CMS

2012/06/01
Archive Id:
Archive Date:

The Strip Tracker Spy Channel Readout System.

R. Bainbridge, N. Cripps, J. Fulcher, G. Hall, A.-M. Magnan, T. Whyntie
Imperial College London (UK)

Abstract

A parallel readout system has been developped to capture the raw data frame coming out of the APV25 readout chip of the CMS Silicon Strip Tracker detector. With about 10 million strips to readout, the data have to be zero-suppressed. Spying on the raw data is hence useful as a debugging tool for channels with faulty behaviour, as a validation tool for the zero-suppression procedure implemented in the standard readout electronics firmware, and finally as a unique way to provide measurements of the gain and noise during collisions, rather than relying on specific calibration runs. In this note, the hardware and software implementations of the spy channel readout system are described. Spy data taken in 2010 and 2011 are analysed, allowing the monitoring of faulty components and the extraction of the first results on real-time calibration.

1 Introduction

The CMS tracker [1] is composed of three parts: one central part with a cylindrical geometry, and two forward parts with disk geometry. The readout is organised in four partitions, called Tracker Inner Barrel (TIB, which also contains the Tracker Inner Disks from both sides, TID+ and TID-), Tracker End Cap from the plus side (TEC+) and minus side (TEC-), and Tracker Outer Barrel (TOB). The silicon detector modules are readout by two or three pairs of APV25 chips [2], such that each chip reads out 128 silicon strips. The readout chain from the APV25 to the Front-End Driver (FED) boards located off-detector, via optical fibers, is described elsewhere [8]. The raw data frame is then processed through the FED, with two possible paths: the spy channel readout, which captures the frame exactly as sent by the APV pairs, described in section 2.3, and the normal mainstream FED FE FPGA processing performing zero-suppression described in Ref. [8]. The next two sections are dedicated to the description of the software implemented to unpack and analyse the spy data, and to the first results of the analysis of data taken in 2010 and 2011.

2 The spy channel hardware implementation

2.1 Hardware path

The FED board is shown in Fig. 1. Data transmitted by optical fibers from the detector modules are connected to the FED via the eight connectors shown on the left side of the picture. The next row of 24 FPGAs on the right of the connectors are the delay FPGAs, where the spy channel readout is implemented. The next row after that are the eight Front-End FPGAs, in which the zero-suppression is done for the normal readout path. Each delay FPGA reads out four fibers, or eight APV pairs. It is hence often that a detector module (connected to either two or three fibers) is read out by two different delay FPGAs.

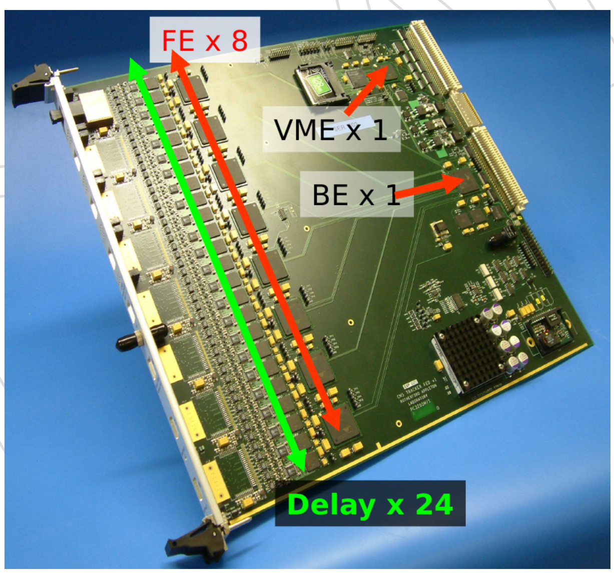

Figure 1: Side view of a FED board, with the FE panel on the left side receiving optical fibers from the detector, and the BE panel connected to the DAQ via S-link cables.

Not all the data can be spied on of course, so an arm signal must be sent to the delay FPGA to start the capture. The signal is sent via TTC, with a low rate of about 0.3 Hz. Once the signal is received, the next frame passing through the delay FPGA will be readout. These data are then propagated via VME, so completely independently of the standard data path going through the S-links. A block diagram of the two possible data paths between the analogue output of the silicon strips and the storage units is shown in Fig. 2. A specific run number is given to the spy data, encoded in the raw data, independent of the global run (although connected with it: it is given by the global CMS run manager, so if the spy run is started just after the gobal run, it will be just incremented by 1). Also, a spy run can cover several global runs, starting and ending the spy process is again independent of the mainstream process.

The FED buffer raw data is made of a header containing e.g. version ID, run number, followed by the data from the 24 delay FPGAs. Each delay FPGA data is made of 4 times 376 bytes for the payload of the four associated channels, plus an additional 8 bytes for encoding two 32-bit counters described in the next section.

The data are then converted to edm format, associated with the corresponding global run, transfered to castor, and registered into DBS under cms_dbs_prod_tier0, dataset name = /SiStripSpy*. The files are located in castor under the path: /castor/cern.ch/cms/store/streamer/SiStripSpy/.

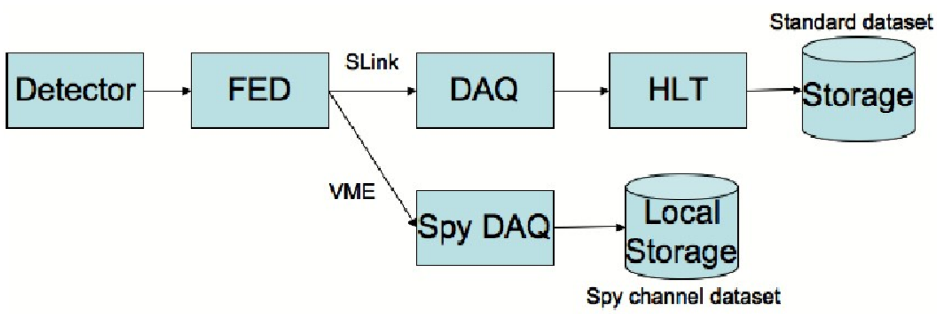

Figure 2: The two possible data paths between the analogue output of the silicon strips and the storage units.

FIXME: what is done in the delay FPGA ? Add schematic view ? And what is done between the delay FPGAs and the VME readout ? Can we have a block description of the content of the spy raw data, similarly to the normal data in Appendix A of Ref. [8] ?

2.2 The Spy Arm mechanism

2.2.1 Description

Two possibilities are implemented:

- arming every N events: convenient to have a sample independent of the data taking conditions, or for studying cosmic triggers.
- arming at a given clock period in the LHC orbit: convenient to study data either during collisions or specifically during the orbit gap e.g. for pedestal and noise monitoring.

Whatever method is chosen (implemented in the firmware), the spy data must keep track of counters to connect the event to a given trigger. This is done by using four counters:

- the APVe address, encoded in the FED buffer header;
- a 32-bit latching event counter incremented when the ARM signal is received, called TotalEventCount. It is encoded in the FED buffer per delay FPGA.
- a 32-bit latching trigger counter incremented when the ARM signal is received, called L1ID. It is also encoded in the FED buffer per delay FPGA.
- a 32-bit orbit counter FIXME: where is it encoded ?

The actual event number for matching to the mainstream event lies between TotalEventCount and TotalTriggerCount. The readout is triggered when these counters do not match their previous value. Together with the orbit counter and APVe address, a unique matching to the mainstream event can be made.

The arming and recording of the frames are FED-dependent: this means that in principle, two FEDs can record different events, even if the arming is made as synchronous as possible. The matching to mainstream event hence has to be done per FED, and there is no way to ensure that a whole event will be recorded consistently. In practice, it is often the case that all FEDs of a given partition record the same event, and often enough that several partitions also record the same event.

2.2.2 TTC B channel signals

FIXME: description ?

2.3 The APV frame

2.3.1 Description

The APV frame is shown in Fig. 3, as it is received by the FED, graphical representation of the spy raw data effectively.

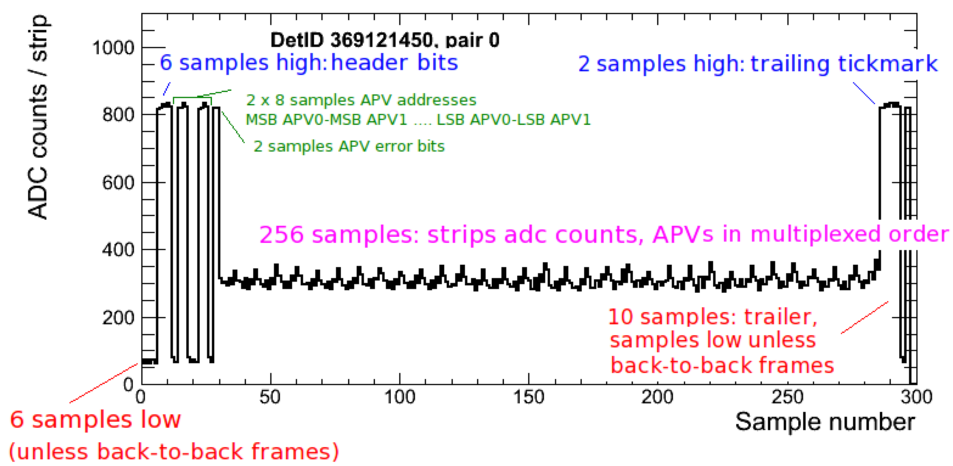

Figure 3: The composition of the APV frame as sent to the FED board.

The spy APV frame contains the data of an APV pair APV0-APV1. It is made of 298 16-bit samples, but processed through 10-bit ADC. Each sample is hence a value between 0 and 1023. The frame has four parts: a digital tickmark, a digital header, an analogue payload, and a digital trailer. As indicated in “digital”, the digital parts are interpreted either as “0” (low) or “1” (high). The lowest and highest values are channel specific, and driven by trimDAC settings, the difference between the two being the range $\mathcal{R}$, connected to the gain $\mathcal{G}$. The discriminating threshold is set to $\frac{2}{3} \times \mathcal{R}$. The gain is $\mathcal{G} = \frac{\mathcal{R}}{640}$. It is defined as the multiplicative factor needed to obtain 800 ADC counts (or 800 mV which is the input to the optical link).

The digital tickmark+header are made of 30 samples. The first 6 are normally low, then the next 6 are high and correspond to the tickmark: this is the signal that a frame has been detected. Next samples encode the APV addresses of both APV, starting with the most significant bits (MSB) for APV0, then for APV1, etc... down to the least significant bits (LSB) for APV0, then for APV1. The last two samples are the APV error bits for APV0 then APV1.

The payload is made of the 256 strips adc count values, starting with (strip 0, APV0), (strip 0, APV1), ..., (strip 127, APV0), (strip 127, APV1). The mean value when there is no signal is called the baseline or pedestal. It is normally situated roughly at $\frac{1}{3} \times \mathcal{R}$.

The digital trailer is made of 12 samples: 2 samples high, followed by 10 samples normally low.

2.3.2 Back-to-back frames

When the APVs register data for two consecutive triggers, the frames are called back-to-back, in which case the last 10 samples of the first frame and the first 6 samples of the second frame are not necessarily low. FIXME: what are they then ? Are the two frames valid anyway or will the second one be out-of-sync and lost ?

3 Processing and analysing the spy data

The software chain is presented in Fig. 4.

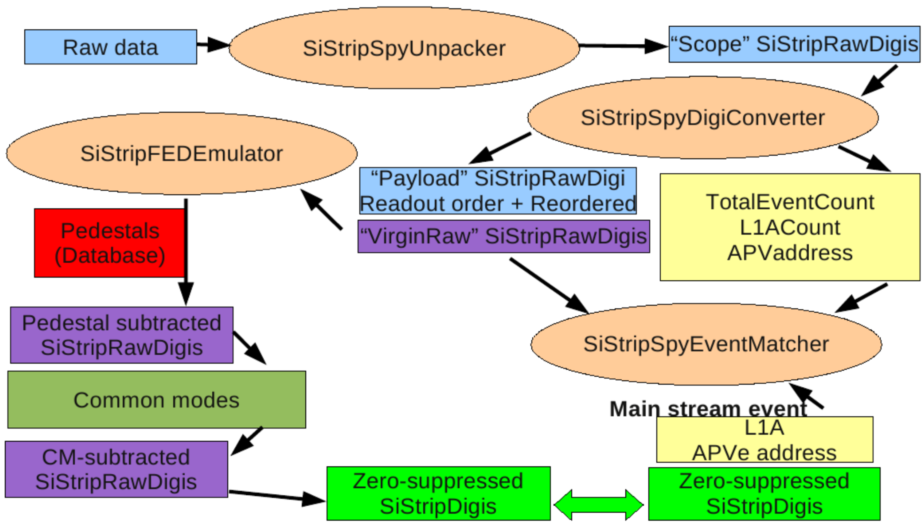

Figure 4: The software chain for processing and analysing spy data.

The code can be found under the CMSSW DQM/SiStripMonitorHardware package. The raw data buffer is first unpacked in order to extract the APV frame per channel, equivalent to data in scope mode (see Ref. [8]). The next step is to reorder the data by detector modules (detid) and extract the counters, so as to obtain data equivalent to the VirginRaw mode. Depending on the usage that will be made of the spy data, it is possible then to perform either or all of:

- matching to the mainstream event: for example to study Highly Ionizing Particles, or validate the FED zero-suppression firmware. The module allows to extract both spy and mainstream events and save them in the same output rootfile.
- emulating the FED zero-suppression algorithm: pedestal subtraction, common-mode subtraction, and zero-suppression.
- monitoring the APV frame quality: e.g. identify trimDAC settings problems, wrong pedestals.

4 Analysis of the data

4.1 Validation of the spy data

Special runs were taken with the mainstream event in VirginRaw mode. Spy data should be exactly identical to virginraw (VR) data. It was found to be mostly the case, but not always. Bit flips (differences in powers of 2) were observed in the adc value in about 20 strips per FED, in only 3 FEDs. The problem was identified to come from the FED FE FPGA firmware, and was since fixed for this issue.

The spy data are hence fully validated. About 1000 channels have empty frames in the spy data (but normal mainstream data) indicating a problem in the spy readout hardware which is not understood yet.

4.2 Validation of the FED firmware

The common mode subtraction sometimes fails in the FED FE FPGA firmware, leading to so-called shot events. The problem is still under investigation. It affects about 0.05-0.2 APV per event, dependent on the trigger rates.

4.3 Monitoring of bad channels

It is interesting to look at the frames given by all channels in order to identify bad components, and understand better what is the problem. Outside of empty frames (about 1000 channels), two categories of quantities can be monitored: the digital levels, and the position of the frame.

4.3.1 Digital levels

The digital levels can be monitored:

- zero-light level: this is the digital low level, normally situated at around 60 ADC counts.
- tickheight: this is the digital high level, normally situated at around 750 ADC counts.
- frame range: this is the difference between the zero-light level and the tickheight, normally around 690 ADC counts.

In a run taken on 19^{th} February 2010 (run number 128646), containing 400 events, the distribution of the tickheight as a function of the zero-light level is shown in Fig. 5 for all channels. The range as a function of the zero-light level is shown in Fig. 6.

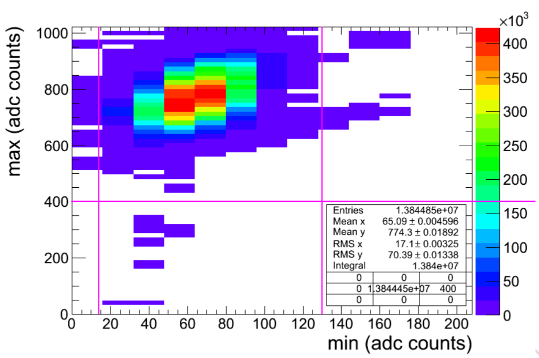

Figure 5: Tickheight as a function of the zero-light level for all channels, in run 128646 (400 events).

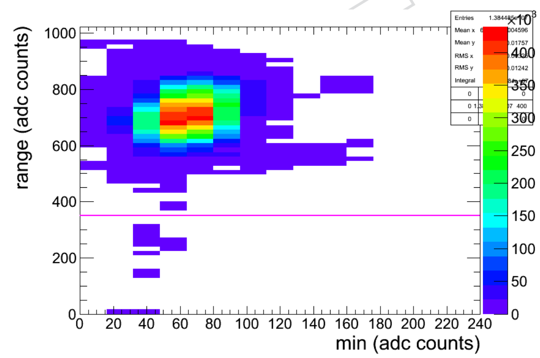

Figure 6: Range as a function of the zero-light level for all channels, in run 128646 (400 events).

These two distributions allow to identify criteria to tag unusual channels. The criteria chosen are: the zero-light level should be between 14 and 130 ADC counts; the tickheight should be above 400 ADC counts but not saturated at the maximum possible (1023 ADC counts); and the range should be above 350 ADC counts. Table 1 summarises the findings in that run for channels which do not pass the quality criteria defined above. Some of the detector modules have all APV pairs affected by the same problem (for example frames with a range close to 0), for others only one APV pair is affected. Examples of such frames are shown in Fig. 7 and Fig. 8.

4.3.2 Position of the frames

Six samples high must be found, and in the right position: in the samples numbered 6 to 11. Similarly, two consecutive high must be found at the end of the frame, in the right position: samples corresponding to the first header bit + 24 (header length) + 256 (adc samples).

A number of frames are found headless, with an example given Fig. 9.

Three delay FPGAs are found to give frames which are displaced compared to their expected position, but otherwise looking normal, as shown in Fig. 10. They are situated in FEDs 52 channels 20-23, 305 36-39 and 448 16-19.

As none of these show out-of-sync errors in the standard DQM analysis, it is probably a spy-related problem inside the delay FPGAs.

Channels showing APVError or APVAddressError in the DQM module (see Ref. [8]) are also confirmed here by identifying pairs with different addresses or error bits in the APV frame header.

Table 1:

Failure	FED ID	Channel ID	Detid	Partition
Tickheight saturated	346	77	470311498	TEC+
	52	15	369121450	TIB
	397	82	436278024	TOB
Low zero-light level	204	all		TEC-
High zero-light level	78	50	369157692	TIB
	97	86	369137062	TIB
	108	43	369142005	TIB
	112	83	369142206	TIB
	116	6	369142317	TIB
	116	33	369142358	TIB
	117	55	369141262	TIB
	149	8	402672909	TID
	159	38	402675344	TID
	182	94	470065648	TEC-
Range close to 0	322	40	47032938	TEC+
	427	90	436299660	TOB
	427	93	436299652	TOB
Low range	108	52	369142061	TIB
	210	20	470116456	TEC-
	245	84	470143914	TEC-
	248	59	470148006	TEC-
	277	5	470423140	TEC+
	282	70	470405608	TEC+
	322	40	470329381	TEC+
	350	89	470307753	TEC+

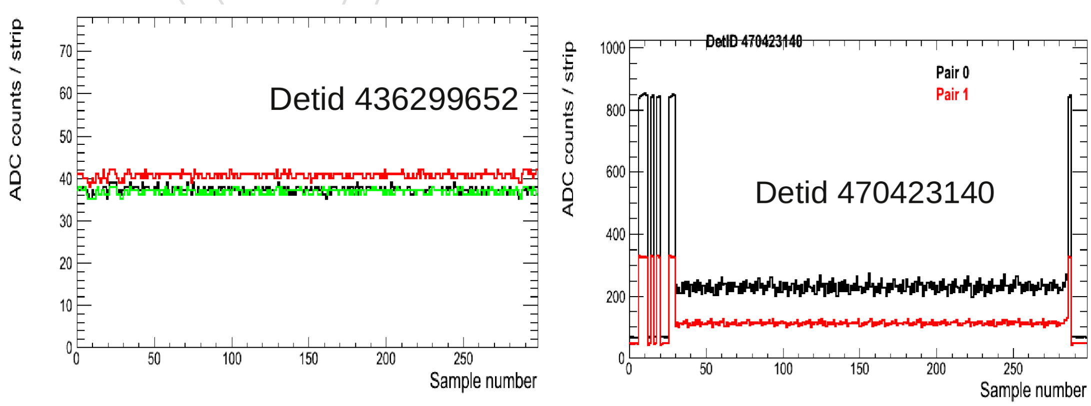

Figure 7: Frames with a range close to 0 for which no header will be detected (left) or with a small range but looking otherwise normal (right).

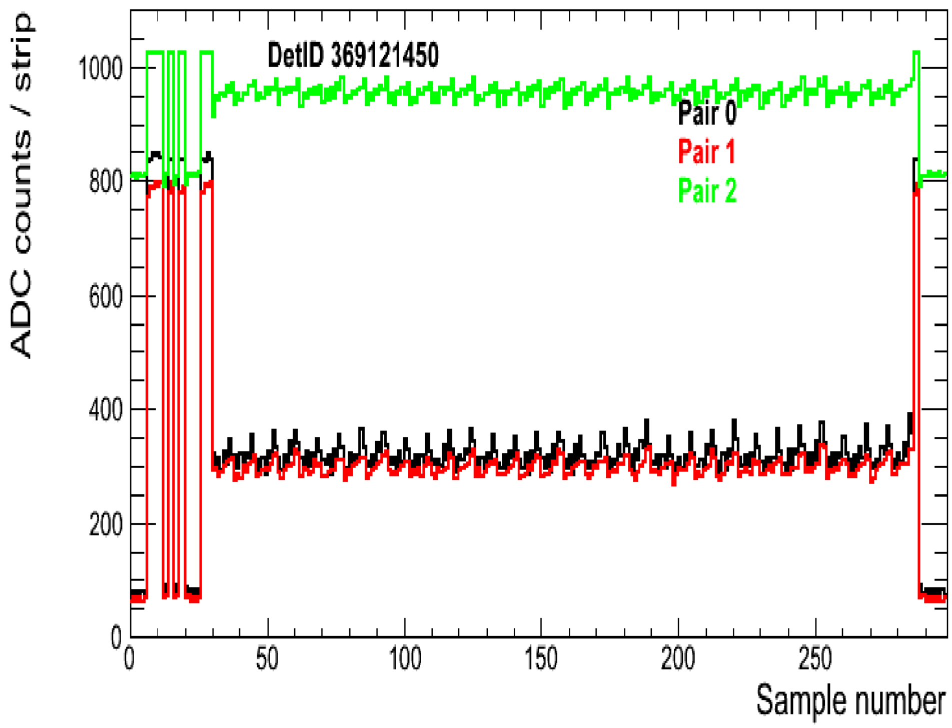

Figure 8: Frame with a high zero-light level and saturated tickheight, with a baseline at two-third of the range instead of one-third expected.

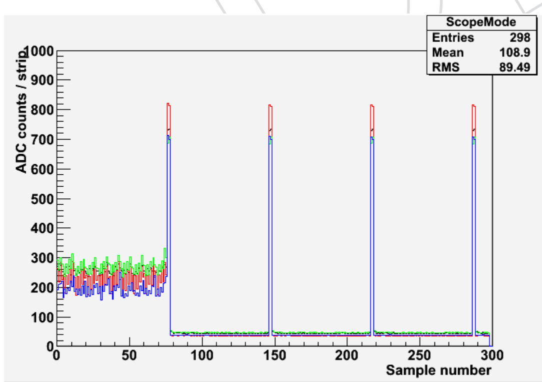

Figure 9: Frame which is missing a header. The two digital-high samples represents the APSP tickmarks spaced by 70 clock cycles, indicating that only the end of the frame has been recorded here.

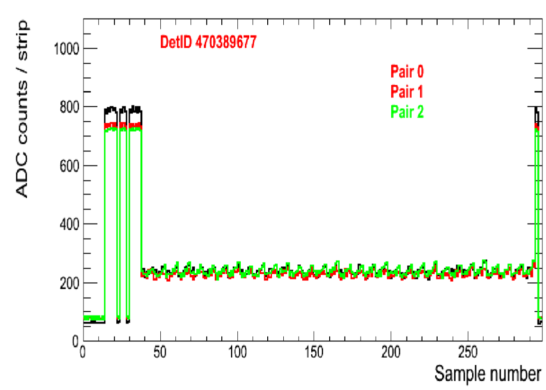

Figure 10: Out-of-sync frames: the frames start at sample number 12 instead of 6, but show the right length afterwards.

4.3.3 Cross-check of the pedestals

When processing the spy data through an emulation of the FED zero-suppression step, it is also possible to identify problems with the pedestals. Very few modules are found to have consistently the wrong pedestals for all strips in the database. In such case, the pedestal subtraction leads e.g. to negative pedestal-subtracted values, i.e. the pedestal value loaded in the database is actually higher than the raw ADC counts. Few strips are found borderline, either noisy (pedestal too small) or inefficient (pedestal too high).

4.4 Gain studies

The gain is studied in several runs taken in 2010, with a very first look at stability in time and uniformity across the detector.

The gain averaged over all channels, measured by the spy channel frame, is shown for three different runs in Fig. 11 as a function of time. Temperature effects can be seen at the beginning of these runs. Runs 131127 and 131268 were taken with the tracker reconfigured for global running just before the run. These temperature effects are clear, but nonetheless small enough (less than 1% variation) such that the average per run will not be affected. The gain per channel averaged over a run is shown for all channels in Fig. 12 for five different runs, showing similar distributions for all runs. Next step is hence to take the average per channel over all runs, to have enough statistics to study the uniformity of the gains across the detector. The results are shown in Fig. 13.

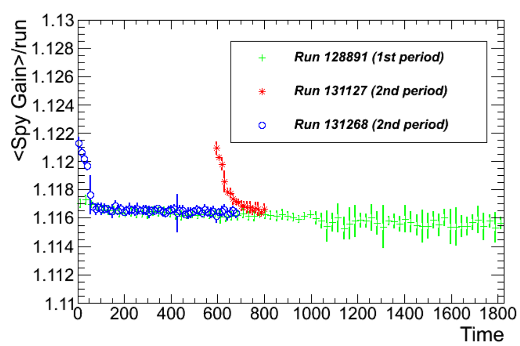

Figure 11: Gain measured by the spy channel data, averaged over all channels, as a function of time for three different runs taken early 2010.

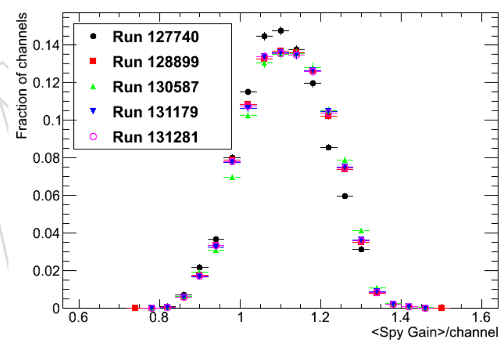

Figure 12: Gain measured by the spy channel data, averaged over a run per channel, shown for all channels for 5 different runs taken early 2010.

The gain obtained using the spy channel data is then compared to the gain recorded in the database from the commissioning runs. A good correlation is obtained, as shown in Fig. 14. The spy gain are however found to be systematically higher than the commissioning values by 0.05 (i.e. roughly 5%). It should be noted however that the spy data mix different runs and environment conditions, whereas the commissioning runs are designed to be taken during stable conditions. The difference between the two is also shown per partition in Fig. 15, and as a function of FED ID and channel id in Fig. 16. A different behaviour can be seen for TIB and TOB, showing quite large tails, compared to TEC+ and TEC-.

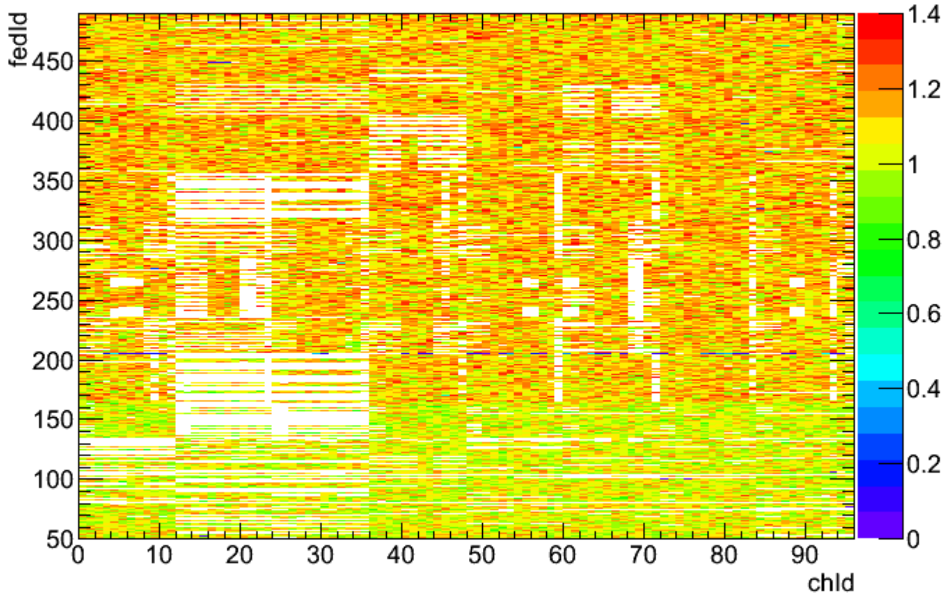

Figure 13: Average gain over all runs analysed (early 2010 data) per channel as a function of the FED id and channel id.

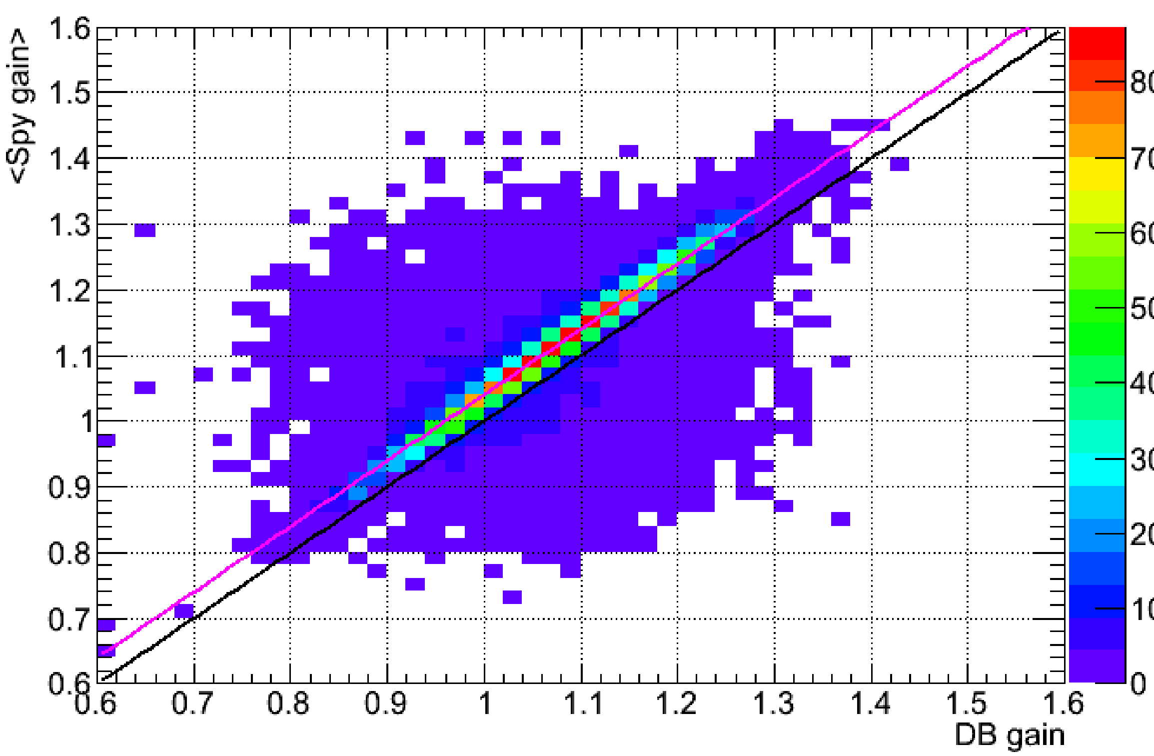

Figure 14: Average gain over all runs analysed (early 2010 data) per channel from the spy data as a function of the values obtained from the commissioning saved in the database (db).

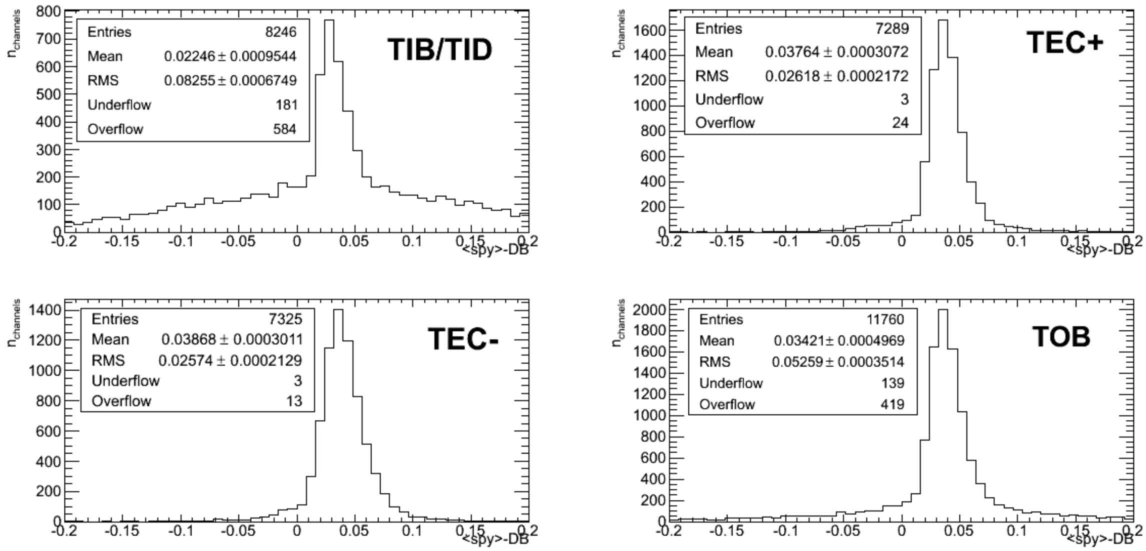

Figure 15: Difference between the results obtained with spy data or commissioning runs, separately for the four partitions.

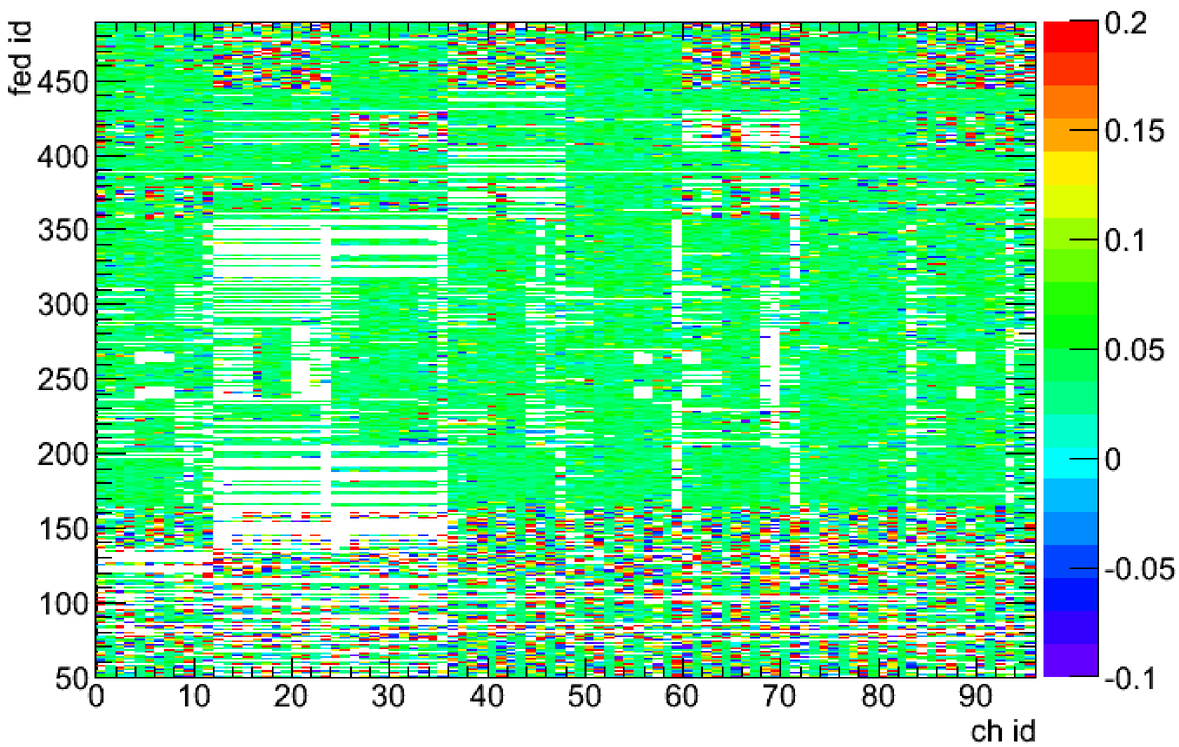

Figure 16: Difference between the results obtained with spy data or commissioning runs as a function of the FED id and channel id.

5 Conclusion

References

[1] The CMS Collaboration, JINST 3 (2008) S08004.

[2] M. Raymond *et al.*, “The APV25 0.25 μm CMOS readout chip for the CMS tracker”, Nuclear Science Symposium Conference Record, 2000 IEEE, 9/113-9/118 (vol.2).

[3] Fed Event Format, see `http://www.technology.stfc.ac.uk/esdg/cms-fed/user-manuals/event-format/FED-Event-Format.doc`.

[4] S-LINK64 protocol, see `http://www.technology.stfc.ac.uk/esdg/cms-fed/daq-link/slink64.pdf`.

[5] DAQ interface guide, see `http://cmsdoc.cern.ch/cms/TRIDAS/horizontal/RUWG/DAQ_IF_guide/DAQ_IF_guide.html`.

[6] L.Benucci *et al.*, ”Data quality monitoring of the CMS Silicon Strip Tracker detector”, Nucl. Instr. and Meth. A (2009), doi:10.1016/j.nima.2009.06.106.

[7] Tracker Maps, see `http://webcms.ba.infn.it/cms-software/cms-grid/index.php/CMSTrackerVisualizationSoftware/TrackerMap`

[8] The CMS Collaboration, “Monitoring of the CMS SiStrip Tracker Data”, CMS-DN-2011/007 (2011).