PROCEEDINGS OF SCIENCE

Commissioning of the Atlas Pixel Detector

Götz Gaycken for the ATLAS Pixel Collaboration
Physikalisches Institut der Universität Bonn
E-mail: Goetz.Gaycken@cern.ch

After several years of research, development and construction, the vertex detector for the ATLAS experiment at the LHC was finally installed in 2007 and commissioning was started in 2008. Details of the commissioning, the problems which were encountered, and the current status are presented.

17th International Workshop on Vertex detectors
July 28 - 1 August 2008
Utö Island, Sweden

PoS(VERTEX 2008)007

In 2007, the vertex detector for the ATLAS experiment [1] at the LHC was finally installed. After some delay caused by ongoing work on surrounding sub-detectors, the connections were finished and the commissioning was started in the beginning of 2008. An important step of the commissioning is the calibration of the 80 million individually tunable pixel. The calibration procedures are described in the following after a brief description of the vertex detector. Finally, some aspects of the assembly, the installation and the commissioning are presented.

1. The ATLAS Pixel Detector

The ATLAS vertex detector [2] is an 80 million pixel detector with 3 layers at radii of 5, 9 and 12 cm and two times three disks in the forward regions, to provide 3 hit coverage for tracks up to $|\eta| < 2.5$. The pixels are realised on 250 μm thick diffusion oxygenated float zone (DOFZ) silicon. The material undergoes type inversion after $0.7 \times 10^{14}\,n_{eq}$ or 1-3 years. A n^+ in n process was chosen, such that the depletion zone develops from the pixelated side after type inversion when applying the bias voltage of 150 to 600 V. The readout electronic is bump bonded directly on top of each pixel. It provides amplification, zero suppression, charge measurement through time-over-threshold (ToT) and hit buffering. A pixel cell is sketched in Figure 2a. The readout electronics handles the expected occupancy of 2 hits per bunch crossing (40 MHz) and cm^2 at the LHC design luminosity of $10^{34}cm^{-2}s^{-1}$ with its 64 hit buffers at the end of each column. The pixels are grouped in 1744 modules of 46 k pixels organised on bi-staves (barrel) or disks (forward region). To meet the required transverse impact parameter resolution of 15 μm, the pixels have a dimension of $50 \times 400\,\mu$m, where the dimension in direction parallel to the beam of 400 μm is limited by the minimum area needed for the readout electronics. The readout electronics is contained in 16 front-end chips per module, which are bump bonded with two different technologies : lead tin solder bumps (IZM) and indium gold tipped bumps (AMS). To minimise dead zones between front-end chips of one module, there are 50% longer pixels at two sides of the front-end chip. Pixel along the longitudinal axis of a module, which would otherwise not be connected to a front-end chip, are ganged with another pixel. The 16 front-end chips of a module are steered by the module controller chip which is located on each module. It distributes trigger, commands and configuration data to the front-end chips, collects the hits and builds the events. The module controller chips is connect via LVDS lines to opto-boards [3] which perform the translation between electrical and optical signals for 6 or 7 modules. The opto-boards receive and emit the light through bundles of 8 optical fibres, where one bundle is devoted to the incoming signals, and one or two bundles to the outgoing signals. Two bundles are needed for the innermost layer (b-layer), to cope with the data rate. The ca. 100 m long fibre bundles are connected to so called BOCs which connect to the back of one of the 132 VME based readout drivers (RODs) each. The light reception (RX) and emission (TX) on the BOC is performed by replaceable plugins. The ROD propagates trigger, commands or configuration data to the modules. The returned event data is either send by the BOC to the ATLAS event building chain or routed to the four digital signal processors (DSPs) which are contained on each ROD for calibration purposes.

The readout electronics produce approximately 10 kW of heat inside the pixel volume, which is dissipated by an evaporative cooling system [4] sketched in Figure 1 . The liquefied coolant (C_3F_8) is distributed to 88 cooling loops. The aluminium cooling pipes are integrated in the carbon

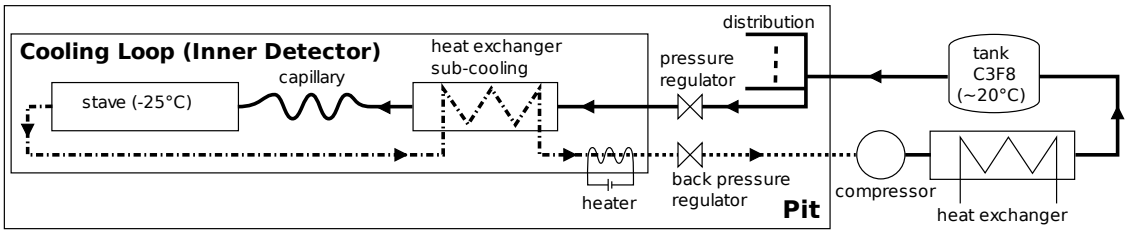

Figure 1: Simplified sketch of the ATLAS evaporative cooling system. The coolant enters the capillary at 13 bar absolute and leaves at 3 bar absolute. This results in a temperature of $-25°C$.

fibre support of the module bi-staves, the disks and the support of the opto-boards. The liquid enters a capillary at the beginning of the cooling loops at a pressure of 13 bar absolute. At the exhaust the back pressure is regulated to ca. 3 bar absolute. To evaporate the remaining liquid, the exhaust lines are equipped with electrical heaters after the heat exchanger. The pressure regulator and the heater are regulated by a programmable logic controller to adjust the cooling loops to a target temperature. At the operation point, the module temperatures are between $-5°C$ and $5°C$, depending on the module position.

2. Calibration

To operate the pixel detector, several registers on the front-end chips and the optical links need to be tuned. There are registers per pixel, per front-end chip and per link. The parameters of all front-end chips of each individual module were calibrated under different environmental conditions, at several different locations and using different setups before the assembly at CERN. These primary tuning parameters are denoted as the production tuning.

The calibration of the 80 million pixels is performed with a distributed system. The calibration is executed by the 132 RODs which are contained in 9 VME crates. Each ROD steers the calibration of up to 26 modules and processes the raw calibration data returned by the modules e.g. histogramming and fitting. Each VME crate is equipped with a single board computer (SBC) which controls the RODs, and which is controlled by a central application.The central application triggers the execution of a sequence of calibration steps. The histograms or tuning parameters which are produced in each calibration step by the RODs are transferred to central servers from where they are available for an automatic analysis and manual inspection. The analysis validates the calibration steps and produces new calibration parameters and summary variables.

2.1 Optical Links

To establish communication with the modules the optical links need to be tuned. The optical links (TX), which are used to transmit the clock, the trigger and commands to the module front-end chips or the module controller chip, encodes commands and clock (40 MHz) using bi phase mark encoding. The encoding, which halves the needed bandwidth, leads to a change of the laser amplitude at least once per clock cycle. This is exploited by the receiver, which is located on the opto-board, to automatically adjust the discrimination threshold and the delay, at which the signal is sampled. To synchronise the trigger distribution to the modules the delay of the sender (TX)

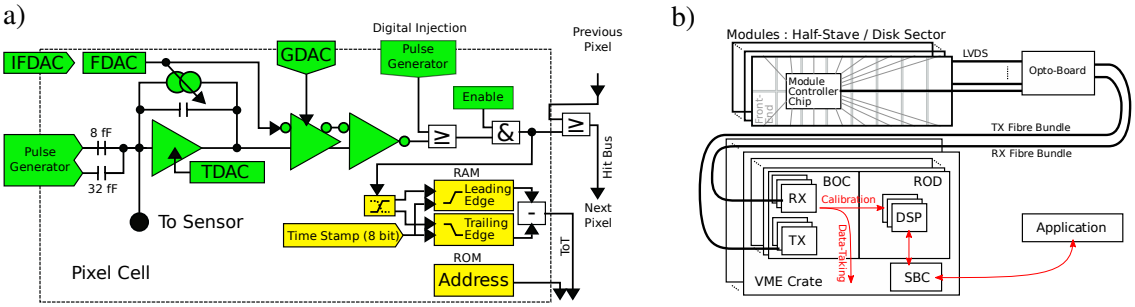

Figure 2: Sketch of a pixel cell a) and the readout system b). Each pixel contains several configurable components (FDAC, TDAC, enable) and some memory cells which contain the pixel address, and store the 8-bit time stamps of the leading and the trailing edge of hits for the calculation of the time-over-threshold (ToT). The hits are buffered in the end-of column buffer. When a trigger arrives, matching hits are pushed to the module controller chip which collects the hits from 16 front-end chips and sends them out to the opto-board. The opto-board sends them via optical links the BOC.

needs to be adjusted according to the fibre length. This can be done with 0.3 ns resolution. For the optical link (RX) which is used to transfer data from the modules, the discriminator threshold and the delay on the receiver need to be tuned manually. Here the raw data bits are sent and there is no guaranteed amplitude change, which could be exploited for an automatic adjustment.

To adjust the RX links, the clock returned by the modules at half the frequency is monitored and the threshold, the delay and the laser power are varied. Finally, an operation point is chosen for which the clock is correctly recovered by the receiver. This method is not sensitive to cases in which a laser slowly ramps up to full amplitude after a long period of silence (slow turn on). The "slow turn on" effect is limited to a small number of laser diodes of the VCSEL[1] arrays of the opto-boards. The second step of the tuning addresses this effect. Instead of returning the clock at half the frequency, the modules sent a pseudo-random pattern which better resembles real data.

2.2 Digital Circuits

After communication has been established, the functionality of the digital circuits is verified. A pulse is injected after the discriminator of each pixel, where one injection step is limited to 3% of the pixels. This test allows to identify faulty front-end chips or individual pixels, which then can be switched off. The test validates the entire readout chain and in particular the tuning of the RX links.

2.3 Threshold

After the communication and the digital circuits have been validated, the analog parts are tuned and tested. The first step is to tune the threshold of the discriminators of each pixel. The threshold is controlled by two DACs, a DAC per front-end chip (GDAC) and one per pixel (TDAC). The DACs are adjusted until each pixel has an occupancy close to 50% when a charge[2] of 4000 e^- is injected into the preamplifier. After the tuning the pixel-to-pixel threshold dispersion amounts to 40 e^-. To verify the tuning, the response of the pixel is measured as a function of the injected charge, where

[1] Vertical Cavity Surface Emitting Lasers

[2] The most probable value of the deposited charge is 20 ke^- for a minimum ionising particle.

the charge is varied from values well below threshold up to values well above (threshold scan). The mean and the width of the resulting Gaussian error curve give the discriminator threshold and the noise of the pixel. Non-ganged pixel show a noise of the order of $170\,e^-$.

The threshold scan is sensitive to the bump connection between the front-end chip and the sensor and the connection of the high voltage. In case of disconnected bumps the noise is slightly decreased and it is independent of the HV. In case of a missing HV connection, the noise increases by a factor of two or more. The two bump bonding technologies show slightly different behaviour due the different resistivity of the bumps which is negligible for the lead tin solder bumps and slightly higher for the indium bumps.

2.4 Cross-Talk Measurement

A more sensitive test for the bump connection is the measurement of the cross-talk between adjacent pixel. A large charge is injected into the preamplifier. A small fraction of the charge will disappear through the bump into the sensor and will be collected by an adjacent pixel. The signal measured by an adjacent pixel is of the order of 1%. In case of disconnected bumps, there is no cross-talk through the sensor. The cross-talk inside the front-end chip itself is negligible.

2.5 Time over Threshold

The next step is the tuning and calibration of the time over threshold (ToT) measurement. The pre-amplifier of each pixel is discharged by a constant current source (feedback current) which leads to a mostly linear discharge. The current is controlled by two DACs, a DAC per front-end chip (IF DAC) and one per pixel (FDAC). The range of the FDAC can be adjusted per front-end chip by a register (IF TRIM). First the pixel-to-pixel dispersion is minimised by tuning the two DACs until an injection of $20\,ke^-$ yields a ToT of 30 bunch crossings (units of 25 ns). For some front-end chips the pixel-to-pixel dispersion is dominated by the outliers. To reduce the number of outliers, the FDAC range is increased by altering the IF TRIM of a front-end chip. Then the tuning is repeated. Since the FDAC step width scales with the FDAC range, the achievable dispersion is degraded for these front-end chips. New FDAC or IF DAC settings impact the threshold tuning. Thus, the latter needs to be refined. Finally, the ToT is measured as a function of the injected charge per front-end chip. The calibration is used in the offline reconstruction.

2.6 In-Time Threshold

Each pixel stores the time of the leading and the falling edge of the discriminator output in units of bunch crossings (25 ns). With decreasing total charge the leading edged crosses the threshold slightly later (timewalk). The correct bunch crossing ID is only assigned to hits whose charge is above a certain minimum threshold, referred to as the in-time threshold. The in-time threshold needs to be determined and can either be used offline to correct the bunch crossing ID of hits with a low charge, in case more than one bunch crossing is read out for each trigger, or it can be used to activate a mechanism on the front-end chip which duplicates all hits with a ToT below a certain threshold and stores them with the bunch crossing ID decreased by one.

3. Assembly

The ATLAS pixel detector was assembled on surface on a 7 m long piece of the beam pipe. The barrel layers and the disks were integrated on the beam-pipe end of January 2007. The next steps were the electrical and cooling circuit connections from the detector to the service panels which are located around the beam pipe in a few meters distance from the interaction region. Each electrical connection was validated first with an electrical circuit tester, then with a series of calibration scans. Since no cooling was available, the calibration scans could only be executed on one front-end chip at the time and for a limited number of pixel (10%) per front-end chip. The opto-links were tuned and the performance of both, the opto-links and the digital circuits, was validated using the method described in Sections 2.1 and 2.2. Finally, the noise, and the HV connection were verified with threshold scans with and without HV. The tests revealed several problems which could be mostly solved. Remaining problems were :

- one module without HV in the intermediate layer,
- modules with 1 broken front-end chip on the outer disks,
- one module with shorted clock on an outermost disk, and
- some modules$\mathcal{O}(4)$ with untunable optical links (RX).

From the system test [5], which took place from October 2006 till January 2007, it became clear that the VCSEL arrays of the opto-board do not always work reliably below an operation temperature of $20°C$. The cooling system which is used for the modules and the opto-boards does not allow for such a high target temperature. Therefore, it was decided to add heaters to the opto-boards and blankets which shield the opto-boards from the environment.

The integration and validation was finished on schedule June 1st 2007.

4. Installation

The pixel detector together with the 7 m long beam pipe was lowered into the ATLAS cavern and inserted into ATLAS at the beginning of July 2007.

In the electrical test that followed, 8 broken HV return lines were discovered. Most likely, the HV connection broke at the location where the cables penetrate into the pixel volume. The problems could be fixed by shorting the redundant return lines outside the detector.

The final connection was delayed until February 2008 to give other sub detectors time to finished their installation. To be ready for the LHC startup, the ID volume had to be closed in April 2008, leaving 2 month to finish the electrical, optical and cooling circuit connections and to validate all connections. In total there are 588 low and high voltage, and sense line connectors with 64 pins each, 4 times 4 fibre bundle arrays of 10 bundles with 16 fibres each on both sides of the detector and 88 cooling circuits.

When the connection of the cooling circuits was started some corrosion was discovered in the exhaust connection pipes. A careful analysis of the pipes surfaced some traces of fluorine, which could be residua from a cleaning oil. The exhaust pipes have a complex geometry to meet

the constraints of the end-plate region. For the bending process , they were filled with sand. It is assumed, that this particular procedure caused the corrosion. No indications of corrosion was found in the other pipes which were machined in a different way. The exhaust pipes were manually rebuild to secure safe operation. The risk for the inlet pipes was considered to be low and no further action was taken. This caused a delay of two weeks.

After the installation of the cooling loops, they were leak tested. The loops were pressurised with helium at 3 bar absolute over several hours. The leak test was sensitivity to pressure drops of $\sim 2\,\mathrm{mbar/h}$. Three leaking loops were identified : one with a leak rate of $150\,\mathrm{mbar/h}$ and the two others with ca. $50\,\mathrm{mbar/h}$. All three loops are cooling disk modules. The leaks are on the inside of the pixel package, either on the inlet or the outlet. In the worst case, this would correspond to coolant loss of $30\,\mathrm{kg/year}$. In 2008, these loops were mostly switched off. It is assumed that these loops can be operated in the future. However, there is the risk that the leaks will increase during operation.

To validate the electrical and optical connections, a series of calibration scans was performed as it was done directly after the assembly : tuning of the optical links, validation of the digital circuits, and a threshold scan. Since the tests were executed before the cooling was available, the tests could only be done on a single front-end chip per module and only for 10% of the pixel of the front-end chips. Six additional modules without HV connection were discovered including one module on the inner most layer (b-layer). The most notable observation was the large dispersion and rather low light which is received on the opto-boards, and a number of dead TX channels. Since then the light which is received on the opto-board was monitored on a regular basis.

5. Commissioning

The initial step of the commissioning was to check the mapping of the cooling controls and the reading, and to verify the regulation of the heater on the exhaust pipes. The cooling circuits were switched on one-by-one and for each circuit and the corresponding pixel modules were switched between three power consumption states : off, $4\,\mathrm{W}$ and $\sim 6\,\mathrm{W}$ per module, where $6\,\mathrm{W}$ corresponds to the power consumption at the end of the lifetime of a module. For most of the circuits, no problems were observed. The heaters regulated their power output within the expectation to compensate for the increased or decreased heat which was produced by the modules in the different states. Though, there were many cooling loops for which the heater regulation did not stabilise at all without powering the modules. These instabilities arise because of an insufficient treatment of the noise of the temperature reading, and because the position of the temperature sensors was not optimal. The cooling system will be refurbished over the winter shutdown in 2008/2009 and these issues will be addressed.

When all the cooling circuits of a quadrant were commissioned, the calibration scans were performed for the first time after the assembly on all front-end chips and pixels. Within the first 6 days 90% of the cooling loops were fully commissioned and calibration scans were performed on half of the detector. In total 9 modules were discovered without HV, and one without clock.

After 6 days of commissioning, four compressors of the cooling plant broke down. It is assumed, that after a restart of the cooling plant, the compressors did not start up correctly, but the magnetic couplers, which limit the torque upon startup, remained slipping for several hours. They

heated up and finally broke apart. Fortunately, when the incident happened all valves closed immediately and no detectable contamination entered the cooling circuits outside of the plant. A large part of the cooling system was disassembled and carefully cleaned. Additional filters and sieves were added to the circuit. The compressors were repaired and equipped with additional diagnostics such that a long time slippage of the torque limiters would be detected. The plant was back in operation July 22, after nearly 3 months. The 6 days of commissioning were considered sufficient to spot fixable problems. So, the end-plates were closed and sealed.

Then the beam-pipe was baked-out to be ready for the LHC start-up in August. Without a full bake-out, the vacuum quality in the interaction region would be several orders of magnitude lower. For ATLAS the bake-out happened in two steps, first the non-NEG coated[3] parts of the beam-pipe would be degassed then the NEG coated parts would be activated. The degassing and activation would be achieved by heating up the corresponding parts to $300^{\circ}C$ and $250^{\circ}C$, respectively. To protect the pixel detector, it was essential to have the cooling system running, in particular the cooling loops of the b-layer. Convection models indicated that the temperature of the modules would stay below $30^{\circ}C$ even in the event of a cooling system shut down which in turn would switch off the beam-pipe heaters. The risk for the pixel detector was considered to be tolerable and the bake-out was executed from July 28 to August 3. No major problems occurred.

After that, the commissioning of the detector was resumed with a strong focus on the preparation for data-taking. The essential steps for data-taking were : tuning of the optical links, identifying non functional front-end chips or modules and masking pixel with high noise to keep the maximum occupancy below the limit. Due to this very short period of commissioning, the program was limited to the tuning of the optical links and the validation. For the module configuration, the so called production tuning was used. Since the time was limited to understand and solve all the problems, approximately 100 modules were not operated for various reasons. The current status is detailed below. The pixel detector joined the ATLAS combined data-taking with cosmic triggers on September 14. During the first beam injections between September 10 and 19, the pixel detector was off to not to be exposed to an unnecessary risk. In total, ATLAS collected 200k and 50k muon tracks with magnetic field on and off, which traversed the pixel volume and contained at least one pixel hit. The combined data-taking was stopped October 20. From then on, the focus shifted towards the understanding of the remaining issues and improving the tuning. A second data-taking period was started end of November. Here only the tracking detectors participated.

6. Detector Status

TX plugins

The light output of the TX plugins has been monitored regularly. Up to now, TX channels keep on dying with about one channel per week. The issue has been investigated in the laboratory and together with the producer. It is assumed that the VCSEL arrays were damaged by electrostatic discharges during the assembly. The theory has been tested in the laboratory with the conclusion, that electrostatic discharge could produce the observed effect. New TX plugins have been produced and assembled more carefully. Currently, 80 of these plugins are in operation since a few weeks

[3] 5μ sputtered coating of Ti-Zr-V

and up to now no dead channel has been discovered on this new production. Currently, there are seven dead TX channels and two more candidates which have not been confirmed yet. Most of them will get fixed by replacing the TX plugins, which is expected to happen in April 2009. For two channels the problem does not seem to be located on the plugin.

RX Opto-links

After the first step of the opto-link tuning (see Section 2.1) 96% of the opto-links could be operated. The second step of the opto-link tuning was fully exploited, after the first round of data-taking. In the first step, the laser was tuned sometimes to too high output power and the slow-turn-on was not taken into account. This was corrected by the second step. The fraction of functional links increased to 99.7%. The remaining 6 links have no operation point which allows an error free transfer. So far, the opto-links have only been operated in the 40 Mbit/s mode. The operation with 80 Mbit/s still needs to be validated.

Digital Circuits and Analog Circuits

After all optical links were tuned, the validation (see Section 2.2) of the digital circuits still failed on a few modules. The problems were traced back to failures on single front-end chips. There are 3 modules with two bad front-end chips and 18 with one and 1 module with a bad column pair in one front-end chip. These modules can be operated after disabling the corresponding front-end chips or pixel. Then there are 12 modules which are not operated due to shorted low voltage supply or clock or high analog or digital currents, or which loose the clock. Two of these modules are still under closer investigation. In total, 0.8% of the pixels are not operated due to problems in the digital or analog circuits or their supply.

Analog Circuits and Connection to Sensor

The next step is to validate the analog parts of the front-end chips (see Section 2.3). Now, charge is injected into the pre-amplifier. The injected charge is varied broadly around the threshold, to measure the selected discriminator threshold and the noise of each pixel. The average threshold per module and the threshold dispersion is shown in Figure 3. After tuning a threshold dispersion of $40\,e^-$ was achieved. Threshold scans exist for 97% of the modules. Some of the modules could not be scanned since the corresponding cooling loops were mostly off, for some modules no scan result exist due to remaining instabilities in the calibration software, and finally some modules cannot be operated due to the reasons mentioned above. No error curve was measured for 2‰ of the pixel, including the broken front-end chips. These pixel cannot be tuned and are most likely dead. They are concentrated on 25 front-end chips. The remaining pixel are distributed randomly. Though there is a systematic concentration at the edges of the modules, and slightly less pronounced, at the edges of front-end chips. For $5 \cdot 10^{-4}$ of the pixels the threshold is more than $500\,e^-$ away from the target threshold of $4000\,e^-$. Looking at the noise determined from the threshold scan, there are 9 modules without HV, which were discovered already directly after the pixel detector was connected. The cross talk scan (Section 2.4) discovered a similar picture. The bumps of 10^5 pixel are disconnected and they are mostly randomly distributed over the modules, but there is also a systematic concentration at the edges of modules and front-end chips. A quantitative correlation between the two scans is still to be done.

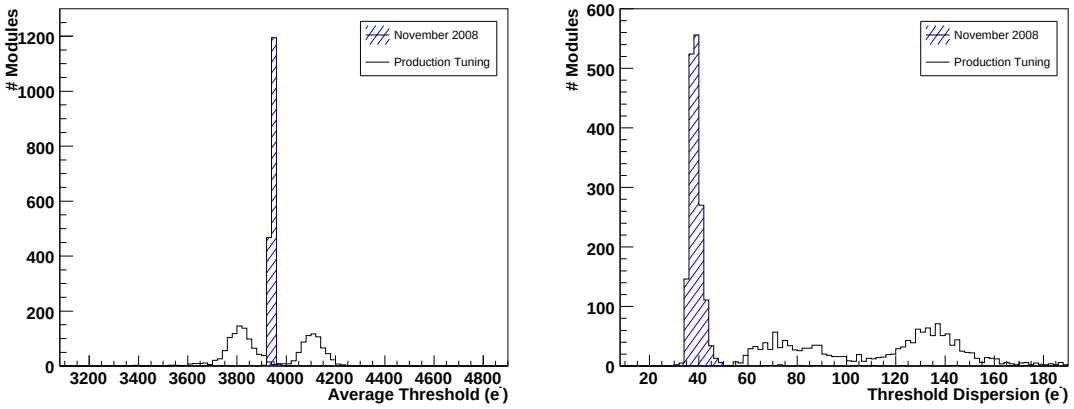

Figure 3: The average threshold and threshold dispersion per module when using the production tuning and the tuning of November 2008. The production tuning was done in different locations with different setups and in a different environment. The two peaks in average threshold tuning are due to two different target thresholds.

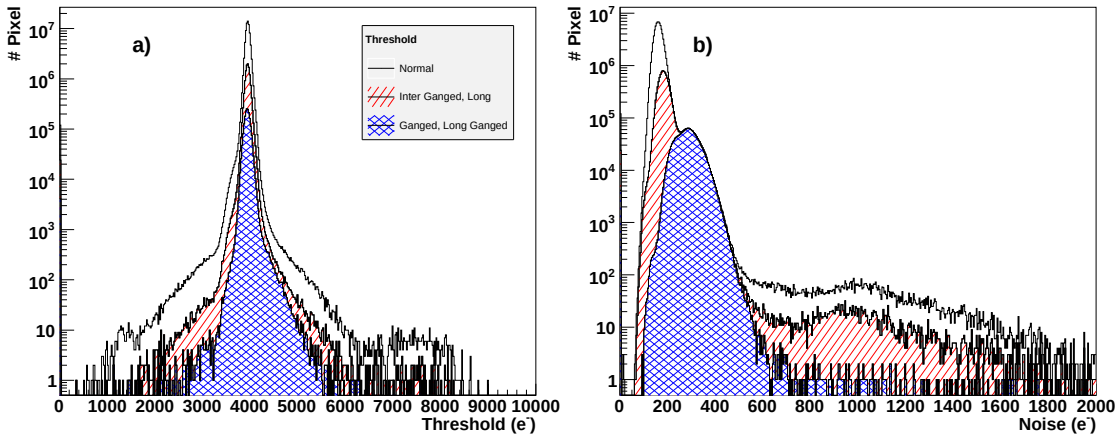

Figure 4: The threshold a) and the noise b) determined from a threshold scan. The scan was performed on 96% of the modules. This also excludes the 9 modules without HV. No response was received for 2‰. The threshold of $5 \cdot 10^{-4}$ pixel differs by more then $500\,e^-$ from the target threshold.

After the threshold tuning, the ToT was tuned and calibrated (see Section 2.5). After tuning, the ToT dispersion is 2%. The analysis of the calibration is still on going.

Occupancy with Pseudo Random Triggers

Finally, the tuning is tested with pseudo random triggers. Until noisy pixel are switched off, the readout system will run easily into saturation. Therefore initially data is taken with low trigger frequency of e.g. 100 Hz. To determine the noise occupancy, pixel with an occupancy higher than 10^{-5}hits/trigger are switched off. The noise occupancy when using the production tuning is shown in Figure 5. About 0.1‰ of the pixel had to be switched off. Then the rate was increased up to 2 kHz. The maximum design readout rate of ca.100 kHz was achieved in a short test, but up to now

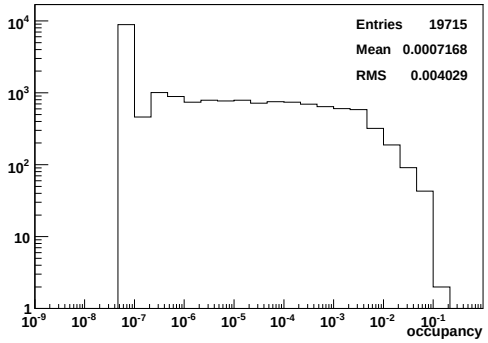

Figure 5: Hit occupancy per trigger with pseudo random triggers when using the production tuning. The hit occupancy was determined from 10^7 events. Pixel with an occupancy larger than $10^{-5}$are switched off.

this was only tried on part of the pixel detector. The data rate for recording events is limited to a few hundred Hz.

Summary

The ATLAS pixel detector was finally assembled, installed, connected, and is being commissioned. After an intense commissioning phase, which is still on going, most pixels of the detector have been tuned and their status has been determined. Currently, there are 12 modules which cannot be operated due to problems with the optical links. For two more modules the problem is unconfirmed. This number will decrease to 10 or 8 after replacing the faulty TX plugins. There are 11 modules which cannot be operated due to low voltage problems, one with a shorted clock and 9 modules which do not receive high voltage. In total 35 modules (2%) cannot be used. Of those modules there are 5 located in the inner most layer (1.7%). Finally, there are 36 modules (2%) which are currently not operated because the corresponding cooling loops are switched off. The location of the non-working modules is visualised in Figure 6. Approximately 2‰ of the pixel on the functional modules are presumably dead, most of them are concentrated on 25 front-end chips (including the 24 broken front-end chips). The threshold and time-over threshold were tuned for all functional pixel. After tuning, the dispersion amounts to $40\,e^-$ for the threshold and 2% for the time-over-threshold.

In a period of several weeks the entire ATLAS detector collected ca. 250k tracks with at least one pixel hit on track. More data is currently collected for the inner tracking detectors only. The hit efficiency looks as expected. Though, the analysis has not been finished yet.

References

[1] **ATLAS** Collaboration, G. Aad *et. al.*, *The atlas experiment at the cern large hadron collider*, *JINST* **3** (2008) S08003.

[2] **ATLAS** Collaboration, M. S. Alam *et. al.*, *Atlas pixel detector: Technical design report*, *CERN-LHCC-98-13*.

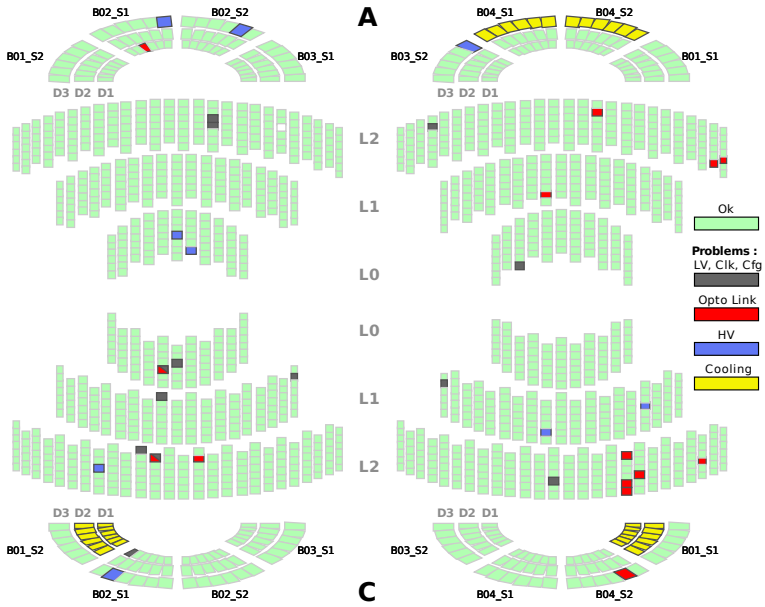

Figure 6: Location of problematic pixel modules in a false perspective, exploded view. There are 11 modules with shorts of the low voltage supply (LV) or clock (Clk), or which cannot be configured or looses their configuration (Cfg). There are 9 modules without high voltage (HV), and for currently 12 or 14 modules the transmission to (TX) or from (RX) the module via the opto-link is broken. The faulty TX links can be mostly repaired by exchanging the TX plugins. This will reduce the number to 8. Finally, 3 cooling loops with small leaks were not operated. This effects 36 modules. It still has to be seen whether these modules can be used.

ATLAS Collaboration, G. Aad *et. al.*, *Atlas pixel detector electronics and sensors*, *JINST* **3** (2008) P07007.

ATLAS Collaboration, G. Aad and other, *The atlas pixel detector mechanics and service*, *to be submitted to JINST* (2008).

[3] K. K. Gan *et. al.*, *Radiation-hard opto-link for the ATLAS pixel detector*, *Nucl. Instrum. Meth.* **A554** (2005) 458–468.

[4] D. Attree *et. al.*, *The evaporative cooling system for the ATLAS inner detector*, *JINST* **3** (2008) P07003.

[5] *ATLAS Pixel Detector Commissioning using Cosmics Rays*, Sept., 2007. 16th International Workshop on Vertex detectors.