Jinst

PUBLISHED BY IOP PUBLISHING FOR SISSA

RECEIVED: *June 10, 2010*
ACCEPTED: *July 8, 2010*
PUBLISHED: *July 30, 2010*

Studies of the CMS tracker at high trigger rate

M. Chan,[a] N. Cripps,[b] P. Everaerts,[a] J. Fulcher,[b] K. Hahn,[a] G. Hall,[b] M. Johnson,[c] Q. Morrissey,[d] S. Nahn,[a,1] M. Noy,[b,2] M. Raymond,[d] K. Sung[a] and S. Tkaczyk[c]

[a] *Laboratory for Nuclear Science, Massachusetts Institute of Technology, Cambridge, MA, U.S.A.*
[b] *Blackett Laboratory, Imperial College, London, England, U.K.*
[c] *Fermi National Accelerator Laboratory, Batavia IL, U.S.A.*
[d] *Rutherford Appleton Laboratory, Didcot, Oxfordshire, England, U.K.*

E-mail: nahn@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 at CERN. During this period the performance of the tracker at trigger rates up to 100 kHz was assessed, and a source of high occupancy events was uncovered, diagnosed, and mitigated.

KEYWORDS: Analogue electronic circuits; Front-end electronics for detector readout; Particle tracking detectors (Solid-state detectors); Large detector systems for particle and astroparticle physics

[1] Corresponding author.
[2] Now at CERN, Geneva, Switzerland.

© 2010 IOP Publishing Ltd and SISSA
doi:10.1088/1748-0221/5/07/P07007

2010 JINST 5 P07007

Contents

1 Introduction **1**
- 1.1 The CMS strip tracker 1
- 1.2 High rate test apparatus 2

2 Tests with the tracker at high rate **2**
- 2.1 Discovery of the high rate Noise effect 2
- 2.2 Investigation of the APV25 4

3 Consequences and potential mitigation **11**

4 Conclusions **12**

1 Introduction

The CMS Tracker [1] is the largest silicon system of its kind ever built. The apparatus was assembled and tested at the Tracker Integration Facility (TIF) at CERN *before* installation at the experimental site. In addition to testing charge collection [2], track reconstruction [3], and alignment [4] performance with cosmic rays, a small slice of the CMS data acquisition (DAQ) hardware was built at the TIF to provide the bandwidth capacity to allow testing at the high acquisition rates expected in proton-proton collisions, of order 50-100 kHz, four orders of magnitude above the rate in cosmic ray studies.

1.1 The CMS strip tracker

The CMS Tracking system is composed of an inner pixel detector as well as an outer tracker based on silicon microstrips. The tracker is further divided up into subdetectors: the inner barrel and inner disks (TIB/TID), the outer barrel (TOB), and the end caps (TEC), all together roughly 10 million strips. Each subdetector is composed of silicon sensors capacitively coupled to front-end electronics in ”modules”. Modules are combined into TIB/TID ”strings”, TOB ”rods”, and TEC ”petals”, which share the same power and readout services. The Tracker Technical Design Report [5, 6] provides more descriptions.

A module measures ionization depositions from charged particles traversing the material. The charge information from the microstrips is sampled every 25 ns and stored in an analogue pipeline by an on-detector ASIC, the APV25 [7]. The data are held pending a decision to read out a particular 25 ns sample, a “trigger”, following which the APV25 readout process starts, converting the analog charge into an optical signal, adding header information, and transmitting the optical signal to a FrontEnd Driver (FED), a VME board which digitizes the data, performs pedestal subtraction and zero suppression, and then passes the data further downstream to the CMS DAQ.

2010 JINST 5 P07007

1.2 High rate test apparatus

The data acquisition hardware at the TIF [8] is a smaller scale version of the CMS DAQ installed at the experiment, which handles data for all CMS subdetectors. There are two interfaces to the CMS DAQ: the data path and control path. The data path starts with FED event fragments, which are routed through a switch to a computer farm responsible for concatenating the fragments into one event, unpacking the payload, and running filtering algorithms before writing the event to disk. The control path is a fan-in monitoring buffers in the FEDs and APV25s to moderate the trigger rate, preventing buffer overflow and subsequent loss of data. Testing at high rate necessitated the implementation of both of these paths. In addition, a Local Trigger Controller VME module generated multi-kHz triggers with programmable time structure, including either Poisson distributed or fixed-frequency triggers. The triggering scheme does not involve the presence of real signals; thus, the data collected represents the noise behavior of the tracker at high rate.

Three different configurations of modules were utilized to study the tracker performance: the "Slice Test", a single-rod testbench, and a single-module testbench. The Slice Test [2–4] consisted of approximately 15% of the full tracker, fully instrumented with power supplies and readout electronics, and operated with either scintillator-based cosmic ray triggers or programmable calibration triggers. As the investigation indicated a systematic effect not related to the scale of system, a spare TOB rod also at the TIF substituted in place of the actual detector, permitting high rate studies in parallel with other programs during TIF operations. This rod was composed of 6 modules with 4 APV25s per module, connected to the DAQ in exactly the same manner as the Slice test. Finally, a second completely different DAQ system was used to probe the APV25 chip and sensor behavior in a more controlled way on the benchtop. This system employed a programmable digital pattern generator to provide the 40 MHz clock and trigger patterns to a single TEC module, and a commercial ADC to digitize the output data streams.

2 Tests with the tracker at high rate

2.1 Discovery of the high rate Noise effect

With the high rate apparatus installed, tests of the noise performance were carried out. First, using the TOB, noise and pedestals were sampled using so-called "Virgin Raw" (VR) mode, where the FED applies no pedestal or common-mode subtraction nor any zero suppression of the data. The noise and pedestal results are used for FED data processing in subsequent runs, so these tests checked that there were no variations due to readout path or readout rate. With rate-independence of the pedestals and noise established, zero-suppression (ZS) was enabled in the FEDs to achieve smaller event size, required for higher rates. In ZS-mode, the FEDs subtract pedestals from the raw ADC data in each event and calculate per-APV25 common-mode levels from the result [9]. After processing, the algorithm forms clusters by requiring either a single strip with an ADC count above 5σ or two or more adjacent strips with ADC counts above 2σ where σ indicates the RMS of the pedestal for that strip. The FED outputs the corrected ADC data of strips in such clusters, and the cluster size and position. This information is used to calculate strip occupancy, defined as the frequency at which a given strip is included in a FED cluster.

2010 JINST 5 P07007

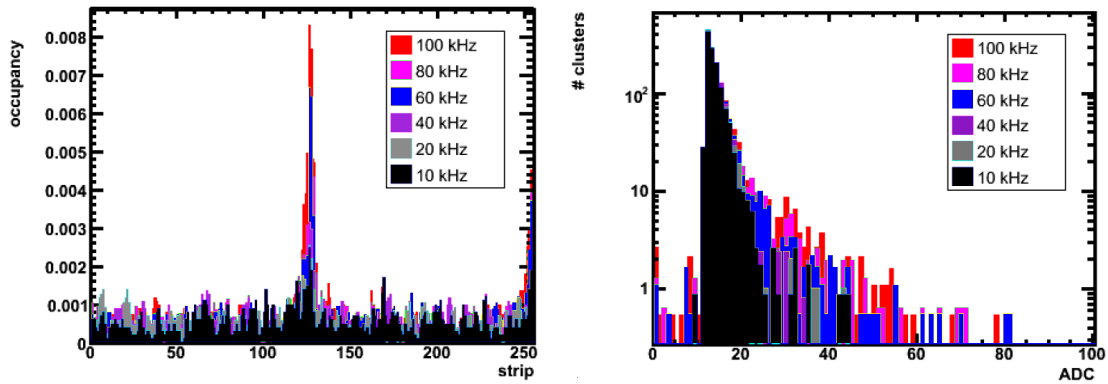

(a) (b)

Figure 1. (a) Occupancy and (b) Cluster charge for a single TEC fiber versus L1A Rate. A clear increase in occupancy with trigger rate is observed again on channels near 127 for both APV25s on the fiber. In addition, the tail of the Cluster charge distribution grows with rate.

The ZS tests revealed an increase in occupancy of edge channels of the APV25 chips with increasing trigger rate, which becomes evident at rates above ~ 30 kHz, both with biased and unbiased sensors. Tests with the TEC and TIB established this high-rate noise effect (HRN) as a universal feature in the tracker. An example of the growth in occupancy with increasing rate from one TEC fiber, which carries data from two APV25s, is shown in figure 1a. In addition to increased occupancy, the edge channel cluster sizes and ADC values also increase with trigger rate, as seen in figure 1b for the same TEC fiber. For reference, a minimum ionizing particle is expected to deposit charge clusters around 100 ADC counts.

Selecting events with large ADC counts in channel 127 in one APV25, and examining the channel 127 output from other APV25s in the same event shows that large clusters comprised of strips with high ADC count appear *simultaneously on every APV25* when HRN occurs, indicating the effect is independent of particular detector, and is triggered by some global signal or that this is a systematic effect from the APV25 or FED which affects all detectors.

To obtain more information on HRN events, a modification was made to the FED firmware to sample VR data while running at 100 kHz, recording ADC values on all strips during high rate operation. Figure 2 plots ADC values of channel 0 and 127 for events in which the channel 127 passes the cluster criteria. The mean of the channel 0 data is ≈ -20, well below the chip average, indicating a strong anti-correlation between the ADC counts of channels 0 and 127. To confirm that anti-correlation, events where either channel has a high ADC count are selected. Figure 2 shows the ADC correlation at different trigger rates, which becomes more pronounced as trigger rate increases.

The HRN effect was uncovered using high-rate Poisson-distributed triggers which closely approximate the triggering scenario expected in CMS. Tests with a fixed-frequency 100 kHz trigger *did not* reproduce the occupancy spikes observed with the Poisson trigger. Under these conditions, the readout of each trigger, which requires 7 μs, is complete before the next trigger is received. The trigger generator was verified not to produce illegal trigger conditions and the time distribution was verified to be Poissonian. It was concluded that the HRN originated in an unexplained interference when data readout coincided with a subsequent trigger.

2010 JINST 5 P07007

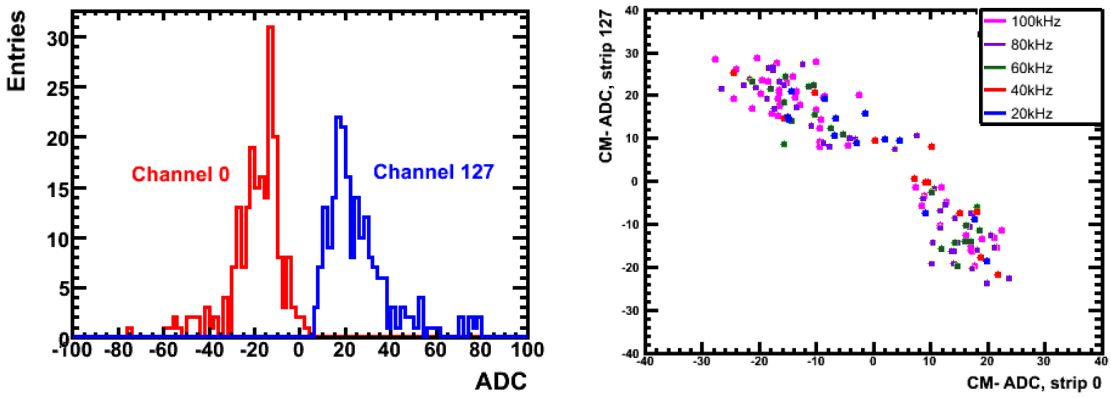

(a) (b)

Figure 2. (a) Common-mode subtracted ADC counts for channel 0 and 127 when channel 127 has a cluster over threshold. The offset from zero for channel 0 indicates a strong anti-correlation between the HRN on these two channels. (b) Average ADC count of channel 0 versus that of channel 127, when one of them is above threshold, at various different trigger rates, confirming the anti-correlation and its dependence on trigger rate.

2.2 Investigation of the APV25

The observation that the effect is independent of module type indicates that the effect arises from a common component, possibly the APV25 chip. To understand further, it is necessary to review how the APV25 functions.

APV25 operation. The APV25 contains a 192×128 cell pipeline that holds charges read from each of 128 silicon microstrips. At the start of a data run, following a master reset the internal logic is initialized and the digital pointers controlling pipeline access start to circulate. A write pointer shifts through the pipeline controlling sampling of the front end amplifier output at 40 MHz. A trigger pointer follows with a time delay equal to the programmed latency. When a trigger occurs one (three) pipeline cells in peak (deconvolution) mode corresponding to the current trigger pointer location are marked for subsequent readout, and are not overwritten until the readout process has completed. The pipeline readout is governed by a separate cycle with a period of 1.75 μs (70 clocks at 40 MHz), and a phase also determined by the master reset signal. The phase of this pipeline readout cycle is reflected in the output data stream as tick marks, large amplitude signal levels occupying one 25 ns clock. The tick marks allow external logic to synchronize to the APV25 output phase.

The APV25 has two modes of operation, peak and deconvolution. In peak mode one sample per channel is read from the pipeline following a trigger, and then transferred to the output via the multiplexer; the sample corresponds to the maximum amplitude from the CR-RC shaped front end amplifier, which has a time constant of 50 ns. In deconvolution mode [10, 11] three samples are read sequentially and a weighted sum formed. This results in an effective re-shaping of the analogue pulse shape to peak at 25 ns and return to the baseline within one clock cycle. These operations take 4 pipeline readout cycles, after which transmission of the output data frame commences. For

2010 JINST 5 P07007

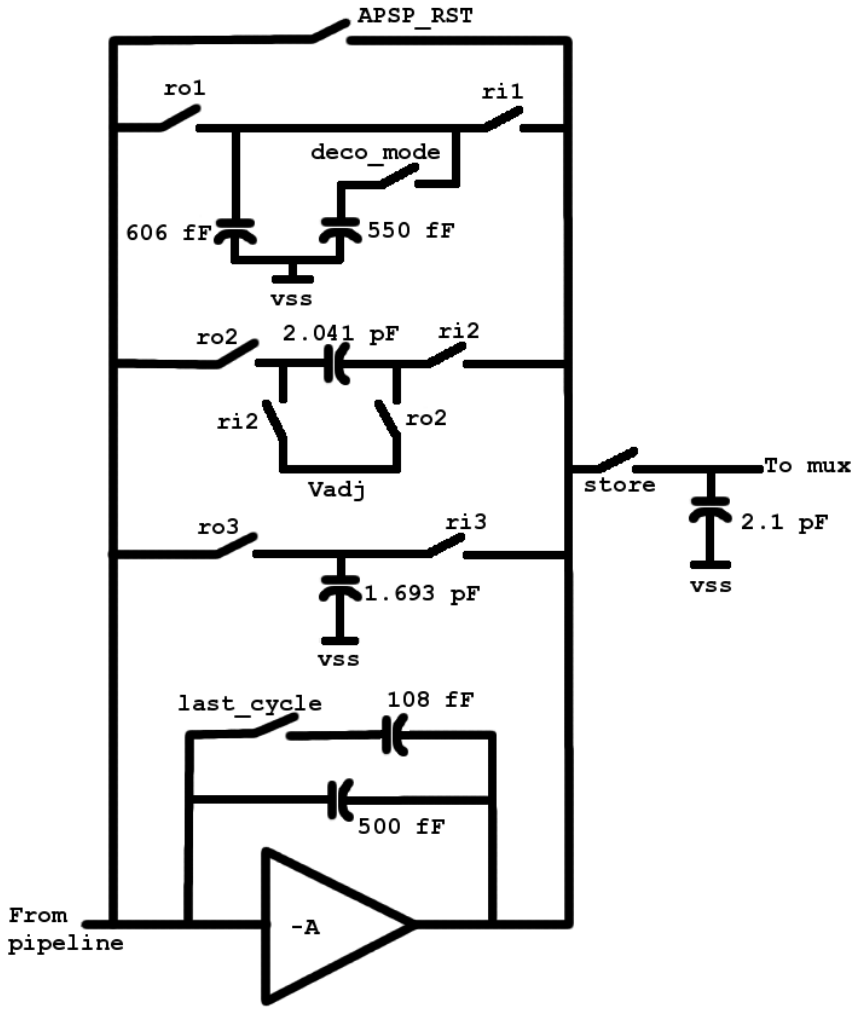

Figure 3. The APSP circuit diagram. The switches labeled ro1,2,3 and ri1,2,3 open and close during the APSP readout cycle for sample processing.

each APV25, a 12-bit digital header precedes the 128 analogue channel samples, creating a data frame of 7 μs total length, the same duration as the 4 pipeline readout cycles. Chip readout is simplified by matching the output data frame and pipeline readout durations, since at high rate triggered data stored in the pipeline can be read out while data from the previous trigger are being multiplexed out.

The Analogue Pulse Shape Processor (APSP) is the part of the APV25 chip which performs the deconvolution operation. The circuit diagram for the APSP is shown in figure 3. During processing, a series of switches in the feedback network of a high gain amplifier are opened and closed in sequence to apply the appropriate weight to each of the three samples, and make the weighted sum. The cycle time of the APSP circuit is chosen to be 1.75 μs so the total processing time matches the 7 μs readout time of the APV25.

APV25 behavior with controlled trigger spacing. To probe APV25 behavior, the programmable trigger generator was used to read data stored in specific pipeline locations by specifying the inter-

2010 JINST 5 P07007

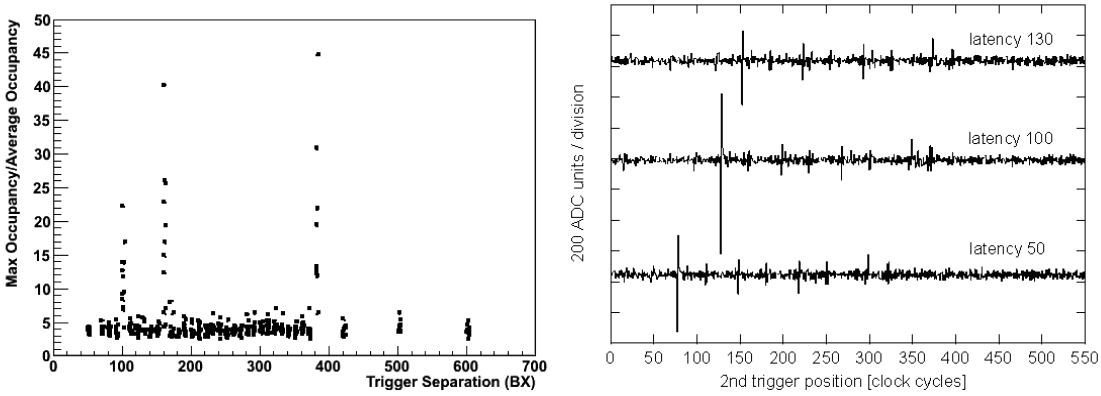

(a) (b)

Figure 4. (a) The maximum occupancy divided by the average occupancy within a fiber (2 APV25s) for all fibers connected to the rod as a function of time separation between two consecutive triggers. Certain trigger intervals show the HRN effect for all fibers simultaneously. (b) Second trigger pedestal data dependence on first and second trigger separation, for channel 127 on the single-module test bench. Data were taken in deconvolution mode for 3 values of latency, and averaged to remove random noise. An arbitrary vertical offset has been applied to the data for clarity.

vals between a master reset (RST) command and trigger (T1). A series of runs scanning T1-RST time separations was taken to assess the occupancy variation with pipeline cell, but no correlation was observed. Similar analysis of occupancy as a function of pipeline address rather than position, as well as a function of readout phase relative to the trigger, and as a function of pipeline cell occupancy, arranging use of adjacent pipeline cells for pairs of triggers, all eliminate these as possible causes of HRN. However, to probe the dependence on few μs timescale trigger intervals, a second trigger (T2) after T1 was added with a variable delay between the triggers. Figure 4(a) shows the maximum occupancy divided by the average occupancy for each fiber from the single-rod test bench, as a function of T2-T1 separations. The plot suggests that several trigger intervals, in particular, 100, 160, and especially 380 clocks, have significantly higher occupancies, for all APV25s simultaneously, indicating that the HRN was caused by interference between two triggers with very specific relative timing.

Additional tests with the single-module test bench (section 1.2) were performed to further probe the observed behavior. The digital pattern generator was programmed to cycle repeatedly through a sequence consisting of a reset, a fixed delay to allow the APV25 pipeline logic to initialize, then two triggers where the first trigger time was fixed relative to the reset, but the delay between first and second triggers was varied. Figure 4(b) shows the second trigger pedestal data dependence on first and second trigger separation, for APV25 edge channel 127. The data were taken in deconvolution mode for different values of programmed latency, and averaged over many triggers, removing the random noise component. Large pedestal disturbances well above noise levels are evident. Most importantly, the excursions occur at trigger separations which depend on the programmed latency, suggesting that they originate in on-chip activity.

Figure 5 shows a similar picture to figure 4(b), but for module strips 127 and 128, edge chan-

2010 JINST 5 P07007

nels on adjacent APV25s, for one value latency, in this case 130 clocks cycles. This demonstrates an anti-correlation in the pedestal disturbances between edge channels, as seen in previous tests. The pedestals associated with the feature at second trigger positions 157 and 158 have been labeled in the figure to illustrate this anti-correlation more clearly.

The APV25 output at actual trigger time represents the state of the APV25 output at the instant the second trigger was applied to the chip. The APV25 output data frame displayed results from the first trigger. There are no obvious correlations between the pedestal patterns arising from the second trigger and the APV25 output sequence, but if the crosstalk is generated at the chip input, as suspected, it is necessary to take the programmed latency into account. From the APSP circuit schematic (figure 3), the simultaneous operation of the switches will draw current from the power supply rails, which is liable to generate significant, impulse-like, current fluctuations throughout the system which may couple back into it and generate noise. This APSP transition switching activity causes interference to couple into the APV25 inputs which will affect data written into the pipeline at that instant, but the subsequent trigger that would sample that data arrives one latency period later.

For example, at second trigger position 380 in figure 5 there is a disturbance visible in the output data retrieved from the pipeline by the second trigger. These data were written into the pipeline 130 clock cycles earlier so should be compared with the APV25 output state at second trigger position ∼250, when the APV25 output data frame resulting from the first trigger was just beginning. The *APV25 output offset by trigger latency* represented in figure 5 is just the *APV25 output at actual trigger time* shifted by the 130 clock latency. It therefore represents the state of the APV25 output when data corresponding to the second trigger were being written into the pipeline. The repetitive patterns and spacing between similar features (often 70 clock cycles) imply a connection with the internal APSP cycle of the APV25. For example there are small spikes at the times of the tick marks, and the APV25 output frame header.

Although the tick marks correspond to features in the APV25 output data, the interference is not produced by a coupling between the output data signals and the chip inputs, but by switching activity within the chip, associated with the pipeline readout phase that is reflected in the tick-marks. The amplitude of the "tick-mark disturbance" is relatively small and would normally be lost in the random noise, averaged out in figures 4(b) and 5. The most prominent feature in the APV25 output data is the effect in clock cycles 157 and 158. This feature corresponds to the closing of switches in the APSP circuit during the period when the first data samples are retrieved from the pipeline, as shown by simulation, discussed below. There are smaller effects 70 and 140 clock cycles later, which correspond to the retrieval of the second and third samples. The later features which correspond to the APV25 header are associated with the APSP readout operation, and the sample/hold stage that precedes the APV25 output multiplexer. Although the other features in the output stream are evident after signal averaging, they do not represent significant additional noise. However the feature in cycles 157-158 is significant and believed to be the main origin of the high rate noise.

The fact that the single-module test bench does not employ the TIF DAQ hardware or software confirms that the HRN is intrinsic to the front end module. A finer grain trigger separation scan was performed with the single-rod apparatus, starting with the T2 arrival coinciding with the readout of T1, achieved with a separation of 362 clocks. Figure 6(a) shows the ZS occupancy calculated from

2010 JINST 5 P07007

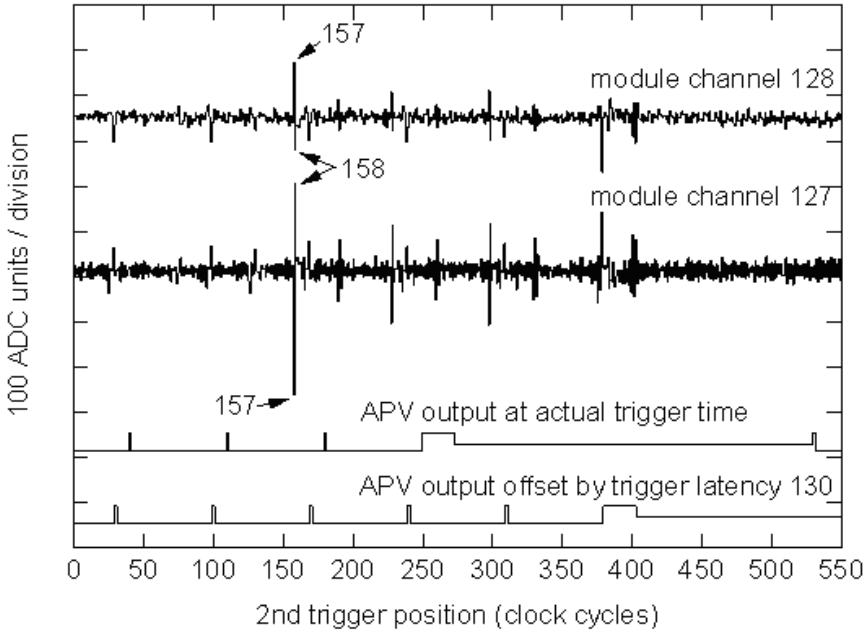

Figure 5. Second trigger pedestal data dependence on first and second trigger separation, for strips 127 and 128 on a TEC module, corresponding to edge channels on neighboring APV25s. Data were taken in deconvolution mode and averaged over many triggers to reduce random noise. The pedestals associated with the feature at second trigger positions 157 and 158 have been labeled to illustrate this anti-correlation more clearly. Two APV25 output data sequences are shown (see text for explanation). An arbitrary vertical offset has been applied to the data for clarity. On this scale, the first trigger was applied at second trigger position -3.

a run taken with this trigger separation, where low-numbered channels are clearly most affected, and figure 6(b) shows a clear anti-correlation in the common-mode subtracted ADC counts of channels 0 and 127 from a VR run taken with the same trigger separation. The T2 position was scanned in steps of 1 clock to produce trigger intervals between 362 clocks and 398 clocks. The ZS data acquired in this procedure show that the peaks of figure 6(a) diminish and shift toward the middle of the chip. The effect begins to reappear on the high-numbered channels starting with a separation of 375 clocks, and is maximal with a 384 clock separation, shown in figure 6(c), when the sampling for T2 is close to the end of header readout. There is again a strong anti-correlation in the pedestal and common-mode subtracted ADC counts for channels 0 and 127 for this separation, but with channel 127 now at positive displacement.

APV25 simulation and coupling mechanism. The measurements with the single-rod test bench were repeated with current probes monitoring the power lines of the APV25, but no obvious excursions were observed. However, detailed APV25 simulations motivated by the observation of the HRN effect find a current change that is correlated with readout. Figure 7 graphs the simulated response of the *GND*, 1.25 and 2.5 V supply and analog output lines during the APV25 readout sequence. The intermediate current spikes appearing on the supply lines are correlated with the falling edges of the three control signals for the APSP input phase and the reset signal to the APSP

2010 JINST 5 P07007

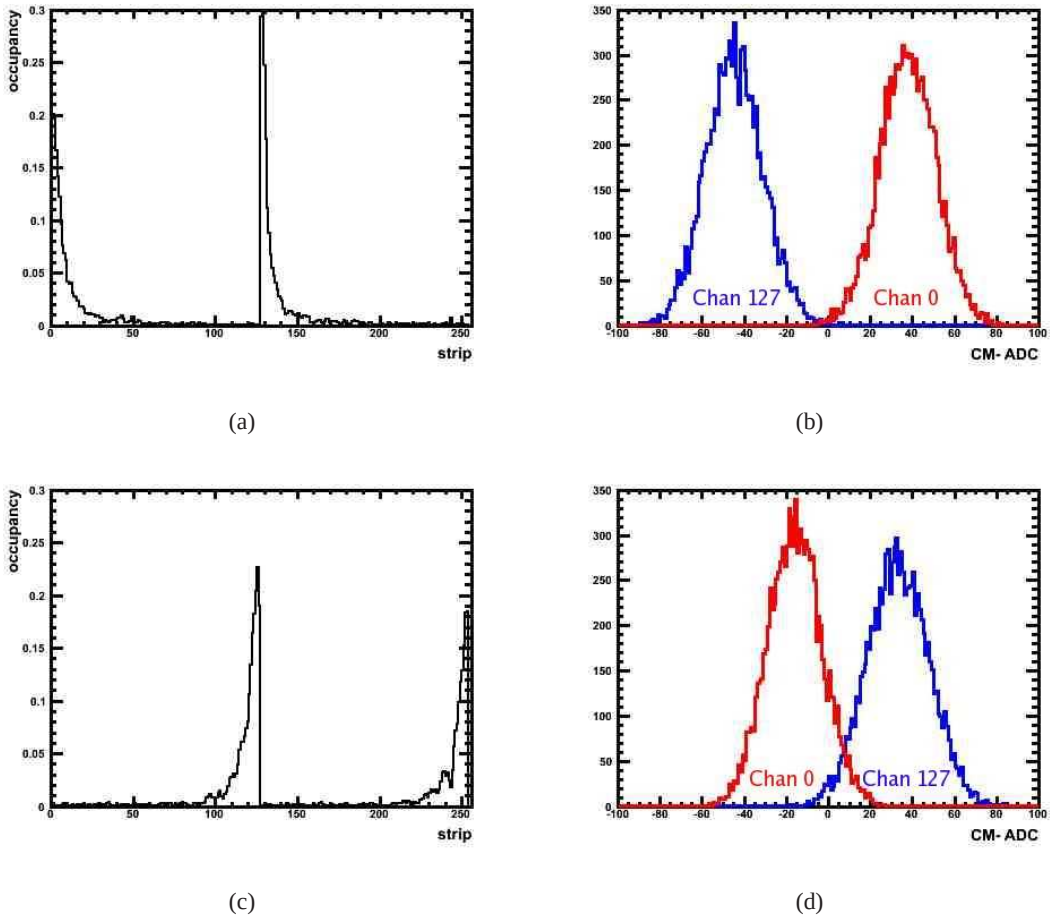

Figure 6. Data for (a) and (b) 362 clock trigger separation and (c) and (d) 384 clock trigger separation
(a),(c): Occupancy spikes occur in the low-numbered/high-numbered channels when setting the trigger interval to 362/384 clocks
(b),(d): Pedestal and common-mode subtracted ADC for channel 0 (red),and channel 127 (blue). Both channels show large displacements from zero and a per-event anti-correlation, but in opposite directions for 362 vs.384 clock separations

amplifiers, *ri1*,*ri2*,*ri3*, and *APSP_RST* in figure 3. The two largest spikes in each cycle come during the output stage and are correlated with the three control signals (*ro1/ro2/ro3*) of the APSP output phase. The positive-going spike on the +2.5V supply is caused by connecting the signal storage capacitor to the APSP output. The negative-going spike occurs when the internal sampling capacitors are disconnected from the APSP. The magnitudes of the largest spikes are below the level of sensitivity achievable with the current probes used in the measurements.

The supply-line current spikes that appear in the APV25 simulations are the likely origin of the HRN; however, an important observation that the HRN *does not occur* when the module was tested without a sensor (a readout hybrid alone) led to investigations of how the HRN couples to the APV25 front-end.

Data was taken with the sensor connection to the APV25 intact, and with the connection

2010 JINST 5 P07007

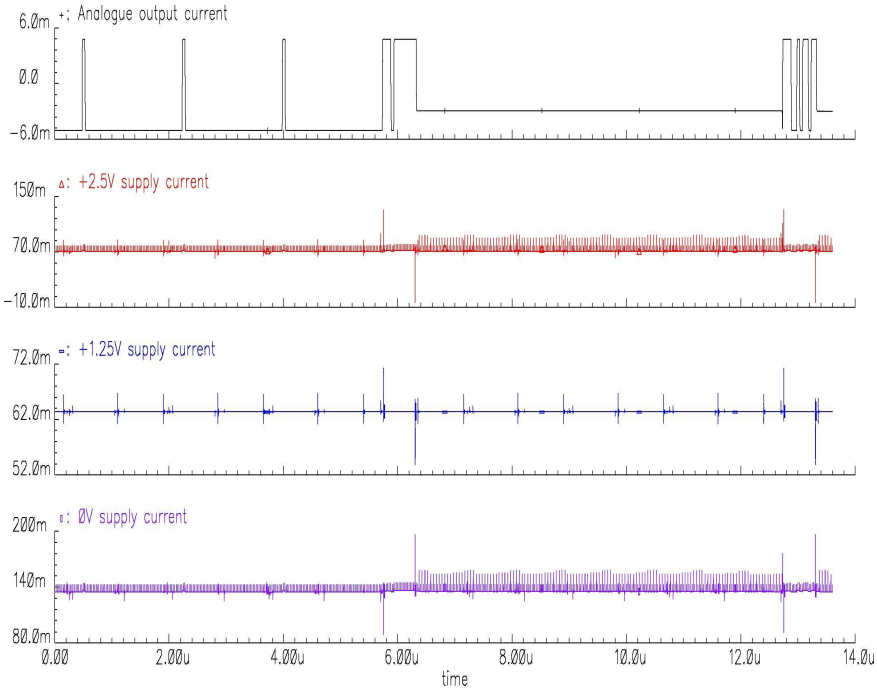

Figure 7. Simulation of APV25 chip behavior, showing short current spikes correlated with APV25 readout activity.

interrupted at different points: the wire bond between the pitch adapter and sensor detached at the sensor side, the bond between the pitch adapter and APV25 detached at the pitch adapter side, and the same bond completely removed. Removing the bond between the pitch adapter and sensor results in a sharp drop in HRN, likely due to a change in capacitance; however not large enough for the effect seen with the connected chip to be entirely capacitive in nature. The disparity in ADC counts between channel 127 on one chip and channel 0 on the next, which are connected to adjacent strips, suggests that it is unlikely that noise is coupling through the silicon sensor itself. If this drop is due to a smaller capacitance on the input, then the minimal difference observed between the pitch adapter connected and the chip with a wire bond alone indicates that most of the signal is picked up by the wire bond. This in turn indicates that the coupling is between the power and ground wire bonds and the adjacent channels.

The leading hypothesis which emerges from these tests is that the HRN is a consequence of inductive coupling to the current spikes produced during the APV25 readout cycle. The physical coupling is between the power bonds and nearby signal bonds, thus affecting only a small subset of channels. The susceptibility of the channel to this coupling is greatly enhanced by the increased capacitance and therefore decreased impedance when the sensor is attached to the APV25, while the pitch adapter and wire bond itself play a less important role. The anti-correlation between the two channels observed is due to the direction of the current flow, which adds to the signal for one channel and subtracts from it for the other.

2010 JINST 5 P07007

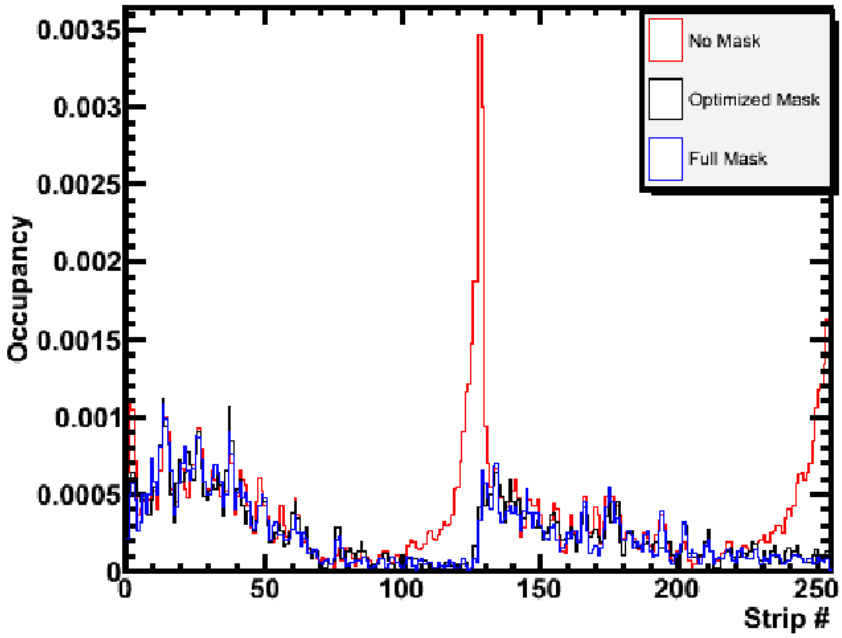

Figure 8. Occupancy of two APV25s in the single-rod testbench demonstrating the effect of the APVe-based mitigation. The red graph has no triggers blocked, the blue has all triggers during the APSP cycle blocked, and black only $\sim 1\%$ of the APSP cycle blocked.

3 Consequences and potential mitigation

The HRN effect generates strip occupancies of order $\sim 1\%$ at trigger rates of 100 kHz, similar to the design maximum subdetector occupancies of $\sim 3\%$ for the tracker, posing a potential bandwidth problem. However, the magnetic field that couples supply line noise to the APV25 front-end falls off quickly and only a few channels are affected on average. The impact of HRN on track reconstruction is more serious. A simple simulation which fluctuated ADC values of edge channels to emulate the addition of HRN into TIF cosmic data showed that the reconstruction of a single, fluctuated event requires 15 minutes to reconstruct using default offline clusterization thresholds, and identifies $\sim 30,000$ phantom track segments. Reasonable event processing times were only achievable by increasing thresholds to unrealistic levels.

CMS will mitigate the consequences of HRN by preventing pathological trigger intervals using an APVe [12], an online hardware emulator of the APV25 normally used to prevent buffer overflow. The APVe receives both trigger and reset signals and can use these to determine the bunch crossings in which data readout will occur. By programming the APVe to block triggers that occur at specific times during the APSP cycle, but only when a previous trigger has been applied, it is possible to prevent triggers which would give rise to a high occupancy event.

The modified APVe functionality has been implemented and tested at the TIF on the single-rod test bench. The possible disadvantage of this solution is the deadtime incurred due to the blocked triggers. The consequent deadtime is very small: 0.25% for each vetoed interval at the maximum 100 kHz trigger rate. Figure 8 demonstrates the effectiveness of this solution, showing the occupancy for one fiber running with 100 kHz Poisson triggers, with three different blocking

2010 JINST 5 P07007

conditions: no blocking at all (red), where the HRN is clearly visible; all 280 clocks of the APSP cycle blocked (blue), which effectively removes the HRN effect but with 40% deadtime at 100 kHz; and only 3 of the 280 clocks blocked (black), which gives the same occupancy profile as the full blocking but with only $\sim$ 1% deadtime at 100 kHz.

4 Conclusions

Tests of the CMS Tracker readout at high random trigger rate have identified an unexpected source of noise which has been traced to short duration current variations in the APV25 front-end chip. The noise enters the system via inductive coupling between wire bonds delivering power to the chip and their neighbors and is greatly enhanced by the sensor capacitance. It is a consequence of certain internal APV25 switching operations which cause short term fluctuations in the supply current. Rate dependent noise is not visible except when a microstrip sensor is connected to the chip.

The synchronous nature of the CMS Tracker readout means that all channels in the system are affected simultaneously which could have an impact on charged particle track identification and reconstruction. As the effect occurs in specific clock cycles following a trigger and readout, it is predictable and can be avoided by vetoing those cycles which coincide with a current spike, with a small impact on the overall deadtime.

Acknowledgments

We thank all our colleagues in the CMS Tracker collaboration for their contributions to the implementation of the system and their support during these studies. In addition, we would like to acknowledge the CMS DAQ team for invaluable help with the installation of the high rate DAQ, in particular Jonatan Piedra, Domenique Gigi, Christoph Schwick, Hannes Sakulin, Frans Meijers, and Sham Sumorok. The studies on the single-rod test bench benefited from the superb technical skills of Evgeny Zverev. We acknowledge support from: FMSR (Austria); FNRS and FWO (Belgium); CNPq, CAPES, FAPERJ, and FAPESP (Brazil); MES (Bulgaria); CERN; CAS, MoST, and NSFC (China); COLCIENCIAS (Colombia); MSES (Croatia); RPF (Cyprus); Academy of Sciences and NICPB (Estonia); Academy of Finland, ME, and HIP (Finland); CEA and CNRS/IN2P3 (France); BMBF, DFG, and HGF (Germany); GSRT (Greece); OTKA and NKTH (Hungary); DAE and DST (India); IPM (Iran); SFI (Ireland); INFN (Italy); NRF (Korea); LAS (Lithuania); CINVESTAV, CONACYT, SEP, and UASLP-FAI (Mexico); PAEC (Pakistan); SCSR (Poland); FCT (Portugal); JINR (Armenia, Belarus, Georgia, Ukraine, Uzbekistan); MST and MAE (Russia); MSTDS (Serbia); MICINN and CPAN (Spain); Swiss Funding Agencies (Switzerland); NSC (Taipei); TUBITAK and TAEK (Turkey); STFC (United Kingdom); DOE and NSF (USA). Individuals have received support from the Marie-Curie IEF program (European Union); the Leventis Foundation; the A. P. Sloan Foundation; and the Alexander von Humboldt Foundation.

References

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

2010 JINST 5 P07007

[2] CMS TRACKER collaboration, W. Adam et al., *Performance studies of the CMS Strip Tracker before installation*, 2009 *JINST* **4** P06009 [arXiv:0901.4316].

[3] CMS TRACKER collaboration, W. Adam et al., *Stand-alone cosmic muon reconstruction before installation of the CMS silicon strip tracker*, 2009 *JINST* **4** P05004 [arXiv:0902.1860].

[4] CMS TRACKER collaboration, W. Adam et al., *Alignment of the CMS silicon strip tracker during stand-alone commissioning*, 2009 *JINST* **4** T07001 [arXiv:0904.1220].

[5] CMS collaboration, *CMS, tracker technical design report*, CERN-LHCC-98-06.

[6] CMS collaboration, *Addendum to the CMS tracker TDR*, CERN-LHCC-2000-016.

[7] M. J. French et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, *Nucl. Instrum. Meth.* **A 466** (2001) 359.

[8] K. Hahn et al., *Studies of the CMS tracker at high trigger rate*, *CMS note 2009/015* (2009).

[9] CMS TRACKER collaboration, K. Bell et al., *User requirements for the final FED of the CMS silicon strip tracker*, *CMS Note 2001/043* (2001).

[10] N. Bingefors et al., *A novel technique for fast pulse shaping using a slow amplifier at LHC*, *Nucl. Instrum. Meth.* **A 326** (1993) 112.

[11] S. Gadomski et al., *The Deconvolution method of fast pulse shaping at hadron colliders*, *Nucl. Instrum. Meth.* **A320** (1992) 217.

[12] G. Iles et al., *The APVE emulator to prevent front-end buffer overflows within the CMS silicon strip tracker*, prepared for 8th *Workshop on Electronics for LHC Experiments*, September 9–13, Colmar, France (2002).

2010 JINST 5 P07007