Fachbereich Mathematik und Naturwissenschaften
Fachgruppe Physik
Bergische Universität Wuppertal

Commissioning of the ATLAS Pixel Detector optical data transmission and studies for readout of the ATLAS IBL and future trackers

Dissertation zur Erlangung des Doktorgrades
vorgelegt von

Jens Dopke

Juni 2011

Contents

To Kerstin Lantzsch

... for always being on call, when I needed her, for falling asleep on her keyboard,
for picking me up in the morning and driving me home at night,
beer, fries, chocolate, the ski club, introducing me to Nethack...

For not yelling at me after my first week with more than 130 h awake at CERN
For letting me have her office chair, (to sleep on it).
For liking the food I prepared and trying everything without any hesitation.

For motivating me to at least try and be better than I am.

Zusammenfassung

Die Hochenergiephysik beschäftigt sich mit der mikroskopischen Struktur der Materie. Das in diesem Rahmen definierte Standardmodell der Teilchenphysik versucht die Bausteine der Materie sowie ihre Wechselwirkungen zusammenzufassen. Seit seiner Niederlegung in den späten 60er Jahren besteht das Standardmodell alle an ihm durchgeführten Tests, zuletzt durch den erfolgreichen Nachweis der Existenz des Top Quarks im Jahre 1995. Seitdem versuchen Teilchenphysiker den letzten Baustein des Standardmodells zu finden, das Higgs-Boson. Das 2. Kapitel der vorgelegten Arbeit wird den, durch das Standardmodell vorgegebenen, Rahmen diskutieren und einen Überblick über aktuelle Themen der Teilchenphysik geben.
Die derzeit grösste Maschine der Teilchenphysik, der Large Hadron Collider am CERN, wird zusammen mit dem ATLAS Experiment in Kapitel 3 eingeführt.
Der Rahmen dieser Dissertation ist das innerste Subsystem von ATLAS, den ATLAS Pixel Detektor. In einem sensitiven Volumen von nur 5 millionstel des gesamten ATLAS Detektors, liefert der Pixel Detektor ungefähr 80% aller Auslesekanäle. Im speziellen wird in Kapitel 4 nicht nur der Pixel Detektor, sondern vor allem sein optisches Auslesesystem behandelt.
Diese Arbeit setzt sich mit der Installation, Konfiguration und Leistung des ATLAS Pixel Optolinks, der optischen Auslese, auseinander und führt dann über zu der optischen Auslese des ersten ATLAS Upgrades, des Insertable B-Layers. Kapitel 5 beschreibt die Mechanismen zur Evaluierung und Optimierung der Pixel Datenübertragung und zieht daraus Erkenntnisse für die Datenübertragung des Insertable B-Layers. In Kapitel 6 wird dann der, noch in Entwicklung befindliche, Insertable B-Layer sowie dessen Auslesesystem diskutiert.
Abschliessend gibt Kapitel 7 einen Überblick der derzeitlichen Entwicklung von Auslesesystemen, sowie der Problematik der Datenauslese an künftigen Collider Experimenten mit stark erhöhter Luminosität.

1 Introduction

Particle physics aims to understand the innermost structure of matter and define the laws of nature from that. In the last fourty years a Standard Model of particle physics has come to great success. After being postulated in the late 60s, it withstood all tests and all of its predictions were proven true, eventually by the discovery of the top quark in 1995. Since then, scientists are trying to tackle the last piece of the Standard Model which is still missing: the Higgs-boson is, to date, yet to be observed. At the same time many new theories have evolved and must be tested for validity. The 2$^{\text{nd}}$ chapter will describe the Standard Model and will outline recent topics in particle physics.
The latest machine, built to investigate in the field of particle physics, is the Large Hadron Collider at CERN, Geneva. One of the Experiments, ATLAS, is the largest machine ever built by the high energy physics community. With an enormous outer dimension of 46 m length and 22 m diameter it observes the highest energy collisions of protons ever produced. Both, the Large Hadron Collider and ATLAS will be described in the 3$^{\text{rd}}$ chapter.
Built from many different detector subsystems, many people were needed to allow for a first successful ATLAS operation in September 2008 and finally stable operation in 2009 and later.
This thesis deals with the readout system of the innermost detector of ATLAS, the Pixel Detector, which is introduced in chapter 4. The Pixel Detector provides about 80% of the total number of readout channels within a sensitive volume, no more than 5 millionth the size of the total ATLAS volume. The core of this thesis describes the layout and commissioning of the ATLAS Pixel Detector readout system, as well as the first stage upgrade of the Pixel Detector towards higher instantaneous luminosities, the Insertable B-Layer.
The author has worked on the Pixel Detector readout since 2003 and made major contributions to the installation and commissioning of the readout system in 2007-2009, most of which are described in chapter 5, as well as to the readout system for the Insertable B-Layer, described in chapter 6.
Chapter 5 outlines the commissioning phase of the ATLAS Pixel Detector optical readout link. Algorithms to measure performance and tune parameters of the optical link are described together with the general performance of the Pixel Detector optical readout system, as well as its shortcomings.
Chapter 6 then introduces the Insertable B-Layer project and how in particular its readout system is designed to surpass the Pixel Detector readout system.

Chapter 7 summarises ideas for even higher luminosities, and the problems that may arise with new readout systems currently under consideration.

2 Physics at the LHC

This chapter introduces the physics useful to know within the scope of this document. The Standard Model subsection explains the context of this thesis. A short introduction will be followed by a section describing recent topics in particle physics, giving reason to how things are done in the following chapter.

2.1 Standard Model of Particle Physics

The Standard Model [Gri08] of particle physics describes the world in words of fundamental particles and their interactions. The interactions include electromagnetic, weak and strong interaction. Gravitation has not yet been included into the Standard Model. The Standard Model was established in the 1960s and 70s in a process of trying to combine all of natures forces within a single combined theory. To date it has withstood all experimental tests, yet it features some "ugly" aspects:

- Hierarchy: The Standard Model defines twelve particles, which basically are three different mass editions of four particles. To date there is no good description of why these families exist and why they are so different in terms of mass. Instead of these masses being defined by the SM, the SM is (partly) defined by these masses.
- Baryon Asymmetry: The universe as we see it contains a large amount of matter, but only very little antimatter. To date there is no good measurement that allows for matter and antimatter to be that different, assuming an initial equal amount of both.

A lot of additions and extensions to the Standard Model exist, trying to explain those and other aspects of it. The most famous, the Higgs mechanism is integrated into the standard model, explaining masses of exchange bosons within the electroweak theory. A particle that should exist within this mechanism, the Higgs boson, is yet to be found.

2.1.1 Matter in the Standard Model

The Standard Model is composed of twelve particles with spin 1/2, the constituents of all matter. As half-spin particles they follow Fermi-Dirac statistics and the Pauli-exclusion principle. Each of them has an anti-particle with inverted quantum numbers.

Table 2.1: *Particle content of the Standard Model*

Type	Charge	1. generation	2. generation	3. generation
Quarks	+2/3	up (u)	charm (c)	top (t)
	-1/3	down (d)	strange (s)	bottom (b)
Leptons	-1	electron (e)	muon (μ)	tau (τ)
-neutrinos	0	ν_e	ν_μ	ν_τ

Table 2.2: *Force mediators of the Standard Model*

Interaction	Mediator	Charge [e]	Spin	Mass [GeV/c^2]
Electromagnetic	**Photon**	0	1	0
Weak	**W$^\pm$**	$\pm$1	1	80.4
	Z	0	1	91.2
Strong	**Gluon**	0	1	0

All particle content of the Standard Model is grouped into quarks and leptons. While the quarks can be separated into up-type and down-type quarks, the leptons can be subdivided by their electric charge into neutrinos and charged leptons. Each of these general types exists in three different flavours or generations (c.f. Table 2.1). All stable matter is composed of first generation particles. The second and third generation are higher mass copies of the first generation and decay into first generation particles.

All particles are subject to some or all of the four fundamental interactions, depending on the charges they carry: electromagnetic, weak and strong interaction, as well as gravitation.

2.1.2 Interactions in the Standard Model

The Standard Model provides precise descriptions for three fundamental interactions: Electromagnetic, Weak and Strong force. Despite many attempts, a unified theory incorporating all four fundamental forces, including gravitation, is not yet available. In contrast to the other three forces, gravitation does not play a major role in particle physics, as it is much weaker than the other forces. Its meaning in every day life comes from the fact that all the other forces allow to cancel out over large distances, whilst gravitation adds up, creating macroscopic phenomena.

In the Standard Model, each force is mediated via force carriers, bosons, which are summarised in Table 2.2.

Quantum Electrodynamics

Quantum ElectroDynamics [Fey49] (QED) is a U(1) gauge theory describing the electromagnetic interaction in accordance with both, quantum mechanics and special relativity. Inclusion of both into a renormalisable theory was a major success of the late 40s and made QED the first working quantum field theory. This lead to subsequent attempts to describe the other observed interactions as gauge theories too, following similar concepts and methods.

Weak Interaction and QED

The most known observation of weak interactions is in nuclear β-decay. Combination of electromagnetism in the form of quantum electrodynamics and the weak interaction lead to a full electroweak theory, based on an SU(2)xU(1) gauge theory. This new theory postulated by Glashow, Weinberg [Wei67] and Salam in the late 60s proposed that not only charged weak currents as observed in β-decay, but also neutral weak currents should exist. Those were first measured at the Gargamelle bubble chamber in 1973 [H^+73]. Only in 1983 a direct observation of both $W^\pm$ [A^+83b, B^+83b] and the Z^0-bosons [A^+83a, B^+83a] succeeded.

The electroweak theory gives a total of four exchange particles. Three of these particles, the force mediators of the weak force, are given mass by the Higgs mechanism [Hig64]. This mass is what leads to the weakness of the interaction; creation of a boson can either happen with the proper amount of energy available (on-shell) or without it (off-shell). Their high masses of $\approx$81 and $\approx$90 GeV/c^2 explains the very short range of the weak interaction, as the gauge bosons are typically created off-shell.

Albeit weak, the weak interaction has a very important feature: exchange of a W-bosons allows one to change the flavour of quarks and thus allows for example the decay of top-quarks.

Another aspect that makes weak interaction very interesting is its symmetry violating behaviour. Weak gauge bosons only interact with left handed particles or right-handed anti-particles, which can be described as a maximal violation of parity symmetry (mirroring of space). Additionally, the weak interaction violates CP-symmetry (where C stands for particle/antiparticle exchange), although this violation is not maximal [CCFT64].

Strong Interaction

The strong interaction, described by Quantum ChromoDynamics (QCD) [N^+10], binds quarks into e.g. protons and neutrons, and further into atomic nuclei. Like electromagnetic interaction happens between electrically charged particles via exchange of a photon, strong

interaction happens between particles carrying colour charge, that is, one out of three colour charges or their corresponding anti-charge.

The carrier of the strong force is the gluon. Gluons couple to colour charge, and carry a colour and an anti-colour themselves. In contrast to the electromagnetic interaction, the strong interaction is therefore self-coupling, allowing for two gluons to fuse into one or two others.

As colour charge does not appear in every day life, it was introduced as forcing itself into a colour neutral state[a], either by pairing of colour and the corresponding anti-colour or by the sum of three different (anti-)colours. The force between colour-charged particles increases with their distance. Tearing two quarks of opposite colour away from each other (assuming they have been in a bound, colour-neutral state before) eventually generates enough energy to create two new quarks in between. The colour connection of the quarks is thus broken by inserting two new quarks. The concept of stronger forces at larger distances is known as colour-confinement, the observation of quasi-free quarks at small distances as asymptotic freedom [GPW04].

2.2 The Particle Physics Frontier

Most measurements within particle physics nowadays are either performed in the frame of astroparticle physics or within a particle accelerator. Astroparticle physics relies on incident cosmic rays hitting a detector. Depending on the energy in question, these rays become extremely rare and thus the detectors need to become very large. Accelerator based experiments happen within a more well-defined environment, allowing for precise measurements of processes. Yet, they lack the very high energies observed in astroparticle physics [Rou10].

Most accelerators use bunches of stable particles, which are accelerated within a machine and brought to collision either with a target or another bunch of accelerated particles. These two techniques are referred to as fixed target or collider experiments. Behind or around the interaction point, depending on the technique used, detectors are registering particles that pass through them. Measuring the momentum and energy of particles, as well as their charge and particle type, it is in theory possible to reconstruct the microscopic process that took place within the interaction point. Yet, multiple processes can fit the same observed "particle tree". Statistics help to figure out whether a process is taking place as expected or not, by observing the cross-section of those tree structures occurring and comparing it with theory predictions.

[a] A kind of white...

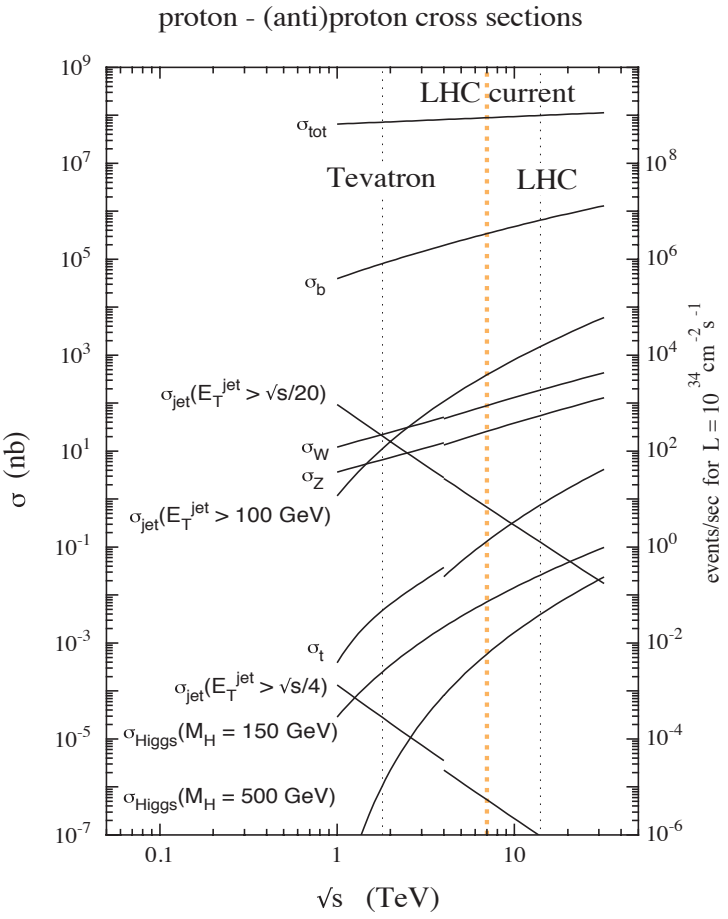

Figure 2.1: *The proton-(anti)proton cross section, plotted for the LHC at right (14 TeV) and Tevatron at left (1.8 TeV here), taken from [AC98]. The recently reached LHC energy of 7 TeV is highlighted by the central orange line.*

2.2.1 Process Cross-section

To observe signatures of new (or known) physics, an evaluation of the probability to observe certain processes needs to be done. Production and decay both play an important role for calculating the cross-section of a process and a corresponding integrated luminosity within which this process will either be observed or disregarded. A plot of the possible Standard Model processes as predicted for the Tevatron[b] and the LHC (c.f. Section 3.1), is given in Figure 2.1.

[b]Tevatron is the proton-antiproton collider at the Fermi National Accelerator Laboratory, situated close to Chicago in the United States.

To sample even very rare processes within a finite time, a very high amount of protons, the initial particles in case of the LHC, needs to be smashed into each other to produce enough luminosity.

2.2.2 Luminosity

Luminosity is the capability of a collider to produce an interaction between particles. The instantaneous luminosity is typically measured in $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and defined by:

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

N_b is the number of particles per Bunch, n_b the number of bunches per beam, f_{rev} the revolution frequency, γ_r the relativistic gamma factor, ϵ_n the normalised transverse beam emittance and β^* the beta function at the collision point. F is a geometric luminosity reduction factor coming from the crossing angle at the interaction point.
The integrated luminosity is typically given in inverse picobarn (pb^{-1}) or femtobarn (fb^{-1}), depending on the order of magnitude. Multiplying it by a processes cross-section, one can calculate the expected number of events of the process in question. E.g. recording $1\,\mathrm{pb}^{-1}$ with the LHC at nominal energy, one can derive from Figure 2.1 that about 50 Z^0s must have been recorded.

Minimum Bias Physics An early subject in all new collider experiments is determination of the total cross-section and the dominating background characteristics. Background here is physics processes happening at lower energies than is interesting. The dominating background is often referred to as Minimum Bias [A+11]. Whilst zero bias is fully random, minimum bias expects an event to have happened. That means there should have been a minimum amount of observable content. The major part of minimum bias content can be pinpointed in Figure 2.1 as the area between σ_{TOT} and σ_b.
Measuring this content and comparing it with results from simulation allows to validate simulation processes. Assuming they do not deliver the exact results, tuning the simulation to data promises more accurate simulation of the more important hard physics processes, that is, processes with a larger energy transfer / heavier particles being involved. A second issue is, that minimum bias events dominate the data taking in recent experiments. Calculating e.g. the ratio of Higgs cross-section with a Higgs mass of 150 GeV over the total cross-section, one comes up with more than $1 \cdot 10^{10}$. That implies one has to produce more than 10 billion interactions at 14 TeV before creating a single Higgs boson (with m_H = 150 GeV).
A third aspect of understanding minimum bias events is that of describing backgrounds: Getting to very high instantaneous luminosities, multiple events start to occur during

a single crossing of projectile bunches. Hence, a recorded bunch crossing contains mostly background and only rarely an interesting event. Therefore, the minimum bias background needs to be well understood to identify a signal within the background events.

Top physics Besides observing known interactions with higher precision, recent particle physics aims for observation of top quark production and decay, as it is the heaviest quark. The top quark was first observed at Tevatron in 1995 [A$^+$95]. Due to their small lifetime, top quarks decay before they start to build hadrons with other quarks. Therefore the top quarks themselves are observed as quasi-free single quarks.

In addition, their high mass plays an important role in searching for new physics. It allows for probing new couplings, which are mass sensitive, and their mass themselves is a mystery to be sorted out.

Top quarks decay with 99.9% probability into a W-boson and a b-quark, the latter of which can then hadronise. The b-quarks within the generated hadrons have a preference to turn into top quarks via radiating a W-boson over turning into anything else. As this is hardly possible due to the lower mass of the b-quark, the lifetime of the b-hadrons is enhanced over lifetimes of other hadronic compounds. This lifetime allows them to move out from the primary vertex, depending on their boost up to a few centimetres. Further decay then rapidly starts forming jets[c] of particles. As the B-hadrons have travelled some distance before decaying, the jets they generate originate from a vertex that is different from the primary vertex. These jets can thus be tagged as b-jets, based on them coming from a secondary vertex.

Higgs Physics Moving towards discoveries, the most interesting recent search is that for the Higgs boson: in the Standard Model the Higgs-boson was predicted with a mass just above 100 GeV, which is already excluded by LEP experiments [P T08]. Yet, there still is a window with low probability between 115 and 160 GeV, that would fit the Higgs boson.

If a standard model Higgs Boson exists, it probably decays into b-quarks, W and Z-pairs or eventually even top quark pairs, depending on its mass. Finding the Higgs boson will depend very much on the detectors built around the interaction points. They need to be able to reconstruct mass very accurately and determine the collision energy. As mentioned before, another major aspect will be identification of B-hadrons, as Higgs-boson production will most probably come with many b-quarks being produced, be it direct, or via decay from top quarks or weak bosons.

[c] A spray of particles coming from the same origin and moving mostly along a common axis, with a typically low separation between the particles.

2.2.3 Conclusion

This chapter was to show, that particle physics in recent days relies on high statistics and high precision measurements at collider experiments, to verify our recent knowledge and reach a better understanding of how subatomic interactions work.

3 Experimental Setup and Boundaries

This chapter introduces the Large Hadron Collider and the ATLAS Experiment, providing the basis for the following chapters. It will shortly discuss the LHC machine and its future. The later subsection gives the basics of the ATLAS machinery.

3.1 The Large Hadron Collider

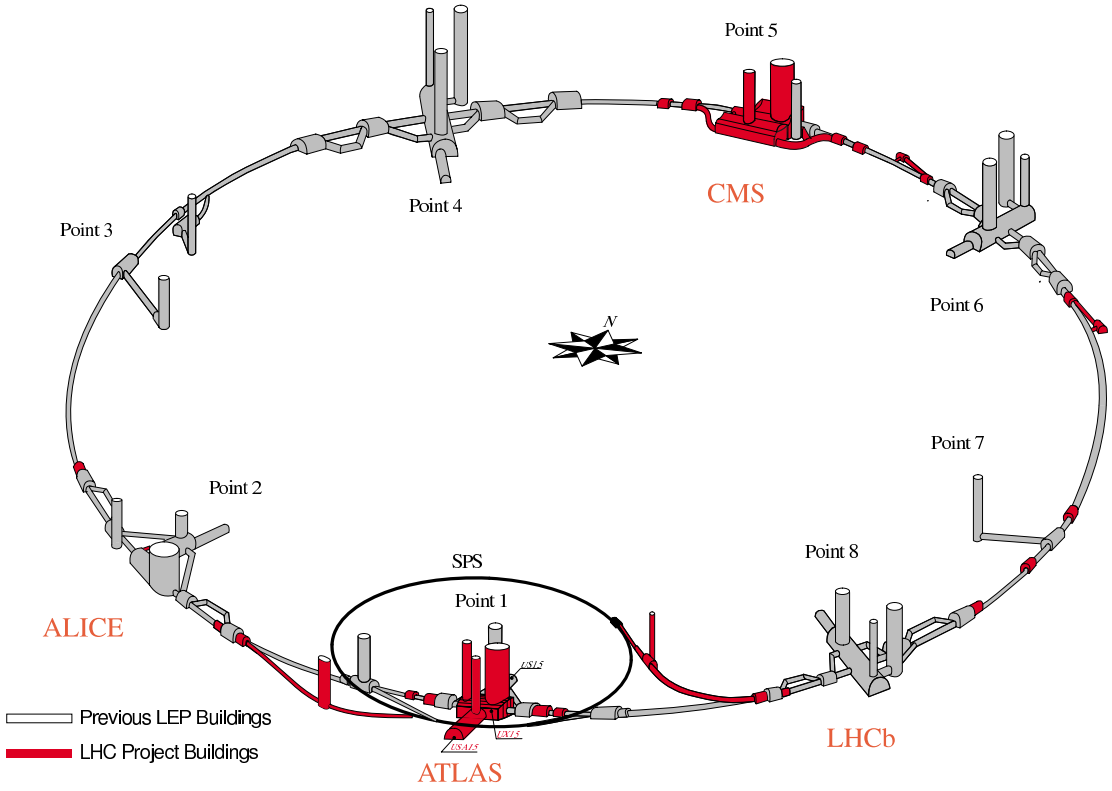

Figure 3.1: *CERNs underground structures with new structures set up for the LHC operation [LHC09].*

The Large Hadron Collider (LHC) [LDR04] (c.f. Figure 3.1) is the highest energy particle accelerator in the world. By design it delivers up to 14 TeV proton and 1148 TeV heavy ion collisions to four large and two small experiments. The LHC machine inherits the 26.7 km long underground tunnel and most service caverns from the previously installed machine,

the Large Electron Positron (LEP) collider. Only in point 1 and point 5, new access shafts and experimental caverns had to be prepared.
Within the LHC tunnel, the previous LEP installation had to be replaced. A new superconducting magnet system incorporating 1232 dipole magnets was installed, delivering a maximum bending field of 8.33 T in normal operation. The system delivers two opposite direction dipole fields to two beampipes. The two beampipes are to carry protons or heavy ions travelling in opposite directions through the LHC. They run in parallel all through the ring. Only within the four experimental areas the beampipes are connected to allow the particle beams to cross each other and collide within the experiments.
Protons are delivered from an initial linear accelerator and a booster, whilst heavy ions are delivered from the low energy ion ring (LEIR) accelerator. Exiting those with an energy of 1.4 GeV, they are accelerated by the Proton Synchrotron (PS) up to 25 GeV to then be injected into the Super Proton Synchrotron (SPS). Injection into the LHC happens with an energy of 450 GeV. The two different beam pipes of the LHC are filled via two different transfer lines from the SPS. The LHC then accelerates the protons to 7 TeV design energy, but is currently limited to a maximum energy of 3.5 TeV. The maximum center-of-mass energy of a collision can therefore be 7 TeV recently, but will reach 14 TeV in the coming years.
Within the LHC, protons are bunched in packages of $1.1 \cdot 10^{11}$ protons, travelling with a minimum distance of roughly 7.5 m or 25 ns. The bunch crossing rate therefore is 40 MHz. A full LHC fill contains at maximum 2808 filled bunches per beam, out of the possible 3564 bunches, as given by their minimum distance, therefore reducing the maximum interaction rate to 31.6 MHz. Using quadrupole magnets around the experimental area, those beams are squeezed to have a tiny diameter within the centre of the experiments (correlated with β^* in Equation 2.1). This allows for the particles to actually collide, as the particle clouds would otherwise just pass through one another.

3.1.1 Experimental Installations

As can be seen in Figure 3.1, the four large experiments are installed in point 1, ATLAS, point 2, ALICE, point 5, CMS, and point 8, LHCb.
ATLAS and CMS are general purpose detectors surrounding the interaction region with silicon trackers, calorimeters and muon chambers. They are built for Standard Model investigations and beyond, seeking to identify the Higgs boson and/or other new physics. ALICE is optimised for heavy ion collisions, able to resolve O(16k) tracks within its inner detector. LHCb is the only large forward experiment. It searches for boosted B-Mesons produced in LHC collisions, to measure CP violation in the B-sector.
The LHC machine is built to deliver proton-proton interactions at a luminosity of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ to ATLAS and CMS. The other experiments are run with a lower luminosity

to not generate multiple interactions per bunch crossing (LHCb) and keep the occupancy within the central detectors low (ALICE). With the number of bunches and protons per bunch being fixed, together with the LHC tunnel length and proton energy, the only parameters left for defining luminosity are the beam emittance, the geometrical reduction factor and β^* (see Section 2.2.2). The minimum β^* reachable with the installed quadrupoles is 0.55 m.

3.1.2 The LHC Luminosity Upgrades

The first stage upgrade already planned for the LHC will not increase the maximum energy, but the luminosity of the accelerator. It is planned for 2016 and will allow to probe high energies within the LHC experiments with better significance, as high energy processes will happen more frequently. The Phase I upgrade is to deliver three times the instantaneous luminosity of the original LHC machine. One increasing factor will come from a new collimator system. This will allow one to shrink β^* (see Equation 2.1) to 0.25 m thereby increasing the LHC luminosity by a factor of 2.2. Additionally a new linear accelerator is being built at CERN now, delivering a lower emittance (ε in the aforementioned equation). As this upgrade increases the number of interactions per bunch crossing by a factor of three, the experiments will need to adjust their trigger mechanisms and ultimately install new hardware components and detector layers to mitigate the effects of increased pileup.

A second stage upgrade is in its R&D phase, to allow an even higher instantaneous luminosity, up to ten times the nominal luminosity of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. As β^* and ε will not be feasible for further increases, this upgrade is planned to change the geometric factor F. Crab cavities [A+09c] allow to create a larger beam overlap, by tilting beam bunches so that they collide head-on, opposed to colliding with a small angle, c.f. Figure 3.2. They also enable one to change the luminosity throughout a run, by having an angled collision initially and then slowly changing for head-on collisions. This process, called luminosity levelling, keeps a constant luminosity throughout a run.

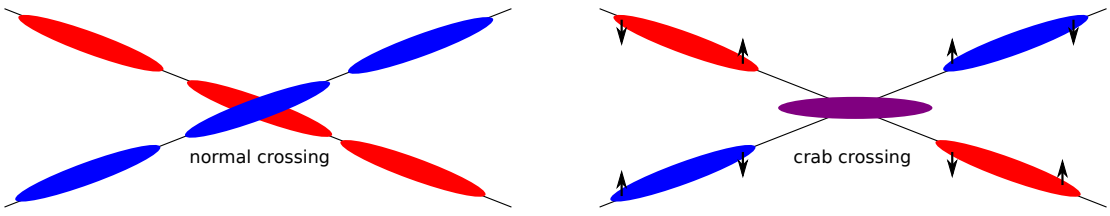

Figure 3.2: *Crab Crossing concept for future LHC upgrades: Instead of crossing bunches with an angle, they are turned to collide head-on.*

3.2 ATLAS

A Toroidal LHC ApparatuS [TAC99], ATLAS, is one of the four large experiments of the LHC. With a total size of 46 m in length and 22 m in diameter, it is the largest experiment at CERN. The biggest part herein is the muon chambers with an air-coil toroidal magnet system.

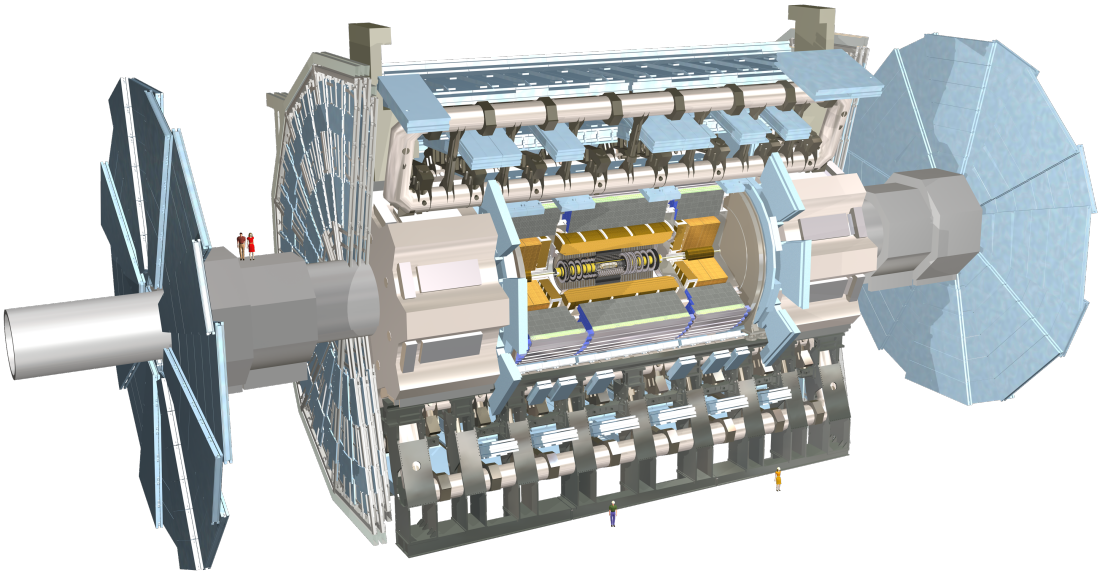

Figure 3.3: *The ATLAS experiment at CERN, a schematic side-view of the detector.*

Describing the detector layout in the following sections, we will move from the inside out, following particles created within ATLAS. A tracking system measures momenta, by delivering spacepoints along the tracks of charged particles that are curved within a solenoidal magnetic field. Afterwards two calorimeter systems absorb those particles, measuring their total energy. Eventually, particles that escaped the calorimeter system (muons) are measured for their momentum within the muon system. Both, calorimeter and muon system, contain fast subsystems that are built to deliver a low latency realtime[a] reading to trigger full event processing.

The ATLAS detector is equipped with subdetectors covering an $|\eta|$ range up to 4.9. Special systems go even closer to the beampipe, e.g. the zero degree calorimeter (ZDC) is set up in the direct forward section, where the ATLAS beampipe splits into the two individual pipes for the LHC. ATLAS is therefore referred to as a 4π detector[b].

[a] i.e. fixed latency

[b] Covering the full solid angle around the interaction point

Table 3.1: *Performance goals of the ATLAS detector as described in [A+08a]*

Detector component	**Required resolution**	η **range**
Tracking System	$\sigma_{p_T}/p_T = 0.05\%\ p_T \oplus 1\%$	$\pm$ 2.5
EM calorimetry	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	$\pm$ 3.2
Hadronic calorimetry (jets)		
barrel and end-cap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	$\pm$ 3.2
forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \lvert\eta\rvert < 4.9$
EM calorimetry	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1\,\text{TeV}$	$\pm$ 2.7

To describe the ATLAS detector, a short description of the ATLAS coordinate system is given here, introducing the aforementioned variable η, which is frequently used in the frame of particle physics.

Coordinate System: The ATLAS coordinate system centre is the nominal collision point within the centre of the detector. The x-axis points from there horizontally towards the centre of the LHC accelerator. The y-axis is defined to point up, so that z- goes along the beam, from ATLAS towards the city of Geneva. The detector cavern is split in two sides which are called A at positive z and C at negative z.

Within the frame of particle physics and functional detector description, cartesian coordinates are not necessarily the best choice. The detector can also be described in terms of R, the radius from the z-axis, φ, the angle around z in mathematical sense and η, the pseudorapidity, which is calculated from the angle Θ against the positive z-axis:

$$\eta = -\ln\tan\Theta/2 \tag{3.1}$$

η is close to rapidity[c], thus angular distributions can be described much better in terms of η than Θ. In the plane perpendicular to the beam axis at $z = 0$, η is 0. Due to the symmetry of the detector around this plane, often things are described in regions of $|\eta|$.

Requirements The ATLAS Detector is designed to meet requirements on energy and momentum resolution in its given subdetectors. These requirements are listed in Table 3.1. They are derived from the precision (ultimately) required for the observation of the physics processes targeted by ATLAS.

[c]See e.g.: Rapidity, Wikipedia

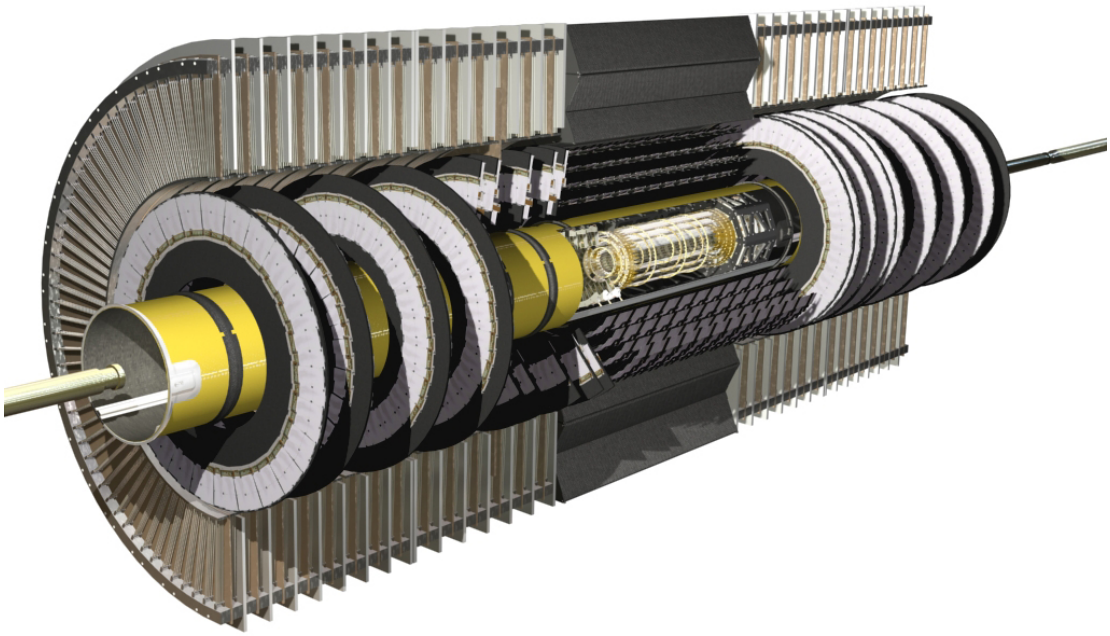

Figure 3.4: *The ATLAS inner detector[AHP]. Parts of TRT and SCT are removed to show the inside objects, SCT and Pixel.*

3.2.1 Inner detector

The ATLAS inner detector measures particle tracks and thereby their momentum, influencing them as little as possible, to not affect that measurement as well as the later total energy measurement in the following detector. The inner detector consists of three subsystems: the Pixel Detector, the SemiConductor Tracker and the Transition Radiation Tracker. Their dimensions are given in Table 3.2.

The central Pixel Detector is to deliver high resolution spacepoints both in $R \cdot \varphi$ and in z-direction. Three layers of pixelised detector deliver a good handle on tracks originating from the ATLAS centre. The detector is capable of resolving multiple interaction vertices, as well as secondary vertices, which are displaced from the beam axis and arise e.g. from the decay of B-hadrons. The Pixel Detector is delivered to ATLAS including a 7 m support structure carrying some on-detector services, as well as the ATLAS Beam Conditions Monitor (BCM) [C+08] integrated with it.

The next detector layer is the the SemiConductor Tracker (SCT). An initial track seed is done using its four (double) layers. Due to its lower granularity it is well suited for seeding tracks, which can then be extended into the Pixel volume to get their final precision.

The outermost tracking system is the Transition Radiation Tracker (TRT) [C+04]. The TRT is a straw tube tracker with transition radiation detection. 36 straw tubes are hit on

Table 3.2: *Parameters of the ATLAS inner detector [A+08a]*

Subsystem	Part	Radial extension [cm]	z-extension [cm]
Beampipe		$2.9 < R < 3.6$	
Pixel	Overall envelope	$4.55 < R < 24.2$	$\|z\| < 309.2$
	Sensitive barrel	$5.05 < R < 12.25$	$\|z\| < 40.05$
	Sensitive endcap	$8.88 < R < 14.96$	$49.5 < \|z\| < 65.0$
SCT	Overall envelope (Barrel)	$25.5 < R < 54.9$	$\|z\| < 80.5$
	Sensitive barrel	$29.9 < R < 51.4$	$\|z\| < 74.9$
	Overall envelope (Endcaps)	$25.1 < R < 61.0$	$81.0 < \|z\| < 279.7$
	Sensitive endcap	$27.5 < R < 56.0$	$83.9 < \|z\| < 273.5$
TRT	Overall envelope (Barrel)	$55.4 < R < 108.2$	$\|z\| < 78.0$
	Sensitive barrel	$56.3 < R < 106.6$	$\|z\| < 71.2$
	Overall envelope (Endcaps)	$61.7 < R < 110.6$	$82.7 < \|z\| < 274.4$
	Sensitive endcap	$64.4 < R < 100.4$	$84.8 < \|z\| < 271.0$

average by every track originating from the centre of ATLAS. Also, each particle is to pass through layers of foil, emitting transition radiation that is registered by the TRT and used to distinguish charged pions from electrons/positrons.
Pixel and SCT as high density detectors need a triggered readout, as full rate readout would only be possible at multiple GBit/s, not available in radiation hard technology today. The TRT on the contrary supports a quasi-realtime readout path allowing for a special trigger signal used in cosmic data taking, the TRT FAST-OR. With this it is possible to record the full cosmics rate available to the ATLAS inner detector.

3.2.2 Calorimeter system

The calorimeter system is a hermetic structure composed of several subsystems to cover individual areas and purposes.
The electromagnetic calorimeter consists of a barrel system covering $|\eta| < 1.475$ and two endcaps reaching out to $|\eta| = 3.2$. The sensitive material is liquid Argon here, interleaved with lead absorber plates. Particles passing through the liquid Argon ionise it, creating charge signatures that can be read out.
The hadronic calorimeter is a mixture of a scintillator tiles based system in the barrel and a liquid argon based system in the forward region. Whilst the barrel system uses steel absorbers, the end cap absorbers are built from copper.
A very forward coverage is established with the forward calorimeter (FCAL), based on liquid argon with copper and tungsten absorbers. Due to the regular exchange of the

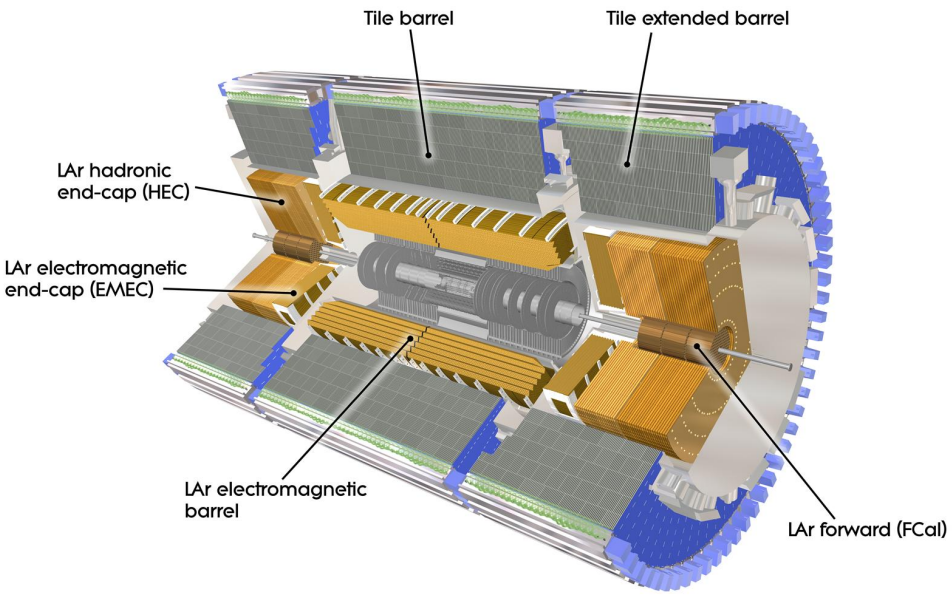

Figure 3.5: *The ATLAS calorimeter system (taken from [AHP]). Removed parts allow to see the inner electromagnetic subsystem and the forward calorimeter.*

detecting material, the FCAL is very radiation hard, allowing it to cope with very high expected particle fluxes. Depending on Luminosity, the FCAL might yet develop problems due to high energy deposition, which could boil the Liquid Argon.
Opposed to the inner tracking detectors, most parts of the calorimeters come with large frontend electronics mounted within the detector. These allow for a low latency and a standard readout path, the former delivering a coarse granularity information to the trigger system. The standard readout path on contrary contains full details of information, but needs triggered readout just like the inner detectors. In case of the electromagnetic calorimeter for example, the trigger subsystem of the frontends builds energy towers along the depth of the calorimeter as seen from the interaction point [B+08]. This allows for fast total energy measurements and jet identification.

3.2.3 Muon system

Outside the calorimeter, the muon spectrometer is located. The only charged particles escaping from the calorimeters are muons, thus the naming of this detector part. To reach the target physics performance the ATLAS muon system must be capable of measuring muons up to 1 TeV with an energy error of less than 10%. Four different detector types are used:

- Monitored Drift Tubes (MDT): Are a precise measurement system, delivering a 30 μm position information in η (z) and R for the barrel and endcap regions respectively.
- Resistive Plate Chambers (RPC): Are fast tracking chambers with a coarse resolution. They deliver φ-Information for the MDT system in the barrel and are part of the Level 1 muon trigger system.
- Cathode Strip Chambers (CSC): These are the very forward muon chambers, replacing the MDT system due to higher expected hit rates.
- Thin Gap Chambers (TGC): Those deliver the same information as the RPCs, but in the endcap region.

A large technical effort within the muon system, especially for the MDTs, is required for good alignment. The maximum chamber length is six meters. Being built from aluminum, the chambers react to temperature changes, hence monitoring of the muon chambers is needed to measure the alignment actively. A laser alignment system, based on interferometry, is used to measure deformation within each muon chamber individually.

3.2.4 Magnet system

To measure charged particle momenta within the central tracker and the outer muon system, both need to be immersed in a strong magnetic field, bending the particles tracks. The ATLAS magnet system splits into a solenoid system for the tracker and a toroidal magnet system for the muon spectrometer.
The solenoid magnet contains just the tracker system, giving good momentum resolution up to very high momenta. Putting the calorimeter outside the high-field solenoid reduces jet widening due to Lorentz forces on charged particles. Thus the ATLAS calorimeter is expected to deliver a good spatial resolution, not just of jets as a whole, but also of a jets sub-structure.
The toroidal field for the muon chambers is provided by an air-core toroidal magnet, built from a barrel section and two endcaps. The field strength varies between 0.5 and 2 T throughout the magnet system.

3.2.5 Trigger and Data Acquisition System

The systems described above deliver a total of about 100 million channels. Reading out each of those for every event and assuming binary information content, ATLAS would have to deliver a total readout bandwidth of about 363 Terabyte per second. This being impossible nowadays, ATLAS reduces the amount of data recorded in three trigger stages, Level 1, 2 and finally the Event Filter.

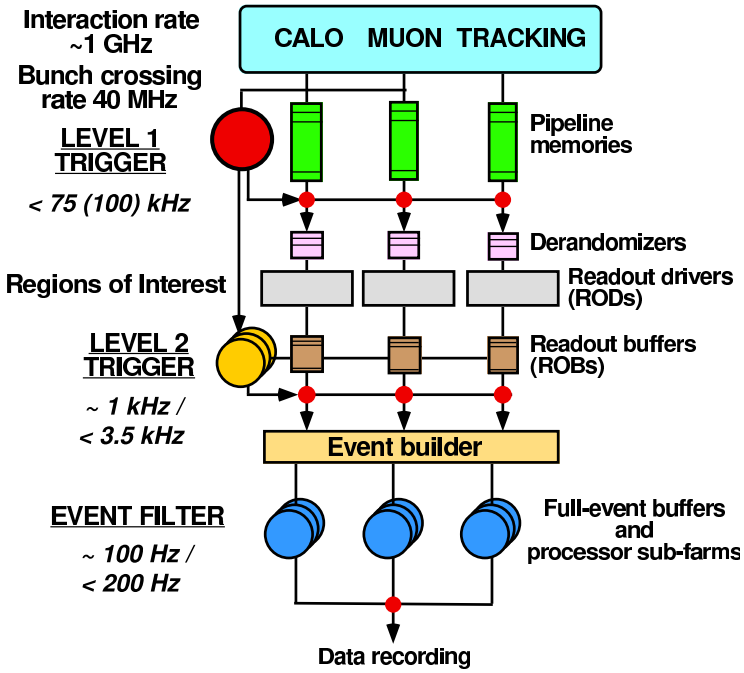

Figure 3.6: *The ATLAS TDAQ concept [AC98]: Data is generated in the frontends at 40 MHz bunch crossing rate with a maximum interaction rate of 1 GHz. Coarse data are checked by the L1 trigger system. In case a trigger is fired, data is read out from the frontend buffers and stored into the readout buffers. From those regions of interest are retrieved and processed by the L2 trigger, which then decides about events that should be fully processed. Eventually the Event Filter (L3 trigger) decide to store information to disk. The total number of about 30 million filled bunch crossings per second is reduced to about 100 events per second which get stored.*

Readout Structure

Shown in Figure 3.6 is the data acquisition structure of ATLAS. Each detector system has to implement a frontend to read out data. Systems included into the level 1 trigger need to deliver that data at full interaction rate, but not necessarily at full granularity, to the central trigger system. Selection of events is the performed in the Central Trigger Processor (CTP), dropping 99.8 % of all events and reducing the readout rate for all systems to a level 1 triggered rate of 75 kHz.

In case an event is selected by the Level 1 system, it is read out from all detector subsystems into generalised readout buffers. Multiple regions of interest can be defined by the L1 trigger and are then processed by the Level 2 Trigger system to perform a more detailed event selection. The output rate here is at maximum 3.5 kHz.

Finally, Level 2 triggered data is fully processed by the Event Filter system, which then writes out selected events to disk at a maximum rate of 200 Hz.

Timing, Trigger and Control

ATLAS defines an experiment wide Timing, Trigger and Control (TTC) system, which allows the subdetectors to synchronise with a global clock source, and distributes common commands. The system splits ATLAS into multiple partitions, each of which can decide to run with the global ATLAS system or in a local mode, allowing independent calibration of subdetectors. The target clock source eventually is the LHC delivering a bunch crossing synchronous clock (BC), whilst commands within the TTC system are typically generated by the Central Trigger Processor, e.g. a level 1 trigger.

Level 1 Trigger

Reducing the initial event number happens based on the Calorimeter and Muon systems. Both deliver full rate output from selected parts into a subsystem trigger processor, that can fire different types of triggers. The muon system can quickly detect high energetic muons, and, depending on thresholds selected, request an event to be triggered, e.g. a muon with more than 30 GeV transverse momentum. The calorimeter system splits into several subprocesses, electromagnetic and hadronic measurements, as well as measurements of total transverse energy, missing transverse energy and jet energy. All of these can also be basis for a selection of a particular event, e.g. a jet in the electromagnetic barrel with more than 30 GeV. These subsystem trigger criteria are handed to the central trigger processor, who can prescale them and selects to raise an L1 accept signal or not. An illustration of the L1 system is given in Figure 3.7.

High Level Trigger

The high level triggers serve a much better event selection process, e.g. based on tracking information which can pinpoint B-hadrons and multiple interactions with certain limits. Both higher level trigger stages, level 2 and event filter, are software based implementations within PC-farms. An easy update and upgrade is thus possible within these systems.
The level 2 trigger system only processes Regions Of Interest (ROI). Those are defined by the L1 trigger system, for example due to a high energetic jet or muon found. Processing a single ROI takes between 10 and 100 ms on a level 2 processing unit, depending on the content of the ROI. Therefore the level 2 trigger system is built from multiple CPUs, to be able to process all 75000 events incoming per second.
The same number of CPUs as are present in the level 2 trigger system are also used within the Event Filger (EF). Here all event data is processed, creating even better defined cut criteria.

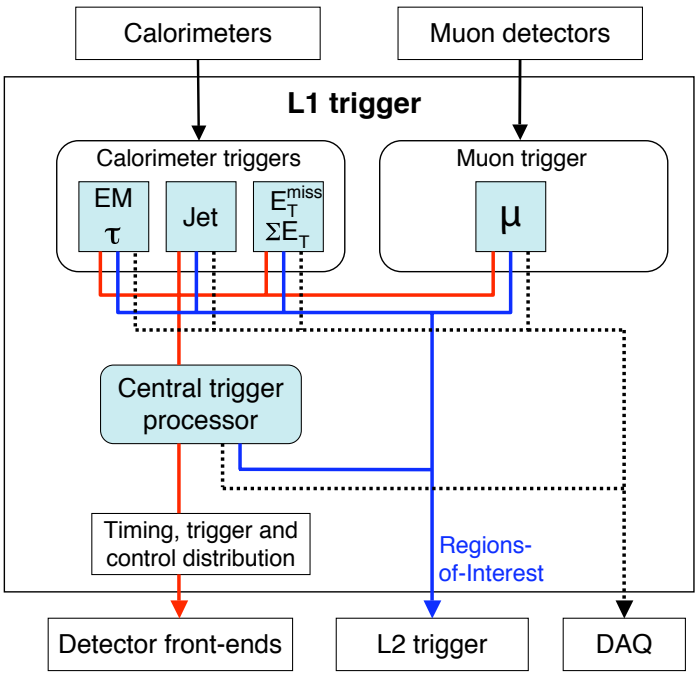

Figure 3.7: *A schematic view of the ATLAS Level 1 Trigger system, as given in [A+08a]*

4 ATLAS Pixel Detector

The innermost subdetector of ATLAS is the Pixel Detector, shown in Figure 4.1. It consists of 1744 modules, glued to a mechanical and thermal support structure. The Pixel Detector delivers a total of more than 80 million channels. It comes in a barrel and two endcap sections, each consisting of three layers. By construction it delivers at least three spacepoints for each charged track with $|\eta| < 2.5$ originating in the centre of ATLAS. Given the spatial resolution of $115\,\mathrm{x}\,14\,\mu$m in z and $R \cdot \varphi$ it serves for highest vertex resolution, down to less than $12\,\mu$m in x and y, allowing for B-tagging as well as good separation of multiple primary vertices.

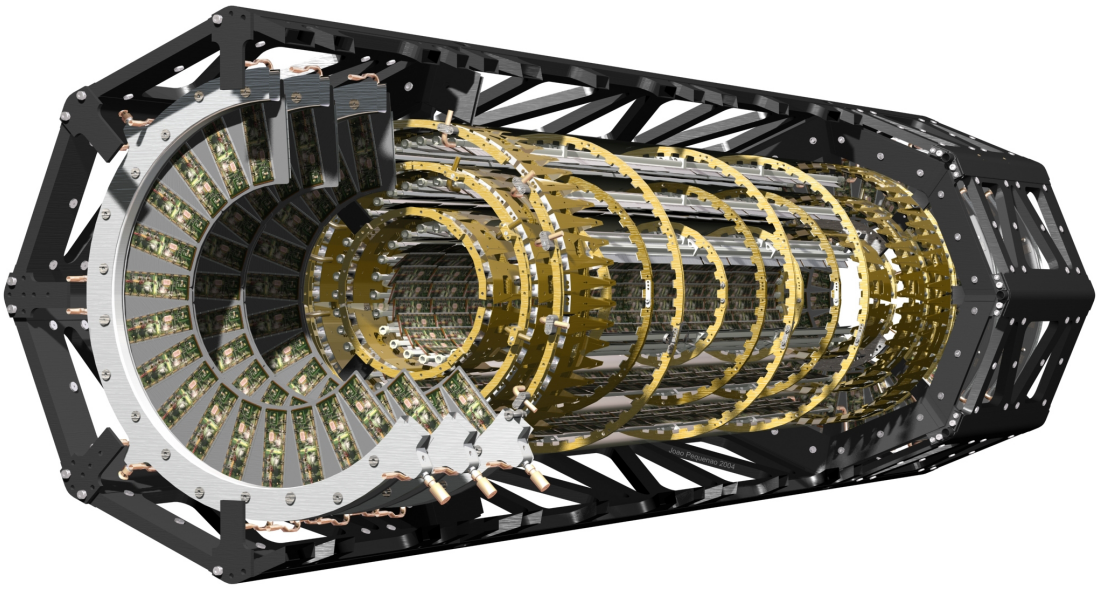

Figure 4.1: *The ATLAS Pixel Detector [AHP]. The 3D model has parts of the detector removed, so that the Barrel and Disk structure is visible*

This chapter will first introduce the ATLAS Pixel Detector, starting with a system outline, it will then provide a module description to explain how very high resolution data is generated within the detector. After a short summary about semiconductors in optical data transmission, the chapter will describe the Pixel Detector readout system and eventually the detector control system.

4.1 System Overview

The ATLAS Pixel Detector is mounted into a 7 m long structure, including 8 service quarter panels (SQPs, 4 per side) to deliver connections for the Pixel Detector modules. These connections are shortly outlined, following Figure 4.2.

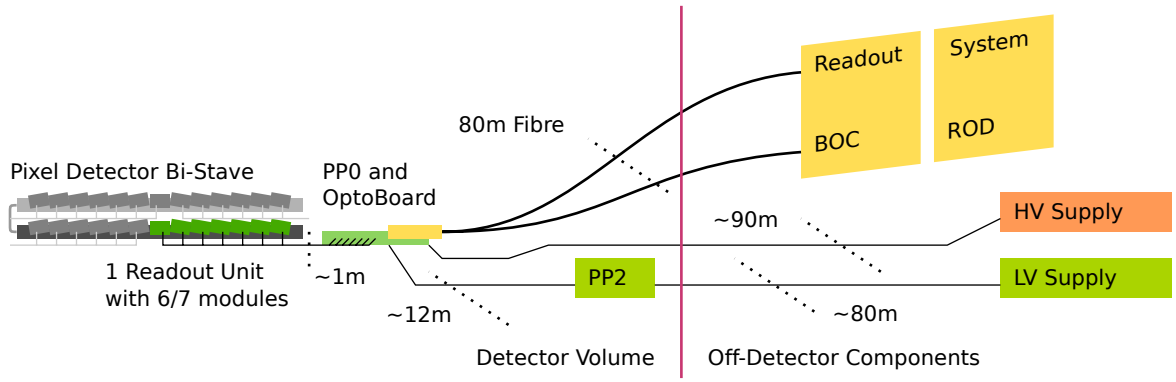

Figure 4.2: *A Pixel System outline, showing readout and supply connections through the most important patch panels. To simplify the picture, monitoring and interlock components are not shown, as well as passive patch-panels and the cooling system.*

The Pixel Detector basic powering and readout unit consists of either 6 or 7 detector modules in the barrel, respectively 6 modules in the discs. These connect to a single Patch Panel 0 (PP0) mounted on the service quarter panels. The patch panel then provides connections for low voltage, high voltage and readout. The former two come in via electrical cables, whilst the latter is delivered from an optoboard. The optoboard receives and transmits optical data from and to the off-detector components through 80 m of optical fibre.
The low voltage supply line needs regulation close to the Pixel Detector package, due to cable resistances. This is done at PP2, which is located 12 m away from the Pixel Detector, still inside the ATLAS Detector. The high voltages on the contrary can be supplied directly from the off-detector electronics, as the very low current suffers no significant voltage drop. The following sections will move through Figure 4.2 from left to right, focussing on the readout connection.

4.2 The Pixel Detector Module

An ATLAS Pixel Module (c.f. Figure 4.3) is built from a single silicon sensor that is bump-bonded to a total of 16 frontends. Bump bonding allows each Pixel cell within the sensor to be connected to a corresponding frontend cell, back-to-back. The backside of the sensor connects to high voltage bias, but is also the platform for a flexible printed circuit

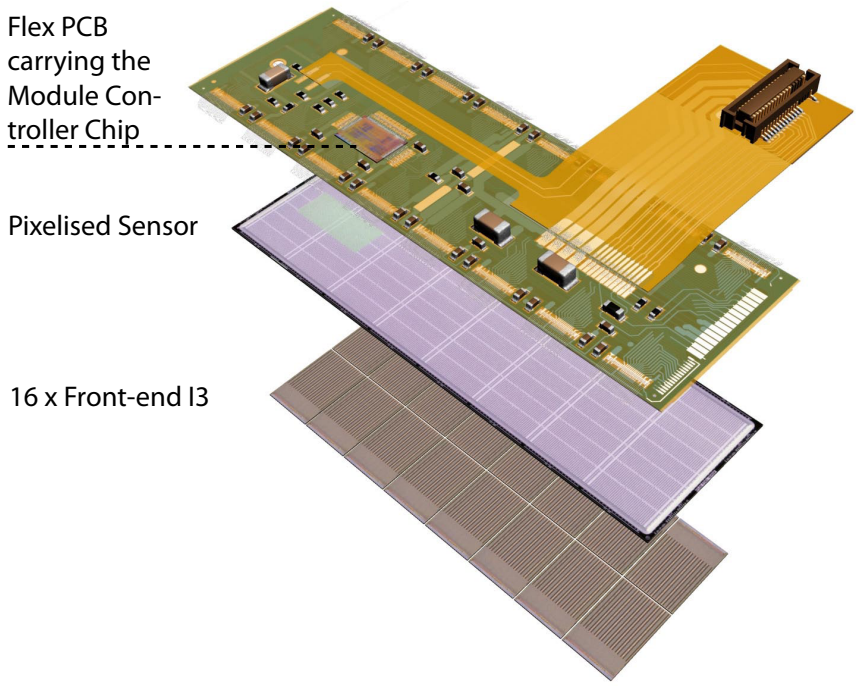

Figure 4.3: *An ATLAS Pixel Detector module in vertically exploded view [Dob07]. Visible are the 16 frontends at the bottom, the sensor in the centre and the flexible printed circuit board at the top, carrying the module controller chip and the pigtail with the Type 0 connector.*

board (flex-pcb). This flex-pcb connects to the frontend via wirebonds at the long edge of the sensor. In addition it carries a module controller chip, which enables the readout system to communicate with the 16 frontend ICs.

4.2.1 The Pixel Sensor

The Pixel Detector active volume is composed of a 250 μm thick oxygen enriched n^+ in n silicon sensor. The sensor layout contains a total of 47232 pixels per sensor, split in 144 columns and 328 rows. To read out the full sensor, 16 frontend chips have to be connected to it, each contacting 164 rows and 18 columns, 2952 pixel cells per frontend. The pixel rows in the sensor center are not directly accessible, as a gap is needed between the different frontend chips. The sensor layout still allows for full coverage over the sensor area, by either implementing long pixel cells or via a pairing scheme.
2496 pixel cells per frontend are the standard 400x50 μm^2 in z and $R \cdot \varphi$ and cover the inner area of each of the 16 frontend chips. A long pixel column at each side of a frontend covers the long edge of the frontend IC with pixels sized 600x50 μm^2. Coverage of the inner area is done using an interconnection scheme of pairs of pixels in the edge area,

occupying a total of 8 rows with pixel pairs. Identification of which pixel was hit can only be done using neighbouring pixels within a cluster here. Therefore the interconnection scheme is a pairing with interleaved single pixels. The pixels using connections are called ganged pixels, whilst the standard pixels filling the area in between are called inter-ganged. Different pixel types deliver different sensor areas and capacitances. Therefore the special pixels incorporate a different leakage current compensation within the frontend and differ in electronic noise performance at the pre-amplifier.
The sensor is built to withstand a total fluence of $1 \cdot 10^{15}$ 1 MeV neutron equivalent. Before irradiation and with a sensor bias voltage of 150 V it delivers 19.6 ke[a] per Minimum Ionising Particle (MIP). After irradiation up to its full design dose, the sensor still delivers about 16 ke, using an increased bias voltage of 600 V.

4.2.2 Pixel frontend cell electronics

The ATLAS Pixel Detector sensor cells are read out using a FrontEnd I3 (FE I3) integrated circuit [P+06]. 2880 pixel frontend cells are bump bonded to the sensor, enabling the frontend chip to read the total 2952 pixel cells from the sensor. All frontend cells are equal as for their readout scheme. A minor difference exists in circuits for leakage current compensation, as the outer and ganged pixel cells need a larger size compensation to keep the preamplifier inputs low.
The frontend cell is shown in Figure 4.4. Charge gathered at the bump-bond pad decreases the voltage level on the preamplifier input next to it. The output of that preamplifier drives a feedback current (above the preamplifier) that removes the charge and thereby clears the pixel again. Also a threshold is applied to that preamplifier output, which allows for distinguishing high and low charge and reducing noise. The discriminator output eventually strobes timestamps into registers for rising and falling edges, not shown within this picture. Both threshold charge and feedback current can be adjusted by local and frontend wide settings. Whilst the threshold is offset by the four bit global component (the GDAC), the feedback current is regulated in range by the global adjustment (not shown in Figure 4.4). A total of 14 configuration bits exist locally for each pixel cell. Seven bits adjust the threshold locally, three bits the feedback current. The remaining four configuration bits allow for switching on and off functional blocks:

- Pixel Enable: Switches on the preamplifier.
- Pixel Mask: Masks the digital output of the discriminator.
- Hitbus: Connects the pixel cell with the global hitbus[b].

[a] Peak value of the the delivered charge per minimum ionising particle.
[b] The hitbus is a system to allow the frontend to trigger itself as soon as a pixel shows a hit.

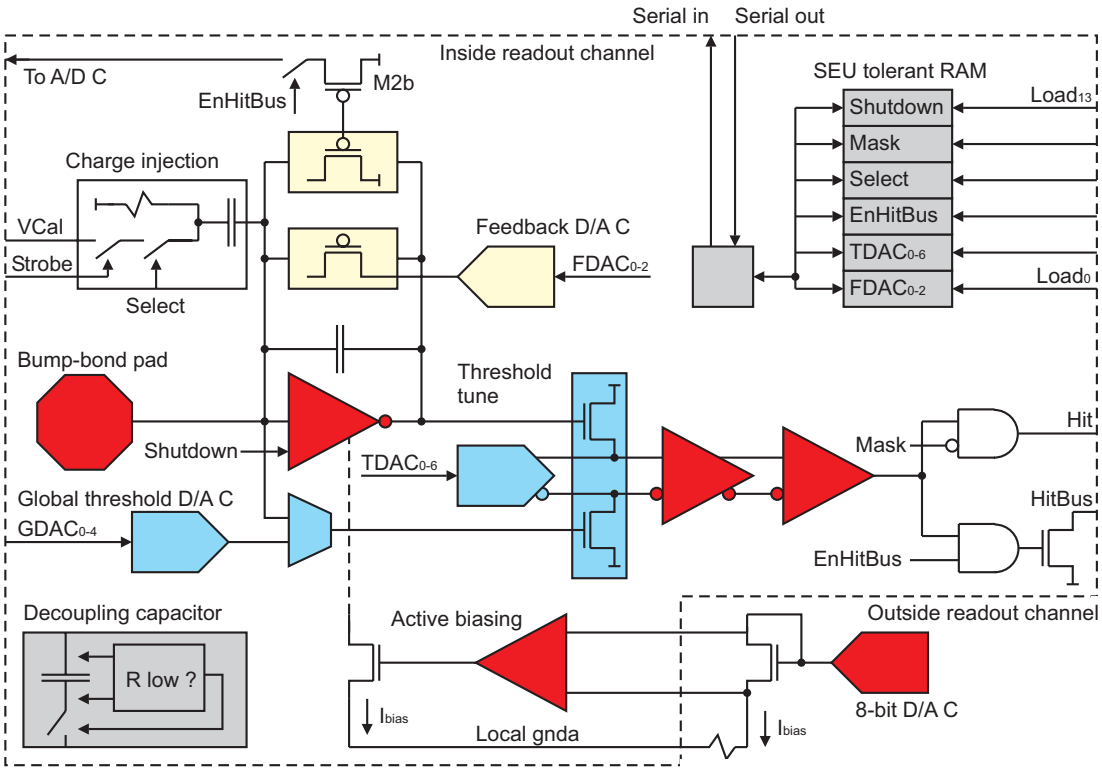

Figure 4.4: *The pixel cell electronics, taken from [A+08b]. Shown are the Pixel connection pad with preamplifier, charge injection and current feedback at left, and discriminator and hit logic at right.*

- Select: Enables charge and digital injection circuitry.

Using these bits, a pixel cell can be included in or excluded from calibration scans or data taking.
Scans in particular are run using the Select bit. It connects the strobe line received from the Module Controller Chip (MCC) with the cells internal charge injection circuitry. Thereby, charges between 0 and up to 200 ke can be injected before the preamplifier. In case of digital injection[c], the strobe line replaces the discriminator output within the pixel cell.

Timewalk

The Pixel Detector preamplifier is constructed such that it always uses nearly the same time to rise to the peak voltage value. Therefore, high charge values have a very steep initial slope before reaching the current feedback, as shown in Figure 4.5, they rise above threshold faster than small charges, which take a longer time. The effect of low charges

[c]Digital injection is a frontend configuration option.

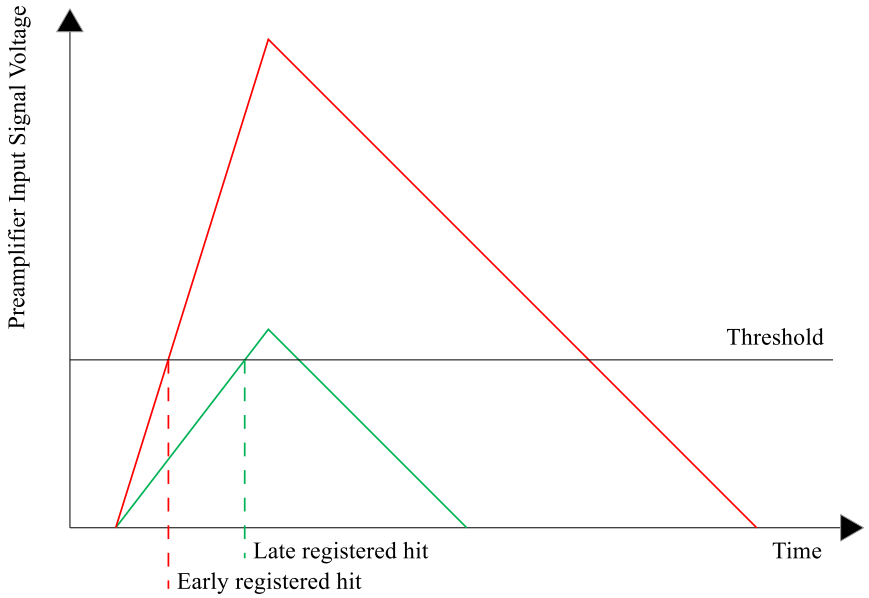

Figure 4.5: *Illustration of timewalk in the Pixel Detector: A hit with low charge rises slower to the peak voltage than a large hit, thereby shifting the registration time.*

being registered later than large charges is called timewalk. The size of the effect can be as large as multiple bunch crossings.
A monitoring histogram from an early ATLAS run is given in Figure 4.6. The Pixel Detector was configured to sample a total of 8 bunch crossings after receiving a trigger signal. The data is sorted within the monitoring, depending on which bunch crossing it was sampled in, relative to the trigger bunch crossing. The beam induced hits are, timewise, placed in the 3rd bin (bunch crossing number 2). The majority of hits with a TOT value below 8 (about 8 k electrons) is still recorded in the 4th bin or even beyond.
As is shown, effects are very harsh in the low charge regime. Readout only takes into account single bunch crossings at highest luminosities. Therefore, to compensate the timewalk, the frontend allows to copy low charge hits into a previous bunch crossing. This helps increasing resolution again, but puts higher load onto the readout system.

4.2.3 Frontend

All Pixel Detector cells are read out using a double column logic, cf. Figure 4.7. Pixels from even and odd columns have a common connection to an End Of Column (EOC) logic. This EOC logic freezes the readout requests as soon as a single Pixel has data available, and then reads out all double column data from high to low row numbers. An arithmetic unit then encodes rising and falling edge counters into a BCID and a Time Over Threshold (TOT) value and stores them in the end of column buffer, a content adressable memory.

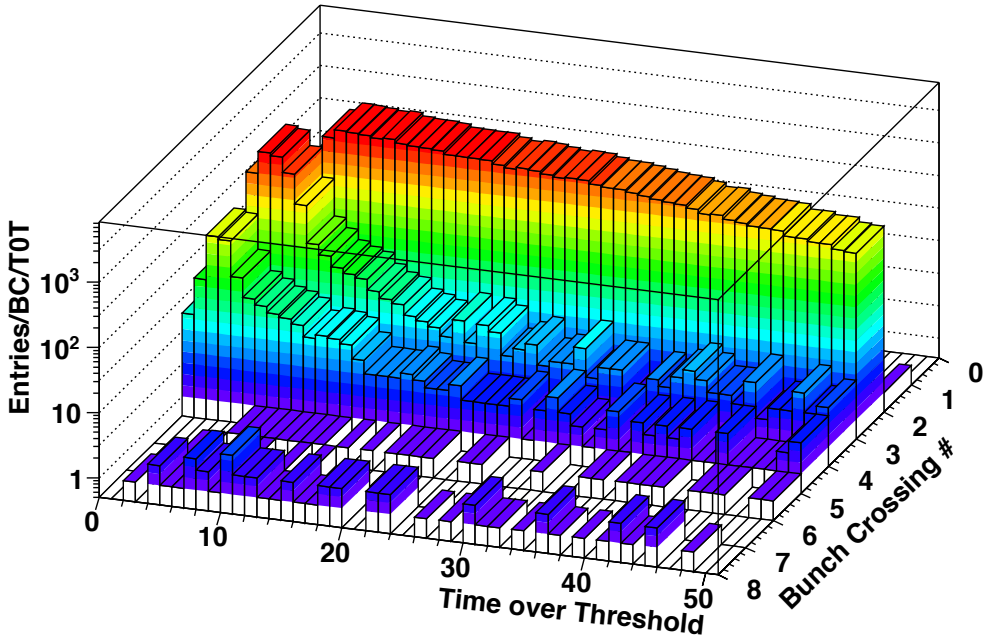

Figure 4.6: *TOT Monitoring from a 7 TeV run (152409). The ATLAS Pixel Detector was configured to record 8 bunch crossings, with the beam induced events being timed into the 3rd recorded bin. Visible is a shift of hits with low TOT values into the 4th and 5th recorded bin, whilst they are missing in the 3rd bin that only records high TOT values.*

A level 1 trigger signal arriving at the frontend chip marks a BCID for readout, which is then used for addressing the content addressable memory. As the trigger signal arrives only after a while, an offset between the BCID of L1A arrival and the BCID being read out can be configured within the frontend chip. Data which is older than the recent BCID minus that offset is removed from the EOC buffers to free space for new data arriving from the pixel cells. This logic gives the freedom to decide between long readout with higher TOT precision, or short readout allowing for higher per-pixel occupancies.

4.2.4 Module Controller Chip

The Module Controller Chip [B+02] (MCC) is the communication device of a module. It connects to each of the 16 frontend chips and provides the external command and data interface(s) of the module. Running on the 40MHz LHC bunch crossing clock, it can deliver data to the off-detector electronics in four configurable transmission modes, using two transmission lines:

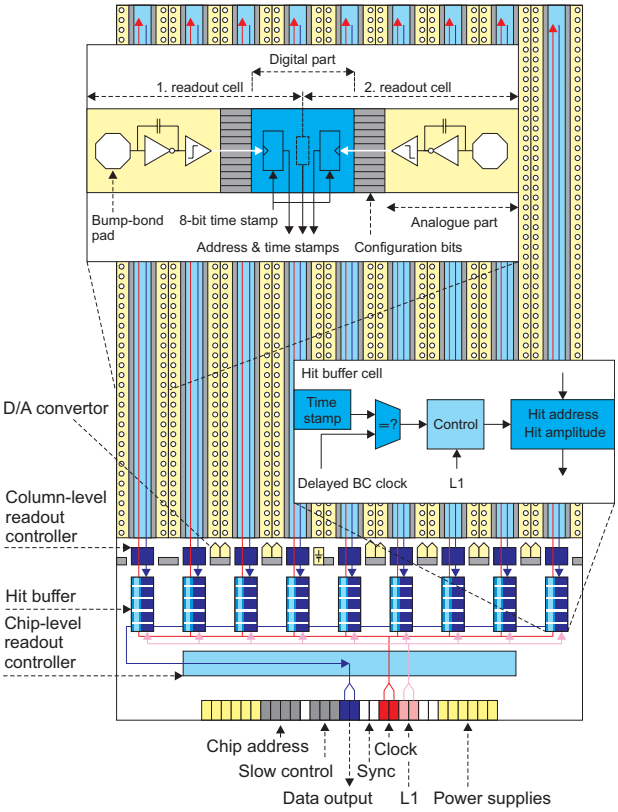

Figure 4.7: *The Frontend I3 column drain architecture, shown with two zoomed Pixel cells sharing the double column readout logic. [A+08b]*

Mode	Description
Single 40	Delivers redundant 40 Mbit/s output data on both links
Dual 40	Delivers two streams of 40 Mbit/s output data with one carrying odd, the other carrying even bits of a combined 80 Mbit/s data stream
Single 80	Delivers redundant 80 Mbit/s output data on both links
Dual 80	Delivers two streams of 80 Mbit/s output data with one carrying odd, the other carrying even bits of a combined 160 Mbit/s data stream

The speed modes are foreseen for different data rate needs throughout the detector layers. The Pixel Detector B-Layer will be read out using two streams at 80 Mbit/s. The other layers and the discs are only given a single transmission line for readout, running either at 80 Mbit/s in Layer 1 and the discs, or at 40 Mbit/s in Layer 2.

Configuration and Calibration

The MCC contains a set of registers to allow module wide configuration, forwarding of configuration streams to selected frontends and eventually strobing pulses into frontends to

perform calibration scans. Using these pulses, the threshold, as well as timewalk behaviour of the frontend can be measured. As the pulse going to the frontend, only the timing of the pulse can be controlled by the MCC. The frontend on the contrary defines the analogue amplitude of the pulse.
Additional to the analogue calibration of the frontend, the MCC holds registers and buffers that allow to calibrate the data transmission between the module and the off-detector electronics. FIFOs that hold data coming from the module can be set to arbitrary values and triggered to be read out to the off-detector electronics, allowing a simple check for data transmission errors. At an even simpler stage, the MCC can generate a clock-like pattern on its output lines.

Data Taking Mode

In data taking mode, the MCC strobes an L1A line for each activated frontend chip on arrival of an L1A pattern[d] at the MCC input, triggering the frontend chips to push data from a certain BCID into the MCC input buffers. Event building then starts from the first frontend chip (FE-0) and continues to the last until all data from a single BCID is sent. In case multiple BCIDs are to be read out by the frontend chip, the MCC generates multiple event data packets with different BCIDs but the same L1ID.

4.3 Optical Semiconductor Devices

This section will give a short introduction to optical semiconductor devices used within the optical communication of the ATLAS Pixel Detector. This will be the electro optical converters, as well as the circuitry used to control these or read their output.
The section is not an integral part of the Pixel Detector, yet it is needed to explain future behaviour.

4.3.1 PiN Diodes

PiN diodes are semiconductor diodes with three layers. They can be utilised as photoreceivers. Between the normal p- and n-type doped semiconductors is an intrinsic layer, that receives a very low dopant concentration of either type. The intrinsic layer affects a PiNs photoreceiver characteristics in two ways:

- Capacitance: Due to the intrinsic layer, a depleted PiN diode has a very large distance between its two poles, giving a low capacitance, even at low voltage biases. Capacitances of PiN diodes utilised within this thesis are in the order of 0.1 pF.

[d] The L1A Bit pattern is 11101 and is detected even if 1 bit is not correctly transmitted

- Conversion: The intrinsic layer is part of the active conversion medium, and thus contributes to the total efficiency of the diode. A typical responsivity of a PiN diode used within the following chapters is 0.6 A/W.

PiN amplifiers

PiN diode output signals in telecommunication are typically at levels of a few hundred micro Amperes. They need to be amplified to allow for threshold application, in particular if the signal swing is very low. Several concepts for amplifiers are available, some of which can only be used to amplify signals with a minimum switching frequency given. Two are presented here, the former reflecting the currently used system, whilst the latter gives insight on future components.

- Limiting Amplifier: A limiting amplifier is an operational amplifier with no or very little feedback into the input stage. Connecting a biased PiN diode through a resistor to ground, gives a voltages drop across the resistor, corresponding to the current flowing through the PiN diode. This current is routed into the positive input of an operational amplifier, where the negative input is given by a threshold setting. The circuit has a major issue, that is the RC component at its input, given by the PiN diode capacitance and the input resistance. This gives a maximum switching frequency, that can still be observed with a fixed threshold application.

- Transimpedance Amplifier (TIA): Circumventing the maximum frequency limitation of the resistive amplifier is done by building the amplifier such that the PiN diode is kept at the same bias voltage, independent of the current that flows. This needs a fast operational amplifier with negative feedback into the PiN diode. If the PiN diode is given a negative bias, it has to be connected to the negative input of the OpAmp. The positive bias will then be given by the feedback, delivered from the OpAmp output. The positive input can be defined to give the bias level of the PiN diode.

Transimpedance amplifier feedbacks are typically designed with an inductance and a capacitance to reduce the input resistance, optimising it for a particular frequency range. To ideally couple into later digitising amplifier stages, a capacitive coupling should then be used. This implies a differentiating behaviour of the amplifier circuit and thus generates a lower cutoff frequency, at which the later amplifier stages have neither positive nor negative input and therefore generate noise. Using a digital integrating stage as the later stage amplifier, this noise generation can be suppressed, and the TIA can be utilised to maximise input frequency response.

4.3.2 Lasers

Light Amplification by Stimulated Emission of Radiation (LASER) has first been realised in the late 40s. Initially postulated by Einstein in his 1917 *Zur Theorie der Strahlung* was the fact that an atom could not only be raised from one energy level to another by an incident photon (thereby absorbing it), but also be lowered by that exact amount of energy[e]. This lowering results in what is called stimulated emission of a photon with the exact same energy, direction and phase of the incident photon.
Using a material with only two energy levels typically results in a low efficiency of the laser. That is due to the higher energy level spontaneously emitting a photon very quickly, which does not leave enough time for an incident photon to pass by and stimulate emission. Most lasers nowadays are based on materials with three or four energy levels included in the laser process. Both types have a pump state, that is driven by input energy from the outside. This pumped level directly emits light, decreasing into a meta-stable state which has a much longer lifetime than the pumped state. As soon as enough atoms are in that meta-stable state, the laser process can start, by directing photons with the correct energy through that medium. Within laser constructions this is typically done by enclosing the pumped volume within two parallel mirrors. As soon as the metastable level starts spontaneous emission along the mirror axis, the volume will be passed by photons with the correct energy multiple times. Having one of the mirrors only 99% active, allows to emit light, while keeping up the stimulated emission.

Diode Lasers: Diode Lasers are a specialised type of Laser. The active volume is a light emitting diode. Recombination of electrons and holes within its junction region leads to spontaneous emission of photons. To enforce stimulated emission, mirrors are placed around the diode junction. The type of mirrors and their placement varies for different types of lasers. Within this thesis, Vertical Cavity Surface Emitting Lasers (VCSELs) are being used. These are built from a vertical diode structure, which is embedded into multiple horizontal layers of Bragg reflectors. A schematic outline is shown in Figure 4.8 depicting the different layers within a VCSEL. It shows the slightly emphasised active layers, called the quantum wells. Three quantum wells were chosen in this schema to match up with the focused ion beam image at right, which shows a VCSEL as used within the context of this thesis. The quantum wells are the optically active region, delimiting the emitted wavelength by their thickness, allowing fewer energy levels to be confined within. Another limitation of the wavelength is given by the bragg reflectors, above and below the active region. The aperture limits the lasers angular emittance.
Due to their small active and total volume, diode lasers and VCSELs in particular have very small turn-on and turn-off times, of typically less than a nanosecond. Yet the small

[e] Assuming the energy levels available within the atom fit that of the incident photon

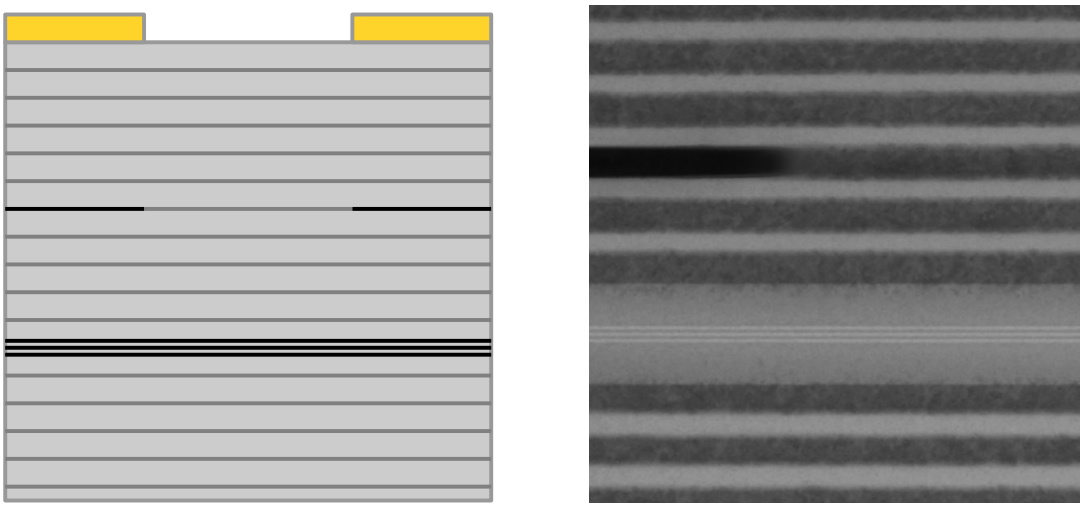

(a) Schematic view of a VCSEL

(b) Focused Ion Beam (FIB) image of a multi quantum well (MQW) VCSEL

Figure 4.8: *The schematic side-view of a quantum well VCSEL is given in (a). The same view, but zoomed into the aperture and quantum well region is given in the photograph of a multi quantum well laser in (b).*

active volume compared with the large reflector volume also causes the laser to be less effective.

Laser Drivers

Diode lasers in general are driven by current sources. Two different currents are typically available, modulation and bias current. The modulation current is to switch between on ("bright") and off ("dim") states of the laser. The bias current should keep the laser just below lasing threshold as shown in Figure 4.9, so that turn-on times become smaller. The modulation current then just needs to lift the laser current above threshold, which also allows one to have the laser working at a lower amplitude, yet delivering a high speed output. Recent high-speed laser drivers also include a pre-emphasis circuitry. This feeds an additional small modulation current through the laser in case a transition from dim to bright is performed, which decreases switching times even more.

4.3.3 Semiconductor Lifetimes

This section will give a short outline on how to define and measure semiconductor devices lifetime. It will be important to understand the general idea of lifetime within the later chapters, thus a little theory background is given here.

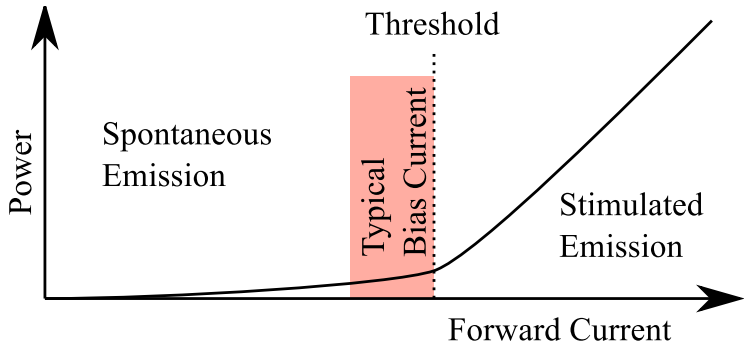

Figure 4.9: *A typical laser L/I curve. Shown is a threshold, as well as the typical bias current region.*

A component's time dependent failure rate splits into three different subsections:

- Early Life Mortality: After a production, some level of internal failure in semiconductor devices is expected. This will cause the corresponding device to fail very early.
- Random Error Period: Assuming the early life mortality is over, a constant, low, and random failure rate is expected in standard devices. Failures in this period should neither be production, nor lifetime dependant, but instead describe what is called the Mean Time Before Failure (MTBF)
- Wear-out Period: When devices have reached their lifetime, they will start failing heavily. The error rate thus increases again.

The curve that describes this error rate is referred to as bathtub curve. It is typically described by three or four different parametrisations of a Weibull Distribution:

$$f(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{(\beta-1)} e^{\left(\frac{t}{\eta}\right)^{\beta}} \tag{4.1}$$

Where t is a Weibull distributed time, β is the shape parameter and η is the scale parameter. Using $\beta = 0.1$, 0.8 and 10, three different Weibull functions have been plotted in Figure 4.10 to generate the sum of them, the "bathtub" curve. As soon as this curve is known an MTBF can be retrieved from the random error period, and a lifetime from either the MTBF, or the wear-out time, depending on which hits earlier[f].

[f]Industrial devices are typically constructed such that their wear-out time is longer than their MTBF.

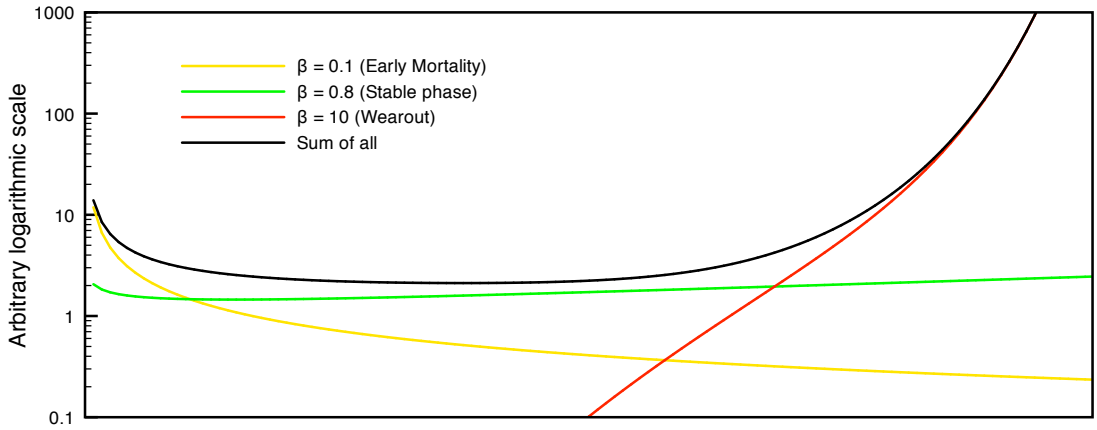

Figure 4.10: *A bathtub curve describing a devices failure rate as a function of time, generated from three different Weibull distributions ($\eta = 1$, $\beta = 0.1$, 0.8 and 10). The x-axis refers to a linear time, but has intentionally no scale given.*

Accelerated Ageing

To quickly obtain a components lifetime, it is necessary to accelerate the process of device ageing using different kinds of stress. That can for example be ambient temperature or forward current in a laser diode. The stress induces early failures, which might otherwise only show up after years. A proper understanding is then needed to know the acceleration factor.

Temperature dependent ageing is most times modelled using the Arrhenius equation:

$$A_f = e^{\left(\frac{E_a}{k_B}\left[\frac{1}{T_R} - \frac{1}{T_T}\right]\right)} \tag{4.2}$$

Where A_f is the acceleration factor, E_a is the activation energy of the material in use (in eV) and k_B is the Boltzmann constant ($\approx 8.617 \cdot 10^{-5}$ eV / K). The temperatures are measured at a reference T_R and a tested T_T sample with typically increased temperature.

The only unknown factor in this equation, the activation energy, needs measurement. It can be derived from two accelerated ageing samples at different temperatures. Lifetime data gained from one and the other gives the acceleration factor, from which the activation energy can then be calculated.

4.4 Optical Readout Links

The Pixel Detector is read out using custom VME[g] based hardware that connects to the detector via an optical connection. The full optical link as depicted for a single channel in Figure 4.11 is composed of a total of three electro-optical components, the optoboard on-detector and the TX- and RX-plugins off-detector. The connecting fibre is a mixture of 12 m radiation hard fibre within the ATLAS calorimeter volume and about 65 m of radiation tolerant fibre. Two different connections are placed within this length. The first is a pluggable connection at the end of the Pixel Detector volume. A second connection is a splice to establish the transition from radiation hard to radiation tolerant fibre.

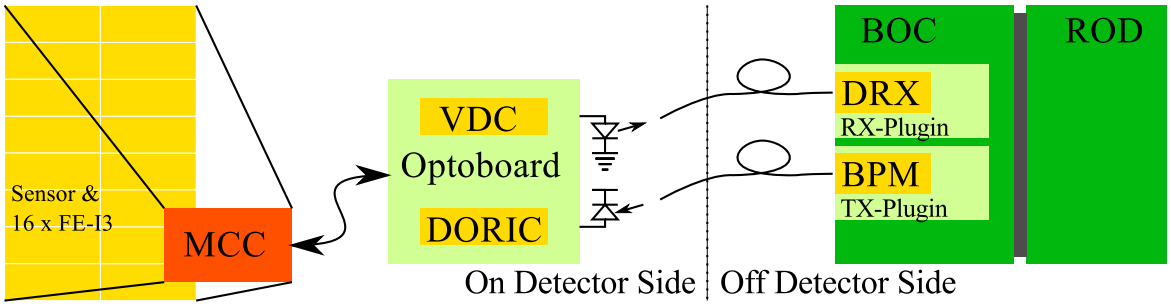

Figure 4.11: *A schematic sketch of the ATLAS Pixel Detector readout link and readout system, reaching from the frontend chips to the Readout Driver. Only a single link is shown here.*

The on-detector electro optical converter, the optoboard [Nde07] as seen in Figure 4.12, is a transceiver component. It contains both, transmitters for 8 or 16 channels, depending on its type, and receivers for 8 channels. The receivers are constructed to receive an encoded stream from the off-detector transmitters and decode clock and data from it. The transmitters are pass-through components, receiving LVDS[h] data from the detector and converting it into light output.

A single optoboard connects to half the modules of a stave[i] (either 6 or 7 modules) or a single disc sector (6 modules). Two different types of optoboards are produced, holding either 8 or 16 transmitter channels. 16 channels are needed to establish 160 Mbit/s transfer rate in two 80 Mbit/s channels per module for the B-Layer. The rest of the detector is read out using either 1x40 or 1x80 modes (c.f. table 4.2.4), hence only given 8 transfer channels on their Optoboards. As the system utilises a maximum of 7 module

[g] Versa Module Eurocard, originally designed for the Motorola 68k processor family, standardised later as ANSI/IEEE 1014-1987

[h] Low Voltage Differential Signaling - A differential signal standard

[i] A stave is the basic mechanical unit of the Pixel Detector in the barrel section

channels per optoboard (14 in case of B-Layer Optoboards), the 8th channel (and 15th/16th correspondingly) are disabled to not introduce noise.

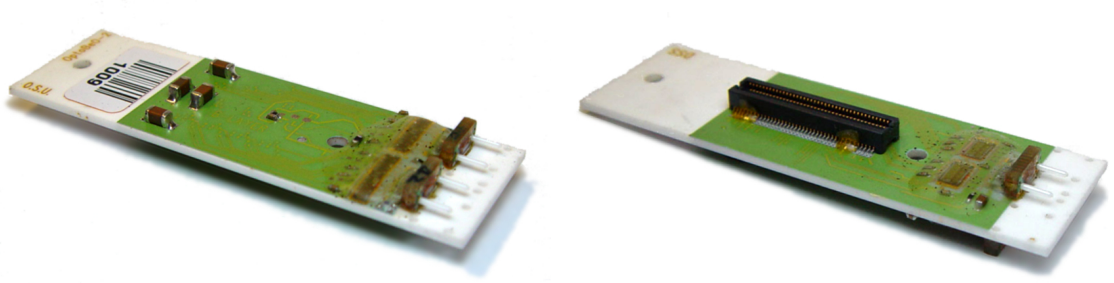

(a) Optoboard Top Side (b) Optoboard Bottom Side

Figure 4.12: *A B-Layer Optoboard, built from beryllium oxide. Shown is the top-side in (a), with mounted VCSEL Driver Chips (VDC) under epoxy (right hand side), connecting to two laser arrays (at the very right). (b) shows the bottom side with two DORIC 4A (Digital Optical Receiver IC) circuits covered in epoxy and a PiN array (same positioning as the top side components).*

Off-detector components are split into the transmitter, the TX-plugin shown in Figure 4.13(a), and the receiver, the RX-plugin shown in Figure 4.13(b). The transmitter encodes twelve data streams with a clock using BiPhase Mark encoding (BPM, see Section 4.4.1) and is able to adjust phases and mark to space ratio of each of the twelve transmission lines. Electro-Optical conversion is done by an external laser array with 12 lasers. The output light power of each transmitter channel is controlled individually via external voltages coming from Digital to Analog Converters (DAC). The receiver converts light into currents using a 12-way PiN diode and applies a threshold to these currents. The threshold can be adjusted between 0 and 255 uA by an external DAC.
12-channel modularities within the optical plugins are inherited from the semiconductor tracker, which uses the same off-detector electronics for readout. The ATLAS Pixel Detector initially had its TX- and RX-plugins produced with 8-channel arrays, to reduce the cost per plugin.
The integrated circuits, as well as the encoding in use within the aforementioned components, will be described in the following subsections. The description splits into the timing, trigger and command connection, and the data connection. Eventually, the off-detector readout system electronics will then follow to complete the readout system description.

4.4.1 Timing, Trigger and Command

The Timing, Trigger and Command (TTC) transfer is established between two Integrated Circuits, the BPM-12 sending data out and the DORIC-4A, receiving that data. The BPM

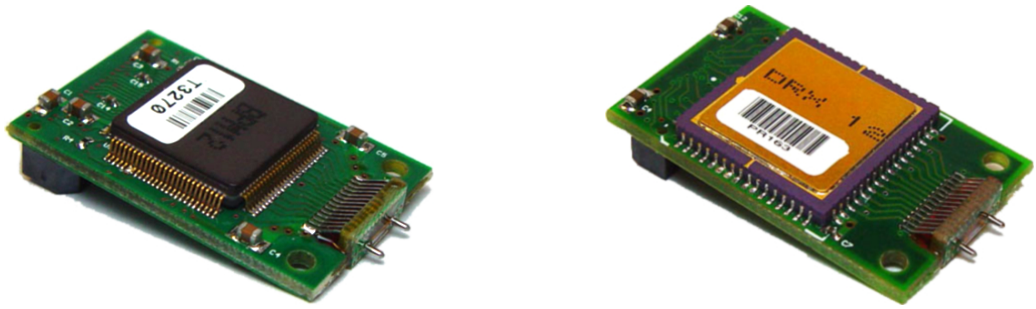

(a) TX-plugin (b) RX-plugin

Figure 4.13: *The off-detector components of the Pixel Detector optolink: The transmitter, including the driver IC and a laser array, and the receiver including the an amplifier and threshold IC.*

drives data and clock, encoded into a single stream, through a laser. Infrared light then propagates through about 65 m of 50 μm core GRaded INdex (GRIN) multimode fibre into twelve meters of 50 μm Stepped Index MultiMode (SIMM) fibre, arriving at a PiN diode on detector. The effective diameter of the SIMM fibre is slightly larger than the GRIN fibres diameter, hence light travelling from GRIN fibre to the SIMM fibre experiences almost no losses due to the transition. The DORIC-4A then reconstructs clock and data from the encoded streams and sends both as low voltages differential signals to a connected module.

BiPhase Mark Encoding

BiPhase Mark encoding encodes data and clock not into output states, but into output state transitions. On every even input clock transition (be it rising or falling edge) the encoder encodes a transition into the output stream. Input data is encoded by embedding a second transition on the output stream synchronised with the odd clock edge (falling or rising now). This logic delivers a balanced link with maximum frequency of the input clock and a minimum frequency of half the input clock. The encoding logic is shown in Figure 4.14(a).

BPM-12 The BPM-12 [W$^+$04] IC is the main device of the TX-plugin. It encodes the LHC bunch crossing clock and twelve data inputs into twelve streams, using BiPhase-Mark (BPM) encoding described above. The encoded streams are driven into a laser array using an internal laser driver. The output high-current can be controlled per stream using an external 12-bit DAC. Its upper 8 bits can be set by the user, thus allowing settings from 0 to 255. A setting of 160 was measured to deliver an output current of 10 mA through the driven load, which is the recommended forward current for the laser in use. The output

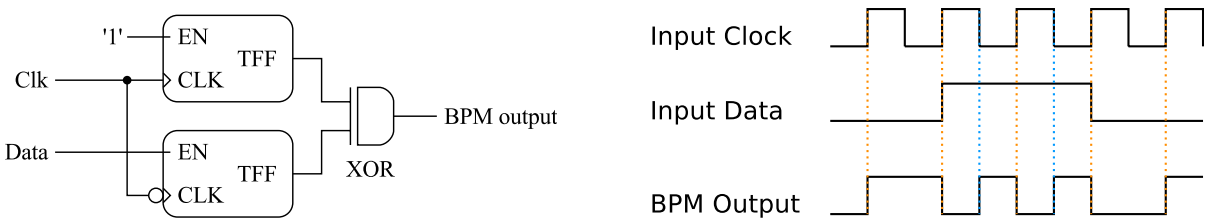

(a) BPM Encoder (b) BPM Encoding Example

Figure 4.14: *A simple BPM Encoder is shown in (a), encoding data on the falling clock edge of the input clock. The rising edge can be chosen by exchanging the CLK input inversion on the two toggle flip flops incorporated. An example output for given input data and clock is shown in (b), with orange lines indicating clock induced transitions and blue lines data driven transitions.*

power is close to linear, compared to the forward current DAC setting, as soon as the laser is driven above threshold (c.f. Figure 4.15).

An internal delay circuit allows for adjusting the phase of each output stream in steps of 320 ps for precise placement of the detector timing with respect to the bunch crossing. This BPM finedelay, as well as the output laser driver of the BPM, transform the shape of the outgoing signal. The latter is particularly strong when operating near threshold. Therefore another circuitry is embedded into the BPM-12, allowing for adjustment of the Mark-to-Space Ratio (MSR). Figure 4.16 shows that circuit. The initial incoming signal is run through two different paths, one being delayed by a fixed delay. Combining those two streams using an AND gate, every incoming positive pulse gets shortened by the propagation delay of the fixed delay. Using the same mechanism with a variable delay and an OR gate, allows for adding length to the signal. The variable delay is available to the user of the BPM-12 as the MSR register, allowing to change the length of positive pulses by ± 5 ns. Assuming either 20 or 40 MHz clock-like output signals, this refers to a change of duty cycle of ± 10 or ± 20 % respectively.

DORIC-4A The DORIC-4A [A+05] integrated circuit connects to a PiN diode, reading its current and decoding clock and command from the incoming BPM stream.

The optoboard is equipped with a PiN diode array, with an adjustable common bias voltage, typically set to 10 Volts. That voltage is delivered through a filtering circuitry with a 1 kΩ series resistance and a 1 μF filtering capacitance to ground. Despite filtering out high frequencies for the supply lines, the circuitry lowers the general bias voltage due to the constant current flowing through the PiN diodes. That can be as high as 6.5 mA, hence delivering a voltage drop of 6.5 V through the series resistance.

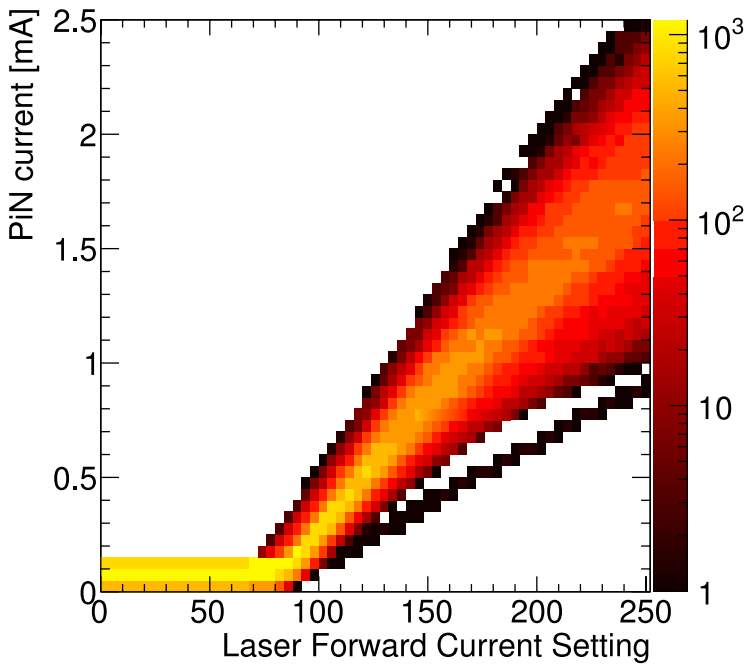

Figure 4.15: *Indirect measurement of laser output power [Hei09]. The output of light power of 283x8 TX-plugin channels was measured by shining their light through a 62.5 µm core GRIN fibre onto an electrically biased PiN diode. The aroused PiN current was measured versus the laser forward current DAC setting of the TX-plugin. A factor of 0.6 A/W was previously measured for the PiN diodes used. An mean threshold can be observed at a forward current setting of 80 and close-to linear behaviour is given above that.*

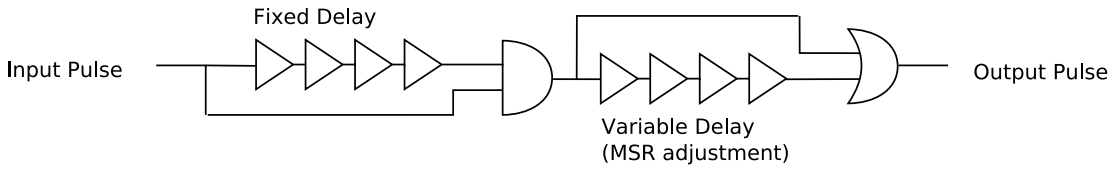

Figure 4.16: *Circuitry to shorten and elongate positive pulses of a digital signal.*

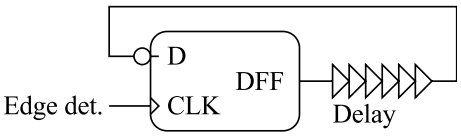

Figure 4.17: *The basic Delay Locked Loop circuit. An edge detector controls data forwarded by a data flipflop. The data itself is the inverted output of the flipflop, routed through a delay circuit. The minimum period or maximum frequency of switching is thus given by the delay.*

The PiN diode output is AC coupled into the DORIC input. An edge detection then generates short pulses from every transition. A special Delay Locked Loop (DLL) is then used to derive the on-detector clock. The clock coming out of that circuit (shown in Figure 4.17) is limited in its maximum frequency, depending on the internal delay. During a reset of the delay control circuitry a maximum frequency gets locked. If this happens with no data being sent on the BPM stream, the maximum internal switching frequency gets locked to 20 MHz. Clock doubling then generates 40 MHz frequency. Once out of its reset state, the DORIC can hold that maximum clock frequency, but as pulses are generated based on the input, it needs an edge input at least every 25 ns. If no input is given, the DORIC might either go out of clock or generate a clock based on noise input, which will have a random behaviour.
Extra transitions within the BPM stream compared to the 20 MHz empty transmission stream provide pulses which are not registered in the DLL. Instead those are sampled using the generated 40 MHz clock and delivered as command output for the connected module. This command output is very much affected by the input signal quality, as well as the stability of the internally generated clock. In case the input signal is lost and the DORIC generates a random output clock, the command stream will follow the same behaviour.

4.4.2 Data Link

The optical data link of the ATLAS Pixel Detector runs with data rates of up to 80 Mb/s per channel using an Non Return to Zero (NRZ) [Wik11] encoding. The VCSEL Driver Circuit (VDC) drives a laser with data coming from a differential input. The optical signal is then routed through the same length of fibre as is for the TTC system, incorporating this time a 50 μm core SIMM fibre spliced with a 62.5 μm core GRIN fibre, thereby, again, keeping the light loss small. Light is received from that fibre using a PiN diode, biased with 10 V. The DRX-12 IC amplifies the incoming signal and applies a threshold that can be set via an external voltage. Phase alignment of the incoming data happens on the Back of Crate Cards (c.f. section 4.5.1).
NRZ encoding was chosen on the return link, as the initial setup was intended for use of LEDs as optical transmitters [DP98, WW96]. These would have suffered from being operated with a 50% balanced code. Using NRZ and assuming a maximum link occupancy with data, the bright time goes down below 45%. The less data is transmitted, the less the LEDs would have been used. Eventually, when lasers were chosen for transmitters, the codec decision was kept.

VDC The VCSEL Driver Circuit [A$^+$05] (VDC) drives up to 4 lasers with individual signals, within the ATLAS Pixel system. The circuit provides the lasers with a fixed

bias current of 1 mA and a modulation current[j] that can be controlled with an external reference current, I_{set}. The modulation current is either routed through the laser, delivering the "bright" output state, or through an internal load, the "dim" state. This switchable load layout effectively reduces noise of the optoboard power consumption.
A full Optoboard uses either two or four VDCs, depending on the type of the Optoboard (Disk or B-Layer). I_{set} is given as a single voltage controlled signal, V_{Iset}, to a full optoboard, hence all laser channels are given the same modulation current. Decoupling of V_{Iset} in front of every VDC circuit is forseen on the Optoboard PCBs, but has not been done, as there were doubts about the capacitors withstanding the total radiation dose.

DRX-12 The DRX-12, c.f. Figure 4.18, is the off-detector receiver circuit of the ATLAS silicon trackers. Its specifications are given in [AW01]. It serves 12 PiN diode inputs, needing high-side bias. The design is meant to work for input currents ranging from 500 μA down to 57 μA amplitude. The inputs are routed to ground via a 3.2kΩ resistance thereby producing a maximum time averaged voltage drop of about 0.8 V. A second input comes from a threshold DAC, which is divided by a voltage divider with a ratio of 1:2. The maximum output voltage of 2.5 V hence produces the same voltage of about 0.8 V. Both voltages drive a comparator, comparing the signal input voltage with the threshold input voltage. An LVDS output of the comparison result is then generated. The bias block, shown in Figure 4.18, limits the maximum input and threshold voltages generated, but also adds capacitance to the inputs, thus decreasing the speed of the comparator.
On the RX-plugin, a bias voltage of 10 V is given for all PiN diodes. The threshold voltages for the twelve channels are delivered by two high precision 6-channel DACs.

4.5 Readout System Backend

The ATLAS Pixel Detector readout system backend is set up within 9U VME crates. Nine crates in total are needed to read out the Pixel Detector. Those are equipped with 132 readout building blocks, built from a ReadOut Driver (ROD) and a Back of Crate (BOC) card. The ROD delivers intelligent logic to generate commands for the Pixel Detector and identify returned data. It also allows the operator to run calibration scans on the detector and monitor the data stream during data taking. A Single Board Computer (SBC) per crate controls all crate electronics in terms of slow control via the VME64x bus. Trigger information is received and distributed via a single TTC Interface Module (TIM), located in Slot 13 of each crate. This includes the LHC clock, as well as external commands like the Level 1 Accept (L1A) signal from the trigger and Event Counter and Bunch Counter Resets (ECR and BCR).

[j] The difference between dim and bright current.

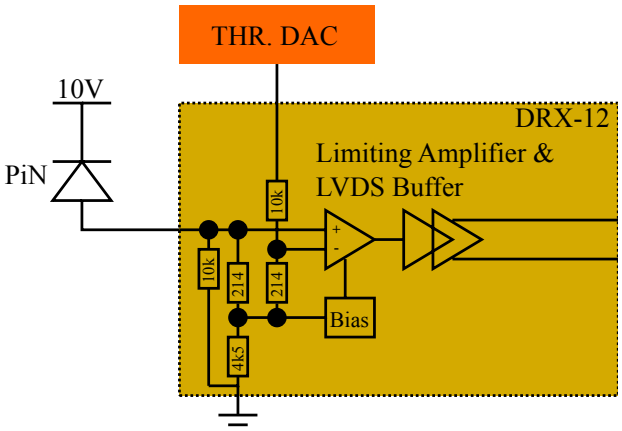

Figure 4.18: *A simplified view of a single channel of the DRX-12 IC and its inner functionality. Important to mention is that the bias current circuitry allows for large currents flowing*

4.5.1 The Back of Crate Card

The ATLAS Pixel Back of Crate [Fli06] (BOC) card shown in Figure 4.19 delivers the optical detector I/O interfaces, a TTC line per module for sending optical data to the detector and one or two data lines per module from the detector to receive data. Signal timing adjustments for the TTC path, as well as optical output power and input threshold adjustments are contained within the BOCs optical plugins described above. The mechanical connector for optical fibres is a different component mounted to the BOC, to which the plugins are screw-mounted. Connection quality thus depends on two components.

The BOC itself delivers timing distribution and adjustment for the readout system. It can run from two clock sources, the external system clock, delivered from the TTC Interface Module via the VME crate's custom backplane, or an internal 40 MHz oscillator. The BOC monitors the external input clock and falls back to the internal clock automatically in case of large differences between the two. This mechanism was introduced to deliver a quasi permanent clock to the detector modules, as otherwise their power consumption could change drastically.

First system tests showed a problem occurring when no TIM is installed in the crate. The clock input of the BOC has no high impedance pull to either logic 0 or 1, but is instead fully floating. Noise induced into this input is seen as a clock and allows for the BOC to go from the internal to the external clock source. As this external clock source is only noise, it causes unstable operation and lets the BOC soon switch back to internal clocking. This behaviour leads to unstable BPM encoding, which drives the detector modules into

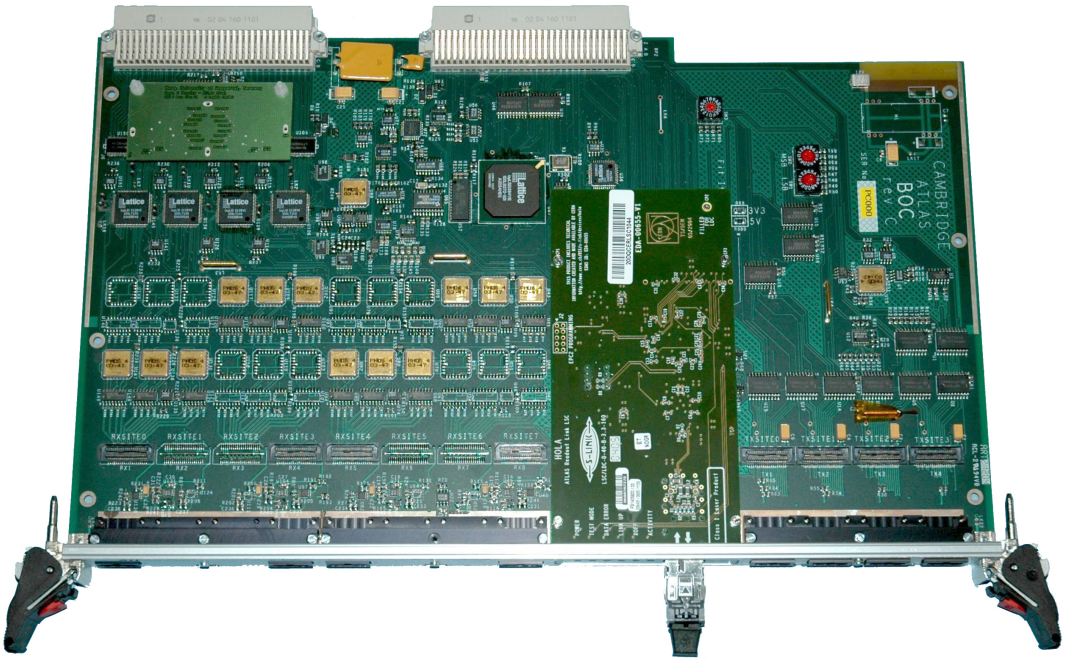

Figure 4.19: *An ATLAS Pixel Back of Crate (BOC) card, shown without electro-optical converters. Attached is a piggy back in the upper left for routing data, and a TDAQ interface card, the HOLA (see text).*

random states. It was therefore allowed to enforce internal clocking of the BOC using a register bit.
From the chosen clock source, three different clock trees with a total of four different clocks are driven:

- TTC Clock: To clock the TX-plugins of the BOC, a particular copy of the system clock, the P-clock, is used. The frequency is always equal to the input clock frequency and it can be shifted in phase using a PHOS4 (see below) delay circuit.

- Delay Source Clocks: To run the PHOS4 delay chips in use on the BOC, a clock reference is needed. All PHOS4 chips receive the same reference, called the A-clock. No adjustments can be done on this clock

- Sampling Clock: To sample data in 40 Mbit/s mode and transfer it to the ROD, the B-clock is used. Like the P-clock it can be shifted in phase, to stably deliver data to the ROD.

- Fine Sampling Clock: To run 80 Mbit/s sampling, a second clock is needed. The V-clock is used in this case. It is also to serve for timing measurements with good

resolution and is therefore routed through a Phase Locked Loop (PLL). This allows to set the phase of the V-Clock in steps as small as 40 ps. Additionally it can be halved and inverted.

All clock circuitry is set up differentially, until they need to be run through or delivered to a PHOS4 circuit.
An additional fifth clock is handed to the readout driver through the backplane. Thereby the BOC is the only device within the system adjusting clocks and hence delivers running conditions for a readout block.

PHOS-4 The PHOS-4 [Gro00] delay circuit delivers a total of 5 inputs and outputs. It receives single ended 3.3 V low voltage TTL signals and is controlled via an I2C bus. Up to 16 different I2C addresses can be set using address inputs. One of the five inputs must be driven with a base clock, serving a 40 MHz frequency, whilst the other 4 can be driven with random data. Pulses above 3 ns length can be delayed by up to 24 ns in steps of 1 ns.

Sampling The BOC itself only samples data with an off-detector clock, such that it can be handled within the following electronics. Besides that, it serves as an input switch board, supplying all parts of the ROD with an equal input load. This is not dynamic but comes with a selected routing, which splits data coming from a number of modules to a number of ROD inputs.
In case of the ATLAS Pixel Detector, incoming data can be sampled in different modes after reception. Using Computer Programmable Logic Devices (CPLDs), either of two clocks or both can be used for sampling data. The dual clock option allows to sample with twice the data rate of standard sampling and requires the clocks to run with opposite phases. The two individual sampling points are then split into two different data lines that are handed to the ROD at 40 Mbit/s rate[k], as shown in Figure 4.20. Running in 40 Mbit/s mode, either of the two lines (depending on the mode choice) is disregarded; only the chosen one is forwarded to allow for more individual modules to be connected to the ROD.

HOLA Interface: To connect the ATLAS Pixel Detector readout system with the ATLAS Data Acquisition system, a Highspeed Optical Link for ATLAS [RvdBH03] (HOLA) is installed as a daugtherboard on the BOC. It delivers an SLink[l] protocol

[k]The Pixel Detector ROD can only handle 40 Mbit/s input data per data line

[l]A CERN standard system for high speed data transmission, based on Texas Instruments' TLK SerDes devices. More information can be found on the S-Link Homepage at CERN: http://www94.web.cern.ch/HSI/s-link/.

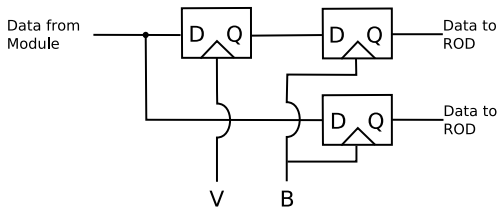

Figure 4.20: *Sampling 80 Mbit/s input data with two 40 MHz clocks, V- and B-clock, into two 40 Mbit/s data streams.*

implementation to the readout driver, so that processed data can be handed over to the next stage of processing.
The SLink comes in two fashions, a Link Source Card (LSC) and a Link Destination Card (LDC). Whilst the HOLA LSC is implemented on the BOC, the LDC is held by an ATLAS ReadOut Buffer INput (ROBIN), which buffers data and interfaces to the high level trigger systems.

4.5.2 The Pixel Readout Driver

The Pixel Readout Driver (ROD) is a 9U VME card, delivering realtime stream processing for Pixel Detector module data and calibration processing capabilities within arithmetic units. Each ROD is composed of 12 FPGAs and 5 DSPs. The DSPs split into an integer DSP and 4 floating point DSPs. The integer DSP controls the RODs operation, stores module configurations and can utilise and alter them locally for running scans without host intervention. The floating point DSPs are indirectly controlled by the integer DSP and used for calculations on calibration data and monitoring. The general naming convention refers to the integer as the master DSP and the floating point as slave DSPs.
The FPGAs deliver a realtime data path, capable of feeding up to 160 MB/s of data into the higher level readout connection, the HOLA. Data passes through:

- Eight Formatter FPGAs: The Formatters (FMT) each receive data from 1, 2 or 4 modules, depending on their data rate. They parallelise the incoming data streams into 32 bit words and, in case of timeouts or skipped events, add empty events to fill the data stream.

- Two Event Fragment Builders: The two Event Fragment Builders (EFB) are contained within a single FPGA. Each is gathering data from a block of four formatters, adding event ID information and summarising errors.

- The Router: Data from the EFBs A and B are merged into a single event fragment when handed to the router and then, depending on configuration, handed to the router. The router can also trap data witthin four FIFOs, each connected with

a floating point DSP. Conditions for trapping can be trigger IDs and error flags, delivering comparable histograms for different conditions.

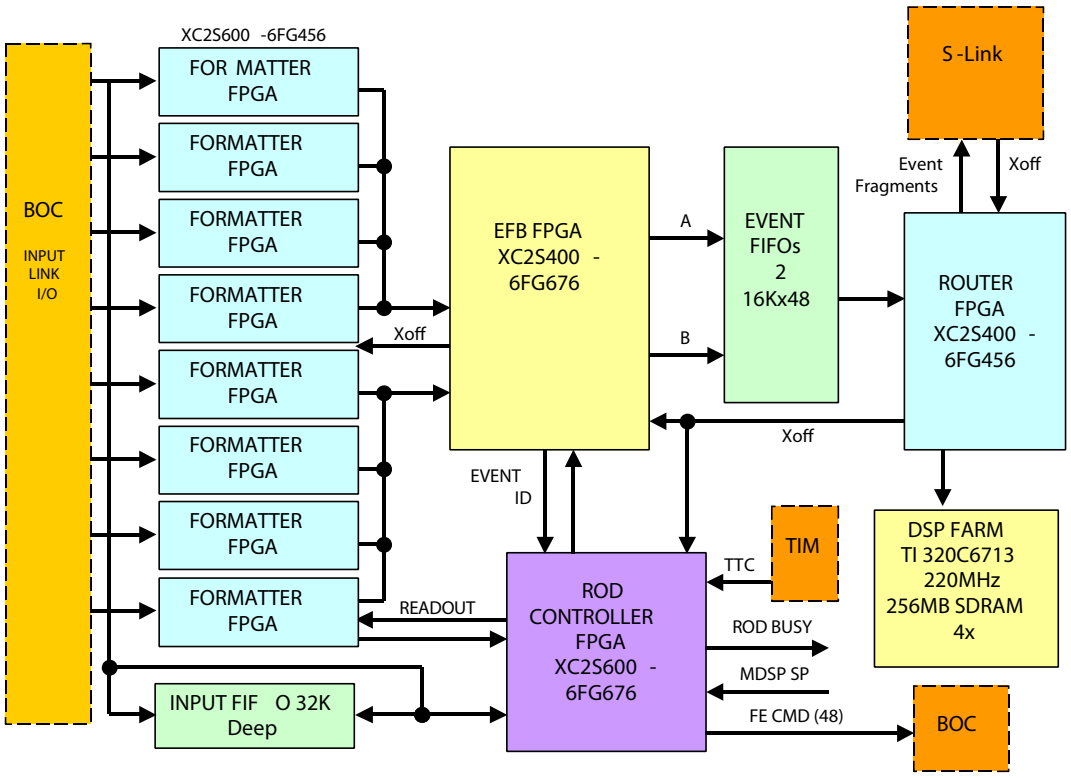

Figure 4.21: *Schematic outline of the ATLAS Silicon ROD datapath.[JJF+06]*

The ROD controller FPGA (RCF) is the operational heart of the ROD, giving access to all ROD components, as well as BOC registers. It also serves the module command connection for each module, allowing for running data taking driven by the TIM input or calibration driven either by the Master DSP or an internal scan engine.
A last FPGA, the Program Reset Manager (PRM), controls bus interfaces and firmware loading. It connects to a flash memory containing firmware for all FPGAs on the ROD but the PRM. As this particular FPGA can only be reprogrammed by a special interface, it serves as a stable control interface. This split mechanism between bus interface and data processing logic serves an easy to upgrade, yet save and stable firmware policy.

4.5.3 The TTC Interface Module

The TTC Interface Module connects to the trigger system input and delivers ATLAS commands and a regenerated LHC clock and for all readout blocks in a crate. It also

holds the ATLAS Level 1 ID and Bunch crossing ID counters and broadcasts them to the Readout Blocks as needed. The connection serving clocks to each readout block is de-skewed, so that all readout blocks within a crate run with the same LHC clock phase. Eight direct command lines, running with LHC-clocked data serve to control data taking operation of the readout drivers. Four lines are direct, delivering Level 1A, Event Counter Reset (ECR), Bunch Counter Reset (BCR) and CALibration (CAL) commands to the readout drivers. Two more signals are used as serial transmission lines, transferring the ATLAS event ID and bunch crossing ID to the RODs.

4.6 Control

To operate the Pixel Detector, before mentioned components like modules and optoboards need to be supplied with power, analogue control and monitoring. The Pixel Detector Control System (DCS) allows to do so, communicating with various supplies and monitoring units, mostly via CAN bus and also Ethernet.
Each Pixel detector module is supplied with two high current voltages for digital (VDD) and analogue (VDDA) supply. Both voltages are sourced from a WIENER power supply for a full set of 6 or 7 modules, connected to the same optoboard. The supply is distributed at Patch Panel 4 (PP4) and current consumption can be measured on both, supply and return lines. A last voltage regulation takes place at PP2, 12 m away from the Pixel Detector modules. Under normal conditions the currents vary between 200 and 900 mA for the digital supply per module. 80 and 1300 mA are standard limits for the modules analog supply.
A high voltage source is to bias the sensor. No near regulation is needed here, as the currents flowing are much lower than in the low voltage domain, so the cable induced voltage drop is low.
For pixel optoboards, a dedicated power supply has been developed, providing not only the low voltage power (V_{VDC}) at more than 180 mA, but also a PiN diode bias voltage (V_{PiN} up to 20 V, 20 mA) and a voltage regulating the optoboards output power (V_{Iset}). The SC-Olink[m] [KKFN05] connects directly with the detector on its low-current outputs and through a redundant PP2 channel on the VVDC line. A single digital connection serves for controlling the optoboard reset. It allows for holding the reset line low for 10 μs.
In addition, the detector modules, the optoboards and other components each carry an NTC[n] to measure the temperatures of all components. An interlock system checks these temperatures and, depending on a switchable matrix, shuts down components which are endangered by too much heat. The same interlock system takes care of laser safety, being

[m] SC-Olink: Supply and Control for the Optical link
[n] NTC: Negative Temperature Coefficient

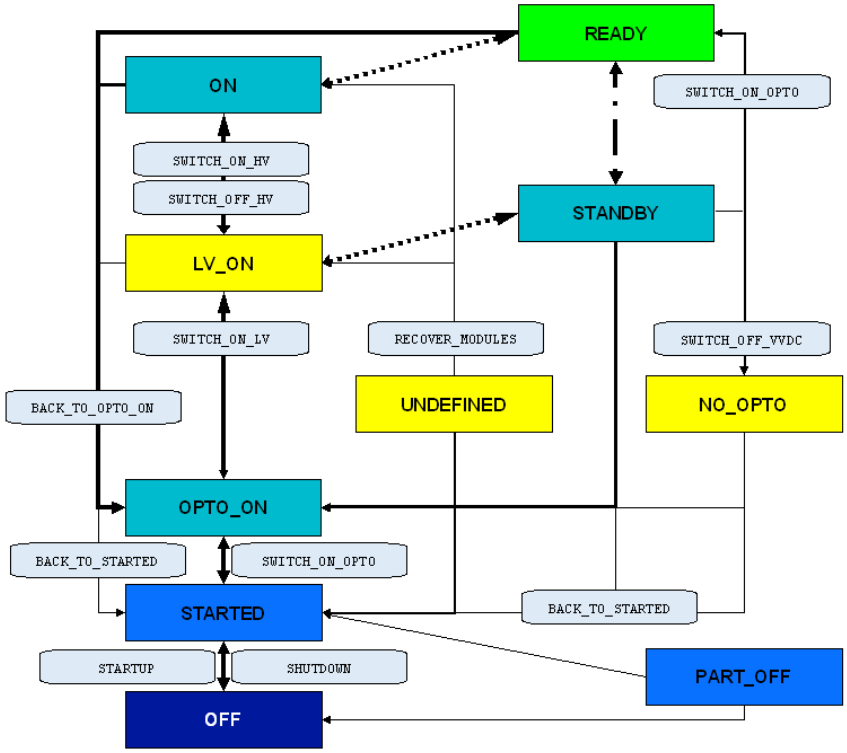

Figure 4.22: *Shown is the Pixel Finite State Machine state diagram for a Readout Unit° in a simplified version.*

connected to all doors that contain data transmission lasers (Class 3B). In case rack doors are opened, both the readout crates inside as well as the SC-Olinks connected to the corresponding optoboards, are interlocked.

4.6.1 Finite State Machine

The user interface for the Pixel DCS system is the Finite State Machine (FSM). It allows switching the detector into states, e.g. OFF, STANDBY or ON, and categorises the monitoring values of each module into a status. A state and status for the full detector is calculated from the individual module states and statuses. Problems can easily be identified from the top level (the detector view) and traced to the particular source.

To first order, the user wants to know whether the detector is ON or OFF and whether it is OK or not concerning that state. As these two binary states are very limiting, additional information was included and other states were defined. As can be seen in Figure 4.22, the state diagram of the ATLAS Pixel Detector is much more detailed.

State Walkthrough: Roughly outlining the operation of the detector, it is brought from OFF to STARTED by switching on all low voltage power supplies and checking monitoring voltages once. This helps to detect broken sense lines, which would drive the regulators into higher voltages and risk destroying the connected module. As soon as STARTED is reached, the optical data connection is powered up by switching on the optoboards. The resulting state is OPTO_ON which supplies the modules with clock and command, although they are not powered yet. The low voltage for the modules can be switched on afterwards, as they're now given stable operation conditions. Switching on modules at the same time as the optoboards might result in random configuration of the modules and therefore high current consumption - which in turn leads to a direct shutdown of the module. As soon as the low voltage was switched on, the modules and in combination the full PP0 ends up in LV_ON. Eventually switching on the high voltage gets the modules into ON state. To get the modules into READY state, which is needed for data taking, the DAQ systems have to configure the modules for data taking. The new configuration then changes the current consumption which, in case the currents are well defined, makes the DCS system assume the modules are in READY state.
The time the modules are operating, an expert can switch off the low voltage supply for the optoboards. This results in a NO_OPTO state, in which the modules do not receive clock and command anymore. This state is needed for some special investigations on the data transmission lines and removes any risk of accidentally configuring modules.

4.6.2 DAQ-DCS communication

The DCS and DAQ systems are run on individual machines and within individual software frameworks. A DAQ to DCS Communication (DDC) component is to handle data and commands from either of the two systems to the other. This way, the DAQ system is able to determine the state of the detector and switch it on or off as may be needed. Commands are in general only sent from the DAQ to the DCS systems, and in case data should be posted to the DCS system, this is not pulled from the DCS system, but pushed from the DAQ side.
Due to monitoring intervals of more than 5 seconds within the DCS system, setting of values is very slow via DDC. A set value arriving is written to hardware, but then the DDC component on the DAQ system side only acknowledges the command when the set value is (within limits) read back from the DCS monitoring system. This latency is to be taken into account in all interactions of DAQ and DCS. It was partially bypassed, by allowing to set multiple values at a time via DDC.

5 Optolink Calibration and Performance

This chapter is to describe the ATLAS Pixel Detector optical link setup procedures, its performance and particularities. An introduction in the first section will explain how the system parameters can be evaluated for functionality and performance. The second section will then describe tuning targets for the optical link. Eventually an overview of the performance of the link will be given in the last subsection, pointing to required changes in a new setup.

5.1 Calibration Mechanisms

The Pixel Detector optical readout link needs manual adjustment of many parameters. General communication only needs tuning on the data link, whilst the timing, trigger and command connection automatically configures. Tuning of the TTC link (changes of the BPM finedelay) affects each module's timing and thus the quality of data taken (c.f. Section 4.2.2). In the process of calibrating each module's timing, a connectivity map is needed, to tell which module is connected to which link.
In theory this map can be prepared by knowing all connections within the system. Practically it has been problematic to identify errors in the mapping and thus a method to detect connections and pinpoint single modules was needed.

5.1.1 Connectivity Scans

The Pixel connectivity map was verified during installation by switching on each device after it got connected and checking whether the expected response was showing up in the detector control system. This eventually lead to a functional mapping of the detector control system, which was then used to run automatic scans within the DAQ.

InLink Scan: The InLink scan is to find working connections into the detector, assuming a given connectivity map. Malfunctioning single channels can be identified easily. Additionally the scan allows the system to automatically flag bad cable connections or wrong connectivities. The latter could either be a cable plugged in a false location, or the connectivity database pointing to the wrong optoboard for a particular TX-plugin.

Parts of the detector that are scanned, have to be in OPTO_ON or NO_OPTO states within the DCS finite state machine. In the first state, the modules are off, in the second, the optoboard is partly off and therefore no commands are sent to the detector modules which, in that case, are on. In both states it is impossible to misconfigure the detector due to random output of the DORIC and hence produce stress due to high thermal load or sudden shutdown by the DCS FSM.
With these starting conditions, the scan tries to evaluate the light power of each TX plugin channel, by switching channels on (or off) one at a time and reading the optoboards PiN current. As the optoboard only has a single supply for the PiN Voltage, the current reading has to be done one channel at a time. The scan is by default executed at a laser forward current setting of 160, corresponding to about 10 mA forward current, which gives comparable results on all links. Badly plugged connectors can thereby be identified and fixed.

InLinkOutLink Scan: The very same scan in the opposite direction is done by having the module control chip returning half its input clock. The modules must be on and operational for the scan to work, thus OPTO_ON or NO_OPTO are not to be chosen here. Instead the detector should either be LV_ON, ON or READY. Thereby, the optoboard lasers are driven with 50 % duty cycle, returning a signal to the off-detector receiver. Here, the PiN current produced is measured. Half-clock return is switched on for one module at a time, allowing for currents to stabilise in between. The scan is normally performed with V_{Iset} set to 900 mV. This allows a direct comparison between different optoboards in terms of optical power, allowing again to identify bad connections.
Despite measuring the output power for each link, the scan tries to read on which channel data is coming in. It can thereby build a map of the interconnections between TTC- and data links[a] for a single BOC.

5.1.2 Data Link Scans

Opposed to the previously described Connectivity Scans, which only perform functional measurements, data link scans try to validate whether chosen parameter sets give stable operation for the optical data link.
Adjustment of the pixel data links incorporates setting of three parameters for each link:

- Light power (On-detector): The on-detector light power can and must be adjusted for a full optoboard (6/7 modules), such that all corresponding receiver thresholds

[a] As the scan activates half-clock return by sending a command on a TTC link and then reading back on data links, this gives the TTC to data link map. It is not supposed to provide the mapping inside the detector.

can be tuned to an ideal setting. The regulation is done via the V_{Iset} input of the optoboard.

- Threshold (Off-detector): The receiving threshold needs to be tuned such that a stable identification of ones and zeroes is guaranteed. The amplitude, and thereby the ideal threshold, of the incoming signal was chosen previously, but is affected by noise on supply lines and on the optoboard power regulation line.

- Phase (Off-detector): The data is to be aligned in phase with the sampling clock off-detector, because no clock is transmitted within the detectors NRZ return signal.

To identify ideal settings for all of these parameters, different types of scans have been developed in the frame of this thesis [DFH+10], verifying functionality of the link setup for a given parameter set.

Pattern Based BOC Scans: The first implementation of BOC scans is based on the MCCs frontend input FIFOs. These can be written via the TTC connection and read out again via the data link, by sending a trigger to the module. To receive the triggered pattern from the module, a set of FIFOs inside the ROD is armed and filled as soon as the readout trigger is sent to the module.
The returned data is compared with the expected return pattern and a number of transmission errors is calculated from that. This is done multiple times for each set of parameters being scanned, delivering a resulting histogram as shown in Figure 5.1. White areas have no error in transmission, other areas have an error count according to the colour scale given.
Each scan output has some characteristics to be introduced:

- Reversed Time Order: BOC scans show a physical quantity, the signals amplitude versus phase, but with a reversed time axis opposed to an oscilloscope for example. That is due to the fact that not the sampling clock is delayed, but the data itself. Hence, at higher delays, the scan shows an earlier phase of the incoming signal. Other than that a BocScan shows a standard eye-pattern[b], the white area being error free.

- Delay Error Band: The vertical error band within the scan reflects the transition phase of the data stream. When the data delay is adjusted such that data is being sampled during the transition, the result is not predictable due to jitter and edge slopes, therefore the delay error bar exists. The delay error band is heavily influenced

[b]Eyepatterns are a standard way of displaying transfer characteristics. They allow for simple mask tests to verify whether a transmission is of appropriate quality. A description can, for example, be found in [Chr94]

by the on-detector clock and can show bad adjustments on the TTC link. It also delivers a possible measurement for changes in the on-detector clock-phase versus the LHC bunch crossing clock.

- Bottom Error Band: The bottom error band exists due to a minor design flaw. It reflects a (dis-) charge current that flows from the PiN diode through the input resistor of the DRX-12. The current is influenced by the received amplitude. It can cover the full threshold range, in case the input light power is too large.

- Top Error Band: The top error band is driven by the incoming signal amplitude. Due to the RX-threshold being a discriminator for the current coming in, the lowest threshold that is within the top error band precisely measures the signal amplitude. In scans with a very fine grained timing, the top error band allows rising and falling edge slopes of the input amplifier to be measured.

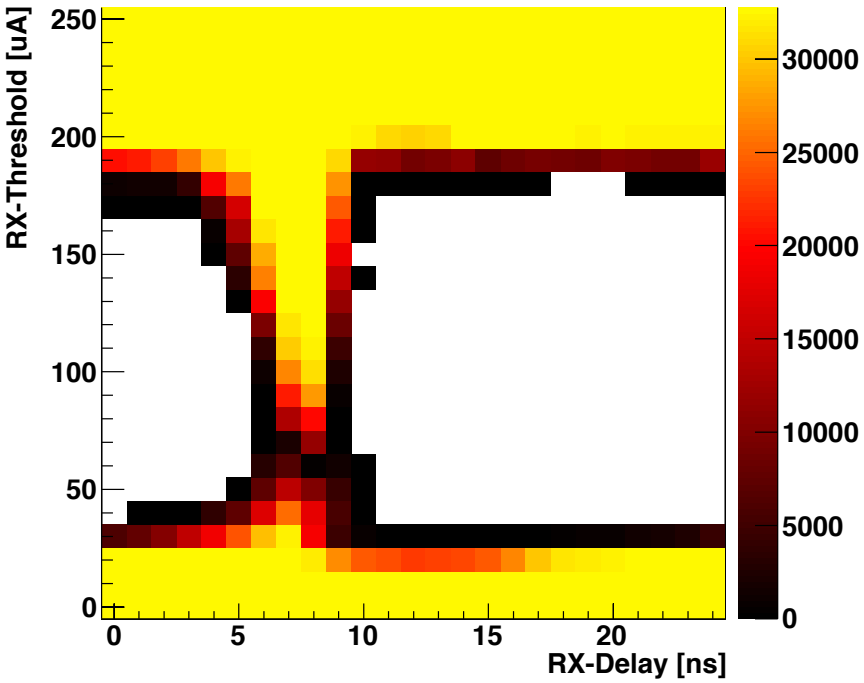

Figure 5.1: *A BOCscan example, showing bottom and top, as well as the delay error band.*

Clock Based BOC Scans: A later implementation of BOC scans incorporates an MCC embedded pattern generator. For the scan, it is set up to return a 20 MHz clock-like pattern, half the on-detector clock. Off-detector, the signal is received using counters in the ROD formatters that only count up on receipt of a '1'. The counters run during

a definable number of clock cycles, at maximum 65536 40 MHz clock cycles. A working set of RX-threshold and -delay will allow the system to sample a '1' at the input of that counter during half the total cycles sampled.
The major advantage over the previously described scan is speed. A single sample with maximum counter limit is done on all of a RODs input lines (at maximum 32, out of which 26 are at maximum connected to modules) within 1.6 ms. Therefore a full BOC scan with 25 RX-delays and 13 RX-threshold values[c] takes as long as half a second, not including preparation and reading back of values. The scan is only limited by the RODs readout bandwidth towards the controlling single board computer via VME. This severely limits the scan, when run in a fully equipped crate. Due to interference between the different readout blocks, it leads to 25 minutes of scan time in a crate equipped with 16 slots and the maximum number of modules connected.
A disadvantage of this scan mode is the "unnatural" scan pattern: Half clock return keeps the laser on 50 % of the time, hence a kind of stable state is delivered. With pseudorandom data, the turn-on and turn-off behaviour of the laser is seen more clearly. Evaluating speed over accuracy, the clock based scans are mostly used for initial setup of all links and for timing related measurements, which profit of the stable return signal.

DSP Based BOC Scans: To build a fast scan that shows the behaviour of the link, as it is during data taking, the initial pattern-based scan has been ported into readout driver DSP code. The same type of scan is run, but takes much less time, as no raw data has to be transferred from the ROD to the SBC. Instead of capturing data within the RODs input FIFOs, they are routed to the slave DSPs, where they are being processed for error counts. Histograms are built within the slave DSPs and read out after the full scan is run. Both, the distribution of data processing as well as the data reduction within the SBC drastically increase the speed of these scans over the above mentioned slow BOC scans (from minutes to seconds). Recently the BocScan run on the DSPs is basically limited by the setup-time to little less than a minute of total run time.

Shapescans Another type of counter mechanism based scan was created during the Pixel Detector commissioning phase. As is for fast BOC scans, shapescans utilise the MCC pattern generator in half-clock mode, but instead of sampling the return data every clock cycle, shapescans use the BOCs V-Clock in half-clock mode to sample data. The difference compared to 40 MHz sampling is shown in Figure 5.2. Sampling every 50 ns does not merge 0 and 1-phase of the incoming signal. A total of 50 phases are measured here, delivering not an error count, but a signal count. This means a shapescan can show the high and low bits returned, instead of "only" showing the deviations from half-clock.

[c] Both are standard values for a BocScan.

An example is shown in Figure 5.3. Features from the previously discussed BocScan can be observed here too: in the vertical axis a noise floor below a threshold of about 30 μA and an extinction level at about 210 μA, and in the horizontal axis two transition regions at about 6 and 33 nanoseconds V-Clock phase. The scan allows to measure individual features in turn-on and turn-off behaviour of bits and therefore holds more information than a standard BocScan.

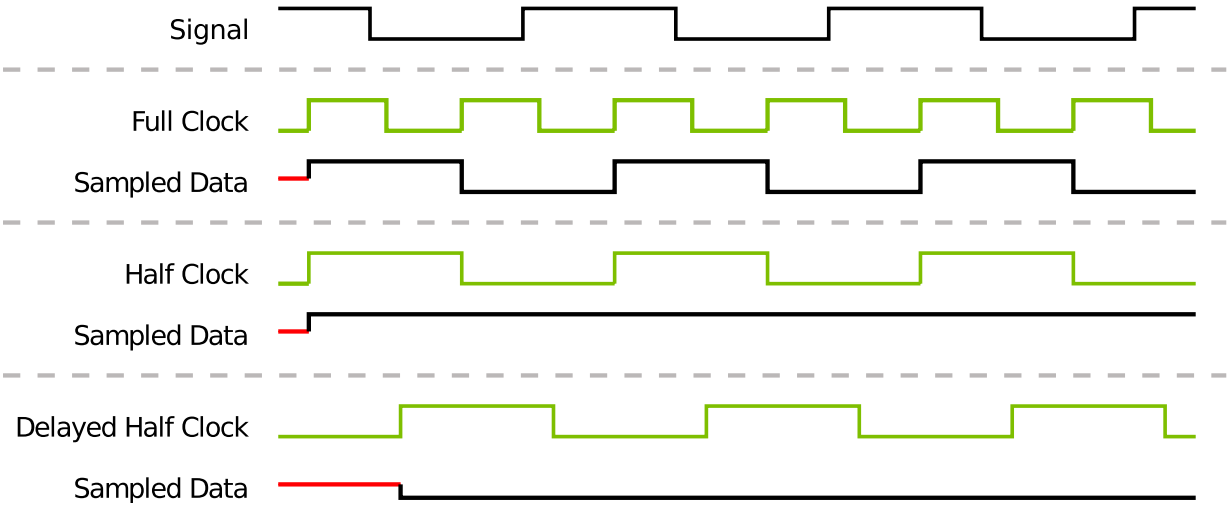

Figure 5.2: *This schema illustrates how a 50 % duty cycle clock-like pattern that repeats every two clock cycles can be sampled. Sampling always happens on rising clock edges. In case the full clock is used, every second clock cycle samples a logical 1, every other a logical zero. Using half that clock to sample the signal, the output is dependant of the phase between clock and signal. Half-clock sampling thus allows to scan signals that repeat with a half-clock period for their content.*

5.1.3 TTC Link Scans

As parameters of the TTC system only affect the on-detector clock reconstruction and thus timing, scans for these parameters need to identify the timing conditions on-detector. The only measures about on-detector timing can be derived from physics content in data (c.f. section 4.2.2) and transitions in returned data streams. The technical approach of checking the transition timing allows for tuning Mark to Space ratio, without beam and a complicated offline analysis. The finephase of the detector modules on the contrary can only be tuned with beam [Ibr10]. A profile of the delivered Time Over Threshold values in different bunch crossings reflects the quality of the tuning. An example is shown in Figure 4.6. Re-establishing tuning conditions without beam is possible, by tuning the TTC link delay to arouse an RX-Delay that is measured with beam in a tuned state.

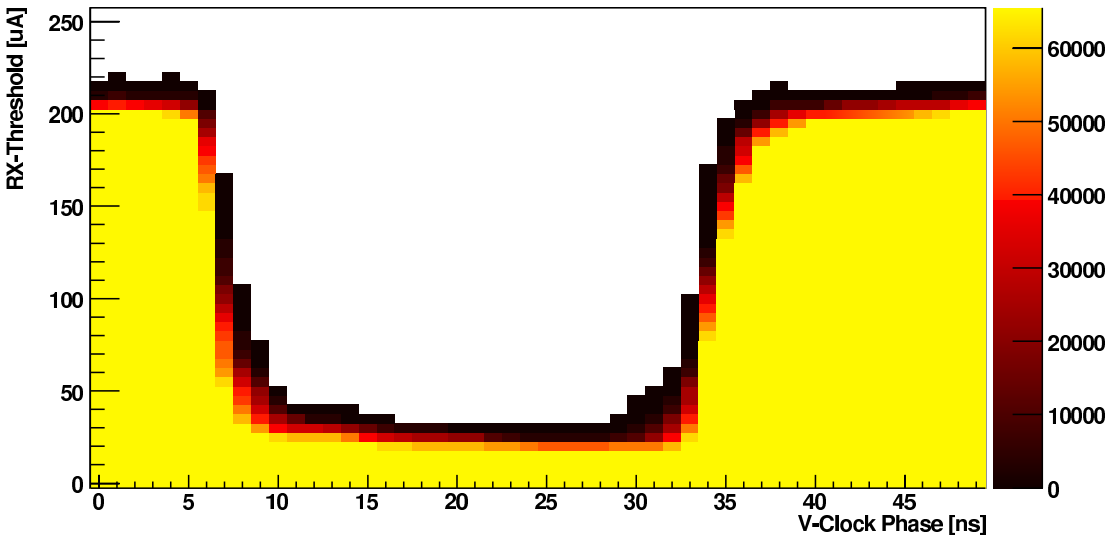

Figure 5.3: *A ShapeScan example. Yellow areas have always seen a binary '1', white areas a '0'. All other colors refer to scan points that had mean values between 0 and 1 (depending on amplitude noise and jitter of the incoming signal). Natural time order is kept in this scan - "later" phases are at right.*

Mark to Space Ratio scans Mark Space Ratio (MSR) is subject to a special scan, as it incorporates indirect measurement of the TTC tuning state, first mentioned in [Dop07]. The scan must measure the on-detector clock quality from data.

To do so, the BOCs V-Clock is used to sample a returned Half-clock stream with the ROD counters for clock based scans. A full signal sampling with all V-Clock coarse phases is done, using a 0-24 ns phase shift of the 20 MHz V-Clock and inversion (to produce a phase shift of half a clock period, 25 ns). Using the full sampling, signal edges can be found and the V-Clock is set such that a sampling of both data edges using the fine phase is possible[d]. This way both rising and falling edges of the signal can be sampled and an accurate position of both can be determined, relative to the V-Clocks fine phase 0 adjustment. This allows to measure the duty cycle of an incoming 20 MHz clock-like signal that has 50 ± 20 % duty cycle.

The on-detector clock is determined by measuring the return signal before and after a "BPM flip", as can be seen from Figure 5.4. Measuring the pulse length in a stable state before the flip and after the flip, there should be no difference in case of an ideal tuning state. The on-detector clock would have a constant period of 25 ns. Two measurements are needed, as pulses could be deformed by other effects, e.g. slow laser turn-on on-detector, hence the returned signal pulse width is not necessarily equal to the corresponding clock-cycle length on-detector.

[d]The V-Clocks fine phase adjustment allows for phase shifts of up to 10 ns in steps of 40 ps.

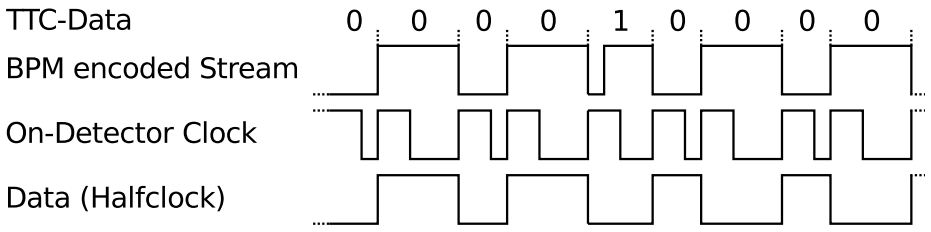

Figure 5.4: *This method is to measure the on-detector clocks. Measuring the returned pulse width before and after a "BPM flip" (the single 1 being sent) allows one to sense the difference in on-detector clock-cycles as given in this figure. This also measures the (mis-) tuning state of Mark to Space Ratio.*

An MSR scan is now done by cycling through all MSR settings and always measuring the returned half-clock duty cycle before and after a BPM flip. The resulting absolute difference, as shown in a scan example in Figure 5.5, allows the ideal setting to be found, which gives the lowest difference in returned duty cycles before and after the BPM flip. Yet, the on-detector clock reconstruction might malfunction at drastic settings of the MSR register and therefore run into a 20 MHz clock lock, instead of the expected 40 MHz. This could be detected off-detector, but then needs DAQ-DCS communication to be fixed, which is the primary reason why it has not been implemented into the scan yet. Also, the particular scan step has to be re-done, as the clock loss might happen during the phase flip and a reset of the Optoboard clock reconstruction might lead to misalignment of the phase of the BPM signal and the returned half-clock.

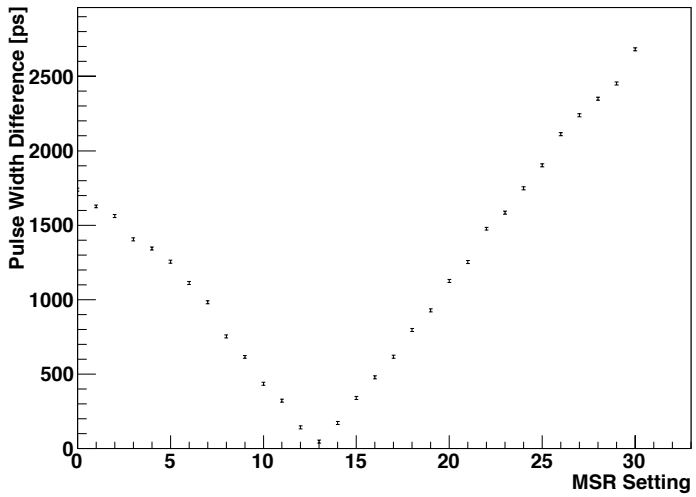

Figure 5.5: *Example of an MSR tuning scan. For each MSR setting, the absolute difference of positive pulse lengths in the returned data stream before and after the BPM flip are plotted.*

With this scan in place, the full detector was scanned for the ideal MSR versus the BPM finedelay setting, to verify and quantify the dependence between BPM finedelay adjustment and MSR, mentioned in Section 4.4.1. The resulting plot is given in Figure 5.6. Error bars are statistical, but directly show that large BPM finedelays and MSR settings lead to unstable operation of the detector, as modules are only included in case they provide a sane scan result.

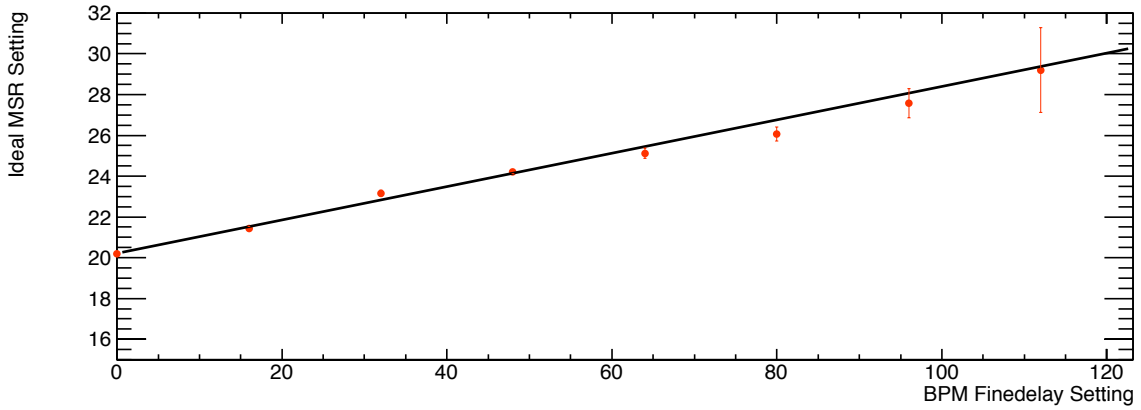

Figure 5.6: *The dependence of an ideal MSR setting (from a tuning scan) of the used BPM finedelay values. The average change of ideal MSR per step in BPM finedelay is* 0.08. *Error bars become visible, as the higher delays result in fewer modules passing the scan successfully.*

5.2 Tuning of the Optical Data Links

As mentioned above, multiple settings affect the operation of an optical link. For the ATLAS Pixel Detector, tuning strategies have been developed in the frame of this thesis to guarantee good performance of the links over an extended period of time.

RX-Delay Tuning Tuning the RX-delay is the most simple step, as it only relies on finding the delay error band. The ideal delay is then set 7 ns higher (modulo 25 ns), as this is far away from the last, and right before the next signal transition. Hence the signal is most stable in that phase as it had the longest time possible to settle.

RX-Threshold Boundaries Whilst the output light power on the optoboard has a very large range of adjustment, the RX-Threshold is limited within 0 - 255 μA. The amplitude of the incoming current signal changes with the output light power. Not only the maximum level of error-free reception is affected by light power, the minimum is as well. Therefore

the threshold needs to be chosen keeping in mind possible changes of the optical light power due to random influences[e].

Optical Power Tuning Two aspects drive the choice in optical power. The first is stability of the signal as such: when V_{Iset} is low, setting the lasers too close to their threshold, they tend to have low switching speeds or might not even switch on at all. The second reason is driven by the threshold setting being stable versus changing conditions on-detector: The bottom error band increases slower with increasing light power than the top error band does. Also, a part of the detector suffers from a problem referred to as Slow Turn-On (STO):
In case the on-detector lasers have not sent data for some time, they switch on slowly. This implies that the first bits sent have a lower incoming power than the consecutive bits. Such the choice of threshold should always be closer to the bottom error band, as that one is more stable with respect to power fluctuations and does not cut out the initial low amplitude bits in a problematic link.
To incorporate both arguments into a light power tuning, a measurement must be given that is sensitive to the light power, even if the top error band exceeds the BOC scan and can thus not be clearly located (which is due to the first tuning argument, power). Within a fully visible[f] BOC scan, the threshold with the highest available range of RX-delays (without errors) is typically located at about one 4th of the full amplitude of the error free region, one 4th of the way from bottom error band towards the top error band. It therefore delivers a good measure for the optical amplitude. The exact same threshold setting is also located in good balance between the top error band and bottom error band to incorporate fluctuations of the light power amplitude and therefore is the set tuning target of a threshold tuning. It can easily be found as the minimum value in the y-axis projection of a BocScan as shown in Figure 5.7. The light power is to be chosen such, that this minimum is close to the upper end of the threshold range.
An example of a slow BOC scan resulting from a properly tuned link is shown in Figure 5.8. The RX-Threshold would here be set to the maximum value, 255 uA, and the RX-Delay to 14 ns, as the Delay error bar ends, to the right, at an RX-Delay of 7 ns.

Self-calibration

The very same tuning targets were embedded into a scan based on the fast BOC scan mechanism, the OPTO_TUNE scan. It automatically leaves the system in the best tuning

[e] If the minimum error free region was always at zero input power, the threshold could be chosen just above zero, disregarding changes of input power.

[f] This refers to the top error band being visible.

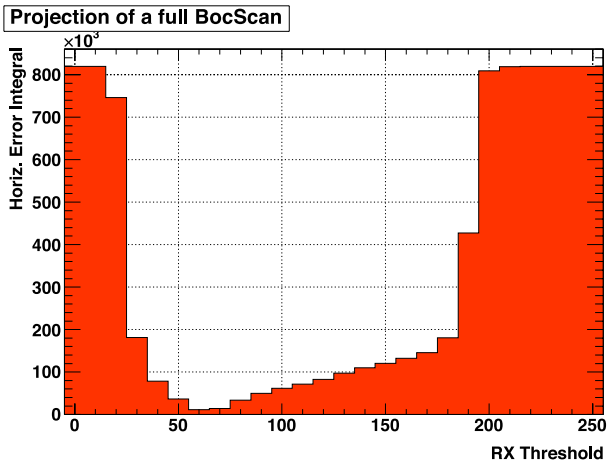

Figure 5.7: *The projection to the y-axis of a BocScan as shown in Figure 5.1. The ideal threshold is defined as the one bin with the smallest entry, 60 in the given case.*

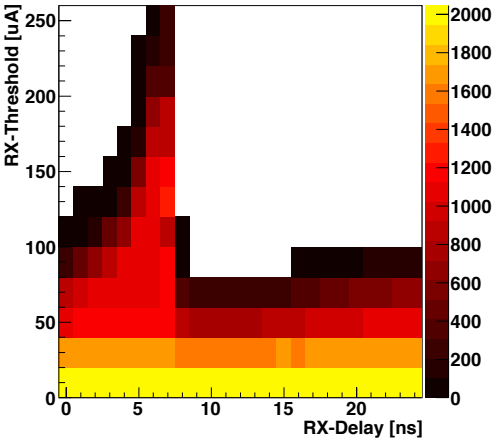

Figure 5.8: *An example of a slow BOC scan for a tuned linked. This scan shows no top error band and delivers the smallest horizontal sum of errors in the top-bin.*

state achievable within the given parameter ranges and that configuration can be written to disk for later use.

The scan is based on horizontal slices of a ShapeScan (c.f. Section 5.1.2). Instead of taking a full two-dimensional scan, a slice is sampled at the central threshold. Depending on the number of phases above threshold, the threshold is in- or decreased to reach a 50 % occupancy of ones within the sampled slice. This is done, until the threshold is optimised - in best cases to the last bit, hence 8 slices of a shapescan are sampled (This is generally not required, as optical power fluctuation and other effects change the ideal threshold by 10-20 settings easily). Reaching a final threshold for a given optical power setting of an optoboard, an evaluation is done to check whether the optical power could be in- or should be decreased.

The scan iterates over all settings with the very high granularity[g], using far less time than a corresponding three-dimensional BOC scan (both slow and fast). In comparison with a 3D DSP based BocScan, it still has drawbacks in terms of speed, which would need tuning of the scan sequence.

[g]In case of RX-Threshold and -Delay, the highest granularity available is used, while V_{Iset} is limited to 12.5mV resolution, which is better than any 3D BocScan delivers.

5.3 Commissioning and Performance

In 2007, the Pixel package construction was finished. The package includes not only the detector, but also 8 service quarter panels (SQP) carrying a Patch Panel 0 for each half-stave (disc sector), as well as cables and cooling pipes up to PP1 at the end of the Pixel package. Finishing up the package included final connection of all cables, optical and electrical, within the package, which needed testing. This was done in a first connectivity test in early 2007 using final support electronics and cables. A second connectivity test was run in 2008, during connection of the detector with its final cables inside the ATLAS cavern, to guarantee good connections within the calorimeter volume.

5.3.1 Connectivity Test 2007

The first connectivity test in 2007 was to check whether all connections inside the package were done, such that the detector package itself was fully functional before being delivered to the ATLAS cavern. The detector was positioned on the integration tool inside SR-1 that was used to assemble the package. Each service quarter panel was tested right after installation, to ease reconnection in case a malfunction was detected.
A simple connection test checked that no connections were open, by safely powering the corresponding lines and reading current responses. As the detector had undergone no irradiation up to that point, there was only dark current on the high voltage (HV) supply lines, which could not be significantly measured to be non-zero. As HV was also served for a full PP0, the conclusion of a measured current could only be that at least one module has an HV connection. Therefore the detector had to be started up and a small threshold scan to be run, showing whether modules were noisy or quiet (this being the absolute indication of HV bias).
Optical connection was checked by running an InLink scan as described above. If a channel was found to be dead, it was checked for a functional TX channel by running a LoopBack test using 1 m fibres, designated for testing. Only if the TX was measured ok, opening of the SQP or unplugging a fibre at the Optoboard was taken into account. An InLinkOutLink scan pinpointed malfunction of the data path. As for the TTC path, initial checks in case of malfunction were done on all parts outside the detector, until no option but reworking of the SQP was left. Both these optical tests were the first pieces of software to use DAQ DCS Communication.
As the threshold scans required the full readout chain to be up and running, that opto-connection was not only tested, but also tuned. The first automatic optolink tuning procedure was therefore set up during the 2007 connectivity test to tune all parameters of the data link. For each PP0 a scan over three parameters was run, V_{Iset}, RX-threshold and RX-delay. The clock based BOC scan mechanism was utilised, as DSP based methods were

not available at that time. From the shape of the non-functional parameter sets within the BOC scan, the ideal (most stable) parameter set was extracted for each module. The very same algorithms were used during the second connectivity test in 2008 and provided tunings that were found stable for months.

5.3.2 Installation and Early Testing

The Pixel Detector package rested unconnected within the Pixel support tube for about 6 months, as connecting the Pixel Detector needed the other inner detectors to be ready with cabling. Therefore the Pixel Detector could not join the cosmics runs in 2007. To still be able to commission the higher level readout system, a simulation of a Pixel Detector trigger response was embedded into the ATLAS Pixel RODs [DD07, JJF+06]. It ran during the 2007 M5(x) cosmics runs, delivering first Pixel Detector data to the higher level readout system. It also enabled the Pixel community to run checks of the monitoring processes, as the simulation delivers specific output patterns. Eventually it moved into the stable release of the ROD firmware and is now a standard part of the ROD's functionality, available for quick testing of data taking and monitoring. An example of the occupancy produced across a module is given in Figure 5.9

Pixel Simulator

The Pixel Simulator is embedded into the ROD's formatter FPGAs. It provides four serial inputs to the ROD formatter logic itself. Depending on the simulation mode it can be sourced by four engines at 40 Mb/s, two engines at 80 Mb/s or a single engine at 160 Mb/s. For every trigger signal received by the formatters, the simulator stores locally generated BCIDs, depending on the number of trigger repeats, and a single L1ID. As soon as these are available, a state machine starts generating a header and a number of randomised hits. The occupancy, as well as some special flags and error generation can be configured within two 32 bit configuration registers.

Initial tests with an occupancy of 7 hits per module and Bunch Crossing (with single Bunch Crossing readout) allowed the Pixel Detector readout system to run at trigger rates of up to 175 kHz, thereby proving that the readout system has headroom for the expected occupancy up to full LHC expectation[h].

[h]The LHC Occupancy expectation from [B+02] is 0.4% in the central B-Layer, giving about 18 hits per module with 75 kHz triggern rate. Including module headers and trailers, this is less transfer and processing rate that given by 7 hits per module at 175 kHz.

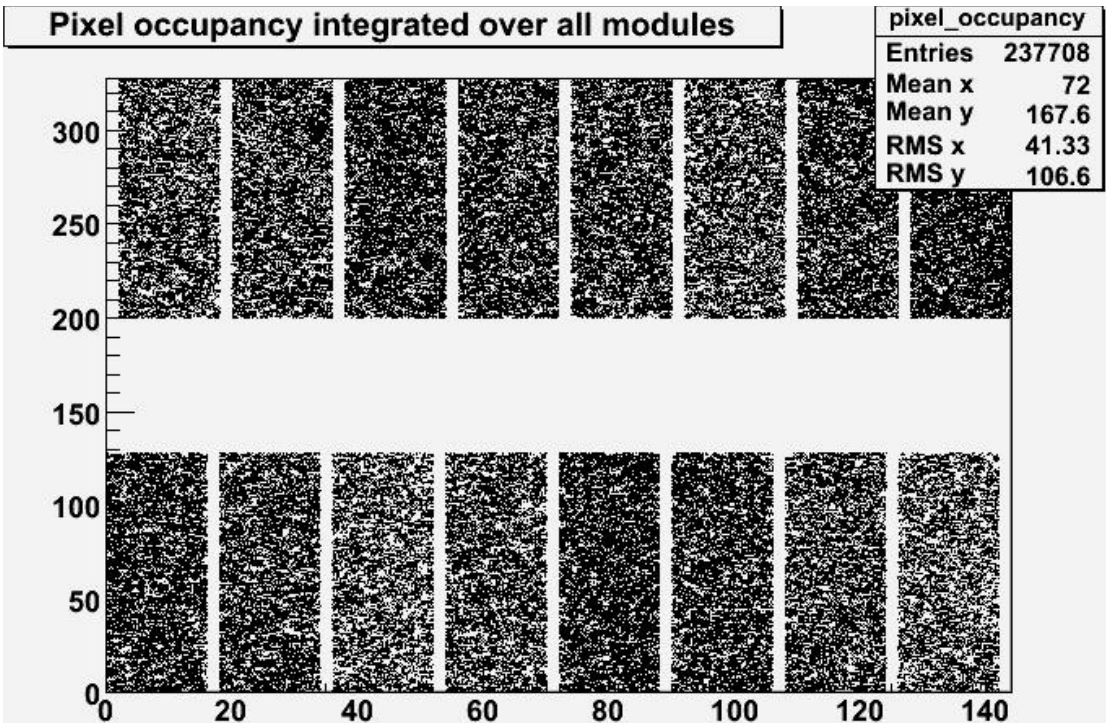

Figure 5.9: *A typical occupancy plot of the ATLAS Pixel Detector modules during the M5 and M5x cosmics runs, utilising the Pixel Simulator. Coordinates shown here are the module wide column and row. The simulator only generates data on columns 0-15 and rows 0-127 per frontend, hence the empty columns and rows which are visible.*

5.3.3 System Tuning

As the connectivity test software was hard to maintain, not scaleable and barely user friendly, a new system had to be set up for permanent usage within the ATLAS Pixel DAQ software framework.

A branch of DAQ software was enhanced during 2007 and 2008 to serve a new user interface, the Calibration Console, together with a well structured calibration backend, the Calibration Analysis framework (CAN). The latter allows plugins to receive both scan configurations and results, and run an analysis that can store results within a globally accessible database. Analysis plugins are written by the corresponding expert, delivering options for the user to adjust and outputs to easily understand a scan[i].

An implementation of the aforementioned Optolink tuning procedure within the CAN framework was done, named *OptoLinkAnalysis*. It can be run on any scan and will try to

[i] e.g. a Pixel threshold scan analysis delivers the mean threshold per module as an output value, which can then be viewed for the full detector within the Calibration Console, allowing easy identification of malfunctioning modules.

collect histograms that fit the BOC scan type, RAW_DATA_DIFF_2[j]. The analysis is called per PP0, so to acquire histograms for a full block that is to be tuned at the same time, due to their common V_{Iset} setting.
Receiving histograms, it checks for the parameters in use, which should be RX-delay and -threshold. In case a 3D BOC scan is performed, the analysis receives multiple 2D histograms for different values of the third variable, which should be V_{Iset}. Detection of 2-dimensional or 3-dimensional scans and corresponding tuning is done automatically.

2D Tuning

Running a 2D analysis, the plugin places all RX-thresholds at the minimal horizontal error sum. In case the mean RX-threshold value chosen for all modules is below a minimal value, a warning is generated. A different warning is generated in case the thresholds are all very high and the scans seem to be saturating. RX-delays are tuned by placing them at a definable distance away from the delay error band, which is identified as the bin with the highest error sum. In case multiple bins are holding the same maximal error sum, the right-most of them is chosen as the base, including a wrapping condition in case the error bar spans between minimal and maximal RX-delay.
After running this tuning based on error projections, a final check is done on whether the tuning has located values that are within definable distances from any error. Minimum distances to right and left as well as above and below can be given.
A 2D tuning is generally used to verify recent tuning settings and perform a simple retuning.

3D Tuning

The 3D Tuning checks for the ideal V_{Iset} while, at every step in V_{Iset}, repeating the steps of the 2D tuning. This is done by checking the low and high threshold warning conditions, generated from the 2D tuning, for every step of V_{Iset}. Going from lower to higher V_{Iset} values, the highest value without high threshold warning is chosen. Should high and low warning exist at the same time, then the individual modules generating that warning will fail. The V_{Iset} tuning however will disregard them, as long as they exist in equal numbers, so as to not prefer one module over the other, e.g. for a PP0 as given in 5.10.
The 3D tuning delivers better quality results, retuning every adjustable parameter of the optical link. Yet, with stable system, the 2D tuning performs well and takes less time, where the speed increase is given by the number of steps in the 3D scan.

[j]The Pixel DAQ Software framework defines different types of histograms, depending not only on the number of dimensions and the numerical content, but also on the logical content. RAW_DATA_DIFF_2 is a 2D BOC scan, typically running on RX-delay and -threshold and delivering an error count.

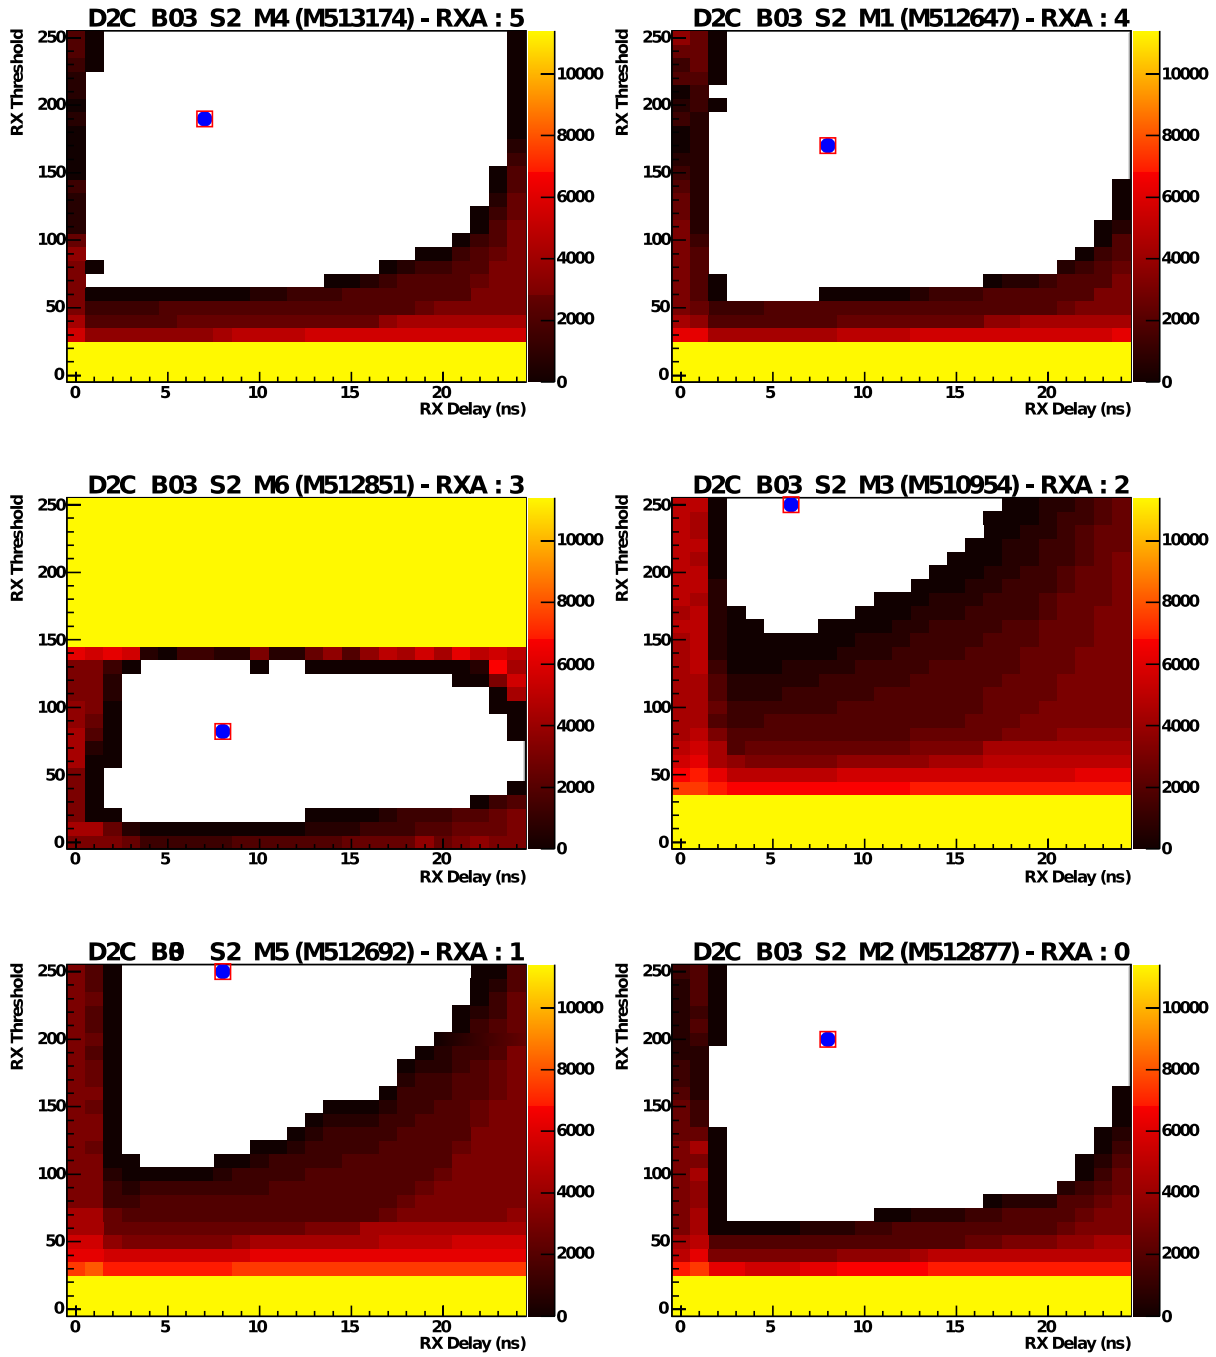

Figure 5.10: *DSP based BocScans for all Modules of a PP0 that is very hard to tune as a full functional Block. Whilst Module M6 is particularly low, M3 is very high in terms of output power.*

MSR Tuning

An MSR tuning happens through a dedicated analysis plugin, *MSRtune*, which finds the minimum value in a scan as shown in Figure 5.5.
The BPM Mark to Space Ratio has been tuned for the first time in November 2008, with only partial success due to many modules loosing clock during the scan. Software changes allowed for a later scan in summer 2009 to generate a full detector tuning, as seen in Figure 5.11. Whilst Figure 5.11(a) shows the ideal MSR settings before setting BPM finedelays for the full detector, Figure 5.11(b) shows the same plots after setting the finedelays. The average change in finedelay was 12, delivering a higher change in the mean ideal MSR value than expected from Figure 5.6.

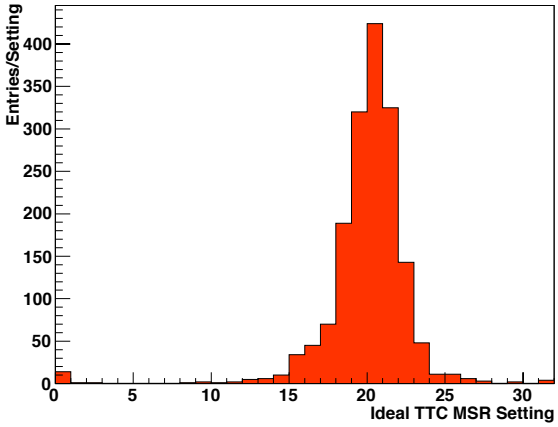

(a) Ideal MSR setting before BPM finedelay tuning

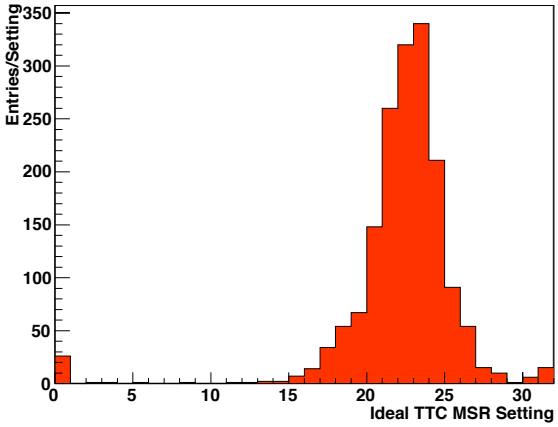

(b) Ideal MSR setting after BPM finedelay tuning

Figure 5.11: *Ideal MSR settings [DFH+10], as measured from an MSR tuning. Figure (a) shows the ideal MSR before setting any specific finedelay values. Figure (b) shows the ideal MSR after setting the BPM finedelay - the average change in finedelay was 12 settings higher than default.*

Additions

Besides tuning the optical link, the OptoLinkAnalysis allows for debugging the optical link. The BOC configuration (RX-delay and -threshold) that was recent when running the scan is checked for being functional or not, by checking the surroundings of the settings for errors. In case any are found, a flag is set that calls for re-tuning the link. Also, the analysis flags single error spots that are surrounded by an error free region. These have been observed rarely with BOC scans and could not yet be traced to any particular software or hardware source. Even if they are close to a tuned setting, the links are functional with

those settings. Hence these standalone spots are not an issue. Bookkeeping of how many of those spots were observed per module is done, which allows the user to easily identify whether errors happened within single links or a whole PP0/ROD.

5.3.4 Performance

After tuning the ATLAS Pixel Detector optical readout link, stable operation was given for an initial period of several months. A tuning based on the OPTO_TUNE scan succeeded in tuning 95% of the Pixel Detector optical links without further attention [DF08]. After a pattern based BOC scan and running the OptoLinkAnalysis, only two out of 1674 operational modules had to be tuned by hand in November 2008. Ever since then, the only change to the general tuning procedure was introduced by the DSP based BOC scan, which now allows the full detector to be tuned without an initial OPTO_TUNE scan in less than 20 minutes.
Two systematic problems exist within the data link:

- Slow turn-on: Slow turn-on behaviour is to date the biggest hindrance in tuning the ATLAS Pixel Detector optical data transmission links. An example of the troubles it brings is given in Figure 5.12 showing a link that runs well, when running with a 50% duty cycle. Using pseudorandom data on the contrary, it only operates well at a much lower ideal threshold and has a smaller operating range. To date no scan exists to characterise this behaviour and choose a best threshold based on the slow turn-on characteristics.

- DRX-12: The DRX-12 features a large time-constant when incoming light changes its power. This is most crucial when the optoboard delivers a large power spread and when lines are driven at 80 Mbit/s, and leads to saturation of a link. It is most probably due to the very high input resistance of about 3.2 kΩ, in combination with a relatively large input capacitance and a slow section of the connected PiN diode. Setting up a new link, this receiver should be replaced with a transimpedance amplifier stage. This would remove any charge or discharge effects of the small PiN diode capacitance, as it would be kept at a constant voltage level. Later amplifier stages could then apply a threshold as is now.

TX Plugin Failures

During the first operation period in 2008, off-detector transmitters started failing, due to their lasers dying. This continued in 2009 and first investigations showed compatible behaviour in devices killed using ESD [Wei08]. The forward I/V curve of dead lasers is

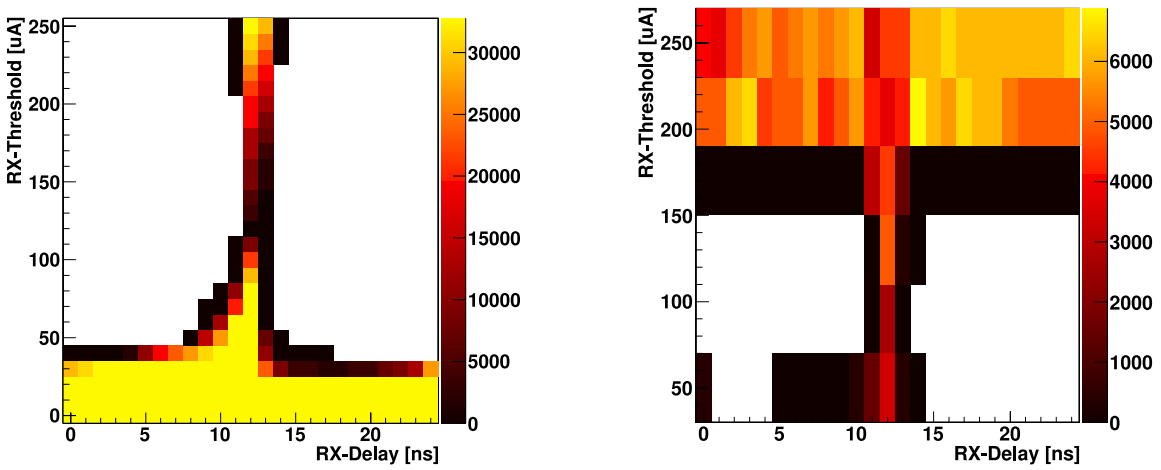

(a) Fast BOC scan

(b) Slow BOC scan

Figure 5.12: *Shown are two BOCscans on the same link, with the exact same settings and environment. On the left is a "fast" mechanism based scan, showing a good signal. On the right is a "slow" mechanism based scan for the same link, showing a very different picture. The top error band is due to Slow-Turn-On effects. (Mind the change in vertical scale due to different scan setups)*

different from a functional laser, c.f. Figure 5.13. This behaviour is observed directly after zapping healthy laser devices with an ESD gun, whilst they are not connected to a BPM[k]. Eventually seeing the zapped devices die in terms of light output power, it was concluded that ESD had been problematic during the first plugin production cycle.

The full production was exchanged in summer 2009, using much better ESD precautions and initial checks on all lasers. Yet, failures on this production appeared in spring 2010 leaving the conclusion that the devices themselves are problematic. Closer investigations followed and are continuing, together with other subdetectors and even other LHC detectors (LHCb). A general observation on ageing devices is that their spectrum changes with the time they are operated. It is tried now to include the spectral width into observations and thereby identify how quick ageing progresses.

Investigations about laser deaths are ongoing, but the general assumption now is that the laser lifetime is much shorter than expected[l]. A recent state of the investigations is given in Appendix A.2.

[k]The BPMs circuitry catches the ESD zap and thus drastically reduces the risk of ESD damage.
[l]Be it due to environmental parameters or the laser devices themselves

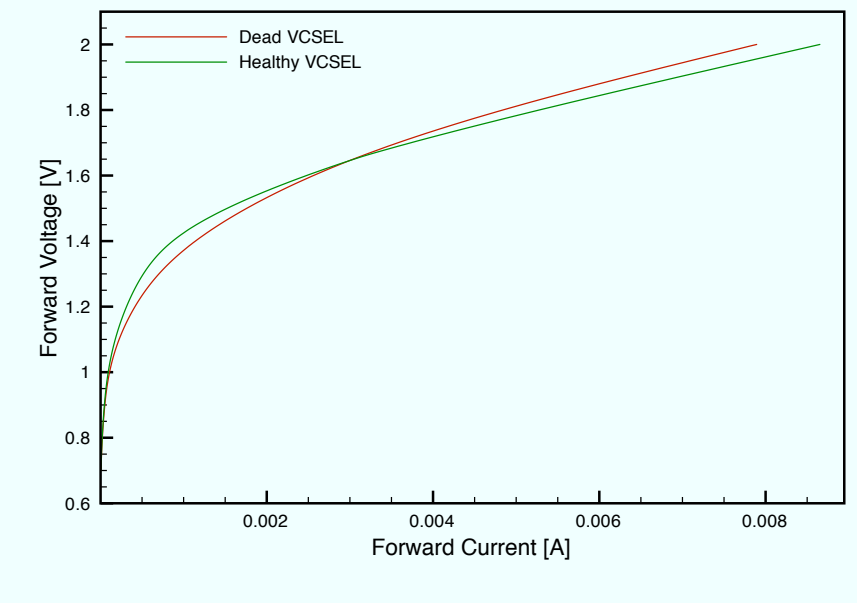

Figure 5.13: *I/V Curve of a functional, and a dead laser within the same array*

Optoboard "Bright" Time Calculation

Due to continuing failures of TX plugins, concerns about the on-detector VCSELs were raised. ATLAS Pixel on-detector lasers are the same lasers as used within the Pixel and SCT off-detector transmitters. There is only one confirmed dead laser on-detector[m], which was dead since very early on. That difference in mortality is probably due to differences in the links and their environment:

- Temperature: An evaporative cooling is attached to the optoboards, as is a heater system. These two keep the optoboards slightly above their minimum stable operating temperature, in most cases around 20°C
- Atmosphere: The optoboards on-detector are placed within the Pixel volume, which is flushed with nitrogen at all times, whilst the off-detector lasers are exposed to humidities between 15% and 55%..
- Link Utilisation: As described above, there is an NRZ encoding on the return link for data. Thus the general state of the Laser is "dim", meaning the Laser runs at low forward current (1 mA), opposed to the off-detector lasers that run 50% of the time at there "bright" current of 10 mA.

[m]Module L2_B22_S1_C7_M6C shows no light, but at the same time seems to produce heat through this channel when being powered, hence the assumption here is that the laser is non-functional.

- Amplitude: The average optoboard "bright" forward current adjustment is 8 mA instead of 10 mA.

In the process of understanding these (de-)acceleration factors, the average bright bits have been calculated [Dop10], assuming a flat random hit location distribution and using occupancies and TOT spectra from online monitoring (As shown in Chapter 4). Including the data stream format, a total of 9.20 bright bits are on average transferred per hit. A header and trailer, plus a single frontend chip number deliver an average of 16 additional bright bits per event. Based on this a calculation was done, extrapolating occupancies by using the instantaneous luminosity forecast given by Steve Myers during ICHEP 2010 [Mye10].
Given in table 5.1, are the numbers used for the calculus, as well as the outcome of the bright time ratio calculation. Coloured columns contain numbers acquired from others, magenta from Steve Myers [Mye10], cyan from the ATLAS trigger community and yellow from the ATLAS Pixel Detector community. All of these numbers are best guesses as of summer 2010. White columns contain the derived calculation results. The occupancy is known for certain beam configurations of 2010. This has been used to estimate a factor between luminosity per bunch crossing and the module occupancy. This number is then assumed as an event calculation basis. Once the total laser bright time per event is calculated, it is multiplied by the trigger repeats and the trigger rate, leading to a total bright time ratio while the LHC is operating. Multiplication of the last two columns and integration over the rows results in the accumulated bright time of the on-detector lasers, as plotted in Figure 5.14.
Based on this estimate, a best and a worst case scenario have been sketched, which change numbers of trigger repeat and the readout speed to better and worse conditions. Both are shown together with the standard case in Figure 5.14. Also the plot includes the initial assumption of 0.5 (Best: 0.25, Worst: 1) month total bright time, which the detector might have seen over the years 2007-2009 of testing and commissioning. Also a per year amount of 2 (Best: 0) bright days for calibration was added, which is, as far as estimates are concerned, more than needed. Yet, it does not really affect the total numbers. The Hübner factor[n] assumed for the LHC operation is 0.6 in all these calculations, which is probably worse than will be[o]. Recent LHC runs mostly show very good beam lifetimes, hence the large estimates.
A task force to understand the problems of the Pixel and SCT data transmission system has been called to investigate on possible scenarios to circumvent major losses of detector parts [GGP10]. Assuming the worst case scenario, the only LHC shutdown allowing for

[n]The Hübner factor describes the effective ratio between a colliders on-time and it's actual run-time. For LEP the Hübner factor was 0.2, where as for the LHC, 0.4 is expected due to longer beam lifetimes.

[o]Worse in speaking for the optical transmitters here

Table 5.1: *Calculation of Laser bright times: The LHC luminosity outline given by Steve Myers at ICHEP [Mye10] in magenta, the ATLAS Level 1 trigger rate estimate in cyan and the Pixel trigger repeat estimate in yellow. All of these are conservative numbers. The data rate was assumed to be 40 MBit/s in [2010;2013], 80 MBit/s in [2014;2018] and finally 160 MBit/s beyond 2018. The estimates were done for B-layer modules. The green box is highlighted for it shows the module occupancy at 2.2·$10^{34} cm^{-2} s^{-1}$, calculated from recent monitoring data and luminosity increase. It fits very well with numbers coming from simulation, assuming about 42 hits/module.*

Year	Instant. luminosity [10^{30}/cm^{2}s]	Colliding bunches #	Module occupancy [Hits/BC]	Trigger rate [Hz]	Trigger repeat #	Bright ratio	Run-time [Days]
2010	10	48	1.25	10000	3	0.0206	60
2011	180	796	1.36	25000	2	0.0356	240
2013	2600	796	19.60	45000	1	0.2213	180
2014	5000	1404	21.37	75000	1	0.1997	240
2015	10000	2808	21.37	75000	1	0.1997	210
2017	10000	2808	21.37	75000	1	0.1997	240
2018	17000	2808	36.32	75000	1	0.3290	240
2019	21800	2808	46.58	75000	1	0.2088	210
2021	40000	2808	85.47	75000	1	0.3769	150

exchange of parts is in 2012. Therefore a project was set up aiming to be ready for replacing the Pixel Detector readout electronics on-detector in 2012.

Laser Tests and Optoboard Lifetime Measurements

A series of investigations has been started to investigate on the early laser death and see whether this is a pure off-detector effect or whether it happens to the on-detector components too.

To compensate for the different conditions on- and off-detector, multiple test are to be carried out to identify behaviour of the very same laser under different environmental conditions:

- Duty Cycle Dependance: Dependance of the ageing on the duty cycle used to drive the lasers is needed. Spurious bit patterns like they are now sent from the on- to the off-detector components might lead to much faster or slower ageing.

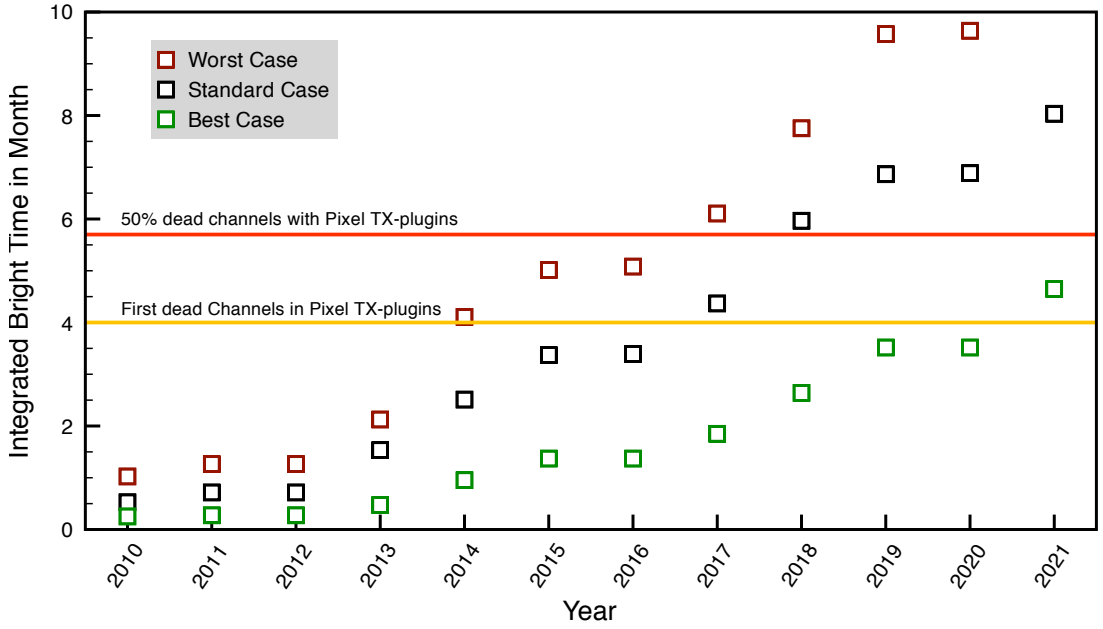

Figure 5.14: *Bright time plotted versus years of LHC runtime. Numbers were calculated from Table 5.1 and indicators are given for when the second TX-plugin installation batch started dying. A 50% dead estimate is given too.*

- Temperature Dependance: The activation energy[p] of the laser devices is yet unknown, but is needed to apply standard ageing formulas like the Arrhenius factor. Therefore the temperature dependant lifetime of the devices has to be measured. This should be possible within multiple high-temperature tests, running for a much shorter time than a lifetime test under realistic conditions.

- Humidity Dependance: First indications on humidity being an affecting factor were observed. Thus studies have to be done within high, medium and no humidty to identify whether humidity is a source or accelerating factor of defects.

As long as no clear indicator of ageing is identified, all of these tests have to be carried out, until a lifetime can be estimated. Therefore multiple devices need to be tested to their end-of-life. Due to the second charge of TX-plugins only starting to die after 4 months of accumulated bright time, estimates say that first indicators of tests started in October 2010 should soon be available. Yet, due to the very different conditions on-detector, especially the much lower humidity and the, in general lower, forward current, final results about whether the Pixel On-detector lasers will die before a possible upgrade in 2016 will probably not be available before the decision has to be taken to remove the Pixel Detector in 2012.

[p]For ageing effects in semiconductor devices see Section 4.3.3

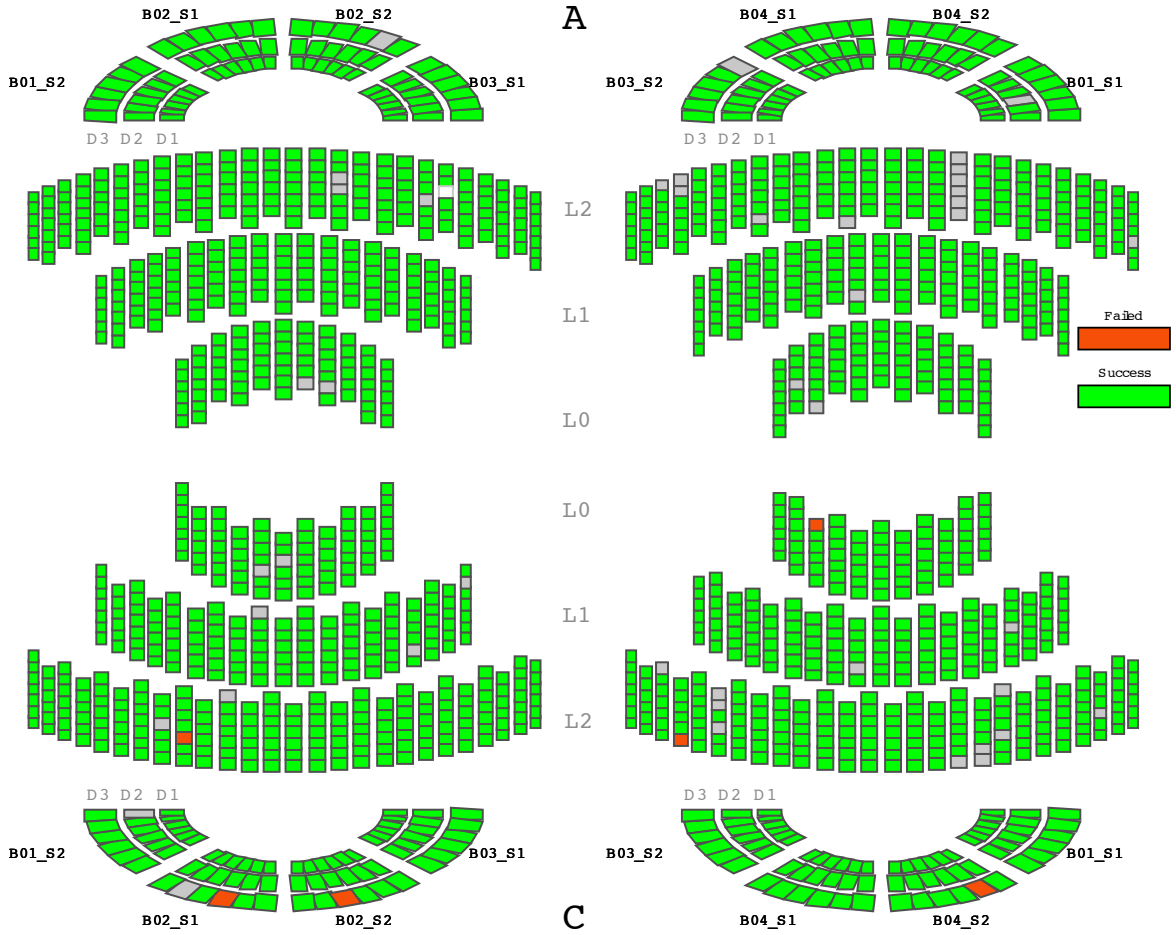

Figure 5.15: *Optolink calibration status of the ATLAS Pixel Detector as given out by the CAN based OPTOLINKanalysis before the LHC turnon in 2009. Green modules pass the tuning analysis, in this case run in 3D with 40 Mbit/s, whilst red ones fail. Grey modules were not operated.*

Assuming only the different currents directly affect the ageing, there is a bright-time factor of 0.64 between TX-plugins and the on-detector lasers ageing, the latter ageing slower.

5.4 Conclusion

The ATLAS Pixel Detector optical links successfully operate 98% of the ATLAS Pixel Detector. Calibration software is in place to automatically set up all but a few links. Typically less than 10 links fail the automatic tuning procedure. Those are for example shown in Figure 5.10 or suffer from bad slow turn-on defects, which cannot be coped with, with the recent tools in place (c.f. Appendix B.3).

The optical link fails to operate a total of 32 modules[q]. For some of these errors the location cannot be given, as it cannot be pinpointed from outside the detector. TTC link failures are seen on 13 modules, 3 of which cannot be configured, whilst the others do not receive clock. The Data link fails on 18 modules, 6 of which are due to a broken V_{Iset} line. Nine of the remaining failing data links send out no light at all, while on 3 an LED mode can be observed. Out of the given 9, only 1 has a failure location that is probably inside the transmitting laser. A single last module is giving issues with calibration scans, the failure location is unknown.
The early observation of dying channels within the off-detector electronics lead to a calculation estimating when the on-detector lasers would fail. Though the failure mechanism is not yet understood, there are good indications that the situation within the detector is much better than outside in terms of environmental influence and utilisation. Though the calculus might be too conservative, the estimated start of dying channels is before the intended phase 1 shutdown in 2016. Therefore measures are taken to produce a replacement of the recent readout system that moves the optical components away from the detector package towards the inner detector endplate. There, replacement should be easily possible in a relatively short phase 1 shutdown, in case newly built devices start failing.
Observations made throughout design, installation and initial operation, allow one to expect certain changes for a future system:

- Balanced Links: As demonstrated by the TTC system, balanced links with automated clock / phase reconstruction are operated much simpler than links that ought to be tuned manually.
- On Detector Conditions Monitoring: The actual Pixel readout system makes it quite hard to tune the On-Detector clock reconstruction. This should be overcome within a coming system, by regenerating the on-detector clocks.
- Component Lifetime Data: The general statement on data transmission Lasers in ATLAS should be that a supplier should always hand over reliability data for their devices.

[q]Most recent value from June 30th, 2011.

6 Insertable B-Layer

As described in Chapter 3.1.2, the LHC will undergo a first stage upgrade, planned for 2016, to increase the instantaneous luminosity. This will deliver a higher discovery potential throughout the following years, by delivering a higher integrated luminosity, and also prepare the LHC for an upgrade towards HL-LHC.
With a planned luminosity of up to three times the LHC luminosity, pile up of events will increase up to a mean of 60-70 events per bunch crossing. The total number of tracks will go up and start to decrease the tracking systems efficiency due to high occupancy, especially within the innermost layer of the Pixel Detector. The column drain architecture of the FE-I3 (c.f. chapter 4.2.3) will lead to a steep increase of readout inefficiencies at an occupancy of 0.5 hits per double column and bunch crossing.
To enhance separation of tracks and primary vertices, a new detector layer will be inserted into the recent ATLAS Pixel Detector. This Insertable B-Layer [C$^+$10] (IBL) will be able to withstand higher occupancies effectively and deliver an increased resolution. It will therefore aid in solving ambiguities, disentangling multiple primary vertices and increasing b-tagging performance.
This chapter will discuss the experimental boundaries and how IBL fits into the concept. Particular focus will then be put on IBL electronics and its readout system. An update on the latest schedule is given in Appendix A.1, whilst here, the schedule as of late 2010 is being reported.

6.1 Experimental Setup

The ATLAS Pixel Detector B-Layer had a replacement foreseen since 2003. The total acquired dose over full LHC runtime was supposed to have the B-Layer remaining too inefficient to actually profit from it being in place. Yet, the replacement is an installation problem [CM09]: Running the LHC for at least 5 years, the detector will need to "cool down"[a] before it can be opened. The Pixel Detector would then need to be transported out of the ATLAS cavern, refurbished, tested and re-installed into the experiment. All of this does not fit into the shutdown schedule of 2016 and is thus disregarded.
Instead it was decided to build a new beampipe, delivering an outer radius of less than 3 cm. This allows one to install a first detector layer at a radius of 3.3 cm. The full package

[a]The detectors material will be activated

of beampipe and new detector layer could have a radius below 4.8 cm. Hence it could be inserted into the recent Pixel Detector, as shown in Figure 6.1, including a separation tube, the Inner Support Tube (IST).

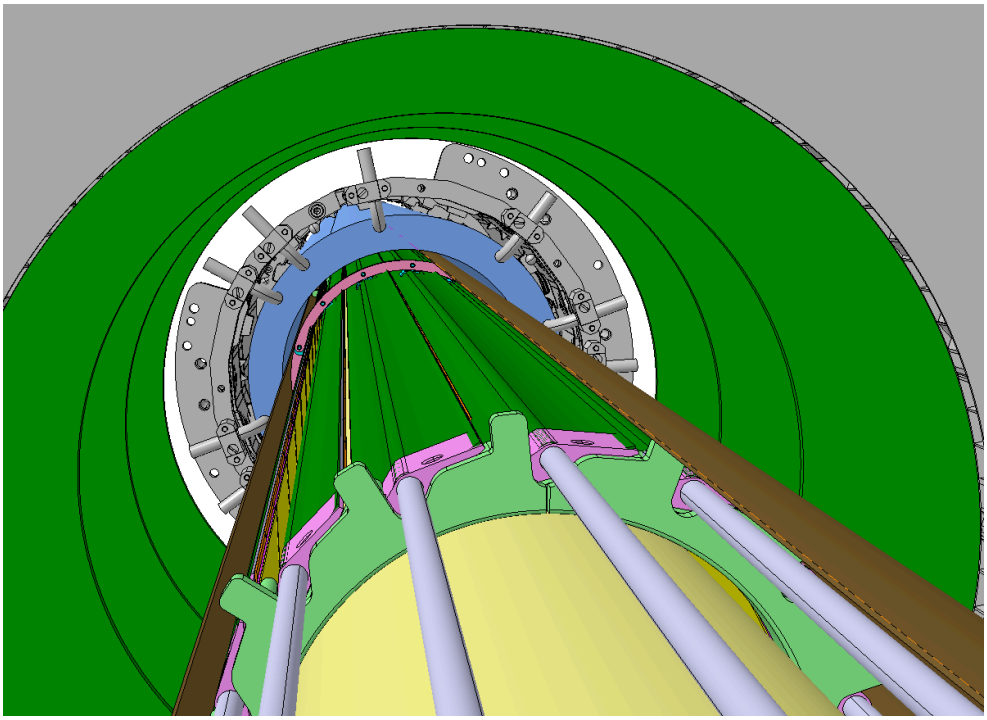

Figure 6.1: *Schematic drawing[C+ 10] showing the insertion of IBL, coming from the lower right hand side, into the Pixel B-Layer (light blue).*

The IBL will only serve a barrel part with a total length of 64 cm, covering an η-range up to $|\eta| < 2.9$. It is planned with 14 staves, each carrying 16 modules. A module is to be built from one or two sensors and two frontend-ICs, where the number of sensors depends on the chosen sensor technology. A module is therefore not defined by sensor or frontend chip, but by the TTC link it connects to - Each module is given a single TTC link. The IBL module surfaces will be parallel to the beam axis, as shingling modules in z is not an option within the very tight space, given by the new beampipe on the inside, and the IST on the outside of IBL.

6.1.1 Frontend Electronics

The IBL will be placed at a smaller radius than the recent B-Layer, therefore suffering higher occupancies. The FE-I3 is unsuited for this situation, as inefficiencies increase drastically. The FE-I3 column drain architecture tries to transfer every hit to the end of column buffer. For each hit, two clock cycles are needed. As higher row hits are

transferred first within the FE-I3 priority system, the lower rows will not get cleared, in case the occupancy increases beyond 0.5 hits per double column and bunch crossing. A secondary source of inefficiency will then be the readout data rate, that limits the clearing of internal buffers. The total inefficiency is given from frontend simulation and shown in Figure 6.2.

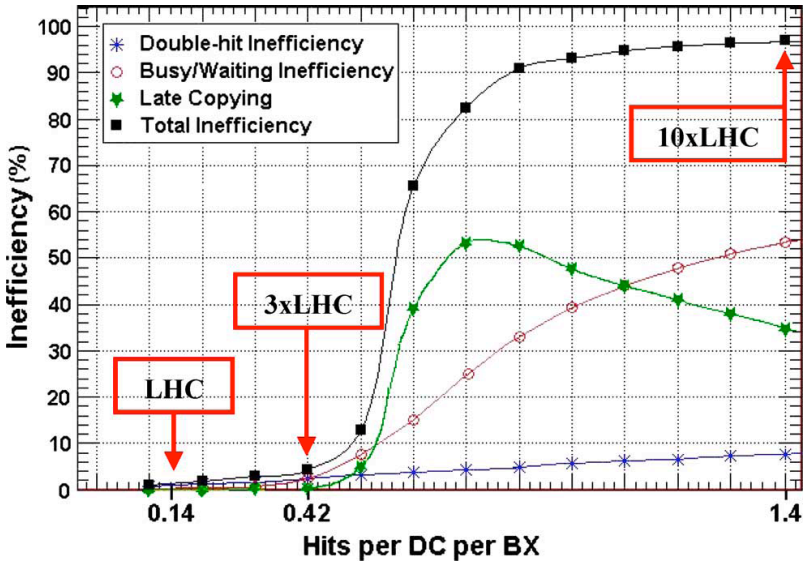

Figure 6.2: *FE-I3 inefficiency [C⁺10] dependent of the double column hit occupancy per bunch crossing. Indicators are given for the estimated occupancy at LHC luminosity, as well as three and ten times the design luminosity.*

Using Table 5.1 from Section 5.3.4, it can be assumed that inefficiencies will arise earlier than 2016, as the B-Layer occupancy is extrapolated to about 20 hits per bunch crossing. Assuming all LHC bunches are filled, which only happens in 2015 [Mye10], this occupancy would have the FE-I3 column drain logic saturated.
Both inefficiency aspects are addressed within a re-design of the Pixel frontend chip, the FrontEnd I4 [A⁺08c] (FE-I4).
A single new FE-I4 provides 26880 pixels in 336 rows and 80 columns. Opposed to the FE-I3 the new frontend works without a module controller chip. Instead it has a direct TTC and clock input and a data output. To build modules from two chips, sharing a TTC connection, the FE-I4 command decoder has an embedded addressing scheme allowing for multiple chips to be connected to the same TTC link. Two data links are given per module, one from each frontend. All connections to the readout chain are low voltage differential signals, running via micro twisted pair cables.
The frontend columns are designed as 40 double columns with a common readout connection to an end of digital column logic (EODCL). Opposed to the previous frontend, hits are stored locally within the pixel cell. Five buffers hold hits locally until either a

trigger signal is received or a maximum time has passed. An incoming trigger signal is distributed across the full frontend and readout into the end of column is only performed if hits are triggered out. This reduces the bandwidth need on the double column readout logic from transferring the induced occupancy per event to a transfer per trigger, going down from about 30 MHz to 75 kHz. A schematic overview of the frontend is given in Figure 6.3
The EODCL within the frontend processes the incoming hits from the pixel array into an event data package that is encoded using an 8B10B encoding scheme. Processing includes a so-called φ-pairing of hits, which encodes hit pairs into a single hit word, in case they are in neighbouring rows. The final package contains a Start of Frame (SOF) special word in the beginning, followed by event header information to identify the event, the hit data itself and an End of Frame (EOF). It is sent to the off-detector electronics using a single 160 MBit/s differential transmission line.

8B10B encoding

8B10B [WF83] encoding is a scheme to transmit 8-bit data using 10-bit encoded words. The encoding is intended for data transmission lines. It is DC-balanced in the long term, with a maximum difference of 0- and 1-count of ± 2 in every 20 bits. The code allows for a maximum of 5 bits in a row being equal. Hence it has a lower bandwidth limit, allowing for clock reconstruction from the code, as well as DC-rejecting coupling of the transmission line.
DC balancing is established by the line encoder. An internal counter counts up on reception of a '1' and down on reception of a '0'. It is allowed to be either +1 or -1 at the end of each code word transmitted and is called Running Disparity. The standard starting value is -1. Every code word exists in two versions; in case the first version of the coded word has a disparity of +2, the second version has a disparity of -2. Thus the encoder can choose to transmit data such that its internal running disparity ends up with a correct value of ± 1.

6.1.2 Optoboards

A differential connection is planned to go from each module out of the IBL support tube onto the inner detector endplate. From there on, IBL will be read out optically, using a re-designed version of the optoboard. Optoboards will be equipped like B-Layer optoboards currently are, with one receiver array and two transmitter arrays, to accommodate for the asymmetric module connection layout. Eight IBL modules will be connected electrically per optoboard. In contrary to the current optoboards, the optical and internal layout should support a total of twelve module connections. The unused 4 channels will serve as spare optical connections, usable via an input and output switching matrix.

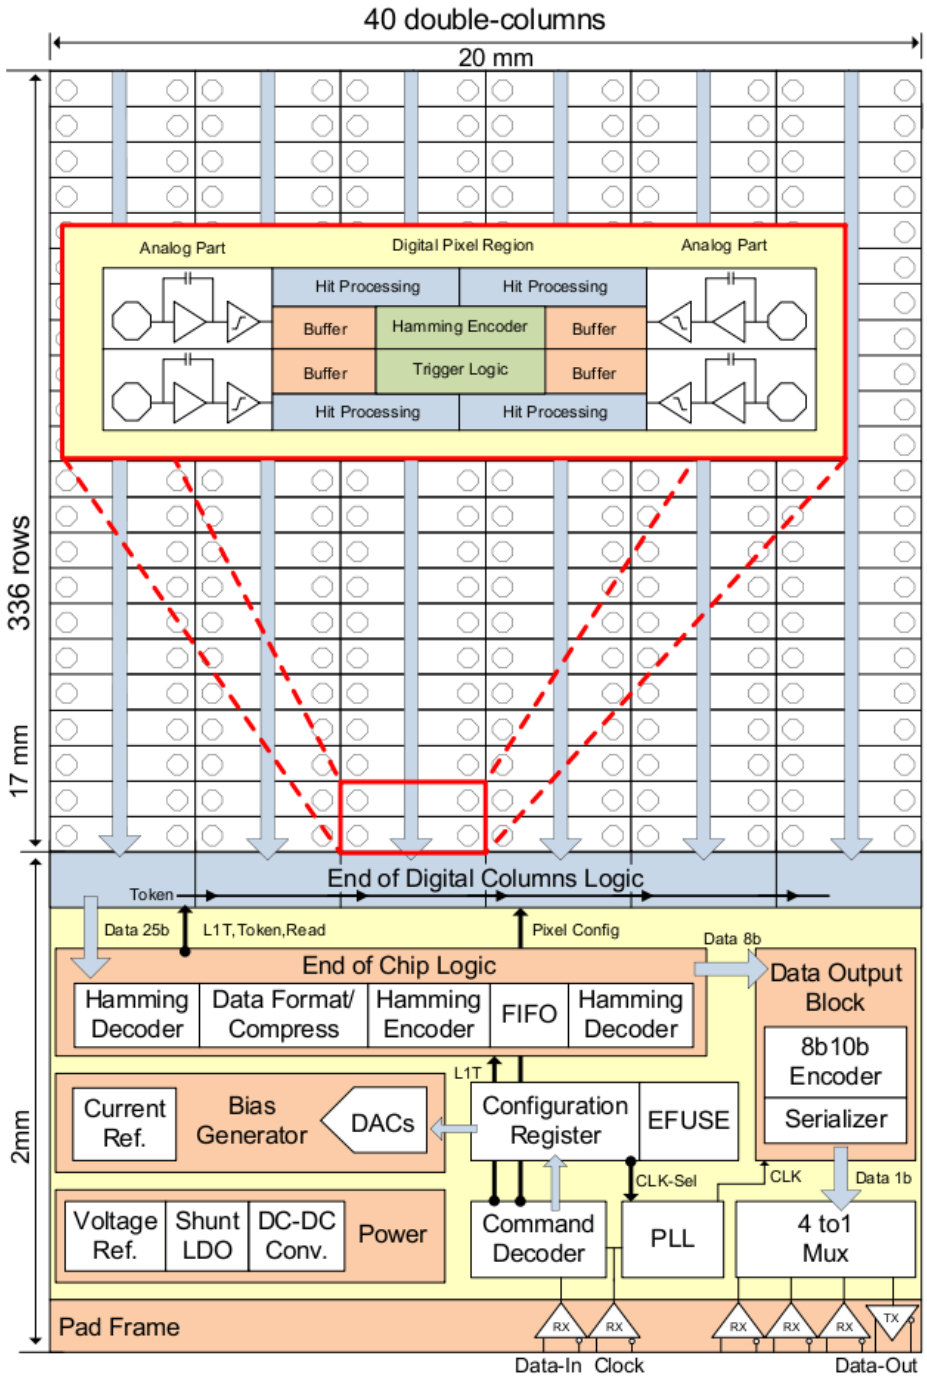

Figure 6.3: *The FE-I4 readout circuit in schematic view. The new pixel area given in 2x2 pixel cell blocks contains the analog pixel cells, buffers for hit storage, as well as a local trigger logic reducing noise occupancy. The end of chip area reads data from the ends of column, encodes φ-pairing and then sends out data using 8B10B encoding. E-Fuses allow to store a permanent chip id and address within each FE-I4. Power regulation circuitry is intended for testing only in the first chip version, but should eventually allow to locally generate voltages needed for the chip.*

A new version of the DORIC-4A is under development, amplifying 12 PiN diode inputs and decoding BPM from those lines. A special version of the FE-I4 command decoder will be attached to the incoming lines, allowing to send commands to the optoboard itself. This will enable the optoboard to receive a mapping of input to output channels, both for the TTC and data links.
The new version of the transmitter IC can be controlled via the same control path used in the new DORIC. In addition to mapping data link channels, the transmitter will allow to adjust laser currents per channel, using an internal digital to analog converter (DAC). The analog part of the laser driver allows to activate pre-emphasis, which means pre-driving the laser before transmitting a "bright" bit, resulting in a much better switching characteristic from 0 to 1. All of this will allow the next generation optoboard to run each optical channel with ideal transmission characteristics.
The receiver and transmitter redundancy schema will help bypassing failing channels. In case of substantial failure rates, as happened with the recent TX-plugins (c.f. Section 5.3.4), this system will allow the exchange rate to be reduce, hence cost and dead time.
As space is very limited within the IBL package, the optoboards are to move outside the pixel support tube frame onto the inner detector endplate. The electrical connection to the modules will be established via micro twisted pair cables that have been tested to data rates of 160 MBit/s.

6.2 Readout System

Due to the IBL being run as part of the Pixel detector subsystem, it is planned to use the same architecture for DAQ hardware, VME. This is primarily to provide hardware compatibility on the lowest level, bus communication and drivers, and the TTC interface. At the same time an extension to the previous VME64X specification exists. VME Double edge Source Synchronous Transfer (VME2eSST[b]) allows data to be transfered on the very same bus that utilises a VME64X connection, but with a transmission rate up to 320 MByte/s. This allows for old and new hardware being used in the same crate, as well as starting the new hardware with old firmware interfaces that already exist.
In the process of upgrading towards higher calibration speed, a communication sideway will be implemented. By allowing the hardware to communicate via Ethernet, not on the bus backplane but outside the VME crate, an additional bandwidth of about 200 MByte/s will be available per slot. This will allow for higher bandwidth parallel readout of histograms and storage into the corresponding software applications. Also it will allow programmers to start moving firmware away from the VME basis towards Ethernet, which currently poses the safest standard for the future.

[b]Standardised as ANSI/VITA 1.5-2003

The hardware composition is to stay the way it is used within the ATLAS Pixel Detector readout system, c.f. Figure 6.4. A ReadOut Driver (ROD) will serve as the intelligent device, communicating with the detector, packaging data for the high level trigger and supporting calibration scans. The Back Of Crate (BOC) card will serve I/O functionality to the ROD, delivering a detector and an HLT interface. The Pixel TTC Interface Module (TIM) will be re-used for IBL, including the third backplane by which it transfers the TTC information to individual slots.

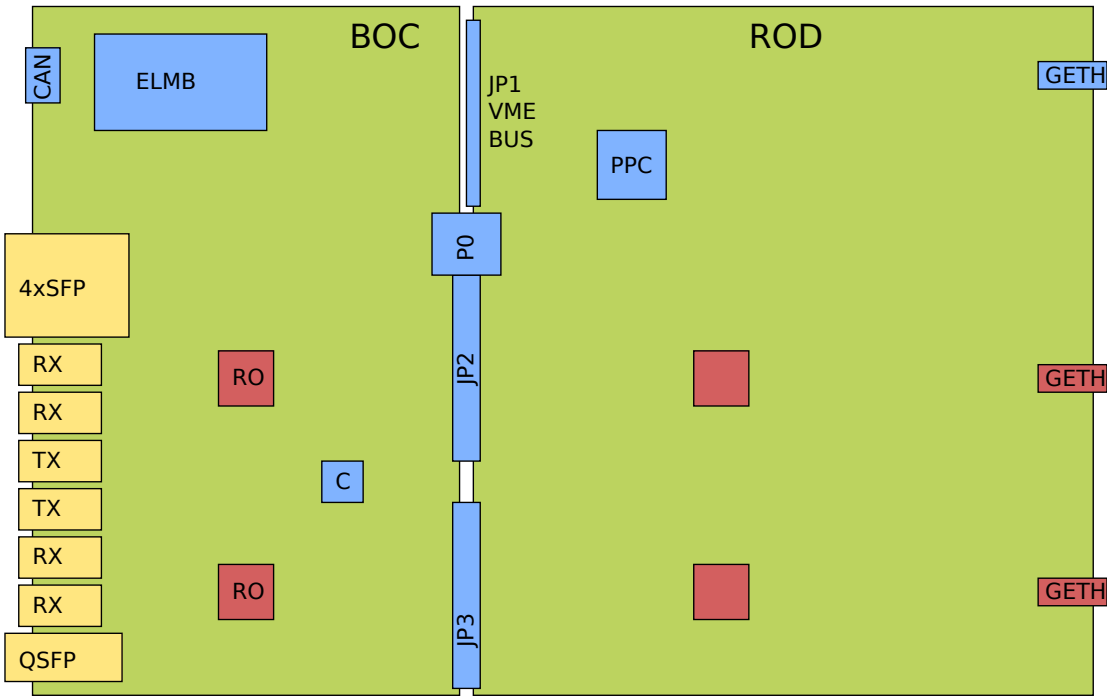

Figure 6.4: *A schematic view of the IBL readout system. All optical I/O is highlighted in yellow, all data processing logic in red, and all (slow) control components in blue. RX and TX plugins connect to the detector, whilst the 4xSFP and the QSFP are transceiver formfactors connecting to the ATLAS TDAQ systems. Memories and electrical buffers are not shown. The image should clarify that the physically dominating component in the readout system is input and output connectors.*

6.2.1 IBL BOC

The Back of Crate [D+11] card is the I/O card to the readout driver. It is to deliver detector side interfaces for transferring Timing, Trigger and Control data into, and physics data out of the detector. Also the BOC is the RODs connection towards the higher level of Data Acquisition systems, namely the ReadOut Buffer INput (ROBIN). Therefore it is to house an SLink interface, connecting to a ROBIN.

Conceptional Layout

In contrast to [D+11], recently the IBL BOC is planned to base on two central programmable logic devices, FPGAs, which connect to all detector, ROD and SLink I/Os (devices marked with "RO" in Figure 6.4). Only the basic register access to the BOC will happen mediated through an alternative programmable device ("C"). This bus interface chip will have a fixed firmware to deliver a stable interface to the BOC. The central programmable devices on the contrary will be re-programmable via the interface device, by programming a local flash with firmware and delivering it to the central devices. Thus the new BOC should be both, save for operation and easy to upgrade in terms of functionality. The BOC will deliver all I/O interfaces within software based (soft-)codecs, such that later exchange is possible. By connecting to all interfaces, re-routing of data for changes in processing will be available from the BOC side.

Soft-Codec TTC Connection

The IBL system will transfer TTC data using the same encoding as was used in the previous system, BiPhase Mark. The major issues with the previous system evolved with time and showed a weakness in the optical transmitter layout. Therefore the IBL is to implement a decoupling between codec and laser driver, utilising the aforementioned software encoder and an off-the-shelf industrial transmitter.
BiPhase Mark encoding itself can easily be implemented within a small logic block, shown in Figure 4.14(a). The crucial part of delaying signals, which has happened within analogue electronics before, is to move into an FPGA too, but will need careful adjustment and checking. Delays within integrated circuits can be based on slow-switching inverter pairs. Within recent FPGAs all logic is optimised for fast switching, which complicates utilisation of logic for generating delays. Yet, hooks to delay signals exist in form of large logic elements serving complex functions like full-adders with carry logic. This carry logic can for example be utilised to generate internal delays in the order of less than a nanosecond, but more than 100 ps.
Performance issue here is to select the right delays, such that a uniform step size can be reached by incorporating multiple blocks of that type. The process to force compilers into compiling with a uniform, instead of a very small delay is still under investigation.

Self-adjusting Receiver

The IBL data path is to be run using 8B10B encoded data with a rate of 160 Mbit/s. Therefore the signal will deliver a minimum switching frequency of 16MHz, allowing to use industrial optical receiver components. These will automatically adjust the input threshold by coupling the first amplifier stage, a TransImpedance Amplifier (TIA), via a capacitance

to the second amplifier stage, the Limiting Amplifier (LA). That AC-coupling delivers an automatic bias current removal, therefore emphasising the dynamic component of the signal.
A conceptual view of the receivers digital logic circuit is shown in Figure 6.5. For each incoming stream, a sampling clock needs to be adjusted or selected, depending on which mechanism is chosen. Using that clock, a deserialized stream can be decoded with an 8B10B decoder. At this stage, losses of synchronisation, bad transferred 8B10B words and running disparity errors can be checked and summarised.
Using the decoder output, Start Of Frame (SOF) and End Of Frame (EOF) special words can be recognised easily and serve to start and stop data processing. As only data within these two special words is needed, the total transmitted data in the later system can be kept low by removing IDLE and also SOF/EOF words at this early stage. As at least 2 bytes are removed from the stream, the data stream has additional capacitance to transport more data than is given by the module. Even more headroom is given by the 8B10B decoding.
In groups of four, these "decoding and shrinking" stages will then feed data into an arbiter that takes care of pushing data into the ROD. As no buffering is set up within the BOC, the ROD will be the first instance to generate timeout events. In case the logic utilisation of the ROD increases, this stage can be moved into the BOC.

Clock alignment: Clock Data Reconstruction (CDR) is the most crucial part of the receiver. It allows to connect the detector without knowing details of its clock source. A first implementation is using clock phase selection as described in [Saw05]. The method uses a local clock, that is within certain limits equal to the remote clock used. It generates four phases of that clock, each differing from the next by 90° and samples the incoming data with each of these four phases. Monitoring transitions between the different phase samples of the stream, the transitions can be tracked and the opposite phase is chosen for sampling. Crucial is to monitor a wraparound in phase selection, going from 270° to 0° or vice versa.
That phase selection has been implemented [Neu11] in combination with an 8B10B decoder and a data buffer. It shows good performance using an FE-I4 emulator as input source, even when running from individual clock sources. The codec implies a minimum and maximum transmission of 4.75 and 5.25 in 5 received bits[c], therefore allowing the transmitter to be off of 160 MHz, the receiver baseline, by as much as 8 MHz without running into any trouble.
A second algorithm for phase selection exists [Die07], based on an external voltage controlled oscillator, driving a phase locked loop. Despite exactly reconstructing the incoming clock, this method has the major disadvantage of using a single clock per receiver channel. As clocks require special routing within FPGAs, this would lead to a very high

[c]The maximum bit length between two transitions is 5 bits within 8B10B. A deviation of .25 refers to a phase step that is not done properly, thus running away from the ideal sampling point.

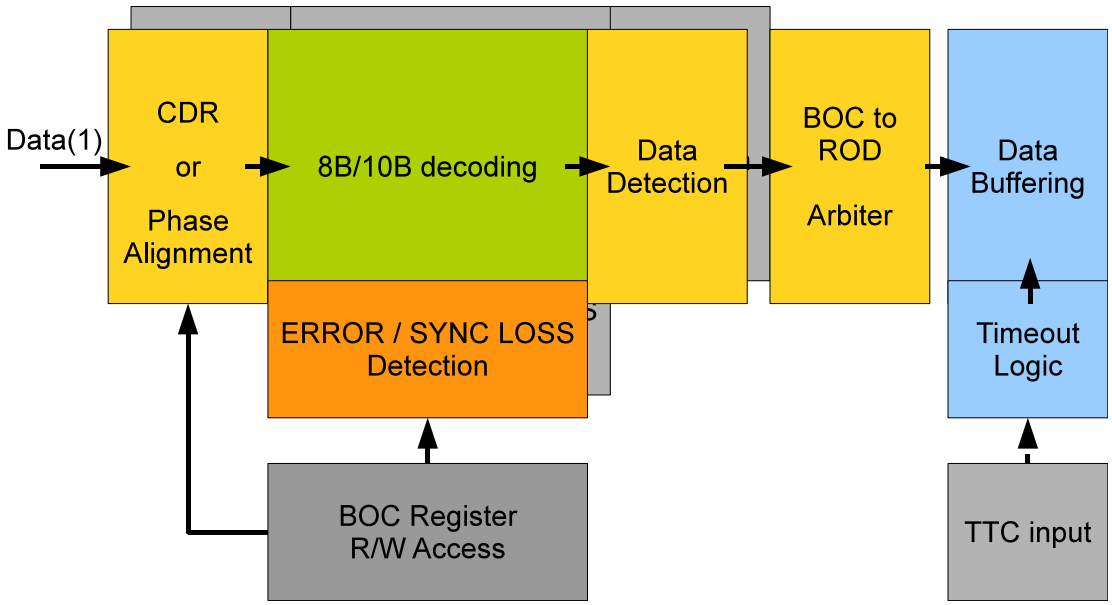

Figure 6.5: *A single channel receiver concept for the IBL. Data flows from left to right, coming in at a rate of 160 Mbit/s. A first block is to align a local clock with the data stream received. Afterwards, the data is sampled and de-serialized. A decoder removes the 8B10B encoding and allows for simple error checks like wrong 8B10B words and running disparity. If a header is detected, the data word is buffered and transmitted into the ROD. Transmission happens through an arbiter that mixes four data receiver streams into one parallel output.*

occupancy within the FPGA. Also this method requires several 100 mW power per channel, ending up in the order of 10 W increased power consumption per readout building block.

Higher Level Trigger Connection

Due to the high total data density, the IBL will not utilize the Highspeed Optical Link for Atlas [RvdBH03] (HOLA) Link Source Card (LSC) to connect with the HLT. Instead the SLink protocol will be implemented within an FPGA using embedded high speed serial links and external optical transceivers[d]. Reprogramming with new configurations will enable upgrades to at least 3.75 Gb/s[e] per link. A single ROD SLink output will be routed into two SLink blocks on the BOC. Both will send data out to different destinations, allowing to produce backpressure. The first destination will be the high level trigger, while

[d]Transceivers will either come as SFP or QSFP: (Quad) Small Formfactor Pluggable - A formfactor for single- and quad-channel optical transceivers.

[e]Recent high speed SerDes blocks within FPGAs support between 3.75 and more than 20 Gbit/s, subject to the chip chosen. Minimum rates differ between about 200 Mbit/s and 3 Gbit/s.

a copy will source a future Fast TracKer [A+09b] (FTK) system, to support the second level trigger system in identifying tracks.

Optical Sub Assemblies

Moving from the Pixel Detector towards IBL, the plan is to replace all optical components needed on the BOC with off-the-shelf industrial components. Based on SNAP12 plugins as defined in a Multi-Source Agreement specification [SNA02] (MSA), 12 channel array links can be utilised within the IBL readout system, as they could be within the Pixel Detector. The plugins come in transmitter and receiver layouts, and serve a limited control capability. The outside form factor as well as the electrical interfaces are defined within the MSA. A plugin for a compliant socket should thus be available from more then one industrial source. Not only is the production path easier - these plugins ease the plug in procedure of the fibre, as the connector housing is part of the plugin itself. Hence the alignment of the plugin versus the motherboard (the BOC) must not be perfect to plug in a fibre.
Typical transmitter assemblies come with laser arrays that include a monitoring PiN diode within each laser. The plugins themselves monitor their output power and can raise single channel error flags in case of malfunction. Reading the plugin status this can be used to automatically re-route data for that particular module via a different fibre, using the IBL optoboards switching capabilities.
Opposed to the transmitters, the receivers have a minmum supported input frequency, probably due to the coupling between first and second amplifier stage. Typical lower cut-off frequencies are in the order of 100 kHz to 1 MHz. Therefore these plugins will work for IBL, delivering a minimum input frequency of 16 MHz[f].

Power Budget

Industrial receiver assemblies are very sensitive and will therefore pose no critical lower limit on the input power. Standard components allow one to operate from input powers as low as 0.05 μW. The upper boundary is a little more stringent, as the on-Detector lasers have shown very good performance. Yet, hitting the boundary of 755 μW with at least 3.5 dB attenuation from the fibre transmission can be circumvented, particularly easy, in case the IBL optoboard will be built to allow single channel regulation and pre-emphasis, both of which are in preparation.
The transmitter assemblies on the contrary are very limited in their range, mostly to be laser class 1M compatible and hence not dangerous. A recent SNAP12 plugin from Avago technologies offers an output power of almost 800 μW. Assuming the DORIC is

[f]Replacing the recent Pixel Detector off-detector readout electronics, in case of failure, with the corresponding IBL electronics will not be possible.

performing well at 67 μA input swing and that the translation factor between optical power and electrical current stays at 0.6 A/W, using 10 V bias voltage, this plugin would allow for a total attenuation of more than 10 dB, expected within the Pixel Detector setup. With the IBL fibres being short and the optoboards being placed outside the very high radiation area, this is underestimated.

6.2.2 IBL ROD

The IBL ReadOut Driver [B$^+$11] (ROD) will connect with up to 16 frontend I4 based modules. It will provide functionality for data taking and calibration, as well as for communicating with the ATLAS IBL BOC card. Three different types of FPGAs will serve for these operations:

- Program Reset Manager (PRM): The PRM is to guarantee stable operation and deliver the basic control interface to the VME bus. It will enable the VME host to reset the other FPGA circuits and load new firmware into the readout driver. The firmware of that device itself is to be frozen and only changed by expert intervention.
- ROD Controller: This device ("PPC" in Figure 6.4) will contain an internal processor, a PowerPC^{TM}, to run primitives[g] requested by the DAQ system. It will also deliver the interface towards the detector and handle inputs from the TTC Interface Module.
- Event Fragment Builder (EFB): A total of four event fragment builders with equal functionality are to be delivered (red data path devices on the ROD in Figure 6.4). These will serve data gathering and forwarding to the HLT interfaces. Each will also come with a histogramming engine to enable monitoring of data and calibration of the IBL.

Data taking

As the data format is much easier to handle with the 8B10B transmission protocol, the data taking logic will have reduced complexity. Additionally, serving a single output speed will reduce the formatter logic effort by about a factor of 2. The IBL BOC will serve the ROD with an input interface delivered as a FIFO write bus:
Eight parallel data lines, an enable flag and two address lines will control data transmission from BOC to ROD. The address should permanently change, allowing the BOC to only register data, but not store it in any buffers deeper than 1 word. Using two times those lines, two receivers will be connected with one event fragment builder. Additional inputs coming from the ROD controller will deliver the ATLAS Event IDs and allow for generating

[g] parametrisable macros

timeout data within the ROD, in case event data is missing. The EFB itself will have to care for gathering data from those two FIFOs and embedding it into a single event fragment with the corresponding IDs. That event fragment should then be routed towards the HLT interface on the BOC.

Calibration

In calibration, the IBL readout driver will read data from the detector, but instead of forwarding it, keep it for internal use. Data will be sampled within a histogramming unit [D$^+$10]. Part of the detector, the size depending on available memory, will be available for realtime sampling, allowing shortest calibration periods. Advanced processing of calibration data (e.g. fitting of threshold curves) will be done after downloading calibration histograms into a standard PC. Opposed to the Pixel ROD DSP environment, the PC will allow standard users to analyse calibration data. Downloads will be done via two ethernet ports, one serving for two EFBs, thereby circumventing bottlenecks on the VME backplane. The histogramming engine incorporates Multiply-Add computational units within recent FPGAs. This results in a high frequency, low logic occupancy firmware running 100 MHz pipelined hit histogramming. In data taking, the histogramming engine will be available to check the full realtime occupancy per module and record a TOT histogram.

6.3 Conclusion

The IBL will help to deliver an increased performance of the ATLAS tracking system within the high luminosity environment of the LHC phase I upgrade. In case of catastrophic losses within the Pixel Detector subsystem, the IBL will help to compensate for those. Aside of that, development done for IBL on-detector data transmission is now returned to the Pixel Detector, in case it comes to replacement of the on-detector readout components as described in 5.3.4.
In contrast to the current Pixel Detector readout system, the IBL system will set up automatically. The operator will thus not be bothered with the technical details of data transmission within the IBL system, therefore bypassing most of the tuning necessities presented in the previous chapter .

BOC Status: To date the firmware blocks to establish TTC communication including analog functionality are well established. A first version of the data receiver block is functional and can retrieve data from an FE-I4 emulator. A slow control bus interface as well as HLT connectivity are being prepared, thereby delivering the later half of BOC functionality needed within the IBL readout system. The BOC PCB is finished in terms of schematic, the layout almost is and prototype production is soon to start. The prototype

should enable early testing in combination with an FE-I4, allowing to control the BOC with a linux core and an ethernet interface installed for debugging and early testing.

ROD Status: The IBL ROD is in the process of being designed. Firmware blocks to communicate histograms to the calibration processing farm, as well as for communicating with the PowerPC^{TM}are being tested these days and show promising results. The finished schematic is handed over for layout and production and should soon be available for programming.

7 Future Readout

As discussed in previous chapters, two measures are of importance when it comes to data transmission within tracking detectors for particle physics:

- Data transmission speed - to deliver data at the given rate, not just through the optical data path, but everywhere in the detector, starting at the hit recording within the frontend.

- Transmission stability / Bit Error Rate - to not affect data quality, in particular links should be stable throughout a data taking run and not require interventions of any kind.

In recent tracking detector designs, data transmission speed is driven by outside parameters. The collider's instantaneous luminosity defines the amount of the triggered content to be transmitted. The event- or trigger rates then give the number of times per second this is to be done. Now the transmission rate can be defined by applying the physical layout of the detector and data encoding sizes. Compared to an estimate based on simulation, the maximum sustainable data rate is typically not chosen much higher. E.g. in the ATLAS Pixel Detector, the FE-I3 frontend chip was only suited for a factor of 3 more hits than estimated within simulation. Higher safety margins would complicate layouts and increase the power drawn by the electronics. It would also lead to higher material effort and is therefore disregarded within inner detector environments.
Stability in terms of Bit Error Rate does not seem an issue yet [Gre01], but is considered to be at particle fluxes of a future High Luminosity (HL-)LHC and beyond [A$^+$01]. Recently, transmission stability has been an issue and needs drastic improvement. A system can not rely on the user adjusting data transmission variables, nor a static configuration due to an analysis. Instead a transmission system should adapt to the given signal. A first approach to self-adjusting systems is being made with the IBL readout system and initial tests show good results.
This chapter will investigate possible designs of readout systems and their usefulness. An initial theoretical approach will summarise capabilities of a link to transmit an amount of physics within a certain time.

7.1 Speed Requirements in Transmission Systems

Talking about transmission systems, it is important to understand what data amount exactly needs to be transferred. The following subsections will elaborate on the content of data transmission within a pixelised tracking detector.

7.1.1 Pileup contribution

Data within high energy physics has to date been dominated by the hard event that generated a trigger signal. Within recent luminosities of accelerators, pileup starts being the main contribution, even and particularly in triggered events. Tevatron has started observing a mean of two additional interactions within data and LHC has recently measured mean pileups of 3.28 events. Assuming in the very same run a per-crossing luminosity[a] of $4.03 \cdot 10^{29} cm^{-2} s^{-1}$, the expected design luminosity[b] of $3.56 \cdot 10^{30} cm^{-2} s^{-1}$ per bunch crossing gives about 29 events on average (which is close to the original estimate of about 25). The number of pileup vertices scales directly with the instantaneous luminosity, thus aiming for $10^{35} cm^{-2} s^{-1}$ in HL-LHC leads to a total of 290 interactions per bunch crossing. Different crossing schemes for the proton bunches, as well as luminosity levelling[c], lead to estimated mean event numbers between 150 and 400.

Those numbers are vague, but the general idea is that future data taking will be dominated by the background generated, not the hard physics process. Data produced will thus follow minimum bias distributions. Due to the higher occupancies within the detector, trigger processes will get harder and need more detailed analysis than is given now. Options here are to have a major upgrade of the L1 trigger system to allow for better precision or to move to higher initial trigger rates and thus allow for a more sophisticated analysis in the offline farms.

7.1.2 Geometry Dependence

Depending on the chosen geometry of the detector, data density coming from events or background can be influenced. The following discussion of geometry dependence is based on Figure 7.1. It shows the charged particle multiplicity from minimum bias events in ATLAS, plotted versus η. The density of charged particles from background is thus almost flat over the range of the inner-detector and therefore the geometry observations can be

[a] That is to say the generated luminosity per bunch pair colliding.

[b] Per bunch, implying $10^{34} cm^{-2} s^{-1}$ / 2808 bunch crossings.

[c] Luminosity levelling refers to the concept of adjusting beam focussing throughout a run, such that the instantaneous luminosity stays constant throughout a run.

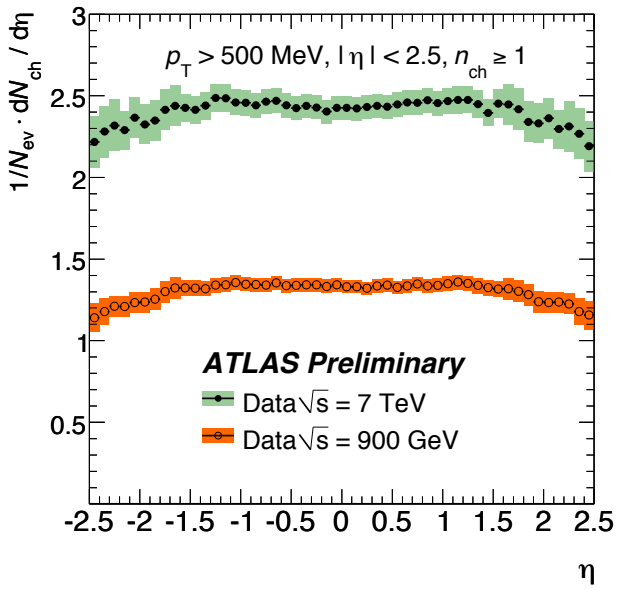

Figure 7.1: *Charged particle multiplicity versus η from first data of the ATLAS experiment [A+10]. The p_T cut has been lowered to 100 MeV in [A+11], giving an even higher multiplicity and a flattened distribution.*

made dependant of η itself[d]. Recent data only allows measurements at 900 GeV and 7 TeV. Estimates for 14 TeV LHC target energy predict a 20 % increase with respect to 7 TeV data. Figure 7.2 shows the covered surface in $\eta \cdot \varphi$ versus the radial positions of a module. This is done here for three different types of areas, Pixel module, IBL module and IBL frontend. All of these are starting at $\eta = 0$ with one edge. The plot gives the installation radius for Pixel and the planned radius for the IBL. As can be seen, the radius of an inner barrel layer will soon be limited by the total track density per module area, which have to be separated and read out.

Not only the radius, also the z dependence has a large influence on interaction driven track occupancy. Figure 7.3 shows Pixel and IBL Modules at Pixel and IBL radius. For IBL, a scaled 12-stave radius is plotted too, showing the increase in density in the innermost module. The effect of lower η coverage at large |z| is compensated by cluster widening: A track originating at z = 0 produces a very wide cluster when hitting modules at the end of staves. E.g. in case of IBL, an end-of-stave cluster induced by a track from the primary vertex is at least 10 pixels wide. This effect is increased in IBL opposed to the ATLAS Pixel Detector, as the latter has a shingling angle in all but the centre module, which slightly reduces cluster sizes at large z.

The equalization of data transmission load is no positive effect, though it eases the design of the readout system: It decreases the efficiency of information transfer in the higher $|\eta|$.

[d] This is limited to the observed range of $|\eta| < 2.5$

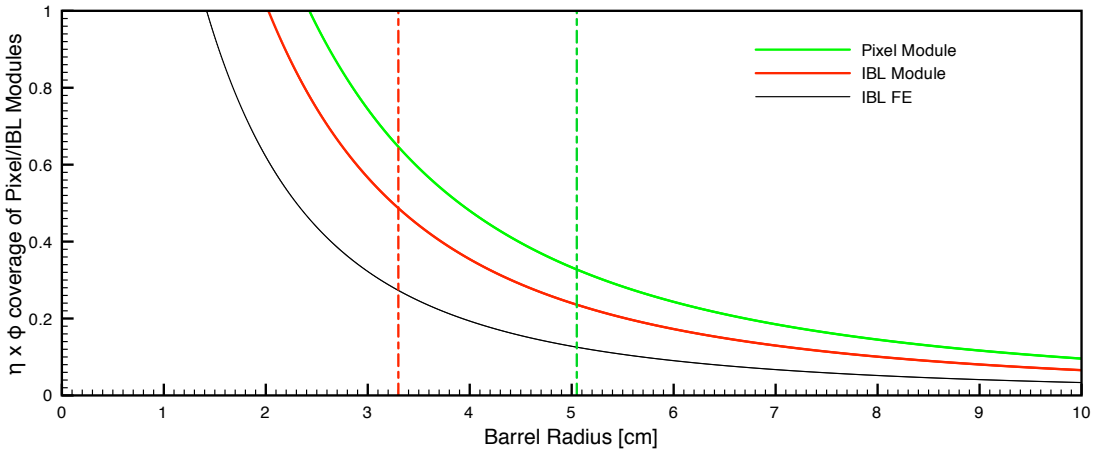

Figure 7.2: *$\eta \cdot \varphi$ coverage of a single module, plotted for Pixel (green) and IBL (red) modules versus the barrel radius at which they are mounted. The dotted vertical lines represent the B-Layer (green) and IBL (red) installation radii.*

Also, the material amount within a particles path is increased at higher $|\eta|$. This increases the amount of multiple scattering and reduces the event data quality in that region.

7.1.3 Parasitics within data transmission/Reducing data

It is seen from first LHC data, that a non-negligible amount of data is halo / beam-gas events, that need to be coped with. The typical signature of these events is an event with very large hit count within the silicon detectors, and particularly large z-width clusters within the barrel part. A big problem of these events is there size, which is abnormally high and eventually leads to problems in the higher level readout chain.
As the hits come from particles passing through the silicon, they deliver a realistic signature in terms of time over threshold, which could just as well come from low p_T tracks originating at the interaction point. A cluster level filter can be applied to reduce the amount of data transmission needed for those parasitics by encoding the cluster width in z.

φ-pairing: Within FE-I4, a double-hit encoding scheme is implemented, transmitting two hits as a single hit word, in case they are neighbouring in φ. Thereby, the total transmission length per single hit is increased by a factor of 1.2, but the same amount of data can transmit two hits, reducing the transmission size for two hits by a factor of 0.6. With larger cluster sizes, the reduction factor approaches 0.6. Using a cluster size spectrum as seen in recent data monitoring (cf. Figure 7.4), a total reduction factor can be calculated. Estimates for IBL assume a total reduction factor of 0.8, whilst for an outer HL-LHC tracker at more then 20 cm radius it could reach 0.74.

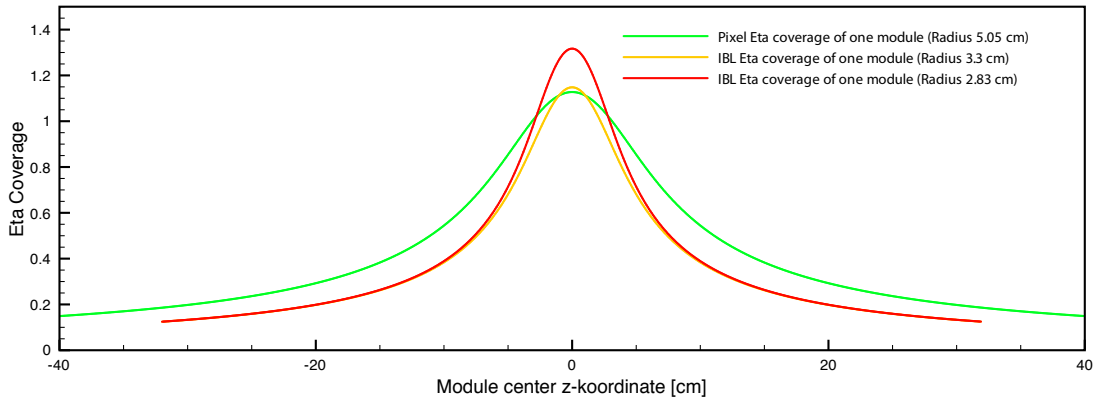

Figure 7.3: *η coverage of a single module, plotted for Pixel (green) and IBL (red) modules versus the z-positon at which the module is centered. A second installation option for IBL is plotted to compare the shape at two different radii.*

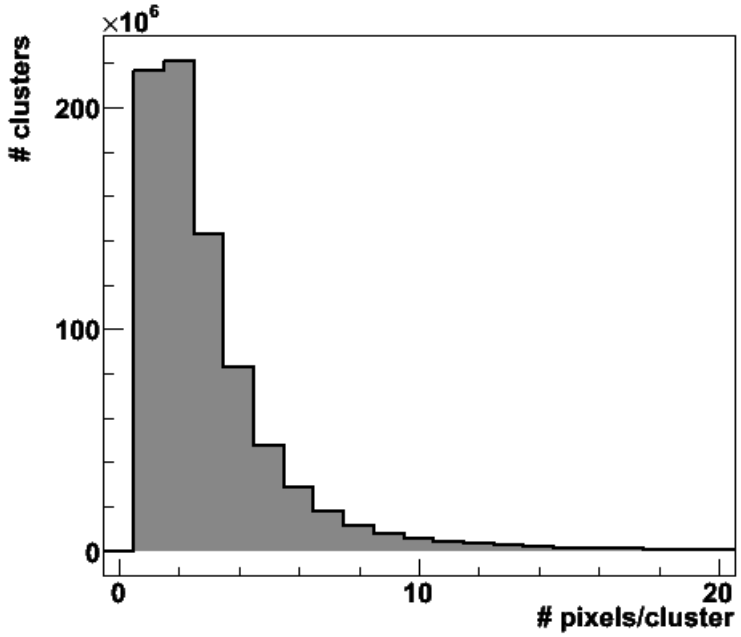

Figure 7.4: *Cluster size histogram from a proton physics run (184022) at 7 TeV. The majority of clusters is larger than one, which makes online cluster compression a reasonable approach for data reduction.*

Within IBL a z-pairing would offer much better compression, as cluster length in z will increase by a factor larger than 1.6 compared to Pixel due to the smaller radius and the smaller pixel size in z. Cluster widths in columns and rows recorded in a recent run are shown in Figure 7.5.

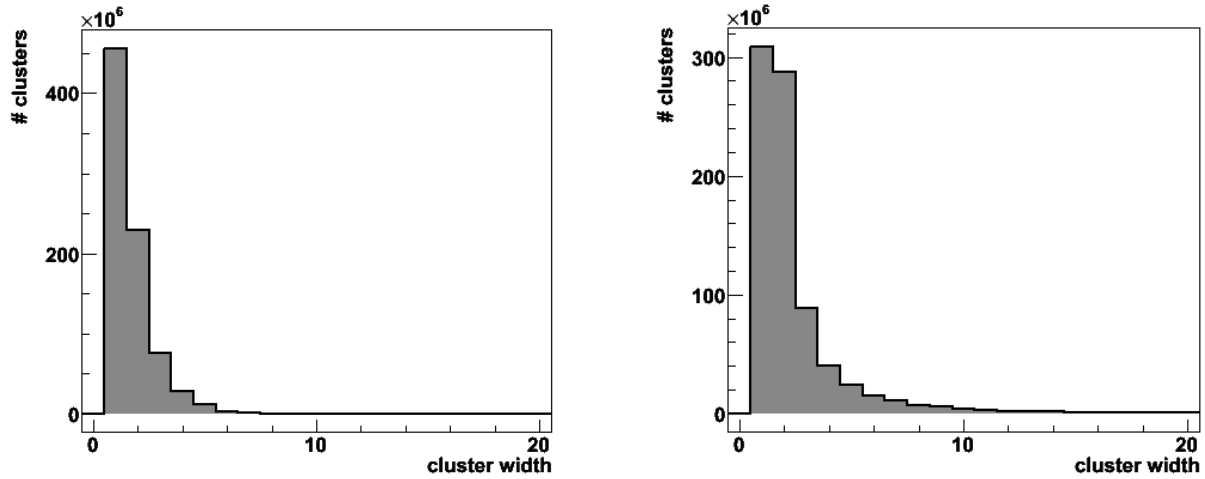

(a) Cluster Column Width (b) Cluster Row Width

Figure 7.5: *Cluster width in columns and rows recorded with the ATLAS Pixel Detector in a recent data taking run (184022), taken from [atl11]*

7.1.4 Higher Level Transmission

Data transmission nowadays does not end at the off-detector interfaces. The current use of custom interfaces puts a hard limit on the maximum amount of data transmitted towards the higher level trigger and DAQ, which in turn might delimit the readout bandwidth eventually.
Data transmitted from the detector to the readout driver is typically expanded whilst being merged into event fragments. This is to deliver a fragment frame, a global ID and some on-the-fly informational content within the event data, to ease processing within the higher level DAQ systems. Also, compressed encoding is decoded on the fly to ease processing by the level two trigger.

7.1.5 Stability requirements

Coming from the current system, it is obvious that transmission systems need more thinking in terms of reliability. Single points of failure create inefficiencies and should thus be

reduced, to zero if possible. This expectation calls for redundancy, in particular in high-speed, high-density systems, where a single point of failure could cause a large part of the detector to malfunction.

As the philosophy in inner detectors is to reduce material to the absolute minimum, embedding a redundant connection becomes more than a technical issue. Redundancy itself can thus not be full (increasing the material needed by more than 100%). Instead a number of spare channels per total channel number should be available, including a switching matrix that allows to move from malfunctioning channels to good ones.

7.2 Gigabit Transceiver

With higher data rates, development effort for custom components becomes larger. Recently experiments developed individual links for most subdetectors. With increased data rates, the effort of design and testing will become overwhelming for the experiments. Therefore a common development is started within the CERN microelectronics group. The goal is to produce an IC family capable of transferring future data loads via an SEU[e] safe protocol.

The project, named GigaBit Transceiver [MMK07] (GBT), splits up the load to develop a full transceiver system into several sub-projects: line/laser driver, transimpedance amplifier and a serialiser/deserialiser unit. Additional sub-components are being developed to support easy connection with the chipset and provide straight forward system interfaces for DAQ and DCS. An outline depicting the basic sub-projects is given in Figure 7.6.

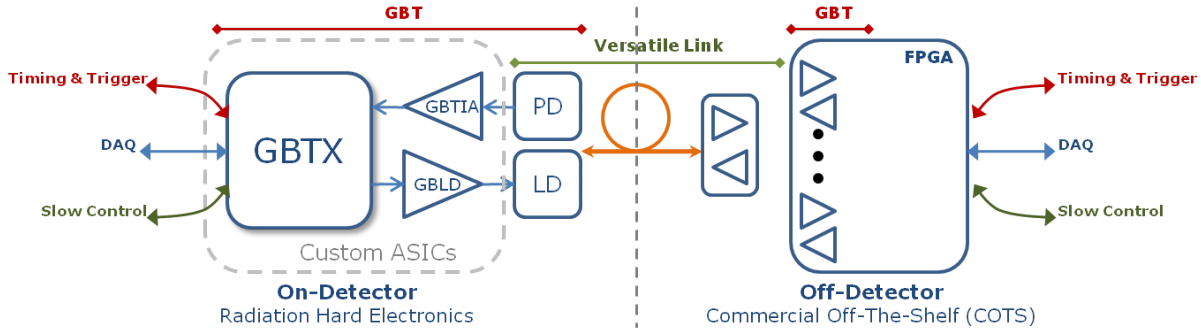

Figure 7.6: *GBT block diagram given in [M+10]. Shown is not only the GBT project, but also the placement of the Versatile Link project within the GBT frame.*

[e]Single Event Upset

7.2.1 GBTx

The GBTx [M+10] is the core design of the GBT project, a 4.8 Gbit/s SerDes[f] chip including internal clock generation, forward error correction en- and decoding, and multiple configurable I/O ports. The high speed link protocol uses 120 bit blocks including 32 bit error correction code and 4 control bits. The link efficiency is thus 70% providing a user bandwidth of 3.36 Gbit/s. Four out of 84 user bits are reserved for a slow control path. Opposed to the IBL layout using 8B10B encoded transmission, the error correcting code allows to detect and correct up to 16 consecutive bit errors within 120 transmitted bits.
The circuit supports local clock generation including phase adjustment of each clock output. For now, the number of clock outputs is limited to three, but has to eventually be increased to one per detector unit connected. The previous transmission systems allowed transfer of the LHC clock per link, which needs to be overcome within the multi-Gbit/s frame.

E-Ports

To exchange data with the GBT, multiple configurable "e-ports" are available. Each of them can run at 80, 160 and 320 Mbit/s using a DC-balanced protocol. Frontend implementations of the e-port counterpart are available from the GBT project to ease the implementation of a GBT based transmission system.
An initial limitation of the GBT project is the escape bandwidth of the chip, given by these e-ports. It keeps connections to individual frontends at a rather low rate and will only deliver a minor upgrade compared to the current ATLAS inner detector systems. Availability of multiple transmission lines from inner detector modules to the GBT is questionable due to material budgets within the ID.
Load sharing would offer a way of spreading the total bandwidth between multiple high speed connections. The e-ports could all run at 320 Mbit/s, but instead of feeding straight through into the GBT transmit section, they must be buffered in the GBT input. Packaging of data and a small input buffer would be needed within the GBT, but would allow to run with a maximum readout rate for a larger area of detector than is typically occupied.

Slow Control

Despite the high speed transfer, the GBT chipset supports monitoring and control via a dedicated Slow Control Adapter [G+09] (GBT-SCA). The project aims to reduce effort of the experiments in this field, too. A detailed risk analysis will be needed in this field, to apply the slow control aspect. The past has shown severe problems with the optical link setup, hence the control of detector components should use a sideway instead of the optical readout connection.

[f] Serializer - Deserializer

Assuming a Pixel B-Layer, or IBL layout, loss of a single GBT would imply losses as large as half a parallel cooling circuit (In the B-Layer that would be 4.5%, in IBL 7%). Yet loss of control of that one stave could affect the counterpart on the second stave. Hence a slow control connection should as least be layed out redundantly here.

Physical redundancy

The GBT chipset offers no means of redundancy within the chipset itself. The optical path can be laid out redundantly, yet the driver and GBTx chipsets will be unique. A dual output can drive two lasers, in case the laser fails. As mentioned before, loss of a single GBT causes loss of major detector parts. Therefore it might be considered to lay out a redundant GBT path that, in case of loss of a GBT, can be activated.

Off-detector components

Present off-detector components start moving towards firmware based implementations of hard logic within FPGAs. This allows simple upgrades, both in terms of firmware and hardware. New functionality can be integrated by updating the firmware, whilst replacing one FPGA board with another, the same firmware can be used to guarantee a smooth transition towards an upgraded readout system.

With the GBT project, FPGA firmware source is delivered [B$^+$09]. It can be implemented into current standard FPGAs with high speed serial ports, which then allow to form the counting room end of a GBT protocol link. Therefore, fewer custom made components can be used. Also, the project can be set up purely within FPGAs, running preliminary tests in a table top setup.

7.2.2 Versatile Link Project

In common with the GigaBit Transceiver project, a versatile link project [A$^+$09a] was set up. This project aims to define the physical layer of a Gbit/s link, including a transfer medium, as well as the driver- and receiver-interface. Whilst possible candidates for radiation hard PiN and laser diodes are selected within this project, the driver and receiver chips are to be delivered by external projects, e.g. the GBT. Characterisation of those components will be done within the versatile link project, checking resistance against SEU in combination with e.g. a selected PiN diode.

Within the project, developments for packaging of lasers and PiN diodes are ongoing. A first version of a plastic package with good alignment has been shown in 2010. Yet,

the development aims for single links and will thus always deliver a relatively high space occupancy per link[g].

Thermal Effects

A subset of creating a new package is taking care of thermal issues within the package. As is given by the Arrhenius factor introduced in Section 4.3.3, the temperature has a large influence of the ageing of components. Hence, a laser needs to be packaged such that its core temperature can be kept low, with the lowest thermal resistance possible through the package. This gets even more important as an irradiated device looses output power and is thus to be driven at higher powers, generating more heat.

A way of measuring that thermal resistance of a package has been evaluated in the frame of the Versatile Link project: A lasers output spectrum shifts with its core temperature. This change can now either be raised by raising the external temperature or by running more power through the laser, which will shift the spectrum of the laser in both cases. When managing to raise the power of the laser and at the same time lowering the environmental temperature so that the lasers central wavelength stays the same, one can measure the thermal resistance as the change of temperature over change of power [F$^+$11]. Setups to perform these measurements are available, to qualify lasers and packages for an upgraded LHC environment.

7.2.3 Wireless transmission

Despite using wired and fibre based transmission systems, first systems for wireless transmissions are being evaluated. A low mass, high frequency transceiver can be built from a custom ASIC [M$^+$09]. The system uses antennas embedded into a circuit, to transmit data from one detector layer to the next. Transmission from the inside to the outside of the detector cannot be done, as the silicon layers absorb the microwaves. Instead layouts are proposed which either increase the number of transmitters per layer, using different transmission bands, or decode the data per layer and multiplex it into a new transmission stream to the next layer.

Both layouts have the advantage of transmitting through a tower, which, as discussed above, should deliver the same track rate. Yet, proper testing of the transceiver for radiation hardness is missing, as well as a field test to verify transmission performance in 2 T magnetic field, encapsulated in large silicon, carbon and metal structures.

[g]The recently presented package is a slightly smaller version of an SFP housing, occupying about 7.5 cm^2.

7.2.4 Market Comparison

High energy physics suffers from multiple performance decreasing factors, opposed to the open market: Material and space requirements limit the maximum utilised performance, as does radiation hardness.
Yet, computing recently faces a bottleneck at the chip escape bandwidth - board to board communication in general works well, but the bandwidth, that can be transferred from one integrated circuit to another is limiting the maximum computing performance of single chips. Hence, industry leaders start moving to integration of very small optical transceivers on a chip level. This development could offer new possibilities for high energy physics, reducing the transceiver size and material amount, as well as the power consumption per Gbit/s.
Figure 7.7(a) gives an idea of where recent technologies are going. Shown is an optical array, packaged for connection of a multimode fibre ribbon with 12 channels, either transmitting or receiving. The needed laser driver or transimpedance amplifier is included within the package, covering less than $1\,\mathrm{cm}^2$ board space. High density optical links like these would remove large packaging and could be connected to driving electronics like multiple GBTs.

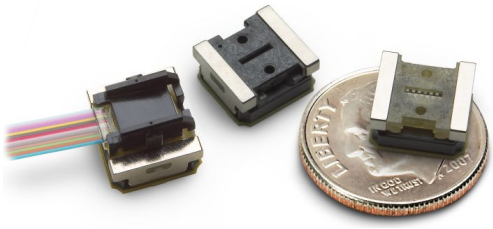

(a) Avago Technologies MicroPod™

(b) Array placement of MicroPod™s

Figure 7.7: *A recent product line from Avago Technologies [Ava10], giving transceivers with 4+4 channels, or 12 channel transmitters or receivers. The array of 4x4 modules shown gives a total bandwidths of 960 Gbit/s (Receiving and transmitting) within less than 20 cm^2.*

Active Optical Cables

In addition to higher density applications, another trend comes up in industrial applications, e.g. in computing centres, which is "ease of use": Fibre optic installations tend to be more sensitive to installation issues like plugging in cables correctly, as well as repetition

cycles. Also, one of the driving costs in cable and electro-optical component assembly is proper alignment. This is nowadays circumvented by encapsulating the electro-optical converters within the cable, an active optical cable. The user only has to obey restrictions concerning the bending radius, but is otherwise given a standard cable.

Whilst the same outline is given to the Versatile Link project, the use for building a radiation hard active optical cable is questionable: The cables themselves would become more expensive. Due to an increased component count, the failure rate is supposedly increasing, and thus the cables would give an even higher total cost. This could only be prevented in case the full cables including transmitters and receivers become a lot cheaper than individual cables, and transmitters and receivers.

7.3 Advanced Trigger Systems

As denoted in section 7.1.1, one HL-LHC requirement is better trigger quality. Recently, two systems evolved following this suggestion, which both support the recent trigger process and should deliver a higher-quality triggered content beyond the level 2 trigger. Additionally an idea of a possible trigger upgrade that reduces the total data rate used without decreasing the read out quality is given.

7.3.1 Level 1.5 Trigger System

Event selection within the Level 2 Trigger will be increasingly difficult within an sLHC environment. With increased pileup, track fitting and pattern recognition will take longer time to be processed, hence increasing L2 latency. The biggest issue here is the initial track seed generation, which tries to estimate the possible number of tracks from combining hits into first tracks. This number decreases as good quality track fits are done and therefore remove hits from the seed sample. Yet, the initial set gets increasingly larger with the occupancy driven into the level 2 trigger.

To increase track parameter quality and fitting speed, a hardware track fitter is proposed, the Fast TracKer [A$^+$09b] (FTK). The technical proposal is a structure comparable to the L1.5 Track Trigger system of CDF. It utilises large content addressable memories with previously stored track patterns to generate a track fitter seed within a very short timescale. The systems output is comparable with offline track fitting qualities. ROI processing is about 1000 times faster than when running within the level 2 software implementation. The FTK can run in parallel with level 2 processes and deliver a track sample for the final processing steps.

7.3.2 Level 0 Trigger System

Additional to the FTK system, an L1 preceding level 0 trigger system is being thought of. Opposed to the recent L1 system, the L0 would run with a 10 times higher rate, but only read out regions of interest from the inner detector frontends. With a supporting fast tracker system, events can then be selected from their track content. This should provide "better" event data to the offline based systems for analysis of (e.g.) b-tags. An L1 trigger could thus depend on tracking information from the ROIs.
To use this trigger system, inner detectors need to support a two staged readout and long-time ($> 10\,\mu$s) buffering of data. Data is to be stored in those buffers and sent out on both, L0 or L1 triggers. Due to L1 arrival depending on L0 processing, the total latency is four times the transfer length from ATLAS into the counting rooms, plus the processing of L0 and L1 trigger systems. With more recent technology solutions being available, an increase of all FE buffers by a factor of at least 4, if not more, can easily be thought of.

Poblems with ROI Based Readout

Moving towards the inner detectors as described above, a module at $\eta = 0$ covers a much wider range of eta than a module at $\eta = 2.9$. Due to interaction induced tracks being longer at larger η this is not very emphasised in the total amount of data that needs to be transmitted per full event[h]. Yet this feature is problematic for ROI based readout, as the innermost module covers the largest range of ROIs. It will thus be triggered much more frequently than a module at large η. Therefore, an ROI based readout for the innermost detector layers will need a very different readout structure, including a per-module link that allows one to read out a module at full L0 rate, at least up to PP0 where it can be interleaved with other modules[i].

7.3.3 Level 1 Track Trigger

A somewhat different approach to tracking within trigger processes is the proposal of a track trigger in form of a specialised detector layer. At high tracker radii, a double-layer microstrip or pixelised detector could locally pinpoint tracks coming from the coordinate system origin or close by. This would allow the number of high p_T tracks to be counted instead of just the E_T measured within the calorimeter, and give the L1 Trigger system a better handle on data selection.
This trigger detector approach is the only one that might allow the L1 trigger rate to be lowered using tracking information. Whilst systems like the L0 trigger can only be utilised

[h] Higher $|\eta|$-ranges have more data per track, lower ones have less data per track, but more tracks.

[i] Here a packaged transfer would be ideal, supporting load induced by the L0 trigger being shared amongst the modules that get triggered.

knowing the next hardware capabilities, the L1 track trigger detector will support a simple extrapolation, using the data volume needed. Here, the major hindrance will be to estimate the trigger quality, in particular its dependency of the detector alignment[j].

7.4 Advanced Readout Systems

Current off-detector readout systems within ATLAS are mostly based on VME components, as the standard allows simple control of components, is reliable and has continuos support. As data rates recently increase, new standards are formed within the telecommunication community.

7.4.1 Serial Protocol Systems

The nature of low speed transmission in VME systems is due to parallel routing of multiple signals within the VME backplane. Up to 64 data lanes need to run absolutely parallel with the system clock, for VME2eSST[k] transfers. This bottleneck can be circumvented in the way that is utilised in module data links - a serial transmission.
Recent system developments start moving towards serialised bus systems, which are found for example in ATCA[l]. An upgrade of VME, VME Switched Serial (VXS), supports various serial transmission standards, e.g. Ethernet or PCIe. As VXS inherits the form factor as well as connectors and supplies from VME, it is not perfectly adapted to recent technologies, yet it does not hinder recent implementations.
Whilst the choice of form factor is basically up to likings of the user, the bus system utilised within has some aspects to decide on:

- Maintainability/Future Safety: Throughout the past 20 years, a change in personal computer bus systems went from ISA via VESA to PCI and, after an intermediate hype for AGP, ended up with PCIe, which dominates the market right now. A bus system change this frequently is not encouraged within high energy physics. From start of development to the end of a runtime, a system has to stay maintainable. Opposed to common high speed bus systems that change quite frequently, Ethernet has proven a very stable protocol on the market and will be, as long as the lowest level protocol of the internet will be based on ethernet standards.

- Ease of Use: Using Ethernet as a bus system obliges the user to define serial hardware access protocols or implement higher level network communication like an

[j] Displacements of this detector might lead to misidentified track p_Ts or origins.
[k] VME double edge source synchronous transfer
[l] Advanced Telecommunications Computing Architecture

IP-stack. Though being generally available, implementing these layers both within hardware and software poses an initial effort that is much lower in protocols like PCIe. The protocol itself defines how registers are accessed and, within a PC framework controlling PCIe hardware, the hardware itself can be accessed via memory maps, as if it was part of the PC's memory. The user doesn't know about the background processes doing the mapping of an address access to the actual hardware. Usage of memory mapped protocols versus communication layer protocols will lead most users to falling for the memory mapped access, which should not be overlooked[m].

7.4.2 System Monitoring

Recent systems allow standardised IPMI[n] interaction with a crate. Therefore, attachement and detachment of hardware can be monitored, malfunctioning can be detected via temperature and power consumption monitoring, and firmware can be upgraded in a native way (opposed to implementing a specialised component as was done within the ATLAS Pixel Detector RODs). Maintainability increases and software can control configurations which previously had to be hardcoded. E.g. inserting a readout driver into a particular crate or using it with a particular back of crate card version could automatically change its configuration and firmware.

7.5 Conclusion

Recent developments like for the previously shown IBL, as well as the future GBT, will provide systems with increased stability and offer a much lower operational effort than was given in the ATLAS Pixel Detector optical readout link.
Future data has to be simulated and trigger scenarios need evaluation. Depending on that, systems like the GBT might well be able to handle the data load given for a stave-like structure of a future pixelised detector, or just a module. Promising is the rise of error correction code within transmission systems, most likely reducing the need for expert interventions. The same is true for self-calibration of transmission systems.
A subject yet to be looked at is redundancy within those high-speed transmission systems. Loss of a single transmission line might reduce a Pixel Detector layer by as much as a full stave, whereas in the recent system, as single module would be lost, or even just reduced in its readout speed.
Another major aspect is link density, which is particularly worrisome within tracking detector environments. Industry is moving towards highly integrated electronics, and albeit

[m] Of course, implementations of a memory mapped hardware access via Ethernet will remove this argument, but for now the market is lacking a good implementation.

[n] Intelligent Platform Management Interface

without radiation hardness, the level of integration should be a design target within high energy physics, too.

8 Concluding Remarks

The work done during this thesis contributed to the installation, commissioning and operation of the ATLAS Pixel Detector optical readout link. As was reported in Chapter 5, the ATLAS Pixel optical readout link as such performs remarkably well. As a result of this work, a tuning algorithm was developed and refined, leaving the optical readout link in stable operation for months. Tuning problems with individual links exist due to features of the used hardware that could only be circumvented by changing components and thus have to be dealt with manually.

In addition, an algorithm to tune the on-detector clock conditions was successfully set up. Both, the data transmission tuning, as well as the timing related tuning, are now the standard calibrations of the Pixel Detector optolink, and can easily be executed by trained calibration shifters.

A major issue with the TTC transmitters, in particular with the lasers in use there, have only been found while operating the link. These issues have, to date, only been mitigated by replacing faulty TX-plugins with new ones. The Pixel On-Detector electro-optical converters utilise the very same lasers. Part of this work describes the extrapolation done to estimate the on-detector impact of this issue. Though unlikely, as argued in Section 5.3.4, the on-detector lasers might need replacement during the next long shutdown. The new Service Quarter Panel (nSQP) project was therefore set up to re-build the on-detector services, changing the design to move the electro-optical converters out of the Pixel Detector volume for future serviceability.

A new detector layer, the Insertable B-Layer, is well under way. In this thesis, an off-detector optical interface was presented, that will cope with higher data loads coming from the IBL. A first prototype will soon be available for testing, and firmware is prepared and tested for operation with that prototype.

Eventually an outline of future readout systems was given, emphasising possible problems and solutions for higher rate readout of data. Whilst outer tracker readout is well under preparation, an inner tracker solution with radius smaller than 20 cm is yet to be found.

9 Acknowledgements

I am writing this for I think it should be said: Writing a PhD thesis is more than one person can possibly bare with by themselves and I will be as blunt as to say, the ones who've been through it alone have missed the best thing about it. I therefore have to thank quite a few people for the best years I have had so far:

Peter Mättig, for being my supervisor during my PhD, also during my diploma, my mentor since I started off at university in 2001 and letting me work with a great team for finalising the installation of the greatest machine ever seen.

Tobias Flick, for being the post-doc behind a not-quite-simple-to-protect-from-himself PhD student. He had the great pleasure of sorting out my miscommunications and cleaning up behind me after I left CERN end of 2008.

Claus Gößling, for reading my thesis and asking me questions that I can actually answer.

Timon Heim and Manuel Neumann, for being my good s....tudents. Manuel needs to work on his coffee addiction though, which I didn't manage to control but rather made use of it...

More students around me, Phillipp, Astrid, Julia, Kathrin, Jenny, Thorsten, as well as the "grown ups", Susanne, Klaus, Charlie, Monsieur Henß, Markus, Torsten and Jojo.

Sebastian Weber, for getting us through lab courses, helping my students with highest level programming languages and sharing an affection for hardware that is hard to find, even in nerdy people

David J. Griffiths, for writing the most awesome book on particle physics. Though me, I do not understand quite everything about particle physics theory, his "Introduction to Elementary Particles" essentially is what made me pass my exams and I have never read anything alike again.

I have to thank Apple for producing my wonderful Mac, which made my computerised life easier and much more productive. Also tools like InkScape and Plot created part of the more beautiful Figures used within this thesis. The uglier ones are ROOT output, but got much fancier due to a colour palette given to me by Goetz Gaycken - who is also my personal god of C++ programming!

Claudia Gemme, for some wonderful words, dinner, quotes like “It works, Champagne!”, “Those young people can never be happy” and “This is *not* funny”, and finally shift-planning in December 2008...
My closest friends: Graziella; Nils, Jenny and Finja; Julia and Dirk-Daniel; Meike; Matthias. They’re my backup and get me back on solid ground.
Special thanks here goes to Thomas Scheler, for studying physics too, therefore not leaving me alone in my topic. I had awesome times visiting him in Grenoble and Edinburgh and sure hope he’ll move to some other sweet locations in the world, for me to be able to visit him there. I’ll always envy him for his lasers.
My Family, this being (grand-)parents and my sister in law, for pushing me further, always being there for me and showing interest in everything I ever did - Besides other things and albeit sinful, making you proud just a little, is driving me.

Fridolin Dittus and Daniel Dobos, for Motivation (+2).
Folks at CERN that “lived” with me, besides many others, Anna, Simone and Kristof!
Corallie Roullier, for dancing with me in the snow

Lucia Masetti, for being my mentor in my bad days, being upset with me when leaving at 2 am without everything being finished, and a very good evening before the LHC crashed in 2008, introducing me to my favourite Whiskey - Lagavulin.
Tatjana, for keeping me focussed and not hating me over what she knows about me.

Sarah Reckefuss, for the off-times, driving me home, a glass of water after a little running and being there when needed the most.

Lea, for sharing a life with me over quite a distance.

My twin and twins, Sven, Jan and last but not least Stefan

A Recent Developments

A.1 Machine Schedule and Effects

Recently, the LHC schedule has shifted, running on up to a 2013 long shutdown (≈1.5 years). Things affected by these are both, the project started to replace the Pixel On-Detector optical readout services, as well as the IBL, which is now to be inserted in 2013. The latter implies that the replacement services have to still be ready in 2012, as the new Service Quarter Panel project and the Insertable B-Layer share a large amount of manpower.
Also, this shift in schedule affects the schedule for the 2016 shutdown, which is now more probably to happen in 2017 or even 2018. Luminosities have recently (May 2011) increased to about $1.2 \cdot 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, using 1042 colliding bunches. At the same time, the Pixel Detector occupancy has increased to about $1.5 \cdot 10^{-4}$. Eventually Pixel will thus probably suffer a little lower occupancy than was simulated. Yet, the Pixel Detector double column readout rate has not been increased to 40 MHz, which means that inefficiencies will probably hit Pixel earlier than expected from simulation. 40 MHz double column readout have shown side effects in calibration and were thus disregarded for readout in data taking.

A.2 Pixel TX deaths and recent results

The Pixel Detector TX plugins, as well as the SCT plugins, are continuously dying and being replaced. Death rates have decreased recently, but that is an artefact of the replacement of plugins. A task force was established investigating on VCSEL lifetime dependencies and possible migitation. Good evidence is given for a dependency on humidity in the environment. This seems to be a design feature of the Truelight Lasers in use, in combination with an enclosing epoxy on the laser surface which is to protect the wirebonds. Plugins with different Lasers are being produced since late 2010 and have been installed for the Pixel Detector B-Layer in spring 2011.
As far as the on-detector Lasers are concerned, the evidence of humidity playing a major role is a relaxing factor. Also, channels connected to disabled modules have been found to transmit light, sometimes with 100% duty cycle. A calculation of their on-time statistically excludes a lifetime compatible with the off-detector Lasers and seems a lot higher.
A document [TAPC11] describing failures within the Pixel Detector and narrowing down their sources is under preparation by the Pixel Collaboration.

A.3 The IBL Readout System

In the context of the IBL readout, initial thoughts assumed the same readout driver as is present in the ATLAS Pixel Detector readout system. As this would have been a bottleneck in terms of processing capability, thinking was to replace the BOC with a card, more functional than the previous one, allowing to route the full data path inside the BOC and only using the ROD for calibration control and slow processing.
This culminated in the idea of a standalone BOC card, which would nowadays be outfitted for PCIe, allowing it to be plugged into a standard PC, shown in Figure A.1. The similarities of PCI with VME, concerning the access via memory maps, would allow to use them same type of software to run on it (writing to registers instead of communicating with a smart device). The currently available technology in terms of FPGAs would allow to not use any extra devices on the board, other than external memory. Connection to an external TTC system would have to happen via an extra card and PC internal connections, as space on the faceplate is very limited.
Processing like S-Curve fitting could be done inside a graphics card here, allowing to reduce load on the PC and the programming effort for the embedded environment of the FPGA. First testbenches for this have been run in 2010 [D+10] and show promising results

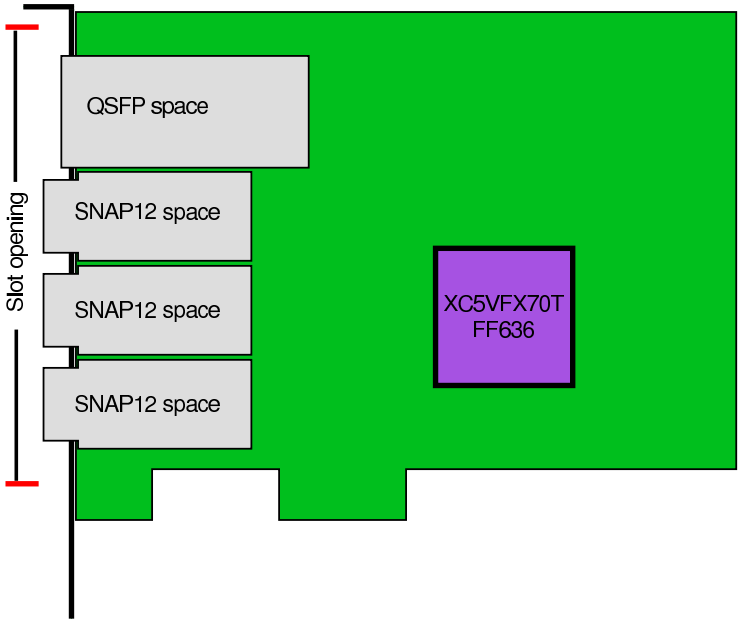

Figure A.1: *A PCIe schema of a Back of Crate card*

B Technical Issues and Solutions

This appendix is to give an outline of things that were introduced, albeit they are not necessarily about the optical communication itself.

B.1 TOT calibration

In the process of understanding the resolution of the ATLAS Pixel Detector in terms of charge sharing, time was invested into understanding the low charge Time Over Threshold (TOT).
The TOT resolution for each pixel can be adapted using a local feedback current trim DAC. It allows for small changes to the global feedback current with a range that can be selected by a global switch. To verify that tuning, a TOT calibration scan can be run, which injects multiple (typically 6) high charges between 10 and 200ke. For each of those injected charges a mean and a sigma are calculated. Those are then fitted versus injected charge, using:

$$TOT = A * \frac{Q - B}{Q - C} \tag{B.1}$$

The Pixel Readout Drivers deliver the mean values to be fitted, caculating them as the mean value of every hit they record from the detector. As soon as the TOT calibration scan moves to lower charges, it should (at Threshold) dip down to 0 TOT, as half the charges was not recognized, c.f. Figure B.1. Instead, by calculating the mean of all hits observed, that very value is artificially increased. As the fit function has a preference to behave like a straight line, it then extrapolates from the TOT(Q) curve at upper values towards zero charge. The offline software has therefore been able to read a TOT of 2 as 2500 electrons, with thresholds being set to 3500e...
A different approach to caluclation of mean and sigma, taking into account the number of hits that was observed, plus a different fit function (which is derived from the voltage triangle at the preamplifier output) serve to describe the behaviour close to threshold much better and can probably enhance resolution...

$$TOT(Q) = (Q - Q_{thr}) \cdot \left(\frac{t_{top}}{Q} + \frac{1}{I_{fb}} \right) \tag{B.2}$$

Within this equation, physical sizes are used, that can be estimated from previous calibration or data taking with multiple bunch crossing readout and histogramming:

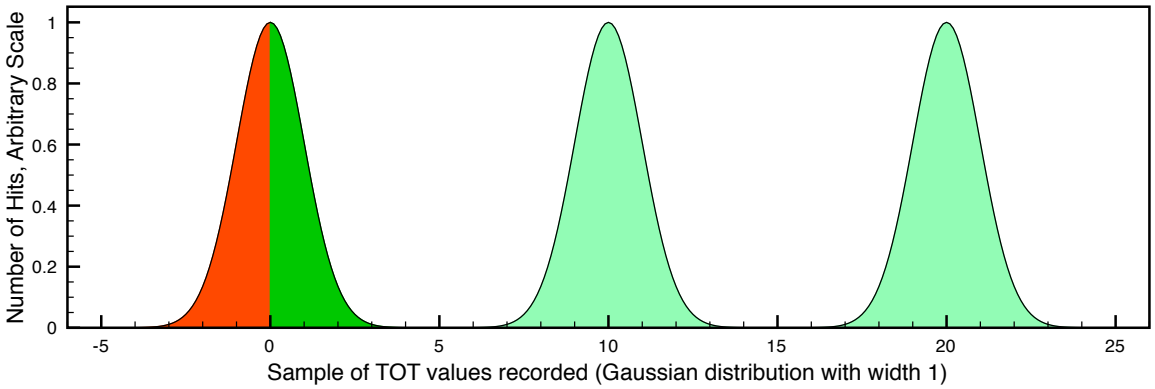

Figure B.1: *TOT values and mean calculation range (green). For different charges injected, a response is plotted. Values below zero (red) are not recorded by the Pixel Detector readout system*

Q_{thr} is the threshold charge (compareable to the threshold scan result)

t_{top} is the time it takes, in BCs, to get from 0 voltage to peak before the preamplifier

I_{fb} is the feedback current in electrons per BC

An example of the Fit results compared with the Default Pixel Detector TOT calibration fit are shown in Figure B.2. Though not very easy to track down, there is a visible difference in the lower end of the two fits: Whilst the default fit tends to pull towards (0, 0), the fits proposed here decreases when reaching the threshold value.

B.2 Crosstalk Scans

During the commisioning phase I delivered the first functional implementation of a Crosstalk scanmask for the new DSP code. The mask allows to read out one Pixel, while injection happens in the neighbouring two. This scan has proven substantial for detecting disconnected bump bonds between the Pixel sensor and the frontend. In the process of setting up this scan, the maskstage[a] configuration was speeded up by two orders of magnitude, using functions for configuration shifting that already existed.

[a]As not all Pixels can be read out at the same time, the Pixel calibration implements a mask of Pixels that are tested at the same time. The steps performed within this mask are referred to as maskstages.

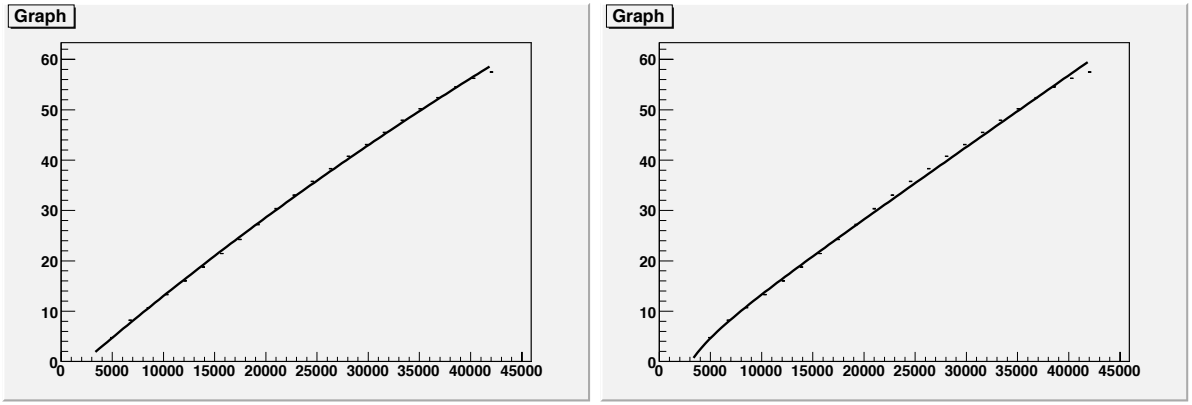

(a) Default TOT fit

(b) TOT fit with physical quantities

Figure B.2: *Two different TOT fit approaches. Though hard to see, the main differences are at the tails of the fit function. Whilst the default fit function tends to better match the upper end of the data, the suggested fit function gives a better modelling of the threshold behaviour and should thus help increasing resolution. Data used for this fit was generated using the standard histogramming and TOT_MEAN calculation.*

B.3 Slow Turn-On Scans

Slow Turn-On is still hard to quantify, as no scan exists that records the changes in the signal amplitude over time. If a scan was to be set up for that, it would have to be based on the slow BocScan mechanism. Recently, trials were made to properly synchronise the pseudorandom pattern transmission used within slow BocScans, with the off-detector pattern recording. Instead of identifying to first incoming "1", the scan is to assume a fixed latency and just histograms the recorded input buffers bit-by-bit. An example is shown in Figure B.3.
As soon as synchronisation works properly, which might be a matter of checking the ROD's FPGA firmware, the scan should allow to not only see bit-wise slow turnon, but also record that data for different data delays. This would then allow to perform an analogue recording of the returned signal with 1 ns resolution, a kind of sampling-oscilloscope shot. It already allows to observe the slow time-constants of the receiving component, but proper quantification can not be done yet.

B.4 VME readout bandwidth

The recent ATLAS Pixel Detector system is often referred to as being very slow. This is a feature of the VME bus, but can be reduced drastically, using software available now, as well as some structuring of the access to the Bus. The best example here is a BOC

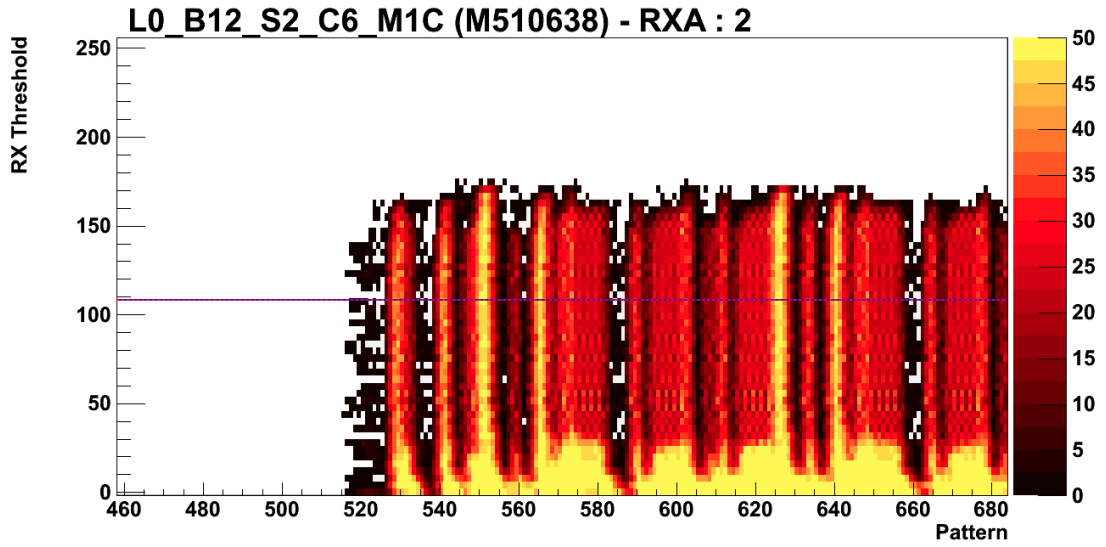

Figure B.3: *A recorded bit pattern of 50 repeated scans with the same pseudorandom sequence for different thresholds. The Scan still has problems with synchronisation, but has gotten a lot better than used to be the case. Growth of the yellow regions at the bottom of the scan shows a slow time-constant, either in the DRX-12 preamplifier or the PiN diode.*

fast scan that was derived from a simple tool checking one single BOC. The tool used to trigger recording of all possible data and then check the status of counting within the formatters, until these were done. Afterwards, counter values were read out and the next setting was scanned. Transferring this into a crate level software, 16 adjacent devices are being read continuously for their counting state, loading the bus with a very lot of transfers containing zero information - the number of counts requested defines the length the counters are running. Adding some microseconds, one can easily estimate the full counting time and read out afterwards without even checking for whether the counters have finished. Scan speedups of more than 60% have been observed using a simple *usleep* instruction within the threaded code.

The VME bandwidth is ≈20-25 MByte/s in 32-bit data transmission mode. This is initially reduced by the ROD who cannot handle requests any faster than with 11 MByte/s. As the communication is typically not done by the program reset manager FPGA, but by the Master DSP, a handover has to be done, which introduces an additional decrease in speed. The MDSP can be read out with a maximum rate of about 7 MB/s. This rate is decreased down to less than 4 MB/s, by the access that is typically done in form of single word write accesses. The VME controller allows to write multiple words as a block at the same time.

The Master DSP as the target of the transfer also allows for a direct address increase and can hence accept such a VME transfer[b].

B.5 Optolink Tuning Problems

Users of the ATLAS Pixel Detector optical readout mostly consider it hard to tune, that is, in about 1% or less of the modules, the automated tuning procedure fails to set a functional operating point. Changing the tuning procedure probably leads to higher losses in the tuned 99% than can be gained from the 1% of non-functional modules.
A difference could probably be made with a different type of receiver circuit. As mentioned previously, c.f. Chapter 5.3.4, the DRX-12 features a relatively slow response time, most probably due to its RC component. At the same time the the PiN diode in use features a slow time constant, which leads to overdriven inputs at the DRX. Both features can be overcome by building the DRX-12 as a transimpedance amplifier with differentiating and integrating stage following. Coupling the TIA into a next stage via a capacitance already delivers the differentiating stage. The second stage needs a defined input level, which can then be influenced due to the capacitive coupling. An upper and a lower threshold can switch a following fast Flip-Flop - with high signals incoming into 1, with low signals into 0. This thresholded binary integrating stage could probably serve as a simple way of transferring transients on the input into a binary output stream.

[b]The automatic address increase given by the VME bus is filtered by the PRM within this transfer mode - it basically refers certain address ranges to fixed addresses on the MDSP.

List of Tables

List of Figures

Bibliography

[A+83a] G. Arnison et al. Experimental observation of lepton pairs of invariant mass around 95 gev/c2 at the cern sps collider. *Physics Letters B*, 126(5):398 – 410, 1983.

[A+83b] G. Arnison et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at $\sqrt{s} = 540GeV$. *Physics Letters B*, 122(1):103 – 116, 1983.

[A+95] S. Abachi et al. Observation of the top Quark. *Phys. Rev. Lett.*, 74:2632ff., 1995.

[A+01] M. L. Andrieux et al. Single-event upset studies of a high-speed digital optical data link. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 456(3):342 – 351, 2001.

[A+05] K. E. Arms et al. ATLAS pixel opto-electronics. *Nucl. Inst. Meth. A*, 554(1-3):458–468, 2005.

[A+08a] G Aad et al. The ATLAS Experiment at the CERN Large Hadron Collider. *J. Instrum.*, 3:S08003. 437 p, 2008.

[A+08b] G Aad et al. ATLAS Pixel Detector electronics and sensors. *J. Instrum.*, 3:P07007, 2008.

[A+08c] David Arutinov et al. Digital architecture and interface of the new atlas pixel front-end ic for upgraded lhc luminosity. In *Nuclear Science Symposium Conference Record, 2008. NSS '08. IEEE*, pages 1923 –1928, October 2008.

[A+09a] L Amaral et al. The versatile link, a common project for super-LHC. *Journal of Instrumentation*, 4(12):P12003, 2009.

[A+09b] A Annovi et al. Development of FTK architecture: a fast hardware track trigger for the ATLAS detector. Technical Report arXiv:0910.1126, Oct 2009.

[A+09c] Ralph Assmann et al. Beam dynamics aspects of crab cavities in the CERN Large Hadron Collider. *Phys. Rev. ST Accel. Beams*, 12(10):101002, Oct 2009.

[A+10] Georges Aad et al. Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC. 2010.

[A+11] Georges Aad et al. Charged-particle multiplicities in pp interactions measured with the ATLAS detector at the LHC. *New J. Phys.*, 13:053033, 2011.

[AC98] The ATLAS Collaboration. Atlas trigger performance: Status report. Technical Report CERN-LHCC-98-015, CERN, Geneva, Jun 1998.

[AHP] ATLAS Homepage. http://www.atlas.ch.

[atl11] The ATLAS online monitoring homepage. 2011.

[Ava10] Avago Technologies. Avago Fiber Optics: Breaking Bandwidth and Performance Barriers in Supercomputing, October 1st 2010.

[AW01] Anthony Weidberg. ATL-IS-ES-0045. Engineering Specification of the DRX-12 circuitry on EDMS, September 2001.

[B+83a] P. Bagnaia et al. Evidence for $Z^0 \rightarrow$ e+e- at the CERN $\bar{p}p$ collider. *Physics Letters B*, 129(1-2):130 – 140, 1983.

[B+83b] M. Banner et al. Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN $\bar{p}p$ collider. *Physics Letters B*, 122(5-6):476 – 485, 1983.

[B+02] R. Beccherle et al. MCC: the Module Controller Chip for the ATLAS Pixel Detector. *Nuclear Instruments and Methods A*, 492(1-2):117 – 133, 2002.

[B+08] H. M. Braun et al. ATLAS liquid argon calorimeter front end electronics. *Journal of Instrumentation*, 3(09):P09003, 2008.

[B+09] S. Baron et al. Implementing the GBT data transmission protocol in FPGAs. 2009.

[B+11] G Bruni et al. ATLAS IBL: integration of new HW/SW readout features for the additional layer of Pixel Detector. *Journal of Instrumentation*, 6(01):C01018, 2011.

[C+04] M. Capeans et al. The transition radiation tracker of the ATLAS experiment. *Nuclear Science, IEEE Transactions on*, 51(3):994 – 1000, june 2004.

[C+08] V Cindro et al. The ATLAS Beam Conditions Monitor. *Journal of Instrumentation*, 3(02):P02004, 2008.

[C+10] M Capeans et al. ATLAS Insertable B-Layer Technical Design Report. Technical Report CERN-LHCC-2010-013 ATLAS-TDR-019, CERN, Geneva, Sep 2010.

[CCFT64] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π Decay of the $K2$ Meson. *Phys. Rev. Lett.*, 13(4):138–140, Jul 1964.

[Chr94] Christopher M. Miller. High-Speed Digital Transmitter Characterization Using Eye Diagram Analysis. 1266 Hewlett-Packard Journal 45, August 1994. pp. 29-37.

[CM09] A. Clark and G. Mornacchi. ATLAS B-Layer Task Force Final Report. Internal Report ATU-SYS-MR-0001, 2009.

[D+10] J Dopke et al. Study of FPGA and GPU based pixel calibration for ATLAS IBL. May 2010.

[D+11] J Dopke et al. The IBL Readout System. *Journal of Instrumentation*, 6(01):C01006, 2011.

[DD07] Daniel Dobos and Jens Dopke. ATLAS Pixel Detector Simulator for the ROD formatter FPGAs, October 2nd 2007.

[DF08] Jens Dopke and Tobias Flick. Optolink Status, November 24th 2008. Presentation at the ATLAS Inner Detector Week.

[DFH+10] J Dopke, T Flick, S-C Hsu, R Madaras, U Mallik, S Mehlhase, D Miller, J Moss, and S-J Park. Commissioning of the atlas pixel optical readout link. Technical Report ATL-COM-INDET-2010-047, CERN, Geneva, Mar 2010.

[Die07] Leonard Dieguez. Clock and Data Recovery with Coded Data Streams. Xilinx Application Note 250, v1.3.2, May 2nd 2007.

[Dob07] Daniel Adam Dobos. *Commissioning Perspectives for the ATLAS Pixel Detector. oai:cds.cern.ch:1092108.* PhD thesis, Univ. Dortmund, Dortmund, 2007. Presented on 1 Sep 2007.

[Dop07] Jens Dopke. Studies for automated tuning of the quality of the Back-of-Crate card for the ATLAS Pixel detector. Diplomarbeit, Universität Wuppertal, WU D 07-02, March 2007.

[Dop10] Jens Dopke. Dead Channel Estimate. Presentation during the Pixel Taskforce, August 3^{rd}, 2010.

[DP98] John D Dowell and M Pearce. ATLAS Front-end Read-out Link Requirements. Technical Report ATL-ELEC-98-001, CERN, Geneva, Jul 1998.

[F+11] T. Flick et al. Measurement of the Thermal Resistance of VCSEL Devices. Jan 2011.

[Fey49] R. P. Feynman. Space-time approach to quantum electrodynamics. *Phys. Rev.*, 76(6):769–789, Sep 1949.

[Fli06] T. Flick. Studies on the Optical Readout for the ATLAS Pixel Detector - Systematical Studies on the Functions of the Back of Crate Card and the Timing of the Pixel Detector. Dissertation, Universität Wuppertal, WUB-DIS 2006-05, July 2006.

[G+09] A Gabrielli et al. The GBT-SCA, a radiation tolerant ASIC for detector control applications in SLHC experiments. 2009.

[GGP10] B. Di Girolamo, D. Giugni, and H. Pernegger. ATLAS Pixel Taskforce 2010. Technical Report ATL-IP-MR-0002, CERN, Geneva, Sep 2010.

[GPW04] David J. Gross, H. David Politzer, and Frank Wilczek. The Discovery of Asymptotic Freedom and the Emergence of QCD. The Nobel Prize in Physics 2004, December 8th 2004.

[Gre01] Ingrid-Maria Gregor. *Optical Links for the ATLAS Pixel Detector*. PhD thesis, University of Wuppertal, Wuppertal, 2001. WUB-DIS 01-03.

[Gri08] D. J. Griffiths. *Introduction to Elementary Particles*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, 2nd edition, 2008.

[Gro00] CERN Microelectronics Group. PHOS4 - 4 Channel delay generation ASIC with 1 ns resolution, 2000.

[H+73] F.J. Hasert et al. Search for elastic muon-neutrino electron scattering. *Physics Letters B*, 46(1):121 – 124, 1973.

[Hei09] Timon Heim. Investigations on ATLAS Pixel TX Plugin Parameters. to be published, 2009.

[Hig64] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13(16):508–509, Oct 1964.

[Ibr10] I Ibragimov. Timing Behavior of the ATLAS Pixel Detector in Calibration, Cosmic-Ray and Collision Data. Jun 2010.

[JJF+06] John Joseph, Richard Jared, Douglas P. Ferguson, Andreas Korn, and Trevor Vickey. ATLAS Silicon ReadOut Driver (ROD) Users Manual, 2006.

[KKFN05] S. Kersten, P. Kind, T. Flick, and S. Kirichu Nderitu.
The Supply and Control System for the Opto Link of the ATLAS Pixel Detector. CERN ATL-IP-ES-0106, June 2005. Revision 1.6.

[LDR04] LHC Design Report Homepage, 2004.

[LHC09] LHC Project Underground Works. CERN Civil Engineering Group Homepage, June 26th 2009.

[M+09] C. Marcu et al. A 90 nm CMOS Low-Power 60 GHz Transceiver With Integrated Baseband Circuitry. *Solid-State Circuits, IEEE Journal of*, 44(12):3434 –3447, dec. 2009.

[M+10] P. Moreira et al. The GBT-SerDes ASIC prototype. *Journal of Instrumentation*, 5(11):C11022, 2010.

[MMK07] P. Moreira, A. Marchioro, and K. Kloukinas. The GBT: A proposed architecure for multi-Gb/s data transmission in high energy physics. 2007.

[Mye10] Steve Myers. LHC Commissioning and First Operation. Report on the LHC, ICHEP 2010 Paris, July 26th 2010.

[N+10] (Particle Data Group) K Nakamura et al. Review of Particle Physics. *Journal of Physics G: Nuclear and Particle Physics*, 37(7A):075021, 2010.

[Nde07] S. Kirichu Nderitu. Atlas Pixel Optoboard Production and Optolink Simulation Studies. Dissertation, Universität Wuppertal, WUP-DIS 2007-03, 2007.

[Neu11] Manuel Neumann. Readout components for the Insertable B-Layer of the ATLAS Pixel detector based on Field Programmable Gate Arrays. Diploma thesis, Bergische Universität Wuppertal, March 2011.

[P T08] P Teixeira-Dias. Higgs boson searches at LEP. *Journal of Physics: Conference Series*, 110(4):042030, 2008.

[P+06] I. Peric et al. The FEI3 readout chip for the ATLAS pixel detector. *Nucl. Instrum. Meth.*, A565:178–187, 2006.

[Rou10] Esteban Roulet. Latest results from the Pierre Auger Observatory. *PoS*, HRMS2010:034, 2010.

[RvdBH03] Aurelio Ruiz, Erik van der Bij, and Stefan Haas. HOLA S-LINK Interface. Data Sheet v1.0 Website, 2003.

[Saw05] Nick Sawyer. Data Recovery. Xilinx Application Note 224, v2.5, July 11th 2005.

[SNA02] SNAP12 Multi-Source Agreement, Specification Appendix, May 15th 2002.

[TAC99] The ATLAS Collaboration. Technical Design Report, Volume 1. ATLAS TDR 14, CERN/LHCC/99-14, Mai 1999.

[TAPC11] The ATLAS Pixel Collaboration. Analysis of the Pixel Module and FE Failures. EDMS Document: ATL-IP-ON-0005, 2011.

[W+04] A. R. Weidberg et al. The off-detector opto-electronics for the optical links of the ATLAS Semiconductor Tracker and Pixel detector. *Nucl. Inst. Meth. A*, Volume 530, Issue 3:293–310, 11 September 2004.

[Wei67] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19(21):1264–1266, Nov 1967.

[Wei08] Anthony Weidberg. ESD Studies, September 3rd 2008.

[WF83] A. X. Widmer and P. A. Franaszek. A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code. *IBM Journal of Research and Development*, 27(5):440 –451, sept. 1983.

[Wik11] Wikipedia. Non-return-to-zero — Wikipedia, The Free Encyclopedia, 2011. [Online; accessed 26-June-2011].

[WW96] D. J. White and A. R. Weidberg. LDC Engineering Specification. unpublished, 1996.