

CMS Tracker commissioning and first operation experience

Christophe Delaere*, on behalf of the CMS collaboration.

CERN, Geneva, Switzerland.

E-mail: christophe.delaere@cern.ch

The CMS silicon strip tracker is the largest device of its type ever built. There are 24244 single-sided micro-strip sensors covering an active area of nearly 200 square meters. After a short introduction on the tracker, the program at the CMS tracker integration facility will be described. The strategy and results from the commissioning will be presented together with results on low-level detector performance. The general experience gained by operating the Tracker at different temperatures will be presented. This includes hardware aspects, acquisition software and infrastructures choices, or distributed data processing.

*16th International Workshop on Vertex detectors
September 23-28, 2007
Lake Placid, NY, USA*

*Speaker.




Figure 1: Schematic cross-section through the CMS tracker. Each line represents a detector module.

1. Introduction

The inner tracking system of CMS is designed to provide a precise and efficient measurement of the trajectories of charged particles encountered during LHC operations[1]. It surrounds the interaction point and has a length of 5.8m for a diameter of 2.5m, so to cover pseudorapidities up to 2.5. A detector technology featuring high granularity and high readout speed is needed to cope with the 1000 particles produced every 25ns, when about 20 p-p collisions occur for each bunch-crossing at high luminosity. These requirements led to a tracker design entirely based on silicon detector technology. Hence, the CMS silicon tracker is the largest device of its type ever built. It is based on pixel technology close to the interaction point, while the reduced particle flux allows the use of silicon micro-strip detectors from a radius of 20cm.

The silicon strip tracker (SST) is divided into four main subsystems (Figure 1). The central region is made of the Inner Barrel (TIB) that extends from $r=20\text{cm}$ to $r=55\text{cm}$ and is composed of four layers, and the Outer Barrel (TOB) that extends to $r=116\text{cm}$ and consists of six layers. In the forward region, the Inner Disks (TID) and the Endcaps (TEC) are made of respectively three and seven disks, up to $|z|=282\text{cm}$. There are 24244 single-sided micro-strip sensors covering an active area of 198m^2 . Throughout the tracker, the strip pitch varies from the inner to the outer layers (from $80\mu\text{m}$ to $205\mu\text{m}$) in order to cope with the anticipated occupancy and to grant a good two-hit resolution.

The size of the device has led to a design where the basic unit, called a module, houses the silicon sensors and the readout electronics. There are 15148 modules. All modules are supported by a frame of carbon fiber or graphite, with a capton layer used to insulate the sensor(s) from the frame, and to provide the electrical connections and temperature probe readout (Figure 2).

Charges are collected every 25ns and stored in an analog pipeline (APV25) on the front-end hybrid [2]. The APV25 has been designed in an IBM $0.25\mu\text{m}$ bulk CMOS process. It has 128 readout channels, each consisting of a low noise and power charge sensitive pre-amplifier, a 50ns CR-RC shaper and a 192 elements deep analogue pipeline where samples of the shaped signal are kept up to $4.8\mu\text{s}$. The CMS tracker is not able to produce a trigger signal by its own. An




Figure 2: (a) Exploded view of a module housing two sensors; (b) Photograph of a TEC ring 6 module mounted on a carrier plate.




Figure 3: Typical pulse shape in peak and deconvolution mode. From [5]

external (Level-1) trigger [3] generated from the information collected by other subdetectors is fed by dedicated optical links from the front-end controllers to the APV25 chips. Upon reception of a trigger signal (typically at 100kHz), the stored charge can either be passed as is (peak mode) or preprocessed by a deconvolution circuit that reduces the signal width by returning a weighted sum of three adjacent cells (deconvolution mode). Data is sent analogically via optical links to the front-end drivers (FED) [4] where the analog-to-digital conversion is done. An internal calibration circuit allows to inject charge with programmable amplitude and delay into the amplifier inputs in order to be able to calibrate and monitor the pulse shape. Figure 3 shows the pulse shape in both peak mode and deconvolution mode.

2. The tracker integration facility

The tracker integration facility (TIF) is the place where the final integration of tracker sub-detectors into the tracker support tube took place. It's a large clean room environment equipped with all needed services to fully commission and validate the detector prior to its transportation to the CMS pit and its installation in the underground cavern. Since extensive test periods were part of the integration process, the TIF program was tremendous opportunity to gain experience with

all aspects of detector operations in an accessible environment. In addition to the TIF, a Tracker Analysis Center (TAC) has therefore been installed on the CERN Meyrin site, in the vicinity of the TIF. That center provided a small computing farm and storage facility, and allowed both the tracker online operations (data acquisition, detector safety, data quality monitoring and online event display) and the offline analysis of collected data. It is envisaged that the TAC will evolve into the tracker operation center for experts and users alike during CMS operations.

The TIF program can be divided into three phases, lasting from November 2006 to July 2007.

Phase I: Integration (From November to March)

Each subdetector arrived at CERN and was inserted in the tracker support tube. TOB was the first as it was directly integrated while seated into the tube. It was followed by the TIB/TID, themselves followed by each of the endcaps.

At each integration step, a full set of preliminary tests and optical survey were performed, followed by the mechanical integration itself. It was followed by the connection of all services (electrical, optical and cooling) and another full set of quality tests.

During that period, the installation of general services was also carried on. It includes a gradual commissioning of all the needed power supplies (both LV and HV), commissioning of the detector safety system, dry air system and cooling plant.

Phase II: Warm developments (From March to May)

The second phase of TIF operations was dedicated to the development of all subsystems to achieve the operation of 12.5% of the tracker. Weekdays were therefore dedicated to the development of DAQ procedures, and to the refinement of commissioning sequences with bigger and bigger systems. It was also the occasion to perform various detector tests such as (cross-)noise studies and to implement hardware fixes when needed.

Weekends were then dedicated to cosmic data taking, mainly to assess the system stability and to collect data for tracking and alignment studies. All that was performed at room temperature.

Phase III: Cold Operation (From May to July)

The last months were then dedicated to the operation at various colder temperatures. Starting at 10 degrees, the system was then brought successively to zero degrees, -10 degrees and finally -15 degrees. The focus was on performances and safety, as well as on the study of potential temperature-dependent defaults. Cosmic data and Laser Alignment data were accumulated at each temperature point for offline processing.

3. The CMS tracker readout and DAQ

Figure 4 shows a schematic of the control and readout systems for the silicon strip tracker.

The control system [6] comprises 300 “control rings” that start and end at the off-detector Front-End Controller (FEC) boards and are responsible for distributing slow control commands, clock and trigger signals to the front-end electronics. The control signal emitted by the FEC is first transferred over digital optical links to digital opto-hybrids and then electrically via “token rings” of Communication and Control Units (CCU) to the front-end electronics.

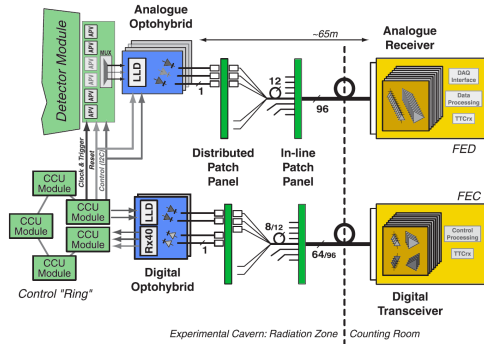


Figure 4: Schematic view of the silicon strip tracker control and readout system. From [8].

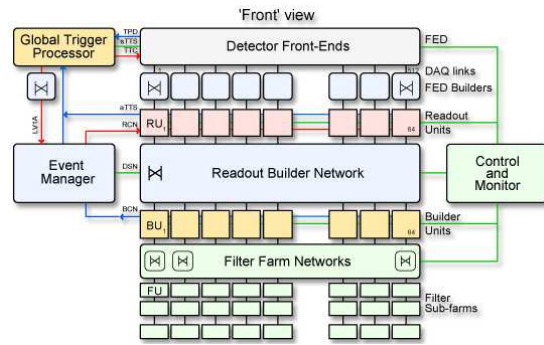


Figure 5: The CMS event builder and related services.

The readout system is based on the APV25 and on the off-detector Front-End Driver (FED) boards. They are interconnected by analog optical fibers. The signal from two APV25 chips is multiplexed onto a single line, converted to an analog optical signal and transferred to the FED boards. The FEDs digitize, process and format the pulse-height data from up to 96 pairs of APVs.

The data acquisition scheme at TIF has been designed to be as close as possible to the final CMS data acquisition system. Most of the differences are consequences of the unavailability of some hardware. The event data enters the event builder [7] as a set of fragments distributed over 65 FEDs (in CMS there will be about 700 FEDs, 440 of which are used for the SST). The task of the event builder, shown schematically on Figure 5, is to collect all the fragments, assemble them into a complete event, and forward them to a processing node on the computing farm.

The CMS event builder uses two stages of switch network. The first stage groups data from 16 FEDs into 64 super-fragments via small switch networks connected to the readout units. In the second stage, complete events are build using readout builders that connect each of the 64 builder units to the 64 sources of super-fragments. The Builder units buffers and forward complete events to the filter units which apply the High-Level Trigger (HLT) algorithm and perform other online data processing and monitoring. That two-stage architecture gives scalability and flexibility to the system.

For the smaller scale operations at TIF, the event builder was build around a single-stage conventional Gigabit Ethernet switch. This was one of the limiting factors for the data acquisition rate. There were six readout units, six builder units and twelve filter units. All twelve filter units were connected to a single storage manager. Archived data was then published on DBS (the Data Set Bookkeeping System, [9]) and shipped to the Fermilab Tier-1 and to the Bari Tier-2 for offline processing. Data were later distributed to Bari, CERN, Pisa and FNAL for analysis.

An external (Level-1) trigger has to be provided to readout data from the analog pipelines and transfer it to the FEDs. Most of the commissioning is performed using a cyclic (random) trigger at about 10Hz generated by the local trigger controller (LTC), but a physics trigger was desired to allow to record cosmic tracks crossing the tracker volume. A scintillator-based cosmic trigger has been installed for that purpose. It is made by the coincidence between scintillators above and below the tracker tube (Figure 6). 5cm of lead blocks were installed on top of the lower scintillators to cut the low-energy tail. With that setup, a trigger rate of 6.5 ± 0.1 Hz has been achieved.

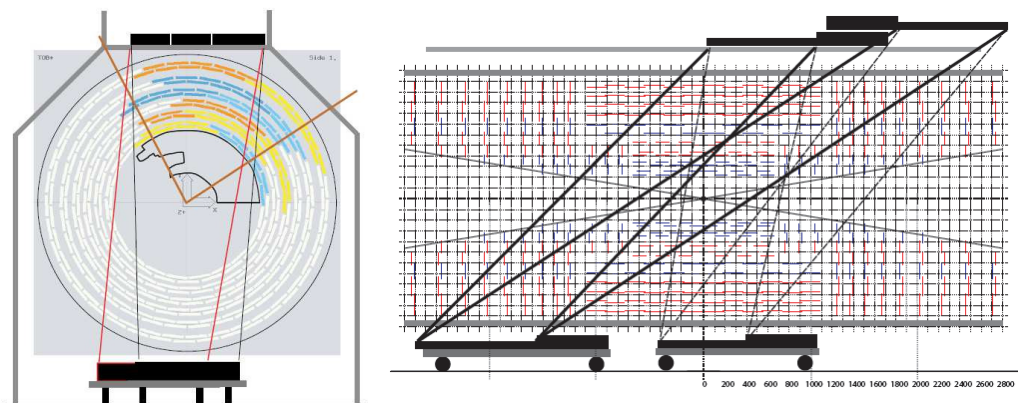


Figure 6: Schematic view of the trigger configuration at the TIF.

4. Commissioning the CMS tracker

Bringing the detector into an operational state suitable for physics is a complex task that comprises several independent *commissioning procedures* in order to configure, synchronize and calibrate the various components of the readout system. During TIF operations, various subsets of the silicon strip tracker have been repeatedly commissioned to form setups of various sizes (culminating with the operation of 15% of the tracker readout) at various temperatures.

The commissioning of the tracker is subdivided into the following steps:

Detector partitioning and cabling Hardware scans are performed to identify all front-end devices and off-detector FEDs that share a common trigger source. It is verified that every component of the front-end can be accessed via the token-ring and that no electrical power cable is swapped.

Control-Readout map The cabling of the optical link system is determined by sending unique patterns from each laser driver to the FEDs. That way, the one-to-one matching between APVs and readout channels is established.

Internal Timing This is the first data-related commissioning step. Internal synchronization of the SST front-end is achieved using dedicated digital pulses (“tick marks”) generated by the APV25. This procedure effectively accounts for signal propagation delays in the control system. The readout is synchronized using a programmable phase-locked loop chip on the front-end hybrid so that all tracker modules are responding at the same time. The resulting uncertainty is better than 0.5ns.

Optical Gain The laser bias and input signal scale factors for each channel of the front-end laser devices are tuned in order to optimize the usage of the available dynamic range.

Analog baseline tune The level of the analog baseline is adjusted to maximize the dynamic range while ensuring a good signal stability.

Tuning the APV pulse shape The pulse shape is tuned to have the design characteristics. That operation is done in peak mode, where the signal must correspond to a 50ns proper time RC-CR curve. That ensures a good homogeneity of the response over the detector, and that these performances can be maintained over the years. By controlling the time behavior of the signal, one also ensure that the response for out-of-time particles (for example beam halo particles, cosmics or heavy stable particles) is as expected/simulated.

Pedestal and noise Pedestal and noise are measured and stored in the online database for later reference (for example during zero-suppression or cluster building at HLT). Bad strips are also flagged at that stage.

APV latency scan Every 25ns, collected charges are stored at the APV level in the analog pipeline, and retrieved upon reception of a level-1 trigger signal. The right bunch-crossing in the analog pipeline is located by finding the optimal latency between the trigger signal and the readout using physics events. The latency that maximizes the signal of clusters belonging to reconstructed tracks is chosen. The tracker is operated in peak mode.

Fine tuning of the pulse shape sampling A correct global timing alignment is crucial, as signal is attenuated by about 4% per nanosecond misalignment (using the deconvolution circuit). Even more, the smallest misalignment affects the “fake” occupancy from adjacent bunch-crossings. The pulse shape sampling is optimized to 1ns taking into account experimental effects like time of flight. That fine tuning is achieved by reprogramming phase-locked loops on each front-end hybrid to maximize further the signal in clusters belonging to reconstructed tracks. The tracker is operated in deconvolution mode.

While commissioning the silicon tracker, lots of data has been obtained. Detailed analysis of that invaluable data has been performed and delivered many results at all levels. There is far too much to detail everything here. In the following, some selected results will be presented to illustrate the overall procedure, its ability to spot problems, and the overall detector quality.

As it is the first data-based commissioning procedure, the internal timing allows to flag most of the badly responding front-ends or under-powered components and to spot broken or dirty optical fibers. Figure 7 shows the delay computed for each module as a function of the position on the control token ring. The signal from devices positioned later on the control ring naturally arrives later, and should be less delayed. This is the observed behavior. That Figure allows to spot easily problematic devices. In this case, one bad TEC module and 3 TOB rods purposely unpowered are clearly spotted.

The optical gain is optimized by scanning over the laser bias values while measuring the tick mark amplitude (that is the digital 1 level). Figure 8 shows the three key quantities monitored during the optical gain scan. In addition to the optimal gain that comes out of the procedure, the amplitude of the input and output tick marks is monitored, together with the slope of the scan curve. That information is complementary to that provided by the internal timing run when studying the optical fibers quality.

The APV pulse shape has to be tuned for each chip using the signal generated by the internal calibration circuit and adjusting programmable registers of the APV25. That guarantees homogeneity of the tracker response and stability of the performances over time. It is also important to

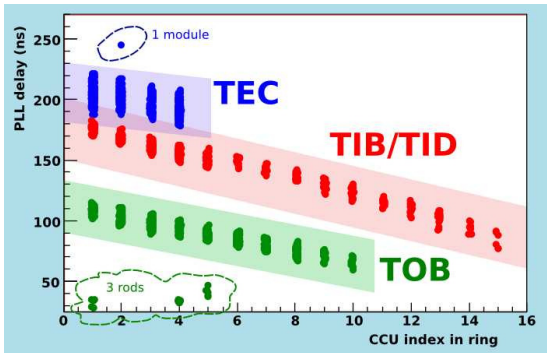


Figure 7: Delay computed for each module as a function of the position on the control token ring.

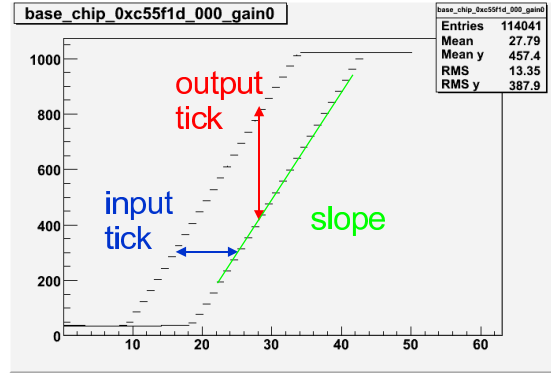




Figure 8: Key quantities monitored during the optical gain scan. The curves correspond to the level of the baseline and the top of the tick mark, as a function of the laser bias, for one value of input signal scale factor.



(a)



(b)

Figure 9: Normalized pulse shape for 800 APV, respectively before optimization (a), and when using individual optimal VFS values for each APV (b). The abscissa is time expressed in nanoseconds.

reduce the systematics in dE/dx studies and impacts the offline absolute calibration. To illustrate this, Figure 9 shows the impact of the “VFS tune” (tune of the shaper feedback voltage bias) on the pulse shape from 800 modules in the TEC. The reduction of the spread of the tail from one module to another is clear. The few non-nominal APVs are also easily spotted, and that information is made available for offline data processing.

Measures of pedestal and noise are necessary to feed the configuration database. That information is required later on to configure the zero-suppression algorithm in the FED and to perform the clustering, already in the HLT. One aim of the TIF operation was also to monitor the stability of the noise over time. It has been shown that the stability is excellent, indeed, the only variations observed over the 6 months of operation corresponding to the steps in temperature. It has also been confirmed that the noise scales as predicted with the strip length, as illustrated in Figure 10.

Since the TIF was equipped with a cosmic trigger, it was possible to take data containing actual signal clusters, and to study at large scale the signal performances. Very good S/N values are achieved, with a typical signal-to-noise ratio of 30/1, stable with time. The signal to noise ratio for each subdetector is shown on Figure 11. As it can be seen on that snapshot from the online data quality monitor, the signal is nicely Landau-shaped, and there is a very low amount of noise clusters. That opens the road to tracking and alignment studies, presented elsewhere in these




Figure 10: Equivalent noise charge after common mode subtraction versus strip length for all TOB and TEV module types, in peak mode.

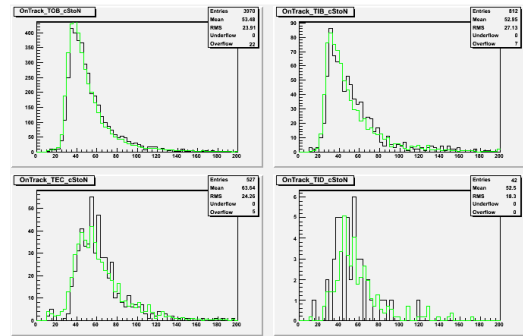


Figure 11: Signal to noise ratio for each subdetector as seen online by the data quality monitor client. From left to right and up to bottom: signal to noise ratio in TOB, TIB, TEC, and TID.

proceedings [10, 11].

All in all, the complete commissioning of the CMS tracker requires about eight hours, plus the time needed to perform the fine tuning of the APV sampling. That last step relies on the rate of high-Pt tracks and is expected to take up to one hour with the initial LHC luminosity.

Most of the commissioning will remain valid for a long period. The latency, for example, remains fixed as long as no modification is done in the cabling and logic of the trigger system. Pedestal, noise and gain will be continuously monitored during operations, eventually triggering in the LHC abort gaps. It is foreseen to perform typically one additional pedestal run per day, which takes about 20 minutes and can be done during a fill.

5. Conclusions: General experience

Not only was the TIF a great tool to exercise the tracker readout and commissioning, it was also a decisive large scale source of general experience in all fields.

The CMS tracker has been completed while maintaining a very high quality of production, and (pre)commissioned at CERN during six months. Standard data acquisition, detector control and data quality tools have been developed and used for the first time with more than one million channels. Substantial datasets have been obtained to evaluate tracker performances at various temperatures.

Important “engineering” tests have been carried out. A deep study and optimization of the grounding scheme for all subdetectors has been done, including many permutations of interference noise tests. Thermal performance were measured at various temperatures, and there were several tests of the Laser Alignment System. Dedicated runs were also performed to cross-check zero suppression algorithms by comparing zero-suppressed data to virgin raw data.

Cosmic data taking was a big achievement. 5,025,043 events were recorded and analyzed using the GRID infrastructures prepared for the final data-taking. Recorded data are used for further study, going from low-level detector performances to tracking and alignment. The stability of noise, gain, and other key parameters was studied over time. Tracker performances are excellent,

with a typical signal to noise ration around 30, 0.2% bad strips (including noisy strips and strips from low-gain modules) and an overall module efficiency exceeding 99.7%.

As a byproduct of daily operations, bookkeeping and communication tools have been deployed (and developed when needed). That includes run summary servers, Elog, snapshot server, database for run Quality flags, data transfer and access procedures.

The Data Quality and Monitoring (DQM) infrastructure was working very well to follow the daily activities and understand online the outcome of commissioning, from monitoring the module behavior (noise, signal, occupancy, ...) till the tracking (angular distribution, efficiency, ...). This was shown to be an invaluable tool to spot early any problem during physics runs.

Following all that, we got the confidence that we can run the tracker safely and effectively, and most importantly several new experts to run the experiment next year! All the CMS tracker community is now looking forward to have data as part of CMS global runs.

References

- [1] CMS collaboration, *Tracker Technical Design Report*, CERN/LHCC 98-6; April 1998; and CMS collaboration, *Addendum to the CMS Tracker TDR*, CERN/LHCC 2000-016, February 2000.
- [2] M.J. French et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, Nucl. Instr. Methods Phys. Res. A 466, Pages 359-365 (2001).
- [3] J. Troska et al., *Implementation of the timing, trigger and control system of the CMS experiment* IEEE Trans. Nucl. Sci. 53/3 834 (2006).
- [4] C. Foudas et al., *The CMS Tracker Readout Front End Driver* IEEE Trans. Nucl. Sci. 52 (2005) 2836-2840.
- [5] C. Delaere, L. Mirabito, *Timing of the CMS tracker. Study of module properties.*, CMS NOTE-2007/027.
- [6] F. Drouhin et al., *The CERN CMS Silicon Strip Tracker Control System*, CERN-CMS-CR-2004-032.
- [7] M. Pieri et al., *CMS DAQ event builder based on gigabit ethernet*, in proceedings of *Computing in High Energy and Nuclear Physics 2006*, Mumbai, India.
- [8] R. Bainbridge et al., *Data acquisition software for the CMS strip tracker*, in proceedings *CHEP 2007*, September 2007.
- [9] CMS collaboration, *Computing Technical Design Report*, CERN/LHCC-2005-023, 20 June 2005.
- [10] C. Noeding et al., *Track Reconstruction and Experience with Cosmic Ray Data in CMS* in proceedings of *16th International Workshop on Vertex detectors*, PoS(Vertex 2007)026.
- [11] G. Flucke et al., *Tracker alignment strategy in CMS and experience with cosmic ray data* in proceedings of *16th International Workshop on Vertex detectors*, PoS(Vertex 2007)029.