

Experience with Commissioning the CMS Pixel Detector

Anders Ryd*[†]

Cornell University, Newman Laboratory, Ithaca NY 14853, USA

E-mail: anders.ryd@cern.ch

The CMS pixel detector consists of three barrel layers and two forward disks on each side of the interaction region. The pixel detector has a total of almost 66 million channels. In this presentation an overview to the pixel DAQ system and the commissioning of the detector prior to installation is given. Some issues experienced with the operation of the detector are discussed.

17th International Workshop on Vertex detectors

July 28–August 1, 2008

Utö Island, Sweden

*Speaker.

[†]On behalf of the CMS pixel group

1. Introduction

The Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider has an all silicon based charged tracking system [1]. Closest to the interaction point is the pixel detector with $100 \times 150 \mu\text{m}^2$ pixels. The pixel detector consists of a three barrel layers (BPIX) and two forward disks (FPIX) on each side of the interaction point as shown in Fig. 1. The barrel layers are at radii of 4.3, 7.2, and 11 cm. The forward disks are at $z = \pm 34.5$ and $z = \pm 46.5$ cm with respect to the interaction point. There are a total of 15,840 readout chips (ROCs) in the CMS pixel detector. Each ROC controls 4,160 pixels for a total of almost 66 million channels.




Figure 1: The CMS pixel detector. The three barrel layers are 52 cm long and at radii of about 4.3, 7.2, and 11 cm respectively. The forward disks are placed at $z = \pm 34.5$ and $z = \pm 46.5$ cm.

Figure 2 illustrates the different components in the Pixel DAQ system [2]. At the top left are modules which consists of ROCs and a token bit manager (TBM). A TBM controls between 8 and 24 ROCs depending on the location in the detector. The ROCs are programmed through the frontend controller (FEC). A serial 40 MHz protocol is used over an optical connection from the off-detector VME modules to the portcard on the detector. This protocol has to be initialized by the so called tracker FEC (TkFEC) in order to set up the proper delays for the serial communication. This serial protocol reflect the data back at the TBM so that we can check that the communication is working properly.

The readout from the pixel detector is done over an analog line at 40 MHz. This analog line encodes pixel addresses using discrete levels. There are 6 distinct levels used for pixel address encoding. These levels are decoded in the frontend driver (FED) to reconstruct the pixel that was hit.

The online software is a distributed system based on the the standard CMS xdaq [3] toolkit. Figure 3 shows the different software components. The PixelFEDSupervisors (one for each VME crate) controls the FEDs, similarly the local trigger control (LTC), trigger and timing control (TTC), FEC, and TkFEC supervisors control the corresponding trigger and pixel VME boards. The PixelSupervisor application coordinates the activities among these components. In a typical online

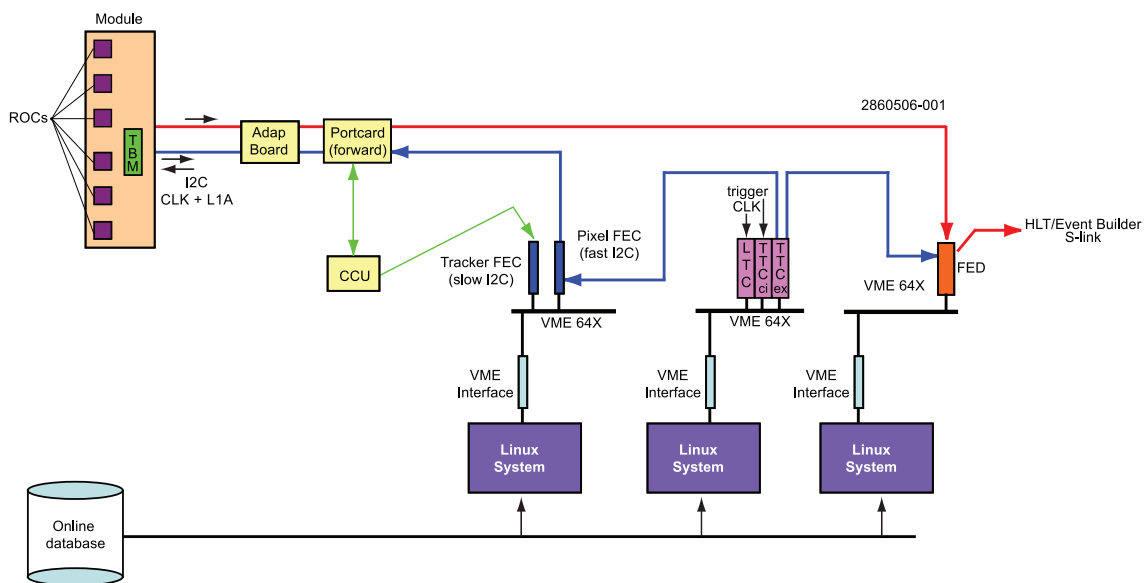


Figure 2: The main components of the CMS pixel DAQ system.

calibration a loop is performed in the pixel supervisor to 1) inject charge on a given set of pixels, 2) use the LTC or TTC to generate the trigger, and 3) use the FED to read out the data.

The interface to run control (RCMS) is done via the so called pixel function manager. In global running the pixel detector receives commands from global run control via the function manager. In local calibrations we operate the pixel detector directly via the pixel supervisor. In addition we have the detector control system (DCS) supervisor. This supervisor acts as an interface to the DCS system and is used to turn on and off power for the detector.

2. Commissioning before installation

The forward detector was assembled at Fermilab and transported to CERN. The last half cylinder arrived at CERN in early 2008. At CERN the detectors were tested. These tests included a complete set of calibrations at room temperature and at about -10 °C. Similar tests were done at PSI¹ with the barrel detector.

The set of calibrations we ran included

- Adjustment of the delay settings on the portcard to allow the serial 40 MHz protocol for the configuration of the frontend devices to work.
- Adjustment of the sampling point (delay and phase setting) in the FED for the digitization of the analog pulse.
- Analog optic hybrid (AOH) bias and gain adjustments.

¹Paul Scherrer Institute, Villigen, Switzerland

2860506-002

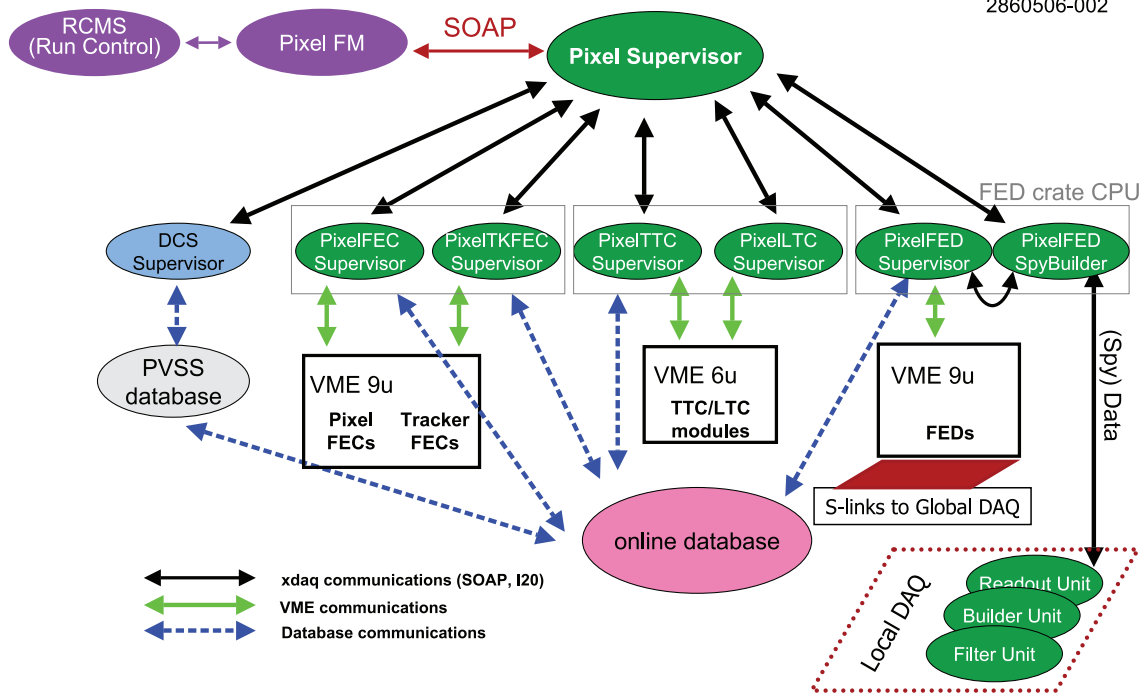


Figure 3: The online software components.

- Baseline adjustment in the FED to set the “black level” for the digitized input to a fixed target level.
- Settings of delays and thresholds in the ROC to allow reading out hits from charge injections.
- Calibration of address levels in the FED for decoding of pixels addresses.
- PixelAlive in which charge is injected on each pixels to check that the pixels are responding.
- Pulse height and linearity optimization. This adjusts DAC settings on the ROC to give a linear response up to about 1.5 mipis.
- Gain calibrations, in which, for each pixel we scan the injected charge and measure the response to calibrate the pedestal and gain. The analysis of the gain calibration data is described in Ref. [5].
- Trimming to unify the thresholds of the pixels.

In the process of testing the detectors a few issues were encountered and fixed. This included replacements of a panel and fixes to high voltage and low voltage connections. Overall the detector performance was excellent. For the FPIX all 4,320 ROCs were working and similarly for the BPIX only a small number of modules (40 ROCs out of 11,520) had problems at the time of installation. In Fig. 4 the threshold distributions are shown for a subset of pixel before trimming and after trimming. As seen the threshold dispersion is large before the thresholds were adjusted. The relatively large threshold here was targetted for initial operation to simplify the commissioning.




Figure 4: Threshold distributions before trimming a), and after trimming b). The units are in the Vcal units corresponding to approximately 65 electrons. The target threshold here of 80 Vcal units is fairly large (close to 5000 electrons). After trimming the pixels the rms of the threshold distribution is about 1.4 Vcal units, or 90 electrons.

3. Installation and checkout

The CMS pixel detector was installed in late July 2008. The complete pixel detector consists of six separate components. The BPIX detector was installed as two half shells and the FPIX detector as four half cylinders. The BPIX detector was installed first and a quick connectivity test was performed to make sure that all connections were made and that basic functionality such as distribution of clocks etc. was established. After the BPIX detector was installed the FPIX detectors were inserted, one side at the time. The two half cylinders on each side are installed at the same time.

In the first pass of the checkout of the FPIX detector it was verified that we could program and read out all 4,320 ROCs. This was done by turning on one sector (a sector is one 32nd of the FPIX detector) at the time. This initial checkout was completed within a few days for both the FPIX and BPIX. After this checkout had verified that all connections, optical and electrical, were working the CMS detector was starting to close and we lost access to the pixel detector.

4. Dead modules, dead and noisy pixels

Before installation all modules in the FPIX part of the detector worked correctly. During installation a HV wirebond was damaged and this took out 10 ROCs. During the initial checkout of the FPIX all ROCs were working properly otherwise. However, about a week after installation one detector segment developed a short in the digital power. Another group of 101 ROCs developed a high voltage short of the sensor bias after about a month, and a group of 24 ROCs developed a problem with the analog output signal. Together these failures represent about 6% of the FPIX detector.

For the BPIX detector there were a few known problems already after assembly. Attempts to fix these problems were made, but these interventions generated new failures and it was decided to install the detector with the following known bad modules

- Broken sensor high voltage connection, 8 ROCs.
- Broken signal wire-bonds in 3 modules, 32 ROCs.

However after installation in CMS a few more modules were discovered to have problems

- No high voltage for sensor bias on 3 modules, 32 ROCs.
- Broken signal wire-bond on one module, 8 ROCs.
- One module can not be programmed, 16 ROCs.

There are an additional 4 ROCs that don't generate an output signal and were disabled. For the BPIX this adds up to 100 ROCs, or 0.87% of the BPIX detector that is non performing. The plan for the FPIX detector is to take it out during the winter access and repair these problems. The BPIX will not be removed.

For the ROCs that work, the number of dead pixel cells is very small. In the FPIX the number of bad pixel cells is about 0.015% while for the BPIX we have about 0.010% dead pixels. This does not include the number of bad bump bonds. From module testing this is known to be about 0.01%. Therefore, the permanently dead pixel fraction is very small, about 2×10^{-4} . This is much smaller than the expected dynamical losses. At the LHC design luminosity this is expected to be a few percent [1].

Some pixels are noisy and have been masked off (disabled). A pixel is defined to be noisy if it has a hit in more than 0.1% of all events read out. With this criteria 263 pixel for the BPIX were masked off and 17 pixels for the FPIX. The fraction of noisy pixels is very small, below 0.5×10^{-5} .

5. Technical issues

Though the installation and initial checkout was successful there are a few issues that will require special attention in the future to simplify the operation. Here I will point out some issues we had with the analog link and the 40 MHz serial protocol for configuring the frontend electronic.

The readout of the CMS pixel detector is done with analog optical links as described above and discussed in more detail in Ref. [4]. The signal consists of an event header in which the signal goes low, ultra black (UB), for 3 clock cycles and back to the black level for one cycle. The next 4 clock cycles encode an event counter. Each clock cycles encodes four levels which gives a counter from 0 to 255. This concludes the TBM header, next follows the ROC header for the first ROC. This consists of an UB followed by a black level. The third clock shows the last programmed DAC or the ROC temperature. Each hit on the ROC is encoded in six clock cycles. The first five encodes, in base 6, the pixel address on the ROC and the the last cycle is the analog charge. This is repeated until there are no more hits on the ROC and then a new ROC header follows for the next ROC. After all ROCs are processed there follow a TBM trailer consisting of two UB levels and 2 black levels. The next 4 clock cycles encode a status word from the TBM. This again uses base 4 as in the event counter in the header. It is crucial that the address levels are kept stable in order to correctly decode the pixel addresses.

In order to decode the correct pixel addresses and event counters we need to determine the address levels. This is done as part of the online calibrations. However, these address levels are

sensitive directly to the light yield and this has to be kept stable. One of the issues we have is that the laser diodes are very temperature sensitive. We have measured that the levels seen in the FED after digitization change by about 45 ADC counts per degree Celsius of temperature change of the AOH. This can be compared to a typical separation between the address levels of about 80 ADC counts. As we will not be able to keep the temperatures sufficiently stable an automatic correction in the FED has been implemented. This correction adjusts the black level continuously when no event data is sent to be at a fixed value. However, if the offset is too large this correction mechanism breaks down as the FED can not identify when event data is sent.

For the barrel detector a thermal connection was made between the AOH and the cooling pipes while for the forward pixel detector the AOHs are just cooled by the air in the support cylinder. The forward detector channels shows much larger variation in the baseline than the barrel. We are currently investigating if we can install cooling for the AOHs in the winter 2008–2009 shutdown.

Programing of the front ends are done via a serial 40 MHz link. This link consists of a clock and a data line. In order to decode the data, the phase of clock and data has to be adjusted. There is also a return data and return clock. This is used to return the data so that we can verify that the communication worked. The return data is generated by the TBM. This means that we are not actually verifying that the data was received by the ROC with this mechanism. But it does verify the most crucial part of the link, the optical connections and that the decoding of the signal in the TBM worked.

The 40 MHz protocol for programming the frontend devices, ROCs and TBMs, have basically worked well. It allows us to configure the complete pixel detector in about 45 s. This includes downloading trim and mask bits for every pixel in the detector. We have had some problems with the reliability of this communication. Part of this might be related to problems with the optical connection. Some number of failures has been fixed by carefully cleaning the fibers. However, the region of delay settings that makes the communication work is very small for some links. This can be as small as about 1 ns out of the 25 ns clock cycle. For the barrel detector the delays needed for different modules along the barrel ladders varies and there is no return data-clock delay setting that works simultaneous for all modules. Hence the return data is only check during running for the FPIX.

6. Conclusions

The CMS pixel detector was successfully installed in the summer of 2008. The installation took about one week including initial checkouts. We have operated the detector successfully both in local calibrations and in global runs with the full CMS experiment. Recently a three week long run to collect 300 M cosmic triggers were completed. This data will allow us to align and perform other calibration of the pixel detector. Some initial results of are presented in Ref. [5]. The pixel detector performance so far is excellent. In the shutdown this winter a few minor problems will be addressed.

7. Acknowledgments

This work was supported by the NSF grant PHY-0757894 and the Alfred P. Sloan Foundation.

I also wish to thank all my CMS collaborators who has contributed to this project and the organizers of the VERTEX 2008 conference.

References

- [1] R. Adolphi *et al.* [CMS Collaboration], JINST **3** (2008) S08004.
- [2] D. Kotlinski, Nucl. Instrum. Methods A 565 (2006) 73.
- [3] J. Gutleber and L. Orsini, Cluster Computing 5, (2002) 55.
- [4] H.Chr. Kästli *et al.*, Nucl. Instrum. Methods A 565 (2006) 188.
- [5] V. Chiochia, These proceedings.