Conceptual design and testing of new trigger hardware for the ATLAS experiment

Dissertation
zur Erlangung des Grades

Doktor der Naturwissenschaften

am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg - Universität
in Mainz

Stefan Rave

geboren in Worms

Mainz, den 28. Juli 2025

Erster Berichterstatter: Prof. Dr. Stefan Tapprogge

Zweiter Berichterstatter: Prof. Dr. Wolfgang Gradl

Datum der mündlichen Prüfung: 26. November 2025

Kurzfassung

Mit dem zunehmenden Verständnis des Standardmodells der Teilchenphysik sowie der darin beschriebenen Teilchen und Wechselwirkungen steigen die Anforderungen an Schwerpunktsenergie und Luminosität, um die Forschung auf diesem Gebiet weiter voranzutreiben. Für den Betrieb des ATLAS Experiments in Run 3 und Run 4 wird daher neue Ausleseelektronik entwickelt, die den gestiegenen Anforderungen hinsichtlich Eingangsbandbreite und Rechenkapazität gerecht wird. Sie ermöglicht die Nutzung einer feineren Granularität der Kalorimeterdaten, die erforderlich ist, um einzelne Objekte im Detektor zuverlässig voneinander zu unterscheiden.

Im Rahmen dieser Arbeit wird das Design des Jet Feature Extractors (jFEX), der hauptsächlich zur Identifikation von Jets und zur Berechnung von Energiesummen dient, entwickelt und getestet. Dieses Design basiert auf der neuesten Generation besonders leistungsfähiger FPGAs und nutzt deren hohe Bandbreite sowie Verarbeitungskapazität, um die strengen Anforderungen des ATLAS-Trigger-Systems zu erfüllen.

Die Verifikation konzentriert sich vor allem auf die dicht gerouteten Hochgeschwindigkeitsverbindungen, die für Übertragungsraten von bis zu $12,8$ Gb/s ausgelegt sind. Zudem erfolgt eine erste Abschätzung der mit dem Design realisierbaren neuen Jet-Trigger-Algorithmen sowie ihrer Leistungsfähigkeit anhand simulierter Ereignisse.

Abstract

As our understanding of the Standard Model of particle physics and its described particles and interactions continues to advance, higher centre-of-mass energies and luminosities are required to further drive the research in this field. For the operation of the ATLAS experiment in Run 3 and Run 4, new trigger hardware is designed that meets the increased requirements in terms of input bandwidth and processing capabilities. This enables the use of higher-granularity calorimeter data required to reliably distinguish between separate objects within the detector.

This thesis presents the design and verification of the jet Feature Extractor (jFEX), mainly intended for the identification of jets and the calculation of energy sums. This design is based on the latest generation of high-end FPGAs, utilising their available bandwidth and processing capabilities to fulfil the strict boundary conditions of the ATLAS trigger systems.

The verification of the design focuses on the densely routed high-speed links, intended for transmission rates of up to 12.8 Gb/s, and an initial estimation of the provided possibilities for new jet trigger algorithms and their performance based on simulated events.

Contents

Chapter 1

Introduction

While the Standard Model of particle physics is highly successful in describing various particles and their interactions, it is also considered incomplete as it cannot be used to calculate gravitational effects and does not provide answers to currently open questions like the existence of dark matter and dark energy. Additionally, parameters of particles and interactions described by the Standard Model, like properties of the only relatively recently discovered Higgs boson, are further studied to test the known model against high precision measurements.

The Large Hadron Collider (LHC) is regularly upgraded as both, detailed measurements of rare particles and the discovery of new physics processes, can benefit from an increase in instantaneous luminosity and centre-of-mass energy. This also increases the demand towards the surrounding experiments which must cope with higher numbers of interactions per bunch crossing which leads to more detectable signals. In order to maintain a high sensitivity for the physics processes that are intended to be studied, dedicated systems are required that control the data acquisition. These trigger systems are designed to be operated under pre-defined conditions. Therefore, in the context of an upgrade of the LHC, the surrounding experiments, like the ATLAS experiment, also require an update of their trigger systems.

The first stage of the ATLAS trigger system uses data from the calorimeters and parts of the muon system to provide a rough estimate of the underlying physics event. As a part of this upgrade, a new subsystem called jet Feature Extractor (jFEX), mainly intended for the identification of jets and the calculation of energy sums, is introduced. This thesis describes design process and verification of the design with respect to its suitability to be used as a part of the ATLAS trigger system.

This thesis is structured as follows: Part I addresses the theoretical background in chapter 2. Part II provides an overview over the underlying experimental setup consisting of the Large Hadron Collider with its affiliated experiments as shown in chapter 3, followed by a more detailed description of the ATLAS experiment in chapter 4 and an abstract of the corresponding first two upgrade stages in chapter 5. Part III discusses the development of the jFEX starting with a detailed description of the requirements in chapter 6. A short introduction to the principles of data processing in hardware is given in chapter 7 to provide the foundation of the jFEX design as discussed in chapter 8 which is then verified in chapter 9. Finally, part IV provides a short summary of the discussed results and further outlook for the jFEX as a part of the ATLAS trigger system.

Part I

Theory

Chapter 2

Theory Foundations

This chapter provides an overview over the theoretical background as the foundation of this thesis. A brief description of the Standard Model of particle physics including its particles, fundamental forces, and limitations is given in section 2.1 followed by a general discussion of possibilities to detect these particles and interactions in section 2.2. The concept of jet algorithms is presented in section 2.3 with a focus on the anti-k_t algorithm. Finally, the physics motivation for the use of a trigger system in an experimental setup like the ATLAS experiment is given in section 2.4.

Unless stated otherwise, the following descriptions are based on 'Elementare Teilchen' by Jörn Bleck-Neuhaus [1] and Marc Thomson's 'Modern Particle Physics' [2].

Throughout this thesis, the convention $\hbar = c = 1$ is used, which leads to notations in natural units using traditional units of energy, usually electron volts (eV), for the denotation of masses, momenta, and other physical units.

2.1 The Standard Model

The Standard Model of particle physics (SM) combines three of the four known fundamental forces in the universe in a single model. It describes the electromagnetic, weak, and strong interaction in the form of a Quantum Field Theory via an exchange of mediating particles. Additionally, the Standard Model classifies the known fundamental particles. Only gravitation cannot be described using the Standard Model. As fundamental particles and their composites have a relatively low mass, gravitational effects can typically be neglected compared to the other forces in calculations for physics processes described by the SM.

At sufficient energies starting around the GeV level, the world of atoms und nuclei, consisting of electrons (e), electron neutrinos (ν_e), protons, and neutrons, responsible for most physics phenomena observed in the daily life, can be breached to reveal an underlying substructure. While the electron to this day is considered a fundamental particle that cannot be split into constituents, proton and neutron are discovered to be bound states of new particles called quarks. The proton consists of two up quarks and one down quark, and the neutron consists of one up quark and two down quarks. These two quark types together with the two mentioned leptons (electron and electron neutrino) are categorised as the first generation of fermions by

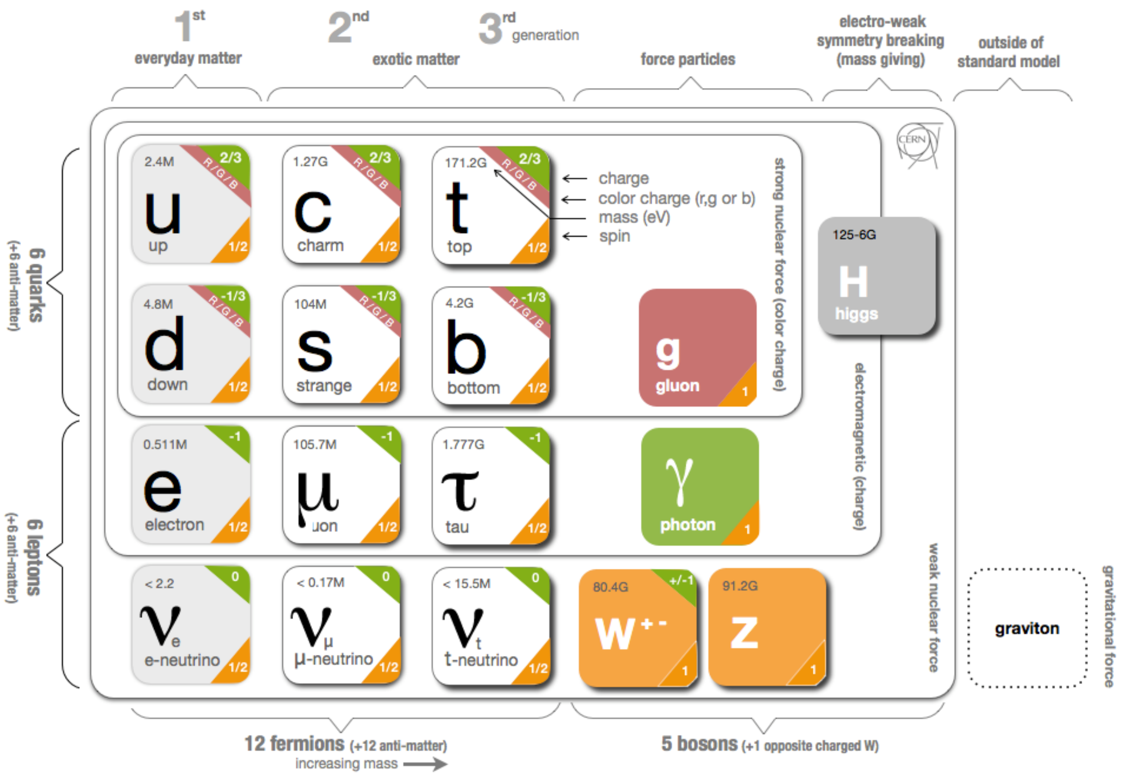

FIGURE 2.1: Fundamental particles of the Standard Model subdivided into fermions (left) and bosons (right). For each particle the mass, charge (in units of elemental charges), potential for colour charges, and spin is given. The gravitational force with a hypothetical graviton is not described by the Standard Model. Figure taken from [3].

the Standard Model as shown in figure 2.1.

As shown in figure 2.1, the common characteristics of all fermions is a half integer spin opposed to bosons with an integer spin. The group of bosons consists of the four gauge bosons, which mediate the fundamental forces between the fermions, and the Higgs boson, which was predicted as a vital part of the Standard Model. The Higgs boson was discovered in 2012 by the ATLAS [4] and CMS [5] experiments at the LHC[1]. This was a valuable confirmation of the concepts behind the Standard Model of particle physics. In our current understanding, it is the interaction between the Higgs field and particles which attributes them their particular masses and provides them many of their observed properties.

Each fermion has an anti-particle with the opposite charge and otherwise identical properties. Anti-particles are typically denoted with a bar over the sign. For example, the anti-particle of the up quark (u) is written as $\bar{u}$. For charged leptons, the notation of the opposite charge can also be used to indicate anti-particles. The anti-electron, or positron, can therefore also be written as e^+ compared to e^- for the electron. Within this thesis, the differentiation between particles and anti-particles is not required in most cases as within this context, their properties, except for the opposite charge, are believed to be identical. Therefore, both forms will be

---

[1] See chapter 3 for more details on the experimental setup.

referred to as the particle, for example as an electron (e), unless explicitly differentiated.
While only the first generation of fermions can be found as stable matter in the universe, high-energy conditions caused e.g. by particle colliders are required[2] to produce the second and third generations for fermions. Each fermion is distinguished from its counterparts in other generations only by its mass. For the electron (e), the following generations are called μ and τ with their corresponding neutrinos ν_μ and ν_τ. In the same manner, the up quark (u) is followed by charme (c) and top (t). These three quark types are therefore known as up-types while the down-types consist of down (d), strange (s), and bottom (b). These 12 fermion types are also known as the flavour states.
The Standard Model describes all fundamental forces via an exchange of mediating particles. Within interactions, particles can occur as virtual particles, allowing them to carry a mass different from their regular properties. As shown in figure 2.1, the W and Z bosons, acting as mediators of the weak nuclear force, couple to the weak isospin which is carried by all known fermions. The photon, as described in Quantum Electrodynamics, only interacts with particles that carry an electric charge, namely the charged leptons and quarks. The strong nuclear force, mediated by gluons as described in Quantum Chromodynamics (QCD), couples to colour charges which can only be found in quarks and gluons. Each quark carries a single colour charge known as red, green, or blue (r, g, b), with corresponding anti-colour for antiquarks ($\bar{r}$, $\bar{g}$, $\bar{b}$). The colour-charge of a gluon consists of a colour and an anti-colour in one of the 8 possible states of coloured combinations according to group theory[3]. The fact that gluons themselves carry a colour charge leads to the possibility of gluon-gluon coupling.

2.1.1 Fundamental Interactions

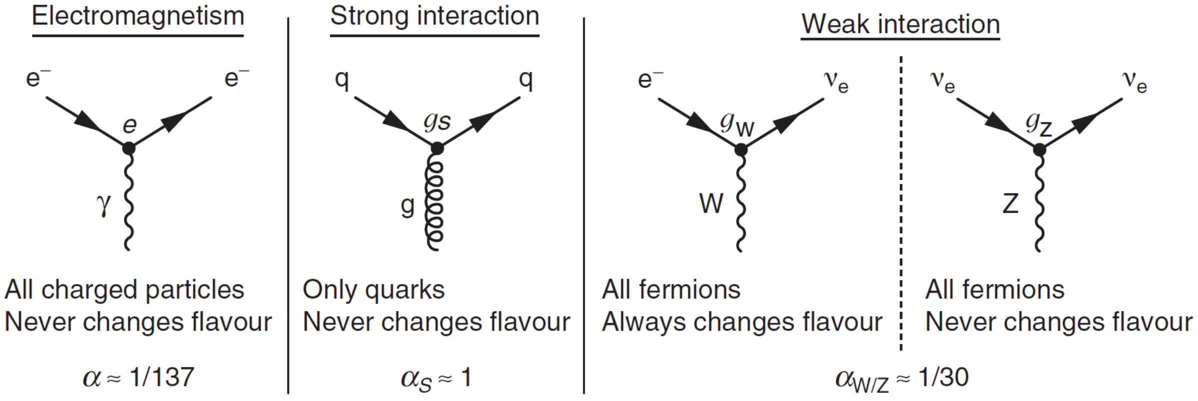

Figure 2.2: Overview of interactions described by the Standard Model of particle physics. The graphical representation for each interaction is shown together with their corresponding coupling constant α. Figure taken from [2].

Interactions between fermions via the fundamental forces are represented by vertices, as shown in figure 2.2, indicating the coupling of the mediating bosons with the fermions. For each interaction type, the coupling constant is noted, expressed as a dimensionless factor (α), which gives the relative probability for processes with n couplings to bosons as $\propto \alpha^n$.

---

[2] Except for neutrino oscillation which is not further discussed within this thesis.

[3] The remaining combination is colour-neutral and therefore not an allowed state for a gluon.

The charged-current of the weak interaction mediated by W bosons is the only interaction that does not only transport energy and momentum but can also couple to different flavour states playing an important role in particle decays. This allows for the transition between a charged lepton and its corresponding neutrino or between quarks with a charge difference of 1 elemental charge.

As transitions between quark generations are possible, the eigenstate linked to a u quark is not written as $|d>$ but as $|d'>$ consisting of a linear combination of $|d>$, $|s>$, and $|b>$ eigenstates. The Cabibbo–Kobayashi–Maskawa (CKM) matrix links the mass eigenstates to the weak eigenstates as

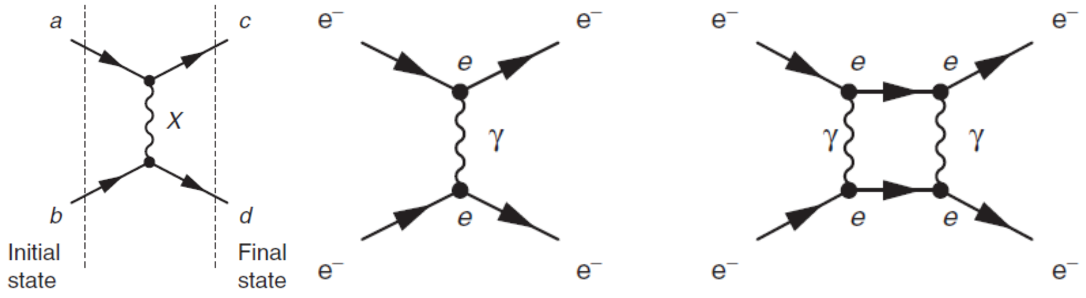

(A) (B)

Figure 2.3: Basic form of a Feynman diagram showing initial (a, b) and final (c, d) state connected via the occurring interaction (A) and an ee scattering process (B). Figures taken from [2].

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{2.1}$$

The transition probability between quark q_i and q_j is given via the corresponding CKM element as $|V_{ij}|^2$. Within the Standard Model, the values of individual CKM matrix elements can only be experimentally determined. The measurements can be interpreted as

$$\begin{pmatrix} |V_{ud}| & |V_{us}| & |V_{ub}| \\ |V_{cd}| & |V_{cs}| & |V_{cb}| \\ |V_{td}| & |V_{ts}| & |V_{tb}| \end{pmatrix} \approx \begin{pmatrix} 0.974 & 0.225 & 0.004 \\ 0.225 & 0.973 & 0.041 \\ 0.009 & 0.040 & 0.999 \end{pmatrix} \tag{2.2}$$

showing a strong correlation between quarks of the same generation. For the t quark, this leads to a probability of about 99.8% for the transition to a b quark.

Feynman diagrams can be used to visualise the interactions between fermions. The basic principle of such a graph is shown in figure 2.3a for a process of unspecified particles $a + b \to c + d$ via the exchange of a boson X with the time propagating from the initial state on the left to the final state on the right. This example shows only the most basic interaction. The two graphs in figure 2.3b, showing the same electron-electron scattering process, represent the circumstance that for every physics process an infinite number of Feynman diagrams can be drawn with increasing number of involved boson-fermion couplings. As each additional coupling reduces the relative probability, only the lowest orders of Feynman diagrams need to be considered according to perturbation theory.

2.1.2 Hadronisation and Particle Decays

An important observation regarding quarks is that no free quarks as fractionally charged particles have ever been discovered in experiments. This behaviour is explained by the hypothesis of colour confinement stating that no objects with non-zero colour charge can exist as free particles. In this concept, the separation of quarks leads to a field of the strong nuclear force between the 'free' quarks that increases in energy as the quarks are separated further. The creation of such a field is visualised in figure 2.4a. This effect leads to the confinement of quarks in colourless hadrons[4]. Hadrons can be observed with a particular mass obtained from their constituent quarks and to a significant extend from the QCD binding energy required to form this bound state. The only quark with no known hadronic states is the t quark as it decays faster than the hadronisation process requires to take place.

The quarks as constituent parts of a hadron, as discussed so far, are also known as the valence quarks. The QCD binding of these quarks additionally produces a sea of virtual gluons that constantly produces quark-antiquark pairs called sea quarks which after a short time are recombined and eliminated. At sufficient energies, this substructure of hadrons can be measured via so-called deep inelastic scattering processes.

The main types of hadrons are mesons and baryons. While mesons consist of a $q\bar{q}$ pair, baryons are made up of a triplet of quarks (or antiquarks), like proton and neutron. Other combinations of four or more quarks are known as exotic hadrons. The existence of these objects has been postulated for a long time but only in recent years, observations of these states were made that are widely accepted in the scientific community [6].

If quarks are moving away from each other with sufficient momentum, for example a quark-antiquark pair traveling back-to-back in their centre-of-mass frame after creation in a particle collider experiment, the energy stored in the field between them becomes sufficient to create new $q\bar{q}$ pairs as shown in figure 2.4b. This continues until the remaining quarks have sufficiently low energy to be combined into colourless hadrons as required by the confinement. The process of creating (multiple) new quarks and antiquarks is known as hadronisation[5]. The initial quark and antiquark each create a hadronic shower following their original direction.

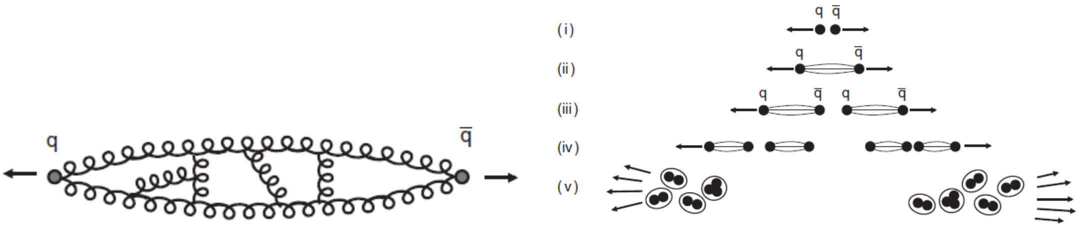

(A) Visualisation of the strong interaction field formed by separated 'free' quarks.

(B) Hadronisation process in multiple steps where an initial $q\bar{q}$ pair produces showers of hadrons.

FIGURE 2.4: Figures taken from [2].

Many of the discussed fundamental particles are not stable but can decay into a final state with a lower rest mass. These processes all involve the weak charged current as required to change the flavour of a fundamental fermion. Figure 2.5a shows the Feynman diagram of the decay

---

[4] A hadron is any combination of quarks held together via the strong interaction.

[5] Hadronisation can also be caused by high-energy gluons.

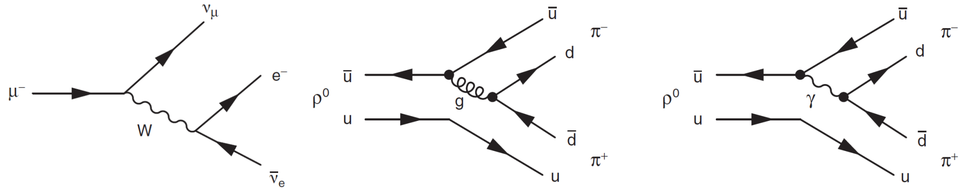

(A) (B)

FIGURE 2.5: Feynman diagrams of the muon decay (A) and $\rho^0(u\bar{u})$ decay channels based on the strong and electromagnetic interaction (B). Figures taken from [2].

process $\mu^- \to e^- \bar{\nu}_e \nu_\mu$. Please note the convention to draw anti-particles with an inverted arrow. Hadrons, consisting of bound quarks, can decay by splitting into their constituent quarks without changing their flavours while releasing its QCD binding energy. These processes do not involve an exchange of W bosons as shown in figure 2.5b. The lifetime of a hadron depends, among other factors, on the strength of the involved interactions[6] and the difference in the resting mass to its final state. The proton is the only hadron assumed to be stable as a free particle while the neutron can become stable if bound in a stable nucleus.

2.1.3 Proton-Proton Collisions

Within this thesis, the most relevant special case of particle collisions are collisions between hadrons where their substructures lead to additional effects. As the LHC mainly uses proton-proton collisions and the system described in this thesis is mainly designed for the conditions created by this type of collision, the following descriptions are focused on *pp*-collisions.

At high energies, the deep inelastic scattering process dominates where sub-particles called partons, namely (anti-)quarks and gluons, are involved in the interactions instead of the hadron reacting as a single object. The high energies remove constituents from the hadron, breaking the bound state and releasing additional particles. As shown in figure 2.6, for each proton, only a single parton is directly involved in the collision process. This parton carries an unknown and largely random momentum fraction (x) of the total momentum of the hadron. As the momentum fractions for the two involved partons are independent, the centre-of-mass frame of their interaction is generally not identical with the frame of reference of the surrounding detector, leading to a so-called boost for the created particles.

The distribution of the momentum fraction for the individual quark types is given by parton distribution functions (PDF). As calculations of PDFs based on perturbation theory are not possible due to the high coupling constant of the strong interaction (α_S), they can only be experimentally determined for example by measuring deep inelastic scattering processes. The current understanding of PDFs of the proton are shown in figure 2.7 with leading contributions for high x coming from the valence quarks and dominating factors from gluons (xg) at low x. As pair production of s quarks is possible with the energy available in the QCD binding of the proton, a small contribution from xs can be seen.

These properties of *pp*-collisions lead to advantage for the use in particle collider experiments

---

[6] Hadrons with available decay processes involving the strong interaction typically have the shortest lifetime of $< 10^{-20}$ s and hadrons only decaying via the weak interaction have the longest lifetime from $\sim 10^{-12}$ s to about 15 minutes for free neutrons.

as well as affiliated challenges. Due to the variable momentum fraction involved in the collision, the exact conditions vary for each event. This effect can be utilised in the search for unknown particles by automatically varying the energy available for the creation of particles while making it more difficult to precisely measure known particles that could be more reliably created with stable collision conditions.

The broken bound state of a proton after the interaction of a parton in a deep inelastic scattering process leads to the creation of additional particles that can cause so-called background signals in the detector. Though, as these particles carry the remaining momentum of the proton, they are typically highly boosted and traveling almost in the same direction as the remaining beam which allows them to leave the detector without noticeable interactions. Most particle collider experiments also aim at causing multiple hadron collisions at the same time to increase the yield of rare physics processes occurring. This leads to the main source of background signals in the detector from simultaneous hadron interactions.

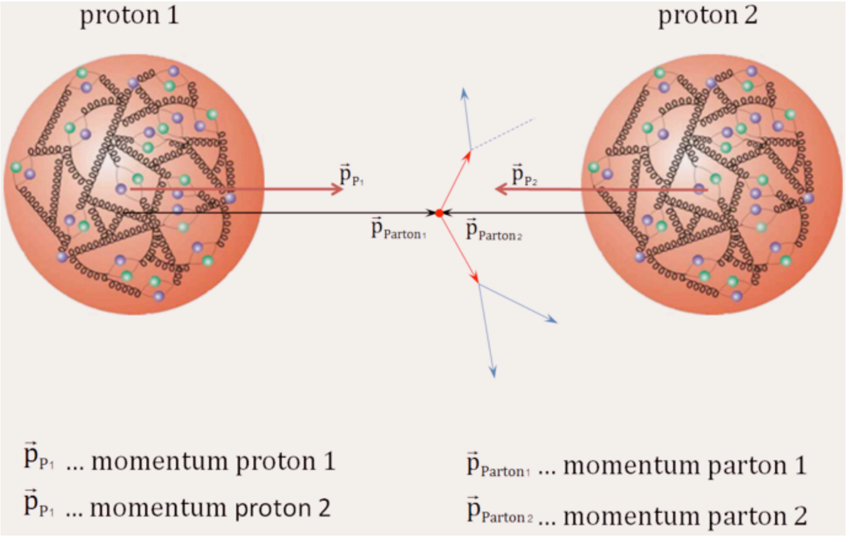

FIGURE 2.6: Representation of a deep inelastic scattering process between protons showing the interaction of partons carrying a fraction of the overall momentum. Figure taken from [7].

2.1.4 Limitations of the Standard Model

The Standard Model of particle physics has been extremely successful as a quantitative description of the observations made in experiments over the last decades as well as predicting the existence of particles like the top quark and the Higgs boson. However, the Standard Model is not able to describe all physics phenomena as for example it does not include gravitational effects. Some of the further limitations of this model are discussed in this section.

Astrophysical observations predict the existence of so-called dark matter and dark energy which is required to explain the rotation velocity of stars in galaxies that exceeds the predicted speeds based on visible matter alone. Based on our current understanding of gravity, this behaviour can only be explained by assuming halos of additional matter that does not interact with photons and is therefore not visible using telescopes. Hence, the naming as dark matter. So far, there has been no evidence of dark matter based on particle physics experiments. Also, the Standard

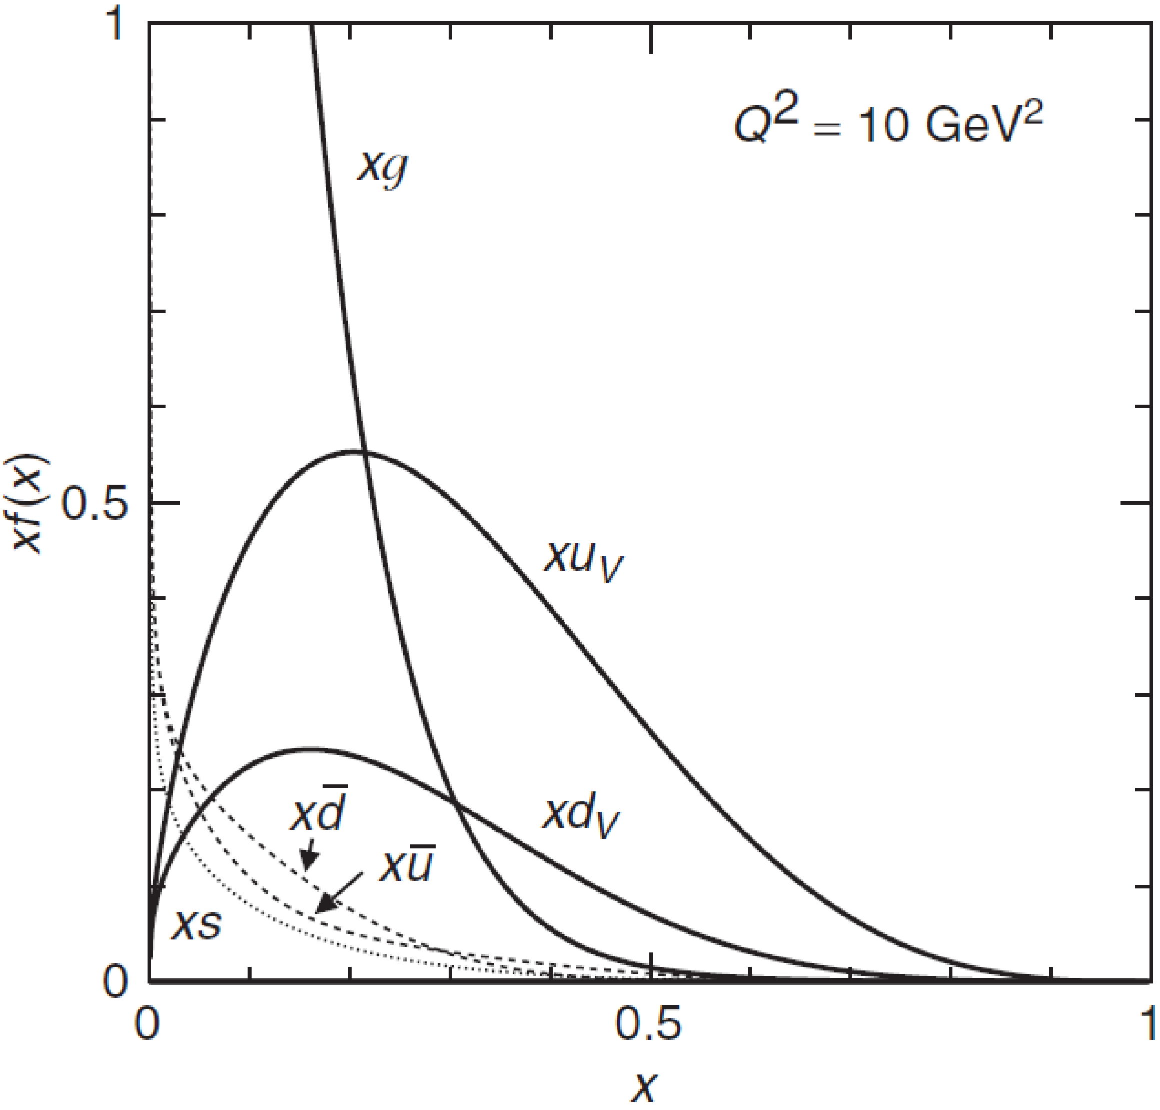

Figure 2.7: Current understanding of the parton distribution functions for the proton at the squared four-momentum of the interaction $Q^2 = 10$ GeV^2 as determined from the MRST fit. Figure taken from [2].

Model does not describe particles that could be used to explain the astrophysical observations. Furthermore, the Standard Model only provides descriptions of the processes around fundamental particles. It can neither explain the observed structure of three generations of fermions nor does it provide a prediction of their masses. Overall, there are at least 18 parameters in the SM that need to be experimentally determined which is unsatisfactory from a scientific point of view. Also, it has not yet been possible to summarise the interactions within the Standard Model into a single Grand Unified Theory (GUT) describing the three types of interactions as different aspects of a single force.

These points and many more, which have not been listed here, lead to believe the Standard Model of particle physics might be an effective theory, useful to describe processes on a certain energy scale, much like Newtonian mechanics can be used in good approximation for processes with low relative speeds with only small derivations from relativistic calculations. Yet, these

points indicate that the Standard Model might not be the final answer to everything.
There are theoretical approaches to solve some of the issues described above. These are often referred to as theories beyond the Standard Model. To this day, none of these extended models were able to provide a more cogent explanation of physics processes or predict particles or interactions that were later discovered and not described by the Standard Model, which therefore remains the leading model in particle physics. Yet the limitations of the SM remain valid, requiring further study of fundamental particles and their interaction in hope of finding new hints that could lead to more complete theories.

2.2 Particle Detection

This section gives a brief overview of the interaction of high-energy particles with matter and the resulting possibilities to detect them, including a short overview of the most common detector types in particle physics experiments. As addition to the sources used for the remaining chapter, this part is also based on 'Teilchendetektoren: Grundlagen und Anwendungen' by Hermann Kolanoski and Norbert Wermes [8].
As only a fraction of the known particles exists naturally in sufficient quantities to study their behaviour, systematic creation of states in which they can be produced is required. Different kinds of particle accelerators have been designed to achieve this, depending on the exact physics processes to be studied. The highest centre-of-mass energies, required to create and study heavy particles, are only achieved by collider experiments using two oncoming beams of particles that cross at a pre-defined interaction point around which the detectors are built. In most setups, the beams are divided into separated bunches. Each crossing of two bunches, potentially leading to interactions between the contained particles, marks an event for the surrounding detector.
For most particle physics experiments, only a relatively small part of the produced particles can be measured directly due to their physics properties. These are the stable or relatively long-lived and light fermions, except for neutrinos who are almost impossible to measure directly, and photons. While μ, neutron, and some of the lighter mesons like $\pi^+(u\bar{d})/\pi^-(d\bar{u})$ or $K^+(u\bar{s})/K^-(s\bar{u})$ can reach the detector[7], other particles will typically decay before traveling a significant distance, leaving only their decay products for detection.
Interactions between particles and materials can roughly be classified into three categories, leading to different types of detectors: charged particles, electromagnetic interaction of electrons and photons, and strong interaction of hadrons. Combinations of different detector types can reveal further information about the encountered particles.
As neutrinos do not fall under any of these categories, the only chance is to indirectly measure them as missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$). This is derived from the conservation of transverse momentum. In high-energy collisions, for example in the ATLAS experiment, momentum and energy of a particle can be assumed to be almost identical, as briefly discussed in section 4.2. For head-on collisions, the total transverse momentum of all produced particles should cancel out. If some particles leave the experiment undetected, the overall measured transverse momentum

---

[7] The exact set of particles from which at least a large proportion reaches the detector heavily depends on the experimental setup. Largest influence comes from the production processes in the collision and the remaining kinetic energy leading to the relativistic effect of time dilation, prolonging the lifetime of particles moving near the speed of light.

will be off by the momentum carried by the unnoticed particles, leading to an $E_{\mathrm{T}}^{\mathrm{miss}}$ component required to cancel out the sum.

2.2.1 Detection of Charged Particles

Charged particles passing through materials interact electromagnetically with the bound electrons, leaving traces of ionised or excited atoms. This effect can be used to reconstruct the tracks of charged particles or extract other information as they pass through the material. There are a number of detector types, called tracking detectors, specialised in measuring the paths of charged particles. For collider experiments, tracking detectors are typically build in a cylindric shape around the beam axes enclosing the interaction point.

Detection of ionisation is possible for example with the use of doped silicon or other semiconductors. The ionisation process induced by charged particles passing at high speeds creates electron-hole pairs which drift in an applied electric potential and can be recombined at p-n junctions which create measurable signals. By forming the semiconductors in pixels or strips and placing them systematically, the path of traversing particles can be reconstructed. If operated within a magnetic field, the extrapolated trajectory of a charged particle can also be used to determine its momentum by measuring the curvature of the track.

In scintillators, excited atoms subsequently release their energy in forms of photons typically in the ultraviolet spectrum. These photons can be converted using fluorescent materials that absorb the UV light and re-emit photons that can be detected using photomultipliers.

Other variants of detectors for charged particles are based on Cherenkov radiation of transition radiation. Cherenkov radiation is emitted when a charged particle passes through a dielectric substance with a velocity higher than the speed of light in that medium. This can be used to compare the speed of particles to thresholds defined by the refractive index of the chosen material or differentiate between the masses of particles with known momenta. Additionally, measurement of the angle at which Cherenkov radiation is emitted can also be used to determine the exact velocity. Transition radiation is created when charged particles pass between materials with different refractive indices. Due to the low intensity of transition radiation, multiple layers of different materials are needed to reliably detect passing particles.

2.2.2 Calorimeters

Both other categories, e/γ and hadrons, can be efficiently stopped in calorimeters to measure their deposited energy. As their leading interactions differ, other physics processes are involved in their detection.

At high energies above $\sim$ 10 GeV, the dominating factor in the energy loss for electrons, as the lightest known charged particles, is bremsstrahlung as a result of their electromagnetic interactions with nuclei. High-energy photons mainly interact via e^+e^- pair production in the electromagnetic field of a nucleus required to guarantee conservation of momentum and energy. These interactions are characterised for electrons by the radiation length (X_0), after which their energy is reduced via bremsstrahlung to $1/e$ of their initial value, and for photons by the mean free path, as the average distance for e^+e^- pair production, of about $7/9\ X_0$. The radiation length of an electron depends on the atomic number of the nucleus, with higher numbers leading to a higher absorption rate.

These closely related characteristics of electrons and photons in matter lead to similar behaviour in detectors as both initial particles start recurring processes of bremsstrahlung and pair production until the original energy is distributed over an electromagnetic shower of ultimately low-energy e/γ and completely absorbed by the surrounding material.

Calorimeters for electromagnetic showers use scintillation light to measure the deposited energy using photon detectors. In order to keep the detectors relatively compact, alternating layers of transparent, active layers, in which the deposited energy can be measured, and high-density passive layers, to stimulate the shower creation, can be used.

For hadrons, energy losses from processes involving the strong interaction dominate their behaviour in matter. As strong interactions are relatively short ranged compared to electromagnetic reactions, the nuclear interaction length (λ_I) is significantly longer than the radiation length for electrons.

While the basic principle of cascading reactions for a hadronic shower is similar to an electromagnetic shower, the details differ not only from the longer interaction length but also regarding the larger number of final states that can be produced. Varying fractions of the energy can also go into electromagnetic showers formed by neutral hadrons decaying into photons, like the neutral pion (π^0).

Due to the (on average) longer range of hadronic particles and their showers in matter, hadronic calorimeters need to provide a larger volume to fully enclose them even if higher-density materials are used. A common way to construct hadronic calorimeters is a sandwich structure of relatively thick absorber layers to develop the shower and thin active layers of scintillating material to measure the deposited energy. This method leads to a smaller fraction of the energy being measured, thus increasing the uncertainties due to the extrapolation. Additionally, about 30% of the initial energy of an incoming hadron is used for nuclear excitation and break-up resulting in processes that cannot be measured by hadronic calorimeters. These effects combined with fluctuations in the compositions of the shower lead to an overall worse energy resolution ($\frac{\sigma_E}{E}$) for hadronic calorimeters of

$$\frac{\sigma_{E,had}}{E} \gtrsim \frac{50\%}{\sqrt{E/GeV}} \tag{2.3}$$

compared to

$$\frac{\sigma_{E,em}}{E} \gtrsim \frac{3\% - 10\%}{\sqrt{E/GeV}} \tag{2.4}$$

for electromagnetic calorimeters.

The basic principle of calorimeters stopping most of the incident particle types can also be used in the identification of muons. Their increased mass compared to electrons significantly reduces their losses from bremsstrahlung, allowing them to escape most detectors in high-energy experiments. Adding layers of tracking detectors in the particle path behind the calorimeter system allows to almost exclusively detect μ particles. So-called spill-out of other particles into the μ system can be identified by the amount of energy deposited in the calorimeters. If no measurement of other particles is required, simple blocks of lead or other high-density materials are also suited to filter for muons.

While normally it is not possible to determine which flavour of quarks were initially produced which formed a hadronic shower, decay processes of b quarks have unique characteristics that makes it possible to identify them with relatively high success rates. Some of the lighter hadrons containing b quarks, like the $\bar{B}^0(b\bar{d})$, are relatively long-lived as their decay processes are based on the weak interaction due to the required change in flavour. If produced in high-energy collisions, this allows the hadrons to travel at least a few mm before decaying. Due to the relatively large mass of b quarks, their decay products can be produced at a measurable angle to the direction of the original quark. If the further decay involves charged particles, their trajectories can be reconstructed in the tracking system and identified to be originated from a position with a certain distance to the interaction point.

The resolution of this secondary vertex against the primary vertex from the original direction of the b quark requires an extremely high resolution in the innermost layers of the tracking system. This process, the so-called b-tagging, plays an important role in many modern particle physics experiments including the ATLAS experiment.

2.3 Jet Algorithms

In particle physics, a jet represents a collection of particles produced from an interaction via the strong force that created high-energy partons. As these partons carry a colour charge, they interact via gluons creating a field of the strong force between them. Given sufficient energy, new partons are created from this interaction leading to successive branching and emission of additional partons. From each initial parton, a so-called parton shower is created while continuously loosing energy from QCD radiations due to the creation of sprays of partons.

Once the partons lost enough energy to reach the confinement scale, the hadronisation process starts, leading to the formation of colour neutral hadrons and thus a hadron shower.

Particles in these hadron showers can decay further and interact with surrounding materials allowing their detection and measurement in calorimeters. These decay processes and interactions can involve purely electromagnetically interacting particles like protons, introducing non-hadrons into the jets.

Measurement and study of jets play an important role in the understanding of reactions in high-energy collisions, quark and gluon properties, and therefore also of the QCD itself.

A schematic representation of the creation of a jet is shown in figure 2.8. Jet algorithms define a section within a calorimeter from where data are combined to calculate the energy deposited

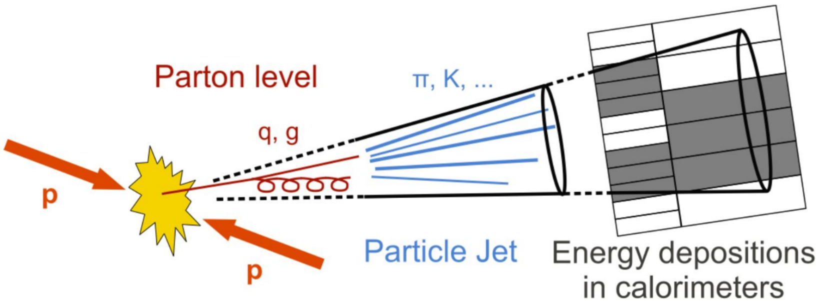

FIGURE 2.8: Schematic representation of the creation of a jet from a proton-proton collision showing the parton level (red), hadronic shower (blue), and energy deposition in a calorimeter (grey). Figure taken from [9].

by a hadronic shower. Most basic algorithms use the data from portions of the calorimeters with fixed sizes and shapes, e.g. a cone, but flexible positions. Other methods, like the k_t and anti-k_t algorithms use a sequential recombination approach of smaller objects to define a jet.
The main jet algorithm used in the context of the ATLAS experiment is the anti-k_t algorithm as introduced in [10]. This algorithm uses an iterative approach to cluster input cells into larger objects which are then defined as jets.
The anti-k_t algorithm uses input objects, for example calorimeter cells with defined energies and positions within the detector, and introduces distances (d_{ij}) between objects i and j as well as for every object its distance to the beam (d_{iB}). The distances are defined as

$$d_{ij} = min(k_{ti}^{-2}, k_{tj}^{-2})\frac{\Delta_{ij}^2}{R^2} \tag{2.5}$$

and

$$d_{iB} = k_{ti}^{-2}, \tag{2.6}$$

where $\Delta^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ with k_{ti}, y_i, and ϕ_i respectively being the transverse momentum, rapidity, and azimuth[8] of object i. Additionally, the radius parameter R is defined that impacts the size of created objects. If multiple versions of this algorithm for jets of different sizes are calculated to account for different physics processes leading to different expected sizes of hadronic showers, they are for the purpose of their distinction typically referred to anti-k_tR. For example, the ATLAS experiment uses anti-k_t4 and anti-k_t10[9].
In every step, the distances from the object with the highest energy to all other objects are calculated. If another object is closest, they are combined to a single new object and all distances are recalculated based on the k_t-weighted position of the newly created object. If the shortest distance is found between object i and the beam, this object is identified as a jet and removed from the process. This continues until no objects are left.
Further refinement to this algorithm can be added by 'trimming' as introduced in [11]. In this process, every jet is again split into its original objects from where the anti-k_t algorithm started. This subset of objects is then reclustered using the same algorithm with a smaller parameter for R. Afterwards, every identified subjet under an adjustable threshold is removed and the remaining subjets are again combined into a trimmed jet.
The reclustering approach of the anti-k_t algorithm and the added trimming account for the often irregular structure of hadronic jets by individually considering or rejecting contributions from small subregions within the overall area of the jet. It also allows to some degree for the elimination of contamination from other objects overlapping with the hadronic shower.

2.4 Physics Motivation for Triggers

The physics processes in particle colliders are largely subject to chance where every possible interaction is represented by a cross section, in the unit of 'barn' with 1 b = 10^{-28} m^2, which depends on the colliding particles and their centre-of-mass energy. This cross section combined with the instantaneous luminosity of a collider experiment, as discussed in section 3.3, determines the average rate in which these interactions occur.

---

[8] For further definitions of these kinematic variables, please refer to section 4.2.

[9] This notation uses R/10 as parameter in the jet definition.

For proton-proton collisions, the conditions for individual interactions are unpredictable as only a single parton with an unknown fraction of the overall momentum from each hadron is directly involved. This uncertainty in the collision conditions comes as a trade-off for the higher centre-of-mass energies achievable with circular hadron accelerators. However, the occurring variation in collision energy can also be used as an advantage in the search for new particles with unknown masses opposed to lepton colliders with an exactly defined collision energy. For example, this was the case during the search for the Higgs boson in the early years of the LHC operation.

For most modern colliders, researchers are not interested in all physics processes that take place. Typically, the main focus is on only a few relatively rare processes with small cross sections or rare decay processes as interactions with larger cross sections are typically already well studied. In order to increase the rate at which the 'interesting' interactions occur, collider experiments are designed to produce multiple collisions of particles per bunch crossing. The remaining collisions, that are not intended to be studied, are referred to as pile-up and present a significant challenge in the data analysis as they overlap with the desired processes.

Modern particle physics experiments provide a fine structure in their detectors for improved resolution. This also leads to an increase in read-out channels and therefore increases the data size per event. For the ATLAS detector, the data collected for a single event amount to about 1 MB. At an average collision rate of 30 MHz, as discussed in the description of the LHC in chapter 3, this would lead to a total data rate of $\sim$ 30 TB/s which is infeasible in both processing and storage.

As the storage of all events is impractical and only a small fraction of events is expected to contain data from interesting physics processes, filters are designed to reduce the data rate while keeping most of the interesting events. These so-called triggers use a subset of the data available from the detector for a quick and rough estimate of its contents to decide which events are recorded for later analysis. For a detailed description of the ATLAS Trigger System, please refer to section 4.4.

Part II

The Experimental Setup

Chapter 3

The Large Hadron Collider

The Large Hadron Collider (LHC) [12] is currently the most powerful particle accelerator in terms of energy and instantaneous luminosity. It is located at the European Organization for Nuclear Research, CERN, near Geneva underneath the France-Switzerland border. The accelerator was designed for the discovery of new particles and interactions and for the precise measurement of Standard Model parameters.

The LHC contains two adjacent parallel beam pipes to travel beams around the ring in opposite directions. It can be operated with beams consisting of up to 2808 bunches of protons or heavy ions, usually lead colliding at a bunch-crossing rate of 40.079 MHz[1]. The main focus is on proton-proton collisions where the design foresees energies of up to 7 TeV per proton beam and luminosities up to 10^{34} cm^{-2}s^{-1}.

An overview of the LHC is shown in section 3.1 including its surrounding complex of pre-accelerators. A short introduction to the main experiments is given in section 3.2 followed by an overview of the operations so far in section 3.3 and plans for future accelerator upgrades in section 3.4.

3.1 Accelerator Complex

The LHC has a length of about 27 km and is built in a tunnel up to 175 m underground. 1232 superconducting dipole magnets are used to bend the beams. The magnetic field within the dipoles is designed to be anti-parallel for the two beam pipes to allow for both beams to be bent by the same magnets. The magnetic field reaches up to $8,3$ T, which is the limiting factor for the possible beam energy in a given accelerator geometry.

On linear parts of the beam line, the bunches are accelerated in eight radiofrequency cavities per beam and focused by 392 superconducting quadrupole magnets. Collisions are possible in four interaction points where the beam lines are crossed at a small angle.

Like any other circular accelerator, the LHC can only be operated using beams within a certain energy range. Due to the configuration of its dipole magnets, the LHC is able to handle beam

---

[1] The average collision rate is reduced by empty buckets to about 30 MHz. The bunch-crossing rate will be referred to as 40 MHz for the remaining thesis as the derivation from this value does not contribute significantly to the design of the experiments.

energies between 450 GeV and 7 TeV for protons and about 24 TeV to 1148 TeV for lead ions[2]. In order to achieve the minimal energy for injection into the LHC, multiple stages of pre-accelerators are required. The complete CERN accelerator complex is shown in figure 3.1.

For proton beams, bunches of ionised hydrogen are accelerated in the LINAC 2, the only linear accelerator in this process, up to an energy of 50 MeV per proton. Subsequently, the bunches are running through the Proton Synchrotron Booster (BOOSTER, also known as PSB), the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS). These three circular accelerators had been in use long before the beginning of the LHC construction but have been repurposed and adapted so they can serve for injecting beams into the LHC. There are two different points where bunches can be injected from the SPS into the LHC feeding the beams for clockwise and counter-clockwise circulation, respectively.

In order to produce ion beams, high purity lead[3] is heated and vaporised to extract $^{208}\text{Pb}^{29+}$-ions, which are accelerated to 4.2 MeV per nucleon in the LINAC 3. Afterwards, a carbon foil is used to further reduce the number of electrons on the ions. After filtering, the remaining ions are accelerated in the Low Energy Ion Ring (LEIR) and the PS. Before entering the SPS and the LHC, the ions are completely ionised by another carbon foil.

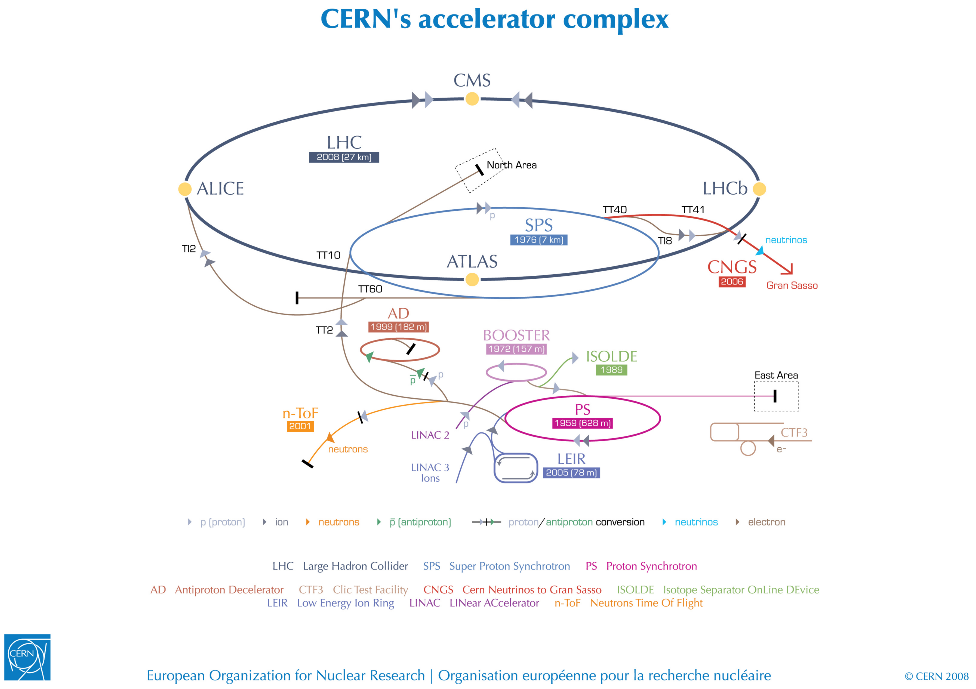

FIGURE 3.1: The complete CERN accelerator complex, including experiments that are not related to the LHC operation. For circular accelerators, the circumference and first year of operation is given. A detailed description of the pre-acceleration steps is given in the text. Figure taken from reference [14].

---

[2] Completely ionised ^{208}Pb-ions at 117 GeV to 2.76 TeV per nucleon.

[3] The use of lead ions is most common for the LHC operation. Yet, for a short period in 2017 xenon was used instead [13]. The injection process for other heavy ions is similar to the described one based on lead.

3.2 Experiments at the LHC

Built around the interaction points of the LHC beams, there are four main experiments located in individual caverns. The ATLAS[4] experiment [15] and the CMS[5] experiment [16] are general-purpose particle detectors to cover a wide range of physics like the Higgs boson search. The two experiments use different setups while targeting a similar physics programme to allow to verify each other's results. A detailed description of the ATLAS detector is given in section 4.

The LHCb experiment [17] is specialized in physics involving bottom quarks, primarily measuring the CP violation in interactions with b-hadrons. As the experiment was designed for lower luminosities of 2×10^{32} cm^{-2}s^{-1}, the LHC beams at the LHCb interaction point are less focused.

The ALICE[6] experiment [18] is optimised to measure the results of heavy-ion collisions and therefore focuses on the lead-filled runs where quark-gluon plasmas are produced to simulate the early stages of the universe.

Additionally, there are three smaller experiments that do not have a dedicated interaction point. The TOTEM[7] experiment [19] is allocated near the CMS experiment. Its goal is to monitor the LHC luminosity and to measure protons from elastic collisions with angles from the beam axis that are too small to be detected by the CMS experiment. The LHCf[8] experiment [20] is placed at a distance of 140 m on either side of the interaction point centred within the ATLAS detector. The main goal of this experiment is to study neutral pions (π^0) to help understand cosmic rays observed in other non-LHC-related experiments. The MoEDAL[9] experiment [21] is primarily used in the search for magnetic monopoles and shares an interaction point with the LHCb detector.

3.3 Performance

Beside the beam energy, the instantaneous luminosity is the key parameter for particle accelerators. It can be described by: $R = \mathcal{L} \cdot \sigma$ as the correlation between the event rate of a physical process R and its cross section σ. The luminosity can be calculated from accelerator parameters as

$$\mathcal{L} = \frac{N_p^2 n_b f_r \gamma}{4\pi \epsilon_n \beta^*} F. \tag{3.1}$$

The numerator consists of the number of particles per bunch N_p, the number of bunches per beam n_b, the relativistic γ-factor, and the revolution frequency f_r. The normalized transverse emittance ϵ_n and the value of the beta function at the interaction point β^* make up the denominator. The integrated luminosity $\mathcal{L}_{int} = \int \mathcal{L} dt$ therefore can be used to determine the expected number of events for physics processes with a given cross section.

The nominal values of the LHC [22] are $N_p = 1.15 \times 10^{11}$, $n_b = 2808$, $f_r = 11.245$ kHz, $\epsilon_n = 3.75 \mu$m, and $\beta^* = 55$ cm. The factor F takes geometrical corrections from the crossing

---

[4] A Toroidal LHC Apparatus

[5] Compact Muon Solenoid

[6] A Large Ion Colliding Experiment

[7] Total Elastic and Diffractive Cross Section Measurement

[8] LHC forward

[9] Monopole and Exotics Detector a the LHC

angle, of about $300\mu\text{rad}$[10] between the two beams, into account reducing the luminosity by a factor of about 0.835. Combined, these values lead to a design luminosity of $\mathcal{L} = 1.0 \times 10^{34}$ $\text{cm}^{-2}\text{s}^{-1}$.

Since 2017, peak luminosities of over twice the design value have been reached. This was achieved by continuous optimisation of the parameters. Main impact came from the reduction of the crossing angle and increased focus of the beams at the interaction points [24].

3.4 Future LHC Upgrades

In order to further increase the instantaneous luminosity and thereby increase the event rate of rare physics processes, the LHC is being upgraded about every $4-5$ years. An overview over the planned upgrades over the next 10 years is given in figure 3.2. Please note that this schedule is based on the assumptions during the development phase of the project presented in this thesis and have since been updated in official sources.

The first upgrade phase (PHASE-I), which already started in 2019 at the end of the data taking period called Run 2, is intended to increase the proton energy to 7 TeV while also increasing the luminosity to 2.5×10^{34} $\text{cm}^{-2}\text{s}^{-1}$ for the duration of Run 3. The second upgrade phase (PHASE-II), which was originally planned to be finished in 2026, includes larger changes while installing the so-called High Luminosity Large Hadron Collider (HL-LHC) [25] to be operated in Run 4. The beam energy will remain unchanged, but the nominal luminosity will be increased to 1.0×10^{35} $\text{cm}^{-2}\text{s}^{-1}$. To enable the increased luminosities, the injectors will be upgraded alongside the LHC. In order to cope with the rising event rates from an upgraded LHC, the experiments located at the interaction points require upgrades to the detectors as well as the

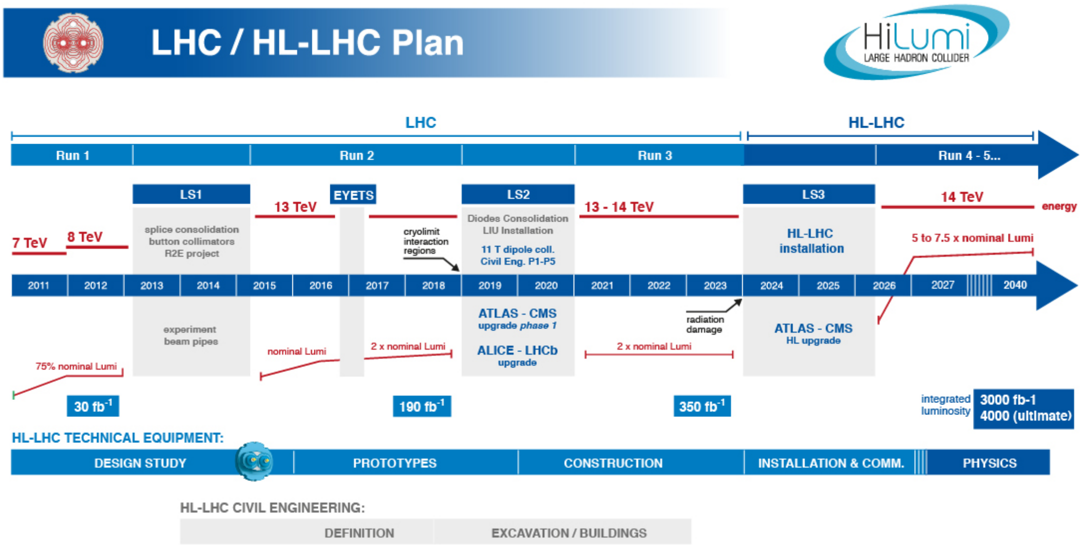

Figure 3.2: Plan for upgrades of the LHC including the centre-of-mass energy for collisions and the expected integrated luminosities per run during the development phase of this thesis. Mentioned are also the required upgrades of the injectors and experiments. Figure taken from reference [26].

---

[10] As measured in 2016 [23].

Trigger and Data Acquisition (TDAQ) systems. A detailed description of the ATLAS upgrades is given in chapter 5.

Chapter 4

The ATLAS Experiment

The ATLAS experiment [15] is a general-purpose particle detector built around one of the interaction points of the LHC. It consists of multiple subsystems that surround the beam axis in layers. The physics programme of ATLAS contains the precise measurement of high energy charged leptons, photons, jets, and missing transverse energy to search for new particles and precisely measure Standard Model parameters. A cut-away view of the detector is shown in figure 4.1.

An overview of the ATLAS experiment is given in the following sections including the coordinate system in section 4.1, commonly used kinematic variables in section 4.2, and the different detector components in section 4.3. Additionally, a general overview of the ATLAS trigger system is given in section 4.4 with a special emphasis on the first stage presented in more detail in section 4.4.1. For a deeper understanding, a short description of the applied algorithms for the jet candidate identification and the summation of transverse energies are described in section 4.4.1.1 and section 4.4.1.2, respectively.

4.1 The Coordinate System

The coordinate system of the ATLAS experiment is a right-handed cartesian coordinate system with its origin in the centre of the detector where the two LHC beams collide in the interaction point. The x-axis points towards the centre of the LHC, the y-axis points towards the surface, and the z-axis, therefore, points counter-clockwise along the beam axis.

The azimuthal angle ϕ is defined within the η-ϕ plane as the derivation from the positive x-axis. The range from 0 to π describes the upper half of the detector whereas $-\pi$ to 0 describes the bottom half. Instead of the polar angle θ , which is measured from the positive z-axis, the pseudorapidity, defined as

$$\eta = -\ln\left(\tan(\theta/2)\right), \tag{4.1}$$

is the most commonly used variable to describe detector dimensions for the ATLAS experiment.

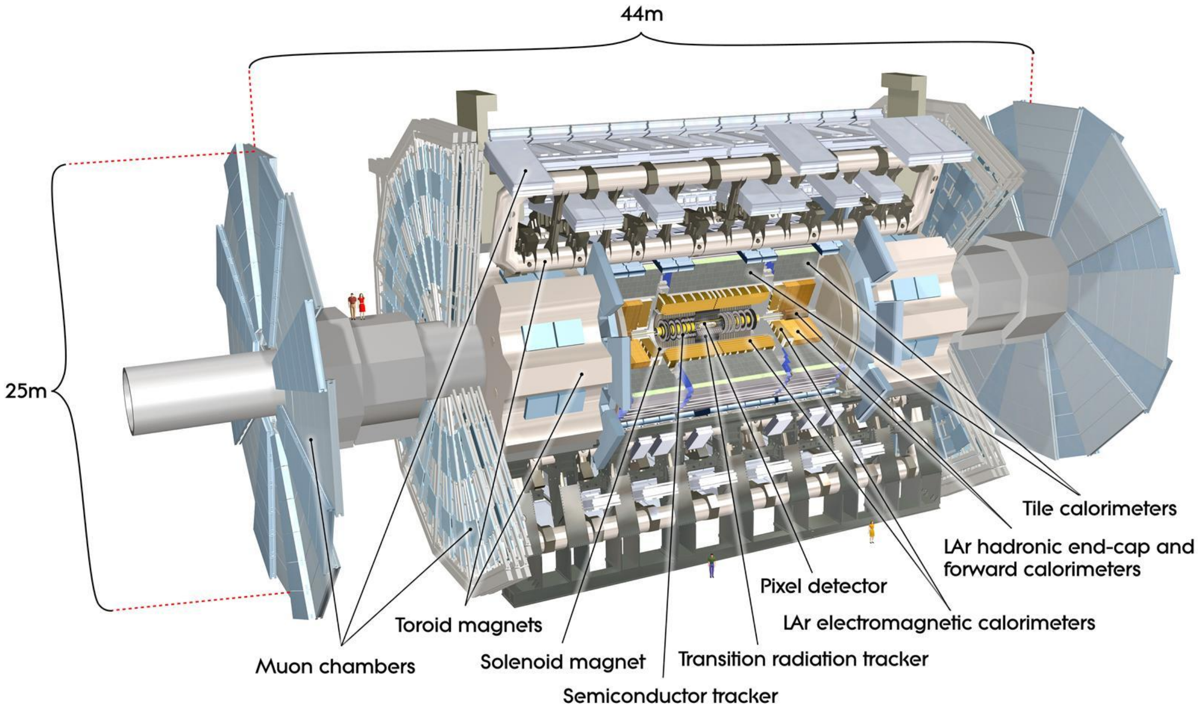

FIGURE 4.1: Cut-away view of the ATLAS experiment. The dimensions of the detector are 25 m in height and 44 m in length. The overall weight of the detector is approximately 7000 tonnes. Figure and description taken from reference [15].

4.2 Commonly Used Kinematic Variables

Using the energy E of a particle and its longitudal momentum p_z, the rapidity of that particle can be calculated as

$$y = \frac{1}{2} \ln \left(\frac{E + p_\mathrm{z}}{E - p_\mathrm{z}} \right), \tag{4.2}$$

measuring the velocity in units of c which this particle would have in classical mechanics without relativistic effects. Differences in the rapidity are Lorentz invariant under longitudinal boosts. For massless particles, the pseudorapidity equals the rapidity. As the LHC beam energy is large compared to the rest masses of most particles that are directly measured by the ATLAS experiment, pseudorapidity differences $\Delta\eta$ are also invariant under boosts along the beam axis in good approximation.
For some aspects, such as describing the size of an object in the detector, the distance ΔR is used and defined as

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \tag{4.3}$$

The interactions caused by the collisions of the LHC beams are in good approximation symmetrical around the beam axis. Therefore, in most cases the separation into components along the x- and y-axis is not required and the transverse momentum p_T and the transverse energy

E_T can be used instead. They are defined as

$$p_\mathrm{T} = \sqrt{p_x^2 + p_y^2} \tag{4.4}$$

and

$$E_\mathrm{T} = \sqrt{p_\mathrm{T}^2 + m^2}. \tag{4.5}$$

For massless particles, p_T and E_T are identical. This can also be used in good approximation for particles with a rest mass that is small compared to their transverse momentum.

For the indirect detection of particles like neutrinos, which hardly interact with the detector, the conservation of momentum is utilised. If the sum over all measured momenta does not equal 0, this indicats undetected particles leaving the detector. As the ATLAS experiment does not cover the complete sphere around the interaction point, particles traveling close to the beam axis will always leave ATLAS undetected, leaving the four-momentum unsuitable for this purpose. The reacting partons of the incoming protons in the collisions have, in first approximation, no transverse momentum. Therefore, momentum conservation requires the vectorial sum over all momenta in the x-y plane to be zero which makes the missing transverse momentum suitable for the indirect detection of particles. This method does not allow to differentiate between multiple such particles. The measured missing transverse momentum and missing transverse energy are defined as

$$\vec{p}_\mathrm{T}^\mathrm{\,miss} = -\sum_i \vec{p}_\mathrm{T,i}^\mathrm{\,rec} \text{ and} \tag{4.6}$$

$$E_\mathrm{T}^\mathrm{miss} = |\vec{p}_\mathrm{T}^\mathrm{\,miss}|, \tag{4.7}$$

respectively. There are multiple methods to calculate these values that differ in which sources are summed up. The method usually used in the ATLAS experiment utilises reconstructed particles such as jets, electrons, photons, and muons as they include well studied energy corrections and noise suppression. Energy deposits in the calorimeters that are not covered by reconstructed particles are also included in the sum.

4.3 Detector Components

The ATLAS experiment consists of multiple detector components, which can be roughly divided into the inner detector, calorimeters, and the muon system.

Closest to the interaction point is the inner detector. It is used to distinguish between charged and neutral particles, like electrons and photons, respectively, and serves as the ATLAS tracking system to identify the paths of charged particles. A description of its subcomponents is given in section 4.3.1. The inner detector is permeated by a solenoidal magnetic field of 2 T to allow for the calculation of the momentum of charged particles from the curvature of their tracks.

The calorimeter system surrounds the inner detector including the solenoid magnet and is used to measure the energy of most kinds of particles except for muons and neutrinos, which can typically escape the detector. This is achieved by stopping these particles and absorbing their energy. The different parts of the systems are described in detail in section 4.3.2. As within the scope of this thesis the only direct detector input is from the calorimeter system, this part of the experiment is described in more detail.

The outermost part of the experiment is the muon spectrometer. It consists of tracking chambers and a toroidal magnet that is separated into a long barrel and two endcaps. The muon spectrometer is described in section 4.3.3.

4.3.1 Inner Detector

The inner detector [27], which comprises the ATLAS tracking system, symmetrically surrounds the beam axis starting only a few centimetres from the beam and extending to a radius of about 1 m at a length of 6.2 m around the interaction point. A cut-away view showing its individual components is shown in figure 4.2. Not shown in this figure is the superconducting solenoid [28] that creates a magnetic field of 2T permeating the inner detector. This field is used to calculate the momentum of charged particles by identifying only a few points of its track to determine the curvature of its flight path in the magnetic field. The inner detector achieves a momentum resolution of $\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} \approx 0.05\%\ p_{\mathrm{T}}\,[\,\mathrm{GeV}] \oplus 1\%$. The inner detector covers the region of $|\eta| < 2.5$.
The innermost part of the ATLAS experiment is the Pixel Detector. It consists of four barrel layers and three disks per endcap to provide at least three space points per track. The innermost barrel layer was added after the 2012 data taking period as the so-called insertable b-layer (IBL) [29]. A single pixel measures $50 \times 400\ \mu\mathrm{m}^2$ of $250\ \mu$m thick silicon. The small size of the pixels compared to other detector components is required to distinguish between tracks close to the interaction point. This information is used to identify secondary vertices of relatively long-lived particles by reconstructing tracks that show an origin with an offset to the interaction point. The IBL is composed of 1744 modules, each measuring $2 \times 6\ \mathrm{cm}^2$ and carrying roughly 47000 pixels.

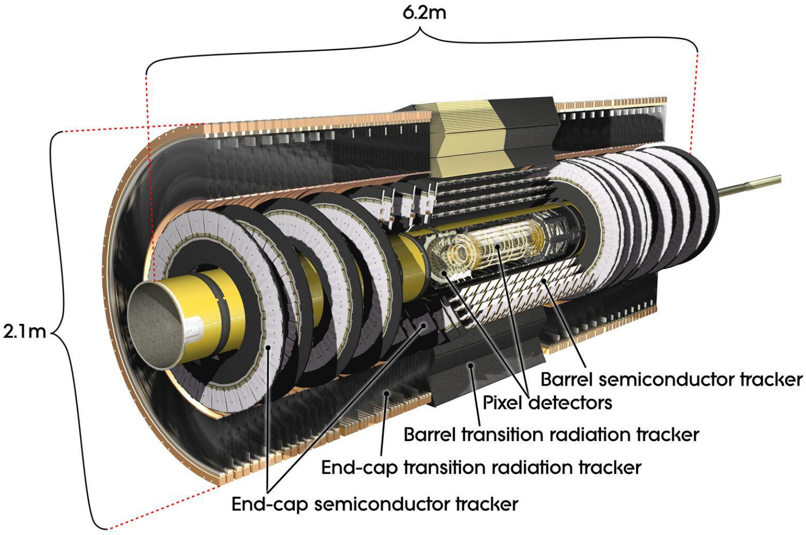

FIGURE 4.2: Cut-away view of the ATLAS inner detector. Figure taken from reference [15].

The Semi-Conductor Tracker (SCT) serves as the second layer of the inner detector. It is built of silicon microstrips, each measuring 80 μm by 12 cm, instead of individual pixels to allow for the coverage of a larger area. In order to narrow the path of the particle down from a microstrip to the equivalent of a pixel, a second layer of microstrips is added, shifted by a small angle to the

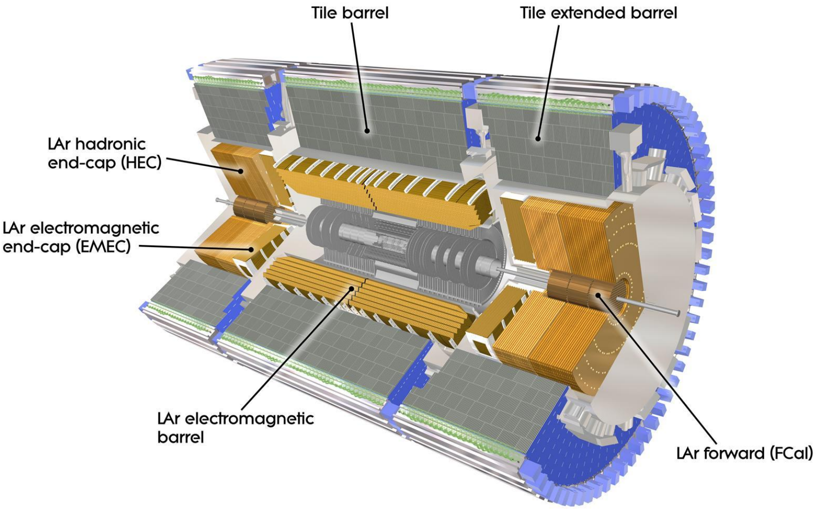

FIGURE 4.3: Cut-away view of the ATLAS calorimeter system. Figure taken from reference [15].

first layer. Barrel and endcaps each consist of four double layers of strips. Detectors based on silicon or other semi-conductors utilise the property of these materials of producing free charge carriers upon exposure to ionising radiation which are absorbed by metal-based electrodes using an electric field.

The Transition Radiation Tracker (TRT) is the outermost component of the inner detector. It is composed of straw drift tubes with a diameter of 4 mm and a length of 144 cm and 37 cm for the barrel and endcap regions, respectively. The use of straw drift tubes allows for the coverage of an even larger volume. Additionally, the straws are embedded in materials with widely varying refraction indices causing particles with a high Lorentz factor $\gamma = E/m$ to produce transition radiation. This is used to distinguish between light and heavy charged particles and serves as an additional criterion in the identification of electrons. The TRT is the only subsystem of the inner detector that only covers a reduced range of $|\eta| < 2.0$

4.3.2 Calorimeters

The ATLAS experiment uses the calorimeter system to measure the energy of most types of particles. The system is subdivided into an electromagnetic and a hadronic calorimeter that are designed to measure the energy based on interactions of particles with material caused by different fundamental forces. Both parts are built as sampling calorimeters consisting of alternating layers of active and passive materials. Incident particles interact with the detector material to create particle showers which are absorbed mainly in the high-density passive material but are periodically sampled in the active layers. In these layers, energy information is obtained by measurement of either low energy photons emitted from scintillating materials or from inductive signals caused by ionisations in drift chambers. Both these methods create signals that allow for the calculation of the deposited energy. This sampled energy information is sufficient to calculate the original particle energy while the passive material used allows for a more compact calorimeter design. An overview of the calorimeter system is given in figure 4.3.

The electromagnetic calorimeter mainly measures photons, electrons, and positrons. The largest proportion of the incoming energy is absorbed by passive material, which consists mainly of lead. Ionisations of liquid argon in the active layers create electromagnetic signals that allow to calculate the absorbed energy. The use of liquid argon requires this part of the detector to be surrounded by a cryostat. Electrons lose the biggest fraction of their energy due to the radiation of photons via bremsstrahlung which produce electron-positron pairs. Given sufficient energy, these pairs again radiate bremsstrahlung leading to a cascade of electromagnetically interacting particles. Incident photons interact like the photons created via bremsstrahlung leading to similar cascades in the calorimeter.

While hadrons are able to interact via the strong force, leading to multiple inelastic scattering processes, they also require more material to be stopped within the detector as the cross-section for these processes are small compared to the interactions dominating electromagnetic cascades. Therefore, a different setup is used to measure the energy of hadronic showers. The hadronic calorimeter uses layers of steel as passive material with scintillating tiles to measure the energy. The endcap and forward sections on both sides of the detector use copper and tungsten as absorber and liquid argon as active material. They are therefore also contained in the cryostats of the electromagnetic calorimeter. As hadronic showers already consist of multiple cascading particles when they reach the detector, the showers created usually spread over a larger area compared to electrons and photons. In accordance with the ATLAS physics programme, the granularity in the hadronic calorimeter can be lower than in the electromagnetic calorimeter to reduce the cost and the number of channels.

Since hadronic showers usually only deposit a small fraction of their energy in the electromagnetic calorimeter while electrons and photons would also be stopped in the hadronic calorimeter, the electromagnetic calorimeter is built around the inner detector with the hadronic calorimeter surrounding it.

A parameter often used to describe calorimeters is the radiation length X_0. It is used for particles that predominantly lose energy via bremsstrahlung, electrons and positrons, and for photons undergoing pair production. It is a material-specific constant that describes the mean distance over which a high-energy electron reduces its energy to $1/e$ of the initial value and is directly related to the mean free path for pair production of high-energy photons. Equivalently, the absorption or attenuation length λ is defined for hadronic particles. These two parameters are used to describe the depth of electromagnetic and hadronic calorimeters, respectively.

4.3.2.1 Electromagnetic Calorimeter

The electromagnetic calorimeter consists of four regions separated in η. The central component, called the barrel calorimeter, covers a range up to $|\eta| = 1.475$. Adjacent to the barrel is the Electromagnetic EndCap (EMEC) calorimeter, which is divided into an outer wheel and an inner wheel. These two parts cover the range $1.375 < |\eta| < 2.5$ and $2.5 < |\eta| < 3.2$, respectively. Thus, the barrel calorimeter and the outer wheel share an η range between $1.375 < |\eta| < 1.475$ where they overlap. Closest to the beam axis, the Forward Calorimeter (FCal) covers the range $3.1 < |\eta| < 4.9$.

The barrel and outer wheels are each separated into three layers and are built in an accordion geometry as shown in figure 4.4. This layout is chosen to provide full azimuthal coverage for

both, active and passive layers. Also shown is the η and ϕ granularity for both, the regular read-out and the trigger towers, and the radiation length per layer which add up to more than 22 X_0 for the complete module. Exemplary, a module from the barrel calorimeter is shown. The endcap modules provide an even higher radiation length of up to 38 X_0.
The achieved relative energy resolution of the electromagnetic calorimeter can be described via

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c. \tag{4.8}$$

The parameters a, b, and c are η-dependent. For low energies, the sampling or stochastic term, a, is the dominant factor representing statistical fluctuation of the shower development in the calorimeter. For low values of $|\eta|$, a reaches about $10\%/\sqrt{E[\text{GeV}]}$ which worsens at higher η-values. The noise term, b, is measured for a mean number of interactions per bunch crossing of $\mu = 20$ as about $350 \times \cosh\eta$ MeV and describes the deterioration of the collected signal by read-out electronics. For high energy particles, the constant term, c, with a design value of 0.7% from e.g. leakage of particle showers and calibration of the read-out system becomes dominating.

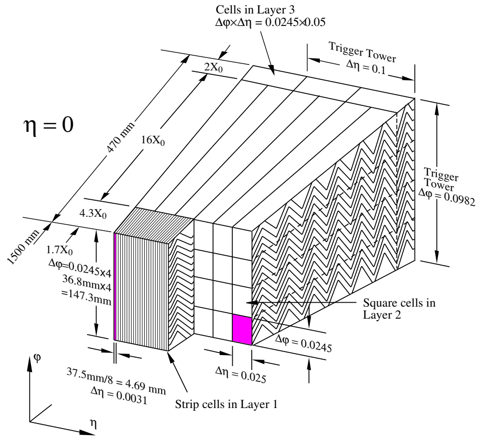

FIGURE 4.4: Sketch of a barrel module where the different layers and the accordion geometry are clearly visible. The radiation length and granularity in η and ϕ of the cells of each of the three layers and of the trigger towers is also shown. Figure taken from reference [15].

4.3.2.2 Hadronic Calorimeter

The central part of the hadronic calorimeter is the tile calorimeter, which covers the range $|\eta| < 1.7$ and is divided into the barrel and extended barrel. Each barrel is divided in 64 modules consisting of steel plates and scintillating tiles. Each edge of a row of scintillating tiles is read out by wavelength shifting fibres into two separate photomultiplier tubes to provide redundancy. A sketch of such a module is shown in figure 4.5. The Hadronic EndCap (HEC) covers the range $1.5 < |\eta| < 3.2$ and, like the electromagnetic calorimeter, uses liquid argon as active component. Additional to the first, electromagnetic layer, the FCal adds two layers as a hadronic calorimeter, which also utilises liquid argon, covering the same range of $3.1 < |\eta| < 4.9$. Combined with the electromagnetic parts, the hadronic calorimeter reaches a radial depth of about 10λ for every component. The granularity is given by the size of the read-out cells, which

provide a resolution of $\eta \times \phi = 0.1 \times 0.1$ in the range $|\eta| < 1.7$ and $\eta \times \phi = 0.2 \times 0.2$ for $1.7 < |\eta| < 3.2$. The FCal provides an irregular η granularity that varies with each layer and a coarse $\phi = 0.4$ granularity.
The relative energy resolution of the hadronic calorimeter in the barrel and endcap regions is

$$\frac{\sigma_E}{E} = \frac{50\%}{\sqrt{E[\text{GeV}]}} \oplus 3\%. \tag{4.9}$$

The forward regions have a reduced relative energy resolution of

$$\frac{\sigma_E}{E} = \frac{100\%}{\sqrt{E[\text{GeV}]}} \oplus 10\%. \tag{4.10}$$

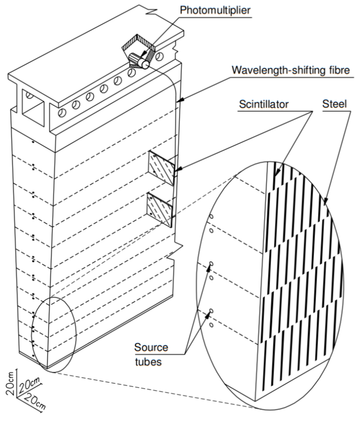

FIGURE 4.5: Schematic view showing how the mechanical assembly and the optical read-out of the tile calorimeter are integrated together. The various components of the optical read-out, namely the tiles, the fibres, and the photomultipliers, are shown. Figure taken from reference [15].

4.3.3 Muon System

As muons are usually able to pass the calorimeter system, they need to be detected and measured separately. For this task, large air-cored toroid magnets [30] combined with separate trigger and high-precision tracking chambers are used to measure the momentum of muons leaving the detector [31]. Tracks measured by the muon systems are also aligned with tracks detected in the inner detector as an additional criterion in the muon identification. A conceptual layout of the muon spectrometer is shown in figure 4.6. The description of the muon system largely follows chapters 1.4 and 6 of reference [15].
Tracks within the range $|\eta| < 1.4$ are bent by the barrel toroid [32] in whose end regions smaller endcap magnets [33] are inserted to cover the range $1.6 < |\eta| < 2.7$. In the transition regions of $1.4 < |\eta| < 1.6$, tracks are bent by a combination of barrel and endcap fields.
For high precision momentum measurements, Monitored Drift Tube (MDT) chambers are used. They cover the range $|\eta| < 2.7$ in three layers except for the innermost layer in which the coverage is reduced to $|\eta| < 2.0$. In this layer, Cathode-Strip Chambers (CSC) are used to cover

the range $2.0 < |\eta| < 2.7$ as they provide higher rate capability and time resolution.
To additionally add the capability to trigger on muon tracks, fast trigger chambers were added to the system, which are capable of delivering track information within a few tens of nanoseconds after the passage of the particle. In the barrel region with $|\eta| < 1.05$, Resistive Plate Chambers (RPC) are used to receive the required information. In the endcaps at $1.05 < |\eta| < 2.4$, Thin Gap Chambers (TGC) were chosen. The outermost parts of the muon system in the range $2.4 < |\eta| < 2.7$ do not include any fast trigger chambers.
The achievable momentum resolution of the muon system depends on the momentum of the particle. For a muon with $p_{\mathrm{T}} = 1$ TeV, a relative momentum resolution of $\sigma_{p_{\mathrm{T}}}/p_{\mathrm{T}} = 10\%$ is achieved.

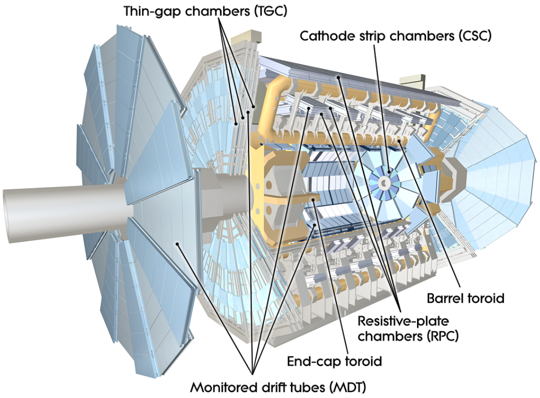

Figure 4.6: Cut-away view of the ATLAS muon system. Figure taken from reference [15].

4.4 ATLAS Trigger and Data Acquisition System

The data collected from a single read-out of the ATLAS detector amount to about 1 MB after zero suppression. At an average collision rate of about 30 MHz[1], the data rate from saving every single event would amount to approximately 30 TB/s. As most interactions that are relevant for offline analysis have a relatively small cross-section, the interaction rate is dominated by background processes. Therefore, it is possible to reduce the event rate by multiple orders of magnitude without losing too many of the desired events. Besides other restrictions in the read-out electronics, this leads to the necessity of a system that controls which events are stored for offline analysis, the so-called trigger system. The main challenge for the trigger system is to efficiently identify the desired interactions in an environment strongly dominated by pile-up and low-energy collisions while providing good rejection for background events.
The ATLAS experiment utilises a two-staged trigger system consisting of a hardware-based first level trigger (Level-1) [34] and a software-based High-Level Trigger (HLT) [35]. Each level runs dedicated algorithms to achieve the best possible differentiation between signal and background events. A schematic overview of the ATLAS trigger system used in Run 2 is shown in figure 4.7.

---

[1] The LHC bunch crossing rate of 40 MHz is reduced by empty bunches to $\sim$ 30 MHz effective rate. Yet, as the system must be able to handle consecutive bunches at the peak rate of 40 MHz, this number is typically used to describe the requirements to the trigger system.

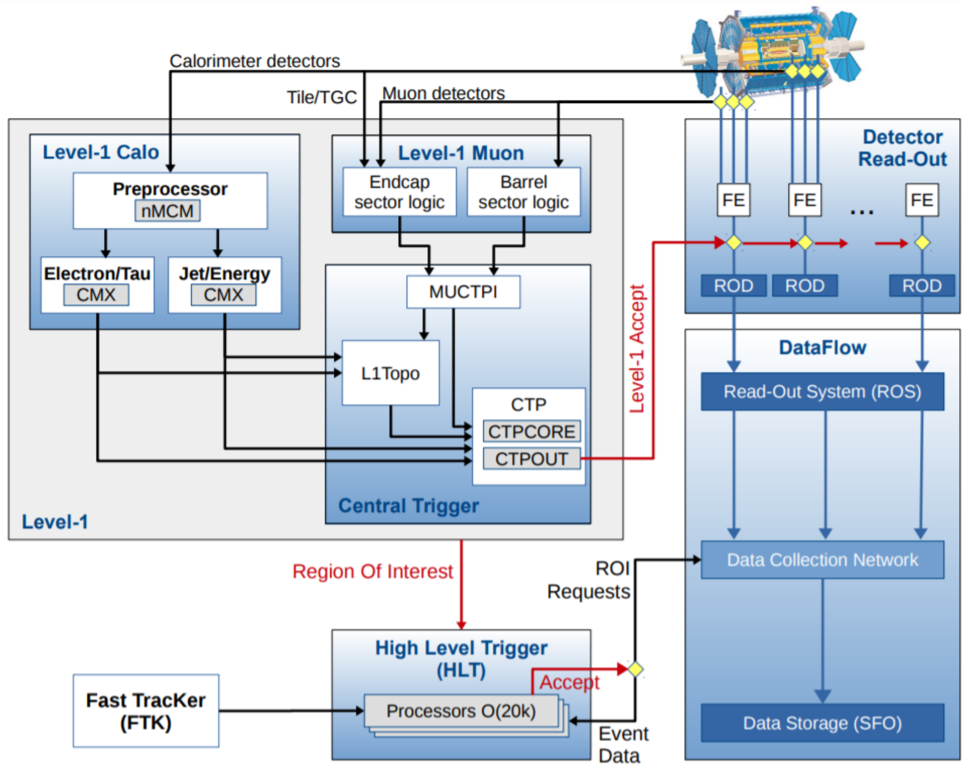

FIGURE 4.7: Schematic layout of the ATLAS trigger and data acquisition system as operated until the end of the data taking period 2018. Figure taken from reference [36].

This already includes several upgrades and conceptual changes compared to the original systems used in the first years of operation. These upgrades were installed during the first long shutdown (LS1, 2013-2014).

The Level-1 trigger is used to reduce the input event rate from 40 MHz to ∼ 100 kHz. For this selection, only input data from the calorimeters and the muon trigger chambers are used. These data are sufficient for a rough estimate of the underlying event while the computation of tracks from the inner detector is too time consuming and the high-precision muon chambers are not able to provide hit information in time for the trigger system. A detailed description of the L1 trigger is given in section 4.4.1. Upon receipt of a Level-1 Accept signal, a more extensive read-out of all detector components is started. This decision is required within 2.5 μs after the collision as the remaining detector information is stored in temporary memory and would irretrievably be lost if the decision signal was not received in time. The Read-Out System collects the event data from the detector, buffers them and distributes them to the High-Level Trigger and to permanent storage in case of a HLT accept signal. Additionally, the Level-1 trigger sends information to the HLT about which detector regions caused the algorithms to accept the event. The HLT further reduces the event rate to ∼ 1.5 kHz utilising the full detector granularity from the affected regions while running offline like algorithms for particle identification. This calculation is performed in ∼ 200 ms. 20000 processors are used to parallelise this last step. Given an acceptance signal from the HLT, the data from the complete detector are sent to the data storage system.

4.4.1 ATLAS Level 1 System

The strict timing and rate requirements to the L1 trigger are met by a pipelined, dead time free[2] system consisting of multiple stages of customized hardware modules that process their input synchronously to the collisions in the centre of the ATLAS detector. One of the main challenges to this system is that the available time span of 2.5 μs also includes the read-out of the detector components, digitisation of the received signals, signal propagation time from the detector to the electronic boards and back to the detector for a potential acceptance signal, and a safety margin in the system design to securely stay within the given time frame. This leaves less than a microsecond for the L1Calo electronics and a processing tome of about 100 ns for the actual trigger algorithms [34]. An overview of the system is shown in the top-left corner of figure 4.7.

Throughout the description of the trigger system including section 5.2, an approximation is used in the form that ϕ value which would precisely be a fraction of π are instead written as decimal number, e.g. $\pi/32$ is expressed as 0.1.

The calculations are done in three steps. The first step is separated between the muon system (L1Muon) and the calorimeters (L1Calo), whereby the calorimeter is again divided into an electromagnetic and a hadronic part. All these components aim to identify candidates for their respective particles as Trigger Objects (TOBs). In the second step, the Topological Processor (L1Topo) is used to combine the results from multiple modules of the first stage. The last stage is the Central Trigger Processor (CTP), which uses the result from topological algorithms as well as lists of TOBs compared to thresholds to make the final decision if a Level-1 Accept signal is to be given. Due to the short processing time and required pipelining abilities of the L1 components, these modules rely on application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs) for their calculations. The basic principle of these is explained in chapter 7.

The basis of the calculations for L1Calo [37] is obtained from on-detector electronics providing analogue signals that are summed in depth over each of the electromagnetic and hadronic calorimeters. The central parts of the detector are covered by trigger towers of $\eta \times \phi = 0.1 \times 0.1$ that become coarser towards the higher η-regions. The analogue signal is transferred from the experiment to a shielded neighbouring cavern hosting the trigger electronics where it is received and digitised by the PreProcessor (PPr).

The PreProcessor Modules (PPMs) themselves serve as carrier boards while the main tasks, such as digitisation and transmission, are executed by smaller modules mounted on top. In a first step, the New Multi-Chip Modules (nMCM) are used to digitise the signal to 10-bit precision. The result is afterwards, among other steps, calibrated, assigned to the correct bunch crossing, and converted to an 8-bit E_{T}-value at a 1 GeV resolution. In a second step, the LVDS Cable Drivers (LCDs) receive the digital signals from the nMCMs and bundle them as required by the Cluster Processor (CP) and Jet/Energy Processor (JEP), respectively.

The CP contains 56 Cluster Processor Modules (CPM) and processes trigger towers in full granularity to identify energy deposits indicative of isolated electrons[3], photons, and τ particles. The JEP consists of 32 Jet/Energy Modules (JEMs) receiving input data in a reduced granularity of $\eta \times \phi = 0.2 \times 0.2$ for the central region to identify jet candidates and calculate transverse

---

[2] Meaning it can receive input every 25 ns while still processing the input from earlier events in different parts of the same hardware components.

[3] As the trigger does not differentiate between electron and positrons, the term electron refers to both particles.

energy sums. CP and JEP modules both form TOBs consisting of a type, energy, and position information. The Common Merger Extended Modules (CMX) reside in each CPM and JEM crate to collect the results from eight modules, count objects above programmable thresholds, and send out counts, sorted TOB lists, and energy sums to the Topological Processor (L1Topo) and Central Trigger Processor (CTP). As not all types of information are required in every L1Topo- and CTP-FPGA, it is also programmable which information is sent per receiving unit. In parallel to the L1Calo system, L1Muon utilises hit information from the fast trigger chambers of the muon system. Muon candidates are identified by requiring coincidences of detected hits in multiple layers of the system, which are consistent with combinations of hits expected from a charged particle originating from the interaction point. As the curvature of the track of a muon in a magnetic field strongly depends on its momentum, up to six different coincidence windows are available. These coincidence windows were determined based on the simulated flight paths of muons with a given momentum in Monte Carlo simulations. The trigger logic is divided between barrel and endcap region for the calculation of the coincidences. Afterwards, a common Muon Central Trigger Processor Interface (MUCTPI) is used to aggregate the results and send it to L1Topo and the CTP.

L1Topo uses identified particle candidates and missing transverse energy information from the previous stage and applies a variety of topological algorithms based on different variables, e.g. to check for angular correlations between objects, to recluster small jets into a single jet that is more similar to the objects used in offline analysis, or to calculate the invariant mass of a potential common mother particle from the transverse energy and direction of its detected decay products. By end of 2018, a total of 113 topological algorithms have been implemented on the L1Topo modules and validated [38]. This option to take additional information of the particle candidates into account for the Level-1 trigger decision was the main upgrade of first trigger stage installed during LS1.

As a final stage, the CTP aggregates information from all subsystems to make the final decision regarding the Level-1 Accept (L1A) signal based on a configurable trigger menu. For the trigger decision, multiple aspects are taken into account like the existing rate limitations and balancing between detected signatures hinting towards different physics processes used in offline analysis. During the first data taking period Run 1, the trigger menu consisted solely of multiplicity triggers of different object types above certain energy thresholds. Starting in Run 2, the menu has been extended by integrating the results of the topological algorithms. In case of an L1A signal, multiple components of the L1 trigger send Regions-of-Interest (RoIs) as input to the next trigger stage and a programmable amount of additional data to the data acquisition system for debugging and analysis of the running algorithms.

4.4.1.1 Jet Identification Algorithm

Due to their constitution of multiple particles with different directions and interactions based on the strong force, hadronic jets typically create a broader shower in the calorimeters compared to electromagnetic particles. Therefore, data from a larger area is required for the identification of jet candidates compared to the algorithms used in the CP [37]. To reconcile this with the limited input capacities of the JEMs, the input data from the central regions of the electromagnetic and hadronic calorimeters are pre-summed in the respective nMCMs to a granularity of $\eta \times \phi = 0.2 \times 0.2$ with a truncated transverse energy resolution of 9 bit. After receipt of the input data, the separate layers are combined to a so-called jet element. Towards the forward regions of the endcaps, some jet elements cover 0.3×0.2 due to a slightly irregular granularity of the detector. Data from the FCal are not binned in η and has a granularity of 0.4 in ϕ.

For the jet candidate identification, a so-called sliding window algorithm is used. An illustration of the different clusters formed for the sliding window algorithm is shown in figure 4.8. Via summation of adjacent jet elements, clusters covering 0.4×0.4 and 0.8×0.8 are formed[4]. The clusters consisting of 2×2 jet elements are formed in a way to create adjacent cells, each of them symmetrically surrounded by a corresponding cluster of 4×4 elements. Each of the smaller clusters is compared to its eight neighbours to search for a local maximum. For neighbouring clusters with identical energy sums, a combination of $>$ and $\geq$ is used in the definition of the maximum to avoid double counting of local maxima or missing potential jet candidates.

If a maximum is found, the energy sums for the two cluster sizes are compared to multiple thresholds and are considered as jet candidates of the associated type if they surpass it. During summation of input data, a saturation of the energy values is monitored due to the limited length of the values. Clusters containing saturated elements automatically pass all thresholds[5]. The availability of multiple window sizes offers the opportunity to choose between the higher efficiency of a large window and the higher ability to differentiate between neighbouring jets provided by a smaller size. For both window sizes, the 2×2 cluster is used as RoI.

As the input rate of the modules used for the calculation is limited by the available FPGAs during the construction phase of the ATLAS experiment, it is not possible to cover the entire calorimeter range with a single module. Therefore, the experiment is virtually split in ϕ-quadrants, which are each subdivided into eight η slices. The regions obtained by this division are each assigned to a single module, which is responsible of identifying jet candidates with their associated RoI within this so-called core area spanning 4×8 jet elements. The JEMs are housed in two crates each containing 16 modules covering opposing ϕ-quadrants and the complete η-range of the detector. The calculation of the larger jet sizes and the search for local maxima on the edges of the assigned region requires additional information. Therefore, data from the environment surrounding the core is also fed into the module extending its input range to 7×11 jet elements. To achieve this, the duplication of a part of the input data is required as the regions from which modules receive data overlap. For the overlap in ϕ, the fan-out is handled by the PPMs, which send copies to both JEP crates. The duplication in η is handled within each JEP crate by sending copies of the received input via the backplane to the neighbouring modules.

The efficiency of the sliding window algorithm normalised to the offline reconstruction depends on multiple parameters such as the utilised window size, the applied Level-1 thresholds, and the offline algorithm. As an example, the efficiency of $\eta \times \phi = 0.8 \times 0.8$ jets with a threshold of 100 GeV compared to the results of an L1Topo reclustering algorithm is shown in figure 4.9 based on a subset of the data collected in 2017. The efficiency is calculated as the fraction of offline anti-k_t jets per p_T-bin that are also identified by the trigger algorithms as jet candidates.

---

[4] A window size of 0.6×0.6 has also been discussed and planned in the original concept but has not been used in the final trigger menu of Run 2.

[5] The CTP automatically sends an L1A signal in this case as the Level-1 system can no longer correctly describe the object.

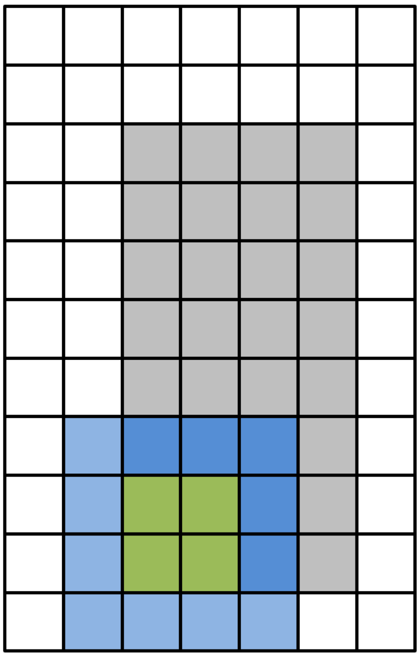

FIGURE 4.8: Illustration of the clusters calculated for the sliding window algorithm. The core area is shown in grey with a surrounding, asymmetric environment in white. A cluster of 0.4×0.4 and 0.8×0.8 are shown in yellow and blue, respectively. The overlap of the blue area and the extended information required for the determination of a local maximum shows the need of an overlap region (white) of data processed as core area of a different FPGA. This illustration represents the regular granularity of input data from the central detector regions.

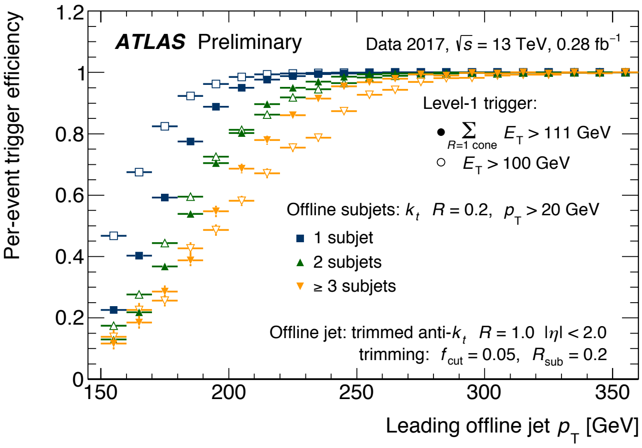

FIGURE 4.9: Level-1 jet identification efficiency for jet candidates identified by the JEP using the $\eta \times \phi = 0.8 \times 0.8$ algorithm (open) compared to the result of the reclustering of these jets at a lower threshold performed by the topological processor (filled). For objects containing multiple subjets, the reclustering algorithm shows a better performance. Data taken from reference [39].

4.4.1.2 Energy Sum Calculation

The calculation of total E_{T} and $E_{\mathrm{T}}^{\mathrm{miss}}$ is split between the JEP and the corresponding CMX. For the calculation of E_{T}, each JEM sums up all jet elements within its core area and sends the result to the CMX. For the calculation of E_{T}, each jet element from the core area is multiplied with two pre-calculated values based on its geometrical position within the detector to calculate the E_x and E_y values, respectively. The results are summed per JEM for each coordinate and sent to the CMX where they are further combined with the input from other JEMs and passed on to L1Topo where they are used to calculate a single global value. Overflows within any of the summations or saturated jet elements are represented by setting the output value to its maximum to force a L1A signal.

Chapter 5

ATLAS Upgrades

As described in 3.4, the LHC is upgraded in multiple stages. To cope with increasing pile-up produced by planned and ongoing upgrades of the LHC, the ATLAS experiment also receives further improvements in the same time spans. This is required as otherwise a sufficient differentiation between different types of particles would no longer be possible. This chapter gives a brief overview of the upgrades to the ATLAS experiment subdivided into changes to the detector itself and the trigger system in 5.1 and 5.2, respectively. Large parts of the ATLAS physics programme require a high sensitivity on different kinds of jets. This includes isolated jets from a multitude of heavy particles, specialised jet algorithms including b-tagging [40] for Higgs bosons decaying into bottom quarks, or boosted top quarks creating large R jets. Further examples and a more detailed discussion of the physics programme for the ATLAS experiment in the coming decades can be found in the letter of intent for the Phase-II upgrade [41] and the scoping document [42].

5.1 Detector Upgrades

This section only covers upgrades that have an impact on the detector configuration, divided into installation planned for the LS2 and LS3, in 5.1.1 and 5.1.2, respectively. Renewal or exchange of read-out electronics or detector materials due to its exposure to intense radiation is not described.

5.1.1 Phase-I

Though there is a certain emphasis of the Phase-I upgrade on the Level-1 trigger, the detector itself is also improved during LS2. The muon system receives a partial replacement called the New Small Wheel (NSW) [43]. While this component is designed for the Phase-I upgrade, it also has to be fully operational in Phase-II at even higher luminosities, which gives important requirements to the performance of the NSW.

The New Small Wheel provides a combination of trigger and precision tracking detectors covering the same range as the previous innermost part of the muon endcaps of $1.3 < |\eta| < 2.7$ with a reduced coverage of the trigger capabilities to $|\eta| < 2.4$. Its aim is to provide a better angular resolution to increase both the trigger efficiency and the resolution of the muon momentum measurement in higher luminosities. In this region, the previously used CSCs and MDT chambers are replaced by Micromegas and small-strip TGCs for precision measurements and trigger functionality, respectively.

5.1.2 Phase-II

To cope with the even higher luminosity rates in Phase-II, several upgrades to the ATLAS detector will be installed during LS3 as introduced in 2012 [41].
During this phase, the entire tracking system of the ATLAS experiment will be replaced by a new all-silicon detector. The new Inner Tracker (ITk) consists of an inner pixel detector [44] and a surrounding strip detector [45], both consisting of a barrel and two endcap regions. The covered range of the tracking system is thereby extended to $|\eta| < 4$.
While large parts of the calorimeter are capable of maintaining their performance after the Phase-II upgrade of the LHC, the forward calorimeters were subject to detailed studies about a possible replacement as they are most affected by the increased particle and energy density. Due to potential irreparable damage to the endcap calorimeters sharing the same cryostats, these plans were discarded and it was decided to instead operate the forward cryostats at a reduced temperature to diminish the effects of the increased luminosity on the energy resolution [46].
In contrast to the calorimeter system, the muon spectrometer will require several upgrades to achieve a sufficient performance in Phase-II [47]. Within the innermost layer of the muon barrel, new RPC and MDT chambers will be installed. A first installation of some of these new detector components is already foreseen for the Phase-I upgrade in a limited range of $1.05 < |\eta| < 1.3$. Overlapping with this pilot installation, TGC doublet chambers in the so-called endcap Inner Large 4 region will be replaced by TGC triplets to achieve a better efficiency in the endcap muon triggers. To extend the coverage of the muon system in the forward region, a high-η tagger is currently under consideration in order to expand the range for muon identification by the region $2.7 < |\eta| < 4$.

5.2 Trigger System Upgrades

The pile-up produced by collisions at the increased luminosities of the upgraded LHC will compromise the isolation of single particles and soften the trigger turn-on of the TDAQ system as it is used in Run 2. Without any updates, the trigger efficiency would therefore drop to a level unsuitable for the ATLAS physics programme. Therefore, it is required to improve the rejection of background events and further optimise the trigger selection for the intended physics programme. One of the most suitable ways to achieve this is an increase of the granularity of calorimeter data available in the first trigger stage.
The descriptions in this chapter are mainly based on the Technical Design Reports [TDRs] for the upgrades of the ATLAS TDAQ System [48] and [49] for Phase-I and Phase-II, respectively. These documents are typically written early in the process of developing the new trigger system to show that it is possible to build a TDAQ system that copes with the increasing demand by a higher instantaneous luminosity. The described systems are therefore based on technology

already available at an early stage. During the development, many aspects are changed due to e.g. the availability of new, more powerful hardware components or conceptual changes. Therefore, the final system can differ significantly from the description in the TDR. Some of these changes and more detailed elaborations are already included in this section but it will not necessarily represent the final design of the system in all details.

The upgrades planned for the HLT and Read-Out System determine the boundary conditions for the first trigger stage. As the boundary conditions for the Level-1 trigger are mainly dictated by the front-end detector electronics, these values remain identical to Run 2 for the Phase-I upgrade with an average output rate of about 100 kHz and a latency of 2.5 μs including read-out electronics. The HLT further reduces this rate to 1kHz. For the operation in Phase-II, the latency is raised to 10 μs. The Event Filter, as the new last stage of the trigger system, and the upgraded read-out electronics are designed to handle an input rate of 1 MHz and reduce it to about 10 kHz for storage. An overview of the Phase-I trigger system is given in 5.2.1 followed by a brief description of the Phase-II system in 5.2.2. Upgrades to the final trigger stage and surrounding components like the Read-Out System are not described in detail within this thesis. Detailed description for these systems can be found in [48] and [49] for the respective phases of the upgrade. Updated documents covering the changes to each component after the TDR will be published for each subsystem during its development.

5.2.1 Phase-I

Like in the trigger system operational during Run 2, the Phase-I trigger relies on a two-staged architecture using a first stage based on custom electronics and a second stage utilising large CPU farms. An important requirement for the development of the Phase-I upgrade was to be compatible with the Phase-II conditions where feasible. Most of the following descriptions are based on the Phase-II Technical Design Report [49] as this document is more recent and provides a good overview over the Phase-I TDAQ system.

The overview of the trigger system in figure 5.1 shows the Level-1 trigger in light purple with subsystems newly introduced in Phase-I. The system consists of two parallel parts, L1Calo and L1Muon. L1Calo uses data from the calorimeter system at a finer granularity compared to the Run 2 system to identify energy deposits from particles that are stopped within the ATLAS experiment. It consists of 3 new subsystems called Feature Extractors, which are installed during the Phase-I upgrade. The electron Feature Extractor (eFEX) receives data mainly from the electromagnetic calorimeters to identify deposits from electrons, taus, or photons replacing the Cluster Processor. The jet Feature Extractor (jFEX) and the global Feature Extractor (gFEX) receive data from all parts of the calorimeter. The gFEX uses input data in a coarse granularity to identify large area jets and to calculate global quantities like missing transverse energy. Designed as a replacement for the Jet/Energy Processor, the jFEX utilises a finer granularity of calorimeter data to calculate smaller jets and global quantities while providing more flexibility in the definition of the algorithms used. A detailed description of L1Calo is given in 5.2.1.1. While L1Muon receives an additional component to utilise the triggering capabilities of the NSW, the task for the overall system remains identical to L1Muon in Run 2: processing data from the fast trigger chambers of the muon system to identify tracks from charged particles leaving the detector. The Muon CTP Interface (MUCTPI) gathers the output from the different muon trigger subsystems, combines them, and sends the results to the Topological Processor

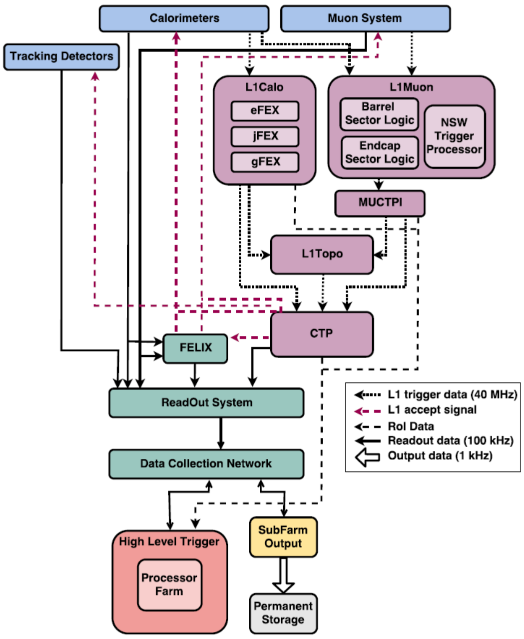

Figure 5.1: The architecture of the ATLAS trigger and data acquisition system after the Phase-I upgrade. The Level-1 trigger system shown in light purple only uses input from the calorimeters and the muon system. The final trigger decision is made in the High-Level Trigger shown in red, which has access to data from all detector subsystems. For simplicity, not all subsystems and connections are shown in this figure but rather a reduced view focusing on new components. Figure taken from the Phase-II TDR [49].

(L1Topo) and the Central Trigger Processor (CTP). The new L1Topo is a revised version of the Run 2 counterpart utilising latest generation FPGAs who provide a higher input bandwidth and more processing resources. The task of the CTP also remains identical to Run 2: it receives results from L1Calo, MUCTPI, and L1Topo and makes the final trigger decision for Level-1 based on a programmable trigger menu.

Systems used to collect and distribute data are shown in green. This includes the Front-End Link Exchange (FELIX), the Read-Out System (ROS), and the Data Collection Network. FELIX is a new component used as an interface between the front-end electronics and the remaining data acquisition system. The ROS collects Regions of Interests (RoIs) from the Level-1 electronics upon the receipt of a Level-1 Accept Signal (L1A), initiates the complete read-out of the detector data corresponding to the trigger signal, and buffers this data for the HLT and potentially for permanent storage.

The HLT software is updated during the Phase-I preparations to adapt to changed input from updated detector components and new RoI formats from the Level-1 trigger. These RoIs are used to focus the HLT calculations on detector regions where the first stage identified candidates in order to reduce the required processing time per event.

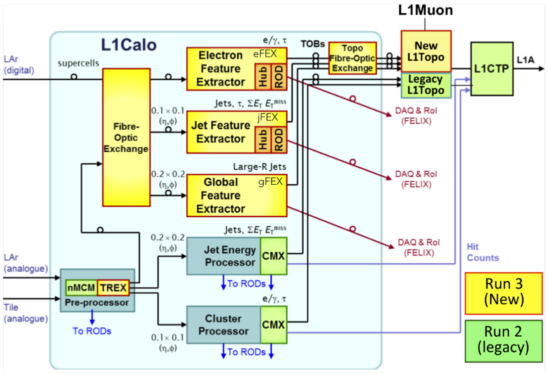

FIGURE 5.2: The architecture of the ATLAS L1Calo trigger system as intended for the start of Phase-I. Components shown in yellow will be added during LS2, while the legacy system shown in green and blue will be shut down and removed once the new system is fully commissioned and its reliability has been verified. Figure taken from reference [50] and modified.

5.2.1.1 L1Calo

A more detailed overview of the Level-1 trigger in Phase-I is shown in figure 5.2. Components shown in yellow are newly developed systems for Phase-I while the legacy system from Run 2 is shown in green. The legacy system is intended to run in parallel to the new components for some time to verify the correct function and reliability before the old system is decommissioned. As in Run 2, the pre-processor receives and digitises analogue signals from the ATLAS calorimeters. The new Tile Rear Extension (TREX) gathers the digitised signals and bundles them according to the needs of the FEXs. The input of LAr signals shown in figure 5.2 is exclusively used for the legacy system. The new Phase-I Feature Extractors are directly supplied with digital signals from the new LAr Trigger prOcessing MEzzanines (LATOME) which provide an extension of the LAr Digital Processing System (DPS). This system provides optical signals directly from the front-end electronics. Both input stages, TREX and LATOME, provide data to the FEXs at a finer granularity compared to their respective Run 2 counterparts while utilising higher data rates per link to compensate the increased amount of data per event. Please note that these systems were not described in the TDR but were included in the processing concept at a later stage.

The aim of the eFEX is an efficient cluster formation to identify electrons, photons, and taus while also providing an improved jet rejection. This is achieved by cluster algorithms that use a size which is just large enough to contain most of the energy deposited by an average electromagnetic shower, but smaller than a typical jet. To reduce the background rates to an acceptable level, further parameters have to be taken into account like distribution of energy within the cluster or energy depositions around it. In order to calculate these parameters, the eFEX receives data from the barrel region at $|\eta| < 2.5$ in 4 separate layers from which the medium layers have a granularity of $\eta \times \phi = 0.025 \times 0.1$ and the outer layers $\eta \times \phi = 0.1 \times 0.1$. Data from the hadronic calorimeters is only received from the first layer at a granularity of $\eta \times \phi = 0.1 \times 0.1$.

The jFEX is primarily responsible for the identification of jets that are reconstructed in offline analysis using the anti-k_t algorithm with a radius of 0.4. In order to further approximate the offline algorithm and thereby increasing the trigger efficiency, a finer granularity is required compared to the JEP. This is achieved by reception of trigger towers at a granularity of $\eta \times \phi = 0.1 \times 0.1$ for $|\eta| < 2.5$, $\eta \times \phi = 0.2 \times 0.2$ for $2.5 < |\eta| < 3.2$, and a varying granularity in the FCal where super cells are received. Within each calorimeter, the layers are summed in depth as no use of parameters depending on the energy distribution is intended. Improvements of the baseline algorithm like approximations of round windows or application of weights to each trigger tower are subject to detailed studies. Further use cases for the jFEX are the identification of tau particles decaying mainly hadronically into sizes that are too large for the eFEX, electrons beyond $|\eta| = 2.5$, and large area jet potentially including the identification of a sub structure. Additionally, the jFEX is capable of calculating global values like total and missing transverse energy. A detailed description of the requirements for the jFEX and the resulting design is given in chapter 6 and 8, respectively.

The gFEX receives data from the entire calorimeter at a coarser granularity of $\eta \times \phi = 0.2 \times 0.2$ on a single module. This allows for the calculation of global values as well as large area jets, approximating the offline algorithm with R = 1. A specific physics case for larger jets are boosted hadronic topologies that are characteristic of many new physics scenarios containing extremely heavy particles leading to multiple highly collimated jets.

On the backplane, eFEX and jFEX share a common approach of 2 Hub modules with Read-Out Driver (ROD) daughter boards in each shelf. The purpose of these modules is to aggregate read-out data from the FEX modules and provide timing and control functionality. As the gFEX comprises of a single module, the Hub and ROD functionality is directly implemented in this module.

All FEXs and the MUCTPI send Trigger Objects (TOBs) containing energy and location information and an identifier for the used algorithm to the Topological Processor. The L1Topo is replaced by new hardware providing essentially the same functionality as in Run 2 while offering significantly increased input bandwidth and processing resources. This is required as the new FEXs are designed with more outgoing links and the capability of higher link speeds to cope with the larger number of expected TOBs due to the increased pile-up.

5.2.2 Phase-II

An overview of the Phase-II TDAQ is shown in figure 5.3. While the structure of the overall system remains similar to the architecture in Phase-I, a key component for improved trigger performance capable of handling the HL-LHC luminosities is not directly shown in this overview. The front-end electronics for the detector components are significantly improved to allow for higher latencies in the first trigger stage and an increased read-out rate without excessive deadtimes. This upgrade allows the first trigger stage to run more elaborate algorithms that are closer to the offline reconstruction algorithms. A higher read-out rate of the entire detector also allows for the construction of a new last trigger stage that handles 10 times the input rate of the Phase-I system, which used to be limited by increasing deadtimes for read-out rates above 100 kHz. Detailed descriptions of the improved front-end electronics can be found in the corresponding documentations for each detector part.

The first trigger stage, which is now labelled as Level-0, mainly consists of the legacy system from Run 3. The Phase-I FEXs remain in operation though their limited granularity is no longer sufficient to single-handedly identify particle candidates as a first trigger stage. Therefore, new

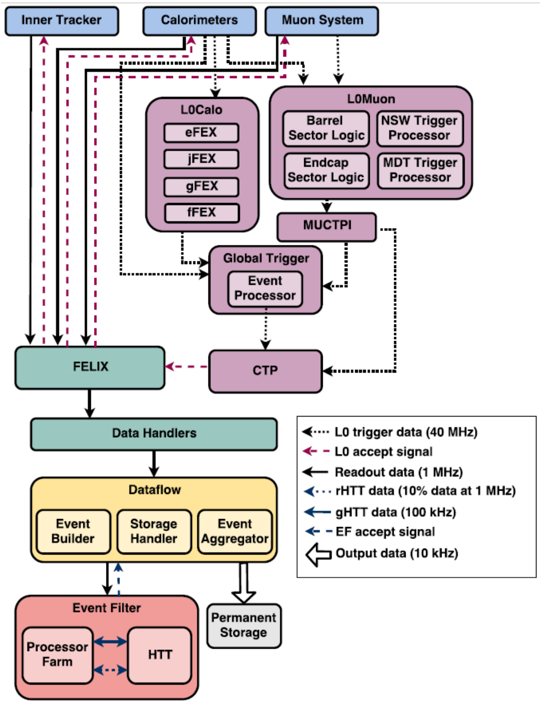

FIGURE 5.3: The architecture of the ATLAS trigger and data acquisition system after the Phase-II upgrade. The first trigger stage is marked as Level-0 due to temporary plans to introduce a new Level-1 that were discarded during the further development of this upgrade. The overall structure remains similar to the system used during Phase-I. An overview of the Phase-II trigger system is given in 5.2.2. For simplicity, not all subsystems and connections are shown in this figure but rather a reduced view focusing on new components. Figure taken from the Phase-II TDR [49].

algorithms are designed to efficiently seed further calculations by the Global Processor. As the legacy system remains almost untouched during the Phase-II upgrade, the impact of Phase-II operational conditions on the Phase-I FEX architecture is limited. To make better use of the ITk by increasing the trigger capabilities in the forward region, a new forward Feature Extractor (fFEX) is built. This component receives data from the forward regions at full granularity and identifies electromagnetic showers in the region of $2.5 < |\eta| < 4.9$ and jets in the region of $3.2 < |\eta| < 4.9$. The L0Muon system receives a new processor to utilise the trigger capabilities of the MDT chambers.

To compensate the limitations of the L0Calo, a new Global Trigger is introduced replacing the Topological Processor, which receives full-granularity calorimeter data to run offline-like algorithms. These results are combined in topological algorithms and multiplicities with the TOBs from L0Calo and the MUCTPI. Results are sent to the CTP, which no longer receives direct input from L0Calo, for the final Level-0 trigger decision.

Upon the receipt of an accept signal from the first trigger stage, the detector components send their data to the DAQ. As the first element in this system, FELIX gathers these data and accumulates them on a common interface for the Data Handlers, which are used for detector-specific processing of the input. From there, data are sent to the Dataflow system which buffers them for the Event Filter and permanent storage. As the new final stage of the trigger system, the Event Filter consists of a large processor farm with additional hardware-based support for track calculation by the HTT. Internally, the Event Filter reduces the input rate of 1 MHz in multiple steps down to the expected 10 kHz output rate at which data is written to permanent storage.

Part III

The Jet Feature Extractor

Chapter 6

jFEX Requirements

In this chapter, the detailed requirements towards the jFEX on interfaces and processing capabilities are presented. The main points were initially defined in the Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System (TDR) [48]. During the development phase, many of these requirements have been altered due to newer concepts for either the jFEX or the remaining Trigger System. These changes are included in this section up to the end of the design phase for the first jFEX prototype. New requirements that were introduced and taken into account after the production of the first prototype are not discussed in the scope of this thesis. In some cases, no official documentation is available. Therefore, the exact descriptions in this chapter can differ from those which can be found in the TDR or in later specifications.

An overview of the required interfaces of the jFEX system is presented in figure 6.1. Please note that this represents the first draft from the Phase-I TDR. Several details within the design have been changed, like the number of FPGAs per module. Yet, the overall structure regarding interfaces remains relevant. The represented real-time data path is described in detail in section 6.1 and section 6.2 for the input and output side of the jFEX, respectively. During the early development stages of the Phase-I trigger upgrade, the ATLAS L1Calo community agreed on a maximum bit error ratio of 10^{-14} for all transmissions in the real-time data path.

Algorithms are only briefly discussed in section 6.3 based on the specifications from the Phase-I TDR. More detailed descriptions are given in section 9.1 where restrictions from the implementation in firmware can be taken into account. Requirements imposed by the chosen form factor are described in section 6.4. Additional specifications from the ATLAS Trigger System on control functionalities and read-out interfaces are discussed in section 6.5.

As described in chapter 5, the jFEX system consists of multiple modules, each covering a part of the ATLAS detector in terms of processing calorimeter data. Within each module, the covered part is again split between single FPGAs. In contrast to the TDR, the presented design consists of 6 jFEX modules each containing 4 FPGAs for processing detector data. Each module receives data from a part of the detector that is limited in η but covers the whole detector in ϕ forming a so-called 'ϕ ring'. Within each ring, a single FPGA covers the complete η-range of the module and a quadrant of the detector in ϕ.

An estimate of the available latency on the real-time path per subsystem in the Level-1 Trigger System is given in the TDR. For the jFEX system, a latency of 14 LHC bunch crossings (BCs)[1]

---

[1] 14 BCs $\approx$ 350 ns

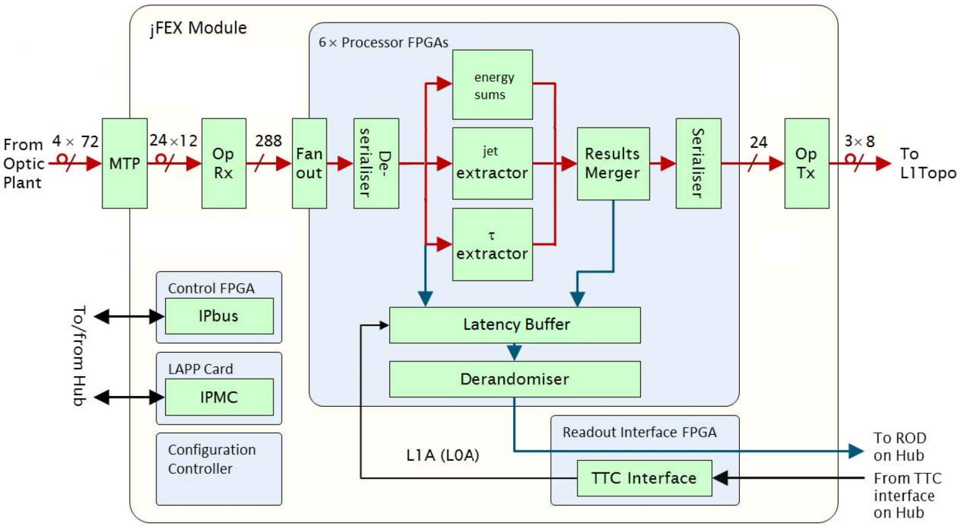

FIGURE 6.1: A block diagram of a jFEX module showing the real-time, read-out and control data paths. Please note that this represents the first draft from the Phase-I TDR. Several details within the design have been changed, like the number of FPGAs per module. Yet, the overall structure regarding interfaces remains relevant. Figure taken from the Phase-I TDR [48].

is intended. Most of this time is expected to be required by input data reception and output transmission, leaving an estimated 100 ns to run the algorithms.
The main focus of the presented requirements and following jFEX design lies on the ATLAS detector, its data output, and Trigger System as it will be available during Run 3. Compatibility with the Trigger System following the Phase-II upgrade is already taken into account in the specifications for Phase-I and otherwise considered in the design of the Run 4 system.

6.1 Real-Time Input

As described in chapter 5, the jFEX system receives input data from both, TREX and LATOME, on optical links, sorted by the Fibre-Optic eXchange Plant (FOX). These modules are designed as extensions of the legacy digitisation system and are not described in the TDR. One of the main requirements the jFEX system has to meet is given by the number of input links these systems provide. This depends on the used link speed, required resolution per tower, and the coverage of the different detector regions by the DPS system. The ATLAS Level-1 Trigger System is designed as a pipelined system without buffers. Therefore, the jFEX is required to receive and process data sets at the LHC clock rate of about 40 MHz, which constitutes an important requirement towards the used algorithms.
The processing of the calorimeter data in multiple FPGAs requires the duplication of input data. Details regarding the intended algorithms are described in section 6.3. The splitting of the jFEX system into separate modules also requires most of the received links to be duplicated at source, namely TREX or LATOME, or via optical splitting in the FOX. While there are no connections between the separate modules, which results in the necessity of input data duplication before entering the jFEX system, the duplication required by the subdivision of a module into 4 FPGAs can be done on board to reduce the required input bandwidth.

This leads to a difference between the number of unique fibres, carrying data combinations different from each other, and total number of links. An overview of the $|\eta|$ dependent duplications required per link in Phase-I is given in Table 6.1. This setup can change in Phase-II if further jFEX modules are added to cover the forward regions with a finer input granularity. The HEC region within $1.5 < |\eta| < 1.6$ is a special case since it overlaps with the TileCal. Since TileCal and HEC data are processed in separate systems, the input data cannot be merged before sending the data to the jFEX. For modules that cover the region $2.4 < |\eta| < 3.2$, this overlap data can be sent on spare capacity on links covering this outer region. A separate link is only required for the two modules covering the innermost part of the detector around $\eta = 0$, increasing the total fibre count for these modules.

$\lvert\eta\rvert$ range	required duplication
$\lvert\eta\rvert < 1.6$	3
$1.5 < \lvert\eta\rvert < 1.6$ (HEC only)	1
$1.6 < \lvert\eta\rvert < 2.4$	2
$\lvert\eta\rvert > 2.4$	1

TABLE 6.1: Overview of the $|\eta|$ dependent duplication of input links to the jFEX system for Phase-I. The single copy listed for the HEC overlap region only applies to a single copy per side of this data being sent on a separate fibre. Two additional copies are included in fibres transmitting data from the region of $2.4 < |\eta| < 3.2$. This duplication is handled by the respective input stages.

Deviating from the TDR description, the baseline link speed for all links received by the jFEX system is considered to be 11.2 Gb/s. Though subject to tests to verify the reliability of data transmission at higher rates than the TDR baseline of 6.4 Gb/s, large parts of the L1Calo community shared the assumption that all newly developed boards will allow for stable operation at higher rates. Therefore, the higher transmission speed and simultaneously increased possibilities for the trigger setup were included in the jFEX requirements and design from an early stage on. The verification of compatibility of all newly designed modules with data transmission at rates of 6.4 Gb/s, 9.6 Gb/s, 11.2 Gb/s, and 12.8 Gb/s were scheduled with the first prototypes. The results of the test featuring the jFEX prototype are discussed in section 9.3.

While no final format decision was made during the jFEX design phase, the main impact for the design comes only from the number of trigger towers per link and the $\eta \times \phi$ coverage of individual LATOME and TREX modules. The combination of both factors implies possible combinations of trigger towers on individual optical links and therefore defines the possible configurations of the jFEX in terms of coverage of the ATLAS detector.

Within a bunch tick of 40 MHz, a total of 280 bits can be sent at 11.2 Gb/s. While the exact data formats were undefined for most of the jFEX development phase, the use of the 8b/10b encoding scheme was defined in the TDR. This scheme uses 10-bit symbols to encode 8 bits of data. The included overhead from the transmission reduces the effective packet size per link to 224 bits per bunch crossing. Advantages of this scheme include good possibilities to recover a clock signal from received data streams since frequent transitions between low and high state are ensured and the DC balance[2] is ensured as on average the same number of logical 0s and 1s are received. Additionally, some bits per data packet are reserved to perform a cyclic redundancy

---

[2] Maintaining the DC-balance can help increase the transmission quality in capacitor-coupled links.

check (CRC), used to detect transmission errors, and information on the saturation of trigger towers.

For efficient use of the available input links, a regular mapping of the towers per DPS module is required. The optimal number of towers per link was found to be 16, providing a high density of towers per link to allow the modules to receive data from a bigger part of the detector while allowing a reasonable energy resolution per tower of at least 10 bits.

An overview of the η dependent input granularity is shown in figure 6.2. In the central region with $|\eta| < 2.5$, a regular resolution of 0.1×0.1 $(\eta \times \phi)$ is available. Data in a resolution of 0.2×0.2 $(\eta \times \phi)$ is available for $2.5 < |\eta| < 3.1$ followed by a single strip of 0.1×0.2 $(\eta \times \phi)$ towards the end of the region covered by EMEC and HEC. Data from the FCal at $3.2 < |\eta| < 4.9$ is received in a 0.4 binning in ϕ. The η binning consists of 12, 8, and 4 bins for the first, second, and third layer, respectively, with each layer covering the complete range of 1.7 in η.

The TREX system digitises data coming from the hadronic tile calorimeter. Each module covers 0.4×1.6 in $(\eta \times \phi)$ requiring a total of 32 TREX modules to cover the TileCal region up to $|\eta| < 1.6$. Therefore, 4 unique links are required per module, each containing data from 0.4×0.4 in $(\eta \times \phi)$.

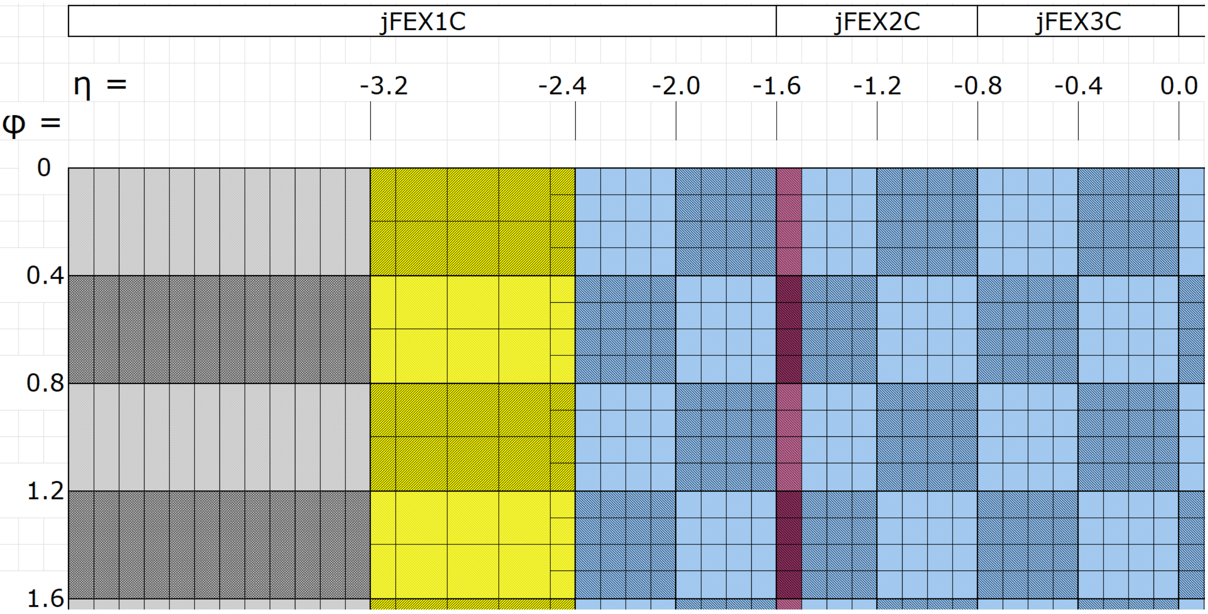

Figure 6.2: Mapping of the different input regions for the jFEX system in the η-ϕ-plane. The blue area represents the central detector region up to $\eta = 2.4$ with its regular granularity and a red stripe representing the HEC-TileCal overlap. The yellow region represents the slightly irregular region of $2.4 < |\eta| < 3.2$ containing different granularities and the grey area shows the FCal. Each box represents a trigger tower, and the shading indicates the separation between the fibres used to send this data. The η binning in the FCal shows the finest available granularity. This scheme does not differentiate between electromagnetic and hadronic calorimeters within the detector though they are received separately. The coverage is limited to 1.6 in ϕ and the detector side with $\eta < 0$ for reasons of readability. Remaining regions are mapped analogously.

The LATOME system covers detector parts with various granularities. Each module covers 0.8×0.4 in $(\eta \times \phi)$ for EMB and EMEC up to $|\eta| < 2.4$, 0.8×0.8 in $(\eta \times \phi)$ for the remaining part of the EMEC and 1.7×1.6 in $(\eta \times \phi)$ for the complete HEC. The first FCal layer is covered by a single LATOME module per side that processes data from 1.8×6.4 in $(\eta \times \phi)$ starting at

$|\eta| = 3.1$ and one module per side for the second and third layer combined covering 1.7×6.4 in $(\eta \times \phi)$ starting at $|\eta| = 3.2$. Since the main part of energy deposition at the overlap of the endcaps and Forward Calorimeter is expected in EMEC and HEC rather than the first FCal layer, the jFEX systems only considers FCal inputs to start at $|\eta| = 3.2$.

In the forward regions, a larger area in $(\eta \times \phi)$ can be covered by a single LATOME module since the granularity is reduced[3] for output of the electromagnetic calorimeter beyond $|\eta| = 2.5$ and due to the generally fixed requirement of 0.1×0.1 in $(\eta \times \phi)$ or coarser for the hadronic calorimeters.

Fibres received from LATOME contain 16 trigger towers covering 0.4×0.4 in $(\eta \times \phi)$, 12 trigger towers covering 0.8×0.4 in $(\eta \times \phi)$, and 12 trigger towers covering 1.7×0.4 in $(\eta \times \phi)$ for $|\eta| < 2.4$, $2.4 < |\eta| < 3.2$, and $|\eta| > 3.2$, respectively. For the special case of the HEC-TileCal overlap, 4 additional fibres per side of the detector are sent each containing 16 trigger towers from a strip covering 0.1×1.6 in $(\eta \times \phi)$.

Overall, the jFEX system receives 128 and 392 unique fibres and 384 and 712 total fibres from TREX and LATOME, respectively.

A comparison of the core region of a jFEX module covering the central detector region, the splitting between the individual FPGAs, and the size of the environment received by a single FPGA is represented in figure 6.3. The core regions are shown as blue boxes with the yellow background representing the environment received by FPGA 2. This visualises the need for data duplication as the environment is significantly larger than the core area.

The Run 4 ATLAS Trigger System described in the Phase-II Technical Design Report [49] introduces large differences compared to the Run 3 system. Yet, the requirements for subsystems that have already been in development while the Phase-II TDR was written and are expected to be still in operation during Phase-II, like the jFEX, remain mainly unchanged. The available finer granularity of input data from the FCal is intended to be covered by the fFEX system. Other adaptions required during the transition to the Phase-II Trigger System, like new trigger algorithms, are expected to be handled in the firmware.

6.2 Real-Time Output

The real-time output of the jFEX system consists of TOBs and global values sent to L1Topo on optical links routed through the front panel. Non-real-time output, such as read-out and other information, e.g. for debugging purpose, is described in section 6.5. The exact format of TOBs and global values has not been conclusively defined within the scope of this thesis as they depend on the exact combination of FEX, Muon and L1Topo algorithms, which require detailed studies to determine the best use of the available bandwidth. To reduce the required latency, each processor FPGA has separate links to L1Topo even though this reduces the possibility to merge results and therefore does not allow for the most efficient options for data transmission. Draft data formats for regular jet TOBs, fat jet TOBs, and τ TOBs are shown in figure 6.4.

---

[3] In the central region, LATOME processes data at a finer granularity than the jFEX input to provide the required resolution to the eFEX system.

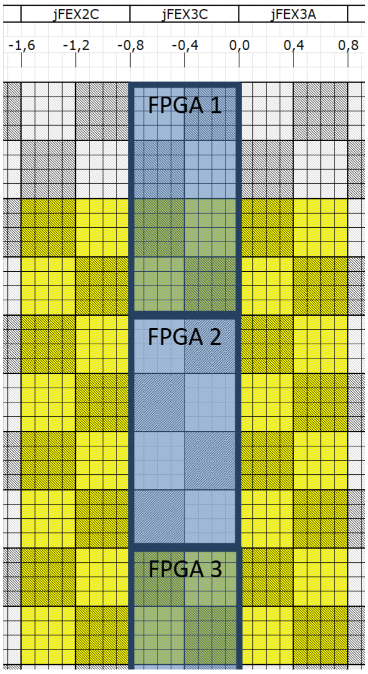

FIGURE 6.3: Size comparison between the core region of individual FPGA (blue boxes) compared to the environment received by FPGA 2 for a jFEX module covering the central detector region. The overlap with neighbouring FPGAs results in onboard data duplication and the extension beyond the core area implies the duplication of data before entering the jFEX system. Data from the corner regions are duplicated both by the input stage and on board. The same representation of granularity in η and ϕ is used as in figure 6.2.

The regular or small jet TOB consists of 5 bits η and 6 bits ϕ local coordinates, two energy sums, each 13 bits wide, and 3 bits Trigger Object Type (TOT) to identify the object as a jet TOB. The local coordinates are sufficient to precisely identify a position within the core area covered by a single FPGA and can be reconstructed to global coordinates on L1Topo where the information of the source FPGA for each incoming link is available once the fibre cabling is finalised. Two energy sums, as the results of different algorithms, are combined to a single TOB to make the data transmission more efficient for regular jets, which are assumed to be the most common TOB type in the jFEX system. This scheme requires both algorithms to use the same method to calculate the position of a jet candidate and therefore to always identify jets with the exact same coordinates. Once a possible jet position is found, the algorithms can differ e.g. in window size, shape or applied tower weighting. The same approach is used in the JEP system where energies from 0.4×0.4 $[\eta \times \phi]$ and 0.8×0.8 $[\eta \times \phi]$ are sent per TOB and can be used for different algorithms in the Topological Processor or trigger decisions in the CTP.

The fat jet TOB uses the same local coordinate format of 5 bits η, 6 bits ϕ, and 3 bits for TOT as the regular jet TOB does. Since the energy is expected to be higher, 15 bits are assigned to the energy sum. 11 bits are described as spare in this first draft that can later be used e.g. for substructure information.

The τ TOB format also uses the conventions of 5 bits η, 6 bits ϕ, and 3 bits TOT. The energy is coded on 13 bits with another 13 spare bits. The decision, how spare bits of τ objects can be used for further information like isolation or energy ratios, requires detailed studies.

The draft of jFEX TOBs used in this thesis proposes to use the same size of 40 bits for all

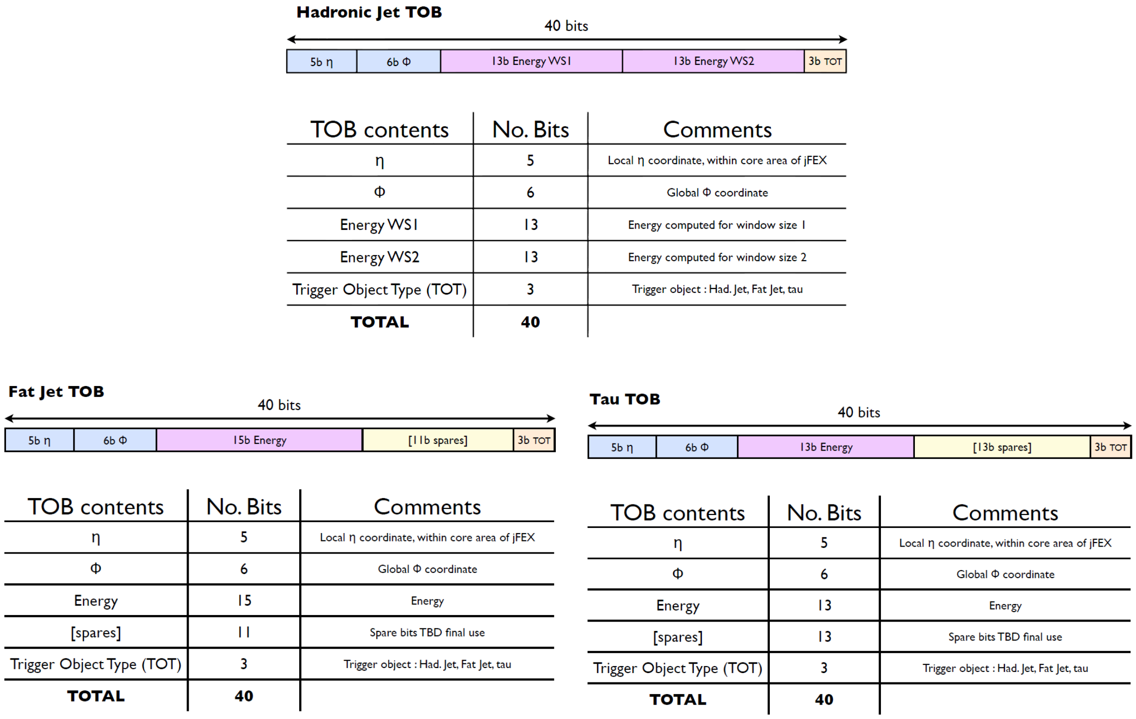

Hadronic Jet TOB

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy WS1	13	Energy computed for window size 1
Energy WS2	13	Energy computed for window size 2
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

Fat Jet TOB

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy	15	Energy
[spares]	11	Spare bits TBD final use
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

Tau TOB

TOB contents	No. Bits	Comments
η	5	Local η coordinate, within core area of jFEX
Φ	6	Global Φ coordinate
Energy	13	Energy
[spares]	13	Spare bits TBD final use
Trigger Object Type (TOT)	3	Trigger object : Had. Jet, Fat Jet, tau
TOTAL	**40**	

FIGURE 6.4: Draft data formats for jet, fat jet, and τ TOBs.

types. This approach was chosen to simplify the assignment of TOBs in the incoming data stream for the Topological Processor. In events where the number of found TOBs exceeds the available bandwidth, prioritising between TOBs will be required to determine which objects can be sent on the real-time data path. Further options regarding the TOB formats are discussed in section 8.3.

Global values are sent separately for E_T, E_T^X, and E_T^Y as 13-bit energy values.

As for the input requirements, the presented requirements on output capabilities for the jFEX are similar in the descriptions from the Phase-II TDR [49]. As system specifications can be altered due to improved capabilities of detailed designs during the actual development of the Run 4 Trigger System, it is considered beneficial to include spare output capabilities, if this does not interfere with the functionalities required for Run 3, to provide flexibility for the use as a part of the Phase-II Trigger System.

6.3 Jet Algorithm

The initial design for the jFEX from the TDR lists algorithms for jets, large-area τ, $E_{\mathrm{T}}^{\mathrm{miss}}$, and total E_{T}. For the jet algorithms, multiple versions are available with jet sizes between 0.9×0.9 and 1.7×1.7 $[\eta \times \phi]$. Out of all algorithms mentioned, the jet algorithms were expected to be most challenging in implementation and resource consumption and therefore constitute the biggest implicit requirement for the jFEX design.

The jet algorithm specified for the jFEX in the TDR is, like the JEP implementation described in section 4.4.1.1, based on the sliding window algorithm. Centred around a single trigger tower, a virtual window is defined within which all included trigger towers are summed up to calculate the E_{T} of the jet. In order to calculate all possible jets in an event, every trigger tower is used as the central trigger tower in the calculation[4]. This 'sliding' of the calculation window over the entire detector probably inspired the naming of this algorithm. If the granularity of input cells in the outer detector regions does not allow forming windows with the exact sizes, an approximation is used and cells that reach over the intended region can optionally be prescaled by factors of 1/2 using a simple bit shift.

The jet algorithm additionally requires the use of a smaller window of 0.3×0.3 in $\eta \times \phi$. Based on this so-called seed, jet candidates are defined by the requirement that the seed of a trigger tower is larger than those of trigger towers with an overlapping seed. This means that a comparison is done against nearest and next-to-nearest neighbours covering an effective area of 5×5 trigger towers for the central region. To avoid an efficiency loss in case of a stalemate, an approach similar to the JEM solutions, using a combination of greater ($>$) and greater-or-equal ($\geq$) relations according to

$$condition = \begin{cases} \geq & if \quad (\Delta\eta + \Delta\phi < 0) \vee [(\Delta\eta + \Delta\phi = 0) \wedge (\Delta\eta < 0)] \\ > & else \end{cases},$$

is used. Here, $\Delta\eta$ and $\Delta\phi$ refer to the difference between the central trigger tower of the current jet window compared to the positions of its neighbours. For the calculation of $\Delta\phi$, the values of ϕ can exceed the specified range for the ATLAS detector to avoid irregularities.

If a jet candidate is identified, its transverse energy is compared to a predefined threshold. A jet is found if a candidate exceeds the threshold and will be forwarded by the jFEX module to the topological processor. Please note that this threshold is different from the trigger threshold described in section 9.1.3. For the coordinates of an identified jet in η and ϕ, the position of the corresponding central trigger tower is used.

Due to the tight latency requirements and the pipelined nature of the data processing in the ATLAS Level-1 Trigger System, all window positions need to be calculated concurrently as sets of detector data, corresponding to an entire event, are received at an input rate of about 40 MHz. Therefore, each step in the calculation may only take about 25 ns as afterwards the resources need to be available for the next set of data.

As mentioned in section 6.1, the coverage of the ATLAS detector is split between multiple FPGAs each covering a disjunct part as their core region. Each FPGA uses only the trigger towers within the core region as central towers for the sliding window algorithm. A schematic representation of the sizes of the different jet definitions compared to the core region and environment is shown in figure 6.5. The size of the largest jet window is the basis for the specification which data each individual FPGA has to receive. The largest defined jet size of 1.7×1.7 $[\eta \times \phi]$ leads to the previously mentioned environment of 0.8 in both η and ϕ to be fully enclosed within the environment for every central trigger tower in the core region of an FPGA.

A variation of the sliding window algorithm suggested in the TDR uses the weighting of trigger towers prior to their summation based on the work of Yue-Shi Lai and Brian Cole [51]. As the applied weight decreases as a Gaussian function with the distance from the central trigger tower, this approach is also referred to as Gaussian weighting or Gaussian filtering. This variation is further discussed in section 9.2.1.

---

[4] Depending on the implementation of the algorithm, trigger towers in the forward detector region may not be used as central trigger towers if the calculation window would expand beyond the coverage of the detector as physics objects located in these positions cannot be reliably reconstructed.

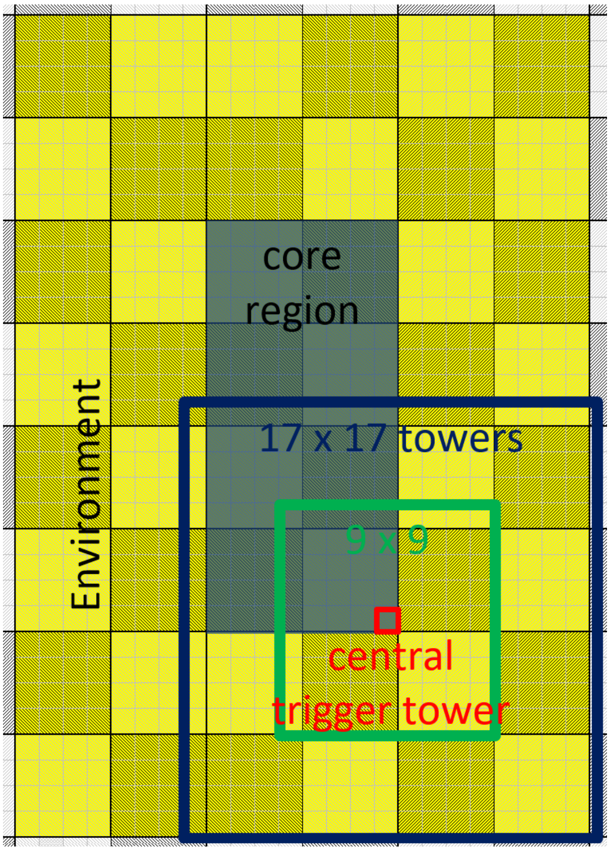

FIGURE 6.5: Graphical representation of the sizes of jets calculated by the sliding window algorithm in the jFEX compared to the core region and environment for an FPGA covering the central detector region. Each box represents a trigger tower after summation of overlapping detector layers. The jet sizes of 9 × 9 shown in green and 17 × 17 trigger towers (blue) illustrate the calculations for small R and large R jets, respectively. This scheme uses the same dimensions as figure 6.3.

6.4 ATCA

Most of the ATLAS off-detector electronics used in Run 2 are based on the VME64x standard. This standard has been used for more than 30 years and has reached its limits in several domains such as power availability and distribution, cooling capability, and bandwidth between modules for certain applications. Therefore, a working group was established to identify a new suitable standard for the next generations of trigger hardware for ATLAS. The Advanced Telecommunications Computing Architecture (ATCA) [52] was recommended by the ATLAS TDAQ group due to many advantages such as available space for signal routing and placement of components, power supply, and cooling capabilities, over other modern systems like μTCA, and was therefore chosen for all newly developed Trigger Systems starting with the Phase-I upgrade. Following this recommendation, the jFEX system is designed as a Front Board based on the PICMG® standard 3.0 Revision 3.0. A shortened version of the specifications is made available by CERN [53].

All jFEX modules are hosted in a single ATCA shelf in a dual star topology with 2 Hub modules in Logical Slots[5] 1 and 2. The Hub modules are centrally developed for the L1Calo Trigger System and are identical for all feature extractors. Their main functionality is the aggregation of read-out data and the distribution of control information to the FEX modules. For the electrical connectivity between Hubs and FEXs, standard connectivity in commercially available ATCA crates can be used.

Control functions as defined by the ATCA specifications are also used and partially included in the ATLAS requirements on module control functionalities. An overview of the predefined backplane connectivity is given in figure 6.6. Please note that this figure is focusing on the Hub modules and therefore includes functionalities that are not directly linked to the jFEX and are therefore not described within the scope of this thesis.

Zone 1 is reserved for power connection and shelf management, including dual redundant -48 V DC power input and hardware addressing, which allows each module to identify its position in a crate. In systems with multiple crates, the hardware address can also be used to distinguish between different crates. Power management functionalities used e.g. to control the maximum allowed power consumption of a module depending on its current state and adjust the cooling of the ATCA crate accordingly, are implemented using the Intelligent Platform Management Interface (IPMI) [54].

Zone 2 is used for the Data Transport Interface. Up to 5 ZD[6] connectors can be used for electrical connections within the crate. It complies with the official specifications to only assemble a subset of these connectors if the connectivity on some channels is not required. Since no communication between the separate modules of the jFEX system is intended, only a single ZD connector, indicated as J23 in the PICMG specifications, is used. Zone 2 links are grouped into Base Channels and Fabric Channels. The 4 Base Channels per Hub are used to control proper operation of each module by setting operation modes and monitoring parameters like temperature and power consumption. The 4 links are separated into 2 unidirectional links and 2 bidirectional links to support different ethernet protocols with up to 1 Gb/s. The 8 Fabric Channels, usually operated at a higher speed, are used to transmit clock signals and data required for communication functionalities as described in section 6.5.

The mechanical dimensions and connections described in the PICMG specifications are essential to build a Front Board that can be used in a standard ATCA shelf. Viewed from the Front Panel, the board is 322.25 mm high and 280 mm deep. The width per slot is defined as 6 HP[7] or 30.48 mm. The maximal height of components on Component Side 1 is 21.33 mm, leaving 6.61 mm for the PCB and Component Side 2 and a 2.54 mm pitch for electrical clearance to the next Front Board. This defined height poses a considerable limitation in the choice of components to be mounted on the jFEX. Yet, with only 6 modules in Node Slots, it is possible to extend the height of each module to 'double slot width' of 12 HP. The extended width allows for higher heat sinks while still having more unblocked room for airflow, fibre routing, and mezzanine boards. The standard ATCA form factor is shown in figure 6.7a. As mandatory features, the board includes a Front Panel to direct the airflow, two handles, and several guidance pins to ensure a stable position in the crate including reliable electrical connections on the backplane. The exact positions for the three Backplane Zones are defined in the standard to ensure compatibility.

---

[5] Referring to positions within a shelf to comply with the standardised ATCA dual star topology.

[6] A connector type designed for pressfit assembly. Exact requirements are described in the ATCA specifications.

[7] Horizontal pitch is a unit of length commonly used to measure the horizontal width of rack-mounted electronical equipment. 1 HP is 0.2 inches or 5.08 mm wide.

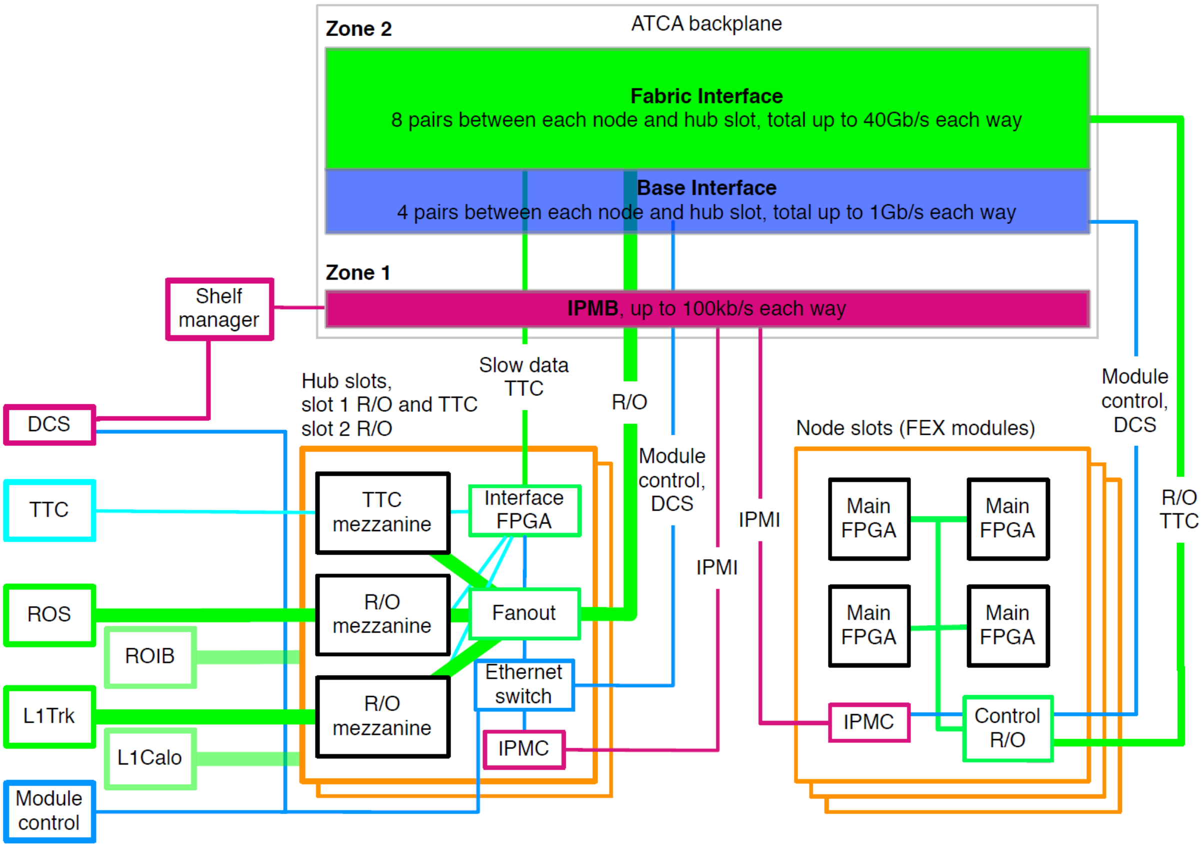

FIGURE 6.6: Control functions via backplane connections for the jFEX as managed by the Hub modules. JFEX modules are represented by the orange rectangles described as Node Slots. Focus points for the jFEX are module control, DCS, read-out, TTC, and IPMI connections. The remaining functionality of the Hub modules is not further discussed in the scope of this thesis. Figure taken from the Phase-I TDR [48].

Within the L1Calo community, it was agreed upon to modify the standard form factor by extending the main board towards Zone 3, as shown in figure 6.7b, where the custom built Rear Transition Modules (RTMs) are used. With this modification, the overall system still stays within the specified dimensions to fit in a regular ATCA shelf while it allows to mount the connectors for Zone 3 on the extended part of the board, leaving more space for electrical components and routing.

While the exact connector types for Zones 1 and 2 are defined in the general specifications, there is a broad variety of possible connectors for Zone 3. Blind mate MTP[8] connectors were chosen for this region for all Phase-I modules. This type offers the advantage that these connectors align and connect automatically when the Front Borad is inserted into the crate. Connectors of this type were commercially available with up to 72 individual fibres prior to the jFEX design phase, allowing for a high input bandwidth to each jFEX module.

---

[8] Multi-fiber Termination Push-on. Connector type for multi-fibre ribbons. Compatible with Multiple-fiber Push-On (MPO) connectors.

6.5 Control and Read-out

The main communication outside the real-time data path between the jFEX and the remaining trigger system is done centrally through the Hub modules providing various control functionalities to the jFEX. The Hub contains several daughter cards which can be replaced to provide flexibility for required changes when upgrading from the Phase-I system to Phase-II. This allows for the jFEX hardware to be compatible with future control functionalities and handle upgrades by changes in the firmware and the Hub module. The only direct communication for a jFEX module is an IPbus implemented via Ethernet on a control mezzanine which is further described in chapter 8.

The Base Channels connected to Hub 1 are used for module control via IPbus including the collection of information for the ATLAS Detector Control System (DCS). The Base Channel interface to Hub 2 is connected to the IPMI Management Controller (IPMC) card[9] to provide power management capabilities as required by the ATCA specifications. One Fabric Channel per Hub module is used to receive the Trigger, Timing & Control (TTC) data. This stream includes the L1A signal to trigger read-out data and other control functionalities. Another link delivers a clock signal to keep the trigger system synchronous to the LHC bunch crossings. The remaining 6 links are used to send read-out data to the Hub modules. As read-out data is sent only upon receipt of a dedicated signal on the non-real-time data path, the data are temporarily stored in the FPGAs.

The original design of the jFEX to Hub/ROD read-out interface was defined under the assumption of 6 FPGAs per jFEX module. This required to redefine the direction of 2 links in ATCA Zone 2 with respect to the PICMG specifications. Since this redefinition has no impact on the backplanes, connectors within an ATCA crate, and the Hub module, that only routes these signals to the ROD mezzanine, this change only needs to be considered on the ROD cards. Because this was already adapted by the remaining system, the jFEX design still uses all 6 available links, equipping two of the FPGAs each with a second link. This additional connection can be used to increase the total bandwidth if required by the read-out system. Since on average an even distribution of read-out relevant data over all FPGAs is expected, inter-FPGA connectivity would be required to distribute these data to match the available output bandwidth per FPGA.

---

[9] CERN-developed mezzanine card used to provide the ATCA-specified power management functionalities to most modules used in the ATLAS Level-1 trigger.

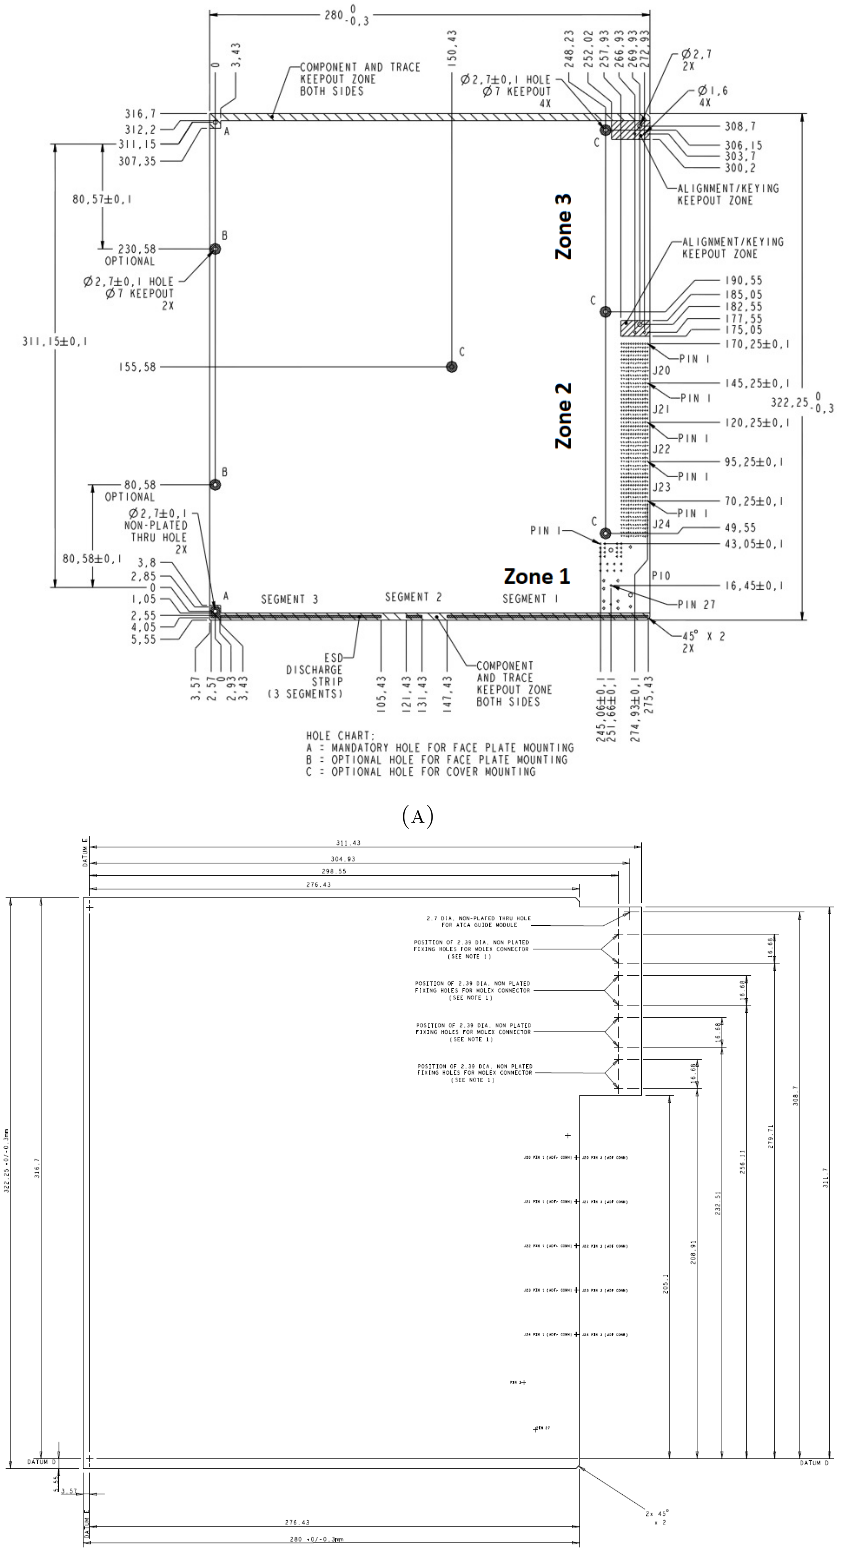

(A)

(B)

Figure 6.7: Standard ATCA Front Board form factor including detailed positions of connectors and other mandatory pins with added indications for the backplane zones (A). Figure taken from the ATCA specifications [53] and modified. Modification of the ATCA form factor (B) by extending the Front Board in Zone 3 towards the RTM [55].

Chapter 7

Hardware Basics

In this section, suitable types of hardware are discussed that meet the requirements on interfaces, as described in chapter 6, while providing the best possibilities to run the ATLAS trigger algorithms. The key components for this are the central devices responsible for the real-time data processing for which there are multiple types that can be considered with each of them offering certain advantages and disadvantages. Available types would be Central Processing Units (CPUs), Graphics Processing Units (GPUs), Field-Programmable Gate Arrays (FPGAs), or Application-Specific Integrated Circuits (ASICs). Though chapter 6 already describes the use of FPGAs as the devices used for data processing in the jFEX, this decision is briefly explained in the paragraphs below.

Systems based on CPUs or GPUs offer more flexibility and easier programming compared to the other options. CPUs are designed to perform tasks sequentially at a very high speed. Even with these higher rates per task, the overall processing power is not sufficient for the required parallel processing of trigger algorithms in a pipelined system. GPUs are designed for parallel processing, typically using more processing units at a lower frequency compared to CPUs. Yet both are more complex systems where a control unit supervises the system, distributing tasks to specialised units within the system. This increases the latency and for most systems makes the latency input-dependent. Additionally, it is typically not an intended mode of operation to drive CPUs or GPUs using an external clock signal which would be required for synchronous operation to the LHC bunch crossing. Therefore, both systems are not suitable to be used in the ATLAS Level-1 trigger. For the HLT on the other hand, a predictable latency per event and clock synchronous operation are not required, making these processing units a suitable choice for the last trigger stage.

FPGAs are designed for massive parallel processing of input data at a low, fixed, and deterministic latency. These devices can be (re-)programmed after manufacturing by configuring the programmable connections between digital logic circuits. ASICs, as the name suggests, are Integrated Circuits (ICs) that are similar to FPGAs in their basic operation but are customised for a specific use case without the option of reprogramming. While ASICs can excel FPGAs in performance and power efficiency, their development requires the final specifications to be available, which is so far not possible for the jFEX, where details for algorithms, interfaces, and requirements for Phase-II remain largely unspecified during the design phase. Additionally, it is an active design decision to keep the system flexible for future changes to the algorithms and interfaces if experience gathered during the operation or the physics programme require an adaption of the first trigger stage. Also, required development effort and manufacturing costs for a small number of ASICs is usually rather high compared to an implementation based on standardised FPGAs. Overall, FPGAs are the only suitable option to cover all the requirements for real-time data processing in the jFEX system. The basic functionality of an FPGA is described in section 7.1. Additionally, a short description of fibre optic transceiver and printed circuit boards which play an import role in the jFEX design is given in section 7.2 and section 7.3, respectively.

7.1 FPGA

To understand the possibilities and limitations of an FPGA-based system, knowledge of their basic structure is required. Therefore, a short description of the overall structure of an FPGA is given in section 7.1.1 followed by an introduction to clocking resources and high-speed transceivers in section 7.1.2 and section 7.1.3, respectively.
While basic operation principle and designs are similar for FPGAs from various manufacturers, there are differences in the detailed layout and naming schemes. Many FPGA series also include a variety of digital or analogue features that are not required in the jFEX system and are therefore not described. The jFEX design, that will be presented in section 8, is based on Xilinx® UltraScale™[56] and UltraScale+ [57] FPGAs. Therefore, examples and terminology from this manufacturer and series are used throughout this section. More detailed descriptions of the FPGA features are provided by Xilinx and publicly accessible on their homepage.
Very High-Speed Integrated Circuit Hardware Description Language (VHDL) [58] is one of the hardware description languages[1] that can be used to implement functionality in an FPGA. Xilinx provides a software suite called Vivado Design Suite for synthesis and analysis of designs. This tool can optimise the implementation of a design within the scope of the textual description, to make most efficient use of the targeted hardware. The exact implementation is done by the design tool while some of the architectural details of the FPGA are not known to the user. As the final implementation of the jFEX firmware lies beyond the scope of this thesis, the structure of hardware description languages and functionality of the required tools are not discussed further.

7.1.1 Fabric

The main structure of an FPGA consists of an array of Configurable Logic Blocks (CLB), Input/Output pads (I/O), and routing channels which can be connected using switching matrices as shown in figure 7.1. Large FPGAs can contain several millions of CLBs [59]. While CLBs can be used for general-purpose combinatorial and sequential circuits, there are so-called hard blocks which are specialised blocks optimised for more efficient implementation of certain tasks. Examples are digital signal processing (DSP) slices, designed for multiplication of longer bit sequences, and memory blocks.

---

[1] In electronics, a Hardware Description Language is a specialised language used for textual descriptions of the structure and behaviour of electronic circuits, allowing for its automated analysis, simulation, and synthesis. The result can then be converted into an IC, like an ASIC, or implemented in a suitable FPGA.

CLBs consist of logic elements such as Look-Up Tables (LUT), flip-flops[2] as 1-bit registers, multiplexers and carry logic for the cascaded use of logic blocks in more complex arithmetic functions. In the UltraScale architecture, eight LUTs including their storage elements are grouped in a slice. A schematic representation of a LUT with its storage elements is shown in figure 7.2. Look-up tables, which provide 6 inputs and 1 output for the UltraScale series, represent truth tables where for every possible input combination the result of the required algorithm is stored in a mask. Algorithms that require more than 6 inputs or multiple outputs can be implemented by combining several LUTs. The UltraScale architecture also allows to use one input as an output, implementing two functions with the same 5 inputs in a single LUT.

The Input/Output pads of an FPGA are organised in banks each containing 52 connections to the outside of the FPGA. These can be used for single-ended transmissions or for up to 24 differential pairs. The challenges and advantages of differential signal routing are discussed in section 7.3. If operated at a transmission rate around the clock frequency used in the FPGA fabric, these links can be used for low latency connections. Increased line rates in the low Gb/s range are possible but this use case requires additional serialisation/deserialisation which increases the latency [62, 63]. Significantly higher transmission rates can be achieved by so-called multi gigabit transceivers (MGT) which are described in section 7.1.3.

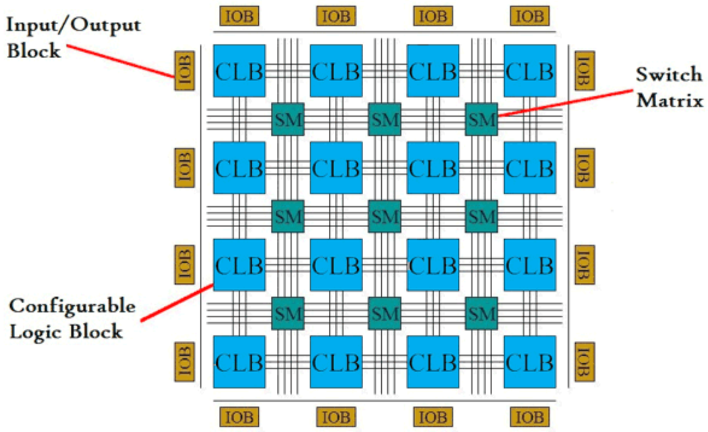

FIGURE 7.1: Basic structure of an FPGA containing Configurable Logic blocks, switching matrices, and Input/Output pads. Figure taken from reference [60].

In the UltraScale architecture, two different types of I/O banks are present. High Performance (HP) banks are designed to provide higher transmission rates while High Range (HR) banks can be configured to support a wider range of signalling standards.

Some of the larger FPGAs available on the market, including some of those suitable for the jFEX design, are built in a non-monolithic way consisting of multiple dies which are produced separately and connected via a silicon interconnect [64]. In the Xilinx terminology, the separate dies are called Super Logic Regions (SLR) connected by a Stacked Silicon Interconnect (SSI). Calculations that require large parts of the logic resources of a non-monolithic FPGA will be spread over multiple SLRs. This might impose additional requirements for the implementation of the algorithms e.g. due to a limited bandwidth between the SLRs or an increased latency on these connections. As during the design phase no experience or external information regarding

---

[2] An analogue circuit with two stable states that can be used to store state information.

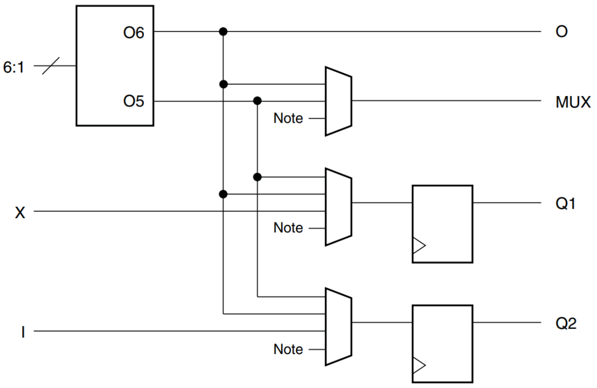

FIGURE 7.2: Schematic representation of a 6-input Look-Up Table with a signal propagation from left to right. The LUT is displayed as the rectangle in the upper left corner. The trapeziums represent multiplexers to choose between the output of the LUT or the direct inputs (X,I) to be routed into the storage elements displayed as the rectangle in the bottom right corner. The 'note' input to the multiplexers represent additional input e.g. for carry logic that can be used to implement arithmetic functions like adders using more than one LUT. Available outputs are a direct output (O), a multiplexer (MUX) merging all LUTs from a logic slice, and delayed outputs via the storage elements (Q1, Q2). Figure taken from the configurable logic block user guide [61].

the limitation of the SSI technology in the UltraScale series was available, this aspect was not taken into further account for the jFEX design.

7.1.2 Clocking

In most applications, FPGAs are operated synchronously driven by an external clock signal. During a clock tick, signals propagate within the FPGA possibly altering the outcome multiple times before converging to the final result. Each step of calculations and required connections adds latency of typically a few nanoseconds, limiting the maximum number of steps that can be included in a single clock cycle. If a design has too many steps within a single tick or the interconnects between utilised logic resources are too long, the design suite used to implement the code is supposed to spot a so-called timing violation and indicate the requirement to restructure the calculation. How many steps can be implemented in a clock cycle depends on the clock frequency which therefore needs to be specified in the design suite. Running the implemented design at a higher clock frequency than the one specified in the design suite can result in nondeterministic outputs.

There are multiple types of clocks present in large FPGAs. The main fabric is driven by one or multiple global clocks which are typically provided by an outside source. Some parts, like high-speed transceivers, require dedicated external reference clocks independent from the global clocks.

Asynchronous operation of FPGAs is possible by simply cascading logic elements without the use of clock-driven storage elements. While this mode potentially decreases the latency of calculations by removing the requirement to wait for a clock signal, it impedes predicting when to expect a final result on any of the outputs especially for rapidly changing inputs. Also, asynchronous systems with calculations that take longer than the time between switching of inputs might not converge to the final result at all. While this operation is theoretically possible, it is rather uncommon, discouraged by many manufacturers, and not intended to be used in any part of the jFEX system.

7.1.3 MGT

Multi-Gigabit transceivers (MGT) are specialised hard blocks that serve as an addition to the regular I/O pads when higher transmission rates are required. In the Virtex UltraScale architecture, two different types of MGTs are available, denoted as GTH and GTY. The main differences between the two types are the increased transmission rate for the GTY of 30.5 Gb/s instead of 16.3 Gb/s for the GTH [59] which comes at the cost of increased power consumption even at the same bit rate [65, 66]. In Xilinx architectures, all MGTs are organised in quads which share certain resources. Each quad has two pin pairs for external clock signals to be used as reference clocks and an additional option to forward the clocking signal to and from neighbouring quads withing the same SLR [67, 68].

A schematic representation of an MGT is shown in figure 7.3. The intended data paths for reception and transmission of data are shown in teal and red, respectively. A feature offered by the Xilinx MGTs is the so-called PMA loopback which provides the option to exchange duplicates of the processed data between the receiving and transmitting side of an MGT in both directions. With this functionality, the MGT can either receive a copy of its transmitted data via the near-end PMA loopback or retransmit an analogue copy of an incoming bit stream via far-end PMA loopback. The latter option is intended to be used in the jFEX with the additional data path shown in green. The use of this operation mode leaves the receiving side unchanged while the right-hand side of the transmission data path between the TX interface and the multiplexer is replaced by the PMA loopback signal. Far-end PMA loopback provides the system with a low-latency copy of a high-speed signal without the use of additional electronic components. A disadvantage of this mode is the expected increase of noise on the retransmitted signal as the analogue copy of the input also includes potential interferences. The impact of the PMA loopback scheme on the quality of received high-speed signals will therefore be verified as a part of the jFEX design verification.

Further components in the MGT hard block that are relevant for the jFEX design are the equaliser (EQ) to improve incoming signal quality, the polarity control allowing the swap of differential signals, the generation (Pattern Generator) and verification (PRBS Checker) of pseudorandom binary sequences, comma detection and alignment to synchronize the receiving FPGA with the correct start of data words in the incoming binary stream, and the native support of 8b/10b encoding. As decided by the ATLAS L1Calo group, all high-speed links to other components in the trigger system use 8b/10b encoded data words. This scheme sends symbols containing 10 bits instead of 8-bit data words, producing an overhead of 25 %. The advantage of this scheme is the achieved DC-balance and ability to recover clock signals and data word boundaries from the arriving data stream. If multiple options are available, the required functionality can be chosen via the configuration of a multiplexer shown as a trapezoid.

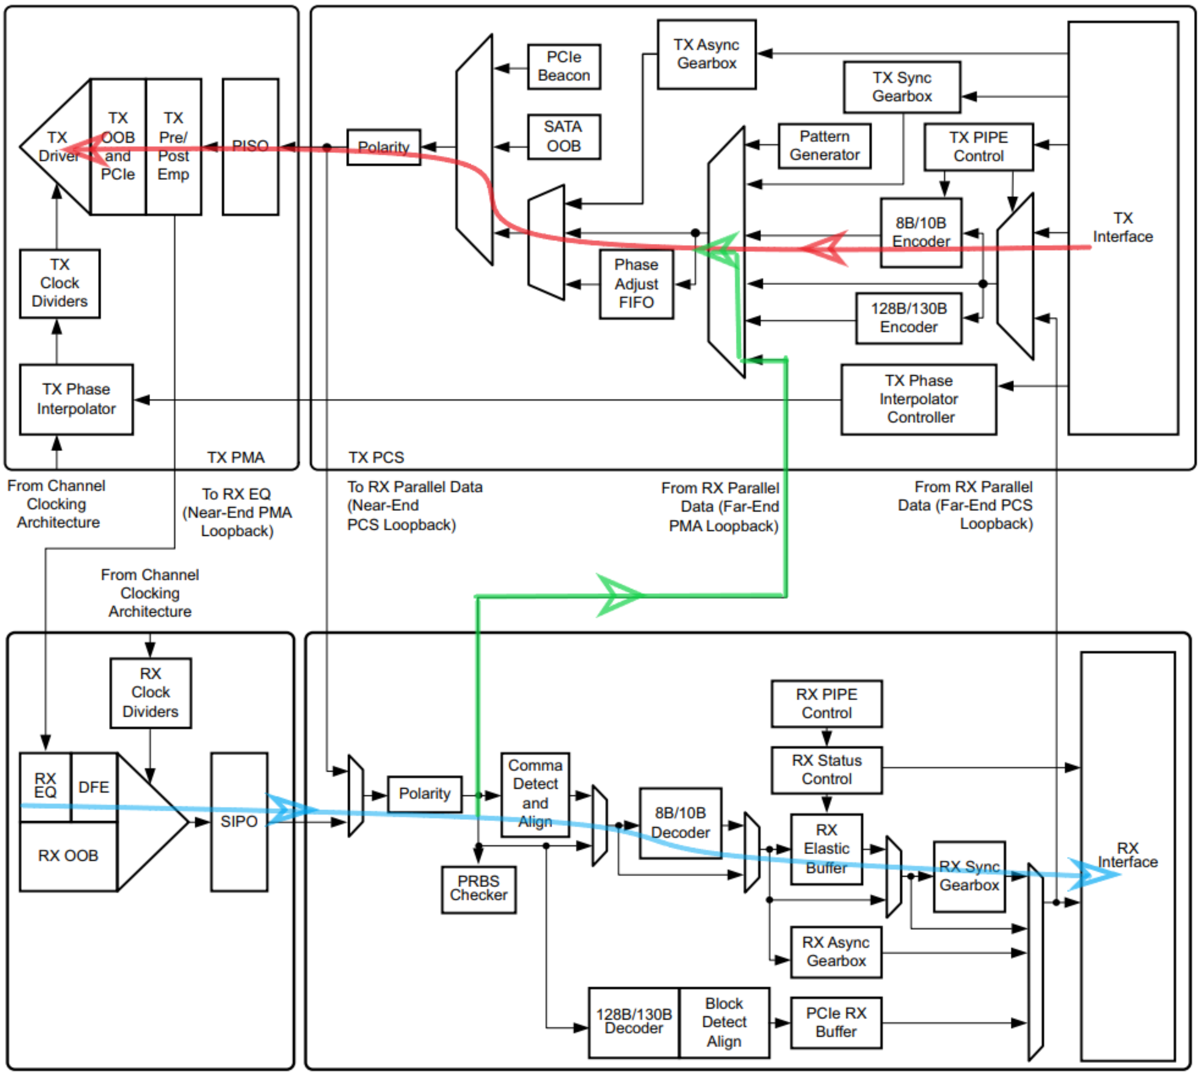

Figure 7.3: Schematic representation of an UltraScale GTH Multi-Gigabit Transceiver. Highlighted are the data paths for the reception (teal), transmission (red), and the far-end PMA loopback feature (green). Figure taken from the GTH transceivers user guide [67].

7.2 Fibre Optic Transceiver

For the transition between the fibre optical input and output signals that connect the jFEX to the real-time data path of the remaining trigger system and the eletronic devices, opto-electronic converters are required. These devices receive an optical signal and convert it into an electronic signal or vice versa. For the reception of optical signals, incoming light is converted into an electronic signal using semiconductors like photodiodes. The transmission of optical signals is done in a similar way based on light-emitting diodes. For the electronic side of the connection, sockets are soldered onto the printed circuit board to provide the required connectivity and power supply to the converting IC. The optical end of the transition is typically fitted with matching connectors to an optical fibre.

7.3 Printed Circuits Boards

A printed circuit board (PCB) is a laminated sandwich structure of conductive layers and insulating substrate designed and built to hold electronic components while also providing the required connectivity between them. Attachment of the electronic components is usually done by soldering them to the PCB either on conductive pads available on the surfaces or via through-hole mounting. On larger PCBs with multiple layers, the conductive layers are usually divided in signal layers, ground planes, and power planes.

The signal layers provide the connectivity between ICs to transfer information. This can be done using single-ended connections or using differential signalling. In differential signalling, two complementary signals are sent simultaneously on balanced lines of conductors. In the case of differential signalling on a PCB, the balanced lines consist of the conductive traces with the same length that are routed close to each other. As signal on the receiving side, the difference between the two signals is used, usually measured at a terminating resistor between both lines. This technique is used in many signalling standards and offers higher resilience against electromagnetic noise and therefore allows for increased data rates compared to single-ended transmission where only individual connections are used. This resilience is achieved as the traces on the PCB are routed close to each other which leads to a similar exposure to electromagnetic fields from external sources or other components and traces on the PCB which is automatically eliminated by using the difference between both lines. Disadvantages of differential signalling include the doubled number of pins on the ICs required per connection and increased complexity to the PCB design due to the mandatory parallel routing for the balanced lines.

The insulating layers of a PCB are used to separate the conductive layers and severely affect the electromagnetic behaviour of the conductive planes, especially for high-speed signals, depending on their dielectric characteristics. Another critical aspect in the PCB design is the matching of the impedance between the transmitter, transmission line, and receiver of high-speed signals in order to avoid reflections that would interfere with the data transmission. For source and sink of the signal, the impedance is given by the design of the IC or can be adjusted using terminating resistors. The impedance of the transmission line is given by the trace dimensions and the surrounding substrate. The selection of a suitable material for the non-conductive substrate is therefore an important design choice in the early stages of the PCB design process.

Chapter 8

jFEX Design

In this chapter, a design is presented that meets the requirements shown in chapter 6 while providing best possible opportunities for the implementation of trigger algorithms as required for the ATLAS physics programme. A model of the board is shown in figure 8.1. This chapter describes some of the design challenges and the approaches taken to identify suitable solutions. As PCB material, Megtron6 [69] was chosen which is a relatively newly developed material, designed specifically for high frequency circuit boards. A visual representation of the design described in this chapter as block diagrams is shown in appendix B. The final version of these block diagrams was created in October 2015 and was used as the main input for the development of schematic entries and board layout by Bruno Bauss. Details that are entirely based on the PICMG specifications, like board dimensions and connector standards, are not discussed again in this section.

The complete development of such a complex board is done in multiple stages from prototypes to the final production version to gain the opportunity to test design and layout and thereby identify possible faults. The schedule of this project included four iterations of the board: two prototypes, a pre-production, and the final production version. This thesis mainly focuses on the first prototype based on the original specifications from the Phase-I trigger upgrade [48] and only some of the later adjustments to the design are contained. A more recent version of the jFEX pre-production board was presented in July 2019 [50].

8.1 FPGA Choice

Since the concept of the system is centred around the FPGAs used for data processing, the FPGA choice is the most critical decision that needs to be taken at the beginning of the design phase. To meet the requirements on input bandwidth, a minimum of 104 high-speed receivers running at up to 12.8 Gb/s are required. Furthermore, it would be preferable if the devices had some additional links available for control functions or spare input for Phase-II. During the early design phase, two FPGA series were announced: Altera®/Intel®[1] Stratix® 10 TX and

---

[1] This decision was taken in 2013. Altera was acquired by Intel in 2015. Therefore, even though at the time of the decision it was called Altera, this series will be referred to as Intel Stratix 10 for the remainder of this chapter.

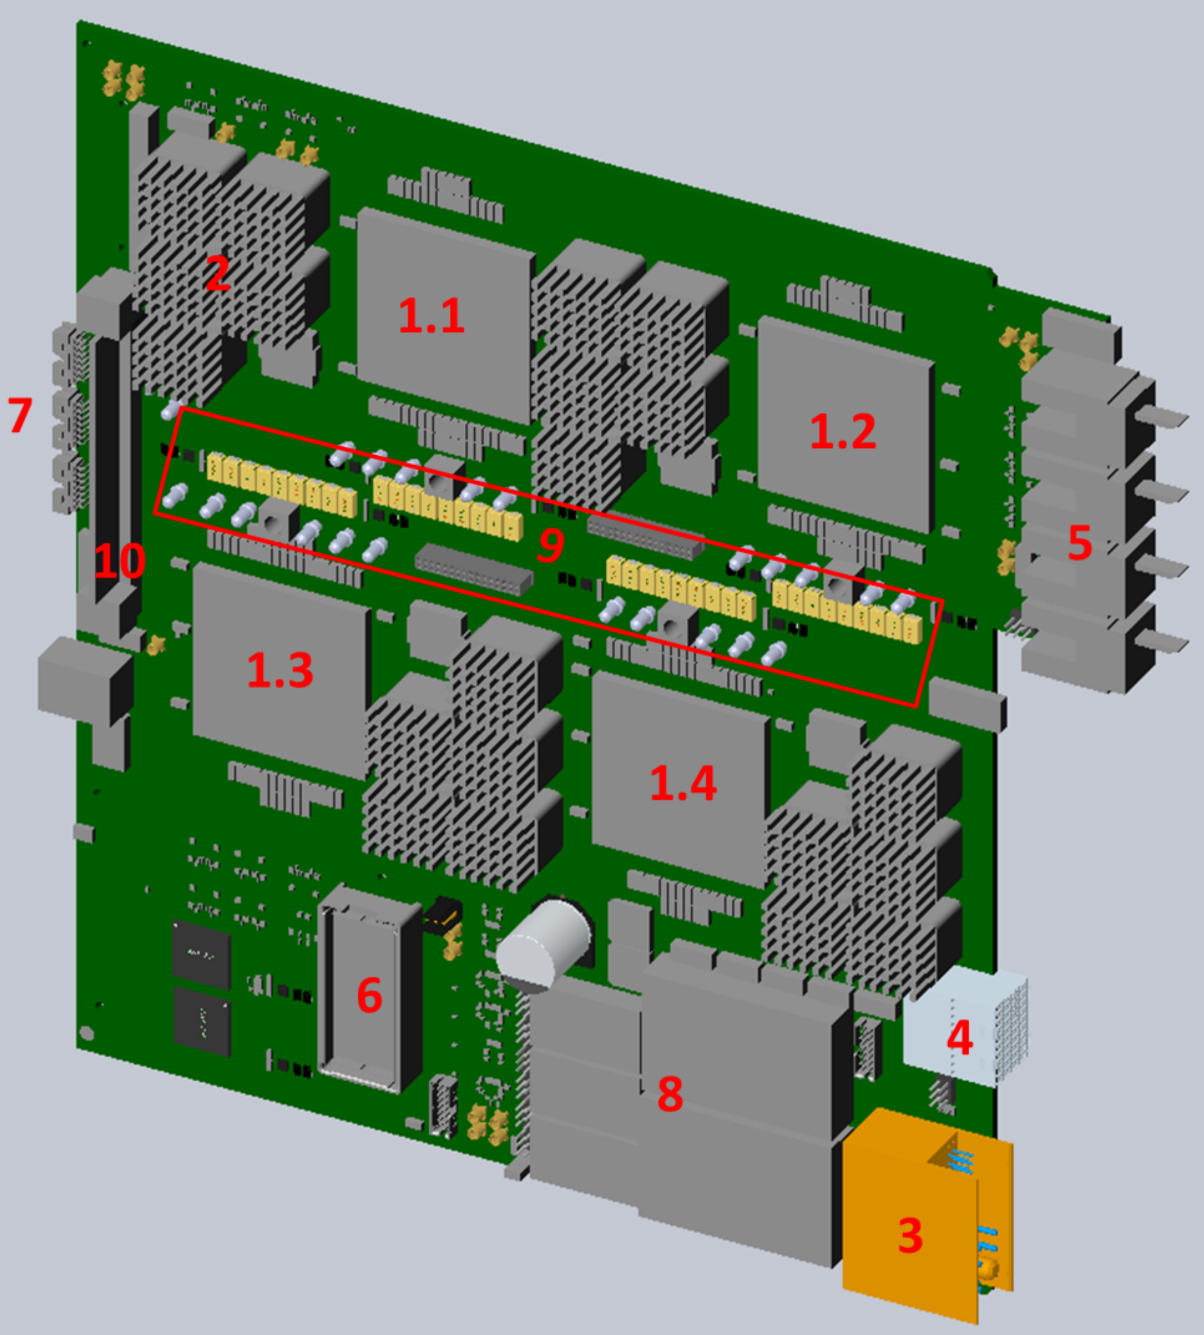

Figure 8.1: Model of the first prototype. The processing FPGAs are numbered from (1.1) to (1.4) for U1 to U4, each accompanied by five receiver MiniPODs and one transmitter, which is only shown as a socket, (2). The right-hand side shows the backplane with standard ATCA connectors in Zone 1 (3) and Zone 2 (4) and connectors for fibre ribbon handling (5). The connector for the extension mezzanine is shown on position (6), left of which the clocking circuitry is placed. Initial configuration of the FPGAs is possible via JTAG headers (7). DC/DC converters for board wide voltages are placed at (8), while the area around (9) is reserved for local converters for each FPGA, to be placed on a mezzanine. A socket for IPMC cards is placed at (10). The front panel including status LEDs is not shown. Board layout by Bruno Bauss.

Xilinx® Virtex® UltraScale™. Both series offer devices that meet the requirement on available bandwidth. An overview of the FPGA series from Intel and Xilinx is given in appendix C. From these series, a total of three FPGA types meet the requirements on input bandwidth: Intel Stratix 10 TX 2500, Intel Stratix 10 TX 2800 and Xilinx UltraScale XCVU190. Please note that the choice was made in the early design phase of the jFEX in 2013. Only devices that were available at that time or announced with a launch date sufficient for the prototype production were considered. More recent FPGAs series have therefore not been taken into account and their capabilities are not discussed in this thesis.

resource (Intel)	Stratix TX 2500	Stratix TX 2800	XCVU190	resource (Xilinx)
Logic elements	2,422,000	2,753,000	2,349,900	System Logic Cells
8-input LUTs	821,150	933,120	1,074,240	CLB LUTs
ALM registers	3,284,600	3,732,480	2,148,480	CLB Flip-Flops
18 × 19 multiplier	10,022	11,520	1,800	DSP Slices
user I/O	296	296	448	HP I/O
LVDS pairs	144	144	204	LVDS pairs
transceiver pairs	144	144	120	transceiver pairs

TABLE 8.1: Comparison of resources from different architectures that serve a similar purpose with a focus on the resource types most relevant for the jFEX design. Full tables are shown in appendix C. Numbers taken from [70] and [59] for Intel and Xilinx resources, respectively.

Table 8.1 shows an overview of the computing resources available in the considered FPGA types with a focus on the resource types most relevant for the jFEX design. A direct comparison of the possible performance between devices from different manufacturers is difficult since they are based on different architectures. Therefore, an exact prediction of the relative performance based only on the numbers of logic elements and other resources between the Intel FPGAs and the Xilinx devices is not possible. The most precise prediction could be made by implementing preliminary versions of algorithms on both types and comparing the relative resource consumption per device. However, with the available manpower on this project this was not possible, especially since Intel FPGAs have never been used in this group before and implementing a design using a completely unknown design suite requires additional effort for learning. Additionally, the development time for the board would most likely have been longer as the inclusion of unknown electronic devices also requires intensive studying of the corresponding documentations which can cover multiple thousands of pages. In contrast, for Xilinx FPGAs, some parts from old designs could be copied or were easier to adapt due to documentations available from the manufacturer that highlights differences between older devices and the latest series.

Most algorithms are not able to utilise every resource of the logic elements they require and leave spare capacities that can, in some cases, be used by other functions. The efficiency of an implementation depends heavily on the exact combination of algorithms and how they can be combined by the design suits to utilise the same logic blocks for multiple functions. This makes the efficiency of a design on a given architecture even more difficult to predict since it depends not only on a preliminary implementation but the final version. Additionally, no FPGAs from the suitable series were available for test implement as both had just been announced by the manufacturers. Therefore, an exact comparison of processing capabilities for the FPGAs could not be given during the design phase of the jFEX. Furthermore, a conflict of the timeline for the first prototype production and the availability of the Intel FPGAs was expected. Therefore, the decision to use the XCVU190 FPGAs was made to minimise the required engineering effort and reduce potential complications by using technology which is better known to the group and that is already closer to its release.

More precisely, the XCVU190-2FLGA2577E was chosen, where the extension describes the exact device package.

The extension '−2' describes a medium speed grade out of a scale from '−1' (slow) to '−3' (fast). Chips rated for higher frequencies are more expensive while the slowest version only supports MGT link speeds of up to 12.5 Gb/s. Therefore, '−2' is the cheapest version that meets the requirements on data transmission rate as described above.

'FLGA2577' combines a description of socket material, format, spacing (FLG), and a footprint in which 2577 pins are used. 'A' is a designator in case multiple packages with the same pin count are available. 'E' indicates the extended temperature grade, which classifies devices usable in a range from 0 °C to 100 °C [59].

8.2 Real-Time Data

Due to the large number of links running in parallel and the highest link speed of all onboard connections, the real-time input and output signals are considered the most crucial part of the design. Therefore, the placement of components on the board was optimised to keep those links as short as possible and these links were prioritised in routing, reserving the first two PCB layers for this purpose. A schematic representation of the input real-time data path is shown in figure 8.2. The scheme shows the main routing channels for the real-time input data consisting of the direct input from the opto-electrical receivers and the inter-FPGA communication. Crossings of routing channels are avoided to reduce the number of signal layers required without using transitions between signal lanes for any connection.

Input data is received via the backplane, as described in section 6, on four fibre ribbons, two 48-way and two 72-way bundles. They are directly split into four, respectively six bundles of 12 fibres, each ending in a PRIZM® connector. The PRIZM connector provides the required form factor for insertion into an Avago MiniPOD™ device which can be used for opto-electrical conversion of signals up to 14 Gb/s. This allows for an input bandwidth of up to ≈ 3 Tb/s per module if operated at the highest considered data rate of $12,8$ Gb/s.

Six uni-directional 12-channel MiniPODs per FPGA, five receivers and one transmitter[2], provide 240 input links and 48 output links per module. The overhead compared to the minimal required input links is used as spare for future detector upgrades or in case individual links might be non-functional.

A sketch of fibre routing channels, including the assignment of backplane connectors to certain MiniPOD positions, is shown in figure 8.3. This sketch is based on a preliminary layout of the board from an early stage in the development. Yet, since most of the available space on the PCB is already densely populated by the main components, no large movement of the components relevant for the fibre routing is possible during the layout phase.

The XCVU190 FPGA has two different types of MGTs, GTH and GTY. The GTY offer the option to use significantly higher data rates, which exceed the link speed options defined for

---

[2] A receiver MiniPOD is defined as the type that receives optical input and converts it to an electrical signal even though this makes the device a transmitter from the electronical point of view. A transmitter is correspondingly the type that receives an electrical signal and converts it to an optical.

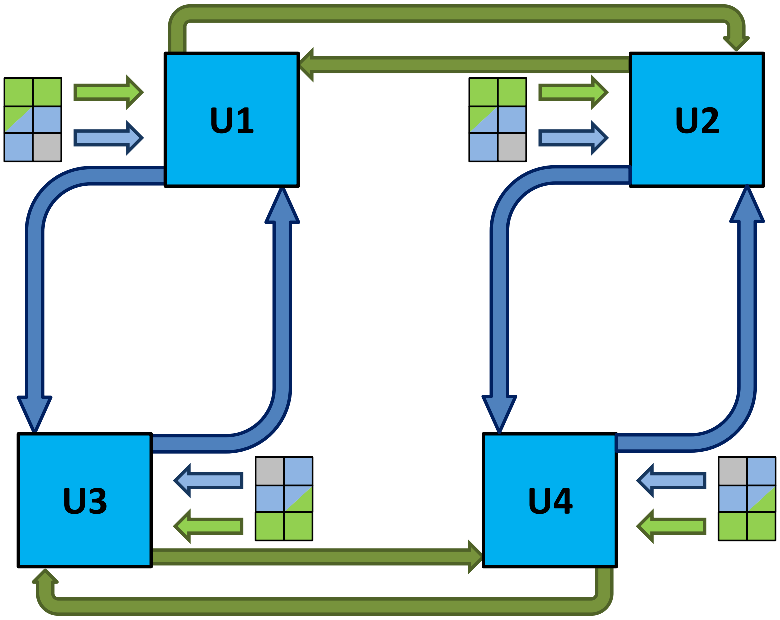

Figure 8.2: Arrangement of the four FPGA/MiniPOD blocks with schematic representation of the routing channels for PMA loopback. Direct connections from the MiniPODs and loopback channels are routed in separate layers of the PCB. Each arrow represents a bus of 29 differential signals. This scheme avoids crossing of buses and therefore allows the complete loopback channels to be routed in a single layer.

L1Calo, at the cost of higher power consumption[3]. Since this better performance is not required, the design tries to minimise the use of GTY transceivers. Therefore, the six MiniPODs for each FPGA are placed on the GTH side of the device, with all receiver MiniPODs connected to GTH channels. The utilisation of the corresponding MGT transmitters is described later in this section. With all links on this side being already in use, the transmitter MiniPOD therefore had to be connected to GTY channels, even though the higher bit rates are not supported by the MiniPOD and can therefore also not be used in Phase-II, if no faster, compatible devices are introduced by the manufacturer. Since careful planning of the connections between FPGA and MiniPODs is required, it is beneficial to define this as a block that can be designed once and reused for all FPGA positions on the board to reduce design work and guarantee an identical layout for all devices. The pre-routed block is shown in figure 8.4.
For each FPGA on the four jFEX modules covering the central part of the ATLAS detector

---

[3] In the chosen speed grade, GTH and GTY transceivers can be operated at 16.375 Gb/s and 28.21 Gb/s, respectively. For the full set of 60 channels per side operated as transceivers at 12.8 Gb/s, the Xilinx Power Estimator 2016.3 gives an expected power consumption of 15 W for GTH channels and 33 W for GTY channels. Due to the varied MGT architectures, the ratio between $\mathrm{MGTAV_{TT}}$ and $\mathrm{MGTAV_{CC}}$ is significantly different for the two types. The issue of power consumption and supplies in general is discussed in section 8.8.

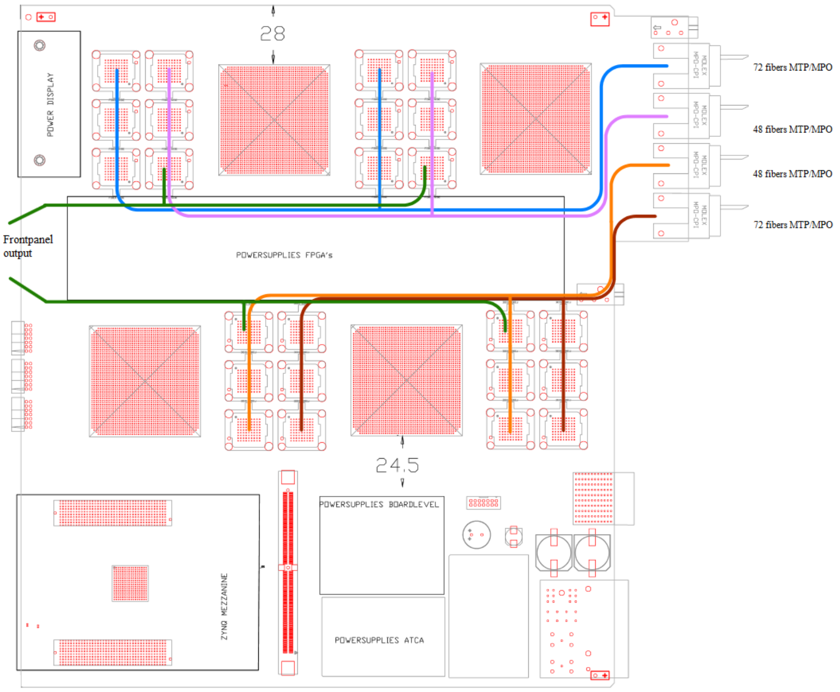

FIGURE 8.3: Sketch of the fibre routing channels for input and output signals, based on a preliminary board layout.

in terms of received calorimeter data, a core region of 0.8×1.6 in $\eta \times \phi$ is defined. The two remaining modules covering the forward regions of the detector are able to cover a core area of 3.3×1.6 in $\eta \times \phi$ due to the reduced granularity of the input data. Every FPGA is intended to run algorithms to define trigger objects that are centred within its core regions. Similar to the JEM design, additional data from calorimeter cells surrounding the core region is required on each FPGA to calculate these objects on every position of the core. The increased maximum jet size of 1.7×1.7 in $\eta \times \phi$ requires an environment of 0.8 in both η and ϕ. Therefore, the total data received for a single FPGA covers a region of 2.4×3.6 in $\eta \times \phi$ or 4.1×3.6 in $\eta \times \phi$ for the central and forward modules, respectively. Data corresponding to its core for each FPGA is received directly from the MiniPODs while the environment in ϕ is being forwarded by neighbouring FPGAs. In term of pseudorapidity, the outer 0.8 in η in both directions are considered environment except for the two forward modules where the core region is extended towards the edge of the η-detector coverage.

The onboard data duplication of the input signals, which is required due to the overlapping detector region covered by FPGAs on the same module, is done via PMA loopback. Four FPGAs on a module, covering a ϕ-ring of detector cells and a maximum jet size of 1.7×1.7 in $\eta \times \phi$, lead to an overlap of 0.8 in ϕ for neighbouring FPGAs[4]. This requires half of the incoming data to be duplicated and sent to each of the neighbouring devices. Since a small number of high-speed links are required for control purpose, 29 links per neighbour are connected in either direction, leaving two input links from the MiniPOD block without the option for PMA loopback and therefore prohibiting its use in the regular real-time data path. Overall, every

---

[4] A neighbour in this sense is an FPGA, which covers a core area with the same η coordinates and that is adjacent in ϕ to its own core area. The four FPGAs are placed on the board in a scheme to minimise the distance between neighbouring FPGAs.

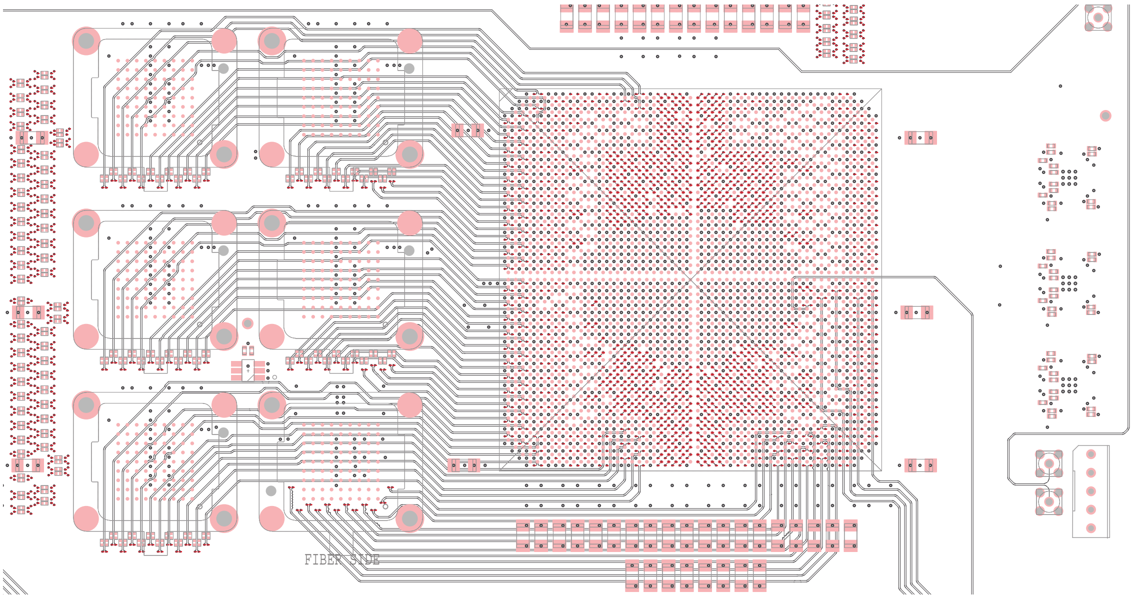

Figure 8.4: Pre-routed block of one FPGA and six MiniPODs, designed by Bruno Bauss. The breakout of other signals like differential pairs from parallel I/O pins is not shown in this view, however, was also done before the duplication of the layout.

FPGA can receive 116 signals with calorimeter data. Half of this is received directly from the MiniPODs and the other half via PMA loopback with some delay in the signal that needs to be considered in the algorithmic firmware.

The four blocks of FPGA and MiniPODs are placed on the board in a diamond shape with two of the blocks being rotated by 180° to simplify the routing of loopback channels.

Since MGTs for high speed signals are usually designed to transmit data over longer electrical traces, like backplane connections, they can easily transmit data over several meters. In this design, the longest data path for real-time signals is about 34 cm. All devices in use are built to support data rates of at least 14 Gb/s, leaving enough safety margin to not operate them close to the edge of their individual capabilities.

The only remaining uncertainty is the permanent use of the PMA loopback mode in the real-time data path since this mode is designed for testing purpose only. Possible effects of the loopback compared to normal transmission utilising the complete TX structure are addressed in initial tests using the first prototype and on the Xilinx VCU110 [71] and VCU118 [72] evaluation boards. The results are shown in chapter 9.

The XCVU190 has eight complete HP Banks, each able to support 24 differential links, and a half bank, providing up to 12 differential links, for general purpose I/O supporting bit rates of up to 1600 Mb/s in LVDS[5] mode. The remaining pins per bank, which do not support differential signalling, can be used for low-speed single ended signals or for static control lines. The majority of the general purpose I/O is connected to the other FPGAs to be utilised in the real-time data path. 72 differential links from three complete banks are connected to each of the neighbouring FPGAs and 60 links from two complete banks and the half bank are routed to the opposite FPGA. Since highest speeds on parallel I/O links will likely increase the latency due to required serialisation and deserialisation, those links are not expected to be operated at more than 640 Mb/s. Yet, latency restrictions might require an even lower rate which needs to be evaluated once first tests with parallel I/O are done and latency can be tested in combination with stability of the links. Therefore, a total bandwidth of 46.08 Gb/s and 38.4 Gb/s, for neighbouring and opposite FPGAs, respectively, can be achieved in both directions combined. Possible use cases for this bandwidth are additional inputs, that cannot be shared via PMA loopback, or intermediate results from algorithms like subsums for pile-up corrections. If the real-time data path does not utilise all connections, spare connectivity can be used for the read-out functionality as discussed in section 8.5. An overview of all high-speed connections and use of the parallel I/O links in the jFEX design is shown in appendix B on page 1 and 6, respectively.

8.3 Output Data Formats

As described in section 6.2, all draft TOB data formats use a common size of 40 bits. If the L1Topo firmware allows for a more flexibility regarding its input with different sizes being received on the same link in different events, this requirement may be dropped. In this case, the size of fat jet and τ TOBs can be reduced by removing unused spare bits. Alternatively, the size of the regular jet TOB can be reduced by removing the second energy level if this information is not used in the following trigger decisions. This allows for a more efficient use of the available bandwidth. If, due to the distribution of topological algorithms over the separate L1Topo FPGAs, only certain TOB types are required on individual L1Topo processors, it is also feasible for the jFEX to send different data to each FPGA.

In case the available TOBs exceed the output bandwidth in individual events, excess TOBs can be stored along with those being sent to L1Topo to be optionally added to the read-out data in case of an L1A signal. The exact scheme of prioritising is subject to detailed studies, though a prioritisation based on the transverse energy is most likely. Additional bits might be added to the data stream beneath the TOBs to indicate TOBs of different categories being suppressed. These categories can be defined to differentiate between objects of different TOB types, η positions or energy thresholds.

Assigning fixed positions on the output data streams to specific TOB types seems to be suboptimal use of the available bandwidth when frequently events are expected where no TOBs for a certain type are found, while other types are suppressed to the limited number assigned. Yet a combination of fixed positions for a single type and positions with flexible content might be reasonable to help reduce the complexity of the L1Topo firmware.

---

[5] Low-Voltage-Differential-Signaling is a technical standard used for high-speed signalling based on copper connections.

All described considerations on TOB formats can be implemented in the jFEX firmware and have therefore no direct influence on the hardware design.

8.4 Control

Except for the IPMI, for which a CERN IPMC card is used, control functionality is implemented on an extension mezzanine. This option is chosen to maintain flexibility on how these functions are implemented by redesigning a small module instead of the main board. In this setup, it is easily possible to adapt to different testing environments or, e.g. in case errors occur, by replacing the mezzanine. Also, fewer components need to be placed on the main board where a single large connector from the Amphenol FCI GIG-Array® series[6] replaces a larger number of devices. For the first iteration, a PicoZed™[7] is used to minimise the effort required to develop the mezzanine card. The contained Zynq allows to choose between programmable logic and processing system when implementing functionality. For the remainder of this section, unless specified otherwise, the PicoZed or its contained Zynq is assumed to provide the described control functionality on the mezzanine. Even though the PicoZed board will be replaced by either a larger Zynq or an FPGA in later iterations of the board.

Backplane signals for incoming TTC data from the Fabric Channels and the IPbus link from the Base Channels are routed directly to this mezzanine. The Zynq serves as IPbus master, which maintains the communication to the surrounding system and operates point-to-point connections to the processing FPGAs to pass on relevant information and collect data from them. For testing purpose, the IPbus can also be operated via an RJ45 connector on the mezzanine placed towards the front panel. The TTC data can either be fanned out by a device qualified for the link speed, which was not yet defined during the design phase of the jFEX or received by the Zynq and sent to the FPGAs on Point-to-Point links. In the second case, link speed and data format can be changed to fit possible limitations on the processor FPGAs.

Additional control functions are operated by the Zynq, using I^2C[8] or PMBus[9] for the control of MiniPODs or power supplies, respectively. Via these connections, the Zynq can read out parameters and change setting for these devices, according to their individual commands, as described in the corresponding datasheets.

In case the IPMC card fails or not all signals are received correctly, these signals are also

---

[6] The largest footprint with 296 signal pins is used. Based on the same footprint, multiple versions with different heights are available. This leaves flexibility to bring the mezzanine closer to the main board for single slot width or leaving more space for airflow in dual slot width applications without any changes to the PCB layout. Since this type, as most large array connectors, is only designed for a limited number of mating cycles, the mezzanine cannot be replaced arbitrarily often. Yet the given number of 25 cycles should be more than sufficient for both prototype testing and operation of the final production series.

[7] A System-on-Module based on a device from the Zynq-7000 SoC family. These modules can be plugged onto the mezzanine using a single connector and run on a single supply voltage. This greatly reduces the design work compared to placing an FPGA or a Zynq device directly on the board.

[8] Inter-Integrated Circuit, commonly pronounced as I-squared-C, is a serial bus typically used to attach lower-speed peripheral ICs to processors. It consists of two lines, one for a clock signal provided by a master and a data line for bidirectional communication. On the jFEX, a single master, implemented in the control unit of the extension mezzanine, is used to read out multiple ICs using commands defined by the slaves.

[9] Power Management Bus is a communication protocol based on I^2C specialised on communication with power supplies. Unlike the I^2C standard, in PMBus a larger number of application specific commands are defined in the protocol. Furthermore, an additional alert line can be used by the slaves to signal certain conditions and initialise communication with the master.

routed to the extension mezzanine[10]. There, a control according to the IPMI standard can be implemented or an alternative IPMC card can be mounted. If a mezzanine is designed including IPMI functionality, this must not be operated in parallel to an IPMC card connected to the main board.

Some low-priority control functions are implemented using an XC2C256 CPLD from the Xilinx CoolRunner™-II series [73] residing on the main PCB. This type has two separate I/O banks. One is supplied with 3.3 V to communicate with onboard chips operated at 3.3 V or 2.5 V. The other one operates at 1.8 V to support 12 single-ended connections to each FPGA. Therefore, the CPLD can also be used as level translator to allow the FPGAs to indirectly communicate with devices that would not be voltage-compatible for direct connections. Additionally, 20 spare links from the 1.8 V bank of the CPLD are routed to the mezzanine. These connections can, flexible with every iteration of the mezzanine, be routed to the Zynq, connected to links from any FPGA or left open. Besides clock control, which will be described in section 8.6, two signals per MiniPOD are routed to the CPLD. These are connected to the IntL and ResetL[11] pins on the opto-electrical converters. Control functions using these connections can either be implemented on the CPLD or forwarded via the CPLD from one of the FPGAs or the Zynq.

Since no communication via Zone 2 of the backplane is required except for the connections to the two Hub slots, the number of connectors can be reduced to unblock their space on the PCB. Only the connector described as J23/P23 in the PICMG specifications will be assembled. Furthermore, on this position a connector of a newer type, the HM-Zd+ from TE connectivity is used. While it is based on a modified footprint, the mechanical characteristics towards the backplane are identical. Therefore, the jFEX maintains the compatibility to the originally designed system while gaining the option to operate backplane links on the higher link speed of up to 12 Gb/s as supported by the new connector type. Overviews of control functionalities are represented in the block diagrams in appendix B on pages 9 and 11.

8.5 Read-Out

The read-out functionality is primarily implemented as Point-to-Point connections from each FPGA to both Hub modules. These backplane connections are qualified for up to 10 Gb/s in standard ATCA crates, which requires MGT transmitters on the FPGAs to be operated at highest possible data rates. With all GTH links being in use by the real-time data path, GTY links are used for this functionality. Additional to one regular link, two FPGAs, U1 and U2, have a second link to each Hub slot. This connectivity serves as spare in case the read-out requirements are increased in Phase-II and can no longer be transmitted on a single link. The excess read-out data from U3 and U4 can be shared on spare parallel I/O links that are not used in the real-time data path. If there is not enough spare connectivity between neighbouring FPGAs available, data sharing can instead be done diagonally across the board. So far it is assumed that both Hub slots receive the same read-out data. Yet, if required, it is possible to send different data or formats to each Hub module. An overview of the read-out connections is given in figure 8.5.

---

[10] This option has been removed in later iterations of the board design.

[11] The Interrupt signal to Host, with an active low state, is used to flag any fault conditions on the MiniPOD. The Reset signal to module, asserted low, disables the optical outputs, and resets all non-volatile settings.

Since the latency requirements on read-out is not nearly as critical as in the real-time data path, a smaller number of links can be used running at a higher speed than 640 Mb/s. With no expected difference in average volume of read-out data produced by the processor FPGAs, half of the additional bandwidth could be assigned to the excess data from an FPGA without the extended read-out connectivity. Therefore, four parallel I/O links, operated at a maximum of 1.25 Gb/s, can transmit the complete excess data, coming in from another device, that can fit on a shared backplane link. In case different read-out data is required to be sent to the two Hub slots, a maximum of eight links would be required. If different encoding schemes, like 8b/10b on the MGT links and no encoding on the parallel I/O links, are used, seven connections would be sufficient.

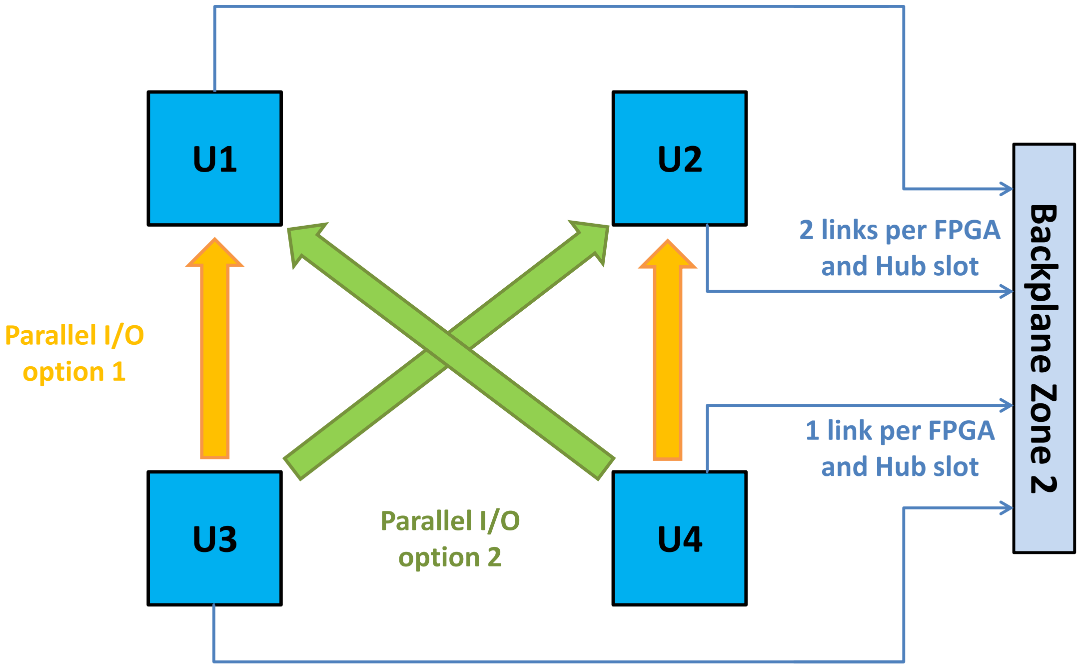

FIGURE 8.5: Schematic representation of the connections used for read-out functionality showing both possible setups for data sharing between FPGAs with different backplane bandwidth via parallel I/O.

8.6 Clock

In order to operate all various connections and functionalities on the processor FPGAs in parallel, a large number of different clock signals is required. In total, there are five MGT reference clocks (refclks) and two global clocks (GCKs). Additionally, there are two special MGT reference clocks provided to U1 exclusively to allow for further test setups. The global clocks are operated at the LHC frequency of 40 MHz and the reference clocks at higher rates. In most cases, the reference clocks are operated at low multiples of the LHC clock depending on the transmission speed of the MGT signal.

Since multiple copies of the same clock are required on the board, so-called clock trees are used to distribute the signals. Multiplexer chips[12] are used to provide the flexibility to switch between varying sources. Derived from the chosen input, multiple stages of fan-out chips[13] can be used to provide the required number of copies.

For the prototype design, most of these clock trees can choose between multiple inputs including onboard sources to provide stand-alone capabilities. Please note that this flexibility required for testing purposes is expected to be reduced in the final design. A proposal for such a simplified clock distribution focused on the relevant usecases during operation as a part of the ATLAS Trigger System is presented in section 9.5.2.

To guarantee best possible quality, a high amplitude differential signal is preferred for clock signals. Therefore, CML signals operated at 2.5 V with a swing of 0.8 V are used wherever possible. While the MGT reference clocks are received on dedicated pins, global clock inputs on the FPGAs are integrated in regular parallel I/O banks, requiring their input level to be compatible with general purpose pins. The conversion from CML to LVDS can be done in fan-out chips that can handle multiple input levels by adjusting the input termination externally.

Even though there are two connections for global clock input to each FPGA, they cannot be operated completely independently. As shown in figure 8.6, there are four possible sources. A local crystal oscillator[14] is placed on the main board, providing a stable 40.0787 MHz clock on LVPECL level intended for stand-alone tests where no other clock inputs are available. The LHC clock can be received from either of the Hub modules even though only the module in Hub slot 1 is expected to send a clock signal. Both of these inputs come from the Fabric Channels of the backplane on CML level. The fourth input arrives from U1 after CMOS to LVDS level conversion. This option is intended for test setup where a fixed phase relationship between an MGT reference clock and a global clock is required without the option to generate the refclk from the GCK. These four input channels are connected to a 4:2 multiplexer, where the local crystal oscillator is selected as default. The two select pins on the multiplexer, used to define which input option is active, are pulled to 2.5 V to ensure the default setting. Other inputs can be chosen by pulling a combination of these lines to ground. For the control of this multiplexer, as well as all other multiplexed solutions in the clock trees, spare links on the CPLD used for control functions are utilised. Since the connected CPLD bank is operated at 3.3 V, the output signals of this bank must not be set to high to avoid damage on the clocking chips. This signal state is set as a default by the board via pull up resistors of 4.7 kΩ against the 2.5 V supply voltage. Only low states are to be handled actively by the CPLD.

The output of the multiplexer is fed into a clock generator[15] and a 2:8 LVDS fan-out with an

---

[12] Chips with multiple inputs that allow to choose any of them as source for the output.

[13] Chips with multiple output links that duplicate the received input. They can be combined with multiplexer chips to create n:m devices. In case of devices combining multiplexing and fan-out functionalities, they are referred to by their main task in the description of the jFEX design.

[14] Also described as quartz in the block diagrams and schematic representations of the clock distribution scheme.

[15] For the global clock tree, this device serves only as a jitter cleaner providing a cleaned version of the input signal. Since the chip used in the jFEX design, the Si5338 has a total of four outputs, the remaining links provide multiples of this clock to three of the MGT reference clock trees.

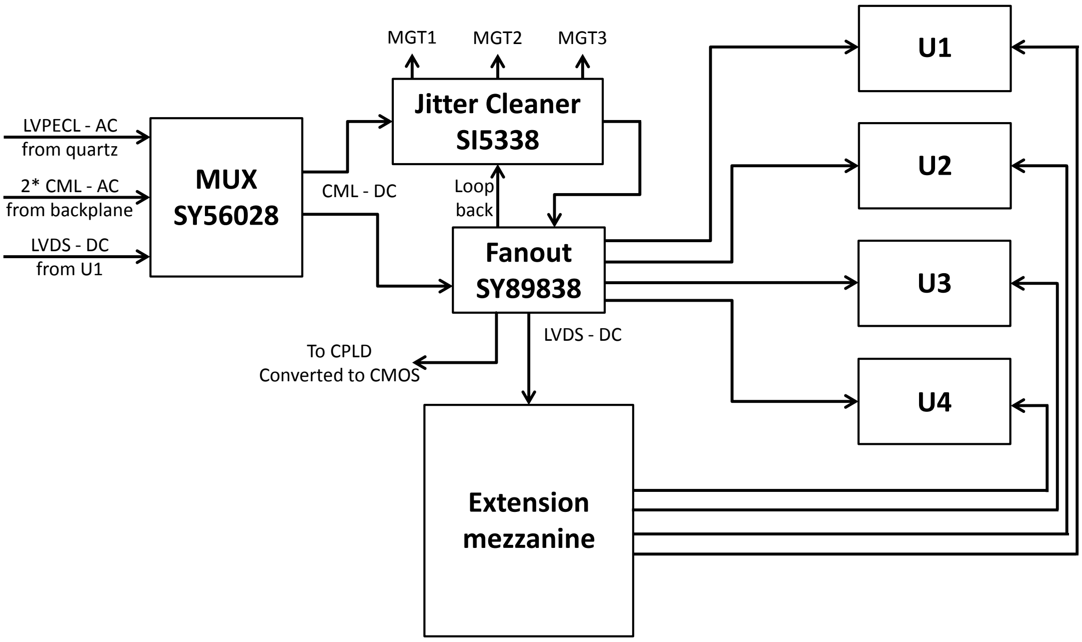

Figure 8.6: Schematic representation of the global clock tree on the jFEX PCB including the utilised signalling standards.

integrated two channel input multiplexer[16]. The jitter cleaner is used because so far it is not known if the quality of the clock signal from the backplane connection is sufficient for stable operation of the FPGAs. Therefore, a cleaned version of the chosen global clock is provided to the second input of the fan-out chip.

The fan-out provides a global clock signal to each FPGA, the extension mezzanine, a loopback signal to the jitter[17] cleaner, used for a special operation mode with a stable phase relationship between input and output, and the CPLD, which is converted from LVDS to a single-ended CMOS level. The unused output is terminated with a 100 Ω resistor. A more detailed version of the distribution of the global clock signal is shown in appendix B on page 5. On the mezzanine, a separate jitter cleaner is placed, which also provides GCK inputs to each FPGA and several multiples of it as MGT reference clocks. By placing additional clocking resources on the mezzanine, flexibility for testing during the prototype phase is gained, allowing to test varying devices and providing stand-alone operation capabilities using the onboard circuits.

Since the GCK signal going to the mezzanine is the same as the one going directly to the FPGAs, these clocks cannot be driven from mixed sources. If two independent global clocks

---

[16] Even though all used devices have this integrated input multiplexer, they are referred to as fan-out in this chapter due to their main functionality.

[17] Derivation of the received signal from an ideal clock signal caused by fluctuations. Static shifts of the phase are not considered jitter and can be compensated in the FPGAs.

are required, an additional clock can be connected directly to the mezzanine from an external source via cable or created using a crystal oscillator. Appropriate connectors are placed on the first iteration of the mezzanine and connected to spare inputs of the jitter cleaner.

An overview of the MGT reference clock trees that include multiple sources and onboard fan-out is given in figure 8.7 and figure 8.8. All multiplexing control, which is not shown in the simplified schematic representations, is handled by the CPLD under the same conditions as in the global clock tree. A more detailed version is shown in appendix B on pages 3 and 4.

Throughout the complete design, the same 2:6 fan-out chips are used in all positions that require CML level signals to reduce the number of unique components on the board and thereby reducing the layout work and development time. With a 5 mm × 5 mm 32-pin QFN[18] package the possible saving of space on the PCB by replacing few of these devices with smaller versions is negligible.

The first MGT reference clock, intended for the reception of real-time input data (input clk), requires the widest fan-out and highest flexibility in terms of available sources. A simplified representation of the distribution of the input clock signal is given in the upper part of figure 8.7. Four onboard clock sources for stand-alone tests are available: three crystal oscillators, pre-programmed to 120 MHz, 160 MHz, and 280 MHz, respectively, and a programmable clock source, controlled via I^2C. For tests in combination with other modules, the two clock generators on the main board and the mezzanine provide adjustable multiples of the global clock. Varying clock frequencies are required for testing the module at the variety of transmissions rates mentioned in chapter 6. From these input options, a single source is selected and provided to all FPGAs via a two staged fan-out scheme. With the option of forwarding the reference clock between neighbouring quads within an SLR, only 12 clock signals per FPGA are required to supply all 120 MGTs. Since the option to forward a reference clock from one MGT to up to two neighbours on either side is new in the UltraScale series, the jFEX design does not rely on this option but only utilises this feature for direct neighbours as it has already been implemented in previous FPGA series. Yet, it is possible to test this operation mode on a subset of the links to verify its reliability for future designs.

The second MGT reference clock, responsible for driving the transmitters in the real-time data path (output clk), is supplied independently to every quad that is connected to a transmitter MiniPOD and additionally to a quad where, for U1 only[19], two transmitters are connected to signal test points with different connection lengths allowing to measure high speed signals with an oscilloscope. An independent clock supply to every quad was chosen to maintain the option of operating the output links at bit rates above 16.375 Gb/s, which requires a special mode in the GTY transceivers that prohibits the use of forwarded reference clocks. Higher speeds can be used to test the performance of the FPGAs when operating close to its limits and as an option for the actual output data path if suitable opto-electrical converters become available in the future. A simplified representation of the distribution of the output clock signal is given in the lower part of figure 8.7.

The third reference clock, driving the output links for the read-out data (read-out clock), is distributed in a single staged scheme via one fan-out chip, receiving inputs from both clock generators.

---

[18] Quad Flat Package is a common surface mounted package for integrated circuits with pins sticking out on each side.

[19] To simplify the design, the clock input is available for all FPGAs and only the data lines are exclusive for U1.

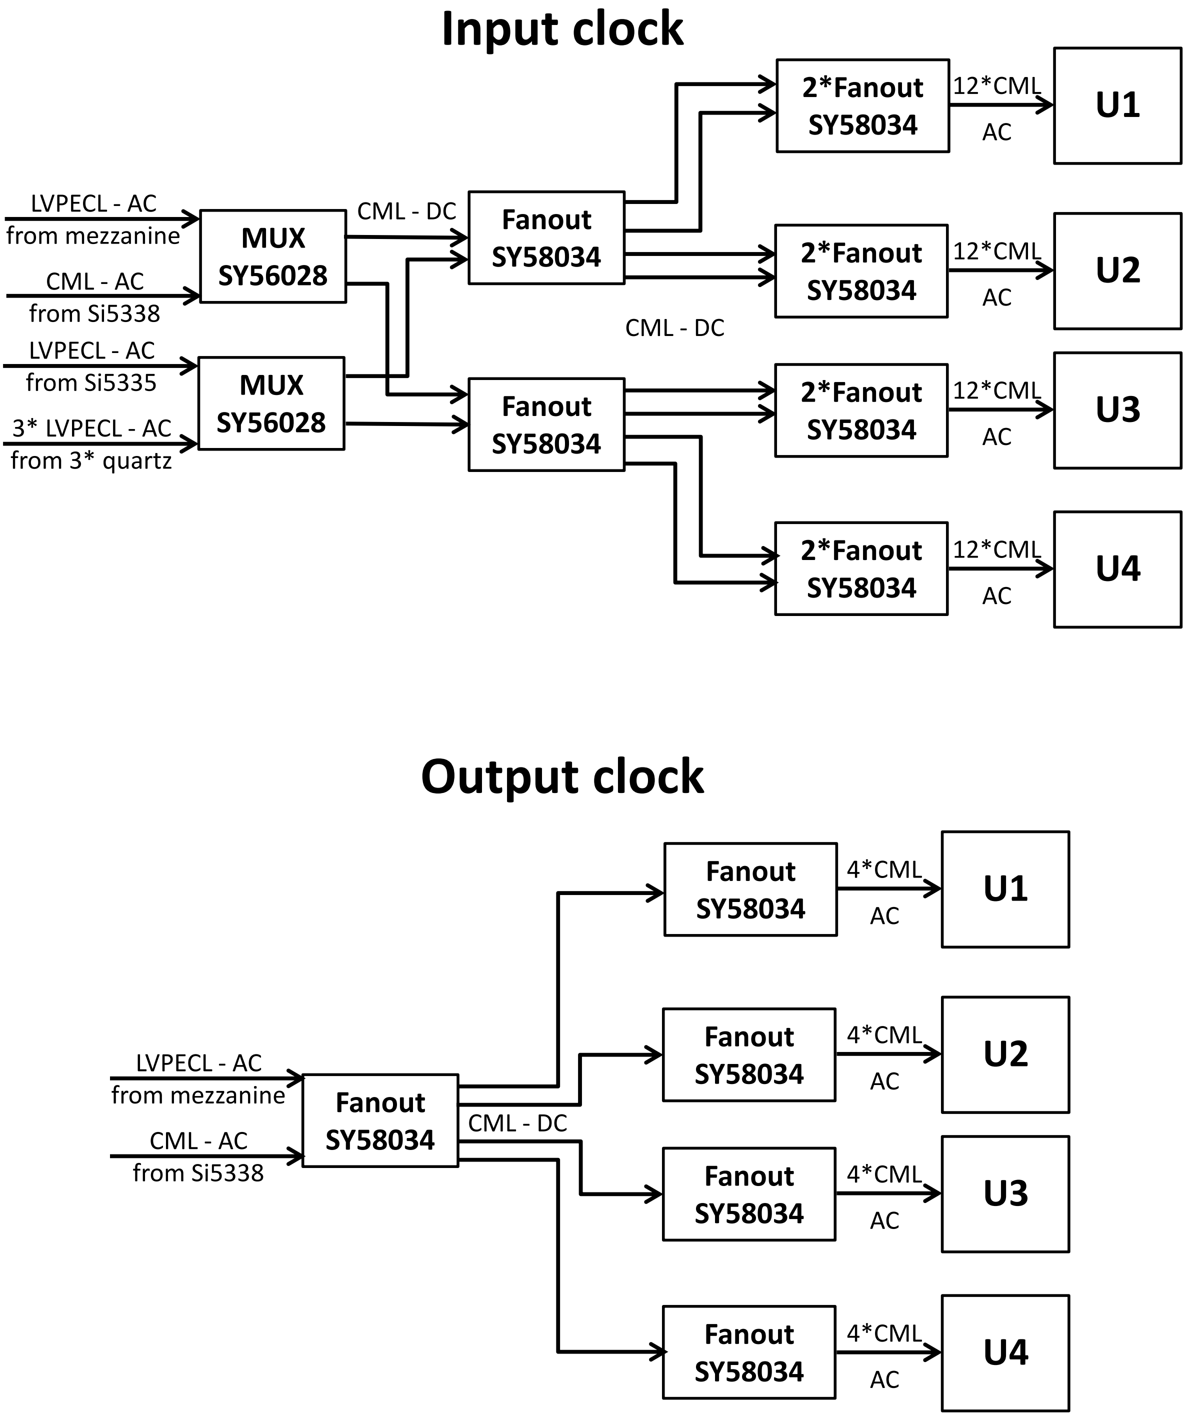

FIGURE 8.7: Schematic representation of input (top) and output (bottom) MGT reference clock trees on the jFEX main board. A more detailed representation is given in appendix B on page 3.

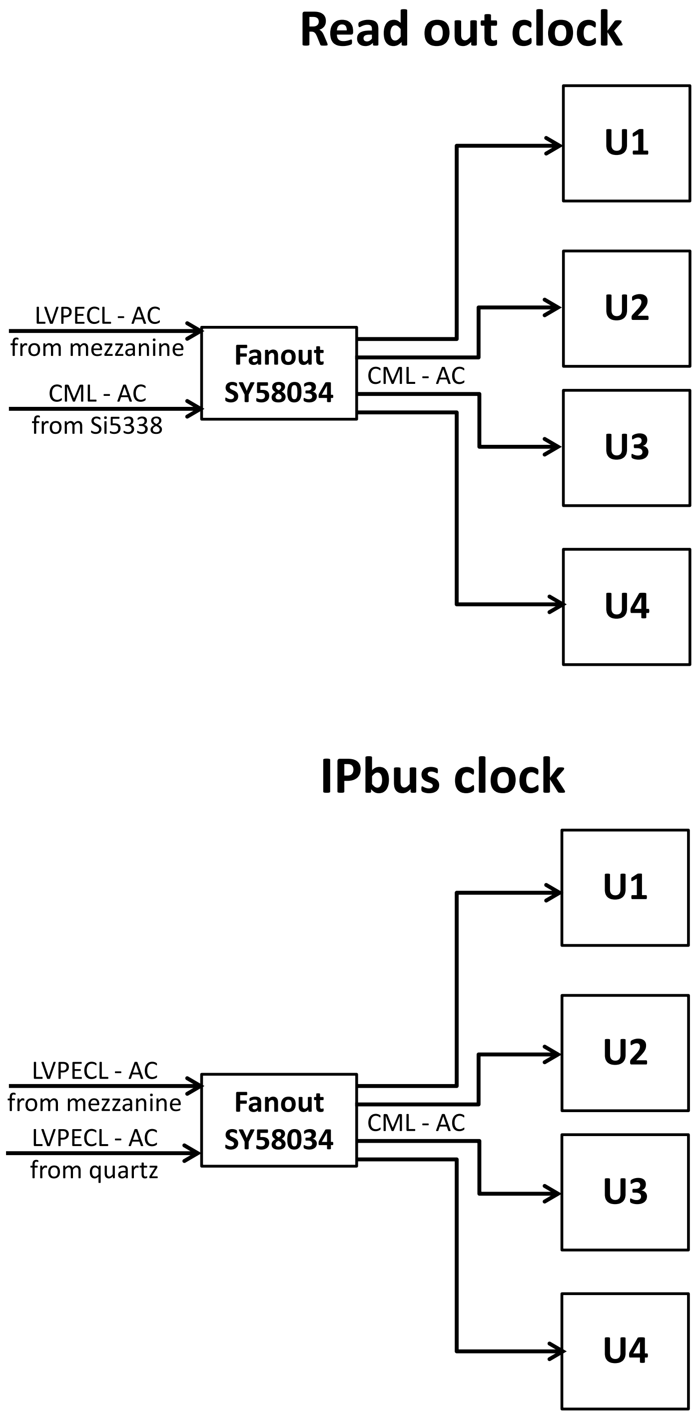

FIGURE 8.8: Schematic representation of read-out (top) and IPbus (bottom) MGT reference clock trees on the jFEX main board. A more detailed representation is given in appendix B on page 4.

The fourth reference clock is meant to drive the links for the IPbus connections to the mezzanine (IPbus clock). It is provided for each FPGA to the quad from where the signals are routed to the mezzanine. Additional fan-out is routed to U1 to allow for external connections via optical links in test setups where no mezzanine is connected. The available sources are a local crystal oscillator, set to 125 MHz, and an input from the mezzanine used as default. Independent from this clock tree, U1 receives input from another local 125 MHz crystal in a separate quad to provide the option of receiving control commands via another optical link. This option is not indented to be used during operation as a part of the ATLAS Trigger System but can be used in test setups in parallel to the complete real-time data path, since it can be operated on one of the spare links with no PMA loopback functionality.

Schematic representations of the distribution of the read-out and IPbus clock signals is shown in figure 8.8.

The last reference clock (TTC clock), used for the potential reception of TTC data as described in chapter 6, is received by each FPGA via a Point-to-Point connection from the mezzanine in the same quad as the IPbus clock. Since this clock is intended for the reception of TTC data redirected from the mezzanine, using a multiple of the LHC clock is advisable. As the onboard clock generator has no unused outputs, no suitable clock source on the main board is available. Therefore, only an input coming from the mezzanine is used.

The CPLD receives two additional single-ended clock signals. One coming from a Schmitt trigger based oscillator[20] set to 1 kHz and the other one from U1, where the frequency can be defined in the FPGA firmware.

8.7 Configuration

Since the FPGAs have no non-volatile memory, they require configuration after every power-up. There are multiple ways to configure the FPGA type in use, controlled by combinations of logical 0 and 1 on dedicated M-lines, M0, M1 and M2, as described in detail in the UltraScale Architecture Configuration User Guide[74]. The configuration modes chosen to be made available in the jFEX design are Master SPI, Master and Slave SelectMAP, and JTAG. Serial configuration modes are also available since they use a subset of the links required for SPI[21] mode but offer no advantage over other modes and are therefore not intended to be used. Also, the Master Serial mode is not recommended by the manufacturer since it will no longer be supported in future FPGA series.

In general, master modes describe an active state where the FPGA will automatically start sending a clock signal from the CCLK pin after power-up and expect incoming data on the input pins corresponding to the selected mode. In a slave mode, the FPGA uses the CCLK pin as an input, waiting for an external device to provide a clock and start the configuration process by sending the configuration data.

The maximum possible clock frequencies, as described in the DC and AC Switching Characteristics [75], for the different configuration modes are between 100 MHz and 150 MHz. As the

---

[20] A comparator circuit which applies positive feedback to its own noninverting input to implement hysteresis. When a single RC integrating circuit is placed between the inverted output and input, it can be used as a simple oscillator.

[21] Serial Peripheral Interface, serial bus system with multiple signal lanes providing data synchronous to a single clock signal.

frequency tolerance for the internal clock sources in master mode is relatively high at ±35 %, the configured link speed needs to be adjusted for this as to not exceed the defined maxima. If more stable frequencies are desired, e.g. to increase the clock rate towards the specified limits, the FPGA-internal clock source can be replaced with an input from the EMCCLK input. The size of the bitstream, which is loaded in the FPGA during the configuration, is about 600 Mb for the XCVU190. Configuration time can be reduced for some of the modes by using multiple data lines in parallel.

The bitstream can either be stored on memory cards that offer an appropriate interface or accessed indirectly via a controller chip. Updates of the stored firmware images are intended to be written via an interface provided by the FPGA via IPbus commands. With this option, no direct access to the storage from outside the jFEX system is required.

In the Master SPI mode, the FPGA sends a clock signal and establishes a bidirectional interface with one output pin. If multiple data lines are used, the output pin is converted to an input during the reception of input data. In the mode with ×8 bus width, a dual ×4 can alternatively be used, allowing the use of two flash memories in parallel. This mode increases the maximum storage capacity, potentially allowing to store multiple bitstreams. This option is useful when, while testing new firmware, a fallback version of a firmware, known to be operational, can be made available. Another requirement on the flash memories is that after only a single command from the FPGA, the whole bitstream needs to be sent without interruption. This prohibits the use of most of the larger flash memories available during the design phase due to built-in limitations on the size of delivered bitstreams per command. As the memory cards are placed on the mezzanine, it was decided to rely on JTAG-based programming of the FPGAs for the first tests and include the other methods at a later stage of the development.

SelectMAP is a configuration mode available in Xilinx Virtex and most Spartan® families, that is similar to SPI but offers up to 32 parallel connections to reduce configuration time. Since only four data lines are available on dedicated configuration pins and every further link reduces the bandwidth during operation of the FPGAs, only the pins required for the ×8 SPI interface are used for configuration purpose, limiting the SelectMAP bandwidth to one byte per clock pulse.

Independent of the chosen configuration mode, the FPGAs always respond to an incoming JTAG[22] signal. This mode continues to be the best supported access to the FPGAs since commercially available USB to JTAG cables can be used to control the device from a connected PC using the Vivado Design Suite. The four JTAG pins are routed to a header matching the one from the USB cable. To allow for boundary scans[22], all FPGAs are cascaded in a so-called daisy-chain. To reduce the configuration time for multiple FPGAs, this chain can be disbanded and connected to separate headers by removing and adding 0 Ω resistors. A representation of the different configuration of the JTAG lines is shown in appendix B on page 8.

All configuration lines, except for the separate JTAG connection, are routed to the mezzanine to maintain flexibility in the choice of configuration mode. The M-lines are set via 1 kΩ resistors to a default of $M[2:0] = 101$ setting the FPGA in a state where it waits for an incoming JTAG signal. This default can be overwritten by placing resistors with lower values on the mezzanine.

---

[22] Short for Joint Test Action Group, a standard defined by IEEE in 1990. Additional to the option of storing bitstreams in the device, boundary scans can be performed. During this scan including multiple complex integrated circuits, the electrical connections on a PCB can be tested. This scan is performed in the commissioning phase for every jFEX module from the prototype phase to the production series. Four mandatory pins are required for the daisy-chained JTAG, TDI (Test Data In), TDO (Test Data Out), TCK (Test Clock) and TMS (Test Mode Select).

Therefore, during tests without a mezzanine, no onboard configuration mode is available. In most test setups JTAG is used to monitor the FPGA during the run.

For the operation in the L1Calo system, Master SPI mode is expected to be used.

A schematic representation showing all links that can be used for configuration interfaces, except for the separate JTAG interface, is shown in appendix B on page 7.

8.8 Power

Most requirements on the power input and distribution within the module are set by the PICMG specifications and the FPGAs. The ATCA dual -48 V input is received by a PIM400 power input module and converted to a stable $+48$ V, a management 3.3 V supply, alternatively referred to as 3V3_up, and a 5 V output, which remains unused in the jFEX design. In a second step, the 48 V are converted to a 12 V supply voltage using a QBVW033A0B step down converter which feeds all remaining circuits on the jFEX. According to the ATCA specifications, 3V3_up and the 5 V supply are limited to 3.6 A and 0.15 A, respectively. In total, no more than 200 W per node slot or 400 W per double width slot shall be used. Already during the jFEX design phase, there were crates available that support power input and cooling for significantly higher consumptions. The crates used for the operation of the ATLAS Trigger System had not been selected prior to the jFEX design phase. As the crates are purchased centrally by ATLAS for all new trigger hardware, requirements for cooling and power supply can be considered. The jFEX modules are designed with an input stage and 12 V converter capable of delivering 400 W independent of the chosen slot width. The different supply voltages and the estimated power consumption are described in section 8.8.1.

Only circuits required for basic control functions in the power management are supplied via 3V3_up. The main control functionality is handled by the IPMC card. Basic monitoring of onboard temperature sensors is also handled by the IPMC module, powering down the module if the temperature exceeds a threshold. This functionality is generally supported by the ATLAS IPMC card but needs to be implemented in firmware. According to the ATCA standard, a signal in the IPMB to the shelf management is required to assure the board is properly monitored before full power is available. Until then, each slot is not allowed to require more than 12 W and the 12 V supply is disabled. For test setups where no IPMC is available or the firmware is not ready, a handle switch is placed on the board to manually overwrite the signal on the IPMB and constantly enable full power to the module. By doing this, all safety functionality from the crate is disabled except for fuses on the -48 V inputs.

Additional circuitry supplied by the management power are a CPLD[23], which handles the sequencing of various DC/DC supplies, and the control functions in the board level power converters for 2.5 V and 3.3 V. However, the payload for all DC/DC converters is taken from the 12 V rail. The sequencing of power supplies is described in 8.8.2.

---

[23] The CPLD also requires a 1.8 V input, which is drawn from 3V3_up using a small local converter.

8.8.1 Consumption

To determine suitable devices for DC/DC converters, an estimate of the maximum expected currents per voltage level is required. Except for supplies to FPGAs and CPLDs, only two different supply voltages are required: 2.5 V and 3.3 V. This includes all smaller ICs for clocking and monitoring. CPLDs and FPGAs operate at a common 1.8 V level for all I/O functionalities. The FPGAs require a larger number of lower voltage inputs for operation of specialised parts. Even though the input pins and internal metal interconnects are separate, some supplies at the same voltage level can be tied together or are even required to be connected on the PCB for optimal performance of the FPGA. This can be used to reduce the number of converters and to simplify the board layout by minimising the separately required power planes. A compromise needs to be found for this against adding unnecessary noise to crucial supplies and while remaining the option of powering certain voltages in a given order. A detailed description of which voltage can or need to be connected to each other is given in the UltraScale Architecture PCB Design User Guide [76]. For this design, an optimum was found by using six separate supplies.

A first estimate on the expected currents per supply voltage and overall heat dissipation was obtained using the Xilinx Power Estimator (XPE) [65, 66]. Version 2016.3 was used as the most recent version of the XPE available during the design phase that supported UltraScale devices. Before power tests using the actual hardware were available, multiple confirmations of this initial estimate with newer versions of the power estimator were done, resulting in only small deviations.

While interfaces for the FPGAs were quite well known, the currents drawn by the fabric are harder to predict. With no available implementation of firmware or detailed knowledge on final algorithms, only an appraisal for a worst case can be used. The estimate is based on the empirical knowledge from implementations of firmware for the Topological Processor [77], where usually no more than 70 % of the logic resources are utilised. To compensate for potential improvements in architecture and synthesis tools, it was assumed that no more than 75 % of any resource will be used in jFEX firmware. Designs with higher utilisation of resources are assumed to fail due to the tight timing constraints set for jFEX and L1Topo alike. The toggle rate[24] was set to 50 % except for memory resources which are expected to be written and read less frequently with an assumed toggle rate of 12.5 %.

The assumed clock rates are set to 320 MHz for DSP slices as the most heavily utilised resource in the weighted jet algorithms. Memory resources are expected to be clocked at 40 MHz. All other internal resources are assumed to run at 160 MHz. MGT links are taken into account with a line rate of 12.8 Gb/s for the real-time data path and lower rates depending for control and read-out connections. The power consumption for general purpose I/O links are estimated based on a link speed of 1.28 Gb/s even though it seemed unlikely that these speeds could be reached for connections between the FPGAs.

This approach is assumed to overestimate the actual power consumption of the final firmware design and thereby include some safety margin in the design. Initial tests of the power consumption were scheduled with the prototypes to verify the estimate and confirm the reliability of XPE predictions once sufficient firmware is available.

---

[24] Average transitions per clock cycle of the inputs. 50 % corresponds to random input while the actual physics related data from the detector would be assumed to stay below this rate due to the highest energy bits only being used for areas of the detector being hit with extremely high-energetic particles.

The separate FPGA voltage supplies in the jFEX design are listed in table 8.2 including their maximum expected currents.

The tolerance for all supply voltages is ±3 %. To minimise deviations from the central values, most DC/DC converters support remote sense lines that can be placed near or under the FPGA to compensate voltage drops over the PCB. This also allows for current dependent voltage drops to be automatically compensated with a minimal delay. Such a consumption dependent scheme can only be implemented efficiently when separate converters are used for each FPGA except for V_{CCO}, where a board wide consistent voltage level is preferable. Additionally, for voltages with high expected currents like V_{CCINT}, the combined demand from multiple FPGAs would exceed the capabilities of converters available during the design phase.

voltage level	maximal current	naming in official documentation
0.95 V	50 A	V_{CCINT}, $V_{\mathrm{CCINT\_IO}}$, V_{CCBRAM}
1.0 V	16 A	$\mathrm{MGTAV_{CC}}$
1.2 V	18 A	$\mathrm{MGTAV_{TT}}$
1.8 V	1 A	$\mathrm{MGTAV_{CCAUX}}$
1.8 V	2 A	V_{CCAUX}, $V_{\mathrm{CCAUX\_IO}}$
1.8 V	1 A	V_{CCO}

TABLE 8.2: Overview of the expected currents drawn from separate supply voltages per processing FPGA. The detailed XPE sheet is shown in appendix D. Total on-chip power is given as 90.5 W, which does not match with a simple summation of the separate contributions leading to 91.3 W. V_{CCINT} and connected voltages are used to power the digital circuits in the FPGA fabric and connection to I/O banks. Auxiliary voltages, such as V_{CCAUX}, are required for analog circuits supporting internal clocking and external data transmission. V_{CCO} is used to operate the parallel I/O banks. MGTs are powered by separate analog supply voltages. $\mathrm{MGTAV_{CC}}$ and $\mathrm{MGTAV_{TT}}$ supply the general MGT circuits and termination circuits, respectively. While only the chip wide consumption is listed the ratio of currents for the two supply voltages differ significantly due to the different architecture of GTH and GTY circuits.

With the exception of the low current voltages, $\mathrm{MGTAV_{CCAUX}}$ and V_{CCAUX}, where linear converters from the MIC68400 series from Micrel operating at an output level of 1.8 V are used, all supplies are placed on power mezzanines connected to the main board with special high current pins. Once a suitable set of DC/DC converters is found, it is possible to place them directly on the main board instead of a mezzanine. Whether this step is taken depends on many additional factors like available height for components, which is obviously decreased for the converters by placing them on mezzanines, and the impact of routing channels that are running under the power supplies. Modifying the design from a mezzanine-based solution to placement of these components on the main PCB might require rerouting of the signal layers.

Extended testing of power converters is required since the Xilinx documents specify low noise input, measured peak-to-peak at a high sampling rate, while most datasheets for supplies only cover noise measurements at a significantly lower sampling rate. To compensate unknown noise spikes, coupling capacitors are placed on the power mezzanines and near the FPGAs for all voltages. Most of the more complex converters, able to provide high currents, provide adjustable output levels. The exact way of adjusting the voltage is manufacturer dependent. This leads to the possibility of setting a wrong value for the supplies and potentially damaging the FPGAs. Placing the converters on mezzanines also allows to test different types in a test setup without the possibility of damaging an FPGA. The most reliable supplies can then be mounted on the jFEX in a tested configuration. A fuse is placed on the power mezzanine in the 12 V input, providing additional protection in case of failing converters. The value of the fuse can be adjusted to fit different setups to adapt to the respective expected power consumption.

The highest currents from the board wide 2.5 V supply are drawn by the MiniPODs, coming to a total of ∼ 12 A for receivers and transmitters combined, and various clocking chips, requiring ∼ 7 A. The total board consumption is estimated to be slightly below 20 A.

The overall power consumption supplied by the main 3.3 V is dominated by MiniPODs with a total ∼ 2.5 A and the remaining DC/DC converters. Since the converters are not fixed on the board and can be exchanged, an exact number for their power consumption cannot be given at this point. An estimate was based on the control current required for the MAX20751 from Maxim Integrated™, which is used on the Xilinx VCU110 evaluation board and was therefore also considered for the jFEX design. This chip serves only as a controller, monitoring up to four separate DC/DC converters. The MAX20751 has two voltage inputs for 1.8 V and 3.3 V. However, it is possible to supply both of them from the 3.3 V input in order to avoid including an additional power level on the main board. Combined, a maximum of 21 mA is drawn. Therefore, even with about 20 DC/DC converters being operated on the board in parallel, the total current for power control will be considerably below 1 A. The CPLD power consumption is highly dependent on the programming and clock frequencies. The standby current for the XC2C256 is maximum 150 μA and dynamic current reaches 27 mA at 50 MHz, which is faster than any clock received by the chip. Therefore, no relevant contribution to the overall current is expected from the CPLDs. Overall, no more than a few Ampere are expected to be drawn from the 3.3 V supply.

In order to simplify the design, at least for the first iteration of power supplies on the prototype, the same 20 A converter were used for 2.5 V and 3.3 V though it is expected to be significantly oversized for the 3.3 V consumption.

device	maximun current [A]	maximum power [W]
clocking circuits	7.1	17.9
MiniPOD transmitter	1.5	3.7
MiniPOD receiver	10.5	26.3

Table 8.3: Devices on the jFEX prototype drawing a significant current from the 2.5 V power supply including their maximum expected board-wide current and power consumption.

device	maximun current [A]	maximum power [W]
JTAG level translator	0.3	0.7
MiniPOD transmitter	0.7	2.4
MiniPOD receiver	1.8	5.9

Table 8.4: Devices on the jFEX prototype with highest contribution to current drawn from 3.3 V supply including their maximum expected board-wide current and power consumption.

8.8.2 Sequencing

Besides the information which voltage supplies can be connected to each other, the PCB Design User Guide also provides recommendations for sequences in which these supplies should be ramped to minimise current drawn and avoid damaging the FPGAs during this process. While

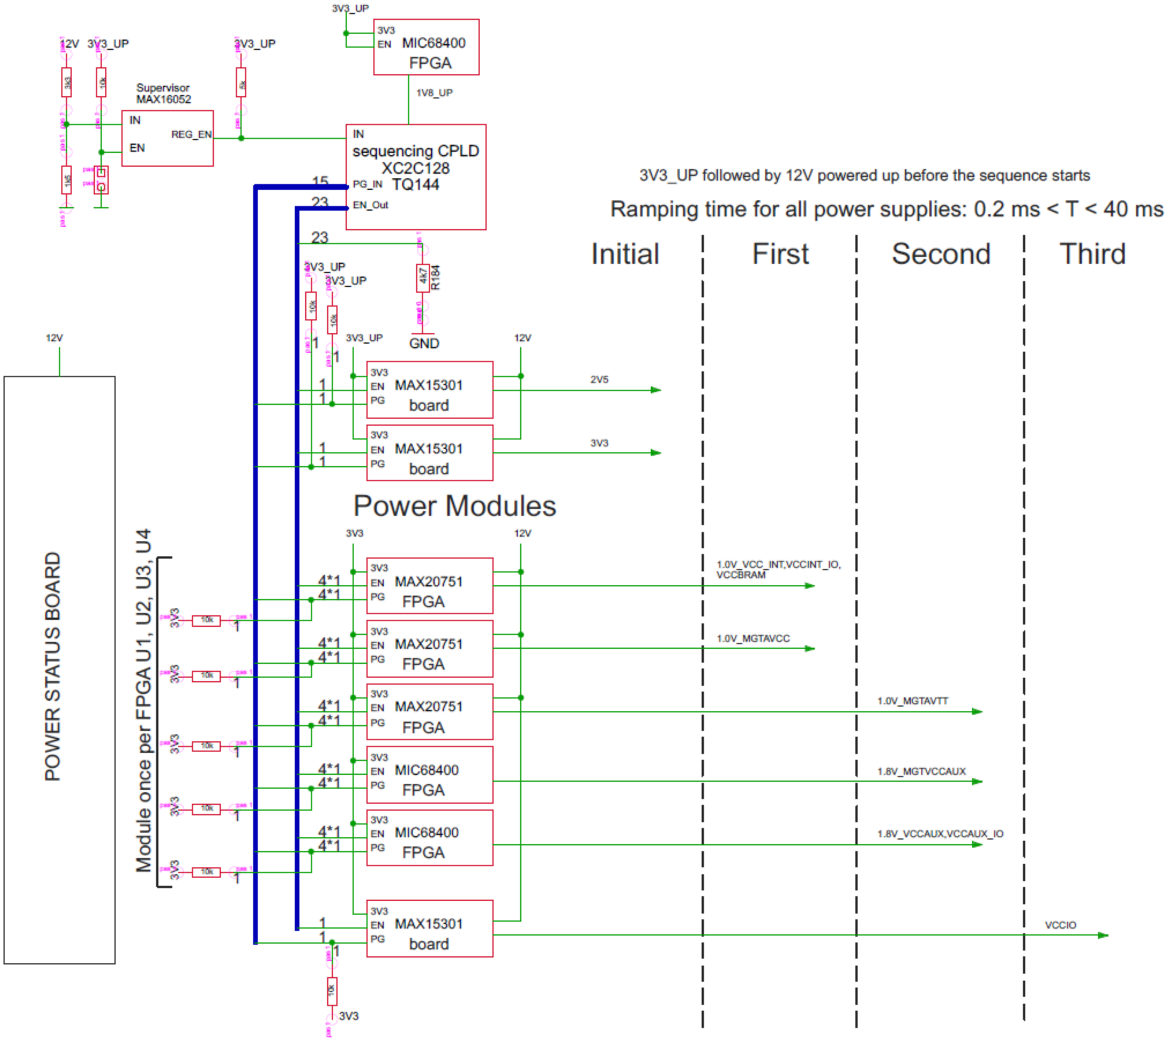

Figure 8.9: Sequencing of the separate power supplies for a single FPGA. Multiple FPGAs can be powered independently but are synchronised by the sequencing CPLD. 3V3_up and 12 V are available before this sequence starts. The used names in some cases, e.g. MAX15301 or MAX20751, refer to devices that were taken into consideration during the design phase and may not represent the actual design. If power supplies are replaced with other models, the sequencing remains identical. This sequence replaces the original scheme included in appendix B on page 2.

there is more than one suitable sequence, figure 8.9 shows the one chosen for the jFEX design, replacing the original sequencing as shown in the block diagrams in appendix B on page 2. This scheme is started and controlled by the CPLD after the 12 V supply is in stable operation and signals this condition accordingly. For every DC/DC converter, a separate signal is sent by the CPLD for power-up and a power good signal is received when the output reaches a threshold. This feedback is given directly from some converter types. For those who do not support this, an external comparator is placed on the corresponding mezzanines. The FPGA power-up sequence is split into 3 rounds with globally specified ramping times and an initial phase to start the remaining infrastructure on the board. For each power supply, figure 8.9 shows the communication channels originating from the blue bus comprising multiple separate connections. Additionally, power input and output including respective voltages and supplied FPGA functionalities are defined.

Chapter 9

Design Verification and Validation

In order to prove the suitability of the presented jFEX design to provide the required functionality in the trigger system of the ATLAS detector, multiple tests have been conducted addressing both, the ability of the hardware to perform as intended in the design as well as the aptitude of the logical design to implement the trigger algorithms. These tests include preliminary simulation of the performance of the algorithms as a part of the ATLAS Trigger System in 9.1, an initial estimation of the resource consumption in the FPGAs for the largest algorithms in 9.2, and tests targeting the most critical part of the hardware design, the high-speed data transmission between the separate stages of the L1 trigger in 9.3.
The scope of this thesis covers only the early development phase up to initial tests with the first prototype around mid 2018. Later stages in the development until the successful commissioning of the jFEX as a part of the ATLAS Trigger System were carried out by my colleagues and are only summarised and referenced.
Based on the results of the tests described below and altered requirements for the modules from the ATLAS group after the first prototype, some suggested changes to the design are presented in 9.5. These changes have partially already been included in the final design of the jFEX.

9.1 Simulation of Jet Algorithms

The goal of the simulations presented in this section is to show that the intended algorithms for the jFEX are suitable for the first trigger stage of the ATLAS experiment. Only preliminary results from an early stage of the development are presented. As by that time only relatively small samples of simulated data for Run 3 were available, detailed studies have not been conducted within the scope of this thesis. Samples used for the simulation of the performance of trigger algorithms in Run 2 could not be used as they simulated different operating conditions in terms of instantaneous luminosity and resulting pile-up. Therefore, these samples would give incorrect results at least for the expected trigger rates as they are based on conditions with fewer interactions per second. Additionally, the Run 3 samples stored in Mainz and the analysis code were lost due to the shutdown of the Mogon1 cluster where the code was stored and could not be recovered for further refinements. Therefore, this performance study occasionally contains incomplete results. To compensate the lack of detailed simulation results within the scope of

this thesis, an overview of the simulation results of the implemented algorithms on the jFEX is given in 9.1.4 based on the thesis of Marcel Weirich [78].
The jFEX is designed to run multiple algorithms in parallel including global energy sums and the identification of missing transverse energy or large τ signatures[1]. The main focus however is on the identification of events that contain detector signatures which are reconstructed in the offline analysis as jets. Therefore, this simulation study is focused on the detection of events containing offline jets. As the main goal of the trigger is to correctly identify events, this initial study also only considers complete events and does not differentiate between distinct objects. An exact mapping between trigger jets and offline jets is left for more detailed studies. The identification of the remaining objects and global values have also not been studied within the scope of this thesis. For further refinement, a direct mapping between trigger jets and offline jets can be used instead of the identification of complete events. This step is also left for more detailed studies. A brief overview of the implemented algorithms in Run 3 is given in 9.1.4.

9.1.1 Jet Algorithm Definitions

Two main jet definitions are used by the ATLAS community. Depending on the physics objects to be studied, anti-k_t4 or anti-k_t10 are chosen. The underlying algorithm for both versions is identical and described in section 2.3. As the iterative nature of the offline algorithms cannot be implemented within the tight time budget of the first trigger stage [79], simplified algorithms are required for the trigger.
The approach used for the jet trigger algorithm is based on the so-called sliding window algorithm as implemented in the JEP which moves a virtual window over the detector, calculates energy sums within the window, and compares the results to pre-defined thresholds and the results from adjacent window positions. A more detailed description of the JEP implementation and the first draft of the version intended for the jFEX in the Phase-I TDR can be found in section 4.4.1.1 and 6.3, respectively. The adaption of this algorithm for the jFEX utilises the finer input granularity and the coverage of larger detector sections per FPGA mainly to gain flexibility in the definition of window sizes and shapes allowing e.g. for approximations of round windows and overall larger windows to better fit the anti-k_t4 and anti-k_t10 offline jets.
The following description represents the implementation used for the simulations which are executed on regular computers. Therefore, details can differ compared to an optimised implementation on an FPGA that allows for efficient parallelisation but provides less flexibility. The results of both implementations are identical as the final version of the detailed studies are used to verify the implementation on the jFEX during the commissioning phase of the upgrade.
The simulated jet trigger algorithms are based on the algorithm implemented in the JEP as described in 4.4.1.1 while utilising the improved capabilities of the jFEX in terms of input granularity and processing resources. The main approach to the jet definition remains unchanged as a similar sliding window algorithm is applied where every trigger tower is used as the central tower. The energy values of all trigger towers within a defined area around the central tower are cumulated and compared to pre-defined thresholds and the sums around neighbouring central towers. Due to the finer input granularity, the jFEX allows for a higher flexibility in the algorithm definition. These initial simulations focus on quadratic window sizes in $\eta \times \phi$ where the

---

[1] The identification of large τ signatures using the jFEX was considered during the development phase but was later decided to not be included in the Run 3 trigger menu.

granularity allows it, as they offer the highest potential of resource- and latency-efficient implementation in an FPGA as demonstrated by the JEP. Within the scope of this thesis, window sizes of 9×9 trigger towers for small R jets and 17×17 trigger towers for large R jets in the central region were chosen corresponding to offline anti-k_t4 and anti-k_t10, respectively. Other shapes like an approximation of a round window and other sizes have been implemented to some extend but have mainly been left for more detailed studies later on. The focus of the simulation in the scope of this thesis lies on the anti-k_t4 algorithm and therefore the smaller window size. For the coarser and irregularly shaped regions towards higher values of $|\eta|$, a trigger tower is considered inside the window if the position of the centre of this trigger tower lies within the area that the window spans in the central region. This area measures roughly 0.9×0.9 in $\eta \times \phi$ for the smaller jet algorithm.

Except for the size, the simulated algorithms are designed analogously to the JEP algorithms. A smaller core area is defined around the central tower which is required to be a local maximum compared to surrounding core areas, following the JEP definition as illustrated in figure 4.8, in order to be considered a trigger jet. For this simulation, a core area of 3×3 trigger towers is used and compared to its eight direct neighbours. In more detailed studies, larger core areas as well as larger regions in which the core area of a trigger jet is required to be the local maximum may be considered. As a further parameter in the algorithm, a separate energy threshold for the core area, the seed threshold, can be defined and varied to optimise the trigger performance. For large area jets corresponding to the offline anti-k_t10 algorithm, additional parameters can be defined, like the existence of a substructure containing multiple core areas with a certain distance that each surpass a threshold and are identified as a local maximum. These extensions of the jet trigger algorithm, as well as the combination of multiple objects like it is already done in the topological algorithms are not discussed within the scope of this thesis but are left for more detailed studies. Without a fixed definition for these possible extensions of the jet algorithm, it is difficult to estimate the impact latency and resource consumption. At least for the versions discussed above, an extended algorithm would mainly utilise calculations, like the summation of the core areas, that are already included in the basic jet trigger algorithm and just add additional comparisons that could be done in parallel to the already existing comparisons. Therefore, for these variations the additional resource consumption would presumably be small with no expected impact on the latency.

For identified trigger jets, the sums over the complete windows are calculated and compared to thresholds that are defined as a part of the simulation study. A jet that surpasses this so-called trigger threshold is sufficient to trigger the read-out for the higher trigger stages and therefore constitutes a so-called single jet trigger[2]. Despite the name, this definition does not exclude events with multiple such jets. The η and ϕ value of a trigger jet is defined as the middle of the η and ϕ position of its central trigger tower.

9.1.2 Simulation Approach

In order to determine the performance of an algorithm in the ATLAS Trigger System, two separate sets of simulation data based on Monte Carlo simulations are required.

The first set contains so-called minimum bias events. The physical processes simulated in these

---

[2] The final comparison to the trigger threshold is not implemented on the jFEX but needs to be simulated for the study of the jFEX algorithms.

events are intended to have as few criteria applied as possible within the given conditions of e.g. colliding particles, centre-of-mass energy, and instantaneous luminosity. These samples are used to simulate the data taking period of the ATLAS experiment without applied triggers and can therefore be used to determine the trigger rate of an algorithm including a variation of its adjustable parameters. This allows the implemented algorithms to stay within the design parameters of the read-out system. For this thesis, a minimum bias sample containing 40000 events is used. For a detailed study, a significantly larger sample would be required to reduce the statistical uncertainties on the expected trigger rates.

The second set may consist of one or multiple signal samples. A signal sample represents a physical process that decays into objects for which the trigger algorithm is designed. In the case of the jFEX jet algorithm, this includes any particle that decays at least partially into hadronic showers which are reconstructed in offline analyses as jets. Within this thesis, a $t\bar{t}$-sample containing 8000 events was used as a signal sample. Like the minimum bias sample, a larger signal sample is required for detailed studies to reduce the statistical uncertainties and to allow for a better distinction between algorithm versions that are close to each other in terms of trigger performance. However, with a dedicated physics process as the basis of the signal sample, a higher proportion of events with high energetic jets is expected, which allows for smaller sample sizes in signal samples compared to minimum bias samples. The choice of $t\bar{t}$ production as the physics process was made based on the potential of $t\bar{t}$ events to either decay into small-R jets or into heavily boosted objects that can be reconstructed using the anti-k_t10 algorithm as well as the availability of this sample during the early development stages.

Both sets rely on Monte Carlo simulations to generate the underlying events in terms of created particles as well as the simulation of the detector response to these energetic particles. The code used for this analysis is based on an analysis framework developed for and by the ATLAS community based on ROOT [80]. For the demonstration of the separate steps to determine the performance of different versions of trigger algorithms, the standard 0.9×0.9 algorithm is used and compared to versions with different seed thresholds.

The overall trigger rate of the first trigger stage is limited to about 100 kHz. As the ATLAS experiment is designed to detect a broad variety of different particles and physics processes, the overall trigger rate is split up between different types of algorithms which will typically have a certain degree of overlap regarding which events are selected. For single jet triggers, an allowed rate of about 10 kHz is estimated as part of the trigger menu [48]. For a single jet trigger algorithm, the main performance indicator is its ability to identify high p_T offline jets while staying within the allowed trigger rate. The lower and sharper the border between identified high-p_T jets and non-identified low-p_T jets is, the better the performance of the trigger algorithm. This can be shown in so-called turn-on curves which are described in 9.1.3.3. To achieve this, the trigger algorithm is optimised to define jets as close to the offline jets as possible.

9.1.3 Jet Algorithm Simulation

9.1.3.1 Threshold Estimation

This leads to the trigger threshold as one of the main boundary conditions for jet trigger algorithms. Therefore, the first step when comparing different trigger algorithms or different configurations is typically the estimation of the threshold for each algorithm that results in the

same trigger rate. Given the effective bunch crossing rate of the LHC of about 32 MHz and the size of 40000 events in the minimum bias sample, a trigger rate of 10 kHz corresponds to ~ 13 events. Though not studied and discussed in detail, as an example of the impact on the overall performance of the algorithm from the choice of a single parameter, a variation of the seed threshold is shown in figure 9.1. The basic shape remains similar for all variations with a plateau starting at a transverse momentum of 0 GeV followed by a steep decline with increasing p_{T}. Yet, the height and width of the plateau strongly depends on the applied seed threshold as more events can be found in which no window position fulfils this requirement. This effect reduces the height of the plateau from the complete size of the sample of 40000 events with no seed threshold to below 70 events at the 70 GeV threshold. Please note that the resolution of the image combined with the logarithmic scale on the y-axis does not allow to determine whether the plateau in figure 9.1a - and to a reduced extend for the other versions - remains at the full sample size or slowly diminishes. The same effect also increases the width of the plateau as the minimum transverse momentum of any identified jet will be the seed threshold. Only in more detailed studies with applied energy corrections leading to potential negative energy values for individual trigger towers, it would be possible to construct jets below the seed threshold. For highest p_{T} jets, the impact of the seed threshold deminishes leaving no noticeable difference above 150 GeV in this sample. The derived p_{T} threshold for the different variations of the sliding window algorithm is indicated by the red line.
Through variation of this parameter, the trigger threshold ranges from 135 GeV for the standard version without a dedicated seed threshold to 120 GeV with a seed threshold of 70 GeV.

9.1.3.2 Control Plots

For a deeper understanding of the behaviour of the trigger algorithms and also as a cross check for their implementation in the simulation, control plots showing kinematic variables of the created jets can be used. For the following plots, only the jet with the highest p_{T} per event was used for each algorithm regardless of the trigger jet reaching the defined trigger threshold. The control plots containing the η, ϕ, and p_{T} distributions are shown in figure 9.2 and figure 9.3 for the trigger jets and the anti-k_{t}4 jets, respectively.
For the trigger algorithm, the histograms for η and ϕ have a significantly increased value in the first positive bin compared to their surroundings. This originates most likely from a binning error caused by different rounding rules for the borders between the bins and the position of the trigger towers. This leads to a single bin in the histogram representing two neighbouring rows of trigger towers while the remaining bins all represent a single row in ϕ and η for the η and ϕ distributions, respectively. Also, please note that statistical fluctuations appear larger due to offset on the y-axes for both ϕ distributions.
Except for the discussed bin, the ϕ histograms show a similar, flat distribution within the statistical uncertainties of the small signal samples for both algorithms. The η distribution for the trigger algorithm shows approximately the same shape as the anti-k_{t} algorithm with a maximum around $\eta = 0$ and falling slopes towards higher values of $|\eta|$. The distribution for trigger jets in regions of higher $|\eta|$ remains at a higher average value of about 40% of the peak value except for individual bins which are most likely caused statistical fluctuations while for the offline algorithm, the region around the peak appears to be slightly wider with only about 20% of the peak value on the sides. The reason for this difference, caused either by the

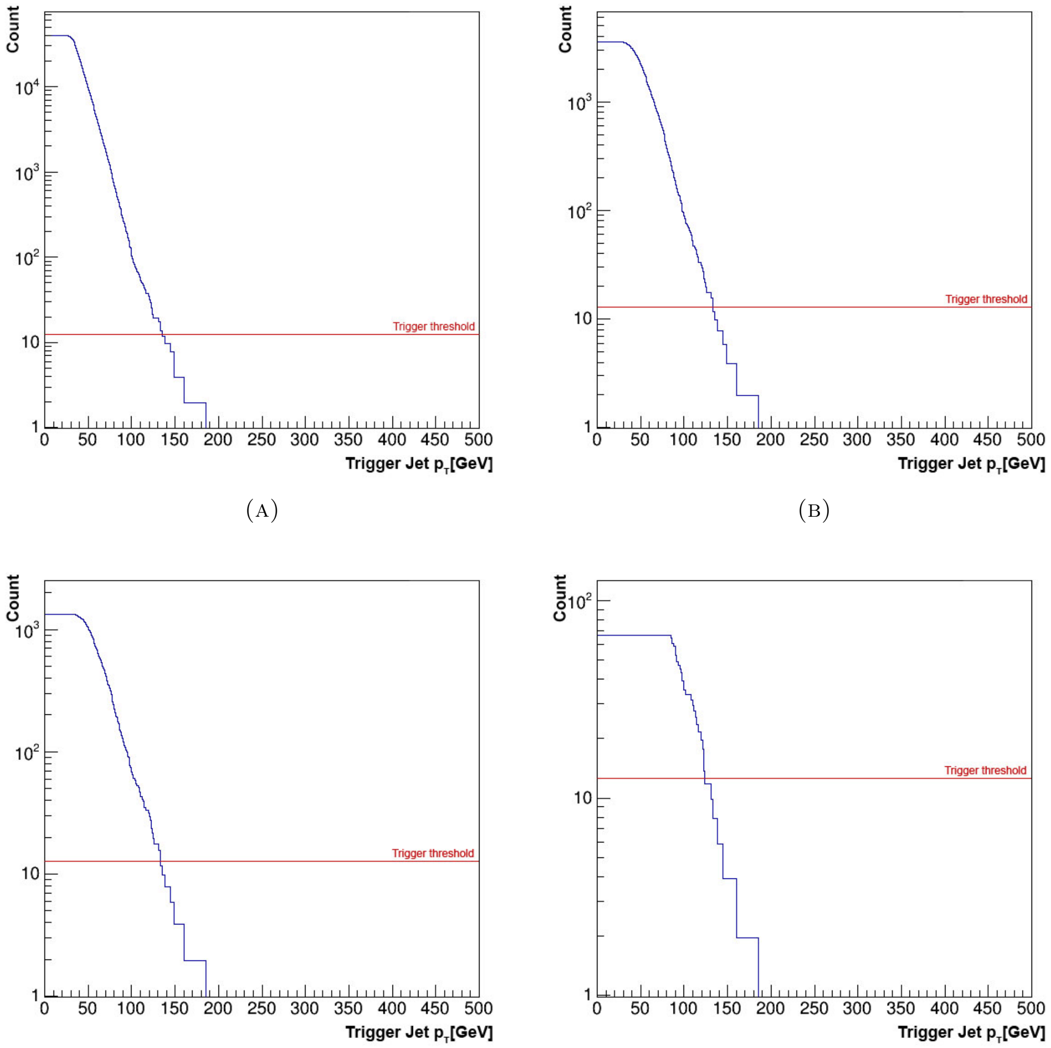

(A)

(B)

(C)

(D)

Figure 9.1: Number of events in minimum bias sample containing at least one small R jet above the p_T value for the basic sliding window algorithm (A) compared to an additional seed threshold of 20 GeV (B), 40 GeV (C), and 70 GeV (D) as a function of the p_T threshold. Simultaneously, the scaling of the y-axis changes with increased seed threshold as the plateau widens with increased thresholds but the total number of events with identified jets is reduced.

implementation or the different algorithms, could not be determined within this thesis.
The p_T plots show a difference in the distributions with the peaks around 45 GeV and 70 GeV for the trigger and offline algorithm, respectively. As the algorithms use different jet definitions and the offline jets are additionally calculated using fully calibrated calorimeter data, this effect is not unexpected. Also please note the different binning of 5 GeV for the offline algorithms compared to 10 GeV for the trigger algorithm.

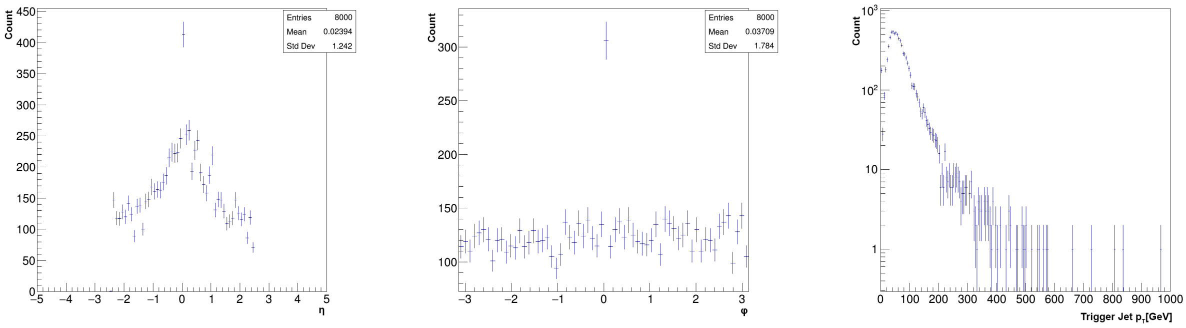

(A) η (B) ϕ (C) p_{T}

FIGURE 9.2: Distribution of η, ϕ, and p_{T} of leading trigger jets per event. The binning artifact in the η distribution and the offset on the y-axis of the ϕ distribution could not be removed due to missing source files.

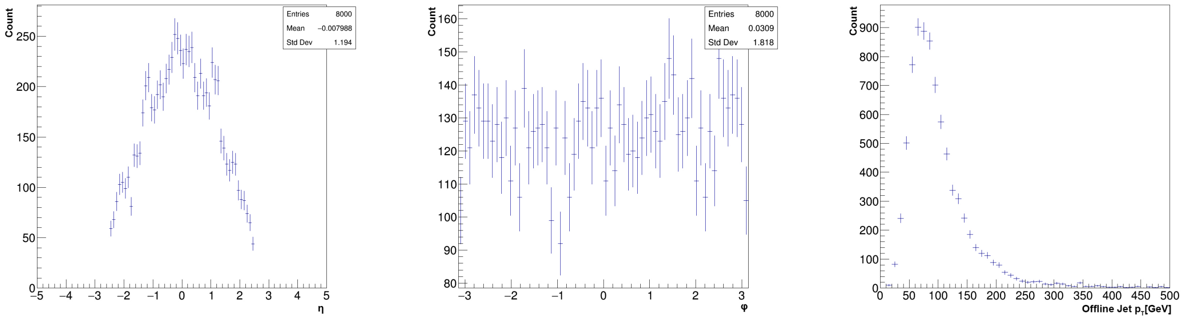

(A) η (B) ϕ (C) p_{T}

FIGURE 9.3: Distribution of η, ϕ, and p_{T} for leading jets based on the anti-k_{t} algorithm. The offset on the y-axis of the ϕ distribution could not be removed due to missing source files.

9.1.3.3 Algorithm Performance

Once the energy threshold for each version of the algorithm is found, this value can be used to determine the efficiency of the algorithms in terms of the identification of events using the $t\bar{t}$ signal sample. For each event, the anti-k_{t}4 jet with the highest transverse momentum is identified and its p_{T}-value is filled into a histogram. Depending on whether a certain algorithm identified a trigger jet above its individual threshold, a separate histogram is also filled with the transverse momentum of the offline jet. Afterwards, the ratio between the histogram containing the events identified by a trigger algorithm and the histogram containing all events gives the efficiency of this algorithm as a function of highest p_{T} offline jet. These so-called turn-on curves are shown in figure 9.4 for variations of the standard algorithm in terms of size and shape of the window without applying additional requirements like the seed threshold. Please note that the p_{T} values of trigger jets are calculated with uncalibrated data and for this reason, among other things, are typically lower than the p_{T} values of their corresponding anti-k_{t}4 jets.
This plot shows a wide range of different configurations of the sliding window algorithm with sizes of 7×7 and 9×9 trigger towers as well as approximations of round windows with radii of $\Delta R < 0.3$ and $\Delta R < 0.4$. Additionally, implementations of the legacy trigger algorithm from the Run 2 trigger system are implemented to compare the capabilities of the jFEX system with

the JEP. Due to a lack of a data type representing the exact input of the legacy system, the implementation of the JEP algorithm was not exactly identical to the legacy trigger system but rather recreated using pre-summations over the jFEX input data. This was done using two approaches verifying each other. The first approach used initial pre-summation of the entire input creating an input format similar to Run 2 data. The second approach, meant to increase the performance of the simulation code, incorporated the summation in the calculation of the jFEX algorithms. In order to estimate the impact of the finer positioning of the jets rather than their window size, an additional version of the JEP algorithm was implemented using the same dimension but utilising the full granularity of jFEX input data. In this simulation setup, no additional criteria like seed thresholds were applied to the algorithms.

The performance of a trigger algorithm is mainly evaluated using the turn-on curves where the ideal trigger algorithm would be represented by a Heaviside step function. Desirable features in the turn-on curve of an algorithm are a steep slope at a comparatively low p_T value with few triggered events at lower p_T and a stable plateau with a 100 % efficiency towards higher transverse momentums. A steep slope in the turn-on curve is reached if the trigger algorithm identifies the same objects as jets as the offline algorithm while also keeping the ratio of the p_T for different jets similar. The absolute p_T values are not directly relevant for the performance of the trigger algorithms as they contain an offset due to missing energy calibrations applied to offline data sets. The final evaluation of trigger algorithms and their configurations is made by representatives of the ATLAS community by defining the trigger menu. As other aspects such as physics programme and limitations in the following trigger levels are also accounted for in this decision, no strict criteria for algorithm performance can be defined. Overall, trigger algorithms with a smaller window size like 7×7 trigger towers or the round windows showed a better performance than larger window sizes such as the version using 9×9 trigger towers as shown in figure 9.4. This is assumed to be an effect of the significantly lower trigger thresholds possible for the smaller sizes. Both round windows also show an increase in trigger performance by lowering the offline event identification threshold by about 10 GeV compared to the JEP algorithm. Potential issues with the implementation of the legacy sliding window algorithms using the Phase-I input data can be seen in the large discrepancies between the different versions of the JEP algorithm. The origin of this could not be clarified due to the loss of the simulation code. Still, it can be assumed that the standard version of the JEP algorithms shown in figure 9.4 as green circles is closest to the actual trigger algorithm of the Run 2 system, as it is the most straightforward implementation of the three versions using the separate implementation on the pre-processed input data.

Whether the smaller window sizes used by the algorithm running on 7×7 trigger towers and the round windows lead to decreased performances in other signal samples or deficiencies in the plateau region of the turn-on curve could not be clarified within the scope of this thesis. All trigger algorithms show a deficiency around 300 GeV where most likely a single event with a relatively high p_T offline jet was not identified. Due to the loss of the simulation code, the origin of this deficiency could not be clarified. Please note that due to the small sample size and the resulting statistical uncertainties, a reliable evaluation of the performance of the trigger algorithms is not possible. Results from more detailed studies comparing the capabilities of the legacy trigger system with the Phase-I upgrade are shown in figure 9.7.

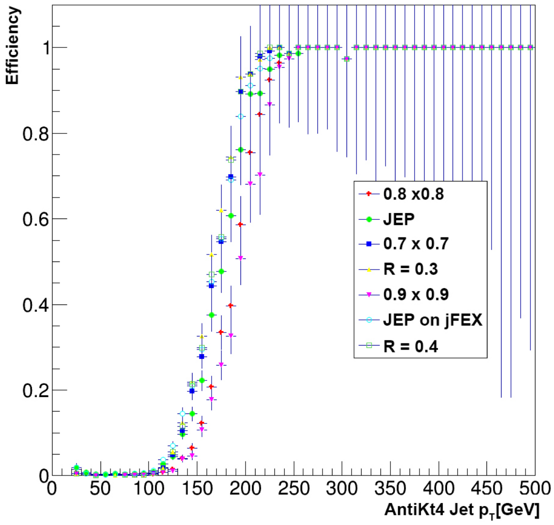

FIGURE 9.4: Identification efficiency of events as a function of their highest p_{T} anti-k_{t}4 jet for different sliding window algorithms in the $t\bar{t}$ signal sample normalised to an expected trigger rate of 10 kHz. Results from sliding window algorithms with a window size of 7×7 and 9×9 trigger towers are marked as blue squares and pink triangles, respectively. Approximations of round windows with radii of $\Delta R < 0.3$ and $\Delta R < 0.4$ are represented by yellow triangles and green non-filled squares, respectively. An implementation of the Run 2 jet trigger algorithm using Phase-I input data is shown as green circles with an alternative implementation approach shown as teal non-filled circles. An additional version of the legacy algorithm using the same window size while utilising the increased resolution of the jFEX system is marked with red crosses.

9.1.3.4 Large R Jets

The same simulation approach can also be used for the large R sliding window algorithms used as a trigger for the offline anti-k_{t}10 jets. Without further refinements of the algorithm, the larger window sizes lead to significantly higher thresholds compared to the smaller window sizes. As shown in 9.5, the threshold corresponding to an approximate trigger rate of 10 kHz for the 11×11 and 17×17 sliding window algorithm lies around 175 GeV and 380 GeV, respectively. These high thresholds leave too few remaining events in the available signal sample for further studies of this algorithm due to prevalent statistical uncertainties as shown in 9.6. Therefore, in-depth simulations of trigger algorithms for large R offline jets including the use of substructure requirements to reduce the trigger threshold are left to more detailed studies.

9.1.3.5 Conclusion of Preliminary Simulations

Even with the limited statistics that were available in the early development stages, it is possible to demonstrate that the jFEX setup is capable of running suitable jet trigger algorithms for the ATLAS experiment. Additional use cases for the jFEX jet algorithms, like the combination of multiple jets or with other objects or quantities like muons or missing transverse energy in

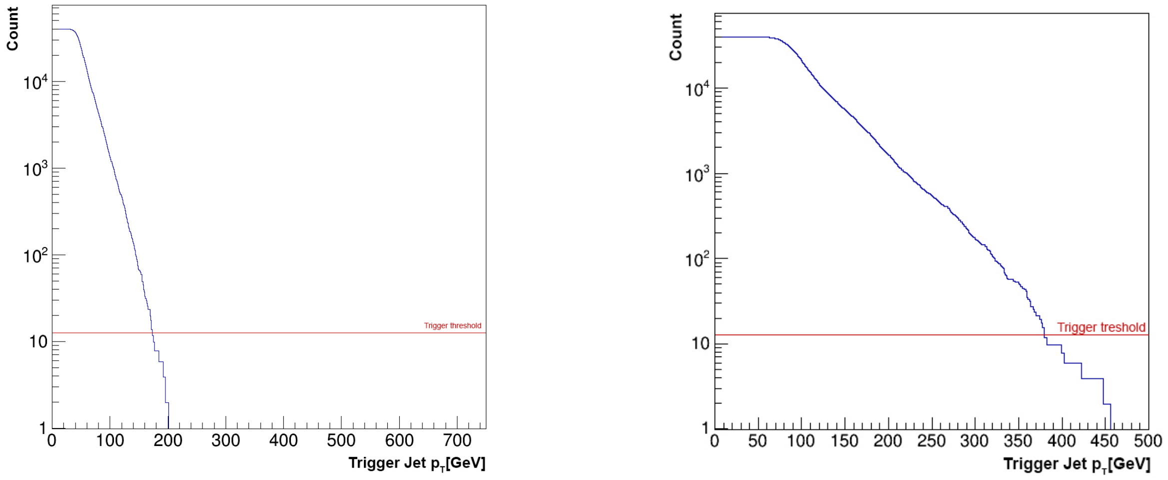

(A) 11×11 (B) 17×17

Figure 9.5: Number of events in minimum bias sample containing jets above threshold for the large sliding window algorithm of 11×11 and 17×17 trigger towers.

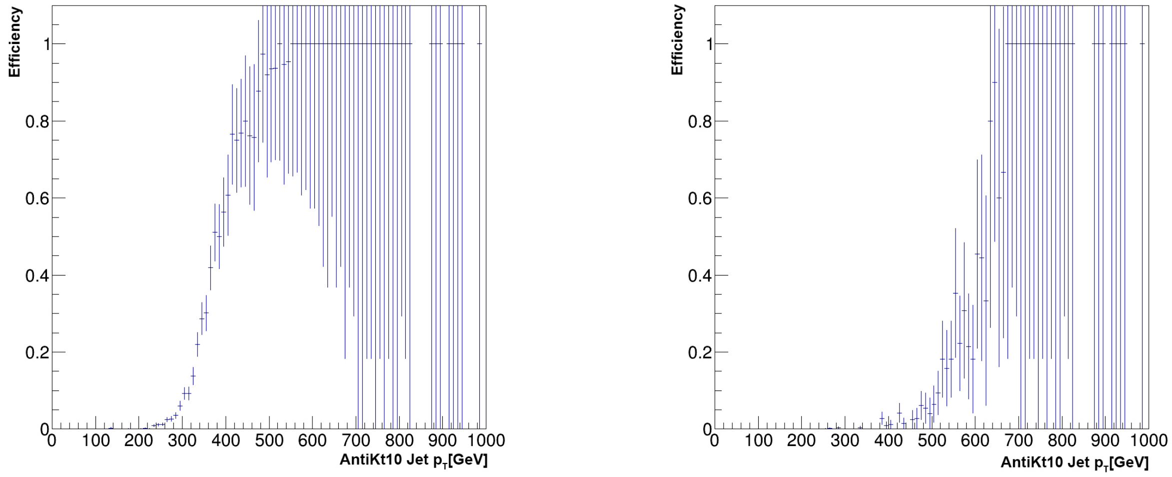

(A) 11×11 (B) 17×17

Figure 9.6: Trigger efficiency of events as a function of their highest p_{T} anti-k_{t}10 jet for window sizes of 11×11 (A) and 17×17 (B) trigger towers in the $t\bar{t}$ signal sample without additional criteria like substructure applied. While the sample size is sufficient for the 11×11 to show an expected turn-on curve, the version using the larger window size is completely dominated by a lack of selected events. Though no exact prediction is possible based on these results, the smaller sized windows show a better performance in this signal sample.

topological algorithms, have not been studied within the scope of this thesis as these studies require detailed simulations for different object types as well as larger signal samples of different physics processes.

For a short discussion of the full potential of the jFEX, an overview of the implemented trigger algorithms for the start of Run 3 is given in the following section. Due to these detailed studies, the suitability of the jFEX hardware design, as presented in this thesis, to serve as a part of the ATLAS Level-1 Trigger System can be shown regardless of the mentioned loss of simulation code and data samples which limited the extend of the studies shown above.

9.1.4 Implemented Algorithms in Run 3

To further demonstrate the capabilities of the jFEX, this section provides a brief overview of the currently implemented algorithms for its use in the ATLAS Trigger System during Run 3. During the first year of Run 3 in 2022, the legacy trigger system was still used for data acquisition as the Phase-I system remained in commissioning with plans to replace the legacy system during the 2023 data taking period after further tests of the functionality of the Phase-I system [81]. By mid 2024, the new Level-1 Trigger System had replaced the legacy system in the data acquisition functionality for the ATLAS experiment. The legacy system is meanwhile operated in parallel while minor optimisations of the Phase-I firmware in terms of performance are being implemented [82]. Descriptions of the algorithms below are based on the thesis of Marcel Weirich [78] with updates regarding the final implementation as a part of the ATLAS Trigger System [83].

As intended in the original design, the jFEX runs algorithms for the calculation of energy sums and the identification of discrete objects. The final version of the jet trigger algorithm utilises the finer input granularity of the jFEX by implementing a round window geometry of $\Delta R < 0.4$ using the definition of ΔR from section 4.1. The larger window size has been left out due to resource constraints in the FPGAs prohibiting its implementation alongside all other features and was replaced by a combination of small R jets and a reclustering algorithm in the Topological Processor. The implemented jet algorithm uses a seed of $\Delta R < 0.2$ which is required to be a local maximum compared to other potential seeds in an area of 5×5 trigger towers for the central region. The option to use a fixed seed threshold has not been implemented in the Run 3 algorithms.

Figure 9.7a shows the per-jet efficiency[3] for jFEX jets, JEM jets, and an implementation of the anti-k_t algorithm based on trigger towers where the legacy system overall shows a comparable performance to the jFEX. For close-by jets, where another offline jet with a transverse momentum of at least 30 GeV within $\Delta R < 0.6$ and at least three trigger jets in the event are required, the jFEX shows a significant increase in performance over the JEM algorithm as shown in figure 9.7b. This advantage in the identification of close-by jets can be further utilised in combination with the Topological Processor.

As an extension and improvement to the eFEX, the jFEX also runs a forward electron algorithm in the region of $|\eta| > 2.3$, where the eFEX coverage of the calorimeter ends, and provides isolation information for τ objects that require a larger window size than what the environment of the eFEX supports. Either algorithm utilises only the input data received from the electromagnetic calorimeters. The isolation information and the eFEX τ objects are combined on the L1Topo. Both algorithms are not discussed in detail in the scope of this thesis as they are mainly used in combination with the eFEX.

As of mid 2024, the Level-1 acceptance rates based purely on jet algorithms implemented on the jFEX average around 4 kHz [84] for single jet triggers and 9 kHz for a multi jet trigger providing at least four objects in a limited η-range. The lowest trigger thresholds used without rate limitations are set to 160 GeV for the single jet trigger. The multi jet trigger requires one

---

[3] In contrast to the simulation studies carried out in the scope of this thesis which used per-event efficiencies.

object with at least 85 GeV and $\eta < 2.1$ as well as three objects above 40 GeV within $\eta < 2.5$. Additional trigger algorithms implemented on the jFEX including missing transverse energy calculations, τ objects, and dedicated forward jet triggers are not discussed in detail within the scope of this thesis.

The calculation of energy sums consists of a scalar value of total transverse energy E_{T} and a vector quantity of missing transverse energy $E_{\mathrm{T}}^{\mathrm{miss}}$, split into two scalar values in the direction of the x-axis and y-axis of the ATLAS detector as defined in section 4.1. All energy sums are calculated by each jFEX module over its respective core region and transferred to the L1Topo for the final summation. For a more detailed description of the algorithms, please refer to the original thesis [78]. The performance of two versions of the $E_{\mathrm{T}}^{\mathrm{miss}}$ algorithm containing noise cuts (NC) or pile-up subtractions (PUS) compared to the Run 2 algorithm as implemented on the JEM as a function of the efficiency of identifying events with a high truth $E_{\mathrm{T}}^{\mathrm{miss}}$ in simulated events is shown in figure 9.7c.

The implemented data formats for Run 3 TOBs mainly follow the structure shown in figure 6.4. As shown in table 9.1, the overall size has since been reduced to 32 bits for all TOB types. The local coordinates in η and ϕ remained identical to the initial draft. All TOB types use 1 bit to indicate the saturation of input towers used in their calculation. Energy values have been reduced to 13 bits for large R jets and 11 bits for all other objects. For jets, the remaining bits are left as spare. Electrons use 6 bits for various information on isolation and fraction of the energy deposited in the electromagnetic parts of the calorimeters, leaving 5 spare bits for future use cases while τ objects use all remaining 11 bits to indicate isolation. No information on the TOB types is required as they are identified by being sent on fixed positions on their respective fibre.

logic resource	jets	large R jets	τ	electrons
η	5 bits	5 bits	5 bits	5 bits
ϕ	4 bits	4 bits	4 bits	4 bits
E_T	11 bits	13 bits	11 bits	11 bits
Others	-	-	11 bits	6 bits
Spare	11 bits	9 bits	-	5 bits
Saturation	1 bit	1 bit	1 bit	1 bit

Table 9.1: Overview of the jFEX TOB data formats in the implementation for Run 3. Input taken from [78].

9.2 Initial Capability Estimations

This section provides an overview of estimations regarding the capabilities of the design on various aspects that were calculated during the design phase of the first prototype and therefore before direct hardware tests with the jFEX design were possible. These were a part of the design phase to avoid producing a prototype that does not meet the requirements or cannot be operated within its expected range. An evaluation of the required logic resources for the implemented functions is presented in section 9.2.1 followed by a power estimation calculating the expected electrical currents for the main internal supply voltages shown in section 9.2.2.

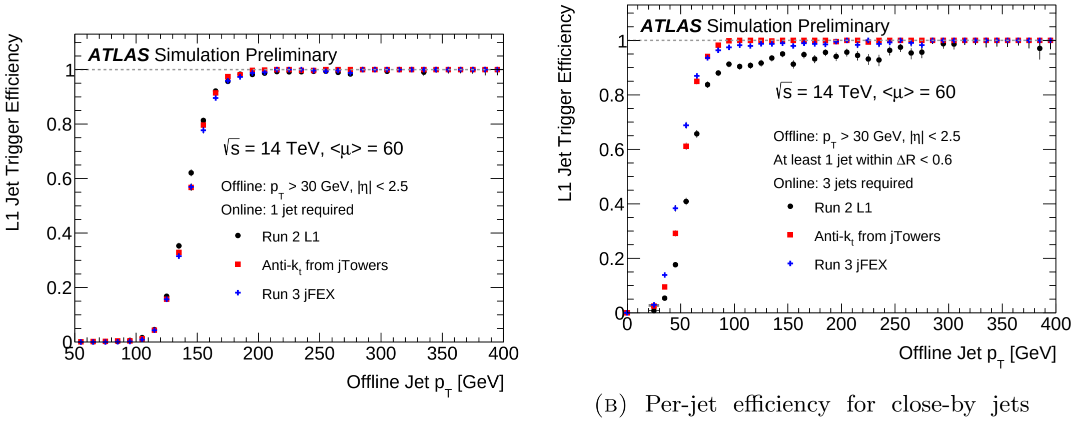

(A) Per-jet efficiency for small jets against offline anti-k_t4 jets.

(B) Per-jet efficiency for close-by jets with another jet identified within $\Delta R <$ 0.6.

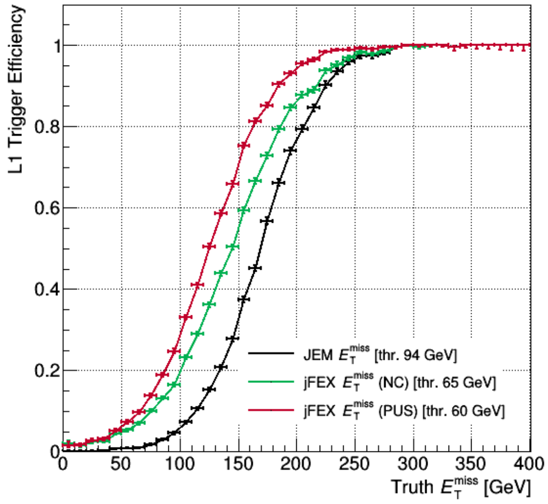

(C) Per-event efficiency for the E_T^{miss} trigger.

FIGURE 9.7: Turn-on curves derived from HH → bb(bb) Monte Carlo simulations. The trigger thresholds of all shown algorithms have been tuned to have the same rate. Plots taken from [85].

9.2.1 Implementation of Jet Algorithms

The sliding window algorithm in its implementation as it is in operation in the JEP is highly optimised for this use case in order to reduce resource consumption and latency by minimising the number of additions. This is achieved by performing the addition step by step for every possible position concurrently and using the pre-added sum from a suitable different position in the next step, as illustrated in figure 9.8a. This leads to a doubling of the window size in each step, allowing the algorithm to reach the targeted size in a minimal number of steps while also minimising the required resources. This concept is relatively easy, yet the exact implementation in hardware is already a highly complex combination of cascaded LUTs. Therefore, only the overall utilised logic resources per algorithm are mentioned in this section to provide an overview of the capabilities of the system instead of demonstrating the exact implementation in hardware.

The following descriptions of the algorithms for the jFEX focus on single FPGAs as the con-

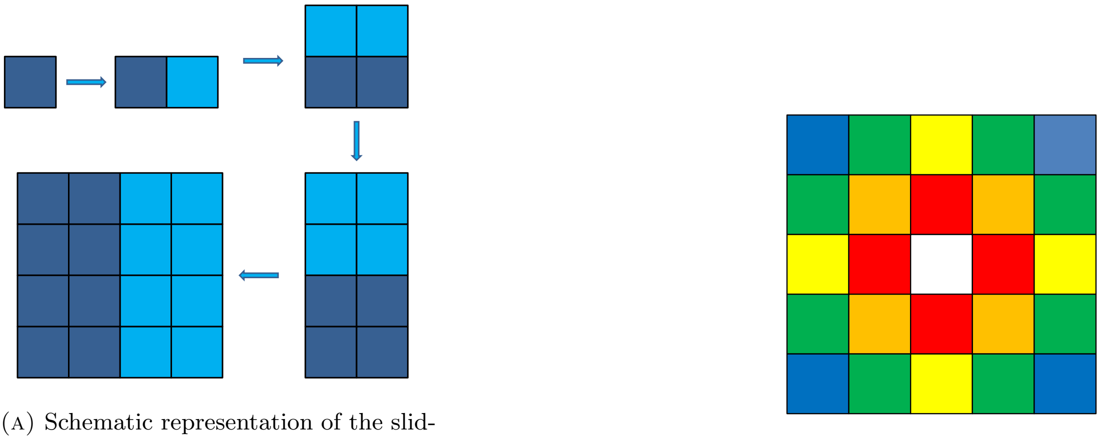

(A) Schematic representation of the sliding window algorithm implementation in the JEP utilising pre-added sums (light blue) in each step. For this system, the maximum size of 0.8 × 0.8 in $\eta \times \phi$ can be reached within four steps, based on 0.2 × 0.2 ($\eta \times \phi$) input granularity.

(B) Schematic representation of a sliding window algorithm with distance-dependent weighting. Towers weighted with the same factor are in the same colour.

FIGURE 9.8

straint for resource availability. Based on the mapping of input data to the modules, as described in section 8.2, the overall highest resource utilisation is expected in the modules covering the central detector regions with both, their core and environment. For these FPGAs, the core region covers a total of 128 towers, requiring the same calculation for each position in parallel. For the Gaussian weighted algorithm, each trigger tower in the window is multiplied with a weighting-factor. Since the distribution of these weights over the towers is unique for each central tower, an efficient way of sharing resources with those required for additions from neighbouring towers, as it is used in the basic sliding window algorithm, is not possible.

Multiplications are most efficiently done in DSP slices. The highest resource consumption is expected for the largest window size of 17 × 17 towers, calculated for up to 128 central towers. This leads to a total of 36864 multiplications that are required, assuming that the central tower has a weight of 1. Since there are only 1800 DSP slices available in the XCVU190, this cannot be implemented even at higher clock rates without optimisations. As there is no easy way of combining calculations required for different central towers, the following optimisations are intended to be done for every window position separately unless stated otherwise.

A first reduction in the number of utilised DSP slices can be achieved by adding up all towers that are multiplied with the same weight. Due to the regular structure of the input data from the central detector regions, at a given distance around every central tower, at least four towers with the same distance can be found. A schematic representation of this scheme is shown in figure 9.8b. For the maximal window size of 17 × 17 towers, this reduces the number of multiplications from 288 to 42 per window position. With an increased clock rate of 320 MHz all multiplications can be done in 6 DSP slices per central tower. For the entire core region covered by a single FPGA, 768 of these elements are required corresponding to an estimated use of 42.7 % of this logic resource. For a smaller window size of 9 × 9, 12 multiplications are required which can be done in 2 DSP slices or 14.2 % of the DSP slices per chip. This optimisation is already included in the initial implementation of the algorithm as used in 9.2.2. Optimisation for the outer regions with a coarser granularity are not discussed in the scope of the thesis as

with a reduced granularity, the number of trigger towers in a weighted sliding window is already reduced and therefore considered to be less critical than the central region.
Further optimisations can be made depending on the exact algorithm. For instance, depending on the definition of the local maximum condition there can be only a single trigger tower within a certain area whose surrounding window is considered a jet candidate in any event. Therefore, the core region can be divided into blocks of a size depending on the local maxima definition. With some multiplexing on the input side, the usage of DSP slices can thereby be reduced to one set per block. Since this can only be done once the result from the maximum finding algorithm is available, this step could significantly increase the latency. Hence, the possibility of implementing this reduction of resources depends on the remaining algorithms and might only be suitable for a system with a longer available processing time such as the first trigger stage after the Phase-II upgrade.
As preliminary versions of the trigger algorithms were implemented by Marcel Weirich [83] during the design phase of the jFEX, their utilisation of logic elements in the XCVU190 could be used to indicate their suitability for the complete set of algorithms. Due to a lack of available MGT firmware, BRAM blocks are used to simulate inputs and outputs in early versions of the firmware. Additional flip-flops are used around the DSP slices in the weighted algorithms to simplify the routing by storing the signal for the next clock cycle even though this might not meet the timing requirements for the final firmware. An IPbus firmware, written by Christian Kahra [86], is included based on the version used by the Run 2 L1Topo. The resource consumption per algorithm is shown in table 9.2. So far, the basic sliding window algorithm is only implemented in the simplest version without any optimisations to estimate an upper limit of the resource allocation for these algorithms. Therefore, the utilisation is likely to decline in the final version, making spare resources available for new algorithms or further improvements of the basic versions. With preliminary versions of the implementation of the full-scale jet algorithms in a basic version, which have the highest resource consumption of all so far known algorithms, already available and functional on the targeted FPGA, it can be assumed that the chosen FPGA meets the requirements on processing capabilities.

logic resource	weighted 9 × 9	weighted 17 × 17	basic SW
LUT	79740 (7.4 %)	262609 (24.4 %)	37706 (3.5 %)
flip-flops	61544 (2.9 %)	181784 (8.5 %)	0 (0 %)
BRAM	24 (0.6 %)	168 (4.4 %)	20 (0.5 %)
DSP	256 (14.2 %)	768 (42.7 %)	0 (0 %)
logic resource	IPbus	E_T	E_T^{miss}
LUT	5320 (0.5 %)	1237 (0.1 %)	1012 (0.1 %)
flip-flops	10035 (0.5 %)	0 (0 %)	0 (0 %)
BRAM	130 (3.4 %)	0 (0 %)	0 (0 %)
DSP	0 (0 %)	0 (0 %)	0 (0 %)

TABLE 9.2: Resource consumption of preliminary implementations of jFEX algorithms and infrastructure, implemented by Marcel Weirich and Christian Kahra [83, 86]. Relative consumption is calculated as proportion of the total resources available in the XCVU190 FPGA.

9.2.2 Power Estimation

A first test was done utilising the VCU110 development kit during the jFEX design phase using an implementation of a simple version of the Gaussian weighted sliding window algorithm done by Marcel Weirich. This evaluation board hosts the same FPGA as used in the jFEX design just in a different package which should have no impact on the power consumption of the internal logic. The implementation is varied in the number of calculations done in parallel which scales linearly with the number of utilised DSP slices. For each implemented design, two predictions from Vivado on the V_{CCINT} current was given, assuming 50 % and 75 % toggle rate of the randomised input data being read from BRAM blocks. The power estimate from Vivado is assumed to be based on the same calculations as the Xilinx Power Estimator referenced in the design phase. The actual current is read out via PMbus from the onboard power supply, compared to the predictions and extrapolated to 100 % DSP slice utilisation. As extrapolation functions a linear function and a quadratic function are chosen. The linear function represents the assumption that the power consumption increases linearly with the number of resources in use. The quadratic function is used to simulate the possibly increasingly complex designs where the usage of other resources grows faster than for DSP slices due to a more difficult and wide-spread routing when larger parts of the FPGA are utilised. The measured and estimated electrical currents are shown in Table 9.3 with the extrapolations in figure 9.9.
The overall agreement between prediction and measurement with the measurements lying between the expected values for 50 % and 75 % toggle rate seems to confirm the assumptions. As an average toggle rate of no more than 50 % is expected, this allows for the assumption that the power estimation of Xilinx underestimates the actual consumption of the FPGA at least for versions of the tools used during the early development which were among the first versions

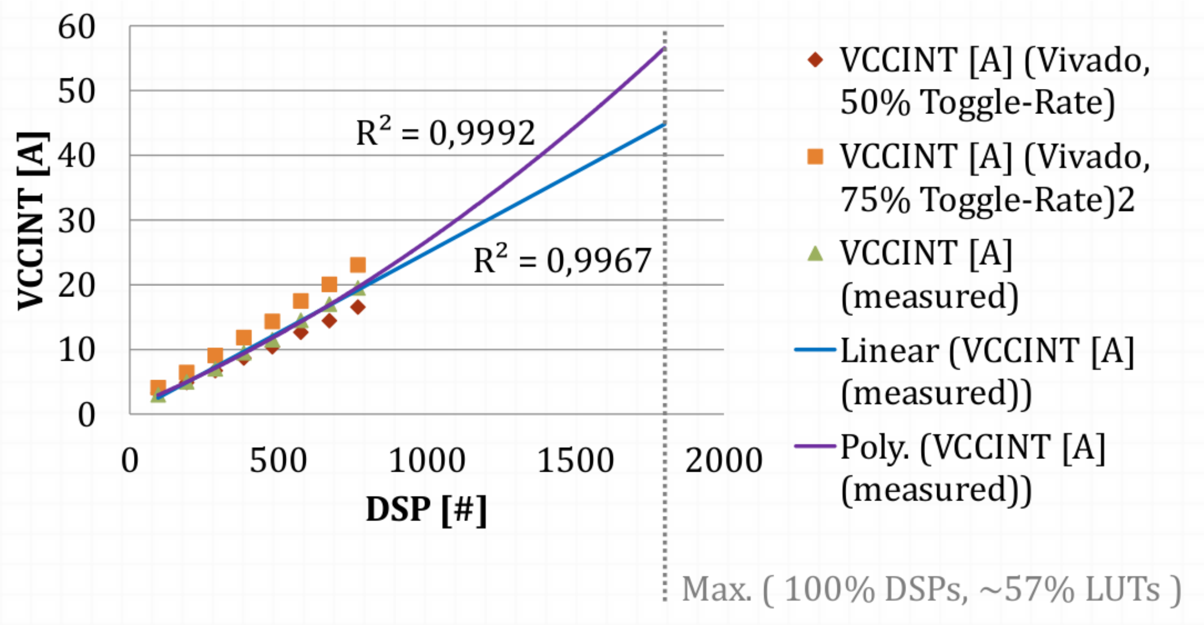

FIGURE 9.9: Measured currents on V_{CCINT} in dependence of the number of utilised DSP slices by the design against the expected currents by Vivado. Linear and quadratic functions are used to extrapolate the current to 100 % utilisation. Tests were done using the VCU110 evaluation board.

# used DSP slices	prediction 50 % rate [A]	prediction 75 % rate [A]	current measured [A]
96	3, 3	4, 1	3, 0
192	4, 9	6, 4	5, 0
288	6, 7	9, 1	7, 0
384	8, 7	11, 8	9, 5
480	10, 4	14, 3	11, 5
576	12, 7	17, 5	14, 5
672	14, 4	20, 0	17, 0
768	16, 5	23, 1	19, 5

Table 9.3: Measured currents on V_{CCINT} for the Gaussian weighted sliding window algorithm using randomised input data for different numbers of utilised DSP slices compared to the expected currents as predicted by Vivado for toggle rates 50 % and 75 %.

to support the new UltraScale family[4]. The extrapolations reach 45 A and 57 A at 1800 DSP slices for linear and quadratic functions, respectively. Based on the ratio of utilised LUTs and DSP slices in the firmware used for these tests, full utilisation of DSP slices corresponds to 57 % of LUTs in use, assuming a constant ratio between the two resources.
As routing within the FPGA will most likely prevent full utilisation of the resources, the expected maximum current drawn from V_{CCINT} is assumed to remain below the conservatively estimated 57 A. This could e.g. be provided by the TDK-Lambda IJB-series which support currents of up to 60 A. With power supplies like this already available in the prototype phase and more powerful parts announced, this part of the jFEX design appears to be in-line with the requirements.

9.3 Hardware Tests

This section describes the tests of the high-speed data connections while focusing on aspects related to PCB design and layout. The described tests were done using the first prototype as shown in figure 9.10b and figure 9.13. Stand-alone tests on this prototype and conceptual changes between this first iteration and the general jFEX design as described in section 8 are documented in section 9.3.1. The integrated tests at CERN with the LATOME prototype and the FEX test module are described in section 9.3.3. The firmware for all high-speed link tests described in this thesis was created based on Integrated Bit Error Ratio Tester (IBERT) IP cores as provided in the Vivado software suite. For the integrated tests at CERN, version 2016.4 was used for firmware creation as well as monitoring and evaluation of the tests. The local tests targeting the PMA loopback scheme, as described in section 9.3.2.2, used the version 2017.1 which was published after the integrated tests.

9.3.1 Prototype

The first jFEX prototype has some design differences to the full-scale design described in section 8. These adjustments were done without altering the PCB design to avoid changing the

---

[4] Xilinx regularly updates the provided tools which might lead to a better agreement with more recent versions.

exact layout tested on this prototype compared to later iterations. This was intended to make test results from a smaller prototype directly applicable to a full-scale board. One of the reasons for this step is the high price for the large processing FPGAs. Also, by the time of the decision, the FPGAs were only available as engineering samples[5] in limited quantities. Therefore, it was decided to mount only two of the FPGAs, U1 and U2, on the board. This setup still allows for testing all critical connections including PMA loopback functionality between the FPGAs.

Besides the reduced number of FPGAs, there are some further changes to the design. Since the prototype is expected to be used only in relatively small setups compared to the complete L1Calo system in operation, the module can in any case be built with double slot width. This simplifies the design of heat sinks on the FPGAs which can use a higher profile to guarantee sufficient thermal conductivity. This is especially useful for test setups where no ATCA crate with standardised cooling system is available. Furthermore, the extension mezzanine can be built with a higher connector allowing airflow between the two PCBs to avoid any interference with the cooling of U1. If the setup as a part of the L1Calo Trigger does not allow for double slot width, the different prototype stages can also be used to test heatsinks that provide the required cooling capacities in single slot width.

With U3 and U4 not assembled, the extension mezzanine can also be designed larger in its first iteration and overlap with U3 or the corresponding MiniPODs. This simplifies the design and allows for the use of a PicoZed board which hosts, alongside the Zynq, required for control purposes, many components that are not required in the application on the jFEX. Yet, as this board comes as a commercially available component, it can be used to reduce development time for a first functional setup.

During the purchase of electrical components, it was not possible to purchase the clock generator Si5335, planned as a frequency-flexible onboard clock source, in a version without predefined frequency plans or with one suitable for jFEX operation. Therefore, this component is present in the schematics but was not assembled and will be removed in later iterations.

After the production of the PCB in late 2016, the manufacturer and the assembly company noticed certain faults on the board. These were not severe enough to exclude the possibility that the assembled board would be functional but a guarantee for this could not be given by either company. Therefore, two options were presented to choose from: waiting for a new production of the PCB or ordering of assembly at own risk. Due to a tight deadline for the integration tests with other groups at CERN, it was chosen to order the assembly of the board, but further reduce the number of assembled FPGAs to one. This allowed at least for the test of possible link speeds as the jFEX was the last system to take the tests that should determine the link speed used in the Phase-I Level-1 trigger. However, the intended tests of the PMA loopback functionality between neighbouring FPGAs required an additional adaptor as described in 9.3.2.2.

For all initial link speed tests, the configurable firmware called Integrated Bit Error Ratio Tester (IBERT) cores provided by Xilinx were used which offer direct read-out of results via Vivado Design Suite, which was used in version 2017.1. A pseudorandom binary sequence as provided by IBERT as PRBS31-bit without 8b/10b encoding was used for the link speed tests. This sequence allows for a correctly configured receiver to recognise the received pattern without additional information. 8b/10b encoding is not supported within IBERT. As the main advantages of the encoding scheme for the jFEX system are the improved clock recovery and DC balance,

---

[5] Engineering samples are early access beta versions of integrated circuits, which can be used in development processes for new designs or as demonstrators but are not yet guaranteed by the manufacturer to be fully functional.

which both are intended to stabilise the connection, it was considered that for successful transmission tests the results will be applicable to encoded data as well. For unsuccessful tests, a rerun using encoded data would be required to correctly assess the link quality. Further tests with custom firmware including encoding and different data sequences were performed at a later stage of the development.

The official design goal agreed upon by the Level-1 trigger community for real-time data transmission is a bit error ratio (BER[6]) of 10^{-14} defined as the number of incorrectly transmitted bits divided by the total number of transferred bits. To achieve this with a sufficient margin, tests for the verification of the jFEX design aimed at a measured BER of 10^{-15} which provides a 95 % confidence level for a BER of about 3.3×10^{-14}.

9.3.2 Stand-alone Tests in Mainz

Upon arrival of the prototype in Mainz in early 2017, the remaining assembly steps and initial tests were done by Bruno Bauss. Required assembly includes a custom heatsink for the FPGA built in-house, that can easily be attached and removed using thermally conductive tape, and several quartz oscillators. The first test of this series is usually a so-called smoke test where an electrical circuit is connected to a power supply for the first time to check for major faults in the wiring. To exclude errors coming from the onboard power supplies, this test was done using separate lab power supplies for each required voltage while manually sequencing those. Afterwards, a boundary scan is performed where connections between different devices, which support this test, can be verified. For the reduced prototype, this test cannot be carried out to the usual extent as there is only a single FPGA present and therefore the communication lines between different devices cannot be tested, leaving only the connection between the mounted processing FPGA and the mezzanine as test subjects. For further tests, the board was equipped with a preliminary power mezzanine to simplify the power input by only requiring two different voltage levels. As the intended power converters for the high current applications of the FPGA were not yet available, smaller versions have been used which limited the possible utilisation of FPGA resources in the coming tests.

After successful verification of the basic functionality of the PCB and assembled components, the board was cleared for further tests. These tests included two in-house tests presented in 9.3.2.1 and 9.3.2.2, respectively, and the integrated tests at CERN discussed in 9.3.3.

9.3.2.1 General Functionality

First functional tests with the prototype board were done in Mainz in a stand-alone setup using the loopback functionality of the MGTs. The main goal of this test was to ensure the general operationality of high-speed links before setting up more extensive tests at CERN together with prototypes from the actual data sources in the L1Calo system.

Due to the limited output capabilities of the jFEX design with only a single 12-way transmitter per FPGA, the input in this test was also limited to a single receiver MiniPOD at any time.

---

[6] In this thesis, the term is used to describe the bit-error ratio. This should not be mistaken as the bit error rate, a measure of bit errors per unit time, which is also commonly described with the same acronym but includes a dependency of the transmission speed.

Therefore, switching between five setups to cover every position of MiniPOD receivers is required. This further increases the time needed to reach low BERs. This first test only aims towards proving that all links are generally operational which could successfully be verified. As the remaining tests carried out later on provide more information on all aspects of data transmission, the results of this test are not further discussed in the scope of this thesis.

9.3.2.2 Loop Back Test

The second set of local tests was done in October 2017 after the integrated tests at CERN. The scope was the evaluation of the data duplication scheme based on PMA loopback. To compensate for the lack of neighbouring FPGAs, a passive board was designed by Bruno Bauss as an adapter to connect the outgoing links from U1 to each of its neighbours with the receiving links coming from the respective side. On the bottom side of the adapter, two grids of electrical contacts, as shown in figure 9.10a, are used to establish a connection with the jFEX board. The maximum available size of the connector of 10 × 20 pins limits the possible coverage of links that can be tested by this setup to 19 links. The position of the pad over the FPGA footprint was chosen to cover the highest number of complete connections. A picture of the adapter is shown in figure 9.10a and the complete assembly with both adapters on the jFEX is shown in figure 9.10b. As the grid of electrical contacts has a spacing of only 0.1 mm, correct alignment of the boards to each other is a critical task. To achieve this, the adapter is provided with holes for guidance pins which are aligned with existing holes on the jFEX board intended for the assembly of components that are not used on the prototype. Glued-on plastic screws and an aluminium block are used to apply the required force upon the pads to make reliable contact. As can be seen in figure 9.10b, the adapter covering the position of U2 has 4 guidance pins on both sides of the pads whereas the adapter utilising the empty spot of U3 only has 2 guidance pins on one side as the other side is blocked by assembled components. The lack of guidance pins for the U3 adapter is assumed to be the reason that no links could be established on that side while all expected links were available on the side of U2. Therefore, the following tests are only based on the links between U1 and U2.

As no additional modules were available during local tests in Mainz, the optical input to the FPGA was limited to a single MiniPOD being fed by its own output. Therefore, the test was divided into three parts focusing on the different positions of the optical connection. Separate firmware versions were required for each setup as the preliminary power supplies could not support the activated high-speed transceivers for all positions in parallel.

For each setup, the MGTs which received an input from the optical connection were operated in PMA loopback mode. Additionally, the neighbouring MGT quad on each side was also operated as regular transceivers to create an environment as similar as possible to regular operation of a full-scale module both within the FPGA logic and for the signal routing on the PCB. The available connections were grouped into the 12 connections using the path via the MiniPODs and the purely electrical group consisting of the PMA loopback signals and neighbouring noise channels. The optical links were monitored to identify from which part of the two-staged connection potential transmission errors originated.

The scope of the test for each PMA loopback link was a BER of 10^{-15} at a data rate of 12.8 Gb/s followed by an eye diagram scan. Eye diagrams are commonly used to graphically analyse the signal quality for digital data transmissions. The principle behind an eye diagram scan is

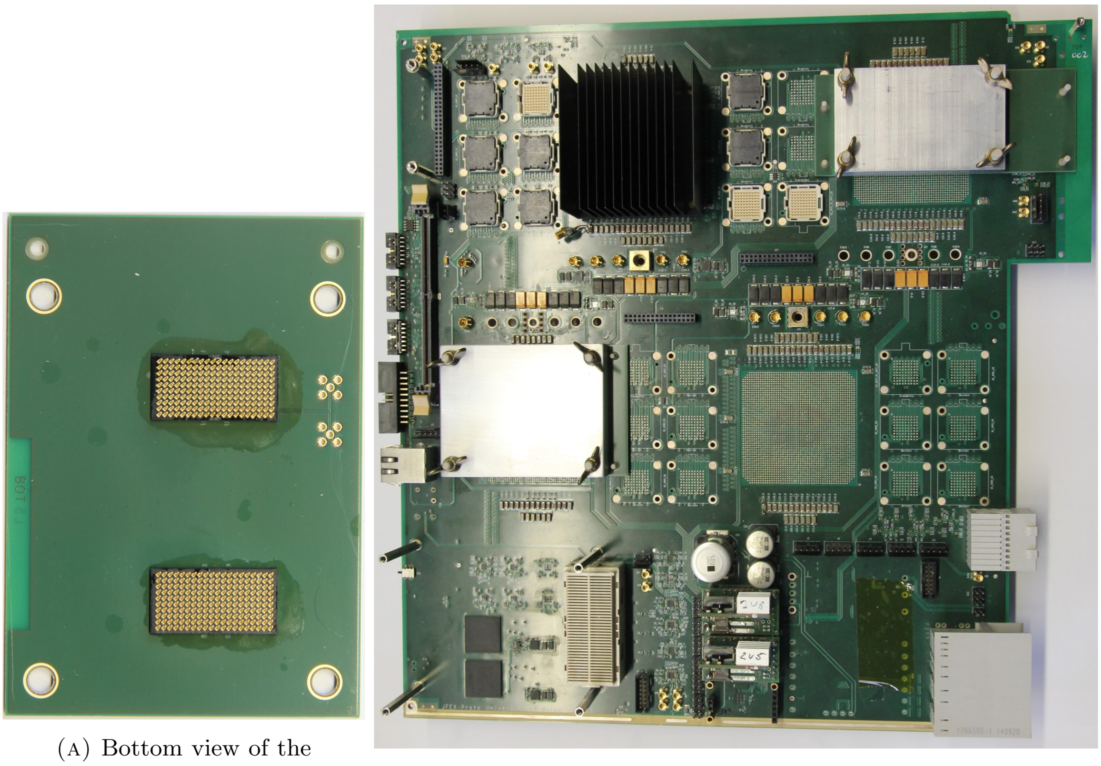

(A) Bottom view of the PMA loopback adapter.

(B) Prototype board with both PMA loopback adapters.

FIGURE 9.10: Pictures of the PMA loopback adapter (A) and the assembled jFEX prototype (B) including the two adapter boards under aluminium block on the left and in the top right corner of the main PCB.

illustrated in figure 9.11. Typically, this is done by stacking multiple recordings of an equalised periodical signal on an oscilloscope screen leading to a result similar to figure 9.11a. As the electrical signals relevant for this test are routed on internal layers of the jFEX PCB, direct access to the signals with an oscilloscope probe is not possible without a severe risk of damaging the board. Therefore, a digital version of the eye diagram scan is used. In this method, the sampling point is intentionally shifted from the ideal position in the middle of the eye labelled as 'Data Sample' to the 'Offset Sample' by adding an offset to the sampling time and the voltage threshold used for comparison of the detected signal to create the horizontal and vertical offset, respectively. Multiple of these sampling points are used to measure the BER which is then plotted using a colour coding to represent the measured BER as shown in figure 9.11b. This feature is offered directly by Vivado containing the utilised firmware package including automatic variation of the sampling time and the voltage threshold against which the input signal is compared. Via this variation, results similar to the oscilloscope-based measurement can be reproduced. The x-axis and y-axis of the digital eye diagram refer to units of time and voltage, respectively. Vivado normalises both of these to a unitless value based on the possible variation[7] of the offset. The size and shape of the open area of this scan, which is the parameter space in which no transmission errors have been recorded, can be used as a measure for the stability of the connection. Due to the normalisation of both axes, the resulting area for scans recorded in Vivado is also unitless. Further information on this method is provided by Xilinx, e.g. in

---

[7] Vivado automatically limits the possible variation range to half the period length and voltage range of the chosen signalling standard.

their technical specifications and the 'Signal Integrity Simulation and On-Chip Evaluation for Low-Cost FPGA Transceivers' [87].

During the tests in all configurations, the 19 connections using PMA loopback transferred at least 10^{15} bits without recorded errors leading to a measured BER of $< 10^{-15}$ per link thus achieving slightly more than the defined goal. The complete results from the described measurements are shown in appendix F.1.

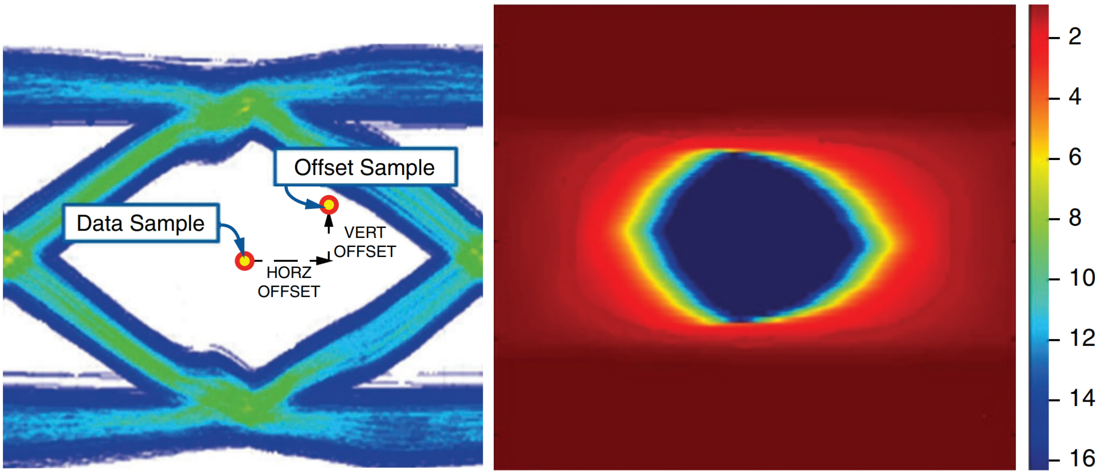

(A)

(B)

FIGURE 9.11: Example pictures of an oscilloscope-based eye diagram (A) showing the variation of the sample point using offsets in horizontal and vertical directions to scan an area. This principle leads to a digital eye diagram as shown in (B). Figures taken from [87].

After successful verification of the reliability regarding data transmission in the used test setup, coarse eye diagram scans were performed on all PMA loopback links with a targeted BER of 10^{-10} per measured point to quantify the margin with respect to the expected performance of the connections in an environment with increased noise. The results of the scans are shown in table 9.4. As the total number provided by Vivado for these tests is difficult to interpret, a second set of scans was performed with these links operating in direct transmission to identify the impact of the PMA loopback on the signal quality.

The eye scans show an alternating pattern between one link with a relatively large open area and one link with a significantly smaller open area. Comparison of this pattern with the footprint of the FPGA shows that MGT positions on the outer edge of the footprint led to smaller eyes. If the root cause for this was to be found in the PCB design, the opposite behaviour would have been expected as the corresponding connections had a shorter path routed through the densely populated region under the devices. This leads to the assumption that the measured difference in signal quality is not caused by the PCB design but rather by the adapter used for this test which might have provided an impaired contact on the outer edge of the connector pads. This assumption was confirmed by tests with the next prototype generation where no such pattern was seen in PMA loopback-based connections between the FPGAs [88].

Except for the connection ending in quad 129 channel 0, all links showed a small deterioration in signal quality when operated in PMA loopback mode. This effect was significantly smaller than the effect described above presumably caused by the contact pads on the adapter. For a deeper understanding of both effects, additional eye diagram scans with a higher resolution were done for the links ending in quad 127 channel 0 and quad 130 channel 1 as the links with

the smallest and largest open area during operation in PMA loopback mode, respectively. The graphical representations of these scans are shown in figure 9.12 for both channels each in direct transmission and PMA loopback mode. The scans for both links show a small but opposite asymmetry in terms of positive and negative offsets in the sampling time represented by a shift along the x-axis. The eye diagram recorded in quad 129 channel 0 in regular resolution did not show any noticeable distortion. Detailed scans were therefore not performed for this channel.
The root cause for this effect could not be identified and was not further studied, as the transmission error for all links already satisfied the defined goal. Additionally, if similar asymmetries lead to transmission errors in the final design, a shift of the eye diagram can be compensated by fine adjustments in the MGT firmware.
Even the combination of both deterioration effects from operation in PMA loopback mode and the stronger deterioration from the connector pads still leaves the transmission signals in a condition sufficient to meet the BER test requirements of $< 10^{-14}$. Therefore, it is concluded that the scheme for input signal duplication using PMA loopback is suitable for the jFEX design from a hardware perspective. The impact on the latency for the real-time algorithms was already considered in the design phase. The comparison of the open area of quad 127 channel 0 and quad 130 channel 1 also suggests that without the use of a connector pad each link has a sufficient margin to allow for stable operation in environments with increased temperature and crosstalk from fully populated boards.

MGT number	PMA loopback	direct transmission	MGT number	PMA loopback	direct transmission
126_3	3648	4416	129_1	4800	5888
127_0	2240	2880	129_2	2560	3072
127_1	4864	5376	129_3	4800	6720
127_2	3136	3712	130_0	2624	2880
127_3	4672	5568	130_1	5248	6336
128_0	2944	3776	130_2	2944	3200
128_1	4736	5312	130_3	5184	6080
128_2	3136	3712	131_0	2944	3264
128_3	4992	5696	131_1	4864	6144
129_0	3008	2688			

TABLE 9.4: Open area of eye scans with and without activation of PMA loopback mode. Higher numbers represent a less interference-prone connection. The eyes were sampled in 8 steps for both sampling time and voltage threshold with a measured BER for the open area of 10^{-10}. The connections are denoted by the receiving MGT as Quad_Channel.

9.3.3 Link-Speed Tests at CERN

Two sets of integrated tests were done at CERN around the second quarter of 2017 to test the compatibility of the jFEX prototype with a LATOME prototype and the FEX Test Module (FTM)[8].
A picture of the jFEX prototype including assembly of power converters, control mezzanine,

---

[8] The FEX Test Modul is a board developed in Birmingham specifically for high-speed signal tests with ATLAS L1Calo Phase-I hardware.

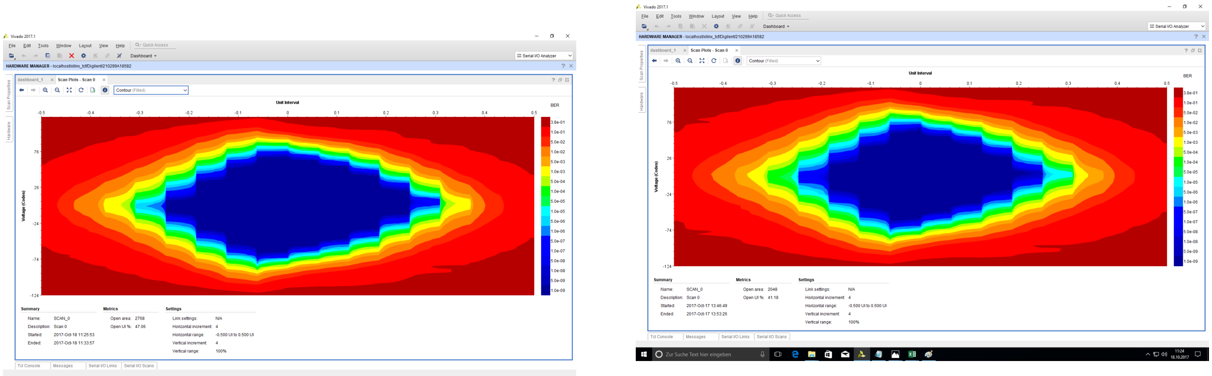

Summary
Name: SCAN_0
Description: Scan 0
Started: 2017-Oct-18 11:25:53
Ended: 2017-Oct-18 11:33:57

Metrics
Open area: 2768
Open UI %: 47.06

Settings
Link settings: N/A
Horizontal increment: 4
Horizontal range: -0.500 UI to 0.500 UI
Vertical increment: 4
Vertical range: 100%

Summary
Name: SCAN_0
Description: Scan 0
Started: 2017-Oct-17 13:46:49
Ended: 2017-Oct-17 13:53:26

Metrics
Open area: 2048
Open UI %: 41.18

Settings
Link settings: N/A
Horizontal increment: 4
Horizontal range: -0.500 UI to 0.500 UI
Vertical increment: 4
Vertical range: 100%

(A) MGT 127 Quad 0 in direct transmission

(B) MGT 127 Quad 0 in PMA loopback mode

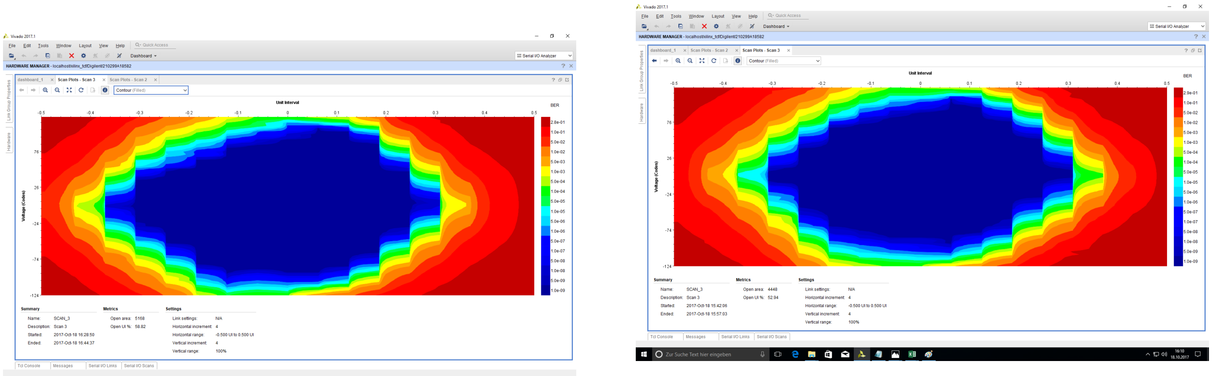

Summary
Name: SCAN_3
Description: Scan 3
Started: 2017-Oct-18 16:28:50
Ended: 2017-Oct-18 16:44:37

Metrics
Open area: 5168
Open UI %: 58.82

Settings
Link settings: N/A
Horizontal increment: 4
Horizontal range: -0.500 UI to 0.500 UI
Vertical increment: 4
Vertical range: 100%

Summary
Name: SCAN_3
Description: Scan 3
Started: 2017-Oct-18 15:42:06
Ended: 2017-Oct-18 15:57:03

Metrics
Open area: 4448
Open UI %: 52.94

Settings
Link settings: N/A
Horizontal increment: 4
Horizontal range: -0.500 UI to 0.500 UI
Vertical increment: 4
Vertical range: 100%

(C) MGT 130 Quad 1 in direct transmission

(D) MGT 130 Quad 1 in PMA loopback mode

FIGURE 9.12: Eye diagrams of the links from the PMA loopback test with the smallest (top) and largest (bottom) open area. For comparision, both links show the scan from the direct transmission where the MGT generates the output itself (left) versus the operation in PMA loopback mode forwarding an incoming signal (right).

MiniPODs, and two fibre bundles of respectively 12 and 72 links is shown in figure 9.13. For direct connections to the LATOME, the 72-way connector was replaced with a similar cable providing 48 inputs. The control mezzanine utilises the PicoZed board which can be seen in red between the main PCB and the mezzanine. While the backplane connectors were mounted, the main power converters were still missing. Therefore, the jFEX board was powered by an external lab power supply providing 12 V and 3.3 V on separate inputs with a common ground connection.

9.3.3.1 LATOME

The setup used for the LATOME tests is shown in figure 9.14. Multiple variations of this setup were used in terms of fibre connections between the modules and clock sources. For any tests featuring the LATOME prototype, the link speed was limited to 11.2 Gb/s due to the FPGA choice and the firmware available for LATOME. The main tests were performed to verify the bit error ratio in the data transmission between the modules as required by the L1Calo group. For this, a first test was done with a direct connection of a 48-way trunk cable and local clock

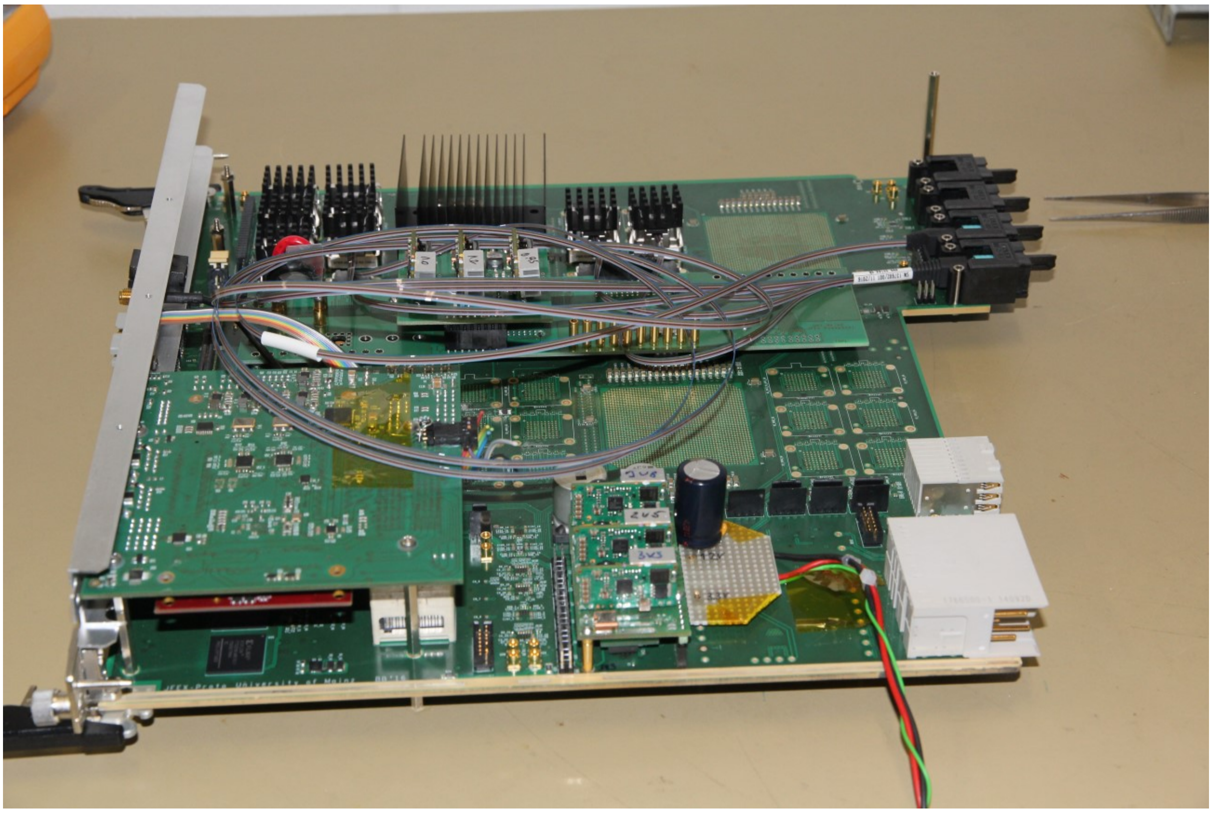

FIGURE 9.13: The jFEX prototype featuring the assembly as used for the integrated tests at CERN. The red board under the mezzanine on the left is the PicoZed, which was used as a quick solution for a control board. The closest of the backplane connectors hosts the 72-ways fibre connection.

sources on each modules aiming at a high number of transmitted bits per channel followed by a second test to verify the fibre connection via the FOX prototype including the 72-way connectors used in the jFEX design and the use of a common clock source to have a setup which is closer to the conditions during the operation as a part of the ATLAS Trigger System. The complete results of the described measurements can be found in appendix F.2. Additionally, short tests were performed featuring dedicated adaptors for controlled attenuation of the signal on a single link. This modification is shown in figure 9.15.

During the first test, a total of about 6×10^{14} bits were transmitted on each of the 48 links without any recorded transmission errors leading to a measured BER of about 1.7×10^{-15} per link which already exceeded the official goal set for this test. For the second test, less time was available due to a tight schedule by the LATOME group. Therefore, a shorter test was chosen providing a BER of about 10^{-13} per link just to verify the basic functionality in this more realistic setup.

Additionally, for a single, randomly selected link the stability of the connection was tested while using optical attenuators to simulate the effects of longer cables and more connectors as used in the complete L1Calo setup. While an attenuation of 5 dB was tolerated, 10 dB led to measurable bit errors at a BER of about 3.4×10^{-11}. As even 5 dB is more than the expected loss in the final design, the detected error rate was not assumed to appear in L1Calo operation.

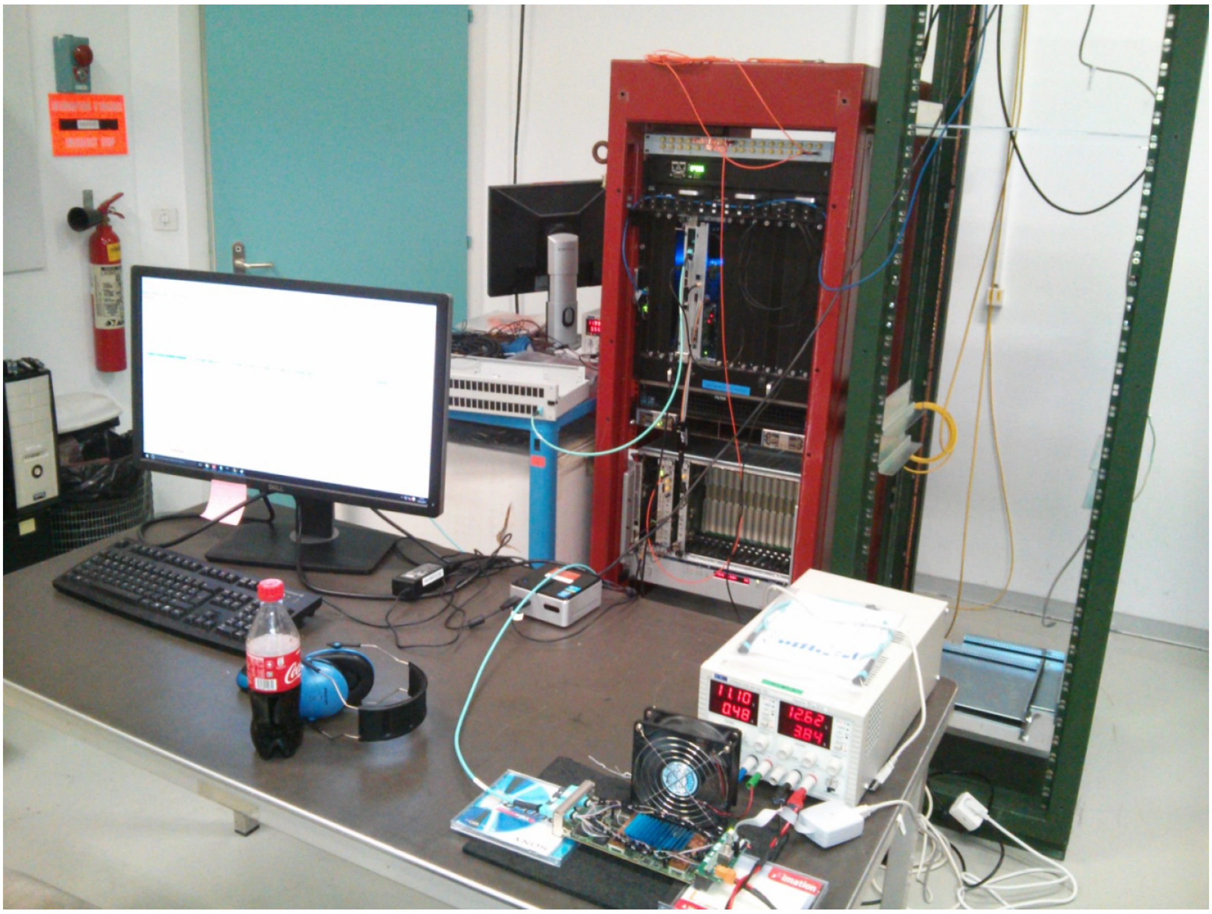

FIGURE 9.14: Setup used for the integrated tests with the LATOME prototype. The jFEX board was vertically inserted in the upper level of the red crate while provided with a 48-way optical input through the open front plane. The LATOME prototype was operated on the table. Multiple setups were used for the fibre connection. Shown in this picture is a single stage of the FOX prototype to the left of the jFEX.

9.3.3.2 FTM

The tests featuring the FEX test module were intended to extend the functional tests for the jFEX beyond the range supported by LATOME. This mainly included an increase of incoming optical signals to at least 60 to completely fill the available side of the assembled FPGA and the provision of a clock signal via the backplane as expected during operation as part of L1Calo. The setup used for these tests is shown in figure 9.16.

Two long term tests were carried out using the FTM as data source with a link speed of 11.2 Gb/s and 12.8 Gb/s, respectively. For both tests, no errors were detected leading to a measured BER of about 10^{-15} per link and test. The complete results for both tests can be found in appendix F.3.

Additional tests with the FTM were carried out to test different features in the firmware like 8b/10b encoded data based on firmware developed by Christian Kahra instead of the IBERT cores for MGT operation. As the goal of these tests was the verification of the firmware instead of the jFEX design, they are not covered within this thesis.

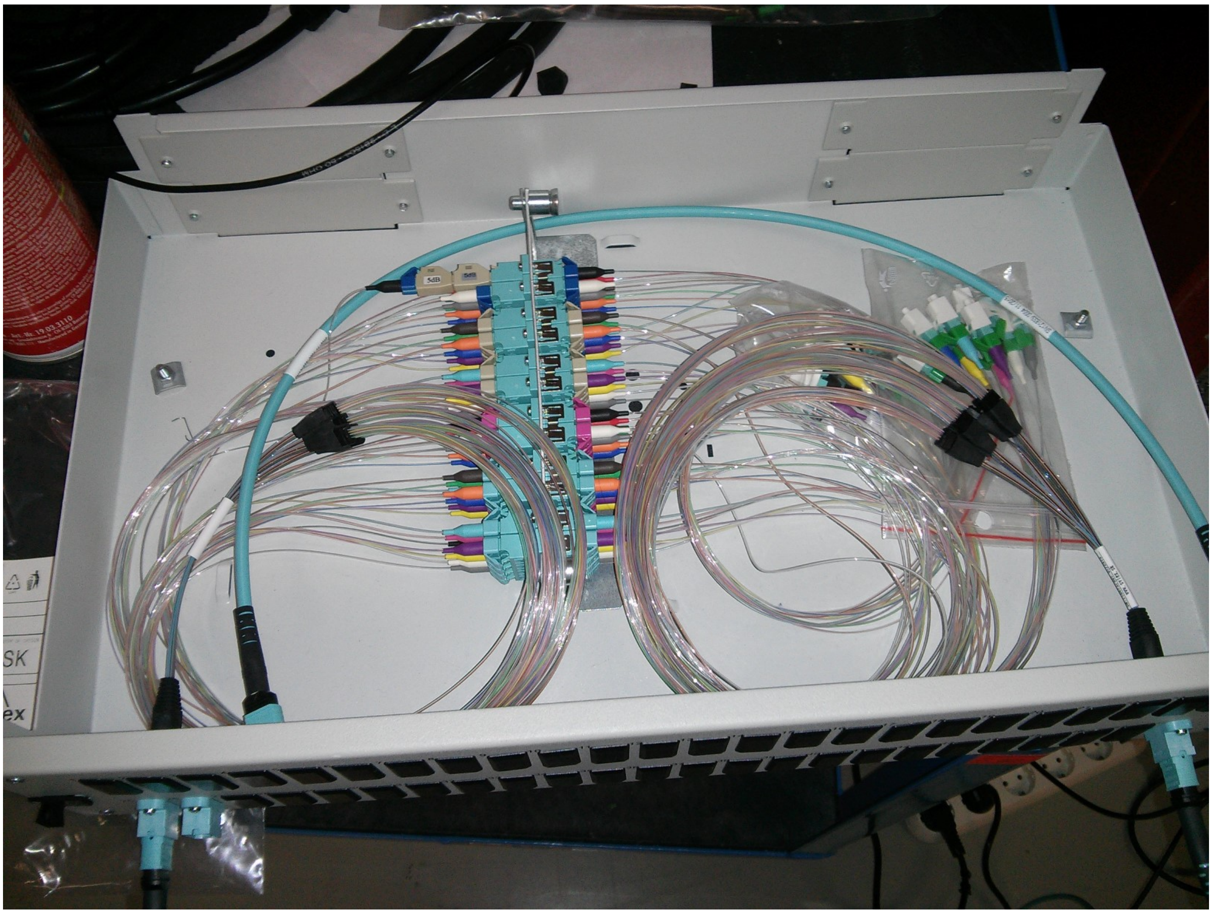

Figure 9.15: FOX prototype featuring two additional attenuation modules on a single link to simulate further stages of potentially lossy connections.

9.4 Summary

During the initial tests and algorithm simulations, the logical design was verified for its suitability to run the intended algorithms as well as their ability to constitute their part of the trigger system for the ATLAS experiment.

Stable data transmission from the LATOME prototype and the FTM showed no transmission errors while already exceeding the defined goals of a measured bit error ratio below 1×10^{-15} per link. Capabilities regarding real-time data output were also indirectly verified by tests using the jFEX as source and sink as the jFEX prototype also serves as a demonstrator for the Phase-I L1Topo design which is intended to be designed using the same technology. While showing minor deterioration in signal quality, the use of the PMA loopback scheme for input data duplication was also successfully verified even though these tests were carried out under conditions that are considered to be more challenging than during regular operation. Therefore, the jFEX hardware design regarding requirements of the ATLAS physics programme on trigger performance and capabilities of the L1Calo system is considered successful.

Regardless of the successful verification of the jFEX board design, a list of small changes is recommended in 9.5. Many of these changes are already included in later iterations of the jFEX

Figure 9.16: Setup used for the integrated tests with the FTM. Both boards were inserted in the upper level of the red rack with up to two FOX prototypes on top of it as the entire fibre handling was done via the backplane.

starting with the second prototype. These fully equipped boards have been used for more recent and more detailed tests, some of which have been published in [50].

9.5 Suggested Changes for Later Iterations

During the design phase, there were already some possible modifications of the design known which could not be implemented in the prototype design for various reasons. In most cases this includes new components offering certain advantages over the old versions but for which a sufficient availability or documentation was not given when layout phase and prototype production started. Therefore, it was at that time not possible to base a design in a project within a tight, given timeline, such as the jFEX, on these chips. Provided their availability during later iterations, some of these components may be taken into consideration. In other cases, more suitable solutions were found while preparing test setups. Also, the increased flexibility of the prototype allowing for it to use it in a broad variety of test setups including stand-alone operation will not be required for production boards and can be removed. This applies mainly to the multiple potential clock sources required for the prototype. Some of these changes are already included in the next iterations of the jFEX.

9.5.1 FPGA Choice

Replacing the large FPGAs with a version from the more recent series called UltraScale+™(USP) means a major change in the jFEX design, yet implies minimal to no difference for the board layout. Amongst the FPGAs that are available in the A2577 package is one device, the XCVU9P, that is exactly pin compatible and can be placed on the board without requiring significant changes to the PCB design. Only some layout recommendations differ between UltraScale and UltraScale+ series. As an example, a larger amount of decoupling capacitors for the power planes are recommended for the newer version. However, this is already covered by the layout for the prototype since a more conservative design was used, based on the VCU110 evaluation board rather than the user guide.

Advantages of the USP series are the reduced power consumption and increased logic resources. Both are mainly achieved by changing the structure size from 20 nm to 16 nm, allowing a higher density of gates and reducing the V_{CCINT} voltage level to 0.85 V for medium speed grades. A comparison of the logic resources relevant for the jFEX system between the two FPGA types is shown in table 9.5. Overall, the UltraScale+ contains about 10 % more logic resources, except for DSP slices where the difference is 280 % for the considered FPGAs. The USP series uses a new distinction between different types of memory. For operation during Phase-I, this resource is not assumed to be critical for the jFEX design since it cannot be used in the real-time functionality of the jFEX and will therefore only be used for control and read-out functions that are not expected to consume a critical amount of resources in the FPGAs. Exact requirements in terms of available memory for the operation as a part of the Phase-II trigger system were not available during the decision process. Therefore, a detailed comparison of the memory capabilities can not be decisive. Due to the identical package, both types support the exact same number of MGT and parallel I/O connections.

For the estimation of the power consumption, the same number of resources in use is assumed for the XCVU9P as it was assumed in 9.2.2. Due to the larger logic resources available, this leads to a drop of the utilisation below the 75 % mark assumed for the XCVU190. Implementing the exact same algorithms on two different types of FPGAs might lead to designs with differences in logic utilisation, due to different placements and routing by Vivado based on the architectures. For the comparison of the power efficiency between the two series, this is neglected. The comparison of the estimated currents drawn per supply voltage is shown in table 9.6. The total power consumption decreases from 92.5 W on the UltraScale device to 77.7 W on the UltraScale+ chip mainly due to the reduced current drawn via $\mathrm{MGTAV}_{\mathrm{CC}}$ and the reduced supply voltage on V_{CCINT}, which has by far the highest contribution to the overall power

resource	XCVU190	XCVU9P
System logic cells	2,349,900	2,586,150
CLB Flip-Flops	2,148,480	2,364,480
CLB LUTs	1,074,240	1,182,240
Max. Distributed RAM (Mb)	14.5	36.1
Block RAM (Mb)	132.9	75.9
DSP slices	1800	6840

Table 9.5: Comparison of selected logic resources between XCVU190 and XCVU9P. Data taken from the UltraScale Architecture Overview [59].

supply voltage	XCVU190	XCVU9P
V_{CCINT}, $V_{\mathrm{CCINT\_IO}}$, V_{CCBRAM}	52 A (0.95 V)	57 A (0.85 V)
$\mathrm{MGTAV}_{\mathrm{CC}}$	16 A (1.0 V)	7 A (0.9 V)
$\mathrm{MGTAV}_{\mathrm{TT}}$ (1.2 V)	18 A	17 A
$\mathrm{MGTAV}_{\mathrm{CCAUX}}$ (1.8 V)	1 A	1 A
V_{CCAUX}, $V_{\mathrm{CCAUX\_IO}}$ (1.8 V)	2 A	2 A
V_{CCO} (1.8 V)	1 A	1 A

TABLE 9.6: Comparison between XCVU190 and XCVU9P of the estimated currents drawn from different supply voltages, based on Xilinx Power Estimator 2016.3 and 2017.2, respectively. The same number of resources in use is assumed for both chips, even though the utilisation for the USP device thereby drops below the targeted 75 % from 8.8.1. Total on-chip power is reduced from 92.5 W on the US device to 77.7 W on the USP chip. The detailed XPE sheets are shown in appendix D and appendix E for the XCVU190 and XCVU9P, respectively.

consumption for both families.

Since Virtex UltraScale+ FPGAs use only GTY transceivers instead of the combination with GTH, the utilisation of different link types with respect to power efficiency for UltraScale devices has no impact on the XCVU9P. While this decreases the reduction in power consumption, it simplifies the firmware which will no longer be required to handle two types of MGT in a single device. Due to the higher resource count in the UltraScale+ FPGA, the difference in the final system might be smaller than estimated in case the larger logic resources are used to implement more - or more complex - algorithms.

If required, an evaluation board hosting the XCVU9P, the VCU118, is available to test whether the UltraScale+ series supports all required features from the UltraScale FPGA.

9.5.2 Clocking

The main focus of jFEX iterations after the first prototype will move from flexibility for different test setups for the design verification to the actual functionality as a part in the L1Calo system. Therefore, parts in the design can be reduced that were required only for stand-alone operation or flexibility of input sources. Such a simplification of the clocking scheme is shown below.

The clocking scheme was simplified by outsourcing clock signal processing completely to the mezzanine and removing quartz oscillators and clock generators from the main board. This also reduces the amount of required fan-out and multiplexing on the jFEX. A single jitter cleaner on the mezzanine is sufficient to clean the global clock and derive multiples of it as MGT reference clocks.

The scheme for most of the MGT reference clocks remains almost untouched since only additional inputs to the fan-out chips from local clock sources need to be removed. Larger changes are required for the input reference clock where the complete multiplexer stage from figure 8.7 is removed, including the local sources. The reference clock for TTC data reception remains unchanged.

As fine delay in parallel I/O might require an independent clock[9], a new scheme for global clock

---

[9] A clock source independent from the one used to drive the parallel I/O links is recommended by Xilinx for Virtex-7 FPGAs. Tests using the second prototype will be required to show if this can be done without an additional clock in the UltraScale or UltraScale+ series as information regarding this topic by the manufacturer were not available during the design phase.

distribution is proposed as shown in figure 9.17. As input, a clock from any of the hub slots can be chosen via a multiplexer and is routed to a jitter cleaner on the extension mezzanine which is not represented in the diagram. From the mezzanine, two independent differential signals are sent to LVDS fan-out chips which distribute the clock to the FPGAs and in one case to the CPLD. On spare output pins of fan-out chips in every clock tree, MMCX connectors are added to provide an easy way to measure the signals.
The focus on mezzanine-based clocking was chosen due to the non-volatile memory of the programmable jitter cleaner[10], which only allows to be programmed twice during its entire lifetime. Therefore, replacing this device might be required even in the final production in case of changes in the firmware, using different frequencies on the MGT reference clocks or to cope with a modified environment after converting to the Phase-II setup.

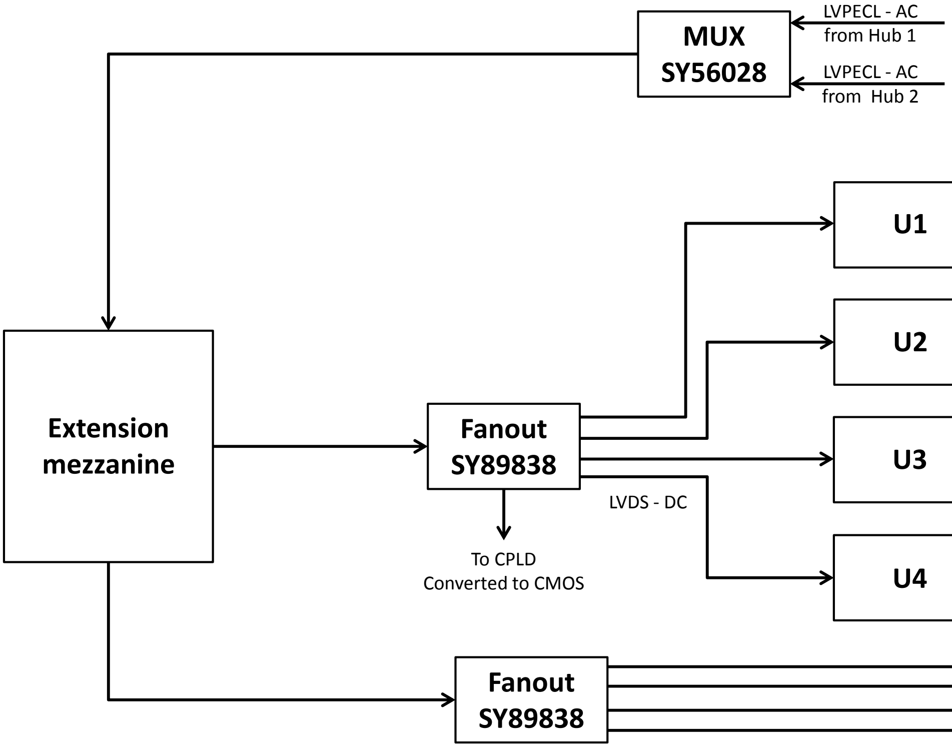

FIGURE 9.17: Proposed global clock distribution scheme for the final production series with the added option of having completely independent global clocks coming from the mezzanine. The complete scheme relies more on the mezzanine which also hosts the jitter cleaner and potentially multiple fan-out chips.

---

[10] It is assumed that the Si5345 which was already used on versions of the extension mezzanine operated on the first prototype will be used.

9.5.3 Miscellanea

For fully populated versions of the jFEX, a new design for the extension mezzanine will be required with respect to the PicoZed. While relying on a simplified layout by using an already completely designed environment for the Zynq was convenient to get an operational board in the shortest time, this setup with only four high-speed links available does not provide enough connectivity for four FPGAs and an external IPbus. Since control firmware development so far only uses the programmable logic part of the Zynq, which allows to base this on the firmware already available for the Run 2 L1Topo, a pure FPGA can be used instead. Fitting devices can be found in the Xilinx Artix®-7 series with the XC7A200T in the FFG1156 package appearing to be most suitable. Alternatively, to preserve the additinoal flexibility provided by the use of a device from the Zynq family, larger devices of this type could be used to provide more high-speed links.

The positioning of the real-time output on the first prototype was designed to simplify the layout. The firmware side might recommend a repositioning within the central Super Logic Region (SLR) in order to reduce SLR crossings. Without further experience in firmware development on UltraScale FPGAs, it is not yet possible to judge whether or not this step will be necessary. Therefore, this change was not considered for the second iteration but could be reconsidered for the final production even though usually no large changes in the design would be introduced at this late stage of the project.

The default setting for configuration can be changed to Master SPI mode for the second prototype board or at least for the final production. Also, the JTAG scheme can be simplified by removing the additional connections and resistors used to bypass potentially missing or non-functional FPGAs as only fully assembled and operational boards will be relevant in later iterations.

9.5.4 Production Design

The final design of the jFEX has since been produced, tested, started commissioning in spring 2022, and is currently serving in its intended role as a part of the ATLAS Trigger System.

Some of the changes to the design used for the first jFEX prototype that are discussed in this chapter have been included in later iterations of the board. The main features are the upgrade to UltraScale+ FPGA and the adaptation of a clocking scheme similar to the one presented in section 9.5.2. Other suggestions like the rerouting of the real-time output signal to utilise the central SLRs were not implemented as no sufficient need from the firmware perspective was seen to justify the included risk in changing the board layout at a late development stage.

The only larger modification to the board design was the addition of termination resistors used in the calibration of MGTs during the configuration of an FPGA according to the termination resistor calibration circuit [67]. In earlier designs, only 2 out of 6 of these pins per FPGA were connected. Therefore, the specified termination input was only provided to MGTs in the central SLR. As new boards were produced after the changes were included, no direct tests could be done to compare the performance of the high-speed links before and after - as the tests would have most likely been dominated by individual characteristics per board. Also, the eye diagram scans shown in table 9.4 show no pattern to indicate measurable deterioration of the links lacking the termination resistors, partially leading to the late discovery of the missing

resistors.
Due to the system design using 6 jFEX modules, an installation in double slot width was chosen for the operation in the ATLAS Trigger System allowing for the use of well-tested cooling and cable management solutions from the prototype and pre-production versions.

Part IV

Summary

Chapter 10

Summary and Outlook

Due to the increased luminosity delivered by the upgraded LHC, a redesign of the trigger system for the ATLAS experiment was required to limit the Level 1 acceptance rate to a manageable level while also maintaining a good sensitivity for certain physics processes. After completion of the ATLAS Phase-I upgrade, the first trigger stage, which is implemented in specialised hardware designs, will be based on three different types of Feature Extractors used to identify energy deposits likely to be linked to certain physics processes in the detector. The jet Feature Extractor (jFEX) is mainly used to calculate hadronic jets and global quantities such as missing transverse energy.

This thesis describes the design, development, and test of the jFEX hardware and its capabilities to run the intended algorithms as a part of the ATLAS trigger system.

The design for operation in Run 3, which started in the year 2022, consists of six modules each carrying four processor FPGAs from the latest FPGA series. Each module provides a maximum input bandwidth of ≈ 3 Tb/s received via 240 links operated at up to 12.8 Gb/s. For the operation during Run 4, currently expected to start in 2029, the jFEX system can be extended by adding further modules to cope with the increased input bandwidth caused by the utilisation of a finer granularity from the forward detector regions in trigger algorithms.

The processing capabilities of the jFEX modules are designed to cope with the required input rate of 40 MHz and tight latency budget of about 100 ns. Several control functionalities are included in the design to integrate the jFEX into the ATLAS trigger system which requires the communication via multiple control interfaces and reception of e.g. TTC signals to align the separate modules and trigger subsystems.

Within the scope of this thesis, conceptual work for the hardware design of the first prototype has been carried out followed by tests mainly to verify the hardware design and board layout regarding its suitability for routing of high-speed signals on densely populated PCBs. All performed tests were successful with most even exceeding the defined goals in terms of an achieved bit error ratio of $< 10^{-15}$.

The intended verification of possible algorithms for the identification of jets could only partially be achieved within this thesis but have since been confirmed by other, more detailed studies. The main cause for the limited extent of studies within the thesis was the loss of analysis code and data samples.

In the later course of the project for the ATLAS community, further iteration of the jFEX board

beyond the scope of this thesis have been developed and tested, leading to the final production design.
By mid 2024, the new Phase-I Trigger System replaced the legacy system in the ATLAS experiment's data aquisition and it will continue to provide the trigger decisions for the remainder of Run 3 and, with further adaptations during the Phase-II upgrade, well beyond 2030.

Appendix

Appendix A

Acronyms

ASIC	Application-Specific Integrated Circuit
ATCA	Advanced Telecommunications Computing Architecture
ATLAS	Toroidal LHC ApparatuS
BER	Bit-Error Ratio
BRAM	Block Random-Access Memory
CLB	Configurable Logic Block
CML	Current Mode Logic
CMOS	Complementary Metal-Oxide-Semiconductor
CMX	Common Merger module eXtended
CP	Cluster Processor
CPLD	Complex Programmable Logic Device
CPM	Cluster Processor Module
CPM	Cluster Processor Module
CTP	Central Trigger Processor
DPS	LAr Digital Processing System
DSP	Digital Signal Processing
eFEX	electromagnetic Feature EXtractor
EMEC	electromagnetic End-Cap
FCAL	Forward Calorimeter
fFEX	forward Feature EXtractor
FTM	FEX Test Module
FOX	Fibre Optical Exchange
FPGA	Field-Programmable Gate Array
GCK	global Clock
gFEX	global Feature EXtractor
HEC	hadronic End-Cap
HLT	High Level Trigger

IBERT	Integrated Bit Error Ratio Tester
IPMI	Intelligent Platform Management Interface
IPMC	IPMI Management Controller
ITk	Inner Tracker
JEM	Jet/Energy Module
JEP	Jet/Energy Processor
jFEX	jet Feature EXtractor
L1A	Level-1 Accept
L1TOPO	Level-1 Topological Processor
LATOME	LAr Trigger prOcessing MEzzanine
LHC	Large Hadron Collider
LUT	Lookup Table
LVDS	Low-Voltage-Differential-Signaling
LVPECL	Low Voltage Positive Emitter-Coupled Logic
MDT	Monitored Drift Tube
MGT	Multi-gigabit transceiver
MTP	Multi-fiber Termination Push-on
MUCTPI	MUon to Central Trigger Processor Interface
nMCM	New Multi-Chip Module
NSW	New Small Wheel
PCB	Printed Circuit Board
PPM	PreProcessor Module
PPr	PreProcessor
QCD	Quantum Chromodynamics
refclk	reference clock
ROD	Readout Driver
RoI	Region of Interest
RPC	Resistive Plate Chambers
RTM	Rear Transition Module
SLR	Super Logic Region
SPI	Serial Peripheral Interface
TDAQ	Trigger and DAQ
TDR	Technical Design Report
TGC	Thin Gap Chamber
TOB	Trigger Object
TREX	Tile Rear Extension
TRT	Transition Radiation Tracker
TTC	Trigger, Timing & Control
US	UltraScale
USP	UltraScale+
XPE	Xilinx Power Estimator

Appendix B

jFEX Block Diagrams

These block diagrams were created in their final version in October 2015 and used as the conceptual basis for the development of the jFEX PCB. The power sequencing shown on page 3 was later reworked during the development of the first prototype.
For improved readability, please refer to the digital version.

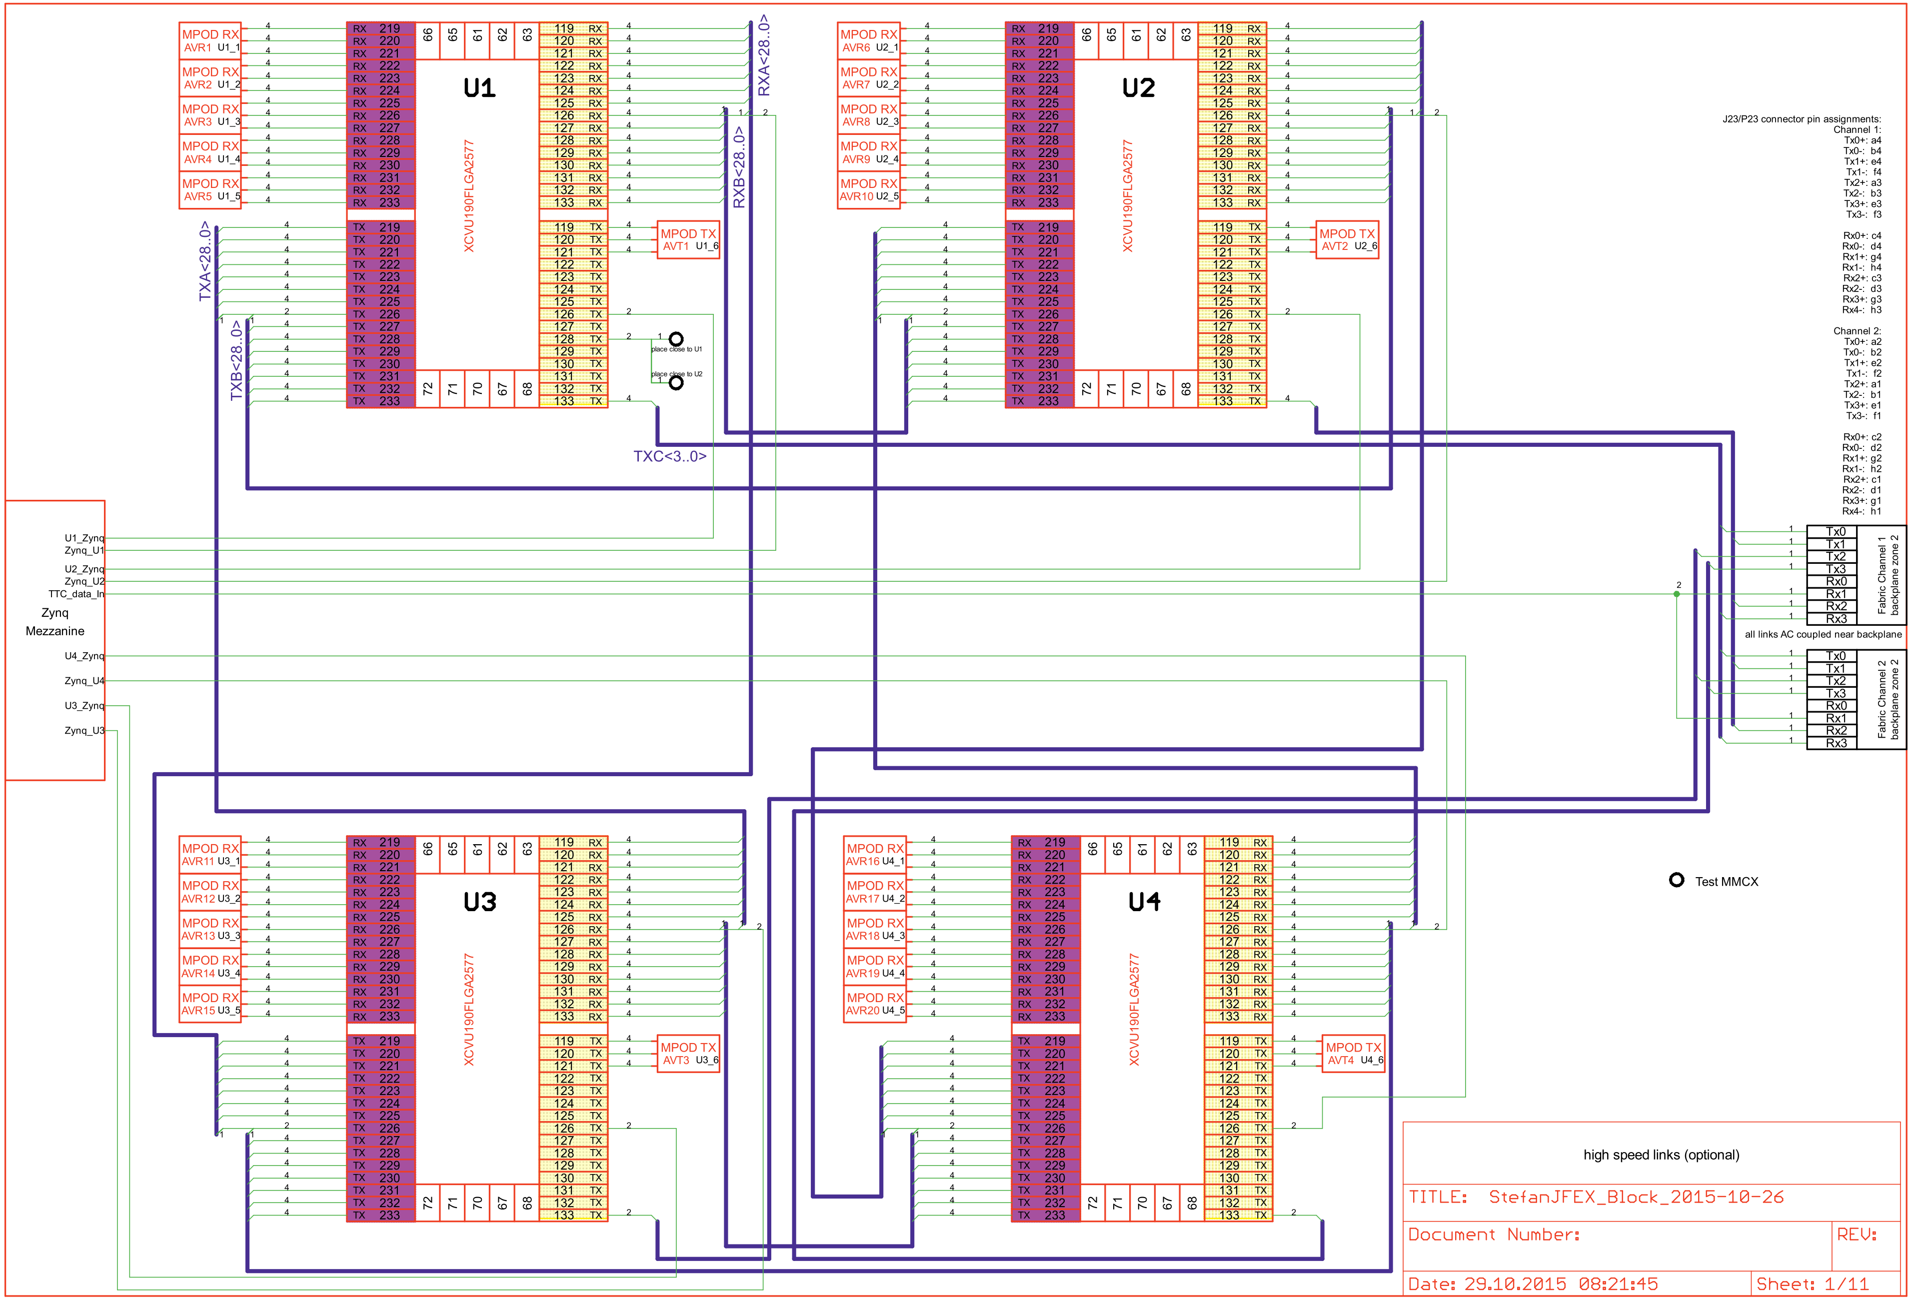

J23/P23 connector pin assignments:
Channel 1:
Tx0+: a4
Tx0-: b4
Tx1+: e4
Tx1-: f4
Tx2+: a3
Tx2-: b3
Tx3+: e3
Tx3-: f3

Rx0+: c4
Rx0-: d4
Rx1+: g4
Rx1-: h4
Rx2+: c3
Rx2-: d3
Rx3+: g3
Rx4-: h3

Channel 2:
Tx0+: a2
Tx0-: b2
Tx1+: e2
Tx1-: f2
Tx2+: a1
Tx2-: b1
Tx3+: e1
Tx3-: f1

Rx0+: c2
Rx0-: d2
Rx1+: g2
Rx1-: h2
Rx2+: c1
Rx2-: d1
Rx3+: g1
Rx4-: h1

high speed links (optional)

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 1/11

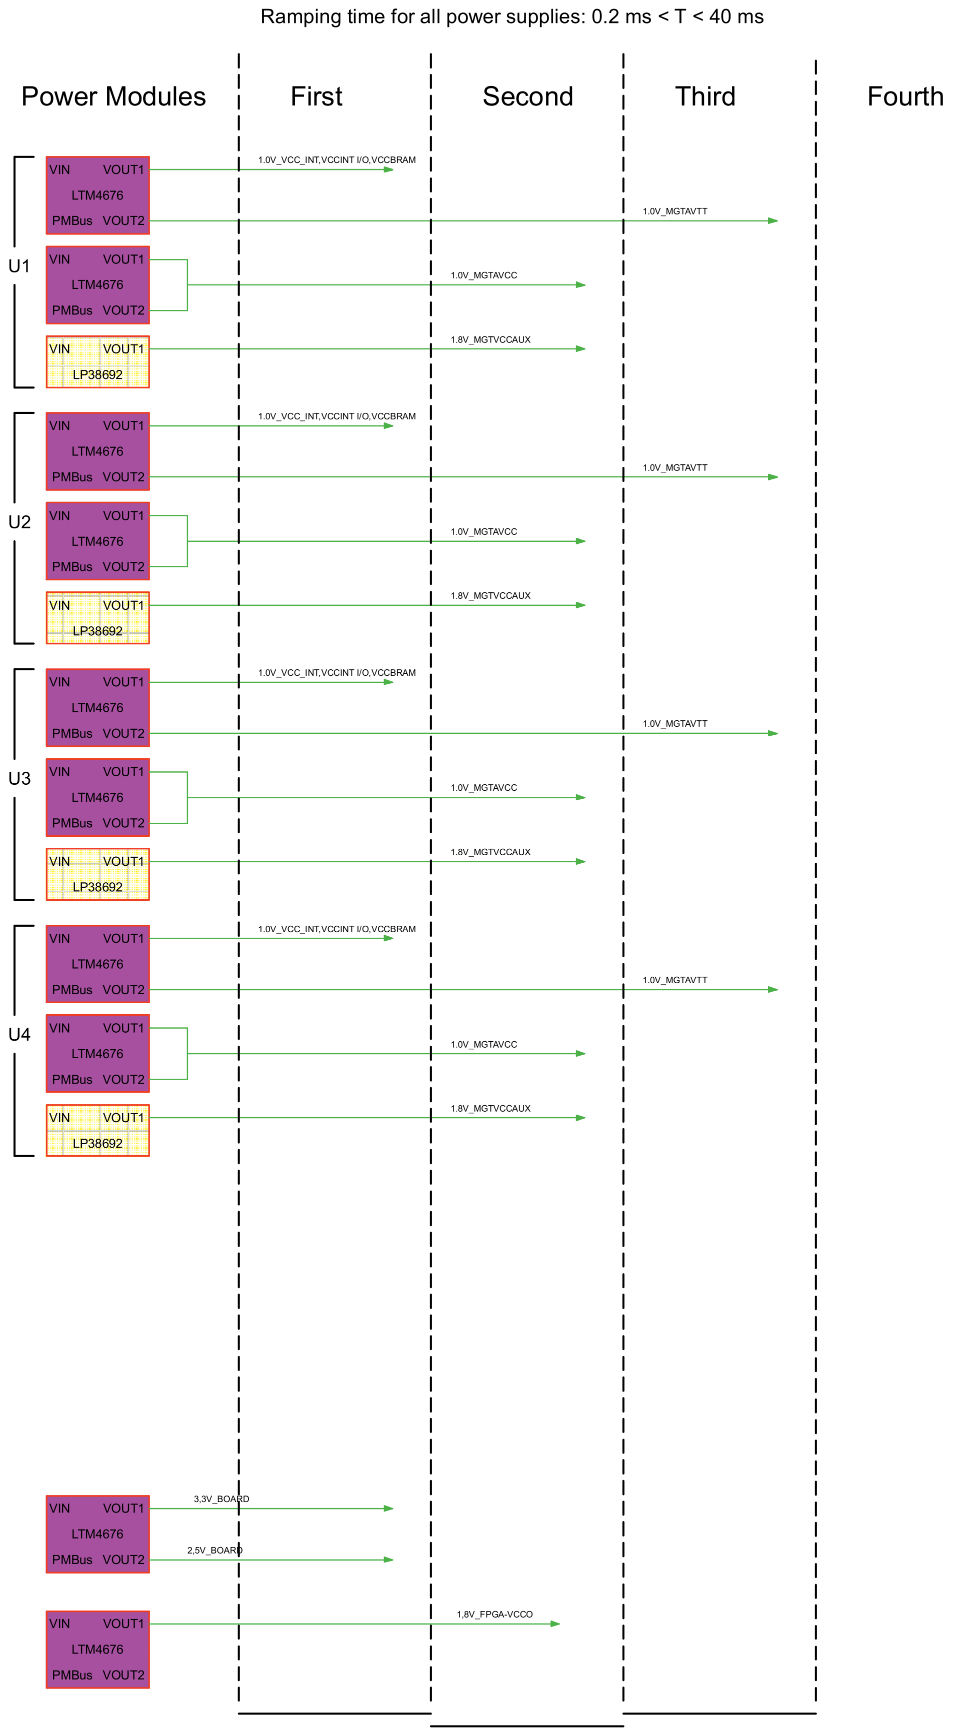

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 2/11

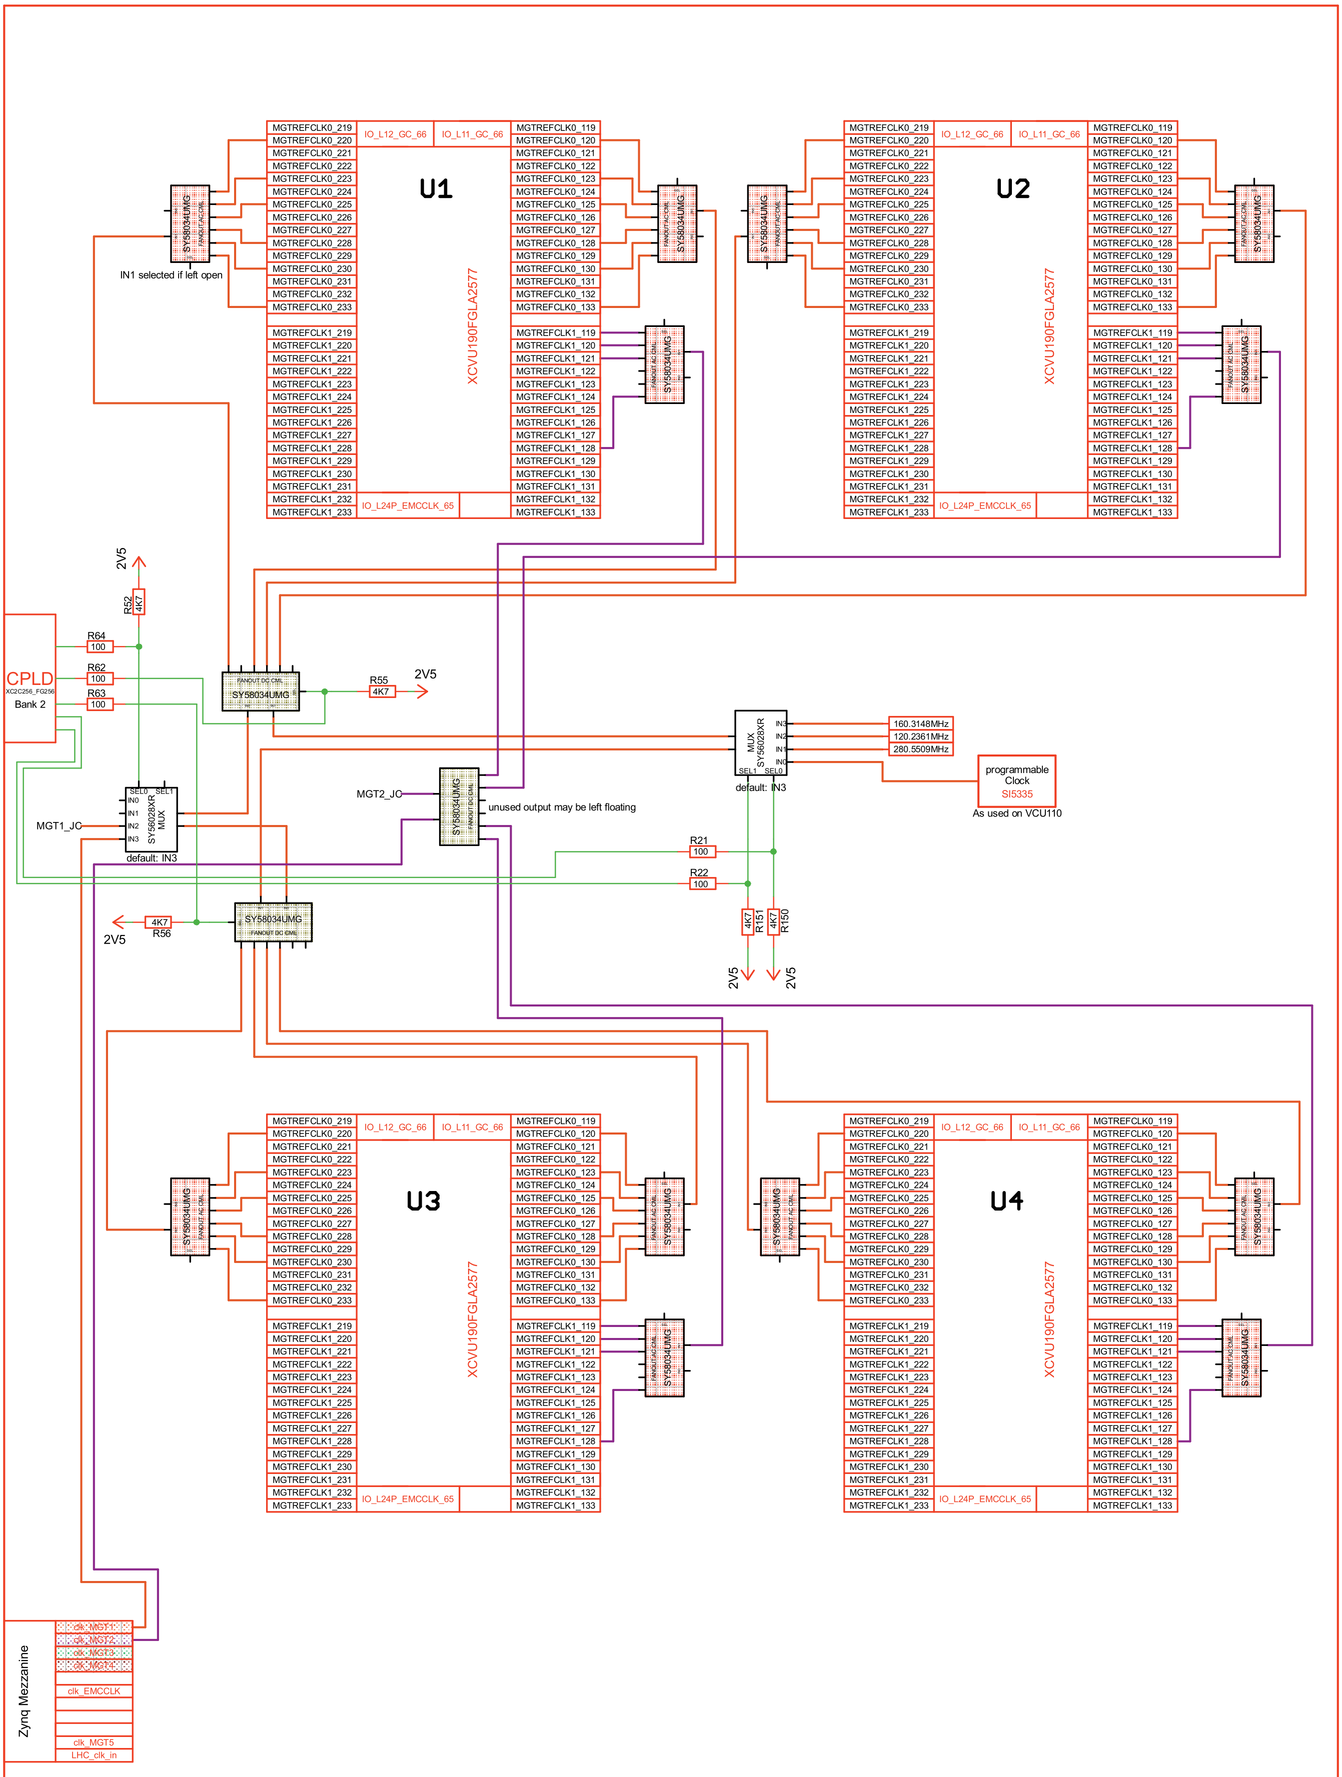

MGT1 + MGT2 CLOCK

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 3/11

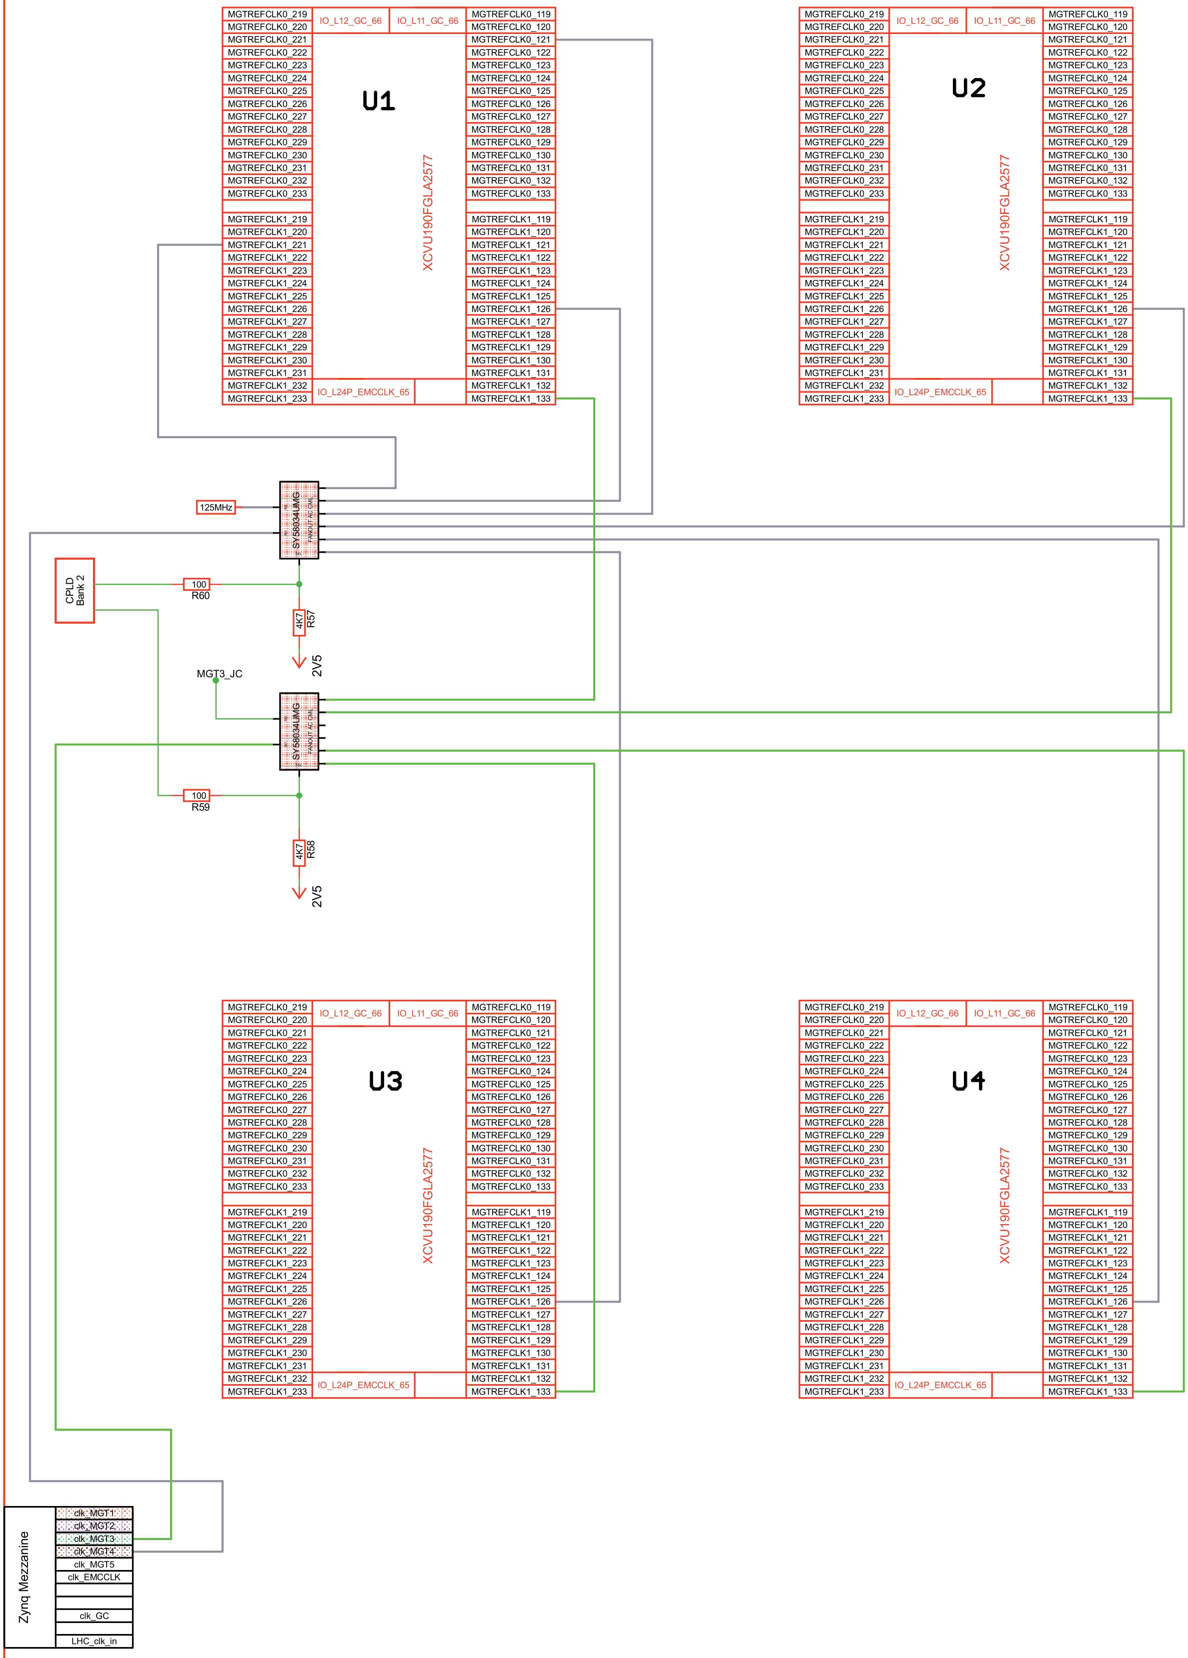

MGT3 + MGT4 +IPBUS CLOCK

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 4/11

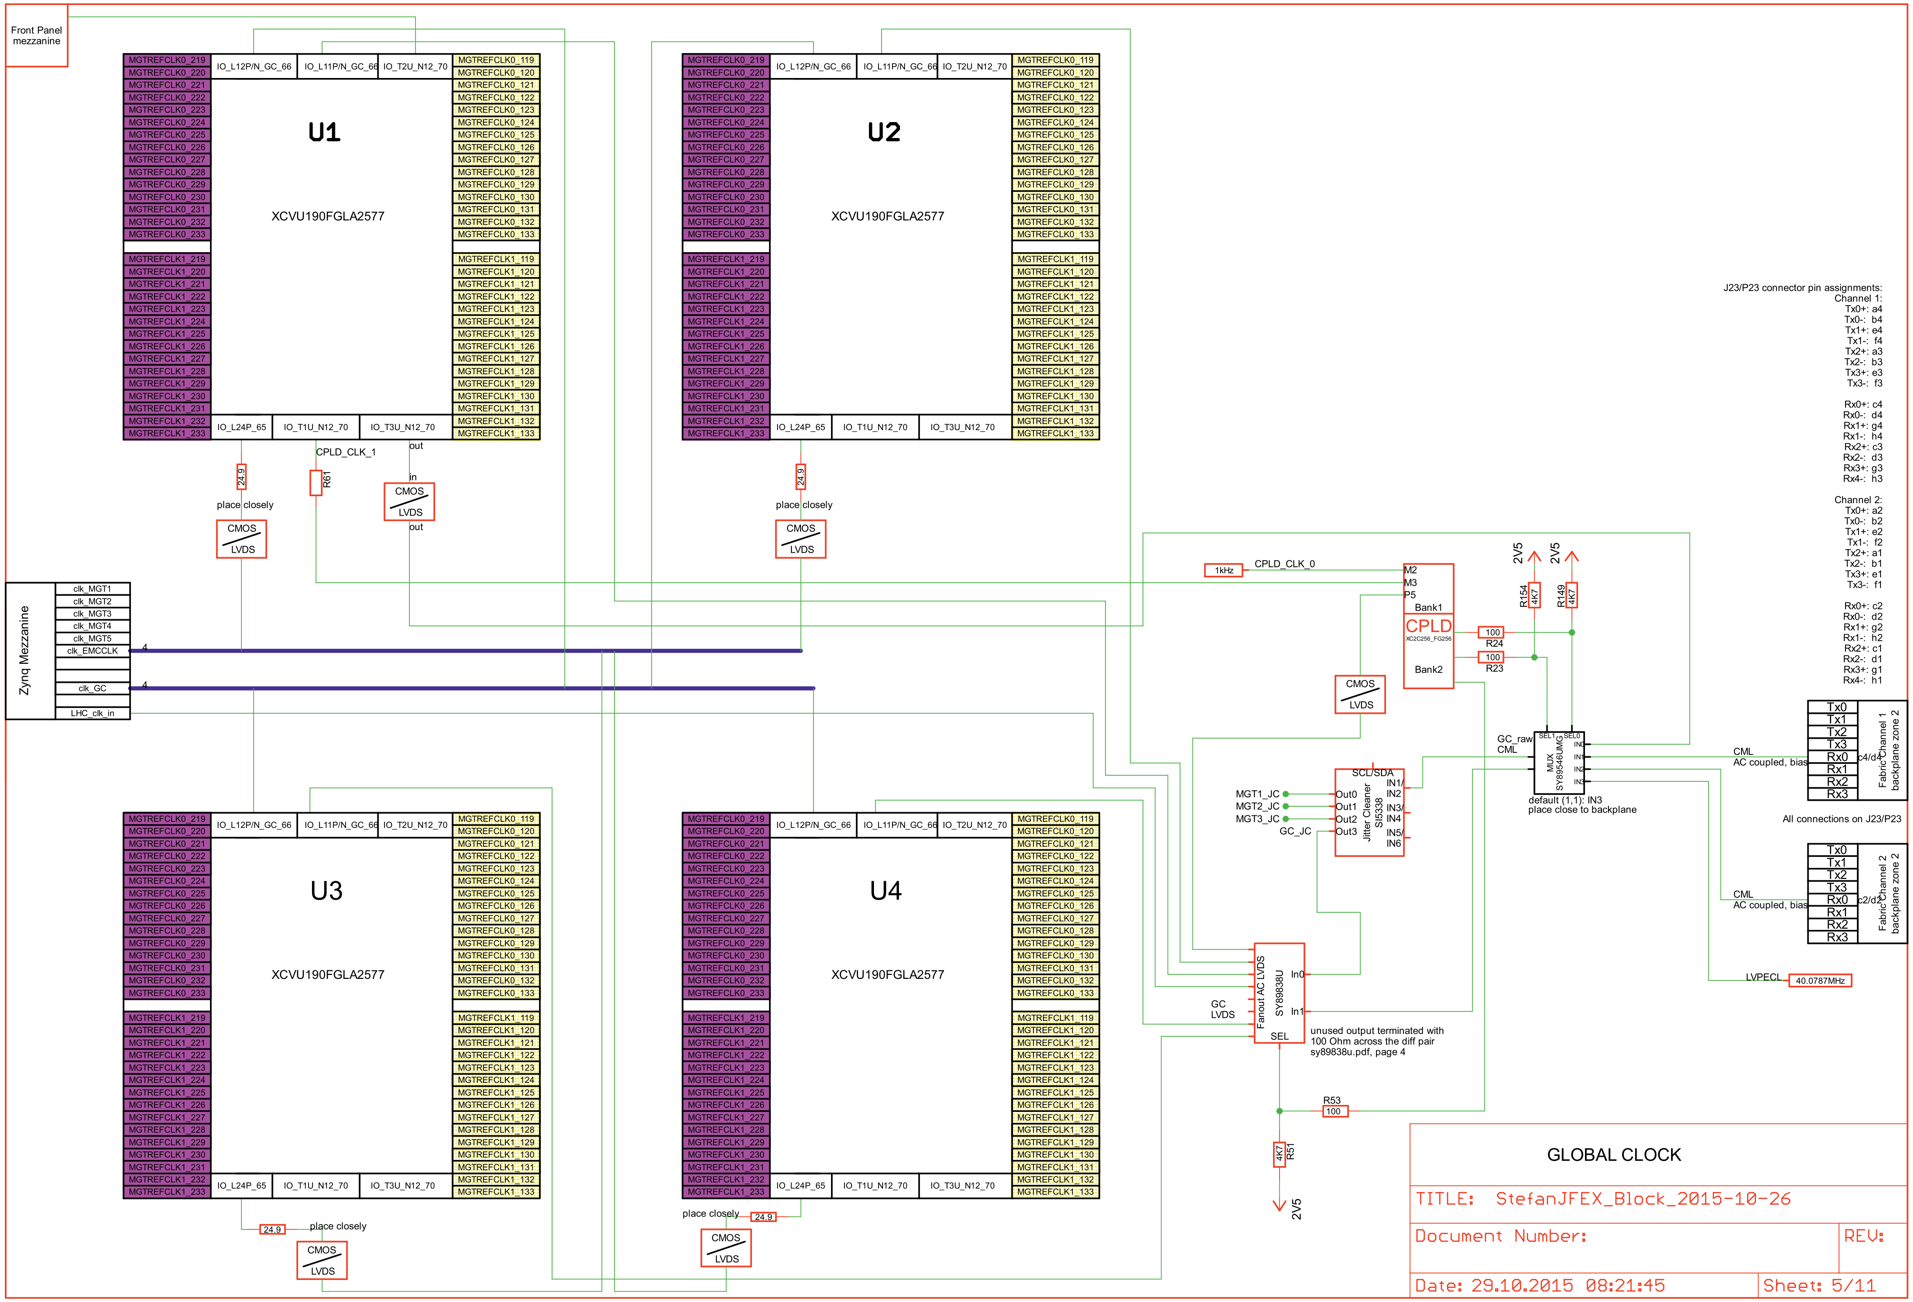

J23/P23 connector pin assignments:
Channel 1:
Tx0+: a4
Tx0-: b4
Tx1+: e4
Tx1-: f4
Tx2+: a3
Tx2-: b3
Tx3+: e3
Tx3-: f3

Rx0+: c4
Rx0-: d4
Rx1+: g4
Rx1-: h4
Rx2+: c3
Rx2-: d3
Rx3+: g3
Rx4-: h3

Channel 2:
Tx0+: a2
Tx0-: b2
Tx1+: e2
Tx1-: f2
Tx2+: a1
Tx2-: b1
Tx3+: e1
Tx3-: f1

Rx0+: c2
Rx0-: d2
Rx1+: g2
Rx1-: h2
Rx2+: c1
Rx2-: d1
Rx3+: g1
Rx4-: h1

GLOBAL CLOCK

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 5/11

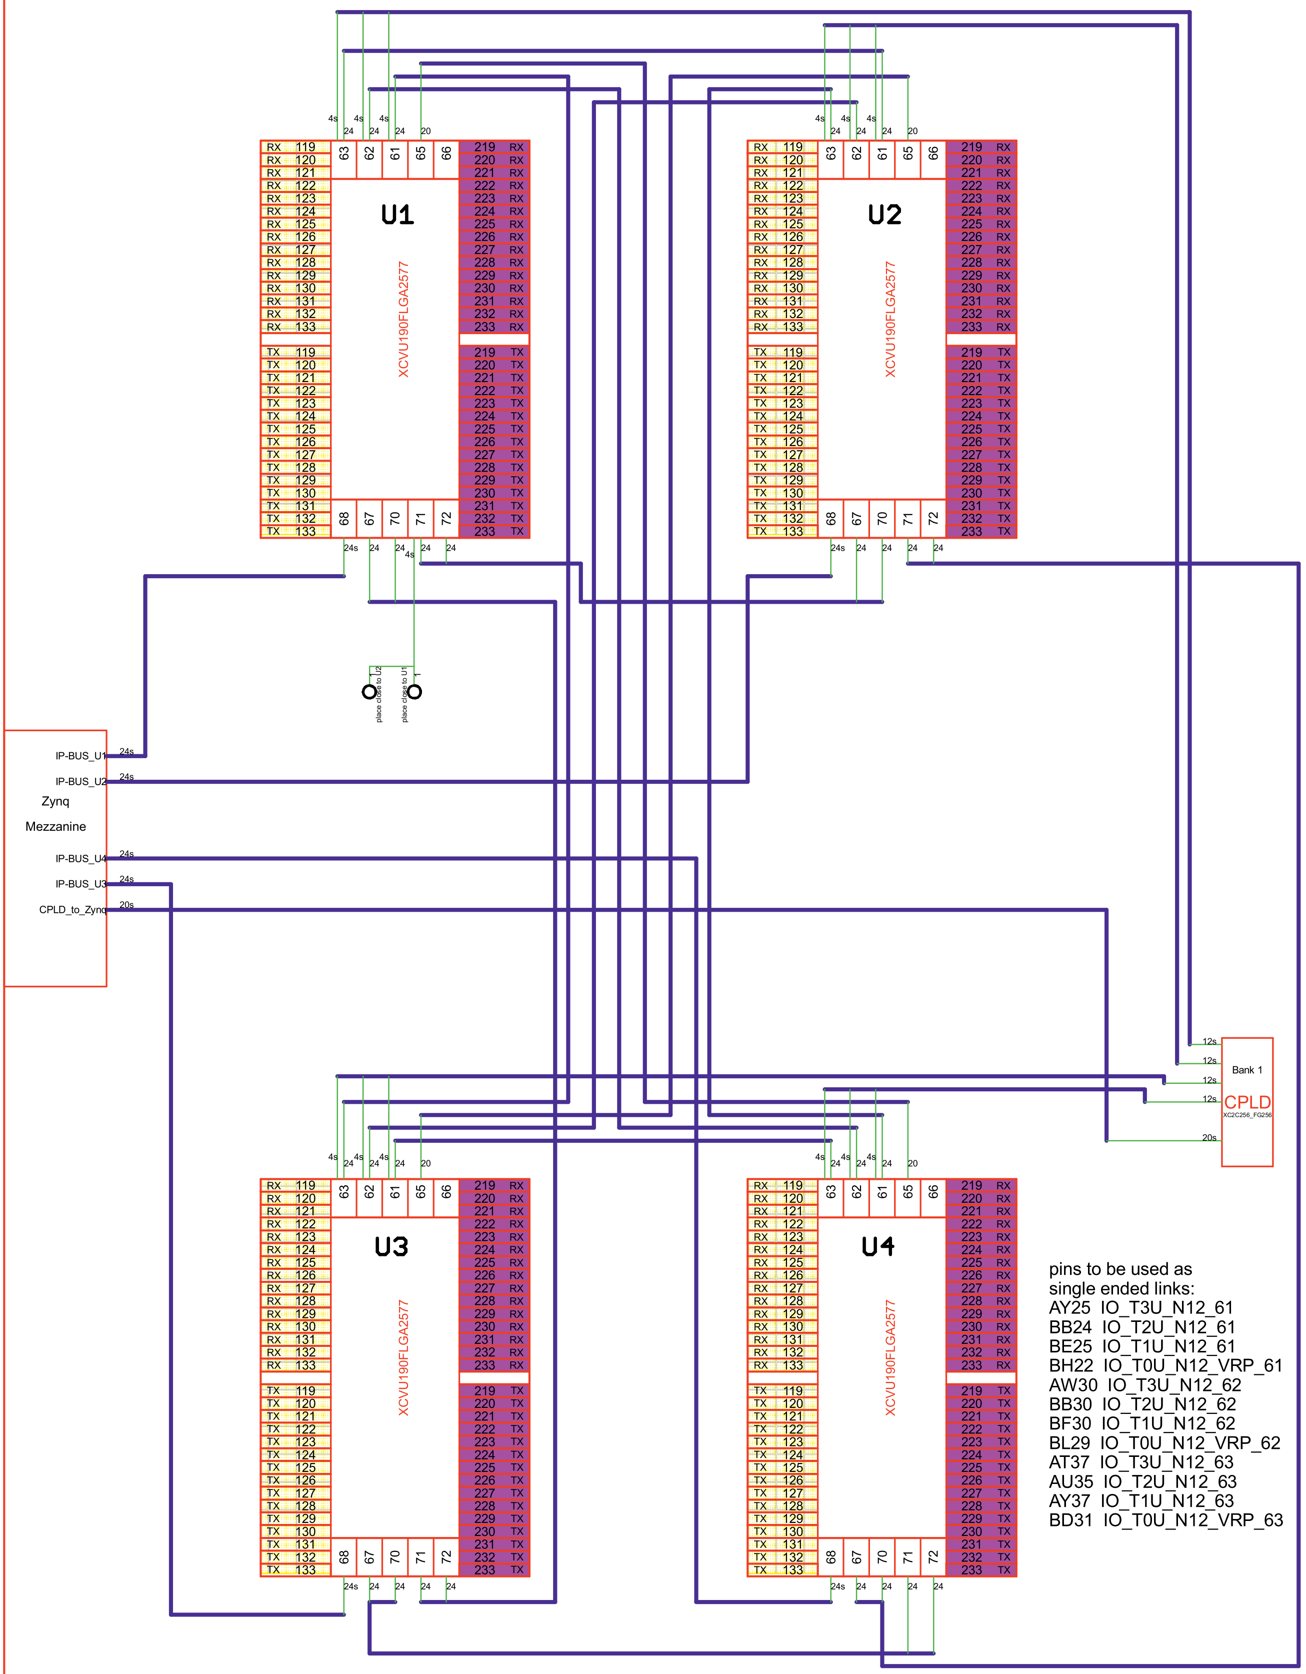

pins to be used as
single ended links:
AY25 IO_T3U_N12_61
BB24 IO_T2U_N12_61
BE25 IO_T1U_N12_61
BH22 IO_T0U_N12_VRP_61
AW30 IO_T3U_N12_62
BB30 IO_T2U_N12_62
BF30 IO_T1U_N12_62
BL29 IO_T0U_N12_VRP_62
AT37 IO_T3U_N12_63
AU35 IO_T2U_N12_63
AY37 IO_T1U_N12_63
BD31 IO_T0U_N12_VRP_63

connect complete banks
swapping of complete banks allowed
swapping within banks allowed
if possible swap identically for all FPGAs
connection of banks within busses:
bank 67 to 71
bank 70 to 72

Bank 68 (XCVU190) single-ended pins on :
IO_L(1-12)P/N 24 single

Bank 65 (XCVU190) differential pairs on :
IO_L1P/N, IO_L(3-20)P/N,IO_L23P/N 20 diff

PARALLEL I/0

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 6/11

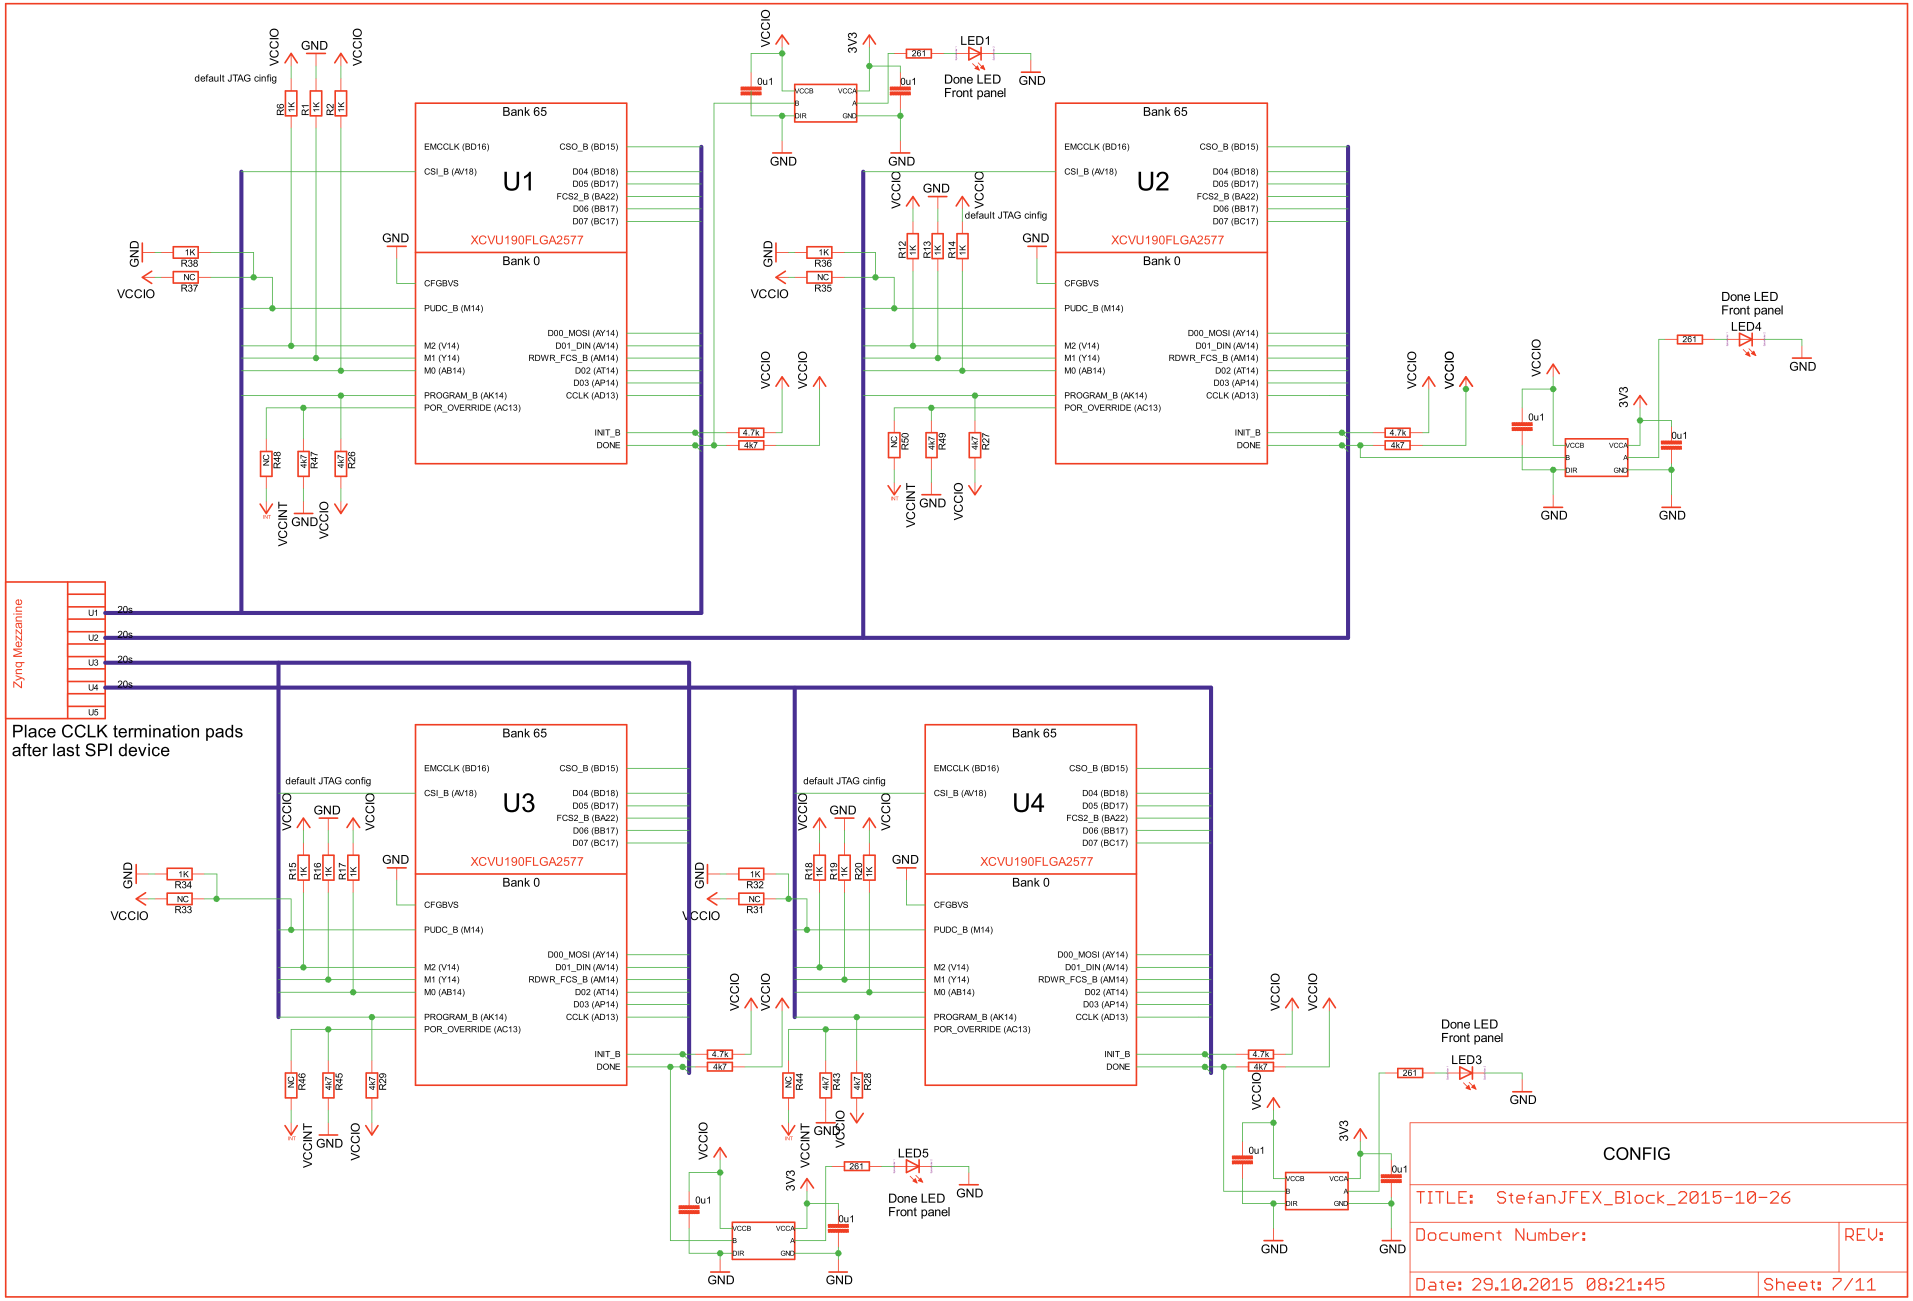

CONFIG

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 7/11

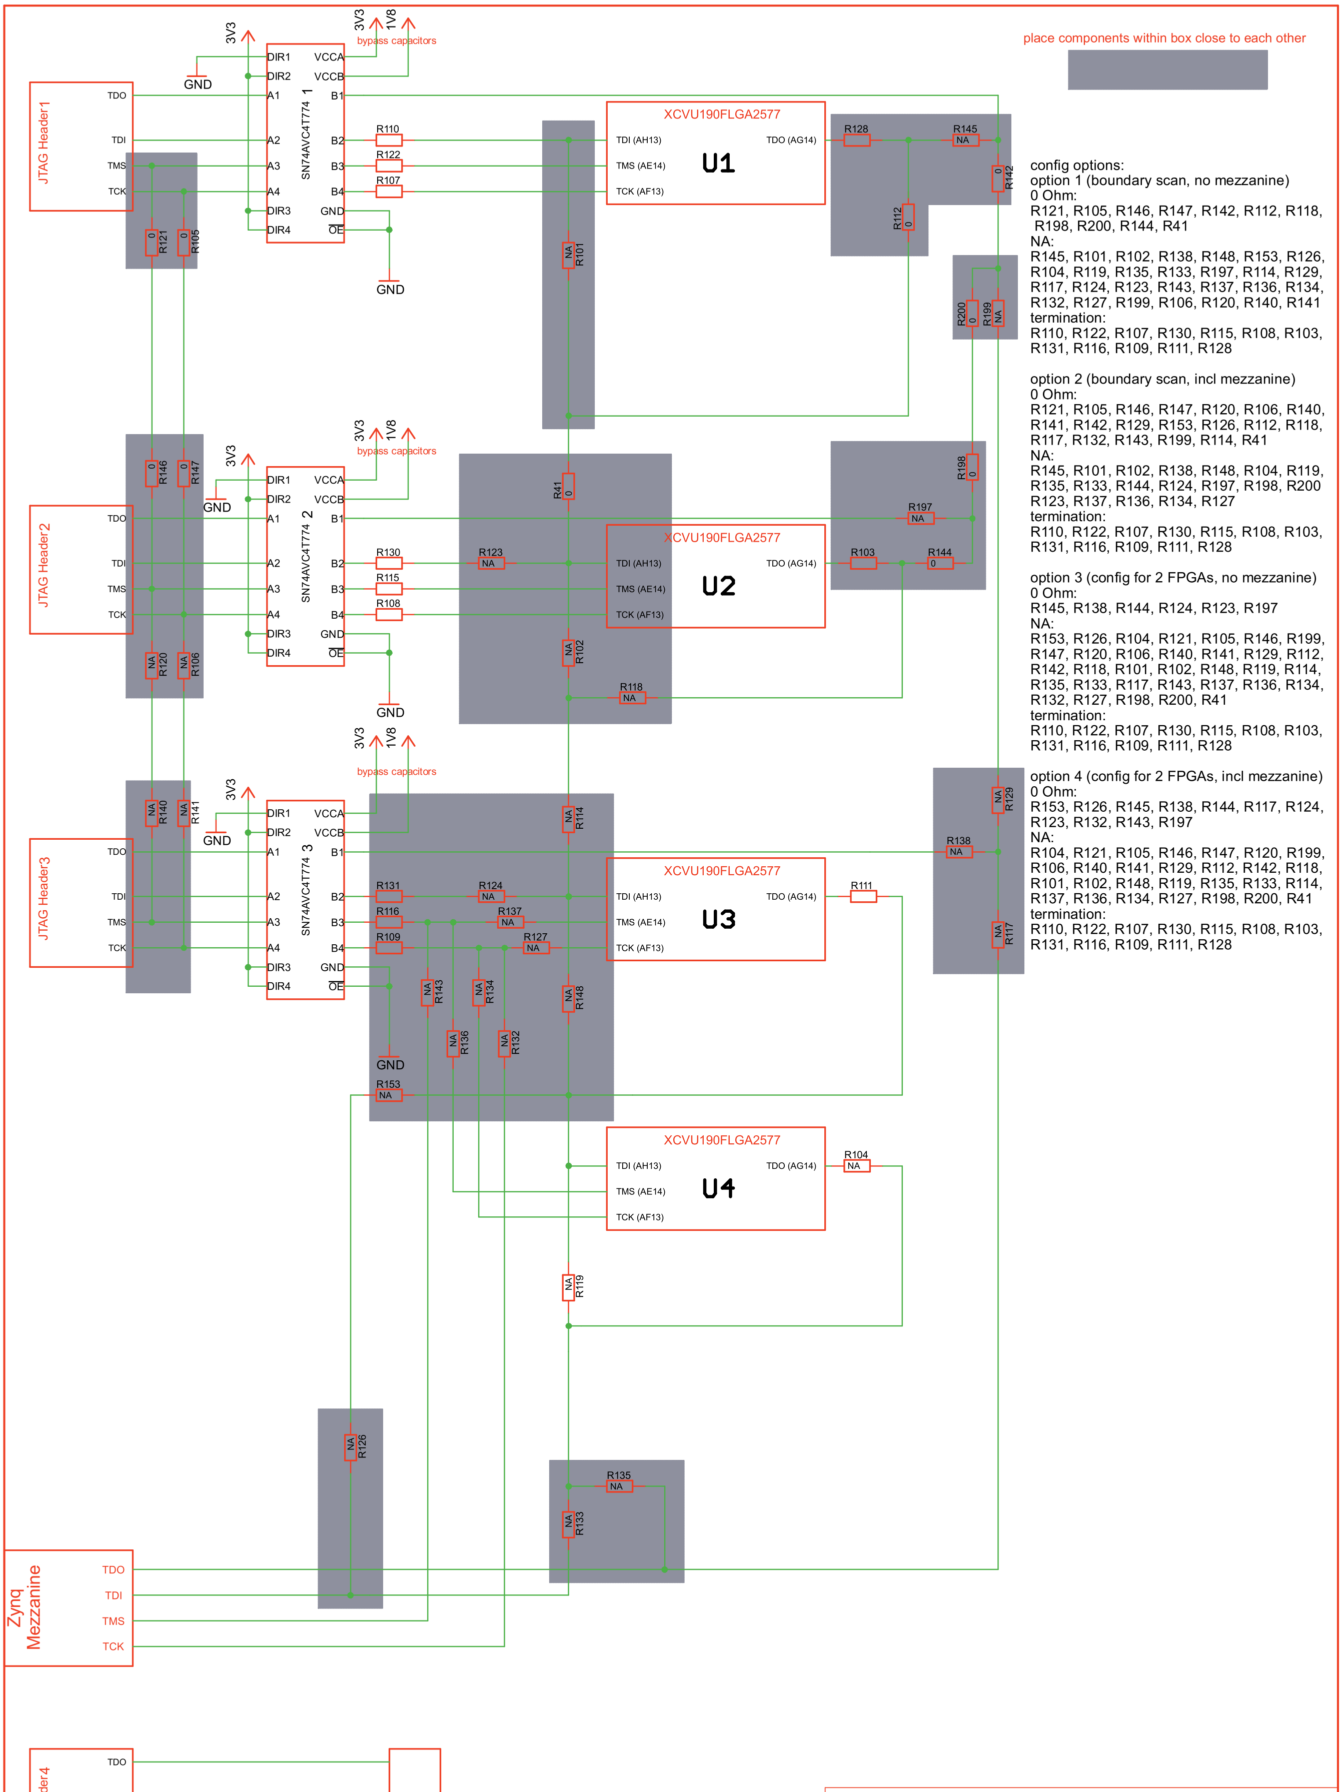

config options:
option 1 (boundary scan, no mezzanine)
0 Ohm:
R121, R105, R146, R147, R142, R112, R118, R198, R200, R144, R41
NA:
R145, R101, R102, R138, R148, R153, R126, R104, R119, R135, R133, R197, R114, R129, R117, R124, R123, R143, R137, R136, R134, R132, R127, R199, R106, R120, R140, R141
termination:
R110, R122, R107, R130, R115, R108, R103, R131, R116, R109, R111, R128

option 2 (boundary scan, incl mezzanine)
0 Ohm:
R121, R105, R146, R147, R120, R106, R140, R141, R142, R129, R153, R126, R112, R118, R117, R132, R143, R199, R114, R41
NA:
R145, R101, R102, R138, R148, R104, R119, R135, R133, R144, R124, R197, R198, R200 R123, R137, R136, R134, R127
termination:
R110, R122, R107, R130, R115, R108, R103, R131, R116, R109, R111, R128

option 3 (config for 2 FPGAs, no mezzanine)
0 Ohm:
R145, R138, R144, R124, R123, R197
NA:
R153, R126, R104, R121, R105, R146, R199, R147, R120, R106, R140, R141, R129, R112, R142, R118, R101, R102, R148, R119, R114, R135, R133, R117, R143, R137, R136, R134, R132, R127, R198, R200, R41
termination:
R110, R122, R107, R130, R115, R108, R103, R131, R116, R109, R111, R128

option 4 (config for 2 FPGAs, incl mezzanine)
0 Ohm:
R153, R126, R145, R138, R144, R117, R124, R123, R132, R143, R197
NA:
R104, R121, R105, R146, R147, R120, R199, R106, R140, R141, R129, R112, R142, R118, R101, R102, R148, R119, R135, R133, R114, R137, R136, R134, R127, R198, R200, R41
termination:
R110, R122, R107, R130, R115, R108, R103, R131, R116, R109, R111, R128

JTAG

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 8/11

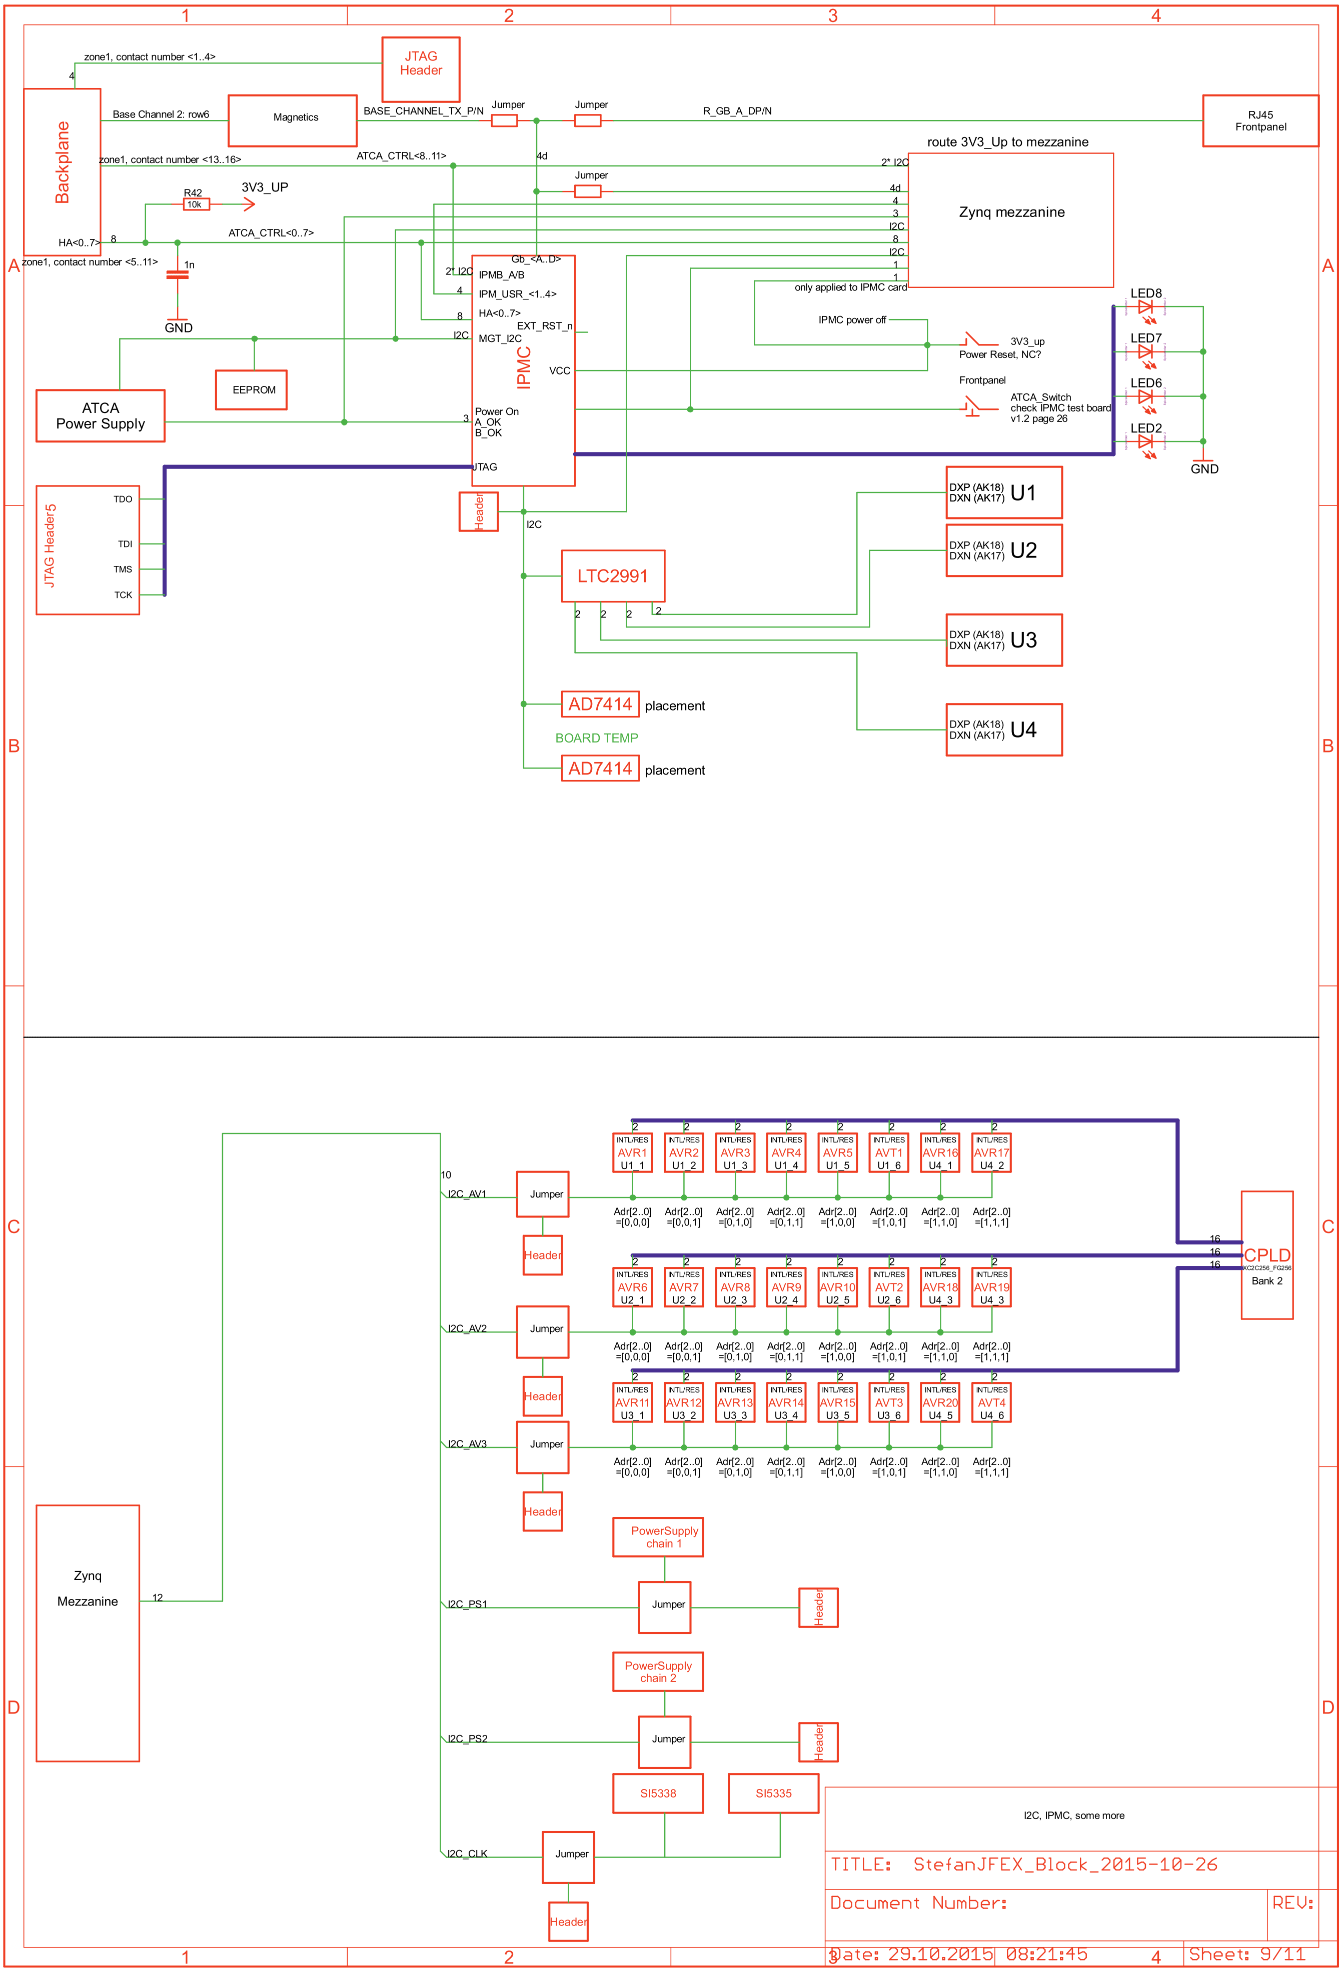

I2C, IPMC, some more

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 9/11

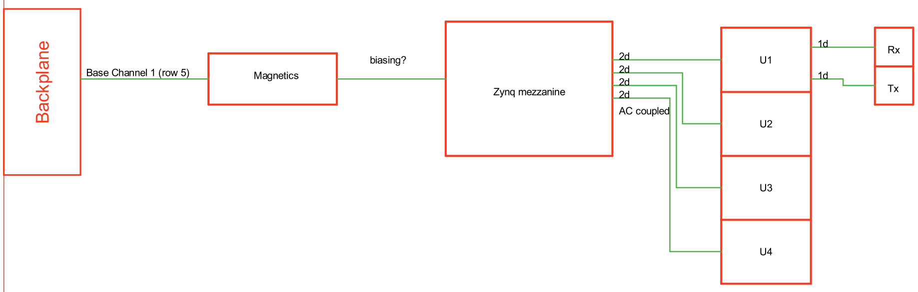

IPBus (summary)

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 10/11

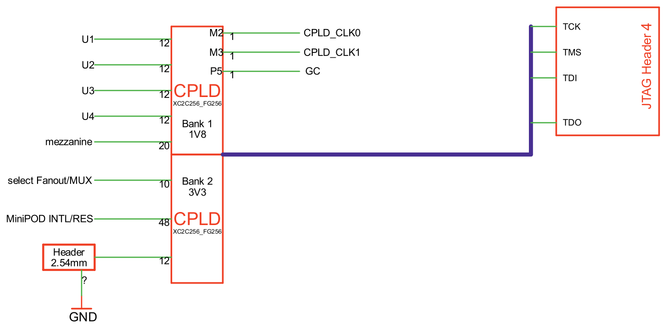

CPLD (summary)

TITLE: StefanJFEX_Block_2015-10-26

Document Number:

REV:

Date: 29.10.2015 08:21:45

Sheet: 11/11

Appendix C

FPGA Overview

Complete overview of the considered FPGA families including their respective provided resources. The available FPGAs are shown in figure C.1 and figure C.2 for Intel and Xilinx devices, respectively.

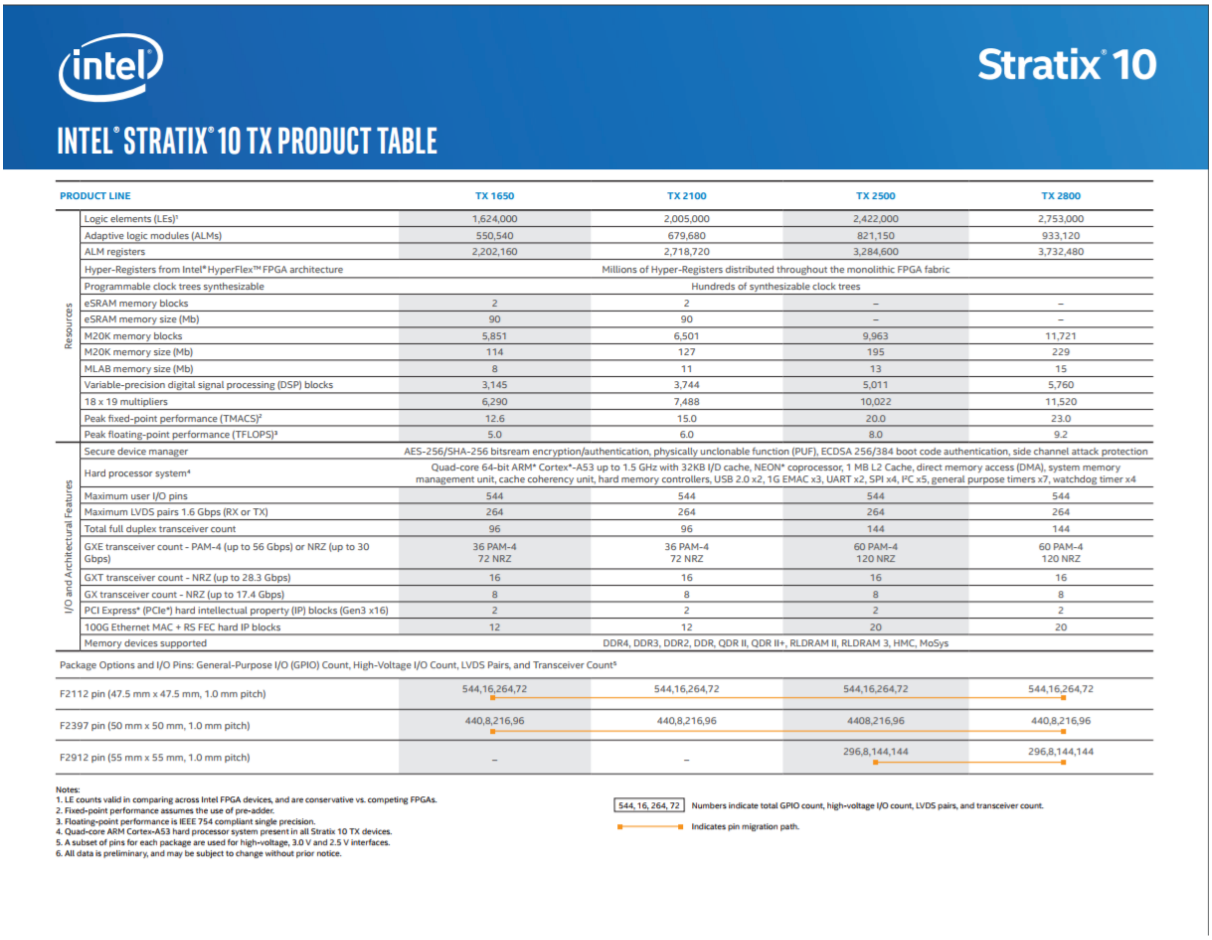

intel

Stratix 10

INTEL® STRATIX® 10 TX PRODUCT TABLE

	PRODUCT LINE	TX 1650	TX 2100	TX 2500	TX 2800
Resources	Logic elements (LEs)[1]	1,624,000	2,005,000	2,422,000	2,753,000
	Adaptive logic modules (ALMs)	550,540	679,680	821,150	933,120
	ALM registers	2,202,160	2,718,720	3,284,600	3,732,480
	Hyper-Registers from Intel® HyperFlex™ FPGA architecture	Millions of Hyper-Registers distributed throughout the monolithic FPGA fabric			
	Programmable clock trees synthesizable	Hundreds of synthesizable clock trees			
	eSRAM memory blocks	2	2	–	–
	eSRAM memory size (Mb)	90	90	–	–
	M20K memory blocks	5,851	6,501	9,963	11,721
	M20K memory size (Mb)	114	127	195	229
	MLAB memory size (Mb)	8	11	13	15
	Variable-precision digital signal processing (DSP) blocks	3,145	3,744	5,011	5,760
	18 x 19 multipliers	6,290	7,488	10,022	11,520
	Peak fixed-point performance (TMACS)[2]	12.6	15.0	20.0	23.0
	Peak floating-point performance (TFLOPS)[3]	5.0	6.0	8.0	9.2
I/O and Architectural Features	Secure device manager	AES-256/SHA-256 bitsream encryption/authentication, physically unclonable function (PUF), ECDSA 256/384 boot code authentication, side channel attack protection			
	Hard processor system[4]	Quad-core 64-bit ARM* Cortex*-A53 up to 1.5 GHz with 32KB I/D cache, NEON* coprocessor, 1 MB L2 Cache, direct memory access (DMA), system memory management unit, cache coherency unit, hard memory controllers, USB 2.0 x2, 1G EMAC x3, UART x2, SPI x4, I²C x5, general purpose timers x7, watchdog timer x4			
	Maximum user I/O pins	544	544	544	544
	Maximum LVDS pairs 1.6 Gbps (RX or TX)	264	264	264	264
	Total full duplex transceiver count	96	96	144	144
	GXE transceiver count - PAM-4 (up to 56 Gbps) or NRZ (up to 30 Gbps)	36 PAM-4 72 NRZ	36 PAM-4 72 NRZ	60 PAM-4 120 NRZ	60 PAM-4 120 NRZ
	GXT transceiver count - NRZ (up to 28.3 Gbps)	16	16	16	16
	GX transceiver count - NRZ (up to 17.4 Gbps)	8	8	8	8
	PCI Express* (PCIe*) hard intellectual property (IP) blocks (Gen3 x16)	2	2	2	2
	100G Ethernet MAC + RS FEC hard IP blocks	12	12	20	20
	Memory devices supported	DDR4, DDR3, DDR2, DDR, QDR II, QDR II+, RLDRAM II, RLDRAM 3, HMC, MoSys			
Package Options and I/O Pins: General-Purpose I/O (GPIO) Count, High-Voltage I/O Count, LVDS Pairs, and Transceiver Count[5]					
	F2112 pin (47.5 mm x 47.5 mm, 1.0 mm pitch)	544,16,264,72	544,16,264,72	544,16,264,72	544,16,264,72
	F2397 pin (50 mm x 50 mm, 1.0 mm pitch)	440,8,216,96	440,8,216,96	4408,216,96	440,8,216,96
	F2912 pin (55 mm x 55 mm, 1.0 mm pitch)	–	–	296,8,144,144	296,8,144,144

Notes:
1. LE counts valid in comparing across Intel FPGA devices, and are conservative vs. competing FPGAs.
2. Fixed-point performance assumes the use of pre-adder.
3. Floating-point performance is IEEE 754 compliant single precision.
4. Quad-core ARM Cortex-A53 hard processor system present in all Stratix 10 TX devices.
5. A subset of pins for each package are used for high-voltage, 3.0 V and 2.5 V interfaces.
6. All data is preliminary, and may be subject to change without prior notice.

544, 16, 264, 72 Numbers indicate total GPIO count, high-voltage I/O count, LVDS pairs, and transceiver count.

Indicates pin migration path.

Figure C.1: Overview of FPGAs in the Intel Stratix 10 TX series. The exact Intel FPGA architecture and naming scheme differs from the description in 7 yet the basic principle remains the same. Taken from [70].

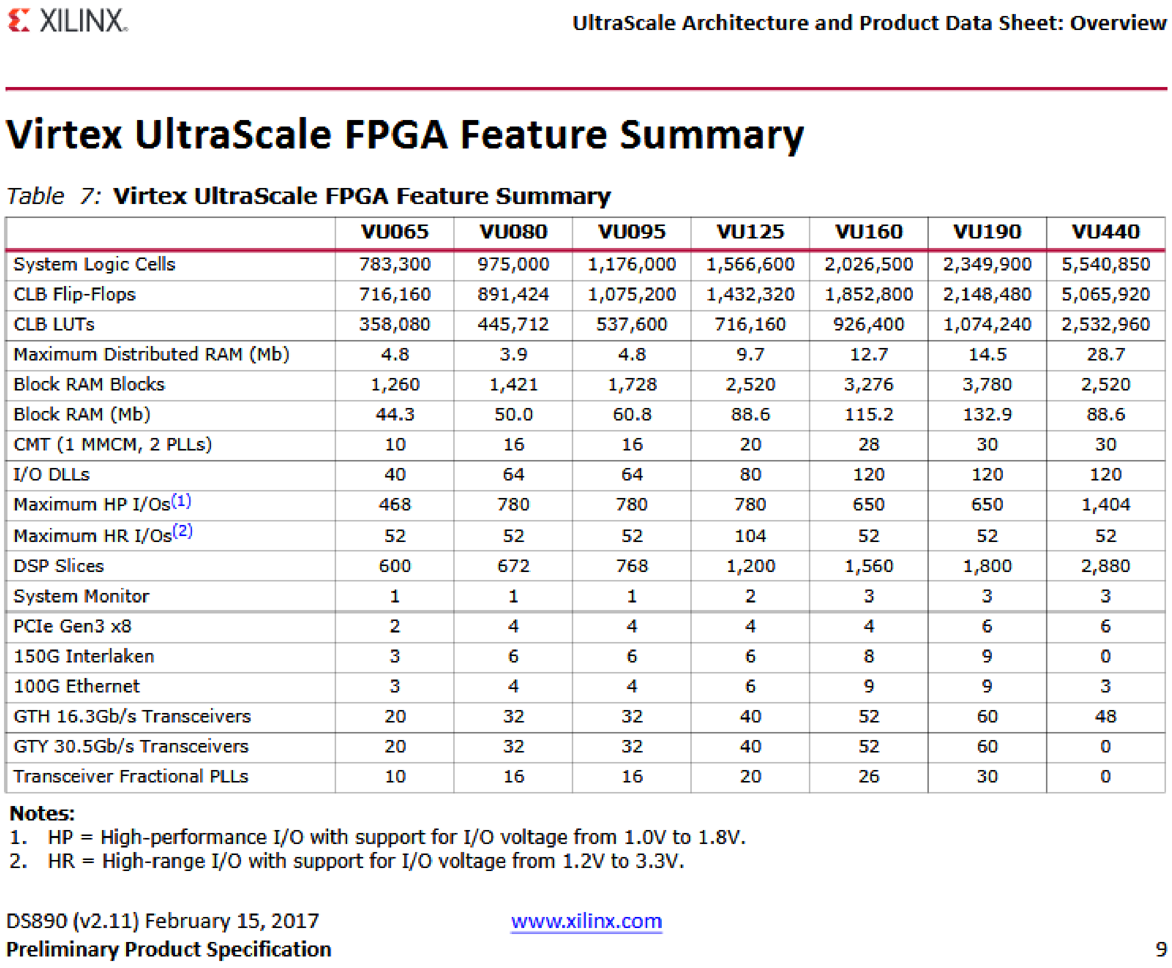

XILINX. **UltraScale Architecture and Product Data Sheet: Overview**

Virtex UltraScale FPGA Feature Summary

Table 7: **Virtex UltraScale FPGA Feature Summary**

	VU065	VU080	VU095	VU125	VU160	VU190	VU440
System Logic Cells	783,300	975,000	1,176,000	1,566,600	2,026,500	2,349,900	5,540,850
CLB Flip-Flops	716,160	891,424	1,075,200	1,432,320	1,852,800	2,148,480	5,065,920
CLB LUTs	358,080	445,712	537,600	716,160	926,400	1,074,240	2,532,960
Maximum Distributed RAM (Mb)	4.8	3.9	4.8	9.7	12.7	14.5	28.7
Block RAM Blocks	1,260	1,421	1,728	2,520	3,276	3,780	2,520
Block RAM (Mb)	44.3	50.0	60.8	88.6	115.2	132.9	88.6
CMT (1 MMCM, 2 PLLs)	10	16	16	20	28	30	30
I/O DLLs	40	64	64	80	120	120	120
Maximum HP I/Os(1)	468	780	780	780	650	650	1,404
Maximum HR I/Os(2)	52	52	52	104	52	52	52
DSP Slices	600	672	768	1,200	1,560	1,800	2,880
System Monitor	1	1	1	2	3	3	3
PCIe Gen3 x8	2	4	4	4	4	6	6
150G Interlaken	3	6	6	6	8	9	0
100G Ethernet	3	4	4	6	9	9	3
GTH 16.3Gb/s Transceivers	20	32	32	40	52	60	48
GTY 30.5Gb/s Transceivers	20	32	32	40	52	60	0
Transceiver Fractional PLLs	10	16	16	20	26	30	0

Notes:

1. HP = High-performance I/O with support for I/O voltage from 1.0V to 1.8V.
2. HR = High-range I/O with support for I/O voltage from 1.2V to 3.3V.

DS890 (v2.11) February 15, 2017 www.xilinx.com
Preliminary Product Specification 9

Figure C.2: Overview of FPGAs in the Xilinx Virtex UltraScale series. Taken from [59].

Appendix D

XPE 2016.3 for XCVU190

Excerpt showing the main contributions to the power consumption calculated for the XCVU190 using the Xilinx Power Estimator 2016.3 [65]. Basis for the used numbers of utilised resources were a combination of known use cases like high-speed links and worst-case scenarios based on maximum possible utilisation for internal computing resources where exact numbers were unavailable.
An overview of resulting currents per input voltage and total power drawn is shown in figure D.1. Detailed views of the assumed settings are shown for GTH and GTY high-speed links in figure D.2 and figure D.3, respectively. Computing resources are listed for DSP slices (D.4), BRAM (D.5), and remaining logic resources (D.6). General purpose I/O links for low-speed and non-real-time connections between the FPGAs are shown in (D.7).

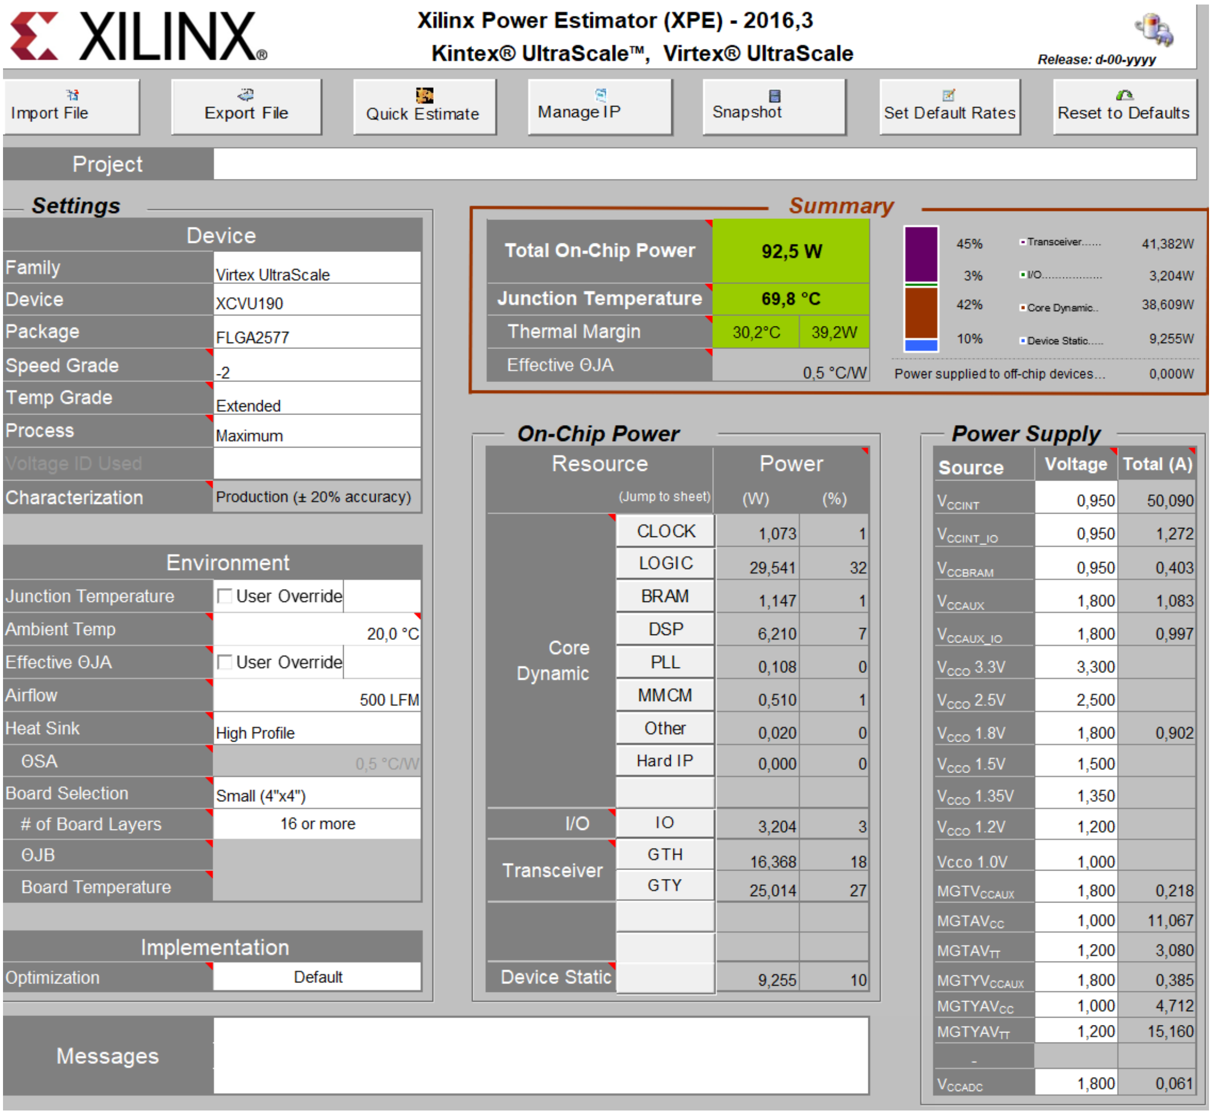

Settings

Device	
Family	Virtex UltraScale
Device	XCVU190
Package	FLGA2577
Speed Grade	-2
Temp Grade	Extended
Process	Maximum
Voltage ID Used	
Characterization	Production (± 20% accuracy)

Environment	
Junction Temperature	☐ User Override
Ambient Temp	20,0 °C
Effective ΘJA	☐ User Override
Airflow	500 LFM
Heat Sink	High Profile
ΘSA	0,5 °C/W
Board Selection	Small (4"x4")
# of Board Layers	16 or more
ΘJB	
Board Temperature	

Implementation	
Optimization	Default

Summary

Total On-Chip Power	92,5 W	
Junction Temperature	69,8 °C	
Thermal Margin	30,2°C	39,2W
Effective ΘJA	0,5 °C/W	

%		
45%	Transceiver......	41,382W
3%	I/O................	3,204W
42%	Core Dynamic..	38,609W
10%	Device Static.....	9,255W

Power supplied to off-chip devices... 0,000W

On-Chip Power

Resource	(Jump to sheet)	Power (W)	(%)
Core Dynamic	CLOCK	1,073	1
	LOGIC	29,541	32
	BRAM	1,147	1
	DSP	6,210	7
	PLL	0,108	0
	MMCM	0,510	1
	Other	0,020	0
	Hard IP	0,000	0
I/O	IO	3,204	3
Transceiver	GTH	16,368	18
	GTY	25,014	27
Device Static		9,255	10

Power Supply

Source	Voltage	Total (A)
V_{CCINT}	0,950	50,090
V_{CCINT_IO}	0,950	1,272
V_{CCBRAM}	0,950	0,403
V_{CCAUX}	1,800	1,083
V_{CCAUX_IO}	1,800	0,997
V_{CCO} 3.3V	3,300	
V_{CCO} 2.5V	2,500	
V_{CCO} 1.8V	1,800	0,902
V_{CCO} 1.5V	1,500	
V_{CCO} 1.35V	1,350	
V_{CCO} 1.2V	1,200	
Vcco 1.0V	1,000	
$MGTV_{CCAUX}$	1,800	0,218
$MGTAV_{CC}$	1,000	11,067
$MGTAV_{TT}$	1,200	3,080
$MGTYV_{CCAUX}$	1,800	0,385
$MGTYAV_{CC}$	1,000	4,712
$MGTYAV_{TT}$	1,200	15,160
-		
V_{CCADC}	1,800	0,061

FIGURE D.1: Overview of resulting currents drawn per input voltage and total power drawn.

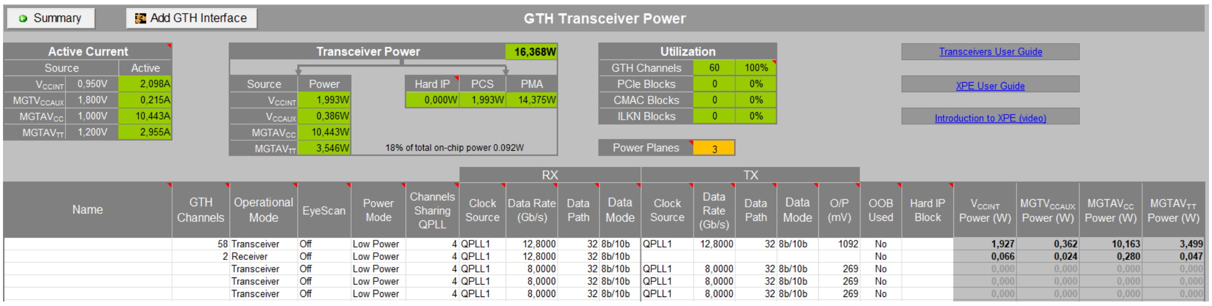

Active Current		
Source		Active
V_{CCINT}	0,950V	2,098A
$MGTV_{CCAUX}$	1,800V	0,215A
$MGTAV_{CC}$	1,000V	10,443A
$MGTAV_{TT}$	1,200V	2,955A

Transceiver Power 16,368W

Source	Power
V_{CCINT}	1,993W
V_{CCAUX}	0,386W
$MGTAV_{CC}$	10,443W
$MGTAV_{TT}$	3,546W

Hard IP	PCS	PMA
0,000W	1,993W	14,375W

18% of total on-chip power 0.092W

Utilization		
GTH Channels	60	100%
PCIe Blocks	0	0%
CMAC Blocks	0	0%
ILKN Blocks	0	0%

Power Planes 3

						RX				TX										
Name	GTH Channels	Operational Mode	EyeScan	Power Mode	Channels Sharing QPLL	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	O/P (mV)	OOB Used	Hard IP Block	V_{CCINT} Power (W)	$MGTV_{CCAUX}$ Power (W)	$MGTAV_{CC}$ Power (W)	$MGTAV_{TT}$ Power (W)
	58	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL1	12,8000	32	8b/10b	1092	No		1,927	0,362	10,163	3,499
	2	Receiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b						No		0,066	0,024	0,280	0,047
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	269	No		0,000	0,000	0,000	0,000
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	269	No		0,000	0,000	0,000	0,000
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	269	No		0,000	0,000	0,000	0,000

FIGURE D.2: Settings used to determine the power consumption of GTH high-speed links including the data duplication using PMA loopback mode.

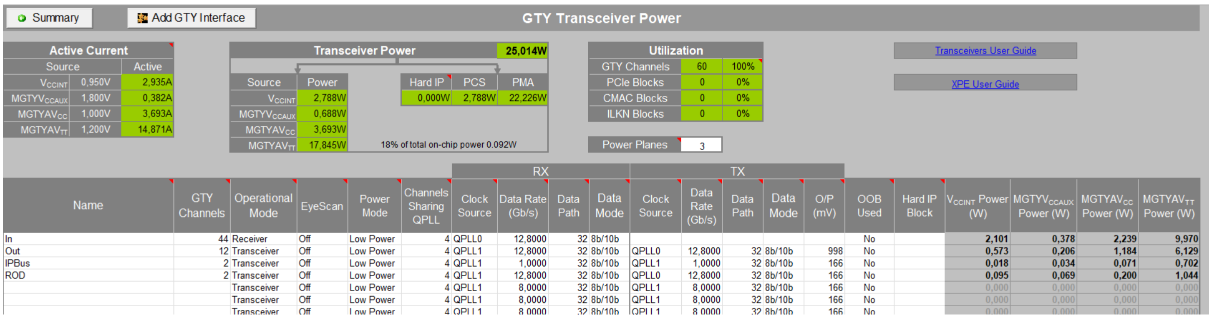

GTY Transceiver Power

Active Current		
Source		Active
V_{CCINT}	0,950V	2,935A
$MGTYV_{CCAUX}$	1,800V	0,382A
$MGTYAV_{CC}$	1,000V	3,693A
$MGTYAV_{TT}$	1,200V	14,871A

Transceiver Power	25,014W
Source	Power
V_{CCINT}	2,788W
$MGTYV_{CCAUX}$	0,688W
$MGTYAV_{CC}$	3,693W
$MGTYAV_{TT}$	17,845W

Hard IP	PCS	PMA
0,000W	2,788W	22,226W

18% of total on-chip power 0.092W

Utilization		
GTY Channels	60	100%
PCIe Blocks	0	0%
CMAC Blocks	0	0%
ILKN Blocks	0	0%

Power Planes 3

						RX				TX										
Name	GTY Channels	Operational Mode	EyeScan	Power Mode	Channels Sharing QPLL	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	O/P (mV)	OOB Used	Hard IP Block	V_{CCINT} Power (W)	$MGTYV_{CCAUX}$ Power (W)	$MGTYAV_{CC}$ Power (W)	$MGTYAV_{TT}$ Power (W)
In	44	Receiver	Off	Low Power	4	QPLL0	12,8000	32	8b/10b						No		2,101	0,378	2,239	9,970
Out	12	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL0	12,8000	32	8b/10b	998	No		0,573	0,206	1,184	6,129
IPBus	2	Transceiver	Off	Low Power	4	QPLL1	1,0000	32	8b/10b	QPLL1	1,0000	32	8b/10b	166	No		0,018	0,034	0,071	0,702
ROD	2	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL0	12,8000	32	8b/10b	166	No		0,095	0,069	0,200	1,044
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	166	No		0,000	0,000	0,000	0,000
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	166	No		0,000	0,000	0,000	0,000
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	166	No		0,000	0,000	0,000	0,000

FIGURE D.3: Settings used to determine the power consumption of GTY high-speed links including input, output, and non-real-time functionalities.

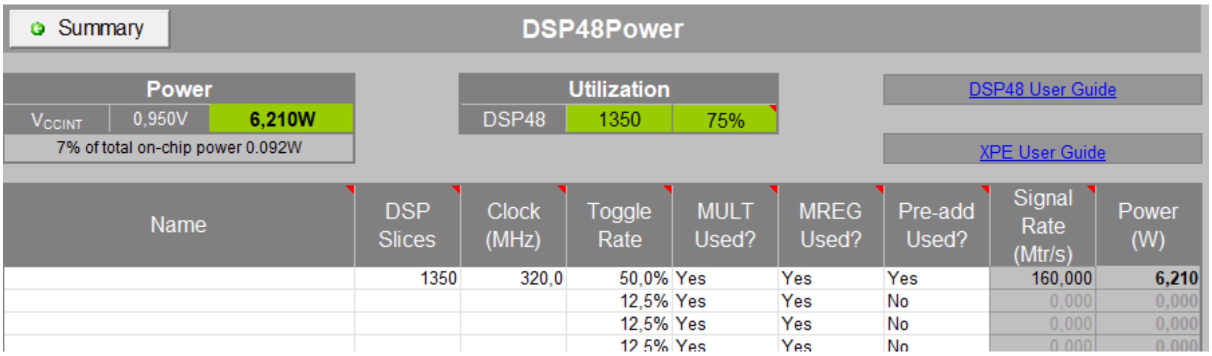

DSP48Power

Power		
V_{CCINT}	0,950V	6,210W
7% of total on-chip power 0.092W		

Utilization		
DSP48	1350	75%

Name	DSP Slices	Clock (MHz)	Toggle Rate	MULT Used?	MREG Used?	Pre-add Used?	Signal Rate (Mtr/s)	Power (W)
	1350	320,0	50,0%	Yes	Yes	Yes	160,000	6,210
			12,5%	Yes	Yes	No	0,000	0,000
			12,5%	Yes	Yes	No	0,000	0,000
			12,5%	Yes	Yes	No	0,000	0,000

FIGURE D.4: Settings used to determine the power consumption of DSP slices assuming a 75 % resource utilisation.

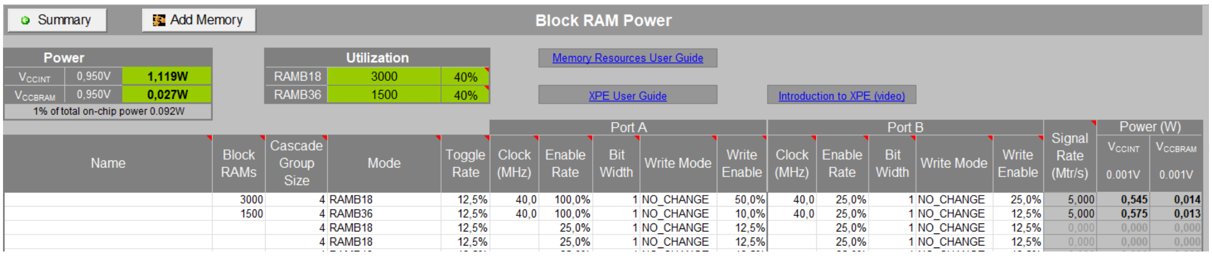

Block RAM Power

Power		
V_{CCINT}	0,950V	1,119W
V_{CCBRAM}	0,950V	0,027W
1% of total on-chip power 0.092W		

Utilization		
RAMB18	3000	40%
RAMB36	1500	40%

					Port A					Port B						Power (W)	
Name	Block RAMs	Cascade Group Size	Mode	Toggle Rate	Clock (MHz)	Enable Rate	Bit Width	Write Mode	Write Enable	Clock (MHz)	Enable Rate	Bit Width	Write Mode	Write Enable	Signal Rate (Mtr/s)	V_{CCINT} 0.001V	V_{CCBRAM} 0.001V
	3000	4	RAMB18	12,5%	40,0	100,0%	1	NO_CHANGE	50,0%	40,0	25,0%	1	NO_CHANGE	25,0%	5,000	0,545	0,014
	1500	4	RAMB36	12,5%	40,0	100,0%	1	NO_CHANGE	10,0%	40,0	25,0%	1	NO_CHANGE	12,5%	5,000	0,575	0,013
		4	RAMB18	12,5%		25,0%	1	NO_CHANGE	12,5%		25,0%	1	NO_CHANGE	12,5%	0,000	0,000	0,000
		4	RAMB18	12,5%		25,0%	1	NO_CHANGE	12,5%		25,0%	1	NO_CHANGE	12,5%	0,000	0,000	0,000

FIGURE D.5: Settings used to determine the power consumption of DSP slices assuming a low utilisation as the resource is not used by the algorithms and only required for read-out and control functionalities.

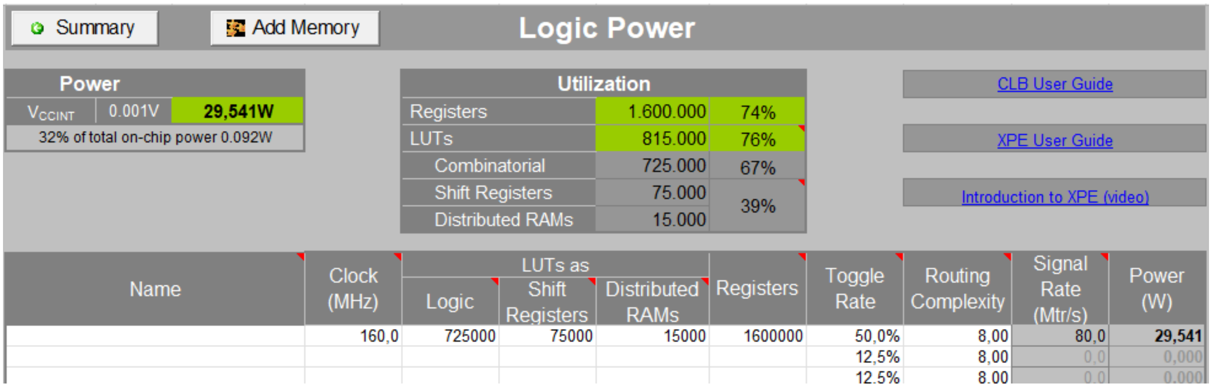

Logic Power

Power		
V_{CCINT}	0.001V	29,541W
32% of total on-chip power 0.092W		

Utilization		
Registers	1.600.000	74%
LUTs	815.000	76%
Combinatorial	725.000	67%
Shift Registers	75.000	39%
Distributed RAMs	15.000	

Name	Clock (MHz)	LUTs as Logic	LUTs as Shift Registers	LUTs as Distributed RAMs	Registers	Toggle Rate	Routing Complexity	Signal Rate (Mtr/s)	Power (W)
	160,0	725000	75000	15000	1600000	50,0%	8,00	80,0	29,541
						12,5%	8,00	0,0	0,000
						12,5%	8,00	0,0	0,000

FIGURE D.6: Settings used to determine the power consumption of generic logic resources.

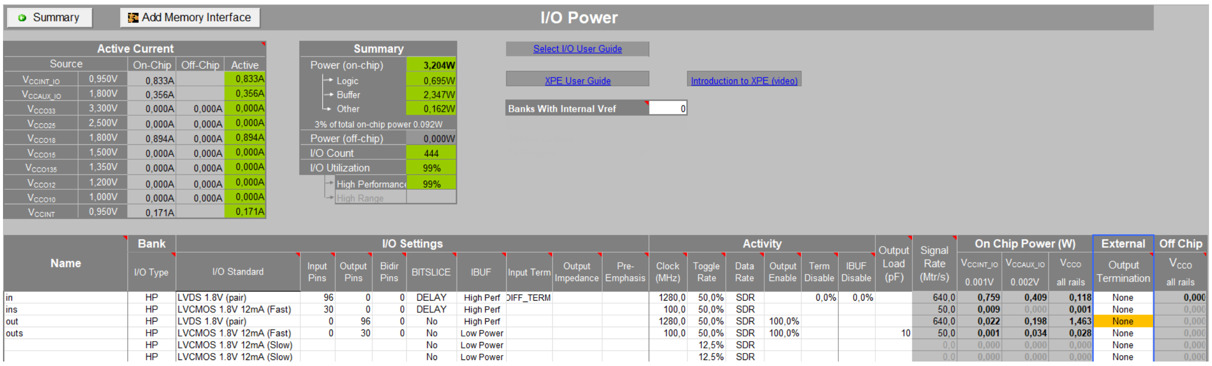

I/O Power

Active Current				
Source		On-Chip	Off-Chip	Active
V_{CCINT_IO}	0,950V	0,833A		0,833A
V_{CCAUX_IO}	1,800V	0,356A		0,356A
V_{CCO33}	3,300V	0,000A	0,000A	0,000A
V_{CCO25}	2,500V	0,000A	0,000A	0,000A
V_{CCO18}	1,800V	0,894A	0,000A	0,894A
V_{CCO15}	1,500V	0,000A	0,000A	0,000A
V_{CCO135}	1,350V	0,000A	0,000A	0,000A
V_{CCO12}	1,200V	0,000A	0,000A	0,000A
V_{CCO10}	1,000V	0,000A	0,000A	0,000A
V_{CCINT}	0,950V	0,171A		0,171A

Summary	
Power (on-chip)	3,204W
→ Logic	0,695W
→ Buffer	2,347W
→ Other	0,162W
3% of total on-chip power 0.092W	
Power (off-chip)	0,000W
I/O Count	444
I/O Utilization	99%
→ High Performance	99%
→ High Range	

Banks With Internal Vref 0

Name	Bank: I/O Type	I/O Settings: I/O Standard	Input Pins	Output Pins	Bidir Pins	BITSLICE	IBUF	Input Term	Output Impedance	Pre-Emphasis
in	HP	LVDS 1.8V (pair)	96	0	0	DELAY	High Perf	DIFF_TERM		
ins	HP	LVCMOS 1.8V 12mA (Fast)	30	0	0	DELAY	High Perf			
out	HP	LVDS 1.8V (pair)	0	96	0	No	High Perf			
outs	HP	LVCMOS 1.8V 12mA (Fast)	0	30	0	No	Low Power			
	HP	LVCMOS 1.8V 12mA (Slow)				No	Low Power			
	HP	LVCMOS 1.8V 12mA (Slow)				No	Low Power			

Name	Activity: Clock (MHz)	Toggle Rate	Data Rate	Output Enable	Term Disable	IBUF Disable	Output Load (pF)	Signal Rate (Mtr/s)	On Chip Power (W): V_{CCINT_IO} 0.001V	V_{CCAUX_IO} 0.002V	V_{CCO} all rails	External: Output Termination	Off Chip: V_{CCO} all rails
in	1280,0	50,0%	SDR		0,0%	0,0%		640,0	0,759	0,409	0,118	None	0,000
ins	100,0	50,0%	SDR					50,0	0,009	0,000	0,001	None	0,000
out	1280,0	50,0%	SDR	100,0%				640,0	0,022	0,198	1,463	None	0,000
outs	100,0	50,0%	SDR	100,0%			10	50,0	0,001	0,034	0,028	None	0,000
		12,5%	SDR					0,0	0,000	0,000	0,000	None	0,000
		12,5%	SDR					0,0	0,000	0,000	0,000	None	0,000

FIGURE D.7: Settings used to determine the power consumption of general purpose I/O links assuming a high data rate between the FPGAs.

Appendix E

XPE 2017.2 for XCVU9P

Excerpt showing the main contributions to the power consumption calculated for the XCVU9P using the Xilinx Power Estimator 2017.2 [66]. To allow for a direct comparison between the FPGA series, the same resource allocation is used as for the XCVU190 in appendix D even though in some cases, like the utilisation of DSP slices, the UltraScale+ design could offer significantly more resources.
An overview of resulting currents per input voltage and total power drawn is shown in figure E.1. A detailed view of the assumed settings for high-speed links is shown in figure E.2. Computing resources are listed for DSP slices (E.3), BRAM (E.4), and remaining logic resources (E.5). General purpose I/O links for low-speed and non-real-time connections between the FPGAs are shown in (E.6).

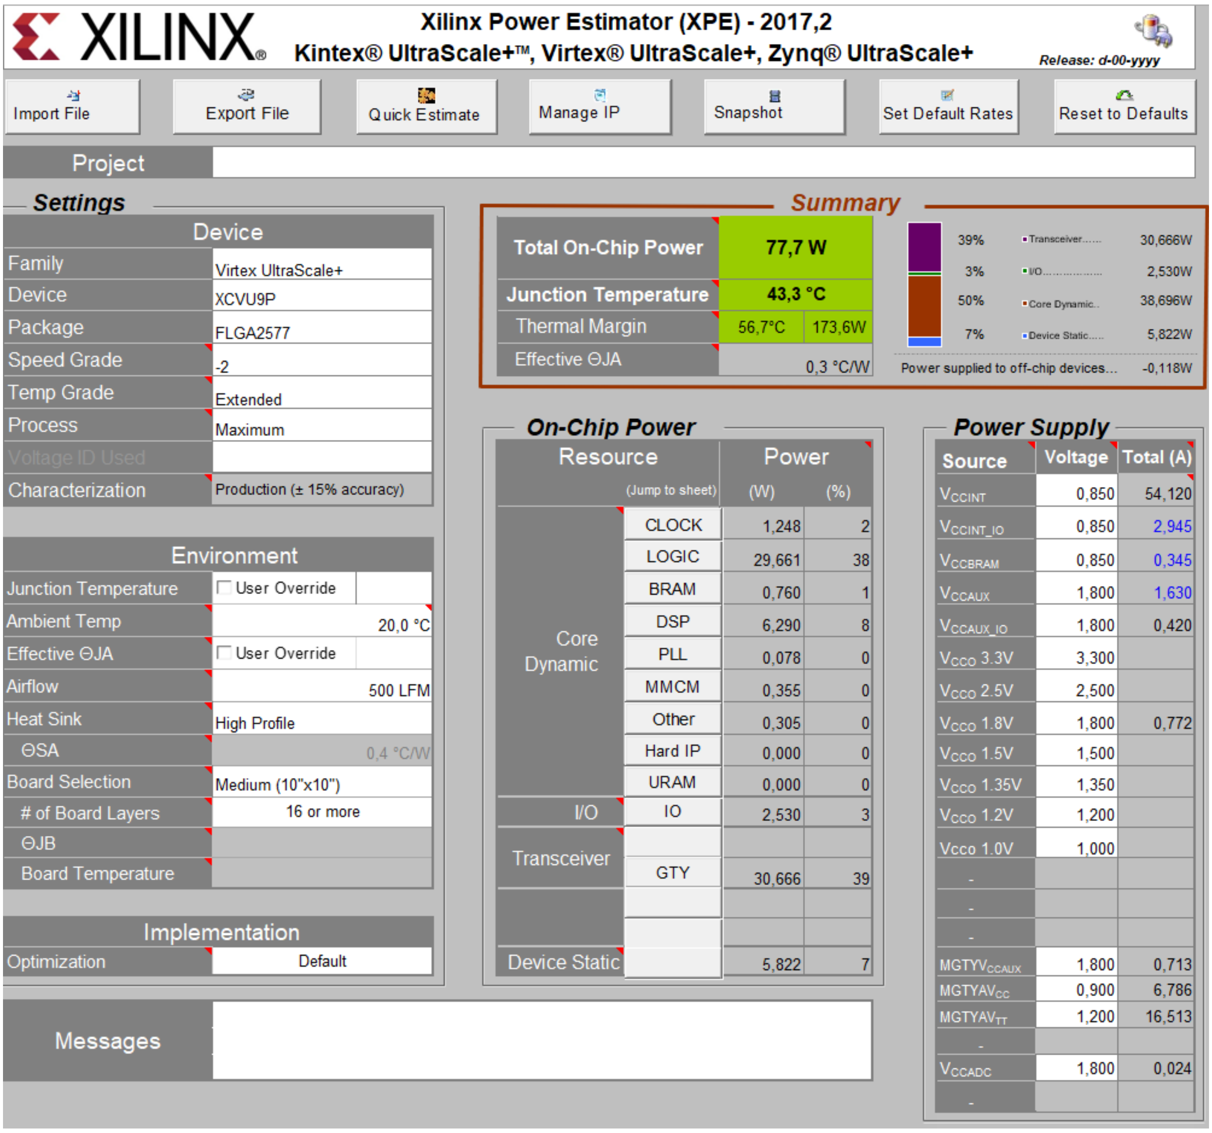

Xilinx Power Estimator (XPE) - 2017,2
Kintex® UltraScale+™, Virtex® UltraScale+, Zynq® UltraScale+

Project

Settings

Device	
Family	Virtex UltraScale+
Device	XCVU9P
Package	FLGA2577
Speed Grade	-2
Temp Grade	Extended
Process	Maximum
Voltage ID Used	
Characterization	Production (± 15% accuracy)

Environment		
Junction Temperature	☐ User Override	
Ambient Temp		20,0 °C
Effective ΘJA	☐ User Override	
Airflow		500 LFM
Heat Sink	High Profile	
ΘSA		0,4 °C/W
Board Selection	Medium (10"x10")	
# of Board Layers	16 or more	
ΘJB		
Board Temperature		

Implementation	
Optimization	Default

Summary

Total On-Chip Power	77,7 W	
Junction Temperature	43,3 °C	
Thermal Margin	56,7°C	173,6W
Effective ΘJA	0,3 °C/W	

39%	Transceiver......	30,666W
3%	I/O................	2,530W
50%	Core Dynamic..	38,696W
7%	Device Static.....	5,822W
Power supplied to off-chip devices...		-0,118W

On-Chip Power

Resource (Jump to sheet)		Power (W)	(%)
Core Dynamic	CLOCK	1,248	2
	LOGIC	29,661	38
	BRAM	0,760	1
	DSP	6,290	8
	PLL	0,078	0
	MMCM	0,355	0
	Other	0,305	0
	Hard IP	0,000	0
	URAM	0,000	0
I/O	IO	2,530	3
Transceiver	GTY	30,666	39
Device Static		5,822	7

Power Supply

Source	Voltage	Total (A)
V_{CCINT}	0,850	54,120
V_{CCINT_IO}	0,850	2,945
V_{CCBRAM}	0,850	0,345
V_{CCAUX}	1,800	1,630
V_{CCAUX_IO}	1,800	0,420
V_{CCO} 3.3V	3,300	
V_{CCO} 2.5V	2,500	
V_{CCO} 1.8V	1,800	0,772
V_{CCO} 1.5V	1,500	
V_{CCO} 1.35V	1,350	
V_{CCO} 1.2V	1,200	
V_{CCO} 1.0V	1,000	
-		
-		
-		
$MGTYV_{CCAUX}$	1,800	0,713
$MGTYAV_{CC}$	0,900	6,786
$MGTYAV_{TT}$	1,200	16,513
-		
V_{CCADC}	1,800	0,024
-		

Messages

FIGURE E.1: Overview of resulting currents drawn per input voltage and total power drawn.

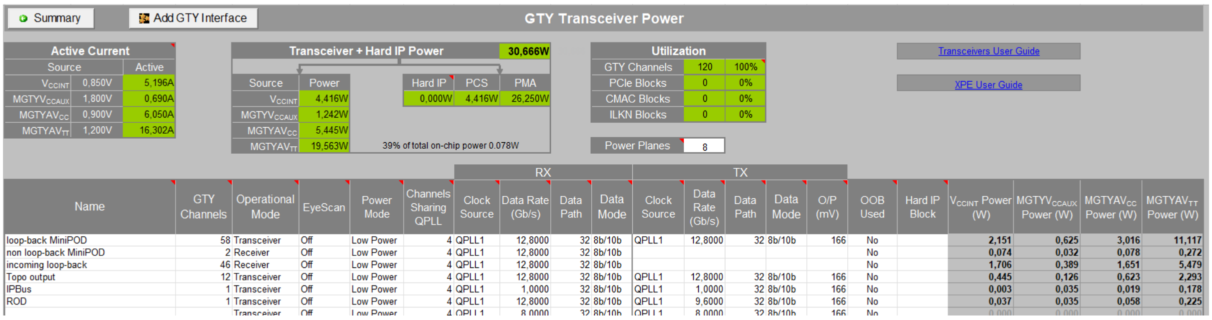

GTY Transceiver Power

Active Current		
Source		Active
V_{CCINT}	0,850V	5,196A
$MGTYV_{CCAUX}$	1,800V	0,690A
$MGTYAV_{CC}$	0,900V	6,050A
$MGTYAV_{TT}$	1,200V	16,302A

Transceiver + Hard IP Power 30,666W

Source	Power
V_{CCINT}	4,416W
$MGTYV_{CCAUX}$	1,242W
$MGTYAV_{CC}$	5,445W
$MGTYAV_{TT}$	19,563W

Hard IP	PCS	PMA
0,000W	4,416W	26,250W

39% of total on-chip power 0.078W

Utilization		
GTY Channels	120	100%
PCIe Blocks	0	0%
CMAC Blocks	0	0%
ILKN Blocks	0	0%

Power Planes: 8

						RX				TX										
Name	GTY Channels	Operational Mode	EyeScan	Power Mode	Channels Sharing QPLL	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	Clock Source	Data Rate (Gb/s)	Data Path	Data Mode	O/P (mV)	OOB Used	Hard IP Block	V_{CCINT} Power (W)	$MGTYV_{CCAUX}$ Power (W)	$MGTYAV_{CC}$ Power (W)	$MGTYAV_{TT}$ Power (W)
loop-back MiniPOD	58	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL1	12,8000	32	8b/10b	166	No		2,151	0,625	3,016	11,117
non loop-back MiniPOD	2	Receiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b						No		0,074	0,032	0,078	0,272
incoming loop-back	46	Receiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b						No		1,706	0,389	1,651	5,479
Topo output	12	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL1	12,8000	32	8b/10b	166	No		0,445	0,126	0,623	2,293
IPBus	1	Transceiver	Off	Low Power	4	QPLL1	1,0000	32	8b/10b	QPLL1	1,0000	32	8b/10b	166	No		0,003	0,035	0,019	0,178
ROD	1	Transceiver	Off	Low Power	4	QPLL1	12,8000	32	8b/10b	QPLL1	9,6000	32	8b/10b	166	No		0,037	0,035	0,058	0,225
		Transceiver	Off	Low Power	4	QPLL1	8,0000	32	8b/10b	QPLL1	8,0000	32	8b/10b	166	No		0,000	0,000	0,000	0,000

FIGURE E.2: Settings used to determine the power consumption of GTY high-speed links covering all high-speed links as no GTH transceivers are available on the XCVU9P.

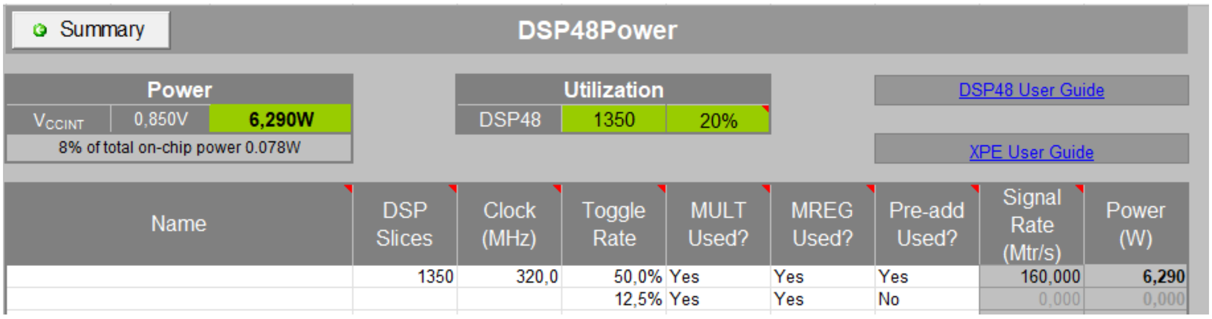

Power		
V_{CCINT}	0,850V	6,290W
8% of total on-chip power 0.078W		

Utilization		
DSP48	1350	20%

Name	DSP Slices	Clock (MHz)	Toggle Rate	MULT Used?	MREG Used?	Pre-add Used?	Signal Rate (Mtr/s)	Power (W)
	1350	320,0	50,0%	Yes	Yes	Yes	160,000	6,290
			12,5%	Yes	Yes	No	0,000	0,000

FIGURE E.3: Settings used to determine the power consumption of DSP slices with settings identical to the XCVU190 estimate.

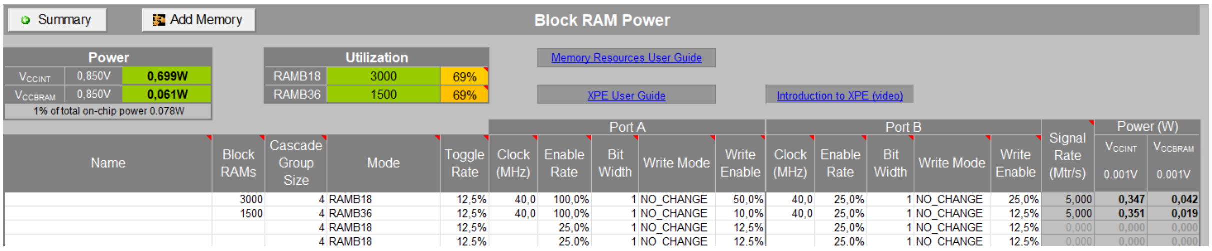

Power		
V_{CCINT}	0,850V	0,699W
V_{CCBRAM}	0,850V	0,061W
1% of total on-chip power 0.078W		

Utilization		
RAMB18	3000	69%
RAMB36	1500	69%

					Port A					Port B						Power (W)	
Name	Block RAMs	Cascade Group Size	Mode	Toggle Rate	Clock (MHz)	Enable Rate	Bit Width	Write Mode	Write Enable	Clock (MHz)	Enable Rate	Bit Width	Write Mode	Write Enable	Signal Rate (Mtr/s)	V_{CCINT} 0.001V	V_{CCBRAM} 0.001V
	3000	4	RAMB18	12,5%	40,0	100,0%	1	NO_CHANGE	50,0%	40,0	25,0%	1	NO_CHANGE	25,0%	5,000	0,347	0,042
	1500	4	RAMB36	12,5%	40,0	100,0%	1	NO_CHANGE	10,0%	40,0	25,0%	1	NO_CHANGE	12,5%	5,000	0,351	0,019
		4	RAMB18	12,5%		25,0%	1	NO_CHANGE	12,5%		25,0%	1	NO_CHANGE	12,5%	0,000	0,000	0,000
		4	RAMB18	12,5%		25,0%	1	NO_CHANGE	12,5%		25,0%	1	NO_CHANGE	12,5%	0,000	0,000	0,000

FIGURE E.4: Settings used to determine the power consumption of DSP slices assuming a low utilisation as the resource is not used by the algorithms and only required for read-out and control functionalities.

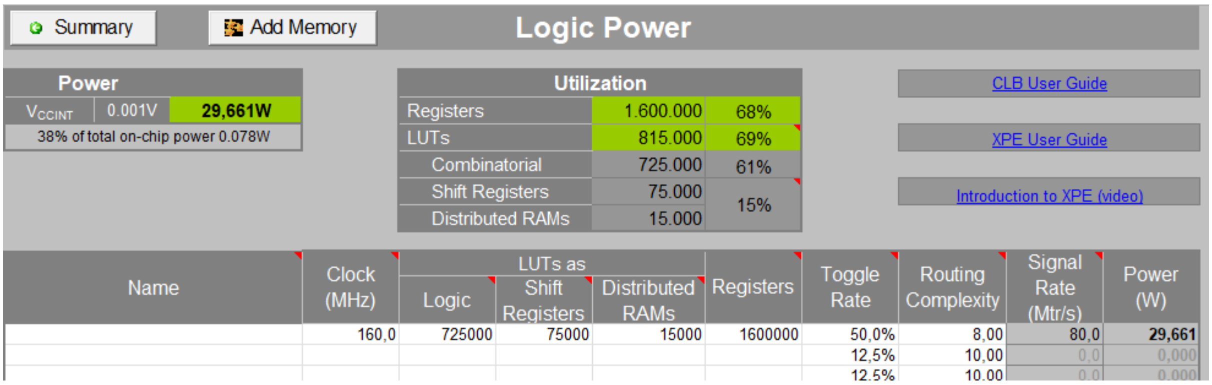

Power		
V_{CCINT}	0.001V	29,661W
38% of total on-chip power 0.078W		

Utilization		
Registers	1.600.000	68%
LUTs	815.000	69%
Combinatorial	725.000	61%
Shift Registers	75.000	15%
Distributed RAMs	15.000	

		LUTs as							
Name	Clock (MHz)	Logic	Shift Registers	Distributed RAMs	Registers	Toggle Rate	Routing Complexity	Signal Rate (Mtr/s)	Power (W)
	160,0	725000	75000	15000	1600000	50,0%	8,00	80,0	29,661
						12,5%	10,00	0,0	0,000
						12,5%	10,00	0,0	0,000

FIGURE E.5: Settings used to determine the power consumption of generic logic resources.

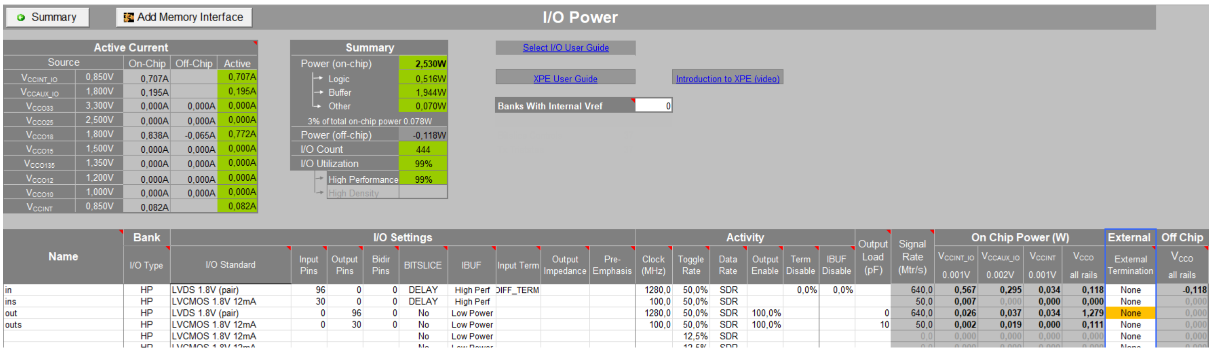

I/O Power

Active Current				
Source		On-Chip	Off-Chip	Active
V_{CCINT_IO}	0,850V	0,707A		0,707A
V_{CCAUX_IO}	1,800V	0,195A		0,195A
V_{CCO33}	3,300V	0,000A	0,000A	0,000A
V_{CCO25}	2,500V	0,000A	0,000A	0,000A
V_{CCO18}	1,800V	0,838A	-0,065A	0,772A
V_{CCO15}	1,500V	0,000A	0,000A	0,000A
V_{CCO135}	1,350V	0,000A	0,000A	0,000A
V_{CCO12}	1,200V	0,000A	0,000A	0,000A
V_{CCO10}	1,000V	0,000A	0,000A	0,000A
V_{CCINT}	0,850V	0,082A		0,082A

Summary	
Power (on-chip)	2,530W
→ Logic	0,516W
→ Buffer	1,944W
→ Other	0,070W
3% of total on-chip power 0.078W	
Power (off-chip)	-0,118W
I/O Count	444
I/O Utilization	99%
→ High Performance	99%
→ High Density	

Select I/O User Guide

XPE User Guide

Introduction to XPE (video)

Banks With Internal Vref: 0

	Bank	I/O Settings									Activity						Output Load (pF)	Signal Rate (Mtr/s)	On Chip Power (W)				External	Off Chip
Name	I/O Type	I/O Standard	Input Pins	Output Pins	Bidir Pins	BITSLICE	IBUF	Input Term	Output Impedance	Pre-Emphasis	Clock (MHz)	Toggle Rate	Data Rate	Output Enable	Term Disable	IBUF Disable			V_{CCINT_IO} 0.001V	V_{CCAUX_IO} 0.002V	V_{CCINT} 0.001V	V_{CCO} all rails	External Termination	V_{CCO} all rails
in	HP	LVDS 1.8V (pair)	96	0	0	DELAY	High Perf	DIFF_TERM			1280,0	50,0%	SDR		0,0%	0,0%		640,0	0,567	0,295	0,034	0,118	None	-0,118
ins	HP	LVCMOS 1.8V 12mA	30	0	0	DELAY	High Perf				100,0	50,0%	SDR					50,0	0,007	0,000	0,000	0,000	None	0,000
out	HP	LVDS 1.8V (pair)	0	96	0	No	Low Power				1280,0	50,0%	SDR	100,0%			0	640,0	0,026	0,037	0,034	1,279	None	0,000
outs	HP	LVCMOS 1.8V 12mA	0	30	0	No	Low Power				100,0	50,0%	SDR	100,0%			10	50,0	0,002	0,019	0,000	0,111	None	0,000
	HP	LVCMOS 1.8V 12mA				No	Low Power					12,5%	SDR					0,0	0,000	0,000	0,000	0,000	None	0,000

Figure E.6: Settings used to determine the power consumption of general purpose I/O links assuming a high data rate between the FPGAs.

Appendix F

High-speed test results

Complete results from the BER tests for the high-speed links on the first jFEX prototype. The different setups featuring the PMA loopback mode, the LATOME prototype, and the FTM are shown in section F.1, section F.2, and section F.3, respectively.

F.1 PMA loopback Tests

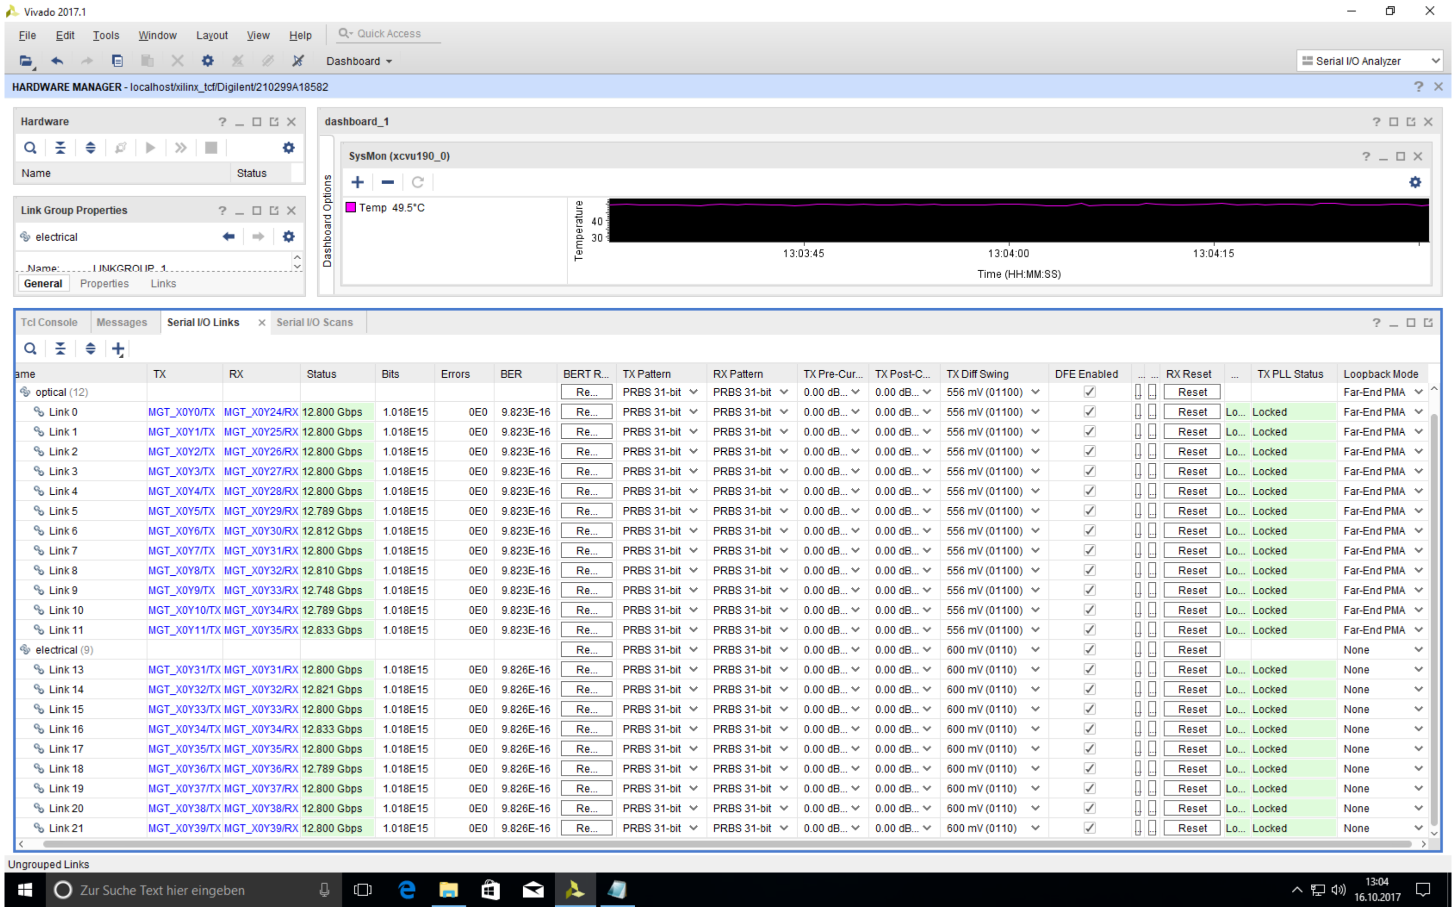

ame	TX	RX	Status	Bits	Errors	BER	BERT R...	TX Pattern	RX Pattern	TX Pre-Cur...	TX Post-C...	TX Diff Swing	DFE Enabled	...	...	RX Reset	...	TX PLL Status	Loopback Mode
optical (12)							Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset			Far-End PMA
Link 0	MGT_X0Y0/TX	MGT_X0Y24/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 1	MGT_X0Y1/TX	MGT_X0Y25/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 2	MGT_X0Y2/TX	MGT_X0Y26/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 3	MGT_X0Y3/TX	MGT_X0Y27/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 4	MGT_X0Y4/TX	MGT_X0Y28/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 5	MGT_X0Y5/TX	MGT_X0Y29/RX	12.789 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 6	MGT_X0Y6/TX	MGT_X0Y30/RX	12.812 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 7	MGT_X0Y7/TX	MGT_X0Y31/RX	12.800 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 8	MGT_X0Y8/TX	MGT_X0Y32/RX	12.810 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 9	MGT_X0Y9/TX	MGT_X0Y33/RX	12.748 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 10	MGT_X0Y10/TX	MGT_X0Y34/RX	12.789 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
Link 11	MGT_X0Y11/TX	MGT_X0Y35/RX	12.833 Gbps	1.018E15	0E0	9.823E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	556 mV (01100)	☑			Reset	Lo...	Locked	Far-End PMA
electrical (9)							Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset			None
Link 13	MGT_X0Y31/TX	MGT_X0Y31/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 14	MGT_X0Y32/TX	MGT_X0Y32/RX	12.821 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 15	MGT_X0Y33/TX	MGT_X0Y33/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 16	MGT_X0Y34/TX	MGT_X0Y34/RX	12.833 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 17	MGT_X0Y35/TX	MGT_X0Y35/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 18	MGT_X0Y36/TX	MGT_X0Y36/RX	12.789 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 19	MGT_X0Y37/TX	MGT_X0Y37/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 20	MGT_X0Y38/TX	MGT_X0Y38/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None
Link 21	MGT_X0Y39/TX	MGT_X0Y39/RX	12.800 Gbps	1.018E15	0E0	9.826E-16	Re...	PRBS 31-bit	PRBS 31-bit	0.00 dB...	0.00 dB...	600 mV (0110)	☑			Reset	Lo...	Locked	None

FIGURE F.1: First setup with the PMA loopback adapter and the receiving MiniPOD on AVR3 position providing a signal to MGT Quads 225 to 227, grouped as 'optical' links. The 'electrical' link group receives signals via the loopback adapter board. Links 13 to 17 receive the actual loopback signal based on the incoming MiniPOD signal while links 18 to 21 monitor the direct transmission created for a more realistic noise environment.

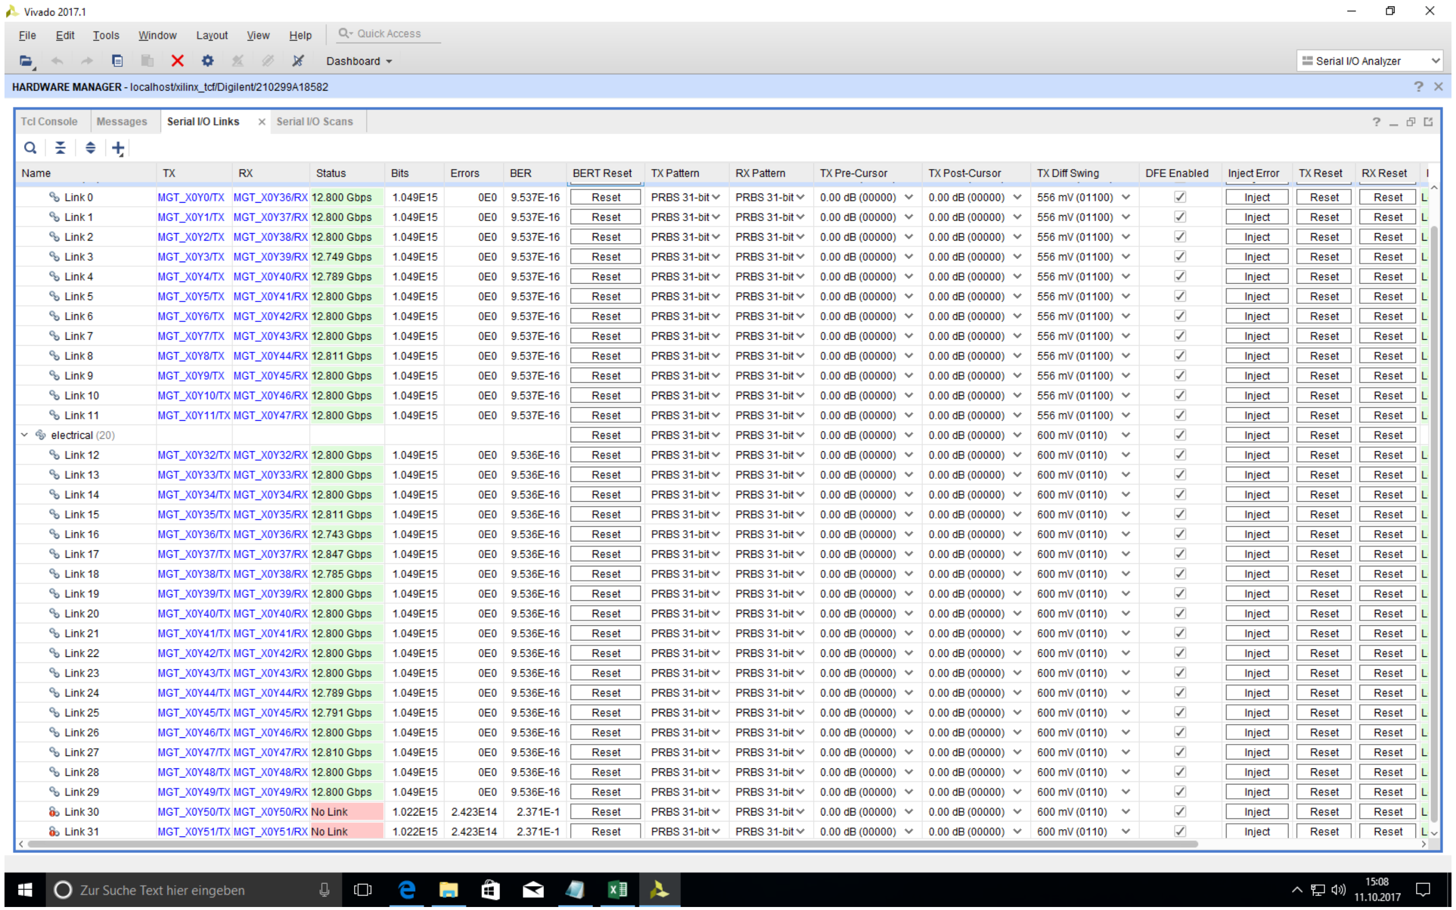

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern	TX Pre-Cursor	TX Post-Cursor	TX Diff Swing	DFE Enabled	Inject Error	TX Reset	RX Reset
Link 0	MGT_X0Y0/TX	MGT_X0Y36/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 1	MGT_X0Y1/TX	MGT_X0Y37/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 2	MGT_X0Y2/TX	MGT_X0Y38/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 3	MGT_X0Y3/TX	MGT_X0Y39/RX	12.749 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 4	MGT_X0Y4/TX	MGT_X0Y40/RX	12.789 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 5	MGT_X0Y5/TX	MGT_X0Y41/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 6	MGT_X0Y6/TX	MGT_X0Y42/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 7	MGT_X0Y7/TX	MGT_X0Y43/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 8	MGT_X0Y8/TX	MGT_X0Y44/RX	12.811 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 9	MGT_X0Y9/TX	MGT_X0Y45/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 10	MGT_X0Y10/TX	MGT_X0Y46/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
Link 11	MGT_X0Y11/TX	MGT_X0Y47/RX	12.800 Gbps	1.049E15	0E0	9.537E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	556 mV (01100)	☑	Inject	Reset	Reset
electrical (20)							Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 12	MGT_X0Y32/TX	MGT_X0Y32/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 13	MGT_X0Y33/TX	MGT_X0Y33/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 14	MGT_X0Y34/TX	MGT_X0Y34/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 15	MGT_X0Y35/TX	MGT_X0Y35/RX	12.811 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 16	MGT_X0Y36/TX	MGT_X0Y36/RX	12.743 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 17	MGT_X0Y37/TX	MGT_X0Y37/RX	12.847 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 18	MGT_X0Y38/TX	MGT_X0Y38/RX	12.785 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 19	MGT_X0Y39/TX	MGT_X0Y39/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 20	MGT_X0Y40/TX	MGT_X0Y40/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 21	MGT_X0Y41/TX	MGT_X0Y41/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 22	MGT_X0Y42/TX	MGT_X0Y42/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 23	MGT_X0Y43/TX	MGT_X0Y43/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 24	MGT_X0Y44/TX	MGT_X0Y44/RX	12.789 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 25	MGT_X0Y45/TX	MGT_X0Y45/RX	12.791 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 26	MGT_X0Y46/TX	MGT_X0Y46/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 27	MGT_X0Y47/TX	MGT_X0Y47/RX	12.810 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 28	MGT_X0Y48/TX	MGT_X0Y48/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 29	MGT_X0Y49/TX	MGT_X0Y49/RX	12.800 Gbps	1.049E15	0E0	9.536E-16	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 30	MGT_X0Y50/TX	MGT_X0Y50/RX	No Link	1.022E15	2.423E14	2.371E-1	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset
Link 31	MGT_X0Y51/TX	MGT_X0Y51/RX	No Link	1.022E15	2.423E14	2.371E-1	Reset	PRBS 31-bit	PRBS 31-bit	0.00 dB (00000)	0.00 dB (00000)	600 mV (0110)	☑	Inject	Reset	Reset

FIGURE F.2: Second setup with the PMA loopback adapter and the receiving MiniPOD on AVR4 position providing a signal to MGT Quads 228 to 230, grouped as 'optical' links. The 'electrical' link group receives signals via the loopback adapter board. Links 16 to 27 receive the actual loopback signal based on the incoming MiniPOD signal while links 12 to 15 and 28 to 29 monitor the direct transmission created for a more realistic noise environment. The connections for the links 30 and 31 are not covered by the loopback adapter.

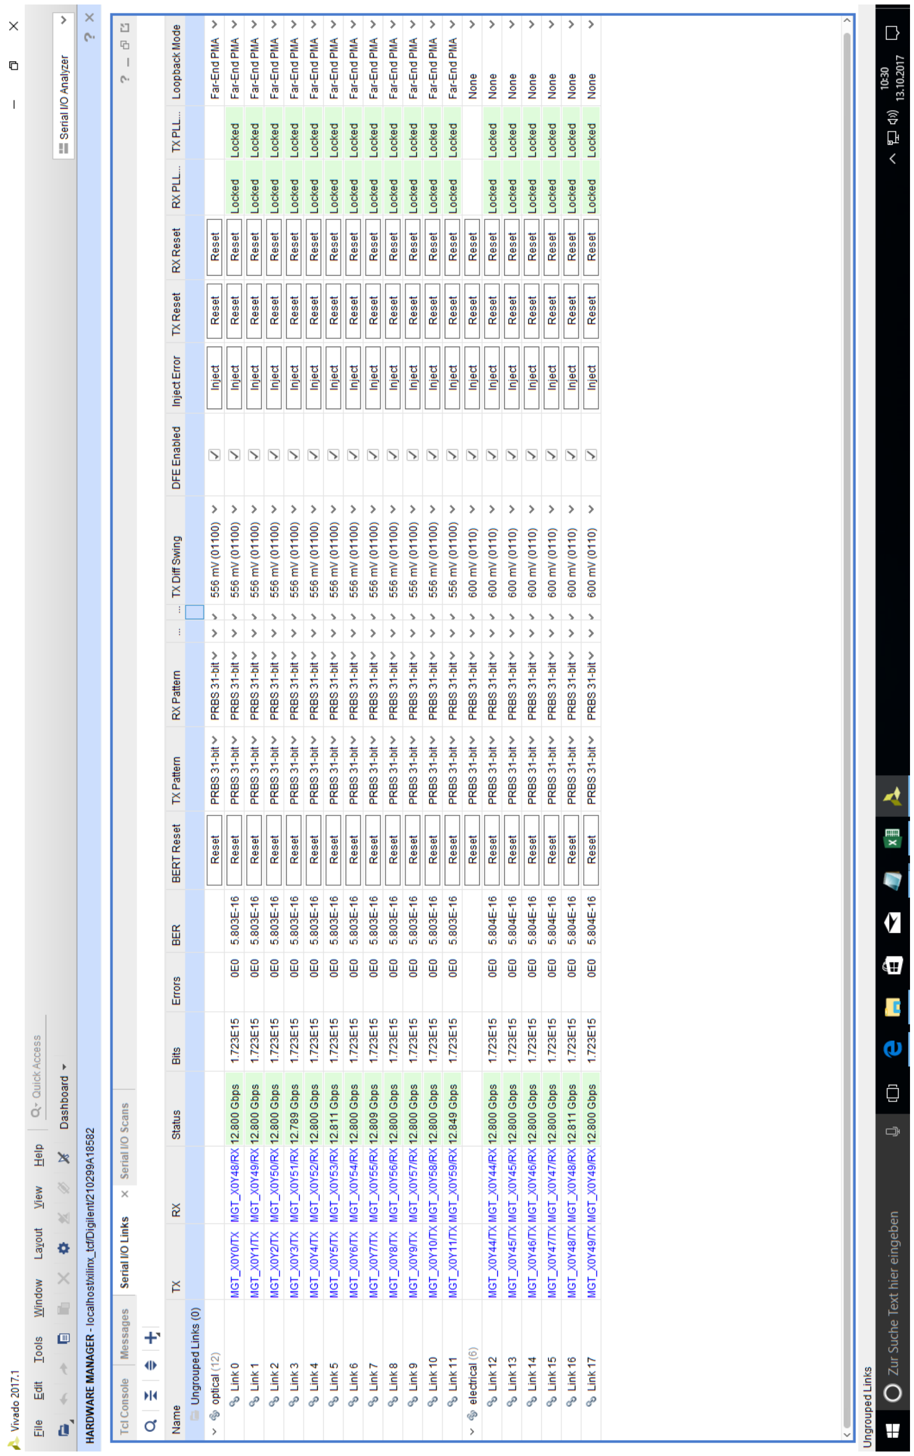

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern	...	TX Diff Swing	DFE Enabled	Inject Error	TX Reset	RX Reset	RX PLL...	TX PLL...	Loopback Mode
Ungrouped Links (0)																		
optical (12)							Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset			Far-End PMA
Link 0	MGT_X0Y0/TX	MGT_X0Y48/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 1	MGT_X0Y1/TX	MGT_X0Y49/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 2	MGT_X0Y2/TX	MGT_X0Y50/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 3	MGT_X0Y3/TX	MGT_X0Y51/RX	12.789 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 4	MGT_X0Y4/TX	MGT_X0Y52/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 5	MGT_X0Y5/TX	MGT_X0Y53/RX	12.811 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 6	MGT_X0Y6/TX	MGT_X0Y54/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 7	MGT_X0Y7/TX	MGT_X0Y55/RX	12.809 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 8	MGT_X0Y8/TX	MGT_X0Y56/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 9	MGT_X0Y9/TX	MGT_X0Y57/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 10	MGT_X0Y10/TX	MGT_X0Y58/RX	12.800 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
Link 11	MGT_X0Y11/TX	MGT_X0Y59/RX	12.849 Gbps	1.723E15	0E0	5.803E-16	Reset	PRBS 31-bit	PRBS 31-bit		556 mV (01100)	☑	Inject	Reset	Reset	Locked	Locked	Far-End PMA
electrical (6)							Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset			None
Link 12	MGT_X0Y44/TX	MGT_X0Y44/RX	12.800 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None
Link 13	MGT_X0Y45/TX	MGT_X0Y45/RX	12.800 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None
Link 14	MGT_X0Y46/TX	MGT_X0Y46/RX	12.800 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None
Link 15	MGT_X0Y47/TX	MGT_X0Y47/RX	12.800 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None
Link 16	MGT_X0Y48/TX	MGT_X0Y48/RX	12.811 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None
Link 17	MGT_X0Y49/TX	MGT_X0Y49/RX	12.800 Gbps	1.723E15	0E0	5.804E-16	Reset	PRBS 31-bit	PRBS 31-bit		600 mV (0110)	☑	Inject	Reset	Reset	Locked	Locked	None

FIGURE F.3: Third setup with the PMA loopback adapter and the receiving MiniPOD on AVR5 position providing a signal to MGT Quads 231 to 233, grouped as 'optical' links. The 'electrical' link group receives signals via the loopback adapter board. Links 16 and 17 receive the actual loopback signal based on the incoming MiniPOD signal while links 12 to 15 monitor the direct transmission created for a more realistic noise environment.

F.2 LATOME Tests

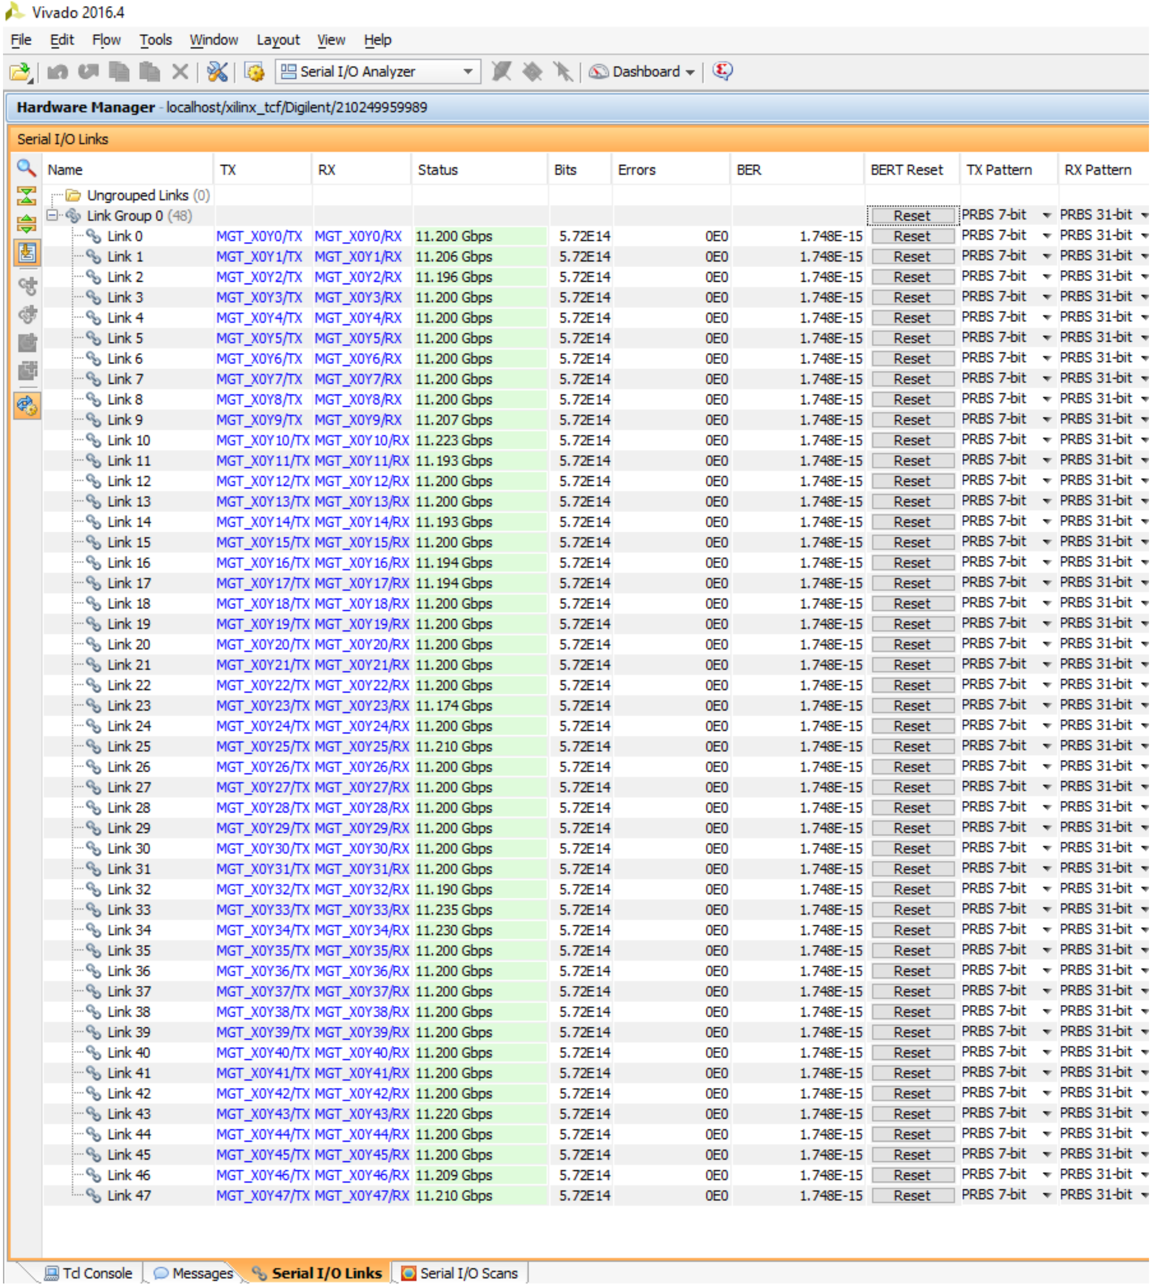

Hardware Manager - localhost/xilinx_tcf/Digilent/210249959989

Serial I/O Links

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Ungrouped Links (0)									
Link Group 0 (48)							Reset	PRBS 7-bit	PRBS 31-bit
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	11.206 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	11.196 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	11.207 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	11.223 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	11.193 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	11.193 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	11.194 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	11.194 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	11.174 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	11.210 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	11.190 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	11.235 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	11.230 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	11.220 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	11.200 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	11.209 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	11.210 Gbps	5.72E14	0E0	1.748E-15	Reset	PRBS 7-bit	PRBS 31-bit

FIGURE F.4: First bit error ratio test between the first jFEX prototype and the LATOME prototype at a transmission rate of 11.2 Gb/s targeting at a high number of transmitted bits. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

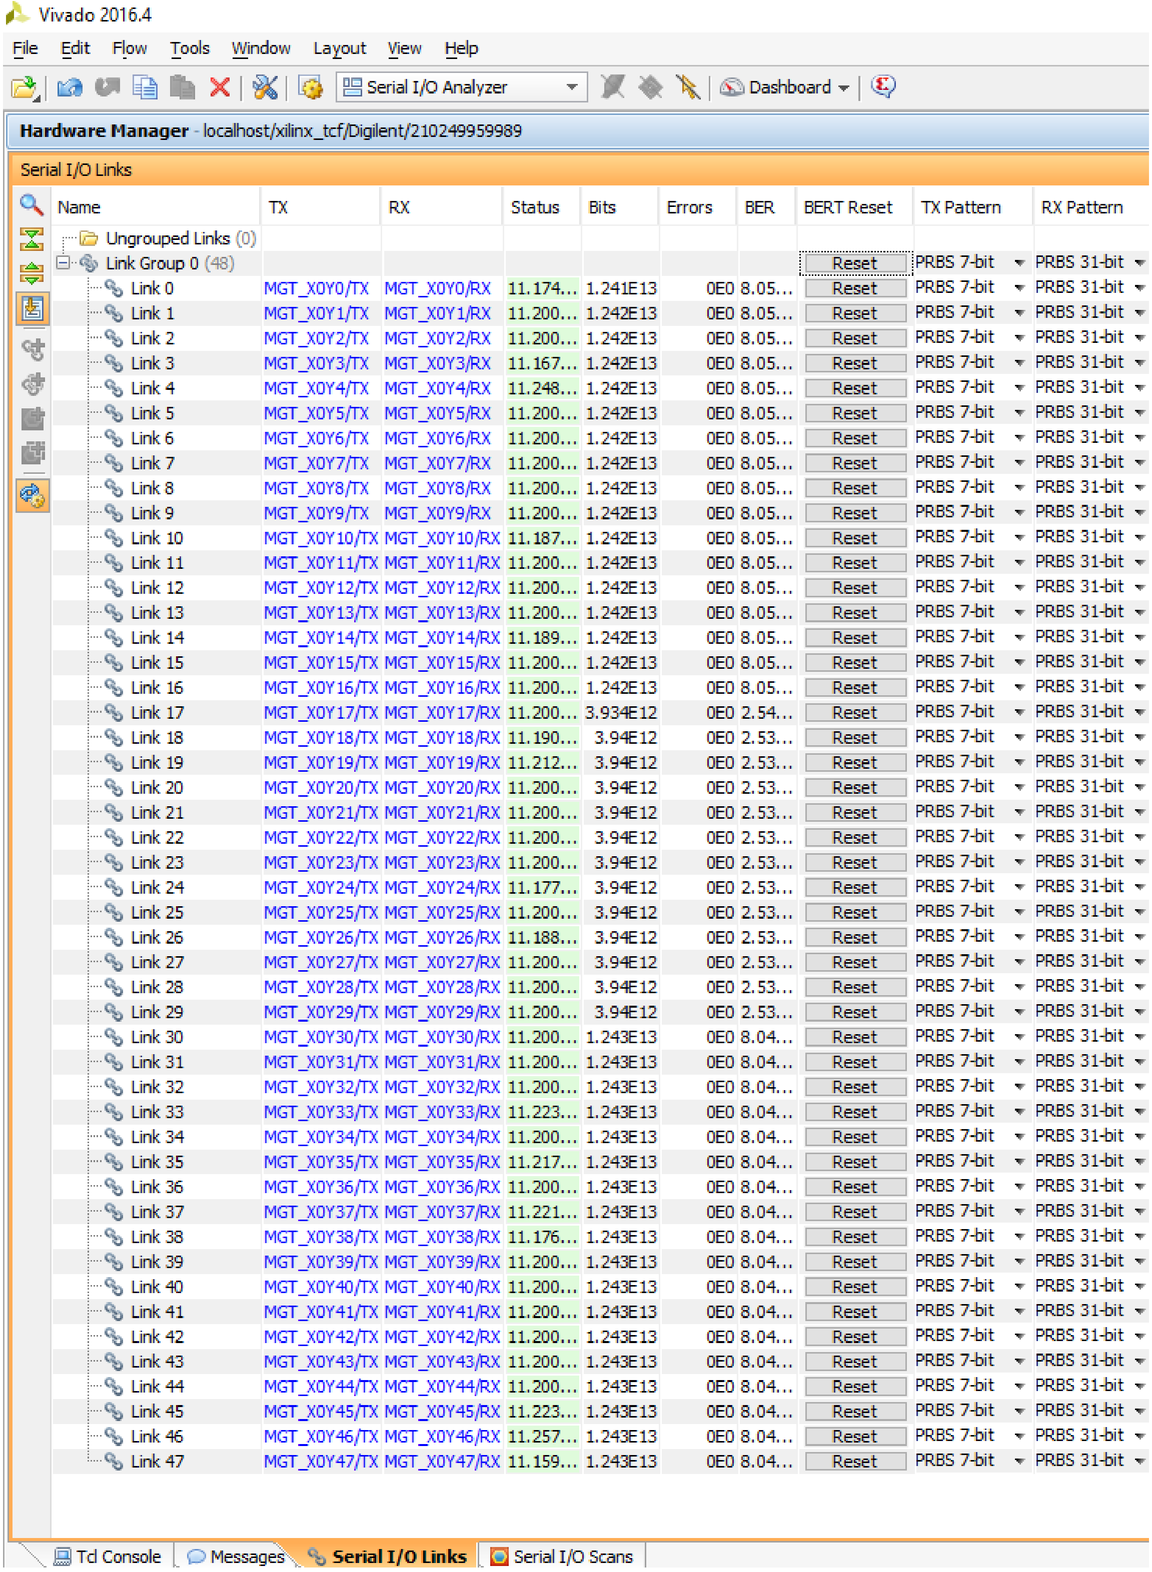

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Ungrouped Links (0)									
Link Group 0 (48)							Reset	PRBS 7-bit	PRBS 31-bit
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	11.174...	1.241E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	11.167...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	11.248...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	11.187...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	11.189...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	11.200...	1.242E13	0E0	8.05...	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	11.200...	3.934E12	0E0	2.54...	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	11.190...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	11.212...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	11.177...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	11.188...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	11.200...	3.94E12	0E0	2.53...	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	11.223...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	11.217...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	11.221...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	11.176...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	11.200...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	11.223...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	11.257...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	11.159...	1.243E13	0E0	8.04...	Reset	PRBS 7-bit	PRBS 31-bit

FIGURE F.5: Second bit error ratio test between the first jFEX prototype and the LATOME prototype at a transmission rate of 11.2 Gb/s running on a more realistic setup including FOX prototypes and a common clock source. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

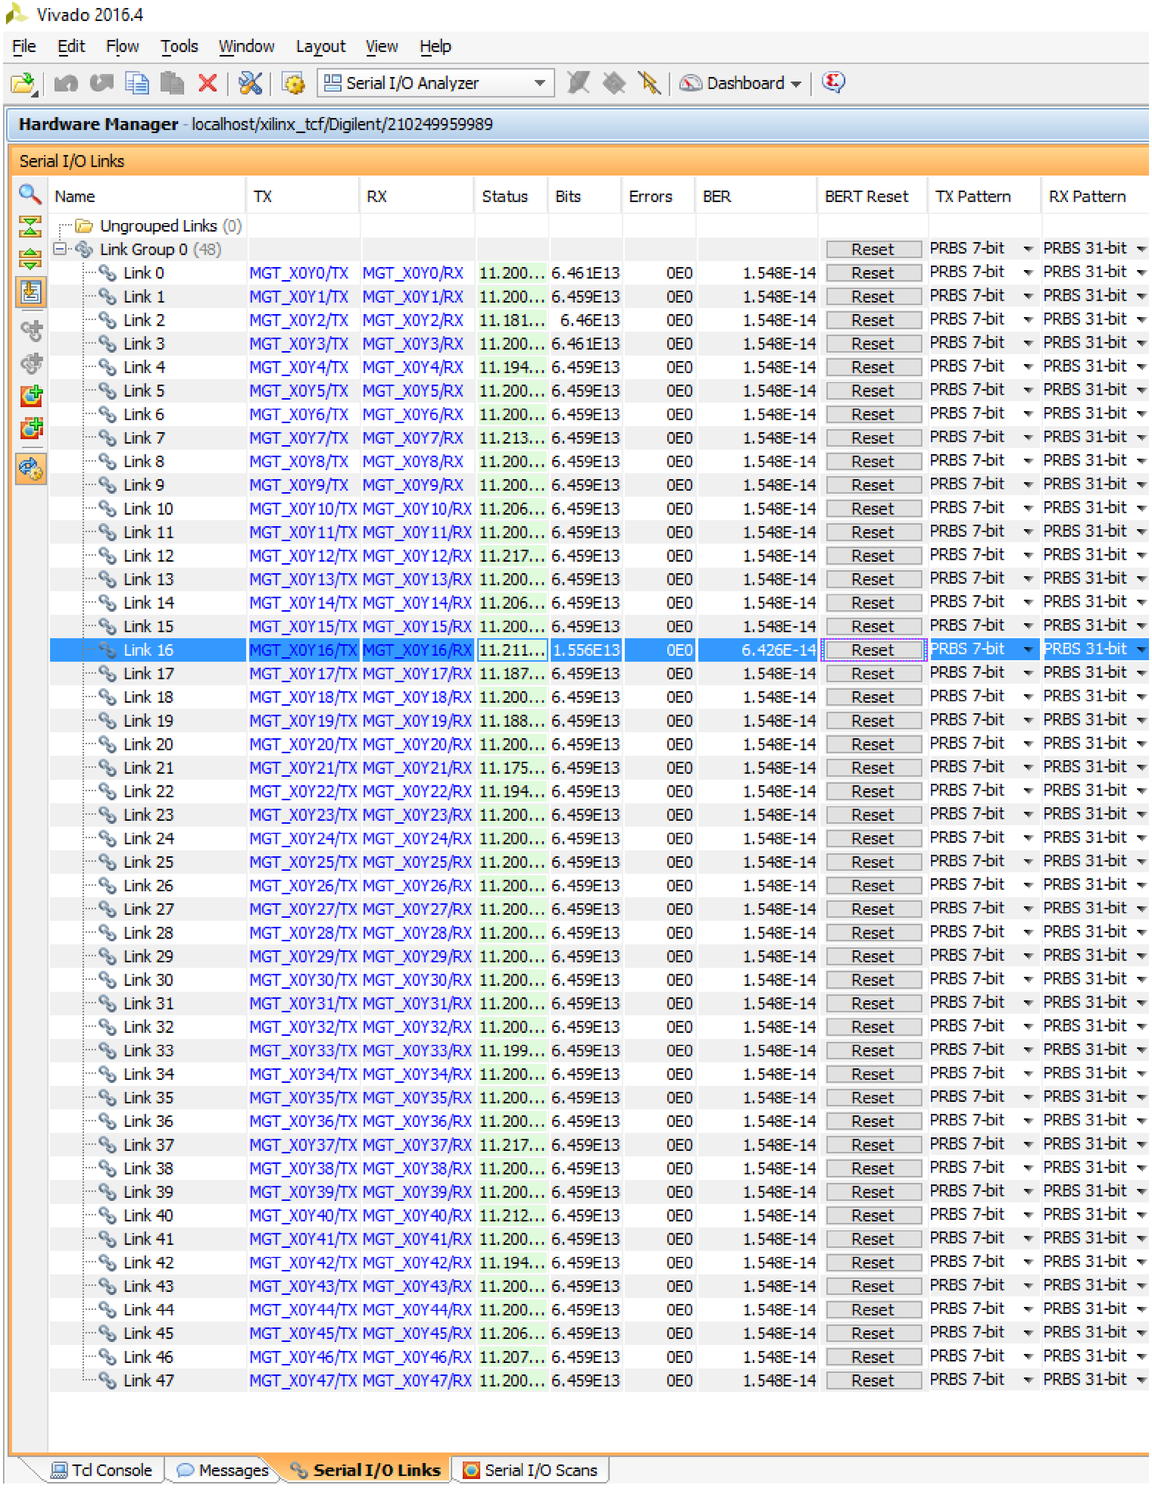

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Ungrouped Links (0)									
Link Group 0 (48)							Reset	PRBS 7-bit	PRBS 31-bit
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	11.200…	6.461E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	11.181…	6.46E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	11.200…	6.461E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	11.194…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	11.213…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	11.206…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	11.217…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	11.206…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	11.211…	1.556E13	0E0	6.426E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	11.187…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	11.188…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	11.175…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	11.194…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	11.199…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	11.217…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	11.212…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	11.194…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	11.206…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	11.207…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	11.200…	6.459E13	0E0	1.548E-14	Reset	PRBS 7-bit	PRBS 31-bit

FIGURE F.6: Bit error ratio test between the first jFEX prototype and the LATOME prototype at a transmission rate of 11.2 Gb/s with a 5 dB artificial attenuation of the optical signal on link 16. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

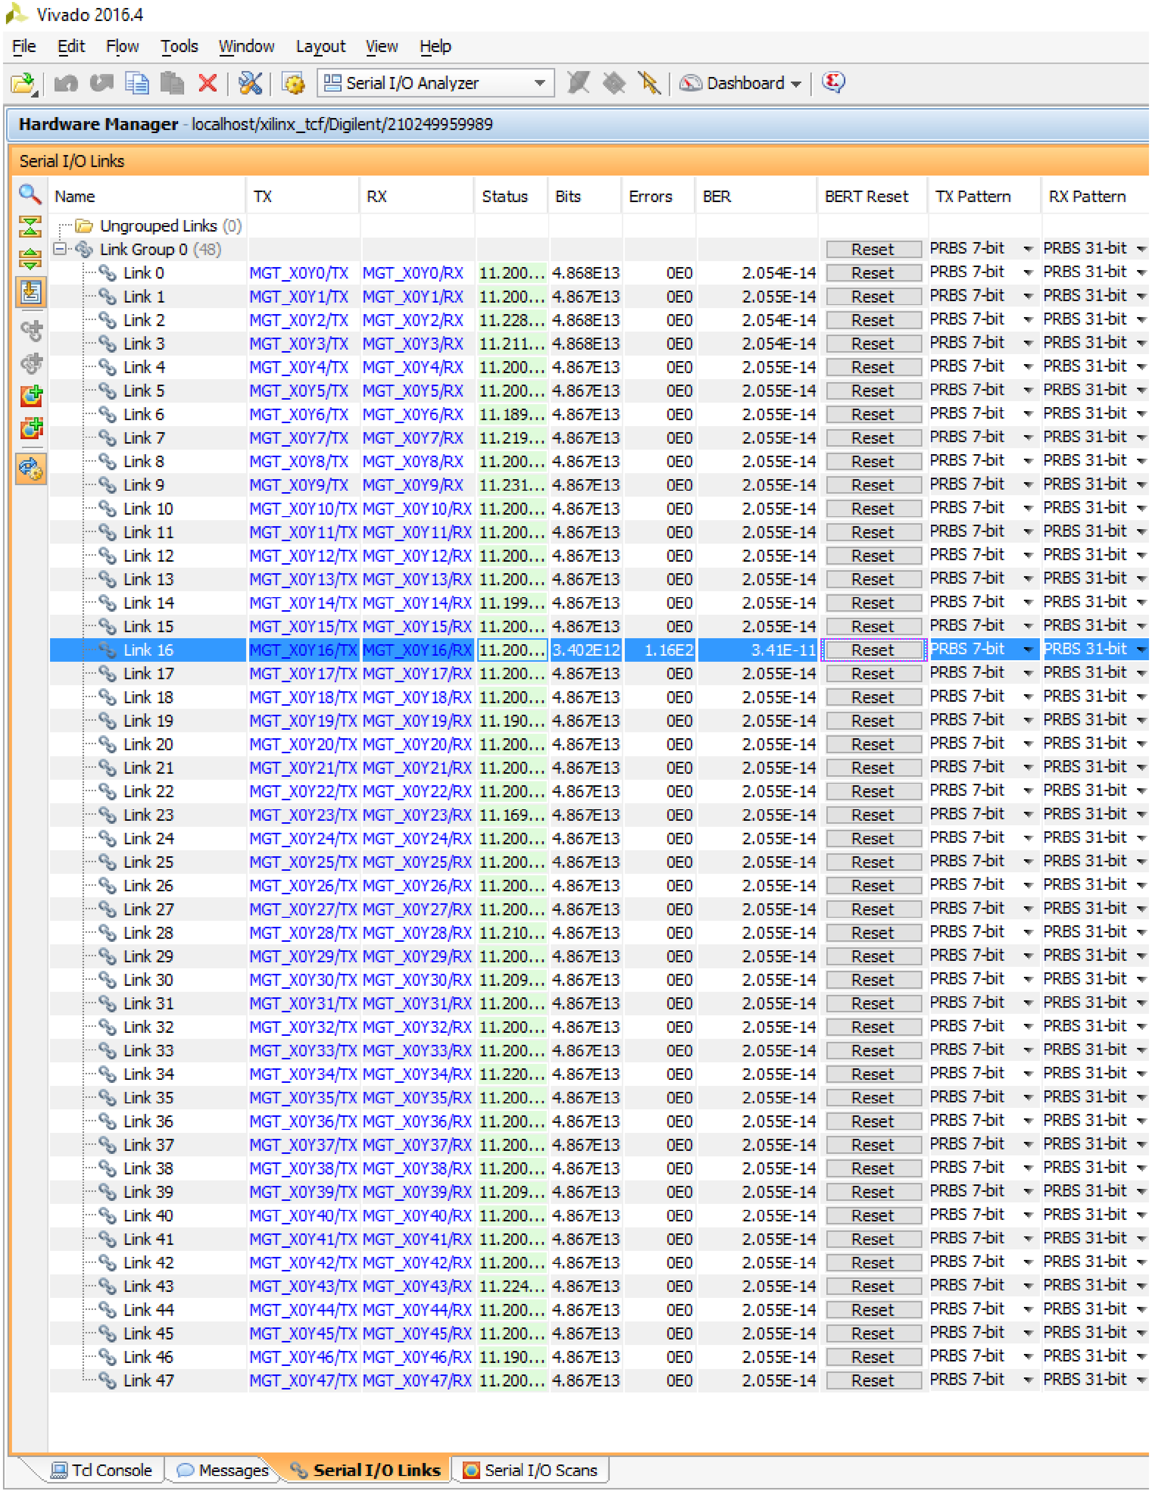

Hardware Manager - localhost/xilinx_tcf/Digilent/210249959989

Serial I/O Links

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Ungrouped Links (0)									
Link Group 0 (48)							Reset	PRBS 7-bit	PRBS 31-bit
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	11.200...	4.868E13	0E0	2.054E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	11.228...	4.868E13	0E0	2.054E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	11.211...	4.868E13	0E0	2.054E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	11.189...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	11.219...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	11.231...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	11.199...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	11.200...	3.402E12	1.16E2	3.41E-11	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	11.190...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	11.169...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	11.210...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	11.209...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	11.220...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	11.209...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	11.224...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	11.190...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	11.200...	4.867E13	0E0	2.055E-14	Reset	PRBS 7-bit	PRBS 31-bit

FIGURE F.7: Bit error ratio test between the first jFEX prototype and the LATOME prototype at a transmission rate of 11.2 Gb/s with a 10 dB artificial attenuation of the optical signal on link 16 showing a measurable bit error rate. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

F.3 FTM Tests

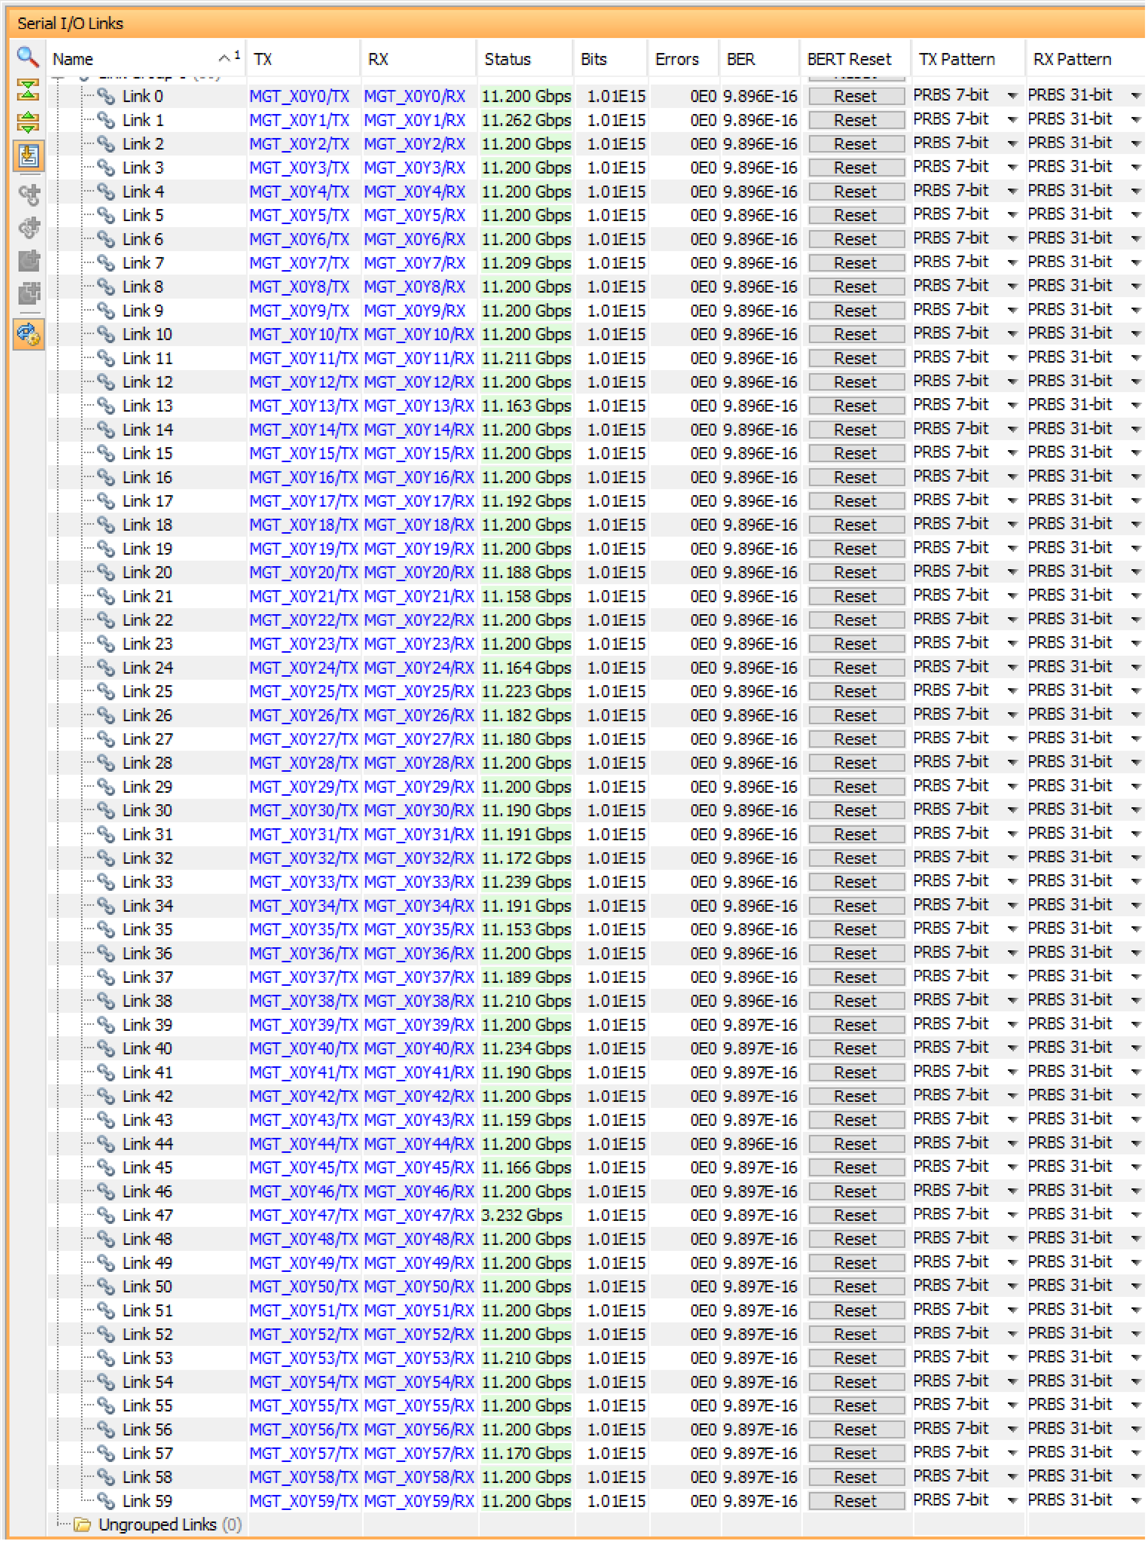

Serial I/O Links

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	11.262 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	11.209 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	11.211 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	11.163 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	11.192 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	11.188 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	11.158 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	11.164 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	11.223 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	11.182 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	11.180 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	11.190 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	11.191 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	11.172 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	11.239 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	11.191 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	11.153 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	11.189 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	11.210 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	11.234 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	11.190 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	11.159 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	11.200 Gbps	1.01E15	0E0	9.896E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	11.166 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	3.232 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 48	MGT_X0Y48/TX	MGT_X0Y48/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 49	MGT_X0Y49/TX	MGT_X0Y49/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 50	MGT_X0Y50/TX	MGT_X0Y50/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 51	MGT_X0Y51/TX	MGT_X0Y51/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 52	MGT_X0Y52/TX	MGT_X0Y52/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 53	MGT_X0Y53/TX	MGT_X0Y53/RX	11.210 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 54	MGT_X0Y54/TX	MGT_X0Y54/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 55	MGT_X0Y55/TX	MGT_X0Y55/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 56	MGT_X0Y56/TX	MGT_X0Y56/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 57	MGT_X0Y57/TX	MGT_X0Y57/RX	11.170 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 58	MGT_X0Y58/TX	MGT_X0Y58/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Link 59	MGT_X0Y59/TX	MGT_X0Y59/RX	11.200 Gbps	1.01E15	0E0	9.897E-16	Reset	PRBS 7-bit	PRBS 31-bit
Ungrouped Links (0)									

Figure F.8: Bit error ratio test between the first jFEX prototype and the FEX test module at a transmission rate of 11.2 Gb/s. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

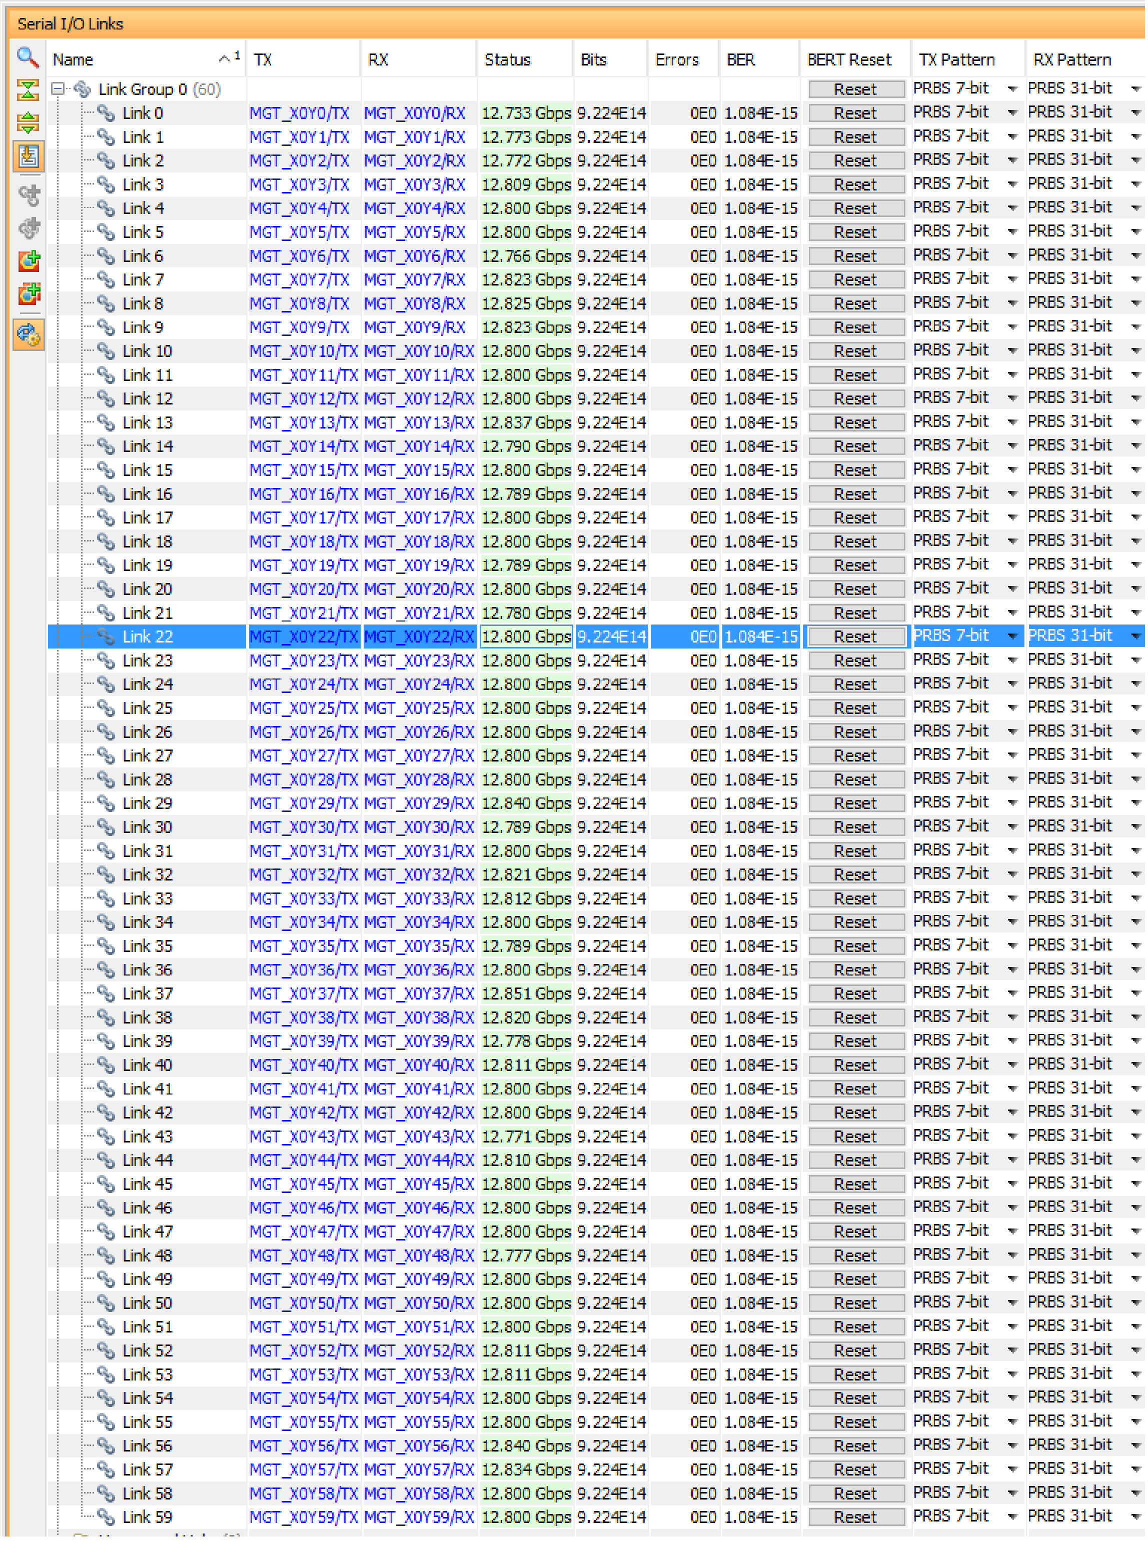

Name	TX	RX	Status	Bits	Errors	BER	BERT Reset	TX Pattern	RX Pattern
Link Group 0 (60)							Reset	PRBS 7-bit	PRBS 31-bit
Link 0	MGT_X0Y0/TX	MGT_X0Y0/RX	12.733 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 1	MGT_X0Y1/TX	MGT_X0Y1/RX	12.773 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 2	MGT_X0Y2/TX	MGT_X0Y2/RX	12.772 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 3	MGT_X0Y3/TX	MGT_X0Y3/RX	12.809 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 4	MGT_X0Y4/TX	MGT_X0Y4/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 5	MGT_X0Y5/TX	MGT_X0Y5/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 6	MGT_X0Y6/TX	MGT_X0Y6/RX	12.766 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 7	MGT_X0Y7/TX	MGT_X0Y7/RX	12.823 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 8	MGT_X0Y8/TX	MGT_X0Y8/RX	12.825 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 9	MGT_X0Y9/TX	MGT_X0Y9/RX	12.823 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 10	MGT_X0Y10/TX	MGT_X0Y10/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 11	MGT_X0Y11/TX	MGT_X0Y11/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 12	MGT_X0Y12/TX	MGT_X0Y12/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 13	MGT_X0Y13/TX	MGT_X0Y13/RX	12.837 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 14	MGT_X0Y14/TX	MGT_X0Y14/RX	12.790 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 15	MGT_X0Y15/TX	MGT_X0Y15/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 16	MGT_X0Y16/TX	MGT_X0Y16/RX	12.789 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 17	MGT_X0Y17/TX	MGT_X0Y17/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 18	MGT_X0Y18/TX	MGT_X0Y18/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 19	MGT_X0Y19/TX	MGT_X0Y19/RX	12.789 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 20	MGT_X0Y20/TX	MGT_X0Y20/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 21	MGT_X0Y21/TX	MGT_X0Y21/RX	12.780 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 22	MGT_X0Y22/TX	MGT_X0Y22/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 23	MGT_X0Y23/TX	MGT_X0Y23/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 24	MGT_X0Y24/TX	MGT_X0Y24/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 25	MGT_X0Y25/TX	MGT_X0Y25/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 26	MGT_X0Y26/TX	MGT_X0Y26/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 27	MGT_X0Y27/TX	MGT_X0Y27/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 28	MGT_X0Y28/TX	MGT_X0Y28/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 29	MGT_X0Y29/TX	MGT_X0Y29/RX	12.840 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 30	MGT_X0Y30/TX	MGT_X0Y30/RX	12.789 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 31	MGT_X0Y31/TX	MGT_X0Y31/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 32	MGT_X0Y32/TX	MGT_X0Y32/RX	12.821 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 33	MGT_X0Y33/TX	MGT_X0Y33/RX	12.812 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 34	MGT_X0Y34/TX	MGT_X0Y34/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 35	MGT_X0Y35/TX	MGT_X0Y35/RX	12.789 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 36	MGT_X0Y36/TX	MGT_X0Y36/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 37	MGT_X0Y37/TX	MGT_X0Y37/RX	12.851 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 38	MGT_X0Y38/TX	MGT_X0Y38/RX	12.820 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 39	MGT_X0Y39/TX	MGT_X0Y39/RX	12.778 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 40	MGT_X0Y40/TX	MGT_X0Y40/RX	12.811 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 41	MGT_X0Y41/TX	MGT_X0Y41/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 42	MGT_X0Y42/TX	MGT_X0Y42/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 43	MGT_X0Y43/TX	MGT_X0Y43/RX	12.771 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 44	MGT_X0Y44/TX	MGT_X0Y44/RX	12.810 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 45	MGT_X0Y45/TX	MGT_X0Y45/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 46	MGT_X0Y46/TX	MGT_X0Y46/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 47	MGT_X0Y47/TX	MGT_X0Y47/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 48	MGT_X0Y48/TX	MGT_X0Y48/RX	12.777 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 49	MGT_X0Y49/TX	MGT_X0Y49/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 50	MGT_X0Y50/TX	MGT_X0Y50/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 51	MGT_X0Y51/TX	MGT_X0Y51/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 52	MGT_X0Y52/TX	MGT_X0Y52/RX	12.811 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 53	MGT_X0Y53/TX	MGT_X0Y53/RX	12.811 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 54	MGT_X0Y54/TX	MGT_X0Y54/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 55	MGT_X0Y55/TX	MGT_X0Y55/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 56	MGT_X0Y56/TX	MGT_X0Y56/RX	12.840 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 57	MGT_X0Y57/TX	MGT_X0Y57/RX	12.834 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 58	MGT_X0Y58/TX	MGT_X0Y58/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit
Link 59	MGT_X0Y59/TX	MGT_X0Y59/RX	12.800 Gbps	9.224E14	0E0	1.084E-15	Reset	PRBS 7-bit	PRBS 31-bit

FIGURE F.9: Bit error ratio test between the first jFEX prototype and the FEX test module at a transmission rate of 12.8 Gb/s. As the output channels of the jFEX were not connected to another board in this test, the properties transmitters were not adjusted to match the testing conditions, leading to the displayed difference between RX and TX pattern.

Bibliography

[1] J. Bleck-Neuhaus.
Elementare Teilchen. Moderne Physik von den Atomen bis zum Standard-Modell.
Berlin, Heidelberg: Springer, 2010. ISBN: 978-3-540-85300-8.

[2] M. Thomson. *Modern Particle Physics.* Cambridge University Press, 2013.
DOI: `10.1017/CBO9781139525367`.

[3] A. Purcell. *Go on a particle quest at the first CERN webfest.*
URL: `http://cds.cern.ch/record/1473657` (visited on 07/03/2022).

[4] G. Aad et al. 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'.
In: *Physics Letters B* 716.1 (2012), pp. 1–29. ISSN: 0370-2693.
DOI: `https://doi.org/10.1016/j.physletb.2012.08.020`.
URL: `http://dx.doi.org/10.1016/j.physletb.2012.08.020`.

[5] S. Chatrchyan et al. 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Physics Letters B* 716 (2012), pp. 30–61.
DOI: `10.1016/j.physletb.2012.08.021`. arXiv: `1207.7235 [hep-ex]`.

[6] Y.-R. Liu et al. 'Pentaquark and Tetraquark States'.
In: *Progress in Particle and Nuclear Physics* 107 (2019), pp. 237–320.
DOI: `10.1016/j.ppnp.2019.04.003`.
URL: `https://doi.org/10.1016%2Fj.ppnp.2019.04.003`.

[7] K. Shaw et al. *Physics at LHC.* Presentation in ATLAS MasterClasses.
Kigali, Rwanda, 2016.
URL: `https://indico.cern.ch/event/528094/contributions/2172873/`.

[8] H. Kolanoski and N. Wermes. *Teilchendetektoren. Grundlagen und Anwendungen.*
Springer, 2016. ISBN: 978-3-662-45349-0, 978-3-662-45350-6.
DOI: `10.1007/978-3-662-45350-6`.

[9] E. Hansen et al. *Jet Reconstruction and Calibration With ATLAS.*
Presentation during the ICHEP conference on high energy physics. 2020.
URL: `https://indico.cern.ch/event/868940/contributions/3813697/attachments/2083345/3499532/ehansen_ichep.pdf`.

[10] M. Cacciari, G. P. Salam and G. Soyez. 'The Anti-k(t) jet clustering algorithm'.
In: *Journal of High Energy Physics* 04 (2008), p. 063.
DOI: `10.1088/1126-6708/2008/04/063`.

[11] D. Krohn, J. Thaler and L.-T. Wang. 'Jet trimming'.
In: *Journal of High Energy Physics* 2010.2 (Feb. 2010). ISSN: 1029-8479.
DOI: `10.1007/jhep02(2010)084`.
URL: `http://dx.doi.org/10.1007/JHEP02(2010)084`.

[12] L. Evans and P. Bryant. 'LHC Machine'.
In: *Journal of Instrumentation* 3 (2008), S08001.
DOI: `10.1088/1748-0221/3/08/S08001`.

[13] M. Schaumann et al. 'First Xenon-Xenon Collisions in the LHC'.
In: CERN-ACC-2018-126 (2018), MOPMF039. 4 p.
DOI: `10.18429/JACoW-IPAC2018-MOPMF039`.
URL: `https://cds.cern.ch/record/2648699`.

[14] C. Lefèvre. 'The CERN accelerator complex. Complexe des accélérateurs du CERN'. 2008. URL: `https://cds.cern.ch/record/1260465`.

[15] G. Aad et al. 'The ATLAS Experiment at the CERN Large Hadron Collider'.
In: *Journal of Instrumentation* 3 (2008), S08003.
DOI: `10.1088/1748-0221/3/08/S08003`.

[16] S. Chatrchyan et al. 'The CMS experiment at the CERN LHC'.
In: *Journal of Instrumentation* 3 (2008), S08004.
DOI: `10.1088/1748-0221/3/08/S08004`.

[17] A. A. Alves Jr. et al. 'The LHCb Detector at the LHC'.
In: *Journal of Instrumentation* 3 (2008), S08005.
DOI: `10.1088/1748-0221/3/08/S08005`.

[18] K. Aamodt et al. 'The ALICE experiment at the CERN LHC'.
In: *Journal of Instrumentation* 3 (2008), S08002.
DOI: `10.1088/1748-0221/3/08/S08002`.

[19] G. Anelli et al. 'The TOTEM experiment at the CERN Large Hadron Collider'.
In: *Journal of Instrumentation* 3 (2008), S08007.
DOI: `10.1088/1748-0221/3/08/S08007`.

[20] O. Adriani et al. 'The LHCf detector at the CERN Large Hadron Collider'.
In: *Journal of Instrumentation* 3 (2008), S08006.
DOI: `10.1088/1748-0221/3/08/S08006`.

[21] J. Pinfold et al. *Technical Design Report of the MoEDAL Experiment*. 2009.
URL: `https://cds.cern.ch/record/1181486`.

[22] W. Herr and B Muratori. *Concept of luminosity*. 2006.
DOI: `10.5170/CERN-2006-002.361`. URL: `https://cds.cern.ch/record/941318`.

[23] M. Hostettler et al. *Impact of the Crossing Angle on Luminosity Asymmetries at the LHC in 2016 Proton Physics Operation*. 2017.
DOI: `10.18429/JACoW-IPAC2017-TUPVA005`.
URL: `https://cds.cern.ch/record/2289120`.

[24] Jorg Wenninger and Michi Hostettler for the Operations group.
LHC Report: colliding at an angle. 2017.
URL: `https://home.cern/news/news/accelerators/lhc-report-colliding-angle`
(visited on 03/11/2019).

[25] G. Apollinari et al.
High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. 0.1.
CERN Yellow Reports: Monographs. Geneva: CERN, 2017.
DOI: `10.23731/CYRM-2017-004`. URL: `https://cds.cern.ch/record/2284929`.

[26] CERN. *Project Schedule.*
URL: `https://project-hl-lhc-industry.web.cern.ch/content/project-schedule`
(visited on 05/12/2019).

[27] S. Haywood et al. *ATLAS inner detector.* Technical design report. ATLAS.
Geneva: CERN, 1997. URL: `http://cds.cern.ch/record/331064`.

[28] *ATLAS central solenoid.* Technical design report. ATLAS. Geneva: CERN, 1997.
DOI: `10.17181/CERN.ZZVJ.2JYE`. URL: `http://cds.cern.ch/record/331067`.

[29] M Capeans et al. *ATLAS Insertable B-Layer Technical Design Report.* 2010.
URL: `https://cds.cern.ch/record/1291633`.

[30] *ATLAS magnet system.* Technical design report. ATLAS. Geneva: CERN, 1997.
DOI: `10.17181/CERN.905C.VDTM`. URL: `https://cds.cern.ch/record/338080`.

[31] *ATLAS muon spectrometer.* Technical design report. ATLAS. Geneva: CERN, 1997.
URL: `https://cds.cern.ch/record/331068`.

[32] J. P. Badiou et al. *ATLAS barrel toroid.* Technical design report. ATLAS.
Geneva: CERN, 1997. DOI: `10.17181/CERN.RF2A.CP5T`.
URL: `https://cds.cern.ch/record/331065`.

[33] *ATLAS end-cap toroids.* Technical design report. ATLAS. Geneva: CERN, 1997.
DOI: `10.17181/CERN.P03D.WQLV`. URL: `https://cds.cern.ch/record/331066`.

[34] *ATLAS level-1 trigger.* Technical design report. ATLAS. Geneva: CERN, 1998.
URL: `http://cds.cern.ch/record/381429`.

[35] P. Jenni et al.
ATLAS high-level trigger, data-acquisition and controls: Technical Design Report.
Technical design report. ATLAS. Geneva: CERN, 2003.
URL: `http://cds.cern.ch/record/616089`.

[36] A. Ruiz-Martinez and ATLAS Collaboration. *The Run-2 ATLAS Trigger System.*
ATL-DAQ-PROC-2016-003. Geneva, 2016. DOI: `10.1088/1742-6596/762/1/012003`.
URL: `http://cds.cern.ch/record/2133909`.

[37] R Achenbach et al. *The ATLAS Level-1 Calorimeter Trigger.* Geneva, 2008.
DOI: `10.1088/1748-0221/3/03/P03001`. URL: `http://cds.cern.ch/record/1080560`.

[38] R. Simoniello.
The ATLAS Level-1 Topological Processor: from design to routine usage in Run-2.
Geneva, 2019. DOI: `10.1109/NSSMIC.2018.8824280`.
URL: `https://cds.cern.ch/record/2649959`.

[39] C. W. Kalderon et al. *Jet Trigger Public Plots – October 2017.*
ATL-COM-DAQ-2017-144. Geneva, 2017.
URL: `https://cds.cern.ch/record/2288217`.

[40] K. Mochizuki.
'b-jet identification algorithms and performance in the ATLAS experiment'.
In: *Nuclear and Particle Physics Proceedings* 273-275 (2016), pp. 2536–1538.
DOI: `10.1016/j.nuclphysbps.2015.09.450`.
URL: `https://doi.org/10.1016/j.nuclphysbps.2015.09.450`.

[41] ATLAS. *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment.*
CERN-LHCC-2012-022. LHCC-I-023. Geneva, 2012.
URL: `https://cds.cern.ch/record/1502664`.

[42] *ATLAS Phase-II Upgrade Scoping Document.* Geneva, 2015.
DOI: `10.17181/CERN.7CRX.AJHP`. URL: `https://cds.cern.ch/record/2055248`.

[43] T Kawamoto et al. *New Small Wheel Technical Design Report.*
CERN-LHCC-2013-006. ATLAS-TDR-020. 2013.
URL: `https://cds.cern.ch/record/1552862`.

[44] *Technical Design Report for the ATLAS Inner Tracker Pixel Detector.* Geneva, 2017.
DOI: `10.17181/CERN.FOZZ.ZP3Q`. URL: `https://cds.cern.ch/record/2285585`.

[45] *Technical Design Report for the ATLAS Inner Tracker Strip Detector.* Geneva, 2017.
URL: `http://cds.cern.ch/record/2257755`.

[46] *ATLAS Liquid Argon Calorimeter Phase-II Upgrade: Technical Design Report.*
Geneva, 2017. DOI: `10.17181/CERN.6QIO.YGHO`.
URL: `https://cds.cern.ch/record/2285582`.

[47] *Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer.*
Geneva, 2017. URL: `https://cds.cern.ch/record/2285580`.

[48] G. Aad et al.
Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System. 2013.
URL: `http://cds.cern.ch/record/1602235`.

[49] *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System.*
Geneva, 2017. DOI: `10.17181/CERN.2LBB.4IAL`.
URL: `https://cds.cern.ch/record/2285584`.

[50] R. Wang. *New Level-1 jet feature extraction modules for ATLAS phase-I upgrade.*
Geneva, 2020. DOI: `10.22323/1.364.0201`.
URL: `http://cds.cern.ch/record/2692468`.

[51] Y.-S. Lai and B. A. Cole. *Jet reconstruction in hadronic collisions by Gaussian filtering.*
2008. arXiv: `0806.1499 [nucl-ex]`. URL: `https://arxiv.org/abs/0806.1499`.

[52] *AdvancedTCA Base Specification.* PICMG-3.0-R3.0, 2017.

[53] *AdvancedTCA PICMG 3.0 Short Form Specification.* PICMG-3.0, Jan. 2003.

[54] I. Corporation et al. *IPMI: Intelligent Platform Management Interface Specification, second generation, Document Revision 1.1.* Oct. 2013.
URL: `https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/ipmi-second-gen-interface-spec-v2-rev1-1.pdf`.

[55] I. Brawn. *private communication.* 2015.

[56] Xilinx. *UltraScale FPGA Product Tables and Product Selection Guide.* XMP102 (v1.7), 2016.

[57] Xilinx. *UltraScale+ FPGAs Product Tables and Product Selection Guide.* XMP103 (v2.0), 2021.

[58] 'IEEE Standard VHDL Language Reference Manual'. In: *IEEE Std 1076-2008 (Revision of IEEE Std 1076-2002)* (2009), pp. 1–640. DOI: `10.1109/IEEESTD.2009.4772740`.

[59] Xilinx. *UltraScale Architecture and Product Data Sheet: Overview.* DS890 (v4.1), Jan. 2022.

[60] N. Brown. *Reconfigurable architectures.* URL: `https://nickbrown.online/?page_id=25` (visited on 07/03/2022).

[61] Xilinx. *UltraScale Architecture Configurable Logic Block User Guide.* UG574 (v1.5), Feb. 2017.

[62] Xilinx. *Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics.* DS923 (v1.19), June 2021.

[63] Xilinx. *UltraScale Architecture SelectIO Resources User Guide.* UG571 (v1.13), Oct. 2021.

[64] Xilinx. *Large FPGA Methodology Guide.* UG872 (v14.3), Oct. 2012.

[65] Xilinx. *Xilinx Power Estimator.* v2016.3. Oct. 2016.

[66] Xilinx. *Xilinx Power Estimator.* v2017.2, May 2017.

[67] Xilinx. *UltraScale Architecture GTH Transceivers User Guide.* UG576 (v1.7.1), Aug. 2021.

[68] Xilinx. *UltraScale Architecture GTY Transceivers User Guide.* UG578 (v1.3.1), Sept. 2021.

[69] Panasonic. *Ultra-low transmission loss Highly heat resistant Multi-layer circuit board materials MEGTRON6.* URL: `https://industrial.panasonic.com/ww/products/pt/megtron/megtron6` (visited on 06/03/2022).

[70] Intel. *Intel Stratix 10 TX Product Table.* 652580 (v1.9), Sept. 2019.

[71] Xilinx. *VCU110 Evaluation Board User Guide.* UG1073 (v1.6), Jan. 2023.

[72] Xilinx. *VCU110 Evaluation Board User Guide.* UG1224 (v1.4), Oct. 2018.

[73] Xilinx. *CoolRunner-II CPLD Family.* DS090 (v3.1), Sept. 2008.

[74] Xilinx. *UltraScale Architecture Configuration.* UG570 (v1.16), Jan. 2022.

[75] Xilinx. *Virtex UltraScale FPGA Data Sheet: DC and AC Switching Characteristics.* DS893 (v1.12), May 2019.

[76] Xilinx. *UltraScale Architecture PCB Design User Guide.* UG583 (v1.24), July 2022.

[77] S. Artz. *private communication.* 2014-2024.

[78] M. Weirich. 'Development of New ATLAS Trigger Algorithms in Search for New Physics at the LHC'. PhD thesis. Johannes Gutenberg-Universität Mainz, 2021. URL: `https://doi.org/10.25358/openscience-6407`.

[79] N. Nottbeck.
Implementierung und Optimierung des Anti-kT-Jet-Clustering-Algorithmus' auf FPGAs.
Johannes Gutenberg-Universität Mainz. 2016.

[80] R. Brun and F. Rademakers. 'ROOT: An object oriented data analysis framework'.
In: *Nucl.Instrum.Meth.* A389 (1997), pp. 81–86.
DOI: `10.1016/S0168-9002(97)00048-X`.

[81] G. A. et al. *Internal workshop.* 2022.

[82] U. Schaefer. *private communication.* 2014-2024.

[83] M. Weirich. *private communication.* 2014-2024.

[84] G. A. et al. *TDAQ Dashboard.* 2023.
URL: `https://atlasop.cern.ch/tdaq/pbeastDashboard/d/iItbciwWk/full-trp-dashboard` (visited on 03/05/2023).

[85] B. T. Carlson et al.
Performance studies of the ATLAS L1Calorimeter trigger upgrade for run 3.
Geneva, 2018. URL: `https://cds.cern.ch/record/2309479`.

[86] C. Kahra. *private communication.* 2014-2024.

[87] H. Fu and R. Mayder.
Signal Integrity Simulation and On-Chip Evaluation for Low-Cost FPGA Transceivers.
2017. URL: `https://www.xilinx.com/support/documentation/white_papers/wp428-7Series-Serial-Link-Signal-Analysis.pdf`.

[88] J. V. D. Souza. *private communication.* 2014-2023.

Danksagung

Personenbezogene Daten sind aus Datenschutzgründen aus der elektronischen Version gelöscht.

Lebenslauf

Personenbezogene Daten sind aus Datenschutzgründen aus der elektronischen Version gelöscht.