PROCEEDINGS OF SCIENCE

CMS Silicon Strips Operations and Performance

Erik Butz*
Massachusetts Institute of Technology
E-mail: ebutz@mit.edu

The CMS silicon strip tracker is the largest silicon detector ever built with almost 10 million readout channels and an active area of close to 200 m^2. With more than 15,000 individual micro-strip silicon modules are powered by almost 1000 power supply modules and produce more than 60 kW of power while operating at low temperatures. Results from the successful operation of the tracker at the first LHC collisions at 0.9, 2.4, and 7 TeV, including environmental control, calibration, detector performance, and monitoring, are discussed. The detector performance is excellent manifested in a nearly 100 % functional tracker with high single hit efficiency, good S/N performance, and high quality track resolution. This is made possible by a fine-grained calibration process and a monitoring of all important quantities for the detector performance during different stages of the operation.

19th International Workshop on Vertex Detectors
June 6 -11 2010
Loch Lomond, Scotland, UK

*Speaker.

PoS(VERTEX 2010)011

1. Introduction

In late Fall 2009, the LHC started operation after almost two decades of planning, construction and commissioning. The CMS detector[1] is one of two large multi-purpose detectors at LHC. At the center of the CMS detector, surrounded by the calorimeters and the large superconducting coil which provides an axial magnetic field of 3.8 T, lie the two tracking detectors: the silicon strip and the silicon pixel detector. The silicon strip detector has an active surface area of 198 m^2, with 24244 individual silicon sensor that are grouped in 15148 modules. The strip tracker employs *p*-in-*n* silicon wafers with 320 or 500 µm thickness and with 512 or 768 micro-strips each. More than 30 different sensor types are used to achieve best possible geometric coverage. With increasing radius, the strip length of the modules increases, in the outermost parts, two individual sensors are connected using wire-bonds.

The strip tracker is partitioned into different substructures that were built separately by different groups of the collaboration: the tracker inner barrel (TIB) and the tracker inner disks (TID) are the innermost part of the strip tracker with 4 layers for TIB and three disks in the TIDs on either side of the interaction point. The TID disks are further subdivided in ring structures of equal radius; each disk consists of three rings. Surrounding the TIB/TID is the large tracker outer barrel (TOB) which consists of 6 layers of silicon modules. On both sides, the strip tracker is completed by two tracker end caps (TEC(s)) which have nine disks of silicon modules each, grouped into up to seven rings. The first two layers of TIB and TOB and the rings 1 and 2 (1, 2 and 5) of the TID(TEC) contain stereo-modules: two silicon modules mounted back-to-back with a 100 mrad stereo angle to provide 2D hit resolution. The cross section of one quadrant of the silicon strip tracker is displayed in Figure 1.

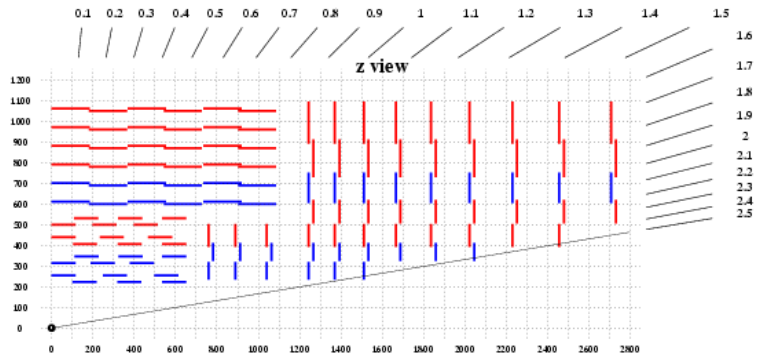

Figure 1: *rz* view drawing of one quadrant of the CMS silicon strip tracker. Single-sided modules for 1D hit reconstruction are shown in red, double-sided modules for 2D reconstruction are shown in blue.

2. Detector Control and Readout System

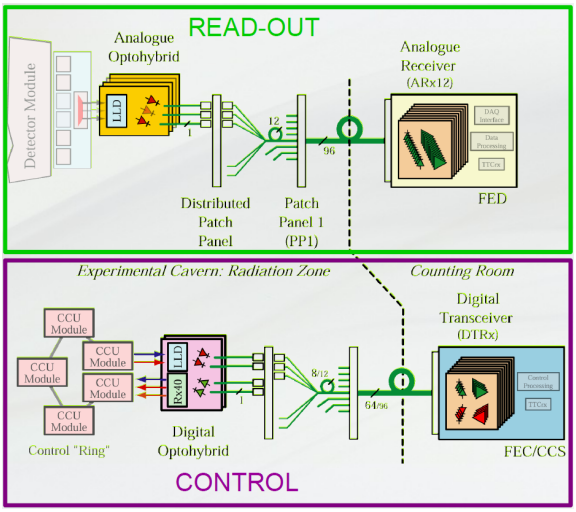

Figure 2: Readout and control schema of the CMS silicon strip tracker. Both the Front-End-Drivers (FED) and the Front-End-Controllers (FEC) are located in the service cavern about 60 m off the detector.

The CMS silicon strip tracker utilizes a fully analog readout and control schema. Both timing, trigger and control (TTC) signals to the detector, as well as data coming from the detector modules are transferred using analog optical links. The strip tracker is controlled by so-called Front End Controller (FEC) cards [2] which distribute commands to token ring networks (*control rings*). The building blocks of the control rings are Communication and Control Units (CCU) which control several detector modules each. The data from the detector modules is collected by the APV25[3] ASIC each of which reads out 128 strips at once and performs signal shaping and buffering on the detector module. Each channel of the APV25 has a 192 cell deep pipeline buffer which is used to store the data for up to 5 µs awaiting a Level 1 (L1) trigger accept signal. The APV25 has two different readout modes, the *peak* mode in which the CR-RC shaped signal with a rise time of 50 ns from the detector is directly stored in the pipeline buffers, or the *deconvolution* mode, a three sample finite impulse response filter which reduces the rise time of the signal pulse to 12.5 ns and contains the whole signal in a single bunch crossing. This pulse shortening comes at the expense of higher noise. The peak mode is robust against timing misalignments due to the broader pulse shape. It is suitable for the commissioning and low luminosity data taking. The deconvolution mode will be the default readout mode during most of the data taking. A comparison of the pulse shapes for peak and deconvolution mode is shown in Figure 3.

After a L1 accept has been received, the analog data from the respective pipeline cell are readout and passed to a multiplexer unit (APVMUX) on the module that combines the data from two APV25 chips and encodes an identifier for the pipeline address and status bits along with the data payload. This *data frame* is then transferred out from the tracker volume via a single optical fiber which is driven by a Linear Laser Driver (LLD) on the module. The laser drivers have four different gain settings which can be used to account for difference in the transmission gain of the individual links. The data are transmitted over a distance of about 100 m, out of the experimental cavern to the Front End Driver (FED) units [4] for further processing. The CMS strip tracker has

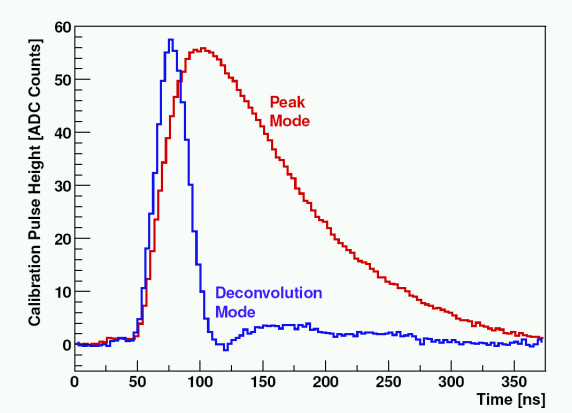

Figure 3: Pulse shapes as seen in the APV25 readout chip in peak (red) and deconvolution mode (blue).

a total of 440 FEDs which process the data from 96 fibers each. The data are digitized by a 10 bit ADC, pedestal values are subtracted and a correction for common mode noise is performed. Cluster finding is done and only strips exceeding certain defined thresholds are kept. The data volume is further reduced by only transmitting the 8 most significant data bits.

A schematic view of the strip tracker control and readout schema is displayed in Figure 2.

3. Powering and Environmental Conditions

The 15148 modules of the strip tracker are powered [5] by a total of 983 individual power supply modules. These are grouped into 29 crates and are controlled by 8 CAEN mainframes. The total power delivered to the tracker varies between 30 and almost 60 kW depending on operation conditions. The power supplies provide supply voltages of 1.25 V and 2.5 V to the modules as well as bias voltages up to 500 V. The transition from HV off to HV on can be done in 75± 5 seconds. For safety of the detector, the high voltage of the system is only raised after the LHC has declared the beam conditions to be stable.

Due to the high amount of power generated by the tracker an efficient cooling of the modules of the strip tracker is very important. It is equally important to mitigate the effects of radiation damage caused by the large particle fluences at the LHC. The detector cooling circuit operates with C_6F_{14} coolant. The cooling plant is designed for operation temperatures between room temperature and -20°C. During the 2010/11 LHC running period, the cooling plant is operated at a temperature of +4°C. To keep the humidity in the tracker volume always well below the dewpoint for any given operational temperature, the detector volume is flushed with nitrogen.

4. Detector Commissioning

To ensure the best possible performance of the detector, several calibration steps are performed prior to data taking. After communication with the detector modules is established and a cabling map is built, a coarse internal time alignment of the detector modules is performed. For this information about the length of the optical fibers and calibration pulses, so-called *tick marks* that are emitted by the APV25 chip every 70 clock cycles in the absence of a data frame are used. The

position of the rising edge of the tick mark is used to synchronize the APVs across the tracker by adjusting a *phase lock loop* (PLL) delay setting which can be adjusted in steps of 1.04 ns.

After this first time alignment, the gain settings of the laser drivers are adjusted to equalize possible differences of the optical links and the laser performance. In Figure 4, the height of the calibration pulses for all modules in the strip tracker is displayed after the calibration procedure. The target value is slightly below 700 ADC counts to make best use of the dynamic range of the optical links.

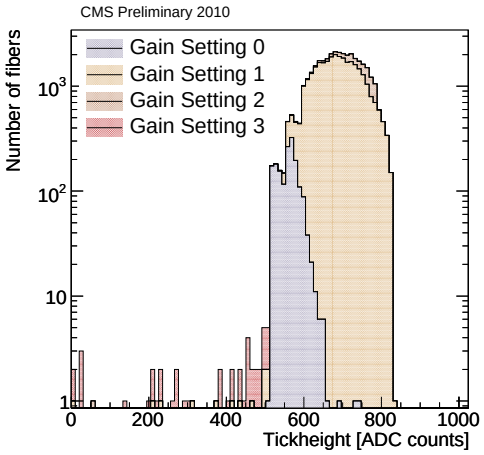

Figure 4: Calibration pulse (tick mark) height for all optical links in the strip tracker separated by the four possible gain settings of the laser driver.

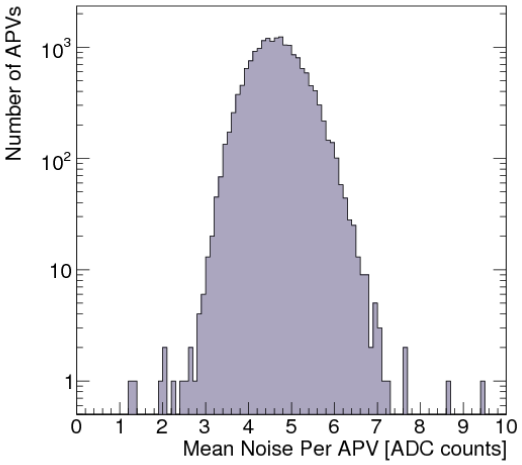

Figure 5: Distribution of mean noise per 128 strips for all active APV25 readout chips in the strip tracker.

A second, more precise time alignment procedure with the gain established, allows for a time alignment to about 1 ns precision within the partitions *TIB/TID, TOB, TEC+, TEC-*. After a tuning of several parameters of the APV25, the final step of the commissioning is the recording of pedestal and noise values for the modules. These values are stored in a database and are loaded into the FEDs at the beginning of a run to perform the zero-suppression and the cluster finding. An example of noise distribution with data from all active APV25 chips in the strip tracker is shown in Figure 5. In this distribution, the noise averaged over the 128 strips of an APV25 chip is shown.

The overall fraction of active detector channels at the beginning of the 2010 data taking period is 98.1 %. These can be subdivided into the different partitions as follows:

	Percentage	Modules total
TIB/TID	96.25	3540
TOB	98.33	5208
TEC-	99.13	3200
TEC+	98.81	3200
Strip Tracker	98.1	15148

Most of the inactive modules are due to correlated failures e.g. two control rings which are non-functional for the current data taking period.

5. Low-Level Detector Performance

After the commissioning of the detector described in the previous section, the detector is ready for data taking with CMS. To provide the optimal sampling point for the readout, a number of steps are necessary to align the timing of the strip tracker with the other subsystems of CMS. First, a coarse timing step needs to be done in order to adjust the strip tracker to read out the same bunch crossing as the rest of the detector. This step can be done using tracks from cosmic muons (see e.g. [6]). After this, the sampling point needs to be adjusted to have the optimal signal height for a maximal signal-to-noise ratio. This adjustment needs to be done with tracks from particle collisions in order to reach a precision of about 1 ns.

To precisely time align the detector, data are taken in peak mode to be largely insensitive to possible timing misalignments. One layer of the detector is then put into the deconvolution readout mode and the timing of this layer is then varied by ±10 ns around the current setting. With this, the deconvolution mode pulse shape is effectively scanned, providing information about the optimal sampling point. Two such time alignment runs were made: one using the collisions at injection energy of the LHC in December 2009 and a second one in April 2010 with collisions at 7 TeV. In the first run, the timing of only a single layer of the TOB was scanned. In the second scan, one layer(disc) of each partition was scanned. The results are shown in Figure 6. It can be seen that at the time of the first scan in 2009, the optimal sampling point of the detector was shifted by 10.1 ns with respect to the optimal point. Since the bulk of the data in 2009 were taken in the peak readout mode, this shift did not affect the data quality or signal-to-noise ratio. In 2010 the detector was continuously read out in deconvolution mode. The second scan showed that the timing of the detector was closer to ideal after the 2009 scan, however shifts for parts of the detector that had not previously been scanned can be seen. Also for the TOB that was already scanned in 2009, a small additional shift of about 2.5 ns can be seen. The reason for this shift is under investigation.

5.1 Signal-to-Noise Ratio

One of the key quantities for a good detector performance is a high signal-to-noise ratio as it is crucial to unambiguously separate signals generated by charged particles from fluctuations in the detector noise. The signal-to-noise ratio is influenced by many factors as for example the sampling point of the signal from the detector (cf. previous section), the operating temperature or the bias voltage. Due to its importance for the pattern recognition, the signal-to-noise ratio is constantly monitored during detector operation (see Section 6). In the following, two examples of

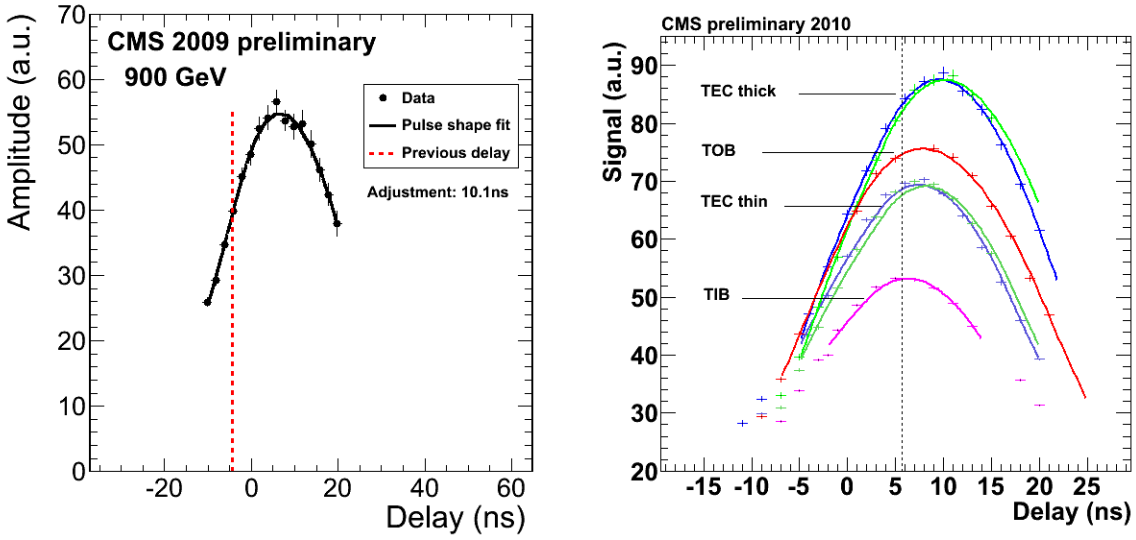

Figure 6: Timing scans performed with collisions at 900 GeV for a single layer of the outer barrel (left) and with 7 TeV collisions for the different partitions separately.

signal-to-noise ratios with different operating conditions are shown. In Figure 7 and 8, the signal-to-noise ratio of the different partitions is shown for data taken in peak mode during the 2009 LHC run. It can be seen that the signal-to-noise ratio is very high for the different partitions with most-

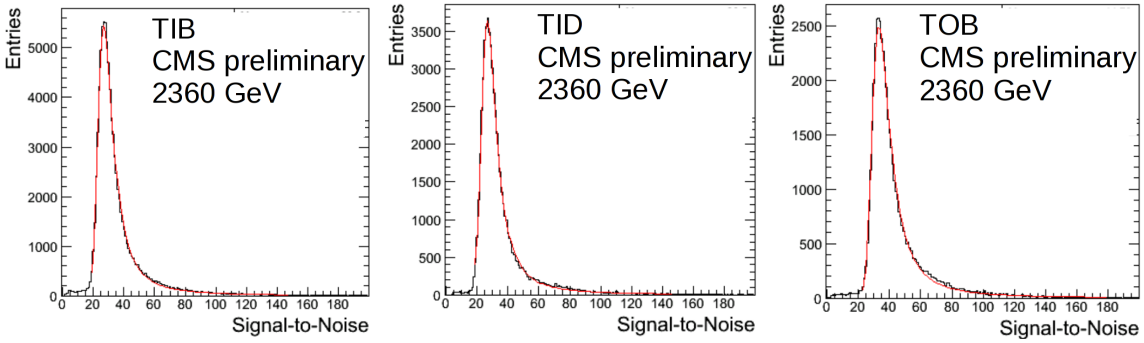

Figure 7: Signal-to-noise distribution for TIB (left), TID (middle) and TOB (right) for data taken at 2.36 TeV in peak mode.

probably value between 27 and 34 for thin and thick sensors, respectively. The signal-to-noise ratio is corrected for the difference in the path length in the silicon due to differences in the track angle.

In Figure 9 and 10, the corresponding distributions are shown for data collected during the early 2010 LHC data taking. These data were taken in deconvolution mode and prior to the second timing scan described above. The signal-to-noise ratio is lower by about 33 % with most-probably values between 19 and 26. This is in good agreement with the expected increase of about 50 % in the detector noise in deconvolution mode.

5.2 Hit Efficiency

As shown in the previous section, the signal-to-noise ratio of the modules of the strip tracker

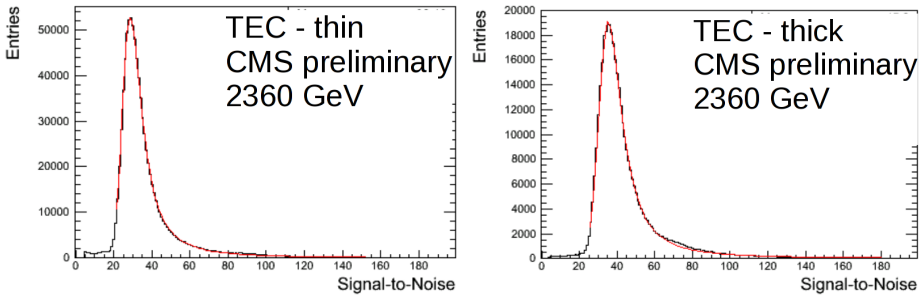

Figure 8: Signal-to-noise distribution for TEC- separated by sensor with 320 µm (left) and 500 µm (right) for data taken at 2.36 TeV in peak mode.

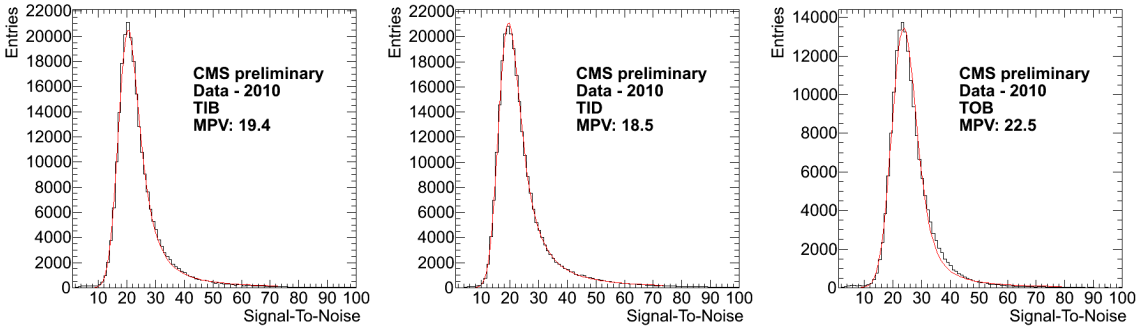

Figure 9: Signal-to-noise distribution for TIB (left), TID (middle) and TOB (right) for data taken at 7 TeV in deconvolution mode.

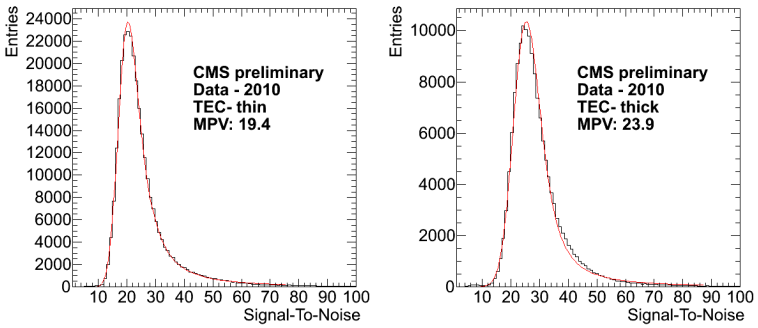

Figure 10: Signal-to-noise distribution for TEC- separated by sensor with 320 µm (left) and 500 µm (right) for data taken at 7 TeV in deconvolution mode.

is very high under different operating conditions. From this it can be expected that the probability or efficiency to find hits from the crossing of charged particles is very high. This efficiency is measured as follows: The pattern recognition is repeated excluding all hits from the module under study. For tracks passing through the module in question, it is checked if a hit consistent with the expected position from the extrapolation of the track direction is found. This procedure is done twice. In the first step, all modules of the strip tracker are used. In the second step, modules which are known to be faulty and hence are not expected to provide hit information are excluded. The results are summarized in Figure 11 where the hit finding efficiency as defined above is shown for

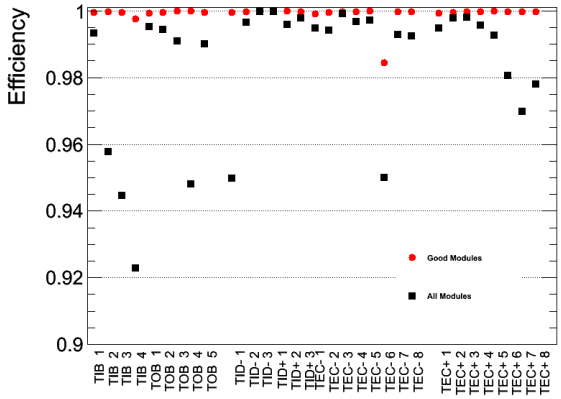

Figure 11: Hit efficiency as defined in the text for all modules in the strip tracker (black) and only considering modules that are good (red).

the various layers of the strip tracker for all modules and only for those modules that are considered as *good*. It can be seen that the hit finding efficiency is above 90 % in all cases. After the exclusion of faulty modules, the hit finding efficiency is essentially 100 % everywhere. The remaining deviations from unity come from a total of 11 modules that could be identified as faulty in addition to the previously known modules with this method.

6. Monitoring

To ensure good performance of the strip tracker and a high data quality at all times, several monitoring mechanisms are available, two of which will be described in the following. The standard *data quality monitoring* relies on normal event data while the *spy channel* has access to raw data coming from the detector.

6.1 Data Quality Monitoring

The standard CMS *data quality monitoring* (DQM) uses a common framework for all subsystems of the CMS detector. The subsystems can define quantities to be monitored, granularities and summary information to be generated. The DQM is divided into two separate instances, the so-called *online* and *offline* DQM. The online DQM uses events from the data stream of the ongoing data taking to compute the various quantities. However it only samples 10 % of the events due to bandwidth limitations. The selection of events to be used for the monitoring can be defined e.g. according to trigger requirements or other criteria. The offline DQM makes use of the entire data set of a given data taking segment. Due to the high computational requirements, it is run with a delay of about 24 hours. The monitoring information is generated by so-called *producer* applications. These process the data and generate and fill the individual monitoring elements. For the viewing the monitoring elements are retrieved by a *consumer* application which can generate summary information and forwards the information to a web interface for user interaction. The consumer application can also display the results from automated quality tests which can alarm the user in case of deviation of distributions from the expected behavior. The layout of the DQM process is displayed in Figure 12.

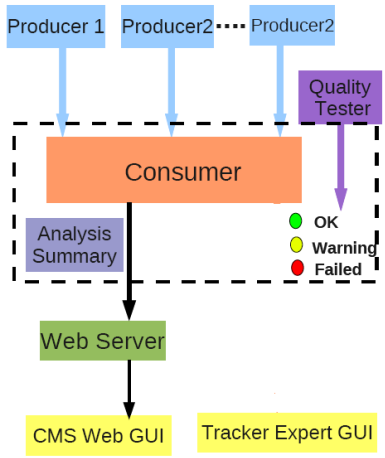

Figure 12: Schematic layout of the data quality monitoring used in the CMS strip tracker.

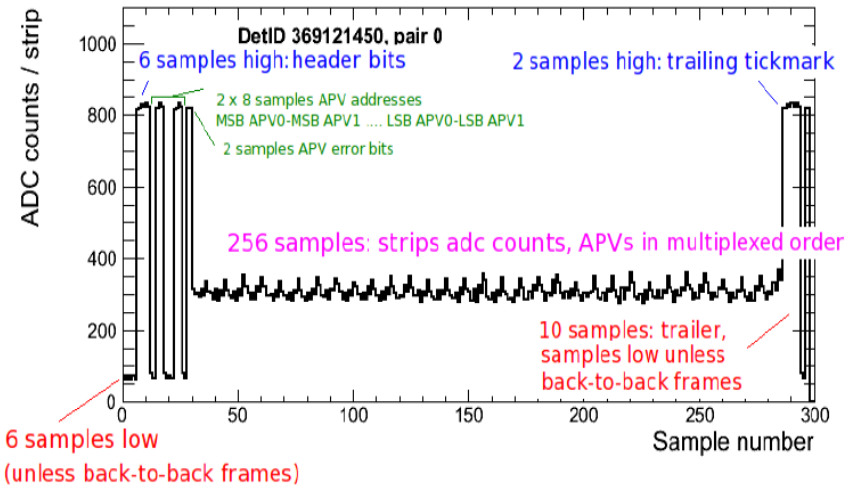

Figure 13: Full data frame from two APVs as seen by the FEDs and recorded with the spy channel

Both DQM instances use the same reconstruction framework which is the standard CMS software framework known as CMSSW. This makes it possible to use the same algorithms and spot features that would be visible in the same way in an offline analysis of the data.

6.2 Spy Channel

The so-call *spy channel* is a separate, non-zero-suppressed data stream in parallel to the normal physics data stream. The spy mechanism enables the FEDs to capture data frames coming from the detector modules at a low rate (about 0.3 Hz). These frames are read out via VME bus and then stored to a local disk for further processing. An example of a spy frame as seen by the spy channel readout can be seen in Figure 13. The data taken can be chosen to be either associated to physics triggers or in the absence of particle collisions e.g. in the abort gap.

With the spy channel all information coming from the detector modules is available for processing and monitoring. The spy channel can be used for determination of calibration constants for

pedestals and noise during running. With this, separate calibration runs can be made obsolete.

The gain and the dynamic range of the system can be monitored, ensuring that there is no time variation of these quantities which would otherwise not be visible in the normal data stream directly.

7. Conclusions and Outlook

The operation and performance of the CMS silicon strip tracker have been described. For the 2010/11 LHC data taking, a total of 98.1 % of the strip tracker is operational. Numerous commissioning steps have been performed to ensure best possible performance of the strip tracker. Scans have been performed to achieve best possible synchronization of the strip tracker with the other CMS subsystems. The relative timing misalignment could be determined to about 1 ns. The signal-to-noise ratio is very high and in agreement with expectations under different operation conditions. The single hit efficiency is close to 100 % and components with sub-standard performance can be reliably identified. A multi-stage data quality monitoring is in place, providing rapid feedback to the operation. The spy channel is a powerful tool for the calibration and monitoring of many key quantities for the operation of the tracker in the coming months and years.

References

[1] R. Adolphi *et al.* [CMS Collaboration], *"The CMS experiment at the CERN LHC,"* JINST **3** (2008) S08004.

[2] F. Drouhin *et al.*, *"The control system for the CMS tracker front end,"* IEEE Trans. Nucl. Sci. **49** (2002) 846.

[3] L. L. Jones *et al.*, *"The APV25 deep submicron readout chip for CMS detectors," Prepared for 5th Workshop on Electronics for the LHC Experiments (LEB 99), Snowmass, Colorado, 20-24 Sep 1999*

[4] C. Foudas *et al.*, *"The CMS tracker readout front end driver,"* IEEE Trans. Nucl. Sci. **52** (2005) 2836 [arXiv:physics/0510229].

[5] S. Paoletti *et al.*, *"The powering scheme of the CMS silicon strip tracker",* in 10th Workshop on Electronics for LHC and future experiments, CERN-2004-010, CERN-LHCC-2004-030, http://cdsweb.cern.ch/record/814088.

[6] S. Chatrchyan *et al.* [CMS Collaboration], *"Commissioning and Performance of the CMS Silicon Strip Tracker with Cosmic Ray Muons,"* JINST **5** (2010) T03008 [arXiv:0911.4996 [physics.ins-det]].