CMS Phase-1 Pixel Detector: Operational Experience, Performance and Lessons Learned

Atanu Modak[*][†]

Kansas State University, USA

E-mail: atanu.modak@cern.ch

The CMS Phase-1 pixel detector was designed to cope with an instantaneous luminosity of $2\times10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at 25 ns bunch spacing with very small efficiency loss. The Phase-1 pixel detector features four barrel layers and three disks per side, with in total 124 Million pixels, while the original CMS pixel detector had one layer and disk less in the barrel and endcap respectively, and only half the number of channels. DC-DC converters were introduced to deliver more power to the detector without the need of replacing the cable plant. CO_2 based cooling was implemented and carbon-based structures were constructed to reduce the material in the tracking volume. The data acquisition system was upgraded to accept higher event rates and a new, digital data format from the detector front-ends. The detector was installed in early 2017 and has been successfully operated since. The LHC is going through a planned long shutdown period during 2019-2020. The pixel detector was extracted in early 2019 after the end of Run 2 data taking and has been kept cold to protect the silicon sensors. The innermost barrel layer will be replaced during this shutdown period and will feature improved ASICs and circuit boards to rectify issues discovered during data taking. This paper focuses on the operational experience of the detector in 2018, highlighting the detector performance and addressing the lessons learned. The current status of the detector and the refurbishment plan will also be discussed.

[*]Speaker.

[†]On behalf of the CMS Collaboration.

1. Introduction

The Compact Muon Solenoid (CMS) [1] experiment at the Large Hadron Collider (LHC) features a silicon based pixel detector at the innermost region close to the interaction point. The pixel detector is a critical component providing high-precision space point measurements for efficient reconstruction of the charged particle trajectories near the interaction region and for the reconstruction of interaction vertices. The original CMS pixel detector (in the following referred to as "Phase-0 detector") was designed for instantaneous luminosities of up to 1×10^{34} cm^{-2}s^{-1}. However, the LHC performed beyond expectation, with the instantaneous luminosity projected to double in 2017. The original CMS pixel detector would not have been able to operate efficiently at the corresponding high data rates as a result of losses in the front-end readout chip. The Phase-0 detector has been replaced with the Phase-1 detector during the LHC winter shutdown 2016-2017 and has been successfully operated since. In this paper, the salient features of the Phase-1 upgrade will be briefly introduced, and the operational experience and tracking performance in 2018 will be described. Currently, the LHC is going through a two-year shutdown period (Long Shutdown 2 (LS2), 2019-2020) for maintenance. The plan for the pixel detector refurbishment during LS2 will also be addressed.

2. Phase-1 Pixel Detector

Compared to the original design, the Phase-1 pixel detector [2] features one additional barrel layer and endcap disk per side. This upgraded design places the innermost barrel layer even closer to the interaction point owing to a narrower beam pipe that was installed during Long Shutdown 1 of the LHC in 2014. Four concentric barrel layers are placed at radii of 29 mm, 68 mm, 109 mm, and 160 mm, and three endcap disks on each end of the detector are placed at distances of 291 mm, 396 mm, and 516 mm from the interaction point. There are in total 1856 hybrid silicon sensor modules in the pixel detector, where 1184 modules are used in the barrel layers and 672 modules are used for the endcap disks. The pixel modules in the disks are arranged in two rings. Each module consists of 16 readout chips and each chip features 80×52 pixels. Figure 1 shows a schematic comparison of the Phase-1 and Phase-0 detector.

The pixel detector upgrade aimed to handle an instantaneous luminosity of 2×10^{34} cm^{-2}s^{-1} by improving the front-end readout chip (ROC) [3], and by migrating from a 40 MHz analog to a 160 Mb/s digital readout scheme [4], without changing many key features, such as the sensor
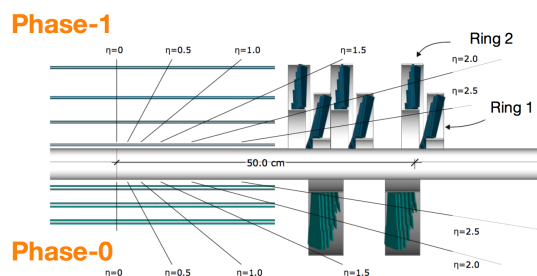
Figure 1: A comparison of the Phase-0 and Phase-1 designs in the longitudinal plane.

specification and readout architecture. The expected data rate in the innermost barrel layer was significantly higher than in the other layers and disks. Therefore layer 1 modules use more data links and feature a special readout chip. Instead of a full double column drain as used in the ROC for layers 2-4, layer 1 chips are equipped with a 2×2 dynamic cluster draining architecture. An ASIC named Token Bit Manager (TBM) [4] organizes the data flow from the readout chips and plays an important role in data transmission.

The addition of one more barrel layer and endcap disk per side almost doubled the number of pixel channels, increasing the total power consumption and the need for more efficient cooling. Since it was impossible to accommodate extra cables to account for this power, a new DC-DC conversion based powering scheme was adopted [5]. Another aspect of the Phase-1 upgrade is the replacement of the mono-phase C_6F_{14} cooling with a more efficient two-phase CO_2 cooling. Finally, the use of thinner, light-weight, stainless steel cooling pipes and carbon fiber based support structures reduced the material budget considerably and contributed to more robust tracking and vertex reconstruction. Figure 2 compares the material budgets of the Phase-0 and Phase-1 detectors. In spite of adding one additional barrel layer, the material budget is not increased in the central pseudorapidity region. Auxiliary electronic boards in the service cylinders were moved towards higher pseudorapidity regions to reduce the material budget in the sensitive tracking volume.
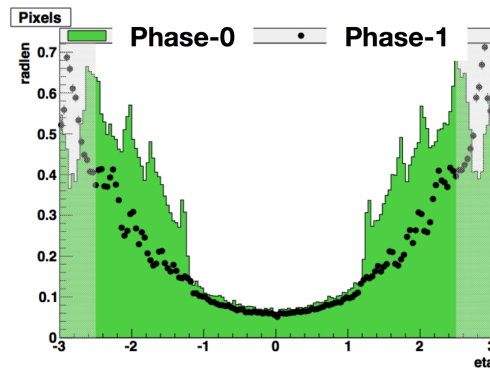
Figure 2: Comparison of the Phase-0 and Phase-1 material budgets in units of radiation lengths as a function of pseudorapidity, η.

Figure 3 illustrates the data acquisition architecture, including the pixel modules, the service cylinder with auxiliary components and optical hybrids [6], and the back-end electronics. Upon reaching the service cylinder, the electrical signal from the module is converted into an optical signal, and is subsequently transmitted to the back-end data acquisition system situated almost 80 m away. The back-end data acquisition system has been upgraded to the μTCA standard [7]. There are two front-end controllers (FEC); PixelFEC and TrackerFEC. The PixelFEC is responsible for programming the modules through the Digital Opto Hybrid (DOH) and also sends clock, trigger, and fast signals to the front-end chips. The TrackerFEC controls the powering and communication units through slow commands. The modules are connected to the support electronics in the service cylinder and readout data are collected by the front-end drivers (PixelFED), which are in turn connected to the central DAQ system through a high speed 10 Gb/s link. The LHC clock and CMS triggers are sent to the pixel system by the Trigger Control and Distribution System (TCDS) [8], which is a central service in CMS.
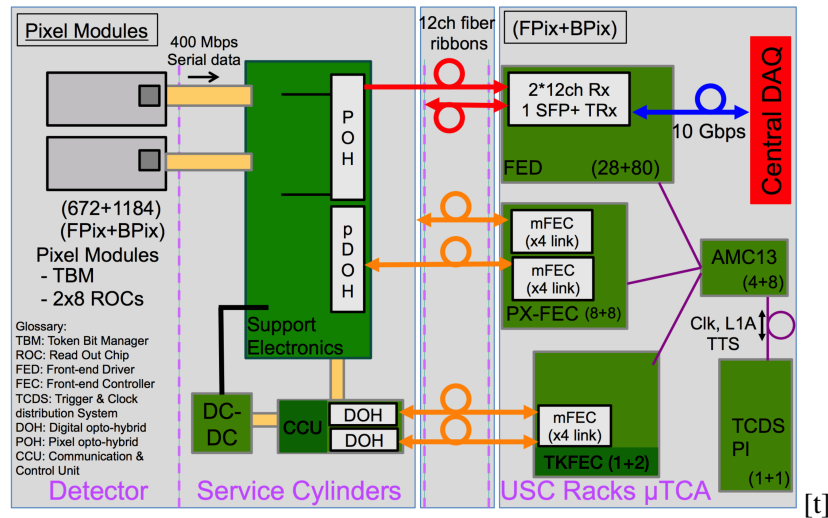
Figure 3: Overview of the μTCA-based data acquisition (DAQ) system of the CMS Phase-1 pixel detector.

3. Operational Experience

Many DC-DC converters have experienced a failure during the 2017 operation, which resulted in non-active detector modules. By the end of 2017, 5% of the total converters were non-functional, with almost 11% of the detector being unresponsive. Therefore the detector was extracted during the 2017 year-end technical stop (YETS) of the LHC, and all the DC-DC converters were replaced with new ones featuring a larger on-board fuse compatible with a lower operating voltage and tolerant to higher current. This refurbishment improved the active detector fraction to ∼97% at the beginning of the 2018 operation. However, the reason of failure of the DC-DC converters was still unknown at that time. In May 2018, the failure was reproduced for the first time in a laboratory outside the CMS detector. It was determined that a fault in the FEAST [9] chip design was responsible for the breakdown due to a charge build-up in a circuit when the chip is in a disabled state and once irradiated to doses above 10 kGy [10]. Enable/disable cycles of the DC-DC converters were necessary because the TBM has a design flaw in a circuit that gets stuck due to single event upsets (SEU) and recovers only after a power reset. Once the problem of the FEAST ASIC was understood the DC-DC converters were always kept in the enabled state, and instead the power supply system, which has coarser granularity, was used to recover from latched TBMs.

One of the major improvements in 2018 was the upgrade of the PixelFEC firmware [11] which is responsible for programming the modules. The new firmware utilizes a local memory on the back-end DAQ card, which has a segmented structure. It allows to pre-store the configuration data and to program the modules in a more parallelized manner. This upgrade reduced the total configuration time of the detector from ∼30 s to ∼3 s, as can be seen in Fig. 4, left. The pixel detector, being situated very close to the interaction point, suffers from a significant rate of SEUs on front-end electronics and optical components. The general strategy to recover from the SEUs has been to reprogram the detector at regular intervals with a fast configuration scheme. The online software has a functionality that constantly looks for SEU-induced errors and activates a recovery procedure whenever required. A full pixel-level programming during this recovery was not possible
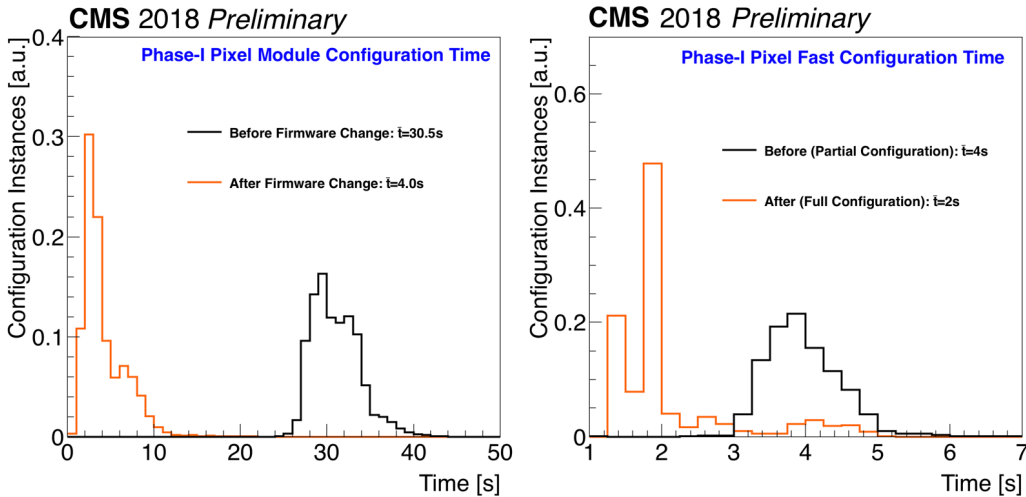
Figure 4: (Left) Module configuration time comparing the old and new firmware, including loading of data into the local memory. (Right) Fast configuration time of the pixel detector before and after the firmware update. With the old firmware only a partial configuration was possible, while the new firmware can perform a full pixel-level programming much faster with pre-loaded configuration data on the local segmented memory.

with the old firmware as it would have taken too much time, and would have resulted in data loss. Therefore only a partial, readout chip level reprogramming was being performed. However, the new firmware, with local pre-stored data, reprograms the detector with the full pixel-level configuration in less time than needed for the partial configuration using the old firmware. This in turn leads to a quicker and more complete recovery of the detector from SEU-induced errors. Figure 4, right, shows the timing performance of the old and new firmware for fast configuration of the detector during SEU recovery.

The distributions of the readout chip thresholds in number of electrons for the barrel and endcap detectors are depicted in Fig. 5, left and right, respectively. These distributions are normalized to the number of ROCs in the respective layer or ring. The thresholds for layers 2, 3, and 4 have been adjusted to ∼1300 electrons and those of the endcap disks to ∼1600 electrons. The modules in layer 1 feature a special readout chip, which was found to suffer from rate-dependent cross-talk. Thresholds in layer 1 have been set above 2000 electrons to suppress this effect. Modules in layer 2 and ring 1 that were damaged in 2017 and could not be replaced are also shown, with thresholds being kept higher to reduce noise.

4. Detector Performance and Radiation Monitoring

The pixel detector continued to deliver an excellent tracking performance and position resolution throughout 2018. Figure 6 shows the hit reconstruction efficiency as a function of the instantaneous luminosity for all layers and disks. Layer 2, 3, and 4 along with the endcap disks exhibit a hit efficiency of more than 99% up to an instantaneous luminosity of 2×10^{34} cm^{-2}s^{-1}. However, layer 1 starts to show some inefficiency at low ($< 1 \times 10^{33}$ cm^{-2}s^{-1}) and high ($> 1.4 \times 10^{34}$ cm^{-2}s^{-1}) luminosities due to a glitch in the buffer logic in the double column periphery of the
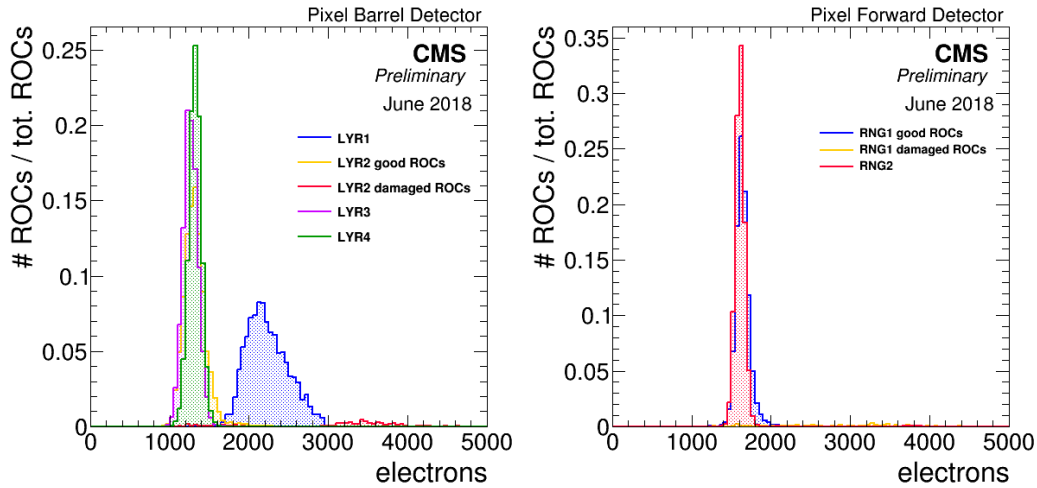
Figure 5: Distribution of readout chip thresholds in number of electrons for the pixel barrel (left) and endcap (right) detectors.

readout chip. This problem has been rectified in the latest version of the chip which is now being used for the construction of the new layer 1 modules, and will be deployed for operation in LHC Run 3 (2021).

Along with an excellent hit efficiency it is also important to maintain a high resolution for hit reconstruction. In Figure 7, the hit residuals are shown as measured using a triplet reconstruction method for layer 3 and disk 2. In this method, hits in the neighboring layers or disks are used to fit a trajectory and then the residuals are calculated in the concerned layer or disk by comparing the expected hit from the fit with the actual one. Residuals are not exactly the same as the intrinsic hit resolution but can be assumed as an upper limit.

Monitoring of the radiation effects on the detector is necessary for smooth operation and plan-
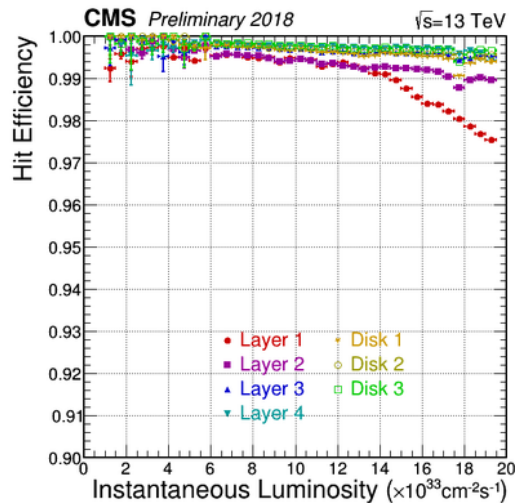

Figure 6: Pixel hit efficiency as a function of the instantaneous luminosity for all barrel layers and endcap disks.
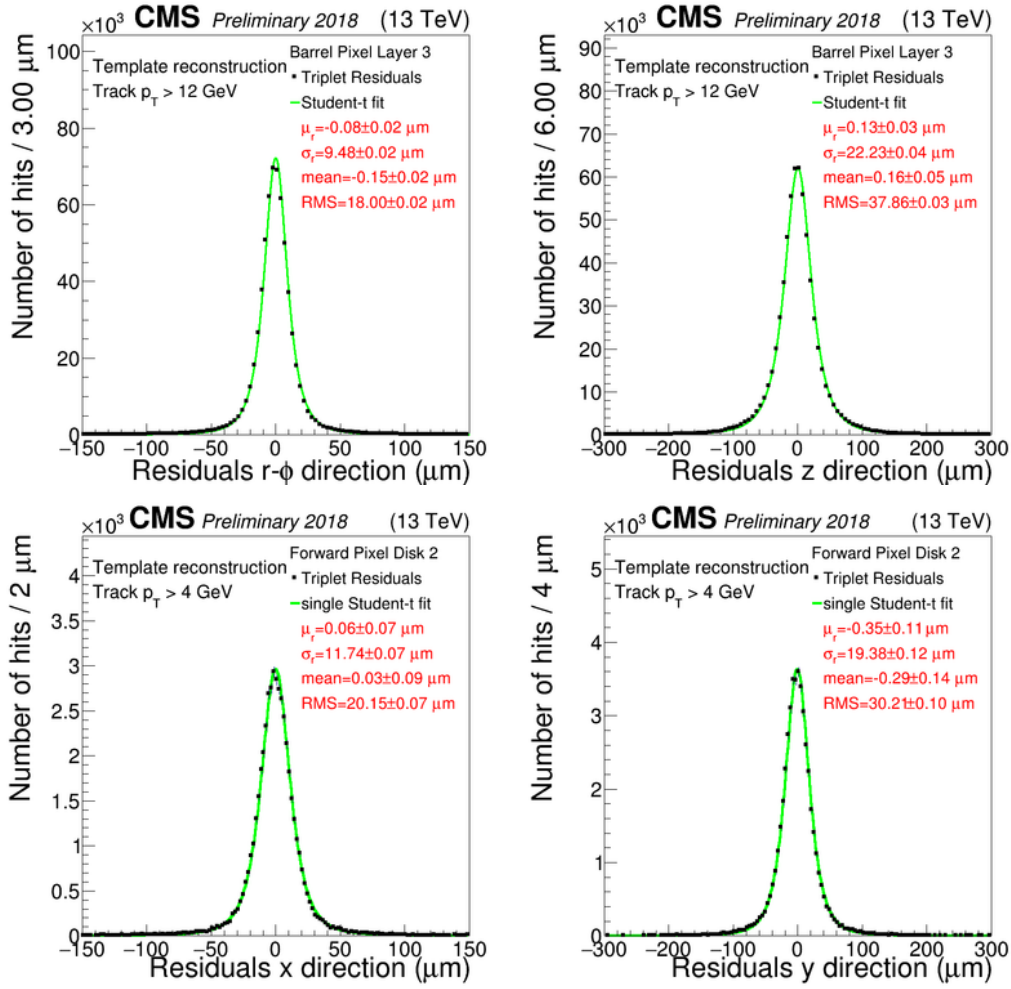
Figure 7: Hit residuals for layer 3 (top) and disk 2 (bottom) using the triplet reconstruction method. The variables *x* and *y* refer to local module coordinates and correspond to the short (100 μm) and long (150 μm) directions of the pixels, respectively.

ning. Figure 8, left, presents the measured mean leakage current for layer 1, which matches well with simulation. Leakage current projections for Run 3 indicate that the leakage current will not hit the limitation from the power supplies. A lot of effort has been put into monitoring the depletion voltage of the silicon sensors. During 2018 operation, bias voltage scans have been performed almost every other week on selected detector modules to monitor the depletion voltage. Figure 8, right, shows the evolution of the depletion voltage for all the four barrel layers as measured using data. The Hamburg model [12] simulates the depletion voltage quite well and predicts a bias voltage of ~800 V for layer 1 by the end of LHC Run 3 operation. The present power supplies need a refurbishment since they can only deliver a bias voltage up to 600 V.
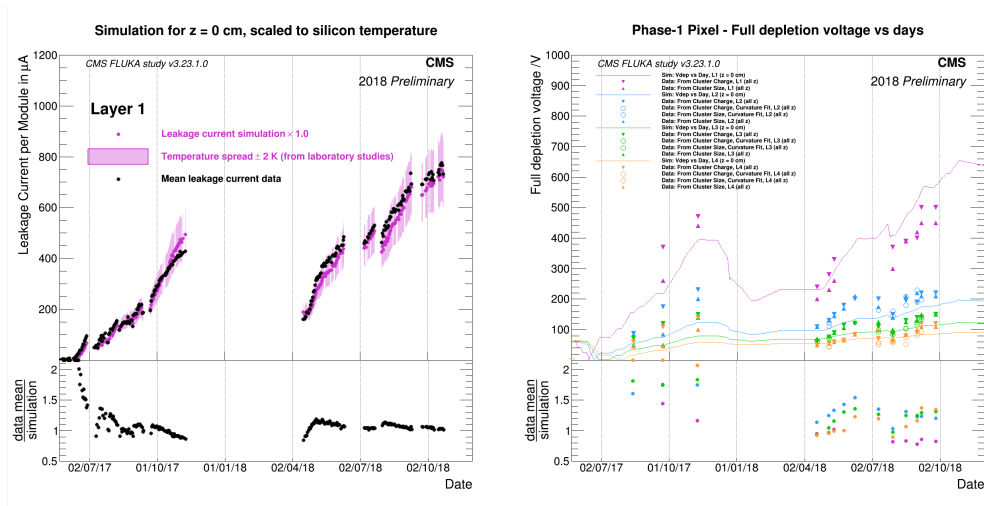
Figure 8: (Left) Measured leakage current for layer 1 and prediction from simulation. (Right) Evolution of the depletion voltage over time comparing the measured (dots) and simulated (line) values.

5. Detector Refurbishment

After the end of Run 2, in December 2018, the LHC entered into a two-year shutdown period (LS2). The detector was extracted in January 2019 from the underground cavern, and stored in a cleanroom facility maintaining a temperature between 0-2 °C to protect the silicon sensors from damage. Figure 9 shows a picture of a cold box, hosting a part of the pixel endcap detector. Temperature and humidity conditions of the boxes are constantly being monitored.

During this LS2 period, several refurbishment projects will be conducted. Layer 1 modules are being replaced with the ones featuring a new ROC and a new TBM. This replacement should fix
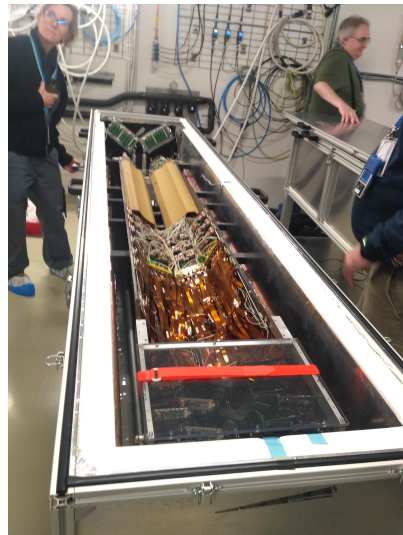
Figure 9: Part of the pixel endcap detector stored inside a cold box after extraction. The storage box is cooled and kept dry to protect the silicon sensors.

the inefficiencies observed in layer 1 and also clear the SEU-induced latching problem in the TBM. Furthermore, all the DC-DC converters will be replaced with new ones featuring a new version of the FEAST chip. Modules in layer 2 that were damaged because they were kept under bias while the DC-DC converters were off will be replaced. Power supplies are being upgraded to deliver a bias voltage of 800 V bearing a maximum current up to 15 mA.

The online software for data acquisition and monitoring is being restructured to improve the maintainability, and additional monitoring capabilities are included. A fast-signal based SEU recovery scheme is being developed to reconfigure the detector at regular intervals. This is expected to perform more efficiently than the software based recovery scheme.

6. Summary

The Phase-1 pixel detector was successfully commissioned and operated in 2018 without encountering any major issues. It was an operational achievement to run the pixel detector without breaking a single DC-DC converter during the year. Overall, the detector delivered good data with high precision tracking and maintaining $> 97\,\%$ hit efficiency across all layers and disks up to an instantaneous luminosity of $2 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The installation of a new innermost layer and other refurbishment projects will ensure an excellent detector until the end of Run 3.

References

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

[2] CMS Collaboration, *CMS Technical Design Report for the Pixel Detector Upgrade*, CERN-LHCC-2012-016, CMS-TDR-11.

[3] H.-C. Kaestli, *Frontend electronics development for the CMS pixel detector upgrade*, Nucl. Instrum. Meth. A 731 (2013) 88.

[4] R. Stringer, *A digital readout system for the CMS phase-1 pixel upgrade*, 2015 JINST 10 C04037.

[5] L. Feld et al., *The DC-DC conversion power system of the CMS Phase-1 pixel upgrade*, 2015 JINST 10 C01052.

[6] J. Troska et al., *Optical readout and control systems for the CMS tracker*, 2003 IEEE Transactions on Nuclear Science, 50 4 pp 1067-1072.

[7] G. Auzinger, *Deployment of the CMS Tracker AMC as backend for the CMS pixel detector*, 2016 JINST 11 C01056.

[8] J. Hegeman et al., *The CMS Timing and Control Distribution System*, 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), San Diego, CA, 2015, pp. 1-3.

[9] CERN, *https://project-dcdc.web.cern.ch/project-dcdc/*

[10] F. Faccio, *https://project-dcdc.web.cern.ch/project-dcdc/public/Documents/Report_IRRAD_tests.pdf*

[11] A. Modak, *Firmware Upgrade of the CMS Phase-1 Pixel Front-End Controller: Impact on Pixel Operation and DAQ*, CMS CR-2019/081.

[12] M. Moll, *Radiation damage in silicon particle detectors – Microscopic defects and macroscopic properties*, PhD thesis, Universität Hamburg, 1999.