UNIVERSITÄT BONN
Physikalisches Institut

System Test and Noise Performance Studies
at
The ATLAS Pixel Detector

von
Jens Weingarten

Abstract: The central component of the ATLAS Inner Tracker is the pixel detector. It consists of three barrel layers and three disk-layers in the end-caps in both forward directions. The innermost barrel layer is mounted at a distance of about 5 cm from the interaction region. With its very high granularity, truly two-dimensional hit information, and fast readout it is well suited to cope with the high densities of charged tracks, expected this close to the interaction region. The huge number of readout channels necessitates a very complex services infrastructure for powering, readout and safety. After a description of the pixel detector and its services infrastructure, key results from the system test at CERN are presented.

Furthermore the noise performance of the pixel detector, crucial for high tracking and vertexing efficiencies, is studied. Measurements of the single-channel random noise are presented together with studies of common mode noise and measurements of the noise occupancy using a random trigger generator.

ii

UNIVERSITÄT BONN
Physikalisches Institut

**System Test and Noise Performance Studies
at
The ATLAS Pixel Detector**

von
Jens Weingarten

ii

System Test and Noise Performance Studies
at
The ATLAS Pixel Detector

Dissertation

zur

Erlangung des Doktorgrades (Dr. rer. nat.)

der

Mathematisch-Naturwissenschaftlichen Fakultät

der

Rheinischen Friedrich-Wilhelms-Universität

zu Bonn

vorgelegt von

Jens Weingarten

aus

Neuwied

Bonn 2007

ii

Angefertigt mit Genehmigung der Mathematisch-Naturwissenschaftlichen Fakultät der Universität Bonn

Ich versichere, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt sowie Zitate kenntlich gemacht habe.

Contents

Chapter 1

Introduction

The search for the fundamental structure of matter has been driving scientists and philosophers since the fifth century BC. Nowadays one describes this fundamental structure and the laws governing the interactions of matter in the *Standard Model* of particle physics. It groups matter into two categories of particles, and the interactions between these particles into three fundamental forces.

The Standard Model has been thoroughly tested and while its predictions hold true to a very high precision for almost all properties of particles and interactions the one shortfall of the Standard Model is that the particles are assumed to be massless. To explain the origin of mass a mechanism has been proposed by the British physicist Peter Higgs. The Higgs mechanism predicts the existence of a massive boson, the so called Higgs boson.

To date physicists have no direct evidence for the existence of the Higgs boson and thus for the Higgs mechanism. This is one of the central tasks for the two multi-purpose particle detectors, ATLAS and CMS that will be operated at the CERN proton-proton Large Hadron Collider (LHC) in Switzerland. The LHC and its experiments will enable physicists to search not only for the origin of mass but also for further fundamental symmetries of nature and new laws of physics that are proposed to solve deficiencies of the standard model, like *Supersymmetry* or extra dimensions of space-time. This will provide deep insight into Nature and will significantly influence our view of the world.

The ATLAS pixel detector will play a crucial role in this enterprise as it provides very precise measurements of the trajectories of particles that are created in the collision of two protons, allowing an accurate determination of the point in space where these particles originate. This is necessary to separate jets of particles that stem from heavy flavor quarks, like b-quarks, from

those that are created in the decay of light quarks (u-,d-,s- and c-quarks). The pixel detector is built in three concentric barrel layers and three disk layers in the forward/backward directions. It consists of 1744 identical units called modules. A module is a hybrid assembly of a silicon sensor and 16 readout chips that are connected to every pixel cell on the sensor. With this layout and the segmentation of the sensor in 50 x 400 $\mu\mathrm{m}^2$ pixels it provides 80 million readout channels.

To operate a detector system of this size an intricate services infrastructure is needed. Power supplies and off-detector electronics, monitoring, interlock and cooling systems have been brought together in a large scale system at CERN for the first time. During the subsequent system test the interplay of the numerous complex components was studied. A lot of experience in operating the setup, that roughly corresponds to 10% of the final detector, was collected. This was the first step towards comissioning the pixel detector system. The system test culminated in several successful data taking runs, measuring cosmic muons.

The second topic of this thesis is the noise performance of the pixel detector. Controlling the noise in the readout electronics was one of the most stringent requirements during the development of the detector. Keeping the noise level low is very important, as high noise not only leads to inefficiencies due to disabled pixels and to possible loss of events because of too high data rates. High noise occupancy during data taking also decreases the single-point resolution, deteriorating the vertexing and b-tagging performances. Therefore measurement techniques for common mode noise in the pixel detector were developed and tested, and extensive studies of the noise performance were carried out during the system test.

In chapter 2 an overview is given of the LHC, the ATLAS detector with its sub-detector components, and the physics that will be studied there. Chapter 3 gives a detailed description of the pixel detector, including the pixel modules, readout and monitroing systems, power supplies and cooling. The system test mentioned above was used for qualification measurements of the system performance. The measurements and results are presented in chapter 4. Chapter 5 is dedicated to measurements of the noise performance of the pixel modules in stand-alone tests and during the system test. Data taking runs using random triggering were studied to determine the noise occupancy in the system and the characteristics of the noise hits produced. The noise performance was also studied during a cosmics data taking run. Chapter 6 gives a summary of the results presented in this thesis.

Chapter 2

LHC and the ATLAS detector

2.1 Physics at the LHC

In this chapter two of the most important physics goals are shortly mentioned to set the stage for a description of the Large Hadron Collider and the ATLAS detector.

2.1.1 The standard model of high-energy physics

The Standard Model of particle physics describes three of the four fundamental interactions between the elementary particles which make up matter. It is a quantum field theory developed in the early 1970s, which is consistent with both quantum mechanics and special relativity. It is invariant under local $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge transformations.
The Standard Model has been thoroughly tested during the last 30 years, and most of the measurements agree with the model predictions to an unprecedented accuracy. This is shown in figure 2.1 exemplarily for the measurement of the Z boson mass.

The basic constituents of matter are spin-1/2 fermions interacting through the exchange of spin-1 force-carrier particles. Due to their integer spin these are bosons. The fermions are grouped into three generations of two leptons (particles not interacting strongly) and two quarks (interacting via all four forces) each. The fermion generations are summarized in table 2.1. These twelve fundamental particles, together with their respective anti-particles (having the same mass but opposite electric charge), form the complete matter content of the Standard Model. Stable everyday matter is made of particles of the first family, while particles of the latter families can only be created artificially in high-energetic collisions.

Figure 2.1: Data from LEP and other e$^+$e$^-$ experiments agree perfectly with the predictions of the Standard Model. [1]

Family I		Family II		Family III		
symbol	mass [MeV]	symbol	mass [GeV]	symbol	mass [GeV]	charge [e$^-$]
leptons						
ν_e	$< 3 \times 10^{-6}$	ν_μ	$< 0.19 \times 10^{-3}$	ν_τ	0.0182	0
e	0.511	μ	0.1057	τ	1.777	-1
quarks						
u	1.5 - 3	c	1.25$\pm$0.09	t	174.2 $\pm$ 3.3	+2/3
d	3 - 7	s	0.095 $\pm$ 0.025	b	4.2 - 4.7	-1/3

Table 2.1: The three generations of leptons and quarks with their most important properties.

The force-mediating particles are the photon, $W^\pm$ and Z^0 for the electroweak interaction, coupling to the electric charge and the weak hypercharge respectively, and eight gluons for the strong interaction, coupling to the color charge. The couplings between the particles of the Standard Model are shown in figure 2.2.

One of the drawbacks of the Standard Model is that all particles are massless. In an extension to the Standard Model, first proposed in 1964 by the British physicist Peter Higgs, mass is given to the particles by coupling to a scalar field. The Higgs field breaks electroweak symmetry spontaneously, due to its non-zero vacuum expectation value. The Higgs mechanism predicts the existence of a massive spin 0 boson, called the Higgs boson, which has

Figure 2.2: Interactions between the fundamental particles of the Standard Model. [2]

yet to be discovered. As electroweak precision measurements and theoretical predictions prefer masses of the Higgs boson of ~ 1 TeV, one of the main goals of the LHC collider is the direct discovery of the Higgs.

2.1.2 Supersymmetry

As we have seen the Standard Model is very successful. Although its predictions have proven to be very accurate it still leaves plenty of questions unanswered [3]:

- Why are some fundamental particles massless, while others have masses up to the weight of a gold nucleus in the case of the top quark?

- Why are the masses of the fundamental particles so small compared to the fundamental mass scale of particle physics, the Planck scale $M_P \sim 10^{19}$GeV ?

- Are the different fundamental forces unified?

- Why are there so many types of 'elementary' particles? Could they all be made out of more fundamental constituents?

- How to explain all the different parameters of the Standard Model: 6 quark masses, 3 charged-lepton masses, 2 weak-boson masses, 4 weak mixing angles and phases, 3 interaction strengths and a non-perturbative strong-interaction vacuum parameter, not to mention parameters describing neutrino oscillations?

- How do the values for the particle quantum numbers, like electric charge Q, weak isospin I, hypercharge Y and color, come about?

These questions necessitate the search for an underlying theory that would present itself in the form of what is referred to as Physics Beyond the Standard Model or New Physics.

At present the most promising model of physics beyond the Standard Model is Supersymmetry. Firstly, it provides answers to the hierarchy problem as will be shown in the next section. Secondly, it yields a development of the coupling strengths for the electro-magnetic, the weak and the strong interactions such that the forces unify at high energy scales. Finally, Supersymmetry seems to be essential for the consistency of string theory, which at the moment is the best candidate for an underlying theory that includes gravity.

Introducing a new symmetry between bosons and fermions, Supersymmetry postulates a superpartner for every fundamental particle. These *sparticles* have the same masses as well as the same couplings and quantum numbers as the standard model particles, except for their spin, which differs by 1/2. Thus, one fermion and one boson each are grouped into what is called a *supermultiplet*. It is not possible for the fermions and bosons of the Standard Model to be superpartners, because their various charges differ. Therefore unknown particles must exist, that form the superpartners to all Standard Model particles. None of the superpartners have been observed so far, which implies that Supersymmetry must be broken.

The Minimal Supersymmetric Standard Model (MSSM) is the minimal extension to the Standard Model that realizes Supersymmetry. It was originally proposed in 1981 to solve the hierarchy problem. It introduces R-parity and adds supersymmetry breaking by introducing explicit soft supersymmetry breaking operators into the Lagrangian.
The MSSM is the leading candidate for a new theory to be discovered at the LHC. It solves the hierarchy problem quite naturally, it provides running coupling constants that unify at high energies, and it provides a promising candidate for cold dark matter (CDM) particles. If R-parity is preserved then the lightest supersymmetric particle (LSP) is stable and it is a weakly interacting massive particle (WIMP) making it a very good CDM candidate.

The hierarchy problem

A lot of theoretical models predict new physics to appear at energy scales of the order of 1 TeV and above. From the precision measurements done at

LEP[1] we know that the Standard Model Higgs boson should have a mass between 114.4 and 1000 GeV. Supersymmetrical loop corrections to the masses of the Higgs and the $W^\pm$ masses as well as a possible unification of the coupling strengths of the three forces at the Planck scale indicate supersymmetric particles weighing about 1 TeV.

One of the main theoretical reasons to expect Supersymmetry at an accessible scale of about 1 TeV is provided by the hierarchy problem:
Why is $m_W << M_P$?
This question arises when considering that m_W and the mass of the Higgs boson m_H contain quantum corrections that are quadratically divergent in the Standard Model:

$$\delta m^2_{H,W} \simeq O\left(\frac{\alpha}{\pi}\right)\Lambda^2 \tag{2.1}$$

If the cutoff parameter Λ, which represents the scale where new physics appears, is comparable to the Planck scale, these corrections are huge ($>> m^2_W$). The fine-tuning required to obtain the measured W mass and the expected Higgs mass on the order of 100 GeV seems quite unnatural. A mechanism that keeps the corrections comparable to the physical masses is preferred by most physicists.
Supersymmetry provides such a mechanism. Since the numbers of bosons and fermions, that contribute to the corrections are equal, and bosonic and fermionic loops have opposite sign, the residual one-loop correction is:

$$\delta m^2_{H,W} \simeq O\left(\frac{\alpha}{\pi}\right)(m^2_B - m^2_F), \tag{2.2}$$

which is $\le m^2_{H,W}$ if the supersymmetric partners have similar masses:

$$\left|m^2_B - m^2_F\right| \le 1\,\mathrm{TeV}^2 \tag{2.3}$$

2.2 The Large Hadron Collider

The cross section of interesting new physics processes decreases at higher energies with E^{-2}_{CM} [1]. Therefore the luminosity of particle accelerators has to increase with E^2_{CM}. Comparing to the performance achieved by LEP, with a maximum luminosity $\mathcal{L} \sim 10^{32}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at a maximum energy $E_{CM} \sim 200$ GeV, a TeV scale collider should have a luminosity $\mathcal{L} \sim 10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

---

[1]Large Electron Positron collider

This is where the LHC[2] enters the stage. Colliding bunches of protons with a center of mass energy $\sqrt{s} = 14$ TeV, it is designed to reach a luminosity of 10^{34} cm^{-2}s^{-1}. This is achieved by colliding bunches of 1.15×10^{11} protons each at a rate of 40 million collisions per second.

The LHC is constructed at the European Laboratory for Particle Physics (CERN) in a tunnel between 50 and 175 m underground with a circumference of 26.7 km. The tunnel, which formerly housed the LEP collider, now accommodates the very complex LHC two-in-one superconducting magnet structure that surrounds the two separate beampipes (see figure 2.3). Each of these beampipes is filled with 2808 bunches of protons, resulting in a total beam current of 0.582 A and a total energy of 664 MJ stored in the collider. Dumping this amount of energy in an instant would evaporate an equivalent of 300 kg of water.


Figure 2.3: Cross sectional view of a LHC dipole magnet. [4]

The collider consists of 1232 main dipoles, which are cooled to 1.9 K using superfluidic helium. They provide magnetic fields of up to 8.4 T. The beams are injected from the SPS[3] at an energy of 450 GeV and are subse-

---

[2]Large Hadron Collider
[3]Super Proton Synchrotron

quently accelerated by superconducting RF[4] cavities operating at a comfortable 4.5 K [5]. There are eight single-cell cavities per beam, each delivering 2 MV (5.3 MV/m) accelerating voltage at 400 MHz. The cavities are situated at Point 4 and will accelerate the beams up to 7 TeV in a ramp-up phase that takes about half an hour. As about 10^{-9} of the total beam power is sufficient to quench a superconducting dipole the beam has to be very well collimated. Therefore sophisticated collimator systems are situated at Points 3 and 7. The beam dump system at Point 6 is able to remove the two beams from the collider in a single turn. Within 86 μs the beam energy is dumped into two external graphite beam dumps. This will be done in regular intervals about every 10 hours, when the beam quality is deteriorated, or whenever quenches in the magnets or other malfunctions happen.

Four experimental caverns are located at Points 1, 2, 5 and 8. The two multi-purpose experiments ATLAS and CMS are designed to measure a wide variety of physics processes and will be looking for Higgs bosons, SUSY particles, heavy W' and Z' particles and the likes. The LHCb experiment will concentrate on the rich field of b-physics, while ALICE will analyze heavy ion collisions, e.g. in Pb-Pb beams that can be accelerated to 5.5 TeV/nucleon with a luminosity of 10^{27}cm^{-2}s^{-1}.

2.3   The ATLAS detector

The ATLAS detector [6] has been designed to exploit the full physics potential of the LHC. This includes the search for the origin of mass, searches for heavy W- and Z-like objects, supersymmetric particles, for compositness of the fundamental fermions as well as the investigation of CP violation in B-decays and detailed studies of the top quark. The ability to cope well with a broad variety of possible physics processes is expected to maximize the detector's potential for the discovery of new, unexpected physics.
A lot of different benchmark processes were used in formulating the basic principles of the detector. These can be summarized as following:

- Very good electromagnetic calorimetry for electron and photon identification and measurements, complemented by hermetic jet and missing transverse energy calorimetry.

- Efficient tracking at high luminosity for lepton momentum measurements, b-quark tagging, enhanced electron and photon identification,

---

[4]radio frequency

Figure 2.4: Schematic view of the LHC collider complex.

as well as tau and heavy-flavor vertexing and reconstruction capability of some B decay final states at lower luminosity.

- Stand-alone, precision, muon-momentum measurements up to highest luminosity, and very low-p_T trigger capability at lower luminosity.

2.3.1   Detector subsystems

The overall detector layout is rather typical for collider experiments in high-energy physics. A central tracking detector close to the interaction point is situated in a 2T solenoidal magnetic field to allow momentum measurement of charged particles. It is surrounded by an electromagnetic and a hadronic calorimeter system. These are followed by a magnet system generating a toroidal magnetic field which allows an independent momentum measurement of the muons measured in the muon chambers.

The Inner Detector

The Inner Detector *(ID)* [8] is situated in a 2T solenoidal magnetic field, generated by about 10 km of superconducting cable, wound into a coil at the inner surface of the calorimeter cryostat. The outer wall of the cryostat acts

Figure 2.5: Schematic view of the ATLAS detector. The detector subsystems are indicated. [7]

as a return yoke for the magnetic field. The coil, weighing about 4 tons, is cooled to 4.5 K by liquid helium.

The Inner Detector is contained within a cylinder of about 7m length and 1.15m radius. Pattern recognition, momentum and vertex measurements, and enhanced electron identification are achieved through a combination of high-resolution pixel and strip (SCT) detectors in the inner part and continuous straw-tube tracking detectors with transition radiation capability (TRT) in the outer part of the tracking volume.

- **The Pixel Detector**
  The innermost sub-detector of the Inner Detector is the pixel detector. Being the main topic of this thesis it is described in full detail in chapter 3.

- **The SemiConductor Tracker** *(SCT)*
  The microstrip detector [9] forms the middle layer of the Inner Detector. It consists of 4 concentric barrel layers and two end-caps of nine disks each. The barrel layers carry 2112 detector modules, while 1976 end-cap modules are mounted on the disks. The whole SCT occupies a cylinder of 5.6 m in length and 56 cm in radius with the innermost

layer at a radius of 27 cm. Its geometry is designed in a way that the
SCT on average provides four precision space-points per track up to
$|\eta| \leq 2.5^5$.

The silicon modules consist of one or two pairs of single-sided p-in-
n microstrip sensors glued back-to-back at a 40 mrad stereo angle to
provide two-dimensional hit information. Figure 2.6 shows an exploded
schematic view of a module with two pairs of sensors. The 285 μm thick
sensors are divided into 768 AC-coupled strips with an active length of
123.2 mm and a pitch of 80 μm for the barrel modules and 57 - 94 μm for
the wedge shaped end-cap modules. The sensor pairs are mounted on
a thermally highly conductive carbon baseboard that provides cooling.
The SCT sensors are operated at -7°C to prevent reverse annealing of
radiation damage. Modules in the barrel part are all identical while
the end-cap modules come in four different geometries, depending on
their radial position on the disks.

Figure 2.6: Exploded schematic view of a SCT module. [10]

The detector has an active area of 61 m^2 of silicon detectors, with 6.2
million readout channels. The spatial resolution is 16 μm in Rϕ and
580 μm in z direction. Tracks can be distinguished if separated by more
than $\sim 200\mu$ m.

The signals from the sensors are read out by binary front-end elec-
tronics, data is converted into optical signals directly on the module.
One incoming optical link is used for configuration and triggering while
data is transmitted via two outgoing links. One driver is sufficient for

---

^{5}pseudorapidity $\eta = -\ln(\tan(\theta/2))$

data transmission, thus ensuring redundancy and failure safety. The off-detector electronics is similar to that used in the pixel detector (see chapter 3.3.4).

- **The Transition Radiation Tracker *TRT***
  The TRT is the outermost part of the Inner Detector. It provides charged-particle tracking, using 420 000 readout channels, as well as electron identification through transition radiation measurements.
  The TRT consists of a central barrel and three end-caps on either side, referred to as end-caps A,B and C with the installation of the end-caps C being postponed. With this geometry the TRT delivers 36 hit measurements on average up to $|\eta| \leq 2.1$.
  Straw drift-tubes, as shown in figure 2.7, are used for tracking. The

Figure 2.7: Drawing of a TRT straw. [11]

straws have a diameter of 4 mm and a length of 39 cm in the end-caps A and B, and a length of 144 cm in the barrel. They are operated in proportional mode in a gas mixture of 70% Xe, 27% CO_2 and 3% O_2, with an avalanche gain of 2.5×10^4. The sense wires are 30 μm in diameter, gold-plated tungsten wires. The straw wall is held at a voltage of -1530 V with respect to the sense wire. A single-point resolution of 170 μm in rϕ direction can be achieved.
Particle identification is done by efficiently converting the transition radiation photons, that are emitted when a charged ultra-relativistic particle crosses the interface between media with different dielectric constants ϵ. As the number of photons emitted per transition is on the

order of $\gamma\alpha_{em}$, with α_{em} being the electro-magnetic coupling constant and $\gamma = E/m$, the number of transitions needs to be large. The energy of the emitted photons depends on the materials used as it is proportional to the squared difference of the plasma frequencies, $(\omega_1 - \omega_2)^2$. For the TRT the energy is of the order of 10 keV as the materials are polypropylene (in fibers or foils to increase the number of transitions) and CO_2 gas. Thus the energy deposition in the TRT is the sum of ionization losses of charged particles ($\sim$2 keV on average) and the larger deposition due to transition radiation photons (typically >5 keV).

Two thresholds are applied to the signals from the straws. A low threshold at about 250 eV to detect all hits, and a high threshold around 5 keV for the transition radiation hits. The high threshold is optimized to distinguish between electron and pion hits as electrons produce more high threshold hits than pions.

In order to reduce occupancy in the barrel the straws are divided in two halfs that are read out separately at either end of the barrel. To provide a z coordinate the drift time in each straw is measured. In test-beam studies a maximum drift time of 48 ns was measured, resulting in a resolution of $\sim$130μm with an efficiency of $\sim$87% at a threshold of 250 eV. Single-straw hit efficiencies have been measured at 96%.

The calorimeter system

An excellent knowledge of the electron or photon energy is necessary for precision measurements like couplings within and beyond the SM, or to resolve possible narrow resonances of new particles over a large background. A good energy resolution and a good linearity are needed for energies ranging from a few GeV up to a few TeV.

An important example is the possible discovery of the Higgs boson [13]. If the Higgs mass is below 130 GeV the decay H $\rightarrow \gamma\gamma$ is the most promising discovery channel. If the Higgs mass is larger and, in particular if it is at least twice the mass of the Z^0-boson, the Higgs can be discovered in the H$\rightarrow Z^0 Z^0 \rightarrow e^+e^-e^+e^-$ decay channel. Even in this case the energy of one of the electrons can be as low as about 10 GeV. The possible observation of the Higgs boson therefore requires excellent measurements of electrons and photons from low to high energies.

- **The electromagnetic calorimeter** [14] is a sampling calorimeter made of lead absorbers bent to an accordion geometry and immersed in a liquid argon bath. Between two consecutive absorbers, a three-layer kapton electrode is dividing the 2.1 mm wide gap into two equal

Figure 2.8: Material thickness in the ATLAS calorimeter system, given in absorption lengths λ, as a function of pseudorapidity [12].

parts and plays a twofold role: it distributes the 2000 V high voltage through the two external layers, and it collects the signal in the inner one.

Liquid argon calorimetry has been chosen for ATLAS because of its intrinsic linear behavior, stability of the response in time, spatially homogeneous response and radiation tolerance. The accordion geometry has been chosen because it allows very good hermeticity. Additionally, it minimizes inductances in the signal paths, allowing the use of the fast shaping which is needed to cope with the 25 ns interval between bunch collisions at the LHC. The total thickness is 24 radiation lengths X_0 in the barrel and 26 X_0 in the end-caps.

In the barrel region ($|\eta < 1.4|$), a readout electrode is segmented laterally in cells pointing to the interaction point; longitudinally, a cell is separated in three parts, front, middle and back, with granularities of $\Delta\eta \times \Delta\phi = 0.003 \times 0.1$ for the front section to $\Delta\eta \times \Delta\phi = 0.050 \times 0.025$ for the back section.

The barrel is constructed from two wheels, each one made of 16 modules, with a module covering a range of $0 < |\eta| < 1.4$ and $2\pi/16$ in ϕ direction. It shares its cryostat with the superconducting solenoid. The

end-caps, covering the region of $1.4 < |\eta| < 3.2$ share their cryostats with hadronic and forward liquid argon calorimeters. To evaluate the amount of energy lost in front of the calorimeters, both the end-caps and the barrel are complemented with presampler detectors, covering the $0 < |\eta| < 1.8$ range. These presamplers basically are thin layers of argon equipped with readout electrodes but no absorber.

The electromagnetic end-cap calorimeter has a mechanical structure similar to the barrel calorimeter, but with absorbers arranged like the spokes of a bicycle wheel. However, while the barrel calorimeter uses only one type of absorbers and has a constant gap thickness, the end cap uses two types of absorbers, one for the outer wheel ($1.4 < |\eta| < 2.5$) and one for the inner wheel ($2.5 < |\eta| < 3.2$), with varying gap thicknesses requiring different HV values as a function of η to maintain a constant response with η.

Test-beam measurements [15] using barrel modules show an energy resolution fulfilling the requirements of $\frac{\sigma_E}{E} = \frac{10\%}{\sqrt{E}} \oplus 0.4\% \oplus \frac{0.3}{E}$.

Figure 2.9: Schematic view of a LArg module [12].

- **The hadronic calorimeter** [12] is also a sampling calorimeter. In the barrel part, where the environment is less harsh in terms of radiation dose, it consists of iron absorbers interleaved with plastic scintillator tiles. It is therefore referred to as *TileCal*. At larger rapidities, in the extended barrel sections, liquid argon technology is used, again because of its intrinsic radiation hardness. The hadronic end-cap calorimeter uses copper as absorber in a parallel-plate geometry while the forward calorimeter uses rod electrodes immersed in liquid argon in a tungsten matrix.

  In total, the hadronic calorimeter covers the range $|\eta| < 4.9$ with a thickness[6] between 11 and 20 interaction lengths depending on η (see figure 2.8). The granularity is $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$, the total number of channels is on the order of 10000.

  The principle of the TileCal is illustrated in figure 2.10. The iron tiles have a constant thickness of 3 mm and are placed perpendicular to the colliding beams and staggered in r direction. Both ends of the scintillating tiles are read out by wavelength shifting fibers into two separate photomultipliers. The signals produced by the scintillating tiles have a typical rise time of a few ns and a width of about 17 ns. The shaper transforms the current pulses from the photomultipliers into unipolar pulses with FWHM[7] of 50 ns. The electronic noise is of the order of 20 MeV per cell.

  Between the barrel cylinder and the extended barrels a 68 cm wide gap provides a passage for the cables from the inner detector and the piping from the EM calorimeter. It also houses the readout electronics for the EM calorimeter.

  The liquid argon calorimeters in the end-caps use 25 - 50 mm thick copper absorbers with gaps of 8.5 mm in between. These gaps are equipped with three electrodes. Thus the maximum drift space is ~1.8 mm. The electronic noise ranges from 200 to 1100 MeV per channel.

  The forward calorimeter is integrated in the end-cap cryostat, with the front face about 5 m from the interaction point. It has to provide at least nine interaction lengths of active detector in a very short longitudinal space. Therefore it uses a high density technique of rods, regularly spaced in a metal matrix. In the front region this matrix is made from copper while at the back tungsten is employed. The 250 - 500 μm wide gaps between matrix and rod is filled with liquid argon. The electronic

---

[6]This is the amount of material in front of the muon system, including support structure of the Tile calorimeter.

[7]Full Width at Half Maximum

Figure 2.10: The principle of the Tile calorimeter design [12].

noise in a jet cone of $\Delta R=0.5$ is $\sim$1 GeV in E_T at $|\eta| = 3.2$ and drops quickly to 0.1 GeV at $|\eta| = 4.6$.

The magnet system

The magnet system for ATLAS [16] is consists of four superconducting magnets: the Central Solenoid (CS), and the Barrel Toroid (BT) and two End-Cap Toroids (ECT) that generate the toroidal field of about 1T. The location of the parts of the magnet system is shown in figure 2.5.

- **The Central Solenoid** generates a 2T field with a current of 7600 A in the superconducting cable. The 5.3 m long coil has an inner radius of 1.2 m. It is a layer coil made of a doped Al stabilized NbTi/Cu rectangular conductor wound on its thin edge inside a 12 mm thin Al support cylinder. It is integrated in the liquid argon cryostat. It was tested up to 8100 A.

- **The End-Cap Toroids** consist each of $5 \times 5\,\mathrm{m}^2$ coil modules made of $41 \times 12\,\mathrm{mm}^2$ NbTi/Cu conductor, interlinked by eight so called keystone boxes, see figure 2.11, left. The castellated shape of the end-cap cryostat makes it possible to rotate the ECTs by 22.5° with respect to the Barrel Toroid. Since the magnetic windings overlap, the ECT is

pulled towards the Barrel Toroid with a force up to about 300 tons.

- **The Barrel Toroid** has eight racetrack coils of $25 \times 5\,\mathrm{m}^2$ size. The coils are kept in place by an inner and an outer set of eight rings of struts positioned along the coils. The assembled toroid is resting on 18 feet connected to the floor of the cavern. The coils are wound from $57 \times 12\,\mathrm{mm}^2$ NbTi/Cu conductor. They have been tested up to a current of 22 kA.

Figure 2.11: Left: Schematic view of an End-Cap Toroid magnet; Right: Schematic view of the Barrel Toroid showing the 8 coils and the positions of the two End-Cap Toroids, one inserted, one retracted. [16]

The muon system

The muon spectrometer [17] includes some precision tracking chambers for accurate momentum resolution and an effective trigger system based on fast response chambers. The Barrel chambers ($|\eta| < 1$) are arranged in three cylindrical layers, while the End-Cap chambers ($|\eta| \leq 2.7$) are mounted on wheels normal to the detector axis. The size of the muon spectrometer is about 22 m in diameter and 44 m in length.

- **The muon trigger** looks for high p_T muons by reconstructing tracks that point roughly at the interaction point, both in the $r-\phi$ and the r-z projection. Two thresholds are applied (about 6 and 20 GeV/c) with the trigger acceptance extending to $|\eta| = 2.4$. Because of the different rates in the barrel and the end-caps, two different kinds of detectors are employed:

Figure 2.12: Schematic view of the muon system. The trigger chambers are shown in red together with the coincidences required for low-p_T and high-p_T triggers respectively [17].

Resistive Plate Chambers ($RPCs$), operated in avalanche mode, are used in the barrel. Each chamber uses two gas volumes, Bakelite plates, and four planes of read-out strips. Two layers of chambers are installed in the middle station to provide the low p_T trigger. A third layer is mounted on the outer chamber station and is used, together with the other layers, for the high p_T trigger. The time resolution of the RPCs has been measured to be below 2 ns, corresponding to trigger resolution better than 3 ns.

Thin Gap Chambers ($TGCs$), multi-wire chambers, operated in saturated mode, are used in the end-caps. Their anode-to-anode pitch is 1.8 mm with the anode to cathode gap being 1.4 mm. The cathode is coated with graphite. External pick-up strips provide the coordinate along the sense wires. The chamber gas is a mixture (55%-45%) of carbon dioxide and n-pentane. Three multi-layers of chambers (one triplet and two doublets) are located in the middle tracking station. Additional TGCs are part of the inner station and are used to increase the tracking ability. Tests performed at a high rate have shown single-plane time resolution of about 4 ns rms, with 98% efficiency, corresponding to a trigger efficiency of 99.6%.

- **The precision chambers** are used for high resolution tracking and

momentum measurement of the muons.

Monitored Drift Tubes (*MDTs*) are used in the pseudorapidity region of $|\eta| < 2$. These are drift chambers formed by aluminum tubes with 3 cm diameter and lengths ranging from 0.9 to 6.2 meters. On each chamber the tubes are arranged in two multi-layers, each formed by three or four layers of tubes. They use a gas mixture of 93% Ar and 7% CO_2, kept at 3 bar(a)[8], and are operated at a gas gain of 2×10^4. They can thus sustain high rates without aging and with very little sensitivity to space-charge. The single-tube resolution is 100 μm for most of the range in drift distance, and the multi-layer resolution is $\sim 50\mu$m. To exploit such tracking accuracy on chambers covering up to 10 m^2, an extremely acurate mechanical construction is needed. In addition to a sophisticated mounting structure, the aluminum frames supporting the multi-layers are equipped with straightness monitors, temperature sensors (needed to correct for thermal expansion of the tubes) and with magnetic field sensors to predict the $\vec{E} \times \vec{B}$ effect on drift time. In total 1174 MDTs are used.

Cathode Strip Chambers (*CSCs*) are used in the region of $2 < |\eta| < 2.7$. These multi-wire proportional chambers *MWPCs*, with a sense wire pitch of 2.54 mm, are read out by 5.08 mm pitch readout strips. The track resolution in the bending plane is 60 μm. 32 CSCs are employed.

Figure 2.12 shows a schematic view of the muon system. The positions of the trigger chambers are highlighted in red with tracks indicating the low p_T and high p_T trigger coincidences.

2.3.2   The data acquisition and high-level trigger system

With a bunch-crossing frequency of 40 MHz and about 25 interactions per crossing in the low-luminosity phase, a very sophisticated trigger and data-acquisition system is required to select the very rare processes, of main interest to physicists.

The ATLAS Trigger and Data Acquisition system (*TDAQ* [18, 19]) is based on a multi-level selection process and a hierarchical acquisition tree. The system consists of a combination of custom made electronics and commercial products. It is required to reduce the data rate by a factor of 10^5
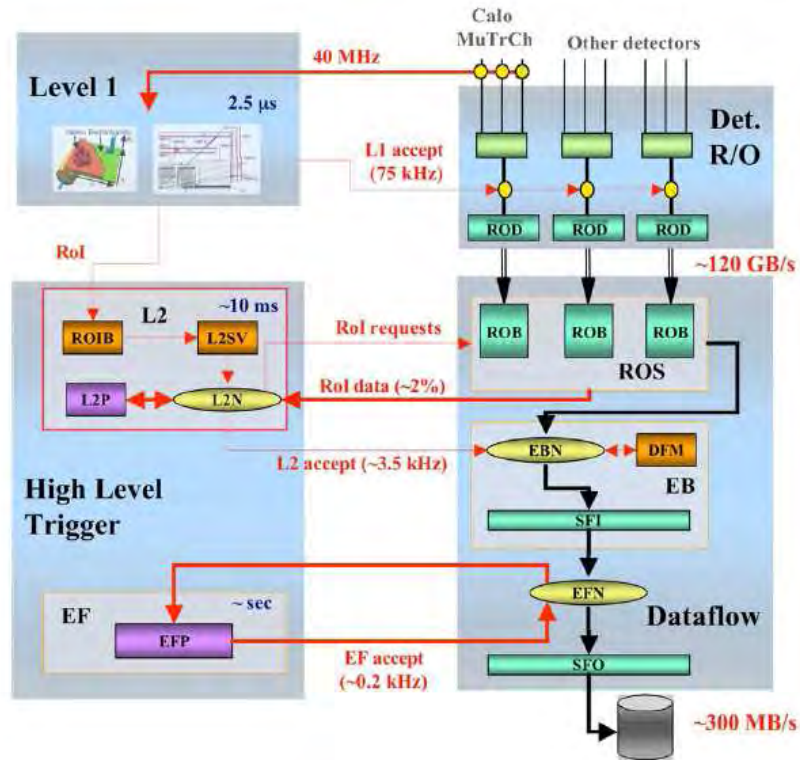
---

[8]Absolute pressure in bar

Figure 2.13: ATLAS trigger and data acquisition architecture. Thinner arrows indicate the flow of control messages, thicker ones indicate the flow of data fragments. The black arrows show the main data path.

from the initial collision frequency of 40 MHz to a rate of events written to mass storage of about 200 Hz, and to provide a total data throughput in the range of Terabit/s. The system is designed to take maximum advantage of the physics nature of events. A Region-of-Interest (RoI) mechanism is used to significantly reduce the amount of data needed to calculate the trigger decisions.

Figure 2.13 shows the functional elements of the ATLAS TDAQ system. It can be divided logically into a fast first level trigger (Level 1), a High Level Trigger system and a Dataflow system.

- **Level 1 trigger.**
  The first trigger level [20] is a hardware based system that reduces the event rate from 40 MHz to about 75 kHz (later on an upgrade to 100 kHz is projected)[9] within a fixed latency of less than 2.5 μs. The

---

[9]Trigger studies estimate the LVL1 rate required to meet the ATLAS physics program

Level 1 (*LVL1*) system is composed of three parts: the Calorimeter Trigger (L1Calo), the Muon Trigger (L1Muon) and the event-decision part implemented in the Central Trigger Processor (CTP).

Events in the calorimeter system (see chapter 2.3.1) are digitized by the Preprocessor which also performs bunch-crossing identification. A Cluster Processor identifies e^-/γ and hadron$/\tau$ candidates using a sliding-window algorithm. The transverse energy E_T of the candidates is discriminated against up to 16 programmable thresholds. A Jet/Energy Processor identifies jet candidates and discriminates their E_T values against eight programmable thresholds. It also evaluates several global energy sums (e.g. missing E_T). The L1Calo sends multiplicities of the e^-/γ, hadron$/\tau$ and jet candidates as well as the global energy information to the Central Trigger Processor.

The algorithms of the muon trigger are based on hit coincidences in different stations within a geometrical track whose width is related to the p_T threshold applied. The coincidence logic allows up to six thresholds to be applied. L1Muon forwards the multiplicities of muon candidates for each threshold to the CTP.

The Central Trigger Processor combines the input signals logically to up to 256 different trigger types, according to a trigger menu. It applies deadtime and prescale factors to each trigger. The L1A signal is a logical OR of all triggers and is distributed to the sub-detectors via TTC[10] partitions including one Local Trigger Processor (*LTP*) each. A BUSY tree allows the sub-detectors to block the generation of LVL1 triggers.

During the latency of the Level 1 trigger selection algorithms the complete event data is kept in the pipeline memories of the sub-detector front-end electronics. Only data of events selected by the LVL1 system is then transferred to the ROBs[11], where it is temporarily stored and provided to the High Level Trigger system on request. The components of the Level-1 system also send information regarding the trigger type and the η and ϕ coordinates of the objects, that caused the event to be accepted, to the Region-of-Interest Builder (*RoIB*) via an optical connection. The RoIB collects the Region-of-Interest information from the Level-1 system and passes it to the Level-2 system. The Region-of-Interest Builder is implemented as a custom VMEbus[12] system con-

---

needs to be about a factor two lower than this.

[10]Timing, Trigger and Configuration

[11]Read-Out Buffers

[12]VersaModule Eurocard bus, ANSI/IEEE 1014-1987 standard.

sisting of multiple input and builder cards controlled by a SBC[13].

- **High Level Trigger and Dataflow.**
  The second and third trigger levels, Level-2 trigger and Event Filter (*EF*), together form the High Level Trigger (*HLT*) system. It is running on farms of commercially available PCs, executing an offline-like software package called HLT Selection Software (*HLTSSW*). This software is based on the ATLAS offline Athena framework. During operation of the Level-2 and EF algorithms the event data is stored in the ROBs. Events can be rejected at any point in the decision process and are only deleted from the ROBs if rejected or after fully building the event upon acception by the Level-2 trigger.
  The Level-2 trigger system is designed to provide a data reduction factor of 20 - 30 within a latency of $\approx$ 10 ms. As it must work at the Level-1 accept rate, and required latency is to be minimized, the algorithms used provide modest precision. They request data from the ROBs identified from the Region-of-Interest. Therefore typically only about 2% of the event data is processed, minimizing network traffic. The expected rate of events accepted by Level-2 is about 3.5 kHz.
  If an event is accepted by the Level-2 trigger algorithms the Data Flow Manager (*DFM*) assigns the event to the event-building nodes (Sub-Farm Interface, *SFI*) according to load-balancing criteria. One SFI per event collects the full event data from the ROBs and fully builds the event before passing it to the Event Filter. More sophisticated reconstruction and trigger algorithms, adapted from the offline reconstruction software, process the event with an average treatment time of $\approx$ 1 s. Events selected by the EF are then sent to mass storage by the Sub-Farm Output (*SFO*) node. An event rate of 200 Hz is expected after the EF. Given a mean ATLAS event size of about 1.6 MB, data is written to mass storage at a rate of about 300 MB/s.

Online software

The Online Software system is responsible for configuring, controlling, and monitoring the TDAQ system. It provides the services that enable the TDAQ and detector systems to start up and shut down. It is also responsible for the synchronization of the the entire system and the supervision of processes. During data taking, access is provided to monitoring tasks, histograms produced in the TDAQ system, and the errors and diagnostics messages sent by different applications. One or more user interfaces display the available

---

[13]Single Board Computer

information and enable the user to configure and control the TDAQ system, including transitions between the states of the TDAQ system described below.

TDAQ states

The high-level operation of the experiment, and the relationships between the main systems (the detector itself, the LHC machine, and the TDAQ system) are described in terms of logical and functional states. A given state determines what operations are allowed while the system is in that state. For example, the operator will not be allowed to turn off high voltages while the detector is taking data, whereas he will be allowed to stop data taking.

Three principal states can be used to define the main operations of the TDAQ system:

- *Initial.* This state represents the case when TDAQ sub-systems are started up but idle. It is the earliest state in which it is possible to send commands directly to the TDAQ control elements. The only operations which are allowed are those which bring the overall system to a situation where data-taking can be performed. The TDAQ system may also revert to this state in order to re-initialize significant parts of ATLAS.

- *Configured.* In this state the TDAQ system is ready, provided other conditions related to the state of the detector or the LHC machine are met, to initiate a data taking session. This means that the various components have been properly initialized. In some initialization procedures, the configured state may be reached by passing through one or several sequentially-dependent intermediate states, for example, loading different layers of software into processors, configuring interdependent components, and so on.

- *Running.* In this state the TDAQ system is data-taking.

Chapter 3

The pixel detector

3.1 Physics performance requirements

Pixel technology has been chosen for the innermost layers of the ATLAS Inner Tracker because it is the one technology that can cope with the physics requirements with respect to tracking in the high radiation environment close to the interaction point. With the small pixel size of 50 μm by 400 μm and the large number of channels, the single-channel occupancy is low enough to ensure proper readout, even in the high luminosity phase of the LHC, where on average 1000 charged particles are expected per collision. As a pixel detector delivers truly two-dimensional measurements of the tracks, there are no ambiguities or 'ghost hits' for multiple hits per detector unit as they are inherent in detector systems yielding only one-dimensional hit information. The exceptional radiation hardness of the pixel detector allows placing it very close to the interaction point. Together with the exemplary single-point resolution this allows to provide robust pattern recognition and excellent vertexing capability at the design luminosity of the LHC.

The capability to distinguish high p_T b-jets from light quark- and gluon-jets *(b-tagging)* is most important for the detection of Higgs and SUSY signals. In particular the production of SUSY particles would result in several high p_T b-jets per event, requiring a b-tagging efficiency of 50% with a rejection factor against light-quark jets of about 100. To achive this performance it is necessary to measure secondary vertices with good precision, as b-tagging algorithms identify b-jets by the small distance the b quark travels before hadronizing.

In addition to the detection of Higgs and SUSY particles the study of the decays of B-hadrons are of great importance. This includes the search for CP-violation, B_S^0-mixing, rare B-decays and B-hadron spectroscopy. Most of the

B-physics topics require efficient detection of low p_T leptons and hadrons together with good vertex resolution in Rϕ- and z-direction. Precise secondary vertex finding is very important in controlling combinatorial background, for instance in the process $B_d^0 \rightarrow \pi^+\pi^-$.
Most of the current tracking algorithms are based on seed points in the pixel detector. Therefore it is necessary to provide a maximum number of high-resolution measurements over a large region in pseudorapidity. The geometry of the pixel detector, consisting of 1744 individual pixel modules arranged in a central barrel region with three layers, and three disk layers in both forward directions, together with a single-hit efficiency of 99.9% [21] enables the pixel detector to provide three hits per track down to a pseudo-rapidity of $|\eta| = 2.5$

In this chapter the basic components of the pixel detector will be described, including the service and readout connections, off-detector electronics, mechanical structures and the cooling system.


Figure 3.1: Technical drawing of the pixel detector mounted in the global support structure. For more information on the mechanics see chapter 3.9.

3.2 The pixel module

The basic functional unit of the pixel detector, the pixel module, is a hybrid assembly consisting of 16 readout chips, arranged in two rows, connected to a silicon sensor. For the interconnection between the pixel implants on the sensor side and the circuitry on the readout chip two different bump-bonding techniques are used with a pitch of 50 μm between adjacent bump bonds. One of the bumping vendors[1] used solder bumps deposited on the readout chip. The readout chips are then positioned atop the corresponding pixel cells on the sensor using a flip-chip technique. The assembly is heated above melting temperature of the lead-tin solder, which during this reflow process connects to a metalization on the sensor. The reflow process leads to self-alignment of the sensor and readout chip as the surface tension of the solder pulls it into a spherical shape. The resulting connections are mechanically very stable, yet a chip is easily disconnected in the event of a large number of bad bump connections.

Figure 3.2: Lateral cut through a sensor/readout chip hybrid with IZM bumps.

The second bumping vendor[2] used indium studs with a tip of gold deposited on sensor and readout chip that are merged under moderate pressure and moderate heating. These bumps have a slightly higher ohmic resistance and are mechanically not quite as stable as the solder bumps.

---

[1] Fraunhofer Institute for Reliability and Microintegration, IZM, Berlin
[2] SELEX Sistemi Integrati, Rome; former Alenia Marconi System, AMS

Figure 3.3: Schematic view of the cross-section of a hybrid pixel module.

To route the supply voltages and TTC signals to the front-end chips, and the hit data to the readout system, a flexible kapton PCB[3] (*flex*) is glued to the back-side of the sensor, which also houses the module control chip (MCC). The front-end chips and the MCC are connected to this flex by 22 μm diameter wire-bonds.

The module is connected to the outside world through a custom-made aluminum cable. This is connected to the flex directly for the disk modules or plugged into a connector on another little flex kapton PCB, the *pigtail*, in case of the barrel modules. These cables are plugged into connectors on the PP0[4], which is why they are called Type0 cables.

Figure 3.3 shows a cross-sectional view of a module with the readout chips bump-bonded to the sensor, the flex glued to the backside of the sensor, housing the MCC. Figure 3.4 shows a schematic view of a module (left) and of a single readout chip (right), indicating the separate module- and chip-coordinate systems. In a single readout chip rows are numbered from 0 to 17 and columns from 0 to 159. In the module coordinate system rows are numbered from 0 to 143, columns from 0 to 319, regardless of chip boundaries.

---

[3]Printed Circuit Board
[4]Patch-Panel 0, see chapter 3.4

Figure 3.4: Schematic view of a module. Indicated is the module coordinate system.

3.2.1 The sensor

The sensor of the ATLAS pixel module has an active area of $60.8 \times 16.4 \, \text{mm}^2$ and a thickness of $250 \, \mu m$. It is divided into 40960 pixels with an area of $400 \times 50 \, \mu\text{m}^2$. These are directly bump-bonded to the electronics circuitry on the readout chips. Adjacent readout chips are $400 \, \mu m$ apart in all directions. In order not to loose the area of the sensor that is not covered by a readout chip, different strategies are used. Pixels between neighboring chips in z direction are elongated to $600 \, \mu m$ instead of the usual $400 \, \mu m$, while pixels between neighboring chips in ϕ direction are connected by traces on the passivation on the sensor to pixels that can be read out. Therefore the former are referred to as *long pixels*, while the latter are called *ganged pixels*. The connection between the sensor implantations forming a ganged pixel is done in a way that every other of the topmost pixels in a row is connected to one of the inter-chip pixels (see figure 3.5). This connection scheme allows for the pattern recognition algorithms to dissolve the ambiguity between the original and the inter-chip pixel.

Signal generation in semi-conductors

Charged particles traversing matter loose energy primarily by ionization and atomic excitation. The mean energy loss is described by the Bethe-Bloch formula: [22]

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} ln \frac{2 m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \tag{3.1}$$

with:

and the kinematic variables $\beta = \frac{v}{c}$ and $\gamma = \frac{1}{\sqrt{1-v^2/c^2}}$.

Symbol	Definition	Units or value
K/A	$4\pi N_A r_e^2 m_e c^2/A$	0.307075 MeV g^{-1}cm^2 for A=1 g mol^{-1}
N_A	Avogadro's number	6.0221415 x 10^{23} mol $^{-1}$
r_e	classical electron radius	2.817940325 fm
ze	charge of incident particle	
Z	atomic number of absorber	
$m_e c^2$	electron rest mass x c^2	0.510998918 MeV
T_{max}	maximum kinetic energy that can be imparted to a free electron in a single collision	
I	mean excitation energy	
$\delta(\beta\gamma)$	density effect correction to ionization energy loss	

Figure 3.6 shows the mean energy loss as a function of $\beta\gamma = p/Mc$, where M is the mass of the incident particle. A minor dependence of M is introduced through T_{max}, but for all practical purposes dE/dx in a given material is only a function of β with a minimum in energy loss at $\beta\gamma \approx 3.5$. For higher energies the electric field of the traversing particle extends, increasing the contribution of distant collisions to the energy loss. This is called the *relativistic rise*. In this region dE/dx rises again with $\ln(\beta\gamma)$.

Due to the statistical nature of the processes involved, energy loss is not uniform along the path of the incident particle. Energy is deposited in clusters. For thick layers of matter the energy loss thus follows a gaussian distribution, whose mean value is given by the Bethe-Bloch formula. For thin layers of matter the distribution is better described by a Landau distribution with a long tail towards high energy loss. This tail is due to direct collisions with electrons. The large kinetic energy transferred to these δ-electrons enables them to travel through the sensor, creating avalanches of free charge carriers by secondary ionization. Because of the high-energy tail the most probable energy loss lies below the mean value as given by the Bethe-Bloch formula. For relativistic particles ($\gamma \gg 100$) it can be calculated by [23]:

$$\Delta E = t(0.1791 + 0.01782\ln(t)) \tag{3.2}$$

with ΔE being the most probable energy loss in keV and t the thickness of the absorber in μm. For the 250 μm thick sensor the most probable value of 69.4 keV is found, as shown in figure 3.7.

Although the band-gap energy in silicon is 1.1 eV, the mean energy required to generate an electron/hole pair is 3.62 eV due to phonon excitations which absorb part of the deposited energy. Thus, for a minimum ionizing particle traversing the 250 μm thick silicon sensor, a charge of about 19 000 e$^-$

Figure 3.5: Schematic view of the region between four readout chips. Ganged pixels including the traces connecting two pixels are shown as well as long pixels.

is deposited.

Radiation damage in silicon sensors

The effects of irradiation on silicon can be classified in two categories: non-ionizing energy loss, mainly affecting the bulk material, and ionizing energy loss, mainly affecting the oxide layer and the Si-SiO$_2$ junction. [24]

In the case of non-ionizing energy loss (NIEL) the traversing particle transfers energy to an atom in the crystal lattice, thus kicking it free from the lattice. If the transfered energy is large enough, this so-called *Primary Knock on Atom* can travel through the crystal removing secondary atoms from the lattice thus creating a cascade of lattice displacements. The probability to create these cluster defects depends on the kind of traversing particle and its energy. The effects of different particles can be normalized to the damage created by a neutron with energy 1 MeV. Therefore the fluence of particles is expressed in equivalents of 1 MeV neutrons (n_{eq}).

Cluster defects create new states in the band-model of the semiconductor

Figure 3.6: Stopping power for μ^+ in copper as a function of $\beta\gamma = p/Mc$. [22]

which are electrically active if they lie in the band-gap. Depending on the position of the new state in the energy gap these clusters have different effects on the sensor:

- *Generation/Recombination Centers*
  States roughly in the middle of the energy gap act as generation or recombination centers. Their probability for emission or absorption of a free charge carrier is higher than for the shallow levels close to the valence or the conduction band. These defects act on the effective concentration of charge carriers, by thermally creating free electron-hole pairs or absorbing electrons or holes close to the edges of the bands. This increases the leakage current, leading to increased heating of the sensor. As leakage current is strongly dependent on temperature, this can lead to thermal runaway, further increasing leakage current. Therefore sufficient cooling of the sensor is necessary.

- *Shallow Donor/Acceptor Levels*
  Shallow states act as trapping centers, temporarily trapping free charge carriers. This leads to a decrease in signal charge. For highly irradiated sensors the charge loss for signal charges is dominated by the smaller charge collection distance due to partial depletion. Therefore radiation induced increase of trapping centers can be treated as a small effect.

Figure 3.7: Most probable energy loss in silicon as a function of sensor thickness

- *Doping Compensation*
  States deep in the energy gap can permanently capture free charge carriers thus changing the effective doping concentration N_{eff}. Figure 3.8 shows the dependance of N_{eff} from the accumulated fluence in n_{eq}. This affects the depletion depth of the material and finally leads to type-inversion. The number of acceptor-like defects is directly proportional to the fluence. Thus initially n-type material can accumulate so many acceptor-like defects that it effectively acts as p-type material. The effective doping concentration can reach even higher values than the initial doping concentration. Therefore the bias voltage needed to fully deplete the irradiated sensor increases.

In the case of ionizing energy loss the traversing particle can knock an atom from its position in the lattice, as described above, and it can ionize an atom, creating free charge carriers. Ionizing energy loss predominantly affects the oxide layer and the silicon-silicon dioxide junction (*surface defects*). As SiO_2 is amorphous the lattice displacement are reversible. If a voltage is applied over the oxide layer the electron-hole pairs are separated and drift through the oxide. The electrons quickly move to the electrode while the holes very slowly drift towards the Si-SiO_2 junction where they are bound in deep trapping levels. This leads to a shift in the threshold voltage in case of a MOS transistor. In the case of the sensor surface the major effect is an increase in the dark current of the sensor that adds to the leakage current. This is a small effect compared to the NIEL induced increase in the leakage current.
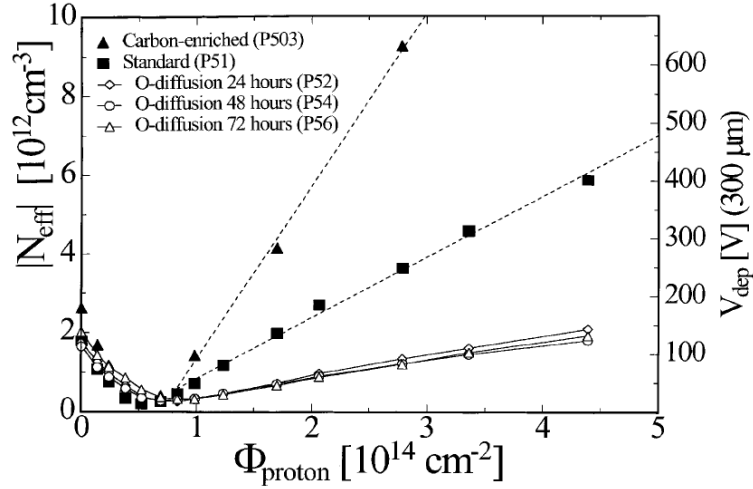
Figure 3.8: Dependence of N_{eff} on the accumulated fluence in n_{eq} for standard and oxygen enriched float-zone silicon. [25]

The structure of cluster defects and of surface defects changes with time. This *annealing* effect leads to a partial reduction of the increased leakage current and the increased effective doping concentration. After some time (*beneficial annealing*) both leakage current and doping concentration reach a minimum value. If the annealing process continues, different types of defects can react and form more electrically active states in the band-gap than were created during irradiation. Leakage current and doping concentration increase beyond the levels without any annealing. This is called *reverse annealing*. As annealing time constants depend on temperature, the effects of annealing can be suppressed by cooling the silicon crystal. Therefore the ATLAS pixel sensors are operated at a temperature of -7° C.

Sensor design

The ATLAS pixel sensors will be irradiated to a fluence of about 10^{15} n_{eq} cm^{-2} within 10 years of operation. For the sensor material diffusion oxygenated float zone (DOFZ) silicon was chosen, as studies by the ROSE collaboration [26] have shown oxygenated silicon to suffer less from the increase in N_{eff} due to radiation damage. Therefore the bias voltage needed to fully deplete the sensor should stay below the maximum voltage of 600 V as determined by the high voltage services.

The segmentation of the sensor is implemented as n^+ pixels[5] in n-type bulk material (see figure 3.9, [24]). Initially the depleted region grows from the backside of the sensor, which is p^+ doped, towards the pixel implantation. As the depletion voltage is low the sensor can be fully depleted. After type-inversion the bulk-material is effectively p-doped, and the depleted region grows from the pixel side towards the backside of the sensor. Even if the sensor cannot be fully depleted after irradiation, charges generated in the depleted region can still reach the pixel electrodes. The sensor can thus be operated in partial depletion.


Figure 3.9: Schematic view of the development of the depletion region before (left) and after (right) type-inversion.

Due to the n-in-n design, without further implantations the n^+ pixels are shorted by an electron accumulation layer underneath the SiO_2 layer. This would render the sensor unusable. Therefore a separate p implantation has to be placed between the pixels to isolate them from each other. To avoid high electric fields at the edges of the implantations, these are implemented in a moderated p-spray design. The complete gap between two n^+ pixels is filled with a p implantation with the doping concentration decreasing from the middle of the gap towards the edges of the n^+ implantations. This reduces the gradient of the doping concentration, reducing the amplitudes of the lateral electric field maxima. Thus the sensor is less prone to avalanche breakdown.

To allow testing the sensors, a bias network is implemented which allows to define the potential of every single pixel. This is done using the *punch-through effect*. In every pixel a small n^+ implantation is implemented, separated from, but surrounded by the pixel implantation. This so-called *bias-dot* is

---

[5]donor concentration of up to 10^{17}-10^{18}cm^{-3}

connected to a metal trace on top of the oxide layer, which in turn connects to a n^+ ring around the whole sensor. Thus all pixels can be held at a defined potential without the need to connect them separately. This bias network also prevents pixels, that are not held at ground potential due to a missing bump-bond, to float to an arbitrary potential. By keeping these pixels close to ground potential sparkover from the sensor to the FE electronics is prevented. As in this design the bias voltage is connected at the back-side of the sensor, the p^+ implantation, to which it is applied, has to be isolated from the edge of the sensor crystal. Otherwise current would be flowing to the pixel side via the conductive crystal edges. Therefore a number of electrically floating guardrings around the p^+ implantation are implemented to gradually lower the potential to that of the pixel implantations on the readout side. In case of the ATLAS pixel sensor a multi-guardring structure of 17 rings with a pitch increasing from 20 to 50 μm is used. It covers a total width of 0.5 mm around the active area. Measurements have shown that the structure can cope with a bias voltage of more than 1000 V after irradiation to $10^{15} n_{eq} cm^{-2}$.

This need for back-side processing is one of the main disadvantages of the n-in-n sensor design, as it increases the production cost significantly.

Studies with irradiated prototype modules [27] have shown that the 250 μm thick sensor can be fully depleted after irradiation to 1.1×10^{15} n_{eq} cm^{-2}. The average signal charge was measured to be $(87 \pm 14)\%$ of its pre-irradiation value.
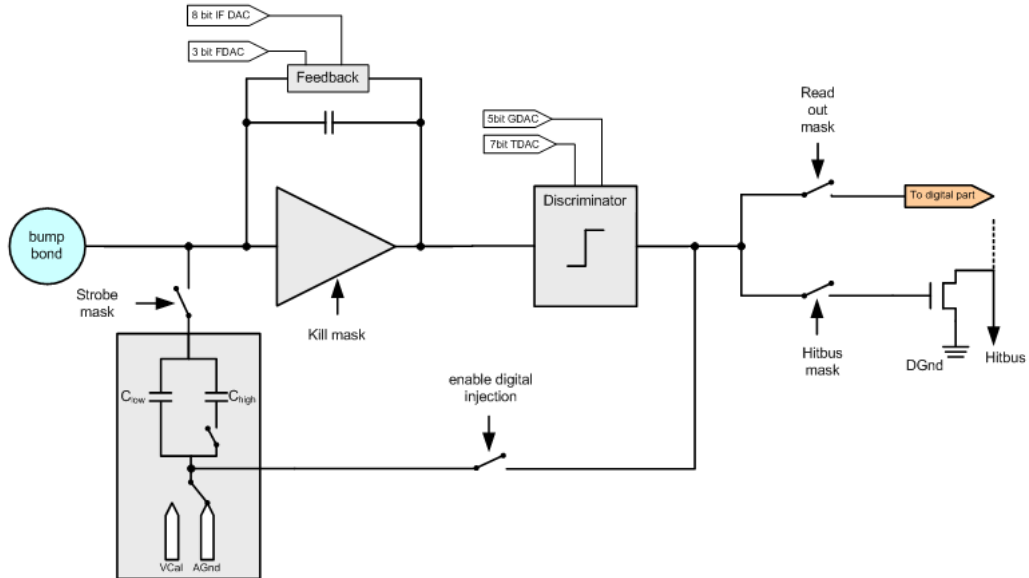

Figure 3.10: Schematic view of the analog part of a pixel cell in the FE-I3 chip including control signals.

3.2.2 The front-end chip

The readout chip of the pixel detector is simply called front-end or FE-chip. It is implemented in a 0.25 μm CMOS process at IBM and consists of 18 x 160 readout cells on an active area of 7.2 x 10.8 mm^2. The readout cells have an area of 400 x 50 μm^2, covering one sensor pixel. They contain two cascaded amplifiers, followed by a fast differential discriminator featuring a detection threshold in the range of 0 to 6000 electrons [28]. The thresholds can be adjusted globally for the chip, using a 5 bit global DAC (GDAC), and individually using the 7 bit trim DAC (TDAC) to reduce pixel to pixel variations in threshold. The first amplifier stage, the preamplifier, is implemented as a charge sensitive amplifier. The charge on the input is collected on a feedback capacity of about 10 fF, which in turn is discharged by an adjustable constant current source. The amplitude of the feedback current can be adjusted globally using the 8 bit IF DAC, and individually using the 3 bit feedback DAC (FDAC). Thus the deviations in amplifier recovery time can be minimized over the chip. Figure 3.10 shows a schematic view of the analog part of a pixel cell.

The point in time when the amplifier signal crosses the discriminator threshold (leading edge time, LE) is used as time stamp for this hit. As the falling edge of the amplifier signal is linear, the difference of the leading and the trailing edge *(time over threshold, TOT)* of the signal is directly related to the charge deposited in the sensor. Thus the front-end provides a measurement of the energy loss of the traversing particle in the sensor by means of the 8 bit TOT information. This information can be used to increase spatial resolution in the case of charge sharing between two pixels, hit by a particle. The probability for a particle, to create a hit in two adjacent pixels, is ∼97%.


Figure 3.11: Drawing of output signals of the amplifier and the discriminator. Shown are signals for a high charge (green) and a low charge (red) at the input of the amplifier.

It is a feature of the preamplifier that the output signal always peaks after

the same time, independent of the injected charge. Therefore the LE time of a hit with small charge deposition can be later than of a hit physically occurring at the same time but with larger charge deposition. This *timewalk effect* can delay a hit by several 40 MHz clock. The front-end employs a hit rejection and a hit doubling mechanism to be able to recover hits that are recorded with a timestamp one bunch crossing too late due to timewalk. An adjustable threshold on the TOT of the hit can be applied to copy it, with the timestamp decremented by one unit. As the probability for a hit with small TOT to be a noise hit is rather large, hits with a TOT smaller than a threshold can be deleted directly to reduce data volume and noise occupancy. The digital part of the pixel cell consists of two 8 bit RAM cells that store the leading and trailing edge timestamp of a hit. It also contains a RAM cell holding the 8 bit pixel address. When a rising edge of the discriminator output signal is detected the gray coded[6] timestamp is stored in the LE RAM cell. Only when the falling edge of the discriminator output signal is detected the TE timestamp is stored and the pixel is marked as hit. This hit flag triggers the readout logic to ripple through the pixels in a column pair. If multiple pixels are hit, the topmost one is read out first by sending its address and the two timestamps to the end-of-column logic. There the TOT is calculated from the difference of the timestamps and the hit information is stored in one of the 64 buffer cells per column pair. Only then the hit flag for this pixel is reset and the next pixel can be read out. During this readout cycle a *freeze* signal, issued by the end-of-column logic, prevents new pixels to set their hit flags. The hits are not lost, but are stored within the pixel and the hit flag is set as soon as the readout cycle is finished.

The LE timestamps of the hits stored in the end-of-column buffers is constantly compared to a readout timestamp counter, which is a copy of the hit timestamp counter reduced by the adjustable *latency*. When the timestamps are equal, i.e. the time between the leading edge of the hit and the LVL1 equals the latency, and the LVL1 signal is active, the hit is sent to the serializer to be written to the MCC. If the LVL1 signal is not active the hit is deleted. The hit information sent to the MCC includes the four least significant bits of an internal bunch-crossing counter (*BCID*), used to identify the event the hit belongs to. This is needed for event building in the MCC.

---

[6]Binary numeral system where two successive values differ in only one digit. Reduces digital activity in the chip.

The Control Circuitry

The configuration of the FE chip, including DAC settings in the analog part, various control bits per pixel and the setting of the output multiplexer, is stored in three different registers. With the extreme fluences close to the interaction point, one of the main design issues is to harden these registers against single-event upsets (SEU). A charged particle can also deposit charge in the readout electronics and so change the potential inside a register cell, as to change the stored logic state. In order to maximize the intervals between reconfigurations of the FE chip, the registers were designed to minimize the abundance of SEU-events.

The settings for individual pixels, like TDAC and FDAC values, as well as four control bits, are stored in 14 latch cells implemented in every pixel. The settings for one of the latches is clocked into a shift register connecting all 2880 pixels. The values are then written to the latch cells in parallel, through a command entered in the command register. For these latch cells very little room was available in the design of the pixel cell. A special SEU hard latch design was used for these cells, maximizing the distance between critical nodes in the cells to reduce the probability for SEU events.

The 231 bit global register and the 29 bit command register are implemented using triple redundancy and various parity checks, as well as special SEU-hard latch cells in some places. The global register stores values for global DACs, like the supply currents for the amplifiers, the global threshold DAC, the output multiplexer and latency settings. The command register handles the write and read signals for the latches of the pixel register, sets the input multiplexer for pixel or global register and issues reset signals to the digital readout logic.

Test functionality

The functionality of the FE-chip can be tested at various stages. To test the communication with the FE-chip and the functionality of the control registers, the contents of the global register and the 14 pixel register latches can be read back.

To provide the possibility of calibration measurements a chopper circuit is implemented in each pixel to generate a calibrated voltage step. An adjustable dc voltage, provided by a global DAC (VCAL DAC), is chopped to the analog power supply voltage when a digital strobe signal is given. The output of the chopper can be connected to different parts of the pixel analog and digital circuitry as shown in figure 3.10.

The voltage step can be applied to a selectable injection capacitor of either

8 or 40 fF. In this way, a calibrated charge is injected into the preamplifier. This function is used for measuring the actual setting of the threshold and the noise, or for calibrating the TOT.

The voltage step can also be applied at the output of the discriminator. In this way the digital circuitry can be tested stand-alone, including the buffer cells and the circuit that calculates the TOT.

For details on the most basic scans using this functionality see chapter 4.3.3.

3.2.3   The Module Control Chip
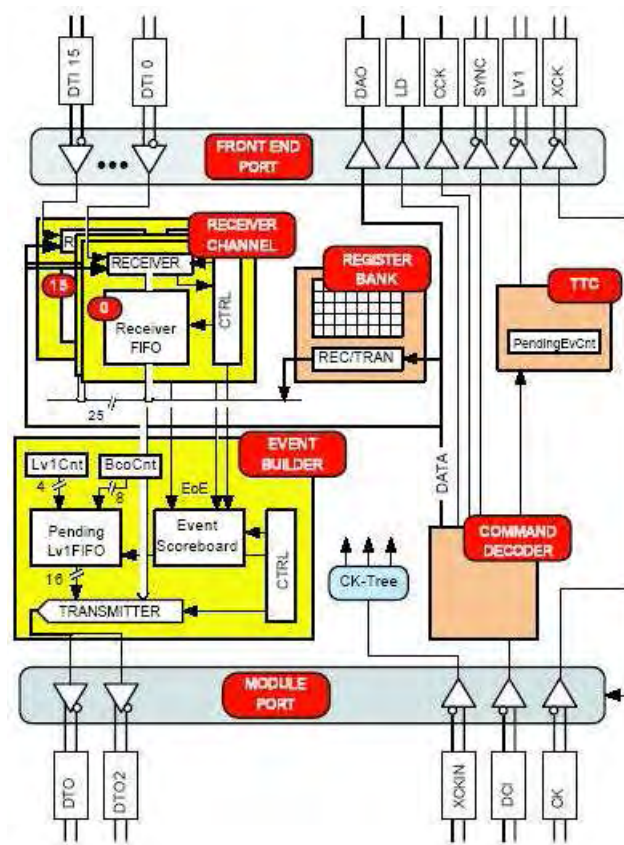
Figure 3.12: Diagram of the building blocks of the module control chip. [29]

The module control chip (MCC) [29] handles the distribution of configuration data, trigger and clock to the 16 FE chips on a module, as well as event readout and event building. It is implemented in a 0.25 μm CMOS process as is the FE chip. Configuration data for the FEs and the MCC itself, as well as the 40 MHz clock, are received from the optoboard (see chapter

3.3.1) and event data is sent to the optoboard via four differential lines using a low-voltage differential signal. To further reduce possible EMI[7] coupling to the sensitive FE chips, data out lines from the FE (FE-DO) use 500 μA current drivers instead of the 3 mA LVDS standard.

From the external clock the MCC creates the 40 MHz module readout clock XCK and the 5 MHz configuration clock CCK. The connections to the FE chips are implemented in a star topology using unidirectional serial lines. The configuration data for a given FE chip is distributed to all FEs in parallel together, with a 4 bit address to identify the FE chip on the module[8]. Configuration data and clock, and the LD signal used in the configuration process of the FE, are distributed as full-swing single-ended CMOS signals. Figure 3.12 shows a simplified block diagram of the MCC internal architecture. The blocks represented are:

- *Front End Port.* This block implements the interface of the MCC core to the FE chips including LVDS and CMOS drivers/receivers. It houses the outputs for the front-end CMOS signals (DAO, LD, CCK) and the differential outputs (SYNC, LV1, XCK) as well as the 16 differential data inputs (DTI 0 ... DTI 15) from the front-end chips.

- *Module Port.* This block implements the interface of the MCC core to the ROD[9] including LVDS drivers/receivers. It houses the differential clock and data inputs (CK and DCI) and the two differential data outputs (DTO and DTO2).

- *Command Decoder.* All commands sent to the MCC are received and decoded in this block. This includes fast commands like the LVL1 trigger signal as well as slow commands like read/write operations to internal MCC registers.

- *Register Bank.* This block contains the internal configuration and status registers necessary for the operation of the MCC. In total it contains 11 16bit registers.

- *TTC (Trigger, Timing and Control).* This block generates the L1 trigger for the FEs upon reception of the LV1 command. It can stop the generation of level 1 triggers if more than 16 events are pending transmission to the ROD. Calibration signals for the FEs and various reset signals for MCC or FEs are also generated here.

---

[7]Electro-Magnetic Interference
[8]The address of a FE-chip is set through wire-bond connections.
[9]Read-Out Driver, see chapter 3.3.4

- *Receiver Channel.* This block is implemented separately for every FE and receives the serial data stream from the FE. After the hit is completely received it is written to the receiver FIFO in parallel. Each FIFO is 128 words deep. The controller in this block handles the read/write pointers for the FIFOs and can detect end-of-event (EOE) words sent by the FE to signal event completion.

- *Event Builder.* This block contains the EventScoreboard and the PendingLv1FIFO. The latter stores the content of the L1 trigger counter which is incremented each time a LV1 command is decoded and does not exceed the storage capacity of the MCC, meaning that the value of the L1 counter is smaller than 16. If the value equals 16, the generation of the L1 trigger for the FEs is suppressed. End-of-event words from the FEs are stored in the EventScoreboard. Once the scoreboard contains EOE words from all 16 FEs for the same event, the event builder will process it.

The MCC has three main system tasks: loading parameter and configuration data in the FE's and the MCC itself, distribution of timing signals like XCK, L1 trigger and resets, and FE chip readout and event building.

- *System configuration.*
  The FE chips and the MCC have to be configured after power-up or before starting a data taking run. It is possible to write and read all MCC registers and FIFO's. This is used to configure, read status information or to test the functionality of the chip. For the latter function it is possible to write simulated events to the receiver FIFO's, and then run the EventBuilder on these events. This is used to test the MCC functionality once it is embedded in the pixel detector module. All the FE registers are written through the MCC.

- *Trigger, Reset, Timing.*
  Each time the MCC receives a LV1 trigger command the TTC block checks if the content of the pending event counter (PendingEvCnt) is less than 16. This means that there are up to 15 events which were triggered and have to be sent to the ROD. In this case the trigger is issued to the FE's. If the PendingEvCnt equals 16, the L1 is not issued and the event is lost. The last event before the lost one informs the ROD of how many events were actually lost.
  The TTC block also generates a hierarchy of reset signals for the MCC and the FE's: BCR (Bunch Counter Reset), ECR (Event Counter Reset), SyncFE (Reset's for the FE's).

The BCR resets the bunch crossing counter in the MCC, whose least significant 8 bit are recorded in the event headers transmitted to the ROD. The ECR resets the whole event data not yet reconstructed and sent to the ROD. This command is a fast way of resuming normal data taking. BCR and ECR signals will be sent periodically by the data acquisition software to ensure synchronization. The SyncFE command creates reset signals for the FE's. The width of the reset signal is a parameter of the command. The FE's interpret the width of the reset signal and generate different levels of internal resets. These are the following:

- *Sync Reset* resets the internal trigger FIFO and deletes all events currently stored in the FE. Width < 4 x 25 ns.

- *Soft Reset* resets all digital readout circuitry to 'empty' state. This includes resetting the hit logic in the individual pixels, the logic that transfers hits to the end-of-column region and the EOC buffers themselves. Also the BCID counters are reset to zero. No configuration data is affected. Width < 8 x 25 ns.

- *Hard Reset* performs all the features of a Soft Reset and additionally resets the Global Register. Thus it resets the FE to the power-on state. Width ≥ 8 x 25 ns.

Finally the TTC logic issues the calibration strobes to the FE's. This is used to calibrate the FE analog part on a pixel-by-pixel basis. The delay of the strobe with respect to the clock (XCK) phase can be adjusted with 6-bit resolution (roughly 0.5 ns steps). Also the delay range can be selected by a 4-bit value.

- *Event Building.*
  Data received from the FE's, in response to a L1 trigger, are deserialized and buffered in one 128 word deep FIFO per FE chip. Thus the 16 data flows are derandomized and the event builder can extract ordered hits to prepare them for transmission to the ROD.
  Each FE sends data as soon as they are available with two constraints: the events are ordered by event number and for each event an end-of-event (EoE) word is generated. EoE is also sent in case of an empty event to preserve synchronization in the data readout. During normal operation each Receiver writes hit data from the FE to the Receiver-FIFO. When an EoE is recognized it is written to the ReceiverFIFO and a flag is set in the corresponding Scoreboard in the Event Builder.

If the Receiver FIFO becomes full while storing data the partially writ-
ten event will be deleted from the FIFO and a 'truncated event' flag is
stored in the FIFO instead. Still the EoE is written to the Receiver-
FIFO and flagged in the Scoreboard.

As soon as the Event Builder finds all EoE words for one complete
event in the Scoreboard register, it starts building and transmitting
the event. The output data stream starts with the bunch crossing ID
(BCID) and the L1 ID (L1ID) corresponding to this event. Then it
starts fetching data from the ReceiverFIFO's. When the Event Builder
finds an EoE in the data it deletes the flag in the Scoreboard register
and moves on to the FIFO for the next FE.
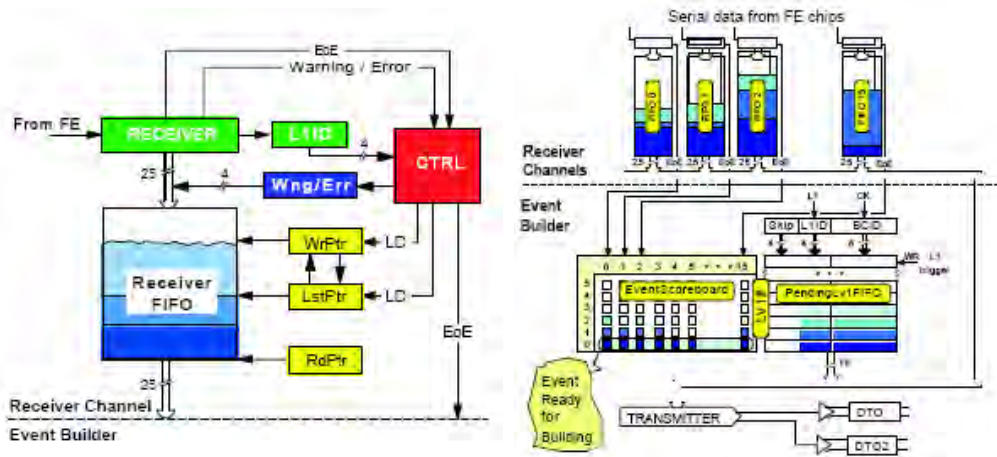
Figure 3.13: Logical schemes of the Receiver block (left) and the Event
Builder (right) as mentioned in the text. [29]

The MCC has two separate data output channels which can be used to
transmit data at different rates. This is necessary as the occupancy in the
B-Layer is significantly higher than in layers 1 and 2. There are four different
modes in which the MCC can utilize these output channels to transmit data:
(see also figure 3.14)

- 40 MBit/s Single-link:
  The same data is sent on both links. Transitions happen on the rising
  edge of the XCK clock. The total data rate is 40 MBit/s. This is the
  readout mode for the modules in Layer 2.

- 40 MBit/s Dual-link:
  The data stream is split over the two links. One link transmits the

even bits of the stream (bit 0,2,4,...) and the second transmits the odd bits (bit 1,3,5,...). Transitions happen on the rising edge of the XCK clock. The total data rate is 80 MBit/s.

- 80 MBit/s Single-link:
  The same data is sent on both links. Transitions happen on the rising and the falling edge of the XCK clock. The total data rate is 80 MBit/s. This is the readout mode for the modules in Layer 1 and the disks.

- 80 MBit/s Dual-link:
  The data stream is split over the two links. One link transmits the even bits of the stream (bit 0,2,4,...) and the second transmits the odd bits (bit 1,3,5,...). Transitions happen on the rising and the falling edge of the XCK clock. The total data rate is 160 MBit/s. This is the readout mode for the modules in the B-Layer.

Only event readout uses the higher bit rates, readout of configuration data is always performed with 40 MBit/s.
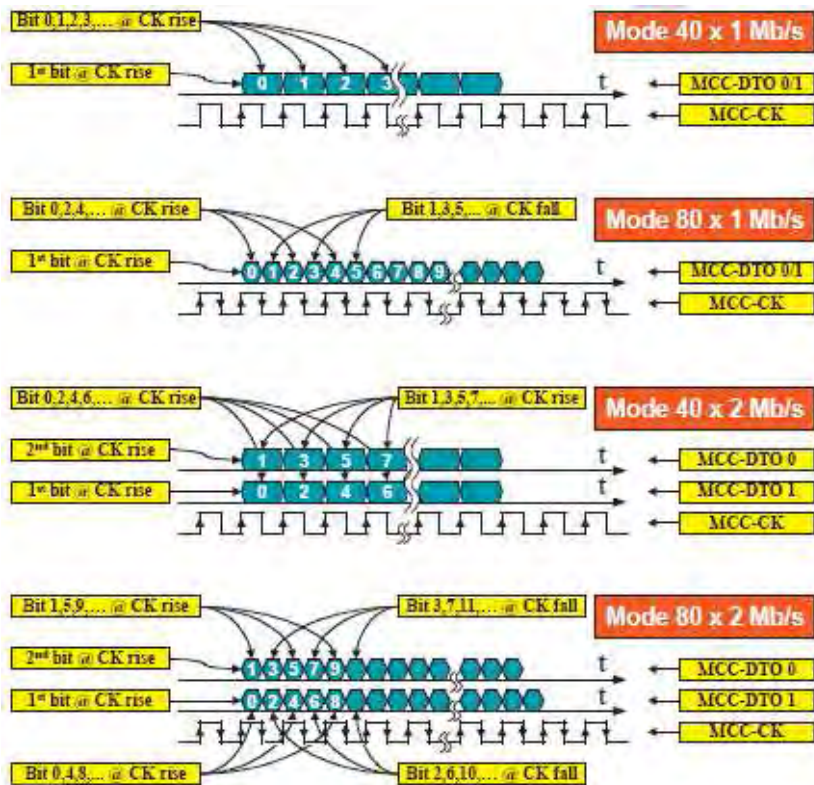

Figure 3.14: Data encoding for the different MCC bandwidth settings. [29]

3.3   The readout and control system

The readout and control system [19] for the pixel detector employs optical
data transmission between the detector and the readout system, located in
the USA15 cavern, approximately 80 meters away from the detector. Optical
data transmission was chosen because it provides high bandwidth together
with low noise and little signal attenuation. The optical links are based on
GaAs Vertical Cavity Surface Emitting Lasers (*VCSELs*), emitting light at a
wavelength around 850 nm, and epitaxial silicon *PiN* diodes. The on-detector
opto-electrical interface is the *optoboard* which receives and transmits the
signals for 6 or 7 modules (the actual number depends on the position in the
detector, see chapter 3.9). The optoboard is connected via about 80 meters
of optical fibers to the off-detector interface, the Back Of Crate card (*BOC*).
One fiber per module is used to transmit clock and control signals and one
or, as the B-Layer modules transmit data in 80 MBit/s dual-link mode, two
fibers are used to transmit event data to the off-detector electronics. Data
transmission on the fibers is unidirectional.

Each BOC is paired with a ReadOut Driver (*ROD*) that is responsible
for configuring the modules, propagating triggers and clock from the ATLAS
trigger system and for formatting the event data returned by the module.
The second purpose of the ROD is detector calibration and monitoring. Up
to 26 detector modules are connected to a single ROD, depending on the read-
out bandwidth of the modules. The RODs and BOCs are located in nine 9U
VME crates which are controlled by a Single Board Computer (*SBC*) each.
As the SBC controls a crate filled with up to 16 RODs it is also referred to
as the ROD Crate Controller (*RCC*).
After formatting the data collected from the modules, the ROD transmits
these event fragments to the ATLAS read-out subsystem (*ROS*) via an op-
tical ethernet connection (*S-Link*). The event fragments are stored in the
Readout Buffers (*ROBs*) contained in the ROS until they are provided on
request to the following stages of the event selection.
Additionally in each of the ROD crates there is a TTC Interface Module
(*TIM*) installed, which propagates the LHC machine clock and the level 1
trigger signals to the RODs and BOCs in the crate.
Figure 3.20 shows a schematic overview of the readout and control system.

3.3.1   Optoboard

The optoboard [30] is a $2 \times 6.5 \, \text{cm}^2$ printed-circuit board, implemented on
a beryllium-oxide (BeO) substrate for heat dissipation reasons. The opto-
boards are located at PP0 in the detector volume, about 1 meter away from

Figure 1: BiPhase Mark encoding

Figure 3.15: BPM encoding principle.

the interaction region. Depending on the exact position inside the detector, the optoboards are connected to 6 or 7 modules by special aluminum cables (see chapter 3.4). Supply voltages and electrical signal connections are made through a 80-pin fine-pitch connector that plugs in to the PP0 kapton PCB. LVDS signals received from the modules are converted by the VCSEL Driver Chip (*VDC*) into single-ended signals suited to drive the VCSEL channels. The optical signals are then transmitted to the RODs via optical fibers using a Non-Return to Zero (*NRZ*) signal.

As only one fiber is used to transmit the clock signal and the configuration data from the ROD, the signals have to be encoded into one data stream. This is done by using BiPhase Mark (*BPM*) encoding. BPM encoding maps one bit of the original data stream onto two bits in the encoded stream. A logical zero in the data is represented as two equal bits (00 or 11) while a logical one in the data is represented as two different bits (10 or 01). The logical level of the BPM encoded stream changes at every rising edge of the original clock. Therefore a long stream of zeros in the data produces a signal that looks like a 20 MHz clock signal, while a stream of ones produces a 40 MHz clock-like signal (see figure 3.15).

The BPM encoded clock and configuration signals are received by a 8-channel PiN-diode array. It transmits the signal electrically to a special ASIC[10], the Digital Opto-Receiver Integrated Circuit (*DORIC*, see below) which decodes the BPM signal. The clock and data signals are transmitted to the MCC differentially.

The VCSEL and PiN arrays mounted on the optoboards are chosen for an operating temperature of 10°C. It has been found that reliable operation of the optical link strongly depends on optoboard temperature. When mounted on the service panels, the optoboards are connected to a separate cooling circuit

---

[10]Application Specific Integrated Circuit

Figure 3.16: Optoboard mounted on PP0. Visible are the two VDC chips and the optical fibers plugged into the connector housing.

and to active heaters to keep their temperature stable. The heaters consist of 4 resistors glued to the optoboard with a total resistivity of 600 Ω. Heaters for six optoboards are connected in parallel on a flexible kapton printed circuit board.

Two different flavors of optoboards are used in the pixel detector. The majority of the boards are referred to as D-Boards. They are equipped with one 8-channel VCSEL array and provide one data channel per module. Thus they can only be used for data transmission up to bandwidths of 80 MBit/s single-link (see chapter 3.2.3). These boards are used to read out the modules of layers 1 and 2 and the disk modules. For the B-Layer modules, transmitting data in 80 MBit/s dual-link mode, so-called B-Boards have to be used. These are equipped with two 8-channel VCSEL arrays providing two data links per module. Therefore the B-Boards are connected to two outgoing fibers per module. Figure 3.16 shows a photograph of a D-Layer optoboard, view from the top side. The two VDC chips are visible, as well as the optical fibers connected to the board.

In total, 272 optoboards - 44 B-Boards and 228 D-Boards - are installed in the detector. See figure 3.17 for the layout of the two flavors of optoboards.

VDC circuit

The VDC is used to convert the differential signals from the MCC to single-ended signals, appropriate to drive a VCSEL in a common cathode array. The output current of the VDC is adjustable from 0 to 20 mA by means of an external control voltage. This voltage (Viset) is internally converted into a control current (Iset) and amplified by a factor of 10. The dim current, defining the light output for a logical zero, is approximately 1 mA. This increases switching speed for the VCSEL. The nominal operating current is 10 mA, the signal rise time is about 1 ns. In order to minimize power supply noise, the current consumption is kept constant, independent of whether the VCSEL output is in dim (logical zero, OFF) or bright (logical one, ON) state. This is done by means of a dummy driver circuit that compensates the different current consumption in ON and OFF state of the VCSEL.
The VDC is a four channel ASIC, therefore a D-Layer optoboard contains two VDCs and one VCSEL array, while a B-Layer board contains four VDCs and two VCSEL arrays.

DORIC

The DORIC decodes the BPM encoded signals received by the PiN diode and sends the clock and data streams to the MCC using differential signals. The input current ranges from 40 to 1000 μA per channel. The decoded clock signal is required to have a duty cycle of (50 ± 4) % and a jitter of less than 1 ns. The bit error rate is required to be less than 10^{-11} after irradiation with a dose of 50 Mrad.
The amplifier/discriminator circuitry of the DORIC incorporates an internal feedback to adjust the threshold, so that a duty cycle of 50 % can be maintained over the complete dynamic range of the input current. Also a delay-lock-loop circuit is implemented to generate the 40 MHz clock signal. In case this circuitry does not work as expected, the output is a 20 MHz clock signal as decoded from an all-zero BPM signal. A Reset line is implemented to reset the delay-lock-loop circuit in case of a failure.
The DORIC, just like the VDC, is a four channel ASIC. As clock and command signals are treated the same way for both optoboard flavors, every optoboard contains two DORICs and one PiN diode.

3.3.2 Optical fibers

The data transmission between the optoboard and the off-detector readout electronics employs optical fibers. As the requirements in terms of radiation

Figure 3.17: Layout of D-Layer and B-Layer optoboards. Top side layouts are different while bottom side layout is the same. [30]

hardness vary along the length of the fibers, a two-stage fiber system is used. Radiation hard *stepped index multimode* (*SIMM*) fibers are used inside the detector volume. These fibers have a core diameter of 50 μm. At a point roughly 9 m away from PP1, close to the muon chambers, the SIMM fibers are spliced to radiation tolerant *graded index multimode* (*GRIN*) fibers. There are two different kinds of GRIN fibers in use. The fibers for data transfer from the optoboard to the BOC have a core diameter of 62.5 μm, while the fibers in the other direction have a core diameter of 50 μm.

The fibers are routed along the backbone of the service quarter panel (see chapter 3.9) in 8-fiber bare ribbons up to PP1. Two 8-fiber ribbons are combined into one MT16 connector attached to one 16-fiber ribbon. The ribbons are bundled into 80 m long cables of four ribbons each.

3.3.3   The Back-of-Crate card

The BOC [31] is the off-detector interface of the optical readout chain. It has been developed at Cavendish Laboratory in Cambridge for the SCT and the pixel detector. The board layout is the same for both detectors, the only difference being the number of used channels. The SCT BOCs handle 48 modules each, while for the pixel BOCs the number of modules connected to one BOC depends on the position in the detector. One BOC will be connected to 26 Layer 2 modules, 13 Layer 1, 12 disk modules or to 6 or 7 B-Layer modules. The mapping is driven by the readout bandwidth used.

For more information see next section.
The BOC card provides the following functionality:

- Reception and distribution of the ATLAS clock signal as provided by the TIM[11].

- Reception of digital control signals from the ROD, BPM encoding and optical transmission to the modules.

- Reception of readout data from the modules, conversion into electrical signals and transmission to the ROD.

- Reception of the parallel event data stream from the ROD and transmission to the *ReadOut Buffer* (*ROB*) via S-Link.

- Adjustment of parameters of the optical interface (timing, laser current, threshold, masking of channels) on the sending and the receiving side.

- Implementation of an off-detector laser interlock.

The layout of the back-of-crate card can be divided into several blocks that provide parts of the functionality mentioned above. The main functions are performed in seven CPLDs[12]. One of them performs the main control of the BOC card and the communication with the ROD. Two CPLDs handle the communication with the TX- and RX-plugins and four CPLDs are operating in the receiving section, registering data and splitting data streams (see below).

- **Clock Section**

clock signal	purpose	delay capability
system clock	clock received from TIM	-
A-clock	clock for internal chip functionality	-
P-clock	clock encoded into BPM signal	0 to 24 ns, 1 ns steps
B-clock	clock for normal data recovery	0 to 24 ns, 1 ns steps
V-clock	clock for 80 Mb/s sampling	0 to 49 ns, 1 ns steps
		0 to 10.2 ns, 40 ps steps
ROD-clock	copy of system clock for the ROD	-

Table 3.1: Overview of the clock signals on the BOC card.

---

[11]TTC Interface Module, see section 3.3.5
[12]Complex Programmable Logic Devices

parameter	control range	purpose
BPM inhibit	ON/OFF	prohibiting data to be encoded
BPM fine delay	0 - 35.56 ns	delay for the signal
	280 ps steps	to the module, channel wise
BPM coarse delay	0 - 775 ns	delay for the signal
	25 ns steps	to the module, channel wise
mark space ratio	settings: 0 to 31	adjust MSR between 30:70
		and 70:30, channel wise
laser current	settings: 0 to 255	adjust laser forward current
		between 0 and 18 mA

Table 3.2: Overview of the control parameters in the transmission section of the BOC card.

The BOC receives the 40 MHz system clock from the TIM, located in the same crate. The clock signal is demultiplexed into 5 clock signals used for different tasks on the BOC card. Most of these clock signals can be delayed with respect to the received clock to optimize timing. Table 3.1 shows details on the clock signals.

- **Transmission Section**
  Commands from the ROD are transmitted to four specialized plugin PCBs on the BOC that hold BPM encoder chips and VCSEL arrays. These *TX-plugins* are referred to as TxA, TxB, TxC and TxD, depending on the slot on the BOC card they are plugged into. On the TX-plugins [32] the control data for each module is BPM encoded separately and transmitted to the optoboard. The BPM encoding is shown in figure 3.15.
  The clock signal used for BPM encoding is the P-clock which is multiplexed for all module channels. It can be delayed separately for each channel inside the BPM chip. There is a fine delay between 0 and 35.56 ns adjustable in steps of 280 ps and a coarse delay between 0 and 775 ns adjustable in steps of 25 ns. Thus every module can be timed individually with respect to the LHC clock. There are also adjustable parameters for the laser output power, the mark-to-space ratio *(MSR)* of the BPM signal and an inhibit signal blocking the data stream from being encoded in the BPM signal. Adjustable parameters are summarized in table 3.2.

  The parameters of the BPM chip have been found to influence each other. Therefore a rather complicated interactive tuning procedure has been defined. Details on the procedure can be found in [31].

- **Receiving Section**
  The optical signals from the modules are received by specialized plugin PCBs on the BOC. These *RX-plugins* hold a 12-way PiN diode array, of which 8 channels are used for pixel detector readout. The remaining four channels are not connected. The RX-plugins are named RxA - RxD, just like the TX-plugins.
  The electrical signals from the PiN diodes are transmitted to an amplifier chip, the Driving and Receiving IC (DRX). The DRX features an adjustable threshold, set channel wise by a 8-bit DAC. It amplifies the received signals and transmits them differentially to the BOC card. Here the phase of the received data is adjusted with respect to the BOC clock, and the data is passed on to the ROD as 40 Mb/s streams. Therefore a 80 Mb/s stream coming from a Layer 1 or a Disk module is split into two streams, a 160 Mb/s stream from a B-Layer module is split into four data streams. The phase adjustment is done using Phos4 chips that can apply a delay between 0 and 24 ns to each individual data stream in steps of 1 ns. The splitting of the data streams into 40 Mb/s streams is done in one CPLD per RX-plugin. To split a 80 Mb/s data stream into two separate 40 Mb/s streams the CPLDs use the B-clock to sample one bit of the incoming stream and the V-clock, which is inverted with respect to the B-clock, to sample the other bit. Therefore one of the resulting bitstreams contains the odd bits of the original stream (bits 1,3,5,...), while the other bitstream contains the even bits (bits 0,2,4,...). These bitstreams are then passed onto the ROD.
  As the number of data streams being passed to the ROD is fixed to 32, the splitting of the data streams implies that, due to the different MCC bandwidths, differing numbers of modules are connected to one BOC card. A single BOC can read out either 26 modules in Layer 2, 12 Disk or 13 Layer 1 modules or 7 B-Layer modules. Due to the routing of the data streams on the BOC card, only RxA and RxD can be used to receive data in 80 or 160 Mb/s. In order to achieve a balanced load in 80 and 160 Mb/s mode a routing board is attached to each BOC card. It distributes the incoming data streams evenly over the CPLDs in case not all RX-plugins are used, and synchronizes the data streams (see figure 3.18). The resulting data streams are passed directly to the formatters on the ROD.
  The adjustable parameters in the receiving section are summarized in table 3.3.

parameter	control range	purpose
RX threshold	settings: 0 to 255	threshold inside the DRX chip
		0 - 250 μA, channel wise
RX data delay	0 to 24 ns	channel wise synchronization
	1 ns steps	for the data signal
B-clock	0 to 25 ns	global data sampling
	1 ns steps	clock
V-clock	0 to 49 ns	global inverted clock
	1 ns steps	for sampling at 80 Mb/s
V-clock fine delay	0 to 10.2 ns	fine adjustment for
	40 ps steps	the V-clock, global

Table 3.3: Overview of the control parameters in the receiving section of the BOC card.

- **S-Link Section**
  Data is sent to the ReadOut Buffers via an optical gigabit ethernet connection using the S-Link protocol [33]. The BOC card carries the HOLA[13] card that drives the 32 bit parallel bus at 40 MHz. No other functionality is required on the BOC side.

Apart from the functionality concerned with data transmission and synchronization, the BOC also provides some monitoring and interlock features. The voltages and currents on the 12 V supply lines for the PiN diodes on the RX-plugins can be measured and used to tune the light output of the optoboard. Also the temperatures of the BOC card underneath the RX- and TX-plugins can be measured using a NTC resistor each.

To provide laser safety, the BOC card has to switch off the lasers on the TX-plugins whenever work is ongoing close to the plugins. Two interlock lines are distributed to each BOC card in a crate via the back-plane. One of the lines is connected to the door of the rack and triggers an interlock when the door is opened. This is called the 'local interlock'. The second line is connected to the on-detector laser interlock system, and triggers an interlock when work is being done in the PP1 region where the optical fibers are accessible. This is referred to as the 'remote interlock'. The two lines are active low, meaning that a voltage of 3.3 V on the lines corresponds to 'interlock not set'. The lines are combined on the BOC with a logical AND, such that the lasers can only be switched on if both lines are at 3.3 V.

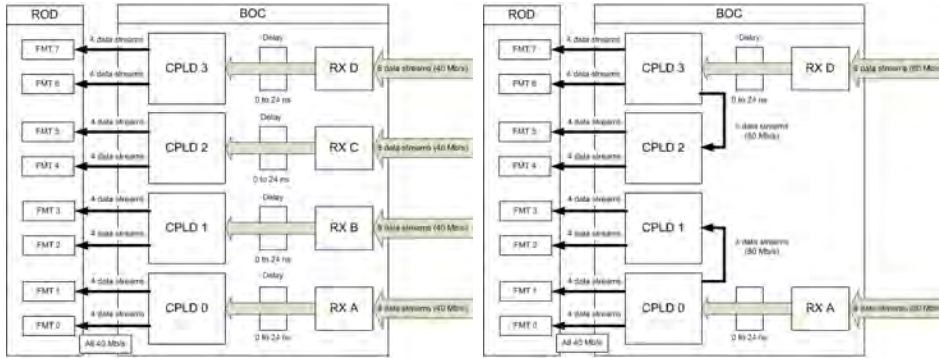[13]High speed Optical Link for ATLAS

Figure 3.18: Dataflow in the BOC in 40 Mb/s mode (left) and 80 Mb/s mode (right).

3.3.4 The Read-Out Driver

Like the back-of-crate card the ReadOut Driver (*ROD* [34]) is common to the SCT and pixel detector. The ROD is a hybrid design of Field Programmable Gate Arrays (FPGAs) and Digital Signal Processors (DSPs). Although the hardware of the ROD card is the same for SCT and pixels, it should be noted that differences in the firmware exist.

The ROD is the main part of the readout and control system. It passes the trigger received from the TIM, formats the event data from the modules before sending it to the ROS, and it provides monitoring capabilities during data taking. It also generates the commands for the modules and sends them as a serial stream. Commands include Level-1 triggers, bunch-crossing and event counter resets, module configuration data and calibration commands. Thus the ROD can also be used during detector calibration and testing.

The ROD has two distinct modes of operation, data-taking and calibration mode. During data-taking, trigger signals are received from the TIM and being passed onto the pixel modules, while event data is processed in the ROD data-path before sending out to the ROS via S-Link. In calibration mode the ROD generates the trigger signals internally, and the event data is processed and histogrammed in a farm of four Slave DSPs.

The ROD layout can be separated into different functional blocks that are shown in the block-diagram in figure 3.19:

- **Controller FPGA and Master DSP**.
  The main function of the Controller FPGA and the Master DSP (*MDSP*) are to setup the ROD data-path, to execute BOC and pixel module configuration and to process and propagate triggers.
  The MDSP has 64kB fast internal memory and additional 32 MB of

slower external memory. It has BOC and ROD registers connected to one of its external memory interfaces. Thus all BOC and ROD registers are set through the MDSP. Additionally it has two serial output ports through which configuration data is sent to the pixel modules. In calibration mode the serial ports are also used for sending trigger signals.

During data-taking the trigger and event description information (Level-1 ID, bunch-crossing ID and trigger type) are received from the TIM. The Controller FPGA processes these informations and creates the trigger signals required by the pixel modules.

- **Formatter FPGAs.**
  The eight identical Formatter FPGAs receive the 40 MHz data-streams from the BOC, that contain the event data. The subsystem specific firmware of the Formatter FPGAs contains four links per formatter in the pixel version (the SCT version contains 12 links). Thus in 160 Mb/s mode one Formatter is used to process the data from one module. In this way the number of Formatter FPGAs and links per Formatter determines the number of modules that can be connected to one ROD. At a bandwidth of 40 Mb/s (Layer 2) up to 32 modules could be connected to one ROD (due to the geometry of the detector only 26 modules are connected). In 80 Mb/s mode (Layer 1 and Disks) 16 modules can be connected (13 are connected) and in 160 Mb/s mode (B-Layer) only 8 modules can be connected (depending on the position in the detector 6 or 7 modules are connected).
  The Formatters carry out the serial to parallel conversion of the event data and the de-randomizing of the event fragments. They are also used to check for data errors and to inform the Controller FPGA when a complete event has been received.

- **Event Fragment Builder FPGA.**
  The Event Fragment Builder (*EFB*) collects the output of the Formatters, checks the Level-1 and Bunch-Crossing IDs (*L1ID* and *BCID*) and generates event headers and trailers. It receives information from the Controller FPGA regarding trigger type, L1ID and BCID for header and trailer generation, and compares them with the data collected from the Formatters by each of two engines. The EFB flags events that contain errors using dedicated bits in the event header and trailer.
  Once an event is completely processed and headers and trailers are generated the EFB passes the data to the Router FPGA.

- **Router FPGA and Slave DSP farm.**
  The purpose of the Router FPGA is to send formatted event data to the ROS via S-Link and/or to the Slave DSPs (*SDSPs*). The four SDSPs are used for error counting, as well as event capture and histogramming. The SDPSs are capable of doing floating-point calculations, they are equipped with 128 kB of fast internal memory and 256 MB of slower external memory. In calibration mode event data is passed to the SDSPs using DMA[14] and the SDSPs are used for histogramming and further data processing (including averaging and fitting).

Once the data-path in the ROD is setup, it is completely handled by the FPGAs and does not require any intervention from the Master or Slave DSPs. Therefore during data-taking the SDSPs can be used for event capture and ROD-level online monitoring of various quantities that can give detector experts information on the performance of different parts of the system.
In the event that the S-Link is receiving data from the ROD faster than it can be transfered to the ROS, the S-Link applies back-pressure to the ROD. This halts the output of event data from the ROD EFB. When the EFB memory is full and back-pressure is still applied, it is propagated to the Formatters, stopping transmission of event data from the Formatters. If a critical limit in the formatter memory is reached a ROD BUSY signal is issued to the TIM to halt triggers from ATLAS. The probability for this to happen is reasonably low up to occupancies of 2.5 times the projected one at high luminosity.
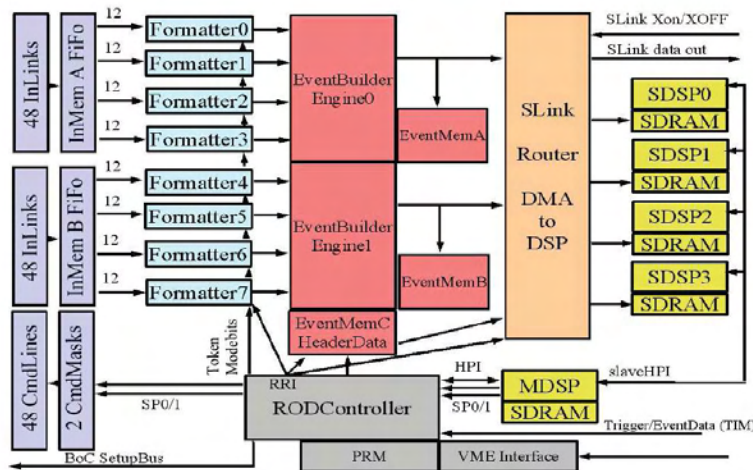

Figure 3.19: Block diagram of the readout driver.

---

[14]Direct Memory Access

type	acronym	purpose
Clock	BC	Bunch-Crossing Clock
Fast command	L1A	Level-1 Accept
	ECR	Event Counter Reset
	BCR	Bunch Counter Reset
	CAL	Calibration signal
Event ID	L1ID	24-bit Level-1 trigger number
	BCID	12-bit Bunch-Crossing number
	TTID	8-bit Trigger type (+2 spare bits)

Table 3.4: Overview of TIM commands.

3.3.5   The TTC Interface Module

The TIM interfaces the ATLAS Level-1 trigger system with the RODs. Its primary functions are:

- To transmit fast commands and event ID from the ATLAS Level-1 trigger system to the RODs with minimum latency. The clock is first transmitted to the BOCs which pass it to the RODs.

- To generate a crate BUSY signal from a gated OR of the ROD BUSY signals. This is passed to the Central Trigger Processor (*CTP*) to stop it from sending triggers.

- To generate and send clock, fast commands and event ID to the RODs when the CTP is not active.

The information sent to the RODs is summarized in table 3.4. The commands are distributed via the back-plane of the 9U VME crate.

  The TTC signals are either received optically by a TTCrx chip [35] from the ATLAS Level-1 trigger system, or they can stem from various sources in stand-alone mode. It is also possible to connect multiple TIMs for stand-alone multi-crate operation.

In stand-alone mode the clock and fast commands can either be generated internally (80.16 MHz crystal oscillator on-board, commands generated on request by the crate controller), or they can be input from external sources in either NIM or differential ECL signal standard. Furthermore all back-plane signals are available on the front-panel of the TIM card as differential ECL signals to allow TIM interconnection.

3.3.6 The Read-Out Subsystem

The ReadOut Sub-system (*ROS*) is a PC that typically houses 4 custom PCI cards that implement the ReadOut Buffers (*ROBs*). These cards are referred to as *ROBINs*. Each ROBIN houses 3 independent buffers. The ROS handles the data requests issued by the High-Level Trigger (*HLT*) system for all of the 12 buffers. Upon reception of a data request the ROS collects the event data from the relevant ROBIN modules, combines them into a ROS fragment and sends it to the requester. For further information see chapter 2.3.2.
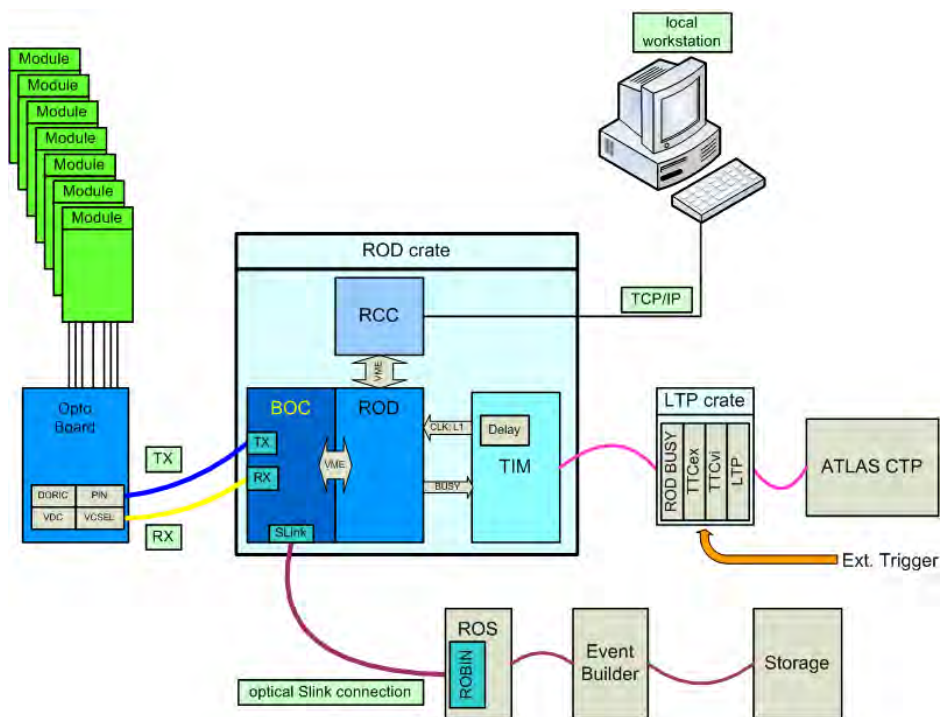

Figure 3.20: Schematic view of the readout system as described in the text above.

3.4   General layout of Detector Services

Various requirements drive the highly modular layout of the detector services:

- integration with services of other subdetectors in common gaps along the routing

- services have to penetrate the thermal barrier and the Faraday cage

- overall geometry of the pixel detector

- the pixel detector is the last component to be installed in ATLAS

Thus services can be disconnected at a couple of breakpoints throughout the routing. These *Patch Panels (PPs)* are situated in different parts of the detector.
The naming of cables between these patch-panels is motivated by the naming of the patch-panels themselves. For example the patch-panel closest to the detector is PP0 and the cables connecting it to PP1 are called Type1 cables. Figure 3.21 shows a simplified schematic view of the services layout, including cable lengths.
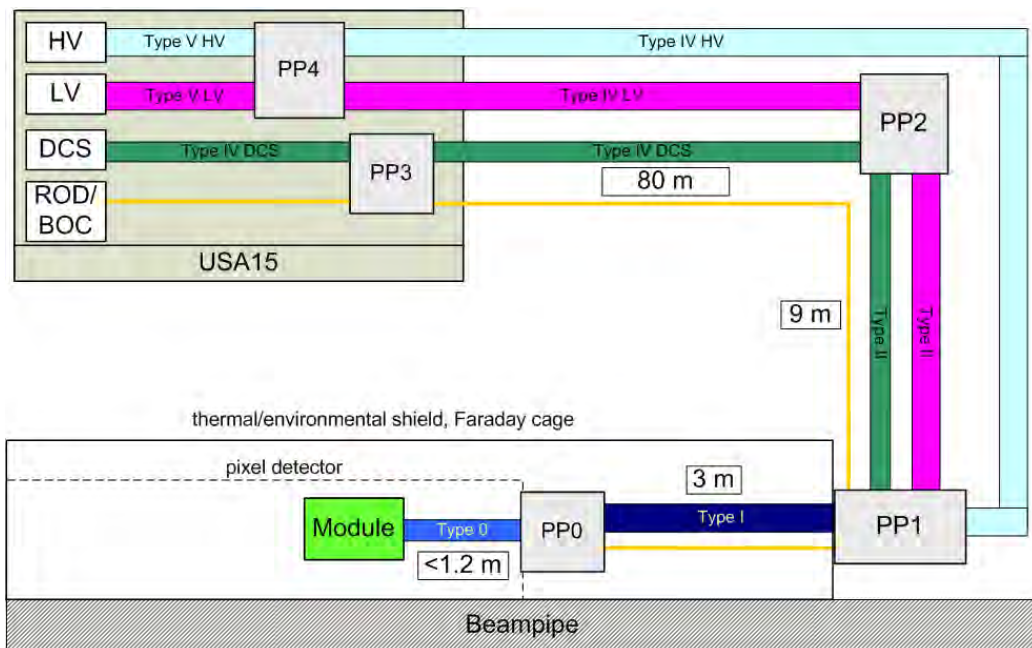

Figure 3.21: Simplified schematic view of the layout of the services for the pixel detector.

- **PP0**
  Situated on the detector side of the about 3m long service panels, that
  are part of the pixel package inside the enclosed pixel volume, PP0 is
  implemented as a flexible kapton PCB. One PP0 holds connectors for
  the cables coming from up to seven modules. These custom aluminum
  cables (*Type0 cables*) contain the power lines for the module low and
  high voltages, as well as the lines for the differential signals and the
  wires connected to the temperature sensor on the module (NTC[15]).
  The power and NTC lines are routed through PP0 electrically. PP0
  also houses one optoboard each (see chapter 3.3.1). The number of
  modules connected to one PP0 depends on the position in the detector.
  Either 6 disk modules or 6 or 7 barrel modules can be connected. For
  more detail on the grouping of modules per PP0 and of PP0s per service
  panel see chapter 3.9.

- **PP1**
  At the endplate of the pixel volume all services have to penetrate the en-
  vironmental barrier and the Faraday cage without breaking the thermal
  and gas seals. Therefore an endplate for the Pixel Support Tube (*PST*)
  has been designed that provides gas tight feed-throughs for electrical
  and optical cables, as well as pipes for dry gas and coolant. This end-
  plate is referred to as PP1 [36]. The number of feed-throughs required
  for the services inside the pixel package is very large. Consequently
  the packing density is very high. Therefore the power and NTC lines
  in the Type1 cables are connected to kapton PCBs, which outside of
  the pixel volume are terminated by round 64-pin LEMO-F connectors
  where the Type2 cables are plugged in. Special requirements concern
  the feed-throughs for the cooling pipes, as they have to move axially
  by few millimeters due to thermal contraction when the cooling starts.
  The fiber ribbons, running along the back-bone of the service panel as
  uncoated 8-way ribbons, are grouped together into 16-way MT con-
  nectors at PP1. All connectors for one quadrant are combined in a
  connection plate. This way all optical cables outside the detector vol-
  ume (corresponding to Type2 cables) can be connected at once to PP1.
  The optical cables contain four 16-way ribbons per cable.
  The grounding scheme for the pixel detector has a star topology, where
  all potentials are connected to one common ground-point. This point is
  located at PP1 on the C side of the PST and has one single connection
  to the common Inner Detector ground (see chapter 3.6).

---

[15]Negative Temperature Coefficient resistor

Down to PP1 all services are pre-installed in the experimental cavern before the pixel detector is lowered and installed inside the Inner Detector. Furthermore this is where the services are disconnected for maintenance work on the pixel detector. It should be noted that PP1 is the only break-point where all services for the pixel detector break.

- **PP2**
  Patch-panel 2 [37] is separated into different blocks, located in different parts of the detector. The power, environmental and NTC/opto boxes are sitting just outside the first layer of muon chambers, approximately 9 m away from PP1. The environmental and the NTC/opto boxes make passive connections between the radiation hard Type2 cables inside the detector volume to radiation tolerant Type3 cables. The power boxes are integral parts of the pixel powering scheme. Not only do they provide the transition from Type2 to Type3 cables, but they also house the boards that regulate the high current supply voltages and compensate for the voltage drop over Type1 and Type0 cables. This is the main reason for PP2 to be situated that close to the pixel detector. The voltage regulators in the power boxes dissipate about 800 W of heat. This heat is removed via the room-temperature C_4F_{16} cooling system that is also used for the TRT. For more details on voltage regulation at PP2 see chapter 3.5.
  A total of 26 power boxes, 24 NTC/opto boxes and 8 environmental boxes are used for the pixel detector.

- **PP3**
  The PP3 crates [38] are situated on the service platforms around AT-LAS. Patch-panel 3 houses front-end components of the Detector Control System (*DCS*), as well as a passive component where rerouting of services for the optical link is done. The components of PP3 are referred to as the Building Block Interlock and Monitoring (*BBIM*) crates, the same object without interlock functionality (*BBM*) and the passive Opto PP3.
  One BBIM box consists of an Embedded Local Monitoring Board (*ELMB* [39]) and four Interlock-Boxes (*IBoxes*). The ELMB is a common AT-LAS development. It provides a 64 channel ADC, 16 digital I/O signals and a controller for CAN-bus communication. In the BBIM it is used to measure the temperatures of NTC thermistors mounted on the pixel modules, opto boards and voltage regulator boards at PP2. The ELMB sends the values to the DCS control station for monitoring. Additionally the I-Boxes receive the analog NTC voltages and after comparison

with fixed thresholds generate digital temperature interlock signals. These are sent to the interlock system via Type4 cables. Four BBIMs can be mounted in one 6U high 19$''$ crate. For more details on the interlock system see chapter 3.7.

The BBM consists only of an ELMB used to read out various kinds of environmental sensors (temperature sensors scattered throughout the pixel detector volume, humidity sensors). The information is sent to the DCS control station via CAN-bus. Four BBMs can be mounted in one 3U high crate.

As mentioned above, the opto PP3 is a purely passive component used to extract the OPTO-NTC wire from the NTC/OPTO Type3 cable to be used in the BBIMs. The remaining opto services (low current supply voltages and DORIC reset line) are combined in a Type4 cable. The opto PP3 is mounted in a 3U high crate.

- **PP4**

  Patch-panel 4 provides power distribution for module high and low voltages as well as current measurements for the individual modules. It is divided into low voltage (LV) and high voltage (HV) PP4 crates. Both types of crates are mounted in the US15 and USA15 caverns.

  The layout of the interlock system and the interlock inputs on the low voltage power supplies (WIENER) drive the connection of one channel of the WIENER power supply, per supply voltage needed (VDD and VDDA), to 6 or 7 modules connected to one PP0. LV-PP4 [40] provides the fanout as well as a measurement of the current in each of the output lines. Current is measured as voltage drop over a 0.1 Ω resistor, and the measurement circuit is galvanically isolated from the ELMB ADC input by means of linear opto-couplers. The increasing sensor leakage current during the lifetime of the experiment makes it necessary to change the multiplicity of the high voltage connection dynamically. In the beginning high voltage for all modules connected to one PP0 will be supplied by one channel of the ISEG high voltage power supply. As the supplied current is limited to 4 mA eventually every module will be supplied by one ISEG channel. The change in multiplicity will be made in HV-PP4, which also provides measurement of the leakage current of the sensors.

For more details on the functionality of the different patch panels see chapters 3.5, 3.7 and 3.10.

3.5   The power supply system

The FE chips and the MCC require two supply voltages (*VDD* and *VDDA*) with their separate returns. These supply lines have to be able to deliver the fairly high currents needed during operation of the module. A separate bias voltage is required to deplete the sensor. Furthermore the optoboard requires one high current supply voltage (*Vvdc*) and two low current ones (*Viset* and *Vpin*). For typical values of the supply voltages and currents see table 3.5.

supply voltage	typical voltage [V]	typical current [mA]
on the module		
VDD	2.0	150 - 900
VDDA	1.6	90 - 1300
Vdep	150 - 600	0.001 - 2
on the optoboard		
Vvdc	2.5	200 - 300
Vpin	10	0.5 - 5
Viset	0.7 - 1.2	0.6 - 1.5
at the power supplies		
SC-OLink (Vvdc)	6	
WIENER (VDD)	7	
WIENER (VDDA)	8	

Table 3.5: Typical values for voltages and currents for the pixel detector

All of the power-supplies are located in the counting room cavern USA15, about 100m away from the devices that are powered. This long distance together with the relatively high currents on some of the lines lead to a significant voltage drop over the supply cables. To compensate for this, the voltages are regulated at PP2 using sense lines on the high and the low sides of the circuit, that run up to the Type0 connector on the module.

3.5.1   Module Low Voltage

The low voltages for the modules (*VDD* and *VDDA*) are generated by power supplies built by WIENER Plein & Baus, generally referred to as WIENER power supplies. Every power supply has 12 channels that are floating with respect to each other. Two of these channels together form the low voltage

supply for one half-stave or sector. One channel supplies VDD with an output voltage of 7V, the second supplies VDDA with an output voltage of 8V. The output voltages are chosen such, that despite the voltage drop over the Type4 cable, the input voltage for the regulators at PP2 is high enough for the regulators to work as designed, while being as low as possible to minimize power dissipation in the regulators as well as in the cables. Voltage and current on the single WIENER channels are measured at the output of the WIENER and can be monitored in the detector control system (*DCS*, see chapter 3.7).

Although the WIENER power supply offers the possibility of remote sensing to compensate for voltage drop, this functionality is not used because it would have doubled the number of wires per Type4 cable. This is undesirable for material budget as well as cost reasons. Moreover doing the voltage drop compensation at the LV power supply directly would provoke the risk of applying a voltage well above the breakdown voltage for the ASICs in the detector. The breakdown voltage for the 0.25 μm DSM process, in which the chips are produced, is specified to be 4V. Voltage much above the breakdown voltage would be applied, in case the current drawn by a group of ASICs would be reduced to zero over a time shorter than the response time of the power supplies. Therefore voltage regulation was chosen to be done at PP2, closer to the detector.

The WIENER supplies are installed in the counting rooms together with the corresponding low voltage PP4 boxes (*LV PP4*). The connections are made via very short pigtail cables on the backside of the WIENER supply. The current measurement on the single channels provided by PP4 is by default done on the high side of the circuit. The low side measurement can be implemented by adding another opto coupler board to the PP4 box to connect the existing measurement circuit to the ELMB.

Voltage regulation at PP2

Voltage regulators, housed in the PP2 power boxes, are employed to compensate for the voltage drop over the last 9 m of cable down to the modules. For the barrel modules the sense lines used for the regulation are connected to the supply lines at the Type0 cable connector (*ELCO*[16] connector) on the Pigtail. The wires are connected directly at the solder joint on the Flex for disk modules. In principle no current is flowing over these wires, so the voltage measured between the high and low side sense lines equals the voltage applied to the module. This voltage is being used for regulation by the

---

[16]The connectors used were produced by Kyocera ELCO.

radiation hard LHC4913 regulator ASIC developed by STMicroelectronics,
referred to as *ST-regulator*. It is a feature of this ASIC to draw a small
current ($\approx$ 40 mA) on the sense return line. This is being addressed in the
regulation circuitry.

16 regulators are loaded on a printed circuit board, referred to as the
motherboard. The output voltage is controlled by a digital trimmer, located
on another PCB which plugs into the motherboard. This PCB is referred
to as the *daughterboard*. It also contains circuitry to monitor output volt-
ages and currents of the galvanically isolated channels. One pair of mother-
and daughterboard regulates the voltages for the modules connected to one
PP0, as well as the VVDC voltage for the corresponding optoboard. Twelve
mother/daughterboard pairs, usually referred to as a *regulator board*, are
mounted into a crate that houses one controller board per crate and provides
cooling for the regulators. The Type2 cables connect to the back-plane of
the crate that acts as an output board, Type3 cables connect to the input
board mounted at the side of the crate. A total of 26 power boxes will be
installed in ATLAS.

The controller board provides two independent signals to disable the out-
puts of the regulator board. One is a global signal affecting all channels. It is
referred to as the KILL signal. Single channels can be switched off by means
of a channel wise INHIBIT signal. Voltage is only applied to the output of a
given channel if neither signal is active.
As mentioned above the output voltage is set by one digital trimmer per
channel which is supplied by the input voltage for this channel. Thus regu-
lators can only be set if the WIENER channels connected to this board are
switched on. As a safety measure a protection diode is implemented. In case
the high side sense line for one channel is disconnected, a LED[17] defines the
potential at the input of the ST-regulator, thereby limiting the maximum
output voltage. Further the regulation circuitry is implemented in a way
that limits output voltage to a narrow range around the required voltages.
The controller board is supplied by separate external voltages. Currently
studies are being carried out to determine whether it is possible to use spare
WIENER channels to power the controllers. In case the supply voltages for
the controller board switches off during operation, latches on the daughter-
board store the present status of the global KILL and the individual INHIBIT
signals. The latches are clocked by a signal issued to change the setting of
the digital trimmers. Thus detector operation is not disturbed by a power

---

[17]Light Emitting Diode

cut to the controller.

The regulation circuitry is implemented such that, in case a WIENER input voltage is being switched off, the regulator channel is inhibited. The latch for this channel stores the active INHIBIT signal. Thus when the input voltage is switched on again the regulator channel stays off until the operator manually issues the clock signal to the latch.

The controller board also houses an ELMB to provide measurements of output voltages and currents and for communication with the detector control system via CANbus. Most of the control signals for the regulator boards are generated and issued by an FPGA on the ELMB.

3.5.2 Module high voltage

The depletion voltage for the sensors is generated by 16-channel high voltage supplies, produced by ISEG. The power supplies can generate voltages up to 700 V with a maximum current of 4 mA. The ramping speed of the output voltage is adjustable and is usually set to around 10 V/s to avoid large electric field gradients in the sensor. If a channel reaches the current limit, the power supply can either switch it off or operate it as a current source. This protects sensors with low breakdown voltages from large currents, that could damage the sensor further and lead to thermal runaway.

At the beginning of the operation, sensors will be biased with 150 V. As the leakage current of the sensor chosen for module assembly is below 10 μA, it is possible to bias multiple sensors with one ISEG channel. Thus during initial operation the modules connected to one PP0 (6 or 7 modules) are biased by one ISEG channel. Due to radiation damage the voltage needed for full depletion will rise and the leakage current increases to a projected 2 mA. The maximum applicable depletion voltage is 600 V, driven by the specifications of the insulation of the HV cables provided by the vendor.

3.5.3 Supply voltages for the optoboard

For the operation of one optoboard three different supply voltages and a digital reset signal for the DORIC are required. The voltages and the reset are provided by a custom power supply, referred to as Supply and Control for the Optical Link (*SC-OLink*). The SC-OLink is implemented as a printed circuit board, that generates the required voltages using a commercial DAC set by an ELMB. The channels are galvanically separated up to voltage differences of 10V, and the output voltages and currents can be monitored. Linear opto-couplers are used to separate the measurement circuit from the ELMB. In total 272 SC-OLink cards are required. Four boards, together with

an ELMB, forma building block of the SC-OLink. Four of these building blocks are mounted in a 3U 19″ crate. The crates are situated in the US15 and USA15 caverns.

Vvdc supplies the 2 controller chips VDC and DORIC, and the VCSEL array. As the current drawn is rather large, the output of the SC-OLink is set to 6V and the voltage is regulated to 2.5V in the PP2 power boxes. Vpin is the depletion voltage for the PiN diode. As the current is very low, voltage drop is negligible and the output voltage of the SC-OLink is set to 10V. Viset is used as a reference current for the VDC that applies the current, multiplied by a factor of 10, to the VCSEL to set the light output power. The two low current supply voltages together with the DORIC reset signal are supplied together by one set of service lines. These services are grouped into one Type4 cable for two optoboards. Vvdc is routed to PP2 in a separate Type4 cable. The maximum ratings for the four channels of one SC-OLink are summarized in table 3.6, the typical values can be found in table 3.5.

The SC-OLink is powered via a transformer (65 VA) to have the individual

	voltage	current
Vpin	20 V	20 mA
Viset	5 V	20 mA
Vvdc	6 V	800 mA

Table 3.6: Specifications for the SC-OLink channels

channels floating, which is a requirement for all power supply channels for the pixel detector.

The low current services for two optoboards are grouped into one Type4 cable. They are separated in the opto PP3 box and instead grouped together with the NTC line for the corresponding optoboard. The low current services are then routed directly to the optoboard.

As mentioned above Vvdc is regulated at PP2. For redundancy, two channels on a regulator board (channel 0 and 15) are used for Vvdc which are set by a single trimmer.

3.6   The grounding and shielding scheme

Grounding and shielding is a key issue for the noise performance of any electrical system. Any system is subject not only to the noise sources inherent in the circuitry, but also to influences by other nearby electronics systems. Apart from the susceptibility to external noise sources, the noise emission of the system has to be minimized. [41]
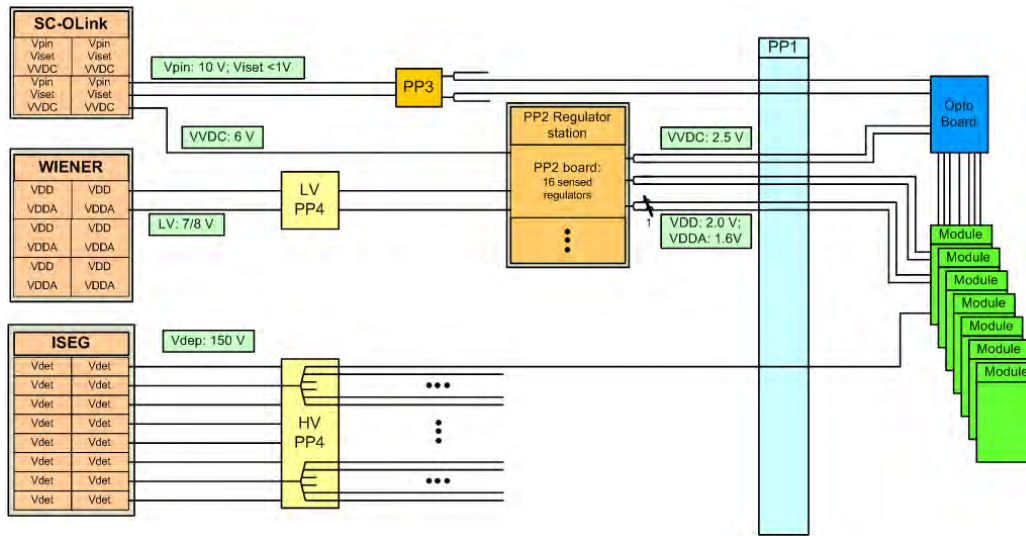
Figure 3.22: Schematic view of the power supply system as described in the text above.

3.6.1 General issues

Typical external noise sources are nearby power lines for other electrical systems, switching electronics like electric motors or pumps, electromagnetic fields from radio or TV broadcasts, as well as natural sources like lightning or solar activity. There are numerous ways to couple noise into electric circuitry. They are referred to as *noise paths* and need not be as simple as wire connections in the wrong places. Methods of noise coupling, leading to different kinds of noise paths, are the following:

- **Conductively coupled noise.**
  One of the most obvious ways to couple noise into a circuit is on a conductor. A wire run through a noisy environment may pick up noise, and then conduct it to another circuit. There it causes interference[18]. The solution is to prevent the wire from picking up noise (*shielding*), or to remove the noise from it, by decoupling before it interferes with the susceptible circuit.

- **Coupling through Common Impedance.**
  Common impedance coupling occuts when currents from two different circuits flow through a common impedance. The voltage drop across the impedance seen by each circuit is influenced by the other. The

---

[18]The undesirable effect of noise in a circuit.

classic examples of this type of coupling is a common ground connection for multiple circuits, or multiple circuits supplied in parallel by one power supply. Here either the cable or the internal resistance of the power suply form the common impedance. Therefore some improvement can be obtained by splitting the cables close to the power supply, thus decreasing the magnitude of the common impedance.

- **Electric and Magnetic Fields.**
  All circuit elements including conductors radiate electromagnetic fields whenever charge is moved. The same conductors also work as antennas, picking up fields radiated by other, near-by circuitry. Electromagnetic fields can couple into conductors in different ways. A time-varying electric field between two conductors can, according to network theory, be represented as a capacitor connecting the two conductors. It is said to couple capacitively to the conductors. A time-varying magnetic field can be represented by a mutual inductance, and is therefore said to couple inductively to the capacitors.

The two primary ways to minimize unwanted noise pickup are *shielding* and *grounding*. The two techniques are closely related, as any shield usually needs to be grounded to effectively suppress electric fields.
When properly used, shields can reduce the amount of noise coupling considerably. They can be placed around components or complete assemblies, as well as cables and transmission lines. The noise voltage, capacitively coupled to a conductor, is determined by the stray capacitance between the conductor acting as noise source, and the receiving conductor. Is the receiver shielded by a conductive shield, the capacitive coupling takes place between the noise source and the shield. If the shield is properly grounded, the noise voltage, coupled into the receiving conductor, can be reduced to zero. In the case of inductive coupling, a shield can reduce the noise voltage coupled into the receiving conductor, significantly if the ohmic resistance of the shield is high. As the voltage coupled into the receiver depends on the area of the magnetic field, the receiver encloses, it is best to keep this area small. Therefore twisted pair cables are inherently less susceptible against inductively coupled noise voltage.

There are two basic objectives involved in designing a good grounding system. The first is to minimize the noise voltage generated by currents from multiple circuits flowing through a common ground impedance. The second is to avoid ground loops which are susceptible to magnetic fields and differences in the ground potentials. In a general sense a ground is defined as a potential which serves as a reference voltage for a circuit or system. It is not

necessarily earth potential, and furthermore two physically separated ground points are seldom at the same potential. Signal grounds can be implemented as single-point grounds, with the circuits either connected in series or in parallel to the ground point. Or they can be implemented as multi-point grounds. Independently of which configuration is used, it is necessary to disconnect signal grounds from high-current and chassis grounds, as these tend to not be very stable. This could induce noise voltages into the circuits through the ground connection.

3.6.2 Grounding and shielding in the pixel detector

Controlling the noise level is very important in the pixel detector, as an increase in noise occupancy has a direct impact on the tracking and vertexing performance of the pixel detector. Therefore the grounding and shielding scheme has been carefully designed to avoid issues like ground loops, parasitic current paths and multiple uncontrolled connections of the pixel detector common ground to ATLAS ground.

All Inner Detector subsystems are electrically independent and use floating power supplies. They are connected to a single ground connection (*IDGND*). Each subsystem is connected to IDGND by a single point connection from each of its EMI[19] shields. Thus the pixel detector has one IDGND connection on the A side. All internal grounds in the subsystem are referenced to the relevant EMI shield. The pixel shield is formed by the outer shield of the beam pipe (50 μm Al foil), the PP1 endplates and the shield surrounding the Pixel detector Support Tube (PST, again 50 μm Al foil).

Electrical grounding hierarchy

The first reference point is the connection between the grounds for the digital and analog supply voltages (*DGND* and *AGND*) on the flex hybrid on the pixel module. It is made via a 0Ω connection between ground traces on the Flex. The return potential for the sensor bias voltage is connected to VDDA via a 10kΩ resistor.

On the next level all digital module grounds are connected via 0Ω to the corresponding Vvdc ground to provide a good reference for the LVDS signals between MCC and DORIC/VDC.

The PP0 ground potential is then connected to the overall pixel ground at PP1. As all power supply channels, including the regulator channels at PP2, are floating, this scheme of connections forms a parallel, single-point ground connection.

---

[19]Electro-Magnetic Interference

Grounding of mechanical structures

Most of the support structures are made of carbon fiber and carbon-carbon composite materials, which are highly conductive. Although all structures are coated with a non-conductive parylene material, they need to be held at a well-defined potential ('grounded'). This is especially true for any structures that are strongly capacitively coupled to the pixel modules, like the local support structures (see chapter 3.9).

The grounding scheme for mechanical structures was carefully designed not to provide any low-impedance current paths through the detector as this could shunt noise currents from the EMI shield into the sensitive pixel detector volume.

The local support structures are build from carbon-carbon material incorporating the tubes of the cooling system. As mentioned above the carbon-carbon structures are parylene coated, as are the cooling pipes. As the coating is not perfect, the ohmic resistance between the sector plates or the stave TMTs[20] and the corresponding cooling tube is on the order of a few $M\Omega$.

Sectors and staves are insulated from each other and from the disk support rings or the barrel half-shells by means of PEEK[21] mounting hardware. Sectors of one disk and staves of one half-shell are daisy-chained together and connected to the disk/half-shell at a single point. Each daisy-chain is connected to the common ground at PP1 individually. While disks are insulated from the global support structure, the half-shells are mounted using conductive hardware. As the global support structure is uniformly conductive and there are many connections between half-shells and global support the whole assembly is at the same potential. The global support structure is grounded using several connections to the same PP1 endplate. Global support is isolated from the Pixel Support Tube (*PST*). The service panels are made up of three electrically separate parts. While the PP0 panels are grounded to PP1 by individual wires, the corrugated carbon-carbon panels and the aluminum backbone tubes are grounded directly at PP1.

A parallel cooling circuit consists of two sectors or two staves. The inlet and exhaust tubes (see chapter 3.10) are electrically connected to the tubes in the staves/sectors. Via the U-Link (see chapter 3.9) the two staves/sectors are electrically connected to each other. As both ends of the cooling loop are grounded at PP1 it forms a conductive loop. These loops could induce noise currents in the cooling tubes, but as the area of the loop is small and the

---

[20]Thermal Management Tile, see chapter 3.9

[21]A very heat and radiation resistant polymer material, electrically insulating, with low gas emission and high wear and abrasion resistance.

loop is contained inside the EMI shield this effect is negligible. The cooling tubes are fed through the PP1 endplate using non-conductive bellows. Thus the ground connection is made through a wire connecting the detector side of the bellows with the PP1 endplate. Outside PP1 the cooling tubes are grounded together in the cooling plant.

Shielding

Inside the EMI shield no shielding is employed. The use of twisted pair cables and LVDS standard signals reduces crosstalk to a minimum.
Cables outside of PP1 have their shield DC-coupled to the internal ground on the respective detector side (meaning at PP1 for Type2 cables) and AC-coupled to the chassis ground on the counting-room side (the PP2 boxes for Type2 cables).

3.7 The interlock system

The interlock system [42] for the pixel detector is designed to assure the safety of the detector as well as the operators. It is a purely hardware based system operating independent of the detector control system. The main risks handled by the interlock system are:

- Over-heating of modules, optoboards or voltage regulators.

- Danger to the entire detector due to failure of the cooling system, radiation, humidity or fire.

- Danger to the operator caused by laser radiation.

The interlock system is designed using an active-low logic, meaning that a voltage of 0V on an interlock input trips the interlock of the respective device. It is also required that, in case of a power cut, the power supply channels stay switched off until the operator has ensured that all settings are correct. A schematic overview of the interlock logic can be found in figure 3.23.

3.7.1 Components of the interlock system

The interlock system consists of various components. The temperature information is converted into digital signals in the BBIM, which are analyzed and acted upon by the Interlock-Matrix mounted in the Logic Unit. The generated interlock signals are distributed to the corresponding power-supplies by the Interlock Distribution Boxes.
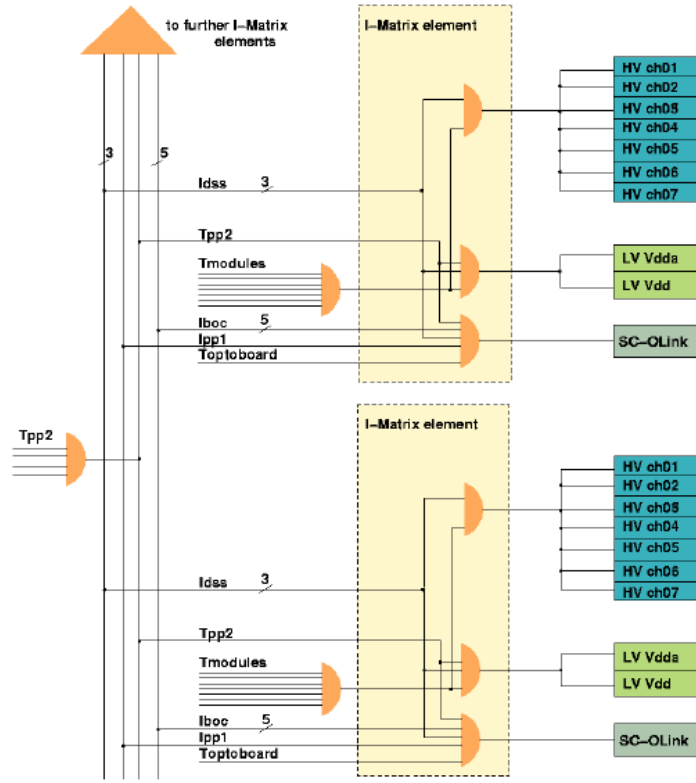
Figure 3.23: Schematic view of the interlock logic for two I-Boxes. [42]

- **The NTC**
  All temperature sensitive devices are equipped with high precision
  10 kΩ Negative Temperature Coefficient (NTC) resistors. The resis-
  tance of a NTC is a non-linear function of temperature which to a first
  approximation can be written as:

$$R(T) = R_{25} \exp\left(B \left(\frac{1}{T} - \frac{1}{T_{25}} \right) \right) \tag{3.3}$$

  R(T) Resistance at temperature T [K]
  R_{25} Resistance at T = 25°C
  T_{25} 298.15 K
  B   3435 K

  The B-value is typical for the type of NTC resistor chosen. It is calcu-
  lated from the resistances at 25°C and 85°C. Thus the formula above
  is exact only for these two temperatures. The maximum error of the

temperature measurement is ± 0.5 K at -10°C.

To determine the temperature T [°C] for a measured resistance R the Steinhart-Hart equation is applied:

$$\frac{1}{T} = A_0 + A_1 \cdot \ln(R/\Omega) + A_3 \cdot (\ln(R/\Omega))^3 \qquad (3.4)$$

A_0 [1/K]	0.66103×10^{-3}	0.7825×10^{-3}
A_1 [1/K]	0.28957×10^{-3}	0.2722×10^{-3}
A_3 [1/K]	0.30692×10^{-7}	0.8659×10^{-7}

Table 3.7: Coefficients for the Steinhart-Hart equation (determined for the purchased batch of NTC resistors)

For more details on the spread of the NTC values and long-term stability see reference [43].

- **The BBIM**
  As mentioned above the BBIM consists of an ELMB and a specialized printed circuit board, the Interlock-Box or I-Box. The ELMB monitors a reference voltage that the I-Box applies across the connected NTC resistors, and measures the actual voltage on the NTC. Thus the resistance can be calculated and, using equation 3.3, converted into the actual temperature.

  The NTC lines are routed to the I-Box in parallel to the ELMB. The I-Box uses a discriminator to generate interlock signals from the NTC voltage. The thresholds for the discriminator can be set by a small PCB referred to as the *shooting-point board*. If the NTC voltage corresponds to a temperature below a lower threshold, the T_{low} signal is generated. If it exceeds an upper threshold the T_{high} signal is generated. These signals are sent to the Interlock Matrix to be acted upon. Using combinations of these two signals a disconnected NTC line can be found which also trips the interlock.

- **The Interlock Matrix**
  The Interlock Matrix (*I-Matrix*) is a specialized PCB, twelve of which are mounted in a 3U crate called *Logic Unit*. Interlock Matrix inputs are equipped with 10 kΩ pull-down resistors. A disconnected cable thus trips the interlock. The Interlock Matrix handles the module temperature signals for one sector/half-stave as well as the temperature signals of the associated optoboard and PP2 regulator boards. It also

handles the local and remote laser interlock signals and the Detector
Safety Signals (*DSS*).

From these input signals the I-Matrix creates interlock signals for the
corresponding power supplies.

There is a special distribution box for the remote laser interlock com-
ing from PP1, the Ipp1-Box. It is a passive unit splitting the incoming
Type4 interlock cables and routing the signals to the respective Logic
Unit. A special kind of interlock matrix handles BOC laser interlocks,
the BOC-I-Box. It generates the interlock signals for the laser interlock
and distributes them via the Ipp1-Boxes to the logic units.

The routing of the interlock signals, generated due to an over-temperature
on a given module, optoboard or PP2, as well as the laser-interlock from
a given BOC, is done in the FPGA inside the Interlock Matrix or BOC-
I-Box. The program is the same for all Interlock Matrices. The right
programming for a given kind of input signals (Tmodule, Topto etc.)
is chosen via dip switches on the I-Matrix boards.

The I-Matrices are also equipped with latches that store an active in-
terlock status for all channels. These are just for readout by DCS.
They do not have any influence on the present interlock status. This is
a nice feature for debugging, as not all power supplies store an active
interlock status.

The I-Matrix has some test functionality and it provides the possibility
to exclude single channels from the calculation of the interlock status.
This is necessary to account for the number of 6 or 7 modules on a
half-stave connected to one I-Matrix.

- **The Interlock Distribution Boxes**

  The interlock signals from the I-Matrices are transferred to the Inter-
  lock Distribution Boxes (*IDBs*) by cables in the DCS racks in the US15
  and USA15 counting rooms. These cables are called Type5, while the
  cables connecting the IDBs to the corresponding power supplies are
  called Type6 cables.

  There are three kinds of IDBs: IDB-LV, IDB-SCOL and IDB-HV that
  passively distribute the interlock signals from the Logic Units to the
  WIENER low voltage power supplies, the SC-OLinks and the ISEG
  high voltage power supplies, respectively. As each IDB can handle 84
  input signals, a total of four boxes of each type are used for the pixel
  detector. The routing of the interlock signals is again done via a FPGA.

3.7.2 Interlock conditions

- **Temperature interlock.**
  The temperatures for which a temperature interlock is issued (*shooting points*) depend on the monitored device.
  For the pixel modules the shooting point (Tmodule) is set to 40°C. This is a safe temperature for continuous operation of the modules and allows to do tests with the modules if no cooling is available.
  For the optoboards the shooting point (Topto) is also set to 40°C. The optoboard was found to work best at temperatures between 20°C and 40°C.
  The shooting point (Tpp2) for the PP2 temperature is set to 60°C.

- **Laser interlock.**
  When the cover of the BOC rack is opened a so-called local interlock (IBOC) is issued that switches off power to the VCSELs on the TX plugins, as well as the SC-OLink for the optoboards connected to the respective BOC.
  The fiber connections at PP1 are monitored by a micro-switch that, when opened, issues a remote interlock (Ipp1) . This switches off all BOCs and all SC-OLinks.

- **Detector Safety Signal.**
  The Detector Safety Signal (*DSS*) is issued by ATLAS in any of the following cases:

  - Failure of the evaporative cooling main or backup system.
  - Failure of any environmental gas system.
  - Failure of the monophase cooling system used for cooling cables.
  - Radiation levels detected by the Beam Conditions Monitor implying beam misplacement or imminent beam loss.
  - Smoke in the PP2 region.

  The DSS interlock (IDSS) will switch off all power supplies.

Table 3.8 shows the actions taken by one I-Matrix for the different interlock conditions.

3.8 The Detector Control System

The Detector Control System (*DCS*) is a software project that controls the power supplies, and monitors and displays parameters like module/optoboard

	HV	LV	SC-OLink	BOC
Tmodule	X	X		
Tpp2	X	X	X	
Topto			X	
IBOC			X	X
Ipp1			all	all
IDSS	all	all	all	

Table 3.8: Actions taken by one I-Matrix for different interlock conditions. Here 'X' means switch off only affected devices, 'all' means switch off all devices in the detector.

temperatures, voltages and currents, environmental sensors and interlock conditions. It is implemented in the SCADA[22] framework PVSSII that is used throughout ATLAS to make integration of the sub-detector DCS projects into the common ATLAS DCS project feasible. PVSSII provides all necessary functionality to supervise and control the numerous systems needed to operate the pixel detector.

Communication to hardware happens mostly through CAN bus[23]. Only the WIENER low voltage power supplies are connected via TCP/IP. On the software side a server/client system is implemented, based on the OPC[24] interface.

The DCS system consists of various sub-systems that provide different functionality.

- **Front-End Integration Tool (*FIT*)**
  To integrate the hardware devices (power supplies, ELMBs, etc.) into the DCS project the Front-End Integration Tool is used. Once a kind of hardware device is defined (e.g. a PP2 regulator board) in the FIT, it provides scripts that automatically integrate all instances of this device needed in the pixel detector. All required parameters like CAN bus node ID, set values for voltages, voltage and current limits, can be set in the FIT. It also provides graphical interfaces to the hardware components to control them and to display monitored values.

- **System Integration Tool (*SIT*)**
  Although every single channel of every device is accessible via the FIT

---

[22]Supervisory Control And Data Acquisition

[23]CAN: Controller Area Network, a serial bus system

[24]Openness, Productivity, Collaboration; former OLE for Process Control, a standardized software interface.
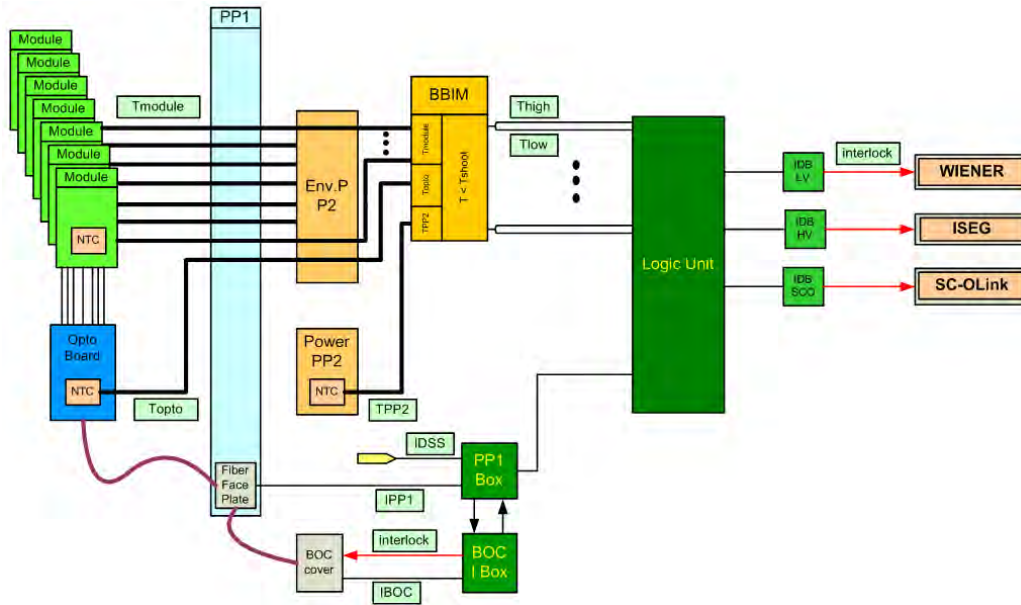
Figure 3.24: Schematic view of the interlock system as described in the text above.

panels, this functional approach is not very user friendly. The user has to know the complete connectivity of the system to find the right values in the FIT panels. Therefore the System Integration Tool (SIT) is used to map hardware channels to representations of the detector components. Thus all relevant data can be displayed in one panel representing a module or optoboard. Panels are grouped to resemble detector parts like half-stave or sectors.

- **Finite State Machine (FSM)**
  To monitor the state of the complete pixel detector a significant reduction in the information provided to the user is necessary. Therefore the Finite State Machine (FSM) calculates quantities called 'STATE' and 'STATUS' from the available information. These quantities are then used to display the condition of the pixel detector. 'STATE' and 'STATUS' are calculated using module temperature, currents on the VDD and VDDA supply and the interlock status of the power supplies.
  The FSM handles the system on different hierarchy levels. The user can monitor and control half-staves and sectors as well as disks and barrel layers or the complete detector system. Following this hierarchy it provides an easy way to switch on or off the complete detector or only parts of it. The FSM uses the mapping of module/optoboard to

power supply and temperature sensor that is set in the SIT.

3.9   Mechanics

The support structures used in the pixel detector package are designed to provide maximum strength using low mass materials. Carbon fiber compound materials are used wherever possible, as they provide very high thermal conductivity, mechanical stability and low interaction cross-section for traversing particles. As they are also very porous and highly conductive, all carbon fiber structures are coated with parylene.

- **Stave/Sector**
  The modules are mounted on local support structures that also provide cooling. For the disks these are called *sectors*. They are made of two carbon-carbon faceplates glued to a carbon foam core containing the aluminum cooling tube. Three disk modules are mounted on each of the faceplates such that the modules on the two faceplates overlap at the inner radius of the disk and cover the disk completely in ϕ direction at the outer radius. This is shown schematically in figure 3.25.

Figure 3.25: Schematic view of modules mounted on a sector. Light green are the modules on the front side, dark green the modules on the back side. Module naming and overlapping regions between modules are indicated.

  The local support structure for the barrel layers is called a *stave*. A stave is made of two pieces of carbon-carbon material glued together in the middle of the structure. Together they form the Thermal Management Tile (*TMT*) which is about 1.3m long. 13 pixel modules are glued to the TMT with a small tilting of 1.1° in the direction of the stave axis. This allows for the modules to have a small overlap. This

is necessary to achieve a full coverage of the active sensor area. On the backside of the TMT the aluminum cooling pipe is held in place by another carbon compound structure, called the *omega profile* because of its shape. The cooling pipe is shaped to have maximum contact surface with the TMT. The omega profile is glued to the TMT and filled with thermal grease to improve heat transport to the cooling pipe.


Figure 3.26: Top: Drawing of a stave with 13 modules. Indicated are dimensions and positions in z direction of the modules as well as the naming described in the text. Bottom left: cross-sectional view of a stave [44]. Bottom right: drawing of two staves assembled into a BiStave.

The staves are tilted along the stave-axis by 20° to increase the incident angle of particles on the sensor. This increases the probability of charge sharing between pixels, which is used to increase spatial resolution. Sectors and staves also house two geometrical reference markers, called *ruby balls*, which allow to survey the position of the respective structure in the global coordinate system.

The modules are identified by their position on the support structure, the *assembly position* (see figures 3.25 and 3.26). On a sector the modules on the faceplate facing the interaction point are named M1 to M3 and on the opposite faceplate M4 to M6. On a stave the modules are named M6A to M1A on the A side, with M6A being furthest away from the interaction point. The middle module is named M0 and the modules on the C side are named M1C to M6C with, again, M6C furthest away from the interaction point.

Every module is connected to the corresponding optoboard via a Type0

cable. The cables are grouped into bundles of 6 or 7 cables. The modules whose cables are grouped into a bundle and which are connected to the same optoboard are referred to as a *half-stave*.

- **Parallel Cooling Circuit**
  A Parallel Cooling Circuit (*PCC*) is formed from two local support structures with their cooling tubes connected. The connection is made using a short bent piece of aluminum tube, referred to as the *U-Link* because of its shape. The PCCs in the disks are called *BiSectors* while the PCCs in the barrels are called *BiStaves*. They form the basic unit of the evaporative cooling system (see chapter 3.10).
  The two staves making up a BiStave are fixed to small PEEK support structures on both ends and in the middle of the BiStave. They are tilted with respect to a perpendicular on the beam pipe surface. This not only provides an incident angle for the traversing particles, but also allows for an overlap between the modules on both staves. This again is necessary to achieve full coverage of the solid angle. The two staves of a BiStave are referred to as Stave1 and Stave2. Type0 cable bundles with seven cables are routed towards side A for one stave and side C for the other stave, to balance the material budget. The inlet and the exhaust of the cooling circuit is always on the same side of the BiStave.

- **Disk/Half-shell**
  Four BiSectors are mounted on a disk support ring to form a *disk* assembly.
  The BiStaves are loaded into *half-shells*, which are later on clamped together to form a barrel *layer*. The number of BiStaves for the B-Layer is 11, for Layer 1 it is 19 and for Layer 2 it is 26. The BiStaves are arranged such that the cooling pipes are connected on alternating sides of the detector. The only exception is the B-Layer, where all cooling tubes are routed to the C side.

- **Endcap/Barrel**
  The half-shells are clamped around the central part of the beryllium beam pipe to form a barrel layer. As the B-Layer radius is smaller than the endflange of the beam pipe, it has to be clamped around the beam pipe during integration. The radii of the layers are summarized in table 3.9. The outermost layer (Layer 2) is surrounded by the barrel frame.
  Three disks are integrated into a disk frame, forming an endcap. The barrel and the endcap frame are screwed together to form the mechanical contact between barrel and endcap section. They are then referred

to as *global support*. Figure 3.1 shows a schematic view of the barrel and disk layers integrated in the global support structure.

| Layer	Radius [mm]	No. of staves
B-Layer	50.5	22
Layer 1	88.5	38
Layer 2	122.5	52

Table 3.9: Radii of the barrel layers

Figure 3.27: Drawing of staves mounted in the barrel layers, viewed from side A.

- **Service Panels**
  A Service Panel, also known as Quarter Service Panel (*QSP*), consists of two PP0 regions and two corrugated carbon-carbon panels on which the electrical and cooling services are routed from PP1 to the detector. The two parts of the service panel are mounted to an aluminum backbone-tube. There are two service panels to a complete QSP, an Inner Service Panel (*ISP*) and an Outer Service Panel (*OSP*). As the names suggest, the radius of the ISP is smaller than that of the OSP. The QSP is supported by the BPSS[25] longerons (see section below).
  The kapton printed circuit boards, which form PP0, are mounted in 2x6 rows. On the OSPs two PP0 PCBs are mounted back-to-back per

---

[25]BeamPipe Support Structure

row, while on the ISPs only one PP0 PCB per row is mounted. Thus an OSP can serve 24 half-staves, an ISP only 12.

The electrical services are routed through the corrugations while the optical fiber ribbons are placed along the aluminum backbone. The cooling tubes are situated in between the inner and the outer service panel with the inlet and exhaust tubes glued together to form the heat exchanger (see chapter 3.10).

The optoboards mounted on the PP0 rows are equipped with copper covers to prevent convective cooling by the cold gas in the detector volume. Furthermore heaters are glued to the optoboards. With the heaters and the separate opto cooling circuit it is possible to adjust the optoboard temperature to the optimum value.

- **Beam Pipe Support Structure**
  The Beam Pipe Support Structure (*BPSS*) is a cruciform assembly connected to carbon-carbon tubes, called *longerons*. It is located close to the detector on each side and is connected to an aluminum ring around the beam pipe by means of four wires. These are routed up to the PP1 region. Using these wires, the beam pipe is supported and adjusted to be coaxial with the pixel package once the service panels are installed. An additional pair of support rings is installed at the ID endplate. The longerons also support the service panels.

- **Pixel Support Tube**
  The Pixel Support Tube (*PST*) consists of a 7m long cylinder extending coaxially all the way through the Inner Detector. This allows independent installation and removal of the pixel detector while the rest of the Inner Detector is in place. The PST consists of a barrel and two forward sections connected through mating flanges. It has rails to allow the pixel package to slide into position inside the Inner Detector. The PP1 faceplates act as endplugs for the PST. When closed, the PST provides a thermal and environmental barrier and is part of the EMI shield.
  The barrel section is supported by the SCT barrel while the forward sections are supported by the barrel cryostat at the end of the Inner Detector volume. All supports are electrically isolating. The PST is fixed in all directions at the barrel support on side C while on side A it is free to move in z-direction to account for thermal contraction coefficient mismatch between the PST and the SCT barrel. The mounts allow motion in z-direction of up to 1.5 mm.
  After the pixel detector is slid into the PST on the rails, it is transferred

to the pixel detector supports. These are integrated into the PST barrel flanges. The supports are electrically isolated from the PST.

Barrel and forward sections are equipped with heaters on the outside. These are necessary to prevent condensation on the gas barrier in case the gas in the gap between PST and SCT is not dry.

To form the EMI shield the PST is covered by 50 μm thick aluminum foil that is electrically connected to the PP1 faceplate. This, in turn, is connected to the 50 μm aluminum foil around the beam pipe to close the faraday cage.

3.10 The cooling system

As mentioned in chapter 3.2.1 the sensors of the pixel modules will be operated at -7°C to prevent reverse annealing and thermal runaway of the radiation damaged sensors. Therefore the heat dissipated by modules and cables has to be removed from the pixel volume in a very efficient way introducing as little extra mass as possible into the pixel volume. Table 3.10 shows a calculation of the worst case power dissipation inside the pixel support tube.

The cooling system for the ATLAS Inner Detector [45, 46] is a bi-phase

Summary of heat loads in the pixel volume [W]		
Part	Power per cooling circuit	power total
B-layer	171	1883
Layer 1	171	3252
Layer 2	171	4450
Total barrel		9586
Disks	79	316
Total disks		1896
Type0 barrel		2174
Type0 disks		364
Opto links & Vvdc cables		577
Type I inside PST		1604
PST heaters		2509
Total inside pixel volume		18710

Table 3.10: Worst case power dissipation in the pixel volume

evaporative fluorocarbon system. An evaporative cooling system has been chosen mainly because of the high heat transfer coefficient between the cooling fluid and the device to be cooled, the small temperature gradients along

long cooling tubes and the smaller size required for the cooling channels. The
smaller size, and thus the lower coolant mass flow, is due to the larger cooling
capacity per unit volume as compared to a mono-phase system. This is a
result of utilizing the latent heat of vaporization rather than a liquids specific
heat capacity. For the coolant, C_3F_8 was chosen because it is non-flammable,
non-conductive, radiation resistant and has a small vapor specific volume al-
lowing minimal tube sizes. From an environmental point of view C_3F_8 has
zero ozone depletion potential. However, it has a high global warming poten-
tial and long atmospheric lifetime. Leak tightness and minimum liquid loss
rate are major requirements for the cooling system, as C_3F_8 is also rather
expensive.

The cooling system can be divided into internal and external parts. The
internal part includes all components inside the ATLAS detector: recuper-
ative heat exchanger, capillaries and the detector structures. The external
part includes the cooling plant and the distribution rack. Figure 3.28 shows
a schematic view of the cooling system.

3.10.1   The external part

The cooling plant, located in USA15, consists of a buffer tank, four com-
pressors, a condenser and a 800 liter storage tank. The buffer tank receives
the coolant in vapor phase. It is held below atmospheric pressure by the
compressors. These compress the low pressure vapor that is at room tem-
perature. The compressed hot vapor enters the condenser that is cooled by
an external chilled water supply. The liquid coolant is stored in a tank that
is cooled by the chilled water. The pressure in the liquid tank is around 17
bar(a)[26], while the system is operating, so that it is above the saturation
point of C_3F_8 for a worst case fluid temperature of 40° C.

The cooling plant is connected to the distribution racks through a main pneu-
matic valve which is operated by the control system. In case of an alarm
condition in the cooling plant, this valve is closed, cutting off the liquid flow
on the inlet side of the distribution racks.

The distribution racks, located on the service platforms inside the experi-
mental cavern UX15, are manifolds connecting the internal part, including
the detector, to the main inlet and exhaust lines. On the inlet side they can
be disconnected from the main inlet line through a pneumatic shut-off valve.
Each of the four racks has inlets connected to pixel cooling circuits as well
as SCT and thermal barrier circuits. Therefore each inlet line is equipped
with a pressure regulator (PR) which also allows to reduce the inlet pres-

---

[26]absolute pressure in bar

sure. This is necessary, as the SCT cooling circuits operate at a lower input pressure than the pixel circuits. The pressure at each exhaust line in a distribution rack (*back-pressure*) is set and regulated via back-pressure regulators (*BPRs*).

3.10.2 The internal part

Each cooling circuit consists of one recuperative heat exchanger, one capillary, the detector structure (one BiStave or BiSector) and one heater on the return vapor line. Each cooling circuit has its own pressure and back-pressure regulator to set the evaporation conditions.

The circuits are operated in fixed-flow mode. The mass flow rate required to cool a given detector structure is determined by the power dissipation in that structure, the efficiency of the heat exchanger and the vapor quality[27] at the exhaust of the circuit. Thus the mass flow rate is specific for each kind of cooling circuit and is tuned by changing the geometry of the capillary (length and inner diameter). However, although each circuit is operated in fixed-flow mode, the flow rate can be adjusted within a small range by changing the pressure regulator setting.

To avoid condensation on the return tubes, the remaining fluid has to be boiled off and the vapor heated up above the cavern dew point. This is achieved through a heater installed on the return tube that is controlled so that the heat provided to the cooling circuit is constant even if the heat dissipation of the detector structures changes.

The recuperative heat exchanger decreases the vapor quality in the inlet line which in turn decreases the required mass flow through the detector structures. It is implemented by simply bringing the inlet and exhaust tubes in good thermal contact throughout the length of the service quarter panel.

3.10.3 Functionality and Control

Figure 3.29 illustrates the operation of the system with the use of the coolants pressure-enthalpy diagram. The parts of the cooling system discussed here, are pointed out in figure 3.28.

In the cooling plant the superheated coolant vapor returns from the detectors (point A), is compressed by an oil-free compressor (A$'$ - B) and condensed in a water cooled condenser (B - C). The liquid C_3F_8 from the condenser is pressurized to 16 bar(a), (C - D) and transferred along 120 m of pipe work

---

[27]Vapor quality describes the ratio of vapor and liquid mass in a mixture: $x_v = \frac{m_v}{m_l+m_v}$ with m_v being the vapor mass and m_l the liquid mass in the mixture.
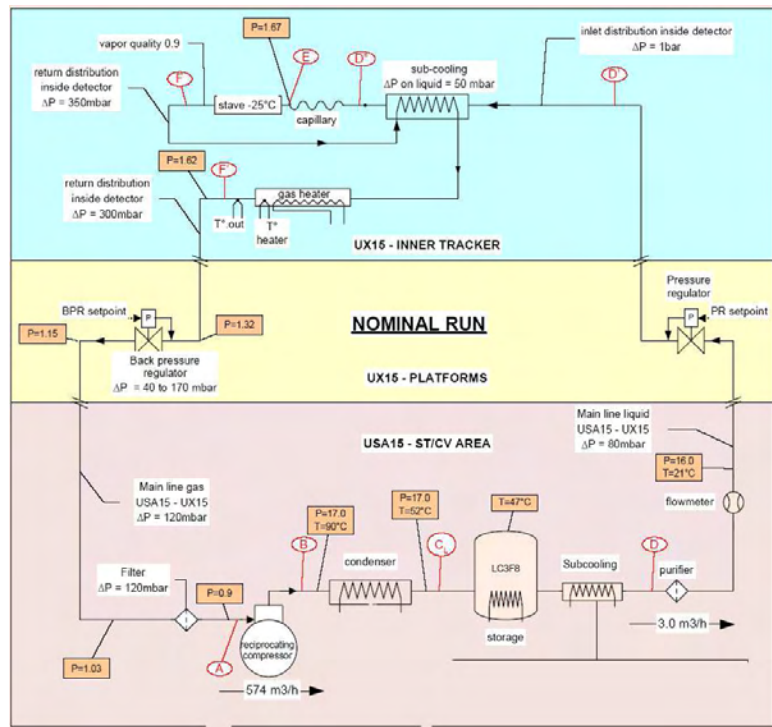
Figure 3.28: Schematic view of the evaporative cooling system.

to the distribution racks on the platforms in the experimental cavern. The
high pressure of the coolant liquid is dictated by the requirement to keep
the coolant above its saturation point along all the pipe work between the
condenser and heat exchangers (C - D$'$). The maximal liquid temperature ex-
pected in the inlet tubes is 40°C. At this temperature the saturation pressure
is 12.8 bar(a). Therefore 13 bar(a) has been set as the minimum pressure
on the liquid side of the system. The liquid flows through a recuperative
heat exchanger and is cooled down by the boiling liquid/cold vapor mixture
returning from cooling channels of the detector (D$'$ - D$''$). The sub-cooled
liquid expands in the capillaries just before the cooling channels (E) and
then remains in saturation conditions (that is boiling), in the detector cool-
ing channels. The evaporation temperature is determined by the saturation
pressure set by the back pressure regulator at the manifold on the external
platforms. A pressure of 1.67 bar(a) is required for a coolant temperature of
-24°C. The C_3F_8 mass flow rate is set to obtain a maximum vapor quality
of 0.9 on the outlet of the cooling channels (F) at the maximum heat load
produced by the detectors. The residual liquid in the mixture at the ex-
haust of the cooling channels performs the sub-cooling of the inlet liquid in

the recuperative heat exchanger. Any remaining liquid is evaporated by the heater, which also raises the vapor temperature above the cavern dew point. Superheated vapor flows across the back pressure regulator on the external platform to the cooling plant situated at a free access area in the ATLAS experimental site.

Figure 3.29: Pressure - enthalpy diagram of C_3F_8 showing the principle of operation of the evaporative cooling system.

The control system is based on CERN recommended Programmable Logic Controllers. It incorporates the following components in a well-proven model architecture:

- The Programmable Logic Controllers (PLCs)

- The Supervisory Control and Data Acquisition (SCADA) system

- The Archiving Computer

- The Experiment SCADA system

- The Technical Data Server (TDS)

The PLCs, the general SCADA system and the Archiving PC are supplied and maintained by CERNs Technical Support - Cooling and Ventilation group (*TS/CV*).

Chapter 4

The system test

> *Be of good cheer. If science teaches us anything, it teaches us to accept our failures, as well as our successes, with quiet dignity and grace.*
> Gene Wilder, 'Young Frankenstein', 1974.

The previous chapter may have left the reader under the impression that the pixel detector is a very complex system. This is why, although every single component undergoes thorough testing during its production, the interplay between the components needs to be tested in a large scale system operating under realistic conditions. This so called *System Test* was conducted at CERN. Starting in spring 2006, a small system was used, including a single BiStave and a few prototype services components, until the final versions became available. In september 2006 the BiStave was replaced with one of the two endcaps of the pixel detector (endcap A), which was delivered to CERN after its assembly and production testing was finished. To operate the endcap, a significant upgrade of the available service infrastructure was necessary. The number of detector and supply channels used in this final stage of the system test roughly corresponds to 10 % of the complete detector system.

Although smaller scale system tests were already conducted in several laboratories, these were mainly focused on the parts of the system produced by the respective institute. The purpose of the system test at CERN was rather manifold:

- Understand the difficulties in getting the system operable.

- Gain experience with the operation of the system.

- Train operators for real detector operation.

93

- Demonstrate operability and functionality of the system by measuring cosmic muons.

4.1   The system test setup

The system test setup in the SR1 building in CERN's ATLAS area consisted of endcap A of the pixel detector, a prototype of an outer service quarter panel (*PSQP*), a small scale version of the biphase cooling system and the services required to power, monitor and control these devices. The endcap and the PSQP were mounted in a large box, flushed with dry air, with the PP1 region of the PSQP outside the main dry volume. The dry air system available in the building provided air at a dew point of approximately -20°C. As the temperature of the C_3F_8 in the cooling pipes is around -25°C this dew point was marginal for cooling operation. Therefore an additional desiccant dryer was employed, which lowered the dew point of the dry air to -40°C. The dew point was monitored using an external hygroscope, measuring the humidity of the air mixture at the point farthest away from the dry air inlet lines.

The endcap was mounted with the disks parallel to the floor, disk 1 facing upwards. Cooling connections to the PSQP were made using production capillaries on the inlet side and special copper adapter pipes to account for the angle of 90° between endcap and PSQP with respect to the final mounting. The modules were connected to the PP0s following the final connection scheme (one sector or half-stave per PP0 row).

The PP1 region of the PSQP was separately insulated and flushed with dry air. This was necessary to keep it accessible with the main volume dry.

This chapter will provide information on the system test setup, where it differs from the final experiment as described in chapter 3, as well as some measurements done to characterize the setup in terms of measurement accuracy, power supply stability, cooling performance etc. Details on the system test setup, that lie beyond the scope of this chapter, can be found in appendix A.

4.1.1   Test of the service chain

Before powering the detector, the service chain is thoroughly tested using a specialized setup. Cables at PP1 are connected to a GPIB controlled setup, including a switching matrix, a programmable load generator and digital multimeters. The service test software is implemented in the PVSSII

framework and integrated with the DCS project. The setup is designed to qualify the full service chain from the power supplies in the US15 and USA15 counting rooms to PP1 and back, including interlock functions. It allows testing of one flavor of Type2 cables at a time: low voltage, high voltage and NTC/Opto-cables. The tests include basic continuity checks, cross-checks of the current measurements at PP4 via the programmable load, adjustment of the voltage regulators at PP2 and a test of the temperature interlock function. The tests are done for every single channel of the service chain, using the mapping provided by the DCS SIT project (see chapter 3.8). If the channels pass the service tests it is considered safe to connect the detector. For more details on the service test setup see [47].

4.1.2 Power supply and interlock system

Power was supplied using production power supplies and patch-panels (see chapter 3.5) except for a missing HV-PP4, of which no prototype existed at the time of the system test. A modularity of one channel per module was used for the high voltage supply instead. The power supplies as well as the patch-panels 4 and 3 were situated in a rack area outside the clean room. The environmental, NTC-OPTO and power boxes at PP2 had to be set up inside the clean room because of the length of the Type2 cables connecting them to the PP1 region of the PSQP.

Great care was taken to implement the interlock system, to ensure the safety of the detector (see chapter 3.7). Temperature interlocks were implemented for the modules and the optoboards, whereas the PP2 temperature could not be read out at that time. Laser interlocks were implemented in every place where connections between optical fibers were made. Microswitches at the PP1 fiber endplate and switches on the doors of the BOC rack generated the local and remote interlocks.
Special attention was paid to protect the detector against failures in the cooling system. A hardware interlock ,triggering on the back-pressure measured in the exhaust lines at PP1, as well as a software interlock, triggering on alarm states of the PLC that controlled the cooling system, and a missing connection to the cooling SCADA system, were set up. These interlock signals were connected to the DSS signal input of the Logic Unit that handled the interlock system.

Figure 4.1: Decrease of liquid weight in the course of one week

4.2   Qualification measurements

The system test setup was used to qualify the cooling, power supply and monitoring systems for the experiment.

4.2.1   The cooling system

The input pressure measured at the inlet manifold at PP1 was around 12 bar(a) while the set value in the plant was around 13 bar(a). The pressure drop happened along the main inlet line, in the distribution rack and the smaller inlet lines going to the manifold. During normal operation the back-pressure was set to 2 bar(a), resulting in an evaporation temperature of the C_3F_8 of -20°C. The average temperature of the modules configured (i.e. dissipating about 3.5W in heat) was around -15°C to -18°C.
 As the system test was also used to test the performance of the cooling system itself, one of the most closely monitored quantities was the coolant loss rate. It was calculated from the liquid weight in the storage tank a certain time after shut-down of the plant, as well as from the liquid weight during stable cooling operation with no other parameters changing and constant heat dissipation from the detector. After the last intervention with the cooling system, a liquid loss rate of (0.08 ± 0.005) kg/h was measured over a period of 14 days.

Figure 4.2: Module temperatures averaged over all modules (left) and over each disk (right) versus power dissipation.

Cooling performance

Only two out of the twelve BiSectors in endcap A had all modules configuring as expected during the studies on cooling performance. The rest had one or more modules that, due to problems with the optical link or the modules themselves, did not draw the expected currents on the supply voltages. Thus it is rather difficult to obtain a realistic picture of the temperatures on a BiSector.
On disks 1 and 3 the inlet of a cooling circuit is at modules 1/6 on sector 1 of a BiSector and the exhaust is at modules 3/4 on sector 2. For disk 2 the coolant circulates the other way around. Generally, the modules at the inlet (modules 1/6, sector 1) and the modules close to the U-Link between the sectors (modules 3/4, sector 1) have the highest temperatures on the BiSector. This is probably due to a locally increased pressure because of the fittings on the cooling pipes. Modules 2/5 on sector 1 and modules 1/6 and 2/5 on sector 2 are slightly cooler. The modules at the exhaust of the cooling circuit (modules 3/4, sector 2) have the lowest temperatures due to locally decreased pressure owing to the larger diameter of the exhaust pipe with respect to the cooling tube in the sector.

Dedicated measurements were carried out to study the cooling performance. Module temperature was examined as a function of dissipated power to quantify the performance of the cooling system. The power dissipated by the modules was increased by raising the supply voltages to the maximum values provided by the voltage regulators, and by setting a configuration that draws maximum current on the digital and analog supply voltages. Temperatures were measured using the readings of the module NTC resistors provided by DCS. After every change in the conditions, a settling time of 30 min was allowed.

Figure 4.3: $\left\langle \frac{dT}{dP} \right\rangle$ for all sectors

Figure 4.2 left shows the dependence of the average module temperature on the average dissipated power. The error bars represent the standard deviations in temperature and power when averaging over all modules on the endcap. Figure 4.2 right shows module temperatures averaged per disk versus power dissipation.

The curve can be approximated by a linear dependence of the form:

$$T = \frac{\overline{dT}}{dP} \times P + T_0 \tag{4.1}$$

Table 4.1 summarizes the average parameters for the disks in endcap A. Disks 2 and 3 are colder than disk 1. This is because the endcap is mounted above the PSQP such that disk 1 is the upper most disk. Thus the cold exhaust lines are routed close to disk 3. Disk 2 is the middle disk of the endcap and thus it is additionally cooled by disks 1 and 3. The values for $\frac{\overline{dT}}{dP}$ match within the measurement errors.

During production tests, the specific thermal resistivity[1] $\frac{dT}{dP}$ is measured to quantify the thermal contact of the modules to the local support structures, i.e. the sectors. Structures with individual modules exceeding the limit of 2.5 °C/W were rejected from integration, the values for $\frac{dT}{dP}$ per module are between 1.5 and 2 °C/W for most of the tested staves. Table 4.2 summarizes the $\left\langle \frac{dT}{dP} \right\rangle$ values for the three disks measured in the system test. The values for each sector are shown in figure 4.2.1. The measured values for $\left\langle \frac{dT}{dP} \right\rangle$

---

[1]i.e. the temperature increase as a function of the increase in dissipated heat, measured in K/W

correspond very well with the values expected from measurements done on single sectors during production.

disk	$\overline{\frac{dT}{dP}}$ [°C/W]	T_0 [°C]
all	1.94	-24.0
D1A	0.33	-23.4
D2A	0.35	-24.4
D3A	0.36	-24.3

Table 4.1: Average module temperature per sector versus power dissipation.

sector	$\left\langle \frac{dT}{dP} \right\rangle$ [°C/W]
D1A	1.99 ± 0.32
D2A	2.31 ± 0.16
D3A	2.28 ± 0.15

Table 4.2: Average specific thermal resistivity per module.

4.2.2 Power supply and interlock system

The power supply system in the system test was used to study the stability of the voltages and possible EMI[2] pick-up in the cables. The rack area was equipped with all the electronics that are installed in the US15 and USA15 caverns, where the power supplies are installed for the final detector. Thus the conditions in the system test in terms of electro-magnetic interference were close to the conditions in the experiment.
Measurements of voltage stability were performed using a Keithley 2001 digital multimeter. Values were recorded at a frequency of about 30 Hz over periods from 30 minutes up to 2 hours.

- **Voltage stability at power supply outputs**
  Voltage stability was measured at the output terminals of the power supplies to determine their performance in the system under different load conditions. Measurement of the output voltages with an oscilloscope showed fast variations of the output voltage below 100 μV. As the load does not change on a SC-OLink channel (current consumption on Vvdc is very stable) the measurements have been done under normal operation conditions.

---

[2]Electro-Magnetic Interference

Comparison of the externally measured values and the values archived in the DCS project showed that, for the WIENER channels the monitored values are quite correct, while for the SC-OLink the values differ by about 50 mV (see table 4.3).

supply	load current [A]	set value [V]	monitored value [V]	measured value [V]
WIENER	4.3	10	9.998	$9.995 \pm 3*10^{-4}$
	5.9	10	9.996	$9.995 \pm 2*10^{-5}$
	7.1	10	9.999	$9.990 \pm 3*10^{-5}$
SC-OLink	0.22	6	6.002	$5.953 \pm 3*10^{-4}$

Table 4.3: Output voltages of the power supplies set via DCS, monitored by DCS and measured at the output terminals of power supplies

- **Voltage stability on the modules**
  As the stability of the supply voltages has a direct influence on the noise performance of a module, the voltage was measured at the Type0 connector on the pigtail. An adapter PCB was used to connect a second Type0 cable in series with the original one, because the module pigtails were not accessible for measurements. The adapter PCB provided measurement pins for the supply voltages but not for the sense lines. Thus the voltage measured on the PCB includes the load dependent voltage drop over the second Type0 cable. To compare these measurements with the expectations from the ohmic resistance of a Type0 cable, the voltage difference is calculated between the different loads.
  It turns out, that the voltage difference for VDDA is slightly larger than expected, while for VDD the two values agree very well. This could be an effect of a bad solder joint in the adapter PCB, which has occasionally been observed.

supply voltage	load current [A]	measured voltage [V]	voltage difference [V] measured	voltage difference [V] expected
VDD	$0.313 \pm 2 \times 10^{-2}$	$2.067 \pm 8 \times 10^{-3}$		
	$0.709 \pm 5 \times 10^{-4}$	$2.314 \pm 1 \times 10^{-3}$	0.21	0.18
VDDA	$0.085 \pm 4 \times 10^{-3}$	$1.704 \pm 3 \times 10^{-4}$		
	$1.312 \pm 2 \times 10^{-3}$	$2.037 \pm 3 \times 10^{-3}$	0.33	0.42

Table 4.4: Supply voltages at the module

- **Voltage and current reading at PP2**

  The output voltage of the regulator channels at PP2, as well as the delivered current, is monitored via DCS. For the calculation of voltage and current from the ADC readings of the ELMB situated on the controller board (see chapter 3.5), a standard calibration is used for all ADC channels. This introduces a slight difference between actual and measured values. Figure 4.4 left shows, that the monitored voltage (Vmon) is systematically higher than the output voltage measured at a programmable load (Vload). The difference of up to 0.4V is less severe than the inaccuracy of the current measurement, which is too low by up to 300 mA. As current consumption gives very useful information on the status of a module, this inaccuracy renders the current measurement at PP2 all but unusable. For the time being, the more reliable current measurement at PP4 is used for monitoring.

Figure 4.4: Difference between regulator sense voltage (left) and output current (right) measured with the ADC at PP2 and with an external programmable load.

- **Current overshoot at power up.**

  When a module is powered, a peak in the digital current consumption is observed. This is probably due to increased current through inverter blocks in the readout chips that can occur during power up of the chips. Currents of up to 1.6A have been observed. Duration and amplitude of these peaks depend on module temperature. For modules at room temperature peak currents around 1A with durations of the overcurrent on the order of 1 second are measured, while for modules at -20°C the duration is on the order of 10 seconds and the amplitude can go up to 1.6A. Duration and amplitude of the peak are also influenced by the order and the timing of switching on the two supply voltages. Switching on the analog supply shortly after the digital supply (as opposed to

switching on both simultaneously) reduces amplitude and duration of
the current peak. This has been taken into account when the power-up
procedure for the finite state machine (see chapter 3.8) was defined.

Interlock timing

The timing of the interlock mechanism was measured using the service test
setup. The setup switches between two different resistors connected to the
NTC lines at PP1 to simulate temperatures below and above the shooting-
point of the temperature interlock. The voltage drop between the high side
and low side NTC lines was used to trigger an oscilloscope. The second chan-
nel of the oscilloscope was connected to the low voltage supply lines. Figure
4.5 shows the output of the oscilloscope with the NTC voltage displayed on
channel 2 and the low voltage supply line displayed on channel 3. The re-
sponse time of the interlock is measured to be 1.1s (between changing the
resistor and switching off the low voltage).

Figure 4.5: Interlock response time. Channel 2 shows the voltage across NTC
lines, channel 3 shows the voltage on low voltage supply lines.

Measurements of the response time of the interlock circuits in the power
supplies alone show response times on the order of 1 second for the WIENER
power supply. Thus the time measured is dominated by the time it takes for
the WIENER to act on an interlock signal.

Accuracy of the temperature measurement

The accuracy of the temperature measurement in the BBIM is tested during the service tests. A resistor is connected to the NTC lines at PP1 and the temperature calculated in the BBIM is logged. Figure 4.6 shows the difference between the temperature expected from the known resistance and the temperature measured at the BBIM. The peak at -0.22°C shows the difference for measurements done at a set temperature of 22°C, while the peak at 0.28°C arises from measurements with a set temperature of -11°C. These differences reflect the characteristic error of the NTC temperature measurement (see [43]). The requirement for accuracy of the temperature measurement was 1°C.

Figure 4.6: Difference of temperatures set and measured in the service tests.

4.3 Procedures from the system test

4.3.1 The power-up sequence

During powering up the detector system great care has to be taken to keep all components in a well-defined state. If, for example, the optoboard is powered without a light input signal on the PiN diode, the DORIC tries to adjust its threshold to a 50% duty cycle. When the threshold ends up in the noise floor, the DORIC will send random output signals, which result in an oscillating current consumption of the modules. Therefore a procedure was fixed in the system test that circumvented this kind of problem. The procedure described below is based on the voltage regulators at PP2 being set to the right output voltages and all channels being inhibited individually and globally (see chapter 3.5).

1. **Initialize RODs**
   This loads the configurations for all FPGAs and DSPs on the RODs and the BOCs. The BOCs will begin immediately to transmit the bunch-crossing clock. This gives a known input signal to the optoboard.

2. **Switch on WIENER and SC-OLink channels**
   As the daughterboards of the voltage regulators at PP2 are powered by their input voltage, it is necessary to switch on the corresponding WIENER channels as well as the SC-OLink channel for Vvdc before starting to operate the regulators. In the SC-OLink usually Vpin, Viset and Vvdc are switched on at the same time. Thus the PiN diode and the reference voltage Viset are applied to the optoboard before Vvdc is switched on at PP2. This turned out not to cause any problems.

3. **Ramp up ISEG channels**
   This provides depletion voltage to the sensors on the modules. This isolates the pixel implantations, reducing the noise at the pre-amplifier input in the readout chips. The high voltage ground is connected to VDDA on the Flex, which is not defined yet, as the voltage regulators are still switched off. This is not a problem because the ISEG channels are floating. They keep the voltage between their output terminals stable, when VDDA is switched on.

4. **Uninhibit regulator channels at PP2**
   At this point it is important to first switch on the channels for VDD and then the ones for VDDA. Otherwise the various DACs in the readout chip are not powered and the settings are unknown. This can lead to a significantly increased current consumption on VDDA before the power-on reset circuit in the FE has set all DACs to zero. The delay between switching on digital and analog supply voltage for a module should be a few seconds, as this reduces amplitude and duration of the current peak described above.

This power up sequence has been implemented in the finite state machine in DSC and has been studied thoroughly during the system test. It proved to work very well. The FSM could reliably be used to power up large parts of the detector automatically.

4.3.2　Optolink tuning

For both, the downstream TTC link and the upstream data link, various parameters have to be adjusted to ensure reliable operation.

Figure 4.7: Result of a BOC scan. The number of bit-erros in the datastream is shown color coded as a function of RX delay (x axis) and RX threshold (y axis).

For the TTC link the main parameter, that needs to be set, is the laser output power of the TX plugin. This proved to be very easy as the DORIC's dynamic threshold setting is very reliable.

For the data link the main parameters are the laser output power of the optoboard VCSEL array, adjusted through Viset for all eight channels, as well as the threshold setting for the DRX chip on the RX plugin and the phase between the received data and the 40 MHz BOC clock. These are usually referred to as the *RX threshold* and the *RX delay*, respectively.

The BOC scan

The *BOC scan* is the main tool to determine the best working point in the three dimensional parameter space for the data link. A well defined bit pattern is sent to the MCC, which returns it to the ROD. The number of bitflips in the returned data stream is color coded in a 2D histogram of RX delay and RX threshold, in figure 4.7 on the x and y axis respectively. It shows the results of a BOC scan at a MCC bandwidth of 40 MBit/s. The main features of this plot are:

- The lower error-band
  The RX threshold is set close to the 'dim' level, the light output corresponding to a logical zero. Bitflips happen from zero to one. It grows towards lower thresholds for lower Viset.

- The upper error-band
  The RX threshold is set higher than the 'bright' level, a logical one.
  Bitflips happen from one to zero. It grows towards higher thresholds
  for higher Viset. The upper error band does not occur in the scan result
  shown in figure 4.7.

- The delay band
  For these delay settings the sampling of the data stream happens dur-
  ing the rising and falling edge of the signal. Bitflips happen in both
  directions, the bit error rate is independent of the threshold between
  the lower and upper error band. Because of the pulse shape of the sig-
  nals, sent by the VCSEL on the optoboard, the delay band is rounded
  for small delays and has a sharp and stable edge for high delays. The
  lower part of the rounded edge grows towards higher thresholds and
  lower delays for increasing Viset, while the upper part grows towards
  lower thresholds and lower delays for decreasing Viset.

- The error-free region (*EFR*)
  The region between the error-bands, where no bitflips happen, is re-
  ferred to as the error-free region. It is to be maximized for reliable
  operation of the data link. As described above it depends strongly on
  Viset.

Viset is used as a third parameter in a BOC scan. The optimal setting
is chosen as the one that maximizes the EFR for most channels of the op-
toboard. RX threshold and delay are set to values slightly off the middle of
the error-free region towards higher thresholds and delays closer to the stable
edge of the delay-band.
For more details on BOC parameter tuning see [48].

The bit pattern used in the BOC scan is about 3000 bits long, there-
fore the bit error rate can only be measured down to about $1\text{x}10^{-3}$. Thus
for longer bit streams the probability of bitflips is not zero. Therefore the
settings obtained in the BOC scan are always verified by running a digital
scan, involving longer bitstreams. In case of bitflips the error-monitoring task
of the ROD would issue a warning or even interrupt the execution of the scan.

For a MCC bandwidth of 80 MBit/s, the V0 clock on the BOC must also
be tuned. It is inverted with respect to the normal bunch-crossing clock and
an additional delay between the two can be set. It is used to sample the data
that is sent at the falling edge of the XCK clock. The delay between V0 and
bunch-crossing clock, referred to as *V0 delay*, is a global BOC parameter.

It turned out that the standard delay setting is in most cases sufficient for operation at the higher MCC bandwidth. Nevertheless, in the BOC scan a second, threshold-independent delay-band is observed significantly reducing the EFR. Thus obtaining reliable settings for RX threshold and delay is somewhat harder for 80 MBit/s than for 40 MBit/s.

The optoboards installed on the PSQP used in the system test were low quality boards, not suited for installation in the detector. This resulted in a number of channels that could not be tuned correctly. Table 4.5 shows a summary of such channels for both MCC bandwidths, out of the 144 channels available. Based on these numbers and the problems found in the system test (see [48]) the qualification procedure for the optoboards was changed and the performance of the optoboards mounted in the final experiment is expected to be significantly better than that observed in the system test.

bandwidth	non-tunable channels	percentage
40 MBit/s	10	7%
80 MBit/s	30	21%

Table 4.5: Number of optoboard channels that could not be tuned in the system test

Although the optimal settings for RX threshold and delay are sensitive to the temperature of the optoboard, they proved very stable with time, once set reliably. During the entire operation of the system test, repeated tunings under the same operating conditions yielded very similar values.

4.3.3 Calibration scans

Digital scan

In a digital scan the voltage strobe generated by the internal chopper circuit of the pixel is applied to the output of the discriminator. Thus a hit is always generated with the TOT defined by the duration of the strobe. This tests the functionality of the RAM cells in the pixels, the circuitry that calculates TOT, and the end-of-column buffers. For a known performance of the module, the digital scan is also used to verify the operation of the optical link. As the data stream is longer than in the BOC scan, the sensitivity on bit errors is higher (see above).

Figure 4.8: Simulated raw data of a threshold scan as described in the text.

Threshold scan

As mentioned in chapter 3.2.2, a chopper circuit is implemented in the FE-chip that allows injection of a known charge into the input of the pre-amplifier. A threshold scan measures the number of hits versus the injected charge, repeating the injection 100 times.

Ideally, the result is the transfer function of a perfect discriminator, which is a step-function (i.e. 0 hits for charges below the discriminator threshold and 100 for charges above the threshold). In reality, noise in the readout circuitry accounts for a smearing of the ideal transfer function. As the noise is random and uncorrelated, the data can be described by a convolution of the step-function and a gaussian distribution, i.e. the gaussian error-function. This function has two parameters: the discriminator threshold, defined as the charge at which 50 % of the injections cause hits, and the electronics noise, defined as the width of the gaussian (see figure 4.8, the height of the plateau is given by the number of injections, and thus not treated as a free parameter).

The slave DSPs on the ROD fit a gaussian error-function to the raw data. The fit-results for every pixel are displayed in two-dimensional histograms representing maps of the respective module. Thus after fitting three 2D histograms, containing the threshold, noise and the χ^2 value of the fit, are downloaded from the ROD to the host PC. Figure 4.9 shows typical examples of the results of a threshold scan for one module. The topmost plots are the 2D map of the module, color-coding the threshold/noise for every single
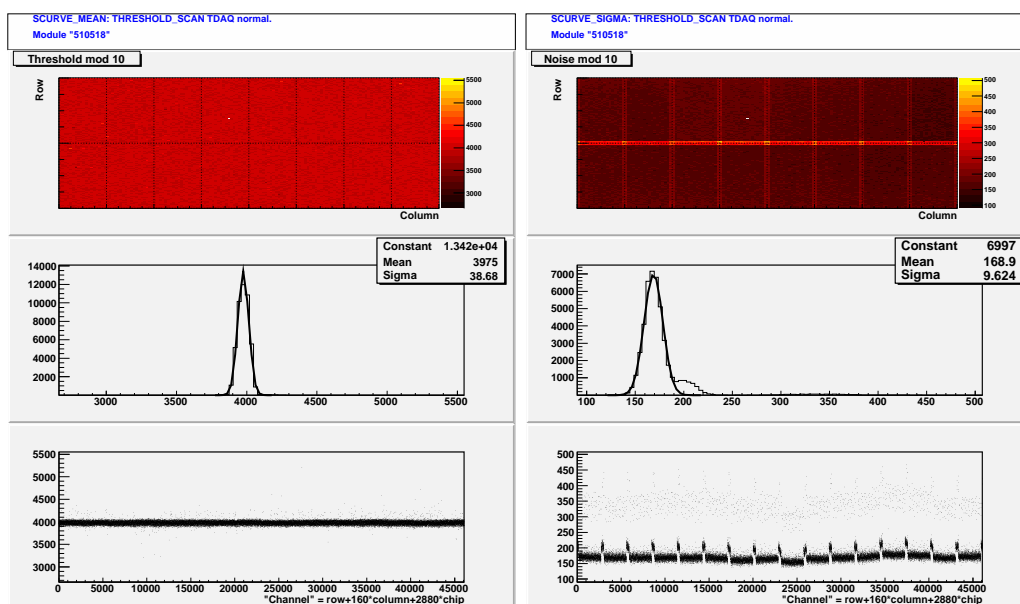
Figure 4.9: Distribution of thresholds (left) and noise (right) on a module.

pixel. In the middle the values are histogrammed and fitted with a gaussian distribution. On the bottom a scatter plot shows the threshold/noise values versus the pixel ID.

In the plots of the noise one can clearly see the 'ganged' pixels along the short edges of the FE-chips with their significantly increased noise and the 'long' pixels along the long edges of the FE-chips that have slightly higher noise than the average pixel (see chapter 3.2.1).

From the values of threshold, dispersion of the thresholds on the module, and the noise, a lot of information can be deduced on the status of a module (see table 4.6). The scan provides information on the module ID, on the high voltage supply for the sensor, on the quality of the bump-bonds, on mechanical damage to the sensor and more. Therefore the threshold scan is one of the most important tools to characterize the modules and the operation of the pixel detector system. Running a threshold scan on all 48 modules connected to four RODs in parallel took about 30 minutes. Saving the scan results took approximately one quarter of that time. As the number of RODs that can be operated in parallel is going to increase while the time it takes to save the data is going to decrease, it might be feasible to scan the thresholds of all modules in the pixel detector during every filling of the LHC ring.

condition	threshold [e$^-$]	dispersion [e$^-$]	noise [e$^-$]
system test tuning	4000	35	165
production tuning	4000	65	165
untuned	X	≈ 600	165
wrong configuration	3500 - 4500	1000	165
no HV (IZM)	no threshold measurable		
no HV (AMS)	≈ 4000	X	400

Table 4.6: Error diagnosis using threshold, dispersion and noise. All values ± 10e$^-$.

Threshold tuning

The threshold can be fine-tuned for every pixel on a FE-chip using the 7 bit TuneDAC (*TDAC*) as mentioned in 3.2.2. The goal is to have a homogeneous threshold distribution over the 46080 pixels on a module, at a target threshold of 4000 e$^-$. The distribution of the TDACs after the fine-tuning is characteristic for the module and if set in another module usually results in a threshold distribution roughly around the target threshold, with a very high dispersion (see 'wrong configuration' in table 4.6).

The algorithm used to determine the optimal setting of the TDACs in the production tests (*production tuning*) done in the module test labs, is different from the algorithm used in the system test (*system test tuning*). During production, a set of threshold scans is done for fixed TDAC values and the data is fitted with the expected curve from simulation of the circuit. From this fit the optimal settings for the TDACs is calculated. In the system test, the TDAC values are incremented or decremented, starting from a default value, depending on the result of a threshold scan. Thus the optimal setting is arrived at iteratively.

As shown in table 4.6, the algorithm used in the system test yields a better uniformity of the threshold over the module (threshold dispersion around 35 e$^-$ instead of 65 e$^-$), although the TDAC distributions for both tuning algorithms are very strongly correlated (the average correlation coefficient is 0.992$\pm$0.002). As the number of threshold scans involved is smaller than for the production tuning, the algorithm is also faster in tuning one module than the one used during production.

Running a threshold tuning on all modules in the system test took about 24 hours, which is about the time expected for tuning of the whole pixel detector.

4.4 Module performance

After verifying the digital functionality of the modules in the *digital scan*, various calibration measurements were done to analyze the performance of the modules. The most important quantities to be studied are the threshold and the electronics noise of the modules. Both are measured by a *threshold scan*.

During these scans not all pixels are tested at the same time. To avoid loss of hits due to overflow of the end-of-column buffers, only 10 pixels per double-column are tested at a time. The mask selecting these pixels is referred to as the *32-step mask*, as it needs to be changed 32 times to test all pixels on the FE-chip. For special scans the number of pixels scanned can be increased (e.g. to test all end-of-column buffer cells) or decreased (e.g. to measure crosstalk).

4.4.1 Threshold and noise on endcap A

Threshold scans were done repeatedly for the modules on the endcap. Configurations using the TDAC tuning obtained during the production of the modules, were used as well as configurations with the TDACs tuned during the system test. In the following, these are referred to as 'production tuning' and 'system test tuning' respectively.

A number of the modules on endcap A could not be tested for various reasons:

- 10 channels of the optical link were not operable

- one module was disconnected because of a missing NTC line

- one module was disconnected because of a missing sense line connection at PP2

Figures 4.10 and 4.11 show the distributions of threshold, threshold dispersion and noise of all scanned modules, and the results of a fit with a gaussian distribution. With standard deviations of about 1 e$^-$ for thresholds and threshold dispersions, and about 10 e$^-$ for noise, they show a very good uniformity of the analog performance of the modules. This is important to obtain a uniform single hit probability throughout the pixel detector.

The outliers in the distributions are due to bad TDAC tuning in case of threshold and threshold dispersion, which could be improved manually. The single module showing very high noise is due to a missing high voltage connection.

Figure 4.10: Distribution of threshold and threshold dispersion values for the modules on endcap A. Outliers are due to not perfectly tuned modules.

Tuning performance

To study the performance of the different tuning algorithms described above, threshold scans with the different TDAC configurations, taken with the system test setup, are compared. Figures 4.12 and 4.13 show the distributions of threshold, threshold dispersion and noise for all tested modules. Figure 4.12 left shows the distribution of the mean threshold values for all modules. Thresholds are centered around the target threshold of 4000 e$^-$ for both tuning algorithms, but the spread of the distribution for the system test tuning is significantly smaller.

Figure 4.12 right shows the values for the dispersion of the thresholds per module. As described above the dispersions for the system test tuning algorithm is centered around 35 e$^-$ while for the production tuning dispersions are around 80 e$^-$.
As is to be expected, figure 4.13 shows that the noise values are not signifi-

Figure 4.11: Distribution of noise values for the modules on endcap A.

Figure 4.12: Comparison of thresholds and threshold dispersions between production and system test tuning.

algorithm	threshold [e⁻]	dispersion [e⁻]	noise [e⁻]
production	4009 ± 34	83 ± 15	169 ± 5
system test	3988 ± 15	35 ± 2	162 ± 8

Table 4.7: Comparison of tuning algorithms

cantly different for scans with the two different configurations. Typical noise values are around 170 e⁻.

The values in table 4.7 show very good homogeneity of threshold, threshold dispersion and noise over all modules of endcap A. This is important to obtain a constant single hit efficiency throughout the pixel detector.

Figure 4.13: Comparison of noise values for production and system test tuning.

Figure 4.14: Distribution of mean noise values for production and system test scans.

Comparison with performance during production

The only quantity independent of the threshold tuning is the noise of the pixel electronics. Figure 4.14 shows the distributions of the noise values for scans taken during module production and during the system test.

scan	noise $[e^-]$
production	191 ± 7
system test	161 ± 10

Table 4.8: Mean noise values for the modules on endcap A from production and system test threshold scans.

A difference in the noise values of about 30 e^- is observed between the scans during production and the system test scans. No masks were applied to noisy pixels and the temperature is about the same for the two sets of scans. Possible reasons for the difference include:

- Different power supplies and cables.

- Operation of several modules mounted on the local support structures as opposed to single module operation.

- Different data-acquisition soft- and hardware are known to cause differences.

A comparable reduction in noise has been measured in specialized noise scans and will be described in the next chapter.

Figure 4.15: TOT mean and sigma calculated from TOT calibration fit results for system test and production scans.

TOT performance

The information on the amplitude of the charge signal registered in the front-end electronics can be used to improve the spatial resolution [21]. Using the TOT information a spatial resolution of $8.4\,\mu$m can be achieved for clustered hits. It is calibrated using the charge-injection circuit in the FE-chip. The calibration data from scans done during module production is compared to similar data from system test measurements.
The data is fitted with a function of the form:

$$TOT = A + \frac{B}{C + Q} \tag{4.2}$$

Using this formula, the fit-parameters A,B and C are huge numbers which cannot be compared easily. Therefore the TOT for a fixed charge of 20000 e^-, calculated pixel by pixel from the fit-parameters, is compared in figure 4.15. The increased spread of TOT mean values and the slightly higher width of the TOT distributions is due to the new TDAC tuning. As the threshold changes, so does the TOT.

4.5   Cosmics data-taking with endcap A

The final goal of the system test was to demonstrate the full functionality of the pixel detector system, by measuring muons from cosmic radiation. This data is used to exercise the offline data analysis software on real data, taken with the full data acquisition chain in place. As the cosmics data-taking was part of the system test, the data acquisition setup is described in this chapter. The cosmics data was analyzed, regarding the influence of noise in the setup. The results are presented in chapter 5.5.

4.5.1   Data acquisition system

For data acquisition, as well as for calibration, twelve pairs of RODs and BOCs were used, distributed over two ROD crates. One crate was equipped with eight RODs, reading out 96 modules, while the other crate held four RODs, reading out 48 modules. Because of problems in memory allocation on the SBC, only four RODs per crate could be operated at a time. Each crate was equipped with a TIM to provide clock synchronisation and triggering during the data acquisition runs. All BOCs used to read out the endcap modules were equipped with two sets of RX- and TX-plugins, to allow readout using the 80 MBit/s mode of the module MCCs. Thus a Disk/Layer 1 configuration could be tested.

In the data acquisition runs, triggers were either generated by a pulse generator or by a coincidence of two planes of scintillators mounted above and below the position of the endcap outside the dry box (see figure 4.16). The trigger signals were provided to a local trigger processor (*LTP*) crate. The LTP crate was used to generate the Level1- and Bunch-Crossing identifiers (L1ID, BCID), several reset signals (ECR, BCR), as well as to handle the ROD busy signals and to provide calibration and trigger signals to the TIMs. The delay of the trigger signals, necessary to account for the trigger latency in the final experiment, was set in the TIMs.

Event data from the modules was transmitted to a ROS PC via S-Link. The data fragments from the ROS were then transmitted to an EventBuilder PC, which stored the completely processed events on a hard disk. From there the data was written to CASTOR[3] for long-term storage.

The rate of Level1 triggers, that could continuously be processed, depended on the number of modules connected to this setup, and to the noise occupancy on the modules. Under normal operating conditions, trigger rates of 50 kHz were achieved with the random trigger generator. If the trigger rate was too high individual RODs would issue BUSY signals, disabling trigger generation.
To suppress low-momentum muons, that suffer the most from multiple scattering, an iron absorber was placed between the endcap and the bottom scintillator. Due to very limited space available, the block of iron was only about 12 cm thick, corresponding to a momentum cut-off of 140 MeV. The iron is shown in figure 4.16, directly above the bottom scintillator. The large

---

[3]Cern Advanced STORage manager, a hierarchical storage management system developed at CERN.

Figure 4.16: Schematic view of the scintillator trigger setup used for cosmics data taking. [48]

scintillators on either side of the bottom one improved the trigger efficiency by increasing the angular acceptance [49].
Monte Carlo simulation of the scintillator trigger setup yielded expected rates of coincidences between the top scintillator and any of the bottom ones, of 16-18 Hz. This agrees very well with the effective trigger frequency, measured in the data taking run, of about 13 Hz.

Chapter 5

Noise performance studies

> *You think that's noise - you ain't heard nuttin' yet!*
> Al Jolson

Noise is an unwanted signal superimposed to the signal that is to be analyzed [50]. It can be produced by external noise sources (*pickup noise*), for example radio transmitters, or by the activity of digital components on a mixed-mode ASIC (digital and analog circuitry in one chip, like the FE readout chip). These types of noise can be reduced by proper design of the circuitry. Unfortunately noise is also an inherent property of any electronics components, like radiation sensors, transistors and resistors. This intrinsic noise is a statistical phenomenon related to thermal movement of the charge carriers and to statistical fluctuations in the generation and transport of charge carriers.

The mean value of a noise voltage or current over a long time is zero. Therefore the mean is not suited for a quantitative description of noise. Instead it is described as rms noise:

$$\langle i^2 \rangle = \sigma^2 = \lim_{T \to \infty} \frac{1}{T} \int_{-T/2}^{T/2} (i - \langle i \rangle)^2 dt \tag{5.1}$$

The quantity $\langle i^2 \rangle$ is called the noise power. Thevenin's theorem states that a noise current source connected in parallel to a load is equivalent to a noise voltage source connected in series with a load.

$$\langle v^2 \rangle = \langle i^2 \rangle \cdot R^2$$

if R is the load resistance. Therefore noise power can also be written in terms of $\langle v^2 \rangle$.

The overall noise in the readout electronics was one of the most stringent requirements in the design of the front-end chip. Keeping the noise level low is necessary, as high noise occupancy can lead to loss of event data due to the high data rate. A noise occupancy above a certain level will also deteriorate tracking performance and thus vertex resolution and b-tagging efficiency.
As dispersion of the discriminator thresholds over a module changes the hit probability, it can be treated as an additional noise component and is added quadratically to the electronics noise. Therefore noise requirements as defined in the Technical Design Report [51] are:

| noise after irradiation | $< 200 \ e^-$ |
| maximum threshold dispersion | $< 200 \ e^-$ |

Table 5.1: Noise requirements for the pixel electronics.

5.1 Single channel noise

5.1.1 Noise in analog electronics

The sources of noise in a system consisting of a sensor, i.e. a diode in reverse-bias, and a charge sensitive amplifier (the analog part of the pixel cells in the FE readout chip) are the following:

- **Shot noise**
  The sensor leakage current is caused by the generation of electron/hole pairs over an energy barrier in the sensor. As this is a statistical process the current is not constant at the quantum level. The result is shot noise at the input of the amplifier. Therefore shot noise can only appear if a voltage is applied that generates a current. It can be shown that the noise power of shot noise is:

$$\langle i^2 \rangle = 2e \cdot \langle i \rangle \cdot \Delta f \tag{5.2}$$

  with the elementary charge e, the frequency range Δf and the average current $\langle i \rangle$.

- **Thermal or Johnson's noise**
  If the NMOS transistor at the input of the pre-amplifier is operated in strong inversion, as is the case for the FE-chip, the channel can be treated as a resistor with nonuniform resistance. In a resistor the free

charge carriers move thermally, creating a noise current independent of whether or not an external voltage is applied. The noise power is:

$$\langle i^2 \rangle = 4k_B \frac{T}{R} \cdot \Delta f \tag{5.3}$$

with Boltzmann's constant k_B and temperature T.
For the thermal noise in the channel of a NMOS transistor operated in strong inversion this becomes

$$\langle i^2 \rangle = 4k_B T \cdot \frac{2}{3} g_m \cdot \Delta f \tag{5.4}$$

with the transconductance g_m of the transistor.

- **Flicker or 1/f noise**
  1/f noise results from a variety of effects, such as impurities in a conductive channel of a transistor, generation and recombination noise due to a base current in the transistor and so on. One example is charge trapping:
  Charge trapping in the gate insulator leads to fluctuations in the oxide/bulk charge, which influences the charge carrier density and mobility. Due to the exponential decay of trapped charges, the noise frequency spectrum has a characteristic 1/f dependence. Hence its name. The noise power is:

$$\langle i^2 \rangle = \frac{K}{C_{ox} W L} \cdot \Delta f \tag{5.5}$$

  Here K is a constant, depending on the geometry of the transistor and the production technology. C_{ox}, W and L are the gate capacitance per unit area and the effective transistor width and length, respectively.

One can see that the noise powers for shot noise and Johnson noise do not depend on frequency, but only on the frequency interval accessible to measurement. This means that noise power is constant for all frequencies. This is called 'white' noise. This is not the case for 1/f noise, hence it is said to be 'pink' noise.
Using the transfer function of the charge sensitive amplifier for serial noise current and parallel noise voltage, one can calculate the noise level at the output of the amplifier due to the noise sources mentioned above. It is usually given in equivalent noise charge (ENC), which is the input charge signal that produces an output pulse with an amplitude equal to the root-mean-square (RMS) noise value at the output. A lengthy derivation, see for

example [50], yields:

$$(ENC)^2 = \frac{1}{4}\frac{a^{\left(\frac{2a}{a-1}\right)}}{1+a}\left(\tau_f S_{i_{ns}} + \frac{C_i^2}{a\tau_f}S_{v_{nT}} + 4C_i^2 k_f \frac{1}{a-1}\ln(a)\right) \qquad (5.6)$$

where $a = \tau_r/\tau_f$, τ_r being the signal rise time constant, τ_f the feedback time constant, C_i the capacitance at the input of the amplifier, k_f a typical constant for the 1/f noise and $S_{v_{nT}}$, $S_{v_{n1/f}}$ and $S_{i_{ns}}$ the power spectral densities $\frac{\langle i^2\rangle}{\Delta f}$ or $\frac{\langle v^2\rangle}{\Delta f}$ of thermal, 1/f and shot noise respectively.

5.1.2 Digital crosstalk

The front-end chip is a mixed-mode integrated circuit, meaning that every pixel contains digital and analog circuitry. Crosstalk between these circuits can happen through the bulk material of the chip. Each change of the logic state in a digital circuit implies a change of its node potentials. In the case of a CMOS digital circuit, the potential swings between the ground potential DGND (logic 0) and the positive supply voltage VDD (logic 1). Each node has a certain capacitance. In order to charge this capacitance C, the power supply lines have to provide a charge $Q = VDD \cdot C$. If many circuits change state at the same time the large transient current causes a voltage drop in the supply lines, due to their finite resistance. Thus digital activity causes changes in the DGND and the VDD potentials. Through the gate-bulk and drain-bulk capacitances these changes can inject charge in the bulk material. The short charge pulses generate noise voltage that propagates through the bulk.

In order to reduce digital crosstalk, analog and digital circuits are powered by separate supply voltages. Furthermore the transistor at the input of the pre-amplifier is implemented as a PMOS transistor sitting in its own n implantation, thus separated from the bulk.

Still an influence of digital crosstalk is measurable. If one displays the analog pre-amplifier output signal of the test pixel on a single-chip assembly on an oscilloscope, a clear influence of the 40 MHz bunch-crossing clock is visible.

5.1.3 Noise in a threshold scan

The threshold scan measures the electronics noise via the width of the gaussian error-function, fitted to the number of hits versus injected charge (see chapter 4.3.3). Typical noise values can be found in table 5.2. Noise in the long pixels is slightly increased due to their larger area and thus higher

capacitance at the pre-amplifier input. The ganged pixels, with two sensor pixel implants connected together via a metal trace on the sensor, have a still larger capacitance at the pre-amplifier input node. For them the input transistor sizes and bias currents have been adopted to this large capacitance. Still the noise is about 1.8 times higher than for normal pixels.

Noise measured in a threshold scan provides a very good way to determine whether or not a sensor pixel is connected to a readout circuit. Usually this is done irradiating a module with a radioactive source (^{241}Am was used during module production), measuring real charge deposition in the sensor. If this is not possible, e.g. in the integrated detector, the noise measurement is consulted. In the case of a broken or missing bump-bond connection, the input capacitance is significantly reduced and the noise drops to about 120 e$^-$. Using this method it was found, that the number of broken bumps increased during every thermal cycle (cooling the module and heating it up again) for one module. Figure 5.1 shows a map of the noise values for this module after the last cooling during the system test. The disconnected regions can be identified on the left-hand side of the module.
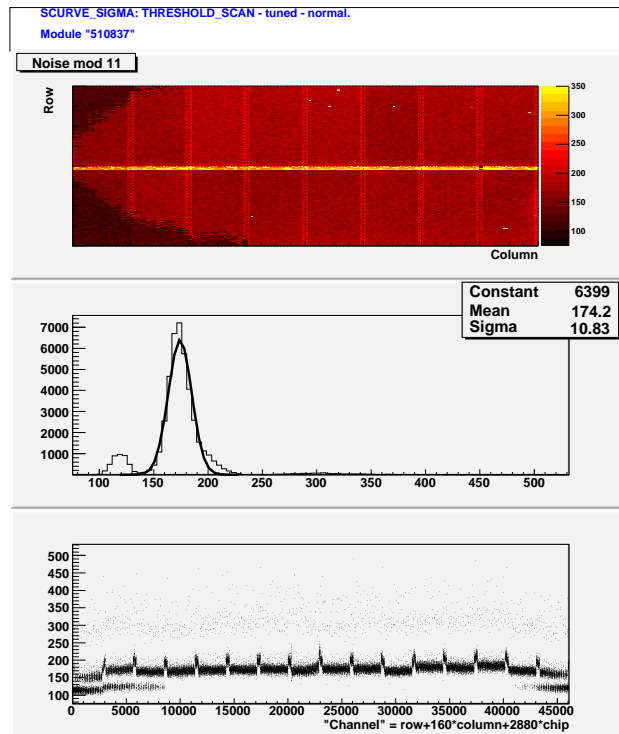
Figure 5.1: Map of noise values for module 510857. Note the regions of low noise at the left-hand side of the module, where the bump-bonds are broken.

	threshold			dispersion	noise		
	normal	long	ganged		normal	long	ganged
	normal threshold scan						
mean	4002.4			34.0	166.6	194.6	298.1
deviation	1.2			0.8	8.2	10.6	30.0
	antikill threshold scan						
mean	3976.3	3963.4	3910.0	36.9	154.8	177.4	241.2
deviation	3.9	6.0	18.7	1.3	5.6	7.7	17.4
	difference between above scans						
mean	26.1	39.0	92.5	-2.9	11.1	16.5	56.2
deviation	4.2	6.3	18.7	1.2	1.6	2.8	14.4

Table 5.2: Difference in threshold, threshold dispersion and noise between a normal threshold scan and one using the antikill option. Given are the mean values and standard deviations for 45 modules. Values in electrons ENC.

Antikill threshold scans

As described in chapter 3.2.2, the amplifiers in the pixels can be switched off individually. Running a threshold scan using this option, digital activity on a front-end chip is reduced to a minimum as no discriminator fires except for the ones being tested at that particular moment.

  The difference between a scan using this option, generally referred to as an *antikill threshold scan* or simply an *antikill scan*, and a normal threshold scan, are twofold: the threshold over all pixels decreases by typically about 30 e$^-$ and the noise over all pixels decreases by about 10 e$^-$.
Studying the normal, long and ganged pixels individually shows that the decrease, both in threshold and in noise, is largest for the ganged pixels, as shown in table 5.2. As the modules were tuned to a threshold of 4000 e$^-$ in normal mode, the threshold values obtained in the normal threshold scan do not differ between normal, long and ganged pixels.
This behavior indicates that ganged pixels are the ones most influenced by digital crosstalk. Threshold scans in normal and antikill mode were performed on modules that have a glass-substrate instead of a silicon sensor. These scans do not show the behavior described above. Noise does not differ significantly at all while the difference in threshold values is small but the same for normal, long and ganged pixels. This indicates that the crosstalk happens via the sensor bulk.

Figure 5.2: Raw data of a high statistics threshold scan, fitted with a gaussian error-function (red curve). Injected charge is given in VCal DAC value.

Gaussian noise distribution

The noise measurement during a threshold scan is based on the assumption, that the noise overlaying the discriminator step-function follows a gaussian distribution. To verify this hypothesis, threshold scans were done with high statistics to enhance possible systematic differences between the data and the fitted error-function. Figure 5.2 shows the fitted raw data for one pixel. The fit underestimates the data for charges below the 50% point, while overestimating it for high charges. If the noise is gaussian, the hit-probability for

Figure 5.3: Integral of the difference between the differential of the s-curve and a gaussian distribution. Shown are the values for all working pixels in the high statictics threshold scan.

a given charge can be calculated using this formula:

$$P(Q_{inj}) = \int_{\mu}^{\infty} \frac{1}{\sigma_{noise}\sqrt{2\pi}} \, \exp\left(-\frac{(q-Q_{inj})^2}{2\sigma_{noise}^2}\right) \tag{5.7}$$

Here σ_{noise} is the width of the gaussian distribution, and the mean value μ corresponds to the threshold charge. Differentiating the number of hits versus injected charge should therefore yield a gaussian distribution. To determine whether this is true, the differential was fitted with a gaussian distribution, and the difference Δ between the value of the fit-function and the data was calculated for every bin in steps of VCal DAC:

$$\Delta(Q) = Gauss(Q) - Data(Q) \tag{5.8}$$

The integral over the $\Delta(Q)$ distribution is expected to be zero, if no systematic differences between the differential and the gaussian distribution exist. Figure 5.3 shows the integral for all working pixels in the high statistics threshold scan. All values are negative, showing that the gaussian distribution systematically underestimates the data.
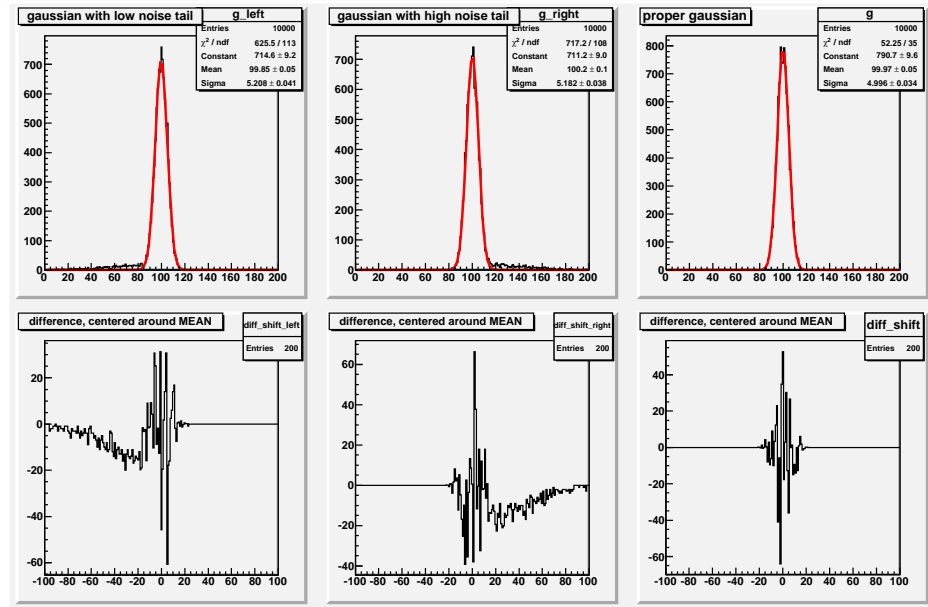
Figure 5.4: Top: Simulated gaussian distributions with low-noise and high-noise tails and without any tail. The red curve shows a fitted gaussian. Bottom: $\Delta(Q\text{-}Q_{thr})$ for above cases. High/Low-Noise tails are clearly visible.

Figure 5.5: ΔI (see text) for all working pixels in the high statistics threshold scan.

Calculating $\Delta(Q)$ for charges below and above the mean value of the fitted gaussian provides a measure to find asymmetric noise tails. The quantity

$$\Delta I = \int_0^{Q_{thr}} \Delta(Q)dQ - \int_{Q_{thr}}^{Q_{inj}^{max}} \Delta(Q)dQ$$

is the integral of the differences $\Delta(Q)$ below and above the mean value of the distribution. It is smaller than zero in the presence of a low-noise tail, larger than zero for a high-noise tail and approximately zero for a proper symmetric distribution without tails. ΔI was calculated for simulated gaussian noise distributions with tails of different amplitudes (see figure 5.4). The simulations show that $\Delta I \leq -900$ for low-noise tails and $\Delta I \geq 900$ for high noise tails. For a gaussian distribution without tails it is $-200 \leq \Delta I \leq 200$. Figure 5.5 shows the distribution of ΔI for all working pixels in the high-statistics threshold scan. No clear indication of a tail towards high or low noise is visible.

Although a gaussian distribution dos not seem to be a perfect description of the electronics noise, no asymmetric noise tails can be observed. Therefore the interpretation of the parameter σ of the gaussian error-function, fitted to threshold scan data, as the electronics noise is valid.

5.2   The minimum operating threshold

The minimum operating threshold is a very good test of the noise performance, because it incorporates all possible noise effects. It is sensitive to coherent noise, as well as the single-channel random noise. Although it does not yield a total number for the noise amplitude, it is an easy way to quanlify the noise performance under varying circumstances. It was, for eaxmple,
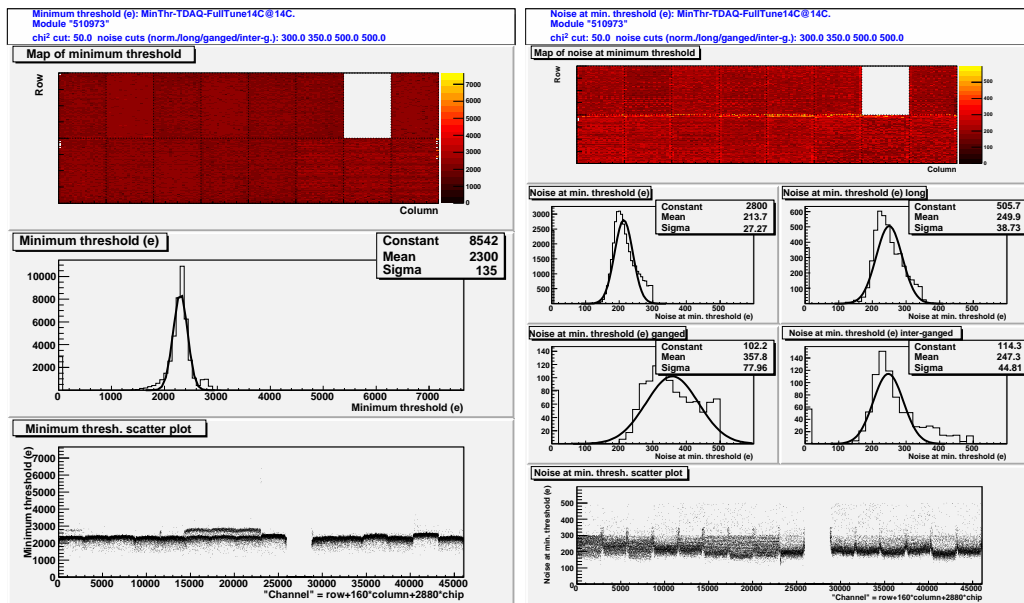
Figure 5.6: Left: Distribution of minimum operating thresholds; Right: Distribution of noise values at the minimum operating threshold. Chip 9 is known to not be working for this module.

used to compare the noise performance of modules in a small scale system test setup to the single module noise performance.

To determine the lowest possible threshold a set of threshold scans is done, decrementing the TDAC value for each pixel. If the χ^2 value of a fit to the data is 50 or higher the pixel is defined to be noisy. Additionally, cuts on the noise, measured in the threshold scans, can be defined. This determines the minimum operating threshold for this pixel.

Measurements of the minimum threshold were carried out in the small scale system test at Wuppertal University [52], using an electrical readout system, as used in the production tests. Figure 5.6 shows typical results for minimal threshold and noise at the minimal threshold. For modules configured to an operating threshold of 4000 e$^-$, the typical minimal threshold is around 2300 e$^-$. Plotting noise, measured in the threshold scans, versus the number of steps the TDACs are decremented, a significant increase in the noise values for low-TDAC configurations can be observed.

The correlation plot for tuned TDACs and TDAC settings at the minimal threshold (figure 5.8) shows that the later is basically the original TDAC configuration shifted by about 25 steps. This difference is approximately the same for all pixels of the module. This indicates a coherent effect as the dominating factor in the generation of the minimum operating threshold. The

Figure 5.7: Noise versus TDAC decrement in a minimum threshold scan. One step in TDAC corresponds to a change in the threshold of about 75 e$^-$. Thus the x axis displays thresholds of 1800-3000 e$^-$

most likely effect is capacitive coupling of the 40 MHz clock signal through the bulk of the FE-chip [53].

Due to the limited time available during the large scale system test at CERN, and some problems with the software used for the calibration scans, these measurements could not be done there.

5.3 Common mode noise measurements

If the noise in multiple readout channels is correlated, this noise contribution is referred to as common mode noise. Unlike detector systems employing an analog readout, the zero-suppressed (quasi-)binary readout[1] of the pixel detector makes it rather difficult to detect a common mode noise component. It is not possible to measure pedestal levels that can be subtracted from the charge signal. Thus common mode noise cannot be corrected for on an event-by-event basis, requiring that it must be reduced to a negligible level. To study the common mode noise performance of the pixel detector a few techniques were developed and tested in the framework of this thesis. These are based on methods developed for the SCT [54, 55] and on the correlations between noisy channels.

Common mode noise in the readout electronics can be generated by various sources, like noise on supply lines or common electromagnetic pick-up. Electromagnetic fields can be generated by nearby switching electronics, like

---

[1]The TOT provides a digitized information on the analog signal.

Figure 5.8: Correlation between TDACs adjusted to an operating threshold of 4000 e$^-$ and the TDAC setting at the minimal threshold.

the readout circuitry of the surrounding detectors, or digital activity of neighboring modules in the pixel detector. These fields can propagate through the surrounding air or through the highly conductive carbon-carbon support structures.

If the noise levels are normalized to the single-channel random noise, denoted as σ_{noise}, the distribution of the random noise in every pixel is given by the gaussian distribution

$$P(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right),\qquad(5.9)$$

and the common mode noise by

$$P(z) = \frac{1}{\sqrt{2\pi\sigma_{cmn}^2}} \exp\left(-\frac{z^2}{2\sigma_{cmn}^2}\right),\qquad(5.10)$$

where σ_{cmn} is the magnitude of the common mode noise. As these contributions are uncorrelated they can be convoluted to give the total noise contribution

$$P(a) = \frac{1}{\sqrt{2\pi(1 + \sigma_{cmn}^2)}} \exp\left(-\frac{a^2}{2(1 + \sigma_{cmn}^2)}\right),\qquad(5.11)$$

where $a = x + z$. The number of hits above a certain threshold τ is given by the complementary error function

$$O(\tau) = \int_\tau^\infty P(a)da = \frac{1}{2}\text{erfc}\left(\frac{\tau}{\sqrt{2(1 + \sigma_{cmn}^2)}}\right), \qquad (5.12)$$

which is also referred to as an 's-curve'. It can be seen from equation 5.12 that the common mode noise adds quadratically to the single-channel random noise and thus shows up in the noise measurement in a threshold scan.

The TOT of a hit is directly proportional to the amplitude of the charge signal from the sensor. It is measured during production for an injection with a large charge to qualify the module. The distribution of TOT values for an injected charge of 20000 e$^-$ is gaussian centered around TOT=30, with $\sigma \approx 0.75$. In the presence of common mode noise a group of channels will show TOT values different from 30 and will thus broaden the TOT distribution quickly. Consequently, common mode noise can also be measured using the TOT information.

For both methods mentioned above the measurement of the common mode noise is rather straightforward. A disadvantage of these methods is, that one needs to compare the results of the dedicated common mode noise measurements with the same measurements done with no common mode noise present. As such measurements are very hard to obtain, more sophisticated techniques are required without the need for a reference.

5.3.1 Occupancy per event

Common mode noise causes coherent fluctuations in the occupancies of groups of channels from one event to the next [54, 55]. The distribution of the number of pixels registering a hit, N_e, is therefore sensitive to common mode noise. If no common mode noise component is present, the distribution of N_e is binomial, denoted $P(N_e) = Bin(N_e; O(\tau))$. Here $O(\tau)$ is the hit probability at a given threshold τ. The mean number of hits is given by

$$\overline{N_e} = nO(\tau)$$

with the number of pixels n. The variance of the distribution is

$$Var(N_e) = nO(\tau)(1 - O(\tau)) = \overline{N_e}(1 - \overline{N_e}/n)$$

Thus the expected variance of the N_e distribution is given by the observed mean number of hits per event. The variance of the binomial distribution

reaches its maximum value of $Var(N_e) = n/4$ for $\overline{N_e} = n/2$. Assuming the common mode noise magnitude is the same for all pixels in a sample, the hit probability for all pixels in this sample varies coherently. Therefore the N_e distribution will be wider than a pure binomial distribution. It should be noted that even large pixel-to-pixel variations of the pixel occupancy generally do not widen the N_e distribution. Such variations could for example originate from large pixel-to-pixel threshold variations. N_e for a mixture of channels with different occupancies is the sum of values taken from narrow distributions and the resulting distribution itself is narrow. Figure 5.9 shows simulated N_e distributions for various magnitudes of the common mode contribution. Simulated were 100000 events in a sample of 1000 pixels.

Figure 5.9: Simulated N_e distributions for different magnitudes of the common mode noise (CMN) contribution.

To derive a measure for the common mode noise contribution, one has to calculate the variance of the N_e distribution in the presence of common mode noise. Noting, that the effect of a shift in the signal due to common mode noise, is equivalent to a shift of the threshold by the same amount in the opposite direction, one can implement common mode noise by choosing

the threshold for each event from a probability distribution

$$g(\tau) = \frac{1}{\sigma_{cmn}\sqrt{2\pi}} exp\left(-\frac{(\tau - \tau_0)^2}{2\sigma_{cmn}^2}\right).$$

With the hit-probability for a given threshold p(τ) the mean number of hits per event $\overline{N_e}$ is now

$$\overline{N_e} = \int_{-\infty}^{\infty} np(\tau)g(\tau)d\tau$$

and the variance of N_e is

$$\begin{aligned} Var(N_e) &= \overline{N_e^2} - \overline{N_e}^2 \\ &= \overline{N_e}(1 - \overline{N_e}) + n(n-1)\int_{-\infty}^{\infty} p^2(\tau)g(\tau)d\tau. \end{aligned} \quad (5.13)$$

If the hit probability $p(\tau) = 0.5$ the integral in equation 5.13 can be solved:

$$\int_{-\infty}^{\infty} p^2(\tau)g(\tau)d\tau = \frac{1}{4} + \frac{1}{2\pi}\tan^{-1}\left(\frac{\sigma_{cmn}^2}{\sqrt{1 + 2\sigma_{cmn}^2}}\right). \quad (5.14)$$

Solving equation 5.13 for σ_{cmn}, we obtain

$$\sigma_{cmn}^2 = \frac{\sin(\alpha)}{1 - \sin(\alpha)} \quad (5.15)$$

where

$$\alpha := 2\pi\frac{Var(N_e) - n/4}{n(n-1)}. \quad (5.16)$$

If the mean hit probability is 50%, the common mode noise can be calculated from the observed mean and variance of the N_e distribution.
As the variance of a binomial distribution for a hit-probability of 50% equals n/4, equation 5.15 can be generalized to define the observable Γ:

$$\Gamma^2 := \frac{\sin(\tilde{\alpha})}{1 - \sin(\tilde{\alpha})} \quad (5.17)$$

where

$$\tilde{\alpha} := 2\pi\frac{Var(N_e) - \overline{N_e}(1 - \overline{N_e}/n)}{n(n-1)}. \quad (5.18)$$

Figure 5.10 shows Γ^2 as a function of the common mode noise amplitude, calculated from simulated N_e distributions. Simulated were 5000 events in a sample of 1000 pixels. The errorbars indicated show that Γ is sensitive down to very small common mode noise amplitudes of few e$^-$.

Figure 5.10: Γ as a function of the common mode noise amplitude.

Implementation

The occupancy per event method was implemented in a calibration scan. The thresholds are adjusted beforehand to minimize their dispersion. Charge corresponding to the threshold charge is injected in a number of pixels (maximum 112 pixels per FE chip, given by the MCC receiver FIFOs), thus ensuring the hit-probability to be close to 50%. The number of pixels registering a hit, N_e, is histogrammed for repeated injections in the same pixels. This scan is referred to as a *HitOccupancy scan*.

5.3.2   Average TOT per event

The distribution of the average of the TOT values for all hits in one event can also be used to determine the common mode noise. Let Q_i^{inj} be the injected charge in event i and σ_{noise}^{ij} the single-channel random noisem in pixel j. The common mode noise contribution can be denoted by σ_{cmn}^{j} for event j, assuming it is the same for all pixels i. The average charge $\overline{Q^j}$ measured for an event j is then:

$$\overline{Q^j} = \frac{1}{n} \sum_{i=0}^{n} \left(Q_i^{inj} + \sigma_{noise}^{ij} + \sigma_{cmn}^{j} \right) \tag{5.19}$$

As for a large number of pixels, n, the average single-channel random noise vanishes, this becomes

$$\overline{Q^j} = \overline{Q^{inj}} + \sigma_{cmn}^{j} \tag{5.20}$$

In the absence of common mode noise the width of the distribution of $\overline{Q_j}$ is zero if the number of pixels is large. As the possible mismatch of the FDAC settings for the measured pixels is not time-dependent, it affects the mean value of the distribution but leaves the width unchanged. Therefore the width is a good measure for common mode noise. It is directly proportional to the common mode noise amplitude σ_{cmn} in units of TOT.

Figure 5.11: Left: Simulated TOTAVERAGE distributions for various common mode noise amplitudes. Note that histogrammed is $128 \times TOTAVERAGE$ like in the calibration scan. Right: RMS width of simulated TOTAVERAGE distributions as a function of the common mode noise amplitude. The width is not zero without common mode noise because of the residual contribution of the single-channel random noise in a small sample of pixels.

Implementation

A large charge is injected in a number of pixels and the average TOT per event is histogrammed. This histogram is referred to as a TOTAVERAGE histogram.

To calculate the common mode noise in ENC electrons a mean TOT calibration is used to convert the average TOT into charge, bin by bin. As the fit parameters of the TOT calibration are very large and have a high dispersion a simple average cannot be used. Therefore each bin of the TOTAVERAGE histogram is converted into charge using the exact TOT calibration function of every pixel that was scanned, histogramming the values. The mean value of the histogram is used as the charge value, its width as an error. This gives rise to a systematic error of about 10 e$^-$.

Also the influence of the single-channel random noise not completely averaging to zero for a small sample of pixels gives rise to a systematic error. The

Figure 5.12: Left: Type0 transformer setup; Right: Setup for noise injection into the carbon-carbon support structure.

simulation of a sample of pixels of the same size, as is used in the calibration scan, yields a value of $\sigma_{totaverage} \approx 0.025$ without common mode noise, corresponding to about 10 e$^-$.

The two kinds of systematic error can be added quadratically to give a total error of approximately 14 e$^-$. Therefore this scan is sensitive to common mode noise levels down to about 45 e$^-$.

5.3.3   Verification measurements

The measurement techniques developed above were tested with deliberate injection of noise into modules using different injection mechanisms. Figure 5.12 shows setups used in the techniques discussed below.

The Type0 transformer

A Type0 cable was coiled around a ferrite core to act as the secondary side of a transformer. The primary side was connected to an arbitrary function generator.

This method simulates inductive coupling into the cables running along the backside of the local-support structures through a region densely packed with digital electronics.

Injection in the support structure

The voltage signal from the function generator was applied directly to a carbon-carbon shingle glued to the bulk-side of the readout chips of a module.

Figure 5.13: Measurement of common mode noise amplitude for noise injection using the Type0 cable transformer.

This simulates capacitive coupling between the conductive carbon-carbon material of the local support structures and the readout electronics.

Results

For both injection methods the frequency of the injected signal was varied from 1 kHz up to 10 MHz. For comparison the difference between the noise measured in a threshold scan with noise injection and a reference scan without injection is plotted.

Figure 5.13 shows that the results of the *Occupancy per event* method and the *Average TOT per event* method are consistent with the noise measurement in a threshold scan. It also shows that the readout circuitry is most sensitive to noise signals between 2 and 7 MHz. This frequency region has been shown to be most sensitive to noise interference in previous measurements. See for example [56].

An interesting feature of the methods, observed during the verification measurements, is shown in figure 5.14. A symmetric sine-wave in the frequency range between 3 and 5 MHz with an amplitude of $1V_{pp}$ on the secondary side of the Type0 transformer was used to simulate noise pickup. Depending on the frequency, the resulting N_e distribution shows a clear double-peak structure. This can be interpreted as the effects of noise pickup during the positive and the negative half of the sine-wave. If a DC

Figure 5.14: N_e distribution with noise injection. Note the double-peak structure that is generated by the amplitude of the injection signal being symmetric around zero.

offset is applied to the sine-wave, so that the injected signal is unipolar, the double-peak structure cannot be observed. A very similar structure could be observed in the TOTAVERAGE distribution.

5.3.4   Results for endcap A

The system test setup described in chapter 4.1 was used to measure the amplitude of the common mode contribution to the total noise in endcap A. For the first time the scans described above were executed on a large number of modules in parallel. Due to problems in the implementation of the scans, only a limited number of modules could be tested. As these are distributed across the whole endcap assembly, the scans still give a pretty good picture of the common mode noise in the endcap.

Occupancy per event

To achieve a hit-probability of 50% on the modules, a charge is injected in each pixel that corresponds to the discriminator threshold charge. As the thresholds of the endcap modules are very similar after fine-tuning (see figure 4.12), and a small variation of the thresholds does not have an influence on the measurement, the same charge was injected in all modules. It was optimized so that the occupancy per event for most modules was as close to 50% as possible. 1440 pixels per module were scanned. Figure 5.15 left shows an example of the N_e distributions measured in these scans. In red the gaussian distribution is shown, that is fitted to the data to obtain $\overline{N_e}$

and $Var(N_e)$. Equations 5.17 and 5.18 yield

$$\Gamma = \sigma_{cmn}/\sigma_{ran} = 0.136 \pm 0.030.$$

With a typical total noise of 160 e$^-$ this leads to a common mode noise contribution of 22 ± 5 e$^-$.
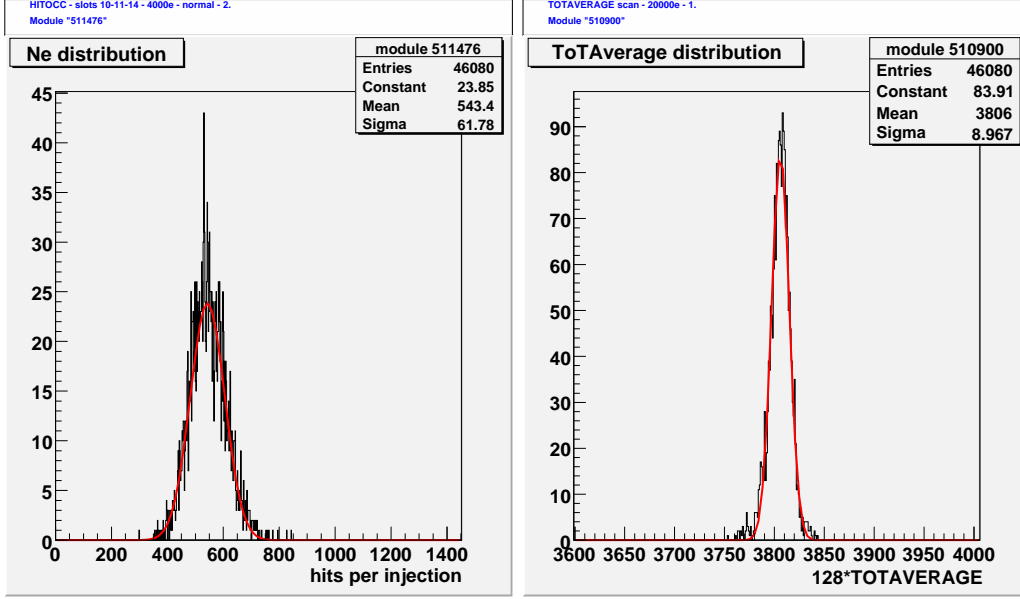
Figure 5.15: Left: Exemplary result of a HitOccupancy scan; Right: Exemplary result of a ToTAverage scan.

Average TOT per event

For this measurement a charge corresponding to 20000 e$^-$ was injected into 1440 pixels per module. Figure 5.15 right shows a typical TOTAVERAGE histogram obtained from these scans. The average TOT in an event is multiplied by 128 before histogramming. The red curve shows the fitted gaussian distribution. The average width of the TOTAVERAGE distribution on the modules in units of 25 ns is measured to be

$$\sigma_{totaverage} = 0.076 \pm 0.01.$$

Using an average TOT calibration this results in a common mode noise amplitude of 41 ± 18 e$^-$, including statistical uncertainty and systematic errors. The results from the two measurements match within these uncertainties.

These results confirm, that the common mode contribution to the overall noise in modules set to standard operating thresholds is indeed negligible.

5.4 Noise Occupancy studies

The most realistic test of the noise performance of the pixel detector is to measure the occupancy due to noise hits. This measurement is done by triggering at random, reading out the hits. This was done in the final stage of the system test at CERN, using the complete readout system together with endcap A of the final detector (see chapter 4.5.1).
The requirement for the occupancy per pixel due to noise hits, generally referred to as noise occupancy per pixel, was to be below 10^{-5}, in order not to influence the tracking efficiency. With the noise levels of the modules, as for example given in table 5.2, the threshold charge is about 20σ above the noise. Therefore the noise occupancy is expected to be very low.

In the scope of this thesis a decoding tool for the stored ROS fragments has been developed, as well as an analysis tool. The following analysis of the noise occupancy data has been done with these tools.
A list of analyzed runs and their main features can be found in table 5.4.

5.4.1 Noise occupancy with standard module configuration
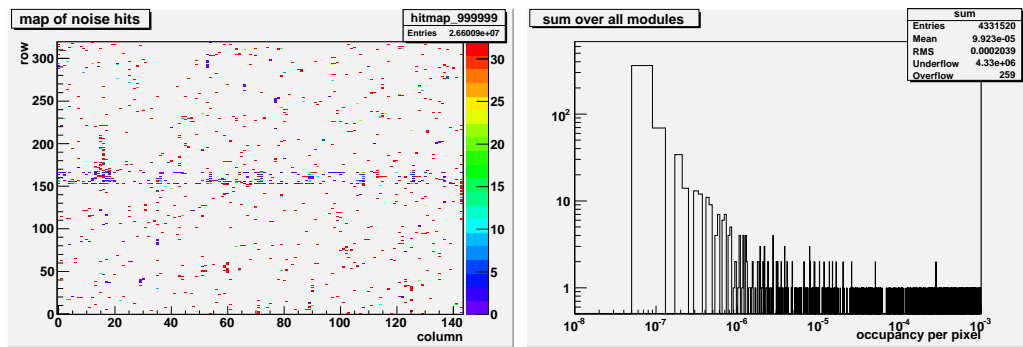
Figure 5.16: Left: Map of noise hits summed up over all active modules in run 1138. Right: Noise occupancy per pixel for all active modules.

Occupancy

The noise occupancy run with the runnumber 1138 was taken, using standard operating settings for the endcap modules, at an effective trigger frequency of 13kHz. A total of 16.8 million triggers were sent, resulting in a sensitivity

on the noise occupancy per pixel of 6×10^{-8}. In 94 modules in endcap A a total of 48 million hits was registered. As a first rough estimate this results in a noise occupancy per pixel of 6.5×10^{-7}. Figure 5.16 left shows a map of the noise hits of all modules summed up in one module map. Apart from the ganged pixels, no clusters of noise hits are visible that would imply common mode noise effects. On the right, the distribution of the noise occupancy per pixel for all active modules in this is shown. One can see that very few pixels show occupancies above the requirement of 10^{-6}. Only 1310 pixels have noise occupancies above the sensitivity limit mentioned above.

In offline analyses, a pixel is defined as being 'noisy' if it shows a noise occupancy above 10^{-4}. Sticking to this definition the number of noisy pixels in run 1138 is 495, corresponding to a fraction 0.011% of all active pixels. This agrees very well to the number of pixels with higher noise, expected from production scans.

Figure 5.17 shows the logarithm of the average noise occupancy, defined as the number of hits in a module divided by the number of pixels and the number of triggers, for all active modules in run 1138. Most of the modules show noise occupancies between 10^{-6} and 10^{-7}.

Figure 5.17: Average noise occupancy for all active modules in run 1138.

TOT spectrum

The TOT spectrum of the noise hits shows a distinct peak around TOT=5. Figure 5.18 shows the TOT spectrum for all hits in the endcap, fitted with a Landau distribution.

The tail to very high TOT values stems from noisy pixels. When multiple noise hits in the same pixel happen before the falling edge of the first hit occured, the TOT is summed up. This results in unphysically high TOT values. The minor peak at TOT=32 and the broad peak around TOT=120

Figure 5.18: TOT spectrum of all noise hits in endcap A during run 1138.

are artifacts from modules with strange TOT spectra. Figure 5.20 shows examples of TOT spectra leading to this shape.

Cluster size

Figure 5.19 shows the distribution of cluster sizes in run 1138. The probability for two-hit clusters from uncorrelated noise hits in two pixels is given by

$$
\begin{aligned}
P(two-hit) &= P(one-hit) \times P(neighbor) \\
&\approx 46080 \cdot 10^{-7} \times 2 \cdot 10^{-7}.
\end{aligned}
$$

(5.21)

The ratio of two-hit clusters to one-hit clusters is then

$$
\frac{N(two-hit)}{N(one-hit)} = \frac{P(two-hit)}{P(one-hit)} \approx 2 \cdot 10^{-7}.
$$

(5.22)

The measured ratio is around 8.6×10^{-4}.

This high rate is not due to crosstalk between pixels. In the case of crosstalk,

Figure 5.19: Distribution of cluster sizes for all modules in run 1138.

Figure 5.20: Left: TOT spectrum with high-TOT tail; Right: TOT spectrum showing artifacts. This module was found to have a local damage in the sensor, causing the strange spectrum.

one would expect the number of two-hit clusters along the long edge of a pixel to be significantly larger than along the short edge. This is because crosstalk is proportional to the inter-pixel capacity, which along the long side is higher than along the short side. Measuring the clustersizes for clusters in both directions separately yielded similar values.

Studying the maps of noise hits for individual modules one finds, that the high number of two-hit clusters is due to small regions of pixels with higher noise occupancy than their neighbors. These are mostly correlated with pixels, that show slightly higher noise in the production threshold scans.

5.4.2 Noise Occupancy vs trigger frequency

Random trigger runs were taken at trigger frequencies varying from 100 Hz up to 50 kHz. This allows to look for a possible influence of trigger frequency on noise occupancy. Due to problems with configurations in this early phase of the data taking with random triggers, the number of modules participating in these runs is limited to 29.

Detailed comparison of noise occupancies per pixel, TOT spectra and clustersizes for runs 1065, 1021 and 1138 (for details on these runs see table 5.4) show no statistically significant differences. This is not surprising, as the window during which the modules are sensitive to noise hits is the same in all runs, regardless of the trigger frequency. In all runs, single LVL1 signals were issued. With the bunch-crossing clock of 40 MHz this results in a sensitive interval of 25 ns.

5.4.3   Noise Occupancy vs threshold

Starting from module configurations with thresholds tuned to 4000 e$^-$ the TDAC settings were reduced for all pixels by a number of steps to reduce threshold. Due to mismatch in the dependenc of threshold vs TDAC between pixels, this led to an increase in threshold dispersion. As long as the threshold is significantly higher than the the minimum threshold dicussed above, noise is independent of the TDAC settings and is not expected to change. TDACs were reduced by 10, 20 and 25 steps respectively. The distributions of the thresholds, threshold dispersions and noise values for the configurations with reduced TDACs are shown in figure 5.21. Mean values for threhold and threshold dispersion for all modules are summarized in table 5.3. The noise values are around 170 e$^-$.
Information on the analyzed runs is summarized in table 5.4.

TDAC setting	threshold	threshold dispersion	avg. noise occupancy	standard deviation
tuned	3974 ± 36e$^-$	34 ± 2e$^-$	$3.5{\times}10^{-7}$	$8.1{\times}10^{-7}$
TDAC - 10	3371 ± 15e$^-$	52 ± 2e$^-$	$3.5{\times}10^{-7}$	$8.1{\times}10^{-7}$
TDAC - 20	2722 ± 35e$^-$	73 ± 4e$^-$	$1.9{\times}10^{-5}$	$1.1{\times}10^{-4}$
TDAC - 25	2392 ± 54e$^-$	88 ± 6e$^-$	$4.8{\times}10^{-4}$	$3.1{\times}10^{-4}$

Table 5.3: Mean threshold and threshold dispersion values for all modules for the TDAC reduced configurations together with the average noise occupancy per pixel over the whole endcap.

These configurations were used to study the evolution of noise occupancy with lowering threshold. The same number of events was taken in the noise occupancy runs with TDAC settings reduced by 10 and 20 steps, as in run 1138.

TDAC - 10

For the run with the TDAC settings reduced by 10 steps, the measured noise occupancy and the number of noisy pixels are similar to the values measured in run 1138. The same is true for the TOT spectrum and the ratio of two-hit clusters to one-hit clusters.

TDAC - 20

In the run with TDAC settings reduced by 20 steps, at a threshold of about 2700 e$^-$, the noise occupancy per pixel for a few modules is much higher
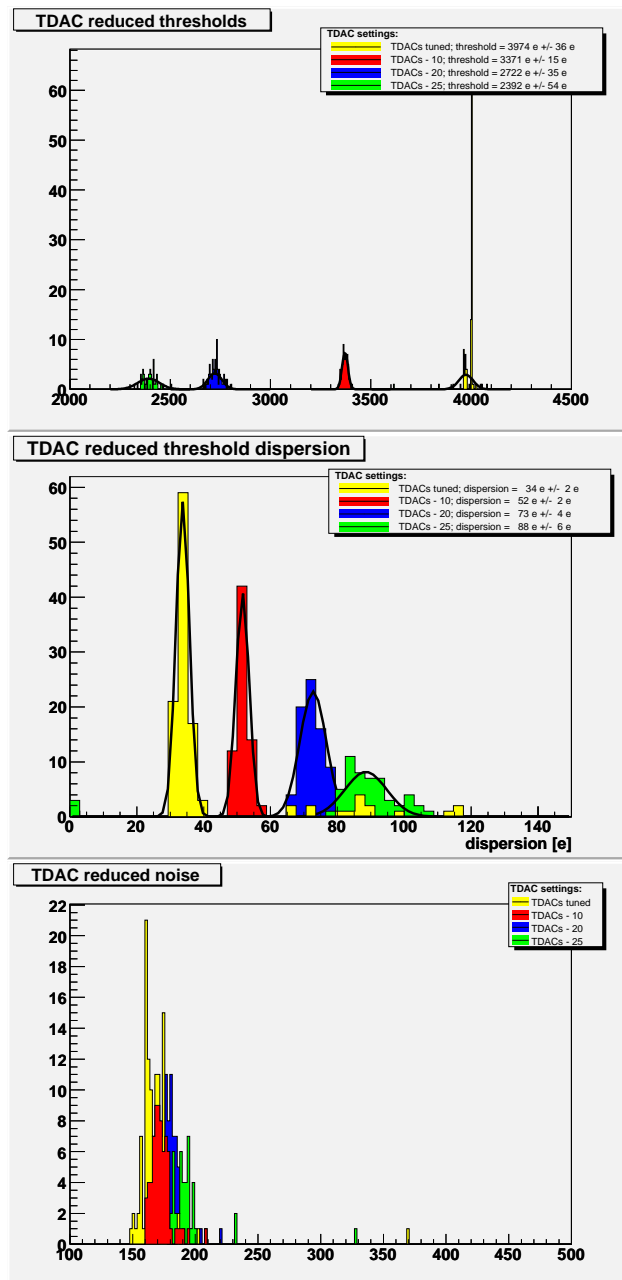
Figure 5.21: Top: Distribution of thresholds with decreased TDAC settings. Middle: Distribution of threshold dispersions. Bottom: Distribution of noise values with decreased TDAC settings.

than before. The TDAC settings are close to the minimal thresholds of these modules. This is studied in more detail in the analysis of run at TDAC-25.
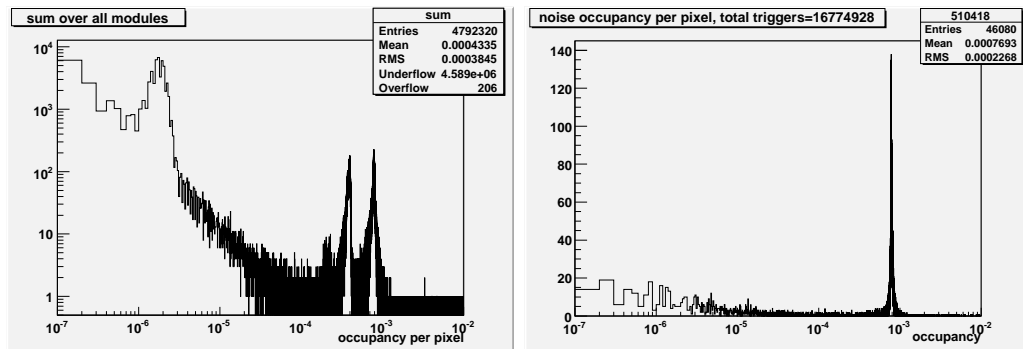
Figure 5.22: Left:Noise occupancy per pixel for all active modules in the TDAC-20 run. Right: Noise occupancy per pixel for one of the noisy modules.

TDAC - 25

In this run, a number of modules had to be disabled from data-taking. They produced so many noise hits that the readout chain could not comply. Several RODs were always busy, hampering the trigger flow from the LTP crate. These were mostly the modules already showing an increased noise occupancy in run 1144.
All remaining modules show high noise occupancy per pixel. The occupancy values range from 2×10^{-7} to 8×10^{-4}, allowing to study the module behavior over a wide range.

The noise performance of the modules can be split into three noise occupancy regimes.

- **Noise occupancy $\leq 3 \times 10^{-4}$:**
  Most modules in this 'low-occupancy' regime show a non-uniform dis-

Figure 5.23: Map of noise hits for a low-occupancy module.

Figure 5.24: Left: Distribution of noise hits per event; Right: TOT spectrum for a low-occupancy module.

tribution of noise hits. The number of hits is higher on the low-column side of the module[2] for a few FE-chips. Figure 5.23 shows an exemplary map of noise hits for a low-occupancy module. The number of readout chips showing this behavior increases with noise occupancy.

Figure 5.25: Left: Map of TOT values averaged over all hits; Right: Map of standard deviations of the average TOT values for a low-occupancy module.

The plots shown in the following are taken from a module with an average noise occupancy of 1.4×10^{-5}. It has proven to be a good example for the general performance of modules in this occupancy regime.

The maximum number of hits per event, see figure 5.24 left, is around 160. For hits homogeneously distributed across the module, this amounts to an average of about 9.5 hits/FE-chip. The average clustersize is 1.1. The TOT spectrum shows a few hits with TOT values larger than expected, but no clear tail towards high TOT. It peaks around TOT=4.5,

---

[2]see figure 3.4 for the coordinate system

as expected.

The map of average TOT values (figure 5.25 left) shows a slight variation between outer and inner chips on a module. Average TOT values for chips close to the short edge of the module are between 3 and 5, while center chips show values of 2 to 3. These account for the smaller peak in the TOT spectrum. This behavior is visible in the average TOT values for all noise occupancy regimes. Superimposed here is a row-dependence of the number of hits. In the FE-chip, two adjacent pixels in the same row share a bias-compensation circuit for two high-current DACs, that supply the two amplifier stages. This leads to small row-dependent differences in the amplifier behavior.

The map of standard-deviations of the TOT values (figure 5.25 right) shows the same variations between chips on the short edges of a module and between adjacent rows. Neither the average TOT values, nor the standard-deviations, show an influence from the pattern in the number of hits visible in the hitmap.

Figure 5.26: Map of noise hits for a medium-occupancy module.

- $3 \times 10^{-4} \leq$ **Noise occupancy** $\leq 7 \times 10^{-4}$**:**
  The modules in this occupancy regime show a distribution of noise hits different from the low-occupancy regime. Each FE-chip shows an increase in noise hits on the low-column side of the module, but in the uppermost columns of the FE-chips the high-occupancy regions also extend towards high row indices (see figure 5.26). The ganged pixels become very noisy as well. Part of the modules in this regime show a

Figure 5.27: Left: Distribution of hits per event; Right: Distribution of clustersizes for a medium-occupancy module.

strong asymmetry in noise hits between the two rows of chips on the module.

The maximum number of hits per event is around 880, resulting in an average of 55 hits per FE and event. Together with the maximum clustersize of 32 this is a clear indication for the presence of common mode noise (see figure 5.27).

The TOT spectrum (figure 5.28 right) shows a long tail towards high TOT values. This is due to the coherent fluctuation of the charge signal. By creating a large number of hits with high TOT values, common mode noise increases the width of the distribution of TOT values per pixel. This shows up in the map of standard deviations of
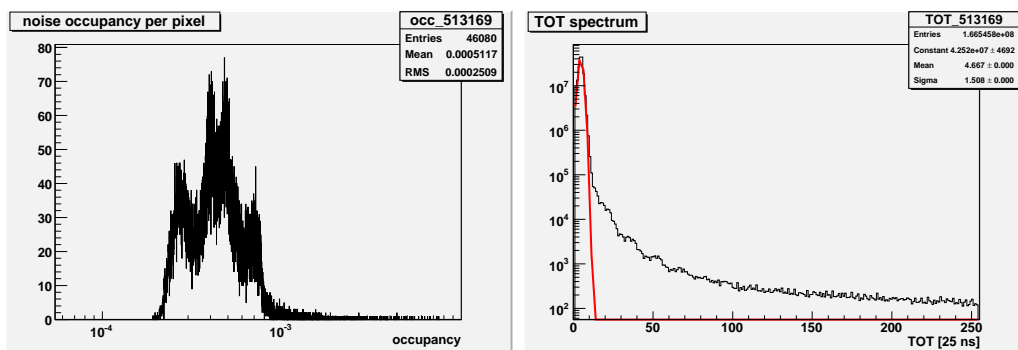
Figure 5.28: Left: Distribution of noise occupancy per pixel; Right: TOT spectrum for a medium occupancy module.

the TOT values (figure 5.29 right). Measuring the width of the TOT
distribution is therefore a quick way to determine areas with increased
common mode noise amplitude.

Figure 5.29: Left: Map of TOT values averaged over all hits; Right: Map of
standard deviations of the TOT values for a medium-occupancy module.

- $7 \times 10^{-4} \leq$ **Noise occupancy:**

Figure 5.30: Map of noise hits for a high-occupancy module.

At these very high noise occupancies the hits are distributed homo-
geneously across the modules. No asymmetries between low and high
rows on the module are observable. The hitmap is dominated by the
row-dependence already described above, and a decrease in the number
of hits by about a factor of 2 in the lower- and uppermost rows of a
module. This *edge effect* is visible in all modules with a noise occu-
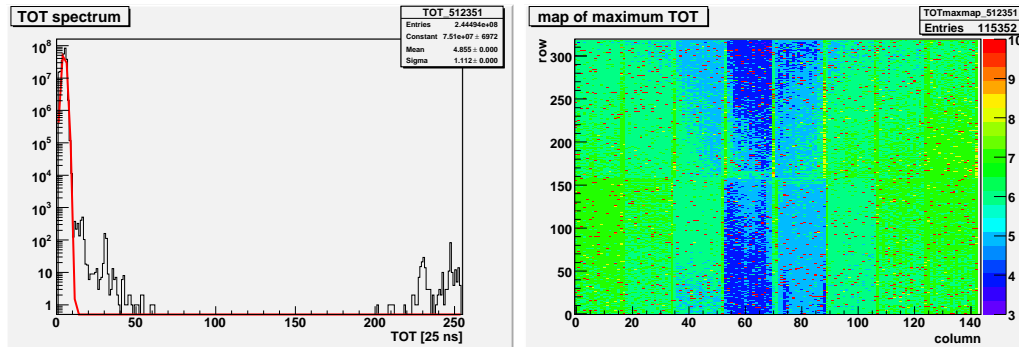pancy above 7×10^{-4}.

Figure 5.31: Left: TOT spectrum of a high-occupancy module; Right: Map of maximum TOT values.

The maximum number of hits per event is about 200, the maximum clustersize is about 3 pixels. There is no coherence visible in the noise hits. The TOT spectrum shows a distinct peak at TOT = 4.9 and few hits up to TOT $\approx$ 50. The few hits at very high TOT values are scattered across the module instead of clustering in any region (see figure 5.31).

The absence of a tail in the TOT spectrum, the small number of hits per event and the small clustersize, as well as the standard deviations of the TOT values, uniform except for the well-known row-dependence (figure 5.32), clearly indicate that no common mode noise is acting on the modules. Noise occupancy in this regime is clearly dominated by random single-channel noise.


Figure 5.32: Left: Map of average TOT values; Right: Map of standard deviations of the TOT values for a high-occupancy module.

The increase in the number of hits on the low-column side of the modules with low and medium noise occupancy cannot be explained by an asymmetry in the FE-chip itself. For chips 0 to 7 it occurs at low chip-columns, and for chips 8 to 15 it occurs at high chip-columns. Therefore an influence transmitted through the sensor is highly likely.

Modules showing increased numbers of hits in low columns are distributed all across the endcap, no sectors are preferred. Therefore asymmetric influences, like light shining on the modules through an opening in the cold-box, can be excluded. The modules are mounted on the disks such that low module-columns point towards the outer disk-radius. This is where the Type0 cables are routed towards PP0. It seems likely that electromagnetic interference from the digital activity in the Type0 cables is the main source of the observed hit pattern.

Studying the distribution of modules with low, medium and high noise occupancy across the disk, a strong correlation between the noise occupancies of modules on the same faceplate of one sector is observed. In most cases, all modules on the same faceplate of the sector show similar noise occupancies. Modules with one or both neighbors switched off[3], usually show low and medium noise occupancies. Modules with both neighbors active show medium and high noise occupancies.

Crosstalk between the modules on opposite faceplates of the disks is not very strong. Noise occupancies are determined by the neighboring modules on the same faceplate, rather than the modules on the opposite side.

The crosstalk between the modules is based on capacitive coupling of noise voltage into the conductive carbon-carbon support structure. This voltage induces noise hits in the neighboring modules, enhancing the capacitive coupling. Therefore all modules of a sector faceplate show similar noise occupancies. As the two faceplates of a sector are insulated from each other, no direct coupling happens between them. The same is true for the boundaries between to sectors, which are also isolated.

5.4.4   Noise induced inter-module crosstalk

We have observed that groups of modules with high digital activity influence neighboring modules. The same could be true for single noisy modules. To study this effect the thresholds of individual modules were set to the values known from the high-noise occupancy run 1150. Detailed comparison of the

---

[3]This was necessary during the noise occupancy runs due to too high occupancy on a number of modules

Figure 5.33: Greyscale images of special TDAC masks applied to four modules in run 1151. White corresponds to TDAC+25, black corresponds to TDAC-25.

performance of neighboring modules, both on the same side of a disk and directly beneath the low-threshold module on the opposite side of the disk, with data from run 1138 yielded no measurable differences. The sensitivity limit on noise occupancy, given by the number of triggers in the run, is 6×10^{-8}. Within this accuracy, neither noise occupancy, nor clustersize or TOT spectrum show any statistically significant deviation from the reference data.

Four modules were set to special TDAC configurations, to study the evolution of the noise occupancy with the reduction of TDACs. Figure 5.33 left shows the greyscale image of a mask with regions of continously decreasing TDAC settings, from TDAC+25 (very high threshold) to TDAC-25 (very low threshold). The two modules configured like this did not show an increase in the noise occupancy per pixel correlated with the TDAC settings.

Figure 5.33 right shows the greyscale image of a mask, that in the lower region again has a continous decrease in TDAC values, while in the upper right corner of the module all TDACs are decreased by 25 steps. The small black rectangle in the upper left region corresponds to a few pixels with TDACs also reduced by 25 steps. Figure 5.34 shows maps of the noise hits for the two modules with the TDAC settings shown in figure 5.33 right. The large regions, where the TDACs are reduced by 25 steps, clearly show an increased noise occupancy.

Figure 5.34: Maps of noise hits for the two modules with the TDAC settings described above.

5.5   Noise in cosmics data taking

The integral intensity of muons, generated by cosmic rays impinging on the earths atmosphere, with an energy above 1 GeV at sea level is $70 \, \mathrm{m}^{-2} \, \mathrm{s}^{-1} \, \mathrm{sr}^{-1}$. Upon detection of a coincidence between the top and bottom scintillators, described in chapter 4.5, 16 consecutive LVL1 trigger signals are sent to the modules. During the time the modules are sensitive, about 4.8×10^{-6} additional muons are expected to traverse any disk of the endcap. Therefore the occupancy per pixel for a single muon producing a hit in one of the 114 active modules is expected to be about 1.9×10^{-7}.


Figure 5.35: Left: Average occupancies per pixel for all active modules; Right: TOT spectrum for a high-occupancy module in run 1129.

As mentioned above, 16 consecutive LVL1 signals were sent for every scintillator coincidence. This was done to make setting the correct delays between the coincidence signal and the sending of the trigger signal to the

modules easier. In the ATLAS experiment, the delays will be adjusted very carefully, such that only one signal needs to be sent for every Level1 trigger. Hit data contains information about which LVL1 signal the hit is assigned to. Figure 5.36 shows the distribution of all hits in run 1129 over the 16 LVL1 signals. A distinct peak is visible at LVL1A=5. As noise hits do not have a fixed phase with respect to the scintillator coincidence, they are uniformly distributed across the LVL1A signals. The excess at LVL1A between 4 and 6 is therefore due to cosmic muons.


Figure 5.36: Distribution of hits within the 16 LVL1 signals sent for every scintillator coincidence.

Analyzing cosmics data taking run 1129, with about 743000 events in total, average occupancies per pixel from 2×10^{-7} up to 1.9×10^{-4} are measured (see figure 5.35). The TOT spectrum for one of the modules with high occupancy shows a tail towards high TOT values that is characteristic for noise hits.

Disregarding hits outside the peak in the LVL1A distribution leaves a very pure sample of cosmic muon hits. Comparing the occupancies before and after this cut, one finds ratios of noise hits to cosmic hits between 2.1 and 46.5.

To be able to compare the characteristics of the noise hits in the cosmics data taking run with the noise occupancy runs, where only one Level1 trigger signal was sent, a pure sample of noise hits is required. To have the same sensitivity to noise hits, only hits with LVL1A=10 are chosen.

Comparing the occupancies in this sample with random trigger run 1138 one finds that the ratio of noise occupancies in the cosmic run and in run 1138 is 0.39 ± 0.28. Taking into account all hits outside the peak in the LVL1A distribution (LVL1A<4 or >6), the ratio becomes 4.24 ± 3.07 (see figure 5.37). This ratio corresponds very well with ratio calculated for the

noise hits at LVL1A=10 times the increased number of LVL1A bins. Despite the trigger being generated by a physical event, the recorded hits are clearly dominated by noise hits for a train of 16 LVL1 signals. This shows that less LVL1 signals per trigger event are desirable.

The TOT spectra of the noise hits, as well as the clustersizes and number of hits per event, in the noise hit sample from the cosmic data taking run correspond very well to these of the random trigger run. No differences in the characteristics of the noise hits are observable.
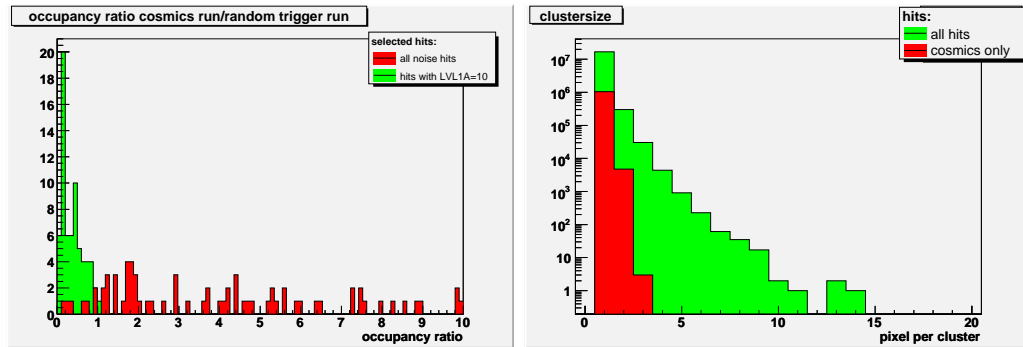

Figure 5.37: Left: Ratio of occupancies with cosmics trigger and with random trigger; Right: Clustersizes in the cosmics run before and after cutting on LVL1A=5.

5.6   Summary of most important results

In this chapter it was shown that, although the single-channel random noise is not perfectly described by a gaussian distribution, no high-noise tails can be observed. The minimum operating threshold, which is a good measure of the total noise in a module, was found to be about 2300 e$^-$ in a small scale system test setup, not using the final readout system. Unfortunately, due to software problems and time constraints, it could not be measured in the large scale system test.

The noise performance of the modules was tested by measuring the noise occupancy when triggering the modules with a fixed frequency. Decoding and analysis tools were developed to study the noise characteristics in these data taking runs. The measurements showed a noise occupancy below the sensitivity limit of the data taking runs of 6×10^{-8} for 99.98% of all pixels of the endcap. This measurement was repeated for decreased thresholds, showing that the noise occupancy rises significantly for thresholds lower than 2700 e$^-$. Below this threshold, coherent noise effects could be observed on

some modules. The distribution of these modules across the endcap pointed towards electromagnetic interference from the Type0 cables as the most probable source.

The noise occupancy was also studied in a data taking run triggering on cosmic muons. It was found that the higher number of Level1 trigger signals issued during this run, compared to the random trigger runs, where only one Level1 trigger was issued, increased the noise occupancy significantly. When analyzing only hits recorded in for one Level1 trigger, the noise occupancy agreed to the one measured in the random trigger runs.

run number	threshold setting	trigger frequency	total number of events
1138	tuned	~ 13kHz	16.8×10^6
1140	TDAC-10	~ 13kHz	16.8×10^6
1144	TDAC-20	~ 11kHz	16.8×10^6
1150	TDAC-25	5kHz	7×10^6
1065	tuned	100 Hz	29000
1021	tuned	30 kHz	6.1×10^6
1129	tuned	cosmics	7.4×10^5

Table 5.4: Overview of analyzed noise occupancy runs and settings used.

Chapter 6

Summary

The ATLAS pixel detector is the largest hybrid pixel detector system ever built. With its nearly 80 million pixels it provides more than half of all readout channels, that ATLAS has in total. A detector this big needs an intricate infrastructure of services and readout. Additional requirements arise from the close proximity of the detector to the interaction point, and the need to minimize multiple scattering by using little and lightweight material.

Although all components, needed to operate the pixel detector, are thoroughly tested during their production, the interplay of the different systems involved could only be tested in a large scale system test. This is where this thesis sets in. The test was done at CERN, starting in spring 2006 with a small scale system until in september 2006 a complete endcap of the pixel detector and enough service components to operate it were available. The upgraded system corresponded roughly to 10% of the final pixel detector.

Within the framework of this thesis, key characteristics of the system were measured and qualified for the final experiment. The results of measurements of the long-term stability of the power supply system, of the reliability of the monitoring and interlock systems and of the performance of the cooling system are presented in chapter 4. All systems perform as required or better and proved to be well suited for long-term operation of the pixel detector under realistic conditions.
Furthermore the analog performance of the detector modules in the system test setup was studied in terms of the distribution of the discriminator thresholds for all modules on the endcap, of threshold dispersions on individual modules and the calibration of the TOT information. Comparison of these quantities with measurements done during the production of the modules yielded no significant differences in the performance of the modules.

The system test culminated in the exercise of the full detector system, including readout and triggering systems, measuring cosmic muons passing through the endcap. This data was used to test offline analysis and alignment algorithms.

Controlling the noise level in the pixel detector was one of the major issues in the developement of the readout electronics and the design of the grounding and shielding scheme. It is very important to keep it as low as possible, because high noise can lead to inefficiencies due to disabled pixels. If the occupancy due to noise hits is high it can also cause data loss or degrade the tracking and vertexing efficiency due to fake hits. Therefore a profound knowledge of the noise performance of the pixel detector is necessary.

In the second part of this thesis the noise performance of the pixel detector electronics was studied. An excellent noise performance was found at the nominal operating threshold, without any trace of common mode noise effects. A detailed summary of the noise performance studies can be found in chapter 5.6.

Overall it was shown in this thesis, that the pixel detector performs well within the requirements and the systems are ready to be installed in the ATLAS experiment.

After the end of the system test, the setup was used to test the connections in the pixel detector, when its assembly was finished. After all connections were proven to be working correctly, the installation of the pixel detector in the heart of ATLAS was finished by the end of june 2007. The pixel detector group is now looking forward to the final comissioning phase and to a successfull operation of the pixel detector.

Appendix A

Details on the system test setup

A.1   Layout of the cooling system

The layout of the cooling system was quite different from the final system described in chapter 3.10. The SR1 cooling plant was designed to provide cooling for an assembly area and for a test area for the SCT, as well as for the pixel system test area. Inlet and outlet lines in the cooling plant split into the lines for the three distribution racks in the areas behind a main inlet valve controlled by the supervisory PLC. Thus the three areas could be operated independently, sharing only the available cooling fluid. Although the input pressure was set by the plant for all inlet lines, it could be regulated for each distribution rack by means of a manual pressure regulator. Each distribution rack had a separate back-pressure regulator (BPR) controlled by the PLC.
A single inlet line was used between the distribution rack and a manifold at PP1. There the fluid was distributed to the cooling pipes in the PSQP. The pipes in one half of the PSQP could be valved off. In front of the manifold an electrically controlled valve (solenoid valve) was installed as part of the interlock system for the cooling. The valve was set to close if the pressure on the exhaust side of the cooling circuits exceeded 3 bar(a), if the network connection to the PC running the SCADA system for the cooling plant was lost, or when the cooling plant entered an alarm state and shut down. When the valve was closed a signal was issued to the DSS input of the logic unit to cut switch off the power of the modules.
The exhaust lines were routed back to the distribution rack separately. Each exhaust line was equipped with an over-pressure relief valve set to open at a pressure of 4 bar(a). The circuit for the optoboard cooling was equipped with

a separate manual back-pressure regulator to adjust the optoboard temperature. The heater, which in the final system will be implemented on part of the return vapor line, in the system test setup was implemented as a heating bath in the distribution rack. The rack also held a sub-cooling unit, referred to as the *heat exchanger*, that cooled the inlet fluid by means of the cold vapor in the return line before the heater. This sub-cooling unit is omitted in the final system. It turned out to cause problems under special flow conditions. When the complete endcap was cooled, the pressure drop over the heat exchanger increased to about 0.8 bar. The back-pressure at the detector would then be close to 3 bar(a), which was the upper limit of the safe operation pressure. Therefore the heat-exchanger was removed.

A.1.1   Operation of the cooling system

No cooling operation was started with a dew point in the dry box above -30°C to prevent condensation on the cold exhaust pipes running along the side of the endcap and through the PSQP. Lowering the dew point in the main volume of the dry box to this value after an access took about 4h.
During operation the critical system parameters (liquid weight in the storage tank, input pressure and back-pressure in the cooling circuits) were constantly monitored, using the SCADA system running on one of the workstations in the control room. Also the inlet and exhaust pressures at the PP1 manifold were monitored via the pixel DCS project.

When operating the cooling system with very low coolant flows, oscillations in the input pressure of about 1 bar were observed, that were reflected in the liquid weight in the tank. They could be traced back to a regulation function of the system. Under normal operating conditions these oscillations could be observed neither in the back-pressure nor in the module temperatures. Therefore they do not influence the performance of the cooling system. They could be stopped by increasing the coolant flow, e.g. by opening a bypass valve on another distribution rack. When running at high back-pressure settings (e.g. the optoboard cooling loop), oscillations developed in the back-pressure. Because of the very small coolant flow, bubbles of coolant vapor tended to grow at constrictions in the tubes until the pressure was high enough to push the vapor through, resulting in a sudden drop of the pressure. These oscillations could not be stopped, because the flow in the respective circuit could not be increased. Thus the back-pressure in the optoboard cooling loop could be adjusted only very roughly.
The influence of heat dissipation of the detector on the parameters of the cooling system was analyzed and turned out to be negligible. The back-pressure

regulators quickly compensated any increase of flow due to increased coolant evaporation in the detector.

The cooling system proved to be very stable during operation and very seldom required intervention from the operators.

A.2 DCS and environmental control

The DCS project, used to monitor and control the detector system, was distributed among four dedicated PCs, referred to as dcs1 to dcs4. Different sub-projects were running on each machine, connected to the main project.

- dcs1
  This machine had a PEAK CANbus interface installed and was used to run the OPC server and client for the ISEG high voltage power supplies. It also ran the corresponding FIT project for the ISEGs, and later on the FSM project.

- dcs2
  This machine had a KVASER CANbus interface installed used to communicate to the WIENER low voltage power supplies and the ELMBs. It ran the OPC server/client for the WIENER supplies as well as the DCS main project. It was used to connect the service test setup and for monitoring during work inside the clean room (e.g. to monitor pressure readings while starting up the cooling system).

- dcs3
  This machine was also used to run the DCS main project, and it was the main user interface to all sub-projects during work in the control room. Furthermore it ran the OPC server/client for the ELMBs, and the RDB manager responsible for archiving the monitored values in an Oracle database. 1500 to 2000 values were constantly monitored and archived.

- dcs4
  This machine ran the project to monitor the latches of the LogicUnits and Interlock Distribution Boxes. It was mainly used to allow software development without interference with system test operation, as the monitoring of the LogicUnits and IDBs was seldom necessary.

These machines were set up in the rack area of the SR1 building. A workstation in the control room was used to connect to dcs3 via Windows

Figure A.1: Block diagram showing the distribution of the DCS project components over the PC described in the text.

Remote Desktop in the beginning. Later on a reduced version of the DCS main project was installed that was connected to the main project on dcs3 via the network.

The DCS system was also used to monitor various environmental parameters. A number of additional NTC thermistors were mounted along the service panel and in the PP0 region to monitor the temperature in the vicinity of the optoboards. A couple of humidity sensors were mounted throughout the dry box. Humidity sensors by Honeywell and Xeritron were mounted in pairs to qualify the Xeritron sensors, which at that time were supposed to be used in the final detector. Unfortunately neither type of sensor operated reliably at the humidities required in the dry box. The Xeritron sensors showed a fast response to changes in the humidity but they took time on the order of hours to settle at the final value of the relative humidity. Furthermore the measurements of absolute humidity were unreliable, at the humidity required for a dew point of -35°C, that was necessary to safely operate the cooling system. Therefore the monitoring of the dew point of the air inside the dry box was done using a very sensitive external hygroscope. Based on this experience, a new type of radiation hard humidity sensors has been chosen for the experiment.

The DCS system was also used to monitor input pressure and back-pressure

in the cooling pipes in the PP1 region and to control the electrical valve on the main cooling inlet. The electrical valve was operated by an ELMB. A software interlock was implemented to close the valve when the back-pressure exceeded 3.5 bar(a). Furthermore the status of the valve was used as an input to the interlock system to power down the detector in case the valve was closed.

Bibliography

[1] J. Ellis. Particle Physics at future colliders. *arXiv:hep-ex/0210052v1*, 2002.

[2] http://en.wikipedia.org/wiki/Image:Interactions.png.

[3] J. Ellis. Limits of the Standard Model. *hep-ph/0211168*, 2002.

[4] CERN Document Server: Photos. *http://cdsweb.cern.ch/collection/Photos*.

[5] D. Boussard and T. Linnecar. The LHC superconducting RF system. Technical Report LHC Project Report 316, 1999.

[6] ATLAS collaboration. ATLAS, Letter of Intent for a General-Purpose pp Experiment at the Large Hadron Collider at CERN. *CERN-LHCC-92-4*, 1992.

[7] J. Pequenao. http://pdgusers.lbl.gov/ pequenao/.

[8] ATLAS Inner Detector community. Inner Detector Technical Design Report, Volume 1. *CERN-LHCC-97-016*, 1997.

[9] V.A. Mitsuo. Readout Electronics Test and Integration of the ATLAS Semiconductor Tracker. *physics/0610114*, 2007.

[10] P. Bell. ATLAS SCT End-Cap Integration. *Talk at the particle physics seminar, Bonn University*, 2007.

[11] C. Rembser. The ATLAS Transition Radiation Tracker. *Talk at the particle physics seminar, Bonn University*, 2007.

[12] ATLAS collaboration. Calorimeter Performance: Technical Design Report. *CERN-LHCC-96-040*, 1997.

[13] M. Aharrouche et al. Energy Linearity and Resolution of the ATLAS Electromagnetic Barrel Calorimeter in an Electron Test-Beam. *Nucl. Instr. and Meth. in Phys. Res. A*, 568:601–623, 2006.

[14] S. Baffioni et al. Electrical measurements on the ATLAS electromagnetic barrel calorimeter. *ATL-LARG-PUB-2007-005*, 2007.

[15] G. Sauvage. The ATLAS electromagnmetic calorimeter: construction, commisioning and selected test beam results. *Nucl. Phys. B*, 150:98 – 101, 2006.

[16] H.H.J ten Kate. ATLAS Superconducting Toroids and Solenoid. *IEEE Transactions on Applied Superconductivity*, 15, 2005.

[17] S. Palestini. The Muon Spectrometer of the ATLAS Experiment. *Nucl. Phys. Proc. Suppl.*, 125:337–345, 2003.

[18] The ATLAS collaboration. ATLAS High-Level Trigger, Data Acquisition and Controls - Technical Design Report. *CERN-LHCC-2003-022*, 2003.

[19] B. Gorini et al. The ATLAS Data Acquisition and High-Level Trigger: Concept, Design and Status. *ATL-DAQ-CONF-2006-016*, 2006.

[20] J. Haller. The first-level trigger of ATLAS. *International Europhysics Conference on High Energy Physics*, PoS(HEP2005)391.

[21] G. Alimonti et al. Test-beam measurement of the detection efficiency of ATLAS Pixel production modules. *ATL-INDET-INT-2005-007*, 2005.

[22] Particle Data Group. Review of Particle Physics. *Journal of Physics G: Nuclear and Particle Physics*, 33, 2006.

[23] H. Bichsel. Straggling in thin silicon detectors. *Rev. Mod. Phys.*, 60:663 – 699, 1988.

[24] F. Huegging. *Der ATLAS Pixelsensor*. PhD thesis, Dortmund University, https://eldorado.uni-dortmund.de/handle/2003/2350, 2001.

[25] G. Lindström et al. Radiation hard silicon detectors - developments by the RD48 (ROSE) collaboration. *Nucl. Inst. and Meth. in Phys. Res. A*, 466:308 – 326, 2001.

[26] RD48. The ROSE collaboration. *CERN / LHCC / 96-23*, 1996.

[27] T. Lari. Radiation hardness studies of silicon pixel detectors. *Nucl. Instr. and Meth. in Phys. Res. A*, 560:93–97, 2006.

[28] L. Blanquart, J. Richardson, K. Einsweiler, P. Fischer, E. Mandelli, G.Meddeler, and I. Peric. FE-I2: A Front-End Readout Chip Designed in a Commercial 0.25-μm process for the ATLAS Pixel Detector at LHC. *IEEE Transactions on Nuclear Science*, 51(4):1358–1364, August 2004.

[29] G. Darbo et al. MCC: the Module Controller Chip for the ATLAS Pixel Detector. *Nucl. Instr. and Meth. in Phys. Res. A*, 492:117–133, 2002.

[30] K.E. Arms et al. ATLAS pixel opto-electronics. *Nucl. Instr. and Meth. in Phys. Res. A*, 554:458–68, 2005.

[31] T. Flick. *Studies on the Optical Readout for the ATLAS Pixel Detector.* PhD thesis, Wuppertal University, http://nbn-resolving.de/urn/resolver.pl?urn=urn%3Anbn%3Ade%3Ahbz%3A468-20060600, 2006.

[32] M.L. Chu et al. The off-detector opto-electronics for the optical link of the ATLAS SemiconductorTracker and Pixel Detector. *Nucl. Instr. and Meth. in Phys. Res. A*, 530:293 – 310, 2004.

[33] H.C. van der Bij, R.A. McLaren, O. Boyle, and G. Rubin. S-LINK, a Data Link Interface Specification for the LHC era. *IEEE Transactions on Nuclear Science*, 44:398 – 402, 1997.

[34] T. Vickey. A Read-out Driver for Silicon Detectors in ATLAS. *Proceedings of the 12th Workshop on Electronics for LHC and Future Experiments (LECC 2006), Valencia, Spain.*

[35] J. Christiansen, A. Marchioro, P. Moreira, and T. Toifl. TTCrx Reference Manual: A Timing, Trigger and Control Distribution Receiver ASIC for LHC Detectors. *http://www.cern.ch/TTC/TTCrx_manual3.9.pdf.*

[36] E. Anderssen et al. ATLAS PIXEL PATCH PANEL 1. Technical Report ATL-IP-ES-0072, 2004.

[37] M. Citterio and C. Meroni. Pixel PP2 Patch Panels design. Technical Report ATL-IP-ES-0074, 2003.

[38] S. Kersten and P. Kind. Pixel Patch Panel PP3 Design. Technical Report ATL-IP-ES-0075, 2005.

[39] Embedded Local Monitoring Board ELMB128. Technical Report http://elmb.web.cern.ch/ELMB/elmb128_1.pdf.

[40] S. Kersten and P. Kind. The LV-PP4 of the ATLAS pixel detector. Technical Report ATL-IP-CD-0002, 2007.

[41] Henry W. Ott. *Noise Reduction Techniques In Electronic Systems.* John Wiley & Sons, 1976.

[42] S. Kersten. The Interlock Matrix of the Pixel Detector Control System. Technical Report ATL-IP-ES-0110, 2005.

[43] S. Kersten and P. Kind. Technical Description of the Interlock Circuit and System of the ATLAS Pixel Detector. Technical Report ATL-IP-ES-0041, 2004.

[44] K.-W. Glitza et al. Local Support Assembly of the ATLAS Pixel Detector. *Nucl. Instr. and Meth. in Phys. Res. A*, 565:102–105, 2006.

[45] M. Stodulski M. Olcese. Inner detector evaporative cooling system: Requirements and general design issues. *ATL-IC-ES-0006*, 2004.

[46] B. Gorski, R.L. Bates, M. Olcese, and Q. Vacek. The Design and Prototyping of the ATLAS Inner Detector Evaporative Cooling System. *IEEE Nuclear Science Symposium Conference Record*, 2:677 – 681, 2004.

[47] M. Leyton and M. Keil. Service Tests for the ATLAS Pixel Detector. Technical Report ATL-IP-QA-0026, 2007.

[48] D. Dobos. *First experience with the ATLAS Pixel Detector.* PhD thesis, University of Dortmund, in preparation.

[49] The ATLAS Pixel System Test Group. Pixel Offline Analysis for End-capA Cosimc Data. *ATLAS Note in preparation*, 2007.

[50] I. Perić. *Design and Realization of Integrated Circuits for the Readout of Pixel Sensors in High-Energy Physics and Biomedical Imaging.* PhD thesis, Bonn University, 2004.

[51] ATLAS Pixel Detector Collaboration. ATLAS pixel detector: Technical Design Report. *CERN-LHCC 98-013*, 1998.

[52] J. Weingarten and K. Reeves et al. System Tests of the ATLAS Pixel Detector. *Proceedings of the 11th Workshop on Electronics for LHC and Future Experiments (LECC 2005), Heidelberg, Germany*, 2005.

[53] I. Peric. private communication.

[54] P.J. Bell and J. Grosse-Knetter. Techniques for the Measurement of Common Mode Noise in Binary Read Out Systems. *Nucl. Instr. and Meth. in Phys. Res. A*, 516:153–159, 2004.

[55] L. Feld et al. Measurement of common mode noise in binary read-out systems. *Nucl. Instr. and Meth. in Phys. Res. A*, 487:557, 2002.

[56] D.B. Ta et al. Serial Powering: Proof of principle demonstration of a scheme for the operation of a large pixel detector at LHC. *Nucl. Instr. and Meth. in Phys. Res. A*, 557:445–459, 2006.

Acknowledgements

Sorry for the disaster. And thanks for your patience!
Chris Thyberg

This thesis could not have been completed without the generous support by many people during the years and I want to use this opportunity to express my gratitude to all of them.

First and foremost I want to thank my family for their lasting patience. Without your support I couldn't have done this. As simple as that...

I also want to thank Prof. Norbert Wermes for giving me the chance to go to CERN to work on this thesis.

I enjoyed working in the SiLab as much as one can enjoy work at all. Thank you guys for the great atmosphere and the many discussions about anything, including even physics. Many thanks to Robert, Lars, Johannes, Edgar, Manuel, Duc, Michael, Hans and all the rest of you guys, for all the fun, the shared misery in the mensa and last, but not least, all the shared coffee.
Special thanks to Markus Mathes, whom I bothered with lots of questions on ROOT, C++ and lately LaTex, and to Jörn Grosse-Knetter, who suffered the pain of proof-reading the whole of this thesis. Thanks for your patience.
I also want to thank Walter Ockenfels and Wolfgang Dietsche for their help with any kind of mechanical, bonding or glueing problems.
There is quite a long list of former colleagues, I owe gratitude to. It includes Fabian Hügging, Tobias Stockmanns, Lasse Klingbeil, Ivan Peric, Jaap Velthuis, and a whole lot more people. I want to thank all of you.

I got to know a number of great people at CERN, all of which made work, and what little life was left, a very interesting experience and a lot of fun. I want to thank Kevin, Tom, Marian, Maurice, Claudia, Götz, Lucia, Iskander, Tobias, Jean-Francois, David, Kendall and Georg. I want to thank

So long, and thanks for all the fish.
Douglas Adams, The Hitchhikers Guide to the Galaxy.