FACHBEREICH MATHEMATIK UND NATURWISSENSCHAFTEN
FACHGRUPPE PHYSIK
BERGISCHE UNIVERSITÄT WUPPERTAL

The Finite State Machine for the ATLAS Pixel Detector

and

Beam Background Studies with the ATLAS Beam Conditions Monitor

Dissertation zur Erlangung des Doktorgrades
vorgelegt von

Kerstin Lantzsch

May 2011

Die Dissertation kann wie folgt zitiert werden:

urn:nbn:de:hbz:468-20130206-154031-3
[http://nbn-resolving.de/urn/resolver.pl?urn=urn%3Anbn%3Ade%3Ahbz%3A468-20130206-154031-3]

Contents

Einleitung

Der Large Hadron Collider am CERN bei Genf hat Ende 2009 mit den ersten Proton-Proton Kollisionen seinen Betrieb aufgenommen. Der ATLAS-Detektor ist einer der Universaldetektoren, und wurde entwickelt und gebaut um neue physikalische Phänomene zu studieren und das Wissen und Verständnis der zugrundeliegenden Theorie zu verbessern. Der Pixeldetektor ist der innerste Subdetektor von ATLAS und wird mit seinen 80 Millionen Auslesekanälen eine wichtige Rolle spielen bei der Identifikation und Vermessung der interessanten physikalischen Phänomene. Komplexe Kontrollsysteme sind erforderlich, um ein System wie den Pixeldetektor sicher betreiben zu können. Letztendlich wird ein einziger Operator für die Kontrolle und Überwachung des gesamten ATLAS-Detektors zuständig sein.

Diese Arbeit besteht aus drei Teilen. Der erste Teil gibt einen Überblick über das Experiment. Kapitel 1 gibt mit einer kurzen Zusammenfassung des Standardmodells der Teilchenphysik eine Motivation zum Bau großer Beschleuniger und Detektoren. In Kapitel 2 wird der Beschleuniger beschrieben, und der ATLAS-Detektor mit seinen Subsystemen wird vorgestellt. Da der Pixeldetektor und der Beam Conditions Monitor von spezieller Bedeutung für diese Arbeit sind, werden sie in größerem Detail in Kapitel 3 und 4 beschrieben.

Der Schwerpunkt dieser Arbeit liegt im zweiten Teil, der das Kontrollsystem des ATLAS Pixeldetektors behandelt. In Kapitel 5 werden die Hardware- und Software-Komponenten beschrieben. Kapitel 6 und 7 stellen den Hauptteil der zugrundeliegenden Arbeit dar. Der Betrieb der ATLAS Subdetektoren wird hauptsächlich durch die Zustandsmaschine (FSM) gewährleistet. Kapitel 6 beginnt mit einer Einführung der zugrundeliegenden Software-Komponenten. Die Anforderungen für die Pixel-FSM werden dargelegt, und die Pixel-spezifische Implementierung diskutiert. In Kapitel 7 werden Skripte vorgestellt, die ergänzend zur FSM entwickelt wurden, um sicherheitsrelevante Aspekte abzusichern, sowie die Automatisierung von Prozeduren zur Interaktion mit externen Systemen handhaben.

Der letzte Teil beschäftigt sich mit Untergrundereignissen aufgrund von strahlbezogenen Belangen. Kapitel 8 beschreibt potentielle Szenarios zu Strahlverlust, sowie Untergrund während des regulären Betriebs. Während der frühen Betriebs-Phase mit wenigen Protonen-Paketen und einer Strahlenergie von bis zu 3500 GeV wurden Daten vom Inneren Detektor und des Beam Condition Monitors korreliert.

Eine Zusammenfassung und Bewertung dieser Arbeit im Hinblick auf den Betrieb des ATLAS Pixeldetektors wird im Abschluß dieser Arbeit gegeben.

Introduction

The Large Hadron Collider at CERN near Geneva has started operation with the first proton-proton collisions in the end of 2009. The ATLAS detector is one of the general purpose experiments designed to take full advantage of the possibilities to study new physics phenomena and improve the knowledge and understanding of the currently known state of theory. The Pixel Detector is the innermost sub-detector of ATLAS and with its 80 million readout channels will play an important role in the identification and measurement of the interesting physics phenomena. To operate a complex detector like the Pixel Detector in a safe and coherent way, sophisticated control systems are necessary. In the end, the control and supervision of the Pixel Detector has to be handled by a single shift operator integrated into the ATLAS control system and in coordination with the data acquisition and the accelerator.

This thesis is divided into three parts. In the first part an overview of the experiment is given.
In chapter 1 the motivation for building large accelerators and detectors is given with a short summary of the Standard Model of particle physics and the phenomena to be studied at the LHC. In chapter 2 the accelerator is described, and the ATLAS detector and its subsystems are introduced. As the Pixel Detector and the Beam Conditions Monitor are of special relevance to this thesis, they are described in more detail in chapters 3 and 4.

The focus of this thesis lies in the second part, covering the Control System of the ATLAS Pixel Detector.
In chapter 5 the hardware and software components are described. Chapter 6 and 7 constitute the main part of the performed work. The Finite State Machine (FSM) is the main tool for the operation of the ATLAS sub-detectors. Chapter 6 starts with an introduction to the underlying (common) software components and prior work. The requirements for the Pixel FSM originating from various directions are stated, and the pixel specific FSM implementation is discussed. Complementary to the FSM, background scripts have been developed to address safety and automatization aspects. The first part of chapter 7 presents the implemented safety actions to protect the Pixel Detector from danger, while the second part deals with automatization of complex interactions with external systems.

Finally, the third part treats beam background related issues.
Chapter 8 describes background related to potential beam loss accidents and accident scenarios, as well as background during regular operation. During the early running with few bunches and beam energies of up to 3500 GeV data from the inner detector and the beam conditions monitor has been correlated for different energies and beam optics.

Concluding this thesis is a summary and evaluation of this work in the larger context of the operation of the ATLAS Pixel Detector. Achieved results are presented and an outlook for possible future developments and work is given.

Chapter 1

Physics at the LHC

Particle physics deals with the basic building blocks of matter and the forces that act between them. Four fundamental interactions are observed in nature – the electromagnetic, weak and strong force, and gravity. The Standard Model describes the elemental particles and three of the four forces. It does not describe gravity which is many orders of magnitude weaker than the other forces, and therefore does not play a role in processes of high energy physics.

1.1 The Standard Model

In the Standard Model [1] [2], all elemental particles can be divided into three generations of leptons and quark and the gauge bosons which transmit the forces between them. The interactions are described by gauge theories based on symmetry groups. Table 1.1 lists the leptons and quarks and some of their properties. The corresponding anti-particle has different sign for the inner quantum numbers, like e.g. charge.

Leptons (Spin $\frac{1}{2}$)			
	1st Generation	2nd Generation	3rd Generation
Name	e	μ	τ
Charge [e]	-1	-1	-1
Mass	0.511 MeV/c^2	106 MeV/c^2	1777 MeV/c^2
Name	ν_e	ν_μ	ν_τ
Charge [e]	0	0	0
Mass	m($\bar{\nu}_e$) < 2 eV/c^2	< 0.19 MeV/c^2	< 18.2 MeV/c^2

Quarks (Spin $\frac{1}{2}$)			
	1st Generation	2nd Generation	3rd Generation
Name	u	c	t
Charge [e]	$+\frac{2}{3}$	$+\frac{2}{3}$	$+\frac{2}{3}$
Mass	1.5 - 3.3 MeV/c^2	1.27 GeV/c^2	171.2 GeV/c^2
Name	d	s	b
Charge [e]	$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$
Mass	3.5 - 6.0 MeV/c^2	104 MeV/c^2	4.20 GeV/c^2

Table 1.1: *The three Generations of Leptons and Quarks [3].*

Particles of the three generations are distinguished through their masses: the mass of the muon is about 200 times that of the electron, and the tau is more than 3000 times heavier. The mass of the neutrinos is extremely small, and they have no electric charge, so that they only interact through weak interaction. Quarks and leptons are fermions with spin 1/2, and they obey the Pauli principle.

Table 1.2 lists the four interactions and the corresponding force carriers. Bosons have integer spin, and do not obey the Pauli principle.

Bound states of three quarks or quark and anti-quark are classified as baryons and mesons. They are subject to strong interaction, and both fall into the group of the hadrons.

The Standard Model comprises the field theories of electromagnetic, strong and weak interaction, where electromagnetic and weak interaction were unified in the electroweak theory.

Interactions

Interaction	el.-magn.	strong	weak	Gravity
boson	Photon	Gluon	$W^{\pm}, Z^0$	Graviton ?
affects				
Quarks	√	√	√	√
charged Leptons	√		√	√
Neutrinos			√	√

gauge bosons (Spin 1)

interaction	el.magn.	weak		strong
gauge boson	Photon	$W^{\pm}$	Z^0	Gluon
charge [e]	0	±1	0	0
mass [GeV/c^2]	0	80.398	91.1876	0

Table 1.2: *Fundamental Interactions and the gauge bosons in the Standard Model [3].*

1.1.1 Quantum Chromo Dynamics

The strong interaction is described by Quantum Chromo Dynamics (QCD). It is based on SU(3) symmetry group, and is a non-abelian field theory. The force carriers are 8 gluons, and the gluon couples to the "color charge" – "red", "green" or "blue" – of quarks and gluons. Gluons can directly couple as they carry color. No free quarks can be observed, but only color singlets ("confinement").

1.1.2 Electroweak Theory

Electromagnetic and weak interactions are unified by the electroweak theory based on SU(2)xU(1) symmetry group. The gauge bosons are the photon and the heavy W- and Z-bosons. Unification requires differentiation between left- and right-handed fermions, and the weak current only couples to left-handed states.

1.1.2.1 The Higgs Mechanism

While in electroweak theory the massless isospin triplet W_μ^1, W_μ^2, W_μ^3, and the singlet B_μ are gauge bosons, in reality W and Z have mass. For this reason the mechanism of "spontaneous symmetry breaking" is introduced, where the introduction of a scalar Higgs doublet gives the real gauge bosons as linear combinations from initial massless gauge bosons connected by the Weinberg angle θ_W. Three bosons (identified as W and Z) acquire masses proportional to the coupling to the Higgs field, while one (identified as the photon) stays massless. The remaining neutral scalar "Higgs" boson H^0 has not yet been observed.

1.2 Physics at the LHC

Figure 1.1 shows cross sections for pp interactions.

In pp interactions, what is referred to as the "underlying event" is the effect from the partons not involved in the hard scattering. There is some influence on the hard interaction, due to conversion of quantum numbers.

The small relative cross sections of interesting processes versus total cross sections leads to the requirement of high luminosity and good trigger and selection mechanism.
Interactions from separate protons in the same bunch as the hard interaction are referred to as "pileup".

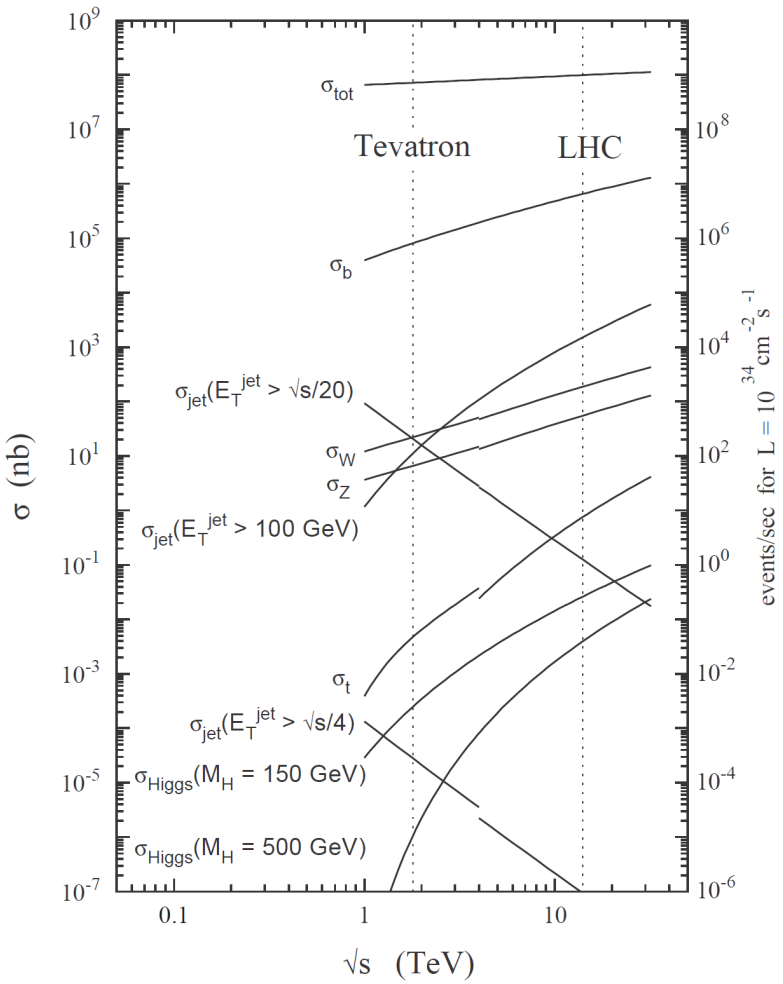

Figure 1.1: *pp cross-section [5].*

The first thing the LHC provides is the possibility to further study already known processes [4]. One motivation is that these processes are background to interesting processes, and therefore must be well understood, and on the other side deviations from expectation might give hints to new physics phenomena. The measurement program concerning the Standard Model will include Minimum Bias interactions, as charged multiplicity distribution, or measurement of W and Z properties. Well known properties like the Z mass or width can serve also for detector calibration.

A large number of top quark events will be produced at the LHC, which will make it possible to measure its properties with greater precision.

One of the main goals is the search for the Higgs boson. Main production mechanisms are shown in figure 1.2. Figures 1.3 (a) and (b) show respectively cross sections for production and branching ratios for the different decay mechanisms.

Figure 1.4 shows the discovery potential for some channels.

Finally, the goals include also the search for new physics, or the possibility to exclude it, including phenomena like super-symmetry, or extra dimensions.

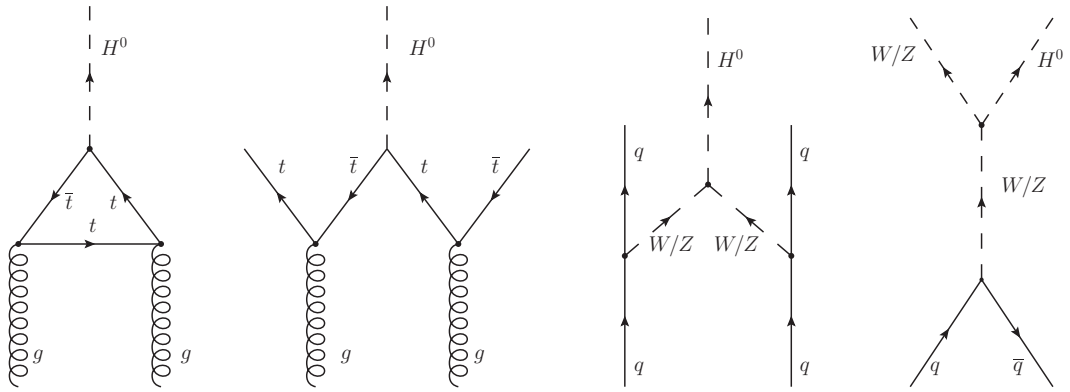

Figure 1.2: *Higgs Production Mechanisms: Gluon-Gluon Fusion, associated production with a $t\bar{t}$ quark pair, Vector-Boson Fusion, Higgs-Strahlung.*

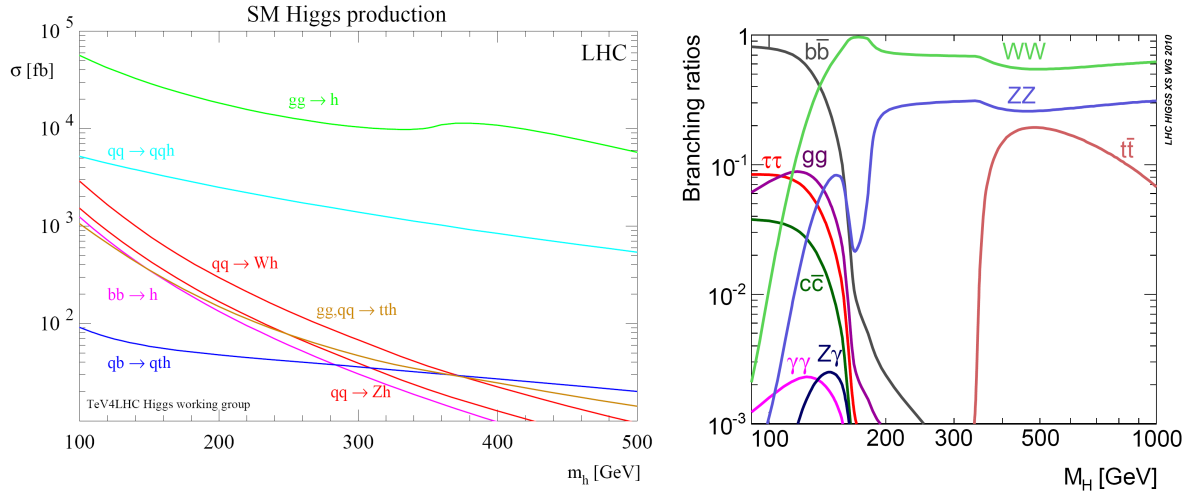

(a) (b)

Figure 1.3: *(a) Cross-section for Higgs production [6], and (b) Branching Ratios for Higgs Decay Mechanisms [7].*

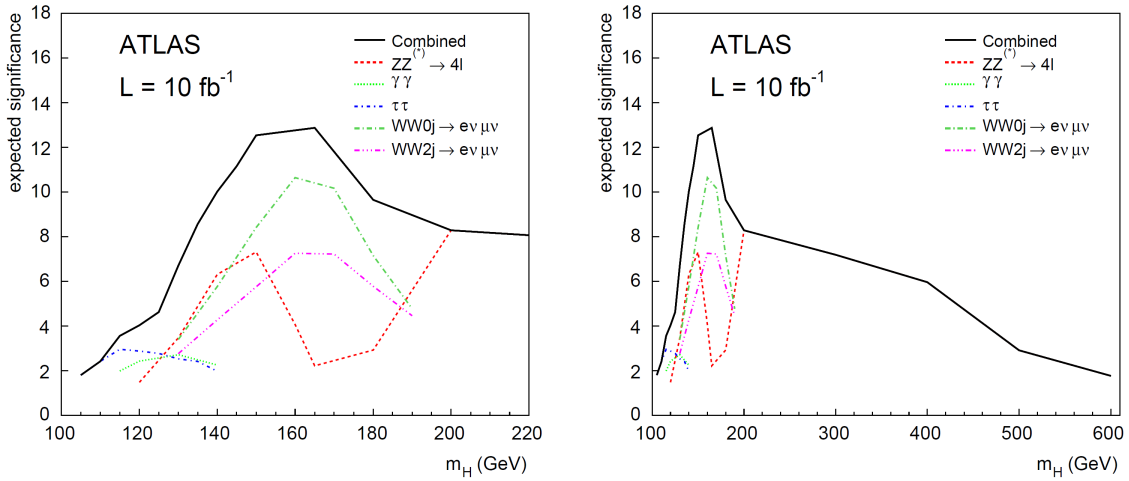

Figure 1.4: *Higgs Discovery Potential for selected channels for 14 TeV [8].*

Chapter 2

LHC and ATLAS

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) accelerator was put into operation 2008 at CERN[1)] located on the Swiss French border near Geneva. It is the biggest accelerator currently in operation. It has a circumference of 26.7 km and is located in the old LEP[2)] tunnel between 50-175 m below the surface. The design center of mass energy is 14 TeV [9].

2.1.1 Overview

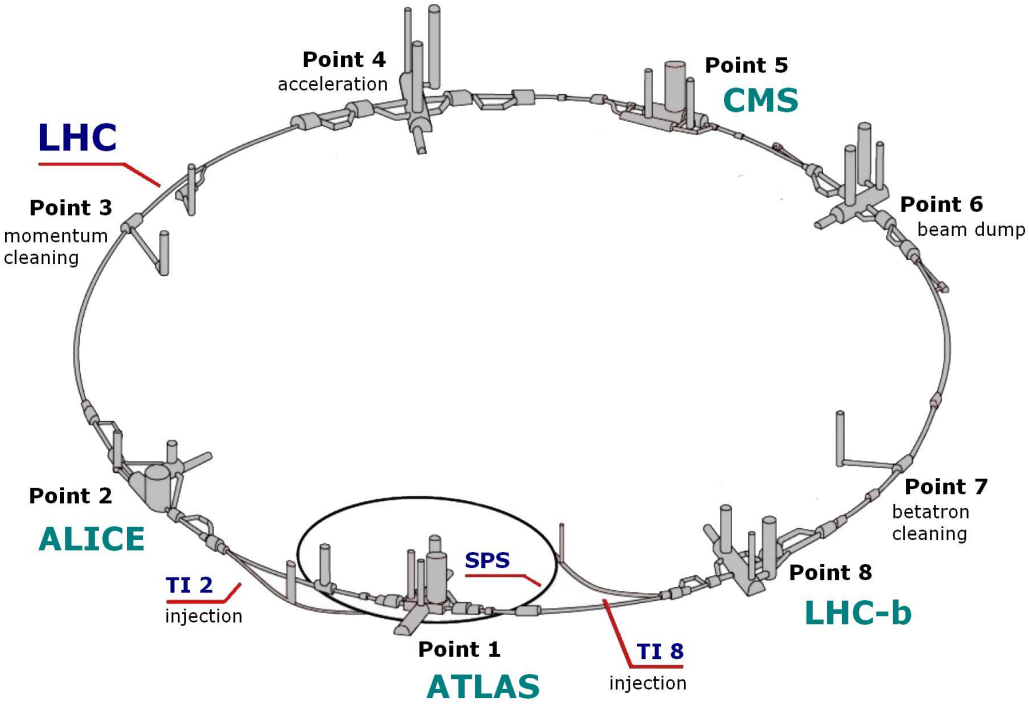

Figure 2.1: *Overview of the LHC complex.*

[1)] **C**onseil **E**uropéen pour la **R**echerche **N**ucléaire
[2)] **L**arge **E**lectron **P**ositron Collider

An overview of the LHC accelerator complex is shown in figure 2.1. The beams are injected from the SPS into the LHC through the injection lines TI 2 (“beam 1”) and TI 8 (“beam 2”). The injection energy is 450 GeV. Superconducting magnets are forcing the protons on their orbit. The necessary field is 8.3 T. Due to the fact that protons are accelerated in each ring, two opposite magnetic fields are needed, which are generated by 1232 dipole magnets. 386 quadrupoles are focusing the beam. sextupole-, octupole-, and decapole-magnets are used to correct field imperfections. The beams are accelerated by two independent RF systems, one for each beam, located at Point 4. Collimators are installed at Point 3 and Point 7 for momentum and betatron cleaning. Protons will collide at four interaction points where the experiments ATLAS, ALICE, CMS and LHCb are located. ATLAS is located at Point 1, and on the opposite side of the ring at Point 5 CMS. LHCb and ALICE are located at Point 7 and Point 2. At the end of each fill, or in case of bad beam conditions, the beam is extracted from the LHC at Point 6. For the beam to be extracted safely, it needs to be deflected from its course which is done by kicker magnets. These have to fire during the so called abort gap, i.e. the part of the beam where there are no protons. The protons are defocused and steered into a graphite block beam dump where their energy is absorbed. The requirement to build up the kicker field is what determines the maximum possible number of 2808 filled bunches out of 3564 positions, with a bunch spacing of 25 ns.

The first critical operation where sensitive hardware could be damaged is during injection. Therefore, a low intensity pilot bunch is injected first, to verify that all the settings of the magnets and collimators are correct. The intensity of 10^9 particles of this pilot bunch is low enough not to harm any equipment or cause the quench of magnets in case of beam loss. Beam Loss Monitors of the accelerator and the experiments are constantly monitoring the beam, and can trigger a beam dump.

2.1.1.1 Nominal/Peak Operation Conditions

The LHC can accelerate protons to the energy of 7 TeV. At nominal operation condition it will collide up to 2808 proton bunches with $1.15 \cdot 10^{11}$ protons in each bunch. This corresponds to a beam current of 0.582 A and a stored energy of 362 MJ. The actual design filling scheme is shown in figure 2.2. The bunch length is 7.55 cm which corresponds to a time of 0.25 ns. Bunches will be squeezed to a beam profile of 16.7 µm The peak luminosity will be 10^{-34} $\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The luminosity L is given by

$$L = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \varepsilon_n \beta^*} F \tag{2.1}$$

with
n_b: number of bunches per beam
N_b number of particles per bunch
f_{ref}: revolution frequency
γ_r: relativistic gamma factor
ϵ_n: normalized transverse beam emittance
β^*: beta function at the collision point
F: geometric luminosity reduction factor due to the crossing angle at the interaction point

Besides proton-proton collisions, also Heavy Ion (Pb) collisions are performed.

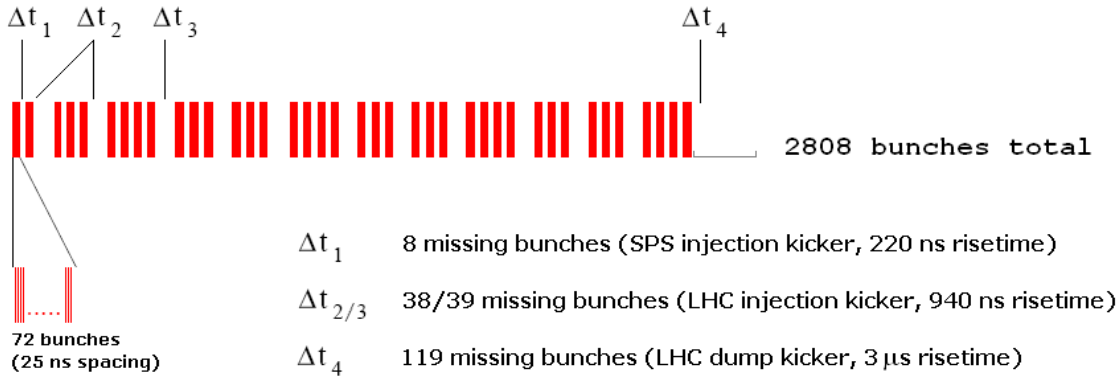

Figure 2.2: *Nominal Bunch Scheme (from [11], [12]). 39 · 72 = 2808 bunches are filled, with varying gaps in between. The 119 bunch "abort gap" of 3 μs needs to be kept free for the dumping of the beam. The 8 missing bunches are caused by SPS constraints, the 38/39 missing bunches by LHC injection constraints.*

2.1.2 Operation

2.1.2.1 Actual Operation Conditions

In the beginning of 2010, the LHC is operated with a beam energy of 3500 GeV. Instead of the nominal bunch size, bunches are squeezed to $\beta^* = 2$ m or 3.5 m. Not all possible bunch positions are filled, but only certain fill patterns. Bunches will collide in ATLAS and CMS if the bunch positions, as denoted by the BCID (the Bunch Crossing IDentifiers), in beam 1 and beam 2 are identical. They will collide in ALICE if they differ by 891, and in LHCb by -894. Parasitic collisions are possible if bunches meet within a given distance of the interaction point, since there is a common beam-pipe section for beam 1 and beam 2 [11]. This should be avoided by applying a (half) crossing angle between the beams of 142.5 μrad.

2.1.3 Absorbers and Collimators

The collimation system is responsible for beam cleaning and machine protection, and in the worst case to absorb high energy deposition in case of beam loss into the material. Figure 2.3 shows the main collimators placed in the LHC ring.

Betatron cleaning refers to the cleaning of the transverse beam extension, momentum cleaning to the longitudinal losses due to off-momentum particles. Specialized collimators are used for magnet protection from beam background and back-splash.

For the cleaning system in sectors 3 and 7, a three stage collimation scheme is used, illustrated in figure 2.4 for sector 7 for injection and collision setup. The collimator opening or "half-gap" is given in units of σ allowing to easily correlate independent from beam settings the collimator distance and beam profile . The "jaws" of the primary collimator ("TCP", Target Collimator Primary) are placed between the beam and the machine aperture to intercept halo. Secondary halo, i.e. particles scattered from the primary collimators, have to be intercepted by secondary collimators ("TCSG", Target Collimator Secondary Graphite) which must be always opened wider than the primary collimators. Tertiary halo must be minimal enough not to cause quenching of the magnets, or further absorber and protection mechanisms must be installed. For the betatron and momentum cleaning regions, absorbers ("TCLA", Target Collimator Long Absorber) are installed after the secondary collimators, and tertiary collimators

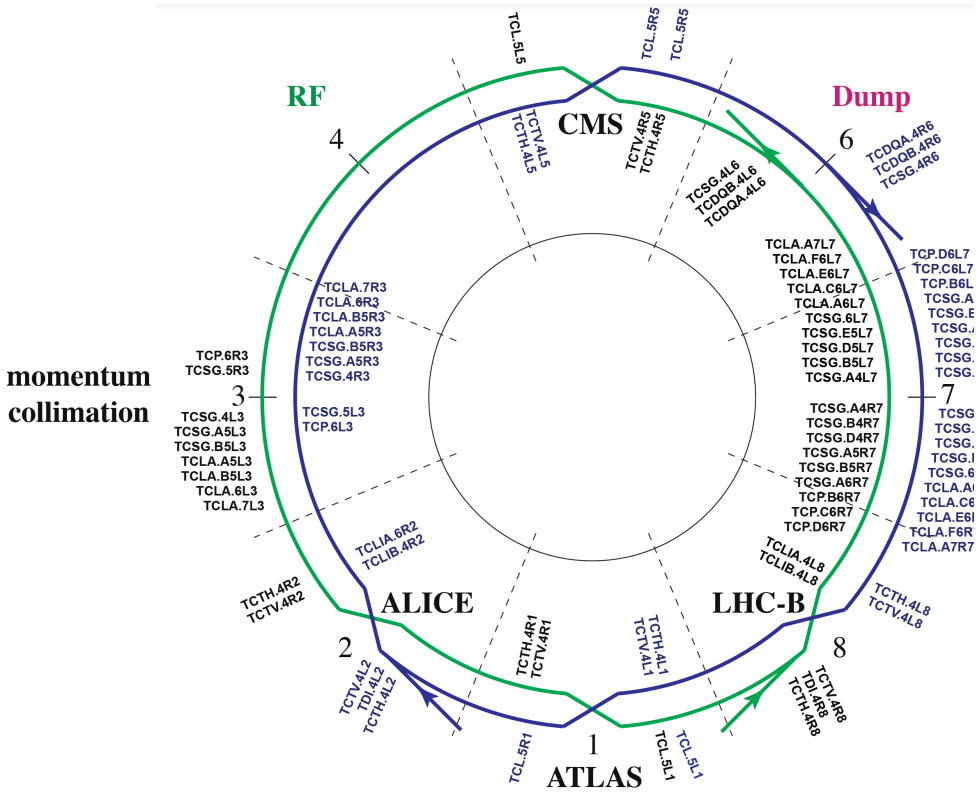

Figure 2.3: *The LHC Collimation System [21].*

("TCT", Target Collimator Tertiary) are installed near the interaction points to protect the superconducting magnet triplets for the focusing of the beam and the experiments.

Most collimators have an active length of 1 m, and a width and depth of 80x25 mm^2. Minimal gap openings correspond to sub-millimeter distances. For the material reinforced carbon fiber is used where the priority lies on the robustness (primary and secondary collimators), while high Z material (Cu/W) is chosen for good absorption qualities [13].

2.1.4 Beam Loss Monitors

Around the ring in the regions most exposed to the risk of beam losses, about 4000 Beam Loss Monitors are installed as part of the protection system. There are four types with different purposes [17]:

1. BLMCs are installed after collimators and cleaning insertions and should detect losses due to beam manipulation, unclean dumps or magnet powering failures. They have a time resolution of about an orbit.

2. BLMS should detect localized losses due to collision products and are installed near particular aperture limits. The time resolution is one orbit.

3. BLMA monitor the losses in the arc where they are located near the quadrupoles. The minimum integration time is 2.5 ms, since these losses are normally slower.

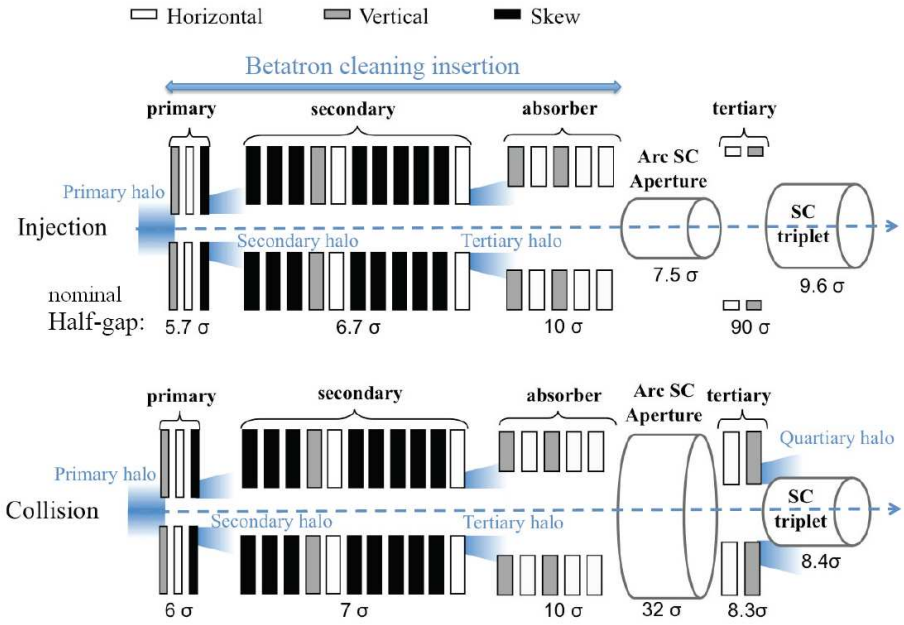

Figure 2.4: *The three stage Collimation System for the betatron cleaning [13]. The nominal collimator settings are indicated.*

4. The BLMB are used mainly for beam studies, not for the protection system. They are located after the primary collimators.

The BLMs must at least cover a dynamic range from pilot beam intensity to levels where magnet quenches are possible.

Further systems are the Beam Current Monitors, which measures the beam lifetime, and the Beam Position Monitors' task is to detect fast changes in the closed orbit. In case of bad condition, they can trigger a beam dump.

For normal conducting magnets with small time constants where in case of powering failure very fast beam losses would occur that could not be detected in time by the BLM system, a system of Fast Magnet Current Change Monitors (FMCM) has been installed [19].

2.1.5 Beam Interlock System

17 beam interlock controllers (BIC) are installed in the LHC to collect beam permit from the experiments and the LHC protection system. The beam permit is indicated by the presence of a signal which is transmitted in two fibres around the ring (for redundancy, and with different directions to minimize the time). The experiments and the LHC protection system can signal to the controllers the loss of beam permit, in the following the signal is removed by opening the loop, which will signal to the Beam Dump System in IR6 to dump the beam.

2.1.6 LHC Beam Dumping System

In the case of a required beam dump either at the end of a physics run, or due to the loss of beam permit, the two beams are extracted by the LHC beam dumping system (LBDS) located at Point 6. The LBDS region is shown in figure 2.5. The dump should only be initiated when there are no protons at this location i.e. during the "abort gap". Otherwise, protons will not get extracted but may be deflected into equipment.

At least 14 out of the 15 kicker magnets (MKD) must fire synchronously to deflect the beam horizontally, then the beam is deflected vertically by the 15 septum magnets (MKS), and diluted by 10 dilution kicker magnets (MKB) to spread the impact over a larger area. The beam dump TDE consisting of graphite, is located 750 m away.

The system is highly redundant, as a partial deflection will not extract the complete beam, but will be driven into the next aperture limiting point. At a misfire of one magnet, the others will be re-triggered, to extract the beam. However, this will most probably not be during the abort gap.

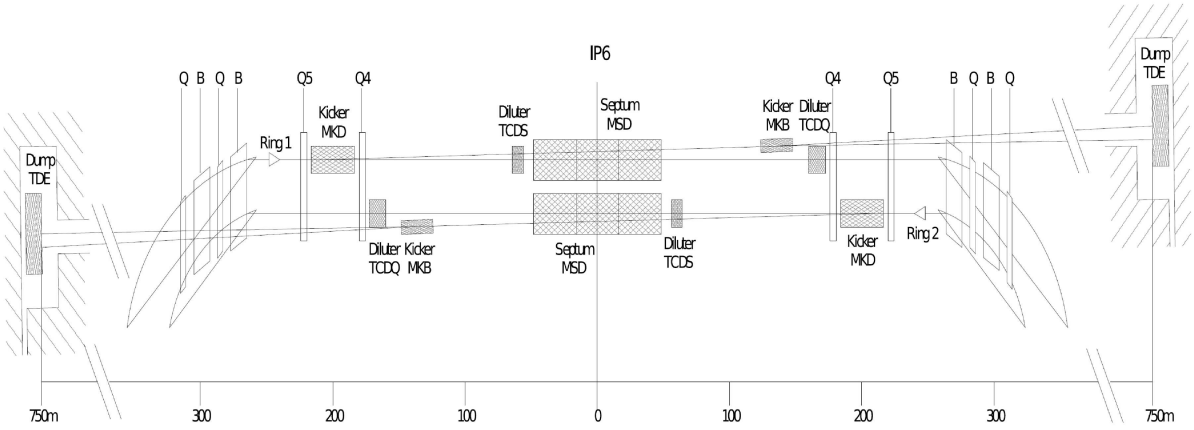

Figure 2.5: *The LHC Beam Dumping System at Point 6 [9].*

2.2 The ATLAS Detector

The ATLAS Detector is a general purpose detector, designed to cover the wide range of expected physics at the LHC accelerator. It is located at Point 1, 80 m underground. ATLAS is built cylindrically around the interaction point, with a length of 44 m, a diameter of 25 m and a weight of 7000 t (figure 2.6). Its design follows the general pattern of detectors in high energy physics collisions experiments which is illustrated in figure 2.7.

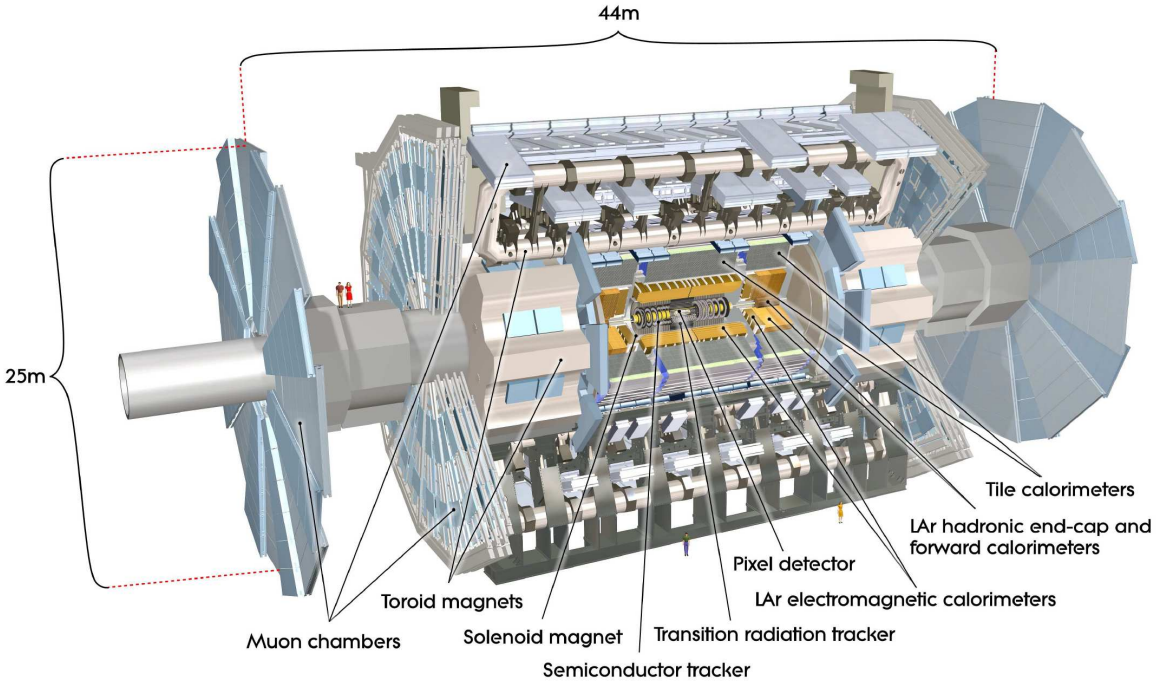

Figure 2.6: *The ATLAS detector [21].*

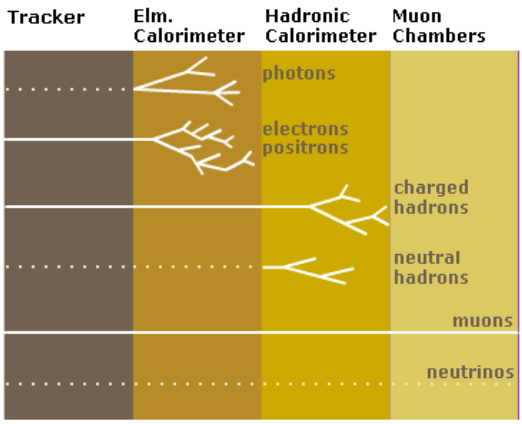

Figure 2.7: *Particle Detection in different sub-detector components (original figure from [20]). Magnetic fields which allow to measure particle momentum and charge are not shown.*

Nearest to the beam-pipe is the tracking system which determines tracks of charged particles. Magnetic fields allow the measurement of the detected particles' momentum. The material used must be kept to a minimum to avoid that particles lose much energy already in the tracker. The next layer consists of calorimeters. Their task is to measure the particles' energy which is deposited in the detector material. The electromagnetic calorimeter measures the

energy of photons, electrons and positrons, which they lose by shower production through bremsstrahlung and pair production. Hadrons are not stopped in the electromagnetic calorimeter, but only in the next layer, the hadronic calorimeter. Hadronic showers are created by strong interactions with the nuclei of the calorimeter material. The only particles that are able to transverse all the previous layers are muons, so the outermost layer is dedicated to the measurement of muon tracks. Neutrinos will not be detected in any layer. The corresponding sub-detectors of ATLAS will be described in the next sections.

ATLAS coordinate system

The ATLAS coordinate system is a Cartesian system with the x-axis pointing from the interaction point towards the center of the LHC ring, y going upwards, and z in the direction of beam 2. Accordingly, the C-side of the detector is at negative z, and the A-side at positive z.
Alternatively, more adapted to the physics event, a radial coordinate system is used, with z the distance from the interaction point along the beam-pipe, φ the azimuthal angle around the beam-pipe with $\varphi = 0$ at $y = 0$ and positive x, and the polar angle θ. Instead of θ the corresponding pseudo-rapidity $\eta = -\ln\tan(\theta/2)$ of mass-less particles is generally used, with negative η at the C-side, $\eta = 0$ at $z = 0$, and positive η at the A-side. Figure 2.8 depicts the used coordinate system.

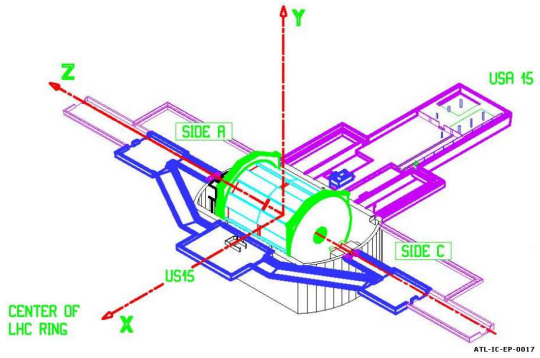

Figure 2.8: *The ATLAS coordinate system. Also shown is the location of the counting rooms and side definition [22]).*

2.2.1 The Inner Detector

The inner detector provides precise tracking information in a $|\eta|$ range up to 2.5.

To provide charge and momentum measurement, the inner detector is placed inside a superconducting solenoid magnet producing a 2 T magnetic field parallel to the beam axis, bending the tracks of charged particles originating from the interaction point in φ direction.

The two innermost detectors are silicon detectors with readout segmentation in pixels or strips, while the outermost sub-detector is a straw tube detector with additional particle identification capabilities using transition radiation.

The high granularity and the close proximity to the interaction point allow for good spatial resolution, but impose additional constraints regarding the high radiation environment.

An evaporative cooling system is responsible for dissipating the power generated by the silicon detectors and keeping them at their operation temperature of about -10°C to prevent negative effects from high radiation doses. Thermal enclosure heaters outside the SCT allow the TRT to operate at room temperature.

Finally, near the beam pipe, a Beam Conditions Monitor (chapter 4) is installed inside the pixel support structure, and a Beam Loss Monitor is located at the ID end plates.

Figures 2.9 and 2.10 show the layout for the inner detector .

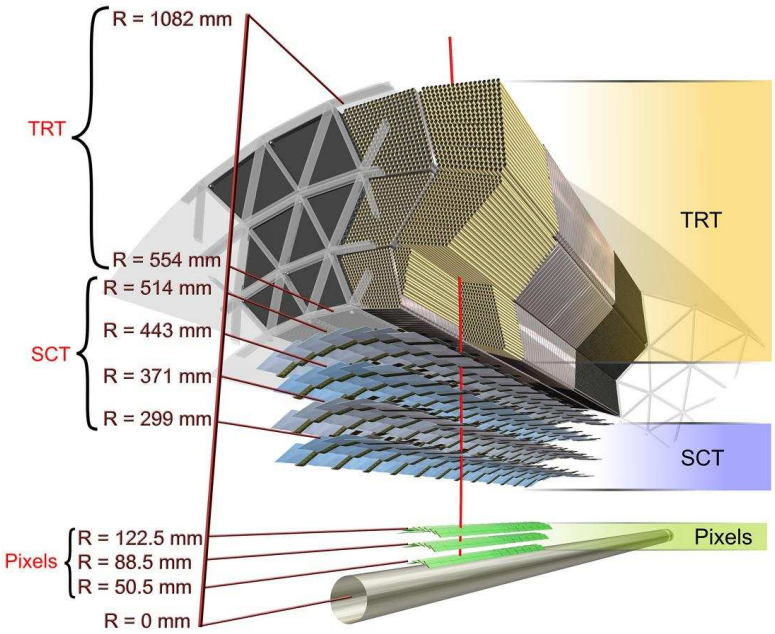

Figure 2.9: *3D view of the Inner Detector [21].*

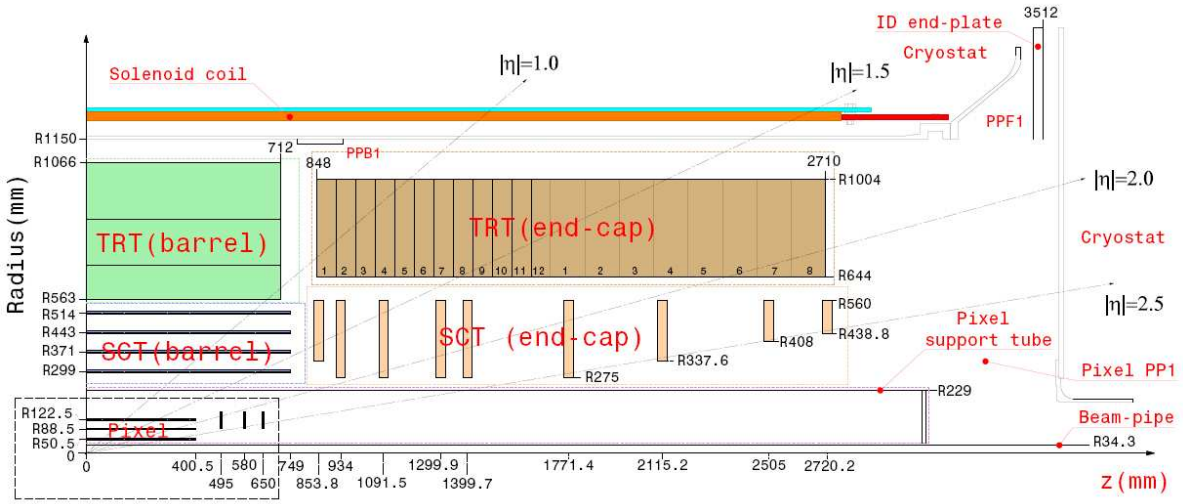

Figure 2.10: *Layout of the Inner Detector [21].*

2.2.1.1 The Pixel Detector

The Pixel Detector is a semiconductor tracking detector with 80 million readout channels, providing precise tracking information near the interaction point. It will be described in more detail in chapter 3.

2.2.1.2 SCT

The SemiConductor Tracker (SCT) is a silicon based microstrip detector. It consists of four layers in the barrel region and nine endcaps on each side with a total of 4088 modules. The barrel layers contains 2112 6.4 cm long double-sided modules with 768 strips each, parallel to the beam axis, but with a stereo angle of 40 mrad between them to allow a more precise determination of the z-coordinate, and a pitch of 80 µm. The nine disks are equipped with 1976 modules of four types – with different lengths depending on the mounting radius. The achieved spatial accuracy in φ is 17 µm, and still 580 µm in the coordinate parallel to strip direction.

Each module is equipped with six readout chips. For data transfer to the counting room, signals are converted from electrical to optical directly on the module, with one control link per double module and two data lines. The SCT has about 6 million readout channels.

2.2.1.3 TRT

The Transition Radiation Tracker (TRT) is a straw tube detector. It provides additional tracking information up to $\eta = 2.0$, and particle identification.

The individual tubes have a diameter of 4 mm and are filled with a xenon-based gas mixture. The 35 µm thick wall lies at a potential of -1530 V with respect to the central wire. In the barrel region, the drift tubes are arranged in three layers of modules, containing 144 cm long straw tubes parallel to the beam axis divided at $\eta = 0$. Polypropylene fibres are located between the tubes to enable transition radiation. The TRT end caps start from z = 85 cm, each containing wheels with four layers of tubes, and polypropylene foil between layers. The tubes are radially arranged with a length of 37 cm, and two layers staggered against each other. A low threshold and a high threshold in the front-end electronics are for hit registering of minimum ionizing particles and identification of low energy transition radiation photons generated by electrons at the transition from the CO_2 medium inside the drift tubes to polypropylene, with different dielectric constants. With drift time measurement, the achieved accuracy per straw is 130 µm. Per track about 36 tubes are transversed.

2.2.2 Calorimeters

The calorimeters have to measure deposited energy with a good resolution in energy and space, to provide enough sensitivity to interesting physics processes. Two possible channels for Higgs production are e.g. $H \rightarrow \gamma\gamma$ or $H \rightarrow ZZ \rightarrow e^+e^-e^+e^-$, defining requirements for electromagnetic calorimetry. A good coverage of a maximum part of the solid angle is necessary to reliably determine missing E_T. The calorimeters provide input to the (Level-1) trigger decision (see section 2.2.5). Figure 2.11 shows an overview of the ATLAS calorimeters.

Both, electromagnetic and hadronic calorimeters are sampling calorimeters, i.e. consisting of alternating layers of active material and absorber. The calorimeter system can be subdivided in barrel, endcap, and forward part, each region containing electromagnetic and hadronic calorimetry. The active material used is liquid argon for all subsystems, except in the barrel region of the hadronic calorimeter, where scintillating tiles are used.

2.2.2.1 The Electromagnetic Calorimeter

The electromagnetic calorimeter uses liquid argon as active material, with lead as absorber and kapton electrodes. It has an accordion structure to avoid gaps in the coverage, as shown in figure 2.12.

The liquid argon calorimeter is located in a cryostat to keep the active material at a temperature of 80 K.

In the barrel region, the liquid argon calorimeter is divided into two parts at $\eta = 0$, and is segmented in three divisions in radial direction, with a total thickness of more than 22 X_0 radiation length, dependent on η. The barrel

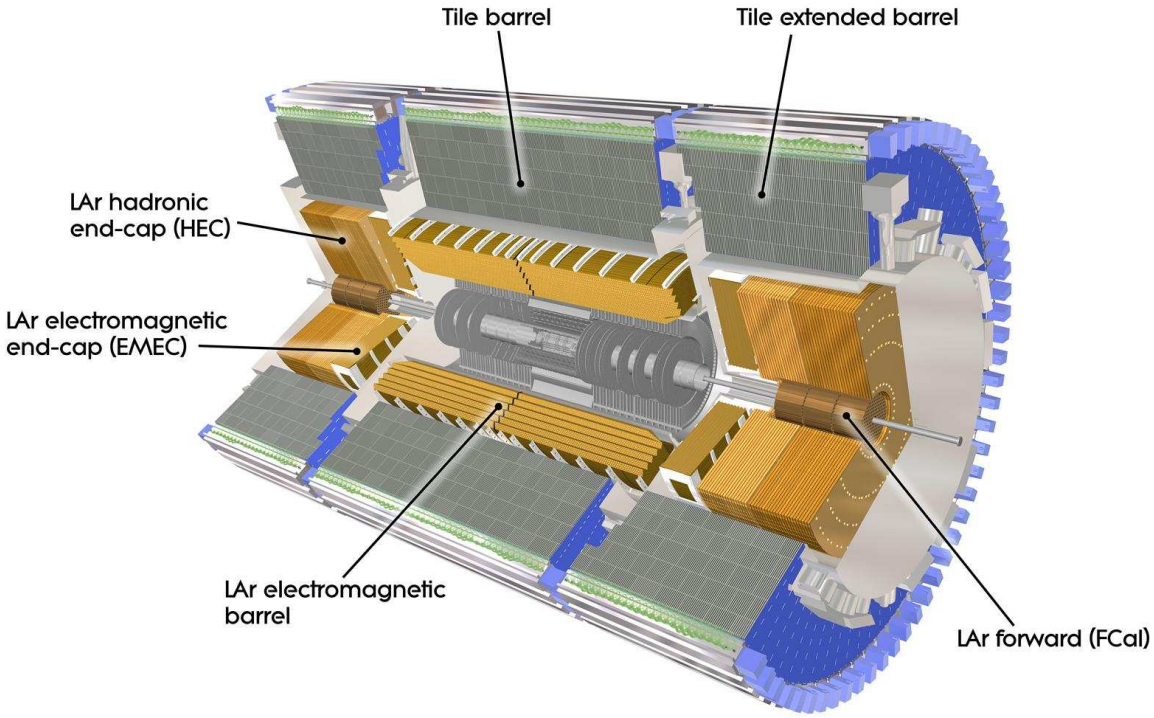

Figure 2.11: *The ATLAS calorimeters [21].*

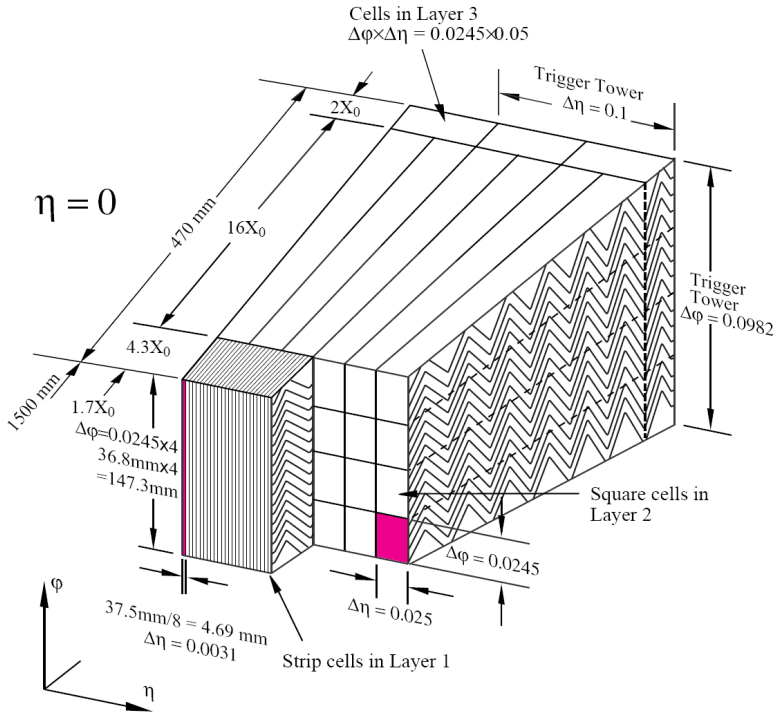

Figure 2.12: *ATLAS Liquid Argon Calorimeter Structure [21].*

region covers $|\eta| < 1.475$. Depending on $|\eta|$ and layer, the granularity ranges from $\Delta\eta \times \Delta\varphi = 0.025/8 \times 0.1$ to 0.075×0.025.

In the endcap, there is one wheel on each side, providing coverage for $1.375 < |\eta| < 3.2$. The electromagnetic calorimeter provides more than 150000 readout channels.

The design energy resolution of the electromagnetic calorimeter in the barrel and endcap region is $\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$ [23].

The forward calorimeter covers $3.1 < |\eta| < 4.0$. The electromagnetic part of the forward calorimeter uses copper as absorber and liquid argon as active material. It provides 1000 readout channels.

2.2.2.2 The Hadronic Calorimeter

The Hadronic Calorimeter in the barrel region uses steel as absorber and scintillating tiles as active material. The emitted light created by hadronic interactions is proportional to the deposited energy. It is collected by wavelength shifting fibres at the edge of each tile, into photomultiplier tubes located at the top of the Tile Calorimeter modules. Figure 2.13 shows one of 64 azimuthal symmetric modules.The Tile Calorimeter is partitioned into barrel and extended barrel region on A- and C-side each.

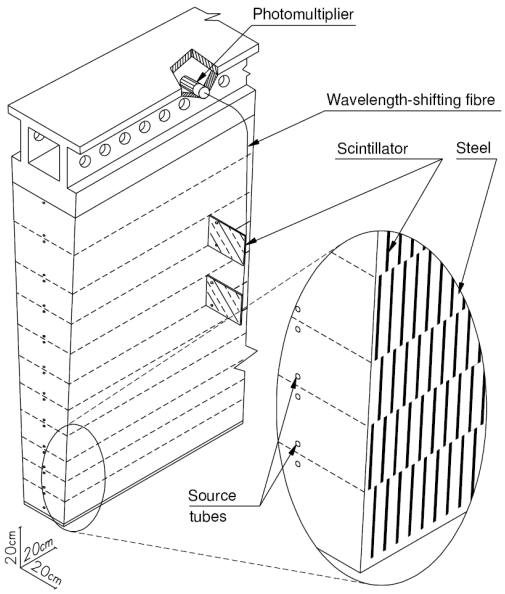

Figure 2.13: *ATLAS Tile Calorimeter Module [21].*

The endcap part of the hadronic calorimeter uses Liquid Argon as active material to achieve the required radiation hardness. It consists of two endcaps behind the electromagnetic calorimeter endcap, using copper as absorber. It is located at $1.5 < |\eta| < 3.2$.

The design energy resolution of the hadronic calorimeter in the barrel and endcap region is $\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$ [24].

The hadronic Forward Calorimeter consists of two sections, using tungsten as absorber and liquid argon as active material.

2.2.3 Muon Chambers

The Muon spectrometer is the outermost subsystem of ATLAS. Figure 2.14 shows the general layout of the Muon chambers. They are arranged into a barrel region with three layers, and an endcap region with inner, middle, and outer layer. Middle and outer layers consist of the big wheels located at z > 13 m and z > 20 m, while the inner layer with the small wheel is located before the endcap toroid. The Muon system has two tasks, to provide a precise measurement of muon tracks, and to provide a fast muon identification for the trigger. For this end, four different technologies are used. For the tracking chambers, three layers of Monitored Drift Tubes (MDTs) are used in the barrel and endcap region. In the endcaps for high η, the expected event rate and radiation level is however too high for MDTs, therefore Cathode Strip Chambers (CSCs) are used here. For the trigger, two different technologies are used, Resistive Plate Chambers (RPCs) in the barrel region, and Thin Gap Chambers (TGCs) in the endcaps.
To allow for the momentum measurement, the Muon chambers are located inside the magnetic field of the toroid magnet system.

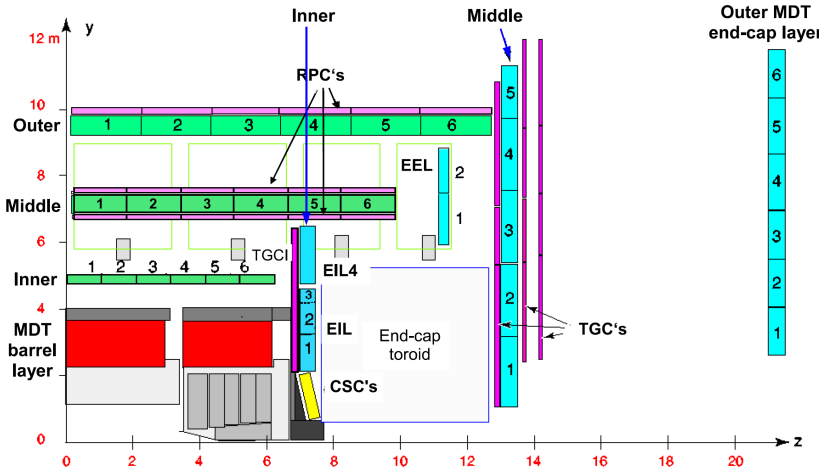

Figure 2.14: *The ATLAS Muon Chambers [21].*

2.2.3.1 The Toroid Magnet

The Toroid Magnet consists of the barrel toroid and two endcap toroids. Each part is built of eight superconducting air core coils, with the endcap structure rotated by 22.5° with respect to the barrel, as can be seen in figure 2.15.

The dimensions of the barrel toroid define the size of the ATLAS detector, with a length of 25.3 m, and inner and outer diameter of 9.4 m and 20.1 m. The field has no component in z-direction. The maximum field of 3.9 T in the barrel and 4.1 T in the endcaps is generated by currents of 20500 A, with 2 hours for a slow dump and 2 minutes for a fast dump. The normal operation temperature lies at 4.6 K. In case of a fast dump, the energy taken by the cold mass results in a temperature of 58 K.

2.2.3.2 Monitored Drift Tubes (MDT)

The MDT subdetector is responsible for precision measurement of muon tracks in a region of $|\eta| < 2.7$. It consists of aluminum drift tubes with a 3 cm diameter with an anode wire in the middle. The tubes differ only in length, for the different regions of the detector, with a maximum of up to 5 m length in the barrel and 6 m in the endcaps. The volume is filled with an Ar/CO_2-gas mixture in a ratio of 93:7, operating at a pressure of 3 bar. Transversing muons create electrons that drift towards the anode. The drift time of up to 700 ns is determined by the minimal

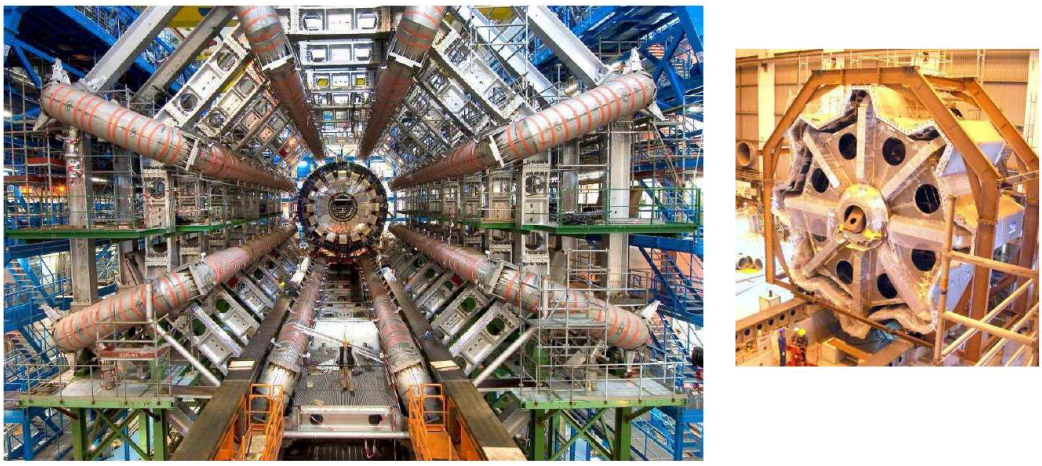

Figure 2.15: *The Toroid Magnet systems. Left: Barrel Part. Right: Cold Mass of Endcap Toroid. [21].*

distance of the track to the wire, with a strong radial dependence of the drift velocity ranging from 10 to 50 µm/ns. A chamber is built of 2 multilayers of 3 or 4 layers of tubes each with a spacing structure in the middle. The tubes are arranged tangentially, i.e. along the magnetic field. In the barrel region, there are three layers of MDTs divided into 16 sectors, with alternating large and small chambers to provide overlap. Each sector consists of 6 chambers per side, with a slightly different geometry in the feet region to account for the detector support structure. In the endcap regions, the trapezoidal chambers are arranged into 16 sectors pointing towards the beam axis. The outer and middle endcap consist of 6 chambers each, while the innermost layer consists of 4 MDT chambers covering up to $|\eta| < 2.0$. At the intersection region between barrel and endcap region in the middle barrel layer an extra endcap layer of 2 chambers provide the coverage.

The achieved spatial resolution in z direction is 80 µm per tube or 35-40 µm per chamber, requiring a precise knowledge of the drift time properties. In addition, a dedicated high precision optical alignment system constantly monitors the position of the chambers. This allows for a standalone muon momentum measurement of better than 10% for 1 TeV tracks.

2.2.3.3 Cathode Strip Chambers (CSCs)

Cathode Strip Chambers (CSCs) cover the region of $2 < |\eta| < 2.7$ due to their tolerance to high rates and better time resolution with an electron drift time of only 40 ns. The CSCs are multi wire proportional chambers with an Ar-CO_2 gas mixture. The wires are aligned in radial direction. The CSCs provide track coordinates in both directions, η and φ.

The CSCs consist of two endcaps with a geometry similar to MDT, of 16 sectors with alternating large and small chambers (see figure 2.16). Each chamber consists of four layers built from 19 mm polyurethane foam with 17 µm copper cladding forming cathode strips. The tungsten-rhenium anode wires are at a potential of 1900 V. The distance between the individual anode wires, and between the anode wires and the cathode strips is both 2.5 mm for the η direction, resulting in $4 \cdot 192$ readout strips per chamber. The track location is determined by the relative signal distribution of 3-5 cathode strips around the maximum pulse height.The cathode strip orientation alternates inside the four CSC layers between perpendicular and parallel to the anode wires, which allows for track measurement in η and φ and two track separation according to pulse height. For better resolution, the CSCs are inclined at an angle to be perpendicular to incoming tracks. The resulting spatial resolution in η is 60 µm per plane. The readout in φ is deliberately wider spaced, with 192 readout strips per chamber (48 per plane), resulting in a resolution of 5 mm.

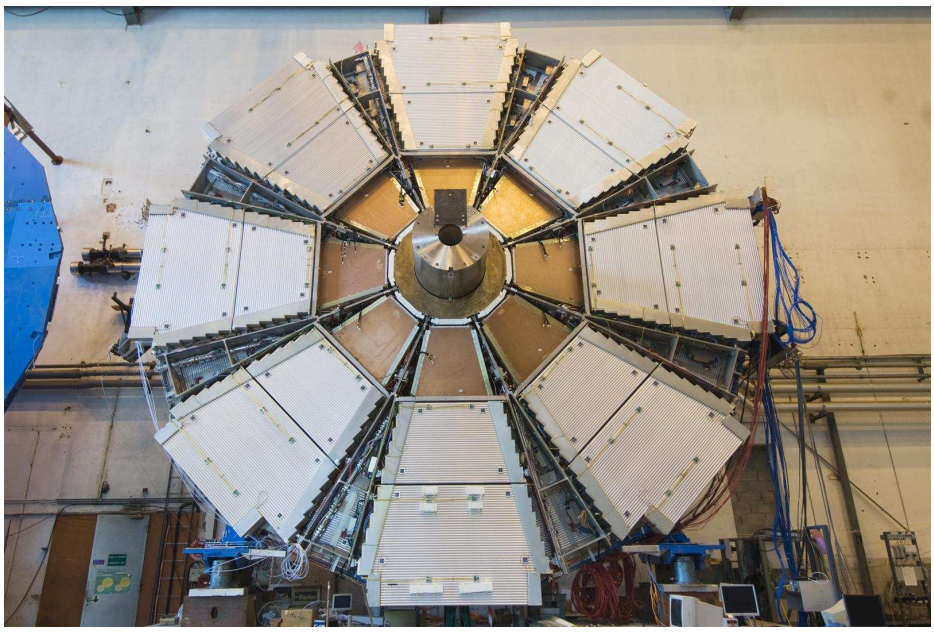

Figure 2.16: *The ATLAS inner endcap muon layer of the small wheel [21]. The innermost chambers are built of CSCs, while for lower η MDTs are used.*

2.2.3.4 Resistive Plate Chambers (RPCs)

To provide fast trigger information as well as the φ information that the MDTs can not provide, in the barrel region Resistive Plate Chambers are installed. They are built from two parallel resisistive plates with a distance of 2 mm between them and a field of 4900 V/m. The volume is filled with a $C_2H_2F_2$-IsoButan-SF_6 gas mixture. The avalanches created by the primary ionization electrons are read out via capacitive coupling to copper readout strips on the outside of each resistive plate. Figure 2.17 shows the meachanical layout of the RPCs.

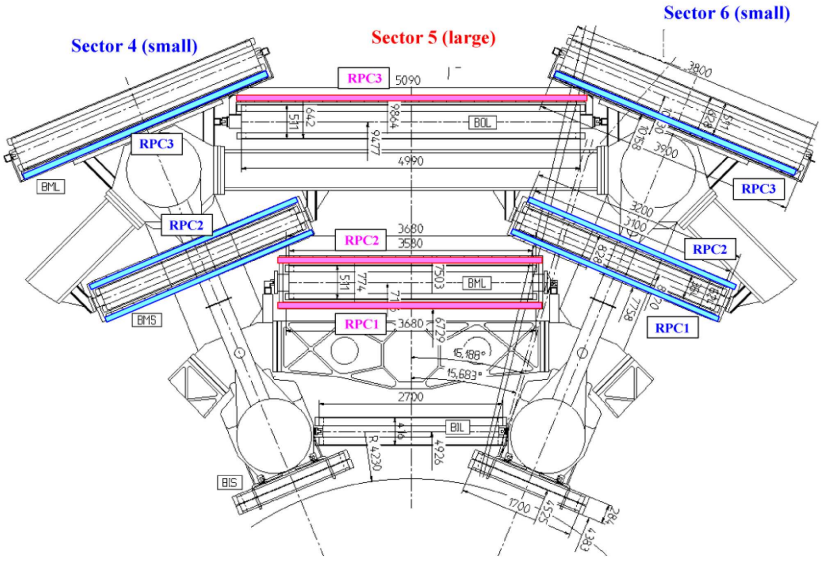

Figure 2.17: *ATLAS RPC layout [21].*

In the middle layer, each MDT chamber is coupled with two RPC chambers, one on each side. In the outer layer

one RPC chamber is coupled to each MDT chamber. For the small sectors the RPCs are on the outside, for the large sectors on the inside. RPC chambers which span more than 1200 mm in z-direction consist of two detector units , which are staggered to avoid inactive regions. Each unit consists of two gas volumes, i.e. four resisistive plates and four planes of readout strips. The readout strips are alternating between z- and φ-direction. Therefore each transversed chamber provides information with two φ and two η readings. The resisistive plate chambers have a time resolution of a few nanoseconds, and a spatial resolution of 1 cm [25]. Coincidences between the first two RPC layers are used to trigger on low-$p_T$muons, coincidences on all three layers for high p_T.

2.2.3.5 Thin Gap Chambers (TGCs)

In the endcap region, Thin Gap Chambers (TGC) provide the muon trigger with a good time resolution of a few nano seconds. TGC consists of multi-wire-proportional chambers with a distance between anode wires of 1.8 mm and a distance of anode wire to cathode plane of only 1.4 mm. The volume is flushed with a flammable mixture of CO_2 and n-pentan. The wires lie at a potential of 2900 V. Figure 2.14 shows the location of the TGCs, in figure 2.18 the geometry of the big wheel can be seen. The big wheel has three TGC layers, where the innermost layer forms a triplet, i.e. has three planes of anode wires, and the outer two doublets. All layers have two layers of segmented copper cladding forming readout strips in φ direction, while each wire plane provides readout information in η. The layers are divided into 12 sectors.

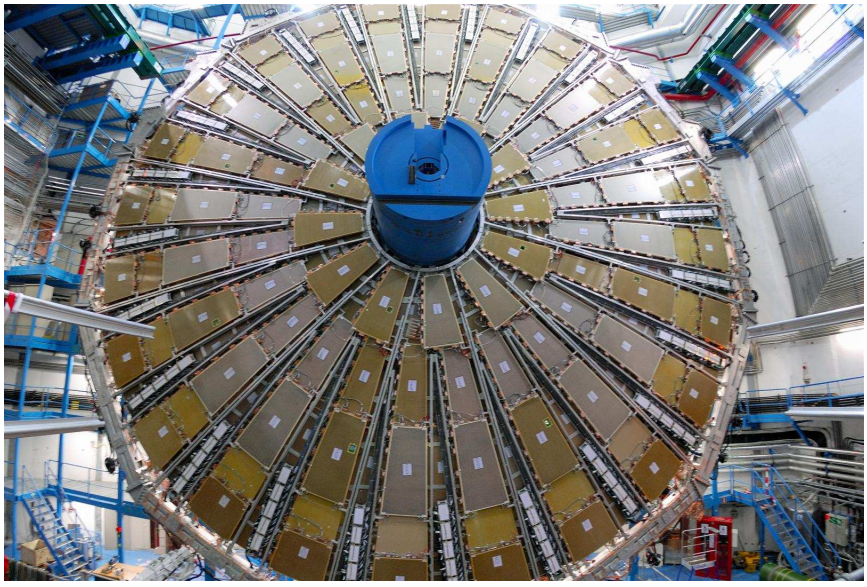

Figure 2.18: *ATLAS TGCs. [20].*

To account for varying required granularity with increasing η, different numbers of wires are grouped together to form a readout channel.

2.2.4 Forward Detectors

At high η three additonal subdetector systems are installed. LUCID and ALFA are dedicated to luminosity measurement, while ZDC is important for heavy ion studies due to its ability to measure extremely forward neutrons.

LUCID

LUCID (LUminosity measurement using a Cherenkov Integrating Detector) is mainly dedicated to the online monitoring of the instantaneous luminosity in ATLAS with a time resolution good enough to measure on a bunch by bunch basis. It is located at z =17 m from the interaction point at a radius correlating to $5.6 < |\eta| < 5.9$. LUCID measurements are calibrated for determination of the absolute luminosity using the knowledge of LHC machine parameters and known physics processes like muon pair production from two photon interactions or leptonic decays of W and Z bosons. LUCID operates under the premise that the number of detected particles is proportional to the number of interactions per bunch crossing. It consists of one module at each side, each containing twenty 1.5 m long aluminium tubes around the beam axis pointing toward the interaction point with a diameter of 15 mm. The tubes are filled with C_4F_{10}, so that Cherenkov light will be created by incoming particles, which is measured by photomultiplier tubes.

ZDC

The Zero Degree Calorimeter (ZDC) is installed 140 m from the interaction point. Its main task is the detection of forward neutrons in heavy ion collisions. By measuring the number of ("spectator") neutrons basically not involved in the collision, it provides information about the centrality of the event. ZDC is achieving a time resolution of 100 ps. On each side of the interaction point four modules are installed, one electromagnetic and three hadronic. The Cherenkov light generated in a matrix of quartz rods inside tungsten plates provides position and energy information.

ALFA

The ALFA (Absolute Luminosity For ATLAS) Detector has been designed for precise absolute luminosity measurement by determining the rate for elastic pp scattering at very low angles. For this, special beam conditions with high β^* and low beam emittance are required. ALFA will be installed about 240 m from the interaction points, with two stations at each side. The installation in Roman Pots will allow ALFA to be moved to a distance of 1 mm to the beam. The detector uses plastic scintillating fibres with a resolution of 30 µm.

2.2.5 Trigger and Data Acquisition (TDAQ)

With a nominal bunch crossing rate of 40 MHz and an interaction rate of 25 collisions per bunch at nominal luminosity, it is impossible to record every single event. To reduce the data rate while keeping the number of interesting events high, a three stage trigger system is in place, shown in 2.19.

The first level – the Level-1-Trigger (L1) – has to do a fast trigger decision reducing the rate to a maximum of 75 kHz. The Level-2-Trigger (L2) rejects L1-triggered events based on more detailed information of the event in the region of interest (RoI) denoted by the L1-trigger, with an output rate of 3 kHz. The Event Filter (EF) bases the final decision on information of the complete event, reducing the rate to 200 Hz.

The Level-1 trigger decision is based on input from the calorimeter system and muon trigger chambers. Events can be selected for high $p_T$muons, or electrons, photons, jets, or taus. The calorimeter can also provide trigger on high transverse energy or high missing transverse energy indicating neutrinos. Additional inputs – up to a total of 160 – from subsystems like LUCID, TRT, BCM, or in particular for early running from the MBTS [3)] are possible. This inputs are processed in the central trigger processor CTP with a resulting trigger menu of 256 different trigger conditions. The CTP sends the Level-1 triggers with the identification of the event bunch crossing (BCID) to the subsystems for the readout of buffered data. The time till the decision is made to readout an event as well as signal

[3)] The **M**inimum **B**ias **T**rigger **S**cintillators (MBTS) are mounted on the inner surface of the LAr endcap cryostat. The system consists of sixteen 2 cm scintillator counters on each detector side, divided into two eta regions each: $2.12 < |\eta| < 2.83$ and $2.83 < |\eta| < 3.85$.

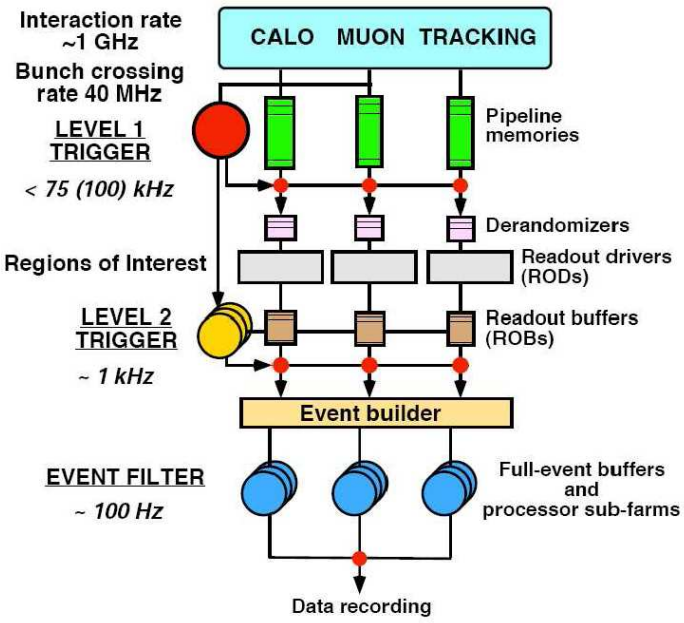

Figure 2.19: *The ATLAS Trigger system [26].*

propagation, the event data has to be kept in buffers on the detector readout electronics. For this reason, the trigger latency has to be kept low, with the required maximum below 2.5 µs, which is realized by custom built hardware and optimised cable routing.

While Level-1 trigger is based on custom built hardware, the High Level Trigger (HLT) consists of off-the-shelf processor farms. In case of a LVL1 trigger being fired, the complete event is read out to the subdetector RODs (Read-Out Drivers). From the ROD the data is transferred via the ROL (Read-Out Link) to the ROS (Read-Out System) where it is held in the ROB (Read-Out Buffer).

Only the data relevant to the Regions of Interest defined by the Level-1 input is requested by the Level-2 trigger from the corresponding ROSs. While the L1 Trigger only has available as information the multiplicities and flags for passed thresholds, the L2 trigger can refine by the trigger decision e.g. by the access to finer granularity data from the calorimeters or by access to inner detector data. After the Level-2 selection – which takes normally in the order of 40 µs – the event is either discarded, or sent from the ROSs to the Event Filter. Here, the event fragments are transferred to one of the SFI (Sub-Farm Input) nodes, where the event is then assembled. From the SFI, the complete event is sent to the Event Filter where the event is further analyzed using the same software framework – Athena – as the later offline analyses. Also, at this point the event fragments get deleted from the ROBs. Based on the used trigger, the events are categorized into streams and in case of selection passed on to the SFO (Sub-Farm Output) from where they are transferred to further storage and offline reconstruction. The event filter decision time can be up to 1 s.

Complementary to this dataflow management, the Data Acquisition has a component for the online coordination monitoring and control of all the subsystems.

2.2.6 The ATLAS Detector Control System

The Detector Control System (DCS) is responsible for the supervision and control of the ATLAS detector where safety and operation are concerned, excluding data taking related aspects for which the DAQ is responsible [27], [28]. For the combined DCS operation of all subsystems, a scheme exclusively based on the use of two interfaces – FSM (Finite State Machine) and alarm screen – is envisaged [27]. While the FSM allows the supervision of

operational conditions and transitions between them, the alarm screen is a redundant but complementary way to monitor all error conditions appearing in the FSM. The responsibility for low-level implementation and operation is in the hand of the subdetectors, while the central DCS is responsible for the overall integration, the interface to the LHC, as well as monitoring of common issues like infrastructure components (e.g. supervision of the racks), and provision of central services (like access control server or DIM (Distributed Information Management, see 5.2.6.1) name server). Figure 2.20 shows the resulting hierarchical FSM structure integrating the subdetectors into a common system. The interface to the DAQ, i.e. TDAQ/Run Control, provided by DDC (DAQ DCS Communication), is foreseen on the TTC partition level, i.e. inside the subdetector branches, below the subdetector nodes. A detailed description for the implementation of the corresponding Pixel subsystem components will be given in chapters 5 to 7.

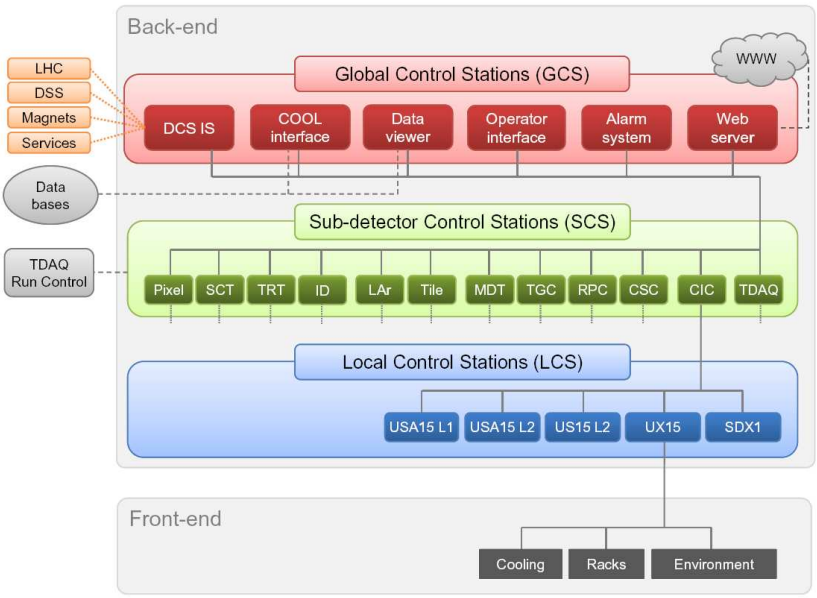

Figure 2.20: *ATLAS DCS structure [27].*

2.2.6.1 The ATLAS LHC Interface

The interface between the ATLAS subdetectors and the LHC provides the stable beam signal from LHC to the subdetectors, and delivers Injection Permit and Beam Permit to the LHC. The only ATLAS sub-detectors providing an input to the beam permit are BCM and BLM, while in addition all subdetectors sensitive to beam conditions are providing an input to the Injection Permit. These are the inner detector silicon detectors and all muon systems.

To facilitate interaction, various software flags are provided.

Each subdetector defines a software “SafeForBeam”-Flag, signifying whether the subdetector is in a safe state with respect to imminent non-stable beam conditions. For most subdetectors this is coupled to the state of the high voltage.

A handshake between LHC and the experiments takes place when LHC has the intention to inject beam into the empty machine (Injection Handshake), change from stable beam mode to non-stable conditions (Adjust Handshake), or dump the beams (Dump Handshake) [101].

2.2.7 The ATLAS Detector Safety System

The ATLAS Detector Safety System (DSS) is providing safety on a detector global level [29], [30]. Its task is to bring the detector and equipment into a safe state in case of detected error conditions. Due to the requirement of highest level availability, reliability and robustness, while keeping a level of programmability, it makes no use of any software components, but relies instead on PLCs (Programmable Logic Controller). The structure of the DSS system is divided into three parts:

- sensor **input** detecting alarm conditions
- activation of an **alarm** depending on logical combination of sensor signals
- **action** taken on the alarm to shut down equipment

The relation between alarms and actions is described in an "Alarm Action Matrix". For each entry there are defined two corresponding time intervals determining the action. The first is the latency, the required duration for which a sensor must show an error condition before it triggers the alarm to avoid reacting on glitches or spikes. I.e. if an error condition disappears before that interval, no alarm will be raised. In the other direction, the conditions must be good for some amount of time, to be considered valid. Once the alarm is active, it can only be reset when the error condition has disappeared and has been inactive for the same time required for the raising of the alarm, and a human operator – normally the SLIMOS (Shift Leader In Matters Of Safety) – acknowledges and resets the alarm. The second time interval is a delay between the raising of the alarm and the execution of the action. This allows DCS to take actions before the hard DSS shutdown, e.g. ramp down power supplies before the rack power is cut. This second number is a delay, i.e. once the alarm has been raised, there is no regular way to stop the action from being executed, even if the error condition disappears and the alarm gets acknowledged in the meantime.An example of this structure is shown in figure 2.21.

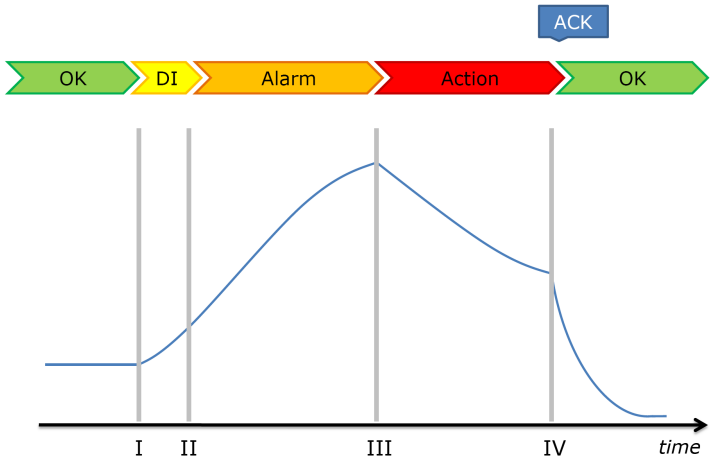

Figure 2.21: *Schematic structure of a DSS action. Time and value in arbitrary units, with high values, e.g. a temperature of a power supply, corresponding to bad conditions. At time I the dedicated DSS-sensor registers an error condition, e.g. due to a cooling failure. After this error condition has been present for a defined time, at time II the DSS-Alarm becomes active. For each pair of alarm and action a delay is defined, after which the corresponding action is triggered at time III. When the error condition is gone – at time IV – the digital input must have been FALSE again for the defined activation time, and at this point it is possible to acknowledge the action. Only after the acknowledging the affected device can be switched on again.*

Even though the DSS is not relying on PCs for its operation, the backend system for supervision and operator interface is implemented on PCs using the same SCADA PVSS II as the DCS systems. Also the Alarm Action Matrix can be programmed using tools implemented for this DSS backend system. Though there is no direct

connection between the systems, the relevant DSS parameters are exported via DIP (Data Interchange Protocol[4)]) to ATLAS DCS to allow the subdetectors to implement monitoring and actions on detected error conditions.

[4)] Protocol based on DIM, with attention to high reliability of the data transfer.

Chapter 3

The Pixel Detector

The Pixel Detector is the innermost sub-detector of ATLAS. Its main task is tracking, and in particular secondary vertex reconstruction. It has to provide two dimensional information with very good spatial resolution of 15 µm in r-φ and 120 µm in z direction. Three space points per track are desired up to $|\eta| < 2.5$. Due to the extreme proximity to the interaction point, high radiation hardness to a lifetime dose of 10^{15} n_{eq} cm^{-2} is required, as well as a low single channel occupancy. To reduce radiation damage, the Pixel Detector is operated at -13 °C. An evaporative cooling system (chapter 3.5) is installed to remove the 15 kW power generated by the detector electronics and services from the system. Figure 3.1 shows the Pixel Detector and its support structure, the Pixel Support Tube (PST).

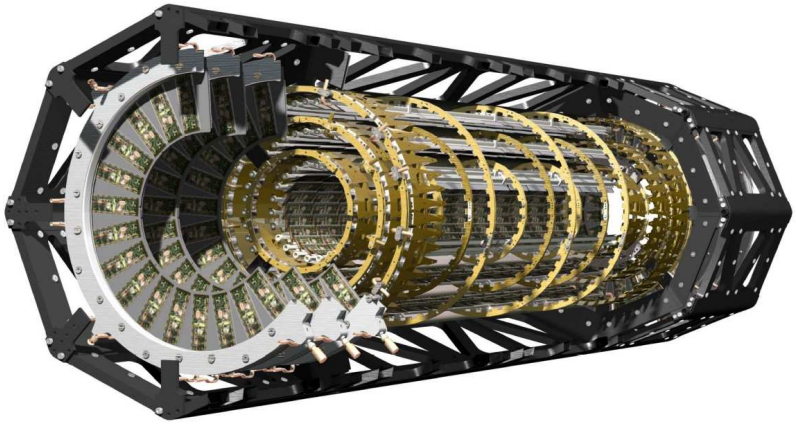

Figure 3.1: *The ATLAS Pixel Detector [21].*

3.1 Layout

The Pixel Detector is built of three concentric layers around the beam pipe and two endcaps with three identical disks. The beam pipe extends to a radius of 3.5 cm, including heaters and isolation. The layer nearest to the beam-pipe, the B-Layer, has a radius of 5 cm, the two outer layers, Layer 1 and Layer 2, are located at 9 cm and 12 cm. All three layers together are referred to as the barrel. They each have a length of 80 cm. The disks have a distance in z from the interaction point of 50 cm, 58 cm and 65 cm. The disks closest to the interaction point are called Disk 1, while Disk 3 is farthest away. The closest distance of the disk modules to the beam axis is 9 cm. The Pixel Detector is built of identical modules, and each module contains 46080 independent pixel cells. A description will be given in section 3.2.

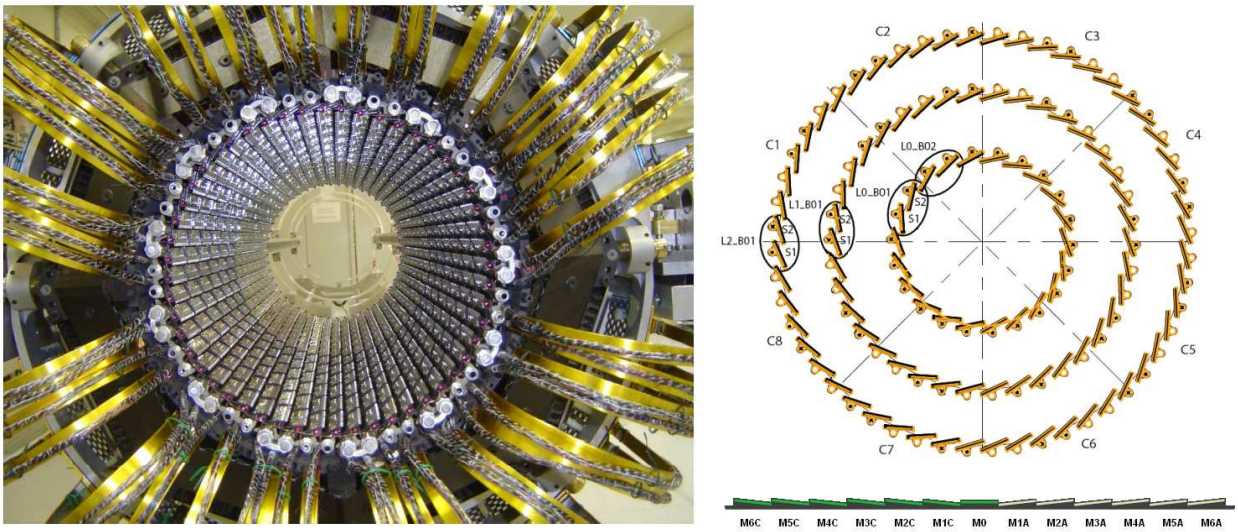

Figure 3.2: *The barrel region of the Pixel Detector. On the left, a photograph of layer 2 [21]. On the right, all three layers can be seen [32]. The staves are tilted by an angle of 20° against each other, to compensate for the Lorentz angle. The layout of the modules on a stave is shown in the bottom right. They are tilted by 1° to give full coverage (not-to-scale for better visibility).*

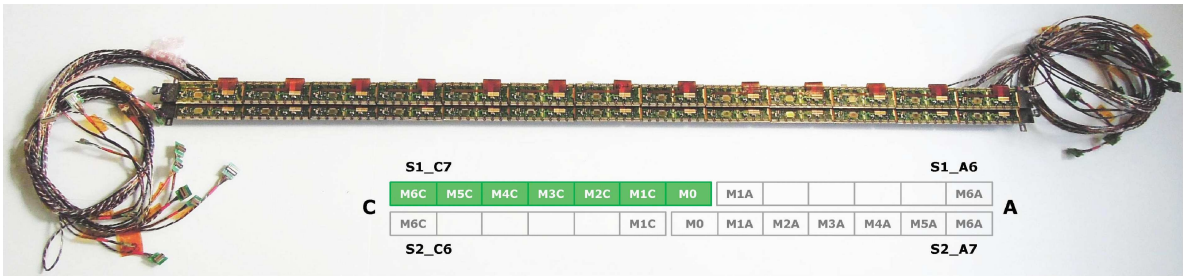

Figure 3.3: *Photograph of a bi-stave, and naming scheme. The type 0 cables for the supply and readout of the four readout units, and the cooling pipe can be seen emerging at the end of the staves. In the bottom, the naming scheme for the modules and readout units is pictured schematically.*

Due to their good characteristics for thermal conductivity, low radiation length, high mechanical stability, low thermal expansion coefficient and radiation hardness, carbon composite materials were chosen for the local support structure.

In the barrel region, the local support structures are called staves. Each stave is loaded with 13 modules. The middle one, “M0”, is oriented in a plane parallel to the beam pipe. The outer modules are tilted by 1.1° each. To cover the whole space around the interaction point, and to minimize the number of hit pixels in a module for tracks from the interaction point inside the magnetic field, i.e. to compensate for the Lorentz angle, each stave is tilted by an angle of 20°. Modules are numbered from inside to outside with a suffix indicating on which side of the stave they are mounted, i.e. M1A to M6A on the A side, and M1C to M6C on the C side. Figure 3.2 shows a photograph of Layer 2, and the overall layout of the barrel region.

In the endcaps, six modules are mounted on every disk sector, three on the front side, and three on the back side. To cover the complete disk area, the modules overlap at lower radius. A photograph of a pixel disk, as well as the overall layout of the disks is shown in figures 3.4 and 3.5. In each disk sector, modules M1 to M3 are facing towards the interaction point, while modules M4, M5, and M6 are facing the opposite direction [33]. With respect to the power supply system, in the endcaps one sector is the smallest independent unit.

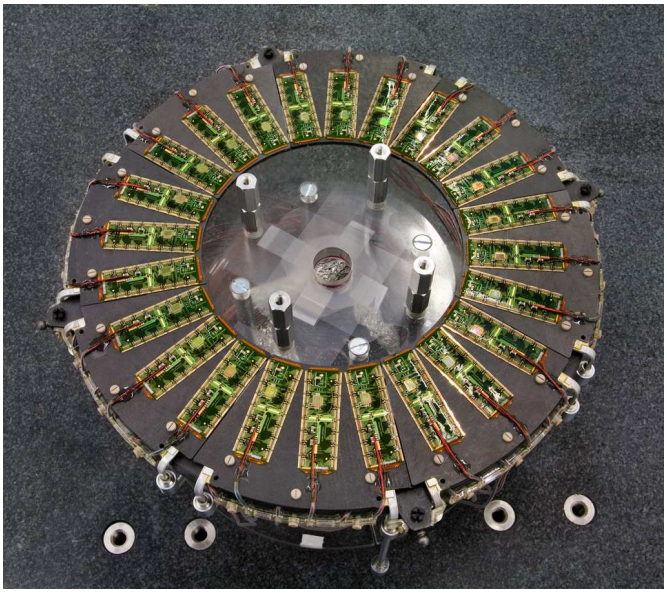

Figure 3.4: *Photograph of one Pixel Disk.*

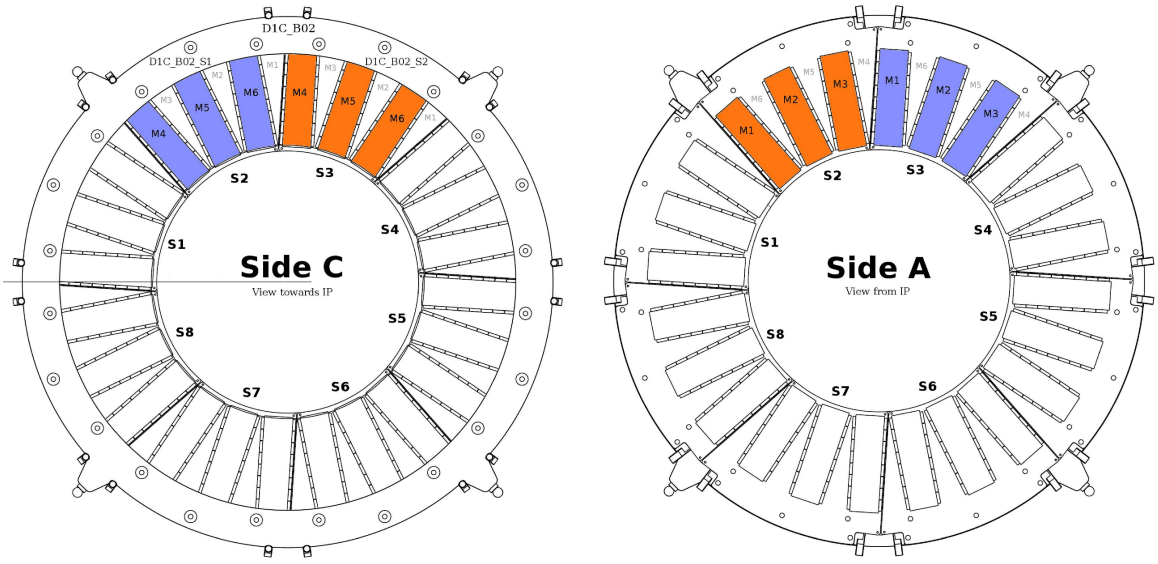

Figure 3.5: *Layout of the Pixel Detector disk modules. The different orientation of the module naming scheme when mounted is due to the fact that all sectors are physically identical.*

Integrated into the staves and disk sectors are the aluminium pipes for the evaporative cooling system. In the barrel region two staves, or a “bi-stave”, form a parallel cooling circuit (PCC), in the endcaps two disk sectors, or a “bisector”. The B-Layer consists of 11 bi-staves, Layer 1 of 19, and Layer 2 of 26 bi-staves. Each disk has four bisectors.

The naming convention for the detector elements follows this geometry. Each module is identified in the following way:
<LayerOrDisk>_<PCC>_<HalfstaveOrSector>_<Module>
The Layers are called L0, L1, L2, and the Disks D3C, D2C, D1C, D1A, D2A, D3A.
The PCCs go from B01 to B26, with a maximum of four PCCs for the disks, and 26 for Layer 2. The first PCC is

located near $\varphi = 0$, and for the disks $\varphi = 0$ is located between sector S1 and S2 of the first PCC.
For the Disks, the sectors are called S1 and S2, while in the Barrel the half-staves are called S1_A6, S1_C7, S2_A7, S2_C6. The "readout units" on the level of the sectors and half-staves are the smallest possible units with respect to the optical readout which can be independently operated.
The modules go from M0 to M6, or M1 to M6. The barrel modules except for M0 have a suffix indicating the side, e.g. M1A or M6C. Disk modules always go from M1 to M6 without suffix.

3.2 The Pixel Module

The Pixel Detector is built from 1744 identical modules. The modules consist of the sensor, the front-end chips, and a flexible kapton PCB ("flex") with the Module Control Chip (MCC). The pixel implants of the sensor are bump bonded to one of 2880 readout cells of a front-end chip. Each module has 16 front-end chips arranged in two rows of 8, resulting in 46080 readout channels per module. The module is glued from this side to the local support structure.

Figure 3.6 shows an exploded view of a Pixel module.

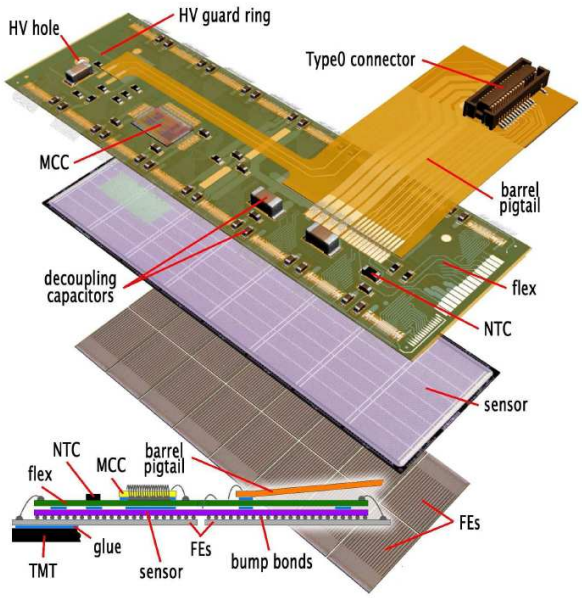

Figure 3.6: *Exploded view of a Pixel module [34].*

The flex is glued to the sensor from the backside of the module, which is the side pointing towards the interaction point for barrel modules. Besides the MCC, a temperature sensor for the monitoring of the module temperature, and passive components like decoupling capacitors are mounted on the flex, and the data transfer lines and the supply lines are routed to the readout electronics. The front-end chips and the MCC are connected to the flex through 25 µm thick wire-bonds.
The connection via the "type 0" cables from the modules to the electrical services is handled differently for barrel and disk modules, which are otherwise identical. For barrel modules, another kapton foil, the pigtail, houses a connector for the type 0 cable. For disk modules, the type 0 cables are directly connected to the flex.

3.2.1 The Pixel Sensor

3.2.1.1 Semiconductor Sensors

The Pixel Detector is a semiconductor detector. Particles are detected by the fact that they leave electron hole pairs along their track in the semiconductor sensors. The generated charge carriers induce a signal in the readout electronics. The energy loss by traversed material length, dE/dx, is described by the Bethe-Bloch formula [35]. The behaviour for silicon is shown in figure 3.7. The minimal energy loss is independent of the particle type, but occurs at different energies for different particles. Particles with that energy are called minimum ionizing particles. For a sensor width of 250 µm, the mean energy is corresponding to a charge of 27000 electron hole pairs, while the most probable energy deposition is corresponding to 19000 electron-hole pairs.

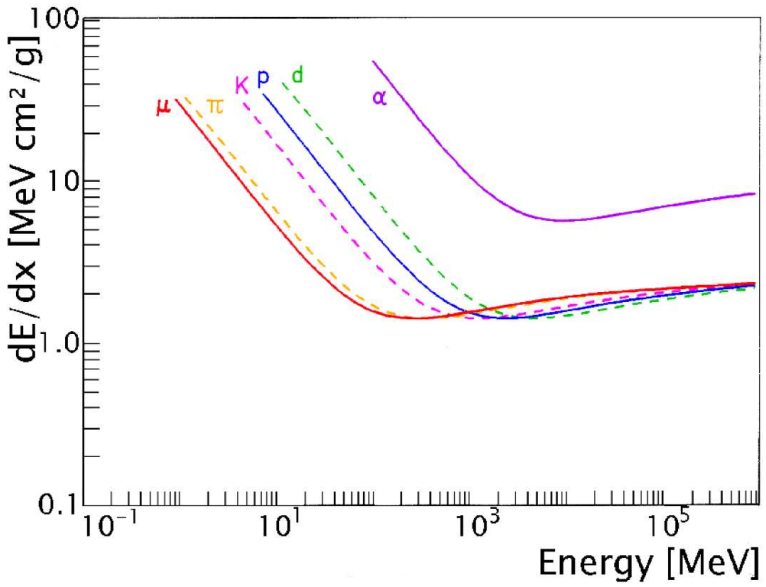

Figure 3.7: *Bethe Bloch formula for Silicon [37].*

3.2.1.2 The Pixel Sensors

Figure 3.8 shows the schematic layout of the Pixel sensors.

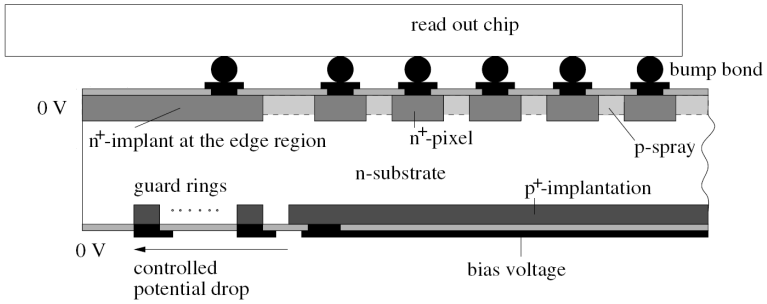

Figure 3.8: *Schematic layout of the Pixel sensors [31].*

The Pixel sensor bulk consists of n doped silicon, with the actual pixel segmentation provided by n^+ implants on the front side of the sensor. The back side is p^+ doped. The sensor has an active area of $60 \cdot 16\,\text{mm}^2$ and a thickness of 250 µm. Most pixels have an area of $50 \times 400\,\mu\text{m}^2$, with the exception of the pixels covering the inter-chip region (section 3.2.2.1), where the pixels have a size of $50 \times 600\,\mu\text{m}^2$.

The pixels are isolated by a moderated p-spray layout, where the surface between pixels is implanted with p doped material with increasing concentration towards the centre of the implant. This has the additional advantage of decreasing the field gradient at the interface between pixel and isolation implants.

With the reverse bias voltage applied at the p^+ side, the depletion zone grows towards the segmented side.

3.2.1.3 Radiation Damage

Irradiation of the sensors will result in the creation of lattice displacements which can create new states in the band structure of the semiconductor. Defects in the band gap will lead to an increased leakage current. Due to permanent capture of free charge carriers, the effective doping concentration changes. This effect can be mitigated by keeping the sensor at higher temperature (beneficial annealing), but only for a limited time. After that, reverse annealing takes place. The time dependence of required depletion voltage and effective doping concentration for a typical temperature profile is illustrated in figure 3.9 for pure silicon and O_2 enriched silicon.

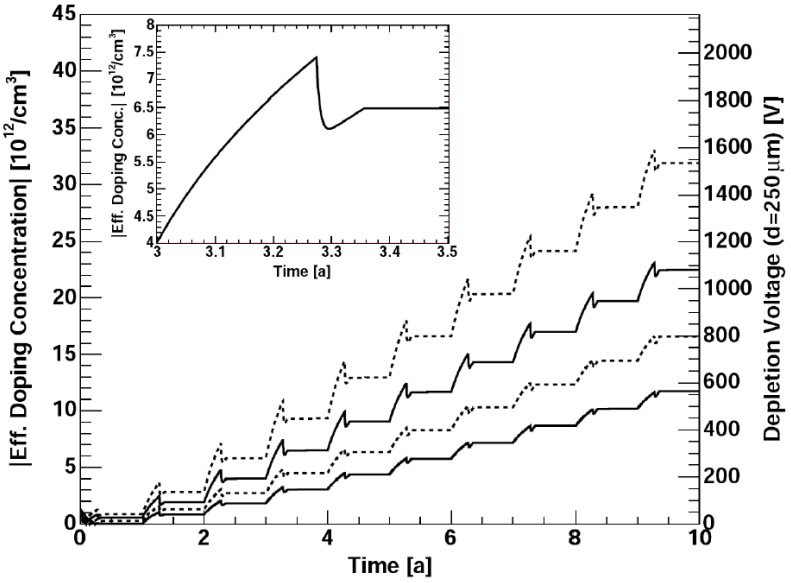

Figure 3.9: *Effective doping concentration for pure and O_2 enriched Silicon, and required depletion voltage for B-Layer and Layer 1 with two different assumptions for the total fluence (from [34]). The inner picture reflects the temperature dependent effects during one year of operation of operation with beam (100 days, 0°C), access (30 days, 20°C), and cooling down during the rest of the year (-20°C).*

After a proton fluence of less than $1 \cdot 10^{14}\ n_{eq}\,cm^{-2}$ the type inversion of the pixel sensor bulk takes place [36]. With the chosen design of the pixel sensor, this results in a "p" bulk sensor, and the depletion zone grows from the n^+ pixel side. This allows an operation even with only partly depleted modules. This fact is depicted in figure 3.10.

3.2.2 Electronics

In this section a short description of the module electronics will be given. This includes the front-end chips and the module control chip. The connection of the module towards the off-detector elements will be described in section 3.3.

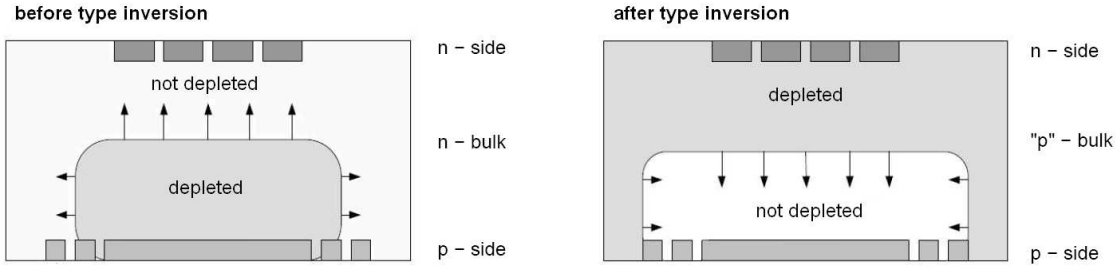

Figure 3.10: *Depletion of the Pixel sensor, before and after type inversion. The chosen design allows an operation even with only partly depleted sensor.*

3.2.2.1 The Front-End Chip

The main part of the front-end chip area is taken by the individual pixel cells with an area of $50 \times 400\ \mu m^2$ each. The pixel cells are addressed by their row and column number. Each front-end chip consists of 18 columns with 160 rows.
A block diagram of the pixel cell is shown in figure 3.11.

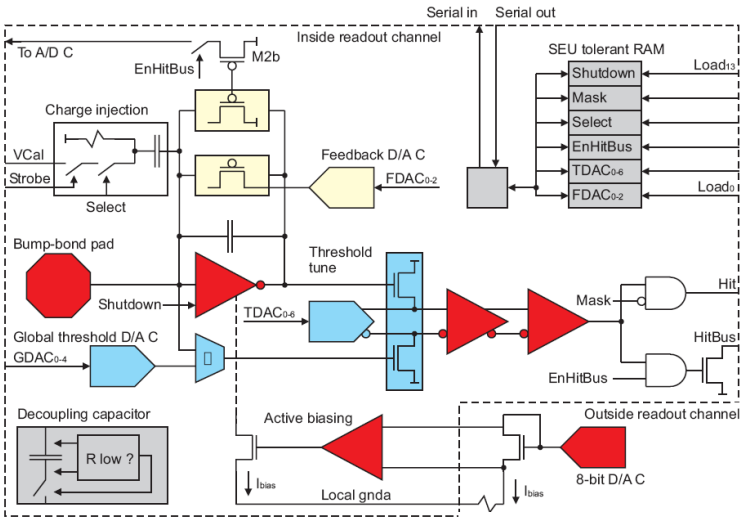

Figure 3.11: *Pixel cell [34].*

The signal charge created by a passing particle is collected at a feedback capacitance of the pre-amplifier. An adjustable constant current source is used to discharge the capacitance. A leakage current compensation circuit can compensate for leakage currents of up to 100 nA. A known charge can be injected before the pre-amplifier to test and calibrate the functionality of the complete analogue circuitry. After the pre-amplifier, the discriminator consists of a second stage of amplification and a differential comparator with a tunable threshold. The time that the signal stays over the threshold ("Time over Threshold", ToT) is used to determine the charge deposited by the incident particle.

The timestamps of leading edge and trailing edge of the discriminator output, used to determine the Time over Threshold, are passed to the digital part of the pixel cell. They are measured in units of the 40 MHz bunch crossing clock. The readout is handled by column-pairs. When the trailing edge of the signal registers, the time information together with the pixel address is passed to the 64 End Of Column (EoC) buffers, where the 8 bit ToT is computed, and the leading edge information is compared to arriving triggers. When a Level-1 trigger arrives, all hits where

the time-stamp of the leading edge, corrected for trigger latency, matches the trigger time, are flagged as belonging to that “L1ID”, and sent to the MCC. It is possible to read out up to 16 consecutive bunches per Level-1 trigger. The relative bunch crossing is given by the “Level-1 Accept”, L1A.

On the level of the pixel cell, 14 bits can be set to configure its behaviour:

- The “Kill” bit is used to enable or disable the pre-amplifier without changing its power consumption.
- The “Mask” bit prevents that the output of the discriminator is passed to the digital part of the pixel cell.
- The “HitBus” can be used to automatically trigger a readout of the front-end when a pixel is hit.
- The “Select” bit marks the corresponding pixel for injection of test pulses.
- The 7 bit TDAC is used to configure the threshold of the discriminator.
- The 3 bit FDAC is used to tune the amplitude of the feedback current.

Additionally, 231 global configuration bits affect all pixel cells of a front-end. These include

- the 5 bit G-DAC, for the tuning of the global threshold setting.
- the 8 bit IF-DAC, for the amplitude of the feedback current.

A distinctive feature of the pre-amplifier is the fact that the position of the peak of the resulting signal is independent of the deposited charge. Therefore, low charge signals are registered later than simultaneous signals with higher charge. This “time-walk” effect is illustrated in figure 3.12. To correct for this effect, it is possible to copy hits below a certain threshold to earlier time bins. As low charge hits are also more probable to be due to noise, the front-end can be configured to ignore these hits.

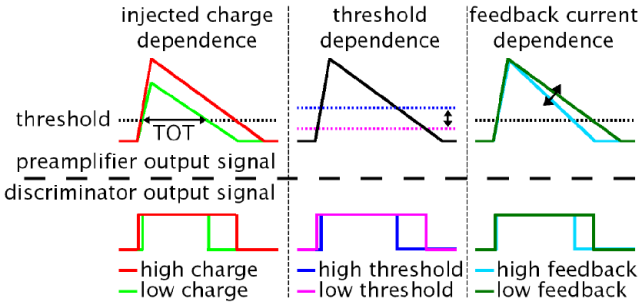

Figure 3.12: *Time walk [37].*

To avoid dead space between the front-end chips, these spaces are filled with 600 µm “long” pixels in z-direction, while in r-φ-direction always two “ganged” pixels are read out together. The actual hit position for the ganged pixels has to be determined by pattern recognition. Therefore, ganged pixels are alternating with individually read out “inter-ganged” pixels. The location and naming of these special pixels is shown in figure 3.13.

3.2.2.2 MCC

The MCC is responsible for the configuration of the front-end chips, the distribution of timing, trigger and command data, and event building from the physics data received from the front-end chips. A schematic diagram of the MCC is shown in figure 3.14.

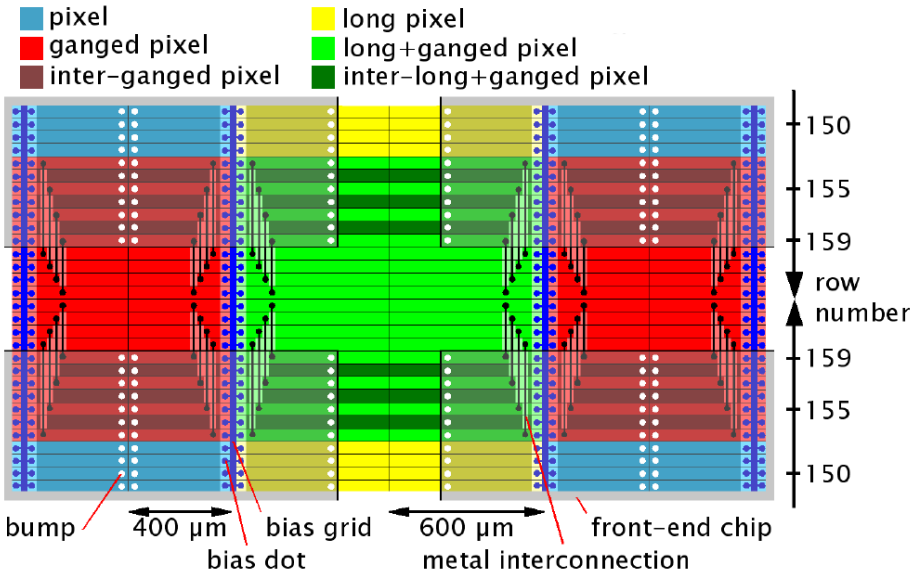

Figure 3.13: *Inter-chip region [37].*

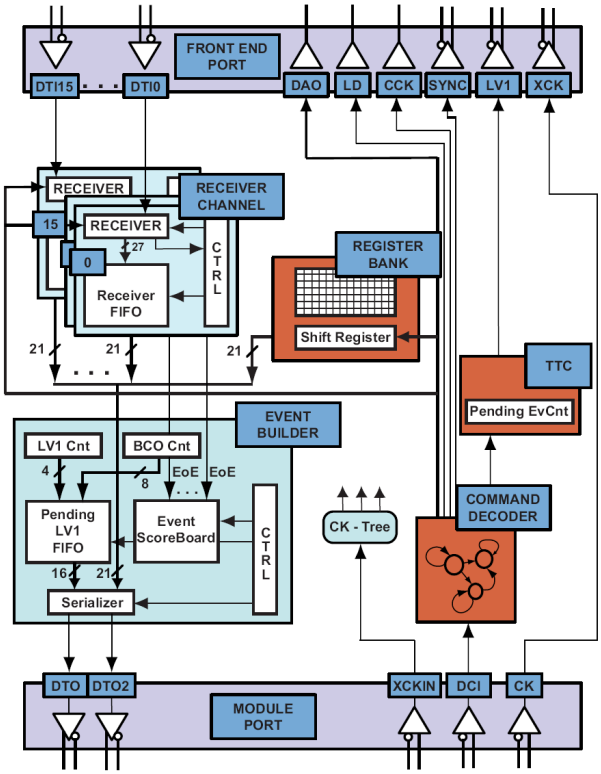

Figure 3.14: *The MCC [34].*

The different tasks correspond to blocks on the MCC. The Module Block represents the interface to the opto-boards (see chapter 3.3), with incoming data, clock and command signals and outgoing data lines. Commands are decoded in the Command Decoder block. The TTC block generates Level-1 triggers for up to 16 concurrently processed events. The Register Bank stores the configuration data for the MCC. The Front-End Block is the interface between

MCC and front-end chips. After a Level-1 trigger is issued to the front-end chips, the corresponding data is sent to the MCC's 16 receiver FIFOs (128×27 bit). Each FIFO is receiving the data from one front-end chip for the up to 16 events. The event builder organizes the hit data to build the event as soon as all information (signalled by the End Of Event Words) is stored in the FIFOs, and sends it out to the off-detector readout electronics.

3.3 Readout

Figure 3.15 shows the readout chain for the Pixel Detector. Optical data transmission was chosen, as it avoids the issue of electromagnetic disturbance of the transferred signals, and, additionally, it provides high bandwidths with low radiation length of the used material. The transformation between electrical and optical signals is handled on-detector by the opto-boards and off-detector on the Back Of Crate (BOC) cards in the counting room. For the electro-optical conversion, Vertical Cavity Surface Emitting Lasers (VCSEL) are used for light emission, and PiN (Positive intrinsic Negative) diodes as light receivers.

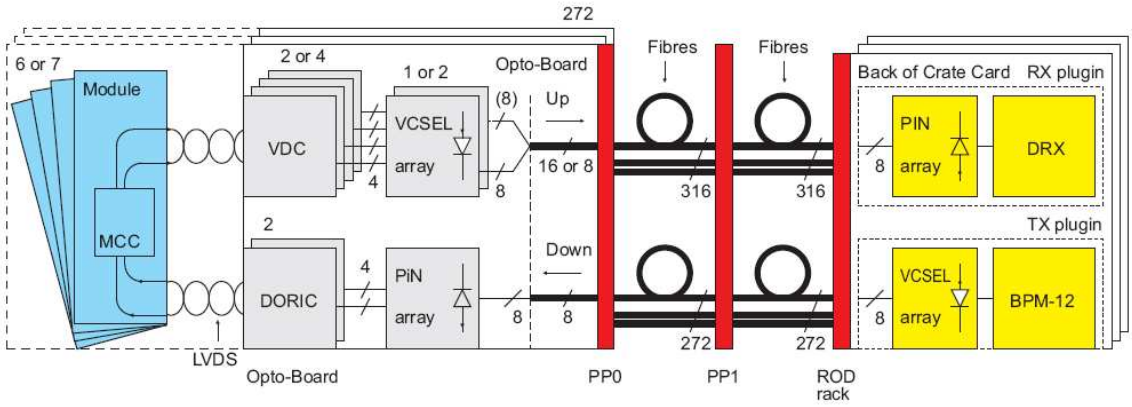

Figure 3.15: *Readout chain of the Pixel Detector [34].*

3.3.1 The Optoboard

Six or seven modules of one disk sector or half-stave are connected to one opto-board. The Timing, Trigger and Control (TTC) signals from the off-detector readout electronics are transferred over eight separate ribbons of optical fibres, and received by an 8-way PiN diode array which converts the laser signals into electrical signals. The DORIC (Digital Opto Receiver Integrated Circuit) chip decodes those signals, and sends them via LVDS (Low Voltage Differential Signal) lines to the modules. The DORIC has four channels, so each opto-board carries two DORICs. For the up-link, the VDC (VCSEL Driver Chip) controls the conversion of electrical signals carrying the physics data in the 8 way VCSEL arrays. Two four-channel VDCs are needed for one VCSEL array. To allow higher readout speed for the innermost detector layer with the highest occupancy, the "B-boards", connected to B-Layer modules, are equipped with two VCSEL arrays and four VDC chips, while all other opto-boards – the "D-boards" – carry only one VCSEL array and two VDC chips.

The opto-board needs three supply voltages (see chapter 5.3.1), VVDC to power VDC and DORIC, VPin for the depletion of the receiving PiN diode, and VISet for emitting laser power. The current measurement at the PiN diode gives information about the total received light in all eight channels. This can be used to verify the connectivity and functionality of the optical link from the counting room to the detector. The corresponding procedure of measuring consecutively all eight channels of one link is known as "Inlink-Scan". At the opto-board, the laser power is tunable by adjusting the corresponding supply voltage. An external reset signal is needed to allow the

DORIC to correctly lock onto the received clock, which is also provided by the power supply crate. The threshold for decoding incoming signals is self-adjusting dynamically, and does not need to be tuned.

272 opto-boards (plus 16 spares) are installed in the Pixel Detector, defining the granularity of one readout unit.

3.3.2 The BOC card

The BOC cards represent the off-detector part of the optical link. They are installed in nine readout crates located in the five DAQ racks in the counting room. Each crate is equipped with up to 16 ROD-BOC pairs, one TIM (TTC Interface Module), and an SBC (Single Board Computer). The BOC cards carry RX- and TX- plugins which handle the electro-optical conversion. The TTC data for the detector is sent by the TX-plugins, which carry a VCSEL array and a BPM-12 chip for the encoding of data and clock in the BPM format (Bi Phase Mark). The laser power is adjustable in 256 steps between 0 and 18 mA. The RX-plugins are equipped with a PiN diode array of eight channels receiving the data from the detector. The DRX (Digital Receiver IC) amplifies the signal and passes them on if they are above a certain threshold. The threshold is one of the tunable parameters off-detector, while the second important parameter is the delay that determines the clock phase at which the arriving signal is sampled. The optimal settings for delay and threshold are depending on the on-detector laser power. They are determined by a procedure known as “2D-” and “3D-BOC-Scan” respectively.

Each BOC can handle the data streams of one to four readout units, depending on the transmission frequency. Data from Layer 2 is sent with a 40 MHz frequency, therefore a maximum of four readout units can be handled by one BOC. Layer 1 and Disks are read out with 80 MHz allowing two readout units per BOC, and for the B-Layer two fibres with 80 MHz each are necessary to enable readout at 160 MHz. Therefore one BOC, equipped with two RX plugins, is needed for each B-Layer readout unit. Disks and Layer 1 require one BOC equipped with two RX plugins to handle two readout units.

The BOC generates, from the ATLAS TTC signal, the clock that is sent to the modules. The delay between clocks is adjustable. The received data, after the event building on the ROD, is sent via the BOC’s S-Link to the ROS (see chapter 2.2.5) of the ATLAS DAQ.

Besides the formatting and event building of the received data, the RODs are responsible for generating commands for the modules, pass trigger and timing, and reset information to the modules as well as the module configurations. The RODs have the capability to allow the monitoring of the incoming data. The RODs have two operation modes, data taking and calibration.

3.4 Overview of the Pixel Detector Services

As the layout of the services, i.e. power cables, optical fibres and cooling pipes, is one of the determining elements for the modularity of the control system, and the corresponding terminology will be used in the following chapters, an overview, mainly with respect to the electrical services, is given in this paragraph.

The power supply system for the Pixel Detector will be described in detail in chapter 5.3.1, and the cooling system in section 3.5.

To supply the detector with all the required supply voltages, the electrical services have to be routed from the counting rooms where the power supplies are located to the Pixel Detector, which is located at the innermost center of ATLAS [39]. Constraints due to exposure to radiation, the desire to minimize material budget, assembly and installation requirements, and available space for the cable routing have to be considered. To account for this, services can be disconnected at breakpoints distributed throughout the cavern, the “patch panels”. If the cable type is changed along the routing, e.g. from lower to higher cable radius or to different modularity, this will happen at the patch panels.

Figure 3.16 shows the location of these break points and the routing of the different cable types between them.

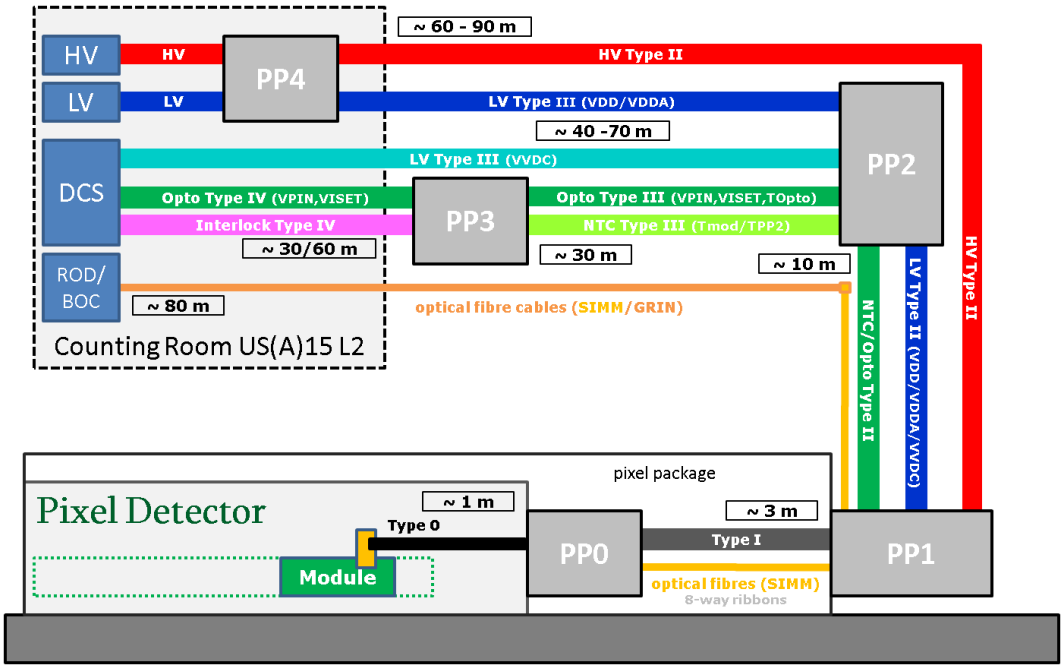

Figure 3.16: *Schematic overview of the patch panel locations and cable routing for the main supply and signal types. As the cable lengths between the patch panels are varying largely by octant, only approximate cable lengths are indicated, to provide an idea of the dimensions.*

The patch panels are numbered from on-detector to off-detector, starting at PP0, where the opto-boards are located. At PP0 the six or seven type 0 cables carrying the module supply lines and data signals are connected, and the electro-optical transformation of data and TTC signals is handled. Bare 8-way-ribbons are carrying the optical signals related to one readout unit from PP0 to PP1, while power lines and temperature information are routed via type I cables. At the end of the service quarter panels, all cables have to pass the airtight barrier sealed by the ID end-plates isolating the Pixel package environmentally and electrically. All services break at PP1, as the Pixel package was the last sub-detector installed, and therefore all cables up to PP1 had already been installed.

Type II cables are connecting to PP1, and here the services are split according to their functionality. The low current opto-board supply voltages (VPin and VISet), the reset line, and the temperature information of the opto-board and modules are carried by the NTC/OPTO type II cables, which are broken passively at PP2 – located inside the barrel Muon chambers – and redistributed to the corresponding type III cables. At the BBIM crates (chapter 5.3.2) at PP3, installed in racks on the service platforms outside the detector, the analogue temperature information is digitized and binary interlock signals are generated from module, opto-board and regulator station temperatures. While the temperatures of modules and regulator stations are directly routed from PP2 to the BBIM crates, the opto-board temperatures have to be extracted from the OPTO type III cables at the Opto-PP3 boxes, where the opto-board supply voltages VISet and VPin are changing in modularity from two to one (seen from off-detector to on-detector), i.e. each OPTO type IV cable carries the supply voltages for two opto-boards from the power supplies (“SC-OLink”, see chapter 5.3.1) in the counting rooms to PP3. The racks where interlock crates and opto-board supply crates are installed are referred to as “DCS racks”.

While NTC/Opto cable routing is split only passively at PP2, the high current supply voltages for the modules and the corresponding opto-board carried by the Low Voltage type II cables are actively regulated at PP2 (see chapter 5.3.1). Besides the need of voltage regulation, PP2 is used as a break point, as the type II cables need to be radiation hard, but it is sufficient to use radiation tolerant cables outwards from PP2. The VVDC supply lines are routed directly between PP2 and the counting room. The module supply voltages are provided by a single power channel each, and therefore are routed through fan-out crates at PP4, located in the counting rooms in direct

proximity to the power supplies. The same is true for the high voltage channels. However, since these are low-current, no active regulation is necessary, allowing to route the HV cables directly from PP1 to the HV-PP4 crates located in the HV-racks in the counting rooms.

The optical data transmission is the only case where the transition from radiation hard SIMM (Stepped Index Multi-Mode) fibres to radiation tolerant GRIN (GRaded INdex Multi-Mode) fibres is not handled at a patch panel. Instead, the fibres are spliced near PP2, but routed directly from PP1 to the readout racks in the USA counting room. At PP1, the transition is made from the bare ribbons with 8 fibres each, to optical cables with 8 ribbons each.

The modularities of the control system are determined to a high degree by the organization of the services at the patch panels, and predominantly by the modularity of Patch Panel 0 located at the detector side of the service quarter panels. On each detector side, one SQP corresponds to one octant. Figure 3.17 shows the layout of the PP0 region.

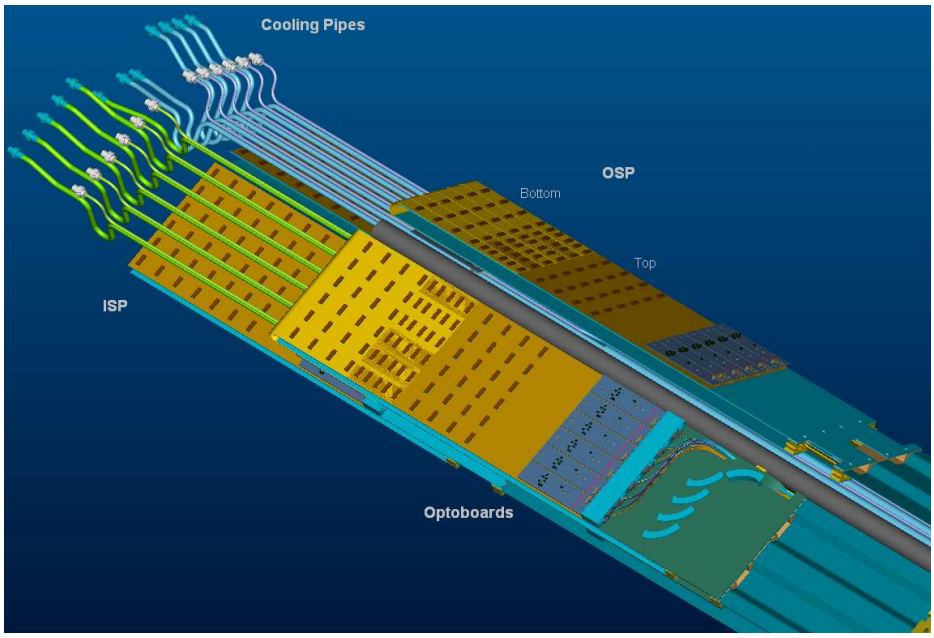

Figure 3.17: *Service Quarter Panel, PP0 Layout [47].*

Each SQP is subdivided into two identical sides connected by a backbone, and each side into six rows, housing the connectors for outer and inner service panel. The outer service panel has opto-boards on top and bottom, so that in each octant 18 opto-boards are installed. The location of the opto-board defines therefore the octant of the readout units with respect to detector geometry. The opto-boards of two interconnected SQP sides are cooled by one dedicated “opto” cooling loop, and on each octant the cooling pipes for six detector cooling loops are routed. It is possible that readout units belonging to one cooling circuit are mounted on separate, not necessarily interconnected, SQPs, therefore belonging to different opto cooling loops. Additionally, the cooling pipes follow a separate modularity defined by which distribution racks they are connected to.

This results in different modularity schemes regarding octants/quadrants:

1. the physical location of the opto-board and services on the SQP.
2. the quadrant of the opto cooling loop determined by the location of the distribution rack.
3. the physical location of the detector cooling pipe.
4. the quadrant of the detector cooling loop.

Due to the fact that all opto-boards of two interconnected service quarter panels belong to one opto cooling loop, items 1 and 2 are identical. Also items 3 and 4 are identical, as all detector cooling pipes that belong to one SQP are routed to the same distribution rack. Figure 3.18 shows the resulting quadrants for opto- and detector cooling loops.

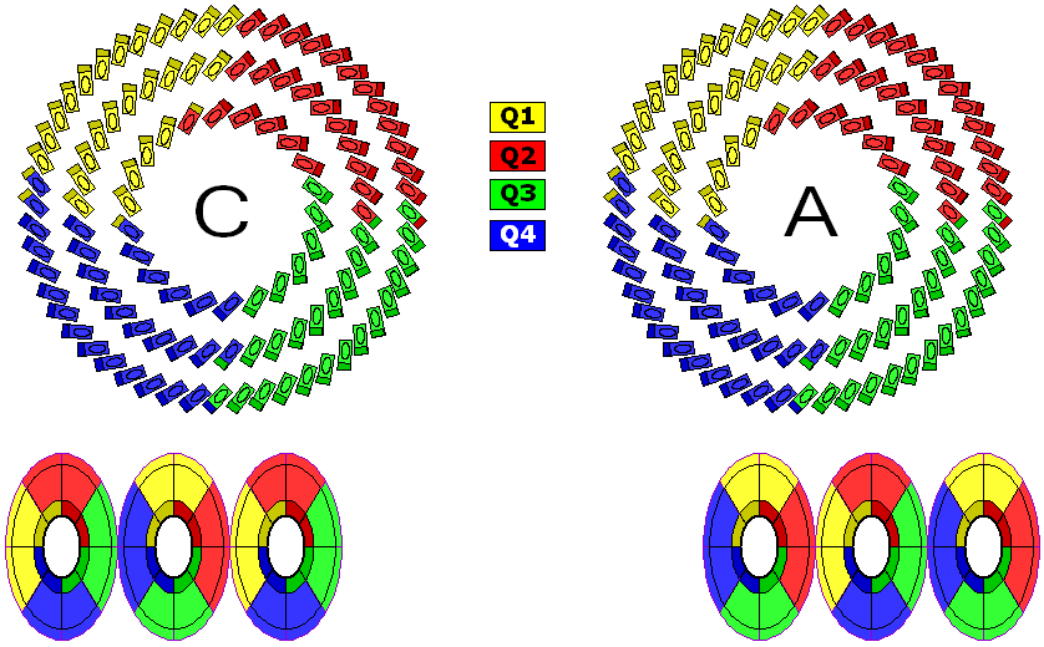

Figure 3.18: *Detector and Opto Cooling Quadrants.*

Additional dependencies in the control system arise e.g. from the fact to which PP2 crate the services are routed, or to which readout rack the optical fibres. Another important modularity for the control system, the interlock quarters, are due to which DCS rack the corresponding services are routed to.

3.5 The Evaporative Cooling System

Pixel modules need to be operated at low temperature to prevent performance decrease due to radiation damage. To remove the foreseen 15 kW of heat generated mainly by the front-end electronics and the increased leakage current in irradiated modules, an evaporative cooling system, based on C_3F_8, is installed [48]. The system has a total cooling capacity of 60 kW.

The Coolant

An evaporative cooling system was chosen, as the exploitation of the phase transition provides a larger cooling capacity per volume compared to mono-phase cooling. This reduces the amount of material to be brought into the inner detector. It enables good heat transfer between liquid and cooling structure and provides a low temperature gradient along the cooling structure. Fluorocarbons are radiation hard, non-toxic, non-flammable, and non-corrosive. C_3F_8 has a saturation vapor pressure which is still above atmospheric pressure for the required -25°C. Operation above atmospheric pressure prevents contamination of the fluid in case of leaks.

Overview of the cooling system

The cooling system can be divided into three functional parts, described in the following sections: The cooling plant and its control, located in the USA 15 area, the distribution racks inside the experimental cavern on the platforms outside the Muon chambers, and the individual cooling loops. It serves the two silicon tracking detectors, SCT and Pixel Detector.

A basic overview of the cooling system layout is shown in figure 3.19.

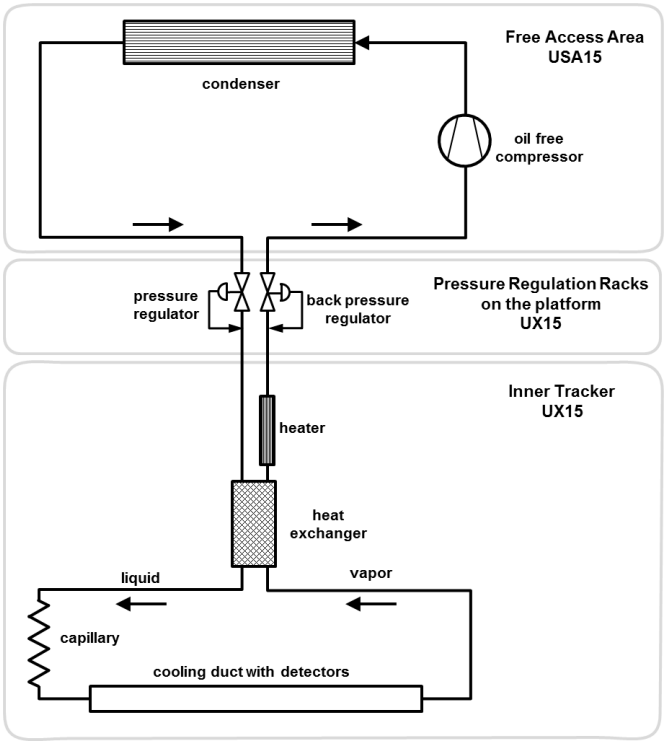

Figure 3.19: *Schematic diagram of the evaporative cooling system [48].*

The cooling plant is located in the USA15 area. From there, liquid is delivered to four distribution racks located inside the cavern outside the Muon chambers. The system is operating with a constant flow, which is regulated by the pressure regulators. The evaporation temperature is determined by the back pressure on each cooling loop. For each loop, one pressure regulator (PR) on the liquid side and one back-pressure regulator (BPR) on the vapor side are located in the distribution racks. The basic layout of the individual loops is the same for all Pixel loops, although they differ in technical details.

All Pixel cooling loops consist of one capillary, a heat exchanger, a heater, the cooling pipe contained in the detector local support structure, and the pipework connecting the loop to the pressure and back pressure regulator in the distribution rack.

After passing the cooling loop, the vapor returns to the distribution racks and to the plant, where the gas is first compressed and then condensed.

The direct control of the system is managed by PLCs (Programmable Logic Controller), while the higher level control and monitoring are integrated into the ATLAS Detector Control System.

The Cooling Plant

Seven compressors are located in USA15 to compress the returning gas from 800 mbar to 17 bar. To keep the pressure at the inlet constant, six compressors are operated in on/off mode depending on load, and only the first has a bypass-line to regulate the flow in a continuous mode. The gas is condensed in a mixed water heat exchanger.

Three control PLCs are also located in the cooling plant area. One PLC is responsible for the control of the plant and the loops, while the second and third PLC are controlling the heater regulation for USA and US side heaters. An additional fourth PLC, the "Communication PLC" is managing the communication to the Detector Control System via the Modbus protocol.

To ensure the recovery of the liquid in case of a power failure, the critical elements (compressor 1, the chiller for mixed water, and the PLCs) are powered by UPS (Uninterruptible Power Supply).

The Distribution Racks

The four distribution racks are located in the cavern on levels 1 and 7 on each side of the detector. They are connected to the plant by four independent pipes, therefore defining four independent quadrants. The pressure and back-pressure regulators for the cooling loops are located at the distribution racks, as well as a pneumatic safety valve on the liquid side and a manual valve in the gas side. Although pressure and back-pressure regulators are pneumatically controlled, they can not be controlled individually. Instead, they are arranged in groups which operate at the same set value. For the Pixel Detector there are two groups of pressure regulators (PR) controlling the flow, and 18 groups of back-pressure regulators (BPR) controlling the temperature. The range for the back-pressure is between 1.7 bar and 6.4 bar corresponding to a temperature between -25°C and 14°C in the two phase liquid. For each loop the pressure reading is available on the liquid and the gas side. The Cooling groups and operational back-pressure set-points are shown in figure 3.20. The resulting module temperature distribution for the barrel is given in figure 3.21. For comparison, the module temperature for un-powered modules without cooling is shown in figure 3.22.

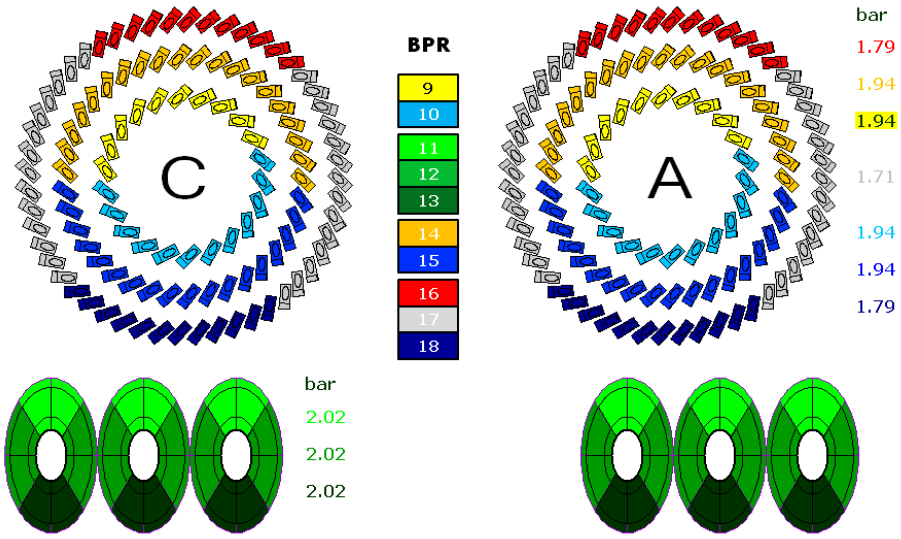

Figure 3.20: *The cooling groups for the 80 Pixel Detector loops. Indicated are also the typical back-pressure value at which the loops were operated in 2010. Layer 2 is operated at the lowest nominal temperature of -24°C, all the disks at -20°C. Each opto-loop is controlled by a separate back-pressure regulator. The opto-loops are operated at 4.15 bar, corresponding to 0°C.*

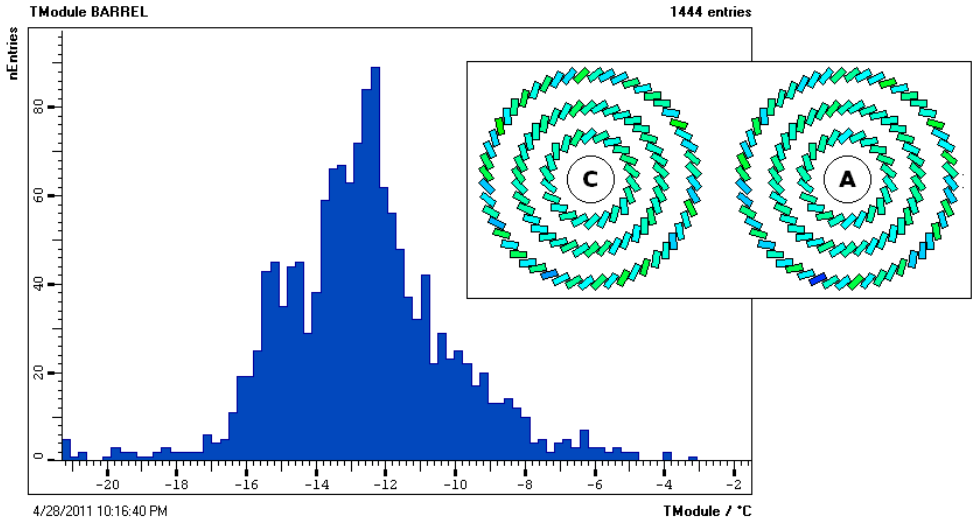

Figure 3.21: *Temperature distribution for barrel modules. In the top right, the average temperature for each readout unit is given. The inhomogeneity in layer 2 is due to the fact that for the inlet staves a smaller diameter pipe was inserted inside the original pipe to recover from corrosion in the original pipes. One readout unit in layer2 (L2_B22_S1_A6) is not operated, resulting in an average module temperature of less than -20° C.*

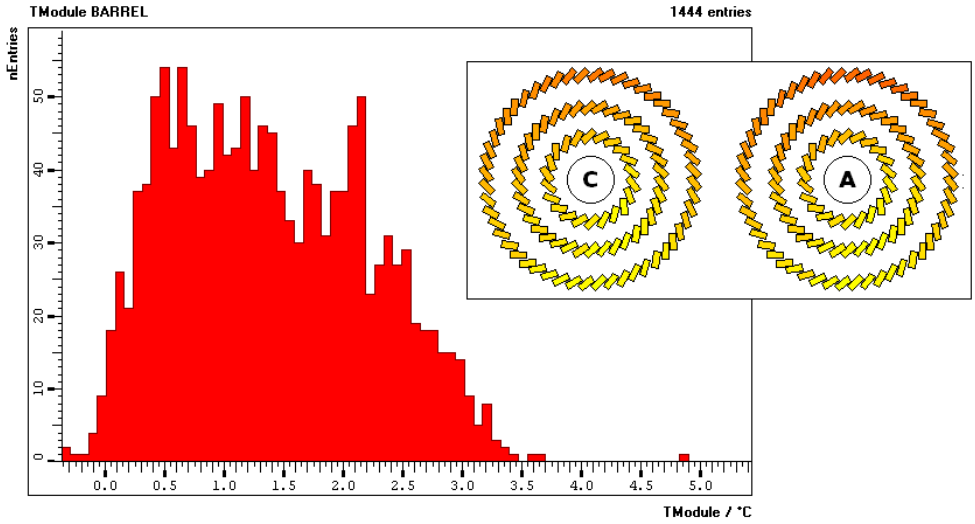

Figure 3.22: *Temperature distribution for un-powered barrel modules during warm-up after a failure of the cooling system.*

The Cooling Loops

The Cooling System consists of 212 individual cooling loops, 88 for the Pixel Detector and 116 for the SCT. Eight of the Pixel loops are dedicated for the cooling of the opto-boards. Each quadrant has two opto-loops, one for the A-side and one for the C-side, and each loop cools 36 opto-boards including spares. 80 loops are dedicated to the cooling of the modules. In the barrel, each loop corresponds to a bi-stave, i.e. 26 modules, and in the disks to a bisector, i.e. 12 modules. Detector loop and opto loop of one readout unit are not necessarily belonging to the same quadrant.
The layout of the cooling loop is shown in figure 3.23.

The incoming liquid is pre-cooled by the returning vapor in a heat-exchanger, located on the Pixel service quarter

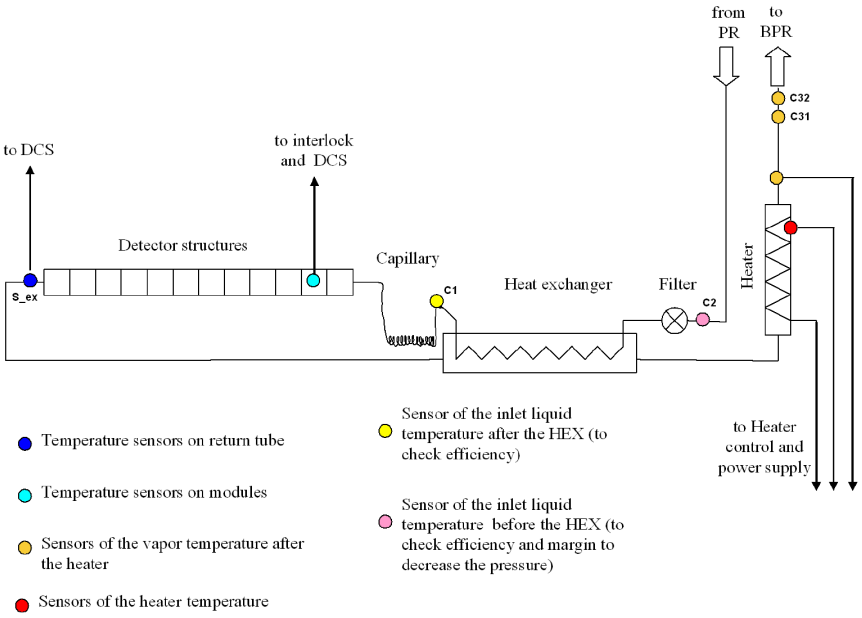

Figure 3.23: *Schematic diagram of a cooling loop [48].*

panels. Passing the capillary, the liquid starts to boil, absorbing the heat from the detector structure. The returning vapor traverses the heat exchanger again, and a heater ensures that all remaining liquid is evaporated and the vapor temperature is above the cavern dew point of 16°C while returning to the distribution racks.

The heat exchanger and the capillary are located inside the Pixel volume, while the heaters are outside the ID-end-plate in the dedicated heater trays. Various sensors for control and monitoring are installed. For Pixel these are

- The "S-Sensors": one NTC per loop monitoring the temperature at the exhaust. For the opto-loops there are three additional sensors located on the heat spreader and upper heat spreader [91]. These NTCs are read out by the Pixel DCS.
- The "C-Sensors" (C1, C2, C31 and C32): the C1-NTC is installed before the capillary but after the heat exchanger, monitoring the inlet temperature. It is read out by Pixel DCS. C2 is installed before the heat exchanger. C31 and C32 are installed after the heater, monitoring the temperature of the returning vapor. All but C1 are monitored by ID DCS.
- A thermocouple (plus one spare) on the exhaust tube a few centimeter after the heater, for the control of the heater temperature. Read out by the PLC.
- A thermocouple (plus one spare) on the heater surface. This serves as input to an interlock box to switch off the heaters if the temperature exceeds 80°C. In addition, the value as well as the interlock state is monitored by an ELMB.

The heater is operated in switching mode by 110 V which are controlled by the PLC using a PID (Proportional Integral Derivative) algorithm [51]. The feedback is provided by the thermocouple located after the heater, but for a few problematic loops this was changed to the C31 sensor. The signal of the thermocouple located on the heater surface is used to generate an interlock signal to temporarily switch off the heater when the temperature exceeds a threshold of 80°C. If the temperature rises to 120 °C, a further interlock (the so-called "super-interlock") cuts

the power to the racks. At high temperature (200°C), the decomposition of the fluid starts. 16 ELMBs are used to monitor interlock temperature and status as well as heater current. The interlock system and corresponding ELMBs for monitoring are located on USA side, the current monitoring is split between USA and US side.

The Control System

The control system for the evaporative cooling system is described in [49], [50]. It provides the interface for the sub-detector DCS projects with respect to the monitoring of the cooling system, and it allows the ID cooling experts to issue commands to the plant by providing a protocol for sending commands to the PLC. The cooling project only connects directly to the communication PLC via Modbus, and all internal data exchange to the three control PLCs is handled by the communication PLC. The high level control and monitoring of the DCS is handled by an FSM according to ATLAS DCS standards (see chapter 2.2.6), but redundant means for sending commands to the PLC are provided. As shown in figure 3.24 the system is organized for the SCT by sub-detector partition and then by quadrants, while for the Pixel Detector the DAQ partitions are not compatible with the cooling quadrants. Background scripts provide additional safety, e.g. by preventing the operation of a loop with a faulty interlock sensor.

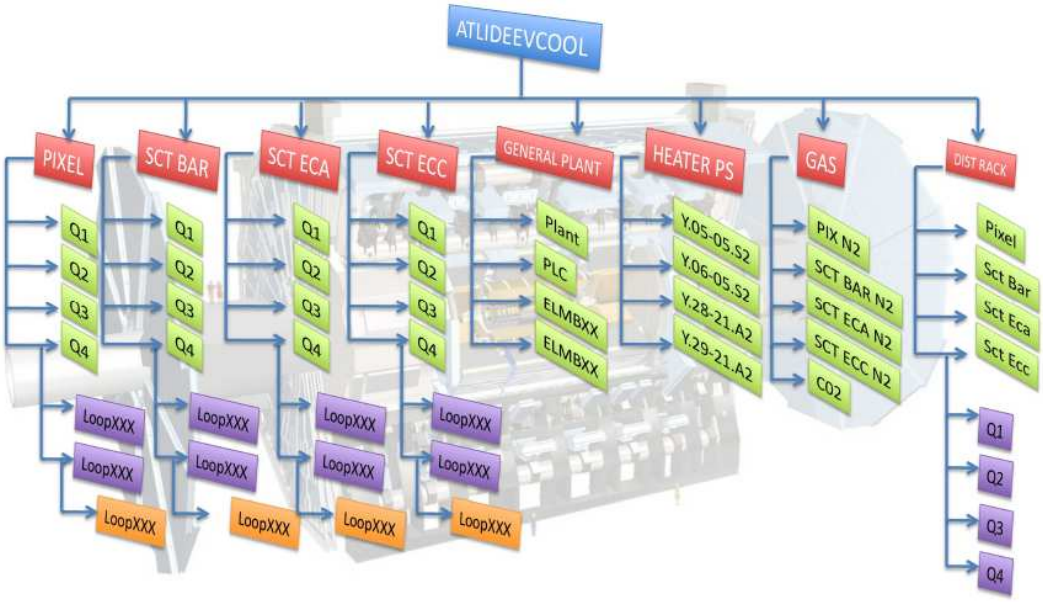

Figure 3.24: *Hierarchy of the Cooling FSM [49].*

3.6 Pixel Detector Performance

The Pixel detector was installed in the ATLAS pit in Summer 2007 and operated the first time in Summer 2008, after the final connection to the services.

After a long period of operation with cosmic rays and an extensive calibration period, the Pixel Detector has been switched on at December 6th of 2009, when LHC provided stable beam conditions for the first time. Since then it routinely takes collision data when LHC delivers stable beams. The following sections give a short overview of the performance of the Pixel Detector [52].

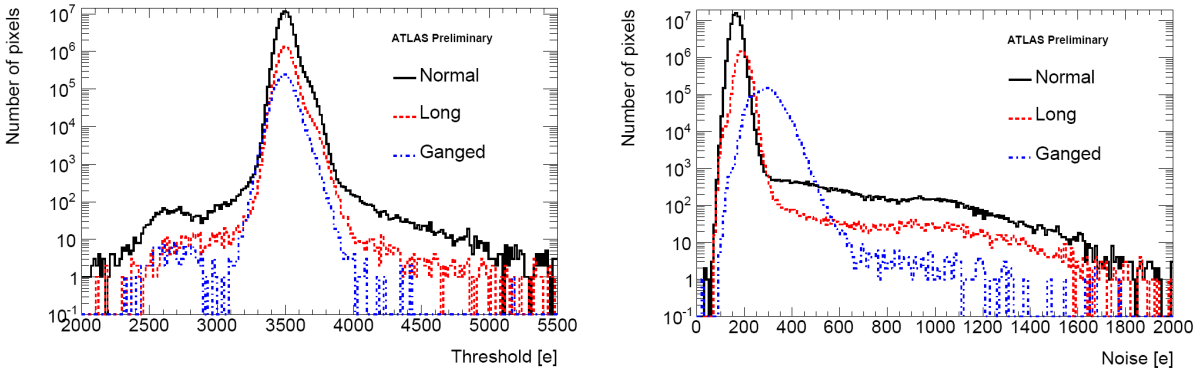

(a) (b)

Figure 3.25: *Pixel threshold (a) and noise (b) distribution. The target threshold is 3500 electrons [52].*

3.6.1 Detector Tuning and Calibration

Optical Scans

Regular checks for the calibration of the Pixel Detector ensure its optimized operation. First, optical communication is established and verified, by the execution of the optical scans. The first scan, the "Inlink-Scan", checks that the off-detector transmitting VCSELs in the BOC crates are operational. The next step is to select a working point for the receiving data-link, i.e. the threshold and delay for the sampling of the incoming data from the opto-boards. For the Inlink-Scan, all 7 laser channels of each off-detector VCSEL-array are consecutively switched on, and the produced current in the on-detector PiN diode is measured. For the 2D-BOC scan, the errors in the optical transmission while scanning through threshold and delay settings, are used to determine the optimal working point.

Module Scans

After the optical communication is established the module front-end settings can be tuned. The tunable parameters are the threshold (set by the 7 bit T-DAC) for the discrimination of signals, with a target threshold of 3500 electrons starting from 2010. The tuning of the Time over Threshold (by the 3 bit F-DAC), for the charge calibration, is performed for the target of 20000 electrons, the most probable charge deposited by a MIP, corresponding to a ToT of 30 bunch crossings. This corresponds to a charge resolution of approximately 660 electrons given by the charge corresponding to a ToT of 1. The threshold is determined by the injection of known charges. For each injected charge, the number of registered hits is counted. While running through different charges, this results in a Gaussian S-curve, where the 50% level gives the threshold, and the noise is determined by the width of the gaussian noise, which is determined from the 30% and 70% s-curve fit. Figure 3.25 shows the threshold and noise distribution for 2010.

3.6.2 Detector Operations

3.6.2.1 Noise Occupancy

In order to minimize the number of noisy pixels, a noise mask is created in special runs without beam using random triggers. Pixels with a noise occupancy (the fraction of measured hits uncorrelated to detector activity) higher than 10^{-5} hits per bunch crossing are masked "online" at the hardware level, excluding those pixels from being read out at all. Thermal cycling of the modules or retuning requires the re-measuring of the noise mask. About 0.01% of the pixels are masked online.

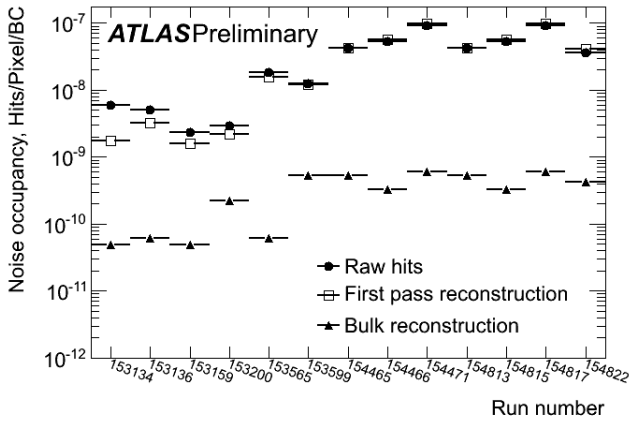

Figure 3.26: *Noise occupancies by run [52].*

In addition to the online noise mask, an "offline" noise mask is determined from reconstructed data during a "calibration loop", and removed for the offline data reconstruction. Figure 3.26 shows the evolution of the noise occupancy during runs taken in April/May 2010. The noise rate of the Pixel Detector is less than 10^{-9} hits bunch crossing per pixel, for the combination of online and offline noise mask.

3.6.2.2 Operation with Beams

During periods with unstable beams, the Pixel Detector must be kept in a safe state. This requires that no high voltage is supplied, and the sensor is not depleted. To prevent large number of noise hits, the pre-amplifiers are kept disabled during these periods. To allow the enabling of the pre-amplifiers during an ongoing data taking run, a warm start procedure has been implemented in the ATLAS DAQ.

When stable beams are declared, the triggers are paused, the high voltage is switched on, and the modules can be re-configured with enabled pre-amplifiers. This leads to an increased data taking efficiency, as the reconfiguration of the Pixel modules during the warm start is much faster than the time it would take to restart the DAQ system. As the time needed for the warm start is the dominating part of the Pixel data taking inefficiency, a gain here is especially relevant.

To recover from problems reducing the data taking efficiency, like high noise rates, hardware problems for the modules or readout link, loss of the clock, or corrupted module configuration, the Pixel Detector DAQ has implemented the possibility to automatically remove and re-include modules from the data taking during a run.

To recognize and fix problems in a timely manner, without waiting for offline reconstruction of the data, online monitoring of the data quality is essential. Events are sampled directly from the data stream, and e.g. occupancy, ToT, or readout errors are monitored on a module level.

3.6.2.3 Timing

In March 2010, at the beginning of the operation at 3.5 TeV beam energy, the Pixel Detector was using a readout window of four consecutive bunch crossings, and the nominal collision event in the second one, i.e. at an L1A of 1, as shown in figure 3.27. Time delay scans performed in spring 2010, to optimize the time delay setting per module, ensured a uniform timing distribution with an average L1A of 1.001 for clusters on track. This made it possible to reduce the readout window to 1 BC in October 2010. Hit-doubling, i.e. copying of hit data to the previous bunch crossing, is activated for hits with a ToT $\leq$ 7 BC, to prevent the loss of hits with low deposited charge due to time-walk (cf. section 3.2.2.1).

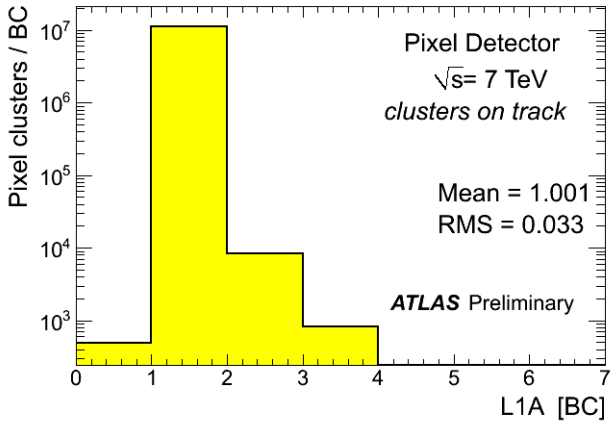

Figure 3.27: *Clusters on track measured in the the Pixel Detector per bunch crossing, with a readout window of 4 bunch crossings [52].*

3.6.3 Data taking and performance

97.3% (1697 out of 1744 modules) have been included in ATLAS data taking during the run periods of 2010 and 2011. For the remaining 47 inactive modules, 28 are located in Layer 2, five in Layer 1, six in the B-Layer, and four in each endcap. Failures are due to open high voltage connections, low voltage shorts, problems in the retrieval of data, no clock, or the failure of an opto-board.

Additionally, 44 front-end chips are removed from data-taking. Recoverable failures are, to the largest part, due to failing off-detector VCSELs.

The Pixel Detector has been included in ATLAS data taking with beam since the first stable beam period in December 2009. The instantaneous luminosity provided by the LHC has now surpassed 10^{33} cm^{-2} s^{-1}, and an integrated luminosity of over 500 pb^{-1} has been recorded by ATLAS during stable beam conditions. In tis time, the Pixel data taking efficiency has been improved to currently around 99%.

Chapter 4

The ATLAS Beam Conditions Monitor

The main task of the ATLAS Beam Conditions Monitor (BCM) [53] is the detection of beam instabilities and to trigger a beam abort if necessary to prevent damage to the sensitive silicon detectors. The radiation hard diamond sensors of the BCM are placed near the beam pipe inside the inner detector, where the LHC does not have its own beam monitoring capabilities. BCM provides bunch-by-bunch coincidence information with an excellent time resolution. It also provides a complementary luminosity measurement.

4.1 Overview

Figure 4.1 shows the layout of the BCM Detector. It consists of four pCVD (polychristalline Chemical Vapor Deposition) diamond modules on each side of the interaction point, mounted on the beam pipe support structure (BPSS) at z = 184 cm from the interaction point and 5.5 cm from the beam axis at $\varphi = 0°$, 90°, 180° and 270°, corresponding to $\eta = 4.2$. To reach single MIP sensitivity, the modules are inclined by 45° with respect to the beam axis. Due to the fast electronics a precise measurement of time of flight difference between the two sides is possible. This measurement principle is used to distinguish collision and background events. Particles originating from the interaction point hit both stations at the same time, so the resulting time of flight difference is zero. Particles created upstream or downstream of ATLAS due to beam background or in the worst case by beam scraping the TAS collimators, pass first one station and after 12.5 ns the station on the other side, giving a time difference of 12.5 ns, resulting from the time of flight needed for the distance of 1.84 m between the two stations. This setup was chosen to optimise the capability to distinguish between collision and background events at the minimal bunch distance of 25 ns in the LHC.

The expected occupancy from minimum bias interactions is 0.5 particles / cm^2 of sensor for each bunch at LHC design luminosity.

4.2 The BCM Modules

The sensor material of the BCM modules is pCVD diamond with Ti-Pt-Au electrodes on each side. Figure 4.2 shows the basic principle of the diamond sensor and a photo of an actual BCM module. At the passage of a charged particle through the sensor, a track of electrons and holes will be created, on average 36 electron-hole-pairs per micrometer. Their drift through the applied field induces current pulses that are collected at the electrodes. The high drift velocity (> 10^7 cm/s) and short charge life time in the material gives high and narrow signals which are a requirement for the bunch-by-bunch measurement of the BCM with the bunch spacing of 25 ns. With a signal rise time of 1 ns, signal width of 3 ns and baseline restoration of 10 ns this is well ensured. The used material is

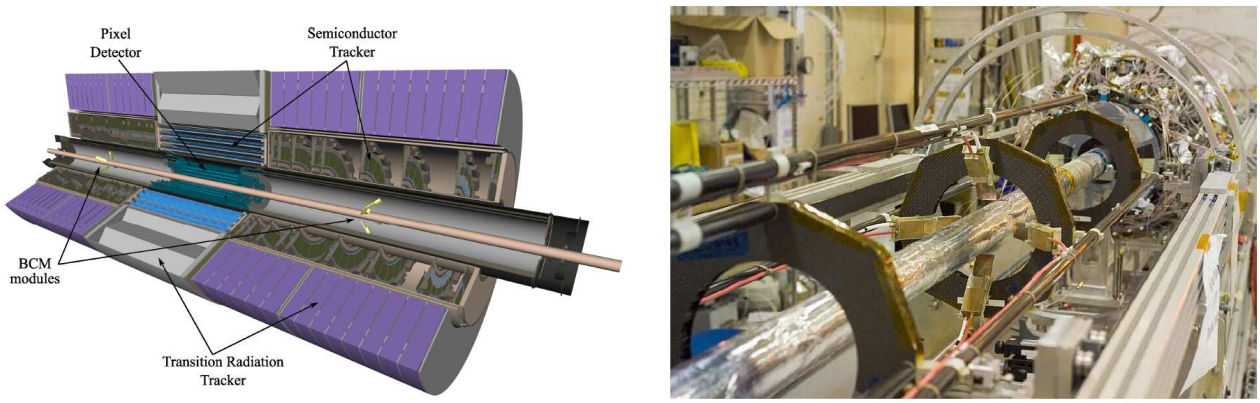

Figure 4.1: *ATLAS Beam Conditions Monitor. Left: Schematic view showing the position of the BCM inside the inner detector. Right: photograph of 4 C-side modules. [21] and [56]*

radiation hard to the required fluences for the experiment lifetime with order of 10^{15} particles /cm and a dose of 500 kGy over ten years.

The size of the sensors is 1 cm^2 with an electrode area of $0.8 \cdot 0.8\ cm^2$. Each module contains two 500 µm thick sensors that are mounted back to back, with the high voltage applied on the outsides, and the readout in the middle. The signals are merged to get a better signal to noise ratio with an achieved ratio of 7-10. The upper sensor is mounted slightly staggered to enlarge the overlap region of the two sensors in the beam direction. With an applied field of 1000 V, a charge collection distance of up to 250 µm can be achieved. The extremely low leakage current – below 1 nA allows operation without the need of external cooling.

Figure 4.2: *BCM Module [53]*

The merged signal of the two sensors is routed through a two stage RF current amplifier[1)] with an amplification of about 20 dB in each stage. The size of the module of 14 cm makes it possible to place the frontend electronics at a reduced radiation level compared to the sensors. The module and the compartments for the two amplifier stages are shielded by copper foil against radio frequency (RF). The analog signal is then transmitted over 15 m to a lower radiaton zone outside the calorimeters.

4.3 Readout and Control

Figure 4.3 shows the Readout and Control Chain of the Beam Conditions Monitor.

The analog signals from the modules are transferred over 15 m to the NINO boards [55]. The signals are split by a voltage divider and fed into two of the 8 input channels of the NINO ASIC (Application-Specific Integrated

[1)] Agilent MGA-62563 GaAs MMIC and Gali 52 In-Ga-P HBT

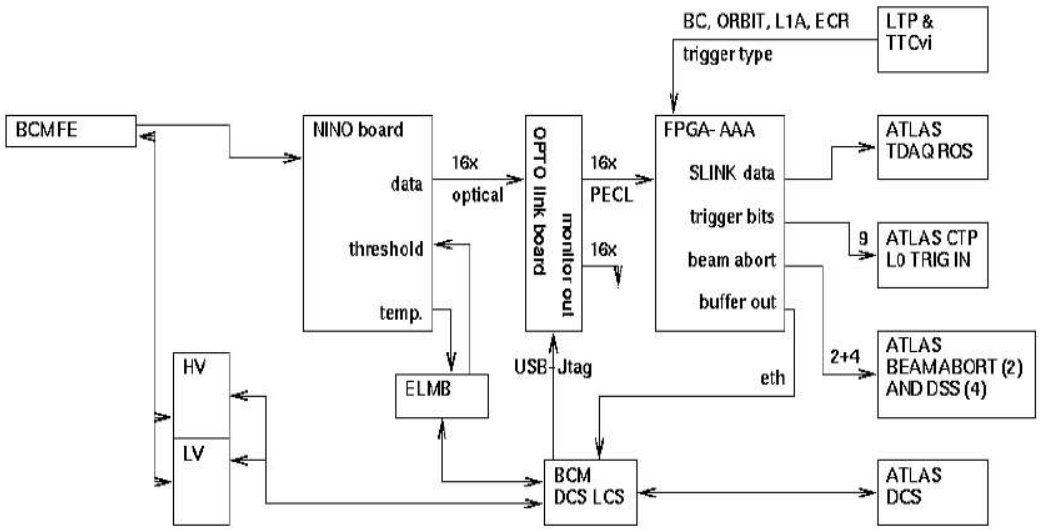

Figure 4.3: *The BCM Readout chain [53].*

Circuit). These channels are referred to as “High Gain” and “Low Gain”. The different amplification makes it possible to use the same signal for different purposes, thus effectively increasing the dynamic range of the NINO, namely the determination of rate and luminosity with the knowledge of all hits, and on the other side the beam abort decision only taking into account higher thresholds for the deposited charge in the sensor. The analog amplitude specifying the Time over Threshold of the signal is converted into a digital pulse. The generated pulse length is proportional to the original ToT-information. The NINO output is transformed into optical signals and transferred via 70 m of optical fiber to one of the two opto receiver board located in the USA15 counting room where they are converted to electrical differential PECL (Positive Emitter-Coupled Logic) signals.

The electrical signals are sent to an FPGA board. The modules are connected to the FPGAs in a scheme to ensure a maximum level of redundancy shown in figure 4.4. The signals are sampled at a speed of 2.56 GHz resulting in a typical time unit of 390 ps/bit for the “time-to-digital” conversion. The boards are equipped with RAM sufficient to store the data of the last 1177 orbits for post mortem analysis in case of a beam dump, triggered by BCM or any other user.

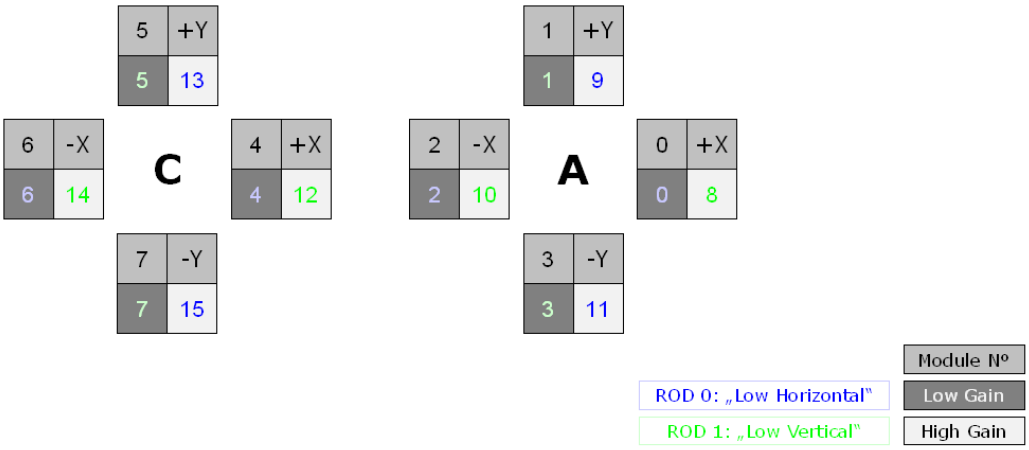

Figure 4.4: *Connection of BCM channels to FPGAs.*

Besides receiving the data, the FPGA board acts as a ROD, i.e. it receives the TTC signals from the Central Trigger Processor (CTP), and sends the data to the ATLAS ROS. Besides this, there are outputs to five other systems:

- ATLAS DAQ
- ATLAS CTP
- ATLAS BIS
- ATLAS DSS
- BCM DCS.

The BCM provides an input for the ATLAS CTP. To be able to provide a trigger input, the information must be provided within 1.5 µs after the actual collision. For this, the FPGA does a basic coincidence determination, e.g. it is possible to trigger on background events in A to C or C to A direction as well as on multiplicity-based events. BCM provides Beam Permit and Injection Permit to the Beam Interlock System (BIS) and Injection Permit System. The Beam Permit condition is evaluated for each ROD separately, and the condition can be programmed by logical requirements. In the initial running phase, this was given by a simple "3+3" condition: the beam permit is removed when both RODs register more than 3 out of 4 low gain hits per bunch crossing. This scheme implies that the beam permit can only be reliably removed when BCM uses the LHC clock. Due to the high sensitivity, the abort condition was changed in 2010 to a "x out of y" scheme, requiring x abort events in y bunches per FPGA for a beam abort. For a beam abort, both FPGAs must drop the beam permit, and the BCM input to the ATLAS BIS must not be masked.
BCM provides inputs to the ATLAS DSS system, two warning and alarm signals for each ROD, and can trigger two DSS alarms: AL_BCM_BeamConditionWarning and AL_BCM_BeamConditionAlarm. This could serve at some point to trigger a high voltage interlock for the ATLAS silicon detectors, however currently no action is defined.
An average rate is transferred via ethernet to the BCM detector control system, i.e. the number of hits per second for each channel. While the low time resolution does not allow monitoring on an event by event basis, the advantage is the independence from the trigger decision and from the availability of the general DAQ system.

In the counting room in USA15 a High Voltage crate is providing the high voltage for the eight BCM modules. A modified SCT supply crate provides the low voltage for the front-end amplifiers and monitors the module temperatures. High and low voltage are routed together in eight cables from the counting room to each of the modules. The hit rate per second is transferred to DCS for each channel, providing information independent of DAQ status and trigger setup.

Chapter 5

The Detector Control System of the ATLAS Pixel Detector

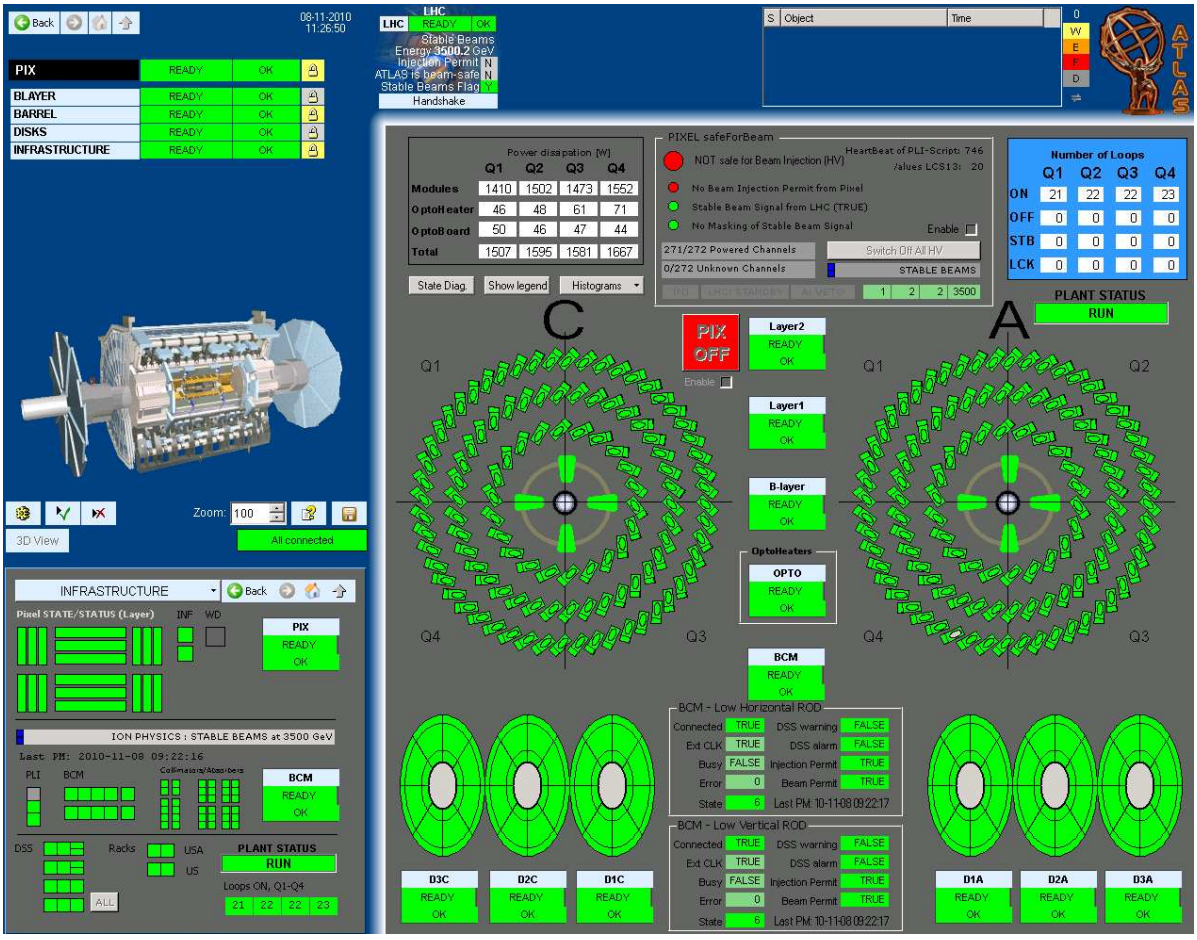

Figure 5.1: *Main User Interface for the Pixel Detector Control System.*

The Pixel Detector Control System is providing the basis for the operation and control of the sub-detector, and in the end allows it to fulfill its task inside ATLAS during data taking. At the same time it has to ensure the safety of the sub-detector.

The control system is responsible for the control of the power supplies providing the required operation voltages needed by the detector hardware, in particular the modules and the opto-boards. It provides the means to configure the set-points, and simultaneously monitor the actual values. In addition to the control of the power supplies, monitoring of the detector's environmental conditions is one of the tasks of the Pixel Detector Control System. To operate the detector by shift personnel, a high degree of automatic procedures and abstraction is required, as well as a clearly arranged and organized way to display the data to the operator. Without this, the large number of channels would not be manageable.

An independent hardware interlock system guarantees the safety of the detector, even in case the software-based control mechanisms would not be available.

To enable a posteriori analysis of error situations or to understand the conditions during data taking, the parameters and conditions have to be stored. For this, an interface to a database has to be provided.

This chapter describes the hardware and software of the control system of the ATLAS Pixel Detector.

5.1 Overview of a Control System

A generic – but with emphasis on the relevance to the Pixel DCS – control system is shown in figure 5.2.

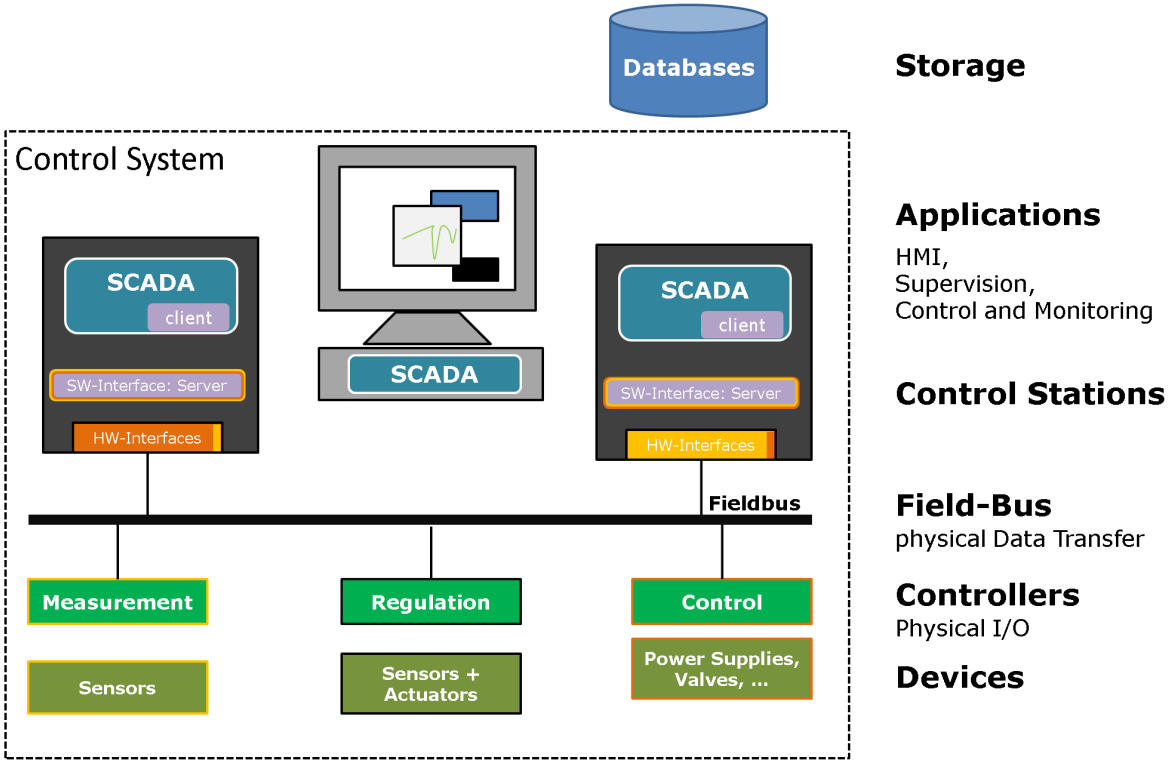

Figure 5.2: *Schematic view of a generic control system.*

The basic hardware components can be roughly divided by their task into measurement, control and regulation type. These hardware components are the physical devices like sensors, valves, or power supplies, and a local instance that is capable to transmit analog measurements converted into digital data, or to receive and translate commands into analog actions. Local intelligence can be utilized for device types, like an automated regulation, e.g. for flow or temperature control. There is however no sharp distinction between the categories, as control

devices often also provide a measurement of the controlled value, and regulation type devices often provide also the monitoring of the regulated value and remote control of the basic regulation parameters.

The Pixel Detector Control System consists of temperature and humidity sensors for the monitoring of the environmental conditions (sections 5.3.2.1, 5.3.2.2), the power supply system (sections 5.3.1), and the interlock system (section 5.3.3) which controls the outputs of the power supply system based mainly on monitored temperatures. Examples for automated regulation can be found in the Inner Detector System, e.g. the control of the thermal enclosure heaters, or the heaters in the evaporative cooling system (chapter 3.5).
The devices used for the physical I/O are ELMBs (Embedded Local Monitor Board) (section 5.2.3) where non-commercial solutions are applied.

The physical transfer of the data to the remote control stations is predominantly handled by bus systems. For Pixel DCS this is CAN (Controller Area Network) (section 5.2.2), but also Ethernet is used.

For higher level control and monitoring, standard PCs are used. They can provide an overall supervision and the human machine interface. For this end, the data must be made available to the final application. Therefore hardware interfaces to the bus are required, and a software interface to make this data available to the Control Software or SCADA (Supervisory Control and Data Acquisition) system. The software interface, e.g. OPC [1)] (section 5.2.4), is often based on a client server architecture, and it is desirable to achieve a degree of independence from actual hardware type and vendor specific implementation. In this case only a generic client is required to enable data access, independent of the actual hardware. As the development environment for the SCADA, PVSS II (ProzessVisualisierungs- und SteuerungsSystem), produced by the Austrian company ETM[2)] is used for all LHC experiments.

The final visualisation of process parameters for the operator does not need to be located where the actual hardware interfaces are installed, but can be run on dedicated PCs (section 5.3.5), accessing the data e.g. by TCP-IP. For the task of high level visualisation and control, e.g. an FSM can be used.

Finally, for offline analysis of historical data, values are transferred and stored to databases which can or can not be embedded in the SCADA itself. For ATLAS, Oracle databases[3)] (section 5.4.6) are employed.

5.2 Common DCS Components

5.2.1 The Joint Controls Project

The Joint Controls Project (JCOP) was set up in 1998 between the four LHC experiments and the CERN IT division, with the goal to reduce the manpower needed for development and maintenance by selecting commonly used interfaces and employing a central development of control software [57]. This includes CAN as one of three recommended options for the field bus and CANopen as protocol as well as OPC as standard software interface. Also the common SCADA system (Supervisory Control And Data Acquisition), PVSS, was selected by a JCOP task force.

5.2.2 CAN

The Controller Area Network (CAN) [59] is used for communication for the majority of the Pixel Detector devices. It is a serial bus that is insensitive to magnetic fields and offers good industry support, high flexibility, error detection and some recovery capabilities. The CAN protocol was developed for the automotive industry and is very reliable and robust. The CAN network comprises multiple nodes with a host processor, a CAN controller,

[1)] formerly **O**LE for **P**rocess **C**ontrol
[2)] Elektrotechnik Mühlgassner, Eisenstadt, Austria (since 2007 subsidiary of the Siemens AG)
[3)] Oracle Corporation, Redwood Shores, CA, US

and a CAN transceiver each. All nodes are constantly monitoring the bus and can send and receive all messages. Bus arbitration is done by Node ID, where each node sends its messages bit wise until one node sends a higher prioritized bit. This node will then continue to send its message without any introduced delay, while the lower priority node will stop sending and retry afterwards.

For error detection and error avoidance multiple mechanisms are employed: the usage of check-sums, introduction and checking of stuffing bits into the NRZ (Non-Return-to-Zero) signal, checking of data frame format, and acknowledgement of received messages. In case of an error detection, error frames are sent.

The data transmission rate is restricted by the cable length. Table 5.1 shows nominal values calculated on the basis of signal transit times. In the Pixel DCS with cable lengths of up to 140 m (from the US counting room) and up to 32 nodes per bus, a bit rate of 125 kbit/sec is used, to allow for an additional safety margin.

Bit rate	cable length
10 kbits/s	6.7 km
20 kbits/s	3.3 km
50 kbits/s	1.3 km
125 kbits/s	530 m
250 kbits/s	270 m
500 kbits/s	130 m
1 Mbits/s	40 m

Table 5.1: *Transmission rate and cable lengths on a CAN bus [58].*

The higher layer communication protocol used in the Pixel DCS is mostly CANopen [59] [60].

5.2.3 ELMB

The Embedded Local Monitor Board (ELMB) [61] was developed by the ATLAS central DCS group in collaboration with NIKHEF[4)] as a flexible low cost, radiation tolerant I/O device with a high number of channels. Communication to the control system is handled via CAN. Figure 5.3 shows a block diagram of the ELMB. It consists of a digital part, the CAN part, and an optional analog part, all of which are galvanically isolated by means of opto-couplers. Voltage regulators are used to adjust the supply voltage to the required voltage for each component, and provide current limitation for protection.

The digital part of the ELMB consists of the Atmel ATmega128 processor, the Infineon SAE 81C91 CAN controller, and digital I/O capabilities. The ELMB provides three default digital ports with eight individual channels each, one input ("Port F"), one output ("Port C"), and one ("Port A") which is configurable by the user either as an in- or output. The input voltage for the digital part ranges from 3.5 to 12 V and is internally regulated to 3.3 V.

The CAN part comprises the CAN transceiver chip PCA82C251 by Philips, and is operated at 5 V which can be provided as 8 to 12 V to the corresponding regulator.

Finally, the "analog" part of the ELMB provides a 16 (+7) bit ADC with four physical differential inputs, multiplexed to 64 channels. The required voltages regulator for the ADC provides ± 5 V. It can be supplied by an input voltage of 8 to 12 V.

The node ID and CAN bit rate for the ELMB is user configurable via DIP-switches. The sampling rate of the ADC (2 Hz to 100 Hz) and its range (25 mV to 5 V, bipolar or unipolar) are remotely configurable. They are set to 15 Hz and 5 V for all Pixel ELMBs. An SPI (Serial Peripheral Interface) interface enables communication to additional components. The ELMBs are provided to the sub-detectors ready to use with a default firmware [62], but can be adapted by the user.

[4)] Nationaal instituut voor subatomaire fysica, Nikhef, previously **N**ationaal **I**nstituut voor **K**ernfysica en **H**oge **E**nergie-Fysica (National Institute for Nuclear Physics and High-Energy Physics, Netherlands).

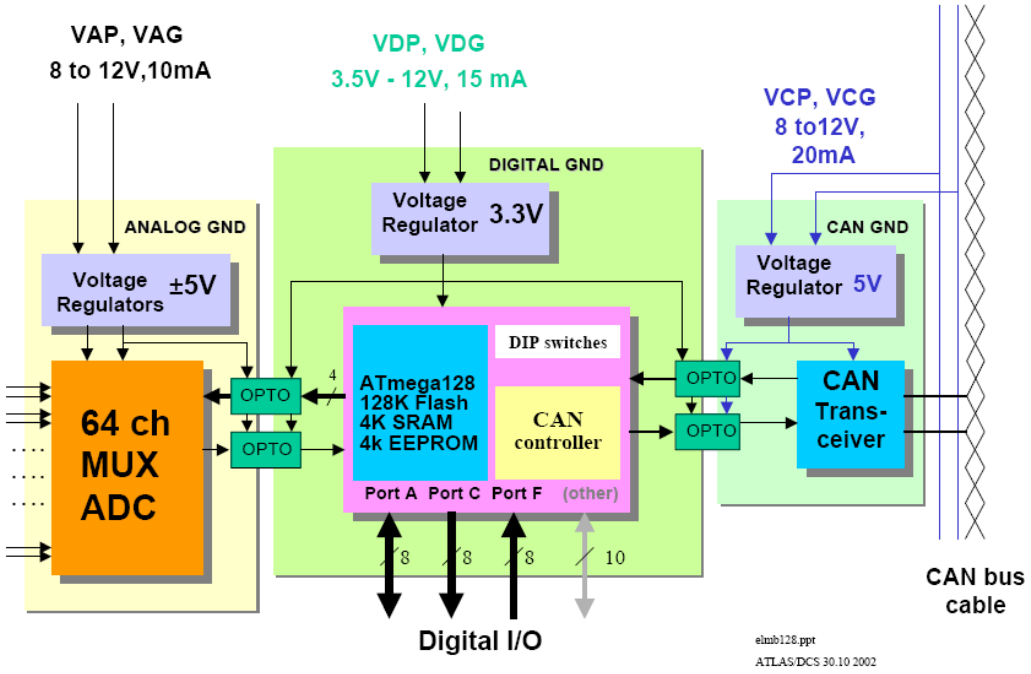

Figure 5.3: *Block Diagram of the ELMB [61].*

CAN-PSU

For the ELMBs a Power Supply Unit (CAN-PSU, Fig. 5.4) was developed by the central CERN PH-ESS group [63]. Internally this crate also makes use of an ELMB itself to monitor (analog channels) and control (on/off via digital ports) the fixed output voltages. Therefore one crate is able to power the CAN-buses as well as the analog and digital part of up to two times 8 buses. The control software for these CAN-PSU crates is also centrally provided [64].

Due to their preeminent importance for the communication to all ELMB nodes, the CAN-PSUs in use for the Pixel Detector Control System are powered by UPS.

5.2.4 OPC

The software interface which is used for nearly all communication between the components of the Pixel DCS is OPC. It is a software standard for vendor independent communication to hardware devices. The original standard OPC-DA (OPC Data Access) is based on Microsoft's OLE (Object Linking and Embedding), COM (Component Object Model) and DCOM (Distributed Component Object Model) technology. The advantage of OPC is the fact that each vendor provides an OPC server optimized for his hardware devices. For the applications it is only necessary to have one type of OPC client to access the data from all different servers. Due to the fact that the OPC standard is based on Microsoft technology this imposes constraints to the choice of the operation system for the PC where OPC server, the client, and the actual hardware interface are installed. The actual control application can however run within these limitations on any other PC.

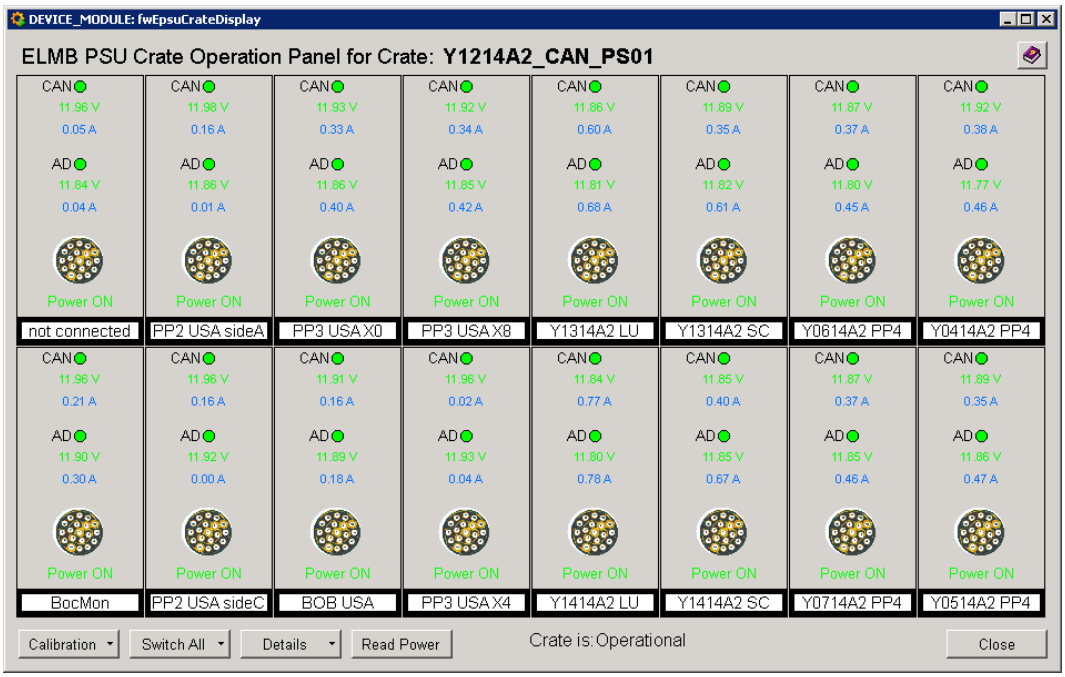

Figure 5.4: *Control Panel for one CAN-PSU crate illustrating its usage.*

5.2.5 PVSS

PVSS II was chosen by the JCOP group as the common SCADA system for the development of the control systems for the LHC experiments. It was chosen after an extensive selection process for its scalability, platform independence, openness, performance, adaptability, and support of many industrial standards among other things.

In this chapter the general concepts of PVSS and its distinctive advantages will be described. Also some terminology of the technical implementations will be introduced, which will be needed for the discussion of implementations specific to the Pixel DCS.

System and Manager Concept

A main concept of PVSS is to divide a system into processes, the “managers”, which are responsible for specific tasks. This is shown in figure 5.5. The internal communication between the processes is handled via TCP/IP. The managers can be grouped by their tasks into four layers. The drivers are located in the process connection layer responsible for the communication with the hardware. For the Pixel DCS this layer consists almost exclusively of OPC clients. The second layer is responsible for internal communication, alarming mechanisms, and history. Its central component is the event manager which is the core of each system. It receives and distributes the data collected from the hardware by the driver layer, and holds an image of all process variables of the system. The data manager provides the interface to the data base. This layer also handles the connections to other systems via the distribution manager. Archiving can be done either by use of the PVSS specific value archives which are stored locally on disk, or, as in case in the ATLAS experiment, by use of an external relational Oracle database for the storage of conditions data. The relevant manager is the RDB manager.

In the third layer, the Control managers allow user defined scripts to run and operate in the background. Also, external applications can be integrated by API (Application Programming Interface) managers.

The final layer is the actual interface to the users, where predefined “panels” can be used to monitor and interact with the system.

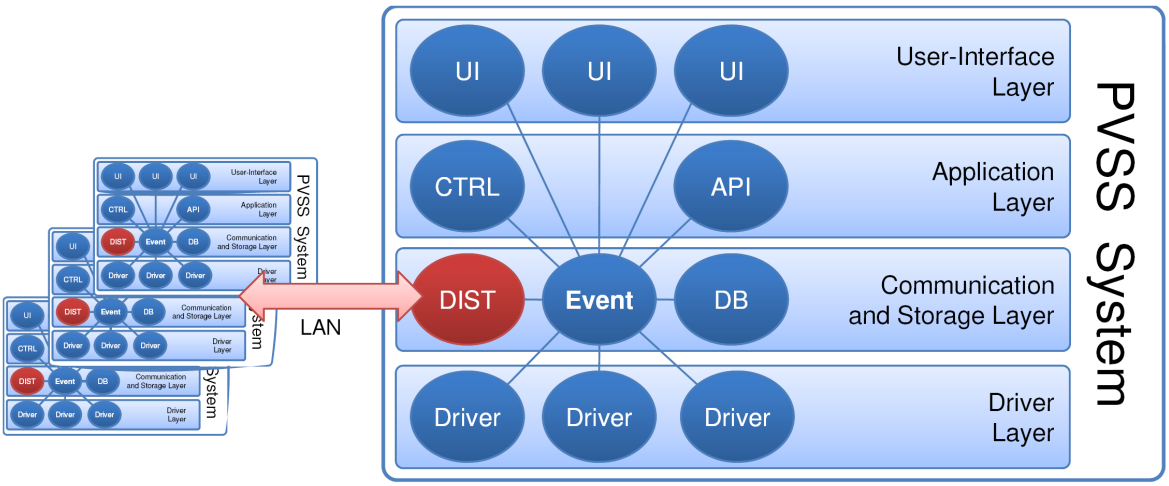

Figure 5.5: *PVSS architecture [27], [65].*

PVSS also offers the possibility to setup systems in a redundant way. For this, two redundant projects are connected via the redundancy manager. It defines which of the projects is the active one and can send messages to the connected hardware and the user interface, and which one is in passive mode but ready to take over in case of a failure. Therefore, both projects need the connection to the hardware, but in case of a PC crash the system downtime can be reduced significantly.

Each self-consistent system – identified by a unique system number – consists of the managers necessary to fulfill the specific tasks of the corresponding system. This architecture is an important factor for the scalability of PVSS. The use of separate processes allows one to split up tasks into multiple managers and make use of multi-core processors, even if a system runs completely on a single PC. In addition, it is possible to connect a number of up to 255 independent systems, which will generally run on separate PCs in the ATLAS Detector Control System. However, there is a limit imposed by the fact that additional managers come with a certain overhead. Also the performance of the complete system is strongly correlated with the total number of datapoints in the connected systems.

The Datapoint Structure

For the storage of the data, PVSS employs the concept of datapoints. This is shown in figure 5.6. The user can define the Data Point Type (DPT) to match the structure of the represented device. Then instances – the datapoints – can be created to handle the individual devices. This device oriented approach fits well to large and complex systems such as the LHC experiments. Further properties of Data Point Elements (DPEs) are stored in the so called “configs” which can be set up individually for each node. Amongst other things this can be the information about which hardware channel is connected (periphery address), whether and how the value is passed between the periphery and the event manager (smoothing, message conversion, command conversion), or the storage of the value (archiving). On each level an alias can be defined, which is inside its own system an equivalent way to the actual addressing via the DPE name.

For alerting functionality, PVSS has implemented a concept of general alert classes and individual alert handling on a node level. The alert class defines e.g. the severity of the alert, whether the operator has to acknowledge it, and whether it should be stored into the database. In the alert handling config the number of alert ranges and their limits can be configured, as well as the corresponding alert class defining what type of alert corresponds to which value range. Summary alerts can be used for alert reduction, in case that a large number of alerts are triggered by the same underlying cause.

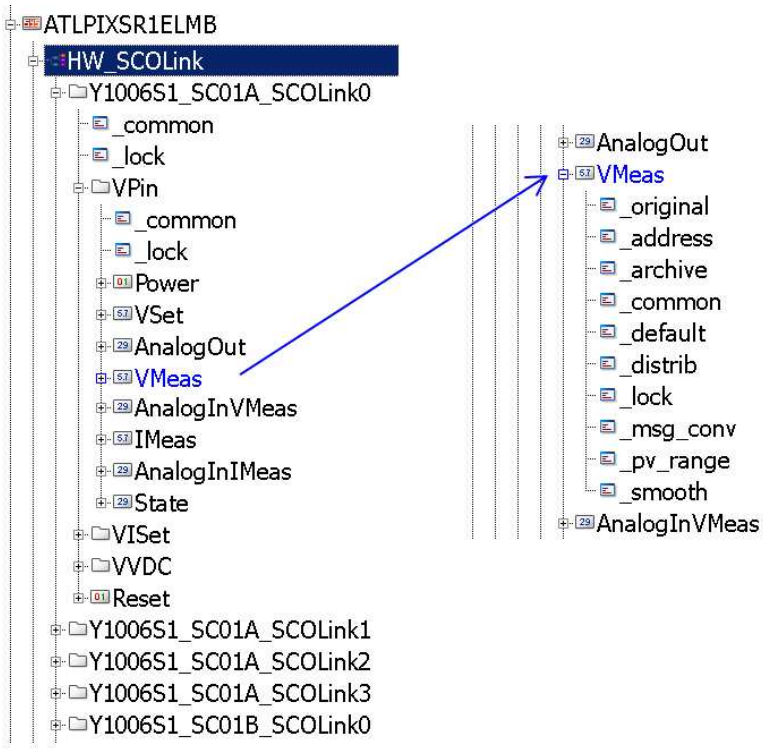

Figure 5.6: *PVSS Data Point Structure. On each system (here ATLPIXSR1ELMB), different data point types (e.g. HW_SCOLink) can be set up. The structure of the data point type defines what DPEs the corresponding data points (Y1006S1_SC01A_SCOLink0) possess. For each node a number of data point configs can be added, the defined configs for VMeas are shown on the right.*

A feature of PVSS is the fact that process variables are not only stored in memory but also on hard disk. While this increases the time to access the values, it guarantees that the state of the system is not lost in case of a system crash, but can be properly recovered.

A summary of the most important definitions and items used in the Pixel DCS is listed here for reference:

- Data Point Type (DPT): The DPT defines the structure of the data. In case the DPT is supposed to manage data from a specific hardware type, it is built in a device oriented way.
- Data Point (DP): DPs are specific instances derived from the DPT, usually representing individual hardware devices.
- Data Point Element (DPE): DPEs represent the internal structure of the DPT, i.e. the “leafs” containing the actual values.
- Data Point Configs: The configs hold additional configuration information for the DPEs (see below).

Every DPE has the configs that hold the “original” value and “common” information like an alternate description or alias, and the possibility to lock that DPE. In additional configs it can e.g. be configured what hardware channel is connected to which DPE (“address”), and in which form it should be handled (“message conversion”, “value range”), when it is updated (“smoothing”), whether it is to be archived, or what values should raise an alert. Another important mechanism is provided by the “dpFunction”, which makes it possible to subscribe to changes of any number of other DPEs. Its actual value is calculated from the values of the other DPEs. This is done by

the event manager, so any function used has also to be known to the event manager, but the advantage is that this mechanism is active as soon and as long as the corresponding project is operational.

User specific function calls

As the Event Manager holds the current image of all process variables all the time, any other manager can be notified on change. The user specific functionality will usually be implemented in CTRL managers. A possible subscription mechanism, where a specified callback function is executed on any change of the variables passed as argument, is the "dpConnect" function. Another way to do a regular evaluation of a callback function is the timed function mechanism. This means that the callback function is executed at regular intervals. The main difference is that this is not event driven. Depending on the use case the most appropriate mechanism – event driven or time driven – can be selected. The parameters of the timed function, e.g. execution interval, can be reconfigured at run time, while the dpConnect is completely depending on the actual refresh rate of the connected values.

5.2.6 The JCOP Framework Components

The JCOP framework provides packages of software components for tasks common to all subsystems. This starts with basic functionality needed by all developers configuring their projects, which is facilitated by the providing of common libraries e.g. for the setup of archiving or alert handling. Tools for the integration and control of commonly used hardware types, e.g. the ELMB, but also commercial power supplies, exist. Generally, tasks which are not experiment specific, but are faced by all users in a similar way, are taken care of by the framework, so ensuring a high degree of conformity and therefore ease of maintenance.

5.2.6.1 The Framework FSM

To handle the control of complex hierarchical systems like the LHC experiments, an FSMs (Finite State Machine) is a well suited tool. The JCOP framework provides an external tool [66], [67], that can be integrated via the API interface into the control system, even though PVSS does not support an FSM per se. At the core of the FSM component is SMI++ [68], an enhanced version of the State Manager Interface (SMI) developed for the DEPLHI experiment at LEP, but based on C++. It is used by ATLAS and CMS for the high level supervision and control of the DCS related aspects. LHCb and ALICE use SMI++ for DCS as well as for DAQ control.

For SMI++, a real world problem can be described as a series of objects which can be "associated" or "logical". The associated objects describe the interface to the real hardware objects. They can be represented by discrete states and actions. The interface between SMI++ and the software components used to control the actual hardware is situated at this level. In the case of the LHC experiments this is PVSS.

The behaviour of the logical objects on the contrary is entirely determined by SMI. Logical objects are organized in "domains" which correspond to one SMI process. In large control systems it is foreseen to have many domains, which are organized hierarchically and distributed over several computers. For the description of objects and actions, the State Manager Language (SML) is used. For each object a dedicated sml-file is produced, which is read in by the SMI process and defines and drives its behaviour.

In the JCOP terminology, the associated objects are the Device Units (DU). They are described by a simple list of states and actions in the sml file. There are two types of logical objects, Control Units (CU) and Logical Units (LU), the difference being that a Control Unit corresponds to an SMI process, while Logical Units are part of the domain's SMI process. This implies that Logical Units can not have other logical objects in the hierarchy below them, and has consequences for partitioning. For logical objects, besides the pure list of states and commands, the state logic needs to be specified.
Commands can be propagated to objects on lower levels in the hierarchy, even with some basic conditional logic.

In addition, actions can be triggered on certain conditions. This more complex format of the sml file defining state logic and possible commands is containing the so called “when-list”.

Concerning the general behaviour of the controls hierarchy, states are only propagated upwards, commands only downwards. The object which is in the hierarchy on the level above another one is referred to as the “parent”, the objects on the level below as its “children”.

The communication between the processes is handled by DIM (Distributed Information Management, described later in this section. The DIM Name Server, “dns”, keeps track which process runs at which location. For the sml code it is only necessary to know the domain names. The user of the “FSM toolkit” can concentrate on the functional behaviour describing the states of his system. For the internal handling of the communication between SMI processes, libraries are provided. This aspect is therefore transparent to the programmer of the control system.

Partitioning

To allow parts of the control system to be run in parallel, there are different modes of partitioning. Device Units and Logical Units can be enabled or disabled. A disabled object is not taken into account for state calculation and does not accept commands from the SMI hierarchy.

For Control Units there are four modes, which can only be set by the parent unit:

- included: The nominal condition of an object. Commands are accepted, and the state is propagated.
- excluded: An excluded object does not contribute to the state of the parent object, and does not accept any commands.
- locked_out: An excluded object can be locked out, so that it will never be automatically re-included into the tree without an “un-lockout” action before.
- ignored: An ignored object accepts commands, but its state is ignored for the higher level states.
- manual: In this case the state of an object is propagated upwards, but commands are only accepted on the actual level. Commands from objects upwards in the hierarchy are not executed.

Besides the partitioning, the concept of ownership (see figure 5.7) plays a role in the control of the FSM hierarchy. Ownership is not tied to a user, but to a certain user interface, or more precisely the manager number. Only the owner of the FSM can partition the FSM tree, and in addition can determine whether the FSM is operated in exclusive or shared mode. Exclusive mode means that commands can only be sent from the owner user interface, while in shared mode any user who has the right privileges can send commands from any user interface.

Distributed Information Management (DIM)

DIM [69], [70] was developed in the frame of the DELPHI project. It is a communication system using TCP/IP and can run as well on Windows as on Linux (and others). It operates on the client server principle, but the same process can act as client and as server (e.g. the DIM Manager in PVSS). The server provides a service – a set of data – to the client. The client subscribes to the service only at startup. It is then updated either on change or on regular time intervals from a callback routine without further requests from the client. All services are registered by the servers with the DIM name server. At startup the client gets the name of the server which publishes the requested service from the dns and subsequently subscribes directly with the server to the service. If any process dies, the subscribing clients are informed and will reconnect as soon as the server is available again.

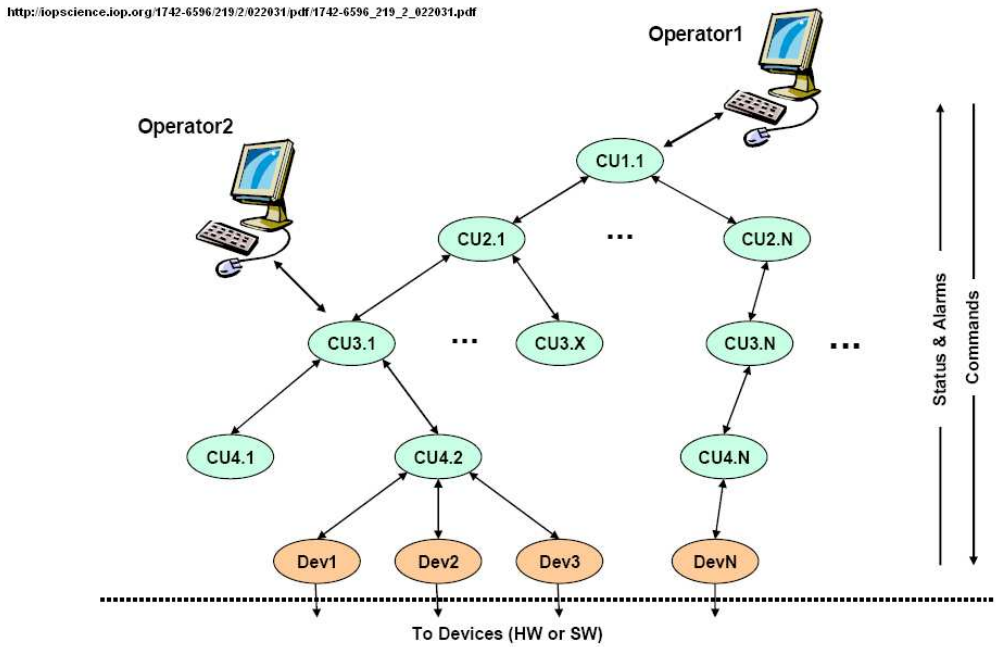

Figure 5.7: *Controls hierarchy [66].*

ATLAS FSM

ATLAS DCS has built its own additions on top of the framework FSM in the fwFsmAtlas package. Besides providing a set of common panels and widgets, the main addition is the concept of two parallel hierarchies to describe the condition of the experiment. These parallel hierarchies are referred to as “state” and “status”, and make it possible to consider to a certain extent separately operational issues from safety issues. The package provides template object types, with a predefined set of (high level) states and commands. The aspects relevant for the implementation of the Pixel Detector Control System will be described in more detail in chapter 6.

5.2.6.2 Access Control

The framework also provides a package for the management of the access control, to prevent unauthorized control actions. Access control is organized in three levels – domains, groups, and users. However, the actual access rights are handled only through privileges. Any privileges are generally assigned to the domain, e.g. the privilege “Control” in the domain “PIX”, but the actual access control is managed at the level of the groups. For each group there is a set of privileges which are automatically granted to all members of this group. A group can also include other groups, i.e. grant its members all privileges of both groups. Each user is granted privileges only based on his group membership.

The main relevant distinction in the DCS roles is the distinction between control and expert privileges. The roles are held in the Access Control Server and are centrally synchronized to the LDAP (Lightweight Directory Access Protocol) managed roles relevant in the system administration environment. Inside the individual DCS subsystems, user information is synchronized regularly with the AC server, but a local cached copy of this information provides also the possibility of limited standalone operation in case of network problems.

There are also two non-DCS roles relevant for DCS operation. The first is the PIX:shifter role, which is automatically assigned and enabled whenever a person is on shift. The PIX:remote role allows someone to gain access to the control system even when not physically present in the control room. However, for getting granted remote

access, a procedure is defined, which generally involves notification of the control room that remote access is requested.

Inside the DCS, the access control acts only on the UI level. In each panel, the developers have to define what action is allowed to which user group. During the initialization, the related items are enabled based on who is logged in as user and the privileges assigned to him. While this provides any desired level of granularity in access control, for the FSM implementation by the framework only two levels of access – expert actions and operator actions – are allowed per node. This can be different for each node, even for nodes of the same type, but on a single node it is never possible to have more than two levels of access control inside the framework FSM functionality.

5.3 Pixel DCS Hardware

The hardware of the Pixel Detector Control System [71] can be divided into three main groups:

- power supply system
- monitoring system
- interlock system

The power supply system (5.3.1) provides the required voltages for the modules and opto-boards.

Module and opto-board temperatures are monitored, as well as temperature and humidity of the detector environment(5.3.2).

The hardware interlock system (5.3.3) protects the detector. The temperatures of the modules and opto-boards provide one type of input to the interlock system, to prevent overheating. The interlock system protects not only detector and equipment from high temperatures, but also shuts down the lasers of the optical link in case that personnel could be exposed to laser light, and accepts inputs from the ATLAS detector safety system to react on global danger situations.

Further devices are the opto-heaters (section 5.3.1.1), which are installed to optimize the operation temperature of the opto-boards.

Besides the crates that are directly responsible for detector devices, the Pixel Detector Control System comprises auxiliary power supplies for DCS and DAQ hardware, and monitors the environmental conditions of the off-detector part of the optical link.

5.3.1 The Power Supply System

Figure 5.8 shows the power supply system for one readout unit, consisting of the six or seven modules and the corresponding opto-board. This includes the required voltages for the depletion of the silicon sensor, the supply of the module front-end electronics, and the control and supply voltage for the opto-board, as well as active regulation near the detector of high current channels.

High Voltage

The depletion voltage (HV) for the pixel modules is provided by iseg[5)] high voltage power supply crates. Each remotely controllable crate can house up to eight 16-channel modules. Communication to the iseg is established via CAN-Bus and a vendor provided OPC server.

[5)] iseg Spezialelektronik GmbH, Radeberg / OT Rossendorf, Germany

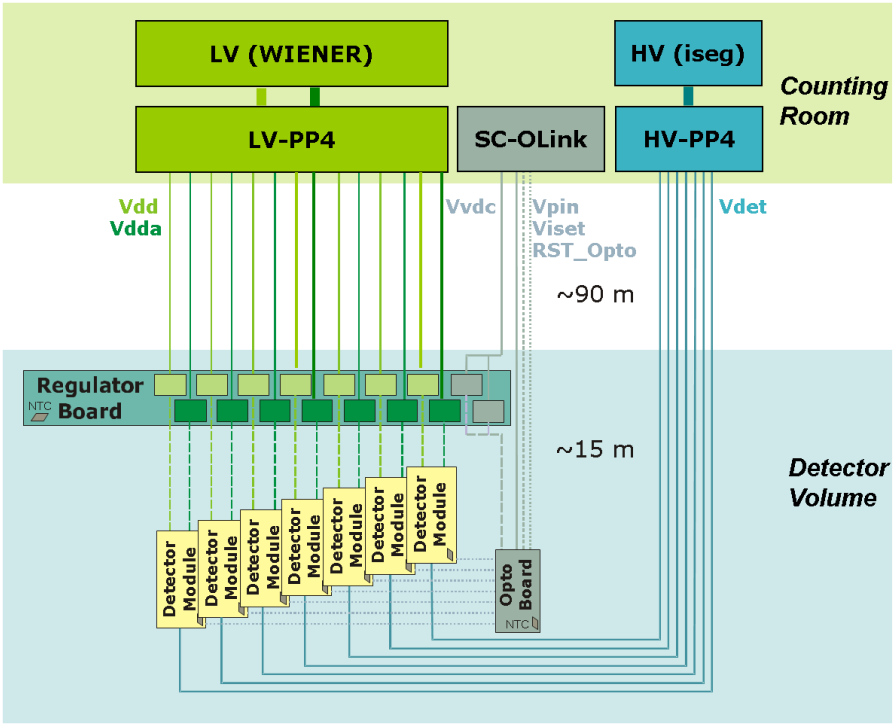

Figure 5.8: *The Pixel Power Supply System, shown for one readout unit.*

In the initial supply scheme, the modules of one readout unit are supplied by one iseg channel, that is split up by fan-out units ("HV-PP4"), installed near the power supplies, into seven separate lines. Once the pixel modules are irradiated, the restriction to a maximum current of 4 mA limits the number of detector modules which can be supplied by a single iseg channel to two.

The High Voltage PP4 crates which are responsible for the fan-out support both these schemes, and can provide current measurement for selected channels.
The current measurement capabilities of the HV-PP4 are integrated into the control system by an ELMB based solution, with communication also via CAN bus, and the CANopen OPC Server.

Low Voltage

Low voltage power supply crates by the company WIENER[6)] provide the power for the front-end electronics of the pixel modules. One channel each is used for the analog (VDDA) and digital (VDD) circuits. LV-PP4 crates are responsible for the fan-out and current measurement of the supply line. The services for each module are routed individually from PP4 to the modules, with active regulation at PP2.

Each WIENER crate has twelve galvanically isolated channels of up to 12 V at max. 15 A. The WIENER crates are the only devices in the Pixel DCS which do not communicate over CAN, but through TCP/IP.

The LV-PP4 crates are built from three identical blocks, with the additional option for each to measure either the currents on the return line or the supply voltage of the crate. De facto all crates used in the production system are identical. Two blocks measure the supply lines to the modules only, while one block is equipped additionally to measure the supply voltage of the crate. The LV-PP4 crates are ELMB based, and communicate via CAN-bus and CANopen OPC.

[6)] W-IE-NE-R Plein & Baus GmbH, Burscheid, Germany

Supply and Control for the Optical Link

The SC-OLinks (Supply and Control for the Optical Link), which are providing the power for the opto-boards, are custom built crates based on ELMBs. For each readout unit one complex channel is required, providing three supply voltages and one control signal to the opto-board. VPin is the depletion voltage for the PiN diode which receives the optical signals. VVDC is the supply voltage for the DORIC and VVDC, which respectively are responsible for decoding the incoming optical "down-link" clock and command signals, and driving the VECSLs for the outgoing "up-link" data signals. VISet determines the laser power for the outgoing signals. RST_opto is a control signal, which induces the DORIC to recover the correct clock frequency. As VVDC is a high current channel, it is regulated at PP2, while for VPin, VISet and RST_Opto there is no further active element between the SC-OLink and the opto-board.

One SC-OLink crate consists of four blocks, and each block can provide power for 4 opto-boards. All channels are hardware current limited.

The Regulator Stations

To compensate for the voltage drops caused by high current values over about 100 m long cables, and for protection of the sensitive front-end electronics against transients, the corresponding high current, low voltage channels are regulated at the service point PP2 inside the detector volume, at a distance of about 15 m from the Pixel Detector. The affected channels are the low voltage supplies Vdd and Vdda for the modules, as well as VVDC for the opto-board. Each regulator station[7)] [79], [78] houses one controller board and twelve regulator boards, and can serve up to twelve readout units. Each regulator board features 16 regulators, seven each for the individual control of the module low voltages for one readout unit, and two for the control of VVDC. The redundancy for the VVDC channel is intended to prevent the loss of a complete readout unit in case of a single failure. The voltage set-point is selected by 100-step digital trimmers. The interface between DCS and PP2 consists of a controller board with an ELMB for communication, and an FPGA (Field Programmable Gate Array) responsible for the actual control.

5.3.1.1 The Opto-heaters

To enable operation of the opto-boards at stable conditions near room temperature, resistive heaters have been installed on the opto-boards [37]. One heater strip serves a row of six opto-boards, as shown in figure 5.9, with a total of 48 heaters. For additional thermal stability, the opto-board region is covered by a "thermal blanket".

With a back-pressure set-point for the opto-loops of 4.15 bar, and a heater temperature set-point at an average of 15 °C in 2010, this corresponds to opto-board operation at 17 °C on average, depending also on board type and position on the service quarter panel.

The opto-heaters were initially controlled by the same switching card system used for the thermal enclosure heaters. During the shutdown in 2008, the system was changed to continuous software regulation, using commercial WIENER power supplies to provide the power for the opto-heaters.

Four twelve channel WIENER power supply crates are employed, capable of providing a maximum power of 40 W per channel, with up to 60 V. One channel is used per heater, limited for safety reasons to 0.5 A. The regulation is done by software, adjusting the output voltage of each channel based on average opto-board or opto-heater temperature (see also section 5.4.3).

[7)] custom built by the INFN Milano group

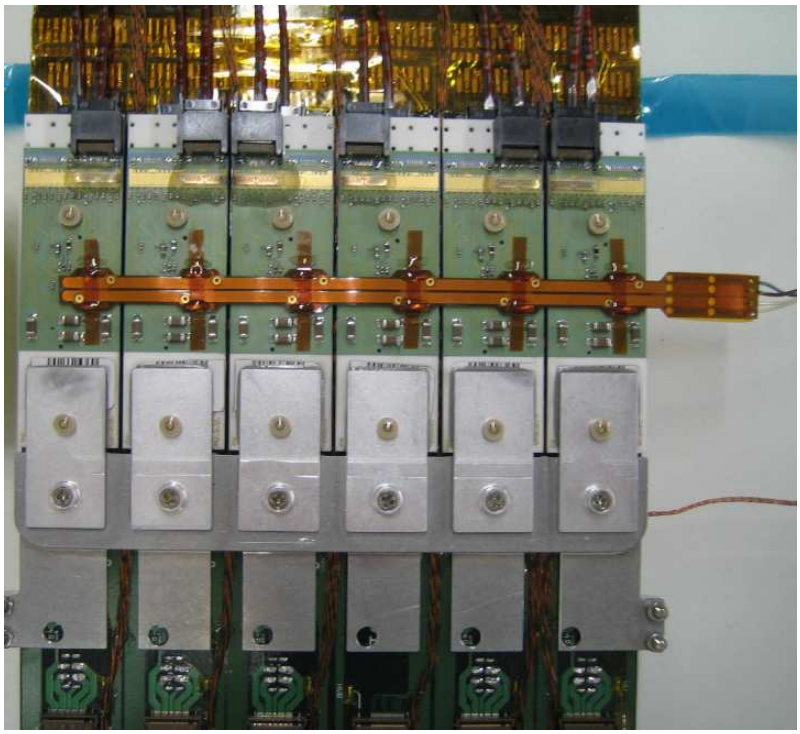

Figure 5.9: *Opto-heaters [80].*

5.3.2 The Monitoring System

Due to the susceptibility to heat of the irradiated modules, the temperature is monitored closely for each detector device. Also, the environmental temperatures and humidity are closely monitored inside the Pixel volume (see 5.4.4.1), since condensation of water may cause damage to the system.

For the temperature monitoring, two types of crates are employed. Environmental temperatures which are not associated to a specific detector device are read out by BBM (Building Block Monitoring) crates. In addition to temperature monitoring, a dedicated block of the BBM crates provides also monitoring of the humidity sensors located inside the pixel volume and in the ID end-plate region. Whenever a temperature signal is also an input to the interlock system, the BBIM (Building Block Interlock and Monitoring) crates are employed for monitoring. In addition to the readout part, the BBIM is equipped to generate the relevant information for the interlock system.

Additionally to the detector, temperatures of DCS and DAQ crates are monitored, namely the regulator stations and the temperatures of the DAQ Back of Crate cards. PP2 temperatures are measured by BBIM crates, while for the monitoring of the BOC temperatures the dedicated BOCMON units are installed inside the DAQ crates. Temperature information is also provided by the commercial power supplies, but is not considered in this context.

5.3.2.1 Environmental Sensors

Temperature Sensors

For the temperature measurement in the Pixel DCS, NTC (Negative Temperature Coefficient) sensors are used, i.e. their conductivity increases with higher temperature. Their advantages include a high relative change is resistance in response to temperature changes, and low self-heating due to a high resistance [8)]. The dependence of temperature and resistance is however not linear, and is described by the Steinhart-Hart equation [72]. For the

[8)] typically 10 kΩ at room temperature for the NTCs used in the Pixel DCS.

Pixel Detector, the resistance is measured over a voltage divider with a series resistor of 10 kΩ and a reference voltage of 2.5 V.

Humidity Sensors

For the humidity measurement of the Pixel Detector environment, three types of humidity sensors are in use. The radiation hard Xeritron [74] is based on the measurement principle of resistivity change, and consists of organic fibre material. This sensor was chosen for its radiation hardness and robustness. However, it is more applicable for the monitoring of humidity trends instead of absolute values, due to the long time needed to reach its saturation value. Therefore, additional Hygrotron (Hygrometrix HMX-2200) [75] sensors are employed. For these, the measurement is based on a piezo-resistive strain gauge. The Hygrotrons have a much faster response time, but less extensive calibration was conducted. To compensate for this, non radiation hard capacitive Honeywell [76] sensors were deployed for "in situ" calibration purposes.

5.3.2.2 the BBIM crates

The BBIM crates (Building Block Interlock and Monitoring) have two tasks, to provide continuous temperature information for monitoring, and to generate binary signals as input for the interlock system. The NTCs of the modules, opto-boards, and the seven sensors placed in each regulator station are read out by the BBIMs. Each crate consists of 4 Blocks with one ELMB and four Interlock "Boxes" ("Ibox"). The temperature monitoring is based on the ADC of an ELMB, while the Ibox is responsible for the interlock signals. Each Ibox generates the signals for 2 times seven temperatures, i.e. it is responsible either for two readout units, two times six opto-boards, or two PP2 crates.

The interlock signals are created by comparator circuits. The interlock threshold, or "shooting point", is determined by an exchangeable resistor, and is set at 40 °C for modules and opto-boards, and at 60 °C for the regulator stations. There are two signals generated per temperature input: "TEMP_HIGH" and "TEMP_LOW"[9)]. This allows the user to distinguish e.g. between interlock generation due to a high temperature, or due to a broken cable. If the BBIM crate itself is un-powered, an interlock will be generated. If the ELMB is un-powered (e.g for the reset of a CAN-Bus), it means that the shooting point will be shifted to lower temperatures. The BBIM power for all BBIM crates is supplied by a dedicated power supply crate ("BBIM-Power") located in the USA15 counting room. For each interlock box the possibility to manually generate a test-signal is implemented. This makes it possible to verify the correct operation of the interlock system, but due to the number of connected sensors only to a best granularity of two readout units.

5.3.2.3 the BBM crates

The BBM crates are monitoring the environmental conditions, which are not directly related to a specific hardware device, and do not provide an input to the interlock system. Each BBM contains three blocks: two BBM-NTC blocks for the monitoring of the environmental temperatures, and one BBM-HS block for the monitoring of the humidity sensors and a number of corresponding temperature sensors for the determination of the dew point.

BBM and BBIM crates are located inside the experimental cavern, outside the Muon chambers.

5.3.2.4 Environmental Monitoring of the DAQ Racks

Six NTCs are placed on each BOC-card. To read out four of them, the BOC-MON units are placed inside the DAQ racks next to the BOC cards. The BOC-MON NTCs are read out by the ADC of an ELMB. As each DAQ

[9)] The threshold for TEMP_LOW will not to reached under normal conditions.

crate can house up to 16 BOC cards, one ELMB is sufficient to monitor one DAQ crate. Unlike BBIM and BBM crates, this leaves no free channels to monitor the reference voltage. The last two NTCs are only read out by the DAQ, and are introduced into the DCS system via DDC (see section 5.4.5. Before the technical stop end of 2010, additional crates ("BOC-HUM") for the humidity monitoring inside the DAQ racks were installed.

5.3.3 The Interlock System

The interlock system for the Pixel Detector protects against three categories of dangers [81]: The Pixel Detector is protected against over-temperature, humans are protected from hazards by laser light, and ATLAS-wide issues that are handled by the detector safety system (DSS). This last category includes different sources of danger like e.g. smoke, power failures, or cooling failures (see 2.2.7). The interlock system reacts to these issued by switching off affected power supplies or lasers.

The "interlock matrix", see figure 5.10, shows how the interlock inputs are connected to the corresponding actions. The hardware to realize this scheme is described in the following.

		Action on Readout Units							
	Reason	LV	HV	SC	Laser	OH	AUX	BO	Granularity
NTCs	TMod (module temperature)	1	1	-	-	-	-	Bakeout Carts [a]	Readout unit
	TOpto (optoboard temperature)	-	-	1	-	All	-		Readout unit
	TPP2 (PP2 temperature)	12	-	12	-	-	-		PP2 crate
	TOH (optoheater temperature)	-	-	-	-	6	-		Optoheater
Contact switches	I-BOC (Rack Door opened)	-	-	116 [b]	116 [b]	-	-		DAQ rack (one or two crates per rack)
	I-PP1 (ID endplate opened)	-	-	138 [c]	All	-	-		PP1 Box [d] (one per counting room)
DSS	IDSS-0 (e.g. cooling failure)	All	All	All	-	All	All		
	IDSS-1 (no stable beams)	-	All	-	-	-	-		
	IDSS-2 (cooling loop [f] failure)	-	-	-	-	-	-	Bakeout Carts	

Figure 5.10: *The interlock matrix. On the left, the inputs to the interlock system are shown, grouped into three categories. Those are high temperature of detector or DCS hardware, danger for the personnel by exposure to laser light, and in the last group detector wide issues handled by the ATLAS detector safety system, DSS. The affected channel types and number of readout units on which the interlock system acts is given. Where there is not a fix granularity, the maximum number of readout units is given. The interlock system can act on low voltage (LV), high voltage (HV) or SC-OLink channels (SC), as well as off-detector lasers (Laser), the opto heater system (OH), and the power supplies for the regulator stations (AUX). The Pixel interlock system also provides an interlock for the beam pipe bakeout system (BO).*
(a) For B-Layer modules. (b) varying strongly by crates. The maximum number of affected readout units is given for the crates L1 and L2. (c) The number of readout units supplied from US and USA side is not identical. (d) The information about opened contact switches is available on a per quarter base, but in any case all readout units belonging to one PP1 Box are switched off. (f) triggered if any of the cooling loops considered as essential failed.

There are three basic types of input to the interlock system:

- NTCs provide temperature information, which the BBIMs convert into binary signals.
- Contact switches in the DAQ racks and at the ID end-plate warn against the danger of exposure to laser light.
- external signals provided by the ATLAS detector safety system or the LHC provide information about various dangers, including among others cooling failures, fire detection, or instable beam conditions.

The interlock system acts primarily on power supplies, and the different types are linked in figure 5.10 to the interlock inputs.

The main part of the interlock system, with respect to channel number, consists of the chain BBIM, Logic Unit, IDB (Interlock Distribution Box). In this chain, the BBIM is responsible for the digitization of the temperature input signals from modules, opto-boards, and regulator stations. The Logic Units have the knowledge of the Interlock Matrix, and the logic to combine the inputs into the decision which hardware needs to be switched off. The interlock distribution box is responsible for the mapping. It contains the information which actual power supply channel belongs to the readout unit affected by a given interlock.

The laser related inputs are handled by the BOC-I-Box and the PP1-Box. The BOC-I-Box provides information about the status of the off-detector door switches, the PP1-Box about the on-detector contact switches. In case any side indicates a potential danger, both on-detector and off-detector lasers must be switched off. The on-detector lasers are powered by the SC-OLink. Therefore, this information is going from the PP1-Boxes and BOC-I-Box via the PP1-Box to the Logic Units, where it is another input to the interlock matrix. The output from the Logic units is distributed by the IDBs to the SC-OLink channels. For the off-detector lasers, the BOC-I-Box handles the information coming from the PP1-Box, and also the information from the door switches of the DAQ racks. This is then distributed to the corresponding BOC cards.

The information from ATLAS DSS is received by the PP1-Boxes. The relevant part of the DSS action matrix, linking dangers to DSS actions, is described in [82], and the currently valid matrix is available online [83]. The pixel interlock system handles three types of DSS actions. The corresponding outputs created by the PP1-Box are called DSS0, DSS1, and DSS2.

The DSS0 signal is created by the PP1-Box if the DSS action related to HV and LV power is triggered. DSS1 only acts on HV power. DSS1 and the detector related part of DSS0 are handled by PP1-Box, Logic Units and IDBs. The third DSS action type, DSS2, does not act on Pixel power supplies but on the beam pipe bake-out carts.

With only two exceptions, all DSS actions which are passing through the pixel interlock system are DSS0 actions. The predominant cause have been failures of the evaporative cooling system. Besides the detector related interlocks, DSS0 also switches off all opto-heaters through the OHI-Box (Opto-Heater Interlock Box), and the auxiliary power supplies for the regulator stations.
DSS1 is acting only on high voltage, and therefore dedicated to beam related issues. There are two possible inputs to the PP1-Box that cause the creation of the DSS1 signal, and the only one actually in use is not generated by ATLAS DSS, but by a signal provided by the LHC. This LHC signal is handled by the PLI-Box (Pixel LHC Interface Box), and joined with the DSS-action before being fed into the PP1-Box.
DSS2 was introduced to protect the Pixel Detector from the potentially fatally high temperatures during the beam pipe bake-out. It was mandatory to immediately stop the bake-out in case of a malfunction of the cooling system. The failure of a critical loop would therefore generate the corresponding DSS action received by the PP1-Box. The DSS2 signal is transferred by a dedicated Bake-out Box (“BOB”) to the beam pipe bake-out system.

The interlock system is arranged in independent interlock quarters. ELMBs are used for the monitoring of the interlock system, and all monitored signals are latched, to allow the user to reconstruct the behaviour of the system even in case of very short interlock signals. The performance of the interlock system is however based purely on hardware components, and not dependent in any way on the software monitoring. All interlock signals follow

negative logic, meaning that in case of a power loss or a broken cable the interlock is active, and the corresponding channels can not be powered.

The following types of crates are employed in the interlock system:

The Logic Unit

The Logic Units are responsible for combining the interlock input signals related to up to twelve readout units. This includes for each readout unit six or seven module temperatures and one opto-board temperature. The readout units that are powered through the same regulator station are also handled by the same logic unit. This allows the interlock system to switch off all power supplies belonging to these readout units, based on the interlock signals of the seven NTCs installed in each regulator station. The DSS and Laser information from the PP1-Box is routed through all Logic Units. The interlock matrix is incorporated in an FPGA program, calculating the required output signals. Each Logic Unit provides 12 inputs for each of the four corresponding IDBs, in total there are seven Logic Units per interlock quarter.

The Interlock Distribution Boxes

The Interlock Distribution Boxes route the interlock signals generated by the Logic Units to the corresponding power supply channels. Therefore, three types of IDB exist, as each type is responsible for either high voltage (IDB-HV), low voltage (IDB-LV), or opto-board supply (IDB-SC). The mapping information is contained in an FPGA program, so that a change in the power supply connectivity does not require a rewiring of the interlock system. The interlocks of each interlock quarter are handled by one IDB of each type, except for the low voltage, where two IDBs are needed, following the modularity of power supply channels per readout unit.

The PP1 Boxes

One PP1-Box in each counting room is responsible for receiving the signals from the detector safety system, the signals related to on- and off-detector lasers, and the input from LHC regarding beam conditions. This information is routed to all the Logic Units in the same counting room. On-detector laser information from both PP1-Boxes is routed to the BOC-I-Box in the USA counting room.

The BOC-I-Box

The BOC-I-Box is installed in one of the DAQ racks and is responsible for the creation of interlock signals in case one of the doors of the racks housing the BOC-crates is opened. These interlock signals are distributed to the BOCs to switch off the off-detector lasers, and to both PP1-Boxes to effect the on-detector lasers to be switched off. In turn, the BOC-I-Box receives the signals from both PP1-Boxes for the distribution to the BOCs.

The OHI-Box

One Opto-heater Interlock Box is installed in each counting room. It receives from each Logic Unit the OR of all binary opto-board temperature signals, in total 14 temperature input signals. In addition, the OHI-Box in the US counting room receives the OR from the USA side as “carry” signal. The US side OHI-Box produces in this way one overall output from the factual ORing of 272 opto-board temperature signals.

The NICO-Box

The NICO (New Interlock and Control of the Opto-heaters) Box is a hybrid box including a BBIM part to generate interlock signals from all 48 analog opto-heater temperatures, and a decision (NICO-IN) and mapping part (NICO-OUT). The evaluation of which channels need to be switched off is based on 48 opto-heater temperatures, the OR of all opto-board temperatures provided by the OHI-Boxes, and DSS0. The interlock threshold for the heater temperature lies at 35 °C and results in the corresponding heater being switched off. This affects up to six readout units. As the opto-board temperature information is retrieved from the logic unit, where the temperature is already transformed into a binary signal, there is only a single interlock threshold for the opto-board temperature, i.e. if any opto-board temperature is higher than 40 °C, all heaters are interlocked. Also in case of an active DSS0 signal, all heaters are interlocked.

The Bake-out Boxes

The Bake-out Boxes are responsible to stop the bake-out of the beam pipe in case the Pixel Detector can not be sufficiently cooled. Two Bake-out Boxes are installed, one in each counting room. The interlock signals related to the B-Layer modules are split in the counting room and routed to the Logic Units and the Bake-out Boxes. The ORed signal generated from all individual module signals readout from the US side is routed to the BOB on the USA side, where a single output is generated from the carry signal, the USA module temperatures , and the DSS2 signal. The DSS2 action is triggered in case one of the loops considered as critical fails. The overall output is sent via the Little BOB to the bake-out system. Next to the adaption of the signal for the bake-out system, Little BOB provides a latching of the interlock. This is the only case where the latching does not affect only the monitoring, but the signal itself is latched. This means that even if the temperature has sufficiently fallen and all loops are running again, the signal still has to be cleared manually for the bake-out to resume.

The Pixel LHC Interface

The PLI-Box was installed to handle the interaction with the accelerator to ensure the safe operation of the detector with respect to unstable beam conditions. As the main consideration was the fear of a high localized charge concentration shorting high voltage to the front-end electronics, it must be ensured that high voltage can only be switched on when the situation allows it. The PLI-Box

- receives one input signal, the STABLE BEAMS signal from the LHC, distributed to the ATLAS sub-detectors by the ATLAS-LHC-Interface.
- provides one signal to LHC via the Injection Permit System, specifying whether the Pixel Detector is in a safe state for beam injection, the INJECTION PERMIT.

The signals exchanged between PLI-Box and ATLAS-LHC-Interface are differential RS485 signals. As the STABLE BEAM signal is not present when there is no beam, which is a situation that is not dangerous to the detector, the possibility to mask this interlock is necessary. Otherwise it would not be easily possible to do any calibration that requires the presence of high voltage. The PLI Box is the only part of the interlock system where a remote overriding of a hardware interlock is possible.

The outputs of the PLI Box are controlled by a digital port of an ELMB which makes it also possible to monitor all relevant signals on analog channels. The monitoring signals are however not latched, but the result is latched in the PP1 Box. The default outputs, which are assumed in case of a reset of the ELMB via the CAN bus, are injection permit granted and interlock unmasked. If the PLI-Box is un-powered, the interlock output is active, and high voltage can not be switched on.

As the hardware interlock system does not foresee high voltage information as an input, all outputs of the PLI-Box have to be generated by software. This has been part of the work for this thesis and is described in more detail in chapter 7.1.4.

5.3.4 Deployment of the Pixel DCS Hardware

All the Pixel DCS Hardware is located inside the USA15 and US15 counting rooms at level2, at PP3, or at PP2.

5.3.4.1 Regulator Stations at PP2

The regulator stations are installed in different locations at PP2, called platforms. 28 PP2 crates are located on six platforms on each side of the detector. Each platform houses two or three PP2 crates, and supplies one or two octants of the Pixel Detector. Figure 5.11 shows the layout and locations.

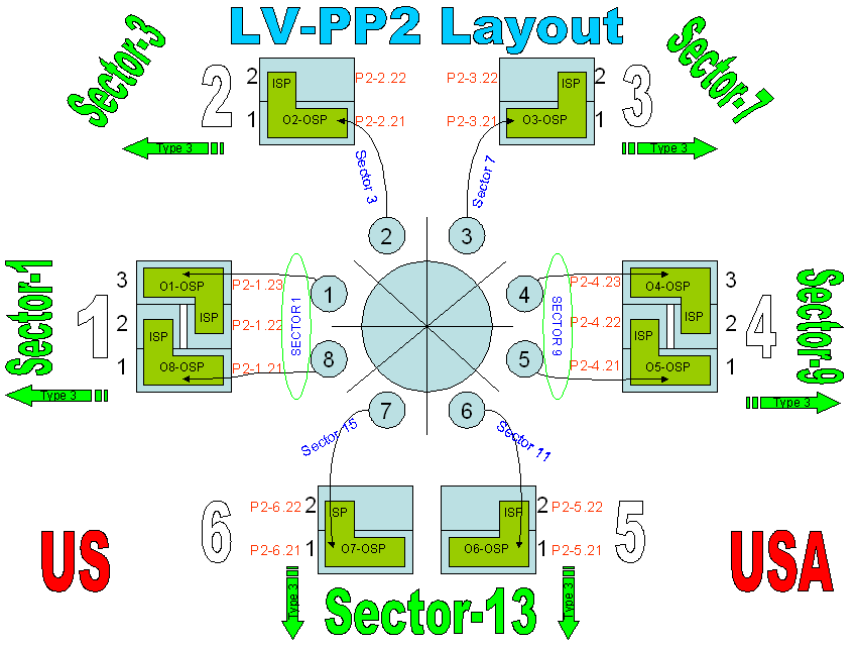

Figure 5.11: *PP2 Platforms [80].*

5.3.4.2 Environmental Interlock and Monitoring Crates at PP3

The BBIM crates are located inside the experimental cavern outside the Muon chambers at PP3. The three BBIM crates each for modules and opto-boards of one quarter are located on Levels 0 and 8, respectively 1 and 7. The PP2 temperatures are partly handled by those crates, and partly by two additional dedicated crates, used for PP2 temperatures only, on Level4. In each quarter, one BBM crate is installed, with two blocks for environmental temperature monitoring and one block for humidity readout. Three additional humidity monitoring blocks are installed for the monitoring of the humidity in the heater trays (see 5.4.4.1). These are located in the racks oriented towards the USA side in level 0 and level 8.

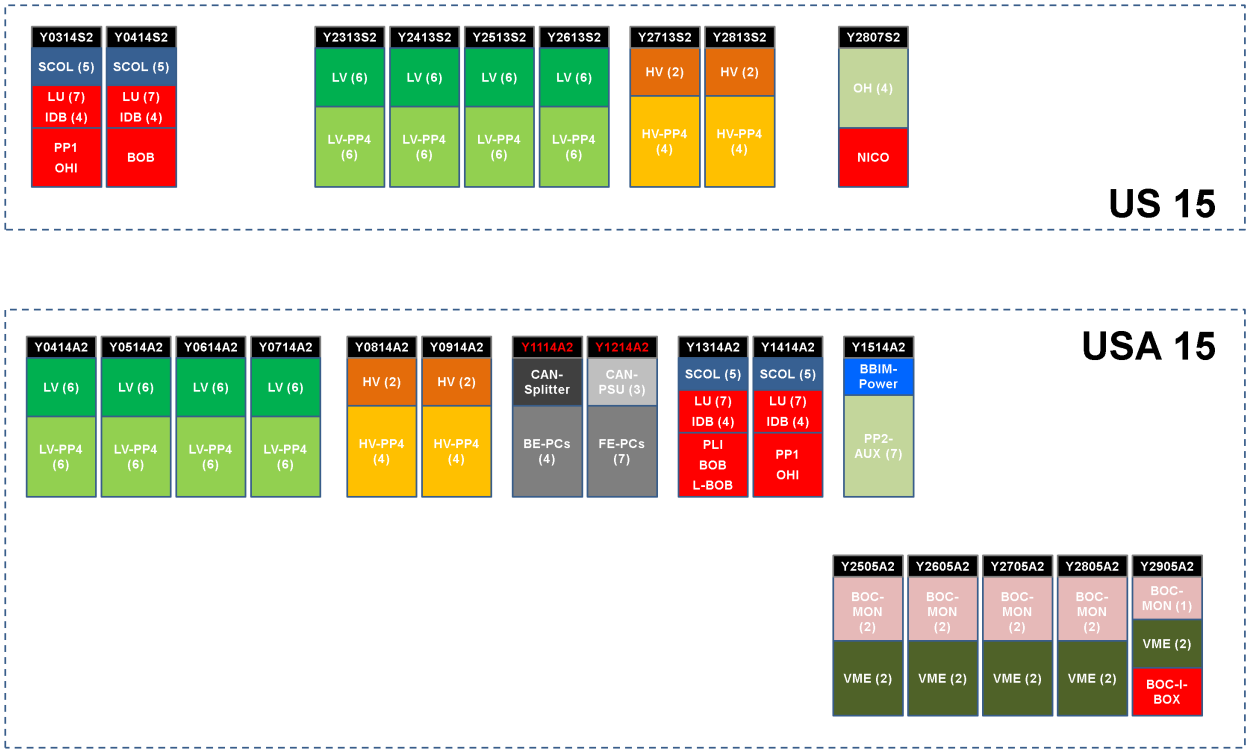

Figure 5.12: *Schematic overview of the Pixel DCS hardware installed in the counting rooms. Used rack space and rack location not to scale.*

5.3.4.3 The Counting Rooms, USA15 and US15.

As the distribution of the pixel DCS hardware inside the counting rooms is relevant for the automatic actions which are one of the focal points of chapter 7, this will be described shortly in this section. The overview of this installation is also important to understand how partial failures will affect the system.

The main power supplies are installed symmetrically inside the two counting rooms: In each counting room are two racks for High Voltage Power Supplies and HV-PP4, four racks for Wiener Low Voltage Power Supplies and LV-PP4 crates, and two racks for the Supply and Control of the Optical Link.

The corresponding interlock crates are installed accordingly: two racks with each seven Logic Units and four Interlock Distribution Boxes in both counting rooms. Additionally, one PP1 Box, one Bake-out Box, and one OHI Box are required.

Some special hardware types are installed only in one of the counting rooms.

US counting room

The opto-heater system started out as a common ID system, therefore the corresponding hardware is installed in the US counting room. These are the WIENER crates for the opto-heater power, and the NICO Box for temperature monitoring and interlock handling of the heaters.

USA counting room

The DCS PCs (section 5.3.5) are located in the USA counting room, to permit continuous access. To ensure a maximum safety against unavailability, they are powered by UPS. For the same reason, the CAN-PSUs are located inside an UPS rack.

Only on USA side are also:

- the BBIM-Power crate, providing the power for all 15 BBIM crates.
- the seven WIENER crates for the auxiliary power of the regulator stations.
- The BOC-I-Box, responsible for the generation and distribution of the off-detector laser interlocks for the nine DAQ crates (located in the USA counting room).
- The BOC-MON units for the temperature monitoring of the DAQ crates.
- The Pixel LHC Interface Box.
- The Little BoB crate.

Figure 5.12 shows the distribution of the different types of DCS hardware.

5.3.5 The Pixel DCS Control Stations

The control of the pixel DCS hardware is distributed over seven "front-end" PCs which are primarily providing the communication to the hardware, as well as hardware oriented low-level control, and four "back-end" PCs which reorganise the information to provide the high level tools for the operation the detector. Following ATLAS terminology, the PC responsible for the integration into ATLAS DCS is called sub-detector control station (SCS), while all other PCs are referred to as local control stations (LCS).

Each of the front-end LCSs is responsible for the control of one or more different types of hardware, as shown in table 5.2. In total, more than 500 CAN nodes on almost 50 CAN buses, and 60 TCP-IP nodes are controlled by the Pixel Detector Control System.

LCS5 is responsible for opto-heater control and regulation. WIENER crates and opto-heater temperatures are read out and controlled from this machine.

LCS6 is controlling the complete high voltage power supply system. It is equipped with three CAN interfaces. 2 CAN buses are providing control of the iseg crates, i.e. one CAN bus for each counting room with four crates each.

LCS9 controls the VME crates. Since 2011, also the BOCMON units are read out from LCS9.

The control of the 48 low voltage WIENER crates is setup on LCS10. The communication is handled via TCP-IP. The LV-PP4 crates monitoring the module currents is located on the same machine. One CAN bus per rack is in use, giving a total of 8 CAN buses. On the same CAN buses, the monitoring cards for the sensor leakage current, HV-PP4, are read out.

The Regulator Stations and the auxiliary WIENER power supplies are controlled on LCS11. Due to the fact that only two or three nodes per CAN bus are connected, always three CAN buses are powered by the same CAN PSU channel. The power is distributed to the three CAN lines in a passive CAN-splitter crate.

LCS12 controls the SC-Olink crates on 4 CAN buses (one per rack), and the complete temperature and humidity monitoring on another 6 CAN buses (one per level and side). One CAN bus per quarter is responsible for three BBIM crates and one BBM crate. The temperature monitoring of the regulator stations in sectors 1 and 9 is handled by a single BBIM block on level 4, readout by a dedicated CAN bus each.

PC	controlled hardware	number of buses	connected nodes
LCS5	Opto-heaters	TCP/IP	4 Wiener Crates
		1 CAN bus	1 BBIM block
LCS6	High Voltage	4 CAN buses	12 iseg half-modules each
		2 CAN buses	4 iseg crates each
LCS9	DAQ crates	1 CAN bus	10 VME crates
	temperature and humidity	1 CAN bus	9 BOCMON units and 1 BOC-HUM unit
LCS10	Low Voltage	TCP/IP	48 Wiener LV crates
	LV current monitoring	8 CAN buses	18 LV-PP4 blocks each (6 crates)
	HV current monitoring		12 ELMBs
LCS11	Regulator Stations	12 CAN buses	3 or 2 crates each (total: 28)
	auxiliary power	TCP/IP	7 Wiener Crates
LCS12	SC-OLink crates	4 CAN buses	18 SC-OLink blocks each (5 crates),
			1 BBIM Power Supply
	Temperature and Humidity	4 CAN buses	12 BBIM blocks each (3 crates),
			3 BBM blocks each, 3 BBM HMX blocks.
	PP2 temperature (sector 1,9)	2 CAN buses	1 BBIM block each
LCS13	Interlock system “LU” buses	4 CAN buses	30/32 nodes
	Interlock system “BOB” buses	2 CAN buses	8/7 nodes (US/USA)
	CAN PSU control	1 CAN bus	3 CAN PSU crates

Table 5.2: *FE-Control Stations and controlled nodes*

The Monitoring of the Interlock system is handled on LCS13. With up to 32 nodes, these are the most populated CAN buses. These are the four “LU” CAN buses, responsible for the Logic Units and IDBs. In one rack in each counting room, a PP1-Box and an OHI-Box are additionally connected to this “LU” CAN bus. One “BOB” CAN bus per counting room connects the BOB crates, the NICO, the PLI Box and the Boc-I-Box.
The three CAN-PSU crates are also controlled from LCS13.

5.4 Pixel DCS Software

The following chapter gives an overview of the pixel DCS software. This overview includes both, software components developed specifically for the use of the pixel control software, and the integration of centrally provided software packages and their adaption to the requirements given by the specifics of the Pixel Detector hardware.
At the lowest level, the FITs (Front-end Integration Tools) provide the tools to integrate hardware devices into the control system software. They enable control and monitoring on a purely “functional” base. The mapping from hardware channels to “geographical” detector devices is the task of the SIT (System Integration Tool). The transition from channel based operation towards detector based operation is provided by the FM. It summarizes the operation condition into states, so making the supervision of thousands of hardware channels by a single operator feasible, and provides the means for well defined transitions between these states. DDC constitutes the interface between DAQ and DCS. Finally, the connection to databases enables storage and retrieval of configuration and conditions data.

5.4.1 The Front-end Integration Tools

In the Pixel Detector Control System, the Front-end Integration Tools (FITs) are the software tools which build the interface between PVSS and the drivers [85] [86]. The FITs provide the software image of the specific hardware devices, which is at the base of all higher level control software. For each device type the corresponding FIT

provides the means to integrate new devices into the system, and to control and monitor these devices. According to these two tasks, all FITs are divided into an “Integration” part, and a “Control” part.

The integration part is responsible for the management and setup of the device configurations. This includes the creation of configuration files, the mapping between driver address (OPC items) and software representation (data point elements), as well as update rates, filtering, and conversion between hardware specific and human readable values.

Once the device is integrated into the system, it can be addressed by the Control part, containing the graphical user interface for basic control and monitoring on a single-channel level. Given the complexity of the system, this is not a feasible tool for detector operation, but can be indispensable or for low-level debugging.

The Pixel DCS has FITs for all custom built ELMB based hardware devices (“FIT CAN-ELMB”), as well as for the commercial WIENER low voltage (“FIT Wiener”) and iseg high voltage power supplies (“FIT iseg”). For all FITs a Watchdog is provided, which checks the alive status of the individual nodes and raises alerts in the event that communication to any node is lost. The integration of the crates of the interlock system is not handled by the Pixel ELMB FIT, but is based on the framework ELMB package [84]. The control part of the Interlock FIT is also based on pixel software developed specifically for this purpose.

The FITs are divided into the following types:

- FIT CAN-ELMB
  - SC-OLink
  - BBIM and BBM
  - BOCMON
  - LV-PP4
  - Regulator Stations
- FIT Interlock System
- FIT iseg
- FIT Wiener

5.4.2 The System Integration Tool

The System Integration Tool (SIT) [87] introduces the mapping between hardware channels and detector devices into the control system. The mapping is handled in the form of aliases, which provide a PVSS internal mechanism to address any DPE (corresponding to a hardware channel) via a second name (used to store the mapping information). This information is extracted from the connectivity database with a HTML/PHP tool (the “Pixel xml extractor”) into intermediate xml files. The SIT reads in these files, and sets the aliases to the specified datapoint elements. At this point, the cabling information is available inside the DCS and can be accessed by any PVSS component. Besides its main task of setting up the mapping information, the SIT also configures the archiving of the used channels and provides the possibility to display information by detector geography. Even though this still relies on single channel operation without predefined procedures, it was nevertheless already of great use for small system test setups at the institutes.

5.4.3 Regulation of the Opto-heaters

The monitoring and low level control of the opto-heaters and their integration into the detector control system is done by the WIENER, ELMB and Interlock FIT [88]. The regulation requires information above the FIT level, namely the mapping between power supply channel, opto-heater, and opto-boards. The regulation is done in software, using a PID algorithm with online adjustable parameters for each separate heater. Either the opto-heater temperature or the average temperature of the corresponding opto-boards can be used as the process variable for the PID on which the regulation is based. The control variable is the output voltage of the WIENER channel.

A software regulation can be justified by the fact that with stable operation conditions, the necessary amount of regulation is very minor, so that the system can survive a short-term unavailability without danger of extreme overshoot. On the other hand, the granted flexibility to regulate on the opto-board temperatures themselves instead of the heater NTC, and the possibility to mitigate the temperature gradients due to different number of operated opto-boards per heater, has proved to be of great benefit. Although it is generally preferable to regulate on average opto-board temperatures, this depends on their availability. Communication problems between the BBIM crate and the opto-heater project, or between the opto-heater project and the project monitoring the opto-board temperatures might therefore require to use the opto-heater regulation as fallback solution.

Multiple safety mechanisms are in place. The output of the WIENER channels is limited to deliver not more more than 0.5 A. In addition, a configurable software limit per heater is introduced of nominally 50 V, and a software interlock if the heater temperature exceeds 30 °C. If all those measures fail, the hardware interlock system (described in chapter 5.3.3) will shut down any opto-heater with a temperature higher than 35 °C, or all heaters if any opto-board temperature exceeds 40 °C.

5.4.4 The Pixel Finite State Machine

The main tool to operate the Pixel Detector and its infrastructure is the FSM. While it is possible to do any control action in the FITs, only the FSM makes it feasible to operate and monitor hardware channels for 272 readout units with 1744 modules in any practicable manner. A limited number of states allow the user to supervise and distinguish all operation relevant conditions. Predefined commands ensure that all transitions happen in a defined way and prevent operational errors.

Branches exist to control the following subsystems:

- Detector
- Infrastructure
  - Environment
  - Interlock System
  - Opto-heater System
  - DAQ crates
  - DCS crates, in particular PP2

The implementation of the FSM branch responsible for detector operation is the main topic of this thesis and will be described in detail in chapter 6. This chapter will address the monitoring of the DCS infrastructure in the FSM. The DCS crates branch will not be described in this chapter, as it will be outlined for the special case of the regulator stations in chapter 6.5.

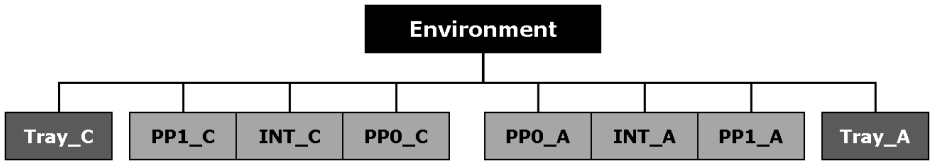

Figure 5.13: *The environmental branch of the FSM.*

5.4.4.1 The Environmental Branch of the FSM

Inside the Pixel volume and the heater trays outside the ID end-plate, a number of NTCs and humidity sensors are installed which are not directly correlated to any detector device. The location of all sensors is described in [91], [92]. The task of the environmental FSM branch is to alert in case of high temperatures, especially in the regions near the cables, or in case of high humidity which can imply the danger of condensation on the electronics.

The volume is divided into four regions on each side: Most sensors are installed in the PP0 region dominated by the number of NTCs on supply and exhaust pipe of each cooling loop. The NTCs most susceptible to register high temperatures are the "type zero gap" NTCs, due to their proximity to type zero cable bundles. In the Intermediate region between PP0 and PP1, and in the PP1 region, NTCs are located near type 1 cable bundles, and in the gas volume. Humidity sensors are installed at the beam pipe support structure near the BCM modules (of type Xeritron), and inside and outside the ID end-plate (Hygrotron and Honeywell). In the heater trays outside the ID end-plate, one hygrotron humidity sensor is installed per octant.

The state and status of this branch is determined by the device units defined at the NTC or humidity sensor level.

5.4.4.2 The Interlock System in the FSM

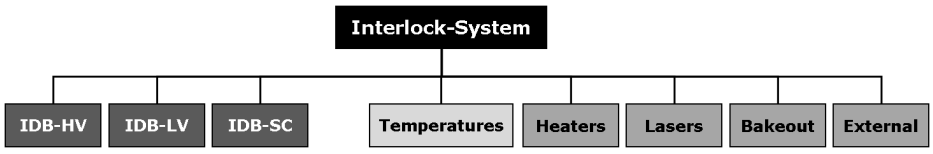

Figure 5.14: *The interlock branch of the FSM.*

The Interlock branch of the FSM has been described in [89]. It monitors the presence of interlock signals as well as the condition of the system itself. The device units correspond mostly to ELMB nodes or crates. The states are determined on the one side by the presence of interlocks resulting in the states READY or INTERLOCK, and on the other side by the state of the monitoring, where a failure is indicated by the fact that signals do not have the expected value. This results in the state UNKNOWN which can be caused by an un-powered crate, or a communication problem on the software or CAN level.
The structure of the tree is indicated in figure 5.14. The aim was to emphasize the effect that a given interlock has for the system, and the origin of the interlock. Therefore the pure crate-related structure was broken up for the Logic Units. As they receive, besides the individual temperature signals, also all global input signals, these device units would for a crate based tree be in INTERLOCK state for nearly any interlock, removing any informative value on this level of the tree hierarchy.

Therefore, for the current tree structure, on the top level the three types of Interlock Distribution Box (IDB) were introduced as nodes, to clearly indicate which type of power supply is affected. This is supplemented by the user interface, which additionally displays the IDB channels according to the affected readout unit.

The Logic Units are reduced to temperatures as cause for an interlock on the detector, grouped by interlock quarter. A high module temperature will cause an INTERLOCK state in the corresponding temperature quarter node, and the effect will be visible in the IDB-LV and IDB-SC nodes.

Additional nodes are also motivated by cause and effect of the interlock.

The bake-out node contains the three bake-out crates (two BOB crates and Little BOB). The INTERLOCK state is triggered by a high B-Layer temperature or an active DSS2 signal, and the effect is the shutdown of the bake-out carts. The Lasers node consists of the BOC-I-Box and laser related PP1 Box signals. The Opto-heater node takes into account the signals from NICO and the OHI-Box for the input, and the NICO output. The External System node monitors the IDSS-0 inputs into the Pixel Interlock System, and the two signals of the PLI-Box for the interaction with the LHC.

5.4.4.3 The Opto-heater FSM

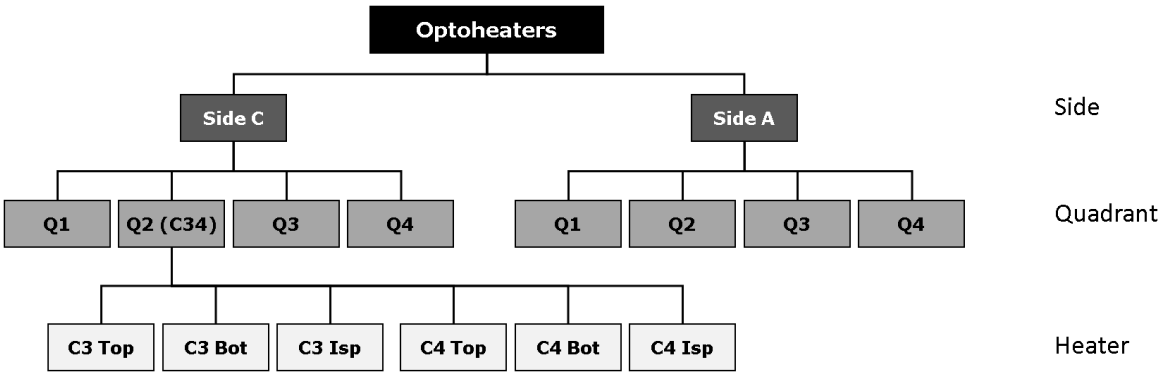

Figure 5.15: *The opto-heater branch of the FSM.*

The opto-heater FSM branch provides the means to operate and monitor the opto-heaters on a high level. The FSM provides the commands to switch the heaters on or off, and change the temperature set points. The states accordingly signify whether the heaters are off or on, and whether stable regulation conditions have been reached. Problems with Control and Communication are treated as UNKNOWN state. The status alerts to deviations from nominal operation conditions. For operation, the heaters are grouped by side and quadrant, as shown in figure 5.15. Each quadrant contains the six heaters of that service quarter panel as device units.

5.4.4.4 The DAQ crates

The DAQ node in the FSM infrastructure branch consists of the nine BOC crate nodes and the TTC crate. The supervision of the DAQ crates includes the monitoring and control of the VME power supplies, and the temperature monitoring of the BOC cards. Therefore, the combination of the crate device and the temperature device builds the crate control units. The VME crate control is based on the standard framework component, while the state and status of the temperature device unit depends on the availability of communication to the ELMB, and the alert settings for the temperature value. Control actions include powering on and off the crates and an reset of the VME bus.

As the detector state depends on the crate state, detector commands on a DAQ crate granularity are provided in this branch.

The mapping of readout unit to readout crate can be seen in figure 5.16.

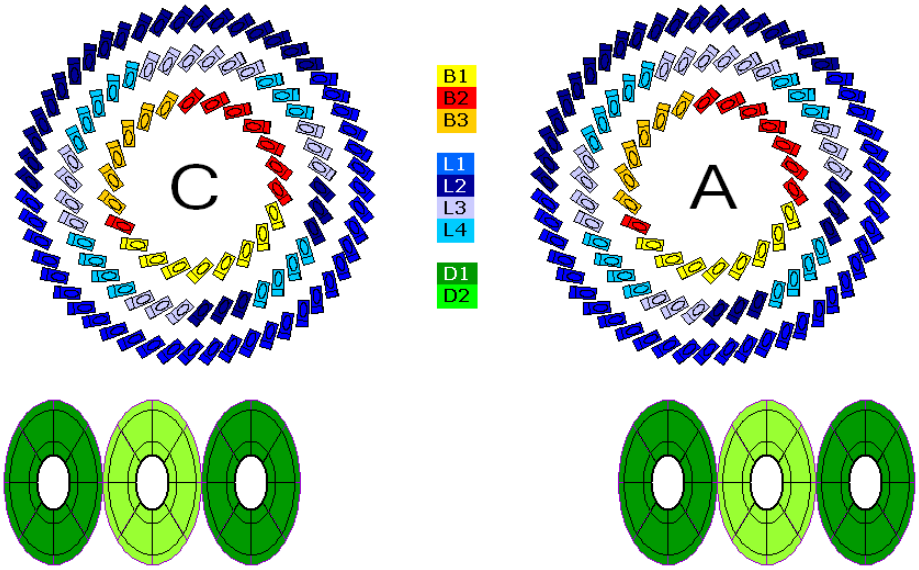

Figure 5.16: *Readout Units per DAQ Crate.*

5.4.5 DAQ-DCS-Communication and Monitoring of the DAQ Hardware

As a way to synchronize DAQ and DCS, a package for the DAQ-DCS-Communication (DDC) is provided. It uses DIM for communication due to its platform independence.

DDC is split into three separate sub-applications

- Command Transfer (CT)
- Message Transfer (MT)
- Data Transfer (DT)

Command Transfer is unidirectional from DAQ to DCS, and is split in transition commands, which are triggered by a state transition of the DAQ FSM, or optional non-transition commands which can be triggered asynchronously from the DAQ FSM. The receiving unit on the DCS side is the DDC-device unit which is obligatory on the TTC partition level in the DCS FSM. The corresponding object type is maintained by ATLAS DCS.

Messages can be passed from DCS to the DAQ system via the DAQ MRS (Message Reporting Service). This is used to report any high level DCS state changes to the DAQ, in particular any time the DCS state deviates from nominal operation conditions.

Data Transfer can be used in both directions, to exchange data between DAQ and DCS. On the ATLAS level, e.g. information regarding the data taking state are subscribed. Also, in the ATLAS LHC project, sub-detector specific DAQ information, like luminosity and detector occupancies are monitored.

For the Pixel Detector, besides the standard uses, DDC is required mainly for the support of Pixel DAQ calibration. Additionally, some of the values read by the DAQ should be transferred to the DCS for archiving. With the beginning of ATLAS data taking with beam, an additional area of application has arisen for the usage of DDC, namely, to manage the synchronization of the pixel pre-amplifiers by the ATLAS DAQ and high voltage by Pixel DCS.

In summary, the following features of DDC are actually in use:

- data transfer from the Pixel DCS to the Pixel DAQ.
- command transfer from Pixel DAQ to Pixel DCS.
- data transfer from Pixel DCS to ATLAS DAQ.
- no direct subscription to values published by ATLAS DAQ (but might be accessed via ATLAS DCS).

For the support of Pixel calibration, interaction with DCS is needed by the DAQ mainly for the tuning of the optical link:

- Inlink scan: value of the Opto-board PiN Current. Handled by standard DT.
- 3D BOC scan: Setting of VISet. Pixel specific implementation of CT.
- other scans: reading of varying values (e.g. temperature) for additional information (standard DT).

DCS information is critical for the analysis of the Inlink scan. During this scan, the laser channels belonging to one opto-board are switched consecutively by the DAQ, and the DCS feedback of the measured opto-board PIN current is used to evaluate if a given channel is transmitting light or not.
Additional information about DCS conditions is read also during the execution of a number of other scans, e.g. during the Inlink-Outlink scan opto-board information is read and saved with the scan.
For the 3D-BOC scan, it is required that the DAQ can change automatically the value of VISet during the scan. While the regular BOC scan determines the best setting for threshold and delay for sampling the signal arriving from the detector at a given VISet, the occurrence of errors depends also on the laser power. Therefore, to determine the optimal settings, this scan steps through VISet values from 0.725 V to 0.9 V in steps of 0.025 V. A pixel specific version of the command transfer was implemented to support this [90].
In addition to the online changing of VISet during scans, it is necessary to apply the nominal settings before data taking. Due to the fact that it is the DAQ that determines the optimal set point, but the DCS that actually sets the value, the current DAQ and DCS configurations have to be synchronized. This is currently still done manually, by sending DDC commands on the VISet level, but it is foreseen to provide a set of commands for database operations that enable the synchronization and loading of specific configurations.

Besides the use for pixel calibration, during data taking, the synchronization with the DAQ after the declaration of stable beam is handled by DDC in a custom solution which was implemented in cooperation with ATLAS DAQ and ATLAS DCS. DDC-DT is used to signal to the ATLAS DAQ that DCS is ready for data taking, and that the module configuration for data taking should be sent.

5.4.6 Databases

In a large experiment, it is not always feasible to store all information inside the DCS itself. Pixel DCS interfaces to three databases, which can be categorized according to their task:

- The ***Connectivity Database*** for the storing of the mapping information.
- The ***Configuration Database*** for the storing and applying of detector settings.
- The ***Conditions Database*** for the archiving of detector conditions for system debugging and offline data analysis.

The Connectivity Database

The Pixel Connectivity Database uses a special implementation [94] based on CORAL[10)], for the storage of connectivity data into an Oracle database. Both, DAQ and DCS information is stored in the same database. Tags provide the possibility for bookkeeping of the time intervals where a given connectivity was in use. The data is stored in a hierarchical way with additional "soft-links" between the hierarchies. A web interface provides a visualization of the tables content into connectivity graphs.

The main interface between the connectivity database and DCS is provided by a tool that "translates" hardware channels into detector elements and creates xml files from this information.

The Configuration Database

Even though most settings are constant during early detector operation, it is nevertheless required to provide different configurations for different scenarios or for changing calibration. These configurations are stored in a configuration database, which is an Oracle database dedicated for each sub-detector. A tool to access the database is provided by the framework. Structures of properties to store – the recipe types – can be defined by the user. They can contain values as well as PVSS configs, like alert limits, setup of archiving, or OPC connection properties.

A saved configuration can then be loaded later at the appropriate time, completely or only partially. The two most important use cases for the Pixel Detector are the default configuration, with all set values except VISet, and the recipe type containing only VISet. The former was used during commissioning of the cooling system to switch between nominal and high power configuration, and the latter is intended to store the optimal VISet values determined by the DAQ. Besides configuration of the hardware, the configuration database is used to store DCS related items like thresholds for software interlocks and the FSM state limits.

The Conditions Database

Two types of values need to be stored for later access. The first type are values which will help to understand and debug detector related conditions. The second type are values that are required for offline analysis of physics data. Relational databases are used for both types of values. All values are first stored into the same Oracle database, the "PVSS Oracle Archive". A process – PVSS2COOL – running at regular intervals is responsible for transferring specified data from the PVSS Oracle archive to the conditions database. This database is accessed by the COOL[11)] API. The data in the conditions database comes with an interval of validity and is grouped in "folders" which are setup based on PVSS data point types. For the Pixel Detector, information about depletion voltage, module temperature and module FSM states are stored into the conditions database, where it can be accessed by Athena based analysis routines.

To access historical values for system debugging from the PVSS Oracle archive, various tools exist. The web-based DCS Data Viewer (DDV) allows the user e.g. to formulate SQL queries by GUI based selection, has inbuilt protection mechanisms to prevent overload of the database by bad queries, and provides output in multiple formats like graphical trend, ascii files, or root files [93].

5.4.7 The Expert System

While the FSM allows the shifter to operate and monitor the Pixel Detector, it still requires a basic level of understanding to diagnose error conditions. The aim of the expert system is to provide some of the expert knowledge to the shifter. For this end, an expert system for the Pixel Detector was designed and implemented [96]. The

[10)] **CO**mmon **R**elational **A**bstraction **L**ayer [95]
[11)] **Co**nditions **O**bjects for **L**CG (LHC Computing Grid)

expert system can be divided into a data taking component, a set of rules, a rule based decision making component (inference machine), and the interface to the user. For the Pixel Advisor it was decided to use DIM as a method to acquire data, due to the fact that it is already used inside DCS and DAQ. It is under considerations to acquire data directly from the database, which would reduce the load on the online system and provide easy access to historical data.
The set of rules can be divided into internal rules for the working of the expert system, and rules for diagnosing detector related problems. It is foreseen that users can insert new rules which will enter the rule base after adaption and checks by experts, and a user based ranking of rules based on successful solutions.

Among the advantages of the expert system is the fact that it can run to a degree independently of the detector control system. This makes it possible to concentrate in its development on intrinsic features, without the fear to disturb detector operation. The expert system can provide explanations for conditions which are not part of the FSM state system. While it is impossible in the FSM to have a separate state for every identifiable operational or error condition, the number of rules in the expert system is not subject to such a restriction. The FSM encapsulates states to a certain degree inside the device units, and has by design only a very limited memory of what happened before an error occurred. For the expert system, it would be possible to access data of related devices after an error occurred from the database.
However, currently the expert has only an extremely limited set of rules available for operation.

Chapter 6

The Finite State Machine

The Finite State Machine (FSM) is the main instrument for the operation of the ATLAS Pixel Detector. As such, it has to fulfill not only the specifications and requirements concerning the safety of the detector, but simultaneously those imposed by the fact that operation has to be handled by non-expert shift-personnel. This concerns not only operation of the detector itself, that will be described in the first part of this chapter, but also expert procedures and recovery actions, described in the second part, 6.5. The design and implementation of the different aspects of the pixel FSM structure under consideration of the given requirements was the core of this thesis.

6.1 Requirements

For the operation of the Pixel Detector through the FSM, several requirements need to be fulfilled. These can be roughly grouped into the following categories:

- operation as a sub-detector of ATLAS
  - safe operation and supervision of the detector by non-experts
  - common abstraction of operation conditions among the sub-detectors
- standalone operation of the Pixel Detector
- the requirements which arise from pixel specific detector hardware and pixel operation
- interaction with external systems which affect pixel operation (e.g. LHC, evaporative cooling)

The details involved in these points will be described in the following paragraphs.

6.1.1 Requirements by ATLAS DCS

One set of requirements are imposed by ATLAS central DCS. Those arise from the fact that all sub-detector Finite State Machines must be integrated into the global ATLAS FSM. In addition to hard requirements, a common "look and feel" is desired to account for the fact that all nodes are foreseen to be operated at one point as a single unit, the ATLAS detector, by one shifter. One central concept of the ATLAS operation model is the setup of two parallel FSM hierarchies. In addition to the state, which represents the operational situation, especially with respect to data taking, the status is supposed to indicate problems related to safety-relevant issues.

6.1.1.1 ATLAS states

To make the operation of the whole ATLAS detector by a single shifter possible, on the top levels of the hierarchy the operation must be very abstract. This means that there is a limited set of operational states, which can not take into account any sub-detector specific peculiarities.

The allowed, respectively required, states are:

- SHUTDOWN: The sub-detector is completely off. No power into the cavern.
- READY: The sub-detector DCS is ready for data taking.
- STANDBY: The sub-detector is as near as possible to the nominal data taking conditions, while still considered safe during unstable beam conditions.
- TRANSITION (optional): A command is currently being executed.
- UNKNOWN: The state of the sub-detector is not known (usually due to loss of communication).
- NOT_READY: A "catch-all" state that does not fall into one of the above categories.

Figure 6.1 shows the state diagram.

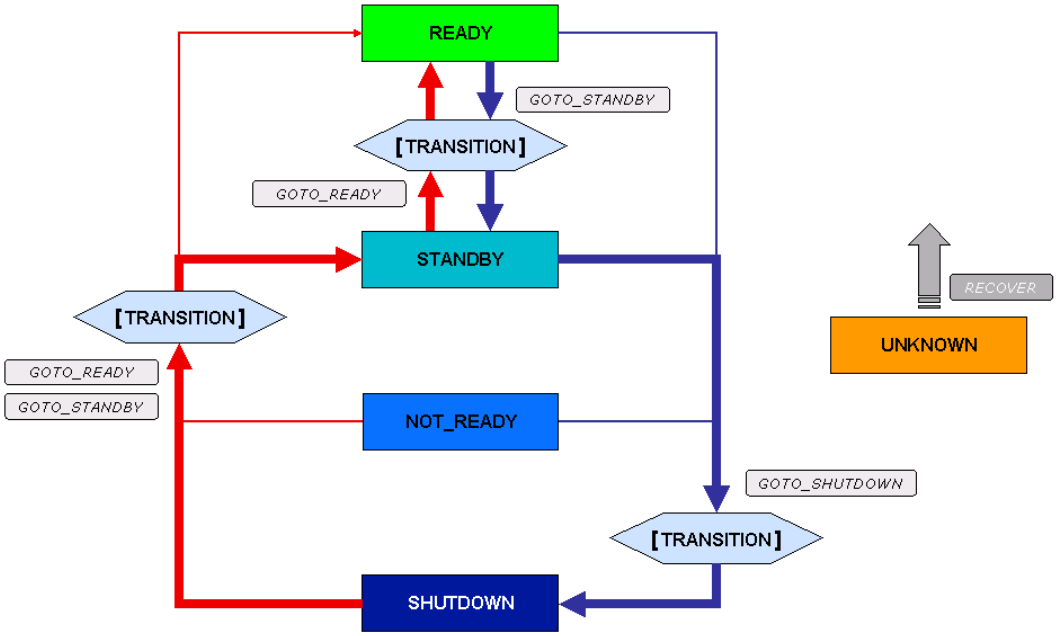

Figure 6.1: *ATLAS States.*

Shutdown is the state where no power arrives at the detector, or even to the cavern. Then the operational procedures vary, depending on whether there is beam or not. When powering up the detector for operation, it is brought to READY state, or, for detectors which are sensitive to beam, to STANDBY state. STANDBY state is not foreseen for sub-detectors not sensitive to beam, while the others can only safely be brought from STANDBY to READY when there is no beam at all or stable beam. Depending on the sub-detector characteristics, a direct transition to READY might also be possible. The TRANSITION state is optional, indicating ongoing command execution. The state NOT_READY means that the sub-detector is for some reason not ready for data taking, but neither in STANDBY, nor completely shutdown. This state encompasses also all sub-detector specific intermediate states. After each data

taking period, the detector is brought back to the safe STANDBY state, and in case of longer shutdown periods, back to the SHUTDOWN state.
UNKNOWN is not a desired state reached by a transition, but means that the state is not known, usually due to loss of communication, e.g. due to network problems or hardware failures. Ideally a RECOVER command is implemented, but might not be feasible if there is no "default" failure mode with well-defined recovery procedures. This is in a large part due to the fact that actions are required which are not under the control of the FSM or even pixel DCS, e.g. a broken cable or a defect network switch.

The sub-detectors which have a STANDBY state are shown in figure 6.2.

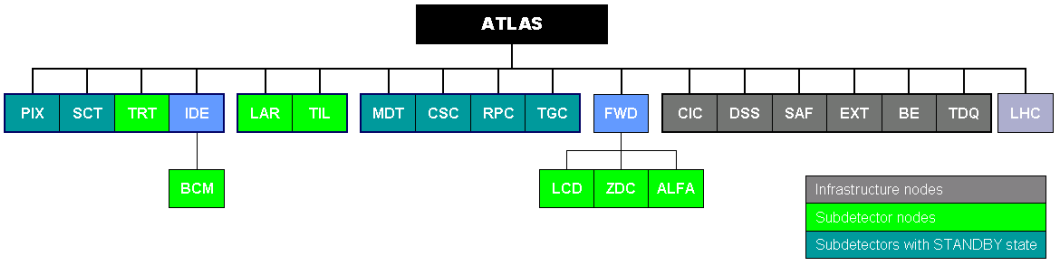

Figure 6.2: *Beam sensitive (teal/blue) and non-sensitive (green) sub-detectors of ATLAS. The children of the ATLAS node can be divided in sub-detectors and infrastructure nodes (grey). For ATLAS to be in* STANDBY*, the silicon detectors and the Muon systems have to be in* STANDBY*, while all other nodes must be* READY*.*

6.1.1.2 The Status

The concept of an independent status hierarchy makes it possible to indicate a problem e.g. due to over-temperature, without immediately leaving the READY for data-taking state. In many cases state and status will not be entirely independent. A high temperature might eventually result in an interlock, making that part of the detector unavailable for data taking after all. In general, a state that does not equal READY can be completely appropriate for a given circumstance, while a status that is not OK will normally require an action to recover a problematic situation.

The values for the status are predefined, and are indicating the severity of the problem. The status OK means that everything is inside normal operation parameters, and no action is required to guarantee the safety of the detector. WARNING defines a mild error condition that should be followed up during working hours. An ERROR in contrast should be followed up as soon as possible, as it severely threatens the functionality of the system. FATAL is the most severe status, and always implies that the system can not work anymore in that condition.

6.1.1.3 Central Operation

Ultimately, the ATLAS detector will be operated by a single shifter. The control of the FSM will be on the Global Control Station (GCS), which includes references to all the sub-detector top-nodes running on the Sub-detector Control Stations (SCS), as well as to the common infrastructure items. The communication to the ATLAS- and sub-detector DAQ is implemented on the partition level, which is for that reason the lowest level up to which standard ATLAS states are required. The control of the hardware happens on the Local Control Stations (LCS), and here the sub-detectors are free in their implementation, to adapt to the sub-detector's needs and special features, as long as the interfacing to ATLAS is guaranteed on the TTC partition level. On the lowest level nodes, the actual interface to the (hardware) devices is implemented. For devices, the default states OFF, ON, and UNKNOWN are foreseen.

The partitioning modes introduced in chapter 5.2.6.1 make it possible to adapt to different operation modes. In case of calibration or standalone running, the sub-detectors can be excluded from the central FSM. During combined data taking, the control over all sub-detector FSM nodes is mandatory at the central DCS desk. During the commissioning phase the control is shared, while ultimately it is supposed to be exclusively at one single place.

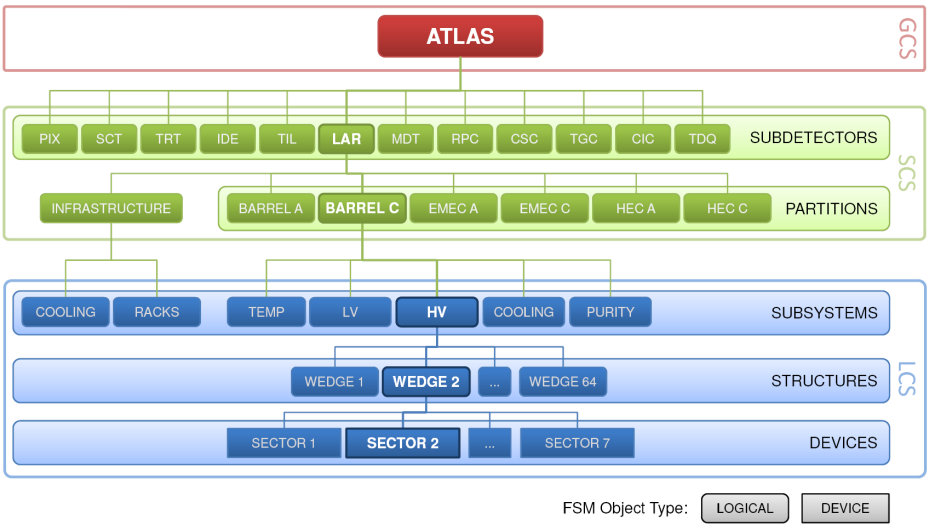

Figure 6.3: *ATLAS FSM architecture (shown for the example of the LAr subsystem) [21]. The DDC connection (not shown) to the respective TTC partition is done on the sub-detector control stations.*

6.1.2 Requirements by Pixel Operation

This chapter describes the general requirements which are determined by the actual (standalone) operation of the Pixel Detector. The details of these requirements and how they are actually implemented will be described in chapter 6.3. There are two major categories, regarding on one side the FSM as the main tool for detector operation by the shift crew, and on the other side the more operator-independent safety relevant aspects.

6.1.2.1 Operation by Shifters

For the operation of the Pixel Detector, the FSM has to fulfill the task of giving a clear overview of the current state of the detector. It must provide the possibility to change that state in a predefined, safe procedure corresponding to the current operational situation. This includes the necessity to prevent that components are powered out of order, and to ensure that all required steps have been executed before a detector component is switched on. A further point is the fact that this has to be done by shift personnel which is usually not composed of detector experts, and may not be aware at all times about the inter-dependencies of the various components. A tool must be provided to make the continuous monitoring of the overall state of the complete system possible, with the possibility to easily recognize and localize error conditions. This includes the required possibility to mask known error conditions, so that new faults can still be spotted easily. A set of commands must be provided for the safe transitioning of the Pixel Detector between the defined states. Ideally, this applies also for easy transitions of only parts of the detector, dictated by the various different modularities of all the components.

6.1.2.2 Interaction with the Pixel Data Acquisition

The actions taken by the pixel DAQ have direct and indirect effects on the detector. Therefore the detector control system can in turn give feedback to the DAQ in case of regular changes, or alert the shifter to the fact that unwanted effects may have occurred. An additional requirement imposed by the way pixel DAQ and DCS interact, is the support for the execution of optical scans.

The effects of the DAQ, namely the module configuration state, are obviously a crucial condition when it comes to issues concerning data taking operation of the Pixel Detector.

6.1.2.3 Interaction with External Systems

The last set of operation requirements is to facilitate the interaction with external systems for the operator. While the base must be laid in the design of the FSM state model, the actual procedures will be described in 7.2.

6.1.2.4 Thermal Stability and Safety

During operation, the thermal stability of the detector should be maintained, and thermal cycling has to be avoided as much as possible. The worst instance of thermal cycling is a failure of the cooling system, in conjunction with its following restart. This leads to large temperature gradients, and has a non-negligible chance of directly causing unrecoverable losses of detector hardware, mainly on the front-end level. Due to the nature of the evaporative cooling system, it is only to a very limited degree possible to avoid the sharp temperature drop. Over-temperature, on the other hand, may worsen effects due to radiation damage. The objective to maintain thermal stability always is a prerequisite for all considerations on implementation.
The hardware interlock system prevents over-temperatures of more than 40°C. The rather high value of the interlock temperature threshold is owed to the fact that part of the early commissioning had to be conducted with the evaporative cooling system not yet fully operational. This is one of the reasons why ultimately the aim is to avoid the intervention of the hardware interlock, and have an additional layer of safety at a much lower configurable software interlock threshold. In addition, the hardware interlock does not ramp down the power supplies, but kills them in an not predefined order. If the interlock is caused by partial power cut, even the granularity may be completely arbitrary with respect to detector geography. After the system was partially switched off, or after a complete shutdown, it has to be ensured that the status of the low-level hardware devices and the information inside the DCS software is synchronized, so that the system is always in a well-defined state.

Due to the design of the hardware interlock, concerning pixel internal conditions regarding detector safety, the interlock system mainly monitors only temperatures. Therefore undesired conditions that are not directly correlated to an over-temperature, like over-voltages or over-currents, can only be handled by software interlocks. These software interlocks will be described in detail in chapter 7.2, while the issue of thermal stability has also direct consequences for the state model.

6.1.3 Summary of the Requirements for the Pixel FSM

In summary the requirements imposed by Pixel operation are as follows:

- prohibit unsafe operation conditions, like over-currents, over-voltages, or over-temperatures
- maintain thermal stability
- safe operation and supervision of the detector by non-experts

- act as a layer between the low level hardware oriented parts of the detector control system and the operation and supervision related part.
  - Ensure the defined state of the low level devices and
  - eliminate the need that the operator should have to know the special internal features of the specific hardware or software implementation.
- provide additional safety layer from information of
  - Evaporative Cooling System
  - Detector Safety System, DSS
  - LHC
- facilitate the interaction with external systems including
  - ATLAS DCS
  - pixel and ATLAS DAQ
  - Evaporative Cooling System
  - LHC

6.2 The Underlying Components

Due to the fact that often the different groups still face similar problems and tasks, many of the components used for the pixel DCS and FSM are actually built on common tools. This starts already with the framework for the SCADA, PVSS, which was chosen by the JCOP group as standard for all the LHC experiments, to focus common software tools and developments for hardware control and detector operation and so saving manpower for development and maintenance[1)]. The components underlying the pixel FSM are described here, starting from the FSM component provided by the JCOP framework, and ending with the infrastructure for the pixel FSM developed and described in 6.2.3.

6.2.1 The JCOP Framework Component "Controls Hierarchy"

For the building of the FSM, the JCOP framework component "Controls Hierarchy" (see also chapter 5.2.6.1) is used. This allows one to build a treelike hierarchy of objects which each represent a Finite State Machine. The behaviour of the objects is defined by SMI code. One process represents each control unit, while the logical units are controlled by internal sub-processes of the parent process. The behaviour of the devices is defined in PVSS scripts, and so their behaviour is independent of the FSM infrastructure overhead. The SMI processes use DIM for communication, and for the initialization of the communication a DIM Name Server is required.
In the tree, commands are always propagated down the tree, from the parent objects to their children, while the states are propagated upwards. It is possible to partition the tree for an independent operation, or to just completely exclude a faulty part of the system temporarily. In addition it is possible to "ignore" an object, i.e. to not propagate its state upwards, or to operate it in "manual" mode, which means that the state will be propagated upwards, but no commands will be received from the parent.

[1)] These common tools provided by the JCOP group are referred to as the "JCOP framework"

6.2.2 The ATLAS FSM

On top of the JCOP framework component, an ATLAS-specific package (described in [97]) has been built. This introduces the concept of parallel trees for state and status - one for the operational state of the system, and one to propagate the error condition. This component defines standard objects and their states, which can be used as the interface between common ATLAS operation and the sub-detectors. The interaction between DAQ and DCS is handled by a special device unit which propagates the state of the partition to the DAQ system via DDC, and which foresees the handling of DAQ transition commands. ATLAS provides a centrally maintained DIM Name Server for the FSM operation of all sub-detectors. Finally, the panel layout is specified to guarantee a common design throughout the community and various common widgets are provided.

6.2.3 Internal Software Infrastructure of the Pixel FSM

A basic concept for an FSM for the Pixel Detector was presented in [98], and validated for a setup representing 10% of the final system, with a rudimentary set of commands and states. Based on a study of performance versus tree structure and depth, the structure of the tree below the partition level is motivated, as well as the definition of the device units below the readout unit level.

Figure 6.4 shows the chosen structure for the pixel FSM hierarchy. Following the requirement by ATLAS DCS to have the TTC partitions on the level below the sub-detector node, the next level in the hierarchy is comprised of the mechanical structure, i.e. the three layers in the barrel region, and an object for each disk. Due to the heavy dependency on the cooling system, the cooling circuits form the next level. As the readout units are the smallest unit which can be independently operated from the DCS point of view, they are the lowest level logical object. The device units are on the module and opto-board level, together with the primary power supply units of iseg and Wiener.

The generation of the tree was automatized by the use of the geographical structure stored in the aliases. This has proved especially useful during the connectivity test, when always-changing sub-parts of the detector had to be tested and operated.

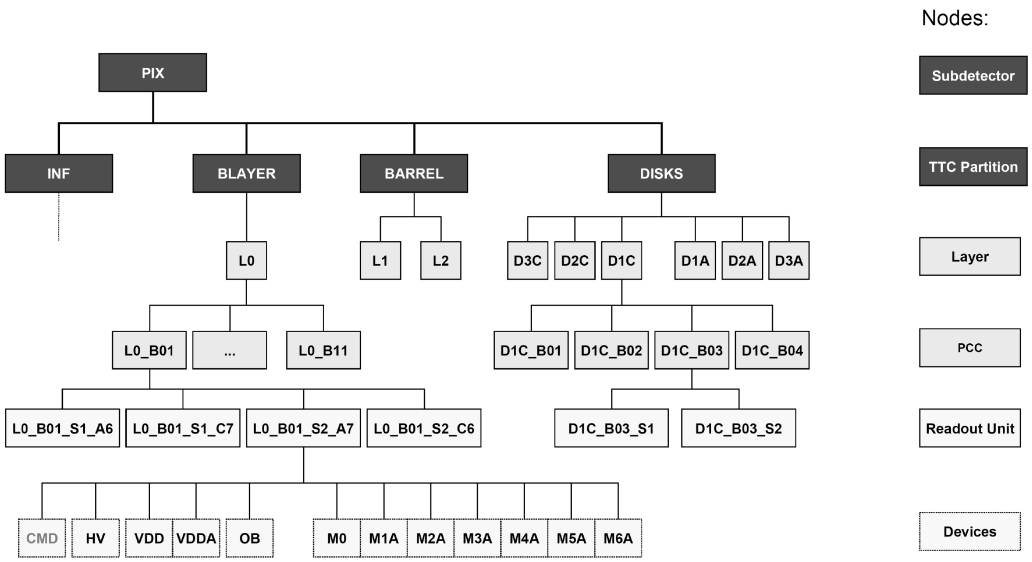

Figure 6.4: *Hierarchy of the Pixel FSM.*

Due to the fact that the device units are rather complex compared to the foreseen case of single hardware channels, the state and status calculation is done in a non-standard way. The first issue to be considered is the sequential readout of the ELMB. A device unit that is depending on more than one channel can not react on every single value change, as this would mean that some channels still have values corresponding to the old state, while others are already updated. To avoid these intermediate undefined states, the values on which the state calculation is based are obtained all at the same time in fixed intervals, with the option to filter out one intermediate undefined state. As the pixel systems are grouped by hardware type, the different values belonging to a detector unit are provided by different systems. Due to PVSS internal limitations, the chosen method of state calculation gives the additional advantage of being more robust against system disconnections and provides a more intuitive access to the value readings which makes the code more maintainable. This relinquishes the possibility for load reduction during quiet periods, but on the other hand the load will behave more evenly over time without extreme peaks during transitions. Last, this decouples the PVSS based state calculation from the FSM-infrastructure, so that changes do not require a regeneration of the complete tree.

The module state-status combination depends on the current consumption of the digital and analog part of the front-end electronics, and can be described in a two-dimensional "matrix", which is easily adjustable to changing requirements. This matrix also contains a descriptive field to give an indication what caused the current state and status to be applicable. To reduce the load on the system, an internal smoothing of state and status was introduced. This means that the control scripts are keeping track of the current state and status internally. Only a real change of operation conditions is passed to the actual FSM mechanism and propagated through the tree. In the other direction, the propagation of commands is handled at the readout unit level in a dedicated device unit. The task of this device unit is the synchronization of commands, since in the standard implementation it is not possible to guarantee an order in which the commands are sent to the individual children.

With this basic structure of the pixel FSM, the system test of an end-cap, which represents about 10% of the final system, was operated in winter 2006/2007. Although the operation mode was mostly a non-permanent operation constantly supervised by detector experts, it was shown that the technical concept for the implementation of the FSM operation provided an effective and promising basis for the operation of the detector, and under the assumption of scalability it would be well feasible to run the FSM with a small number of PCs in the pit.

6.3 Operation Model

The previous chapter has given the context of the software infrastructure building the basis of the Pixel FSM. For the day-to-day operation of the Pixel Detector, a set of states and commands has to be specified and defined on the readout group level, to make safe detector operation possible as well as continuous monitoring of the detector situation. The number of states should be kept as low as possible but still allow the user to distinguish all relevant operational conditions. The basic assumption is that detector operation is based on the collective operation of many readout groups.

For the implementation of the FSM, the first necessity is to identify the states which should be easily distinguishable for the operator due to operational reasons, and the preconditions which have to be met for moving into higher states. The starting point for command definition, which is the second cornerstone of the FSM implementation, lies also at the hierarchy level of the readout groups. The reliable determination of the operational state of the detector, and the commands to safely transition between them is the fundament of the FSM. This was one of the main aspects of the work presented in this thesis.

6.3.1 Definition of the Readout Unit States and Commands

An important conceptual point for the state definition is the fact that a state always refers to the current, actual condition of the detector, and not to the desired or requested state. I.e. even after a "GOTO_READY" command

the state does not necessarily result in “READY”, but instead it takes into account the factual situation. The advantage of using this strategy is that it is always easily possible to monitor the current state of the detector. The knowledge and history about the requested state on the other hand is not as apparent.

The states which are dedicated to pixel operation can be roughly grouped into powered states including STANDBY, un-powered states, and intermediate states. Intermediate states are not supposed to be persistent during data taking periods, but are not purely transitory during startup and shutdown or during calibration. In addition, there are states describing the condition of the detector control system itself rather than the operational state of the detector.

Figure 6.5 shows the state diagram resulting from the operational requirements and considerations which will be discussed in the next paragraphs.

Effectively, the final concept is the result of an iterated process during the period of this thesis, of evolving specifications and requirements.In the beginning, special scenarios had to be supported (connectivity test, partial operation) which lost in importance with time. On the other hand, additional aspects of operation gained in relevance and imposed new requirements. This includes e.g. operation with beam or handling of the optical link during continuous operation.

In the following, only the final implementation will be presented.

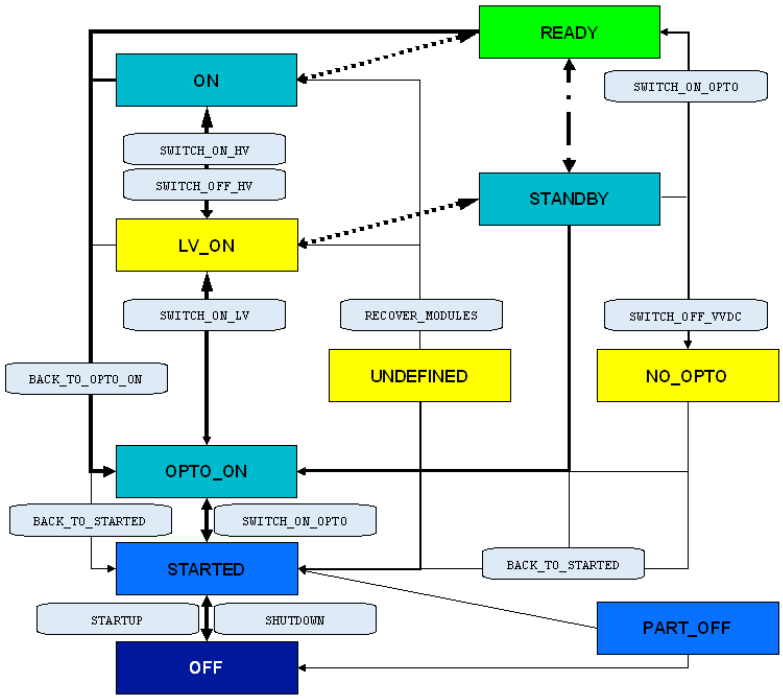

Figure 6.5: *Simplified State Diagram for the Readout Unit.*

States describing the un-powered Detector

From the detector perspective, the most basic state is when there is no or almost no power going into the cavern, which is the state which should be present during longer shutdown periods, or in case of adverse situations, e.g. if there is the danger of fire, indicated by smoke alarm, or if the generated heat can not be dissipated due to a cooling failure. This power-less state is what corresponds to the ATLAS state SHUTDOWN, and is called OFF on the readout group level. In the complete SHUTDOWN state, not for all hardware types control and monitoring

is possible. Namely the ELMB and FPGA of the PP2 crates can be un-powered in this state, and also the iseg crates can be off. At least the primary power supplies of the readout unit are all off. Therefore, when power is reestablished, a startup procedure is necessary. One reason is to bring hardware and software, mainly with respect to the PP2 control, back into a synchronized and defined state. This recovery procedure is implemented into the *STARTUP* command, and brings the detector from SHUTDOWN to STARTED. Besides synchronizing hardware and software settings, for this command a procedure for checking the sense lines is implemented (see also paragraph 6.3.1). In the STARTED state still no power arrives at the detector, but the system is ready to be operated without any further recovery procedure. A hardware interlock can leave the detector in a state where either Wiener or SC-Olink is powered, while the other is un-powered. This is not a desired state, however it is still a safe state. For this reason, this condition is assigned a separated state. This partially un-powered state is called PART_OFF. In this way, all the safe conditions will result in a state that allows the user to recognize that the detector is un-powered.

States describing the nominally powered and configured Detector

From an operational, in particular data taking, point of view, the next relevant state is when the detector is fully powered, and therefore ready with respect to the detector control system. However, in the case of the Pixel Detector, the modules come up in an un-configured state after a power cycle, and need to be configured for data taking. Normally, the configuration state of a module can be reliably determined from the typical current consumption. Even though only the DAQ can configure modules, and therefore this configured state can not be reached by DCS actions, it is nevertheless desirable to give the DAQ shifter feedback whether or not the modules have successfully reached their final state. These two states, where the detector is fully powered, and finally ready for data taking, are called ON and READY respectively.

Due to the sensitivity of the pixel modules against beam related dangers, the state STANDBY must be clearly recognizable. This state is characterized by the fact that the detector is in a safe state during unstable beam conditions, but otherwise ready for data taking, so that no beam time is lost once stable beams are declared. For the Pixel Detector, this is given when the high voltage is off, and in addition the pre-amplifiers in the pixel cells are killed. This is another action that can only be performed by the DAQ in a regular way. [2)]

States describing intermediate or irregular Conditions

In addition to the STANDBY state, there is the corresponding situation where the module configuration does not match the desired one, or the modules are un-configured altogether. The case of a wrong configuration must be clearly distinguishable from STANDBY, and is called LV_ON. The case of un-configured modules is not so unambiguous, as the detector is safe, and the same action that would transform a module from STANDBY to READY will also convert it from the un-configured state. If the module is un-configured because it is disabled in the DAQ, the configure transition will however not act on this module. In this case the condition "safe state" is fulfilled, but "almost ready for data taking" not. As the DCS can not determine reliably why the module is un-configured, this state is treated the same way as a wrong configuration in the LV_ON state.

During powering up of the detector, the threshold implementation of the opto-board (see also 3.3) makes it necessary to have the opto-board powered before switching on the modules. This is necessary due to the fact that the DORIC will adjust the threshold according to the power delivered by the off detector lasers. If no laser power at all is present at the DORIC, the threshold will be lowered to the point where noise is interpreted as signal, which will then be sent to the modules and might be interpreted as actual commands. This will result in erratic current consumption by the modules, and over-currents might occur. This alone does not entirely justify a dedicated state, as this issue could also be handled internally during command execution.

[2)] In an emergency, the modules can be power cycled, and will come up un-configured and with killed pre-amplifiers. However, this is undesirable due to the resulting thermal cycling of the modules. Also, if the DAQ is not available to reconfigure the modules, the total power will be reduced to about a third.

A second use case for a separate state is due to the layout of the cooling system and the handling of the interaction with it. If not all cooling loops are operated, it is nevertheless desirable to let this affect only the minimal possible part of the detector. This is why the software interlock and veto which is triggered by the state of the cooling loops handles modules and opto-boards separately. In turn, the resulting state should be defined, to allow for the requirement of recovery commands. If only the detector loop fails, the modules are off, while the opto-board is still on, which corresponds to the same OPTO_ON state which is also passed during the powering of the detector.

The characteristics of this OPTO_ON state are the fact that the opto-board is fully powered, and from the off-detector side a certain amount of light is arriving at the opto-boards, while the modules are still off.

If only an opto-loop fails while the detector is powered, the modules will stay on, and only the opto-board will be switched off. The resulting state, with no power consumption at the opto-board, is called NO_OPTO. An interlock due to the opening of DAQ rack doors, or the preemptive software command to switch off the corresponding opto-boards, will result in the same state.

The same issue that requires the opto-boards to be powered before the modules, is also relevant during the Inlink-Scan, i.e. the measurement of the PiN current per off detector laser channel. If the detector is powered, it must be prevented that the opto-board sends commands to the modules during the scan while the lasers are off. This can be achieved by not powering the DORIC while the lasers are switched off. As VVDC is what determines the power consumption of the opto-board, this state can also be summarized as NO_OPTO power.

A situation that does not fit into the operation scheme will result in the state UNDEFINED. A readout unit can be in an UNDEFINED state due to an UNDEFINED state of any device unit below. Also a combination of defined states that is not foreseen to happen during normal operation, e. g. an undesired combination of ON and OFF states of the children, results in an UNDEFINED state of the parent. Also a readout unit that has only part of the modules powered, is an UNDEFINED state. This can be the result of a software interlock, which is "unforeseen" only in the sense that it should be dealt with before resuming normal operation.

DCS related States

A last category of states is determined by the state of the detector control system itself, more than the operational state of the detector.

It is critically important that the case where the state calculation is not reliably possible, can be recognized by the shifter. The cause can be the loss of communication to one or more devices, which might be caused purely on the software side in that part or all of a project is not running[3)], it might be due to network or CAN issues, or finally it might be due to defect hardware. The UNKNOWN state signifies this loss of DCS monitoring capabilities, which will in general also mean the loss of DCS control.

To give the operator immediate feedback after the sending of a command, and to guarantee that each command is executed without interference, all commands will activate the TRANSITION state. In this state most commands will not propagate. The previous command should be executed properly, and the well-defined target state reached, before any other command will be executed reliably.

Summary of the Readout Unit States

Pixel specific states are needed supplementary to the ATLAS states SHUTDOWN, READY, and STANDBY. Regarding the ATLAS states, to note is the fact that pixel READY requires a DAQ action, and can not actively be reached by the detector control system. The pixel readout unit state ON is a "ready" in the sense that no additional DCS action is required or even possible, to make the system ready for data taking. STARTED is equivalent to SHUTDOWN in the sense that no power arrives at the detector, however, PP2, which is inside the cavern, is powered and

[3)] The case of an internal failure inside the FSM projects is monitored separately

its state synchronized with the DCS software. PART_OFF is a state in between, with no power at the detector, but no guaranteed synchronization.

In addition to the ATLAS state STANDBY, there is the corresponding situation where the high voltage is off, but the modules are un-configured or even have the pre-amplifiers still on. Also the transition from LV_ON to STANDBY can only be effected by the DAQ.

During regular operation, transitions will only occur between the states READY, ON (transient), and STANDBY. This also means that the power consumption and temperature of the detector are kept stable.

Commands

To handle the transitions between the defined states, a set of required commands has to be defined and implemented, which has been accomplished within the scope of this thesis. After the states have been established and motivated, the commands are the second cornerstone of the operation model. The implementation ranges from simple powering of individual channels to complex procedures to handle the hardware in a well defined sequential and safe way.

The most elaborate and time consuming command is the *STARTUP* command, which brings the detector after a shutdown period back to the STARTED state. It is based on the startup procedure as proposed by M. Citterio [99]. It is necessary to synchronize hardware and software state, as it is not be guaranteed that PP2 channels remember their last setting after being power cycled.
The functionality of the regulators relies on the sense lines which sense the actual voltage at the load. In case of a broken sense line, the measurement for the voltage will behave in a different way for the different scenarios. In Figure 6.6 a measurement done by M. Leyton [99] shows this behaviour for all lines OK, high sense open, low sense open, or both open.

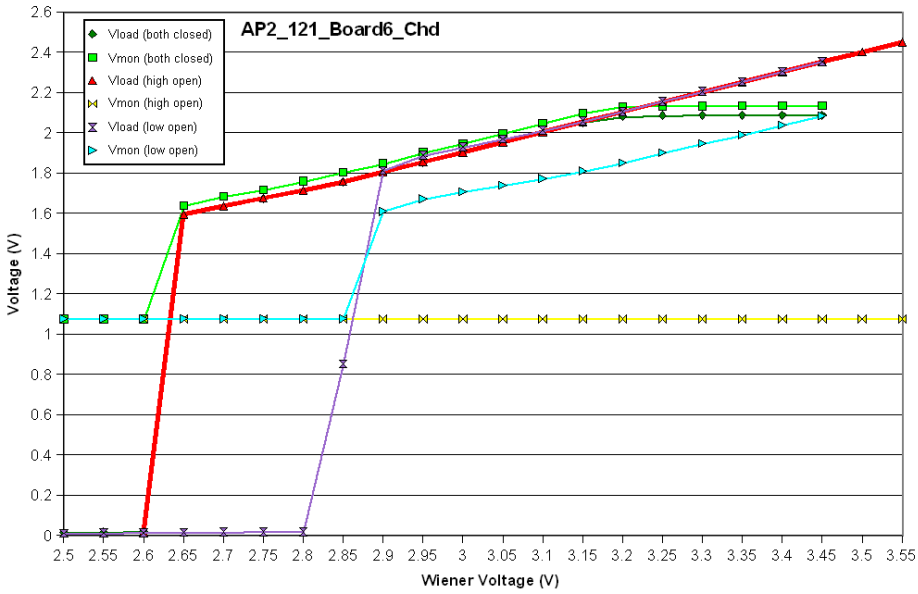

Figure 6.6: *Behaviour of voltage as seen by the module and as reported by PP2 in case of open/closed sense lines.*

It can be seen that there is no good chance to detect a broken low sense due to the small window between Vmon for a good channel and a bad channel. However a broken high sense can be recognized by the fact that Vmon is measured as zero over a wide range of Wiener voltage, even when in reality a voltage is applied to the load. This

can already be recognized at a lower voltage level, where the voltage seen by the module is still low enough to be safe. If directly the full Wiener voltage would be applied, the module would see severe over-voltage, which can be avoided if a failure is detected in time.

From the STARTED state, the command *SWITCH_OFF* will switch off the primary power supplies, and put the readout unit back into the OFF state. For powering the detector, the opto-boards must be switched on first. This is done by the *SWITCH_ON_OPTO* command. It will switch on VISet and VPin, and last put VVDC at the SC-OLink to the nominal voltage and uninhibit the channel at the regulator. This order makes it possible to use this command also while the modules are already powered.

The *SWITCH_ON_LV* command is only available from OPTO_ON. This avoids the above-mentioned issue of wrong and arbitrary signals being sent to the modules. It uninhibits the regulator channels for the module low voltage supply.

High voltage will be switched on last, and only depending on the beam conditions. The command SWITCH_ON_HV acknowledges any remaining interlock at the iseg crates, and switches on the channel belonging to the readout unit. Correspondingly, SWITCH_OFF_HV will switch off this high voltage channel.

BACK_TO_OPTO_ON and *BACK_TO_STARTED* are general commands, that bring the readout unit from any state to OPTO_ON or STARTED. As such, they revert the command order for powering: switch off the high voltage, switch off the modules, and in case of *BACK_TO_STARTED* switch off the opto-board. For improved robustness, but with loss in speed, every internal step will always be executed from any initial state, without check for the actual powering state of the component. In case of a shutdown of the complete detector, these commands can be executed in the time span of about two minutes, which is to be compared with the delay of 5 minutes before the hardware interlock is triggered.

SWITCH_OFF_VVDC is foreseen to support the Inlink scan by the DAQ with a powered detector. It will leave the SC-OLink channels powering the opto-board on, and only inhibits the regulator channel for VVDC. In this way it is prevented that any signals are sent to the modules while the laser channels are switched off successively.

Only a very limited set of commands is available from UNDEFINED, and most of them will not permit any additional power to the detector before the situation is resolved by an expert. Currently the only exception is the *RECOVER_MODULES* command, which switches on modules that have been software interlocked. Massive amounts of modules being switched off due to over-currents can for example be caused by an irregular opening of a DAQ rack door, or power loss of an entire DAQ rack, giving rise to the need of recovery from a higher level.

6.3.2 Device Unit States and Commands

The readout units consist of six or seven module device units, one device unit for the opto-board, two device units for the two Wiener channels (for Vdd and Vdda), and currently one device unit for the high voltage channel. In addition, there is the CMD device unit which handles the synchronization of readout group commands. Figure 6.7 shows the device units of the readout unit and the parameters which are monitored in DCS and by which they can be characterized.

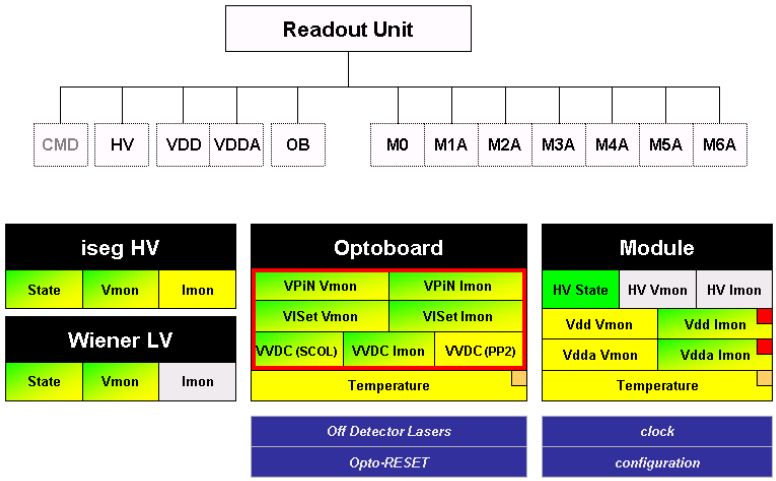

Figure 6.7: *Devices of the readout unit and parameters measured in DCS. Green color signifies dependence of the state, yellow color dependence of the status on the value. Red and orange means the value is the input of a software interlock, triggered by* FATAL *or* ERROR *status. The modules and opto-boards also depend on values not under DCS control or not directly measured by DCS. This is shown by dark blue color. Grey colored values do not influence directly state and status calculation.*

Module Device Unit States

For the modules, the monitored parameters are the supply voltages of the digital and analog circuits provided by the voltage regulators at PP2, and the corresponding currents measured by PP4. The depletion voltage of the sensors and the leakage current is only measured at the iseg, i.e. currently per readout unit. For each module the temperature is monitored. By these values state and status can be determined. From the DCS side what determines the operational state of the module is, whether the supply voltages are provided or not. However, some parameters controlled by the DAQ also have an influence on those values measured in DCS.

The different module states that can generally be identified by the typical current consumption are the following:

- The module is off.
- The module is un-configured.
- The module is configured for data taking.
- The module is configured for data taking, but the pre-amplifiers are killed.
- The module receives only a 20MHz clock.
- The module receives no clock.
- The module is noisy.

A summary of all implemented module states is given in table 6.1.

The considerations and requirements that lead to the choice and definition of the above states will be described in the following.

The module states are what drives the powered readout unit states, which were discussed in paragraph 6.3.1, and from an operational point of view the distinct states are mostly congruent to the general ATLAS states READY, STANDBY, SHUTDOWN (OFF) and NOT_READY. The main conceptual difference is the definition of the READY state, i.e. the decision - due to the almost inextricable relation between DCS and DAQ and the influence of DAQ actions on DCS states in general - to take into account also parameters which are not under DCS control.

While in ATLAS DCS the state READY is commonly defined as ready relating to DCS alone, there are various

State	Description
OFF	Neither low voltage nor high voltage
DISABLED	The module is off, but the state will be ignored for the calculation of the readout unit state. Any non-zero current consumption will trigger a further switch off command.
LOCKED_OUT	Like DISABLED, but different privileges needed to enter or leave this state.
ON	Low voltage and high voltage is applied, module is completely powered from DCS side, but further action is necessary to make the module ready for data taking, i.e. the module needs to be configured by the DAQ, or the module has lost the clock.
READY	Module completely powered and configured for data taking. (Noisy modules will be marked by status only)
STANDBY	Module powered and configured, but not depleted and pre-amplifiers are killed
LV_ON	The module is powered without high voltage, but is not configured or the pre-amplifiers are not killed. Any operational problem like lost clock will also prevent the module from reaching the STANDBY state.
UNDEFINED	Currents not inside normal operation ranges, or inconsistent HV/LV state.
UNKNOWN	No reliable low voltage current measurements
TRANSITION	Command is being executed

Table 6.1: *Module states.*

reasons why this is not the best solution for the Pixel Detector. The first argument to consider DAQ dominated effects is the fact that the DCS states can provide a fast feedback about the success or failure of a DAQ action and the overall configuration state of the detector on a module level granularity, and give hints still to the front-end level. What can not be provided, is diagnostics on the (multiple) pixel level, due to the fact that masked pixels are still powered the same way as unmasked pixel cells.

Another argument to consider DAQ effects, is the fact that it is not always possible to completely disentangle DAQ and DCS responsibility. One example is the clock for the modules, which is generated by the DAQ, and the lasers transmitting it are controlled by the DAQ, but if the clock is lost, the reset signal for the DORIC to lock into the correct frequency is controlled by the DCS. In addition to the effects for data taking, a module that does not receive the correct clock will only have about one third of the power consumption, which affects the evaporative cooling system. Another case is the module configuration, which can only be provided by the DAQ, but if the modules are intentionally or unintentionally power cycled they will be un-configured even without any action by the DAQ. A noisy module will generate heat due to an increased current consumption.

Furthermore, the configuration state of the modules is defined as safety aspect with respect to beam issues, as the pre-amplifiers need to be killed during unstable beam conditions.

These arguments outweigh the aspects which are in favor of a module state purely calculated by supply voltages. The main rationale for the pure DCS approach would be a self-consistency of the two systems. If a module would be in the state READY as soon as the supply voltages are fully provided, then no properties which can not be influenced by DCS go into the state calculation, and all states could be reached by DCS actions. In this way, an incorrect DCS state would have its reason in the detector control system, making it easier for the shifter to identify the failure source. This holds especially true once there is no dedicated pixel shifter anymore, who is not familiar with the unique definition of the READY state for the Pixel Detector.

However, operational experience has proved the value of the current definition of READY as “ready for data taking”. Especially the fact that killed pre-amplifiers are a safety aspect makes it mandatory to monitor at least some configuration aspects inside DCS. Taking into account that STANDBY is also required by ATLAS, the fact that also READY contains DAQ information is only an additional step, and not a completely new aspect.

To mention is the fact that a noisy module will be recognizable by the status, while being in state READY. Incorrect clock as well as no configuration or configuration without pre-amplifiers is summarized by the state ON. In the states without high voltage, only a correctly configured module will be in STANDBY. If the currents are not

consistent with the defined states, the module is in an UNDEFINED state.

Like for the readout unit, also for the module level there are the states UNKNOWN and TRANSITION relating to the detector control system instead of operational conditions.

Though it is generally foreseen that shifter operation does not take place on the device unit level, there are some exceptions, as well as expert commands. Therefore the TRANSITION state for the module device unit indicates a command issued on this level.

Since the module states are mainly driven by the current consumption, and the current consumption is the most critical issue to decide whether a module is powered or not and therefore in a safe state, the communication to the PP4 determines the UNKNOWN state of the module.

Further, there are states that describe an operational condition - the module is off - but with additional features. The first is that this module is not taken into account for the readout unit state, while the status is still propagated, and the second point is that if the module is switched on by any means outside the FSM, a renewed switch off command will keep the module in the desired un-powered state.

The necessity for this states arises due to the fact that a readout unit with powered and un-powered modules has an UNDEFINED state. While the FSM framework provides the possibility to disable a device unit, this will completely ignore also all future failure conditions, while the pixel specific states ignore only the propagation of the state, and provide some kind of measures to guarantee that the state stays as it is desired.

The difference between DISABLED and LOCKED_OUT is only the intention and the needed privileges to execute the command, while the implementation and effects are identical. The DISABLED state is intended in the case that a module is temporarily not supposed to be switched on, either due to a hardware failure or a readout issue, and the command can be issued by the shifter and by the DAQ via DDC. LOCKED_OUT guarantees that the module can be re-included only by experts, if it is considered unsafe to operate the module at all.

In case of modules which do not behave conform to normal modules, but where it has been verified that they provide good data during running, it can be decided to operate them anyway, if it is deemed safe. In that case, it has to be ensured that the detector state will not be influenced by the UNDEFINED state of these modules. The solution is to provide a non-standard matrix for the state calculation of these modules. Due to the fact that it is easy to introduce unwanted effects into the running system, there are various sanity checks in place to avoid the unintentional undermining of the software interlocks and the definition of the OFF state as an un-powered state.

Besides the operational state, there are some values which are relevant only for the safety, and therefore handled in the status. Too high module temperature can cause WARNING or ERROR, where WARNING is normally equivalent to the prohibition to power the module, while ERROR will trigger an automatic switching off. This means that the shooting point of the software interlock is always equivalent to a temperature ERROR, while the veto temperature is only by custom equivalent to the WARNING temperature. Too low temperature is an indication that the module NTC is disconnected, and will also cause a WARNING. Analogous to the hardware interlock, which has to be disabled, a low reading can also be ignored for status calculation in software. A second issue is the supply voltage, which is another input to the status in case it differs to much from the nominal value. The software interlock due to over-currents will be triggered by FATAL status. The different handling of currents and temperatures is due to the fact that the hardware interlock is only triggering on temperatures, which means that by definition no temperatures above the hardware interlock shooting point should ever occur.

Due to the complexity of the input values, the modules status is generally not derived from PVSS alerts, but calculated in the same script as the states. To keep the desired consistency of PVSS alert handling and FSM status, and on the same time avoid redundant alerts, the standard way for pixel device units is to set the alert handling *only* on the FSM status.

By definition, UNKNOWN state goes always with ERROR status. Bad status values during transitions are suppressed for a limited time, to account for the fact that changes will not be read out at exactly the same moment for all

channels, and intermediate undefined situations are certain to be encountered.

Opto-board Device Unit States

For the opto-boards, the monitored parameters are the three voltages supplied by the SC-OLink and the corresponding currents for depleting the PiN diode on the opto-board (VPin), for setting the working point of the on detector lasers (VISet), and for supplying the VDC and the DORIC (VVDC). The regulated VVDC voltage is controlled and measured at PP2. Each opto-board has an NTC for temperature monitoring. The OPTO-RESET signal is also controlled by DCS, but no opto-board related value inside DCS is monitoring its effects. On the other hand, the state of the DAQ controlled off detector lasers can be seen by their effect on the PiN current measured by the SC-OLink.

For the opto-board, there are the following distinct conditions which warrant a dedicated state:

- The opto-board is off.
- The opto-board is off, but the corresponding regulator channel is powered (and inhibited).
- The opto-board is fully powered.
- The opto-board is fully powered, except for the corresponding regulator channel which is inhibited.[4)]
- The opto-board is fully powered, but the off-detector lasers are off.
- The opto-board is only partially powered.

A summary of all implemented opto-board states is given in table 6.2.

Table 6.2 shows the resulting state definitions.

State	Description
OFF	All channels un-powered.
STARTED	VVDC at 5V, regulator channel inhibited.
READY	Opto-board completely powered, and IPin above threshold.
NO_VVDC	All SC-OLink channels powered, but regulator channel for VVDC inhibited.
ON	All channels powered and uninhibited, but off-detector laser current too low.
UNDEFINED	Voltages and/or currents not inside nominal operation ranges.
TRANSITION	Command is being executed.
UNKNOWN	No communication to SC-OLink.

Table 6.2: *Opto-board States.*

For the opto-board, the same decision to include DAQ controlled parameters into the DCS state calculation was taken, as these parameters are relevant for the operational procedures concerning DCS. In case of the opto-board, this factor includes only the off-detector lasers. As described above, if the lasers are off while the modules are powered, it is possible that random noise will be sent as commands to the modules and cause over-currents. Therefore, the state READY requires that all supply voltages are inside the nominal range, and the PiN current exceeds at least a minimal value. If the off-detector lasers are off, the opto-board is in the ON state, in analogy to the modules: all supply voltages are provided, but a DAQ action is necessary to become ready for data taking. The DCS however does *not* keep track of the correct current, which is different for each opto-board, and depending among others on issues like alignment of the TX-Plugin in the BOC, and the laser settings, and in such might be changing.

For the un-powered states, STARTED and OFF are the same state for the opto-board itself. The difference between the states is the fact that in the STARTED state the regulator channel input for VVDC is powered after the *STARTUP* procedure for the readout unit.

[4)] That implies: No commands to the modules, no data to the counting room, almost zero power consumption.

For the Inlink Scan, which verifies that the individual laser channels of one off-detector VCSEL array are functional, it is necessary to only switch on one laser at a time. If this scan should be run with modules powered, an additional state is required. As long as VVDC is off, the DORIC can not send commands to the modules. This makes it possible to run the modules while the opto-board does not receive any clock and command data. The corresponding opto-board state is called NO_VVDC. All other conditions, where not all supply voltages are in the nominal range, are summarized as UNDEFINED. This is also true for the condition where only VVDC is powered, and in case of a functional opto reset line, a voltage can be seen on the VPin channel. While this was used during the connectivity test, it will not normally play a role during regular operation and therefore does not warrant a dedicated state.

Regarding DCS related states, as there are commands defined on the device unit level, also for the opto-board the TRANSITION state exists to signal an ongoing command. For the opto-board Device Unit, the UNKNOWN state is determined only by the SC-OLink channels.

Like for the modules, the temperature is what determines the status of the opto-board. Again, there are two threshold for high temperatures, one for a disconnected NTC, but in addition, for the opto-board also low temperature is relevant: At too low temperatures, the laser output power can decrease enough that the opto-boards would not generally be ready for data taking.
The monitoring of VVDC at the regulator is important for the opto-boards, as the tolerance range for the voltage is smaller. A deviation of 100 mV from the nominal value will generate an ERROR. Over-currents or over-voltages outside the defined value matrix cause a FATAL status, and trigger a software interlock.

High Voltage and Low Voltage Device Unit States

While modules and opto-boards are complex device units depending on many inputs, for the low voltage and high voltage device unit the state depends only on voltage, current, and state of a single hardware channel. The states for the Wiener Low Voltage and iseg High Voltage device units are summarized in table 6.3.

State	Wiener	iseg
OFF	Channel is off.	
ON	Channel is at nominal voltage	
UNKNOWN	No communication to the hardware.	
UNDEFINED	Values not inside nominal operation ranges or inconsistent.	
STARTED	Channel is on and at reduced voltage setting	n.a.
RAMPING	n.a.	Channel is on and ramping up or down.

Table 6.3: *Wiener and Iseg Device Unit States.*

Besides the standard states, the Low Voltage Device Unit has a STARTED state, where a reduced voltage is supplied to the regulator. For both device types, the internal interlock information is handled by the status, if it is consistent with the power status, otherwise by an UNDEFINED state. Due to the internal handling of interlock information, an ERROR status of the iseg DU indicates that an interlock was fired while the detector was powered, alerting to a situation which is not consistent to normal operation procedure. This is primarily true for temperature interlock without the corresponding part of the detector being switched off by software in time, but also a lost NTC, or a beam dump without handshake and therefore without regular ramp down will trigger this reaction.

There is no TRANSITION state, as no operator commands are defined on the level of this device units. However, taking into account that at nominal depletion voltage for an irradiated detector of up to 600 V even with a ramping speed of 10% the required ramping time lies in the range of the state calculation interval, the dedicated RAMPING state is introduced for the iseg device unit.

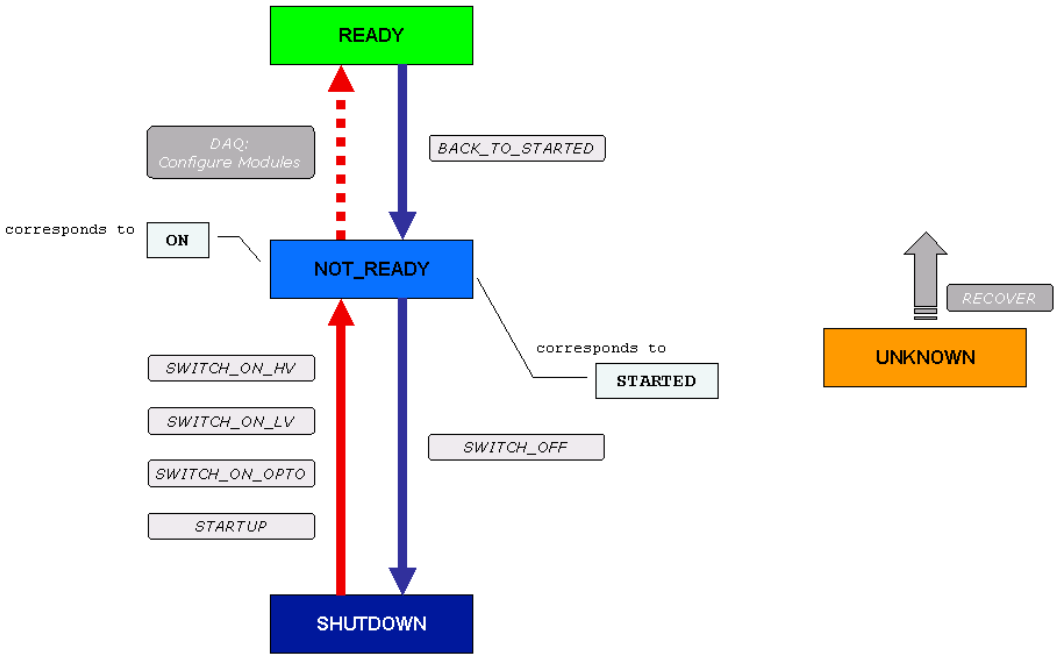

Figure 6.8: *ATLAS state diagram projected on Pixel operation model (Calibration without beam).*

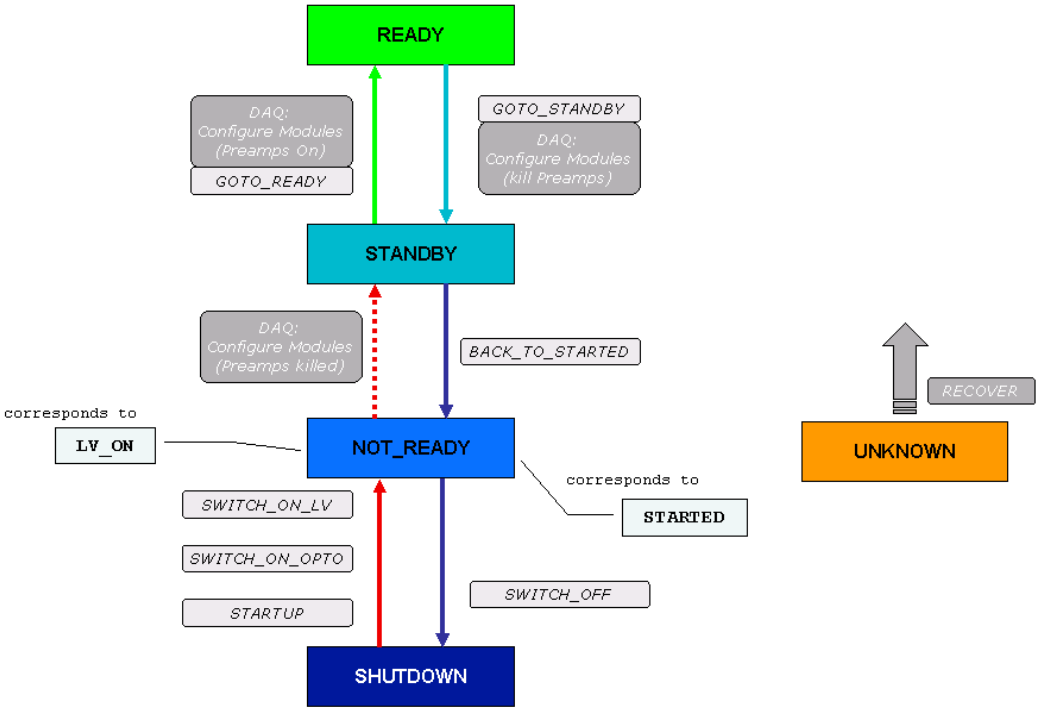

Figure 6.9: *ATLAS state diagram projected on pixel operation model (Data Taking with beam).*

The COMMAND Device Unit states

The COMMAND device unit finally has only two states, IDLE and ACTIVE, to account for the fact if a command is being executed, or if the readout unit is ready to receive another command.

Device Unit Commands

On the device unit level, commands should be generally expert command, with a few exceptions. One exception is the disabling of modules.

6.3.3 State Propagation

On the higher level, it has to be determined how to summarize and propagate the lower level states upwards. For summarizing, most mixed states will be propagated as PART_ON, due to the fact that the states are chosen to be distinctive, and therefore for a common state on the higher level generally the common ground between any two states is missing. An exception is ON and READY which will propagate as ON, due to the definition of ON as "completely powered in DCS", which applies to both states. The downside to a common propagation is that it can not be easily determined if the whole detector is in a state with HV on, but pre-amplifiers already off. To provide this possibility, however a separate state is already required on the module level. Correspondingly LV_ON and STANDBY propagate as LV_ON up to the layer level. Any Mixture of OFF, STARTED, and PART_OFF, propagates as PART_OFF, so that up to the layer level it can be easily seen if the detector is in a safe state. UNKNOWN, TRANSITION, and UNDEFINED will in that order always dominate the state of the parent object.

For the priority, UNKNOWN has the highest priority, but will therefore overwrite anything else, until the issue is solved. Commands might entail unforeseeable results if only part of the hardware reacts. This is alleviated due to the required intermediate states, but should be avoided nevertheless. While a command is ongoing, no other command should be sent, however it may not be unambiguous if one TRANSITION state should be propagated already with very high priority.

The translation of pixel specific states into ATLAS conform states has to happen at the partition level, where DDC is established, in order to allow for a centralized handling of this. This means that all pixel specific states will be summarized as NOT_READY. Figure 6.8 and 6.9 show the state diagram and the typical command sequence if it is reduced to ATLAS states.

6.4 Discussion of the chosen Implementation for the Pixel FSM

The current implementation of the Pixel FSM has proved to be a reliable tool well suited to the operation of the Pixel Detector. Nevertheless, from the start there were items where conflicting considerations had to be weighed against each other, and where the operational experience of many months might give weight to other aspects then what was originally envisaged as the main issues. This paragraph contains a discussion of the implemented model and evaluates the taken decisions in the light of gained experience.

6.4.1 Alternative/Additional Implementation Options

6.4.1.1 The Tree Structure

For the structure of the tree, the sub-detector node and the DAQ partition nodes beneath are defined by ATLAS DCS. For the node below, it was decided to go with the mechanical layers, which combine the lower nodes to a manageable number and allow for a clear geographical layout in the graphical representation. It was considered

necessary to have a dedicated node for the cooling loops, due to the pivotal role the cooling system plays for pixel operation. This node has proved useful during the commissioning of the cooling system with its loop-by-loop operation and gives a natural place to check the power consumption of the pixel modules per cooling loop. The readout group is a required node, the necessity of this node has been argued in [98] and has been proved in operation. For a comprehensible state calculation, the readout unit is not the lowest level, but the modules, opto-boards and power supply channels are separate and independent device units.

This hybrid solution does not fit uniquely to any special use case, but is due to its generality rather universally applicable, and also the clarity for the shifter must not be underestimated. Nevertheless, there are different options which could have been chosen for the tree structure.

One is given by the fact that the current solution is not very well suited for a joint operation of cooling system and Pixel Detector, due to the fact that the cooling system follows the structure determined by the distribution racks in 4 quadrants, and not the pixel TTC partitions. This was especially relevant during the commissioning phase, when only part of the cooling system was operable. However, as the primary tree structure, it is not viable to follow the cooling granularity, if one wants to be conform to the ATLAS requirement of a TTC partition node. Even with some working around these restrictions, the cooling system in itself is not self-consistent with respect to the Pixel Detector readout units. In each quadrant there are readout units for which the opto-loop belongs to a different quadrant than the detector loop. Figure 6.10 shows a tree structure that follows the cooling modularity as much as possible.
In reality, to facilitate operation, an auxiliary tree from references to the main tree was built according to the shown structure, but not as the main branch. Finally, the remaining actions depending on the cooling system, were automatized completely with direct access via the readout units.

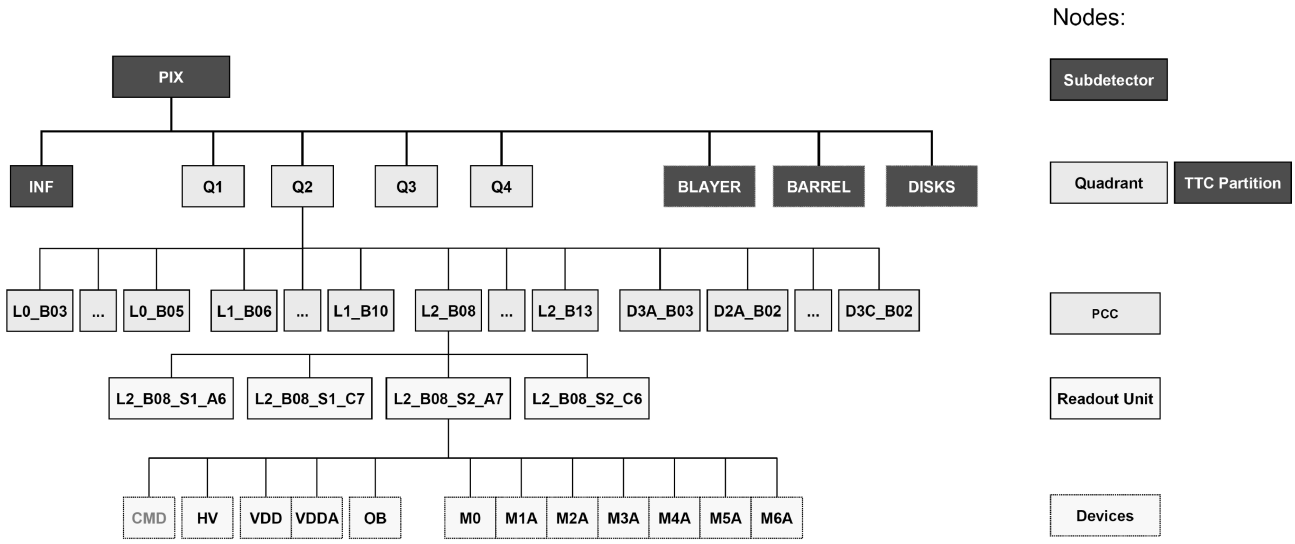

Figure 6.10: *A tree structure for the pixel FSM that follows the modularity of the detector cooling, while trying to remain compliant to ATLAS requirements. The missing direct equivalence to the DAQ partitions is problematic, and would have to be accounted for in a custom workaround. The fact that the opto-board cooling does not fit in all cases into the quadrant scheme by detector cooling loops is an intrinsic drawback of this approach.*

Another possible structure would be to follow almost exclusively the DAQ modularity. This would result in the same number of nodes below the TTC partition nodes as in the mechanical model, i.e. nine nodes for the DAQ crates instead of nine nodes for Layers and Disks. The next level might be the ROD, or directly the readout

unit. The problem with a ROD level node is the fact that this creates either 22 redundant nodes in the B-Layer, or introduces an inconsistent tree structure depending on the geographical position. The obvious advantage of a DAQ-based solution is the possibility to easily see the summary state of each ROD respectively ROD-crate, and the facility to give commands with the corresponding granularity. Especially with frequent interventions on individual DAQ-Racks this allows one to easily prepare the affected readout units and verify the correct state before triggering a hardware interlock. Also the forcible "reset" of the modules by power-cycling would be made easier in case of single unresponsive RODs. In the current system, the commands per DAQ crate are already implemented in the infrastructure branch, the state per DAQ-crate could be provided by another parallel tree, although at the cost of system load. The only feature not easily provided otherwise inside the FSM framework is the command per ROD, as this would exceed the rebuild-by-references solution and require dedicated object types.

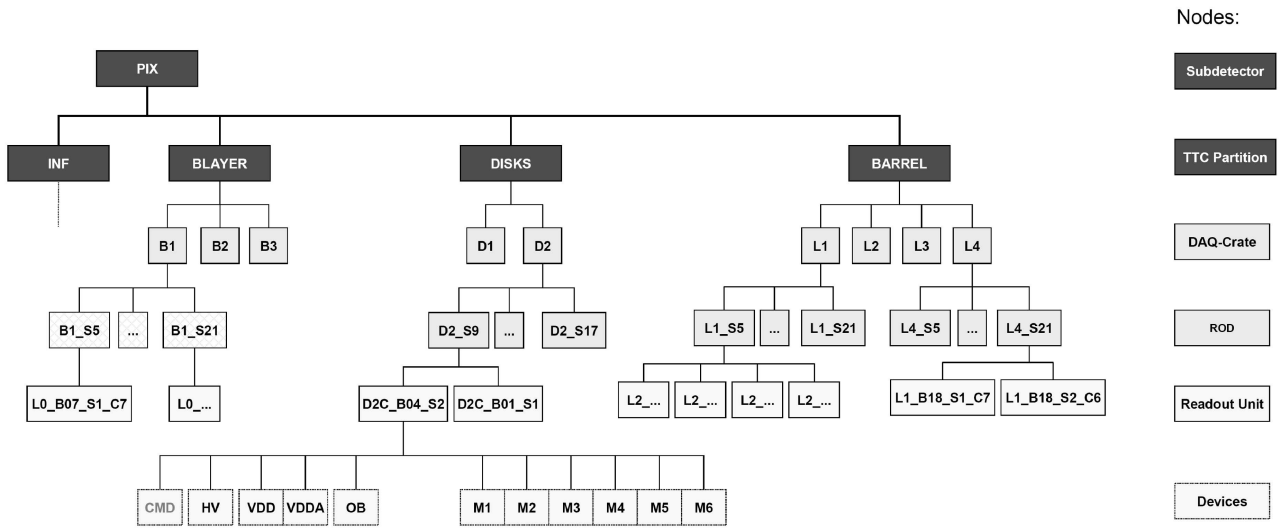

Figure 6.11: *Alternative tree structure following the readout crate modularity. The DAQ partitioning comes for free in this approach, however the cooling loop structure is completely disregarded. The varying number of readout units per ROD leads to redundant nodes in the B-Layer, where there is exactly one ROD per readout unit.*

While a tree structure that follows the cooling granularity does not seem feasible, there is no compulsive structural argument against a tree structure according to the DAQ units. Both discussed alternative tree structures have their origin in operational needs which were not considered critical at the time of the original implementation. However, those needs arise by a limitation of the external system, not of the DCS itself. For the cooling, this was the fact that for a time only part of the system could be operated. For the DAQ the proposed tree structure accommodates the fact that frequent interventions at the DAQ racks are required at the time of the data taking period at beam energy of 3.5 TeV in the second half of 2010. However, the possibility to reset – or configure – modules is clearly a DAQ task. A workaround in DCS to reset modules should be never more than an emergency solution, and might therefore be dismissed as major design criteria for the DCS FSM.

A tree structure by DCS hardware crates for main operation was dismissed before even the first cables were laid. While this may be feasible for other subsystems, it is not feasible for operation of the complex Pixel Detector, with its dependencies between the many different supply types for each readout unit, and well defined operational procedures. It is de facto useful to have this granularity available (see also 6.5), but detector operation can only be handled by detector units.

In summary, the implemented structure works well, and the used non-mandatory nodes have general advantages.

Even though in given situations a more specialized tree structure would often be better suited to the task, this goes at the cost of simplicity and clarity, and the desired feature can often be achieved outside the FSM tree, for the mostly limited time that this need is persisting acutely.

6.4.1.2 State Issues

As described in paragraph 6.3.2, there are more distinctive module conditions than there are states. To keep the system understandable also for non-detector experts, and also for load reasons, the number of states should not become too high. However, it still should be considered if all the use cases are sufficiently covered with the current states.

The most obvious case where a good overview should be provided over the detector state is the question whether the pre-amplifiers are on or off. While the high voltage is off, this is fulfilled for the configured state, where STANDBY is unambiguous. The un-configured state is more dependent on the specific module concerned, and also varying with temperature. Therefore it is not possible to define a global set of limits to uniquely describe this state. On the other hand, the un-configured state normally only occurs during startup of the system. This makes it low priority to identify this condition with a separate state. Another issue is the case where high voltage is on. In case that stable beam conditions are about to end, the desired action is to kill the pre-amplifiers and then switch off the high voltage. Currently there is no easy way to be sure that all modules have reached the desired state, due to the fact that ON comprises multiple conditions on the module level, and also on the readout unit level ON can be a mixture of modules in the states ON and READY. An additional state on either level would simplify the task to decide whether pre-amplifiers are on or off. On the other hand, it is not a very common case that the command to switch off pre-amplifiers works only partially. The exceptions are either complete readout units, respectively RODs, or a single module. In this sense a dedicated state would simplify operation, but in the majority of the time there is no special necessity.

While there are good reasons to have an independent level of device units below the readout unit, as described in [98], there are some cases where an awareness of the global readout unit state would be useful during the calculation of the device unit state. The most frequent case is the desire to suppress an avalanche of alerts during calibration. While an alert should be raised in case a single module loses clock during stable data taking conditions, the same state is the expected one during an Inlink Scan, due to the fact that VVDC will be switched off. Since the readout unit is in the NO_OPTO state in that case, an additional WARNING for the missing clock is redundant. Generally, the suppression of raised alarms during calibration is an issue that has to be addressed in the future.

The display of the interlock condition on the readout unit level is relying on the interlock information delivered by the primary power supplies. Generally, this gives a complementary picture, as the iseg displays latched fired interlocks, and the Wiener actual unlatched ones independent of the power state. However, it is not un-contended for all cases whether a particular alarm should be raised. Therefore, a calculation which is based on all available interlock information for the readout unit could provide a more complete assessment depending on the global situation. However, moving interlock information into the state instead of the status is not feasible due to the many different interlock situations which would overlay and hide real operational issues.

6.4.1.3 Command Issues

Availability/Restriction of commands

The question which commands should be executable from which states and with which privileges has not always been unambiguous, and rather the arguments for each case had to be balanced. This means that with additional or different input from operation the initial reasoning and decisions might have to be reviewed. A controversial issue are the available commands in the UNDEFINED state. The generally pursued line has been to not apply power to an undefined situation, without understanding and resolving the issue before. The underlying philosophy is that an

unpredicted situation should be assessed by an expert, to prevent that harm is being done to the detector. However, in some cases it is the application of power what can resolve the UNDEFINED situation. One example are readout units in the NO_OPTO state. Even with the precautions described in 6.3.1, the probability of modules getting in an UNDEFINED state is only severely reduced, but not eliminated. A global *SWITCH_ON_OPTO* will leave exactly the affected readout units un-recovered. The second use case arises in case that the detector is configured for data taking while the high voltage is off. Part of the modules will get noisy, and the current consumption might increase above the limit for READY and into UNDEFINED. Switching on the high voltage, if possible, would make this situation more and not less stable. This is especially relevant if the DAQ is unavailable for switching off the pre-amplifiers.

Availability of Single Switch On Command

Another issue is the question whether a single switch on command with internal checks for fulfilled preconditions should be made available, or if the single steps should always be visible to the shifter. This question does not arise if detector and cooling system are both started in parallel, due to the discrete nature of this process. However, with a more stable cooling system, it will happen more often in the future that the cooling stays on while the detector is shutdown, as this will extremely limit the thermal stress for the detector. In the long term future this will be necessary to avoid the increase of depletion voltage due to radiation damage. A single command would speed up and facilitate the switching on of the detector.

Availability of Recovery Actions

Currently there is only a very limited number of default recovery commands, and no automatic procedures at all. This is mainly due to the limited number of well-defined "standard failure modes", which can be identified reliably and would allow for a predefined recovery procedure. The most basic recovery command is the OPTO_RESET in case of modules which do not receive the correct clock. An automatic procedure however would be rather complex and involve not only DCS, due to the fact that the readout of the entire ROD might be affected. RECOVER_MODULES is intended to recover the detector from a situation where there was a massive amount of software interlocks, but it is confirmed that the cause for the software interlock has been addressed. This will recover the detector from an UNDEFINED state due to modules without low voltage supply. Also in this case an automatic procedure is not foreseen, as it must be assured that the reason for the software interlocks has been addressed. The implementation of automated recovery procedures in the case that the Pixel Detector leaves the READY state is therefore currently not feasible, and it will in most cases be necessary to consult with a detector expert or intervene on the hardware to find the reason for a given failure.

6.4.2 Conclusion

The current implementation has proved its usefulness for the operation of the Pixel Detector. Even if there is room for improvement in the future, for each change the expected gain has to be carefully weighed against the risk of introducing instabilities into a working system. While there are situations which a more specialized implementation could address better, the pixel FSM has shown its ability to handle diverse situations starting from the system test to the operation of the full detector during calibration and data taking.

6.5 Monitoring of DCS crates in the Infrastructure branch of the FSM

Besides the detector branch of the FSM which handles the TTC partitions, the Infrastructure branch monitors and controls the more low level functional part of the detector control system. Figure 6.12 shows an overview of the Infrastructure node. The environmental conditions, i.e. temperature and humidity, inside the pixel package are monitored passively by the ENV node [92]. The main task of the Interlock node is the monitoring of interlocks in the various parts of the detector, as well as the state of the hardware itself [89]. It has evolved from a almost purely crate based monitoring, into a more functional direction giving the shifter an indication of both cause and effect of a given interlock. The Opto-heater branch focuses on the operation of the opto-heaters more than the monitoring of the hardware, and next to the basic commands of switching on and off, it allows the expert to configure set values and operation mode 5.4.3. The DAQ node monitors VME crates and is responsible for the temperature monitoring of the BOC cards. It offers a basic method to switch on or off detector parts by DAQ crate. The monitoring of the DCS crates will be described in more detail later in this paragraph. Last, the Infrastructure contains references to two external nodes which can affect the operation of the Pixel Detector, the evaporative cooling system and the monitoring of the pixel racks from the Common Infrastructure Controls.

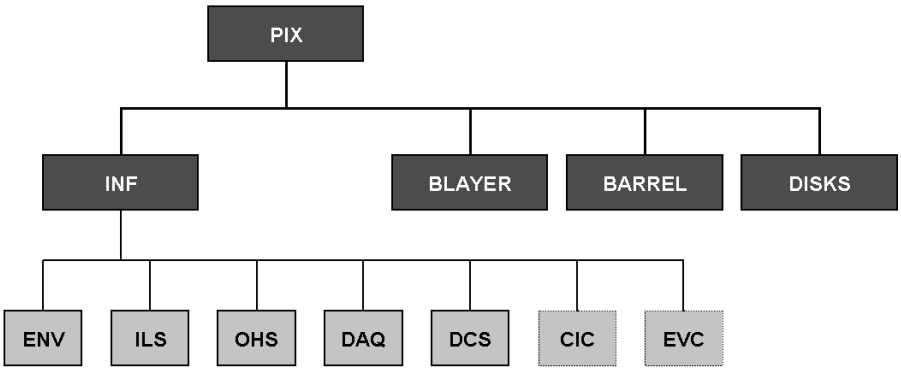

Figure 6.12: *The Infrastructure Branch of the Pixel FSM.*

6.5.1 Requirements

The DCS crates branch of the pixel FSM is responsible for the monitoring of the DCS crate infrastructure on a rather low level very near to the actual FITs, with a modularity determined by the hardware, instead any detector geographical considerations.

Its tasks are:

- The monitoring of the communication state to the individual CAN nodes
- provide controls for the CAN management items outside the low level tools
- temperature monitoring of the hardware, and software interlocks where appropriate
- possibility to control exclusively all detector elements connected to a certain hardware device
- device specific actions

As there are many different types of crates, not all of these items apply in the same degree to all crates.

The monitoring of the communication is utilizing information from the watchdogs, and additionally for all nodes using an ELMB, the provided information about the CAN power supplies can be taken into account. The DCS crates branch is supposed to give access to the control functions of the CAN power supply via the FSM.

Crate temperature is monitored, and where necessary software interlocks are provided. Not only for this, it is necessary to provide the possibility to access all detector elements connected to a certain crate, in order to be able to switch them off in case of over-temperature, or in case of a manual intervention on a single hardware unit.

Finally, for some crates there are actions that have to be performed on a crate or node basis, which do not fit into the detector modularity, or create too much overhead due to the arbitrary grouping of channels. Examples for this are the necessity to acknowledge hardware interlocks for the high voltage on the iseg module level. If this would be done from inside the detector branch, the number of sent commands would be up to 8 times higher then really necessary, or a significant delay in the commands would have to be introduced. The regulator crates have a large set of unique recovery procedures and routines. Due to this fact, this node will be described in more detail in the next paragraphs.

Figure 6.13 shows the DCS crates branch, and in more detail the tree structure for the PP2 crates. Since this is the most intricate part, it was also the one which had to be deployed in the production system first. The interlock crates, DAQ crates and opto-heaters have separate nodes directly in the infrastructure node, due to the fact that the emphasis is on operation instead of crate monitoring, or in case of the DAQ due to the fact that the crate operation is under a different responsibility. For the pure DCS crates, the iseg part is running standalone in the production system, while the dedicated monitoring of all other crate types is still in the development phase, and handled from the local PCs via low level tools.

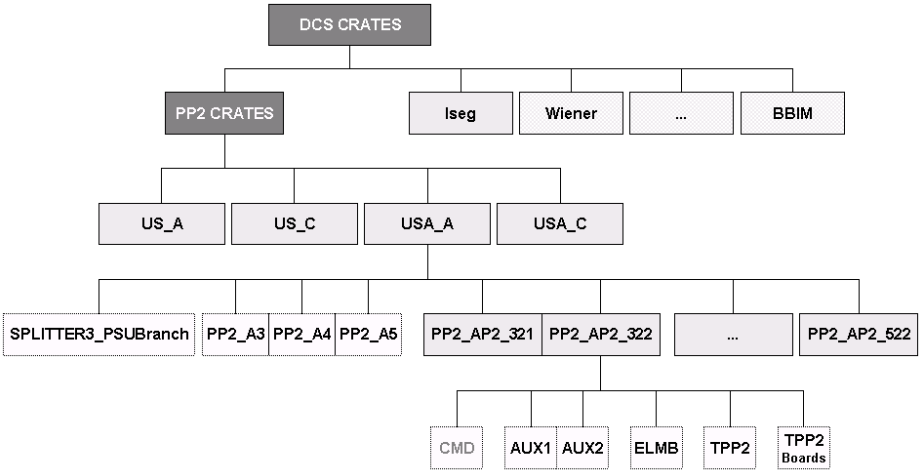

Figure 6.13: *The DCS crates branch. The tree structure of the PP2 crates branch is shown. This is the only part integrated into the production system. The iseg part is running in a standalone project in the production system, while all other nodes are still in the development phase.*

6.5.2 Monitoring and Control of the Regulator Stations in the FSM

The regulator stations installed at PP2 are one of the critical components of the Pixel Detector Control System. A well defined state has to be established at the startup to guarantee the correct and safe operation of the regulators. Therefore, in the course of this thesis, emphasis was put on the integration of the regulator stations into the DCS crates branch of the FSM. The detector branch described in chapter 6.3 covers already common scenarios, so that detector operation by non-experts is possible without resorting to additional resources. In the PP2 branch, access

to additional information is provided, and means are made available for expert operators for coping also with non-standard situations.

The PP2 branch has to deal in particular with the following issues:

- Ensure always correct voltages and channel states by making imperative the recovery procedure when necessary.
- provide alerts and software interlock based on board and crate temperature
- monitor communication status
- translate mapping structure between readout units and PP2 crates to provide a parallel command granularity

The PP2 crates are a special case of DCS crates monitoring due to the elaborate recovery procedures and intricate command handling in addition to standard monitoring. The powering scheme of the CAN bus and the ELMB is different from all other crates. This also implies that communication with the crates is not possible during shutdown, as the auxiliary power supplies necessary for CAN controller and FPGA are shutdown via DSS e.g. in case of a failure of the evaporative cooling plant. In contrast, communication to all other ELMBs – as well as Wiener crates – is guaranteed by UPS. The PP2 crates need to be cooled during operation, so a dedicated software and hardware temperature interlock are in place in case of a failure or loss of the mono-phase cooling system. After a shutdown, dedicated recovery procedures are required.

FSM tree for the PP2 crates

Figure 6.13 shows the hierarchy of the PP2 crates tree.

The second to top level is defined by the fact that there are only 4 CAN PSU branches providing all the required CAN power. The platforms supplied by one CAN PSU branch are grouped by the detector side they are on – A- or C-side – and if they are inside the ring or outside: facing to US-side or USA-side. The location and modularity of the platforms is described in chapter 5.3.4.1. This means that one CAN PSU output supplies e.g. all the platforms facing to US side on side A of the ATLAS detector, and the last one all platforms facing to USA-side on side C. Since this supply is split into 3 individual CAN bus supplies, this node is called the “CANSplitter” control unit.

Each CANSplitter control unit has a corresponding device unit for the actual determination of the power status and for the control of the CAN power, which is by its setup always shared by 3 CAN buses.

To avoid a too deep structure of the tree with only few nodes on each level, there is no dedicated CAN bus level. Instead, a device unit for each of the three CAN buses (in figure 6.13 these are the CAN buses PP2_A3, PP2_A4 and PP2_A5) and the control units for the 7 connected regulator crates (the crates PP2_AP2_321 to PP2_AP2_522) are created equally on the next hierarchy level. The CAN device unit allows for monitoring and setting of the CAN management items (start, stop, and reset of the CAN bus, and the SYNC interval), while the crate control unit (shown for crate PP2_AP2_322) is the central part of the control and monitoring.

The PP2-Crate Control Unit

The crates are the basic unit with respect to the control and monitoring of PP2. The operability is primarily given by the fact that the two auxiliary power channels are supplying the correct voltage. This is monitored by the two “PP2Aux” device units. The communication to the detector control system is handled by the ELMB, which passes commands to the FPGA and sends monitoring values to the PVSS project as long as the so called “monitoring routine” is running. The ELMB state is mainly monitored by the Watchdog. In this way, the “PP2ELMB” device unit contains indirectly also CAN information, as no CAN traffic is possible without CAN power being supplied (modularity of 7 crates) and without the CAN bus being started (modularity of 2 or 3 crates). The temperature, “TPP2”, is monitored by 7 NTCs per crate which are read out through a BBIM, and by 2 NTCs located on each

PP2 board which are read out by the PP2 itself, “TPP2Boards”.
As in the detector branch, there is a dedicated “CMD” device unit to synchronize commands internally. Additionally, this is used to keep track of any situation which might loose the synchronization between hardware and software.
After each shutdown the state of the software and hardware must be synchronized, to avoid undefined behaviour. In addition to the procedure that is already available in the detector branch at *STARTUP*, a similar procedure, i.e. reconfirmation of the inhibit-map and trimmer settings, can also be done while the detector is powered to recover from an intended or unintended power cycle of auxiliary power only.

The state diagram for the PP2 crates is shown in figure 6.14. In the figure also the most relevant combination of device unit states that contribute to a certain crate state is indicated.

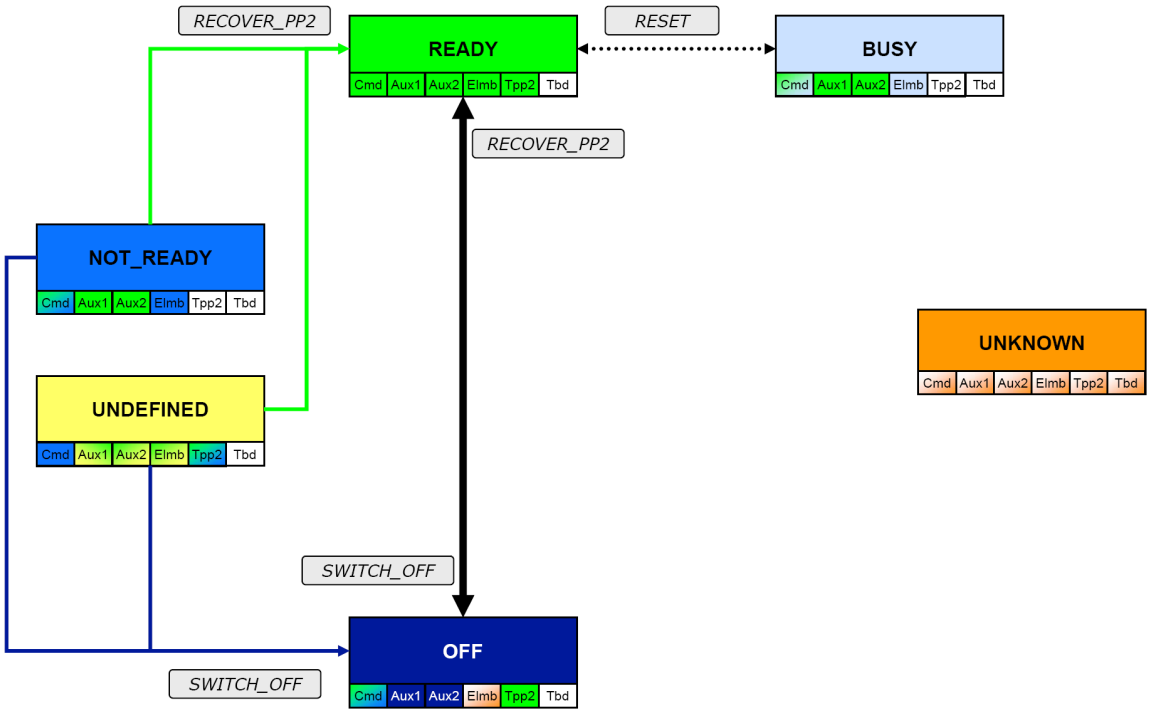

Figure 6.14: *Simplified state diagram for the PP2 crate. Below the state, the most significant combination of contributing device unit states is indicated. Due to the complex nature, not all possible combinations are shown. Only default commands acting exclusively on the PP2 branch are shown.*

In an ideal case, the operation of PP2 only ever goes between OFF and READY, respectively BUSY. In that case, OFF occurs when also the detector is off, e.g. during a Christmas Shutdown, and therefore there is no need for sending commands via PP2 until the system is brought back up. The temperature measured by the BBIM is therefore required to be inside tolerable limits, because the main occurrence of the OFF state is also intended to monitor conditions during long term shutdown periods. During operation of the detector, the PP2 crate should always be READY, or in case of active commands in the detector branch, BUSY. The corresponding procedure to reach READY after any intended or unintended shutdown of auxiliary power is implemented in the *RECOVER_PP2* command. The need for a dedicated recovery procedure arises due to the effect of a power loss for the different supply voltages.

The main aspect for the PP2 branch is how to deal with the loss of the auxiliary power supply. In that case, all

powered channels will become uninhibited at the next clocked command, providing power to modules and opto-boards. In addition, it can not be guaranteed that the original trimmer setting is kept. Therefore, in the case that the primary power supplies are on during loss of controller power, a special inhibit command without clock is provided by ELMB/FPGA, to be able to recover without changing the overall state of the detector. The *RECOVER_PP2* command checks the CAN power, ensures that controller power is supplied, checks that communication to the ELMB is possible, and finally reapplies the inhibit-map using the aforementioned version of the inhibit command without sending clock. The last step is restarting of the monitoring routine.

In case that low voltage or SC-OLink primary power supply is lost, this is handled in the *STARTUP* of the detector branch. Depending on the type of channel and the off-time, in this case the regulator channel could otherwise come back with or without active output. Also, the trimmer setting is kept in most cases, but this can not be guaranteed.

The *RECOVER_PP2* command is also foreseen in case that a controller needs to be power cycled to regain functionality. In this case the auxiliary power would be – shortly – lost while the detector is powered. This procedure would be handled completely by FSM commands inside the PP2 branch of the FSM, and the crate would go via the TRANSITION state to recover READY.

Another issue is the accidental loss of auxiliary power compared to the intentional power-cycling. The recovery procedure is exactly the same, but the software must keep book about this fact. Otherwise the first clocked command will leave all channels uninhibited, regardless of their intended state. Therefore, if any of the auxiliary power channels are detected to be un-powered, the crate will go into the UNDEFINED state. Simple re-powering of the controller power is not sufficient to leave this state. This is only possible by going through the recovery procedure. Temperature problems can also cause the UNDEFINED state, as the implemented interlocks should prevent any over-temperature.

The NOT_READY state is primarily driven by the ELMB. If the ELMB is not responding, even though the auxiliary power channels are on, this leads to the NOT_READY state. Here it depends on the reason why the ELMB is not responding, what action is necessary to recover the system. In some cases an expert intervention is still necessary on a lower level.

If any channel is UNKNOWN, the crate will be generally in an UNKNOWN state. However for some conditions some device units are expected to be unreachable by CAN communication, namely the ELMB device unit and the TPP2Board device unit. Those cases will not contribute to an UNKNOWN state.

Any command initiated from the PP2 branch will result in an TRANSITION state.

Besides the commands acting exclusively on the PP2 branch, the possibility to control detector units by PP2 crate mapping is provided. The use case for this is primarily interventions on single crates. For this purpose the affected crate has to be powerless, i.e. the crate and detector supply voltages have to be switched off. For all other crates the thermo-cycling of the detector should be avoided. The mapping of readout units to PP2 crates is shown in figure 6.15.

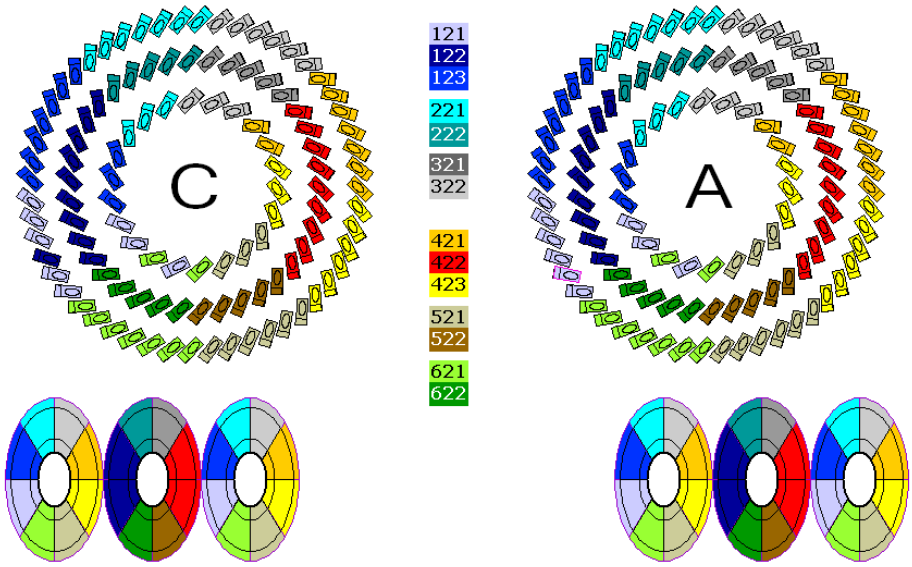

Figure 6.15: *Readout Units per PP2 Crate.*

It can be seen that operating readout units by PP2 crates from the detector branch is not feasible, especially since this mapping is not a priori known to the FSM but has to be accessed indirectly.

6.5.2.1 Conclusion

In summary the main use case of the DCS crates branch in general is the operation of detector units by hardware crate modularity. Without a possibility to access commands in the corresponding modularity from the FSM, any hardware intervention will require time consuming operation to switch off only the corresponding part of the detector, or will result in detector units being unnecessarily thermo-cycled.

While temperature monitoring is useful, this alone would not justify a dedicated branch, as the alarm screen is normally sufficient to point the shifter into the right direction.

Also the handling of the CAN-PSU is in the current state of detector operation not necessarily required to be accessed via the FSM. Problems here happen rarely, and can be handled by experts logging into local machine. This part can be seen as preparation for a future situation, where there might be well-defined recovery procedures that can be done by shifters, who should not be able to access the DCS PCs themselves.

The last point are the device specific actions, where the PP2 crate operation has the most elaborate recovery procedures, going far beyond what is necessary for the other crate types. The inclusion into the FSM has proved extremely useful, as it makes it possible to move operations that were clear expert actions into the range where it is possible to be executed by experienced operators.

Chapter 7

Automatic Procedures

The FSM, described in chapter 6, represents the interface between detector control system and shifters for the daily operation of the Pixel Detector. In case of emerging non-regular or error conditions, the FSM alerts to their presence by the change of state and status. If these issues are safety relevant, it is not feasible to require the shifter to react in a timely and appropriate way. But also complicated regular scenarios can lead to operational errors by shifters unfamiliar with the required procedures. Therefore, the logical and consequential direction to follow in the context of this thesis, was to automatize safety actions and procedures based on the fundament provided by the FSM. In this chapter, these automated scripts invisible to the operator, dealing with safety issues and the introduction of automatic procedures, are presented.

The first part of this chapter will discuss the various risk scenarios where an additional layer of safety must be provided by the DCS software. These aspects always require a fast action to avert risks to the Pixel Detector, and can therefore not rely on the shifters' reaction time. This safeguard functionality is provided by permanently active scripts, continuously running in the background.

The second part will concentrate on making the operational procedures automatic. If complicated shifter actions can be avoided, this will also reduce the probability of operational and user errors. In this way, automatic procedures can also contribute to the safety of the system. Downtime can be avoided or reduced. Depending on the use case, these procedures can be completely automatized by background scripts, or are made available only when required.

7.1 Safety

The implemented safety measures fall into three categories.

The first part are the software interlocks, safety actions deeply integrated into the structure of the FSM. These are based on low level information originating from inside the Pixel DCS projects. The next step is to consider higher level information directly provided by an external PVSS system, and finally, in the last category, information directly provided by the ATLAS Detector Safety System, DSS. In this way, the actions cover the spectrum from highly available information with relatively slow sensitivity to the given danger on the one side, to a fast alerting to failure conditions with lower availability of the information on the other side.

7.1.1 FSM based Software Interlocks

The FSM generates two kinds of software interlocks directly from the state calculation script. Voltage and current of the modules and opto-boards are continuously monitored, and if any value stays in the critical range for more than one reading – to filter out spikes – the corresponding device will be switched off. The temperature based

software interlock is composed of an active and a passive part. The passive part prevents any device from being switched on if its temperature exceeds a VETO temperature. The active software interlock switches the corresponding detector part off, once a device temperature passes the ERROR threshold. Both thresholds as well as a WARNING level are configurable, but there is only one value for all modules and one value for all opto-boards of the detector, so that the temperature spread has to be taken into account when defining the limits. During commissioning, modules were operated for limited periods of time with the software limit near the hardware interlock, while for a cooled down detector the threshold should be as low as possible. However, configuration of the shooting point is a manual action. The switching between cold and warm state is not automated, and also the back pressure set point of the cooling loops is not automatically taken into account for the cold state.

The affected detector part is different for opto-boards and modules. If an opto-board reaches the ERROR temperature, it is switched off, but in case of a high module temperature the action is analog to the hardware interlock i.e. all modules of that readout unit are switched off. This is done to protect modules without NTC against overheating. These modules can not trigger a hardware interlock, as the whole readout unit could not be operated otherwise. The only way an overheating of those modules can be detected is by the fact that the neighbouring modules will also show an increased temperature. If those are switched off separately by software, the whole readout unit will seem to have temperatures inside the accepted range. The risk to damage a module without NTC was considered higher than the desire to avoid power cycling of the other modules in the same readout unit. Furthermore, a single module should only overheat if, for any reason, the thermal contact becomes bad. Normally, all modules of one readout unit are threatened equally by overheating, in case the cooling is not working.
The analogy to the hardware interlock does not extend to the unavailability of temperature information. In this case no action is taken, as the hardware interlock reacts first. Once it is disabled, no software action is desired anymore.

7.1.2 The Cooling Script

The most probable reason for over-temperature of the detector is the failing of the evaporative cooling, described in chapter 3.5. The evaporative cooling system is monitored by a dedicated PVSS project, "ATLIDEEVCOOL" [49]. Using information from this project, it is possible to react on a prospective temperature rise due to a failure in the cooling system before it is actually picked up by the temperature sensors of the Pixel Detector.

The primary objective of the cooling script is to prevent over-temperature. But while the script should react on all cooling failures, on the other hand it should not be the cause of any unnecessary thermal cycling for the detector, i.e it should only react in case of a real failure, act on the smallest sensible detector unit, and bring it to the highest possible safe state. Due to safety relevance of the script, it should be kept as simple as possible, to minimize the danger of software failure. Additionally, it must provide also a passive interlock, i.e. prohibit the re-powering of detector units for which the interlock condition is active.

Finally it must be possible to mask the script, for the cases where it is necessary to temporarily operate parts of the detector without cooling.

7.1.2.1 Requirements

The requirements for the cooling script can be summarized in the following basic points.

- Provide a software interlock that is triggered earlier than the hardware temperature interlock, by acting on information provided by the ID cooling project.
- Prevent the powering of non-cooled detector units.
- Avoid unnecessary thermal cycling.

- Avoid causing downtime for operation.
- Provide a possibility to mask the actions
- Due to the safety relevance, the implementation has to be fail-safe to a high degree.

7.1.2.2 Trigger and Action Matrix

In the first step for the implementation of the cooling script, two issues have to be addressed. It has to be defined on what values a action needs to be triggered, and what form the action takes.

Triggering of an Action

The cooling project already provides a summary of loop and plant state, as well as the analog readings of corresponding sensors. Table 7.1 shows the states of the cooling loops and their definition.

Accordingly, possible triggers for an action are:

- plant state or status
- loop state or status
- analog values

state	meaning
ON	the loop is in operation, liquid flow is supplied the valve on the liquid side is open, and the pressure regulator at nominal value the valve on the gsa side is open, and the back pressure regulator is at nominal
STANDBY	the pipe is at low pressure, there is no flow the valve on the liquid side is closed, the pressure regulator is at minimal setting (“vacuum”) the valve on the gas side is closed, the back pressure regulator is at the nominal setting
OFF	the pipe is isolated the valve on the liquid side is closed, the pressure regulator is at minimal setting (“vacuum”) back pressure regulator is at minimal setting (“vacuum”), the gas is evacuated fast from the line
LOCKED	the pipe is isolated the valve on the liquid side is closed, the pressure regulator is at minimal setting (“vacuum”) the valve on the gas side is closed, the back pressure regulator is at the nominal setting the heater is shut down the loop can not be put in operation without explicit confirmation

Table 7.1: *General Cooling Loop States [100].*

It was decided that an action would have to take place only if a loop is not anymore in the ON state. No action should be taken on the plant state in the framework of this software interlock.

Timing of the Action

For the actual time of the execution of the action, the following question have to be considered:

- Action depending on final state of the cooling loop? (STANDBY vs. OFF)

- Start of action time based or temperature based?

Figure 7.1 shows the temperature evolution for the Pixel modules for the different modes of stopped cooling loops without change in the module power at the example of one loop. Figures 7.2 and 7.3 show the distribution in the behaviour for all loops, for final loop state OFF and STANDBY, with powered modules.

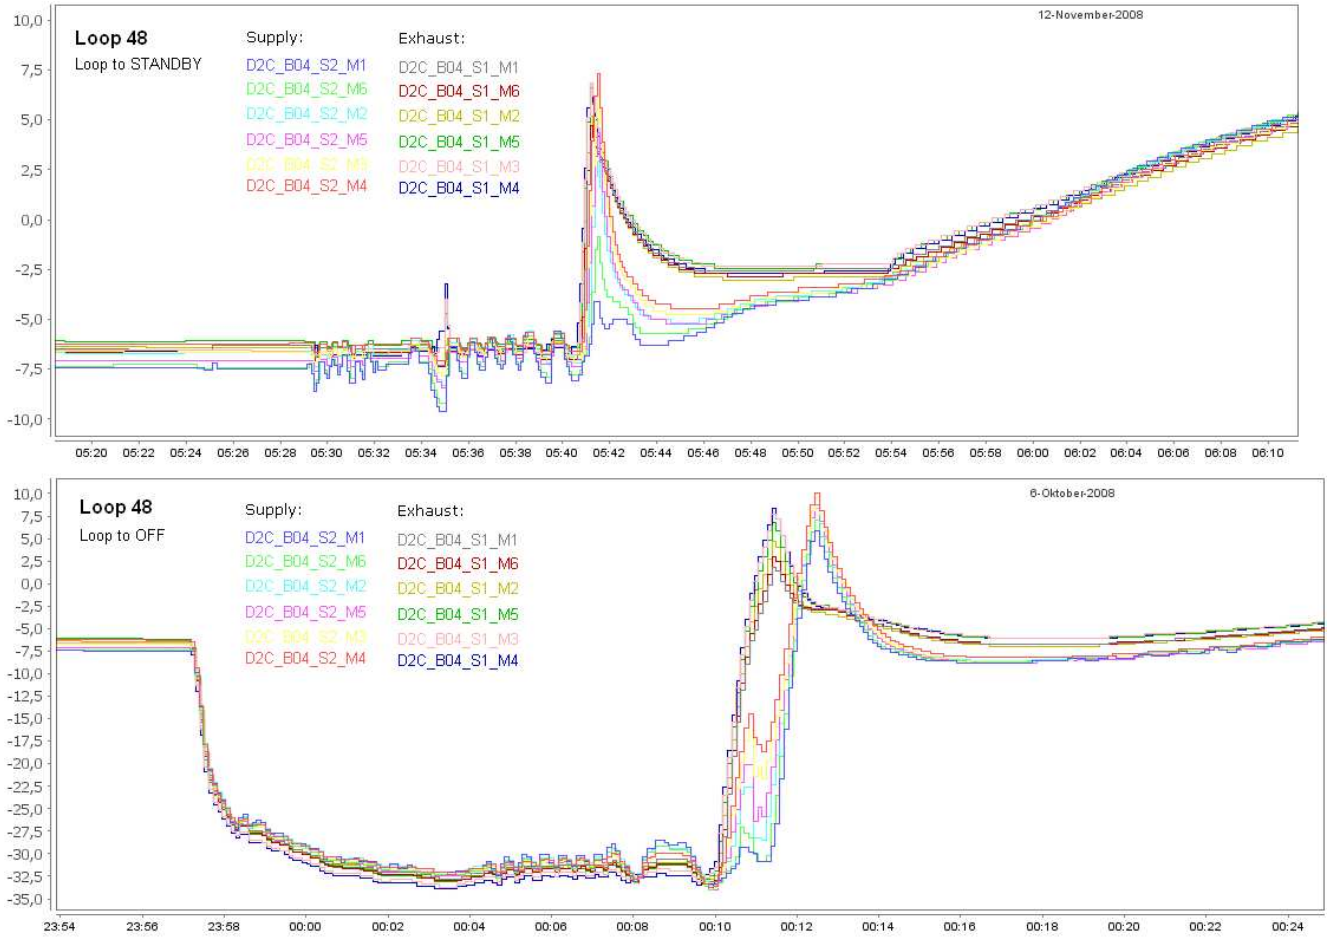

Figure 7.1: *Temperature evolution for powered Pixel modules for the example of loop 48 during a shutdown of the cooling system. The upper plot shows this for loop-state* STANDBY*, the lower for* OFF*. After the loop goes to* STANDBY*, the temperature stays almost constant until the remaining liquid is completely evaporated. At this point the temperature starts to rise fast due to the module power which is not dissipated anymore. The subsequent drop in temperature is due to the fact that the modules are switched off at that point. In contrast, for the loop going to* OFF*, there is an immediate temperature drop caused by the missing back pressure. Also in this case, there is a steep rise in module temperature when the remaining liquid is fully evaporated. Modules in the supply and exhaust sector can be clearly distinguished by their relative temperature evolution.*

From these distributions it is clear that, in case of the loop going to OFF, it is not desired to add an additional temperature gradient to the already steep drop after the failure by switching off the modules immediately. On the other hand, it can also be seen that during the steep rise when all the fluid is evaporated, big temperature gradients in the positive direction can occur, especially if the powering down of the modules is not instantaneous. Currently, the nominal behaviour at a failure is for the loops to go into STANDBY instead of OFF to avoid the fast temperature drop. This being the case, it was decided to switch off the modules immediately after a detected failure. The excursion to low temperatures is present in this case only due to the removal of module power, but not as pronounced. Extremely high temperatures can be avoided completely.

There is however no implemented check for the target loop-state, but instead the action of the cooling script uses the extrinsic knowledge about the programmed reaction of the cooling plant in case of a failure.

Summarizing, for the triggering of the action, this is given by the fact that the loop is not in the state ON anymore. The desired execution time of the action depends on the final loop state. For the loop state OFF, this is after the initial steep temperature drop, and as near as possible to the point when the fluid is evaporated. This is realized by requiring one module of each readout unit to rise above 5°C at an operation back-pressure of 3 bar. However, on request from the Pixel community, the action of the cooling system does not foresee a loop going to OFF anymore. Therefore, the remaining case which is of interest, is when the loop goes to STANDBY. For the timing it was decided to react immediately on the failure, also in light of the fact that the worst temperature gradients happen when switching on the cooling, and not during the shutdown. A fast reaction at the point of the temperature rise was dismissed, due to the requirement of keeping the script simple. A fixed delay does not seem feasible, as there is no way to adjust to unforeseen behaviour, and, as can be seen in figure 7.3, would not be the same for all cooling loops.

After defining the trigger and the point of reaction, what remains to be defined is the action itself. Considering the requirement to avoid unnecessary thermal cycling and minimize downtime for operation, this means the action should switch off the smallest possible detector unit to the highest safe state.
There are different ways of defining the "smallest possible detector unit". For one, this can be defined as the complete cooling loop with all readout units belonging to it. This view is supported by the fact that the operation of any readout unit is only possible if both, modules and opto-board, are powered. However, while this approach has the advantage of simplicity, it has the disadvantage of little flexibility, and thermal changes in detector units where it is not necessarily required.
The chosen granularity is therefore below the readout unit level, treating modules and opto-boards separately. With this granularity, the "highest safe state" is to switch off the opto-board in case of a failure of the opto-loop, and the modules for a detector loop failure.
This offers additional advantages, e.g for the thermal stability for the cooling system: an unstable detector loop is not made more unstable in case an opto loop fails. This was of high concern when the stability of the system was still more fragile. If on the other hand a detector loop fails, all the opto-boards can be kept on, which is of concern for the opto-heater regulation, especially if the regulation uses the opto-heater NTC for feedback. As this NTC is located near to the first opto-board, the regulation by switching card would overcompensate if this opto-board is un-powered. Keeping all detector devices belonging to active cooling loops powered decreases the risk of triggering an interlock on the heater temperature. While the system as a whole is more stable at the current time, during early operation there was the danger of causing a trip of the whole cooling plant by removing the heatload from the loops.

From the operational point of view of the detector, only this high granularity approach allows one to operate part of the detector by quadrant. This concerns the edges of the cooling quadrants, where detector loops and optoloops do not always match (as shown in figure 3.18). Furthermore, this separation is what makes it possible to leave the detector loops off until the optoboards are on and the detector is in OPTO_ON. In this way it is possible to first switch on all optoloops, check that the detector reaches the correct state, proceed to switch on detector loops, and immediately power the corresponding part of the detector.

These considerations can be summarized in the action matrix given in Table 7.2. The only exception to the rule of "highest possible safe state" is the case of a simultaneous failure of optoloop and detector loop. Since this always affects a larger part of the detector and is with high probability symptom for a global failure, the time needed for reaching the final state should not be prolonged artificially. To not counteract the reaction of the DSS script, which will be covered in the next section, in this case the readout unit will go to OFF instead of STARTED. For switching off the loops temporarily, this results in a two step procedure to ensure a uniform state to switch off first the detector loops and subsequently the optoloops.

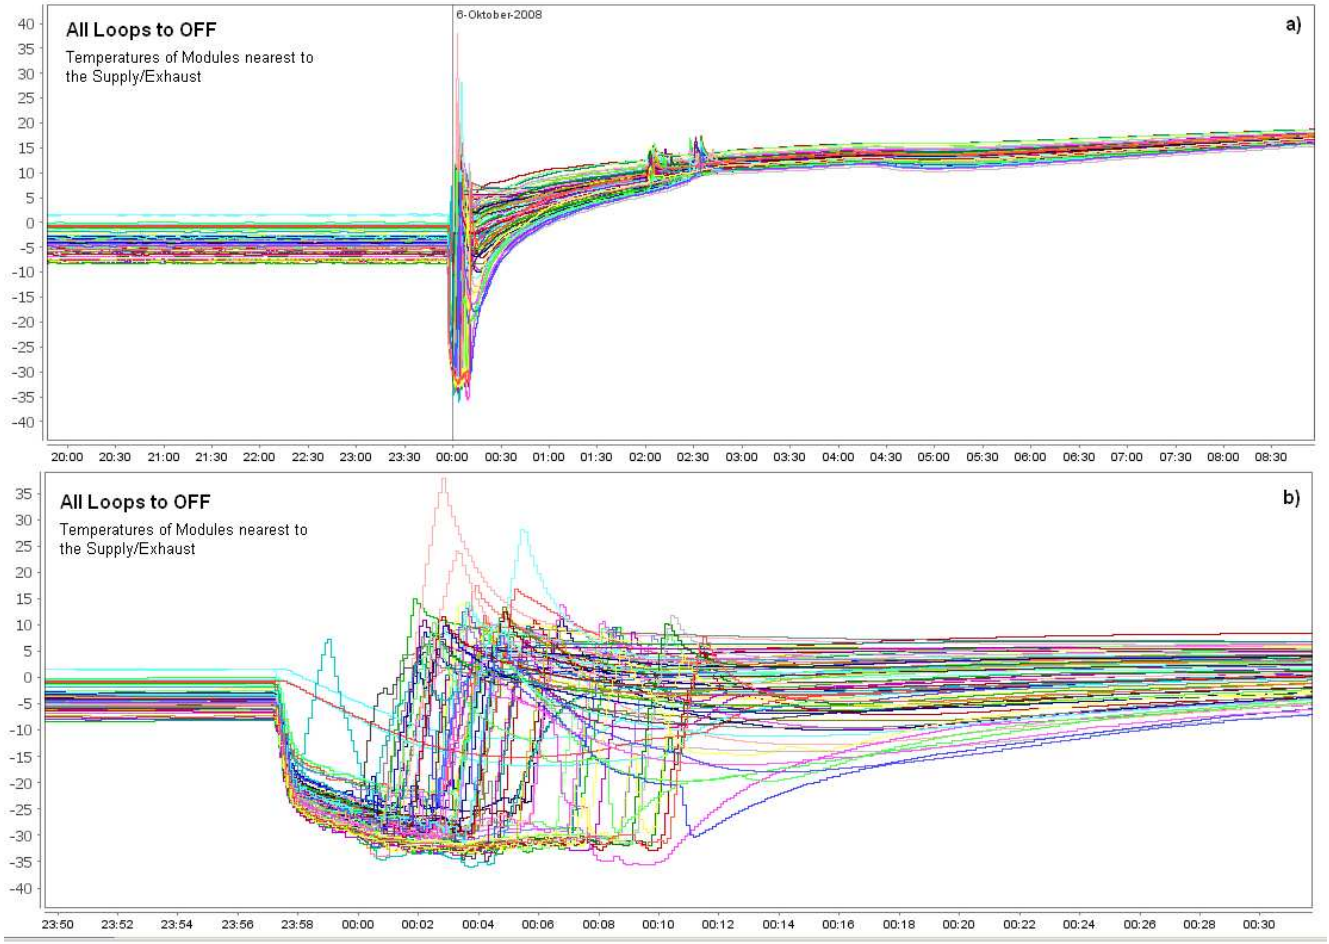

Figure 7.2: *Trend of module temperatures for a cooling shutdown with all cooling loops going to* OFF*, in October 2008. The module temperatures for the modules closest to the exhaust of each loop are shown. In (a), the evolution over a longer timescale is shown, (b) focuses on the time around the incident.*

The Action Matrix

All the actions can be masked, to permit limited operation without cooling, or only with partial cooling.

Technical Implementation

The actual technical implementation is based on regular checks in intervals of 10 s of the plant state and cooling loop states corresponding to a given readout unit. The granularity is chosen as the readout unit is the natural unit of the Pixel FSM, and due to the fact that a given module loop always has more than one associated optoloop. The time based check gives an additional measure of robustness, especially in case of connection problems between the systems. Also, as there is no action in case of an already reached safe state, it makes sense to have the same interval for FSM state calculation and cooling script, as the cooling script is embedded in the state calculation script. This has the additional advantage that the cooling script has access to the internal variables which manage the state. Therefore it is not dependent on the overhead of the FSM framework with DIM communication, which even might not be running at a give time, but relies only on PVSS. Furthermore, the FSM can access the resulting evaluation of loop state per readout unit, and can veto accordingly commands that would apply power to an uncooled detector unit.

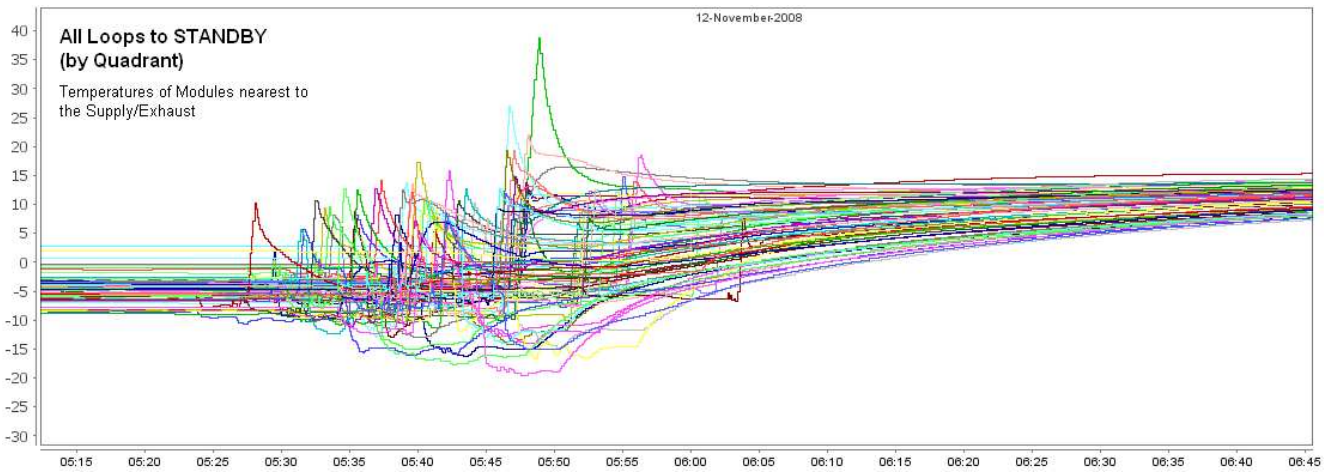

Figure 7.3: *Trend of module temperatures for a cooling shutdown with all loops going to* STANDBY*, in November 2008. The module temperatures for the modules closest to the exhaust of each loop are shown. Due to the fact that this was a manual switch off, the loops were not switched off all at the same time, but by quadrant. Problematic loops in Quadrant 4 were switched off separately.*

modules state	opto state	FSM state	detector loop failure	opto loop failure	detector and opto loop failure
on	on	READY	BACK_TO_OPTO_ON	GOTO_NO_OPTO	SW_INTERLOCK
off	on	OPTO_ON	no action	GOTO_NO_OPTO	GOTO_NO_OPTO
on	off	NO_OPTO	BACK_TO_STARTED	no action -	BACK_TO_STARTED
off	off	STARTED	no action	no action	no action

Table 7.2: *The Action Matrix of the Cooling Script, showing the initial FSM- and cooling loop state, and the corresponding action in case of a failure of one or both loops*

Unavailability of Cooling Information

Also the question of how to handle the unavailability of information provided by the Cooling project has to be addressed. In this first section, only the case is treated where the cause for the communication loss does not lie within the Pixel project, which has to be considered separately.
In the initial approach, it was assumed that in case of communication problems, the connection would be re-established rather fast, in a timespan of two minutes. This has proven to be extremely aggressive, and even more so in the current mode of operation where there is not necessarily an expert available in a comparable timespan. While it is expected that these problems are primarily pure control system problems on the level of the ID project or the network, and normally not related to the plant itself, it is not possible to recover that fast. However, it was also considered unsafe to enlarge this timespan to aribitrary limits, especially as the remaining software and hardware interlock threshold is too high for comfortable emergency protection for the complete detector. To avoid

both, switching off the detector with an operational cooling system, as well as operating the detector for a longer timespan with one of its safety mechanisms unavailable, a separate solution had to be found.

For these reasons, in case of a disconnection, the module and optoboard temperatures are monitored continuously, and a change above a certain threshold will be taken as indication of a cooling problem. The chosen threshold is a temperature change of more than 5°C with respect to the temperature at the connection loss, and the implemented reaction is the same as if the loop failure would be indicated by the cooling-project loop-state itself, i.e. the action matrix applies.

In the case that the problem is not in the ID project, the limitations for an automated actions are met. If the FSM project is disconnected, evidently no actions are possible. If the connections are lost to the cooling project and the Pixel temperature monitoring project at the same time, no action is taken, as the initial temperature and all following will be marked as invalid. In this case an expert needs to evaluate the situation, as any automated action would only work on part of the system.

7.1.2.3 Summary and Conclusion

In summary, the cooling script meets the requirements outlined in 7.1.2.1. Based on the loop states of module loop and optoloop, for each readout unit the command to remove the power from the affected detector unit – modules or optoboard – is issued immediately in case of a failure, and each command that would re-apply power is vetoed. Next to the increased flexibility in operating the detector in case only part of the cooling system is available, this avoids thermal cycling of the non-affected detector part and helps to keep the cooling system stable by not removing heat load from an operated loop.

The question of how to react on an unavailability of the external cooling information is also strongly correlated to the desire to avoid thermal cycling and operational downtime. While in the first days of operation, both of the silicon detectors – Pixel and SCT – had the same action of a complete shutdown after 2 minutes, this was changed for the Pixel Detector to a temperature based software interlock reacting on a temperature change instead of a fixed time based threshold.

Minimizing completely the thermal cycling in both directions is not a trivial task. This is given mainly by the complexity of both systems, and the desire to keep the implementation simple. Nevertheless, both types of temperature gradients are reduced with respect to early operation. The most important point is the fact that the cooling loops will not fail into the state OFF, so that it is possible to switch off the detector devices immediately without worsening an already steep temperature drop.
As the detector will be switched off 5 min after a plant failure, any attempt at further optimization would only apply to scenarios where the plant stays ON. Given the fact that the temperature gradient of the detector modules at switching on of the cooling plant is much more severe than during the failure itself, further optimization in the production system is currently not envisaged.

Figure 7.4 shows the reaction of the system as it is currently in place.

The temperature gradient due to the immediate removal of module power can be clearly seen. It amounts to a temperature drop by about -10°C, compared to -25°C when keeping module power but putting the loop to the OFF state. Even though the software action is supposed to act in an immediate way, the trend shows different reaction times for different readout units. This is explained by the fact that the commands are partly serialized in the FSM due to properties of the lower level software. It can be seen that after about five minutes the the first readout units show a rising temperature, indicating the complete evaporation of the cooling fluid, while for some loops this takes as long as 15 min. This indicates that it is not desirable to reduce the delay of the DSS action.

Currently efforts are ongoing to reduce the temperature gradient at switching on, by making this a two step procedure, swiching on the loops with a high back pressure setting, and reducing the back pressure to the nominal value in the second step after some stabilization time. In case this brings the temperature gradient down to a level

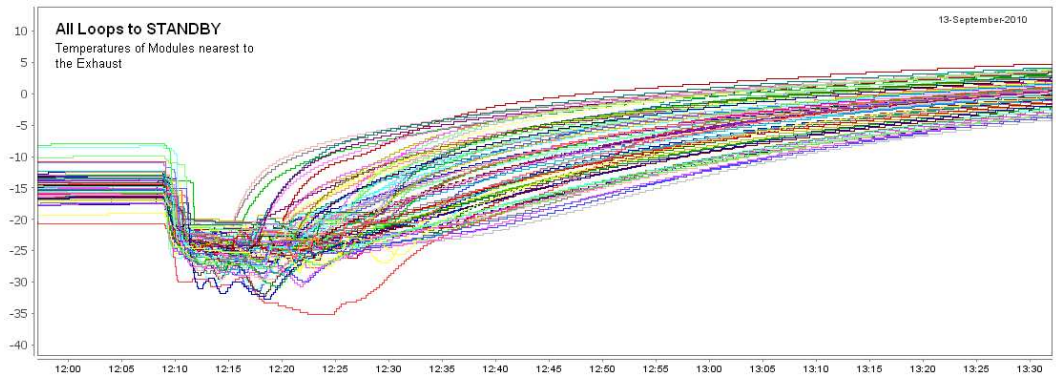

Figure 7.4: *Trend of module temperatures for a cooling shutdown with all loops going to* STANDBY*, in September 2010. The module temperatures for the modules closest to the exhaust of each loop are shown. Seen is the initial drop of module temperature due to the immediate switch off at the time of failure, the period of constant temperature until the remaining fluid is evaporated, and the subsequent rise in temperature.*

comparable to the switching off, this might result in a re-evaluation of the switching off action, within the above constraints.

In summary, the cooling script provides an additional layer of safety in case of a cooling failure by switching off the affected detector units with a high granularity before the temperature threshold of the hardware interlock is reached.

7.1.3 Software Actions on Inputs from the Global Detector Safety System

As described in chapter 2.2.7, detector wide safety actions are defined within the DSS matrix. These Alarms are provided via DIP to ATLAS DCS, and are from there available to the subdetectors.

This makes it possible to implement a software action triggered by the DSS Alarm, with the goal of ramping down the affected power supplies in a synchronized and coherent way before the hardware interlock shuts them down forcibly in an unpredicted order. While this high level information provides information not available inside the subsystems, e.g. smoke, and allows one to assess a situation faster and on a more abstract level than e.g. based on information gathered only from subsystem internal sensors, there is no defined requirement concerning the availability for DCS.

All DSS alarms concerning the Pixel Detector are described in the document [82]. There are two classes of alarms: Alarms which are considered critical enough to require an immediate hardware action, and on the other side alarms with a delayed action. The alarms belonging to the first group with an immediate hardware action will not be discussed here, as their effects can only be monitored by the detector DCS, with no possibility for a preemptive software action. This category includes e.g. smoke detected in the inner detector.
On the other side are the alarms with delayed actions, which can with respect to their relevance to the Pixel DCS also be grouped into two categories: alarms with an action concerning the evaporative cooling plant, and all others.

The DSS alarm regarding the evaporative cooling is the most relevant software action to be considered.

For the Pixel Detector, the structure of any DSS related item (see also see chapter 2.2.7) consists of

- sensors providing the inputs for the Action Matrix.
- an alarm, triggered by a logical combination of input conditions.
- the action, which is executed with a given delay after the alarm was triggered.

These informations are provided to the ATLGCSIS1 project, and the Pixel SCS connects to this project. For the implementation of software actions, the alarms on which to react have to be specified, the desired software action defined, and the technical implementation which is best suited to these requirements determined. A clean way to mask the action is required in addition.

7.1.3.1 Pixel DSS Interlocks

The Pixel interlock system can distinguish three DSS inputs (see chapter 5.3.3) called here DSS0, DSS1, and DSS2. The results are defined in the Pixel interlock matrix:

- **DSS0** switches off all High Voltage, Low Voltage and Optolink supply voltages for the complete detector. All Optoheaters are interlocked. The auxiliary supply voltage for the PP2 voltage regulators is switched off. CAN power, BBIM power (this makes a continued monitoring of detector temperatures possible) and off detector lasers are not affected.
- **DSS1** switches off all High Voltage Supplies.
- **DSS2** interlocks the bakeout carts used for the bakeout of the beampipe.

7.1.3.2 Relevant Alarms

The following DSS alarms are candidates for having defined a DCS software action:

- AL_COL_ID_PlantNotRunning
  - triggered immediately by cooling plant non operational
  - action is delayed by 5 min
  - DSS action: DSS0
- AL_COL_Cables_CoolingFailure
  - triggered after 5 min by cable cooling non operational
  - action is delayed by 20 min
  - DSS action: DSS0
- AL_INF_CoolingUS15_RackFailure
  - triggered after 300 s by rack cooling failure for the US counting room
  - action is delayed by 30 min (1799 s).
  - DSS action: all racks on US side will be switched off.
- AL_INF_CoolingUSA15_RackFailure
  - triggered after 300 s by rack cooling failure for the USA counting rooms.
  - action is delayed by 30 min (1799 s).
  - DSS action: all racks on USA side will be switched off.
- AL_INF_Power_USA_UPSonBattery
  - triggered after 5 min if the diesel generator fails to start.

  - action is delayed by 5 s.
  - DSS action: DSS0.

- AL_INF_Power_US15_UPSonBattery
  - triggered after 5 min if the diesel generator fails to start.
  - action is delayed by 5 s.
  - DSS action: DSS0 for power supplies located in the US15 counting room.

Alarms which fall into the "indirect DCS action via cooling DSS alarm" are predominantly smoke alarms detected by sniffers in areas not directly related to the inner detector.

AL_COL_ID_PlantNotRunning

For the required DCS action, the same arguments as described for the cooling script in section 7.1.2 are valid. Accordingly, the implemented action is to switch off all readout units immediately when the alarm is triggered. After the detector has been switched off, the auxiliary power supplies for PP2 are also switched off.

AL_COL_Cables_CoolingFailure

In case of a cable cooling failure, the danger is overheating of the cables and mainly of the regulator stations at PP2. The Pixel hardware interlock due to overtemperature is triggered at 60°C, the software interlock at 45°C. The implemented software action on the DSS alarm is analog to the hardware DSS action, i.e. switch off all readout units and PP2 supplies, while the time of the software action is chosen to stay below the software interlock threshold and well below the hardware temperature limit. This is given for a reaction time 15 min.

AL_INF_CoolingUS15_RackFailure and AL_INF_CoolingUSA_RackFailure

The two concerns at these alarms are the temperature of the crates installed in the affected counting room, and a clean ramping down of the affected detector devices by software. For the temperature aspect, this is mostly covered by the defined delay of the DSS action, so the remaining concern is to switch off the corresping detector parts before the hardware interlock is triggered. This implies a software action after 25 min. Concerning the affected detector parts, the behaviour is different for US and USA side, due to the fact that the BBIM power supply for the complete detector is located in USA15. This will cause a hardware temperature interlock for the complete detector as soon as the DSS action switches off rack Y1514A2 which houses the power supply crate for the temperature monitoring and interlock crates (BBIM). Additionally, loss of rack power on USA side implies loss of control of the regulator, and switching off of all off-detector lasers. A power cut on US side affects readout units powered from USA side only by the fact that the corresponding optoheaters are switched off, which prevents data taking, but does not require an additional action.
In summary, for a rack cooling failure on US side, all readout units powered from there are switched off after 25 min. For a rack cooling failure on USA side, the implemented action is to switch off the complete detector.

AL_INF_Power_US15_UPSonBattery

In case of a power failure, diesel generators will provide power for the UPS. In case of a failure to start the diesel generator, backup supplies guarantee 10 minutes for a clean shutdown. In the US counting room, the network switch for the Wiener power supplies is supplied by UPS, to prevent the loss of control in case only the rack power of the rack housing the network switch is cut, while the Wiener power supplies might be still providing power to the detector. Even in this case, however, the power supply of the detector can still be remotely removed, by

switching the corresponding regulator channels. Therefore, the main concern for this alarm is clean shutdown by software, before the hardware interlock is triggered. This is given by the fact that equipment of the evaporative cooling system is affected by this failure, and so AL_COL_ID_PlantNotRunning will also be triggered with 300 s latency and 0 s delay. No additional software action on this specific alarm is implemented.

AL_INF_Power_USA_UPSonBattery

While a UPS failure in the US counting room mainly requires a clean software shutdown of the detector, this is a more serious issue on the USA side. Here, all important DCS control stations are powered by UPS, and the CAN power supply. In case of a UPS failure, this means that the complete DCS monitoring and control of a potentially powered detector is lost. For this reason, the DSS action of shutting down the complete detector is executed 5 min after the alarm is raised. While the DSS action takes care of the detector safety, the DCS PCs will experience a hard shutdown once the power stops. This bears the risk of a corruption of the operating system or the PVSS internal database, causing additional downtime for recovery or even re-installation of the system. This can be prevented by shutting down the PVSS projects and the PCs within 5 min after the alarm. There is currently no action implemented to preemptively shutdown the DCS PCs, as a dual powering of the PCs is under discussion.

7.1.3.3 Implementation

The DSS information is used in two ways: a passive "veto" prohibits commands that would provide power to the detector, and the active part – referred to as the "DSS script" – that executes the defined action.

No action is taken when the information is unavailable. For the evaporative cooling, there is another source of information provided by the ID cooling project, and temperature sensors read out by the Pixel DCS are also available with respect to monophase cooling failure. The hardware DSS action will take place regardless, so a preemptive DCS software action is not justified. As far as the veto is concerned, i.e. prohibiting the powering of the detector when there is no DSS information available, there is no automatic check in place. Instead, this has to be decided on a case by case basis. Information might be unavailable if the information is not published by the DSS system, or if the client subscribing to the information is not running in the central DCS project (ATLGCSIS1), or if the connection between the Pixel PVSS project and the central DCS project is lost.
Additionally, it is possible to deliberately ignore an occuring DSS alarm, both for the passive veto and the active software interlock.

Due to lower level of guaranteed availability for the DIP-publications, for the DSS script an event driven approach was chosen to trigger the DCS action on the appearance of a DSS alarm. This provides slightly faster reaction with less load for the system, but on the other hand there is less robustness in case of a PVSS project disconnection. The structure of the cooling (evaporative, monophase, or rack-cooling) related DCS actions is shown in figure 7.5:

Any update of the DSS alarm or its validity information will trigger the DCS callback function. In case of an active alarm, it is checked wether the information is invalid (and therefore potentially obsolete), and wether the action is masked. If this is the case, no further action is taken. If there is an active, valid alarm, and an action is required, the script determines the affected detector parts (mainly divided by US or USA counting room), the expected time of the DSS action, and the delay for the DCS software action. The DSS action will take place once the alarm is active, therefore the DCS action should be executed regardless wether the DSS alarm is still active or not. However, while the delay is running, it is still possible to mask the DCS action. At the end of the delay, the shutdown is initiated. To avoid interference with ongoing actions and redundant commands, the state of all affected readout units is checked. If no command is active, and a safe state not yet reached, the shutdown command is sent, until all readout units are unpowered. If applicable, the auxiliary PP2 power supplies are switched off afterwards. Under

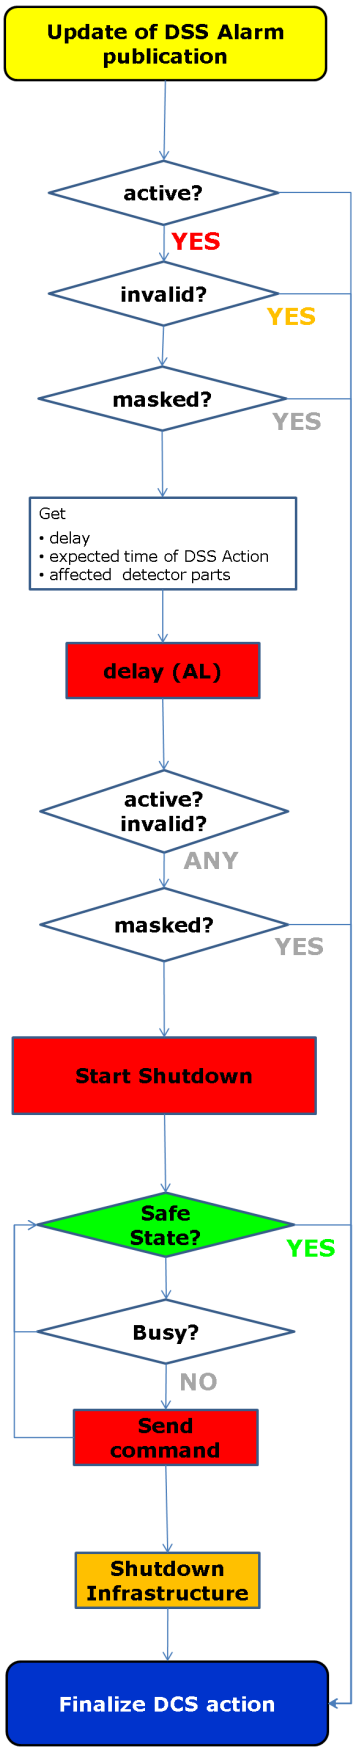

Figure 7.5: *DCS actions on active DSS alarm.*

regular conditions, the DCS action should be finalized before the DSS action takes place.

The DSS script is running centrally for the complete Pixel Detector on the sub-detector control station. This makes it possible to react in a global way to DSS alarms, which concern large parts of the detector. If the need should arise, this simplifies the task to minmize CAN traffic and command execution time. It is traded off by the fact that the DSS script does not have access to the internal state calculation done by the FSM, but must get it from information saved into datapoints. The synchronization of ongoing actions becomes more difficult, but as all processes – FSM and cooling script – respect the TRANSITION state, this affects only commands that are ongoing at the moment of the alarm. For this reason, the most time intensive *STARTUP* command has periodical checks wether the conditions for execution still apply. An *ABORT_COMMAND* flag set by global scripts and checked by the FSM at key points during command execution is under consideration.
In summary, it is still considered more advantageous to handle the DSS alarms globally.

7.1.3.4 Conclusion and Outlook

The publishing of DSS information via DIP provides an additional high level criterion to trigger software actions, albeit with an accepted less than 100% availability. A central event driven action was chosen, which reflects the global nature of the affected issues. The script provides a clean rampdown of affected detector parts versus a hard shutdown of possibly only part of the system.
What remains to be done, is a definition of desired software actions for additional DSS alarms, and an ATLAS wide streamlining of the implemented software actions with respect to other subdetectors. This concerns mainly the supervision of triggered actions. Especially in view of planned reduction of shift crew, active or impending actions must be easily traceable at any given moment, and successful or failed completion should be automatically logged. Information about the masking of action must be easily accessible.

7.1.4 Handling of Beam Related Information

Although it was tested that a production Pixel module survives an instantaneous radiation dose [102] higher than for the pilot beam, studies from CMS with high intensity lasers have shown that it is possible to significantly damage the frontend electronics. Large quantities of injected charge are flooding the sensor with charge, causing a "short circuit" through which high voltage is directly applied to the sensitive frontend electronics and destroying it. These results have moved the Pixel community to define "high voltage off" as the safe state in presence of beam.

The aim of the work described in this section was to implement safety mechanism that can enforce a safe state of the detector in all upcoming situations, while avoiding the disruption of accelerator operation.

7.1.4.1 The Hardware Signals of the PLI Box

To handle beam related safety issues, the PLI box (described in chapter 5.3.3) was installed in the Pixel interlock system. A signal from the LHC indicates if there are *stable beam* conditions or not, and an interlock to the high voltage power supplies, DSS1, is generated in case of absence of stable beam conditions. However, the LHC signal distinguishes only between "stable beam" and "no stable beam", and there is no signal provided indicating the "stable absence of beam". Therefore it is necessary to provide a very elementary possibility to mask the interlock, as it would be otherwise impossible to calibrate the detector with depleted detector modules during no beam periods.

To ensure that LHC can not inject beam while ATLAS is not in a safe state, the ATLAS Injection Permit System (IPS) provides an *injection permit* if ATLAS is in a safe state. The ATLAS injection permit is a logical AND of the individual subdetectors' injection permits. It is possible to mask the individual subdetector inputs. For the Pixel Detector, the signal to the ATLAS IPS is provided by the PLI box. It is a hardware signal, but not generated from hardwired logic inputs, as it is not foreseen to have a high voltage output as input into the Pixel interlock system. Instead, the injection permit has to be generated by software. The number of unneccessary vetoes should be kept to an absolute minimum.
Due to its special status and rather late addition to the system, much of the basic safety aspects had to be implemented in software. Also, the PLI box is the only Pixel interlock box where the hardware interlock can be circumvented by software settings with the so-called *"mask"* signal.

The PLI-Box provides three *hardware* signals:

- The Stable Beams signal from the LHC (''STABLE BEAM"/''NO STABLE BEAM").
- The possibility to mask the hardware interlock (''MASKED"/''UNMASKED").
- The Injection Permit to the ATLAS IPS (''PERMIT"/''VETO").

7.1.4.2 Software generated PLI Box Outputs

The requirements for the PLI software can be summarized very generally in the following way:

Requirements for the PLI software

- Ensure that the outputs of the PLI-Box are always set in a way that does not compromise the safety of the Pixel Detector.
- Do not delay LHC operation (unneccesarily).
- Do not constrict the flexibility for switching on high voltage when it is considered safe.

- Robustness of the software protection.

The first three points are listed in order of priority. The first point includes an automatic generation of output signals to prohibit unsafe conditions. LHC operation should not be unneccesarily delayed. The main relevant causes in this perspective that could result in an interference with LHC operation is the failure to grant the injection permit and lateness in reaching a safe state. While it should be prevented to switch on high voltage in unsafe conditions, the highest compatible amount of flexibility to switch on high voltage should still be given. All this should be implemented in a highly reliable and robust way.

The resulting software safety measures can be divided into two parts. The first part is implemented on a very low level, and is automatically present and active as soon as the interlock project is up and running and receiving correct signals from the PLI hardware.

Event driven Output Generation

The event driven part of the PLI software is implemented on a very low level[1)], and is automatically present and active as soon as the interlock project is up and running and receiving correct signals from the PLI hardware. This part checks if the stable beam signal is present, and ensures that the outputs for MASK and PERMIT are set to the corresponding desired values. By implementating the following logic, also illustrated in figure 7.6, already a certain amount of safety is provided which can only be sidelined to cause an unsafe situation by human action.

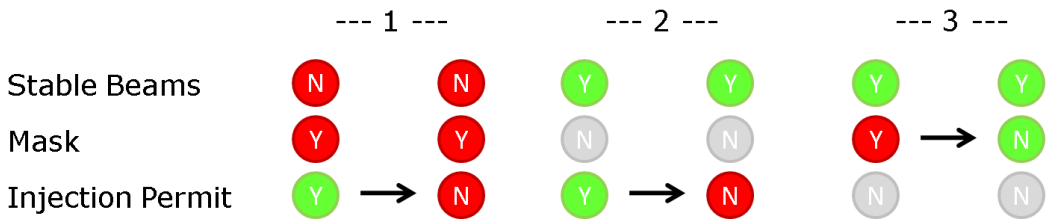

Figure 7.6: *Logic of the automatized PLI-Signals. This is implemented on a very low level of the PVSS project, guaranteeing that this part is active as long as the event manager is running.*

1. NO STABLE BEAM and MASKED and PERMIT: set PERMIT to FALSE
   (No injection permit while masked, as it would be at any time possible for the operator to switch on high voltage)
2. STABLE BEAM and PERMIT: set PERMIT to FALSE (No beam should be injected during stable beam)
3. STABLE BEAM and MASKED: set MASKED to FALSE (No masking of LHC signal during stable beam. This is necessary to guarantee that the hardware interlock can fire when the beam conditions get unsafe.)

This guarantees the following minimum safety conditions in a very basic “hardwired” fashion:

- When the Injection Permit is given, high voltage is interlocked.
- The hardware interlock can not be deactivated when there is stable beam.

However, since there is only limited information by hardware, these items can not cover all possible situations. It can only ensure that dangerous settings can not occur without human intervention. Figure 7.7 shows the nominal settings, and a comparison to possibly occuring situations. The table only considers the desired logical combinations with an abstract danger potential, while the actual risks are left for a later discussion.

[1)] as a dp-Function (see chapter 5.2.5)

Nominal Settings	No Beam, Calibration	Beam Injection	Ramp	Stable Beams	Adjust Handshake or Beam Dump	Unstable Beam or No Beam
Stable Beam Flag	No	No	No	Yes	Yes	No
Masking of Stable Beam Flag	Yes	No	No	No	No	No
Injection Permit	No	Yes	Yes	No	No	Yes (if STANDBY)
HV On?	Yes	No	Yes	Yes	No (Handshake) Yes (Unplanned Dump)	No

Possible Settings	No Beam, Calibration	Beam Injection	Ramp	Stable Beams	Adjust Handshake or Beam Dump	Unstable Beam or No Beam
Stable Beam Flag	No	No	No	Yes	Yes	No
Masking of Stable Beam Flag	Yes	Yes (Not prevented)	Yes (Not prevented)	No	No	Yes (Not prevented)
Injection Permit	No	No (Beam already in?)	No	No	No	No
HV On?	Yes	Yes	Yes	Yes	Yes	Yes

Figure 7.7: *Logic of PLI-Signals and possible HV states. The table columns follow roughly the LHC operation cycle and its meaning for the detector. The upper table shows nominal settings, the lower the worst combination of settings that could occur. The color coding is according to a "good" state for each item. In any situation by definition not all signals can be in their "good" state at the same time. This is marked as yellow or grey color. E.g. during stable beam (TRUE = good), Injection permit (TRUE = good) is not given. However, this is expected and accepted and therefore marked as grey. Masking of the stable beam flag should only happen when there is no beam in the machine, where it is acceptable, and therefore marked as yellow. Dangerous situations are marked as red, and can only happen by human error, therefore not appearing in the nominal case. Here, the only dangerous situation is during unplanned, unclean dumps, when high voltage was not switched off before. More details for each point are given in the text.*

- **NO BEAM (Calibration during long access periods)**: If there is a longer period without beam (machine mode ACCESS and beam mode NO BEAM), no injection permit is required, and the stable beam signal (which is FALSE) can be safely masked to switch on high voltage. The only non-regular issue can originate from LHC, where a fake stable beam signal could at its removal cause an interlock to the Pixel high voltage. If this is anounced while calibrations are planned, the Pixel software safety measures have to be deactivated.

- **Beam Injection**: This phase refers to beam modes INJECTION PROBE BEAM and INJECTION PHYSICS BEAM. The nominal scenario for beam injection is to bring the Pixel Detector into STANDBY, and give the injection permit to LHC. The main worry is to delay LHC operation by withholding the injection permit. However, there is no software mechanism to prevent masking of the stable beam signal in this stage of operation. The worst scenario is that the stable beam signal is masked by human error during injection, with the beam already partly in the machine. The injection permit is then automatically removed, preventing LHC from finishing injection. It is then possible to switch on high voltage, with beam circulating in the machine, or even during the injection of the pilot beam.

- **Ramping and preparation for stable beams**: The same arguments as in the injection phase are valid, with the difference that the injection permit is not required anymore. This phase includes normally the beam modes PREPARE RAMP, RAMP, FLAT TOP, SQUEEZE, and ADJUST.

- **STABLE BEAMS**: At stable beams, it is possible and safe[2)] to switch on high voltage. Injection permit is automatically removed, but not required by the LHC. If the stable beam signal gets masked, it will be

[2)] the operational procedure to verify safe beam conditions is defined in [104].

immediately returned to the unmasked state, so ensuring that the Pixel high voltage interlock is "armed" in case of a change of beam mode[3)].

- **Adjust Handshake and Beam Dump**: In case of a planned change from STABLE BEAMS to ADJUST or BEAM DUMP mode, a handshake between LHC and the experiments takes place. LHC will warn of the impending beam mode change, the experiments will move to STANDBY, and inform LHC that a safe state is reached. Only then the stable beam flag will be dropped, and the beam dumped or major beam adjustments made. In case of an unexpected beam dump, high voltage can not be removed fast enough. For this reason, the stable beam flag is kept even after the beams are dumped, to allow the experiments to ramp down their high voltage without beam in the machine. Once the stable beam flag is dropped and the detector is in STANDBY, the Pixel injection permit will be automatically granted to prepare for the next operation cycle. The result of the online analysis of the BCM post mortem data is only considered for granting or withholding the global ATLAS injection permit, not the Pixel or BCM permit.
  While the conditions are safe once the beams are dumped, in case of a handshake for beam mode change, potentially dangerous scenarios are not prevented. STANDBY is the nominal condition during a planned beam dump and the change to adjust mode. The rampdown of the high voltage was a manual operation during the 2010 running period, and the ATLAS confirmation for having reached a safe mode is still left to human discretion.

- **Unstable beam conditions or no beam**: After the beams are dumped and before the next injection phase (beam modes RAMP DOWN, NO BEAM, CYCLING, SETUP), in the time needed for ramping down the magnets, it is generally foreseen to keep the Pixel Detector high voltage off[4)], and already grant the injection permit, even if it is not yet immediately required. Also with beam still in the machine after an adjust handshake, STANDBY and granted injection permit are the nominal conditons, but here it is mandatory to keep high voltage off.
  For both cases, there is no mechanism in place to prevent the masking of the stable beam signal. It is therefore theoretically possible to mask the interlock after a beam dump or after an ADJUST handshake and to switch on the high voltage with potentially unstable and dangerous beam conditions.

While it is possible to create potentially dangerous situations by masking the stable beam signal, this can not be prevented without access to external information. The access for masking the interlock is highly restricted, in fact only three people[5)] have the privileges to do this. For all other cases, the described safety measures provide a relatively high level of protection, especially given the fact that even if the control is lost, without active intervention there is no possibility to make any situation more dangerous: If high voltage is on, the injection permit is not granted, and if the injection permit is granted, it is not possible to switch on high voltage. If high voltage is on with stable beam, it is guaranteed that the hardware interlock is not masked.

Additional Safety Checks of the PLI Script

In addition to the low level, event driven part, a supplemental safety layer is provided by the "PLI-script". A function is executed at regular time intervals (10 s), to

- check that the current output settings are consistent with the desired logic described above.
- take into account additional information, namely high voltage outputs
- remove the potentially dangerous masking, if it is unused

[3)] this is relevant only for a change into ADJUST mode. in case of fastly deteriorating beam conditions and a resulting dump, the interlock is not fast enough.

[4)] Depending on the LHC turnaround time, and the ATLAS policy of allowing the subdetectors calibration with high voltage only in case of downtimes with a foreseen minimum length

[5)] The Pixel Project Leader, Run Coordinator and Deputy Run Coordinator

- check that the event driven logic is not disabled
- provide a heartbeat for external alerting if the safety measures are inactive
- automatically generate the injection permit, if it is safe

The script does not introduce new items, but complements the basic mechanism. It regularly executes a cross-check that the state of the output signals is consistent with the defined signal logic, and corrects it if this is not the case. Furthermore, it accesses the high voltage information provided by a different project of the Pixel DCS[6)]. It is theoretically possible in some very special cases that high voltage is on even though the PLI outputs are, or seem to be, set correctly. For this check, a high voltage channel is considered ON when the voltage exceeds 5 V for more than one reading.If any channel is considered as powered, the injection permit is removed, and an alert is raised. The PLI script does not take any action on the high voltage control. Even though the state of the preamplifiers are considered a safety aspect, the PLI script itself checks only HV channels in the main routine, as the module configuration is nothing that can be controlled inside DCS.
A channel is considered UNKNOWN in case communication problems are detected by the watchdog in the iseg project. To avoid an unnecessary removal of the injection permit, no action is taken in case of UNKNOWN channels. Only an extremely obscure combination of circumstances would result in powered channels in a case where no communication to the iseg power supplies is available.

As an additional safety measure, the masking of the hardware interlock is removed if it is unused for a specified amount of time. I.e. if the signal is masked without switching on high voltage, or when high voltage is switched off without removing the mask, the script will do this after a certain timeout automatically.

The script provides some mechanisms against unintended disabling of the software protections. It monitors the presence of the event driven part of the PLI logic, and generates a heartbeat to provide the possibility for external monitoring of the PLI-script. This is the task of a dedicated watchdog, that alerts shifter and expert in case of the protection is disabled.

With gained operational experience, it was decided that the fact that injection permit can only be set by run coordinators and project leader is not feasible. Therefore, an additional functionality has been implemented in case the FSM top node is in STANDBY. If no internal check of the PLI script is preventing it, the injection permit will be generated automatically.

This main part of the PLI-script is an additional safety layer, relying only on Pixel internally available measured information of the interlock and high voltage system. Since 2009 there is an extension of the PLI-script to cover two of the remaining ways to reach an unsafe situation. For this it has however rely on external information published via DIP, or provided by the ATLAS LHC Interface.

7.1.4.3 Extension of the PLI script

Before the ATLAS injection permit is granted, the operator has to evaluate wether all subdetectors are in a safe state. For this end, all subdetectors provide very simple software flags that are propagated to the ATLAS GUI for LHC handshake interaction. For the Pixel Detector, this global flag consists of two local flags. The first one indicates “safe for beam” if the state of all high voltage channels is off, and the second flag signals that Pixel is not safe for beam if the hardware interlock signal is masked at any time.

[6)] For this reason, this information can not directly influence the calculation by the dp-Function.

One way to reach an unsafe state is the situation where ATLAS is masking the input to the injection permit signal provided by the Pixel Detector. It is then possible for LHC to start beam injection while the Pixel Detector high voltage is still on. The aim is to detect this situation and react in a way to reach a safe condition, even if the ATLAS injection permit is granted without Pixel being in a safe state. As the masking of the Pixel injection permit is not under the sub-detector's control, a safe situation can only be ensured by removing high voltage.

The masking of individual sub-detectors is done in the ATLAS IPS, and the setting and readback value are available via PVSS distributed connection to the ATLAS LHC project (ATLGCSLHC). The script takes no action during stable beam. In case the Pixel injection permit is masked by ATLAS, it provides the following functionality:

- masking of the stable beam signal is prevented.
- high voltage is ramped down, and an existing mask of the stable beam signal is removed afterwards.
- in case of a longer unavailability of the information, high voltage is ramped down and the masking of the stable beam signal is removed.

This alleviates one risk scenario, which relied previously exclusively on correct human perfomance, and where no hardware protection is given against reaching a potentially dangerous situation.
Another danger is given by the fact that it is not possible to decide by information available to the Pixel DCS hardware if it is safe to mask the stable beam signal. A possible gain in safety could be achieved by preventing that for certain beam modes which are not considered definitely safe. This information is available from the ATLAS project which subscribes for information transmitted via the DIP protocol from the LHC. However, during operation with 1018 GeV beam energy in the end of 2009, the case already occured that the Pixel Detector was operated without the stable beam flag present. On the other hand, no incident of an unintended masking has happened yet. Therefore, this possibility to restrain the use of the interlock masking is currently not considered to be of immediate importance.

7.1.4.4 Conclusion and Outlook

The main task of the PLI software is to ensure that the Pixel Detector is in a safe state during beam injection, while keeping the occurences of fake vetoes at an absolute minimum. Furthermore, the overriding of the interlock caused by the absence of the stable beam signal has to be restricted where possible, while not preventing its use where it is adequate. The resulting implementation has been described. A high robustness is given by splitting up the monitoring into separate entities, with complementary advantages. The lowest level provides a basic control of the output signals which ensures safe operation under nominal conditions. It is extremely self-contained, providing a high level of availability and robustness. A higher level script complements this mechanism with additional information sources, internal supervision functionality and a different execution mechanism.
Remaining risk scenarios have been discussed. Partly, they are addressed by taking into account information from outside Pixel DCS, and partly by restricting access to related functionality.
In more than one year of operation, the software has worked very reliably in the intended way, with only minor instances of too much caution related to fake spikes in the high voltage measurement readout. No instances of removing a requested injection permit have occured.

7.2 Procedures

In the course of the work for this thesis, the automatizing of semi-frequently occuring procedures, dealing with complex – and therefore error-prone – scenarios, was approached and implemented.

The aim of automatizing these procedures is to minimize user errors or operational delay, due to the fact that shifters are not familiar with the corresponding operation.

Since these procedures are normally considered as rather high level operation, and the need to execute the procedures is depending on the fact that a fully operational DCS is already available, they can in their implementation rely on FSM availability and functionality to a very high degree.

7.2.1 Cooling Startup

For the recovery from a cooling shutdown, the systems – the silicon detectors and the cooling system – have to be brought up in a coordinated and synchronized way. The aim is to minimize on one side the time where the cooling system operates without the heatload from the detector, as some loops are very sensitive to this, and on the other side also optimize the required recovery time.

However, the partitioning of both the ID FSM (by subsystem[7)] and quadrant), and the Pixel FSM (by Pixel DAQ partitions and mechanical layers), is not optimized to deal with the actual procedure for switching on the cooling. Normally, this happens by quadrant, and for SCT and Pixel Detector in parallel, as both subdetectors desire a fast recovery. On the ID side, there is a fix delay of 5 min between two sets of commands, driven by the PLC to ensure a stabilization of the conditions during switch-on, before another change request is accepted. This delay determines the startup time, especially as commands for Pixel Detector and SCT can not be given in parallel through the standard ID FSM interface.

Once the loops are started, the other side of the issue is that at this point the corresponding detector parts should be powered as fast as possible, to avoid that unstable loops have to be operated without heat load for too long. Even though this was made easier by regrouping the Pixel PCCs inside an additional FSM branch according to quadrants, certain state combinations could still prevent the possibility to give commands at the quadrant level. In the worst case, it was necessary to issue commands on the readout unit level.

To simplify and accelerate the startup by automatic procedures, a protocol was defined by ID, Pixel Detector and SCT. On ID side, this provides the possibility to switch on Pixel and SCT loops in parallel, while on the detector side the script takes care of powering all detector elements belonging to started loops.

In summary, the aim of automatizing the common startup with the cooling is the following issues:

- minimize the number of steps in the cooling (reduces time for bringing up the system, by avoiding delay between consecutive PLC commands).
- minimize number of commands sent by the Pixel shifter (otheriwse the non-matching granularity will make this slow and error prone).
- minimize the time between starting a loop and powering the corresponding detector part (to stabilize the cooling loops).
- utilize predefined steps to switch on the loops.

In the implementation, the following issues have to be addressed:

[7)] Subsystems in this context refer to the Pixel Detector and the three SCT DAQ partitons

- Prevent an accidental switch on (The Pixel shifter must confirm the detector startup).
- Take into account different operational situations (beam/no beam, initial temperature).
- Give feedback to the shifter about success, or any required action.
- A DAQ action is required in the course of the procedure (configure modules).

Figure 7.8 shows a diagram depicting the resulting scheme implemented for operation.

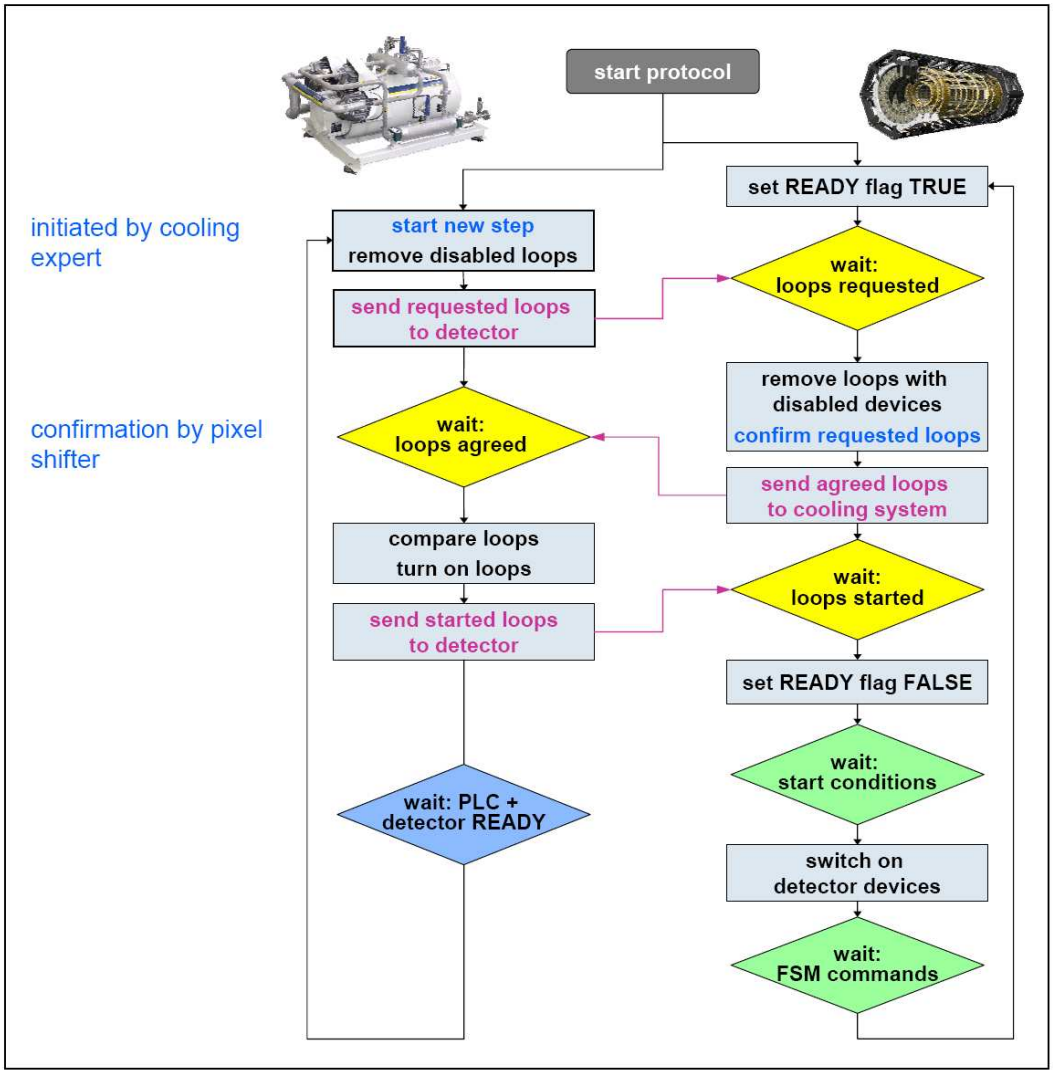

Figure 7.8: *Handshake procedure for the automatic cooling startup.*

The implemented protocol between SCT, Pixel and ID is:

- The detector DCS sets the 'DcsReady' Flag.
  This is done by the Pixel shifter for the initialisation of the protocol.
- ID sends the predefined list of 'requested loops' for the current step.
- The detector removes unavailable loops and sends a list of 'agreed loops'.
  For the Pixel Detector, unavailable loops are PCCs which are excluded from the FSM, or are not yet through

the *Startup* procedure. Opto loops will always be agreed to start and are the first step in the procedure. The target state for the detector after the first step is therefore OPTO_ON. In the following steps, PCCs which are not in OPTO_ON are not considered for the automatic start. They have to be addressed manually by the shifter after addressing the reason for the failure to reach the target state. If the FSM is not started, the requested loops will also be rejected, as it is not foreseen to operate the Pixel Detector without the FSM. The last reason to reject the start of a loop is the case that the autostart is globally disabled. On the ID side, the procedure can handle asynchronous management of Pixel Detector and SCT, i.e. it is not necessary for both subdetectors to execute the same step at a time.

- Optionally a confirmation from the Pixel shifter can be required, before sending the agreed list to ID.
- ID checks that all loops in the agreed list were actually in the requested list, and sends the corresponding command to the PLC.
  ID reports these loops as 'started' to the detectors, once the timeout counter for the PLC is zero. This should lead to a delay of less than a minute between reporting the loops as started and the actual start. All loops in one step are handled by one PLC command, so they are started at the same time.
- The sub-detector sets the DCSReady Flag to FALSE and powers up, as soon as detector side preconditions are fulfilled.
  For the Pixel Detector that would be the actual loop DPE reporting 'ON' (the cooling script would veto the command otherwise, see 7.1.2) and a temperature below the VETO threshold of the software interlock. In case of SCT, the requirement is a sufficient decrease of temperature.
- When all detector units are in the final state, or a timeout has expired, the detector sets DCSReady to TRUE. For the Pixel Detector the target state is ON or LV_ON for module loops, as long as no protocol with the DAQ is implemented. For opto loops the target state of the corresponding readout units is OPTO_ON.
- ID begins the next step of the protocol. For the Pixel Detector there are up to 12 steps available, while the default procedure involves 5 steps, 1 for all optoloops and one for each quadrant.

The automated startup procedure is implemented in a standalone script running on the subdetector control station, as it is necessary to have simultanous communication to the ID-project as well as to all FSM projects. As this protocol is addressing operational procedures, it can be assumed that the DCS (and DAQ) infrastructure is fully available. The underlying script can therefore use high level FSM functions, and run on a central machine with distributed connections to the involved systems.

7.2.1.1 Conclusion

The automated startup is in use since the restart of operation in 2010, and has proven to be very useful. It formalizes and automatizes the procedure to start up the cooling and the two silicon subdetectors in parallel and so reduces the risk of operational errors. It optimizes the recovery time by combining steps of the two subdetectors into one PLC command. The time for powering and configuring the Pixel modules associated to one step is overshadowed by the 5 min delay between PLC commands, so it is important to keep the number of PLC commands low for a fast recovery.
The interaction with the subdetector DAQ is still a manual action, however the shifter is alerted by the script at the appropriate time. This manual element in the procedure was acceptable, as it would represent a considerable effort on both DAQ and DCS side to automatize for rather small gain in operation.

Since the first implementation, additional conditions and checks for operation with beam have been implemented, providing a flexible use of the procedure when the beam conditions do not permit a full startup including high voltage.

7.2.2 Procedures for the Interaction with Beam Related Systems

In the nominal ATLAS data taking scenario, a data taking run is started well before stable beam conditions are reached, typically before beam injection. During this phase the Pixel Detector high voltage must be off for safety reasons. Fully configured, this would result in many noise hits, leading to an increased power consumption of the Pixel modules and a blocking of the ATLAS DAQ by the busy Pixel modules. It is therefore necessary to have the Pixel preamplifiers off while the sensors are not depleted. To be able to nevertheless include the Pixel Detector into the ATLAS DAQ from the beginning of the run, the so-called "warm start" was implemented by the ATLAS DAQ, allowing one to reconfigure the Pixel Detector during an ongoing run. Information from Pixel DCS is provided to initiate the configuration of the modules. At the end of a stable beam period the opposite procedure takes place, the so-called "warm stop".

7.2.2.1 Warm Start

The goal of the warm start automatic procedure is to minimize the delay for the detector to reach the READY state once stable beam conditions are established, and therefore maximize data taking time. Once LHC declares stable beam and the Pixel Detector high voltage is ramped up, the expert system of the ATLAS DAQ pauses the triggers, to allow the Pixel preamplifiers to be switched on. Then, the triggers are resumed and the warm start is completed. The run is then continued with trigger prescales for physics conditions.

In this procedure there are two steps where an automatization can be considered on the DCS side. The first is the ramping up of the high voltage. Currently, this is preceded by a careful expert-check of the beam conditions with respect to intensity, background levels, and collimator and absorber position, described in [104].

While the ramping up itself is a manual operation, a large part of the warm start has been automatized by the DAQ. To contribute to that effort, on the DCS side a flag is generated and published, once the conditions for switching on the preamplifiers are met. For each partition the *GOTO_READY* command triggers a check whether high voltage is ramped up successfully. A corresponding flag is publihed via DIM for the ATLAS DAQ, which uses it to trigger the warm start procedure.

In summary, the necessary steps upon reaching stable beams are:

- LHC: declaration of stable beam
- check of beam conditions (expert action)
- switch on HV (shifter action)
- signal to DAQ that HV is on (automatic)
- Warm Start (automatic)
  - hold triggers
  - configure modules (switch on preamps)
  - resume triggers
- Finalize Warm Start and change prescales (RC shifter action)

7.2.2.2 Outlook: Warm Stop

The goal of automated actions for the warm stop is also to maximize operation time with stable beam, and to prevent unsafe situations during the adjust handshake with the LHC. Automatizing actions at the end of a stable

beam period is intended to remove an arbitrary factor included in manual, human action, and should allow in the end a tuning to adjust delays to a common factor for all the experiments and subsystems. On ATLAS side, the DAQ action of the warm start is already automatized triggered by the WARNING state in the relevant LHC handshake. Accordingly, the DCS has to verify that the DAQ action was fully executed, for an automatic ramping down of the high voltage. Also, in case of a beam dump, an automatic rampdown could be triggered on the post mortem counter. The implementation requirements for automatizing this action are currently under discussion.

Chapter 8

Beam Background Correlations in Pixel, SCT and BCM Data

8.1 Beam Accidents

8.1.1 Scenarios for Accidents

Scenarios for beam accidents can be grouped in two categories: fast losses, single-turn losses, and losses with time constants ranging from multiple turns to seconds. While in the first case, there is no chance for the machine protection system to react, the slower losses should be picked up by the machine protection system.

8.1.1.1 Fast and Single Turn Losses

In case of injection or dump errors, the beam will normally be lost without completing the turn, while the beam dump system can only safely extract the beam in the 3 µs abort gap in every 89 µs orbit. An additional delay of up to 50 µs is caused by the time the dump decision needs to propagate to the Beam Dump System, and a third turn is used to extract all the bunches from the ring. The first possibility leading to the failure mode of very fast beam loss in less than one turn can occur during injection, in case of a misfiring of an injection kicker magnet or of a transverse mismatch during injection. In this case the beam will be immediately deflected on the nearest downstream aperture limiting equipment without the possibility for a clean extraction of the beam. In case of wrong magnet settings, the injected beam will also not be able to complete an orbit. For this reason, the first injected bunch into the empty accelerator has a low intensity of $0.5 \cdot 10^{10}$ protons, and it is not allowed to have any higher intensity as it might cause magnet-quenching.

Another possible mechanism for fast losses is a random misfiring of a beam-dump kicker magnet. Even though the other 14 magnets will be re-triggered and fire within a short amount of time, it is very unlikely for the initial firing to have happened during the abort gap. This case is similar to a regular dump, with all 15 magnets triggered, but happening asynchronously to the abort gap. The only solution in these cases is the protection function of absorbers and collimators.

8.1.1.2 Multi-Turn Losses

Once the beam has successfully completed the orbit, fast losses should only be possible in case of misfiring of the beam-dump kicker magnets. Other possible failure sources include trips of magnet power converters, wrong positioning of movable equipment, magnet quenches, vacuum leaks, or problems with the RF system. The quenching of a magnet is directly monitored and would lead to an immediate beam abort, while the time constant for the

quenching lies in the order of 10 ms. Also the other failure modes should be picked up by the beam protection system, leading to a beam dump as soon as the acceptable beam loss level is passed.

8.1.2 The Situation for the Pixel Detector

Effects of accidental beam losses for Pixel modules were studied by directly exposing production modules to high instantaneous radiation doses of up to $3 \cdot 10^{10}$ protons/cm^2 in high intensity beam at the CERN PS facility [102]. It was found that the module can survive with minimal or no influence on performance except for the effect of configuration loss at the highest instantaneous dose and for the expected effects by the total dose. However, other damage mechanisms have not been excluded. It has been suggested that the creation of highly localized charge densities created near the gates of a transistor might destroy it, however no tangible signs were observed during the tests.

The exposure of CMS modules to higher charge density by means of lasers did show damage to the front-end electronics with an assumed damage mechanism of shorting the high voltage to the front-end electronics. For this reason, the safety of the Pixel modules in beam loss scenarios can not be taken as fully granted. However the possible scenarios during which a very high instantaneous dose is delivered to Pixel modules is extremely rare at LHC and related to multiple machine failures that would also affect LHC magnets in a catastrophic way.

The Pixel Detector is located far from the injection and the dump region in the so-called Insertion Region 1, and so it is in the safest position with respect to single turn losses. Figure 8.1 shows the layout of this region.

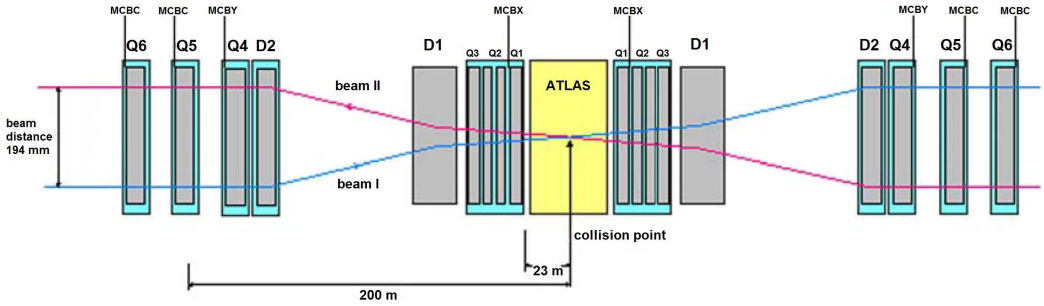

Figure 8.1: *Insertion Region 1 [103].*

Beam loss scenarios in the ATLAS region are discussed in [103]. Special focus lies on the settings of D1, D2 and MCBX magnets, as those have a large bending power with a fast response (since they are warm magnets) and may result in a scraping of the TAS (Target Absorber Secondary) absorber at injection. During injection it is therefore mandatory for the Pixel Detector to switch off high voltage, in particular as there is no reason to operate the detector at that time.

8.2 Background During Operations

8.2.1 Beam Degradation Mechanisms

Even during regular operation, the intensity of the circulating beams will steadily decrease. Among others, the following phenomena contribute to the continuous loss of protons from the beam [13]:

- intra-beam scattering of protons belonging to the same bunch may lead to emittance growth and finally protons escaping from the beam.
- scattering of protons with residual gas molecules in the beam pipe vacuum (beam gas).

- elastically scattered protons from collisions can end up in the halo.
- unavoidable losses due to regular operations like orbit correction, tune optimization, or chromaticity change.
- collisions.

8.2.2 Beam Halo and Beam Gas as Background in ATLAS

The main sources of beam background in ATLAS are beam halo and beam gas. Beam gas originates from the interaction of beam particles with residual gas, mostly hydrogen, in the beam pipe near the experiments. On the other hand, beam halo results e.g. from protons scattered at the cleaning collimators in IR3, potentially circulating for many orbits before eventually being lost, or tertiary halo scattering at the TCT.

Due to its time of flight measurement capabilities, BCM can distinguish events resulting from background, passing one station 12.5 ns after the other, and collision events originating from the interaction point, giving a detector response at the same time in both stations. Due to the low acceptance, only a relatively small fraction of recorded events however contain at least one hit on each side, which are necessary to determine a time of flight difference.

Statistically, the ratio of collision to background events should hold information about beam conditions.

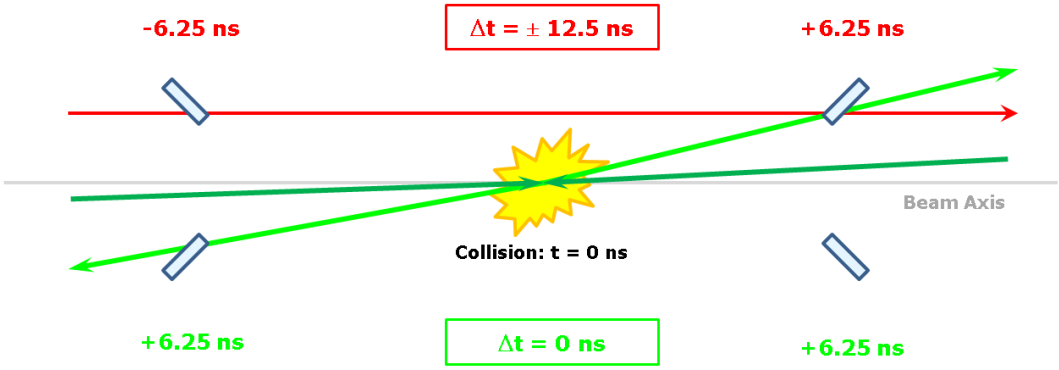

Figure 8.2: *Working principle of the BCM time of flight measurement for the classification of background and collision like event. In time only hits, respectively in- and out-of-time hits, are corresponding to collision- and background-like events (Not to scale).*

8.2.3 Measurements

During the initial phase of LHC operation with beam energies of 450 GeV and 3500 GeV and few bunches circulating in the machine, data for BCM, Pixel and SCT were studied.

8.2.3.1 BCM Data Structure

Before introducing the measurements, a short overview over the BCM measurement principle and data structure is presented.

Due to its excellent time resolution capabilities, the ATLAS Beam Conditions Monitor provides the possibility to distinguish collision and background-like events. For this, time of flight difference between detector modules on the two sides of the interaction point is used. The underlying principle is illustrated in figure 8.2. Events with "out-of-time-hits", registered at a time $t_B = c \cdot 1.84$ m before the time t_0 of the nominal collision, can be attributed to background events originating e.g. from particles moving parallel to the beam forming the beam halo. These events are characterized by the fact that modules on opposite z-sides of the interaction point are separated

by 12.5 ns in time. Positive sign is attributed to beam 1 ("A to C"), i.e. side A is passed before side C, negative sign to beam 2. For collision events on the other hand, modules will register "in-time-hits" on both sides at the same time t_C, 6.25 ns after the collision, resulting in a time difference of 0.

Hit Distribution inside the default Readout Window

The BCM can record a total of 31 bunch-crossings or Level-1 accepts (L1A) around the triggered bunch-crossing. With nominal timing setting, hits with BCM L1A of 18 are in the triggered bunch. As described in chapter 4, the Pulse Position encodes the time the hit occurred, with a resolution of 390 ps per bin. The pulse width is proportional to the Time over Threshold. The shown pulse position is "centered" at zero for the nominal collision bunch crossing, with pulse position sampling units of 390 ps, i.e. 64 pulse position "time" bins per bunch crossing. In the following, this time bin with L1A = 18 will be referred to as the "central" bin with respect to the triggered event, even though the readout window is not symmetric around this L1A. The timing is tuned in a way that in-time hits are in the pulse position window of 36-50 [390 ps], while out of time hits fall into the window of 4-18 [390 ps].

To understand how the recorded BCM data should be utilized, for one run the distribution of all recorded hits inside the full readout window was analyzed. In the given run, the filling scheme has 13 bunches per beam, with 8 colliding bunches in ATLAS:

Filled bunches beam 1: 1, 101, 201, 301, 601, 895, 995, 1095, 1195, 1786, 1886, 1986, 2086
Filled bunches beam 2: 1, 101, 201, 301, 501, 892, 992, 1092, 1192, 1786, 1886, 1986, 2086

The bunches are categorized into "paired bunches", i.e. bunches that are filled in both beams, and "unpaired bunches", i.e. bunches that are filled only for one beam. Unpaired bunches are further divided into "isolated" and "non-isolated". For the given filling scheme, bunches 501 and 601 are "unpaired isolated", as there is no other filled bunch near, and bunches 892/895, 992/995, 1095/1092 and 1195/1192 are "unpaired non-isolated", as there is a filled bunch only displaced by three BCIDs in the other beam.

Figure 8.3 shows the distribution of number of recorded BCM hits for each pulse position for the given run.

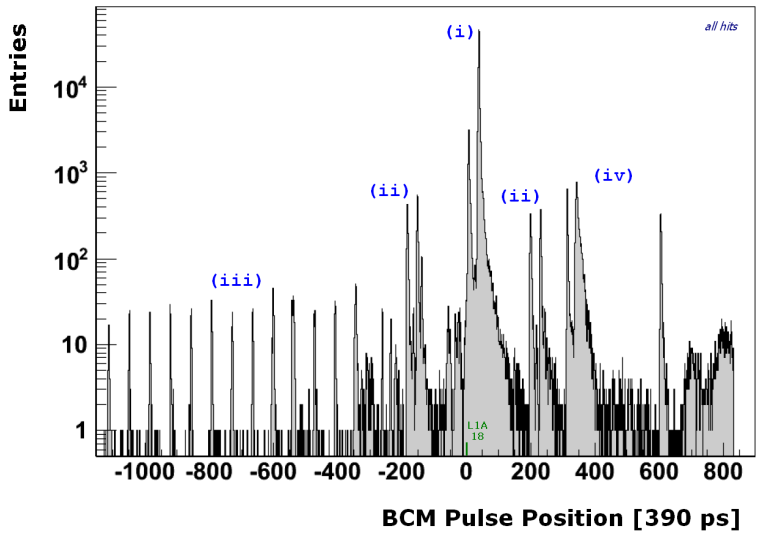

Figure 8.3: *BCM hit distribution for one run in the maximum BCM readout window of 31 bunch crossings for each triggered event.*

The readout window can be divided according to four categories:

1. The central bin contains the highest number of hits, and has a large excess of in-time-hits vs out-of-time hits.
2. The bins that are three BCIDs away from the nominal collision time contain a significant number of hits, with the number of in-time and out-of-time hits in the same order of magnitude.
3. The first 13 BCIDs have a few hits in the in-time-window.
4. significant number of hits are present in some higher BCIDs.

These can be interpreted as follows:

The central bin contains the highest number of hits, as this is the event where the nominal collision is triggered. As collision events have only in-time-hits, the excess of in-time-hits indicates collisions. The presence of out-of-time-hits indicates background events, although at a lower rate.

To understand the hits registered in region (ii), the BCIDs not originating from paired bunches were studied further, as shown in figure 8.4. The three central bins contain in-time and out-of-time hits in a comparable number, as expected for background events.

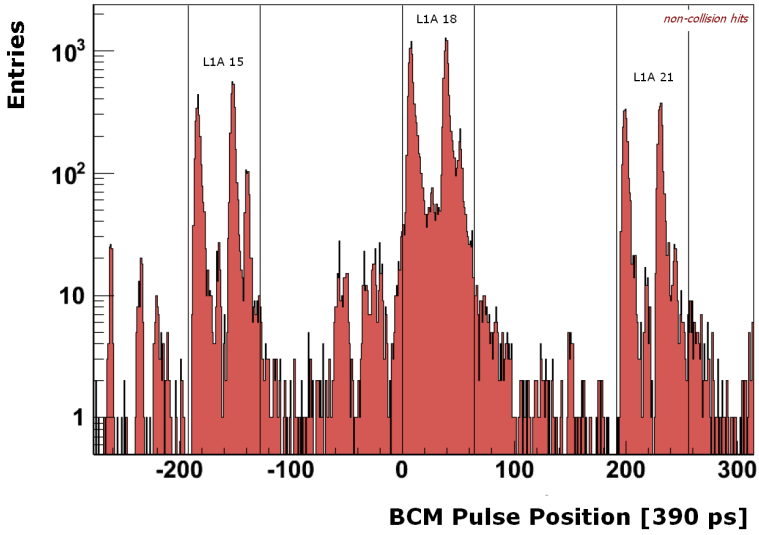

Figure 8.4: *BCM Hits for unpaired and unfilled bunches around the nominal L1A of 18. The filling scheme has unpaired bunches displaced by 3 BCIDs.*

The origin of the contributions to the three bins with L1A = 15, 18, 21 can be illustrated by separating the entries according to triggered BCID, which is shown in figures 8.5 for the non-isolated bunches. For the unpaired-non-isolated BCIDs of beam 1, one sees the expected background events in the triggered time-bin, and additionally a distribution of background related hits in the time-bin three BCIDs "early", at the time when a filled bunch from beam 2 passes ATLAS. Analogously, for beam 2, there are the triggered hits in the central bin, and additional recorded background events three BCIDs "late", originating from the passing of a filled beam 1-bunch, resulting in the same distribution, but shifted by three bunch crossings.
The number of recorded events in the earlier one of the unpaired non-isolated bunches is slightly higher than for the second bunch. This might be explained with the simple trigger dead time of 5 bunch crossings, which means that a second trigger only three bunch crossings after a previous one is not allowed.

For the isolated bunches 501 and 601, hits are visible only in the triggered bunch, as shown in figure 8.6.

If a trigger arrives some time after a collision event, the hits from the real collision event will appear in the recorded data earlier by the same amount of time. A ringing of the MBTS after high pulses triggers events after the nominal collision bunch crossing, which can be seen as recorded BCM hits for L1As between 0 and 12. If the collision

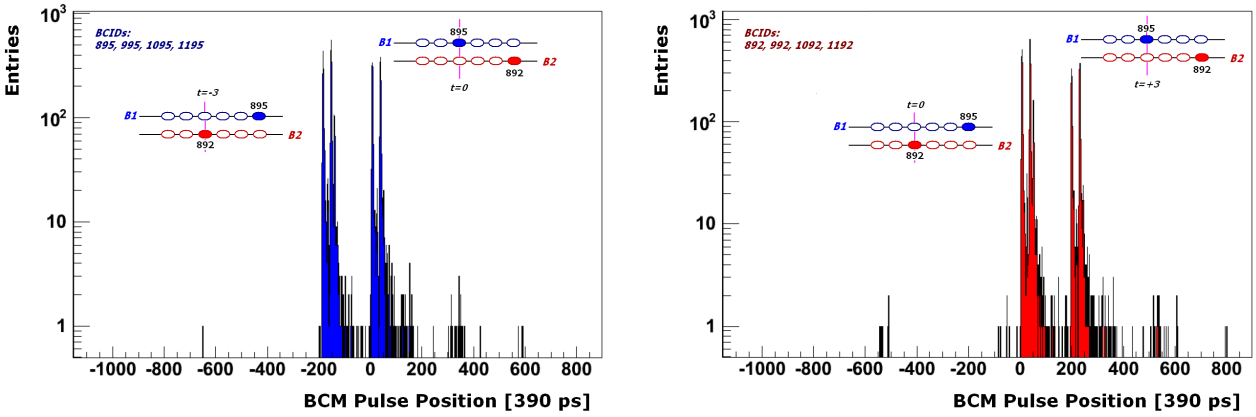

(a) (b)

Figure 8.5: *Unpaired non-isolated bunches, (a) in beam 1 and (b) in beam 2.*

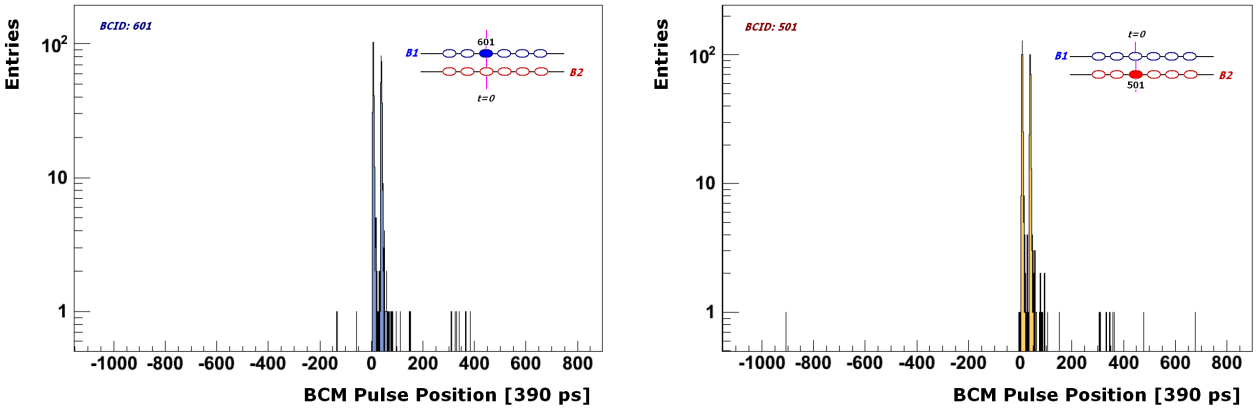

(a) (b)

Figure 8.6: *Unpaired isolated bunches, (a) in beam 1 and (b) in beam 2.*

event is triggered, the next possible trigger can only occur 6 BCIDs later, due to trigger dead time. The number of hits per triggered event of these type is, for the given run, of the order of 10, which explains why in-time hits are recorded, while almost no out of time hits are seen, as this is consistent with the in-time to out-of-time ratio as it is seen for the nominal L1A.

In addition to these physical events, hardware effects were found by the BCM collaboration and addressed during the winter shutdown of 2010/2011.

In the following, only hits recorded in the nominal filled bunch with L1A = 18 are considered.

Pulse Position and Pulse Width

The pulse position of a hit is equivalent to the time it was registered. The pulse width is correlated with the deposited charge in the sensor. Figure 8.7 (a) shows a histogram of pulse position vs pulse width for all BCM hits recorded in one run. Pulse position of 1 is an overflow bin. The majority of hits is expected to be inside the in-time or out-of-time window, and a deviation from this can be used to study the time-walk effect. Time-walk is caused by the fact that for a constant peak position of time over threshold signals, signals with a lower peak value pass the threshold later than a signal generated at the same time with a higher peak value would do.

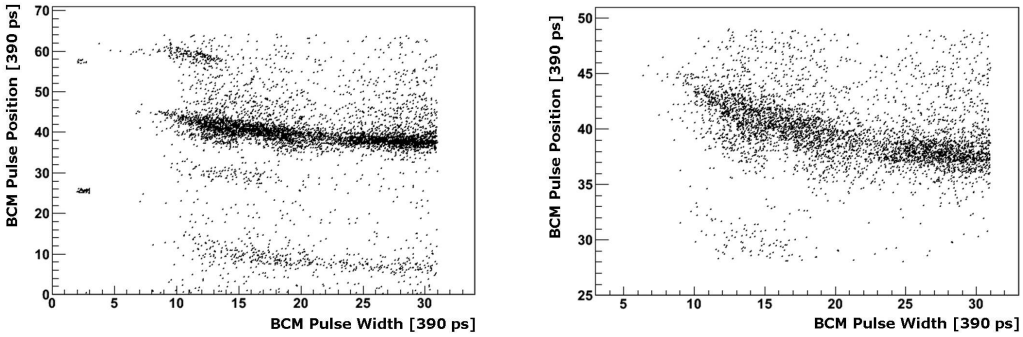

(a) (b)

Figure 8.7: *BCM Pulse Position vs. Pulse Width for (a) all hits, and (b) in-time-hits.*

Besides the effect explainable by time-walk, a significant number of hits fall outside the in-time value with a pulse position higher than $50 \cdot 390$ ps. There is no obvious dependence from the pulse width for these hits, making it unlikely to originate from time-walk.

Figure 8.7 (b) shows only hits in the in-time-window, so the effect of the time-walk can be seen more clearly. The effect is more pronounced for signals with a low deposited charge, while it decreases for higher pulse widths.

Time Difference

Two algorithms to determine the Δt are considered. The main algorithm checks for hits on A- and C-side, and determines the time difference as the pulse position difference between single hits on each side. Hits from different type of channels (i.e. low gain vs. high gain) can not be compared as the difference in pulse shape would distort the resulting value for the time difference (time walk). Therefore, events can have no defined time difference if there are BCM hits only on one side (or none at all), or a number of time difference values related to the number of hits on each side: $n_{\Delta t} = n_A \cdot n_C$. For the sign of the time difference, the convention is that background from beam 1 ("A to C") has a positive time difference, i.e. hits in the A-side modules occur earlier than in the C-side modules: $\Delta t = (P_C - P_A) \cdot (0.390\, ns)$, where P is the dimensionless pulse position in bins of 390 ps. Background from beam 2 has, by definition, a negative time difference. The aligned time difference emphasizes the separation between collision and background, as the impact of ambiguous information is decreased. This is achieved by taking into account only hits from A- and C-side modules with the same x- and y- coordinate. The downside of this approach is that there are less events where the the aligned time difference has a defined value. A typical distribution of time difference values for both algorithms is shown in figure 8.8.

Background from beam 2 (seen at $-32 \cdot 390$ ps) is slightly higher than background from beam 1, which is expected from the accelerator properties and compatible with results from other sub-detectors.

Figure 8.9 shows the time difference per event according to the event BCID. The filling scheme is the introduced filling scheme with 13 bunches per beam, and 8 colliding bunches. Non-colliding non-isolated bunches were vetoed here, and no BCM hits are recorded.

The resulting distribution of time difference per BCID shows accordingly collision-like distributions for the paired bunches at BCIDs 1, 101, 201, 301, 1786, 1886, 1986, and 2086. For BCID 501, a distribution with peak at $-32 \cdot$ 390ps is obtained, and at $+32 \cdot 390$ ps for BCID 601, which is consistent with background events from beam 2 or beam 1. No events are recorded for bunches 895/892, 995/992, 1095/1092, 1195/1092, as triggers are vetoed.

Figure 8.10 shows the obtained time difference values depending on the pulse position of the individual hits, and an overview classifying the different types of events is shown in figure 8.11.

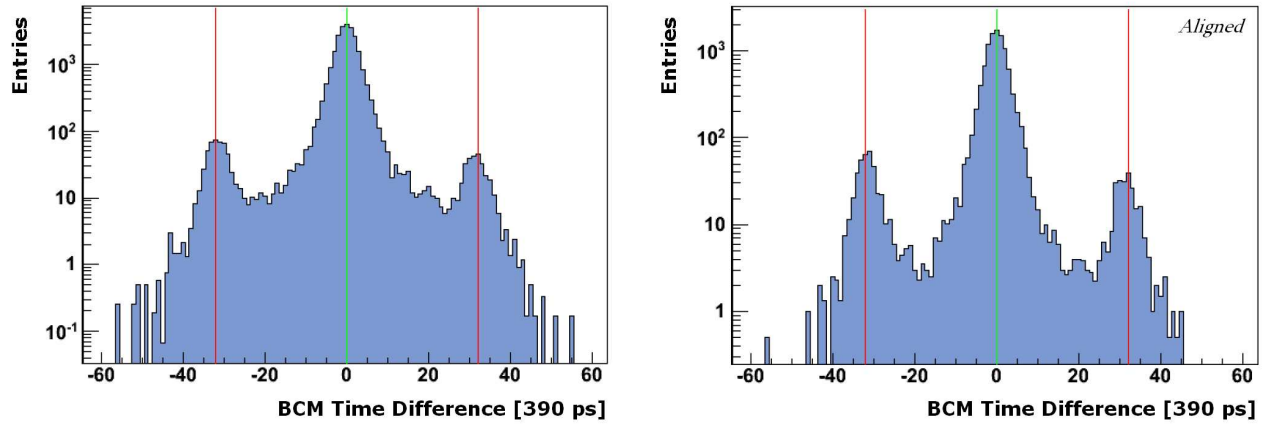

(a) (b)

Figure 8.8: *Δt distribution for two different algorithms. (a) considering all permutations. (b) only taking into account Δt between hits in aligned modules. Separation between collision and background is more pronounced in the latter.*

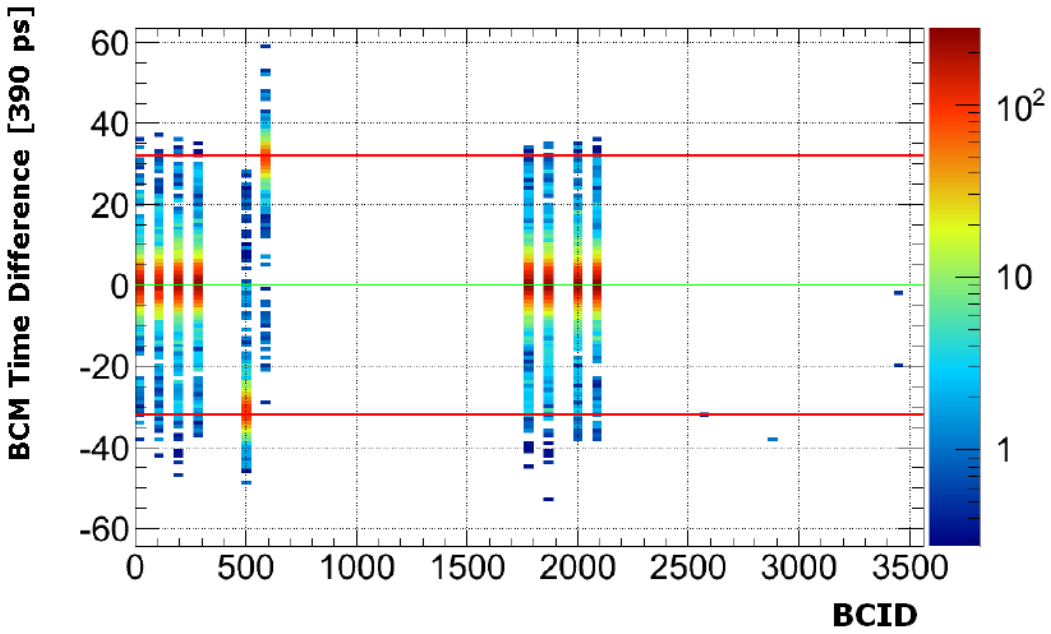

Figure 8.9: *Δt vs. BCID. Paired and unpaired bunches can be distinguished, and beam1/beam2. Trigger effect can be seen by veto for filled non-isolated bunches in BCID range between 800 and 1200 (exact BCIDs given in text).*

- Collision events cause two histogram entries at $\Delta t = 0$ and in the in-time-window (green circle in figure 8.11).

- Background from beam 1 causes two entries at $\Delta t = +12.5$ ns, with an A-side hit in the out-of-time window and a C-side hit in the in-time window (blue circle).

- Background from beam 2 causes two entries at $\Delta t = -12.5$ ns, with an C-side hit in the out-of-time window and an A-side hit in the in-time window (red circle).

In addition to these well-defined time differences, hits can be delayed, resulting in time differences outside the well-defined bands.

- Collision events with a late in-time hit on the C-side give a positive Δt between 0 and + 12.5 ns. Background events with a late out-time hit on the C-side give Δt between + 12.5 ns and 0 (light blue).
- Collision events with a late in-time hit on the A-side give a negative Δt between 0 and - 12.5 ns. Background events with a late out-of-time hit on the A-side give Δt between - 12.5 ns and 0. (orange).
- Background events from beam 1 (beam 2) with a late in-time-hit on the C-side (A-side) give a positive (negative) Δt with the absolute value higher than 12.5 ns (dark blue/dark red).

Late hits can be due to time-walk, i.e. hits with a low charge deposition registering late, while de facto the module hits occurred with a physical time difference compatible with 0 or $\pm$12.5 ns. Another explanation for the occurrence of late hits would be that they do not originate from the primary collision, but from afterglow still present in the experiment for a time.

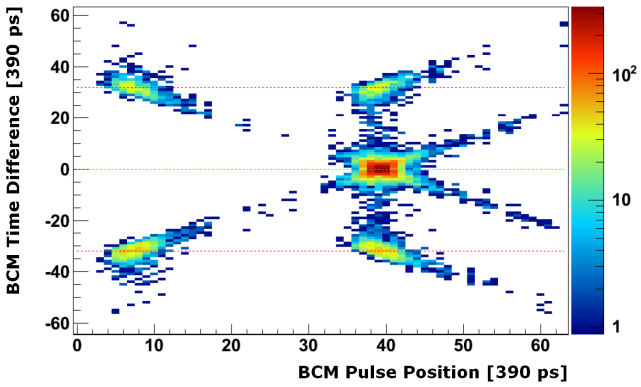

Figure 8.10: *Dependence of Time Difference and Pulse Position.*

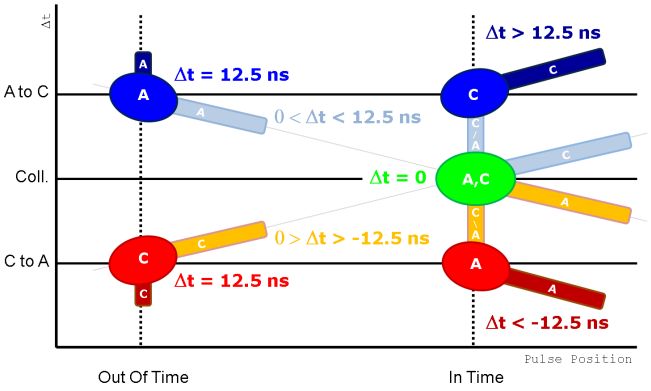

Figure 8.11: *Schematic overview of dependence of Time Difference and Pulse Position.*

Beam Optics Dependence

The behaviour of the BCM time difference measurement for different LHC beam optics was studied, in particular the β^* dependence. β^*, i.e. the β-function at the interaction point, is a measure of the focusing of the beam. The

process of reducing β^* is called "squeezing". The time difference distribution for $\beta^* = 11$ m and $\beta^* = 2$ m was measured. Figure 8.12 illustrates how the beam profile evolves with changing beam optics [105]. With reduced β^*, the beam size becomes smaller at the interaction point, but larger in the superconducting magnet triplet, requiring a setting of the TCT collimator that ensures protection of the triplet. Evidently, the primary goal of minimizing β^* is increasing the luminosity at the interaction point according to eq. 2.1.

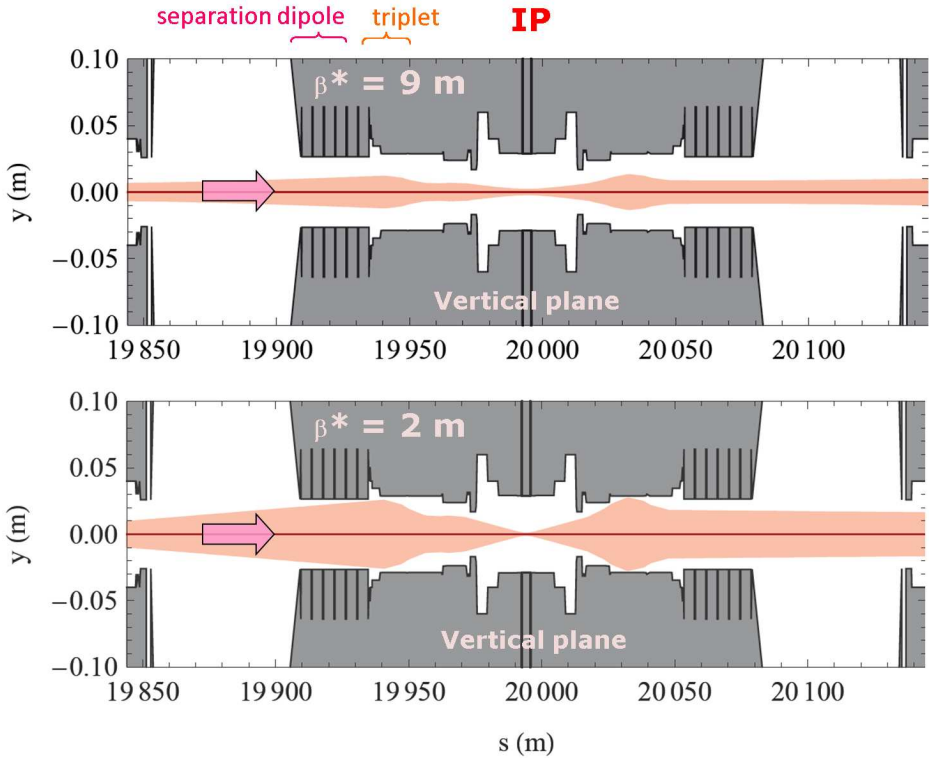

Figure 8.12: *Beam Profile at IR1 for different beam optics. Top: $\beta^* = 9$ m, bottom: $\beta^* = 2$ m [105].*

For the measurement of the time difference, only events during stable beams from paired BCIDs were used, to ensure comparability between different runs. Only hits with L1A consistent with the nominal collision were taken into account, as the focus of this study is the ratio of collision events to beam background, i.e. halo events. Figure 8.13 shows the time difference distributions for run 152409 with β^* of 11 m, i.e. un-squeezed beam, and for run 154822 with β^* of 2 m, i.e. squeezed beam.

The green curve shows all possible Δt values from all recorded events. In order not to increase artificially the background by considering an amount of events, varying by run, while the beams are not yet colliding, only events recorded during stable beam periods are considered (yellow curve). As the unpaired bunches only contribute to background and not to physics events, events from unfilled or unpaired bunches are disregarded (orange curve). The light blue curve shows the time difference distribution for all Δt values from paired bunches during stable beams. The final distribution (dark blue) is obtained by correcting for the number of hits per event, as otherwise events with higher number of hits would be over-weighted than events with few hits. The final correction has a high effect for almost completely removing the Δt peaks corresponding to background, indicating that for squeezed beam the ratio of number of recorded background events vs collision events is smaller, while indicating to background events with a higher number of hits per event.

To understand the behaviour with respect to the beam optics the following points can be considered. While the beam profile for squeezed beam with the larger size away from the interaction point might lead to the expectation of a higher number of halo like background, the collimator settings, especially for the TCT, have to be taken into account. Also a higher luminosity for squeezed beam would favor collision like events.

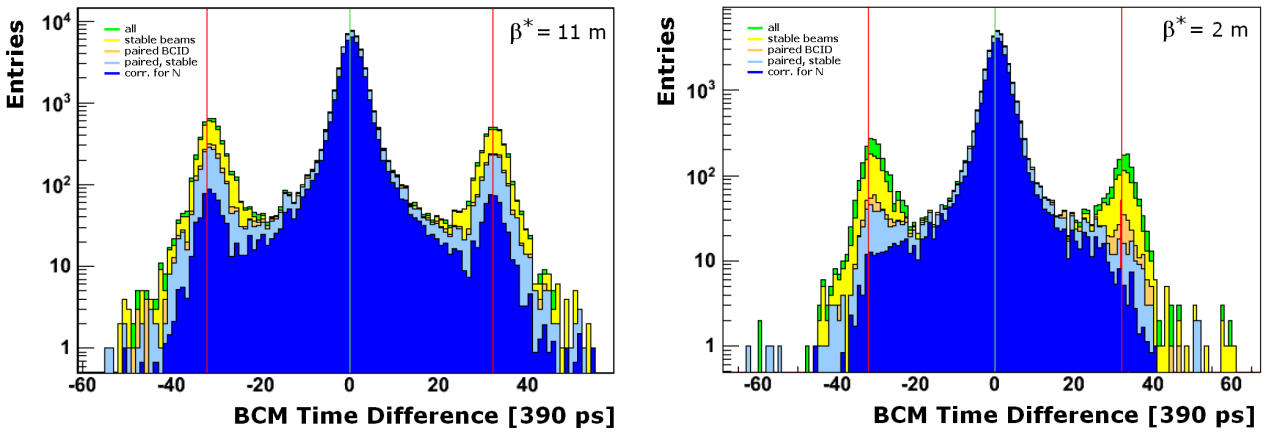

(a) (b)

Figure 8.13: *Δt for different beam optics, for (a) un-squeezed and (b) squeezed beams. The corrected (dark blue) distribution takes into account the correction factor for the number of hits per event. The ratio of collision to background-like events is higher for squeezed beam.*

Correlations

The correlation of hit multiplicity vs. Δt distribution was studied, to understand how the time difference distribution algorithm is affected by the number of hits. A simple averaging of the time difference values calculated for each event would obviously take care of the issue of unevenly weighted events. However, it is evident that there are events where this fails to correctly depict the real processes, e.g. events recording collision data, and in addition hits due to beam halo would give a time difference value between collision and background window. Therefore the BCM hit multiplicity was analyzed with respect to the resulting time difference, to verify if there is a bias in treating events with different number of hits with the same weight. The result is shown in figure 8.14. Only events with L1A = 18 during stable beam were taken into account. A correlation between number of hits and time difference can be clearly seen.

After establishing the dependence of hit multiplicity and time difference, for the time difference calculation algorithm it is clearly necessary to take this into account, consistently with the results seen in figure 8.13.

This dependence was studied also for the Pixel Detector. Figure 8.15 shows the corresponding distribution. Also here a clear correlation can be seen for events flagged as background having higher number of pixel hits. In collision events, for the barrel region, tracks are roughly perpendicular to the modules, giving a low number of hits per track. The beam halo parallel to the pixel modules leads to high number of hits per particle compared to collision events, as particle tracks parallel to the pixel surface give a high number of hits per track, as illustrated in figure 8.16.

Finally, it was studied how the hit multiplicity behaves in dependence of the radial distance from the beam axis. For this, it was studied how the two innermost sub-detectors, the Pixel Detector and SCT behave for different events, classified as indicated by the BCM time difference. The result is presented in figure 8.17. The correlation between number of Pixel hits and number of SCT hits was studied per event. In Figure 8.17 (a), all events are considered without any further restriction. Two classes of events can be distinguished. The majority of entries has a well defined ratio of Pixel hits vs. SCT hits consistent with 0.375. A line with that slope is introduced to guide the eye and help comparison between the different subplots. Another category of events has a higher number of Pixel hits with respect to the number of SCT hits. These events are less frequent, and the ratio-distribution is not so clearly defined to a narrow line.

To further understand how these two classes are defined, the dependence from the BCM time distribution was studied in figure 8.17 (b). The remaining events, for which a BCM time difference can be calculated, are considered,

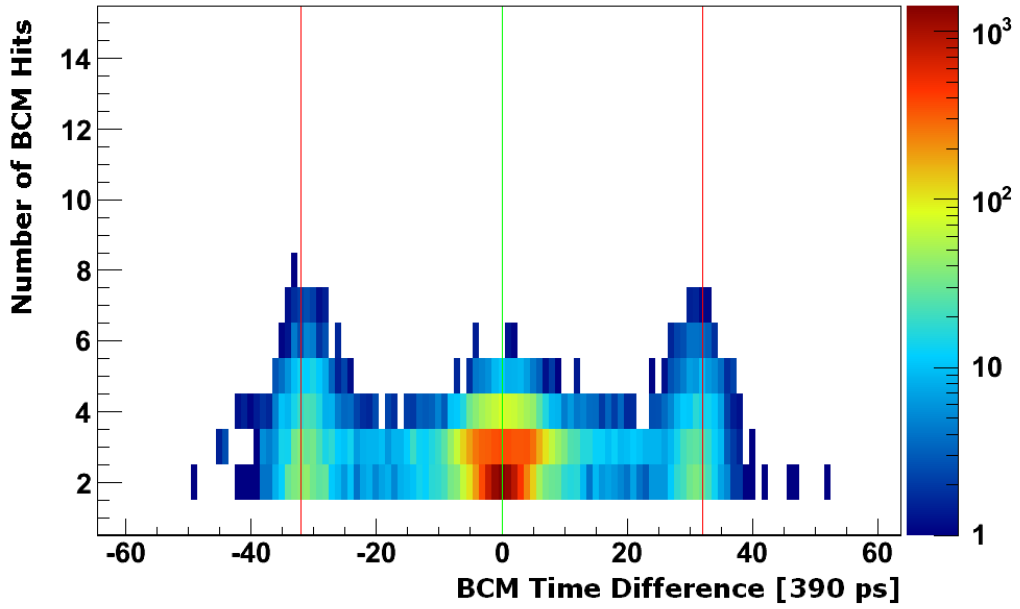

Figure 8.14: *Number of BCM hits vs. BCM Δt distribution. A clear correlation can be seen for events flagged as background having higher number of hits.*

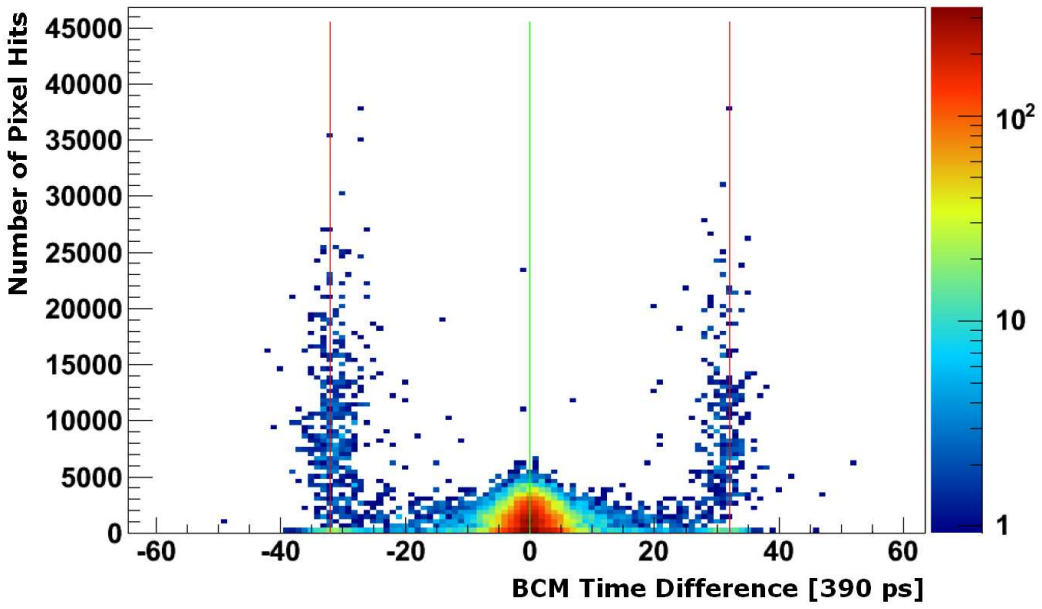

Figure 8.15: *Number of Pixel hits vs. BCM Δt. Clear correlation can be seen for events flagged as background having higher number of Pixel hits.*

leading to a lower total number of events, as the acceptance of the BCM is not so high to expect at least two hits in every triggered event. While the number of considered events decreases, the basic distribution does not change noticeably.

The situation changes significantly with a cut to events flagged by BCM as collision-like, as shown in fig-

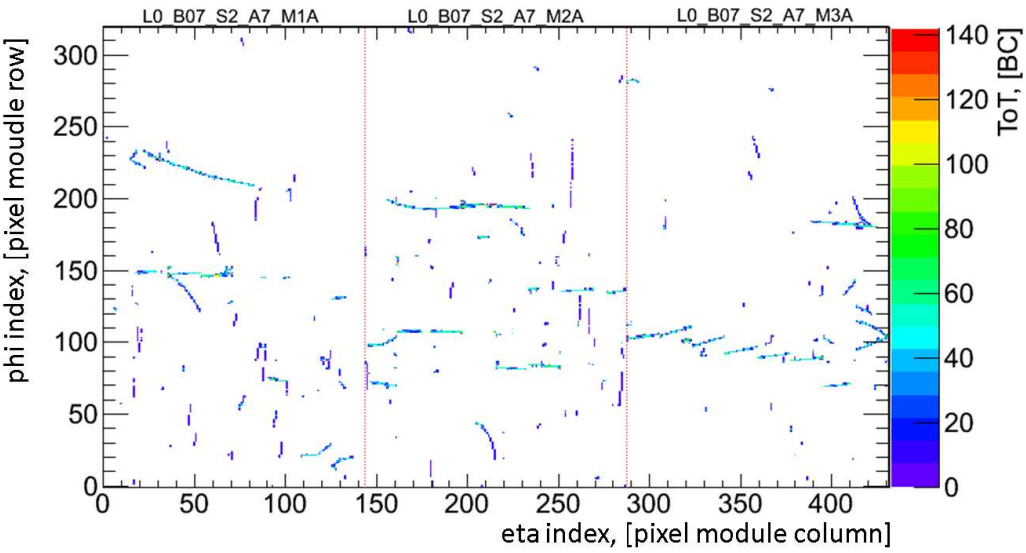

Figure 8.16: *Pixel hits given by their ToT in neighbouring modules for one event. Track-like hit distribution originating from beam halo can be observed.*

ure 8.17 (c). Almost exclusively events with the lower, narrower ratio distribution of Pixel hits vs. SCT hits remain.

For background-like events selected in figure 8.17 (d), the other class of events remains, with the higher ratio of Pixel hits vs. SCT hits, and with a wider distribution of the ratio. A low number of collision-like events can be seen also for Δt indicating background. This should be studied by BCID: a possible explanation would involve collision-like events with underlying BCM background.

A possible explanation for this clear correlation is that the ratio for collision event is defined by the geometry of the detector, i.e. almost any track will nominally pass each layer once. This means, e.g for the barrel region, that it is most likely to record three Pixel hits for every four SCT layers equivalent to 8 SCT hits. On the other hand, background events produce more hits nearer to the beam pipe, and in a less deterministic way. This would increase the ratio in favor of Pixel hits, while simultaneously account for the wider ratio distribution.

Studying this correlation of the two innermost sub-detectors together with the timing capabilities of the BCM, shows a good correlation of hit multiplicity ratio to the type of observed event, originating from collisions or beam halo.

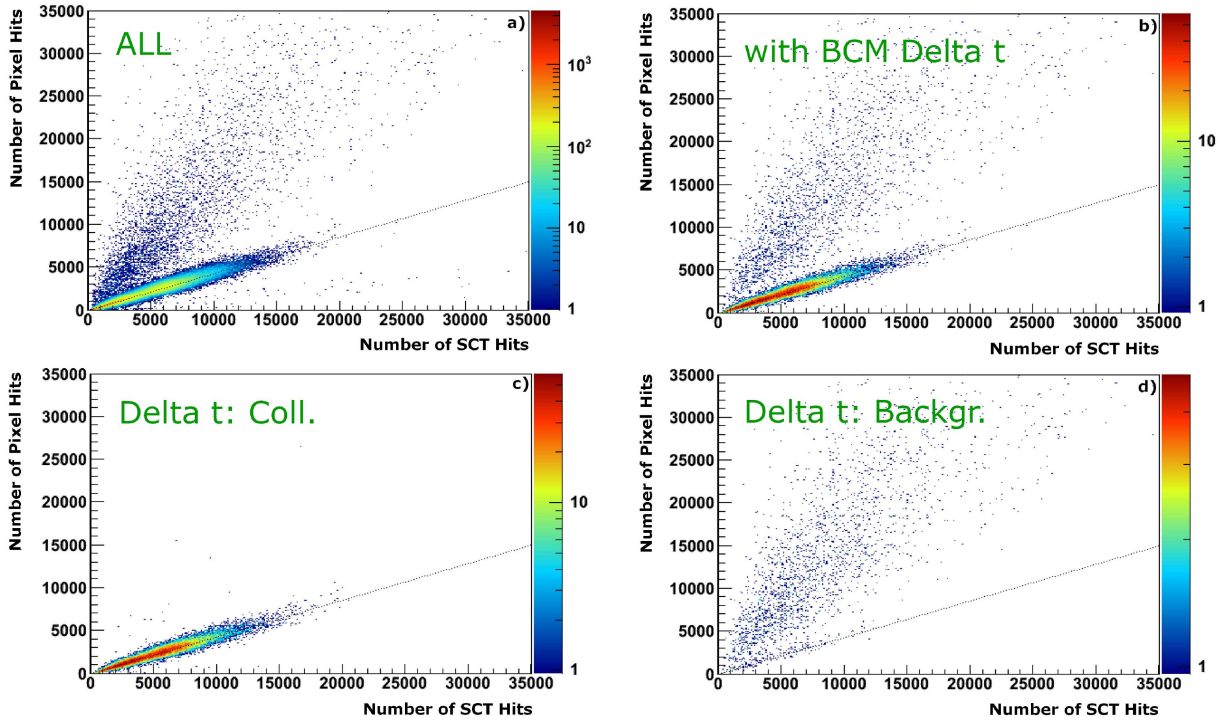

Figure 8.17: *Number of Pixel hits vs. number of SCT hits. (a): all events, no requirement for BCM hits. (b): enough BCM hits to calculate Δt. (c): Δt indicating collision events. (d): Δt indicating background events. Collision events are thought to produce a roughly fixed ratio of Pixel vs SCT Hits, determined by detector geometry. Background events produce more hits nearer the beam pipe in a less deterministic way.*

8.2.4 Conclusions

The situation for the Pixel Detector with respect to safety issues for beam loss scenarios was discussed. Beam background in ATLAS due to beam halo and beam gas was studied. For this, data recorded by the ATLAS Beam Conditions Monitor was analyzed. The measurement principle and the data structure of the BCM was presented. The recorded hit distribution was analyzed also in view of data used for further studies. The contributions to different regions of the readout window were examined, and the timing properties with respect to the deposited charge in the sensors. After this preparatory study for the understanding of the used data, the time difference Δt as a highly useful criterion for classification of events was established, and different algorithms for its calculation were presented. The distribution of this quantity was examined, and on the basis of its BCID dependence for a given run, the expected effectiveness for the classification of events with respect to their origin from collisions or background events from either beam was verified.

In particular the behaviour with respect to different beam optics settings was studied in more detail. It was seen that the focusing of the beam in the interaction point improves the ratio of collision to background events. Two mechanisms are considered for their influence on the ration of collision to background events. For the higher focused beam this is the widening of the beam outside the interaction point giving rise to more halo events, and the higher achieved luminosity, favoring collision events. This study gives indication that collision and background events have different properties with respect to BCM hit multiplicity. This was verified by regarding the number of BCM hits vs. the observed time difference. Higher number of BCM hits is seen for Δt consistent with background events. This correlation was also studied for the Pixel Detector, showing that events with an extremely high multiplicity are almost exclusively background events, while the majority of collision events produces a distribution with a lower number of hits. This was also illustrated by the hit distribution for a single event showing tracks

parallel to barrel module surface, as expected from beam halo.

The radial dependence of background and collision events was studied using the correlation of the hit multiplicity for the Pixel Detector and the SCT. Using the classification provided by the BCM Δt measurement, the observed two classes of events can be well separated, consistent with a geometrical interpretation of the expected hits for halo or collision events.

Finally the qualitative correlations between data of the ID silicon detectors and the Beam Conditions Monitor are clearly visible. Most of the results obtained in this field have been presented in several ATLAS meetings and are systematically described for the first time in this thesis.

Chapter 9

Conclusions and Outlook

The Pixel Detector Control System has a critical role in the safe operation of the detector. All the major milestones of Pixel operation: System Test in SR1, Service Test, Connectivity Test on the surface and in the pit, cosmic-ray data taking, and finally the first proton collisions with stable beam in December 2009 at 450 GeV and in March 2010 with the unprecedented energy of 3500 GeV – were relying on a reliably working DCS. In turn, the increasing understanding of operational procedures and conditions led to an evolution and refining of requirements and specifications.

In this time the DCS was evolving from small test systems operated by detector and DCS experts to the continuous operation of the production system by single non-expert shifters, and will in the future evolve to a single shifter controlling not only Pixel DCS but complete ATLAS. This process was only possible by the usage of a dedicated tool – the FSM – that provides an easily accessible overview of the detector state and allows to change between states with well defined procedures. The compilation of the requirements and translation of the operational model into the state model of the FSM was a main part of this thesis. Even though the FSM was in a constant optimization process, this did not have negative impact on operational availability and reliability of the system.

Complementary to the implementation of the Pixel FSM, an increasing demand and focus was put into additional safety and automatic procedures. A number of dedicated scripts, running in background, have been implemented to continuously check the state of the cooling system, and, in case of a failure, to switch off the relevant parts of the Pixel Detector. These scripts have been working very reliably, leaving the hardware interlock as a last line of defense that is rarely needed. Complex interactions with external systems have been automatized, reducing the danger of operational errors and resulting in a greater efficiency. This includes the common startup of Pixel Detector and Cooling System that was challenging to realize due to the different modularities of the systems. Additional scripts for the interaction with the LHC and ATLAS Data Acquisition have been put in place or are under development and finalization. They address beam related safety issues and the start-stop of data taking periods at stable beams.

Finally, beam background events have been studied for the first proton-proton collisions with data from the ATLAS Beam Conditions Monitor and the Inner Detector silicon detectors, in particular the Pixel Detector. It was seen that the time of flight difference measured by the Beam Conditions Monitor is a good criterion to distinguish between collision and background events and can serve as an indication for the quality of the beam conditions. Using this information, also correlations between Pixel and SCT data and beam conditions have been studied. Qualitative correlations between data of the Beam Conditions Monitor and SCT and Pixel Detector are clearly visible and documented in a systematic manner in this thesis.

Bibliography

[1] F. Halzen, A. Martin, Quarks and Leptons: An Introduction Course in Modern Particle Physics, Wiley (1984)

[2] D. Griffiths, Introduction to Elementary Particles, Wiley (2008)

[3] Particle Data Group, Review of Particle Physics, Physics Letters B 667 (2008)

[4] ATLAS Collaboration, Expected Performance of the ATLAS Experiment. Detector, Trigger and Physics, CERN-OPEN-2008-020

[5] ATLAS Collaboration, ATLAS high-level trigger, data-acquisition and controls : Technical Design Report, CERN/LHCC/2003-022

[6] U. Aglietti et al., Tevatron-for-LHC Report: Higgs, hep-ph/0612172

[7] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, R. Tanaka (Eds.), Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables, CERN-2011-002 (CERN, Geneva, 2011)

[8] Juan A Valls Ferrer, Higgs Discovery Potential with the ATLAS Detector at the LHC, Journal of Physics: Conference Series 171 (2009) 012026

[9] Lyndon Evans and Philip Bryant (editors), LHC Machine, 2008 JINST 3 S08001

[10] Reyes Alemany, Mike Lamont, Stephen Page, Functional specification: LHC Modes, LHC-OP-ES-0005, rev 1.0 (2009)

[11] T. Pieloni, A Study of Beam-Beam Effects in Hadron Colliders with a Large Number of Bunches, CERN-THESIS-2010-056

[12] W. Herr et al., LHC bunch Filling schemes for commissioning and initial luminosity optimization, CERN-LHC-Project-Note-415, 2008

[13] C. Bracco, Commissioning Scenario and Tests for the LHC Collimation System, CERN-THESIS-2009-031

[14] V. Kain, Machine Protection and Beam Quality during the LHC Injection Process, CERN-THESIS-2005-047

[15] G. Robert-Demolaize, Design and Performance Optimization of the LHC Collimation System., CERN-THESIS-2006-069

[16] G. Guaglio, Reliability of the Beam Loss Monitors System for the Large Hadron Collider at CERN., CERN-THESIS-2006-012

[17] B. Holzer et al., Beam Loss Monitoring System for the LHC., CERN-AB-2006-009

[18] Bernard Jeanneret, Helmut Burkhardt, On The Measurement of the Beam Losses in the LHC Rings., EDMS: LHC-BLM-ES-0001

[19] R. Schmidt, Fast Magnet Current Change Monitors., EDMS: LHC-CI-EC-0001 (628602)

[20] ATLAS public web site, http://www.atlas.ch

[21] G. Aad et al., The ATLAS experiment at the CERN Large Hadron Collider, JINST 3 (2008) S08003

[22] L. Bouvet et al., ATLAS Inner Detector Services: Parameters describing the orientation naming and layout of Inner Detector Services, EDMS: ATL-IC-EP-0017

[23] ATLAS Collaboration, Liquid Argon Calorimeter Technical Design Report, CERN LHCC 96-41

[24] ATLAS Tile Calorimeter Collaboration, Tile Calorimeter Technical Design Report, CERN LHCC 96-42

[25] ATLAS Muon Collaboration, ATLAS Muon Spectrometer Technical Design Report, CERN LHCC 97-22

[26] ATLAS Trigger Performance Group, ATLAS Trigger Performance Status Report, CERN-LHCC-98-015

[27] A. Barriuso Poy et al., The Detector control system of the ATLAS experiment, 2008 JINST 3 P05006

[28] A. Barriuso Poy et al., ATLAS DCS Integration Guideline, ATL-DQ-ON-0013

[29] S. Lüders et al., The CERN Detector Safety System for the LHC experiments, proceedings of ICALEPCS 03, Gyeongju, Korea, 2003

[30] S. Franz et al., The Detector Safety System of the ATLAS experiment, 2009 JINST 4 P09012

[31] ATLAS Pixel Detector Community, ATLAS Pixel Detector Technical Design Report, CERN/LHCC/98-13

[32] M. Olcese et al., Barrel Layout Parameters, EDMS, ATL-IP-EP-0015

[33] Pixel Naming and Convention, EDMS, ATL-IP-MN-0004

[34] G. Aad et al., ATLAS Pixel Detector electronics and sensors, JINST 3 (2008) P07007

[35] W. R. Leo, Techniques for Nuclear and Particle Physics Experiments, Springer, 2nd Rev edition (February 1994)

[36] F. Lemeilleur, G. Lindström, S. Watts, Third RD48 status report, CERN-LHCC-2000-009

[37] D. Dobos, Commissioning Perspectives for the ATLAS Pixel Detector, PhD. Thesis, Universität Dortmund, September 2007

[38] J. Weingarten, System Test and Noise Performance Studies at The ATLAS Pixel Detector, PhD. Thesis, Universität Bonn, September 2007

[39] M. Olcese et al., Requirements for Pixel Detector Services, EDMS, ATL-IP-ES-0007

[40] M. Olcese et al., Pixel Detector Active Area Layout, EDMS, ATL-IP-EP-0004

[41] M. Citterio, Pixel High Voltage Cables, EDMS, ATC-TB-ER-0018

[42] M. Citterio, Pixel Electrical Signal Type II Cables, EDMS, ATC-TB-ER-0019

[43] M. Citterio, Pixel Electrical Signal Type III Cables, EDMS, ATC-TB-ER-0029

[44] M. Citterio, Pixel LV Type II Cables, EDMS, ATC-TB-ER-0024

[45] M. Citterio, Pixel LV Type III Cables, EDMS, ATC-TB-ER-0017

[46] M. Citterio, Pixel Type IV Cables: Electrical Services from PP3 to the counting room, EDMS, ATC-TB-ER-0020

[47] Service Panel Master Document, EDMS, ATL-IP-QA-0022

[48] D. Attree et al., The Evaporative Cooling System for the ATLAS Inner Detector, 2008 JINST 3 P07003

[49] L. Luisa, Detector Control System (DCS) for the Evaporative Cooling Plant for the ATLAS Inner Detector and Tools for the DCS Offline Monitoring., Tesi di Dottorato, Università degli Studi di Udine, 2009

[50] L. Luisa, Description of the Detector Control System for the Inner Detector evaporative cooling., ATL-IC-ER-0010

[51] H. Mann, H. Schiffelgen, R. Froriep, Einführung in die Regelungstechnik, Hanser, München (2000)

[52] J. Moss, Commissioning and Operation of the ATLAS Pixel Detector, ATL-INDET-PROC-2010-038

[53] V. Cindro et al., The ATLAS Beam Conditions Monitor, 2008 JINST 3 P02004

[54] A. Gorišek et al., Commissioning and first operational Experience of the ATLAS Beam Conditions and Loss Monitors based on pCVD Diamond Sensors, Nuclear Instruments and Methods in Physics Research A617 (2010) 444-448

[55] NINO web site (CERN Knowledge & Technology Transfer), http://technologytransfer.web.cern.ch/technologytransfer/en/Technology/NINO.html

[56] I. Dolenc, Development of Beam Conditions Monitor for the ATLAS Experiment, PhD. Thesis, University of Ljubljana, 2008

[57] O. Holme, M. Gonzalez Berges, P. Golonka and S. Schmeling, The JCOP framework, proceedings of ICALEPCS 05, Geneva, Switzerland, 2005

[58] B. Hallgren, Controller Area Network - Hardware Description, http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/DCS/CANINFO/canhard.htm

[59] CAN in Automation Homepage, http://www.can-cia.org

[60] CANopen Website (IXXAT), http://www.canopensolutions.com/english/about_canopen/about_canopen.shtml

[61] ATLAS DCS, Embedded Local Monitor Board ELMB128, Oct 2002, http://elmb.web.cern.ch/ELMB/elmb128_1.pdf

[62] H. Boterenbrood, CANopen Application Software for the ELMB128 v2.2, Jan 2005, http://www.nikhef.nl/pub/departments/ct/po/html/ELMB/ELMBresources.html

[63] PH-ESS group, CAN Power Supply – Homepage, Jan 2006, http://ess.web.cern.ch/ESS/canpsuProject/index.htm

[64] J. Cook, ELMB PSU Introduction, Sep 2006, ATLAS-DQ-ON-0025

[65] ETM, PVSS Online Help

[66] C. Gaspar, B. Franek, Tools for the automation of large distributed control systems, IEEE Trans. Nucl. Sci. 53 (2006) 974

[67] C. Gaspar, Hierarchical Controls, Configuration & Operation, http://lhcb-online.web.cern.ch/lhcb-online/ecs/fw/FSMConfig.pdf

[68] C. Gaspar, B. Franek, SMI++ - Object Oriented Framework for Designing and Implementing Distributed Control Systems, IEEE Trans. Nucl. Sci. Vol. 52, No. 4 (2005)

[69] C. Gaspar, M. Dönszelmann, Ph. Charpentier, DIM, a portable, light weight package for information publishing, data transfer and inter-process communication, Comput. Phys. Commun. 140 (2001) 102

[70] C. Gaspar, DIM - Distributed Information Management System, http://dim.web.cern.ch/dim/

[71] S. Kersten, Overview on the Pixel DCS Hardware, EDMS ATL-IP-CD-0002 (2006)

[72] J. Steinhart, S. Hart, Calibration curves for thermistors, Deep Sea Research and Oceanographic Abstracts, Volume 15, Issue 4, August 1968

[73] semitec, Thermistor 103KT1608, Data sheet available online at http://www.semitec.co.jp/english2/products/pdf/KT_Thermistor.pdf

[74] Hygrometrix, Humidity Sensor Xeritron XNAM 10205-H, Data sheet available online at https://twiki.cern.ch/twiki/pub/Atlas/IDHumiditySensors/xeritronds2005.pdf

[75] Hygrometrix, Humidity Sensor HMX-2000, Data sheet available online at https://twiki.cern.ch/twiki/pub/Atlas/IDHumiditySensors/HMX2000DSv2d.pdf

[76] Honeywell, Humidity Sensor, http://sensing.honeywell.com

[77] S. Kersten, P. Kind, Pixel Patch Panel PP3 Design, EDMS: ATL-IP-ES-0075 (2005)

[78] M. Citterio et al., Voltage-regulation stations for the distributed power supply of the ATLAS pixel detector, IEEE Nucl. Sci. Symp. Conf. Rec. 1 (2007) 704

[79] M. Citterio, C. Meroni, Pixel PP2 Patch Panels design, EDMS: ATL-IP-ES-0074

[80] S. Mehlhase et al., A Summary of Thermal Measurements from the Pixel Commissioning Program, ATL-COM-INDET-2011-002

[81] S. Kersten, The Interlock Matrix of the Pixel Detector Control System, ATL-IP-ES-0110

[82] R. Boyd, F. Pedrosa, The DSS usage by the ATLAS Pixel Detector, ATC-TY-ON-0014

[83] The Glance project, http://atglance.web.cern.ch/atglance/DSS/

[84] J. Boek, Das Interlocksystem des ATLAS Pixeldetektors, Diploma Thesis, Universität Wuppertal, WU D 07–07, April 2007

[85] J. Schultes, Das Backendsystem für die ATLAS-Pixeldetektorkontrolle, Diploma Thesis, Universität Wuppertal, WU D 03–07, May 2003

[86] K. Lantzsch, Integrationssoftware für das Kontrollsystem des ATLAS Pixeldetektors, Diploma Thesis, Universität Wuppertal, WU D 06–08, August 2006

[87] T. Henß, Benutzerschnittstellen des ATLAS-Pixeldetektorkontrollsystems, Diploma Thesis, Universität Wuppertal, WU D 05–04, May 2005

[88] J. Moss, private communication

[89] T. Boek, Die Finite State Machine für das Interlock System des ATLAS Pixel Detektors, Diploma Thesis, Universität Wuppertal, WU D 08–02, February 2008

[90] S. Weber, Kommunikation zwischen Datennahme- und Kontrollsystem des ATLAS Pixeldetektors, Diploma Thesis, Universität Wuppertal, WU D 07–15, October 2007

[91] PIXEL Humidity Sensors, EDMS: ATL-IP-ER-0021

[92] A. Micelli, Integration of the Environmental Monitoring into the Pixel Detector Control System for the ATLAS experiment at the Large Hadron Collider, Tesi di Laurea, Universita degli Studi di Udine, 2008

[93] C. Tsarouchas et al., ATLAS Detector Control System Data Viewer, proceedings of ICALEPCS 2011, Grenoble, France

[94] J.-F. Arguin, J. Beringer, A. Gaponenko, M. Zdrazil, CoralDB : The Connectivity Database for the ATLAS Pixel Detector, ATL-COM-INDET-2008-001

[95] R. Chytracek et al., POOL development status and production experience,
IEEE Trans. Nucl. Sci. 52 (2005) 2827

[96] T. Henß, Entwurf und Implementation eines Expertensystems für das Detektorkontrollsystem des ATLAS-Pixeldetektors, Dissertation, Universität Wuppertal, WUB DIS 2008-08, December 2008

[97] A. Barriuso Poy, S. Schlenker, FSM Integration Guideline, ATL-DQ-ON-0010

[98] J. Schultes, Qualifizierungsmessungen des Spannungsversorgungssystems sowie Konzeptionierung einer Zustandsmaschine für die Detektorkontrolle des ATLAS-Pixeldetektors,
Dissertation, Universität Wuppertal, WUB-DIS 2007-04, February 2007

[99] M. Citterio, M. Leyton, Private Communication

[100] S. Berry, Analyse Fonctionnelle Atlas Pixel, SCT C3F8 Evaporative cooling system, EDMS Id: 830871

[101] Reyes Alemany, Mike Lamont, Kris Kostro, LHC - Experiments Handshake Protocol over DIP,
LHC-OP-ES-0019

[102] A. Andreazza et al., Effect of accidental beam losses on the ATLAS pixel detector,
Nucl. Inst. and Meth. A, 565 (2006) 50-54.

[103] D. Bocian, Accidental beam losses during injection in the interaction region IR1.,
LHC Project Note 335, January 2004.

[104] B. di Girolamo, Pixel Detector Turn-on Strategy: LHC conditions dependency,
EDMS, ATL-IP-ON-0004

[105] R. Bruce, LHC collimation at IP1 Presentation at ATLAS Pixel General Meeting, 08-06-2010