The CMS Silicon Strip Tracker: Status and Performance

Natale Demaria*

INFN Torino, v.P.Giuria 1, 10125 Torino, Italy

E-mail: `demaria@to.infn.it`

The CMS Silicon Strip Tracker is the largest silicon detector ever built. The high precision measurement of the track hit position, over a large radius range, determines the unprecedented performance of CMS tracking. The operational experience matured on this very complex detector up to the first half of 2011 and the current status are described here, including the powering system, cooling and the back-end electronics. The commissioning sequence to calibrate the detector is illustrated and results on the detector performance are presented.

---

*Speaker.

1. Introduction

The CMS Tracker [1] is the tracking detector built for the CMS experiment at the CERN Large Hadron Collider. It is a unique instrument in both size and complexity: it contains two systems based on silicon sensor technology, one employing pixels and another using silicon microstrips. The Silicon Strip Tracker (referred to simply as SST) is the subject of this paper, and surrounds the pixel system: it is the largest silicon detector ever built, with 9.3 million sensor channels covering a surface area of $198\,\text{m}^2$.

The SST has been thoroughly tested before and after the installation in the experimental cavern, using cosmics rays [2]. The first collisions at CMS were recorded in December 2009 at energies of $\sqrt{s} = 900\,\text{GeV}$ and $2.36\,\text{TeV}$ and from 2010 onwards at an energy of $\sqrt{s} = 7\,\text{TeV}$.

In the first half of 2011 LHC has entered a new luminosity regime of $10^{33}cm^{-2}s^{-1}$ with an inter-bunch spacing of $50\,\text{ns}$ and a pile-up that is half of nominal. This paper reports on the Silicon Strip Tracker status, operational experience and performance up to the first half of the 2011 data-taking period.

2. The CMS Silicon Strip Tracker

The SST covers a radius from $24\,\text{cm}$ to $110\,\text{cm}$, providing a high precision momentum measurement for $|\eta|<2.5$. There are four major subsystems: in the low $|\eta|$ the Tracker Inner Barrel (TIB) with four layers and the Tracker Outer Barrel (TOB) with six layers; in the forward region the Tracker Inner Disks (TID) and the Tracker End Caps (TEC) with three small and nine large disks on each side respectively. Each TID and TEC disk is organized in three and four to seven rings respectively, where each ring corresponds to a different radial range. The layout of the detector is shown in Fig. 1. The fundamental building block of the detector is the silicon detector module, displayed in Fig. 2 for a TIB layer. All modules use single sided silicon sensors but the first two
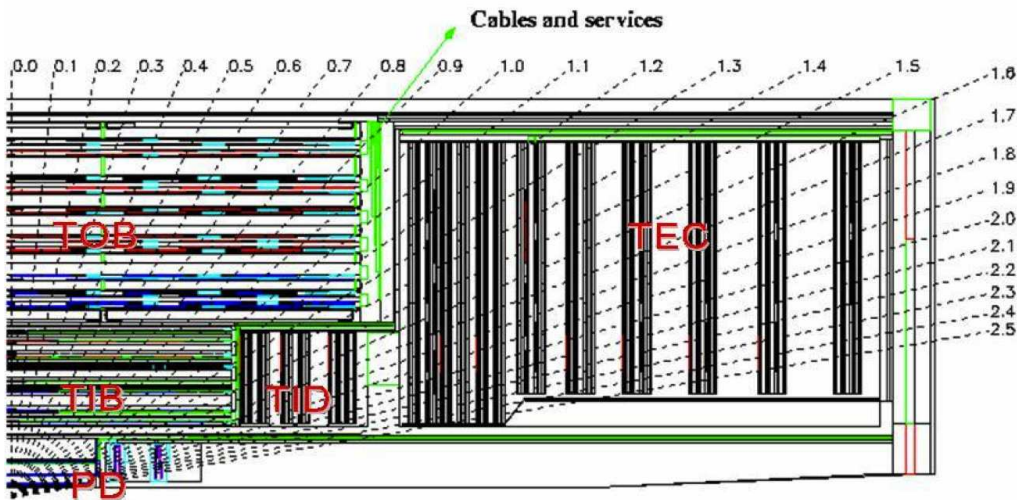

Figure 1: Rz view of one quadrant of the CMS Silicon Tracker.

layers of TIB and TOB and rings (1,2) of TID and (1,2,5) of TEC use sandwiches of back to back modules with both $R\phi$ (in the barrel) or $Z\phi$[1] (in disks) and stereo measurement in order to provide a three dimensional measurement of the track hit element.

All silicon sensors are single sided p^+ on n microstrip detectors, AC coupled with polysilicon bias resistors of 1.5 MΩ. Different silicon microstrip geometries have been adopted, with increasing strip pitch (80 to 200 μm) and length (9 to 20 cm) moving to the higher radius in order to optimize the occupancy. Noise increase determined by the higher capacitance load of outer silicon sensors have been compensated with a larger signal collected in a thicker substrate: 300 μm for $R < 58$ cm and 500 μm for $R > 58$ cm. Silicon resistivity has been chosen according to expected radiation fluences: lower (1.5-3 kΩ cm) for thin and higher (3.5-7 kΩ cm) for thick silicons. There are in total 14 different sensor designs.

The silicon sensor is connected to the APV25 front-end chip [3], which is mounted on a hybrid mechanically supported by a carbon fiber frame. The analog signal from the silicon sensors is amplified, shaped and stored in a 192 element deep analogue pipeline every 25 ns. A subsequent stage can either pass directly the pipeline signal (Peak mode) or form a weighted sum of three consecutive samples effectively reducing the rise time to 12.5 ns and containing the whole signal in 25 ns (Deconvolution mode).

The readout and the control chain are shown in Fig. 3. The signals from each APV25 are then converted to light by an Analog Opto Hybrid (AOH) which sent it via optical fibers to the Front End Driver board (FED) [4] where they are digitized and further processed, prior to transmission to the central DAQ system. Data can be taken in two different modes: Virgin Raw (VR) or Zero Suppressed (ZS). In VR mode, all channels are read out with the full 10 bits ADC resolution. In ZS mode, the FED applies pedestal subtraction, common mode rejection and a fast clustering algorithm, using signal height with a reduced 8 bits resolution: only channels forming a cluster are output. Clock, trigger signals, and slow control communication with the front-end electronics
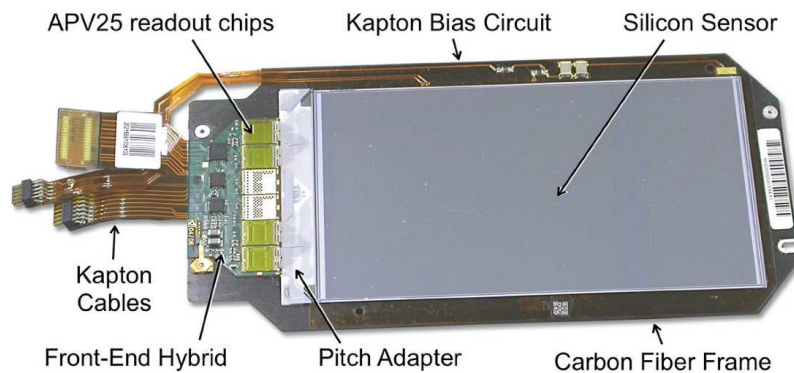

Figure 2: Silicon Strip module of the TIB.

are managed by the Front End Controller (FEC) boards[5] and sent via optical fibers to the Digital Opto Hybrid (DOH) for each control ring of SST: signals are distributed by the DOH to every

---

[1]The coordinates are refered to the centre of CMS, that is where collisions take place. Z is the coordinate along the beam axis, while R and ϕ are the radial distance and the azimutal angle in the transverse plane.

Communication Control Unit (CCU) [6] in a control ring. Finally each CCU sends signals to SST modules, in particular clock and trigger via a Phase Locked Loop (PLL) circuit on each module, and receives monitoring data provided by each module from the Detector Control Unit (DCU) chip [7].
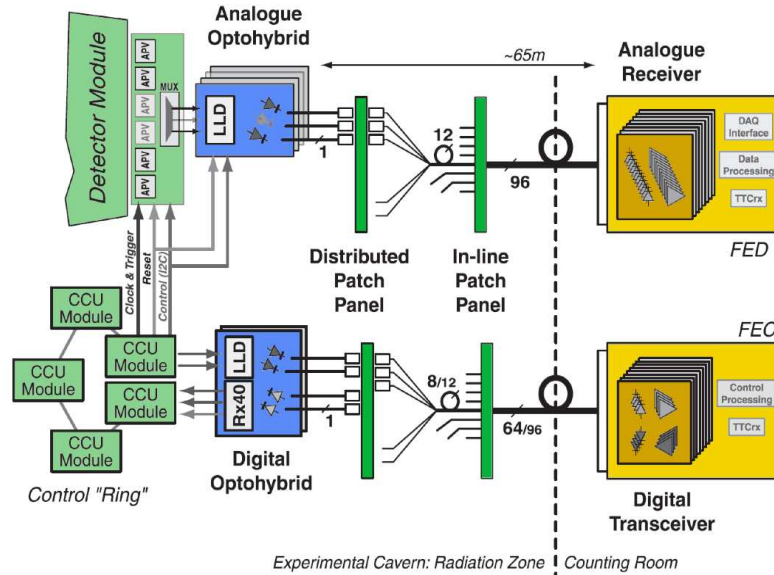

Figure 3: Readout and control chain of the Silicon Strip Tracker.

2.1 Cooling

Liquid cooling is used for the SST front end electronics. The liquid is circulated through aluminum pipes that are in contact with the silicon modules. The chosen liquid is C_6F_{14}, which is highly volatile and can be used at low temperature.

There are two main cooling plants distributing liquid to the SST via 180 lines. During 2010 two leaky lines were kept closed all the time on one cooling plant, nevertheless the leak rate increased during the year. In 2010-2011 winter stop a detailed leak search was performed and eventually three additional leaky lines were found and closed. Changes on the cooling system were done in order to improve the stability of the system, increase the pump and motor lifetime and minimize potential stress for the aluminium pipes. A Variable Frequency Drive (VFD) was commissioned, allowing a smoother start up of the cooling; the outgoing cooling plant pressure was reduced from 9 to 7 bars, and safety pressure switches were installed. Finally, pressure delivered to detector lines was decreased by almost 1 bar with no change on silicon module temperatures. An acceptable low leak-rate (0.7 kg/day) has been reached and is stable since six months. The front end electronics without cooling show a higher temperature but can be operated safely and reliably, they only needed to be recalibrated, mainly because of the change in AOH gain.

The coolant temperature used during 2010 and 2011 is $4\,^\circ$C which is compatible with the low irradiation level received in this period. A map of the silicon sensor temperature is shown in Fig. 4. The SST will go to the coolant nominal temperature of $-20\,^\circ$C after the first long shutdown on LHC

(starting 2013) , when the machine will achieve the nominal luminosity and the low temperature will be needed to limit the effects of radiation damage to the detector.
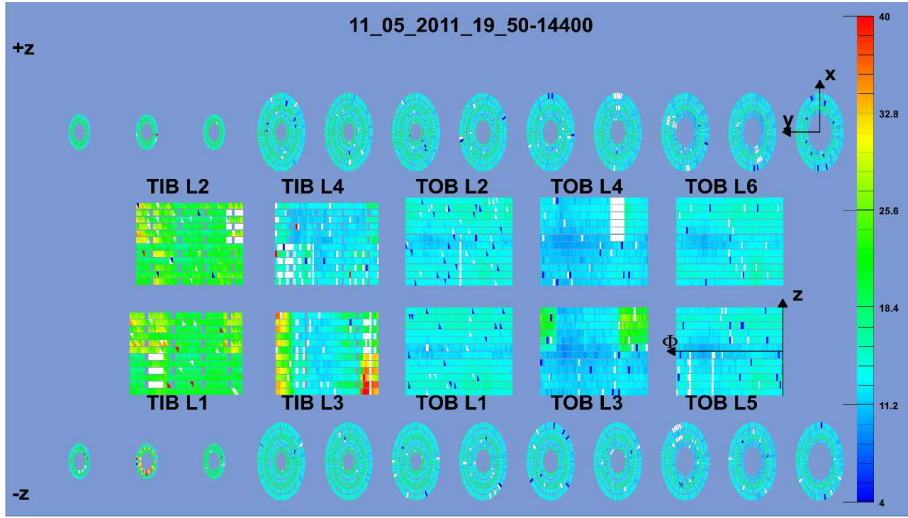

Figure 4: Map of silicon temperatures inside SST as measured by the DCU. White regions correspond to parts excluded from the readout. Parts that are not cooled are visible in TIB L3 (3 lines), TOB L3 (1 line), TID- middle disk (1 line).

2.2 Power Supply System

The SST is consuming a power varying between 30 and almost 60 kW, depending on the operation conditions. A custom power supply system was developed for the SST [8]. The power to the modules is provided by 1000 floating power supply modules, each one hosting two units. Each unit is connected via one 30-50m long cable to a part of the SST and provides low voltage lines for 2.5V and 1.25V and two high voltage bias lines for the silicon sensors.

The power to the control rings is provided from 356 control power supply units that are grouped in modules of four units.

All power modules are arranged in 29 crates situated on balconies of the experimental cavern and are controlled by 8 CAEN mainframes situated in an accessible area in the service cavern.

The level of failure during the year has been of the order of 1% during 2010 and has further decreased during 2011. The replacement of a power supply unit is normally done at first opportunity of down time of the LHC but can also wait for the next period of technical stop of the machine.

2.3 DAQ

The SST DAQ is made out of 440 FED modules, each processing the data coming from 96 fibers, and 48 FEC modules, each controlling 8 control rings. FED and FEC modules are arranged in 35 VME 9U crates. SST DAQ has been running stably. The only interventions required since 2008 have been: the early replacement of 5% of VME-PCI boards due to failures in the opto-receiver during the 2010; the replacement of two FEDs and of a few VME crate power supplies.

Improved FED firmware was deployed to increase the robustness of the readout against spurious extra frames coming from the detector. When a FED detects bad data, it goes into Out-Of-Sync status, waits for a resynchronization command (resync) that is sent automatically by the L1 trigger, and then cleans up the culprit data and restarts. The dead-time, caused by recoveries from these events accounts for less than 0.1%. The status of FEDs is monitored constantly during data taking, by accessing the FED information via VME bus. A modification of the firmware was needed to manage correctly a request of VME access contemporary to a request of FED resync.

During 2011 the reading of slow control data provided by the DCU at module level (temperatures, leakage current, low voltage) was improved. This is an important ingredient to understand detector instabilities, working conditions and radiation damage.

3. Detector Operation: Running and Commissioning

The SST has demonstrated to be a very reliable and stable detector. It has collected high quality physics data with an uptime greater than 98.5% during pp running in 2010-2011 and 99.7% during the 2010 heavy ion running. Sources of downtime have been identified and problems were properly addressed. Important steps have been the commissioning of the detector and achieving good stability of the detector systems.

3.1 Commissioning

In order to bring the SST detector into an operational state suitable for data-taking, several commissioning procedures are required to configure, calibrate, and synchronize the various hardware components of the control and readout systems. The full commissioning sequence of the SST is explained in Ref [2]. It is mostly based on calibrations done without an external trigger and consists of the following sequence: internal synchronization of all analog signals from front end chips to the FED; gain equalization of all AOHs transferring the signal out via optical fibers; optimization of the average baseline at the FED; tuning of the pulse shape of each chip and finally the measurement of the average baseline (pedestal) and of the noise of each channel. Some of these steps are shown in Fig. 5. The measured pedestal and noise values are then uploaded to the FED and used to perform the zero suppression of the data. In absence of a real signal, the SST is very quiet: occupancy due to noise is of the other of 10^{-5} in deconvolution and 10^{-6} in peak mode, to be compared with about 1-4% occupancy during collisions at full luminosity. Gain, pedestal and noise have been monitored by taking periodically timing and pedestal runs and they have been confirmed to be very stable: updates where done only occasionally in order to recuperate few per mil of the detector, following minor hardware interventions done during LHC technical stops.

Detector parts that are malfunctioning are mainly detected during SST commissioning. Only good channels are selected to be readout and during the year 2010 they constitute 97.8% of SST. This has to be compared with the 98.5% obtained after the installation of the SST in CMS and first commissioning. The fraction of good channels in the several sub-detectors is as follow: 94.3% for TIB, 98.1% for TID, 98.2% for TOB and 98.9% for TEC. The major component to bad channels comes from two control rings: the control cable has a short at the level of a patch panel located outside the SST, but inside CMS. These two shorts are responsible for 0.7% of bad channels and will be recovered during the next LHC shutdown in 2012. Other bad channels are due to one control

ring (0.35%) HV lines missing (0.1%); HV lines shorts (0.7%); bad fibers and other problems (0.4%). It has to be underlined that the SST has been designed with high redundancy and can accept a high level of dead channels before tracking performance is affected.

The signal in deconvolution mode is very fast and therefore a precise timing with respect to the trigger is necessary. A preliminary scan with 20 steps of 2 ns each was done during 2009 for one TOB layer and was repeated in early April 2010 scanning one layer per sub-detector. These measurements take few hours since they require several run stops in order to reconfigure modules parameters. A new procedure, taking half an hour, was developed: one detector configuration is created and uploaded once, setting a randomized delay within $\pm$ 10 ns across the nominal values for each module, allowing an analysis at layer level. Results are shown on the left side of Fig. 6 where the signal is plotted as function of the delay: the expected width of 12ns of the deconvolution signal is visible. From this measurement, the time position of the maximum signal with respect to the nominal setting point, can be computed for each layer: results are shown on the right side of Fig. 6 and deviations are within 1.5 ns, determining a variation of 1% in signal height.

3.2 Operation with pp and Heavy Ion collisions

The high voltage of the SST is raised once LHC has declared stable beams and after verifying that background conditions measured by the CMS beam radiation monitor are fine: the turn on takes 70 sec and has negligible impact on the data taking efficiency. The high voltage is turned off again at beam dump handshake. During inter-fill periods, when there is no beam in the machine, the SST is turn on again to collect cosmic-ray data that, together with collisions data, are crucial ingredients for measuring the SST alignment. Periods with several hours of beam downtime were also used to increase the statistics of cosmic rays.During pp running the FEDs are readout in ZS at a L1 trigger rate of several tens of kHz: the SST event size is about 250 kbytes.

The FED firmware has an algorithm to estimate the baseline per event that is optimized for pp collisions. The Heavy Ion collisions are characterized by a much higher occupancy and therefore it was decided to readout the whole tracker data, without zero suppression, i.e. in VR mode: the optimal baseline estimation is made by offline analysis and then applied event per event. The SST event size is of about 20 Mbyte and CMS DAQ needed an optimization in order to run with large events at lower L1 rates. The required L1 rate of about 250 Hz was achieved, while the data
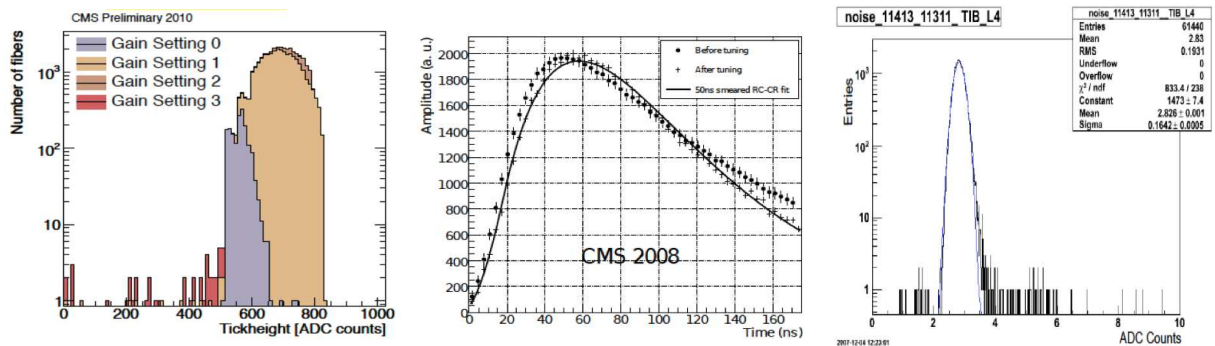

Figure 5: Commissioning of SST. Gain equalization, Pulse shape tuning, a noise profile for TIB layer 4

passing the High Level Trigger (HLT) were recorded at about 100 Hz, determining a data rate above 1 Gb/sec, that is about ten times higher than during pp running.

The offline zero suppression algorithm has been studied in detail and will be applied at HLT level during 2011 heavy ion running: that will allow CMS to take data at a L1 rate of about 3 kHz. Dry runs have been done already and were successful.

3.3 Radiation damage

As expected, silicon detector leakage currents have started to show a linear increase as a function of the integrated luminosity. This is due to well-studied radiation-damage effects induced by charged and neutral hadrons. Normalization of current to the silicon volume and sensor temperature allows to determine the distribution of the radiation dose in the detector. A preliminary analysis based on high-voltage currents in power supplies demonstrates that the radial dependency and absolute values are compatible with expectations, more details can be found in [9]. High voltage scans were done to determine the depletion voltage: they have shown so far no change with respect to the initial values. More quantitative studies are in preparation.

3.4 Performance Studies

The signal-to-noise ratio is an important ingredient to guarantee high hit reconstruction efficiency and good spatial resolution. Values for deconvolution running mode are very good and fit with expectations: results normalized to an incident angle of 90° to the silicon sensor for TIB and TOB are shown in Fig. 7. The most probable value of the Landau has been measured to 18.5, 19.4, 23.9, 18.4 and 22.4 respectively for TID, TEC thin, TEC thick, TIB and TOB sub-detector modules.
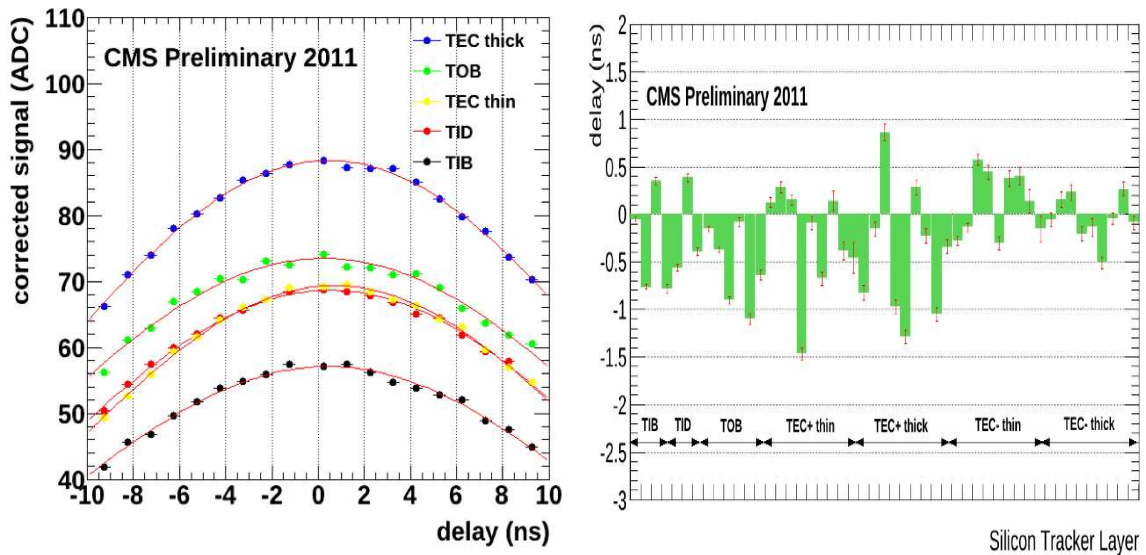

Figure 6: On the left: results on measurements of signal time profile versus the delay with respect to the nominal trigger timing. On the right: time synchonization of the maximum signal versus the layer/disk of the SST, where displacements of max 1.5ns are visible.
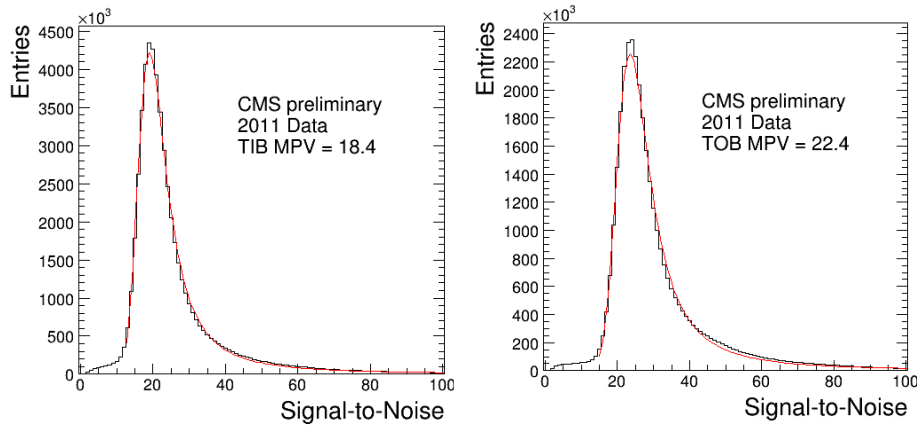
Figure 7: Signal-to-noise ratio for TIB and TOB modules.

The efficiency to reconstruct a hit is measured counting how many times a good track extrapolated to the layer produces a hit. Cuts have been applied, to avoid crossing the layer at the border of the acceptance region, given the extrapolation error of the track, and known bad modules have been excluded. The overall measured efficiency is 99.9% and the results per SST layer/disk is shown in Fig. 8.
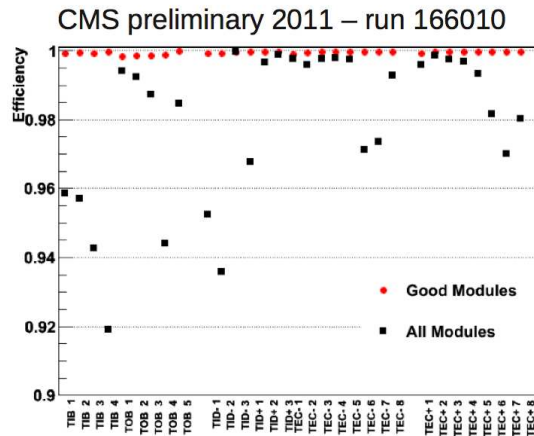

Figure 8: Hit efficiency for all SST layers/disks: considering all modules (black) and considering only modules that are good (red).

The hit resolution has been measured using tracks crossing module overlap regions. The small distance and amount of material between the overlapping modules makes the comparison of local coordinates of the two modules a good estimate of the hit resolution. The measurement is almost insensitive to local misalignment, except for uncertainties on the relative angles between the two modules, which have a negligible contribution at first order. Results are shown in Fig. 9 for tracks almost perpendicular to the silicon surface. The spatial resolution is well below the binary resolution determined by the strip pitch.
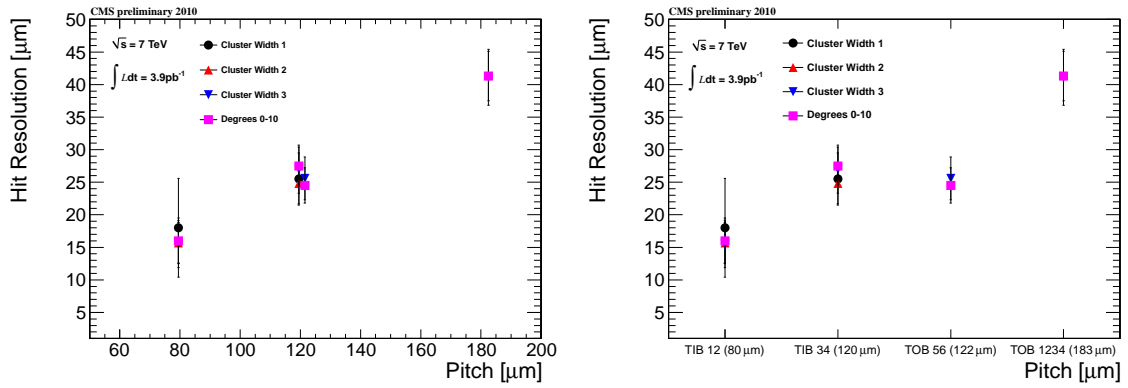
Figure 9: Detector hit resolutions measured using overlaps, for different cluster widths. On the left is shown the dependence with strip pitch; on the right the different module types are specified.
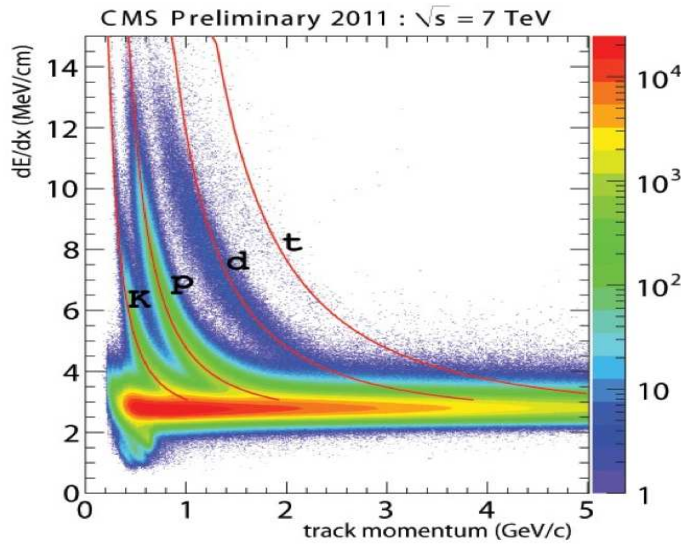

Figure 10: Measurement of dE/dx in the Silicon Strip Tracker.

For each track the SST provides not only a measurement of momentum, but also a multiple measurement of the ionization energy loss dE/dx and therefore provides some particle identification capabilities. A barrel track can include 10 independent measurements from the SST. The dE/dx distribution versus the momentum can be seen in Fig. 10: kaon, proton, deuteron and tritium bands are visible. Solid lines are Bethe-Bloch expectations extrapolated from a fit to the proton line: deviations at large dE/dx are due to saturation effects.

3.5 Conclusions

The Silicon Strip Tracker of CMS has demonstrated excellent performance during 2010-2011 collisions, with little downtime due to the stability of all detector systems, in particular the cooling, power supplies and back end electronics. The signal-to-noise ratio, cluster reconstruction

efficiency, spatial resolution and dE/dx measurement are all extremely good. This has been achieved thanks to a detailed detector commissioning and calibration procedures. The large increase of instantaneous luminosity did not adversely affect the detector performance, and the first effects of radiation damage do follow expectations.

References

[1] The CMS Collaboration, *The CMS experiment at the CERN LHC*, **JINST 803** (2008) S08004.

   The CMS Collaboration, *CMS: The Tracker Project Technical Design Report*, **CERN-LHCC-98-06** (1998).

   The CMS Collaboration, *Addendum to CMS Tracker TDR*, **CERN-LHCC-2000-016** (2000).

[2] The CMS Collaboration, *Commissioning and Performance of the CMS Silicon Strip Tracker with Cosmic Ray Muons* **JINST 5** (2010) T03008.

   W.Adam, N.Demaria et al., *Performance studies of the CMS Strip Tracker before installation* **JINST 4** (2009) P06009.

[3] M.J. French et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, Nucl. Inst. and Methods **A466** 359–365 (2001).

   M. Raymond et al., *Final Results from the APV25 Production Wafer Testing*, 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany (2005) 453–457.

[4] C. Foudas et al., *The CMS Tracker Readout Front End Driver*, IEEE Trans. Nucl. Sci. **52** 2836–2840 (2005).

   G. Iles et al., *Performance of the CMS silicon tracker front-end driver*, 10th Workshop on Electronics for LHC and Future Experiments, Boston, USA (2004) 222–226.

[5] K. Kloukinas et al., *FEC-CCS: A common front-end controller card for the CMS detector electronics*, 12th Workshop on Electronics For LHC and Future Experiments, Spain (2006) 179–184.

[6] C. Paillard, C. Ljuslin, A. Marchioro, *The CCU25 : a network oriented communication and control unit integrated circuit in a 0.25 μm CMOS technology*, 8th Workshop on Electronics for LHC Experiments, Colmar, France (2002) 174–178.

[7] G. Magazzù, A. Marchioro, P. Moreira, *The detector control unit : an ASIC for the monitoring of the CMS silicon tracker*, IEEE Trans. Nucl. Sci. **51** 1333–1336 (2004).

[8] S. Paoletti et al., *The Powering Scheme of the CMS Silicon Strip Tracker*, Proceeding of the 10th Workshop on Electronics for LHC and future Experiments, Boston, USA (2004).

[9] F.Hartmann, *Evolution of Silicon Parameters due to Irradiation at the LHC*, See proceeding of this conference.