Chapter 7

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 material thickness crossed by muons, as a function of pseudorapidity, is shown in figure 7.1.

The CMS muon system is designed to have the capability of reconstructing the momentum and charge of muons over the the entire kinematic range of the LHC. CMS uses 3 types of gaseous particle detectors for muon identification [132]. 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 m^2 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

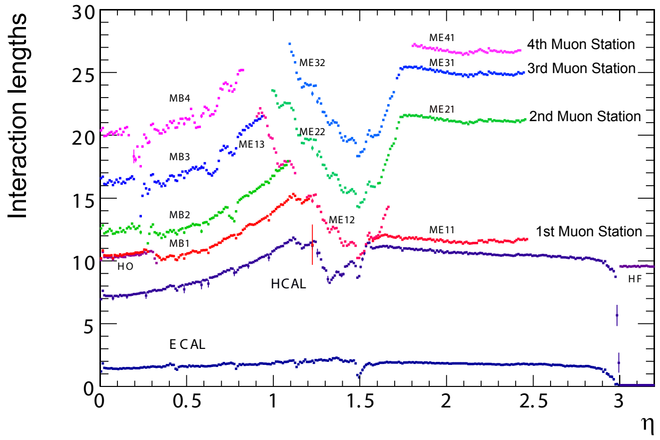

Figure 7.1: Material thickness in interaction lengths at various depths, as a function of pseudorapidity.

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 (figure 7.2) 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 [132].

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

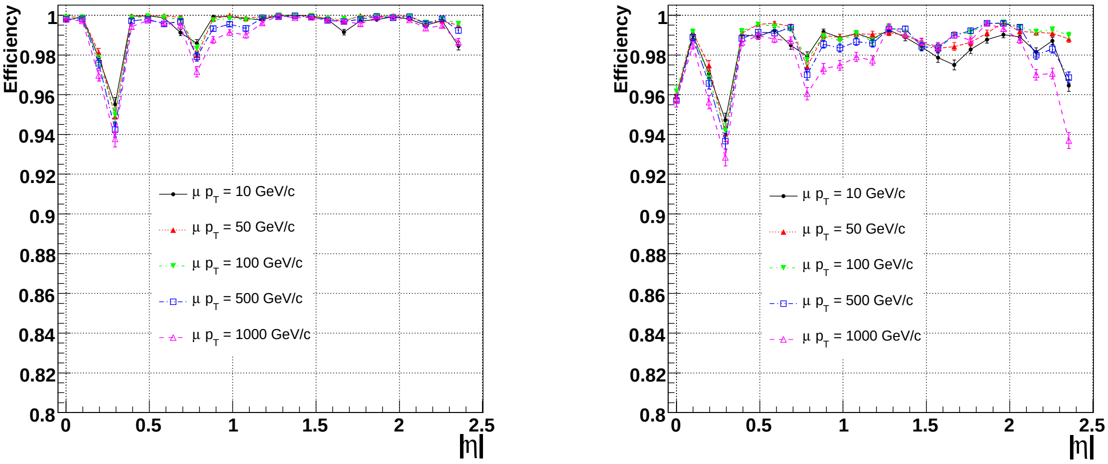

Figure 7.2: Muon reconstruction efficiency as a function of pseudorapidity for selected values of p_T. Left panel: standalone reconstruction (using only hits from the muon system with a vertex constraint). Right panel: global reconstruction (using hits from both the muon system and the tracker).

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% (figure 1.2). 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 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.

7.1 Drift tube system

7.1.1 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 85% Ar + 15% CO_2). 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.

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 (figure 7.3). 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 figure 7.4), each made of 4 layers of rectangular drift cells staggered by half a cell. The SL is the smallest independent unit of the design.

The wires in the 2 outer SLs are parallel to the beam line and provide a track measurement in the magnetic bending plane (r-ϕ). 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 station, 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 ($\geq$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

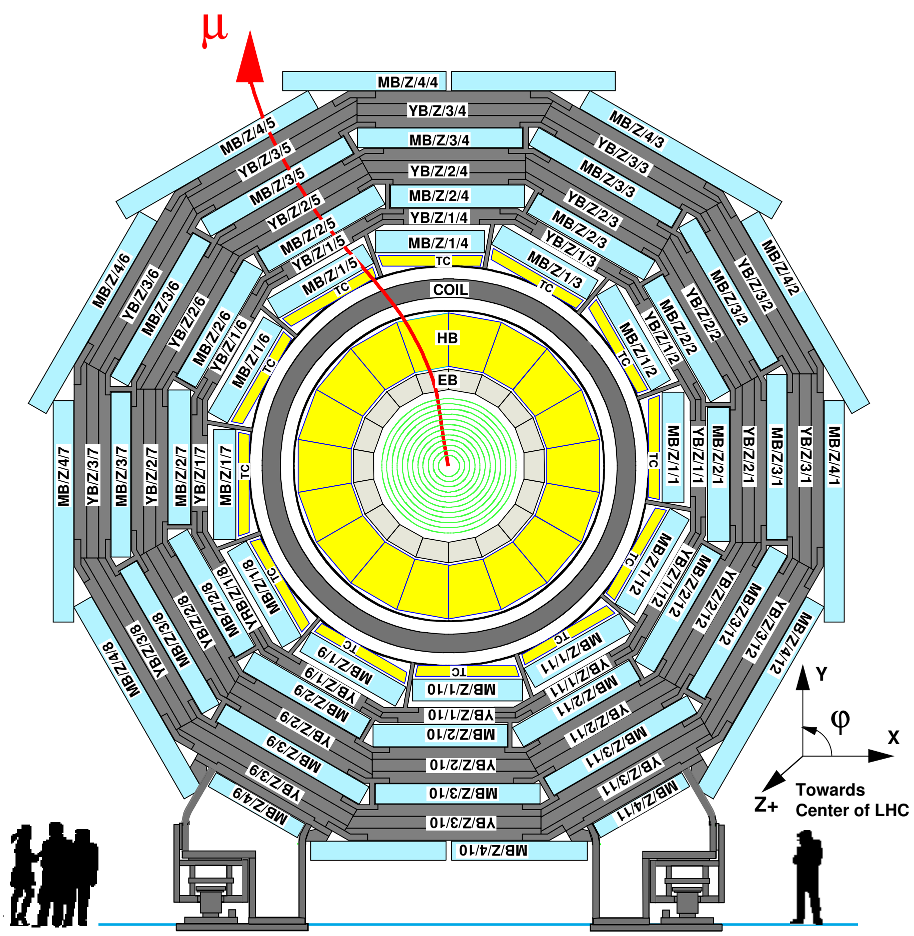

Figure 7.3: Layout of the CMS barrel muon DT chambers in one of the 5 wheels. The chambers in each wheel are identical with the exception of wheels –1 and +1 where the presence of cryogenic chimneys for the magnet shortens the chambers in 2 sectors. Note that in sectors 4 (top) and 10 (bottom) the MB4 chambers are cut in half to simplify the mechanical assembly and the global chamber layout.

the 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 mm^2.

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 (figure 7.4). The SLs are glued to the outer faces of the honeycomb. In this design, the honeycomb serves as a very light spacer,

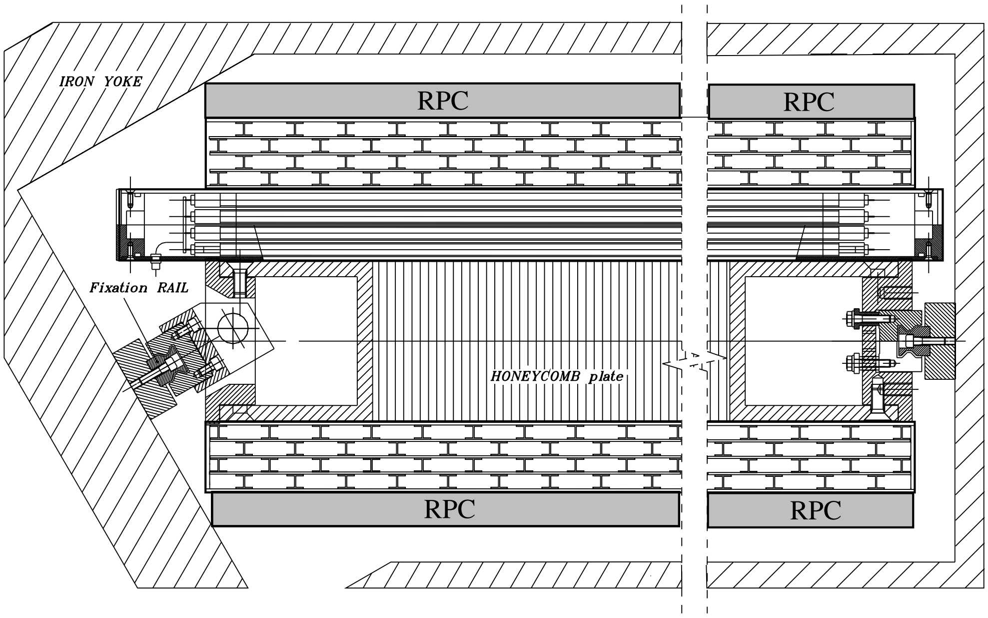

Figure 7.4: A DT chamber in position inside the iron yoke; the view is in the (r-ϕ) plane. One can see the 2 SLs with wires along the beam direction and the other perpendicular to it. In between is a honeycomb plate with supports attached to the iron yoke. Not shown are the RPCs, which are attached to the DT chambers via support plates glued to the bottom and/or top faces, depending on chamber type.

with rigidity provided by the outer planes of tubes. A thick spacer also helps to improve angular resolution within a station. Table 7.1 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 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 $p_T = 200\,\text{GeV}$. The 100-μm target chamber resolution is achieved by the 8 track points measured in the two ϕ SLs, since

Table 7.1: 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 figure 7.3). The SLs of type $\Phi(\Theta)$ measure the $\phi(z)$ coordinate in the CMS coordinate system.

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

the single wire resolution is better than 250 μm. To avoid corrections to the primary TDC data (section 7.1.3), 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 (section 7.1.3) identifier. The cell design includes 5 electrodes, 1 anode wire, 2 field shaping strips, and 2 cathode strips (figure 7.5 and section 7.1.2). The requirements of 250-μm resolution and 150-μm nonlinearity can be obtained by operating the tubes at atmospheric pressure with an Ar/CO_2 gas mixture and by keeping the CO_2 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.

7.1.2 Technical design

Drift cell

Figure 7.5 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

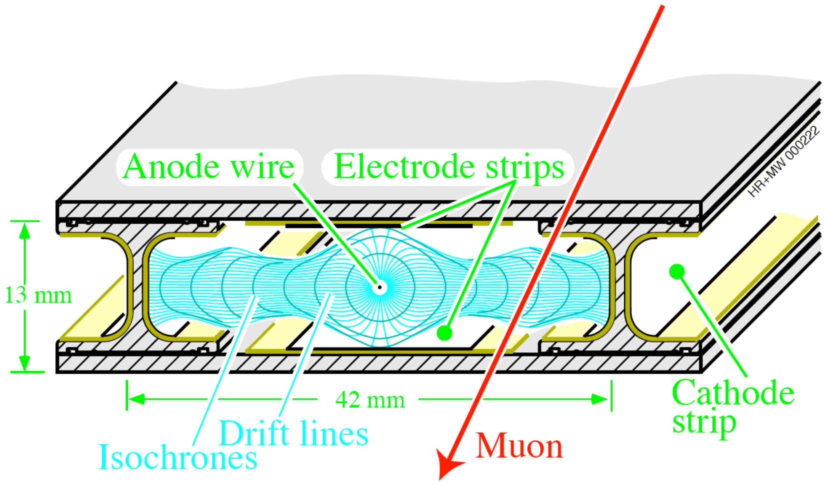

Figure 7.5: Sketch of a cell showing drift lines and isochrones. The plates at the top and bottom of the cell are at ground potential. The voltages applied to the electrodes are +3600V for wires, +1800V for strips, and −1200V for cathodes.

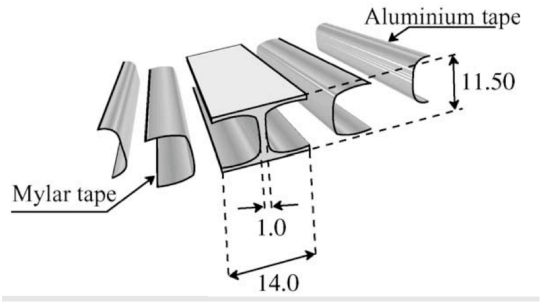

Figure 7.6: Exploded view of the cathode electrodes, glued on the I-beams.

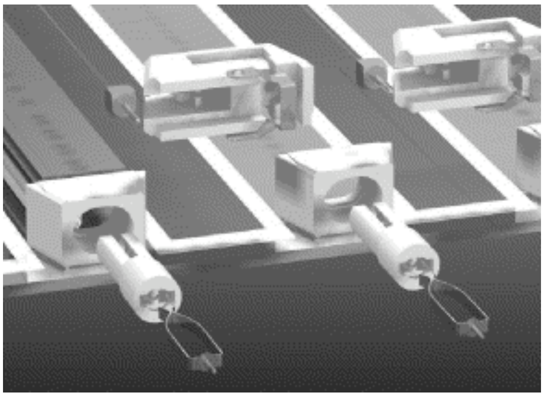

Figure 7.7: Exploded view of the end part of the drift cells showing the different end-plugs and spring contacts for high voltage connections.

are placed on both sides of the I-beams (figure 7.6) 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 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 figure 7.7. The Ar-CO_2 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 [133]. A calculation of the drift velocity using GARFIELD [134] showed that drift velocity saturation occurs between 1 and 2 kV/cm (figure 7.8). This may be compared to the the drift velocity as measured with the Drift Velocity Chamber (VdC) (figure 7.9).

The drift cells will operate at a gas gain of 10^5, 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 [136] (figure 7.10) 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 10^5. 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.

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 section 7.1.2. 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

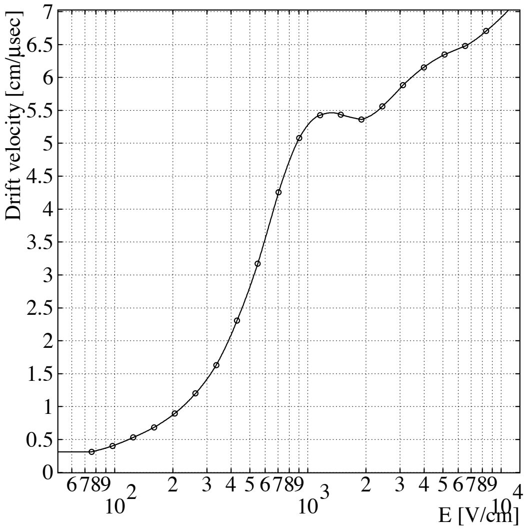

Figure 7.8: Calculated drift velocity (in cm/μs) as a function of the electric field (in V/cm) for a gas mixture Ar/CO_2 (85%/15%). The values obtained are very similar for the 2 different environment pressures, 973 and 1027 hPa.

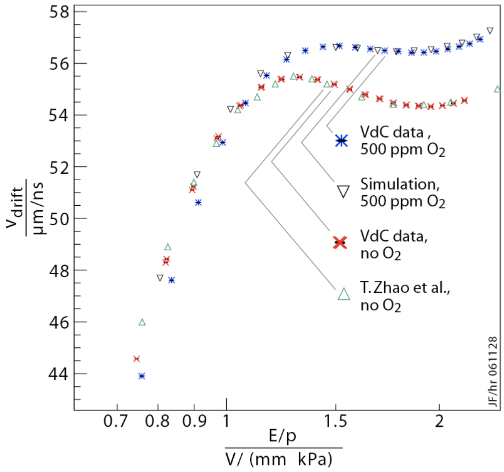

Figure 7.9: Measured and calculated drift velocities (in μm/ns) as functions of electric field and gas pressure for a pure gas mixture Ar/CO_2 (85%/15%) and for a gas mixture with air impurities corresponding to 500-ppm O_2. The measurements were obtained with the VdC, a dedicated reference drift chamber that will be used for drift-velocity monitoring during CMS running. For comparison, results of measurements from [135] for a pure gas mixture, and a simulation with Magboltz [134] for a mixture with impurities are also shown.

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-MΩ 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 mm^2 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 mm^2) outlet hole to guarantee that the same amount of gas is circulated to each

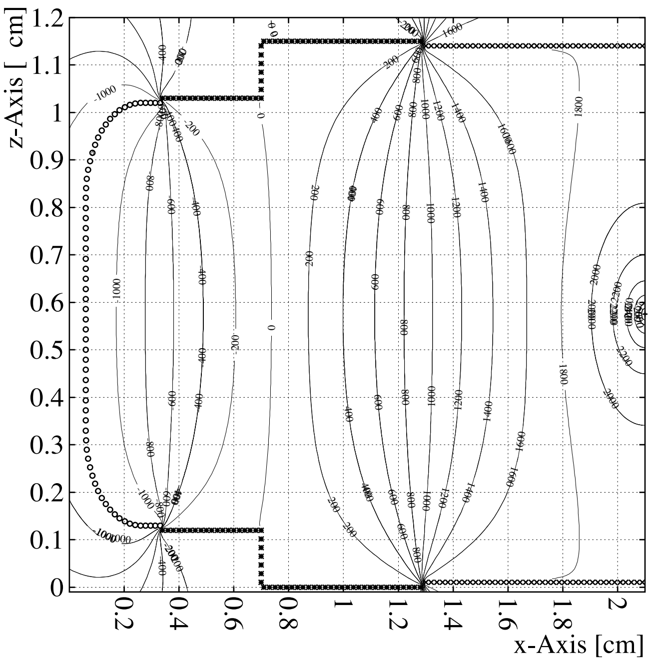

Figure 7.10: Equipotential lines in half of a drift cell. The anode wire is on the right side. The lines are labeled with the potentials in volts (the *x*-axis is perpendicular to the wires on the wire plane, while the *z*-axis is orthogonal to the wire plane).

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 ≈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 (figure 7.11), 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 figure 7.12. 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

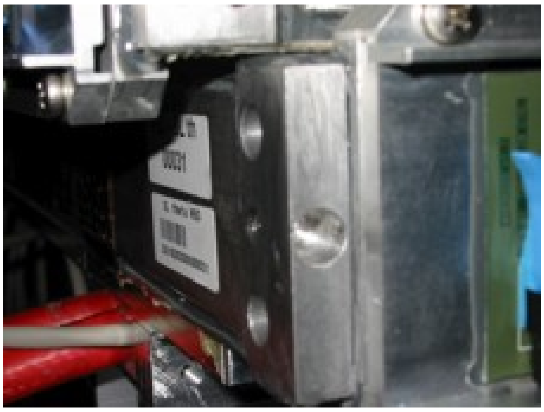

Figure 7.11: Corner blocks of an SL. These pieces also carry the reference marks with respect to which the wire positions are measured.

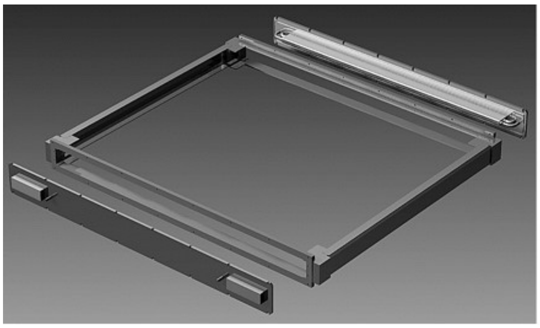

Figure 7.12: A 3-dimensional computer model of the gas enclosure of the SLs.

2008 JINST 3 S08004

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; hence an SL can be fully tested before it is glued to form a DT chamber. SLs are glued to a honeycomb panel (figure 7.4) 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.

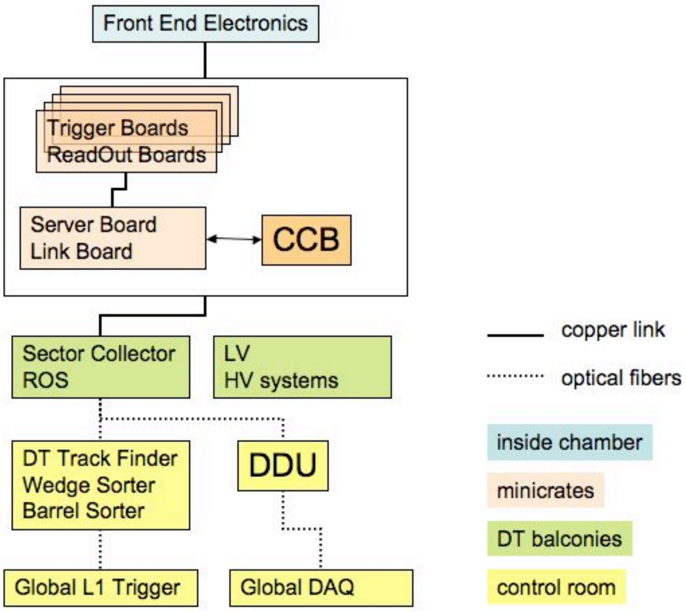

Figure 7.13: Block diagram of the DT electronic system.

7.1.3 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 figure 7.13 and briefly summarized, whereas detailed information will be described in the following sections.

Front-end electronics and HV distribution are 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 section 8.2. 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, section 8.2) and to the Read Out Server (ROS, section 7.1.3) 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) (section 7.1.3) and to the CMS L1 system through the Sector Collector (SC) and the Drift Tube Track Finder (DTTF, section 8.2).

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 μm BiCMOS technology [138]. This chip integrates signal processing for 4 channels (4 drift tubes) plus some ancillary functions in a 2.5×2.5 mm^2 die and 80 000 pieces were produced with a fabrication yield better than 95%. Figure 7.14 shows the ASIC block diagram and the pinout of the TQFP44 package used for it. Each 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 Ω 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)/pF

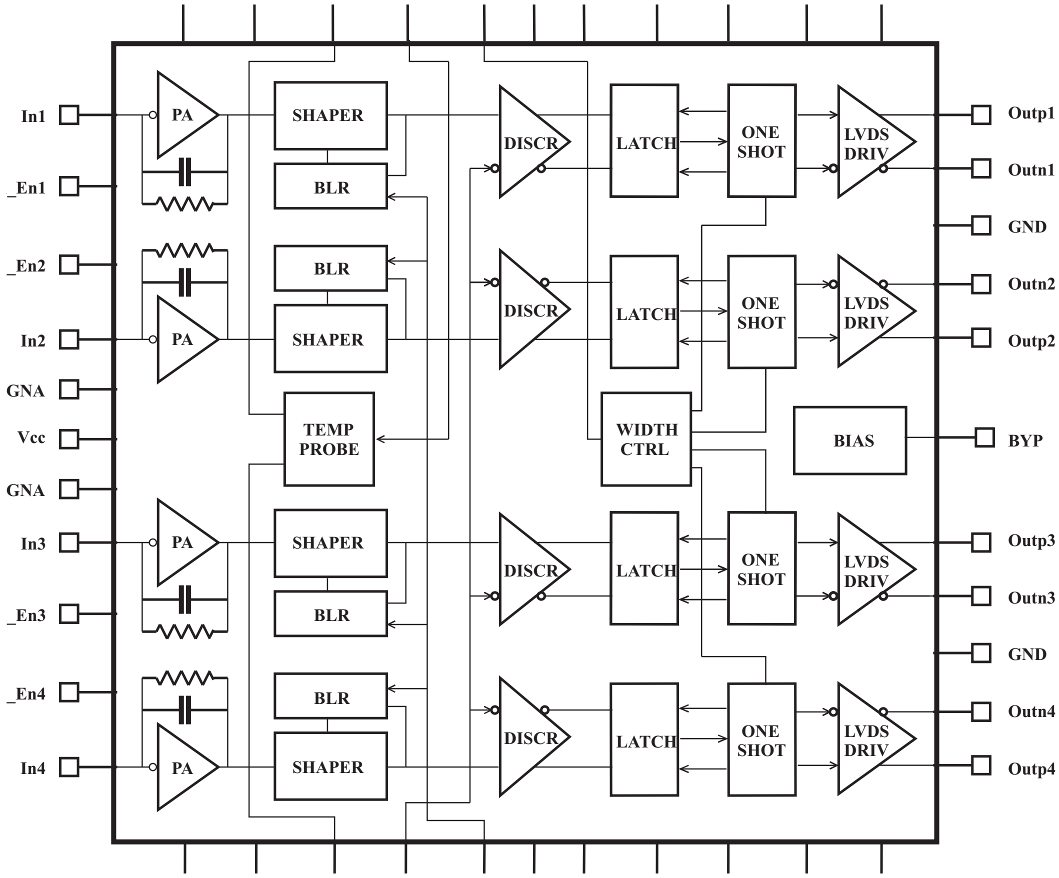

Figure 7.14: Block diagram of the MAD ASIC.

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 (figure 7.15): the gain reduces to 3.4 mV/fC and noise and crosstalk increase to (1850±60)/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).

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.

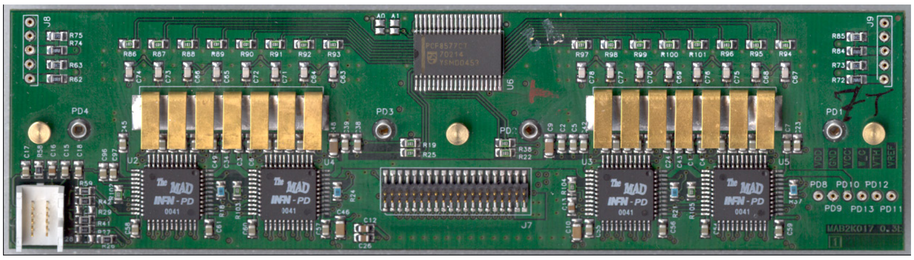

Figure 7.15: Front-end board (FEB).

Finally, the radiation tolerance and overall reliability of the front-end board and associated electronics were investigated [139]. 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°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 [140]), located in minicrates, as described in section 7.1.3, together with the DT muon trigger electronics. Two FTP cables are used to send digitized data from each minicrate to the rack 30 m away 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. The 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 [141], 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 ns = 0.78 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

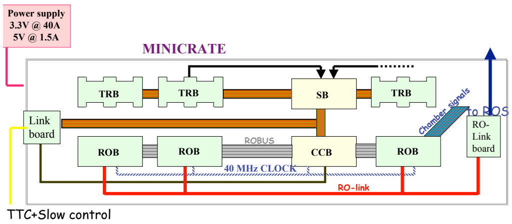

Figure 7.16: Sketch of the read-out and trigger electronics located inside a minicrate.

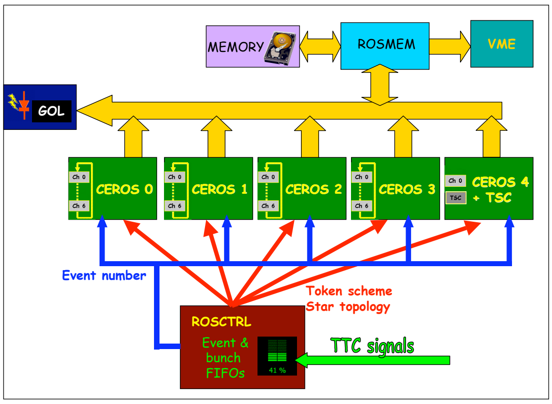

Figure 7.17: Sketch of the splitting of the 25 channels in four groups.

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 ROBs per minicrate. The smallest minicrate has 3 ROBs 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 figure 7.16, Trigger Boards (TRB), located inside the minicrate are connected to the ROBs to receive TTL translated hit signals. As described at the beginning of this section, FTP cables connect the output of the ROBs to the ROS boards [142]. 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 ROBs, 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 Figure 7.17 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 [143]) 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 figure 7.18.

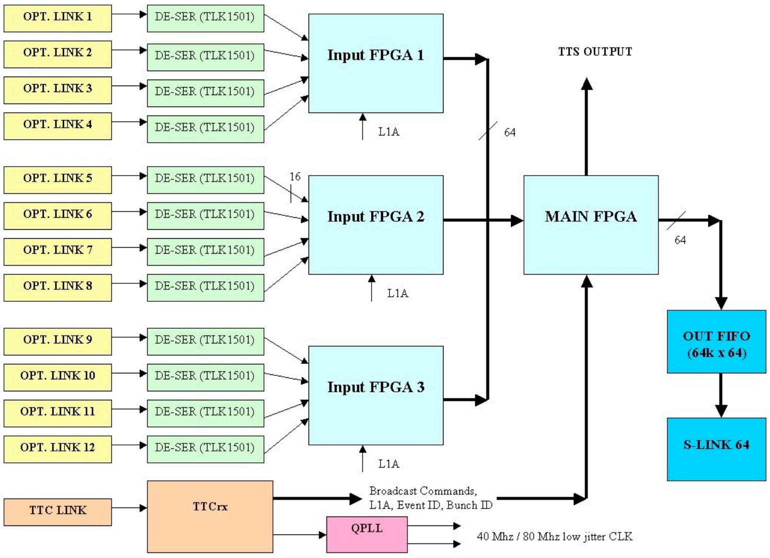

Figure 7.18: Scheme of the DDU architecture.

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 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 output current), a dual LV of ±15V (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 μA.

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

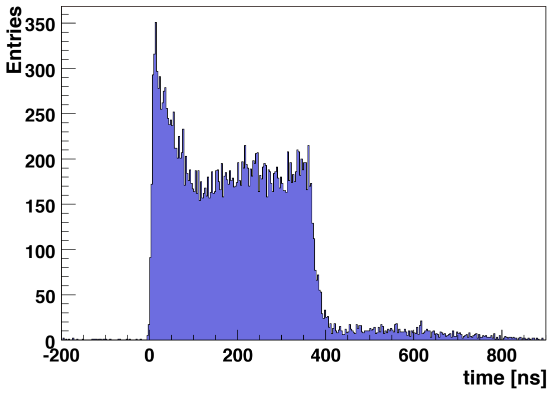

Figure 7.19: Drift Time distribution of a good cell. Wire position corresponds to time 0.

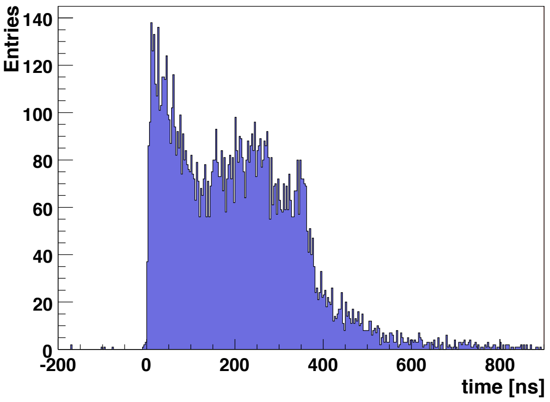

Figure 7.20: Drift Time distribution of a cell with a disconnected cathode.

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:

- VCCMC = 5.8 V, software current limit i_0 = 3 A;
- VDDMC = 4 V, software current limit i_0 = 30 A;
- VCCFE = 5.2 V, software current limit i_0 = 3 A;
- VDDFE = 2.6 V, software current limit i_0 = 4 A;
- VSC = 2.6 V, the current limit depends on the number of SC and ROS boards plugged in the crate.

7.1.4 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 (figures 7.19 and 7.20) and disconnected strips (figure 7.21).

Beside efficiencies, other relevant working parameters are measured from cosmic-ray data, as calibration stability (drift velocity measured to be stable within ±5%), deviation from linear drift parametrization, measured to be well within 100 μm (figure 7.22), wire positions and comparison with CCD measurements during assembly, relative alignment of layers, and noise (typically below 100 Hz per cell, see figure 7.23).

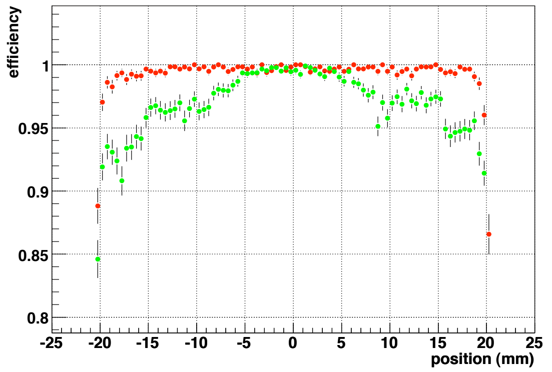

Figure 7.21: Single cell efficiency for a good cell (red dots) and for a cell with a disconnected strip (green dots).

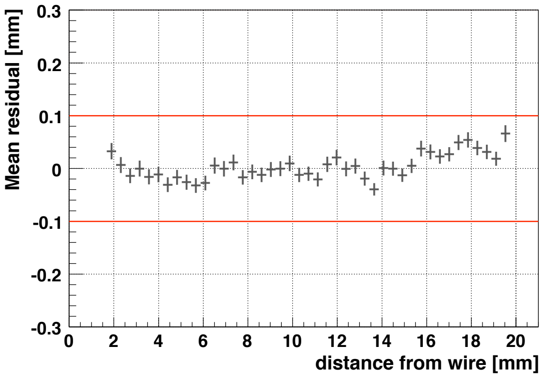

Figure 7.22: Residuals as a function of the distance from the wire, indicating a linear dependence well within ±100 μm in the full cell range.

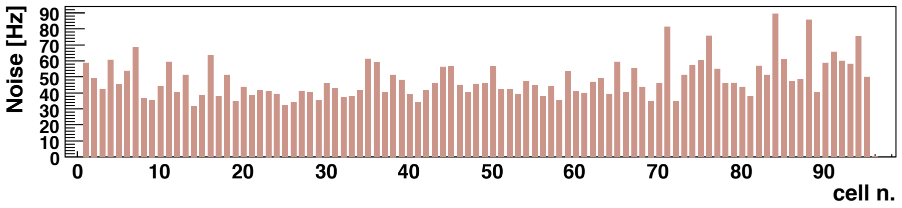

Figure 7.23: Typical hit rate distribution in one layer, as measured during the test of a SuperLayer with cosmic rays. This rate is dominated by noise and is typically below 100 Hz per cell.

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 [144]), 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 section 7.1.4. 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 (section 7.1.3), 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 (section 7.1.3). 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 (section 7.4). 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 mm^2 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 I^{2}C 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 (figure 7.24). 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 was $\sigma < 40\ \mu$m and the position of the forks relative to the chamber

Figure 7.24: The chamber calibration bench in the CERN ISR Lab.

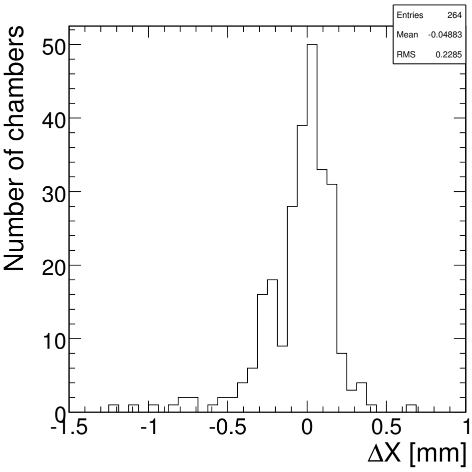

Figure 7.25: Distribution of the ϕ-deviation from the nominal design value for the two Φ-type superlayers.

was σ <70 μm. Both values are within the acceptable range defined by physics requirements. Figure 7.25 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/CO_2 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 unaccessible 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

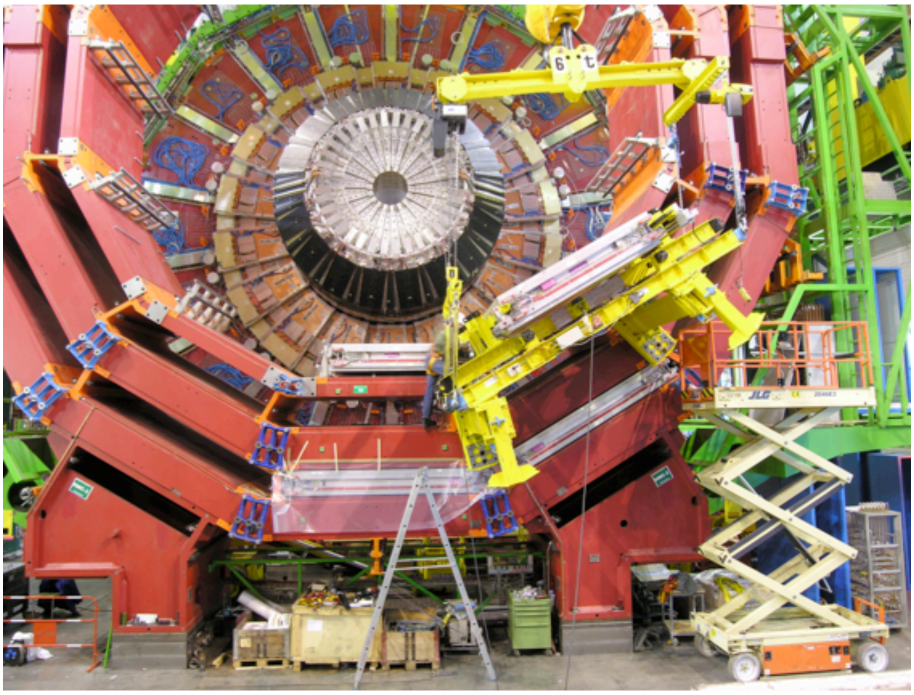

Figure 7.26: Installation of MB1 station on Wheel -2. The yellow frame is the *cradle* used to insert the chamber in its location inside the iron slot.

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 figure 7.8 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 (figure 7.26). 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 figure 7.26.

The first chambers were installed in the bottom sectors of Wheel +2 in July 2004 and the surface installation was completed in Wheel -2 in December 2006. Installation completion, for sectors 1 and 7, which could not be filled on the surface since these parts of the wheel were reserved for handling during the heavy lowering operation, was made underground.

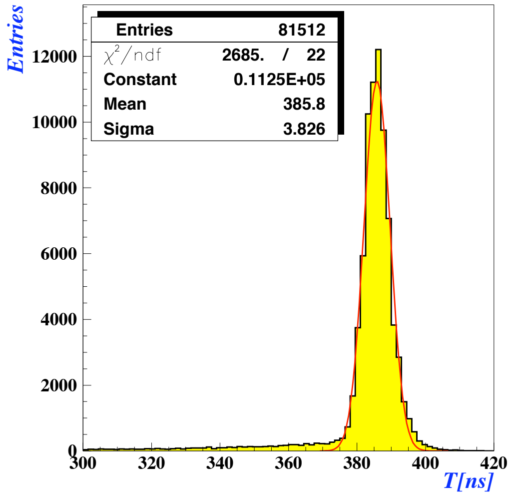

Figure 7.27: Gaussian fitted *MT* distribution with $\sigma_t = 170\ \mu$m. The position of the *MT* peak allows the determination of an average drift velocity of 54.4 μm/ns.

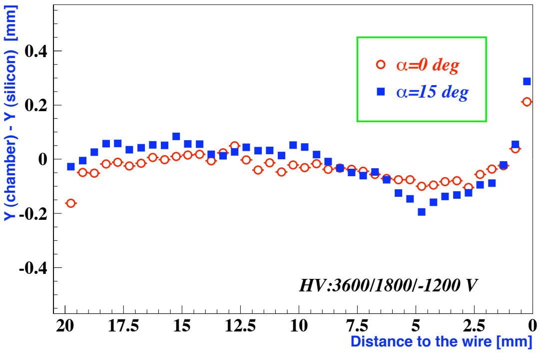

Figure 7.28: Deviation from linearity as a function of the distance to the wire for tracks with angles of incidence α=0° and 15°. The cathode is centred at 21 mm.

7.1.5 Chamber performance

Chamber and trigger performances have been thoroughly analyzed at various stages, on prototypes before mass production (with and without external magnetic field [145]), on final chambers with test beams and with the CERN Gamma Irradiation Facility [144, 146, 147], 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 performance under different conditions. Single cell spatial resolution could be determined [144] 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 3 consecutive, staggered layers. The smoothness of the drift time box and the fast drop of the trailing edge (figure 7.19) are both signs of the saturation of the drift velocity. Under the assumption that the time resolution is the same in all layers $\sigma_t = \sqrt{\frac{2}{3}} \cdot \sigma_{\mathrm{MT}}$, one can easily observe an average spatial resolution of 170 μm (figure 7.27). 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 figure 7.28 deviations from linearity are within ± 100 μm, which corresponds to a trigger jitter smaller than 5 ns.

Silicon Beam Telescope data can also be used to measure precisely the chamber efficiency (figure 7.29). 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

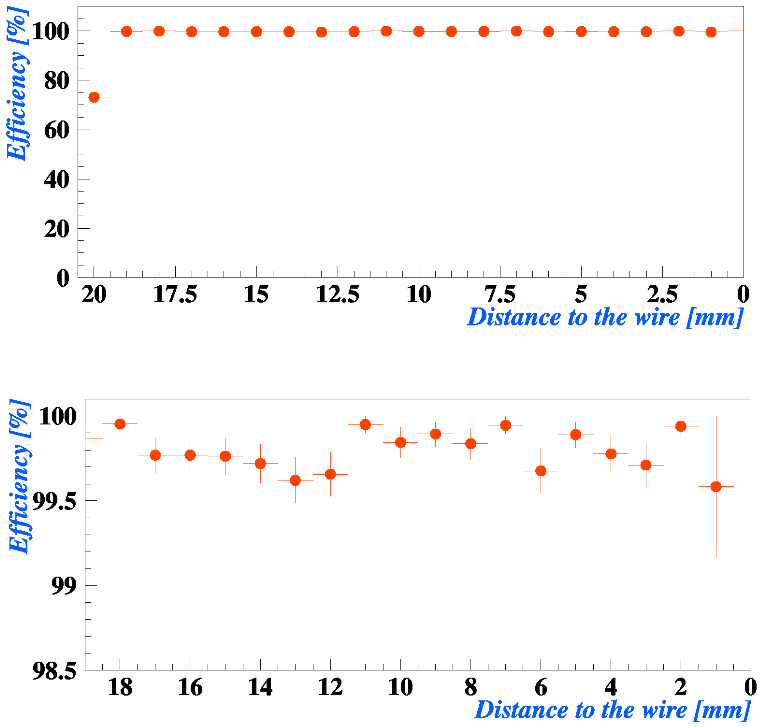

Figure 7.29: Efficiency as a function of the distance to the wire (top), for tracks orthogonal to chamber surface; (bottom) with an expanded scale excluding the I-beam region.

99.5%. The typical intrinsic average noise, as measured during chamber construction or during a test beam in dedicated random trigger runs [146], is shown in figure 7.30. It is reasonably stable, at values of ≈50 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 figure 7.31 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 any DT chamber during normal LHC operations. Since the chambers are operated in the iron yoke of CMS, where 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, figure 7.32), and tested both at dedicated muon test beams where the chambers were operated inside a magnetic field [145, 146], and during the CMS Magnet Test and Cosmic Challenge, with comsic rays. The distortion of the electron drift lines caused by a field of 0.5 T parallel to the wires can be seen in figure 7.33. 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

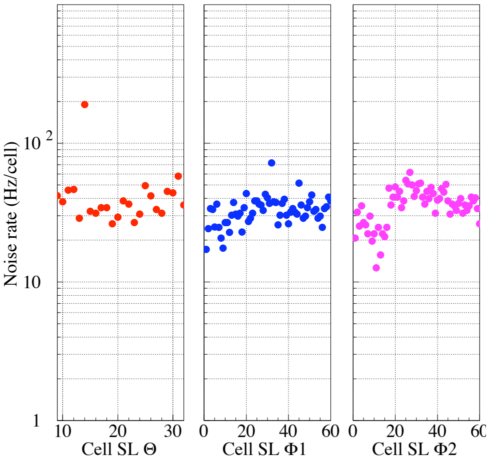

Figure 7.30: The noise cell occupancy as a function of the cell number, for the three SLs.

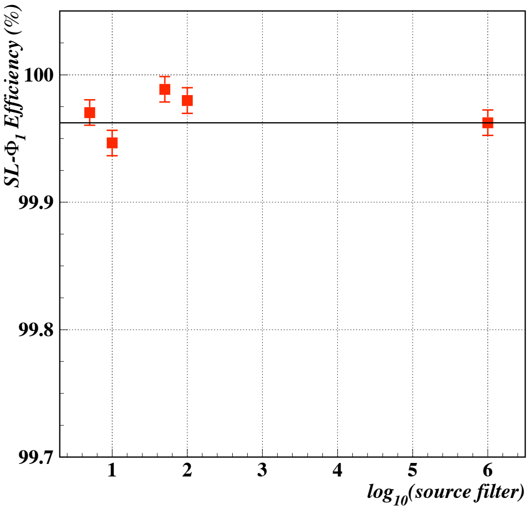

Figure 7.31: The cell efficiency in superlayer Φ_1 as a function of the irradiation filter value. The smallest filter value provides a background a factor 2 larger than the maximum one to be expected during LHC operation in any of the DT chambers in CMS.

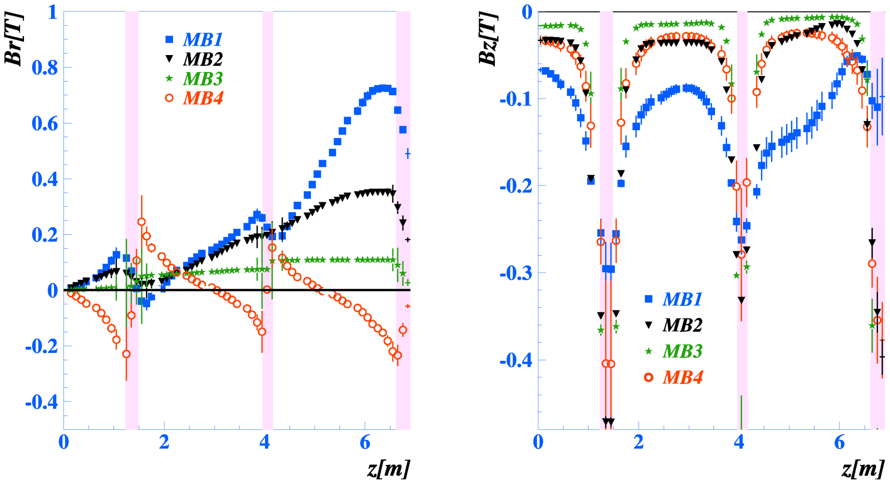

Figure 7.32: Radial (B_r) and longitudinal (B_z) components of the CMS magnetic field in the regions where the barrel chambers are placed as a function of the position along the beam direction (the centre of the detector is at z=0). Vertical bands indicate the separation between chamber wheels (in these particular regions B_z becomes significant). The biggest B_r values (0.7–0.8 T) occur in the MB1 region near the endcaps.

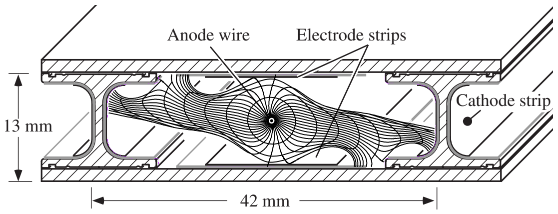

Figure 7.33: Simulation of the distortion produced in the drift lines by a 0.5 T magnetic field parallel to the wires.

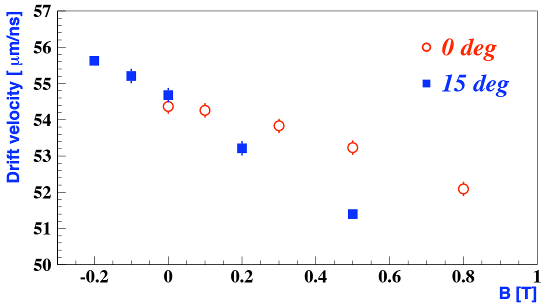

Figure 7.34: Drift velocity for several magnetic fields for perpendicular (0°) and inclined (15°) tracks.

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 7.34 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 μm 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 [147, 148]. 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 figure 7.35. 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 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 figure 7.35, superimposed on results without iron absorber. Ghosts are copies 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

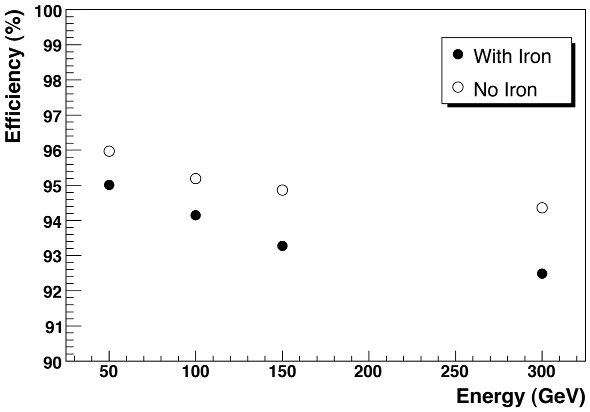

Figure 7.35: BX identification efficiency in a muon station in single muon events, as a function of the incident muon momentum, for events with and without the iron absorber placed in front of the muon station.

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 (section 8.2). 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 figures 7.36 and 7.37 as a function of the incident muon momentum, for events with and without the iron absorber. The performance of the Φ Track Finder (PHTF, section 8.2) was also tested within the same muon test beam at CERN [148]. The PHTF was used to reconstruct muon trigger candidates using both muon stations, using ϕ-view local trigger primitives. Figure 7.38 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 7.39 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

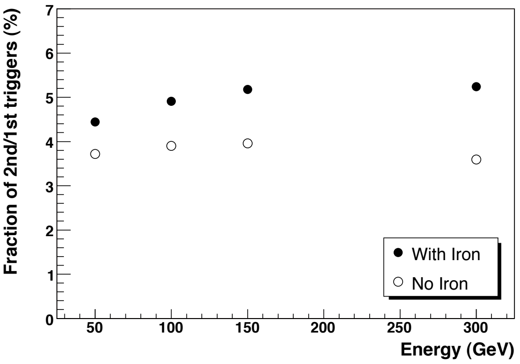

Figure 7.36: Fraction of ghost triggers observed in a muon station in single muon events, defined as the ratio of the number of second tracks over the number of first tracks, delivered by the local trigger at the correct BX, as a function of the muon momentum, for events with and without the iron absorber.

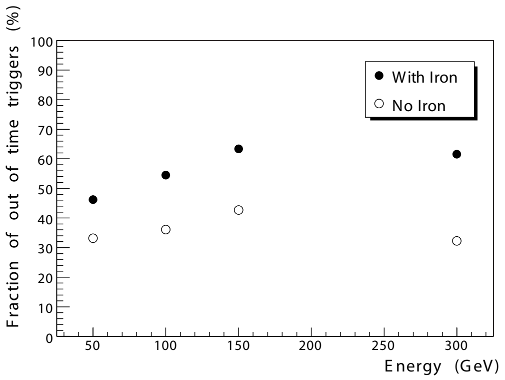

Figure 7.37: The fraction of out-of-time triggers in a muon station in single muon events, defined as the number of out-of-time trigger segments divided by the number of selected single muon events, as a function of the muon momentum, for events with and without the iron absorber.

respectively. The correct BX is 24. The PHTF efficiency for HH coincidences is 99.7±0.1% and is practically 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 figure 7.37), 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 of identifying potential problems before final cabling. Given the previous testing stages, it was mainly a test of the electronics and connections. Possible damage 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 consisted in the following steps:

1. connection to the power supplies and to the test-stand, which handled the trigger and read-out, in order to verify the functionality of the connectors and cabling of the chamber;
2. check of the minicrate performance by testing the internal connections and by monitoring the electronics boards inside the minicrate;

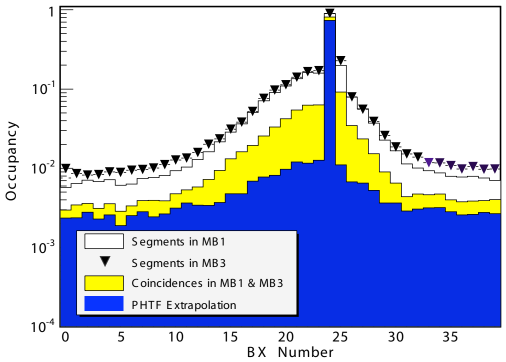

Figure 7.38: Distribution of the BX assigned to the tracks found by the PHTF. The BX is correctly identified when its value is 24. Superimposed are the distributions of the same quantity determined independently by the local trigger in MB1 and MB3, as well as the distribution of the determined BX when a trigger segment at the same BX was delivered in coincidence in MB1 and MB3.

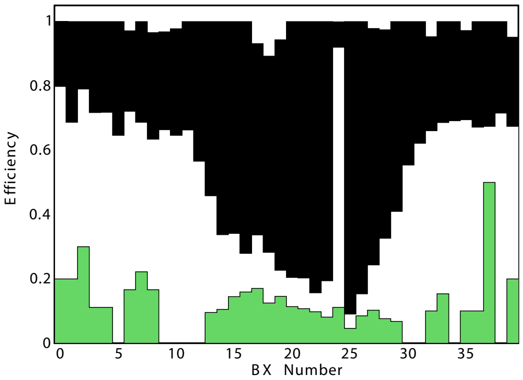

Figure 7.39: Efficiency to reconstruct a trigger track by the PHTF, as a function of the BX (white), in events with a two stations (MB1-MB3) coincidence. The correct BX is 24. Superimposed are the efficiency to reconstruct a trigger track when there is a coincidence of two trigger segments both of HH quality (black), and a coincidence of two trigger segments both of L quality (light green).

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 ≈1–2 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 varied 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 were not tested at this stage since they required full cabling to the tower electronics.

The chamber orientation with respect to incoming cosmic-rays is purely horizontal only in sectors 4 and 10. This yields a reduction in the 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 figure 7.40, where the occupancy per wire is summed over the four layers of superlayer Φ2. The reduced geometrical acceptance at the chamber edges is caused by the combination of two effects: i) the direction of cosmic rays, which is mainly vertical, and ii) the shielding of the iron yoke on soft muons.

Data recorded with 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

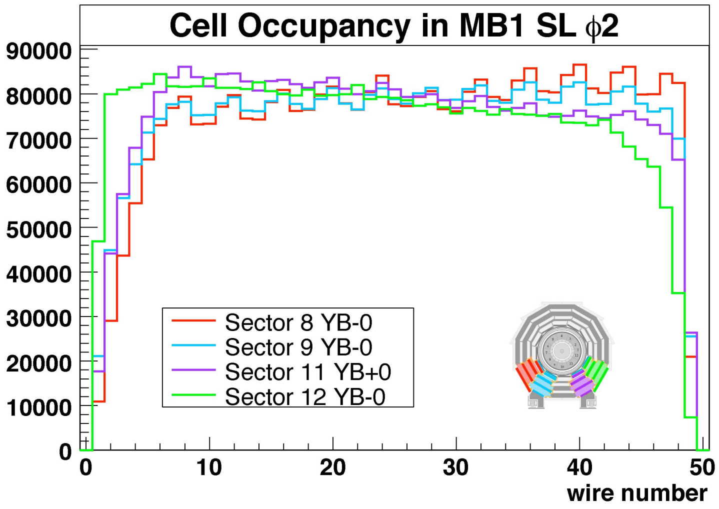

Figure 7.40: Sum of the occupancy of the four layers in MB1 Φ2 superlayer for sectors 8, 9, 11 and 12. The sectors have different inclinations as shown on the right. Material inside the yoke along with the iron yoke itself shield partly the soft cosmic ray muons.

from reconstructed tracks with hits found in the traversed cell or its 2 neighbours (N_h), normalized to the total reconstructed tracks traversing the considered layer (N_track):

$$\varepsilon_\mathrm{Layer} = N_\mathrm{h}/N_\mathrm{track} \tag{7.1}$$

requiring $\geq$5 hits in ϕ and $\geq$3 hits in z. Figure 7.41 (left panel) 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 figure 7.41 (right panel).

Similarly the reconstruction efficiency is determined as the fraction of reconstructed tracks when requiring $\geq$7 hits per track, an example of which can be seen in figure 7.42. 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).

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 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

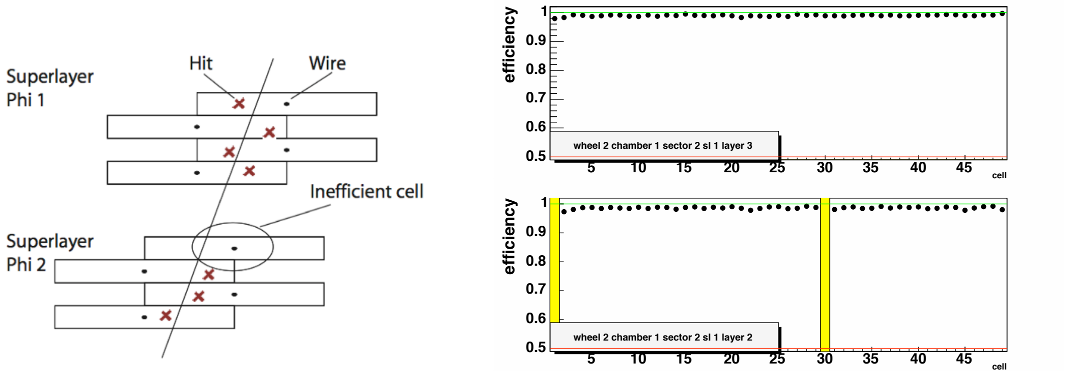

Figure 7.41: Cell efficiencies for two of the four layers of Φ1, MB1 sector 2.

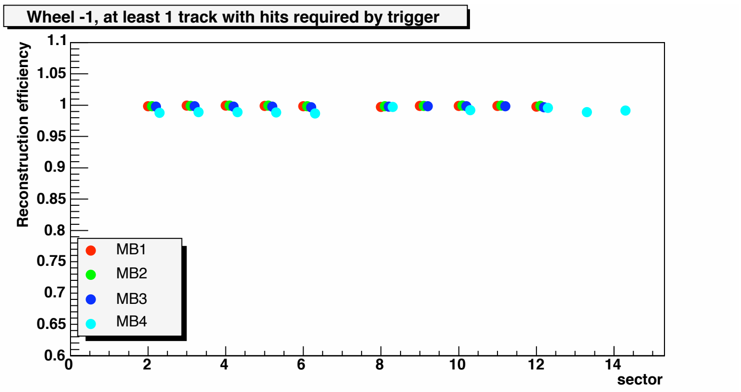

Figure 7.42: Reconstruction efficiency with 7 points in both Φ-type superlayers, according to the trigger condition, requiring either two 4-hit segments in both of the two Φ-type superlayers or one 4-hit and one 3-hit segment.

- check the effect of the 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 the effect of magnetic field 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;

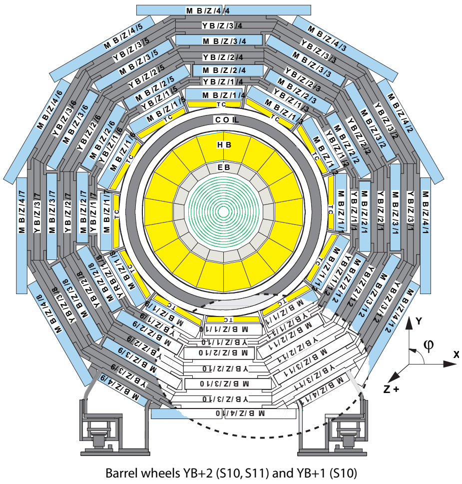

Figure 7.43: The MTCC exploited in the barrel region three sectors in wheels YB+1 (sector 10) and YB+2 (sectors 10, 11) instrumented with DTs and RPCs.

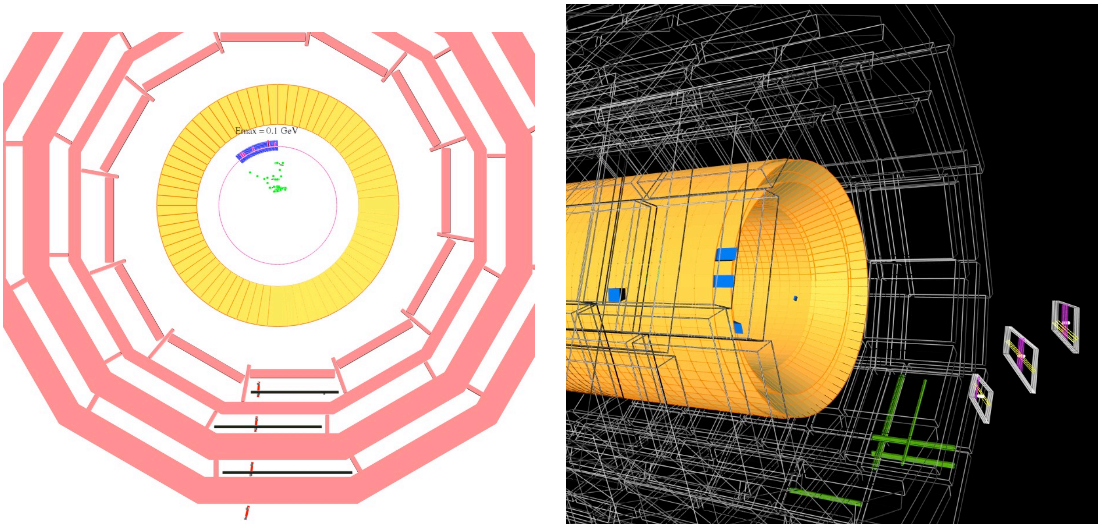

Figure 7.44: Event displays from the MTCC. Left panel: muon reconstructed in a DT sector in conjunction with Tracker activity. Right panel: muon track passing through both DT and CSC chambers.

- 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 (figure 7.43). This accounted for 14 DT chambers, corresponding to about 10 000 channels. Beside the cross-check of chamber performance previously carried out in test beams, MTCC data provided a unique opportunity to test the reconstruction algorithms for different magnetic field strengths (figure 7.44(left)) and to observe for the first time tracks combined in different detectors (figure 7.44(right)). As described in section 7.1.5, because of the radial component of the magnetic field between the 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) were collected at several values of the B field. Data were collected at 0 T (as a reference), then at 2, 3, 3.5, 3.8, and 4 T (93 million events at 3.8 and 4 T), which allowed a detailed mapping of the Lorentz angle effect in an MB1 and an 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 centre of the CMS barrel (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 magnetic field mapping operations) effort was put in tools for fine inter-synchronization 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 desynchronization by a few nanoseconds.

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 section 7.1.5, 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 [148] 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 (section 8.2) failing

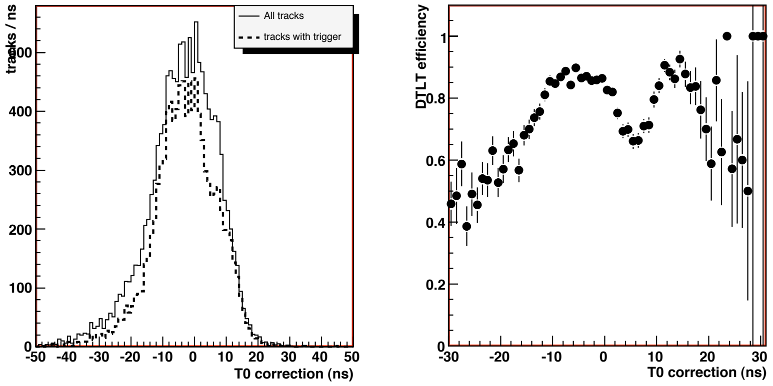

Figure 7.45: Left panel: distribution of t_0 for all track segments in a station, for events triggered by RPC (solid line). The distribution of the same quantity for events also triggered by the DT is superimposed (dashed line). Right panel: DT local trigger efficiency as a function of the quantity t_0, obtained by the ratio of the two histograms shown on the left.

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 7.45 (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 figure 7.45 (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 section 7.1.5. 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 7.46 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

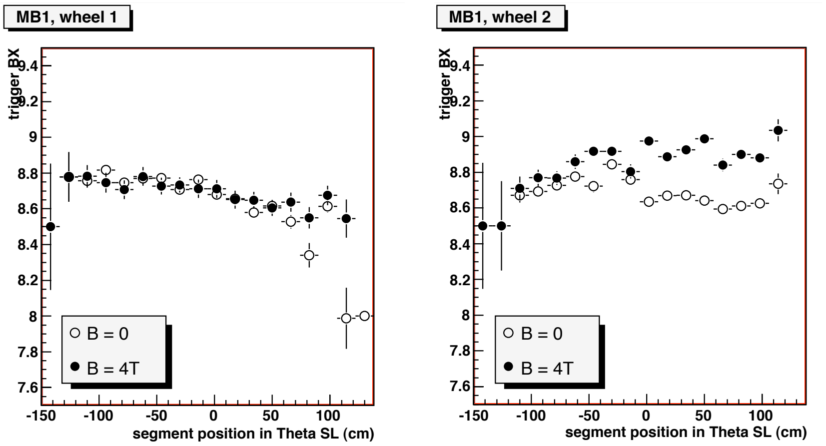

Figure 7.46: BX determined by the DTLT as a function of the track position in the *z*-direction of the muon station, with and without magnetic field, for MB1 in Wheel 1 on the left, where no effect of the B-field is expected. On the right, the same quantities are shown for MB1 in Wheel 2, where an influence of the magnetic field on the drift velocity is expected.

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 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.

7.2 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 (figures 7.47 and 7.48). 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 $|\eta| < 1.6$.

The CSCs are multiwire proportional chambers comprised of 6 anode wire planes interleaved among 7 cathode panels (figure 7.49). 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 [149], the muon coordinate along the wires (ϕ in the CMS coordinate system) is obtained by interpolating charges induced on strips (figure 7.50). The largest chambers, ME2/2 and ME3/2, are about 3.4×1.5 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 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.

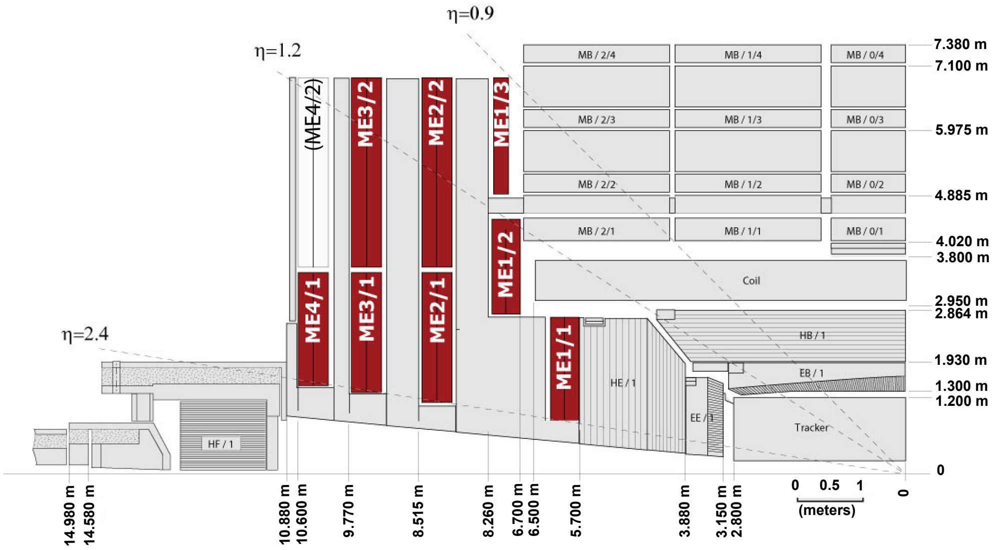

Figure 7.47: Quarter-view of the CMS detector. Cathode strip chambers of the Endcap Muon system are highlighted.

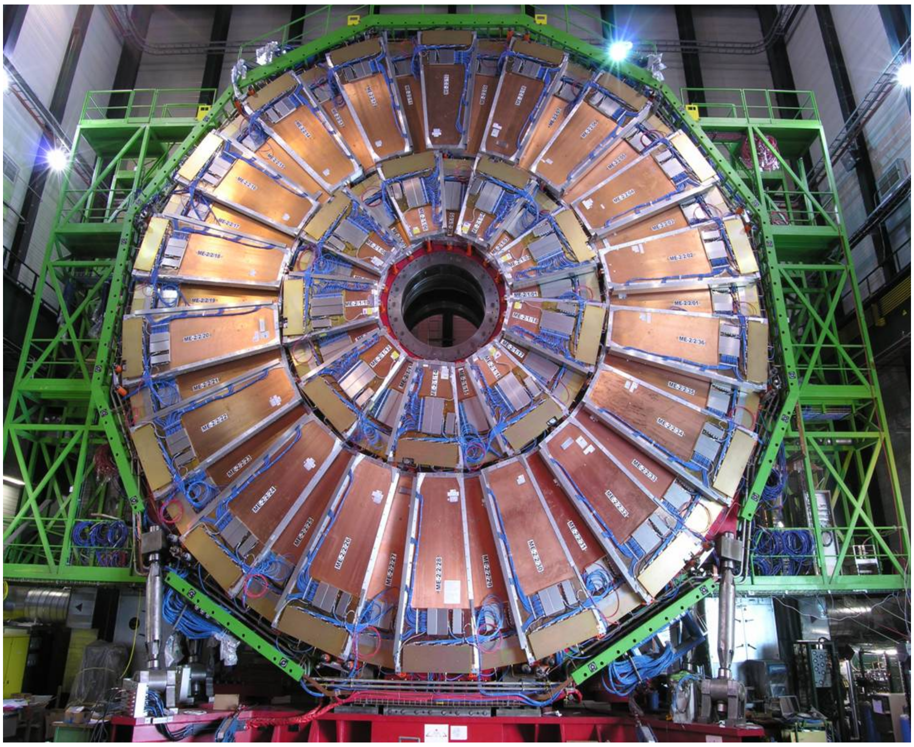

Figure 7.48: The ME2 station of CSCs. The outer ring consists of 36 ME2/2 chambers, each spanning 10° in ϕ, and the inner ring of eighteen 20° ME2/1 chambers. The chambers overlap to provide contiguous coverage in ϕ.

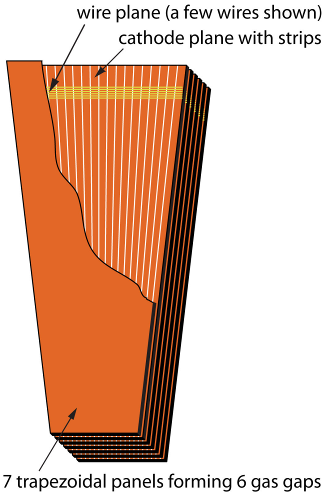

Figure 7.49: Layout of a CSC made of 7 trapezoidal panels. The panels form 6 gas gaps with planes of sensitive anode wires. The cut-out in the top panel reveals anode wires and cathode strips. Only a few wires are shown to indicate their azimuthal direction. Strips of constant $\Delta\phi$ run lengthwise (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.

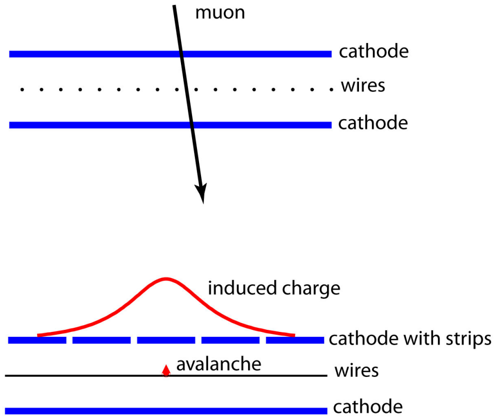

Figure 7.50: A schematic view of a single gap illustrating the principle of CSC operation. By interpolating charges induced on cathode strips by avalanche positive ions near a wire, one can obtain a precise localisation of an avalanche along the wire direction.

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.
- About 75 μm off-line spatial resolution in r-ϕ for ME1/1 and ME1/2 chambers and about 150 μm for all others.

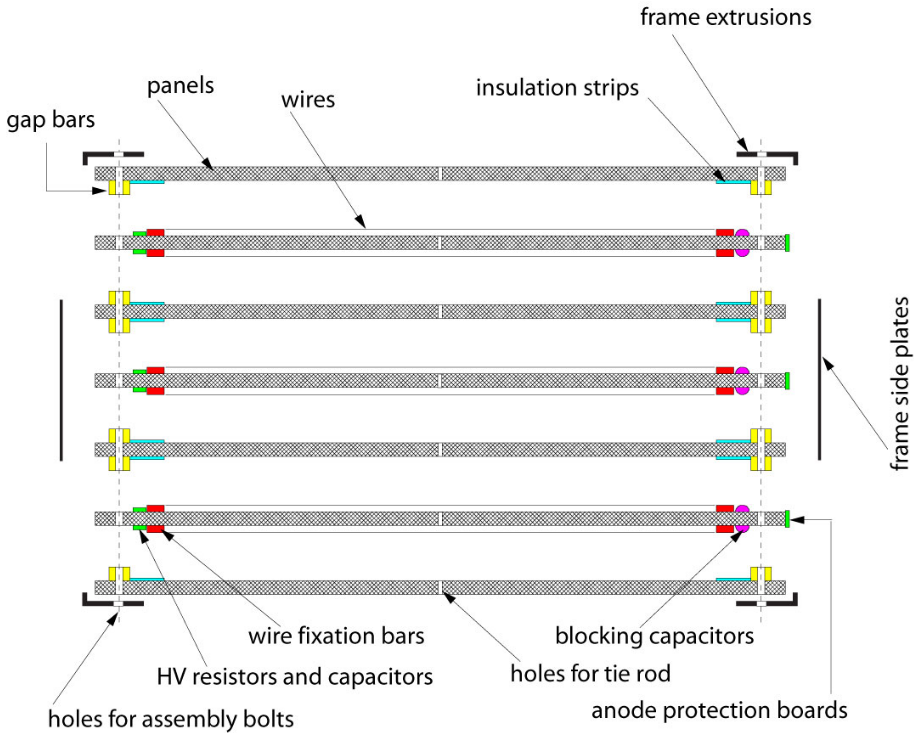

Figure 7.51: Mechanical design of the CMS cathode strip chambers (exploded view).

7.2.1 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 (figure 7.51). 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-μm-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 figure 7.51) 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 ±300 μm 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 μm (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 figure 7.51) 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 μm 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 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 ± 150 μm. 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 cm^3/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 $40\%\mathrm{Ar} + 50\%\mathrm{CO_2} + 10\%\mathrm{CF_4}$. The CO_2 component is a non-flammable quencher needed to achieve large gas gains, while the main function of the CF_4 is to prevent polymerisation on wires. A detailed discussion of the gas optimisation can be found elsewhere [150].

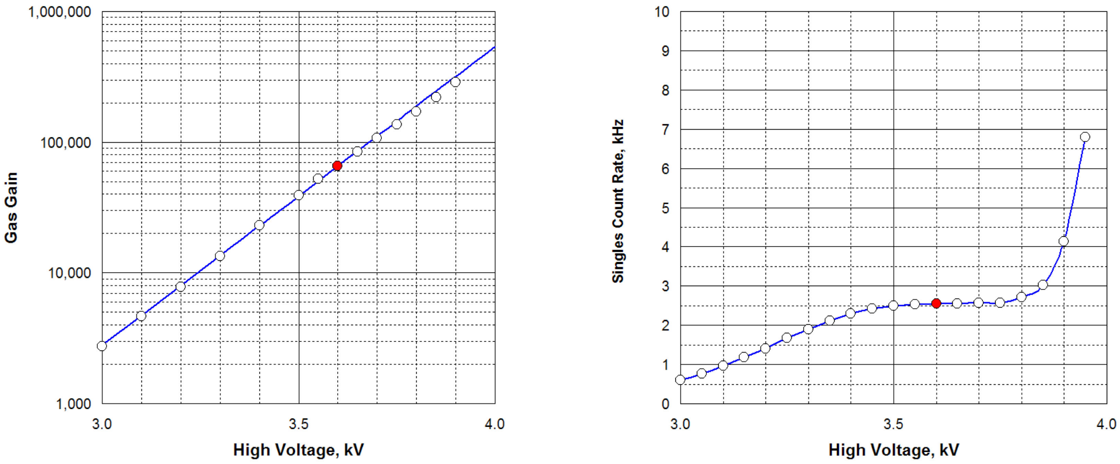

Figure 7.52: Left panel: CSC gas gain vs. high voltage. Right panel: ME2/1 chamber singles rate vs. high voltage (the overall sensitive area of all 6 planes in this chamber is ≈9.5 m^2).

Figure 7.52 (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 (figure 7.52 (right)).

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 μm, 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 $\alpha_L = 29^\circ$ (figure 7.53). 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.

7.2.2 Electronics design

Figure 7.54 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 figure 7.55 (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 [151].

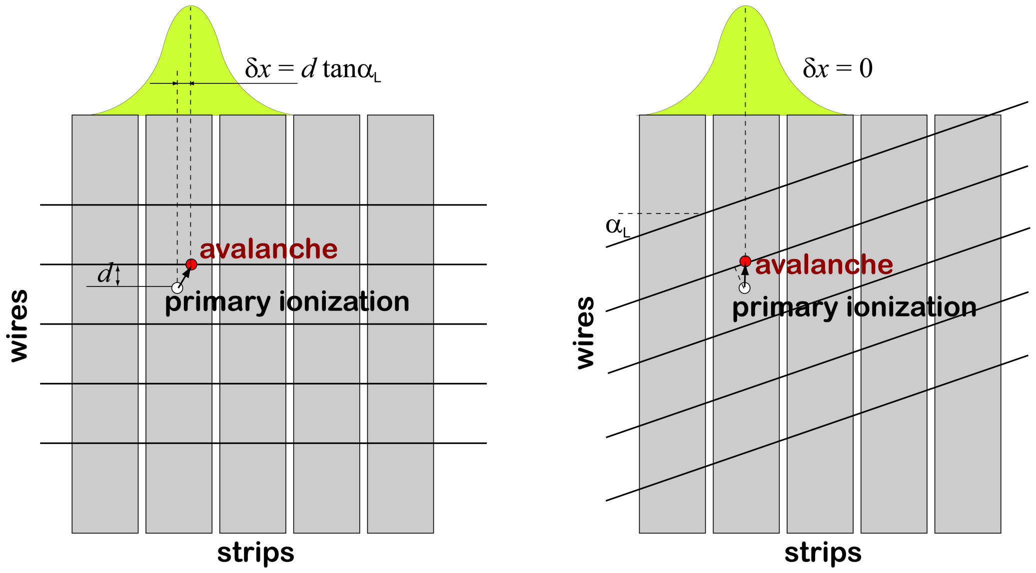

Figure 7.53: Left panel: if the ME1/1 wires were not tilted, ionisation electrons, as they drift toward the anode wires in the strong magnetic field normal to the plane of the drawing, would be carried sideways by the Lorentz force. The direction and size of the shift would depend on whether the electrons 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.

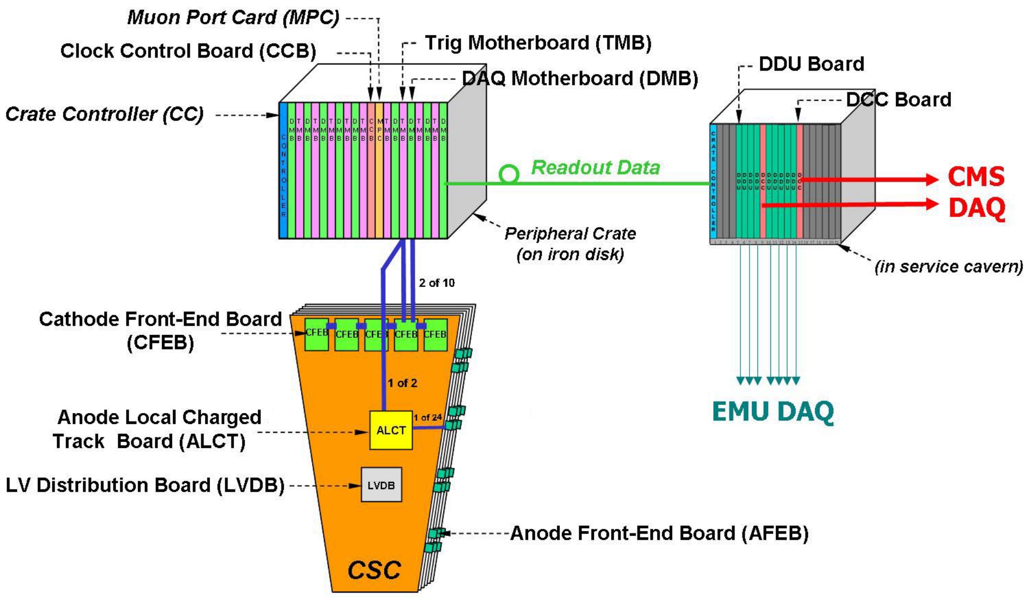

Figure 7.54: Schematic layout of the CSC trigger and read-out electronics.

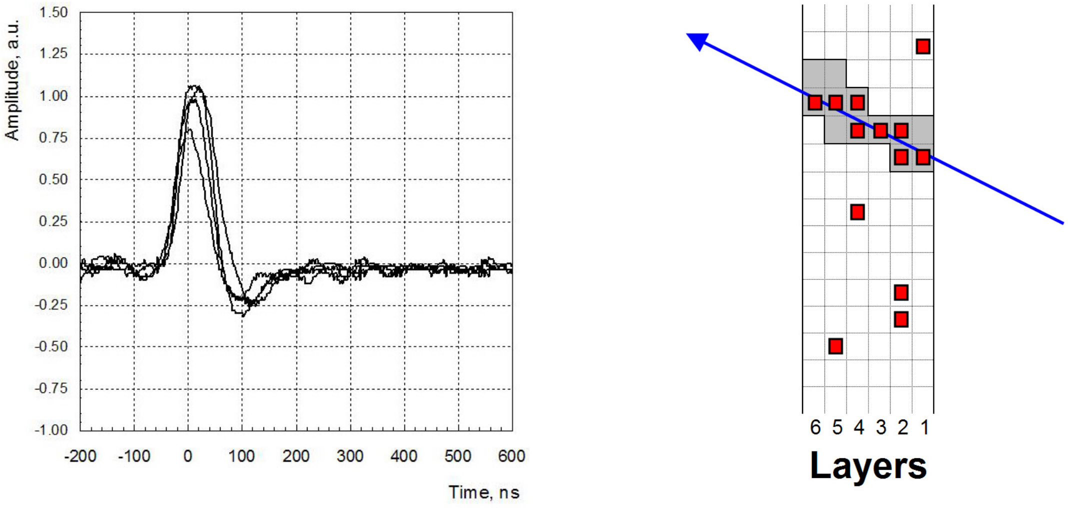

Figure 7.55: Left panel: muon signals as seen at the AFEB amplifier output. Right panel: a schematic event display showing anode signals in the 6 planes of a CSC (small dark squares). The ALCT board FPGA logic is programmed to scan the chamber and search for hits falling inside predefined patterns (grey 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.

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 μs. 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 7.55 (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 is expected. However, these hits, being completely uncorrelated, would not typically line up to form track-like patterns. Found patterns, called ALCTs, are trigger primitives. They are transmitted further downstream to the muon Level-1 trigger electronics that builds muon track candidates from these primitives. The time it takes to form an anode track trigger primitive is 225 ns (including drift time). Each ALCT board can find up to 2 such patterns per bunch crossing, which is adequate for the expected chamber track occupancy at the nominal LHC luminosity.

For the cathode strips, 1 cathode front-end board (CFEB) serves (6 planes)×(16 strips) = 96 channels and has 6 parallel chains of the following chips (figure 7.56 (left)): 16-channel amplifier-shaper ASIC, 16-channel switched capacitor array (SCA) ASIC, 12-bit 1-channel ADC, and 16-channel comparator ASIC. There are 4 to 5 CFEBs per chamber.

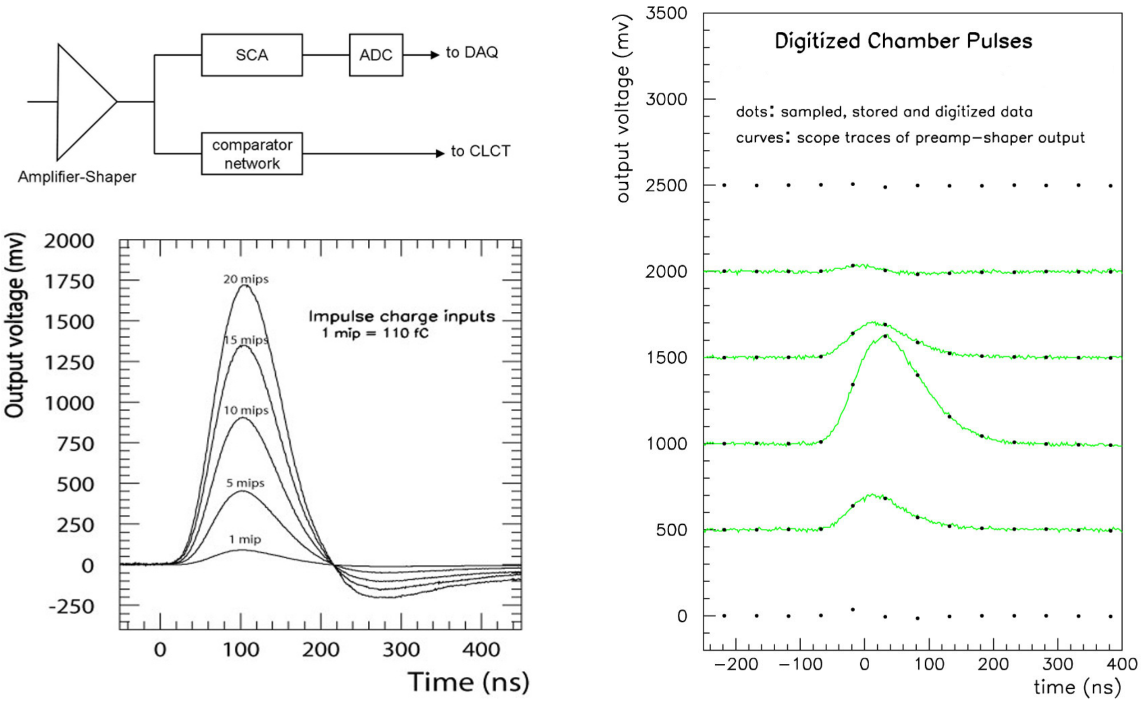

Figure 7.56: Left panel: basic functional diagram of a CFEB and the CFEB amplifier-shaper response to a δ-function input pulse. The undershoot is intended to compensate for the long tail present in muon hit signals resulting from the slow drift of positive ions away from the anode wires. Right panel: 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.

The amplifier-shaper ASIC has 100-ns shaping time and a sensitivity of 0.85 mV/fC over a linear range up to 1 V. The equivalent noise level at ≈300-pF strip capacitance is typically 1.5 fC. The shaping is based on a semi-Gaussian transfer function with an undershoot designed to compensate for the $1/t$ signal tail due to the slow drift of positive ions. After convolution with the current pulse produced in a chamber by a muon, the amplifier-shaper signal peaks around 150 ns and has no tail (figure 7.56 (right)). The CFEB sees about 8% of the total avalanche charge, i.e., about 100 fC on average.

The output from this chip is split into 2 pathways. One leads to the SCA chip [152], which samples the waveform of each strip signal every 50 ns in sync with the LHC clock and stores this analog information on its capacitors. The depth of this analog memory is 96 capacitor cells per channel, or 96 × 50 ns = 4.8 μs. Upon receiving the L1A command 3.2 μs following a collision, 8 or 16 consecutive samples from the proper time range among the SCA capacitors are retrieved and digitised individually by the 12-bit flash ADCs. The digital information is passed to the DAQ via an intermediate digital data buffer. For the digitisation and subsequent read-out by the DAQ to happen, the L1A signal must be in coincidence with the *cathode local charged track* (CLCT) primitive decision described below.

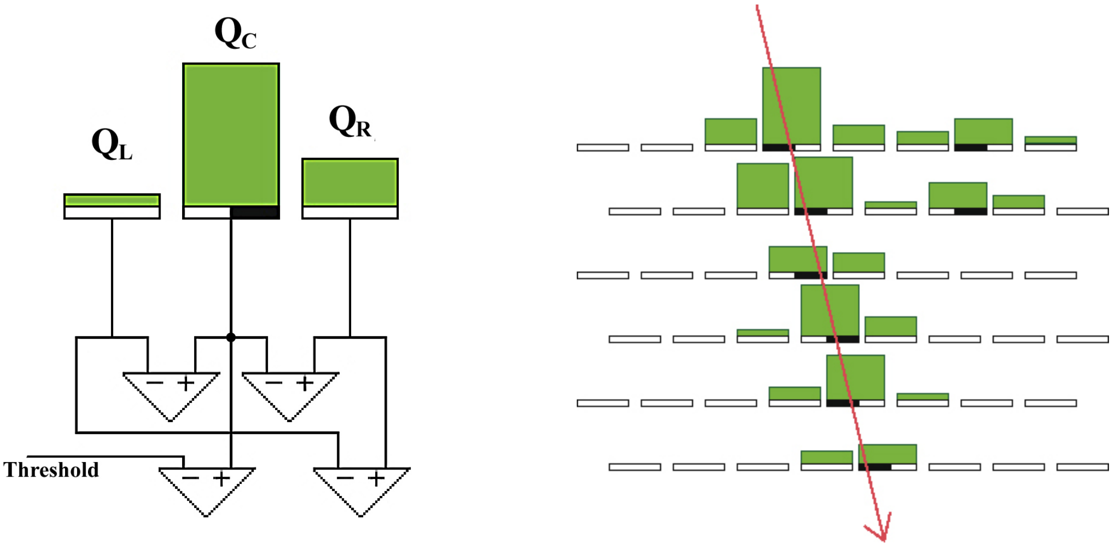

Figure 7.57: Left panel: a simplified schematic of the idea behind the comparator network. For each group of 3 adjacent strips, comparators compare the central strip signal Q_c with a threshold and with the central-to-left $Q_c - Q_l$, central-to-right $Q_c - Q_r$, and right-to-left $Q_r - Q_l$ strip signal differences. If $Q_c >$ 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.

The second amplifier-shaper output goes to the comparator network. This chip compares signals on triplets of adjacent strips at the time when signals reach their maximum amplitude. By means of such comparisons, the comparator network can identify a muon hit location to within one half of a strip width, independent of the signal amplitude, the induced charge shape (as long as it is "bell"-like), and the strip width itself [153] (figure 7.57 (left)).

Comparator half-strip hits are sent to the trigger motherboard (TMB). Like the ALCT board, the TMB searches for patterns of half-strip comparator hits that would be consistent with muon tracks of interest (figure 7.57 (right)). There is 1 TMB per chamber and up to 2 CLCTs per bunch crossing can be found per TMB. As in the ALCT pattern 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 mounted on 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×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×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 μA 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 μA current for individual channels as long as the average consumption does not exceed 40 μA per channel. The maximum expected current at the design LHC luminosity for the most-loaded HV segment is $<$10 μA.

7.2.3 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 7.58 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

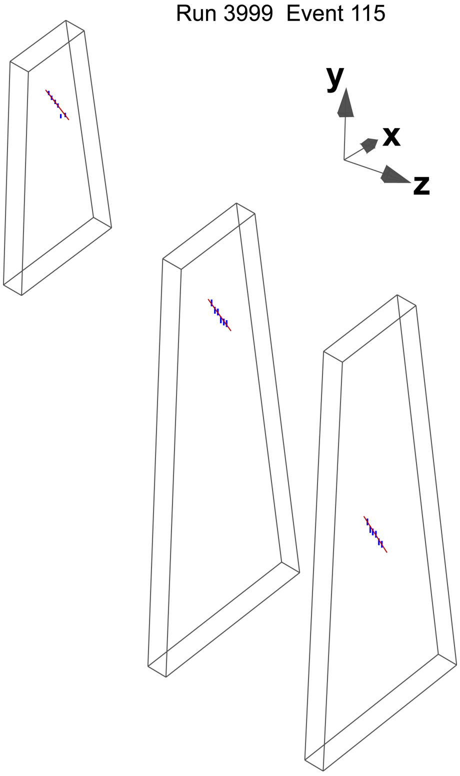

Figure 7.58: Part of a CMS event display showing a muon event detected by CSCs during the MTCC. Only those chambers containing muon hits are displayed.

Cs-137 source of 0.7 MeV photons with an intensity of $\approx 0.7 \times 10^{12}$ 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×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 figure 7.59 (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, 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×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 (99.93±0.03)%. For the 0.07% of events with missing 2D-LCTs in ME2/2 chambers, the majority of tracks (reconstructed

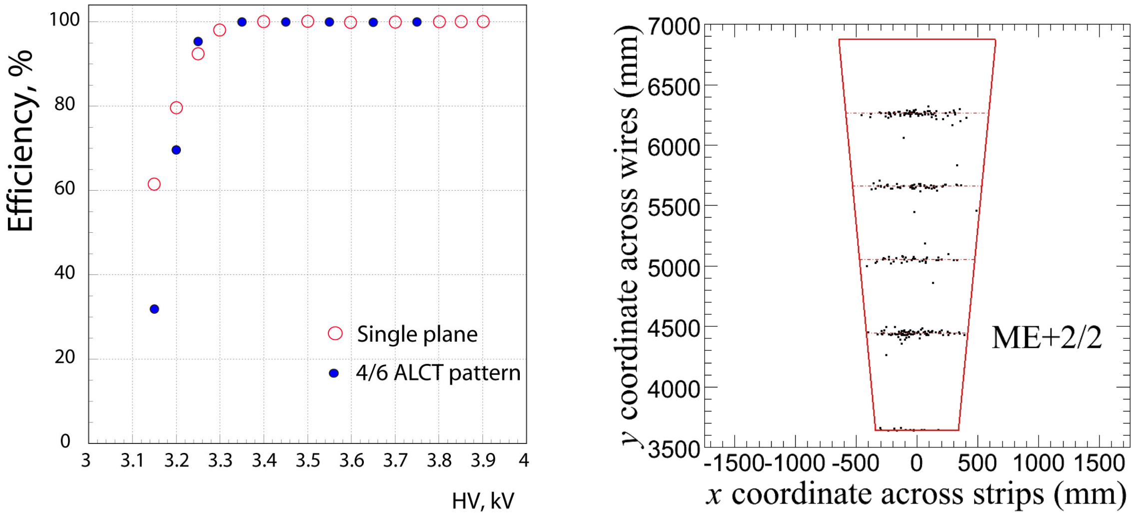

Figure 7.59: Left panel: single-plane anode signal efficiency (open circles) and ALCT pattern finding efficiency (filled circles) vs. high voltage. Right panel: predicted position of muon tracks in ME2/2 chambers for events when no LCT was found in these chambers (superimposed results for 6 chambers). The dashed lines indicate where wire planes of the ME2/2 chambers have inefficient bands separating independent HV segments.

using the ME1 and ME3 LCT stubs) were found to cross ME2/2 chambers in inefficient bands separating the chamber high-voltage segments (figure 7.59 (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 figure 7.60 (left). The observed spread of ≈0.5 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 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 7.60 (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 (figure 7.61 (left, top)) has an RMS of about 11 ns. Clearly this is too wide for a chamber hit to be assigned unambiguously to

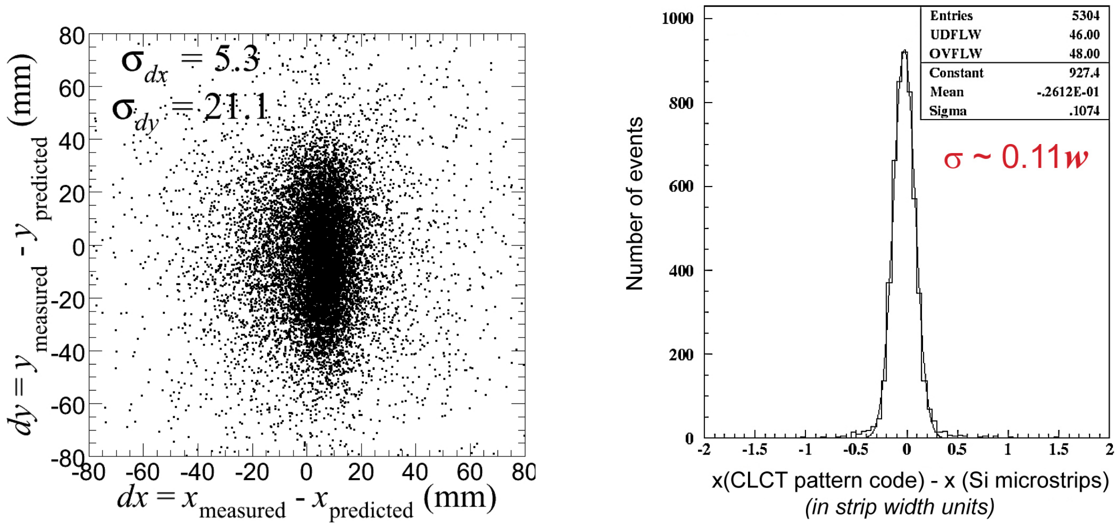

Figure 7.60: Left panel: 2D-coordinate residuals between LCTs found in ME2/2 chambers and muon-track positions. The *x* axis runs along the wires and the *y* axis along the strips. The observed spread is consistent with multiple scattering of cosmic-ray muons in the steel disks and the expected CLCT 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.

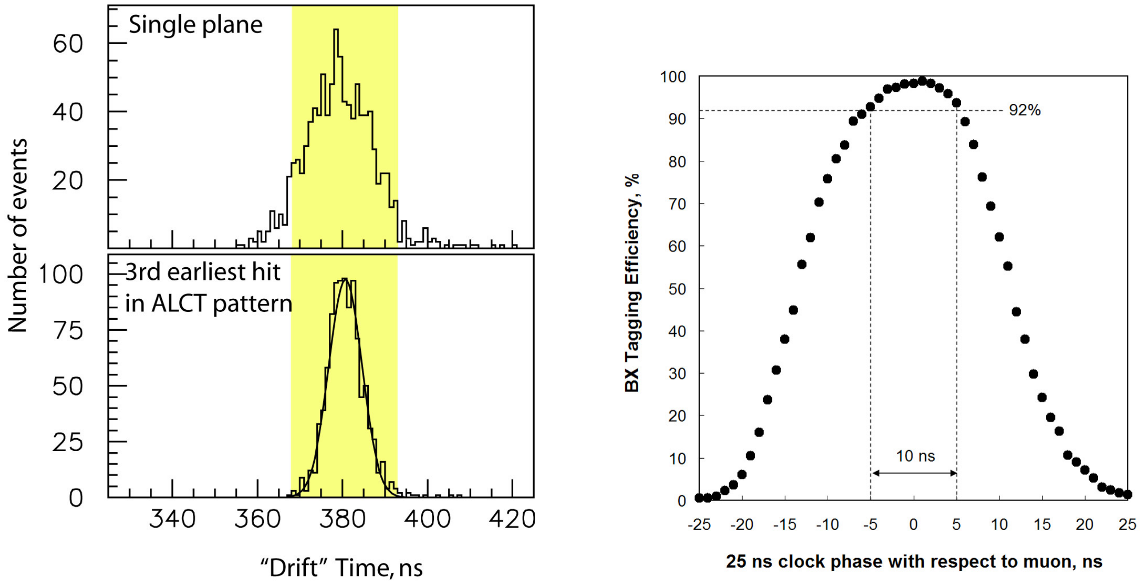

Figure 7.61: Left panel: time distributions of the response of a single plane to a passing muon (top) and for the 3rd earliest hit in an ALCT pattern (bottom). The horizontal scale has an arbitrary offset. The shaded band indicates the 25-ns window, the time between bunch crossings at the LHC. Right 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.

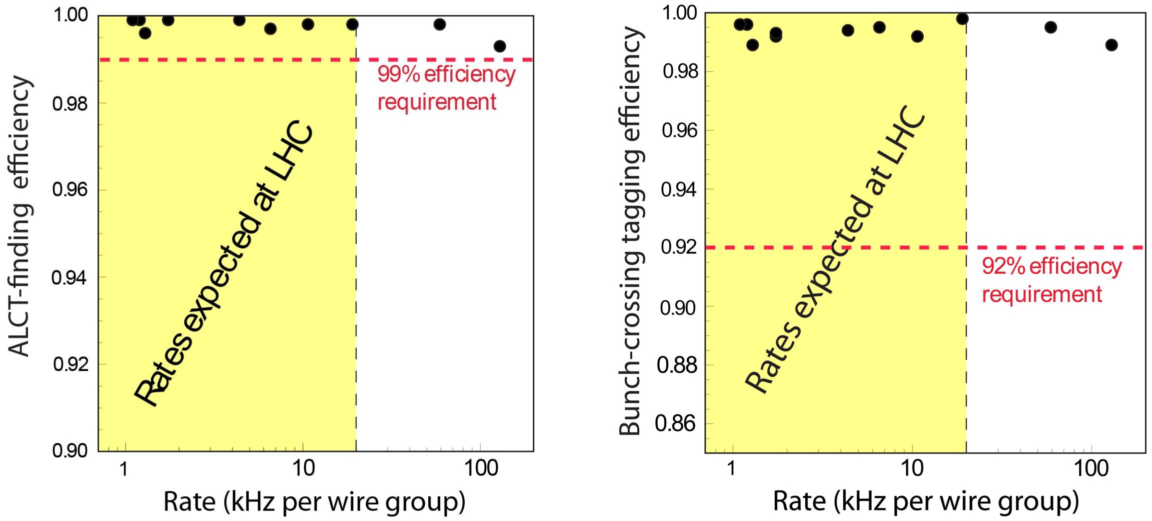

Figure 7.62: Left panel: ALCT-finding efficiency vs. rate of random hits per wire group. Right panel: efficiency of correct bunch tagging vs. rate of random hits per wire group. The shaded areas show the range of rates expected in different chambers at full LHC luminosity.

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 (figure 7.61 (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 7.61 (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.

Results obtained from a CSC irradiated with 0.7 MeV photons in a muon beam at the GIF (figure 7.62) 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 7.63 shows an example of a badly “contaminated” muon where the muon track is nevertheless successfully identified.

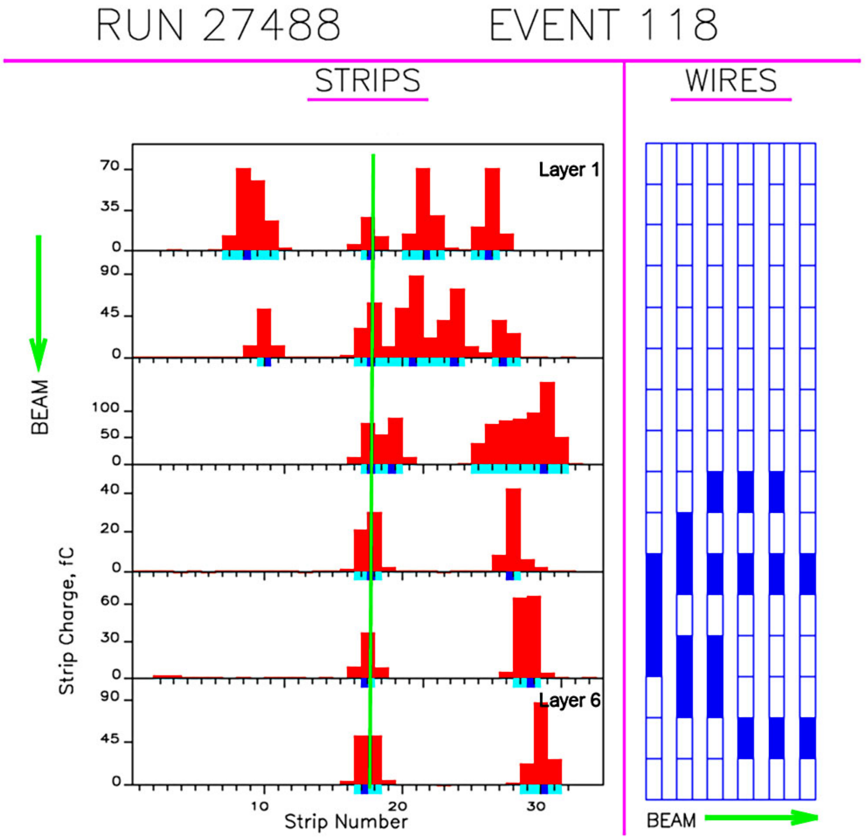

Figure 7.63: A sample CSC event of a muon accompanied by substantial bremsstrahlung radiation. The left side of the plot shows the charge (blocks just above the axes) on each of the 32 strips in each of the 6 chamber layers, while the right side shows the information from the anode wire groups. Strips 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 $\lambda = 17.59$.

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 [155]:

$$\frac{1}{Q}\frac{dQ}{d\lambda} = K_1 \left[\frac{1-\tanh^2(K_2\lambda)}{1+K_3\tanh^2(K_2\lambda)} \right], \tag{7.2}$$

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

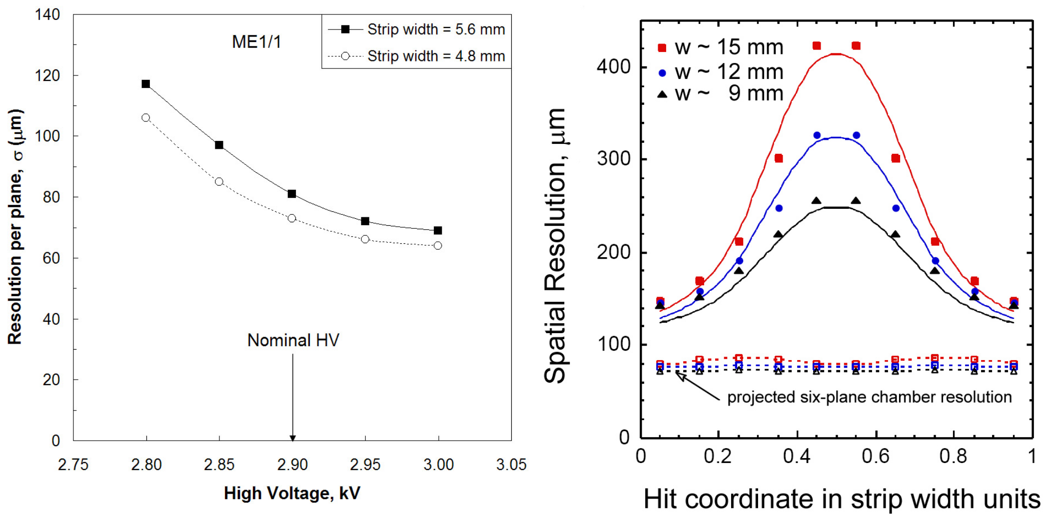

Figure 7.64: Left panel: ME1/1 chamber single-plane resolution vs. HV. Right panel: ultimate large CSC offline resolutions per plane for different muon passage points across a strip for areas with different strip widths for data (closed symbols) and simulation (solid lines). The expected overall 6-plane CSC resolutions are shown by open symbols and dashed lines.

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 (figure 7.64 (left)) is about 80 μm at nominal HV. The 6-plane chamber resolution is estimated to be $\approx 80/\sqrt{6} = 33$ μm, 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 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 7.64 (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

$$\frac{1}{\sigma^2_{\mathrm{CSC}}(x/w)} = \sum \frac{1}{\sigma^2_i(x/w)} \,. \tag{7.3}$$

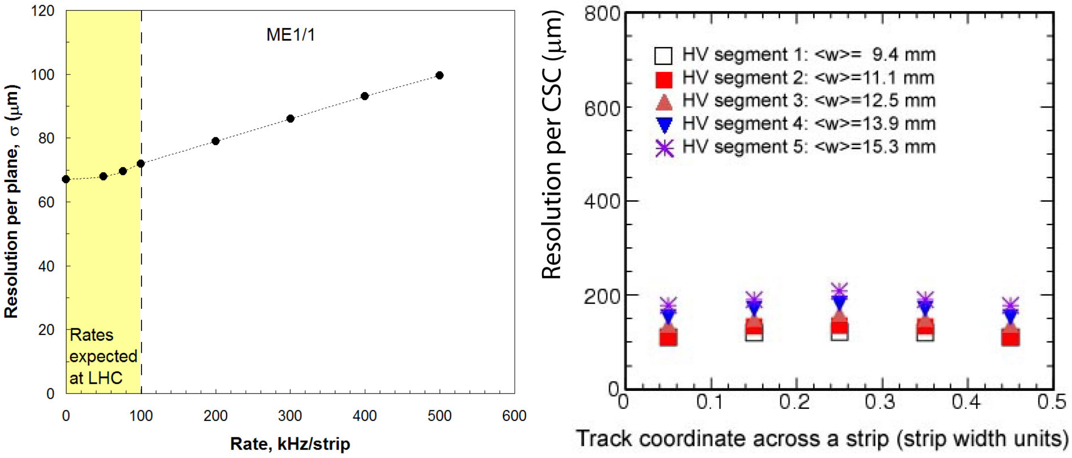

Figure 7.65: Left panel: deterioration of spatial resolution (ME1/1 chambers) with increasing rate of signals. The resolution remains well within the design specs even at rates far exceeding those expected at the LHC. Right panel: expected 6-plane chamber resolution for nearly perpendicular tracks vs. muon coordinate across a strip as evaluated from the single-plane resolution obtained with a simplified and fast reconstruction specifically targeted for the HLT.

The expected combined resolution for a 6-plane chamber is ≈80 μm almost independent of the hit position in a chamber, better than the 150-μm goal.

Even at the highest rates expected at the LHC, the CSC resolution will stay well within the design specifications (figure 7.65 (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):

$$(x/w)_{\text{measured}} = f_w(r), \quad \text{where } r = \frac{1}{2}\left[\frac{Q_{\text{right}} - Q_{\text{left}}}{Q_{\text{centre}} - min(Q_{\text{right}}, Q_{\text{left}})}\right]. \tag{7.4}$$

This algorithm localises muon stubs in a chamber with a precision of <200 μm (figure 7.65 (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.

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 [156] 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 [157]).

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 $\approx$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 [158]. 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 $\approx 3\times10^{11}$ 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 $\approx$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 $\approx$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-

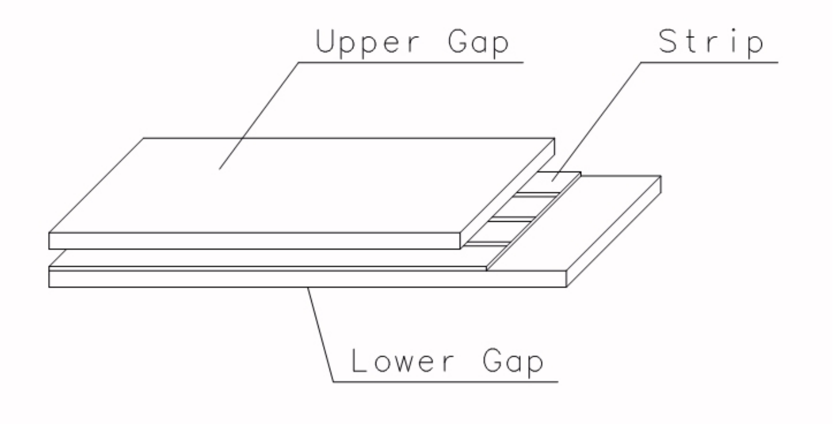

Figure 7.66: Layout of a double-gap RPC.

Table 7.2: Basic construction parameters.

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

track segments (input to the Level-1 trigger) and the spatial resolution attainable at the high-level trigger and offline was $\approx$150 μm.

7.3 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 [159, 160]. 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 (figure 7.66) [161, 162]. 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 7.2 lists the basic construction and operating parameters of the CMS double-gap RPCs.

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 [163, 164]. 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 are 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 allows 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 η = 2.1. However, in the first phase, due to budget limitations, only 3 layers up to η = 1.6 are 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.

7.3.1 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 (figure 7.67).

In the first and second muon stations there are 2 arrays of RPC chambers located internally and externally with 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 7.3), 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 7.4.

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 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° in ϕ). Figures 7.68 and 7.69 show schematic views of chamber modules with 2 and 3 double-gaps, respectively. Table 7.5 lists some global information regarding the barrel detector.

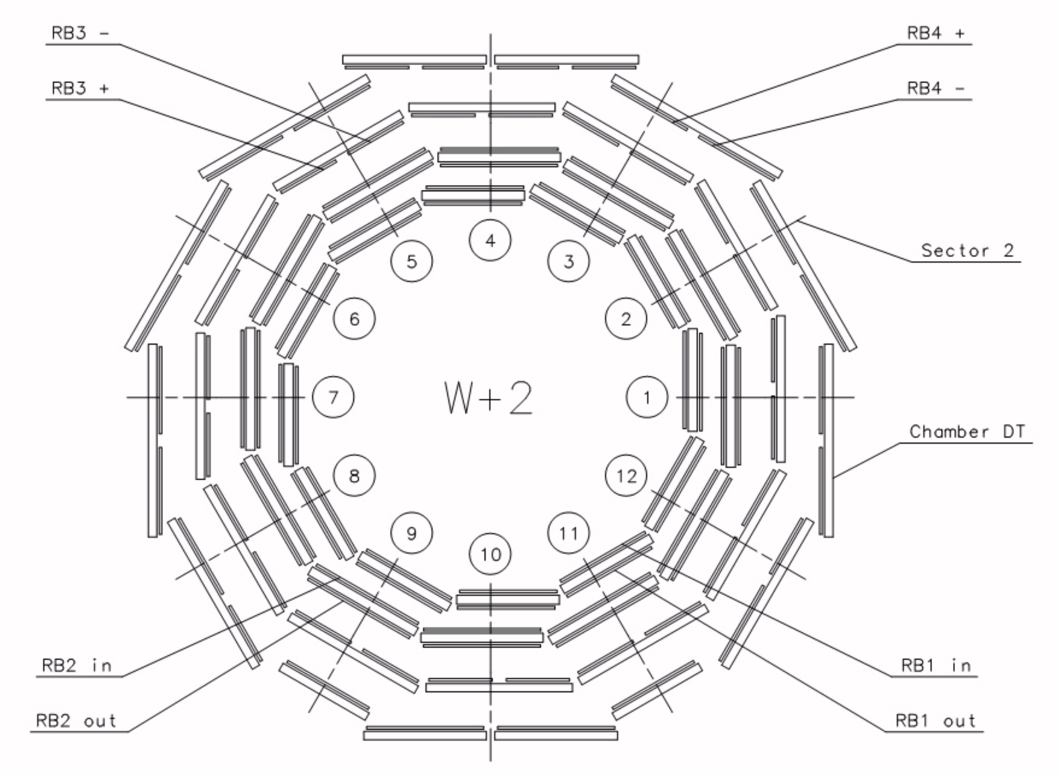

Figure 7.67: Schematic layout of one of the 5 barrel wheels, which are labeled –2, –1, 0, +1, and +2, respectively. Each wheel is divided into 12 sectors that are numbered as shown.

Table 7.3: 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

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 identifi-

Table 7.4: 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	

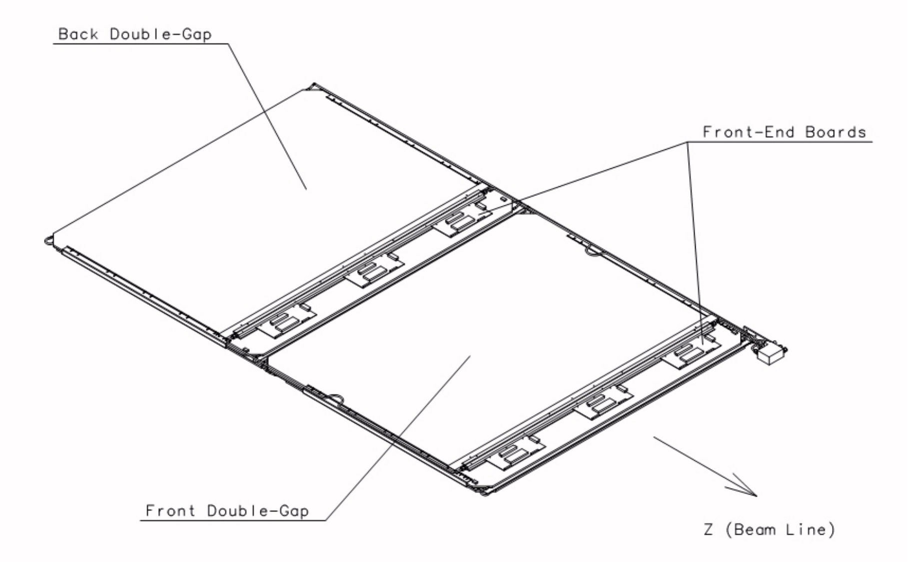

Figure 7.68: Schematic layout of chamber module with 2 double-gaps.

Table 7.5: 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

cation: 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 figure 7.70. They overlap in ϕ as to avoid dead space at chamber edges. Except

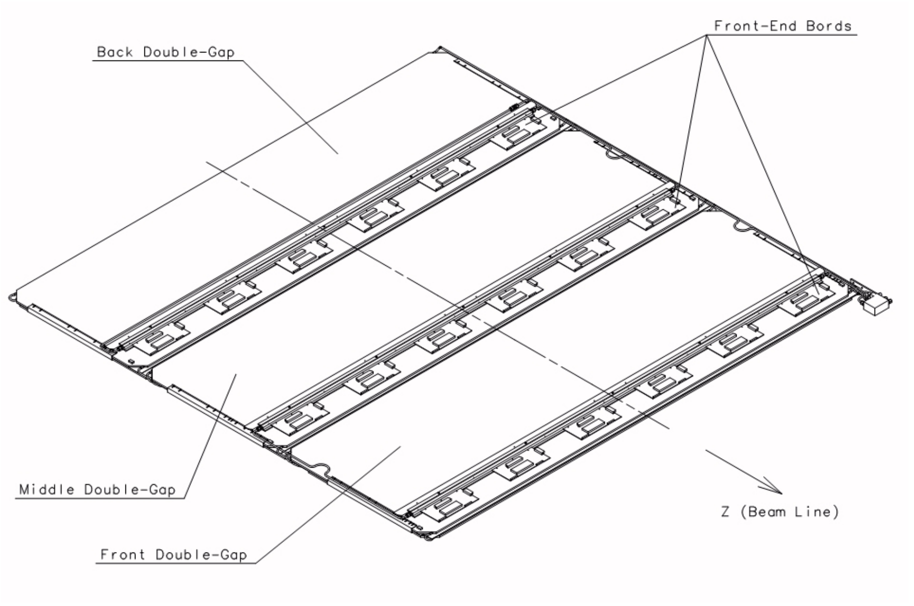

Figure 7.69: Schematic layout of chamber module with 3 double-gaps.

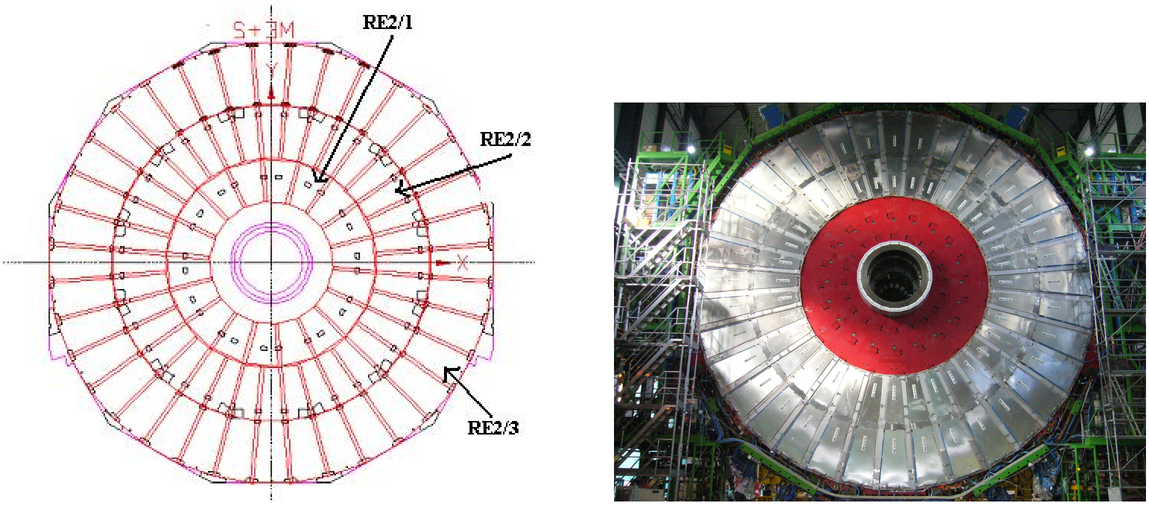

Figure 7.70: Left panel: schematic r-ϕ layout of RPC station RE2 on the back side of the first endcap yoke. Right panel: RPC station RE2 on the back side of the YE-1 yoke. The inner ring has been staged and is absent here.

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} cm^{-2}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, respectively. They remain uncovered since the corresponding CSC stations 2 and 3 are mounted on both faces of YE2. Figure 7.71 shows a schematic layout

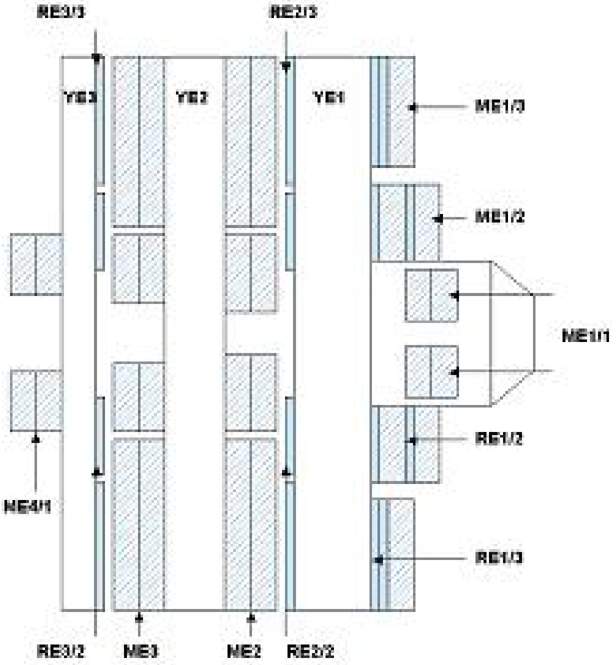

Figure 7.71: Schematic layout of the CMS endcap for the initial muon system.

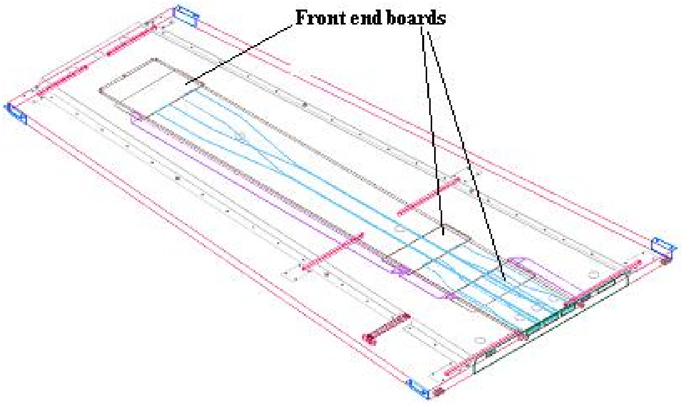

Figure 7.72: A view of an endcap RPC chamber.

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 16×16 mm^2 section spacer frame (figure 7.72). 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.

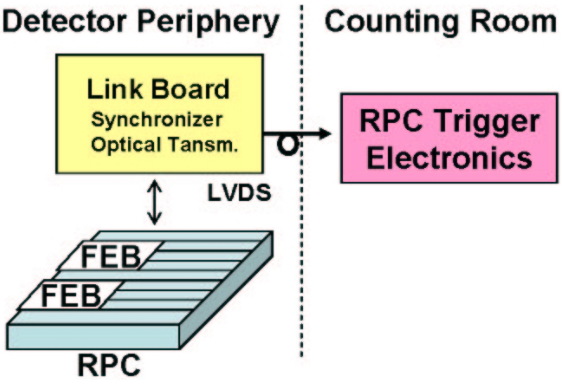

Figure 7.73: Block diagram of RPC read-out electronics.

7.3.2 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 unsynchronized to Link Boards (LB) placed around the detector. The LBs synchronize 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 figure 7.73.

The FEBs house two (barrel version) or four (endcap version) front-end chips, which are custom ASICs designed in AMS 0.8 μm CMOS technology [165]. 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-Ω 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 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 [166]. 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 [167].

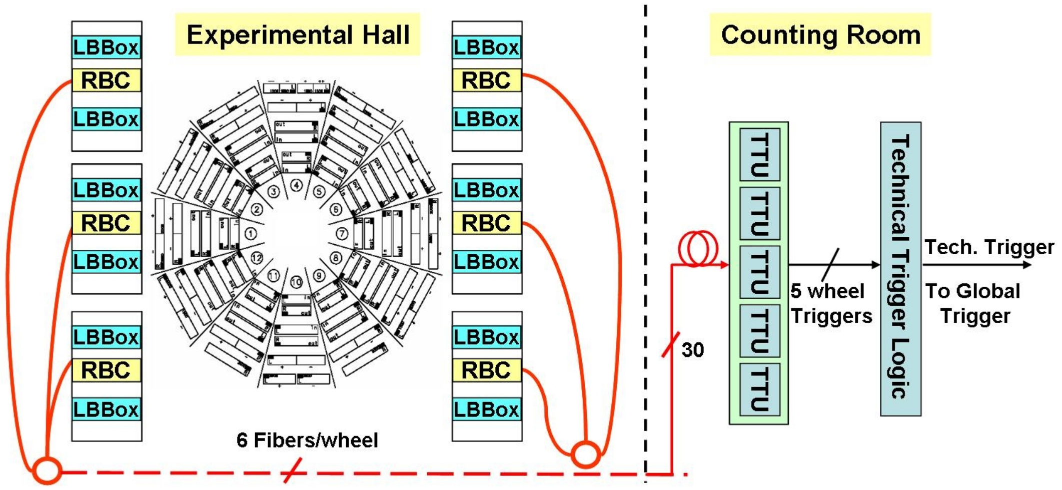

Figure 7.74: Technical Trigger schematic layout.

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 (figure 7.74) [168].

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.

7.3.3 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 are 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/cm^2 and $5{\cdot}10^{11}$ neutrons/cm^2. 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 7.6.

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

Table 7.6: 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 μA	100 mA
Voltage monit. precision	<10 V	100 mV

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 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 ETHERNET3 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 μA. 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 μA 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 7.7.

Table 7.7: Summary of HV systems.

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

Table 7.8: Summary of LV systems.

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

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 ($\mathrm{LV_a}$) and digital ($\mathrm{LV_d}$) 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 $\mathrm{LV_a}$ and 1 $\mathrm{LV_d}$ channel per chamber. However, for the endcap detector, at the start-up 1 $\mathrm{LV_a}$ and 1 $\mathrm{LV_d}$ are distributed between 2 chambers. A summary of the LV systems is given in table 7.8.

7.3.4 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 are read out through the LB electronics to monitor the temperature.

7.3.5 Gas system

Test results [169] showed that RPCs are suitably operated with a 3-component non-flammable mixture of 96.2% R134a ($C_2H_2F_4$), 3.5% iC_4H_{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

Table 7.9: Main gas parameters of the CMS RPC system.

Gas volume	14 m^3
Gas mixture composition	96.2% R134a, 3.5% iC_4H_{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)

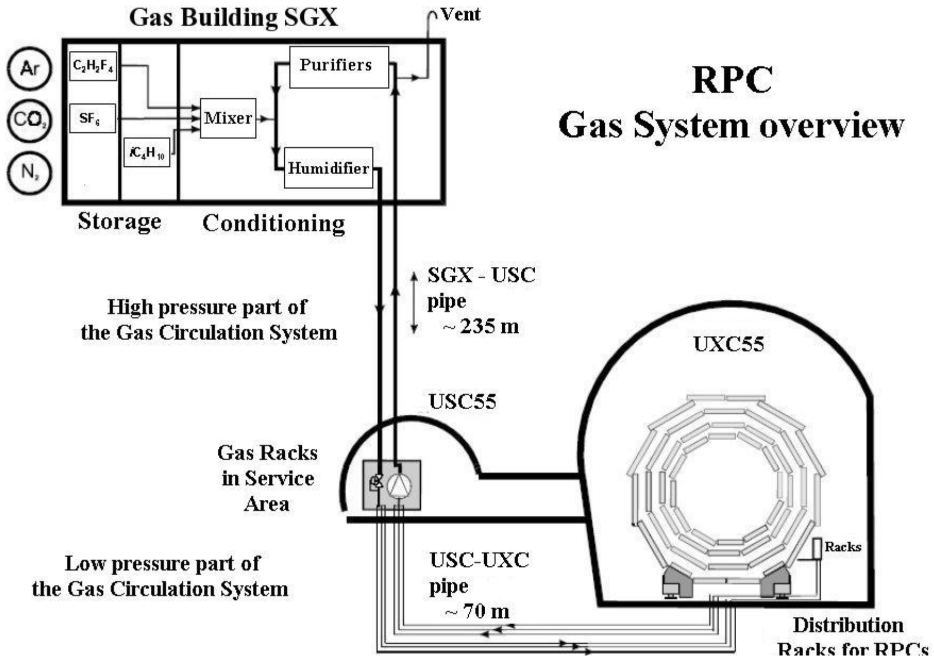

Figure 7.75: Closed-loop circulation system.

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 7.9.

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 [132]. The full closed-loop circulation system (figure 7.75) 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,

Table 7.10: 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

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 iC_4H_{10} fraction increases beyond the flammability limit. For fast detector filling, parallel rotameters are 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 (figure 7.76). 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 7.10).

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 REn: REn/2 and REn/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 (figure 7.77a). 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 (figure 7.77b). In RE2 and RE3

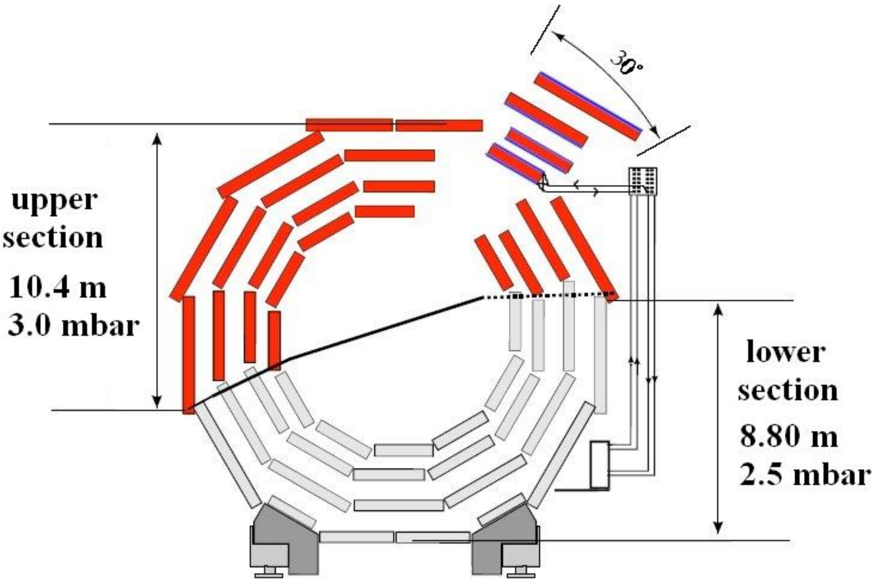

Figure 7.76: The 2 zones into which a wheel is sub-divided. Each station (2 chambers) is supplied by a gas line.

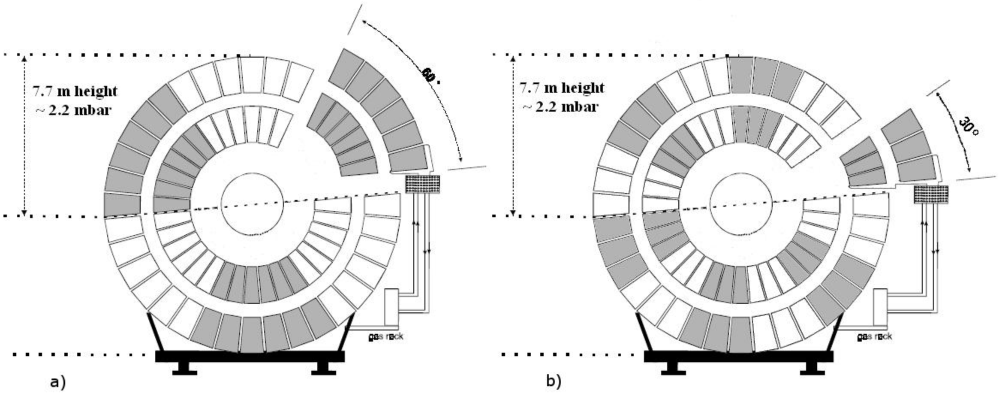

Figure 7.77: In each endcap disk the RPC detectors are divided in 2 rings. RE1 (a) is divided into 60° sectors, while in RE2 and in RE3 (b) sectors are composed of 3 chambers of the internal ring and the corresponding 3 of the external ring. Every sector is supplied by 2 independent gas lines.

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 7.11.

Table 7.11: 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 channels	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

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.

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 (50×50 cm^2) single-gap RPCs whose working points (gain and efficiency) are 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] performs 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 allow the analysis of the gas mixtures that return from every half wheel. In the surface gas building (SGX), sampling points are available to monitor the effectiveness as well as the status of each cartridge in the purifier module.

7.3.6 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 ($T_0 = 293$ K and $P_0 = 1010$ mbar).

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. Rigourous 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 figure 7.78 the distribution of the maximum efficiency for all the barrel RPCs is shown. The mean value of the distribution is 97.2%. In figure 7.79 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

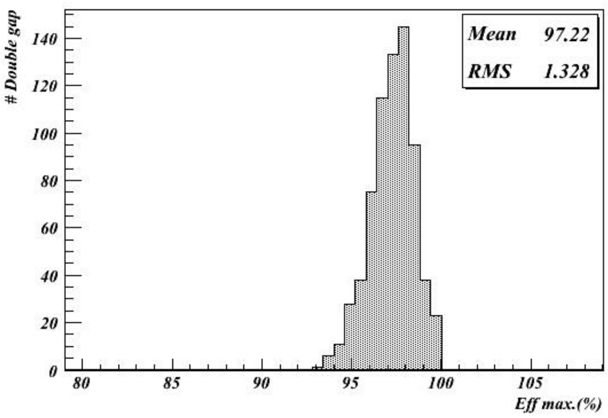

Figure 7.78: Distribution of plateau efficiency for all the barrel chambers.

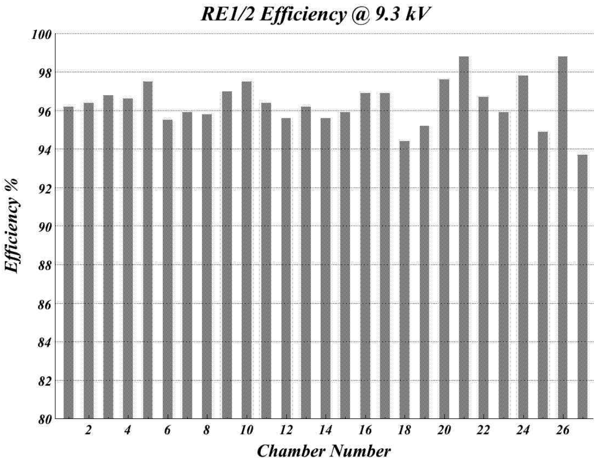

Figure 7.79: Efficiency at HV = 9.3 kV for the first 27 endcap chambers.

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 figure 7.80.

The chamber cluster size is defined as the average value of the cluster-size distribution. In Figure 7.81 the profile histogram of the cluster-size distribution as a function of the $\mathrm{HV_{eff}}$ 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.

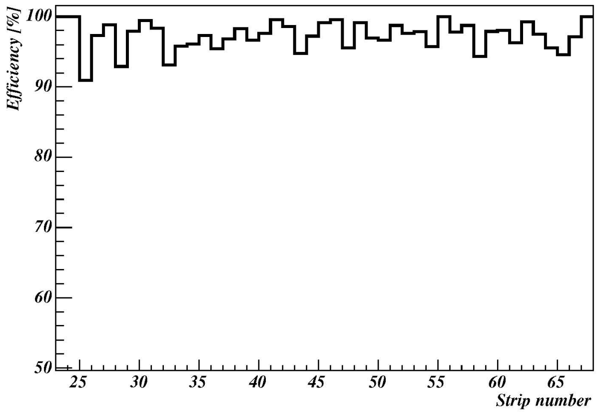

Figure 7.80: Local efficiency at HV_{eff} = 9.6 kV for a barrel chamber.

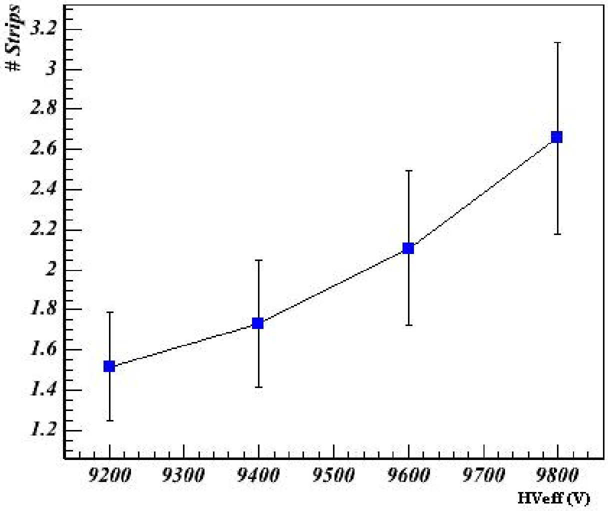

Figure 7.81: Profile histogram of the chambers' cluster-size distribution as a function of HV_{eff}. The dots and bars are the average and the root-mean-square of the cluster-size distributions, respectively.

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 synchronized with the other muon detector (DT and CSC) triggers.

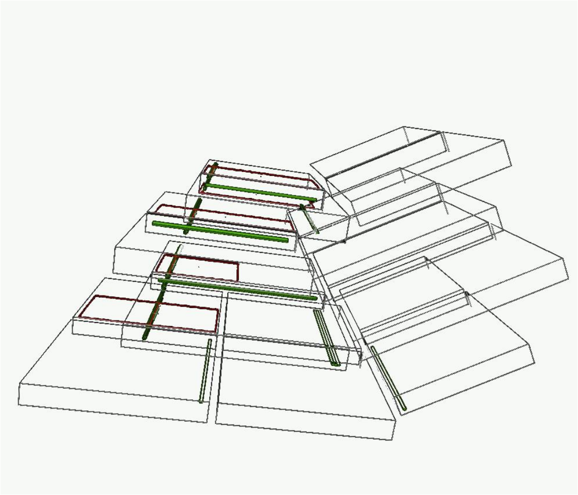

Figure 7.82: Iguana muon reconstruction: RPC-fired strips are in red and DT hits in green.

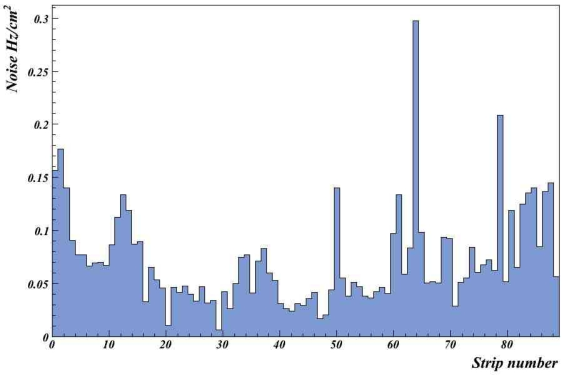

Figure 7.83: Barrel chamber noise profile at $\mathrm{HV}_{\mathrm{eff}}$ = 9.6 kV.

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 7.82 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 figure 7.83 at $\mathrm{HV}_{\mathrm{eff}}$ = 9.6 kV, while figure 7.84 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 figure 7.85 the chamber efficiency as a function of the $\mathrm{HV}_{\mathrm{eff}}$ is shown for some RPC stations.

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

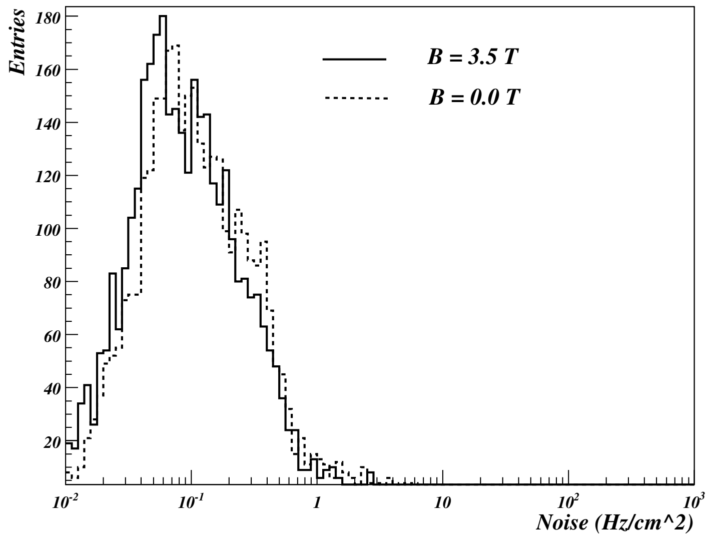

Figure 7.84: Noise distributions in 2 different magnetic fields. All strips of the barrel stations are included in the distribution.

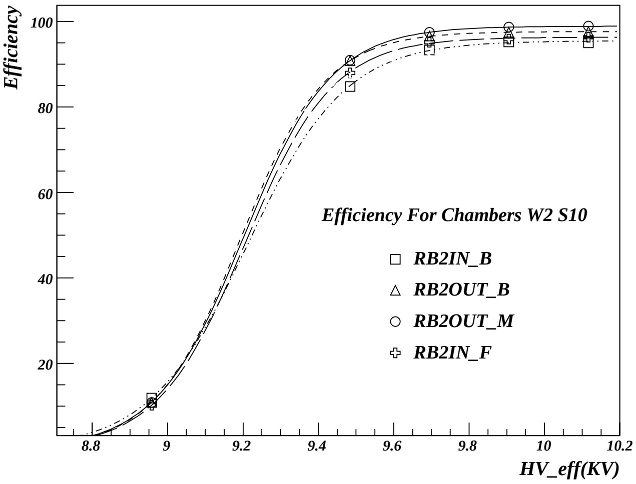

Figure 7.85: Global efficiency vs. HV_{eff}, estimated by means of DT segment extrapolation.

7.4 Optical alignment system

For optimal performance of the muon spectrometer [132] 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 (section 7.1 and 7.2). 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:

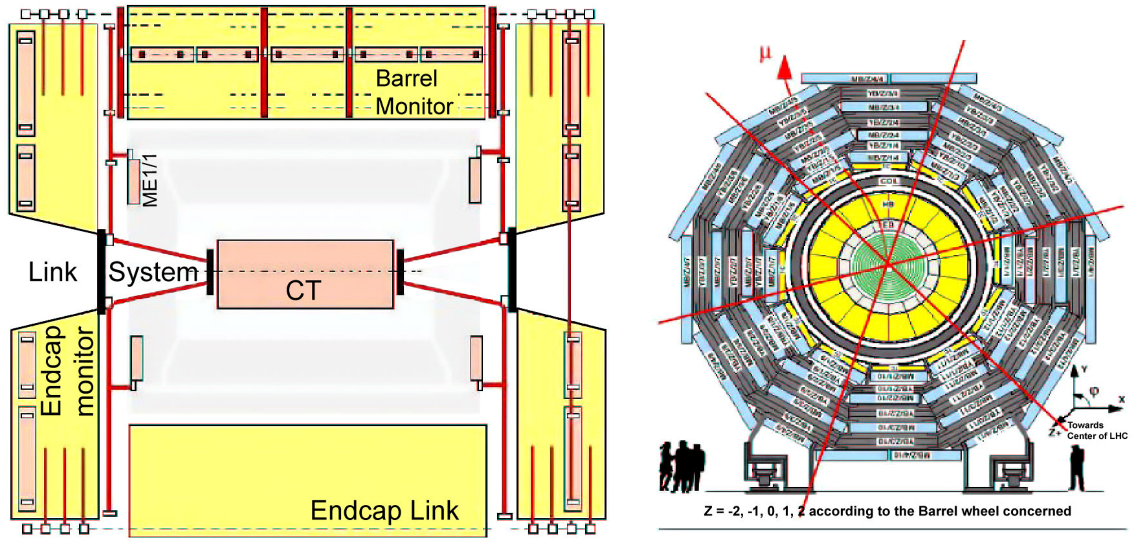

Figure 7.86: Schematic view of the alignment system. Left panel: longitudinal view of CMS. The continuous and dotted lines show different optical light paths. Right panel: transverse view of the barrel muon detector. The crossing lines indicate the r-z alignment planes with 60° staggering in ϕ.

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 7.86 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.

7.4.1 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 $\approx$ 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 (figure 7.87) 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 μm and 50 μrad [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 50×65 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 μm. 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 μm precision, as described in section 7.1.4. 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 μm. Long term measurements showed good stability of the centroids and light intensity distributions. CMOS miniature video sensors, containing 384×288 pixels with 12×12 μm^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 μm and 50 μrad.

Once MABs were installed (figure 7.88), the initial MAB positions on the barrel wheels were determined by photogrammetry measurements.

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 synchronizes the operation of the light sources mounted on

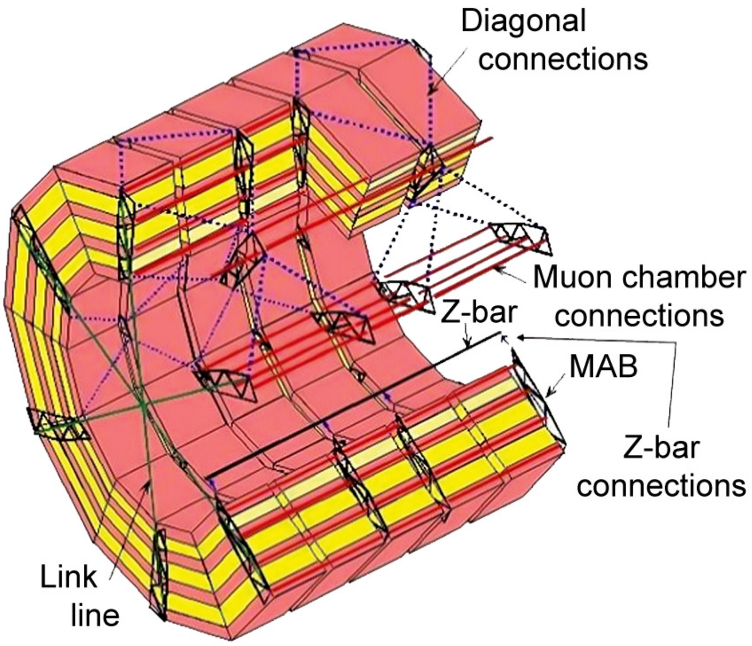

Figure 7.87: Schematic view of the barrel monitoring system showing the optical network among the MAB structures.

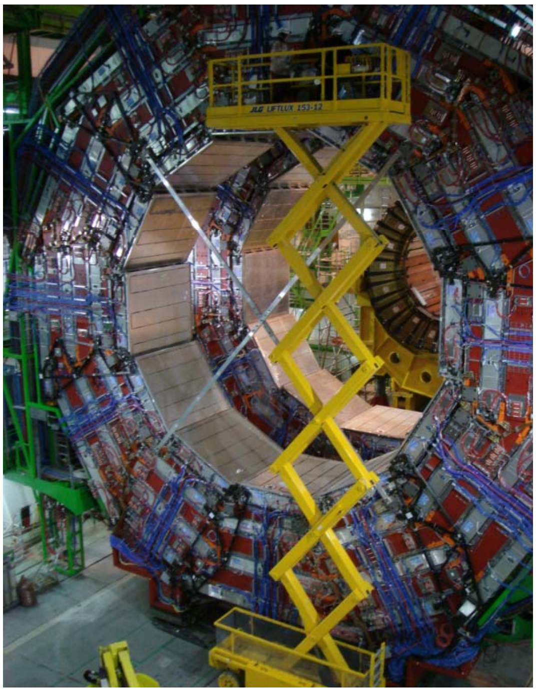

Figure 7.88: Installation of the MABs on wheel YB+2.

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 μm for chambers in the same sector and about 250 μm between barrel sectors. The current understanding of its performance is discussed in section 7.4.3.

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 μm in r-ϕ. Because the r and r-ϕ accuracies are directly coupled, the required accuracy in the r-position is $\approx$400 μm. 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 cm with an accuracy of $\approx$1 mm.

The system uses a complex arrangement of 5 types of sensors for the transferring and monitoring of ϕ, r, and z coordinates (figure 7.89). 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 7.90 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 figure 7.90 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

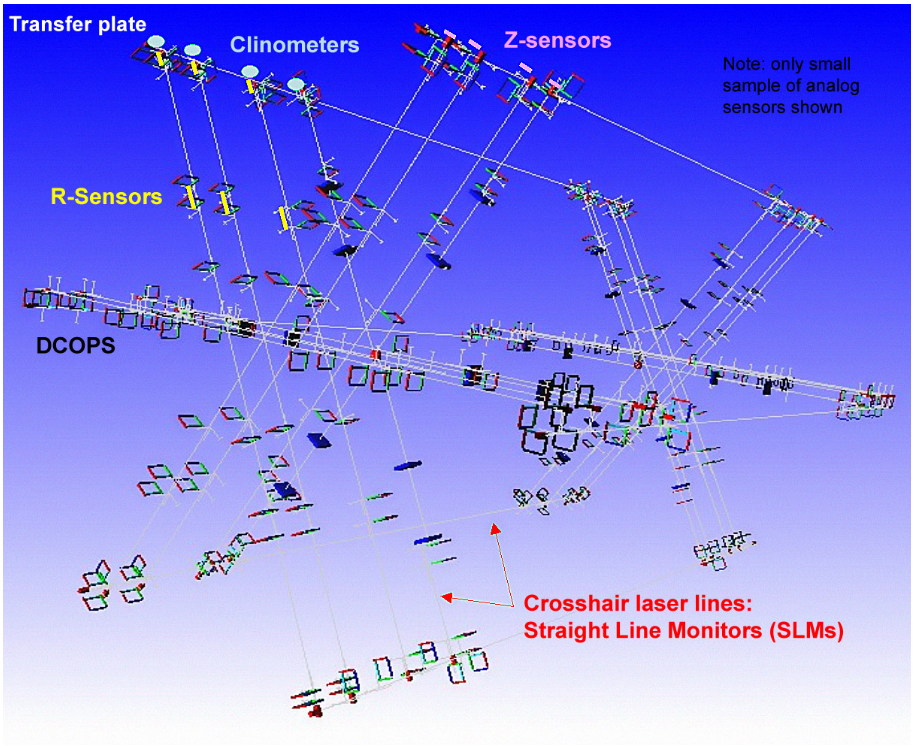

Figure 7.89: Visualisation of the geometry and components of the muon endcap alignment system. The square objects represent optical sensors (DCOPS) for monitoring 3 straight laser lines across each endcap station.

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.

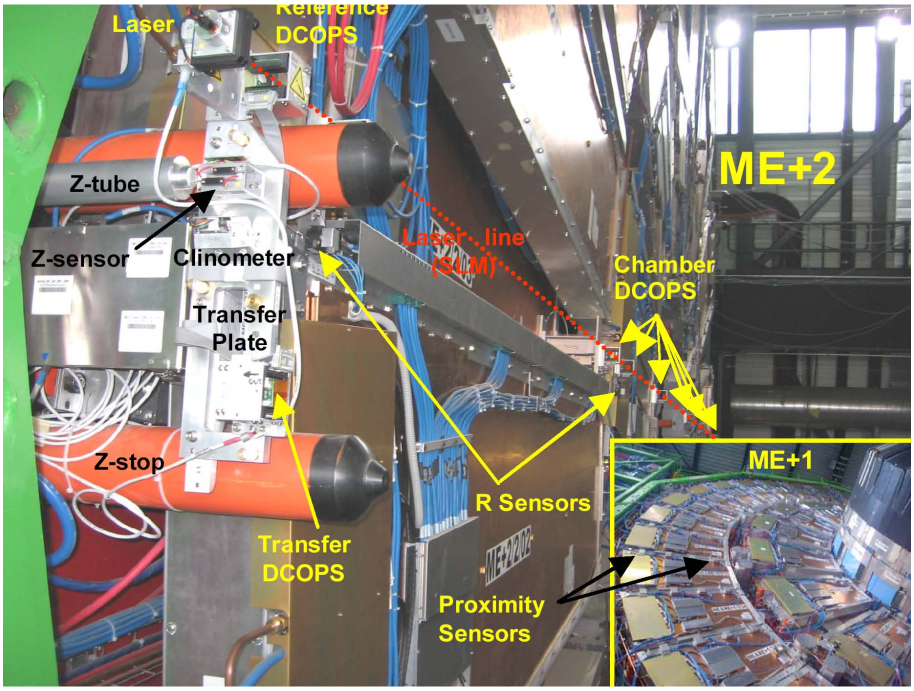

Figure 7.90: Close-up of one of the 3 Straight Line Monitors (SLM) on the ME+2 station with cross-hair laser, DCOPS, and analog sensors (r, z, and Tiltmeter). The insert indicates the location of proximity sensors on ME+1.

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 (figure 7.86) 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±1 wheels of the endcap muon spectrometer (link disks, LD); and the MAB structures attached at the external barrel wheels, YB±2. Figure 7.91 (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.

The ARs are rigidly attached to the endcap tracker detectors, TECs, through a purely mechanical connection with the instrumented silicon volume (section 3.3.7). Three pillars, acting as support fixations, connect the last instrumented disk of each TEC with the corresponding AR, at

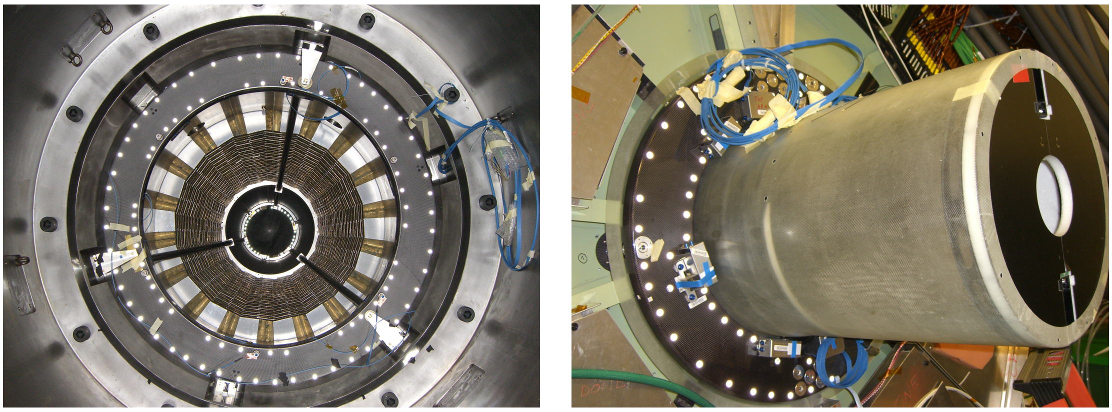

Figure 7.91: Left panel: Link Disk and Alignment Ring installed in the inner $\eta = 3$ cone during the first closing of the detector in summer 2006, MTCC period. Right panel: Alignment Ring mounted in the TEC end flange.

both ends of the tracker volume, see figure 7.91 (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 μrad. Long term measurements were performed after beam adjustments. Beam pointing stability, including temperature effects, was found to be better than 30 μrad. The adjustment and calibration accuracy was limited to 30–100 μrad due to the finite dimension of the structures combined with the intrinsic accuracy of the survey and photogrammetry measurement techniques of 70 μm.

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.

7.4.2 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.

7.4.3 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 (section 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 μm 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 figure 7.92. 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 7.92 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 figure 7.92, the nose is pulled towards the interaction point, the magnitude of the compression is perfectly correlated with the magnet current, reaching up to $\approx$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 $\approx$600 μm, while expanding azimuthally by $\approx$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 (figure 7.93). 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 figure 7.93. Note that the

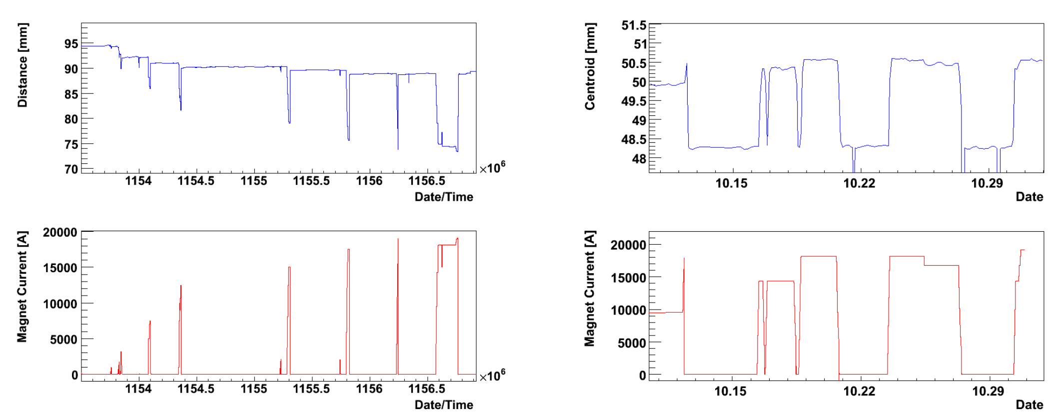

Figure 7.92: Deformations of endcap disks and barrel wheels vs magnet current cycling. Left panel: The bottom plot shows the magnet powering cycle exercised during the first phase of the Magnet Test period. The top plot shows the measured YE+1 nose compression towards the interaction point. Right panel: 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.

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 figure 7.94. 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.

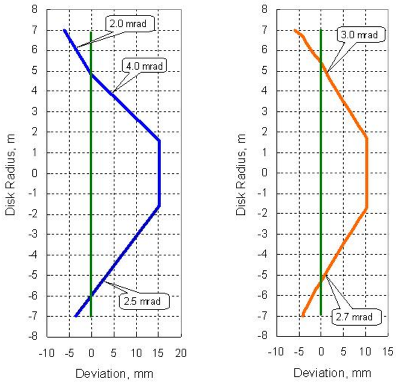

Figure 7.93: Comparison of the YE+1 disk deformations in the r-z plane at full magnetic field (4 T) measured by the alignment system (left panel) and predictions from finite element analysis (right panel). The vertical lines correspond to 0 magnetic field.

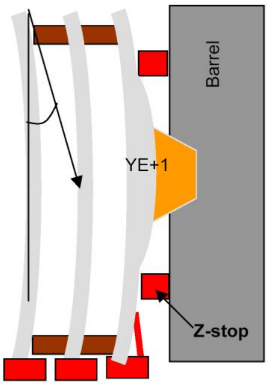

Figure 7.94: Current understanding of disk deformation due to magnetic forces based on alignment system measurements. The z-stops (red) prevent the disks from getting pushed into each other. Note that the indicated bending angle is exaggerated for illustrative purposes. Its measured magnitude is 2.5 mrad.