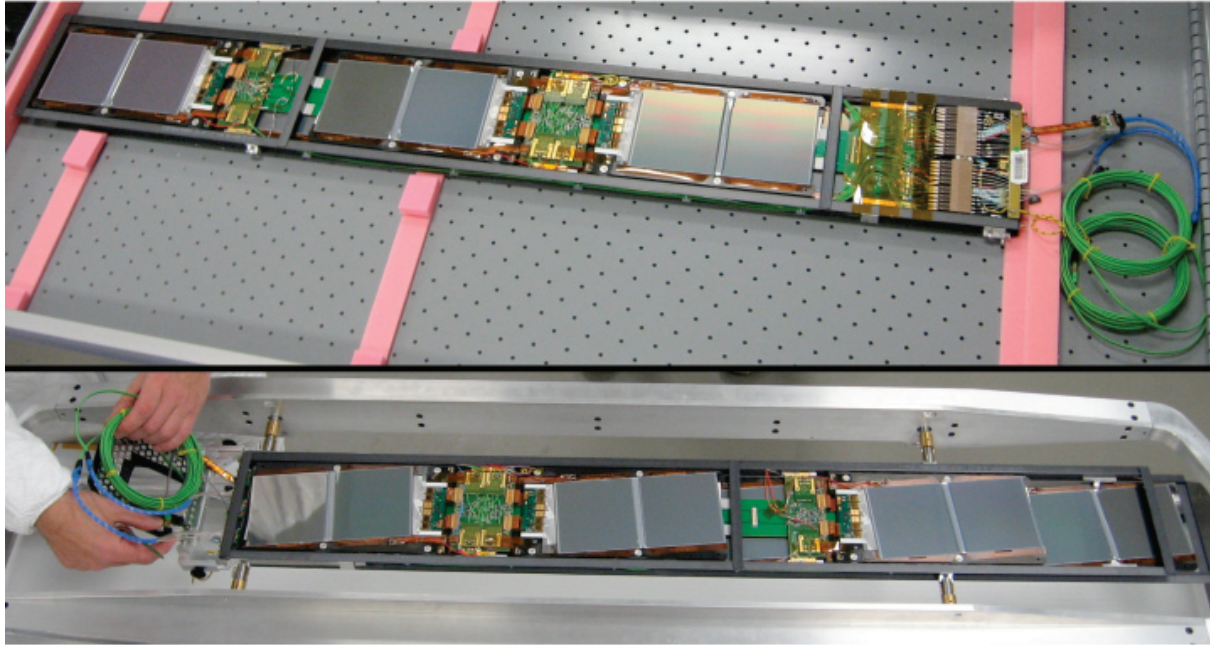


Figure 57: Top panel: photo of an assembled double-sided rod, showing the CCUM side, with the 12-way optical ribbons connected to the AOH fibers. Bottom panel: double-sided rod being prepared for insertion in the TOB mechanics; the side opposite to the CCUM is shown.

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/ICC 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 (Fig. 57). 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.

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 (Fig. 58), are protected by alodined aluminium plates of 0.5 mm thickness, which are locally connected to the power return line.

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 alodined aluminium bar of $10 \times 10 \text{ mm}^2$ 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 Fig. 55); 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.

4.3.6 Tracker EndCaps (TEC)

Mechanical structure The endcaps extend radially from 220 mm to 1135 mm and from $\pm 1240 \text{ mm}$ to $\pm 2800 \text{ mm}$ 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 Figs. 59 and 60. 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

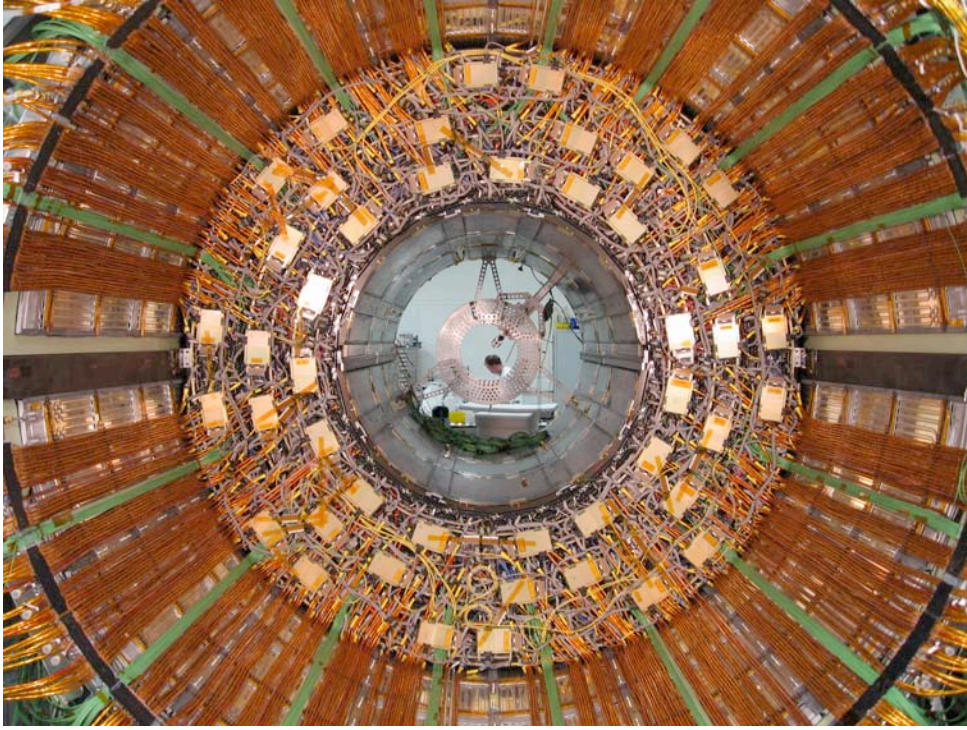


Figure 58: Photo of the completed $z+$ side of the TOB. The DOHMs form the outer layer of the services on the TOB end flange. Optical ribbons (green) run out, grouped in 16 channels. Power cables and feeding pipes run parallel to each other on the thermal screen panels.

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 (Figs. 61 and 62). 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 — (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 ϕ , 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

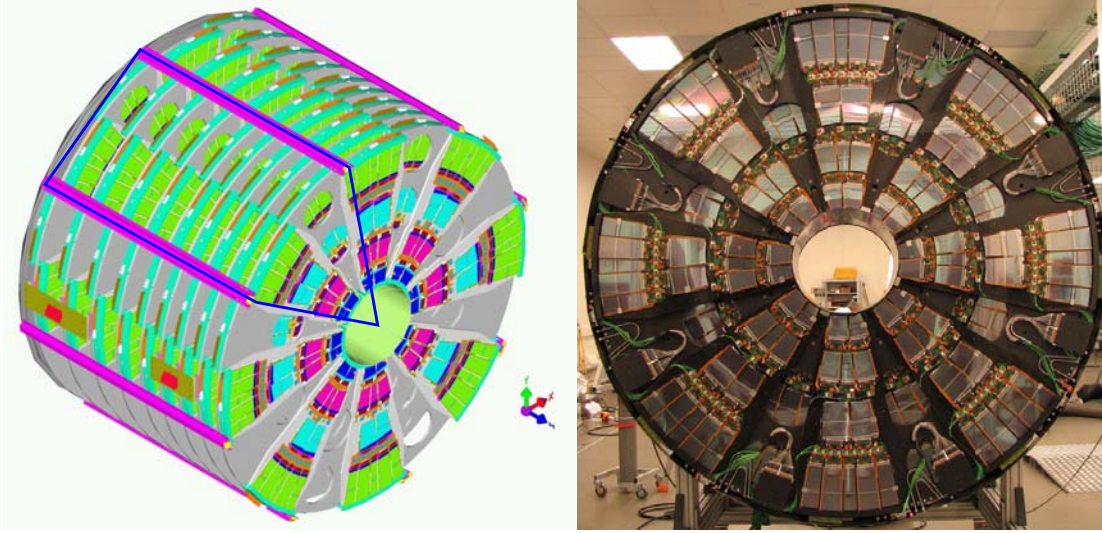


Figure 59: Left panel: Sketch of one tracker endcap. Modules are arranged in rings around the beam axis. They are mounted on trapezoidal sub-structures called *petals*. One sector, indicated with a line, consists of nine front petals mounted on the disk sides facing the interaction point (3 FD13, 3 FD46, 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 60: Side view of a TEC.

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 Fig. 61 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 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 I²C 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. 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 k Ω 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 k Ω 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

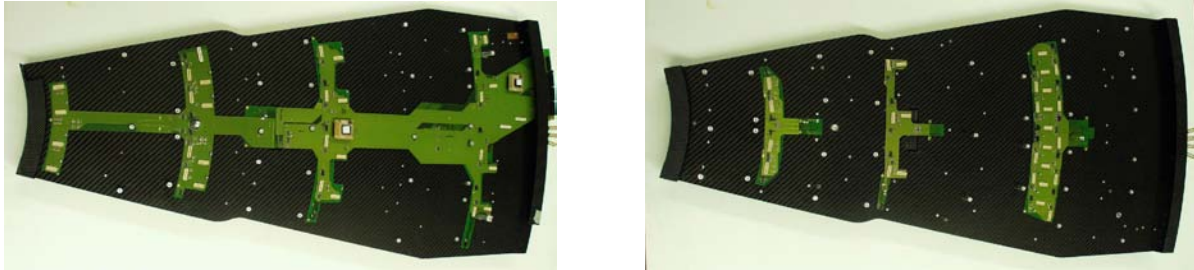


Figure 61: The different ICBs on the two sides of a front petal: ICB_2 and ICB_46 on side B, and ICB_1, ICB_3 and ICB_57 on side A (from left to right). On ICB_46, the two CCUMs are plugged.

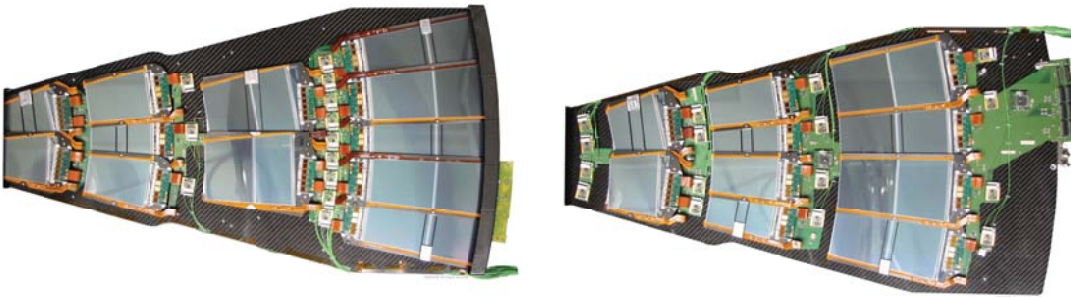


Figure 62: Left photograph: front side of a TEC Petal. Right photograph: back side.

Table 4: 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 $10\ \mu\text{m}$.

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	

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 $20\ \mu\text{m}$ 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.

4.3.7 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 (Fig. 33), 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 [35] and achieved sensor placement accuracy (Fig. 52).

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 photogrammetry with an accuracy of $150\ \mu\text{m}$ ($80\ \mu\text{m}$) under good (optimal) conditions for relative measurements. The measured and expected mounting precision from those data are summarized in Table 4. 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, Fig. 63) uses infrared laser beams with a wavelength $\lambda = 1075\ \text{nm}$ 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

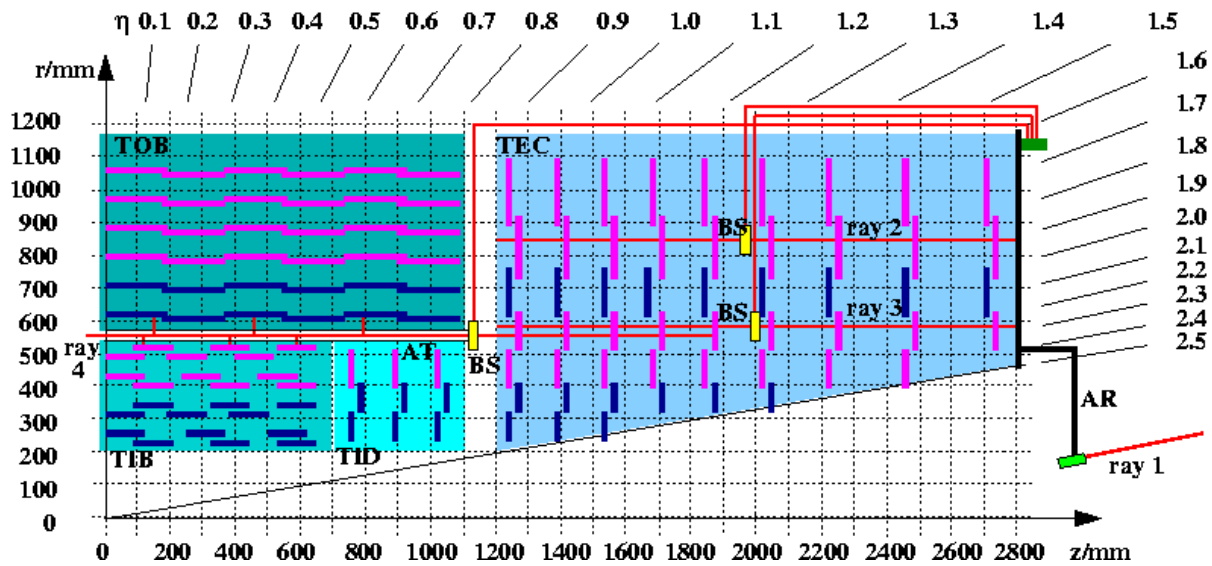


Figure 63: Overview of the CMS Laser Alignment System.

is to generate alignment information on a continuous basis, providing geometry reconstruction of the tracker sub-structures at the level of $100 \mu\text{m}$, 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 $10 \mu\text{m}$, 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 alignment 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 [51], 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 [52], which bridges the gap between global and local algorithms by taking into account the most important correlations. In addition the HIP [53] 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.

During integration of the TEC+, deviation of disk positions and rotations from nominal values have been deter-

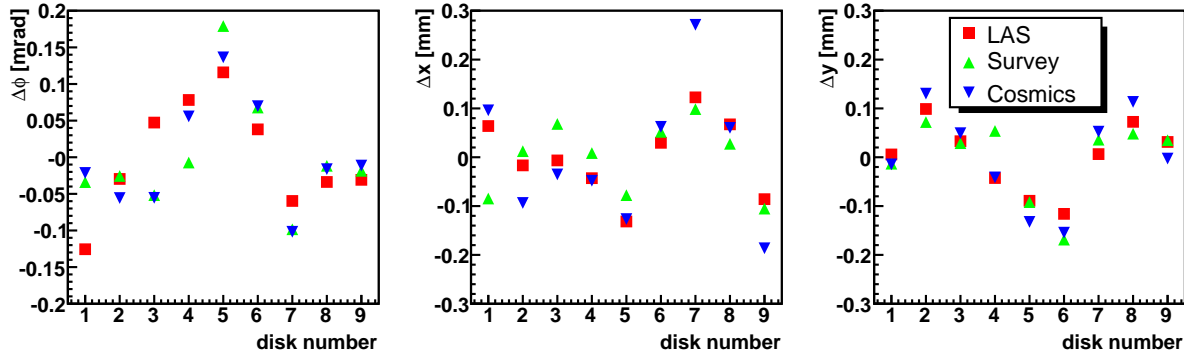


Figure 64: TEC+ disk rotation $\Delta\phi$ (around the beam axis) and displacements $\Delta x, \Delta y$ (in the disk plane) as determined from survey, LAS and cosmic muon tracks.

mined from survey with photogrammetry, the LAS, and tracks from cosmic muons. Figure 64 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 $60 \mu\text{m}$ and $80 \mu\text{rad}$ with each other, which can be taken as an upper value on the precision of each method.

4.3.8 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, Sect. 10) 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. All parameters are archived to ORACLE. The TDSS (tracker DCS) system is fully implemented in the global CMS DSS (DCS) and Run Control system.

4.3.9 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 [54]. 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/cm².

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 5. 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 [54]. 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 μ m 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 5. 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 Fig. 65. A linear fit to these data yields the following dependence of the ENC on the strip length L :

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

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 I²C 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 com-

Module type	Pitch [μm]	Strip length [mm]	S/N	S/N	ENC [e^-]	ENC [e^-]
			Peak mode	Dec. mode	Peak mode	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

Table 5: 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.

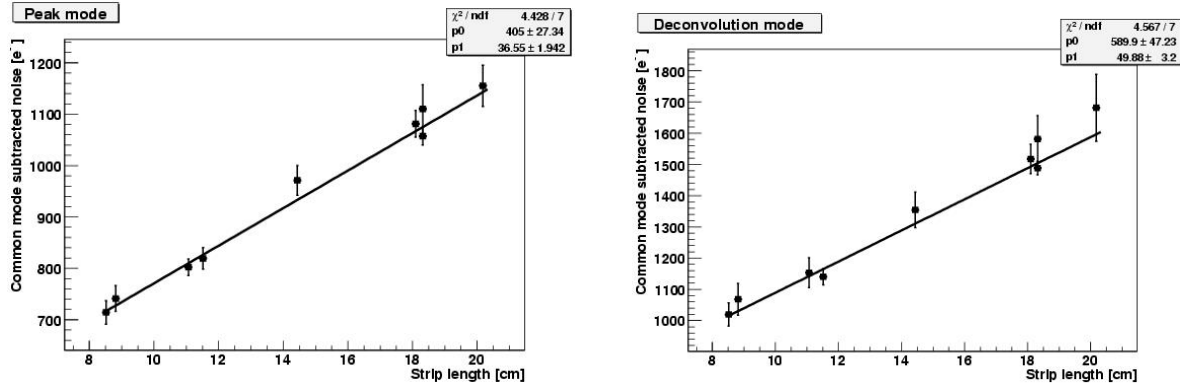


Figure 65: Equivalent noise charge after common mode subtraction versus strip length for all TOB and TEC module types, in peak (left panel) and deconvolution mode (right panel).

ments 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 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 I²C communication 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.

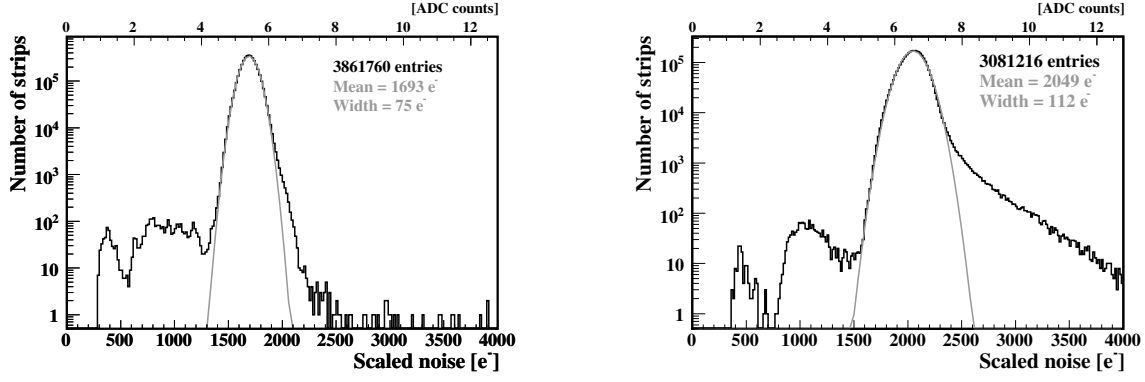


Figure 66: Normalized common mode subtracted noise of all strips (scaled to the strip length of ring 1 sensors) of both TECs (left panel) and the TOB (right panel), in deconvolution mode. Details are described in the text.

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 Fig. 66, 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 Fig. 66). 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 ± 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 $(22 \pm 4) \%$ and $(21 \pm 3) \%$ of the mean intrinsic noise in peak and deconvolution mode, respectively (Fig. 67, 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 $(2.5 \pm 0.2) \%$ in both read-out modes, as shown in Fig. 67, 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 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 [55], 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.

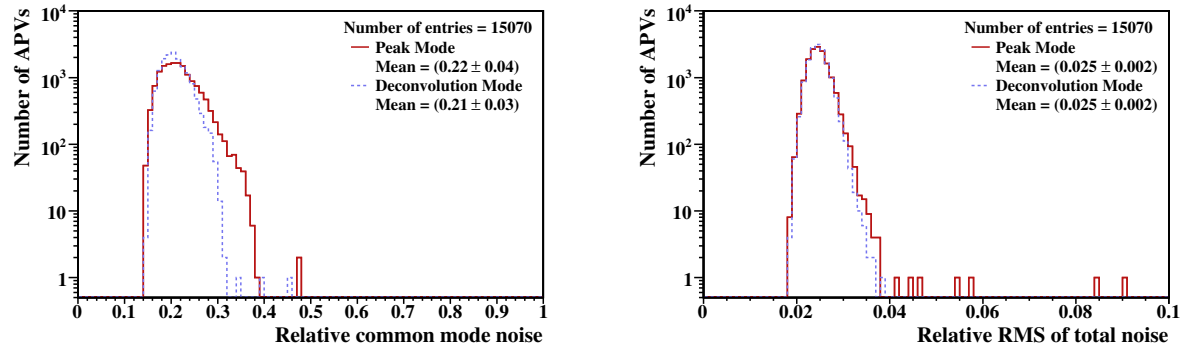


Figure 67: Ratio between common mode noise and mean intrinsic noise (left panel) and ratio between the RMS of the total noise and the mean total noise (right panel), calculated per APV, in peak and deconvolution mode for all non-defective APVs of TEC+.

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 Fig. 68 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 $+10^\circ\text{C}$ and hybrid temperatures of $+10^\circ\text{C}$ to $+30^\circ\text{C}$ 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 $+24^\circ\text{C}$ and the hybrid temperature about $+30^\circ\text{C}$.

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 Fig. 66, 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-.

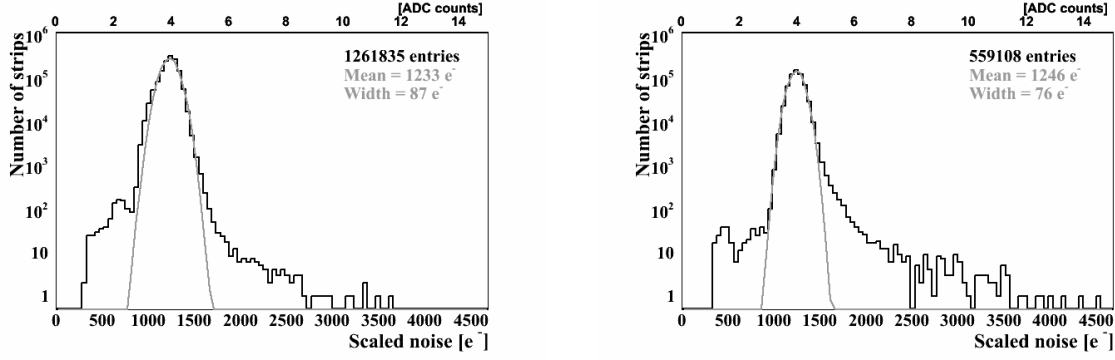


Figure 68: Normalized common mode subtracted single strip noise for TIB (left panel) and TID (right panel), in deconvolution mode. Details are described in the text.

Irradiation studies As already discussed in detail in Sect. 4.1.1, the silicon strip tracker will suffer from a severe level of radiation during its 10 year long lifetime: up to $1.8 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$ for TIB/TID and TEC and up to $0.5 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$ for TOB, assuming an integrated luminosity of 500 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 [57]). 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 $100 \mu\text{A}$ and a beam spot diameter of 1 cm is available.

To study the performance of complete irradiated modules, several OB1 and OB2 modules (Table 5 for explanation) were irradiated with a proton fluence ranging from $0.1 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$ to $0.7 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$, and one OB2 module was subjected to a neutron fluence of about $1.2 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$ [56]. Two TEC W5 modules were irradiated with a proton fluence of about $1.1 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$, and three TIB IB1 modules were subjected to a proton fluence of $0.5 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$ to $2.1 \times 10^{14} \text{ n}_{\text{eq}}\text{cm}^{-2}$. 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 Fig. 69 (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 excellent agreement with the Hamburg model [58], 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 Fig. 69 (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 $(3.79 \pm 0.27) \times 10^{-17} \text{ A/cm}$, 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 Fig. 70. 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 [59, 60].

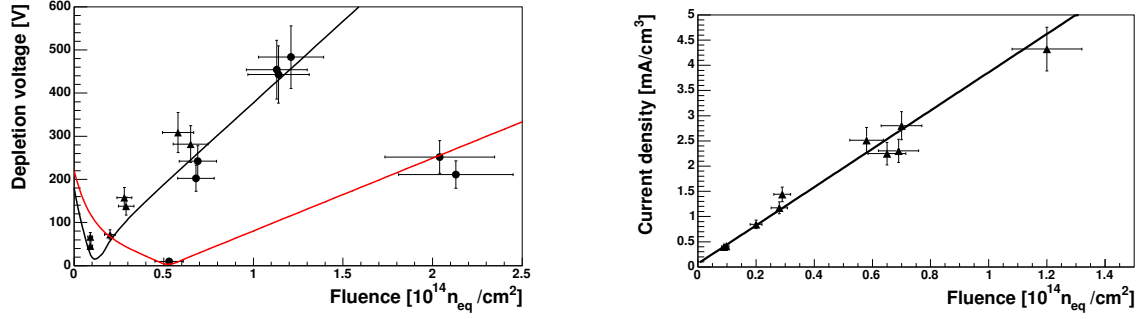


Figure 69: Left panel: Variation of depletion voltage with fluence for OB1 (triangles), OB2 and W5 (dots on upper curve) and IB2 (dots on lower curve) modules, after an annealing time of 80 minutes after each irradiation step. The curves correspond to calculations for 500 μm (upper curve) and 320 μm (lower curve) sensors for an annealing time of 80 minutes at 60 $^{\circ}\text{C}$. Right panel: Current density, scaled to 20 $^{\circ}\text{C}$, versus fluence after annealing for 80 minutes at 60 $^{\circ}\text{C}$.

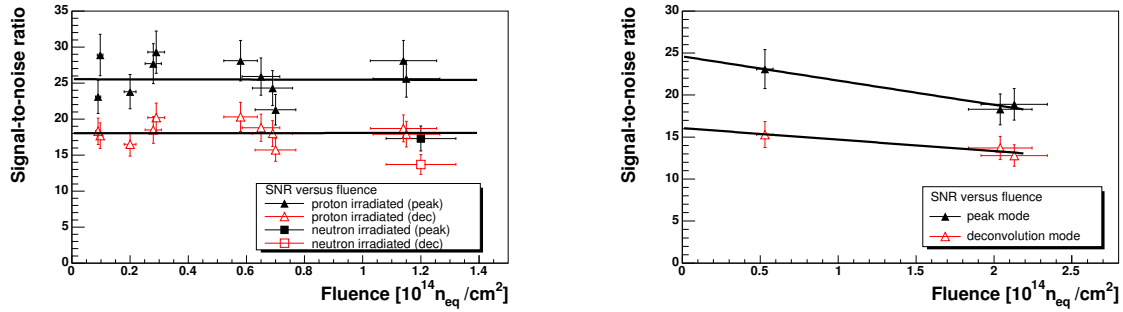


Figure 70: Signal-to-noise ratio versus fluence for modules with 500 μm (left panel) and 320 μm thick sensors (right panel) in peak (filled symbols) and deconvolution mode (open symbols).

5 Electromagnetic calorimeter

The electromagnetic calorimeter of CMS (ECAL) is a hermetic homogeneous calorimeter made of 61 200 lead tungstate (PbWO_4) crystals mounted in the central barrel part, closed by 7 324 crystals in each of the two end-caps. A preshower detector is placed in front of the endcaps crystals. Avalanche photodiodes (APDs) are used as photodetectors in the barrel and vacuum phototriodes (VPTs) in the endcaps. The use of high density crystals has allowed the design of a calorimeter which is fast, has fine granularity and is radiation resistant, all important characteristics in the LHC environment. One of the driving criteria in the design was the capability to detect the decay to two photons of the postulated Higgs boson. This capability is enhanced by the good energy resolution provided by a homogeneous crystal calorimeter.

5.1 Lead tungstate crystals

The characteristics [61] of the PbWO_4 crystals make them an appropriate choice for operation at LHC. The high density (8.28 g/cm^3), short radiation length (0.89 cm) and small Molière radius (2.2 cm) result in a fine granularity and a compact calorimeter. In recent years, PbWO_4 scintillation properties and other qualities have been progressively improved, leading to the mass production of optically clear, fast and radiation hard crystals [62, 63]. The scintillation decay time of these production crystals is of the same order of magnitude as the LHC bunch crossing time: about 80% of the light is emitted in 25 ns. The light output is relatively low and varies with temperature ($-2.1\%^\circ\text{C}^{-1}$ at 18°C [64]): at 18°C about 4.5 photoelectrons per MeV are collected in both APDs and VPTs. The crystals emit blue-green scintillation light with a broad maximum at 420–430 nm [63, 65]. Optical transmission and radioluminescence spectra are shown in Fig. 71.

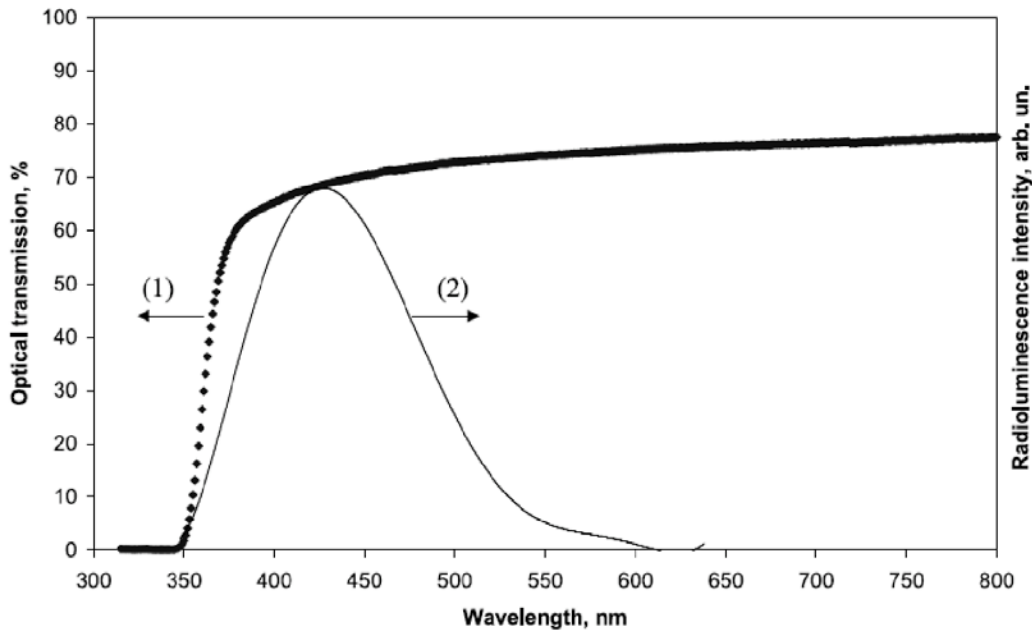


Figure 71: Optical transmission (1, left scale) and radioluminescence intensity (2, right scale) for production PbWO_4 crystals.

To exploit the total internal reflection for optimum light collection on the photodetector, the crystals are polished after machining. For fully polished crystals, the truncated pyramidal shape makes the light collection non-uniform along the crystal length. The effect is large because of the high refractive index $n = 2.29$ [66] and the needed uniformity [67] is achieved by depolishing one lateral face. In the endcaps, the light collection is naturally more uniform because the crystal faces are nearly parallel.

The crystals have to withstand the radiation levels and particle fluxes [68] anticipated throughout the duration of the experiment. Ionizing radiation produces absorption bands through the formation of colour centres due to oxygen vacancies and impurities in the lattice. The practical consequence is a wavelength-dependent loss of light transmission without changes to the scintillation mechanism, a damage which can be tracked and corrected for by monitoring the optical transparency with injected laser light (Sect. 5.9). The damage reaches a dose-rate dependent

equilibrium level which results from a balance between damage and recovery at 18 °C [63, 69]. To ensure an adequate performance throughout LHC operation, the crystals are required to exhibit radiation hardness properties quantified as an induced light attenuation length (at high dose rate) greater than approximately 3 times the crystal length even when the damage is saturated. Hadrons have been measured to induce a specific, cumulative reduction of light transmission, but the extrapolation to LHC conditions indicates that the damage will remain within the limits required for good ECAL performance [70, 71].

5.2 The ECAL layout and mechanics

The barrel part of the ECAL (EB) covers the pseudorapidity range $|\eta| < 1.479$. The barrel granularity is 360-fold in ϕ and (2×85) -fold in η , resulting in a total of 61 200 crystals. The crystals have a tapered shape, slightly varying with position in η . They are mounted in a quasi-projective geometry to avoid cracks aligned with particle trajectories, so that their axes make a small angle (3°) with respect to the vector from the nominal interaction vertex, in both the ϕ and η projections. The crystal cross-section corresponds to approximately 0.0174×0.0174 in η - ϕ or $22 \times 22 \text{ mm}^2$ at the front face of crystal, and $26 \times 26 \text{ mm}^2$ at the rear face. The crystal length is 230 mm corresponding to $25.8 X_0$. The barrel crystal volume is 8.14 m^3 and the weight is 67.4 t.

The centres of the front faces of the crystals are at a radius 1.29 m. The crystals are contained in a thin-walled alveolar structure (submodule). The alveolar wall is 0.1 mm thick and is made of an aluminium layer, facing the crystal, and two layers of glass fiber-epoxy resin. To avoid oxidation, a special coating is applied to the aluminium surface. The nominal crystal to crystal distance is 0.35 mm inside a submodule, and 0.5 mm between submodules. To reduce the number of different type of crystals, each submodule contains only a pair of shapes, left and right reflections of a single shape. In total, there are 17 such pairs of shapes. The submodules are assembled into modules of different types, according to the position in η , each containing 400 or 500 crystals. Four modules, separated by aluminium conical webs 4 mm thick, are assembled in a supermodule, which contains 1700 crystals (Fig. 72).

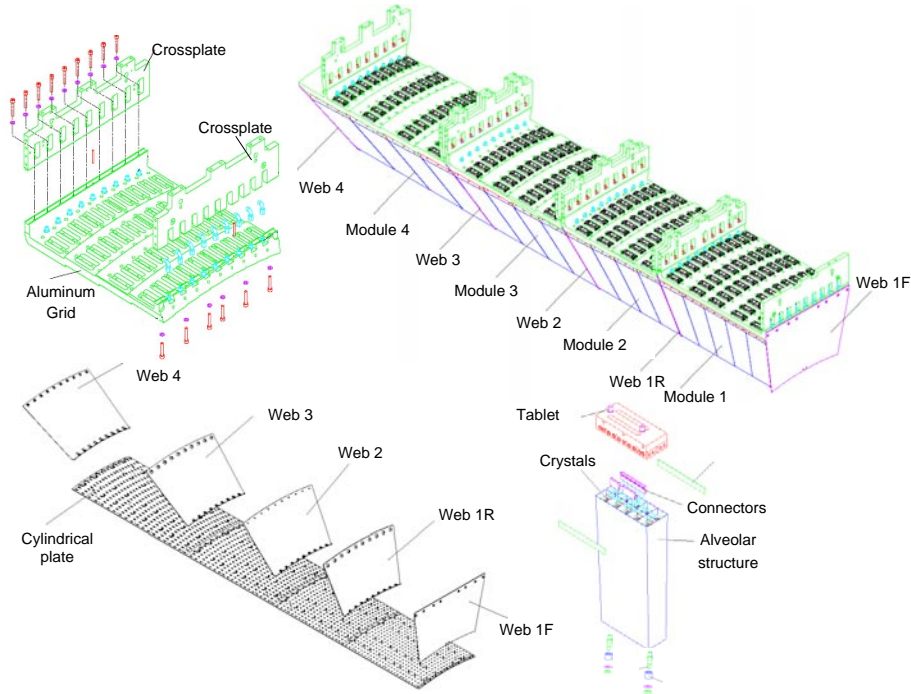


Figure 72: Layout of the ECAL barrel mechanics.

In each module, the submodules are held in partial cantilever by an aluminium grid, which supports their weight from the rear. At the front the submodule free ends are connected together by pincers that cancel the relative tangential displacements. The submodule cantilever is reduced by the action of a 4 mm thick cylindrical plate where the front of the submodules are supported by setpins. Not all the submodules are connected to the cylindrical plate

but only four rows in ϕ from a total of ten. The portion of the submodule load taken at the front by the cylindrical plate is transmitted to the aluminium grids of the different modules via the conical webs interspaced between the modules [72]. Each module is supported and positioned in the supermodule at the rear end through the grid by a spine beam. The spine is provided with pads which slide into rails housed on the front face of the HCAL barrel, allowing the installation and support of each single supermodule. The cylindrical plate in front of the supermodule also provides the fixation of the monitoring system (see below) and the holes for its optical fibers.

All services, cooling manifolds and cables converge to a patch panel at the external end of the supermodule. Eighteen supermodules, each covering 20° in ϕ , form a half barrel.

The endcaps (EE) cover the rapidity range $1.479 < |\eta| < 3.0$. The longitudinal distance between the interaction point and the endcap envelope is 3 154 mm, taking account of the estimated shift toward the interaction point by 1.6 cm when the 4 T magnetic field is switched on. The endcap consists of identically shaped crystals grouped in mechanical units of 5×5 crystals (supercrystals, or SCs) consisting of a carbon-fibre alveola structure. Each endcap is divided into 2 halves, or *Dees*. Each Dee holds 3 662 crystals. These are contained in 138 standard SCs and 18 special partial supercrystals on the inner and outer circumference. The crystals and SCs are arranged in a rectangular x - y grid, with the crystals pointing at a focus 1 300 mm beyond the interaction point, giving off-pointing angles ranging from 2 to 8 degrees. The crystals have a rear face cross section $30 \times 30 \text{ mm}^2$, a front face cross section $28.62 \times 28.62 \text{ mm}^2$ and a length of 220 mm ($24.7 X_0$). The endcaps crystal volume is 2.90 m^3 and the weight is 24.0 t. The layout of the calorimeter is shown in Fig. 73.

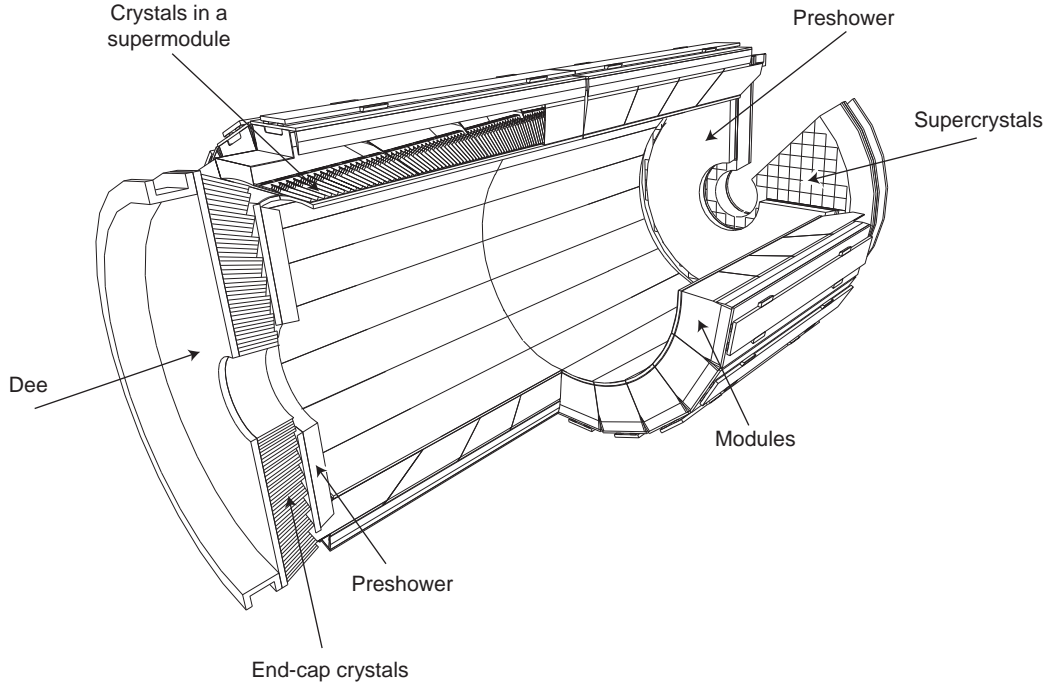


Figure 73: Layout of the CMS electromagnetic calorimeter showing the arrangement of crystal modules, supermodules and endcaps, with the preshower in front.

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent. Both variations are negative with increasing temperature. The overall variation of the response to incident electrons with temperature has been measured in test beam [73] to be $(-3.8 \pm 0.4)\%^\circ\text{C}^{-1}$. The temperature of the system has therefore to be maintained constant to high precision, requiring a cooling system capable of extracting the heat dissipated by the read-out electronics and of keeping the temperature of crystals and photodetectors stable within $\pm 0.05^\circ\text{C}$ to preserve energy resolution. The nominal operating temperature of the CMS ECAL is 18°C . The cooling system has to comply with this severe thermal requirement. The system employs water flow to stabilise the detector. In the barrel, each supermodule is independently supplied with water at 18°C . The water runs through a thermal screen placed in front of the crystals which thermally decouples them from the silicon tracker, and through pipes embedded in the aluminium grid, connected in parallel. Beyond the grid, a 9 mm thick layer of insulating foam (Armaflex) is placed to minimise the heat flowing from the read-out electronics

towards the crystals. Return pipes distribute the water through a manifold to a set of aluminium cooling bars. These bars are in close contact with the very front end electronics (VFE) cards and absorb the heat dissipated by the components mounted on these cards. A thermally conductive paste (gap filler 2000, produced by Bergquist) is used to provide a good contact between the electronic components and a metal plate facing each board. This plate is coupled to the cooling bar by a conductive pad (ultrasoft gap pad, also produced by Bergquist). Both the gap pad and the gap filler have been irradiated with twice the dose expected in the ECAL endcaps after 10 years at the LHC and have shown no change in character or loss of performance.

Extended tests of the cooling system have been performed with good results [73]. Residual effects caused by a possible variation of the power dissipated by the electronics were measured in the extreme case of electronics switched on and off. The conclusion is that contributions to the constant term of the energy resolution due to thermal fluctuations will be negligible, even without temperature corrections.

5.3 Photodetectors

The photodetectors need to be fast, radiation tolerant and be able to operate in the longitudinal 4 T magnetic field. In addition, because of the small light yield of the crystals, they should amplify and be insensitive to particles traversing them (*nuclear counter effect*). The configuration of the magnetic field and the expected level of radiation led to different choices: avalanche photodiodes in the barrel and vacuum phototriodes in the endcaps. The lower quantum efficiency and internal gain of the vacuum phototriodes, compared to the avalanche photodiodes, is offset by their larger surface coverage on the back face of the crystals.

5.3.1 Barrel: avalanche photodiodes

In the barrel, the photodetectors are Hamamatsu type S8148 reverse structure (i.e., with the bulk n-type silicon behind the p-n junction) avalanche photodiodes (APDs) specially developed for the CMS ECAL. Each APD has an active area of $5 \times 5 \text{ mm}^2$ and a pair is mounted on each crystal. They are operated at gain 50 and read out in parallel. The main properties of the APDs at gain 50 and 18°C are listed in Table 6.

Table 6: Properties of the APDs at gain 50 and 18°C .

Sensitive area	$5 \times 5 \text{ mm}^2$
Operating voltage	340–430 V
Breakdown voltage - operating voltage	$45 \pm 5 \text{ V}$
Quantum efficiency (430 nm)	$75 \pm 2\%$
Capacitance	$80 \pm 2 \text{ pF}$
Excess noise factor	2.1 ± 0.2
Effective thickness	$6 \pm 0.5 \mu\text{m}$
Series resistance	$< 10 \Omega$
Voltage sensitivity of the gain ($1/M \cdot dM/dV$)	$3.1 \pm 0.1\%/V$
Temperature sensitivity of the gain ($1/M \cdot dM/dT$)	$-2.4 \pm 0.2\%/^\circ\text{C}$
Rise time	$< 2 \text{ ns}$
Dark current	$< 50 \text{ nA}$
Typical dark current	3 nA
Dark current after $2 \times 10^{12} \text{ n/cm}^2$	5 μA

The sensitivity to the nuclear counter effect is given by the effective thickness of $6 \mu\text{m}$, which translates into a signal from a minimum ionizing particle traversing an APD equivalent to about 100 MeV deposited in the PbWO_4 .

For ECAL acceptance each APD was required to be fully depleted and to pass through a screening procedure involving 5 kGy of ^{60}Co irradiation and 1 month of operation at 80°C . Each APD was tested to breakdown and required to show no significant noise increase up to a gain of 300. The screening and testing aimed to ensure reliable operation for 10 years under high luminosity LHC conditions for over 99% of the APDs installed in the ECAL [74]. Based on tests with hadron irradiations [75] it is expected that the dark current after such operation will have risen to about $5 \mu\text{A}$, but that no other properties will have changed. Small samples of APDs were irradiated with a ^{251}Cf source to monitor the effectiveness of the screening procedure in selecting radiation resistant APDs.

The gain stability directly affects the ECAL energy resolution. Since the APD gain has a high dependence on the bias voltage ($\alpha_V = 1/M dM/dV \simeq 3.1\%/V$ at gain 50), to keep this contribution to the resolution at the

level of per mille, the APDs require a very stable power supply system: the stability of the voltage has to be of the order of few tens of mV. This requirement applies to all the electrical system characteristics: noise, ripple, voltage regulation and absolute precision, for short and long term periods. A custom high voltage (HV) power supply system has been designed for the CMS ECAL in collaboration with the CAEN Company [76]. To remain far from high doses of radiation, the HV system is located in the CMS service cavern, some 120 m away from the detector. The HV channels are floating and use sense wire to correct for variations in the voltage drop on the leads. The system is based on a standard control crate (SY1527) hosting 8 boards expressly designed for this application (A1520E). The SY1527 integrate a PC capable to communicate with the board controller via an internal bus and different interfaces are available to integrate the SY1527 on the ECAL detector control system (DCS). The board design is based on a modular concept so that each HV channel is implemented on a separate module and up to 9 channels can be hosted on a single A1520E board. Each channel can give a bias voltage to 50 APD pairs from 0 to 500 V with maximum current of 15 mA. In total, there are 18 crates and 144 boards. Temperature drift compensation is possible due to presence on the crate of temperature probes that can be used to monitor the environment temperature for adjustments of the voltage setting.

The operating gain of 50 requires a voltage between 340 and 430 V. The APDs are sorted according to their operating voltage into bins 5 V wide, and then paired such that each pair has a mean gain of 50. Each pair is mounted in parallel in a *capsule*, a molded receptacle with foam, which is then glued on the back of each crystal. The capsules are connected to the read-out electronics by Kapton flexible printed circuit boards of variable length, dictated by the capsule's position within the submodule. Each capsule receives the bias voltage through an RC filter network and a protection resistor.

One 100 k Ω negative temperature coefficient thermistor from Betatherm, used as temperature sensor, is embedded in every tenth APD capsule. There are twenty-two different types of capsules, differing for the Kapton length and for the presence of the thermistor.

5.3.2 Endcap: vacuum phototriodes

In the endcaps, the photodetectors are vacuum phototriodes (VPTs) (type PMT188 from National Research Institute Electron in St. Petersburg). Vacuum phototriodes are photomultipliers having a single gain stage. These particular devices were developed specially for CMS [77] and have an anode of very fine copper mesh (10 μ m pitch) allowing them to operate in the 4 T magnetic field. Each VPT is 25 mm in diameter, with an active area of approximately 280 mm²; one VPT is glued to the back of each crystal. One Betatherm thermistor is embedded into each supercrystal. The VPTs have a mean quantum efficiency of the bialkali photocathode (SbKCs) of 22% at 430 nm, and a mean gain of 10.2 at zero field. When placed in a strong axial magnetic field, the response is slightly reduced and there is a modest variation of response with the angle of the VPT axis with respect to the field over the range of angles relevant to the CMS endcaps (6° to 26°). The mean response in a magnetic field of 4 T, with the VPT axis at 15° to the field direction, is typically > 90% of that in zero field [78].

All VPTs are tested by the manufacturer before delivery, without an applied magnetic field. All VPTs are also tested on receipt by CMS to determine their response as a function of magnetic field up to 1.8 T. Each device is measured at a set of angles with respect to the applied field, spanning the range of angles covered by the endcaps. In addition, at least 10% of the tubes, selected at random, are also tested in a 4 T superconducting magnet, at a fixed angle of 15°, to verify satisfactory operation at the full field of CMS.

The estimated doses and particle fluences for 10 years of LHC operation are 0.5 kGy and 5×10^{13} n/cm² at the outer circumference of the endcaps and 20 kGy and 7×10^{14} n/cm² at $|\eta| = 2.6$. Sample faceplates from every glass production batch were irradiated with a ⁶⁰Co source to 20 kGy. The faceplates were required to show a transmission loss, integrated over the wavelength range corresponding to PbWO₄ emission, of less than 10%. Irradiation of VPTs in a nuclear reactor to 7×10^{14} n/cm² showed a loss in anode sensitivity entirely consistent with discolouration of the faceplate caused by the accompanying gamma dose (100 kGy) [79]. Irradiations of tubes biased to the working voltage, with both gammas and neutrons showed no adverse effects, apart from an increase in anode current, attributable to the production of Cerenkov light in the faceplates.

The VPTs are operated with the photocathode at ground potential and the dynode and anode biased at +600 V and +800 V respectively. The high voltage system is based (like the APD system) on CAEN SY1527 standard control crates, although for the VPTs, the crates are equipped with standard 12-channel A1735P boards, each channel rated at 1.5 kV and 7 mA. At the operating bias, the VPT gain is close to saturation thus the voltages do not have to be controlled very precisely. However, care must be taken to minimise ripple and noise, since these would feed directly into the input of the sensitive preamplifier that is connected to the anode. Filtering is achieved

with RC networks mounted inside the supercrystals (SC), close to the VPTs. An entire endcap is biased using one SY1527 crate equipped with just two A1735P boards. On each board, only eight of the twelve output channels will initially be used, leaving four spare channels. The spare outputs may be used at a later stage, if noisy channels develop which can be recovered by operating at a lower bias voltage. The HV from the CAEN power supplies is transmitted to the SCs via a custom designed HV distribution system which provides hard-wired protection against over-voltage and over-current, and sensitive current monitoring. For each endcap, this system is housed in five crates. Each crate hosts up to five input cards, receiving the HV from the power supplies, and up to six output cards, with each output card serving up to twelve SCs. The HV supplies and distribution system are mounted in two racks (one for each endcap) located in the Service Cavern. Each SC is served by two coaxial cables (one for the anode, one for the dynode) running from the Service Cavern to the detector, via intermediate patch panels. The total cable length is approximately 120 m and the cable capacitance forms part of the filter network. Inside an SC the HV is distributed to the VPTs via five filter cards, each serving five VPTs. The spread in anode sensitivity among the VPTs is 25% (RMS). They are therefore sorted into six groups which are distributed on the endcaps with the highest sensitivities at the outer circumference grading to the lowest sensitivities at the inner circumference. This arrangement provides a roughly constant sensitivity to the transverse energy across the endcaps.

The anode sensitivity of a VPT may show a dependence on count rate (anode current) under certain conditions. For example, in the absence of a magnetic field, if the count rate falls to a few Hz, following a period of high rate operation, the anode sensitivity may rise suddenly and take several hours to return to the nominal value. The magnitude of the effect may vary from a few percent to a few tens of percent. In the presence of a strong magnetic field (as in normal CMS operation), the effect is strongly suppressed or absent. Nevertheless, it has been judged prudent to incorporate a light pulser system on the ECAL endcaps. This delivers a constant background rate of at least 100 Hz of pulses of approximately 50 GeV equivalent energy to all VPTs, thus ensuring that they are kept “active”, even in the absence of LHC interactions.

The system consists of a control and trigger unit located in the Service Cavern, and sets of pulsed light sources mounted on the circumference of each Dee. The light is produced by Luxeon III light emitting diodes (type LXHL-PR09), whose peak emission wavelength is 455 nm. The LEDs are driven by high output current op-amps (LT6300 from Linear Technology). The drive pulses have amplitudes of 1.2 A and a widths of 80 ns. A single light source consists of a cluster of seven LEDs and associated drive-circuits. These are configured singly or in pairs, with the drive-circuits and LEDs mounted on double-sided printed circuit boards housed within metal enclosures. There are four such enclosures distributed around the circumference of each Dee, housing 19 light sources. A schematic representation of the system for distributing the light pulses is shown in Fig. 74.

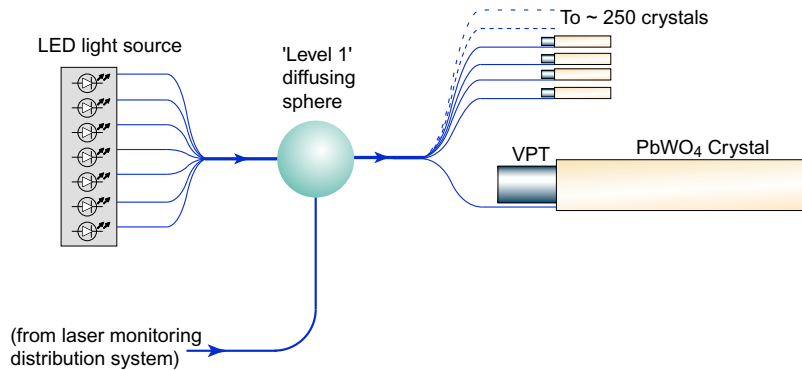


Figure 74: Distribution system for VPT stabilisation light pulses.

An all-silica optical fibre (CF01493-43 from OFS (Furukawa)) is inserted into a hole drilled into the lens of each LED and collects light by *proximity focusing*. The seven fibres from a given light source are combined into a single bundle that transports light to a diffusing sphere which has a dual role, acting also as part of the distribution network of the laser monitoring system. Light from each diffusing sphere is distributed to up to 220 individual detector channels through the set of optical fibres that also carry the laser monitoring pulses. Light is injected via the rear face of a crystal, which carries the VPT, and reaches the VPT via reflection from the front of the crystal.

The system is synchronised to pulse during a fraction of the $3\ \mu\text{s}$ abort gaps that occur during every $89\ \mu\text{s}$ cycle of the LHC circulating beams.

5.4 On-detector electronics

The ECAL read-out has to acquire the small signals of the photo-detectors with high speed and precision. Every bunch crossing digital sums representing the energy deposit in a trigger tower are generated and sent to the trigger system. The digitized data are stored during the Level-1 trigger latency of $\approx 3\ \mu\text{s}$.

The on-detector electronics has been designed to fully read a trigger tower (5×5 crystals in $\eta \times \phi$) or a super-crystal for EB and EE respectively. It consists of five Very Front End (VFE) boards, one Front End (FE) board, two (EB) or six (EE) Gigabit Optical Hybrids (GOH), one Low Voltage Regulator card (LVR) and a motherboard.

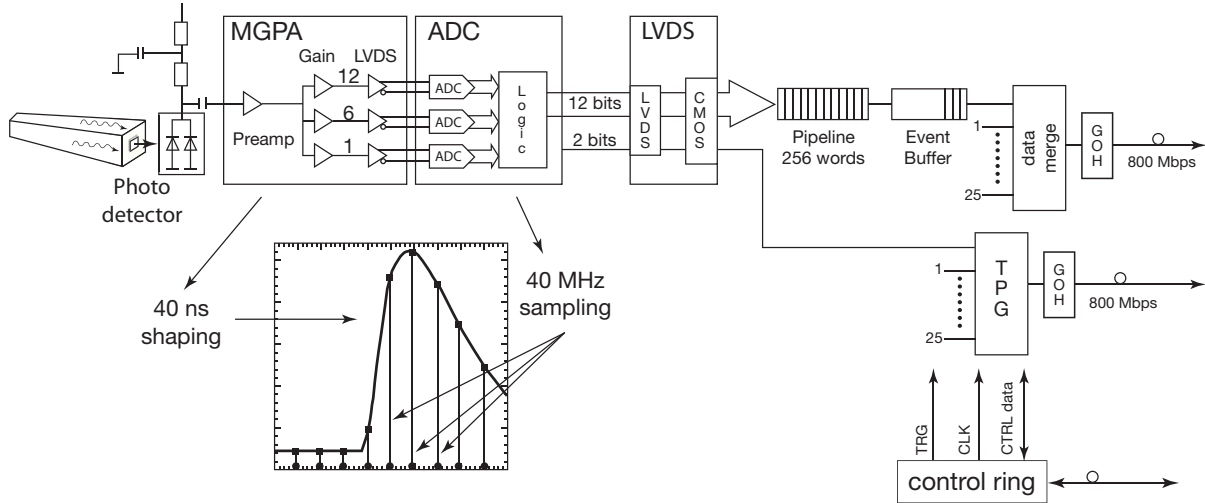


Figure 75: Schematic view of the on detector electronics: the scintillation light is collected by photodetectors (in the figure the case of APD is presented), the signal is shaped by a Multi-Gain Pre-Amplifier and digitized by 40 MHz ADC; a radiation hard buffer (LVDS) adapts the ADC output to the FE card, where data pipeline and Trigger Primitives Generation (TPG) are performed; trigger words are sent at 25 ns rate, while data are transmitted on receipt of a Level-1 trigger; GOHs provide in both cases the data serializer and the laser diode, sending the signals on a fibre to the off detector electronics over a distance of about 100 m. A control token ring connects groups of FE cards, providing Level-1 trigger (TRG) and clock (CLK) signals, together with control data in and out (CTRL data).

The motherboard is located in front of the cooling bars. It connects to 25 photo-detectors and to the temperature sensors using Kapton flexible printed circuit boards and coaxial cables for EB and EE respectively. In the case of the EB the motherboard distributes and filters the APD bias voltage. Two motherboards are connected to one CAEN HV supply located at a distance of about 120 m using remote sensing. In the case of the EE the operating voltages for the VPTs are distributed and filtered by a separate HV filter card, hosting as well the decoupling capacitor for the anode signals. Five of these cards serving five VPTs each are installed into each super-crystal. One LVR and five VFE cards plug into the motherboard.

Each LVR card [80] uses 11 radiation hard low voltage regulators (LHC4913) developed by ST-microelectronics and the RD49 project at CERN. The regulators have built in over-temperature protection, output current limitation and an inhibit input. The output voltages of 2.5 V are distributed to the FE card and via the motherboard to the VFE cards. Three Detector Control Unit (DCU) ASICs on each LVR card, interfaced to the FE card, monitor all input and output voltages. All regulators, excluding the one providing power to the control interface of the FE card, can be powered down remotely by an external inhibit. Four LVR cards are connected by a passive low voltage distribution (LVD) block to one radiation and magnetic field tolerant Wiener low voltage power supply located about 30 m away in racks attached to the magnet yoke.

The signals are pre-amplified and shaped and then amplified by three amplifiers with nominal gains of 1, 6 and

12. This functionality is built into the Multi Gain Pre-Amplifier (MGPA) [81], an ASIC developed in $0.25\ \mu\text{m}$ technology. The full scale signals of the APDs and VPTs are 60 pC and 12.8 pC corresponding to $\approx 1.5\ \text{TeV}$ and 1.6–3.1 TeV for EB and EE respectively. The shaping is done by a CR-RC network with a shaping time of $\approx 40\ \text{ns}$. The MGPA has a power consumption of 580 mW at 2.5 V. The output pulse non-linearity is less than 1%. The noise for gain 12 is about $8000\ e^-$ for the APD configuration and about $4000\ e^-$ for the VPT configuration. The MGPA contains three programmable 8-bit DACs to adjust the baseline to the ADC inputs. An integrated test-pulse generator with an amplitude adjustable by means of an 8-bit DAC allows a test of the read-out electronics over the full dynamic range.

The 3 analog output signals of the MGPA are digitized in parallel by a multi-channel, 40 MHz, 12-bit ADC, the AD41240 [82], developed in $0.25\ \mu\text{m}$ technology. It has an effective number of bits of 10.9. An integrated logic selects the highest non-saturated signal as output and reports the 12 bits of the corresponding ADC together with two bits coding the ADC number.

If the read-out switches to a lower gain as the pulse grows, it is prevented from immediately reverting to the higher gain when the pulse falls: once the pulse has declined to the point where it could be read out at the higher gain again, the read-out is then forced to continue reading out at the lower gain for the next five samples.

A radiation hard buffer (LVDS.RX) developed in $0.25\ \mu\text{m}$ technology, adapts the low voltage differential output signals of the AD41240 to the single ended CMOS inputs on the FE card. Five identical read-out channels are integrated into a VFE card, together with a Detector Control Unit (DCU) for the measurement of the APD leakage currents and the read-out of the thermistors. The noise obtained with the VFE cards installed into super-modules is typically 1.1, 0.75 and 0.6 ADC counts for gains 12, 6 and 1 respectively. This corresponds to $\approx 40\ \text{MeV}$ for gain 12.

The FE card [83] stores the digitized data during the Level-1 trigger latency in 256 word deep dual ported memories, so called pipelines. Five such pipelines and the logic to calculate the energy sum of the 5 channels once every bunch crossing are integrated into an ASIC developed in $0.25\ \mu\text{m}$ technology called FENIX. Each VFE card is serviced by a FENIX chip. Thus the energy is summed in strips of 5 crystals along phi. In the case of the EE the five strip sums are transmitted by five GOHs (see below) to the off-detector electronics Trigger Concentrator Card (TTC), while in the case of the EB a sixth FENIX sums the five strip sums and calculates the “fine grain” electromagnetic bit, set to identify electromagnetic shower candidates on the basis of the energy profile of the trigger tower. The trigger tower energy sum together with the fine grain bit is transmitted using one GOH to the TCC. On receipt of a Level-1 trigger the corresponding data, ten 40-MHz samples per channel, are transmitted in $\approx 7.5\ \mu\text{s}$ to the off-detector electronics Data Concentrator Card (DCC) using an identical GOH. The Clock and Control Unit (CCU) ASIC together with the LVDS.MUX ASIC provide the interface to the token rings.

The ECAL serial digital data links are based on the technology developed for the CMS Tracker analog links (Sect. 4.3). The GOH consists of a data serializer and laser driver chip, the GOL, and a laser diode with an attached fiber pigtail. Fibers, fiber interconnections and a 12-channel NGK receiver module complete the optical link system. It uses single mode fibers operating at 1310 nm wavelength over a distance of about 100 m. The fiber attenuation of $\approx 0.04\ \text{dB}$ is negligible. The optical links are operated at 800 Mbit/s.

The VFE and FE electronics are controlled using a 40-MHz digital optical link system, controlled by the off-detector Clock and Control System (CCS) boards. A 12-fiber ribbon is connected to the token ring link board, generating an electrical control ring, the token ring. Each super module has 8 token rings which connect to groups of eight to ten FE cards including the two FE cards of the laser monitoring electronics module (MEM). The system has redundancy, as long as there are no two consecutive FE cards malfunctioning, by means of two independent bi-directional optical links, using 4 fibers each. It provides fast and slow control functions. While the fast control transmits the level one trigger information and the 40-MHz clock, the slow control comprises the configuration of the FE and VFE electronics as well as the read-out of status information, temperatures, voltages and APD leakage currents.

5.5 Off-detector electronics

5.5.1 Global architecture

The ECAL off-detector read-out and trigger architecture [84, 85] is illustrated schematically in Fig. 76. The system is composed of different electronic boards sitting in 18 VME-9U crates (the CCS, TCC and DCC modules) and in 1 VME-6U crate (the selective read-out processor, SRP, system). The system serves both the DAQ and the trigger paths. In the DAQ path, the DCC performs data read-out and data reduction based on the selective read-out flags

computed by the SRP system. In the trigger path, at each bunch crossing, trigger primitive generation started in the FE boards is finalized and synchronized in the TCC before transmission to the regional calorimeter trigger.

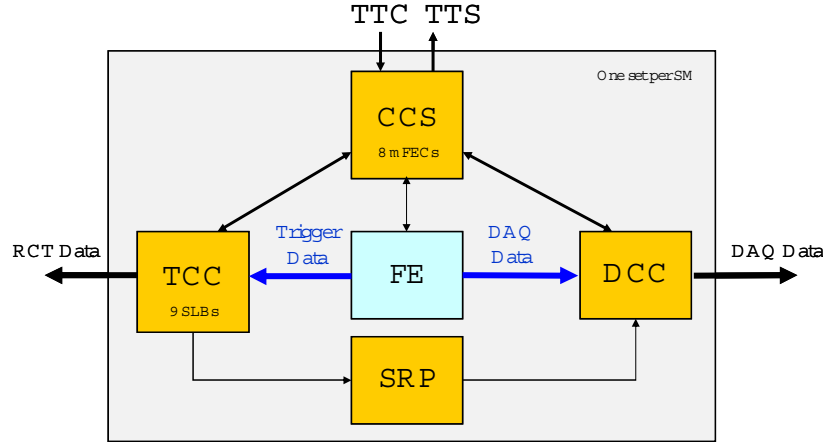


Figure 76: Schematic view of ECAL off-detector electronics.

The clock and control system (CCS) board distributes the system clock, trigger and broadcast commands, configures the FE electronics and provides an interface to the trigger throttling system. The TTC signals are translated and encoded by suppression of clock edges and sent to the *mezzanine* Front End Controller cards (mFEC). The mFEC interfaces optically with a FE token ring. The 8 mFECs of the CCS board control a supermodule.

The trigger concentration card (TCC) [86] main functionalities include the completion of the trigger primitive generation and their transmission to the synchronization and link board (SLB) mezzanines [87] at each bunch crossing, the classification of each trigger tower and its transmission to the Selective Read-out Processor at each Level-1 trigger accept signal, and the storage of the trigger primitives during the Level-1 latency for subsequent reading by the DCC.

Each TCC collects trigger data from 68 FE boards in the barrel, corresponding to a supermodule, and from 48 FE boards in the endcaps corresponding to the inner or outer part of a 20° ϕ sector. In the endcaps, trigger primitive computation is completed in the TCCs, which must perform a mapping between the collected pseudo-strips trigger data from the different supercrystals and the associated trigger towers. The encoded trigger primitives (8 bits for the nonlinear representation of the trigger tower E_T plus the fine grain bit) are time aligned and sent to the regional trigger processors by the SLB. The trigger primitives are stored in the TCC during the Level-1 latency for subsequent reading by the DCC. In the barrel region a single TCC is interfaced with 1 DCC. In the endcap region, a DCC serves 4 TCCs covering a 40° sector.

The data concentration card (DCC) [88, 89] is responsible for collecting crystal data from up to 68 FE boards. Two extra FE links are dedicated to the read-out of laser monitoring data (PN diodes). The DCC also collects trigger data transmitted from the TCC modules and the selective read-out flags transmitted from the SRP system. A data suppression factor near 20 is attained using a programmable selective read-out algorithm. When operating in the selective read-out mode the SRP flags indicate the level of suppression that must be applied to the crystal data of a given FE read-out. For the application of zero suppression, time samples pass through a finite impulse response filter with 6 consecutive positions and the result is compared to a threshold. If any time sample of the 6 has been digitized at a gain other than the maximum, then zero suppression is not applied to the channel.

Data integrity is checked, including verification of the event-fragment header, in particular the data synchronization check, verification of the event-fragment word count and verification of the event-fragment parity bits. Identified error conditions, triggered by input event-fragment checks, link errors, data timeouts or buffer memory overflows are flagged in the DCC error registers and incremented in associated error counters. Error conditions are flagged in the DCC event header.

Input and output memory occupancy is monitored to prevent buffer overflows. If a first occupancy level is reached, the Trigger Throttling System (TTS) signal *Warning Overflow* is issued, requesting a reduction of the trigger rate. In a second level a TTS signal *Busy* inhibits new triggers and empty events (events with just the header words and trailer) are stored. DCC events are transmitted to the central CMS DAQ using a S-LINK64 data link interface at

a maximum data rate of 528 MB/s, while an average transmission data flow of 200 MB/s is expected after ECAL data reduction. Laser triggers (for crystal transparency monitoring) will occur with a programmable frequency and synchronously with the LHC gap. No data reduction is applied for these events, which are read-out following a TTC test enable command. A VME memory is used for local DAQ, allowing VME access to physics events and laser events in spy mode.

The selective read-out processor (SRP) [90] is responsible for the implementation of the selective read-out algorithm. The system is composed by a single 6U-VME crate with twelve identical algorithm boards (AB). The AB computes the selective read-out flags in different calorimeter partitions. The flags are composed of 3 bits, indicating the suppression level that must be applied to the associated read-out units.

5.5.2 The trigger and read-out paths

The ECAL data, in the form of trigger primitives, are sent to the Level-1 calorimeter trigger processor, for each bunch crossing. The trigger primitives each refer to a single trigger tower and consist of the summed transverse energy deposited in the tower, and the fine grain bit, which characterizes the lateral profile of the electromagnetic shower. The accept signal, for accepted events, is returned from the global trigger in about $3\mu s$. The selected events are read out through the data acquisition system to the Filter Farm where further rate reduction is performed using the full detector data.

The read-out system is structured in sets of 5×5 crystals. The FE card stores the data, in 256-clock cycles deep memory banks, awaiting a Level-1 trigger decision during at most 128 bunch crossings after the collision occurred. It implements most of the Trigger Primitives Generation (TPG) pipeline (Sect. 5.5.3).

In the barrel, these 5×5 crystal sets correspond to the trigger towers. Each trigger tower is divided into 5 ϕ -oriented strips, whose energy deposits are summed by the FE board trigger pipeline to give the total transverse energy of the tower, called the main trigger primitive. Each FE is served by two optical links for sending the data and trigger primitives respectively and a third electrical serial link which transmits the clock, control and Level-1 trigger signals.

In the endcaps, the read-out modularity maps onto the 5×5 mechanical units (supercrystals). However the sizes of the trigger towers vary in order to approximately follow the η , ϕ geometry of the HCAL and Level-1 trigger processor. The supercrystals are divided into groups of 5 contiguous crystals. These groups are of variable shape and referred to as pseudo-strips. The trigger towers are composed of several pseudo-strips and may extend over more than one supercrystal. Since the read-out structure does not match the trigger structure, only the pseudo-strip summations are performed on the detector. The total transverse energy of the trigger tower is computed by the off-detector electronics. Hence, each endcap FE board is served by 6 optical links, 5 of them being used to transmit the trigger primitives. As in the barrel an electrical serial link transmits the clock, control and Level-1 trigger signals.

After time alignment the ECAL trigger primitives are sent at 1.2 Gb/s to the regional calorimeter trigger, via 10-m-long electrical cables, where together with HCAL trigger primitives, the electron/photon and jets candidates are computed as well as the total transverse energy.

5.5.3 Algorithms performed by the trigger primitive generation

The TPG logic implemented on the FE boards combines the digitized samples delivered by the VFE boards to determine the trigger primitives and the bunch crossing to which they should be assigned. The logic must reconstruct the signal amplitude to be assigned to each bunch-crossing from the continuous stream of successive digitizations.

The TPG logic is implemented as a pipeline, operated at the LHC bunch crossing frequency. The trigger primitives are delivered to the regional calorimeter trigger after a constant latency of 52 clock cycles, of which 22 are used for transmission over the optical fibres and cables. The signal processing performed in the VFE and FE barrel electronics has a total duration of only 17 clock cycles. The remaining part of the latency is mainly due to formatting and time alignment of the digital signals. Ideally, the output of this processing should be a stream of zeroes, unless there is a signal in the tower resulting from a bunch crossing exactly 17 clock cycles before. In this case the output is a word encoding the summed transverse energy in the tower together with the fine grain bit. The endcap pipeline is split between the on-detector and off-detector electronics and implements very similar algorithms. The trigger primitives are expected to be delivered to the regional calorimeter trigger in 50 clock cycles in the endcap case.

5.5.4 Classification performed by the selective read-out

About 100 kB per event has been allocated for ECAL data. The full ECAL data for an event, if all channels are read out, exceeds this target by a factor of nearly 20. Reduction of the data volume, *selective read-out*, can be performed by the Selective Read-out Processor [85, 90] so that the suppression applied to a channel takes account of energy deposits in the vicinity. For the measure of the energy in a region, the trigger tower sums are used. In the barrel the read-out modularity corresponds exactly to the 5×5 -crystal trigger towers. In the endcap, the situation is more complex. The simplified and illustrative description below is given for the barrel case.

The selective read-out algorithm classifies the trigger towers of the ECAL into 3 classes using the Level-1 trigger primitives. The energy deposited in each trigger tower is compared to 2 thresholds. Trigger towers with an energy above the higher threshold are classified as high interest trigger towers, those with an energy between the 2 thresholds as medium interest, and those with an energy below the lower threshold as low interest trigger towers.

These classifications can be used flexibly to implement a range of algorithms by using different thresholds to define the classes, and different suppression levels for the read-out of the channels within each class. The algorithm currently used in the simulation provides adequate data reduction even at high luminosity. The algorithm functions as follows: if a trigger tower belongs to the high interest class ($E_T > 5$ GeV) then the crystals of this trigger tower and of its neighbour trigger towers (225 crystals in total in the barrel case) are read with no zero suppression. If a trigger tower belongs to the medium interest class ($E_T > 2.5$ GeV), then the crystals of this trigger tower (25 crystals in the barrel case) are read with no suppression. If a trigger tower belongs to the low interest class and it is not the neighbour of a high interest trigger tower, then the crystals in it are read with zero suppression at about $3\sigma_{\text{noise}}$.

For debugging purposes, the selective read-out can be deactivated and either a global zero suppression (same threshold for every channel) or no zero suppression applied. Even when the selective read-out is not applied the selective read-out flags are inserted into the data stream and can be used offline for debugging purposes.

5.6 Preshower detector

The principal aim of the CMS Preshower detector is to identify neutral pions in the endcaps within a fiducial region $1.653 < |\eta| < 2.6$. It also helps the identification of electrons against minimum ionizing particles, and improves the position determination of electrons and photons with high granularity.

5.6.1 Geometry

The Preshower is a sampling calorimeter with two layers: lead radiators initiate electromagnetic showers from incoming photons/electrons whilst silicon strip sensors placed after each radiator measure the deposited energy and the transverse shower profiles. The total thickness of the Preshower is 20 cm.

The material thickness of the Preshower traversed at $\eta = 1.653$ before reaching the first sensor plane is $2 X_0$, followed by a further $1 X_0$ before reaching the second plane. Thus about 95% of single incident photons start showering before the second sensor plane. The orientation of the strips in the two planes is orthogonal. A major design consideration is that all lead is covered by silicon sensors, including the effects of shower spread, primary vertex spread etc. For optimum Level-1 trigger performance the profile of the outer edge of the lead should follow the shape of the ECAL crystals behind it. For the inner radius the effect of the exact profiling of the lead is far less critical, and thus a circular shape has been chosen. The lead planes are arranged in two Dees, one on each side of the beam pipe, with the same orientation as the crystal Dees.

Each silicon sensor measures $63 \times 63 \text{ mm}^2$, with an active area of $61 \times 61 \text{ mm}^2$ divided into 32 strips (1.9 mm pitch). The nominal thickness of the silicon is $320 \mu\text{m}$; a minimum ionizing particle (MIP) will deposit 3.6 fC of charge in this thickness (at normal incidence). The sensors are precisely glued to ceramic supports, which also support the front-end electronics assembly (see below), and this is in turn glued to an aluminium tile that allows a 2 mm overlap of the active part of the sensors in the direction parallel to the strips. In order to improve noise performance the tile is constructed in two parts, with a glass fibre insulation in between. The combination of sensor + front-end electronics + supports is known as a micromodule.

The micromodules are placed on baseplates in groups of 7, 8 or 10 that, when coupled to an electronics system motherboard (SMB) placed above the micromodules, form a ladder. The spacing between silicon strips (at the edges) in adjacent micromodules within a ladder is 2.4 mm, whilst the spacing between strips in adjacent ladders is normally 2.5 mm. For the region where the two Dees join this spacing is increased to 3.0 mm.

Figure 77 shows a complete ladder (*Type-1* for 8 micromodules) and an inset shows the micromodule.

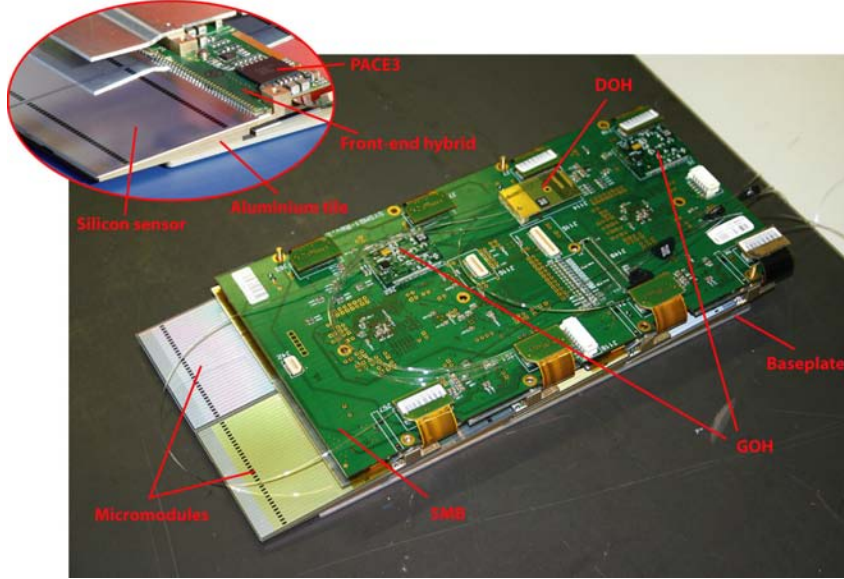


Figure 77: Photograph of a complete type-1 ladder, with an inset showing details of a micromodule

The ladders are attached to the radiators in an x - y configuration. Around 500 ladders are required, corresponding to a total of around 4 300 micromodules and 137 000 individual read-out channels. Further details of the layout can be found in [91].

5.6.2 Preshower electronics

Each silicon sensor is DC-coupled to a front-end ASIC (PACE3 [92]) that performs preamplification, signal shaping and voltage sampling. Data is clocked into an on-chip high dynamic range 192-cell deep analogue memory at 40 MHz.

For each Level-1 trigger received, 3 consecutive cells of the memory, corresponding to time samples on the baseline, near the peak and after the peak, are read-out for all 32 channels through a 20 MHz multiplexer. The PACE3 has a switchable gain:

- Low gain: For normal physics running with a large dynamic range (0-1600 fC) with a S/N of around 3 for a single MIP;
- High gain: For MIP calibration purposes [93], with a reduced dynamic range (0-200 fC) but with a S/N approaching 10 for a single MIP.

The PACE3 are soldered to front-end hybrids that contain embedded polyimide cables to connect to the SMBs. The SMBs contain AD41240 12-bit ADCs that digitize the data from 1 or 2 PACE3. The digital data are then formatted and packaged by a second Preshower ASIC known as the K-chip [94]. The K-chip also performs synchronization checks on the data, adds bunch/event counter information to the data packets and transmits the data to the Preshower-DCC (see below) via gigabit optical hybrids (GOH). The SMB also contains an implementation of the CMS tracker control system.

Groups of up to 12 ladders are connected via polyimide cables to form control rings. Off-detector CCS cards (identical to those of the ECAL except not all FEC mezzanines are mounted for the Preshower) communicate via digital optical hybrids (DOH) mounted on 2 of the SMBs in each control ring. The full Preshower comprises 4 planes of 12 control rings each.

The Preshower-DCC [95] is based on the DCC of the ECAL except it is a modular design incorporating a VME host board mounted with optoRx12 [96] mezzanines. The modular design has allowed a development collaboration with the TOTEM experiment which uses the same components but in a different manner. The optoRx12 incorporates an NGK 12-way optical receiver and an Altera Stratix GX FPGA that performs data deserialization,

pedestal subtraction, common-mode noise reduction, bunch-crossing assignment, charge reconstruction and zero suppression [97]. The sparsified data from up to 3 optoRx12 are merged by another FPGA on the host board that then transmits data packets to the event builder via an Slink64 interface. The host board also provides data spying as well as TTC and VME interfaces. A provision has been made on the host board to allow the plug-in of an additional mezzanine board mounted with FPGAs/processors that could provide more data reduction power if necessary in the future.

5.7 ECAL detector control system

The ECAL Detector Control System (DCS) comprises the monitoring of the detector status, in particular various kinds of environmental parameters, as well as the ECAL safety system (ESS), which will generate alarms and hardwired interlocks in case of situations which could lead to damaging the detector hardware. It consists of the following sub-systems: ECAL Safety System (ESS), Precision Temperature Monitoring (PTM), Humidity Monitoring (HM), High Voltage (HV), Low Voltage (LV) and monitoring of the laser operation, the cooling system and of the parameters (temperatures in capsules, temperatures on the printed circuit boards, APD leakage currents) read out by the DCUs on the VFE and LVR boards. Further details on the ECAL DCS are available [98].

The whole DCS software is based on the commercial SCADA package PVSS II (Sect. 10). A distributed system is built out of several applications dedicated to the DCS sub-systems. Every application is implemented as a Finite State Machine (FSM) and linked to a supervisory level, which summarizes the overall ECAL DCS status and itself incorporates a FSM. Finally, this ECAL DCS supervisor is linked to the general CMS DCS supervisory node, in order to communicate the status and alarms and to receive commands which are propagated down to the relevant sub-systems.

5.7.1 Safety system

The purpose of the ESS [99] is to monitor the air temperature of the VFE and FE environment (expected to be around 25–30 °C) and the water leakage detection cable, which is routed inside the electronics compartment, to control the proper functioning of the cooling system and to automatically perform pre-defined safety actions and generate interlocks in case of any alarm situation. One pair of temperature sensors is placed at the centre of each module. The read-out system, with full built-in redundancy, is independent of the DAQ and control links and based on a Programmable Logic Controller (PLC) situated in the Service Cavern. In case of any critical reading hardwired interlock signals will be routed to the relevant crates in order to switch off the HV and LV and/or the cooling PLC in order to stop the water flow on a certain cooling line. The proper functioning of the ESS PLC itself is monitored by the general CMS detector safety system.

5.7.2 Temperature

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent, as described in Sect. 5.2. Therefore a major task for the ECAL DCS is the monitoring of the system's temperature and the verification that the required temperature stability of (18 ± 0.05) °C of the crystal volume and the APDs is achieved. The PTM is designed to read out thermistors, placed on both sides of the crystal volume, with a relative precision better than 0.01 °C. In total there are ten sensors per supermodule. Two immersion probes measure the temperature of the incoming and outgoing cooling water, whereas two sensors per module, one on the grid and one on the thermal screen side of the crystal volume, monitor the crystal's temperature. The read-out is based on the Embedded Local Monitoring Board (ELMB) developed by ATLAS which functions completely independently of the DAQ and control links. In addition, sensors fixed to the back surface of every tenth crystal in the barrel, and one in 25 crystals in the endcap, are read out by the DCUs placed on the VFE boards. With this temperature monitoring it has been shown that the water cooling system can indeed ensure the required temperature stability [73].

5.7.3 Dark current

The APD dark current will increase during CMS operation due to bulk damage of the silicon structure by neutrons. Part of this damage anneals, but the overall effect will be an increase in electronics noise, due to an increasing dark current, over the lifetime of the detector. The dark current of all APD channels will be continuously monitored.

5.7.4 HV and LV

The DCS system operates the CAEN HV system via an OPC server. The functionalities include the independent configuration of the HV channels with various set of voltages, the monitor of the voltage and the current delivered by each channel and the database recording of the settings. The ECAL Safety System can switch off the HV via the individual board interlocks.

The ECAL amplification and digitization electronics located on the VFE electronics cards require a very stable low voltage to maintain constant signal amplification. The system uses Low Voltage Regulator that guarantee the required stability of the signal amplification. The Low Voltage Regulators Boards are equipped with DCUs that measure the voltages and these measurements are read via the Token Ring. Overall the power is supplied by MARATON crates (WIENER), which are operated and monitored by the DCS.

5.8 Detector calibration

Calibration is a severe technical challenge for the operation of the CMS ECAL. Many small effects which are negligible at low precision need to be treated with care as the level of precision of a few per mille is approached. ECAL calibration is naturally seen as composed of a global component, giving the absolute energy scale, and a channel-to-channel relative component, which is referred to as intercalibration. The essential issues are uniformity over the whole ECAL and stability, so that showers in different locations in the ECAL in data recorded at different times are accurately related to each other.

The main source of channel-to-channel response variation in the barrel is the crystal-to-crystal variation of scintillation light yield which has an RMS of $\approx 8\%$ within most supermodules, although the total variation among all barrel crystals is $\approx 15\%$. In the endcap the VPT signal yield, the product of the gain, quantum efficiency and photocathode area, has an RMS variation of almost 25%. Preliminary estimates of the intercalibration coefficients are obtained from laboratory measurements of crystal light yield and photodetector/electronics response [100]. Applying this information reduces the channel-to-channel variation to less than 5% in the barrel and less than 10% in the endcaps.

All 36 supermodules were commissioned in turn by operating them on a cosmic ray stand for a period of about one week. A muon traversing the full length of a crystal deposits an energy of approximately 250 MeV, permitting intercalibration information to be obtained for the barrel ECAL [101]. In 2006, nine supermodules were intercalibrated with high energy electrons (90 and 120 GeV), in a geometrical configuration that reproduced the incidence of particles during CMS operation. One of the supermodules was exposed to the beam on two occasions, separated by an interval of one month. The resulting sets of inter-calibration coefficients are in close agreement, the distribution of differences having an RMS spread of 0.27%, indicating a reproducibility within the statistical precision of the individual measurements (Fig. 78).

A comparison of the cosmic ray and high energy electron data demonstrates that the precision of the cosmic ray inter-calibration is better than 1.5% over most of the volume of a supermodule, rising to just above 2% at the outer end (corresponding to $\eta \approx 1.5$). The mean value of the precision of the cosmic intercalibration, averaged over all the channels in the nine supermodules for which a comparison with electrons can be made, is 1.5% (Fig. 79).

The ultimate intercalibration precision will be achieved *in situ* with physics events. As a first step, imposing the ϕ -independence of the energy deposited in the calorimeter can be used to rapidly confirm, and possibly improve on, the start-up intercalibration within fixed η regions. The intercalibration method that has been investigated in the most detail uses the momentum of isolated electrons measured in the tracker. These electrons, mainly from $W \rightarrow e\nu$, are abundant ($\sigma \approx 20$ nb) and have a similar p_T to the photons of the benchmark channel $H \rightarrow \gamma\gamma$. A complementary method, not relying on the tracker momentum measurement, is based on $\pi^0 \rightarrow \gamma\gamma$ and $\eta \rightarrow \gamma\gamma$ mass reconstruction. Most methods of intercalibration will be local to a region of the ECAL, and a further step intercalibrating these regions to one another will be needed. This is a consequence of the significant systematic variations that occur as a function of pseudorapidity such as (or including): the large variation of the thickness of the tracker material, the variation of the structure of the ECAL (both the major differences between the barrel and endcap, and the small continuous variation of the geometry along the length of the barrel), and the variation of background characteristics for $\pi^0 \rightarrow \gamma\gamma$.

Over the period of time in which the physics events used to provide an intercalibration are taken the response must remain stable to high precision. Where there is a source of significant variation, as in the case of the changes in crystal transparency caused by irradiation and subsequent annealing, the variation must be precisely tracked by an independent measurement. The changes in crystal transparency are tracked and corrected using the laser

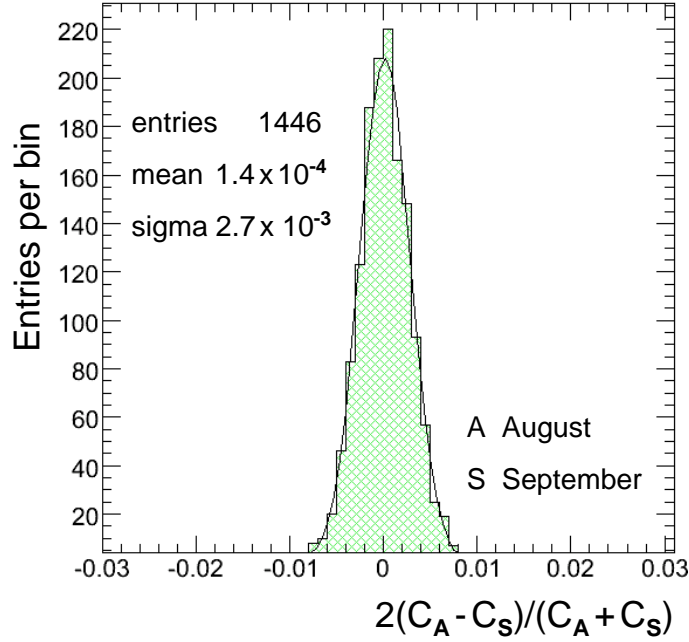


Figure 78: Distribution of differences of inter-calibration coefficients from a supermodule exposed to a high energy electron beam on two occasions, separated by a period of one month. The reproducibility of the intercalibration coefficients ($\text{RMS}/\sqrt{2}$) is measured to be 0.2%.

monitoring system.

The final goal of calibration is to achieve the most accurate energy measurements for electrons and photons. Different reconstruction algorithms are used to estimate the energy of different electromagnetic objects, i.e., unconverted photons, electrons and converted photons, each of them having their own correction functions. At present these “algorithmic” corrections are obtained from the simulated data by accessing the generated parameters of the Monte Carlo simulation. For some of the corrections, for example the containment corrections, this is an acceptable procedure provided that test beam data is used to verify the simulation, so that, in effect, the simulation is being used only as a means of interpolating and extrapolating from data taken in the test beam. In other cases, where the test beam provides no useful information, for example in issues related to conversions and bremsstrahlung radiation in the tracker material, it will be important to find ways of using information that can be obtained from data taken in situ with the running detector. Two particularly useful channels which can be used to obtain such information, and also assist in the step of intercalibrating regions of the ECAL to one another, are under investigation: $Z \rightarrow e^+e^-$, and $Z \rightarrow \mu^+\mu^-\gamma$ (the photon coming from inner bremsstrahlung).

5.9 Laser monitor system

Although radiation resistant, ECAL PbWO_4 crystals show a limited but rapid loss of optical transmission under irradiation due to the production of colour centres which absorb a fraction of the transmitted light. At the ECAL working temperature (18 °C) the damage anneals and the balance between damage and annealing results in a dose-rate dependent equilibrium of the optical transmission, if the dose rate is constant. In the varying conditions of LHC running the result is a cyclic transparency behaviour between LHC collision runs and machine refills (Fig. 80). The magnitude of the changes is dose-rate dependent, and is expected to range from 1 or 2 per cent at low luminosity in the barrel, to tens of per cent in the high η regions of the endcap at high luminosity. The performance of the calorimeter would be unacceptably degraded by these radiation induced transparency changes were they not measured and corrected for.

The evolution of the crystal transparency is measured using laser pulses injected into the crystals via optical fibres. The response is normalized by the laser pulse magnitude measured using silicon PN photodiodes. PN type photodiodes were chosen because of their very narrow depletion zone ($\approx 7 \mu\text{m}$ with +4 V reverse bias), making them much less sensitive to type inversion than the faster PIN photodiodes. Thus $R(t) = \text{APD}(t)/\text{PN}(t)$ is used as the measure of the crystal transparency. The laser monitoring system [68] performing this task is briefly outlined in

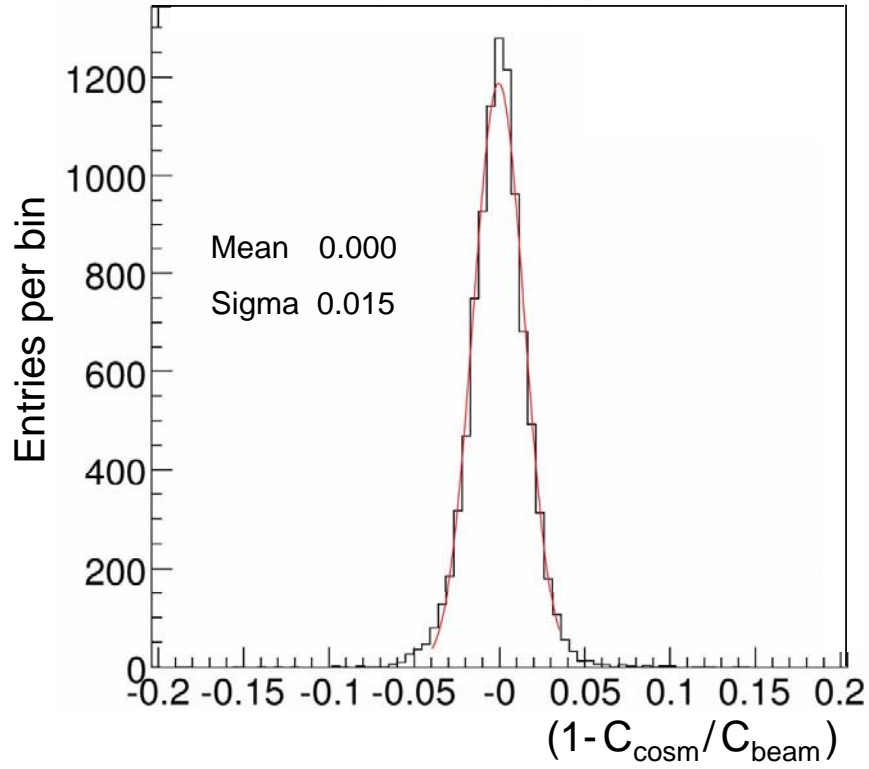


Figure 79: Distribution of the relative differences between the inter-calibration coefficients measured with high energy electrons and those obtained from cosmic ray muons.

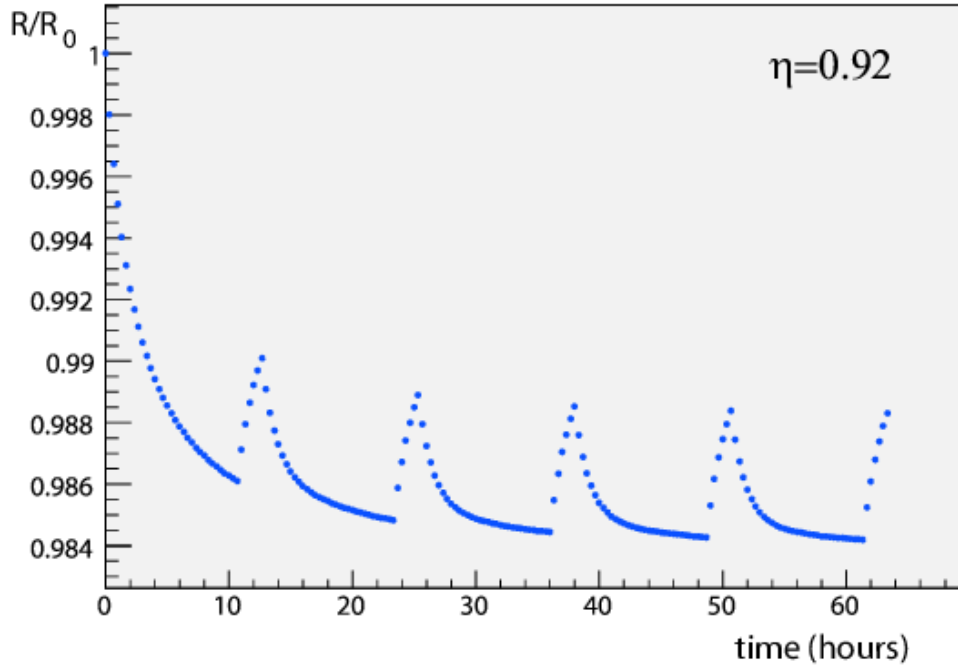


Figure 80: Simulation of crystal transparency evolution at LHC based on test-beam results. For this illustrative example a luminosity of $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ was assumed, together with a machine cycle consisting of a 10 hour coast followed by 2 hours filling time. The crystal behaviour under irradiation was modeled on data taken during a crystal irradiation in the test beam.

the next Section. Because of the different optical paths and spectra of the injected laser pulses and the scintillation light, the changes in crystal transparency cause a change in response to the laser light which is not necessarily equal to the change in response to scintillation light. For attenuations $< 10\%$ the relationship between the changes can be expressed by a power law,

$$S(t)/S(t_0) = [R(t)/R(t_0)]^\alpha, \quad (2)$$

where $S(t)$ represents the response to scintillation light and α is characteristic of the crystal which depends on the production method ($\alpha \approx 1.53$ for BCTP crystals, and $\alpha \approx 1.0$ for SIC crystals). An example of this relationship is given in Fig. 81. This power law describes well the behaviour of all the crystals that have been evaluated in the test beam, and this formula is expected to be valid in the barrel for both low and high luminosity at LHC.

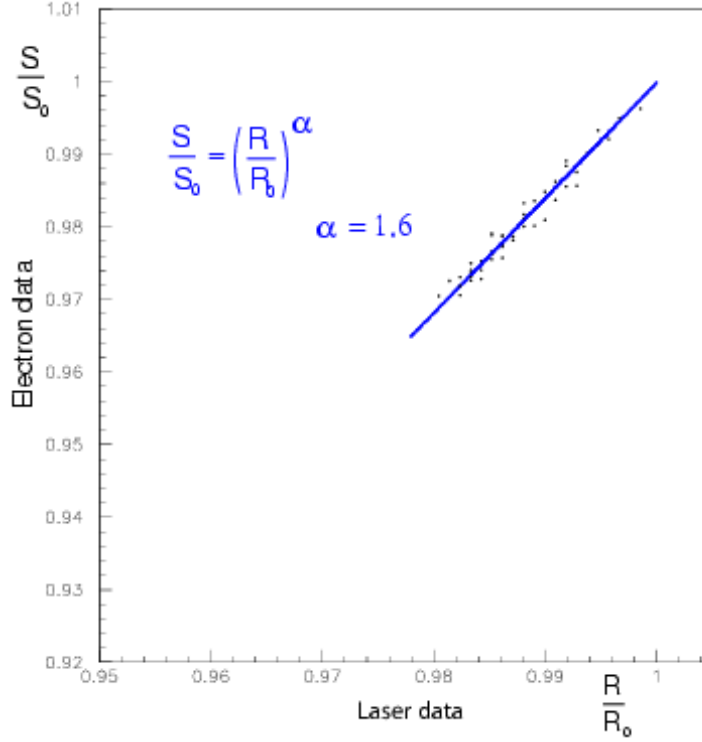


Figure 81: Relation between the transmission losses for scintillation light and for laser light for a given crystal. The signals are followed during the irradiation and the recovery.

5.9.1 Laser-monitoring system overview

Figure 82 shows the basic components of the laser-monitoring system: two laser wavelengths are used for the basic source. One, blue, at $\lambda=440$ nm, is very close to the scintillation emission peak, which is used to follow the changes in transparency due to radiation, and the other, infra-red, at $\lambda=796$ nm, is far from the emission peak, hence very little affected by changes in transparency, which can be used to verify the stability of other elements in the system. The spectral contamination is less than 10^{-3} . The lasers are operated such that the full width at half maximum of the pulses is ≈ 30 ns. The lasers can be pulsed at a rate of ≈ 80 Hz, and the pulse timing jitter is less than 3 ns which allows adequate trigger synchronization with the LHC bunch train and ECAL ADC clock.

The pulse energy of 1 mJ/pulse at the principal monitoring wavelength corresponds to ≈ 1.3 TeV, and a linear attenuator allows 1% steps down to 13 GeV. The pulse intensity instability is less than 10% which guarantees a monitoring precision of 0.1% by using the PN silicon photodiode normalization.

There are 3 light sources, 2 blue and 1 infrared. The duplication of the blue source provides fault tolerance and allows maintenance of one while the other is in use, ensuring that a source at the wavelength used to track changes in transparency is always available. Each source consists of an Nd:YLF pump laser, its power supply and cooler unit and corresponding transformer, a Ti:Sapphire laser and its controller, and a NESLAB cooler for an LBO crystal in the Ti:S laser. Each pair of the YLF and Ti:S lasers and their corresponding optics are mounted on an optical table. Each source has its own diagnostics, 2 fibre-optic switches, internal monitors and corresponding PC

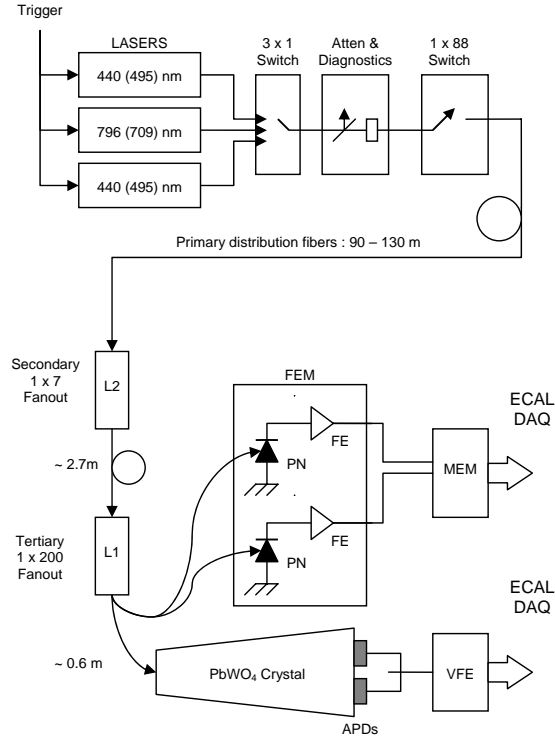


Figure 82: The components of the laser monitoring system.

based controllers. Further details can be found in [102].

The monitoring light pulses are distributed via a system of optical fibres. A fibre optic switch at the laser directs the laser pulses to 1 of 88 calorimeter regions (72 half supermodules in the barrel and 8 regions in each endcap). A two-stage distribution system mounted on each calorimeter region delivers the light to each crystal.

To provide continuous monitoring, about 1% of the $3.17 \mu\text{s}$ beam gap in every $88.924 \mu\text{s}$ LHC beam cycle will be used to inject monitoring light pulses into crystals. The time needed to scan the entire ECAL is expected to be about 30 minutes.

The first laser system was installed in the CERN H4 test beam site in August 2001. The other two laser systems were installed at H4 in August, 2003. All three laser systems have been used in the ECAL test beam program since their installation, and more than 10 000 laser hours have been cumulated.

The relative stability between a pair of reference PN photodiodes monitoring the same group of 200 crystals is shown in Fig. 83. The system achieves 0.0074% RMS over 7.5 days operation.

The response to injected laser light (normalized by the reference PN photodiodes) is presented in Fig. 84 for a group of 200 crystals measured for 11.5 days at the wavelength of 440 nm, showing that a stability of 0.068% is achieved at the scintillation wavelength.

The effect of the monitor correction procedure is presented in Fig. 85, showing that electron signals taken during an irradiation test at H4 are effectively corrected using laser monitor runs taken during the same data-taking period, providing an equalisation of the corrected response at the level of few per mille [103].

5.10 Energy resolution

For energies below about 500 GeV, where shower leakage from the rear of the calorimeter starts to become significant, the energy resolution can be parametrized as in Eq. 1 (Sect. 2), that is repeated for convenience here:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \quad (3)$$

where S is the stochastic term, N the noise term, and C the constant term. The individual contributions are discussed below.

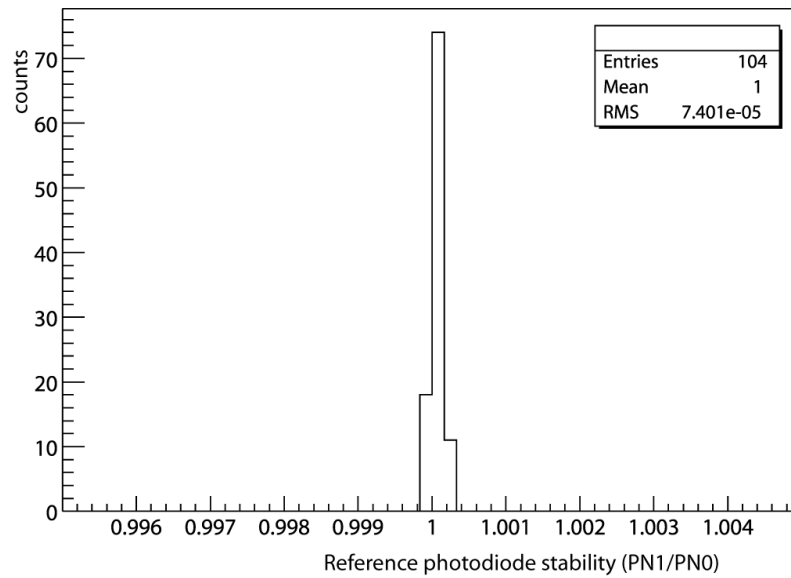


Figure 83: Relative stability between a pair of reference PN photodiodes monitoring 200 crystals measured in autumn 2004 at the CERN test beam facility.

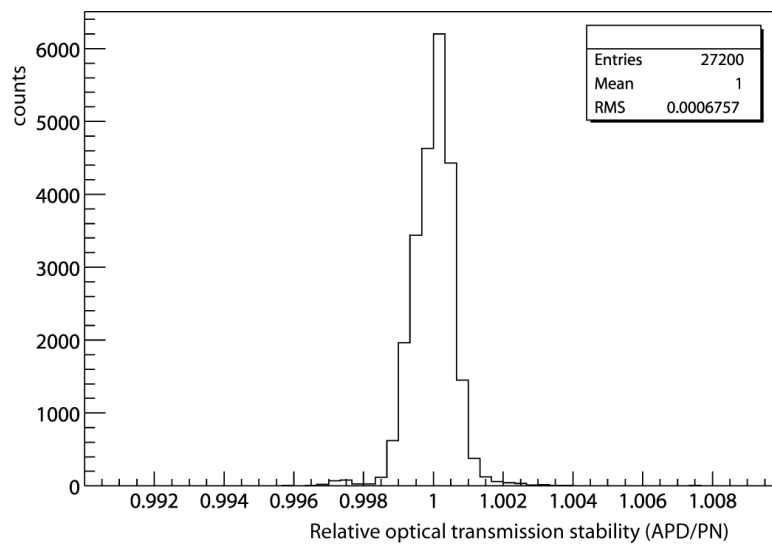


Figure 84: Stability of crystal transmission measurements at 440 nm (blue laser) over 11.5 days operation for a module of 200 crystals.

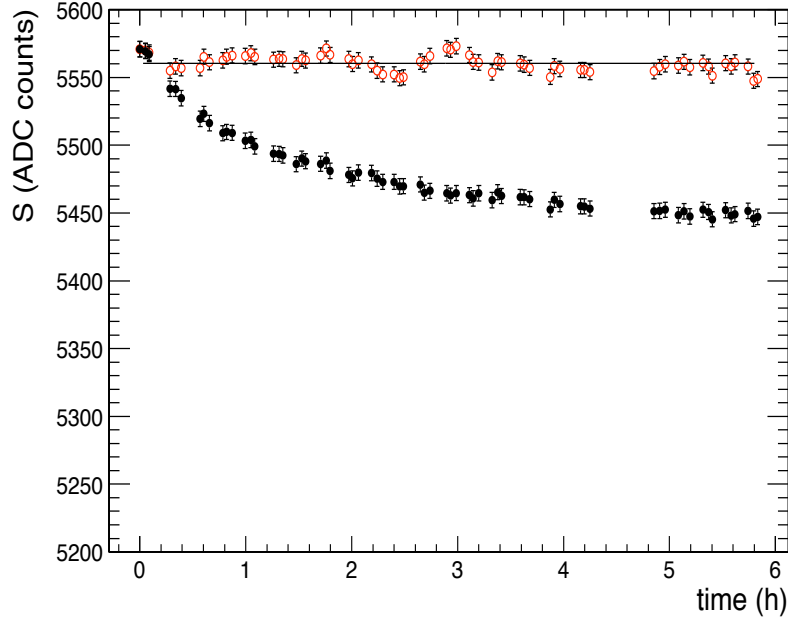


Figure 85: Effect of the monitor correction procedure on test beam data: full black points refer to signals measured during test beam irradiation, open red points are the same after the monitor corrections.

The stochastic term

There are three basic contributions to the stochastic term:

1. event-to-event fluctuations in the lateral shower containment,
2. a photostatistics contribution of 2.1%,
3. fluctuations in the energy deposited in the preshower absorber (where present) with respect to what is measured in the preshower silicon detector.

The contribution to the stochastic term coming from fluctuations in the lateral containment is expected to be about 1.5% when energy is reconstructed by summing an array of 5×5 crystals, and about 2% when using 3×3 crystals.

The photostatistics contribution is given by:

$$a_{pe} = \sqrt{\frac{F}{N_{pe}}} \quad (4)$$

where N_{pe} is the number of primary photoelectrons released in the photodetector per GeV, and F is the excess noise factor which parametrizes fluctuations in the gain process. This factor has a value close to 2 for the APDs, and is about 2.5 for the VPTs. A value of $N_{pe} \approx 4500$ pe/GeV is found for the barrel, giving $\approx 2.1\%$ for the photostatistics contribution to the stochastic term. In the endcap the photostatistics contribution is similar, since the larger collection area of the VPT largely compensates for the reduced quantum efficiency of the photocathode.

The contribution to the energy resolution from the preshower device can be approximately parametrized as a stochastic term with a value of $5\%/\sqrt{E}$, where E is in GeV. But, because it samples only the beginning of the shower, the resolution is, in fact, predicted to vary like $\sigma/E \propto 1/E^{0.75}$. A beam test in 1999 [104] verified this prediction.

The constant term

The most important contributions to the constant term may be listed as follows:

1. non-uniformity of the longitudinal light collection,

2. intercalibration errors,
3. leakage of energy from the back of the crystal.

The effects of the longitudinal light collection curve have been studied in detail. Quite stringent requirements are made on the crystal longitudinal uniformity. Requiring the constant term contribution due to non-uniformity be less than 0.3%, sets a limit on the slope of the longitudinal light collection curve in the region of the shower maximum of $\approx 0.35\%$ per radiation length. A small increase in response towards the rear of the crystal helps to compensate the rear leakage from late developing showers, which would otherwise cause a low energy tail. The required response is achieved in the barrel by depolishing one long face of the crystals to a designated roughness. This surface treatment is incorporated into the crystal production process.

The effect of rear leakage is very small. Charged particles leaking from the back of the crystals can also give a direct signal in the APDs (nuclear counter effect), but test beam data show that this effect is negligible for isolated electromagnetic showers: no tails on the high side of the energy distribution are observed even at the highest electron energy tested (280 GeV).

The noise term

There are three contributions to the noise term:

1. electronics noise,
2. digitization noise,
3. pileup noise.

The signal amplitude in the test beam is reconstructed using a simple digital filter. The noise measured, after this amplitude reconstruction, for channels in barrel supermodules is ≈ 40 MeV/channel in the highest gain range. This noise includes both electronics and digitization noise. The amplitude reconstruction makes use of an event-by-event baseline subtraction using 3 digitization samples taken directly before the signal pulse. This procedure removes the small channel-to-channel correlated noise. Its success is evidenced by the fact that, after this procedure, the noise in the sum of 25 channels is almost exactly 5 times the noise in a single channel [105].

In the endcap it is intended to sort the VPTs in bins of overall signal yield, which includes the photocathode area, the quantum efficiency and the VPT gain. The VPTs with higher overall signal yield are used for the larger radius regions of the endcap. This has the result that the transverse energy equivalent of the noise will be more or less constant, with a value of $\sigma_{ET} \approx 50$ MeV.

Neutron irradiation of the APDs in the barrel induces a leakage current which contributes to the electronics noise. The evolution of the leakage current and induced noise over the lifetime of the experiment has been extensively studied. The expected contribution is equivalent to 8 MeV/channel after one year of operation at $\mathcal{L} = 10^{33} \text{ cm}^{-2}\text{s}^{-1}$, and 30 MeV/channel at the end of the first year of operation at $\mathcal{L} = 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ [68].

The shaped signals from the preamplifier output will extend over several LHC bunch crossings. When using a multi-weights method to reconstruct the signal amplitude [105], up to 8 time samples are used. Pileup noise will occur if additional particles reaching the calorimeter cause signals which overlap these samples.

The magnitude of pileup noise expected at low luminosity ($\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$) has been studied using detailed simulation of minimum bias events generated between -5 and $+3$ bunch crossings before and after the signal. The average number of minimum bias events per bunch crossing was 3.5. Figure 86 shows the reconstructed amplitude observed with and without pileup in the absence of any signal. The fraction of events with a signal beyond the Gaussian distribution of the electronics noise is small, showing that at low luminosity the pileup contribution to noise is small.

Energy resolution in the test beam

In 2004 a fully equipped barrel supermodule was tested in the CERN H4 beam. The energy resolution measured with electron beams having momenta between 20 and 250 GeV/c confirmed the expectations described above [106]. Since the electron shower energy contained in a finite crystal matrix depends on the particle impact position with respect to the matrix boundaries, the intrinsic performance of the calorimeter was studied by using

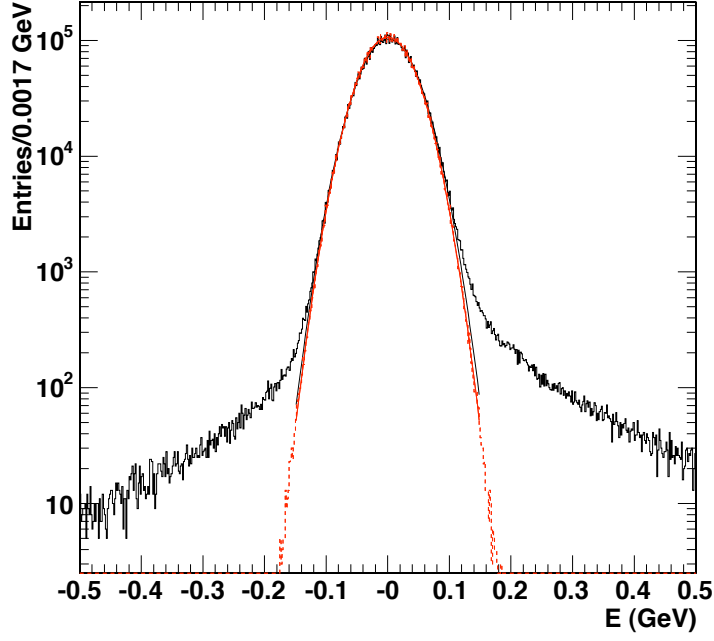


Figure 86: Reconstructed amplitude in ECAL barrel channels in the absence of a signal, without pileup (dashed histogram) and with pileup (solid histogram). A Gaussian of width 40 MeV is superimposed on the dashed histogram.

events where the electron was limited to a $4 \times 4 \text{ mm}^2$ region around the point of maximum containment (*central impact*). Figure 3 shows the resolution as a function of energy when the incident electrons were restricted in this way. The energy is reconstructed by summing 3×3 crystals. A typical energy resolution was found to be:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{2.8\%}{\sqrt{E}}\right)^2 + \left(\frac{0.12}{E}\right)^2 + (0.30\%)^2, \quad (5)$$

where E is in GeV. This result is in agreement with the expected contributions detailed in the earlier part of this Section. (Results from beam-test runs taken in 2006, using the final VFE card, show a 10% improvement of the noise performance.)

The energy resolution was also measured with no restriction on the lateral position of the incident electrons except that provided by the $20 \times 20 \text{ mm}^2$ trigger. The trigger was roughly centred ($\pm 3 \text{ mm}$) on the point of maximum response of a crystals. In this case a shower containment correction was made as a function of incident position, as measured from the distribution of energies in the crystal, to account for the variation of the amount of energy contained in the matrix. For energy reconstruction in either a 3×3 or a 5×5 matrix an energy resolution of better than 0.45% is found for 120 GeV electrons after correction for containment. Figure 87 shows an example of the energy distributions before and after correction for the case of reconstruction in a 5×5 matrix, where the correction is smaller than for the 3×3 case.

The energy resolution has also been measured for a series of 25 runs where the beam was directed at locations uniformly covering a 3×3 array of crystals. In this case a resolution of 0.5% was measured for 120 GeV electrons.

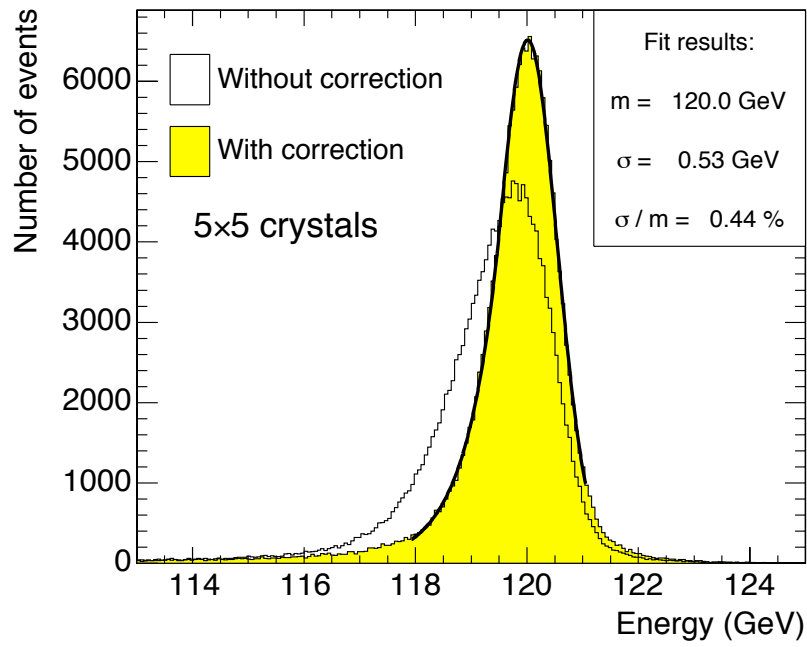


Figure 87: Distribution of energy reconstructed in a 5×5 matrix, before and after correction for containment, when 120 GeV electrons are incident over a $20 \times 20 \text{ mm}^2$ area.

6 Hadron calorimeter

The CMS detector is designed to study a wide range of high-energy processes involving diverse signatures of final states. The hadron calorimeters are particularly important for the measurement of hadron jets and neutrinos or exotic particles resulting in apparent missing transverse energy [1].

Figure 88 shows the longitudinal view of the CMS detector. The dashed lines are at fixed η values. The hadron calorimeter barrel and endcaps sit behind the tracker and the electromagnetic calorimeter as seen from the interaction point. The hadron calorimeter barrel is radially restricted between the outer extent of the electromagnetic calorimeter ($R = 1.77$ m) and the inner extent of the magnet coil ($R = 2.95$ m). This constrains the total amount of material which can be put in to absorb the hadronic shower. Therefore, an outer hadron calorimeter or *tail catcher* is placed outside the solenoid complementing the barrel calorimeter. Beyond $|\eta| = 3$, the forward hadron calorimeters placed at 11.2 m from the interaction point extend the pseudorapidity coverage down to $|\eta| = 5.2$ using a Cherenkov-based, radiation-hard technology. The following Sections describe these subdetectors in detail.

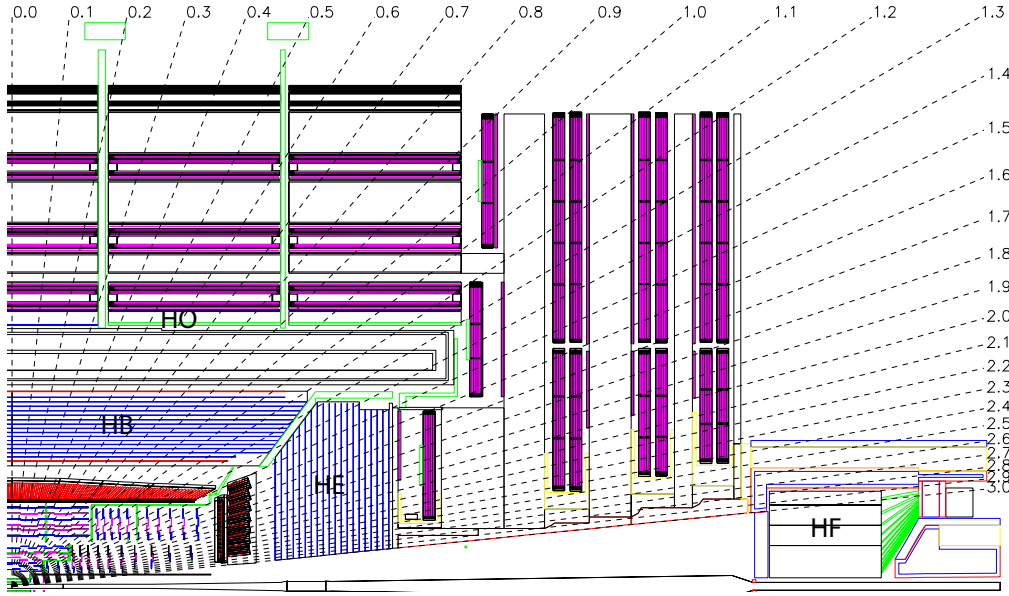


Figure 88: Longitudinal view of the CMS detector showing the locations of the hadron barrel (HB), endcap (HE), outer (HO) and forward (HF) calorimeters.

6.1 Barrel design (HB)

The HB is a sampling calorimeter covering the pseudorapidity range $|\eta| < 1.3$. The HB is divided into two half-barrel sections (Fig. 89), each half-section being inserted from either end of the barrel cryostat of the superconducting solenoid and subsequently hung from rails in the median plane. Since the HB is very rigid compared to the cryostat, great care has been taken to ensure that the barrel load is distributed evenly along the rails [107].

Absorber geometry

The HB consists of 36 identical azimuthal wedges which form the two half-barrels (HB+ and HB-). The wedges are constructed out of flat brass absorber plates (Table 7) aligned parallel to the beam axis. The numbering scheme of the wedges is shown in Fig. 90. Each wedge is segmented into four azimuthal angle (ϕ) sectors. The plates are bolted together in a staggered geometry resulting in a configuration that contains no projective dead material for the full radial extent of a wedge (Fig. 91). The innermost and outermost plates are made of stainless steel for structural strength. The plastic scintillator is divided into 16 η sectors, resulting in a segmentation $(\Delta\eta, \Delta\phi) = (0.087, 0.087)$. The wedges are themselves bolted together, in such a fashion as to minimize the crack between the wedges to less than 2 mm.

The absorber (Table 8) consists of a 40-mm-thick front steel plate, followed by eight 50.5-mm-thick brass plates, six 56.5-mm-thick brass plates, and a 75-mm-thick steel back plate. The total absorber thickness at 90° is 5.82 interaction lengths (λ_I). The HB effective thickness increases with polar angle (θ) as $1/\sin\theta$, resulting in $10.6 \lambda_I$ at $|\eta| = 1.3$. The electromagnetic crystal calorimeter [68] in front of HB adds about $1.1 \lambda_I$ of material.

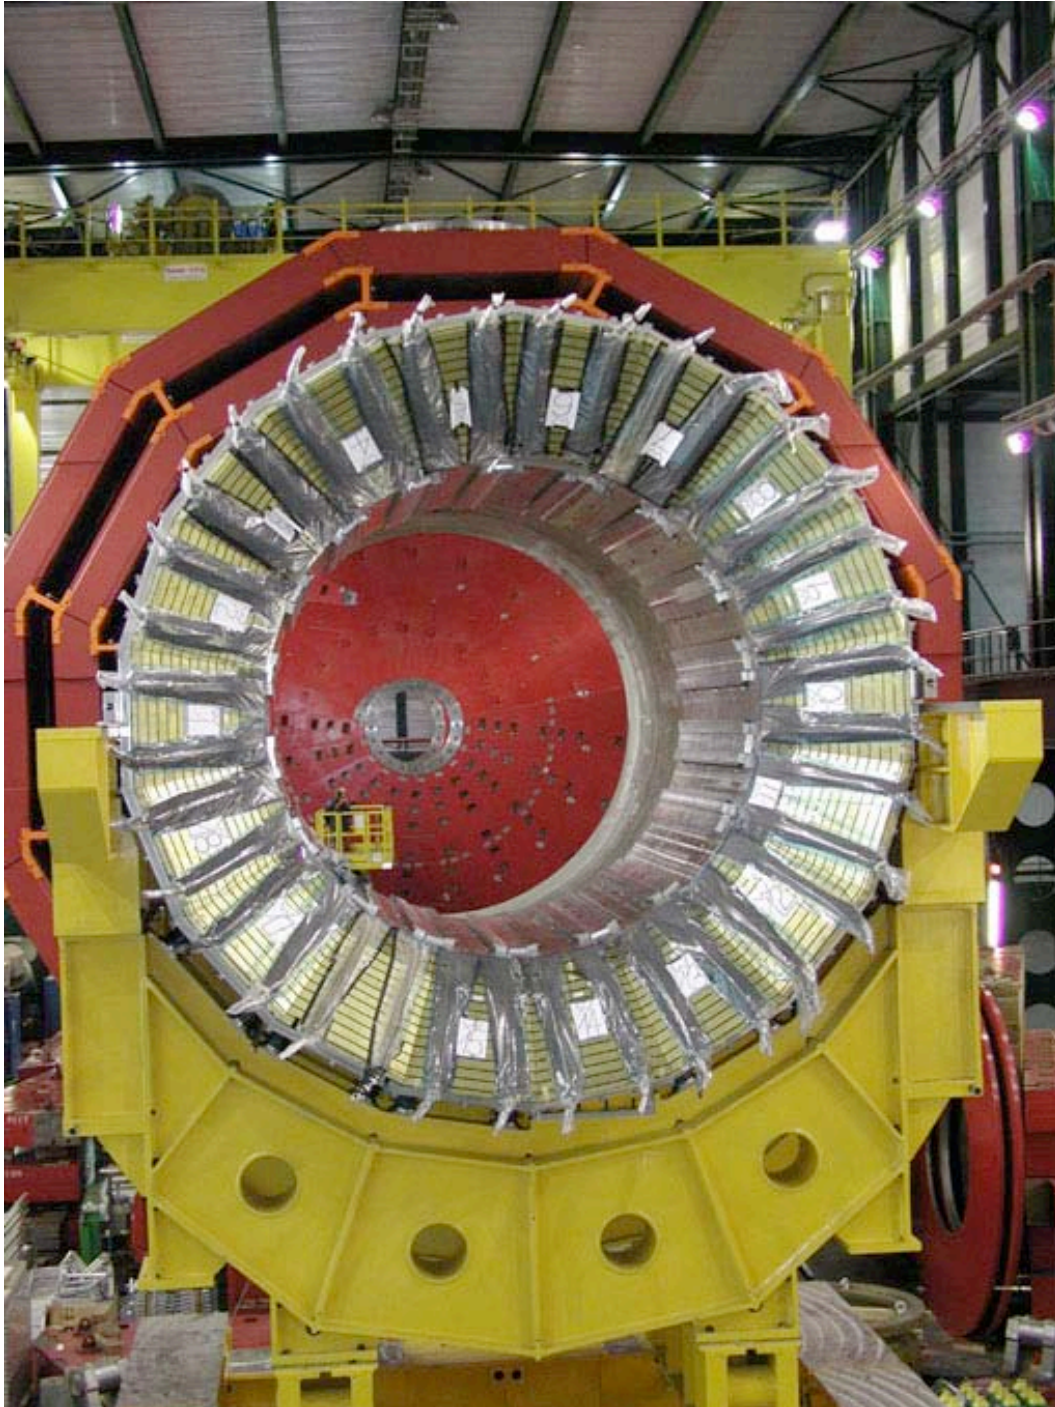


Figure 89: Assembled HCAL half-barrel in SX5, the above ground assembly hall.

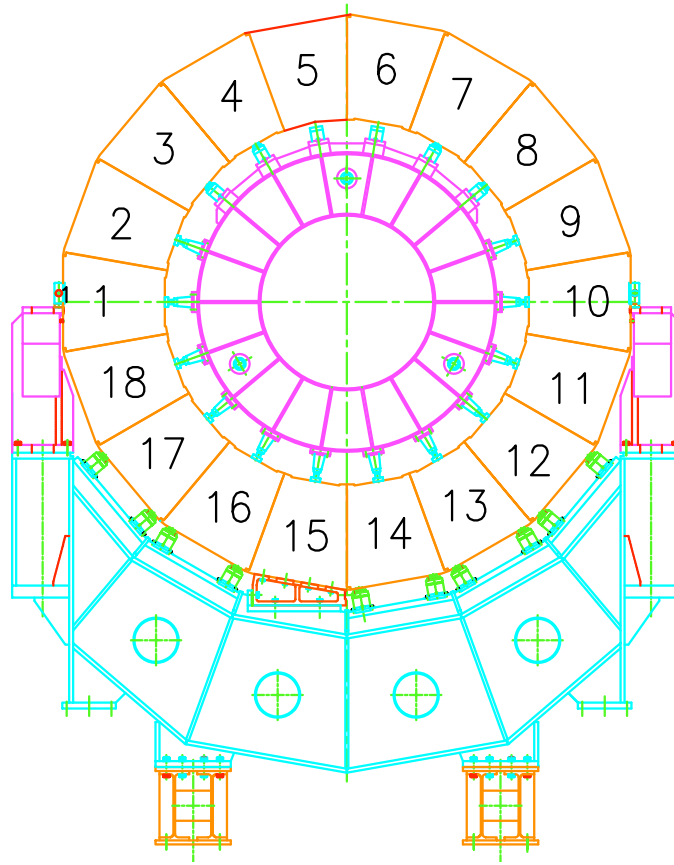


Figure 90: Numbering scheme for the HB wedges. Wedge 1 is on the inside ($+x$ direction) of the LHC ring.

Table 7: Physical properties of the HB brass absorber, known as C26000/cartridge brass.

chemical composition	70% Cu, 30% Zn
density	8.53 g/cm ³
radiation length	1.49 cm
interaction length	16.42 cm

Table 8: Absorber thickness in the HB wedges.

<i>layer</i>	<i>material</i>	<i>thickness</i>
front plate	steel	40 mm
1-8	brass	50.5 mm
9-14	brass	56.5 mm
back plate	steel	75 mm

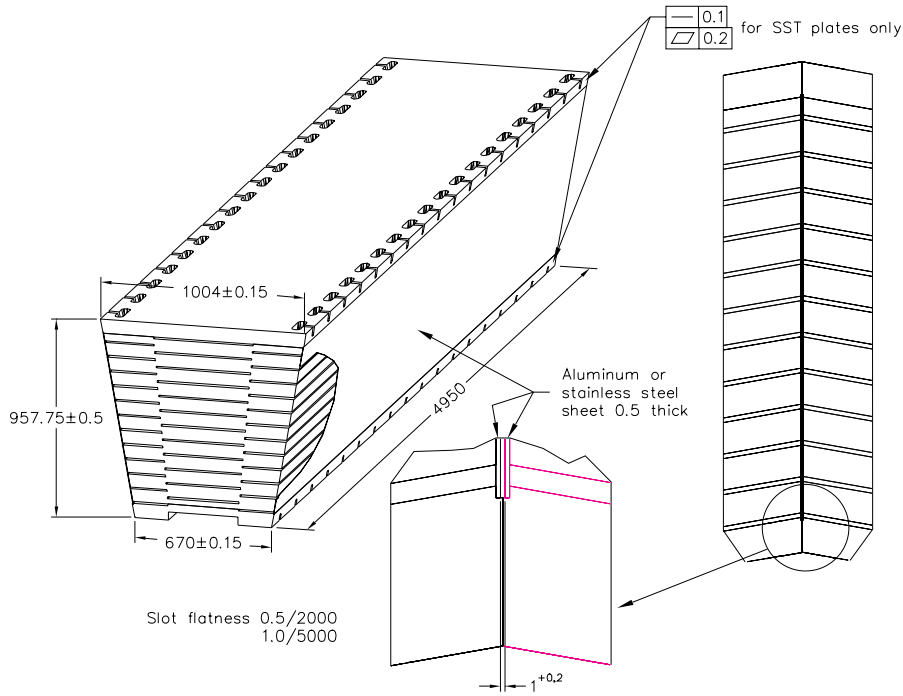


Figure 91: Isometric view of the HB wedges, showing the hermetic design of the scintillator sampling.

Scintillator

The active medium uses the well known tile and wavelength shifting fibre concept to bring out the light. The CMS hadron calorimeter consists of about 70 000 tiles. In order to limit the number of individual elements to be handled, the tiles of a given ϕ layer are grouped into a single mechanical scintillator tray unit. Figure 92 shows a typical tray. The tray geometry has allowed for construction and testing of the scintillators remote from the experimental installation area. Furthermore, individual scintillator trays may be replaced without disassembly of the absorber in the event of catastrophic damage. Each HB wedge has four ϕ divisions (ϕ -index = 1–4). Trays with segmentation of ϕ -index 2 and 3 go into the center of a wedge while trays with segmentation of ϕ -index 1 and 4 go into the edge slots in a wedge (Fig. 91). Each layer has 108 trays. Figure 93 shows a cross section of the tray.

The HB baseline active material is 3.7-mm-thick Kuraray SCSN81 plastic scintillator, chosen for its long-term stability and moderate radiation hardness. The first layer of scintillator (layer 0) is located in front of the steel support plate. It was originally foreseen to have a separate read-out [107] and is made of 9-mm-thick Bicron

Typical megatiles (16 towers)

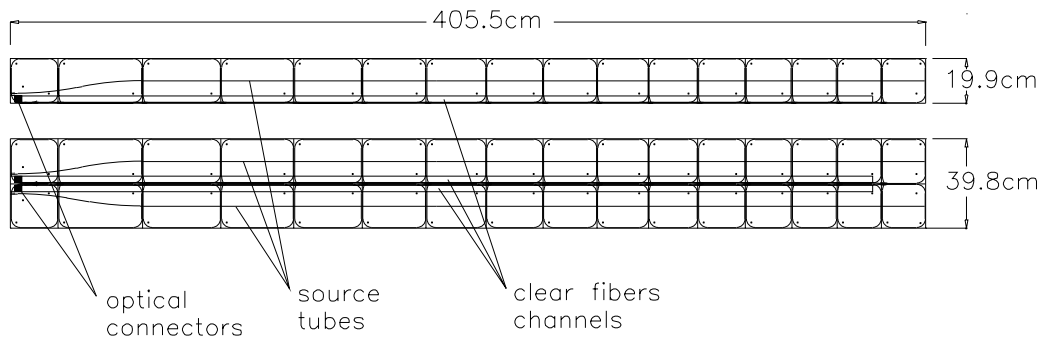


Figure 92: Scintillator trays.

Cross Section

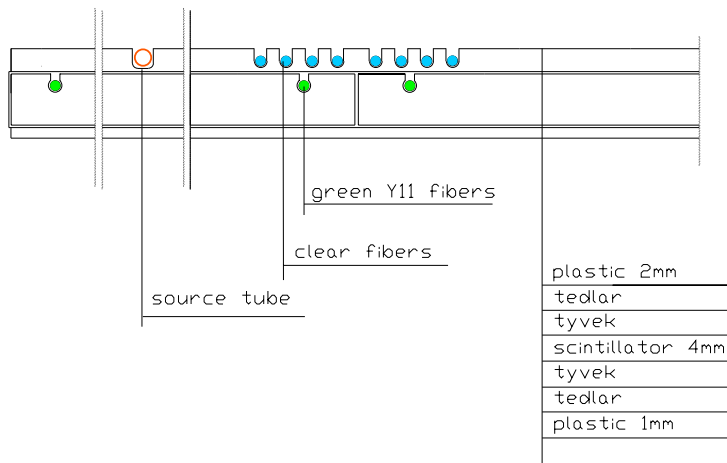


Figure 93: Cross-sectional view of a scintillator tray.

BC408. The scintillators are summarized in Table 9. The purpose of layer zero is to sample hadronic showers developing in the inert material between the EB and HB. The larger thickness of layer 16 serves to correct for late developing showers leaking out the back of HB.

Table 9: Scintillator in the HB wedges.

<i>layer</i>	<i>material</i>	<i>thickness</i>
0	Bicron BC408	9 mm
1-15	Kuraray SCSN81	3.7 mm
16	Kuraray SCSN81	9 mm

A tray is made of individual scintillators with edges painted white and wrapped in Tyvek 1073D which are attached to a 0.5-mm-thick plastic substrate with plastic rivets. Light from each tile is collected with a 0.94-mm-diameter green double-cladded wavelength-shifting fibre (Kuraray Y-11) placed in a machined groove in the scintillator. For calibration purposes, each tray has 1-mm-diameter stainless steel tubes, called *source tubes*, that carry Cs^{137} (or optionally Co^{60}) radioactive sources through the center of each tile. An additional quartz fibre is used to inject ultraviolet (337 nm) laser light into the layer 9 tiles. The top of the tray is covered with 2-mm-thick white polystyrene. The cover is grooved to provide routing paths for fibres to the outside of the tray and also to accommodate the tubes for moving radioactive sources.

After exiting the scintillator, the wavelength shifting fibres (WLS) are spliced to clear fibres (Kuraray double-clad). The clear fibre goes to an optical connector at the end of the tray. An optical cable takes the light to an optical decoding unit (ODU). The ODU arranges the fibres into read-out towers and brings the light to a hybrid photodiode (HPD) [108]. An additional fibre enters each HPD for direct injection of light using either the laser or a light emitting diode (LED). A schematic of the fibre optics is shown in Fig. 94 and the actual cabling is shown in Fig. 95.

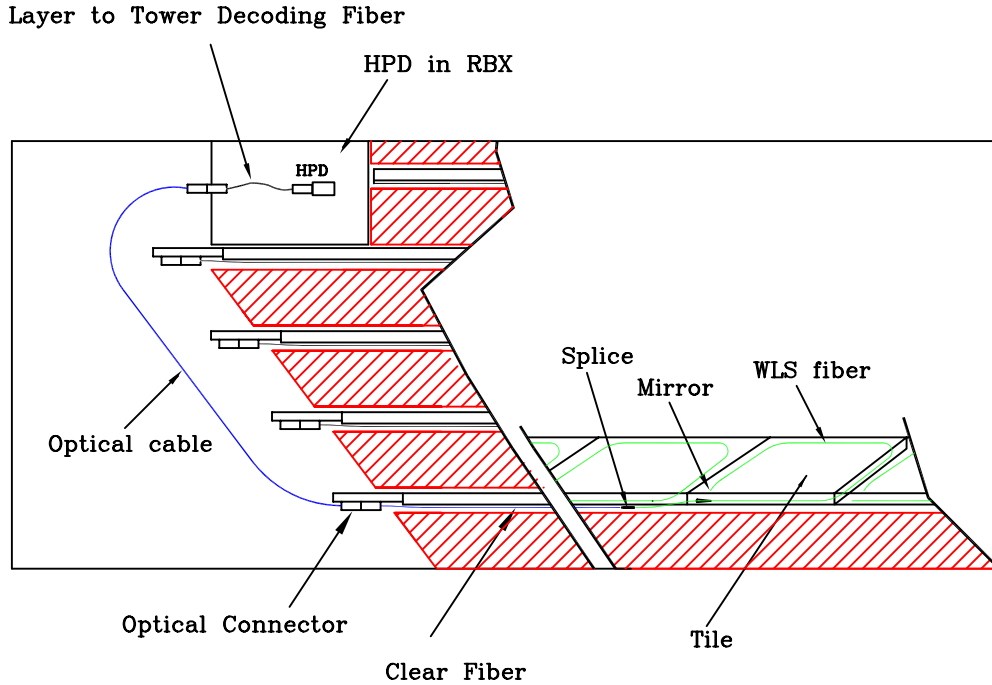


Figure 94: Schematic of the HB optics.

The HPD consists of a photocathode held at a HV of -8 kV at a distance of approximately 3.3 mm from a pixelated silicon photodiode. The ionization produced by the accelerated photoelectron in the diode results in a gain of the HPD of approximately 2000. There are 19 hexagonal 20-mm² pixels in a single HPD, the centermost of which is

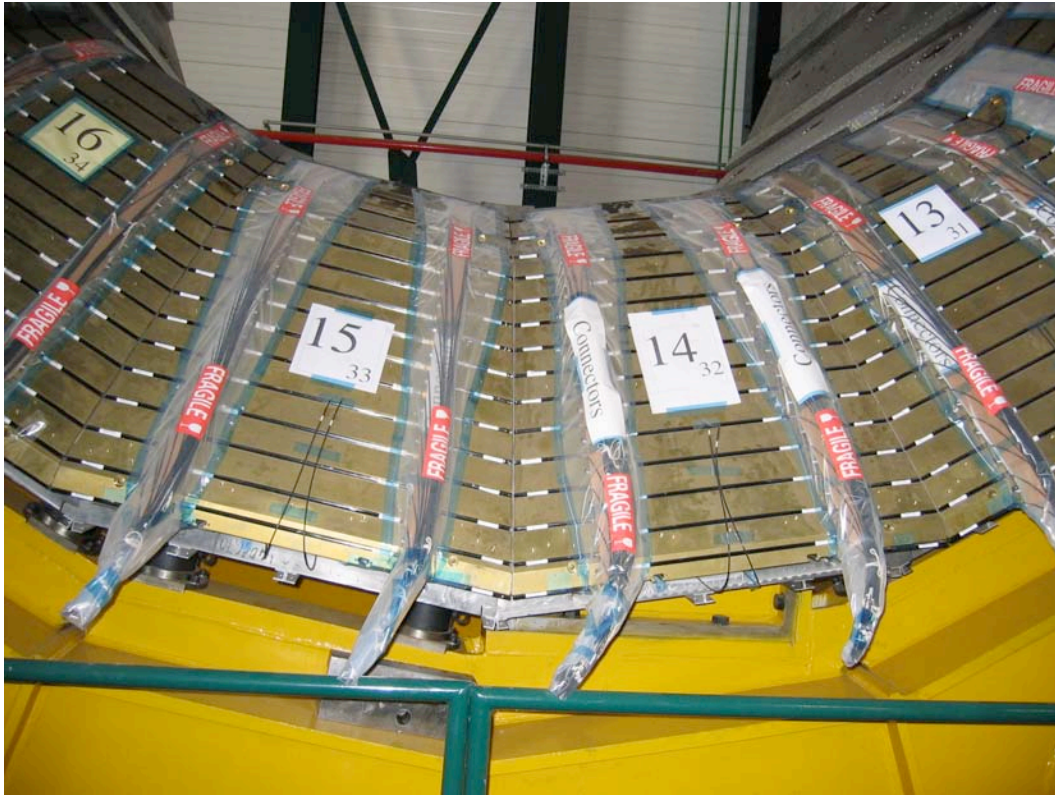


Figure 95: Close up view of the assembled HB wedges, showing the optical cabling.

not read-out. A cross sectional view of an HPD is shown in Fig. 96.

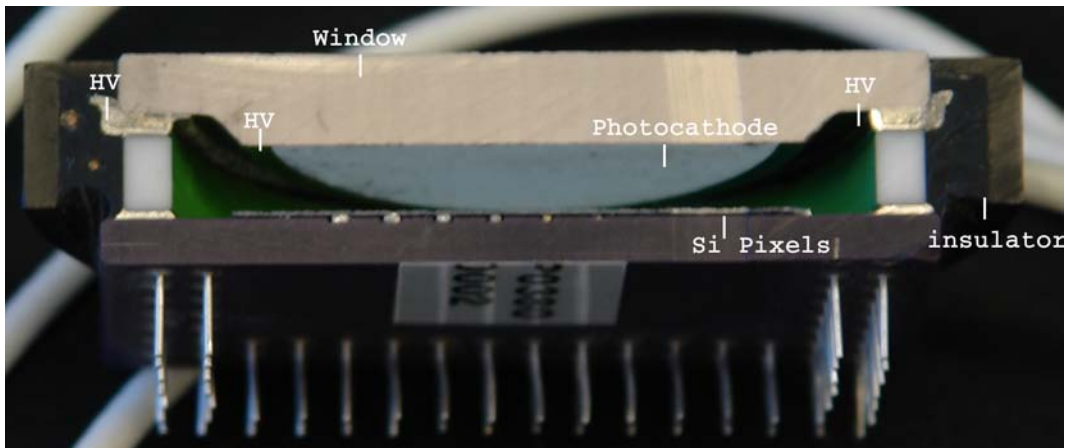


Figure 96: Cross sectional view of an HPD.

During the production and assembly process, the WLS fibres are cut, polished, and mirrored. The reflectivity of the mirror is checked by measuring test fibres which are mirrored along with the fibres used in the calorimeter. Measuring the reflectivity of the mirror is done with a computer controlled UV scanner with the fibres read out by photodiodes. Clear fibres are spliced onto WLS fibres with a fusion splicer. The transmission across the splice is checked by splicing a sample of WLS fibres onto WLS fibres. The splice region is measured with the UV scanner. The transmission across the splice is 92.6% with an RMS of 1.8%. Next, the optical fibres are glued into a 10 fibre connector. This configuration is called a pigtail. In order to get the fibre lengths correct, the pigtail is assembled in a template. The connector is diamond polished. The fibres are measured with the UV scanner. The scanner checks the green fibre, clear fibre, splice, and mirror. The RMS of the light from the fibres is 1.9%. After the pigtail is inserted into the tray, the completed tray is checked with an automated source scanner using a Cs^{137} source inside

a lead collimator. This yields a 4 cm diameter source spot on the tray. The collimator is moved with a computer controlled x-y motor. From the scanner we determine the relative light yield of each tile and the uniformity of each tray. The light yield of the individual tiles has an RMS of 4.6%, while the transverse uniformity of the tile is 4.5%. A Cs^{137} wire source is run through the 4 source tubes and the light yield is measured. The RMS of the ratio of collimated source to wire source is 1.3%. This means the line sources, which can be used when the calorimeter is completely assembled, can calibrate individual tiles to better than 2%. In addition to the moving wire source, there are laser and LED light injection systems.

Longitudinal segmentation

The η towers 1–14 have a single longitudinal read-out. The η towers closest to the endcap transition region (15 and 16) are segmented in depth. The front segment of tower 15 contains either 12 or 13 scintillators, due to the placement of the read-out box and the staggering of the layers (layers 0–11 for the middle two ϕ sectors and 0–12 layers for the outer two ϕ sectors). The rear segment of tower 15 has three scintillators. Tower 16, which is in front of the endcap (HE) has five scintillators in the front segment and three in the rear. The front segment of tower 16 does not have a layer-0 scintillator. The tower segmentation is summarized in Fig. 97 and Table 10.

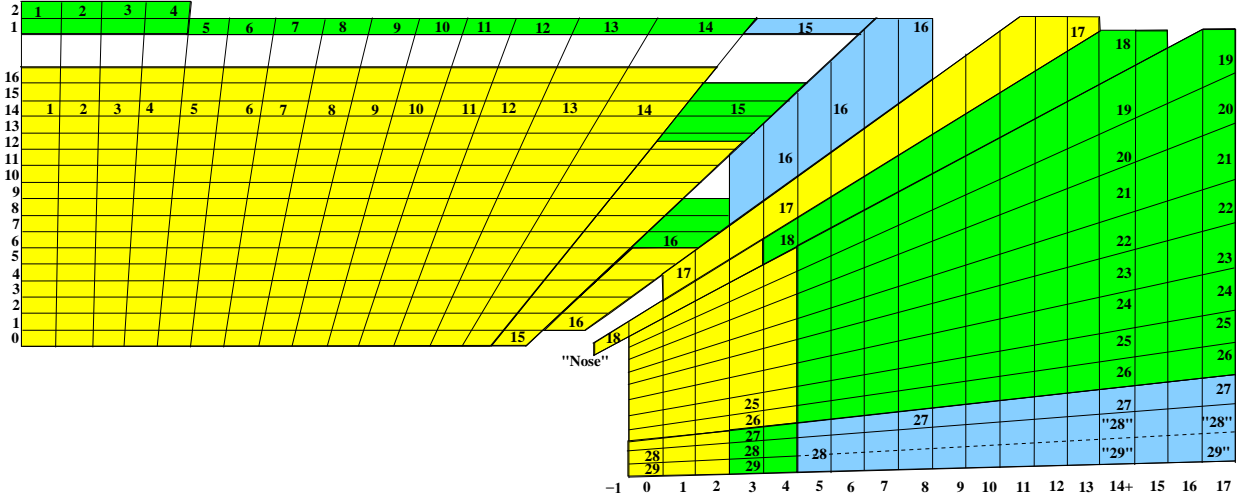


Figure 97: The HCAL tower segmentation in the r, z plane for one-fourth of the HB, HO, and HE detectors. The numbers on the left (0-16) indicate the scintillator layers. The shading represents different longitudinal readouts.

6.2 Endcap design (HE)

The hadron calorimeter endcaps (HE) [107] cover a substantial portion of the rapidity range, $1.3 < |\eta| < 3$ (13.2% of the solid angle), a region containing about 34% of the particles produced in the final state. The high luminosity of the LHC ($10^{34} \text{ cm}^{-2} \text{ s}^{-1}$) requires HE to handle high (MHz) counting rates and have high radiation tolerance (10 MRad after 10 years of operation at design luminosity) at $|\eta| \simeq 3$. Since the calorimeter is inserted into the ends of a 4-T solenoidal magnet, the absorber must be made from a non-magnetic material. It must also have a maximum number of interaction lengths to contain hadronic showers, good mechanical properties and reasonable cost, leading to the choice of C26000 cartridge brass. The endcaps are attached to the muon endcap yoke as shown in Figs. 98 and 99. Only a small part of the calorimeter structure can be used for the fixation to the magnet iron, because the majority of the space between HE and muon absorber is occupied with muon cathode strip chambers. A 10-t electromagnetic calorimeter (EE) with a 2-t preshower detector (ES) is attached at the front face of HE. The large weight involved (about 300 t) and a strict requirement to minimize non-instrumented materials along particle trajectories, has made the design of HE an unprecedented challenge to engineers. An interface kinematic scheme was developed in order to provide precise positioning of the endcap detectors with respect to the adjacent muon

Table 10: Tower data for HB. The given thicknesses correspond to the center of the tower. Note that tower 16 overlaps with HE.

<i>tower</i>	<i>η range</i>	<i>thickness (λ_I)</i>
1	0.000 – 0.087	5.39
2	0.087 – 0.174	5.43
3	0.174 – 0.261	5.51
4	0.261 – 0.348	5.63
5	0.348 – 0.435	5.80
6	0.435 – 0.522	6.01
7	0.522 – 0.609	6.26
8	0.609 – 0.696	6.57
9	0.696 – 0.783	6.92
10	0.783 – 0.870	7.32
11	0.870 – 0.957	7.79
12	0.957 – 1.044	8.30
13	1.044 – 1.131	8.89
14	1.131 – 1.218	9.54
15	1.218 – 1.305	10.3
16	1.305 – 1.392	overlaps with HE

station, and to minimize the influence of deformation under magnetic forces. The interface kinematic contains a sliding joint between the interface tube, and HE back-flange and the hinge connection between brackets and the iron disk (YE1). Structural materials used in the interface system are non-magnetic in order not to distort the axial magnetic field of up to 4 T.

Absorber geometry

The design of the absorber is driven by the need to minimize the cracks between HB and HE, rather than single-particle energy resolution, since the resolution of jets in HE will be limited by pileup, magnetic field effects, and parton fragmentation [109, 110]. The plates are bolted together in a staggered geometry resulting in a configuration that contains no projective “dead” material (Fig. 100). The design provides a self-supporting hermetic construction. The brass plates are 79-mm-thick with 9-mm gaps to accommodate the scintillators. The total length of the calorimeter, including electromagnetic crystals, is about 10 interaction lengths (λ_I).

The outer layers of HE have a cutout region for installation of the photodetectors and front-end electronics. To compensate for the resulting reduction of material, an extra layer (-1) is added to tower 18 [111]. The outer layers are fixed to a 10-cm-thick stainless steel support plate. The optical elements are inserted into the gaps after the absorber is completely assembled; therefore, the optical elements must have a rigid structure to allow insertion from any position.

Scintillator trays

The scintillation light is collected by wavelength shifting (WLS) fibres [112, 113]. The design minimizes dead zones because the absorber can be made as a solid piece without supporting structures while at the same time the light can be easily routed to the photodetectors. Trapezoidal-shaped scintillators (Fig. 101), 4-mm-thick SCSN81 or 9-mm-thick Bicron BC408 for layer 0, have grooves in which the WLS fibres are inserted. The ends of the fibres are machined with a diamond fly cutter and one end is covered with aluminium to increase the light collection. The other end is spliced to a clear fibre, which is terminated in an optical connector. The connector with the glued fibres is also machined by a diamond fly cutter. The scintillator is painted along the narrow edges and put into a frame to form a tray. The total number of tiles for both HE calorimeters is 20 916 and the number of trays is 1368. The design of a tray is presented in Fig. 102. The numbering scheme in η is shown in Fig. 103, and the CMS convention for ϕ as applied to HE is shown in Fig. 104. The scintillators are wrapped with Tyvek and sandwiched between sheets of duraluminum. The stack contains holes for fibres which are terminated with optical connectors. The gap between the duraluminum plates is fixed by brass spacers screwed together. The granularity of the calorimeters is $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi \approx 0.17 \times 0.17$ for $|\eta| \geq 1.6$.

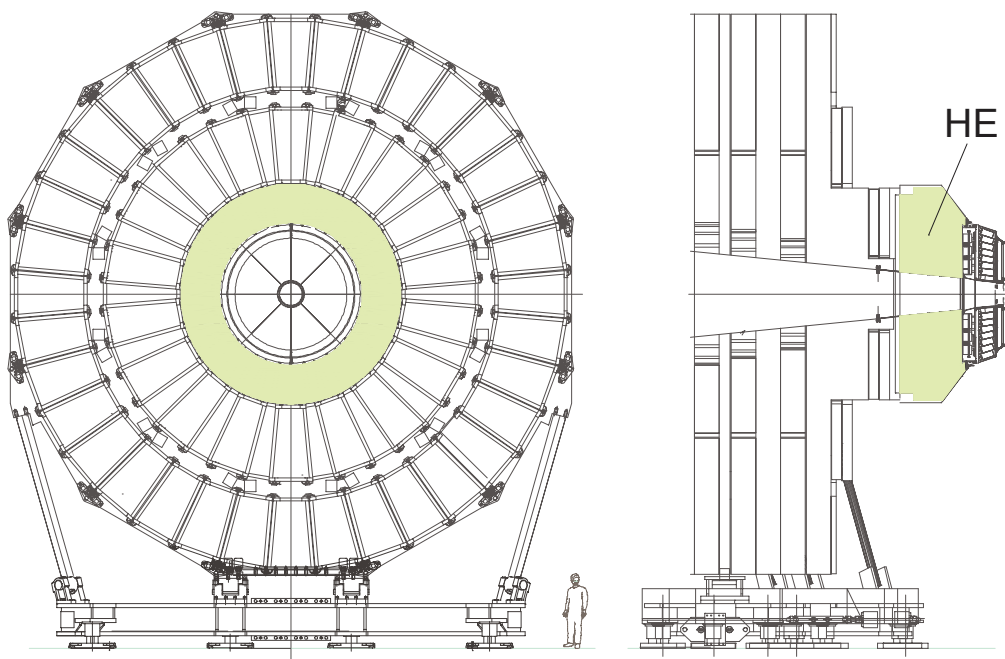


Figure 98: Hadron endcap (HE) calorimeter mounted on the endcap iron yoke.

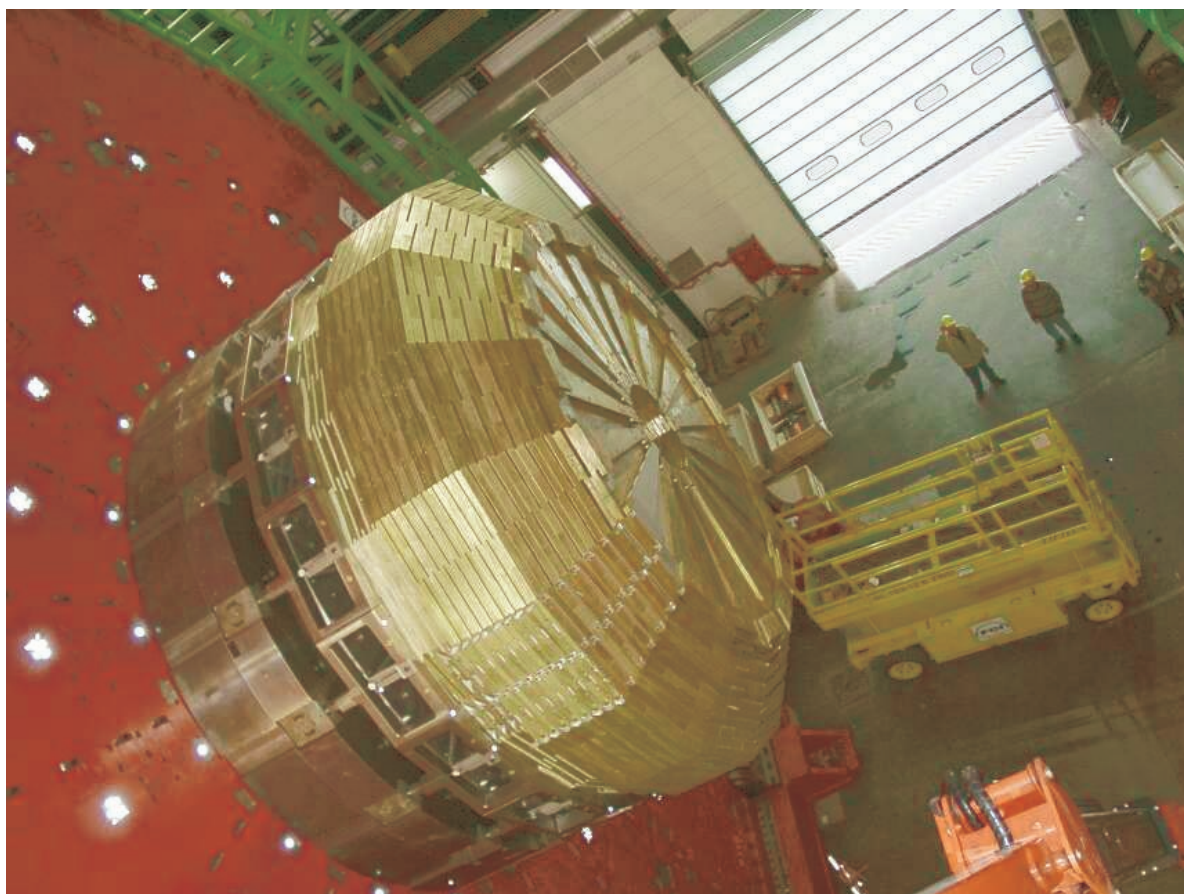


Figure 99: Partially assembled HE-minus absorber in the CMS surface hall (SX5). Scintillator trays can be seen to be inserted in some of the outer sectors.

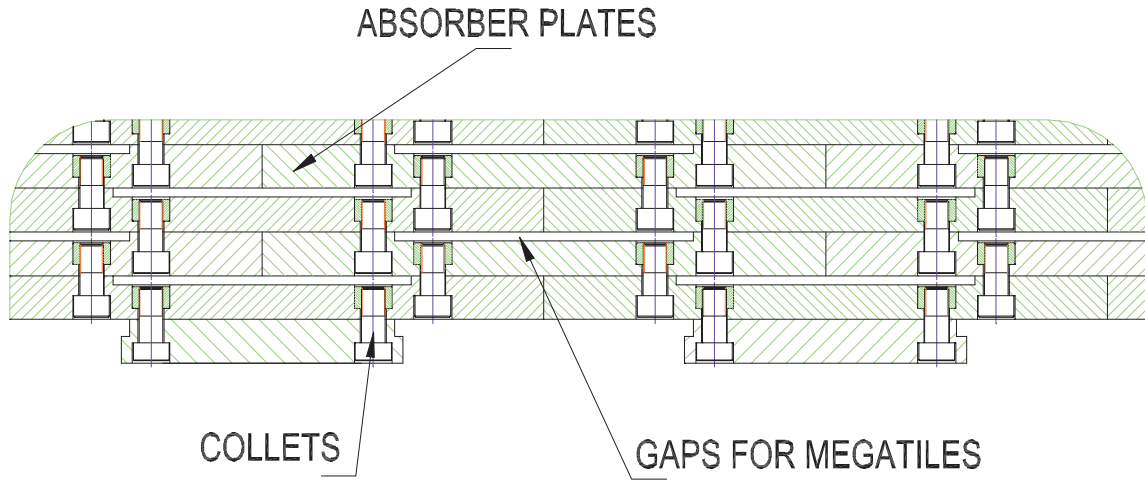


Figure 100: Mechanical structure of the HE absorber. Particles enter the calorimeter from the bottom.

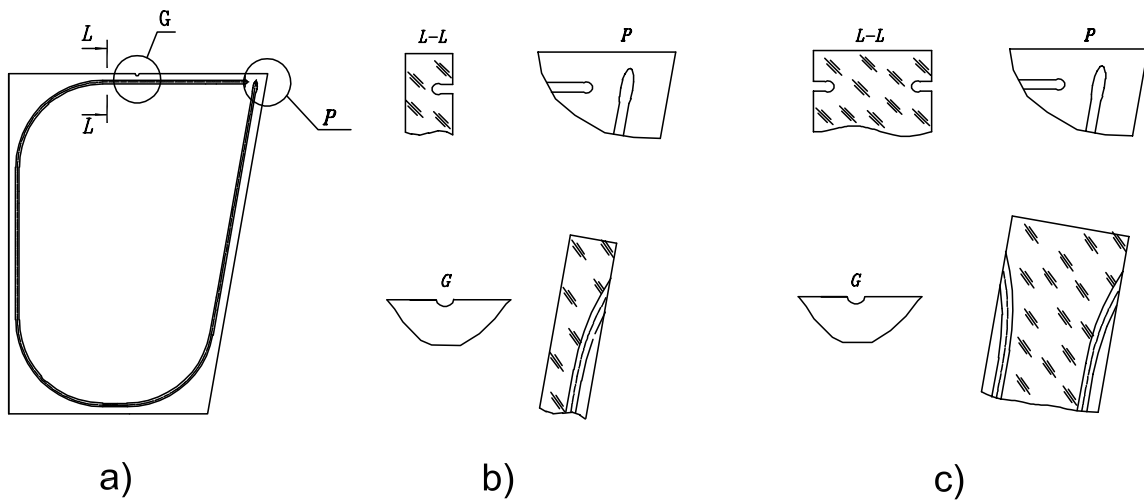
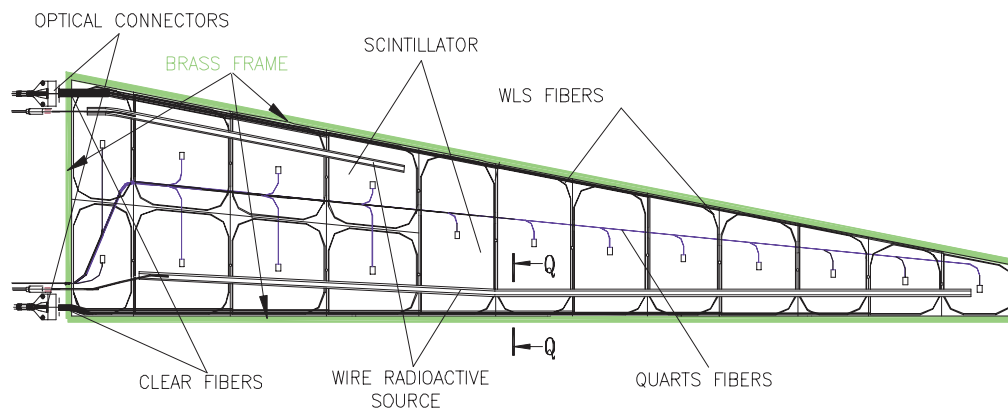
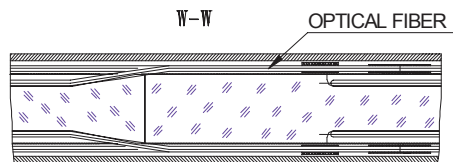
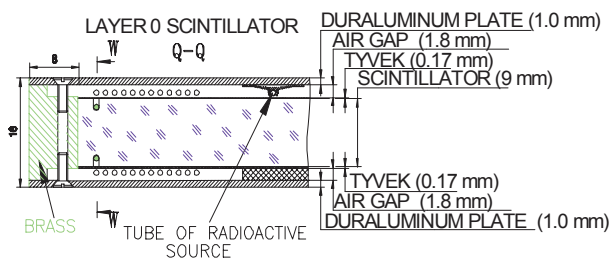


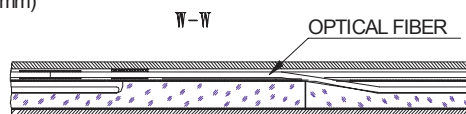
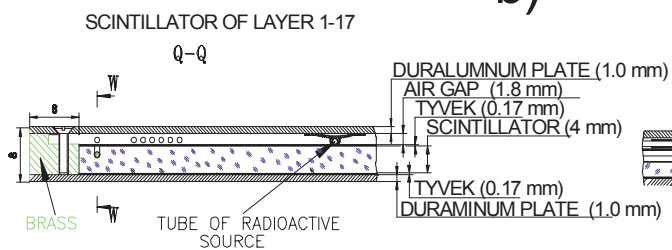
Figure 101: a) Basic structure of a scintillator tile with a groove to fix wavelength shifting fibre, b) cross section of the 4-mm-thick scintillator for layers 1–17, and c) cross section of the 9-mm-thick scintillator for layer zero. Two layers of reflecting paint cover the side surfaces of the tile.



a)



b)



c)

Figure 102: The design of the calorimeter scintillator trays: a) front view of a tray without upper aluminium cover, b) cut out view of the layer-0 tray with two fibres from a tile, c) cut out view of a tray for layers 1–17.

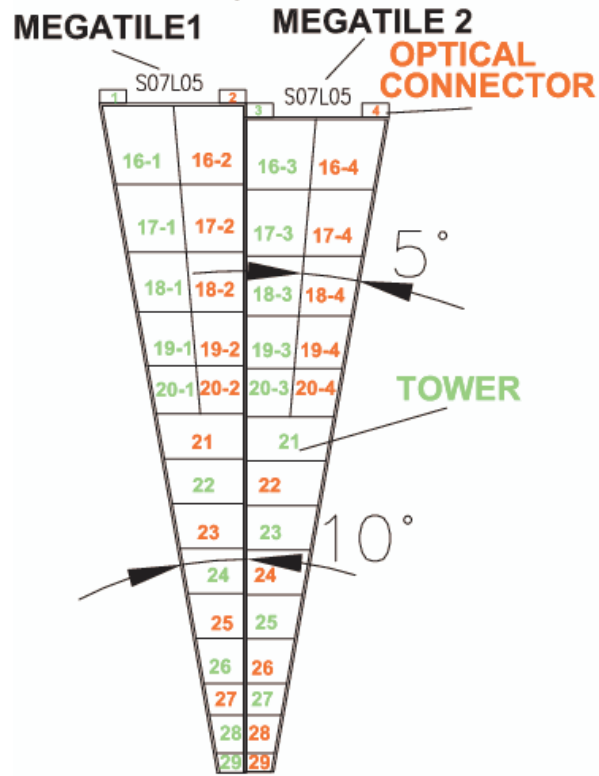


Figure 103: Numbering scheme for the tiles in adjacent scintillator trays.

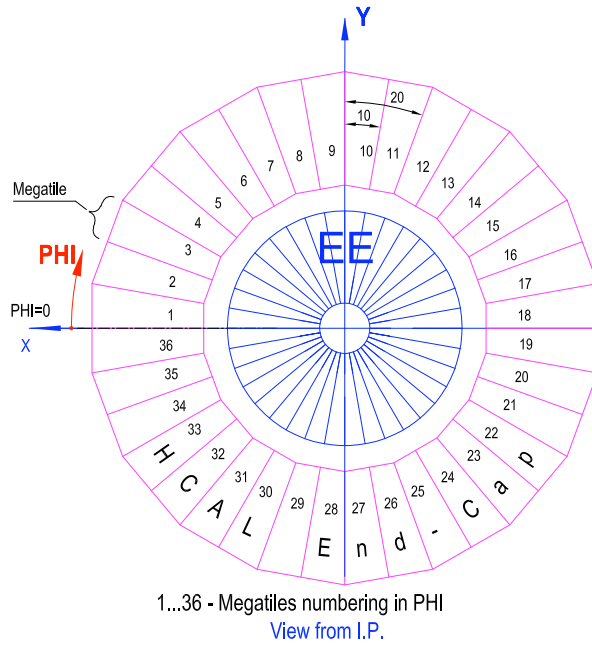


Figure 104: Numbering scheme for the HE wedges as viewed from the interaction point. The $+x$ direction points to the center of the LHC ring.

The tray design is very robust and reliable. The trays are relatively stiff which is very important for insertion into the absorber. To control the scintillator tray quality, a UV nitrogen laser was used to excite the scintillators. The light is fed by quartz fibres to the connector and is fanned out as shown in Fig. 102. These fibres are terminated with aluminium reflectors and distribute the light to all tiles. The light signal produced by a UV flash in the scintillator is similar to the signal induced by a charged particle. This allows a performance check of the entire optical route from scintillator to electronics, providing an important technique to track possible degradation of transparency due to radiation damage. For further calibration and monitoring, a radioactive source moving in a stainless steel tube is used to study the time-dependence of calibration coefficients.

The trays are inserted into the gaps in the absorber and fixed by screws. At the back of the calorimeter, boxes with photodetectors and electronics are located in the notch shown in Fig. 105. Optical cables transfer signals from the scintillator trays to the photodetectors. The partially assembled HE is shown in Fig. 99. Multipixel hybrid photodiodes (HPDs) are used as photodetectors due to their low sensitivity to magnetic fields and their large dynamical range.

Longitudinal segmentation

The longitudinal segmentation of HE (Fig. 97) is, in part, motivated by the radiation environment. Correction of the calibration coefficients after scintillator degradation can be applied, in order to restore the energy resolution. The towers nearest the beam line (27 and 28 plus guard ring “29”) have 3 divisions in depth which are read-out separately. The other towers (except 16 and 17 which overlap with the electromagnetic barrel calorimeter) have two longitudinal readouts for potential use during the time period when the electromagnetic endcap calorimeter (EE) may not yet be available. A special scintillator layer of 9 mm BC408 (layer 0) is installed in front of the absorber to partially correct for the different response of EE to electrons and hadrons and for particle absorption in the mechanical structure supporting EE.

6.3 Outer calorimeter design (HO)

In the central pseudorapidity region, the combined stopping power of EB plus HB does not provide sufficient containment for hadron showers. To ensure adequate sampling depth for $|\eta| < 1.3$, the hadron calorimeter is extended outside the solenoid with a tail catcher called the HO or outer calorimeter. The HO utilises the solenoid coil as an additional absorber equal to $1.4/\sin\theta$ interaction lengths and is used to identify late starting showers and to measure the shower energy deposited after HB.

Outside the vacuum tank of the solenoid, the magnetic field is returned through an iron yoke designed in the form of five 2.536 m wide (along z -axis) rings. The HO is placed as the first sensitive layer in each of these five rings. The rings are identified by the numbers $-2, -1, 0, +1, +2$. The numbering increases with z and the nominal central z positions of the five rings are respectively -5.342 m, -2.686 m, 0 , $+2.686$ m and $+5.342$ m. At $\eta = 0$, HB has the minimal absorber depth. Therefore, the central ring (ring 0) has two layers of HO scintillators on either side of a 19.5 cm thick piece of iron (the tail catcher iron) at radial distances of 3.82 m and 4.07 m, respectively. All other rings have a single HO layer at a radial distance of 4.07 m. The total depth of the calorimeter system is thus extended to a minimum of $11.8 \lambda_I$ except at the barrel-endcap boundary region.

The HO is constrained by the geometry of the muon system. Fig. 106 shows the position of HO layers in the rings of the muon stations in the overall CMS setup. The segmentation of these detectors closely follows that of the barrel muon system. Each ring has 12 identical ϕ -sectors. The 12 sectors are separated by 75-mm-thick stainless steel beams which hold successive layers of iron of the return yoke as well as the muon system. The space between successive muon rings in the η direction and also the space occupied by the stainless steel beams in the ϕ direction are not available for HO. In addition, the space occupied by the cryogenic “chimneys” in sector 3 of ring -1 , and sector 4 of ring $+1$ are also not available for HO. The chimneys are used for the cryogenic transfer lines and power cables of the magnet system. Finally, the mechanical structures needed to position the scintillator trays further constrain HO along ϕ .

In the radial direction each HO layer has been allocated a total of 40 mm, of which only 16 mm is available for the detector layer, the rest being used for the aluminium honeycomb support structures. In addition, the HO modules are independently supported from the steel beams located on either side of each ϕ sector. The thickness and position of the iron ribs in the yoke structure further constrain the shape and segmentation of the HO.

The sizes and positions of the tiles in HO are supposed to roughly map the layers of HB to make towers of granularity 0.087×0.087 in η and ϕ . The HO consists of one (rings ± 1 and ± 2) or two (ring 0) layers of scintillator

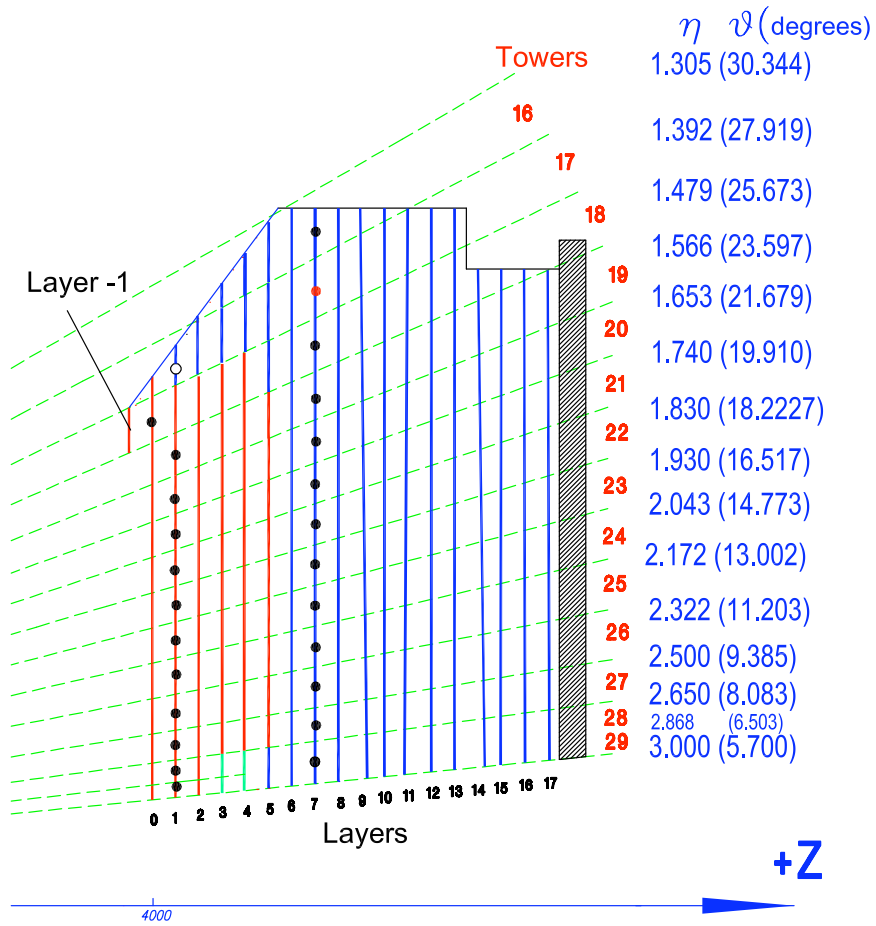


Figure 105: Longitudinal and angular segmentation of the HE calorimeter. The dashed lines point to the interaction point.

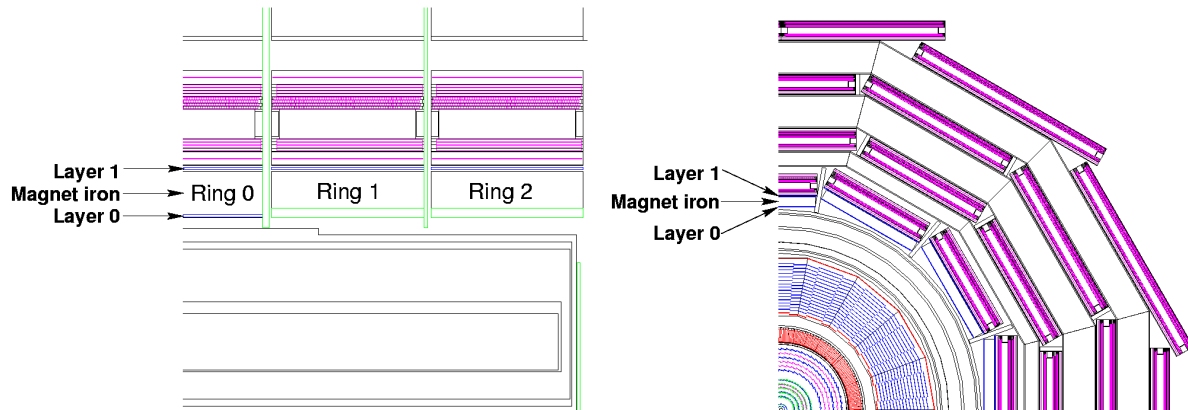


Figure 106: Longitudinal and transverse views of the CMS detector showing the position of HO layers.

tiles located in front of the first layer of the barrel muon detector. Scintillation light from the tiles is collected using multi-clad Y11 Kuraray wavelength shifting (WLS) fibres of diameter 0.94 mm, and transported to the photo detectors located on the structure of the return yoke by splicing a multi-clad Kuraray clear fibre (also of 0.94 mm diameter) with the WLS fibre. In order to simplify installation of HO, the scintillator tiles are packed into a single unit called a tray. Each tray corresponds to one ϕ slice (5° wide in ϕ). However, along the z (η) direction, a tray covers the entire span of a muon ring. Figure 107 shows a schematic view of a HO tray where one tile is mapped to a tower of HB and the optical cable from the tray is connected to the read-out box.

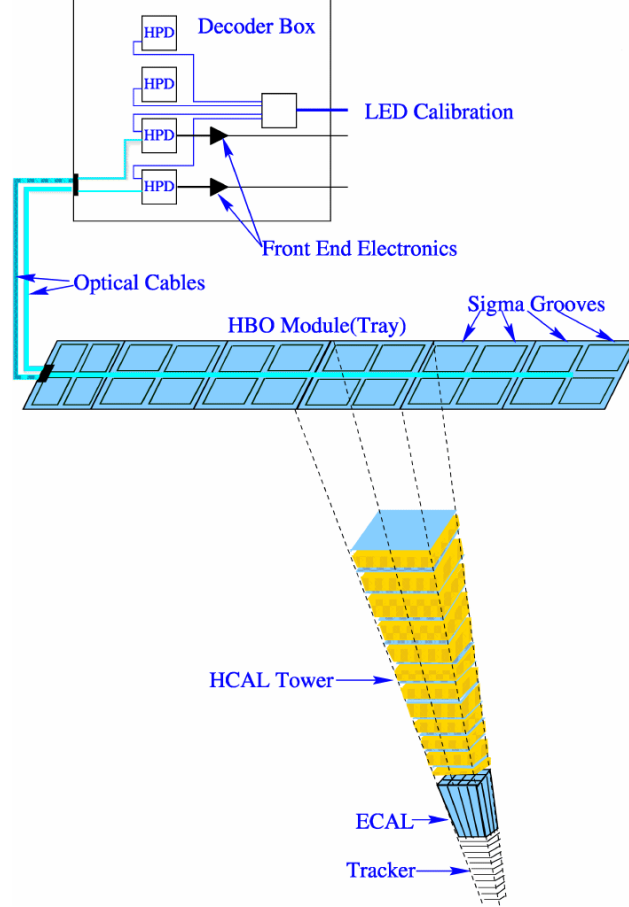


Figure 107: Schematic view of a HO tray shown with individual tiles and the corresponding grooves for WLS fibres. Each optically independent (4 WLS fibres) tile is mapped to a tower of HB. Optical fibres from the tray are routed to the decoder box which contains the photodetector and read-out electronics.

The physics impact of HO has been studied [114] using a simulation of the CMS detector. Single pions of fixed energies are shot at specific η values and the resulting energy deposits in the electromagnetic calorimeter and in the layers of the hadron calorimeter are combined to measure the energy. Figure 108 shows distributions of the measured energy scaled to the incident energy for 200 GeV pions at $\eta = 0$ and 225 GeV at $\eta = 0.5$ (pointing towards the middle of ring 1). The solid and dashed lines in the figure indicate measurements without and with HO, respectively. As can be seen in Fig. 108, there is an excess in $\text{Energy}/E_{\text{incident}} < 1$ for measurements without HO, because of leakage. The measurements with HO are more Gaussian in nature indicating that the addition of HO recovers the effect of leakage. The effect of leakage is visible at $\eta = 0$ (ring 0) from 70 GeV, increasing with energy. The mean fraction of energy in HO increases from 0.38% for 10 GeV pions to 4.3% for 300 GeV pions. There is some evidence of leakage without HO in ring 1 but it is reduced due to the greater HB thickness at larger $|\eta|$. The amount of leakage in ring 2 is found to be negligible at energies below 300 GeV.

The effect of shower leakage has a direct consequence on the measurement of missing transverse energy (E_T^{miss}). Study of QCD events shows that the cross section for those events, where at least one particle has E_T above 500 GeV, is several pb. For these events the HO is useful to decrease the leakage and improve the E_T^{miss} measurement. Figure 109 shows the dijet integrated cross section for E_T^{miss} above a certain value. It is clear from the figure that the inclusion of HO reduces the dijet rate by a factor of 1.5 or more for moderate E_T^{miss} values, a region important

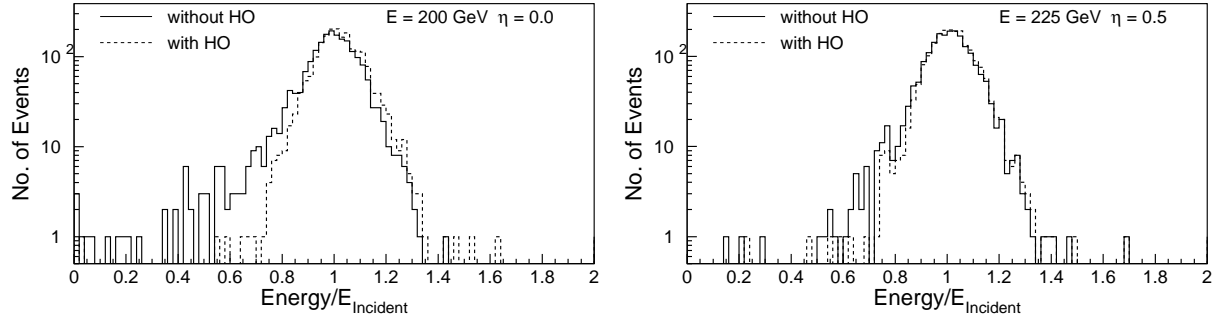


Figure 108: A simulation of the distribution of the measured energy scaled to the incident energy for pions with incident energies of (left panel) 200 GeV at $\eta = 0$ and (right panel) 225 GeV at $|\eta| = 0.5$. The solid and dashed histograms are measurements without and with HO, respectively.

for searches of supersymmetric particles.

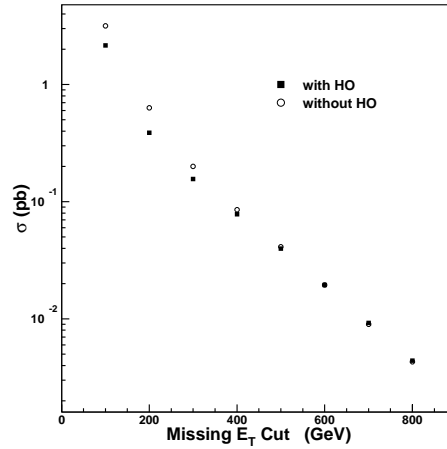


Figure 109: Integrated cross section above threshold for intrinsically balanced QCD dijet events as a function of missing E_T with or without HO.

Module specification

HO is physically divided into 5 rings in η conforming to the muon ring structure. The rings are numbered -2 , -1 , 0 , $+1$ and $+2$ with increasing η . Each ring of the HO is divided into 12 identical ϕ sectors and each sector has 6 slices (numbered 1 to 6 counting clockwise) in ϕ . The ϕ slices of a layer are identical in all sectors. The widths of the slices along ϕ are given in Table 11. In each ϕ slice, there is a further division along η . The smallest scintillator unit in HO thus obtained is called a tile. The scintillator tiles in each ϕ sector belong to a plane. The perpendicular distance of this plane from the z -axis is 3.82 m for layer 0 and 4.07 m for layer 1. The tiles in each ϕ slice of a ring are mechanically held together in the form of a tray.

Table 11: Dimension of tiles along ϕ for different trays. Each tray corresponds to one ϕ -slice in a ϕ sector.

Ring	Layer	Width along ϕ in mm					
		Tray 1	Tray 2	Tray 3	Tray 4	Tray 5	Tray 6
0	0	274	343	332	327	327	268
0	1	300	364	352	347	347	292
$\pm 1, \pm 2$	1	317	366	354	349	349	406

Both layers of ring 0 have 8 η -divisions (i.e. 8 tiles in a tray): -4 , -3 , -2 , -1 , $+1$, $+2$, $+3$, $+4$. Ring 1 has 6 divisions: $5 \cdots 10$ and ring 2 has 5 divisions: $11 \cdots 15$. Ring -1 and ring -2 have the same number of divisions as rings 1 and 2 but with $-ve$ indices. The η -dimension of any tile with $-ve$ tower number is the same as the one

with +ve number. The tile dimensions along η are shown in Table 12.

Table 12: HO tile dimensions along η for different rings and layers. The tile sizes, which are constrained by muon ring boundaries, are also given.

Tower #	η_{\max}	Length (mm)	Tower #	η_{\max}	Length (mm)
Ring 0 Layer 0			Ring 0 Layer 1		
1	0.087	331.5	1	0.087	351.2
2	0.174	334.0	2	0.174	353.8
3	0.262	339.0	3	0.262	359.2
4	0.326	248.8	4	0.307	189.1
Ring 1 Layer 1			Ring 2 Layer 1		
5	0.436	391.5	11	0.960	420.1
6	0.524	394.2	12	1.047	545.1
7	0.611	411.0	13	1.135	583.3
8	0.698	430.9	14	1.222	626.0
9	0.785	454.0	15	1.262	333.5
10	0.861	426.0			

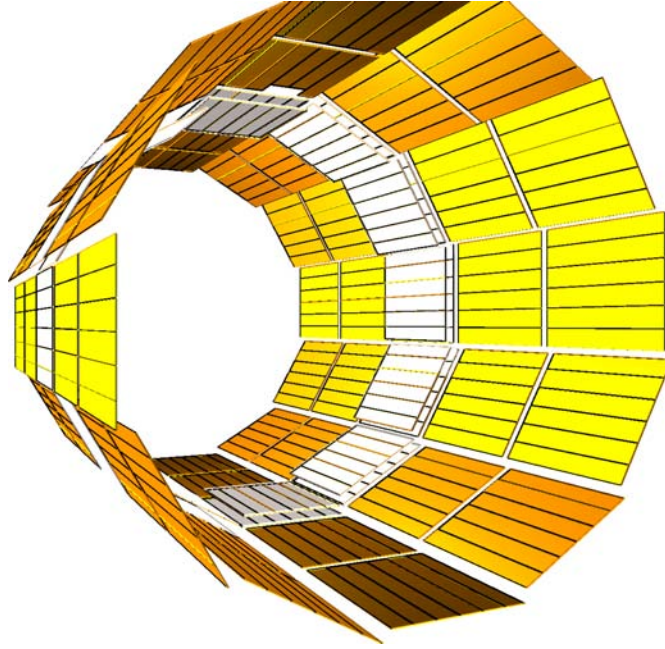


Figure 110: Layout of all the HO trays in the CMS detector.

Figure 110 shows the final layout of all the HO trays in the CMS detector. The length of a full tray is 2510 mm whereas the shorter trays, the sizes of which are constrained because of the chimney (trays 4 and 5 in sector 4 of ring +1 and trays 3, 4, 5 and 6 in sector 3 of ring -1), are 2119-mm long. The shorter trays are constructed without the tile corresponding to tower number ± 5 . Because of the constraints imposed by the gap between ring 0 and rings ± 1 , the η boundaries of HO tower 4 do not match the barrel η boundaries; therefore, part of HO tower 5 overlaps with tower 4 in the barrel.

Tiles

Scintillator tiles are made from Bicron BC408 scintillator plates of thickness 10^{+0}_{-1} mm. Figure 111 shows a typical HO scintillator tile. The WLS fibres are held inside the tile in grooves with a key hole cross section. Each groove

has a circular part (of diameter 1.35 mm) inside the scintillator and a neck of 0.86 mm width. The grooves are 2.05-mm deep. Each tile has 4 identical grooves, one groove in each quadrant of the tile. The grooves closely follow the quadrant boundary. The corners of the grooves are rounded to prevent damage to the fibre at the bend and to ease fibre insertion. The groove design is slightly different for the tile where the optical connector is placed at the end of the tray. Since the tiles are large, 4 grooves ensure good light collection and less attenuation of light.



Figure 111: View of a typical tile of HO with WLS fibres inserted in the 4 grooves of the tile.

The HO has 95 different tile dimensions, 75 for layer 1 and 20 for layer 0. The total number of tiles is 2730 (2154 for layer 1 and 576 for layer 0).

Trays

All tiles in each ϕ slice of a sector are grouped together in the form of a tray. Each tray contains 5 tiles in rings ± 2 ; 6 tiles in rings ± 1 and 8 tiles in ring 0. The edges of the tiles are painted with Bicon reflecting white paint for better light collection as well as isolating the individual tiles of a tray. Further isolation of tiles is achieved by inserting a piece of black tedlar in between the adjacent tiles. The tiles in a tray are covered with a single big piece of white, reflective tyvek paper. Then they are covered with black tedlar paper to prevent light leakage. This package is placed between two black plastic plates for mechanical stability and ease of handling. The top plastic cover is 2-mm-thick and the bottom one is 1-mm-thick. Figure 112 shows a cross section of a tray to illustrate the different components. The plastic covers (top and bottom) have holes matching with the holes in the tiles. Specially designed countersunk screws passing through these holes fix the plastic covers firmly on the tiles.

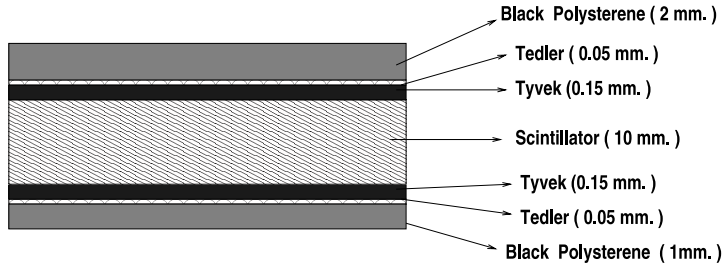


Figure 112: Cross section of a HO tray showing the different components.

The 2 mm plastic sheet on the top has 1.6 mm deep channels grooved on it (on the outer side) to route the fibres from individual tiles to an optical connector placed in a groove at the edge of the tray. A 1.5-mm-wide straight groove runs along the edge of the top cover to accommodate a stainless steel tube. This is used for the passage of a radioactive source which is employed in calibrating the modules. Each connector has two holes and they are fixed to the scintillator-plastic assembly through matching holes. Each ϕ sector in each ring has 6 trays. There are 360 trays for layer 1 and 72 trays for layer 0.

Pigtails

The light collected by the WLS fibres inserted in the tiles needs to be transported to photodetectors located far away on the muon rings. The captive ends of the WLS fibres, which reside inside the groove, are polished, aluminized and protected using a thin polymer coating. The other end of the WLS fibre comes out of the tile through a slot made on the 2-mm-thick black plastic cover sheet. To minimise the loss of light in transportation, the WLS fibre (attenuation length of ≈ 1.8 m) is spliced to a clear fibre (attenuation length of ≈ 8.0 m). A fibre is spliced only if

the potential WLS light loss is larger than the light loss at a spliced joint. Thus depending on tile length (along η) 2–3 fibres in each pigtail are made only of WLS fibres. The clear fibres from each tile follow the guiding grooves on the top plastic to the optical connector at the end. Each tray has two optical connectors mounted on one end of the tray. In a tray, the grooves of the tiles form two rows along η . The fibres from all grooves on one row terminate on one connector (Fig. 113). The number of fibres from trays in different rings are given in Table 13.

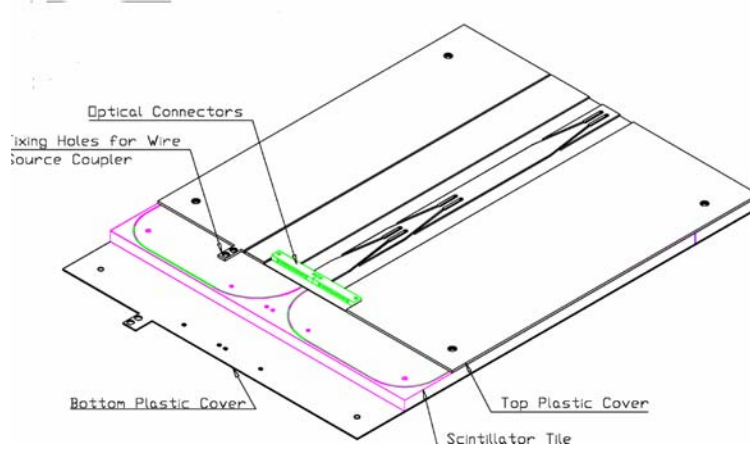


Figure 113: The arrangement of scintillation tiles, plastic covers and connectors in a tray. The components are slightly displaced from their true positions to show their matching designs.

Table 13: Tray specifications for different rings of HO

Ring #	Tiles/tray	Fibres/tray	Fibres/connector
0	8	32	16
± 1	6	24	12
± 2	5	20	10

The bunch of fibres fixed to the optical connector is called a pigtail (Fig. 114). Each tray has 2 pigtails and there are 864 pigtails in total: 720 for layer 1 and 144 for layer 0. Each fibre in a pigtail is cut to the proper length to match the groove length in the scintillator plus the distance from the scintillator to the optical connector at the end of the tray.

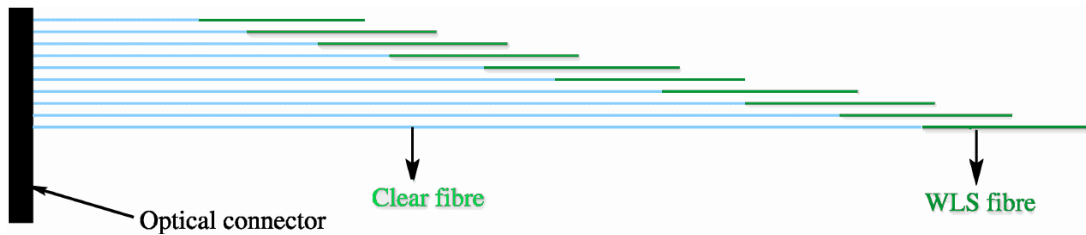


Figure 114: Illustration of an assembled pigtail (not drawn to scale).

6.4 Forward calorimeter design (HF)

The forward calorimeter will experience unprecedented particle fluxes. On average, 760 GeV per proton-proton interaction is deposited into the two forward calorimeters, compared to only 100 GeV for the rest of the detector. Moreover, this energy is not uniformly distributed but has a pronounced maximum at the highest rapidities. At $|\eta| = 5$ after an integrated luminosity of $5 \times 10^5 \text{ pb}^{-1}$ (≈ 10 years of LHC operation), the HF will experience $\approx 10 \text{ MGy}$. The charged hadron rates will also be extremely high. For the same integrated luminosity, inside the HF absorber at 125 cm from the beam-line, the rate will exceed 10^{11} per cm^2 [107]. This hostile environment presents a considerable challenge to calorimetry, and the design of the HF calorimeter was first and foremost guided by the necessity to survive in these harsh conditions, preferably for at least a decade. Successful operation critically depends on the radiation hardness of the active material. This was the principal reason why quartz fibres

(fused-silica core and polymer hard-cladding) were chosen as the active medium.

The signal is generated when charged shower particles above the Cherenkov threshold ($E \geq 190$ keV for electrons) generate Cherenkov light, thereby rendering the calorimeter mostly sensitive to the electromagnetic component of showers [115]. A small fraction of the generated light is captured, $f_{\text{trap}} = \text{NA}/2n_{\text{core}}^2$, where NA is the numerical aperture ($\text{NA} = 0.33 \pm 0.02$) and n_{core} is the refractive index of the quartz core. Only light that hits the core-cladding interface at an angle larger than the critical angle (71°) contributes to the calorimeter signal. The half-angle $\theta = 19^\circ$ is determined by the refractive indices of the core (n_{core}) and the cladding (n_{clad}), $\sin \theta = \sqrt{n_{\text{core}}^2 - n_{\text{clad}}^2}$. The fibres measure $600 \pm 10 \mu\text{m}$ in diameter for the fused-silica core, $630_{-10}^{+5} \mu\text{m}$ with the polymer hard-cladding, and $800 \pm 30 \mu\text{m}$ with the protective acrylate buffer. Over 1000 km of fibres are used in the HF calorimeters. The fibres are cleaved at both ends by a diamond cutter. The attenuation length of these fibres is measured to be ≈ 15 m using high energy electrons at 90° to the fibres.

The optical attenuation at a wavelength λ in these types of fibres scales as $a(\lambda)(D/D_0)^{b(\lambda)}$ where D is the accumulated dose, which is normalized to a reference dose ($D_0 = 1$ MGy) for convenience. For example, at a wavelength $\lambda = 450$ nm at the accumulated dose of $D = 1$ MGy, the induced attenuation is ≈ 1.5 dB/m, thus defining a . The a and b parameters characterize the radiation hardness of a given fibre. For high OH^- (300–500 ppm) HF fibres at 450 nm, the measured values are $a \approx 1.5$ and $b \approx 0.3$ [116, 117, 118]. An accumulated dose of 10 MGy will result in a loss of optical transmission by a half, which is the worst case for HF after a decade.

The calorimeter consists of a steel absorber structure that is composed of 5 mm thick grooved plates. Fibres are inserted in these grooves. The detector is functionally subdivided into two longitudinal segments. Half of the fibres run over the full depth of the absorber ($165 \text{ cm} \approx 10\lambda_I$) while the other half starts at a depth of 22 cm from the front of the detector. These two sets of fibres are read out separately. This arrangement makes it possible to distinguish showers generated by electrons and photons, which deposit a large fraction of their energy in the first 22 cm, from those generated by hadrons, which produce nearly equal signals in both calorimeter segments on average. The long fibre section is referred as L (measuring the total signal), and the short fibre section as S (measuring the energy deposition after 22 cm of steel). The absorber has grooves ($0.90_{-0}^{+0.12}$ mm wide and $1.06_{-0}^{+0.6}$ mm in depth) which make a square grid separated by 5.0 ± 0.1 mm center-to-center. Long and short fibres alternate in these grooves. The packing fraction by volume (fibre/total) in the first 22 cm is 0.57% and is twice as large beyond that depth.

The forward calorimeter is essentially a cylindrical steel structure with an outer radius of 130.0 cm. The front face of the calorimeter is located at 11.2 m from the interaction point. The hole for the beam pipe is cylindrical, with radius 12.5 cm from the center of the beam line. This structure is azimuthally subdivided into 20° modular wedges. Thirty-six such wedges (18 on either side of the interaction point) make up the HF calorimeters. A cross sectional view of the HF is shown in Fig. 115. The fibres run parallel to the beam line, and are bundled to form 0.175×0.175 ($\Delta\eta \times \Delta\phi$) towers (Fig. 116 and Table 14). The detector is housed in a hermetic radiation shielding which consists of layers of 40 cm thick steel, 40 cm of concrete, and 5 cm of polyethylene. A large plug structure in the back of the detector provides additional shielding.

Bundled fibres are held in ferrules which illuminate one end of the air-core light guides that penetrate through 42.5 cm of the shielding matrix (steel, lead, and polyethylene). This shielding is necessary to protect the photo-multipliers and the front-end electronics housed in the read-out boxes. The air-core light guide consists of a hollow tube lined on the inside with highly reflective custom-made sheets. These metal-coated reflectors are designed to be very efficient ($> 90\%$) in the visible spectrum at the relevant angles (≈ 70 degrees from normal). Light typically makes five bounces before reaching the photocathode and nearly half the light is lost in this transport. Each light guide is coupled to a standard bialkaline, 8-stage photomultiplier tube with a borosilicate glass window. A read-out box (RBX) houses 24 PMTs and services half of a wedge (10° in ϕ).

The entire calorimeter system, with its shielding components, is mounted on a rigid table which supports more than 240 t with less than 1 mm deflection. The absorber alone weighs 108 t. The table is also designed for horizontal separation of the detector into two sections to clear the beam pipe at installation and removal. It is possible to align the forward calorimeters within ± 1 mm with respect to the rest of the CMS experiment.

6.5 Read-out electronics and slow control

The overview of the full HCAL read-out chain is shown in Fig. 117. The read-out consists of an optical to electrical transducer followed by a fast charge-integrating ADC. The digital output of the ADC is transmitted for every bunch

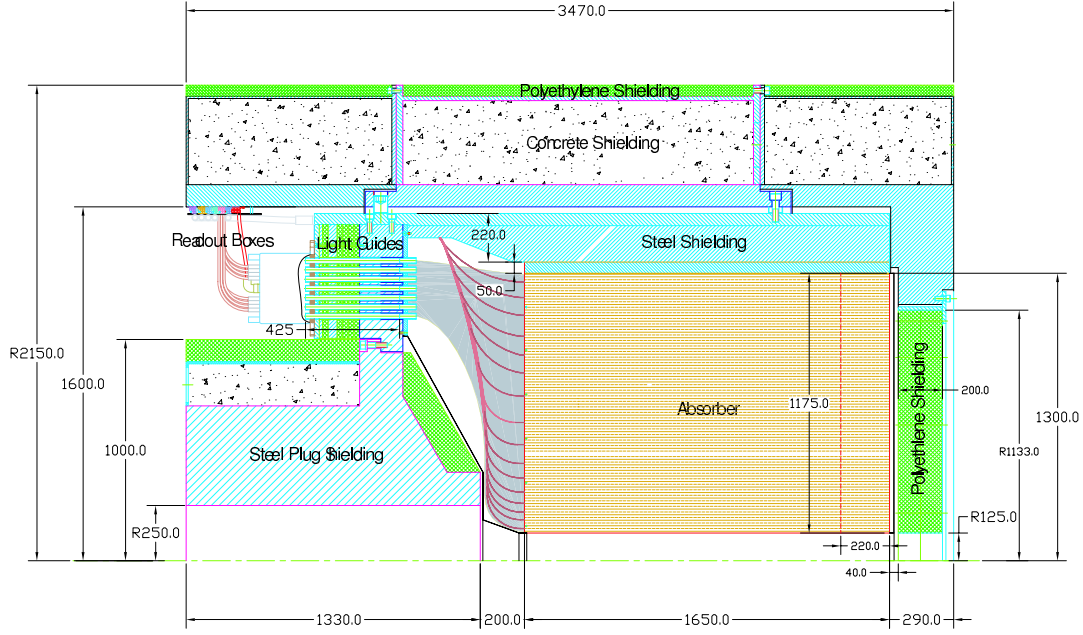


Figure 115: The cross sectional view of the HF calorimeter shows that the sensitive area extends from 125 to 1300 mm in the radial direction. The absorber in the beam direction measures 1650 mm. Bundled fibres (shaded area) are routed from the back of the calorimeter to air-core light guides which penetrate through a steel-lead-polyethylene shielding matrix. Light is detected by PMTs housed in the read-out boxes. Stainless steel radioactive source tubes (red lines) are installed for each tower and are accessible from outside the detector for source calibration. The interaction point is at 11.2 meters from the front of the calorimeter to the right. All dimensions are in mm.

Table 14: The tower sizes, number of fibres, bundle sizes and the percentage of photocathode area utilized are listed below for each tower. The air-core light guides are tapered to better match the photocathode area for towers 1, 2 and 3.

Ring No	(r_{in}, r_{out}) [mm]	$\Delta\eta$	$\Delta\phi$ [degree]	N_{fib}	A_{bundle} [mm ²]	$\frac{A_{bundle}}{A_{photocathode}}$
1	(1162–1300)	0.111	10	594	551	1.14
2	(975–1162)	0.175	10	696	652	1.33
3	(818–975)	0.175	10	491	469	0.96
4	(686–818)	0.175	10	346	324	0.66
5	(576–686)	0.175	10	242	231	0.47
6	(483–576)	0.175	10	171	167	0.34
7	(406–483)	0.175	10	120	120	0.25
8	(340–406)	0.175	10	85	88	0.18
9	(286–340)	0.175	10	59	63	0.13
10	(240–286)	0.175	10	41	46	0.94
11	(201–240)	0.175	10	30	35	0.71
12	(169–201)	0.175	20	42	52	0.11
13	(125–169)	0.300	20	45	50	0.10

over a gigabit digital optical fibre to the counting house. In the counting house, the signal is deserialized and used to construct trigger primitives which are sent to the calorimeter trigger. The data and trigger primitives are also pipelined for transmission to the DAQ upon a Level-1 Accept (L1A) decision.

The optical signals from the scintillator-based detectors (HB/HE/HO) are converted to electrical signals using multichannel hybrid photodiodes (HPDs) which provide a gain of ≈ 2000 . A detailed view of the scintillator-based

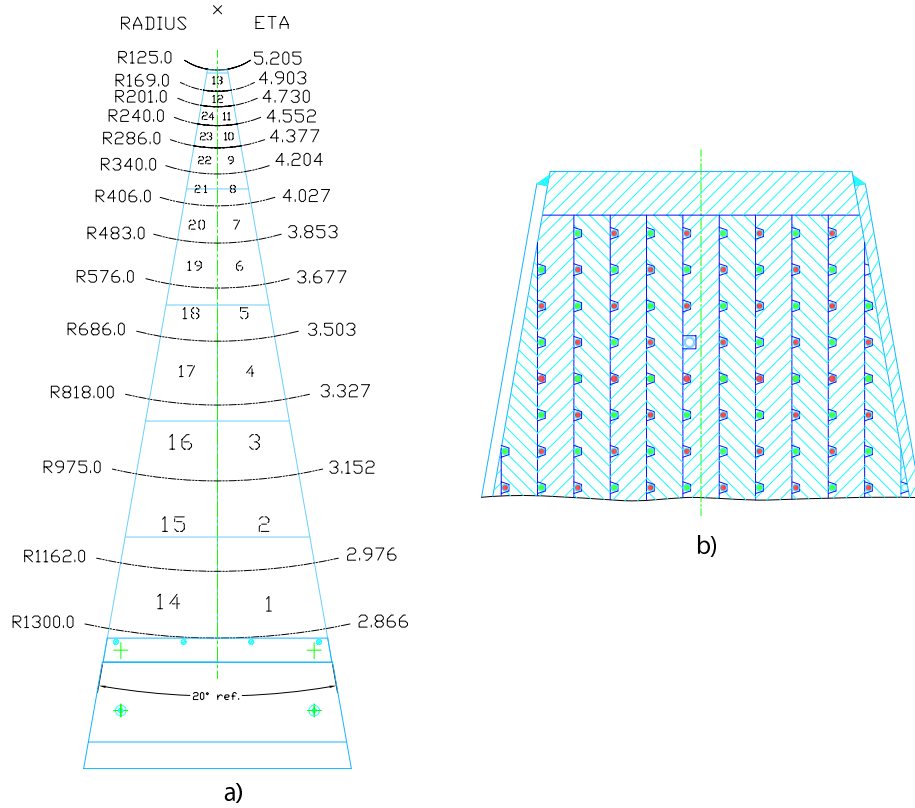


Figure 116: a) Transverse segmentation of the HF towers. b) An expanded view of the wedge shows the squared out groove holding the radioactive source tube.

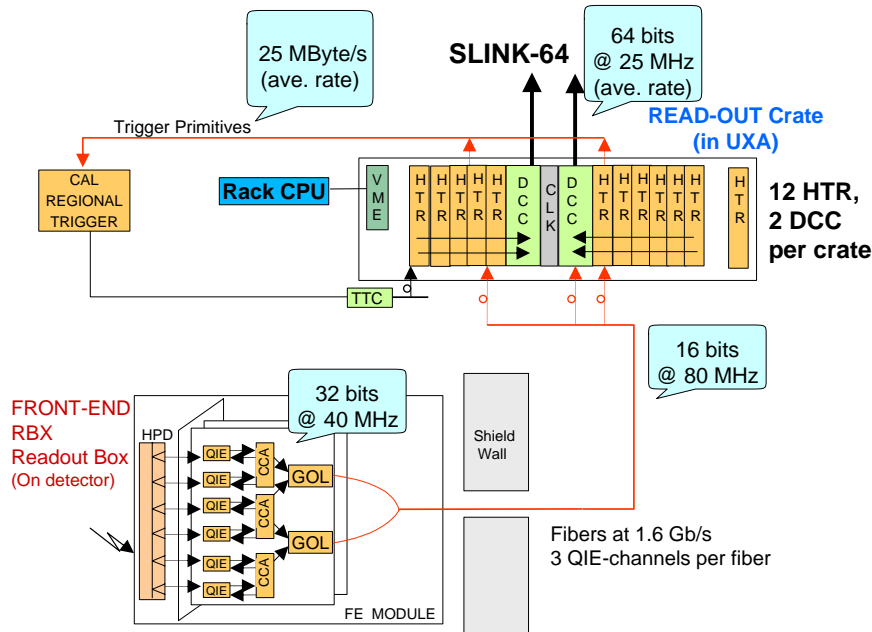


Figure 117: Overview of HCAL read-out electronics.

front-end read-out chain is given in Fig. 118. The optical signals from individual sampling layers are brought out on clear fibres. The fibres corresponding to a projective calorimeter tower are mapping via an optical decoding unit (ODU) to a cookie that interfaces to individual pixels on the HPD. In the forward calorimeter, where the magnetic fields are much smaller than in the central detector, conventional photomultiplier tubes (Hamamatsu

R7525HA) are used and quartz-fibre bundles are routed directly to the phototube windows.

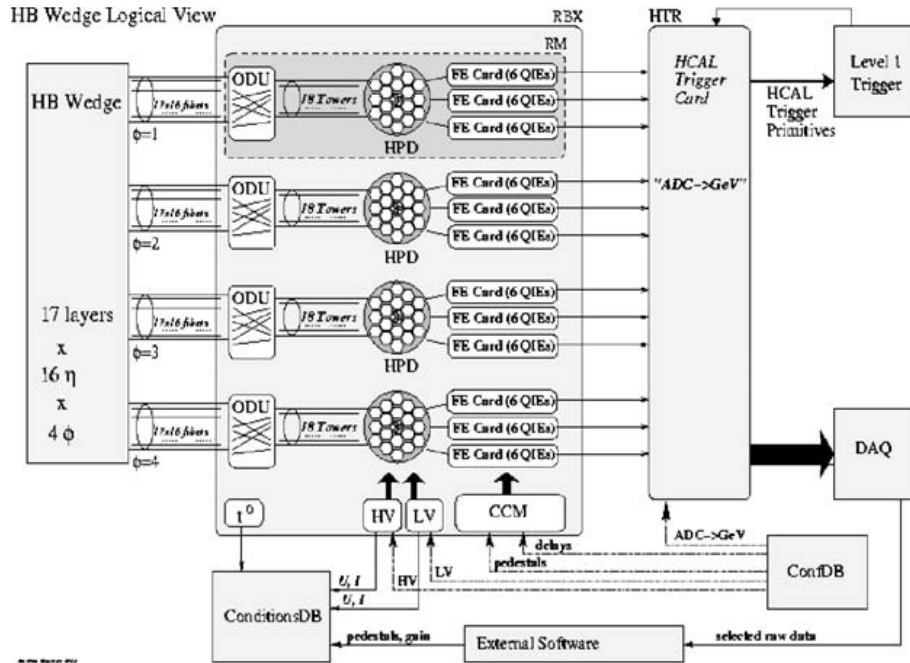


Figure 118: Overview of HCAL read-out/trigger chain and connections to database.

An overview of the HCAL controls is given in Fig. 119. Several PCs in the CMS control room operated through PVSS are used to control high and low voltages. The control system also downloads pedestal DAC and timing parameters to front-ends and controls many of the calibration and monitoring systems including the source calibration drivers, the LED pulsers, and the laser system. These systems record temperature, humidity and other constants useful for correlation studies of detector/calibration stability.

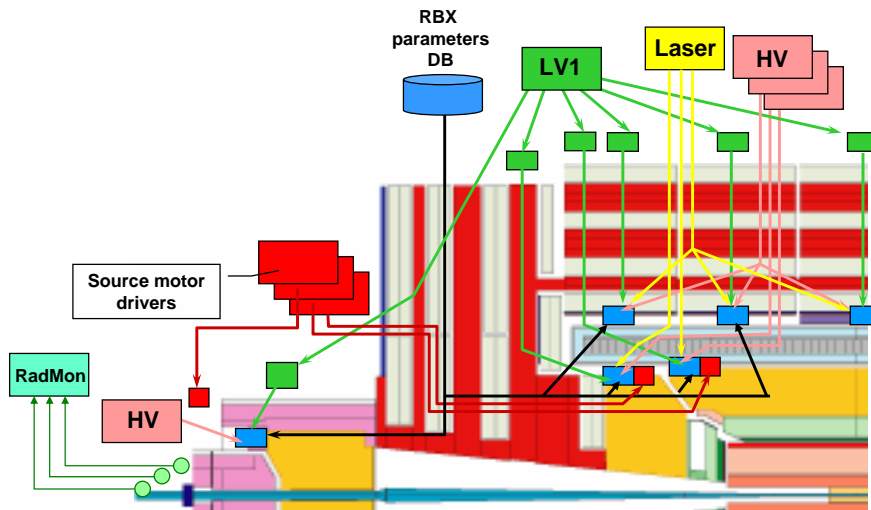


Figure 119: Overview of HCAL detector controls.

The configuration database contains the relationships or mapping for all HCAL detector components: wedges, layers, read-out boxes (RBX), cables, HCAL Trigger (HTR) cards, and calibration parameters for various components e.g. RBX, QIE, source types and strength. The conditions database has the slow-controls logging, the calibration

constants (pedestals, gains, timing information, etc.) and the configuration database downloaded to the read-out system during the initialization.

The analogue signal from the HPD or photomultiplier is converted to a digital signal by a charge-integrating ADC ASIC called the QIE (Charge-Integrator and Encoder). The QIE internally contains four capacitors which are connected in turn to the input, one during each 25 ns period. The integrated charge from the capacitors is converted to a seven-bit non-linear scale to cover the large dynamic range of the detector. The ADC is designed so its contribution to the detector energy resolution over its multi-range operation is negligible, as shown in Fig. 120. The QIE input characteristics were chosen from test beam data to optimize speed and noise performance. This resulted in a per channel RMS noise of 4600 electrons (0.7 fC) corresponding to about 180 MeV.

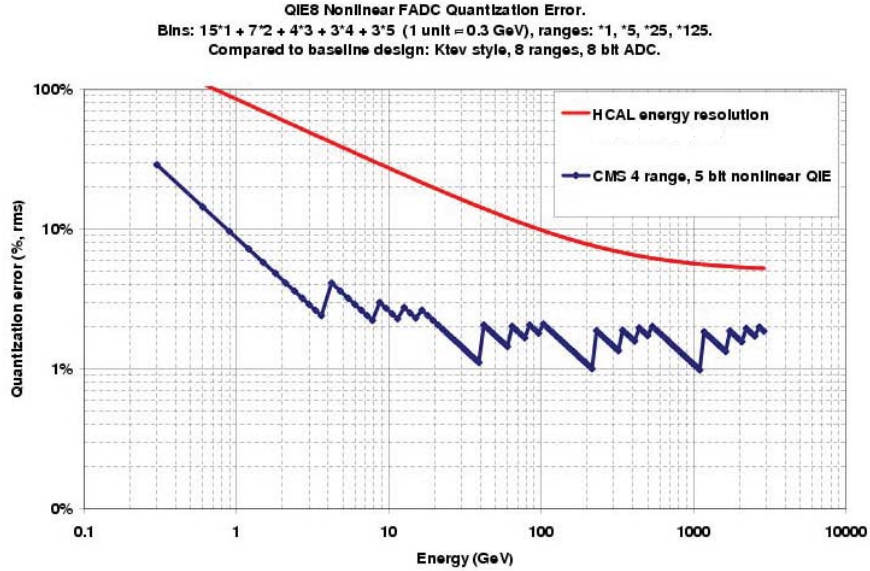


Figure 120: Contribution of the FADC quantization error to the resolution, compared with a representative HCAL resolution curve.

The digital outputs of three QIE channels are combined with some monitoring information to create a 32-bit data word. The 32-bit data, at a rate of 40 MHz, is fed into the Gigabit Optical Link (GOL) chip and transmitted using 8b/10b encoding off the detector to the counting house. In the counting house, the data is received by the HCAL Trigger/Read-out (HTR) board. The HTR board contains the Level-1 pipeline and also constructs the trigger primitives for HCAL. These trigger primitives are sent to the Regional Calorimeter trigger via Serial Link Board mezzanine cards. The HTR board receives data for 48 channels (16 data fibres) and may host up to six SLBs.

When a L1A is received by the HTR through the TTC system, it prepares a packet of data for the DAQ including a programmable number of precision read-out values and trigger primitives around the triggered bunch crossing. For normal operations, the HTR will transmit 7 time samples for each non-zero channel and a single trigger primitive for every trigger tower. These packets of data, each covering 24 channels, are transmitted by LVDS to the HCAL Data Concentrator Card (DCC). The DCC is the HCAL Front-End Driver (FED) and concentrates the data from up to 360 channels for transmission into the DAQ.

The Level-1 trigger primitives (TPG) are calculated in the HTR modules. The QIE data are linearized and converted to transverse energy with a single look up table. Two or more consecutive time samples are summed. A sum over depth is made for those towers having longitudinal segmentation. A final look up table is used to compress the data before sending the data across the trigger link to the regional calorimeter trigger. Table 15 summarizes the geometry of the trigger towers. The HF towers are summed in η and ϕ before being sent to the trigger.

Timing and synchronization

The QIE integration clock is controlled by the Channel Control ASIC (CCA) which allows for fine-skewing of the integration phase of each tower relative to the machine clock. This allows each channel's integration phase to correct for differences in the time-of-flight from the interaction region as well as differences in the optical pathlength within the detector.

Tower index	$ \eta_{max} $	Detector	Size	
			η	ϕ
1–15	$0.087 \times \eta$	HB	0.087	5°
16	1.392	HB, HE	0.087	5°
17–20	$0.087 \times \eta$	HE	0.087	5°
21	1.830	HE	0.090	5°
22	1.930	HE	0.100	5°
23	2.043	HE	0.113	5°
24	2.172	HE	0.129	5°
25	2.322	HE	0.150	5°
26	2.500	HE	0.178	5°
27	2.650	HE	0.150	5°
28	3.000	HE	0.350	5°
29	3.314	HF	0.461	20°
30	3.839	HF	0.525	20°
31	4.363	HF	0.524	20°
32	5.191	HF	0.828	20°

Table 15: Sizes of the HCAL trigger towers in η and ϕ .

Figure 121 shows that scintillator tile signals produce relatively fast pulses such that 68% of the pulse is contained within a 25 ns window. Figure 122 shows the pulse shape for the forward calorimeter. The Cerenkov process and the phototubes used in the forward calorimeter are extremely fast, so the pulse in HF is only 10 ns wide. The HF is thus subject only to in-time pile-up which is important in the highly active forward region of CMS.

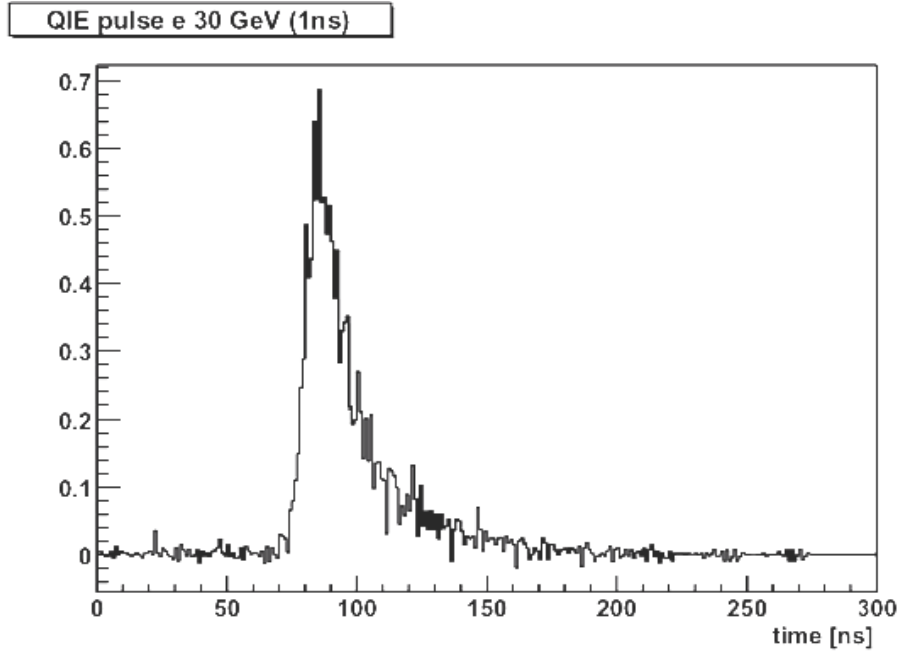


Figure 121: Measured single event pulse shape from the scintillator tiles, representative of HB/HE/HO pulse shapes.

An additional important effect on the HCAL pulse timing in HB/HE/HO comes from the input stage of the QIE. The QIE has an amplitude-dependent impedance which implies a faster pulse shape for large signals than for small ones, as seen in Fig. 123. The amount of time slewing is dependent on the noise characteristics of the QIE, so the final QIE ASICs for the barrel and endcap were chosen to limit the timeslew to the “medium” case in exchange for somewhat increased noise. In the outer calorimeter, the noise level is a critical factor for muon identification and pile-up is much less important so the quieter “slow” characteristics were chosen for the HO QIEs.

The in-situ synchronization of HCAL is performed using the HCAL laser system. The laser system consists of a single UV laser which can illuminate an entire half-barrel of HB or a single endcap at once through a series

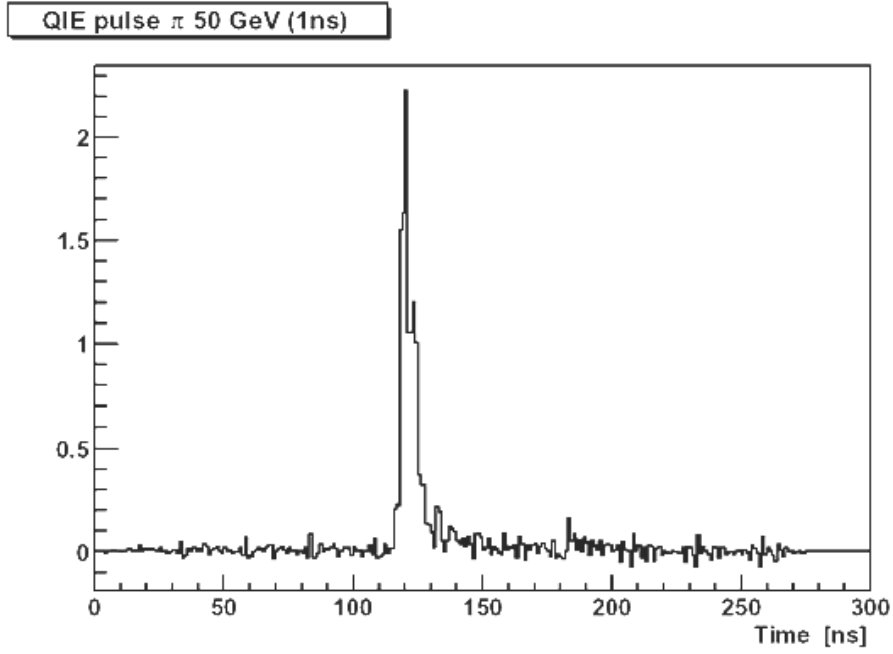


Figure 122: Measured pulse shape, energy collected vs. time, for HF.

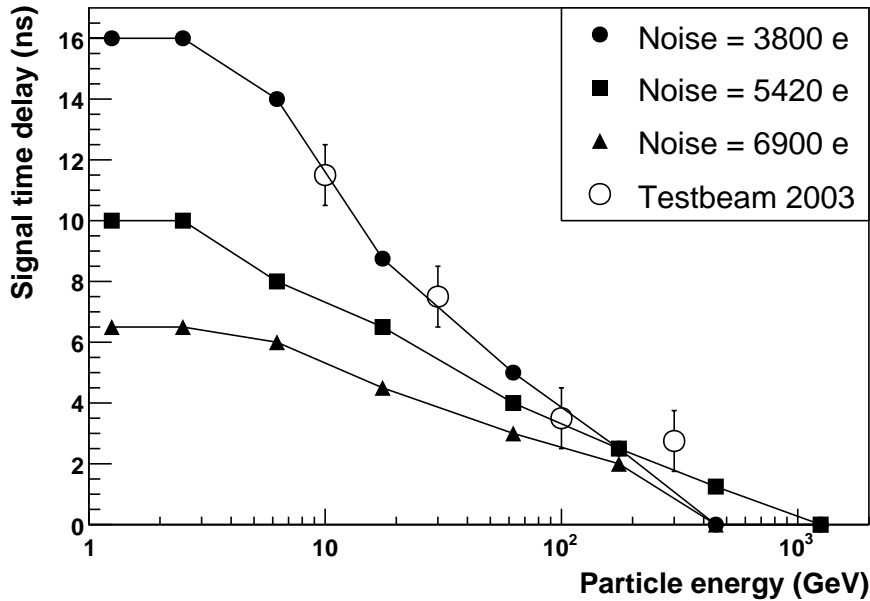


Figure 123: Pulse time variation as a function of signal amplitude as measured on the bench (solid points) for several input amplifier configurations compared with test beam measurements from 2003.

of optical splitters. The quartz fibres which lead from the laser to the detector have been carefully controlled to equalize the optical path length to each wedge. The laser can be directed either straight onto a scintillator block connected to the HPD or into the wedge. Within layer 9 of each wedge is an arrangement of optical fibres which mimic the time-of-flight from the interaction region. This arrangement allows the timing of HCAL to be flattened and monitored, as has been demonstrated in test beam data taking, which verified the timing determined by the laser using the synchronized beam. In the HO and HF detectors, only the photodetector can be illuminated so the alignment will be based on construction and test beam data.

The channel-by-channel bunch synchronization of HCAL will be determined using a histogramming procedure in the serial link boards (SLBs) which determine the bunch synchronization using the beam structure of the LHC. The event and bunch synchronization is monitored using fast control signals originating from the TTC system which are transmitted in the data stream between the front-ends and the HTR. On a global scale, the bunch and event

synchronization between the HCAL and other detector subsystems is determined using muons and other correlated physics signals.

6.6 HF luminosity monitor

The CMS luminosity measurement will be used to monitor the LHC's performance on a bunch-by-bunch basis in real time and to provide an overall normalization for physics analyses. The design goal for the real-time measurement is to determine the average luminosity with a 1% statistical accuracy with an update rate of 1 Hz. For offline analyses, the design goal is a systematic accuracy of 5%, although every reasonable effort will be made to produce a more accurate result. Both of these requirements must be met over a very large range of luminosities, extending from roughly $10^{28} \text{ cm}^{-2}\text{s}^{-1}$ to $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and possibly beyond.

A number of techniques capable of providing suitable luminosity information in real time have been identified [17]. One technique employs signals from the forward hadron calorimeter (HF) while another, called the Pixel Luminosity Telescope (PLT), uses a set of purpose-built particle tracking telescopes based on single-crystal diamond pixel detectors. At the time of writing, the PLT has not been formally approved, but is under study. The methods based on signals from the HF are the ones being most vigorously pursued.

Two methods for extracting a real-time relative instantaneous luminosity with the HF have been studied. The first method is based on *zero counting* in which the average fraction of empty towers is used to infer the mean number of interactions per bunch crossing. The second method exploits the linear relationship between the average transverse energy per tower and the luminosity.

Outputs of the QIE chips used to digitize the signals from the HF PMTs on a bunch-by-bunch basis are routed to a set of 36 HCAL Trigger and Read-out (HTR) boards, each of which services 24 HF physical towers. In order to derive a luminosity signal from the HTR, an additional mezzanine board called the HF luminosity transmitter (HLX) is mounted on each of the HTR boards. The HLX taps into the raw HF data stream and collects channel occupancy and E_T -sum data and transmits them to a central collector node over standard 100-Mbps Ethernet. The HLX boards have the same form factor as the Synchronization and Link Boards (SLBs) used to interface the ECAL and HCAL readouts to the Regional Calorimeter Trigger (RCT) system.

Although all HF channels can be read by the HLX, MC studies indicate that the best linearity for occupancy histograms is obtained using just two η rings. Hence two sets of two rings are used for the occupancy histograms. Four rings are combined to form the E_T -sum histogram. The algorithm has been optimized to minimize sensitivity to pedestal drifts, gain changes and other related effects. Each of the two sets of rings sends 12 bits of data to the HLX. There are three occupancy histograms dedicated to each of the following possible states for each tower: enabled-below-threshold, over-threshold-1, over-threshold-2. In addition, a 15-bit E_T sum value is sent to the HLX and a further histogram based on all 13 HF η rings is filled for use by the LHC. As a result, the input to the HLX is used to create eight histograms: two sets of three occupancy histograms, one E_T -sum histogram, and one additional occupancy histogram.

Each histogram has 3564 bins, one for each bunch in the LHC orbit. Each occupancy-histogram bin uses two bytes, and there are four bytes per bin in the E_T sum histogram. The baseline design is to add the results from all desirable channels into a single set of histograms.

The histograms are transmitted as UDP packets from the HLX cards once roughly every 0.37 s, which is safely within the 1.45 s (worst case) histogram overflow time. The Ethernet core in the HLX automatically packages the data to make optimal use of network bandwidth. Each histogram spans several Ethernet packets, the precise number depending on the type of histogram. The eight sets of histograms comprise about 70 kB of data, which is transmitted at a rate of approximately 1.6 Mbps to an Ethernet switch that aggregates the data from multiple HLX boards. The switch multi-casts the data to a pair of luminosity server nodes. One of the servers is responsible for publishing the luminosity information to various clients, such as the CMS and LHC control rooms and the Fermilab Remote Operations Center (ROC). The second server archives the data for each luminosity section (one luminosity section corresponds to 2^{20} orbits, or about 93 s). An XDAQ layer on this server makes it possible to communicate with other CMS DAQ systems.

7 Forward detectors

7.1 CASTOR

The CASTOR (Centauro And Strange Object Research) detector is a quartz-tungsten sampling calorimeter [119], designed for the very forward rapidity region in heavy ion and proton-proton collisions at the LHC. Its physics motivation is to complement the nucleus-nucleus physics programme [121], developed essentially in the baryon-free mid-rapidity region, and also the diffractive and low- x physics in pp collisions [122]. CASTOR will be installed at 14.38 m from the interaction point, covering the pseudorapidity range $5.2 < |\eta| < 6.6$. Figure 124 shows the location of CASTOR in the CMS forward region. The calorimeter will be constructed in two halves surrounding the beam pipe when closed, as shown in Fig. 125. The calorimeter and its readout are designed in such a way as to permit the observation of the cascade development of the impinging particles as they traverse the calorimeter. The typical total and electromagnetic energies in the CASTOR acceptance range (about 180 TeV and 50 TeV, respectively, according to HIJING [120] Pb-Pb simulations at 5.5 TeV) can be measured with a resolution better than $\approx 1\%$.

The main advantages of quartz calorimeters are radiation hardness, fast response and compact detector dimensions [123], making them suitable for the experimental conditions encountered in the very forward region at the LHC. The typical visible transverse sizes of hadronic and electromagnetic showers in quartz calorimeters are 5–10 cm and about 10 mm respectively (for 95% signal containment), i.e. are a factor 3 to 4 times narrower than those in “standard” (scintillation) calorimeters [123]. A detailed description of the operation principle (including optimal geometrical specifications of the quartz and tungsten plates, and performances of light-guides, reflectors and photodetectors) can be found in references [124, 125].

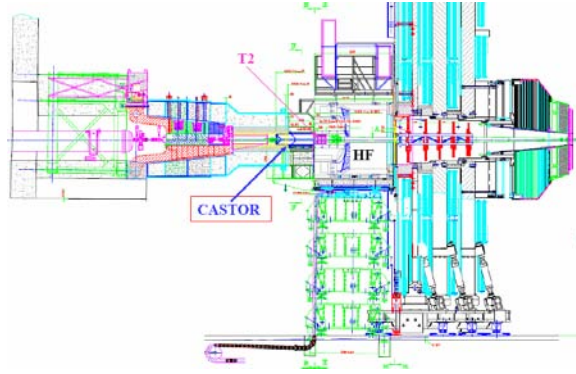


Figure 124: Location of CASTOR in the CMS forward region.

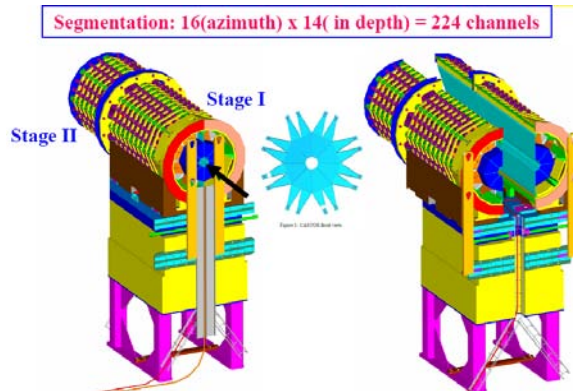


Figure 125: CASTOR calorimeter and support.

Tungsten-Quartz plates

The CASTOR detector is a Cerenkov-based calorimeter, similar in concept to the HF. It is constructed from layers of tungsten (W) plates (density $\approx 19 \text{ g/cm}^3$) as absorber and fused silica quartz (Q) plates as active medium. For the electromagnetic (EM) section, the W plates have a thickness of 5.0 mm and the Q plates 2.0 mm. For the

hadronic (HAD) section, the W and Q plates have thicknesses of 10.0 mm and 4.0 mm, respectively. The W/Q plates are inclined 45° with respect to the direction of the impinging particles, in order to maximize the Cerenkov light output in the quartz. The combination of one W and one Q plate is called a sampling unit (SU). Figure 126 shows the complicated geometry of the W/Q plates, due to their 45° inclination.

In the EM section, each sampling unit corresponds to $2.012 X_0$ ($0.0792 \lambda_I$). Each readout unit (RU) consists of 5 SUs and is $10.06 X_0$ ($0.396 \lambda_I$) deep. The EM section is divided in two successive RUs and has a total of $20.12 X_0$ ($0.792 \lambda_I$). In the hadronic section, a sampling unit corresponds to $0.1584 \lambda_I$. Each readout unit consists of 5 SUs and is $0.792 \lambda_I$ deep. The HAD section has 12 RUs, corresponding to $9.504 \lambda_I$. In total, the calorimeter has $10.3 \lambda_I$. The total number of channels is 224. The calorimeter will be built and installed in two stages, as shown in Fig. 126.

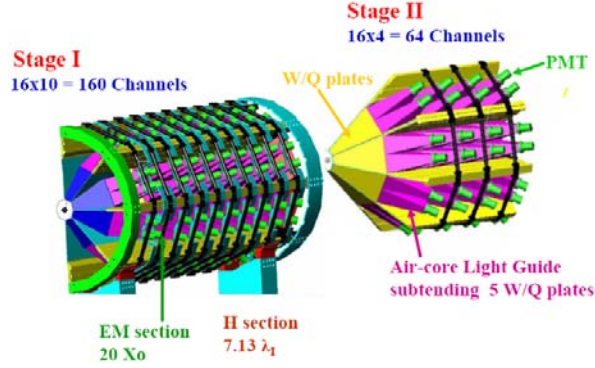


Figure 126: Details of the components and geometry of the CASTOR calorimeter.

Light-guides and photodetectors

The Cerenkov light, produced by the passage of relativistic charged particles through the quartz medium, is collected in sections (RUs) along the length of the calorimeters and focused by air-core light guides onto the photo-multiplier (PMT), as shown in Fig. 126. The inside surfaces of the light guides are covered with Dupont $[\text{AlO}_2 + \text{SiO}_2 + \text{TiO}_2]$ reflective foil. The light guide is made out of a 0.8 mm stainless steel sheet. Each light guide subtends 5 SUs in both the EM and HAD sections. The PMT is located in the aluminium housing on the top. Two types of PMTs are currently under consideration: (i) a Hamamatsu R7899 PMT, and (ii) a radiation-hard multi-mesh, small-size PMT FEU-187 produced by Research Institute Electron (RIE, St. Petersburg), with cathode area $\approx 2 \text{ cm}^2$. Both PMTs allow the muon MIP peak to be separated from the pedestal, an important feature for calibration purposes.

Beam tests results

The energy linearity and resolution as well as the spatial resolution of two CASTOR prototypes have been studied at CERN/SPS tests in 2003 [124] and 2004 [125] (as well as in tests end-of-summer 2007, for the final prototype). The response of the calorimeter to electromagnetic and hadronic showers has been analysed with $E = 20\text{--}200 \text{ GeV}$ electrons, $E = 20\text{--}350 \text{ GeV}$ pions, and $E = 50, 150 \text{ GeV}$ muons. Good energy linearity for electrons and pions in the full range tested is observed. For the EM section, the constant term of the energy resolution, that limits performance at high energies, is less than 1%, whereas the stochastic term is $\approx 50\%$. The measured spatial resolution of the electron (pion) showers is $\sigma_{\text{EM(HAD)}} = 1.7 (6.4) \text{ mm}$.

7.2 Zero degree calorimeter (ZDC)

A set of two zero degree calorimeters [126, 127], with pseudorapidity coverage of $|\eta| \geq 8.3$ for neutral particles, are designed to complement the CMS very forward region, especially for heavy ion and pp diffractive studies. Each ZDC has two independent parts: the electromagnetic (EM) and hadronic (HAD) sections. Two identical ZDCs will be located between the two LHC beam pipes at $\approx 140 \text{ m}$ on each side of the CMS interaction region at the detector slot of 1 m length, 96 mm width and 607 mm height inside the neutral particle absorber TAN [128]. The TAN is located in front of the D2 separation dipole. It was designed to protect magnets and detectors against debris generated in the pp collisions, and against beam halo and beam losses. During heavy ion running the combined

(EM + HAD) calorimeter should allow the reconstruction of the energy of 2.75 TeV spectator neutrons with a resolution of 10–15%. Sampling calorimeters using tungsten and quartz fibers have been chosen for the detection of the energy in the ZDCs with a design similar to HF and CASTOR. The quartz-quartz fibers [126] can withstand up to 30 GRad with only a few percent loss in transparency in the wavelength range 300–425 nm. During the low-luminosity pp ($10^{33} \text{ cm}^{-2}\text{s}^{-1}$) and design-luminosity Pb-Pb ($10^{27} \text{ cm}^{-2}\text{s}^{-1}$) runs, the expected average absorbed radiation doses is about 180 MGy and 300 kGy, respectively, per data-taking year.

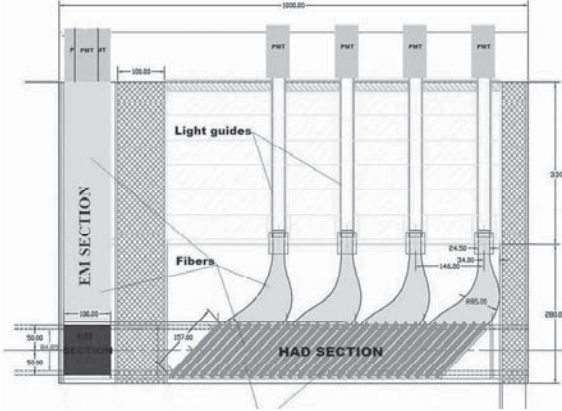


Figure 127: The side view of the ZDC showing the EM and HAD sections.



Figure 128: Photograph of the ZDC HAD section.

Figure 127 shows a side view of the ZDC with the EM section in front and the HAD section behind. A photo of the HAD section is shown in Fig. 128. The total depth of the combined system is ≈ 7.5 hadronic interaction lengths (λ_I). The configuration includes 9 mm Cu plates in the front and back of each section. For the TAN's final detector configuration an LHC real-time luminosity monitor (BRAN, Beam RATE of Neutrals [129]) will be mounted in the 120 mm space between the ZDC's calorimetric sections. The HAD section consists of 24 layers of 15.5 mm thick tungsten plates and 24 layers of 0.7 mm diameter quartz fibers ($6.5 \lambda_I$). The tungsten plates are tilted by 45° to optimize Cerenkov-light output. The EM section is made of 33 layers of 2-mm-thick tungsten plates and 33 layers of 0.7-mm-diameter quartz fibers ($19 X_0$). The tungsten plates are oriented vertically. The fibers are laid in ribbons. The hadronic section of each ZDC requires 24 fiber ribbons. After exiting the tungsten plates the fibers from 6 individual ribbons are grouped together to form a readout bundle. This bundle is compressed and glued with epoxy into a tube. From there, an optical air-core light guide will carry the light through radiation shielding to the photomultiplier tube. The full hadronic section will consist of four identical towers divided in the longitudinal direction. For the electromagnetic section, fibers from all 33 fiber ribbons will be divided in the horizontal direction into five identical fiber bundles. These 5 bundles will form five horizontal towers and each fiber bundle will be mounted with a 0.5 mm air gap from the photocathode of a phototube. The EM and HAD sections will be instrumented with the same type of phototube as the HF: Hamamatsu R7525 phototubes with a bi-alkali photocathode, resulting in an average quantum efficiency for Cerenkov light of about 10%.

There are a total of 18 readout channels for the two ZDCs. The signals from the ZDCs are transmitted through a long (210 m) coaxial cable to the front-end HCAL VME crates in the underground counting room (USC55). The signal from each channel will be split, with 90% going to the QIE (Charge Integrator and Encoder) while 10% will be used for making trigger signals. An analog sum, proportional to the total energy deposition in each detector, will provide the basic Level 1 trigger in the heavy-ion running mode: the coincidence of (neutron) signals from both sides of the interaction point is sensitive to most of the nuclear and electromagnetic cross section. A left-right timing coincidence will also be used as a fast vertex trigger, to suppress beam-gas events in the heavy ion runs. Information from scalers will be used for tuning the interaction of beams and for defining the real-time luminosity. Finally it may be possible to improve the overall energy resolution of the system by looking at the correlation between the ZDC and the BRAN detector, which sits between the electromagnetic and hadronic sections, near the shower maximum.

The response of the ZDC EM and HAD sections has been studied in beam tests at the CERN/SPS in 2006 [130] and 2007. The calorimeter is found to be linear within 2% in the range from 20 GeV to 100 GeV. The energy

resolution obtained for the different positron energies can be parametrized as

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{70\%}{\sqrt{E}}\right)^2 + (8\%)^2 \quad (6)$$

where E is in GeV. Positive pions with energies of 150 GeV and 300 GeV were used to measure the response of the combined EM+HAD system. The pion energy resolution, obtained by a Landau fit, can be parametrized as

$$\frac{\sigma}{E} = \frac{138\%}{\sqrt{E}} + 13\% \quad (7)$$

where E, again, is in GeV. The width of EM showers is ≈ 5 mm. Such a good position resolution will allow measurement of the beam crossing angle with a resolution of ≈ 10 mrad.

The performance of both the left and right ZDCs has been studied with electron, pion and muon beams in 2007. Figure 129 shows online plots for positrons entering the electromagnetic section of one calorimeter.

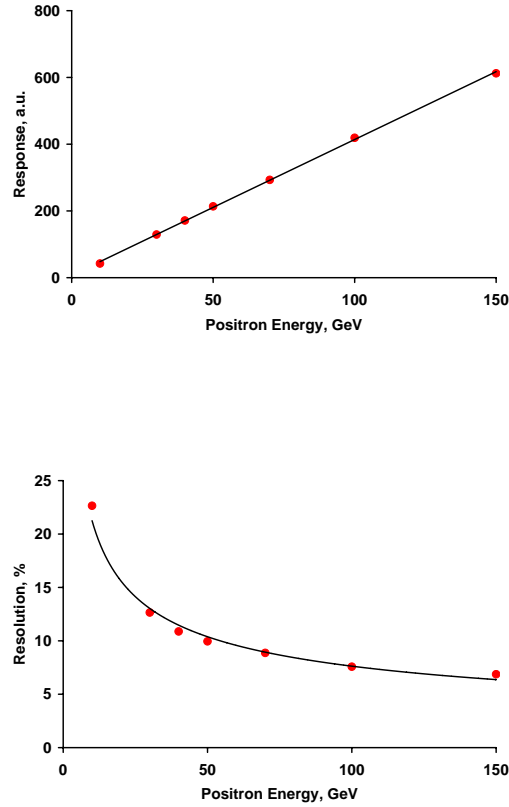


Figure 129: Online results for positrons from the 2007 test beam. The top panel shows the response linearity, while the bottom panel gives the energy resolution as a function of the incoming positron beam energy.

8 The muon system

Muon detection is a powerful tool for recognizing signatures of interesting processes over the very high background rate expected at the LHC with full luminosity. For example, the predicted decay of the Standard Model Higgs boson into ZZ or ZZ^* , which in turn decay into 4 leptons, has been called “gold plated” for the case in which all the leptons are muons. Besides the relative ease in detecting muons, the best 4-particle mass resolution can be achieved if all the leptons are muons because they are less affected than electrons by radiative losses in the tracker material. This example, and others from SUSY models, emphasize the discovery potential of muon final states and the necessity for wide angular coverage for muon detection.

Therefore, as is implied by the experiment’s middle name, the detection of muons is of central importance to CMS: precise and robust muon measurement was a central theme from its earliest design stages. The muon system has 3 functions: muon identification, momentum measurement, and triggering. Good muon momentum resolution and trigger capability are enabled by the high-field solenoidal magnet and its flux-return yoke. The latter also serves as a hadron absorber for the identification of muons.

The CMS muon system is designed to have the capability of reconstructing the momentum and charge of muons over the entire kinematic range of the LHC. CMS uses 3 types of gaseous particle detectors for muon identification [131]. Due to the shape of the solenoid magnet, the muon system was naturally driven to have a cylindrical, barrel section and 2 planar endcap regions. Because the muon system consists of about 25 000 m² of detection planes, the muon chambers had to be inexpensive, reliable, and robust.

In the barrel region, where the neutron-induced background is small, the muon rate is low, and the 4-T magnetic field is uniform and mostly contained in the steel yoke, drift chambers with standard rectangular drift cells are used. The barrel drift tube (DT) chambers cover the pseudorapidity region $|\eta| < 1.2$ and are organized into 4 stations interspersed among the layers of the flux return plates. The first 3 stations each contain 8 chambers, in 2 groups of 4, which measure the muon coordinate in the r - ϕ bending plane, and 4 chambers which provide a measurement in the z direction, along the beam line. The fourth station does not contain the z -measuring planes. The 2 sets of 4 chambers in each station are separated as much as possible to achieve the best angular resolution. The drift cells of each chamber are offset by a half-cell width with respect to their neighbor to eliminate dead spots in the efficiency. This arrangement also provides a convenient way to measure the muon time with excellent time resolution, using simple meantimer circuits, for efficient, standalone bunch-crossing identification. The number of chambers in each station and their orientation were chosen to provide good efficiency for linking together muon hits from different stations into a single muon track and for rejecting background hits.

In the 2 endcap regions of CMS, where the muon rates and background levels are high and the magnetic field is large and non-uniform, the muon system uses cathode strip chambers (CSC). With their fast response time, fine segmentation, and radiation resistance, the CSCs identify muons between $|\eta|$ values of 0.9 and 2.4. There are 4 stations of CSCs in each endcap, with chambers positioned perpendicular to the beam line and interspersed between the flux return plates. The cathode strips of each chamber run radially outward and provide a precision measurement in the r - ϕ bending plane. The anode wires run approximately perpendicular to the strips and are also read out in order to provide measurements of η and the beam-crossing time of a muon. Each 6-layer CSC provides robust pattern recognition for rejection of non-muon backgrounds and efficient matching of hits to those in other stations and to the CMS inner tracker.

Because the muon detector elements cover the full pseudorapidity interval $|\eta| < 2.4$ with no acceptance gaps, muon identification is ensured over the range corresponding to $10^\circ < \theta < 170^\circ$. Offline reconstruction efficiency of simulated single-muon samples (Fig. 130) is typically 95–99% except in the regions around $|\eta| = 0.25$ and 0.8 (the regions between 2 DT wheels) and $|\eta| = 1.2$ (the transition region between the DT and CSC systems), where the efficiency drops. Negligible punchthrough reaches the system due to the amount of material in front of the muon system, which exceeds 16 interaction lengths [131].

Due to multiple-scattering in the detector material before the first muon station, the offline muon momentum resolution of the standalone muon system is about 9% for small values of η and p for transverse momenta up to 200 GeV [17]. At 1 TeV the standalone momentum resolution varies between 15% and 40%, depending on $|\eta|$. A global momentum fit using also the inner tracker improves the momentum resolution by an order of magnitude at low momenta. At high momenta (1 TeV) both detector parts together yield a momentum resolution of about 5% (Fig. 2). Note that the muon system and the inner tracker provide independent muon momentum measurements; this redundancy enhances fault finding and permits cross-checking between the systems.

A crucial characteristic of the DT and CSC subsystems is that they can each trigger on the p_T of muons with good efficiency and high background rejection, independent of the rest of the detector. The Level-1 trigger p_T resolution

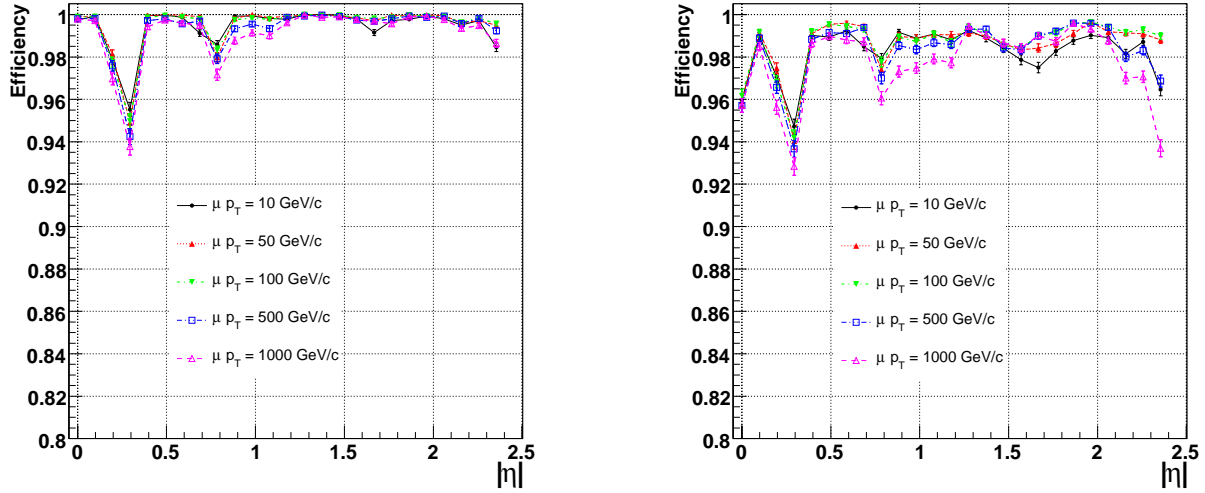


Figure 130: Muon reconstruction efficiency as a function of pseudorapidity for selected values of p_T . Left panel: standalone reconstruction (using only hits from the muon system with a vertex constraint). Right panel: global reconstruction (using hits from both the muon system and the tracker).

is about 15% in the barrel and 25% in the endcap.

Because of the uncertainty in the eventual background rates and in the ability of the muon system to measure the correct beam-crossing time when the LHC reaches full luminosity, a complementary, dedicated trigger system consisting of resistive plate chambers (RPC) was added in both the barrel and endcap regions. The RPCs provide a fast, independent, and highly-segmented trigger with a sharp p_T threshold over a large portion of the rapidity range ($|\eta| < 1.6$) of the muon system. The RPCs are double-gap chambers, operated in avalanche mode to ensure good operation at high rates. They produce a fast response, with good time resolution but coarser position resolution than the DTs or CSCs. They also help to resolve ambiguities in attempting to make tracks from multiple hits in a chamber.

A total of 6 layers of RPCs are embedded in the barrel muon system, 2 in each of the first 2 stations, and 1 in each of the last 2 stations. The redundancy in the first 2 stations allows the trigger algorithm to work even for low- p_T tracks that may stop before reaching the outer 2 stations. In the endcap region, there is a plane of RPCs in each of the first 3 stations in order for the trigger to use the coincidences between stations to reduce background, to improve the time resolution for bunch-crossing identification, and to achieve a good p_T resolution.

Finally, a sophisticated alignment system measures the positions of the muon detectors with respect to each other and to the inner tracker, in order to optimize the muon momentum resolution.

8.1 Drift tube system

8.1.1 General description

The CMS barrel muon detector consists of 4 stations forming concentric cylinders around the beam line: the 3 inner cylinders have 60 drift chambers each and the outer cylinder has 70. There are about 172 000 sensitive wires. It is possible to use drift chambers as the tracking detectors for the barrel muon system because of the low expected rate and the relatively low strength of the local magnetic field.

The wire length, around 2.4 m in the chambers measured in an r - ϕ projection, is constrained by the longitudinal segmentation of the iron barrel yoke. The transverse dimension of the drift cell, i.e., the maximum path and time of drift, was chosen to be 21 mm (corresponding to a drift time of 380 ns in a gas mixture of 15% Ar + 85% CO₂). This value is small enough to produce a negligible occupancy and to avoid the need for multi-hit electronics, yet the cell is large enough to limit the number of active channels to an affordable value. A tube was chosen as the basic drift unit to obtain protection against damage from a broken wire and to partially decouple contiguous cells from the electromagnetic debris accompanying the muon itself.

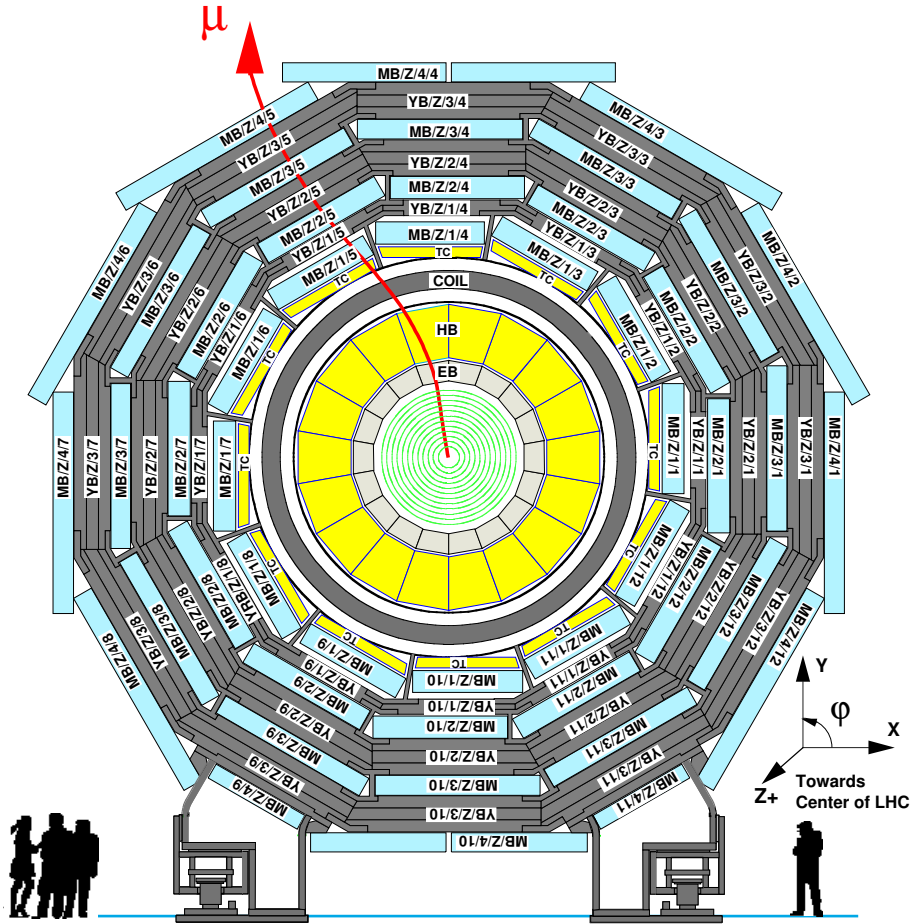


Figure 131: Layout of the CMS barrel muon DT chambers in one of the 5 wheels. The chambers in each wheel are identical with the exception of wheels -1 and $+1$ where the presence of cryogenic chimneys for the magnet shortens the chambers in 2 sectors. Note that in sectors 4 (top) and 10 (bottom) the MB4 chambers are cut in half to simplify the mechanical assembly and the global chamber layout.

The amount of iron in the return yoke was dictated by the decision to have a large and intense solenoidal magnetic field at the core of CMS. Two detector layers, one inside the yoke and the other outside, would be insufficient for reliable identification and measurement of a muon in CMS. Therefore, 2 additional layers are embedded within the yoke iron (Fig. 131). In each of the 12 sectors of the yoke there are 4 muon chambers per wheel, labeled MB1, MB2, MB3, and MB4. The yoke-iron supports that are between the chambers of a station generate 12 unavoidable

dead zones in the ϕ coverage, although the supports are placed so as not to overlap in ϕ .

A drift-tube (DT) chamber is made of 3 (or 2) superlayers (SL, see Fig. 132), each made of 4 layers of rectangular drift cells staggered by half a cell. The SL is the smallest independent unit of the design.

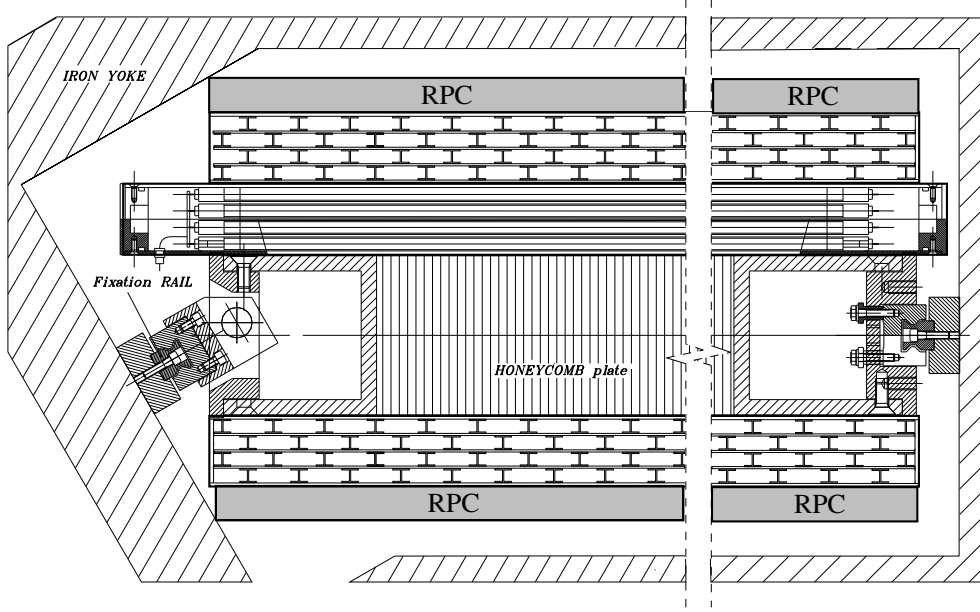


Figure 132: A DT chamber in position inside the iron yoke; the view is in the $(r-\phi)$ plane. One can see the 2 SLs with wires along the beam direction and the other perpendicular to it. In between is a honeycomb plate with supports attached to the iron yoke. Not shown are the RPCs, which are attached to the DT chambers via support plates glued to the bottom and/or top faces, depending on chamber type.

The wires in the 2 outer SLs are parallel to the beam line and provide a track measurement in the magnetic bending plane $(r-\phi)$. In the inner SL, the wires are orthogonal to the beam line and measure the z position along the beam. This third, z -measuring, SL is not present in the fourth layer of chambers, which therefore measures only the ϕ coordinate. A muon coming from the interaction point first encounters a ϕ -measuring SL, passes through the honeycomb plate, then crosses the z -measuring SL and the second ϕ -measuring SL. In this scenario, there still exist limited regions of η in which the combined effect of the ϕ and z discontinuities limits to only 2 (out of 4), the number of stations crossed by a muon.

At high momenta (≥ 40 GeV), the probability of electromagnetic cascades accompanying the parent muon becomes relevant. A reliable way to cope with this effect in the regions where only 2 stations are available is to have a good tracking efficiency in each station even in the presence of electromagnetic debris. Redundancy is also needed to cope with the uncorrelated background hits generated by neutrons and photons whose rate is much larger than that from prompt muons. Redundancy is obtained by having several layers of separated drift cells per station. The separation, i.e., the thickness of the tube walls, should be large enough to decouple the basic units against low-energy electrons. The relatively thick wall of the DTs, 1.5 mm, gives an effective decoupling among the several layers of tubes inside the same station. With this design, the efficiency to reconstruct a high p_T muon track with a momentum measurement delivered by the barrel muon system alone is better than 95% in the pseudorapidity range covered by 4 stations, i.e., $\eta < 0.8$. The constraints of mechanical stability, limited space, and the requirement of redundancy led to the choice of a tube cross section of 13×42 mm².

The many layers of heavy tubes require a robust and light mechanical structure to avoid significant deformations due to gravity in the chambers, especially in those that lie nearly horizontal. The chosen structure is basically frameless and for lightness and rigidity uses an aluminium honeycomb plate that separates the outer superlayer(s) from the inner one (Fig. 132). The SLs are glued to the outer faces of the honeycomb. In this design, the honeycomb serves as a very light spacer, with rigidity provided by the outer planes of tubes. A thick spacer also helps having good angular resolution within a station. Table 16 provides a summary of the general DT chamber parameters.

One SL, that is, a group of 4 consecutive layers of thin tubes staggered by half a tube, gives excellent time-tagging capability, with a time resolution of a few nanoseconds. This capability provides local, stand-alone, and efficient bunch-crossing identification. The time tagging is delayed by a constant amount of time equal to the maximum

chamber type	No. of chambers	No. of SL Φ	No. of SL Θ	No. of ch. SL Φ	No. of ch. SL Θ	Wire length Φ (mm)	Wire length Θ (mm)	Sum of ch.
MB/W/1/S	58	2	1	196	228	2379	2038	35960
MB/1/1/4	1	2	1	196	190	1989	2038	582
MB/-1/1/3	1	2	1	196	190	1989	2038	582
MB/W/2/S	58	2	1	238	228	2379	2501	40832
MB/1/2/4	1	2	1	238	190	1989	2501	666
MB/-1/2/3	1	2	1	238	190	1989	2501	666
MB/W/3/S	58	2	1	286	228	2379	3021	46400
MB/1/3/4	1	2	1	286	190	1989	3021	762
MB/-1/3/3	1	2	1	286	190	1989	3021	762
MB/W/4/S	29	2	0	382	0	2379	0	22156
MB/-1/4/3	1	2	0	382	0	1989	0	764
MB/W/4/4	8	2	0	286	0	2379	0	4576
MB/1/4/4	2	2	0	286	0	1989	0	1144
MB/W/4/8,12	10	2	0	372	0	2379	0	7440
MB/W/4/9,11	10	2	0	190	0	2379	0	3800
MB/W/4/10	10	2	0	238	0	2379	0	4760
total	250							171852

Table 16: Chambers of the CMS DT system. Notation: MB/wheel/station/sector. W stands for all 5 wheels (numbered $-2, -1, 0, 1$, and 2) and S means any sector (1 to 12, see Fig. 131). The SLs of type $\Phi(\Theta)$ measure the $\phi(z)$ coordinate in the CMS coordinate system.

possible drift-time, which is determined by the size of the tube, the electrical field, and the gas mixture. Within the angular range of interest, the time resolution was shown to be largely independent of the track angle, but this requires the cell optics to maintain a linear relationship between the distance from the wire of the crossing track and the drift-time of the electrons along the entire drift path. Bunch-crossing tagging is performed independently in each of the 3 SLs by fast pattern-recognition circuitry. Together with the bunch-crossing assignment, this circuit delivers the position of the centre of gravity of the track segment and its angle in the SL reference system with precisions of 1.5 mm and 20 mrad, respectively. This information is used by the first-level muon trigger for the time and transverse momentum assignment.

The goal of the mechanical precision of the construction of a chamber was to achieve a global resolution in r - ϕ of 100 μm . This figure makes the precision of the MB1 chamber (the innermost layer) comparable to the multiple scattering contribution up to $p_T = 200$ GeV. The 100- μm target chamber resolution is achieved by the 8 track points measured in the two ϕ SLs, since the single wire resolution is better than 250 μm . To avoid corrections to the primary TDC data (Sect. 8.1.3), the deviation from linearity of the space-time relation in each drift cell must be less than 100–150 μm . This figure matches well with the requirements of linearity for the bunch-crossing (Sect. 8.1.3) identifier. The cell design includes 5 electrodes, 1 anode wire, 2 field shaping strips, and 2 cathode strips (Fig. 133 and Sect. 8.1.2). The requirements of 250- μm resolution and 150- μm nonlinearity can be obtained by operating the tubes at atmospheric pressure with an Ar/CO₂ gas mixture and by keeping the CO₂ concentration in the 10–20% range. The multi-electrode design also ensures this performance in the presence of the stray magnetic field present in some regions of the chambers. It is worth noting that to reach this local performance in a single tube, the precision requirement on the position of the field-shaping electrodes, including the wires, is about 300 μm , which is considerably less demanding than the 100 μm required for the mechanical construction.

8.1.2 Technical design

Drift cell Figure 133 shows the drift lines in a cell. The anode is a 50- μm -diameter gold-plated stainless-steel wire. The field electrode is made of a 16-mm-wide, 50- μm thick aluminium tape, glued on a 100- μm thick, 23-mm-wide mylar tape that insulates the electrode with respect to the aluminium plate set to ground. Both the conductive and insulating ribbons are self-adhesive with a pressure-activated glue. Field electrodes are positioned at the top and bottom of the drift cell. Cathodes are placed on both sides of the I-beams (Fig. 134) following a technique similar to that used for the strip electrodes on the aluminium plates. A cathode consists of a 50- μm -thick, 11.5-mm-wide aluminium tape insulated from the I-beam by 19-mm-wide, 100- μm -thick mylar tape. This design allows for at least 3.5 mm separation of the electrode from the sides of the grounded I-beam. At the extremities

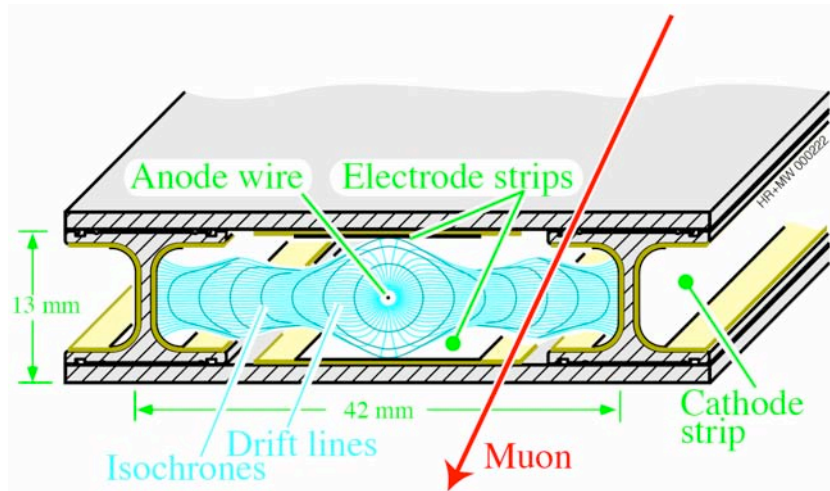


Figure 133: Sketch of a cell showing drift lines and isochrones. The plates at the top and bottom of the cell are at ground potential. The voltages applied to the electrodes are $+3600\text{V}$ for wires, $+1800\text{V}$ for strips, and -1200V for cathodes.

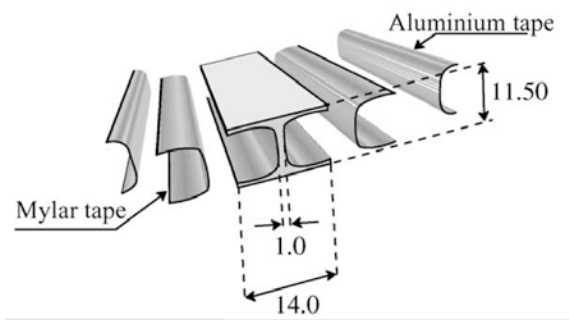


Figure 134: Exploded view of the cathode electrodes, glued on the I-beams.

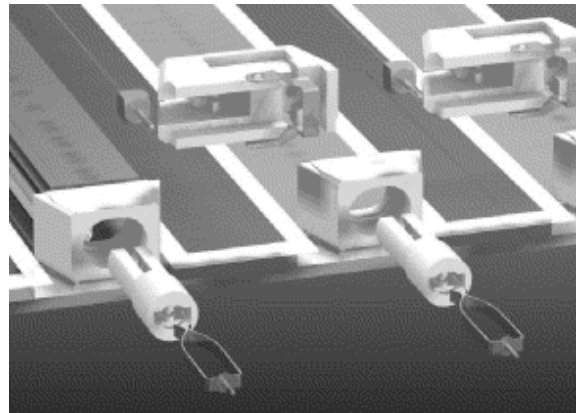


Figure 135: Exploded view of the end part of the drift cells showing the different end-plugs and spring contacts for high voltage connections.

the mylar tape is cut flush with respect to the I-beam ends while the aluminium tape is recessed by 5 mm. Special tools were designed and built to glue the electrode strips to both the plates and the I-beams. The only difference between the tapes used for the electrode strips and the ones just described is the width: the mylar tape used for the electrode strips is 23-mm wide and the aluminium tape is 16-mm wide. These strips are set to a positive voltage and help to improve the shaping of the electric field and the linearity of the space-time relation, most noticeably in the presence of magnetic fields.

The cathode and wire end-plugs were designed to protect against discharges from the border at the end of the cathode strips and to house the wire holder, which is crucial for the wire position precision. The wire holders protrude inside the cell providing 12 mm of additional protection around the wire. The I-beam and wire end-plug pieces, as well as the springs connecting the electrodes to the high voltage, are shown in Fig. 135. The Ar-CO₂ mixture and the drift-cell optics described above provide a linear relationship between time and drift path. This is an essential requirement for the use of the chamber as a first-level trigger device [132]. A calculation of the drift velocity using GARFIELD [133] showed that drift velocity saturation occurs between 1 and 2 kV/cm (Fig. 136). This may be compared to the drift velocity as measured with the Drift Velocity Chamber (VdC) (Fig. 137).

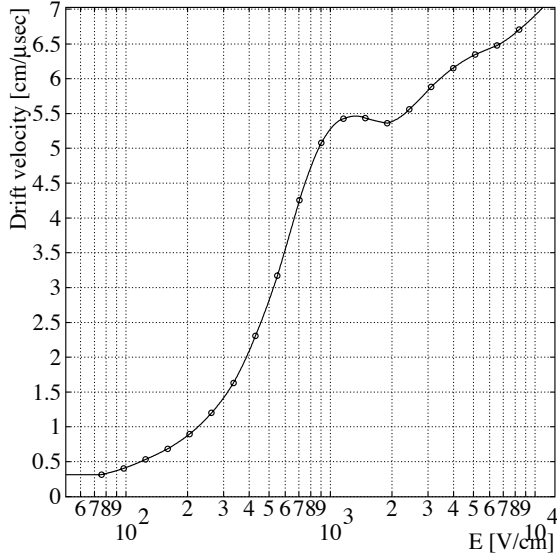


Figure 136: Calculated drift velocity (in cm/μs) as a function of the electric field (in V/cm) for a gas mixture Ar/CO₂ (85%/15%). The values obtained are very similar for the 2 different environment pressures, 973 and 1027 hPa.

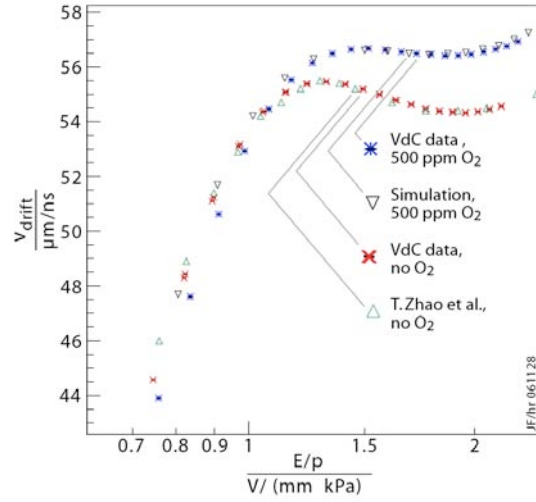


Figure 137: Measured and calculated drift velocities (in μm/ns) as functions of electric field and gas pressure for a pure gas mixture Ar/CO₂ (85%/15%) and for a gas mixture with air impurities corresponding to 500-ppm O₂. The measurements were obtained with the VdC, a dedicated reference drift chamber that will be used for drift-velocity monitoring during CMS running. For comparison, results of measurements from [134] for a pure gas mixture, and a simulation with Magboltz [133] for a mixture with impurities are also shown.

The drift cells will operate at a gas gain of 10⁵, allowing them to work within an efficiency plateau with a wide threshold range, which is convenient for the operation of large chambers in the environment expected at CMS. A computation of equipotential lines [135] (Fig. 138) is useful for better understanding of the role of each electrode. The position of the 0 V equipotential in the region between the central strips and the cathodes is mainly determined by the size of the electrodes and not by their voltage values. The gas gain is mainly determined by the voltage drop from the wire to the nearest electrode, the strips. The wire/strip voltage difference must be kept between 1.75 and 1.85 kV to achieve a gain not far from the expected value of 10⁵. Under the rough assumption of a uniform drift field of 1.5 kV/cm, the distances between the various electrodes imply that the strips should be set to a voltage larger than or equal to 1.7 kV and the cathodes to around -1 kV. As described below, during the chamber

commissioning in laboratories and at CERN (with $B=0$), satisfactory performance was obtained with the voltages of cathodes, strips, and wires set to -1.2 , $+1.8$, and $+3.6$ kV, respectively.

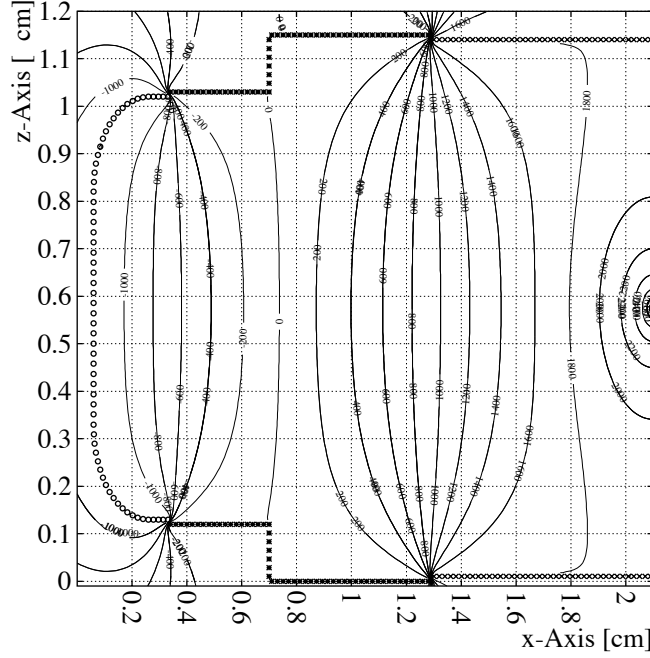


Figure 138: Equipotential lines in half of a drift cell. The anode wire is on the right side. The lines are labeled with the potentials in volts (the x -axis is perpendicular to the wires on the wire plane, while the z -axis is orthogonal to the wire plane).

Chamber mechanics and services A chamber is assembled by gluing 3 (or 2) SLs to an aluminium honeycomb plate to ensure the required stiffness. Each SL is made of 5 aluminium sheets, 1.5-mm thick, separated by 11.5-mm-high, 1-mm-thick aluminium I-beams, as described in Sect. 8.1.2. The cell pitch is 42 mm, while the layer pitch is 13 mm. For the construction of the SLs, a full layer of cells is built at the same time by gluing together 2 aluminium plates separated by an array of parallel aluminium I-beams. The pitch and height of an I-beam determine the larger and smaller dimensions of a cell, respectively. An SL has an independent gas and electronics enclosure. Each SL is assembled and tested individually before being glued to the honeycomb plate and/or to the other SL to form a chamber.

HV connections to the cells and the front-end electronics are located at opposite ends of the wires. The HV is fed into each SL via two 52-pin custom connectors and distributed to the drift cells via printed-circuit HV boards (HVB). Each HVB is mounted along the edge of the aluminium plate separating 2 layers of drift cells and serves the 8 cells above and the 8 cells below it. One HV channel is dedicated to each group of 8 anode wires, while for the other 2 voltages there is 1 HV channel for 16 cells in the same layer. The strip and the cathode voltages can be daisy chained from an HVB to the next one. On the HVB there is 1 capacitor for filtering for each group. As a current limiter, a 50-M Ω ceramic resistor is used in series. There is 1 resistor for each anode and 1 for each group of 4 strips or cathodes. Due to the restricted space as required to minimise dead space, the size of the HVB is only about 307×37 mm² and special care had to be taken to maximise the distance between the printed HV lines and to avoid any embedded gas pocket within the HVB volume.

The gas enclosures are different: on the HV side they contain the HV distribution system and gas connector, whereas on the front-end sides there are the HV decoupling capacitors, the front-end circuitry, the pulse distribution system, the gas distribution, and the necessary cooling for the electronics. Inside the SL, gas is distributed as a “shower”, with each cell being fed through a small-diameter (1 mm²) outlet hole to guarantee that the same amount of gas is circulated to each channel. The distance between the position where the wire enters the end plug and the outer face of the gas enclosure, which determines the SL dead area, corresponds to 60 mm on both the HV and the front-end side (corresponding to $\approx 10\%$ dead space).

It is very important that the individual SLs of the DT chambers are gas tight because contamination by nitrogen (from air) changes the drift velocity by a sizeable amount, while oxygen reduces the signal efficiency, when its contamination exceeds 2000 ppm. Contamination by air including 1000 ppm of O_2 changes the maximum drift time by about 2% with respect to no contamination, with a sizeable effect on the trigger performance of the detector. In the DT chambers, the gas tightness of the SLs is obtained by gluing profiles to the outer aluminium skins. Along 2 sides of the SL, C-shaped profiles are used and the ends of these profiles are glued to reference blocks (Fig. 139), forming the corners of the SL box. The front and back of the box have L-shaped profiles glued along the plate border to form an open frame, which is then closed with removable long cover plates that contain all necessary gas connectors, HV connectors, and signal outputs, equipped with O-rings that seal the structure. A 3-dimensional computer model of the gas enclosure for one SL, where the outer aluminium plates have been removed to expose all details of the gas enclosure, is shown in Fig. 140. With this type of gas enclosure we can obtain a level of oxygen contamination of 10–20 ppm, downstream of the 3 SLs flushed in parallel with about 1 volume change per day.

During SL assembly, before the fifth aluminium plate is glued closing the structure, reference blocks are glued such that their positions with respect to the wires can be measured precisely. Thus, when the chamber is completed, the wire positions may be determined by measuring the reference marks on the blocks. Pressure and temperature monitoring probes, ground straps that connect all the aluminium planes to form a unique ground reference, and a Faraday cage for the signals, front-end electronics, and HV distribution complete the equipment that is in the gas enclosures of each SL. Each SL is fully independent with respect to gas tightness, HV, and front-end electronics;



Figure 139: Corner blocks of an SL. These pieces also carry the reference marks with respect to which the wire positions are measured.

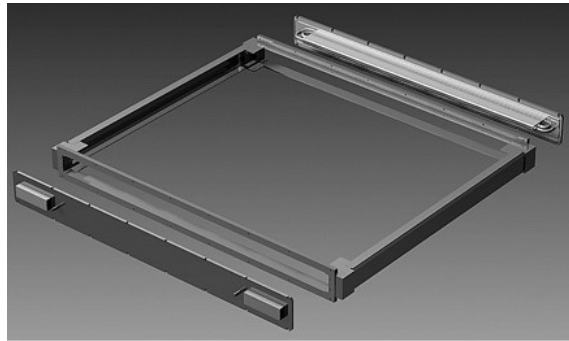


Figure 140: A 3-dimensional computer model of the gas enclosure of the SLs.

hence an SL can be fully tested before it is glued to form a DT chamber. SLs are glued to a honeycomb panel (Fig. 132) that sustains and gives rigidity to the chamber and provides the fixation points from which it is suspended in the CMS barrel steel yoke (two Φ -type SLs and one Θ -type SL in the case of layers 1, 2, and 3; and two Φ -type SLs for layer 4). The panel thickness varies from 125 mm for the first three stations, to 178 mm for the fourth station. It is delivered with the correct dimensions and equipped with the C-shaped profiles at the periphery that are used for the supports and for part of the electronics.

The space for the chamber supports and attachments, the passages for alignment, and the local read-out and trigger electronics is provided by a channel running around the border of the honeycomb plate. The channel is approximately as wide and deep as the honeycomb plate thickness. The 2 channels parallel to the beam line and to the yoke steel supports house the kinematic fixations to the yoke supports themselves, and the longitudinal alignment passages. One of the 2 remaining sides houses the read-out and trigger electronics that collect the full chamber information (*minicrates*). To ease chamber handling, all services are connected on the same side of the chamber. All the general services for the chambers are located around each barrel wheel on the 4 balconies along the walls of the CMS cavern where there is space for the racks and crates. Each wheel is thus an independent, large subsystem.

8.1.3 Electronics

The DT electronics is a complex, heavily integrated system, which includes L1 trigger logic, read-out data handling, and service electronics, such as the LV and HV systems. A description of the electronic system layout together with the functions associated to each sub-task is shown in Fig. 141 and briefly summarized, whereas detailed informations will be described in the following Sections. Front-end electronics and HV distribution are

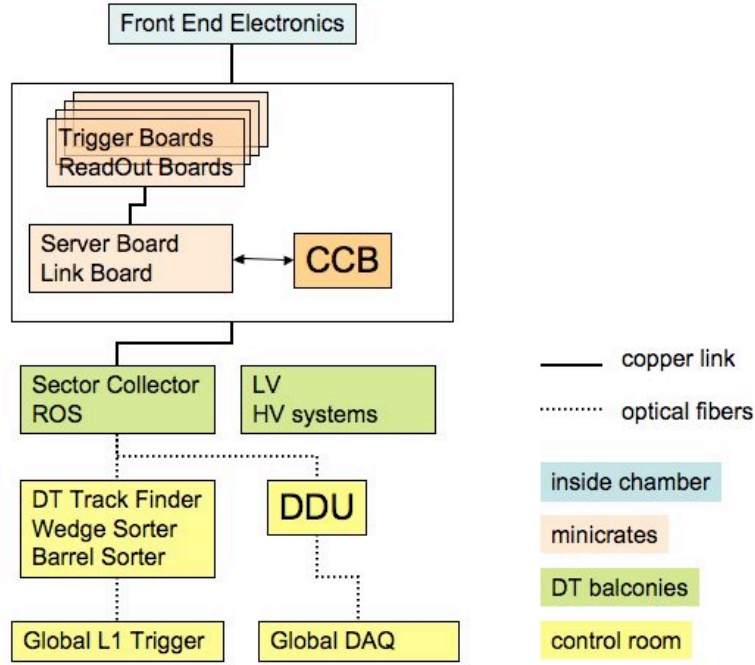


Figure 141: Block diagram of the DT electronic system.

physically embedded in the chamber gas volume. Amplified and shaped signals are directly fed to the *minicrates*. A minicrate, as described previously, is an aluminium structure attached to the honeycomb of the drift tube chambers that houses both the first level of the read-out and of the trigger electronics. The trigger boards located in the minicrates are the Trigger Boards (TRB) and the Server Boards (SB), as described in detail in Sect. 9.2. In each TRB are located the Bunch Crossing and Track Identifier (BTI), which provides independent segments from each chamber SL, and the Track Correlator (TRACO), which correlates ϕ segments in the same chamber by requiring a spatial matching between segments occurring at the same bunch crossing (BX). TRB output signals are fed to the Server Board (SB) which selects the best two tracks from all TRACO candidates and sends the data out of the minicrate. In parallel to the trigger signals, chamber data are fed to the read-out system through the Read Out Boards (ROB), which are in charge of the time digitization of chamber signals related to the Level-1 Accept (L1A) trigger decision and the data merging to the next stages of the data acquisition chain. The Chamber Control Board (CCB) located at the centre of the minicrate, allows ROB and TRB configuration and monitoring. It works together with the CCB link board, on one of the minicrate ends, that receives data from the Slow Control and global experiment Timing and Trigger Control (TTC) system. Among other tasks, the CCB distributes the LHC 40.08 MHz clock and other TTC signals to every board in the minicrate.

Trigger and data signals coming out of the minicrates are collected by VME electronics installed in the iron balconies attached to the DT wheels, respectively to the Sector Collector (SC, Sect. 9.2) and to the Read Out Server (ROS, Sect. 8.1.3) where data merging is performed. From the wheel balconies, data are sent via optical links, both to the CMS central acquisition system through the Detector Dependent Unit (DDU) (Sect. 8.1.3) and to the CMS L1 system through the Sector Collector (SC) and the Drift Tube Track Finder (DTTF, Sect. 9.2).

Front-end electronics The front-end electronics for the barrel muon detector must satisfy many stringent requirements. Its functions are to amplify the signals produced by the detector, compare them with a threshold, and send the results to the trigger and read-out chains located on the chamber. Analog signal processing must use a short shaping time to achieve a high spatial resolution while introducing minimal noise; this allows low-gain operation of the drift tubes, thus improving reliability and chamber lifetime. The downstream comparator has to be very fast and precise to allow the use of low threshold values, which reduce the influence of the signal amplitude on the time response. The output driver also must be very fast, and it must deliver differential levels that minimise mutual interferences and can be transmitted through low-cost cables. Besides the above functions, several features that simplify the control and monitoring of the data acquisition have been implemented. The large number of channels and the resulting need for both high reliability and low cost, limited space, and concerns about power consumption led to the necessity to integrate the front-end electronics as much as possible.

The resulting custom front-end application specific integrated circuit chip (ASIC), named MAD, was developed using $0.8\ \mu\text{m}$ BiCMOS technology [139]. This chip integrates signal processing for 4 channels (4 drift tubes) plus some ancillary functions in a $2.5 \times 2.5\ \text{mm}^2$ die and 80 000 pieces were produced with a fabrication yield better than 95%. Figure 142 shows the ASIC block diagram and the pinout of the TQFP44 package used for it. Each

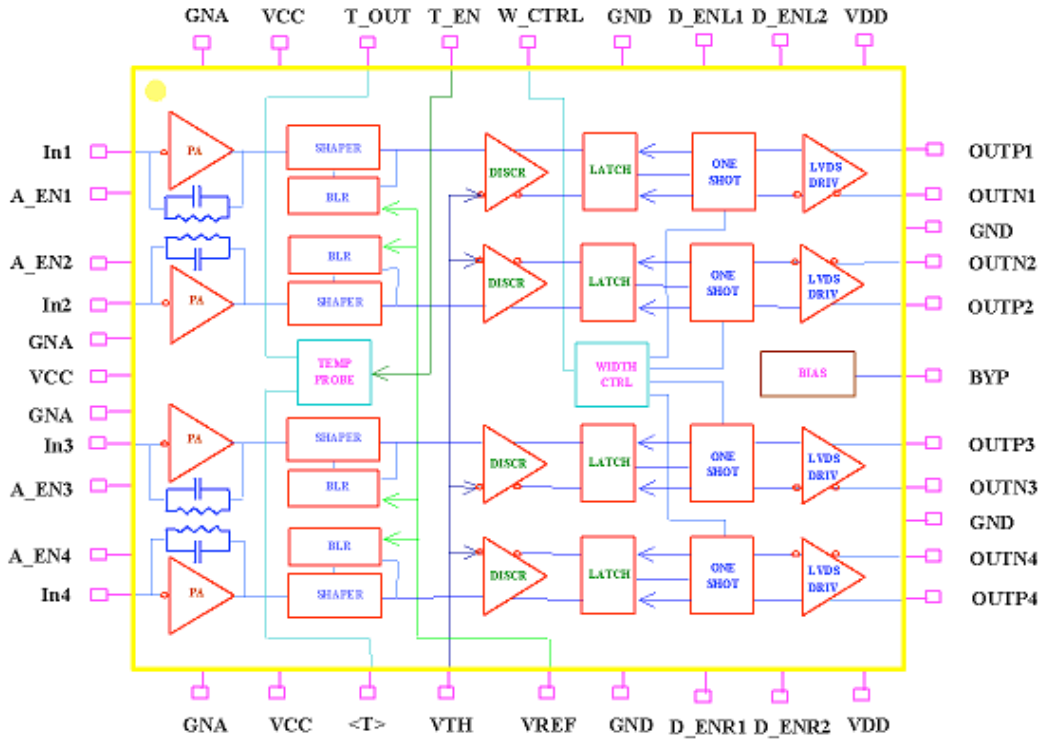


Figure 142: Block diagram of the MAD ASIC.

of the 4 identical analog chains begins with a charge preamplifier that uses a single gain stage, folded/unfolded cascode, having a GBW product in excess of 1 GHz (result from simulation). The feedback time constant is 33 ns while input impedance is $\approx 100\ \Omega$ in the range 5–200 MHz.

The shaper that follows is a low-gain integrator with a small time constant. Its output is directly connected to 1 input of a latched discriminator made of 2 differential gain stages, the other input being connected to the external threshold pin V_{th} , common to all channels. Auxiliary circuits allow the masking of individual channels at the shaper input (pins A_ENn in high state) thus stopping the propagation of excessive noise background to the trigger and DAQ electronics. A similar but faster enable/disable function was implemented on the cable-driver stage to select channels that output signals in response to a test input. A temperature probe was also included for monitoring the operating conditions of the detector.

The ASIC operates with 2 distinct supply voltages, 5 V for the analog section and 2.5 V for the output stage, with a total power consumption of 100 mW (25 mW/ch) equally split between the 2 voltages and almost independent of the temperature and signal rate. Several tests have been carried out on the MAD ASIC both on the bench and in the field in various configurations. For the analog section an average gain of about 3.7 mV/fC was found for bare

chips with $(1370 \pm 48)/\text{pF}$ electrons ENC. Another key characteristic for operation with low signals is the crosstalk, which is less than 0.1%; moreover, the baseline restorer and the comparator offsets sum up to less than ± 2 mV total error.

The chip performance is somewhat degraded when it is mounted on a front-end board (Fig. 143): the gain reduces to 3.4 mV/fC and noise and crosstalk increase to $(1850 \pm 60)/\text{pF}$ electrons and 0.2%, respectively. These effects are caused by the input protection network, which is made of an external resistor and diodes that together with 100 μm gaps included in the PCB are capable of dissipating the energy stored in the 470 pF capacitors that connect the detector wires (biased at 3.6 kV) to the ASIC inputs. This protection is effective even in the case of repeated sparks (ASICs survive >105 sparks at 3.6 kV amplitude with 1 spark/s on all channels).



Figure 143: Front-end board (FEB).

The above figures enable front-end operation at a threshold well below 10 fC (the value used during test beams was 5 fC) when connected to the detector, which has a maximum capacitance of 40 pF. The propagation delay of the chip is less than 5 ns with little dependence on signal amplitude (time walk is less than 7 ns). The rate capability of the MAD ASIC largely exceeds demand: 800 fC charge pulses (just below saturation) at 2 MHz rate do not affect the efficiency in detecting 5 fC interleaved signals, so there is a wide safety margin with respect to the total rate (about 10 kHz) foreseen per drift tube during CMS operation.

Finally, the radiation tolerance and overall reliability of the front-end board and associated electronics were investigated [140]. Radiation testing involved a series of tests with thermal and high-energy neutrons, protons, and γ -rays to simulate the behaviour in a CMS-like environment. The results can be summarized in latch-up immunity (undetected SELs even with heavy ions on naked dies), very little sensitivity to SEUs (only a few thousand spurious counts/channel calculated for the whole detector lifetime), and tolerance to total integrated dose orders of magnitude higher than foreseen in 10 years of CMS operation. In addition, accelerated ageing in a climatic chamber at 125 $^{\circ}\text{C}$ was carried out for >3000 hours on 20 FEBs and related circuits without revealing any fault.

Read-out electronics The electronics of the read-out system of the CMS DTs is responsible for the time digitization of the signals generated in the drift chambers and for the data transmission to higher levels of the DAQ system. The time digitization of the signals is performed at the Read-Out Boards (ROB [141]), located in minicrates, as described in Sect. 8.1.3, together with the DT muon trigger electronics. Two FTP cables are used to send digitized data from each minicrate to the 30 m far away rack in the towers beside the CMS wheels where the Read Out Server (ROS) boards are located. Each ROS merges data coming from chambers of one wheel sector through a 70 m optical link to the Detector Dependent Units (DDU) in the USC55 control room, performing the multiplexing of 1500 copper links into 60 optical links. All Read Out Systems have been developed according to the requirements both of the expected trigger rates (100 kHz) at the high luminosity of LHC, with an average occupancy of 0.76% in the whole detector, a L1 trigger latency of 3.2 μs , and of the need of operating in an environment where the integrated neutron fluence will reach 10^{10} cm^{-2} in 10 years of activity.

Read-Out Boards are built around a 32-channel high performance TDC, the HPTDC, which is the third generation of TDC's developed by the CERN Microelectronics group [137], and it has been implemented in IBM 0.25 μm CMOS technology. This highly programmable TDC is based on the Delay Locked Loop (DLL) principle, providing a time bin of $25/32 \text{ ns} = 0.78 \text{ ns}$, which corresponds to 265-ps resolution, when it is clocked at the LHC 40.08-MHz frequency. This time resolution is enough to obtain a single wire position resolution of 250 μm .

The number of HPTDCs per ROB has been decided following a compromise between the number of unused channels when the granularity is too small and the multiplication of common components when it is too big. Finally, each ROB has 4 HPTDCs connected in a clock synchronous token ring passing scheme, where one of them is configured as a master to control the token of the read-out data chain. The token ring scheme is designed following

a failsafe mechanism, which avoids that the failure in one of the TDCs interrupts the whole ROB operation. Both hardware and software bypassing systems have been implemented.

Depending on the chamber type, the number of channels is different and accordingly, the number of ROB per minicrate. The smallest minicrate has 3 ROB and the biggest has 7. They are all connected to the Control Board (CCB) that manages, among others, the Timing and Trigger Control (TTC) signals. As can be seen in Fig. 144, Trigger Boards (TRB), located inside the minicrate are connected to the ROB to receive TTL translated hit signals. As described at the beginning of this Section, FTP cables connect the output of the ROB to the ROS boards [138].

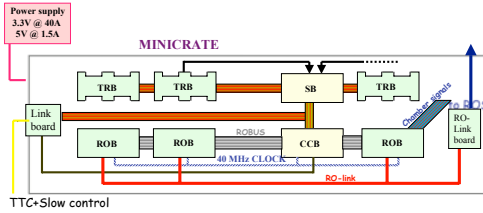


Figure 144: Sketch of the read-out and trigger electronics located inside a minicrate.

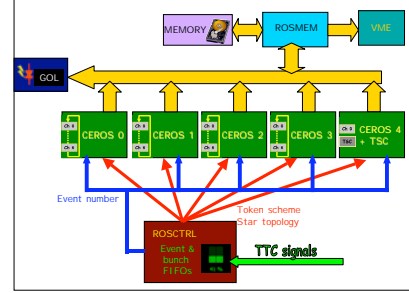


Figure 145: Sketch of the splitting of the 25 channels in four groups.

Located in the barrel tower racks there are 60 ROS boards, 12 per wheel, 1 per sector (four to five minicrates), so each ROS receives 25 channels of the LVDS copper ROB-ROS link. These 9U boards have to multiplex data coming from the ROB, adding necessary information of ROB number, link status and other information, and send them to the DDU through a fast link. Another feature of the ROS board, is that it also includes a power supply protection circuitry, current and temperature monitoring, and a 512 kB memory to test and perform data flow snapshots for traceability in case of transmission errors. In Fig. 145 it can be seen how the 25 channels are split in four groups of six channels each, so-called CEROS, controlled by an FPGA that manages the FIFO read-out performing a pooling search for the next event to be read. These FPGAs also filter the events, discarding headers and trailers of those channels without timing or error information, reducing accordingly the data overhead.

A test performed on 10 prototypes, keeping them in an oven at 125 °C for 2000 hours, in order to simulate 10 years of CMS activity, gave no faults. The DT Front-End-Driver (FED, also called DDU [142]) is the last component of the DT read-out electronics. The DT FED system consists of 5 VME64X 9U boards housed in the CMS service cavern; each board collects data through serial optical connections from 12 ROS, corresponding to an entire DT wheel, and transmits a formatted event fragment to the CMS common DAQ through a S-Link transmitter module. The synchronization with the trigger system is guaranteed by the TTC network, providing the LHC timing signal, the L1 trigger accept and fast commands, that are distributed to the different parts of the FED board. The layout of the board is depicted in Fig. 146.

The data rate in each DT FED board is limited by the maximum rate the CMS DAQ can accept from an S-Link connection (about 200 Mbytes/s). The number of boards has been chosen to deal with the expected DT event size (7 kbytes/event) at 100 kHz trigger rate.

High Voltage and Low Voltage systems The CAEN SY1527 universal multichannel power supply systems are used to supply high and low voltages (HV and LV) to the muon DT chambers. The basic modules of the DT HV system consist of the A876 master boards and A877 remote boards. A maximum of 8 A876 master boards can be housed in the SY1527 mainframe. Each of them supplies high voltages and low voltages, controls and monitoring, to a maximum of 4 independent A877 remote boards, each one powering one DT chamber (two in the case of MB4 of Sectors 4 or 10). The A877 remote boards are located in a separate non-powered mechanical crate sitting in racks on the towers next to the wheels. The A876 delivers to each A877 HV board: a positive HV in the range from 0 to +4.2 kV (2.5 mA maximum output current), a negative HV in the range from 0 to -2.2 kV (1 mA maximum

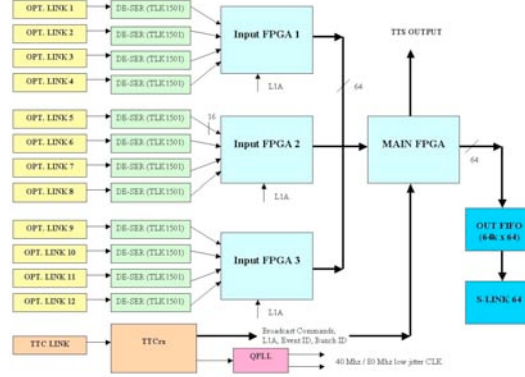


Figure 146: Scheme of the DDU architecture.

output current), a dual LV of $\pm 15\text{V}$ (1.5 A maximum output current). The A877 HV outputs are subdivided into 12 groups (8 for the special A877 boards powering the MB4 chambers) conventionally called macro-channels. Each macro-channel supplies 4 HV channels per layer: 2 anodes (the wires of each layer are divided into 2 groups), 1 strip and 1 cathode. For all HV channels, the maximum output current is hardware limited to $100\text{ }\mu\text{A}$.

The DT LV system uses three different types of CAEN Easy3000 modules: the A3009 to provide V_{CC} and V_{DD} voltages to the chamber Front End electronics and the V_{CC} voltage to the mini-crates, the A3050 for the mini-crate V_{DD} and the A3100 to power the Sector Collector Crates. The control of the Easy3000 power supply system is done remotely using a branch controller (Mod. A1676A) plugged in a SY1527 mainframe located in the control room. Each A1676A branch controller can handle up to 6 Easy3000 crates. The Easy3000 crate is powered by external 48 V DC that is provided by the CAEN AC/DC converter A3486S module. The following voltages are delivered:

- $V_{CCMC} = 5.8\text{ V}$, software current limit $i_0 = 3\text{ A}$;
- $V_{DDMC} = 4\text{ V}$, software current limit $i_0 = 30\text{ A}$;
- $V_{CCFE} = 5.2\text{ V}$, software current limit $i_0 = 3\text{ A}$;
- $V_{DDFE} = 2.6\text{ V}$, software current limit $i_0 = 4\text{ A}$;
- $V_{SC} = 2.6\text{ V}$, the current limit depends on the number of SC and ROS boards plugged in the crate.

8.1.4 Chamber assembly, dressing, and installation

Chamber assembly Mass chamber assembly was started in January 2002 and was fully completed (spares included) in June 2006, with a constant production rate for all the four production sites involved. The collection of a good sample of cosmic muons allowed full testing of a constructed chamber, before sending it to CERN. In each laboratory cosmic-ray events were triggered by an external scintillator system which covered the full acceptance of the chamber. Since final minicrate electronics was not available during chamber assembly, drift times were measured with external TDCs and a custom DAQ. Typical trigger rates were 50–100 Hz, resulting in $\approx 10^6$ events in a few hours. With such a large data sample it was possible to spot and cure problems which could not be detected in previous tests, like disconnected cathodes (Figs. 147 and 148) and disconnected strips (Fig. 149). Beside efficiencies, other relevant working parameters are measured from cosmic-ray data, as calibration stability (drift velocity measured to be stable within $\pm 5\%$), deviation from linear drift parametrization, measured to be well within $100\text{ }\mu\text{m}$ (Fig. 150), wire positions and comparison with CCD measurements during assembly, relative alignment of layers, and noise (typically below 100 Hz per cell, see Fig. 151).

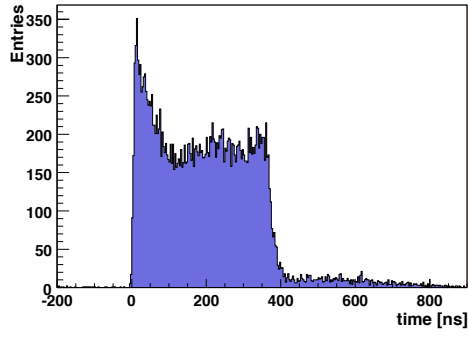


Figure 147: Drift Time distribution of a good cell. Wire position corresponds to time 0.

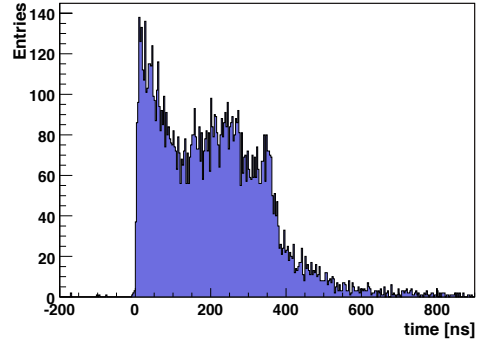


Figure 148: Drift Time distribution of a cell with a disconnected cathode.

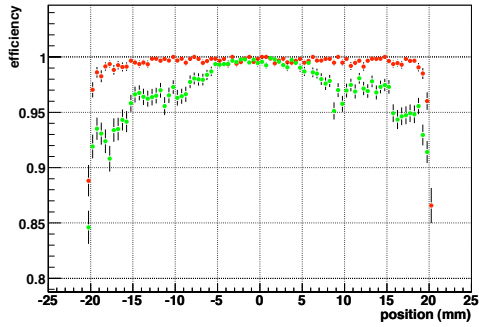


Figure 149: Single cell efficiency for a good cell (red dots) and for a cell with a disconnected strip (green dots).

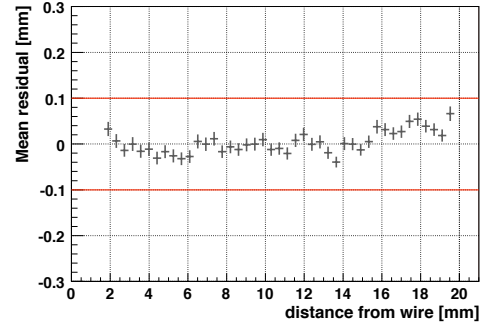


Figure 150: Residuals as a function of the distance from the wire, indicating a linear dependence well within $\pm 100 \mu\text{m}$ in the full cell range.

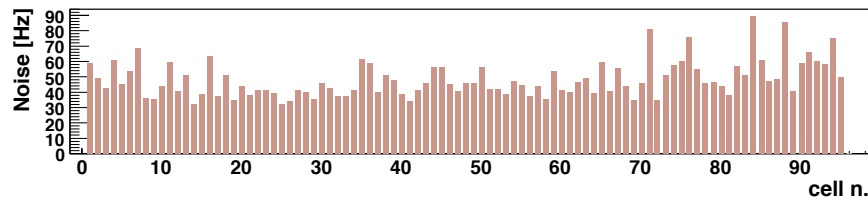


Figure 151: Typical hit rate distribution in one layer, as measured during the test of a SuperLayer with cosmic rays. This rate is dominated by noise and is typically below 100 Hz per cell.

Chamber dressing All chambers, built and fully tested at the production sites, were sent to CERN for final testing and commissioning prior to installation in the experiment. Since the arrival of the first chamber at CERN (an MB2 type chamber arrived in summer 2000, prepared for a test beam [143]), a total of 272 more chambers (including spares) have been received from all four production sites, leading to a continuous workflow of dressing and testing.

At a first stage chambers are assigned to a particular position in CMS, depending on their orientation. Before any test, each chamber undergoes the optical alignment procedure described in Sect. 8.1.4. After going through the alignment procedure, chambers are equipped with gas components (cooling pipes, gas manifolds, PADC pressure meters), HV cables and additional items like stickers, protectors, grounding straps, etc. Basically all components except for the minicrates (Sect. 8.1.3), minicrates-related items and external protections are installed at this stage.

After dressing the chamber the following tests were performed:

- High voltage long term tests;
- Gas tightness tests;
- Cosmic-muon tests.

The high voltage long term test consists of a continuous monitoring of the high voltage performance (electric current) under the nominal values for all components (3600 V for wires, 1800 V for strips, and -1200 V for cathodes) for a minimum of 6 weeks. The time constant of the chamber with the final gas connections is also computed as a measurement of gas tightness. No significant degradation has been observed with respect to the values measured at the sites.

Finally a cosmic test stand has been set up with trigger scintillators, independent cabling, LV and HV supplies and several HPTDCs, capable of measuring one chamber at a time, and registering several millions of triggers in a few hours. The later analysis of these data allows the recognition of almost all kind of problems related to the chamber itself and its internal electronics.

Once this first certification step is passed, the chamber dressing is completed and the chamber is declared ready for minicrate installation (Sect. 8.1.3). All signal cables from the chamber to the minicrate are installed and tested, and then the minicrate itself is inserted. The performance of the minicrate is tested at this stage for the first time together with a real chamber. All internal connections are checked, as well as configurability and data processing performance of the full local electronics chain.

At this point the chamber has passed all tests and can be considered *ready to install*. Last dressing steps are performed (installation of carters and additional protections) and then DT chambers are coupled together to RPCs, forming an installable barrel muon package.

Chamber survey To determine the chamber positions in the CMS coordinate system and to follow their movements a position monitoring Alignment System was built (Sect. 8.4). All the 250 DT-chambers positions are recorded by this system via optical connections using LED light sources mounted on the chambers and specially designed video-cameras fixed to the return yoke of the barrel. Four LED-holders called forks are mounted on the side-profile of the honeycomb structure (two per side), using the rectangular 50×65 mm² tube as the light-passage. Each fork has 10 LEDs, 6 and 4 respectively on each side of the fork. The control of the LEDs (on-off, current) is performed via I²C bus system integrated in the minicrate.

The primary aim of the alignment is to give the positions of the anode wires but this is not directly possible. On the other hand all the wire positions are measured with respect to the SL corner during the chamber construction. To establish the connection between the LEDs and the corner blocks a calibration bench was built at the CERN ISR site (Fig. 152). This bench had two functions. The first one was to measure the corner block positions with respect to each other, allowing the full-chamber geometry, including the relative positions of the superlayers in 3D and their planarity, to be measured. The second function was to measure the LED positions in the chamber coordinate system stretched on the corner blocks. The bench contained video-cameras that could observe the LEDs and photogrammetry targets to measure the corner blocks with respect to the LEDs by photogrammetric methods. The bench allowed us to measure all the types of chambers from MB1 to MB4. The full bench was calibrated and recalibrated before each chamber calibration campaign by the CERN survey group. Also additional LEDs were mounted on the bench to detect any significant deformation of the bench itself. The precision of the bench-measurements for the locations of the corner blocks are $\sigma < 40$ μ m, the position of the forks relative to the



Figure 152: The chamber calibration bench in the CERN ISR Lab.

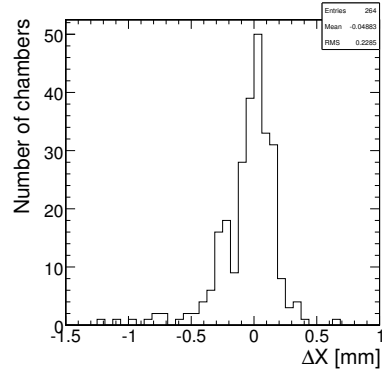


Figure 153: Distribution of the ϕ -deviation from the nominal design value for the two Φ -type superlayers.

chamber is $\sigma < 70 \mu\text{m}$. Both values are within the acceptable range defined by physics requirements. Figure 153 shows the result of the residual distributions of Φ -type SL corner block positions for all measured chambers.

Gas system For the DT chambers a safe and inexpensive gas mixture is used, namely Ar/CO₂ in the ratio 85/15 volume. The gas is distributed in parallel to all drift cells in four steps: (1) the main line is split into 5 lines to feed each of the 5 barrel wheels; (2) on the wheel it is split into 50 lines to feed the 50 chambers on the wheel; (3) on the chambers it is split into 3 lines to feed the 3 SuperLayers; (4) within the SL a long tube with small holes distributes the gas over the drift cells. The nominal flow is 50 l/h for each chamber. Due to the large total number of 250 chambers, a closed loop circuit with a cleaning station is used. It is foreseen to add about 10% fresh gas daily. The gas system is run at constant absolute pressure inside the chambers, to avoid any variation of the drift velocity inside the chamber. The pressure is regulated for each wheel. There are flowmeters at the inlet and outlet of each line at the gas distribution rack on the wheel. The gas pressure is also measured with two sensors at the inlet and outlet gas manifolds on the chambers, amounting to 1000 sensors in total. They should ensure a safe and redundant measurement of the pressure at every chamber, as needed for inaccessible chambers. To be able to analyze the gas actually present in every chamber, a return line brings a sample to the gas room. There is one such line per wheel and a remotely controlled multiway valve permits the selection of the desired chamber or the gas arriving at the wheel, for analysis.

The gas is also analysed independently for each of the 5 wheels and consists of a measurement of the oxygen and of the humidity content of the gas, as well as a direct measurement of the main parameter of the DT chambers, the drift velocity. The drift velocity is measured with a small drift chamber (VdC), which features a very homogeneous, constant, known and adjustable electric field in a region where two thin beams of electrons from a beta-source cross the chamber volume and trigger a counter outside the chamber. The distribution of times between the trigger signal and the signal from the anode wire of the chamber is recorded. The distance between the two beams being well known by construction, by measuring the distance in time between the peaks from the two sources, one reads directly the drift velocity. Variations of the drift velocity can be monitored accurately by accumulating data for about 5 min. The data shown in Fig. 136 demonstrate that the absolute values measured with the VdC agree with the expectation. The special merit of a direct monitoring of the drift velocity is that one does not need to know which impurities are affecting the drift velocity to monitor it.

Chamber installation in CMS The main installation tool is a platform (*cradle*) with the same support rail as in the iron pockets that can be anchored to interface pads mounted on the wheel (Fig. 154). Pneumatic movements allow the precise alignment of the chamber with respect to the iron pockets. Once the rails on the cradle and those in the iron pockets are aligned, the chamber is pushed into position with an electric motor. The installation of the chamber in the MB1 station of Wheel -2 sector 11 is shown in Fig. 154.

The first chambers were installed in the bottom sectors of Wheel +2 in July 2004 and the surface installation was



Figure 154: Installation of MB1 station on Wheel -2. The yellow frame is the *cradle* used to insert the chamber in its location inside the iron slot.

completed in Wheel -2 in December 2006. Installation completion, for sectors 1 and 7, which cannot be filled on surface since these parts of the wheel are reserved for handling during the heavy lowering operation, is made underground.

8.1.5 Chamber performance

Chamber and trigger performances have been thoroughly analyzed at various stages, on prototypes before mass production (with and without external magnetic field [144]), on final chambers with test beams and with the CERN Gamma Irradiation Facility [143, 145, 146], with cosmic-rays both at production sites and at the commissioning of the installed chambers, and finally with the so-called *Magnet Test and Cosmic Challenge* (MTCC) in 2006, where part of the DT system, completely installed and equipped with final hardware, was tested together with the final CMS DAQ system.

Test beam data: chamber performance Several dedicated muon test beam runs were set up in order to test chamber performances under different conditions. Single cell spatial resolution could be determined [143] simply by the dispersion of the $MT = (t_1 + t_3)/2 + t_2$ distribution, MT being the meantime obtained from the time of the signals (t_1, t_2, t_3) generated by the incoming muon in three consecutive, and staggered layers. The smoothness of the drift time box and the fast drop of the trailing edge (Fig. 147) are both signs of the saturation of the drift velocity. Under the assumption that the time resolution is the same in all layers $\sigma_t = \sqrt{\frac{2}{3}} \cdot \sigma_{MT}$, one can easily observe an average spatial resolution of $170 \mu\text{m}$ (Fig. 155). Using a Silicon Beam Telescope, it was possible to measure the deviation from the extrapolated hit on the SL and the reconstructed position. As can be seen in Fig. 156 deviations from linearity are within $\pm 100 \mu\text{m}$, which corresponds to a trigger jitter smaller than 5 ns. Silicon Beam Telescope data can also be used to measure precisely the chamber efficiency (Fig. 157). The geometrical acceptance associated with the presence of the I-beam is clearly seen by the drop of efficiency in that region. In the rest of the cell the efficiency is always higher than 99.5%. The typical intrinsic average noise, as measured during chamber construction or, during a test beam in dedicated random trigger runs [145], is shown in Fig. 158. It is reasonably stable, at values of $\approx 50 \text{ Hz}$, and does not vary much with channel number. Also the effect of higher noise levels, generated at the CERN Irradiation Facility, by photo-conversion in the chamber aluminum walls at chosen rates, both on reconstruction and trigger efficiency, was studied. The SL segment reconstruction efficiency is shown in Fig. 159 for various filter values of the gamma source, showing no significant dependence of the reconstruction algorithm on the gamma irradiation level, even at noise rates higher than the maximum levels expected in

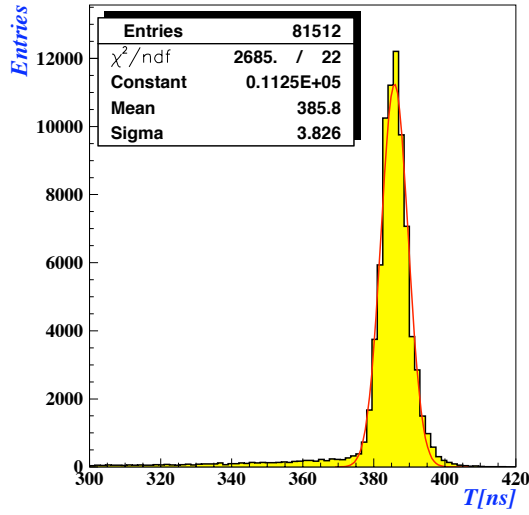


Figure 155: Gaussian fitted MT distribution with $\sigma_t = 170 \mu\text{m}$. The position of the MT peak allows the determination of an average drift velocity of $54.4 \mu\text{m/ns}$.

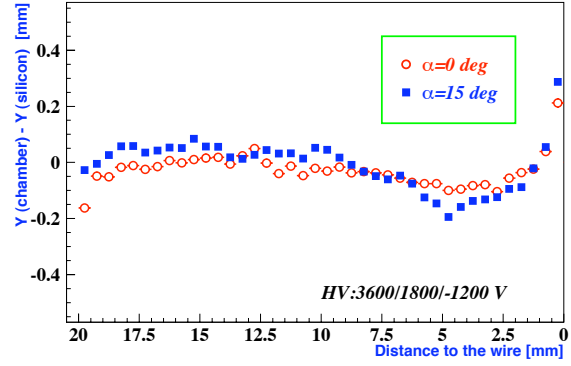


Figure 156: Deviation from linearity as a function of the distance to the wire for tracks with angles of incidence $\alpha=0^\circ$ and 15° . The cathode is centred at 21 mm.

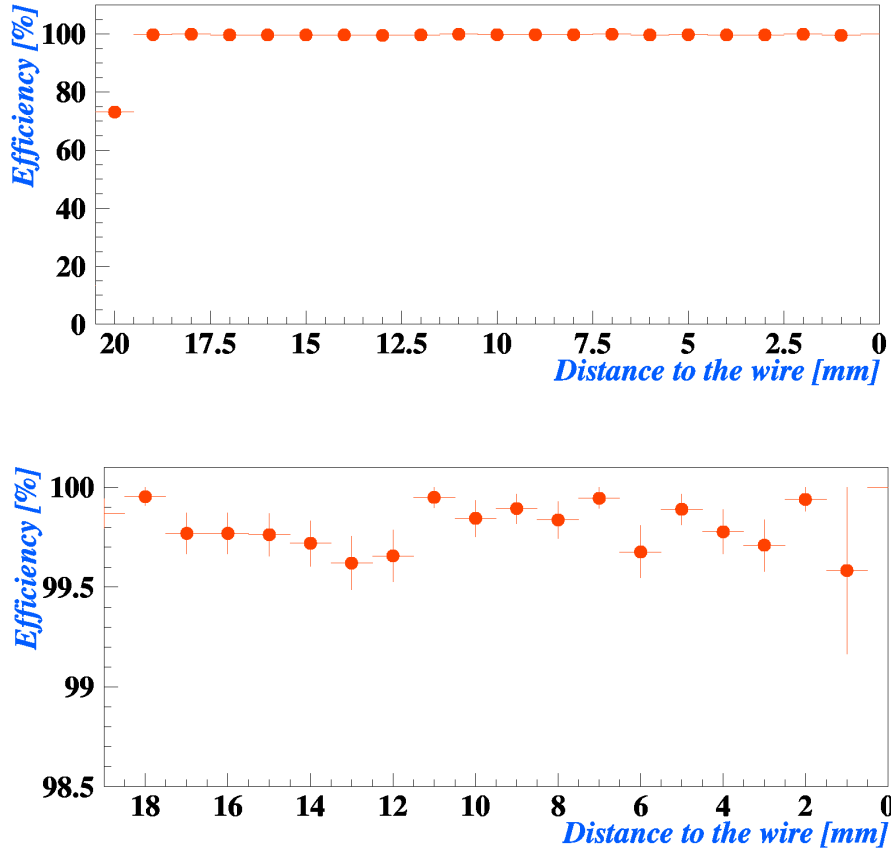


Figure 157: Efficiency as a function of the distance to the wire (top), for tracks orthogonal to chamber surface; (bottom) with an expanded scale excluding the I-beam region.

any DT chamber during normal LHC operations. Since the chambers are operated in the iron yoke of CMS, where

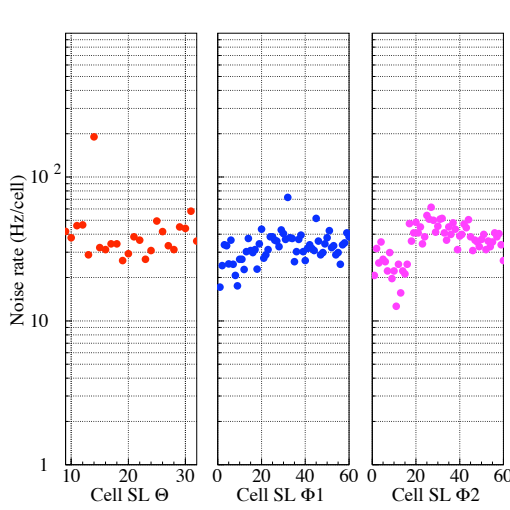


Figure 158: The noise cell occupancy as a function of the cell number, for the three SLs.

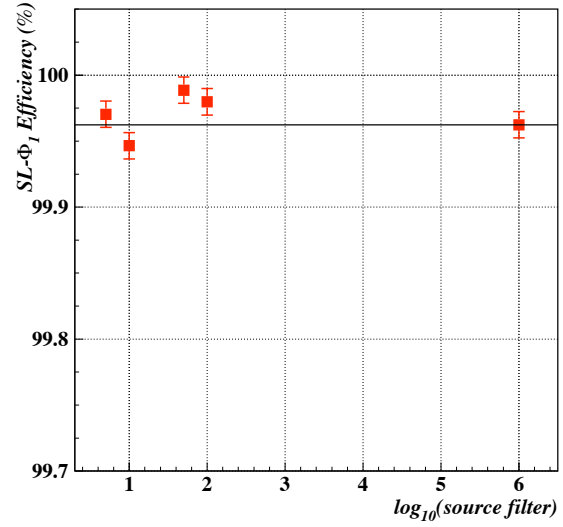


Figure 159: The cell efficiency in superlayer Φ_1 as a function of the irradiation filter value. The smallest filter value provides a background a factor 2 larger than the maximum one to be expected during LHC operation in any of the DT chambers in CMS.

important stray magnetic fields are present, the impact of the radial and longitudinal components of the field were carefully simulated (finite element analysis program ANSYS, Fig. 160), and tested both at dedicated muon test beams where the chambers were operated inside a magnetic field [145, 144], and during the CMS Magnet Test and Cosmic Challenge, with cosmic rays. The distortion of the electron drift lines caused by a field of 0.5 T parallel to the wires can be seen in Fig. 161. This distortion can be roughly approximated by a rotation of the drift lines around the wire, simulating a rotation of the drift cell with respect to the direction of the incident particles. A change in the maximum drift path and time is generated, together with a drop of efficiency for inclined tracks which go through the I-beam region where the drift lines do not reach the wire. The staggering of the four layers minimizes the impact on track measurement of this last effect. In the case of a homogeneous magnetic field along the wire, the main consequence would be an effective lower drift velocity, but in CMS the magnetic field is not homogeneous. Figure 162 shows the drift velocity values obtained for several magnetic field values in the case of perpendicular tracks. A variation in the drift velocity of around 3% is observed from $B=0$ to 0.5 T. This corresponds to a change in the maximum drift time of less than 12 ns, which is acceptable both for reconstruction and trigger efficiency.

The effects of the magnetic field on linearity are not very important below 0.3 T but increase dramatically for higher fields, mainly near the I-beams. The resolution is also slightly deteriorated by low magnetic fields, but it is still better than $300 \mu\text{m}$ below 0.3 T. In CMS the magnetic field component parallel to the wires measuring the coordinate in the bending plane is expected to be below 0.1 T, with only very limited regions reaching 0.3 T. In such conditions, the results obtained confirm that the performance of the drift tube chambers fulfil the requirements.

Test beam data: trigger performance Like the chamber performance, the DT local trigger has also been tested extensively using test beam facilities at CERN [146, 147]. In particular, to fully test the performance of the trigger electronics, bunched beams having the same time structure as the LHC were used at the CERN SPS, producing high momentum muon tracks separated by multiples of 25 ns.

The bunch crossing (BX) identification efficiency is defined as the fraction of selected single muon events for which the local trigger delivered at least one trigger segment at the correct BX. This quantity was measured as a function of the muon momentum, and results are shown in Fig. 163. The measurement was also performed after inserting iron slabs, for a total depth of 15 cm, in front of the muon chamber. The effect of the iron absorber is the

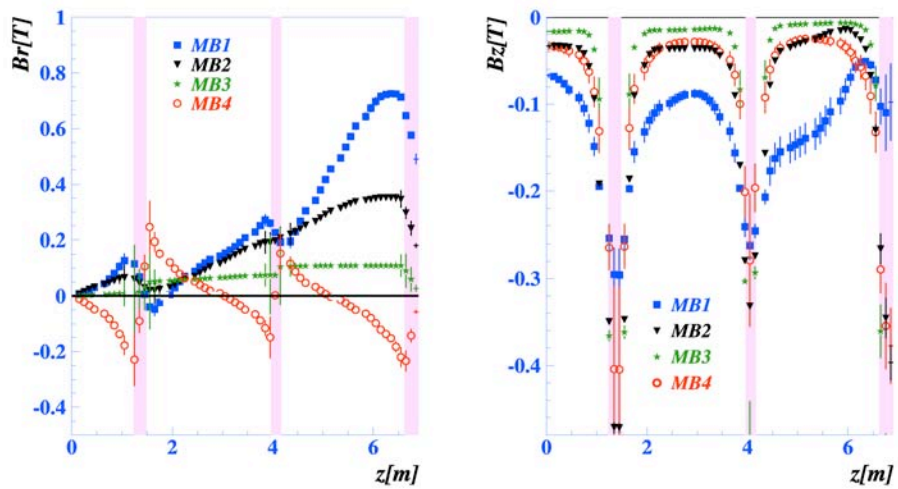


Figure 160: Radial (B_r) and longitudinal (B_z) components of the CMS magnetic field in the regions where the barrel chambers are placed as a function of the position along the beam direction (the centre of the detector is at $z=0$). Vertical bands indicate the separation between chamber wheels (in these particular regions B_z becomes significant). The biggest B_r values (0.7–0.8 T) occur in the MB1 region near the endcaps.

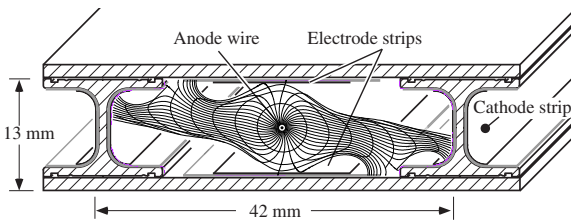


Figure 161: Simulation of the distortion produced in the drift lines by a 0.5 T magnetic field parallel to the wires.

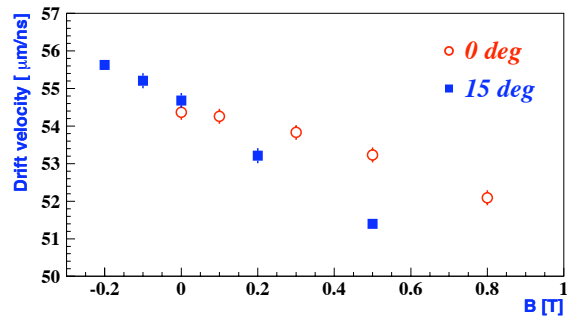


Figure 162: Drift velocity for several magnetic fields for perpendicular (0°) and inclined (15°) tracks.

enhancement of the probability for a high momentum muon to produce electromagnetic showers. Such a probability also increases as a function of the muon momentum. This has the effect to decrease the BX identification efficiency. Results are also shown in Fig. 163, superimposed on results without iron absorber. Ghosts are copies

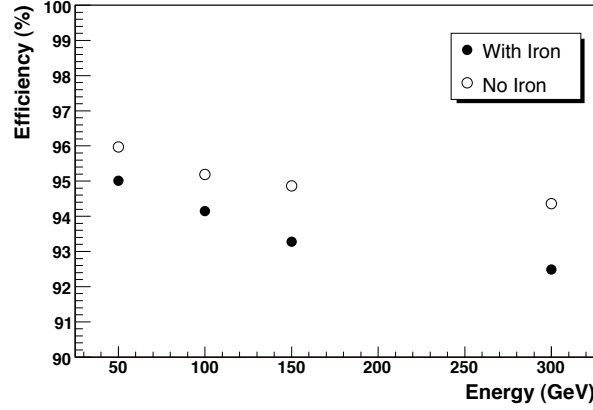


Figure 163: BX identification efficiency in a muon station in single muon events, as a function of the incident muon momentum, for events with and without the iron absorber placed in front of the muon station.

of the trigger segment at the correct BX, as well as fake triggers at the wrong BX. They may originate from wrong alignment of hits in a DT muon station, due to the presence of extra hits produced by electromagnetic cascades and δ -rays, or from redundancies in the trigger electronics. In the case of single muon events, if two trigger segments are delivered at the same BX by the local trigger system in a muon station, the second trigger is considered a ghost copy of the first one. Although generally with a poorer quality, ghosts at the correct BX reproduce the characteristics of the main trigger segment in terms of position and angle. The production of segments associated to a wrong BX, arising from wrong hit alignment, is intrinsic to the BTI algorithm (Sect. 9.2). In addition there are also cases in which the hit alignment is spoiled by δ -ray production or electromagnetic showering. Such fake triggers, which are called out-of-time ghosts, are almost entirely uncorrelated low quality segments, and are distributed over a wide range of BXs. The fraction of ghost triggers at the correct BX as a function of incident muon momentum, and the fraction of out-of-time triggers, as defined above, are shown respectively in Figs. 164 and 165 as a function of the incident muon momentum, for events with and without the iron absorber. The performance of the Φ Track

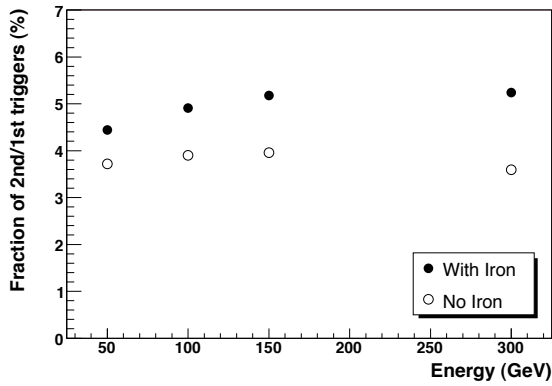


Figure 164: Fraction of ghost triggers observed in a muon station in single muon events, defined as the ratio of the number of second tracks over the number of first tracks, delivered by the local trigger at the correct BX, as a function of the muon momentum, for events with and without the iron absorber.

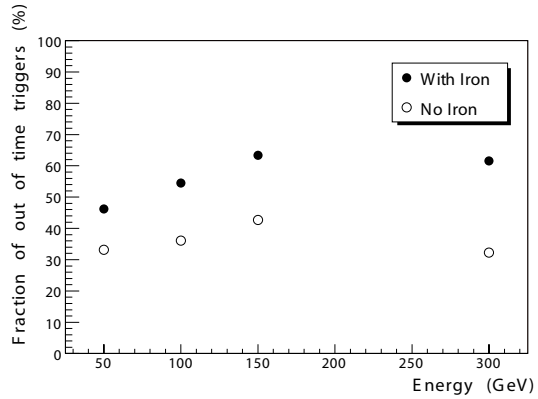


Figure 165: The fraction of out-of-time triggers in a muon station in single muon events, defined as the number of out-of-time trigger segments divided by the number of selected single muon events, as a function of the muon momentum, for events with and without the iron absorber.

Finder (PHTF, Sect. 9.2) was also tested within the same muon test beam at CERN [147]. The PHTF was used to

reconstruct muon trigger candidates using both muon stations, using ϕ -view local trigger primitives. Figure 166 shows the distribution of the BX assigned to the tracks found by the PHTF. The BX is correctly identified when its value is 24. Superimposed are the distributions of the same quantity determined independently by the local trigger in MB1 and MB3, as well as the distribution of the determined BX when a trigger segment with the same BX was delivered in coincidence in MB1 and MB3. It can be seen that the PHTF is fully efficient to deliver track candidates at the correct BX, whereas for out-of-time triggers the corresponding PHTF trigger rate is suppressed at the level of 1% or less.

A large fraction of the out-of-time triggers is due to real muons crossing the experimental apparatus at a BX different from 24, and which are correctly reconstructed by the PHTF. This is confirmed by the fact that the trigger segments that are matched together to form such tracks, are mainly of the type HH (four hits in both Φ -type SL), thus indicating a real muon track. Figure 167 shows the PHTF efficiency to reconstruct a trigger track in events with a MB1 and MB3 coincidence as a function of the BX. Superimposed are the efficiency to reconstruct a trigger track when the two trigger segments are both of HH type, and both of L (three out of four hits in a SL) quality respectively. The correct BX is 24. The PHTF efficiency for HH coincidences is $99.7 \pm 0.1\%$ and is practically

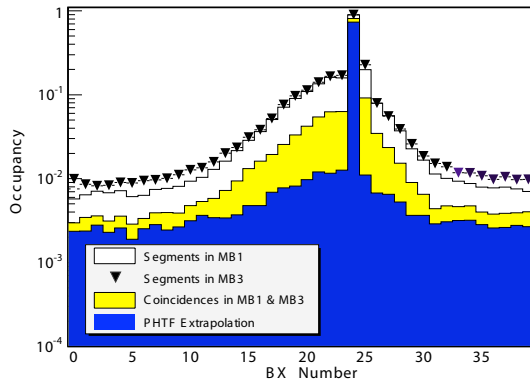


Figure 166: Distribution of the BX assigned to the tracks found by the PHTF. The BX is correctly identified when its value is 24. Superimposed are the distributions of the same quantity determined independently by the local trigger in MB1 and MB3, as well as the distribution of the determined BX when a trigger segment at the same BX was delivered in coincidence in MB1 and MB3.

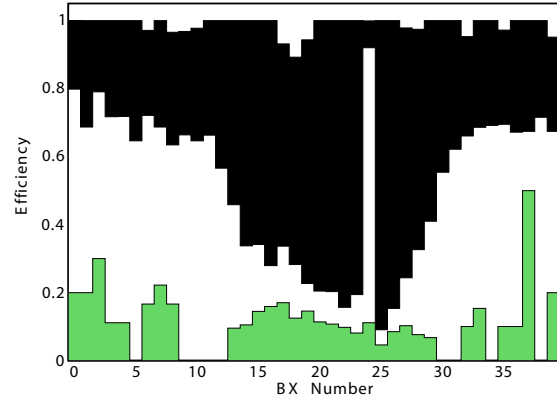


Figure 167: Efficiency to reconstruct a trigger track by the PHTF, as a function of the BX (white), in events with a two stations (MB1-MB3) coincidence. The correct BX is 24. Superimposed are the efficiency to reconstruct a trigger track when there is a coincidence of two trigger segments both of HH quality (black), and a coincidence of two trigger segments both of L quality (light green).

constant for any BX. This fits with the expectations, as such tracks are real muons crossing the apparatus. On the other hand, when the trigger segments have a low quality, which is typical for fake triggers, the PHTF ghost suppression is very effective. The rejection power for ghosts (L coincidences at $BX \neq 24$) is 9.5 ± 0.4 . Therefore, although the out-of-time local trigger rate in a single station is rather high (as shown for example in Fig. 165), the PHTF is very effective in ghost rejection.

Commissioning of installed chambers After installation in their final positions in the five CMS barrel wheels in the CMS surface hall, the chambers, including read-out and trigger electronics, were tested again with the goal to identify potential problems before final cabling. Given the previous testing stages, it was mainly a test of the electronics and connections. Possible damages or loose connections as a consequence of installation (where chambers undergo some mechanical movement) could also be detected at this stage. This commissioning step was performed before final cabling (since the cables cover the minicrates and would prevent access for potential repairs) and involved all chambers in the ten sectors (84% of the full system). The commissioning consists the following steps:

1. connection to the necessary voltages and to the test-stand, which handles the trigger and read-out, while verifying the functionality of the connectors and cabling of the chamber;

2. check of the minicrate performance via test of the internal connections and functionality monitor of all electronics boards inside the minicrate;
3. T0 determination (the starting point) of the drift time spectrum with test pulses. The T0 is specific for every cell, its cell-to-cell variation within a chamber is of the order of $\approx 1\text{--}2$ ns;
4. cosmics data taking in different trigger configurations;
5. analysis of the cosmics data and verification of the chamber and electronics performance.

Depending on the amount of repairs, between two and five chambers per week were commissioned. Cosmic muon tracks were recorded in auto-trigger mode in different trigger configurations. Trigger rates vary from 80 Hz to 600 Hz per chamber depending on the sector inclination, the trigger configuration, and the chamber type. Higher level trigger components (tower electronics) as well as RPC connections are not tested at this stage since they require full cabling to the tower electronics.

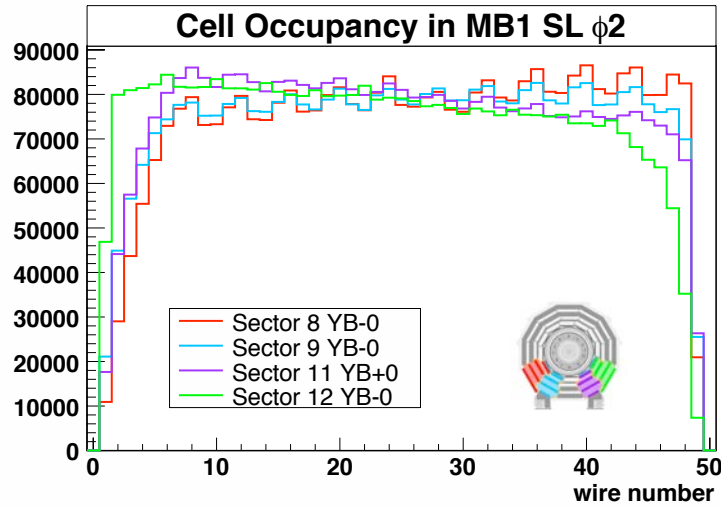


Figure 168: Sum of the occupancy of the four layers in MB1 $\Phi 1$ superlayer for sectors 8, 9, 11 and 12. The sectors have different inclinations as shown on the right. Material inside the yoke along with the iron yoke itself shield partly the soft cosmic ray muons.

The chamber orientation with respect to the incoming cosmic-ray flux is purely horizontal only in sectors 4 and 10. This yields a reduction in the uniform occupancy near the edges of the Φ -type superlayer for sectors near the vertical. An example comparing the bottom sectors 8, 9, 11 and 12 is shown in Fig. 168. The occupancy per wire is summed up over the four layers of superlayer $\Phi 1$ with two dead cells causing an obvious drop of $\approx 1/4$ in the occupancy. The reduced geometrical acceptance at the chamber edges is caused by the mainly vertical direction of cosmic ray muons along with the shielding by the iron yoke on the muon soft momentum spectrum.

Data recorded with the highly selective trigger condition, 4 hits in both Φ -type superlayers or 3 hits in one and 4 in the other, are used to calculate the efficiency. The cell efficiency is calculated from the ratio of reconstructed tracks traversing the considered layer (N_{track}) and its subset with hits found in the traversed cell or its 2 neighbours of the considered layer (N_h):

$$\epsilon_{\text{Layer}} = N_h / N_{\text{track}} \quad (8)$$

requiring ≥ 5 hits in ϕ and ≥ 3 hits in z . The left sketch in Fig. 169 illustrates the combined track fit in the case of the ϕ projection. Because of the normalization to the number of tracks, the reduced occupancy near the chamber edges does not play a role. The cell efficiency is almost constant across the chamber, usually $\geq 98\%$, as seen, for example, in Fig. 169.

Similarly the reconstruction efficiency is determined as the fraction of reconstructed tracks when requiring ≥ 7 hits per track, an example of which can be seen in Fig. 170. Here we should remark that the overall number of dead

cells, as measured during chamber commissioning at CERN, amounts to $\approx 0.2\%$ of the total number of channels (298 out of 171 852).

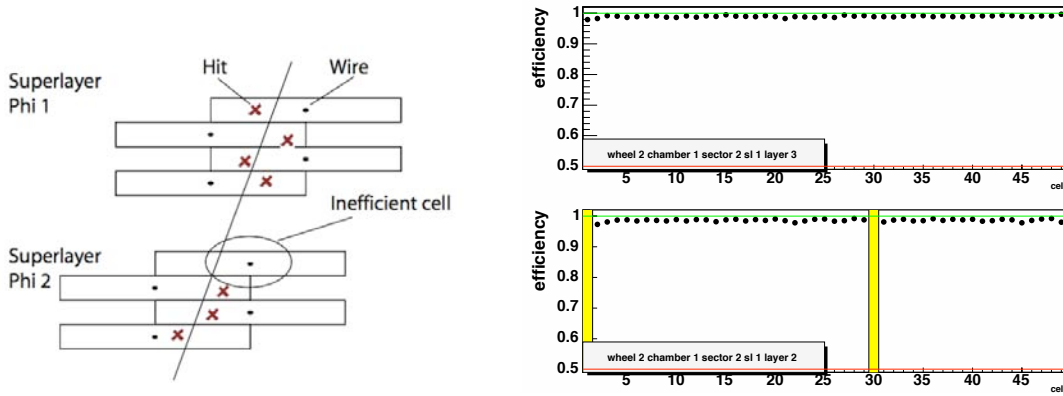


Figure 169: Cell efficiencies for two of the four layers of $\Phi 1$, MB1 sector 2.

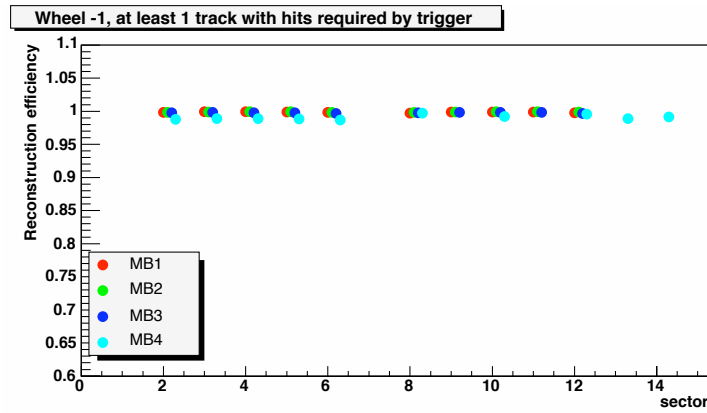


Figure 170: Reconstruction efficiency with 7 points in both Φ -type superlayers, according to the trigger condition, requiring either two 4-hit segments in both of the two Φ -type superlayers or one 4-hit and one 3-hit segment.

Results from Magnet Test and Cosmic Challenge A further important test of the muon system with emphasis on integration into the overall CMS DAQ and Trigger system, is the aforementioned Magnet Test and Cosmic Challenge (MTCC), performed at CERN during summer 2006. For the first time, the three muon subsystems were operated together. The DT system made use of the complete read-out and trigger chain with final hardware, and the recorded data allowed the study of cosmic muon tracks in a magnetic field. Another important task was the generation of a cosmic-muon trigger for the read-out of all CMS subsystems participating in the MTCC. Several goals were accomplished by the DT system during the MTCC:

1. check of chamber performance and read-out;
 - check the effect of the magnetic fringe field on chamber performance in terms of HV behaviour, drift velocity, position resolution, and efficiency;
 - exercise the complete read-out chain from the chamber through the ROS-25 up to the DDU;
 - test of HV and LV in the final set-up. Integrate HV and LV control into central CMS services.
2. Trigger
 - operate the complete trigger chain with final hardware;
 - provide a cosmic-ray trigger to CMS;
 - check B field effect on trigger timing (i.e. bunch crossing identification), requiring the RPC trigger to set a reference T0.
3. Software

- integration of DAQ and Data Quality Monitoring into CMS;
- exercise the reconstruction software under realistic conditions.

4. Take data with other CMS subsystems.

For the DTs (as well as for the barrel RPCs) three instrumented sectors were read-out, the bottom sector 10 and the adjacent sector 11, both in YB+2, along with sector 10 in YB+1 (Fig. 171). This accounts for 14 DT chambers, corresponding to about 10 000 channels. Beside the cross-check of chamber performance previously

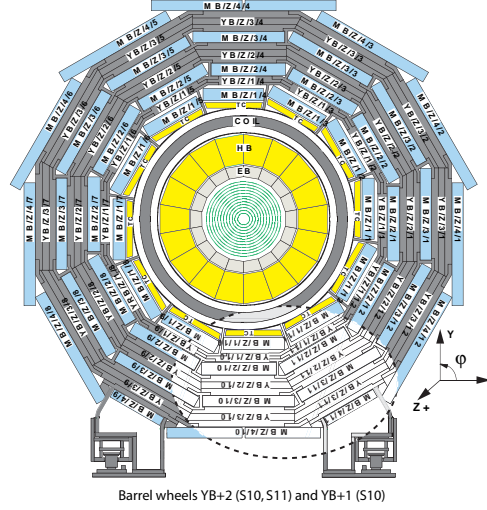


Figure 171: The MTCC exploited in the barrel region three sectors in wheels YB+1 (sector 10) and YB+2 (sectors 10, 11) instrumented with DTs and RPCs.

carried out in test beams, MTCC data provided a unique opportunity to test the reconstruction algorithms for different magnetic field strengths (Fig. 172(left)) and to observe for the first time tracks combined in different detectors (Fig. 172(right)). As described in Sect. 8.1.5, because of the radial component of the magnetic field

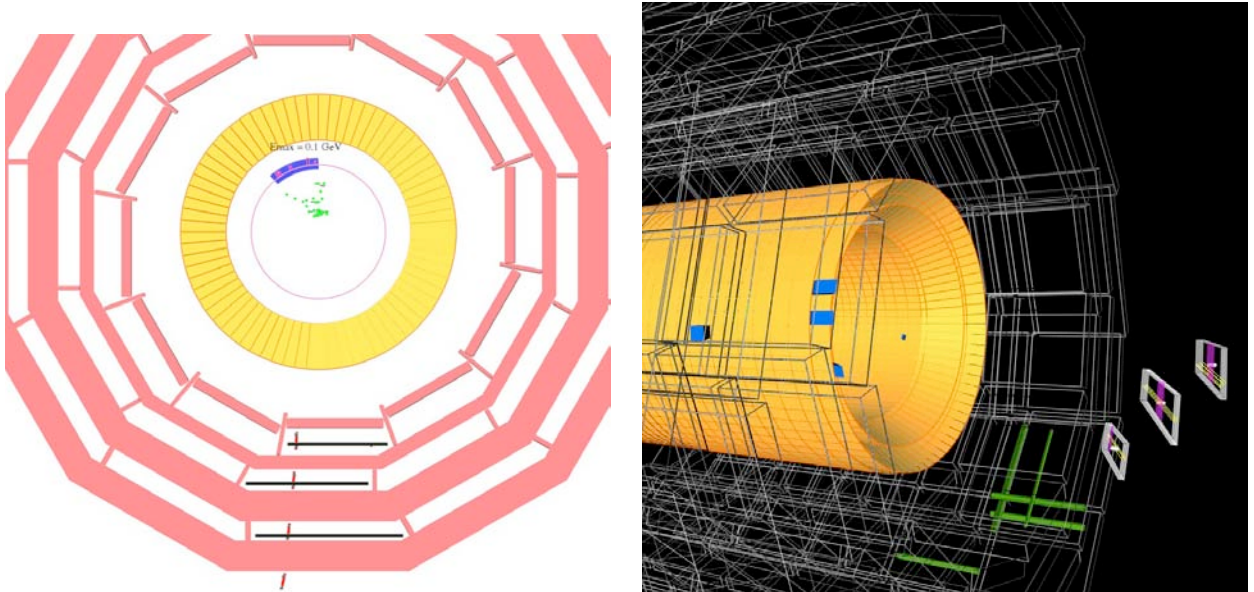


Figure 172: Event displays from the MTCC. Left panel: muon reconstructed in a DT sector in conjunction with Tracker activity. Right panel: muon track passing through both DT and CSC chambers.

between solenoid and endcap disks, the electron drift direction in the r - ϕ view changes, acquiring a Lorentz angle. The angle is increasingly larger as the B radial component increases along z . Thus signals generated by muon hits in a r - ϕ drift tube at a given distance from the wire but at different z positions will appear at different times. The effect has implications both for the trigger synchronization and for the muon track reconstruction and it should be calibrated out before LHC start-up. During the MTCC a total of 159 million cosmic-muon events (48 million DT

triggered) have been collected at several values of the B field. Data were collected at 0 T (as a reference), then at 2 T, 3 T, 3.5 T, 3.8 T and 4 T (93 million events at 3.8 and 4 T) which allowed a detailed mapping of the Lorentz angle effect in a MB1 and MB2 chamber. Some 15 million events at 0 T and 3.8 T (1.6 million DT triggered) have been taken with the MB1 local trigger configured to select only muon segments pointing to the CMS barrel centre (LHC beam interaction point): this sample is specific for trigger timing studies in the DT-CSC overlap, in particular also the muon time-of-flight is the same as in a LHC run.

While procedures for the synchronization of the DT system in stand-alone were studied in Phase I of the MTCC, in Phase II (during the magnet mapping operations) effort was put in tools for fine inter-synchronisation of the muon detectors (DT, CSC and RPC). In particular the analysis of the DT trigger data at the chamber output as function of RPC-originated L1A has proven to be sensitive to a few nanoseconds desynchronisation.

To study the efficiency of the DT Local Trigger (DTLT), events were selected by requiring the presence of the RPC triggers RBC1 (for wheel YB+1) or RBC2 (for RPCs in wheel YB+2) triggers. In such events, track segments were reconstructed in each muon station independently, whenever possible, using TDC hits. If more than one track segment was reconstructed in a given station, the one with the largest number of associated hits was taken. The efficiency of the DTLT was computed for each muon station separately, by counting events with a reconstructed muon segment, and comparing them with events which also had a trigger segment at any BX in the same muon station. Accepting a trigger regardless of its BX position was dictated by the fact that cosmic rays are likely to generate triggers in nearby BXs, as with non-bunched particles the system is intrinsically not synchronized. Only correlated trigger segments, namely of quality High-High (HH), High-Low (HL) or Low-Low (LL), as defined in Sect. 8.1.5, were released by the trigger sector collector and thus used to compute the DTLT efficiency.

This efficiency was found to be about 65–70% in all stations, independent of the magnetic field. The measurements obtained using 40 MHz bunched muon beams [147] provided a much higher efficiency, of the order of 85% or more, and for which only triggers at the correct BX were considered.

The observed lower efficiency of the DTLT can be explained by the fact that, while the trigger system is clocked every 25 ns, cosmic rays occur at any time. In bunched beams the BTIs can be properly synchronized by choosing the best phase which maximizes the number of HH triggers with respect to higher level triggers. In the case of cosmic rays, this is not possible, due to the random arrival time of the muons thus making the BTI synchronization itself meaningless. In such conditions the rate of Low-quality trigger segments released by the BTIs increases, and also the associated BX can easily fluctuate. As only correlated trigger segments were released by the trigger Sector Collector in each station, we also expect an increase in TRACOs (Sect. 9.2) failing to correlate segments among the two superlayers of a given station, with respect to a perfectly synchronized system, which will turn into a DTLT efficiency loss.

Consequently, one expects that muons crossing the detector at the “correct time” (for which the BTIs behave as perfectly synchronized to the clock) will be detected with the highest efficiency, as such a condition is the same as in the bunched beam tests. On the other hand, muons crossing the detector out of such a “correct time” will be detected with lower efficiency, as for them the synchronization of the system is not optimised.

The quantity t_0 is the time correction to be added to the t_{Trig} of the event to obtain the effective time at which the given muon crossed the detector. It can be computed event-by-event by minimizing the space resolution of the reconstructed track segment in a station. Figure 173 (left) shows the distribution of the t_0 correction of the reconstructed track segments in events triggered by the barrel RPC, superimposed to the one for events which also had a DT local trigger. The best trigger efficiency is obtained only at some preferred t_0 values. This can be seen in Fig. 173 (right) which shows the DTLT efficiency as a function of the t_0 -correction time, obtained as the ratio of the two distributions previously described. Two peaks at efficiency around 90% are visible. They correspond to the case in which the muon crosses the detector at the “correct time” for which BTIs are synchronized. The two peaks correspond to two adjacent BXs. For other t_0 values, the efficiency can be very low, as in this case the system is not synchronized. The observed DTLT efficiency is therefore explained by the fact that cosmic rays occur at random time with respect to the “correct time” at which the BTIs have maximum efficiency.

The magnetic field modifies the shape of the field lines in the drift cell, thus affecting the effective drift velocity, as discussed in Sect. 8.1.5. The largest effect in the barrel is expected to occur in station MB1 in Wheels +2 and -2. From the point of view of the DTLT, a change of the effective drift velocity, if large enough, could make the BX determination less precise, and shift its value by one unit.

Figure 174 shows this effect for the two MB1 stations in Wheel 2, displaying the BX value determined by the DTLT as a function of the z -position of the track in the chamber, with and without magnetic field. While no clear effect is visible in MB1-Wheel 1, in MB1-Wheel 2 there is a slight delay of the average BX value which tends to

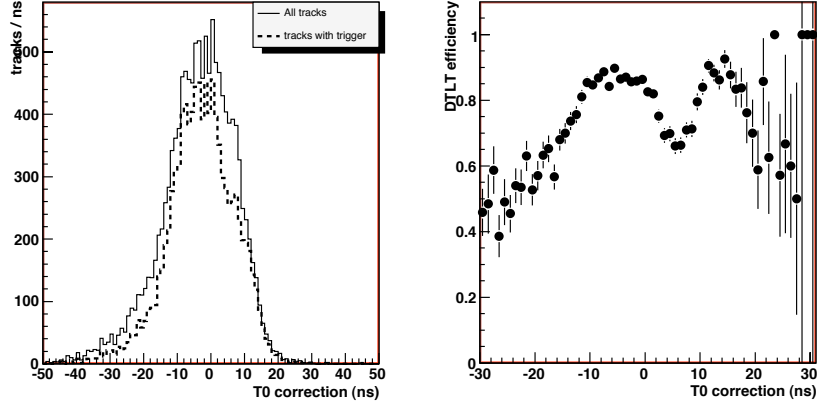


Figure 173: Left panel: Distribution of t_0 for all track segments in a station, for events triggered by RPC (solid line). The distribution of the same quantity for events also triggered by the DT is superimposed (dashed line). Right panel: DT local trigger efficiency as a function of the quantity t_0 , obtained by the ratio of the two histograms shown on the left.

increase as the track approaches the edge of the wheel, corresponding to larger values of z , where the stray field components are larger. This delay is at most of the order of 0.3 units of BX.

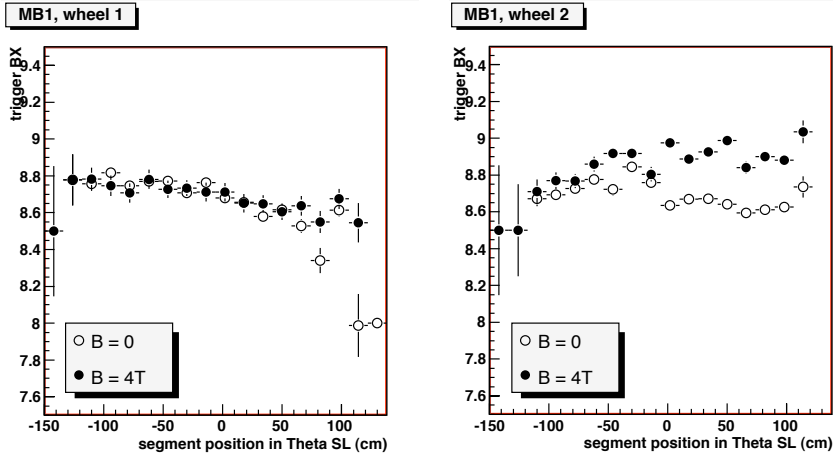


Figure 174: BX determined by the DTLT as a function of the track position in the z -direction of the muon station, with and without magnetic field, for MB1 in Wheel 1 on the left, where no effect of the B-field is expected. On the right, the same quantities are shown for MB1 in Wheel 2, where an influence of the magnetic field on the drift velocity is expected.

8.2 Cathode strip chambers

At the time of the LHC start-up, the CMS Endcap Muon system will consist of 468 cathode strip chambers (CSC) arranged in groups as follows: 72 ME1/1, 72 ME1/2, 72 ME1/3, 36 ME2/1, 72 ME2/2, 36 ME3/1, 72 ME3/2, and 36 ME4/1 (Figs. 175 and 176). The de-scoped 72 ME4/2 chambers will not be available during early years of CMS operation. The chambers are trapezoidal and cover either 10° or 20° in ϕ ; all chambers, except for the ME1/3 ring, overlap and provide contiguous ϕ -coverage. A muon in the pseudorapidity range $1.2 < |\eta| < 2.4$ crosses 3 or 4 CSCs. In the endcap-barrel overlap range, $0.9 < |\eta| < 1.2$, muons are detected by both the barrel drift tubes (DT) and endcap CSCs. In the baseline design, muons with $|\eta| < 2.1$ are also detected by resistive plate chambers (RPC); however, in the initial detector this coverage is reduced to $|\eta| < 1.6$.

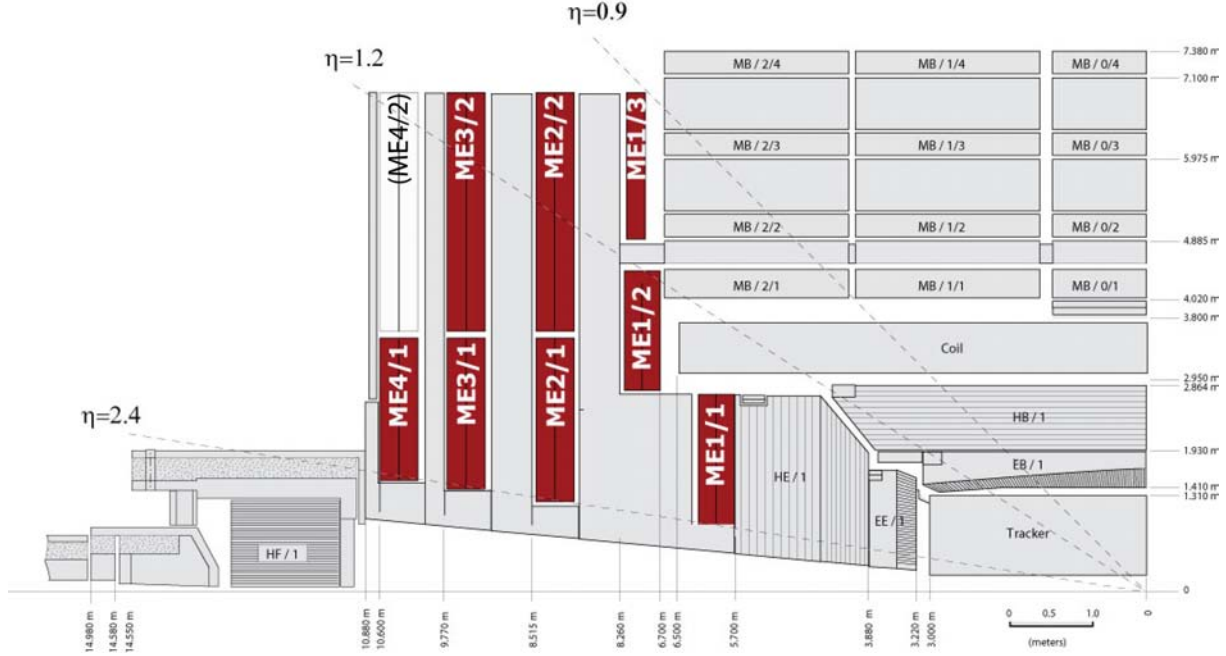


Figure 175: Quarter-view of the CMS detector. Cathode strip chambers of the Endcap Muon system are highlighted.

The CSCs are multiwire proportional chambers comprised of 6 anode wire planes interleaved among 7 cathode panels (Fig. 177). Wires run azimuthally and define a track's radial coordinate. Strips are milled on cathode panels and run lengthwise at constant $\Delta\phi$ width. Following the original CSC idea [148], the muon coordinate along the wires (ϕ in the CMS coordinate system) is obtained by interpolating charges induced on strips (Fig. 178). The largest chambers, ME2/2 and ME3/2, are about $3.4 \times 1.5 \text{ m}^2$ in size. The overall area covered by the sensitive planes of all chambers is about 5000 m^2 , the gas volume is $>50 \text{ m}^3$, and the number of wires is about 2 million. There are about 9000 high-voltage channels in the system, about 220 000 cathode strip read-out channels with 12-bit signal digitisation, and about 180 000 anode wire read-out channels.

The CSCs provide the functions of precision muon measurement and muon trigger in one device. They can operate at high rates and in large and non-uniform magnetic fields. They do not require precise gas, temperature, or pressure control. Moreover, a radial fan-shaped strip pattern, natural for measurements in the endcap region, can be easily arranged on the cathode planes.

The performance requirements for the CMS cathode strip chamber system include the following:

- Reliable and low-maintenance operation for at least 10 years at the full LHC luminosity, i.e., at estimated random hit rates up to 1 kHz/cm^2 ;
- At least 99% efficiency per chamber for finding track stubs by the first-level trigger;
- At least 92% probability per chamber of identifying correct bunch crossings by the first-level trigger. With such an efficiency per chamber and 3–4 CSCs on a muon track path, a simple majority rule ensures that the reconstructed muons will be assigned the correct bunch crossing number in more than 99% of cases;
- About 2 mm resolution in r - ϕ at the first-level trigger.

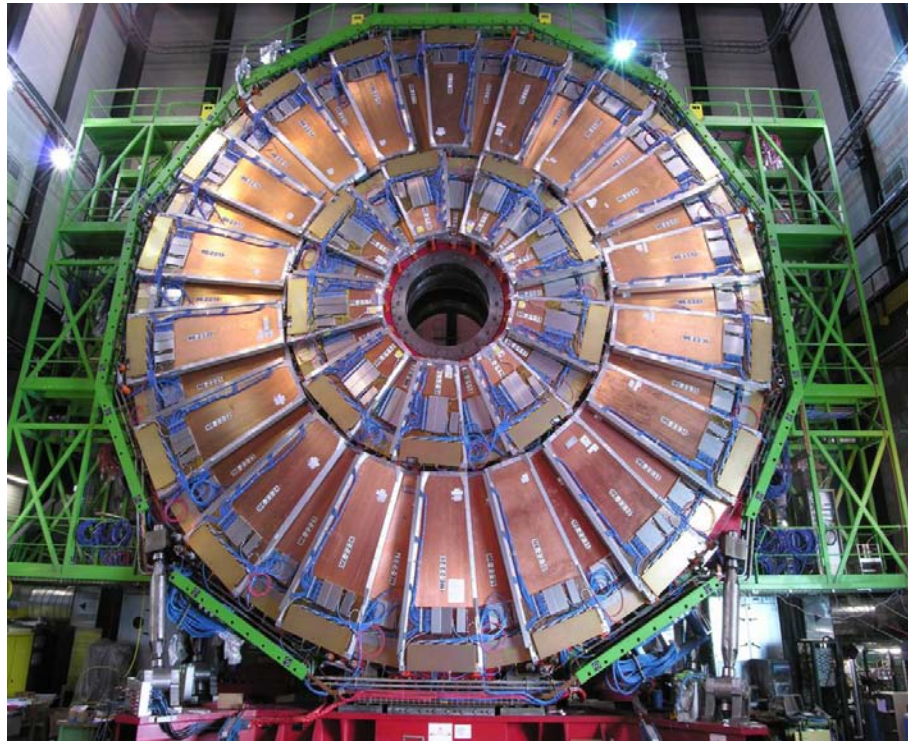


Figure 176: The ME2 station of CSCs. The outer ring consists of 36 ME2/2 chambers, each spanning 10° in ϕ , and the inner ring of eighteen 20° ME2/1 chambers. The chambers overlap to provide contiguous coverage in ϕ .

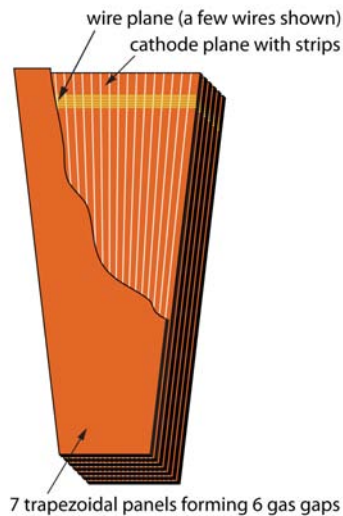


Figure 177: Layout of a CSC made of 7 trapezoidal panels. The panels form 6 gas gaps with planes of sensitive anode wires. The cut-out in the top panel reveals anode wires and cathode strips. Only a few wires are shown to indicate their azimuthal direction. Strips of constant $\Delta\phi$ run lengthwise (radially). The 144 largest CSCs are 3.4 m long along the strip direction and up to 1.5 m wide along the wire direction.

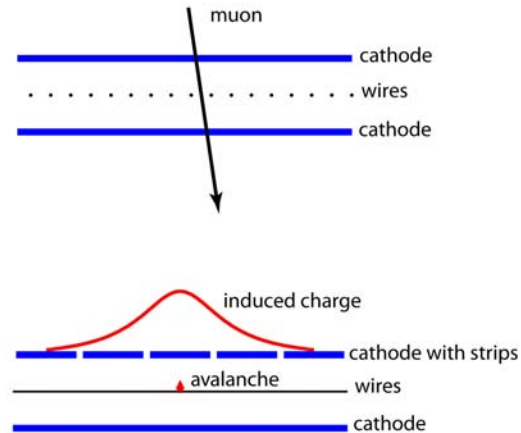


Figure 178: A schematic view of a single gap illustrating the principle of CSC operation. By interpolating charges induced on cathode strips by avalanche positive ions near a wire, one can obtain a precise localisation of an avalanche along the wire direction.

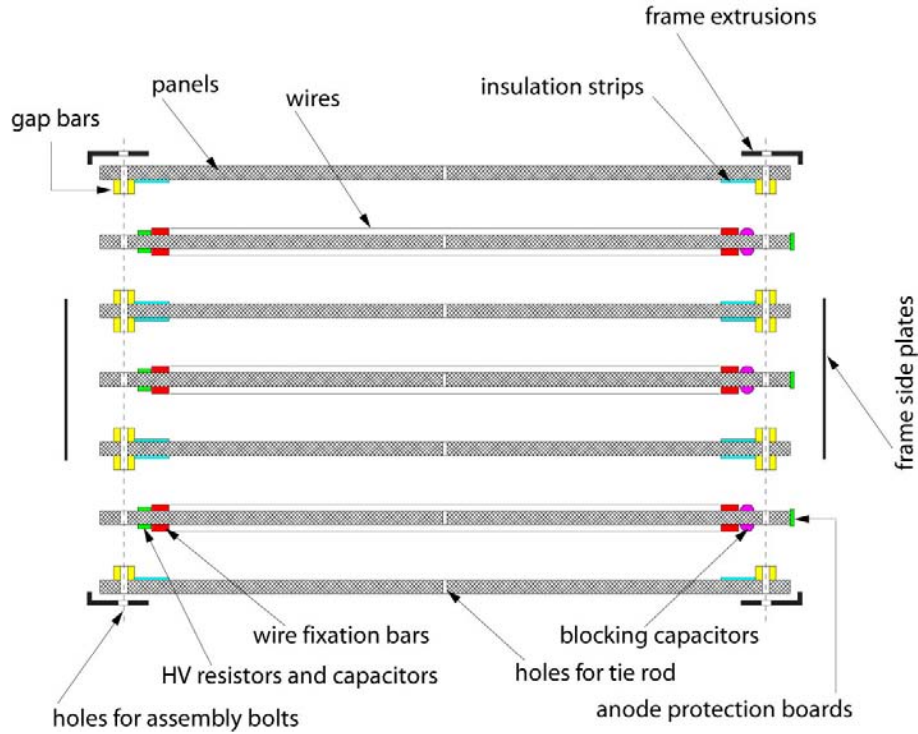


Figure 179: Mechanical design of the CMS cathode strip chambers (exploded view).

- About $75\ \mu\text{m}$ off-line spatial resolution in $r\text{-}\phi$ for ME1/1 and ME1/2 chambers and about $150\ \mu\text{m}$ for all others.

8.2.1 Chamber mechanical design

The 72 ME1/1 chambers and the larger 396 chambers have somewhat different mechanical designs. Below, we describe the design of the larger chambers using ME2/2 as an example and, then, at the end of this Section, summarise the ME1/1-specific features that distinguish them from the other chambers.

The mechanical structure is based on seven 16.2-mm-thick trapezoidal panels (Fig. 179). The panels are made of a 12.7-mm-thick polycarbonate honeycomb core with two 1.6-mm FR4 skins commercially glued on each side. FR4 is a fire-retardant fibreglass/epoxy material widely used for printed circuit boards. The FR4 skins are clad with $36\text{-}\mu\text{m}$ -thick copper on their outer surfaces, forming the cathode planes.

FR4 cathode gap bars are glued to both sides of every other panel (panels 1, 3, 5, 7 in Fig. 179) so that when the panels are stacked together, these cathode bars define 6 gas gaps of 9.5 mm. To provide additional support, there are 4 spacers placed between panels along the chamber centreline. When all 7 panels are put together, the entire stack is tightened down with bolts along the chamber perimeter (through holes in the cathode gap bars) and at 4 points along the chamber centreline (through holes in the spacers). Such an arrangement ensures that no panel has more than 60 cm of unsupported length. Measurements show that most of the panels are flat within the required $\pm 300\ \mu\text{m}$ on such spans. This specification arises from the desire to keep gas-gain variations within a factor of 2.

Six of the panels have a pattern of 80 strips milled on one side. Strips, being radial, have a pitch that varies from 8.4 mm at the narrow chamber end to 16 mm at the wide end. The gap between strips is about 0.5 mm. The precision of milling was better than $50\ \mu\text{m}$ (rms). Milling was done with a cutter tilted at 45° to make the groove edges smoother (otherwise, sharp edges and burrs might provoke sparking and discharges).

Three of the panels are so-called *anode* panels (panels 2, 4, 6 in Fig. 179) around which anode wires were wound (these panels do not have gap bars). A specially designed winding machine wound wires directly on a panel by rotating it around its long axis at a speed of about 5 turns per minute; one panel could be completed (about 1000 wires on each side) in less than 4 hours. The wire spacing of about 3.2 mm was defined by combs: threaded rods running the full panel length and attached to the panel edges during winding. Gold-plated tungsten wires, $50\ \mu\text{m}$ in diameter, were stretched at 250-g tension (about 70% of the elastic limit) and run their full length up to 1.2 m without any intermediate supports. The electrostatic stability limit for the longest wires is above 6 kV (the

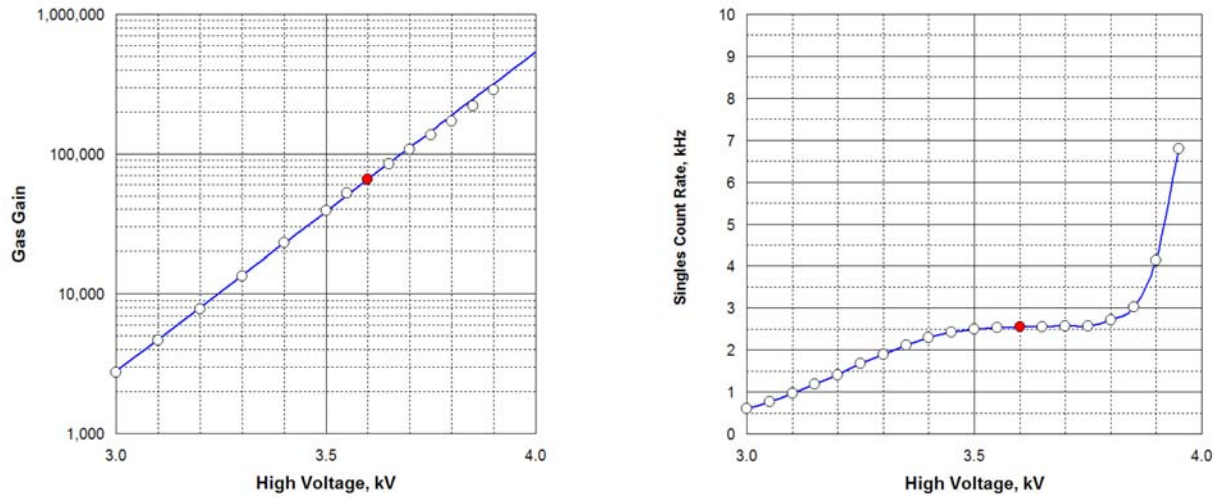


Figure 180: Left panel: CSC gas gain vs. high voltage. Right panel: ME2/1 chamber singles rate vs. high voltage (the overall sensitive area of all 6 planes in this chamber is $\approx 9.5 \text{ m}^2$).

nominal operational point is 3.6 kV). Based on measurements during production, the wire tension non-uniformity does not exceed $\pm 10\%$, while wire spacing variations are less than $\pm 150 \mu\text{m}$. Wires found to fall outside of these specifications were replaced.

After winding, the wires were first glued and then soldered to anode bars 4.75 mm in height (half of the gas gap). The anode bars are made of copper-clad FR4 and carry the electric artwork. An automated soldering machine allowed for soldering at a speed of 3.5 s per joint. Groups of 16 wires make 1 anode read-out channel with a width of about 5 cm. High voltage (HV) is distributed to the wire groups on one end and signals are read out on the other end via 1 nF blocking capacitors.

Each wire plane is sub-divided by spacer bars into 5 independent HV segments, which allows us to independently regulate or turn off HV on any of the 5 sections. In places where the spacer bars were inserted (and prior to their installation), 8 wires were removed. Two gold-plated 200- μm guard wires were inserted in place of the first and eighth thin wires that were removed to eliminate edge effects. The very first and last wires in each plane are also thicker. If the edge thin wires were to be left unguarded, the electric field on them would be much larger than for the rest of the wires, which would provoke discharges. Such plane segmentation, because of the intermediate panel supports and the individual HV control over smaller wire-plane sections, makes the overall chamber performance very robust.

After stacking the panels and tightening the bolts (with O-rings), continuous beads of RTV sealant were applied along the outer seams between the panels and gap bars. The O-rings around the bolts and the RTV seal make the chambers hermetic. Should it be necessary, a chamber can be opened, serviced, and resealed. Gas enters into one of the outer gas gaps via an inlet in a cathode gap bar, flows from one plane to another in a zigzag manner via special holes in the panels, then exits from the last gas gap via an outlet in a gap bar. The leak rate, measured during production and after installation of the chambers, was required to be $< 1\%$ of the chamber volume per day at an over-pressure of 7.5 mbar (e.g., $< 2 \text{ cm}^3/\text{min}$ for the largest chambers whose gas volume is about 200 litres).

Side plates made of 3.2-mm-thick Al extrusions were attached around the chamber perimeter. They stiffen the chamber and connect the top and bottom copper skins to form a Faraday cage.

The nominal gas mixture is 40%Ar + 50%CO₂ + 10%CF₄. The CO₂ component is a non-flammable quencher needed to achieve large gas gains, while the main function of the CF₄ is to prevent polymerisation on wires. A detailed discussion of the gas optimisation can be found elsewhere [149].

Figure 180 (left) shows the chamber gas gain vs. high voltage. The nominal operating HV was chosen to be 3.6 kV, which corresponds to a gas gain on the order of 7×10^4 . A minimum ionising particle (MIP) produces about 100 electrons in a gas gap, thus the charge per MIP in an avalanche is about 1 pC. As is shown below, at this operational point, the cathode and anode electronics have a very high efficiency and an adequate signal-to-noise ratio. The operational range of the chambers extends to 3.9 kV. Typically, we start seeing a sharp rise in the rate of spurious pulses at about 3.9–4.0 kV (Fig. 180 (right)).

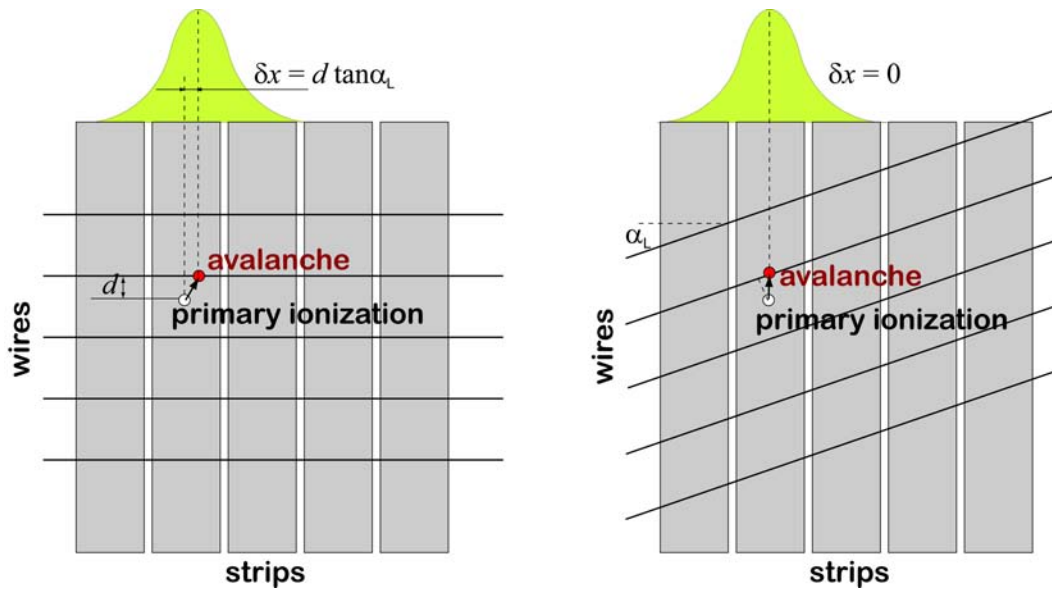


Figure 181: Left panel: if the ME1/1 wires were not tilted, ionisation electrons, as they drift toward the anode wires in the strong magnetic field normal to the plane of the drawing, would be carried sideways by the Lorentz force. The direction and size of the shift would depend on whether the electrons drift upwards or downwards and on how far away they were from the wires to begin with. These sideways displacements would spread the charge over the anode wires. Right panel: by tilting the wires at the Lorentz angle α_L , all ionisation electrons arrive near the same point.

The 72 ME1/1 chambers have differences in their mechanical design with respect to the other CSCs. The gas gap is 7 mm, wire diameter is $30\ \mu\text{m}$, and wire spacing is 2.5 mm, so the nominal HV for these chambers is somewhat lower: 2.9 kV. Most notably, the ME1/1 anode wires are not azimuthal, but are tilted by an angle $\alpha_L = 29^\circ$ (Fig. 181). Unlike the other CSCs, the ME1/1 chambers are inside the CMS solenoid and see its strong and uniform 4 T axial field. The wire tilt compensates for the Lorentz angle so that electrons drift parallel to the strips, enabling a precise measurement of the r - ϕ -coordinate.

8.2.2 Electronics design

Figure 182 shows a schematic layout of the custom-made trigger and read-out electronic boards developed for the CSC system.

An anode front-end board (AFEB) has one 16-channel amplifier-discriminator application-specific integrated circuit (ASIC). The amplifier has a 30-ns shaper (semi-Gaussian with 2-exponent tail cancellation designed to suppress the slow signal component associated with a drift of positive ions away from the anode wires), about 7 mV/fC sensitivity, and 1.4 fC noise at a typical wire group capacitance of 180 pF for the largest chambers. With the 30-ns shaping time, an AFEB sees about 12% of the total avalanche charge, i.e., an average of about 130 fC. A typical chamber signal as seen at the output of this amplifier is shown in Fig. 183 (left). The constant-fraction discriminator has a threshold nominally set at 20 fC (input equivalent charge) and its slewing time is less than 3 ns for the 60–600 fC signal range. Depending on chamber size, there are 12 to 42 AFEBs per chamber. Further details on the AFEB design and performance can be found elsewhere [150].

Every 25 ns, in synchronization with the LHC collisions, all AFEB outputs, 40-ns-long step pulses, are sampled by an FPGA-based anode local charged track (ALCT) board, 1 board per chamber. The recorded *yes/no* information is stored in a FIFO. Upon receiving a CMS-wide Level-1 Accept (L1A) trigger command, the recorded information within the proper time window is extracted and reported to the DAQ. The latency of the L1A command with respect to the time of a collision is $3.2\ \mu\text{s}$. The temporal length of the raw-hit record transmitted to the DAQ can be as large as 800 ns.

The ALCT board has another important function. Based on the information from all anode channels, the FPGA code constantly (every 25 ns) searches for patterns of hits among the 6 planes that would be consistent with muon tracks originating from the interaction point. For a pattern to be valid, we require that hits from at least 4 planes be present in the pattern. Figure 183 (right) illustrates how patterns are identified in the presence of spurious single-plane hits. Due to a large neutron-induced photon background, a substantial rate of such single-plane hits

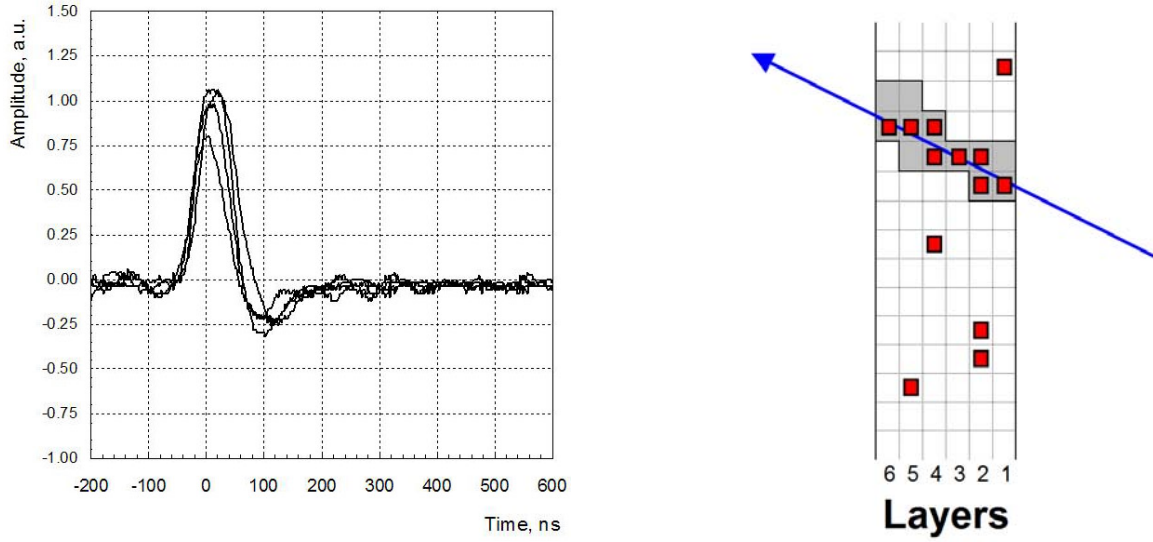


Figure 183: Left panel: muon signals as seen at the AFEB amplifier output. Right panel: a schematic event display showing anode signals in the 6 planes of a CSC (small dark squares). The ALCT board FPGA logic is programmed to scan the chamber and search for hits falling inside predefined patterns (grey cells) consistent with muons originating from the interaction point. Hits must be present in at least 4 planes for an ALCT pattern to be found.

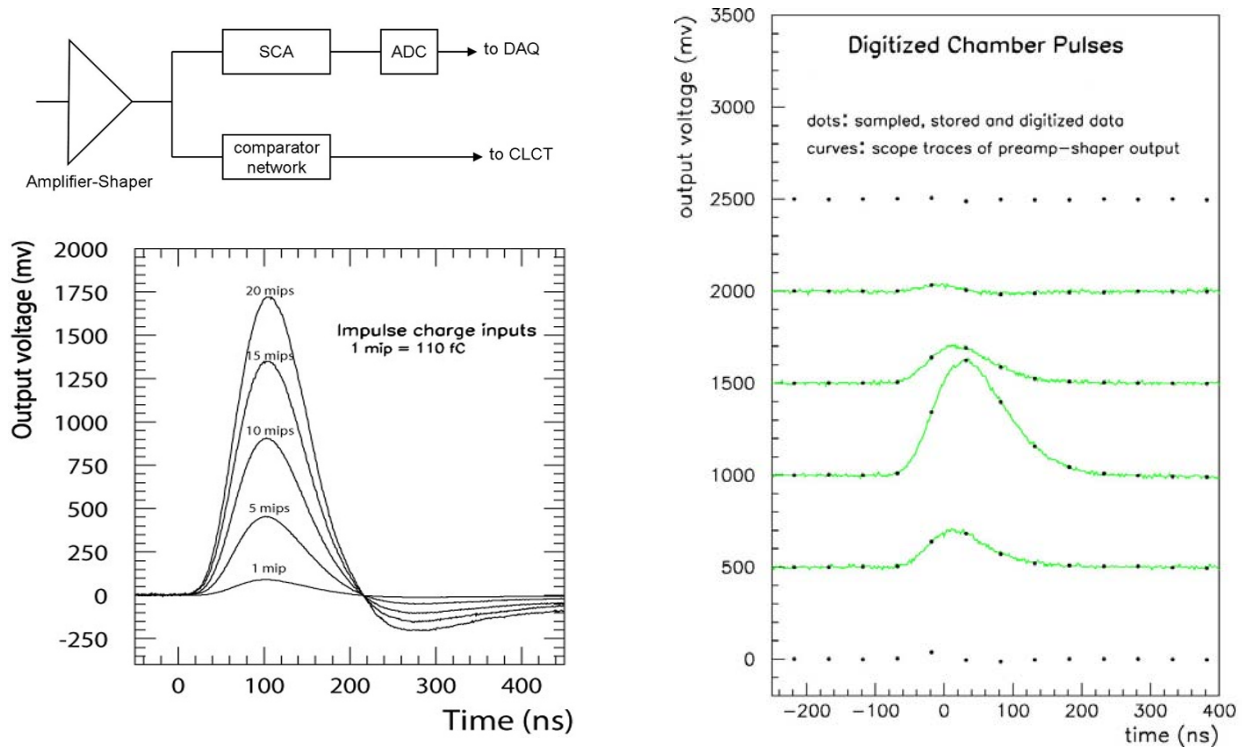


Figure 184: Left panel: basic functional diagram of a CFEB and the CFEB amplifier-shaper response to a δ -function input pulse. The undershoot is intended to compensate for the long tail present in muon hit signals resulting from the slow drift of positive ions away from the anode wires. Right panel: signals from a muon track on 6 contiguous strips in a layer. The 4 curves are oscilloscope traces and the 6 lines of dots are digitised outputs. The signals have an arbitrary vertical offset for ease of viewing.

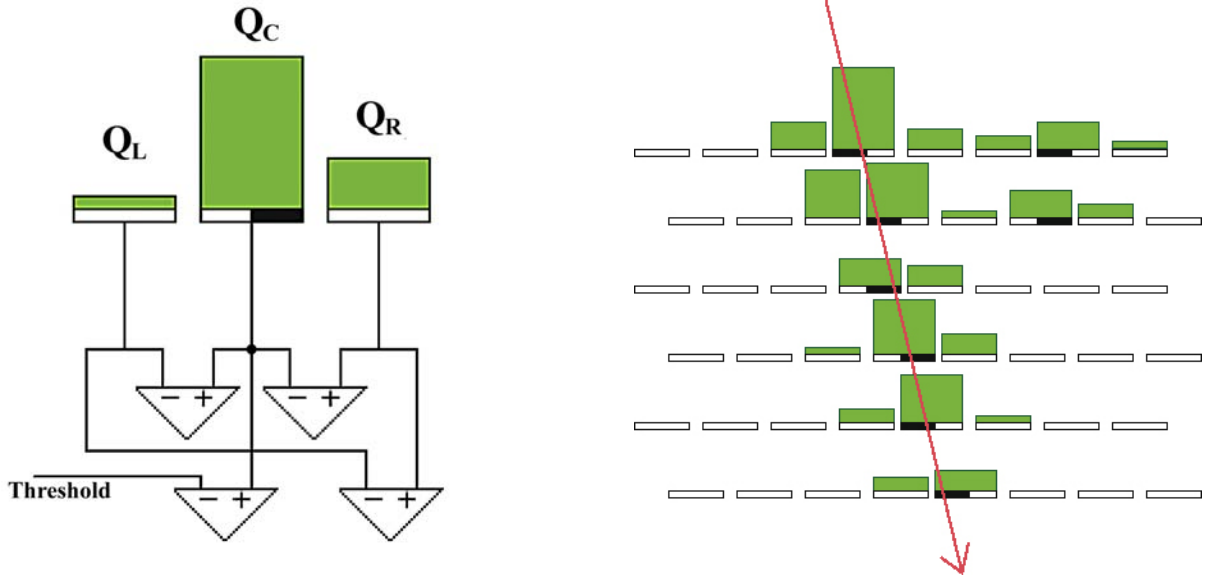


Figure 185: Left panel: a simplified schematic of the idea behind the comparator network. For each group of 3 adjacent strips, comparators compare the central strip signal Q_c with a threshold and with the central-to-left $Q_c - Q_l$, central-to-right $Q_c - Q_r$, and right-to-left $Q_r - Q_l$ strip signal differences. If $Q_c > \text{threshold}$, $Q_c - Q_l > 0$, $Q_c - Q_r > 0$, and $Q_r - Q_l > 0$ (as shown here), the hit position is somewhere within the right half of the central strip. Right panel: a *cathode local charged track* is a pattern of half-strip hits consistent with a muon track.

search, for a CLCT pattern to be found, half-strip hits must be present from at least 4 planes. Unlike the ALCT boards, the TMBs are not affixed to the chambers, but are in peripheral VME crates mounted along the outer rim of the endcap steel disks.

The TMB also matches ALCT and CLCT patterns found within a chamber to make correlated 2-dimensional LCTs ($2D-LCT = ALCT \times CLCT$), up to 2 per bunch crossing. These 2D-LCTs are sent to muon port cards (MPC), each of which serves 9 chambers covering either 60° - or 30° -sectors in ϕ . For each bunch crossing, an MPC performs a preliminary sorting of all received correlated 2D-LCTs and finds the 3 best-quality candidates—these are then sent further upstream to the muon L1-trigger electronics.

Raw data are collected by the DAQ motherboards (DMB) located in the peripheral crates. There is one DMB for each chamber. The data consist of anode and comparator hits within a time window up to 32 bunch crossings long, ALCT and CLCT decisions in the same window, and digitised strip signal waveforms (eight or sixteen 50-ns time samples). The status of the various electronic boards is also a part of the event record. The data collected by the DMB are passed to a detector-dependent unit (DDU) board, then to a data concentration card (DCC), and finally to the CMS filter farm to be processed by the CMS high-level trigger (HLT) software. The expected event size per chamber is about 5 kBytes.

It is important to note that the CSC read-out is intrinsically zero-suppressed. The anode raw data in a particular chamber are passed downstream only if there is an ALCT pattern in coincidence with the L1A signal. Likewise, the cathode information, comparator hits and digitised strip signal waveforms, are passed downstream to the DAQ only if there was a similar $CLCT \times L1A$ coincidence. The coincidence window is programmable, but is nominally set at 75 ns, i.e., ± 1 bunch crossing.

At the design LHC luminosity, we expect on average to find track stubs in 2 chambers for each L1A signal. With the maximum CMS L1A rate of 100 kHz, the data flow rate from CSCs to HLT is estimated to be around 1 GB/s.

Operation of the peripheral VME crates is supported by clock-control boards (CCB) and custom crate controllers. As its name implies, the CCB distributes the LHC clock and all CMS control commands (like L1A signals).

The HV system is custom made and provides channel-by-channel regulated voltage up to 4.0 kV with about 10 V precision. Currents of less than $10 \mu A$ can be measured with a precision of 100 nA, while the precision for larger currents is about 1%. The system can provide more than $100 \mu A$ current for individual channels as long as the average consumption does not exceed $40 \mu A$ per channel. The maximum expected current at the design LHC luminosity for the most-loaded HV segment is $< 10 \mu A$.

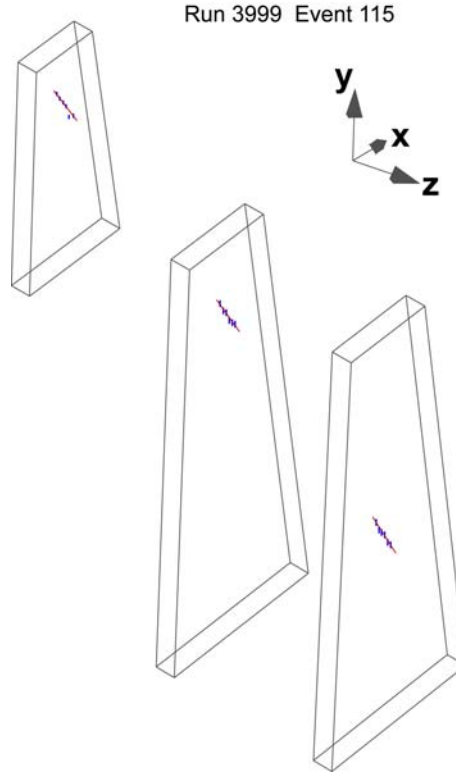


Figure 186: Part of a CMS event display showing a muon event detected by CSCs during the MTCC. Only those chambers containing muon hits are displayed.

8.2.3 Performance

The results presented in this Section come from tests conducted with final-design CSCs in high-energy muon beams at CERN, with cosmic-ray muons in a lab or in situ after installation, and at the Gamma Irradiation Facility (GIF).

A high-energy muon beam provides a test environment with maximum control, but it can expose only a small portion of a chamber. We typically used 100–300 GeV beams, which also allowed us to study chamber performance in the presence of bremsstrahlung radiation. To study performance over the entire chamber area, for many years we have tested individual large chambers with cosmic-ray muons at various research laboratories.

During the CMS Magnet Test and Cosmic Challenge (MTCC) in 2006, a substantial part of CMS was operated as a unified system. The CSC subsystem was represented by a 60° sector, or 36 chambers. This allowed us to obtain in situ performance results for a large number of CSCs operating simultaneously with other CMS subsystems. Figure 186 shows a muon event as detected by CSCs at the MTCC.

At the design LHC luminosity, we expect a large neutron flux in the underground cavern, which upon thermalization and capture is predicted to result in a substantial flux of ≈ 1 MeV photons. Of the photons that enter a chamber, about 1% will convert to electrons. These electrons will give rise to large rates of random hits up to 1 kHz/cm^2 in the CSCs. The GIF at CERN has a Cs-137 source of 0.7 MeV photons with an intensity of $\approx 0.7 \times 10^{12} \text{ Bq}$. Tests at the GIF allowed us to study the chamber performance in an environment of high random-hit rates. We also used these facilities for chamber-ageing studies.

Trigger primitives It is important to note that the efficiency of finding trigger primitives (LCTs) directly affects not only the muon trigger, but also the DAQ path. As was described earlier, the CSC read-out is driven by an LCT \times L1A coincidence. If an LCT is not found, there will not be a coincidence, and no raw hits will be recorded and available for the offline reconstruction.

The anode signal efficiency of a single plane is shown in Fig. 187 (left). The same figure also shows the efficiency for finding ALCTs, patterns of hits in 6 planes consistent with a muon track. The desired ALCT-finding efficiency of 99% is reached above 3.4 kV. At 3.6 kV, the ALCT-finding efficiency is about 99.9%. These results were obtained for test-beam muons going through a small area of a chamber free of dead zones. For CLCT patterns,

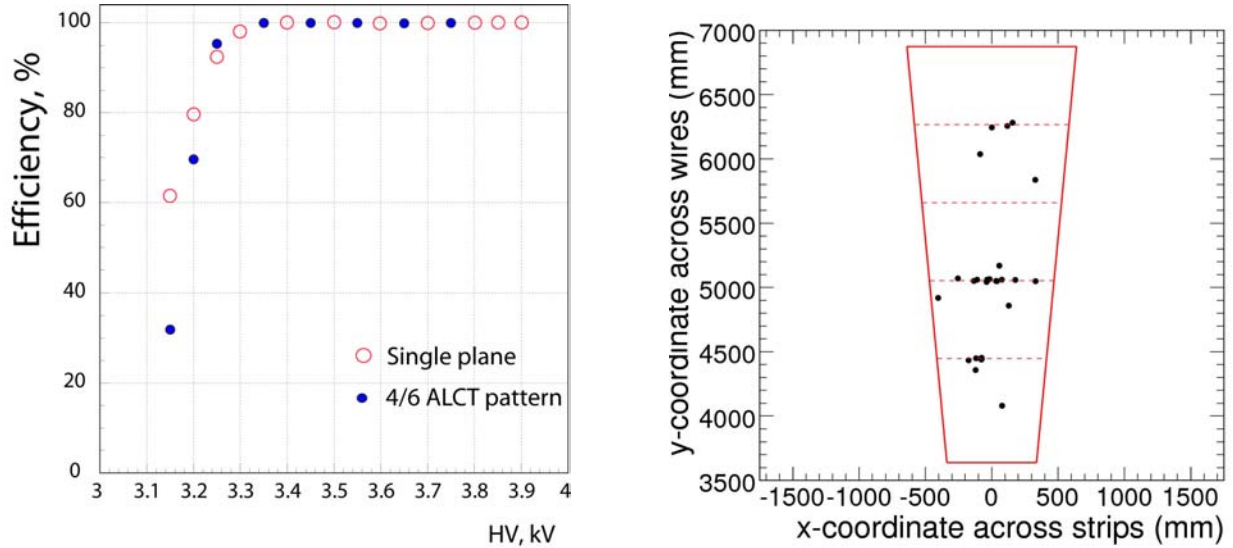


Figure 187: Left panel: single-plane anode signal efficiency (open circles) and ALCT pattern finding efficiency (filled circles) vs. high voltage. Right panel: predicted position of muon tracks in ME2/2 chambers for events when no LCT was found in these chambers (superimposed results for 6 chambers). The dashed lines indicate where wire planes of the ME2/2 chambers have inefficient bands separating independent HV segments.

similar results are achieved at about 50 V higher. This is because the cathode signal is somewhat smaller than the anode signal.

The overall efficiency of finding 2D-LCT patterns (ALCT \times CLCT) averaged over the entire area of many chambers was studied with cosmic-ray muons at the MTCC. At the nominal HV of 3.6 kV, the average 2D-LCT efficiency in 6 ME2/2 chambers was found to be $(99.6 \pm 0.1)\%$. For the 0.4% of events with missing 2D-LCTs in ME2/2 chambers, the majority of tracks (reconstructed using the ME1 and ME3 LCT stubs) were found to cross ME2/2 chambers in inefficient bands separating the chamber high-voltage segments (Fig. 187 (right)).

To test whether the found LCTs are indeed associated with the muons going through the chambers, we looked at the relative distance between the (x, y) coordinates of 2D-LCTs found in ME2/2 chambers and the muon track (x, y) coordinates in the ME2 station as predicted from the 2D-LCTs in the ME1 and ME3 stations. (Here x and y are the local chamber coordinates across the cathode strips and anode wires, respectively.) The 2D-residuals between the measured and predicted (x, y) coordinates are shown in Fig. 188 (left). The observed spread of ≈ 1.8 cm along the x axis is consistent with the expected multiple scattering of cosmic-ray muons penetrating the endcap steel disks. As is shown below, the intrinsic precision of CLCT localisation is much better than that. The distribution along the y axis is noticeably broader due to a much coarser wire group segmentation of 5 cm, which defines the precision of ALCT localisation in these chambers.

For studying the intrinsic CLCT-localisation precision, we used a test chamber in a muon beam and a telescope of Si micro-strip detectors to precisely determine the position of a muon going through the test chamber. To achieve the best results, a given CLCT pattern is assigned an x coordinate corresponding to the average of all muons that can generate such a pattern. Figure 188 (right) shows the residuals between the Si-based track coordinate and the CLCT-based coordinate. The distribution is Gaussian and has $\sigma \approx 0.11$ in strip width units, which is better than the desired 2 mm for even the widest 16-mm strips. In the more conservative approach currently implemented in the muon trigger firmware, CLCT patterns are localised within a half-strip. Therefore, in this approach the CLCT spatial resolution is approximately $(w/2)/\sqrt{12} \approx 0.14w$, where w is the strip width.

The time distribution of anode signals from a single chamber plane (Fig. 189 (left, top)) has an RMS of about 11 ns. Clearly this is too wide for a chamber hit to be assigned unambiguously to the correct bunch crossing. We overcome this problem by making use of all 6 planes in a chamber. The time distribution for the 3rd earliest hit in an ALCT pattern (Fig. 189 (left, bottom)) is a much narrower Gaussian with $\sigma < 5$ ns, the use of which results in a bunch-tagging efficiency of 98–99%, well above the desired 92% level. Figure 189 (right) shows the accuracy required for aligning the phase of the 25-ns clock on an ALCT board with the LHC clock. The acceptable range of phase misalignment is ± 5 ns. CLCTs tend to have slightly worse timing properties due to the slower CFEB shaping time and smaller amplitude of strip signals, so we assign the time tagged by the ALCT pattern to the matched 2D-LCT.

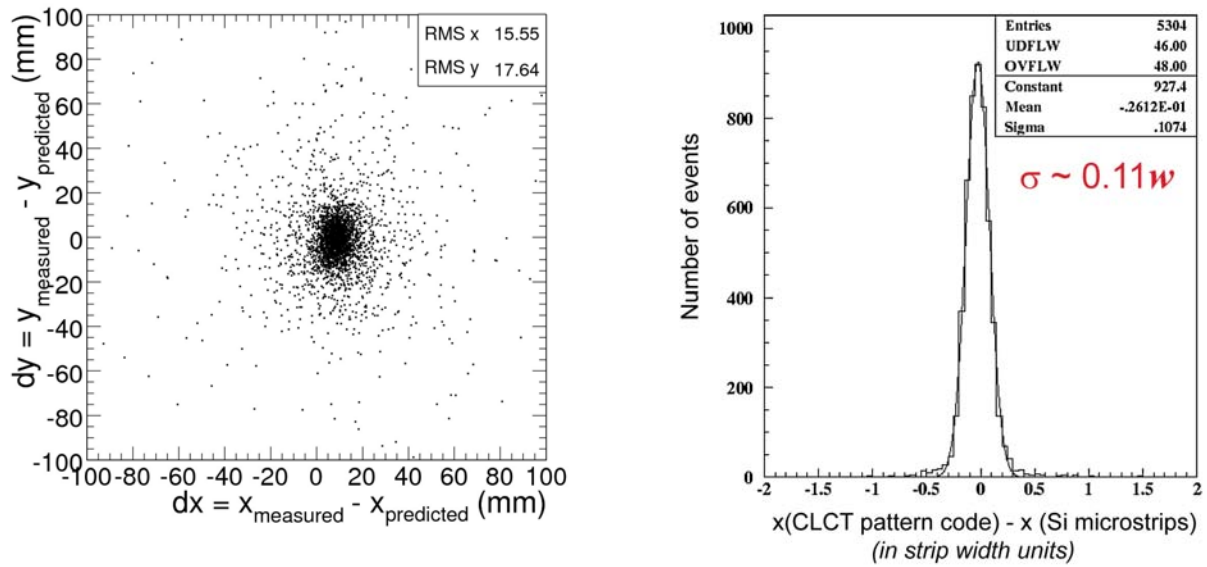


Figure 188: Left panel: 2D-coordinate residuals between LCTs found in ME2/2 chambers and muon-track positions. The x axis runs along the wires and the y axis along the strips. The observed spread is consistent with multiple scattering of cosmic-ray muons in the steel disks and the expected CLCT and ALCT spatial resolutions. Right panel: Residuals between the CLCT pattern-defined muon coordinate and the coordinate predicted by the Si beam telescope in a 300-GeV muon beam. The residuals are shown in units of strip width.

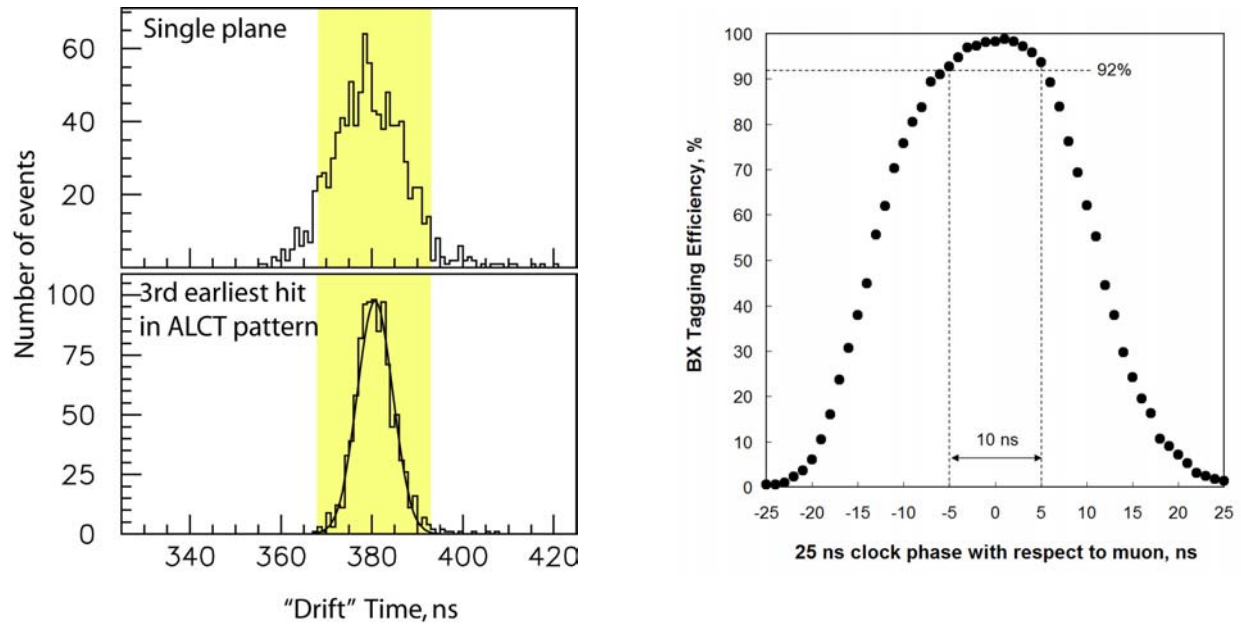


Figure 189: Left panel: Time distributions of the response of a single plane to a passing muon (top) and for the 3rd earliest hit in an ALCT pattern (bottom). The horizontal scale has an arbitrary offset. The shaded band indicates the 25-ns window, the time between bunch crossings at the LHC. Right panel: Probability for correct bunch-crossing tagging vs. relative phase shift between the 25-ns clock on an ALCT board and the LHC 25-ns clock.

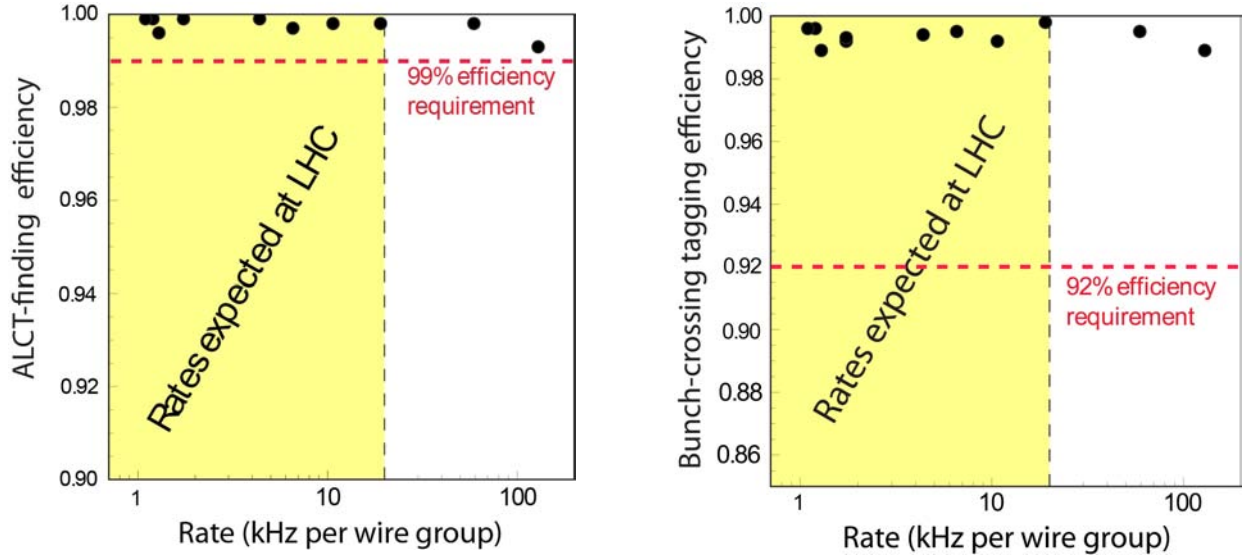


Figure 190: Left panel: ALCT-finding efficiency vs. rate of random hits per wire group. Right panel: efficiency of correct bunch tagging vs. rate of random hits per wire group. The shaded areas show the range of rates expected in different chambers at full LHC luminosity.

Results obtained from a CSC irradiated with 0.7 MeV photons in a muon beam at the GIF (Fig. 190) show that the ALCT-finding and bunch-tagging efficiencies remain very robust even at random-hit rates far exceeding those expected at full LHC luminosity.

During 300-GeV muon-beam tests, a 30-cm-thick steel slab was moved in front of the test chamber to study the effect of bremsstrahlung radiation on the reconstruction of muon stubs at the trigger level. In offline analysis of strip data, we classified each muon as either “clean” (multiple charge clusters observed in only 1 plane) or otherwise “contaminated.” Without the steel slab, the fraction of “clean” muons was 94% and CLCT patterns were formed from half-strips with a 99.5% efficiency. With the steel slab, the fraction of “clean” muons dropped to 80%, while the CLCT-finding efficiency remained very high (98.9%). Figure 191 shows an example of a badly “contaminated” muon where the muon track is nevertheless successfully identified.

Spatial resolution based on digitised strip signals An avalanche on a wire induces charge on a cathode plane. In a first approximation, the shape of the induced charge can be parameterized by the so-called Gatti function [154]:

$$\frac{1}{Q} \frac{dQ}{d\lambda} = K_1 \left[\frac{1 - \tanh^2(K_2\lambda)}{1 + K_3 \tanh^2(K_2\lambda)} \right], \quad (9)$$

where $\lambda = x/h$, in which x is the coordinate across a strip and h is the cathode-anode spacing, and the coefficients K_1 , K_2 , and K_3 are defined by the chamber geometry.

Given the geometry of the CSCs, most of the induced charge is shared among 3–4 strips. As described earlier, a strip signal waveform is sampled and digitised every 50 ns. The signal peaks in about 150 ns and comes back very close to the baseline within the next 150 ns so that the overall pulse duration is roughly 300 ns. Such a 2-dimensional charge cluster can be fit to obtain the spatial coordinate, time, and cluster charge. To achieve the best possible resolution, we take into account empirical corrections for the induced charge shape, the time structure of the signal waveform, strip-to-strip cross-talk, electronic-noise correlations between nearby time samples, and electronic pedestal and gain calibrations.

By design, ME1/1 and ME1/2 chambers have narrower strips and thus deliver better resolution. The ME1/1 single-plane resolution (Fig. 192 (left)) is about 80 μm at nominal HV. The 6-plane chamber resolution is estimated to be $\approx 80/\sqrt{6} = 33 \mu\text{m}$, plus alignment errors. Clearly, the desired resolution of 75 μm per 6-plane chamber is within reach.

The single-plane spatial resolution of the larger CSCs (with very wide strips up to 16 mm) depends very strongly on the muon coordinate across a strip. Muons that pass through a strip centre will be measured poorly (and the wider the strip, the worse the measurement). On the other hand, muons hitting between strips will be measured

RUN 27488

EVENT 118

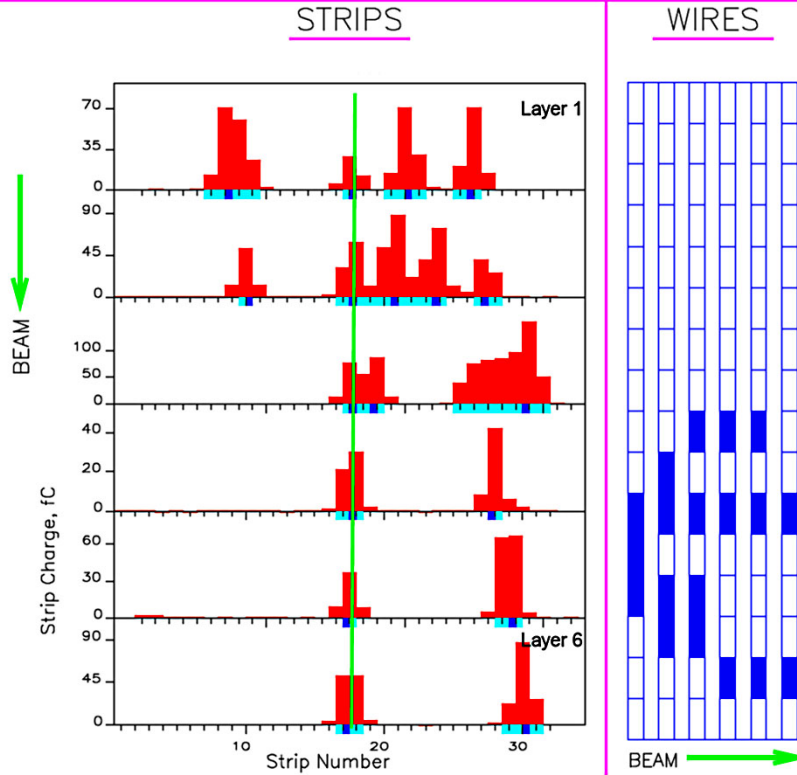


Figure 191: A sample CSC event of a muon accompanied by substantial bremsstrahlung radiation. The left side of the plot shows the charge (blocks just above the axes) on each of the 32 strips in each of the 6 chamber layers, while the right side shows the information from the anode wire groups. Strips with charge above the trigger threshold are marked with light shading below the axes, while the half-strip “peaks” are marked with dark squares. The vertical line at $\lambda = 17.56$ shows the track position extrapolated from the Si beam telescope. This event was assigned a 6-layer CLCT code that corresponds to an average track position $\lambda = 17.59$.

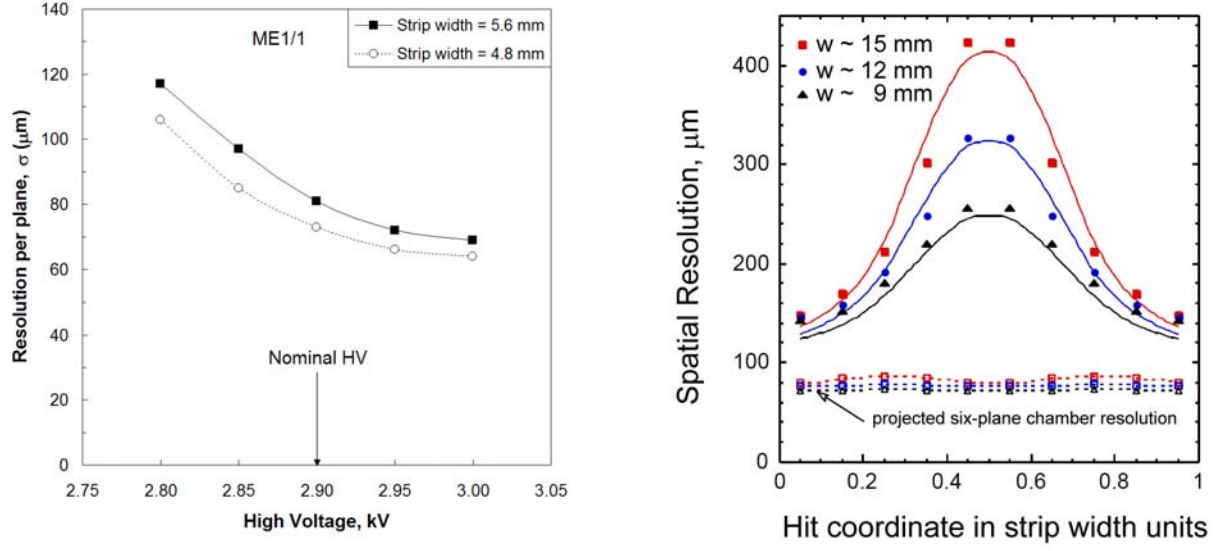


Figure 192: Left panel: ME1/1 chamber single-plane resolution vs. HV. Right panel: Ultimate large CSC offline resolutions per plane for different muon passage points across a strip for areas with different strip widths for data (closed symbols) and simulation (solid lines). The expected overall 6-plane CSC resolutions are shown by open symbols and dashed lines.

nearly equally well for any strip width. We took advantage of this feature in our design. In the larger chambers, strips in adjacent planes are staggered by one half of the strip pitch. High-energy muons, for which we need the best chamber resolution, appear as nearly straight-through tracks. If such a muon goes through areas with poor resolution in odd planes, it will have very good measurements in even planes, and vice versa. Therefore, by combining measurements from 6 planes with proper weighting, a muon track segment is accurately localised. Figure 192 (right) shows single-plane resolutions, σ_i , for chamber regions with different strip widths and the resulting combined 6-plane resolution, σ_{CSC} , which is estimated by

$$\frac{1}{\sigma_{\text{CSC}}^2(x/w)} = \sum \frac{1}{\sigma_i^2(x/w)}. \quad (10)$$

The expected combined resolution for a 6-plane chamber is $\approx 80 \mu\text{m}$ almost independent of the hit position in a chamber, better than the $150\text{-}\mu\text{m}$ goal.

Even at the highest rates expected at the LHC, the CSC resolution will stay well within the design specifications (Fig. 193 (left)). A simplified algorithm for hit-position reconstruction that does not use any fitting, iterative procedures, or chamber- or electronics-specific corrections/calibrations was tested on the 12 largest chambers in the MTCC cosmic-ray runs. Being simple and fast, this algorithm is specifically targeted for the HLT. First, 2D-track segments in the chambers are identified by directly accessing the ALCT- and CLCT-pattern records available in the DAQ. Then, the coordinate is calculated by using a simple analytical function $f_w(r)$ of the ratio r built from the charges Q on 3 adjacent strips (centre, right, and left strips):

$$(x/w)_{\text{measured}} = f_w(r), \quad \text{where } r = \frac{1}{2} \left[\frac{Q_{\text{right}} - Q_{\text{left}}}{Q_{\text{centre}} - \min(Q_{\text{right}}, Q_{\text{left}})} \right]. \quad (11)$$

This algorithm localises muon stubs in a chamber with a precision of $< 200 \mu\text{m}$ (Fig. 193 (right)), which is more than adequate for the HLT. The highest muon p_T threshold used by the HLT is 40 GeV. Due to the muon multiple scattering in the calorimeters and in the steel disks, for muons with transverse momenta $p_T < 40$ GeV, one need not measure muon coordinates with a precision much better than ≈ 0.5 mm. This holds true for a muon momentum measurement in the stand-alone muon system, for associating stand-alone muons with tracks in the central tracker, and for the ultimate muon momentum measurement, which is achieved by means of combining information from the central tracker and the muon system.

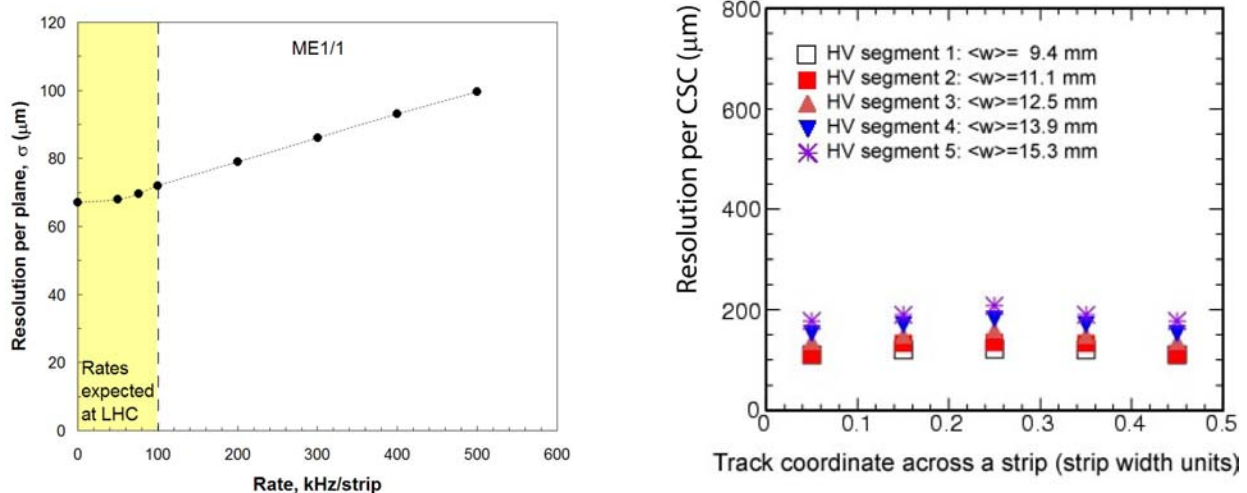


Figure 193: Left panel: deterioration of spatial resolution (ME1/1 chambers) with increasing rate of signals. The resolution remains well within the design specs even at rates far exceeding those expected at the LHC. Right panel: expected 6-plane chamber resolution for nearly perpendicular tracks vs. muon coordinate across a strip as evaluated from the single-plane resolution obtained with a simplified and fast reconstruction specifically targeted for the HLT.

Radiation tolerance The high radiation rates at the LHC could result in devastating problems; thus, the detectors and electronics were designed to be robust. To validate the design, we carried out a series of detailed tests of chamber ageing and electronic board radiation damage.

Ageing studies were conducted [155] by irradiating CSCs at the GIF for several months. The prototype gas system operated in recycling mode as envisioned for full-system operation (2 gas volume exchanges per day with about 5% fresh gas added in each 1-volume cycle). The chambers showed little change in gas gain, dark current, and spurious pulse rate. The total accumulated charge on the wires was about 0.4 C/cm, corresponding to about 50 years of operation at full LHC luminosity in the worst areas closest to the beam line. Upon opening the chambers, we observed a layer of deposits on the cathode surfaces, but not on the anode wires. The deposits on the cathodes, being slightly conductive (established by a small reduction of resistance between strips), did not affect performance (e.g., by the Malter effect [156]).

To test the stability of electronic board performance, we dealt separately with 2 distinct radiation components: total ionisation dose and neutron fluence. The total ionising dose for 10 LHC years is ≈ 20 Gy for on-chamber boards and 2 Gy for peripheral crate electronics. The integrated neutron flux over 10 years ranges from about 10^{10} to 10^{12} cm^{-2} . Analog components of the electronics may suffer a steady and permanent deterioration in performance, while the main danger for digital electronics are Single Event Effects (SEE), including Single Event Upsets (SEU) and Single Event Latching (SEL). Upon an SEE occurrence, the electronics can typically be reset by reloading the FPGAs or cycling the power: SEEs can thus be characterised by the meantime between occurrences.

All electronic chips and components were tested with radiation doses far exceeding the 10 LHC-year equivalent [157]. For final-design boards, no significant deterioration in analog performance was observed (noise, gain, threshold, etc.). All digital-electronic FPGAs were tested for SEEs using typical fluences of $\approx 3 \times 10^{11}$ cm^{-2} . No SEL was observed on any FPGA during testing. SEU rates were dominated by the control logic on the CFEB boards. The SEU rate was lowered significantly by introducing a design with triple-voting logic. The mean time between SEEs on a single CFEB was measured and extrapolated to be 700 h at the LHC neutron fluence. With ≈ 2400 CFEBs in our system, a single CFEB will fail due to an SEU about every 30 min during LHC running, which is an acceptable rate, and will need to be reset.

Reliability Extensive testing of prototypes has shown that the CMS Endcap Muon System based on CSC technology would meet all performance requirements and could be built within the constraints of the construction budget. There are 468 six-plane CSCs in the system, with CSC planes comprising ≈ 5000 m^2 of sensitive area. The total number of read-out channels is about 400 000. During the years of construction and commissioning, the

CMS CSCs have proven to be very reliable in operation (e.g., not a single wire out of about 2 000 000 in the system has ever snapped) and have confirmed the expected performance. As an example, analyses of the first data taken in situ with 36 chambers and cosmic-ray muons showed that chambers had a 99.9% efficiency to detect muon-track segments (input to the Level-1 trigger) and the spatial resolution attainable at the high-level trigger and offline was $\approx 150 \mu\text{m}$.

8.3 Resistive Plate Chamber system

Resistive Plate Chambers (RPC) are gaseous parallel-plate detectors that combine adequate spatial resolution with a time resolution comparable to that of scintillators [158, 159]. An RPC is capable of tagging the time of an ionising event in a much shorter time than the 25 ns between 2 consecutive LHC bunch crossings (BX). Therefore, a fast dedicated muon trigger device based on RPCs can identify unambiguously the relevant BX to which a muon track is associated even in the presence of the high rate and background expected at the LHC. Signals from such devices directly provide the time and position of a muon hit with the required accuracy. A trigger based on RPCs has to provide the BX assignment to candidate tracks and estimate the transverse momenta with high efficiency in an environment where rates may reach 10^3 Hz/cm^2 .

The CMS RPC basic double-gap module consists of 2 gaps, hereafter referred as *up* and *down* gaps, operated in avalanche mode with common pick-up read-out strips in between (Fig. 194) [160, 161]. The total induced signal is the sum of the 2 single-gap signals. This allows the single-gaps to operate at lower gas gain (lower high voltage) with an effective detector efficiency higher than for a single-gap. Table 17 lists the basic construction and operating parameters of the CMS double-gap RPCs.

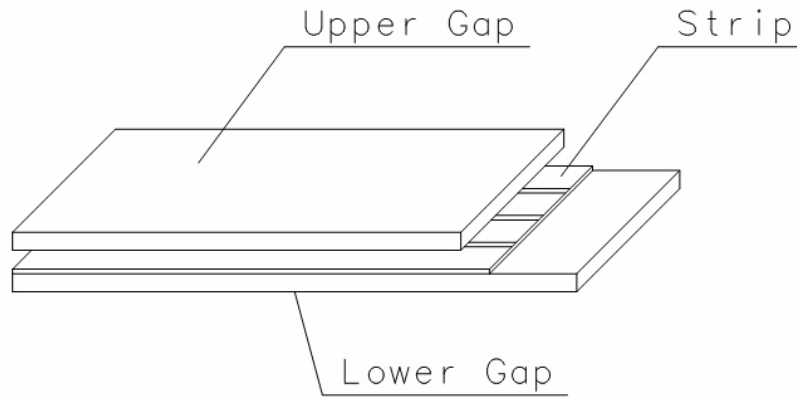


Figure 194: Layout of a double-gap RPC.

Table 17: Basic construction parameters.

Bakelite thickness	2 mm
Bakelite bulk resistivity	$1-2 \times 10^{10} \Omega \cdot \text{cm}$
Gap width	2 mm

Extensive ageing tests have been performed over the past years with both neutron and gamma sources to verify long term detector performance in the LHC background environment [162, 163]. Results confirm that over a period equivalent to 10 CMS-operation years, no efficiency degradation is expected while all other characteristic parameters stay well within the project specifications. Six layers of RPC chambers will be embedded in the barrel iron yoke, 2 located in each of the first and second muon stations and 1 in each of the 2 last stations. The redundancy in the first 2 stations will allow the trigger algorithm to perform the reconstruction always on the basis of 4 layers, even for low p_T particles, which may stop inside the iron yoke. In the endcap region, the baseline design foresees the instrumentation of the iron disks with 4 layers of RPCs to cover the region up to $\eta = 2.1$. However, in the first phase, due to budget limitations, only 3 layers up to $\eta = 1.6$ will be built. In addition, the background rate in the high η region is significantly higher, well beyond the limit reached during the ageing test. Additional R&D to certify the detector performance under such conditions is ongoing.

8.3.1 Detector layout

Barrel system In the barrel iron yoke, the RPC chambers form 6 coaxial sensitive cylinders (all around the beam axis) that are approximated with concentric dodecagon arrays arranged into 4 stations (Fig. 195).

In the first and second muon stations there are 2 arrays of RPC chambers located internally and externally with

Table 18: Number of RPCs for different wheels.

RPC	W+2	W+1	W0	W-1	W-2	Total
RB1in	12	12	12	12	12	60
RB1out	12	12	12	12	12	60
RB2/2in	12	-	-	-	12	24
RB2/2out	-	12	12	12	-	36
RB2/3in	-	12	12	12	-	36
RB2/3out	12	-	-	-	12	24
RB3	24	24	24	24	24	120
RB4	24	24	24	24	24	120
Total	96	96	96	96	96	480

respect to the Drift Tube (DT) chambers: RB1in and RB2in at smaller radius and RB1out and RB2out at larger radius. In the third and fourth stations there are again 2 RPC chambers, both located on the inner side of the DT layer (named RB3+ and RB3-, RB4+ and RB4-). A special case is RB4 in sector 4, which consists of 4 chambers: RB4++, RB4+, RB4-, and RB4--. Finally, in sectors 9 and 11 there is only 1 RB4 chamber.

In total there are 480 rectangular chambers (Table 18), each one 2455 mm long in the beam direction. Exceptions are the chambers in sector 3 of wheel -1 and sector 4 of wheel +1, which are 2055 mm long to allow passage of the magnet cooling chimney. Chambers RB1, RB2, and RB3 have widths 2080, 2500, and 1500 mm, respectively. The widths of the RB4 chambers (which depend on location) are given in Table 19.

Physics requirements demand that the strips always run along the beam direction and are divided into 2 parts for chambers RB1, RB3, and RB4. The RB2 chambers, a special case for the trigger algorithm, have strips divided into 2 parts (RB2in in wheels +2 and -2 and RB2out in wheels +1, 0, and -1) and into 3 parts (RB2out in wheels +2 and -2 and RB2in in wheels +1, 0, and -1). Each chamber therefore consists of either 2 or 3 double-gap modules mounted sequentially in the beam direction to cover the active area. For each double-gap module (up to 96 strips/double-gap), the front-end electronics boards are located at the strip end, which minimises the signal

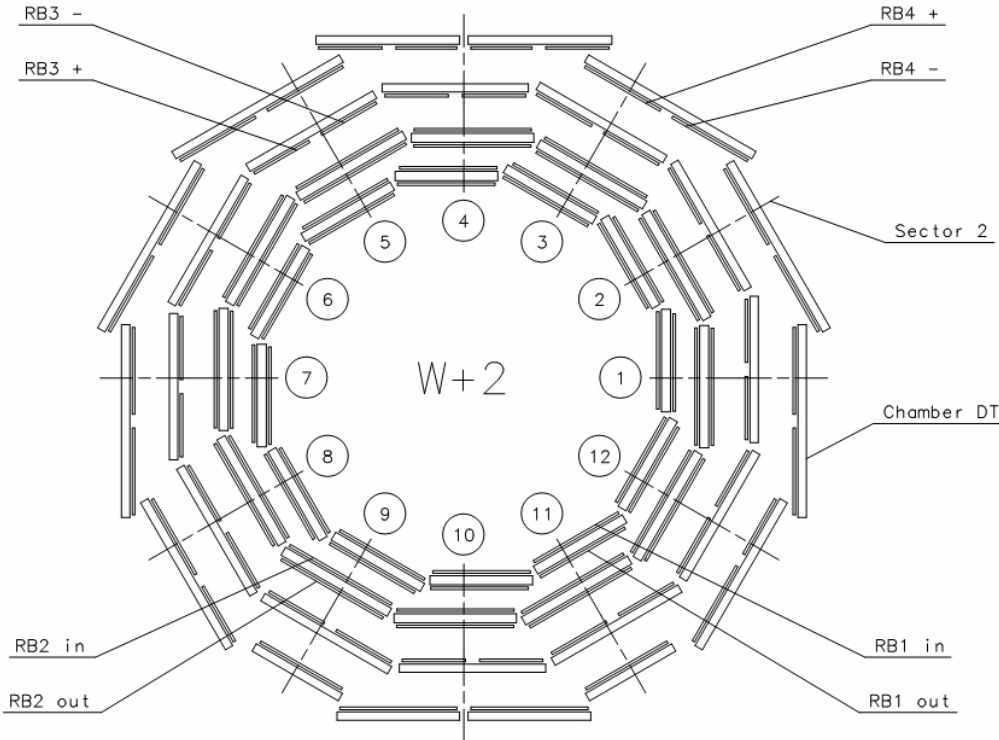


Figure 195: Schematic layout of one of the 5 barrel wheels, which are labeled -2, -1, 0, +1, and +2, respectively. Each wheel is divided into 12 sectors that are numbered as shown.

Table 19: Widths of the RB4 chambers.

Sector	RB4+	RB4++	RB4−	RB4− −	RB4
S1–S3 S5–S7	2000 mm		2000 mm		
S8	1500 mm		2000 mm		
S12	2000 mm		1500 mm		
S9, S11					2000 mm
S10	2500 mm		2500 mm		
S4	1500 mm	1500 mm	1500 mm	1500 mm	

arrival time with respect to the interaction point. The strip widths increase accordingly from the inner stations to the outer ones to preserve projectivity (each strip covers $5/16^\circ$ in ϕ). Figures 196 and 197 show schematic views of chamber modules with 2 and 3 double-gaps, respectively. Table 20 lists some global information regarding the barrel detector.

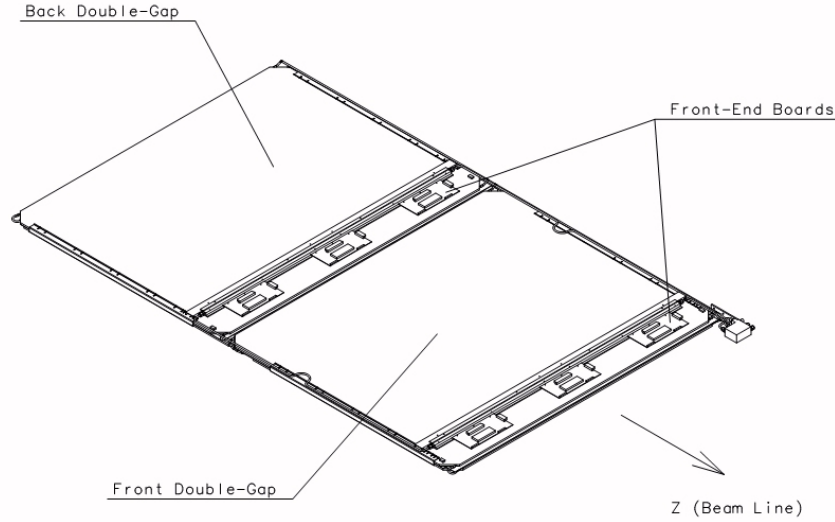


Figure 196: Schematic layout of chamber module with 2 double-gaps.

Table 20: Barrel RPC system global parameters.

Number of stations	480
Total surface area	2400 m ²
Number of double-gaps	1020
Number of strips	80 640

Endcap system In the forward and backward regions of the CMS detector, 3 iron disks constitute the endcap yokes. Like in the barrel, 2 complementary muon detector systems are deployed for robust muon identification: cathode strip chambers (CSC) and RPCs. They are mounted on both faces of the disks to yield 4 CSC stations (ME1–4) and, for the initial detector, 3 RPC stations (RE1–3). The double-gaps in every station have a trapezoidal shape and are arranged in 3 concentric rings as shown in the r - ϕ view of Fig. 198. They overlap in ϕ as to avoid dead space at chamber edges. Except for station 1, the chambers of the innermost ring span 20° in ϕ , all others span 10° . As mentioned before, the high η part of the RPC system (beyond $\eta \approx 1.6$) has been staged until the LHC is scheduled to deliver its design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.

Station RE1 is mounted on the interaction point (IP) side of the first endcap disk (YE1), underneath the CSC chambers of ME1. Stations RE2 and 3 are mounted on the back side of YE1 and on the IP side of YE3, respec-

tively. They remain uncovered since the corresponding CSC stations 2 and 3 are mounted on both faces of YE2. Figure 199 shows a schematic layout of the CMS endcap defining the nomenclature of the muon stations. Each endcap RPC chamber consists of a double-gap structure enclosed in a flat trapezoidal shaped box made of 2 aluminium honeycomb panels of 6 mm thickness each and a $16 \times 16 \text{ mm}^2$ section spacer frame (Fig. 200). The strip panel, sandwiched in between the gas gaps, has copper strip sections on a G10 support. Strips run radially and are radially segmented into 3 trigger sections for the REn/2 and REn/3 chambers ($n = 1-3$). The 32 strips of the 10° RPC chambers are projective to the beam line, following a homothetic pattern. Besides the different mechanical shape and assembly, the front-end electronics, services, trigger, and read-out schemes of the endcap RPC system are identical to the barrel system. To an operator, the CMS barrel and endcap RPC systems look identical.

8.3.2 Readout electronics

Front-end electronics The read-out strips are connected to Front-End Boards (FEB). After having been amplified and discriminated, signals are sent unsynchronised to Link Boards (LB) placed around the detector. The LBs synchronise the signals with the 40-MHz LHC clock and transmit them to the trigger electronics located in the CMS counting room over a 90-m optical link at 1.6 GHz, as shown in the block diagram of Fig. 201.

The FEBs house two (barrel version) or four (endcap version) front-end chips, which are custom ASICs designed in AMS $0.8 \mu\text{m}$ CMOS technology [164]. Each chip is made of 8 identical channels, each consisting of an amplifier, zero-crossing discriminator, one-shot, and LVDS driver. The preamplifier is a trans-resistance stage with $15\text{-}\Omega$ input impedance to match the characteristic impedance of the strips. It is followed by a gain stage to provide an overall charge sensitivity of 2 mV/fC .

Since accurate RPC timing information is crucial for providing an unambiguous bunch-crossing assignment of the event, the zero-crossing discrimination technique was adopted to make the timing response amplitude-independent. In fact, considering that the RPC signals have a wide dynamic range (from few tens of fC to 10 pC), the implemented architecture provides a time walk below 1 ns , while the simpler leading-edge discrimination technique would have provided a time walk of $\approx 10 \text{ ns}$. The discriminator is followed by a one-shot circuit. This produces a pulse shaped at 100 ns to mask possible after-pulses that may follow the avalanche pulse. Finally, an LVDS driver is used to send the signals to the LB in differential mode.

Gamma-irradiation tests showed no performance degradation of either the front-end chip or the control electronics on the FEB [165]. Moreover, tests with thermal and fast reactor neutrons (0.4 eV – 10 MeV) and with more energetic neutrons (65 MeV), have certified that the circuit can sustain the expected CMS operational conditions [166].

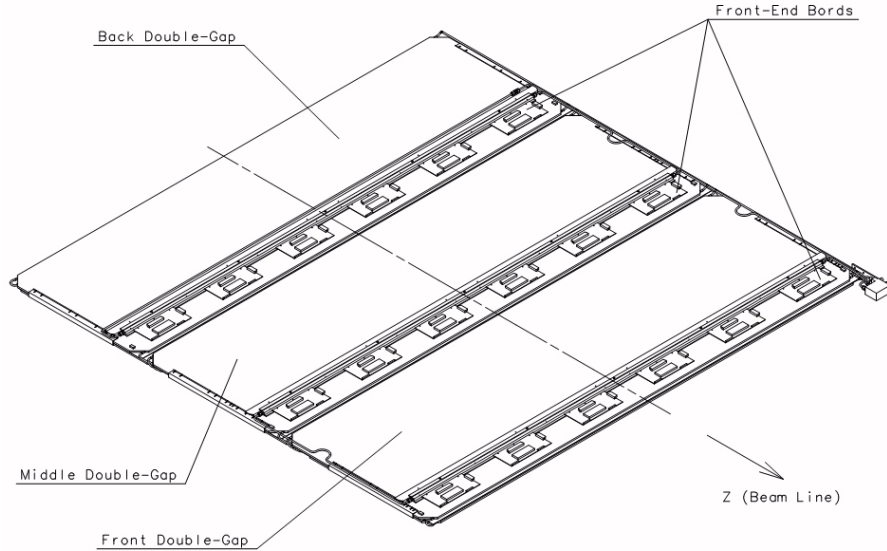


Figure 197: Schematic layout of chamber module with 3 double-gaps.

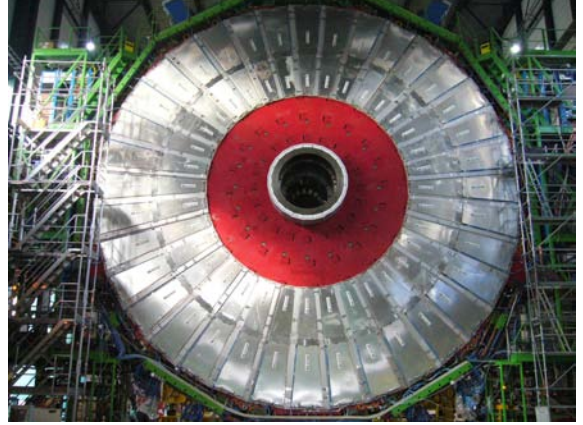
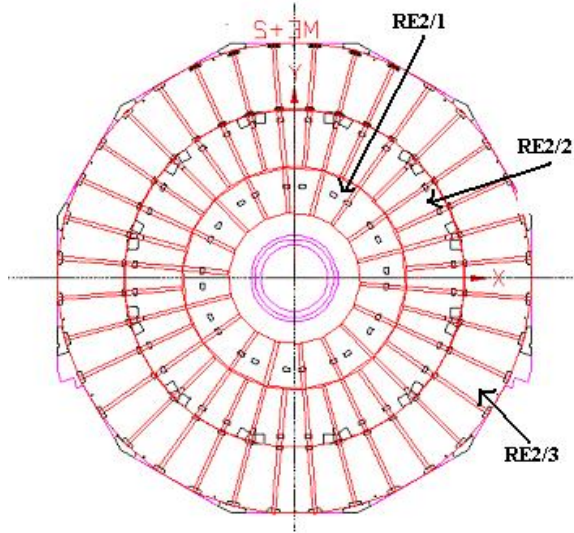


Figure 198: Left panel: Schematic r - ϕ layout of RPC station RE2 on the back side of the first endcap yoke. Right panel: RPC station RE2 on the back side of the YE-1 yoke. The inner ring has been staged and is absent here.

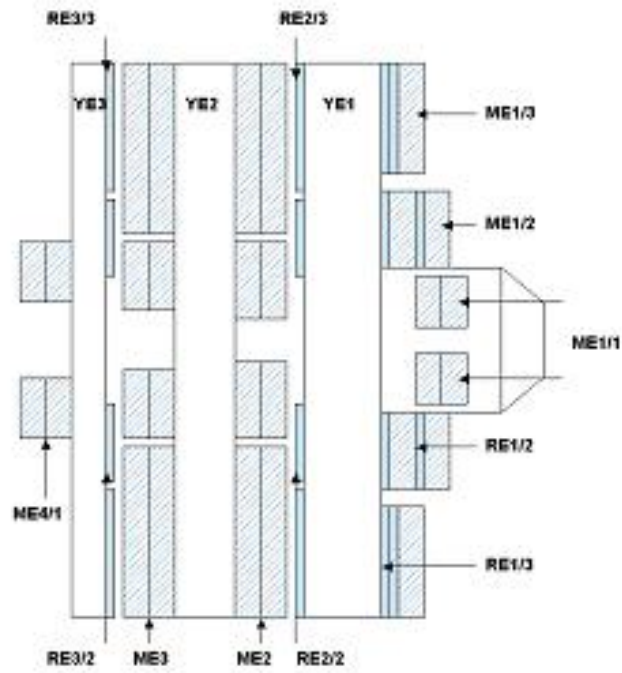


Figure 199: Schematic layout of the CMS endcap for the initial muon system.

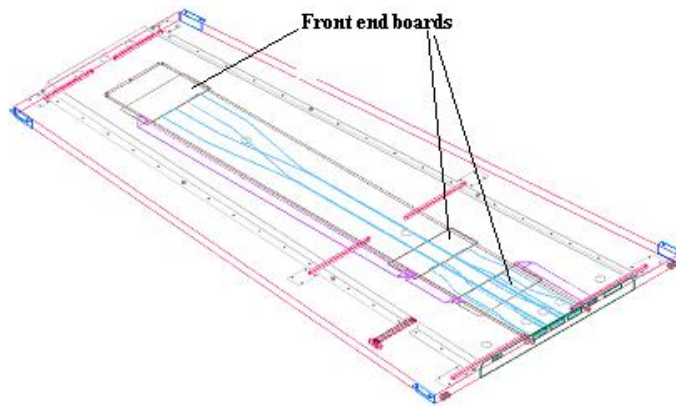


Figure 200: A view of an endcap RPC chamber.

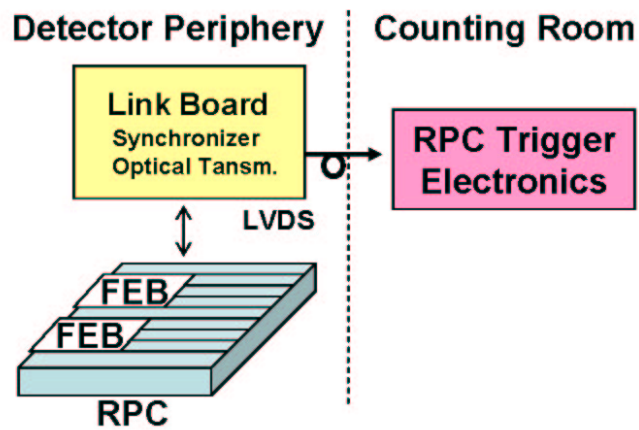


Figure 201: Block diagram of RPC read-out electronics.

RPC technical trigger electronics Study of the detector performance is a crucial aspect during the detector-commissioning phase. The RPC trigger system was designed to identify muon tracks starting from the interaction point. Therefore, all interconnections among the LBs and trigger electronics were optimised to fulfil a projective vertex geometry, not adequate for triggering on cosmic rays. Therefore, an RPC-based trigger (Technical Trigger in the following) has been implemented by means of 2 types of electronic boards: the RPC Balcony Collector (RBC) housed in the cavern and the Technical Trigger Unit (TTU) located in the counting room (Fig. 202) [167].

The RBC collects 96-strip OR signals from the barrel LBs and produces a “local” sector-based cosmic trigger to be used during commissioning or calibration of the detector. The RBC transmits the ORs optically to the TTU boards in the Counting Room (30 fibres in total), where a wheel-level cosmic trigger is produced and sent as a Technical Trigger to the CMS Global Trigger. A proper algorithm for searching for cosmic-muon tracks is implemented in the TTU.

8.3.3 Low voltage and high voltage systems

General requirements The RPC power systems operate in a hostile environment due to the high magnetic field and high radiation flux. Large portions of the power systems will be near the detectors on the balcony racks placed around the barrel wheels and the endcap disks. In these areas the magnetic field can reach 1 T with radiation around $5 \cdot 10^{10}$ protons/cm² and $5 \cdot 10^{11}$ neutrons/cm². In cooperation with the ATLAS, ALICE, and LHCb groups, the CMS collaboration developed a new design for an RPC power system able to operate in such conditions. The main requirements for the RPC HV and LV power supplies are collected in Table 21.

Table 21: HV and LV power supply requirements.

	HV	LV
Maximum Voltage	12 kV	7 V
Maximum Current	1 mA	3 A
Ripple	<100 mV pp at load (freq <20 MHz)	<10 mV pp at load (freq <20 MHz)
Programmable Voltage	from 0 to 12 kV	from 0 to 7 V
Current monit. precision	0.1 μ A	100 mA
Voltage monit. precision	<10 V	100 mV

The HV and LV systems are both based on a master/slave architecture. The master, called the mainframe, is devoted to controlling and monitoring one or more slaves and is placed in a safe and accessible area like the control room. The slaves can be located near the detector and are designed to be modular and multi-functional to accept both HV and LV boards. These have to work in a hostile and inaccessible area and are based on radiation-tolerant and

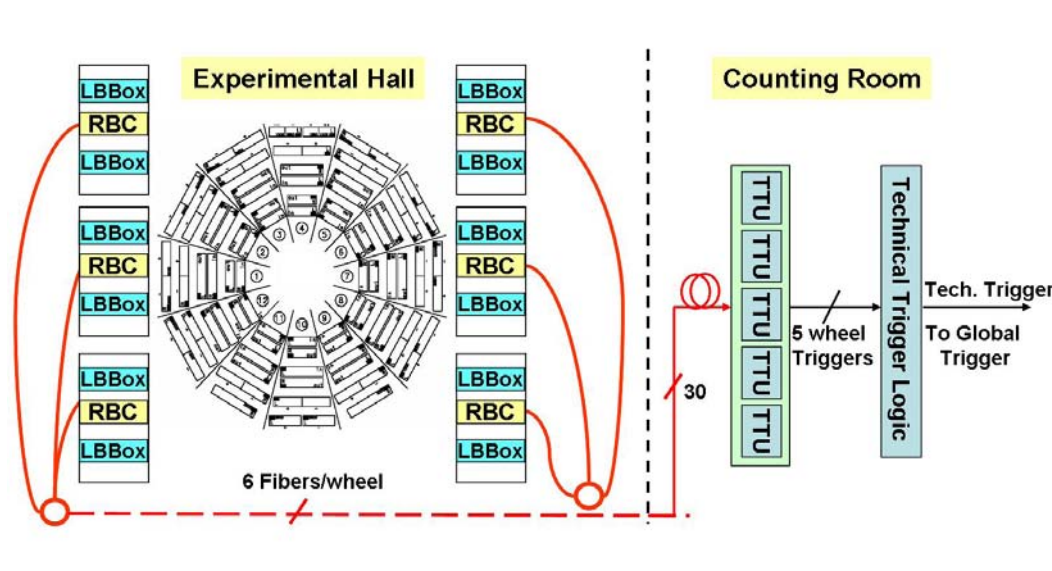


Figure 202: Technical Trigger schematic layout.

magnetic-field-tolerant electronics.

Past experience with RPC detector systems, however, suggested that it is important to have the HV power supplies in an accessible area. In case of unsustainable high current on a detector, the possibility of removing a channel during operation should be available. Therefore, the CMS RPC collaboration decided to keep the master/slave architecture for both the HV and LV systems but to move all HV system components into the control room.

HV and LV system description The system is based on the EASY (Embedded Assembly SYstem) project. It is made of a master SY1527 (mainframe) which houses up to 16 branch controller boards (A1676A) and of EASY3000 crates (slaves). The EASY3000 crate can house different boards (high and low voltage, ADC, and DAC). Each EASY3000 board operates as a channel of the A1676A and can be accessed through the mainframe. The EASY architecture foresees 2 independent 48-V power supplies to power independently the channel regulators and the control logic. The EASY system is connected to the external world through a serial port and an ETHER-NET3 interface that allows the user to monitor and control the whole system with various software from a very easy TELNET interface to a more sophisticated OPC protocol.

HV hardware

The A3512 double-width board is equipped with 6 floating 12 kV/1 mA channels of either positive or negative polarity. The 6 channels have an independent return to avoid ground loops. The board is designed with an output voltage that can be programmed and monitored in the 0–12 kV range with 1 V resolution and with a monitored current resolution of 0.1 μ A. This current resolution allows the Detector Control System (DCS) to study the current behaviour of every chamber with an accuracy of at least 1/10 of the measured current (between 10 and 20 μ A per chamber). In the barrel there is 1 HV channel per chamber, while in the endcap region 1 channel supplies 2 chambers. In this last case, an upgrade of the system will depend on future budget availability. A summary of the HV systems is given in Table 22.

Table 22: Summary of HV systems.

	Barrel	Endcap
HV channels	480	216
HV boards	80	36
Easy3000 Crates	14	6

LV hardware

The CAEN A3009 board is a 12-channel 8V/9A power-supply board for the EASY Crate. It was developed for operation in magnetic fields and radioactive environments. The output-voltage range is 1.5–8 V with 5-mV monitor resolution; channel control includes various alarms and protections. The output current is monitored with 10-mA resolution.

Each chamber is supplied by 2 LV lines for the front-end analog (LV_a) and digital (LV_d) parts. To avoid ground loops on the detector, it is important to preserve, when possible, a 1-to-1 correspondence between LV channel and chamber. This is achieved in the barrel system, where there is 1 LV_a and 1 LV_d channel per chamber. However, for the endcap detector, at the start-up 1 LV_a and 1 LV_d will be distributed between 2 chambers. A summary of the LV systems is given in Table 23.

Table 23: Summary of LV systems.

	Barrel	Endcap
LV channels	720	432
LV boards	60	36
Easy3000 Crates	20	14

8.3.4 Temperature control system

RPC operation is sensitive to both temperature and atmospheric pressure. Therefore, the chambers are constantly monitored to compensate in real time for the detector operating point (HV value). A network of 420 sensors located inside the barrel chambers is available to monitor the temperature. The AD592BN sensor (Analog Devices) can work in a hostile environment with a resolution of about 0.5 °C, better than the CMS requirement (1 °C). Sensors are read out by a 128-channel ADC equipped with a 12-V input stage.

Additional sensors are available on each front-end board; they will be read out through the LB electronics to monitor the temperature.

8.3.5 Gas system

Test results [168] showed that RPCs are suitably operated with a 3-component non-flammable mixture of 96.2% R134a ($\text{C}_2\text{H}_2\text{F}_4$), 3.5% $i\text{C}_4\text{H}_{10}$ and 0.3% SF_6 . Water vapour is added to the gas mixture to maintain a relative humidity of about 45% and to avoid changes of the bakelite resistivity. The basic function of the gas system is to mix the different gas components in the appropriate proportions and to distribute the mixture to the individual chambers. The large detector volume and the use of a relatively expensive gas mixture make a closed-loop circulation system mandatory. The main gas-system parameters are given in Table 24.

Table 24: Main gas parameters of the CMS RPC system.

Gas volume	14 m ³
Gas mixture composition	96.2% R134a, 3.5% $i\text{C}_4\text{H}_{10}$, and 0.3% SF_6
Internal chamber pressure above atmosphere	3 mbar
Nominal flow rate	10 m ³ /h
Fresh gas replenishing rate	0.2 m ³ /h
Number of gas channels	250 (barrel) + 144 (endcaps)

The system consists of several modules: the primary gas supply, mixer and closed-loop circulation system, gas distributors to the chambers, purifier, pump, and gas-analysis station [131, 169]. The full closed-loop circulation system (Fig. 203) extends from the surface gas building SGX to the USC55 service cavern and UXC55 experimental cavern.

Mixer The primary gas supplies and the mixer are situated in the SGX building. The flow of component gases is controlled by mass-flow meters. Flows are monitored by a computer-controlled process, which constantly calculates and adjusts the mixture percentages supplied to the system. The gas mixture is maintained non-flammable by permanent monitoring. The gas flow is stopped automatically if the $i\text{C}_4\text{H}_{10}$ fraction increases beyond the flammability limit. For fast detector filling, parallel rotameters will be used, instead of the mass-flow controllers, yielding a complete volume renewal in about 8 hours.

Closed-loop circulation The mixed gas is circulated in a common closed loop for the barrel and both endcaps. The circulation loop is distributed among 3 locations:

- the purifier, gas input, and exhaust gas connections are located in the SGX building;
- the pressure controllers, separation of barrel and endcaps systems, compressor, and analysis instrumentation are located in the USC (accessible at any time);
- the manifolds for the chamber-gas supplies and channel flow meters are mounted in the distribution racks near the detector.

The high density of the used mixture generates a hydrostatic pressure of about 0.3 mbar/m above atmospheric pressure. Since the total RPC detector height is about 15 m, the barrel detector is split into 2 zones (top and bottom) that have independent pressure regulation systems (Fig. 204). Each barrel muon station has an independent gas line. The 2 RPC chambers located in a station are supplied in parallel from the same patch panel sitting nearby. This configuration leads to 250 gas channels (50 per wheel) for the full barrel detector (Table 25).

Table 25: Chamber volumes and gas flow for a single wheel of the barrel detector.

Station	RPCs in that station	Volume per RPC (l)	number of gas channels	Volume per gas channel (l)	Operating channel flow (l/h)	total flow per station type (l/h)
RB1	24	20.6	12	41.2	27.5	330
RB2	24	25.4	12	50.8	33.9	406
RB3	12	31.4	12	31.4	20.9	251
RB4	12	43.4	14	43.4	28.9	347
Total	72		50			1334

Each endcap detector consists of 3 disks, RE1, RE2, and RE3, with a total of 216 double-gap chambers. Each disk is composed of 2 concentric rings (i.e., for RE n : RE n /2 and RE n /3) of 36 chambers each. In the RE1 rings the chambers are divided in 6 ϕ sections of 60°. A section contains 6 chambers and is supplied with 2 gas lines for the up and down gaps (Fig. 205a). The gas flow in the up gap is in the opposite sense to that in the down gap to improve the average gas quality. In the RE2 and RE3 stations, the chambers are divided into 12 ϕ sections of 30°. Each section contains 3 chambers of the external ring and the corresponding 3 chambers of the internal ring, i.e., an RE2 section includes 3 RE2/2 and 3 RE2/3 chambers (Fig. 205b). In RE2 and RE3 as well there are 2 gas lines per section (for the up and down gaps) and the flows are in the opposite sense between the two. The total number of channels and the relative gas flows are summarised in Table 26.

Pressure regulation system and gas distribution in UXC Pressure regulation is achieved in the USC area for each of the 2 zones. Each height section has its own pressure control and protection system consisting of bubblers located in the distribution racks at the bottom of the wheels/disks. The oil level is adjusted to account for the hydrostatic pressure differences in the 2 zones. The distribution racks are installed at the bottom of each wheel/disk. The supply and return lines for each station are equipped with a mass-flow meter and a needle valve (only at the inlet). The flow measurements allow the detection of possible leaks, while the needle valves are used for the flow adjustment between different stations.

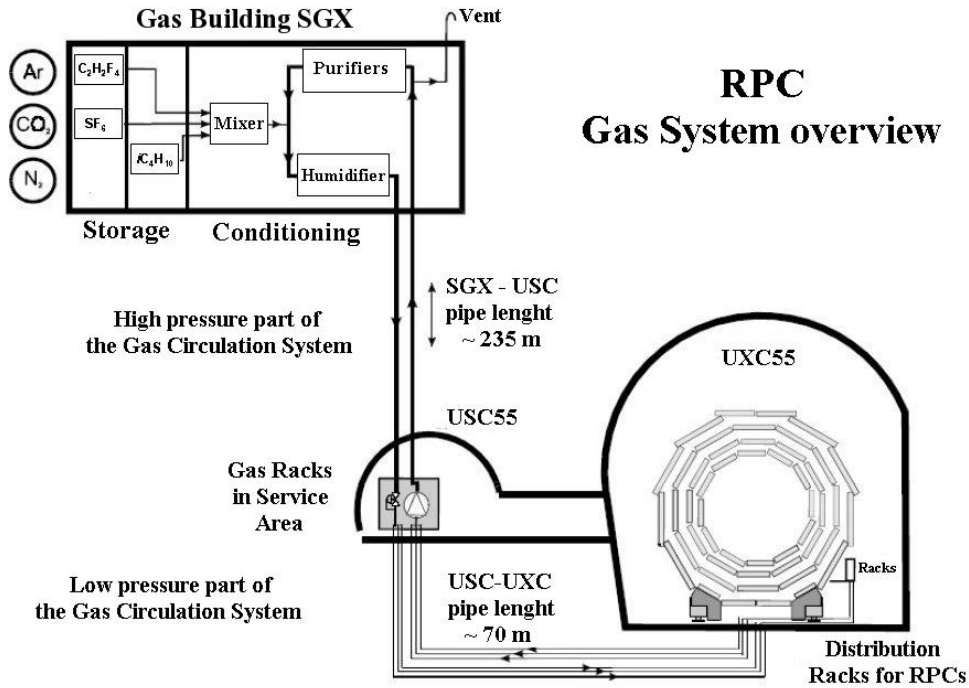


Figure 203: Closed-loop circulation system.

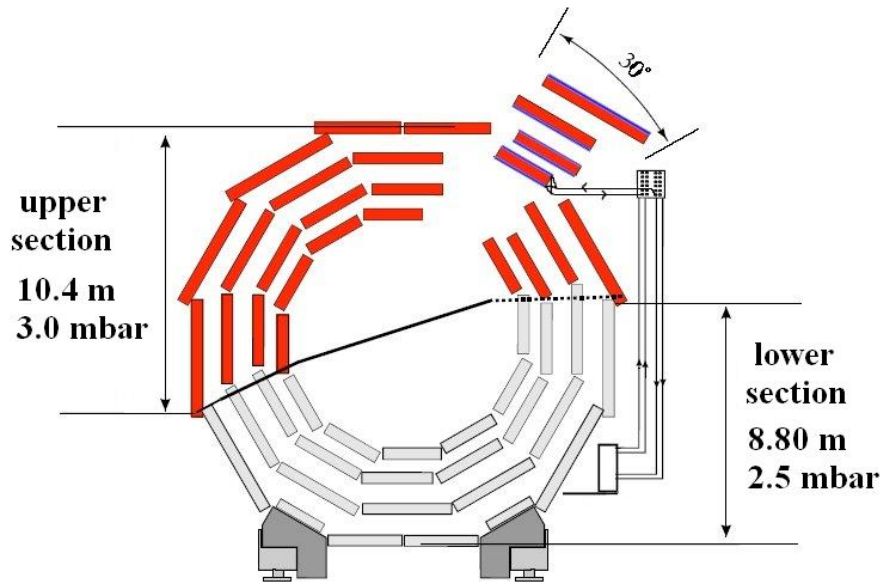


Figure 204: The 2 zones into which a wheel is sub-divided. Each station (2 chambers) is supplied by a gas line.

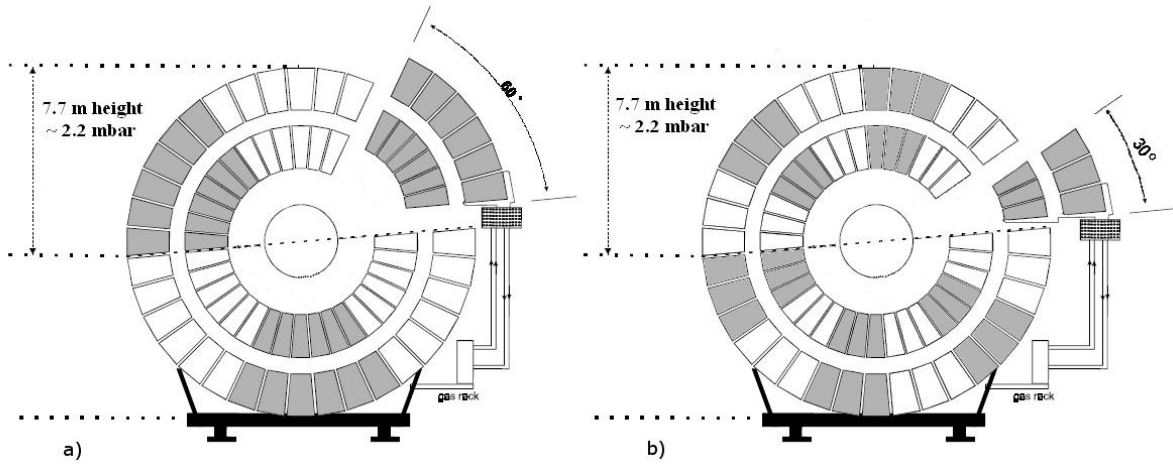


Figure 205: In each endcap disk the RPC detectors are divided in 2 rings. RE1 (a) is divided into 60° sectors, while in RE2 and in RE3 (b) sectors are composed of 3 chambers of the internal ring and the corresponding 3 of the external ring. Every sector is supplied by 2 independent gas lines.

Table 26: Chamber volumes and gas flows for a single endcap of the CMS RPC system.

Module	RPCs in that module	Volume per RPC (l)	number of gas channel	Volume per gas channel (l)	Operating channel flow (l/h)	total flow per module type (l/h)
RE/1/2	36	5.1	12	30.6	20.4	122
RE/1/3	36	7.4	12	44.4	29.6	178
RE/2/2	36	5.1	24	20.3	13.5	324
RE/2/3	36	8.4				
RE/3/2	36	5.1	24	20.3	13.5	324
RE/3/3	36	8.4				
Total	216		72			948

Purifier Results from long term tests performed by CMS showed that the impurity concentrations produced in the RPC chambers are high enough to influence the detector performance if they are not properly removed from the mixture. Therefore, to achieve a high recycling rate the closed-loop circulation system is equipped with a purifier module containing 3 cleaning agents. In the first running phase the 3 cleaning agents are contained in 2 purifiers. Both the purifiers are 24-l cartridges. The first is filled with a 0.5-nm molecular sieve, while the second is filled with the following combination: 25% Cu-Zn filter (type R12, BASF), 25% Cu filter (type R3-11G, BASF), and 50% Ni-Al₂O₃ filter (type 6525, Leuna). During the high luminosity running period the second purifier will be split into 2 separate 24-l cartridges, the first containing the R12 and R3-11G cleaning agents and the second containing the 6525 Leuna filter. Each purifier is equipped with an automatic regeneration system: 2 identical cartridges are present allowing the regeneration of a cartridge while the other is in use.

Gas-quality monitoring Two independent systems are in preparation to continuously monitor the gas quality. The gas-gain monitoring system [170] is based on several small ($50 \times 50 \text{ cm}^2$) single-gap RPCs whose working points (gain and efficiency) will be continuously monitored online. The system is designed to provide a fast and accurate determination of any shift in the working points. The small single-gap RPCs are divided into several sub-groups supplied with gas coming from different parts of the full system (i.e., fresh gas mixture, input to the chambers in closed-loop circulation, and return from the closed-loop circulation). The second gas monitoring system [171] will perform both qualitative and quantitative gas chemical analyses with a set-up that includes a gas chromatograph, pH sensors, and specific fluoride electrodes. In the underground service cavern (USC), many sampling points equipped with manual valves will allow the analysis of the gas mixtures that return from every half wheel. In the surface gas building (SGX), sampling points will be available to monitor the effectiveness as well as the status of each cartridge in the purifier module.

8.3.6 Chamber construction and testing

In view of the extremely large-scale production (a factor of 10 greater than in past experiments), impressive quality control and certification protocols were set along the production chain at many different levels:

- selection of electrodes and resistivity certification;
- certification of single-gaps and double-gaps;
- chamber testing.

Details regarding the quality certification procedures have been reported elsewhere [172]. Only a short summary of the chamber testing results is given below.

Chamber performance at the test sites Several RPC test stands were in operation. Each telescope consisted of a tower in which several detectors could be placed horizontally and read out in coincidence with the passage of the crossing of cosmic muons. Two sets of scintillators, at the top and the bottom of the telescope, were used for triggering purposes. Atmospheric and environmental conditions were continuously monitored during the tests. These values were used to scale the applied HV for temperature and pressure variations to evaluate the effective high voltage (HV_{eff}) [173] at given reference values ($T_0 = 293 \text{ K}$ and $P_0 = 1010 \text{ mbar}$).

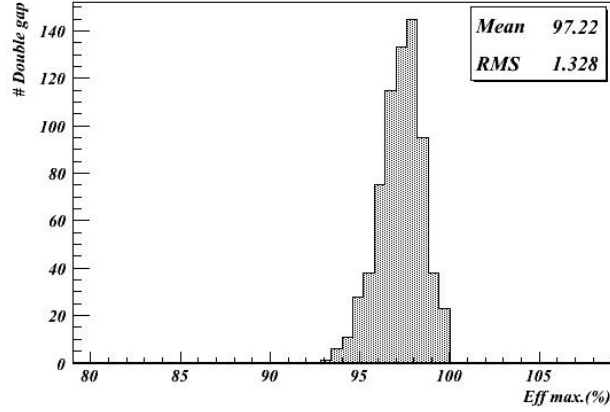


Figure 206: Distribution of plateau efficiency for all the barrel chambers.

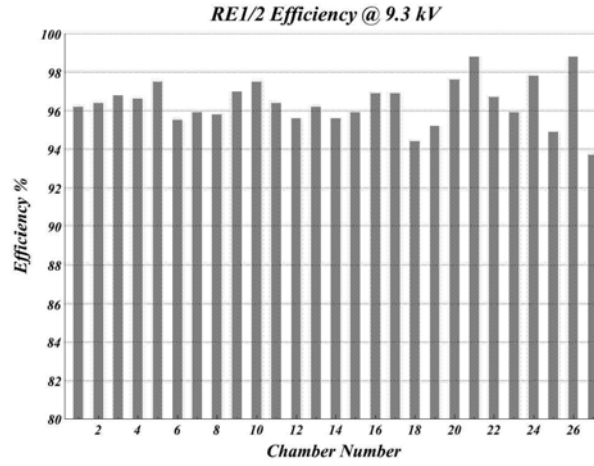


Figure 207: Efficiency at HV = 9.3 kV for the first 27 endcap chambers.

The final gas mixture (96.2% $C_2H_2F_4$, 3.5% iC_4H_{10} , and 0.3% SF_6) was used and water vapour was added to keep the gas relative humidity at a value of about 45%.

The tracking capabilities of the test telescope provided a full characterisation of the detectors in terms of efficiency, cluster size, and noise rate. Also the chambers' local efficiency and the spatial resolution were studied. Rigorous and automatic protocols were developed and systematically applied at all test sites in order to accept chambers that satisfied the CMS requirements.

First, the chamber efficiency was studied with the “coincidence” method by evaluating the ratio between the number of events in which an RPC module had at least 1 fired strip in the trigger window (100 ns) and the total number of recorded events. In Fig. 206 the distribution of the maximum efficiency for all the barrel RPCs is shown. The mean value of the distribution is 97.2%. In Fig. 207 the efficiency distribution at HV = 9.3 kV for the first 27 endcap chambers is shown.

The chamber response uniformity was also studied by performing track recognition through the telescope. Muon trajectories were reconstructed in 2-dimensional views, where the x coordinate is defined by the strip position along the chamber and the y coordinate by the chamber position in the tower. Details about the pattern recognition algorithm have been presented elsewhere [174]. The track-impact point on the chamber under test was also determined and the distance to the nearest cluster centre was evaluated. A chamber was considered efficient if the reconstructed muon trajectory matched the fired strip. A typical strip-by-strip efficiency plot is shown in Fig. 208.

The chamber cluster size is defined as the average value of the cluster-size distribution. In Fig. 209 the profile histogram of the cluster-size distribution as a function of the HV_{eff} is shown for all the barrel chambers. A chamber was accepted if the cluster size was below 3 strips at the knee of the efficiency plateau.

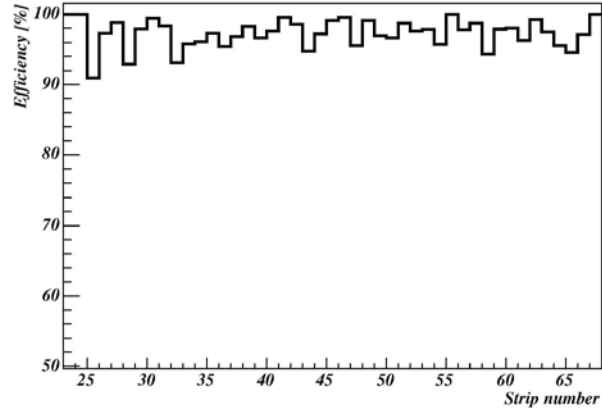


Figure 208: Local efficiency at $HV_{\text{eff}} = 9.6$ kV for a barrel chamber.

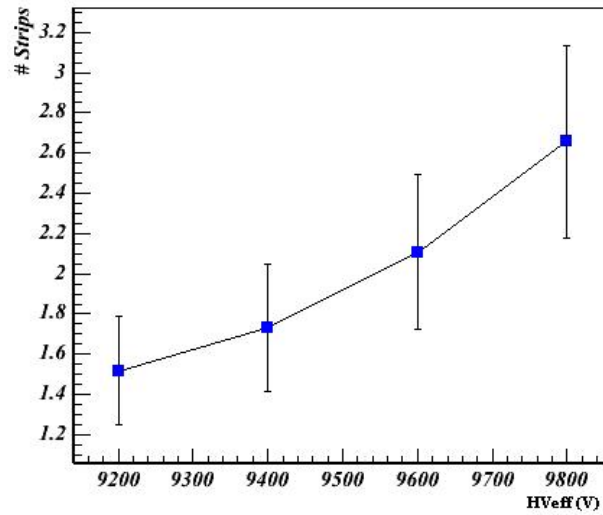


Figure 209: Profile histogram of the chambers' cluster-size distribution as a function of HV_{eff} . The dots and bars are the average and the root-mean-square of the cluster-size distributions, respectively.

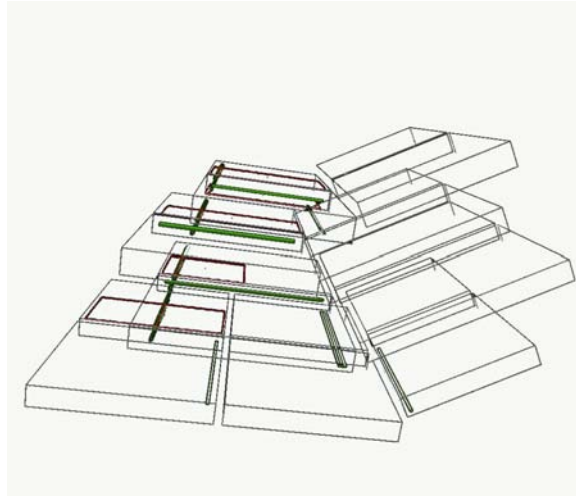


Figure 210: Iguana muon reconstruction: RPC-fired strips are in red and DT hits in green.

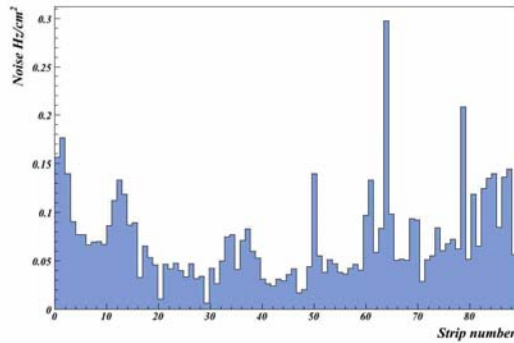


Figure 211: Barrel chamber noise profile at $HV_{\text{eff}} = 9.6$ kV.

Magnet Test and Cosmic Challenge (MTCC) During the summer and fall of 2006 a first integrated test of an entire CMS “slice” was performed in the SX5 experimental surface hall. For the RPC system, 3 barrel sectors and a 60° portion of the first positive endcap disk were involved in the test. The chambers were operated with their final power system configuration, and CMS DAQ software, data quality monitor (DQM), and detector control system (DCS) were implemented for the detector read-out and control.

The RPC Balcony Collector (RBC) board provided a cosmic trigger with a selectable majority level of signals from the 6 RPC barrel chambers. A trigger rate of about 30 Hz/sector for a majority level of 5/6, and 13 Hz/sector for a 6/6 majority was found while operating the detector on the surface. The RBC trigger was well synchronised with the other muon detector (DT and CSC) triggers.

Several millions of events were collected with different trigger configurations. The DQM was used successfully during the MTCC. It allowed the online checking of the quality of the data and the chamber behaviour in terms of cluster size, number of clusters, etc. Figure 210 shows the event display of a typical cosmic muon triggered by the RBC and crossing both RPCs and DTs.

Specific runs were taken before and during the test to evaluate the noise rate. Preliminarily, all the threshold values on the front-end electronic discriminators were set to achieve the best noise configuration with higher efficiency. The chamber noise rate profile for a barrel station is shown in Fig. 211 at $HV_{\text{eff}} = 9.6$ kV, while Fig. 212 shows the overall noise distribution for all the barrel strips involved in the test.

The RPC efficiency can be studied by extrapolating DT segments to the corresponding RPC layer and by requiring matching hits within an appropriate width. In Fig. 213 the chamber efficiency as a function of the HV_{eff} is shown for some RPC stations.

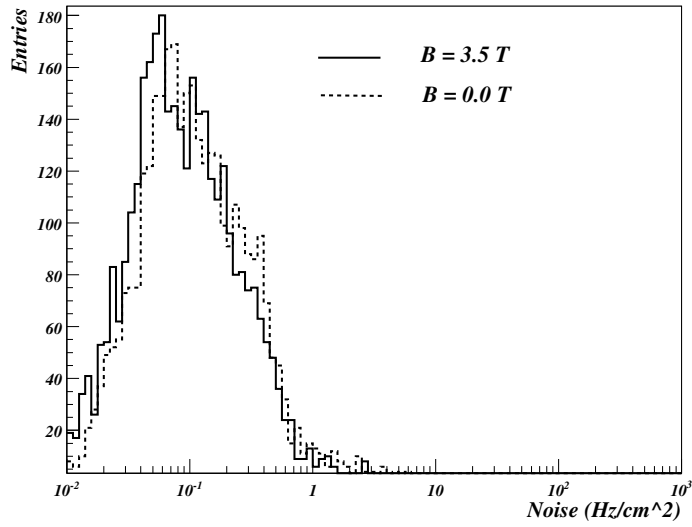


Figure 212: Noise distributions in 2 different magnetic fields. All strips of the barrel stations are included in the distribution.

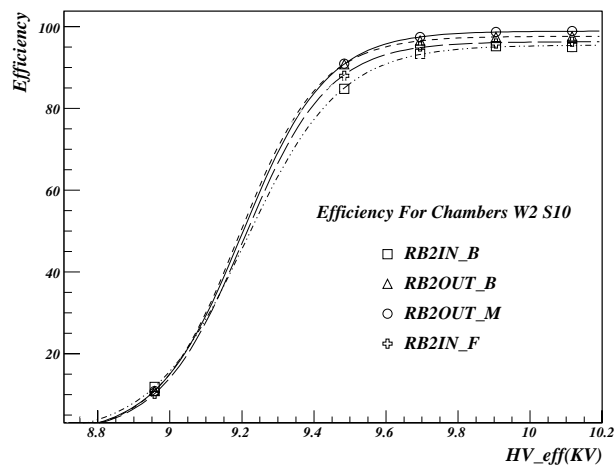


Figure 213: Global efficiency vs. HV_{eff} , estimated by means of DT segment extrapolation.

Results are in agreement with those obtained during testing at construction sites and fully meet the design specifications.

8.4 Optical alignment system

For optimal performance of the muon spectrometer [131] over the entire momentum range up to 1 TeV, the different muon chambers must be aligned with respect to each other and to the central tracking system to within a few hundred μm in $r\phi$. The required alignment precision for the endcap chambers is 75–200 μm , while for the barrel the precision varies from 150 μm for the inner chambers of Station 1 to 350 μm for the outer chambers of Station 4. To this end, after following strict chamber construction specifications, CMS combines precise survey and photogrammetry measurements, measurements from an opto-mechanical system, and the results of alignment algorithms based on muon tracks (both from cosmic rays and from pp collisions) crossing the spectrometer.

There are several potential sources of misalignment in the muon spectrometer, from chamber production to final detector operating conditions, including:

- Chamber construction tolerances. These are unavoidable geometrical tolerances in the production of the chamber parts, such as mis-positioning of wires or strips within a layer and relative shifts in the layer-superlayer assembly. The relative positioning of the different internal components of a chamber was measured during construction to be within the required tolerances (Sect. 8.1 and 8.2). After assembly, all chambers were tested with cosmic muon data and showed good correlation between those measurements and the results of muon track fits. Furthermore, the geometry of the DT chambers was measured at the CERN ISR assembly hall using optical and survey techniques. These data are compared with construction drawings and cosmic data to provide corrections to the nominal chamber geometry when necessary.
- Detector assembly, closing tolerances. Gravitational distortions of the return yoke lead to static deformations of the steel support. This effect, together with the installation tolerances, results in displacements of the chambers in the different barrel wheels and endcap disks of up to several millimetres with respect to their nominal detector positions. After chamber installation, survey and photogrammetry measurements were performed for each wheel and disk. These measurements provide an initial geometry—position and orientation of each muon chamber in the different yoke structures—which absorbs installation tolerances and static steel deformations [175].
- Solenoid effects. Magnetic field distortions lead to almost perfect elastic deformations of the return yoke, at the level of a few centimetres. They result in further displacement of the chambers. The new detector geometry resulting from the magnetic forces is accessed with measurements of the optical system and track-based alignment techniques.
- Time-dependent effects. During operation, thermal instabilities and other time-dependent factors can cause dynamic misalignments at the sub-millimetre level.

The Muon Alignment (MA) system was designed to provide continuous and accurate monitoring of the barrel and endcap muon detectors among themselves as well as alignment between them and the inner tracker detector. To fulfil these tasks the system is organized in separate blocks: local systems for barrel and endcap muon detectors to monitor the relative positions of the chambers, and a link system that relates the muon and central tracker systems and allows simultaneous monitoring of the detectors.

The system must generate alignment information for the detector geometry with or without collisions in the accelerator. The dynamic range of the system allows it to work in the solenoidal magnetic field between 0 and 4 T. Its goal is to provide independent monitoring of the CMS tracking detector geometry with respect to an internal light-based reference system. This will help to disentangle geometrical errors from sources of uncertainty present in the track-based alignment approach, e.g., knowledge of the magnetic field, material description, and drift velocity.

The basic geometrical segmentation consists of 3 r - z alignment planes with 60° staggering in ϕ . This segmentation is based on the 12-fold geometry of the barrel muon detector. Within each plane, the 3 tracking sub-detectors of CMS (central tracker, barrel and endcap muon detectors) are linked together. Figure 214 shows schematic longitudinal and transverse views of CMS, with the light paths indicated. Furthermore, the barrel and endcap monitoring systems can work in stand-alone mode, in which they provide reconstruction of the full geometry of each independent sub-detector. The layout of the optical paths allows the monitoring of each of the 250 DT chambers, while only one sixth of selected CSCs in the 4 endcap stations are directly monitored. Alignment sensors located in the region between the muon barrel wheels and endcap disks allow the tracker and muon detectors to be aligned with respect to each other.

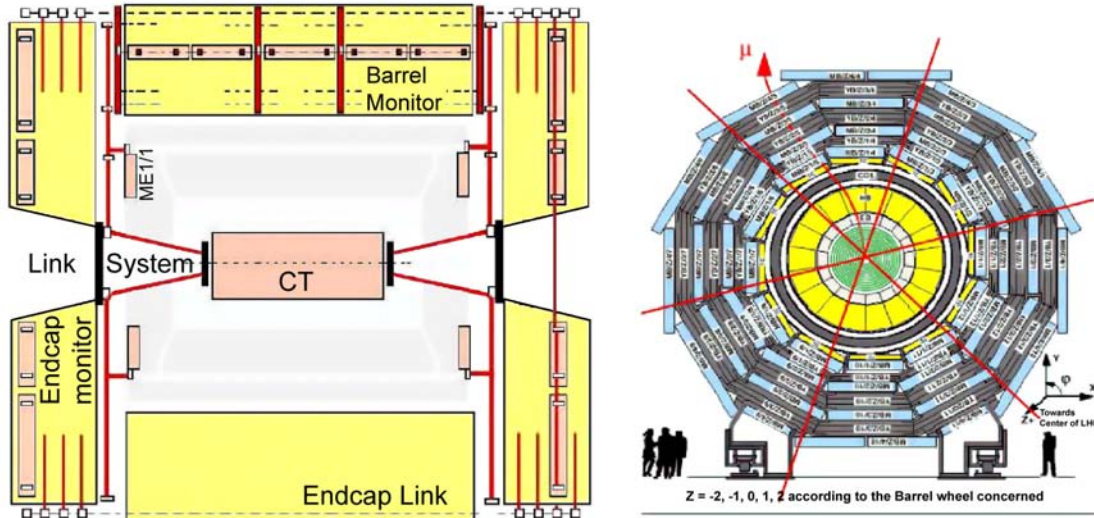


Figure 214: Schematic view of the alignment system. Left panel: longitudinal view of CMS. The continuous and dotted lines show different optical light paths. Right panel: transverse view of the barrel muon detector. The crossing lines indicate the r - z alignment planes with 60° staggering in ϕ .

8.4.1 System layout and calibration procedures

The optical network uses two types of light sources: LEDs and laser beams. It is composed of 10 000 LEDs and 150 laser beams together with precise measuring devices: ≈ 900 photo-detectors and ≈ 600 analog sensors (distance sensors and inclinometers), complemented by temperature, humidity and Hall probes. The system is structured into three basic blocks whose main features are described below.

Muon barrel alignment The monitoring of the barrel muon detector (Fig. 215) is based on the measurement of all the 250 DT chamber positions with respect to a floating network of 36 rigid reference structures, called MABs (Module for the Alignment of Barrel). The MAB design was optimised to achieve adequate mechanical rigidity of the structures under load and in thermal and humidity gradients. Long term measurements showed deviations below $100 \mu\text{m}$ and $50 \mu\text{rad}$ [176]. The MABs are fixed to the barrel yoke forming 12 r - z planes parallel to the beam line and distributed in ϕ every 60° . Each structure contains 8 specially designed video cameras that observe LED sources mounted on the DT chambers. Extra light sources and video-cameras in specific MABs serve to connect MABs in different planes forming a closed optical network (called diagonal connections). The MAB positions in the z coordinate are measured with respect to 6 calibrated carbon-fibre bars (z -bars) sitting on the outer surface of the vacuum tank of the solenoid. The MABs in the external wheels, $YB \pm 2$, are equipped with extra alignment sensors and light sources that connect the barrel monitoring system with the endcap and tracker detectors.

The 4 corners of the DTs are equipped with LED light sources. Four LED-holders, or forks, are rigidly mounted on the side-profile of the honeycomb structure (2 per side) and use the rectangular $50 \times 65 \text{ mm}^2$ tube as a light passage. Each fork contains 10 LEDs, 6 and 4, respectively, on each side. There are 10 000 light sources mounted on the DT chambers. The position of the forks with respect to the chamber geometry was measured on a dedicated bench with a precision of $< 70 \mu\text{m}$. As an important by-product, the calibration also provides the full geometry, including the planarity, trapezoidity, and the relative positions of superlayers for each DT chamber with $< 40 \mu\text{m}$ precision, as described in Sect. 8.1.4. Each LED-holder and video-sensor was individually calibrated before its assembly on the DT chambers or MABs and z -bars. LED-holders were measured and the position of the light centroid was determined with respect to the holder mechanics with an accuracy of $10 \mu\text{m}$. Long term measurements showed good stability of the centroids and light intensity distributions. CMOS miniature video sensors, containing

384×288 pixels with $12 \times 12 \mu\text{m}^2$ pixel size, were calibrated to absorb residual response non-uniformities and the intensity nonlinearities. The video cameras, consisting of a video sensor and a single-element lens assembled in an aluminium box, were also calibrated to determine their inner geometrical parameters. Fully instrumented MABs containing the necessary number of survey fiducials were calibrated on a special bench, where the whole geometry of the structure, positions, and orientations of elements were determined with overall accuracies of $70 \mu\text{m}$ and $50 \mu\text{rad}$.

Once MABs were installed (Fig. 216), the initial MAB positions on the barrel wheels were determined by photogrammetry measurements.

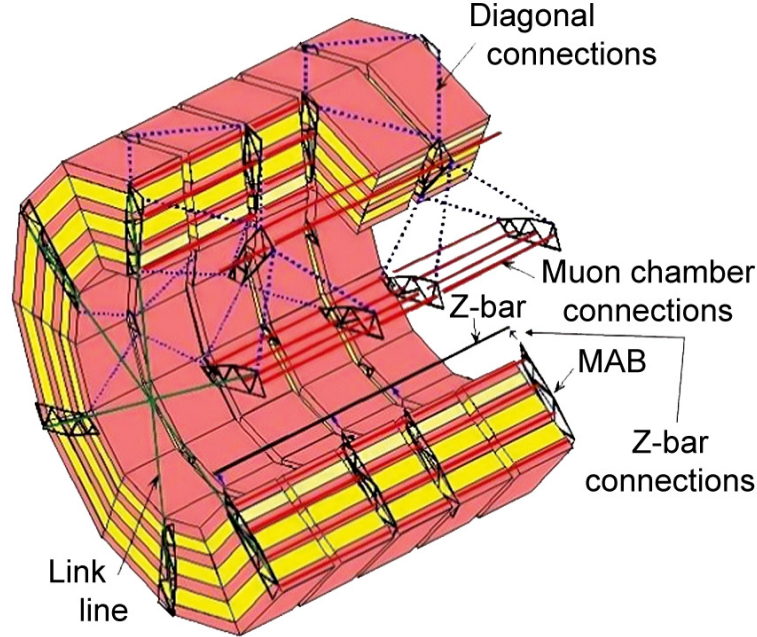


Figure 215: Schematic view of the barrel monitoring system showing the optical network among the MAB structures.

The control, read-out, and data preprocessing [177] are performed by a network of local minicomputers (1 per MAB, 36 in total) that makes it possible to run the full system in parallel. The minicomputers are connected to the main control PC via an Ethernet network capable of working in magnetic fields. The main control PC synchronises the operation of the light sources mounted on the DT chambers and the read-out of the images taken by the cameras. The minicomputers control the light sources mounted on the MABs and the z -bars, read out the temperature and humidity sensors, perform the image read-out and digitisation, and calculate the image centroids of the light sources. The results are transferred to the main control PC, which is connected to the corresponding central CMS units.

Based on simulation, the barrel monitoring system should provide a stand-alone measurement of the barrel chambers with an average r - ϕ position accuracy of $100 \mu\text{m}$ for chambers in the same sector and about $250 \mu\text{m}$ between barrel sectors. The current understanding of its performance is discussed in Sect. 8.4.3.

Muon endcap alignment The muon endcap alignment system [178] is designed to continuously and accurately monitor the actual positions of the 486 CSCs relative to each other, relative to the tracking system, and ultimately within the absolute coordinates of CMS. Due to the large magnetic field, the chambers mounted on the endcap yoke undergo substantial motion and deformation, on the order of a few centimetres, when the field is switched on and off. The alignment system must measure the disk deformation and monitor the absolute positions of the CSCs in the r - ϕ plane and in z . From simulations, the required absolute alignment accuracies were found to run from 75 to $200 \mu\text{m}$ in r - ϕ . Because the r and r - ϕ accuracies are directly coupled, the required accuracy in the r -position is $\approx 400 \mu\text{m}$. The z displacement due to the deformation of the iron yoke disks caused by the strong and non-uniform magnetic field in the endcaps requires the alignment sensors to be able to accommodate a dynamic range of $\approx 2 \text{ cm}$.

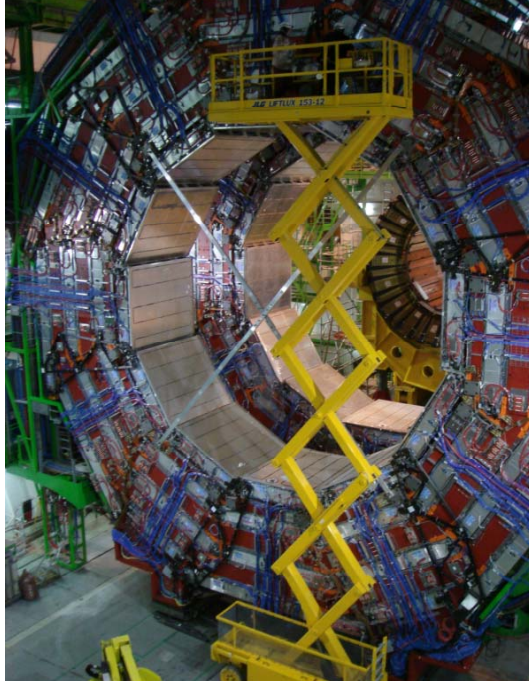


Figure 216: Installation of the MABs on wheel YB+2.

with an accuracy of ≈ 1 mm.

The system uses a complex arrangement of 5 types of sensors for the transferring and monitoring of ϕ , r , and z coordinates (Fig. 217). The system measures one sixth of all endcap chambers. The main monitoring tools within the r - ϕ plane are the Straight Line Monitors (SLM). Each SLM consist of 2 cross-hair lasers, which emit a nearly radial laser beam across 4 chambers from each end, and provide straight reference lines that are picked up by 2 optical sensors (Digital CCD Optical Position Sensors, DCOPS [179]). This arrangement provides references for the chamber positions relative to the laser lines. Figure 218 shows a photograph of a complete SLM on station ME+2. The figure also indicates r -sensors for monitoring radial chamber positions, z -sensors for axial distance measurements between stations, and a clinometer for monitoring the tilt of the mechanical support assembly (transfer plate) onto which lasers, reference DCOPS, and z -sensors are mounted. The inset in Fig. 218 shows the location of proximity sensors on the outer ring of the ME+1 station, which monitor the azimuthal distance between neighbouring chambers. These are necessary because the outer ring of ME1 chambers is the only ring for which the CSCs do not overlap in ϕ . Furthermore, every CSC and alignment device is equipped with photogrammetry targets to allow absolute magnet-off measurements.

The ϕ coordinate alignment is handled by optical SLMs and transfer lines. Transfer laser lines run parallel to the CMS z -axis along the outer cylindrical envelope of CMS at 6 points separated by 60° in ϕ . The SLMs run across the surface of one sixth of all the CSCs, along radial directions, and link transfer lines on opposite sides of a disk. Both laser lines have a similar basic configuration: a laser beam defines a direction in space that is picked up by several DCOPS precisely mounted on CSCs or transfer plates to reference their own positions. Mounting accuracies due to tolerances of dowel pins and dowel holes are ≈ 50 μm . Every DCOPS comprises 4 linear CCDs, each with 2048 pixels and 14 μm pixel pitch. The CCDs are basically arranged in the shape of a square and can be illuminated by cross-hair lasers from either side. The r and z coordinate measurements are performed by analog linear potentiometers and optical distance devices in contact with aluminium tubes of calibrated length.

All analog sensors were calibrated with a 1D precision linear mover with 6.4 μm step size. The uncertainty in the absolute distance calibration is 100 μm for r sensors and 53 μm for z sensors [180]. Calibration for optical DCOPS consisted in determining the distance from the surface of the mount hole for a reference dowel pin to the first active CCD pixel and measuring the projected pixel pitch of each of the 4 CCDs. This was done on a calibration bench where a fibre-bundle variable light source at the focus of a parabolic mirror illuminated a mask with 8 optical slits. A simple geometry reconstruction, based on coordinate-measuring-machine data for the calibration mask and sensor mounts, determined the physical pixel positions. Calibration errors were typically 30 to 50 μm .

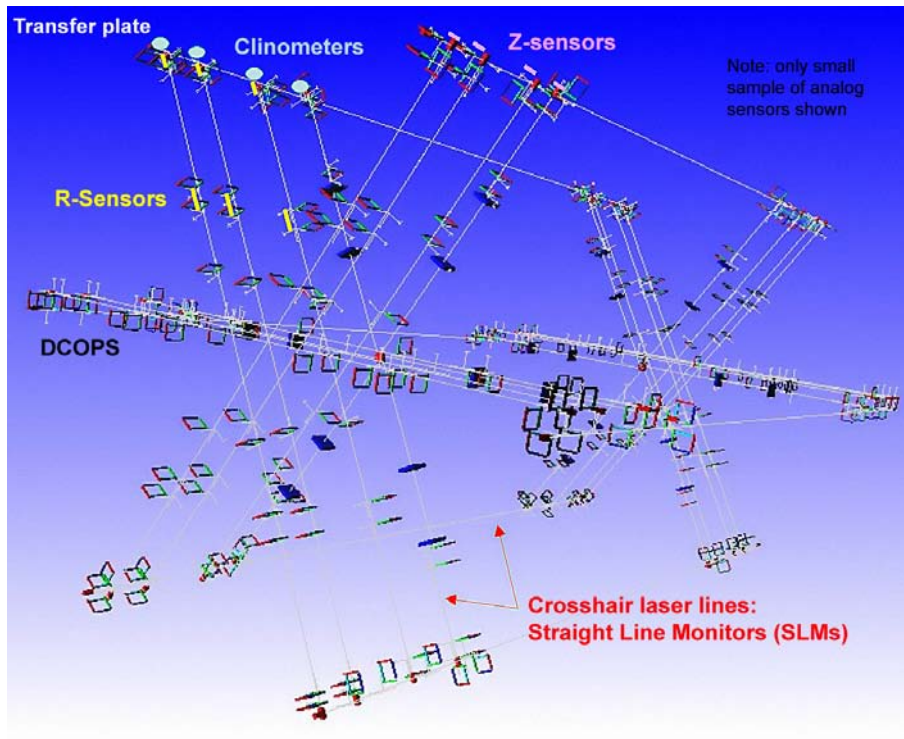


Figure 217: Visualisation of the geometry and components of the muon endcap alignment system. The square objects represent optical sensors (DCOPS) for monitoring 3 straight laser lines across each endcap station.

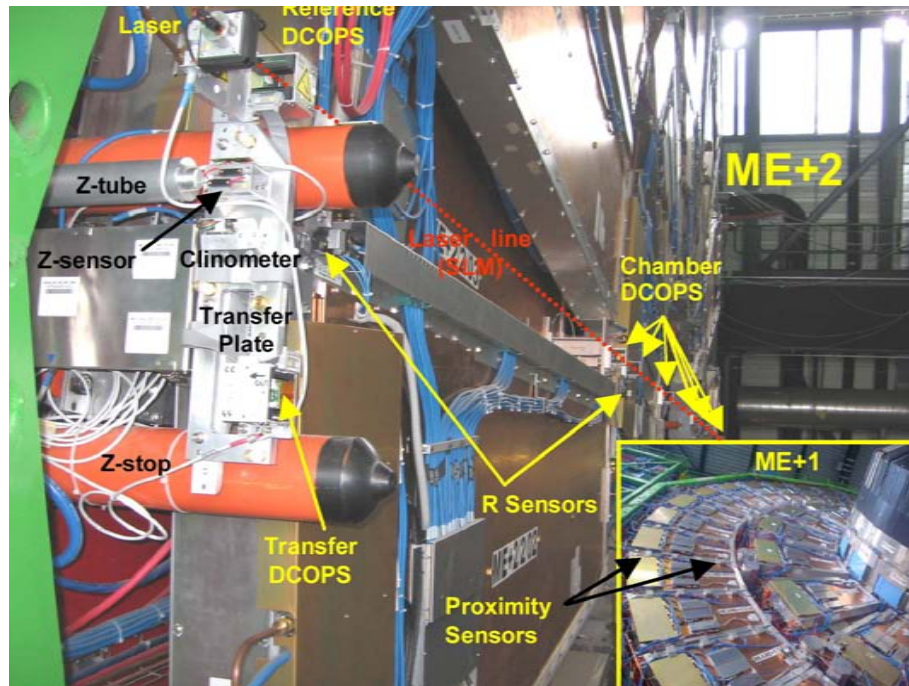


Figure 218: Close-up of one of the 3 Straight Line Monitors (SLM) on the ME+2 station with cross-hair laser, DCOPS, and analog sensors (r , z , and Tiltmeter). The insert indicates the location of proximity sensors on ME+1.

Link system The purpose of the link alignment system is to measure the relative position of the muon spectrometer and the tracker in a common CMS coordinate system. It is designed to work in a challenging environment of very high radiation and magnetic field, meet tight space constraints, and provide high precision measurements over long distances. A distributed network of opto-electronic position sensors, ASPDs (amorphous-silicon position detectors) [181], placed around the muon spectrometer and tracker volumes are connected by laser lines. The entire system is divided into 3 ϕ -planes 60° apart; this segmentation allows a direct reference of each muon barrel sector with the tracker detector and provides a direct reference also to the endcap alignment lines in the first endcap station, ME1. Each plane consists of 4 quadrants (Fig. 214) resulting in 12 laser paths: 6 on each z -side of the CMS detector, and generated by 36 laser sources. The system uses 3 types of reference structures: rigid carbon-fibre annular structures placed at both ends of the tracker (alignment rings, AR) and at the $YE\pm 1$ wheels of the endcap muon spectrometer (link disks, LD); and the MAB structures attached at the external barrel wheels, $YB\pm 2$. Figure 219 (left) shows the LD and AR carbon fibre structures installed in the inner $\eta = 3$ cone. The link measurement network is complemented by electrolytic tiltmeters, proximity sensors in contact with aluminium tubes of calibrated length, magnetic probes, and temperature sensors.

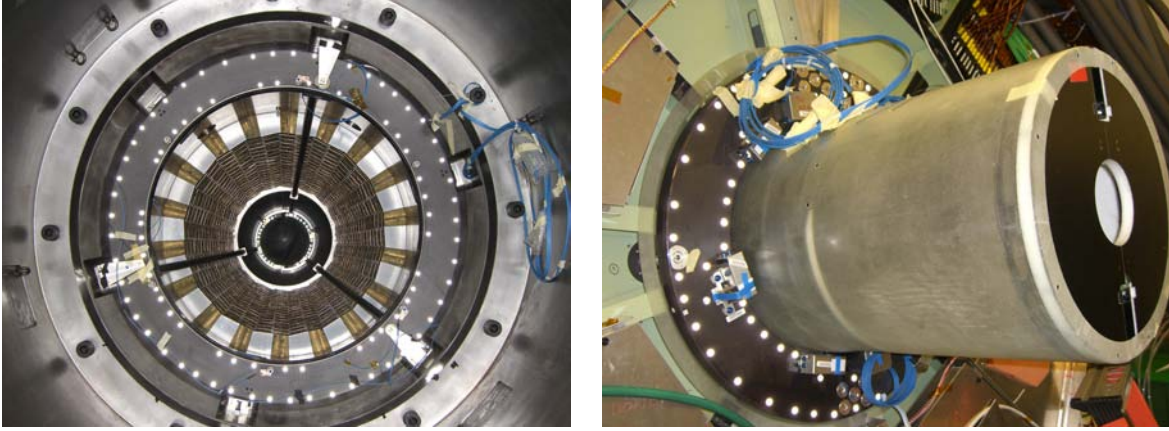


Figure 219: Left panel: Link Disk and Alignment Ring installed in the inner $\eta = 3$ cone during the first closing of the detector in summer 2006, MTCC period. Right panel: Alignment Ring mounted in the TEC end flange.

The ARs are rigidly attached to the endcap tracker detectors, TECs, through a purely mechanical connection with the instrumented silicon volume (Sect. 4.3.7). Three pillars, acting as support fixations, connect the last instrumented disk of each TEC with the corresponding AR, at both ends of the tracker volume, see Fig. 219 (right). The position and orientation of the ARs with respect to the TEC disks 9 and 10 were measured with a coordinate-measurement machine using the external survey fiducials prior to the TEC assembly and instrumentation. Changes in angular orientations are monitored by high precision tiltmeters placed at the AR and TEC disk 10. Laser sources originating at the AR and running along the inner detector boundary reach ASPD sensors on the first endcap disk, ME1, and on the external barrel wheel.

The ASPDs are 2D semitransparent photo-sensors, which consist of 2 groups of 64 silicon micro-strips with a pitch of 430 μm oriented perpendicularly. With $\geq 80\%$ transmittance for the 685 nm wavelength used in the system, they allow multi-point measurements along the light path without significant distortions in the beam direction. The intrinsic position resolution is about 2 μm . The location, centre position, and orientation of the ASPD with respect to reference pins in their mechanical mount are measured with a non-contact CMM (Coordinate Measuring Machine) with an overall accuracy of 15 μm . Distance measurement devices (optical distance sensors and linear potentiometers) already mounted in their final mechanics were calibrated using 2 μm resolution linear movers and pre-measured calibration fixtures. The uncertainties in absolute and relative calibration [182] are below 50 μm and 20 μm , respectively, for the different sensor types. The intrinsic accuracy of the tiltmeters sensors, after calibration, is about 2 μrad ; mechanical offsets inherent to the mechanical mounts and assembly tolerances are determined by survey and photogrammetry techniques.

The light sources (collimators) and specific optical devices housed on the alignment reference structures (AR, LD, and MABs) create the laser beam paths with the layout described above. Each collimator is focused to its working distance to ensure Gaussian beam profiles along the propagation path to avoid beam-shape-induced bias in the position reconstruction. The adjustment and calibration [183] of the laser rays for the AR and LD structures were done on a dedicated bench instrumented with a precise survey network that mimics the nominal detector geometry.

Beams were adjusted to their nominal geometry with a precision better than $100\ \mu\text{rad}$. Long term measurements were performed after beam adjustments. Beam pointing stability, including temperature effects, was found to be better than $30\ \mu\text{rad}$. The adjustment and calibration accuracy was limited to $30\text{--}100\ \mu\text{rad}$ due to the finite dimension of the structures combined with the intrinsic accuracy of the survey and photogrammetry measurement techniques of $70\ \mu\text{m}$.

Survey and photogrammetry measurements are also performed during the installation of the alignment structures in the detector. An installation accuracy of the order of a few millimetres and milliradians is needed to ensure correct functionality of the system, taking into account the standard CMS assembly tolerances of the big endcap disks and barrel wheels.

The control, read-out, and data preprocessing are performed by two types of electronic boards. Analog sensors read-out and laser control use standard ELMB (embedded local monitor board) cards [184]. For the read-out of ASPD sensors, custom made LEB (local electronic board) cards were developed. LEBs are intelligent imaging acquisition boards made to read and control up to 4 ASPD sensors. They are based on Hitachi micro-controllers. ELMB and LEB boards use the CAN communication protocol to connect the front-end electronics and the main control PC unit.

8.4.2 Geometry reconstruction

The DAQ, monitoring, and control software are integrated into the DCS (detector control system) environment. Data are recorded in an online Oracle database and subsequently converted into ntuples by specialised programs that perform database queries and apply calibrations to convert raw values into meaningful physical quantities. This provides the necessary information for global geometry reconstruction, which is handled by COCOA (CMS Object-Oriented Code for Optical Alignment) [185], an offline program to simulate and reconstruct the complex optical alignment system. Due to the unknown movements of different CMS structures, the sensors of the optical alignment systems will not measure the expected nominal values. The aim of COCOA is to analyse the observed changes in these measurements to determine which are the displacements and/or rotations that caused them. The approach adopted by COCOA to tackle this problem is to solve the system of equations that relate the measurement values to all the positions, rotations, and internal parameters of all the objects that make up the system. In fact, to solve the system of equations, one does not need to know the explicit form of the equations, but only the derivatives of each measurement value with respect to each object parameter. COCOA uses a geometrical approximation of the propagation of light to calculate numerically these derivatives and then solves the system of equations through a nonlinear least squares method. Due to the large number of parameters in CMS (about 30 000), big matrices are needed. COCOA matrix manipulations are based on the Meschach Library [186].

COCOA has proved its robustness through its extensive use in CMS for several design studies, as well as for the analysis of several test benches and magnet test results. Its output, the aligned geometry, will be used as input geometry for track reconstruction as well as for further alignment studies based on muon tracks from cosmic rays and from pp collisions.

8.4.3 System commissioning and operating performance

A first test of the large superconducting solenoid magnet in the CMS detector was successfully performed between June and November 2006, during which stable operation at full field (4 T) was achieved (Sect. 3.3). The alignment sensors, read-out, and DAQ software were commissioned during this test period for about one third of the system, instrumented at the $+z$ side of the detector. This allowed the first full-scale dynamic test of the system. The performance of the system as well as the main features of the yoke displacement and deformation were studied. The relevant results are summarised below:

- Measurement of relative movements due to thermal changes.
The effects of thermal changes (day-night variations) for DT and CSC chambers were recorded for the conditions present during the test, with the detector in the surface assembly hall and power on only $\approx 5\%$ of the muon spectrometer. The measured relative movement did not exceed $50\ \mu\text{m}$ over the entire test period, with changes in position showing a good correlation with temperature. Although a movement of this magnitude is not relevant from the physics analysis point of view, its measurement illustrates the good resolution of the alignment system.
- Measurement of the displacements and deformations of the yoke structures.
Two effects were observed. The first is the change in the original positions of the structures (the positions

before any magnet operation). The displacements of the structures along the z direction towards the solenoid seem to stabilise after the first 2.5–3 T are reached. This compression is permanent, meaning it is not reversed/recovered in subsequent magnet off states, and it is interpreted as the final closing of the structures due to the magnetic forces acting on the iron. These measured displacements are specific to the first CMS closing experience and cannot be extrapolated to other scenarios.

The second effect is the almost perfectly elastic deformations between magnet-on and magnet-off states, as illustrated in Fig. 220. At 4 T, the elastic deformation of the barrel yoke, measured at the end of the $+z$ side with respect of the plane of the interaction point, is about 2.5 mm. Figure 220 shows the elastic compression of the barrel wheels versus the magnet current as recorded in the second phase of the Magnet Test period. Despite the large overall compression of the barrel spectrometer, an important measurement was the stability of the barrel chambers during the whole data-taking period. The relative movements in the r - ϕ direction did not exceed 60 μm .

The behaviour of the endcap disk is more complicated. Due to the strong gradient in the magnetic field near the end of the solenoid, strong magnetic forces pull the central portions of the endcap disks towards the center of the detector. As shown in Fig. 220, the nose is pulled towards the interaction point, the magnitude of the compression is perfectly correlated with the magnet current, reaching up to ≈ 16 mm at 4 T. The various z -stops, which prevent the disks from getting pushed into each other and onto the barrel wheels, cause the endcap disks to bend into a cone shape. The z -stops between endcap and barrel, positioned at nearly half the disk radius, cause the side of the YE1 disk facing the barrel to compress radially around them by ≈ 600 μm , while expanding azimuthally by ≈ 800 μm . This explains the radial compression of the face of ME+1 and the larger bending angles at mid-radius than at the outer edge (Fig. 221). Endcap disk deformations are predicted by finite element analysis (FEA) using the ANSYS program [3]. The measurements are in reasonable quantitative agreement for all displacements and deformations, as shown in Fig. 221. Note that the front z -stops, between the ME1 and barrel wheels, were not included in the FEA, which explains the difference. The difference between top and bottom is also explained by the presence of the carts that support the disks.

The rest of the endcap stations on YE+2 and YE+3 experience a maximum bending angle of ≈ 2.5 mrad relative to the vertical, as sketched in Fig. 222. As in the case of the barrel chambers, with stable 4 T field, the observed relative movements were very small.

- Detector closing tolerances and reproducibility.

The test of the magnet was divided into 2 phases, separated by a short period during which the yoke was open to extract the inner detectors, tracker, and ECAL modules. This allowed a test of the reproducibility in the closing procedure and tolerances, as well as the study of the compatibility of measurements between the two phases. Reproducibility in the closing was at the level of a few millimetres for the barrel wheels and about an order of magnitude higher for the endcap disks. The particular conditions of the test did not allow the establishment of a solid understanding of reproducibility for the process of closing the different structures. Instead, the system was able to reproduce the same magnetic-force-induced effects as measured in the first period.

From this test we conclude that the system operates adequately under magnetic fields both in terms of dynamic range and measurement performance. The system precision achieved is ≤ 300 μm and the measurement accuracy has been validated against results from photogrammetry and cosmic ray tracks.

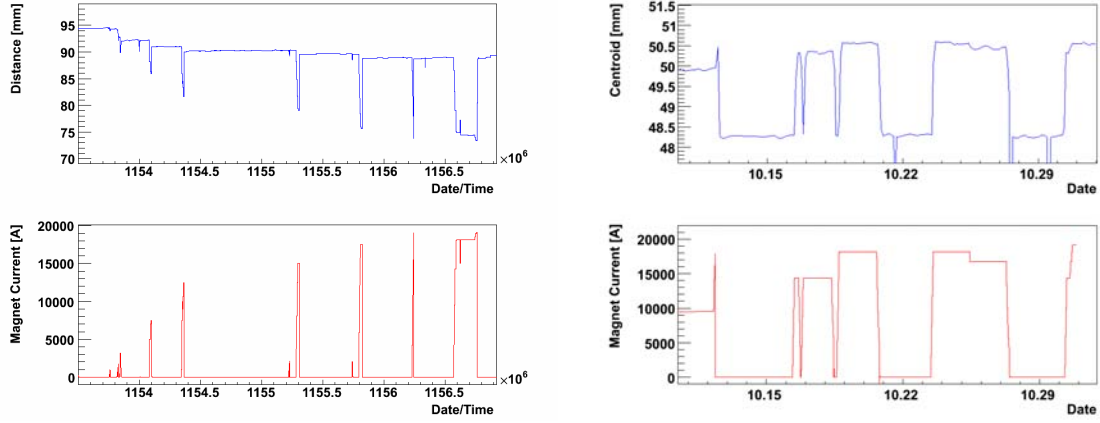


Figure 220: Deformations of endcap disks and barrel wheels vs magnet current cycling. Left panel: The bottom plot shows the magnet powering cycle exercised during the first phase of the Magnet Test period. The top plot shows the measured YE+1 nose compression towards the interaction point. Right panel: The bottom plot shows the magnet powering cycle exercised during the second phase of the Magnet Test period. The top plot shows the calculated approximate YB+2 compression towards the interaction point.

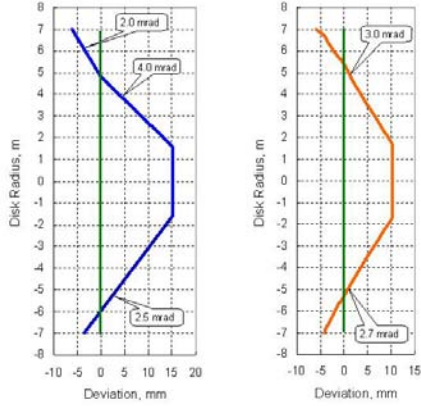


Figure 221: Comparison of the YE+1 disk deformations in the r - z plane at full magnetic field (4 T) measured by the alignment system (left panel) and predictions from finite element analysis (right panel). The vertical lines correspond to 0 magnetic field.

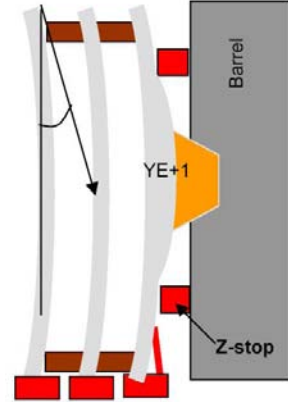


Figure 222: Current understanding of disk deformation due to magnetic forces based on alignment system measurements. The z -stops (red) prevent the disks from getting pushed into each other. Note that the indicated bending angle is exaggerated for illustrative purposes. Its measured magnitude is 2.5 mrad.

9 Trigger

The LHC provides proton-proton and heavy-ion collisions at high interaction rates. For protons the beam crossing interval is 25 ns, corresponding to a crossing frequency of 40 MHz. Depending on luminosity, several collisions occur at each crossing of the proton bunches (approximately 20 simultaneous pp collisions at the nominal design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$). Since it is impossible to store and process the large amount of data associated with the resulting high number of events, a drastic rate reduction has to be achieved. This task is performed by the trigger system, which is the start of the physics event selection process. The rate is reduced in two steps called Level-1 Trigger (L1T) [187] and High-Level Trigger (HLT) [188], respectively. The Level-1 (L1) Trigger consists of custom-designed, largely programmable electronics, whereas the HLT is a software system implemented in a filter farm of about one thousand commercial processors. The rate reduction capability is designed to be at least a factor of 10^6 for the combined L1T and HLT. The design output rate limit of the L1T is 100 kHz, which translates in practice to a calculated maximal output rate of 30 kHz, assuming an approximate safety factor of three. The L1T uses coarsely segmented data from the calorimeters and the muon system, while holding the high-resolution data in pipelined memories in the front-end electronics. The HLT has access to the complete read-out data and can therefore perform complex calculations similar to those made in the the analysis off-line software if required for specially interesting events. Since HLT algorithms will evolve with time and experience they are not described here. More information may be found in [189]. For reasons of flexibility the L1T hardware is implemented in FPGA technology where possible, but ASICs and programmable memory lookup tables (LUT) are also widely used where speed, density and radiation resistance requirements are important. A software system, the Trigger Supervisor [190], controls the configuration and operation of the trigger components.

The L1T has local, regional and global components. At the bottom end, the Local Triggers, also called Trigger Primitive Generators (TPG), are based on energy deposits in calorimeter trigger towers and track segments or hit patterns in muon chambers, respectively. Regional Triggers combine their information and use pattern logic to determine ranked and sorted trigger objects such as electron or muon candidates in limited spatial regions. The rank is determined as a function of energy or momentum and quality, which reflects the level of confidence attributed to the L1 parameter measurements, based on detailed knowledge of the detectors and trigger electronics and on the amount of information available. The Global Calorimeter and Global Muon Triggers determine the highest-rank calorimeter and muon objects across the entire experiment and transfer them to the Global Trigger, the top entity of the Level-1 hierarchy. The latter takes the decision to reject an event or to accept it for further evaluation by the HLT. The decision is based on algorithm calculations and on the readiness of the sub-detectors and the DAQ, which is determined by the Trigger Control System (TCS). The Level-1 Accept (L1A) decision is communicated to the sub-detectors through the Timing, Trigger and Control (TTC) system. The architecture of the L1T is depicted in Fig. 223. The L1T has to analyze every bunch crossing. The allowed L1T latency, between a

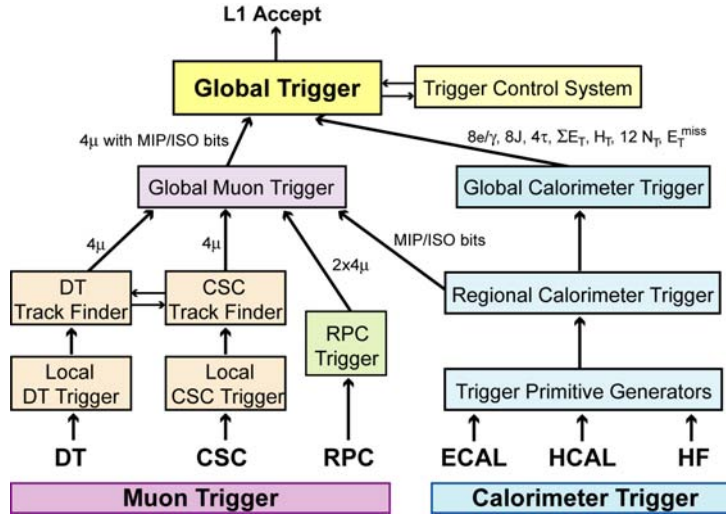


Figure 223: Architecture of the Level-1 Trigger.

given bunch crossing and the distribution of the trigger decision to the detector front-end electronics, is $3.2 \mu\text{s}$. The processing must therefore be pipelined in order to enable a quasi-deadtime-free operation. The L1T electronics is housed partly on the detectors, partly in the underground control room located at a distance of approximately 90 m from the experimental cavern.

9.1 Calorimeter trigger

The Trigger Primitive Generators (TPG) make up the first or local step of the Calorimeter Trigger pipeline. For triggering purposes the calorimeters are subdivided in trigger towers. The TPGs sum the transverse energies measured in ECAL crystals or HCAL read-out towers to obtain the trigger tower E_T and attach the correct bunch crossing number. In the region up to $|\eta| = 1.74$ each trigger tower has an (η, ϕ) -coverage of 0.087×0.087 . Beyond that boundary the towers are larger. The TPG electronics is integrated with the calorimeter read-out. The TPGs are transmitted through high-speed serial links to the Regional Calorimeter Trigger, which determines regional candidate electrons/photons, transverse energy sums, τ -veto bits and information relevant for muons in the form of minimum-ionizing particle (MIP) and isolation (ISO) bits. The Global Calorimeter Trigger determines the highest-rank calorimeter trigger objects across the entire detector.

Calorimeter trigger primitive generators

The ECAL on-detector front-end electronics boards, each serving 25 crystals, receive the ADC signals from the very front-end electronics located at the rear of the detector modules. They contain most of the TPG pipeline in six radiation-hard $0.25 \mu\text{m}$ CMOS ASIC chips named FENIX. An off-detector Trigger Concentrator Card (TCC) collects the primitives from 68 front-end boards in the barrel and 48 boards in the endcaps through optical links. The TCCs finalize the TPG generation and encoding, store the trigger primitives during the L1 latency time and transmit them to the RCT by dedicated daughter boards, the Synchronization and Link Boards (SLB), upon reception of a L1A signal. The SLBs synchronize the trigger data through circuits that histogram the LHC bunch crossing structure. Each trigger tower is aligned with the bunch crossing zero signal. A Data Concentrator Card (DCC) performs the opto-electronic conversion and deserialization of the serial input data streams and sends the read-out data collected from the front-end boards to the DAQ. Clock and Control System (CCS) boards distribute the clock, the L1A and control signals to the TCC, the DCC and the on-detector electronics. The ECAL TPG hardware is contained in twelve 9U VME crates for the barrel and six for the endcaps.

The front-end modules of the hadron calorimeter contain Charge Integrator and Encoder (QIE) ADC chips to digitize the signals from the photo detectors. Optical links transmit the data to the HCAL Trigger and Readout (HTR) boards. Each HTR board processes 48 channels. It linearizes, filters and converts the input data to generate the HCAL trigger primitives. The energy values of front and back towers are added and the bunch crossing number is assigned by a peak filtering algorithm. As for the ECAL, the primitives are sent to the RCT by SLBs, and read-out data are collected by a DCC. The HCAL trigger electronics is contained in 26 9U VME crates. Each crate houses 18 HTR boards and one DCC.

Regional Calorimeter Trigger

The Regional Calorimeter Trigger [191] determines electron/photon candidates and transverse energy sums per calorimeter region. Information relevant for muons about isolation and compatibility with minimally ionizing particles is also calculated. A region consists of 4×4 trigger towers except in HF where a region is one trigger tower. Electromagnetic and hadronic transverse energies are summed in each tower.

The e/γ trigger algorithm (Fig. 224) starts by determining the tower with the largest energy deposit and is applied across the entire ECAL region. The energy of the tower with the next-highest deposit in one of the four broad side neighbours is then added. Isolated and non-isolated e/γ within $|\eta| \leq 2.5$ are determined by the trigger. A non-isolated e/γ requires passing of two shower profile vetoes. The first one is based on a fine-grain crystal energy profile reflecting the lateral extension of a shower. The fine-grain bit is set by the TPG if the shower is contained in a matrix of 2×5 crystals. The matrix is dimensioned such that it also allows for the detection of bremsstrahlung due to the magnetic field. The second one is based on the ratio of the deposited energies in the hadronic and in the electromagnetic sections. A typical maximal value of 5% is allowed for that ratio. An isolated electron/photon candidate has to pass the previous vetoes for all eight neighbouring towers. In addition, at least one quiet corner made of four groups of five electromagnetic towers surrounding the hit tower is required. Four isolated and four non-isolated e/γ per region are forwarded to the GCT.

The RCT also sums the transverse energy in a given region of the central calorimeter (HF is not included) and determines τ -veto bits for the identification of jets from one- and three-prong τ -decays, which are narrower than ordinary quark/gluon jets. A τ -veto bit is set unless the pattern of active towers corresponds to at most 2×2 contiguous trigger towers within a 4×4 tower region. Jets can be classified as τ -jet only at $|\eta| < 3.0$ (not in HF).

The RCT hardware consists of 18 regional 9U VME crates and one 6U clock distribution crate located in the underground control room. Each crate covers a region of $\Delta\eta \times \Delta\phi = 5.0 \times 0.7$. Receiver cards are plugged

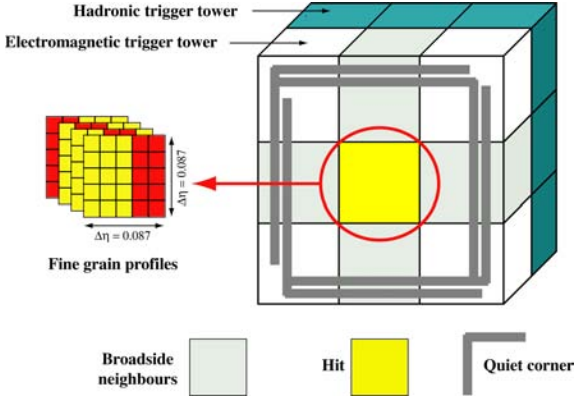


Figure 224: Electron/photon algorithm.



Figure 225: Electron Isolation Card.

into the rear of the regional crates. Seven cards per crate receive the ECAL and HCAL primitives. The HF primitives are directly received on a Jet/Summary card. The serial input data are converted to 120 MHz parallel data, deskewed, linearized and summed before transmission on a 160 MHz custom monolithic backplane to seven Electron Isolation Cards (EIC) and one Jet/Summary Card (JSC) mounted at the front side of the crate. Different ASICs perform the algorithm calculations. An EIC is shown in Fig. 225.

Global Calorimeter Trigger

The Global Calorimeter Trigger determines jets, the total transverse energy, the missing transverse energy, jet counts, and H_T (the scalar transverse energy sum of all jets above a programmable threshold). It also provides the highest-rank isolated and non-isolated e/γ candidates.

Jets are found by a four-stage clustering technique based on jet finders operating in 2×12 cells in ϕ and η , spanning 40° and half the detector, respectively, in these directions. The cell at $\eta=0$ is duplicated. In the first stage mini-clusters are created by summing energies within 2×3 cells if a central cell has more energy than neighbouring cells. In the second stage the three largest mini-clusters in each of the two ϕ -strips are transferred in opposite ϕ -directions. These are compared against the existing mini-clusters on the receiving ϕ -strip. Mini-clusters adjacent or diagonally adjacent to a larger mini-cluster are removed. In the third and fourth stages the received mini-clusters that survive have their three adjacent cells in the receiving ϕ -strip combined to make a 3×3 cell. A jet is classified as a τ jet if none of the corresponding RCT regions had a τ -veto bit set. After sorting, up to four jets and four τ jets from the central HCAL and four jets from HF are forwarded to the GT. The magnitude and direction of the missing energy and the total transverse energy are computed from the regional transverse energy sums in the two coordinates transverse to the beam within $|\eta| < 5$. Twelve jet counts for different programmable E_T -thresholds and optionally also different (η, ϕ) -regions are computed. Muon MIP/ISO bits are received from the RCT along with the e/γ data and are forwarded to the GMT through a dedicated muon processing system. Apart from triggering, the GCT also acts as the read-out system for the RCT. The GCT has in addition the capability to monitor rates for certain trigger algorithms and from those deduce information about the LHC luminosity as seen by the CMS trigger system.

All GCT electronics is located in the underground control room. The large amount of data from the RCT crates are transmitted electronically to 63 Source Cards, which reorder the data onto 252 optical fibres. The core of the GCT processing is performed by Leaf Cards, which can be configured as electron or jet cards. Several Leaf Cards can be connected with each other in order to perform complex tasks such as the jet finding. There are two Leaf Cards for electrons and six for jets. Each electron leaf card receives the e/γ data from one half of the RCT crates on 27 fibres and sorts them. Each jet card receives 30 regional sum fibres from three RCT crates via the source cards. They perform the jet clustering and transmit the jet candidates to two Wheel Cards for sorting and data compression. They also calculate partial energy sums and jet counts and forward them to the Wheel Cards. A Concentrator Card finally collects the data from all Electron Leaf and Wheel Cards and performs the final sorting for electrons/photons, completes the jet finding in the boundaries between groups of three Leaf Cards, sorts all jets, calculates the global energy and jet count quantities and sends the final results to the GT and the DAQ. In addition to the tasks involving e/γ 's, jets and energy sums, the GCT also handles MIP/ISO bits for muons. They are processed

by three muon processing cards, which receive 6 muon fibres each from Source Cards. The processor design is built on an evolution of the leaf concept and uses a modular, low-latency architecture based on the μ TCA industry standard [193]. An active custom backplane based on the principle of a crosspoint switch allows a programmable routing of the 504 MIP/ISO bits, which are then transmitted to the GMT on 24 links.

9.2 Muon trigger

All three muon systems – the DT, the CSC and the RPC – take part in the trigger. The barrel DT chambers provide local trigger information in the form of track segments in the ϕ -projection and hit patterns in the η -projection. The endcap CSCs deliver 3-dimensional track segments. All chamber types also identify the bunch crossing from which an event originated. The Regional Muon Trigger consists of the DT and CSC Track Finders, which join segments to complete tracks and assign physical parameters to them. In addition, the RPC trigger chambers, which have excellent timing resolution, deliver their own track candidates based on regional hit patterns. The Global Muon Trigger then combines the information from the three sub-detectors, achieving an improved momentum resolution and efficiency compared to the stand-alone systems. The initial rapidity coverage of the muon trigger is $|\eta| \leq 2.1$ at the startup of LHC. The design coverage is $|\eta| \leq 2.4$.

Drift Tube local trigger

The electronics of the DT local trigger consists of four basic components (Fig. 226): Bunch and Track Identifiers (BTI), Track Correlators (TRACO), Trigger Servers (TS) and Sector Collectors (SC). While the SCs are placed on the sides of the experimental cavern, all other trigger and read-out electronics is housed in minicrates on the front side of each chamber. All devices are implemented in custom-built integrated circuits. The BTIs are interfaced

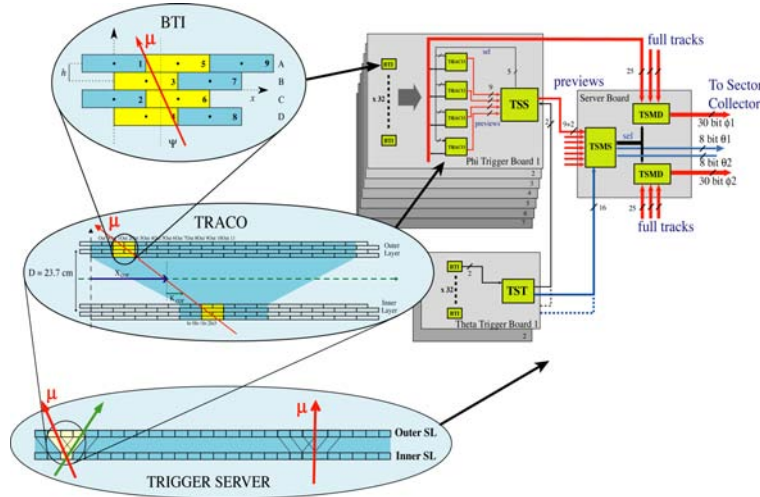


Figure 226: Drift Tube Local Trigger.

to the front-end electronics of the chambers. Using the signals from the wires they generate a trigger at a fixed time after the passage of the muon. Each BTI searches for coincident, aligned hits in the four equidistant planes of staggered drift tubes in each chamber superlayer. The association of hits is based on a meantimer technique [194], which uses the fact that there is a fixed relation between the drift times of any three adjacent planes. From the associated hits, track segments defined by position and angular direction are determined. The spatial resolution of one BTI is better than 1.4 mm, the angular resolution better than 60 mrad. The BTI algorithm is implemented in a 64-pin ASIC with CMOS 0.5 μ m Standard Cell technology. There are a few hundred BTIs per chamber.

The DT chambers have two Φ -type superlayers, measuring ϕ coordinates. The TRACO attempts to correlate the track segments measured in each of them. If a correlation can be found, the TRACO defines a new segment, enhancing the angular resolution and producing a quality hierarchy. Four BTIs in the inner Φ -type superlayer and 12 BTIs in the outer Φ -type superlayer are connected to one TRACO. The number of TRACOs is 25 for the largest muon chamber type. The TRACO is implemented in a 240-pin ASIC with CMOS 0.35 μ m Gate Array technology. The trigger data of at most two track segments per bunch crossing reconstructed by each TRACO are transmitted to the TS, whose purpose is to perform a track selection in a multitrack environment.

The TS has two components, one for the transverse projection (TS_ϕ) and the other for the longitudinal projection (TS_θ). The first one processes the output from the TRACO, whilst the second uses directly the output of the BTIs of the θ view delivered by the Θ -type superlayers present in the three innermost muon stations. The TS_ϕ consists itself of two components, the Track Sorter Slave (TSS) and the Track Sorter Master (TSM). The TSS preselects the tracks with the best quality and the smallest bending angle based on a reduced preview data set coming from the TRACOs in order to save processing time. A select line in the TRACO with the best track is then activated and the TRACO is allowed to send the full data to the TSM. The corresponding preview data are also sent to the TSM for a second stage processing. The TSM analyzes up to seven preview words from the TSSs. The output consists of the two tracks with the highest transverse momentum. There is one TSM per muon station. In the longitudinal view the TS_θ groups the information from the 64 BTIs per chamber. From each BTI two bits are received, a trigger bit and a quality bit. A logic OR of groups of eight bits is applied. The output data consist of 8 bits indicating the position of the muon and 8 quality bits.

The requirement of robustness implies redundancy, which introduces, however, a certain amount of noise or duplicate tracks giving rise to false triggers. Therefore the BTIs, the TRACOs and the different parts of the TS contain complex noise and ghost reduction mechanisms. The trigger and also the read-out data from each of the sixty 30° -sectors of CMS are sent to Sector Collector (SC) units, where the trigger information — the position, transverse momentum and track quality — is coded and transmitted to the DT regional trigger, called the Drift Tube Trigger Track Finder (DTTF), through high-speed optical links.

Cathode Strip Chamber local trigger

The endcap regions are challenging for the trigger since many particles are present and muons at a given p_T have a higher momentum than in the barrel, which gives rise to more bremsstrahlung photons. In addition, photon conversions in a high-radiation (neutron-induced) environment occur frequently. Therefore the CSCs consist of six layers equipped with anode wires and cathode strips, which can be correlated. Muon track segments, also called Local Charged Tracks (LCT), consisting of positions, angles and bunch crossing information are first determined separately in the nearly orthogonal anode and cathode views. They are then correlated in time and in the number of layers hit. The cathode electronics is optimized to measure the ϕ -coordinate, the anode electronics to identify the bunch crossing with high efficiency.

An electric charge collected by the anode wires induces a charge of opposite sign in the cathode strips nearby. The trigger electronics determines the centre of gravity of the charge with a resolution of half a strip width, between 1.5 and 8 mm, depending on the radius. By demanding that at least four layers are hit, the position of a muon can be determined with a resolution of 0.15 strip widths. Due to the finite drift time the anode signals in the six chamber layers are spread out over an interval of more than two bunch crossings. As for the cathodes, at least four coincident hits are required, since in contrast to neutron-induced background a real muon leaves coincident signals in at least four layers with a probability that exceeds 99%. Actually a coincidence of two signals (pre-trigger) is used to identify the crossing, in order to allow for long drift time hits to arrive. A validation of the track occurs if in the following two bunch crossings at least four coincident signals are found. In order to reduce the number of trigger channels 10 to 15 anode wires are ORed. Figure 227 shows the principles of the cathode and anode trigger electronics and the bunch crossing assignment.

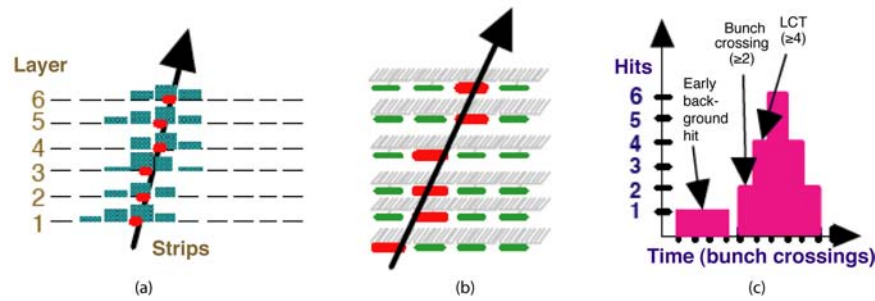


Figure 227: Cathode Strip Chamber Local Trigger: (a) Cathode LCT formation from strips, (b) Anode LCT formation from wire group hits, (c) Bunch crossing assignment.

The track segments from the cathode and anode electronics are finally combined into three-dimensional LCTs. They are characterized by the high-precision ϕ -coordinate in the bending plane, the bending angle ϕ_b , a rough

η -value and the bunch crossing number. The best two LCTs of each chamber are transmitted to the regional CSC trigger, called the CSC Track Finder (CSCTF), which joins segments to complete tracks.

The hardware of the CSC local trigger consists of seven types of electronics boards. Cathode and anode front-end boards (CFEB and AFEB) amplify and digitize the signals. Anode LCT-finding boards (ALCT) latch the anode hits at 40 MHz, find hit patterns in the six chamber layers that are consistent with having originated at the vertex, and determine the bunch crossing. They send the anode information to the Cathode LCT-finding plus Trigger Motherboard (CLTC/TMB) cards. The CLCT circuits look for strip hit patterns consistent with high-momentum tracks. The TMB circuits perform a time coincidence of cathode and anode LCT information. If a coincidence is found, they send the information to the Muon Port Cards (MPC). The TMB selects up to two LCTs based on quality cuts. In order to cancel out ghosts a coincidence with RPC hits is established if two or more LCTs are found. A MPC receives the LCTs from the CLTC/TMBs of one endcap muon station sector, selects the best two or three LCTs depending on the station number and sends them over optical links to the CSC Track Finder. The anode and cathode LCTs and the raw hits are recorded by DAQ motherboards (DAQMB) and transmitted to the CSC detector-dependent units (DDU) belonging to the DAQ system upon reception of a L1A signal. The LHC timing reference, the L1A decision, the bunch crossing number and bunch counter reset signals are distributed by the Clock and Control Boards (CCB). The front-end boards and the ALCTs are mounted directly on the chambers. The rest of the local trigger electronics is housed in 48 peripheral crates on the endcap disks. The optical fibres to the control room depart from there. Except for the comparator-network ASIC implemented in the CLCT module, the CSC trigger electronics is built in FPGA technology.

Resistive Plate Chamber trigger

The RPCs are dedicated trigger detectors. Several layers of double-gap RPCs are mounted on the DT and CSC tracking chambers, six in the central region (two layers on the inside and outside of the two innermost muon stations, one on the inside of the two outermost stations) and four in the forward parts (one layer on the inside of each station). Their main advantage is their excellent timing resolution of about 1 ns, which ensures an unambiguous bunch crossing identification. For triggering purposes the measurement of the momentum of a particle is also important. In the magnetic field, muons are bent in the plane transverse to the LHC beams. It is sufficient to measure the azimuthal coordinate ϕ at several points along the track to determine the bending and thus the p_T . Therefore the RPC strips run parallel to the beam pipe in the barrel, and radially in the endcaps. There are about 165 000 strips in total, which are connected to front-end boards (FEB) handling 16 channels each.

The RPC trigger is based on the spatial and temporal coincidence of hits in several layers. It is segmented in 33 trigger towers in η , which are each subdivided in 144 segments in ϕ . As opposed to the DT/CSC, there is no local processing on a chamber apart from synchronization and cluster reduction. The Pattern Comparator Trigger (PACT) logic [195] compares strip signals of all four muon stations to predefined patterns in order to assign p_T and electric charge, after having established at least three coincident hits in time in four planes. Spatially the PACT algorithm requires a minimum number of hit planes, which varies depending on the trigger tower and on the p_T of the muon. Either 4/6 (four out of six), 4/5, 3/4 or 3/3 hit layers are minimally required. A quality parameter reflects the numbers of hit layers. For six planes there are typically 14 000 possible patterns. The outer section of the hadron calorimeter (HO) consists of scintillators placed after the magnet coil up to $|\eta| < 1.24$. Their signals can also be taken into account by the RPC trigger in order to reduce rates and suppress background [196]. The algorithm requires HO confirmation for low-quality RPC triggers. The optical links from the four HO HTR boards are received by the RPC trigger boards, and the signals are treated and incorporated in the PACT logic like an additional RPC plane, with the required number of planes hit increased by one.

The RPC signals are transmitted from the FEBs, which contain ASICs manufactured in 0.8 μm BiCMOS technology, to the Link Boards (LB), where they are synchronized, multiplexed, serialized and then sent via 1732 optical links to 108 Trigger Boards in 12 trigger crates in the control room. The 1640 LBs are housed in 136 Link Board Boxes. The Trigger Boards contain the complex PACT logic which fits into a large FPGA. There are 396 PACT chips in the system. Since duplicate tracks may be found due to the algorithm concept and the geometry, a *ghost busting* logic is also necessary. The RPC muon candidates are sorted separately in the barrel and forward regions. The best four barrel and the best four forward muons are sent to the Global Muon Trigger. The RPC data record is generated on the Data Concentrator Card, which receives data from the individual trigger boards.

Drift Tube and Cathode Strip Chamber track finders

The regional muon trigger based on the precision tracking chambers consists of the Drift Tube Track Finder (DTTF) in the barrel [197] and the CSC Track Finder (CSCTF) in the endcaps [198]. They identify muon candidates,

determine their transverse momenta, locations and quality. The candidates are sorted by rank, which is a function of p_T and quality. The DTTF and the CSCTF each deliver up to four muons to the Global Muon Trigger.

The track finding principle relies on extrapolation from a source track segment in one muon station to a possible target segment in another station according to a pre-calculated trajectory originating at the vertex. If a compatible target segment with respect to location and bending angle is found, it is linked to the source segment. A maximum number of compatible track segments in up to four muon stations is joined to form a complete track, to which parameters are then assigned. The extrapolation principle is shown in Fig. 228. While the CSCTF in-

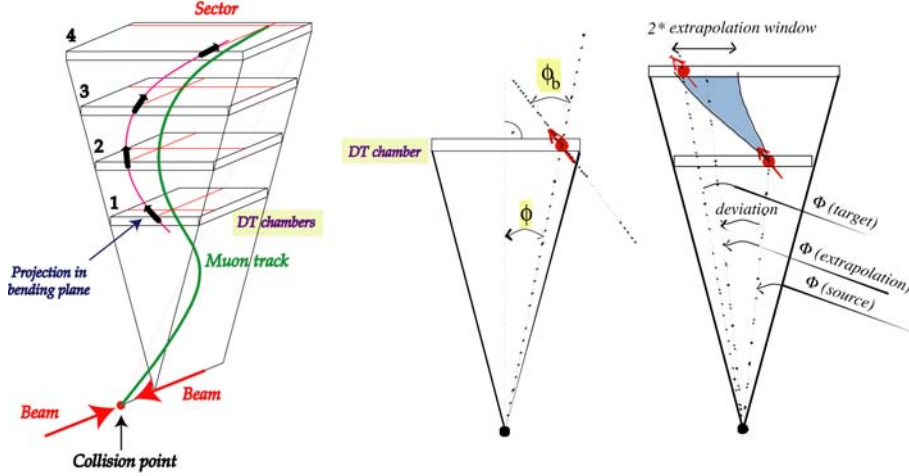


Figure 228: Track Finder extrapolation scheme.

corporates 3-dimensional spatial information from the CSC chambers in the track finding procedure, the DTTF operates 2-dimensionally in the ϕ -projection. A coarse assignment of η is nevertheless possible by determining which chambers were crossed by the track. In most cases an even more refined η -value can be assigned using the information from the θ -superlayers. Both for the DTTF and the CSCTF, the track finder logic fits into high-density FPGAs. For the regional trigger the DT chambers are organized in sectors and wedges. There are twelve horizontal wedges parallel to the beams. Each wedge has six 30°-sectors in ϕ . The central wheel has 2×12 half-width sectors, whereas the four outer wheels are subdivided in 12 full-width sectors each. In the two endcaps the track finding is partitioned in 2×6 60°-sectors. In the overlap region between the DT and CSC chambers, around $|\eta| \approx 1$, information from both devices is used.

In the DTTF the track finding in ϕ is performed by 72 sector processors, also called Phi Track Finders (PHTF). Per chamber they receive at most two track segments from the DT local trigger through optical links. The segment information is composed of the relative position of the segments inside a sector, its bending angle and a quality code. If there are two segments present in a chamber, the second one is sent not at the bunch crossing from which it originated but at the subsequent one, provided that in that crossing no other segment occurred. A tag bit to indicate this *second track segment* status is therefore necessary. The sector processors attempt to join track segments to form complete tracks. The parameters of all compatible segments are pre-calculated. Extrapolation windows, which are adjustable, are stored in look-up tables. Muon tracks can cross sector boundaries, therefore data are exchanged between sector processors and a cancellation scheme to avoid duplicated tracks has to be incorporated.

The track finding in η , with the goal to refine η -values, is performed by 12 η assignment units, also called Eta Track Finders (ETTF). A pattern matching rather than an extrapolation method is used, since for muon stations 1, 2 and 3 the η -information coming from the DT local trigger is contained in a bit pattern representing adjacent chamber areas. The tracks in η are matched with those of the ϕ -projection, if possible. For each wedge, the combined output of the PHTFs and the ETTFs, which consists of the transverse momentum including the electric charge, the ϕ - and η -values and quality for at most 12 muon candidates corresponding to a maximum of two track candidates per 30°-sector, is delivered to a first sorting stage, the Wedge Sorter (WS). There are twelve of these sorters, which have to sort at most 144 candidates in total by p_T and quality. Suppression of remaining duplicate candidates found by adjacent sector processors and track quality filtering is also performed by these units. The two highest-rank muons found in each WS, at most 24 in total, are then transmitted to the final Barrel Sorter (BS). The latter selects the best four candidates in the entire central region, which are then delivered to the Global Muon Trigger for matching with the RPC and CSC candidates.

The DTTF data are recorded by the data acquisition system. A special read-out unit, the DAQ Concentrator Card (DCC) has been developed. It gathers the data from each wedge, through six Data Link Interface Boards (DLI). Each DLI serves two wedges. The DTTF electronics is contained in three racks in the control room. Two racks contain six track finder crates, which each house the electronics for two wedges as well as a crate controller. There is also one Timing Module (TIM) in each of these crates to distribute the clock and other timing signals. The third rack houses the central crate containing the BS, the DCC, a TIM module and boards for interfacing with the LHC machine clock and the Trigger Control System.

As for the DTTF, the core components of the CSCTF are the sector processors. They receive, through optical links, the LCT data from the Muon Port Cards in the peripheral crates. Each sector processor receives up to six LCTs from ME1 and three LCTs each from stations ME2, ME3 and ME4. Up to four track segments are also transmitted from DT station MB2. First the data are latched and synchronized, and the original LCT information is converted to reflect global (η, ϕ) -coordinates. Then nearly all possible pairwise combinations of track segments are tested for consistency with a single track in the processors' extrapolation units. In contrast to the DTTF, no data exchange between neighbour processors is performed. Complete tracks are assembled from the extrapolation results and redundant tracks canceled as in the DTTF. The best three muons per processor are selected and assigned kinematic and quality parameters. The p_T assignment, through SRAM look-up tables, is based on the ϕ -information from up to three muon stations. The data are collected in a detector-dependent unit (DDU) for the read-out. The twelve sector processors are housed in a single crate in the counting room. This crate also contains a Clock and Control Board (CCB) similar to the ones in the local CSC trigger electronics, which distributes the clock, bunch crossing reset, bunch crossing zero and other timing signals. Over the custom-developed GTL+ backplane a maximum of 36 candidate tracks is transmitted to the forward Muon Sorter board, which determines the best four muons in the two endcaps and sends them to the GMT.

Global Muon Trigger

The purpose of the Global Muon Trigger [199] is to improve trigger efficiency, reduce trigger rates and suppress background by making use of the complementarity and redundancy of the three muon systems. It receives for every bunch crossing up to four muon candidates each from the DTs and barrel RPCs, and up to four each from the CSCs and endcap RPCs. The candidate information consists of p_T and charge, η , ϕ and a quality code. From the GCT it also receives isolation and minimally ionizing particle bits for each calorimeter region sized $\Delta\eta \times \Delta\phi = 0.35 \times 0.35$. A muon is considered isolated if its energy deposit in the calorimeter region from which it emerged is below a defined threshold. DT and CSC candidates are first matched with barrel and forward RPC candidates based on their spatial coordinates. If a match is possible, the kinematic parameters are merged. Several merging options are possible and can be selected individually for all of these parameters, taking into account the strengths of the individual muon systems. Unmatched candidates are optionally suppressed based on η and quality. Cancel-out units reduce duplication of muons in the overlap region between the barrel and the endcaps, where the same muon may be reported by both the DT and CSC triggers. Muons are back-extrapolated through the calorimeter regions to the vertex, in order to retrieve the corresponding MIP and ISO bits, which are then added to the GMT output and can be taken into account by the Global Trigger. Finally, the muons are sorted by transverse momentum and quality, first separately in the barrel and forward regions, and then together to deliver four final candidates to the GT. A read-out processor collects the input muon data and the output record. The GMT electronics is housed in the same crate as the GT (Fig. 229). The 16 muon cables are directly connected to the GMT logic board, which has a special four VME slot wide front panel. The logic itself, which is contained in FPGA chips, only occupies one slot. The MIP/ISO bits from the GCT are received and synchronized by three Pipeline Synchronizing Buffer (PSB) input modules, which are also used in the GT. The PSB boards receive the bits via 1.4 Gbit/s serial links and are mounted at the back of the crate, behind the wide logic front panel. The MIP/ISO bits are transmitted from the PSBs to the logic board by GTL+ point-to-point links on the GT backplane.

9.3 Global Trigger

The Global Trigger [200] takes the decision to accept or reject an event at L1 based on trigger objects delivered by the GCT and GMT. These objects consist in candidate-particle, such as e/γ (isolated and non-isolated), muons, central and forward hadronic jets, τ jets, as well as global quantities: total and missing transverse energies, the scalar sum (H_T) of the transverse energies of jets above a programmable threshold, and twelve threshold-dependent jet multiplicities. Objects representing particles and jets are ranked and sorted. Up to four objects are available. They are characterized by their p_T or E_T , (η, ϕ) -coordinates, and quality. For muons, charge, MIP and ISO bits are also available.

The GT has five basic stages: input, logic, decision, distribution and read-out. The corresponding electronics boards use FPGA technology [201]. All of them, as well as the boards of the GMT, are housed in one central 9U high crate, which is shown in Fig. 229. Three Pipeline Synchronizing Buffer (PSB) input boards receive the

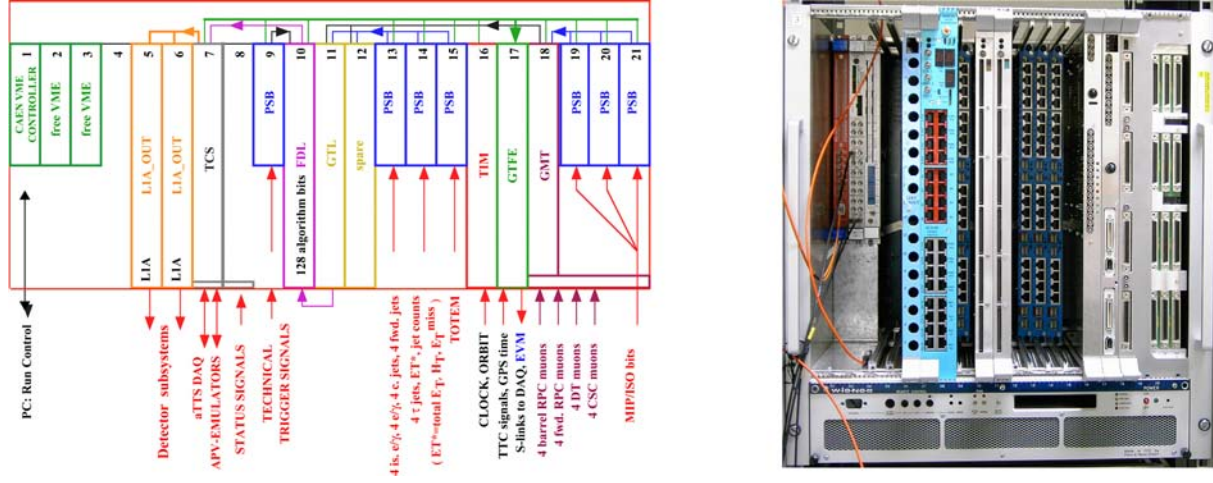


Figure 229: Global Trigger central crate.

calorimeter trigger objects from the GCT and align them in time. The muons are received from the GMT through the backplane. An additional PSB board can receive direct trigger signals from sub-detectors or the TOTEM experiment for special purposes such as calibration. These signals are called *technical triggers*. The core of the GT is the Global Trigger Logic (GTL) stage, in which algorithm calculations are performed. The most basic algorithms consist of applying p_T or E_T thresholds to single objects, or of requiring the jet multiplicities to exceed defined values. Since location and quality information is available, more complex algorithms based on topological conditions can also be programmed into the logic. A graphical interface [202] is used to set up the trigger algorithm menu. The results of the algorithm calculations are sent to the Final Decision Logic (FDL) in the form of one bit per algorithm. The number of algorithms that can be executed in parallel is 128. Up to 64 technical trigger bits may in addition be received directly from the dedicated PSB board. For normal physics data taking a single trigger mask is applied, and the L1A decision is taken accordingly. For commissioning, calibration and tests of individual subsystems up to eight final ORs can be applied and correspondingly eight L1A signals can be issued. The distribution of the L1A decision to the subsystems is performed by two L1A_OUT output boards, provided that it is authorized by the Trigger Control System described in Sect. 9.4. A Timing Module (TIM) is also necessary to receive the LHC machine clock and to distribute it to the boards. Finally, the Global Trigger Frontend (GTFE) board collects the GT data records, appends the GPS event time received from the machine, and sends them to the data acquisition for read-out.

9.4 Trigger Control System

The Trigger Control System (TCS) [203] controls the delivery of the L1A signals, depending on the status of the sub-detector read-out systems and the data acquisition. The status is derived from signals provided by the Trigger Throttle System (TTS). The TCS also issues synchronization and reset commands, and controls the delivery of test and calibration triggers. It uses the Timing, Trigger and Control distribution network [204], which is interfaced to the LHC machine.

The TCS architecture is represented in Fig. 230. Different subsystems may be operated independently if required. For this purpose the experiment is divided into 32 partitions, each representing a subsystem or a major component of it. Each partition is assigned to a partition group, also called a TCS partition. Within such a TCS partition all connected partitions operate concurrently. For commissioning and testing up to eight TCS partitions are available, with their own L1A signals distributed in different time slots allocated by a priority scheme or in round robin mode. During normal physics data taking there is only one single TCS partition. Subsystems may either be operated centrally as members of a partition or privately through a Local Trigger Controller (LTC). Switching between central and local mode is performed by the TTCi (TTC CMS interface) module, which provides the interface between the respective trigger control module and the destinations for the transmission of the L1A signal and other fast commands for synchronization and control. The TTC Encoder and Transmitter (TTCex) module

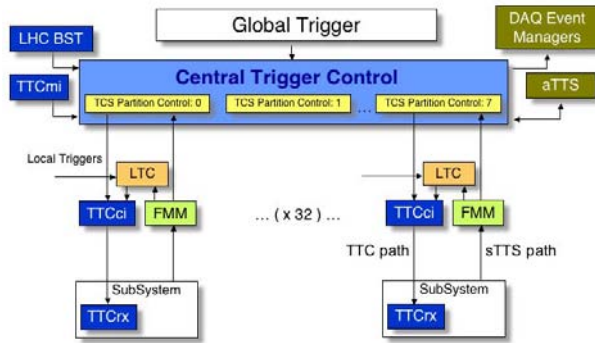


Figure 230: Trigger Control System architecture.

encodes the signals received from the TTCci and drives optical splitters with a laser transmitter. The LHC clock is received from the TTC machine interface (TTCmi). At the destinations the TTC signals are received by TTC receivers (TTCrx) containing low-jitter quartz PLLs.

The central TCS module, which resides in the Global Trigger crate, is connected to the LHC machine through the TIM module, to the FDL through the GT backplane, and to 32 TTCci modules through the LA1_OUT boards. The TTS, to which it is also connected, has a synchronous (sTTS) and an asynchronous (aTTS) branch. The sTTS collects status information from the front-end electronics of 24 sub-detector partitions and up to eight tracker and preshower front-end buffer emulators. The status signals, coded in four bits, denote the conditions *disconnected*, *overflow warning*, *synchronization loss*, *busy*, *ready* and *error*. The signals are generated by the Fast Merging Modules (FMM) through logical operations on up to 32 groups of four sTTS binary signals and are received by four conversion boards located in a 6U crate next to the GT central crate. The aTTS runs under control of the DAQ software and monitors the behaviour of the read-out and trigger electronics. It receives and sends status information concerning the 8 DAQ partitions, which match the TCS partitions. It is coded in a similar way as the sTTS. Depending on the meaning of the status signals different protocols are executed. For example, in case of warning on the use of resources due to excessive trigger rates, prescale factors may be applied in the FDL to algorithms causing them. A loss of synchronization would initiate a reset procedure. General trigger rules for minimal spacing of L1As are also implemented in the TCS. The total deadtime estimated at the maximum L1 output rate of 100 kHz is estimated to be below 1%. Deadtime and monitoring counters are provided in the TCS. The central board sends to the DAQ Event Manager (EVM) located in the surface control room the total L1A count, the bunch crossing number in the range from 1 to 3564, the orbit number, the event number for each TCS/DAQ partition, all FDL algorithm bits and other information.

10 Data Acquisition

The architecture of the CMS Data Acquisition (DAQ) system is shown schematically in Fig. 231. The CMS Trigger and DAQ system is designed to collect and analyse the detector information at the LHC bunch crossing frequency of 40 MHz. The rate of events to be recorded for offline processing and analysis is on the order of 100 Hz. At the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, the LHC rate of proton collisions will be around 20 per bunch crossing, producing approximately 1 MByte of zero-suppressed data in the CMS read-out systems. The first level trigger is designed to reduce the incoming average data rate to a maximum of 100 kHz, by processing fast trigger information coming from the calorimeters and the muon chambers, and selecting events with interesting signatures. Therefore, the DAQ system must sustain a maximum input rate of 100 kHz, for a data flow of 100 GByte/s coming from approximately 650 data sources, and must provide enough computing power for a software filter system, the High Level Trigger (HLT), to reduce the rate of stored events by a factor of 1000. In CMS all events that pass the Level-1 (L1) trigger are sent to a computer farm (Event Filter) that performs physics selections, using faster versions of the offline reconstruction software, to filter events and achieve the required output rate. The design of the CMS Data Acquisition System and of the High Level Trigger is described in detail in the respective Technical Design Report [188].

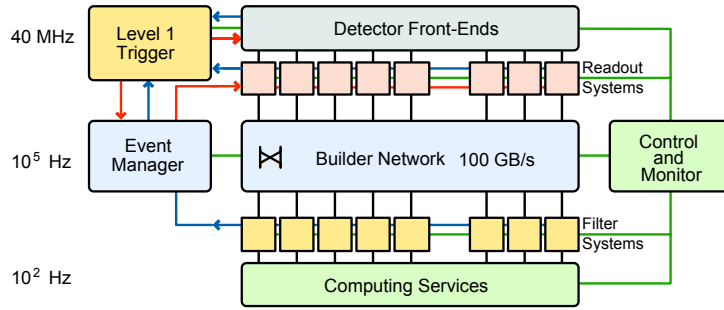


Figure 231: Architecture of the CMS DAQ system.

The read-out parameters of all sub-detectors are summarized in Table 27. Each data source to the DAQ system is expected to deliver an average event fragment size of ≈ 2 kByte (for pp collisions at design luminosity). In some case two data sources are merged in order to reach this nominal size.

Table 27: Sub-detector read-out parameters.

sub-detector	number of channels	number of FE chips	number of detector data links	number of data sources (FEDs)	number of DAQ links (FRLs)
Tracker pixel	≈ 44 M	15840	1504	40	40
Tracker strips	≈ 9.3 M	≈ 72 k	≈ 36 k	440	250 (merged)
Preshower	144384	4512	1128	56	56
ECAL	75848	≈ 21 k	≈ 9 k	54	54
HCAL	9072	9072	3072	32	32
Muons CSC	≈ 500 k	≈ 76 k	540	8	8
Muons RPC	192 k	≈ 8.6 k	732	3	3
Muons DT	195 k	48820	60	10	10
Global Trigger	n/a	n/a	n/a	3	3
CSC, DT Track Finder	n/a	n/a	n/a	2	2
Total	≈ 55 M			626	458

A schematic view of the components of the CMS DAQ system is shown in Fig. 232. The various sub-detector front-end systems (FES) store data continuously in 40-MHz pipelined buffers. Upon arrival of a synchronous L1 trigger ($3.2 \mu\text{s}$ latency) via the Timing, Trigger and Control (TTC) system [204, 207], the corresponding data are extracted from the front-end buffers and pushed into the DAQ system by the Front-End Drivers (FEDs). The data from the FEDs are read into the Front-end Read-out Links (FRLs) that are able to merge data from two FEDs. The number of FRLs corresponding to the CMS read-out parameters of Table 27 is 458. In the “baseline” configuration, there are 512 FRLs. These additional inputs are used for combined operation with the TOTEM experiment [2],

10.1 Sub-detector read-out interface

The design of the FED is sub-detector specific, however, a common interface from the FED to the central DAQ system has been implemented. The hardware of this interface is based on S-Link64 [208]. The FED encapsulates the data received from the front-end electronics in a common data structure by adding a header and a trailer that mark the beginning and the end of an event fragment. The header and trailer partially consists of event information used in the event building process, such as the event number derived from the TTC L1A signals. The trailer includes a CRC (Cyclic Redundancy Check) of the data record. The payload of the event fragments is only inspected in the Event Filter.

The physical implementation of one element of the sub-detector read-out interface and its FED builder port is shown in Figs. 233 and 234.

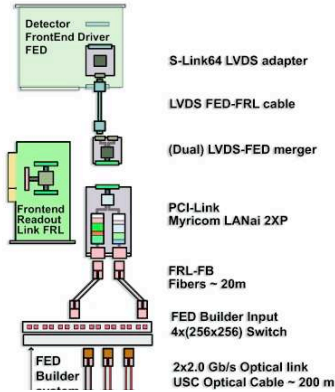


Figure 233: Diagram of a generic sub-detector read-out into the FED builder.

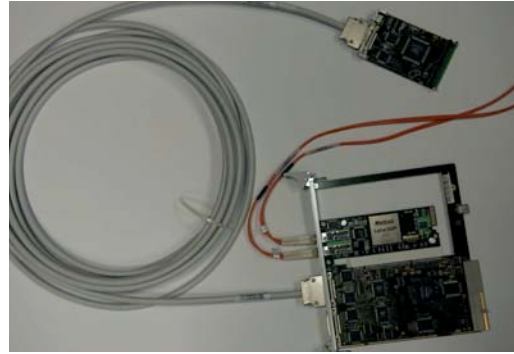


Figure 234: Photograph of S-link64 sender card, LVDS cable and compact-PCI FRL card with embedded LANai2XP NIC. Connected to the NIC are fibres that go to the FED Builder input switch.

The S-Link64 Sender card

The S-Link64 Sender card is a custom developed Common Mezzanine Card (CMC), directly plugged into the sub-detector FED. It receives data from the FED via an S-Link64 port and checks the CRC in order to check transmission errors over the S-Link. The card is able to buffer up to 1.6 kByte of data before generating back-pressure to the FED. The CMC has an LVDS converter to interface with a copper cable which can have a maximum length of 11 m. This cable (Multi-conductor round cable v98 manufactured by 3M) comprises 14 twisted pairs, which are individually shielded. The link provides feedback lines in order to signal back-pressure and to initiate an automatic self test. The nominal data transfer rate over the LVDS cable is 400 MByte/s (50 MHz clock, 64 bits), which is twice the maximum sustained design throughput of the DAQ system.

The Front-end Read-out Link

The FRL is a custom 6U Compact-PCI card (Fig. 235). It has three interfaces: an input interface which handles up to two LVDS cables; an output interface to the FED Builder implemented as a 64-bit/66-MHz PCI connector for a Network Interface Card (NIC); and a configuration and control interface which is a PCI bus interface connected to the Compact-PCI backplane. The function of the FRL board is performed by two FPGAs (Altera EP20K100EFC324-1 for the PCI bridge and EP1S10F672C6 for the main logic).

The FRL receives event fragments and checks the CRC in order to check transmission errors over the LVDS cable. In the case where the FRL receives data from two FEDs, the two data records are merged. Data are buffered in memories of 64 kByte size and pushed into the NIC in fixed size blocks via the onboard PCI bus.

The FRL card also provides monitoring features, such as the ability to spy on a fraction of events via the Compact-PCI bus, and to accumulate histograms of quantities such as fragment size.

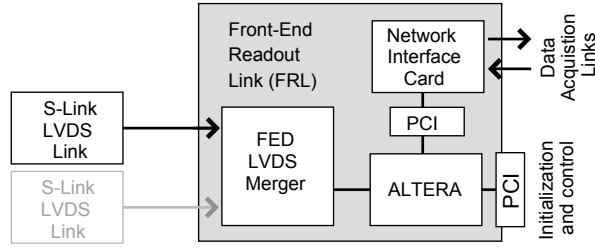


Figure 235: FRL layout.

Up to 16 FRL cards are placed in a crate with a Compact-PCI backplane. Each crate is connected to a PC via a compact PCI bridge (StarFabric CPCI/PCI system expansion board from Hartmann Elektronik), which is used for configuration, control and monitoring. There are 50 FRL crates in total.

10.2 The Trigger Throttling System and sub-detector fast-control interface

The TTS provides the feedback from all FEDs and their associated front-end systems to the GTP. It is a hardwired system, acting on the dataflow with a reaction time of less than $1 \mu\text{s}$. Each FED provides fast signals indicating the state of the read-out. The states *Ready*, *Warning*, *Busy*, *Out-Of-Sync* and *Error* are defined (listed in order of increasing priority). Ready, Warning and Busy are generated according to the amount of internal data buffering available and are used to indicate if more triggers can be accepted. Given the trigger rules (Sect. 9.4) and a TTS latency of $1 \mu\text{s}$, a FED has to be able to accept 2 more triggers after asserting Busy state. Out-Of-Sync and Error indicate that synchronization was lost or an error occurred in the front-end electronics. The GTP attempts to recover automatically from these states by issuing a L1-Resync or L1-Reset command via the TTC system. These fast TTS signals are transmitted over shielded RJ-45 cables with four twisted pairs using the LVDS signaling standard.

The Fast Merging Module

For flexibility, FEDs are grouped into 32 TTC partitions which may be operated independently of each other. The Level-1 Trigger Control System separately distributes triggers to these 32 TTC partitions and separately receives trigger throttling signals for each TTC partition. TTS signals from all FEDs in a TTC partition thus need to be merged with low latency. Dedicated Fast Merging Modules (FMMs), have been designed for this task. These modules can merge and monitor up to 20 inputs and have quad outputs. Optionally, FMMs can be configured to merge two independent groups of 10 inputs with two independent twin outputs. For partitions with more than 20 FEDs, FMMs are cascaded in two layers.

The FMM is a custom-built 6U compact-PCI card (Figs. 236 and 237). It has three main components: a PCI Interface FPGA, a main logic FPGA and an on-board memory block. The 80 MHz internal clock of the FMM is not synchronized to the LHC clock. Input signals are synchronized to the internal clock by requiring two successive samples in the same state. The input signals are then merged by selecting the highest priority input signal from the enabled inputs according to the signal priorities listed above. Optionally, Out-of-Sync input signals are only taken into account if the number of inputs in Out-of-Sync state exceeds a programmable threshold.

The FMM also provides extensive monitoring capabilities in order to diagnose the causes for dead-times. Each state transition at the inputs is detected and stored with a time-stamp (25 ns resolution) in a circular buffer memory that can hold up to 128 k transitions. The times spent in the states Warning and Busy are counted with 25 ns resolution for each input channel and for the output(s).

FMM cards are configured, controlled and monitored by a PC via a compact-PCI interface (StarFabric CPCI/PCI system expansion board from Hartmann Elektronik). The total system comprises 8 FMM crates with up to 9 FMMs in each crate. A total of 60 FMM modules are needed in order to merge the TTS signals of all TTC partitions of CMS.

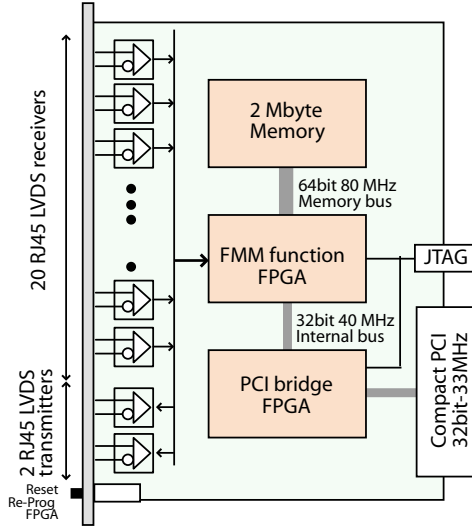


Figure 236: Block diagram of the FMM.



Figure 237: Photograph of a FMM module.

10.3 Testing

In order to test and commission the central DAQ system independently of the GTP and of the sub-detector DAQ systems, a number of additional components have been developed. These are not used in standard data taking.

The Global Trigger Processor emulator (GTPe) [205] emulates the functionality of the GTP (see Fig. 232). It reproduces the LHC beam structure, generates random or clocked triggers up to 4 MHz, respects the trigger rules, applies partitioning, transmits the GTPe data fragment over S-Link64 and receives and handles sTTS and aTTS backpressure signals. The hardware implementation is based on the Generic-III PCI card [206] and an interface module GTPe-IO.

In normal data taking mode, triggers from the global trigger are distributed to the FEDs via the TTC. When using the GTPe, special test triggers are sent directly to the FRL crates via a lemo cable to a trigger distributor card which distributes the trigger over the backplane to all the FRLs in the crate. Because the FEDs are not being used in this mode, busy signals from the FRLs are collected by the trigger distributor card and sent to a dedicated set of FMM modules for fast merging of this subset of sTTS signals. A dedicated mode of the FRL firmware handles the GTPe test triggers and instead of reading out the FEDs, the FRL generates data fragments with sizes according to a predefined table. In this way, the full central DAQ system can be tested.

10.4 The Event Builder

A schematic view of the Event Builder system is shown in Fig. 238. The event builder is composed of two stages: the FED-builder and the RU-builder. Each of the ≈ 512 FRLs generate event fragments with an average size of ≈ 2 kByte and the FED-builder is in charge of transporting these fragments to the surface building (SCX) and assembling them into 72 super-fragments with an average size of ≈ 16 kByte. The super-fragments are then stored in large buffers in Read-out Units (RU), waiting for the second stage of event building (RU-builder), which is implemented with multiple 72×72 networks. There can be up to 8 RU-builders, or DAQ slices, connected to the FED-builder layer. Each FED-builder is in charge of distributing the super-fragments, on an event by event basis, to the RU-builders and ensures that all super-fragments corresponding to one event go to one and only one DAQ slice, and are read by one Builder Unit (BU) of the RU-builder network. The complete event is then transferred to a single unit of the Event Filter. By utilising an architecture with two-stage event building, the full size system can be progressively deployed slice by slice while the traffic load to the second stage is optimized. The event builder is implemented with commodity equipment, including processing nodes, network switches and network interfaces.

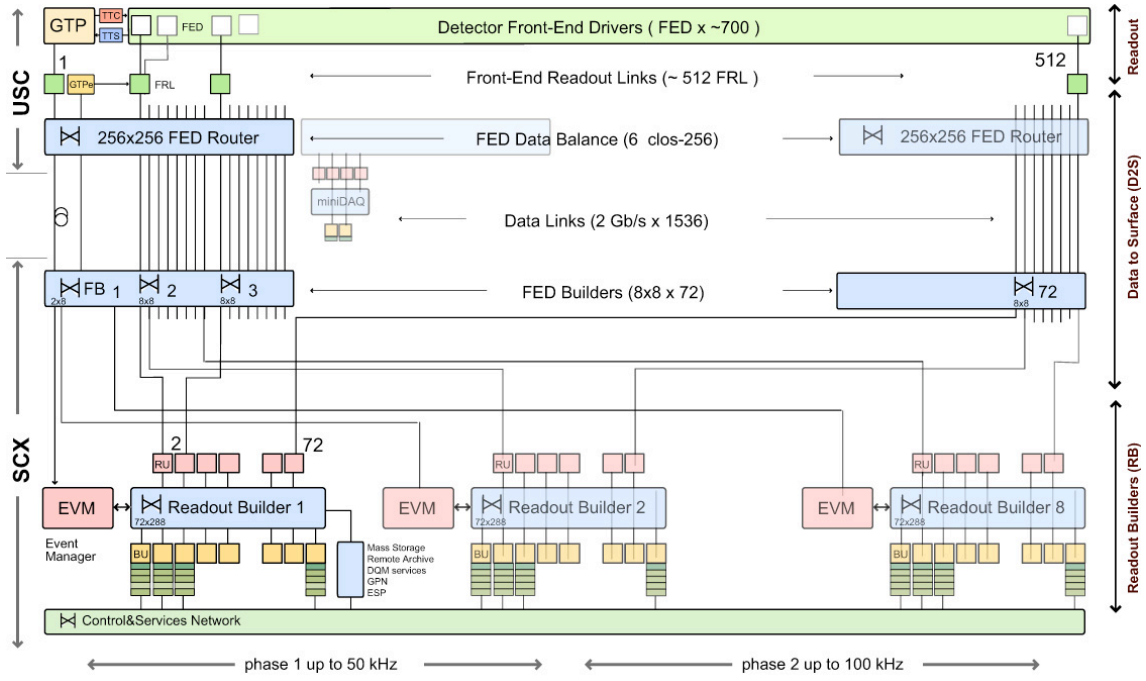


Figure 238: Schematic view of the Event Builder.

The processing nodes are server PC's with PCI busses to host the NICs. They run the Linux operating system.

The event builder is lossless and when necessary back-pressure is propagated up to the FRL and subsequently to the FED, which can throttle the trigger via the TTS.

The FED-Builder stage

The FED Builder is based on Myrinet [209], an interconnect technology for clusters, composed of crossbar switches and NICs, connected by point to point bi-directional links. It employs wormhole routing and flow control at the link level. The LANai ASIC on the NIC contains an embedded RISC processor.

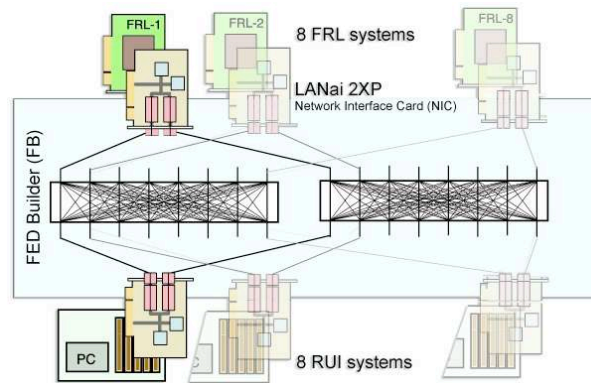


Figure 239: 8×8 FED builder with a two-rail network.

The FED-builder stage is implemented using multiple $N \times M$ FED-builders. In the baseline configuration, $N \leq 8$ and M equals the number of DAQ slices, which is 8 for the full system. In general, one $N \times M$ FED-builder comprises N input NICs hosted by the FRL, M output NICs hosted by the RU PCs, and two independent $N \times M$ switches to form 2 rails (Fig. 239). The NIC (M3F2-PCIXE-2 based on the LANai2XP) has two bi-directional fibre optic ports, with 2 Gbit/s data rate each. Each rail of one NIC is connected to an independent switch. In practice, instead of a large amount of small physical switches, a single large switching fabric is used per rail (see

below). This 2-rail configuration doubles the bandwidth to 4 Gbit/s per FRL and provides redundancy.

The software running on the NICs has been custom developed in the C language. The FED-builder input NICs are programmed to read fragments from the FRL and to send them to the switch fabric, with a destination assigned on the basis of the fragment event number and a predefined look-up table. The FED-builder output NICs are hosted by the RU-builder PCs. They are programmed to concatenate fragments with the same event number from all the connected input cards, in order to build the super-fragments. The FED-builder is lossless due to a basic flow control and retransmission protocol, implemented on the RISC processor on the NIC.

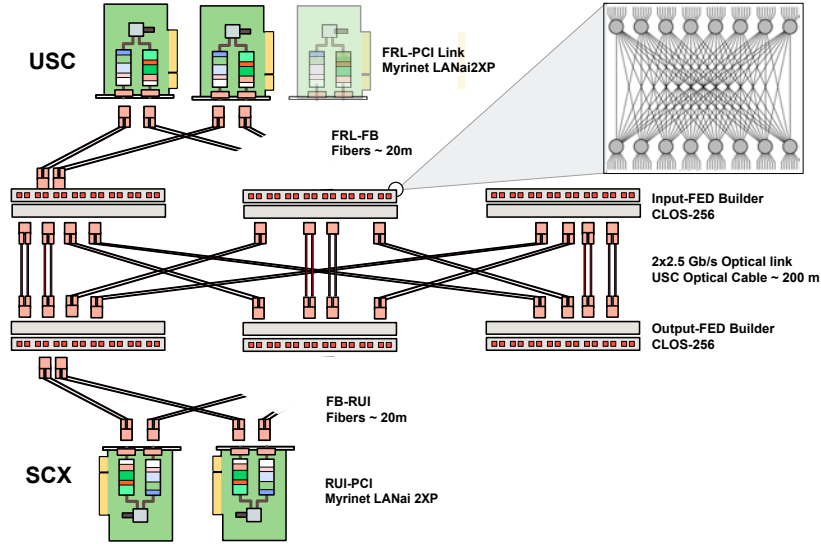


Figure 240: The FED builder stage switching fabric (only 1 rail is shown).

The physical switching fabric is composed of six Clos-256 enclosures per rail (Fig. 240). A Clos-256 enclosure is internally composed of two layers of 16×16 (Xbar32) cross-bars. Three Clos-256 enclosures are located in the underground electronics room (USC), while the other three are in the surface DAQ building (SCX). They are connected by 768 200 m optical fiber pairs per rail, bundled in 12 cables. The Clos-256 enclosures are partly populated with linecards. Currently, the system has ports to accommodate a total of 576 FRLs and 576 RU PCs. In the baseline it is configured as 72 times 8×8 FED-builders. The use of a large switching fabric allows for a high re-configurability of the FED-builders in software, which enables traffic balancing by redefining the super-fragment composition and traffic rerouting in case of a hardware failure.

The Myrinet NIC can transfer 4.0 Gbit/s data rate over the two optical rails. For random traffic the network efficiency is approximately 60%, due to head-of-line blocking. An event building throughput of about 300 MByte/s per node is achieved for variable sized fragments with a nominal average of 2 kBytes by using two rails [210]. Hence, the sustained aggregate throughput through the FED-builder stage is ≈ 1.4 Tbit/s, satisfying the CMS DAQ requirements. A maximum peak throughput of ≈ 2 Tbit/s is possible, if fully exploiting the FRL and Myrinet bandwidth, and using traffic shaping in the FED-builders.

The RU-builder stage

The RU-builder stage assembles super-fragments into complete events. Each RU-builder must collect event superfragments of average size ≈ 16 kByte from 72 data sources and build them into complete events at a rate of 12.5 kHz. A diagram of one slice of the RU-builder, indicating its various components is shown in Fig. 241. A RU-builder is made up of a number of Read-out Units (RU), Builder Units (BU) and a single Event Manager (EVM) connected together by a switching network. The EVM supervises the data flow in the RU-builder and receives a data record from the GTP via a dedicated FED-builder. The EVM allocates events on request to a BU, which subsequently collects the super-fragments from all RUs. From the BUs the complete events are sent to the Filter Units (FU).

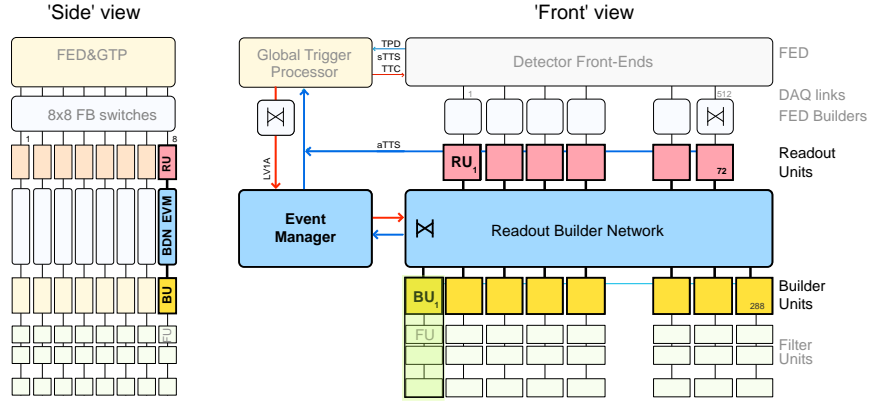


Figure 241: A Read-out Builder slice.

In the baseline configuration the number of RUs per RU-builder is 72 and there is a factor of 4 more FU nodes. The RU, BU and EVM software components can be distributed in various ways on the PC nodes. In the baseline, there are two layers of PCs: RUs and BU-FUs. Here, the events are built by the BU software component in the same PC that runs the filter unit (FU) software, referred to as BU-FU. The RU nodes are server PCs with two 2 GHz dual-core Xeon processors (e5120) and 4 GByte of main memory (Dell PowerEdge 2950). They host a Myrinet NIC (M3F2-PCIXE-2) for the input from the FED-builder and a 4-port GbE NIC (PEG4I-ROHS Quad port copper Gigabit Ethernet PCI Express Server Adapter from Silicom Ltd.). In the baseline, two links of the 4-port NIC are cabled to the network switch, which is sufficient to satisfy the throughput requirement of at least 200 MByte/s per RU node. A single ethernet link for each BU-FU node is sufficient, as the required throughput is 50 MByte/s. The EVB switching network is implemented with eight E-1200 switches (Terascale E-1200 router from Force10), one per DAQ-slice.

The RU-builder is based on TCP/IP over Gigabit Ethernet. The design choice of TCP/IP over Gigabit Ethernet has the advantage of using standard hardware and software. When operating an event builder over Ethernet close to wire speed, typically packet loss occurs because hardware flow control is not propagated from end-point to end-point through the switches. TCP/IP provides a reliable transport service that removes the need for the event building application to detect and deal with lost packets. It also provides flow control and congestion control. TCP/IP uses a substantial amount of host resources. Roughly 20% of a 2 GHz Xeon CPU core is required to transmit 1 Gbit/s, when using jumbo-frames (MTU=9000 Bytes).

For the event builder with 2 Ethernet links per RU node, a throughput of ≈ 240 MByte/s per RU node has been achieved at the nominal super-fragment size of 16 kBytes [211]. This can be increased, if needed for higher trigger rate or larger event sizes, to ≈ 360 MByte/s per RU node by installing a third Ethernet link per RU node. At a nominal throughput of 60 MByte/s, corresponding to 1/4 of the RU throughput, the event building tasks consume roughly 10% of the CPU resources on the BU-FU event filter nodes.

10.5 The Event Filter

The primary goal of the Event Filter complex is to reduce the nominal Level-1 Trigger accept rate of 100 kHz to a manageable level for the mass storage and offline processing systems while preserving interesting physics events and not introducing additional experiment dead-time.

The Event Filter complex:

- collects events accepted by the Level-1 Trigger system from the Event Builder and distributes them to worker nodes for further processing;
- performs basic consistency checks on the event data;
- runs offline-quality reconstruction modules and filters to process and select events for storage (High Level Trigger, “HLT”);
- generates, collects, and distributes Data Quality Monitoring (DQM) information resulting from online event processing in the HLT;

- serves a subset of the events to local and remote online consumers (EventServer, ES) for calibration and DQM;
- routes selected events to local storage in several online streams according to their trigger configuration;
- transfers data from local storage at the CMS site to mass storage in the CERN data centre at the Meyrin site.

Architecture and data flow

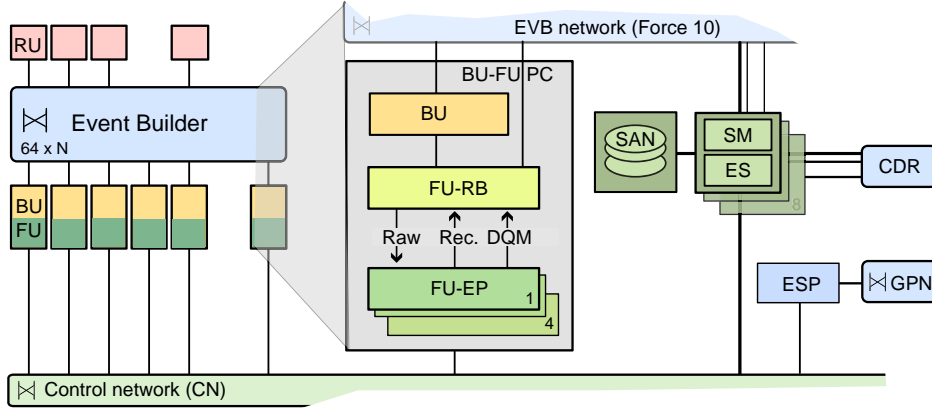


Figure 242: Architecture and data flow of the Filter Farm.

The architecture of the CMS Event Filter is schematically illustrated in Fig. 242. The Event Filter hardware consists of a large farm of processors (on the order of 1000), running the HLT selection (Filter Farm), and a data logging system connected to a Storage Area Network (SAN). The Builder Unit (BU), belonging to the event builder, delivers complete events to one of multiple copies of the Filter Unit Event Processors (FU-EP) via the Filter Unit Resource Broker (FU-RB). A logically separate switch fabric provides the connectivity from the Filter Units to the data logging nodes. These data logging nodes are connected to a Fibre-Channel SAN, that is capable of a peak bandwidth of 1 GByte/s and has a capacity of several hundred TBytes.

The filter farm is logically subdivided into groups of processing nodes (Builder/Filter Unit, BU-FU). Each BU-FU hosts an identical set of software modules in a distributed environment based on the XDAQ online framework (Sect. 10.7). As mentioned above, there are three separate applications, the Builder Unit (BU), the Filter Unit “Resource Broker” (RB) and the Event Processor (EP). The RB is responsible for managing memory resources for input and output to/from the HLT processes, and the communication with the Event Builder and the data logger. A complete event is handed by the RB, upon request, to the Event Processor (EP). The EP process uses the CMS reconstruction framework (CMSSW) [212] to steer the execution of reconstruction and selection code forming the HLT selection. Multiple EP processes can coexist in a single processor to provide concurrent execution and thus saturate the processing CPU resources.

Event processing

The EP reconstruction and selection algorithms are configured at the start of each data-taking run, by a configuration management system based on Oracle and working under run control supervision. The reconstruction, selection, and analysis modules specified in the configuration are instructed to obtain calibration constants and other time-dependent information from an Oracle database using standard access methods supported by the reconstruction framework.

Each reconstruction or selection algorithm runs in a predefined sequence, starting from raw data unpacking modules, which deal with sub-detector specific raw-data formats. The binary raw-data is formatted into C++ objects associated with sub-detector channels through a channel map using the FED block identifiers. Reconstruction modules can attach reconstructed objects to the event data structure. A full history of the execution of the HLT is attached to each accepted event. In addition, bookkeeping information is maintained by each EP process, and periodically collected by a supervising system to provide full accounting of events accepted and rejected by each HLT path. When a decision is reached, accepted events, comprising raw data and event reconstruction information produced by the HLT, are handed back to the RB for transfer to the data logging process.

Monitoring

The operation of unpacking and reconstruction modules running in the Event Processors is monitored using the Physics and Data Quality Monitoring infrastructure (DQM). Additional analysis modules may be executed in the process, outside the normal selection sequences, to provide fast feedback about the quality of the data using the same infrastructure. Information collected by the DQM on the individual FU nodes is periodically forwarded to the data logging system via the RB, providing a DQM Collector functionality (DQMC).

Data logging and Event Server

Events accepted for storage by one of the EP processes are transmitted to the RB, which forwards them to the Storage Manager (SM) process running in the data logger nodes via the data logging network. The SM is responsible of managing the various online streams to provide an appropriate granularity of event data for transfer, processing and bookkeeping. The data logger supports both disk streams for physics or calibration data, and network streams for the usage of consumer processes carrying out calibration or monitoring tasks. The network streams are created “on demand” by a consumer process connecting to the EventServer (ES) function of each SM. In normal operation with multiple SMs, the ES is multiplexed across the various sub-farm SMs by a caching Event Server Proxy (ESP). File and network streams can deal transparently with event data or DQM data.

Each data logger node hosting a SM is responsible for direct management of its disk pool in the storage area network. This includes correct interleaving of write transactions of the SM and data transfer, via the Central Data Recording (CDR) network to the offline systems for analysis and distribution to remote sites.

10.6 Networking and Computing Infrastructure

Networking Infrastructure

The general networking infrastructure of the experiment is based on Ethernet and is separated into:

- CMS Technical Network (CMS-TN);
- CERN general purpose network (CERN-GPN);
- Connection to the Central Data Recording (CDR);
- LHC technical network (LHC-TN).

These four networks are interconnected. In addition there are the dedicated DAQ event building networks (Myrinet and Ethernet) and dedicated networks for equipment control which have no connection to any of those four.

The CMS-TN is a general purpose network connecting all machines directly related to the operation of the experiment. It is used for system administration of all computers and for configuration, control and monitoring of all online applications. This network is not accessible directly from outside, but can be reached through dual-homed Application Gateway (AG) machines, which have one connection to CMS-TN and one to CERN-GPN switches. The CMS-TN is implemented as a distributed backbone with two routers located in SCX, and two in USC. Typically all computers in a rack are served by a switch in that rack. Each switch has two Gigabit Ethernet uplinks to the backbone routers. The desktop computers in the control room are also connected to the CMS-TN.

The CERN-GPN is the CERN-site network. The GPN will be used at all places to connect visitors (wired or wireless), who will use the application gateways to connect to the CMS-TN. This network will typically provide 1 Gbit/s connections to the GPN backbone.

The CDR connects to the Tier0 (Sect. 12) system at the CERN Meyrin site using a minimum of 10 Gbit/s. A set of 8 triple-homed servers (one connection on the DAQ, one to CERN-TN, one on the CDR switches) are dedicated to this task. These are the Storage Manager nodes.

The LHC-TN allows data exchange between some of the CMS equipment and the LHC controls. The CMS-TN interconnects with the CERN-GPN and the LHC-TN using a filter implemented in the backbone routers.

Computing infrastructure

As previously discussed, the DAQ system comprises thousands of computing nodes. These are all rack mounted PCs. All PCs are Intel x86 based, running the CERN distribution of the Redhat Linux OS. The PC cluster includes a

global file server and other services to be able to operate independently from the CERN-GPN. System installation is done with the Quattor toolkit [213]. In addition around a hundred PCs are used for DCS, running Microsoft Windows.

Database services are provided by a 6-node Oracle Real Application Cluster.

10.7 DAQ software, control and monitor

As stated previously, the CMS DAQ is designed in a way such that its hardware implementation can be staged as the LHC accelerator luminosity increases as well as the experiment's need for higher throughput. Thus the CMS DAQ online software must be highly scalable and must also support a diverse hardware base. The online software encompasses a distributed processing environment, data acquisition components, the run control and the detector control system. All subsystems of the DAQ have adopted the central online software frameworks with the philosophy of using common and standardized software technologies in order to reduce the effort associated with the maintenance and evolution of the detector read-out system over the long lifetime of the experiment.

XDAQ Framework

The XDAQ (Cross-Platform DAQ Framework) framework [214] is designed for the development of distributed data acquisition systems. XDAQ includes a distributed processing environment called the "executive" that provides applications with the necessary functions for communication, configuration control and monitoring. Written entirely in C++, it provides applications with efficient, asynchronous communication in a platform independent way, including the use of memory pools for fast and predictable buffer allocation, support for zero-copy operation and a dispatching mechanism for an event-driven processing scheme. A copy of the executive process runs on every processing node in the data acquisition network.

XDAQ Applications are modeled according to a software component model [216] and follow a prescribed interface. They are compiled and the object code is loaded and configured dynamically at run-time into a running executive using the XML schema [221]. Multiple application components, even of the same application class, may coexist in a single executive process.

All configuration, control and monitoring can be performed through the SOAP/http [217] protocol, widely used in Web enabled applications. A rich set of data structures, including lists, vectors and histograms are exportable and can be inspected by clients through the executive SOAP services. Additional utility components provide support for hardware and database access.

XDAQ Applications and Libraries

XDAQ components [219] developed for CMS include applications such as the distributed Event Builder (EVB), sub-detector electronics configuration and monitoring components (FEC and FED), and central DAQ applications.

The generic event builder application consists of the main XDAQ components: a read-out unit (RU), a builder unit (BU) and an event manager (EVM). Data that are recorded by custom read-out devices are forwarded to the read-out unit application. A RU buffers data from subsequent single events until it receives a control message to forward a specific event fragment to a BU. A BU collects the event fragments belonging to a single event from all the RUs and combines them into a complete event. The BU provides an interface to the event data processors that apply event-filtering algorithms and provide data persistency (Sect. 10.5). The EVM interfaces to the trigger read-out electronics and so controls the event building process by mediating control messages between RUs and BUs. For efficient transmission of binary (i.e. event) data the I2O specification [218] is followed. The event builder is a generic application that can run on a wide range of underlying hardware and is also used in local data acquisition systems, such as sub-detector test beams [215].

Data transmission in the XDAQ programming environment is carried out by special application components named peer transports. Peer transports register themselves with the XDAQ executive as being capable of resolving addresses as well as transmitting and receiving data. Communication between XDAQ applications is then accomplished by using an executive function that, when invoked, redirects the outgoing message to the proper peer-transport that in turn delivers the data over the associated medium. In this way the framework is independent of any transport protocol or network and can be extended at any time to accommodate newly appearing communication technologies. Baseline peer transports have been implemented for efficient binary data transmission using an asynchronous TCP/IP protocol and for simple message handling using the SOAP XML message format, however

the framework is independent of the peer transport used, so optimisation of this layer is transparent to the rest of the XDAQ applications.

Libraries and device drivers have been developed that provide generic user access to VME and PCI modules and support Myrinet. Additional XDAQ components include support for the CMS custom front end devices, persistent monitoring message support and a gateway application to interface the XDAQ SOAP environment to the Detector Control System (Sect. 10.8).

HyperDAQ

An extension to the XDAQ framework, HyperDAQ [220], exploits the http protocol which creates an entry point into XDAQ executives. A combination of HyperDAQ and Peer-to-Peer technology, used to discover new XDAQ applications providing data content, presents to the user links to the data content providers as they become available (Fig. 243). In this way, any node in the distributed online cluster can become an access point from which the entire distributed system can be explored by navigating from one application to another as the links become available. The HyperDAQ system has proved to be invaluable in debugging and monitoring the DAQ during full system integration tests.

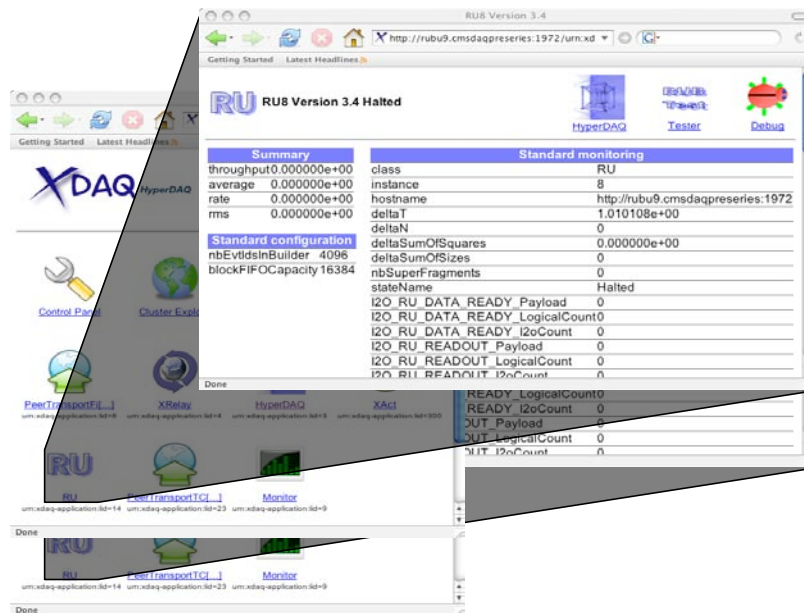


Figure 243: Example HyperDAQ page. Clicking on the RU application in the HyperDAQ page brings up monitoring information for the Read-out Unit application.

Run Control System

The Run Control System configures and controls the online applications of the DAQ components and interfaces to the Detector Control Systems. It is an interactive system furnishing diagnostic and status information and providing the entry point for control. There are $O(10000)$ applications to manage, running on $O(1000)$ PCs.

The Run Control structure is organized into eleven different sub-systems, with each sub-system corresponding to a sub-detector or self-contained component, e.g. the Hadron Calorimeter, central DAQ or global trigger. The Run Control and Monitoring System (RCMS) framework, provides a uniform API to common tasks like storage and retrieval of process configuration data, state-machine models for process control and access to the online monitoring system.

Run Control applications and services are implemented in Java as components of a common web application "RCMS" provided by the framework. The Run Control is designed as a scalable and distributed system to run on multiple machines, thus the system can be easily expanded by adding additional hardware.

In RCMS the basic unit is the Function Manager (Sect. 10.7). The interfaces are specified with the Web Service

Description Language (WSDL) using the Axis [222] implementation of Web Services (WS). Various Web Service clients including Java, LabView and Perl have been implemented to provide access to the Run Control services. The publicly available official reference implementation of the Java Servlet technology Tomcat [223], by the Apache Software Foundation, has been chosen as the platform to run the Run Control web-applications. For persistency both Oracle and MySQL are supported by the RCMS framework.

For the baseline DAQ system, ten PCs running Linux are sufficient to control the experiment. A special copy of the XDAQ executive, the job control, is always running on each online node to accept SOAP commands from the run control to start and configure additional XDAQ executives. One common database (Oracle) is shared by all online processes and RCMS installations. Configuration management across sub-systems is achieved using global configuration keys (Sect. 10.7).

The services and tools provided by RCMS comprise:

- Function Manager Framework;
- Resource and Account Management Services;
- Configurator;
- Log Message Collector.

In the following a few key components of Run Control are discussed.

Function Manager

A hierarchical control tree, with a central controller on the top and one or more controllers for each sub-system, structures the flow of commands and state information. The controllers are written using a common design paradigm, the so-called “Function Manager” (FM).

The FM has available a finite state machine, an interface for remote invocation, and a set of services to start, configure and control remote processes and to access configuration information from a DBMS. The FM is the basic element in the control tree. A standardized state machine model has been adopted by the sub-system for the first level of FMs which are directly steered by the central controller in the control tree (Fig. 244).

Resource and Account Management Services

The Resource Service (RS) stores all information necessary to start and configure the online processes of the DAQ and sub-detectors. The data is represented as Java objects which are made persistent in the DBMS both as blobs and optionally as relational tables. RCS views the experiment as a collection of configurations, where a configuration is one or more groups of resources and one or more function manager implementations for control. The configuration is specific to each sub-system. Each sub-system has its own schema instance of the RS in the DBMS. The resource definition of a run is then the set of configurations of all participating sub-systems. The configuration of a given sub-system is resolved via a key mechanism. The sub-systems register a “configuration key” to a given global key identifying the configuration of the global run. All changes to configurations and global keys are versioned and trackable.

Users have to be authenticated to get access to the RCS resources. The resource service manages the configurations based on RCMS user accounts. Multiple configurations by multiple users can be run simultaneously in the same instance of the RCMS web-application.

Configurator

In order to create central DAQ configurations for different data taking scenarios, the CMS DAQ configurator application has been developed. The configurations can be tailored for reading out specific sub-sets of FEDs with specific throughput requirements, using different FED-Builder configurations and different sub-sets of the available event builder hardware. The configurations can be adapted to different versions and parameter settings of the online software components. The process of parametrising tens of thousands applications on thousands of hosts is largely simplified by factorizing the structure of the DAQ system and of the software settings. The structural representations of the DAQ configurations are stored in the CMS DAQ hardware and configuration database which also holds a representation of the available hardware components and their connectivity. Templates of software

parameters for all software components are stored in a separate software template database. The CMS DAQ configurator application reads from both databases. It automatically calculates application parameters such as those depending on the connectivity information of the underlying hardware and creates the Java objects of the Resource Service. XML configuration files for the XDAQ executive processes are generated from these Java objects and stored in the Resource Service database. The CMS DAQ configurator application can also be used to generate configurations for test-bed hardware.

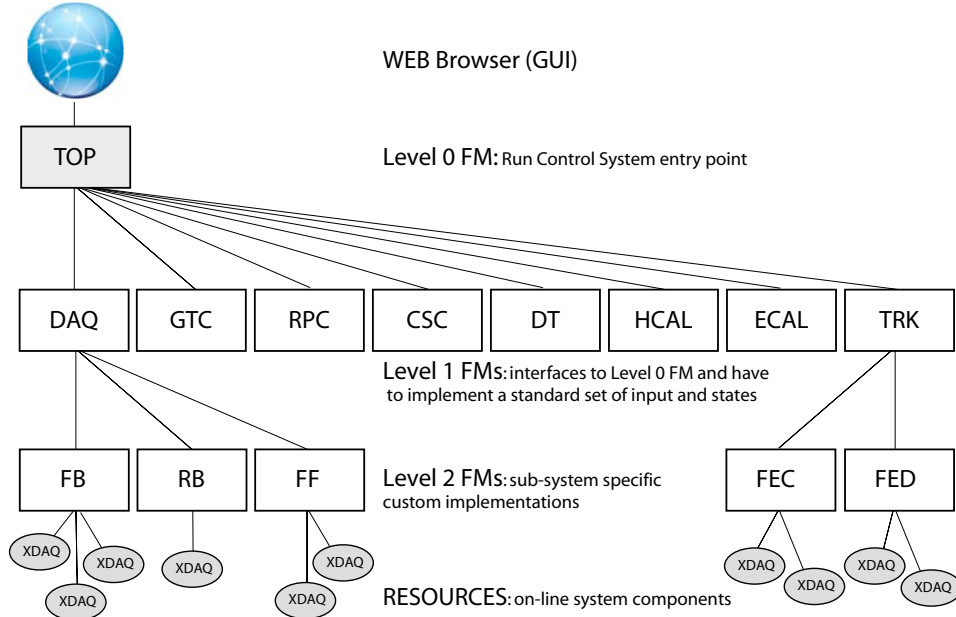


Figure 244: The RC hierarchy showing the full DAQ system. The Top Function Manager controls the next layer (Level 1) of Function Managers who in turn control the Level 2 (sub-detector level) Function Managers. The sub-detector Function Managers are responsible for managing the online system component resources.

Log Message Collector

The Log Message Collector (LMC) is a web application to collect logging information from log4j and log4c compliant applications. The LMC has receiver modules for log4c messages used with C++ applications and log4j messages used with Java applications in XML and in binary format. Appender modules are implemented for TCP socket, TCP socket hub and JMS connections. Log messages can be stored on files with a File Appender, or in a DBMS with a DB Appender. Appenders can be active concurrently. Log messages are filtered by severity in the appender modules.

Each subsystem has its own instance of a LMC. A central LMC concentrates the messages of the subsystems and forms the entry point for the visualization client of messages, e.g. the Apache Chainsaw log message viewer.

10.8 Detector Control System

Function

The main purpose of the Detector Control System (DCS) is to ensure the correct operation of the CMS experiment, so that high quality data is taken with the apparatus. The scope of DCS includes all subsystems involved in the control and monitor of the detector, its active elements, the electronics on and off the detector and the overall environment.

The Detector Control Systems of individual sub-detectors are connected to the central DCS Supervisor (Fig. 245) for combined operation. These sub-detector DCS subsystems handle all the individual detector electronics such as the CAEN high-voltage power supplies and other electronics both commercial and custom made. The low-voltage system and the gas system are common for all sub-detectors whereas the cooling systems are built individually

by each sub-detector. Additional components such as front-end detector read-out links are also monitored by the DCS.

The DCS provides both bookkeeping of detector parameters (Table 28) and safety-related functions, including alarm handling and limiting the control of critical components via a software access control. The alarm handling and automated actions are designed in a way to anticipate major problems that would otherwise initiate Detector Safety System (DSS) actions, and warn the operator in advance that some action is needed. The alarm handling includes an SMS (mobile “text messaging”) system that warns DCS users (for example a sub-detector expert) about abnormal system parameters. These SMS messages may require an acknowledgment by replying to the received alert SMS, and the status of both the alerts and acknowledgment’s is displayed in the control room so that the operators in the control room are aware that experts are investigating the alarms. The DCS also collects relevant information from DSS. Monitoring of DCS parameters is possible via the Oracle Portal web pages that allow users to analyse both real time and archived data.

The DCS has to communicate with external entities as well, in particular the DAQ run control, and serves as the interface between the CMS experiment and the LHC accelerator. Many of the features provided by the DCS are needed at all times, and as a result selected parts of the DCS must function continually on a 24-hour basis during the entire year. To ensure this continuity UPS and redundant software and hardware systems are implemented in critical areas, however even non-critical nodes can be recovered in the order of minutes thanks to a CMS specific automated software recovery system.

Implementation

The DCS software is based on the Siemens commercial SCADA package PVSS-II and the CERN Joint Controls Project (JCOP) framework [224]. Industrial controls hardware is interfaced by PVSS-II via various supported drivers OPC (OLE for Process Automation) [225] protocol, Siemens S7, SNMP, or Modbus. The JCOP framework provides common solutions to similar problems across all LHC experiments. This framework includes PVSS-II components to control and monitor the most commonly used commercial hardware (CAEN and Wiener) as well as control for additional hardware devices designed at CERN like the widely used ELMB (Embedded Local Monitoring Board). It also provides a Data Interchange Protocol (DIP). For hardware not covered by JCOP, PVSS-II offers the possibility of implementing new drivers and components, and CMS has developed sub-detector specific software.

The control application behaviour of all sub-detectors and support services are modeled as Finite State Machine (FSM) nodes, using the FSM toolkit provided by the JCOP framework. The detector controls are organized in a tree-like FSM node hierarchy representing the logical structure of the detector, where commands flow down and states and alarms are propagated up (Fig. 245). The different control systems of the experiment have been integrated into a single control tree, whose top node is referred to as the CMS DCS Supervisor. CMS has put policies into place to ensure a homogeneous and coherent use of the DCS [226].

The DCS is a distributed system and comprises all control applications dedicated to sub-systems, communicating via the PVSS proprietary network protocol. In total there will be around 100 PCs with the majority of them running Microsoft Windows, although Linux is also supported.

PVSS-II includes a proprietary database that is used to store in real time the values of all the parameters defining the current state of the system (e.g. high-voltage settings, alarms, etc.). The configuration of PVSS-II itself is also stored in this database. For static and large amounts of data, an external Oracle database is used to store configuration data, and to archive measured values of parameters from PVSS to Oracle tables. Selected data from DCS is exported to the CMS conditions database, which contains all the data describing the detector environment needed for the offline reconstruction. The DCS access control system uses the LDAP and Oracle identity management tools which has web support for account management.

During normal physics data taking the DCS will act as a slave to run control and will therefore have to receive commands and send back status information. A communication mechanism between DCS and external entities is provided by the CMS specific PVSS SOAP interface (PSX). The PSX is a SOAP server implemented with XDAQ (Sect. 10.7) using the PVSS native interface and JCOP framework, and allows access to the entire PVSS-II system via SOAP.

Table 28: Summary of detector parameters that are specific to each sub-detector.

Sub-detector	Monitored Parameters	Drivers	PCs / PLCs
muon CSCs	HV: 11 k channels, 218 k params LV: 8 k params Peripheral crate controller: 24 k params	Wiener and CAEN OPC custom: HV controller, Peripheral crate ctr	16
muon DTs	HV: 15 k channels - 110 k params LV: 1100 channels - 10 k parameters trigger: ≈ 12 k params	CAEN OPC CAEN OPC custom	7 5
muon RPCs	HV: 900 channels LV: 1850 channels sensors: 360 channels Front end: 50 k params	CAEN OPC CAEN OPC custom	2 4
HCAL	HV: 450 channels, 3500 params LV: 270 channels, ≈ 1600 params HSS: 220 params Front end: 21 k params	custom CAEN OPC custom	
ECAL	LV: ≈ 4500 params HV: ≈ 2600 params Cooling: ≈ 200 params FE monitoring: ≈ 80 k params ESS: ≈ 600 params PTM/HM: ≈ 1 k params	CAN (Wiener) CAEN OPC Simatic S7, PSX custom: PSX Simatic S7 CANopen-ELMB OPC	3 4 2 / 1 1 1 / 2 2
Strip tracker	HV: 4.4 k channels LV: 356 channels 111 k params Temperature: 1200 sensors	CAEN OPC, Siemens S7 custom: PSX	10
Pixel tracker	HV: 192 channels, 384 params LV: 192 channels, 384 params Temperature: 200 sensors	CAEN OPC, Siemens S7 custom: PSX	2 / 1
Alignment	LV: 200 channels 10 k LEDs, 150 lasers ≈ 2 k sensors	CAEN OPC ELMB, custom	5
Central DCS Services	rack control: 10 k params LHC interface: 1 k params cooling and ventilation DSS	Wiener OPC, CAN, SNMP custom: PSX	15

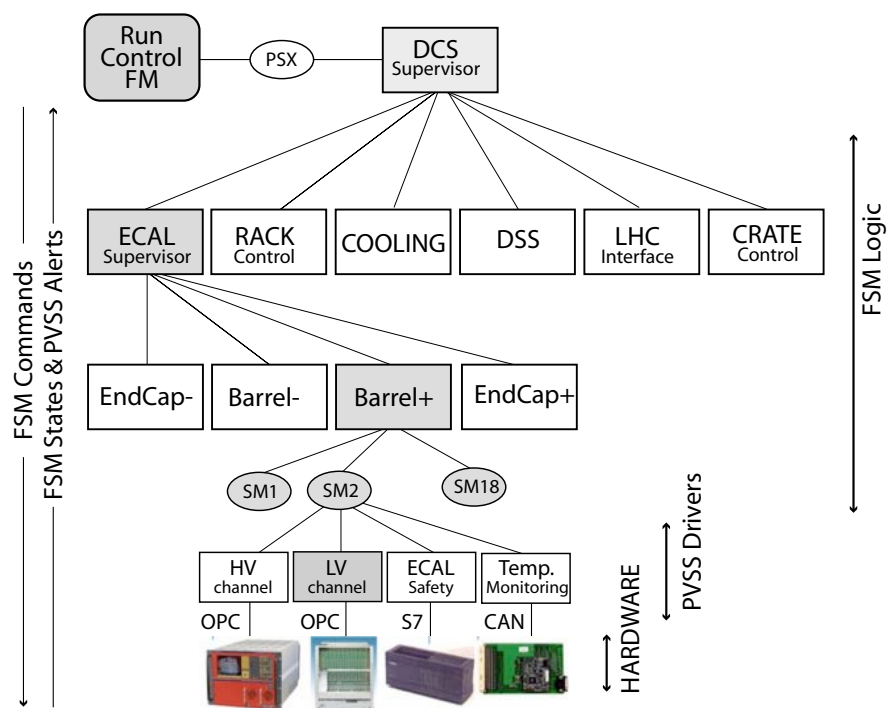


Figure 245: Outline of the Detector Control System hierarchy. Shown are all global services and ECAL as an example of a sub-detector control.

11 Detector infrastructures and safety systems

The common term *infrastructures* includes very different systems, ranging from basic site facilities to more detector-specific and safety-related services. In this Section, the main general systems are described.

11.1 Detector powering

CMS, like any other modern particle physics detector, needs considerable electrical power for its front-end electronics (FEE), for electronics racks in counting rooms and in site control centres, and finally for auxiliary services (cranes, ventilation and cooling stations, lifts and access facilities, etc.). Different power sources are available on site. Uninterruptible Power Systems (UPS), for valuable equipment that must stay on in case of power disruption, secure power for specific users for a short period, before being backed-up by a diesel engine. Common users are connected to standard network power. Table 29 gives an overview of the power requirements for CMS.

Table 29: Power requirements for CMS.

System	Power (kW)
General site services	2200
Electronics racks	2300
Low voltage to front-end electronics	1000
Magnet and cryogenics	1250
Ventilation stations	1250
Surface cooling stations	4000
Underground cooling stations	1500
Total steady-state consumption	9000

With the exception of the cooling stations, the racks system is the most important client in terms of power. Rows of racks are fed by power bus-bars. Each single rack has a main breaker piloted by a dedicated PLC, the whole system being located in a box attached to the bus-bar. A power and a control cable run from this box to the power distributor cabinet inside each rack. Cabinets provide single-phase, three-phase or three-phase + neutral current distribution. The breaker PLC is controlled by the Detector Control System (DCS) via a network connection. Single racks can be switched on-off upon DCS request and the status of each breaker is known by DCS as well. Moreover, a hardwired connection to the Detector Safety System (DSS) secures the system in case of smoke or a high temperature is detected inside a rack. Figure 246 describes the logic behind the power controls.

11.2 Detector cooling

11.2.1 Front-end electronics cooling

The CMS front-end electronics dissipates some 800 kW in the experimental cavern. This huge amount of heat is intercepted by cooling water at 18 °C for the ECAL, HCAL and Muon systems, and by C_6F_{14} fluid at temperature ranging from −15 °C to −25 °C for the Preshower, Pixel and Tracker systems. In addition, some 1600 kW are dissipated by the rack system. Table 30 shows the power dissipated by each system.

Table 30: Cooling power for different sub-systems.

System	Power (kW)
Muon Endcaps	100
Muon Barrel	50
HCAL and Yoke Barrel	60
ECAL	300
Rack system	1600
Tracker, Pixel, Preshower	150

Chilled water at 14 °C is produced at the surface in the SU5 building and then transferred to the USC55 cooling plant, where five independent water circuits, each one with its own heat-exchangers, pumps and controls, produce and distribute water at 18 °C to the experiment cavern. The Tracker, Pixel, and Preshower systems have on their primary side chilled water at 6 °C, and they have their own cooling cabinets in UXC55 to shorten the transfer lines.

system, type, length, starting point, endpoint and routing. Some 30 000 cables are referenced in the data-base.

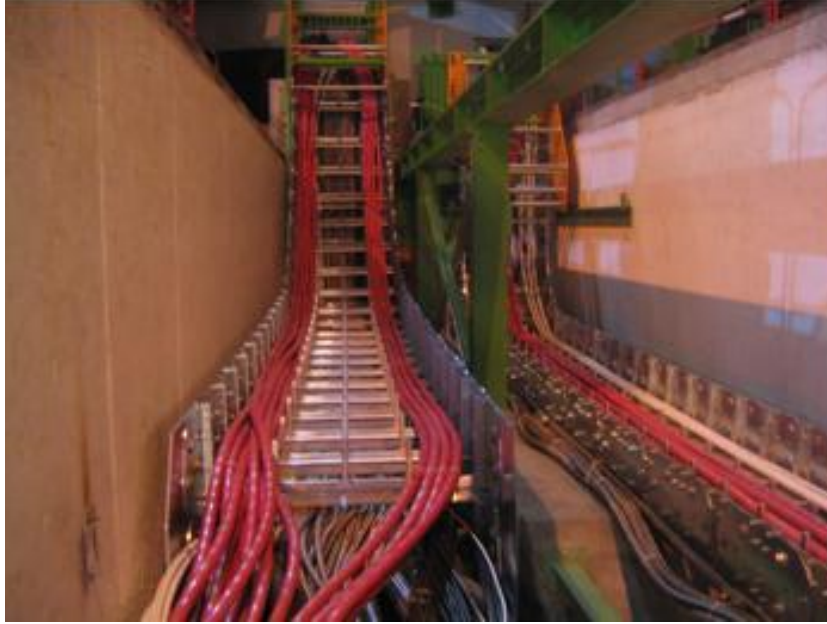


Figure 247: YB+2 and YE+1 cable-chains in UXC55 basement trenches.

11.4 Detector moving system

The CMS moving system has been designed according to the following criteria: affordability, robustness, preciseness, easyness in handling and compactness. The boundary conditions have been determined on the one hand by the weight and dimensions of the assemblies and on the other hand by the friction, the slope and the size of the cavern.

11.4.1 Sliding system

In order to limit friction and thus the power of the pulling system, CMS has chosen a heavy duty air pad system for the long movements (10 cm to 10 m) and a flat grease pad system for the final approach (up to 10 cm). In addition, these systems allow, without any additional structure, movement perpendicular to the beam. The air pads (Fig. 248) have rubber sealing rings that prevent air losses. The system can be used with compressed air bottles only. At the same time, this sealing increases somewhat the friction factor, which lies around 0.8% before moving and goes down to around 0.4% once moving has begun. The grease pad system produces a final approach with practically no friction.

11.4.2 Pulling system

The pulling system consists of a hydraulic strand jack system and includes 6 jacks with strands (of which the two in the center are pivotable) and a strand storage mandrel. Taking into account the slope of the cavern (1.234%) and the friction of the airpads and cable chain, the system must be capable of safely pulling uphill 2.5% of the maximum load, which is 2600 t (3 endcaps together). Whereas going uphill is a pure pulling, going downhill needs a retaining force in order to produce a smooth, constant movement of the load. This was integrated into the design of the hydraulic control unit.

11.5 Beam and Radiation Monitoring systems

11.5.1 Introduction

The Beam and Radiation Monitoring systems (BRM) [227] perform both a monitoring and a protection function for CMS. To this end, multiple and redundant systems have been installed, some of which can be used to initiate LHC



Figure 248: A transport beam for barrel rings with 4 air pads fixed on it.

beam aborts and/or CMS equipment control, others of which can be used for fast beam/detector optimisations. All systems will provide long term monitoring of the received radiation dose in various regions of the CMS detector.

The CMS experiment sits in an unprecedentedly high radiation field for a HEP experiment and much effort has gone into the design and construction of systems with very high radiation tolerance. Nevertheless, the LHC is designed to run with 362 MJ of stored energy in one beam and with proton intensities in excess of 10^{14} per beam. Even very small fractional losses of this beam risk causing serious damage to detector elements. Whilst the LHC itself has extensive instrumentation designed for machine protection, CMS requirements dictate that CMS must be able to detect beam-related problems as they develop and to assert beam aborts if required. In addition, CMS must be able to log data and perform post-mortem analyses in the case of accidents and understand the accumulated dosage and potential longer term damage to the detector elements. To this end CMS has implemented the BRM systems.

While radiation damage can lead to long term effects, the most likely damage scenarios involve very fast bursts of radiation/energy-dissipation in detector elements. Thus the protection systems must be sensitive to very fast changes in beam conditions; the BRM systems can detect changes at the 25 ns level, though the initially deployed protection systems will react in times of order 3–40 μ s. Additionally, the BRM systems provide monitoring and tuning data to permit operator intervention to diagnose and improve beam conditions. In addition, all BRM systems can be used to monitor integrated dose and detector component aging over the years of LHC operation.

In designing the BRM, CMS imposed several design constraints; namely to implement systems which can stay alive at any time when beam may be in the LHC independently of the state of CMS operations; that have readout and post-mortem capabilities extremely close to those of the LHC machine protection systems; and that offer a high

degree of redundancy and a wide dynamic range for protection and monitoring scenarios. Given these constraints, the BRM protection system, summarised in Table 31, has been implemented. The BRM system, its nomenclature and sub-system locations in CMS are also represented in Fig. 249.

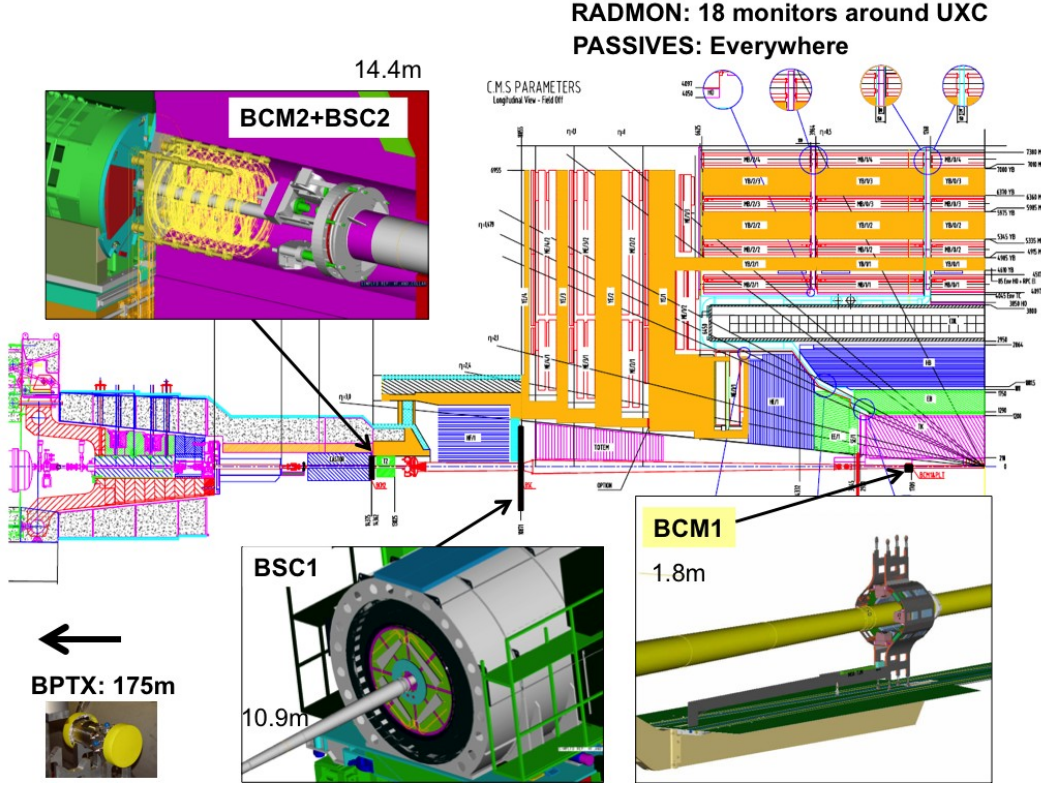


Figure 249: Layout of CMS BRM sub-systems.

11.5.2 Protection systems

The protection systems are based on chemical vapour deposition diamond detectors [228] similar to those that have been widely used in recent collider experiments [229, 230] where they have proven to be radiation hard [231], fast enough to match beam abort scenarios, and small enough to be inserted into areas close to key detector components without adding substantial material or services.

In CMS there are two protection systems foreseen for initial LHC operation. The first is the BCM1L which is four polycrystalline diamonds, each $10 \times 10 \times 0.4$ mm, positioned on either side of the IP at z values of ± 1.8 m close to the beam pipe and the inner-tracker pixel detectors (Sect. 4) at a radius of 4.5 cm. The second protection system is the BCM2L. This is a set of twelve polycrystalline diamonds, each $10 \times 10 \times 0.4$ mm, on either side of the IP behind the TOTEM T2 detector at a z position of ± 14.4 m. On each side of the IP, a set of eight sensors are deployed at an outer radius of 29 cm and an additional four at an inner radius of 5 cm. Here BCM refers to Beam Conditions Monitor, the index 1 or 2 refers to the two locations in z and the final character L indicates that these detectors are used in a leakage current measurement mode as relative flux monitors, typically integrating the leakage current over μ s time scales. The BCM1L diamonds are arranged on the x and y axes. The BCM2L comprise eight diamonds at 45° intervals at large radius and four on the x, y axes at small radius. The BCM1L and inner BCM2L diamonds measure a rate which is dominated by pp interactions at the IP. The outer BCM2L diamonds are hidden from the beam-spot and are expected to be largely sensitive to beam-halo rates.

The diamonds used for BCM1L and BCM2L are essentially identical, but the two systems differ in the readout methods adopted. The BCM2L uses a completely standard LHC Beam Loss Monitor (BLM) electronics and data processing [232, 233] that is read out asynchronously with the LHC machine with 40μ s sampling. The BCM1L readout uses the same LHC BLM back-end electronics, but uses an additional mezzanine card to provide sub-orbit sampling. The readout is synchronised with the $89\text{-}\mu$ s LHC orbit, allowing user-configurable sampling, so that the sampling can be matched to the LHC bunch trains. In addition the BCM1L allows sampling of the LHC abort gap, which must be kept empty to avoid a spray of particles being directed at CMS during a beam dump.

sub-system (Sensor type)	Location Distance from IP (m)	Sampling Time	Function	Readout + Interface LHC or CMS type	Number of Sensors
Passives (TLD+Alanine)	CMS and UXC	\sim months	Monitoring	N/A	Many
RADMON (RadFets+SRAM)	CMS and UXC	1 s	Monitoring	Standard LHC	18
BCM2L (Polycrystalline Diamond)	Behind TOTEM T2 $z=\pm 14.4$ m	$40 \mu\text{s}$	Protection	Standard LHC	24
BCM1L (Polycrystalline Diamond)	Pixel Volume $z=\pm 1.8$ m	$5 \mu\text{s}$	Protection	Standard LHC	8
BCM2F (Polycrystalline Diamond)	Behind TOTEM T2 $z=\pm 14.4$ m	\sim ns	Monitoring	CMS Standalone	8
BSC (Scintillator Tiles)	Front of HF $z=\pm 10.9$ m	\sim ns	Monitoring	CMS Standalone	32
BCM1F (Single Crystal Diamond)	Pixel Volume $z=\pm 1.8$ m	\sim ns	Monitoring	CMS Standalone	8
BPTX (Button Beam Pickup)	Upstream of IP5 $z=\pm 175$ m	200 ps	Monitoring	CMS Standalone	2

Table 31: The sub-systems to be deployed as part of the initial BRM. The table is ordered from top to bottom in increasing time resolution.

Using a set of thresholds in the readout systems and a combinatorial logic to reduce sensitivity to individual noise events, a hardware beam abort signal can be generated and transmitted to the LHC machine via the Beam Interlock System [234], leading to the dumping of the beams within 3 orbits. A lower threshold value can be used to send hardware signals to CMS sub-detector clients to initiate high and/or low voltage ramp-downs.

In the event of a beam abort initiated by CMS, or any of the other LHC (or experiment) protection systems, a full history of the BCM1L and BCM2L signals is produced and transmitted to the LHC control room.

11.5.3 Monitoring systems

Several monitoring systems are listed in Table 31; the BCM1F and BCM2F are also based upon diamond sensors, but with readouts able to resolve the sub-bunch structure, the Beam Scintillator Counters (BSC) are a series of scintillator tiles designed to provide hit and coincidence rates, the BPTX is designed to provide precise information on the bunch structure and timing of the beam, and the RADMON and Passives systems give calibrated information on the radiation field within the CMS cavern.

The BCM1F, BSC and BPTX are sensitive to time structure below the 25-ns level; as such they also provide technical trigger inputs into the global CMS trigger. In particular, the inputs from the BPTX and BSC provide zero- and minimum-bias triggers, respectively. Additionally, all three of these systems are sensitive to all machine intensities including the LHC pilot beam, where a single low intensity bunch is injected for studies or to confirm parameter settings prior to full intensity injection.

The BCM1F consists of four single crystal diamonds, each $5\times 5\times 0.5$ mm, positioned on either side of the IP at z values of ± 1.8 m at a radius of 4.5 cm, in close proximity to the BCM1L detectors. The purpose of the BCM1F is as a diagnostic tool to be able to flag problematic beam conditions resulting in “bursts” of beam loss over very short periods of time. Such beam losses are expected to be one of the principle damage scenarios for the CMS detector systems. The location for the BCM1F is close to the optimal position in terms of timing separation between ingoing and outgoing particles from the IP (i.e. 6.25 ns from the IP). The gated rate information from the BCM1F should therefore give a very good handle on the comparative rate of background from beam halo to that from luminosity products. The sensor is connected to the JK16 radiation hard amplifier [235], after which the signal is transmitted to the counting room over an analog optical link built from the tracker optical components [236].

The detector and sensor is sensitive to one MIP and has a timing resolution for single hits of a few ns. The performance of the front end electronics is shown in Fig. 250. Good separation can be seen between the signal and the noise. The pulse height was found to saturate at 100 V bias voltage across the sensor. The back-end readout produces rate, multiplicity, timing and coincidence information independently of the CMS DAQ. However there is the possibility to feed information into the event stream via a standard CMS SLINK.

In a similar vein to the BCM1F, the BCM2F is composed of four diamonds at the BCM2L location read out by a fast digitiser. The aim of this system is to provide additional diagnostic information at this location, as the digitiser can sample at 1 GHz, hence giving information on the sub-bunch level [237]. Whilst this will not be MIP-sensitive, it will help resolve the timing structure of periods of enhanced background.

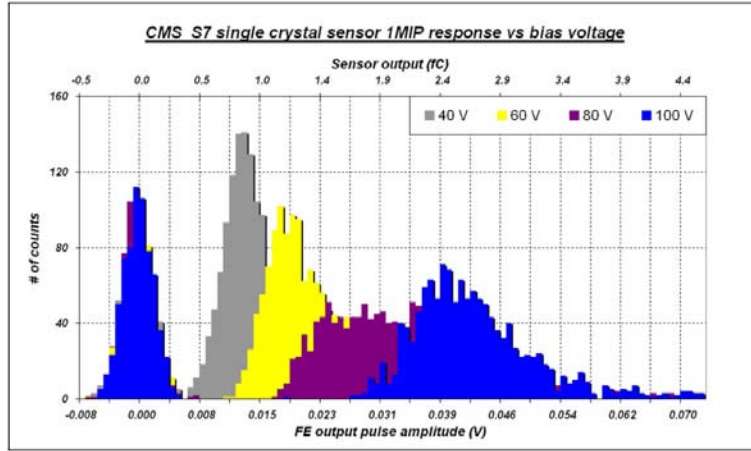


Figure 250: MIP response of BCM1F single-crystal diamond with front-end electronics as a function of bias voltage of the sensor. The superposition of histograms around 0-V output amplitude indicates the noise.

Beam Scintillator Counters (BSC) are a series of scintillator tiles designed to provide hit and coincidence rates, with a design similar to those used at previous experiments [238]. The scintillators and PMTs used for the BSC are recycled from OPAL [239]. The layout and geometry of the scintillator tiles are shown in Fig. 251. The BSC1 is located on the front of the HF, at ± 10.9 m from the IP, and consists of two types of tiles. Next to the beampipe are the disks, segmented into 8 independent slices in ϕ with an inner radius of between 22 cm and an outer radius of 45 cm. The primary function of the disks is to provide the rate information corresponding to the beam conditions. In addition, there are four large area “paddles” further out, at a radial distance of between ≈ 55 cm and ≈ 80 cm, which in addition to providing rate information, will also provide coincidence information which can be used to tag halo muons passing through the detector for calibration purposes. The area covered by the BSC is about 25% of the tracker; therefore these tiles can be indicative of activity within this bunch crossing, and can be used to provide a minimum-bias trigger for commissioning and systematic studies as required. The BSC2 is located behind TOTEM T2 at ± 14.4 m from the IP. The BSC2 consists of two tiles on each side of the IP, with a minimum inner radius of between 5 cm and a maximum outer radius of 29 cm. The primary function of the BSC2 is to distinguish between ingoing and outgoing particles along the beamline, as there is a 4-ns timing difference between them. The rates at this location can therefore be gated as to whether they are incoming (beam halo only) or outgoing (collision products and beam halo).

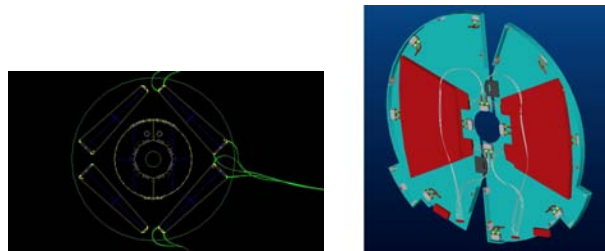


Figure 251: Layout of the Beam Scintillator Counters tiles. The left-hand panel shows the layout for BSC1, the right-hand panel for BSC2. The locations of the BCM2 sensors can also be seen in the right-hand panel.

The Beam Timing for the experiments (BPTX) is a beam pickup device specifically installed to provide the experiments with the timing structure of the LHC beam. This beam pickup is a standard button monitor used everywhere around the LHC ring for the beam position monitors. Two are installed for CMS; 175 m left and right upstream of the IP. At this location, there are two beampipes, and so therefore the timing measurement is only of the incoming beam. To optimise the timing measurement, the four buttons (left, right, up, down) of the pickup have been electrically connected together, so there is no position information but the signal strength is maximised and hence the resolution on the timing.

An oscilloscope-based read out was chosen for the BPTX and developed in common with ATLAS [240]. The BPTX will provide accurate information on the timing and phase of each bunch and its intensity. The phases of all the experimental clocks can be compared to the measured phase of each bunch at a precision better than 200 ps. This will also allow a z position for the interaction point to be calculated from the relative phases of the BPTX measurements on opposite sides of the IP. The BPTX can also detect problems with the bunch structure to be seen, as well as measurements on the proportion of beam which has drifted into the neighbouring RF bucket.

In parallel to the oscilloscope based readout, the signals from the BPTX will also be discriminated and sent as three technical trigger inputs to the CMS global trigger. This will provide three flags on each bunch crossing as to whether: a) bunch in beam 1 is occupied; b) bunch in beam 2 is occupied; c) both beams are occupied. The flag where both beams are occupied is indicative of whether collisions can occur in this bunch crossing, and therefore provides a zero-bias trigger for commissioning of the trigger system.

At 18 locations around the CMS cavern, RADMON [241] detectors are installed. The RADMON detectors each provide well calibrated measurements at these locations of: a) the dose and dose rate using RadFETs; b) the hadron flux with energies above 20 MeV and the single event upset rate using SRAM; c) the 1-MeV-equivalent neutron fluence using pin diodes. RADMON detectors are installed all around the LHC ring, and in the experimental insertions. The RADMON detectors at CMS will be integrated into and read out via the accelerator-wide RADMON system.

The integrated radiation dose throughout the CMS cavern will be measured during each run period with passive dosimetry. This allows to map the radiation field throughout the cavern and will be used to validate the simulations of the anticipated doses. This gives an absolute scale to the other measurements. The dosimeters chosen are TLDs and Alanine.

12 Computing

12.1 Overview

The CMS offline computing system must support the storage, transfer and manipulation of the recorded data for the lifetime of the experiment. The system accepts real-time detector information from the data acquisition system at the experimental site; ensures safe curation of the raw data; performs pattern recognition, event filtering, and data reduction; supports the physics analysis activities of the collaboration. The system also supports production and distribution of simulated data, and access to conditions and calibration information and other non-event data.

The users of the system, and the physical computer centres it comprises, are distributed worldwide, interconnected by high-speed international networks. Unlike previous generations of experiments, the majority of data storage and processing resources available to CMS lie outside the host laboratory. A fully distributed computing model has therefore been designed from the outset. The system is based upon Grid middleware, with the common Grid services at centres defined and managed through the Worldwide LHC Computing Grid (WLCG) project [242], a collaboration between LHC experiments, computing centres, and middleware providers.

The nature of the CMS experimental programme poses several challenges for the offline computing system:

- The requirement to analyse very large statistics datasets in pursuit of rare signals, coupled with the fine granularity of the CMS detector, implies a volume of data unprecedented in scientific computing. This requires a system of *large scale*, supporting efficient approaches to data reduction and pattern recognition.
- The system is required to be *highly flexible*, allowing any user access to any data item recorded or calculated during the lifetime of the experiment. A software framework is required which supports a wide variety of data processing tasks in a consistent way, and which must evolve along with the goals of the experiment. Since the CMS programme centres on discovery of new phenomena, under new experimental conditions, analysis requirements cannot be wholly defined in advance.
- A complex distributed system of such large scale must be designed from the outset for *manageability*, both in the operation of computing resources for physics, and in terms of software construction and maintenance. The *longevity* of the system, of 15 years or more, implies several generations of underlying hardware and software, and many changes of personnel, during the lifetime of the system.

Key components of the computing system include:

- An event data model and corresponding application framework;
- Distributed database systems allowing access to non-event data;
- A set of computing services, providing tools to transfer, locate, and process large collections of events;
- Underlying generic Grid services giving access to distributed computing resources;
- Computer centres, managing and providing access to storage and CPU at a local level.

At each level, the design challenges have been addressed through construction of a modular system of loosely coupled components with well-defined interfaces, and with emphasis on scalability to very large event samples [243].

12.2 Application Framework

The CMS application software must perform a variety of event processing, selection and analysis tasks, and is used in both offline and online contexts. The software must be sufficiently modular that it can be developed and maintained by a large group of geographically dispersed collaborators. The chosen architecture consists of a common framework which is adaptable for each type of computing environment, physics modules which plug into the framework via a well-defined interface, and a service and utility toolkit which decouples the physics modules from details of event I/O, user interface, and other environmental constraints [212].

The central concept of the CMS data model is the *Event*. The Event provides access to the recorded data from a single triggered bunch-crossing, and to new data derived from it. This may include raw digitised data, reconstructed

products, or high-level analysis objects, for real or simulated crossings. The Event also contains information describing the origin of the raw data, and the provenance of all derived data products. The inclusion of provenance information allows users to unambiguously identify how each event contributing to a final analysis was produced; it includes a record of the software configuration and conditions / calibration setup used to produce each new data product. Events are physically stored as persistent ROOT files [244].

The Event is used by a variety of *physics modules*, which may read data from it, or add new data, with provenance information automatically included. Each module performs a well-defined function relating to the selection, reconstruction or analysis of the Event. Several module types exist, each with a specialised interface. These include: *event data producers*, which add new data products into the event; *filters* used in online triggering and selection; *analysers*, producing summary information from an event collection; and *input and output modules* for both disk storage and DAQ.

Modules are insulated from the computing environment, execute independently from one another, and communicate only through the Event; this allows modules to be developed and verified independently. A complete CMS application is constructed by specifying to the Framework one or more ordered sequences of modules through which each Event must flow, along with the configuration for each. The Framework configures the modules, schedules their execution, and provides access to global services and utilities (Fig. 252).

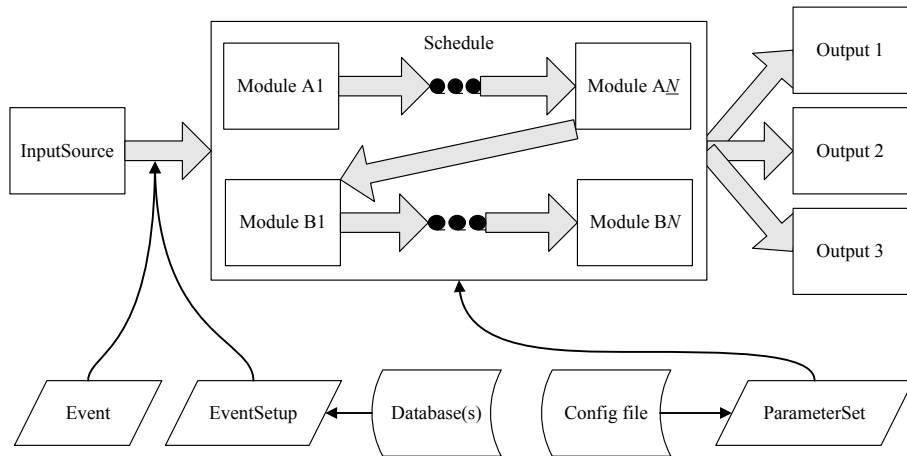


Figure 252: Modules within the CMS Application Framework.

12.3 Data formats and processing

In order to achieve the required level of data reduction, whilst maintaining flexibility, CMS makes use of several event formats with differing levels of detail and precision. Other specialised event formats are used for heavy-ion data. The process of data reduction and analysis takes place in several steps, typically carried out at different computer centres.

RAW format

RAW events contain the full recorded information from the detector, plus a record of the trigger decision and other metadata. RAW data is accepted into the offline system at the HLT output rate (nominally 300 Hz for pp collisions). An extension of the RAW data format is used to store the output of CMS Monte Carlo simulation tools. The RAW data is permanently archived in safe storage, and is designed to occupy around 1.5 MB/event (2 MB/event for simulated data, due to additional Monte Carlo truth information).

The RAW data will be classified by the online system into several distinct *primary datasets*, based upon the trigger signature. Event classification at the earliest possible stage has several advantages, including the possibility of assigning priorities for data reconstruction and transfer in the case of backlog, and balancing of data placement at centres outside CERN. CMS will also define one or more flexible “express streams” used for prompt calibration and rapid access to interesting or anomalous events.

RECO format

Reconstructed (RECO) data is produced by applying several levels of pattern recognition and compression algorithms to the RAW data. These algorithms include: detector-specific filtering and correction of the digitised data; cluster- and track-finding; primary and secondary vertex reconstruction; and particle ID, using a variety of algorithms operating on cross-detector information.

Reconstruction is the most CPU-intensive activity in the CMS data processing chain. The resulting RECO events contain high-level *physics objects*, plus a full record of the reconstructed hits and clusters used to produce them. Sufficient information is retained to allow subsequent application of new calibrations or algorithms without recourse to RAW data, though basic improvements in pattern recognition or event formats will probably require re-production of the RECO data at least once per year. The reconstruction process results in a data reduction of around a factor of six, with RECO events designed to occupy around 0.25 MB/event (0.4 MB/event for simulated data).

AOD format

AOD (Analysis Object Data) is the compact analysis format, designed to allow a wide range of physics analyses whilst occupying sufficiently small storage so that very large event samples may be held at many centres. AOD events contain the parameters of high-level physics objects, plus sufficient additional information to allow kinematic refitting. This format will require around 50 kB/event, small enough to allow a complete copy of the experimental data in AOD format to be held at computing centres outside CERN. AOD data is produced by filtering of RECO data, either in bulk production, or in a skimming process which may also filter a primary dataset into several analysis datasets.

Non-Event data

In addition to event data recorded from the detector, a variety of *non-event data* is required in order to interpret and reconstruct events. CMS makes use of four types of non-event data: construction data, generated during the construction of the detector; equipment management data; configuration data, comprising programmable parameters related to detector operation; and conditions data, including calibrations, alignments and detector status information. We concentrate here on the lattermost category.

Conditions data are produced and required by both online and offline applications, and have a well-defined interval of validity (IOV). For instance, calibration constants for a given run may be derived from prompt reconstruction of a subset of recorded events, and then used both by the HLT system and for subsequent reconstruction and analysis at computing centres around the world. Non-event data are held in a number of central Oracle databases, for access by online and offline applications. New conditions data, including calibration and alignment constants produced offline, may be replicated between the databases as required. Conditions data access at remote sites takes place via the FroNTier system [245] which uses a distributed network of caching http proxy servers.

12.4 Computing centres

The scale of the computing system is such that it could not, even in principle, be hosted entirely at one site. The system is built using computing resources at a range of scales, provided by collaborating institutes around the world. CMS proposes to use a hierarchical architecture of Tiered centres, similar to that originally devised in the MONARC working group [246], with a single Tier-0 centre at CERN, a few Tier-1 centres at national computing facilities, and several Tier-2 centres at institutes. A representation of the dataflow between centres is shown in Fig. 253.

The CMS computing model depends upon reliable and performant network links between sites. In the case of transfers between Tier-0 and Tier-1 centres, these network links are implemented as an optical private network (LHC-OPN) [247]. Data transfers between Tier-1 and Tier-2 centres typically takes place over general-purpose national and international research networks.

Tier-0 centre

A single Tier-0 centre is hosted at CERN. Its primary functions are to:

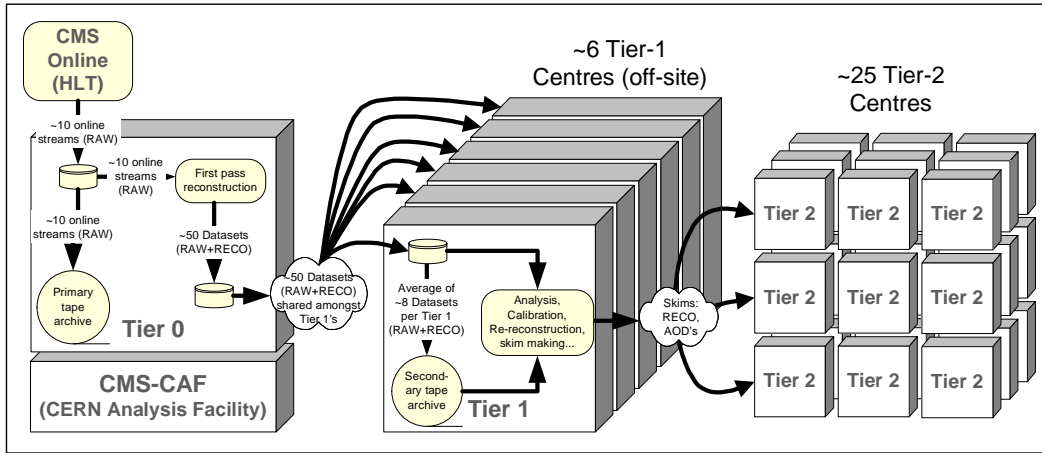


Figure 253: Dataflow between CMS Computing Centres.

- Accept data from the online system with guaranteed integrity and latency, and copy it to permanent mass storage;
- Carry out prompt reconstruction of the RAW data to produce first-pass RECO datasets. The centre must keep pace with the average rate of data recording, and must provide sufficient input buffering to absorb fluctuations in data rate;
- Reliably export a copy of RAW and RECO data to Tier-1 centres. Data is not considered “safe” for deletion from Tier-0 buffers until it is held at at least two independent sites. (One of these is CERN computing centre, playing the role of a Tier-1.)

During the LHC low-luminosity phase, the Tier-0 is intended to be available outside data-taking periods for second-pass reconstruction and other scheduled processing activities. High-luminosity running will require the use of the Tier-0 for most of the year. The Tier-0 is a common CMS facility used only for well-controlled batch work; it is not accessible for analysis use.

Tier-1 centres

A few large Tier-1 centres are hosted at collaborating national labs and computing centres around the world. These centres are operated by a professional staff on a 24/365 basis, with the emphasis on extremely reliable delivery of data-intensive processing services. Each site provides large batch CPU facilities, a mass storage system including a robotic tape archive, and very high speed international network links including a dedicated link to the LHC-OPN. The primary functions of a Tier-1 are to:

- Provide long-term safe storage of RAW data from CMS, providing a second complete copy outside CERN distributed across the centres. Each Tier-1 takes long-term custodial responsibility for a fraction of the CMS dataset;
- Store and serve to Tier-2 centres simulated and derived data. Each Tier-1 holds a fraction of the CMS simulated and RECO data, and a complete copy of the AOD data. It can rapidly transfers these data to any Tier-2 centre which requires them for analysis;
- Carry out second-pass reconstruction: a Tier-1 provides access to its archive of RAW data to allow reproduction of RECO datasets using improved algorithms or calibrations;
- Provide rapid access to very large data samples for skimming and data-intensive analysis: a Tier-1 can support high-statistics analysis projects which would be infeasible at a Tier-2 centre.

Since each Tier-1 centre holds unique RAW and RECO datasets, it must be capable of serving data to any CMS Tier-2. However, for the purposes of Monte Carlo data receipt and AOD data serving, the Tier-1 serves a defined set of a few “associated” Tier-2 centres, usually defined by geographical proximity.

Tier-2 centres

Several Tier-2 centres of varying sizes are hosted at CMS institutes. A Tier-2 centre typically divides its resources between the local user community and CMS as a whole. Tier-2 centres are subject to less stringent requirements on availability and data security than a Tier-1 centre, making them feasible to manage with the resources available to a typical University group. The functions of a Tier-2 centre may include:

- Support of analysis activities, including local storage of data samples transferred from Tier-1 centres, and access to a flexible CPU farm; in particular, the Tier-2 centres are designed to support final-stage analysis requiring repeated passes over a reduced dataset;
- Support of specialised activities such as offline calibration and alignment tasks, and detector studies;
- Production of Monte Carlo data, and its transfer to an associated Tier-1 centre for long term storage.

CERN Analysis Facility

In addition to the Tier-0 centre, CERN also hosts an Analysis Facility which combines flexible CPU resources with rapid access to the entire CMS dataset. This centre supports fast turn-around analysis when required, and a variety of other specialised functions (calibration, performance monitoring) related to the operation of the CMS detector. The centre effectively combines the rapid data access capabilities of a Tier-1 with the flexibility of a very large Tier-2.

12.5 Computing services

Grid computing

The integration of the resources at CMS computing centres into a single coherent system relies upon Grid middleware which presents a standardised interface to storage and CPU facilities at each WLCG (Worldwide LHC Computing Grid) site. The Grid allows remote job submission and data access with robust security and accounting. The detailed architecture of the Grid is described in the WLCG Technical Design Report [242].

A number of CMS-specific distributed computing services operate above the generic Grid layer, facilitating higher-level data and workload management functions. These services require CMS-specific software agents to run at some sites, in addition to generic Grid services. CMS also provides specialised user-intelligible interfaces to the Grid for analysis job submission and monitoring, and tools for automated steering and monitoring of large-scale data production and processing. An overview of the CMS Computing Services components is shown in Fig. 254.

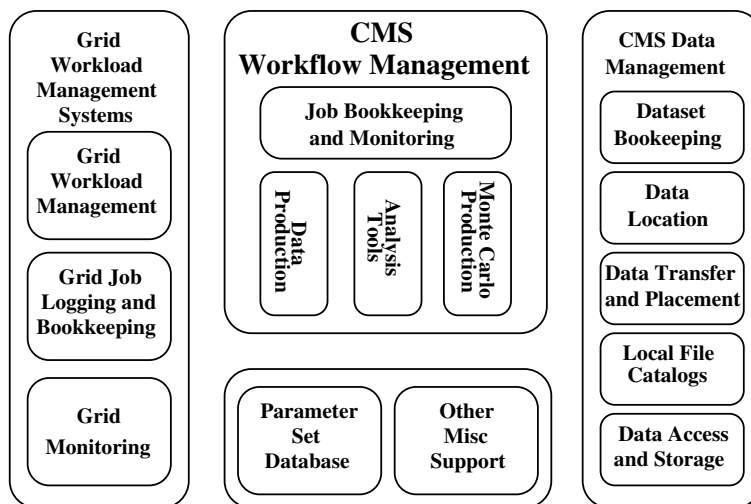


Figure 254: Overview of the CMS Computing Services.

Data management

CMS requires tools to catalogue the data which exist, to track the location of the corresponding physical data files on site storage systems, and to manage and monitor the flow of data between sites. In order to simplify the data management problem, the data management system therefore defines higher-level concepts including: *dataset*, a logical collection of data grouped by physical-meaningful criteria; *event collection*, roughly corresponding to an experiment “run” for a given dataset definition; and *file block*, an aggregation of a few TB of data files, representing the smallest unit of operation of the data transfer system.

To provide the connection between abstract datasets and physical files, a multi-tiered catalogue system is used. The *Dataset Bookkeeping System* provides a standardised and queryable means of cataloguing and describing event data [249]. It is the principle means of data discovery for the end user, answering the question “which data of this type exists in the system ?” A second catalogue system, the *Data Location Service* provides the mapping between file blocks to the particular sites at which they are located, taking into account the possibility of replicas at multiple sites. *Local File Catalogues* at each site map logical files onto physical files in local storage.

The *data transfer and placement system* is responsible for the physical movement of fileblocks between sites on demand; it is currently implemented by the PhEDEx system [248]. This system must schedule, monitor and verify the movement of data in conjunction with the storage interfaces at CMS sites, ensuring optimal use of the available bandwidth. The baseline mode of operation for the data management system is that the collaboration will explicitly place datasets at defined sites, where they will remain for access by CMS applications until removed.

Workload management

Processing and analysis of data at sites is typically performed by submission of batch jobs to a remote site via the Grid workload management system. A standard job wrapper performs the necessary setup, executes a CMSSW application upon data present on local storage at the site, arranges for any produced data to be made accessible via Grid data management tools, and provides logging information. This process is supported by several CMS-specific services.

A *parameter set management system*, implemented with either global or local scope according to the application, allows the storage and tracking of the configuration of CMSSW applications submitted to the Grid. A lightweight *job bookkeeping and monitoring system* allows users to track, monitor, and retrieve output from jobs currently submitted to and executing at remote sites [250]. The system also provides a uniform interface to a variety of Grid-based and local batch-system based submission tools. In addition, a suite of software distribution tools provide facilities for automated installation of standard CMS applications and libraries at remote sites.

Bulk workflow management

For very large-scale data processing (including Monte Carlo production, skimming and event reconstruction), a specialised bulk workflow management tool has been developed. The ProdAgent system comprises a collaborative distributed network of automated job managers, operating at Tier-0, Tier-1 and Tier-2 sites [250]. The system provides facilities for large-scale Grid job submission, interface to the CMS data catalogues and data management system, and handling of large flows of logging and status information. A highly automated system such as ProdAgent is essential in order to allow the CMS data processing system to be controlled and monitored by a moderately-sized data operations team.

User workflow management

For a generic CMS physicist, a dedicated tool (CRAB) for workflow management is available [250]. It allows to submit user-specific jobs to a remote computing element which can access data previously transferred to a close storage element. CRAB takes care of interfacing with the user environment, it provides data-discovery and data-location services, and Grid infrastructure. It also manages status reporting, monitoring, and user job output which can be put on a user-selected storage element. Via a simple configuration file, a physicist can thus access data available on remote sites as easily as he can access local data: all infrastructure complexities are hidden to him as much as possible. There is also a client-server architecture available, so the job is not directly submitted to the Grid but to a dedicated CRAB server, which, in turn, handles the job on behalf of the user, interacting with the Grid services.

12.6 System commissioning and tests

It has been recognised since the very start of preparations for LHC that the construction and organisation of the experiment computing systems would be a key challenge. Each component of the system must be designed with attention to both scalability and flexibility, and rigorously tested at realistic scale. The reliance on distributed computing, using the relatively new Grid approach, has many advantages, but adds further complexity in controlled deployment and testing compared to a system located primarily at a single site.

The relatively large cost of the computing system dictates that centres must build up their resources in a carefully controlled way; the rapidly falling price of hardware dictates that full-scale resources will only become available shortly before they are required, and that efficient use of resources is a strong requirement. The emphasis in CMS has been on a series of increasing scale full-system tests (“data challenges”) over the last three years, exercising all available components in a realistic way.

In 2006, CMS carried out a large-scale Computing, Software and Analysis challenge (CSA06) [251]. The scale of the test was set at 25% of the nominal 2008 performance, with the computing system operated continuously at this level for more than six weeks. The challenge was carried out using realistic application software and computing tools. Amongst the targets for the test were:

- Preparation of a 50 million event Monte Carlo dataset at around twenty CMS Tier-1 and Tier-2 centres in the weeks preceding the challenge, and upload to CERN;
- Playback of the MC dataset for prompt Tier-0 reconstruction at around 40 Hz, including the application of calibration constants from the offline database, and splitting the event sample into around ten datasets;
- Distribution of AOD and RAW to all Tier-1 centres, and subsequent alignment / calibration, reconstruction and skimming operations at several sites;
- Transfer of skimmed data to Tier-2 centres, and successful running of physics analysis jobs.

Overall, the key metrics for success in the challenge were met or exceeded: the reconstruction rate at the Tier-0 reached 300 Hz; an export rate of over 240 MB/s to Tier-1 centres, and a sustained dataflow within the computing system of over 800 MB/s (Fig. 255), was attained; Tier-1 site availability at six out of seven centres remained above 90%; and around 30 individual analysis exercises were carried out during the challenge. A number of important lessons for future operations were also learned. CMS will finalise its data challenge programme with additional scale tests, in the final stages of preparation at the time of writing. In parallel with data challenges, continuous programmes are under way to deploy, commission and test hardware at the computing centres, and to debug and demonstrate reliable and high-speed data network links between them.

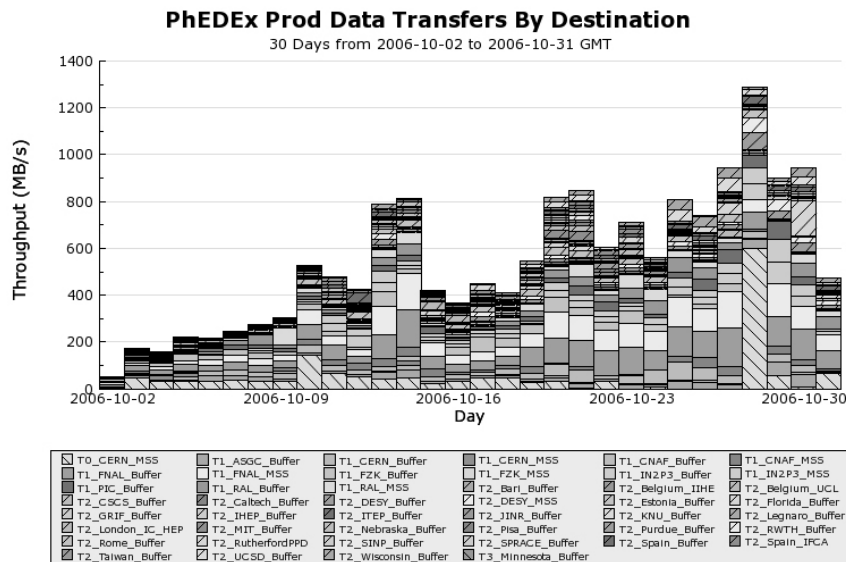


Figure 255: Dataflow during the CSA06 Data Challenge.

13 Conclusions

The Compact Muon Solenoid detector has been described in detail. The expected physics performance of the apparatus has been described elsewhere [17].

At the time of this paper, the apparatus is essentially completed and installed.

After more than 10 years of design and construction, the CMS magnet has been constructed and successfully tested. Most of the magnetic, electrical, mechanical, and cryogenics parameters measured during the tests are in good agreement with calculated values. The CMS magnet is the largest superconducting solenoid ever built for a physics experiment in terms of bending power for physics, total stored energy, and stored energy per unit of cold mass.

The silicon-strip inner tracker, with about 200 m^2 of active silicon, has been integrated into its support tube, commissioned, and thoroughly tested with cosmic rays. Its performance is excellent, fulfilling the design specifications. The silicon tracker was installed into CMS in december 2007. More than 80% of the pixel modules are completed; it is planned to install the Pixel detector into CMS in early 2008.

The ECAL, comprising over 75 000 lead tungstate crystals, is the largest crystal calorimeter ever built. The crystals in the barrel part, comprising over 60 000 crystals, have been inter-calibrated using cosmic rays and about a third in particle beams, demonstrating the ability to measure the energies ranging from those deposited by minimum ionising particles to high-energy electrons. An energy resolution of 0.5% for 120 GeV electrons has been attained. The ECAL barrel has been installed in the experiment and is being commissioned. The endcaps are foreseen to be inserted into the experiment in 2008.

The entire HCAL has been completed and commissioned on the surface. The HCAL modules are currently being commissioned in the experiment proper.

The various components of the Muon System (drift tubes, cathode strip chambers, resistive plate chambers) have been completed. A significant fraction of the Muon System has been commissioned and tested on surface with cosmic rays, and it is now being integrated into the experiment and being commissioned in-situ.

In the very forward region, the Zero Degree Calorimeter has been completed and CASTOR is expected to be completed in 2008.

The off-detector electronics are currently being installed and operations for trigger commissioning are taking place.

Common data-acquisition runs with various sub-detectors, sometimes using cosmic rays, are regularly taking place at the experiment and will continue into spring 2008 in anticipation of collisions at LHC in mid-2008.

Acknowledgements

We are indebted to the staffs at CERN and at our home institutes who made the construction of CMS possible. We acknowledge the support from: CERN; Federal Ministry of Science and Research (Austria); FNRS and FWO (Belgium); CNPq and FAPERJ (Brazil); Ministry of Education and Science (Bulgaria); CAS and NSFC (China); Ministry of Science and Technology (Croatia); University of Cyprus (Cyprus); Estonian Academy of Sciences and NICPB (Estonia); Academy of Finland, Ministry of Education and Helsinki Institute of Physics (Finland); CEA and CNRS/IN2P3 (France); BMBF and DESY (Germany); General Secretariat for Research and Technology (Greece); NKTH (Hungary); DAE and DST (India); IPM (Iran); UCD (Ireland); INFN (Italy); KICOS (Korea); CINVESTAV (Mexico); PAEC (Pakistan); State Commission for Scientific Research (Poland); GRICES (Portugal); JINR (Armenia, Belarus, Georgia, Ukraine, Uzbekistan), Ministry of Science and Technologies of the Russian Federation, Ministry of Atomic Energy (Russia); Ministry of Science and Environmental Protection of the Republic of Serbia (Serbia); Oficina de Ciencia y Tecnologia (Spain); ETHZ, PSI, University of Zurich (Switzerland); National Science Council (Taipei); TAEA (Turkey); STFC (United Kingdom); DOE and NSF (USA).

Acronyms

AB	Algorithm Board
AC	Alternating Current
ADC	Analog to Digital Converter
AFEB	Anode Front-End Board, CSC system
AG	Application Gateway
ALCT	Anode Local Charged Track trigger primitive, CSC system
AOD	Analysis Object Data - a compact event format for physics analysis
AOH	Analog Opto-Hybrid
APD	Avalanche Photo-Diode
API	Application Programming Interface
APV	Analogue Pipeline (Voltage mode), front-end read-out chip of Tracker
ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal LHC ApparatuS experiment
aTTS	Asynchronous Trigger Throttle System
AWG	American Wire Gauge
BMU	Barrel Muon system
BD	Back Disk
BP	Back Petal
BPix	Barrel Pixel
BR	Branching Ratio
BRAN	Beam RAte of Neutrals
BTI	Bunch and Track Identifier trigger primitive, DT system
BU	Builder Unit
BX	Bunch Crossing
BXN	Bunch Crossing Number
CASTOR	CentauRO And Strange Object Research
CC	Cosmic Challenge
CCS	Clock and Control System
CCU	Communication and Control Unit
CCUM	Communication and Control Unit Module
CDR	Central Data Recording
CFC	Carbon Fiber Composite
CFEB	Cathode Front-End Board of the CSC system
CLCT	Cathode Local Charged Track trigger primitive, CSC system
CMM	Coordinate Measuring Machine
CMN	Common Mode Noise
CMOS	Complementary Metal Oxide Semiconductor
CMS	Compact Muon Solenoid experiment
CMSSW	CMS SoftWare framework

COSINE	Consistent Online Software IN- tegration Environment, project integrating online with offline soft- ware
CPU	Central Processing Unit
COCOA	CMS Object-oriented Code for optical Alignment
CRAB	CMS Remote Analysis Builder
CRack	Cosmic Rack, a set of TOB rods
CRC	Cyclic Redundancy Check error detection
CSC	Cathode Strip Chamber muon sys- tem
CSCTF	Cathode Strip Chamber Trigger Track Finder
D2S	Data to Surface
DAC	Digital to Analog Converter
DAQ	Data Acquisition
DAQMB	Data Acquisition Motherboard, CSC L1 trigger
DBMS	Database Management System
DC	Direct Current
DCC	Data Concentrator Card
DCCT	DC Current Transformer
DCS	Detector Control System
DCU	Detector Control Unit
DDD	Detector Description Database
DDL	Data Description Language
DDU	Detector Dependent Unit in DAQ system
DIP	Data Interchange Protocol (CERN)
DMB	DAQ MotherBoard of CSC system
DOFZ	Diffusion Oxygenated Float Zone
DOH	Digital Opto-Hybrid
DOHM	Digital Opto-Hybrid Module
DQM	Data Quality Monitoring
DQMC	Data Quality Monitoring Collector
DSS	Detector Safety System
DT	Drift Tube muon system
DTTF	Drift Tube Trigger Track Finder, DT L1 trigger
EB	Electromagnetic Calorimeter (Bar- rel)
ECAL	Electromagnetic Calorimeter
EDM	Event Data Model
EDMS	Engineering Database Management System
EE	Electromagnetic Calorimeter (End- cap)
EIC	Electromagnetic Isolation Card, regional calorimeter trigger
ELMB	Embedded Local Monitoring Board (ECAL)
EMDB	Equipment Management DataBase
EMU	Endcap Muon system
ENC	Equivalent Noise Charge
EP	Event Processor

ES	Endcap preShower detector, also Event Server
ESP	Event Server Proxy
ESS	ECAL Safety System
ETTF	Eta Track Finder, DT regional muon trigger
EVb	Event Builder
EVF	Event Filter Farm
EVM	Event Manager
FB	FED builder
FD	Front Disk
FDL	Final Decision Logic, L1 Global Trigger
FE	Front-End
FEB	Front-End Board
FEC	Front-End Card, Front End Con- troller
FED	Front-End Driver
FES	Front-End System
FENIX	ECAL front-end read-out ASIC
FEVT	Event format comprising the union of RAW and RECO data
FF	Filter Farm
FMM	Fast Merging Module
FIFO	First In First Out buffer
FP	Front Petal
FPGA	Field Programmable Gate Array
FPix	Forward Pixel
FRL	Front-End Read-out Link
FSM	Finite State Machine
FTP	Foil screened Twisted Pair cables
FU	Filter Unit
Gb	Gigabit (10^9 bits)
GB	Gigabyte (10^9 bytes)
GBW	Gain BandWidth product
GCALEOR	Computer program for hadron shower calculations
GCT	Global Calorimeter Trigger (L1)
GIF	Gamma Irradiation Facility
GMR	Global Muon Reconstructor
GMT	Global Muon Trigger (L1)
GOH	Giga Optical Hybrid
GOL	Gigabit Optical Link
GPN	General Purpose Network (CERN campus)
GT	Global Trigger (L1)
GTFE	Global Trigger Front-end board (L1)
GTL	Global Trigger Logic board (L1)
GTL+	Gunning Transceiver Logic, up- graded version, developed by Fairchild Semiconductor
GTP	Global Trigger Processor
GTPe	Global Trigger Processor emulator
GUI	Graphical User Interface
H2	Beamline at CERN
HCAL	Hadron Calorimeter
HB	Hadron Calorimeter (Barrel)

HDI	High Density Interconnect
HE	Hadron Calorimeter (Endcap)
HF	Hadron Calorimeter (Forward)
HG	High Gain
HI	Heavy Ion(s)
HIJING	Heavy Ion Jet INteraction Genera- tor, Monte Carlo event generator for heavy-ion collisions
HIP	Hits and Impact Point alignment method, also Highly Ionizing Parti- cle
HLT	High-Level Trigger
HM	Humidity Monitoring
HO	Hadron Calorimeter (Outer Barrel)
HPD	Hybrid Photo-Diode
HTML	HyperText Mark-up Language
HTR	HCAL Trigger and Read-out
HV	High Voltage
ICC	Internal Calibration Circuit, preshower
IGUANA	Interactive Graphics for User ANAl- ysis
I²C	Inter-Integrated Circuit
ICB	InterConnect Board (TEC), Inter- Connect Bus (TOB)
ICC	InterConnect Card
I/O	Input/Output
IOV	Interval Of Validity
IP	Impact Parameter, also Impact Point or Internet Protocol
ISO	Isolation bit in muon trigger
ISR	Intersecting Storage Ring collider at CERN
JCOP	Joint controls Project at CERN
JSC	Jet Summary Card, in Regional Calorimeter Trigger
JTAG	Joint Test Action Group
kb	kilobit (10^3 bits)
kB	kilobytes (10^3 bytes)
L1	Level-1 hardware-based trigger
L1A	Level-1 Accept
L1T	Level-1 Trigger
LAN	Local Area Network
LAS	Laser Alignment System
LCG	LHC Computing Grid (a common computing project)
LCT	Local Charged Track trigger primi- tive of CSC system
LDAP	Lightweight Directory Access Pro- tocol
LED	Light Emitting Diode
LEP	Large Electron Positron collider at CERN
LG	Low Gain
LHC	Large Hadron Collider
LSB	Least Significant Bit
LUT	Lookup Table

LTC	Local Trigger Controller
LV	Low Voltage
LVD	Low Voltage Distribution
LVDS	Low Voltage Differential Signaling
LVR	Low Voltage Regulator
MA	Module Alignment
MAB	Module Alignment of Barrel
Mb	Megabit (10^6 bits)
MB	Muon system (Barrel), also Mother Board or Megabyte (10^6 bytes)
MC	Monte Carlo simulation program/technique, also Mini-Crate of DT system
ME	Muon system (Endcap) or Monitoring Element
MEM	Monitoring Electronics Module
mFEC	mezzanine Front End Controller
MGPA	Multiple Gain Pre-Amplifier chip, ECAL
MILLEPEDE	Algorithm for tracker alignment
MIP	Minimum Ionizing Particle
MOS	Metal Oxide Semiconductor
MOSFET	Metal Oxide Semiconductor Field Effect Transistor
MPC	Muon Port Card, CSC L1 trigger
MSS	Magnet Safety System
MT	Mean Time
MTCC	Magnet Test Cosmic Challenge
MTU	Maximum Transfer Unit
NIC	Network Interface Card
NIEL	Non-Ionizing Energy Loss
O2O	Online to Offline
ODBMS	Object Database Management System
OMDS	Online Master Data Storage
OPC	OLE for Process Automation
ORCOF	Offline ReConstruction Offline subset, conditions database
ORCON	Offline ReConstruction Online subset, conditions database
OS	Operating System
P5	Point 5 collision area of LHC
PACE	Preshower front end ASIC
PACT	PATtern Comparator Trigger, RPC system
PB	Petabyte (10^{15} bytes)
PC	Personal Computer
PCB	Printed Circuit Board
PCI	Peripheral Component Interconnect
PD	Pixel Detector
PDF	Parton Density Function, also Probability Distribution Function
PHTF	Phi Track Finder, DT regional muon trigger
PLC	Programmable Logic Controller (ECAL)
PLD	Programmable Logic Device

PLL	Phase-Locked Loop
PP	Patch Panel
PLT	Pixel Luminosity Telescope
PS	Proton Synchrotron
PSB	Pipeline Synchronizing Buffer, L1 Global Trigger and Global Muon Trigger
PSX	PVSS SOAP Interface
PTM	Precision Temperature Monitoring, ECAL
PV	Primary Vertex
PVSS	Prozessvisualisierungs- und Steuerungs-System
QIE	Charge Integrator and Decoder, ECAL frontend electronics
QPLL	Quartz Phase-Locked Loop
RAW	Event format from the online con- taining full detector and trigger data
RB	Read-out Unit Builder, also Re- source Broker
RCT	Regional Calorimeter Trigger (L1)
RCS	Run Control System
RECO	Event format for reconstructed objects such as tracks, vertices, jets, etc.
RH	Relative Humidity
RISC	Reduced Instruction Set Computer
RMS	Root Mean Square
ROB	ReadOut Board, DT system
ROC	ReadOut Chip, pixels
ROS	ReadOut Server board, DT system
RPC	Resistive Plate Chamber muon system
RS	Resource Service
RU	Read-out Unit
SAN	Storage Area Network
SC	Sector Collector, DT muon L1 trigger or Super Crystal, ECAL
SCA	Switched Capacitor Array buffer, CSC system
SCADA	Supervisory Control And Data Acquisition
SCX	Surface Control eXperimental building at P5
SDA	Slow Dump Accelerated
SEL	Single Event Lathcup
SEU	Single Event Upset
SFM	SubFarm Manager
Skim	Subset of events selected from a larger set
SLB	Synchronization and Link Board
SM	SuperModule (ECAL) or Storage Manager (DAQ)
SMB	System Mother Board
SMD	Surface Mounted Device
SMS	Short Message Service (mobile phones)

S/N	Signal to Noise ratio
SNMP	Simple Network Management Protocol
SOAP	Simple Object Access Protocol
SPS	Super Proton Synchrotron
SRP	Selective Read-out Processor
SST	Silicon Strip Tracker
STL	Standard Template Library
sTTS	Synchronous Trigger Throttle System
SV	Secondary Vertex
SX5	Surface hall at Point 5 for CMS
T1, T2	Tracking telescopes of TOTEM
TAG	Event index information such as run/event number, trigger bits, etc.
Tb	Terabit (10^{12} bits)
TB	Terabyte (10^{12} bytes)
TBM	Token Bit Manager
TCA	Telecom Computing Architecture
TCC	Trigger Concentrator Card
TCP	Transmission Control Protocol
TCS	Trigger Control System
TDC	Time to Digital Converter
TDR	Technical Design Report
TDSS	Tracker Detector Safety System
TEC	Tracker EndCap
TF	Track-Finder, muon L1 trigger
TIB	Tracker Inner Barrel
TID	Tracker Inner Disks
TIM	Timing Module, Global Trigger and Drift Tube Trigger Track Finder
TMB	Trigger MotherBoard, CSC L1 trigger
TN	Technical Network
TOB	Tracker Outer Barrel
TOTEM	TOTAL Elastic and diffractive cross section Measurement
TPD	Tracker Pixel Detector
TPG	Trigger Primitive Generator
TRACO	Track Correlator, DT L1 trigger
TriDAS	Trigger and Data Acquisition project
TRLB	Token Ring Link Board
TS	Trigger Server, DT L1 trigger
TSM	Track Sorter Master, DT L1 trigger
TSS	Track Sorter Slave, DT L1 trigger
TTC	Trigger Timing and Control
TTCex	TTC Encoder and Transmitter
TTCmi	TTC Machine Interface
TTCrx	TTC Receiver
TTS	Trigger Throttling System
USC55	Underground Service Cavern at Point 5 for CMS
UXC55	Underground eXperimental Cavern at Point 5 for CMS
VFE	Very Front End
VHDI	Very High Density Interconnect
VME	Versa Module Eurocard

VPT	Vacuum PhotoTriode
WAN	Wide Area Network
WLCG	Worldwide LHC Computing Grid
WLS	WaveLength Shifting
XDAQ	Software framework for CMS Data Acquisition
XML	eXtensible Markup Language
YB	Yoke (Barrel)
YE	Yoke (Endcap)
ZDC	Zero Degree Calorimeter

References

- [1] CMS Collaboration, *The Compact Muon Solenoid Technical Proposal*, CERN/LHCC 94-38 (1994).
- [2] TOTEM Collaboration, paper in this JINST issue.
- [3] CMS Collaboration, *The Magnet Project: Technical Design Report*, CERN/LHCC 97-10 (1997).
- [4] A. Hervé et al., *Status of the Construction of CMS Magnet*, 2004 IEEE Trans. Appl. Superconduct. **14** 524.
- [5] A. Hervé, *The CMS Detector Magnet*, 2000 IEEE Trans Appl. Superconduct. **10** 389.
- [6] F. Kircher et al., *Final Design of the CMS Solenoid Cold Mass*, 2000 IEEE Trans Appl. Superconduct. **10** 407.
- [7] ALEPH Collaboration, *ALEPH: a detector for electron-positron annihilation at LEP*, 1990 Nucl. Instrum. Meth. A **294** 121.
- [8] DELPHI Collaboration, *The DELPHI detector at LEP*, 1991 Nucl. Instrum. Meth. A **303** 233.
- [9] ZEUS Collaboration, *The ZEUS detector*, unpublished, available at <http://www-zeus.desy.de/bluebook/bluebook.html>.
- [10] BABAR Collaboration, *The BABAR detector*, 2002 Nucl. Instrum. Meth. A **479** 1.
- [11] B. Blau et al., *The CMS Conductor*, 2002 IEEE Trans. Appl. Superconduct. **12** 345.
- [12] S. Sgobba et al., *Mechanical performance at cryogenic temperature of the modules of the external cylinder of CMS and quality controls applied during their fabrication*, 2004 IEEE Trans. Appl. Superconduct. **14** 556.
- [13] P. Fabbicatore et al., *The Construction of Modules Composing the CMS Superconducting Coil*, 2004 IEEE Trans Appl. Superconduct. **14** 552.
- [14] P. Fazilleau et al., *Design, construction and tests of the 20kA Current Leads for the CMS Solenoid*, IEEE Trans. Appl. Superconduct. **14** 1765.
- [15] CMS Collaboration, *The Tracker Project Technical Design Report*, CERN/LHCC 98-006 (1998).
- [16] CMS Collaboration, *Addendum to the CMS Tracker TDR*, CERN/LHCC 2000-016 (2000).
- [17] CMS Collaboration, *Detector Performance and Software, Physical Technical Design Report, Volume I*, CERN/LHCC 2006-001 (2006).
CMS Collaboration, *Physics Performance, Physical Technical Design Report, Volume II*, CERN-LHCC-2006-021, 2006 J. Phys. G **34** 995.
- [18] ROSE Collaboration, *2nd RD48 Status Report*, CERN/LHCC 98-39, October 1998.
- [19] G. Bolla et al., *Design and test of pixel sensors for the CMS experiment*, 2001 Nucl. Instrum. Meth. A **461** 182.
- [20] J. Kemmer et al., *Streifendetektor*, Patentoffenlegungsschrift DE 19620081 A1, 21.11.1997, Munich, Germany
- [21] K. Arndt et al., *Silicon sensors development for the CMS pixel system*, 2003 Nucl. Instrum. Meth. A **511** 106.
- [22] Y. Allkofer et al., *Design and Performance of the Silicon Sensors for the CMS Barrel Pixel Detector*, to be submitted to Nucl. Instrum. Meth. A.
- [23] G. Lindström et al., *Radiation Hard Silicon Detectors – Developments by the RD48 (ROSE) Collaboration*, 2001 Nucl. Instrum. Meth. A **466** 308.
- [24] D. Kotlinski, *The Control and Readout System of the CMS Pixel Barrel Detector*, 2006 Nucl. Instrum. Meth. A **565** 73.
- [25] H. C. Kästli, *Design and performance of the CMS pixel detector readout chip*, 2006 Nucl. Instrum. Meth. A **565** 188.

- [26] E. Bartz, *The 0.25 μ m Token Bit Manager Chip for the CMS Pixel Readout, Proceedings of the 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany (2005).*
- [27] K. Kloukinas et al., *FEC-CCS: A common Front-End Controller card for the CMS detector electronics, Proc. of the 12th Workshop on Electronics for LHC and Future Experiments, Valencia, Spain, CERN-2007-001 179 (2006).*
- [28] Manuals CERN ASICs
<http://cmstrackercontrol.web.cern.ch/cmstrackercontrol/manuals.htm>
- [29] G. Bolla et al., *Wire-bond failures Induced by resonant vibrations in the CDF silicon detector*, Nuclear Science Symposium Conference record, 2003 IEEE **3** 1641.
- [30] D. Menasce, M. Turqueti and L. Uplegger, *The Renaissance: a test-stand for the Forward CMS Pixel Tracker assembly*, 2007 Nucl. Instrum. Meth. A **579** 1141.
- [31] S. Paoletti et al., *The Powering Scheme of the CMS Silicon Strip Tracker, 10th Workshop on electronics for LHC and future experiments, CERN 2004-010, CERN-LHCC-2004-030 (2004).*
- [32] R. Fenner and E. Zdankiewicz, *Micromachined Water Vapor Sensors: A Review of Sensing Technologies*, 2001 IEEE Sensors Journal, **1** 309.
- [33] *Statement List (STL) for S7-300 and S7-400 Programming*, Reference Manual, Siemens AG Automation and Drives, 2003.
- [34] *Interface for accessing the PVSS system through SOAP XML protocol*,
<http://xdaqwiki.cern.ch/index.php/PSX> .
- [35] L. Borrello et al., *Sensor Design for the CMS Silicon Strip Tracker*, CMS Note 2003/020 (2003).
- [36] J.-L. Agram et al., *The silicon sensors for the Compact Muon Solenoid tracker: design and qualification procedure*, 2004 Nucl. Instrum. and Meth. A **517** 77.
- [37] S. Braibant et al., *Investigation of Design Parameters for Radiation Hard Silicon Microstrip Detectors*, 2002 Nucl. Instrum. and Meth. A **485** 343.
- [38] M. Raymond et al., *The CMS Tracker APV25 0.25 μ m CMOS Readout Chip, Proc. of the 6th Workshop on Electronics for LHC Experiments, Cracow (2000).*
- [39] A. Marchioro, *Deep submicron technologies for HEP, Proc. of the 4th workshop on electronics for LHC experiments, CERN/LHCC/98-36,40-46 (1998).*
- [40] J. Troska et al., *Optical Readout and Control Systems for the CMS Tracker*, 2003 IEEE Trans. on Nucl. Sci. **50** 1067.
- [41] *Electrical Characteristics of Low Voltage Differential Signaling (LVDS) Interface Circuits*, ANSI/TIA/EIA-644-A-2001, (2001).
- [42] F. Vasey, C. Biber, M. Sugiyama and J. Troska, *A 12-Channel Analog Optical-Receiver Module*, 2005 Journal of Lightwave Technology **23** 4270.
- [43] K. Gill et al., *Progress on the CMS Tracker control system, Proc. of the 11th Workshop on Electronics for LHC and Future Experiments LECC 2005, Heidelberg (2005).*
- [44] Coughlan, J. et al., *The CMS Tracker Front-End Driver, Proc. of the 9th Workshop on Electronics for LHC Experiments, Amsterdam (2003).*
- [45] C. Ljuslin, A. Marchioro and C. Paillard, *The CCU25: a network oriented Communication and Control Unit integrated circuit in a 0.25 μ m CMOS technology, Proc. of the 8th Workshop on Electronics for LHC experiments, Colmar (2002).*
- [46] K. Kloukinas, A. Marchioro, P. Moreira and P. Placidi, *A 40 MHz clock and trigger recovery circuit for the CMS tracker fabricated in a 0.25 μ m CMOS technology and using a self calibration technique, Proc. of the 5th Workshop on Electronics for LHC Experiments, Snowmass (1999).*

- [47] Philips Semiconductors, *THE I²C-BUS SPECIFICATION*, version 2.1, document order number: 9398 393 40011, January 2000.
- [48] G. Magazzu, A. Marchioro and P. Moreira, *The Detector Control Unit: An ASIC for the Monitoring of the CMS Silicon Tracker*, 2004 IEEE Trans. Nucl.Sci. **51** 1333.
- [49] U. Goerlach, *Industrial Production of Front-End Hybrids for the CMS Silicon Tracker*, *Proc. of the 9th Workshop on Electronics for LHC experiments*, Amsterdam (2003).
- [50] M. Axer et al., *The Qualification of Silicon Microstrip Detector Modules for the CMS Inner Tracking Detector*, CMS Note 2006/141 (2006).
- [51] P. Schleper, G. Steinbrück and M. Stoye, *Software Alignment of the CMS Tracker using MILLEPEDE II*, CMS Note 2006/011 (2006).
- [52] E. Widl, R. Frühwirth and W. Adam, *A Kalman Filter for Track-based Alignment*, CMS Note 2006/022 (2006).
- [53] V. Karimäki, A. Heikkinen, T. Lampén and T. Lindén, *Sensor Alignment by Tracks*, *Proceedings of the CHEP2003 - International Conference on Computing in High Energy and Nuclear Physics*, La Jolla (2003), CMS Conference Report CR 2003/022 (2003).
- [54] R. Brauer et al., *Design and Test Beam Performance of Substructures of the CMS Tracker End Caps*, CMS Note 2005/025 (2005).
- [55] CMS TIB Collaboration, *Validation tests of CMS TIB/TID structures*, in preparation.
- [56] W. deBoer et al., *The performance of irradiated CMS silicon micro-strip detector modules*, CMS Note 2006/049 (2006).
- [57] A. Chilingarov et al., 1995 Nucl. Instrum. Meth. A **360** 432.
- [58] E. Fretwurst et al., *Reverse annealing of the effective impurity concentration and long term operational scenario for silicon detectors in future collider experiments*. 1994 Nucl. Instrum. Meth. A **342** 119.
- [59] M. M. Angarano et al., *Study of radiation damage and substrate resistivity effects from beam test of silicon microstrip detectors using LHC readout electronics*, CMS Note 2000/053 (2000)
- [60] S. Albergo et al., *High Voltage Operation of Heavily Irradiated Silicon Microstrip Detectors*, CMS Conference Report CR 1999/019 (1999).
- [61] Particle Data Group Collaboration, S. Eidelman et al., *Review of particle physics*, 2004 Phys. Lett. B **592** 1.
- [62] P. Lecoq et al, *Lead tungstate (PbWO₄) scintillators for LHC EM calorimetry*, 1995 Nucl. Instrum. Meth. A **365** 291.
S. Baccaro et al., *Influence of La³⁺-doping on radiation hardness and thermoluminescence characteristics of PbWO₄*, 1997 Phys. Stat. Sol. (a) **160** R5.
A. A. Annekov et al., *Improvement of several properties of lead tungstate crystals with different doping ions*, 1998 Nucl. Instrum. Meth. A **402** 75.
M. Kobayashi et al., *Improvement of radiation hardness of PbWO₄ scintillating crystals by La-doping*, 1998 Nucl. Instrum. Meth. A **404** 149.
H. F. Chen et al., *Radiation damage measurements of undoped lead tungstate crystals for the CMS electromagnetic calorimeter at LHC*, 1998 Nucl. Instrum. Meth. A **414** 149.
H. Hofer et al., *Afterglow measurements of lead tungstate crystals*, 1999 Nucl. Instrum. Meth. A **433** 630.
M. Kobayashi et al., *Significant improvement of PbWO₄ scintillating crystals by doping with trivalent ions*, 1999 Nucl. Instrum. Meth. A **434** 412.
- [63] A. A. Annenkov, M. V. Korzhik, and P. Lecoq, *Lead tungstate scintillation material*, 2002 Nucl. Instrum. Meth. A **490** 30.
- [64] I. Dafinei, E. Auffray, P. Lecoq, M. Schneegans, *Lead Tungstate for High Energy Calorimetry*, *Proceedings of Scintillator and Phosphor Materials Symposium, San Francisco 1994*, Mat. Res. Soc. Symp. Proc. Vol. 348 (1994) 99.

- [65] X. D. Qu, L. Y. Zhang, and R. Y. Zhu, *Radiation induced color centers and light monitoring for lead tungstate crystals*, 2000 IEEE Trans. Nucl. Sci. **47** 1741.
- [66] S. Baccaro et al., *Ordinary and extraordinary complex refractive index of the lead tungstate (PbWO₄) crystals*, 1997 Nucl. Instrum. Meth. A **385** 209.
- [67] D. Graham and C. Seez, *Simulation of Longitudinal Light Collection Uniformity in PbWO₄ crystals*, CMS Note 1996/002 (1996).
- [68] CMS Collaboration, *The Electromagnetic Calorimeter Project, Technical Design Report*, CERN/LHCC 97-33 (1997). CMS Addendum CERN/LHCC 2002-27.
- [69] R. Y. Zhu, *Radiation damage in scintillating crystals*, 1998 Nucl. Instrum. Meth. A **413** 297.
- [70] M. Huhtinen et al., *High-energy proton induced damage in PbWO₄ calorimeter crystals*, 2005 Nucl. Instrum. Meth. A **545** 63.
- [71] P. Lecomte et al., *High-energy proton induced damage study of scintillation light output from PbWO₄ calorimeter crystals*, 2006 Nucl. Instrum. Meth. A **564** 164.
- [72] M. Lebeau, F. Mossire, H. Rezvani Naraghi, *The Super-Basket : Incorporation of Conical Reinforcements in the CMS ECAL EB Support Structure*, CMS Note 2001/003 (2001).
- [73] P. Adzic et al., *Results of the first performance tests of the CMS electromagnetic calorimeter*, 2006 Eur. Phys. J. C **44** S2 1.
- [74] K. Deiters et al., *Double screening tests of the CMS ECAL avalanche photodiodes*, 2005 Nucl. Instrum. Meth. A **543** 549.
- [75] S. Baccaro et al., *Radiation damage effect on Avalanche Photo Diodes*, 1999 Nucl. Instrum. Meth. A **426** 206.
- [76] A. Bartoloni, *The power supply system for CMS-ECAL APDs, Proceedings of the 7th Workshop on Electronics for LHC Experiments, Stockholm 2001*.
A. Bartoloni et al., *High Voltage System for the CMS Electromagnetic Calorimeter*, 2007 Nucl. Instrum. Meth. A **582** 426.
- [77] K.W. Bell et al., *Vacuum Phototriodes for the CMS Electromagnetic Calorimeter Endcap*, 2004 IEEE Trans. Nucl. Sci. **51** 2284.
N. A. Bajanov et al., *Fine-mesh photodetectors for CMS endcap electromagnetic calorimeter*, 2000 Nucl. Instrum. Meth. A **442** 146.
Yu. Blinnikov et al., *Radiation hardness, excess noise factor and short-term gain instability of vacuum phototriodes for the operation in pseudorapidity range $1.5 \leq \eta \leq 3.0$ at CMS ECAL*, 2003 Nucl. Instrum. Meth. A **504** 228.
- [78] K.W. Bell et al., *The response to high magnetic fields of the vacuum phototriodes for the Compact Muon Solenoid endcap electromagnetic calorimeter*, 2003 Nucl. Instrum. Meth. A **504** 255.
- [79] Yu. I. Gusev et al., *Super radiation hard vacuum phototriodes for the CMS endcap ECAL*, 2004 Nucl. Instrum. Meth. A **535** 511.
- [80] B. Betev et al., *Low Voltage Supply System for the Very Front End Readout Electronics of the CMS Electromagnetic Calorimeter*, CERN/LHCC 2003/55 (2003) 353.
- [81] M. Raymond et al., *The MGPA Electromagnetic Calorimeter Readout Chip for CMS*, CERN/LHCC 2003/55 (2003) 83.
- [82] G. Minderico et al., *A CMOS low power, quad channel, 12 bit, 40 MS/s pipelined ADC for applications in particle physics calorimetry*, CERN/LHCC 2003/55 (2003) 88.
- [83] M. Hansen, *The New Readout Architecture for the CMS ECAL*, CERN/LHCC 2003/55 (2003) 78.
- [84] R. Alemany et al., *CMS ECAL Off-Detector Electronics, Proceedings of the 11th International Conference On Calorimetry In High Energy Physics, CALOR 2004, Perugia 2004*.

- [85] R. Alemany et al., *Overview of the ECAL Off-Detector Electronics of the CMS Experiment*, Nuclear Science Symposium Conference record, 2004 IEEE **2** 1053.
- [86] P. Paganini et al., *Tests of the boards generating the CMS ECAL Trigger Primitives: from the On Detector electronics to the Off-Detector electronics system*, *Proceedings of the 10th Workshop on electronics for LHC experiments, Boston (2004)*.
- [87] N. Almeida et al., *Calorimeter Trigger Synchronization in CMS, Implementation and Test System*, *Proceedings of the 10th Workshop on electronics for LHC experiments, Boston (2004)*.
- [88] N. Almeida et al., *Data Concentrator Card and Test System for the CMS ECAL Readout*, *Proceedings of the 9th Workshop on Electronics for the LHC Experiments, Amsterdam (2003)*. CERN-LHCC-2003-055 (2003).
- [89] J. C. Silva et al., *Test results of the Data Concentrator Card of the CMS Electromagnetic Calorimeter Readout System*, *Proceedings of the 10th Workshop on electronics for LHC experiments, Boston (2004)*.
- [90] N. Almeida et al., *The selective readout processor for the CMS electromagnetic calorimeter*, Nuclear Science Symposium Conference record, 2004 IEEE **3** 1721.
- [91] P. Wertelaers et al., *ECAL Preshower*, CMS ECAL EDR-4, Vol. 2 CMS-GE-ER (2000).
- [92] P. Aspell et al., *PACE3: A large dynamic range analogue memory ASIC assembly designed for the readout of silicon sensors in the LHC CMS Preshower*, *Proceedings of the 10th Workshop on electronics for LHC experiments, Boston (2004)*.
- [93] I. Evangelou, *CMS Preshower in-situ absolute calibration*, *Proceedings of the 9th ICATPP Conference, Villa-Olmo, Como (2005)*.
- [94] K. Kloukinas et al., *Kchip: a radiation tolerant digital data concentrator chip for the CMS Preshower detector*, *Proceedings of the 9th Workshop on electronics for LHC experiments, Amsterdam (2003)*, CERN-LHCC-2003-055 (2003), 66.
- [95] G. Antchev et al., *A VME-Based Readout System for the CMS Preshower Sub-Detector*, 2007 IEEE Trans. Nucl. Sci. **54** 623.
- [96] S. Reynaud and P. Vichoudis *A multi-channel optical plug-in module for gigabit data reception*, *Proceedings of the 12th Workshop on electronics for LHC and future experiments, Valencia 2006*, CERN-LHCC-2007-006 (2007).
- [97] D. Barney et al., *Implementation of on-line data reduction algorithms in the CMS Endcap Preshower Data Concentrator Card*, 2007 JINST **2** 03001.
- [98] P. Adzic et al., *The detector control system for the electromagnetic calorimeter of the CMS experiment at the LHC*, *Proceedings of the 10th International Conference on Accelerator and Large Experimental Physics Control Systems, Geneva 2005*, ICALEPS 2005, CMS Conference Report CR 2005/28 (2005).
- [99] P. Milenovic et al. *Performance of the CMS ECAL safety system for Super Modules SM0 and SM1*, 2005 Nucl. Instrum. Meth. A **554** 427.
- [100] F. Cavallari et al., *Intercalibration of ECAL crystals using laboratory measurements*, CMS Note 2006/073 (2006).
- [101] W. Bertl et al., *Feasibility of intercalibration of CMS ECAL supermodules with cosmic rays*, 2005 Eur. Phys. J. C **41** S2 11.
F. Ferri, P. Govoni, *The CMS Electromagnetic Calorimeter Pre-calibration with Cosmic Rays and Test Beam Electrons*, *Proceedings of VCI 2007, Vienna 2007*, CMS Conference Report CR 2007/12 (2007).
- [102] L. Zhang et al., *Performance of the Monitoring Light Source for the CMS Lead Tungstate Crystal Calorimeter*, 2005 IEEE Trans. Nucl. Sci. **52** 1123.
M. Anfreville et al., *Laser monitoring system for the CMS lead tungstate crystal calorimeter*, CMS Note 2007/28 (2007)
- [103] A. Ghezzi et al., *Analysis of the response evolution of the CMS electromagnetic calorimeter under electron and pion irradiation*, CMS Note 2006/38 (2006).

- [104] P. Aspell et al., *Results from the 1999 Beam Test of a Preshower Prototype*, CMS Note 2000/1 (2000).
- [105] P. Adzic et al, *Reconstruction of the signal amplitude of the CMS electromagnetic calorimeter*, 2006 Eur. Phys. J. C **46** 23.
- [106] P. Adzic et al, *Energy resolution of the barrel of the CMS electromagnetic calorimeter*, 2007 JINST **2** P04004 .
- [107] CMS, *The Hadron Calorimeter Technical Design Report*, CERN/LHCC 97-31 (1997).
- [108] P. Cushman, A. Heering, and A. Ronzhin, *Custom HPD Readout for the CMS HCAL*, 2000 Nucl. Instrum. and Meth., A **442** 289.
- [109] A. Heister et al., *Measurement of Jets With the CMS Detector at the LHC*, CMS Note 2006/036 (2006).
- [110] H. Pi et al., *Measurement of Missing Transverse Energy With the CMS Detector at the LHC*, CMS Note 2006/035 (2006), to be published in Eur. Phys. J. C.
- [111] G. Baatian et al., *Design, Performance, and Calibration of CMS Hadron-Barrel Calorimeter Wedges*, CMS Note 2006/138 (2006), submitted to Eur. Phys. J. C.
- [112] V. I. Kryshkin and A. I. Ronzhin, 1986 Nucl. Instrum. Meth. A **583**.
- [113] M. G. Albrow et al., 1987 Nucl. Instrum. Meth. A **256** 23.
- [114] Sudeshna Banerjee and Sunanda Banerjee, *Performance of Hadron Calorimeter with and without HO*, CMS Note 1999/063 (1999).
- [115] N. Akchurin and R. Wigmans, *Quartz fibers as active elements in detectors for particle physics*, 2002 Review of Scientific Instruments, **74** 2955.
- [116] I. Dumanoglu et al., *Radiation-hardness Studies of High OH⁻ Content Quartz Fibers Irradiated with 500 MeV Electrons*, 2002 Nucl. Instrum. Meth. A **490** 444.
- [117] N. Akchurin et al., *Effects of Radiation Damage and Their Consequences for the Performance of the Forward Calorimeters in the CMS Experiment*, 2002 Nucl. Instrum. Meth. B **187** 66.
- [118] R. Thomas, *Study of radiation hardness of optical fibers*, MS Thesis, Texas Tech University, 2004.
- [119] A. Panagiotou et al., *CASTOR Engineering Design Report*, CMS Note in preparation, 2007.
- [120] X. N. Wang and M. Gyulassy, *HIJING: A Monte Carlo model for multiple jet production in p p, p A and A A collisions*, 1991 Phys. Rev. D **44** 3501,
X. N. Wang, *A pQCD-based approach to parton production and equilibration in high-energy nuclear collisions*, 1997 Phys. Rept. **280** 287.
- [121] D. D’Enterria et al., *CMS physics technical design report: Addendum on high density QCD with heavy ions*, 2007 J. Phys. G **34** 2307, CERN-LHCC-2007-009 (2007).
- [122] M. Albrow et al., *Prospects for diffractive and forward physics at the LHC*, CMS Note 2007/002 (2007), CERN-LHCC-2006-039 (2006).
- [123] G. Mavromanolakis, *Quartz fiber calorimetry and calorimeters*, 1998 [arXiv:physics/0412123].
- [124] X. Aslanoglou et al., *First performance studies of a prototype for the CASTOR forward calorimeter at the CMS experiment*, 2007 [arXiv:0706.2576].
- [125] X. Aslanoglou et al., *Performance Studies of Prototype II for the CASTOR forward Calorimeter at the CMS Experiment*, 2007 [arXiv:0706.2641], to be published in Eur. Phys. J. C.
- [126] A.S. Ayan et al., *CMS Zero-Degree-Calorimeter (ZDC). Technical Design Report*, in preparation.
- [127] O. A. Grachov et al., *Status of zero degree calorimeter for CMS experiment*, 2006 AIP Conf. Proc. **867** 258 [nucl-ex/0608052].
- [128] E. H. Hoyer, W. C. Turner and N. V. Mokhov, *Proc.of the 6th European Part.Accel.Conf., Stockholm (1998)*.

- [129] J.-F. Beche et al., *An Ionization Chamber Shower Detector for the LHC Luminosity Monitor*, Nuclear Science Symposium Conference Record, IEEE 2000 **1** 5-136.
- [130] O. A. Grachov et al., *Measuring photons and neutrons at zero degrees in CMS*, 2007 [nucl-ex/0703001].
- [131] CMS Collaboration, *The Muon Project, Technical Design Report*, CERN/LHCC 97-32 (1997).
- [132] M. De Giorgi et al., *Proceedings of the First Workshop on Electronics for LHC experiments*, CERN/LHCC 95-56 (1995).
- [133] R. Veenhof, *Garfield. A Drift Chamber Simulation Program User's Guide*, CERN Program Library W5050 (1994).
- [134] T. Zhao et al., *A study of electron drift velocity in Ar-CO₂ and Ar-CO₂-CF₄ gas mixtures*, 1994 Nucl. Instrum. Meth. A **340** 485.
- [135] A. Benvenuti et al., *Simulations in the development of the barrel muon chambers for the CMS detector at LHC*, 1998 Nucl. Instrum. Meth. A **405** 20.
- [136] CMS Muon Group, G. Alampi et al., *Electrode Strip Deposition for the CMS Barrel Drift Tube System*, CMS Note 2006/144 (2006).
- [137] J. Christiansen et al., *Proceedings of the 6th Workshop on Electronics for LHC Experiments, Cracow (2000)*.
- [138] P. Moreira et al., *Proceedings of the 7th Workshop on Electronics for LHC Experiments, Stockholm (2001)*.
- [139] F. Gonella, M. Pegoraro, *A prototype ASIC for the readout of the drift tubes of CMS Barrel Muon Chambers*, CERN/LHCC 98-36, 257 (1998).
- [140] L. Barcellan et al., *Single Events Effects induced by heavy ions on the frontend ASIC developed for the muon DT chambers*, LNL Annual Report (2000).
- [141] C. Fernandez Bedoya et al., *Electronics for the CMS muon drift tube chambers: The read-out minicrate*, 2005 IEEE Trans. Nucl. Sci. 52 944.
- [142] G. Dellacasa, V. Monaco and A. Staiano, *Proceedings of Hadron Collider Physics, La Biodola, Isola d'Elba (2007)*.
- [143] C. Albajar et al., *Test beam analysis of the first CMS drift tube muon chamber*, 2004 Nucl. Instrum. Meth. A **525** 465.
- [144] M. Aguilar-Benitez et al., *Study of magnetic field effects in drift tubes for the barrel muon chambers of the CMS detector at the LHC*, 1998 Nucl. Instrum. Meth. A **416** 243.
- [145] M. Aguilar-Benitez et al., *Construction and test of the final CMS Barrel Drift Tube Muon Chamber prototype*, 2002 Nucl. Instrum. Meth. A **480** 658.
- [146] P. Arce et al., *Bunched beam test of the CMS drift tubes local muon trigger*, 2004 Nucl. Instrum. Meth. A **534** 441.
- [147] M. Aldaya et al., *Fine synchronization of the muon drift tubes local trigger*, 2007 Nucl. Instrum. Meth. A **579** 951.
- [148] G. Charpak and F. Sauli, *High-accuracy, two-dimensional read-out in multiwire proportional chambers*, 1973 Nucl. Instrum. Meth. **113** 381.
- [149] C. Anderson et al., *Effect of gas composition on the performance of cathode strip chambers for the CMS endcap muon system*, CMS Note 2004-033 (2004).
- [150] T. Ferguson et al., *Anode front-end electronics for the cathode strip chambers of the CMS Endcap Muon detector*, Nucl. Instrum. Meth. A **539** 386 (2005).
- [151] R. Breedon et al., *Performance and radiation testing of a low-noise switched capacitor array for the CMS Endcap Muon system*, in *Proceedings of the 6th Workshop on Electronics for LHC Experiments*, CERN 2000-010 187 (2000).

- [152] M. M. Baarmand et al., *Spatial resolution attainable with cathode strip chambers at the trigger level*, 1999 Nucl. Instrum. Meth. A **425** 92.
- [153] V. Barashko, *Performance validation tests of the cathode strip chambers for CMS muon system*, Nuclear Science Symposium Conference Record, IEEE 2005 **2** 827.
- [154] E. Gatti et al., *Optimum geometry for strip cathodes on grids in MWPC for avalanche localization along the anode wires*, 1979 Nucl. Instrum. Meth. **163** 83.
- [155] O. Prokofiev et al., *Aging tests of full scale CMS muon cathode strip chambers*, 2003 Nucl. Instrum. Meth. A **515** 226.
- [156] L. Malter, *Thin Film Field Emission*, Phys. Rev. 1936 **50** 48.
- [157] R. Breedon et al., *Results of radiation test of the cathode front-end board for CMS Endcap Muon chambers*, 2001 Nucl. Instrum. Meth. A **471** 340.
- [158] R. Santonico and R. Cardarelli, *Development of Resistive Plate Counters*, 1981 Nucl. Instrum. Meth. **187** 377.
- [159] R. Cardarelli et al., *Performance of a Resistive Plate Chamber operating with pure CF₃Br*, 1993 Nucl. Instrum. Meth. A **333** 399.
- [160] M. Abbrescia et al., *The simulation of Resistive Plate Chambers in avalanche mode: charge spectra and efficiency*, 1999 Nucl. Instrum. Meth. A **431** 413.
- [161] M. Abbrescia et al., *Local and global performance of avalanche operated double gap a Resistive Plate Chamber for the CMS experiment*, 1999 Nucl. Instrum. Meth. A **434** 90.
- [162] M. Abbrescia et al., *Study of long term performance of CMS RPC under irradiation at the CERN GIF*, 2004 Nucl. Instrum. Meth. A **533** 102.
- [163] M. Abbrescia et al., *Neutron irradiation of the RPCs for the CMS experiment*, 2003 Nucl. Instrum. Meth. A **508** 120.
- [164] M. Abbrescia et al., *New developments on front-end electronics for the CMS Resistive Plate Chambers*, 2000 Nucl. Instrum. Meth. A **456** 143.
- [165] M. Abbrescia et al., *Long term performance of double gap Resistive Plate Chamber under gamma irradiation*, 2002 Nucl. Instrum. Meth. A **477** 293.
- [166] M. Abbrescia et al., *Neutron induced single event upset on the RPC front-end chips for the CMS experiment*, 2002 Nucl. Instrum. Meth. A **484** 494.
- [167] M. Abbrescia et al., *An RPC-based Technical Trigger for the CMS Experiment*, Proceedings of the 12th Workshop on Electronics for LHC and Future Experiments, Valencia, Spain, 25 Sept. 2006.
- [168] G. Iaselli et al., *Properties of C₂H₂F₄-based gas mixture for avalanche mode operation of Resistive Plate Chambers*, 1997 Nucl. Instrum. Meth. A **398** 173.
- [169] M. Bosteels et al., *CMS Gas System Proposal*, CMS Note 1999/018 (1999).
L. Besset, F. Hahn, S. Haider, and C. Zinoni, *Experimental tests with a standard closed loop gas circulation system*, CMS Note 2000/040 (2000).
I. Crotty, F. Hahn, G. Iaselli, and S. Nuzzo, *RPC gas distribution layout*, CMS Note 2001/019 (2001).
- [170] M. Abbrescia et al., *Gas analysis and monitoring systems for the RPC detector of CMS at LHC*, Nuclear Science Symposium Conference record, 2006 IEEE **2** 891, [Frascati preprint LNF-06-34(P)].
- [171] M. Abbrescia et al., *HF production in CMS-Resistive Plate Chambers*, 2006 Nucl. Phys. B (Proc. Suppl.) **158** 30.

- [172] M. Abbrescia et al., *The bakelite for the RPCs of the experiment CMS*, 2000 Nucl. Instrum. Meth. A **456** 132.
M. Abbrescia et al., *Production and quality control of the Barrel RPC chambers of the CMS experiment*, 2006 Nucl. Phys. B (Proc. Suppl.) **150** 290.
M. Abbrescia et al., *Quality control tests for the CMS barrel RPCs*, 2006 Nucl. Phys. B (Proc. Suppl.) **158** 73.
A. Ball et al., *Cosmic ray certification of the first 100 CMS endcap RPCs and the corresponding construction database*, 2006 Nucl. Phys. B (Proc. Suppl.) **158** 99.
Z. Aftab et al., *Production and quality control for the CMS endcap RPCs*, 2006 Nucl. Phys. B (Proc. Suppl.) **158** 16.
Z. Aftab et al., *Assembly and quality certification for the first station of the endcap RPCs, RE1*, 2006 Nucl. Phys. B (Proc. Suppl.) **158** 103.
- [173] M. Abbrescia et al., *Resistive plate chambers performances at cosmic rays fluxes*, 1995 Nucl. Instrum. Meth. A **359** 603.
- [174] M. Abbrescia et al., *Cosmic ray tests of double-gap resistive plate chambers for the CMS experiment*, 2005 Nucl. Instrum. Meth. A **550**, no. 1–2, 116.
- [175] CMS Muon Detector Survey Documents, EDMS Doc. CMS-00000083880.
<https://edms.cern.ch/cedar/plsql/cms> .
- [176] C. Carneiro, R. Goudard and C. Humbertclaude, *CMS MAB prototype-deformation test under load and humidity-measurements from July to October 2002*, EDMS Doc. CMS-MA-UR-0001.
<https://edms.cern.ch/cedar/plsql/cms> .
- [177] G. Szekely et al., *Muon Barrel alignment system based on a net of PC/104 board computers*, *Proceedings of the 9th Workshop on Electronics for LHC Experiments, Amsterdam (2003)*, CERN-2003-006, CERN-LHC-2003-055, LHCC-G-061 (2003).
- [178] M. Hohlmann et al., *Design and performance of the alignment system for the CMS Muon Endcaps* Nuclear Science Symposium Conference record, 2006 IEEE **1** 489,
R. H. Lee, *Simulation and study of the CMS Endcap Muon alignment scheme*, Purdue University, PhD Thesis, 2002.
- [179] J. Moromisato et al., *The development of totally transparent position sensors*, 2005 Nucl. Instrum. Meth. A **538** 234.
- [180] M. Ripert, *Calibration of analog sensors for the alignment of muon chambers in the CMS experiment*, Florida Institute of Technology, MS Thesis, 2005.
- [181] A. Calderon et al., *Large size high performance transparent amorphous silicon sensors for laser beam position detection*, 2006 Nucl. Instrum. Meth. A **565** 603.
- [182] A. Lopez Virto, *Caracterizacion y Pruebas de Validacion del Sistema Link de Alineamiento de CMS*, Universidad de Cantabria, PhD Thesis, 2003.
- [183] A. Calderon, *Construccion, Calibracion y Evaluacion del Sistema Link de Alineamiento del Espectrometro de Muones del Experimento CMS*, Universidad de Cantabria, PhD Thesis, 2006.
- [184] *ELMB Boards*
<http://elmb.web.cern.ch/ELMB/elmb128.html>
- [185] P. Arce, *Object oriented software for simulation and reconstruction of big alignment systems*, 2003 Nucl. Instrum. Meth. A **502** 696.
- [186] D. E. Stewart and Z. Leyk, *Meschach Library*
<http://www.netlib.org/c/meschach/readme>
- [187] CMS Collaboration, *The TriDAS Project – The Level-1 Trigger Technical Design Report*, CERN/LHCC 2000-38 (2000).
- [188] CMS Collaboration, *The TriDAS Project – Data Acquisition and High-Level Trigger Technical Design Report*, CERN/LHCC 2002-26 (2002).

- [189] W. Adam et al., CMS Trigger and Data Acquisition Group, *The CMS high level trigger*, 2005. Eur. Phys. J. C **46** 605 [hep-ex/0512077].
- [190] I. Magrans de Arbril, C.-E. Wulz, J. Varela, *Conceptual Design of the CMS Trigger Supervisor*, 2006 IEEE Trans. Nucl. Sci. **53** 474.
- [191] P. Chumney et al., *Level-1 Regional Calorimeter Trigger System for CMS*, *Proc. of Computing in High Energy Physics and Nuclear Physics, La Jolla (2003)*, [hep-ex/0305047].
- [192] <http://www.hep.ph.ic.ac.uk/cms/gct>
- [193] <http://www.picmg.org/v2internal/microTCA.htm>
- [194] F. Gasparini et al., *Bunch crossing identification at LHC using a mean-timer technique*, 1993 Nucl. Instrum. Meth. A **336** 91.
- [195] M. Andlinger et al., *Pattern Comparator Trigger (PACT) for the Muon System of the CMS Experiment*, 1996 Nucl. Instrum. Meth. A **370** 389.
- [196] C. Albajar et al., *Conceptual design of an improved CMS RPC Muon Trigger using the Hadron Outer scintillators*, 2005 Nucl. Instrum. Meth. A **545** 97.
- [197] J. Erö et al., *The CMS Drift Tube Trigger Track Finder*, CMS Note 2007 (2007).
- [198] D. Acosta et al., *Performance of a Pre-Production Track-Finding Processor for the Level-1 Trigger of the CMS Endcap Muon System*, *Proc. of the 10th Workshop on Electronics for LHC and Future Experiments, Boston (2004)*.
- [199] H. Sakulin et al., *Implementation and Test of the First-Level Global Muon Trigger of the CMS Experiment*, *Proc. of the 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg (2005)*.
- [200] C.-E. Wulz, *Concept of the CMS First Level Global Trigger for the CMS Experiment at LHC*, 2001 Nucl. Instrum. Meth. A **473** 231.
- [201] A. Taurok, H. Bergauer, M. Padrta, *Implementation and Synchronisation of the First Level Global Trigger for the CMS Experiment at LHC*, 2001 Nucl. Instrum. Meth. A **473** 243.
- [202] Ph. Glaser et al., *Design and Development of a Graphical Setup Software for the CMS Global Trigger*, 2006 IEEE Trans. Nucl. Sci. 53 Iss. **3** 1282.
- [203] CMS Trigger and Data Acquisition Group, *CMS L1 Trigger Control System*, CMS Note 2002/033 (2002).
- [204] B. G. Taylor, *Timing Distribution at the LHC*, *Proc. of the 8th Workshop on Electronics for LHC and Future Experiments, Colmar (2002)*.
- [205] T. Geralis et al., *The Global Trigger Processor Emulator for the CMS experiment*, 2005 IEEE Trans. on Nucl. Science **52** 1679.
- [206] E. Cano et al., *FED-kit design for CMS DAQ system*, *Proceedings of the 8th Workshop on Electronics for LHC experiments, Colmar, France, 9-13 September 2002*.
- [207] B.G. Taylor, *TTC Distribution for LHC Detectors*, 1998 IEEE Trans. Nucl. Sci. **45** 82, <http://www.cern.ch/TTC/intro.htm>.
- [208] A. Racz, R. McLaren, E. van der Bij, *The S-Link64 bit extension specification: S-Link64*, available at <http://hsi.web.cern.ch/HSI/s-link>.
- [209] N.J. Boden et al., *Myrinet - A Gigabit per Second Local Area Network*, 1995 IEEE-Micro **15** 29.
- [210] G. Bauer et al., *The Tera-bit/s Super-Fragment Builder and Trigger Throttling System for the Compact Muon Solenoid Experiment at CERN*, *15th IEEE Real Time Conference 2007, Batavia (2007)*, CMS CR-2007/020 (2007), Submitted to IEEE Transactions on Nuclear Science.
- [211] G. Bauer et al., *CMS DAQ Event Builder based on Gigabit Ethernet*, *15th IEEE Real Time Conference 2007, Batavia (2007)*, CMS CR-2007/016 (2007), Submitted to IEEE Transactions on Nuclear Science.

- [212] C.-D. Jones et al., *The New CMS Data Model and Framework, Proceedings of the conference on Computing in High Energy Physics, Mumbai (2006)*.
- [213] Quattor is a system administration toolkit,
<http://www.quattor.org> or <http://quattor.web.cern.ch> .
- [214] J. Gutleber and L. Orsini, *Software architecture for processing clusters based on i2o*, 2002 Cluster Computing **5** 55.
- [215] V. Brigljevic et al., *Using XDAQ in Application Scenarios of the CMS Experiment, Proceedings of Computing in High Energy Physics, La Jolla, (2003)*
- [216] O. Nierstrasz, S. Gibbs and D. Tschritzis, *Component-Oriented Software Development*, 1992 Comm. of the ACM **35** 160.
- [217] D. Box et al., *Simple Object Access Protocol (SOAP) 1.1*, W3C Note 08 (2000),
<http://www.w3.org/TR/SOAP> .
- [218] for the I2O standard,
<http://developer.osdl.org/dev/opendoc/Online/Local/I20/index.html> .
- [219] <http://xdaqwiki.cern.ch> , and references therein.
- [220] R. Arcidiacono et al., *HyperDAQ – Where Data Acquisition Meets the Web, 10th ICALEPCS Int. Conf. on Accelerator and Large Expt. Physics Control Systems, Geneva (2005)*
- [221] J. Boyer, Canonical XML version 1.0, W3C Recommendation, 16 August 2006,
<http://www.w3c.org/XML> . FED-kit design for CMS DAQ system
- [222] Apache Axis is an XML based Web service framework,
<http://ws.apache.org/axis/> .
- [223] The Apache Tomcat servlet container,
<http://tomcat.apache.org> .
- [224] JCOP framework, <http://itcobe.web.cern.ch/itcobe/Projects/Framework/welcome.html>
- [225] OLE for Process Control (OPC), <http://www.opcfoundation.org/>
- [226] R. Arcidiacono et al., *CMS DCS design concepts, Proceedings of the 10th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS 2005), Geneva, Switzerland, 10-15 Oct 2005*.
- [227] L. Fernandez-Hernando et al., *Development of a CVD diamond Beam Condition Monitor for CMS at the Large Hadron Collider*, 2005 Nucl. Instrum. Meth. A **552** 183.
A. Macpherson, *Beam Condition Monitoring and radiation damage concerns of the experiment*, Proceedings of the XV LHC Project Chamonix Workshop, Divonne 2006, 198.
D. Chong et al., *Validation of synthetic diamond for a Beam Condition Monitor for the Compact Muon Solenoid Experiment*, 2007 IEEE Trans. Nucl. Sci. **54** 182.
- [228] R.J. Tapper, *Diamond detectors*, 2000 Rep. on Prog. in Phys. **63** 8.
- [229] R. Eusebi et al., *A diamond-based Beam Condition Monitor for the CDF experiment*, 2006 IEEE Trans. Nucl. Sci. **2** 709.
- [230] M. Brunisma et al., *CVD diamonds in the BaBar radiation monitoring system* 2006 Nucl. Phys. B **150** 164.
- [231] W. de Boer et al., *Radiation hardness of diamond and silicon sensors compared*, 2007 Physica Status Solidi A **204** 3004.
- [232] B. Dehning et al., *The Beam Loss Monitor system*, Proceedings of the XIII LHC Project Chamonix Workshop, Chamonix 2004.
E. Effinger et al., *The LHC Beam Loss Monitoring system's data acquisition card*, Proceedings of LECC, Valencia 2006.
C. Zamantzas et al., *The LHC Beam Loss Monitoring system's surface building installation*, Proceedings of LECC, Valencia 2006.

- [233] C. Zamantzas, *The real-time data analysis and decision system for particle flux detection in the LHC accelerator at CERN*, Brunel University, PhD Thesis, 2006, CERN-THESIS-2006-037.
- [234] R. Schmidt et al., *Beam interlocks for LHC and SPS*, Proceedings of the International Conference on Accelerator and Large Experimental Physics Control Systems, Gyeongju 2003, ICALEPCS.
B. Todd, *A Beam Interlock System for CERN high energy accelerators*, Brunel University, PhD Thesis, 2007, CERN-THESIS-2007-019.
- [235] J. Kaplon and W. Dabrowski, *Fast CMOS Binary Front End for Silicon Strip Detectors at LHC Experiments*, 2005 IEEE Trans. Nucl. Sci. **52** 2713.
- [236] J. Troska et al., *Optical readout and control systems for the CMS Tracker*, 2003 IEEE Trans. Nucl. Sci. **50** 1067.
- [237] J. Bol, *Strahlmonitore aus diamant fuer teilchenstrahlen hoher intensitaet*, Karlsruhe University, PhD Thesis, 2006, IEKP-KA/2006-8.
- [238] J. Furltova, *Search for exotic processes in events with large missing transverse momentum in ZEUS at HERA*, Hamburg University, PhD Thesis, 2004, DESY-THESIS-2004-046.
- [239] G. Aguillion et al., *Thin scintillating tiles with high light yield for the OPAL endcaps*, 1998 Nucl. Instrum. Meth. A **417** 266.
- [240] C. Ohm, *Phase and intensity monitoring of the particle beams at the ATLAS experiment*, Linköping University, MS Thesis, 2007, urn:nbn:se:liu:diva-9614.
- [241] T. Wijnands, *Radiation monitoring for equipment in the LHC tunnel*, Functional Specification, 2005 EDMS Document 565013, <https://edms.cern.ch/>.
C. Pignard and T. Wijnands, *Radiation tolerant commercial of the shelf components for the remote readout of PIN diodes and Radfets*, Proceedings of the RADECS Conference, Cap d'Agde 2005.
- [242] J. Knobloch et al., *LHC Computing Grid Technical Design Report*, CERN-LHCC-2005-024 (2005).
- [243] CMS Collaboration, *CMS, The Computing Project: Technical Design Report*, CERN-LHCC-2005-023 (2005).
- [244] R. Brun, F. Rademakers, *ROOT - An Object Oriented Data Analysis Framework*, 1997 Nucl. Instrum. Meth. A **389** 81, <http://root.cern.ch>.
- [245] S. Kosyakov et al., *FroNtier: High Performance Database Access Using Standard Web Components in a Scalable Multi-tier Architecture*, Proceedings of the conference on Computing in High Energy Physics, Interlaken (2004).
- [246] M. Aderholz et al., *Models of Networked Analysis at Regional Centres for LHC Experiments (MONARC) - PHASE 2 REPORT*, CERN-LHB-2000-001 (2000).
- [247] <http://lhcopn.cern.ch>
- [248] J. Rehn et al., *PhEDEx high-throughput data transfer management system*, Proceedings of the conference on Computing in High Energy Physics, Mumbai (2006).
- [249] A. Fanfani et al., *Distributed Data Management in CMS*, Proceedings of the conference on Computing in High Energy Physics, Mumbai (2006).
- [250] D. Spiga et al., *CMS Workload Management*, 2007 Nucl. Phys. B **172C** 141.
- [251] D. Acosta et al., *CMS Computing, Software and Analysis Challenge in 2006 (CSA06) Summary*, CERN-LHCC-2007-010 (2007).