Department of Physics and Astronomy

University of Heidelberg

Master thesis

in Physics

submitted by

Dieter Liebert

born in Stuttgart

2019

Test and Monitoring Procedure for the TREX Module

This Master thesis has been carried out by Dieter Liebert

at the

Kirchhoff-Institute for Physics

under the supervision of

Prof. Dr. Hans-Christian Schultz-Coulon

Zusammenfassung:

Um im Zuge der zukünftigen LHC-Upgrades weiterhin eine hohe Effizienz an Triggerentscheidungen zu gewährleisten, werden einzelne Module des ATLAS Level-1 Kalorimetertriggers ersetzt oder überarbeitet. Eine Erweiterung des Preprocessors dient dabei der optischen Bereitstellung von Tile-Kalorimeterdaten an die neuen Triggerprozessoren bei simultaner Datenübertragung zu den Altsystemen mittels LVDS. Die „Tile Rear Extension“(TREX) Leiterplatten sind dazu mit modernsten FPGAs und optischen Hochgeschwindigkeitssendern ausgestattet.
In dieser Arbeit werden TREX Prototypen zunächst bezüglich ihrer Betriebsbedingungen zur Einhaltung der einzelnen Bauteilspezifikationen und Abschätzung der Leistungsaufnahme unter größtmöglicher Last überwacht. Dies ergibt einen maximalen Verbrauch von 43 W. Weiterhin wird der Datenempfang in deren Hauptverarbeitungseinheit sichergestellt und einem Langzeittest unterzogen. Die Bewahrung der richtigen Strahlkreuzungszuweisung von Daten und deren zeitliche Ausrichtung wird aufgezeigt. Die optische Übertragung wird anhand von Augendiagrammen und standardisierten IEEE Anforderungen in eigenständigen Prüfschleifen und gemeinsamen Tests mit einem Zielprozessor charakterisiert. Dabei wurden Bitfehlerraten von 1.7×10^{-15} auf allen Übetragungsstrecken gemessen. Ein Latenzzeitenvergleich zum bisherigen Datenpfad zeigt keinen kritischen Zuwachs. Ein sicherer Betrieb des TREX-Moduls ist damit gewährleistet.

Abstract:

In order to ensure a high efficiency of trigger decisions in the course of future LHC upgrades, individual modules of the ATLAS Level-1 Calorimeter Trigger will be replaced or revised. An extension of the Preprocessor serves for the optical provision of Tile calorimeter data to the new trigger processors with simultaneous data transfer to the legacy systems via LVDS. For this purpose, the "Tile Rear Extension"(TREX) PCBs are equipped with state-of-the-art FPGAs and high-speed optical transmitters.
In this thesis, TREX prototypes are first monitored with regard to their operating conditions for compliance with the individual component specifications and estimation of power consumption under the greatest possible load. This results in a maximum power consumption of 43 W. Furthermore, the data reception in their main processing unit is ensured and subjected to a long-term test. The preservation of the proper bunch-crossing assignment of data and their temporal alignment is demonstrated. The optical transmission is characterised by eye diagrams and standardised IEEE requirements in stand-alone loopback and joint tests with a target processor. Bit error rates of 1.7×10^{-15} were measured on all transmission lines. A latency comparison to the legacy data path shows no critical increase. Safe operation of the TREX module is thus guaranteed.

Contents

Chapter 1

Introduction

Research on the properties of matter in large and small dimensions has led in the past to a deeper understanding of processes and interrelations in nature. Physical theories mostly focus on single channels of interactions or are only valid under ideal conditions. Experiments are designed to confirm or reject the theories or to restrict them to ranges of validity. One of the most successful theories is the Standard Model (SM) of particle physics which was able to predict the existence of several particles and their properties. Its complete confirmation would not have been possible without the construction of huge accelerators and appropriate particle detectors. But known phenomena exist that can't be explained within the framework of the SM, like the presence of dark matter or the baryon asymmetry problem.

In the search for new physics beyond the SM (BSM), modern experiments produce an overwhelming amount of data. Real-world limitations make it impossible to transfer all this data to mass storage. But this is usually also not demanded. Many experiments are rather interested in rare or so far unseen events. Without an effective pre-selection based on clear criteria, the desired events can't be recorded for an offline analysis. Selection criteria are e.g. energy thresholds or the types of identified particles. Trigger systems realise this pre-selection in hardware, software or in a combination of both and are precisely concerted to the experimental conditions. If they change, also the trigger must adapt to the new situation.

This is the current scenario for the ATLAS (A Toroidal LHC ApparatuS) trigger system, which has to cope with increased event rates and harsher pile-up conditions in the course of the Large Hadron Collider (LHC) upgrades that aim to reach a higher luminosity and energy. The ATLAS trigger is grouped into two levels, with the first level consisting of fast custom-built hardware that processes calorimeter and muon detector data. The second level is the software-based High Level Trigger, that makes the final selections in use of the full detector information.

The LHC upgrades affect mainly a subgroup of the ATLAS first level trigger: the Level-1 Calorimeter Trigger (L1Calo). New front-end calorimeter readout electronics are no longer compatible with the Run-2 design. Dedicated trigger hardware, the so called Feature Extractors (FEXes), will provide optical, digital interfaces to receive information from the calorimeter electronics in finer granularity. Other calorimeter types will continue to send analogue signals until the end of Run-3, which are first routed to the L1Calo Preprocessor where they have to be digitised and converted to optical data for transmission to the FEXes. This is accomplished by Tile Rear Extension (TREX) modules that will be installed in the Preprocessor system.

Three TREX prototype modules were fabricated and fully assembled at the Kirchhoff-Institute for Physics in Heidelberg. This thesis aimed to validate TREX operation concerning different transmission lines and the operating conditions on the basis of the prototypes. Board temperatures, voltages and

currents were monitored at full load comparing the development under two different heat sinks. The reception of electrical signals on the real-time path in the main processing unit of the TREX was tested with the usage of a delay element, that was able to shift the sampling point of incoming serial streams by more than a bit length. Optical transmission was verified in stand-alone and joint tests with a FEX module, including bit error rate measurements. Standardised receiver specifications were used for this purpose. A latency comparison between the former and upgraded Preprocessor system was done to exclude timing issues on the legacy data path.
Chapter 2 starts with an overview of the ATLAS experiment. Individual detector constituents are presented in detail. Chapter 3 explains the functionality of the L1Calo trigger with a focus on the Preprocessor. The fourth chapter summarises the upgrades for Run-3 including the TREX, which makes up the largest part here. Chapter 5 presents and discusses the mentioned tests on different transmission lines and the monitoring of operating conditions. This thesis ends with a summary of the results and a little outlook of possible future modifications in the sixth chapter.

Chapter 2

The ATLAS Experiment

The ATLAS experiment, located at the LHC in Geneva, Switzerland is a particle detector which aims to investigate key issues of modern particle physics experimentally and to observe new phenomena beyond the Standard Model of physics. It is a general purpose detector, which is able to detect multiple hadronic and leptonic final states in proton-proton collisions in the TeV scale. This chapter gives an overview of the detector structure and the tasks of the single detector parts according to [8].

2.1 The ATLAS Detector

The ATLAS detector is with its dimensions of 46m × 25m × 25m and a weight of 7000 tons the largest elementary particle detector at a collider ever built. A layout is shown in figure 2.1. The interaction point (IA) is in the middle of the detector, where the proton beams cross. It defines the point of origin for its coordinate system. The z-axis is aligned with the beam axis. The x-axis points towards the middle of the LHC ring and the y-axis points upwards. Using spherical coordinates, the polar angle θ is usually expressed as pseudorapidity

$$\eta = -ln[tan(\theta/2)]. \tag{2.1}$$

This has the advantage that in the relativistic limit ($m << p$) the pseudorapidity matches the beamline rapidity

$$y = \frac{1}{2} ln \left(\frac{E + p_{\mathrm{L}}}{E - p_{\mathrm{L}}} \right), \tag{2.2}$$

where p_{L} is the longitudinal momentum along the beam axis. Since rapidities are additive and their differences lorentz-invariant, they are a common quantity in particle physics, whereas the pseudorapidity is easier to determine experimentally.
As the LHC delivers bunches of proton beams at a frequency of 40.08 MHz (referred to as 40 MHz from now on) with an average number of 30 to 50 particle collisions, which is even desired to be increased up to 200 within the upcoming upgrades in the near future, the technical requirements are rather ambitious. Also the different behaviour of various particle species in decay or deposition of energy makes it necessary to divide the detector into the following major parts (from inside to outside): the inner tracking detector, the electromagnetic and hadronic calorimeters, the muon spectrometer with toroid magnets and the trigger and data acquisition system, which is not shown in figure 2.1. In the following these subsystems are described in more detail.

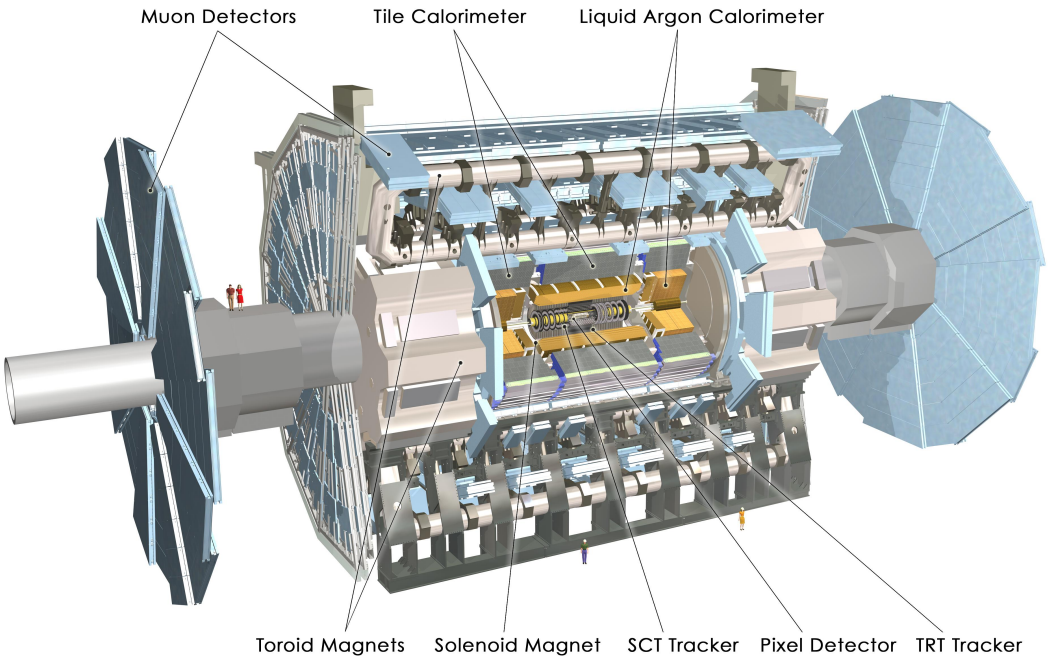

Figure 2.1: Schematic layout of the ATLAS detector. Various components are labelled. Particle beams enter from the sides and interact in the middle of the detector. [1]

2.1.1 The Inner Tracking Detector

The inner detector (ID), closest to the interaction point, has the function to reconstruct trajectories of electrically charged particles which also results in a momentum measurement. Therefore, it is placed inside a 2 T magnetic field created by a solenoid magnet. In addition, the point of origin of produced particles in collisions has to be identified. To achieve an appropriate spatial resolution, the inner detector is made up of the following three stand-alone sub-detectors: the Pixel Detector, the Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). They complement one another to achieve high precision in all coordinates for a pseudorapidity rang of $|\eta| < 2.5$ and a transversal momentum of $p_{\mathrm{T}} > 0.5$ GeV. The relation between bending radius R, magnetic field strength B and transversal momentum p_{T} with unit charge is

$$p_{\mathrm{T}}[\mathrm{GeV}] = 0.3B[\mathrm{T}]R[\mathrm{m}], \tag{2.3}$$

concluding that particles with lower p_{T} are stuck inside the detector. The inner detector and its parts are illustrated in figure 2.2. The sub-detectors, which are installed in two implementations (barrel and end-cap regions), are briefly described in the following.

The Pixel Detector

The Pixel Detector, the innermost part, has the highest spatial resolution because of its small distance to the IA, involving a high track density. The detection method is based upon a p-n junction in reverse bias, resulting in a large depletion zone. Traversing photons or charged particles excite electrons to the conduction band, creating electron-hole pairs, which are separated through the applied voltage and detected afterwards. The in total 1744 modules (15.5 $\mathrm{cm}^2\times$ 250 μm) are arranged in three concentrically (barrel) or disk (end-cap) layers, where the inner barrel layer is placed 5 cm around the beam axis. The smallest size of a single pixel is 50 μm in Φ and 400 μm in z or R direction, which leads to a spatial resolution of 10 μm in the $(R-\Phi)$ plane and 115 μm in z direction (or R for end-cap

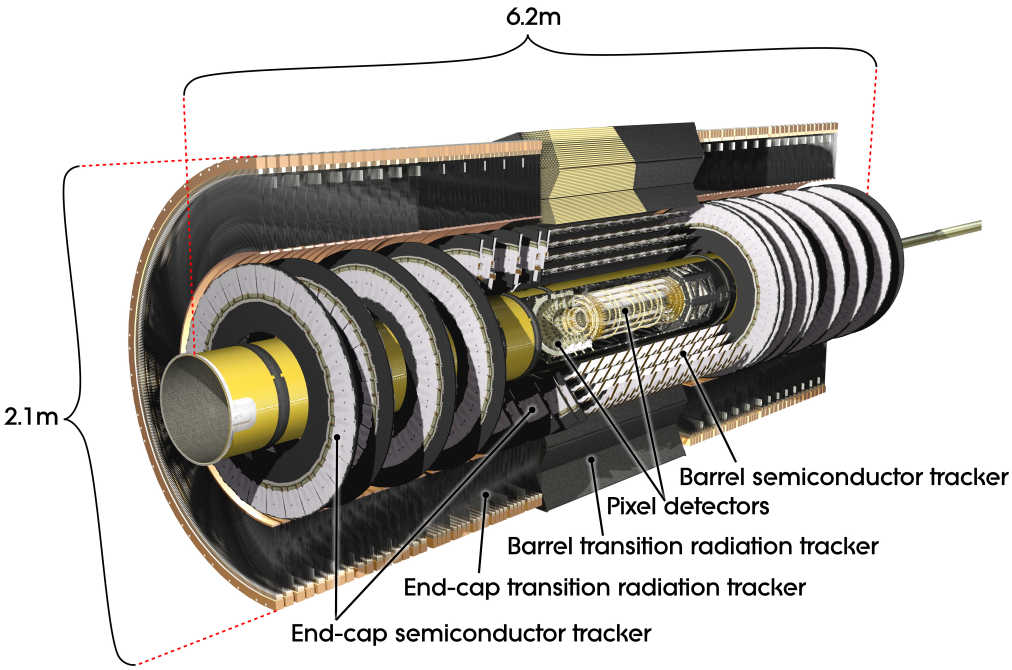

Figure 2.2: Schematic layout of the ATLAS inner detector. Pixel Detector, SCT and TRT are labelled. The detectors are either arranged concentrically around the beam axis or as end-caps at the detector's front and back to avoid particle loss. [1]

detectors). A fourth pixel layer was installed later during LS1 (long shutdown 1) [16].

The Semiconductor Tracker

The SCT consists of four layers in the barrel region and 18 disks in the end-caps of double-sided semiconductor microstrip sensors, where the two sides are rotated by a stereo angle of 40 mrad to get two-dimensional space point information. Two sensors form in total 768 silicon strips with a length of 12 cm, a thickness of 285 μm and a constant pitch of 80 μm [2]. These sensors are grouped to four to form a module and the whole SCT is made up of 4088 modules. This topology results in a resolution of 17 μm in the $(R - \Phi)$ plane and 580 μm in z (barrel) or R (end-cap) direction.

The Transition Radiation Tracker

The TRT is a combination of a transition radiation detector and a drift chamber. The basic detector elements are straw tubes with a 4 mm diameter and a 31 μm thick gold plated tungsten anode wire in the centre. The straw tube wall forms the cathode with a 0.2 μm Al layer. They are filled with a gas mixture of 70 % Xe, 27 % CO_2 and 3 % O_2 with a bit over pressure. The volume between the tubes is stuffed with 19 μm fibres (barrel) or 15 μm foils (end-cap) made out of polypropylene, functioning as a radiator to increase signals of the transition radiation effect. In the barrel region the straw tubes are grouped up to 73 layers with a length of 144 cm parallel to the beam axis, whereas there are 160 radially layers in the end-caps with a length of 37 cm. This design leads to a measurement only in the $(R - \Phi)$ plane with an accuracy of 130 μm, leaving 36 hits per track for charged particles with transverse momentum $p_T > 0.5$ GeV and pseudorapidity $\eta < |2.0|$. Since the transition radiation depends on the Lorentz factor $\gamma = E/m$, the strongest contribution comes from electrons with photon energies up to 30 keV. These photons can be absorbed by the Xe in the straw tubes, giving a characteristical signal for electron identification besides the TRT function as a tracker.

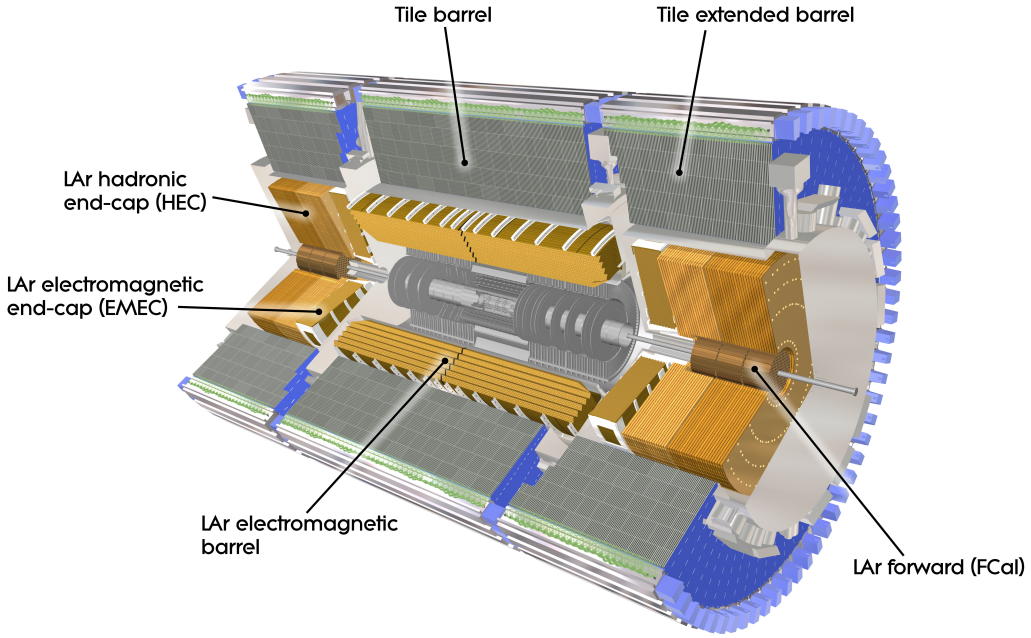

Figure 2.3: Schematic layout of the ATLAS calorimeters. The LAr detectors are shown in orange, where electromagnetic and hadronic parts are indicated by slightly different colours in the end-cap regions. The hadronic Tile calorimeter is shown in blue from the side. The ID is greyed out. [1]

2.1.2 Calorimetry

The calorimeters perform the destructive energy measurement of particles created in the proton-proton collisions. ATLAS uses sampling calorimeters, which are constructed from alternating layers of dense absorber and active sampling material. They cover the pseudorapidity range $\eta < |4.9|$ by their size of 12 m in length and 4.25 m in the outer radius. As shown in figure 2.3, the calorimetry is subdivided into two main parts: the electromagnetic (ECAL) and the hadronic calorimeter (HCAL). Again, both calorimeter types are installed as barrel or end-cap implementations. The functionality of both is explained in the sections below.

The Electromagnetic Calorimeter

The ECAL measures precisely the energy deposition of electrons and photons. All of its modules use liquid argon (LAr) as active sampling and lead as absorber material. The different layers are arranged in an accordion shape resulting in a nearly equal response signal independent of the particle's direction of travel. At energies beginning from 1.022 MeV, electrons lose their energy mainly through bremsstrahlung and pair production in the absorber, generating a cascade of secondary particles. At lower energies ionisation and the photoelectric effect dominate. To collect the ionisation charges, a voltage of 2 kV is applied between two copper electrodes in the LAr layer. A third electrode in the centre does the readout via capacitive coupling.
The electromagnetic barrel (EMB) and end-cap calorimeter (EMEC) together cover a pseudorapidity range of $|\eta| < 3.2$. They have a total thickness of 22 to 33 radiation lengths X_0, varying for different η regions. Because of the big amount of produced particles in collisions and the demand for an appropriate resolution, the $\Delta\eta$ granularity of the EMB increases towards the beam line. It reaches from 0.003 at the inner surface to 0.05 at the outer one. In the EMEC the $\Delta\eta$ granularity varies from 0.003 to 0.1. In $\Delta\Phi$ the granularity reaches from 0.025 to 0.1. In order to take energy losses in the

ID into account a presampler is placed in front of the ECAL, which is in the end just one LAr layer of 11 mm thickness. With all this implementations the ECAL reaches an energy resolution of

$$\frac{\sigma(E)}{E} = \frac{10\%}{\sqrt{E}} \oplus 0.17\%. \tag{2.4}$$

The constant term arises due to non-uniformities in the ECAL response.

The Hadronic Calorimeter

The ATLAS HCAL consist of three parts: the Tile calorimeter, the LAr end-cap calorimeter (HEC) and the LAr forward calorimeter (FCal).
The Tile calorimeter is constructed around the ECAL and has alternating layers of steel absorbers and plastic scintillators. Deposited energy excites molecules in the scintillator, which send out photons at deexcitation. Since they are usually in the UV scale, wave length shifters are added to convert the UV photons to visible light. Additionally, the photons are collected by fibres which also apply a wave length shift and route them to photomultipliers (PMTs). Since the physical processes for hadronic showers differ from electromagnetic showers, even electrically charged hadrons won't lose all their energy in the ECAL, mainly because of their high mass. This demands for a detector that is big enough to contain the whole hadronic shower. Therefore, the Tile calorimeter has a radial thickness of about 7.4 hadronic interaction lengths λ_{had}. But still with these dimensions neutrinos and muons will escape the Tile calorimeter. In the case of the neutrinos (and potentially dark matter) this leads to missing energy in the evaluation of the measurements.
The HEC has copper absorbers and is installed in the range $1.5 < |\eta| < 3.2$ with a thickness if 12 λ_{had}.
The FCal covers the highest η ranges up to $|4.9|$ and uses copper absorbers in the first layer because it is optimised for electromagnetic showers, but also tungsten for the following layers. It reaches a thickness of 10 λ_{had}.
The whole HCAL $\Delta\eta$ x $\Delta\Phi$ granularity varies from 0.1 x 0.1 to 0.2 x 0.2 for different subcalorimeters and η ranges. The energy resolutions are now different for different particles. As a reference the resolution for charged pions are provided.

$$\begin{aligned} \text{Tile calorimeter: } \frac{\sigma(E)}{E} &= \frac{52.7\%}{\sqrt{E}} \oplus 5.7\%. \\ \text{HEC: } \frac{\sigma(E)}{E} &= \frac{70.6\%}{\sqrt{E}} \oplus 5.8\%. \\ \text{FCal: } \frac{\sigma(E)}{E} &= \frac{70.0\%}{\sqrt{E}} \oplus 3.0\%. \end{aligned} \tag{2.5}$$

2.1.3 The Muon Spectrometer

Since muons interact electromagnetically one could naively expect that they are covered by the ECAL when created at the IA. But due to their high mass they lose only little energy via bremsstrahlung, leaving the ID and calorimetry nearly unaffected. The same holds for secondary muons produced in particle showers. Therefore it requires an additional system to take muons or other charged particles into account that eventually pass the calorimetry. In order to measure their momentum, a magnetic field in the range $|\eta| < 2.7$ is created by superconductive toroid magnets with eight coils each for barrel (0.5T) and end-cap regions (1T). Apart from tracking, the muon system does bunch-crossing identification and triggering for defined p_{T} thresholds at $|\eta| < 2.4$. So the entities of the muon system can be divided into precision-tracking chambers and trigger chambers.

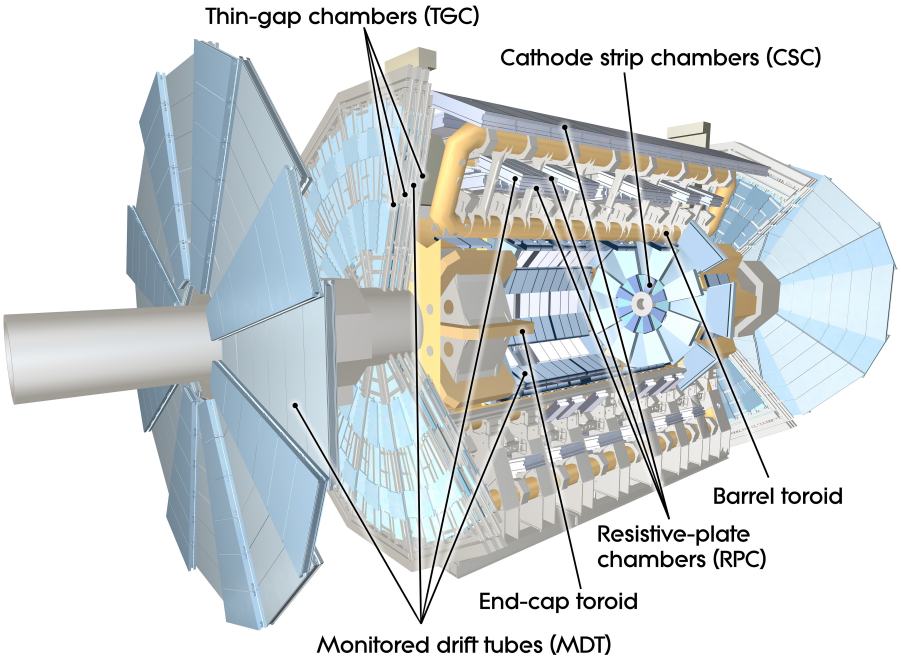

Figure 2.4: Conceptual layout of the ATLAS Muon Spectrometer. Toroid magnets are illustrated as yellow loops. Tracking and trigger chambers are greyish blue [1].

Precision-tracking chambers

Tracking is mainly achieved via Monitored Drift Tubes (MDTs) which are cylindrical tubes filled with a gas mixture of Ar and CO_2 that is ionised by traversing muons. Charges are collected by a central anode wire with an applied voltage of approximately 3 kV between anode wire and cathode tube walls. While the spatial resolution of a single tube is 80 μm, MDTs are grouped to chambers with up to eight consecutive segments of tubes achieving a spatial resolution of 35 μm per chamber. In the barrel region the MDT chambers are installed in three concentric layers between and on the coils of the toroid magnets, while for the end-cap regions they are installed as four disks before and behind the magnets. Only for the innermost end-cap disk Cathode Strip Chambers (CSCs) are used instead of MDTs in the region $2.0 < |\eta| < 2.7$ (compare figure 2.4). CSCs are also gaseous detectors constructed as multiwire proportional chambers with segmented cathode strips orthogonal to the anode wires. They allow a higher rate capability and granularity compared to the MDTs, which is needed for high $|\eta|$ and small distances to the IA. Their spatial resolution is 40 μm. The precision tracking with applied magnetic field allows a momentum resolution of about 10% for 1 TeV tracks.

Trigger chambers

For triggering, the muon system is complemented by Resistive Plate Chambers (RPCs) in the barrel and Thin Gap Chambers (TGCs) in the end-cap region. They are designed for short readout times in the range of tens of nanoseconds with small spreaded signals to identify the bunch-crossing. The design principle of RPCs is a gaseous detector with two parallel plates in a distance of 2 mm and an applied voltage of 9.8 kV. They are mounted together with the MDT chambers in the second, where they surround them, and the third concentric layer. In contrast TGCs are multiwire proportional chambers with a smaller distance between wire and cathode plates then the wire-to-wire distance, resulting in small drift times. One TGC is placed in front of the first tracking disk, while another is placed in front and two are placed after the second disk.

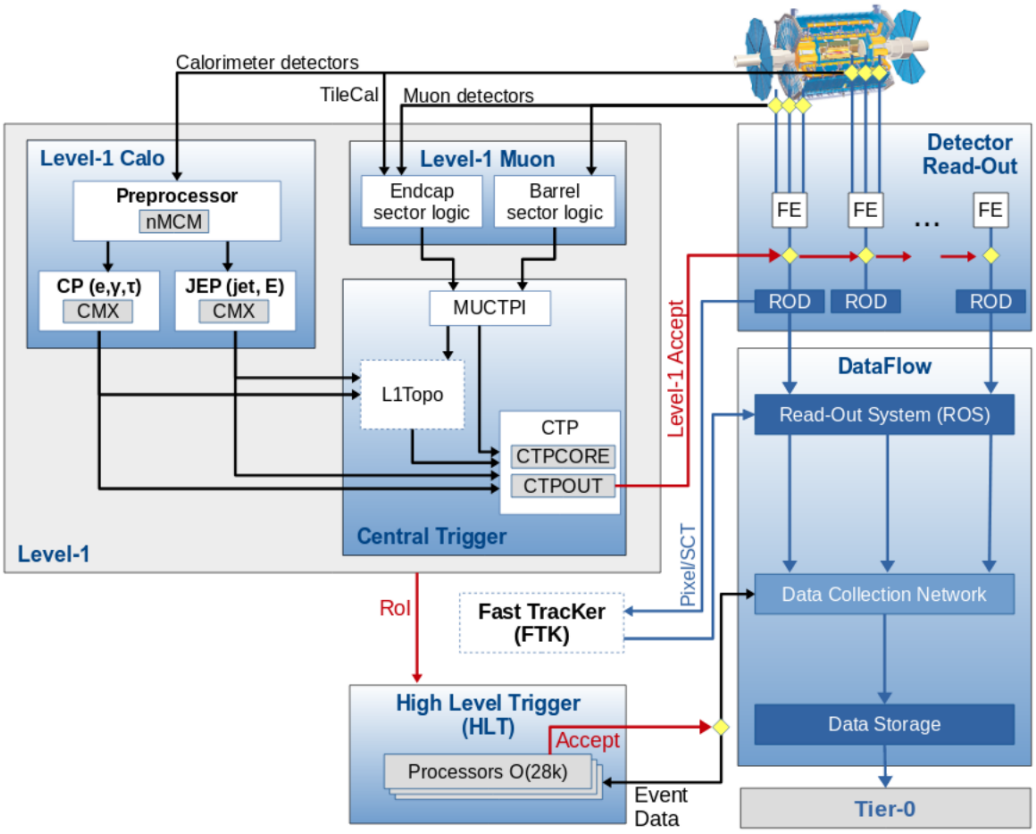

Figure 2.5: Overview of the ATLAS TDAQ system in Run-2. FTK was in Run-2 still under commissioning [3].

2.1.4 Trigger and Data Acquisition System

Trigger and Data Acquisition (TDAQ) is a substantial task in collider experiments as it reduces the initial amount of raw data by separating interesting event candidates from background. With about 10^8 readout channels the overall output of the ATLAS detector is also way too big to write everything to data storage. In Run 2 the whole trigger system, illustrated in figure 2.5, reduces the bunch-crossing frequency of 40.08 MHz to approximately 1 kHz in two stages. The first stage is the Level-1 (L1) trigger. It is realised by custom hardware and receives signals on trigger tower basis from the calorimeters and the muon detector. Upon a selected event an L1 Accept signal (L1A) is send within 2.5 μs to the subdetectors which starts the readout to pipeline memories. Based on the decisions of the second stage, the High Level Trigger (HLT), the buffered data is rejected or transferred to data storage for offline analysis. The HLT is software running on commercial computer farms using calorimeter information in finer granularity but also tracking data from the ID for its decisions.
The L1 Trigger can be further subdivided into L1Calo, L1Muon and Central Trigger. Of main interest for this work will be the L1Calo trigger, especially its Preprocessor and upgrades, that will be the content of the next chapters.

Chapter 3

Level-1 Calorimeter Trigger

As a part of the first level the L1Calo trigger is involved in the primary event selection looking for signatures of electrons, photons, τ leptons and jets with high transversal momentum in the calorimeters. It also computes the total transverse energy ΣE_{T} and the missing transverse energy E_{T}^{miss}, which occurs due to neutrino production or possible dark matter objects. In the following sections an overview of the L1Calo architecture and its input signals is given as described in [8] and [10]. Section 3.3 on the Preprocessor also used [12] and [26] as a reference.

3.1 Input Signals

The L1Calo trigger receives analogue input signals from both calorimeters. The signals are not send by the single calorimeter cells but first summed up to trigger towers in the front-end electronics. At $|\eta| < 2.5$ the granularity of these signals is 0.1×0.1 in $\Delta\eta \times \Delta\Phi$ but gets coarser at larger pseudorapidity, shown in figure 3.1. Such a tower covers the full depth of the corresponding electromagnetic or hadronic detector parts. By this procedure 7168 trigger tower signals are formed in total, each describing E_{T} values in multiple calorimeter cells for ECAL and raw energy for HCAL. The resulting trigger tower

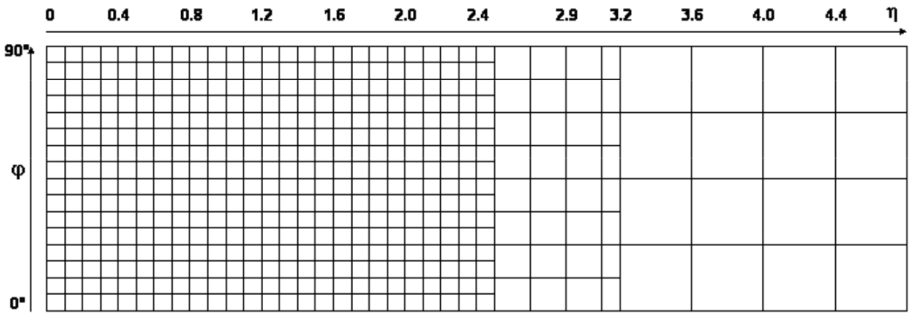

Figure 3.1: Trigger tower granularity for $\eta > 0$ and one quadrant in Φ. [10].

pulses are in general of bipolar shape. The integrals of the positive and negative part match due to charge conservation. Figure 3.2 exemplifies the pulses for EMB and FCAL. The shape inequalities of individual detector parts result from different detector geometries. Pulses of the Tile calorimeter look like the positive part of the EMB with a longer but therefore lower undershoot, so that it can be usually ignored.

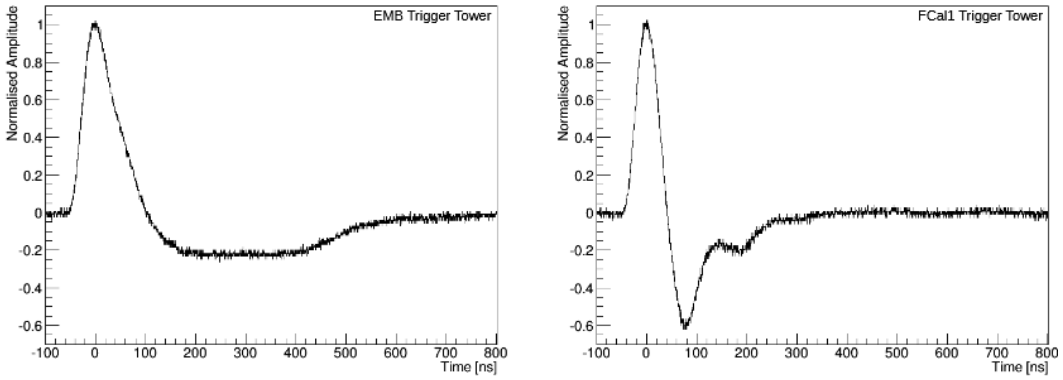

Figure 3.2: Bipolar trigger tower pulses for EMB (left) and FCAL (right) [26].

3.2 Implementation in Run-2

As it can be seen in figure 2.5, the L1Calo trigger consists in Run-2 of three subsystems:

1. The Preprocessor which mainly digitises the incoming analogue calorimeter signals, assigns a transversal energy value E_{T} and determines the bunch-crossing (BC).

2. The Cluster Processor (CP) which identifies electrons, photons, τ leptons and hadrons.

3. The Jet/Energy Processor (JEP) which provides identification of jets and computes the total, the missing and the sum E_{T} values.

Whereas the Preprocessor tasks are self-contained, the CP and JEP have some commonalities. They are both performing sliding window algorithms. For CP such a window is composed by 4×4 trigger towers while JEP uses so called jet elements which are a summation over 2×2 trigger towers, so that one jet element covers a granularity of 0.2×0.2 in $\Delta\eta \times \Delta\Phi$. JEP window sizes vary from 2×2 to 4×4 jet elements. Basically these algorithms look, in steps of one in η or Φ direction, for the exceedance of definable energy thresholds of neighbouring trigger tower or jet element sums across their assigned area. Additional isolation criteria can be applied. By this method, CP and JEP produce results in form of Trigger Objects (TOBs) which contain information about energy, location and the type of particle. TOBs are forwarded to the Central Trigger which adds similar information form L1Muon and makes the final L1 trigger decision.
For a proper performance the whole azimuthal coordinate range is divided into four quadrants. To avoid efficiency loss across quadrant boundaries, some trigger towers and jet elements are duplicated by the Preprocessor. In the end, the CP is realised in four crates each containing 14 modules and the JEP in two crates with 16 modules each. One CP crate and eight JEP modules cover one Φ quadrant respectively.

3.3 The Preprocessor

The complete Preprocessor system is distributed over 124 hardware-identical Preprocessor Modules (PPMs) in eight VME[1] crates, where six crates are assigned for LAr and the others for the Tile calorimeter. Each crate has in addition a Timing Control Module (TCM) that distributes information such as the BC-clock (40 MHz), L1A or counter reset signals over the VME backplane. A single board computer acts as crate controller for configuring and controlling of the PPMs via the VME interface.

[1] Versa Module Eurocard.

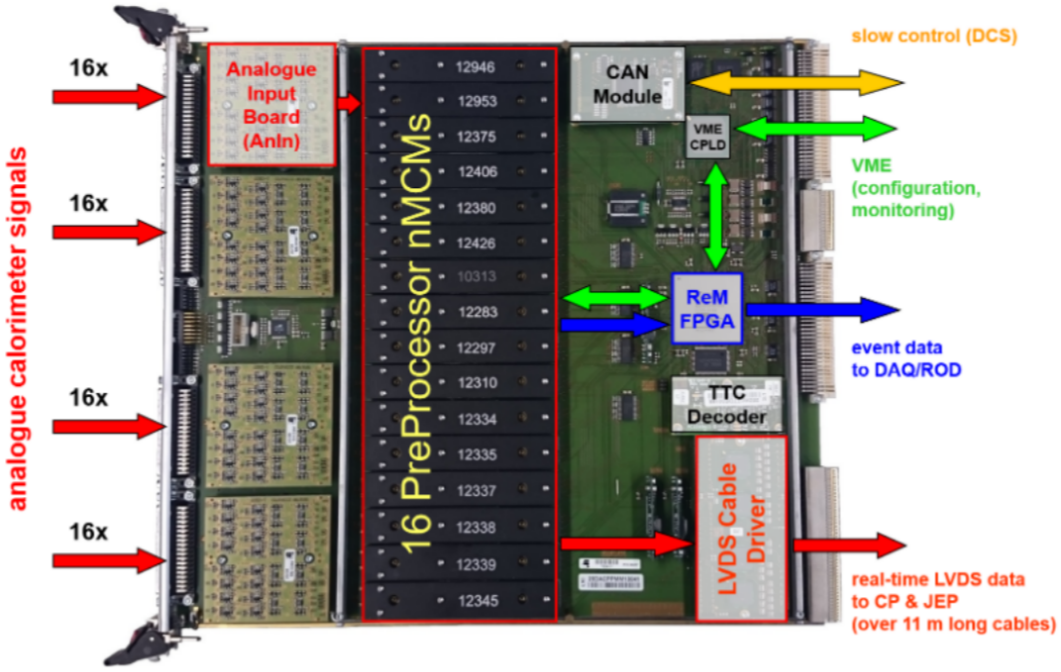

Figure 3.3: The Preprocessor Module. Real-time path elements are marked in red, the readout path in blue, communication to DCS in yellow and the VME path for board configuration or lab monitoring in green.

All signal processing paths on the PPM are illustrated in figure 3.3. The analogue trigger tower pulses are received as differential signals, having the advantage of low electromagnetic susceptibility since interfering signals should have the same impact on both wires. These signals come with amplitudes up to 2.5 V corresponding to a maximum E_{T} value of 255 GeV. In a first step, up to 64 differential signals are converted to single ended form by operational amplifiers on four Analogue Input (AnIn) boards. A fixed gain is applied to fit the 1 V digitisation range and a digital-to-analogue converter (DAC) may add a voltage offset to have an uniform baseline for all channels. Next, the signals are routed to the new Multichip Modules (nMCMs), the main processing units on the PPM. They were installed during LS1, replacing the prior modules based on application-specific integrated circuits (ASICs) by newer ones based on field programmable gate arrays (FPGAs). A PPM contains 16 nMCMs each handling four trigger tower signals. After digitisation with 10-bit resolution at the doubled BC frequency (80 MHz), the results are received by the Calorimeter Information Preprocessor (CALIPPR) FPGA. Figure 3.4 demonstrates its processing paths. An overview of the most important tasks follows.

- Synchronisation of trigger tower pulses: The signals arrive asynchronously due to different time-of-flights from the IA to the particular calorimeter cells and different cable lengths from calorimeters to the PPM. For a proper BC-wise L1 decision, these signals have to be aligned in time. A coarse alignment is achieved by a First In - First Out (FIFO) pipeline using the BC-clock to end up with delays in 25 ns steps. Fine time adjustment is achieved by delaying the sampling points of the analogue-to-digital converters (ADCs) in steps of 1.042 ns with a generated clock-like signal.

- Bunch-crossing identification (BCID): As shown in figure 3.2, the trigger tower pulses are multiple bunch-crossings wide. For a correct E_{T} extraction the peak maximum has to be identified and connected to the right BC for a proper trigger functionality. Several logical blocks are assigned to this task. For unsaturated pulses a finite impulse response (FIR) filter in combination

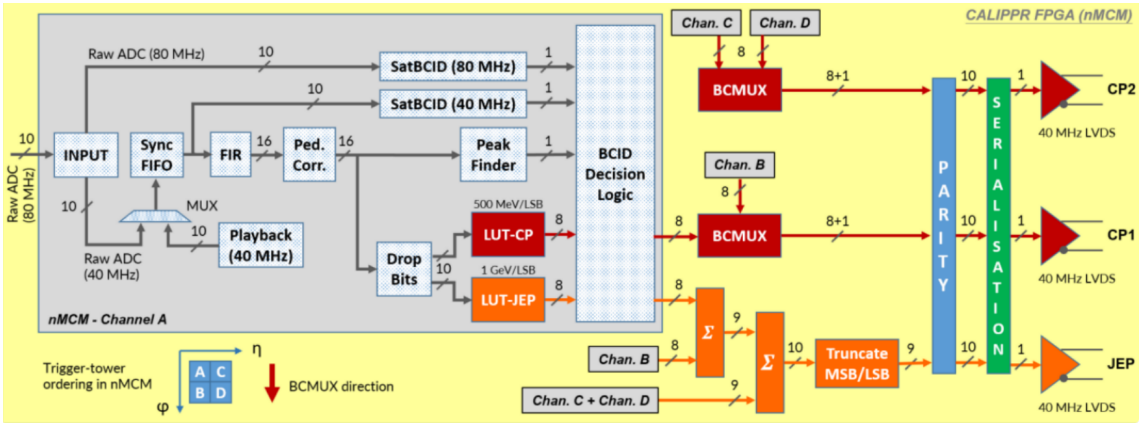

Figure 3.4: Functional behaviour of the CALIPPR FPGA real-time path. Various trigger tower signals are indicated by channel labeling. Transmission paths to CP are marked in red while orange paths show transmission to JEP. The Playback is a configurable memory used for test purposes [12].

with a peak finder algorithm is used. The FIR filter weights five consecutive ADC samples with pulse optimised coefficients and adds them up. The peak finder then simply searches for the first maximal value of three successive FIR outputs. Other values are set to 0. Saturated pulses are treated separately by looking at the rise times of the unfiltered ADC signals. Before reaching saturation, these signals are compared against defined thresholds to estimate the rise time of the pulse and the peak maximum, which is necessary if a signal saturates for several BCs. An additional option, although never been in use, is a voltage comparator on the AnIn boards with just one logical output. CALIPPR then detects the crossing over a defined threshold.

- Noise and pedestal correction: Contribution from thermal noise and pile-up effects have to be considered for a proper trigger functionality, since they adulterate the trigger tower signals. At ATLAS, thermal noise is of the order of 200 to 400 MeV. But the signal-to-noise ratio is optimised by the FIR filter coefficients. Pile-up effects vary with distance to the IA and trigger tower sizes and are considered by a dedicated algorithm, which works BC-wise for each channel. This algorithm computes an average filter output and subtracts a configurable constant from it that stands for the expected pedestal value. The filter outputs are then subtracted by this difference, giving the corrected values.

- Identification of deposited E_{T}: FIR filter outputs are 16-bit values that are reduced to ten bits in the first instance. This reduced value is then mapped to an transversal energy using a look-up-table (LUT), one for CP and one for JEP side respectively. Such a LUT is a 1024 $\times$ 8 memory, so that each reduced FIR input is mapped to an 8-bit energy value up to 255 GeV. These LUTs show in principal a linear behaviour, but a noise cut is applied assigning an output of zero to low energy peaks. Separate LUTs for CP an JEP allow an independent configuration and calibration which effects the trigger selectivity.

- Bunch-crossing multiplexing (BCMux): The E_{T} results for two channels on CP side are multiplexed into one serial stream. An additional bit, the BCMux flag, is sent to indicate which channel is sending the current data.

- Jet Sums: Since JEP uses jet elements instead of trigger towers, the E_{T} results have to be preprocessed before sending them to JEP. For unsaturated values the sum of four trigger tower channels is computed, giving a 10-bit value. If one of the trigger towers is saturated the summed

value is set to maximum, preserving saturation. Before transmission to JEP, the most or least significant bit is dropped to end up with nine bits like for the CP with energy and BCmux flag.

- Low Voltage Differential Signaling (LVDS) transmission: To ensure an error free transmission one odd parity bit is packed together with the actual data. It is set to 1 if the number of 1s in the data is odd, otherwise it is set to 0. A defective transmission can thus be detected on receiver side if the number of changed bits is odd. When it comes to serialisation, the actual data is enclosed by a start (logical 1) and stop bit (logical 0), so that the receiver knows how to pick data. In a last step these twelve bits are converted to differential form and transmitted with 480 Mbps at low voltage.

What is missing up to this point is the duplication of trigger towers or jet elements across Φ quadrants. This tasks is adopted by the LVDS cable driver (LCD) board, more precisely by four Xilinx FPGAs on the LCD. It receives in total 48 LVDS streams from the nMCMs, fans out the required signals and drives them over 11 meter long cables. Furthermore, the LCD compensates transmission losses by amplifying the signals.
In order to validate and calibrate the PPM trigger decisions, the sampled data and assigned energy have to be checked. This is why dedicated pipeline memories are integrated on the nMCM that are buffering the data for 128 BCs. Upon an L1A, which is on-board distributed by the Trigger, Timing and Control (TTC) decoder, the Readout Manager (ReM) FPGA collects the data from both pipelines, formats it and transmits it to the DAQ via a Readout Driver (ROD) module that does not belong to the Preprocessor itself.
The ReM FPGA and also a Complex Programmable Logic Device (CPLD) are furthermore involved in configuration and lab monitoring via VME. They handle the addressing and data transfer of the programmable locations. The single board computer acts as the VME controller.
Finally the PPMs themselves have to be monitored to guarantee smooth functioning. This is done during operation by the ATLAS Detector Control System (DCS), which is even able to power down the boards. A Fujitsu micro-controller, located on the PPM between the two middle AnIn boards, gathers all accessible information like temperatures and voltages and sends it to DCS via an CAN[2]-bus interface.
At the end of this chapter the playback memory in the CALIPPR should be emphasised. It is a 256 register memory and can be used as an alternative input. Data is loaded from VME to configure it channel- and BC-wise for an unambiguous identification of links and is injected in the real-time path upon request. The playback is used for several test applications that will appear in this thesis.

[2]Controller Area Network

Chapter 4

Run-3 Upgrades

The discovery of the Higgs boson at the LHC by the ATLAS and CMS experiment in 2012 opened a new unexplored area to particle physics. In order to increase the potential for discoveries and to improve precision on measurement, the LHC will undergo substantial upgrades in the next years to reach a higher luminosity and energy. For Run-3, starting from 2021, the luminosity will be approximately increased by a factor of two to $3{\cdot}10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$, resulting in higher event rates and harsher pile-up conditions. The ATLAS trigger system has to cope with this new experimental environment to maintain a high efficiency. The Run-3 upgrades focus mainly on improvements of the trigger and readout electronics, satisfying also requirements beyond Run-3. Nevertheless, some parts of the detector will be enhanced too. This chapter gives an overview of the upgrades during the second long shutdown (LS2) period as described in [4] and [7]. Section 4.3 is mainly based on [11].

4.1 Detector Upgrades

The increased event rates and pile-up after LHC upgrades will be especially challenging for the innermost detector parts. At high η, where background rates are high due to slightly scattered or unscattered beam particles, the detector performances is severely affected. For the muon system this affects the so called small wheel region. In figure 2.4, this corresponds to the first end-cap disk where the CSCs are mounted. In order to keep up with the expected particle rates of 15 kHz/cm^2 at High-Luminosity LHC (HL-LHC) and to reduce fake trigger events a New Small Wheel (NSW) will be installed. It uses micro mesh gaseous structures (Micromegas) as tracking and small-strip TGCs as trigger elements. Micromegas use a mesh to divide the detection volume in a cell into a amplification and a conversion gap with a high ratio of their electric fields. Ions created in the conversion gap are then collected by the mesh while electrons reach the amplification gap, starting an avalanche which is detected by electrodes. Small-strip TGCs use the same detection principle as explained in section 2.1.3 but with enhanced performance. Both detector types have proven to satisfy the demanded conditions. Another innovation is planned to be improved: the ATLAS forward proton detector (AFP). It aims to identify events with intact protons at high η. Such processes are associated with the so far not well understood elastic and diffractive scattering. It allows the testing of Quantum chromodynamics (QCD) and BSM models. On each side of the IA two stations at respective distances of 206 and 214 m will be installed, that can be retracted during proton injection and inserted when the beam is stable. The first station is a pure silicon tracker while the second combines tracking and timing detectors. This design allows to measure the momentum, angle and time-of-flight of protons from pp interactions. The AFP was already partially operating in Run-2 with missing time-of-flight systems that were not installed due to lacking vacuum tightness of PMTs. This will be integrated for Run-3.

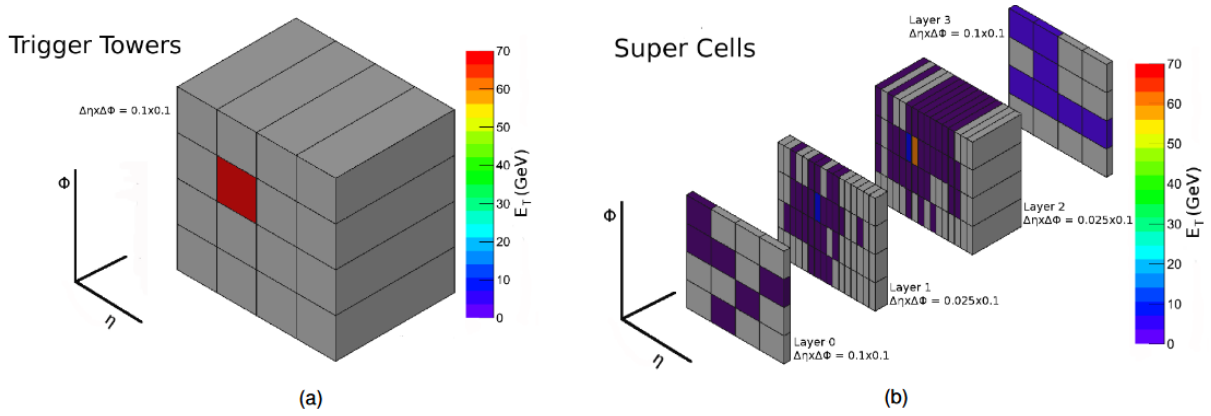

Figure 4.1: Energy deposition and reported signature of a 70 GeV electron in the current trigger towers (a) and in the upgraded super cells (b) [5].

Although it was planned to commission the Fast Tracker (FTK) earlier, it is now being finalised in the Run-3 upgrades. The FTK will be part of the TDAQ system and aims to provide fast tracking information from the ID as input to the HLT for every L1A. It will be realised by custom hardware and will operate for tracks with a transversal momentum p_{T} bigger 1 GeV/c within 100 μs. FTK will implement pattern recognition algorithms that match ID information to simulated tracks, stored in on-board memory chips.

4.2 LAr Calorimeter Upgrade

The LS2 upgrades of the LAr calorimeter will not improve the detector itself but rather aim to increase the quality of calorimeter-based L1Calo trigger decisions. To achieve this, LAr will provide information in finer granularity replacing the former trigger towers by so called super cells. A single tower is replaced by ten super cells by dividing the tower into four layers in transversal direction. The two inner layers are furthermore subdivided into cells with a granularity of 0.025×0.1 in $\Delta\eta \times \Delta\Phi$. Figure 4.1 demonstrates the difference in signature for a 70 GeV electron in trigger towers and super cells. This new architecture will allow to make use of new sophisticated algorithms enhancing jet rejection and pile-up correction. The readout of calorimeter cells in the new design requires a complete rework of the front-end electronics. For a faster transmission data is sent optically via a distribution system, called Fibre Optics eXchange (FOX), directly to L1Calo at 11.2 Gbps.
The L1Calo Preprocessor was designed for electrical input and thus can not receive the new optical signals. Also processing of more granular signals demands for changes, concluding that L1Calo has also to be updated to keep up with the increased requirements.

4.3 Upgrades of the Level-1 Calorimeter Trigger

New electronics must be developed for the reception of fine-granular optical LAr data. L1Calo is expanded by three FEXes (Feature Extractors) during LS2 which receive data directly from the LAr calorimeter. The tasks of CP and JEP are taken over by the electron FEX (eFEX) and the jet FEX (jFEX). A global FEX (gFEX) is searching for large-area jet objects. Figure 4.2 illustrates the old and new architecture.
Since the FEXes need also data from the Tile calorimeter, which continues to send analogue signals on trigger tower basis until the end of Run-3, the Preprocessor will be extended to provide Tile data optically. To achieve this, TREX (Tile Rear EXtension) modules will be installed in LS2. A

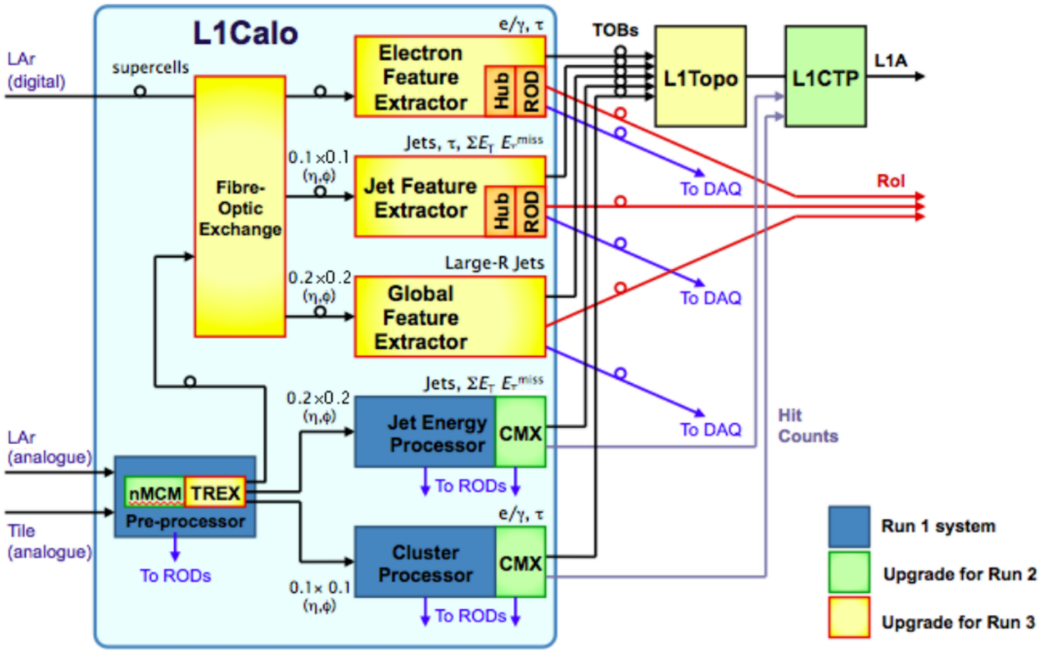

Figure 4.2: L1Calo architecture in Run-3. Elements in yellow show the new L1Calo constituents for Run-3. Elements in blue and green were used in previous runs [6].

TREX module is an 18-layer PCB[1] equipped with state-of-the-art FPGAs and optical transmitters. Until commissioning and validation of the upgrades the new and the legacy system will operate simultaneously. CP, JEP and the PPMs that receive analogue LAr data will be removed accordingly. After Run-3, the current Preprocessor plus extension will be replaced by custom Tile electronics similar to LAr. In the following TREX and its components are discussed in detail.

Tile Rear Extension

A schematic overview of TREX and its components is given in figure 4.3. As the Preprocessor upgrade only affects the PPM crates responsible for the Tile calorimeter, 32 TREX modules need to be installed. These modules are designed in such a way that they are simply plugged together with the PPMs on the backplane of the crate. Therefore, the LCD daughter-card has to be replaced by a passive module which only routes the signals to TREX. The tasks of the former LCD are taken over by four Data In-Out (DINO) FPGAs on TREX. These are Xilinx Artix-7 FPGAs which additionally duplicate the real-time data and send it to the main processing unit of TREX, the Preprocessor Data Collector (PREDATOR) FPGA.
For the prototype, the PREDATOR is a Xilinx Kintex Ultrascale device with many intellectual property (IP) cores and 64 Gigabit transceivers of GTH series that can handle data rates up to 16.2 Gbps [17], whereat the agreement of data rates from TREX to FEXes is 11.2 Gbps. The final TREX version will use a smaller FPGA with less but still sufficient resources and only 56 GTHs, which is more than the maximum required number of 29 during operation. In the PREDATOR the incoming LVDS streams are at first deserialised and then demultiplexed in case of CP streams. LVDS start and stop bits as well as the parity bit are checked and discrepancies are reported to VME registers and FEXes via an error flag. Before optical transmission, the data is duplicated and formatted as required. Also parts of the readout functionality have to be moved from ReM to PREDATOR, since former readout

[1] Printed Circuit Board

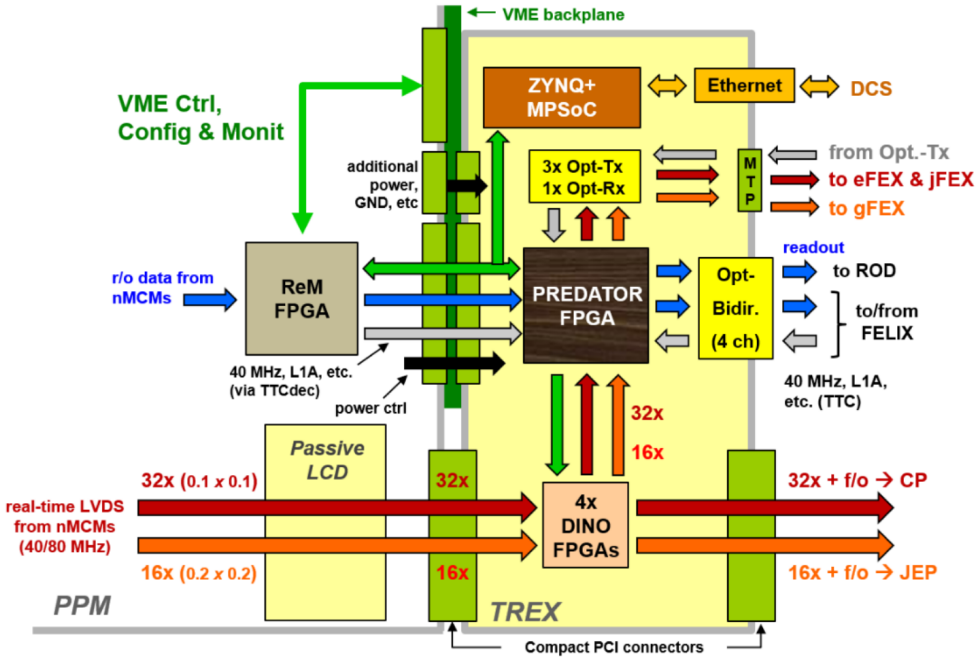

Figure 4.3: Schematics of the TREX module. Real-time paths are marked in red and orange, VME paths in green and readout in blue. Passive LCD and ReM FPGA are components on the PPM. [11]

cards have to be removed to give TREX space. Board configuration works as for the PPM through VME. The connection is established via the CPLD and ReM FPGA for all TREX devices.

The slow control interface to the ATLAS DCS can't be implemented through the CAN module of the PPM because of insufficient physical connections between the modules. A commercial Zynq+ Multi-Processor System-on-Chip (MPSoC) device adopts this task for the TREX. It combines ARM[2] and FPGA architecture. Important information is gathered by Znyq and sent to the DCS via an Ethernet interface. The DCS is then able to power down the TREX separately or together with the PPM, since the power management of PPM and TREX are coupled. Note that the Zynq device is not part of the first TREX prototypes, but will be installed on future versions.

The opto-electrical conversion and their transmission is realised by Samtec FireFly devices. They achieve data rates up to 14.1 Gbps [28]. The board is equipped with four twelve-channel transmitters, where each is responsible for the transmission to one FEX, having one spare in case of defects to avoid module replacements. They are placed close to the PREDATOR on the PCB to keep track lengths short on the electrical high-speed connection. A duplex with four input and four output links is used to send event related data upon an L1A. On its input paths TTC signals are received. The readout interface in use after decommissioning of the legacy system is the Front-End Link EXchange (FELIX), which is in general the system between detector electronics and DAQ. In addition, a FireFly receiver is implemented only on the prototype for testing and debugging purposes of the optical transmission and is not foreseen for the final TREX version.

Apart from the main devices presented, smaller supplementary components are also installed on TREX. PLL[3] chips are used to generate low-jitter clocks as a reference for the Gigabit transceivers. Power Manager ICs[4] enable a safe activation and deactivation of the board and offer a hot-swap functionality to exchange it while under power. They allow further to implement a power sequence to switch on or off the various supply voltages in a fixed order, as recommended by the manufac-

[2] Advanced RISC Machine
[3] Phase-locked loop
[4] Integrated-circuit

turer. Together with the ADCs they are also used to monitor operating conditions, which is essential to guarantee safe operation. These components can provide environmental data for the slow-control functionality to DCS.

Chapter 5

TREX Functional Tests

After development of a prototype board the functionality of its single transmission steps have to be validated, as well as the verification of operating conditions within the specifications of the single devices. In the case of TREX three prototypes were fabricated and intensively tested in a laboratory environment. The setup contains an open VME crate equipped with a few fans, a VME crate controller and a PPM plus TREX in one crate slot. Additional components are added depending on the type of test. This chapter discusses the measurements of the operating conditions, the LVDS data reception in the PREDATOR FPGA, the validation of the optical transmission and a latency comparison of the former and upgraded Preprocessor system.

5.1 Monitoring of Operating Conditions

In order to make a statement about the reliability of TREX operation the operating parameters, i.e. temperatures, voltages and currents, have to be monitored and checked at best under final environmental or extreme conditions. Modern devices like FPGAs are equipped with on-chip sensors that translate the measured value into a digitised voltage and store it in a register. The simplest temperature sensor is for example a temperature-dependent resistance through which a little current is sent. The voltage drop over the resistance can then be converted into a temperature. Similar, currents are measured by the voltage drop they induce in a resistance. Connected with an ADC, the value can be read out via software.
In case of the PPM the amount of parameters is limited. There are 16 temperatures coming from the nMCMs and seven voltages, mainly supply voltages from VME backplane connectors or transformed voltages from DC/DC converters, to be monitored. With the TREX the amount is drastically extended. There are in total 14 temperature sensors on TREX. Each DINO FPGA, each Firefly, two power managers and two sensors on PREDATOR are available. 13 currents are measured indirectly and made accessible by two ADCs. The most important ones are currents from supply voltage pins in order to estimate the power consumption. The number of voltages captured with software tools was chosen to 44. Because of the large amount of components there are many possible voltage values which can be read by software, since some devices need multiple and different supply voltages. Even maximum and minimum long-term voltages are accessible. In the following, the monitored voltages are restricted to supply voltages. Thus, a sum of 94 values is recorded for PPM and TREX.
For the purpose of lab monitoring several applications were created with C++. The first application reads all the desired parameters via VME and writes them to log files as plain text on demand. While data taking the parameters of interest, grouped by temperatures, voltages or currents, are simultaneously displayed in the console. This division is made due to the large amount of parameters. The log

files are picked by a second application which presents them graphically by using ROOT. It creates single and summary plots grouped by boards, devices or quantities. In the end, one ROOT file is generated per measurement that stores all the related plots.
Another application compares actual measured values against definable thresholds. There are upper and lower thresholds for warning and error alarms, based on the device specifications. The limit range is narrower for warnings than for errors. A list of defined thresholds can be found in the appendix A. The outcome for each readout is also displayed in the console and written to log files. Warnings are reported and the operation continues, but in case of an error TREX is switched off by changing configuration bits in an I/O expander chip, that is connected to the power distribution of TREX. This can be undone via VME to power up TREX again.
For both applications, the total testing time and the readout period are configurable at start. The console outputs also indicate the current status of the measurements. But the reported values are not very meaningful in a first instance. The values seem to be rather arbitrary. In order to make them available in conventional dimensions transfer functions have to be applied. These functions are mainly provided by the manufacturer in the datasheet of the single device, as well as the register addresses of the parameters of interest. In case of the measured ADC values some additional factors have to be applied to consider previous voltage division or amplification steps. Also some data must be shifted or equivalently divided/multiplied, since not all provided bits in a register represent the measured value or two registers must be combined to a data word.. The transfer functions for TREX devices are summarised in table 5.1.
The results in the following sections make usage of the presented applications. They discuss the temperature development and power consumption at full load.

Device	Quantity		Data size	Transfer Function
Power Managers LTC2977	Temperature	[°C]	16 bit	$\tilde{b}[15:11] \cdot 2^{\tilde{b}[10:0]}$
Power Managers LTC2977	Voltage	[V]	16 bit	$b \cdot 2^{-13}$ or $\tilde{b}[15:11] \cdot 2^{\tilde{b}[10:0]}$
Samtec FireFlys	Temperature	[°C]	8 bit	b
Samtec FireFlys	Voltage	[V]	Twice 8 bit	$((b1 << 8)\|b2) \cdot 10^{-4}$
DINO FPGAs	Temperature	[°C]	16 bit	$(b >> 4) \cdot \frac{503.975}{4096} - 273.15$
DINO FPGAs	Voltage	[V]	16 bit	$(b >> 4)/4096 \cdot 3$
PREDATOR FPGA	Temperature	[°C]	16 bit	$(b >> 6) \cdot \frac{501.3743}{1024} - 273.6777$
PREDATOR FPGA	Voltage	[V]	16 bit	$b/65536 \cdot 3$
ADC MAX11615 & MAX11617	Voltage	[V]	11 bit	$b \cdot V_{\mathrm{D}}/2000$
ADC MAX11617	Current	[A]	11 bit	$b/(2R \cdot Gain)$

Table 5.1: Transfer functions for TREX devices. The symbol b stand for integer values of the device data in the given size. A tilde expresses a two's complement. R declares the resistance used for current measurements (2 or 5 mΩ) and $Gain$ considers the amplification by an upstream operational amplifier (20 or 50). V_{D} is a voltage division factor (2 or 3) of down transformed voltages to fit the digitisation range of the ADCs [21], [24], [25], [29], [30], [32].

5.1.1 Temperature Development

Heat generation during hardware operation is an unavoidable fact. To ensure a safe running and to avoid material damage, electronic installations need appropriate cooling. This is usually achieved by air- or water-cooling systems and heat sinks. A reference point for a safe temperature operating range is given by the manufacturer's specifications of the single devices. Xilinx recommends e.g. for the PREDATOR FPGA a range of 0-80°C [19]. In order to estimate possible critical heat developments, the board is operated at full load. The maximum possible consumption is simulated by enabling all 64 Gigabit transceivers of the PREDATOR FPGA and by sending stress patterns from the PPM playback memories. These stress patterns are an alternating series of logical 1s and 0s. The permanent toggling is a strain for the hardware. The Gigabit transceivers are very resource-intensive, so that 64 in simultaneous operation bring the PREDATOR FPGA to full use of the capacity. This happens in an open crate with a single occupied slot, equipped with provisional fans for all prototype modules at about 25°C room temperature. Additional heat sinks are mounted on top of the PREDATOR FPGA and FireFlys. Figure 5.1 shows the ones used in these tests. The first one is a monolithic corrugated body made of aluminium with a thermal conductive foil of 2mm thickness placed between the devices and body. The second one consist of three separated parts for a quick exchange of components and a 0.2mm thermal conductive foil. The test procedure is the following: the measurement starts with a cold board and a minimum configuration of the PREDATOR FPGA. After a few readouts the PREDATOR FPGA is reconfigured with all GTHs active. When the maximum temperature is reached the GTHs are again deactivated.
The results for the first heat sink are shown in figure 5.2. The effect of the GTHs is clearly visible by a sharp rise and fall of temperature when they are activated and deactivated. A maximum temperature of about 72 °C is reached on the PREDATOR FPGA and 55 °C for the FireFlys. The DINO FPGAs and the Power Managers seem to remain more or less stable or are only heated because of the material heat flow from PREDATOR. But 72 degrees are rather high compared to the device specification and the open environment. In a fully occupied crate, the inner boards will even be exposed to an increased heat generation due to the air binding and the surrounding boards. Thus, this heat removal may not be sufficient. The temperature development with the second heat sink is given in figure 5.3. It shows that the PREDATOR FPGA is approaching a maximum temperature of about 55 °C, while the hottest FireFly reaches 51 °C. The other temperature curves are similar to those of the previous measurement, whereat the upper DINO FPGAs are a little cooler. The second heat sinks thus offers a better heat dissipation. Note for both measurements that the temperature differences of the single FireFlys and

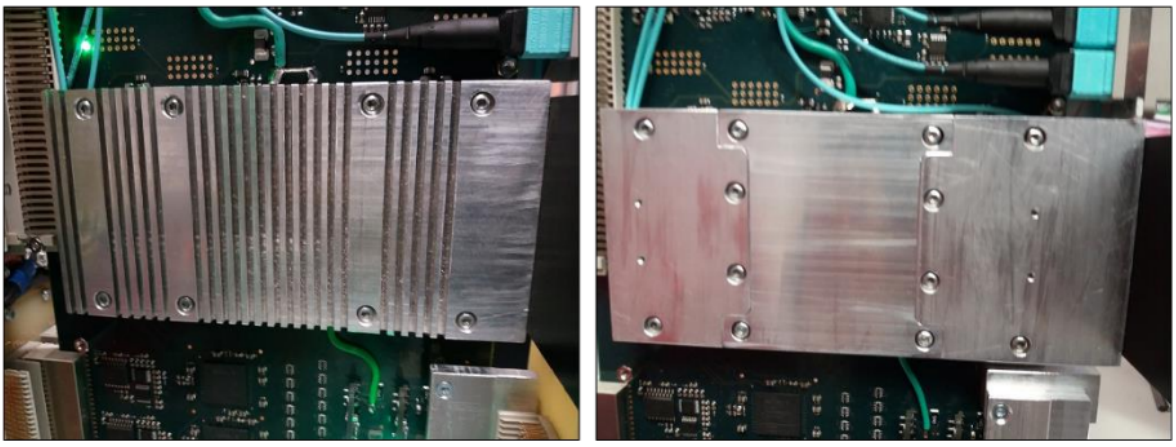

Figure 5.1: Heat sinks used on the TREX prototype modules. The left one is a single corrugated body. The right one consists of three separated parts.

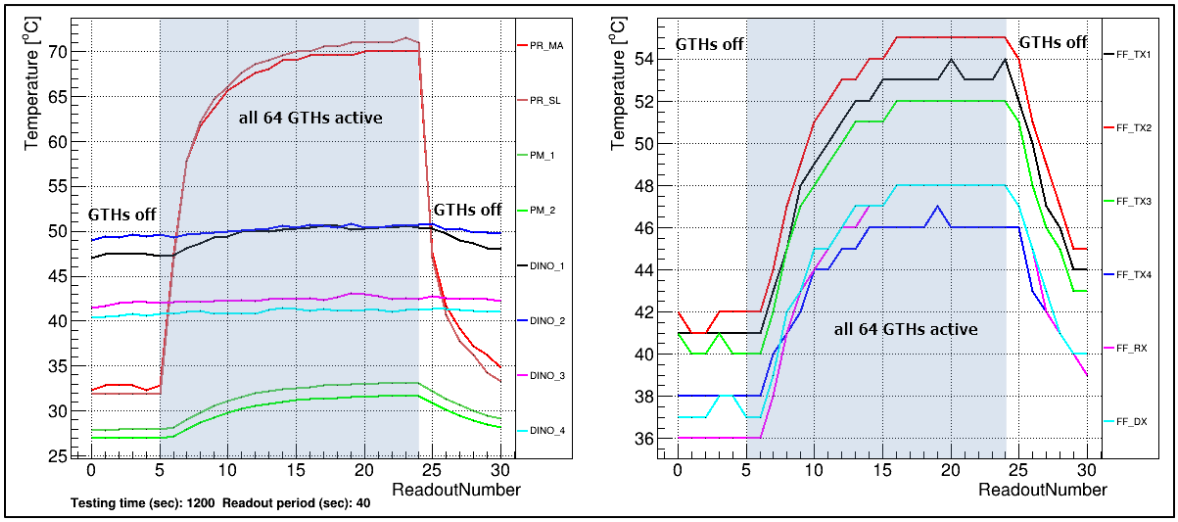

Figure 5.2: TREX temperature monitoring with the first heat sink. It shows a 20 minutes measurement with a 40 second readout period. The left side shows the temperature development of the Power Mangers (PM) , the PREDATOR (PR) and the DINO FPGAS. The right side shows the temperature development of the FireFlys. The blue-shaded areas mark the time with 64 active Gigabit transceivers in the PREDATOR FPGA.

DINO FPGAs are related to the position and orientation of the provisional cooling.
While the first heat sink has a bigger surface area due to its shape, which is usually preferable, the second has a ten times thinner conductive thermal foil that allows a better heat transport from PREDATOR FPGA to the aluminium body. With this heat sink and an adequate air-cooling system it is guaranteed that TREX is able to operate in a safe temperature range. It should also be mentioned that Run-3 operation only requires 29 Gigabit transceivers to be simultaneous active. The heat generation on a single TREX board will therefore be lower compared to these tests.

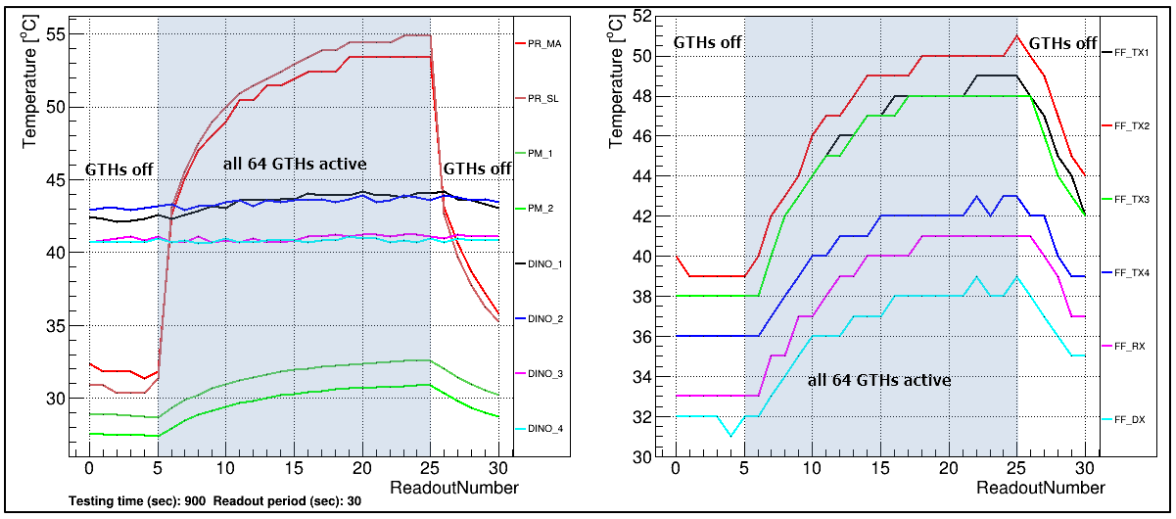

Figure 5.3: TREX temperature monitoring with the second heat sink. It shows a 15 minutes measurement with a 30 second readout period. The left side shows the temperature development of the Power Mangers (PM) , the PREDATOR (PR) and the DINO FPGAS. The right side shows the temperature development of the FireFlys. The blue-shaded areas mark the time with 64 active Gigabit transceivers in the PREDATOR FPGA.

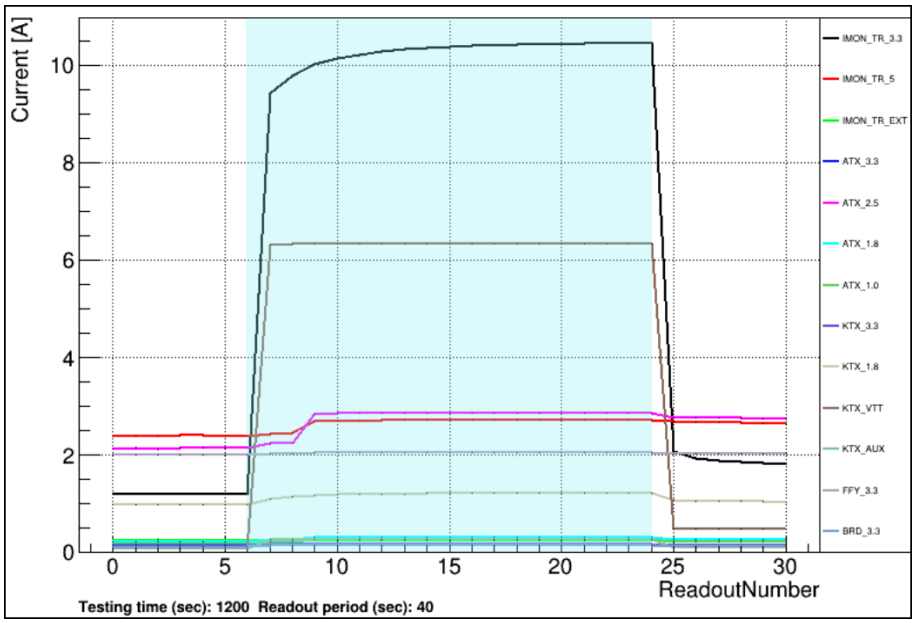

Figure 5.4: TREX current monitoring with the first heat sink. The black (IMON_TR_3.3) and red (IMON_TR_5) curve show currents from the +3.3V and +5V crate's power distribution. The others show various currents to supply pins of DINO (ATX) and PREDATOR (KTX) FPGAs or FireFlys (FFY). The blue-shaded areas mark the time with 64 active Gigabit transceivers in the PREDATOR FPGA.

5.1.2 Power Consumption

In general, power consumption is a key criterion for all electronic hardware. In the present case, TREX is ideally supplied by the same power supply unit as the PPM, since it represents its physical extension, although the possibility of an external supply is kept open. The current prototype is connected to the +3.3V and +5V power distribution lanes of the crate. These lanes provide up to 345A respectively. In Run-2, the PPMs consumed roughly 140A on the 3.3V supply and 160A on the 5V supply per crate [11]. Therefore, it has to be known whether the power demands of a single TREX exceeds the given limits of approximately 100W. This is checked under the some conditions like for the temperature. Current measurements for both heat sinks are shown in figures 5.4 and 5.5. The voltages are stable within their tolerance ranges, that is why the corresponding plots are not presented here. The main focus of both plots lies on the maximum currents of the power distributions. There is a clear increase in current with activation of the GTHs, mainly on the +3.3V distribution. The single current plots show that this trend is the same for all values, but with various characteristics. With the first heat sink the currents rise to approximately 10.5A on the +3.3V lane and 2.7A on the +5V lane, giving 48W of total power consumption. With the other heat sink these currents fall to 9.8A and 2.2A, resulting in 43W. This drop is explained by the temperature difference for both measurements. According to Matthiessen's rule the behaviour of the electrical resistivity of a conductor, in this case copper ρ_{CU}, can be split into a temperature-dependent $\rho(T)$ and -independent part ρ_{D}:

$$\rho_{\mathrm{CU}} = \rho(T) + \rho_{\mathrm{D}}. \tag{5.1}$$

The constant term arises due to defects in the conductors crystal structure. The temperature dependence considers electron-phonon and electron-electron scattering processes, which are reduced at lower temperature [15]. A bigger resistance at higher temperature causes more losses in the board connections, that make the current increase for a still sufficient supply of the single devices. The sec-

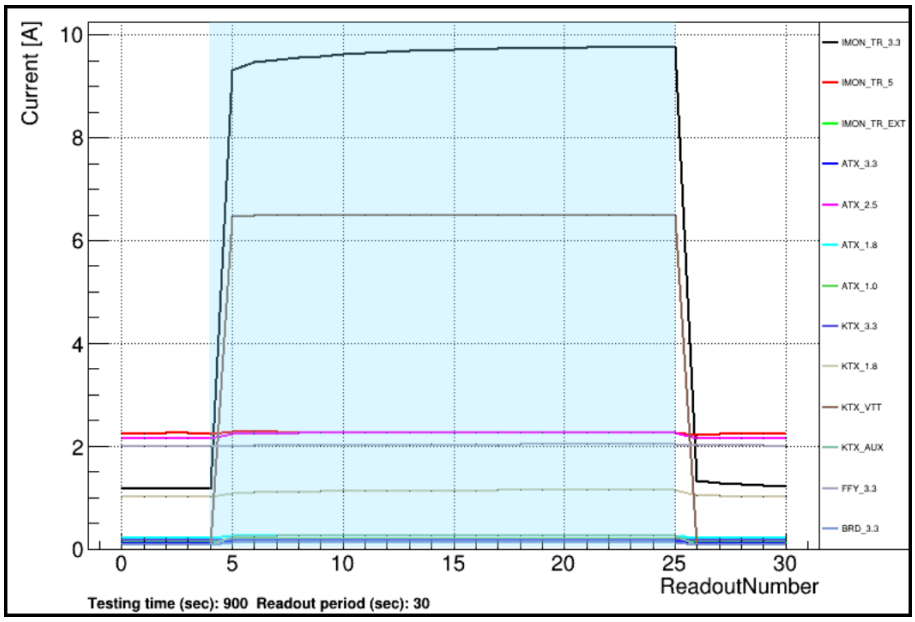

Figure 5.5: TREX current monitoring with the second heat sink.T he black (IMO0_TR_3.3) and red (IMON_TR_5) curve show currents from the +3.3V and +5V crate's power distribution. The others show various currents to supply pins of DINO (ATX) and PREDATOR (KTX) FPGAs or FireFlys (FFY). The blue-shaded areas mark the time with 64 active Gigabit transceivers in the PREDATOR FPGA.

ond highest current in the plots, named KTX_VTT, monitors the supply for the Gigabit transceivers. As expected, they ensure the highest current increase. The results show that TREX consumes less than half of the available power budget at full load. The operation with the same power supply like for the PPM should be safe. Nevertheless the balancing can be optimised. The power budget on the +3.3V lane is 42W per board, which is already reached in the upper measurements. On the other side, the +5V power budget is 58W per board. Some of the DC-DC converters on TREX, used for power distribution, offer the possibility to be supplied from either side. The connection is established by a solder bridge and can therefore be moved to other supply voltage pins. Also the supply of the Zynq+ device will be realised from the +5V distribution lane.

5.1.3 Operation in a new Environment

For test with other L1Calo submodules a TREX prototype has also been installed at CERN in a surface test facility (STF). The whole crate is placed inside of racks, together with the power supply unit and other hardware. The stable operation in this new environment has to be validated. For this reason TREX was monitored with the created application. Figure 5.6 shows the TREX temperature profile for a two hour measurement as a representative. All quantities stayed stable within their tolerances. Here, stress tests like in previous sections were omitted. The alarm application was running afterwards for about 26 hours. No errors or warnings occurred during this period, concluding TREX is running safely at STF. This test can be repeated with more logic or eventually the final firmware in the FPGAs.

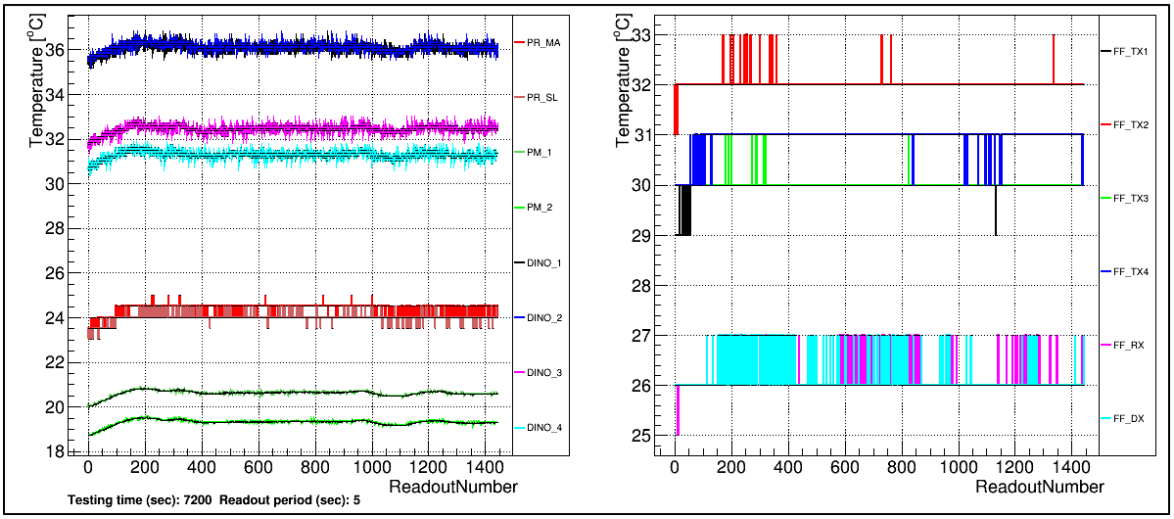

Figure 5.6: TREX temperature monitoring at CERN STF. It shows a two hour measurement with a five second readout period. The jumps in FireFly temperatures are due to the resolution of the temperature sensor.

5.2 LVDS Data Reception and Alignment in PREDATOR FPGA

As explained in section 4.3, the real-time data of the nMCMs is duplicated on 48 physical links in the DINO FPGAs and send to PREDATOR. But the reception of this data has various timing on each link due to different signal path lengths and processing times in the DINO FPGAs or various routing through the PCB. This data has to be aligned on all links for a synchronous operation and correct BC-wise data selection.

To check an error-free transmission and reception, the playback memories in the CALIPPR FPGAs are configured with known data. These memories start sending their contents on command. On PREDATOR side, these signals pass through several processing steps according to figure 5.7. At first, after conversion of the differential signals to single-ended form, they are send through a delay element called IDEALYE3. It is an Xilinx IP core and simply used for delaying signals by shifting the sampling point at the next input stage like figure 5.8 illustrates. ISERDESE3 is another IP core responsible for the parallelisation of the incoming serial streams in packages of four bits. How to pick an combine these packages to the original twelve bit word is decided by the bitslip and synchronisation logic. When setting up the board, various commands are sent to the single devices. The PREDATOR receives synchronisation commands, whereat the nMCMs send training patterns to the bitslip logic. Since the LVDS frame bits are always a logical 1 at the beginning and a logical 0 at the end of the data word, these training patterns are a series of six consecutive 1s and six consecutive 0s. The receiving logic recomposes the incoming bits to this pattern and locks its final decision for future operation.

After this stage, the recombined data is monitored by calculating and comparing the parity bit and checking synchronisation. In case of former CP data, the streams are additionally demultiplexed before further processing and preparation for optical transmission occurs. The received data can also be transferred to a SPY memory for testing purposes. These are 18-bit wide memories with 1024 registers on each transmission channel. At the end, those entries are read via VME and compared to the data that has been sent.

To emulate the processing the data undergoes on its way to PREDATOR, dedicated software functions are used to make playback and SPY data comparable. On CALIPPR side, single logical blocks can be bypassed for simplicity of these tests. Also the synchronisation FIFO, the FIR filter and the pedestal

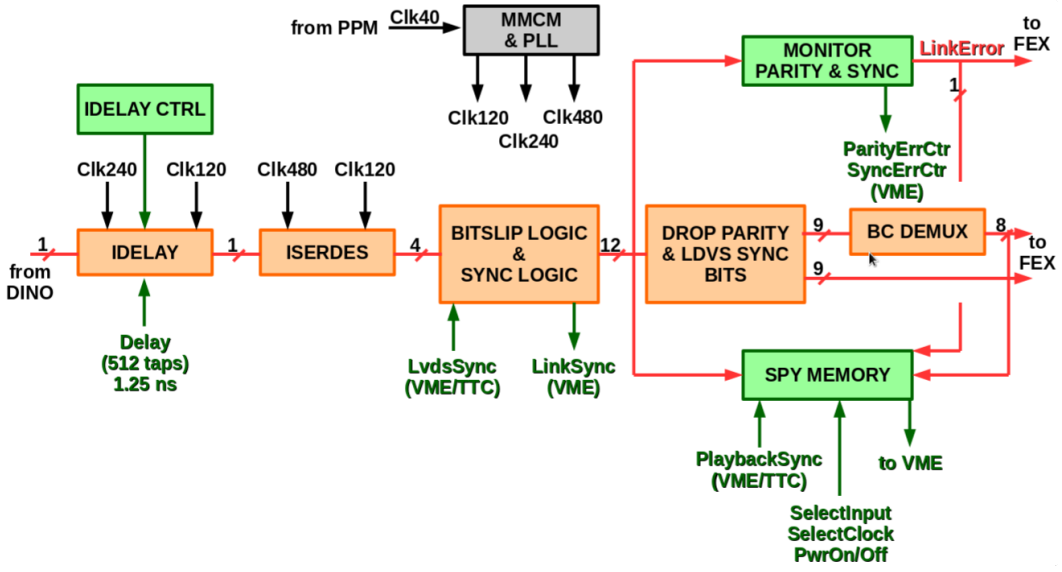

Figure 5.7: LVDS data input to PREDATOR FPGA. IDELAY, IDEALYCTRL and ISERDES are Xilinx IP cores. Components in green are for configuration and monitoring, while orange elements indicate real-time processing steps [27].

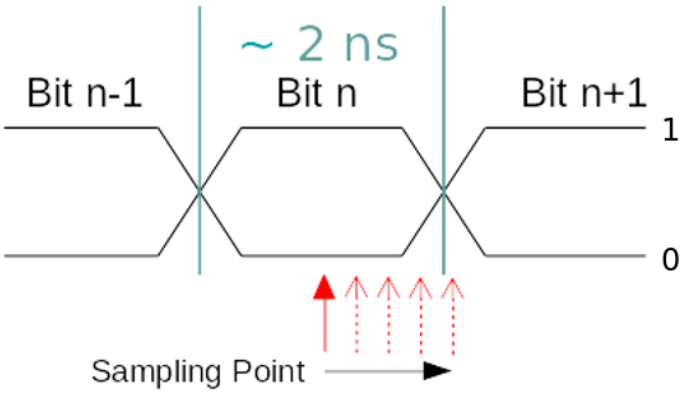

Figure 5.8: Sketch of IDELAYE3 functionality. Logical levels are indicated on the right side. Single bits are depicted as eye diagrams. IDELAYE3 shifts the point in time when incoming signals get sampled. At this stage, transmission occurs at 480 Mbps that correspond to a bit length of 2.083ns.

correction can be configured for an effectless routing of the playback data (compare figure 3.4). The LUT, the BCID decision logic, the BC-multiplexing and the packing of the parity bit can be bypassed for CP links, so that the pure playback data is received by PREDATOR. This does not hold for JEP links. Here, only the bit truncation as well as the parity packing can be bypassed. The simulation can be configured to consider this different processing.

5.2.1 Characterisation of IDELAYE3

Although the functionality of the IDELAYE3 is already known, there are a few characteristics that have to be determined. In general, this component is a 512 tap delay line that offers the possibility for a calibrated time delay with voltage and temperature compensation [22]. This built-in-self-calibration (BISC) process runs permanently in the background and is also responsible for an initial non-zero delay tap value, even if no delay applied, that wants to compensate internal skews between clock and data insertion at the first capturing flip-flops of ISERDESE3. These skews appear because of different routing of clock and data in the FPGA. Typically, the outcome is a tap value between 45 and 65 which is called Align Delay. An IDELAYCTRL component must be instantiated to make this happen. The resolution of a single tap is given by 2.5 to 15 ps [19], ending up with total delays of 1.28 to 7.68 ns. The exact value is also connected to the reference clock of the IDELAYCTRL, but explanations how to adjust the clock for different tap sizes are missing. Nevertheless, the tap resolution and with it the overall delay can be determined [18]. When instantiating the IDEALYE3 component in the firmware, certain attributes have to be set. One of these is an initial Delay Value in the range of 0 - 1250 ps. If it is non-zero, the resulting tap number will be bigger then the Align Delay. Thus, two firmware version with different Delay Values $d_1 > d_0$ are created and the corresponding tap numbers n_1 and n_0 are read via VME. The tap size resolution t_{res} can then be calculated according to

$$t_{\text{res}} = \frac{d_1 - d_0}{n_1 - n_0}. \tag{5.2}$$

Doing this with even more firmware versions, average resolutions of 4.35 - 4.46 ps for different links are calculated. This gives a total delay of 2.22 - 2.28 ns for the whole delay line. With the given transmission rates of 480 Mbps at this stage, the IDELAYE3 component is therefore able to delay incoming signals by more then a bit length. The exact numbers may vary with different boards but they should be comparable.

5.2.2 Test Procedure

The precise insight of the reception process is crucial for the avoidance of incorrect sampling. Signal sampling during the transition between logical levels, corresponding to the line crossings in figure 5.8, leads to transmission errors because the logical state is not clearly defined in this period. The fine-tuning of the sampling point must therefore be carried out on each receiving system. With the use of the introduced delay component two things have to be determined:

1) valid delay tap ranges.

2) the stability of a fixed delay.

The idea is simply to send data from the playback memories, store it in SPY and search for transmission errors in data and LVDS frame bits. This is repeated for each delay tap or for a fixed delay over a configurable period of time and controlled from software. The results are presented in histograms. Because of the different bypass options for CP and JEP links, different playback patterns, that means the series of transmitted data from the playback memory to PREDATOR, are used on the respective links. On CP side either a constant pattern, as for the synchronisation of the receiving logic on PREDATOR, or a ramp pattern with increasing numbers is sent, that provokes a more stressful transmission. The pattern used on JEP side is a series of growing pulses in increment steps of four with a zero as every second entry, since the Peak Finder algorithm only permits a non-zero data transmission at every second bunch-crossing. The decision in favour of four-increments was made because the sequence of LUT and channel summations is effectively truncating the two least significant bits of the 10-bit data word in the applied bypass configuration of the CALIPPR FPGA. This effect is not considered in the existing simulation software.
All patterns start with an unique identifier that flags the beginning of the receiving data in the SPY memory. Activation of SPY and the start of playback transmission can not be controlled from software in perfect timing. The first registers of SPY will be therefore filled with some dummy values at these tests. But since the SPY exceeds the playback in size by a factor of four, the whole playback pattern is at least three times completely contained in SPY. Finally, to make a statement about the alignment of data, the register addresses of the unique identifier will be compared to each other on all links. To cover all receiving logic, a extended version of this test application checks the BC demultiplexing and those alignment in the same manner.

5.2.3 Delay Scans

As explained, the first tasks is the determination of valid delay tap ranges. The software processes the following sequence: set a delay tap number, synchronise the receiving logic, enable playback transmission, read the received data for analysis and start all over with an increased delay tap number. The data is thereby sent in packages of twelve bits. Ten bits are playback data while the other two are the LVDS frame bits. Errors in data and LVDS frame transmission are treated separately for analysis in the first instance. But observations revealed the same error tendency, which is why histograms with the combined results are created. Since the playback pattern is checked up to three times in SPY, the maximum number of errors is 768 for data and LVDS frame bits respectively. If the application is unable to detect the unique identifier, the number of combined errors is set to a maximum of 2048, twice the register size of SPY.
The resulting histograms are presented in figures 5.9 and 5.10. Link numbers up to 31 stand for CP connections, higher numbers for JEP connections. The results can be interpreted in analogy to eye diagrams. These diagrams are usually recorded with an oscilloscope, showing the overlay of the same

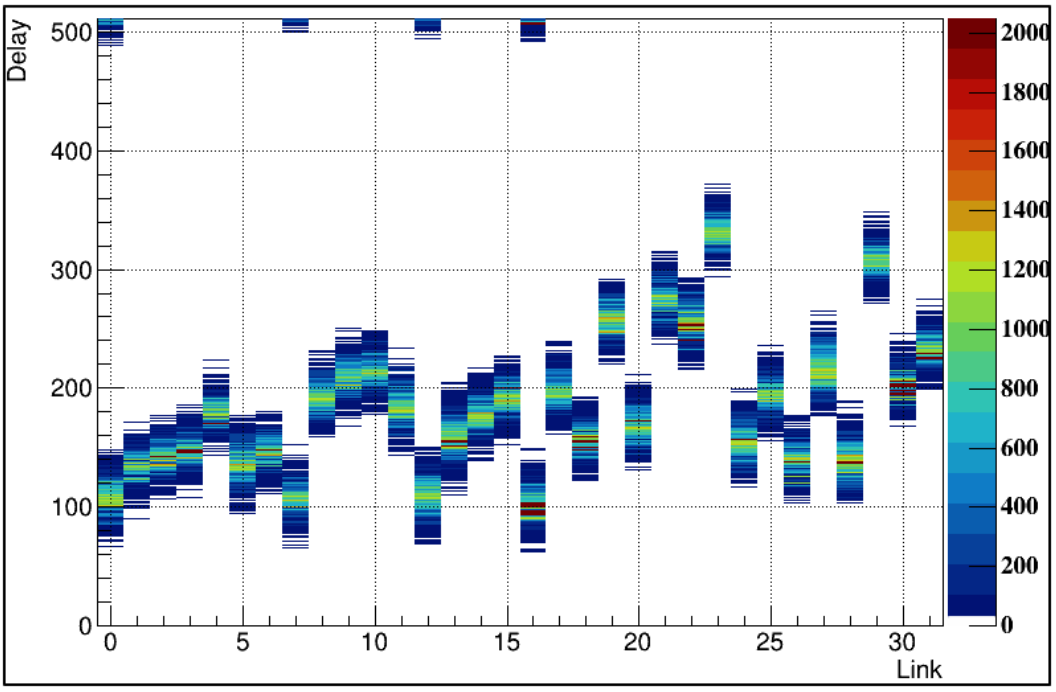

Figure 5.9: Delay scan for CP links. The link number is plotted against the delay tap number of IDELAYE3. Colours indicate the amount of detected transmission errors for a given delay tap. White areas stand for correct transmission.

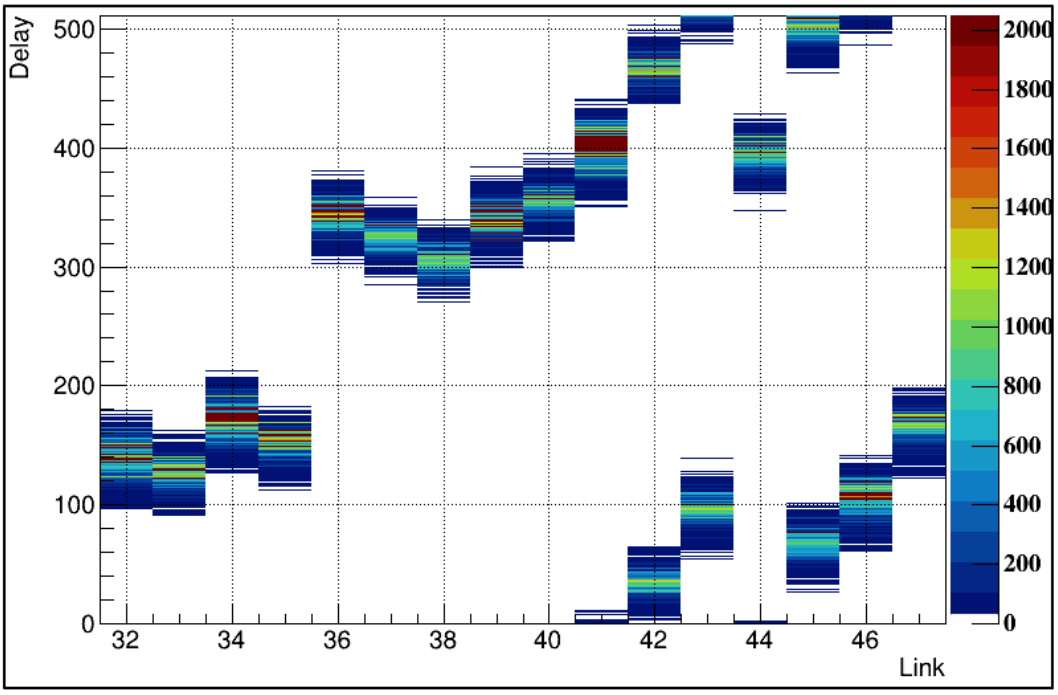

Figure 5.10: Delay scan for JEP links. The link number is plotted against the delay tap number of IDELAYE3. Colours indicate the amount of detected transmission errors for a given delay tap. White areas stand for correct transmission.

digital signal for different points in time. Coloured areas, indicating transmission errors, correspond to the eye margins where transitions of logical levels take place. They also demonstrate nicely by the colour changes that the transmission errors increase towards the middle of these regions. This is conform to an increased intersymbol interference because of the closure at eye margins.
Some links also show the beginning of a second eye margin. This reassures the characterisation of the delay element and the determined total delay. The results vary slightly with the used playback pattern. A more stressful pattern like the ramp broadens the error ranges in comparison to a constant one. But the positions in the histograms stay the same. The results are in good agreement with all TREX prototypes while using the same PPM. Only very slight shifts in the invalid delay ranges were observed, whereat the overall arrangement was the same on all TREXes. These shifts are most likely caused by manufacturing tolerances of the PCB.
The histograms clarify that there are enough valid delay taps that allow a safe sampling of receiving data in PREDATOR. Ideally, sampling takes place in the middle of an eye. Converted to a delay tap number this would be e.g. approximately 320 for link 0. But to assess whether a once configured delay is stable and safe enough for operation, a long-term test must be performed.

5.2.4 Stability of Valid Delays

Valid delay tap numbers with equal distance to the visible or estimated margins outside the diagrams are picked from the previous plots for each link separately and used for configuration of the respective delay elements. The numbers are defined in a text file and read by software for the setup. Afterwards, transmission, reading and analysis of data is repeatedly executed for a adjustable time interval and period. Figures 5.11 and 5.12 show the outcome for an 24 h measurement with a readout every two seconds. This readout cycle was chosen because this time span was sufficient to test all CP links once.

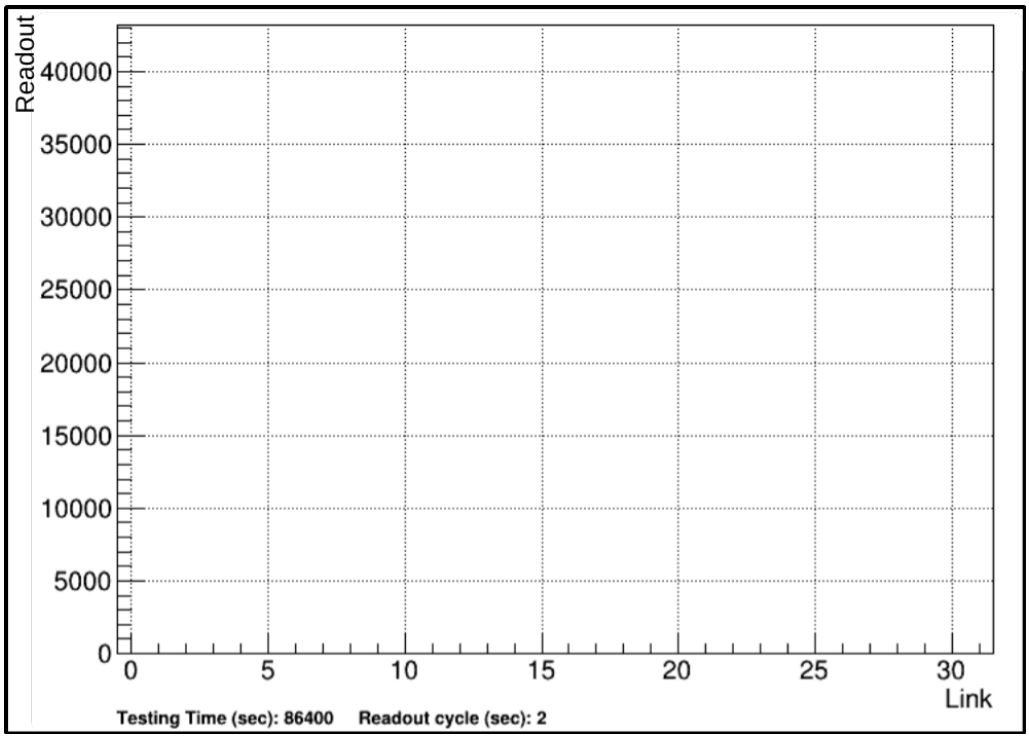

Figure 5.11: Stability of fixed delays for CP links. The stability was tested for 24 hours in a two second readout cycle. The histogram shows the amount of errors for a given link and readout number. No errors were detected.

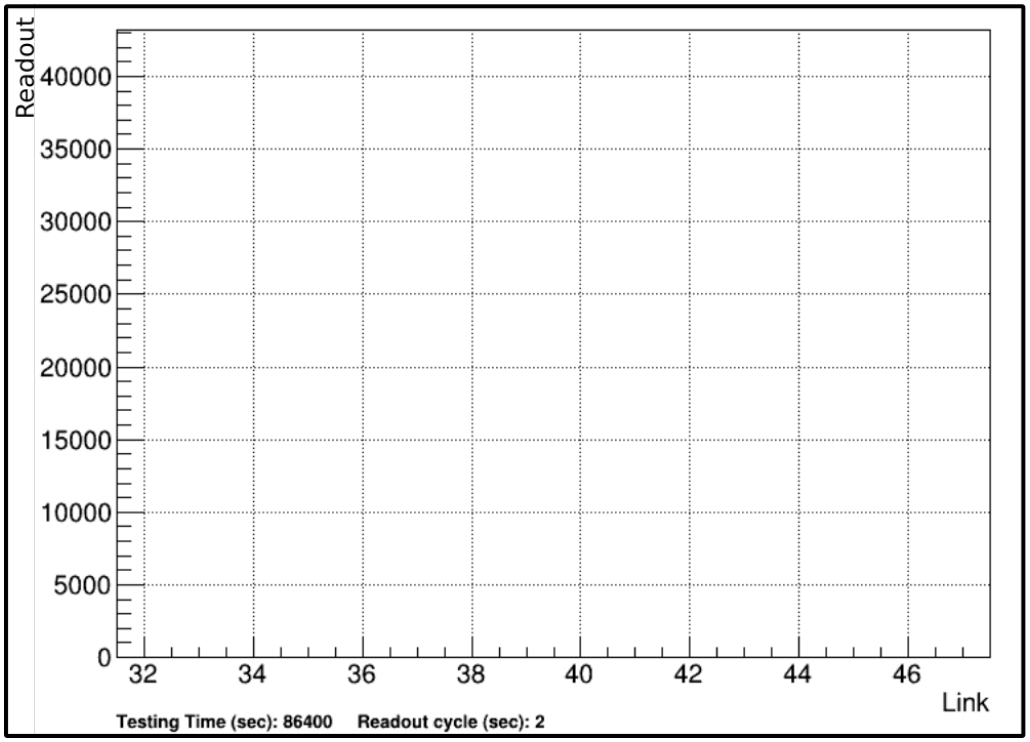

Figure 5.12: Stability of fixed delays for JEP links. The stability was tested for 24 hours in a two second readout cycle. The histogram shows the amount of errors for a given link and readout number. No errors were detected.

No transmission errors occurred during operation. A safe LVDS reception in PREDATOR FPGA is thus guaranteed for longer operation. Setting the delay tap close to a margin results in occasional errors in the transmission. This was seen in other stability tests with shorter measuring time.

5.2.5 Data Alignment

It remains the question whether data is aligned on all links. For this reason, the start addresses of the unique identifier in the SPY memory are presented in histograms for all links and delay tap numbers. Certain links showed a different outcome then the others in the first attempts. All these links, however, only belonged to one I/O bank. A bank is a group of I/O pins that share some commonalities like a common power supply or similar. The identified problem was attributed to internal signal routing in the FPGA. The enable command for the SPY memory is sent simultaneously to all links. But the ones of the identified I/O bank received this command earlier compared to the others. This was overcome with the instantiation of some synchronisation flip-flops, that forward the command signal only with a rising clock edge. After the implementation, the final results in figure 5.13 and 5.14 were generated. They show that data is aligned on all links and for all delay taps in a given test mode. The alignment was also maintained for a fixed delay during the stability tests.
Comparing the register addresses of CP and JEP links, a constant discrepancy of five registers is observed. There are two reasons for this difference. The data on JEP side undergoes a longer processing in CALIPPR because of the various bypass options. There are four more steps in processing than on CP side, consisting of the LUT, the BCID logic and the channel summations. This is clocked, like the writing to SPY, with the BC frequency of 40 MHz. The second reason is pattern related.

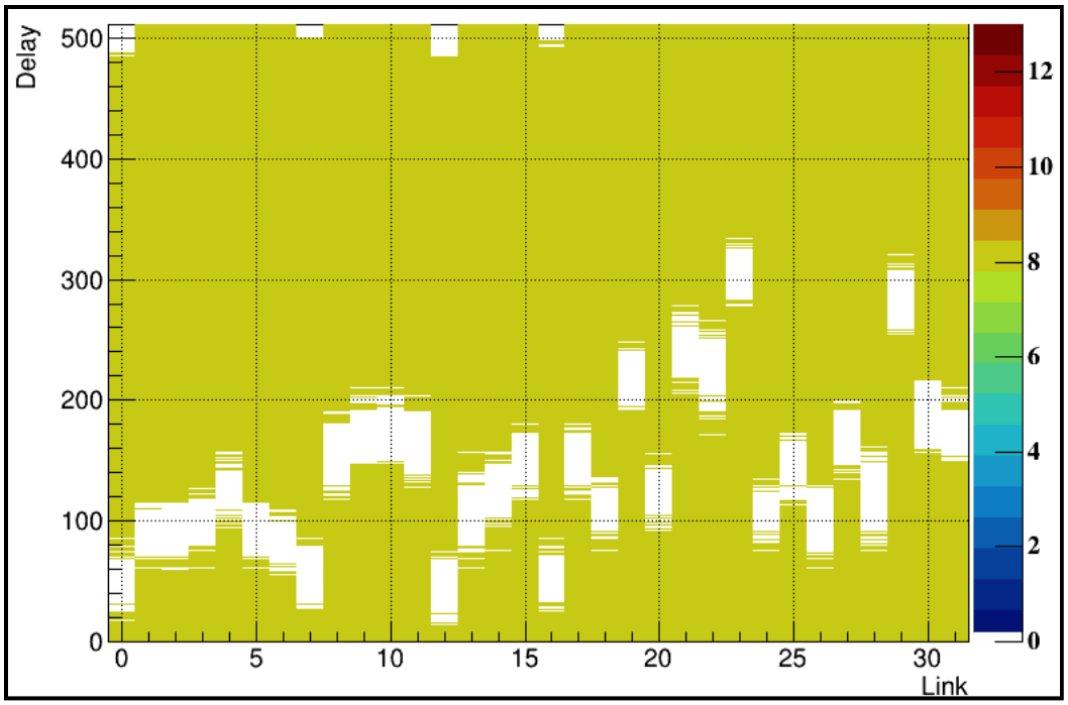

Figure 5.13: Data alignment in PREDATOR FPGA for CP links. Defective delay regions are skipped for this histogram. The unique identifier is found in SPY at register number 8 on all CP links.

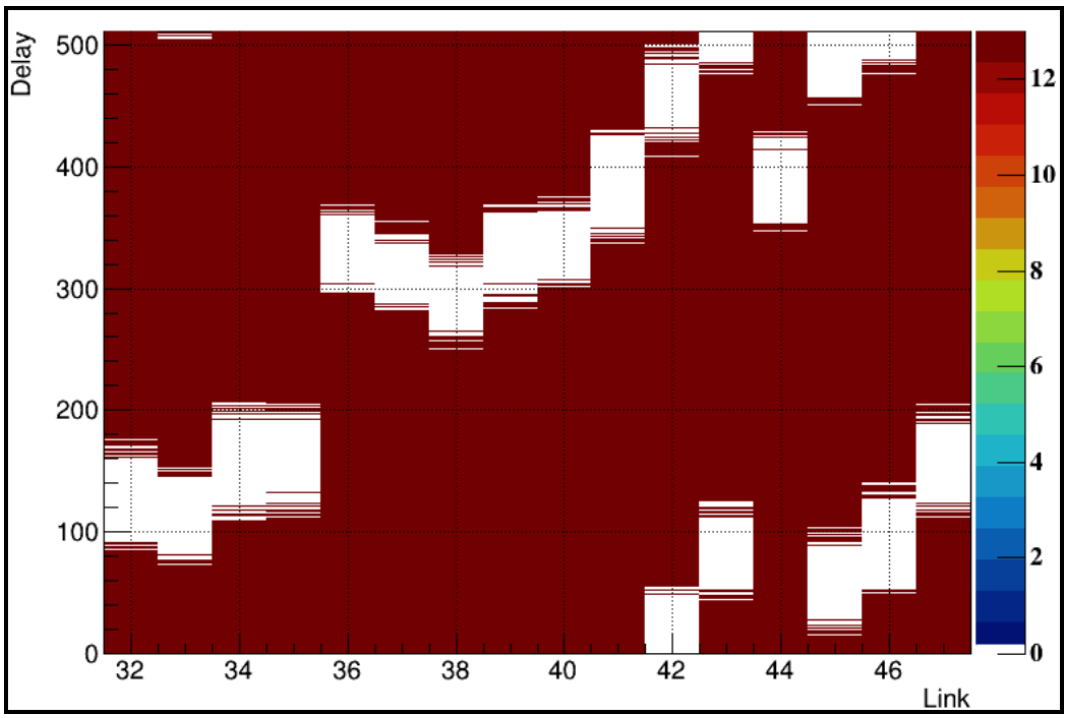

Figure 5.14: Data alignment in PREDATOR FPGA for JEP links. Defective delay regions are skipped for this histogram. The unique identifier is found in SPY at register number 13 on all JEP links.

The unique identifier is not the very first entry in the JEP pattern. A leading number is again blocked by the peak finder algorithm because of missing comparative values. Therefore, the pattern starts with a zero and the start address on JEP side is shifted by another register.

5.2.6 Common Test Mode

The general reception of nMCM signals is guaranteed according to the previous tests. To validate also the BC demultiplexing the test applications are extended by this option. The analysis has to adapt the new input to SPY memory how figure 5.7 indicates. For a correct assignment of data PREDATOR has to check the BCMux flag. During this step other data unrelated information, e.g. the LVDS frame bits, is dropped. However, the monitoring in PREDATOR from 12-bit data allows the readout of a LinkError flag that informs about the loss of synchronisation, i.e. incorrect LVDS frame bit reception, or an inaccurate parity bit, whose transmission is suppressed here. So, the input size switches to nine bit.
Since it is a common test, the bypass options and playback pattern for the JEP side are now also applied for CP links. The pattern is modificated by inserting an additional zero in between two non-zero values to incorporate possible effects of the BCID decision logic on CP links. The transmitting channel varies every bunch-crossing on these links. The BCID sets every second value, amongst other reasons, to zero to allow this multiplexed transmission. BCID and BC-multiplexing can collide at this point resulting in data blocking. As for the single test modes the writing to SPY occurs at the BC frequency of 40 MHz.
Figure 5.15 shows the delay scan in the common mode. Basically, the outcome here is the same as in figures 5.9 and 5.10 in one histogram. The demultiplexing is here expressed in the fact, that two neighbouring link numbers show largely the same characteristics, since it is the same physical connection. The correct bunch-crossing assignment is demonstrated in figure 5.16. The register address is again the same on all channels. It matches the number 13 as for the JEP tests. This can be attributed to the bypass options and the used pattern.

This section demonstrated that LVDS reception in PREDATOR FPGA is safe and stable for appropriate delay taps. The bunch-crossing assignments from the nMCMs and their temporal alignment on all channels are not affected by the transmission to TREX or the processing in its FPGA. Since the valid delay ranges may vary with the use of different TREXes and PPMs, such delay scans have to be run for the final system before operation to find and set a suitable delay tap. It is still under discussion whether this task will be implemented in firmware to set this automatically when the extended Preprocessor system is switched on, or adopted by software resulting in manual adjustment. The duration of such a delay scan, which took 15 minutes in this case including all analysis, has to be considered. A subsequent application already reduced this by a factor of two. But this can still be unpleasant regarding 32 TREX modules. An implementation in firmware is expected to be faster.

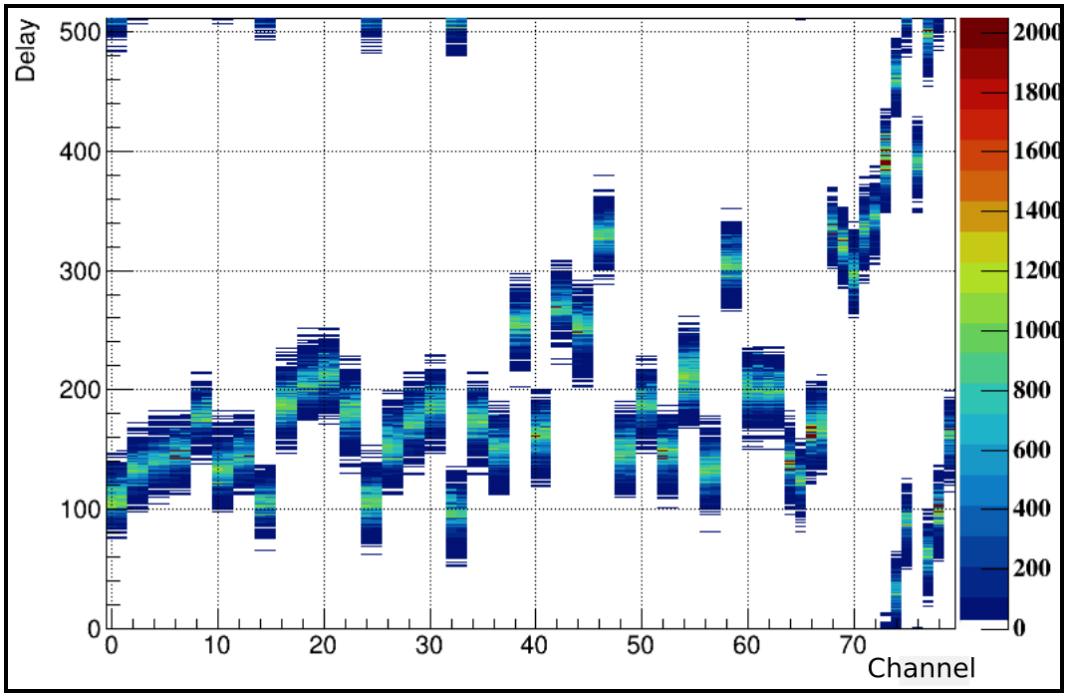

Figure 5.15: Common delay scan. Channel numbers up to 64 are CP related connections. Here, two neighbouring channels use one physical connection. Higher channel numbers are JEP related and have one physical connection per channel.

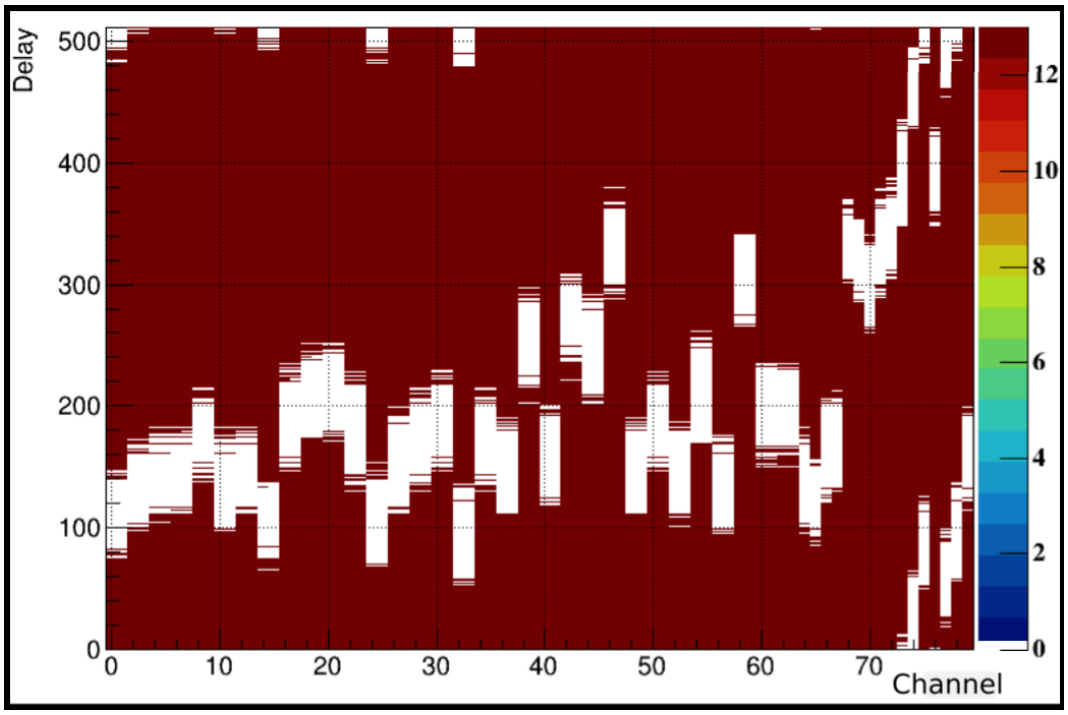

Figure 5.16: Demultiplexed data alignment in PREDATOR FPGA. Defective delay regions are skipped for this histogram. The unique identifier is found in SPY register 13 on all channels.

5.3 Characterisation of Optical High-Speed Transmission

After the test of the general conditions on TREX and the validation of the correct reception and alignment in its main processing unit, the optical transmission to the subsequent systems of TREX has to be evaluated. This is done as stand-alone and joint tests with a jFEX module. At first, the test environment and settings are described together with a more general description of the serial transceivers. The TREX components that participate in the optical transmission are the PREDATOR FPGA and the Samtec FireFlys.

5.3.1 Test Environment

The main tools in use for the characterisation of the transmission are the Vivado Design Suite and an Integrated Bit Error Rate Tester (IBERT). Vivado is the design environment for Xilinx FPGAs. It enables synthesis and analysis of Hardware Description Language (HDL) designs. IBERT is an IP core for monitoring and evaluation of the Gigabit transceivers. It includes pattern generators and checkers and accesses the transceiver ports and their configuration, such as line rates or reference clocks [20]. The combination of these tools allows the creation of statistical eye diagrams to judge the optical connections. These diagrams show in principal the same like eye diagrams on an oscilloscope and are therefore a valuable instrument to estimate the quality of transmission. Only the recording differs from the oscilloscope measurements.
Figure 5.17 clarifies the data taking method. The bit length is referred to as an Unit Interval (UI), since this depends on the clock speeds. The whole picture is split into 255 $\times$ 65 pixels. The final resolution is, however, dependent on the measurement settings. In general, two sampling points are needed for the statistical eyes. The first one is the fixed Data Sample. It stays in the centre of the eye where sampling is valid and acts as a reference. The other point is the Offset Sample. It moves across

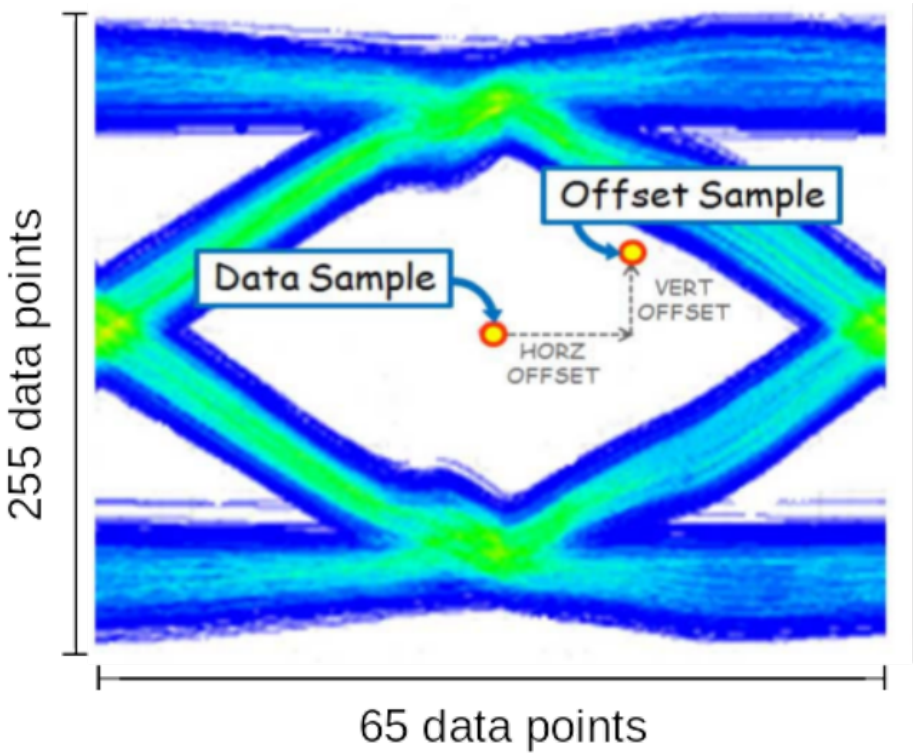

Figure 5.17: Measuring principle of statistical eye diagrams. The Unit Interval is divided into 65 and the voltage amplitude into 255 data points. The whole eye diagram is a map consisting of up to 16575 pixels [23].

the picture by varying the sampling time and voltage threshold. The offset increments in vertical and horizontal direction are adjustable from one to 16. Bigger increment steps result in a bigger pixel size and therefore lower resolution. The outcome of the sampling at the two reading points is permanently compared against each other. This is done until the dwell bit error rate (BER) is reached. The number is set at the start of the measurement and can be adjusted from 10^{-5} to 10^{-15} in power of tens. By this, a BER is calculated for each pixel as a function of the offset.
The results are presented in a coloured picture in common logarithmic scale. Additionally, algorithms for smoother pixel transitions can be applied in Vivado. By means of the final eye shape and opening, statements about noise, jitter or equalization effects can be made. Because of the amount of statistics in these diagrams, they are expected to be more closed then an oscilloscope picture. In order not to have to look at each eye diagram individually for evaluation, quantities are provided for comparability. The first one is the Open Area which represents the opening of the eye by the number of pixels that have reached the dwell BER. The width of the eye is expressed in percentage of the UI.
A little arithmetic example shall illustrate the measurement time for a single eye diagram. Assuming a transmission speed of 10 Gbps and a dwell BER of 10^{-10}, then a single pixel is checked within one second. Choosing increments of one in both directions, the whole measurement would take more than 4.5 hours, because of the amount of pixels. Bigger increments shorten the measurement time by the respective factor. To avoid starting a measurement for each transmitting channel individually, TCL[1] scripts are used in Vivado to run the eye scan in one pass for all detected connections.

5.3.2 Transceiver Architecture

Insight into the architecture of participating transceiver and the effects of possible modifications are useful for a later interpretation of the results. Generally, the present line rates of TREX optical communication to FEXes, i.e. 11.2 Gbps, are quite high for a simple transmission and correct reception without signal enhancements on the PCB. Insertion losses increase with higher frequency and trace lengths. Additional attenuation effects like reflections may appear on the transmission path. To compensate these issues, signals can be boosted on transmitting and receiving side. Enhancement on transmitting side is referred to as pre-emphasis. It aims to balance low and high frequency boosts considering the known attenuation on the transmission path to rebuild the original signal on receiving side. In contrast, signal enhancement on reception is called equalization. There are various equalization techniques optimised for different rates and distances.
The architecture of the GTH transceivers is sketched in figure 5.18. Pre-emphasis is here realised by a FIR filter. Equalization methods on reception are called Continuous Linear Time Equalization (CTLE) and Decision Feedback Equalizer (DFE). CTLE does in principle something similar like the pre-emphasis. Because of the low-pass characteristics of insertion and transmission, CTLE applies a high-pass filter for balance. Improper equalization yields in under-equalized, that means a closed eye, or over-equalized eye diagrams, which are narrow but wide open and can result in high jitter. A disadvantage of CTLE is the simultaneous amplification of noise, if present. DFE avoids this, because it only subtracts, after the decision which symbol is most likely received, the distortions of the known impulse response for the following symbol. Intersymbol interference and noise amplification are prevented in this way, but the behaviour of the transmission path has to be resolved for a appropriate equalization. DFE is usually activated in combination with CTLE. Both methods provide furthermore an auto-adaption that operates continuously to avoid manual tuning.
Figure 5.18 also shows that the eye scans are recorded after equalization. In principle, the characterisation of the optical transmission shall be comparable to other systems. The eye scans in this thesis

[1] Tool Command Language

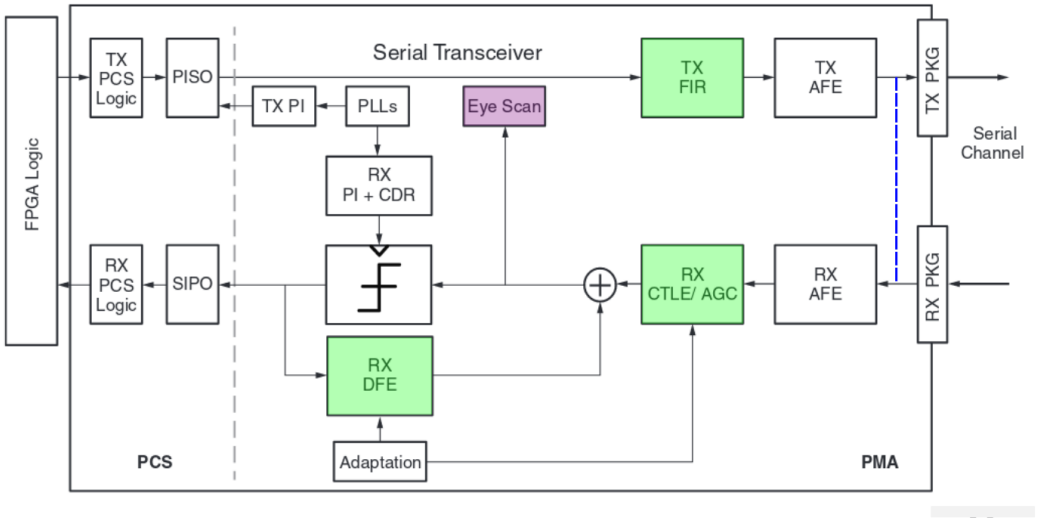

Figure 5.18: GTH transceiver architecture. Elements in green are used for signal emphasis or equalization. The tap of signals for eye scans is marked in purple. Tx stands for transmitting and Rx for receiving logic. The blue dashed line indicates the loopback for near-end PMA tests. Acronyms not explained in the text can be found in the appendix B [13].

are therefore taken under default settings, meaning an active CTLE in auto-adaption and no DFE and pre-emphasis, although these settings should be optimised for the final operation. Anyway, DFE shows to worsen most results on TREX although it is recommended to use at line rates of 10 to 15 Gbps at complex interface applications [14]. A deactivation of CTLE at the given line rate may result in a closure of the eye on PCB paths.
Apart from the enhancements in the FPGA, signals can be boosted in the FireFly transmitters. This option is also neglected here. Generally, signal enhancement would also result in a higher power consumption, that would have to be taken into account for the monitoring of operating conditions.

5.3.3 Test Settings and Evaluation Procedure

The characterisation is carried out under different test settings. Stand-alone tests are performed as loopbacks either on the electrical or optical paths. The electrical loopback is referred to as near-end Physical Medium Attachement (NEPMA) test. The entire signal path is basically pictured in figure 5.18. Signals emitted from the transmitter reach the receiver directly and thus are not leaving the FPGA. NEPMA serves essentially as a reference, since it is expected that it shows the best results due to the short signal paths.
For the optical loopback, signals are sent from the Gigabit transceivers to a FireFly transmitter and led through a optical fibre to the FireFly receiver on TREX and back into the PREDATOR FPGA. A special loopback device adopts this for the duplex. Thus, this measurement is taken device-wise.
Joint tests are performed with a jFEX final prototype module in either direction; from TREX to jFEX with four simultaneously active FireFly transmitters and from jFEX to TREX in parallel use of receiver and duplex. Here, the signals are routed over 30 m long fibres. In all settings PRBS[2]-31 test patterns are transmitted at the final line rates of 11.2 Gbps. These are binary sequences that approximate white noise and are thus hard to predict. Only the bidirectional FireFly duplex is also tested under different data rates of 4.8 and 9.6 Gbps, since it will receive from or transmit to FELIX information at these speeds [11].
The PREDATOR FPGA is configured with an IBERT test firmware during these investigation, drop-

[2]Pseudorandom binary sequence

ping other functionality. With the given amount of 68 high-speed optical links, whereat 52 are output and 16 input links, and the various settings a large amount of eye diagrams is created. The evaluation of each single eye can be overburdened. Thus, this tasks needs to be adopted by software. Luckily, IEEE[3] provides standardised requirements for optical operation up to 100 Gbps [31]. Since the measuring setup only allows the validation at reception, a mask for stressed receive characteristics is applied. This is a diamond-shaped mask which is placed symmetrically in the middle of an eye diagram and covers 44 % of the UI and 34 % of the voltage amplitude. It defines a minimal eye opening for a reliable operation on receiver side.
Eye scans taken with Vivado and IBERT can be exported to text files, where all necessary information about the measurement settings and the single pixel BERs is stored. These text files are used for analysis. At first, IEEE mask boundaries are computed and the contained pixels, that are within or touched by the mask, are checked against the dwell BER. For a more flexible assessment the data points are classified into valid, invalid and almost valid. Valid pixels reach the dwell BER while almost valid ones are of the same magnitude like the dwell BER but with a higher pre-factor. All other points are invalid. Dependent on the increment steps, the quantity of contained pixels varies widely. This can be just 5 for 16 increments or up to 1500 for single increment steps. Results are provided in the console output and in summary histograms that plot the number of invalid and almost valid pixels in two bins against the link number. The exported Vivado files are further used to recreate the eye diagrams with a drawing of the mask.

5.3.4 Results

In order to avoid unnecessarily long measuring times, a dwell BER of 10^{-8} has been agreed upon for all test settings with single increments in both directions. Only for the test from jFEX to TREX the increments are set to eight. The time for the joint tests was limited and the focus lied of course in the transmission to jFEX. The individual test results are presented in the following.

Near-end PMA

The eyes in NEPMA test are broad, wide open and symmetrical for all 64 GTH transceivers. Reported Open Areas are in the range of ten to twelve thousand with UIs bigger then 73 %. Also the evaluation software reported only valid data points. A NEPMA Vivado eye scan is provided in figure 5.19 for

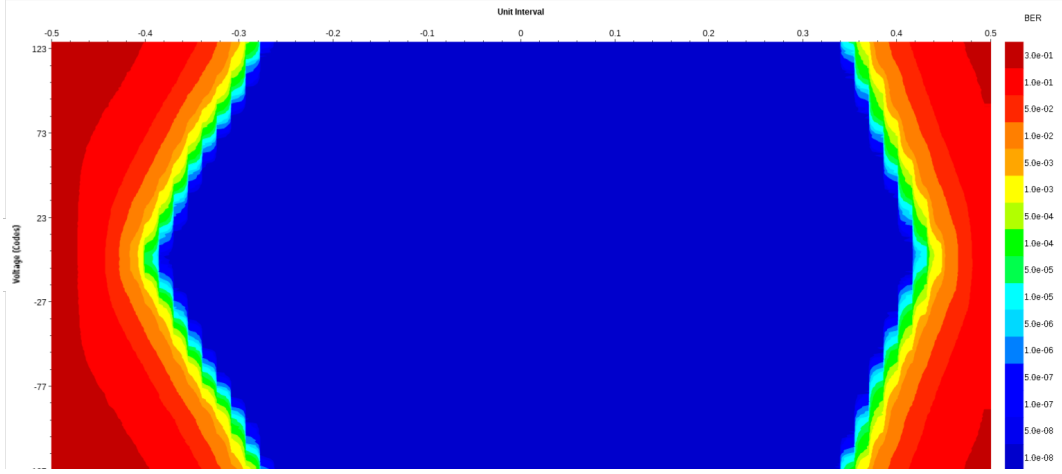

Figure 5.19: Vivado NEPMA eye diagram.

[3]Institute of Electrical and Electronics Engineers

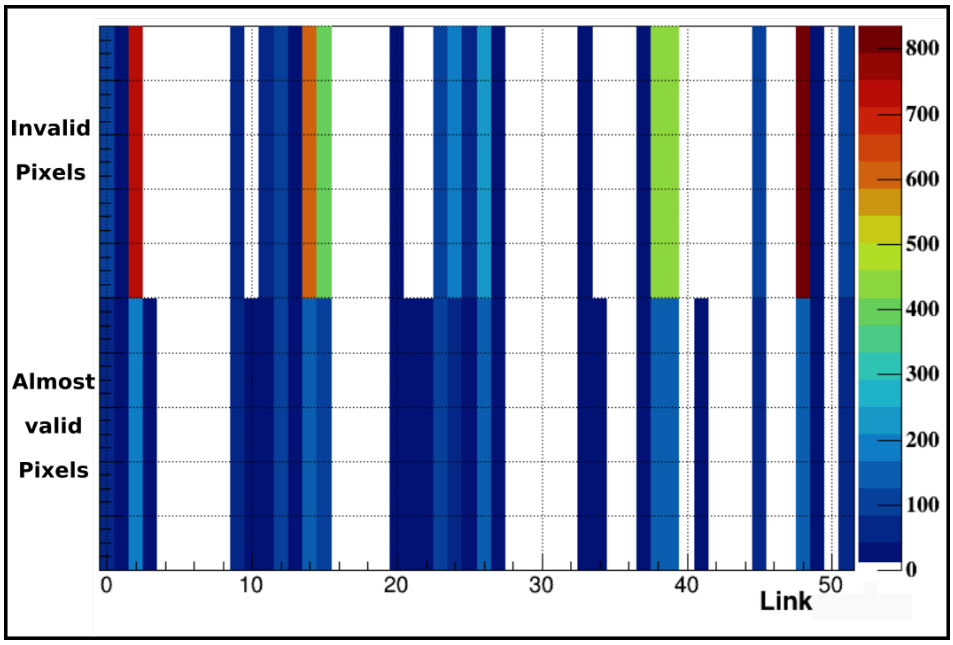

Figure 5.20: Optical loopback summary at 11.2 Gbps. 1500 pixels within the IEEE mask were checked. Link numbers up to 47 are established connections from FireFly transmitter to the Firefly receiver. Links 48 to 51 belong to the duplex.

comparison with the other test modes. Such eyes are desirable for final operation. From here, it can be concluded that the high-speed transmission lines are at least reliable on FPGA level.

Optical Loopback

The results for the optical loopback are not as uniformly as for the NEPMA. Figure 5.20 shows the summary of the analysis software. Although not foreseen to operate at 11.2 Gbps, the duplex is included here. Many links are reported to violate the IEEE mask boundaries because of eye closure. Judging by the results, the recipient side is suspected to distort the signals. Starting from link 2, every twelfth link shows the worst results during the individual component measurements. This is exactly the channel size of one FireFly.
The best and worst eye are further shown in figure 5.21. The deviation to the previous setting is evident. The Open Area is reduced by half even for the good one. The asymmetries those eyes show are also observed on other links. This points to reflections on the PCB connections. Reflections occur if the impedance of a source and a load don't match or if the transmission line is not terminated with its impedance. The routing through Vertical Interconnect Accesses (VIAs) on PCBs may add imaginary parts to the transmission line impedance, which causes distortions [33]. This also holds for the duplex. In its case, the eye openings increase significantly with lower speeds compared to the measurements at 11.2 Gbps, so that the summary histograms report no violation of the IEEE mask boundaries. Also BERs were measured with the outcome of 1.7×10^{-15} at 4.8 Gbps in less than 34 h and 6.1×10^{-15} at 9.6 Gbps in less than 5 h. No errors were detected in both cases, but asymmetries in some eye shapes remained.
The transmitting side is rather excluded as the origin of these distortions, since possible reflections are refused at opto-electrical conversion. To confirm this, a test with off-board transmitters is repeated.

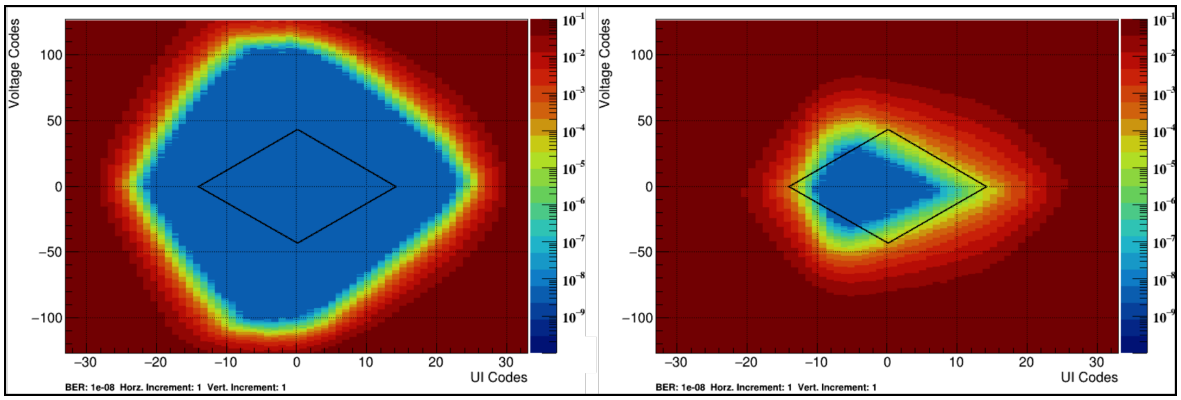

Figure 5.21: Recreated eye diagrams with IEEE mask from optical loopback recorded on FireFly receiver side at 11.2 Gbps. The left picture shows the best (link 6, Open Area ∼5700) , the right the worst eye diagram (link 2, Open Area ∼600).

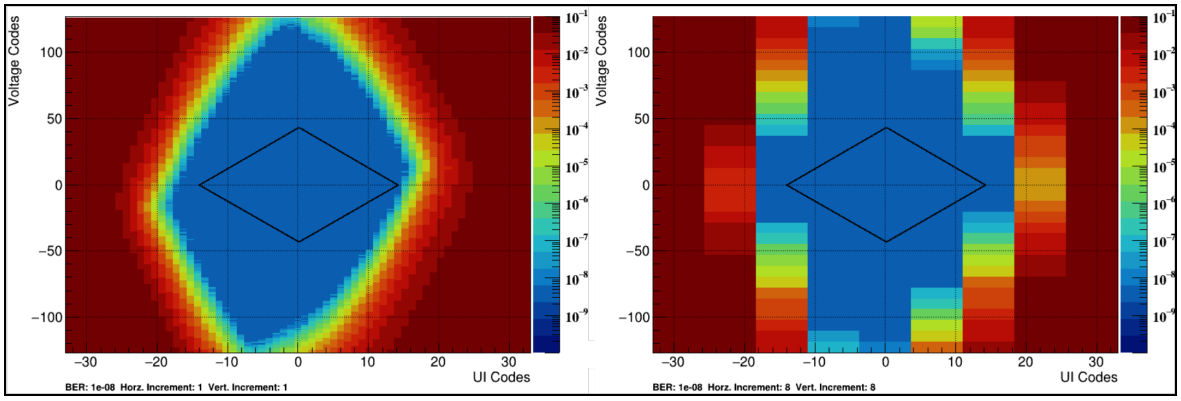

Figure 5.22: Recreated eye diagrams with IEEE mask comparing results from optical loopback and jFEX to TREX test for the second input link of the FireFly duplex at 11.2 Gbps. The left picture was recorded with single increment steps (Open Area ∼4800), the right with increments of 8 (Open Area ∼5800) in both directions.

jFEX to TREX

The coarser resolution of the jFEX to TREX tests makes comparison a bit difficult. Also the analysis software is here only of limited use. Nevertheless, similarities are found in eye openings and shapes where differences should be attributed only to resolution and statistical fluctuations of the eye scans. Figure 5.22 compares the second input link of the duplex, that is link 49 in figure 5.20, in optical loopback and joint test as an example. The position is the same and the shape, taking the resolution into account, too. Reported Open Areas differ by a number of 1000, whereas the increase for the joint test is probably owed, again, to the resolution. Such similarities are observed on many links. This, together with previous and following results, affirms the assumption of reflection. Anyway, the optical reception on TREX is only important for one duplex link, since the FireFly receiver serves as debug device and is not foreseen at final operation. The best duplex link can be selected for this task. Measured Open Areas of 6000 to 9000 at lower speeds and the ability to adjust the equalization allow secure reception of data.

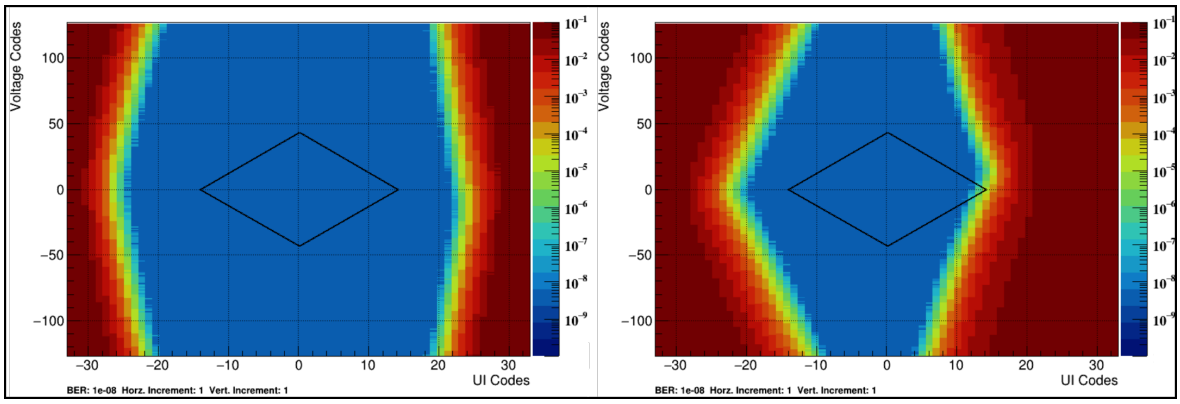

Figure 5.23: Recreated eye diagrams with IEEE mask from TREX to jFEX tests at 11.2 Gbps. The left picture shows the best (link 21, Open Area ∼10500) , the right the worst eye diagram (link 8, Open Area ∼5700).

TREX to jFEX

The TREX to jFEX measurements were performed with two different BC-clock settings. Either both modules were running with independent clocks or a common clock was supplied by the TREX. The results of both clock settings are consistent. The eyes are much more satisfying than those at the reception on TREX. Figure 5.23 shows again the best and the worst eye judged on the basis of the Open Area. Many links achieve NEPMA standards. Also here, a BER of 1.7×10^{-15} was determined over 14 h of operation with no occurrence of errors on all links. However, the analysis software reported violations of the IEEE mask boundaries, as figure 5.24 shows. Unlike for the optical loopback, this happens not because of eye closure but because of the narrowness of these eyes, always extending the eye margins into the mask on the right side like in figure 5.23. Thus, some links appear to be over-equalized. Although the BER tests were successful on all links, this might have an impact on it. A repetition of the BER measurements over a longer period of two to three days is recommended and foreseen to reach BERs of smaller 10^{-15}. Anyhow, since the jFEX also uses Xilinx FPGAs as pro-

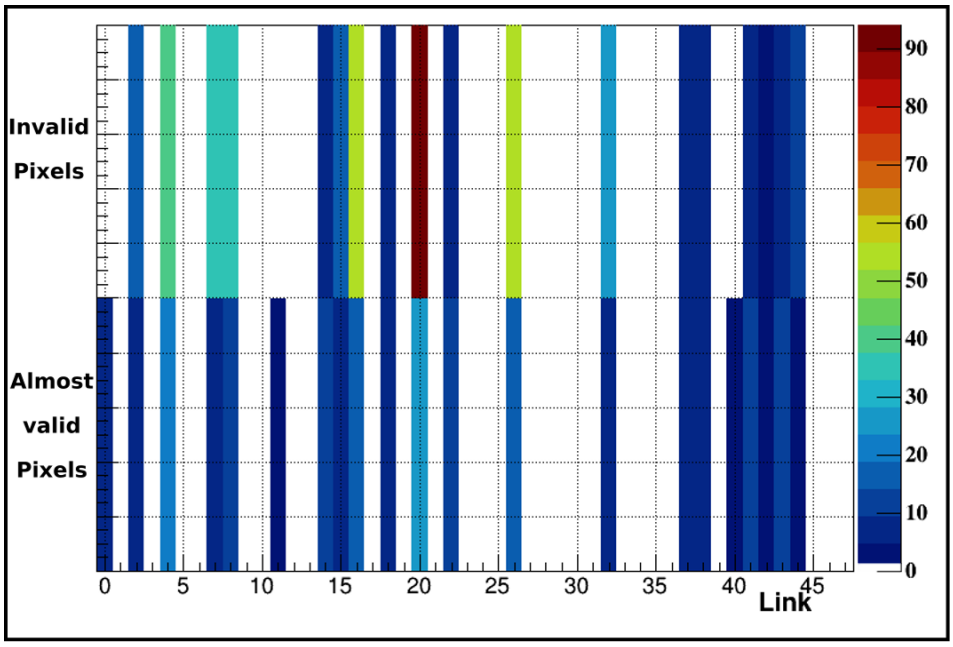

Figure 5.24: TREX to jFEX summary histogram. 1500 pixels within the IEEE mask were checked. Connection with the duplex was not established here.

cessing units [9], the equalization methods presented in section 5.3.2 are also adjustable here. Given this and again the fact that not all of the tested links are needed for final operation, a secure data transmission is guaranteed. In case of the jFEX only one Firefly is reserved for the communication from TREX, limiting the fibre links to twelve.

The results demonstrate that optical high-speed transmission from TREX is operating successfully. Optical reception with the duplex is also safe for the desired line rates. Potential problems that may occur can be overcome with signal enhancements available on Xilinx FPGAs. The provision of optical Tile calorimeter data to the FEXes and a coordinated operation with other trigger modules is thereby feasible.

5.4 Latency Comparison

The single processing steps in L1 trigger functionality are subject to temporal criteria. The front-end buffers of the detector are dimensioned to store data for a maximum of 2.5 μs. Within this time the L1 trigger decision has to be made. The legacy trigger path with the extended Preprocessor system has to maintain this standard. The correct BC-wise assignment in subsequent processors is also time-sensitive. Too long delays may cause the data to be assigned to the next BC. A latency comparison between the new and old system should provide information on these points.
The setup for these measurements is shown in figure 5.25. The latency times are recorded manually with an oscilloscope and the usual configuration of the nMCMs and DINO FPGAs. An analogue signal generator creates the input signals which are used for triggering upon the arrival at the PPM input. A change in LVDS data at the output connectors of the PPM or TREX determines the latency with respect to the input pulse. This is recorded for each physical link including the fan-out individually. Because of the manual cursor adjustment on the oscilloscope, the accuracy of these measurements is limited. They serve only as an estimate to exclude latency issues. A measurement error of 1-2 ns has to be regarded.
TREX is clearly expected to bring an increase of latency simply because it extends signal paths. On the other hand, the replacement of the LCD by a passive component reduces it. The results are presented in figure 5.26 for CP and JEP links separately. The links are denoted by their origin and destination. Comparing the slowest links for the respective destination processors, TREX delays the latency by 1.6 ns on the data path to CP and 1.2 ns on the data path to JEP. The highest increase on a single link is measured to be 2.6 ns for the fan-out of the first channel from the second nMCM (MCM_2_CP1_FOUT). The difference between CP and JEP links is about 25 ns, which is one BC period and agrees with the different processing in the nMCMs. The outcome is further in accordance with single latency test of the passive LCD and processing times in the DINO FPGAs. The passive LCD brings a reduction up to 5 ns, while the routing-delay through the DINO FPGAs is maximally 7.4 ns [11]. The results show that the probability of an incorrect BC assignment in CP or JEP is rather low, since the latency increase is small compared to the BC period. L1 trigger decision are thus unaffected by this increase.

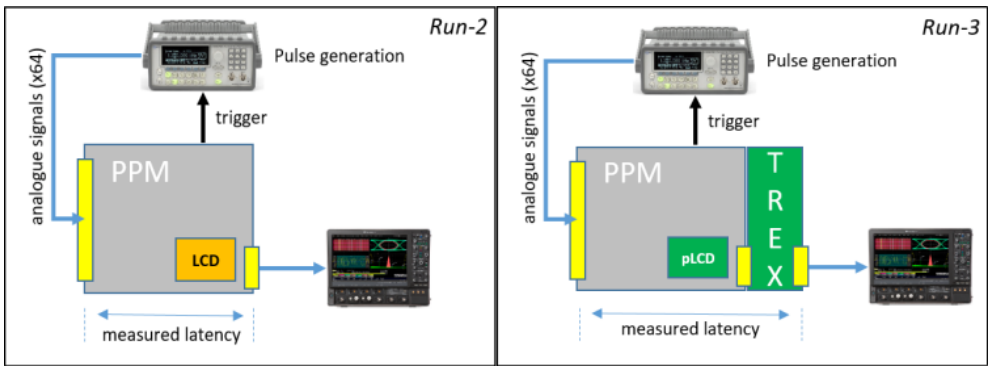

Figure 5.25: Setups for Latency Measurements. The left side shows the configuration for Run-2. The right side is arranged with the upgraded components.

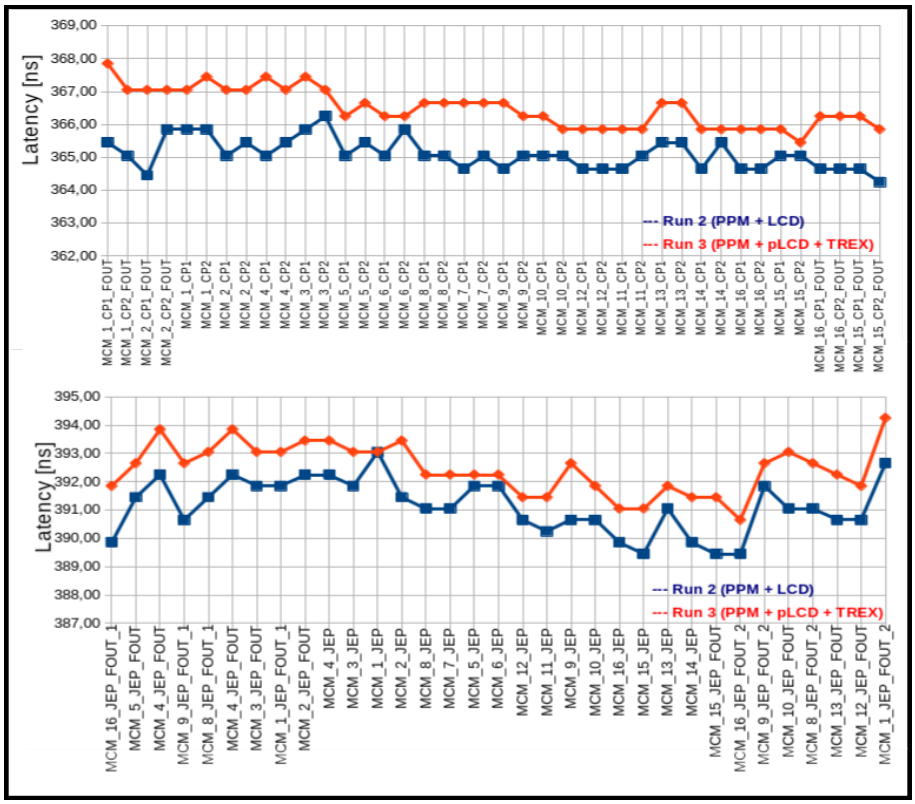

Figure 5.26: Latency measurements Run-2 vs Run-3. The upper chart shows the comparison for CP links, the lower one for JEP links.

Chapter 6

Summary

The ATLAS detector is one of the most powerful detectors at a collider that contributed in the past to essential discoveries and still aims to continue so today. In the course of the LHC upgrades increased requirements are made on the single detector constituents and its trigger system. In order to maintain a high trigger efficiency, the LAr calorimeter revises its front-end electronics to provide the L1Calo FEXes with its data optically and in finer granularity. The TREX completes the existing Preprocessor system to adopt the optical transmission for Tile calorimeter data.

This thesis aimed to assure TREX functionality concerning operating conditions, LVDS reception in PREDATOR FPGA and optical transmission to subsequent L1Calo processors. A latency comparison excluded timing issues on the legacy data path.

The development of operating conditions was monitored in an open environment at room temperature using two different heat sinks. Maximum temperatures of 72 °C and 55 °C were measured respectively, where the second heat sink was able to lower the temperature because of a thinner thermal conductive foil. This also resulted in a reduction of the board currents and thus of the power consumption. A TREX board with a mounted heat sink of the second type consumes 43 W at full load of the PREDATOR FPGA, whereby the use of Gigabit transceivers in final operation is reduced compared to the presented tests. This is half of the remaining power budget for a single board when using the same power supply unit as for the PPM. Similar tests with an adequate air-cooling system will be repeated when sufficient TREX boards are available to equip an entire crate. Long-term monitoring of TREX at CERN STF revealed safe running.

To validate LVDS reception in the PREDATOR FPGA, a delay element was used that is able to shift the sampling point by more then a bit length. Stable delay values were found on all links by sending known data from the nMCMs of the PPM and comparing it on reception deploying the entire receiving logic. The alignment of BC-wise assignments was shown to be preserved. Whether the determination of a safe sampling point will be implemented in the final firmware design or taken over by software has not yet been decided.

The optical high-speed transmission was evaluated under different test settings on the basis of IEEE standards. Optical loopback and joint tests showed that distortions occur at reception, that can be overcome for the FireFly duplex at slower line rates. Results are not alarming in the case of the receiver since it will be discarded for final operation. The transmission to the jFEX was successful on all links, resulting in BERs of 1.7×10^{-15} on all connections. Although some links violated IEEE demands, no errors have appeared during these tests. Longer measurements will indicate whether over-equalization is problematic for the BER. Fine-tuning of equalization settings offer improvements in this case. For TREX, it can be concluded that optical transmission works as required, since eye diagrams in the joint tests were wide open.

All three tested prototype modules of the first TREX version showed an good overall performance. Results were consistent within the manufacturing tolerances. The next version is currently in production and will include the Zynq+ MPSoC device for a slow control interface to the ATLAS DCS.
Once the L1Calo upgrades are validated, CP and JEP will be decommissioned offering possibilities to modify nMCM algorithms and PPM output data. Single trigger tower and jet sum energies could be delivered in finer resolution by enlarging the ratio of bits per value. To handle the increased data volume, the real-time output of the PPM can be configured to clock at 80 MHz, twice the BC frequency. Whether this becomes necessary will be decided when the physics requirements of all FEXes are known.

Appendix

A Warning and Error Alarms

Temperatures [°C]

Device	Lower Warning	Upper Warning	Lower Error	Upper Error
PREDATOR FPGA	-	80	-	85
DINO FPGAs	-	60	-	75
Samtec FireFlys	-	60	-	70
Power Managers LTC2977	-	60	-	70

Supply Voltages [V]

Device	Nominal Value	Lower Warning	Upper Warning	Lower Error	Upper Error
PREDATOR FPGA	0.95	0.93	0.97	0.922	0.979
	1.2	1.18	1.22	1.17	1.23
	1.8	1.765	1.835	1.746	1.854
	3.3	3.234	3.366	3.2	3.4
DINO FPGAs	1.0	0.97	1.03	0.95	1.05
	1.8	1.746	1.854	1.71	1.89
	2.5	2.425	2.575	2.375	2.623
	3.3	3.234	3.366	3.2	3.4
Samtec FireFlys	3.3	3.234	3.366	3.2	3.4
Power Managers LTC2977	3.3	3.234	3.366	3.2	3.4
Board voltages	3.3	3.234	3.366	3.2	3.4
	5.0	4.9	5.1	4.85	5.15

B Lists

List of Figures

List of Tables

List of Acronyms

ADC	Analogue-to-Digital Converter
AFE	Analogue-Front-End
AFP	ATLAS Forward Proton detector
AnIn	Analogue Input
ASIC	Application-Specific Integrated Circuit
ATLAS	A Toroidal LHC ApparatuS
BC	Bunch-Crossing
BCID	Bunch-Crossing Identification
BCMux	Bunch-Crossing Multiplexing
BER	Bit Error Rate
BISC	Built-In-Self-Calibration
BSM	Beyond the Standard Model
CALIPPR	Calorimeter Information Preprocessor
CAN	Controller Area Network
CDR	Clock Data Recovery
CMS	Compact Muon Solenoid
CP	Cluster Processor
CPLD	Complex Programmable Logic Device
CSC	Cathode Strip Chambers
DAC	Digital-to-Analogue Converter
DCS	Detector Control System
DINO	Data-In-Out
ECAl	Electromagnetic Calorimeter
eFEX	Electron FEX
EMB	Electromagnetic Barrel

EMEC Electromagnetic End-cap Calorimeter

FCAL Forward Calorimeter

FELIX Front-End Link Exchange

FEX Feature Extractor

FIFO First-In-First-Out

FIR Finite Impulse Response

FOX Fibre Optics Exchange

FPGA Field Programmable Gate Array

FTK Fast Tracker

gFEX Global FEX

HCAL Hadronic Calorimeter

HDL Hardware Description Language

HEC Hadronic End-cap Calorimeter

HL-LHC High-Luminosity LHC

HLT High Level Trigger

IA Interaction Point

IBERT Integrated Bit Error Ratio Tester

ID Inner Detector

IEEE Institute of Electrical and Electronics Engineers

IP Intellectual Property

JEP Jet/Energy Processor

jFEX Jet FEX

L1 Level-1

L1A Level-1 Accept

L1Calo Level-1 Calorimeter Trigger

LAr Liguid Argon

LCD LVDS Cable Driver

LHC Large Hadron Collider

LS Long Shutdown

LUT Look-Up Table

LVDS Low Voltage Differential Signaling

MDT	Monitored Drift Tube
Micromega	Micro Mesh Gaseous Structure
MPSoC	Multi-Processor System-on-Chip
NEPMA	Near-End PMA
NSW	New Small Wheel
PCB	Printed-Circuit Board
PCS	Physical Coding Sublayer
PISO	Parallel-In-Serial-Out
PLL	Phase-Locked-Loop
PMA	Physical Medium Attachment
PMT	Photomultiplier
PPM	Preprocessor Module
PREDATOR	Preprocessor Data Collector
QCD	Quantum Chromodynamics
ReM	Readout Manager
RPC	Resistive Plate Chamber
SCT	SemiConductor Tracker
SIPO	Serial-In-Parallel-Out
SM	Standard Model
STF	Surface Test Facility
TDAQ	Trigger & Data Acquisition
TCL	Tool Command Language
TCM	Timing Control Module
TGC	Thin Gap Chamber
TOB	Trigger Object
TREX	Tile Rear Extension
TRT	Transition Radiation Tracker
TTC	Timing, Trigger & Control
UI	Unit Interval
VIA	Vertical Interconnect Access
VME	Versa Module Eurocard

Bibliography

[1] ATLAS Collaboration. Webpage. http://atlas.cern/ Accessed: 2018-08-23.

[2] ATLAS Collaboration. *Operation and performance of the ATLAS semiconductor tracker*, arXiv:1404.7473, 2014.

[3] ATLAS Collaboration. *Performance of the ATLAS Trigger System in 2015*, arXiv:1611.09661v2, 2015.

[4] ATLAS Collaboration. *Letter of Intent Phase-I Upgrade*, CERN-LHCC-2011-012, 2012.

[5] ATLAS Collaboration. *ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report*, CERN-LHCC-2013-017, 2013.

[6] ATLAS Collaboration. *ATLAS TDAQ System Phase-I Upgrade Technical Design Report*, CERN-LHCC-2013-018, 2013.

[7] ATLAS Collaboration. *ATLAS Upgrade Status Report 2018–2019*, CERN-RRB-2019-024, 2019.

[8] ATLAS Collaboration. *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST 3 2008 S08003.

[9] M. Weirich et al. *Development of the jet Feature EXtractor (jFEX) for the ATLAS Level 1 Calorimeter Trigger Upgrade at the LHC*, DOI: 10.22323/1.313.0133, 2017.

[10] R. Achenbach et al. *The ATLAS Level-1 calorimeter trigger*, JINST 3 2008 P03001.

[11] V. Andrei et al. *Tile Rear Extension (TREX) ATLAS TDAQ Phase-I Upgrade Final Design Report*, 2019.

[12] V. Andrei et al. *Tile Rear Extension Module for the Phase-I Upgrade of then ATLAS L1Calo PreProcessor System*, JINST 12 C03034, 2017.

[13] B. Jiao from Xilinx Inc. *Leveraging UltraScale Architecture Transceivers for High-Speed Serial I/O Connectivity*, WP458 (v2.0), 2015.

[14] H. Fu from Xilinx Inc. *Equalization for High-Speed Serial Interfaces in Xilinx 7 Series FPGA Transceivers*, WP419 (v1.0), 2012.

[15] S. Hunklinger. *Festkörperphysik*. Berlin, Boston: Oldenburgverlag, 2014.

[16] F. Hügging. *The ATLAS Pixel Insertable B-layer (IBL)*, arXiv:1012.2742, 2011.

[17] Xilinx Inc. *High Speed Serial*. https://www.xilinx.com/products/technology/high-speed-serial.html, Accessed: 2019-06-27.

[18] Xilinx Inc. *Xilinx Community Forum: UltraScale IODELAY - What is the clock alignment delay*, AR# 64743, 2016. Accessed: 2019-06-09.

[19] Xilinx Inc. *Kintex UltraScale FPGAs Data Sheet:DC and AC Switching Characteristics*, DS892 (v1.18), 2019.

[20] Xilinx Inc. *IBERT for UltraScale GTH Transceivers v1.3 - LogiCORE IP Product Guide*, PG173, 2017.

[21] Xilinx Inc. *7 Series FPGAs and Zynq-7000 SoC XADC Dual 12-Bit 1 MSPS Analog-to-Digital Converter - User Guide*, UG480 (v1.10.1), 2018.

[22] Xilinx Inc. *UltraScale Architecture SelectIO Resources - User Guide*, UG571 (v1.10), 2019.

[23] Xilinx Inc. *UltraScale Architecture GTH Transceivers - User Guide*, UG576 (v1.5.1), 2018.

[24] Xilinx Inc. *UltraScale Architecture System Monitor - User Guide*, UG580 (v1.9.1), 2019.

[25] Maxim Integrated. *MAX11612–MAX11617, Low-Power, 4-/8-/12-Channel, I2C, 12-Bit ADCs in Ultra-Small Packages*, 2012. https://datasheets.maximintegrated.com/en/ds/MAX11612-MAX11617.pdf.

[26] J. Jongmanns. *The Upgrade of the PreProcessor of the ATLAS Level-1 Calorimeter Trigger for LHC Run-2*. PhD thesis, Heidelberg University, 2017.

[27] Jan Jongmanns. *PPM & TREX Firmware Designs - TREX FDR*, 2018. https://indico.cern.ch/event/757101/.

[28] Samtec. *FireFly Micro Flyover System*. https://elcon-electronic.de/fileadmin/user_upload/firefly-brochure.pdf, Accessed: 27.05.2019.

[29] Samtec. *FireFly Optical 14G x12 Data Sheet*, 2015.

[30] Samtec. *FireFly Optical ECUO 14G x4 Data Sheet*, 2017.

[31] IEEE Computer Society. *IEEE Standard for Ethernet Amendment 3: Physical Layer Specifications and Management Parameters for 40 Gb/s and 100 Gb/s Operation over Fiber Optic Cables*, IEEE Std 802.3bm, 2015.

[32] Linear Technologie. *LTC2977 - 8-Channel PMBus Power System Manager Featuring Accurate Output Voltage Measurement*, 2013. https://www.analog.com/media/en/technical-documentation/data-sheets/2977fc.pdf.

[33] Keysight Technologies. *Demystifying Vias in High-Speed PCB Design*. https://www.keysight.com/upload/cmc_upload/All/2_Demystifying_Vias_in_High-Speed_PCB_Designs.pdf, Accesed: 2017-06-20.

Erklärung:

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den (Datum) .....