Design and Evaluation of the IBL BOC for the ATLAS Experiment at CERN

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
der Universität Mannheim

vorgelegt von

Dipl.-Phys. Nicolai Schroer
aus Paderborn

Mannheim, Dezember 2012

Dekan: Prof. Dr. Heinz Jürgen Müller
Referent: Prof. Dr. Reinhard Männer
Korreferent: Prof. Dr. Wolfgang Effelsberg

Tag der mündlichen Prüfung: 14. Februar 2013

Abstract

In 2013 during a 20 month long shutdown of the LHC the Pixel Detector of the ATLAS Experiment at CERN will be upgraded by inserting a fourth innermost layer between the beam pipe and the current detector. This so called *Insertable B-Layer* (IBL) will be constructed with 448 of the new FE-I4 chips to handle the readout of the about 12 million pixels provided by the sensors of this layer. The improved architecture and increased bandwidth of these new readout chips requires new off-detector electronics which were decided to be also backwards compatible to the existing system. Hence the VME card pair establishing the optical interface to front-end and data acquisition (BOC) and managing the data processing and calibration (ROD) have been redesigned for the IBL.

In this thesis the redesign of the BOC card is motivated and presented. At first the ATLAS Experiment is described and the need to upgrade the Pixel Detector with a new layer is explained. As the readout chip architecture of the current system has flaws preventing its use for the IBL the new FE-I4 is introduced, and with a look at the current off-detector electronics the need for a redesign of it is justified. Starting with the conceptual planning, the redesign process of the BOC card is presented from hard- and firmware development to testing of the first prototypes. The redesigned BOC is based on modern FPGA technology in conjunction with commercial off-the-shelf optical transceiver modules to provide an integration four times higher than the current system, including the flexibility to adjust to different use cases by simply changing the firmware.

Zusammenfassung

Im Jahr 2013 wird während einer 20 monatigen Betriebspause des LHC der Pixel Detektor des ATLAS Experimentes am CERN durch das Einfügen einer vierten innersten Lage zwischen Strahlrohr und momentanen Detektor aufgerüstet. Dieses sogenannte *Insertable B-Layer* (IBL) wird mit 448 Stück des neuen FE-I4 Chips ausgestattet sein, die die Auslese der etwa 12 Millionen Pixel der Sensor Chips dieser Lage übernehmen. Die verbesserte Architektur und erhöhte Bandbreite dieser neuen Auslesechips erfordert neue Off-Detektor Elektronik. Da beschlossen wurde, dass diese rückwärts kompatibel zu dem bestehenden System sein soll wurde das VME Karten Paar das das optische Interface zu dem Detektor und dem Datenaufnahme System bildet (BOC) und die Daten verarbeitet (ROD) neu entwickelt.

In dieser Arbeit wird die Neugestaltung der BOC-Karte motiviert und detailliert vorgestellt. Zunächst wird das ATLAS Experiment beschrieben und die Notwendigkeit den Pixel Detektor mit einer neuen Lage zu ergänzen erläutert. Da die Architektur des aktuellen Auslesechips Mängel hat die die Verwendung für den IBL verhindern wird der neue FE-I4 Chip vorgestellt und mit einem Blick auf die momentane Ausleseelektronik die Notwendigkeit für eine Neugestaltung derselben gerechtfertigt. Beginnend mit der konzeptionellen Planung wird der Prozess der Hard- und Firmware Entwicklung bis zum Testen der ersten Prototypen vorgestellt. Die neu gestaltete BOC Karte basiert auf moderner FPGA Technologie in Verbindung mit kommerziell erhältlichen optischen Transceiver Modulen was eine viermal höhere Integration als das derzeitige System ermöglicht. Außerdem bietet es die Flexibilität an verschiedene Anwendungsfälle durch einfaches Ändern der Firmware anpassbar zu sein.

Contents

Introduction

Figure 1: Picture from the Earth as seen from outer space with the location of the CERN near Geneva, Switzerland marked with the X. Picture is taken from Google Earth.

So that I may perceive whatever holds - The world together in its inmost folds.[1], This famous quote from Goethe's *Faust* is often used in works about particle physics as it perfectly illustrates humans ambition to study and understand how our world works. With the recent announcement[1] of the discovery of a new particle

[1] Johann Wolfgang von Goethe (1808) - Faust I, in a translation by Georg Madison Priest.

an important milestone in this understanding has been reached. Particle physicists research the composition and structure of matter on a subatomic level. To describe all known particles and their interactions they use a quantum field theory called *Standard Model*[2]. A particle which gives mass to the other elementary particles (except photon and gluon) was predicted[3] in this model, but the experimental proof for the existence of this so called *Higgs Boson* is missing so far. However, the work of thousands of physicists from around the world seems to have payed off, as a good candidate has now been found.

The Higgs Boson can not be observed directly because its lifetime is expected to last only a tiny fraction of a second until it decays into sets of known elementary particles with distinct energy. These channels, as the sets are also called, are predicted by the Standard Model. By detecting them and measuring the energies of the decay products the mass of the origin particle can be determined. In order to prove the Higgs Bosons existence, a powerful particle accelerator is needed to produce collisions with enough energy to create such a particle. Also a precise detector system is required at the interaction point to detect the decay products of the particles created during the collisions. Huge amounts of data needs[1] to be analyzed to narrow down its potential mass and to have significant statistics.

Experimental data obtained by the ATLAS (e.g. Figure 2) and the CMS experiments from about a quadrillion (10^{15}) proton-proton collisions at the Large Hadron Collider (LHC) at CERN indicates a new particle in the mass region between 125 GeV and 126 GeV. A combination of data collected from several main decay channels shows a significance of 5 σ, meaning that only one in about 3 million measurements would yield the same result by random fluctuations (in case this particle would not exist). This is considered accurate enough to justify the declaration of the discovery of a new particle.

But the work is not over yet. More measurements need to be done to determine if the new particle is indeed the Standard Model Higgs Boson or just another particle from a currently unknown family. In any case, data taking will continue at the LHC in 2012 with even higher beam energy and intensity. This will yield multiple sets of the current event data by the end of the year before the system will be shut down until the end of 2014. During this long shutdown period upgrades will be installed to the detectors and the accelerator.

[1] *Data* is a plural term, but the use of singular verbs together with it is "widely accepted in standard English"[4] and the IEEE Computer Society states in their style guide to "follow author preference for use as singular or plural"[5], and my preference is the singular.

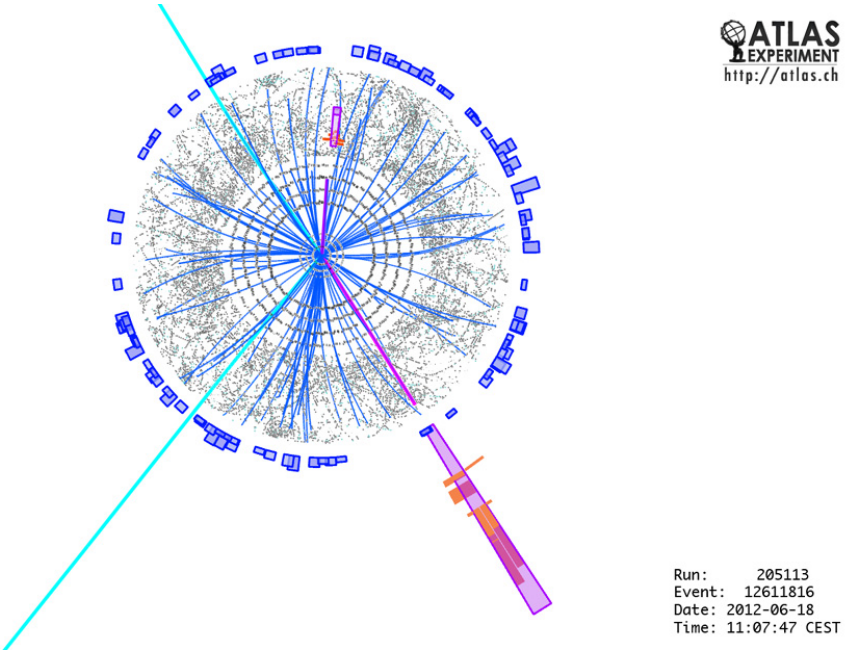

Figure 2: Event display of a H $\rightarrow$ 2e2mu candidate event recorded by ATLAS at CERN. The light blue tracks are from muons, the purple from electrons.[6]

This thesis will focus on one specific upgrade, the *Insertable B-Layer* (IBL) for the Pixel Detector of ATLAS, especially the redesign of the off-detector *Back Of Crate* card (BOC) which is needed for the data readout. An overview will be given first to describe the framework of this work and to motivate its necessity. Starting from a global scale the focus will zoom in on CERN and the LHC, with ATLAS as one of the most significant experiments. Then it continues inward through the other sub-detector systems to the Pixel Detector close to the collision point and its sensor chips.
After the IBL has been motivated and the new sensor chips necessary for it have been introduced, the off-detector electronics currently used for the readout of the Pixel Detector will be described. As its limitations prevent the use for the IBL, new electronics are needed. The planning, design, production and testing of the new BOC which is a part thereof, will be the main part of this thesis.

1

The Large Hadron Collider at CERN

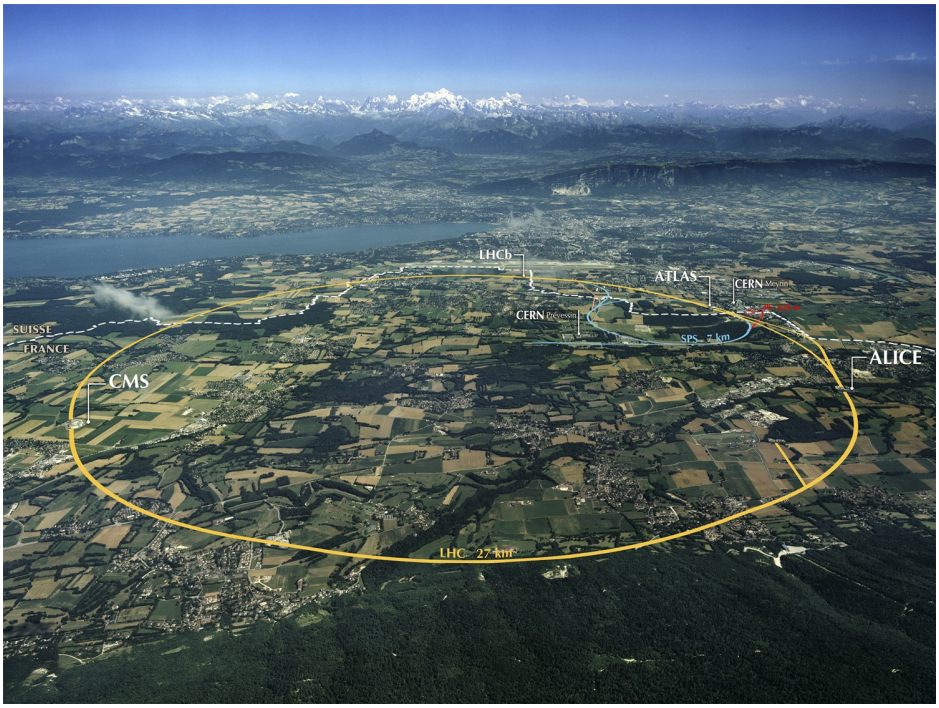

Figure 1.1: Aerial view of the region around Geneva where CERN and LHC are located [7].

Near Geneva the *Conseil Européen pour la Recherche Nucléaire* (CERN) was founded in 1952 as a provisional body to build a research organization for fundamental physics in Europe. When this organization came to life two years later in 1954 it kept the name CERN, even though it is officially called *European Organization for*

Nuclear Research now. Thousands of scientists from more than 600 universities and institutes from the 20 European member states and from around the world work at or for the CERN. Main focus is the fundamental research of the constituents of the nucleus and thus to increase the understanding of the Universe.
To study the structure of matter, particle accelerators are needed to boost the energy of particle beams before they are made to collide to possibly create new particles. As the lifetime of these new particles is usually extremely short, they are observed by detecting their decay products. Hence several particle accelerators and detectors were built at the CERN starting in the late 1950s (see Figure 1.2 for a map of todays accelerators at CERN). One of the larger ones is the *Proton Synchroton* (PS) which started its operation in 1959 and is still in use today to pre-accelerate protons and ions to about 25 GeV. In 1976 the *Super Proton Synchroton* (SPS) was commissioned which is, with its almost 7 km in circumference, still the second largest accelerator at CERN. With its help the Nobel Prize winning discovery of the W and Z bosons was made in 1983[1]. Nowadays the SPS is used to further accelerate the 25 GeV particles from the PS to an energy of 450 GeV before they are injected into the *Large Hadron Collider* (LHC, see section 1.1 for more details). In 1989 the *Large Electron-Positron* (LEP) collider was built underground in a circular tunnel with a circumference of 27 km. This was one of the largest and most powerful accelerators for electrons and positrons. It was used until the end of the year 2000 to determine precise values for the properties of the elementary particles of the Standard Model, e.g. the mass of the Z and W bosons. After the shutdown the LEP was disassembled to make room for the construction of the LHC.

1.1 The LHC

As its name implies, the LHC was built to accelerate hadrons, unlike the LEP which was operated with leptons (electron and positron). While leptons are elementary particles, hadrons are composite ones composed of quarks, which are kept together by the strong force (for more information on the particles of the Standard Model see Appendix A.1). As these hadrons have a significantly higher mass than the leptons used before, they lose considerably less energy due to the synchroton radiation while circulating[2] and hence can reach higher energies. The LHC will be used mainly for proton-proton collisions, but some periods are run with lead nuclei to allow heavy-ion experiments.

[1] Nobel Prize in physics for Carlo Rubbia and Simon van der Meer in 1984 for "their decisive contributions to the large project, which led to the discovery of the field particles W and Z, communicators of weak interaction"[8].
[2] Energy loss per circulation is inversely proportional to the forth-power of the rest mass (see Appendix A.2).

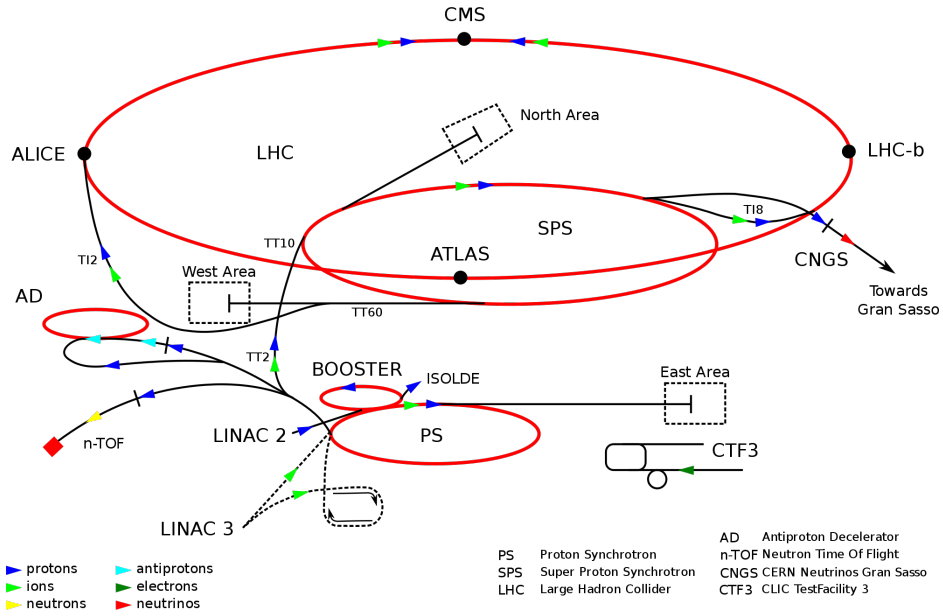

Figure 1.2: Illustration of the CERN accelerator complex[9].

As already mentioned at the beginning of this chapter, a series of accelerators[1] is used to increase the energy of the particles before they are injected into the LHC. There they reach their maximum speed and energy after circulating for about 20 minutes. Two beams of particles are circulating in ultra-high vacuum in opposite directions in adjacent beam pipes, which were built into the same 27 km tunnel the LEP was situated in before. They are held on track by 1,232 dipole magnets which bend the trajectory of the beams to a circular shape, while 392 quadrupole magnets are installed to keep the beams focused. The magnetic field needs to be adjusted to the increasing particle energy, and will reach more than 8 Tesla at its maximum. To achieve these high magnetic field strengths superconducting electromagnets are used. They are made out of coils of niobium-titanium wire and need to be cooled down to 1.9 K (-271.25°C) with liquid helium to ensure a safe superconductive state. The particles will not be accelerated as a continuous beam, but in bunches. At maximum power the LHC is designed to have about 2,800 bunches of ≈100 billion protons each circulating at 11 kHz around the ring. Yielding a peak luminosity of 10^{34} $cm^{-2}s^{-1}$ and a collision rate of 40 MHz at the intersection points in the four experimental caverns where the detectors were built. It is planned to increase the luminosity to ≈3 x 10^{34} $cm^{-2}s^{-1}$ during the Phase-I upgrade which is foreseen for the long shut-down around 2018. After that, the *High Luminosity - LHC* (HL-LHC) upgrade in 2022 is supposed to push it further up to ≈5-10 x 10^{34} $cm^{-2}s^{-1}$[10].

[1]So far not mentioned were the linear accelerators LINAC 2 and 3 as well as the proton synchroton BOOSTER which are used to accelerate the particles before they are injected into the PS (cp. Figure 1.2.)

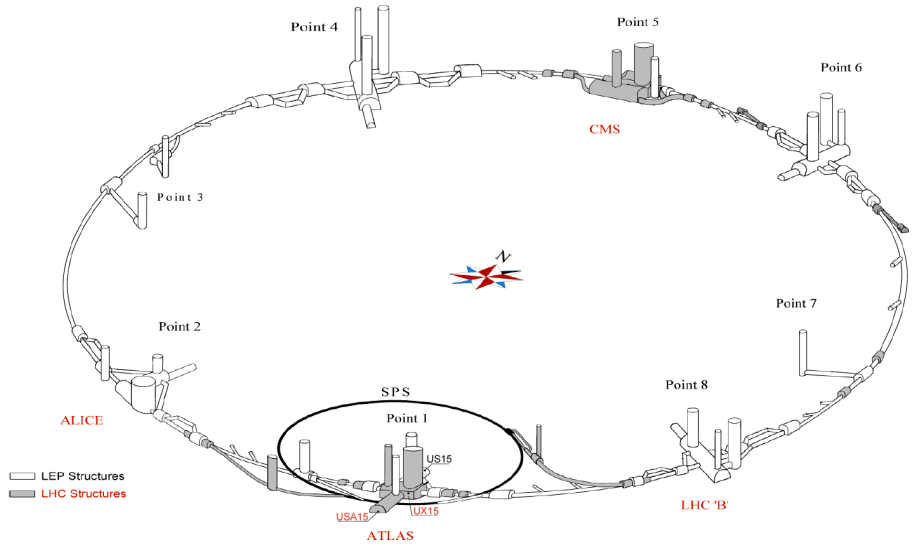

Figure 1.3: Underground structures from LEP and LHC at CERN[10].

1.1.1 Experiments

Six experiments have been built by international collaborations to study the particles produced during the collisions. Two big general-purpose detectors named ATLAS and CMS[1] are built around the interaction points at Point 1 and 5 respectively (see Figure 1.3 for a map of the underground structures of the LHC). They share these caverns with smaller experiments which focus on forward particles, LHCf is located at Point 1 together with ATLAS, and CMS shares Point 5 with TOTEM. ALICE at Point 2 and LHCb at Point 8 are medium sized detectors that are optimized for their special purpose.

CMS

The *Compact Muon Solenoid* (CMS) experiment[11] is based on a general purpose detector which is built around a cylindrical superconducting magnet (see Figure 1.4 for an overview of the detector). Between this magnet and the interaction point are three detector layers that track and measure different kinds of particles. At first, a silicon tracker is used to determine the momentum of the passing particles, before two calorimeter layers measure their energy. While the *Electromagnetic Calorimeter* (ECAL) is designed for photons and electrons, the *Hadron Calorimeter* (HCAL) can measure all particles made up of quarks. In this energy measurement the particles are slowed to a stop depositing all their energy in the calorimeter. On the outside

[1] Denotation of the abbreviations and some information on the experiments will be given in the following sections.

of the magnet four layers each of iron and muon detection chambers are alternating interleaved. The iron layers act as a return yoke to control the magnetic field and to absorb all remaining particles, except muons and weakly interacting particles like neutrinos. Muons penetrate the calorimeters and the iron layers, and are detected in the four layers of muon chambers. This is a very important task as muons are supposed to be produced when some of the predicted new particles decay, the Higgs Boson for example has a decay channel into four muons. The main physics goal of CMS is general exploring of particle physics on a TeV scale. This might shed light on physics beyond the Standard Model and also includes finding proof for the existence of the Higgs Boson.

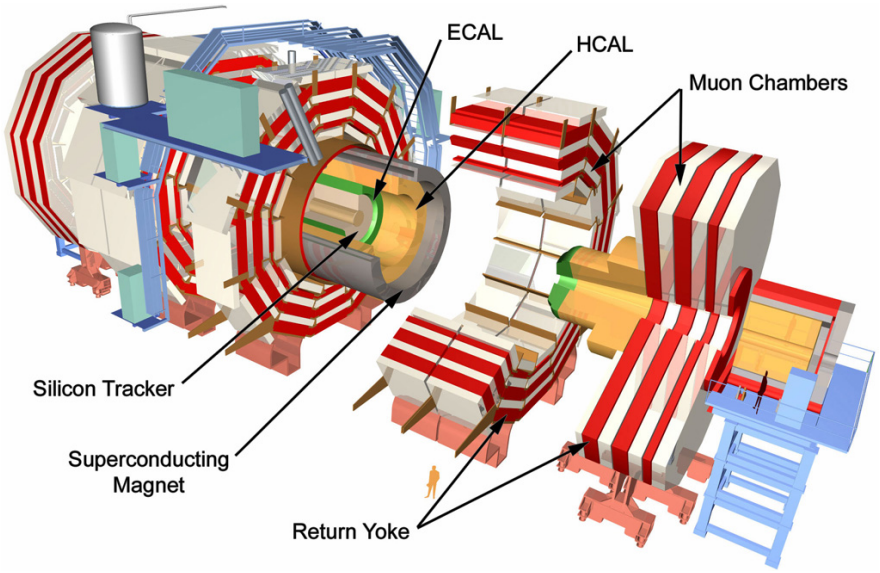

Figure 1.4: Drawing of the CMS detector built around the beam pipe and the interaction point. Picture taken from the CMS website[12].

TOTEM

Around the same interaction point as CMS the *TOTal Elastic and diffractive cross section Measurement* (TOTEM) detector system was constructed[13]. It consists of two sets of telescopes and *Roman Pots* (RP)[1] placed along the beam line, symmetrically in both directions from the interaction point. The telescopes T1 and T2 covering the pseudo rapidity[2] range of $3.1 \leq |\eta| \leq 6.5$ are placed at distances of 9 m and 13.5 m respectively, and are responsible for detecting charged particles produced by

[1] Roman Pots are silicon detectors in small secondary vacuum chambers (pots) that can be moved into the primary vacuum and thus very close to the beam axis.

[2] The pseudo rapidity $\eta = -ln\left(tan\frac{\Theta}{2}\right)$ is used as a coordinate to describe the angle of a particle relative to the beam axis. With Θ being the angle between the momentum of the particle and the beam axis.

inelastic scatterings. The roman pots are placed at insertions that can be moved into the beam pipe at positions 147 m and 220 m from the interaction point. They are responsible for the detection of leading protons only a few mm away from the beam axis. This is necessary as a goal of TOTEM is to measure the total proton-proton cross-section with a method which is independent from the luminosity, but makes use of the inelastic scattering rate and the elastic cross section[14].

ALICE

A Large Ion Collider Experiment (ALICE)[15] was built in the cavern at Point 2 to study quark-gluon plasma[1] and other particles created in lead-ion collisions. Therefore a combination of different detector types is used to gather as much and diverse information as possible. Enclosing the interaction point six layers of silicon detectors form the *Inner Tracking System* (ITS) to measure the location of the particles. This ITS is situated in a gas filled *Time Projection Chamber* (TPC) which provides track finding, momentum measurement and identification of the collision products. Surrounding the TPC is a *Transition Radiation Detector* (TRD) which acts as the main electron detector. Additional detectors like a *Time Of Flight* (TOF) measurement, a photon spectrometer (PHOS), an electromagnetic calorimeter (EMCAL), several muon spectrometer and other systems complete the experiment.

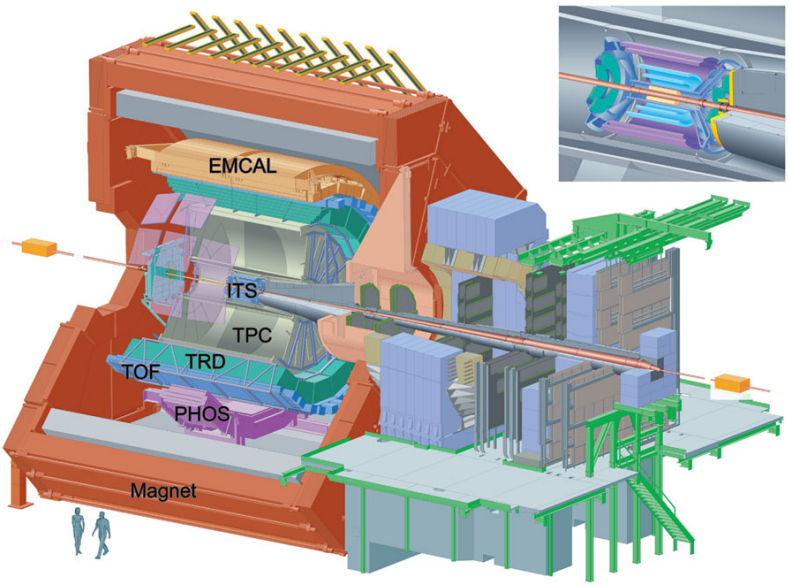

Figure 1.5: Illustration of ALICE with its individual detector systems. Upper right corner shows a closeup of the ITS. Picture was taken from the ALICE webpage[15].

[1] State of matter that is believed to have existed shortly after the Big Bang.

LHCb

The b in LHCb stands for beauty, as the research of this experiment at Point 8 is focused on the study of the beauty (or bottom) quark and its decay products[16]. The aim is to get a deeper understanding of the conditions shortly after the Big Bang and to figure out why matter[1] is dominating in our universe. The detector used therefor is a system of sub-detectors (see Figure 1.6) that does not surround the interaction point, but covers the forward region with several layers of different detector types. Closest to the collisions is the *VErtex LOcator* (VELO) which is measuring the tracks of the particles and separates primary and secondary vertices[2]. It is followed by the first *Ring Imaging CHerenkov* (RICH1) detector responsible for the identification of particles with low momentum, while the second one[3] will do the same with particles with higher momentum. A combination of silicon-strip and straw-tube detectors form the tracker system around a dipole magnet to detect charged particles. Calorimeter (electromagnetic and hadron, ECAL + HCAL) for energy measurement and muon detectors build the back of the detector.

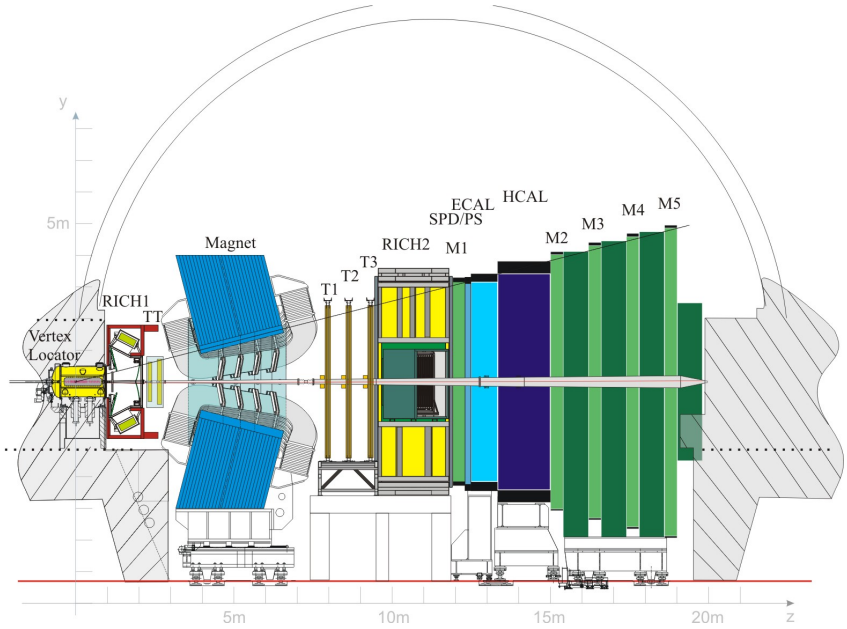

Figure 1.6: Drawing of the LHCb detector. Picture taken from Wikipedia.

[1] Matter and anti-matter were produced in equal quantities in the Big Bang, they annihilate each other while releasing energy.

[2] Vertex denotes the interaction point here, the primary vertex is the initial collision point and secondary vertices are points where particles decay into other particles.

[3] RICH2, located behind the tracker system and the dipole magnet with respect to the position of RICH1.

LHCf

The smallest of the experiments is LHCf (LHC forward) which consists of only two scintillator calorimeters. They are located between the beam pipes in both directions at a distance of 140 m from the collision point inside the ATLAS experiment (see Chapter 2 for details on ATLAS). The goal is to detect particles (mostly neutral pions) emitted at zero degree collision angle to simulate and study ultra-high energy cosmic rays.

ATLAS

A detailed description of the ATLAS Experiment will be given in the next chapter (chapter 2).

2

The ATLAS Experiment

The relevant experiment for this thesis is the *A Toroidal LHC Apparatus* (ATLAS) Experiment which will be introduced here[17]. A collaboration of more than 3,000 physicists from all around the world[1] work together to construct and run this detector. It is 45 m long, has a diameter of 25 m, weighs about 7,000 tons and is situated in the underground cavern UX-15 at Point 1 (see Figure 1.3). About 3,000 km of cables are installed to supply the detector with power and to read out the experimental data from its more than 100 million channels.

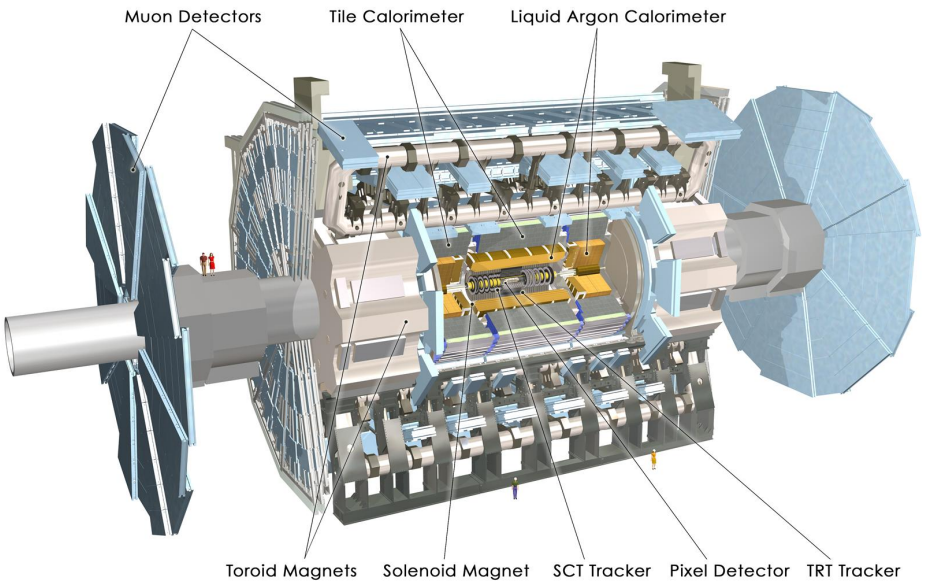

Figure 2.1: Computer generated image of ATLAS and its subdetectors[18].

[1] They belong to 174 universities and institutes from 38 countries.

2.1 Physics Goals

ATLAS was designed as a general-purpose detector to be able to detect a broad range of particles, including so far unknown ones that might be produced during the high energetic collisions. A primary goal is to search for evidence of the Higgs Boson to verify its prediction by the Standard Model, or to prove that it does not exist and that the Standard Model needs to be adapted accordingly. For this search possible decay channels of the Higgs Boson are measured. These channels and the according mass regions are excluded from further search if no statistical indication suggest extraordinary occurrences for them. This is necessary due to the very short lifetime of the Higgs Boson which can not be detected directly before it decays. After its mass region was narrowed down the focus is now on two possible decay channels, either two photons or four leptons.
In another area of research the so called *CP violation* is under investigation. CP is a combination of the charge conjugation and parity symmetry. That means that the laws of physics apply to a particle in the same way as to an antiparticle (C-Symmetry) and the equivalent for flipping the algebraic sign of spatial coordinates (P-Symmetry). The violation of these symmetries is one necessary part of the *Sakharov Conditions*[1] which lead to a particle generating process which produces more matter than antimatter. This is interesting because it is believed that during the Big Bang same amounts of matter and antimatter were produced and as they annihilate each other nothing but energetic radiation should be left, which is obviously not the case as we live in a world of matter.
Furthermore known particles of the Standard Model like the W boson or heavy quarks like the top quark will be measured to determine their properties with more precision than it was possible so far. Besides the consolidation of the known and predicted particles from the Standard Model, potential new physics like supersymmetry (SUSY), additional dimensions or another potential generation of quarks will be researched.

2.2 Sub-Detector Systems

The ATLAS detector is made up of three sub-detector systems and a combination of toroid and solenoid magnets around the collision point (see Figure 2.1). These superconducting magnets create a magnetic field which bends the trajectories of charged particles proportional to their speed[2]. Hence measuring the degree and direction of the curvature of a track allows to determine the momentum and charge of the particle. For this task the Inner Detector is placed in the 2 Tesla strong homogeneous magnetic field inside of the solenoid. Which itself is surrounded by a combination of

[1] Proposed by Andrei Sakharov in 1967, republished in 1991 [19].
[2] Responsible for that is the *Lorentz Force*, see Appendix A.3 for more details.

calorimeters to measure the energy of the particles. While most particles deposit all their energy in the calorimeter and thus are absorbed, muons will penetrate it. To track and measure them a combination of toroid magnets and muon spectrometers around all the other sub-systems complete the ATLAS detector. To describe a location in the detector *R*, ϕ and η are used as coordinates. The beam pipe is the z-axis with its origin at the interaction point in the center of the detector. *R* being the radial distance, ϕ the azimuthal angle and η is the pseudo rapidity which has the following relation to the polar angle θ: $\eta = -ln\left(tan\frac{\Theta}{2}\right)$.
A more detailed description of the individual detector systems will be given in the following subsections. The Pixel Detector being the relevant component for this thesis will follow in its own section.

2.2.1 Muon Spectrometer

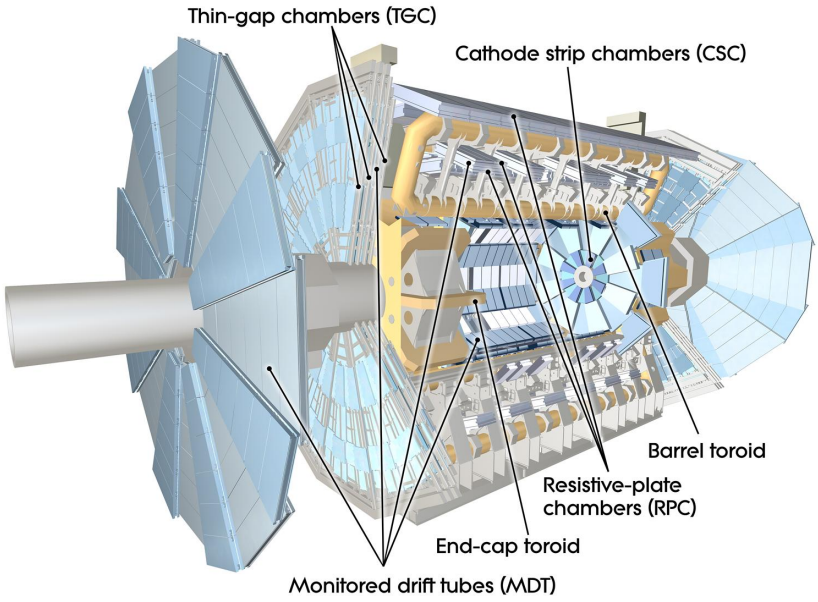

Figure 2.2: Schematic view of the ATLAS muon spectrometer and the toroid magnets[20].

The outermost structures of the ATLAS detector comprise the muon spectrometer. It is designed to precisely measure the muon tracks and utilize their detection to generate trigger signals. For this tasks a combination of four different detector types is employed as can be seen in Figure 2.2. For the precise measurements concentric barrel layers of *Monitored Drift Tube* (MDT) detectors are built around the interaction point. In the forward region MDT discs are installed for the same purpose, but as their occupancy is too high close to the beam pipe they are complemented with *Cathode Strip Chambers* (CSC) there. The generation of the fast trigger signals is

performed by *Resistive Plate Chambers* (RPC) which are built on top of the MDT barrels and the *Thin Gap Chambers* (TGC) in forward direction. All these subsystems combined provide about 1.2 million readout channels which are important to complement the measurement of the total energy of an event. While the fast trigger helps to select interesting events which are characterized by the emission of one or several muons.

2.2.2 Calorimeter

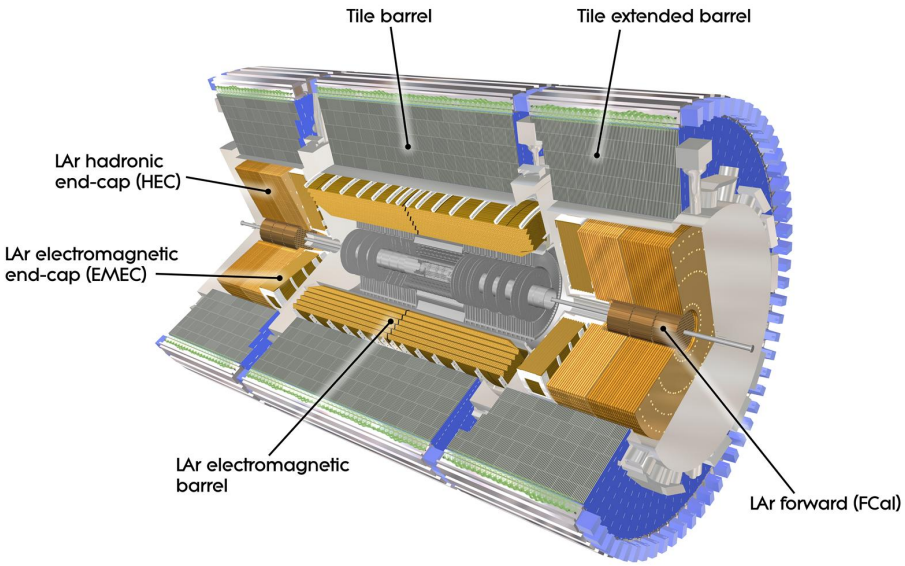

Figure 2.3: Computer generated image of the ATLAS calorimeter system[21].

The layer between the muon spectrometer and the solenoid magnet houses the calorimeter system which is responsible for the energy and location measurement of charged and neutral particles. In Figure 2.3 the internal two layer structure of it can be seen. The inner part is the electromagnetic calorimeter which measures particles that interact electromagnetically, like charged particles and photons. This is done by absorbing them in accordion shaped layers of lead and steel which are interleaved with liquid argon that acts as active material. The same active material, which needs to be kept at the correct temperature (-183°C) via a cryostat, is used in both the electromagnetic and the hadronic calorimeters in the end-caps that cover the forward region. The barrel section of the hadronic calorimeter consists of alternating layers of plastic scintillators as active- and steel as absorber material. It is designed to detect and measure hadrons that interact via the strong force and pass the electromagnetic

calorimeter.

2.2.3 Inner Detector

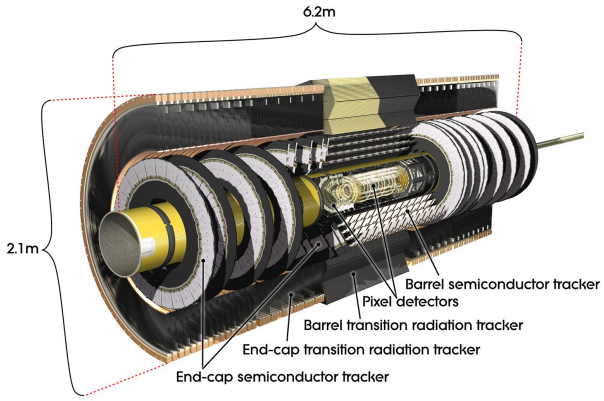

Figure 2.4: Drawing of the Inner Detector of ATLAS[22].

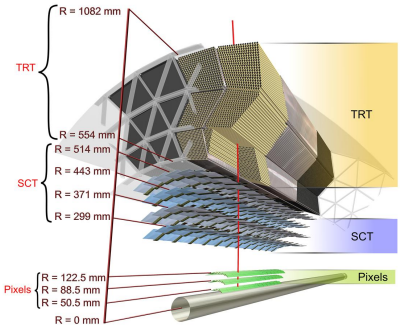

Figure 2.5: Illustration of the layer structure of the Inner Detector[22].

At the very center of ATLAS, enclosing the interaction point and within the magnetic field of the solenoid is the Inner Detector which is made up of three subsystems: the *Transition Radiation Tracker* (TRT), the *Semi Conductor Tracker* (SCT) and the *Pixel Detector*. Figure 2.4 and 2.5 illustrate the position of these subsystems within the Inner Detector and their layer structure around the beam pipe. This setup is a combination of high-resolution detectors surrounded by a continuous tracking system. Primary and secondary vertices can be reconstructed from the precise tracking information as well as the impact parameter[1].

The TRT which forms the outer layer is made up of 370,000 straws. These aluminum coated polyamide tubes are 4 mm in diameter and either 1.44 m or 0.39 m long, depending on if they are used for the barrel or the end caps respectively. They are filled with a gas mix, which is ionized by passing charged particles, and have a sense wire in the middle to collect the produced charges. By measuring the drift time during the readout, the spatial location of the hit can be determined with a resolution of about 200 μm per straw. To get strong signals from ultra-relativistic charged particles, materials with varying refraction indices are placed between the straws causing them to produce transition radiation. Each channel has two different thresholds which allows to distinguish between tracking hits (lower threshold) and hits caused by the transition radiation (higher threshold).

[1] Here the impact parameter denotes the perpendicular distance between the paths of two particles before they collide or interact.

Between TRT and Pixel Detector, the SCT provides four layers of silicon micro-strip detectors. These layers are populated with about 4 thousand modules which are equipped from both sides with silicon sensor chips. Each of these sensors has 780 80 μm wide strips on its 6 x 6 cm^2 active area and is combined with a second sensor to form strips of 12 cm length. Additional to the four barrel double-layers there are nine endcap discs in each forward region, and together this yields a total of more than 6 million silicon strips covering ≈61 m^2 for precise track measurements.

The final sub-system is the Pixel Detector which will be described in more detail in the next section due to its significance to this thesis.

2.3 Pixel Detector

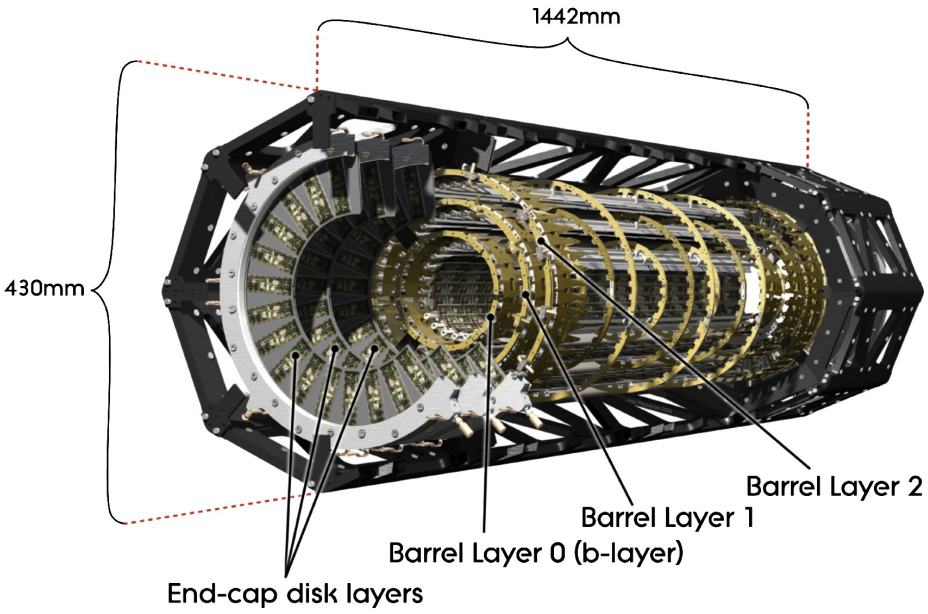

Figure 2.6: Illustration of the ATLAS Pixel Detector[23].

The innermost sub-detector is the Pixel Detector which is built directly around the collision point. Its physical dimensions of about 1.5 m x 0.5 m x 0.5 m seem rather unimpressive compared to the size of the complete ATLAS detector, but it is a very important component as it delivers 3 space-points per hit with a very high resolution close to the interaction point. This is crucial for the detection of very short lived particles like B-Hadrons,[1] and the overall impact parameter resolution of the Inner Detector is determined by it. From the three barrel layers of the Pixel Detector the

[1] Hadrons that contain a bottom- or b-quark.

one with the smallest radius (50.5 mm) is especially important for the reconstruction of secondary vertices of B-Meson decay. This allows to tag the emerging jets as b-jets[1] and thus the layer is called the *B-Layer*. The second and third layers with radii of 88.5 mm and 122.5 mm are called Layer-1 and Layer-2 respectively (see Figure 2.6). In each forward direction three end-cap disks are installed that can detect particles up to a pseudorapidity of $|\eta|$ = 2.5.
The barrel layers are comprised of carbon staves which have 13 sensor modules mounted to one side and a cooling pipe on the backside. By slightly overlapping the modules and arranging the staves in a turbine like structure the non-active area of the resulting barrel is minimized. Instead of staves, so called disk sectors are used to build the end-caps. Each of these sectors has 3 modules on both sides of its support and cooling structure, and 8 sectors are used to build one of the disks. An overview of the amount of staves, modules and channels per layer is given in Table 2.1.

	Staves/Sectors [#]	**Modules** [#]	**FE-Chips** [#]	**Channels** [#]	**Active Area** [m^2]
B-Layer	22	286	4,576	13,178,880	0.28
Layer-1	38	494	7,904	22,736,520	0.49
Layer-2	52	676	10,816	31,150,080	0.67
Total Layer	112	1456	23,296	67,092,480	1.45
Disk 0	8	48	768	2,211,840	0.0475
Disk 1	8	48	768	2,211,840	0.0475
Disk 2	8	48	768	2,211,840	0.0475
Total Disks	48	288	2,304	13,271,040	0.28
Overall Total	160	1,744	25,600	80,363,520	1.73

Table 2.1: Basic parameters of the ATLAS Pixel Detector.

2.3.1 Front-End Electronics

ATLAS Pixel Module

For the modules used in the construction of the Pixel Detector a hybrid approach is employed which uses different silicon substrates for the sensor and the readout electronics[25]. One 63 mm x 18.6 mm silicon sensor with an active area of 60.8 x 16.4 mm^2 which is subdivided into 46,080 pixels of 400 μm x 50 μm size[2] is used per module. Each of these sensor pixels is connected to the corresponding pixels on the readout chips by bump-bonding[3]. A total of 16 of these front-end (FE) chips arranged in two rows of 8 each is needed for the readout of one sensor (see the next subsection

[1] This process is known as b-tagging.
[2] This is the number of effective pixels, and there are special *long* pixel (600 μm) to cover the regions between the readout chips, see section 2.3.1.
[3] For this small solder balls are applied to each contact on one side, then both parts are combined and heated to melt the solder and connect them.

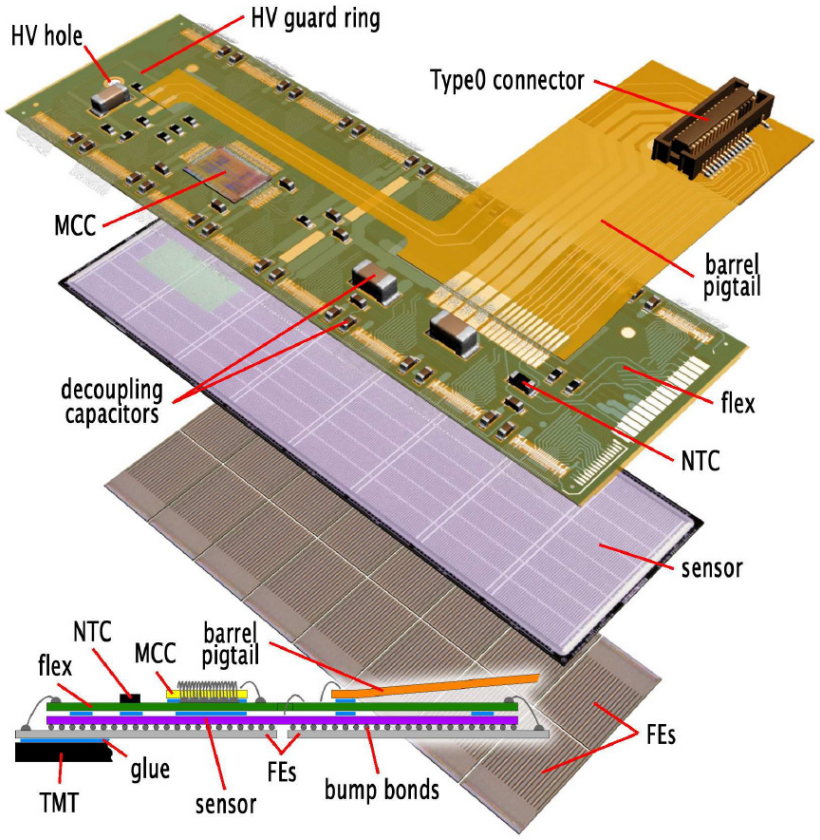

Figure 2.7: Explosion view of the components of the ATLAS Pixel Detector barrel front-end module [24]. The NTC is a resistor with a Negative Temperature Coefficient (thermistor) and is used to monitor the temperature of the module, the other components are described in the text.

for more details on this FE chip). On the opposite side a flexible *printed circuit board* (PCB), referred to as *flex-hybrid*, is glued to the sensor and the readout chips are connected to it via wire bonding. These wires are not only used for communication, but also supply the FE chips with power which is decoupled by passive capacitors on the flex-hybrid. This PCB was also specifically designed to be able to handle the *High Voltages* (HV) of up to 600 V which are needed for the depletion of the sensor. An overview of this assembly is given in Figure 2.7. All the FE chips send their data to the *Module Control Chip* (MCC)[26] after amplifying and digitizing it. They are connected to it in a star-topology via *Low Voltage Differential Signaling* (LVDS) serial links to reduce the amount of traces needed. The MCC is responsible for the configuration of the FE chips and the MCC itself after power-up. During data taking it distributes calibration, timing and trigger signals to the FE, reads out the data and manages the event building. This event data is then sent to the Optoboard[27] which establishes the connection to the off-detector electronics. The complete module is glued with

its FE chip side to the *Thermal Management Tile* (TMT) of the stave to enable heat transfer to the cooling pipe attached to the TMT. While the barrel modules have a flexible foil called pigtail to connect to the electrical services of the detector via micro-cables, the modules for the disk sections do not need pigtails as they have these cables directly attached to their flex-hybrid.

Silicon Sensor

The silicon sensor is the sensitive part of the pixel modules. Charged particles and photons passing through the semiconductor material ionize it and create electron-hole pairs along their way. By applying an external voltage these pairs can be separated and detected by the attached readout electronics. To create one electron-hole pair the silicon band gap of 1.12 eV needs to be overcome, and together with the energy loss due to the emission of photons an average of 3.62 eV is needed to ionize one atom. This small quantum allows a high resolution in measuring the energy, because the total amount of pairs produced is proportional to the energy deposited in the material[1]. Due to the high density of silicon, even very thin sensors produce strong signals which has the advantage that the traversing particles are not affected or deflected much. Todays manufacturing technologies allow structure sizes on a sub-micron[2] scale, so a very good spatial resolution can be reached as well.

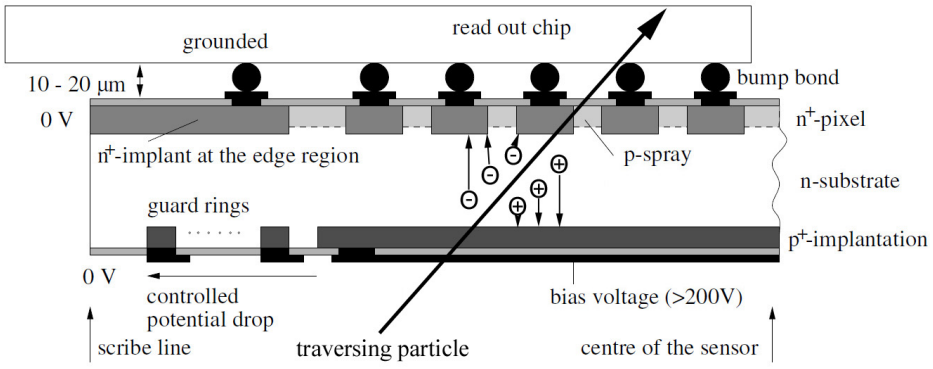

Figure 2.8: Drawing of the structure of the sensor for the Pixel Detector[28]. Between the n^+ pixels a moderated p-spray (with decreasing concentration towards the pixel) was used contrary to the drawing.

For the ATLAS Pixel Detector sensors a 256±3 μm thick n-type substrate was chosen which requires doping from both sides, n^+ on the pixel readout side and p^+ on the back side. On the p^+ side multiple guard rings around the doping zone are

[1] The amount of energy a charged particle loses while traversing matter is given by the *Bethe-Bloch Formula* (see Appendix A.4)

[2] Smaller than 1 μm.

implemented to isolate it from the sensor edges which allows operation with very high bias voltages of up to 600 V [29]. The n^+ implants on the other side need to get isolated from each other as they are set in a n-type substrate. This is done via the moderated[1] p-spray technique which not only covers the whole surface with a weak p-implant (p-spray), but has a higher concentration in the middle between two pixels and a reduced one towards them. While the p-layer is too weak to influence the n^+ implants, it is strong enough to prevent electrons from traveling between adjacent implants. An illustration of these structures can be seen in Figure 2.8. The sensor chips have a total of 47,232 pixels that are arranged in 328 rows with 144 columns each.

When the sensor is unirradiated, the bias voltage builds the depletion zone starting from the backside and extends it all the way to the front (upper part of Figure 2.9), hence completely isolating the pixels and minimizing the leakage current. During operation in the detector, irradiation will lead to a change of the effective doping concentration in the bulk and finally a type inversion will occur. After the inversion the depletion zone grows from the pixel side (see lower part of Figure 2.9), isolating the pixels and thus allowing operation even if the bulk is not completely depleted. The bias voltage necessary to reach complete depletion needs to be adapted to the changes in bulk type and increases with the radiation damage. As a countermeasure, oxygen impurities in the bulk silicon proved to decrease the radiation damage and thus the needed bias voltage increases at a lower rate than in standard silicon[29], this effect is illustrated in Figure 2.10.

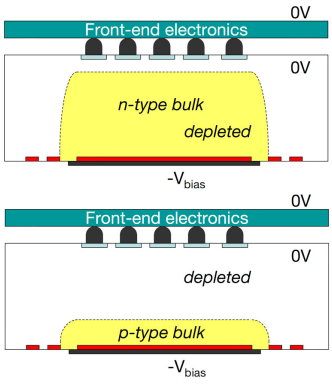

Figure 2.9: Illustration of the depletion zones before (upper picture) and after (lower picture) irradiation[25].

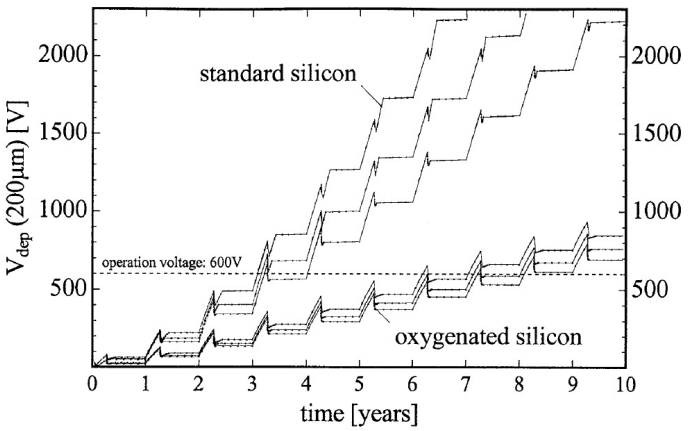

Figure 2.10: Illustration of calculations for the influence of oxygenated silicon on the necessary depletion voltage for the b-layer, three different warm-up scenarios were used[29].

[1] The advantage of the moderated p-spray over regular p-spray is increased high voltage tolerance before irradiation.

As mentioned before, 16 FE-chips arranged in two rows of 8 each are needed for the readout of one sensor tile. For technical reasons there is a pitch of 400 μm between these chips, and to prevent this area from being inactive two types of special pixels are incorporated in this regions[1]. While the common pixels are sized 50 x 400 μm^2 there are two columns of *long pixels* of 50 x 600 μm^2 covering the area between the two rows of FE-chips (the green and yellow pixels in Figure 2.11, the yellow are also ganged). In the gaps between the 8 columns of readout chips are 8 rows of uncovered pixels from which 4 each are connected to the adjacent FE-chip by connecting the pixel to another pixel which is covered by the readout chip. These readout cells with two pixels attached are called *ganged pixels* and alternate with cells that have only one connected (*inter-ganged pixels*). Due to this ganging the effective number of pixels that can be read out is 46,080 and not the 47,232 that are actually present on the sensor tile.

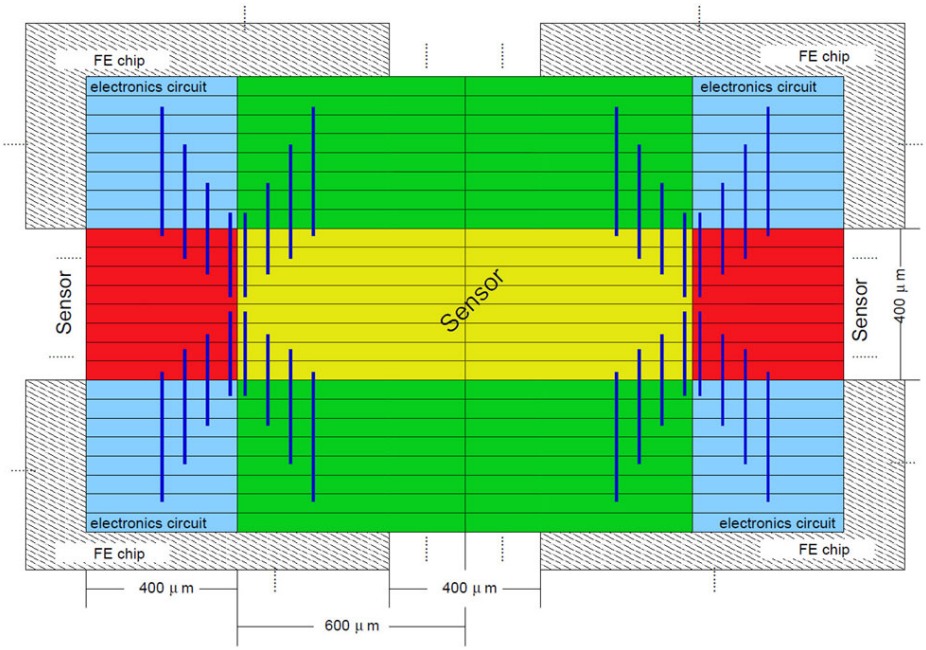

Figure 2.11: Special pixels in the gaps between the readout chips, green: *long pixels* covered by an FE-chip, red: *ganged pixels* that are between two chips and need to be connected to other pixel cells for readout (dark blue lines represent the connections made by metalization), yellow: *long-ganged pixels*, blue: common pixels [30].

[1] Basically there are five different types, but there are only two concepts and the other three types are mixes thereof.

FE-I3 Front-End Chip

For the readout of the pixel sensor the third generation of a special front-end chip called *FE-I3* is used. Its about 3.5 million transistors are implemented in a 0.25 μm radiation hard CMOS technology on a 7.6 x 13.6 mm^2 chip from which the upper 74% (7.6 x 10.8 mm^2) is the active area and the rest is occupied by the readout periphery. To accommodate to the high event rate of up to 40 MHz (bunch crossings every 25 ns) it is necessary to have a readout circuit for each individual pixel, because the shifting of the collected charges to the readout electronics at the edges of the sensor[1] would take too long. Therefore the FE-I3 chip is connected to the pixel sensor by bump-bonding each of its 2,880 readout channels directly to the individual pixel cells. The analog parts of the readout cells are arranged in a matrix of 160 rows with 18 columns from which 2 each are attached to the digital part of the readout architecture (see Figure 2.13).

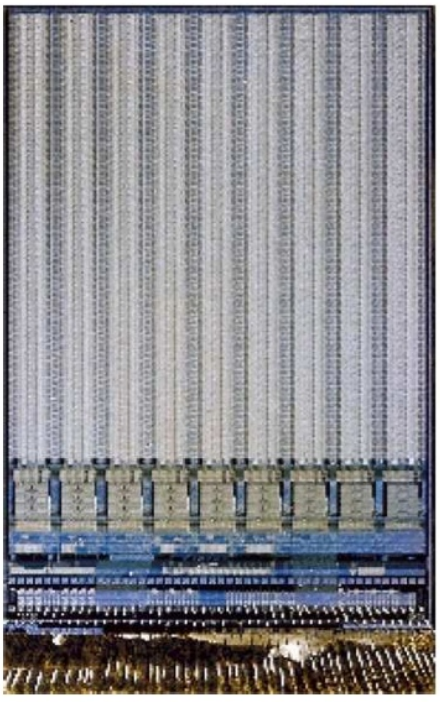

Figure 2.12: Photograph of the FE-I3 readout chip[31].

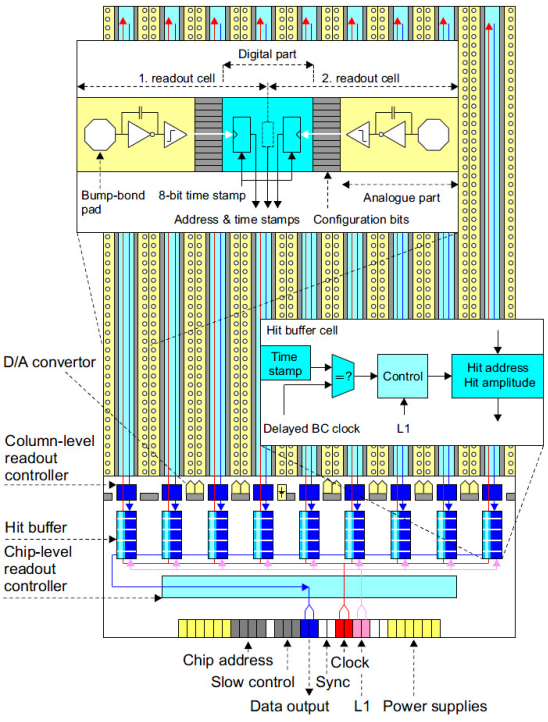

Figure 2.13: Overview of the architecture elements of the FE-I3 readout chip[32].

In these cells the charge collected from the sensor (via the bump bond) or injected by a pulse generator[2] is pre-amplified and loads up a so called feedback capacitor C_{fb} which is then discharged by the current I_{fb}. In the next step a discriminator

[1] This technology is used in charge coupled devices (CCD) which are often used for digital imaging.
[2] This is done to test and calibrate the readout chip.

digitizes this analog signal by outputting a digital pulse as long as the input signal is above a certain programmable threshold. The length of this pulse is measured by an 8-bit counter in multiples of the bunch crossing period of 25 ns and is called *Time over Threshold* (ToT). It is proportional to the amount of charge deposited in the feedback capacitor and can be used to calibrate the individual readout channels to compensate small disparities caused by a non perfect manufacturing process. This ToT value, the address of the corresponding pixel and the time stamp of the hit for two adjacent columns are transferred (with 20 MHz) to the digital part of the chip. There this information is stored in hit buffers at the end of the columns (EoC) (see Figure 2.13). This is called *column drain architecture* as data for all hits is transferred to the buffers, even if there is no trigger signal for it resulting in the data being deleted after the latency of the Level 1 (L1) trigger (3.2 μs) has passed. If a L1 trigger is received its time stamp is compared to the buffered hits, and all corresponding data is transferred with 40 Mbit/s to the MCC which was already mentioned on page 20.

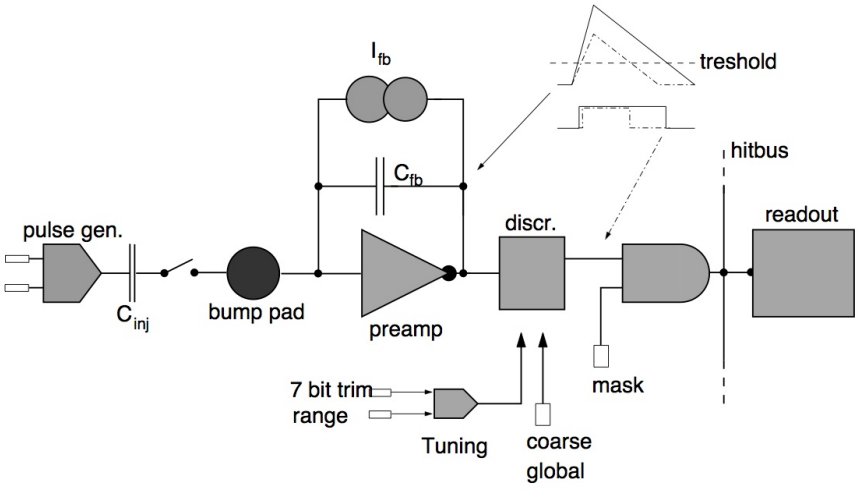

Figure 2.14: Simplified drawing of the analog part of the FE-I3 chips[31], including an example of how the triangular capacitor charge characteristic is translated into a digital pulse with a length corresponding to the Time over Threshold (ToT).

Optoboard

The connection between the detector front-end electronics and the off-detector readout electronics is established via optical fibers driven by custom made Optoboards [27]. These 2 x 6.5 cm^2 beryllium-oxide (BeO) PCBs[1] are plugged, with their 80-pin

[1] This material was chosen over the standard FR-4 (which is usually used for PCBs) for its superior thermal conductivity and thus better heat management.

connector on the bottom side, into the *Patch Panel 0* (PP0). This is connected to the Pixel Modules and is located close to the Pixel Detector. While the layout of the bottom side is the same for all Optoboards, two different types for the top layer were designed (see Figure 2.15). From the total of 272 installed Optoboards 228 are used for the readout of the pixel modules in the disks and in Layer 1 and 2 (D-Boards) while the remaining 44 are responsible for the B-Layer (B-Boards). This diversity is due to the close proximity of the B-Layer to the collisions which causes the hit occupancy to be significantly higher than in the other layers of the Pixel Detector. In detail that results in 1 link at 40 Mbit/s for Layer-2, 1 link at 80 Mbit/s for Layer-1 and the disks and 2 links at 80 Mbit/s (160 Mbit/s combined) for the B-Layer. These links are established by a combination of *Vertical Cavity Surface Emitting Laser* (VCSEL) diode arrays and *VCSEL-Driving Circuit* (VDC) chips. One VDC converts 4 LVDS inputs into single-ended *Non-Return-to-Zero* (NRZ)[1] signals to drive the VCSEL diodes which have a data rate of 80 Mbit/s, thus 2 VDCs are needed per 8-channel VCSEL array. From these 8 channels only 6 or 7 are used to have spare connections. Besides slightly different routing on the top layer of the B-Layer Optoboard the main difference to the other Optoboards is that it is populated with the full set of 4 VDCs and 2 VCSEL arrays, instead of only 2 and 1 respectively. This is needed to be able to transfer the 160 Mbit/s (2 x 80 Mbit/s) coming from the B-Layer modules.

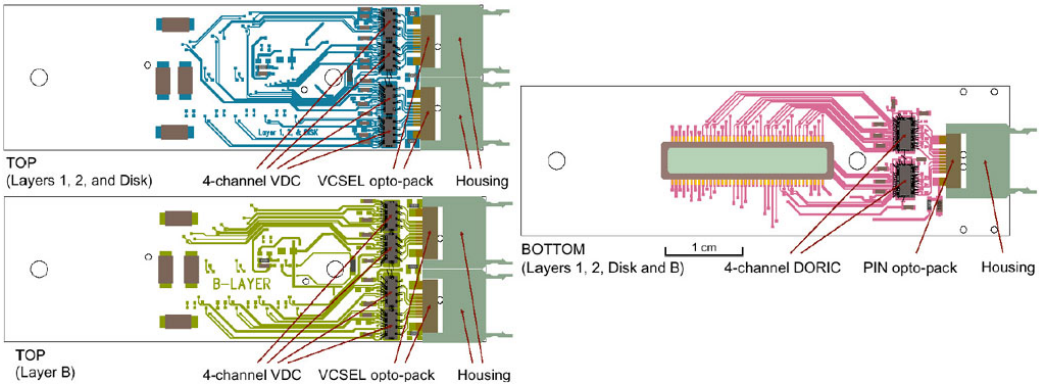

Figure 2.15: Layout of the two different types of Optoboards[27]. While the bottom layer is the same for both, the top layer has small differences in the routing and only half the VDCs and VCSELs are populated on the version for the disks and Layer 1 & 2.

On the bottom layer an 8-way[2] PiN diode[3] array converts the optical transmissions from the off-detector electronics into electrical signals. Each of these data streams

[1] NRZ is the most simple binary code, "1" is represented as one signal condition (typically a high/positive voltage) and "0" as another (low/negative voltage).

[2] Only 7 channels are used to have some redundancy.

[3] A PiN diode has an intrinsic (not or only lightly doped) region between the p- and n-type regions. For the usage as a photo-diode it is reverse biased and basically acts as a silicon detector for photons, or in this case as an optical to electrical signal converter.

contains the *Bi-Phase Mark* (BPM) encoded *Timing, Trigger and Control* (TTC) signals together with a 40 MHz clock[1] for one pixel module. Two *Digital Optical Receiver ICs*[2] (DORIC) with 4 channels each are responsible for the *Clock and Data Recovery* (CDR) of these streams. The single-ended inputs from the PiN array are decoded into LVDS data and clock signals within strict requirements. The duty cycle of the recovered 40 MHz clock needs to be (50±4)% [33] with less than 1 ns total timing error, and the *Bit-Error Rate* (BER) needs to be below 10^{-11} at the end of the expected liftime[3] of the DORIC. The clock and data signals are then sent to the MCCs which distribute them on the pixel modules.

2.4 The IBL upgrade for the Pixel Detector

2.4.1 B-Layer

When the current B-Layer was designed, no technology was available to construct it with enough radiation tolerance to yield good performance while operating it at design luminosity (10^{34} cm^{-2}s^{-1}) for more than 7 years[34]. With time the radiation damage will decrease the performance of the sensors and readout electronics. The latter suffers especially from direct ionizing radiation which changes the threshold voltage of the transistors by trapping charge in the gate oxide. The sensors on the other hand are more sensitive to displacement damage in the silicon bulk which increases the leakage currents and the bias voltage necessary for full depletion (this effect was shown before in Figure 2.10). These effects lead to less charge being collected in the sensor and at the same time more noise[4] is induced. This requires higher thresholds in the discriminator stage of the analog readout and hence degrades the performance even more.
Besides these physical effects, additional problems will arise from the architecture of the readout electronics which were designed to be fully effective at the LHC design luminosity and a Level-1 trigger rate of 100 kHz. But its limitations will already deteriorate the efficiency with the foreseen luminosity increase to ≈3 x 10^{34} cm^{-2}s^{-1} during the Phase-I upgrade. Two components of the readout chain are responsible for this, the FE-I3 chip and the MCC.
The column drain architecture of the FE-I3 which transfers hit data from all 320 pixels of a double column sequentially via a shared bus to the EoC buffers is very sensitive to the hit occupancy. This occupancy scales with the luminosity as it causes

[1] This is a "byproduct" of the BPM encoding which has a transition for each rising clock edge and another one at the falling edge if a data bit (a logic "1") is to be transmitted.
[2] Integrated Circuit
[3] Corresponding to a radiation dose of 50 Mrad.
[4] E.g. the leakage currents cause a higher power consumption which increases the temperature and thus the noise.

more hits to be registered which in turn increases the utilization of the bus. Three effects contribute to the total inefficiency: After it receives a hit a pixel stays busy until the data is transferred to the buffer (*busy/waiting inefficiency*), if this waiting time is longer than the L1 latency it will arrive too late in the buffer (*late copying inefficiency*) and even if the pixel readout is not delayed a second hit can occur so fast after the first that it is not being recognized (*double-hit inefficiency*). How these effects add up can be seen in Figure 2.16. At a hit rate corresponding to approximately the luminosity of the Phase-I upgrade the bus is saturated and the inefficiency increases dramatically.

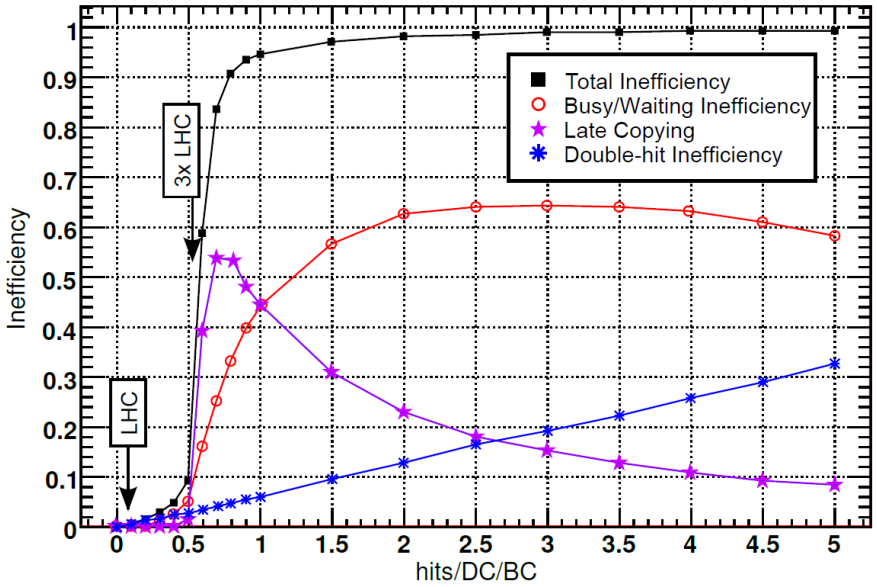

Figure 2.16: Inefficiencies (1 = 100%) of the FE-I3 in the current B-Layer at a distance of 5 cm from the collisions in relation to the number of hits per double-column (DC) and bunch crossing (BC)[35]. 3x LHC notes the inefficiency resulting from the luminosity increase during the Phase-I upgrade.

More hit occupancy corresponds to more data needing to be transferred to the off-detector electronics, assuming the L1 trigger rate stays the same. The MCC which collects the data from the 16 FE-I3 chips on the module has two data connections to the Optoboard, each can transmit either 40 or 80 Mbit/s, hence a maximum bandwidth of 160 Mbit/s. To transfer the data of the B-Layer before the Phase-I upgrade, both links are already required to run with 80 Mbit/s. Hence the only solution to avoid data loss at a higher luminosity is to reduce the trigger rate.
But a good performance of the B-Layer is mandatory for the physics program of the ATLAS Experiment, because it provides very important data for vertex reconstruction and b-tagging. To ensure proper function it was initially planned to adapt to higher luminosities by replacing the B-Layer. But last minute design changes during the

construction of ATLAS prevent an extraction of the B-Layer from the Pixel Detector now. A *B-Layer Task Force* was established and came to the conclusion that the only viable option would be to insert a forth pixel layer between the current B-Layer and the beam pipe[36].

2.4.2 IBL

The first step in the design of the *Insertable B-Layer* (IBL) was to assess the space available for it. Between the current beam pipe (inner radius of 29 mm) and the B-Layer there is only a free radial space of 8.5 mm (see Figure 2.17). To be able to insert an additional pixel layer, a new beam pipe with a 4 mm smaller radius will be used which increases the free space to 12.5 mm. These still tight constraints will require a very delicate extraction and insertion operation, as a 7.3 m long piece of the beam pipe needs to be replaced. Issues to be considered include the bending of the beam pipe when it is separated from its support structures, the toxic material it is made of (beryllium) and the overall radiation environment (due to activated material in the detector) in which this work has to be performed.

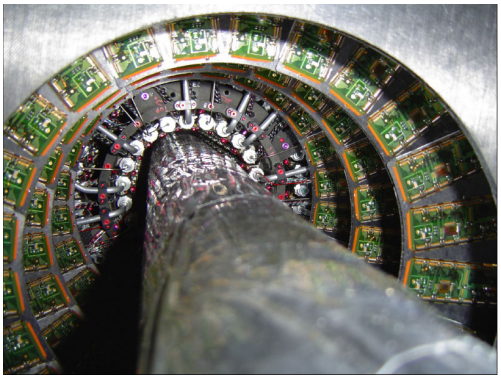

Figure 2.17: Photograph showing the gap between the current beam pipe and the B-Layer of the Pixel Detector[34].

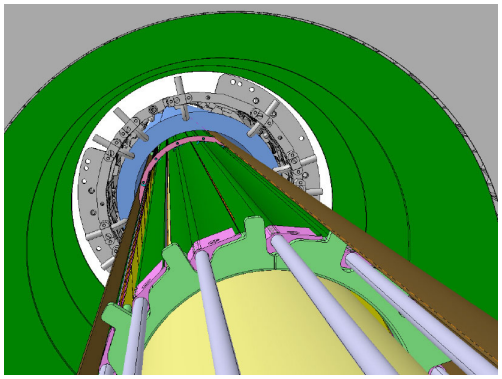

Figure 2.18: Graphical representation of the IBL inserted together with the smaller beam pipe[34].

The IBL will consist of 14 staves with a length of 72 cm, each of them equipped with 20 sensor modules (12 double chip and 8 single chip modules, see section 2.4.2 for further information). These staves will be arranged in a cylindrical form with the sensors at an average radial distance of 33 mm and a tilt of 14° ensuring overlap and thus full coverage in the Φ direction. Along the beam pipe (z-axis) an overlap of the modules is not feasible due to the limited space, hence gaps between the modules and thereby inefficiencies are unavoidable[34].
The limitations of the sensor and FE-I3 readout chip used in the current B-Layer were mentioned in the previous section (2.4.1). Utilization of these components for the IBL

would worsen their issues even more due to the shorter distance to the interaction point and therefore increased hit rate and radiation damage. Hence new sensors and new readout chips were designed for application in the IBL.

The FE-I4 Readout Chip

The FE-I4 is the next generation of the front-end chip and is manufactured in 130 nm CMOS technology. This feature size has a high inherent radiation hardness which was improved even further by a special layout design (e.g. additional redundancy), providing a tolerance of more than 300 Mrad[37]. Of its total area of 20.2 x 18.8 mm^2 only a small strip of 2 mm (x 20.2 mm) is needed for the peripheral readout part. The active area makes up almost 90 % of the chip and is divided into an array of 26,880 pixels (80 x 336) with a size of 50 x 250 μm^2 each. The analog readout is similar to the one of the FE-I3, each pixel has an individualy adjustable amplifiying, shaping and digitizing stage. But the digital part has a new architecture, contrary to the former column-drain technique all hit and ToT (4-bit resolution) information is stored locally, close to the pixels and is only transferred to the periphery if a L1-trigger for it was received.

In a double column 4 pixels are grouped together in a readout block, they store their hit data in 5 entry deep buffers which is kept for an adjustable latency of up to 255 clock cycles (40 MHz). Only if a corresponding trigger is received the data is transmitted to the peripheral logic where it is 8b10b encoded and shipped off directly to an Optoboard with 160Mbit/s. Less than 1 % of the pixel hits are actually transferred to the EoC logic which keeps the occupancy of the bus low and will result in an inefficiency of only $\approx$ 0.6 % for the Phase-I luminosity of 3 x 10^{34} $cm^{-2}s^{-1}$. As a positive side effect, the power consumption of the digital part will be below 10 μW per pixel for the IBL[38]. To control the FE-I4 command and configuration signals from the off-detector electronics are received directly via a serial LVDS connection to an Optoboard and are then decoded locally. Hence the FE-I4 has no need for a MCC anymore.

IBL Sensor Development

For the IBL three different sensor technologies were evaluated, *planar silicon*, *3D silicon* and *diamond*. *Planar silicon* is the option understood best, as it is commonly used and can be produced at relatively low costs with a high yield, but it needs to be cooled down the most and requires the highest bias voltages. The current Pixel Detector is based on planar sensors as described in section 2.3.1. For *3D silicon* sensors the electrodes for the pixels are not realized by doping strips on the surface, but by alternating pixel and bias "pillars" in the bulk as it is illustrated in Figure 2.20.

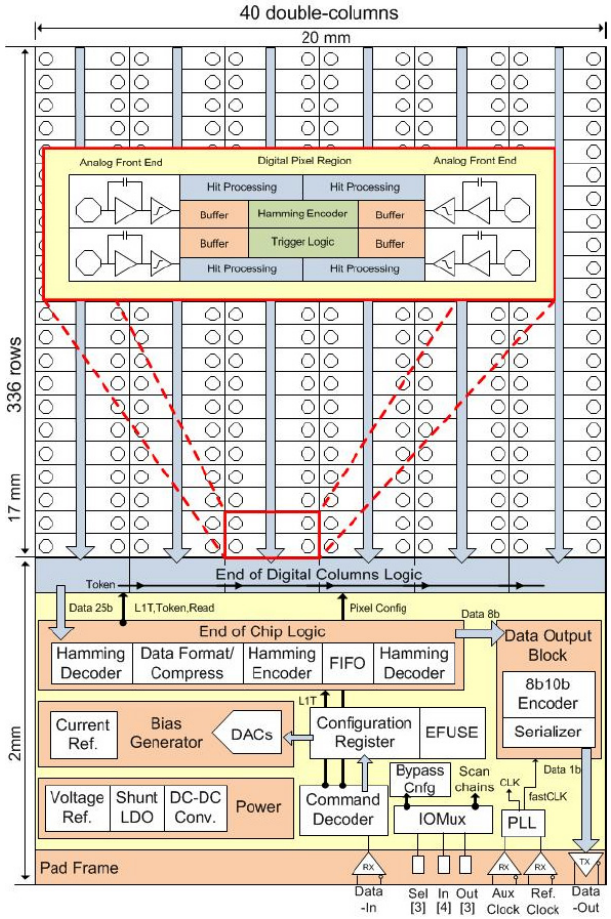

Figure 2.19: Overview of the architecture of the FE-I4 chip, not true to scale[37].

These columns of pillars can be arranged at shorter intervals than the thickness of the bulk which has the advantage that lower depletion voltages can be used. Also the charge can be collected faster as it does not have to travel that far even though the same active length is available for an ionizing particle. This, together with the parallel arrangement of the electrodes to the particle tracks leads to an increased radiation hardness. *Diamond* has a large band gap of 5.5 eV (Si 1.12 eV) and a high displacement energy[1] of 43 eV (Si 13-20 eV) per atom, both are features that ensure a good radiation tolerance. But the trade-off is that only a limited number of charge carriers is produced by a traversing particle. Besides that, other advantageous properties include low leakage currents and capacitances which decrease the noise. The thermal conductivity is so good that the sensor does not need to be cooled, but might even be used as a heat spreader for the readout electronics[39][40]. Both 3D

[1] This is the minimum kinetic energy an atom needs to be permanently displaced from a solid.

and diamond are rather new sensor technologies, so it needs to be demonstrated that they can be produced at acceptable costs with a good yield.

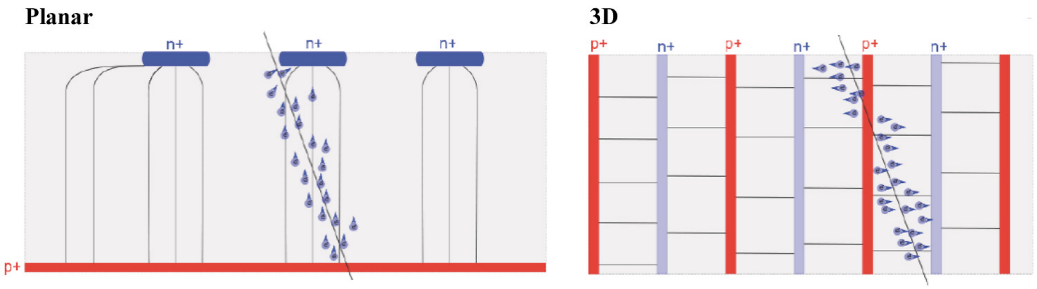

Figure 2.20: Illustration of the difference between planar (left) and 3D (right) silicon sensors[41].

IBL Modules

It was previously planned to insert the IBL in 2016[34], but the schedule changed to 2014 and thus only so far well tested technologies will be used. These are slim-edged n-on-n planar- and double sided 3D silicon sensors. The modules of the IBL are constructed as single- (3D) and two-chip (planar) assemblies. In both cases it is important to keep the gap between two adjacent chips as small as possible as no overlap is available (in z direction) to cover this inactive area.
For the planar sensor this is achieved by using two rows of long-pixels with 450 x 50 μm^2 (instead of 250 x 50 μm^2) to cover the area between the two FE-I4 chips on the double-chip module. The inactive edge on the outer rows is minimized by extending the pixels to 500 x 50 μm^2 reaching beyond the guard rings (this is the reason why they are called slim-edged) and resulting in a gap between two planar modules of about 200 μm.
The low bias voltage needed for the 3D sensors allows to remove some of the guard structures on the sides of the sensors while still being able to safely operate them. This yields about the same gap of about 200 μm, both between two single-chip (3D) modules and between single- and double-chip modules.
How these modules are arranged on a stave is shown in Figure 2.21. The 3D sensors with their column structure perpendicular to the beam axis are especially suitable for regions with high pseudo-rapidity and thus are placed on the outer ends of the stave.

	Staves	Planar Modules	3D Modules	FE-I4	Channel
IBL	14	168	112	448	$\approx$ 12 x 10^6

Table 2.2: List of the number of front-end readout electronics needed for the IBL.

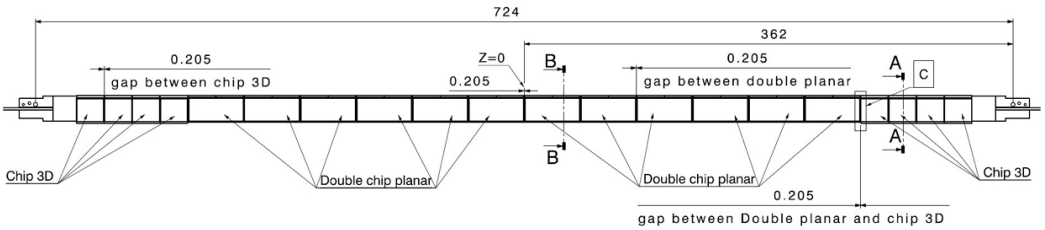

Figure 2.21: Drawing of the IBL stave with its 12 planar double-chip- and 8 3D single-chip-modules[42]. Numbers are given in mm.

IBL Optoboard

Due to the higher data rate of the FE-I4 chips, new Optoboards that can handle 160 Mbit/s need to be developed for the IBL as well. As each stave is equipped with 32 FE-I4s it was decided to manage 16 chips per Optoboard. Two 12-channel VCSEL arrays and one 12-channel PiN array are combined on one Optoboard with new receiver / driver ASICS that utilize 8 channels while keeping the remaining 4 as spare connections (see Figures 2.22 and 2.23). This way two FE-I4 chips share a TTC link, but have one data link each, hence 28 new Optoboards (Figure 2.24) are needed for the IBL yielding a total of 448 data- and 224 TTC links.

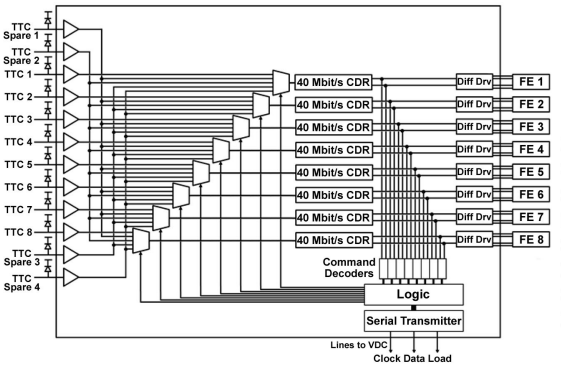

Figure 2.22: Illustration of the PiN receiver and decoder ASIC for IBL managing the 40 MBit/s TTC data[43].

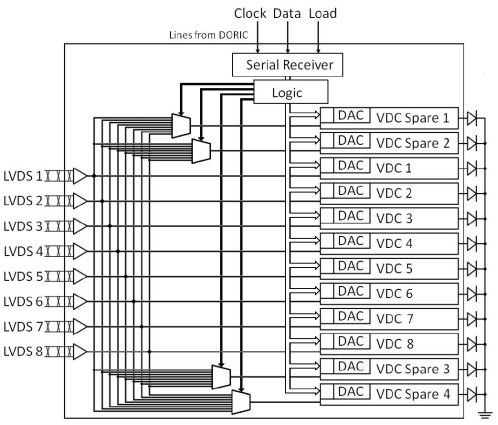

Figure 2.23: VCSEL driver circuit for the IBL Optobard[43].

IBL Performance

The installation of the IBL as a forth layer will not only guarantee the current performance of the Pixel Detector, but will improve it. To estimate its impact on the reconstruction of tracks and vertices as well as on the b-tagging performance several simulation based studies were conducted[34]. For this the IBL detector layer was

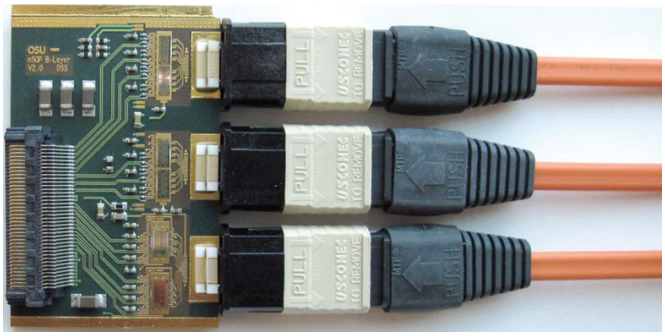

Figure 2.24: Photo of a prototype of the IBL Optoboard with its two VCSEL arrays (top connections) to send data to the off-detector electronics and the PiN array (lower connector) to receive the TTC commands[44].

integrated into the ATLAS Inner Detector software system, and standard reconstruction and analysis tools are used to get results as realistic as possible.
Tracking of single particles and particles in jets gains precision by the improved impact parameter resolution which is due to the additional layer at a closer radius and its smaller pixel size. A precise impact parameter resolution is improving the performance of an algorithm used for b-tagging (IP3D) as well. Reconstruction of primary and secondary vertices also benefits from the IBL and contributes to the gain in b-tagging efficiency via an algorithm based on secondary vertices (SV1). In Figure 2.25 the results of these algorithms are compared for ATLAS with and without the IBL, a combination of both algorithms yields a better performance for the expected pileup during Phase-I than the current system has with no pileup[34].

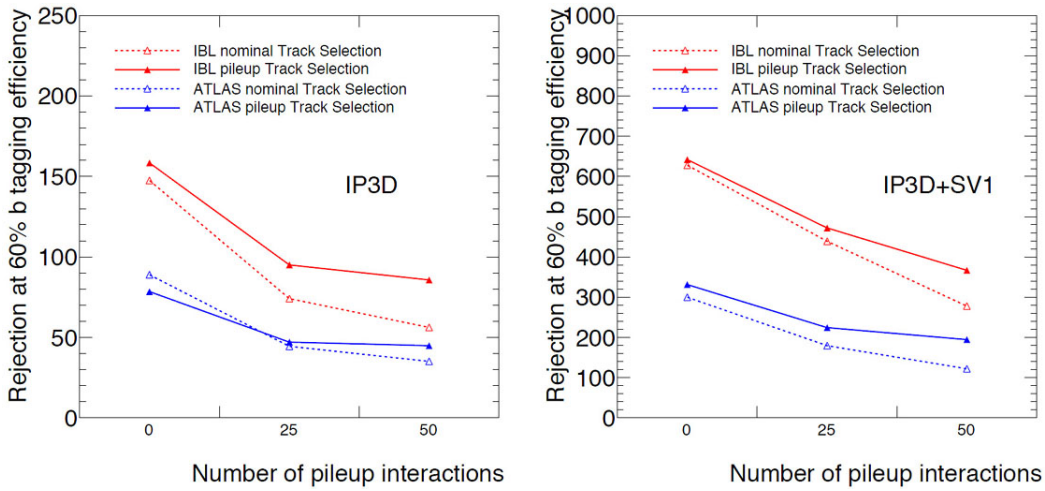

Figure 2.25: Rejection of light jets in $t\bar{t}$ events at 60% b-tagging efficiency[34]. Average number of pileups is expected to be around 50 for a luminosity of 2 x 10^{34} $cm^{-2}s^{-1}$.

3

ATLAS Off-Detector Electronics State of the Art

In the previous chapter the ATLAS Detector was described in detail as it is important to detect the interesting physical events, but the best detector system is useless without its off-detector electronics. These make sure that the detector can be configured and controlled and the data from millions[1] of channels is readout and provided for further analysis and storage. Contrary to the front-end electronics which are situated in or close by the detector in a typically radiating environment the off-detector electronics are placed apart in a so called counting room which is shielded from radiation and thus is accessible during data taking.
In this chapter an overview of the current state of the art of these electronic components for the readout of the ATLAS Pixel Detector is given. At first, the currently installed VME system is introduced and described in detail. Then another already existing readout system based on the more recent ATCA standard is presented as a potential alternative.

3.1 Current VME System

For the ATLAS Experiment the counting room is located in the *Underground Service Area-15* (USA-15)[2] which is a 20 x 62 m^2 cavern, shielded from the ATLAS Detector by a 2 m thick concrete wall. There are two floors with racks for the off-detector electronics from which 9 VME64x[3] crates are used to house the components necessary for the readout of the Pixel Detector.

[1] About 100 million for the ATLAS Detector.
[2] See appendix B.1 for an illustration of the ATLAS cavern.
[3] Versa Module Eurocard in the 64 bit version with extensions specified in ANSI/VITA 1.1-1997.

The connection to the front-end Optoboards (introduced in section 2.3.1) is established by optical fibers with a length of about 80 m. In the counting room these fibers are plugged into the *Back Of Crate* (BOC) card which handles the opto-electrical conversion in both directions, to and from the detector front-end. As the redesign of this component is the topic of this thesis there will be a thorough description of its features and functionalities in section 3.1.2. The BOC[1] is paired with a *Read Out Driver* (ROD) card which is plugged into the same slot, but from the front of the crate. Up to 16 of these card pairs can be installed in each of the 9 crates, and a total of 132 is required to handle the readout of the ATLAS Pixel Detector. Additionally each crate is equipped with one Single Board Computer (SBC) which allows to control the data acquisition in the crate and one interface module to the *Timing, Trigger and Control* (TTC) network of the accelerator complex.
Event data is built on the ROD and shipped off to *Read Out Buffers* (ROB) via an S-LINK[2] connection on the BOC (see Figure 3.1 for an illustration of the readout chain). Details about this ROB as well as further information on the *Data Acquisition* (DAQ) and *Read Out System* (ROS) of the ATLAS Experiment can be found for example in [45] and will not be further elaborated in this thesis.

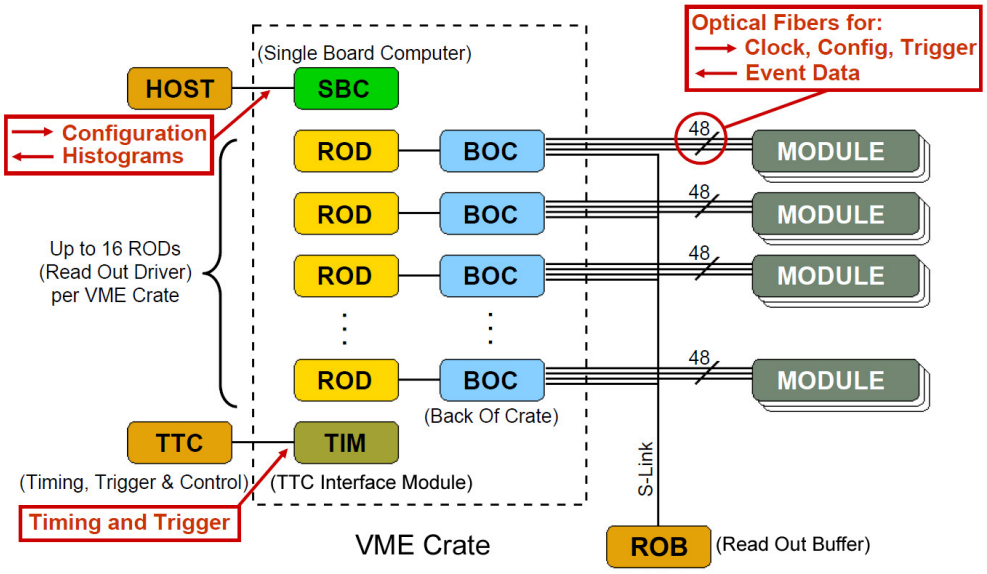

Figure 3.1: Overview of the VME crate structure for the data readout of the ATLAS Pixel Detector[46].

[1] *BOC card* would be the correct expression, but in this thesis BOC always stands for the card and never for the physical location in the crate. Hence from now on the BOC card is referenced as *the BOC*.
[2] A *simple link* specification developed at CERN, more details will be given in the BOC section (3.1.2).

3.1.1 The crate components SBC, TIM and ROD

A short introduction of the individual electronic boards installed in the readout crates will be given here. The BOC will be described in more detail in its own section as the redesign of it is the topic of this thesis and thus more background information is mandatory.

SBC

As crate controller (hence the SBC is also known as *ROD Crate Controller* (RCC)) a commercial VP-110 *Single Board Computer*[47] (SBC) from Concurrent Technologies is used. This 6U[1] single slot VME card features a Pentium III processor running at 933 MHz, 512 MB of RAM, two 10/100 Mbit/s Ethernet connections and it interfaces to the VME bus by a Tundra Universe II chip. It acts as the VME bus master and communicates with the ROD cards and the *TTC-Interface Module* (TIM) in the crate, the BOC cards are not connected to the VME bus. The RCC can manage the crate in a stand-alone mode to perform tasks like configuration and calibration (of the front-end electronics) but usually it is steered by commands from the ATLAS DAQ and *Detector Control System* (DCS) software which are received via its Ethernet connections.

TIM

The TIM card is connected to the TTC system of the ATLAS Detector, it receives the Level-1 trigger and distributes it to the ROD cards in the crate. If necessary, a busy-signal can be sent to the *Central Trigger Processor* (CTP) in order to stop the sending of further trigger signals, this happens for example if one of the ROD cards in the crate is busy (masked OR of all ROD busy-signals in the crate). Additionally fast commands and event IDs are received from the TTC and passed to the ROD cards, the *Bunch Crossing* (BC) clock is distributed as well but it is sent to the BOC cards (more details about the clocking can be found in the BOC section (3.1.2)). A list of the TTC signals is given in the following:

- clock
  - 40 MHz Bunch Crossing clock: BC
- fast command
  - Level-1 Accept (Trigger): L1A
  - Event Counter Reset: ECR

[1] There are 3 standard form-factors for VMEbus cards: 3U, 6U and 9U. 1U corresponds to a height of 1.75 inch or 4.445 cm.

  - Bunch Counter Reset: BCR
  - Calibrate signal: CAL
- event ID
  - 24-bit Level-1 trigger number: L1ID
  - 12-bit Bunch Crossing number: BCID
  - 8-bit trigger type: TTID

All the above signals can be generated for stand-alone operation on the TIM itself if requested by the RCC, which also writes the configuration registers on the TIM for normal run mode. A picture of this 9U VME board can be found in Appendix B.2.

ROD

A central part in the readout chain of the Pixel Detector is the 9U VME ROD card which is responsible for building event fragments from the data of the front end modules and to pass them the TTC signals (mentioned in the previous section). Secondary tasks are the configuration and calibration of the front-end and online monitoring of the received data. To fulfill these requirements a hybrid configuration of *Field-Programmable Gate Arrays* (FPGA) and *Digital Signal Processors* (DSP) is employed[1] as described in the following and illustrated in Figure 3.2.
Event data from the detector is received via 96[2] serial data lines from the BOC through the VME backplane connectors (direct interconnects, the VME bus is not involved). A set of eight *Formatter FPGAs* is used to convert these serial- into two 32-bit wide parallel data streams. They also de-randomize the fragment data and search it for header and trailer information as well as errors. In this and the following error detection steps information about the errors is written to the header word.
The *Event Fragment Builder FPGA* processes both parallel data streams simultaneously. It checks for- and counts the errors before building and storing the event fragments in two 16k word deep FIFOs[3] until they are complete and ready to be shipped to the ROB. The next step in the chain is the *Router FPGA* which sends the fragment data on its way to the S-Link connection to the ROB and also can be used to trap it for further processing in the *Slave DSPs* (SDSP). There are 4 SDSPs which can run several algorithms on the fragment data to allow on-line error monitoring, histogramming, front-end module calibration and correction of the synchronization independent from the normal data taking.

[1] There are two slightly different version of the ROD card, one for the readout of the SCT and another for the Pixel Detector. They use different firmware to handle the readout, more information about the ROD can be found in its users manual [48]
[2] For the readout of the Pixel Detector only 48 of them are used.
[3] First in first out memory.

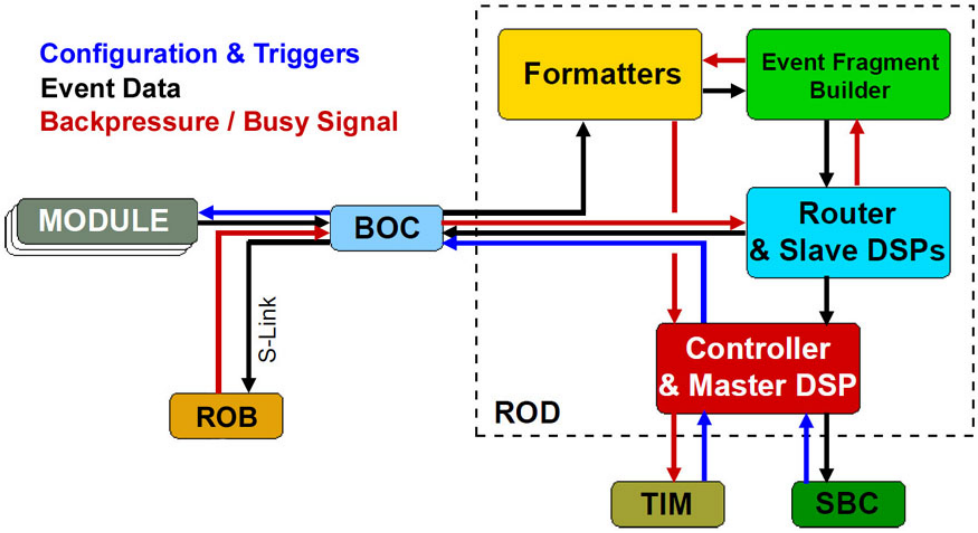

Figure 3.2: Illustration of the data path on the ROD[46].

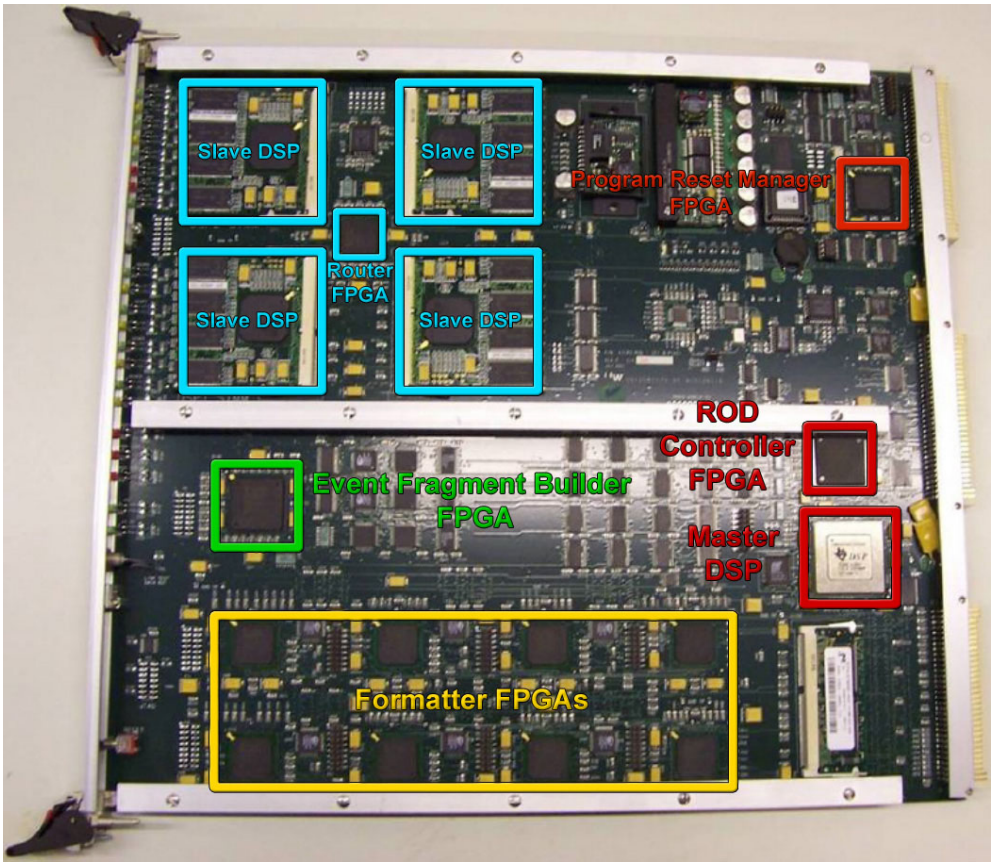

Figure 3.3: Photograph of a ROD card with its main components highlighted[49].

To allow remote control of the ROD card it is connected to the VME bus via the *Program Reset Manager* (PRM) FPGA which provides access to all other peripherals on the board[1]. The PRM also configures and resets the ROD card after power on or reset, this includes the initialization and booting of all the FPGAs and DSPs. The function of the ROD card itself is controlled by a combination of a *Master DSP* (MDSP) which runs software for non real-time operations and a *ROD Controller FPGA* (RCF) performing all functions that need to be done in real-time.
The MDSP manages the operation of the complete ROD card and can be controlled by the user via command registers. But once the ROD card is fully configured for physics data taking it is not needed to perform any crucial functions. These are all performed by the RCF. It has logic blocks to process the TTC information, to send command and trigger signals via the BOC to the front-end electronics and to collect event information for the individual triggers and distribute it to the according FPGAs in the data path (e.g. the Formatter FPGAs). The RCF also interfaces these data path FPGAs, the SDSPs and the BOC to the MDSP allowing for example to transfer calibration data (histograms) via the VME bus to the SBC for further evaluation. But the bandwidth of the VME bus is with 160 MByte/s[2] rather limited as the bus is shared between up to 16 RODs which results in an effective transfer rate of $\approx$ 7 MByte/s per ROD[34].

3.1.2 The BOC

The BOC is used in the readout of two sub-detector systems, the Pixel Detector and the SCT. The PCB layout is the same, but for the *Pixel BOC* not all components are populated in the RX path. This is due to the higher bandwidth of the pixel modules and the overall bandwidth limit of 160 MByte/s of the ROD[3], hence less modules will be readout by a Pixel BOC. An overview of the number of modules connected per Pixel BOC will be given in the RX section.
The BOC provides the interface between the front-end electronics and the off-detector readout components as well as the connection to the DAQ system. It receives the optical data signals from the Optoboards and converts them into electrical data streams which are then passed to the ROD. This requires precise timing and synchronization of the data and an appropriate adjustment of the thresholds in the PiN Diodes used for the conversion. For the downstream to the detector modules the command stream received from the ROD is BPM encoded, and the light power of the laser diodes is regulated to guarantee accurate transmission. All this tasks are performed synchronous to the 40 MHz BC clock which is provided by the TIM board in the crate

[1] A VME to *Host Port Interface* (HPI) bridge is implemented for that.
[2] Bandwidth of the VME64x bus.
[3] This results from the S-Link bandwidth.

and is distributed to the ROD card by the BOC. Additionally, parallel event fragment data streams from the ROD are passed to the readout buffer of the DAQ, and for safety reasons an interlock is provided to switch of the laser diodes.

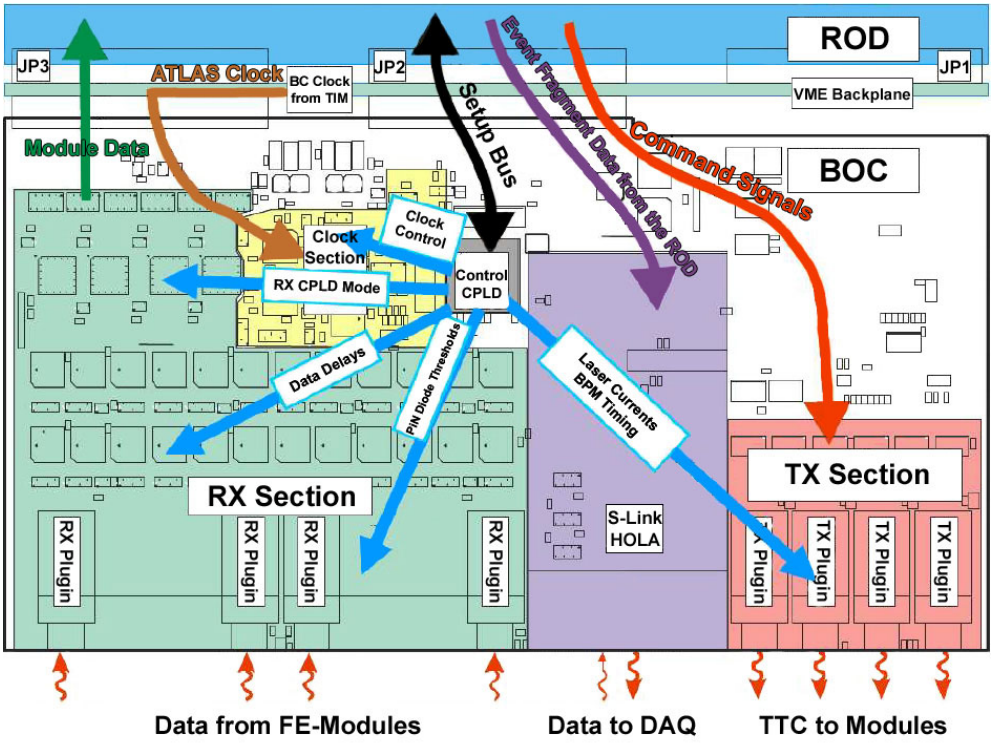

Figure 3.4: Drawing of the Pixel BOC with its individual sections and data paths[30].

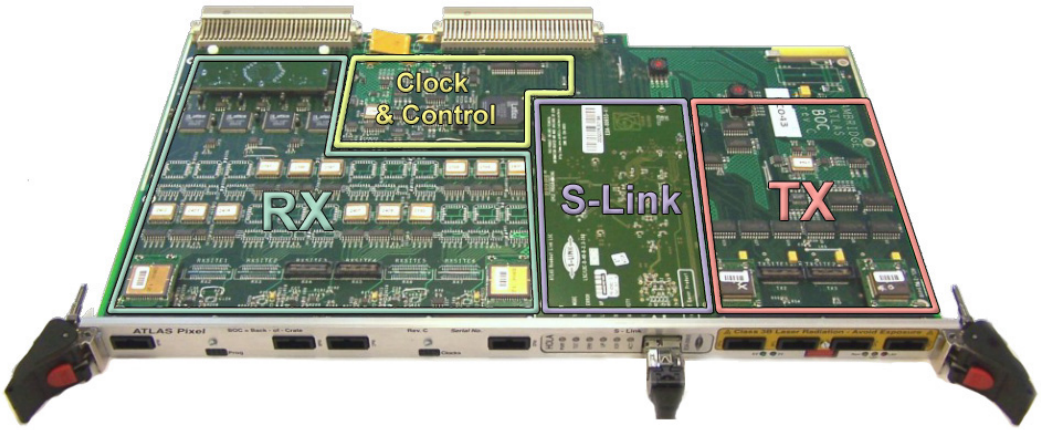

Figure 3.5: Photograph of a Pixel BOC with its sections highlighted[50], equipped with plugins for 80 Mbit/s operation.

All these functions are realized on a 9U VME transition board which is plugged into the back of the crate, hence its name. Contrary to the 9U x 400 mm front-side boards (TIM and ROD) it is 220 mm long, uses only 2 (JP1 & JP2) of the 3 VME backplane connectors and is not connected to the VME bus. Its functionality can be grouped into four sections: RX, TX, Clock & Control and the S-Link. These are illustrated for the Pixel BOC in the Figures 3.4 and 3.5 and will be described in the following.

Clock & Control Section

The 40 MHz BC clock distributed by the TIM in the crate is received via the backplane and duplicated into 6[1] clock signals on the BOC. One is directly sent to the ROD card and another is used to clock the delay elements on the BOC, these are called *ROD-Clock* and *A-Clock* respectively. For the on-board logic the *FP-Clock* is used. The other 3 clock signals have the capability to be delayed to allow adjusting of the phase relationships. For the BPM encoding of the command stream sent to the detector front-end the *P-Clock* is used which can be delayed in 1 ns steps up to a total of 24 ns. The same type of delay can be applied to the *B-Clock* which is used to recover the 40 Mbit/s data streams from the front-end modules. For the sampling of 80 Mbit/s streams a second clock called *V-Clock* is employed, it is the inverse of the B-Clock and has extended delay options of 0 to 48 ns in steps of 1 ns and a fine delay of up to 10.2 ns with a 40 ps granularity. As the period of a 40 MHz clock is 25 ns its full range is covered (more than twice in case of the V-Clock) by these delays.

To control the functionalities of the BOC a *Complex Programmable Logic Device* (CPLD) is connected to the ROD card via the asynchronous Setup Bus with 10 address and 8 data lines. It provides access to the registers of all components on the BOC and controls the *Digital to Analog Converter* (DAC) of the optical plugins. More details will be given when the individual components are described in the according sections.

The temperature of the board can be monitored at two locations by measuring the changing resistance of a NTC thermistor[2] below the RX- and the TX-Plugins. Additionally the voltage and current of the PiN arrays are measured and can be readout like the temperature data by an *Embedded Local Monitor Board* (ELMB)[51] which is housed in the crate as well, and is connected to each BOC in the crate by an 8-wire cable. The ELMB itself interfaces to the *Detector Control System* (DCS) via a *Controller Area Network* (CAN) bus to allow remote supervision from the control room.

As a safety precaution a combination of two interlock signals is used to switch off

[1] These clock signals are multiplexed into several copies if necessary.

[2] The resistance of a thermistor changes significantly with temperature and can thus be used to measure the temperature, it can either increase or decrease, in the latter case it is called *Negative Temperature Coefficient* (NTC)

the lasers on the TX-Plugins if needed. One signal is generated crate internal by the ELMB and by a sensor monitoring the crate door, and the other interlock comes from the front-end interlock system. If one of these signals is set (active low) the lasers are powered down immediately.

TX Section

The TX section is responsible for proper encoding and transmission of the TTC signals to the front-end electronics. These signals are received from the ROD card via 48 data lines over the backplane connector and routed through buffers[1] to the four plugin sites. The corresponding TX-plugins (Figure 3.6) are populated with a custom made *Application Specific Integrated Circuit* (ASIC) called *BPM12*[33], a 12 (8)[2] channel VCSEL array and DACs to regulate the light power of the laser diodes.
The BPM12 chip performs the DC balanced BPM encoding by combining the P-Clock with the data input (on 12 individual channels, hence its name) to an output which has a transition for every rising P-Clock edge and an additional one if the data input is "1". To achieve correctly timed arrival of the TTC signals at the front-end each of the encoded data streams can be delayed in up to 31 course steps of 25 ns[3] (implemented by flip flops clocked with 40 MHz). Additionally its phase can be fine tuned by shifting it up to 127 times by about 280 ps (achieved by a selectable number of inverter pairs in the data path) which allows to cover more than a full clock period with a maximum possible phase shift of $\approx$ 35 ns. As a duty cycle as close to 50 % as possible is needed to reduce the jitter in the clock signal recovered by the DORIC on the Optoboard[33] a *Mark to Space Ratio* (MSR)[4] adjustment circuit is available. It has 32 possible register settings which allow to set the duty cycle to $\approx (50 \pm 10)$ %. For setup purposes after powerup of the front-end chips an inhibit signal can be used which causes the BPM12 to encode a constant "0" which results in a 20 MHz clock[33].
The setting of all these properties is managed by the Control CPLD which writes the appropriate values to the corresponding registers of the BPM12 chip. It also adjusts the DAC values for the laser current, and as the latter are write-only they are also stored in a RAM to be accessible.
For the readout of the Pixel Detector each TX-plugin is connected to an Optoboard

[1] More information on buffering and signal levels will be given in the Signal Standards section.
[2] There was an initial charge with 8 channel VCSEL arrays for the Pixel Detector readout as the 12 channel version is only needed for the SCT readout, but nowadays both sub-detector systems use plugins with 12 channels.
[3] Hence a maximum delay of 775 ns can be achieved.
[4] The MSR is describing the same signal property as the duty cycle, the relation between "high" and "low" times in a periodic signal, but they are determined slightly different: Duty cylce $D = \frac{T_{high}}{T_{period}}$ and MSR $M = \frac{T_{high}}{T_{low}}$, hence a duty cycle of 0.5 (= 50%) corresponds to a MSR of 1.

via 8 fibers[1] from which one is needed to transfer the TTC information to one module. From the possible maximum of 32[2] modules that could be connected only 26 are attached at most due to redundancy and modularity issues.

Figure 3.6: Photograph of a TX-plugin (34 x 21 mm^2) with the prominent BPM12 ASIC and the VCSEL laser array for transmission to the front-end electronics, picture taken from [52].

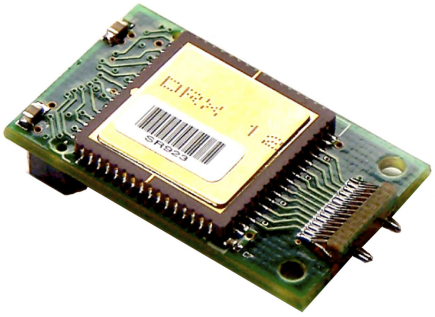

Figure 3.7: Picture of a RX-plugin (35 x 21 mm^2) populated with a DRX12 ASIC and a PiN diode array to receive the module data, picture taken from [52].

RX Section

In the RX section the event data from the front-end modules is received, synchronized and multiplexed into 32 x 40 Mbit/s data streams which are then passed to the formatters on the ROD. The data arrives optically in 8-way[3] fiber ribbons from the front-end and is transformed into electrical signals in a PiN array, which as well as a custom ASIC called *DRX12* is housed on a plugin PCB similar to the TX-plugin (see Figure 3.7). The DRX12 has DC coupled comparator input stages to receive the NRZ data streams from the Optoboards. For each channel an individual threshold can be set for the comparators by applying different reference voltages generated by external DACs. Upon reception the signals are amplified in the DRX12 chip and passed as LVDS streams from the plugins to the BOC. There the differential streams are buffered and converted to single-ended signals which then are fed into 4 channel *PHOS4*[53] delay ASICs (three per plugin). These PHOS4 chips are used for the coarse delay of the clock signals in the clock section as well, and similar to their utilization there they can be used to adjust the phase of the received data by delaying it in steps of 1 ns up to a maximum of 24 ns[4]. After each individual data stream was put into

[1] These 8 fibers are arranged as a ribbon, and 8 of those ribbons are bundled together as a cable. More information on the type of fibers used can be found in [30].

[2] Or 48 for the SCT BOC and the corresponding 12 channel plugins.

[3] From which only 6 or 7 are used as mentioned in the description of the Optoboard.

[4] For the V-Clock a feedback connection of the output of a PHOS4 to the input of another allows to double the delay range to 48 ns.

the correct phase relation with the ROD-Clock it is connected to one of the 4 CPLDs in the data path which will register the data and transmit it clocked with the B-Clock to the ROD. The Control CPLD is responsible to manage delays and thresholds by setting the registers to the appropriate values.
On the Pixel BOC up to four RX Plugins can be attached, and each CPLD will handle the 8 data streams from one plugin. These are then split to a pair of 4 streams and sent with 40 Mbit/s to one of the 8 formatter FPGAs on the ROD (2 (split) x 4 (data streams) x 4 (CPLDs) x 40 Mbit/s = 160 MByte/s). This implies that for modules sending at a higher rate than 40 Mbit/s, which is true for all but the Layer-2 modules, not all four RX Plugins can be used. Hence, for the readout of modules sending at 80 Mbit/s only two RX Plugins will be utilized per BOC which will send their data to one CPLD each. The B-Clock and its inverted copy, the V-Clock, are used to sample this data and register it as two sets of 40 Mbit/s streams from which one is sent to the adjacent so far unused CPLD. Thus, each of the four CPLDs again manages 8 data streams of 40 Mbit/s. A list of the number of BOCs and plugins needed for the readout of the Pixel Detector can be found in the following table 3.1.

Part	**Staves/Sectors** [#]	**Optoboards** [#]	**BOCs** [#]	**TX-/RX-Plugins** [# per BOC]	**Data Rate** [Mbit/s]
B-Layer	22	44 (B)	44	1 / 2	160 (2 x 80)
Layer-1	38	76 (D)	38	2 / 2	80
Layer-2	52	104 (D)	26	4 / 4	40
Disks	48	48 (D)	24	2 / 2	80

Table 3.1: Readout components necessary for the ATLAS Pixel Detector, (B) and (D) denote the type of Optoboard used.

S-Link Section

The task of the S-Link[54] section is to ship the event fragments built on the ROD to the readout buffers of the DAQ system. For this a mezzanine card is plugged onto the BOC and the 32 parallel data lines[1] from the ROD are routed through to its connector without any manipulation besides buffering.
The mezzanine card is a *High-speed Optical Link for ATLAS* (HOLA) interface card which implements the S-Link (See Figure 3.8). The S-Link was specified at CERN to provide a simple link (hence the name) interface between any front-end electronics and readout hardware by just connecting them to a *Link Source-* and *Link Destination* mezzanine *Card* respectively (LSC / LDC). It features simple data movement in one direction with a maximum data rate of 160 MByte/s and includes error detection and self test functions. In a duplex version there is also a small return channel to

[1] There are 7 additional control signals for the S-Link.

send flow control or other commands back to the sender. In the readout of the Pixel Detector only two flags (LDOWN & LFF)[1] are used to communicate back pressure and to pause the data transmission if the ROB is busy.

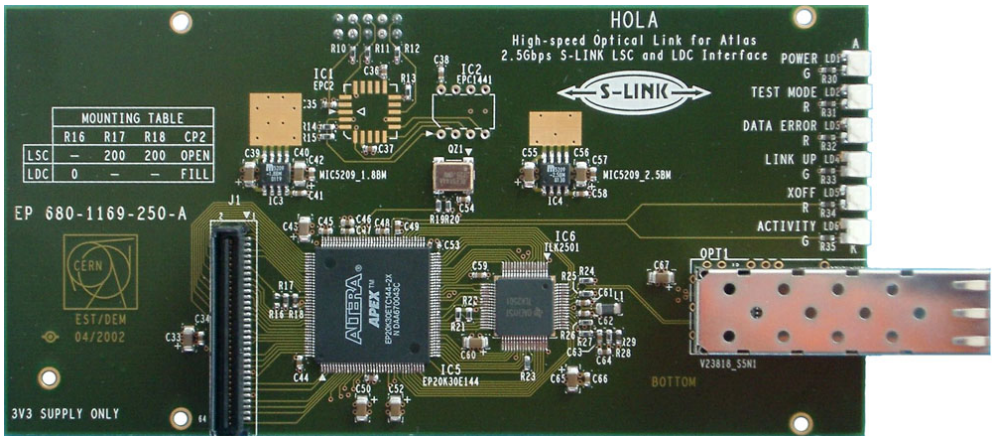

Figure 3.8: Picture of a HOLA mezzanine card for the S-Link connection from BOC to ROB[54].

Power Supply & Signal Standards

The BOC is powered by the VME backplane and is connected to a 3.3 V and 5 V supply. It consumes up to $\approx$ 3 A (3.3 V) and $\approx$ 1 A (5 V) respectively when the firmware is loaded and the plugins are assembled. A 12 V line is also connected to provide a bias voltage for the PiN arrays of the RX-Plugins, its power consumption depends on the light input into these arrays and is $\approx$ 0.1 A at maximum.
For the interface between the BOC and the ROD 3.3 V *Low Voltage Transistor-transistor Logic* (LVTTL) and 5 V *Positive Emitter Coupled Logic* (PECL) signal standards are used. Internally on the BOC more electrical standards like 3.3 V Low Voltage PECL and LVDS are employed, but these do not interfere with the off-BOC communication. A list of the interface signals of the VME connectors of the BOC will be given in the following table 3.2.

[1] Link Down and Link Full.

Signal	Description	Direction	Connector	IO Standard
XC[47:0]	TX commands to FE	ROD -> BOC	P2	LVTTL
RD[95:0]	RX data from modules	ROD <- BOC	P3	LVTTL
SD[7:0]	Setup Bus Data	ROD <-> BOC	P2	LVTTL
SA[9:0]	Setup Bus Address	ROD -> BOC	P2	LVTTL
SWRN	Setup Bus Write	ROD -> BOC	P2	LVTTL
SSTBN	Setup Bus Strobe	ROD -> BOC	P2	LVTTL
SBUSY	Setup Bus Busy	ROD <- BOC	P2	LVTTL
L-UD[31:0]	S-Link Data	ROD -> BOC	P2	LVTTL
L-URESET#	S-Link Reset	ROD -> BOC	P3	LVTTL
L-UWEN#	S-Link Write Enable	ROD -> BOC	P3	LVTTL
L-UTEST#	S-Link Test Mode	ROD -> BOC	P3	LVTTL
L-UCLK	S-Link Clock	ROD -> BOC	P3	LVTTL
L-UCTRL#	S-Link Control	ROD -> BOC	P3	LVTTL
L-ULDOWN#	S-Link Link Failure	ROD <- BOC	P3	LVTTL
L-ULFF#	S-Link Link Full Flag	ROD <- BOC	P3	LVTTL
CLK40+	40 MHz BC Clock	TIM -> BOC	P3	PECL
CLK40-	40 MHz BC Clock	TIM -> BOC	P3	PECL
RCLK40+	40 MHz ROD Clock	ROD <- BOC	P3	PECL
RCLK40-	40 MHz ROD Clock	ROD <- BOC	P3	PECL
TIM_OK	TIM functions OK	TIM -> BOC	P3	LVTTL
ROD_SENSE	Mis-Location Interlock	ROD -> BOC	P3	LVTTL
DET_LAS_EN	Door Open interlock	Backplane <- BOC	P3	LVTTL
BOC_LAS_EN	external interlock	Backplane -> BOC	P3	LVTTL
BOC_OK	BOC functions OK	ROD <- BOC	P2	LVTTL

Table 3.2: Summary of the interface signals of the BOC VME connectors. The LVTTL signals are all single-ended 3.3 V and the 5 V PECL clock signals are differential.

3.2 An ATCA based readout system

3.2.1 The ATCA Crate Specification

An alternative approach for the readout electronics is based on boards for crates designed following the *Advanced Telecommunications Computing Architecture* (ATCA)[55] specifications. An overview of the ATCA crate structure can be seen in Figure 3.9. From the front side of the crate 8U (322 mm) x 280 mm boards are inserted and connected to three backplane connectors in locations called Zone-1 to -3. The connector in Zone-1 establishes the power supply and provides connections for the crate management. Zone-2 is responsible for the data interconnects, up to 5 connectors with 40 signal pairs each can be combined to a total of 200 differential 100 Ω point-to-point connections. There are slight differences in the routing topology of these interconnects and some special lines for clocking and updating are provided, but a thorough description is omitted at this point as it would not add significant information (Specifications can be found in [56]). In Zone-3 custom made connections can be realized, usually this is used to connect to a PCB inserted from the back of the crate which is called *Rear Transition Module* (RTM). Overall this crate structure provides a commercially standardized platform with potential for high speed serial interconnects on the backplane and support for different IO interfaces. Hence, it was used as a substrate to base a readout system on.

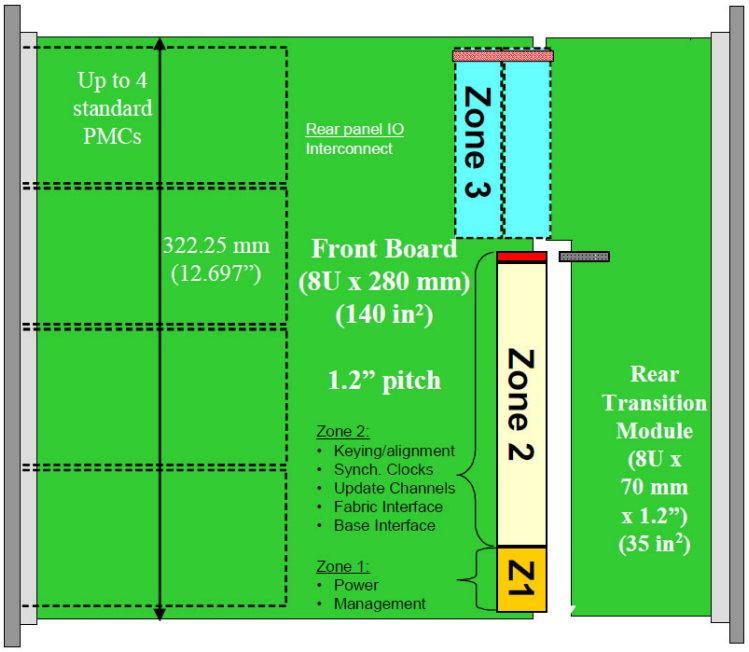

Figure 3.9: Drawing of the ATCA specified crate structure with Front Board and the Rear Transition Module[56]. PMC stands for *PCI Mezzanine Card*.

3.2.2 The SLAC ATCA System

At the *Stanford Linear Accelerator Center* (SLAC)[1] in California (USA) a new generation of DAQ electronics was developed based on generic building blocks with the ATCA system as packaging. As all current *High-Energy Physics* (HEP) experiments have the need for DAQ systems which are capable of dealing with huge amounts of data with very strong time constraints usually custom readout electronics are developed. These have high bandwidth IOs to interconnect the front-end with computational elements, to reduce the readout data, and the rest of the DAQ system. They are often specifically built for only one experiment. So the idea is to have an easily scalable number of cheap (regarding cost, size and power consumption) computational elements as one block, and another block which provides inexpensive interconnects between them with low latency and high bandwidth based on an industrial standard. A combination of these two blocks can then be applied as a readout solution for different experiments and thus no individual custom electronics need to be developed anymore. Only the RTM needs to be adapted to the specific front-end electronics of different experiments by the appropriate optical transceivers.

The Reconfigurable Cluster Element (RCE)

For the computational element an approach based on a Xilinx Virtex-4 FX FPGA as *System on a Chip* (SoC) is employed. This FPGA provides a PowerPC processor core running at 450 MHz with 128 MB configuration memory[2] and 512 MB of Reduced-Latency DRAM (RLDRAM). Further features are 192 DSPs, many combinatoric logic cells and internal block RAM. For high-speed interconnections up to 24 *Multi-Gigabit Transceivers* (MGT) can be used with built-in functions like SERDES (serializer/deserializer), clock and data recovery and several encoding/decoding schemes like 8b/10b for example. By bundling four MGTs together a *10-Gigabit Ethernet* (10GbE) link is created, and two of these are used to connect the RCE to the backplane to establish connections to other RCE-Boards or to the Cluster Interconnect (next section). MGTs are also used to interface to the optical transceivers on the RTM[57]. An RCE-Board with two RCEs under heat sinks and its corresponding RTM as used in the PetaCache project can be seen in Figure 3.10.
The RCE was also built as a mezzanine board to be plugged into an ATCA carrier board (more information in the Cluster On Board section). While the first generation (GEN-1) is equipped with a Virtex-4 FX60 running a single PPC core the second generation (GEN-II) is available as a single-element board with a Virtex-5 FX70T and a dual-element board featuring a Virtex-5 FX130T running either a single- or a dual

[1] Nowadays known as the SLAC National Accelerator Laboratory.
[2] This 128 MB and the 512 MB RAM are external and not included in the FPGA.

550 MHZ PPC core respectively. A SO-DIMM slot is available on the mezzanine card to provide up to 4 GB of DDR-3 RAM and an SD Flash card is used to hold the configuration data. Additional to the 10 Gbit/s Ethernet interface the KR-4 standard with 40 Gbit/s is supported and a TTCrx[58] ASIC was added to all mezzanine RCEs to be able to receive and handle the TTC signals from the LHC system directly without the need of a TIM.

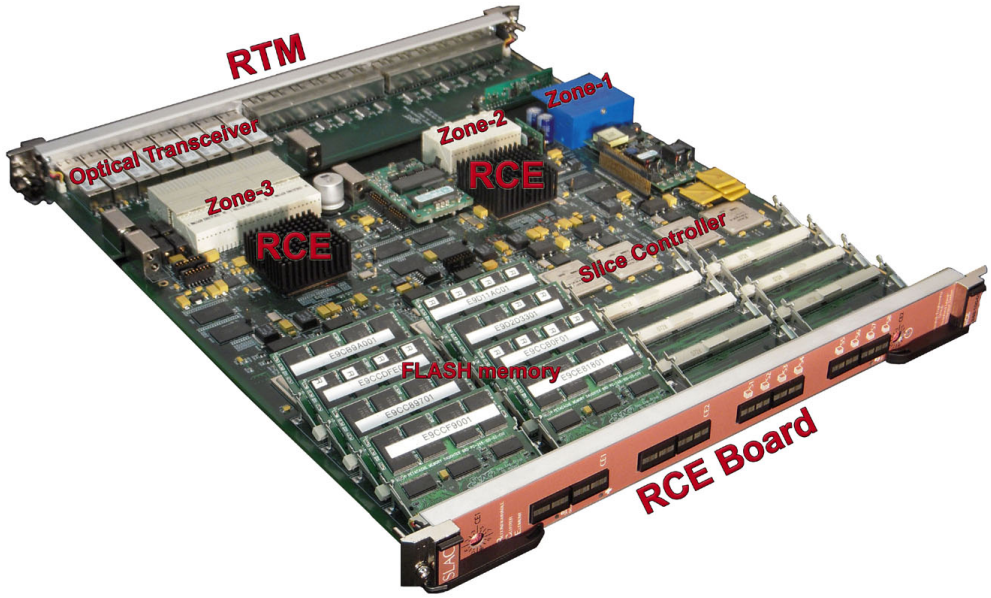

Figure 3.10: Picture of a RCE-Board and its RTM as used in the PetaCache project, Flash memory and Slice Controller FPGAs are specific to this project and otherwise not relevant to the ATCA readout scheme[57].

The Cluster Interconnect (CI)

To connect the computational elements a 10 GbE switch with up to 48 ports is realized on another ATCA PCB. This *Cluster Interconnect* (CI) board is based on two Fulcrum Microsystems FM22xx or FM6xxx ASICS with up to 24 ports each, full layer-2 functionality and several supported physical interfaces including XAUI, KFR, KR-4 and SM-II (100 / 1000 Base-T)[57]. To control the switch, a RCE (Virtex-4) is connected to both switch ASICs by a management bus and is hence called *Switch Management Unit* (SMU). The CI does not only provide connectivity to the backplane in the crate, but also to external systems via GbE and 10-GbE connections on the front plate and especially on the corresponding RTM (see Figure 3.11).

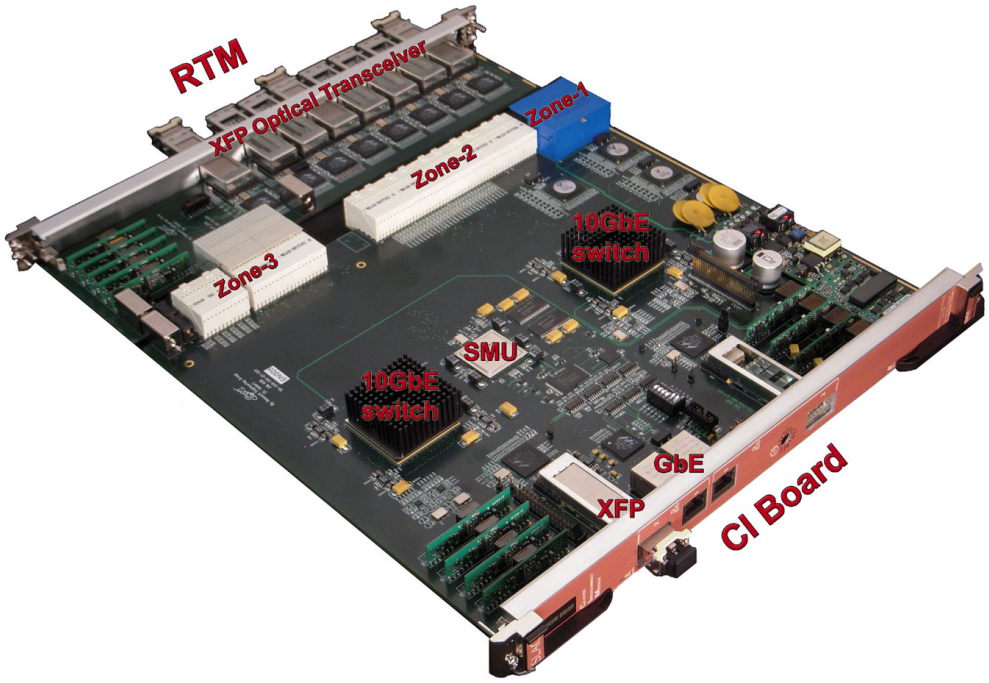

Figure 3.11: Photo of a CI board with its two 10GbE switch ASICS under the heat sinks and a RTM with optical transceivers[57].

The Cluster On Board (COB)

With the development of the already mentioned mezzanine RCEs this system evolved into the *Cluster On Board* (COB) which is a combination of both, the RCE and the CI together on one ATCA board. A schematic view of the COB can be seen in Figure 3.12. Up to four mezzanine RCEs can be plugged onto the board as *Data Processing Modules* (DPM) with either one or two CE cores. They are connected by 40 GbE links to another mezzanine RCE which acts together with a 24-port switch ASIC as the *Data Transportation Module* (DTM) and establishes connectivity to other COBs in the crate and to external systems like the TTC network for example. An *Intelligent Platform Management Controller* (IPMC) is employed to handle the configuration of the board and to monitor its status. A picture of the forth version (v4) of this board can be seen in Figure 3.13, the fifth (v5) which will include some design optimizations (e.g. the SFP+ Ethernet connections will be moved to the front plate) is currently in production.

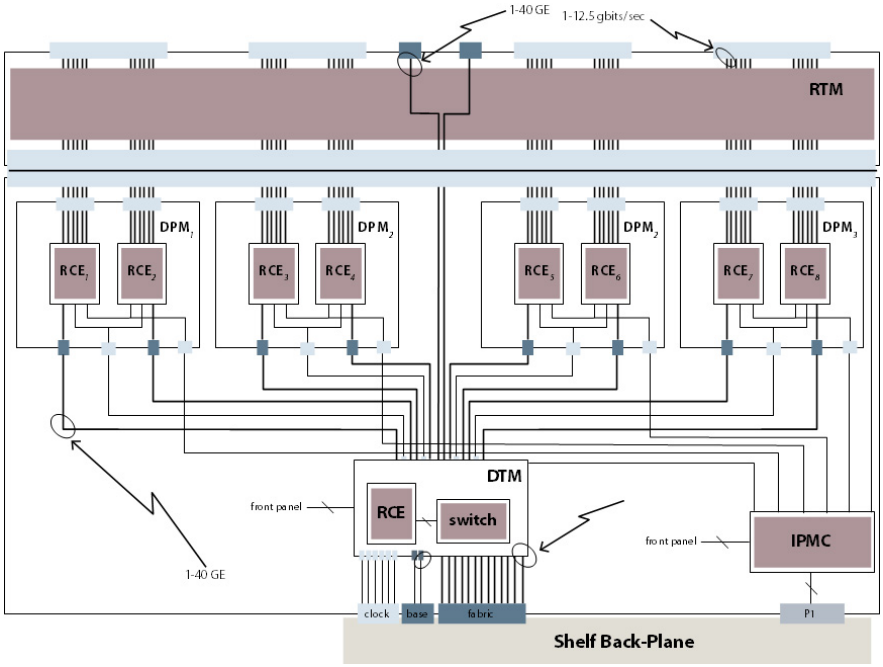

Figure 3.12: Diagram of the structure of the *Cluster On Board* (COB)[57].

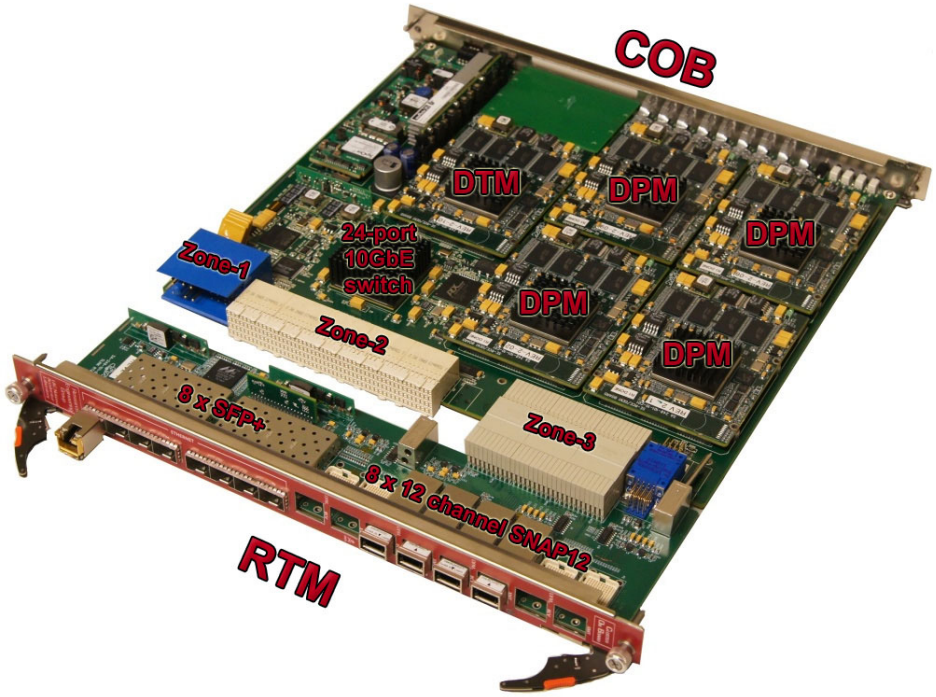

Figure 3.13: Picture of a COB (v4) equipped with five Gen-I mezzanine RCEs and a RTM with 8 SFP+ transceiver and 8 12-channel SNAP12 modules[59].

ATCA for ATLAS

There have been extensive studies on a possible integration of the ATCA readout system into the DAQ system of the ATLAS Detector. Among application use cases for the Pixel Detector and SCT, especially an implementation for the readout of the IBL was evaluated. An overview of this COB based system is given in Figure 3.14. Twelve COBs together with two different flavors of RTMs, managing either the connection to the front-end or to the ROS via S-Link, are housed in one crate and manage the readout of the complete IBL. There is no need for a SBC nor a TIM in the crate and the bottleneck of the VME bus for the configuration and calibration connection is avoided as each COB / RTM pair has two 10 GbE connections. Further plans include an integration of the ROB into the system by providing the necessary buffering and Ethernet connections for the DAQ system on an appropriate COB / RTM pair[57].

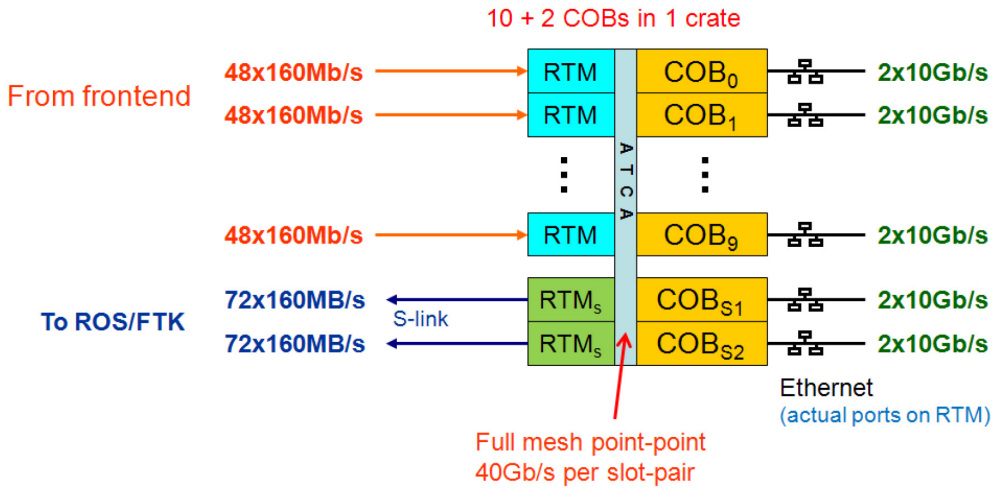

Figure 3.14: Possible IBL readout scheme based on a COB ATCA system[57].

4

IBL Re-BOC - Hardware

The increased bandwidth of 160 Mbit/s of the FE-I4 front-end chips can not be readout by the current VME system as it is designed for a maximum data rate of 80 Mbit/s per channel. Hence, a new readout solution is needed for the IBL, and different options were evaluated. Even though ATCA is an appealing new standard and thus a good candidate for the readout system for the IBL, it was decided by the IBL collaboration to keep the VME standard as a baseline solution. This decision is based on the rather tight schedule for the integration of the IBL into the ATLAS detector and the fact that the ATCA system would not be backwards compatible to the current system at all. Especially the compatibility was regarded as very important as the IBL will only be a part in an existing detector. As a direct conclusion of this issue, the two VME cards that handle the data path (BOC & ROD) need to be redesigned for the IBL and are required to work with both, the old and the new corresponding partner card to guarantee backwards compatibility[34]. The SBC and the TIM will be used as they are in the current system without any modifications.
The redesign of the BOC for the readout of the IBL is the topic of this thesis, hence in the following the focus will be solely on the Re-BOC[1] and the new ROD will only be mentioned briefly.

4.1 Concept

The main task of the Re-BOC will be the same as for the current BOC, it has to establish the interface between front-end electronics, ROD and the off-detector readout system. It has to be capable to achieve this with the new designed front-end components (FE-I4 & new Optoboard) as well as with the old ones (FE-I3 & old Optoboard).

[1] This stands for *redesigned BOC* and will be used from now on in this thesis to denote the new BOC for the IBL.

While also being compatible with the pin assignment on the VME backplane connectors to allow communication with both the current and the redesigned ROD. Hence for the concept of the Re-BOC several constraints have to be met and requirements fulfilled.

- It has to be designed as a 9U x 220 mm VME64x PCB to fit into the current crate structure.
- The pinout on the VME backplane connectors needs to be backwards compatible to the current system regarding the connections to ROD and TIM. This includes the support for the signal standards used at the moment (3.3V LVTTL / 5V PECL).
- The 40 MHz BC clock needs to be received from the TIM, provided to the ROD and distributed on the Re-BOC including enough delay management options.
- For the link to the detector, TTC commands received from the ROD need to be BPM encoded with an appropriate clock and its delay and duty-cycle need to be adjustable[60].
- Module data arriving at 160 Mbit/s has to be received, synchronized, decoded (8b10b) and provided to the ROD. Except the 8b10b decoding this has to work for 40 Mbit/s and 80 Mbit/s data streams as well[60].
- For the event fragment data arriving from the ROD an S-Link interface to the DAQ system is required, additionally a second S-Link connection should be foreseen to send a copy of the event fragment data to a *Fast TracKer* (FTK)[61] system to improve the *High Level Trigger* (HLT).
- Increase the overall integration and bandwidth by a factor of two or four.
- Avoid custom made components like the optical plugins with their ASICs used for data handling or the PHOS4 delay chips, and replace them by *Commercially Of The Shelf* (COTS) available parts[60].

To fulfill all requirements an approach has been chosen that utilizes modern FPGAs together with commercially available optical plugins. The FPGAs support several signal standards with their IO buffers which allows to adapt them to either the current system or the IBL readout. This and the appropriate data path logic can easily be changed by programming the corresponding firmware. To have additional data lines between the BOC and the ROD a so far unused connector between the JP1 and JP2 VME backplane connectors is employed. This so called JP0 connector adds another 95 connections (5 rows of 19 each) to the 2 x 160 links of the JP2 and JP3 connectors from which only 247 are available for data connections as the remaining

73 are needed for power and ground terminals. For the interface between the front-end and the BOC 12-channel optical receiver and transmitter plugins in the SNAP12 form factor will be attached to standard IO pins of the FPGAs driving or receiving the signals. Overall the Re-BOC will feature four times the bandwidth of a current BOC, allowing to read out all 32 FE-I4 chips of an IBL stave at 160 Mbit/s with a single Re-BOC (see Table 4.1 for a list of the readout components, additionally 14 redesigned RODs will be needed to handle the increased throughput from the Re-BOCs).

Part	**Staves** [#]	**FE-I4** [#]	**Optoboards** [#]	**Re-BOCs** [#]	**TX-/RX-SNAP12** [# per Re-BOC]	**RX Rate** [Mbit/s]
IBL	14	448	28 (IBL)	14	2 / 4	160

Table 4.1: Readout components necessary for the IBL for the ATLAS Pixel Detector.

The S-Link will be established by an integrated implementation which will utilize MGTs from the FPGAs to drive the event fragment data via standard *Small Form-factor Pluggable* (SFP) or *Quad SFP* (QSFP) modules to the DAQ system. With eight S-Link connections, there will be twice as many available as needed for the pure data transfer to fully accommodate the requirements for the implementation of the FTK. This way the rather bulky HOLA mezzanine cards can be avoided. Even though there are double link HOLA cards available now that could support the FTK, an implementation of the Re-BOC relying on these cards would still require four of them which would cover most of the area of the PCB and on the front plate and thus leave almost no space available for other crucial components.
A concept drawing of the structure of the Re-BOC prototype can be seen in Figure 4.1. The general composition of the Re-BOC consists of two logically equivalent data paths which are named north and south. The main component of each path is a Spartan-6 LX150T FPGA called *BOC Main FPGA* (BMF) which is responsible for all the data handling. It manages the correct transmission of the TTC signals to the front-end, receives and synchronizes the module data before passing it to the ROD and finally transfers the event fragment data from the ROD to the DAQ system. As optical components each BMF has one SNAP12 transmitter and two receivers attached, an old TX Plugin connector is available as well on the prototype as a fall back option. Additionally four SFP transceivers (north BMF) and a *Quad SFP* (QSFP) (south BMF) module are connected to the MGTs. Each BMF has access to 512 Mbit of DDR2 RAM and a *Low Pin Count FPGA Mezzanine Connector* (LPC-FMC) as an expansion option. Both BMFs are controlled and configured by another FPGA, this Spartan-6 LX75T is therefore known as the *BOC Control FPGA* (BCF). Flash EEPROMs[1] for the configuration data for the Re-BOC, DDR2 RAM and an Ethernet interface are some of the peripheral components attached to the BCF. Also a direct connection to the DCS

[1] Electrically Erasable and Programmable Read Only Memory

is provided by a local ELMB which will monitor the Re-BOC. This rough overview of the Re-BOC concept will be complemented with detailed information in the following sections that will elaborate the individual blocks.

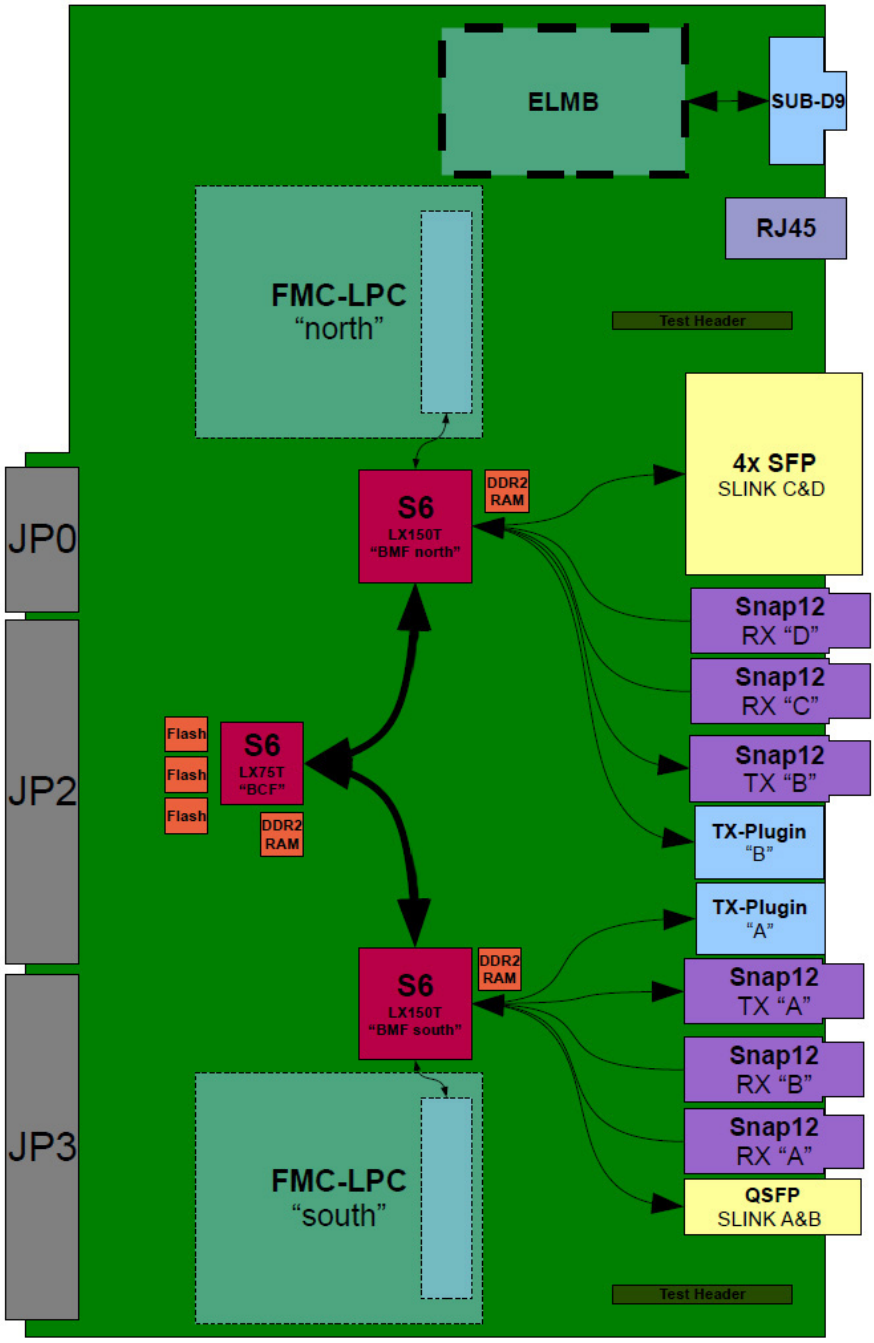

Figure 4.1: Conceptual drawing of the prototype Re-BOC, not all data lines are drawn to provide a clear overview.

4.2 Technology

Contrary to the current BOC the Re-BOC will be based on COTS components rather than custom made ones. In this section the technology of the components used for the redesign will be introduced and described briefly. Most of the functionality of the Re-BOC will be performed by the FPGAs and the appropriate firmware blocks which will be described in the next chapter (Chapter 5).

4.2.1 FPGA

The main work load on the Re-BOC will be handled by FPGAs which offer the performance of custom built hardware without the limitations of ASICs, as they have the capability to change their functional behavior after being soldered to a PCB. For this purpose they consist of an array of rather simple *Configurable Logic Blocks* (CLB)[1] that are connected via a switch matrix to a global routing network in the chip (see Figure 4.2). By loading a configuration file into the FPGA the behavior of the CLBs is chosen and the appropriate connections between them are established. This file is created by programming the desired behavior using a *Hardware Description Language* (HDL) (usually VHDL[2] or Verilog), and then compiling it together with the necessary hardware information in a typically vendor specific development environment.

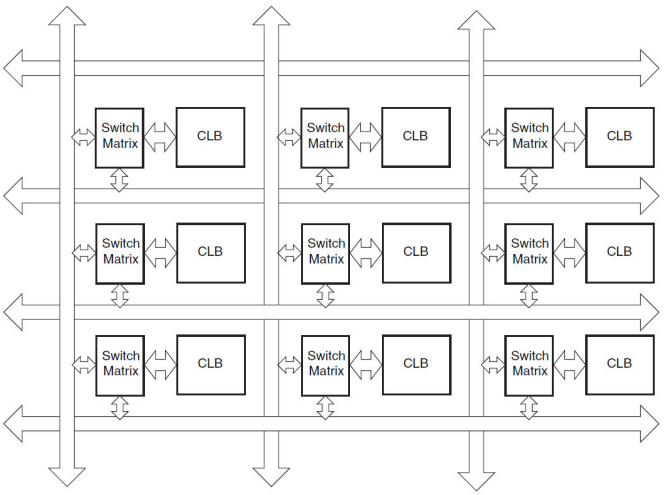

Figure 4.2: Illustration of the array of CLBs with their switching matrices to connect them to the global routing network[62].

[1] They are called like this by Xilinx, the company who invented the first commercial FPGA in 1985. Other vendors use slightly different denotations, but as Xilinx FPGAs are used on the Re-BOC their names will be used in this thesis.

[2] The V stands for VHSIC: Very High Speed integrated Circuit.

Spartan-6 Architecture & Features

On the Re-BOC two different types of Spartan-6[63] FPGAs will be assembled, the LX75T and the LX150T. Therefore the architecture and features of modern FPGAs will be described here on the basis of the Spartan-6 Series[1].
Each CLB of a Spartan-6 consists of two slices from which three different flavors exist. All slices include four 6-input logic blocks based on *Look Up Tables* (LUT) and 8 flip-flops as memory elements. Some slices are more complex and also feature carry logic and wide multiplexers. The most sophisticated slices have all the already mentioned features plus the additional ability to use their LUTs as distributed RAM or shift-register (a diagram of this slice can be found in Appendix B.3).
Besides these basic elements dedicated blocks for special purposes exist to improve their performance. These include slices with DSPs for fast arithmetics and signal processing, block RAM and clock tiles with special routing networks for clock management and distribution. This *Clock Management Tiles* (CMT) include two *Digital Clock Manager* (DCM) which can be used to synthesize different frequencies, remove the clock skew and provide the ability to shift the phase. Each CMT also has a *Phase-Locked Loop* (PLL) to optimize the output deskewing, add a jitter filtering capability and increase the possible range of frequency synthesis.
The IO pins of an FPGA are grouped in banks with individual supply voltages which allows to configure them for different IO signal standards (in Figure 4.3 the layout of the IO banks of the LX150T FPGA is illustrated). As long as the standards are based on the same voltage they can be mixed within the same bank. Some banks have dedicated blocks to allow the easy integration of special interfaces, for example a memory interface to an external DDR[2] DRAM chip or a network connection to an Ethernet-PHY[3].
Two pins each are combined in one IO tile, they both have individual buffers for in- and output which allow not only to choose between IO standards but to vary the output drive strength, slew rate or add on-chip termination to improve the signal quality. They can be either used as two single-ended pins or as one differential connection. Between the *IO Buffers* (IOB) and the FPGA fabric IOLOGIC blocks are available (see Figure 4.4) that can be used to add individual delays to each pin or to integrate a serializer/ de-serializer circuit into the data path.

Other important IO resources available are the high-speed serial transceivers (MGTs[4]) which support several industrial standards with line rates from ≈600 Mbit/s

[1] Other modern FPGAs have a similar architecture and set of features but there are different levels of complexity depending on the field of application.
[2] The Spartan-6 supports DDR, DDR2, DDR3 and LPDDR (Low Power DDR) with a bandwidth of up to 800 Mbit/s.
[3] This chip implements the *physical layer* of the network and is hence called PHY.
[4] These Multi-Gigabit Transceiver were already mentioned in the ATCA RCE section, and are also known as GTP or RocketIO transceiver.

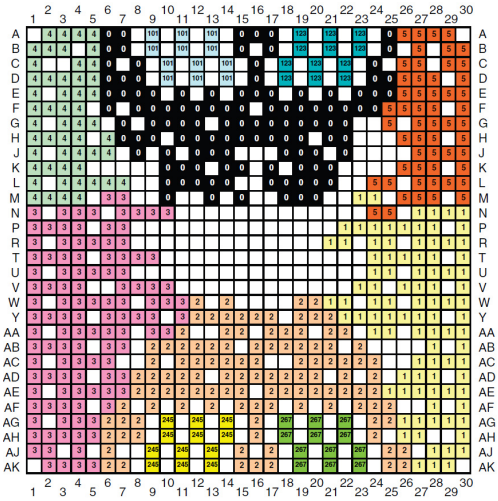

Figure 4.3: Diagram of the IO Banks of the Spartan-6 LX150T in the FG(G)900 package[64]. The white squares are mostly power & ground connections.

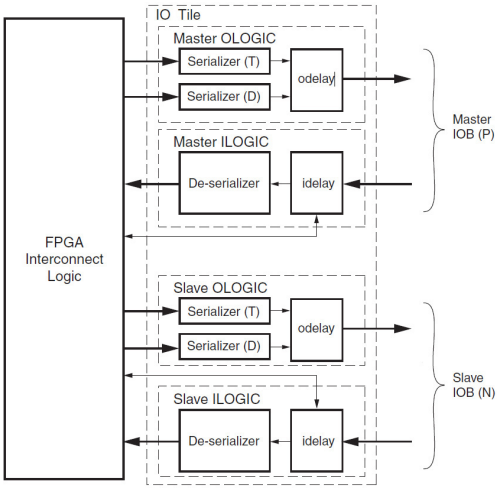

Figure 4.4: Illustration of the resources of an IO tile of a Spartan-6[65], for differential mode the resources of one tile are combined.

up to 3.2 Gbit/s. They provide built-in features like SERDES, clock data recovery including alignment with the help of comma characters and 8b10b de/en-coding. Two transceivers are integrated in one dual-tile together with two PLLs which are responsible to provide an appropriate reference clock for the high-speed part of the transceiver logic.
A summary of the resources of the two types of FPGAs used on the Re-BOC is given in Table 4.2.

Device	Logic Cells [#]	Slices [#]	LUT [#]	Block RAM [kb]	MGTs [#]	Max. IO [# SE pins]
Spartan-6 LX75T	74,637	11,662	46,648	3,096	4*	348
Spartan-6 LX150T	147,443	23,038	92,152	4,824	8	540

Table 4.2: Components of the Spartan-6 FPGAs used on the Re-BOC. *In the FG(G)484 package which is used on the Re-BOC, the FG(G)676 package of the LX75T features 8 MGTs.

There are multiple ways to upload the configuration file into the FPGA, the most common interfaces are JTAG[1] and the *Serial Peripheral Interface* (SPI) which both allow to attach and configure multiple devices with one interface. The typical JTAG setup includes a header on the PCB to which usually a programming cable is connected to transfer the configuration file from a computer or to connect to special cores in the FPGA design that allow to supervise its internal behavior. SPI can be used for example to connect a serial flash EEPROM with the configuration data in

[1] *Joint Test Action Group* (JTAG) is a test and debug interface mainly for PCBs and ICs which is standardized as IEEE 1149.1.

it to the FPGA which will then configure itself with this data after power up. Even though there are parallel interfaces as well, mostly the serial implementations are used and the configuration file is streamed one bit after another into the FPGA. This is reflected in the common use of the terms *bit-file* or *bit-stream* for the configuration data.

FPGA Firmware Programming

Nowadays hardware description languages allow to create complex digital circuits independent from the device the design is targeted for. The languages used to describe the behavior of hardware have similarities to other programming languages used for software regarding the available programming options (e.g. variables, functions, loops...). But there is one big difference, while a software program is executed sequentially hardware works concurrent and includes propagation delays. This and the way how the final implementation with logic gates can be influenced by the coding style needs to be considered during the design process.
After all desired functionality has been programmed the HDL sources need to be synthesized, in this step the written language is translated into blocks of logic gates. Then, these blocks get adapted to the vendor and device specific implementations regarding the CLBs before the design is placed into the FPGA fabric and the interconnects are getting routed. All this information is then written to the configuration file.

4.2.2 Optical Components

To replace the custom made optical plugins used on the current BOC the functionality of their ASICS is moved into the FPGAs, and commercially standardized optical components are attached to establish the connections. For the interface to the front-end transmitter and receiver modules compliant to the SNAP12 *Multi-Source Agreement* (MSA)[66] will be used as already mentioned. They offer 12 independent channels for data rates of up to 2.7[1] Gbit/s per channel with a maximum distance of 600 m depending on which fibers and data rate is used.
The transmitter module has 12 differential and internally terminated inputs compatible with both LVPECL and *Current Mode Logic* (CML) to convert the incoming electrical signals via laser driving circuits and a VCSEL array into optical data streams. Control signals compatible to LVCMOS[2] are available to enable, disable or reset the module, and an output indicates if there is a problem. A two wire serial interface (also known as I^2C-bus[3]) can be used to access an internal control block which is

[1] Some vendors build SNAP12 modules that exceed these specifications and can handle up to 3.5 Gbit/s
[2] Low Voltage Complementary Metal Oxide Semiconductor
[3] Inter-IC-bus[67]

responsible for the operation of the module by setting channel properties like the modulation and biasing of the laser driver.
The receiver module has a PiN diode array to convert the optical signals into electrical currents which are then amplified and transformed to voltages by an *Trans-Impedance Amplifier* (TIA) stage before they are driven differentially to the CML outputs. Similar to the transmitter, LVCMOS control signals are available to enable the receiver or the squelch[1] option. A status output and an I^2C interface complement the connections. Block diagrams of the internal structure of transmitter and receiver can be seen in Figure 4.5. Both types of modules operate with a supply voltage of 3.3 V and are connected to the PCB by plugging them from above into a 10 x 10 *Ball Grid Array* (BGA) connector, hence they can not be replaced while the Re-BOC is still in the crate.
Important for the application of these modules in the IBL readout is their capability to handle the rather low data rates needed, 40 Mbit/s for the transmitter and 160 Mbit/s for the receiver[2]. This issue will be analyzed in section 6.2.1 in the chapter about the testing of the Re-BOC prototype.

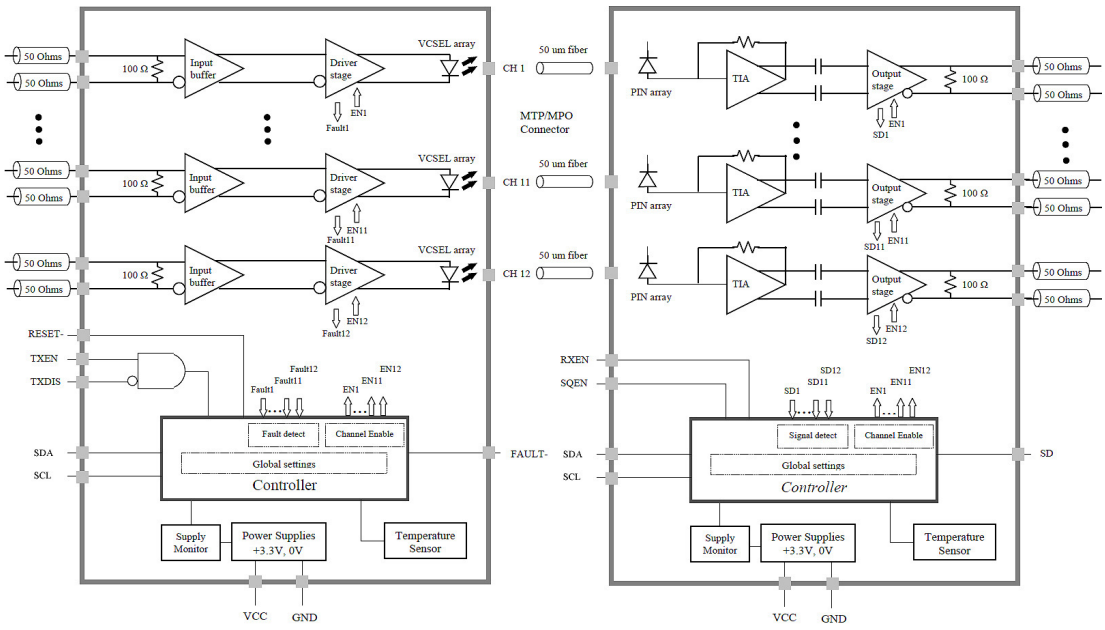

Figure 4.5: Block diagrams of SNAP12 transmitter (left) and receiver (right) from Reflexphotonics[68].

For the S-Link connection one QSFP[69] and four SFP[70] transceiver modules are available. These form factors are widely used and offer a broad range of available line

[1] If this option is enabled the outputs of channels that receive no- or not powerful enough optical signals are driven low.
[2] And down to 40 Mbit/s for the compatibility mode.

rates of up to 10 Gbit/s. Even though they are full-duplex devices the back channel from ROS to BOC will not be used. The HOLA mezzanine cards which establish the S-Link on the current BOCs are equipped with SFP modules as well, making them the natural choice, but to allow the evaluation of a possible replacement by the more compact QSFP transceivers a mix of both was implemented.
They both have CML compatible differential IOs for the data lines and LVTTL control and diagnostic signals. They run with 3.3 V supply voltage and are hot-pluggable in situ as they are inserted from the side and not from above as the SNAP12s.
In Figure 4.6 all optical plugins used on the Re-BOC prototype can be seen.

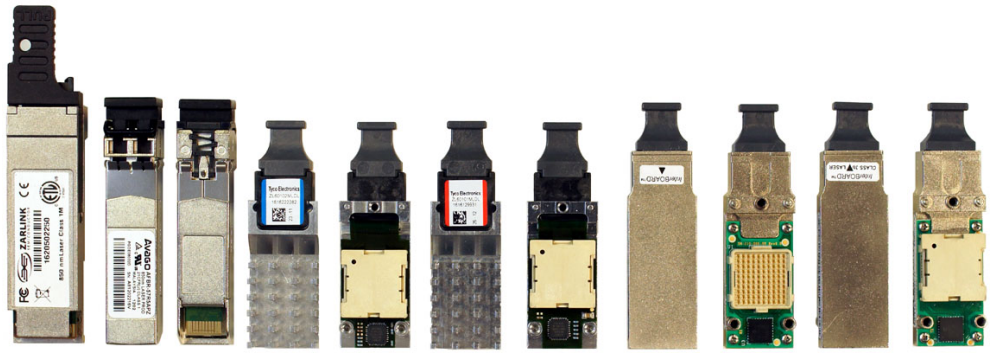

Figure 4.6: Photograph of the optical plugins used on the Re-BOC prototype. From left to right: QSFP, SFP, SNAP12 RX & TX (Tyco) and SNAP12 RX & TX (Reflex Photonics).

4.2.3 Clock Components

On the current BOC a complex combination of ICs is applied to receive the 40 MHz BC clock from the TIM and duplicate it for the different use cases mentioned in section 3.1.2. These ICs include a PLL, a clock divider, a CPLD, several delay ICs and multiple multiplexers. To interface between the different signal standards necessary for these components (5 V PECL, 3.3 V LVPECL, 3.3 V LVDS, 3.3 V LVCMOS) a series of buffers and driver circuits are needed which further complicates things (in Appendix B.4 the schematic of the clock circuits on the current BOC can be seen).
For the clock distribution and handling on the Re-BOC an approach with reduced complexity was chosen. It makes use of advanced timing options provided by the Spartan-6 FPGAs and only requires a minimum of additional components.
To ensure compatibility with both the TIM and ROD (old and new) level translation buffers[1] convert the 5 V PECL input clock from the TIM to 3.3 V LVPECL and vice versa for the ROD clock.

[1] To receive the TIM clock a MC100LVEL92 device from ON Semiconductor is used and to transmit the clock to the ROD a MC100EP16TDG driver from the same company is employed.

Two cascaded programmable delay chips[1] are used to synchronize the phase of the BOC and ROD clock with the TTC clock. Their internal structure consists of a chain of multiplexers causing a fixed minimum delay of 2.2 ns and allowing to add up to 10.2 ns with a resolution of 10 ps which results in a total possible delay range from 4.4 ns to 24.4 ns for the combination of the two ICs. Including the clock management options of the FPGAs (e.g. inversion (180° phase shift) of the clock) a more than sufficient delay range is available for the phase adjustment.
After the delay chips a clock generator and jitter cleaner IC[2] is employed to improve the quality of the clock signal. It is capable of providing clocks with less than 1 ps RMS jitter[71] at its two differential outputs, from which one is used to drive the clock signal to the ROD (via the already mentioned translation buffer), and the other is connected to an 1 to 4 fanout buffer[3] supplying all three FPGAs on the Re-BOC. In case no TTC clock is received from the TIM, the clock generator IC can utilize a local 40 MHz quartz which is connected to its auxiliary input to provide a substitute clock. A picture of the clock components on the PCB can be seen in Figure 4.7. To manage the TTC clock circuit an 11-bit wide bus connection from the delay chips to the BCF is available which allows to set the appropriate delay, while the clock generator IC is controlled via an SPI link (to the BCF as well).
Additional clock elements on the Re-BOC are needed for the Ethernet interface and to provide the FPGAs with a clock for the logical design and reference clocks which are required to operate the MGTs. Therefore a 25 MHz quartz crystal is connected to the Ethernet PHY and five dual frequency crystal oscillators[4] with 100 MHz and 125 MHz are available for the clocking of the FPGAs. The clock for the internal FPGA logic is fixed to 100 MHz but the other frequencies can be set by the BCF which enables the MGTs to be operated at line rates of 1.25 Gbit/s, 2.0 Gbit/s[5], 2.5 Gbit/s or 3.125 Gbit/s. Fanout buffers[6] are applied to duplicate these clocks as each FPGA needs at least one reference clock per MGT dual-tile and a design clock.

4.2.4 Power Supply

Advanced fabrication processes allow to create ICs with shrinking feature size which reduces the leakage currents in the semiconductor material and thus requires lower supply voltages for the operation. The Spartan-6 for example which is manufactured in a 45 nm process needs a core voltage of only 1.2 V. But the VME crate provides only 3.3 V, 5 V and 12 V, hence local step-down voltage converters are necessary to

[1] MC100EP195B from OC Semiconductor.
[2] Texas Instruments CDCE62002
[3] ON Semiconductor NB6N14S
[4] Silicon Labs Si532
[5] This is the default line rate for the S-Link.
[6] The same NB6N14S from On Semiconductors as used for the TTC clock.

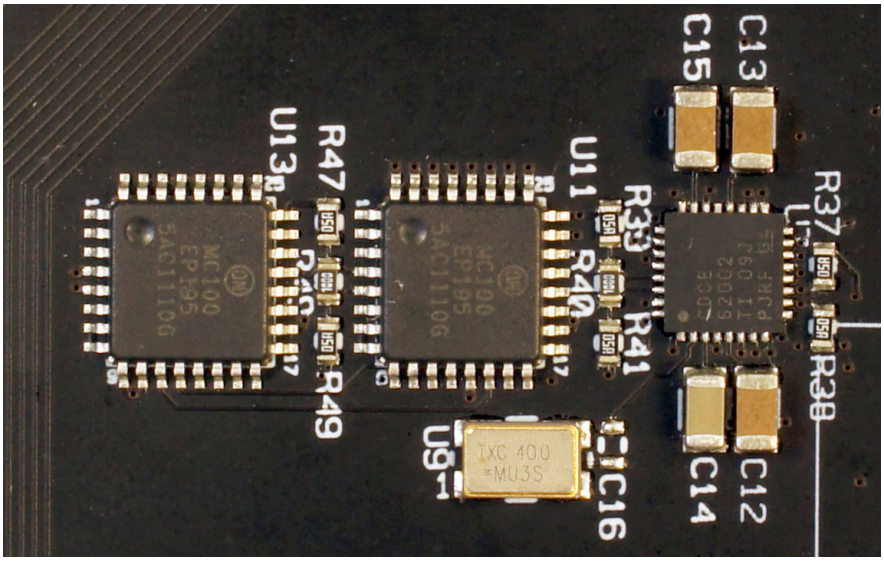

Figure 4.7: Photograph of the cascaded clock delay chips, the jitter cleaner and the 40 MHz quartz on the Re-BOC. The level translation buffers are placed close to the VME connector and the fanout buffer for the on-BOC clocks is on the backside of the PCB.

power all components on the Re-BOC with the appropriate supply voltage. A list of the required voltages which are not covered by the crate supply is given in Table 4.3.

Voltage [V]	Component
0.9	DDR2 RAM Reference (SSTL-18) / Termination
1.2	Spartan-6 Core
1.2	Spartan-6 MGT
1.5	SSTL-3 Reference
1.8	Ethernet PHY
1.8	DDR2 RAM (SSTL-18)
2.5	FMC Mezzanine Connector (incl. the corresponding FPGA banks)

Table 4.3: Components and their voltages that are not provided by the VME crate power supply.

To generate both the 1.8 V and the 2.5 V supplies, highly efficient (up to 95%) power modules[1] from Texas Instruments (PTH05050W) are applied that convert a 5 V input to the required voltages and are able to drive a maximum current of 6 A. For the 1.2 V core voltage supply of the three Spartan-6 devices a similar module is used, the PTH05010W from the same company also generates the output voltage from a 5 V input (up to 96% efficiency), but is able to provide an output current of up to 15 A. To allow the powering of the complete Re-BOC from a single 5 V supply another PTH05010W module is available which is set to an output voltage of 3.3 V. For this

[1] Three of these modules are used.

mode of operation its output needs to be connected to the 3.3 V net of the PCB by populating a 0 Ω resistor.
Noise on the power supply of the MGTs can cause a decrease in their performance, and thus the peak-to-peak input voltage ripple[1] should not exceed 10 mV[72]. Therefore a LMZ10504 power module from National Semiconductor is applied to generate the 1.2 V for the MGTs from a 5 V input with a low output ripple voltage of less than 10 mV and an efficiency of up to 96%.
For the 0.9 V reference and termination voltages of the DDR2 memory chips and the corresponding FPGA banks which establish the memory interface a linear regulator (LP2998MR) from National Semiconductor is used. This device is specifically designed to provide reference and termination voltages for SSTL-18 and SSTL-2 use cases especially in DDR2 and DDR memory applications. To generate the 1.5 V reference voltage for the SSTL-3 signal standard at first a very precise 3 V voltage is created by a REF3330 device from Texas Instruments[2]. This low dropout voltage reference driver utilizes an input of 5 V and provides 3 V with an accuracy of ± 0.15% and less than 100 μV ripple noise[73]. This voltage is then split in half by a resistive voltage divider with two 1k Ω resistors in series.
A standard 20-pin ATX computer power supply connector is available to power the board without the need of a VME crate. This allows easy testing and evaluation of the Re-BOC prototype "on the table" with good access to all components and test headers avoiding the rather cramped and restrictive environment in a crate.

[1] This is the unwanted fluctuation of the voltage around its nominal value.
[2] This component was developed by Burr Brown Products but this company was bought by Texas Instruments in the year 2000.

4.3 PCB Design Process

In this section the process will be described in which the conceptual ideas are at first transformed into circuit schematics and then the actual PCB layout is created. For both sub-processes, the schematic drawing and the PCB layout, brief introductions into design methods and rules will be given. These design tasks were performed in the *Mentor Graphics Expedition Enterprise Flow* development environment which provides all the necessary tools.

4.3.1 Circuit Schematics

A schematic is the simplified graphical representation of an electrical circuit in which the components are depicted as symbols. For simple passive or logic parts standardized symbols are commonly used while more complex components are represented mostly by "black box"[1] rectangles. In Figure 4.8 a part of the Re-BOC schematics is shown to illustrate this. Electrical connections are drawn as lines between the corresponding pins of the components, and if more than one sheet is needed for the schematic, special cross reference symbols represent the connections between the sheets. More details on how these schematics translate into actual circuits on a PCB will be given in the Re-BOC Schematics section.

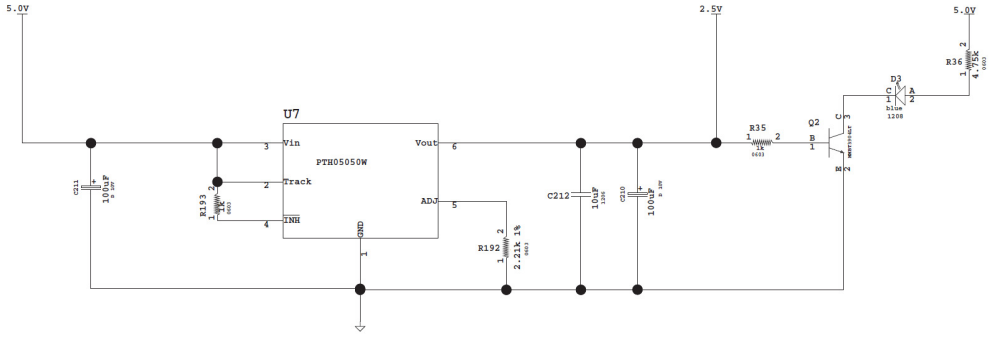

Figure 4.8: Detail of the Re-BOC schematics showing a power supply module, its passive components and a status LED switched by a transistor.

Electrical Circuit Design Basics

When a digital circuit design is planned it is not sufficient to just put all the functional components (e.g. FPGAs, ICs, connectors...) on the schematic and connect them. The analog properties of the electrical signals need to be considered as well

[1] Showing the in- and outputs but not the internal functions.

to ensure good signal quality and a good performance of the circuit. Especially important are the provision of a stable power supply for the active components and a proper termination of the signals between them. Both concepts will be introduced briefly in the following.

Stable Power Supply

A power supply does not only distribute energy to all connected components to allow them to operate, but needs to provide them with stable reference voltages as well. These are required for the digital signal transmissions as the inputs of electrical components compare a received signal to internal references to determine if it is a logic high or low signal. If the reference is noisy it decreases the margin in which a correct decision can be made. Noise on the supply voltage originates from both the power and the ground connections and is often caused by return currents acting on the inductance of the wiring and the power supply. Hence the impedance of the power and ground connections to the components should be as low as possible and there should be low impedance connections between power and ground[74]. A solution to the issue with the component connections is to use *solid copper planes* for power and ground. They are usually situated close to each other and thus have a lot of capacitance which yields a very low impedance at high frequencies. To add the low impedance connections between power and ground *bypass capacitors* are placed close to the components to connect the two planes. To suppress various noise frequencies capacitors with different values are employed, small values like 0.001 μF are used for high-, 0.1 μF for middle- and values of more than 4.7 μF for low frequencies. Hence a combination of two or three different capacitor values can be used to cover a wide noise bandwidth.
Some components like clock drivers produce a lot of high frequent switching noise which needs to be kept from spreading to the power net on the PCB. This is achieved by *decoupling* the device by placing a ferrite bead between the power supply net of the board and the input of the component. These ferrite beads act together with a capacitor as a passive low-pass filter. It has an almost pure inductive behavior for low frequencies, but turns highly resistive at higher ones and thus dissipates the noise as heat.
Both these concepts can be seen in Figure 4.9 which shows the power supply filtering for the PLL used for the TTC Clock on the Re-BOC. It has a low-pass decoupling filter consisting of a ferrite bead combined with a 10 μF capacitor and another 10 μF capacitor which works in conjunction with several 1 μF and 0.1 μF capacitors as bypass to provide noise filtering over a wide frequency range.

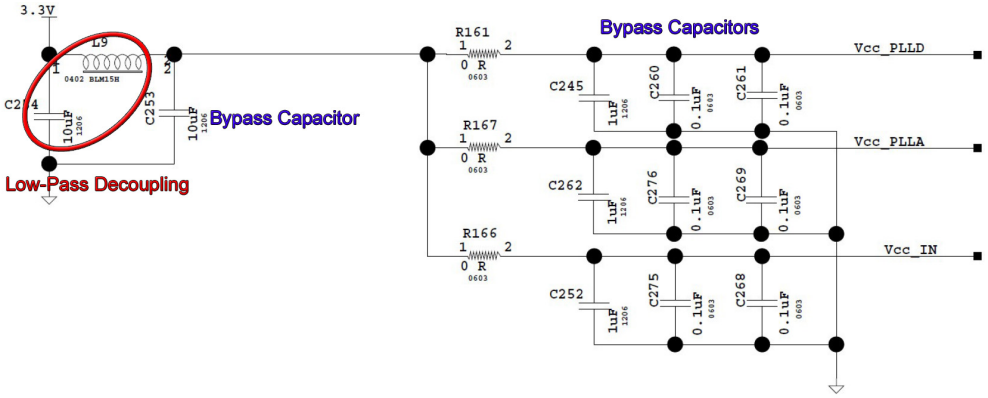

Figure 4.9: Power supply filtering circuit for the PLL handling the TTC Clock on the Re-BOC with low-pass decoupling and bypass capacitors for a wide frequency range (the 0 Ω resistors were placed in case additional decoupling would be needed for the individual voltages of the PLL, then they would be replaced by ferrite beads).

Signal Termination

The lines used to connect parts on the schematic translate to transmission lines on the PCB with a certain impedance $Z_0(\omega)$. Signals traversing this line get attenuated[1], and at its end a part $T(\omega)$ which depends on the load impedance $Z_L(\omega)$ emerges into the connected device. A fraction $R_1(\omega)$ is reflected and propagates back to the source where a part $R_2(\omega)$, depending on the source impedance $Z_S(\omega)$ is reflected again. Each reflected signal gets attenuated again and adds to the emerging signal at the destination, but to simplify things only $T(\omega)$, $R_1(\omega)$ and $R_2(\omega)$ will be considered here. A detailed examination of this phenomenon can be found in [74]. The transmitted and reflected fractions are given by the following equations.

$$T(\omega) = \frac{2Z_L(\omega)}{Z_L(\omega) + Z_0(\omega)} \tag{4.1}$$

$$R_1(\omega) = \frac{Z_L(\omega) - Z_0(\omega)}{Z_L(\omega) + Z_0(\omega)} \tag{4.2}$$

$$R_2(\omega) = \frac{Z_S(\omega) - Z_0(\omega)}{Z_S(\omega) + Z_0(\omega)} \tag{4.3}$$

To avoid these reflections and the resulting signal degradation two termination methods are commonly used, the *end termination* and the *source termination*.
In the end termination the first reflection is suppressed by setting the load impedance to the same value as the transmission line impedance, with $Z_L = Z_0$ Equation 4.1 yields $T = 1$ and Equation 4.2 $R_1 = 0$ which is exactly what is wanted. The source

[1] By the propagation function $H_X(\omega) = e^{-X[(R(\omega)+j\omega L)(j\omega C)]^{\frac{1}{2}}}$ [74].

termination works similar, as the second reflection can be eliminated by a source impedance equal to the transmission line impedance (resulting in R_2 = 0 in Equation 4.3). In this case the signal intensity which propagates the transmission line is halved at the source, but the reflection at its far end (assumed there is no additional end termination) is also of this half intensity and both add up to a full signal level[74]. This reflection is then suppressed by the source termination.

Several techniques can be utilized to terminate a transmission line. Most common are *series* termination on the source side and *parallel*, *split* and *differential* termination on the receiver side, but also combinations thereof are possible. For the series termination a resistor with a value equal to the characteristic impedance of the transmission line minus the output impedance of the driver is used as can be seen in part (a) of Figure 4.10. The voltage V_{TT} to which the resistor is connected for the parallel termination (4.10.b) can vary between 0 V[1] (*pull-down*) and the supply voltage (*pull-up*). The pull-down option requires a strong driver circuit to be able to reach the minimum necessary voltage V_H for a logic 1 as a current of $I_H = \frac{V_H - V_{TT}}{R_P}$ is flowing through the termination resistor, the logic 0 is no issue as the pull-down does the job for the driver. E.g. to drive a 3.3 V LVCMOS signal to a V_H of 2 V with a 50 Ω termination resistor a current of $\frac{2-0}{50}$ A = 40 mA needs to be supplied and the Spartan-6 for example has a maximum drive strength of only 24 mA (which is actually quite a lot). In the pull-up implementation the driver needs to sink[2] a current of $I_L = \frac{V_{CC} - V_L}{R_P}$ to reach at least the maximum voltage V_L recognized as logic 0 while the logic 1 is established by the pull-up. In the 3.3 V LVCMOS example from before a current of $\frac{3.3-0.8}{50}$ A = 50 mA needs to be sunk which is too much for most drivers (24 mA for a Spartan-6).

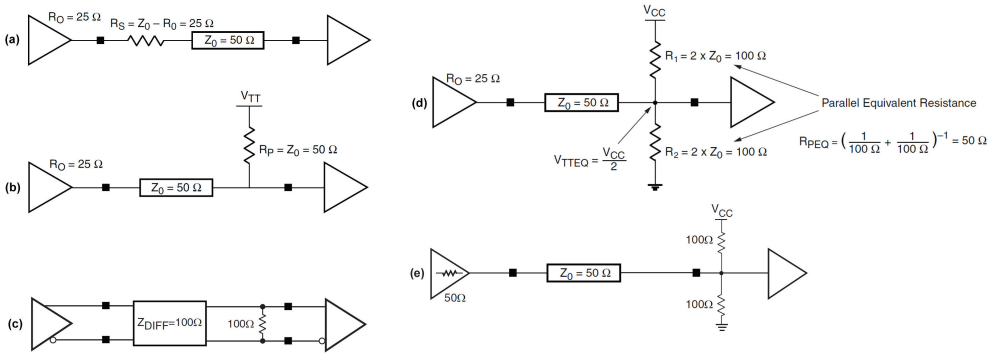

Figure 4.10: Graphical illustration of different transmission line termination techniques[75]. a) Series Source- b) Parallel Receiver-, c) Differential Receiver-, d) Split Receiver and e) On-Chip Series- and Split Termination.

[1] Or more general V_{EE} as negative voltages are also used by some signal standards.
[2] This can vary between different signal standards and driver implementations.

To solve this issue an adequate termination voltage between V_{EE} and V_{CC} can be used which raises the signal level to a point high enough to allows the driver to pull above V_H and still low enough to sink below V_L. If such a voltage is not available it can be provided by a split termination (Figure 4.10.d) which can be transformed into the Thevenin equivalent voltage source $V_{TTEQ} = \frac{R_1 V_{EE} + R_2 V_{CC}}{R_1 + R_2}$ with an equivalent impedance equal to the transmission line $Z_0 = (\frac{1}{R_1} + \frac{1}{R_2})^{-1}$. Differential transmission lines are terminated parallel as well, but the resistor matching the differential impedance is connecting both lines at the receiving end and does not pull to an external voltage (Figure 4.10.c).
Modern FPGAs like the Spartan-6 provide the user with *On Chip Termination* as an option that can make the external termination of transmission lines obsolete. The output buffers feature a *Programmable Output Driver Impedance* which can be used as source termination, and the inputs have *Programmable Input Termination Resistors* that can implement split- and differential termination (Figure 4.10.e shows on chip source- and split termination). But care needs to be taken as the current flowing through the termination resistors is generating heat that needs to be dissipated, in Section 6.2.2 consequences thereof will be presented.

Simultaneously Switching Outputs

Another issue induced by current flow is *ground bounce*. When the output driver switches from high to low, current flows through it to the ground connection (usually a ground plane on a PCB), and if many outputs in close vicinity switch in the same direction the cumulative current can cause[1] the voltage level of the local ground to raise to a level that the output transistor switches off. When the charge is dissipated in the ground plane the transistor switches back on, and the same effect might occur again causing the output to bounce. The same can happen with the supply voltage sagging locally by driving too much outputs high simultaneously causing *power bounce*. Even though the levels might not get shifted enough to force the transistor to switch, noise is induced and the margins for proper signal quality are decreased. These issues are especially critical for highly integrated ICs like FPGAs which have many IOs close together. To avoid these issues and to guarantee an acceptable noise level the data sheets of these devices give a maximum number of *Simultaneously Switching Outputs* (SSO) depending on the signal standard and the drive strength used. Obviously standards with lower voltages and low drive strengths cause less currents to flow and can thus switch more outputs at the same time.

[1] This is due to the inductance of the connection from device to PCB, e.g. package bond wire, trace and PCB impedance.

The Re-BOC Schematics

The schematics drawn in a modern development environment are not only flat graphical representations of the electric circuit, but also include the physical properties of the devices and parts intended to be used on the final PCB. Also all connections and nets have distinct identifications, and differential lines are associated with their corresponding partner line. For the design of the Re-BOC schematics the *Mentor Graphics DxDesigner*[1] was used, but to be able to start drawing the circuits the necessary components need to be created for the part library first.

Creating Parts for the library

Each *part* in the library consists of two components, the *symbol* representing the logic view as it will be seen on the schematic and the *cell* which corresponds to the physical geometry of the part.
The symbol is drawn in the *symbol editor*. Its most important components are the *pins* which represent its available connections. These pins have an individual name and a pin number associated while the name can be chosen freely (but of course reasonable names should be used that indicate the nature of the pin e.g. GND for a ground connection). The pin number has to correspond to the one of the physical package of the component to which the symbol is assigned. The form of the symbol has no influence on its function, but usually (as already mentioned) standardized symbols for passive components (capacitors, resistors...) are used, and more complex devices are represented by rectangles. Parts with a lot of pins are often split into multiples of these rectangles which can be placed on different schematic sheets to separate them by function or to make the schematic more clearly arranged.
In the *cell editor* the physical *foot print* (also known as *land pattern*) of a part is created which includes all information that is needed to connect it to the PCB. Most important are the *pads* which are the contacts to which the pins of the component will be soldered to. These can be edited in the *pad-stack editor* in which form and size of the pad (the copper on the PCB to which the pin will be soldered) and the *solder mask* (opening in the *solder resist*[2] to keep the pad accessible, hence usually a little bit bigger than the pad) can be adjusted. For *vias*[3] and components with pins that need to be stuck through the PCB the diameter of the necessary hole and plane clearances[4] can be set.

[1] This is a part of the Mentor Graphics Expedition Enterprise Flow environment.
[2] This is a non conductive varnish with which the PCB is coated on top and bottom side to prevent short circuits and to protect the copper traces.
[3] Vertical Interconnect Access (VIA) is a hole in the PCB that is plated with copper which allows to connect transmission lines on different layers of the PCB.
[4] When traversing a solid copper plane (e.g. a ground or power plane) around the hole of the via a certain radius is kept free from the copper of the plane to avoid short circuits and unwanted influences on the signal.

In the data sheets of electric components usually recommendations for the size of the pads and their arrangement are given. When all pads are placed and numbered accordingly in the cell editor, additional properties can be added that are useful during the layout of the PCB (e.g. *placement outlines* to mark the space required for the component to avoid conflicts with adjacent parts) and for the final population with components (*silkscreen print* which is used to print the reference designator and pin-1 marker on the PCB to ease the correct assembly of the part).
In the *part editor* a symbol and a cell are combined to a *part* by assigning the pins of the symbol to the ones of the cell (that is why the correct numbering of the pins of the symbol and of course of the cell are so important). It is possible to create multiple parts by adding different cells to one symbol and vice versa. This allows to create one symbol for a standard part like a capacitor and then combine it with different standard package sizes (e.g. 0402, 0603[1]) to the corresponding parts (e.g. capacitor_0603). On the other hand these standard cells can be combined with another symbol to form different sized resistors for example.

Drawing the Re-BOC Schematics

For the Re-BOC schematics the Spartan-6 FPGAs are created with one symbol per IO bank, several others for the MGTs (one symbol per dual tile) and for the individual power supplies (for the FPGA logic and for each IO bank to accommodate the signal standards it is intended to use). Of the six banks of the Spartan-6 LX150T FPGAs four are set to an IO voltage of 3.3 V, one establishes the DDR-2 memory interface (1.8 V) and the last one is connected to the mezzanine (LPC-FMC) connector with 2.5 V signals. The many 3.3 V banks are needed because the crucial point of the required backwards compatibility is the transmission of the event data (*RX Data*) to the ROD and the SSO limits for the necessary 3.3 V LVTTL signal standard are tight. In Table 4.4 the limits of the individual banks are listed together with the IO voltages used on the Re-BOC, and the distribution of the connections to- and from the ROD (Command (CMD) and S-Link connections are mostly input signals to the BMFs which do not affect the SSO limits). To ease this SSO issue the communication between Re-BOC and the redesigned ROD will be based on the *Stub Series Terminated Logic* (SSTL) signal standard which requires a 1.5 V reference voltage (for SSTL-3) connected to dedicated pins of the FPGA and greatly increases the SSO limits (see Table 4.4). To provide enough bandwidth without the need for additional connections the data rate is increased to 80 Mbit/s per line and the 96 lines are split between the two BMFs (BMF South: RXDATA[0:47], BMF North: RXDATA[48:95]).

Something similar is done with the backplane connections for the S-Link data, as

[1] Standardized shapes and sizes for rectangular passive *Surface-Mounted Devices* (SMD) with two terminals, the numbers are width and length of the part in tens of mils, e.g. 0603 corresponds to (0.063 x 0.031) inch.

	Bank 0	Bank 1	Bank 2	Bank 3	Bank 4	Bank 5	Total
Maximum IO	105	84	104	112	49	48	502
LVTTL* SSO	15	14	13	14	7	8	71
SSTL-3-I SSO	75	84	65	112	49	48	433
BMF North	2.5V	3.3V	3.3V	3.3V	3.3V	1.8V	
RX Data to ROD	0	14	13	14	7	0	48
CMD from ROD	0	0	0	8	0	0	8
S-Link from ROD	0	0	0	6	40	0	46
Setup Bus via BCF	0	0	0	24	0	0	24
BMF South	3.3V	3.3V	2.5V	3.3V	3.3V	1.8V	
RX Data to ROD	15	13	0	14	6	0	48
CMD from ROD	0	0	0	8	0	0	8
S-Link from ROD	0	0	0	6	40	0	46
Setup Bus via BCF	0	0	0	24	0	0	24

Table 4.4: Maximum IOs and simultaneously switching outputs (SSO) for the Spartan-6 LX150T (FG(G)900 package)[76]. For backwards compatibility LVTTL (* 3.3V, *24mA* drive strength, *fast* slew rate) and for the communication with a new ROD SSTL-3-I is used as signal standard. To stay within the SSO limits the RX Data connections to the ROD are distributed over several banks of the FPGAs. The other connections are mostly input only so they do not have an influence on the SSO limits.

using the former implementation with a 32-bit wide data bus (+ 7 control and status lines) is not possible due to the lack of available connections. Hence for the required four S-Links the data bus width is reduced to 16-bit, but it is run at the double line rate of 80 Mbit/s. This way two of the data busses and one set of the 7 control and status lines are attached to the previous connections on the VME connector (JP2) and the rest utilizes the newly introduced JP0 connector.
The Setup Bus is connected to the BCF which handles the communication with the ROD and controls the BMFs which are attached to it by duplicated versions of the bus. The address lines of the bus are increased to 16 (compared to 10 on the current BOC) and the additional lines are connected to former command lines on JP2 which are not required on the Re-BOC anymore as only two TX modules connect with 8 channels each to the front-end electronics[1].
This connections to the VME connectors provide the required backwards compatibility and also enable the advanced features of the Re-BOC. For the functionality of the Re-BOC its other active components were created as parts (as described in the previous section) and placed on separate schematic sheets to keep a clear structure as can be seen in the following list of the schematic sheets (37 in total, sheet numbers are given in the parentheses) of the Re-BOC.

[1] The current BOC can be equipped with up to 4 TX plugins with 12 channels each, therefore 48 command lines are available.

- Spartan-6 LX150T - BMF north
  - Bank 0 & 1 (1)
  - Bank 2 & 3 (2)
  - Bank 4 & 5 (3)
  - MGTs (4)
  - Power & Ground (5)
- Spartan-6 LX150T - BMF south
  - Bank 0 & 1 (6)
  - Bank 2 & 3 (7)
  - Bank 4 & 5 (8)
  - MGTs (9)
  - Power & Ground (10)
- Spartan-6 LX75 - BCF
  - Bank 0 & 1 (11)
  - Bank 2 & 3 (12)
  - MGTs (13)
  - Power & Ground (14)
- TTC Clock Distribution (15)
- Reference Clock for FPGAs and their MGTs (16)
- Power Supply for FPGA Core, MGT and 2.5 V banks (17)
- Power Supply for 1.8 V and the reference voltages (18)
- DDR-2 Memory for
  - BMF north (19)
  - BMF south (20)
  - BCF (21)
- FLASH Memories (22)
- Optical Components
  - SNAP12 RX (23)
  - SNAP12 TX (24)
  - QSFP (25)
  - SFP (26)
  - old TX Plugins (27)
- Ethernet PHY (28)
- VME Connectors
  - J0 (29)
  - J2 (30)
  - J3 (31)
- Mezzanine Connector FMC LPC north (32)
- Mezzanine Connector FMC LPC south (33)
- ELMB Connector (34)
- NTC Thermistor Sensors (35)
- Test Outputs, LEDs and push buttons (36)
- USB UART and EEPROM (for the BCF) (37)

The necessary decoupling and bypass capacitors are added to the power nets of these parts on their corresponding sheets as well as termination resistors if required. Then all interconnections are established by drawing lines between the IO pins of the parts on one sheet. If a connection to a part on another sheet is required, properly named inter-sheet connectors are employed, as lines and nets that have an identical name are regarded as connected within a design project. This is important for the next step in which the abstract graphical schematics are translated into the actual physical layout.

4.3.2 PCB Layout

In the production of a PCB the required conductive contacts and transmission lines are etched into (typically) thin copper layers. Therefor stencil masks are generated in the layout process by placing the foot-prints of all components on an area which is true to scale to the final PCB and by establishing all connections with transmission lines of adequate size. For electrical and optical[1] reasons these traces[2] need to be routed within certain rules regarding their width and the distance to other traces. Also the structure of the layer stack of the PCB needs to be designed and taken into account.

Design Rules

The lower limits of the structure sizes that can be used for the layout depend on the manufacturing process and the machines used by the company which produces the PCB. Typical values are $\approx$100 μm for trace width and distance between adjacent traces and about 200 μm for the diameter of drilled holes (e.g. for vias). Based on these constraints the optimal trace widths need to be determined to provide the correct impedance for the individual signal standards used in the design. For this at first the layer stack of the PCB needs to be defined.

The layer stack

A PCB is made of conductive layers for the signal transmission that alternate with insulating dielectric layers. The most common material[3] for the insulator is FR-4 (Flame Resistant-4) which consists of fiberglass and a flame resistant epoxy resin and is prepared in two different ways. For the use as a *core* the resin is cured and copper is laminated to both sides of it, contrary to this the resin is not hardened

[1] Regarding the projection and etching of the layout mask.
[2] Transmission lines on a PCB are called traces.
[3] This is due to its good mechanical and electrical characteristics combined with relatively low costs. For the Re-BOC FR-4 is used as well as dielectric material.

for the so called *prepreg* (pre-impregnated) layers which are used to bond the PCB together. After the inner signal layers are created by transferring the circuit design to the copper of the cores by photo-lithography and subsequent etching, they are arranged in a stack with alternating layers of prepreg which is also used to create the outer layers in conjunction with a copper foil. Then this stack is heated and pressed together, causing the prepreg to harden and glue all layers together. Now the required holes (e.g. for vias and *Plated-Through Hole* (PTH) components) are drilled and plated with copper (if needed) before the traces and pads are transferred to the outer copper layers and the surface is finished by applying a solder resistant coating. The composition of the PCB with 12 copper layers which is used for the Re-BOC can be seen in Figure 4.11. The rather odd numbers for the thickness of the copper layers result from the imperial units which are often used to specify PCBs[1], and in which the thickness is given in ounces[2] of copper with typical values of 0.5 oz (17.5 μm), 1 oz (35 μm) and 2 oz (70 μm).

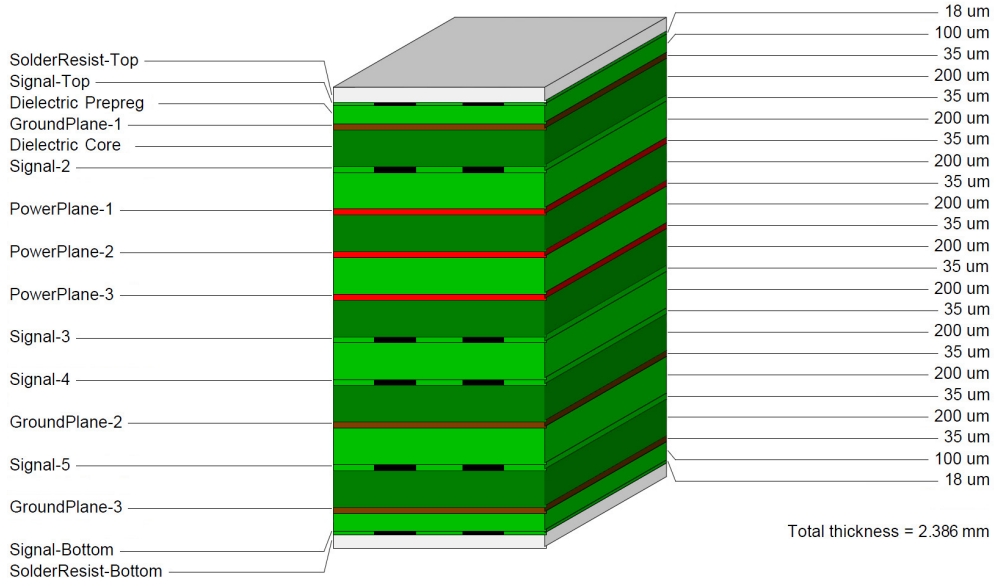

Figure 4.11: Stack structure of the 12-layer (copper) PCB for the Re-BOC (um = μm, picture taken from Mentor Graphics Stackup Editor).

The amount of signal layers of the PCB for the Re-BOC is mostly determined by the Spartan-6 LX150T FPGA and its BGA package with an array of 900 (30x30) balls with a pitch of 1 mm. To connect to all these pins and to route the signals out from underneath the device vias need to be placed between the BGA pads limiting the area which can be used for the so called *breakout* traces. Additional to the signal layers,

[1] This is due to the fact that PCBs were mostly developed and made popular by the USA.
[2] This value is regarding an area of one square foot, so the correct unit is ($\frac{oz}{ft^2}$).

power and ground planes are needed for the power supply of the PCB (as mentioned in Section 4.3.1) which also provide a stable reference plane for adjacent signal layers. This is used to calculate and set the proper impedance of the transmission lines by choosing their dimensions (see next section).
For the layer stack of the Re-BOC PCB 12 layers with standard copper thicknesses are used, and a symmetric structure with solid planes between signal layers is provided to ensure good signal quality. Between Signal Layer 3 and 4 (see Figure 4.11) no such plane is available as it would have required a dielectric core[1] with two additional copper layers causing increased production costs. To prevent any impact on the performance of the Re-BOC these two signal layers (3 & 4) are only used to route traces for noncritical signals such as low-speed control lines or the connections to the test pin-headers. The core and prepreg layer dimensions yield an overall symmetrical PCB with a total thickness within the required range of (2.4 $\pm$ 0.2) mm[77].

Trace Constraints

To ensure good signal quality for transmissions over the traces on a PCB a constant and adequate impedance is required which depends on the layer stack-up and the dimensions of the traces. Software tools are available which ease this process greatly by allowing to vary all parameters until the correct impedance value is calculated. In Figure 4.12 two example calculations for the Re-BOC traces with such a tool can be seen, after the appropriate trace structure is selected in the upper row (or the structure menu) the necessary values can be set and the impedance is calculated.

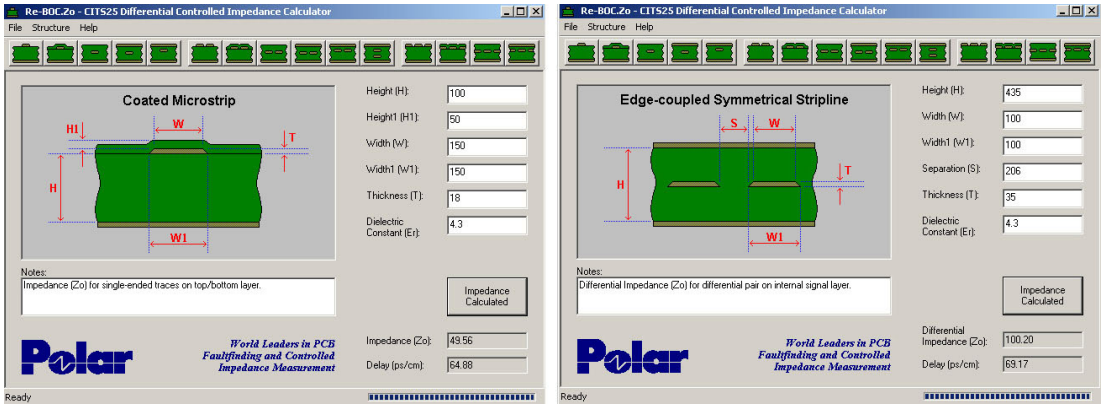

Figure 4.12: Tool from Polar Industries[78] to calculate the impedance of traces in different settings (see upper row of possible trace arrangements). The picture on the left shows the calculation for single-ended 50 Ω traces on an outer layer and the right is for 100 Ω differential impedance in a symmetric inner layer. All dimensional values are given in μm.

[1] This is due to the PCB structure with alternating core and prepreg layers.

In the *Constraint Editor System* of the Expedition Layout tool[1] the layer stack-up can be set and the corresponding impedances for the individual layers calculated. Each electrical net of a circuit design can be assigned to a *net-class* for which individual constraints can be chosen. For example, the width and separation distance of the two traces of a differential pair is set for each layer to the appropriate values yielding an impedance of 100 Ω. When a net of a certain class is routed in the layout tool the corresponding constraints are applied automatically, regarding the former example this leads to both traces being placed at the same time.

The Re-BOC PCB Layout

After the layer stack and the constraints are set, the outline of the final PCB is created in the *Draw Mode* of the layout tool which is used to add geometrical shapes or to edit the silkscreen print items. Then, after all design information is imported from the schematics via the so called *forward annotation*, the parts[2] can be placed in the *Place Mode*, and the connections between them are established in the *Route Mode*.
The placement of the parts on the Re-BOC is mostly determined by the optical components on the right edge and the VME connectors on the left edge of the board. Between them the BMF FPGAs are positioned central on the PCB to allow the routing of all necessary breakout traces and to keep the lengths of traces from the same signal families as equal as possible to reduce runtime skew. Special care was taken for the transmission lines for the RX-Data to the ROD by adding meanders to the short ones of them to increase their length and thus their runtime to adjust them to the longer ones (see lower left corner of Figure 4.13). For the same reason the components for the handling of the 40 MHz TTC clock are placed and routed to have minimal skew in the signals to the individual FPGAs. The BCF is located between the two other FPGAs to provide easy connections for the Setup Bus to the BMFs and the other components controlled by it. A list of the minimal and maximal runtime and the resulting skews for the important transmission lines on the Re-BOC is given in Table 4.5. All skews are below one nanosecond and can be easily balanced by the delay elements in the IO cells of the Spartan6 FPGAs.

The DDR2 memory chips are placed close to the FPGAs and their corresponding IO banks that establish the memory interface. Other, not crucial parts like the ELMB and mezzanine card connectors or the test pin-headers are located on parts of the PCB where they do not disturb the routing of important signals. Components for decoupling, bypassing and termination (as introduced in section 4.3.1) are distributed on both sides of the PCB to allow ideal placement close to their destination.

[1] Which is a part of the Mentor Graphics Expedition Enterprise Flow development environment.
[2] Not the actual part is placed but the corresponding land pattern.

Signal	From	To	Runtime (min - max) [ns]	Skew [ns]
RXDATA[0:47]	BMF south	ROD	0.78 - 1.25	0.47
RXDATA[48:95]	BMF north	ROD	1.21 - 1.66	0.45
RXDATA[0:95]	Re-BOC	ROD	0.78 - 1.66	0.88
Clk40_BOC0	PLL	BMF north	1.67	0.08
Clk40_BOC1	PLL	BMF south	1.65	0.08
Clk40_BOC3	PLL	BCF	1.66	0.08
Clk40_2ROD	PLL	ROD	1.59	0.08
Setup Bus Add	ROD	BCF	0.51 - 0.72	0.21
Setup Bus Data	ROD	BCF	0.62 - 0.73	0.11
Setup Bus Add	BCF	BMF north	0.21 - 0.39	0.18
Setup Bus Data	BCF	BMF north	0.23 - 0.33	0.1
Setup Bus Add	BCF	BMF south	0.47 - 0.74	0.27
Setup Bus Data	BCF	BMF south	0.44 - 0.69	0.25
S-Link Data	ROD	BMF north	0.58 - 0.78	0.2
S-Link Data	ROD	BMF south	0.57 - 1.19	0.62
QSFP RX	QSFP	BMF south	0.59 - 0.63	0.04
QSFP TX	BMF south	QSFP	0.44 - 0.6	0.16
SFP RX	SFP	BMF north	0.53 - 0.72	0.19
SFP TX	BMF north	SFP	0.41 - 0.61	0.2
SNAP12 RX-A	SNAP12	BMF south	0.34 - 0.47	0.13
SNAP12 RX-B	SNAP12	BMF south	0.31 - 0.44	0.13
SNAP12 RX-C	SNAP12	BMF north	0.30 - 0.43	0.13
SNAP12 RX-D	SNAP12	BMF north	0.30 - 0.47	0.17
SNAP12 TX-A	BMF south	SNAP12	0.36 - 0.49	0.13
SNAP12 TX-B	BMF north	SNAP12	0.58 - 0.72	0.14
GMII RX	PHY	BCF	1.64 - 1.72	0.08
GMII TX	BCF	PHY	1.71 - 1.78	0.07
DDR2 ADD	BMF north	RAM	0.05 - 0.15	0.1
DDR2 DATA	BMF north	RAM	0.13 - 0.27	0.14
DDR2 ADD	BMF south	RAM	0.06 - 0.15	0.09
DDR2 DATA	BMF south	RAM	0.1 - 0.28	0.18
DDR2 ADD	BCF	RAM	0.05 - 0.1	0.05
DDR2 DATA	BCF	RAM	0.03 - 0.12	0.09

Table 4.5: Minimum and maximum signal runtime and skew in the transmission lines on the Re-BOC. For differential signals (to the optical components) the runtime difference between individual signals is given and not the one between corresponding differential pairs, those are all maximal 0.01 ns.

After all parts are placed and all necessary connections are established by the appropriate transmission lines (Figure 4.13 shows the finished layout of the Re-BOC (Revision B)) the electrical behavior and signal integrity can be simulated. For this the Mentor Graphics *HyperLynx PCB Analysis Software* is included in the Expedition Design Environment and allows to simulate single nets or complex components and their interactions. To simulate components most vendors provide files with *Input/output Buffer Information Specifications* (IBIS) for their parts which include information about its electrical IO behavior and parasitic values without the need to disclose the internal functionality.
When the design and its proper functionality has been verified an etching stencil is created for each layer and saved in a so called *Gerber*[1] file. These files, including the drilling information for the holes and the silk screen print, are then sent to the PCB manufacturer who builds and electrically tests the board. Not the logical behavior is tested, but the different electric nets are checked for short circuits and the impedance of the transmission lines is controlled. The finished PCB (see Figure 4.14 for a picture of the Re-BOC PCB) is then populated with components by soldering them to the appropriate contacts (Figure 4.15).

[1] This file format is often used to transfer image data for PCB production, it is named after H. Joseph Gerber and his company (Gerber Scientific Instrument Company) who developed it.

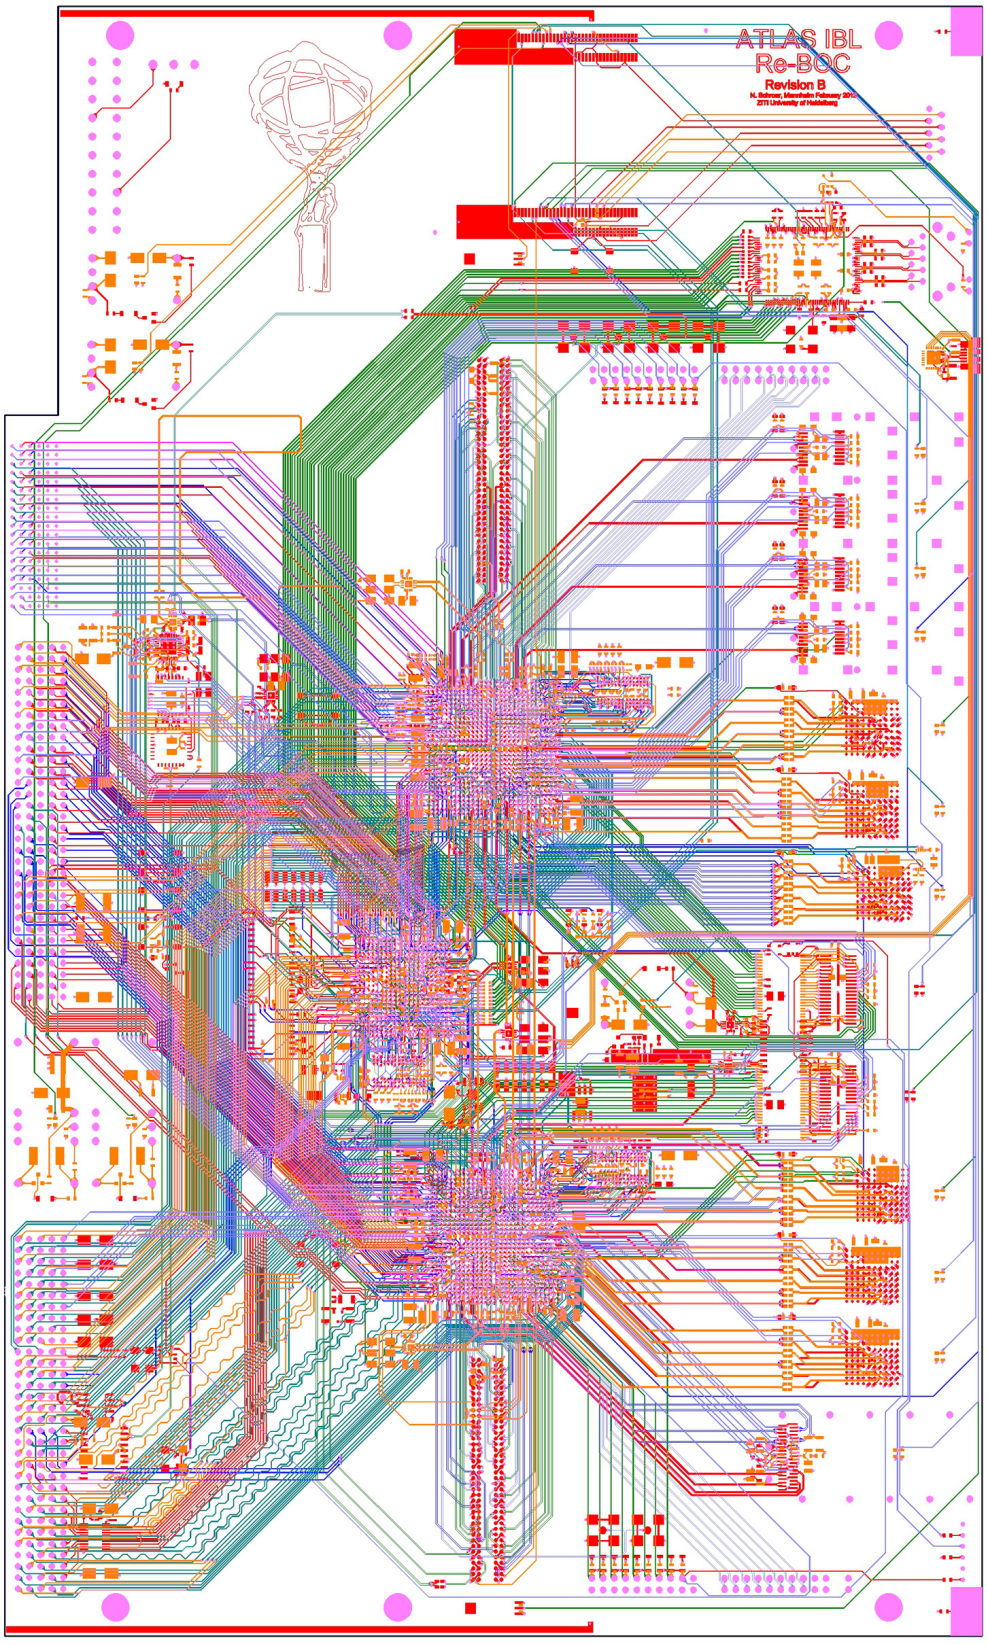

Figure 4.13: Overview of the layout of the Re-BOC with all routing layers visible.

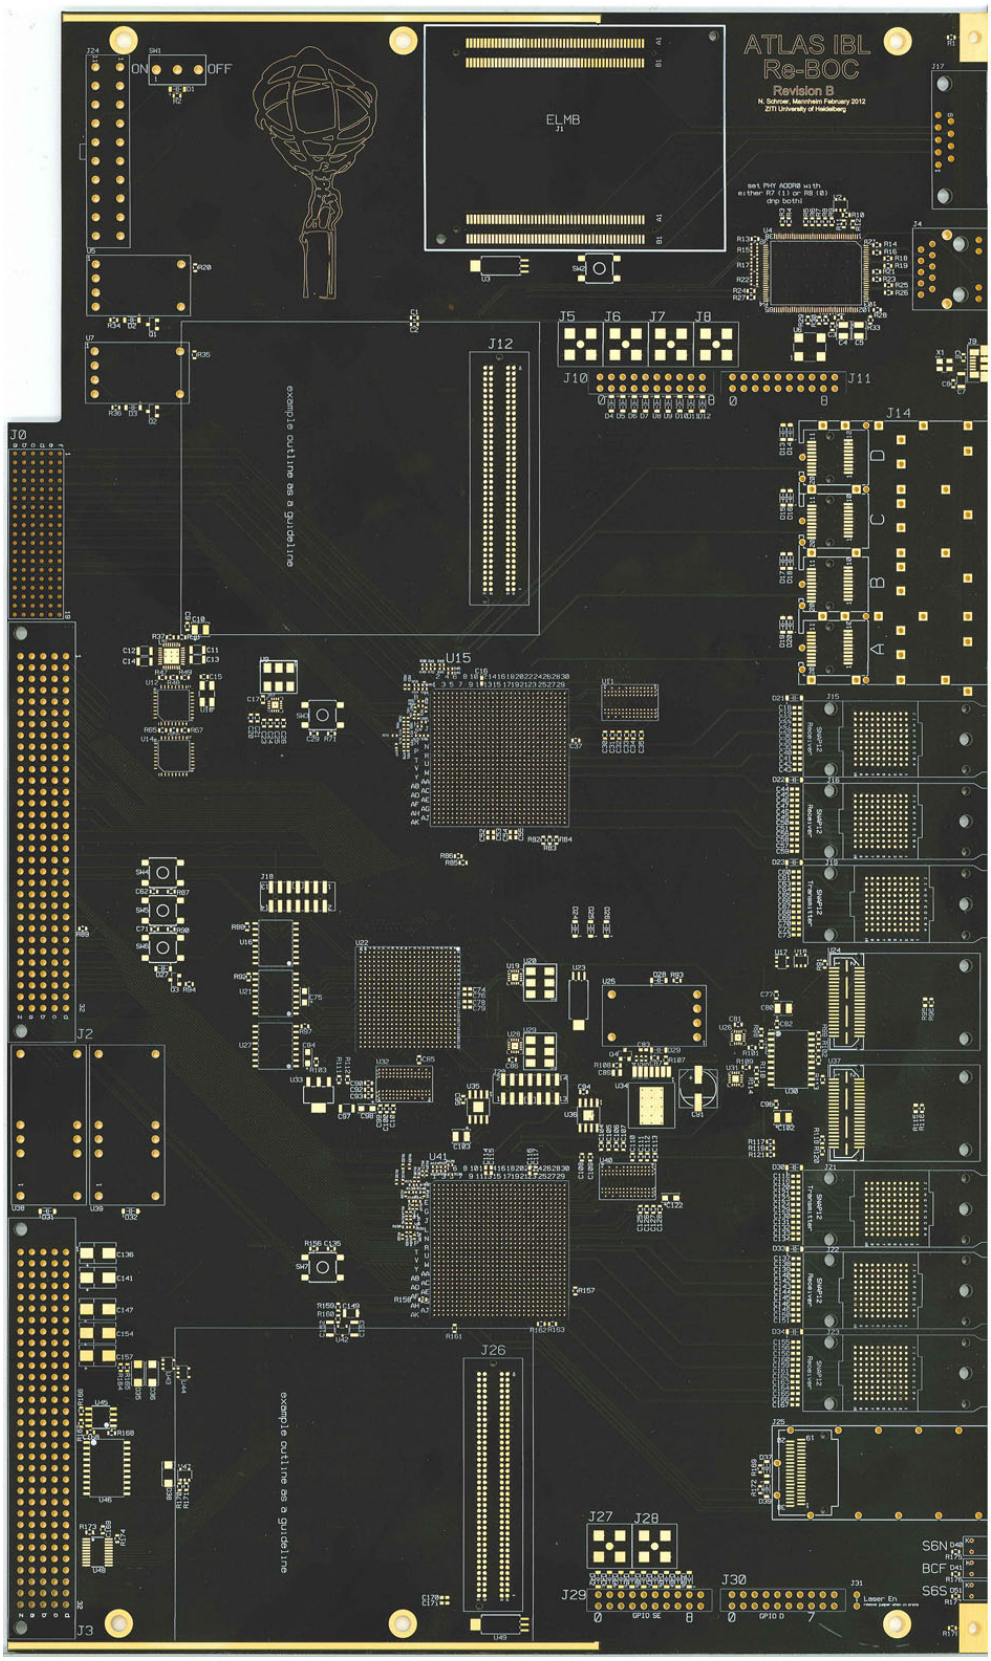

Figure 4.14: Picture of the unpopulated front side of the Re-BOC (Rev. B).

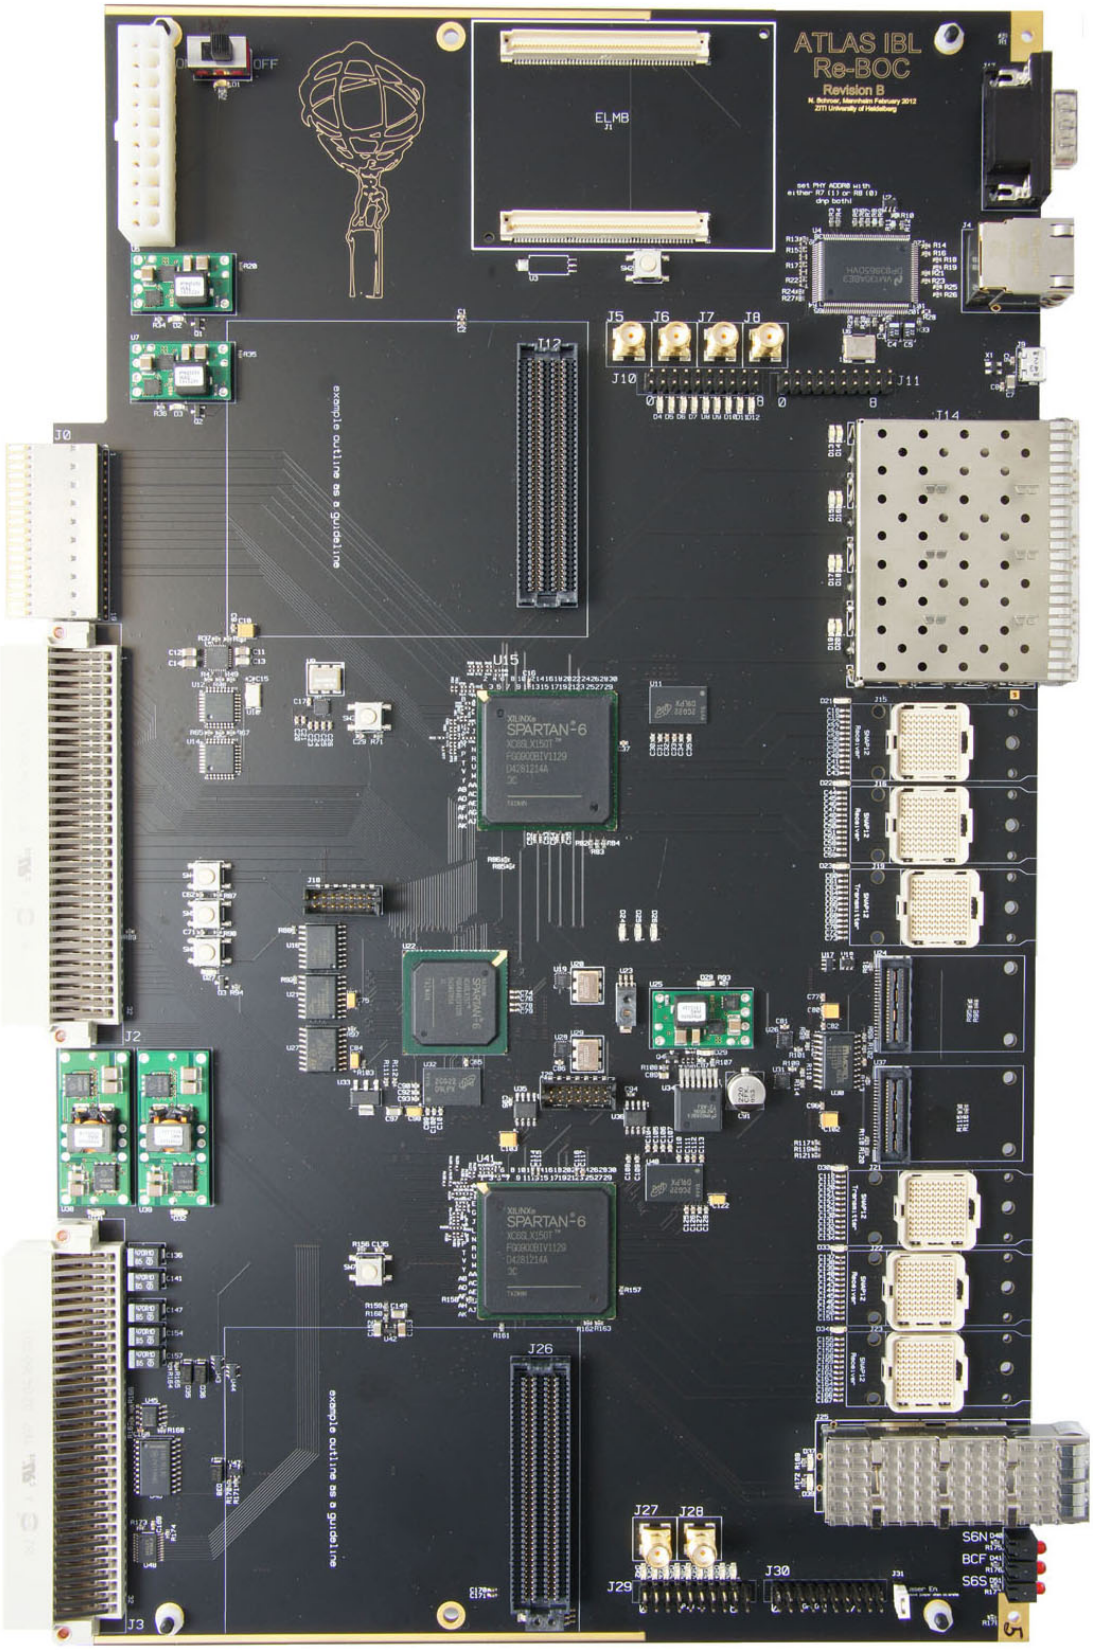

Figure 4.15: Photo of the Re-BOC (Rev. B) populated with components.

5

IBL Re-BOC - Firmware

Due to the concept of the Re-BOC being based on configurable hardware, the firmware required to implement the necessary functionality is an important part of the redesign. The development of the firmware blocks started before the first Re-BOC prototype was at hand, thus a demonstrator setup was used for the evaluation[52, 79]. When the first Re-BOC PCBs became available the design was migrated and is now refined and tested on the prototypes[80].
In this chapter the firmware blocks for the individual data processing and management parts will be introduced and described in detail. Two different firmware sets will be considered, at first the firmware for the BCF managing the control functions on the Re-BOC will be presented, then the part for the data path handling in the BMFs is described. Following that, a short glimpse on the software to connect a computer to the Re-BOC and access an attached FE-I4 chip will be given.

5.1 Control Firmware for the BCF

The main purpose of the BCF is what gave it its name (BOC Control FPGA), controlling the Re-BOC. For this task it is connected to the ROD via the Setup Bus and passes the received commands to the BMFs. Integration of a MicroBlaze soft-core processor into the BCF allows to run the Re-BOC without the need of a ROD to control it. Commands can be received via the Ethernet- or the USB UART interface, but its main application will be to provide diagnostic features for the board. Another important task of the BCF is the management of the 40 MHz clock received from the TIM in the crate, therefore it is connected to the delay buffers and the PLL (introduced in Section 4.2.3) to steer the synchronization and distribution of the necessary clock signals. The firmware for all three FPGAs and the software for the MicroBlaze core is stored in flash memories attached to the BCF which configures itself from these

after power-on or a board reset before it reads and distributes the firmware for the two BMFs. An overview of the BCF firmware blocks can be seen in Figure 5.1.

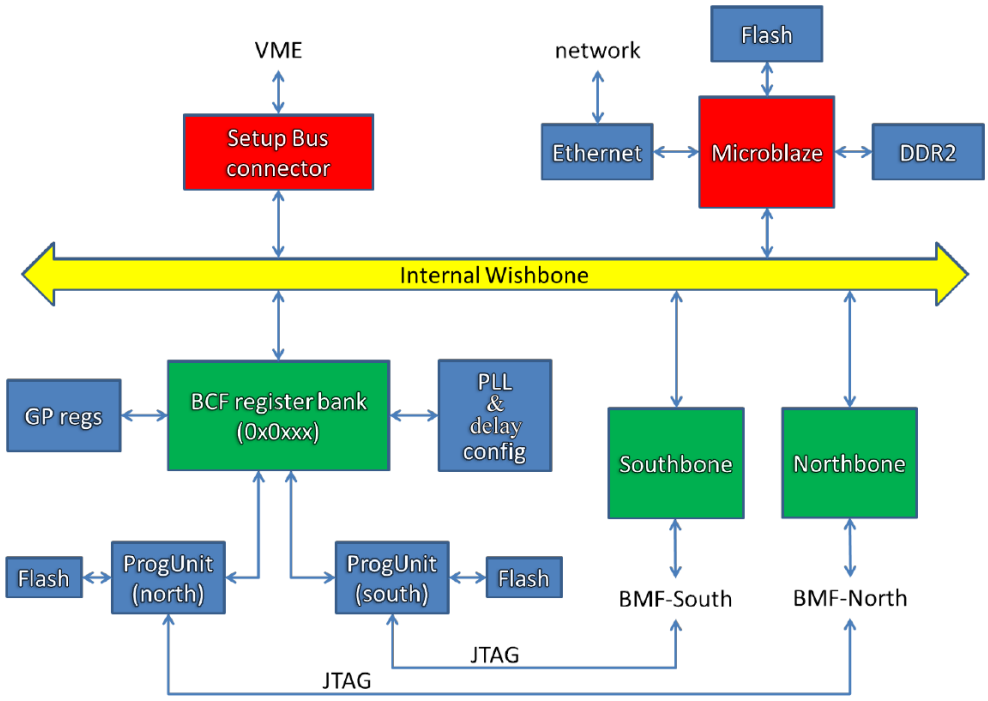

Figure 5.1: Block Diagram of the BCF firmware[80].

5.1.1 Setup- / Wishbone Bus

The Setup Bus connection via the VME connectors to the ROD is established by a 16-bit address- and an 8-bit data bus with three control signals (Write Enable, Strobe and Busy). In the BCF this bus is converted into an internal bus structure based on the Wishbone Interconnection Architecture[81] that interfaces to the MicroBlaze core and to register banks in the BCF and the BMFs.

FPGA	**address bit**															
	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
BCF	F	F	X	X	X	X	X	X	X	R	R	R	R	R	R	R
BMF	F	F	X	X	S	S	C	C	C	C	C	R	R	R	R	R

Table 5.1: Address scheme for the Setup Bus on the Re-BOC, F: FPGA select (00: BCF, 01: BMF-South, 10: BMF-North), X: unused, R: register number (BMF: 0-31, BCF: 0-127), S: register select (00: general registers, 01: S-Link registers, 10: RX path registers, 11: TX path registers), C: channel number (0-31).

The 16 address bits allow to access up to 64 kB[1] of memory which will be used for registers that control and configure all the functionalities on the Re-BOC. So far only a small fraction thereof is utilized. The BCF provides 80 registers to store information about the installed firmware version and the PLL and BMF configuration. Each BMF has 22 general registers to store the firmware version number, and additionally 12 registers for each of the 32 RX channels plus 9 for every TX channel to provide access to control and status information. In Table 5.1 an overview of the address scheme for these registers is given. A complete definition of all registers and its settings can be found in the IBL-BOC Users Guide[82].

5.1.2 Clocking

To control the delay setting of the two cascaded delay chips, an 11-bit wide bus interfaces between the BCF and the chips, while the PLL is controlled via a 4-wire SPI connection. Additionally the BCF can set the frequencies of the reference clocks for the MGTs of all FPGAs on the Re-BOC.
The delay which is added by the programmable delay chips is set by applying a binary vector to the lower 10 bits of the control bus (see Figure 5.2). This results in a delay corresponding to the decimal value of this vector times 10 ps, hence delay values from 0 to 10,230 ps can be set. If more delay is needed the 11th bit is set which controls the $CASCADE$ and $\overline{CASCADE}$ outputs which are connected to the SETMAX and SETMIN inputs of the second delay chip and either set the delay to its maximum value[2] or to zero (the fixed delay of $\approx$ 2.2 ns per delay chip needs to be considered).

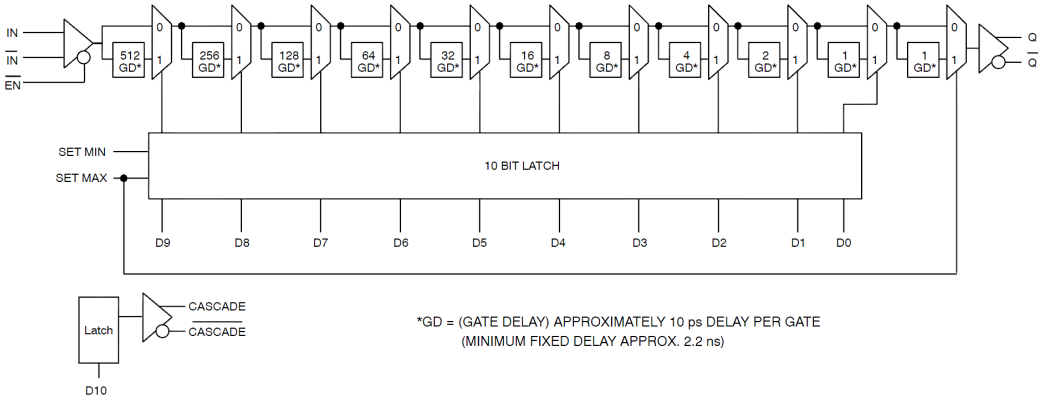

Figure 5.2: Diagram of the the delay selection logic in the programmable delay chip MC100EP195[83].

[1] 64 k addresses for 8-bit register.
[2] The maximum delay value in this case is 10,240 ps as there is an additional 1x GD delay in the chain which is only activated by SETMAX (see Figure 5.2).

The PLL is configured by writing *x"FF7000E"* to register 0 and *x"839F21E"* to register 1, this settings select the appropriate input (LVDS and auto select between REF_IN and AUX_IN (see Figure 5.3)), set the output buffer to LVPECL and choose the correct settings required for an output frequency of 40 MHz. The relation to the input frequency is, with respect to Figure 5.3, given by:

$$F_{out} = F_{in}\frac{F}{R \cdot I \cdot O} \tag{5.1}$$

To fulfill Equation 5.1 with $F_{in} = F_{out} = 40MHz$ and the additional constraints $1.750GHz < O \cdot P \cdot F_{out} < 2.356GHz$ and $40kHz \leq F_{COMP} = \frac{F_{in}}{R \cdot I} \leq 40MHz$ the following values are programmed (by setting register 0 & 1 to the values mentioned before) R = 1, O = 16, I = 16, P = 3, F = 256.

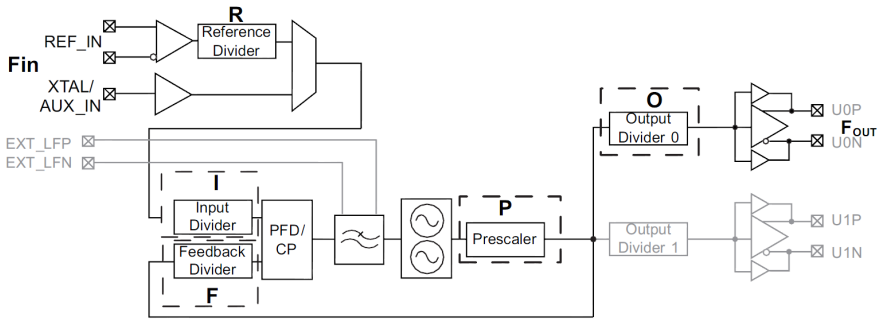

Figure 5.3: Clock path in the CDCE62002 Frequency Synthesizer with the individual dividers that can be set[71]. The components between the input/feedback divider and the prescaler form the PLL and are (from left to right) *Phase Frequency Detector* (PFD)/ *Charge Pump* (CP), loop filter and a dual *Voltage Controlled Oscillator* (VCO)

5.1.3 IP cores

Other firmware components can be included into the design by using *Intellectual Property* (IP) cores that provide FPGA and architecture specific functionality. These cores can be customized and created by the use of the Xilinx Core Generator tool. Thereby a *Gigabit Media Independent Interface* (GMII) core is generated for the BCF to establish the connection to the Ethernet PHY. The memory interfaces to the DDR2 chips attached to each of the three FPGAs is also an IP core generated with the help of this tool. Both these interfaces are attached to the MicroBlaze processor which handles the network communication and utilizes the DDR2 memory. Less complex components that are generated include several FIFOs, clock generators and ChipScope[1] cores for the debugging of the firmware.

[1] ChipScope is a powerful tool from Xilinx that allows to insert and access *Integrated Logic Analyzer* (ILA) and *Virtual IO* (VIO) cores into an FPGA design.

5.2 Data Path Handling Firmware for the BMFs (North & South)

Both BMFs have equivalent firmware with only minor variances in the pinout due to different routing which is caused mainly by the placement and orientation of the individual FPGAs on the PCB and the differing optical transceivers for the S-Link (SFP / QSFP). In Figure 5.4 the individual firmware blocks are illustrated together with their corresponding registers that are accessed by the BCF via the Wishbone connection to configure and control the behavior of the BMF. Its purpose is the management of the command and data streams between ROD, front-end and the ROS as part of the ATLAS DAQ system. For each of these three tasks an individual firmware block is responsible and will be introduced in the following.

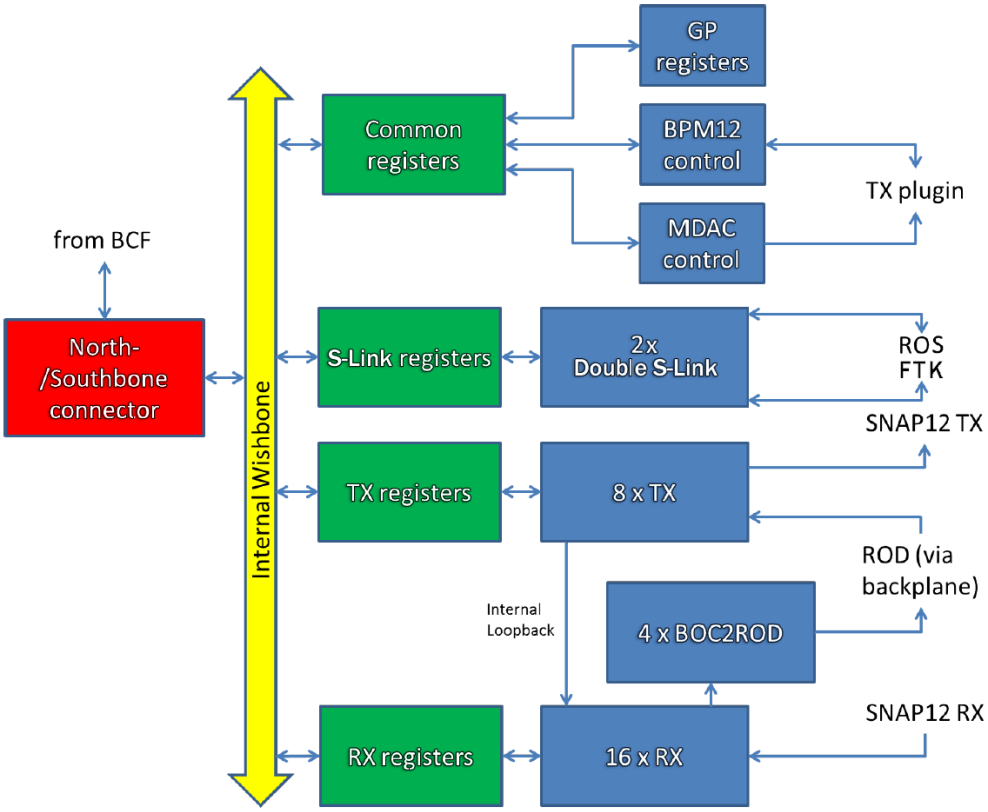

Figure 5.4: Diagram of the firmware blocks for the BMFs[80].

5.2.1 TX Path

This firmware block implements all the functionality of the TX path of the current BOC which was described in Section 3.1.2.
The TTC command streams for the front-end modules (8 per BMF) are received from the ROD and are then BPM encoded with the help of an 80 MHz clock. A transition in the resulting clock-data stream is induced at every or every other rising edge depending on whether the input is high or low respectively. The advantage of this implementation is that the duty cycle of the output is close to 50% which reduces the need to correct it[1]. The timing of the BPM encoded stream can be adjusted

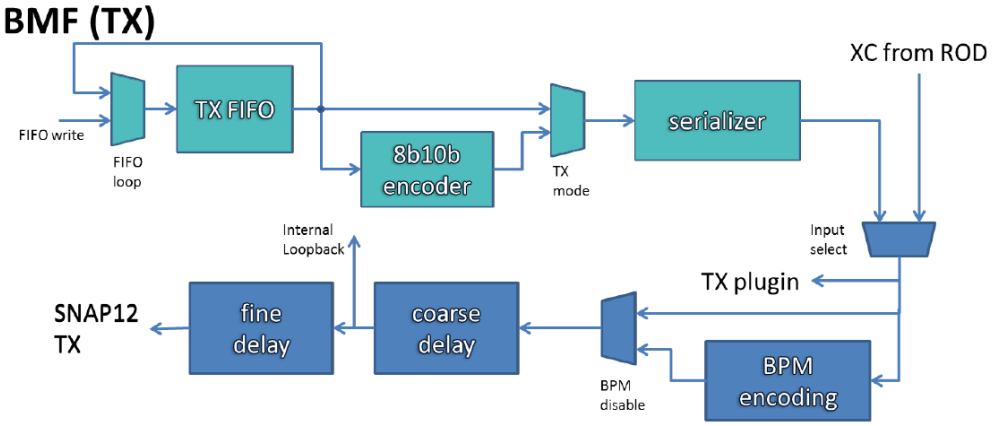

Figure 5.5: Firmware blocks in the TX path of the Re-BOC, the turquoise blocks are for test purposes (standalone without ROD and/or FE-I4)[80].

in the following delay blocks to compensate the different propagation delays in the connections to the front-end. This is necessary to guarantee the synchronization of the readout to the correct bunch crossings. Coarse delay is added by applying a variable number of up to 255 registers to the data path which are clocked with 160 MHz and thus introduce delay in steps of 6.25 ns[2]. The fine delay adjustment is used to tune the phase-relation of the streams and is available in two different implementations. One is based on the IODELAY cells in the buffers of the Spartan-6 and allows to add delay in steps of ≈(10 - 300)[3] ps up to a maximum of ≈10 ns. The second implementation utilizes chains of up to 128 multiplexers as delay elements with a step size of ≈275 ps which allows a total delay of around 35 ns. Both implementations are not perfect as they have an influence on the duty cycle of the delayed signal (see the next chapter). To counter this effect the implementation based on multiplexers has a built-in MSR adjustment circuit before the fine delay is

[1] As already mentioned in section 2.3.1 this is important as the DORIC generates its clock signal from this stream via a CDR circuit and requires a precise duty cycle of (50±4)% to ensure correct function[33].
[2] Hence this coarse delay can delay the data stream by more than 63 40 MHz clock periods (25 ns).
[3] The duration of the tap delay steps depends on the speed grade of the FPGA.

added. This adjustment block has 32 different settings to either shorten or broaden the input pulses. The shortening is done by splitting the data path and adding a variable delay (Settings 0-15, delay increases with number) to one path while the other is delayed by a fix delay which is bigger that the variable delay at setting 15 before connecting both paths to an AND gate. The output of this gate is then split again and another variable delay can be added to one path (Settings 17-31) before this and the other path which is not delayed are connected to an OR gate to allow the broadening of the pulse. In setting 16 the variable delay of the first block is set to the same delay as the fixed delay in its other path and in the second block the variable delay is set to zero, hence the pulse is neither shortened nor broadened.
Additional firmware blocks for standalone tests (the turquoise ones in Figure 5.5) allow to input command data without a ROD connected and to simulate the front-end electronics by providing a 8b10b encoded serial data stream and an internal loopback option.

5.2.2 RX Path

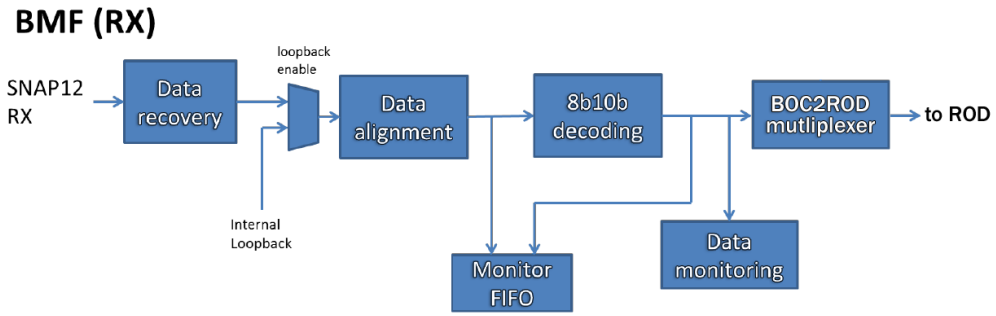

Figure 5.6: Block diagram of the firmware blocks for the RX path in the BMFs of the Re-BOC[80].

This part of the firmware (see Figure 5.6) needs to handle the 160 Mbit/s DC balanced 8b10b encoded serial data streams from the new IBL front-end electronics that are the main reason for the redesign of the BOC as the current version can only manage 80 Mbit/s (as described in Section 3.1.2). The data streams are received by an edge detection and data recovery block which is inspired by an algorithm described in Xilinx application note XAPP224[84]. It utilizes the ISERDES blocks of the Spartan-6 to perform 4-times oversampling and is controlled by a state-machine. Then the serial streams are parallelized with shift registers, and the resulting 10-bit words are aligned correctly with the help of special bit patterns. For these so called comma characters the common 8b10b control symbols K28.1, K28.5 and K28.7[1] are

[1]These characters are: K28.1: 001111 1001 K28.5: 001111 1010 K28.7: 001111 1000 and their inverse respectively.

used as Idle, *End of File* (EoF) and *Start of File* (SoF) markers. For the following 8b10b decoding circuit several straight forward implementations are available[1], and for the Re-BOC an IP core is used which is created by the Xilinx Core Generator. The decoded 8-bit words from 4 data streams are then multiplexed and transmitted to the ROD via a 12-bit wide bus which runs at 80 MHz and consists of 8 lines for the data, 2 address bits to identify the stream number and 2 control bits (valid, k-word). Four of these multiplexer blocks are used in each BMF for the transfer of the data of the 16 FE-I4 chips that are connected to it, hence both BMFs combined utilize 96 connections to the ROD with a cumulative bandwidth of 640 MByte/s[2]. Monitoring blocks will be available to allow the supervision of the data transfer including information about errors therein.

5.2.3 S-Link

Each of the BMFs will handle two sets of S-Link signals which are composed of 16 data and 7 control & status lines and transfer data with 80 MHz, thus offering the S-Link bandwidth of 160 MByte/s. In this firmware block the S-Link source interface is implemented and the resulting data streams are duplicated before they are sent to the MGT interface which serializes the data and sends it to the optical transceivers. Hence two double S-Links are available per BMF to transfer the event fragment data to the DAQ and to the FTK system. The bandwidth will be 160 MByte/s per link to be compatible with the current ROS electronics but higher rates are possible (on the Re-BOC) and can easily be implemented with a firmware update. Same as on the current BOC the link return lines will not be used, but flags are available to indicate if the link is down (LDOWN) or full (LFF) to provide options for flow control.

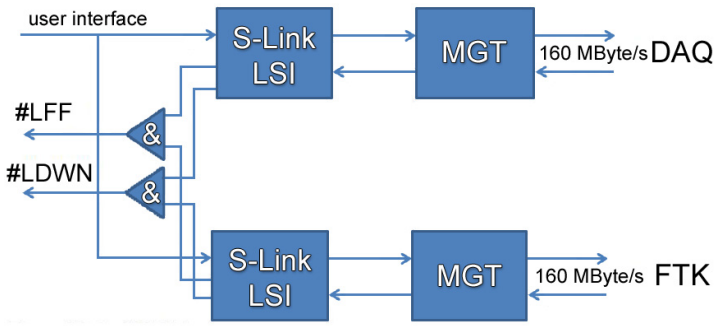

Figure 5.7: Double-S-Link block implementing two Link Source Interfaces (LSI) for the data transfer to DAQ and FTK. Both flag signals (LFF, LDWN) are low-active, thus the logical AND is effectively an OR.

[1] For example on www.opencores.org.
[2] The same 96 VME connections as on the current BOC are used which has a max bandwidth of 160 MByte/s.

5.3 The IBL Re-BOC Library

The *IBL Re-BOC Library* is a collection of C++ classes that can be used to access the IBL Re-BOC from a computer that is connected via the Ethernet interface[82]. This allows to perform standalone tests without a ROD card by providing the Setup Bus functionalities, including access to all registers on the Re-BOC, to the software interface.
All classes use constants definitions etc. from the addressmap header file and the *BOC*, *Options* and *FE-I4* classes (the red rectangles in Figure 5.8) are constructed directly. The BOC class being the central part of the library is constructed with the IP address of the Re-BOC as argument to establish the communication, and it also provides functions to create the BCF and BMF classes (GetBcf and GetBmf respectively). The BCF class gives access to the functions of the BCF and is for example used to flash the BMF firmware (the code for this program can be found at the end of this section). The BMF class can be used to construct the RX- and TX-channel classes which provide access to status information and can be set to read and generate test data as well as communicate with a real FE-I4 chip. For this the FE-I4 class is constructed directly with the attached RX- and TX-channel and the FE-I4 chip ID as argument. The last class to be mentioned is Options, which specifies all the available options for programs using this library.

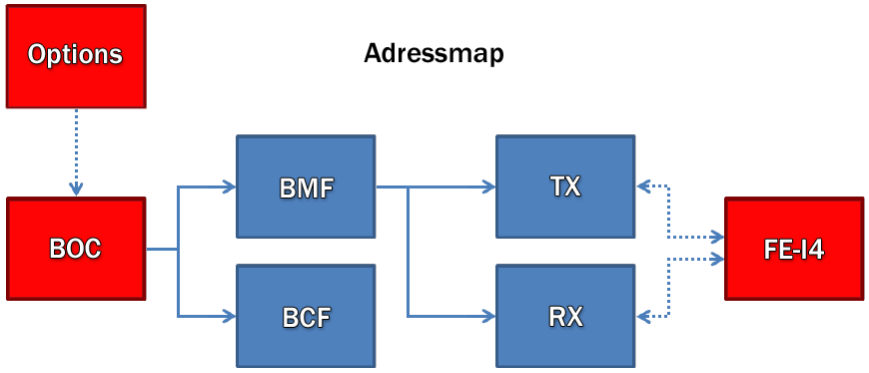

Figure 5.8: Illustration of the classes of the IBL Re-BOC Library. The red classes are constructed directly, the arrows indicate constructions by another class and the dotted arrows illustrate information flow.

```
# include "AddressMap.h"
# include "Boc.h"
# include "Bcf.h"
# include "Options.h"
int main (int argc , char * argv [])
{
        Options o(argc , argv );
        std :: cout << "BOC-IP:" << o. GetIp () << std :: endl ;
        Boc boc(o. GetIp ());
        Bcf bcf = boc . GetBcf ();
        bcf . FlashBmf (o. GetTarget (), o. GetFlashFile ());
        return 0;
}
```

Code example for the IBL Re-BOC Library showing the FlashBmf program[82].

6

IBL Re-BOC Prototype Testing

While the hardware of the Re-BOC was being designed, the firmware blocks for it were developed and tested on an evaluation board setup[52, 79]. When the first prototype was available the testing was extended to the hardware components on the PCB and the firmware was adapted to it. Meanwhile the second revision of the Re-BOC prototype is at hand and the evaluation of its functionalities is ongoing[80]. Standalone tests with local or external optical loopbacks are complemented with system tests. These include FE-I4 *Single Chip Cards* (SCC)[1] to validate the communication with the front-end, and the crate internal connectivity is tested with the ROD prototype[2] which is under evaluation as well.
In this chapter the test process and its preliminary outcomes will be presented.

6.1 Evaluation Board Test-Setup

As no actual Re-BOC hardware was at hand at the beginning of the design process the firmware blocks for it were developed and tested on the basis of a commercial evaluation platform[3]. For this a *Xilinx SP605* board was used which is equipped with a Spartan-6 LX45T and hence provides the same FPGA architecture for testing as will be used on the Re-BOC.
To provide further options to evaluate the firmware design custom made daughter boards could be attached to the mezzanine connector (FMC LPC) of the SP605. With their help optical transmissions over SNAP12 transmitter, receiver and QSFP modules could be integrated into the test setup as well as a FE-I4 SCC.
An USB UART connection to a computer provided the communication interface to the

[1] This are small mezzanine boards equipped with a single FE-I4 chip (hence the name) and the necessary peripheral components and interface connectors to provide access to the chip.
[2] At the time this thesis was written the second revision of the ROD prototype for the IBL was used.
[3] This is described in detail in the Diploma Thesis of T.Heim [52].

internal Wishbone Bus structure to evaluate the Setup Bus to Wishbone conversion and to control and supervise the testing of the other firmware blocks. The proper functionality of this Setup Bus interface was validated by running the corresponding firmware block[1] of the current ROD on a second evaluation board.

Standalone Tests

The first basic tests concerned the TX-path firmware blocks (Figure 5.5), dummy data was written into the TX FIFO and the BPM encoding was verified at first in a simulation (Figure 6.1) and then with an oscilloscope at the output (pins on a mezzanine board). A delay of two cycles of the 80 MHz clock used for the encoding is caused by this block, but the encoding works flawless and yields a BPM encoded data stream with a duty cycle of (50.6 ±1)%[52].

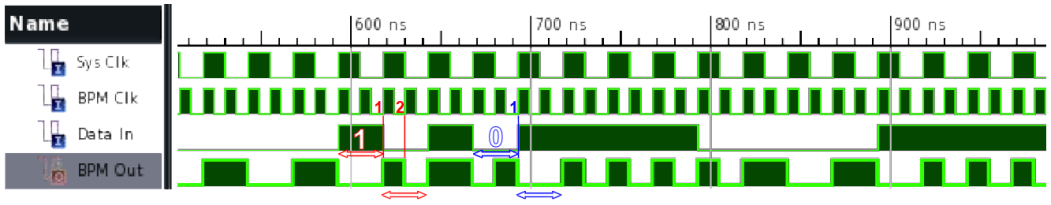

Figure 6.1: Simulated BPM encoding: The input data is clocked with the 40 MHz Sys Clk and encoded with the 80 MHz BPM Clk[52]. The resulting output has transitions on every rising edge of the BPM Clk if a 1 is encoded and a single transition for a 0.

The next step was the fine delay block which was tested on the evaluation board setup in the multiplexer implementation that includes the MSR adjustment circuit. It turned out, that the placement of the multiplexer delay elements in the FPGA is important to optimize the size of the delay steps to achieve a linear relation between the number of steps and the resulting delay. In Figure 6.2 the influence of the placement optimization[2] on the delay steps of the MSR adjustment is illustrated. The width of the pulse can be changed by about ±5 ns around its nominal value (25 ns period at MSR setting 16) with a step size of ≈(339±4) ps per MSR setting. The fine delay step size and the influence of it on the duty cycle of the signal is shown in Figure 6.3, the measurement shows a granularity of ≈(275±3) ps per fine delay step and a strong dependency of the output duty cycle on the added delay.

For the testing of the RX path dummy data is written into the TX FIFO, then it is 8b10b encoded and serialized (the BPM encoding is bypassed) before it is looped back (either externally or internally) with 160 Mbit/s into the RX path (Figure 5.6).

[1] The implementation of the Setup Bus interface on the ROD.

[2] *Relative Location Constraints* (RLOC) are used to automatize the placement of the multiplexer chains by defining the placement of the individual elements in each chain but not the absolute positions of the chains.

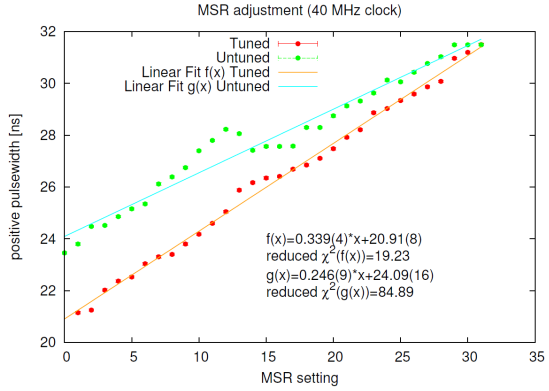

Figure 6.2: Graph of the pulse width measurements in relation to the MSR adjustment setting, contrary to the green- (upper) the red- (lower) series of measurements were obtained from an circuit implementation with optimized multiplexer placement[52].

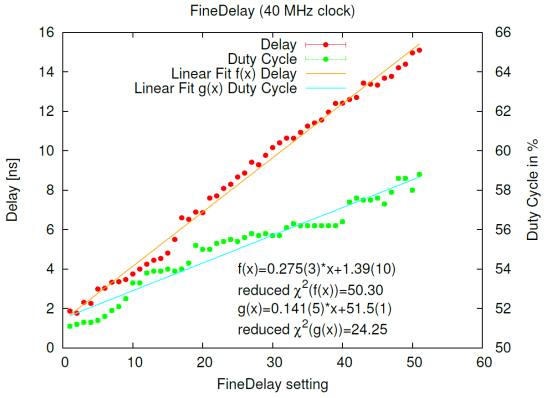

Figure 6.3: Measurement of the fine delay step size (red (upper) measurement points) and its influence on the output duty cycle (depicted with the green (lower) row of dots) in relation to the fine delay setting[52].

There the data is received (in the oversampling circuit for the external loopback), deserialized and properly aligned before it is decoded and passed to the BOC-to-ROD multiplexer. The output thereof is compared to the initial data written into the TX FIFO to estimate the error rate of the transmission. This first loopback test worked without errors, so the phase of the outgoing transmission was shifted by applying fine delay in the TX path to simulate a more realistic scenario in which the data is in an arbitrary phase relation to the receiving circuit and not in a fixed phase with rather low skew, as it was due to the short loopback. This *phase scan* measurements had first transmission errors occurring for a fine delay setting of 35 and increasing errors for higher delay settings, caused by the duty cycle distortion (see Figure 6.4). But this measurement showed that the receiving firmware block is working independently from the phase relation of the incoming data as the output phase is delayed by more than a complete clock period (35 x 275 ps = 9.625 ns compared to a clock period of 6.35 ns) before first errors arise.

In the next test the influence of the MSR adjustment on the data reception was evaluated, and it was found that the transmitted data is received without errors for MSR settings between $\approx$ 10-25, so again the errors are caused by a distorted duty cycle. Finally both tests were combined and the error rate for all combinations of MSR- and fine delay settings were measured. The result is shown in Figure 6.5 and illustrates what was to be expected: The window of MSR settings that yields an output with acceptable duty cycle, and thus no errors (e.g. $\approx$ 10-25 for a fine delay setting of 0), shifts with increasing fine delay and hence increasingly distorted duty cycle towards lower MSR settings. So both blocks can be used together to compensate the

influence of the delay block.
All tests regarding the data paths were initiated via a software interface allowing to access the registers in the firmware design via an UART connection between the SP605 and a computer.

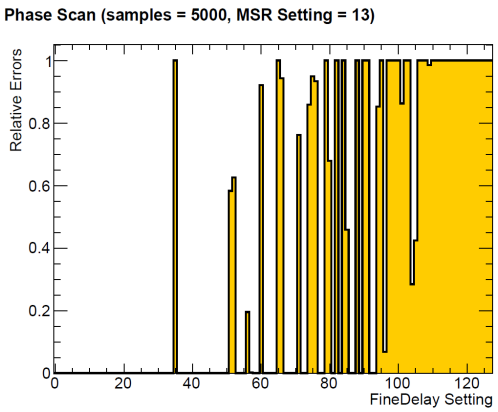

Figure 6.4: Relative error rate of the receiving firmware blocks in relation to the fine delay setting (≈275 ps per step) of the output[52].

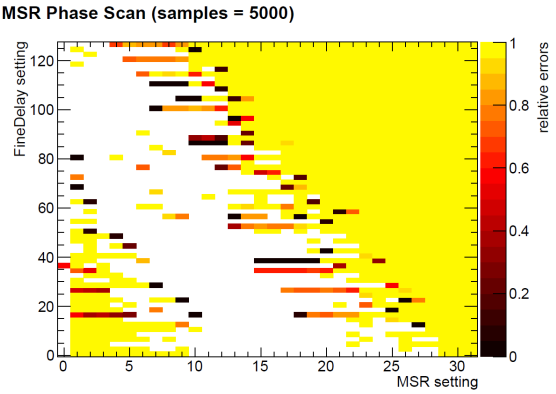

Figure 6.5: Relative error rate for the different combinations of MSR and fine delay setting[52].

System Test

After the basic functionality of the data path firmware blocks had been verified, the test setup was expanded by adding an FE-I4 SCC (Figure 6.6) to test the communication with a real readout chip. The SCC is not equipped with a DORIC[1], so clock and data are transmitted in two individual differential twisted-pair cables to the mezzanine board, and one differential connection is used to transfer data back to the SP605.
At first the registers of the FE-I4 need to be configured with the proper values to enable the chip for the tests. This was performed with the help of the software interface, and the correct configurations were verified by reading the registers after they had been written. Then a first readout test was done by performing a so called *digital scan* in which digital pulses are injected into the readout cells of the FE-I4 before they are read out again to simulate data taking. All cells of the chip were injected with a pulse of the same length before the readout was triggered. The 8b10b encoded data sent from the FE-I4 was received and decoded properly as can be seen in Figure 6.7 which shows the result of the measurement.

[1] The DORIC (see Section 2.3.1) manages the BPM decoding of the data stream from the BOC and passes both the recovered clock and data to the FE-I4 chip.

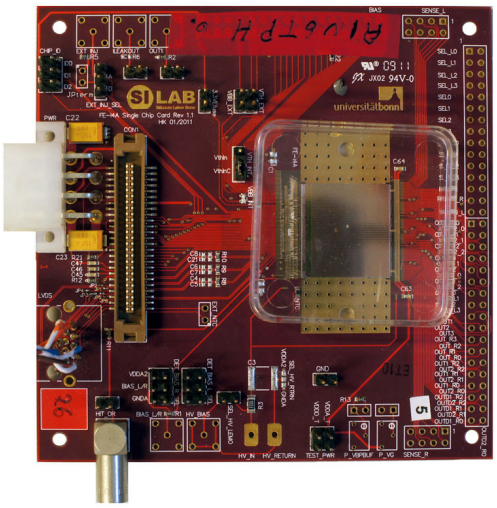

Figure 6.6: Photo of a FE-I4 Single Chip Card.

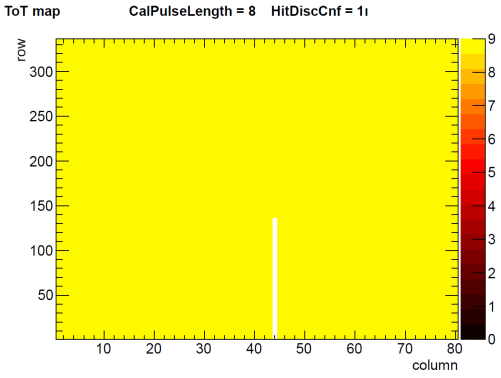

Figure 6.7: Time over Threshold values received during a digital scan of an FE-I4 SCC performed by the firmware running on the evaluation board setup, the missing data in column 44 is caused by a defect of the FE-I4 sample and not by a flaw in the firmware design[52].

6.2 Prototype Testing

The first revision of the Re-BOC had several issues with components (e.g. two diodes placed in the wrong orientation) causing short circuits that had to be located and fixed before the testing could begin. For the second revision all bugs that had been found were fixed and the firmware evaluation and hardware testing that started with the first prototypes is now continued with the five PCBs of the second batch. All measurements performed with the evaluation board setup and the boards of the first revision have been repeated to verify the results with the second revision.

6.2.1 Standalone Tests

At first the basic functionality of all the electric components on the prototypes had to be checked. For this several simple firmware designs were created to test if all FPGAs and their attached peripherals on the PCB are usable. After this was confirmed the evaluation of the Re-BOC firmware blocks and hardware properties that are important for the integration into the ATLAS Experiment started.

Optical Components

The optical interfaces of the Re-BOC are very important and it is crucial that they perform well enough. The SFP and QSFP modules used to establish the S-Link connection to the ROS are commonly used components. In the current system SFP transceivers are also used for the DAQ interface, so no special requirements need

to be fulfilled by them other then working properly. This has been verified on the prototype by connecting the four SFP transceivers of the northern BMF to the QSFP transceiver attached to the southern BMF and transferring test data pattern between them. Checking the received data patterns for errors showed that the transmission with 2 Gbit/s over all four channels worked without problems.
More important is the verification of the functionality of the SNAP12 modules as they replace the custom built plugins used in the current system. Two main concerns need to be checked to proof the feasibility of their integration into the design. At first they have to be able to be operated at the rather low transmission rates of 40 to 160 Mbit/s required by the front-end electronics, and secondly the light output power of the TX modules and the sensitivity of the RX modules need to be sufficient to guarantee stable transmissions for fibers that have been damaged by the radiation.
To test the possible range of transmission rates of the SNAP12 modules, test signals with varying frequencies were generated, and then transmitted via a TX module to a RX module whose output was connected to an oscilloscope[1]. It was found that the transmission works for frequencies down to $\approx$ 300 kHz before errors occur and the connection breaks down[85].

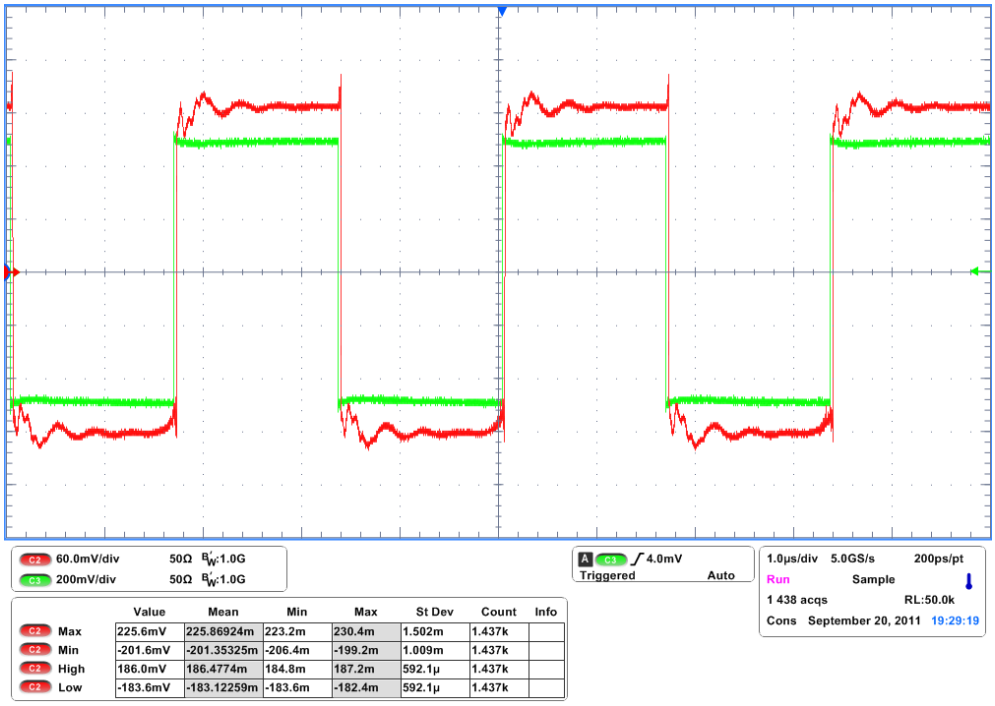

Figure 6.8: Picture of an oscilloscope showing a generated (green) 300 kHz input signal to a SNAP12 TX module and the output (red) of a SNAP12 RX module receiving it[85].

[1] These tests were mainly performed by B. Schneider at the University of Bern, Switzerland.

In Figure 6.8 one of these measurements is shown, the green curve is a generated 300 kHz signal before it is transmitted by the SNAP12 TX module, and the red curve is the output of the RX module receiving this signal. Beside this evaluation of the lower frequency border the target data rates of 40, 80 and 160 Mbit/s were tested and verified without issues.
To estimate the sensitivity of the PiN diode array used in the RX modules a VCSEL array with known changeable output power was employed to measure the amplitude of the outputs of the SNAP12 receiver in relation to the input light power[1]. From the 8 channels of the VCSEL array only 5 were functioning, and each one was used separately for these measurements. RX modules from two different vendors (Reflex Photonics and Tyco Electronics) were evaluated, and the results can be seen in Figure 6.9 and 6.10. Due to the problems of the Reflex Photonics devices at high light power they were excluded from further evaluation, and modules from Avago are now tested instead which proofed to have a very similar behavior than the Tyco receiver.

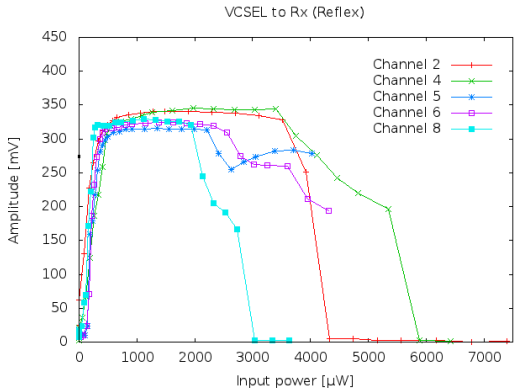

Figure 6.9: Graph of the output amplitude of a SNAP12 RX module from Reflex Photonics in relation to the light input power (measurements by B. Schneider (University of Bern)).

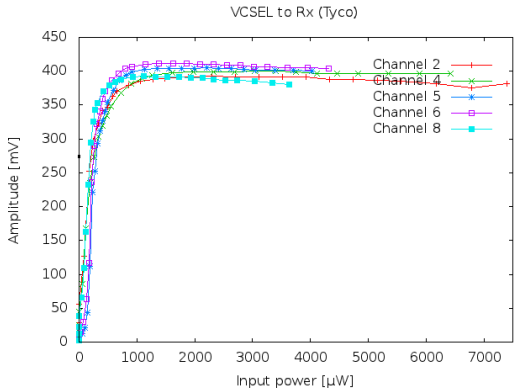

Figure 6.10: Same as Figure 6.9, but for a Tyco Electronics device.

For the measurement of the output power of the SNAP12 TX modules which can be seen in Figure 6.11 an optical power meter was used. The measured lower limit of ≈ 0.65 mW (-1.9 dBm[2]) needs to be above the threshold required by the PiN diodes on the Optoboard (≈ 0.25 mW (-6 dBm)[3]) including the attenuation induced by fibers damaged by the radiation environment. Studies showed that the *Radiation Induced Attenuation* (RIA) that can be expected for the fibers between the Optoboards and the

[1] Measurements were done at the University of Bern by B. Schneider.
[2] The power in dBm is given by $P_{dBm} = 10 \cdot \log_{10} \frac{power}{1mW}$ with $power$ being the one measured and 1 mW as reference power.
[3] This is due to the requirement of an input current of ≈100 μA for the DORIC and the average responsivity of the PiN diodes of ≈0.4 A/W[27, 86].

off-detector electronics is around 0.6 dB[1] which is below the limit of 1 dB that was set for RIA in the fibers[87]. Hence even if the minimum output power of a SNAP12 transmitter gets damped by the maximum attenuation in the fibers there will be enough light power arriving at the Optoboard[80].

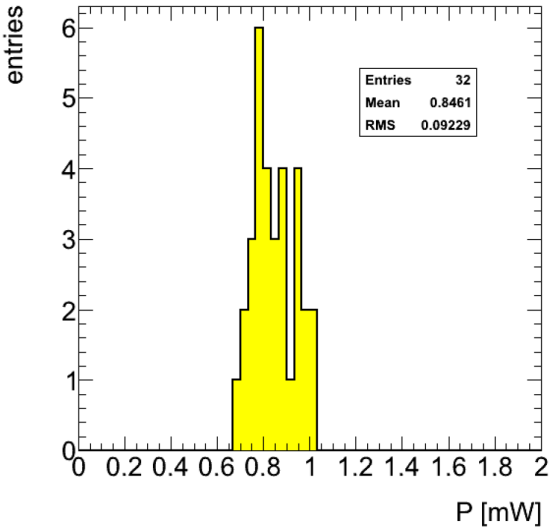

Figure 6.11: Measurement of the output power of the SNAP12 TX modules[80].

Clock Components

For a reliable operation of all the components on the board, stable clock signals are required, thus all clock components on the RE-BOC have been checked. The outputs of all Si532 dual-frequency oscillators on the board were measured with an oscilloscope and the correct frequencies verified (e.g. in Figure 6.12 one of these measurements can be seen). On the prototypes of the first revision fanout buffers with the wrong signal standard (LVPECL) were assembled for the clock signals which required to replace them by buffers with LVDS output in order to use the clocks. The PCBs of the second revision are populated with the correct buffer type.
Programming the PLL to the desired output frequency of 40 MHz was tested with the 40 MHz quartz crystal oscillator on the Re-BOC as input source as the TTC clock was not available during this test. After initial issues with the configuration of the PLL were solved it could be programmed correctly and the output was measured to be stable 40MHz (Figure 6.13).
To be able to test the cascaded delay chips a 40 MHz clock was generated locally and looped back over the VME connector into the TTC clock path on the Re-BOC.

[1] This number is very pessimistic and based on a worst case scenario so the actual attenuation should be lower.

Then the delay was incremented repeatedly by the firmware and the output of the second buffer (delayed clock) was monitored together with the input of the first buffer (fix clock) on an oscilloscope. Changes in the delay settings were clearly visible in the changing phase relation between the two clock signals and therewith the proper function and delay range of the two chips was verified.

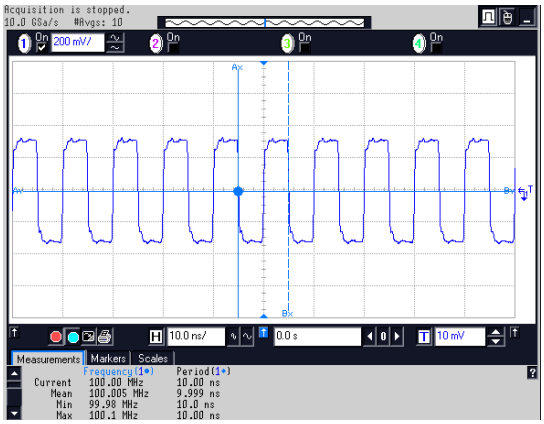

Figure 6.12: Measurement of the 100 MHz clock signal generated by a dual-frequency oscillator (Si532) on the Re-BOC (Measurement by M. Wensing).

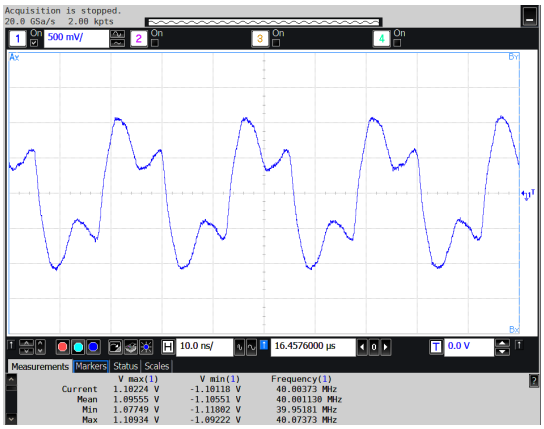

Figure 6.13: 40 MHz clock signal generated by the PLL on the Re-BOC (Measurement by M. Wensing).

6.2.2 System Tests

After the basic functionalities of the Re-BOC prototype had been verified the tests were extended by connecting a FE-I4 SCC, a redesigned ROD or a combination of both. This way system test can be performed in an environment that is very close to the final setup.

FE-I4 Readout

The first approach was to connect the FE-I4 SCC to the Re-BOC in the same way it was attached to the evaluation board to be able to use the same[1] firmware for the tests. Therefore the three differential pair connections to the FE-I4 were established with cables attached to the pin-headers on the Re-BOC (Figure 6.14). Then the FE-I4 chip was configured and its readout cells were injected with digital pulses, this time pulses of varying length (increasing with column number) were used. Everything worked as intended and the result of the readout of the chip can be seen in Figure 6.15.

[1] Including minor changes to adapt the design to the bigger FPGA.

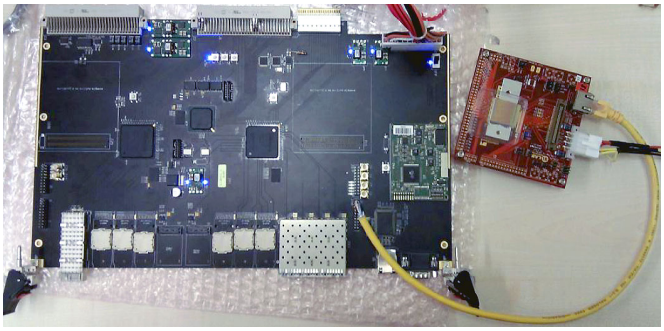

Figure 6.14: Photo of the test setup to readout the FE-I4 SCC via direct electrical connections (Photo by M. Wensing).

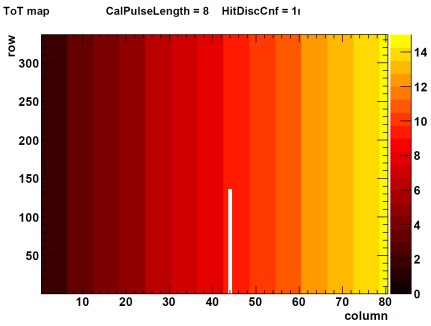

Figure 6.15: Result of the digital scan of the FE-I4 by the Re-BOC prototype, pulses of varying length were used to inject the columns (missing data in Column 44 is caused by an error on the FE-I4).

For a more realistic scenario an Optoboard was integrated into the test setup which allows to also test the TX and RX path of the firmware. The commands to configure, inject and readout the FE-I4 chip were BPM encoded and transmitted optically to the Optoboard. A program of the IBL Re-BOC Library (Section 5.3) was employed that allows to load a bitmap file and write it to the readout cells. The following readout of the cells and the optical data transmission to the Re-BOC worked as well as the sampling and decoding of this data in the corresponding firmware blocks (the result is shown in Figure 6.16). Hence the readout chain between off-detector (Re-BOC) and front-end electronics (Optoboard & FE-I4) has been tested successfully.

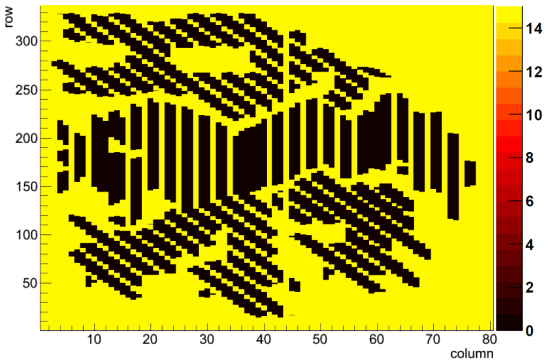

Figure 6.16: Digital scan of a FE-I4 that is connected to the Re-BOC via an Optoboard, the readout cells were injected with the logo of the University of Wuppertal (rotated counter clockwise by 90°)[80].

ROD Prototype

Parallel to the evaluation of the front-end interface the connectivity over the VME connectors to the ROD prototype was tested. For this, both cards were either connected by adapters with VME connectors on each side, which allows to run the tests on a table, or both cards were plugged into a VME crate. While testing on a table has the advantage of easy access to both boards to supervise and measure their behavior, the VME crate provides a test environment equivalent to the system both cards are going to be installed in.

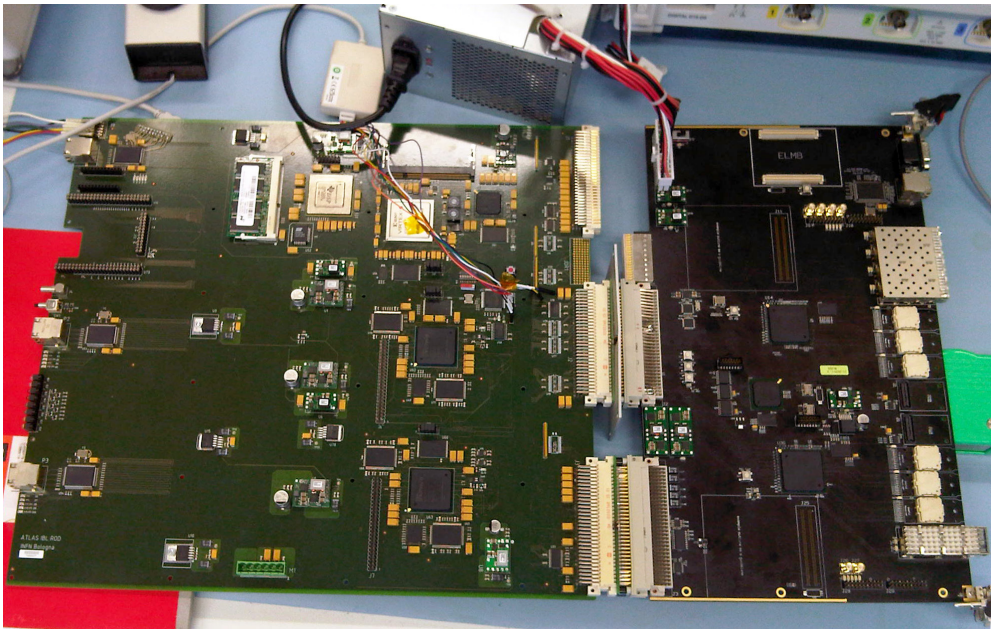

Figure 6.17: Connectivity test setup with prototypes of the first revisions of ROD (left) and Re-BOC (right). Picture taken by A. Kugel.

The main task in this tests was to evaluate the connectivity between both prototypes and to validate that the transmission of data with 80 MHz via the VME connectors is feasible. This also includes measuring the signal quality of the SSTL-3 standard which is intended to be used for the interface between Re-BOC and ROD. Hence a test design was implemented that generates dummy data that is then being transferred via the VME connectors to the other board where it is checked for errors. This test was run with different frequencies (10, 40, 80 and 100 MHz) for the transfer and all connections between the boards were evaluated in both directions, from Re-BOC to ROD and vice versa.
At the target speed of 80 MHz more than 2 Terabyte of data was transmitted between both cards without a single error which corresponds to a bit error rate of less than $5 \cdot 10^{-14}$[80]. Even though the connection worked without special termination of

the signal lines it was tested if the signal quality could be improved by termination. For this the internal on-chip termination feature of the Spartan-6 was employed to implement a split termination (UNTUNED_SPLIT_50) with a parallel equivalent resistance of 50 Ω (see Figure 4.10.d) on the receiving side (on the ROD). Measurements with this termination showed an improved signal integrity, but at the same time the power consumption and temperature of the receiving FPGA increased. The current flow caused by this termination can be estimated with the Thevenin voltage and the termination resistance to $I_{term} = \frac{V_T}{R_{Term}} = \frac{0.5 \cdot 3.3}{50} A = 33mA$ per terminated IO pin. That this causes serious issues was learned the hard way as during initial tests 80 IOs of a Spartan-6 on the ROD prototype were set to internal termination which caused the FPGA to overheat and die. The following Table 6.1 shows the results of the measurements which were done to examine this behavior. The current increase caused by different numbers of terminated IO pins was measured and the temperature of the FPGA casing was supervised with an infrared camera[1].

Terminated IOs [#]	Current increase [mA]	FPGA Temperature [°C]
0	0	28.4
1	30	34.7
10	310	60.0
20	610	75.0
40	1230	95.0
48	1430	>100
80	2460*	>125*

Table 6.1: Current and temperature increase for different amounts of terminated (UNTUNED_SPLIT_50) SSTL-3 IOs. *The values for 80 IOs is only an estimate, during this configuration the Spartan-6 died.

The current increase is consistent with the calculated estimate and the rapidly increasing temperature explains the death of the Spartan-6 in the initial test as its maximum temperature rating is 125 °C[76]. Even though these measurements were performed without a heat sink attached to the FPGA it is obvious that care needs to be taken when large numbers of transmission lines need to be terminated. Therefore the second revision (RevB) of the Re-BOC features the option to implement external termination resistors close to the FPGA to reduce the strain put on the chip.
Besides these connectivity tests the Setup-Bus interface from ROD to Re-BOC was implemented and tested. It was possible to access all the registers on the Re-BOC from a computer connected to the ROD which is a scenario in accordance to the final system.

[1] Measurements performed by D. Falchieri at the University of Bologna and INFN, Italy.

7

Outlook & Summary

At the time this thesis is written the evaluation and testing of the Re-BOC prototypes is still ongoing. A system test including all components involved in the final readout scenario is setup at CERN by equipping a VME crate with TIM, SBC and the prototypes of ROD and Re-BOC. The number of FE-I4 chips connected and readout will be increased until the full design utilization is reached to test the readout hardware under full load. Also the third revision of both prototype boards which includes minor fixes and improvements is currently being prepared for production and will be at hand and available for evaluation soon. In the following sections the schedule for the installation of the IBL into the ATLAS Detector will be given followed by a glimpse on possible upgrades of the off-detector electronics in the future and the concluding summary.

7.1 IBL Installation and LHC Schedule

If there are no major show-stopping flaws in the third revision of the Re-BOC prototype it will be used for the production of the final batch of boards for the installation in the counting room at Point 1. The schedule for the next year (2013) foresees the construction of the IBL in SR-1 which is a building with a large 700 m^2 clean room on the surface above the ATLAS detector. By the middle of the year the full setup of IBL and off-detector electronics is supposed to be fully assembled for complete power and readout tests. By the end of 2013 the system is planned to be ready for the installation of the IBL which will be either performed in SR-1 in case the Pixel Detector is extracted out of the ATLAS detector system and brought to the surface or in situ in UX-15 if it is not extracted. The fully tested off-detector readout hardware will be moved from SR-1 to the counting room in USA-15 during the first two months of 2014. There final tests will be performed to have a commissioned system as soon as

possible to integrate the detector into the tests. By mid 2014 the complete detector and readout system should be fully cabled, tested and commissioned to allow the closing of the Inner Detector of the ATLAS experiment. At the end of the year 2014 all upgrade operations have to be complete to allow the LHC to go into operation again. Data taking will continue with this system until 2018 when the next long shutdown is scheduled and further updates and upgrades will be installed. Finally around 2022 the Long Shutdown-3 will be used to upgrade the LHC to the High-Luminosity LHC which will include the replacement of the complete Inner Detector and potentially other sub-detector systems. An overview of the schedule for the operation of the LHC can be found in Figure 7.1.

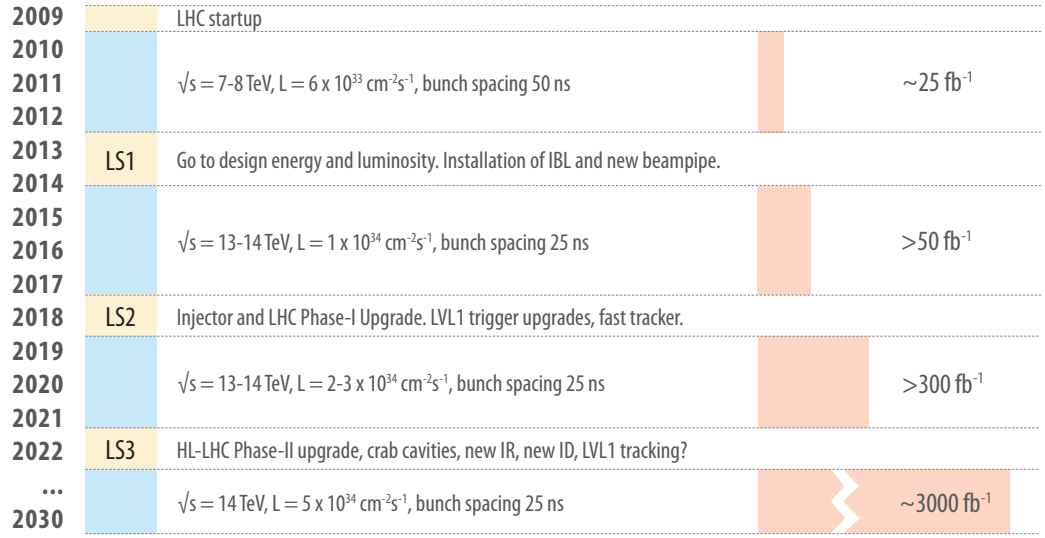

Figure 7.1: Schedule for the long shutdowns (LS) of the LHC and the foreseen upgrades, including estimates for the center-of-mass energy $\sqrt{s}$, luminosity L and the integrated luminosity in fb^{-1}[88, 89].

7.2 Future Readout Concept

It is foreseen that the HL-LHC upgrade with the replacement of the Inner Detector will also require another upgrade of the readout electronics. The readout solution based on the ATCA system which was introduced in Section 3.2 might be a good candidate for this and is therefore extensively studied. Another approach[1] integrates both the BOC and the ROD functionality on a PCIe board which could be housed in standard commodity PCs. A sketch of this concept can be seen in Figure 7.2 for

[1] This is only a concept which is not in development so far.

which the dimensions of a current high-end graphics board were used as baseline for the maximum PCB size that still fits into a regular PC case.

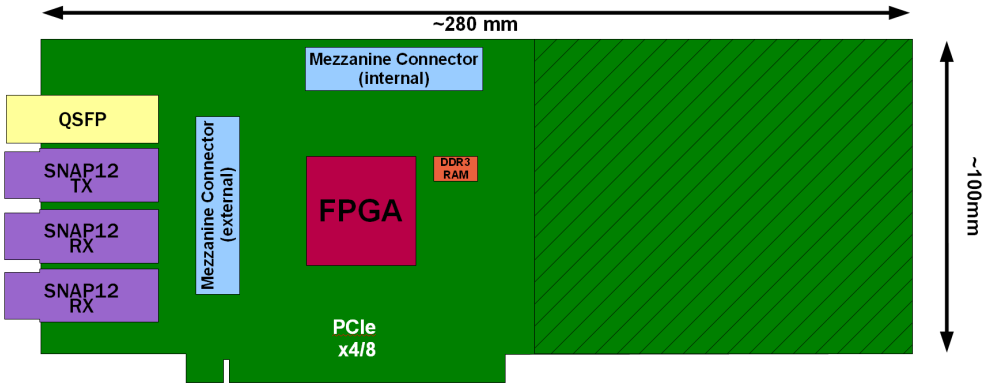

Figure 7.2: Conceptual drawing of a PCIe board combining the functionalities of BOC and ROD. As baseline dimension a high-end graphics adapter was used to estimate the maximum available space which results in the so far unused hatched area. While the internal mezzanine connector can be used to interface a TTC adapter card in the same computer the external one can be used to add more optical interfaces if required.

For the BOC functionality two SNAP12 RX and one TX module interface to the detector front-end while one QSFP transceiver provides the two dual S-Link connections to FTK and ROS. The data path logic which handles both the former BOC and ROD tasks is integrated in one Xilinx FPGA from the 7-Series families (Artix, Kintex or Virtex) or in their potential succeeding models. In Table 7.1 the top models of these FPGA families are compared with the Spartan-6 which is used as BMF[1] on the Re-BOC.

	Spartan-6	Artix-7	Kintex-7	Virtex-7
Logic Cells	148k	215k	478k	1,955k
Block RAM [Mb]	4.8	13	34	68
DSP Slices*	180	740	1,920	3,600
High Speed Transceiver	8	16	32	96
Line Rate** [Gbit/s]	3.2	6.6	12.5	28.05
PCIe Interface	1x (v1.1)	x4 (v2)	x8 (v2)	x8 (v3)
IO Pins	540	500	500	1200

Table 7.1: Comparison of the main features of several FPGA families (their top-of-the-line models)[63, 90]. *The DSP slices of the 7 Series FPGAs (DSP48E1) are more advanced than the ones from the Spartan-6 (DSP48A1), e.g. the maximum operation frequency is now 741 MHz compared to the former 390 MHz. **Line Rates of the high-speed transceivers, the numbers for the 7-Series should be taken with a pinch of salt as the actually achievable speeds are usually lower than the ones announced by the marketing department.

[1] This is actually the top model from the Spartan-6 familiy as well (LX150T).

While the event fragment building and histogramming is performed locally in the FPGA more complex fitting algorithms that are needed during detector calibration can be outsourced via the PCIe interface to either the CPU or GPU of the host computer to speed up the calculations[1]. The connection to the control room is established by the network interface of the computer which also performs all the functionalities of the SBC. To replace the TIM an adapter card based on TTCrx chips[58] receives the TTC signals and distributes them to the mezzanine connectors of the data path boards in the computer. With commercial mainboards providing up to 7 PCIe slots several configurations are possible, without a dedicated graphics adapter (which would occupy two slots) up to 6 of the readout cards could be combined with one TTC interface card yielding a system with the same number of channels as 3 Re-BOC & ROD pairs. With a double-slot GPU 4 readout cards could be used (replacing 2 VME card pairs) so a total of 5 to 7 of these readout computers would be necessary for the readout of the IBL (compare with Table 4.1).
A lot of potential space on the PCB is unused so far (the hatched area in Figure 7.2) if such dimensions would be used which could be utilized to add additional features. An interesting approach would be to not only integrate BOC and ROD, but also the ROBIN[45] (*Read-Out Buffer IN*) on this card. The task of the ROBINs is to receive the event fragment data from the ROD via the S-Link connection on the BOC, buffer this data locally and either send it to the DAQ if requested or delete it if not. Hence another FPGA is utilized to manage the ROBIN related data handling. It receives the event data directly via high-speed on-board links instead of an S-Link connection and buffers it in fast DDR DRAM. Thus the QSFP module for the external S-Link can be replaced[2] by a 10 Gigabit-Ethernet connection to the DAQ system. This concept is illustrated in Figure 7.3.

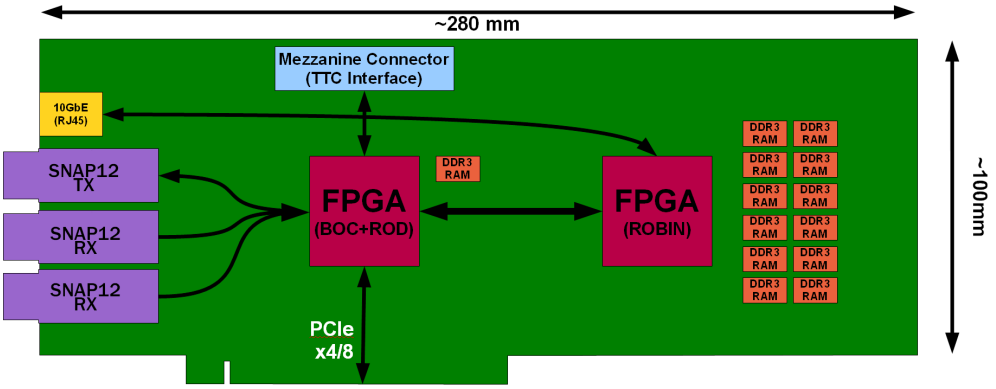

Figure 7.3: Sketch of a PCIe concept card integrating BOC, ROD and ROBIN on a single device.

[1] While the current ROD is equipped with discrete DSPs for these tasks the redesigned IBL-ROD will send the data via Gigabit-Ethernet connections to computer farms which will handle the fitting.
[2] As a drawback thereof the connection to the FTK would have to be implemented in a different way.

7.3 Summary

In this thesis the development of the Re-BOC for the readout of the IBL for the Pixel Detector of the ATLAS Experiment at CERN was motivated and described from scratch to the first system tests with the prototypes. Even though these tests are still ongoing their preliminary results indicate that the Re-BOC performs as intended. The proper functionalities of all the firmware blocks necessary for the design (except the S-Link implementation) have been verified in standalone- and system tests. The S-Link was successfully simulated on the Spartan-6 architecture and is currently being integrated into the CERN S-Link repository, so this will be no show-stopper. Further testing and work on the firmware design will be done to optimize its behavior, for example the duty-cycle distortion introduced by the fine-delay implementation in the TX-Path will be extensively studied.
The conceptual ideas of replacing all custom made parts by commercially available components, simplifying the clock distribution and reducing the utilized signal standards to a minimum were successfully integrated into the design of the Re-BOC. The commercial optical components used on the Re-BOC perform very well and together with the flexibility of the FPGAs many options for possible upgrades are provided. So far the backwards compatibility between the Re-BOC and a current ROD was not evaluated as the schedule until the installation of the IBL is rather tight. If the compatibility is proven additional Re-BOCs will be produced to gradually replace the current BOCs to avoid the need for another batch of the custom made optical plugins. Overall the Re-BOC will be used in the ATLAS Experiment for about 10 years until 2022 when a completely new readout system will be installed.
With the fast advancing technological developments in the areas of FPGAs and optical transceivers this new system will be considerably more integrated than the current one. An idea of such an implementation was given in the previous section. But at the time the development thereof starts many possible options will be available to choose from.

Appendix

A Physics

A.1 Particles and Forces of the Standard Model

This is only a very brief overview of the particles and the forces affecting them as described in the Standard Model.

In the standard model, matter consists of 12 elementary particles which interact by 3 different forces that are mediated by force carriers. All of these 12 elementary particles have a spin of $\frac{1}{2}$ and are thus called *fermions*, while the force carriers have spin 1 and are known as *gauge bosons*. Six of the fermions, the *quarks*, interact by their *color charge*, which is mediated by *gluons* as force carriers of the *strong interaction*. Only quarks and gluons have a color charge, so the other half of the fermions is called *leptons*. They are subdivided into particles with and without an electric charge (*neutrinos*). While the particles with electric charge interact by the *electromagnetic force* which is mediated by *photons*, the neutrinos are only influenced by the *weak interaction* with the *Z and W bosons* as force carriers. As the name implies the weak interaction is the weakest of the three forces mentioned here, the electromagnetic force is about eleven magnitudes (100 billion times) stronger and the strong interaction is with two additional magnitudes the strongest (but the interactions differ in range and behavior).

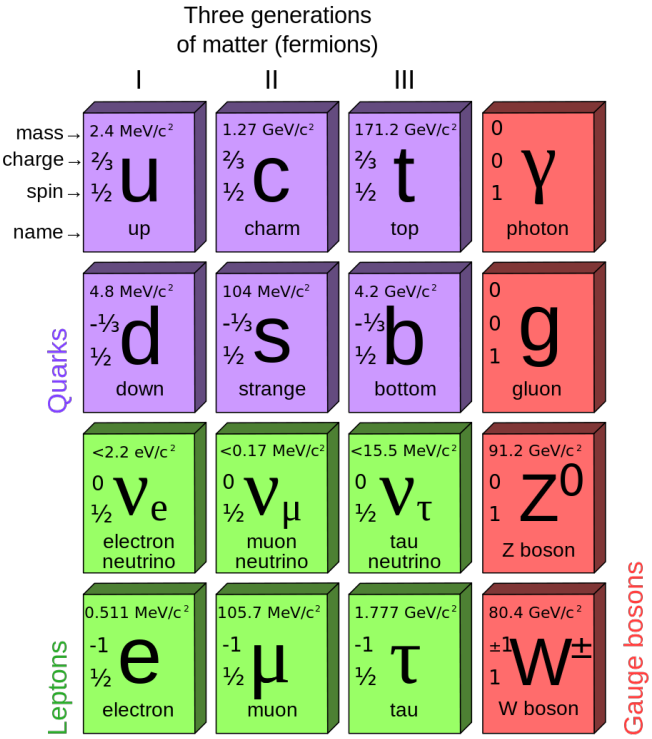

Figure A.1: Elementary particles and gauge bosons of the Standard Model, picture from Wikipedia.

A.2 Synchroton Radiation

Synchroton radiation are electromagnetic waves emitted by charged relativistic (v≈c) particles that are accelerated perpendicular to their direction of motion, e.g. when moving through a magnetic field. The energy that a particle with the charge Z·e loses due to this radiation during one circulation in e.g. a ring accelerator is given by equation 1.

$$\Delta E = \frac{(Ze)^2 \cdot \beta^3 \cdot \gamma^4}{\epsilon_0 \cdot 3R} = \frac{(Ze)^2}{\epsilon_0 \cdot 3R} \cdot \frac{E^4}{(m_0 c^2)^4} \tag{1}$$

$\beta = \frac{v}{c} \approx 1$, $\gamma = \frac{1}{\sqrt{1-\beta^2}} = \frac{E}{m_0 c^2}$ (Lorentz factor), ϵ_0: vacuum permittivity and R: radius of the ring

A.3 Lorentz Force

A charged particle or generally a charge *q* traversing an external electric (*E*) and magnetic (*B*) field with the speed *v* is affected with the Lorentz Force which is given by equation 2 ($\times$ being the vector cross product).

$$F = q(E + v \times B) \tag{2}$$

A.4 Bethe-Bloch Formula

The amount of energy *dE* a particle[1] with the charge *z* (in multiples of the elementary charge *e* = 1.602 x 10^{-19}C) loses while traversing with speed *v* the distance *dx* through a material with the density ρ and the atomic number (Z) and weight (A) is given by the Bethe-Bloch Formula:

$$-\frac{dE}{dx} = 2\pi N_a r_e^2 m_e c^2 \rho \frac{Z}{A} \frac{z^2}{\beta^2} \cdot \left[ln\left(\frac{2m_e \gamma^2 v^2 W_{max}}{I^2} - 2\beta^2 - \delta - 2\frac{C}{Z} \right) \right] \tag{3}$$

with:
N_a = 6.022 x 10^{23}mol^{-1} (Avogadro number)
r_e = 2.817 x 10^{-15}m (classical electron radius)
m_e = 511 keV (electron mass)
ρ = 2.33 g cm^{-3} (density of the absorbing matter, value given here is for silicon)
I ≈ 173 eV (mean effective ionization potential, value for silicon)
$\beta = \frac{v}{c}$
$\gamma = \frac{1}{\sqrt{1-\beta^2}}$
δ: density correction
C: shell correction
W_{max}: maximum energy transfer in a single head-head collision

More details can be found for example in [91].

[1] This formula is valid for protons, alpha particles and atomic ions, for electrons the energy loss is different.

B Additional Pictures

In this appendix pictures are shown that provide additional and interesting information but would have taken too much space in the main part of the thesis. Links in the captions of the pictures allow to "jump" back to the part of the thesis where this picture was mentioned to avoid tiresome scrolling. Needless to say that this only works in the pdf and not in the printed version.

B.1 The ATLAS cavern

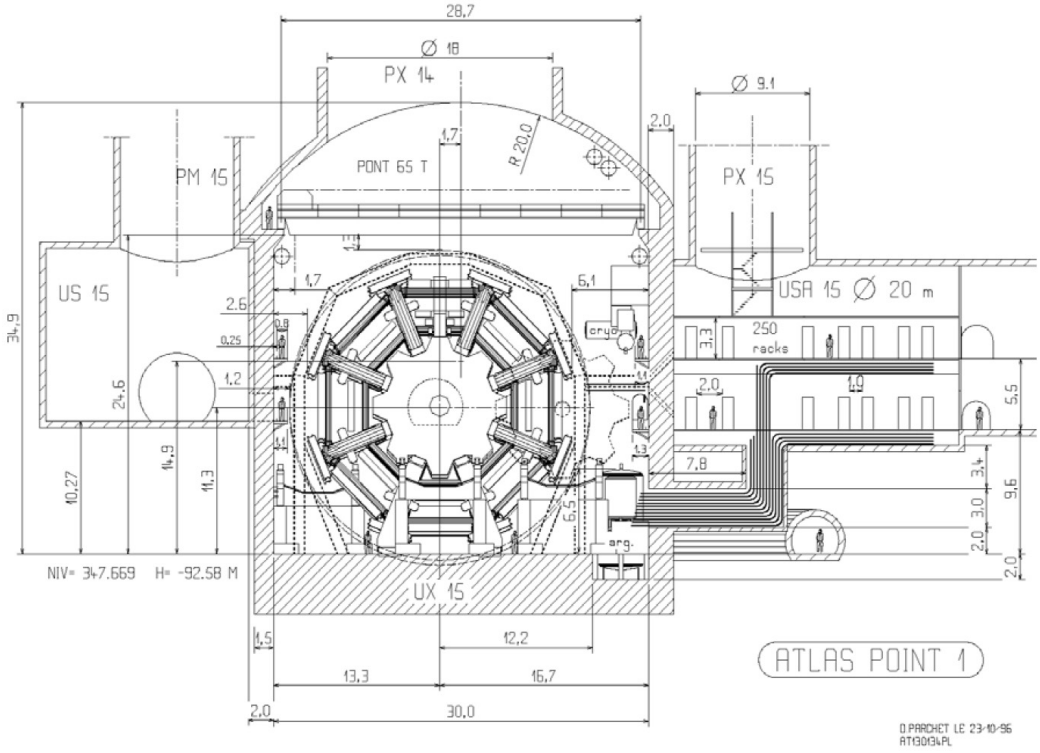

Figure B.2: Cross section of the ATLAS caverns, most important the experimental cavern UX15 and the service area USA15[92]. To get back to the Off-Detector Electronics Chapter click here.

B.2 The TTC Interface Module (TIM)

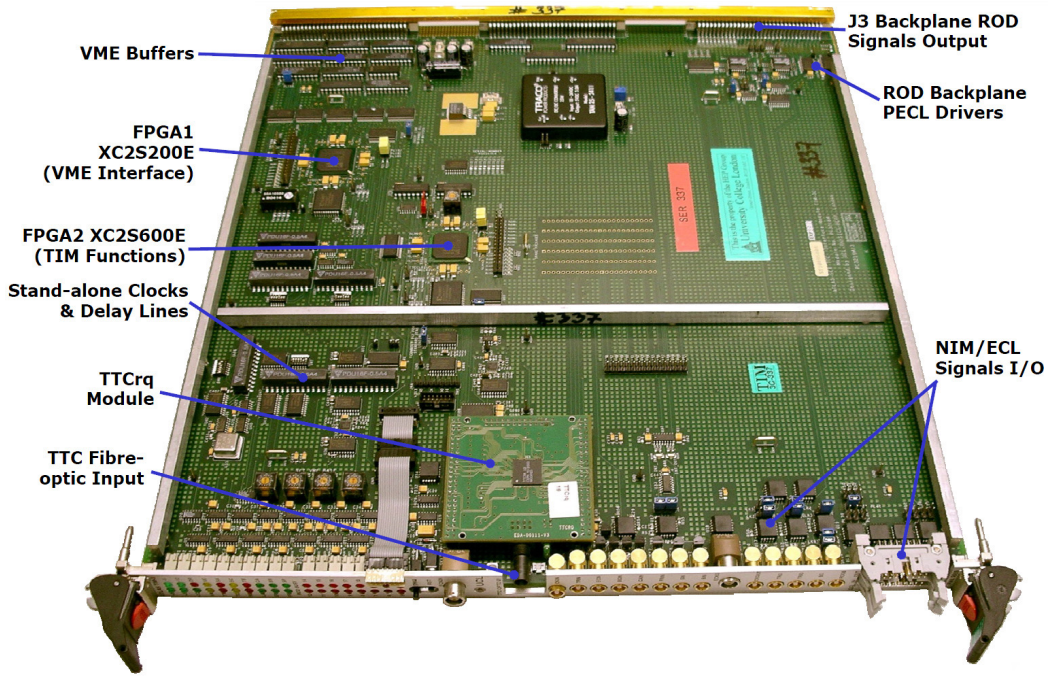

Figure B.3: Photo of the TIM-3C with annotations for its main components[93]. To get back to the TIM section click here.

B.3 Spartan-6 FPGA Slice

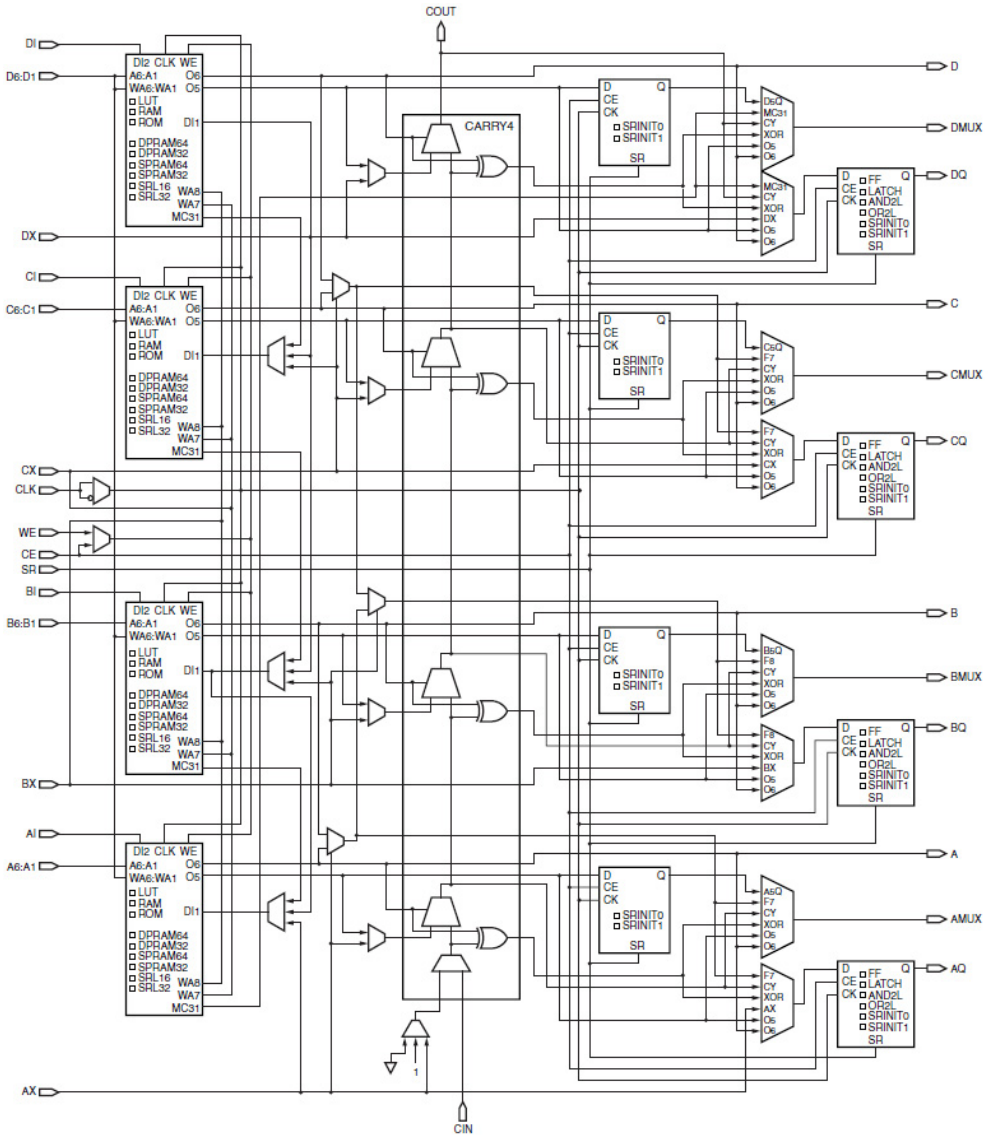

Figure B.4: Diagram of the blocks of the most complex slice (SLICEM) in a Spartan-6 FPGA [62]. To get back to the FPGA architecture section click here.

B.4 BOC Clock Section

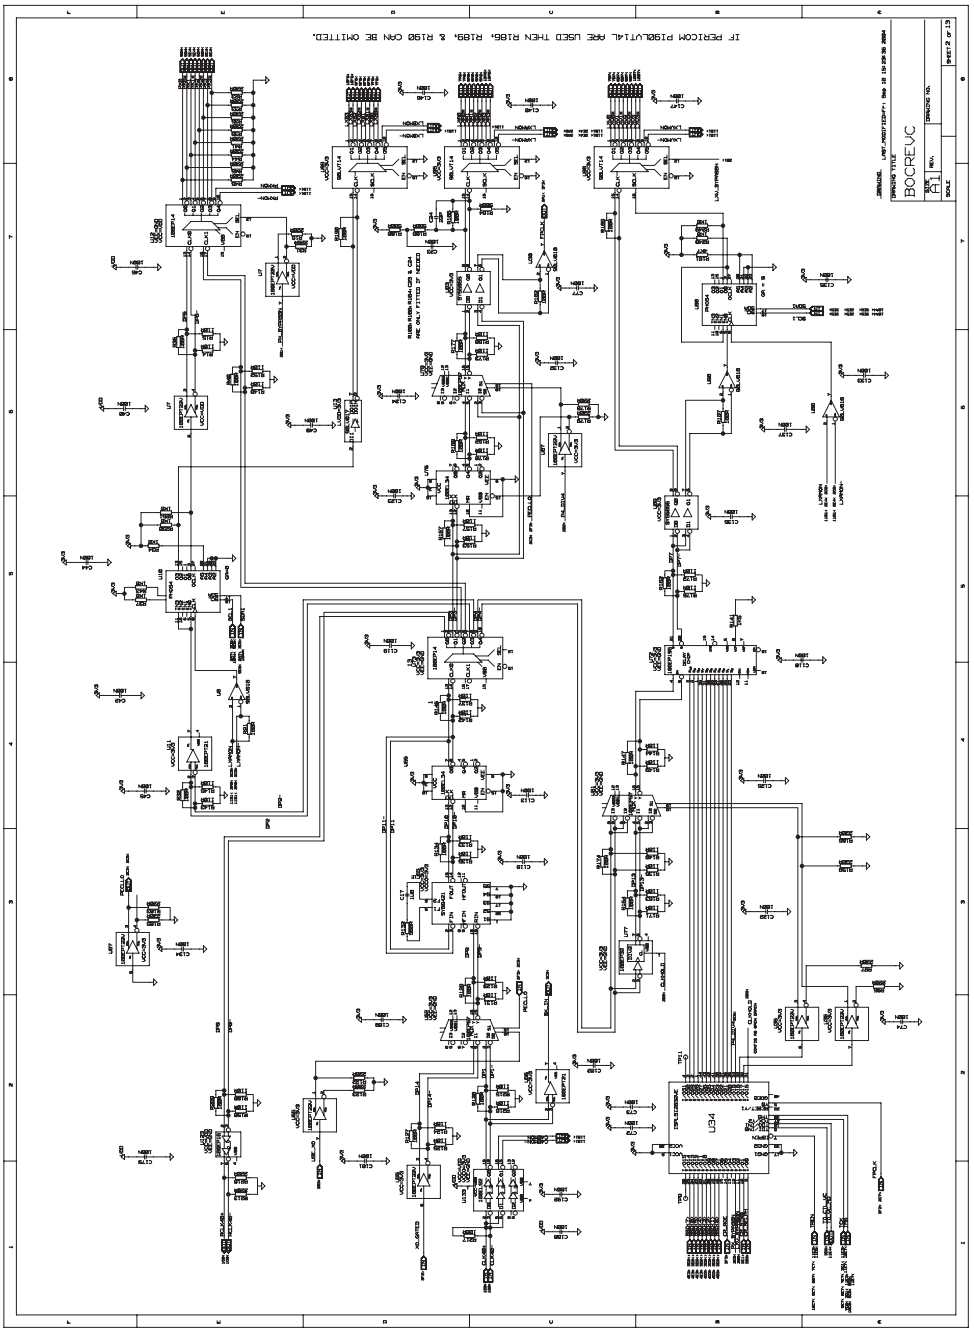

Figure B.5: Schematics of the clock section of the current BOC. To get back to the Clock Components section click here.

References

[1] CERN. *CERN experiments observe particle consistent with long-sought Higgs boson*, **PR17.12**. CERN Press Release, April 2012. 1

[2] M J Herrero. *The Standard Model*, **arXiv:hep-ph/9812242**. Cornell University Library, December 1998. 2

[3] P W Higgs. *Broken Symmetry And The Masses Of Gauge Bosons*, **Phys. Rev. Lett. 13, 508 - 509 (1964)**. American Physical Society, October 1964. 2

[4] Oxford Dictionaries. *Oxford Dictionary : data*. Oxford Dictionary. 2

[5] IEEE Computer Society. *IEEE Computer Society Style Guide*. IEEE. 2

[6] ATLAS Collaboration. *Event display of a H -> 2e2mu candidate event*, **ATLAS-PHO-COLLAB-2012-011**. CERN Document Server, April 2012. 3, 127

[7] Maximilien Brice. *Aerial View of the CERN taken in 2008*, **LHC-MI-0807031**. CERN Document Server, July 2008. 5, 127

[8] The Royal Swedish Academy of Sciences. *The Nobel Prize in Physics 1984*. Nobel Prize Press Release, 1984. 6

[9] CERN. *CERN Accelerator Complex*. CERN / Wikipedia, May 2011. 7, 127

[10] O Brüning, P Collier, P Lebrun, S Myers, R Ostojic, J Poole, and P Proudlock. *LHC Design Report v2*, **CERN-2004-003-V-2**. CERN Document Server, 2004. 7, 8, 127

[11] The CMS Collaboration. *CMS Experiment Homepage*. CERN. 8

[12] CMS Collaboration. *CMS public website*. CERN. 9, 127

[13] The TOTEM Collaboration. *TOTEM Experiment Homepage*. CERN. 9

[14] TOTEM Collaboration. *The TOTEM experiment at the CERN Large Hadron Collider*, **JINST 3 S08007**. Jinst - Institue of Pysics Publishing and SISSA, August 2008. 10

[15] ALICE Collaboration. *ALICE website*. CERN. 10, 127

[16] The LHCb Collaboration. *LHCb Experiment Homepage*. CERN. 11

[17] The ATLAS Collaboration. *ATLAS Experiment Homepage*. CERN. 13

[18] Joao Pequenao. *Computer generated image of the whole ATLAS detector*, **CERN-GE-080301201**. CERN Document Server, March 2008. 13, 127

[19] Andrei D Sakharov. *Violation of CP in variance, C asymmetry, and baryon asymmetry of the universe*, **34(5):392**. Soviet Physics Uspekhi, 1991. 14

[20] Joao Pequenao. *Computer generated image of the ATLAS Muons subsystem*, **CERN-GE-080301701**. CERN Document Server, March 2008. 15, 127

[21] Joao Pequenao. *Computer generated image of the ATLAS calorimeter*, **CERN-GE-080301501**. CERN Document Server, March 2008. 16, 127

[22] Joao Pequenao. *Computer generated image of the ATLAS inner detector*, **CERN-GE-0803014**. CERN Document Server, March 2008. 17, 127

[23] Joao Pequenao. *Computer generated image of the Pixel, part of the ATLAS inner detector*, **CERN-GE-080301302**. CERN Document Server, March 2008. 18, 127

[24] Universität Bonn. *ATLAS Pixel Detector Web Page*. Universität Bonn. 20, 127

[25] M Gilchriese et al. *ATLAS pixel detector electronics and sensors*, **JINST 3 P07007**. Jinst - Institue of Pysics Publishing and SISSA, 2008. 19, 22, 127

[26] R Beccherle et al. *MCC: The Module Controller Chip for the ATLAS Pixel Detector*, **492:117-133**. Nucl. Instrum. Methods Phys. Res. A, October 2002. 20

[27] K E Arms et al. *ATLAS pixel opto-electronics*, **554:458-468**. Nucl. Instrum. Methods Phys. Res. A, 2005. 20, 25, 26, 103, 128

[28] K Einsweiler. *The ATLAS Pixel Detector*. LBNL instrumentation Colloquium, 2005. 21, 127

[29] R Wunstorf. *Radiation tolerant sensors for the ATLAS pixel detector*, **466:327-334**. Nucl. Instrum. Methods Phys. Res. A, 2001. 22, 127

[30] T Flick. *Studies on the Optical Readout for the ATLAS Pixel Detector*, **PhD Thesis**. Bergische Universität Wuppertal, July 2006. 23, 41, 44, 128

[31] TU Dortmund. *ATLAS Pixel Detector Front End Electronics Web Page*. TU Dortmund. 24, 25, 128

[32] I Perić et al. *The FEI3 Readout Chip for the ATLAS Pixel Detector*, **565:178-187**. Nucl. Instrum. Methods Phys. Res. A, September 2006. 24, 128

[33] M L Chu et al. *The off-detector opto-electronics for the optical links of the ATLAS Semiconductor Tracker and Pixel detector*, **530:293-310**. Nucl. Instrum. Methods Phys. Res. A, June 2004. 27, 43, 92

[34] ATLAS Collaboration. *ATLAS Insertable B-Layer Technical Design Report*, **ATLAS TDR 019**. CERN, September 2010. 27, 29, 32, 33, 34, 40, 55, 128

[35] M Barbero et al. *A new ATLAS pixel front-end IC for upgraded LHC luminosity*, **604:397-399**. Nucl. Instrum. Methods Phys. Res. A, June 2009. 28, 128

[36] A Clark and G Mornacchi. *ATLAS B-Layer Task Force Final Report*, **Internal Report ATL-GEN-2005-001**. CERN, September 2009. 29

[37] FE-I4 Collaboration. *The FE-I4A Integrated Circuit Guide*, **Version 11.3**. CERN, March 2011. 30, 31, 128

[38] M Barbero. *FE-I4 pixel readout chip and IBL module*, **ATL-INDET-SLIDE-2011-303**. CERN, June 2011. 30

[39] M Mathes. *Development and Characterization of Diamond and 3D-Silicon Pixel Detectors with ATLAS-Pixel Readout Electronics*, **BONN-IR-2008-15**. Bonn University, December 2008. 31

[40] H Kagan. *Diamond radiation detectors may be forever!*, **546:222-227**. Nucl. Instrum. Methods Phys. Res. A, July 2005. 31

[41] P Grenier. *Pixel Sensors for ATLAS Upgrades*, **ESTB-ESA Workshop**. SLAC, March 2011. 32, 128

[42] F Djama. *Overview of the ATLAS Insertable B-Layer (IBL) Project*, **ATL-INDET-SLIDE-2012-430**. CERN, July 2012. 33, 128

[43] K K Gan, H P Kagan, R D Kass, J Moore, D Pignotti, S Smith, P Buchholz, A Wiese, and M Ziolkowski. *Radiation-Hard ASICs for Optical Data Transmission in the First Phase of the LHC Upgrade*, **NSS2011**. Ohio State University, October 2011. 33, 128

[44] K K Gan, H P Kagan, R D Kass, J Moore, D Pignotti, S Smith, P Buchholz, A Wiese, and M Ziolkowski. *Status of On-Detector Opto-Links*, **IBL General Meeting**. Ohio State University, October 2011. 34, 128

[45] A Kugel. *The ATLAS ROBIN - A High performance Data-Acquisition Module*, **PhD Thesis**. Universität Mannheim, June 2009. 36, 112

[46] T Vickey. *A Read-out Driver for Silicon Detectors in ATLAS*, **ATL-INDET-SLIDE-2009-345**. CERN, November 2009. 36, 39, 128

[47] Concurrent Technologies Inc. *Technical Reference Manual for VP 110/01x VME Pentium III-M Single Board Computer*, **550 00 14 Rev 03**. Concurrent Technologies Inc., December 2003. 37

[48] J Joseph, R Jared, D P Ferguson, A Korn, and T Vickey. *ATLAS Silicon ReadOut Driver(ROD) Users Manual*, **v1.64**. CERN, October 2008. 38

[49] T Vickey. *A Read-out Driver for Silicon Detectors in ATLAS*. CERN, September 2005. 39, 128

[50] T Flick. *Bergische Universität Wuppertal - Projects (webpage)*. Bergische Universität Wuppertal. 41, 128

[51] ELMB collaboration. *ELMB homepage*. CERN. 42

[52] T Heim. *Design and development of the IBL-BOC firmware for the ATLAS Pixel IBL optical datalink system*, **Diploma Thesis**. Bergische Universität Wuppertal, September 2011. 44, 87, 97, 98, 99, 100, 101, 128, 129, 130

[53] T Toifl. *4 Channel Programmable Delay Generation ASIC*, **LHCb Week**. CERN, February 2000. 44

[54] E van der Bij and S Haas. *CERN S-LINK homepage*. CERN. 45, 46, 129

[55] PCI Industrial Computer Manufacturers Group (PICMG). *ATCA Specifications and Resources - webpage*. PICMG. 48

[56] PCI Industrial Computer Manufacturers Group (PICMG). *ATCA PICMG 3.0 Short Form Specification*. PICMG, January 2003. 48, 129

[57] R Bartoldus, R Claus, A Haas, G Haller, R Herbst, M Huffer, M Kocian, E Strauss, S Dong, M Wittgen, E Devetak, D Puldon, D Tsybychev, B Blair, and J Zhang. *High Bandwidth DAQ R&D for ATLAS Upgrade*, **v1.30**. SLAC, February 2011. 49, 50, 51, 52, 53, 129

[58] P Moreira, J Christiansen, A Marchioro, and T Toifl. *TTCrx Reference Manual*, **v3.9**. CERN, October 2004. 50, 112

[59] S Dong and M Huffer. *Prospect of SCT Readout Replacement with RCE/ATCA*, **SCT DAQ Upgrade for High-μ**. CERN, May 2012. 52, 129

[60] J Dopke. *Commissioning of the ATLAS pixel detector optical data transmission and studies for readout of the ATLAS IBL and future trackers*. University of Wuppertal. 56

[61] ATLAS FTK Collaboration. *FTK: Fast TracKing for the ATLAS Trigger*, **ATL-DAQ-SLIDE-2011-260**. CERN, June 2011. 56

[62] Xilinx. *Spartan-6 FPGA Configurable Logic Block - User Guide*, **UG384 v1**. Xilinx, February 2010. 59, 121, 129, 131

[63] Xilinx. *Spartan-6 Family Overview*, **DS160 v2.0**. Xilinx, October 2011. 60, 111, 134

[64] Xilinx. *Spartan-6 FPGA Packaging and Pinouts - Product Specification*, **UG385 v2.2**. Xilinx, August 2011. 61, 129

[65] Xilinx. *Spartan-6 FPGA SelectIO Resources - User Guide*, **UG381 v1.4**. Xilinx, December 2010. 61, 129

[66] SNAP12 Membership. *SNAP12 Specifications - Appendix to SNAP12 Multi-Source Agreement*, **1.1**. SNAP12 Membership, May 2002. 62

[67] NXP Semiconductors. *I^2C-bus specification and user manual*, **UM10204 Rev.4**. NXP Semiconductors, February 2012. 62

[68] Reflex Photonics Inc. *Interboard 40 Gbps SNAP12 Parallel Fiber Optic Transmitter and Receiver*, **Rev. 3.6**. Reflex Photonics Inc., December 2010. 63, 129

[69] SFF Commitee. *Specification for QSFP (Quad Small Formfactor Pluggable) Transceiver*, **INF-8438i Rev. 1.0**. SFF Commitee, November 2006. 63

[70] SFF Commitee. *Specification for SFP (Small Formfactor Pluggable) Transceiver*, **INF-8074i Rev. 1.0**. SFF Commitee, May 2001. 63

[71] Texas Instruments. *Four Output Clock Generator/Jitter Cleaner With Integrated Dual VCOs*, **SCAS882D**. Texas Instruments, February 2012. 65, 90, 130

[72] Xilinx. *Spartan-6 FPGA GTP Transceivers - Advance Product Specification*, **UG386 v2.2**. Xilinx, April 2010. 67

[73] Texas Instruments. *3.9μA, SC70-3, 30ppm/$^\circ$C Drift VOLTAGE REFERENCE*, **SBOS392A**. Texas Instruments, September 2007. 67

[74] H Johnson and M Graham. *High-Speed Digital Design A Handbook of Black Magic*, **ISBN13: 978-0133957242**. Prentice Hall, April 1993. 69, 70, 71

[75] Xilinx. *Viretx-6 FPGA PCB Design Guide*, **UG373 v1.2**. Xilinx, June 2010. 71, 129

[76] Xilinx. *Spartan-6 FPGA DC and Switching Characteristics - Data Sheet*, **DS162 v3.0**. Xilinx, October 2011. 75, 108, 133

[77] CERN. *VME64x Transition Module Dimensions*. CERN, September 1997. 79

[78] Polar Instruments. *Web Page*. Polar Instruments. 79, 130

[79] J Ancu. *The ATLAS IBL BOC Demonstrator*, **ATL-INDET-PROC-2011-038**. CERN, 2011. 87, 97

[80] M Wensing et al. *Testing and firmware development for the ATLAS IBL BOC prototype*, **ATL-INDET-PROC-2012-022**. TWEPP Conference Proceedings, November 2012. 87, 88, 91, 92, 93, 97, 104, 106, 107, 130, 131

[81] Opencores. *Specifications for the Wishbone System-on-Chip Interconnection Architecture for Portable IP Cores*, **Revision B.3**. Opencores, September 2002. 88

[82] M Wensing and B Bergenthal. *IBL BOC Users Guide*, **draft version**. September 2012. 89, 95, 96

[83] ON Semiconductor. *3.3V ECL Programmable Delay Chip*, **MC10EP195/D Rev. 18**. ON Semiconductor, April 2009. 89, 130

[84] N Sawyer. *Data Recovery*, **XAPP224 v2.5**. Xilinx, July 2005. 93

[85] B Schneider. *Electrical SNAP12 tests*, **IBL General Meeting**. CERN, October 2011. 102, 131

[86] S K Nderitu. *Atlas Pixel Opto-Board Production and Analysis and Optolink Simulation Studies*, **WUB-DIS-2007-03**. University of Wuppertal, 2007. 103

[87] D Hall, B Todd Huffman, and A Weidberg. *The radiation induced attenuation of optical fibres below -20°C exposed to lifetime HL-LHC doses at a dose rate of 700 Gy(Si)/hr*, **JINST 7 C01047**. Jinst - Institue of Pysics Publishing and SISSA, January 2012. 104

[88] P Vankov. *ATLAS Upgrade for the HL-LHC: meeting the challenges of a five-fold increase in collision rate*, **ATL-UPGRADE-PROC-2012-003**. CERN, January 2012. 110, 131

[89] M Kretz. *Studies Concerning the ATLAS IBL Calibration Architecture*, **CERN-THESIS-2012-067**. CERN, June 2012. 110, 131

[90] Xilinx. *7 Series FPGAs Overview*, **DS180 v1.12**. Xilinx, October 2012. 111, 134

[91] W R Leo. *Techniques for Nuclear and Particle Physics Experiments*. Springer-Verlag, 1994. 118

[92] I Dawson and V Hedberg. *Radiation in the USA15 cavern in ATLAS*, **ATL-COM-TECH-2004-001**. CERN, 2004. 119, 131

[93] UCL HEP Group. *TTC Interface Module for SCT & Pixel Detectors (webpage)*. University College London (UCL). 120, 131

List of Figures

List of Tables

Glossary

Not all abbreviations used in this thesis are explained here, common- and standard abbreviations used in computer science (e.g. RAM, MB, MHz, ...) are assumed to be known by the reader.

ALICE	A Large Ion Collider Experiment
ASIC	Application Specific Integrated Circuit
ATCA	Advanced Telecommunications Computing Architecture
BCF	BOC Control FPGA
BGA	Ball Grid Array
BMF	BOC Main FPGA
BOC	Back Of Crate
BPM	Bi-Phase Mark
CAN	Controller Area Network
CERN	Conseil Européen pour la Recherche Nucléaire, but known today as European Organization for Nuclear Research
CI	Cluster Interconnect
CLB	Configurable Logic Blocks
CML	Current Mode Logic
CMS	Compact Muon Solenoid
CMT	Clock Management Tile
COB	Cluster On Board
COTS	Commercially Of The Shelf
CP	Charge Pump
CPLD	Complex Programmable Logic Device
DAC	Digital to Analog Converter
DAQ	Data AcQuisition
DCM	Digital Clock Manager
DCS	Detector Control System
DPM	Data Processing Modules
DTM	Data Transportation Module
ECAL	Electromagnetic CALorimeter

ELMB Embedded Local Monitor Board

EoF End of File

FPGA Field-Programmable Gate Array

FTK Fast TracKer

GMII Gigabit Media Independent Interface

HCAL Hadron CALorimeter

HDL Hardware Description Language

HEP High Energy Physics

HL-LHC High Luminosity - Large Hadron Collider

HLT High Level Trigger

HOLA High-speed Optical Link for ATLAS

IBIS Input/output Buffer Information Specification

IBL Insertable B-Layer

ILA Integrated Logic Analyzer

IOB IO Buffers

IP Intellectual Property

IPMC Intelligent Platform Management Controller

JTAG Joint Test Action Group

LDC Link Destination Card

LEP Large Electron-Positron (collider)

LHC Large Hadron Collider

LPC-FMC Low Pin Count FPGA Mezzanine Connector

LSC Link Source Card

LUT Look Up Table

LVDS Low Voltage Differential Signaling

LVTTL Low Voltage Transistor-transistor Logic

MGT Multi-Gigabit Transceivers

MSA Multi-Source Agreement

MSR Mark to Space Ratio

NRZ Non-Return-to-Zero

NTC Negative Temperature Coefficient

PECL Positive Emitter Coupled Logic

PFD Phase Frequency Detector

PLL Phase-Locked Loop

PRM Program Reset Manager

PS Proton Synchroton

PTH Plated-Through Hole

QSFP Quad Small Form-factor Pluggable

RCC Read Out Driver (ROD) Crate Controller

RCE Reconfigurable Cluster Element
RIA Radiation Induced Attenuation
RICH Ring Imaging CHerenkov
ROB Read Out Buffer
ROBIN Read-Out Buffer IN
ROD Read Out Driver
ROS Read Out System
RTM Rear Transition Module
SBC Single Board Computer
SCC Single Chip Card
SFP Small Form-factor Pluggable
SLAC Stanford Linear Accelerator Center
SMD Surface-Mounted Devices
SMU Switch Management Unit
SoC System on Chip
SoF Start of File
SPI Serial Peripheral Interface
SPS Super Proton Synchroton
SSO Simultaneously Switching Outputs
SSTL Stub Series Terminated Logic
TIA Trans-Impedance Amplifier
TIM Timing, Trigger and Control (TTC) Interface Module
TOTEM TOTal Elastic and diffractive cross section Measurement
TTC Timing, Trigger and Control
USA-15 Underground Service Area - 15
VCO Voltage Controlled Oscillator
VCSEL Vertical Cavity Surface Emitting Laser
VDC VCSEL-Driving Circuit
VIA Vertical Interconnect Access
VIO Virtual IO
VME Versa Module Eurocard

Acknowledgements

With the writing of this thesis my involvement in the work for the ATLAS Experiment at CERN will come to an end. It was a very interesting time which was filled with many different challenges to overcome. But there was also a lot of knowledge and experience to be gained and it is satisfying to have worked for such a big project. The thought of having contributed, even if it was only a little bit, to the possibly Nobel Prize winning discovery of the Higgs Boson is thrilling. While my share in that discovery will at the most be enough for a pixel in the dot of the "*et al.*" on the Nobel Prize Diploma, several other people deserve to be mentioned considerably more than that for helping and contributing to my work.

First of all I would like to thank my PhD supervisor *Professor Reinhard Männer* for providing me the chance to work on this exciting project. Without his trust and confidence in me this thesis would not have been possible. The same is true for *Andreas Kugel* who worked with me in Mannheim for the ATLAS Experiment. His supervision and support kept me on the right track and I am very grateful for his help. He also worked on the firmware for the Re-BOC while I was designing the hardware. In this regard *Timon Heim* and *Marius Wensing* from the University of Wuppertal have to be mentioned as they contributed greatly to the development and testing of the firmware for which I am much obliged. Additional thanks go to *Basil Schneider* for performing the evaluation and testing of the SNAP12 modules at the University of Bern. During the debugging of the first prototypes *Andreas Wurz* and his good knowledge of electrical engineering was a great help for which I am very grateful.

Other colleagues from the ATLAS collaboration who earned my gratitude by their help and support are *Tobias Flick*, *Jens Dopke* and *Markus Joos*.

More general thanks go to my friends and coworkers for their moral support and the fun times we had. Especially to be mentioned are *Thomas Gerlach* and *Moritz Kretz* who shared my office in Mannheim with me and *Stefan Kirsch* and *Torsten Alt* for adding some entertainment to my visits at CERN.

Last but of course not least I would like to thank my parents and my girlfriend *Natalie Hoge* for believing in me and giving me emotional support.