

Operation and Performance of the CMS Silicon Strip Tracker

Gero Flucke* †

Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

E-mail: gero.flucke@desy.de

The CMS silicon strip tracker is the largest silicon detector ever built with almost 10 million read-out channels and an active area of close to 200 m². This article describes the status of the detector after operation for two and a half years of LHC running. Results from radiation measurements and detector performance, including alignment, are shown.

The 21st International Workshop on Vertex Detectors

16-21 September 2012

Jeju, Korea

*Speaker.

†On behalf of the CMS Collaboration




Figure 1: Schematic view in the rz plane of a quarter of the CMS silicon tracker. Solid magenta lines indicate the position of single strip modules and hollow blue areas the regions instrumented with $r\phi$ and stereo strip modules mounted back-to-back.

1. Introduction

The CMS tracker [1, 2] is the central tracking device of the CMS experiment [2] at CERN's Large Hadron Collider, situated inside a solenoidal magnet that provides a magnetic field of $B = 3.8$ T. The tracker is based on silicon sensor technology using pixels in its inner part and microstrips in the outer part. Operational issues and performance of this outer part, called silicon strip tracker, are the topic of this article.

Full operation of the CMS silicon strip tracker started in 2008 recording muon tracks from cosmic rays. Since the end of 2009 the LHC provides proton-proton and lead-lead collisions with increasing centre-of-mass energy \sqrt{s} and peak luminosity L . In summer 2012 these values reached $\sqrt{s} = 8$ TeV and more than $L = 7 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in proton-proton operation, respectively.

The z -axis of the coordinate system of the CMS experiment is along the direction of the anticlockwise LHC beam. The polar angle θ is measured from the positive z -axis and the azimuthal angle ϕ is measured in the plane perpendicular to the z -axis, whereas the radius r denotes the distance from it. The pseudorapidity is defined as $\eta = -\log[\tan(\theta/2)]$.

2. The CMS Silicon Strip Tracker

The fundamental building blocks of the CMS silicon strip detector are its 15 148 modules. The layout of the detector can be seen in figure 1. The modules are arranged in four sub-detectors. In the Tracker Inner and Outer Barrels (TIB and TOB), modules at similar radii are placed in four and six layers around the z -axis, respectively. In the Tracker Inner Disks (TID) and the Tracker End-Caps (TEC) in the forward region, modules are mounted orthogonal to the z -axis on three and nine disks, respectively, and are arranged in rings with increasing radii. The microstrips are generally oriented such that the measurement is in the $r\phi$ direction. At some radii, there are additional modules where the strips are rotated by 100 mrad to provide position information in the rz plane. These are called *stereo* modules, in contrast to the normal $r\phi$ modules.

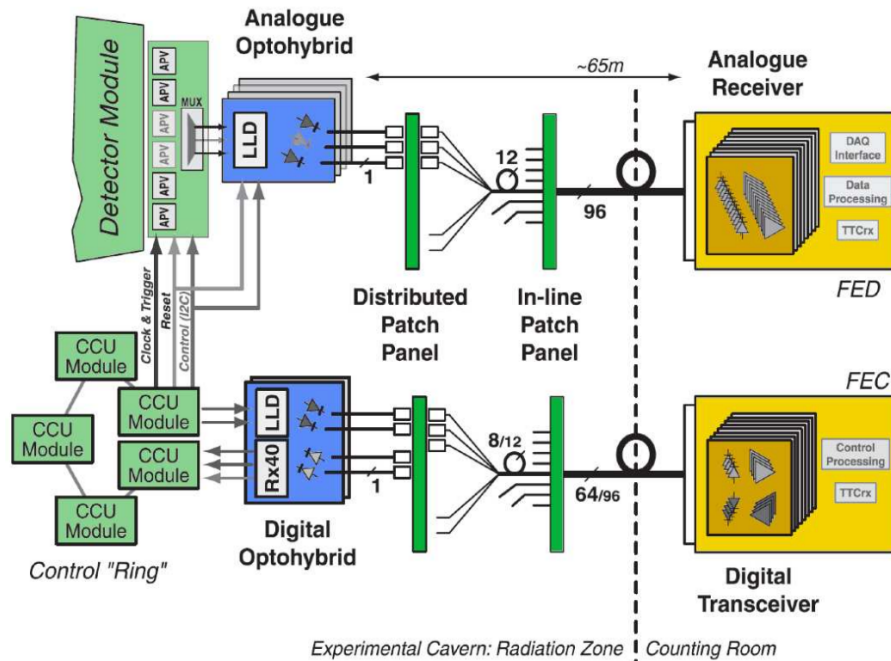


Figure 2: Sketch of the readout and control chains of the CMS strip detector.

Each module contains either one “thin” sensor of 320 μm thickness (TIB, TID, inner TEC rings) or two “thick” sensors of 500 μm in a daisy-chained readout (TOB, outer TEC rings). Their 512 or 768 microstrips are p^+ implants in an n -type silicon bulk. The strip pitch varies from 80 to 205 μm , while the ratio of the implant width and the pitch is 0.25 everywhere.

Readout and Control Chains The readout and control chains are shown in figure 2. The charge of 128 neighbored strips is read out via AC coupling by one APV front-end chip. The analogue values are stored every 25 ns in a pipeline of 192 cells. Two readout modes exist: in *peak* mode the signal shape behaves like a 50 ns CR-RC circuit. The broad signal size makes this readout mode rather robust against time mis-alignment. This is particularly important when taking cosmic ray data that are asynchronous with the 40 MHz clock of the data acquisition system (DAQ). In *deconvolution* mode the signal stored in the pipeline is a weighted mean of three consecutive bare signals. The significantly reduced signal width allows the identification of the correct LHC bunches that are nominally separated by 25 ns. The cost is a slight increase of noise. The APVs send their data if requested by an external trigger. The data from two APVs are multiplexed before the ADC counts are converted to light by an Analog Opto Hybrid (AOH) and sent via optical fibres to the Front End Driver (FED). A total of 440 FEDs receive non-zero suppressed data via the fibres and perform the digitisation, common mode noise subtraction and zero suppression, using pedestal and noise values recorded in dedicated calibration runs in beam-off periods.

Front End Controller (FEC) boards manage the clock, trigger signals and slow control communication. Information is sent via optical fibres to Digital Opto Hybrids (DOH) which serve as entry point to so-called *control rings*. These consist of a number of Communication and Control Units (CCU) arranged in a token ring structure. Each CCU programs a number of individual detector modules connected to it via an I²C bus and in addition provides clock and trigger signals. Among




Figure 3: Map of the silicon sensor temperatures in °C as measured by the DCUs during operation in May 2011. The modules are grouped in the 3 (9) disks of both TIDs (TECs) and the 4 (6) layers of the TIB (TOB). Regions connected to closed cooling loops are indicated by red ellipses.

the programmed devices on each module are the Detector Control Units (DCU) which provide information about module temperatures, leakage currents and low voltages. During LHC operation, every 300 s the module status is sent via CCU and DOH back to the FEC, providing fine grained knowledge of the detector status.

Services While the FEDs and FECs for the readout and control are placed in a service cavern, the 972 power supply (PS) units for the silicon modules are located inside the experimental cavern to reduce the power loss in the copper cables. This requires radiation hardness and B-field tolerance. Each PS module hosts two units, each providing two low voltage lines (1.25 V and 2.5 V) for the APVs, AOHs, etc. as well as two high voltage (HV) lines to supply up to 600 V of bias voltage to 3-12 silicon modules. The current operation point is 300 V. In addition there are 356 PS units that supply the control rings with low voltage. The rate of PS unit failures is well below 1% per year.

About 50 kW of heat have to be removed to keep the strip tracker at constant temperature during operation. Cooling power is provided by two independent cooling plants with 90 lines each. The cooling liquid is C_6F_{14} which is highly volatile and neutral to the electronics. While in 2010 the high leak rate was of big concern, the measures described in [3] successfully kept it below 1 kg per day since spring 2011. Five cooling lines had to be closed, leaving a small fraction of the modules with indirect cooling only. This can be seen in figure 3 that shows the sensor temperature of all strip modules as measured by their DCUs. Till the end of 2012, the cooling plants operate at a cooling liquid temperature of 4°C. The system will be refurbished during the long shutdown one (LS1) in 2013-14. After LS1, the operation temperature will be substantially lower to mitigate effects of irradiation that are not an issue so far.

3. Performance

The vast majority of the silicon strip modules are performing well. Only a small fraction of them was not active in August 2012. Reasons for their exclusion from the DAQ are shorts in control rings or HV lines, open HV lines, control rings with failing communication, and, to a smaller extend, broken readout fibres or failing CCUs. The distribution of non-active modules is illustrated in figure 4. The number of active modules decreased only slightly with time, from




Figure 4: Distribution of silicon strip modules that were active (green), masked (red) or never commissioned (white) in August 2012. Coloured ellipses indicate the main reason for groups of modules not being active: LV problems of control rings (black), communication problems in control rings (blue), and shorts in HV lines (red).

98.5% during cosmic ray data taking in 2008 via 97.75% in 2011 to now 97.61%. In addition to the online masking, an offline run-by-run analysis determines modules temporarily excluded from the DAQ, noisy and dead channels, etc. This algorithm usually flags 0.1–0.6% of the channels (i.e. strips) as bad. Both, online masked modules and offline flagged channels are available for the full reconstruction of the data that is delayed by 48 h.

To provide as much collision data as possible for physics analysis, the CMS detector and all its components should be operational whenever the LHC is in stable operation mode. Until August, 7% of the integrated luminosity delivered by the LHC in 2012 has been lost due to downtime. Despite its large number of FEDs, the strip tracker is only responsible for 16% of the downtime. More operational steps have been automated to decrease the data taking time without a fully functional strip tracker. First there is the raising of the HV when the LHC declares 'stable beams' and the beam conditions are clean as measured by the beam radiation monitor of CMS. This procedure at the start of an LHC fill now takes about 1.1 minutes compared to the 1.5–2 minutes needed when the operation was executed manually. Second, the newly introduced CMS *soft error recovery* procedure also reduces the downtime. For example, control rings that lost communication used to take 4–5 minutes to recover while now they take only 20–30 s.

The leakage currents through the silicon diodes increase with irradiation. Module-by-module measurements of these currents, again provided by the DCUs, can be used to cross check models of the integrated radiation doses and of the radiation damage induced by the irradiation. Figure 5 shows a good correlation between the DCU measurements at the end of 2011 and the prediction of simulations for 5 fb^{-1} integrated luminosity of pp collisions at $\sqrt{s} = 7 \text{ TeV}$ as accumulated until then. This match between data and measurements reassures the use of these models to predict the detector behaviour and its lifetime for the future.

While the peak readout mode of the APVs is rather insensitive to timing misalignment below the LHC bunch crossing distance of 25 ns, small timing offsets have an influence in deconvolution mode. Therefore the delay of each module can be varied in steps of 1.04 ns. This fine delay can be calibrated using only a few minutes of collision data taking. In the beginning of 2011 this procedure determined a time misalignment of 3 ns for most of the TOB modules and the delay has been adjusted accordingly. A recalibration in 2012 showed no need for new adjustments. This can be seen in figure 6: For most layers the optimal delay change is below the step size of 1.04 ns.

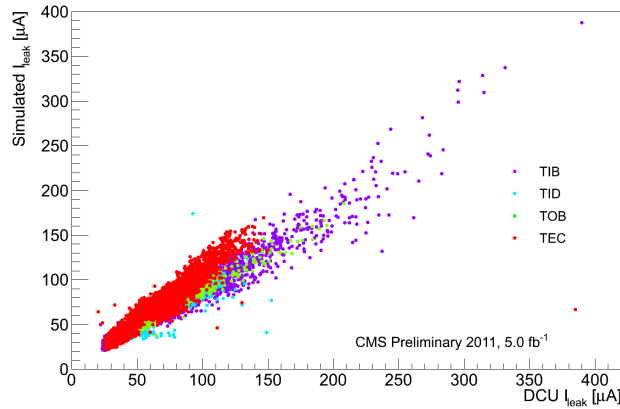


Figure 5: Leakage currents in the silicon strip modules: predicted by simulations for 5 fb⁻¹ of integrated luminosity of *pp* collisions versus measurements at the end of 2011.

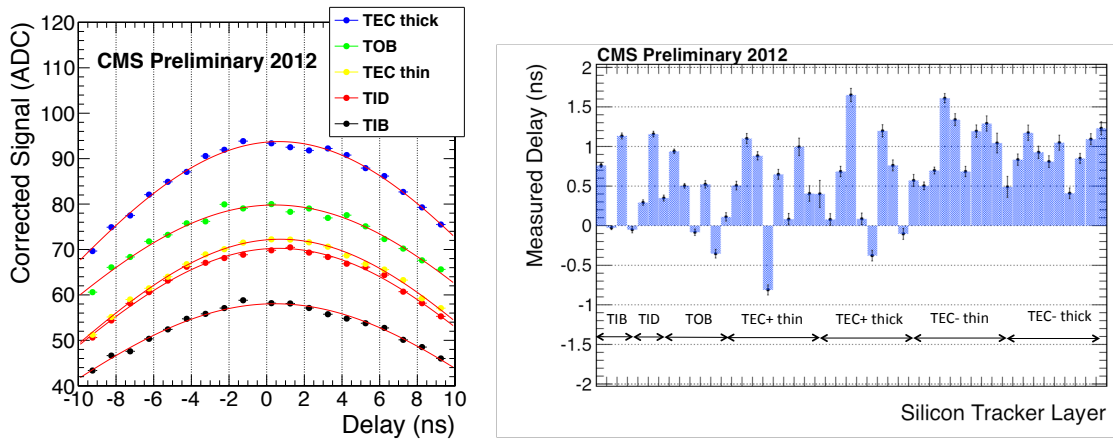


Figure 6: Timing scan results in 2012: Mean signal counts versus delay for groups of modules of same thickness in a subdetector (left) and peak positions of such distributions per layer (right).

The good performance of the strip tracker manifests itself in an excellent signal-to-noise ratio for clusters associated to reconstructed tracks. Figure 7 shows the distributions for TIB and TOB modules working in deconvolution mode. The most probable values (MPV) exceed 18 and 20 in the TIB and TOB, respectively. These high signal-to-noise ratios result in hit reconstruction efficiencies far above 99% in all detector layers for working modules. This is shown in figure 8 together with the efficiencies if all modules are considered. The latter are smaller by definition and reflect the bad module distribution. The resolution of the reconstructed hit positions on the modules has been measured using tracks with two hits in one layer, i.e. in the overlap regions of the modules. The employed double-difference technique is largely insensitive to misalignment effects and profits from the little material between the two modules. The results, shown in figure 9 for TIB and TOB for the different strip pitches, show that values below the binary resolution are reached. This is independent of whether hits of one, two or three strip clusters are investigated, or whether, irrespective of the number of strips, all hits are considered where the track crosses the




Figure 7: Signal-to-noise ratio for modules in the TIB (left) and TOB (right), corrected for the track angle relative to the module surface. The red lines represent fits of a Landau distribution convoluted with a Gaussian.




Figure 8: Hit efficiencies per layer measured with tracks crossing the modules inside its boundaries.




Figure 9: Resolution of the reconstructed hit positions for modules on layers with different strip pitches.

modules almost orthogonally.

4. Alignment

Track-based alignment is employed for the CMS silicon tracker using the global fit approach of the MILLEPEDE II program [4]. The global χ^2 , built from track-hit residuals, is minimised as a function of track and geometry parameters. The alignment procedure [5] applied for the reprocessing of 2011 data made use of about 22 million tracks of different topologies: high p_t isolated muon tracks, low momentum tracks, cosmic ray tracks and pairs of muon tracks from Z^0 boson decays. In total about 200 000 alignment parameters have been determined in a single fit, i.e. up to nine parameters for each of the 24 244 strip and 1440 pixel sensors: three translations, three




Figure 10: Average track angle corrected residuals $\langle (u_{trk} - u_{hit}) / \tan \psi \rangle$ (weighted by $\tan^2 \psi$) as a function of the normalised track position along the measurement (u_r , left) and insensitive (v_r , right) coordinates of all TOB modules, and average χ^2 probability of cosmic ray track fits versus the closest approach to the beamline for different module surface parametrisations (right). Compared are different surface shape parametrisations used in alignment. The determined coefficients are applied in the hit position reconstruction.

rotations and three sensor curvature parameters. The latter incorporate that the sensors are not flat, e.g. due to single-sided processing.

Sensor Curvatures Non-flatness of modules can be investigated using the average track angle corrected residual $\langle (u_{trk} - u_{hit}) / \tan \psi \rangle$ where ψ is the angle between the module normal and the track in the plane spanned by the normal and the measurement direction u . These are shown in figure 10 for cosmic ray muon tracks recorded in peak mode. For an alignment that assumes flat modules (open magenta circles), one can see that the modules are curved along the measurement direction u . Along the insensitive coordinate v one sees a significant kink between the two daisy-chained sensors. The final alignment determines positions and orientations of the two sensors of each module independently and furthermore determines, for each sensor, coefficients of second order polynomials that parametrise the curved sensor surface. If this alignment is applied and the curvature coefficients are used to correct the hit positions depending on the track angle, the modules appear to be flat (full black circles), as expected by the track reconstruction. Figure 11 shows averages of the determined curvature coefficients and of the kinks between the daisy-chained sensors in TOB and TEC.

The determination of the kinks between the sensors and their curvatures is relevant as demonstrated in the right of figure 10. Without these corrections to the hit positions the track fit probability $\text{Prob}(\chi^2, N_{df})$, depending on the number of degrees of freedom N_{df} and the χ^2 of the track fit, significantly degrades for larger distances to the beamline, caused by the larger average track angles on the module surfaces.

Looking at figure 10 one can see an offset of roughly 5 μm of the flat curve for the final alignment. This is related to the module timing shift of 3 ns for the TOB modules in the beginning of 2011 (cf. section 3). If the time delay of a module is slightly off, not all charge from the back side of the module is collected in time in the strips for deconvolution mode readout. The modules are effectively thinner and their position along the direction of the module normal appears shifted. A so called “back-plane correction” was calibrated using 2010 data and applied in hit position

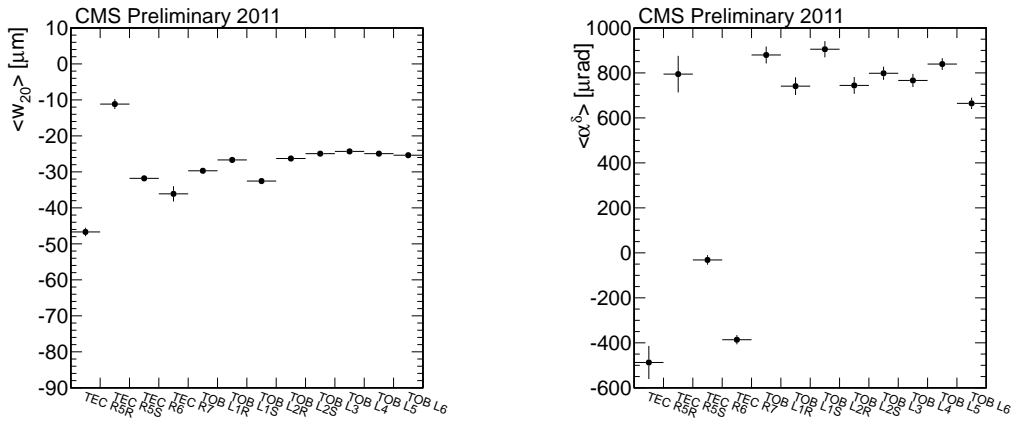


Figure 11: Curvature along the measurement direction (left) and half kink angle between the two sensors (right) for modules with daisy-chained sensors. Shown is the average per layer (L) and ring (R), respectively, separating stereo and $r\phi$ modules.

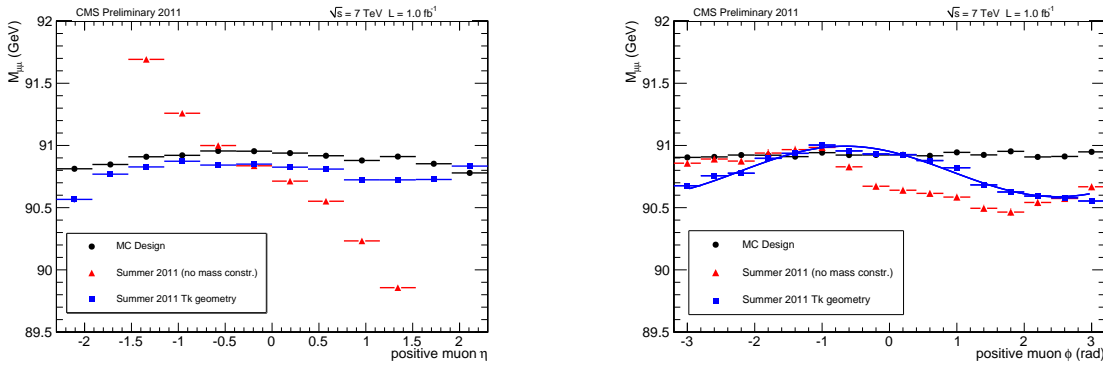


Figure 12: Peak of the distribution of the reconstructed invariant mass for $Z^0 \rightarrow \mu^+\mu^-$ decays as a function of η (left) and ϕ (right) of the μ^+ , comparing the 2011 alignment (blue squares), an alignment without using the Z^0 mass constraint (red triangles), and a simulation with perfect alignment (black circles).

reconstruction for deconvolution mode data to compensate the apparent difference of the module positions between the readout modes. The back-plane correction has not been re-calibrated after the timing shift in 2011, resulting in a tension between the effective module position in peak and deconvolution mode during the alignment procedure. This tension leads to the offset in figure 10.

Momentum Scale A major challenge of the χ^2 minimisation as a function of track and geometry parameters is that geometrical distortions and track parameter biases can compensate each other, leaving the χ^2 unchanged. This can lead to biased track parameter measurements. To overcome this problem, it is essential to utilise datasets with different track topologies as input, especially cosmic ray tracks. In addition, the 2011 alignment adds information of the nominal Z^0 mass for events with $Z^0 \rightarrow \mu^+\mu^-$ decays into the procedure. As figure 12 shows, this controls the momentum measurement: the reconstructed Z^0 mass is almost independent of the direction of the μ^+ . Only a sinusoidal dependence on ϕ remains, with a small amplitude of about 2% of the Z^0 mass peak.

5. Summary

The CMS silicon strip tracker continued to run very efficiently in 2012. Even in the third year of LHC running more than 97.6% of the strip modules are active. Power supplies and cooling services show a low failure and leakage rate, respectively. Automation of the HV raising is introduced and the CMS soft error recovery procedure of CMS is utilised to further reduce the CMS downtime caused by the strip tracker. Leakage currents as predicted by models for the accumulated radiation largely agree with the module-by-module measurements, giving confidence for future operation with even higher luminosities.

The good performance manifests itself in signal-to-noise ratios close to 20 in the barrel, a hit efficiency for active modules above 99% and a hit resolution well below the binary resolution.

The alignment procedure has been extended to determine about 200 000 parameters, including corrections for sensor curvatures caused by e.g. single sided processing. High local precision allows to spot effective module position differences for deconvolution readout caused by a timing shift of 3 ns. Momentum biases of tracks in the momentum range probed by $Z^0 \rightarrow \mu^+\mu^-$ decays are controlled by integrating information on the known Z^0 mass into the alignment procedure.

The long shutdown of the LHC in 2013 and 2014 will be used to consolidate the services and to possibly increase the fraction of active modules by up to 1%. Especially the cooling system will be refurbished to allow operation far below 0°C to mitigate radiation effects expected for the much higher integrated luminosities of the LHC after the shutdown. This will ensure that the CMS silicon strip tracker continues to provide high quality data for physics measurements and discoveries at the high energy frontier.

References

- [1] CMS Collaboration, *The Tracker Project Technical Design Report*, 1998, *CERN-LHCC-98-06*.
CMS Collaboration, *Addendum to CMS Tracker TDR*, 2000, *CERN-LHCC-2000-016*.
- [2] CMS Collaboration, *The CMS Experiment at the CERN LHC*, 2008, *JINST* **3** S08004.
- [3] L. Demaria, *The CMS Silicon Strip Tracker: Status and Performance*, 2011, *PoS (Vertex 2011)* 010 [CMS-CR-2011/181].
- [4] V. Blobel, *Software Alignment for Tracking Detectors*, 2006, *NIM* **A566** 5–13
(see also https://www.wiki.terascale.de/index.php/Millepede_II).
- [5] G. Flucke, *Alignment of the CMS Silicon Tracker*, 2012, *J.Phys.Conf.Ser.*, **368** 012036 [CMS-CR-2011-323].