Studies of the Assembled CMS Tracker

P. Everaerts[a], K. Hahn[a], S. Nahn[a]

[a] Massachusetts Institute of Technology, 77 Mass. Ave, Cambridge MA 02139, USA

peveraer@mit.edu

Abstract

During the latter months of 2006 and the first half of 2007, the CMS Tracker was assembled and operated at the Tracker Integration Facility in Building 186 at CERN. At this time, several dedicated studies were carried out to validate the performance of the tracker after assembly, testing general noise performance, looking at a specific problem showing up for part of the tracker [1], and also looking at the performance at high acquisition rates [2]. We report on the the results of these studies and their consequences for operation of the Tracker at the experiment.

I. The CMS Silicon Strip Tracker.

With its 210 m^2 of silicon, 5.4 m length, 2.4 m diameter and 9.6 million readout channels, the CMS strip tracker is clearly the largest and most complicated silicon detector ever built. It consists of 4 main parts: the endcaps (TEC), the inner barrel (TIB), the outer Barrel (TOB) and the inner disks (TID). All together 15148 modules are distributed amongst these 4 systems. Because of its size and complexity, the collaboration paid meticulous attention to quality control and testing all the way through construction. However, some effects could not be detected during the construction and the final assembly of the subdetectors and the first large-scale tests were needed to point them out. This paper discusses the investigations into the cause of these effects as well as the ramifications for operations at the LHC.

A. TIF and Point 5

The first large scale tests were performed in the Tracker Integration Facility (TIF), a clean room constructed at the Meyrin site specifically for this purpose: the different subdetectors were brought together there and assembled from October 2006 until March 2007. Between March and July 2007, cosmics were taken at the TIF in a 'Sector Test'. All together 4.7 million cosmic triggers were recorded and up to a quarter of the tracker was read out using the final data acquisition (DAQ) electronics. During these tests, the tracker was operated at a range of temperatures, ranging from +15 to -15 degrees Celsius. The results of charge collection [3], track reconstruction [4] and alignment studies [5] agreed well with those expected from the construction studies.

The tracker was then inserted at Point 5 in the CMS Cavern in December 2007. The connection work took from December 2007 until March 2008 and after a short delay due to cooling problems, the tracker was commissioned between half June and end of August 2008. The TIF and the installation and commissioning at Point 5 provided the first possibilities for performing large-scale tests.

B. DAQ chain

To appreciate fully the effects discussed requires a rough idea of the CMS tracker Data Acquisition scheme. Figure 1 depicts the control flow to the front end ASICs, and the data flow back back into the electronics for processing and storage. The Frond End Controller (FEC) VME board sends trigger and clock information to the Digital OptoHybrid Module (DOHM), which performs the optical to electrical conversion and forwards the electrical signals to the Control and Communication Unit (CCU), which distributes them to all front-end ASICs. The front-end ASIC (implemented as the APV chip) samples the strip charge, does the analogue pulse shaping, stores the data locally, and upon request transfers them optically to the Front-End-Drivers (FED) VME board. There the data is digitized and under default operation clusters are formed and zero-suppression applied, dropping clusters below a preset threshold. Finally the FED pushes the data out to the Central CMS DAQ.

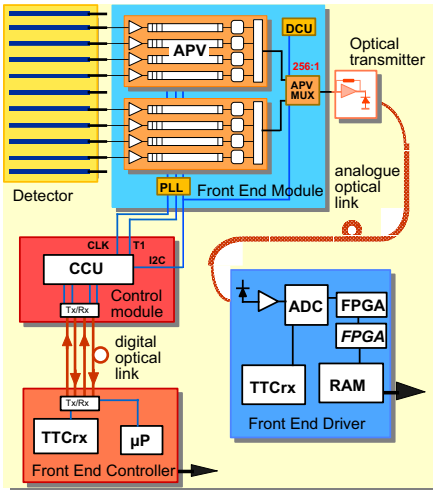

Figure 1: Schematic view of the CMS electronics.

II. Wing noise [1]

The Tracker Outer Barrel (TOB) is made out of 688 "rods". Each rod can contain up to 12 modules, and each module consists of 2 silicon sensors and 4 or 6 APV chips on a hybrid which services the sensors. A schematic set-up of a rod (without the sensors installed) can be seen in fig. 2. Cooling, power, optical readout and control signals all enter the rod from one end. Power and electrical signals are sent over a 51 mm wide circuit board called the interconnect bus, a multilayer board that is

symmetric between the top half and bottom half.

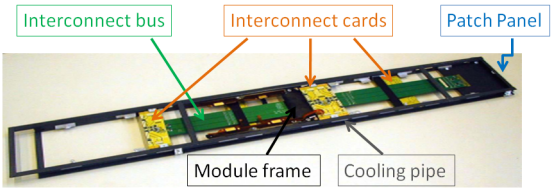

Figure 2: Schematic view of a TOB rod

The noise (RMS) distribution of such a TOB rod shown in figure 3 looks rather strange. This normalized noise distribution from module 6 of a TOB rod shows a dramatic increase towards the edge channels of the APV chips, which resemble wings, from which the effect was named **wing noise**. These peaks reach up to 40% of the average noise value of the module depicted, and can be even more pronounced on other modules. This effect escaped detection during module testing, and was discovered after the assembly of the rods. The observation that the effect is most prominent on module 6 for rods of layers 3 and 4 indicated that this position was special for those particular type of rods. The influence of the extra noise is so big that it would even affect tracking, therefore a solution needed to be found.

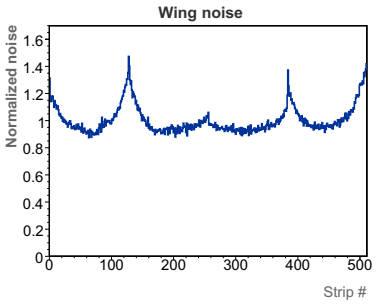

Figure 3: Wing noise on a TOB module

A. *Investigation*

To investigate the origin of the noise, shielding was applied to the different components of the rod. A testbench rod was used for the first investigations because it allowed more flexibility. The rod power bus was considered a first good candidate for sourcing the noise, because it runs under every sensor. Placing a sheet of copper-clad Kapton between the sensor and the interconnect bus eliminated all the noise, this happened independent of whether this sheet was grounded or not. This pointed in the direction of magnetic coupling, because even an ungrounded sheet would then remove all noise. For magnetic coupling, the copper generates eddy currents that would counteract the external field and thus block the external interference. For electrical coupling an ungrounded sheet of copper would act as a floating capacitor and continue to couple the noise into the sensor. In this case it could only reduce the noise, but never fully eliminate it. Similar sheets positioned between the sensor and hybrid or the fiber frame did not show any influence.

Having understood how the noise entered the sensors, the question turned to what the source was. After some investigation, the wing noise source was discovered to be the connection between the CCU and the DOHM on the rod. The noise current runs on top and bottom of the interconnect bus and is returned through the common ground in the middle. The problem lies with the DOHM differential signals to the CCU. These differential electrical signals are never fully balanced, and the imbalance current has to be returned to the DOHM via the ground of the interconnect bus. The excess current then sources the magnetic coupling. The current path for the noise current is from the CCU over the control circuit board to the power bus, through the power cables, supply rack and then the control power cables to the DOHM. This hypothesis was verified by adding a piece of copper tape between the DOHM and the circuit board, providing a direct return path for the current, which caused the noise to disappear. In addition, this was checked by putting ferrites around the power cables, increasing the impedance and thus reducing the noise current, which was observed. In performing this test it was realized that also the cooling pipes can provide an adequate path for this return current.

The most likely coupling is through the field component parallel to the sensor. The loop which creates the signal voltage is the following: aluminum sense lead on the top, the high voltage bias plane on bottom and the capacitance of reverse biased diodes at the two sides. One way of investigating this is through the bias voltage. Under this hypothesis increasing the bias voltage will decrease the capacitance and increase the impedance of the loop, thus reducing the current and for the APV (a charge sensitive amplifier) also lowering the noise signal. Thus the relative size of the wings should decrease when applying bias voltage, as is confirmed by data. The composition of the induction loop also explains why the noise is most predominantly present in layers 3 and 4. These layers have the largest width of strips, which will give the largest capacitance, thus resulting in most pronounced effect, as observed.

B. *Solution and implementation.*

The solution for this problem is to establish a common high-quality ground for the power cables, the control power cables and the cooling pipes. This was done for the tracker on the stainless steel cryostat of the CMS magnet. Small daughtercards with capacitors filtering external noise were designed that grounded everything together at the daughterboards (fig. 4), which were mounted on and grounded to the cryostat. This scheme was first tested at the TIF, and then in the CMS cavern at Point 5 by utilizing the Rod-In-a-Box (RIB). The RIB was the first piece of the final tracker installed at Point 5, brought there months before the tracker as a proxy to test and modify as appropriate the cabling and connection scheme for the tracker. The RIB showed that the daughter card scheme worked: the wing noise was absent initially, with the removal of the daughtercard, it reappeared.

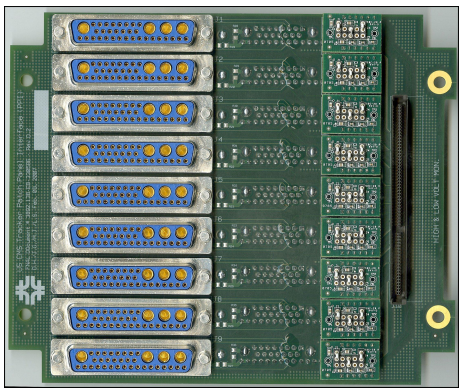

Figure 4: Daughterboard with small daughtercards visible at right

The final results for the TOB noise distribution from the commissioning are shown in fig. 5. At the left the results of the TIF (without daughtercards) are shown [3], and at the right the results of the commissioning are shown, after the scheme was implemented. The shoulder at high noise drops with **two orders of magnitude**. Currently there are only 9 modules on 7 rods (out of 688 rods) that still show some wing noise. This remaining 1% of problems is most likely due to mechanical problems with those specific daughtercards or the grounding of the cables. This was already discovered on 8 other modules that were recovered after investigating the connections. When comparing these two plots one has to be aware that the commissioning results have been normalized, and the bump at low noise values is due to dead strips in the tracker.

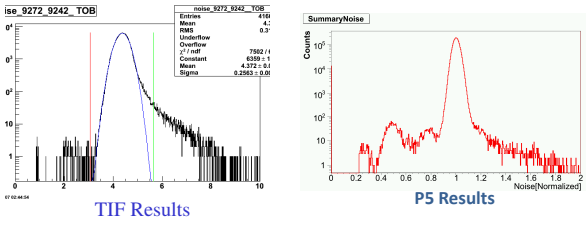

Figure 5: Results at TIF without (left) and at Point 5 with daughtercards (right).

III. High-Rate Noise [2]

At the TIF cosmics were used to sample detector performance, but these cosmics were at low trigger rate, of order 10 Hz, while the collisions in CMS will take place at high rate, up to 100 kHz. Therefore a dedicated study was performed to test the behavior of the system at high trigger rate. For this task, the data-acquisition system was augmented, replacing the FED VME readout which cannot sustain high trigger rates with a small scale slice of the final CMS DAQ called a "Column-DAQ". This Column-DAQ required hardware both for readout to acquire rapidly and trigger control to avoid buffer overflow. Fig. 6 gives an idea of this set-up and can also be used as a schematic for the CMS DAQ. As there was no physics signal which would give such high trigger rates, just noise and pedestals were measured using a random high rate trigger.

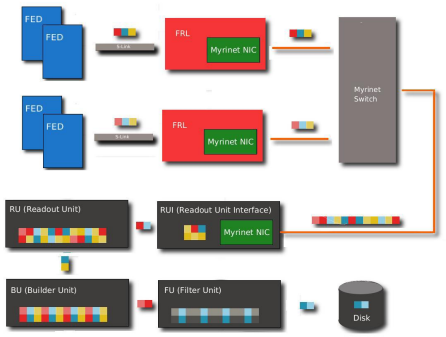

Figure 6: Schematic view of our Column DAQ

A. Column DAQ

Since the final CMS DAQ was not available at the TIF, a special dedicated column-DAQ was built up. Rather than VME readout, data proceeds from the FED through Hardware using S-Link and dedicated network links to a switch and into a small scale computer farm emulating the CMS trigger farm, where it was written to disk. Even with these modifications, the readout was limited to at best 10 kHz. This necessitated further measures to increase the performance of the system. These steps included removing a software bottleneck in writing data to disk, adding an additional rail of network from the FEDs through the switch to the PCs, and the implementation of prescaling at various points in the system, randomly dropping events to relieve bandwidth pressure. With these modifications the 100 kHz level could be reached, but only taking data in Zero-Suppression mode, thus losing all information about channels below threshold. Looking at the occupancies in different channels did point out the effect described below, but not until another prescale was implemented, this time at the FED, the very beginning of the chain in fig. 6, could the system take data from all channels in an event ('Virgin Raw' readout mode). This turned out to be a critical modification to provide a more thorough investigation of the phenomenon.

B. High-rate noise characteristics.

When investigating noise behavior versus trigger rate, nothing special was observed at low rates. The behavior of the pedestals and noise at 100 Hz and 3 kHz is consistent. On the other hand, at rates above 30 kHz, a considerable growth in average occupancy took place at the edge channels of the chip (channels 127, 255,...), as shown in figure 7.

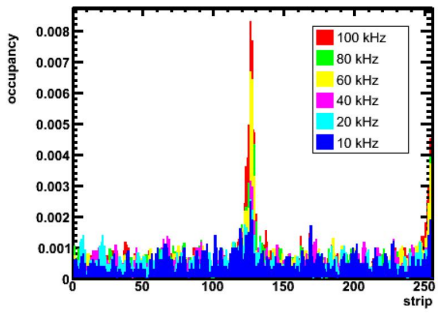

Figure 7: Increase of occupancy at high trigger rates.

Testing on many different types of modules proved that this phenomenon was not isolated to a few modules, but affected every single APV of the tracker simultaneously and thus it was a global problem in the tracker running at High Rate, and not due to some faulty components. Fig. 8 illustrates this by showing two different fibers on the same FED and a fiber on a different FED (thus independent) all showing increased ADC counts near the edge channels on the same particular event. Once the full data was available, the common-mode subtracted pedestals could be calculated, and an anti-correlation between the behavior on strip 0 and that on strip 127 was observed and noted as a distinct signature of the high-rate noise. Another distinctive feature of the effect was that it only occurred in randomly distributed trigger intervals, not in fixed (high) rate trigger situations.

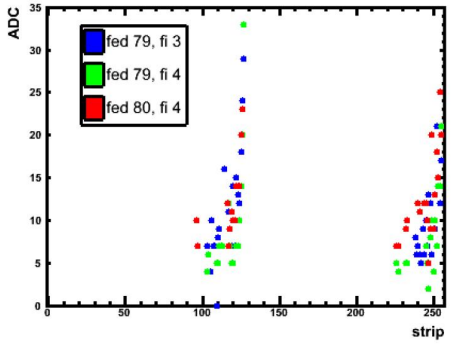

Figure 8: Simultaneous rise on different APVs.

The first attempt to narrow down the source of the noise was to try to shield the sensors, motivated by the wing noise study. However, there was no reduction in the effect. Due to the fixed vs. random trigger effect, the trigger hardware was investigated to ensure it was obeying "trigger rules", not sending commands at the wrong time; no violations of the rules were found. Alternatively, the same effect suggests that perhaps the random trigger scheme was sampling some part of phase space that fixed triggers did not, such as some problem in the channel pipelines in the APV used to store the raw data before readout. By controlling exactly when triggers were the sent the pipeline position dependence was studied and cleared from any blame. However, in the same study the observation was made that independent of the pipeline, certain intervals between two triggers yielded spikes of maximum occupancy as can be seen in fig. 9. These jumps happened for all fibers simultaneously around the magic bunch-crossing numbers 100, 160 and 380, and near these values of trigger interval the noise increases and when it reaches one of these numbers, the rise is the most pronounced.

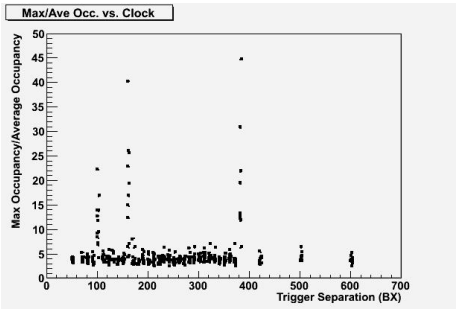

Figure 9: Maximum occupancy versus trigger interval.

C. Explanation

With the knowledge of particular intervals which trigger the effect, a deeper investigation using an oscilliscope to probe both the trigger commands and optical APV data when the effect occurs provided the final explanation for the phenomenon. The effect is caused by cross-talk between the data-acquisition and the read-out. Fig. 10 gives a schematic timeline of readout and acquisition. When a trigger arrives, the data which was buffered in the APV one latency (a programmable time delay) before the trigger arrives is flagged for readout. At the end of the APV's 7 μs readout cycle , this data is then pushed from the APV upstream to the FEDs. This readout takes a full APV cycle, as is shown in the top half of fig. 10. Just before the APV starts sending data, it also draws more current, which can now affect the data-acquisition or the buffering of new data. Thus in case the trigger then also flags this data for read-out as done in the bottom half of the drawing, then the high-rate noise appears, as depicted in the bottom half of fig. 10. In reality, there are even three current rises during the APVs 7 μs read-out cycle, leading to three magic numbers of 100, 160 and 380 bunch-crossings (1 bunch-crossing=25 ns at LHC). Thus the effect is an interplay between two tasks of the APV: the buffering and the sending of the data up-stream. This explanation has been confirmed by independent bench tests at Imperial College and simulations from Rutherford Laboratory.

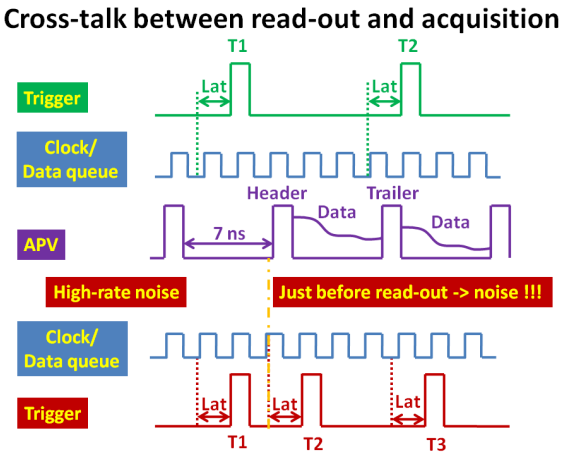

Figure 10: Schematic overview of 'high rate noise'-creation.

D. Possible solutions

There are several ways to mitigate the high-rate noise which are under consideration. The high-rate noise could be identified offline and the corresponding data flagged as corrupted and thus thrown away, leading to a 1 % loss of data, and therefore not favored. A second option would be to change the trigger rules, so that the 'magic' trigger intervals and the surrounding bunchrossings are prohibited. This would lead to considerable deadtime, especially because intervals are hard to define in trigger rules, usually only lower time bounds are used. It would also be possible to deal with the problem further down the DAQ chain and hardware veto the triggers that arrive at the moment the APV is reading out data. This option still generates dead-time but a lot less than the trigger rules, the disadvantage is that it is difficult to implement in the firmware. Finally the anti-correlation observed in the common-mode subtracted pedestal could be used as a signal of the High Rate effect. The values of strip 0 and 127 are of similar magnitude but with an opposite sign in case of high-rate noise. This could be used to reject these events, but that is fairly risky since the algorithm is not optimized yet and it is not known how much good data would be thrown away. A variant on this option would be to flag these events and then deal with them offline. This would also require a firmware change to provide an extra flagging bit in the FED header.

IV. Commissioning results

The overall results during the tracker commissioning are extremely good. The noise problems are under the percent level. In the end only 71 modules out of 15148 (0.5%) show a non-standard noise distribution. 9 of these are wing noise problems, but there are also a few new phenomena that showed up, like the mysterious noise (25 modules), figure 11a) and the MUX noise (18 modules), figure 11 b). The MUX noise owes it name to the spikes occurring every 16 strips, indicating a multiplexing problem on the concerned strings. Both of these new noise problems are correlated with bad supply of power to the modules or a bad control signal. This was noticed by looking at the voltage values at the modules itself, which turned out to be low for several of them and also by looking at odd synchronisation results indicating a problem with the control signals.

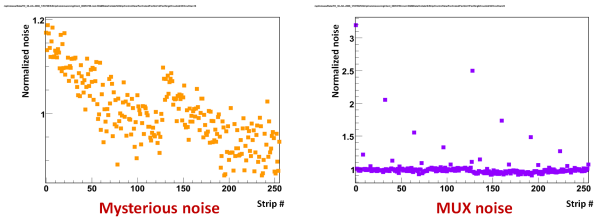

Figure 11: a. Mysterious noise, b. MUX Noise

V. Conclusions

In general the CMS tracker performs extremely well, but there were a few small hiccups when the first large-scale tests were performed. Two new effects were found, the high-rate noise and the wing noise. A dedicated investigation at the TIF lead to the discovery of the origin in both cases and for the wing noise a solution has already been implemented. For the high-rate noise multiple options are still considered. These two cases showed the merit of the TIF as a testbench, as both problems were discovered there at an early stage, long before the tracker was moved to Point 5. This allowed for adequate measures to be taken in time. The daughtercard scheme was implemented before tracker arrival, and the high rate mitigation will take place well before CMS expects to reach high trigger rates. As commissioning ensues, new noise phenomena are showing up on a small scale, and the CMS tracker community highly anticipates what the first pp collisions may bring.

Acknowledgments

The authors would like to thank especially everyone who worked at wing noise and the high-rate noise puzzles as well as the rest of the CMS tracker community.

References

[1] F. Arteche *et al.*, Noise in the CMS Tracker, in preparation.

[2] K. Hahn *et al.*, Studies of the CMS Tracker at High Trigger Rate, in preparation.

[3] CMS Tracker Collaboration, L. Demaria and P. Azzi, Tracker Detector Performance and Simulation Tuning with Cosmic Ray Data at the Tracker Integration Facility, *CMS Note* in preparation

[4] CMS Tracker Collaboration, C. Noeding and W. Adam, Tracker Reconstruction with Cosmic Ray Data at the Tracker Integration Facility, *CMS Note* in preparation

[5] CMS Tracker Collaboration, A. Gritsan and G. Flucke, CMS Tracker Alignment at the Tracker Integration Facility, *CMS Note* in preparation