
An FPGA-based Trigger Processor for a Measurement of Deeply Virtual Compton Scattering at the COMPASS-II Experiment

Sebastian Schopferer



**Fakultät für Mathematik und Physik
Albert-Ludwigs-Universität Freiburg**

An FPGA-based Trigger Processor for a Measurement of Deeply Virtual Compton Scattering at the COMPASS-II Experiment

Dissertation
zur
Erlangung des Doktorgrades
der
Fakultät für Mathematik und Physik
der
Albert-Ludwigs-Universität
Freiburg im Breisgau

vorgelegt von
Sebastian Schopferer
aus
Lahr/Schwarzwald

Freiburg, Oktober 2013

Dekan:	Prof. Dr. Michael Růžička
Leiter der Arbeit:	Prof. Dr. Horst Fischer
Referent:	Prof. Dr. Horst Fischer
Korreferent:	Prof. Dr. Karl Jakobs
Tag der Verkündung des Prüfungsergebnisses:	16.12.2013

Teile dieser Arbeit wurden in folgenden Fachzeitschriften veröffentlicht:

M. Alekseev, et al.: Physics Letters B, **693**(3):227 – 235 (2010). doi:10.1016/j.physletb.2010.08.034
C. Adolph, et al.: Phys. Rev. D, **87**(5):052018 (2013). doi:10.1103/PhysRevD.87.052018
S. Bartknecht, et al.: Nucl. Instr. and Meth. A, **623**(1):507 – 509 (2010). doi:10.1016/j.nima.2010.03.052
S. Bartknecht, et al.: IEEE Trans. Nucl. Sci., **58**(4):1456–1459 (2011). doi:10.1109/TNS.2011.2142195
J. Bieling, et al.: Nucl. Instr. and Meth. A, **672**(0):13 – 20 (2012). doi:10.1016/j.nima.2011.12.104
M. Büchele, et al.: Physics Procedia, **37**(0):1827 – 1834 (2012). doi:10.1016/j.phpro.2012.02.504
M. Büchele, et al.: JINST, **7**(03):C03008 (2012). doi:10.1088/1748-0221/7/03/C03008
T. Baumann, et al.: JINST, **8**(01):C01016 (2013). doi:10.1088/1748-0221/8/01/C01016
M. Alexeev, et al.: Nucl. Instr. and Meth. A, **695**:159 – 162 (2012). doi:10.1016/j.nima.2011.11.079
M. Alexeev, et al.: Nucl. Instr. and Meth. A, **732**:264 – 268 (2013). doi:10.1016/j.nima.2013.08.020
M. Alexeev, et al.: Nucl. Instr. and Meth. A, “*Development of a Thick GEM-based detector for single photon detection*”, to be published (2013).

Contents

1	Introduction	1
2	Theoretical Motivation	3
2.1	The Nucleon Spin	3
2.2	Deep Inelastic Scattering	3
2.3	Parton Distribution Functions	5
2.3.1	Unpolarized PDFs	6
2.3.2	Longitudinal Polarized PDFs	6
2.4	Generalized Parton Distributions	9
2.4.1	Kinematic Variables	10
2.4.2	Relation of the GPDs to Known Distributions	10
2.4.3	Impact Parameter Dependent Parton Distributions	12
2.4.4	GPDs and the Nucleon Spin Structure	13
2.5	Deeply Virtual Compton Scattering	14
2.6	Beam Charge & Spin Asymmetry at COMPASS-II	16
3	The COMPASS-II Experiment	19
3.1	The Beam	19
3.2	The Target Region	22
3.2.1	The Liquid Hydrogen Target	22
3.2.2	The Recoil Proton Detector	22
3.3	The Spectrometer	26
3.3.1	Tracking Detectors	26
3.3.2	Calorimeters	26
3.3.3	Muon Identification	27
3.3.4	RICH-1	29
3.4	The Trigger System	31
3.5	The Data Acquisition System	32
3.5.1	The GANDALF Framework	34
4	CAMERA Proton Trigger	37
4.1	Detector Principle	37
4.1.1	Time-of-Flight Measurement	37
4.1.2	Energy Loss Measurement	38
4.2	Trigger Concept	41

4.2.1	Design Objectives	41
4.2.2	Trigger Conditions	41
4.2.2.1	Geometric Coincidence	41
4.2.2.2	Time Calibration	42
4.2.2.3	Cuts on Energy Deposition	43
4.2.2.4	Correlation with the Beam	45
4.2.3	Further Requirements	46
4.3	Electronics Framework	46
4.3.1	VME64x/VXS Crate	47
4.3.2	Backplane Link	48
5	The TIGER Module	53
5.1	Mainboard	54
5.1.1	Virtex-6 FPGA	54
5.1.1.1	Configurable Logic	56
5.1.1.2	Clock Management	57
5.1.1.3	Block RAM	59
5.1.1.4	DSP Slices	60
5.1.1.5	Inputs and Outputs	60
5.1.1.6	Gigabit Transceivers	62
5.1.2	DDR3 Memory	63
5.1.3	Clock Distribution Network	64
5.1.3.1	PCI Express Reference Clock	64
5.1.3.2	Experiment Synchronous Clocks	65
5.1.3.3	Free-Running Clocks	66
5.1.4	CoolRunner-II CPLD	67
5.1.5	Configuration Scheme	68
5.1.5.1	Configuration at Power-Up	68
5.1.5.2	Online Configuration	69
5.1.5.3	JTAG Configuration and Flash Programming	70
5.1.6	Power Distribution System	71
5.1.6.1	Hot-Swap Capability	71
5.1.6.2	Secondary Voltage Rails	72
5.1.6.3	Power-Up/Power-Down Sequence	74
5.1.6.4	Power Distribution System Summary	75
5.1.7	PCB Layout	76
5.1.7.1	Schematic	76
5.1.7.2	Constraints	77
5.1.7.3	Layout	78
5.1.7.4	Simulation	79
5.2	CPU and GPU Extension Boards	82
5.2.1	COM Express Module Specification	82
5.2.2	Kontron COMe-mTT10	82

5.2.3	COM Interfaces	83
5.2.3.1	External Interfaces	83
5.2.3.2	On-board Interfaces	83
5.2.4	MXM GPU Module Specification	87
5.2.4.1	TIGER MXM Interface	87
5.2.5	AMD Radeon GPU Card	89
5.2.6	PCI Express Interconnection	89
5.2.6.1	PCI Express Specification	89
5.2.6.2	PCI Express Device Layers	91
5.2.6.3	PCI Express Transaction Protocol	92
5.2.6.4	TIGER PCI Express Link Structure	93
5.2.6.5	TIGER Switch Design	94
5.3	Interfaces	97
5.3.1	VXS Interface	97
5.3.1.1	High Speed Signals	97
5.3.1.2	Sideband Signals	98
5.3.2	TIGER-to-TIGER Links	100
5.3.3	SFP Transceiver Sockets	101
5.3.4	TCS Interface	101
5.3.5	General Purpose I/O	101
6	Firmware and Software	105
6.1	FPGA Firmware	105
6.1.1	TIGER Base Design	106
6.1.1.1	CPLD Interface	107
6.1.1.2	Clock Management, Reset Logic and VXS I/O Buffers	107
6.1.1.3	TCS Interface Module	109
6.1.1.4	PCIe Endpoint Block	111
6.1.1.5	ChipScope Cores	112
6.1.1.6	Optional Cores	113
6.1.1.7	Partial Reconfiguration Workflow	114
6.1.2	TIGER Trigger Processor Design	116
6.1.3	TIGER Readout Concentrator Design	119
6.2	Operating System and Device Drivers	120
6.2.1	Kontron Drivers	121
6.2.2	CPLD Driver	121
6.2.3	PCIe Driver	121
6.2.4	JTAG Driver and ChipScope Server	123
6.2.5	AMD Driver and OpenCL SDK	124
6.3	Monitoring and Control Tools	124
7	Commissioning	127
7.1	CAMERA Installation	127
7.1.1	Laser Calibration Runs	129

7.1.2	TIGER Latency Adjustment	132
7.1.3	The StartCounter	134
7.2	Physics Data	135
7.2.1	First Proton Signals	135
7.2.2	Comparison of two CAMERA Calibration Methods	136
7.2.3	Muon Runs	139
8	Summary	143
A	TIGER Board Layout	145
B	FPGA and CPLD Pin Plan	149
C	Connector Pin-Out	157
C.1	VXS Backplane Connectors	157
C.2	VHDCI Connector	161
D	CAMERA Photos	163
	Bibliography	167

*“A poet once said, ‘The whole universe is in a glass of wine.’
We will probably never know in what sense he meant that,
for poets do not write to be understood. But it is true that if
we look at a glass of wine closely enough we see the entire
universe.”*

– Richard Feynman

The Feynman Lectures on Physics, Volume I, 3–7

1. Introduction

When looking at matter closely enough, one can resolve its inner structure. The visible matter in our universe consists of two classes of particles – leptons and hadrons. According to the Standard Model of particle physics, leptons are elementary particles without substructure, whereas hadrons are composed of smaller constituents. The most familiar hadrons are protons and neutrons, which form the atomic nucleus and are therefore called nucleons.

The observed cross-section for deep inelastic lepton–nucleon scattering measured in several experiments in the 1960s led to the development of the *parton model* [1, 2]. In this model nucleons are composed of point-like, quasi-free particles, the so-called partons, which were later identified with the quarks and gluons. According to current knowledge a nucleon is made up of three valence quarks surrounded by gluons mediating the strong force between them. At short time scales the gluons may fluctuate into virtual quark–antiquark pairs, referred to as sea quarks. The composition of the *nucleon's momentum* from its constituents has been determined by several experiments, showing that – under certain conditions – the quarks & antiquarks and the gluons each carry approximately half of the momentum of the nucleon.

However, the *nucleon spin structure* is not clear today. The naive assumption, that the spin of the nucleon is mainly carried by the quarks & antiquarks, has been disproven in the late 1980s by an experiment of the European Muon Collaboration [3], which was confirmed by further experiments at CERN [4], SLAC [5] and DESY [6]. Parton helicity distributions of the nucleon, which were extracted from these data sets, yield a contribution of the quark & antiquark helicities $\Delta\Sigma$ of only about 30%. In addition, the contribution of the gluon helicity ΔG to the nucleon spin seems also to be small. The gluon helicity distribution $\Delta g(x)$ for a mean value of the gluon momentum fraction $\langle x \rangle \approx 0.2$ has been determined recently at the COMPASS¹ experiment at CERN [7], indicating that its first moment is small and in the order of $|\Delta G| \approx 0.25$. Since the quark & antiquark and gluon helicities alone do not suffice to yield the nucleon spin, this raises the question about the contributions of the orbital angular momenta of the quarks and gluons.

¹Common Muon and Proton Apparatus for Structure and Spectroscopy

During the last 20 years the parton model has been advanced in order to clarify the spin structure of the nucleon. The concept of generalized parton distributions (GPDs) has been introduced, which provides for the first time access to the total angular momentum of the partons. Hopefully, this will finally help to solve the last part of the long-standing *spin puzzle*. The GPDs can be extracted from cross section and asymmetry measurements of exclusive scattering processes like deeply virtual Compton scattering (DVCS) or hard exclusive meson production. A theoretical introduction to GPDs and DVCS is given in Chapter 2 of the thesis at hand.

Started in 2012, the COMPASS-II experiment is focusing on a measurement of the DVCS process $\mu p \rightarrow \mu p \gamma$. Several upgrades of the experimental setup have been performed, namely the construction of a long liquid hydrogen target and a surrounding recoil proton detector called CAMERA² as well as an additional large angle electromagnetic calorimeter. Details are given in Chapter 3. The measurement of DVCS is an exclusive measurement and the reconstruction of all particles in the final state is necessary. Therefore, the CAMERA detector is built for the detection of recoil particles, which leave the target under large polar angles and with very low momentum. Based on a time-of-flight measurement between two barrels of scintillators, it allows to detect protons with a kinetic energy down to 35 MeV. At the same time, protons can be distinguished from other particles resulting from background processes by means of an energy loss measurement in the scintillating material.

The signals of the CAMERA detector are read out by photomultipliers followed by fast digitizers, which perform an online pulse shape analysis. This serves two purposes: On the one hand the information is transferred to the data acquisition system (DAQ) upon receiving a first-level trigger (FLT), but on the other hand the CAMERA detector shall also be able to influence the FLT decision itself. This requires the extension of the CAMERA readout framework by a digital trigger system, whose development constitutes the main part of this thesis. The trigger system shall be able to select events with a recoil proton in the final state while suppressing background events, using the particle identification capabilities of the CAMERA detector. Challenging selection criteria based on both the time-of-flight and the energy loss measurement call for a powerful programmable logic board. At the same time, the integration into the existing COMPASS trigger system poses strict constraints on the latency of the trigger decision.

The concept for the new proton trigger system is introduced in Chapter 4. This paves the way for a new FPGA-based trigger and DAQ hardware called TIGER³, which is detailed in Chapter 5. The module is operated in two firmware configurations, serving two distinct purposes. Firstly, the *trigger processor* is responsible for the generation of a trigger signal based on recoil particles, which is included in the global FLT decision. Secondly, a *readout concentrator* allows to multiplex the data streams of up to 18 readout modules into one link to the DAQ. The firmware development is covered in Chapter 6. The CAMERA detector and the corresponding readout and trigger electronics was commissioned during a DVCS test run in autumn 2012. Information on the installation in the experiment and the calibration procedure as well as a first glimpse of the ongoing data analysis is provided in Chapter 7, before the results of this thesis are summarized in Chapter 8.

²COMPASS Apparatus for Measurement of Exclusive ReActions

³Trigger Implementation for GANDALF Electronic Readout

2. Theoretical Motivation

2.1 The Nucleon Spin

Spin is a basic property of any quantum-mechanical system. It is expressed in multiples of the reduced Planck constant $\hbar$. Protons and neutrons, together called nucleons, are spin-1/2 particles. In the *parton model* developed by Bjorken [1] and Feynman [2] the nucleons are composed of point-like particles called the "partons". They were later identified with the quarks, which had been introduced by Gell-Mann [8], and the gluons. While quarks are spin-1/2 particles, the gluons are the gauge bosons for the strong force, and therefore carry spin-1.

The spin of the nucleon can be written as [9, p.536]:

$$\frac{S}{\hbar} = \frac{1}{2} = \frac{1}{2} \Delta\Sigma + L_q + \Delta G + L_g, \quad (2.1)$$

where $\Delta\Sigma$ is the spin contribution from the quarks and anti-quarks, ΔG is the spin contribution from the gluons, L_q and L_g are the orbital angular momenta of the quarks and gluons respectively.

While $\Delta\Sigma \approx 0.3$ and $|\Delta G| \approx 0.2 \dots 0.3$ have been determined by former experiments [10, 11], the orbital angular momenta are not known today. The generalized parton distributions, which can be accessed by measuring deeply virtual Compton scattering or hard exclusive meson production, will soon provide information about the total angular momenta of the quarks and gluons. The theoretical background will be explained in the following sections.

2.2 Deep Inelastic Scattering

Deep inelastic scattering (DIS) can be used to investigate the nucleon structure. In DIS processes an incoming lepton l is scattered off a quark of the nucleon N by exchanging a virtual boson. For small center of mass energies $\sqrt{s}$, like in the COMPASS experiment, this boson is usually a virtual photon γ^* . The struck quark escapes from the nucleon, which consequently fragments into one or several hadrons (Fig. 2.1). If the scattered lepton is the only particle that




Figure 2.1: Schematic diagram of semi-inclusive deep inelastic scattering. A lepton l is scattered off a quark of the nucleon N . The momentum of the nucleon is P , the momentum of the lepton is k before and k' after the exchange of a virtual photon γ^* with momentum q . The helicities of the particles are drawn as wide arrows. [12, p.3]

is detected in the final state, the measurement is called *inclusive*. With at least one additional hadron detected in the final state the measurement is called *semi-inclusive*. For an *exclusive* measurement all outgoing particles have to be detected.

In the parton model DIS processes are best described in the *infinite momentum frame*, where the nucleon carries an 'infinite' momentum opposite to the one of the virtual photon [13, p.77]. In this reference frame the transverse momenta of the constituents are neglected, because they are small compared to the longitudinal components. Furthermore, time dilatation implies that there is no time for interaction between the nucleon constituents and the scattering process can therefore be treated in impulse approximation.

The following Lorentz invariant kinematic variables are used to describe the DIS:

$$q^2 = (k - k')^2 < 0 \quad \text{squared four-momentum of the virtual photon,} \quad (2.2)$$

$$\nu = \frac{P \cdot q}{M} \quad \text{energy loss of the scattered lepton,} \quad (2.3)$$

$$s = (k + P)^2 \quad \text{squared invariant center of mass energy,} \quad (2.4)$$

$$W^2 = (P + q)^2 \quad \text{squared invariant mass of the photon-nucleon system,} \quad (2.5)$$

where k and k' are the four-momenta of the incoming and scattered lepton, M is the mass of the nucleon and P is its 4-momentum. Because q^2 is negative, the virtuality of the photon is defined as $Q^2 = -q^2$.

In the lab frame of a fixed-target experiment, where the nucleon is at rest prior to the interaction, these variables are:

$$Q^2 \stackrel{\text{lab}}{=} 4EE' \sin^2 \frac{\theta}{2}, \quad (2.6)$$




Figure 2.2: Factorization of the deep inelastic scattering into a hard leptonic and a soft hadronic part. The soft part is parametrized by parton distribution functions. Indicated in brackets are the four-momenta of the respective particles.

$$\nu \stackrel{\text{lab}}{=} E - E', \quad (2.7)$$

$$s \stackrel{\text{lab}}{=} 2ME + M^2, \quad (2.8)$$

$$W^2 \stackrel{\text{lab}}{=} M^2 + 2M\nu - Q^2, \quad (2.9)$$

where E and E' are the energy of the incoming and of the scattered lepton and θ is the polar angle of the lepton scattering process.

In addition, one defines the dimension-less quantities

$$x_B = \frac{Q^2}{2P \cdot q} \stackrel{\text{lab}}{=} \frac{Q^2}{2M\nu} \quad \text{Bjorken scaling variable,} \quad (2.10)$$

$$y = \frac{P \cdot q}{P \cdot k} \stackrel{\text{lab}}{=} \frac{\nu}{E} \quad \text{fractional energy loss of the lepton.} \quad (2.11)$$

The distribution of cross sections in dependence of the Bjorken variable x_B is experimentally accessible. In the infinite momentum frame x_B can be interpreted as fraction x of the nucleon's four-momentum, which is carried by the struck quark before the virtual photon is absorbed [12, p.3]. The x_B distribution hence gives information about the momentum distribution of the quarks in the nucleon.

2.3 Parton Distribution Functions

In the Bjorken limit

$$Q^2, \nu \rightarrow \infty, \quad x_B = \text{const.} \quad (2.12)$$

the DIS process (Fig. 2.2) can be separated into a hard leptonic and a soft hadronic part [14] [15, p.183]:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{\alpha_{\text{em}}^2}{Q^4} \frac{E'}{E} L_{\mu\nu} W^{\mu\nu}. \quad (2.13)$$

The leptonic tensor $L_{\mu\nu}$, which describes the scattering of the virtual photon off a quark, can be calculated in QED, while the hadronic tensor $W^{\mu\nu}$ can only be parametrized by parton distribution functions (PDFs).

2.3.1 Unpolarized PDFs

The unpolarized parton distribution functions $q_f(x_B, Q^2)$ are defined as the number densities of partons of type f with momentum fraction x_B in the nucleon, when probed at a scale Q^2 . In the parton model, viewed in the infinite momentum frame, the nucleon is a collection of point-like, non-interacting partons, hence the DIS cross section $\sigma(x_B, Q^2)$ is just the sum of the cross sections for scattering from individual partons weighted by their number density [13, p.78]:

$$\sigma(x_B, Q^2) \propto \sum_f e_f^2 q_f(x_B, Q^2). \quad (2.14)$$

The sum runs over the quark and anti-quark flavors $f = u, d, s, \bar{u}, \bar{d}, \bar{s}, \dots$, and the electric charge e_f of the quarks is expressed in units of the elementary charge.

On the other hand, the DIS cross section can be expressed using the structure functions F_1 and F_2 , which can be measured in lepton scattering on unpolarized nucleons. A detailed calculation yields the relation between the unpolarized PDFs and the structure functions F_1 and F_2 [15, p.192]:

$$F_1(x_B, Q^2) = \frac{1}{2} \sum_f e_f^2 q_f(x_B, Q^2), \quad (2.15)$$

$$F_2(x_B, Q^2) = x_B \sum_f e_f^2 q_f(x_B, Q^2). \quad (2.16)$$

From Eqs. (2.15) and (2.16) the Callan-Gross relation

$$F_2(x_B, Q^2) = 2x_B F_1(x_B, Q^2) \quad (2.17)$$

is derived, whose experimental verification showed that the quarks carry spin-1/2 [16].

Early measurements of the structure functions showed a scaling behavior, i.e. no dependence of Q^2 , but actually a weak logarithmic variation is introduced due to radiative corrections in QCD. The Q^2 dependence of the proton structure function F_2 is shown in Fig. 2.3. The rise for small x_B and the drop for large x_B is explained by the proton's quark composition. The valence quarks of the nucleon are surrounded by a large number of virtual sea quarks with mainly very small momentum fractions. Because the resolution of a virtual photon increases with Q^2 , one detects more sea quarks at higher Q^2 . Consequently, more partons at small x_B and fewer partons at larger x_B can be found with increasing Q^2 [12]. The Q^2 dependence of the parton distributions is suppressed in the following notation, since it is not essential for the further considerations.

2.3.2 Longitudinal Polarized PDFs

Besides the momentum distribution q_f , also the helicity distribution Δq_f of the quarks can be measured with DIS. Longitudinal polarized leptons, i.e. leptons with helicity along or opposite to their momentum, are scattered off nucleons, which are polarized parallel ($\Rightarrow$) or anti-parallel ($\Leftarrow$) to the momentum of the incident leptons. The virtual photon, which inherits the helicity




Figure 2.3: Proton structure function $F_2(x_B, Q^2)$ measured in electromagnetic scattering of leptons on nucleons by various collider and fixed target experiments. For better visibility F_2 is multiplied by 2^{i_x} for the different x_B bins. i_x is the number of the x_B bin from 1 ($x_B = 0.85$) to 24 ($x_B = 5 \cdot 10^{-5}$). Combined statistical and systematic errors are shown by the vertical bars. [17, p.241]




Figure 2.4: Scattering of longitudinal polarized leptons off protons with same or opposite polarization. The virtual photon can only be absorbed by a quark with opposite helicity. [18]




Figure 2.5: Spin-dependent structure function $xg_1(x)$ of the proton, deuteron and neutron measured in DIS of polarized leptons by various experiments. Combined statistical and systematic errors are shown by the vertical bars. [17, p.247]

of the lepton to a certain extent depending on the lepton kinematics, can only be absorbed by a quark with helicity opposite to the photon helicity (Fig. 2.4). Thus quarks with a certain helicity can be selectively probed by choosing the nucleon's polarization.

The quark helicity distribution is defined in the parton model as the difference of the density distribution of quarks with helicity parallel ($q_f^{\vec{}}$) or anti-parallel ($q_f^{\leftarrow}$) to the nucleon's helicity:

$$\Delta q_f(x_B) = q_f^{\vec{}}(x_B) - q_f^{\leftarrow}(x_B). \quad (2.18)$$

The unpolarized or polarization-averaged parton distribution, which was introduced above, can be written as

$$q_f(x_B) = q_f^{\vec{}}(x_B) + q_f^{\leftarrow}(x_B). \quad (2.19)$$

From this follows the positivity limit $|\Delta q_f(x_B)| < q_f(x_B)$.

The spin-dependent structure function g_1 (Fig. 2.5) can be extracted from the difference of the cross sections with same and opposite polarizations of beam and target [13, p.99], measured in deep inelastic scattering:

$$\frac{d^2\sigma^{\leftarrow}}{d\Omega dE'} - \frac{d^2\sigma^{\vec{}}}{d\Omega dE'} = \frac{4\alpha_{\text{em}}^2}{M\nu Q^2} \frac{E'}{E} [(E + E' \cos\theta) g_1(x_B) - 2Mxg_2(x_B)]. \quad (2.20)$$

Because $E \gg M$, one is not sensitive to g_2 in an experiment with a longitudinal polarized target. It can only be determined using a transverse polarized target. The polarized PDFs contribute to $g_1(x_B)$, analogous to Eq. (2.15):

$$g_1(x_B) = \frac{1}{2} \sum_f e_f^2 \Delta q_f(x_B). \quad (2.21)$$

By integrating the polarized parton distributions over x_B , one obtains their first moments:

$$\Delta q_f = \int dx_B \Delta q_f(x_B). \quad (2.22)$$

By summing over all quark and anti-quark flavors one gets

$$\Delta \Sigma = \sum_f \Delta q_f, \quad (2.23)$$

the contribution of the quark helicities to the nucleon's spin in the sum rule (2.1).

2.4 Generalized Parton Distributions

The concept of generalized parton distributions (GPDs) was introduced by Müller [19], Radyushkin [20] and Ji [21]. The factorization in a hard and a soft sub-process of the DIS can in more general terms also be applied to other processes, where momentum is transferred to the target. The deeply virtual Compton scattering (DVCS) $\gamma^* N \rightarrow \gamma N$, which is shown in Fig. 2.6 and explained in detail in the next section, or the hard exclusive meson production (HEMP) are examples for such processes. An extensive overview of the theory for GPDs is given in [22].

In the Bjorken limit (2.12) and for a momentum transfer to the nucleon t , which is much smaller than Q^2 , the DVCS amplitude can be factorized into a convolution of the Compton amplitude for the lepton scattering, and a quark correlation function, which describes the soft part of the process [23, 24]. It is shown that the soft part of the DVCS amplitude is the same as the one of the elastic form factors. In fact, this is true for the soft part of many hard exclusive reactions. Hence the GPDs are introduced as a universal description to parametrize this soft part.

The four spin-1/2 GPDs, which will be used in the following sections, are listed in table 2.1. There are nucleon-helicity conserving ($H^{f,g}$, $\tilde{H}^{f,g}$) and nucleon-helicity inverting ($E^{f,g}$, $\tilde{E}^{f,g}$) GPDs, both polarized ($\sim$) and unpolarized. The polarized GPDs are quark helicity dependent, which is not the case for the unpolarized ones. The index f stands for the quark flavor u, d, s and the index g denotes the gluon GPDs.

Table 2.1: Spin-1/2 GPDs

	unpolarized	polarized
nucleon helicity conservation	$H^{f,g}$	$\tilde{H}^{f,g}$
nucleon helicity flip	$E^{f,g}$	$\tilde{E}^{f,g}$




Figure 2.6: Handbag diagram of the DVCS process $lN \rightarrow lN\gamma$. The bubble represents the inner structure of the nucleon, which here is described by GPDs. Indicated in brackets are the four-momenta of the respective particles. The kinematic variables t , Δ and ξ are explained in 2.4.1.

2.4.1 Kinematic Variables

To describe the GPDs, some additional kinematic variables are introduced.

The Mandelstam variable

$$t = (P - P')^2 = -\Delta^2 \quad (2.24)$$

specifies the transferred four-momentum between the initial and final state of the nucleon.

In the Bjorken limit the dimension-less variable ξ ('skewness')

$$\xi = \frac{x_B}{2 - x_B} \quad (2.25)$$

characterizes the direction of the momentum transfer t relative to the direction of momentum of the nucleon in the infinite momentum frame. For $\xi = 0$ they are orthogonal, for $\xi \neq 0$ the momentum transfer has a component parallel to the virtual photon. The momentum fraction of the interacting quark in the initial and in the final state is given by the sum $x + \xi$ and the difference $x - \xi$ respectively. Please note that x is the not measurable mean value of the quark's momentum fractions before and after the scattering process, and may not be identified with the Bjorken x_B . GPDs are defined for $x \in [-1, 1]$ with negative values describing the momentum fractions for anti-quarks. The dependence of the GPDs on x and ξ is mostly unknown today. Fig. 2.7 shows a model calculation for $H^u(x, \xi, t = 0)$.

2.4.2 Relation of the GPDs to Known Distributions

In the forward limit

$$t = 0 \text{ and } \xi = 0, \quad (2.26)$$




Figure 2.7: Model calculation for $H^u(x, \xi, 0)$ from [25]. The thick (red) line at $\xi = 0$ corresponds to the normal PDFs. The outer region $|x| > \xi$ is accessible via the DVCS process.

the four-momentum and helicity of the nucleon remain unchanged during the scattering process. Hence the helicity conserving GPDs are related to the corresponding PDFs [26]:

$$\text{for } x > 0: \quad H^f(x, 0, 0) = q_f(x), \quad \tilde{H}^f(x, 0, 0) = \Delta q_f(x), \quad (2.27)$$

$$H^g(x, 0, 0) = x g(x), \quad \tilde{H}^g(x, 0, 0) = x \Delta g(x), \quad (2.28)$$

$$\text{for } x < 0: \quad H^f(x, 0, 0) = -\bar{q}_f(-x), \quad \tilde{H}^f(x, 0, 0) = \Delta \bar{q}_f(-x). \quad (2.29)$$

The helicity inverting GPDs are not related to known PDFs, since the nucleon helicity flip requires a transfer of orbital angular momentum, which is possible only with finite transverse momentum transfer.

Furthermore, there is a relation between the first moments of the GPDs and the elastic form factors of the nucleon [21, p.613]:

$$\int_{-1}^1 dx H^f(x, \xi, t) = F_1^f(t), \quad (2.30)$$

$$\int_{-1}^1 dx E^f(x, \xi, t) = F_2^f(t), \quad (2.31)$$

$$\int_{-1}^1 dx \tilde{H}^f(x, \xi, t) = G_A^f(t), \quad (2.32)$$

$$\int_{-1}^1 dx \tilde{E}^f(x, \xi, t) = G_P^f(t). \quad (2.33)$$

The functions $F_1^f(t)$, $F_2^f(t)$, $G_A^f(t)$ and $G_P^f(t)$ are the contributions of the quark flavor f to the Dirac, Pauli, axial and pseudoscalar form factors.

For the GPDs there are higher order corrections in α_s , so they are – like the usual PDFs – depending on Q^2 . But $F_1^f(t)$ and $F_2^f(t)$ are independent from Q^2 , since they are defined as matrix




Figure 2.8: Impact parameter dependent PDFs for 'nucleon tomography'. (a) For a fixed x , $q_f(x, \mathbf{b}_\perp)$ describes the distribution of the transverse distance $b \equiv \mathbf{b}_\perp$ of partons carrying the fraction x of the nucleon's longitudinal momentum P , from the center of momentum of the nucleon $\mathbf{R}_\perp$. (b) Transverse spatial parton distribution in the nucleon at certain parton longitudinal momentum fractions x . [27, p.10]

elements of a conserved current. Hence the Q^2 dependency of the GPDs has to disappear in the first moments (2.30) and (2.31).

By the integration over x in (2.30) - (2.33), the reference to the infinite momentum frame in which ξ is defined is lost. Consequently, Lorentz invariance requires these integrals to be independent from ξ .

2.4.3 Impact Parameter Dependent Parton Distributions

In the limit of $\xi = 0$, where the parton carries the same longitudinal momentum fraction x in initial and final state, the GPDs can be interpreted in a phenomenological way [28]. In this case the momentum transfer is purely transverse, i.e. $t = -\Delta_\perp^2$. Here and in the following, boldface symbols denote 2-dimensional variables in the transverse plane. Like the Fourier transformations of the nucleon form factors show the distribution of charge in the nucleon, the GPDs provide information about the spatial distribution of partons with a momentum fraction x . The impact parameter dependent parton distributions are defined as the Fourier transform of the $-\Delta_\perp^2$ dependence of the GPD $H^f(x, 0, -\Delta_\perp^2)$:

$$q_f(x, \mathbf{b}_\perp) = \int \frac{d^2 \Delta_\perp}{(2\pi)^2} e^{-i\Delta_\perp \cdot \mathbf{b}_\perp} H^f(x, 0, -\Delta_\perp^2). \quad (2.34)$$

The result is a kind of 3-dimensional picture of the parton distributions in the nucleon. One dimension is the longitudinal fractional momentum x , which is parallel to the virtual photon. The two others give the position in the *transverse* plane – the impact parameter $\mathbf{b}_\perp$. A simultaneous measurement of both x and the *longitudinal* position is forbidden by the uncertainty principle.

The impact parameter dependent parton distributions $q_f(x, \mathbf{b}_\perp)$ for a set of different values of x are visualized in Fig. 2.8, which can be interpreted as momentum dissected images of the




Figure 2.9: Qualitative form of the impact parameter dependent PDFs $q_f(x, \mathbf{b}_\perp)$. [25]

nucleon. Depending on the momentum fraction x , either the valence quarks, the pion cloud or the sea quarks and gluons are resolved.

The impact parameter $\mathbf{b}_\perp$ is measured relative to the center of momentum of the nucleon $\mathbf{R}_\perp$, which is defined as the sum over the transverse positions $\mathbf{r}_{\perp,i}$ of all partons, weighted with their fractional momentum x_i :

$$\mathbf{R}_\perp = \sum_{i=q,g} x_i \mathbf{r}_{\perp,i}. \quad (2.35)$$

For $x \rightarrow 1$ the center of momentum is mainly defined by the active quark, and the width of the distribution $q_f(x, \mathbf{b}_\perp) = \sum_f q_f(x, \mathbf{b}_\perp)$ necessarily tends to zero (Fig. 2.9). On the other hand, partons with small x can be found also farther away from the nucleon's center of momentum.

2.4.4 GPDs and the Nucleon Spin Structure

One of the goals of the COMPASS experiment is the investigation of the nucleon spin structure. Up to now, the contribution of the orbital angular momenta L_f to the nucleon spin

$$\frac{1}{2} = \sum_f J_f + J_g, \quad (2.36)$$

$$J_f = \frac{1}{2}(\Delta q_f + \Delta \bar{q}_f) + L_f \quad (2.37)$$

is mainly unknown, since there was no process to measure the orbital angular momenta of the partons directly. In 1997, Xiangdong Ji [21] published his sum rule, which relates the second moments of the GPDs to the total angular momentum of the quarks:

$$J_f = \frac{1}{2} \lim_{t \rightarrow 0} \int_{-1}^1 dx x \left[H^f(x, \xi, t) + E^f(x, \xi, t) \right]. \quad (2.38)$$




Figure 2.10: Leading order processes for lepto production of real photons. (a) deeply virtual Compton scattering. (b) and (c) Bethe-Heitler process.

Using Ji's sum rule it is possible for the first time to determine the orbital angular momenta of the quarks. For this purpose it is necessary to precisely constrain the GPDs for $t \rightarrow 0$. This limit is not accessible in deep inelastic scattering, but it can be investigated in hard exclusive processes like DVCS or HEMP.

2.5 Deeply Virtual Compton Scattering

The GPDs H , $\tilde{H}$, E , $\tilde{E}$ of the u and d quarks can be experimentally investigated in hard exclusive photon or meson production. Observables like cross sections and asymmetries in dependency of beam or target polarization or beam charge contain information about the GPDs. Because the GPDs are typically convoluted with a hard scattering amplitude, they cannot be extracted from the data directly. Instead, the functional dependency of the GPDs on x , ξ and t is parametrized by theoretical models and the parameters are determined by comparison with the data.

One of the processes that will help to constrain the GPDs is the deeply virtual Compton scattering, where a real photon is produced by scattering a lepton probe off a nucleon target, which remains intact:

$$l N \rightarrow l N \gamma. \quad (2.39)$$

Fig. 2.10 (a) describes the DVCS process. A quark with longitudinal momentum fraction $x + \xi$ absorbs a virtual photon with virtuality Q^2 , and then emits a real photon which reduces the longitudinal momentum fraction to $x - \xi$. The quark GPD describes the correlation between the initial and final state quarks for given values of t and Q^2 .

The t dependence of the quark GPDs is directly accessible in DVCS, but this requires a high experimental precision to allow an extrapolation for $t \rightarrow 0$. This information is important for the determination of the total angular momentum of the quarks using Ji's sum rule (2.38).




Figure 2.11: Expected cross section of the reaction $\mu p \rightarrow \mu p \gamma$ as a function of the angle ϕ between the lepton scattering plane and the production plane. BH (blue), DVCS (red) and total cross section (dotted) are plotted for different x_B regions. Simulation was performed for $E_\mu = 160$ GeV, $Q^2 = 2$ GeV² and $t = -0.1$ GeV² (COMPASS kinematic). [30, p.7]

The Bethe-Heitler process (BH), which is shown in Fig. 2.10 (b) and (c), and the DVCS process have the same final state. Therefore their amplitudes interfere, and the differential cross section for BH and DVCS can be written as [29]:

$$\frac{d\sigma(lN \rightarrow lN\gamma)}{dx dQ^2 d|t| d\phi} \propto |\tau_{BH}|^2 + |\tau_{DVCS}|^2 + \underbrace{\tau_{DVCS}\tau_{BH}^* + \tau_{DVCS}^*\tau_{BH}}_I, \quad (2.40)$$

where ϕ is the azimuthal angle between the lepton scattering plane and the production plane, which is defined by the proton and the real photon.

In the DVCS process, the real photon originates from the nucleon and carries information about the GPDs, while the real photon in the BH process resulting from bremsstrahlung does not carry information about the nucleon and therefore can not be used for the determination of the GPDs. There are two possibilities to separate the contributions from BH and DVCS. On the one hand the DVCS cross section can be measured in a kinematic region, where DVCS dominates over BH (Fig. 2.11, right). In this case the cross section is basically given by the square of the DVCS amplitude which, at leading order, has the form [22, p.49]:

$$\tau_{DVCS} \sim \int_{-1}^1 \frac{H(x, \xi, t)}{x - \xi + i\epsilon} dx \dots \quad (2.41)$$

On the other hand one can measure the beam charge difference to separate BH and DVCS. When inverting the charge of the lepton beam, the DVCS amplitude τ_{DVCS} changes its sign, which is not the case for the BH amplitude τ_{BH} . Because τ_{BH} is real, the difference of cross sections is:

$$\sigma(l^+) - \sigma(l^-) \propto \tau_{BH} \Re(\tau_{DVCS}). \quad (2.42)$$

Here and in the following, σ is the short form for $\frac{d\sigma(lN \rightarrow lN\gamma)}{dx dQ^2 d|t| d\phi}$.

The BH amplitude is well known, because it only depends on the elastic nucleon form factors. The difference of cross sections is proportional to the real part of the DVCS amplitude, which for an unpolarized target is [31, p.29]:

$$\tau_{DVCS} \propto \left[F_1 \mathcal{H} + \xi(F_1 + F_2) \widetilde{\mathcal{H}} - \frac{t}{4m^2} F_2 \mathcal{E} \right]. \quad (2.43)$$

Here m is the nucleon mass, F_1 and F_2 are the elastic Dirac and Pauli form factors, $\mathcal{H}, \widetilde{\mathcal{H}}, \mathcal{E}$ and $\widetilde{\mathcal{E}}$ are the Compton form factors. These are convolutions of the respective quark GPDs with functions describing the hard Compton scattering $\gamma^* q$.

2.6 Beam Charge & Spin Asymmetry at COMPASS-II

The COMPASS-II experiment will measure the DVCS process to constrain the GPDs. The kinematic range for x is between ~ 0.01 and ~ 0.1 which has not yet been covered by other experiments. Hence COMPASS-II is bridging the gap between the collider experiments H1/ZEUS and the fixed target experiment HERMES [27, p.12]. Since the muon beam at the COMPASS-II experiment is naturally polarized due to the parity violation of the pion decay (see section 3.1), a change of sign of the beam charge (“+” $\Rightarrow$ “-”) will also invert the polarization (“ $\leftarrow$ ” $\Rightarrow$ “ $\rightarrow$ ”).

The Compton form factor $\mathcal{H}$ is related to the difference $\mathcal{D}$ and the sum $\mathcal{S}$ of the measured cross sections [27, p.16]:

$$\mathcal{D} = \sigma(\mu^{+, \leftarrow}, \phi) - \sigma(\mu^{-, \rightarrow}, \phi) \propto \Re(F_1 \mathcal{H}) \cos \phi \quad (2.44)$$

$$\mathcal{S} = \sigma(\mu^{+, \leftarrow}, \phi) + \sigma(\mu^{-, \rightarrow}, \phi) \propto \Im(F_1 \mathcal{H}) \sin \phi \quad (2.45)$$

Since the cross sections depend on the azimuthal angle ϕ between the scattering and production planes, different combinations of quark GPDs can be determined by either integrating over ϕ or analyzing the angular distribution [27, p.15]:

- The ϕ -dependent analysis of the sum $\mathcal{S}$ provides the imaginary part of the Compton form factor $\mathcal{H}$, which is in leading order given by the GPD H . This allows one to constrain parametric models for the GPD H^f in the subspace ($x = \xi, t$) at a given Q^2 .
- An analysis of the sum $\mathcal{S}$ integrated over ϕ provides information over the transverse size of the nucleon in the COMPASS-II x -domain, as it can be used to obtain the impact parameter dependent parton distributions, as shown in section 2.4.3.
- The ϕ -dependent analysis of the difference $\mathcal{D}$ provides the real part of the Compton form factor $\mathcal{H}$, which is in leading order given by a convolution of the GPD H with a function describing the hard quark-photon interaction.

Fig. 2.12 shows the projected azimuthal distribution of the beam charge & spin difference $\mathcal{D}$ in one exemplary kinematic bin obtained by calculations using two different models as well as a fit to world data. The predicted accuracy at COMPASS-II for a running time of 280 days is also given in this plot.




Figure 2.12: Projected statistical and systematic accuracy for a measurement of the ϕ dependence of the beam charge & spin difference [27, p.23]. One exemplary bin of the COMPASS-II kinematical range is shown ($1 \text{ GeV}^2 \leq Q^2 \leq 4 \text{ GeV}^2$, $0.03 \leq x_B \leq 0.07$). Predictions are calculated using the VGG model [32] and a fit to world data [33].

The analysis of the beam charge & spin difference $\mathcal{D}$ requires a precise determination of the overall incoming $\mu^{+, -}$ and $\mu^{-, +}$ fluxes and the corresponding detection efficiencies at the percent level, since these quantities have a strong impact on the systematics. A relative precision of $\approx 3\%$ or better is required for the luminosity, i.e. incoming flux times target density [27, p.90]. It is easier to measure the beam charge & spin asymmetry $\mathcal{A} = \mathcal{D}/\mathcal{S}$, as certain systematic effects cancel at least partially. Nonetheless the beam flux determination is still one of the leading contributions to the systematic error. The asymmetry is also expected to be less sensitive to theoretical corrections. However, the interpretation of the asymmetry (projections shown in Fig. 2.13) is not straightforward, since suitable models still have to be developed [27, p.23].

In 2009 a DVCS test run was performed at COMPASS to provide a first evaluation of the relative contributions of the $|\text{DVCS}|^2$ and $|\text{BH}|^2$ terms and the interference term. The experimental setup included a 40 cm long liquid hydrogen target surrounded by a 2 m long recoil proton detector, followed by the spectrometer. The beam line provided high-intensity μ^+ and μ^- beams with an energy of 160 GeV. From the recorded data set the exclusive single-photon events have been selected and their ϕ distribution is plotted in three bins in x_B (Fig. 2.14). Additionally, these plots contain the Monte Carlo simulations for the contributions to hard exclusive single-photon production by the Bethe-Heitler process ($|\text{BH}|^2$). In the low- x_B bin, which is dominated by the BH process, 251 events are observed. This is used to normalize the BH predictions to the data. In the DVCS dominated high- x_B bin 54 events are selected, while the BH contribution is estimated to 20 events. The remaining 34 events are due to pure DVCS or π^0 background, where only one of the photons from π^0 decay is detected.




Figure 2.13: COMPASS-II projections for the t dependence of the amplitude of the $\cos\phi$ modulation of the beam charge & spin asymmetry $\mathcal{A}$ for six bins in x_B , assuming a running time of 280 days [27, p.24]. Recent HERMES results (blue triangles) [34] and a fit to world data (green curves) [33] are given for comparison.




Figure 2.14: Preliminary results of the COMPASS DVCS test run in 2009. Distribution in the azimuthal angle ϕ for measured exclusive single-photon events $\mu p \rightarrow \mu' p \gamma$ with $Q^2 > 1 \text{ GeV}^2$ for three x_B bins. For comparison, the expected contributions by the Bethe-Heitler process, obtained by Monte Carlo simulations, are plotted. [27, p.36]

3. The COMPASS-II Experiment

The COMPASS experiment [35] is a fixed target experiment located in the CERN North Area, making use of high-energy, high-intensity muon or hadron beams of the SPS¹ accelerator complex. In the past, numerous measurements on the helicity and transverse spin structure of protons and neutrons have been performed. Starting in 2012, an experimental phase called COMPASS-II pursues a new physics program, including the determination of GPDs by measuring the DVCS amplitude.

The experimental setup, which will be described in this chapter, consists of the beam line, the target region and the forward spectrometer (Fig. 3.1). In the beam line the beam particles are guided to the experiment hall and their tracks and momenta are measured. The target with its surrounding detectors is placed at the interaction point. The outgoing particles are detected in the two-stage spectrometer. The angular acceptance of the first spectrometer stage (LAS²) is up to 180 mrad. Particles that pass through the central hole in the LAS enter the second spectrometer stage (SAS³), which has an acceptance of 30 mrad. Each spectrometer stage includes a dipole magnet (SM1 and SM2 respectively) to deflect charged particles in the horizontal plane, as well as detectors for particle identification, track reconstruction and energy determination [35, p.7]. Fig. 3.1 shows the spectrometer setup for the 2012 DVCS run. Please note that the laboratory coordinate system is defined such that the z-axis points in direction of the beam. The most important components of the experimental setup are introduced in the following sections. In particular, the upgrades that were performed for the DVCS measurements are described in more detail.

3.1 The Beam

The beam for the COMPASS experiment is provided by the M2 beam line, which is part of the SPS accelerator complex. Several configurations for muon or hadron beams are available, al-

¹Super Proton Synchrotron

²Large Angle Spectrometer

³Small Angle Spectrometer




Figure 3.1: Spectrometer setup of the COMPASS-II experiment for the 2012 DVCS run, consisting of a large angle spectrometer (LAS) and a small angle spectrometer (SAS). Additional tracking detectors, which are not visible in this picture, are present in both spectrometer stages. The various detectors are described in the text. [36]

lowing for a wide range of hadron structure and spectroscopy measurements. For the DVCS process polarized muon beams are required.

Protons are accelerated in the SPS to momenta up to 400 GeV/c. They are extracted in so-called spills, which are approx. 9.6 seconds long, at intervals of typically 43 seconds. The protons are deflected to a primary beryllium target (T6), whose length can be chosen depending on the required secondary beam intensity. Thereby mainly pions and kaons are produced, which are then momentum filtered and sent through a 600 m long tunnel, where they partly decay to muons during their transit. At the end of the tunnel the remaining hadrons are removed in a beryllium absorber. The muons are then focused, momentum selected and guided through a 800 m long beam line. Here the beam is bent upwards by dipole magnets to bring it from the underground SPS to the surface, where it is bent back in a horizontal direction before entering the experiment hall [35, p.13].

To reach sufficiently large beam intensity, a deviation of the muon momentum of up to 5% from the nominal value is accepted. This requires a momentum measurement for every single beam particle. The measurement is performed with the Beam Momentum Stations (BMS), which are positioned around the last bending magnet (B6) of the beam line (Fig. 3.2). The BMS are hodoscopes to measure the particle tracks in front and behind the magnet, which allows to calculate the muon momentum [35, p.15].

The muons are naturally polarized due to the parity violation of the pion and kaon decay. The degree of polarization depends on the momentum ratio of muons and hadrons. With a ratio of $p_\mu/p_\pi = (160 \text{ GeV}/c)/(172 \text{ GeV}/c)$ one can reach a polarization of $(-80 \pm 4)\%$ and an intensity of $3.8 \cdot 10^8$ muons per spill for a μ^+ beam. The maximum possible intensity for a μ^- beam is about one third of that value [27, p.19]. The majority of the muons is concentrated in the central beam spot, but there are also some muons with larger distance to the beam axis; they form the so-called halo.

It has already been mentioned in section 2.6 that the luminosity has to be determined with a precision of the order of 3%. This requires an accurate measurement of the muon beam flux, which is done by means of scintillating fiber detectors [31, p.18]. They are used to sample the beam at random points in time to determine incident beam tracks. At the same time these fiber detectors are connected to scaler modules to count the absolute number of beam particles.




Figure 3.2: The Beam Momentum Stations (BM01 - BM06) are placed around the last bending magnet (B6) of the beam line.




Figure 3.3: Sketch of the new liquid hydrogen target (not to scale). [38, 39]

3.2 The Target Region

3.2.1 The Liquid Hydrogen Target

For the COMPASS-II GPD program a new liquid hydrogen target (Fig. 3.3) has been designed and was successfully put into operation for the DVCS test run in 2012. The target is 255 cm long, thus providing a luminosity of up to $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ with the currently available beam intensity. The diameter was chosen to be 40 mm to match the transverse size of the beam. One of the design requirements was to minimize the material that has to be traversed by the DVCS final state particles. Photon absorption has to be kept at a minimum and even low-momentum protons should be able to escape the target cryostat. Therefore, the target cell is made of Kapton film with a thickness of 125 μm , as shown in Fig. 3.4. It is wrapped with a super-insulation foil consisting of 30 layers of 11 μm thick aluminum foil and placed inside a tubular vacuum cryostat. The tube with an inner diameter of 78 mm is made of 1 mm thick carbon fiber and is sealed with a 0.35 mm thick Mylar film window (Fig. 3.4, right). To cool down the target a SHI CH-110 cryocooler [37] is used, which provides a cooling power of 30 W at 20 K. After 15 hours of cooling, the nominal temperature of 18 K is reached and the isolation vacuum is $7 \cdot 10^{-7}$ mbar. For a precise determination of the luminosity it is important to maintain a homogeneous hydrogen density. Therefore the target has been carefully aligned to be perfectly horizontal.

3.2.2 The Recoil Proton Detector

For the DVCS reaction, the determination of exclusive events is mandatory. Since the missing mass method is not sufficient to ensure exclusivity at COMPASS energies, the recoiled proton has to be detected as well [27, p.89]. For this task, a new recoil proton detector – the so-called CAMERA⁴ detector – has been built in 2012. It was commissioned during the DVCS test run, which is detailed in chapter 7. A 3D drawing is shown in Fig. 3.5, some photographs can be found in appendix D. The measurement principle and detector design will be described below.

⁴COMPASS Apparatus for Measurement of Exclusive ReActions




Figure 3.4: Left: Photo of the downstream end of the liquid hydrogen target cell made of Kapton film with a Mylar end cap. Right: Photo of the Mylar film window for the vacuum cryostat. [40]




Figure 3.5: CAD drawing of the CAMERA detector viewed from downstream [41]. The overall dimensions are 4.2 m in length and 3.2 m in diameter. The barrel can be rotated on the platform by $\pm 90^\circ$ for calibration measurements with cosmics.




Figure 3.6: Individual view of the *Ring A* elements of the CAMERA detector, consisting of scintillators, light guides and photo multiplier tubes. [36]

The velocity of the recoil particle from the target is determined by a time-of-flight measurement between two barrels of scintillators, which surround the target. At the same time the energy loss of the particle traversing the scintillating material is measured. Using the Bethe-Bloch formula [42], these information allow to identify protons and distinguish them from other particles, e.g. pions or δ -ray electrons from the target.

The CAMERA detector has been designed for full azimuthal angular acceptance. In z -direction it has to cover the full length of the target, and it has to provide a large polar angular acceptance taking into account the kinematic distribution of the recoil protons. The two barrels of the detector are each segmented into 24 scintillator slats. The inner barrel (*Ring A*) with a diameter of 52 cm is composed of 275 cm long slats with a thickness of only 4 mm. This allows even slow protons (down to $p \approx 260$ MeV/c) to pass the material and reach the outer barrel for a time-of-flight measurement. In contrast, the outer barrel (*Ring B*) with a diameter of 230 cm is made from 5 cm thick material for good energy resolution. The slats are 360 cm long to provide a polar angular acceptance from 45° to 90° with respect to the beam axis. For an enhanced azimuthal angular resolution the two barrels are rotated by half a segment (i.e. 7.5°) against each other. The exact dimensions of the counters are given in Tab. 3.1.

The scintillating elements are made from BC-408 plastic scintillator material [43]. This material was chosen because it has a long light attenuation length, which is important for such long elements, but is still reasonable fast for time-of-flight measurements. Both ends of the elements are coupled via light guides to photo multiplier tubes (PMT), which are placed outside of the spectrometer acceptance. In particular for the downstream side of *Ring A* very long light guides had to be used (Fig. 3.6). Due to space constraints upstream and downstream of the detector, the PMTs of *Ring B* had to be mounted orthogonal to the scintillator elements. Therefore special light guides with a 90° turn have been manufactured (Fig. 3.7).




Figure 3.7: The light guides for *Ring B* have been manufactured at the workshop of the Freiburg institute of physics. Left: The object is milled from a solid block of acrylic glass using a 5-axis CNC machine. Right: A finished light guide. [44]

Table 3.1: Properties of the counter elements of the CAMERA detector. Each Ring (A=inner, B=outer) consists of 24 counters. [41, 47]

Element	Property	Ring A	Ring B
Scintillator	length	275.0 cm	360.0 cm
	thickness	0.4 cm	5.0 cm
	width ⁵	(6.5...6.6) cm	(29.0...30.3) cm
	dist. from beam axis ⁶	25.9 cm	114.7 cm
	material	BC-408	BC-408
PMT	type	Hamamatsu R10533	ET Enterprises 9823B
	window $\varnothing$	51 mm	130 mm
	photo-cathode active $\varnothing$	46 mm	110 mm
combined	time resolution	~ 380 ps	~ 175 ps

For *Ring A* the 2-inch PMT model *Hamamatsu R10533* [45] has been chosen. *Ring B* is read by 5-inch *ET Enterprises 9823B* [46] PMTs. In total the CAMERA detector comprises 96 channels, which are digitized with fast transient analyzer modules. Details of the readout will be described in section 3.5.1. Apart from the readout, the CAMERA detector information shall also be used for a recoil particle trigger that identifies protons in real-time. For this purpose new digital trigger electronics based on powerful programmable logic devices has been designed during the thesis at hand. The development of the trigger system is detailed in the following chapters.

For the calibration and monitoring of the CAMERA detector a laser system has been installed. Optical fibers are connected to the center of each scintillator for the distribution of laser reference pulses. These pulses create very precise time-stamps simultaneously on all detector elements, which are used for timing calibration. The T0-calibration procedure is explained in section 7.1.1. Furthermore, the measurement of the pulse amplitudes simplifies the procedure of finding the correct high-voltage settings for the PMTs, and it allows for the monitoring of the PMT gain stability.

⁵trapezoidal cross section

⁶measured to the center of the scintillator

3.3 The Spectrometer

3.3.1 Tracking Detectors

In the COMPASS spectrometer a large variety of tracking detectors is used to precisely determine the tracks of final state particles. They are also utilized to measure the momentum of particles that are deflected in the dipole field of the spectrometer magnets. Table 3.2 summarizes the properties of the different tracking detectors.

Depending on the distance from the beam axis, different spatial and time resolutions and rate capabilities are required. Hence the tracking detectors are classified in three regions [35, p.10]. Closest to the beam there are the *Very small area trackers (VSAT)*. These are silicon strip detectors, scintillating fiber detectors, Pixel-GEMs⁷ and Pixel-Micromegas⁸. They cover a range of up to 3.5 cm from the beam axis and provide a high rate stability. The scintillating fibers provide a good time resolution of 400 ps, which allows for a correlation between hits and beam tracks without ambiguity even for the high rates of up to $3 \cdot 10^6 \text{ s}^{-1}$ per fiber in the beam spot.

In the intermediate region between 2.5 cm and 40 cm from the beam axis GEM and Micromegas detectors with strip readout represent the *Small area trackers (SAT)*.

Finally, the outer regions of the spectrometer are covered by the *Large area trackers (LAT)*, which are multi-wire proportional chambers (MWPC), drift chambers and straw detectors of various sizes and shapes.

3.3.2 Calorimeters

In each COMPASS spectrometer stage there is one electromagnetic and one hadronic calorimeter (cf. table 3.3) to measure the energy of electrons, photons and hadrons in the final state. Originally the electromagnetic calorimeters ECAL1 and ECAL2 were composed of homogeneous lead glass modules, where photons lose their energy in electromagnetic showers by means of bremsstrahlung and pair production. The Cherenkov light of the shower electrons is detected with PMTs. Since the depth of the calorimeter blocks is approx. 16 radiation lengths, more than 99% of the particle energy is contained in the shower [35, p.47]. In 2008 the central part of ECAL2 was replaced by sampling calorimeter modules with alternating layers of lead and scintillator (the so-called *Shashlyk* modules), providing higher radiation hardness and a better energy resolution than the lead glass modules.

The hadronic calorimeters HCAL1 and HCAL2 are sampling calorimeters with alternating layers of iron and plastic scintillator. Hadrons generate hadronic showers, while interacting with the iron, which are then detected in the scintillating layers. Since the hadronic interaction length is much greater than the radiation length, the depth of the hadronic calorimeters is much larger than the depth of the electromagnetic calorimeters. Therefore the HCAL1 and HCAL2 are located in the spectrometer immediately behind ECAL1 and ECAL2 respectively, in order to not spoil the good energy resolution of the ECALs. As a side effect, the hadronic shower generation already starts in the ECALs.

⁷Gas Electron Multiplier

⁸Micromesh Gaseous Structure

The New ECAL0

For the GPD program an additional electromagnetic calorimeter ECAL0 was developed to cover larger photon angles. On the one hand, this increases the accessible x_B region [27, p.99], on the other hand it is also needed for the π^0 background suppression. ECAL0 is placed directly after the CAMERA detector to reduce its geometrical size. The angular acceptance should match the aperture of the CAMERA detector. This results in a total size of $220 \times 220 \text{ cm}^2$ with a central hole of about $50 \times 50 \text{ cm}^2$. The hole of ECAL0 matches the outer acceptance of ECAL1. ECAL0 is built from modules sized $12 \times 12 \text{ cm}^2$, which requires ≈ 300 modules for the final construction. However, during the 2012 DVCS test run, a smaller prototype ($132 \times 108 \text{ cm}^2$) with 56 modules was used [48].

For the construction of ECAL0 a new type of sampling calorimeter module (Fig. 3.8) has been developed [49]. This was necessary for two reasons. On the one hand, due to space constraints the total length of the module is limited. On the other hand, the close proximity to the spectrometer magnet SM1 required a design that is insensitive to magnetic fields.

The calorimeter part of the Shashlyk module is built from 109 sampling sandwiches with an edge length of 12 cm. Each sandwich consists of a 0.8 mm thick lead plate and a polystyrene-based scintillator⁹ layer with a thickness of 1.5 mm. The scintillator is segmented into 9 tiles of $4 \times 4 \text{ cm}^2$ each. The edges of the tiles are painted white to form 9 light-isolated towers, which are separately read out with wavelength shifting (WLS) fibers¹⁰. The total length of the calorimeter stack is 25 cm, which corresponds to ≈ 15 radiation lengths.

In order to be insensitive to magnetic fields, no PMTs are used for photon detection. Instead, multi-pixel avalanche photo diodes (MAPD) from Zecotek Photonics [50] were chosen to perform this task. The WLS fibers are coupled to a MAPD-3N photo diode separately for each of the readout towers, thus providing 9 individual channels per Shashlyk block. The MAPD-3N has its highest photon detection efficiency in the green region, which matches the emission of the WLS fibers. Because the gain of MAPDs is temperature dependent, a peltier element is used to stabilize the temperature of the photo diodes. The complete module including calorimeter, readout and cooling has a length of 42 cm. Test-beam measurements [51] confirmed, that the MAPD readout is comparable to the traditional PMT readout in terms of linearity and energy resolution, when the temperature was stabilized at 15°C .

3.3.3 Muon Identification

To identify the scattered muons, there are several hadron absorbers, dubbed Muon Filters (MF), placed in the spectrometer. They are enclosed by detector layers to form the so-called Muon Walls (MW). MW 1 at the end of the LAS consists of a 60 cm thick iron absorber (MF 1) between two groups of drift tube detectors. MW 2 at the end of the SAS is built from a 2,4 m thick concrete absorber (MF 2) accompanied by drift tubes and MWPCs. These detectors act as filters for

⁹The used material is PS + 2% pTP + 0.05% POPOP. The primary scintillator para-Terphenyl (pTP) emits light with a wavelength of 340 nm. The secondary scintillator POPOP is used to shift the emission to a longer wavelength (410 nm).


¹⁰Saint-Gobain BCF-91A. Shifts blue to green. Peak emission at 494 nm.

Table 3.2: COMPASS tracking detectors with typical dimensions, spatial and time resolutions. [35, 52]

	Detector type	Active area	Spatial resolution	Time res.
VSAT	Scint. fibers	$(3.9\text{ cm})^2 \dots (12.3\text{ cm})^2$	$(130 \dots 210)\text{ }\mu\text{m}$	0.4 ns
	Silicon strips	$5 \times 7\text{ cm}^2$	$(8 \dots 11)\text{ }\mu\text{m}$	2.5 ns
	Pixel-GEM	$10 \times 10\text{ cm}^2$	$95\text{ }\mu\text{m}$	9.9 ns
SAT	GEM	$31 \times 31\text{ cm}^2$	$70\text{ }\mu\text{m}$	12.0 ns
	Micromegas	$40 \times 40\text{ cm}^2$	$90\text{ }\mu\text{m}$	9.0 ns
LAT	MWPC	$178 \times (90 \dots 120)\text{ cm}^2$	$1600\text{ }\mu\text{m}$	
	Drift chambers (DC)	$180 \times 127\text{ cm}^2$	$190\text{ }\mu\text{m}$	
	Large area DC	$500 \times 250\text{ cm}^2$	$500\text{ }\mu\text{m}$	
	Straws	$280 \times 323\text{ cm}^2$	$190\text{ }\mu\text{m}$	

Table 3.3: Electromagnetic and hadronic calorimeters in the COMPASS-II spectrometer setup. An ECAL0 prototype was installed in 2012. The full detector is currently under construction. [35, 48]

Calorimeter	Active area ($w \times h$)	Channels	Energy resolution $\frac{\Delta E}{E}$
ECAL0 ₂₀₁₂	$132 \times 108\text{ cm}^2$	504	$0.08\sqrt{\frac{\text{GeV}}{E}} \oplus 0.02$
ECAL0	$220 \times 220\text{ cm}^2$	≈ 2700	$0.08\sqrt{\frac{\text{GeV}}{E}} \oplus 0.02$
ECAL1	$397 \times 286\text{ cm}^2$	1476	$0.06\sqrt{\frac{\text{GeV}}{E}} \oplus 0.02$
ECAL2	$245 \times 184\text{ cm}^2$	3072	$0.06\sqrt{\frac{\text{GeV}}{E}} \oplus 0.02$
HCAL1	$420 \times 300\text{ cm}^2$	480	$0.59\sqrt{\frac{\text{GeV}}{E}} \oplus 0.08$
HCAL2	$440 \times 200\text{ cm}^2$	216	$0.66\sqrt{\frac{\text{GeV}}{E}} \oplus 0.05$

**Figure 3.8:** Schematic view of the Shashlyk calorimeter module for ECAL0. [49]

the muons, which are able to pass the absorbers due to their weak interaction, while all other particles are absorbed. Therefore, particles that are detected both in front of and behind the absorbers can be identified with muons.

3.3.4 RICH-1

The RICH-1 detector in the COMPASS spectrometer is a large volume ring imaging Cherenkov detector, which is used to identify particles by means of the Cherenkov effect. The detector is filled with the radiator gas C_4F_{10} with a refractive index of $n = 1.0015$ at atmospheric pressure and a temperature of 25°C . A particle passing the radiator gas emits Cherenkov light under a certain angle ϕ_c , which depends on the velocity $\beta = v/c$ of the particle:

$$\cos\phi_c = \frac{1}{n\beta} = \frac{1}{n} \frac{1}{\sqrt{1 + m^2/p^2}}. \quad (3.1)$$

The light is reflected by spherical mirrors to photon detectors in the focal plane, where the image appears as a ring. From the ring diameter the Cherenkov angle can be calculated, and using the particle momentum p measured in the spectrometer one obtains the mass m of the particle via (3.1). With RICH-1 pions, kaons and protons can be separated in the momentum range between 5 GeV/c and 43 GeV/c.

In the inner region of the detector plane, where the occupancy is highest, multi-anode photomultiplier tubes are used for the photon detection. The outer region is currently covered by MWPCs with CsI coated photo-cathodes. In the course of the COMPASS-II upgrades, these chambers will be replaced by THGEM¹¹ based detectors, which have been recently developed [53–55].

THGEMs are electron multipliers, which are basically derived from GEMs by scaling the geometrical parameters and changing the production technology. Instead of thin foils, as they are used for GEMs, the new detector type is built of standard Cu-coated PCB¹² material with a thickness between 0.2 mm and 1.0 mm. Therefore the THGEM layers are quite stable, allowing to produce large size detectors. The layers with a size of up to $60 \times 60 \text{ cm}^2$ are manufactured with standard PCB technology, by drilling holes with a typical diameter of 0.4 mm and a pitch of 0.8 mm. Finally a clearance ring around the hole is produced by Cu etching (see Fig. 3.9).

A THGEM based detector (Fig. 3.10) is typically constructed of a plane of wires defining a drift field, three THGEM layers for electron multiplication and an anode layer to collect the charge. The detector is covered with a fused silica window to allow UV photons to enter the chamber, and it is operated with a Ar/CH_4 gas mixture. The top side of the first THGEM is coated with CsI to act as a photo-cathode. To each THGEM layer a bias voltage ΔV is applied between top and bottom side, generating an electric field, which is particularly strong in the holes, leading to electron multiplication. Between the layers transfer fields guide the electrons to the next amplification stage, and the induction field finally guides the electrons to the anode pads, where the signals are read out.

¹¹Thick GEM

¹²Printed Circuit Board




Figure 3.9: The THGEM layer is produced of Cu-coated halogen-free fiberglass PCB material. The holes are drilled and the clearance ring (rim) is produced by etching. [54]




Figure 3.10: Structure of a THGEM-based photodetector. [53]

3.4 The Trigger System

With a beam intensity of up to $3.8 \cdot 10^8$ muons per spill it is obvious, that not all events can be recorded. Therefore, it is essential to use a trigger system that is able to identify the events, which are possibly interesting for the analysis. Due to the limited buffer memory of the front-end electronics the time available for a trigger decision is typically in the order of $1 \mu\text{s}$. Several trigger types can be combined to form the COMPASS first-level trigger, which is then distributed via the Trigger Control System (TCS) [56] to all readout modules.

The TCS signal consists of two channels which are time-division multiplexed, encoded with a biphasic mark code and transmitted by means of a unidirectional optical fiber network. Channel A is used to transmit a 1-bit first-level trigger information only, thus providing a fixed latency and low dead-time. All other information is transmitted consecutively on channel B, including event numbers, spill numbers and synchronization signals like *Begin of Spill* or *End of Spill*. Last but not least the experiment's reference clock, which is used to encode the TCS signal, is recovered by the readout modules to obtain a global time reference.

Muon Trigger

For the DVCS measurement a trigger on scattered muons in a large kinematic range in Q^2 and x_B is mandatory. There are four muon trigger systems (Fig. 3.11), each being responsible for a different kinematic region [27, p.84]:

- the Ladder trigger (H4L, H5L) covers low Q^2 and high y ,
- the Middle trigger (H4M, H5M) covers low Q^2 and all y ,
- the Outer trigger (H3O, H4O) covers intermediate Q^2 and all y ,
- the LAS trigger (H1, H2) covers large Q^2 and high x_B .

Each system consists of two scintillating hodoscopes, of which at least one is located behind a muon filter. To detect the scattered muon two different approaches are used. For the *target pointing trigger* the vertical position of the scattered muon is measured with two horizontal hodoscope layers to determine the scattering angle in the y - z plane. In this plane the particle tracks are not bent by the spectrometer magnets, so the track can be extrapolated to $z = 0$, which is the z -position of the target center. The vertical position of the track at $z = 0$ is then tested for compatibility with the target position (vertical target pointing). This method is used by the Outer and the LAS system, which trigger on events with medium and large Q^2 .

At low Q^2 the scattering angle of the muon is too small for target pointing, so these events are triggered by the *energy loss trigger*. The energy loss y of the muon is determined from the deflection in the dipole fields of the spectrometer magnets. Using a coincidence of two vertical hodoscope strips in the Ladder system, muons with a large deflection in the magnets but very small scattering angle can be selected yielding events with a large energy loss, but small Q^2 . Finally, the Middle trigger system consists of both vertical and horizontal hodoscope layers, hence combining the target pointing and the energy loss method [57].




Figure 3.11: Position of the trigger hodoscopes in the COMPASS-II spectrometer. The inner trigger (H4I and H5I) is not used for the DVCS measurement. Adapted from [57, p.223].

Veto System

Although most of the beam muons are concentrated in a small spot, there are still many halo muons which are able to cause a false trigger when crossing the hodoscopes. To suppress these events, several veto elements are installed upstream of the target. They consist of segmented scintillation counters with a central hole, where the beam is allowed to pass [27, p.86].

Proton Trigger

For the DVCS measurements it was decided to develop a new proton trigger system, to enrich the sample of exclusive events with a proton in the final state. This system shall be able to identify recoil protons based on the information from the CAMERA detector in real-time. At the same time, background signals from δ -ray electrons or pions should be distinguished and prevented from generating a false trigger signal, which is only possible by sophisticated digital processing. Since the trigger signal has to be included in the first-level trigger, the available time for the trigger decision is limited. The development of the new CAMERA proton trigger system, which was conducted during this thesis, will be described in detail in the following chapters. The underlying concept is presented in chapter 4. The newly-created hardware platform – the so-called TIGER¹³ board – and the corresponding firmware are detailed in chapters 5 and 6.

3.5 The Data Acquisition System

The COMPASS data acquisition system (DAQ) is responsible for the readout of about 300 000 detector channels at trigger rates of up to 100 kHz. The high beam intensity results in hit occupancies of up to several MHz in single channels. Altogether the whole spectrometer produces several Gigabytes of data per second during a 9.6 s long spill. To handle this amount of data, a custom DAQ system was designed [58], which will be described in the following.

¹³Trigger Implementation for GANDALF Electronic Readout




Figure 3.12: Overview of the COMPASS data acquisition system, showing the number of involved modules and the maximum data transfer rates of the various links.

A schematic overview of the data flow from the detectors to the final storage in the CERN computing center is given in Fig. 3.12. The signals of the various detectors are digitized on front-end cards, which are either mounted directly on the detectors or located next to them. According to requirements the digitizers are of the ADC¹⁴, TDC¹⁵ or Scaler type. Upon receiving a trigger, the digitized data is transferred via HOTLink-based [59] copper or optical links to VME readout modules, which can be either CATCH¹⁶, GeSiCA¹⁷ or HotGeSiCA modules [35, p.54]. Recently, with the GANDALF¹⁸ module (see section 3.5.1 below), a fourth type of readout module has been introduced to the COMPASS DAQ system.

The readout modules collect the data from the front-ends and merge it according to the event numbers, while adding some header information for the experiment-wide unique channel number identification. The final data stream is transmitted via S-Link [60] to readout buffer PCs (ROB), where the data of one spill is buffered in RAM before being distributed via Gigabit Ethernet to event builder PCs (EVB) in a round-robin manner. The EVBs receive the sub-events from the ROB for a specific event number and concatenate them to form the final event, which is then stored in the DATE format [61]. The resulting files are transferred to the central data recording service CASTOR¹⁹ where they are stored on magnetic tapes.

¹⁴Analog to Digital Converter

¹⁵Time to Digital Converter¹⁶COMPASS Accumulate, Transfer and Control Hardware

¹⁷GEM and Silicon Control and Acquisition

¹⁸Generic Advanced Numerical Device for Analytic and Logic Functions

¹⁹CERN Advanced STORage manager

When the current DAQ system was designed back in 2001, it could easily handle the data rate coming from the detectors. Since then it has been upgraded several times to keep up with an increased number of detector channels and higher beam intensities and trigger rates. Nevertheless, after 12 years of successful data taking, the even higher data rates of the upcoming COMPASS-II physics program will require a new event building system, which is currently in development [62].

3.5.1 The GANDALF Framework

With the GANDALF framework [63], a new generation of readout electronics has been introduced to accommodate a growing demand for high-performance digitizers in particle physics experiments. Originally designed at the Freiburg institute of physics for the readout of the CAMERA detector, the GANDALF module is now used in various configurations for several detectors in the COMPASS experiment.

GANDALF is a versatile detector readout system based on the VME64x/VXS form factor²⁰, consisting of a generic mainboard and a set of application specific add-on mezzanine cards²¹. The GANDALF board is capable of implementing sophisticated real-time digital signal processing (DSP) algorithms, owing to its high-performance FPGA²² devices and large on-board memories. Two mezzanine slots per board can be populated with add-on cards of various types, to adapt the signal inputs to the particular requirement. High-speed data buses with a throughput of 48 Gbit/s per mezzanine slot transmit the acquired signals directly to the FPGA on the mainboard.

Currently three types of mezzanine cards are available, which can be chosen depending on the intended functionality and the signals that arrive from the detector front-end:

- The **ADC mezzanine card** [64] hosts eight 12 bit @ 500 MS/s analog-to-digital converters for continuous digitization of up to eight detector channels. Two ADCs at a time can optionally be combined to sample the same signal in time-interleaved mode resulting in a 1 GS/s digitization of four analog channels per mezzanine card. The raw data stream is forwarded to the DSP FPGA on the mainboard for real-time pulse shape analysis and feature extraction. This operation mode is called the *GANDALF transient analyzer* [65, 66].
- The **digital I/O mezzanine card** [67] is capable of reading out 64 detector channels, accepting signals in the differential LVDS or LVPECL standard as they are usually provided by pre-amplifier and discriminator modules. The acquired signals are directly fed into the FPGA, where the user can select from various processing functionalities which have been implemented so far: a time-to-digital converter (called *M1 TDC*) [68, 69], a scaler [70], a combination of both [71], or a mean-timer with subsequent coincidence logic [72].

²⁰More information about the VMEbus architecture can be found in section 4.3.1.

²¹Daughtercards, which are mounted on top of and parallel to the mainboard.

²²Field-Programmable Gate Array

- The **ARWEN optical mezzanine card** [73] is available for receiving data from digital front-end cards, which are typically mounted directly on the detector. The connection is established via optical fibers using a constant-latency serial protocol for the up-link, allowing to synchronously transmit a reference clock and trigger signals to the front-end cards, and a 3 Gbit/s high-speed serial down-link to receive detector data, which has been digitized on the front-end cards. A TDC front-end card utilizing this optical connection is currently developed for the readout of the prospective THGEM chambers for the RICH-1 detector.

CAMERA Readout

For the readout of the CAMERA detector 12 GANDALF transient analyzer modules are employed to digitize the 96 PMT channels with 1 GS/s resolution. An online pulse-feature extraction based on a digital constant fraction discrimination (dCFD) is performed by the on-board FPGAs, yielding a significant reduction of the data volume which has to be transferred to the DAQ. After the signal processing each PMT pulse is represented by a set of three values, namely a time stamp, an amplitude and an integral. For a small fraction of all events the full waveform of the pulse is also sent to the DAQ, which enables a monitoring of the feature extraction algorithm.

Thanks to the interpolation of the dCFD method [74] the timing resolution of the transient analyzers is much better than the 1 ns sampling, but the actual value depends on the signal amplitude. The digitized signal is represented by discrete samples with 4096 possible amplitude values due to the 12 bit ADC. Hence, it is obvious that the interpolation is more precise the higher the pulse amplitude is with respect to the ADC dynamic range. It has been shown [66] that the dCFD timing resolution is inversely proportional to the pulse amplitude, resulting in a timing resolution better than 100 ps for pulses exceeding 2% of the dynamic range, and better than 20 ps for pulses which are larger than 10% of the dynamic range.

Further Detectors

Due to the increased beam intensity for the DVCS measurement, the readout of several detectors has been changed from CATCH modules to GANDALF modules. In particular the scintillating fiber detectors in front of the target, which sustain the highest hit rate, are connected to GANDALF TDCs since 2012. Thereby the limit for the trigger latency could be slightly relaxed without degrading the detectors' efficiency thanks to the much larger hit buffers of the FPGA-based M1 TDCs.

4. CAMERA Proton Trigger

The COMPASS-II GPD program includes the study of the DVCS process, as described in section 2.5. To enable the selection of the rare DVCS events, the development of a new trigger system which is sensitive to the recoil protons was mandatory. The first section of the chapter at hand briefly recalls the measurement principle of the CAMERA detector. In the following section, the concept for the novel proton trigger system is developed. Finally, this chapter is concluded with a description of the electronics framework, which is provided for the readout of the CAMERA detector. The TIGER module, which constitutes the hardware platform for the new trigger system, was designed to operate within this framework. Its development is an integral part of this thesis and is described in detail in chapter 5.

4.1 Detector Principle

Since the detector has already been described in section 3.2.2, its features will only be briefly summarized. Two barrels of scintillating material, each segmented in 24 slats, are surrounding the target. The elements of *Ring A* are located in a distance of 26 cm from the beam axis and the *Ring B* segments are placed 115 cm away from the beam axis. Each slat is read on both ends with PMTs and the signals are digitized and processed by GANDALF modules. For each digitized pulse, a set of features is extracted, containing the pulse arrival time, the maximum amplitude and the pulse integral.

The following nomenclature is used to identify the readout channels of the detector: $X_{i,e}$, where $X \in \{A, B\}$ specifies the ring (inner or outer), $i \in \{0, 1, \dots, 23\}$ specifies the index of the element, and $e \in \{up, dn\}$ specifies the end of the counter (upstream or downstream). A scintillator is referenced by the name X_i . When dealing with a particular counter, the index i may be omitted for easier notation.

4.1.1 Time-of-Flight Measurement

In Fig. 4.1 one segment of the CAMERA detector with two facing elements is shown to illustrate the measurement principle. A particle originating from the target is crossing the two detector

elements A and B , generating scintillation light while traversing the material. This event of light creation in a detector element is in the following called a hit. The light pulses propagate to the ends (up and dn) of the counters, where they arrive at time $t_{A,up}$, $t_{A,dn}$, $t_{B,up}$ and $t_{B,dn}$ respectively.

The time t_A and position z_A of the hit in counter A is then calculated as follows:

$$t_A = \frac{t_{A,up} + t_{A,dn}}{2}, \quad (4.1)$$

$$z_A = c_{A,eff} \frac{t_{A,up} - t_{A,dn}}{2}, \quad (4.2)$$

where $c_{A,eff}$ is the effective speed of light in the A scintillator. For the hit in counter B , time and position t_B and z_B are calculated analogously. The time calibration of the CAMERA detector (cf. 4.2.2.2) is chosen such, that the positions z_A and z_B are given relative to the center of the corresponding counter $z_{0,A}$ and $z_{0,B}$ respectively.

By combining the hit information from the inner and the outer counter, the particle track can be reconstructed. The time of flight t_{ToF} and the distance of flight s is used to calculate the velocity β_{ToF} and the polar angle θ as follows:

$$t_{ToF} = t_B - t_A, \quad (4.3)$$

$$\Delta z = z_B - z_A + k, \quad (4.4)$$

$$s = \sqrt{d^2 + \Delta z^2}, \quad (4.5)$$

$$\beta_{ToF} = \frac{s}{c \cdot t_{ToF}}, \quad (4.6)$$

$$\theta = \arctan\left(\frac{d}{\Delta z}\right), \quad (4.7)$$

where $k = z_{0,B} - z_{0,A}$ is the offset in z -direction of the center of counter B with respect to the center of counter A , and d is the radial distance between the inner and the outer counter.

4.1.2 Energy Loss Measurement

After the track reconstruction has been performed, the measured PMT pulse amplitudes are used to obtain additional information for the particle identification. The geometric mean $E_X = \sqrt{E_{X,up} \cdot E_{X,dn}}$ of the upstream and downstream PMT pulse amplitudes after calibration is proportional to the energy loss of the particle in the scintillator.

The mean energy loss of a charged particle traversing material is described by the Bethe-Bloch equation [17, p.324]

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right]. \quad (4.8)$$

Here $T_{\max}$ is the maximum kinetic energy transfer in a single collision, z is the charge number of the incident particle, $\gamma = (1 - \beta^2)^{-1/2}$ is the Lorentz factor, and $K = 0.307 \text{ MeV g}^{-1} \text{ cm}^2$ combines some constants.




Figure 4.1: Principle of the Time-of-Flight measurement with the CAMERA detector. A recoil particle (p') generates scintillation light when crossing the detector elements A and B (indicated by the red dots). The light pulses propagate to the upstream and downstream ends of the counters (red dashed arrows), where they are detected by PMTs. From the pulse arrival times the particle track can be reconstructed.

In a specific material with charge number Z , atomic mass A and mean excitation energy I , the stopping power (i.e. $-\langle dE/dx \rangle$) for charged particles with mass $M \gg m_e$ is approximately

$$-\left\langle \frac{dE}{dx} \right\rangle \propto \frac{1}{\beta^2} \ln(\text{const} \cdot \beta^2 \gamma^2). \quad (4.9)$$

The Bethe-Bloch equation is valid in the region $0.1 \lesssim \beta\gamma \lesssim 1000$. At the lower limit one has to take the orbital velocity of the electrons into account with the so-called shell correction, while at the upper limit radiative effects become more important than ionization. The δ -term corrects the density effect which is caused by the polarizability of the medium [17, p.326].

In Fig. 4.2 the stopping power in several materials is plotted as a function of $\beta\gamma = p/Mc$. Axes for the muon, pion and proton momenta are also drawn for comparison. The data points for a vinyl-toluene-based plastic scintillator material like the BC-408, which is used for the CAMERA detector, are taken from the PSTAR database [75]. The qualitative behavior of the stopping power functions is similar in most materials. They show a broad minimum around $\beta\gamma = 3.5$, where particles are minimum-ionizing¹. At lower energies $-\langle dE/dx \rangle$ falls with the kinematical factor $\beta^{-5/3}$, while the rise at higher energies is caused by the relativistic extension of the transverse electric field.

The β -dependence of the stopping power can now be utilized to identify recoil particles with the CAMERA detector. The expected energy loss for protons and pions in the CAMERA scintillators has been simulated with the COMPASS-II Monte-Carlo toolkit TGEANT and is shown in Fig. 4.3. One can see, that for $\beta < 0.6$ pions can be separated from protons to a large extent.

¹particles with mean energy loss rates close to this minimum are called minimum-ionizing particles (MIP)




Figure 4.2: Mean energy loss rate in several materials as a function of $\beta\gamma = p/Mc$, with additional axes for muon, pion and proton momenta. Radiative effects are not included [17, p.325]. The data points ($\star$) for plastic scintillator are from [75].




Figure 4.3: Energy loss of protons (red) and pions (blue) in *Ring B* vs. velocity β for tracks with a polar angle θ of 50°, 70° and 90°. Simulation with TGEANT. [76, p.27]

4.2 Trigger Concept

4.2.1 Design Objectives

The existing COMPASS trigger system is mainly based on the detection of scattered muons. However, for the 2.5 m long liquid hydrogen target and the maximum possible beam intensity the anticipated trigger rate would saturate the readout electronics and the DAQ system. Therefore it was decided to develop a proton trigger system which can be integrated in the existing TCS (cf. 3.4). Based on the CAMERA detector information, which is provided by the GANDALF readout modules, the so-called TIGER trigger shall be able to identify events with a proton in the final state.

Simulations [27, p.19] have shown that the projected total number of DVCS events is in the order of 55 events per day, while the number in certain (x_B, Q^2) bins can be as low as one or two events per day. It is therefore mandatory for the TIGER trigger to be very efficient. On the other hand, the rate of particles in the CAMERA detector is expected to be in the order of several MHz, mainly due to the production of δ -ray electrons in the target material and the surrounding walls and due to pions from fragmentation processes. These background signals should be identified and subsequently suppressed in the trigger decision, resulting in a high purity of the proton trigger signal. These two key requirements – very high efficiency and good purity – are best fulfilled by utilizing the advantages of high-performance programmable logic devices. The TIGER module provides a dedicated hardware platform to implement the proton trigger logic based on the following conditions.

4.2.2 Trigger Conditions

4.2.2.1 Geometric Coincidence

A particle from the target is passing through *Ring A* and *B* of the CAMERA detector, producing a pair of hits in the scintillator slats. Due to the azimuthal rotation of the two detector barrels by 7.5° against each other to increase the angular resolution, two *A–B* combinations are possible for every element (see Fig. 4.4). The coincidence logic therefore has to check a total number of 48 combinations between *A* and *B* counters:

$$C_i = (A_i \otimes B_i) \oplus (A_i \otimes B_j), \quad (4.10)$$

$$C = C_0 \oplus \dots \oplus C_{23}, \quad (4.11)$$

for $i \in \{0, 1, \dots, 23\}$ and $j = (i + 1) \bmod 24$.

Here $(X \otimes Y)$ denotes the coincidence between two hits in the elements *X* and *Y*, and *C* is the logical disjunction of all 48 possible coincidences.

The conditions under which two hits are meant to be in coincidence are defined as follows [77, p.55]:

- $\Delta z \geq -\delta_z$ (scattering in forward direction),
- $t_{\min} \leq t_{\text{ToF}} \leq t_{\max}$ (reasonable time of flight).




Figure 4.4: Coincidence logic between the inner and the outer ring of the CAMERA detector. For every inner counter A_i there are two possible outer counters B_i and B_{i+1} , which have to be checked for corresponding hits.

Since protons can only be scattered in forward direction, the lower limit for Δz is principally 0, but to accommodate the finite position resolution of the counters an additional margin δ_z is allowed. The limits for the time of flight are chosen according to the CAMERA geometry and the detectable β range. To be on the safe side, this can be e.g. set to [0 ns, 80 ns].

4.2.2.2 Time Calibration

The description of the time measurement in 4.1.1 has not covered the need for calibrations. Actually, the time stamps that are measured by the GANDALF module have an individual offset for every channel due to differences of the effective speed of light in the scintillators, as well as due to different lengths of the light guides and cables and individual transit times of the PMTs. This is calibrated by means of the CAMERA laser system. Using the optical fibers, which are connected to the center of each scintillator, a laser pulse is coupled into all slats simultaneously. The individual time offsets are then adjusted such that

$$z_{Ai} = 0 \text{ and } z_{Bi} = 0 \quad \forall i \in \{0, 1, \dots, 23\}, \quad (4.12)$$

$$t_{\text{ToF}}(A_i \rightarrow B_j) = 0 \quad \forall i, j \in \{0, 1, \dots, 23\}. \quad (4.13)$$




Figure 4.5: Distribution of the kinetic energy of protons at the vertex from DVCS events, which have been generated with HEPGen [76, p.85]. These are not reconstructed events and no detector response is simulated for this plot.

Dedicated calibration runs have been taken with the laser system during the commissioning phase of the 2012 DVCS test run. From this data t^0 offsets for the individual channels have been determined (cf. 7.1.1). They are needed for an online correction of the time stamps in the trigger module. For the offline analysis the calibration is performed following a different approach, based on physical events instead of laser pulses. The two methods are compared in section 7.2.2.

4.2.2.3 Cuts on Energy Deposition

With the geometric correlation alone it is not possible to trigger only on protons. The majority of the tracks originating from the target are generated by δ -ray electrons or pions, leading to a high rate of false trigger signals. The β -dependent energy loss of charged particles in the scintillator slats according to the Bethe-Bloch equation can be exploited to suppress these undesirable coincidences.

To determine a set of cuts, which can be applied to the energy deposition in the counters, a large sample of DVCS events has been simulated with a Monte-Carlo toolchain [76, 78]. The events are generated with HEPGen² – a Monte-Carlo generator, which is dedicated to studies of exclusive processes in the DIS domain. One can see from Fig. 4.5, that the kinetic energy of the protons at the vertex is mostly below 500 MeV, which corresponds to a velocity of $\beta \lesssim 0.75$. The generated events are forwarded to TGEANT – a Geant4-based simulation of the COMPASS-II spectrometer and in particular the CAMERA detector. Here the detector response is simulated and the events are reconstructed. One obtains that protons with a kinetic energy of less than ≈ 35 MeV at the vertex do not reach the B scintillator. This results in a lower limit for the velocity

²Hard Exclusive Production GENerator




Figure 4.6: Simulated energy loss of the DVCS proton in *Ring A* (left) and *Ring B* (right) of the CAMERA detector vs. the proton velocity β_{ToF} , which is determined from the time of flight [78, p.90]. The continuation of the rising trend above the cusp in the *Ring B* energy loss distribution is due to proton tracks with a small polar angle θ . Events have been generated with HEPGen, and detector response has been simulated with TGEANT.

of detectable protons of $\beta \approx 0.27$ at the vertex. Due to the large energy loss of these slow protons in the target material, the *A* counter and the air, the velocity β_{ToF} that is reconstructed from the time-of-flight measurement between the *A* and *B* counters is significantly lower ($\beta_{\text{ToF}} \approx 0.15$). While this effect is corrected for in the offline analysis (details in [79]), for the real-time trigger processing one is limited to the velocity calculated from the ToF measurement, i.e. the mean velocity between the *A* and *B* counter.

From the simulation one obtains the energy loss of the protons in the *Ring A* and *Ring B* scintillators as well as the reconstructed velocity β_{ToF} . The resulting distributions of energy loss vs. β_{ToF} are shown in Fig. 4.6. The left hand plot shows the energy loss in *Ring A*, which basically follows the fall of the Bethe-Bloch formula. One can clearly see the lower limit of $\beta_{\text{ToF}} \approx 0.15$, since no time-of-flight calculation is possible for protons which do not reach the *B* counter. The right hand plot shows the energy loss in *Ring B*. The rising branch below $\beta_{\text{ToF}} \approx 0.4$ contains the protons that are fully stopped in the *B* scintillator releasing all their remaining kinetic energy, while protons above this limit can pass through the *B* counter, which again results in the characteristic $\beta^{-5/3}$ fall according to Bethe-Bloch.

The broadening of the proton band to higher energy losses is caused by two effects:

- Depending on the polar angle θ of the proton track the distance through the scintillator varies by a factor of up to ≈ 2 .
- The energy loss probability distribution is best described by a highly-skewed Landau distribution, with a most probable value for the energy loss well below the mean $\langle dE/dx \rangle$ that is given by the Bethe-Bloch equation. Due to the very rare high-energy-transfer collisions the energy loss can fluctuate significantly. [17, p.327]




Figure 4.7: Simulated energy loss of the DVCS proton in the inner (A) and outer (B) scintillators of the CAMERA detector. [78, p.89]

Another possibility to show the energy loss of the protons in the CAMERA detector is to plot the energy loss in the inner scintillators versus the energy loss in the outer scintillators (Fig. 4.7), resulting in a representation of the proton band which is not depending on β . The branch above the 'knee' consists of the protons which are stopped in B, while the lower branch contains the protons which are fast enough to leave the detector.

4.2.2.4 Correlation with the Beam

While leaving the target, the recoil protons can lose a large amount of energy. Especially the slow protons ($\beta \lesssim 0.4$) are decelerated significantly by the liquid hydrogen and the material surrounding the target, i.e. the Kapton cell, the super-insulation foil and the carbon tube [79]. For this reason and also due to the moderate time resolution of the A counters, it is not possible to obtain a precise vertex time from the proton alone.

To optimize the timing of the trigger signal, which is required to have a constant time offset with respect to the incident beam particle, additional information from another detector is needed. The scintillating fiber (SciFi) stations FI01 and FI02 located upstream of the target can provide a precise timing of the beam particles. The SciFis use 16-channel multi-anode PMTs for photon detection. Apart from the anode outputs, which are connected to TDCs for detector readout, there is an output of the last common dynode stage of the PMT, which basically provides an 'or' signal of 16 channels. By digitizing these dynode signals one obtains time stamps of the beam particles, which can later be used for a re-timing of the trigger output.

It should be mentioned here that for the offline analysis a precise vertex time is also essential for improvement of the momentum resolution of the CAMERA detector. For this purpose an additional SciFi detector has been built and placed upstream of the target (cf. 7.1.3).

4.2.3 Further Requirements

For a successful implementation of the TIGER system in the existing COMPASS environment, some general requirements have to be fulfilled. Because of the limited buffer memory of the front-end electronics, the latency of the first-level trigger has to be kept at a minimum, typically below 3 μ s. Taking the inevitable cable delays and the transmission through the TCS into account, the trigger decision has to be made within less than a microsecond. This makes high demands on the processing speed of the proton trigger system. It will become apparent in the following sections, how this challenge has been approached with the development of the TIGER system and its tight integration into the GANDALF framework.

The existing COMPASS trigger system is in large part based on standard NIM³ electronics. Therefore, the final proton trigger signal emitted from the TIGER module has to be a simple logic pulse with a fixed latency with respect to the original event. Since the trigger processing has not necessarily a constant run time – it can for example depend on the actual hit rate in the counters –, the calculated trigger time stamp has to be buffered until it is time to assert the trigger signal.

4.3 Electronics Framework

As it has already been described, the readout of the CAMERA detector is carried out with 12 GANDALF transient analyzer modules, which digitize the 96 PMT channels and extract pulse features in real-time using a dCFD algorithm. Thereby each PMT pulse is characterized by a set of three parameters: a time stamp, an amplitude and an integral. This data is subsequently used for two different tasks. On the one hand, the pulse information is buffered until it is eventually transmitted to the DAQ system upon reception of a trigger – this is the classical function of a detector readout module. But on the other hand, the extracted pulse parameters are instantly transmitted to the TIGER trigger processor, where the information from all detector channels are centralized, allowing to generate a trigger signal based on the afore-mentioned conditions.

Fig. 4.8 illustrates the application of the TIGER modules for the CAMERA readout and proton trigger. The PMT signals of the CAMERA detector as well as the dynode signals of the SciFi stations FI01 and FI02 are digitized by GANDALF transient analyzers. The data flow emanating from the GANDALF modules is assigned to two different sub-networks according to the specified tasks: trigger generation and detector readout. Two TIGER modules with different firmware configurations are forming the crosspoints of the respective sub-nets.

The trigger sub-network, depicted on the left-hand side of the diagram, feeds the pulse information – the so-called trigger primitives – to a first TIGER module, which is configured to act as a trigger processor. To cope with the high hit rate of up to several MHz in single detector channels and the low-latency requirement for the trigger generation, a custom protocol (cf. 4.3.2) has been designed for the transmission of the trigger primitives. The trigger processor operates on this information by means of its high-performance FPGA device in order to identify proton

³Nuclear Instrumentation Module




Figure 4.8: Application of the TIGER modules for the CAMERA readout and proton trigger.

candidates and make a trigger decision. The functionality of the FPGA firmware is described in section 6.1.2. The CAMERA trigger signal is finally included in the global COMPASS trigger system, where it is combined with trigger signals from other parts of the spectrometer.

The readout sub-network, depicted on the right-hand side of the diagram, is responsible for the collection of the detector data and its transmission to the DAQ. Instead of connecting each GANDALF readout module via a dedicated S-LINK to the DAQ, a second TIGER module is used as a readout concentrator, thus significantly reducing the number of required S-LINKs and ROB PCs. At the same time, the TIGER readout concentrator distributes the TCS signal to all GANDALF modules, providing a synchronous reference clock and the COMPASS first-level trigger. More information about the FPGA design in this configuration is given in section 6.1.3.

4.3.1 VME64x/VXS Crate

The hardware platform, which has been chosen to implement the above-described architecture, is based on the VME64x/VXS⁴ standard ANSI/VITA 41.0 [80]. It extends the parallel VME-bus [81], which is widely used in nuclear and particle physics experiments since many years,

⁴VME Switched Serial

by high-speed serial interconnections for increased data transfer rates. The VITA 41 standard provides a base specification (VXS.0) for mechanics, power and backplane connections, and it defines two types of VXS boards: payload cards and switch cards. In addition, various protocol layer specifications (VXS.1 to VXS.5) are available, which define the usage of specific serial fabrics such as InfiniBand, Ethernet or PCI Express for a communication between the modules.

A WIENER crate (model UEV 6021 with UEP 6021 power supply [82]) equipped with a Hartmann B18118233I backplane [83] provides 18 payload card slots, 2 switch card slots and 1 legacy VME64x slot (see Fig. 4.9). It can therefore house all the GANDALF modules needed for the CAMERA and SciFi dynode readout, as well as two TIGER modules and a VMEbus CPU board (MEN A20 [84]). The A20 CPU board is connected via Ethernet to the DAQ control network and serves as a bridge to the VMEbus for configuration and monitoring of the GANDALF modules. Since the switch card slots are not part of the VMEbus, the TIGER modules have to be connected separately to the DAQ control network, as described in section 5.2.

The VXS part of the backplane is designed in a dual star configuration, connecting each of the payload cards separately to the two switch cards by direct point-to-point links. According to the VXS protocol layer specifications, these links are defined as four bidirectional high-speed serial interconnections, i.e. eight differential pair connections between each payload and switch card. Due to the specific requirements of the data transmission between the GANDALF (G) and TIGER (T) modules, it was decided to use only the physical layer but not the protocol layer of the VXS standard. Hence the direction of the signal transmission could be defined as shown in Fig. 4.9 to accommodate the data flow, which is mainly unidirectional. For the connection to the trigger processor, all eight lanes are used in the direction $G \rightarrow T$, whereas the connection to the readout concentrator is sub-divided into six lanes for event data transmission ($G \rightarrow T$) and two lanes for the TCS distribution ($T \rightarrow G$).

4.3.2 Backplane Link

The challenge while implementing the custom transmission protocol for the VXS backplane was to find a trade-off between the different requirements that arose from the task to send the trigger primitives from the GANDALF modules to the TIGER trigger processor. The most critical requirement is a predictable latency, because hit information which are not received in time, cannot be considered for the trigger generation anymore. In addition, the different detector channels should not influence each other in terms of transmission rate and latency, to avoid that for example one noisy channel can saturate the link and suppress hits from other channels. Fortunately the number of available transmission lanes of the VXS link to the trigger processor just equals the number of digitization channels of the GANDALF transient analyzer, which allowed a one-to-one mapping.

The pulse processing algorithm in the GANDALF module is continuously looking for hits on the input channels and calculating its parameters with high resolution [74, p.88]. However, for the trigger processing the pulse feature information is not required with full precision. Simulations have been performed [77, p.41] to study the influence of limited accuracy of the hit information on the proton identification capabilities. Thereupon, the precision of the trigger primitives was defined as follows, to optimize the transmission rate over the VXS backplane:




Figure 4.9: The VME64x/VXS backplane with dual star configuration has a VXS point-to-point connection from every payload board slot to each of the two switch board slots. The direction of transmission for the 8 differential pair signals, which is deviating from the definition in the VITA 41 standard, is indicated by the wide arrows. There are 8 TX lanes from every GANDALF slot to the TIGER slot B for transmission of the trigger primitives. The connection to TIGER slot A implements 6 TX lanes for event data transmission and 2 RX lanes for TCS data reception. Slot number 1 is reserved for the VMEbus CPU. The diagram is based on [85, p.41].

- a 9-bit amplitude, by discarding the 3 least significant bits of the ADC digit,
- a 10-bit coarse time, which provides a dynamic range for the time measurement of 1024 ns,
- a 6-bit high-res time, sub-dividing the 1 ns time unit in 64 bins.

The VXS backplane link, which has been developed for the GANDALF framework in [86], is operating at 500 MHz DDR⁵, thus providing a throughput of 1 Gbit/s for each channel. It is implemented using a 10:1 SerDes⁶ logic, which is self-calibrating after establishing the connection to compensate for different trace lengths on the backplane and a possible offset between the transmitter and the receiver clock. Once the link is synchronized it is long-term stable and the bit error rate has been determined to be below $6.7 \cdot 10^{-15}$ with a confidence level of 99% [86, p.93]. The format of the trigger primitive data packet is shown in Tab. 4.1. For every detected

⁵Double Data Rate

⁶Serializer/Deserializer

Table 4.1: Format of the trigger primitive data packet for transmission via the VXS backplane. P_0 and P_1 are parity bits to allow error detection.

Bit	9	8	7	6	5	4	3	2	1	0
Word 1	1	amplitude (9)								
Word 2	coarse time (10)									
Word 3	high-res time (6)						0	0	P ₁	P ₀

hit three 10-bit data words are transmitted, which takes 30 ns. Therefore the mean hit rate on a detector channel may be up to 33 MHz.

The amount of data which is transmitted for every trigger primitive has an immediate effect on both the acceptable mean hit rate and the latency increase in case of pile-up pulses. When a hit is recognized by the dCFD algorithm while the previous one is still being transmitted, the new hit has to be buffered. In an extreme scenario with several near following pulses on a detector channel the latency is significantly increased. Therefore the buffer size is limited to six data packets and further hits are not considered for transmission to the trigger TIGER.

A detailed sketch of the entire data flow that is involved in the CAMERA trigger generation, is shown in Fig. 4.10. The variable t indicates the accumulated latencies for the various stages of the transmission path, starting at $t = 0$ with the primary event in the target. After a delay of at most 150 ns including the time of flight of the recoil particle, light propagation in the scintillator and cable delays, the signal arrives at the GANDALF ADC module. After an additional delay of 80 ns the hit parameter calculation of the dCFD algorithm is finished and the trigger primitive (TP) is available in the TP buffer. Depending on the fill level some additional latency is caused by the buffering. For the exemplary values in this diagram, two immediately preceding hits are assumed. The subsequent transmission to the TIGER module via the VXS link takes 50 ns and another 30 ns later the trigger primitive has been written to the TP memory, where it can be accessed by the trigger logic.

The proton trigger signal, which is generated by the TIGER module, is transmitted to the trigger barrack at the far end of the experimental hall. It has to arrive at the COMPASS trigger system within 1200 ns after the primary event. Taking the propagation delay of the cable to the TCS of roughly 330 ns into account the time available for the trigger logic is approximately 510 ns.

In order to meet the hard real-time requirement for the proton trigger generation, the TIGER board features powerful programmable logic devices. Its development is detailed in the following chapter. The trigger logic is implemented in a custom firmware, which is described in section 6.1.2.




Figure 4.10: Diagram of the trigger primitive (TP) transmission path. The entire data flow of the hit information from the primary event (at time $t = 0$) to the final trigger signal is sketched and exemplary points in time t are given, indicating when the information is available at the respective stages.

5. The TIGER Module

The TIGER module is designed to complement the readout electronics of the CAMERA detector. Hence it has to fit into the existing GANDALF framework, which has recently been developed for the COMPASS-II experiment. As described in section 4.3, the GANDALF framework is based on the VXS (ANSI/VITA 41) standard, which combines the well-proven VMEbus with a high speed point-to-point connection on a single backplane.

While the GANDALF transient analyzers are VXS payload boards, the TIGER module is operated in the VXS crate's switch slots. And although it is designed with the form factor of a VXS switch board, the TIGER module is a good deal more than just a switch. It provides immense computing power for a sophisticated trigger generation, large memories to buffer detector readout data and several interfaces to the existing COMPASS DAQ equipment.

Fig. 5.1 shows the main functional blocks of the TIGER module. The centerpiece of the board is a large field-programmable gate array (FPGA), which carries out the actual high-performance data processing. It also controls most of the high-speed interfaces and memory devices. The FPGA is accompanied by a CPU, which runs a standard Linux operating system (OS) and provides the necessary interfaces for slow-control tasks like configuration and monitoring of the TIGER system. Optionally, a GPU¹ can be added to the system to assist with complex calculations. Both CPU and GPU are located on exchangeable add-on boards complying with the COM Express and the MXM² standards respectively (see 5.2). They are connected to the FPGA with PCI Express (PCIe) links. A CPLD³ is responsible for the power-up and power-down sequences of the board, the FPGA configuration process and the handshaking process between the TIGER and the GANDALF modules in the VXS crate. It also implements the necessary glue logic to connect different parts of the board. All hardware components and interfaces will be described in-depth in the following sections.

¹Graphics Processing Unit

²Mobile PCI Express Module

³Complex Programmable Logic Device




Figure 5.1: Block diagram of the TIGER module, showing the most relevant components and interfaces. All items are detailed in the text.

5.1 Mainboard

The form factor for a 6U-VXS switch board is defined in the VITA 41.0 standard [80]. The TIGER mainboard complies mechanically and electrically with this specification. The dimensions of the printed circuit board (PCB) are $233 \times 160 \text{ mm}^2$. It is equipped with a IEEE 1101.10 compliant [87] EMC⁴ front panel, electrostatic discharge strips and injector/ejector handles. The handles facilitate a safe plug/unplug procedure by providing the necessary insertion force, which can be more than 500 N due to the large amount of backplane connectors. At the same time they control the hot-swap process of the board (cf. 5.1.6.1), ensuring that the main power is shut down before the board is pulled out of the crate. The TIGER module is shown in Fig. 5.2 with labeling for the major components and for the interface connectors. The heat sinks have been removed for this picture to show the devices underneath.

5.1.1 Virtex-6 FPGA

The central component of the TIGER module is a Virtex-6 FPGA from Xilinx [88]. FPGAs are integrated circuit devices which can be configured flexibly to perform complex digital computations. They are field-programmable, which means they can be reconfigured in-system at any time. Since the end of the 1990s, FPGAs are widely used in trigger and data acquisition systems for HEP experiments [89], providing certain advantages over traditional ASIC⁵ devices. For example, the re-programmability allows to fix design problems or to extend the functionality of the design. In addition, due to the strongly increased cost and development time for ASIC productions using the current manufacturing processes, FPGAs provide a better unit cost for low and medium quantities.

⁴Electromagnetic Compatibility

⁵Application-Specific Integrated Circuit




Figure 5.2: Photo of the TIGER module and of its front panel. The FPGA and its related components and interfaces are indicated with blue labels. The CPU with the corresponding interfaces is labeled in orange color. Areas related to the power supply are marked with green borders. All components and interfaces are detailed in the chapter at hand.

While the first devices mainly consisted of programmable logic gates with programmable interconnects, new features were added in the following device generations. Nowadays FPGA devices are much more than just "programmable gate arrays" – in fact they contain complex embedded function blocks for digital signal processing (DSP), clock management or fast serial communication, to mention just a few examples. Recently a new device family [90] based on a System-on-Chip architecture has been introduced combining an ARM⁶ multi-core processing system with traditional programmable logic on the same chip.

The FPGA device which was chosen for the TIGER module is the XC6VSX315T with speed grade -2 in the FF1759 package. Its most important functional elements are briefly described in the following subsections and the key features are summarized in Tab. 5.1. The development of the TIGER FPGA firmware, which implements the specific behavior of the board, is described in chapter 6.

5.1.1.1 Configurable Logic

The unit cell of the Virtex-6 programmable logic is the so-called slice, which consists of four function generators (FG), eight storage elements, some multiplexers and arithmetic carry logic. A simplified schematic without the carry logic is shown in Fig. 5.3. Each FG can be configured either as 6-input look-up table (LUT) or as a 5-input LUT with two separate outputs. Using F7AMUX, F7BMUX and F8MUX it is possible to combine up to four LUTs of a slice to implement functions of 7 or 8 inputs as well as large multiplexers up to 16:1.

The storage elements serve as edge-triggered D-type flip-flops (FF) to register either the LUT outputs (O5, O6) or the bypass slice inputs (AX..DX). Alternatively four of the storage elements can be configured as level-sensitive latches (LAT), but in this case the remaining four (FF*) cannot be used. All storage elements in the slice are controlled by the common clock (CLK), clock enable (CE) and set/reset (SR) signals. The slice provides four asynchronous outputs (A..D), four registered outputs (AQ..DQ) and four multiplex outputs (AMUX..DMUX). The specific function of a slice is defined by the content of the LUTs and the signal routing through the multiplexers, which is selected during the device configuration.

When implementing arithmetic adders and subtractors, each one-bit addition is depending on the carry bit from the next less significant stage. This poses a problem for the fast addition of wide binary numbers, since the carry bit has to propagate through all stages until it is available in the most significant bit position. Carry-lookahead adders overcome this limitation by pre-calculating whether a stage is going to propagate an incoming carry bit and using fast dedicated paths to transmit the carry information. Each slice contains a four bit carry chain made of multiplexers and XOR gates running upward in the FPGA, with direct carry-in and carry-out connections to the slices above and below. The other inputs and outputs of the slices are connected via a switch matrix to the general routing matrix of the FPGA (Fig. 5.4).

The next higher hierarchical unit in the Virtex architecture is the configurable logic block (CLB), containing two slices. In roughly half of the CLBs one of the slices is of the *SLICEM* type. The LUTs in these slices can optionally be used as distributed 64-bit RAM or as 32-bit shift registers.

⁶Advanced RISC Machine




Figure 5.3: Simplified schematic of a Virtex-6 slice. The carry logic is not shown. [91, p.35]

The remaining slices are of the standard *SLICEL* type. The CLBs are arranged in a grid spanning the entire FPGA (Fig. 5.5), apart from an area in the center which holds specialized elements dedicated to e.g. clocking or device configuration. The array of CLBs is also intercepted regularly by columns of block memory and DSP elements.

5.1.1.2 Clock Management

The individual parts of the TIGER FPGA design are clocked with several different frequencies. Some of them are predetermined e.g. by interface specifications, while others may be tuned for maximum performance or to match the required throughput. To provide all the necessary clock frequencies throughout the device, the firmware designer can make use of the mixed-




Figure 5.4: Left: a configurable logic block (CLB) consists of two slices, which are connected to a switch matrix. All slices of the same column are linked vertically by a carry chain (CIN and COUT). Right: the switch matrices have access to the general routing matrix to establish connections between the CLBs.




Figure 5.5: The FPGA floor plan shows the distribution of the CLBs across the chip area, as well as the placement of I/O, RAM, DSP, clocking and high-speed communication elements. Details on these elements are given in the text.




Figure 5.6: The Virtex-6 mixed-mode clock manager is based on a phase-locked loop to maintain a stable phase and frequency relationship of the output clocks to the reference input. The phase-frequency detector (PFD) compares the feedback clock to the input clock and in case of deviations it generates up/down signals, which are converted by the charge pump (CP) and loop filter (LF) into a reference voltage for the voltage-controlled oscillator (VCO). Adapted from [92, p.39].

mode clock managers (MMCMs) which are available in the FPGA. These PLL⁷-based frequency synthesizers are capable of generating eight clock signals with different frequencies and programmable phase offsets from a given reference clock. The principal of operation is sketched in Fig. 5.6 – the output frequencies are defined by selecting a multiplier (M), an input divider (D) and output dividers (O_i) according to the formula

$$f_{\text{out},i} = f_{\text{in}} \frac{M}{D \cdot O_i}. \quad (5.1)$$

Since the distribution of the clock signals is crucial for high-speed FPGA designs, there are a number of dedicated clock lines available. Clocks which are required all over the design are transmitted via global clock lines providing the highest fan-out. Regional clocks can be used to drive logic which is located within a confined area. A clock region is 40 CLB high and half the chip wide. I/O clocks are reserved for use with I/O logic like DDR buffers and SerDes.

5.1.1.3 Block RAM

All data processing algorithms need storage elements to retain data between the clock cycles. The registers and distributed RAM in the slices can be used to store small amounts of data, whereas dedicated Block RAM (BRAM) elements are available for larger data sets. The basic BRAM element (Fig. 5.7) of the Virtex-6 is a 36 kbit dual-port memory with selectable data width⁸, which can be set to different values for port A and port B. The data outputs can make use of an optional pipeline register, which adds one clock cycle of latency but on the other hand allows for higher operation frequencies. Read and write access is possible via both ports using independent clocks, which turns the BRAM element into an ideal data buffer between different clock domains. For first-in/first-out (FIFO) applications the BRAM elements contain dedicated

⁷Phase-Locked Loop

⁸Available values: 1, 2, 4, 9, 18, or 36 bit. For a data width smaller than 9, only 32 kbit of the memory can be used.

control logic for the address counter and status flag generation like *FIFO full* or *FIFO empty*. Either a synchronous or an asynchronous –i.e. dual-clock– FIFO can be implemented using one or several cascaded BRAM elements, depending on the required depth.

FIFO elements and block memories are extensively used in the TIGER design to buffer the PMT, hit and track information between the individual stages of the trigger processing (cf. 6.1.2). This is one of the reasons for the choice of an SXT FPGA, which has considerably more BRAM elements than the LXT models. Another application for BRAM elements is the implementation of read-only memory to store constants or look-up tables for the use by DSP algorithms, e.g. the energy cuts according to section 4.2.2.3 for the proton selection of the trigger.

5.1.1.4 DSP Slices

For fast digital signal processing applications the Virtex-6 FPGA provides dedicated DSP48E1 slices. A simplified diagram of the slice content is given in Fig. 5.8. The central element is a 25-bit $\times$ 18-bit multiplier with a pre-adder on one input, followed by data path multiplexers (X,Y,Z) and a 48-bit 3-input adder/subtractor. The output of the adder/subtractor can be fed back via the Z multiplexer to act as an accumulator. The specific operation of the DSP slice is dynamically selected by control signals (*OPMODE*, *ALUMODE*). The adder/subtractor can also be configured as a logic unit to implement bit-wise functions or pattern detection.

Within each input path and behind the multiplier there are additional pipeline registers, which can be used to increase the maximum possible clock frequency at the expense of some latency. They are denoted with the red rectangles in Fig. 5.8. DSP algorithms like finite impulse response (FIR) filters with several stages are implemented by cascading multiple DSP48E1 slices in a column. For this purpose cascable input and output streams are available between the pipeline registers of adjacent DSP slices to avoid the use of general logic and routing resources.

5.1.1.5 Inputs and Outputs

Most of the I/O pins of the Virtex-6 FPGA are configurable by the designer to match the electrical signal standard of the interfacing components. These are the so-called User-I/O pins – in contrast to the power and ground pins and a few configuration pins, which serve a dedicated purpose. A large number of both single-ended and differential standards is supported. The pins are grouped in banks of 40 I/Os each with additional dedicated V_{CCO} supply voltage pins to power the output buffers. Since different standards may require different supply voltages, this poses some constraints for the combination of multiple signal standards in one bank.

The FPGA on the TIGER module uses the following signal standards on the User-I/O pins:

- LVDS⁹ (2.5 V) for clock signals, for all signals connected to the VXS backplane and for the differential I/O connector on the front panel (see 5.3),
- LVCMOS¹⁰ (2.5 V) for the interface to the CPLD (see 6.1.1.1) and for some general purpose control signals,

⁹Low Voltage Differential Signaling

¹⁰Low Voltage Complementary Metal Oxide Semiconductor




Figure 5.7: Interface signals of the Virtex-6 dual-port Block RAM element. [93, p.14]




Figure 5.8: Simplified diagram of the DSP48E1 slice, adapted from [94, p.18]. The red boxes denote optional pipeline registers.

- SSTL¹¹ / Differential SSTL (1.5 V) for the DDR3 memory interface (see 5.1.2).

Differential LVDS inputs are internally terminated with 100Ω resistors. This saves a large number of discrete components which would have otherwise to be placed very close to the FPGA package. Similarly the bidirectional SSTL pins make use of the internal split-thevenin termination provided by the Virtex-6 DCI¹² feature. This allows to properly terminate the signal trace when the I/O buffer is tri-stated, i.e. used as input, while in output mode the termination resistors are disabled.

During the design phase of a high pin-count FPGA board a careful floor and pin planning is mandatory. Floor planning involves the detailed analysis of the data streams in the FPGA to allow for a wise placement of the logical blocks on the chip area, which in turn supports the routing process during the firmware implementation to meet the timing constraints and reach a higher performance. This procedure is accompanied by the pin planning, since the connection of data buses to certain package pins naturally defines where these data enter or leave the silicon. In practice the result of these two planning processes is always a trade-off between optimal routability of the signal traces on the PCB (see also 5.1.7) and a reasonably arranged data flow through the FPGA logic. The package pin plan for the TIGER FPGA is shown in Fig. 5.9 and the pin-out is tabulated in appendix B.

Located directly behind the I/O drivers and receivers there are dedicated logic resources to assist the interfacing with high-speed signals. Amongst others these are DDR buffers, IODELAY elements and SerDes logic blocks, which are employed for the VXS connection between the GANDALF and TIGER modules as described in section 4.3.2.

¹¹Stub-Series Terminated Logic

¹²Digitally Controlled Impedance




Figure 5.9: Package pin plan for the TIGER FPGA. The association of the pins to the various interfaces is indicated by the color code.

5.1.1.6 Gigabit Transceivers

For fast serial data transmission the Virtex-6 FPGA is equipped with GTX transceivers, i.e. pairs of transmitter and receiver, capable of data rates up to 6.6 Gbit/s. Transmitter and receiver are basically a parallel-to-serial converter and a serial-to-parallel converter respectively, with a programmable conversion ratio between 8 and 40. The serial side is connected to CML¹³

¹³Current Mode Logic

Table 5.1: Key features of the Xilinx Virtex-6 FPGA XC6VSX315T-2 FF1759.

CLB slices	49200
SLICEMs / SLICELs	20360 / 28840
Flip-Flops	393600
Distributed RAM	5090 kbit
Block RAM	$704 \times 36 \text{ kbit} \approx 25 \text{ Mbit}$
DSP48E1 slices	1344
$f_{\max}$ of BRAM / DSP	540 MHz
GTX transceivers	24
User I/O pins	720
Package pins	1759
Package size	$42.5 \times 42.5 \text{ mm}^2$

compliant drivers/buffers, which in turn are connected to dedicated transceiver pins at one edge of the device package. The parallel side is interfacing the FPGA logic fabric. That way the data processing in the internal logic can take place at much lower clock frequencies than the serial transmission. A wide range of industry-standard high-speed serial protocols is supported by the GTX transceivers, i.e. Infiniband, Aurora, PCI Express, S-ATA, Gbit-Ethernet and many more, but of course custom protocols like the CERN-developed S-LINK can also be implemented (see 6.1.3). The built-in 8b/10b and 64b/66b algorithms can optionally be used to manipulate the data according to these line codes to achieve DC balancing and allow clock recovery from the serial bit stream by the receiver.

The GTX transceiver tiles contain internal PLLs to generate the high-speed bit clock and the divided word clock, offering dedicated reference clock inputs for best possible performance. The TX output buffer features a programmable differential voltage swing and a pre-emphasis driver to maximize signal integrity. The receiver contains equalization circuits to compensate for high-frequency losses in the channel. Both TX and RX settings must be optimized to accommodate the transceivers to the specific channel characteristics. This is especially relevant for the higher data rates, in order to reach a satisfying bit error ratio.

On the TIGER module GTX transceivers are utilized for PCI Express connections between the FPGA and the CPU and GPU extension boards (see 5.2), for a high-speed link to the adjacent TIGER module via the VXS backplane (5.3.2) and for optical fiber connections to the DAQ system (5.3.3).

5.1.2 DDR3 Memory

A SO-DIMM socket for a DDR3 memory module is directly connected to the Virtex-6 FPGA on the TIGER board, providing a temporary storage capability, for example to buffer detector data from the readout modules, which may incur at peak rates during the on-spill period. To gain access to the DDR3 memory in the FPGA design, an interface is created using the Xilinx *Memory*

Interface Generator (MIG) wizard [95]. This graphical tool generates IP core files containing the controller and physical layer, which are necessary to support external SDRAM modules. Detailed information about the specific memory module being used is requested during the generation process. This includes the module organization, i.e. the row, column and bank address width, and the SDRAM timings. This information is usually found in the datasheet of the memory module, but may also be read out from the SPD¹⁴ EEPROM.

A 2 GB SO-DIMM DDR3 module from Hynix¹⁵ has been selected for the application on the TIGER board. It is constructed from eight 256M × 8bit components and is a 64-bit wide, single-rank module. Supported timings are DDR3-1066 (533 MHz clock) with CL7 and DDR3-1333 (667 MHz clock) with CL9 [96]. The CL value specifies the column address strobe (CAS) latency in clock cycles between the assertion of the CAS signal and the availability of the requested data. A detailed description of the DDR3 device operation may be found in [97]. For access to the Hynix DDR3 memory module a MIG core has been created according to the FPGA pin-out, as specified in appendix B. Several design rules for the I/O bank selection and pin allocation [95, p.135] had to be considered during the FPGA pin planning, to satisfy the timing requirements of the memory interface. Details about the implementation of the MIG core in the TIGER firmware are given in section 6.1.1.6.

5.1.3 Clock Distribution Network

The TIGER module requires a number of different clock signals for operation. While some frequencies are defined by the various interface standards, others have to be synchronized to the COMPASS experiment's reference clock. But all have in common the requirement for low-jitter clock signals for best possible performance. The clock generation and distribution network shown in Fig. 5.10 has been carefully designed to meet these demands.

5.1.3.1 PCI Express Reference Clock

The PCIe data transfer is based on high-speed serial point-to-point links using 8b/10b encoding. More details about the PCI Express architecture can be found in 5.2.6. No clock signal is transmitted on the link; instead the receiver recovers the RX clock from the bit transitions of the data stream using a PLL. In addition, a 100 MHz reference clock is supplied to each PCIe device to allow for the generation of all required internal clocks.

The PCIe hierarchy on the TIGER board is quite simple. It contains the CPU which forms the root complex, the FPGA acting as endpoint/switch and finally the GPU as another endpoint. Hence the reference clock which is provided by the COM Express CPU module has to be distributed to the two other members of the PCIe structure. This is done with a 1:2 fan-out buffer ICS9DB102 from IDT [98], a PLL-based zero-delay buffer compliant to PCIe Gen1 and Gen2 clocking requirements. The PLL bandwidth is selectable (0.5 MHz or 2.5 MHz) to minimize jitter peaking in the PLLs of the downstream devices. The HCSL¹⁶ outputs are connected via AC-coupling capacitors to the reference clock inputs of the MXM socket and the FPGA's GTX

¹⁴Serial Presence Detect

¹⁵Part number: HMT325S6BFR8C-H9

¹⁶High Speed Current Steering Logic




Figure 5.10: Clock generation and distribution network of the TIGER module. A detailed description is given in the text.

transceiver tiles. The ICS9DB102 clock buffer starts up together with the CPU and the PLL locks within 2 ms after the reference clock is started. The outputs are enabled only when the downstream devices request a clock signal via independent CLKREQ# pins.

5.1.3.2 Experiment Synchronous Clocks

The COMPASS trigger information is distributed experiment-wide by the TCS system via optical fibers. The transmitted signal is a serial data stream with a bit rate of four times the COMPASS reference clock

$$f_{\text{TCS}} = 4 \cdot 38.88 \text{ MHz} = 155.52 \text{ MHz}. \quad (5.2)$$

The optical TCS signal is received on the TIGER board by the *Gimli* card (cf. 5.3.4), where it is converted back into an electrical signal and the TCS clock is reconstructed with a CLC016 clock recovery chip. Thereupon the clock signal is split with a 1:2 LVDS fan-out buffer NB6N11S from ON Semiconductor [99]. One copy of the TCS clock is directed to the FPGA, together with the TCS data from the *Gimli* card, while the second copy is driving the input of a Silicon Labs Si5326A [100] clock multiplier chip.

The Si5326A device (Fig. 5.11) generates two programmable clocks by means of a digitally controlled PLL with programmable loop filter bandwidth. A reference crystal is used to actively attenuate the jitter of the input clock resulting in a sub-picosecond jitter performance. The clock multiplication ratio is defined by four divider values, according to the following formula:

$$f_{\text{out},i} = f_{\text{in}} \frac{N2}{N3 \cdot N1_{\text{HS}} \cdot N1_{\text{LS},i}}. \quad (5.3)$$

All settings are programmed through an I²C interface via the CPU using the software tool described in section 6.3.




Figure 5.11: The Si5326A clock multiplier chip contains a DSPLL, which implements a digital loop filter (dLF) based on a DSP algorithm to regulate the digitally-controlled oscillator (DCO). The phase detector (PD) compares the input and feedback clocks. Adapted from [101].

Altogether the TIGER Virtex-6 FPGA receives three clock signals which are related to the global COMPASS reference clock, in particular the 155.52 MHz TCS clock, which is used to register the TCS data input and to perform TCS-synchronous tasks like the reset of certain counters at the begin of spill (see 6.1.2), as well as two derived clocks with a fixed frequency and phase relationship with respect to the TCS clock. A special procedure called *Si sweep* (see 6.1.1.3) is performed to align the phase of the derived clocks with the reference clock. In practice, one of the Si5326A clock outputs is programmed identically to the ADC sampling clock on the GANDALF modules (505.44 MHz) to obtain the time unit, which is necessary to interpret the time stamps of the digitized PMT pulses that are transferred via the VXS bus. For this reason the TCS-related part of the TIGER clock network is an exact replication of the correspondent GANDALF part.

5.1.3.3 Free-Running Clocks

A Silicon Labs Si5338 [102] clock generator chip is used on the TIGER module to produce four additional clock signals, which are made available to the FPGA. This I²C-programmable device is capable of generating arbitrary frequencies up to 350 MHz either locked to an external reference clock or in free-running mode. Its two-stage synthesis architecture consists of a high-frequency VCO¹⁷ (2.2 to 2.84 GHz) and a PLL with fractional-N feedback and output dividers. From the wide range of supported single-ended and differential output signal standards, LVDS was chosen for the connection to the Virtex-6 FPGA.

On the TIGER board the Si5338 is operated in combination with a 25 MHz reference crystal in free-running mode to generate the frequencies listed in table 5.2. The DDR3 clock and the User clock are fed into the FPGA's logic fabric, while the Aurora and SFP transceiver clock signals are connected via AC-coupling capacitors to the reference clock inputs of the GTX transceiver tiles, which implement the corresponding interfaces. In case of future changes to these interfaces, for example the adoption of a different transmission protocol for the SFP link, a new frequency plan can be calculated using the ClockBuilder software [103]. The resultant register file is then loaded to the device via the I²C programming tool as described in section 6.3. Finally, a separate 40 MHz quartz oscillator is used for an independent clocking of the CPLD.

¹⁷Voltage-Controlled Oscillator

Table 5.2: Clock signals produced with the Si5338 clock generator chip and connected to the TIGER FPGA.

Output #	Purpose	Frequency	Connected to
0	Aurora transceiver clock	125 MHz	GTX ref. clock input
1	User reference clock	200 MHz	FPGA global clock input
2	DDR3 clock	266 MHz	FPGA global clock input
3	SFP transceiver clock	125 MHz	GTX ref. clock input

5.1.4 CoolRunner-II CPLD

The Xilinx CoolRunner-II is a programmable logic device, which consists of up to 32 function blocks (FB) connected via an advanced interconnect matrix (AIM) [104]. The logic in the FB is implemented within a sea-of-gates using the disjunctive normal form ("sum-of-products"). Each FB contains a programmable logic array (PLA) which forms 56 product terms (AND relations) of up to 40 inputs, followed by 16 macrocells, each capable to provide a conjunction (OR gate) of any combination of the product terms. Additionally, each macrocell contains a storage element which can be configured as either D or T flip-flop or transparent latch.

The model XC2C512-7FTG256C used on the TIGER board features 32 function blocks with a total of 512 macrocells and 212 I/O pins in four banks. Compared to an FPGA, a CPLD is rather limited in terms of logic capabilities and speed, but it has the advantage of internal non-volatile (NV) configuration memory allowing it to operate directly after power-up. The time needed for reading the configuration from NV memory during start-up is below 400 μ s [105], and for firmware updates the NV memory is in-system programmable using a JTAG interface.

A CPLD is ideally suited for many low-level board management tasks due to its "instant-on" feature. On the TIGER module it is controlling several vital processes both during start-up and during operation. Due to its tight connection with multiple onboard devices and its diverse purposes, it is reasonable not to detail all the CPLD functions at this point but rather list references to the sections where these functions will be described in conjunction with the respective components:

- power-up and power-down sequencing – the various power converters are enabled and disabled in a safe manner, triggered by the switch handles on the front panel (see 5.1.6),
- FPGA configuration – after switching on the module, the FPGA is initialized with the application specific firmware using one of the methods described in 5.1.5,
- VXS handshake – for safe hot-plugging of the GANDALF and TIGER modules in the VXS crate, a handshaking protocol (see 5.3.1.2) has been introduced which regulates the activation of the VXS bus output drivers.

To enable communication between the CPU and the CPLD, a LPC¹⁸ bus interface (cf. 5.2.3.2) has been implemented in the CPLD firmware, which acts as a LPC-to-parallel bridge support-

¹⁸Low Pin Count




Figure 5.12: Schematic overview of the FPGA configuration interfaces on the TIGER board.

ing *I/O Read* and *I/O Write* cycles. The available registers are listed in section 6.2.2, together with a description of the driver which is provided for a convenient access from the Linux OS.

In addition, the CPLD implements so-called glue logic, i.e. simple logic for joining different parts of the board design. By providing I/O banks with different supply voltages (2.5 V and 3.3 V) it can convert the signal levels between the FPGA, which supports a maximum I/O voltage of 2.5 V, and components with 3.3 V LVTTTL or LVCMOS interfaces. During commissioning of a newly developed module it is also advantageous that the connections may be changed by firmware upgrades in case of unforeseen problems or requested changes in behavior, which would not be possible with hard-wired connections. Finally, all the front panel LEDs and the debug LEDs are driven by the CPLD.

5.1.5 Configuration Scheme

The customized behavior of the Virtex-6 FPGA is defined by the firmware design (see also 6.1). It is specified using a hardware description language (HDL) and implemented using the Xilinx ISE/PlanAhead toolchain [106], resulting in a bitstream file which can be used for the device configuration. The bitstream for the SX315T device has a size of 104.5 Mbit and it contains all the information about the configuration of the slices, the routing and the initial values of the registers and block RAM elements. Since in many designs most of the block RAMs are initialized with '0's and not all FPGA elements are actually used, it is beneficial to enable the so-called bitstream compression in the implementation process to reduce the file size by a factor approx. between 1.5 and 3.

The configuration is written to the FPGA's internal SRAM, which is a volatile memory and must therefore be reloaded after every power-up. There are different configuration options available on the TIGER board – the parts which are involved in the configuration process are shown in Fig. 5.12.

5.1.5.1 Configuration at Power-Up

The Xilinx Platform Flash XL (briefly "Flash" in the following) is a 128 Mbit non-volatile NOR flash-based storage device [107]. It has been developed as a dedicated configuration solution for large Virtex-5 and Virtex-6 FPGAs, providing bitstream transfer rates of up to 800 Mbit/s in




Figure 5.13: Initial FPGA configuration from Platform Flash XL after power-up is controlled by the CPLD.




Figure 5.14: Online full reconfiguration of the FPGA is executed by the CPU using the LPC bus to the CPLD.

burst read mode. The device features a standard NOR-flash interface, which is an interface with full address and data buses for byte-addressable random access operations in contrast to the paged access to NAND flash.

The Flash holds a basic configuration file for the FPGA, the so-called *TIGER Base design* (see 6.1.1). This design contains only the static parts which are available in all FPGA user designs, like the PCIe interface, the I/O structures and the clocking resources. Immediately after the board has been powered up, the CPLD initializes the download of the configuration bitstream from the Flash to the Virtex-6 via the 16-bit SelectMap interface (Fig. 5.13). By using a 40 MHz configuration clock (CCLK) the bitstream transfer takes only 50 ms. Up to 1 ms of additional time is required to allow the FPGA start-up sequence to finish, which includes waiting for DCI matching and MMCM locking. The end of the start-up sequence is indicated by the FPGA's DONE signal. It is necessary to perform the initial FPGA configuration before the COM Express CPU starts to boot, so that the PCIe endpoint in the FPGA is ready when the OS is enumerating the buses and devices (cf. 5.2.6). Hence the CPU is held in the reset state until the configuration is finished. In this way it is guaranteed that the FPGA will be discovered by the OS as a PCIe device.

5.1.5.2 Online Configuration

After the OS has completed its boot process, the FPGA can be reconfigured at any time from the CPU. The necessary software tools are described in section 6.3. For a full reconfiguration, the new bitstream file is transferred via the LPC bus to the CPLD (Fig. 5.14), which forwards the 16-bit data words to the SelectMap interface of the FPGA using non-continuous data loading with a controlled CCLK [108, p.42]. In this mode the clock is halted and it toggles only once when the next word is ready to be loaded into the FPGA. Due to the limited bandwidth of the LPC bus this method is rather slow compared to the configuration from the Flash.

An alternative approach for the online configuration of the FPGA is the partial reconfiguration (PR) via PCIe (Fig. 5.15). For this purpose the Virtex-6 contains an internal configuration access port (ICAP) [108, p.74]. The ICAP provides an interface similar to the SelectMap interface, which is accessible from the FPGA fabric. It can be used to load partial bitstream files to the




Figure 5.15: On partial reconfiguration via PCIe the static part of the FPGA design remains unchanged.




Figure 5.16: The Flash can be updated either via the FPGA by means of a JTAG cable or online via the CPU.

FPGA's configuration SRAM, allowing to selectively reconfiguring a specific area of the device [109]. The remaining part of the design stays operational during the reconfiguration process. However, in order to use the PR method, the FPGA design must be divided into a static and one or more reconfigurable partitions. The static partition contains all I/Os and clock related components like the MMCMs and global clock buffers, as well as the part of the design which receives the bitstream from the CPU via the PCIe interface and forwards it to the ICAP. The development of a PR capable design for the TIGER module is described in section 6.1.1.7.

5.1.5.3 JTAG Configuration and Flash Programming

The Virtex-6 FPGA provides a JTAG¹⁹ interface according to the IEEE standard 1149.1 [110], which defines a test architecture for integrated circuits. It contains a set of registers, which are addressed via the test access port (TAP). The mandatory TAP pins are Test Data In, Test Data Out, Test Mode Select and Test Clock (TDI, TDO, TMS, TCK), allowing to combine multiple devices in a chain via the TDI/TDO pins.

The standard was originally developed to test the interconnections between ICs on the board-level as well as some basic functionalities of the ICs itself, which is known as boundary-scan. Besides this, many devices support vendor-specific instructions, such as configure and verify commands to load firmware data into FPGAs or CPLDs. For this purpose Xilinx provides programming cables (*Platform Cable USB*) and software (*Xilinx iMPACT*) [108, p.67] for convenient prototype development or debugging (see also 6.1.1.5).

Furthermore, the JTAG connection is used to program the Platform Flash XL indirectly through the FPGA, since the Flash has no own JTAG interface. For this task the *Xilinx iMPACT* software first initializes the FPGA with a proprietary firmware [107, p.10] featuring an in-system programming engine, which gains access to the Flash through its standard NOR-flash interface. Afterwards the configuration file is transferred via JTAG instructions to the programming engine, which executes the necessary write commands on the flash interface. This programming path is depicted by the red arrow labeled with ① in Fig. 5.16. An alternative path labeled with

¹⁹Joint Test Action Group




Figure 5.17: Schematic overview of the power distribution system for the TIGER module. The DC/DC converter part is detailed in Fig. 5.18.

② for sending bitstream files to the Flash is drawn from the CPU via the LPC bus to the CPLD, which can likewise access the NOR-flash interface and execute the write commands.

5.1.6 Power Distribution System

A stable and low-noise power distribution system has been designed for the TIGER module, taking into account the specific requirements of all the onboard components. The primary power input of the board is defined by the VXS standard, which specifies a 6-pin backplane connector²⁰ providing a voltage of +5 V with a current of up to 30 A. All other necessary voltages are generated by onboard power regulators.

5.1.6.1 Hot-Swap Capability

An overview of the available power rails in Fig. 5.17 shows, that the power distribution network consists of two branches, which are turned on independently by dedicated hot-swap controllers. The **standby branch** is enabled as soon as the two VME injector/ejector handles on the board's front panel are closed using a Texas Instruments (TI) TPS2420 hot-swap controller with an integrated MOSFET. The rail VCC5V0_S5 is used for the standby power supply of the CPU module, and two more rails (1.8 V and 3.3 V) are generated by low-dropout (LDO) linear regulators for powering the CPLD and Platform Flash XL devices and the controller for the switched-mode DC/DC converters (cf. 5.1.6.2). The power rails of the standby branch are indicated by the suffix "_S5", referring to the ACPI²¹ *Soft Off* power state.

The TPS2420 senses the load current and protects the system in the event of a failure by limiting the output current and shutting down the MOSFET based on a programmable fault threshold. When the load current exceeds this threshold a fault timer is started, which will eventually shut down the device if the over-current condition is still present on expiration of the timer. The

²⁰Positronic VPX6W6M400A1-B1A5A6A

²¹Advanced Configuration and Power Interface

occurrence of such a fault event is indicated by the red fault LED (LD6) on the back side of the TIGER board. It can be reset by either toggling the VME switches or power-cycling the crate. The green power-good LED (LD7) indicates that the three _S5 power rails have been ramped up successfully.

As soon as the CPLD is initialized, it executes the power-up sequence for the **main branch**. Almost all devices on the TIGER board are powered by the various voltage rails of this power distribution branch. The main 5 V rail VCC5V0 is enabled by a TI TPS2331 hot-swap controller linked to a IRF6718 MOSFET, which is rated at 61 A and features a very low on-resistance of $\approx 0.5\text{m}\Omega$ [111]. Using an external sense resistor, the TPS2331 implements a circuit breaker with a programmable overcurrent threshold, to shut down the power rail in case of a short circuit. A potential fault would be indicated by the red LED LD8. Out of the main 5 V rail, all other voltages which are needed on the TIGER board are generated by DC/DC converters.

5.1.6.2 Secondary Voltage Rails

The Virtex-6 FPGA as well as the surrounding auxiliary devices require a number of different supply voltages, which have to be generated by onboard power converters. For efficiency reasons the application of switched-mode DC/DC converters is preferred over linear regulators. Therefore a digital power management solution based on the TI UCD9246 PWM system controller [112] and four TI PTD08A PowerTrain modules [113] has been selected, which is well suited for high-performance FPGA systems. The configuration is shown in Fig. 5.18.

The PowerTrain modules are the actual conversion section of the digital power system and they are available with different current ratings to match the requirements of the specific application. The expected power consumption of the TIGER board was thoroughly evaluated based on the device datasheets and the Xilinx Power Estimator [114] spreadsheet. Since the power consumption of an FPGA largely depends on the implemented firmware design, the estimate was based on a rather busy design utilizing between 75 % and 85 % of the available logic, DSP and BRAM resources running at a clock frequency of 500 MHz. This should provide sufficient margin for even the most power consuming FPGA designs, which may be implemented for future applications of the TIGER module. In conclusion, two PowerTrain modules with current ratings of 20 A are selected for the 1.0 V and 1.5 V rails, and two modules with 10 A each are providing the 2.5 V and 3.3 V rails.

The UCD9246 device is a four-rail digital controller for DC/DC step-down converters. For each rail a set of parameters including the desired output voltage, the maximum allowed current and the behavior in case of failures is programmed into the integrated flash memory during the initial start-up of the system using an external programming cable. A PMBus²² interface is used to communicate with the UCD9246. PMBus is a variant of the I²C-based SMBus²³ protocol, specifically designed to control power supplies. A software tool (*Fusion Digital Power Designer*) is available from TI, which provides a graphical user interface to configure the device parameters and monitor the operation for debug purposes. Once a valid configuration has

²²Power Management Bus

²³System Management Bus




Figure 5.18: Generation of the secondary voltage rails by switched-mode DC/DC converters and LDO regulators.

been written to the flash memory, the device will operate according to these parameters on subsequent switch-ons.

For each of the four rails the UCD9246 integrates circuitry for DC/DC loop management, which generates a PWM²⁴ signal to control the MOSFET driver of the connected PowerTrain module, and it monitors its output voltage and compares it to the programmed reference. On deviation from the reference, the pulse width of the control signal is adjusted to compensate the error. In addition, the output current and the temperature of the power stage is constantly monitored and actions like shut-down or power limiting can be taken, as soon as programmable thresholds are exceeded. Power sequencing, soft-start/soft-stop and voltage tracking can be defined separately for every rail to satisfy the requirements of the powered devices. The start and stop sequences implemented for the TIGER board are shown in Fig. 5.19. Dependencies between rails can also be defined: for example the FPGA I/O voltage rails are immediately shut down in case of a failure on the FPGA core voltage rail, whereas the 3.3 V rail stays unaffected to keep the TIGER board responsive.

It can be seen in Fig. 5.18, that there are some additional rails in the main power distribution branch which are not controlled by the UCD9246. They will be briefly summarized in the following:

²⁴Pulse-Width Modulation




Figure 5.19: The normal ramp-up and ramp-down sequences for the four power rails, which are controlled by the UCD9246. The FPGA core voltage is ramped up first, followed by the FPGA I/O voltages, as soon as VCCINT has reached the power-good threshold. The power-down is performed in reversed order.

- A Linear Technology LTM4605 DC/DC μ Module regulator [115] is operated in boost mode to provide a 12.5 V supply for the MXM GPU card.
- A TI TPS51200 linear sink/source regulator is used to provide the termination voltage for the DDR3 SO-DIMM, and three LDO regulators (TI TPS74401) are employed to generate low-noise analog supply voltages for the Virtex-6 multi-gigabit transceiver (MGT) blocks. The linear regulators are connected to the 1.5 V rail for lowest possible power dissipation and heat production.

5.1.6.3 Power-Up/Power-Down Sequence

The CPLD is responsible for the TIGER board's start-up sequence, which involves enabling the power regulators, configuring the FPGA and booting the CPU in a coordinated manner. No less important is the shut-down sequence, before the module can be safely removed from the crate. A state diagram of the implemented start and stop behavior is shown in Fig. 5.20. The start-up process is executed, as soon as both VME handles are closed. State "s2" emulates a push of the CPU's power button, to wake it up from the ACPI soft-off state. Should one of the handles be opened again during the start-up phase, power is switched off instantly. When the state "running" is reached, the behavior of the VME handles changes insofar as opening one of them will trigger a controlled shut-down process, which is indicated by flashing the Power LED on the front panel. If the second switch is opened before the shut-down is completed, i.e. the user did not wait patiently, power is disabled instantly. In this way it is guaranteed, that all on-board power is certainly off, when the board is ejected from the VXS slot.




Figure 5.20: UML state diagram of the start-up and shut-down sequences implemented in the CPLD.

5.1.6.4 Power Distribution System Summary

All power rails available on the TIGER module are summarized in Tab. 5.3, stating the voltages, nominal currents and current limits. The nominal current value is the maximum permanent current, which the specific power converter is able to supply according to the manufacturer's datasheet. For converters with a programmable current limit a soft cap has been set based on the estimated power consumption. The selected values are given in the soft limit column. The remaining devices contain a non-programmable internal protection circuit, which will limit the current to a fixed cap to avoid damage. For these devices the hard limit column of Tab. 5.3 states the internal limit. This value is inherent to the specific device and cannot be changed from external. For power rails with subordinated LDO converters, i.e. the VCC5V0_S5 and VCC1V5 rails, the programmable current limits have to take into account also the current that is drawn by the LDO converters. Likewise, the measured currents on these rails also include the

Table 5.3: Summary of the nominal voltages and current limits of the power rails on the TIGER module.

Rail name	Voltage / V	Current / A		
	nominal	nominal	soft limit	hard limit
VCC5V0_S5	5.0		2.9	4.3
VCC1V8_S5	1.8	0.4		0.7
VCC3V3_S5	3.3	1.5		2.5
VCC5V0	5.0		33.0	
VCCINT	1.0	20.0	15.0	
VCC1V5	1.5	20.0	15.0	
VTTDDR	0.75	3.0		3.5
MGTAVCC	1.0	6.0		7.6
MGTAVTT	1.2	3.0		3.8
VCC2V5	2.5	10.0	6.0	
VCC3V3_S0	3.3	10.0	6.0	
VCC_GPU	12.5	4.0		5.0

currents through the LDO rails. The main hot-swap controller limits the current of the VCC5V0 rail to 33 A, which includes the current to the DC/DC converters as well as the load on the 5 V rail itself. The total power consumption of the board's main power branch may therefore be up to 165 W.

5.1.7 PCB Layout

The TIGER board has been developed using the *Cadence Allegro PCB Designer* [116], a powerful EDA²⁵ suite, which provides tools for all steps of the PCB design process. The main tasks during the PCB development of the TIGER module will be detailed in this section. These are:

- to capture the schematic diagram of the circuit,
- to define constraints,
- to layout the board,
- to perform simulations of the signal and power integrity.

5.1.7.1 Schematic

The schematic diagram is captured using the *Allegro System Architect*. This tool allows entering circuit diagrams using various methods. The spreadsheet method is particularly useful for defining connections between devices with a large number of pins in a tabular form. This covers for example the buses between the FPGA and the different connectors on the TIGER module. Most of these buses are direct connections, possibly with some additional discrete components involved, like resistors for termination or capacitors for AC coupling. Both differential and single-ended connections of this type are entered very efficiently using the spreadsheet editor (Fig. 5.21). Schematic drawings of the devices and connections captured with this method are created automatically for documentation purposes (Fig. 5.22). The traditional schematic editor on the other hand is used for entering power and analog blocks of the design, by placing device symbols on a canvas and drawing wires between them.

²⁵Electronic Design Automation

i30 (xc6vfx365tff1759) - U8 Functions: 28						
☐ Expand All Pins ☒ Show Differential Pairs ☒ Show Vectors						
Pin Name	Pin Number	Pin Type	Signal	Termination		
* *	*	*	*	*	*	*
DP_MGTREFCLK0_117	G10,G9	Input				
DP_MGTREFCLK1_117	E10,E9	Input	aurora_mgtclk	DPSeriesCap		
mgtrefclk1p_117	E10	Input	aurora_mgtclk+			
mgtrefclk1n_117	E9	Input	aurora_mgtclk-			
DP_MGTRX0_117	H7,H8	Input	aurora_rx<0>			
mgtrxp0_117	H7	Input	aurora_rx+<0>			
mgtrxn0_117	H8	Input	aurora_rx-<0>			
DP_MGTRX1_117	G5,G6	Input	aurora_rx<1>			
DP_MGTRX2_117	F7,F8	Input	aurora_rx<2>			
DP_MGTRX3_117	E5,E6	Input	aurora_rx<3>			
DP_MGTTX0_117	J1,J2	Output	aurora_tx<0>			
DP_MGTTX1_117	H3,H4	Output	aurora_tx<1>			
DP_MGTTX2_117	G1,G2	Output	aurora_tx<2>			
DP_MGTTX3_117	F3,F4	Output	aurora_tx<3>			

Figure 5.21: The FPGA connectivity is captured using the *Allegro System Architect* spreadsheet editor.




Figure 5.22: Auto-generated schematic drawing for the FPGA bank 117.

5.1.7.2 Constraints

Apart from the pure schematic, an extensive set of constraints has to be defined when developing a high-speed digital device. Every component and every interconnect technology has certain requirements to the physical layout of the board, and a multitude of design rules has to be observed, in order to achieve the best possible performance. The *Allegro Constraint Manager* is utilized to define so-called constraint sets (CSets). A CSet collects parameters, which are later used to control the routing of the board. Since different signal classes may have different routing requirements, references to any electrical, physical and spacing CSet can be assigned to each net in the design.

An electrical CSet defines parameters which guarantee in particular adequate signal integrity and correct timings, as required by the various interface technologies. For example, most signal standards require a specific impedance of the transmission line. For differential signals, additional parameters like the phase tolerance between the two lines of a pair and the maximum uncoupled length are defined. The TIGER board contains electrical CSets for single-ended, dif-

ferential and high-speed differential signals (see Tab. 5.4). Additional sets are defined for the buses of the DDR3 SO-DIMM interface, in order to constrain the relative propagation delay between the signals. Within each data byte bus, the signals have to match with an accuracy of 5 ps. The address and control lines are required to match with an accuracy of 20 ps.

Table 5.4: Electrical constraint sets (ECS) defined for the TIGER board.

	Single-ended ECS	Differential ECS	High-speed diff. ECS
Impedance	$(50 \pm 5)\Omega$	$(100 \pm 5)\Omega$	$(100 \pm 5)\Omega$
Static phase tolerance	–	2.5 mm	0.5 mm
Dynamic phase tolerance @ max. length	–	5.0 mm 15.0 mm	1.5 mm 5 mm
Max. uncoupled length	–	20.0 mm	7.0 mm

A physical CSet defines parameters of the copper traces, like the line width and the gap between the two lines of a differential pair (DP), in order to achieve the target impedance of the transmission lines. In addition, a so-called neck mode is specified, allowing to route nets in a more narrow way over a limited distance. This is especially useful for the fan-out region of the FPGA, where the signal density is very high (see Fig. 5.23). Last but not least, the vias²⁶ that can be used for each net class are defined in the physical CSets. For example, vias with a larger diameter may be required for power nets. The most important physical constraints of the TIGER board are tabulated in Tab. 5.5.

Table 5.5: Physical constraint sets (PCS) defined for the TIGER board.

	Single-ended PCS		Differential PCS	Power net PCS
PCB layer	inner	outer	all	all
Line width	0.115 mm	0.140 mm	0.100 mm	0.140 mm
DP gap	–	–	0.200 mm	–
Line width (neck)	0.095 mm	0.120 mm	0.090 mm	0.100 mm
DP gap (neck)	–	–	0.140 mm	–
Via hole diameter	0.200 mm	0.200 mm	0.200 mm	0.300 mm
Via pad diameter	0.500 mm	0.500 mm	0.500 mm	0.600 mm

A spacing CSet finally defines the minimum distances between the different object types (lines, pins, shapes, vias, holes, ...) on the PCB. For each combination a separate constraint is defined, i.e. line-to-line, line-to-pin, pin-to-pin, and so on. The minimal allowed values are specified by the PCB manufacturer, but susceptibility to crosstalk must also be taken into account when selecting these values.

5.1.7.3 Layout

The netlist and the constraints from the previous two steps are exported into a board file for use with the *Allegro PCB Editor*. Before the actual layout process can start, the board cross section has to be defined. The TIGER PCB consists of 14 copper layers, as shown in Fig. A.4. There are eight signal layers available for routing of the interconnections, two power layers for

²⁶metallized drill holes in a PCB to connect between different layers




Figure 5.23: The fan-out of the FPGA signals is conducted in neck mode. This enables to route both traces of a differential pair between two via rows.




Figure 5.24: The propagation delays of the DDR3 memory signals are matched in groups by inserting additional trace length ("snakes" or "trombones").

the supply voltages and four ground planes. The cross section is constructed such that it supports good signal and power integrity. First of all, for each signal layer there is a corresponding reference plane, which allows to define the impedance of the transmission lines and which also carries the return currents. This is important for the high-frequency behavior of the signal traces in order to keep the electronic noise at a low level [117]. The distance between the signal and reference layers influences also the physical constraints for single-ended and differential transmission lines. Proper values for the trace width, DP gap and layer spacing are calculated by the Xsection tool in the *PCB Editor*. Furthermore, the spacing between the power planes and the corresponding ground planes is minimized to 50 μm to form a low impedance power distribution system [118].

The placement of the components and the routing of the signal traces were performed manually for best possible control over the layout. The main challenge while routing the board was the high density of signals below the FPGA. The FF1759 ball grid package exhibits 1759 pins in a 42×42 array with a pitch of 1 mm. To pass through the via array, which is needed to fan out the connections, the neck mode constraint with reduced trace width was used (Fig. 5.23). Design rule checks (DRC) are continuously performed during the routing process to ensure that all constraints are met. As an example, the relative propagation delay of the DDR3 SO-DIMM signals is calculated and a deviation beyond the specified limit is indicated by a DRC marker. The relevant signal traces have thus to be modified with the delay tuning tool by inserting additional segments – the so-called "snakes" (Fig. 5.24). All in all the final TIGER board file contains 847 components, 4163 vias and 4480 routed connections with a total length of 103 m.

5.1.7.4 Simulation

Prior to the production of the board various simulations have been performed, in order to assure best possible signal and power integrity. The power distribution network (PDN) of the

TIGER board has to supply large currents to the devices without creating significant noise. The 1.0 V rail (VCCINT) for example, which provides the core power for the Virtex-6 FPGA, is designed to source up to 20 A. At the same time, the allowed range for the FPGA supply voltage is between 0.95 V and 1.05 V. Two main aspects have been verified with the *Allegro PDN Analysis* tool. Firstly, a static IR drop analysis has been executed to determine the expected voltage drop due to the resistance of the power planes. The obtained color map indicates the voltage drop at any location with respect to the voltage at the output pin of the power regulator. Based on these results, the regions with the highest current densities have been identified and the layout has been adjusted in order to improve the situation, for example by adding more vias, or by widening the copper areas of the power nets. Secondly, a power network impedance analysis was performed to estimate the voltage ripple of the supply nets. The lower the impedance of the PDN is, the lower is the noise that is caused by the high-frequency currents due to the switching of the ICs. To keep the impedance of the PDN low, decoupling capacitors are placed at the board, which quickly provide adequate charge to the ICs. But also the closely spaced power and ground planes of the PCB itself act as a capacitance [117]. Results of this simulation were used to optimize the quantity and position of the decoupling capacitors.

The most sensitive signal traces on the TIGER board in terms of signal integrity are the high-speed serial links, which are connected to the Virtex-6 GTX transceivers. They have been simulated with *Allegro SigXplorer*. With maximum data rates of more than 6 Gbit/s, and also due to the backplane connection, the Aurora lanes (cf. 5.3.2) are the most interesting signals for such a simulation. The procedure will briefly be described in the following.

In a first step, the topology of the probed net is extracted. This creates a model of the transmission line, based on the actual PCB routing including vias and AC coupling capacitors. Thereupon accurate signal models for all involved components are added to the topology. They are usually obtained from the respective manufacturer. For the Virtex-6 FPGA there are models for the GTX transceiver and for the package available from Xilinx [119]. A model for the Tyco backplane connector could unfortunately not be obtained; instead a generic model for a connector with similar characteristics was used. The models are provided in various formats, of which some have to be converted before they can be used in SigXplorer. The model for the Virtex-6 package, which describes the parasitic effects of the bonding wires and the solder bumps, exists in the SPICE²⁷ format. The behavior of the GTX transmitters and receivers is described by IBIS-AMI²⁸ models. The backplane connector model is provided in form of a scattering matrix (S-parameter file). Since an Aurora lane extends from one TIGER board over the backplane to the second TIGER board, several topologies are connected to form a complete transmission channel.

With the *Channel Analysis* function SigXplorer provides a very efficient method for the simulation of high-speed serial links. Thanks to the IBIS-AMI model, which includes the pre-emphasis, equalization and clock recovery algorithms of the GTX transceivers, a large number of bits can be simulated in a reasonable amount of time [120]. The result is presented in form of an eye diagram like in Fig. 5.25, combining the effects of noise, jitter and intersymbol

²⁷Simulation Program with Integrated Circuit Emphasis

²⁸Input Output Buffer Information Specification with Algorithmic Modeling Interface

interference on the signal quality. The simulation can also be used to study how different TX emphasis settings affect the eye opening. The optimum values for the pre-cursor and post-cursor emphasis settings have been determined for the serial link between two TIGER boards for different data rates, as shown in Tab. 5.6.

Table 5.6: Optimized TX emphasis settings and expected insertion loss obtained from a SigXplorer channel analysis simulation.

Data rate	5.0 Gbit/s	6.25 Gbit/s
pre-cursor emphasis	0.61 dB	0.30 dB
post-cursor emphasis	1.94 dB	2.76 dB
differential insertion loss	-4.8 dB @ 2.5 GHz	-6.9 dB @ 3.125 GHz




Figure 5.25: Eye diagram of the differential GTX transceiver input voltage on the RX side of the Aurora channel. Simulated with SigXplorer for a data rate of 6.25 Gbit/s with optimized TX emphasis settings.

5.2 CPU and GPU Extension Boards

The CPU and GPU parts of the TIGER module have been implemented using appropriate add-on cards, which has been decided for several reasons. First, sole embedded CPU and GPU chips are almost impossible to obtain in low quantities since they are usually sold only to certified manufacturers. Second, design complexity is reduced a lot, when pre-assembled modules with standardized form-factors are used, and finally this approach provides the opportunity for future upgrades, when more powerful CPU or GPU models become available.

5.2.1 COM Express Module Specification

For the CPU module the COM Express standard [121, 122] has been chosen, a specification for x86-based computer-on-modules (COM) released by the PICMG²⁹. A COM integrates all necessary components for a very compact PC, including CPU, RAM, non-volatile memory and I/O controllers. However, instead of standard connectors for various PC peripherals the COM features a high-density connector to the baseboard, where the I/O signals can either be directly connected to other on-board components or broken out to external connectors. The COM Express specification defines several module sizes and pin-out types to provide a wide range of products in terms of performance, interface capabilities and power consumption. For the TIGER project the smallest module size ("Mini", $55 \times 84 \text{ mm}^2$) has been selected with pin-out type 10, which is defined since revision 2.0 of the specification.

5.2.2 Kontron COMe-mTT10

The Kontron COMe-mTT10 [123] is a COM Express Mini module with pin-out type 10, and it is hence suited for the application on the TIGER board, where it is mainly used for slow-control tasks like configuration and monitoring. Therefore it does not have to provide lots of computing power but rather a convenient access to the various devices that have to be controlled remotely during run-time.

From the several variants of the COMe-mTT10, which differ in CPU speed, memory size and temperature range, the model COMe-mTT10 E680T 1GB/4GB³⁰ has been chosen for the application on the TIGER module. It features an Intel Atom E680T single-core CPU with a maximum clock frequency of 1.6 GHz and an integrated graphics core, complemented by an EG20T platform controller hub (PCH). The Atom's internal memory controller has access to 1 GB of DDR2-800 RAM, which is shared between the CPU and GPU core, but since the integrated graphics is not used during normal operation of the TIGER module almost the whole memory space is available to the CPU. Besides a number of standard PC interfaces, which will be described in the next subsection, the PCH provides access to a 4 GB solid-state disk (SSD). Both RAM and SSD chips are directly soldered on the COM Express board.

The Kontron COM Express Mini module supports a power supply with a single voltage in the range from 4.75 V to 14 V, which allows it to be directly connected to the primary 5 V rail of the TIGER board without intermediate power conversion. The power consumption of the COM board is between 5 W when idle and 7.5 W under full load.

²⁹PCI Industrial Computer Manufacturers Group

³⁰Part No. 34003-1040-16-1

5.2.3 COM Interfaces

5.2.3.1 External Interfaces

The following interfaces of the COM are made available to the user by breaking them out to external connectors on the TIGER main board:

- A Gigabit Ethernet interface (RJ45) is used to connect the TIGER board to the DAQ control network. The COM can optionally be configured as diskless system and boot from network (cf. 6.2) using the Preboot eXecution Environment (PXE). The Ethernet interface is also used for remote access via SSH by the monitoring and control tools.
- A USB 2.0 port is available on the front panel to connect peripheral equipment like a keyboard, a mass storage device or a JTAG programming cable for remote debugging (cf. 6.2.4). A standard USB hub can be used to connect more than one device at a time. The maximum continuous output current of the USB port is 500 mA.
- For temporary connections of a hard disk or optical disk drive there is a SATA 2.0 port available on the TIGER module. It may be used during OS or software installation.
- The PCH's SDIO interface is connected to a SD card slot, which is located on the TIGER board beneath the COM Express module. SD and SDHC memory cards³¹ with normal speed and high speed bus interfaces as defined in the SD specifications [124] are supported according to the Intel PCH datasheet [125, p.369].
- The LVDS display port of the Intel Atom's integrated graphics unit [126, p.86] is broken out to a Hirose FH12-40S-0.5SV(55) 40-pin FFC³² connector [127]. This type of connector is normally used to attach flat panel displays to the base board, but with an LVDS-to-DVI adapter³³ available from Kontron [128] an external monitor with DVI-D port can be connected. This is mainly used during the OS installation phase and to access the BIOS menu, while under normal conditions the board is operated headless, i.e. without monitor and input devices.

5.2.3.2 On-board Interfaces

Several on-board devices are connected to the appropriate interfaces of the COM Express module. Besides the PCI Express connection to the FPGA, which will be discussed in section 5.2.6, there are I²C bus, SMBus and LPC bus connections used on the TIGER board, as described in the following.

I²C Bus and SMBus

The I²C bus³⁴ is a two-wire serial bus for low-speed communication, consisting of two bidirectional open-drain signals called SDA (serial data) and SCL (serial clock) with pull-up resistors

³¹Secure Digital and Secure Digital High Capacity

³²Flexible Flat Cable

³³Kontron ADA-LVDS-DVI 24bit, Part No. 96007-0000-00-1

³⁴Inter-Integrated Circuit bus

to the supply voltage. A device connected to the I²C bus can either assert its output low, thus forcing the corresponding signal line to GND which indicates a '0', or keep the output floating, which results in the signal line being pulled up to indicate a '1'. A detailed description of the I²C bus protocol can be found in the specification [129].

The implementation on the COMe-mTT10 supports a clock speed of up to 400 kHz (*Fast mode*), 7bit and 10bit addressing modes and clock stretching, which allows slave devices to delay the data transfer. In addition, the Intel Atom CPU features a SMBus (System Management Bus) interface, which is also forwarded to the COM Express connector. The SMBus protocol is derived from the I²C bus specification and it is commonly used on PC mainboards to communicate with power and clock related chips and hardware monitoring devices. The maximum supported clock speed is 100 kHz. From the software point of view the SMBus is handled just like an I²C bus. The internal bus numbers of the COMe-mTT10 are 0 for the SMBus and 1 for the I²C bus.

Fig. 5.26 shows the structure of the SMBus and the I²C bus on the TIGER module. Each bus connects devices which are located on the COM Express module as well as devices which are located on the TIGER mainboard. This separation is depicted by the dashed line, indicating the transition of the buses from the COM to the base board. The device addresses are given in Tab. 5.7 and Tab. 5.8 respectively.

The main I²C bus fans out to four downstream channels using a PCA9546A 4-channel bus switch [130]. These sub-buses can selectively be enabled by writing to a control register of the switch (cf. Fig. 5.27). This allows the connection of multiple devices with identical addresses to one bus structure, as long as they are connected to different sub-buses and only one of the sub-buses is active at a time. On the TIGER board this feature is necessary for the I²C connection of the two SFP transceivers in order to access their enhanced digital diagnostic interface [131] for monitoring purposes. Only one of the two SFP sub-buses may be enabled at a time, otherwise both transceivers would respond to the same address A2h³⁵.

Table 5.7: Slave addresses of the SMBus devices on the COMe module and on the TIGER board.

Address	Device	Comment	Location
98h	Winbond W83771W	COMe hardware monitor	COMe module
A0h	SPD EEPROM	COMe internal use only	
D2h	Clock Generator	COMe internal use only	
9Ch	TI UCD9246	PWM system controller	TIGER board
D4h	ICS9DB102	PCIe clock buffer	

Low Pin Count Bus

The LPC bus [132] has been designed by Intel to replace the parallel ISA bus in modern PCs and it is commonly used to connect low-bandwidth devices like legacy I/O controllers or audio chips. On the TIGER board, it is connected to the CoolRunner-II CPLD to provide various options for the FPGA configuration (cf. 5.1.5).

³⁵The ID interface of the SFP modules can't be used due to an address conflict with the JIDA EEPROM.




Figure 5.26: SMBus and I²C bus structure with connected devices. The dashed line indicates the transition of the buses from the COM to the base board. Devices shown above this line are located on the COMe module, while devices shown below this line are located on the TIGER module.




Figure 5.27: I²C data transfer to write to the control register of the PCA9546A switch with slave address EEh. The four lower bits of the control register indicate which channels are enabled. Any combination of sub-buses may be enabled by writing '1's to the corresponding selection bits B_i . [130]

Table 5.8: I²C bus device addresses. Devices listed in the upper section of the table are located on the COMe module; devices listed in the middle section of the table are located on the TIGER module. The lower section contains the sub-buses, which are controlled by the PCA9546A bus switch.

Sub-bus	Address	Device	Comment
-	58h	Digital potentiometer	COMe internal use only
-	A0h	Kontron JIDA EEPROM	COMe internal use only
-	78h	MI12832DO	OLED display
-	D0h	Si5326A	clock multiplier chip
-	E2h	Si5338	clock generator chip
-	EEh	PCA9546A	4-channel bus switch
0	98h	MXM card	GPU temperature sensor
1	A4h	DDR3 memory	DDR3 SO-DIMM SPD
2	A2h	SFP 1 module	diagnostic interface
3	A2h	SFP 2 module	diagnostic interface




Figure 5.28: Typical I/O read cycle on the LPC bus. Additional wait states could be inserted during the SYNC phase by the peripheral, until it is ready to send data. The last START field in the diagram is already the beginning of the next transfer cycle.

An LPC interface requires only 7 signals: a 33 MHz PCI clock (LCLK), a reset signal (LRESET#), a control line (LFRAME#) to indicate the start of a new cycle, and a 4-bit wide multiplexed command/address/data bus (LAD[3:0]). The data transfer via the LPC bus takes place in so-called cycles. During a cycle, the LAD[3:0] lines carry several fields, each with a length of one or multiple nibbles³⁶ (cf. Tab. 5.9). Depending on the amount of data which has to be transferred, a cycle takes a variable number of clock periods until it is completed. The timing of an exemplary cycle is shown in Fig. 5.28.

Since the Intel Atom LPC controller does not implement DMA or bus mastering cycles or multi-byte firmware memory cycles [126, p.187], only standard memory and I/O cycles can be used for data transmission between the host (CPU) and the peripheral (CPLD). Because of the lower number of required clock periods, it was decided to use I/O cycles for the LPC communication on the TIGER board. The implementation of the LPC interface core in the CPLD is based on an OpenCores project [133]. The following fields are accepted by the LPC core:

- START indicates the start or stop of a transaction. A LAD[3:0] nibble is interpreted as START field, when the LFRAME# signal is asserted. Accepted values are '0000b' (start) and '1111b' (stop/abort), other values are ignored by the core.
- CT/D indicates the cycle type and direction of the data transfer. Only I/O Read ('0000b') and I/O Write ('0010b') cycles are supported by the core.
- ADDR is a 16-bit wide field for I/O cycles, which takes 4 clock periods to transmit. The most significant nibble is driven out first.
- TAR (turn-around) fields are used to turn control over to the peripheral and back to the host. Each turn-around sequence takes two clock periods.
- SYNC fields are sent by the peripheral to indicate it is either ready to complete the cycle or it needs some wait states. However, the CPLD core will always synchronize within 1 clock period and does not require wait states to be inserted.
- DATA is an 8-bit wide field, taking 2 clock periods. Contrary to the ADDR field, in the DATA field the least significant nibble is driven out first. During a read cycle the DATA field is driven by the peripheral, and during a write cycle it is driven by the host.

³⁶half a byte, i.e. 4 bit

Table 5.9: Length of the LAD[3:0] fields for different cycle types (I/O Read, I/O Write, Memory Read, Memory Write) and the minimum cycle duration for each type.

Field	Description	# clock periods			
		I/O R	I/O W	Mem R	Mem W
START	start of the transaction	1	1	1	1
CT/D	cycle type / direction	1	1	1	1
ADDR	I/O or memory address	4	4	8	8
DATA	write data transfer	-	2	-	2
TAR	turn around bus to peripheral	2	2	2	2
SYNC	indicates synchronization	≥ 1	≥ 1	≥ 1	≥ 1
DATA	read data transfer	2	-	2	-
TAR	turn around bus to host	2	2	2	2
	cycle duration	≥ 13	≥ 13	≥ 17	≥ 17

5.2.4 MXM GPU Module Specification

The Mobile PCI Express Module (MXM) standard has been designed to ease the employment of PCIe based graphics modules in small form factor computer systems, including notebooks, blade and rack mount servers, mobile workstations, as well as all-in-one and home theater PCs. The current version 3.0 of the MXM specification has been released in 2008 and consists of two parts: the hardware specification [134] covering the electrical, mechanical and thermal aspects and the software specification [135]. Up to 16 PCIe Gen2 lanes are supported, as well as a wide range of analog and digital display interfaces. Due to the low z-height, MXMs are ideally suited for GPGPU³⁷ applications on low-profile processing boards in VXS or OpenVPX systems. Recently MXM GPUs became more important also for military and aerospace applications [136, 137], since it has been shown that these modules can survive in harsh environments [138].

Two compatible form factors Type A and Type B are defined, covering different performance ranges, space requirements and power consumptions. For the TIGER project the smaller Type A form factor was chosen, which limits the dimension of the GPU card to $82 \times 70 \text{ mm}^2$ and power consumption to 50 W. The MXM electrical interface uses a 0.5 mm pitch, 285-pin card-edge connection system. The module PCB exposes edge finger contacts and is inserted into the baseboard connector similar to a SO-DIMM RAM module, by placing it in the slot at approximately a 25-degree angle and then pushing it down until it is parallel to the baseboard.

5.2.4.1 TIGER MXM Interface

The MXM interface was integrated in the TIGER design following the recommendations of the MXM system design guide [139]. Since the GPU will only be used for general-purpose computing, no display is connected to the MXM card. All signals which are connected on the TIGER board are summarized in Tabs. 5.10 - 5.12. The # suffix denotes an active low signal or the complementary signal of a differential pair.

³⁷General-purpose computing on graphics processing units

³⁸open drain signal

³⁹The specification allows any voltage between 7 V and 20 V and a current of up to 10 A.

Table 5.10: PCIe signal group of the MXM interface.

Signal Name	I/O	Type	Description / connected to
PEX_TX(7..0) PEX_TX(7..0)#	I	Diff	TX lanes from the upstream PCIe device / connected to the Virtex-6 GTX transceivers
PEX_RX(7..0) PEX_RX(7..0)#	O	Diff	RX lanes to the upstream PCIe device / connected to the Virtex-6 GTX transceivers
PEX_REFCLK PEX_REFCLK#	I	Diff	PCIe reference clock / from the ICS9DB102 clock fan-out
PEX_RST#	I	CMOS	PCIe system reset / connected to the CPLD
PEX_CLK_REQ#	O	OD ³⁸	PCIe clock request / connected to clock fan-out via CPLD
PEX_STD_SW#	I	OD	selects voltage swing of differential PCIe signals 0: full swing level 1 or NC: reduced swing level / connected to the CPLD

Table 5.11: Power and thermal management signal group of the MXM interface.

Signal Name	I/O	Type	Description / connected to
SMB_CLK SMB_DAT	I/O	OD	SMBus clock/data for access to the thermal sensor / connected to the SMBus of the CPU via bus switch
TH_OVERT#	O	OD	thermal shutdown request due to over-temperature / connected to the CPLD, which will instantly power down the MXM module when this signal is asserted
TH_ALERT#	I/O	OD	thermal interrupt request / connected to the CPLD, but not used
PWR_LEVEL	I	OD	signals the module to switch to a lower power state / connected to the CPLD, but not used
PWR_EN	I	CMOS	assert to power on the module / connected to the CPLD, used for power sequencing
PWR_GOOD	O	OD	indicates the status of the internal power regulators / connected to the CPLD, used for power sequencing

Table 5.12: System management signal group of the MXM interface.

Signal Name	I/O	Type	Description / connected to
WAKE#	I/O	OD	wake up system from suspend or from soft-off / connected to the CPLD, but not used
VGA_DISABLE#	I	OD	PCI class code select pin for the GPU / connected to the CPLD, tied to GND to select non- VGA device, since the primary display adapter is lo- cated on the COM Express CPU
PRSNT_R# PRSNT_L#	O	OD	MXM module presence detect / connected to the CPLD, used for power sequencing

Table 5.13: Power rails for the MXM module as provided by the TIGER board.

Name	Voltage	Current
PWR_SRC ³⁹	12.5 V	5 A
5V0	5.0 V	2.5 A
3V3	3.3 V	2 A

The MXM card requires three power supply rails as listed in Tab. 5.13. They may be powered up in any sequence, as long as the PWR_EN signal stays low until all voltages are stable. After the assertion of PWR_EN, the module's internal power regulators ramp up and the PWR_GOOD signal indicates that all rails are within the specified tolerance. To shut down the module, PWR_EN must be deasserted before removing power. System reset may be deasserted at the earliest 200 μ s after PWR_GOOD indicates stable power and 100 μ s after a valid PCIe reference clock is provided [134, p.39]. The power-up and reset sequence is controlled by the CPLD on the TIGER board, in order to assure the specified timings and dependencies.

5.2.5 AMD Radeon GPU Card

Although the MXM 3.0 specification was released some years ago, it is still hard to find MXM modules with state-of-the-art GPUs on the market. Only recently, thanks to the increasing interest in GPU-computing for medical, scientific and military applications, the situation seems to improve slowly. For the TIGER project an AMD-based MXM card⁴⁰ from WOLF [140] has been selected. It features a Radeon E6760 embedded GPU running with up to 600 MHz and 1 GB of GDDR5 memory connected via a 128-bit 800 MHz interface. The GPU contains 480 shader processing units providing a peak performance of 576 GFLOP/s (single precision floating point operations).

The AMD Radeon E6760 embedded GPU supports OpenCL⁴¹ [141] for GPGPU computing. OpenCL is a programming framework to perform general-purpose computations on heterogeneous systems, i.e. systems consisting of multi-core CPUs, GPUs and possibly also other processors like DSPs or FPGAs. The C-based language is used to write kernels, which are executed on the cores of the underlying hardware, using the advantages of parallel computing and dynamic scheduling of the tasks across the available cores. OpenCL is a royalty-free open standard and is developed by the Khronos Group. The AMD software development kit and device driver which are needed for developing an OpenCL program and running it on the GPU are described in the software chapter (6.2.5).

5.2.6 PCI Express Interconnection

5.2.6.1 PCI Express Specification

The PCI Express specification [142] has been released in 2002 by the PCI-SIG⁴², defining a new-generation high-performance bus architecture to interconnect devices in computing and communication applications. It was designed to replace the old parallel PCI and PCI-X buses, which

⁴⁰Type: MXM3.0-E6760-VO

⁴¹Open Computing Language

⁴²PCI Special Interest Group




Figure 5.29: Topology of a typical PCIe system, consisting of a root complex, a switch and three endpoints (EP). The bus numbers are assigned during the enumeration process. Adapted from [143, p.48]

were in use since the beginning of the 90's of the last century. PCIe provides vastly increased bandwidths with a lower number of required pins and better scalability, while still maintaining software compatibility with its predecessors by employing the same usage model, address space model and transaction types. An extensive overview about the PCIe architecture is given in [143]. The interconnection of all devices in a PCIe system is called a **hierarchy** (Fig. 5.29). It is implemented using high-speed serial point-to-point links with a scalable number of lanes, usually between x1 (one lane in each direction) and x16 (16 lanes in each direction) in dependence of the required bandwidth. The original specification defined the Gen1 transfer rate of 2.5 Gbit/s per lane, while the Gen2 transfer rate of 5.0 Gbit/s per lane was added in revision 2.0 [144].

The **root complex**, located at the top of the PCIe hierarchy, establishes the connection of the CPU and memory to the PCIe fabric. It contains one or multiple root ports, which can be utilized to connect endpoint or switch devices. The root complex is responsible for the configuration of the PCIe devices immediately after power-up of the system, which includes the enumeration process – i.e. the process of discovering all devices and assigning bus and device numbers – and the assignment of base addresses for each device by writing to the according configuration registers. **Endpoints** are peripheral devices in a PCIe hierarchy, which contain an upstream port only, i.e. a port pointing in direction of the root complex. An endpoint may initiate transactions as a requester or respond to transactions as a completer (cf. 5.2.6.3). For backwards compatibility a legacy endpoint type is defined, which may support IO transactions mapped to the system's IO address space. Since a PCIe link is always a point-to-point connection, multiple devices are connected by using **switches**, which are responsible for the routing of

the data packets to the intended destination. A switch contains one upstream port and several downstream ports pointing away from the root complex. Each port of a switch is associated with a virtual PCI-to-PCI bridge, which are internally connected via a virtual bus (see inset in Fig. 5.29). During the enumeration process, a bus number is assigned to this virtual bus [143, p.50].

5.2.6.2 PCI Express Device Layers

The PCIe communication protocol consists of three layers (Fig. 5.30), namely the transaction layer, the data link layer and the physical layer. For each layer a packet category is defined, i.e. the transaction layer packet (TLP), the data link layer packet (DLLP) and the physical layer packet (PLP). Each packet category originates and terminates at the corresponding layer of the transmitter and receiver device respectively. A **TLP** consists of a header containing the transaction type and routing information, an optional data section and an optional end-to-end CRC⁴³ (ECRC) field. It originates at the transaction layer of the transmitting device and is forwarded downwards through the protocol stack, where additional information is appended. A sequence number and a link CRC (LCRC) field are added by the data link layer and start and end framing characters are added by the physical layer. At the receiving device, the TLP is disassembled stepwise while moving upwards in the protocol stack. **DLLPs** are exchanged between the data link layers of the two link partners for link management purposes, e.g. flow control credit updates and TLP acknowledgement. DLLPs are always related to a direct connection between two devices; they do not contain routing information and do not pass through switches. Finally **PLPs**, which are also called ordered-sets in the specification, are used during link initialization and training and for clock tolerance compensation by regularly inserting SKIP ordered-sets. Like DLLPs, the PLPs also do not contain routing information and do not pass through switches.

On top of the protocol stack the software layer contains the user application, which generates outgoing and processes incoming requests and completions. The **transaction layer** constructs out-bound TLPs and puts them to the TX buffer, and it pulls in-bound TLPs from the RX buffer, checks for errors based on the ECRC field and forwards the packets to the intended destination, which can be either the software layer, another port –if the device is a switch– or the configuration management block. Although not strictly part of the transaction layer according to the specification, the management of the configuration space registers is often implemented at this level in PCIe IP cores, to relieve the user of the burden to handle configuration requests. The flow control of TLPs is based on credits, which indicate how much buffer space is left on the receiving side of a link⁴⁴.

The **data link layer** is responsible for the data integrity of transferred packets. Upon reception of a TLP it is checked for errors based on the LCRC field and it is either acknowledged by sending an ACK DLLP, or the packet is requested again by sending a no-acknowledge (NAK) DLLP. When the transmitting device receives a NAK, it will send the TLP again from its replay buffer. A TLP

⁴³Cyclic Redundancy Check

⁴⁴In fact, up to eight virtual channels may be defined on a link, each comprising dedicated RX/TX buffers and credit information. This allows to implement multiple traffic classes with different priorities.




Figure 5.30: PCIe devices are divided into several layers, each responsible for a specific part of the communication. The physical layer can be further divided into sublayers. In the FPGA, the connection between the MAC layer of the PCIe IP core and the PCS layer located in the GTX transceiver block is made using the PIPE interface. Details can be found in the text.

will be deleted from the replay buffer only after the reception of an ACK, which ensures a high reliability of PCIe data transfers. In order to be able to associate the ACK/NAK DLLPs with the original TLP, a sequence ID is contained in the packets.

The **physical layer** can be divided into a logical and an electrical part. The logical part is further subdivided into the media access layer (MAC) and the physical coding sublayer (PCS). The MAC contains the link training and status state machine, distributes the bytes across the available lanes on the link and scrambles them to reduce EMI noise. The PCS performs 8b/10b encoding/decoding and contains an elastic buffer on the receiving side, which is necessary to compensate for clock differences. The physical media attachment layer (PMA) forms the electrical part of the physical layer, which contains the SERDES converters and the differential analog buffers. Since the PCS/PMA, together also called the PHY, and the MAC often reside in different devices or IP cores, the PIPE⁴⁵ interface has been specified for the connection between PHY and MAC [145].

5.2.6.3 PCI Express Transaction Protocol

Each data transfer in the PCIe fabric is based on a transaction, i.e. a series of packets which are exchanged between a requester and a completer. A transaction is started with a request

⁴⁵Intel PHY Interface for PCI Express

packet and can be either posted, which means it does not require a response, or non-posted. In the latter case it is executed as split transaction, which is finalized with a completion packet at a later time. Several transaction types are defined in the PCIe specification. These include the memory read/write, IO read/write and configuration read/write transactions known from the PCI protocol, as well as the new message transaction for packet-based in-band signaling of errors and interrupts. Only memory write and message transactions are posted, all others are non-posted [143, p.56].

For each transaction type, a corresponding TLP packet type is defined, which is used to issue a request. In addition, two TLP packet types are available for completions with and without data. The packet type is specified in the TLP header, together with routing information to identify the final destination. In contrast to the local link traffic (i.e. DLLPs and PLPs), which is always related to a direct connection between two neighboring devices, TLPs are forwarded by switches⁴⁶ from one link to another in order to carry out a transaction between two devices, which are further apart in the PCIe hierarchy. Depending on the packet type of the TLP one of three available routing methods is chosen: address routing, ID routing or implicit routing. **Address routing** is used for memory and IO requests, referring to the system memory and IO map, and **ID routing** is used for completions and configuration requests, which refer to the target's logical position in the hierarchy, i.e. bus/device/function number. **Implicit routing** is based on the knowledge of a switch, in which direction the root complex is to be found. Message transactions may be routed by any of the three methods [143, p.117].

A special case is the routing of **configuration requests**, which can only be issued by the root complex and can be either of type 0 or type 1. A type 0 configuration request is always consumed by the device which receives it. Hence, if the root complex wants to send a configuration request to a device further down in the hierarchy, it will issue a type 1 request. The only devices paying attention to a type 1 request are PCI-to-PCI bridges. They handle the packet depending on the target bus number, which is specified in the header. If the target bus equals the bus directly connected to the bridge's downstream side (called the secondary bus), the bridge converts the packet from type 1 to type 0 by changing the header and passes it to the secondary bus. If the target bus is located beneath the secondary bus, the packet is forwarded as a type 1 configuration request without modification [143, p.732].

5.2.6.4 TIGER PCI Express Link Structure

COM Express CPU, Virtex-6 FPGA and MXM GPU of the TIGER board are connected by PCIe links, as shown in Fig. 5.31. The CPU contains four PCIe root ports, each providing an x1 Gen1 link with 2.5 Gbit/s in each direction. One link is used internally on the COMe module to connect the PCH. The three remaining links #0 to #2 are available on the baseboard connector. Link #0 of the CPU module is connected to the first PCIe interface of the FPGA at block location X0Y0. Since the bandwidth of the Atom's root complex is rather low, it was decided to interconnect the FPGA and the GPU directly instead of hooking the GPU to another root port of the CPU. This allows for a high-speed, low-latency data transfer between the two devices utilizing

⁴⁶A multi-port root complex may also implement routing between its ports, although this is not required by the PCIe specification [142, p.31].




Figure 5.31: PCI Express structure on the TIGER board: while the COM Express module is connected to the FPGA via a x1 Gen1 link, the MXM GPU module is connected to the FPGA via a x8 Gen2 link providing much higher bandwidth.

the maximum available bandwidth of the Virtex-6 PCIe block. The interconnection between the FPGA PCIe block X0Y1 and the MXM GPU is designed as an x8 Gen2 link, featuring a raw bandwidth of up to 40 Gbit/s in each direction.

5.2.6.5 TIGER Switch Design

Due to the structure of the PCIe links on the TIGER board, the FPGA must act as switch in order to create a PCIe conform hierarchy with the root complex of the COM Express CPU on top, but at the same time the FPGA should implement endpoint functionality. A schematic overview of the switch design is given in Fig. 5.32. It comprises a virtual three-port switch, which is responsible for the TLP routing to the specified destination. The implemented routing mechanisms will be described below. The upstream port of the switch is connected to the COM Express module, the external downstream port is connected to the MXM GPU, and the internal downstream port is connected to the FPGA's endpoint logic. In addition, a bus-mastering DMA⁴⁷ engine is included in the FPGA design, which can be programmed to perform memory read and memory write transactions. The DMA transactions may target the system RAM as well as memory that is located on another PCIe device and has been mapped to an address range in the system memory map. This setup allows to transfer data

- between the CPU and the FPGA, by targeting the internal memory range,
- between the CPU and the GPU, by targeting the memory window of the external downstream port,
- and between the FPGA and the GPU, by initiating a peer-to-peer transfer to be performed by the DMA engine.

The peer-to-peer transfer is the preferred method to exchange data between the FPGA and the GPU, since the CPU and system memory are not involved. Hence, a lower latency is achieved and the transfer rate is only limited by the bandwidth of the PCIe x8 Gen2 link between the two devices. In addition, RAM and CPU resources are saved for other tasks. The CPU is just running

⁴⁷Direct Memory Access




Figure 5.32: The PCI Express hierarchy on the TIGER board requires the implementation of a switch in the FPGA. The upstream port of the switch is connected to the COM Express module, one of the downstream ports is connected to the MXM GPU. The second downstream port is a virtual port, which is internally connected to the FPGA's endpoint logic.

the GPU driver and initializing the peer-to-peer transfer. As soon as the transfer starts, request and completion packets are exchanged directly between the peers.

For each TLP, which is handled by the switch, one of the three available routing mechanisms is employed, depending on the TLP type (cf. 5.2.6.3), as shown in Fig. 5.33. If the TLP type indicates address routing, the header contains a destination address, which is compared by the switch against the memory windows of the downstream ports. These windows are defined by the base and limit registers in the type 1 configuration space headers, which are initialized during the enumeration process. If, on the other hand, the TLP type indicates ID routing, the header contains a destination ID, i.e. bus, device and function numbers. The routing decision is made based on the secondary and subordinate bus number registers in the configuration space headers of the downstream ports. Implicit routing is only based on the direction of the TLP, which is either upstream or downstream.

The concept of the switch design for the TIGER FPGA has been implemented using the *PCI Express XpressRICH3* IP core by PLDA [146]. This highly customizable IP core solution can be configured for rootport, endpoint, bridge or switch operation and is compliant with the PCIe base specification rev. 3.0. On the physical side of the core, a PIPE interface is used to connect to the PHY, which is contained in the Xilinx Virtex-6 GTX transceivers. On the other side of the core a transaction layer interface is provided for the user application.

An endpoint reference design is included with the *XpressRICH3* core, which is able to perform basic TLP processing and supports PIO⁴⁸ memory read and memory write transactions as well

⁴⁸Programmed Input/Output




Figure 5.33: Flow chart for the routing decision of the virtual PCIe switch, which is implemented in the TIGER FPGA.

as DMA transfers. To implement the switching functionality outlined above, two instances of the core are utilized to form the upstream and downstream switch ports. A user application module was developed in the hardware description language Verilog to implement all switch specific functions, since they are neither part of the core nor part of the reference design. All routing mechanisms, which have been described above, are supported by the switch module. In addition, the configuration space interface of the downstream bridge is controlled by the user code, to enable access to the configuration space registers via CfgRd and CfgWr transactions during the enumeration process.

The design was extensively tested in the course of a product evaluation with simulations in ModelSim [147], using PLDA's *PCI Express Bus Functional Model* (BFM) [148]. The BFM emulates a PCIe environment and allows to send and receive packets, monitor bus activity and report errors. Two instances of the BFM have been connected to the user design to verify correct behavior of the switching logic and the configuration space access. Unfortunately, a synthesis of the design for the Virtex-6 hardware was not possible with the evaluation version, since the PHY code is not provided with this package and a full license was not obtained as of this writing. The first generation of FPGA firmware designs for the TIGER module, including all designs which were used during the 2012 DVCS data taking, did not include this switch functionality. A standard endpoint design using the Xilinx *Virtex-6 FPGA Integrated Block for PCI Express* IP core was used instead. Details about the FPGA firmware and the driver software can be found in sections 6.1.1 and 6.2.3 respectively.

5.3 Interfaces

The interfaces provided by the CPU have already been specified in the previous section. The remaining interfaces of the TIGER board, which are directly connected to the FPGA, are described in the following. They can be classified into high speed I/Os (5.3.1 - 5.3.3), an experiment synchronous reference input (5.3.4) and general purpose I/Os (5.3.5).

5.3.1 VXS Interface

Since the TIGER module is designed for the operation in VXS switch slots, its backplane connection consists of a large number of point-to-point links for communication with other modules in the crate. As explained in section 4.3.1, each GANDALF payload board is directly connected to each of the two switch boards via the so-called VXS payload ports (PP). Furthermore, the so-called VXS switch ports (SP) provide direct connections between the two switch slots. In total, the TIGER board exhibits five VXS backplane connectors, four of them (P2 - P5) being Tier 2 high speed connectors and one (P1) being a Tier 1 sideband connector [80, p.18].

5.3.1.1 High Speed Signals

The VXS Tier 2 connectors⁴⁹ are arranged without gaps in between each other to form a monolithic connection area featuring a total of 192 differential pair signals, which are shielded by intermediate ground pins to reduce crosstalk. On the TIGER board the signals are allocated as follows (pin-out tables in appendix C.1):

- $18 \times 8 = 144$ differential pairs to the GANDALF boards,
- 16 differential pairs to the second TIGER board,
- 32 unused differential pairs (not connected).

While the VXS interfaces of the GANDALF and TIGER modules are electrically and mechanically designed according to the VITA 41.0 specification, they do not comply with any VXS protocol layer as defined in the VITA 41.x standards. These standards are packet-switched serial protocols like InfiniBand, Serial RapidIO, 1/10 Gigabit Ethernet or PCI Express. But due to specific requirements of the CAMERA trigger and readout application, a custom transmission protocol was developed for the VXS communication in the GANDALF framework (cf. 4.3.2). The TIGER module provides a dedicated VXS port to every payload slot in the crate, each consisting of eight unidirectional point-to-point links based on a 500 MHz DDR LVDS transmission. The nomenclature of the TIGER VXS links is given by the following pattern, with $n \in \{1, 2, \dots, 18\}$ being the number of the VXS port and $i \in \{0, 1, \dots, 7\}$ being the link index within the port: `VXSPORT_ $n$ < $i$ >`.

The direction of the links is chosen in the FPGA firmware design depending on the application of the TIGER board. In the trigger design (6.1.2) all eight links are configured as inputs to receive the trigger primitives; the readout design (6.1.3) uses six inputs to receive the event data and two outputs to transmit the TCS information. The link assignment is detailed in Tab. 5.14. The VITA 41 documents do not define the mapping between VXS port numbers and crate slot numbers [80, p.27]. This is left to the manufacturer of the backplane and may be chosen to

⁴⁹Tyco 1410137-1, 1410138-1, 1410139-1

Table 5.14: Usage of the VXS links between the GANDALF and the TIGER modules. The backplane signal names in the central columns refer to the VITA 41 specification. This does, however, not reflect the link direction as defined in the GANDALF framework.

GANDALF		VXS backplane signal		TIGER		
usage	signal name	GANDALF	TIGER	link index	usage	
VXS link_clk out	VXS_A<0>	Tx 0	Rx 0	0	VXS link_clk in	VXS switch slot A
S-LINK out<0>	VXS_A<1>	Tx 1	Rx 1	1	S-Mux in<0>	
S-LINK out<1>	VXS_A<2>	Tx 2	Rx 2	2	S-Mux in<1>	
S-LINK out<2>	VXS_A<3>	Tx 3	Rx 3	3	S-Mux in<2>	
S-LINK out<3>	VXS_A<4>	Rx 0	Tx 0	4	S-Mux in<3>	
	VXS_A<5>	Rx 1	Tx 1	5		
TCS data in	VXS_A<6>	Rx 2	Tx 2	6	TCS data out	
TCS clock in	VXS_A<7>	Rx 3	Tx 3	7	TCS clock out	VXS switch slot B
Trigger out<0>	VXS_B<0>	Tx 0	Rx 0	0	Trigger in<0>	
Trigger out<1>	VXS_B<1>	Tx 1	Rx 1	1	Trigger in<1>	
Trigger out<2>	VXS_B<2>	Tx 2	Rx 2	2	Trigger in<2>	
Trigger out<3>	VXS_B<3>	Tx 3	Rx 3	3	Trigger in<3>	
Trigger out<4>	VXS_B<4>	Rx 0	Tx 0	4	Trigger in<4>	
Trigger out<5>	VXS_B<5>	Rx 1	Tx 1	5	Trigger in<5>	
Trigger out<6>	VXS_B<6>	Rx 2	Tx 2	6	Trigger in<6>	
Trigger out<7>	VXS_B<7>	Rx 3	Tx 3	7	Trigger in<7>	

optimize signal routing. The mapping which is used in the GANDALF framework is shown in Tab. 5.15. It specifies the relation between the geographic address of the payload slot, the VXS port number of the switch slot and the I/O bank assignment of the TIGER FPGA.

5.3.1.2 Sideband Signals

The VXS Tier 1 connector⁵⁰ provides single-ended pins used for low speed sideband signals. Apart from a few system management pins inherited from VME64x (SYSRST#, SYSFAIL#, GA0#, ...), there are two signals available for each payload board – PP_n_SCL and PP_n_SDA when referring to the VITA 41.x protocol layer standards. In the custom protocol of the GANDALF framework these signals are employed for a handshaking system between the GANDALF and TIGER boards in the VXS crate. Since they are not used as a I²C bus as the VITA 41.x names would imply, the signals will be called PP_C(*n*) and PP_D(*n*) in the following to avoid confusion.

The handshaking was implemented to ensure the hot-swap capability of the boards, allowing to safely insert them into and remove them from the crate at any time. The need for such a protection system arises from the fact that the FPGAs on the TIGER and the GANDALF board are directly connected to each other by the VXS backplane signals, without any DC blocking capacitors, buffer or driver chips or similar devices. This poses a possible danger for the FPGAs in situations when one of the boards is unpowered. The FPGA's I/O buffers are designed with a conventional CMOS output structure containing two clamp diodes on every pin, one to VCCO and one to GND [149]. When driving an unpowered bank, the clamp diodes will become

⁵⁰Tyco 1410421-1

Table 5.15: Mapping between the VXS port number and the slot number, called geographic address, for the Hartmann backplane [83] used in the GANDALF framework, as well as the assignment of Virtex-6 I/O banks on the TIGER module. Since one FPGA I/O bank with 20 differential pairs provides connections for 2.5 VXS ports, some ports are split between neighboring banks.

VXS port	Slot (GA)	FPGA I/O bank
1	10	14
2	13	14
3	9	15/14
4	8	13/12
5	14	15
6	15	13
7	7	15
8	6	13
9	16	16
10	17	12
11	5	16
12	4	12
13	18	17/16
14	19	22/23
15	3	17
16	2	22
17	20	17
18	21	22

forward biased. For this case the Virtex-6 data sheet specifies a maximum current of 10 mA through any pin in a powered or unpowered bank. A total limit of 100 mA per bank [150, p.2] applies to prevent damage to the device. If there is a low impedance path from VCCO to GND, excessive current flow could lead to device failure [151]. Therefore, the FPGA output buffers which drive the VXS signals to the backplane are enabled only after the link partner has indicated that it is ready to receive data.

The handshake signals $PP_C(n)$ and $PP_D(n)$ are connected to the Virtex-5 FPGA on the GANDALF side, and to the CoolRunner-II CPLD on the TIGER side, and they are pulled to 3.3 V by weak pullup resistors on the VXS backplane. When a GANDALF module is inserted into a slot, connected to port n , and while its FPGA is still unpowered, the corresponding $PP_C(n)$ and $PP_D(n)$ signals are forced to GND due to the clamp diodes of the Virtex-5. As soon as the FPGA is powered, the I/Os enter a high-impedance state, and as a result the handshake signals are pulled again to 3.3 V. When the FPGA firmware is finally loaded, it may drive $PP_C(n)$ low to request TCS signals. The TIGER board detects the state of the GANDALF boards based on Tab. 5.16.

On the other hand, the CPLD on the TIGER module is driving $PP_D(n)$ low, as long as the Virtex-6 is not powered up and configured, which causes the GANDALF FPGA to tri-state its VXS outputs. Furthermore, there are two software commands which have to be issued to the GANDALF board in order to enable the VXS outputs, namely the Fast Registers *FR_ReadoutTiger*

Table 5.16: Truth table to determine the state of the GANDALF boards on the basis of the handshake signals.

PP_C	PP_D	State
1	X	no GANDALF board is placed in the slot, or the design is not requesting TCS signals
0	0	a GANDALF board is placed in the slot, but it is switched off
0	1	the GANDALF board in this slot is functional and requesting TCS signals

Ready and *FR_TriggerTigerReady*. Details about the Fast Register commands can be found in the GANDALF User Guide [152].

Finally, it should be mentioned, that the GA0# signal is used by the CPLD to distinguish between the two switch slots in the VXS crate. Slot A on the left-hand side, which is reserved for the TIGER readout concentrator, is identified by GA0#='0', and slot B on the right-hand side, which is reserved for the TIGER trigger processor, is identified by GA0#='1'. The slot position is also announced to the FPGA, in order to cross check whether the correct design type (i.e. readout or trigger) has been loaded.

5.3.2 TIGER-to-TIGER Links

Two high-speed interfaces between the TIGER modules within a VXS crate are provided:

- Switch port SP4 on connector P2 is connected to FPGA user I/Os in the same manner as the VXS payload ports. It is used for a synchronous data transfer between the TIGER boards following the same protocol as described above. This includes the forwarding of the TCS signal to the trigger TIGER to gain access to the COMPASS reference clock and the SPS spill structure as well as the transmission of event data from the trigger TIGER to the readout TIGER for monitoring purposes. Further details on the exchanged data will be given in section 6.1.2.
- Switch port SP3 on connector P5 is connected to GTX transceiver tiles on the Virtex-6 FPGA with AC-coupling capacitors on the RX side, thus providing a 4-lane high-speed serial link between the TIGER boards. By using the Xilinx IP core generator, an Aurora interface can be customized and integrated in the FPGA user design. Aurora 64B/66B [153] is a lightweight link-layer protocol which is suited for chip-to-chip as well as board-to-board serial communication, providing high throughput of up to 25 Gbit/s for a 4-lane design with very little protocol overhead. The implementation process is detailed in section 6.1.1.6.

In addition, there are eight single-ended signals (SW_SE<7..0>) located on the VXS sideband connector, which can be used to transmit status information between the two TIGER boards. SW_SE<0> indicates the status of the readout TIGER and SW_SE<1> indicates the status of the trigger TIGER. The remaining signals are currently unused.

5.3.3 SFP Transceiver Sockets

Two sockets for small form-factor pluggable (SFP) transceivers reside on the TIGER board. SFP [154] is a de facto standard for compact, hot-pluggable transceiver modules commonly used in optical networking applications. SFP modules provide an electrical interface to the motherboard and an optical interface to connect a fiber optic cable. Modules with various optical transceiver types are available for both multi-mode and single-mode fibers. The most popular optical connector is the LC type [155, p.57]. For communication over short distances of up to 550 m usually multi-mode fibers are used in conjunction with LEDs or laser diodes operating at a wavelength of 850 nm (transceiver type SX). Typical multi-mode fibers have a core diameter of 50 μm and a cladding diameter of 125 μm . A classification system is used to specify the modal bandwidth, i.e. the bandwidth–distance product. To this day, it defines four classes: OM1 to OM4. The most common multi-mode fibers today are OM3 (50/125) and OM4 (50/125), which can be recognized by their cyan jacket, in contrast to the older OM1 and OM2 fibers, which are usually orange [155, p.69].

For a connection to the DAQ system the TIGER board is equipped with up to two SFP modules⁵¹. Currently the CERN-developed HOLA⁵² S-LINK protocol is used at COMPASS to connect the TIGER data concentrator to the readout buffer PCs, providing a peak data rate of 160 MB/s per link (cf. 6.1.3). However, a future upgrade of the DAQ readout PCs could yield data rates of up to 6.6 Gbit/s, which is the limit of the Virtex-6 GTX transceivers.

5.3.4 TCS Interface

The TIGER module features an interface to the COMPASS trigger control system, which is identical to the one of the GANDALF module. The TCS signal, which is distributed by optical fibers to every readout board in the experiment, provides a global reference clock as well as first-level triggers and SPS spill information (cf. 3.4). In order to permit the employment of the TIGER module also in other experiments which may use a different system for the distribution of triggers, the TCS interface is held flexible thanks to exchangeable add-on cards, the so-called *Gimli* cards [64, p.45]. The standard type provides an input for the optical TCS fiber, and performs a clock recovery from the serial data stream. Two LVDS signals are finally transmitted from the *Gimli* card to the TIGER board: the TCS reference clock (155.52 MHz at COMPASS) and the TCS data stream. The connector pin-out is to be found in [63, p.167]. The trigger data and clock signals from the add-on card are available in the Virtex-6 FPGA, where the decoding of the TCS information is performed (cf. 6.1.1.3). Furthermore, a copy of the clock signal is provided to the Si5326A clock multiplier chip in order to generate additional experiment synchronous clocks (cf. 5.1.3.2). An alternative version of the *Gimli* card is available, featuring inputs for clock and trigger signals as well as an oven-controlled crystal oscillator which provides a stable, low-jitter 20 MHz reference clock.

5.3.5 General Purpose I/O

A number of general purpose inputs and outputs have been added to the TIGER board, which may be used for arbitrary tasks within the Virtex-6 firmware. First of all, there are two LEMO

⁵¹for example: Finisar FTLF8524P2BNV SFP Transceiver

⁵²High-speed Optical Link for Atlas

00 sockets⁵³ available on the front panel to provide the output signals LVTTLOUT1 and LVTTLOUT2. These LVTTL-compatible signals are controlled via the FPGA user-design and can provide up to ± 24 mA of continuous output current⁵⁴. The LEMO outputs are most commonly used for the transmission of the generated trigger signal from the TIGER trigger processor to the COMPASS trigger logic.

A dual VHDCI⁵⁵ female connector Honda HDRA-ED136LFZGT [157] is employed to provide 32 LVDS inputs and 32 LVDS outputs. The inputs are located on the left-hand side connector, and the outputs are located on the right-hand side connector (see Fig. 5.35). The connector pin-out is to be found in appendix C.2. Both input and output signals are buffered by LVDS receivers/drivers⁵⁶ before entering the FPGA, in order to protect its I/Os from short circuits and electrostatic discharges. The LVDS I/Os may be used in the FPGA design for example to input additional detector information like veto signals and to output additional trigger related signals.

Finally, also the front panel display and the LEDs can be counted to the I/Os in a sense (Fig. 5.34). The 32×128 pixel OLED display of the type Multi-Inno MI12832DO [158] is used to show general board status information. It is controlled via the I²C bus interface of the COMe module using a software tool, which is described in 6.3. Several LEDs at the front panel and at the rear side of the PCB indicate the current status of the module. Their meanings are listed in Tab. 5.17. While the front panel LEDs are visible from a distance, those on the rear side are rather meant for debugging purposes. Fig. A.1 in appendix A may be used for locating the LEDs on the PCB by means of their reference designator.

Table 5.17: Front panel LEDs and debug LEDs of the TIGER module.

Name	RefDes	Color	Meaning
Front LED 1		red	on: main power failure blinking: CPU suspended
Front LED 2		green	on: power-up completed blinking: shutdown in progress
Front LED 3		green	on: FPGA configured
Front LED 4		green	<i>currently not used</i>
HDD ACT	LD1	red	blinking: SSD harddisk activity
Sby p_fault	LD6	red	on: standby hot-swap controller fault (over-current)
Sby p_good	LD7	green	on: all standby power rails ok
Main p_fault	LD8	red	on: main hot-swap controller fault (over-current)
CPLD 1R	LD9	red	on: GPU power failure
CPLD 1G	LD10	green	<i>currently not used</i>
CPLD 2R	LD11	red	
CPLD 2G	LD12	green	

⁵³LEMO EPL.00.250.NTN [156]

⁵⁴due to the use of a TI SN74LVC2G34 buffer gate

⁵⁵Very-High-Density Cable Interconnect

⁵⁶ON Semiconductor NB4N855S



Figure 5.34: The TIGER module features four LEDs and an OLED display on its front panel. The meaning of the LEDs is explained in Tab. 5.17. The display is used to show general status information.




Figure 5.35: A dual VHDCI connector provides 32 LVDS inputs on its left hand side and 32 LVDS outputs on its right hand side. In addition, two LVTTTL outputs are available using LEMO connectors.

6. Firmware and Software

The hardware of the TIGER module has been detailed in the previous chapter. But since an FPGA based electronic device like this is not operational without the adequate firmware, the three firmware designs, which are available to date, will be outlined in the following section. The so-called *Base design* contains all the necessary building blocks to implement the board interfaces, and to control the internal clock generation and the start-up of the FPGA. Two further designs have been derived from the *Base design*, in order to implement the specific behavior of the *Trigger Processor* and the *Readout Concentrator* applications.

In the subsequent sections the software to control the TIGER board is briefly explained. Starting from a description of the operating system and device drivers running on the COM Express CPU (6.2), the development of various configuration and monitoring tools is highlighted (6.3).

6.1 FPGA Firmware

The FPGA firmware for the TIGER board is written in VHDL¹, a hardware description language for digital systems, which is specified in the IEEE 1076 standard. VHDL is a text-based method of describing the structure and the behavior of a logic circuit. It is adopted to design models for synthesis as well as to write test-benches for simulations. However, only part of the syntax can be synthesized to program actual hardware. VHDL supports the description of concurrent systems by the use of processes. A comprehensive introduction to digital system design with VHDL may be found in [159, 160].

In several places throughout this chapter VHDL keywords, signal or constant names will be mentioned. They will be set in typewriter font for better recognition. Constants are written in uppercase letters.

To implement a firmware design for use on a Xilinx FPGA device, the VHDL source code is translated by the synthesis tool to a register-transfer level (RTL) schematic, which is subsequently converted into a gate-level description, called a netlist. This file is then used as an

¹Very High Speed Integrated Circuit Hardware Description Language




Figure 6.1: Top-level block diagram of the TIGER *Base design* implemented in the Virtex-6 FPGA. The interface signals are indicated by black arrows, and the clock signals are drawn in blue color.

input to the implementation tools, which perform the steps translate, map, place and route, in order to generate the physical layout. The result is finally stored in a bitstream file to configure the device.

6.1.1 TIGER Base Design

A top-level block diagram of the *Base design* is shown in Fig. 6.1. The purpose of the *Base design* is to provide the user defined application with all the necessary interfaces, global clocks and control signals. The goal is to relieve the firmware designer, who will write the actual user application, of the burden to implement all the interface and board management logic, as well as to ensure the compatibility between different user designs in terms of the communication protocols with the outside world. Therefore several VHDL components are instantiated in the *Base design* next to the application module, whose functionalities will be described in the following. All clock inputs as well as I/O signals, which enter the Virtex-6 FPGA, are connected to the top-level entity as shown in the schematic drawing.

The *Base design* is developed in the form of a Xilinx ISE PlanAhead [106] project, which includes a number of source files, specifically

- a top-level VHDL file, containing the FPGA top-level module,

- VHDL files for the *Base design* components instantiated at the top-level,
- UCF constraint files, defining the pin locations, IO standards and timing constraints,
- test-bench VHDL files for behavioral simulation of the design or parts of it,
- the TG_defines package file, which declares constants and types not to be changed by the user,
- and the TG_parameters package file, which declares constants the user may change to customize the design.

To give an example for a user definable constant in the TG_parameters package, the constant TIGER_FUNCTIONALITY may either take the value TIGER_FUNC_TRIGGER or the value TIGER_FUNC_READOUT to specify the intended design type. Different code variants will be instantiated by means of VHDL generate statements at several places in the *Base design*, depending on this constant's value.

6.1.1.1 CPLD Interface

The CPLD interface module is used to transmit status flags as well as some reset and control signals between the CPLD and the FPGA on the TIGER board. The SelectMap A<> and D<> buses (cf. Fig. 5.12), which are routed using a fly-by topology from the CPLD via the Platform Flash XL to the FPGA, are reused for this purpose. Before the FPGA is initialized, the SelectMap signals act as configuration pins, used to transfer the bitstream into the FPGA's configuration SRAM. As soon as the device configuration is finished, these I/Os are accessible in the VHDL top-level module, where they are called CA<> and CD<>. The availability of the CPLD interface signals on the SelectMap buses is indicated by the CA<22:21> flags (see Tab. 6.1).

To coordinate the usage of the SelectMap signals, the CPLD firmware implements a register called flash_mode, which is set by means of the CPLD driver (cf. 6.2.2) to one of the following values:

- FM_STARTUP – flash is enabled for automatic FPGA configuration (default),
- FM_SELECTMAP – flash is disabled, SelectMap signals are driven by the CPLD,
- FM_FLASHPROG – flash is enabled to be reprogrammed by the CPLD,
- FM_NORMAL – flash is disabled, CA and CD signals are used for FPGA communication.

The CPLD interface uses the CD bus to announce the presence of the GANDALF boards in the VXS crate to the FPGA. This information is generated by the CPLD based on the VXS handshake signals (cf. 5.3.1.2), and evaluated by the FPGA in order to decide whether the output buffers of the VXS ports are to be enabled or tri-stated. The CA bus is used to transmit various flags between the CPLD and the FPGA. CA<22:8> contains flags from the CPLD to the FPGA, while CA<7:0> transmits flags in the opposite direction. The assignment of the flags is shown in Tab. 6.1.

6.1.1.2 Clock Management, Reset Logic and VXS I/O Buffers

As seen in the *Base design* top-level schematic, three VHDL blocks are directly connected to the CPLD interface module. The first one is the clock management block, which is responsible for the generation and distribution of several clock signals for use throughout the FPGA design. A

Table 6.1: Assignment of the CPLD interface flags to the CA<> and CD<> bus bits. Bits, which are not listed here, are available for future use.

Bus range	Name and description of the flag
– CPLD to FPGA –	
CD<17:0>	status of the GANDALF boards (CD<i> == '1' ⇔ VXS port (i+1) is ready)
CA<22:21>	"10" (constant) indicates that CA and CD signals are valid (FM_NORMAL)
CA<20>	TigerFunction: switch slot position, which defines the module function (<i>Trigger</i> or <i>Readout</i>)
CA<19>	COMe_cb_reset#: carrier board reset signal, asserted by CPU at boot time
CA<18>	OtherTigerReady: status of the other TIGER board in the crate
CA<17>	UserReset1: optional reset signal, software-controlled via CPLD driver
CA<10>	tcs_lo1: TCS receiver "loss of lock" (from CLC016)
CA<9>	si_lo1: Si5326A "loss of lock"
CA<8>	si_los: Si5326A "loss of signal"
– FPGA to CPLD –	
CA<7>	tcs_rate: rate select pin for the TCS receiver
CA<6>	si_inc: Si5326A input-to-output skew increment pin
CA<5>	si_dec: Si5326A input-to-output skew decrement pin
CA<2>	c2_startup_rst: startup reset for clock domain 2
CA<1>	c1_startup_rst: startup reset for clock domain 1
CA<0>	binfile_mismatch: indicates that the loaded design type does not match the module function

total of eight differential clock signals enter the Virtex-6 FPGA on the TIGER board for various purposes (see Tab. 6.2). Three of them are connected to the clock management block, namely `clk_si1`, `clk_si2` and `clk_ref`. They are routed through global clock buffers (BUFG) in order to make them accessible everywhere in the FPGA. In addition, the following Virtex-6 device primitives are instantiated:

- An IDELAYCTRL module [161, p.112], connected to `clk_ref`, performs the calibration of the individual IODELAY elements to reduce the effects of process, voltage, and temperature variations. IODELAY elements are used in the DDR3 memory interface (5.1.2) as well as in the VXS backplane link interface (4.3.2) to dynamically compensate for the different input delays of the signals in the data buses.
- A MMCM_BASE primitive [92, p.40], fed by `clk_ref`, generates a set of clock frequencies needed for the *Readout* application, in particular 155 MHz, 40 MHz, 300 MHz and 100 MHz clocks.
- A MMCM_BASE primitive, fed by `clk_si2` (505.44 MHz), generates a set of TCS related clock frequencies needed for the *Trigger* application, in particular $0.8 \cdot f_{si2}$, $0.2 \cdot f_{si2}$, $1.0 \cdot f_{si2}$ and $0.5 \cdot f_{si2}$.

The second VHDL module, which is connected to the CPLD interface, is the reset logic. It provides various reset signals, which are deasserted in a coordinated manner after the device

Table 6.2: List of the clock signals which enter the TIGER FPGA. Origin and frequency of the clocks, as well as the destination VHDL block in the FPGA design are specified. The signals are either connected to global clock capable user I/O pins (GC), or to the MGTREFCLK inputs of the GTX transceiver tiles.

Origin	Clock signal	Frequency	Destination	Clock Pins
Gimli	tcs_clk_V6	155.52 MHz	TCS interface	GC
Si5326A	clk_si1	variable	Base design clock mgmt.	
	clk_si2	505.44 MHz		
Si5338	clk_ref	200.0 MHz		
	clk_ddr3	266.6̄ MHz	Memory interface	
	aurora_mgtclk	125.0 MHz	Aurora core	
	sfp_mgtclk	125.0 MHz	User app.	
COMe	pcie_sys_clk	100.0 MHz	PCIe core	MGTREFCLK

start-up. A prompt and a delayed version of the reset signal may be generated for every clock domain. The reset signals carry names of the form *clockdomain_startup_rst[_delayed]*, i.e. the base name "startup_rst" is prefixed by the name of a clock domain and may be suffixed by the string "_delayed". The prompt reset signal is deasserted synchronously as soon as the correspondent clock signal is present. The delayed reset signal is deasserted synchronously a defined number of clock cycles after the CPLD interface enters the flash mode FM_NORMAL. The software-controlled user reset is also synchronized to the various clock domains. It should be noted that the delayed reset signals will not be deasserted in the case of a binfile mismatch, i.e. if the loaded FPGA design file does not match the TIGER function (*Trigger* or *Readout*) as defined by the position of the board in the crate.

Finally, the last top-level elements controlled by the CPLD interface are the VXS I/O buffers. The differential I/O buffers for the VXS ports are instantiated in the *Base design* using VHDL `if...generate` statements, to choose the signal direction based on the constant TIGER_FUNCTIONALITY. In case of the *Trigger* design, eight input buffers (IBUFDS) per VXS port are generated, while a VXS port of the *Readout* design consists of six inputs and two outputs (cf. Tab. 5.14). Differential tri-state buffers (OBUFTDS) are used for the outputs, with the T pin being controlled by the CPLD interface based on the status of the GANDALF board on the other side of the link.

6.1.1.3 TCS Interface Module

The TCS clock and data signals provided by the *Gimli* card (cf. 5.3.4) enter the TCS interface module of the *Base design*. The serial data stream is decoded by first demultiplexing the A and B channels, which carry the first-level trigger (FLT) signal and the TCS commands respectively. The phase information which is needed for an unambiguous assignment of the two channels can be obtained from the TCS data itself due to the biphasic mark encoding scheme and the bit statistics [64, p.50]. This phase information also allows to reconstruct the original 38.88 MHz COMPASS reference clock with a constant latency.

While the FLT signal is immediately made available to the user logic, the B-channel data is parallelized and interpreted as TCS commands, which are 20 bits wide plus an optional 6-bit




Figure 6.2: Schematic diagram of the method to determine the alignment of the 505.44 MHz Si5326A output clock with respect to the 38.88 MHz COMPASS reference clock.

wide checksum. The most common ones are the broadcast commands BC1 and BC2 [162, p.18]. BC1 is used to transmit the *Begin of Spill* and *End of Spill* signals together with the spill number. BC2 contains the event type and event number. Based on this information a so-called event label is generated and written to a FIFO, from where it is read out by the user application at a later time.

As mentioned earlier, the Si5326A clock multiplier chip is utilized to reduce the jitter of the 155.52 MHz TCS clock and to generate some related clock frequencies for use in the TIGER FPGA. These clocks are required to show a fixed phase relationship to the COMPASS reference clock. Due to the operation principle of the Si5326A chip, the phase offset between the input clock and the output clock is not defined per se. This may be illustrated with the help of Fig. 5.11, showing that the output clocks are obtained by dividing the high-frequency clock of the digitally-controlled oscillator ($f_{\text{DCO}} \approx 5\text{GHz}$). This will result in a random phase offset, depending on which clock edge the division counter is started. However, the input-to-output skew is controlled by the INC and DEC pins [100, p.60] in steps of $1/f_{\text{DCO}}$, allowing to align the Si5326A outputs with the COMPASS reference clock, provided that the output frequency is an integer multiple of the reference frequency.

The alignment process – also called *Si sweep* – is an iterative process of shifting the skew value until a minimum phase offset is achieved [163]. In each step, the actual alignment of the output clock and the reference clock is statistically determined, as shown in Fig. 6.2. Two copies of the high frequency output clock are routed through IDELAY elements of the FPGA, in order to delay one of the signals with respect to the other. The delay value should approximately equal the skew increment step of the Si5326A. Both signals are then registered by flip-flops using the lower frequency reference clock. In the case that the edges of both clocks are correctly aligned to each other, a '1' will be seen on the non-delayed branch, but on the delayed branch a '0' is still visible, since the edge has not yet propagated through the delay. Therefore, a '1' at the output of the AND gate will indicate a phase match. Finally, a counter is used to determine the degree of alignment for every skew increment step, in order to detect the optimal setting.




Figure 6.3: Internal structure of the PCIe interface block.

6.1.1.4 PCIe Endpoint Block

A PCIe endpoint interface is included in the *Base design* to enable a high-bandwidth communication between the CPU and the FPGA on the TIGER board. It is based on an open-source PCIe DMA controller ('pcie_sg_dma') available at OpenCores [164], which was extended to provide some project specific features. The main elements of the PCIe interface block (see Fig. 6.3) are

- a Xilinx integrated PCIe IP core, configured as an endpoint with a Gen1 x1 link and three base address registers (BAR0 - BAR2),
- the controller, containing a TLP manager to receive and send transaction layer packets and the DMA and PIO engines,
- and a memory bridge, to direct the read and write transactions to the appropriate targets.

Three target address ranges are located at the slave side of the memory bridge, each related to a specific memory type. The register space (BAR0) contains system registers, status, control and DMA related registers as well as a number of user definable registers. The memory space (BAR1) implements a dual-port Block RAM with a size of 4096×64 bits, and BAR2 contains two FIFO data ports, one in read and one in write direction. FIFO control and status registers are located in BAR0. Analogous to the three address spaces, the PCIe interface block exposes three types of communication ports to the user application: user registers, a configuration memory and a FIFO interface.

User registers are available in the TIGER *Base design* in two different forms: the so-called *Static registers* and *Fast registers*. The *Static registers* are 32-bit registers that can be accessed by PCIe PIO read and write operations. There are read-only registers (`pcie_rx_regs(i)`) and write-only registers (`pcie_tx_regs(i)`). The *Static registers* may be used in the design for arbitrary purposes. However, the register `pcie_tx_regs(1)` has a special function: it controls the *Fast registers*, which are 1-bit write-only registers (`FastRegister(i)`). The *Fast registers* are a convenient way to manipulate control signals in the FPGA design. They can be either set to '1' or

'0' or they can be toggled for one period of the PCIe transaction clock. A typical application for *Fast registers* is to enable, disable or trigger a certain process in the FPGA design.

The configuration memory (*Configmem*) is located in the dual-port BRAM at BAR1. One port of the BRAM is connected to the PCIe memory bridge, enabling read and write access by the CPU. The second port is connected to a 3-port memory bridge [165] featuring a WISHBONE slave interface for each port. This allows up to three WISHBONE [166] compliant masters in the user application to access the *Configmem*. The bridge arbitrates the requests of the masters in a round robin manner. A typical use case for the *Configmem* is the online update of certain registers in a VHDL component. First the new values are written to the appropriate address in the *Configmem*, and then a *Fast register* is toggled, which triggers a VHDL process to transfer the values from the *Configmem* to the registers.

The third interface between the PCIe core and the user application logic are two 64-bit wide FIFOs with configurable depth (16 k by default). A board-to-host FIFO is available for data transmission from the FPGA logic to the CPU's memory, and a host-to-board FIFO is used to transmit data in the opposite direction. For highest performance the FIFOs should be accessed in DMA transfer mode.

6.1.1.5 ChipScope Cores

Xilinx ChipScope Pro [167] is a powerful tool-set for in-circuit debugging and verification of FPGA designs. It provides the firmware developer with a convenient logic analyzer solution, which is able to monitor virtually any signal in the design, without the need to bring it out to an external test pin. The ChipScope solution consists of a set of IP cores, which are customized and generated by means of the *Xilinx CORE Generator* tool and inserted into the VHDL design under test:

- The ILA (integrated logic analyzer) core is the main part of the ChipScope logic. It provides a customizable number of data and trigger inputs, in order to connect the signals to be investigated. Flexible match units are available for elaborate triggering, based on conditions which are configured during runtime, to detect the events of interest. Upon a trigger event the ILA captures data frames of variable length and stores them using on-chip Block RAM.
- The VIO (virtual input/output) core is used to monitor and drive internal FPGA signals. Both synchronous and asynchronous inputs and outputs are available.
- The ICON (integrated controller) core establishes the communication between the aforementioned cores and the host computer via a JTAG download cable. It can connect up to 15 instances of ILA and VIO cores.

To analyze the TIGER FPGA firmware, ChipScope cores have been inserted into the *Base design*, which may be connected to the signals of interest. The design is synthesized and implemented as usual, either with the ISE tool-chain or the PlanAhead design flow. The FPGA is initialized

with the resulting bitstream and hooked up to a *Platform Cable USB* via the JTAG boundary scan port (cf. 5.1.5.3). The *ChipScope Pro Analyzer* tool is a graphical user interface to communicate with the ChipScope cores in the target design. It is used to define the trigger settings and the sampling parameters of the logic analyzer cores, to visualize the acquired data, and to control the virtual I/Os. Automatic data export can be configured to allow an offline analysis of the data.

6.1.1.6 Optional Cores

Two optional VHDL modules, drawn as dashed boxes in the top-level diagram Fig. 6.1, are available in the *TIGER Base design*, which may be enabled by the application designer in order to access the correspondent interfaces. These are the memory interface to the DDR3 SO-DIMM and the Aurora interface for a communication between the two TIGER modules in the VXS crate. When an optional core is instantiated in the *Base design*, it is necessary to add also the respective UCF file to the project, in order to include additional location and timing constraints.

Memory Interface

The DDR3 memory interface core (5.1.2) has been generated with the Xilinx MIG tool [95] by specifying the timing parameters of the Hynix SO-DIMM [96] and the pin-out of the TIGER board FPGA (appendix B). The core provides the controller and PHY modules to access the DDR3 SDRAM. It is also responsible for the initialization of the memory device after power-up and the calibration of the read and write data paths to account for delays. The user interface of the MIG core provides the connections for the application module. It presents a flat address space to the user and maps it to the native bank/column/row addressing scheme of the SDRAM. A description of the relevant user interface signals is found in [95, p.70]. The MIG core instantiates an additional MMCM, to generate required clock signals based on the 266 MHz reference clock.

To determine the performance of the memory interface on the TIGER module a MIG core controlling the Hynix module has been created for the Virtex-6 FPGA. The MIG example design, which is also generated by the wizard together with the memory interface core modules, contains a traffic generator to create various test patterns, which are continuously written to the memory and read back to verify the integrity of the data transfer. The available test patterns include repetitive patterns like "walking 1s" and "walking 0s", which are typically used for memory testing, as well as pseudo-random data. The example design was integrated in the *TIGER Base design* and extended by counters for the number of read cycles and the number of detected bit errors. The counters have been connected to a ChipScope core (see 6.1.1.5) in order to access the counter values during the test using the ChipScope Analyzer software. In addition, the traffic generator and the memory interface could be controlled by ChipScope VIO signals.

The MIG core is specified up to 533 MHz according to the Xilinx datasheet, but during the performance evaluation the interface could even be overclocked to 600 MHz before bit errors started to appear. This might be at least partly due to the fact, that the test design does not use a significant portion of the available FPGA resources and placement and routing can be therefore realized in an optimal way. The maximum clock rate may be reduced in a full design. Several

long term tests have been performed to determine an upper limit for the bit error rate (BER) of the memory interface. The traffic generator was operated for 21 hours in pseudo-random mode and for 24 hours in "walking 1s" mode, thereby transferring 360 TB of data without any bit error. This allows an estimation of the BER to $< 3.5 \cdot 10^{-16}$.

Aurora Interface

A TIGER-to-TIGER high-speed serial interface is available based on the Aurora 64B/66B LogiCORE IP [153]. A 4-lane full-duplex configuration with a streaming interface has been generated for use with the *Base design*. The Virtex-6 GTX transceivers of the Aurora core connect to the RX and TX lanes of the VXS switch port SP3 (cf. 5.3.2) on the backplane. The reference clock is connected to a dedicated MGT clock input of the transceiver tile to minimize jitter. The Aurora module contains an additional MMCM to generate the user clock, the transmission clock and internal synchronization clocks for the serial transceivers. The user interface of the core provides the necessary ports to stream data between the two link partners. In the streaming mode, the Aurora channel is used as a pipe. Words are written into the TX side and get delivered to the RX side of the pipe. The data width is 64 bits per lane, i.e. 256 bits for the x4 link between the TIGER boards. Besides the data ports there are data valid and ready flags available to control the data flow. Gaps in the data transmission may be generated by the sender at any time by pulling the valid flag low. Native flow control signals enable the receiver to pause the data transmission by sending an XOFF message. Finally, some status signals are available to indicate software or hardware errors as well as the channel initialization status.

In order to achieve the highest possible data rate the GTX transceiver settings (cf. 5.1.1.6) must be tuned carefully. They greatly depend on the electrical characteristics of the transmission channel. The optimum TX emphasis parameters have been obtained by channel simulations as described in section 5.1.7.4. The RX equalization parameters are determined using the ChipScope IBERT² core [167, p.18]. With the help of this core a bit error rate test of the high-speed serial channels between two TIGER boards is performed. Each FPGA implements an instance of the IBERT core, which generates a pseudo-random test pattern (31-bit PRBS) for transmission to the link partner, and at the same time it checks the pattern that it receives from the link. An automatic sweep of the RX parameters can be performed to determine the settings with the lowest bit error rate.

After the optimum settings had been obtained, 12-hour test runs have been conducted for data rates of 5.0 Gbit/s and 6.25 Gbit/s. The 5.0 Gbit/s channel did not produce a single bit error in 12 hours, resulting in an error rate of $\text{BER}_{5.0\text{Gbps}} < 5 \cdot 10^{-15}$. The 6.25 Gbit/s channel showed an asymmetric behavior. In the direction from switch slot A to switch slot B no bit error occurred during the test, but in the opposite direction 124 bit errors were detected during a total transmission of $2.53 \cdot 10^{14}$ bits. This results in an error rate of $\text{BER}_{6.25\text{Gbps}} = 4.9 \cdot 10^{-13}$.

6.1.1.7 Partial Reconfiguration Workflow

It has already been mentioned in the last chapter (5.1.5.2) that the Virtex-6 supports dynamic partial reconfiguration (PR). This allows to change a specific part of the FPGA firmware on the

²Integrated Bit Error Rate Tester




Figure 6.4: Sketch of a partial reconfiguration project.

fly while the remaining part of the device keeps operating. The implementation of a PR design poses some additional demands because not all FPGA elements may be reconfigured online. In addition, a well-regulated behavior of the static part must be guaranteed while the reconfiguration is in progress.

To evaluate the partial reconfiguration feature, the *TIGER Base design* was converted into a PR capable test design. This was possible without substantial modifications, since the basic structure already fulfilled most of the requirements. The PR workflow in Xilinx PlanAhead includes the following steps [168]:

- The static top-level design instantiates a *reconfigurable partition* (RP) as a black box (Fig. 6.4). The design is synthesized, which results in a netlist that does not contain the reconfigurable logic.
- One or more *reconfigurable modules* (RM) are synthesized independently as separate projects. A RM netlist contains the reconfigurable logic which is implemented in the RP.
- A project for the actual implementation is created in PlanAhead with the PR option enabled. The top-level netlist is imported and the black box instance is defined as RP. The RM netlists are added to the RP. An area constraint (*Pblock*) is attached to the RP, in order to define the FPGA region that gets reconfigured during the PR process.
- An implementation run, called a *configuration* in PlanAhead, is created for every RM. For the first configuration both the static logic and the RM is implemented. Afterwards the static partition is *promoted* for use in subsequent implementation runs. In the other configurations the static logic is *imported* and only the RM is implemented.
- Finally for each configuration there are full and partial bitstream files generated.

The modules of the *TIGER Base design*, which have been described in this chapter so far, constitute the static top-level design of the PR test project. This includes all the interfaces at the FPGA boundary as well as the clock related parts (MMCM, BUFG). These elements cannot be

reconfigured on the fly.³ The user-defined part of a TIGER design is developed as a RM. It may contain logic that is mapped to CLBs, Block RAM and FIFO elements and DSP blocks. All signals that cross the boundary between the static partition and the reconfigurable partition have to use so-called partition pins. These are dedicated logic elements forming a routing bridge (basically a 1-input look-up table). They are inserted automatically during the implementation process to guarantee that the connection points for any RM are at the same position.

Compared to the standard *Base design*, the PR capable version implements the ICAP component of the Virtex-6 and some additional logic. It forwards the partial bitstream from the PCIe host-to-board FIFO to the SelectMap-like interface of the ICAP. To ensure a well-defined behavior during the reconfiguration process, all signals that enter or leave the RP must be registered with flip-flops at the partition boundary. It is beneficial to do this on both sides for eased timing closure. The static logic must ignore the outputs from the RP during the reconfiguration, since the signals at the partition pins are undefined until the configuration is complete. Hence, the enable signal of the flip-flops is used to isolate the logic of the two partitions.

6.1.2 TIGER Trigger Processor Design

The first application specific project derived from the *Base design* is the TIGER trigger processor. It implements the proton trigger conditions for the CAMERA detector as outlined in section 4.2.2. Details about the development and verification of the trigger logic are to be found in a recent thesis [77]. The main features are briefly summarized in the following.

The trigger processing takes place in four stages, as shown in Fig. 6.5. The first stage is the 'slat_logic' which correlates the PMT pulse information from the upstream and downstream end of each scintillator slat, in order to reconstruct the time and position of a hit. A 'slat_logic' instance exists for every CAMERA scintillator, each containing a 'hit_manager_up' and a 'hit_manager_down' module. The hit managers receive the trigger primitives, i.e. the GANDALF pulse parameters, from the VXS interface and store them in form of 'pmt_info' data objects. They are also responsible for the T0-correction to compensate for the individual delays of the detector channels (cf. 4.2.2.2). The 'slat_logic' operates on the 'pmt_info' elements to detect coincidences and finally outputs 'slat_info' objects, which are further processed in the second stage. Here, a 'track_logic' instance for each possible A-B coincidence is responsible for the reconstruction of possible particle tracks. This process results in a series of 'track_info' data objects, which contain the relevant track parameters. In a third stage, the 'cut_logic' filters the coincidences found in the previous step based on the event topology and on various conditions like energy thresholds or cuts in the energy loss distribution (cf. Fig. 4.7). Coincidences that survive these cuts enter the fourth and last stage – the 'release_logic'. Its purpose is to buffer the trigger decision until it is finally released at the right time to fulfill the constant latency requirement (cf. 4.2.3). As a result a single pulse is generated at the LVTTL trigger output.

In total 12 GANDALF modules digitize the 96 PMT channels of the CAMERA detector and transmit the information about detected pulses via the backplane to the TIGER trigger processor.

³However, some elements (MMCM, GTX transceiver) provide a dynamic reconfiguration port to change their settings.




Figure 6.5: The TIGER trigger logic is arranged in four stages. The trigger primitives arriving from the backplane are combined to form hit information, and these are further combined into tracks. After applying optional cuts, the remaining coincidences cause trigger time stamps being enqueued in the release buffer. Subsequently a trigger pulse is emitted at the specified time. [77, p.46]

The VXS ports that are used for the reception of trigger primitives are located at the left edge of the Virtex-6 device, as shown in Fig. 6.6. The mapping between the CAMERA channels and the FPGA I/O banks has been defined considering an optimized routing of the coincidence logic. Since coincidences can occur between the combinations A_i-B_i and A_i-B_{i+1} , the corresponding 'slat_logic' instances should be located close to the 'track_logic' instances to which they contribute. As it can be seen, this is fulfilled for all combinations but for one ($A_{11}-B_{12}$), which is connecting the top and bottom end of the device. Optimal placement of the logic blocks is ensured by carefully floorplanning the design. During this task the design hierarchy is broken apart into so-called Pblocks⁴, which can be constrained to a specific device area [77, p.62].

Due to the allocation of the CAMERA channels to specific I/O banks in the trigger design on the one hand, and the fixed mapping between VXS ports and backplane slots on the other hand, the position of the GANDALF modules in the crate is implied. The resultant configuration of the CAMERA readout crate is given in Tab. 6.3.

To monitor and control the TIGER trigger processor, the design adopts some of the interfaces provided by the *Base design*.

- The PCIe board-to-host FIFO may be used to spy on the 'pmt_info', 'slat_info' and 'track_info' streams of selectable channels for test and debug purposes. Scaler values are also transmitted via this interface to monitor the hit and trigger rates [77, p.59].

⁴Physical blocks




Figure 6.6: PlanAhead screenshot showing the FPGA floorplan of the trigger design. The detector channel numbers are noted next to the inputs. The hierarchical blocks of the trigger logic are placed at defined locations by the use of Pblocks and area constraints. The number of connections between the 'slat_logic' and the 'track_logic' Pblocks is indicated by the red (many) and orange (fewer) lines. Pblock_2 contains the coincidence $A_{11}-B_{12}$. [77, p.64]

Table 6.3: Configuration of the CAMERA readout crate, as it was used during the 2012 DVCS run.

Slot (GA)	2	3	4	5	6	7	10
IO bank	22	17	12	16	13	15	14
Detector	FI1	CAMERA ring A					
Channel	0..5	8..11	12..15	4..7	16..19	0..3	20..23

Slot (GA)	13	14	15	16	17	20	21
IO bank	14	15	13	16	12	17	22
Detector	CAMERA ring B						FI2
Channel	20..23	0..3	16..19	4..7	12..15	8..11	0..5

- The backplane connection to the TIGER readout concentrator is used to transmit the aforementioned monitoring information in an event-based way for inclusion in the S-LINK data stream. This enables an offline comparison of the trigger data with the detector data recorded by the GANDALF modules [77, p.60].
- The configuration memory (*Configmem*) provided by the PCIe interface is utilized to set a multitude of parameters for the trigger logic. These include the trigger latency, the width of the coincidence windows for the slat and the track logic, exclusion of CAMERA segments from the trigger, and monitoring settings, to name just a few [77, p.61].
- The TCS information received from the backplane is used to synchronize the system time at each *Begin of Spill* signal, in order to allow the interpretation of the coarse time information, which is contained in the GANDALF trigger primitives [77, p.45]. The system time is also needed for the release logic to guarantee a constant trigger latency.

6.1.3 TIGER Readout Concentrator Design

The readout concentrator is the second application of the TIGER module. It is used in the VXS crate to multiplex the S-LINK data received from the GANDALF modules into a single data stream to the DAQ, and it also distributes the TCS signal to all the modules in the crate. The development of this application firmware will be described in full detail in a forthcoming thesis [169]. The most important aspects are summarized in the following.

The COMPASS DAQ system adopts the S-LINK standard to transmit the event data from the readout modules to the spill buffer PCs (cf. 3.5). Each S-LINK connection is able to transmit up to 160 MB/s using the current HOLA protocol, or up to 100 MB/s using the older ODIN protocol [170]. On the receiving side a link destination card (LDC) mounted on a PCI spill buffer card is required for every S-LINK. Up to four of this cards may be inserted in a readout PC. To optimize the utilization of the links in order to minimize the number of required PCs and spill buffer cards, it is desired to multiplex several readout modules into one S-LINK. For the CATCH modules the SMUX card [35, p.61], which is plugged into the rear transition connector of a VME crate, provides up to 4:1 multiplexing of neighboring modules. This is used for detectors producing low or medium data rates in order to save on DAQ equipment.

For the GANDALF readout framework it was decided to develop a similar S-LINK multiplexing system. Thanks to the high-bandwidth point-to-point connections on the VXS backplane, this could even be achieved without additional rear transition cards. Four data signals together with a clock signal are available from each GANDALF module to the TIGER readout concentrator module located in switch slot A using the VXS_A port (cf. Tab. 5.14), allowing to transfer the S-LINK data words chopped up in nibbles. An identical connection is also available from the TIGER trigger module. After the S-LINK words have been received and parallelized again on the TIGER side they are written into buffer FIFOs. Every input port features a dedicated FIFO, deep enough to hold several complete events. Custom FIFO logic keeps track of the S-LINK start and end markers to detect the events' boundaries and performs some data integrity checks.

The S-LINK multiplexer logic of the TIGER readout concentrator design concatenates the data streams of up to 18 GANDALF boards and two TIGER boards to form a so-called SMUX event.

As soon as all active ports indicate the presence of a new event, the multiplexing process is started. It first snoops on the event headers in the input FIFOs in order to determine the total size of the SMUX event and to compare the spill and event numbers of all ports against each other. If the spill/event number of an event does not match, or a GANDALF module does not send data within a certain timeout period, the correspondent port is disabled till the end of the spill. Finally, all sane events are written to an output FIFO, preceded by the SMUX header containing the total event size and the spill/event number.

The S-LINK transmission to the COMPASS DAQ is performed with a HOLA LSC core kindly provided by the CERN PH-ESE group [171]. This core utilizes the Virtex-6 GTX transceivers at a line speed of 2.0 Gbit/s to implement the SERDES functionality of the TI TLK2501 transceiver chips, which are present on the standard S-LINK hardware. Together with a multimode fiber transceiver plugged into the SFP socket of the TIGER board, the HOLA LSC core provides a usable bandwidth of 160 MB/s and is able to interface with the existing DAQ equipment of the experiment.

The COMPASS trigger information and reference clock are distributed via the TCS optical fiber network. It is built up of ten primary laser transmitters with subsequent passive optical 1:32 splitters, yielding a total of 320 available fiber connections [56]. To cut down on the number of required TCS fiber outlets, the readout concentrator is also used to distribute the TCS clock and data signal via the backplane to all GANDALF modules and to the second TIGER module. The VXS lines with indices <6> and <7> of each payload and switch port transmit the TCS data and clock signals respectively. Virtex-6 output DDR registers (ODDR) are operated for this purpose in a clock forwarding mode [161, p.117], in order to minimize the skew.

For locations, where a TCS is not available, e.g. at test beams or in the lab, the functionality of the TCS controller [162, p.11] has been implemented in a VHDL module, which may be instantiated in the readout concentrator design. Different trigger modes are available, including an internal pseudo-random trigger with variable rate, a 1 kHz clock-like trigger and an external trigger provided by a NIM signal. Event and spill numbers are incremented automatically, and the first-level trigger and event label information are multiplexed following the same procedure as in the COMPASS TCS. This allows to run a stand-alone DAQ system consisting of a TIGER readout concentrator module and up to 18 GANDALF boards in one VXS crate.

6.2 Operating System and Device Drivers

Since the COM Express module features a standard x86 CPU, it can run with a conventional Linux operating system. Kontron provides a board support package (BSP) for the COMe-mTT10, which contains board specific drivers and kernel patches for kernel 2.6.37 to support the module's hardware [172]. The BSP was integrated into a Fedora 14 distribution (32-bit) and installed on the SSD of the CPU module. To ease the software maintenance for multiple TIGER boards in the DAQ environment a diskless setup was created, allowing to boot all systems from a single network source using PXE. For this purpose an *initramfs* image was created with the dracut tool and the client file system was copied to a NFS share following the instructions in [173].

Since the Fedora project is the basis for the Red Hat Enterprise Linux distribution, which in turn is the basis for Scientific Linux – the standard Linux distribution at CERN – the Fedora operating system for the TIGER board was integrated into the COMPASS environment without difficulty. The versions of Scientific Linux CERN (SLC) that are currently used for the COMPASS DAQ are SLC5 (kernel 2.6.18) for the online and run control machines and SLC4 (kernel 2.6.9) for most of the frontend CPUs.

6.2.1 Kontron Drivers

The Linux BSP for the Kontron COMe-mTT10 module includes a set of drivers for the on-board PLD⁵, called the `kempld` driver package. It consists of a core driver (`kempld-core`), which provides the driver framework and is necessary for all other `kempld` drivers, and several function drivers. The most important ones are the I²C driver (`kempld-i2c`) and the watchdog driver (`kempld-wdt`) [172].

The `kempld-i2c` driver provides access to the I²C controller in order to communicate with I²C devices both on the COMe module and on the TIGER board. The driver can be used by other kernel drivers as well as by user-space applications. Some configuration tools for the TIGER system (cf. 6.3) utilize this driver for I²C communication.

The `kempld-wdt` driver is used to activate, deactivate and trigger the watchdog of the COMe module. A watchdog timer is a hardware component used to reset the system in case of a software failure. Under normal conditions it is regularly triggered by the software to reset the timer. But if the software becomes unresponsive, the notifications will fail to appear and the watchdog will timeout and reset the system. Since the TIGER modules are located in the COMPASS experimental area, which is unaccessible during beam time, manual resets of the CPU are not easily possible. Therefore the watchdog feature is an important precaution to guarantee minimum downtime of the system. The simplest way to use the watchdog is to open the file `/dev/watchdog` and regularly write some characters to it, in order to notify the watchdog that the user-space program is still alive. If no notification is received within a certain timeout period, the system is reset.

6.2.2 CPLD Driver

The TIGER CPLD is connected to the LPC bus and it is located in the system's I/O map at address 800h. A direct access to the CPLD registers is possible with low-level port input and output functions like `inb()` and `outb()`, however the usage of these functions in user space requires I/O permissions and the user would have to deal with the actual hardware register contents. Therefore, a simple *ioctl* char driver [174, 175] was written, which creates a device file at `/dev/cpld0` to ease the communication with the CPLD. The defined *ioctl*() commands are listed in Tab. 6.4.

6.2.3 PCIe Driver

Communication with the PCIe interface of the Virtex-6 FPGA is accomplished using the Linux `pciDriver` [176] and the MPRACE library [177]. The Linux `pciDriver` is a generic driver for PCI

⁵Programmable Logic Device

Table 6.4: TIGER CPLD ioctl() commands.

ioctl() command	type	description
CPLD_GET_STATUS	_IOR	returns the status register; for the interpretation of the data fields see Tab. 6.5
CPLD_SET_FM	_IOW	sets the flash mode of the CPLD interface (cf. 6.1.1.1)
CPLD_SET_CA	_IOW	sets the CA signals of the SelectMap interface
CPLD_SET_CD	_IOW	sets the CD signals of the SelectMap interface; the CCLK is automatically toggled by the CPLD after this register has been written
CPLD_RESET_FPGA	_IO	resets the FPGA by pulling down the Program_B pin
CPLD_INDICATE_PDOWN	_IO	notifies the power sequencing state machine of the CPLD that a shutdown is in progress
CPLD_RESET_USER1	_IO	toggles the user defined reset signal
CPLD_RESET_TMCRATE	_IO	toggles the RATE pin of the <i>Gimli</i> card; this usually solves a locking problem of the CLC016

Table 6.5: TIGER CPLD status register. The meaning of the status signals is explained in Tab. 6.1.

31	14	13	0
GANDALF Status GS<17:0>		TIGER Status TS<13:0>	
status of all the GANDALF boards in the VXS crate		various status information about the TIGER board	
GS<i> == '1' ⇕ board at VXS port (i+1) is ready		TS<13> – OtherTigerReady TS<12> – c2_startup_rst TS<11> – c1_startup_rst TS<10> – si_los TS<9> – si_lol TS<8> – tcs_lol TS<7> – ucd_pgood (DC/DC power good) TS<6> – ldo_pgood (low drop-out reg. pwr good) TS<5> – INIT_B pin of FPGA TS<4> – DONE pin of FPGA TS<3> – binfile_mismatch TS<2> – TigerFunction TS<1:0> – flash_mode	

devices, which may of course also be used for PCI Express devices, since the two standards are software compatible. The driver maps the PCI BARs to the user space and provides access to the configuration registers of the device. The MPRACE library provides C++ classes for convenient I/O transactions with the PCIe board.

A new class for the TIGER board has been derived from the basic mprace Board class. The TIGER class implements public functions for read and write access to the three BARs of the PCIe endpoint interface. The registers (BAR0) can only be accessed by PIO operations, while the *Configmem* (BAR1) and the FIFOs (BAR2) can be accessed by both PIO and DMA operations.

Two important extensions have been made to the TIGER class in order to allow an online FPGA reconfiguration:

- A helper class (TgCPLD) was created to manage the access to the CPLD registers via the device file `/dev/cpld0`. It implements amongst others the functions `getStatus()` to read the CPLD status register and `configFPGA()` to reset the FPGA and reconfigure it via the SelectMap interface. The TIGER class provides wrapper functions for these CPLD functions.
- Two additional `ioctl()` commands are implemented in the `pciDriver`, which are used to call the kernel functions⁶ `pci_save_state()` and `pci_restore_state()`. These are PCI bus service functions allowing to save the PCI configuration space of a device to a `pci_dev` struct (normally called before suspending the system) and to restore it afterwards.

The wrapper functions `savePciState()` and `restorePciState()` were added to the C++ interface of the `pciDriver` (class `PciDevice`) and to the TIGER class, allowing to explicitly call the kernel save and restore functions. They must be executed before and after the reconfiguration of the FPGA in order to retain the content of the PCIe configuration registers of the Virtex-6 integrated endpoint.

6.2.4 JTAG Driver and ChipScope Server

The Xilinx Platform USB cable drivers were installed on the TIGER system, by following the instructions in [178]. This allows to use the Xilinx Lab Tools, specifically the iMPACT tool for programming a CPLD or Platform Flash XL, and the ChipScope server for FPGA firmware debugging. These tools can run directly on the TIGER CPU, without the need of a dedicated PC next to the VXS crate. A JTAG cable is simply connected to the USB port of the TIGER board and to the JTAG TAP of the device under test. Besides the FPGA on the TIGER board itself, this may also be the JTAG chain of one of the GANDALF boards in the crate. After the ChipScope server has been started on the TIGER CPU, the Analyzer GUI can be launched on any PC in the local network in order to perform a debugging session.

⁶defined in kernel source file `drivers/pci/pci.c`

6.2.5 AMD Driver and OpenCL SDK

The AMD Catalyst 12.10 proprietary Linux x86 display driver has been installed to support the Radeon E6760 for GPGPU applications. It should be noted that the Catalyst Control Center option has to be enabled during the installation and the XServer must be started even though the system is operated headless, otherwise the access to temperature monitoring and performance settings of the GPU is not possible. In addition, DPMS⁷ has to be disabled in the `/etc/X11/xorg.conf` in order to prevent the GPU from entering a power saving state due to the absence of a display, which would result in GPGPU applications running at degraded performance.

For the development of OpenCL applications the AMD Accelerated Parallel Processing (APP) Software Development Kit (SDK) v2.8 has been installed. The APP SDK provides AMD's OpenCL implementation consisting of an API and a runtime, as well as several sample applications. When an OpenCL program is executed, a series of API calls configure the system for execution, an embedded just-in-time compiler compiles the OpenCL code, and the runtime asynchronously coordinates execution between parallel kernels [179].

6.3 Monitoring and Control Tools

A number of command line tools have been developed to control the TIGER module. The most important ones will be described in the following.

tiger_status

This tool prints out the status of the TIGER board, i.e. power-good status of DC/DC converters and linear regulators, FPGA configuration status, TCS and Si5326A lock status, temperature and voltage monitoring information.

```
usage:  tiger_status [-v] [-c]
optional arguments:
  -v :  verbose mode, prints all information, default is to print only errors
  -c :  print the crate overview (indicates which VXS ports are ready)
```

tigersm

The TIGER SelectMap configuration tool is used to load the FPGA firmware. It first saves the PCI state, located in the PCIe configuration registers, then it reconfigures the FPGA with the specified binfile and restores the PCI state afterwards.

```
usage:  tigersm binfile
arguments:
  binfile :  path to the FPGA binfile
```

⁷Display Power Management Signaling

tg_register

A tool to read and write the *Fast registers* and *Static registers* (i.e. User registers) of the PCIe endpoint interface.

```
usage:  tg_register -r|-w -f|-u reg [data]
arguments:
  -r :  register read command
  -w :  register write command
  -f :  Fast register access
  -u :  User register access
  reg  :  register number
  data :  data word, required for write commands; in conjunction with option
          -f only the values 0, 1 or 2 (i.e. toggle) are allowed
  reg and data are either decimal or hexadecimal numbers (with prepended 0x)
```

tg_configmem

A tool for read and write access to the *Configmem* of the PCIe endpoint interface.

```
usage:  tg_configmem -r|-w addr [data]
arguments:
  -r :  read access
  -w :  write access
  addr :  Configmem address
  data :  data word, required for write commands
  addr and data are interpreted as hexadecimal number
```

si5326prog

A tool to program the Si5326A clock multiplier chip. It parses a register file, created with the SiLabs Precision Clock EVB Software, writes the data via the I²C bus to the chip, and waits until the PLL has locked.

```
usage:  si5326prog registerfile
arguments:
  registerfile :  path to the Si5326 register .txt file
```

si5338prog

A tool to program the Si5338 clock generator chip. It parses a register file, created with the SiLabs ClockBuilder Desktop Software, writes the data via the I²C bus to the chip, and initializes the locking process. As soon as the PLL has locked, it writes the current VCO calibration values to the FCAL registers (cf. flow chart in [102, p.21]).

```
usage:  si5338prog registerfile
arguments:
  registerfile :  path to the Si5338 register .txt file
```

displayinit

This tool is used to initialize the front panel display (Multi-Inno MI12832DO) according to the procedure in [158, p.14]. The display contains a SSD1306 controller [180] which is programmed using the Adafruit SSD1306 OLED driver library [181].

`usage: displayinit`

ttconfig and trigger-FIFO

The tools `ttconfig` and `trigger-FIFO` are specific to the TIGER trigger processor design. Their usage is described in [77].

7. Commissioning

7.1 CAMERA Installation

During a changeover period in September 2012, the COMPASS experiment was prepared to perform DVCS test measurements. Among several other changes, the newly-designed liquid hydrogen target and the surrounding CAMERA detector have been put into operation. The CAMERA detector had been completely mounted and tested in advance in a construction hall ("clean area") next to the COMPASS hall, which allowed to crane the whole structure right into the experimental area. The preparations included the following steps:

- mounting the scintillator/light guide elements on the retaining structure,
- mounting the photo multiplier tubes,
- mounting the high-voltage power supply and the VXS crate on the CAMERA platform,
- installing the electronic readout modules in the VXS crate,
- installing all HV and signal cables, as well as fibers for the laser pulser,
- checking the counters for light-tightness,
- performing readout tests with cosmic particles.

The barrel of the CAMERA detector is seated on roller bearings allowing to rotate the detector on its platform by $\pm 90^\circ$ for measurements with cosmic particles. This possible rotation had to be respected also for the fixation of the cables. Therefore the cables were laid out with some additional length, forming a loop below the detector. To avoid disorder, the cables were bundled in groups of 12 each, and the bunches were loosely attached to a cross rail of the platform. While the HV cables are directly plugged into the power supply modules, the signal cables terminate at a patch panel, which is mounted above the VXS crate. This allows to use short patch cables with reduced diameter for the connections between the panel and the analog inputs of the GANDALF modules, while low-loss cables with larger diameter are used for the connection between the PMTs and the panel. The two different cable types are listed in Tab. 7.1. Additionally, the patch panel provides more flexibility in case of a potential module replacement. A photo of the patch panel and the VXS crate is found in appendix D.

Twelve GANDALF transient analyzer modules with eight channels each were placed in the VXS crate in order to read out the 96 CAMERA channels, together with two TIGER modules for the trigger processing and readout concentration tasks. For readout tests in the clean area an autonomous DAQ system consisting of a single PC and the TIGER-internal TCS controller (cf. 6.1.3) was used. Since there are several devices at the CAMERA detector which require a network connection, an Ethernet switch with optical uplink has been mounted next to the readout crate. A list of all devices involved in the CAMERA readout is given in Tab. 7.2. All in all, the number of required connections, which have to be established at the moment of installing the CAMERA detector in the experimental area, is reduced to a minimum. Besides a power connection and a multi-fiber cable, which is used to transmit Ethernet, S-Link and TCS signals, only a coax cable is needed to transmit the proton trigger signal to the COMPASS trigger barrack, which is located at the downstream end of the spectrometer.

Table 7.1: Parameters of the signal cables for the CAMERA detector. Halogen free coax cables from Huber+Suhner [182] are used.

	PMT – panel	panel – GANDALF
Length	10 m	1 m
Connectors	BNC plug / BNC jack	BNC plug / SMC jack
Type	H+S Enviroflex 400	H+S Enviroflex 316D
Impedance	50 Ω	
Signal delay	4.71 ns/m	
Screening	double braid, >70 dB (up to 5 GHz)	
Diameter	5 mm	3.16 mm
Attenuation	0.42 dB/m @ 600 MHz	0.66 dB/m @ 600 MHz

Table 7.2: List of electronic devices for the CAMERA readout.

Quantity	Device type and purpose
1	VME64x/VXS crate (WIENER UEV 6021) with power supply (UEP 6021) and fan tray (UEL 6020) backplane: Hartmann B18118233I
1	VME CPU (MEN A20) to control the GANDALF modules
6	GANDALF transient analyzers, 1 GS/s, 2 V dynamic range for the readout of the <i>Ring A</i> PMTs
6	GANDALF transient analyzers, 1 GS/s, 4 V dynamic range for the readout of the <i>Ring B</i> PMTs
2	GANDALF transient analyzers, 1 GS/s, 2 V dynamic range for the readout of the SciFi dynode signals
2	TIGER modules for trigger processor and readout concentrator tasks
1	HV power supply (CAEN SY1527) for the CAMERA PMTs
1	Gbit Ethernet switch (HP V1810-24G) with SX-LC Mini-GBIC provides network connections for VME CPU, fan tray, TIGER modules, HV supply
1	Readout PC with S-LINK spill buffer card (LDC43_1 in DAQ room)

7.1.1 Laser Calibration Runs

The CAMERA scintillator tiles are equipped with a laser pulser for calibration purposes. Short light pulses are coupled via optical fibers into the counters at their central z-positions. Just like the light from particles crossing the scintillators also the laser light is propagated to the PMTs at the upstream and downstream ends of the counters. The PMT signals (Fig. 7.1) are digitized by the GANDALF modules and the detected pulses are transmitted to the TIGER trigger processor, which performs the usual hit and track finding logic (cf. Fig. 6.5). But in contrast to actual particles, which require a finite time to fly from *Ring A* to *Ring B*, the laser pulses are coupled into the inner and the outer counter at the same instant, thus imitating the track of an "imaginary particle" with infinite speed, i.e. with a time of flight of zero.

Dedicated laser runs have been carried out every now and then, for example at times when no beam was available from the accelerator. The laser was enabled to emit pulses with a rate of approx. 8.5 kHz. The events were triggered by the TIGER trigger, with only the geometric coincidence condition (4.10) between *Ring A* and *Ring B* required. An energy loss condition (cf. 4.2.2.3) was not applied and a correlation with the beam was –of course– not required. The collected laser data allows to perform a T0-calibration for the 96 CAMERA detector channels. A set of t^0 offsets is determined such that the corrected time difference $t_{\text{diff}}^{\text{corr}}$ between upstream and downstream channel of each element is zero and the corrected mean time $t_{\text{mean}}^{\text{corr}}$ is identical on all elements:

$$t_{\text{diff},X,i}^{\text{corr}} = t_{X,i,\text{up}}^{\text{corr}} - t_{X,i,\text{dn}}^{\text{corr}} \equiv 0 \quad \wedge \quad t_{\text{mean},X,i}^{\text{corr}} = 0.5 \cdot (t_{X,i,\text{up}}^{\text{corr}} + t_{X,i,\text{dn}}^{\text{corr}}) \equiv t_{\text{mean},0} \quad (7.1)$$

$$\forall X \in \{A, B\} \text{ and } i \in \{0, 1, \dots, 23\}$$

The time values $t^{\text{corr}} = t^{\text{raw}} + t^0$ denote the raw time stamps from the GANDALF transient analyzers corrected by the T0-calibration constants.

To obtain the t^0 offsets from the laser data, histograms of the uncorrected time difference and uncorrected mean time are created for each counter and fitted with a Gaussian to determine the mean of the distributions. The results are shown in Fig. 7.2. The values from the fits are translated into t^0 offsets for the individual channels, in order to shift the time differences to zero and the mean times to a common value. The target mean value $t_{\text{mean},0}$ is defined as the smallest integer value such that all t^0 offsets are positive. This choice is beneficial for the handling of the calibration offsets in the TIGER trigger processor. The resultant offsets are tabulated in Tab. 7.3.

These t^0 offsets are written into the configuration memory of the TIGER trigger processor with the `ttconfig` tool [77, p.60]. They are used to correct the time stamps of the trigger primitives that arrive from the GANDALF modules. For each of the 96 TIGER input channels the respective calibration offset is added before the 'pmt_info' objects are stored in the 'hit_manager' buffers (cf. 6.1.2).

In addition, the extracted T0-calibration constants for the single channels are converted into calibration offsets for t_{mean} , t_{diff} and t_{ToF} . They are needed in the online monitoring software `coool` to correct the time values of various 'shift plots'. These are some basic plots that are used




Figure 7.1: Typical PMT signal for a laser pulse in a *Ring B* counter of the CAMERA detector. The plot is generated with the *GANDALF Toolbox* [163] – a GUI for parsing COMPASS raw data. 1 ADC bin $\approx$ 1 mV, 1 sample interval $\approx$ 0.989 ns




Figure 7.2: Uncorrected time difference and mean time between upstream and downstream PMTs of the CAMERA scintillator elements. Results of the Gaussian fits (mean and sigma) are plotted for each detector element. Laser data from run 108912. On counter B11 there was a problem with the laser fiber (very low amplitude of the laser signal).




Figure 7.3: CAMERA monitoring plots for run 108898 with pion beam. The plots in the left column are not calibrated; for the plots in the right column the T0-calibration constants are applied. (1) and (2) show the time-of-flight for each possible A–B counter combination. (3) and (4) show the mean time of the hits in the counters; (5) and (6) show the z position of these hits.

Table 7.3: t^0 calibration offsets for the CAMERA detector, obtained from the laser run 108912.

i	A_up	A_down	B_up	B_down
0	32.609	30.454	1.845	2.778
1	33.894	30.848	1.187	1.760
2	32.953	30.671	2.058	0.883
3	34.309	31.030	3.696	0.561
4	36.096	32.965	3.171	4.119
5	36.152	32.368	0.532	3.708
6	36.222	33.175	4.997	5.865
7	36.773	33.554	4.521	5.180
8	34.187	29.454	2.477	4.160
9	35.570	32.045	4.584	1.758
10	35.621	32.373	4.036	4.190
11	34.894	31.941	2.000	2.000
12	36.631	31.969	1.235	3.160
13	36.814	32.330	3.259	2.106
14	36.980	34.009	0.430	3.366
15	37.344	32.291	3.267	3.727
16	35.902	31.675	5.711	4.936
17	35.359	33.398	3.409	3.155
18	36.386	34.540	3.402	6.587
19	36.835	33.252	5.184	4.959
20	36.208	30.595	3.313	2.431
21	35.209	32.034	2.605	2.809
22	35.124	32.158	2.820	2.885
23	35.645	32.075	1.229	0.214

by the shift crew to observe the behavior of the CAMERA detector during the data taking. An example is given in Fig. 7.3 for the run 108898 (pion beam). The time-of-flight distribution for each A-B combination is shown in the topmost plots. The distributions of the mean time and the z position of the hits for every channel are shown in the middle and lower plots. In the left column the uncalibrated values are plotted and in the right column the values were corrected by the T0-calibration constants which have been obtained from the laser data.

7.1.2 TIGER Latency Adjustment

In order to include the CAMERA trigger in the COMPASS trigger system, the latency of the TIGER trigger processor has to be set correctly. This is due to the fact that the readout electronics is continuously buffering the last few microseconds of detector data and – upon arrival of a first-level trigger – only a small time window from the past is sent to the DAQ. For every readout module a latency setting defines, how far in the past the readout window has to start in order to capture the hits corresponding to the triggered event. The actual value for each module depends on the length of the TCS fiber, which transmits the FLT, and on the detector's position in the experiment – the further upstream, the earlier the particles pass. The COMPASS trigger system combines a multitude of trigger signals, including several scattered muon triggers (cf. 3.4), calorimeter triggers, beam triggers, the new proton trigger and a random trigger.




Figure 7.4: Measurement of the effective trigger latency for different settings [77, p.70]. The non-flat plateau is an artifact of the measuring method, which is explained in the text.

All these trigger signals with the exception of the random trigger must arrive at the TCS controller with a fixed delay with respect to the causing event. The reference time of an event is defined by the signal of the beam particle in the scintillating fiber detectors. For most of the trigger types in the COMPASS experiment the latency is adjusted by inserting delay cables to match the timings.

The TIGER trigger processor on the contrary is a digital system calculating the event timestamp based on the trigger primitives received from the GANDALF transient analyzers and buffering the trigger attempt in the release logic (cf. 6.1.2) until a configurable time span has been elapsed since the event. This latency setting is consequently adjusted in order to overlap the TIGER proton trigger signal with the other triggers. In practice a latency scan was performed using a pion beam to bring the CAMERA trigger in coincidence with the beam trigger. The latency of the TIGER release logic can be set in steps¹ of 1 ns.

The effective trigger latency, which includes some additional cable delay, has been measured for different latency settings and is shown in Fig. 7.4. The measurements (a), (b) and (c) are taken with latency settings of 500 ns, 1500 ns and 700 ns respectively. The result of a latency

¹the exact value is $(2 \cdot f_{\text{ADC}})^{-1} = 0.989$ ns

scan is plotted in (d), where the latency setting was increased in steps of 200 ns starting at 800 ns. The rectangular distribution arises from the fact, that the TIGER trigger processor releases triggers with a granularity of 1 ns [77, p.71]. The non-flat shape of the plateau is caused by the measurement setup, which was used to determine the latency. The trigger output signal of the TIGER module was connected to an additional GANDALF channel in order to obtain a time stamp of the trigger signal. The latency was calculated as the difference between the time stamp of the trigger pulse and the time stamp of the input pulse. Because the trigger output of the TIGER and the digitization of the GANDALF are operated with synchronous clocks, the trigger signal will always occur at the same phase of the sampling period. Furthermore, the interpolated time that is calculated by the GANDALF dCFD method is not uniformly distributed for small pulses [74]. This results in the excess, which is seen at the left side of the latency distributions. When measured with an oscilloscope, the plateau becomes flat.

7.1.3 The StartCounter

Since the uncertainty in the CAMERA time-of-flight is dominated by the time resolution of *Ring A*, an alternative time-of-flight between the primary vertex and *Ring B* would result in a more accurate momentum determination for the recoil protons. This requires precise knowledge of the vertex time. For this purpose, a new scintillating fiber detector has been built for the COMPASS-II upgrade – the so-called StartCounter. It is used to determine the time of the beam particles, which is then propagated to the vertex position. The improvement in the momentum resolution of recoil protons by using the StartCounter is shown in Fig. 7.5.

The StartCounter is made up of three stations (FI12Y, FI13U and FI13V), which are placed in the beam upstream of the liquid hydrogen target between the existing FI01 and FI02 detectors. The main design goals were high detection efficiency, good time resolution and high rate stability. Each station consists of two layers with an active detector area of $48 \times 48 \text{ mm}^2$. The layers are built from round scintillating fibers with a diameter of 2.5 mm, which are staggered with a pitch of 1.5 mm. Clear light guiding fibers are welded to the active fibers, forwarding the light to 16-channel multi-anode PMTs. The electrical readout is performed with GANDALF TDC modules. For the FI13 U and V stations 32 fibers per layer are connected to dedicated PMT channels, and the layers are alternately read out to the left and the right hand side. This allows to calculate a meantime between the two corresponding layers in order to minimize the variation of the hit time depending on the hit position along the fiber. The FI12Y station is constructed with triple layers, where columns of three fibers each are coupled to one PMT pad, thus yielding more light for an improved time resolution. The two layers of this station are also read out one to the left and one to the right hand side for meantiming.

Compared to the existing SciFi stations FI01 and FI02, which are built with 0.5 mm fibers and provide a time resolution in the order of 500 ps, the StartCounter exhibits a significantly better time resolution thanks to the larger fibers. First studies show that resolutions of approx. 320 ps can be expected for a single station, resulting in a combined resolution of $\sigma_{\text{SC}} < 200 \text{ ps}$ for the StartCounter [79]. Hence the StartCounter can be used to calculate a precise timestamp for the primary vertex and improve the momentum resolution of the CAMERA detector in the offline analysis.




Figure 7.5: Relative momentum resolution for recoil protons [79]. The red triangles show the resolution, when the time stamp of *Ring A* is used for the time-of-flight determination. Blue diamonds show the resolution, when the StartCounter (SC) is used to calculate a vertex time for the alternative time-of-flight determination. An improvement of up to 20% for large momenta is achieved.

7.2 Physics Data

7.2.1 First Proton Signals

The first physics data with CAMERA have been recorded with a pion beam scattered off the liquid hydrogen target. From these runs a data sample of elastic pion-proton events has been extracted in order to characterize the detector and to obtain precise calibration values, which are required for the later analysis of the pion and muon runs. For each scintillator the effective speed of light c_{eff} and a correction offset for the time difference t_{diff} between upstream and downstream pulse arrival time is determined in order to calibrate the z position (Eq. (4.2)) of the hits in the counter [79].

This method of calibration is more precise than the T0-calibration with the laser system (cf. 7.1.1) and is therefore preferred for the offline analysis. On the other hand, the calibration constants obtained from the laser data allow to almost instantly see CAMERA proton signals in the online monitoring software. Energy-loss-vs.- β_{ToF} plots like in Fig. 7.6 are available for all counters of the CAMERA detector, enabling the shift crew to verify the correct operation of the detector during the data taking. Although time-of-flight corrections and energy calibrations are not applied in the online monitoring plots, the proton band is clearly visible and the distribution is comparable to the results obtained from the Monte-Carlo simulations (Fig. 4.6). The horizontal band below the proton band at high β_{ToF} is caused by the background from δ -ray electrons.

For the offline analysis some more calibration steps are required:




Figure 7.6: Energy loss of particles crossing the CAMERA scintillators A_0 and B_0 versus their velocity β_{ToF} . Time-of-flight corrections and energy calibrations are not applied. The plot is generated with the COMPASS online monitoring software *COOOL* for a run with π^- beam.

- A time-of-flight correction is performed for each of the 48 scintillator combinations A_i-B_i and $A_i-B_{(i+1)}$. Various methods to obtain the correction offsets are compared in [79]. One possibility is to choose the offset such that the position of the cusp in the proton band (i.e. the velocity of protons which have just enough energy to cross the B scintillators) matches the value from the simulation.
- A position offset in z direction is determined for each *Ring A* scintillator element.
- A scaling factor for the energy calibration is determined for every counter, allowing to convert the raw amplitude information of the PMT pulses (given in ADC digits) into an energy loss of the particle in the scintillating material.

After applying these corrections to the raw data, an energy loss distribution for the elastic pion-proton sample is obtained as shown in Fig. 7.7. It contains the data from all *Ring B* scintillator elements, showing the energy loss of the recoil particles in the CAMERA detector versus their velocity. The proton band is nicely visible.

The pion runs were triggered by the condition " $\text{TIGER} \wedge \text{BeamTrigger} \wedge \neg(\text{Veto})$ ". In order to determine the purity of this trigger signal, a histogram of the number of tracks that could be reconstructed in the CAMERA detector for each event is shown in Fig. 7.8. It turns out that only in 3% of the events the TIGER trigger has fired spuriously although no track was found in the CAMERA detector.

7.2.2 Comparison of two CAMERA Calibration Methods

With the laser data and the elastic pion events there are two independent calibration methods available for the CAMERA detector. The results of both approaches are compared in the following. The correction offsets for t_{diff} and t_{ToF} are calculated from the two calibration constant data sets.




Figure 7.7: Energy loss for protons in *Ring B* vs. their velocity β_{ToF} (all scintillator elements). Elastic pion-proton events from run 108898. [79]




Figure 7.8: Number of tracks that could be reconstructed in the CAMERA detector for events of a pion run triggered by the TIGER trigger.

For the laser method the offsets are calculated from the t^0 constants:

$$t_{\text{diff},X,i}^{\text{corr}} = t_{X,i,\text{up}}^{\text{corr}} - t_{X,i,\text{dn}}^{\text{corr}} = t_{X,i,\text{up}}^{\text{raw}} - t_{X,i,\text{dn}}^{\text{raw}} + \underbrace{t_{X,i,\text{up}}^0 - t_{X,i,\text{dn}}^0}_{t_{\text{diff},X,i}^{\text{off,laser}}} \quad (7.2)$$

$$\begin{aligned} t_{\text{ToF},i,j}^{\text{corr}} &= \frac{t_{Bj,\text{up}}^{\text{corr}} + t_{Bj,\text{dn}}^{\text{corr}}}{2} - \frac{t_{Ai,\text{up}}^{\text{corr}} + t_{Ai,\text{dn}}^{\text{corr}}}{2} \\ &= \frac{t_{Bj,\text{up}}^{\text{raw}} + t_{Bj,\text{dn}}^{\text{raw}} - t_{Ai,\text{up}}^{\text{raw}} - t_{Ai,\text{dn}}^{\text{raw}}}{2} + \underbrace{\frac{t_{Bj,\text{up}}^0 + t_{Bj,\text{dn}}^0 - t_{Ai,\text{up}}^0 - t_{Ai,\text{dn}}^0}{2}}_{t_{\text{ToF},i,j}^{\text{off,laser}}} \end{aligned} \quad (7.3)$$

The calibration with physics events is more complex. The procedure, which is detailed in [79], yields a set of calibration constants. They are enumerated in Tab. 7.4 and their values are tabulated in the *CAMERA Helper* calibration file [79, p.110]. The *CAMERA Helper* is a utility class for the COMPASS analysis tool *PHAST* [183]. The correction offsets for t_{diff} and t_{ToF} are also calculated from the *CAMERA Helper* calibration constants, which is indicated by the index 'CH'.

The calculation of the time-of-flight offset is rather straightforward:

$$t_{\text{ToF},i,j}^{\text{corr,CH}} = \frac{t_{Bj,\text{up}}^{\text{raw}} + t_{Bj,\text{dn}}^{\text{raw}} - t_{Ai,\text{up}}^{\text{raw}} - t_{Ai,\text{dn}}^{\text{raw}}}{2} + \underbrace{t_{\text{ToF,off,global}} - t_{\text{ToF,off},i,j}}_{t_{\text{ToF},i,j}^{\text{off,CH}}} \quad (7.4)$$

However, an offset for the *up-down* time difference for the *Ring A* elements is not explicitly calculated in the *CAMERA Helper*; it is rather contained in the equation for the z position of the hits:

$$z_{A,i}^{\text{corr,CH}} = t_{\text{diff},A,i}^{\text{corr,CH}} \frac{c_{\text{eff},A,i}}{2} + z_{\text{off},A,i} = \left(t_{A,i,\text{up}}^{\text{raw}} - t_{A,i,\text{dn}}^{\text{raw}} - \underbrace{t_{\text{diff,off},A,i} + \frac{2z_{\text{off},A,i}}{c_{\text{eff},A,i}}}_{t_{\text{diff},A,i}^{\text{off,CH}}} \right) \frac{c_{\text{eff},A,i}}{2} \quad (7.5)$$

$$\text{with } z_{\text{off},A,i} = z_{\text{cal},A,i} - z_{\text{nom},A}.$$

For the *Ring B* elements no z offset is assumed, so the time difference offset is

$$t_{\text{diff},B,i}^{\text{off,CH}} = -t_{\text{diff,off},B,i}. \quad (7.6)$$

The time-difference offsets $t_{\text{diff},X,i}^{\text{off}}$ and the time-of-flight offsets $t_{\text{ToF},i,j}^{\text{off}}$ obtained from the laser data (run 108912) and from the physics data (run 108898) are compared in Fig. 7.9. Both calibrations are in good agreement. The largest deviation is found for $t_{\text{diff},B,7}^{\text{off}}$ and $t_{\text{diff},B,11}^{\text{off}}$. The reason for the discrepancy between laser and physics calibration at counter *B11* is a problem with the laser signal that is coupled into this element. Due to a broken fiber the signal amplitude was very low and almost no hits were detected, which resulted in a poor fit result (cf. Fig.

Table 7.4: Calibration constants for the *CAMERA Helper* class. The indices for the detector elements are $X \in \{A, B\}$, $i \in \{0, 1, \dots, 23\}$ and $j \in \{i, (i + 1) \bmod 24\}$.

constant name	explanation
$t_{\text{diff,off},X,i}$	offset for the time difference between upstream and downstream PMT
$t_{\text{ToF,off},i,j}$	offset for the time-of-flight between A_i and B_j
$t_{\text{ToF,off,global}}$	global time-of-flight offset
$z_{\text{cal},A,i}$	calibrated z position of the center of the <i>Ring A</i> elements
$z_{\text{nom},A}$	nominal z position of the center of the <i>Ring A</i> elements
$c_{\text{eff},X,i}$	effective speed of light in the scintillator elements

7.2). For the deviation on counter *B7* no reason could be identified. The calibrations for the remaining counters match within a few hundred picoseconds. The distributions of the deviations between laser calibration and *CAMERA Helper* calibration are shown in Fig.7.10. One may notice that the laser offsets are systematically shifted to larger values, but this is still within the RMS error of the distributions.

Summing up, the laser calibration matches the physics calibration quite well, providing a suitable set of calibration offsets for online purposes. Of course the laser method contains some uncertainties, which render impossible its usage for exclusive analyses. For example there are possible differences in the length of the laser fibers. And for the z-position of the counters and the distance between the *A* and *B* counters the nominal values are assumed, while the actual positions may deviate from these values. Therefore, the calibration method with physics events is used to achieve best possible momentum resolution for the *CAMERA* detector.

7.2.3 Muon Runs

The majority of the data that has been recorded during the DVCS test run in 2012 is with muon beam. The total number of good spills (i.e. usable spills for the analysis) is 9120 with μ^+ beam and 20197 with μ^- beam. The intensity of the negative beam is roughly one third of the intensity of the positive beam. Thus the 1:2 sharing of the beam time between μ^+ and μ^- was chosen as a tradeoff between maximizing the total number of events and balancing the statistics for the two polarities. Currently for approx. 30% of the data on tape the events are reconstructed and available for analysis. Data samples have been selected for two analysis channels: exclusive production of photons ($\mu p \rightarrow \mu' p \gamma$) and exclusive production of ρ^0 mesons ($\mu p \rightarrow \mu' p \rho^0$). First analyses are currently ongoing. To demonstrate the capability of the *CAMERA* detector, the energy-loss-vs.- β plots for the exclusive γ and ρ^0 events are shown in Fig. 7.11 and Fig. 7.12.




Figure 7.9: Comparison of the calibration results obtained with the laser method and the *CAMERA Helper* method. The time-difference and time-of-flight calibration offsets are plotted for each counter (counter combination). Both methods are in good agreement.




Figure 7.10: Deviations between the calibration offsets obtained from laser data and physics data. When excluding the deficient counter *B11* from the $t_{\text{diff},B}^{\text{off}}$ distribution, the mean is reduced to 0.33 ns with an RMS of 0.36 ns.




Figure 7.11: Energy loss ΔE_B of protons in the *Ring B* scintillators of the CAMERA detector vs. their velocity β_{ToF} for a sample of exclusive $\mu p \rightarrow \mu' p \gamma$ events, which have been recorded in the DVCS test run 2012. [184]




Figure 7.12: Energy loss ΔE_B of protons in the *Ring B* scintillators of the CAMERA detector vs. their velocity β_{ToF} for a sample of exclusive $\mu p \rightarrow \mu' p \rho^0$ events, which have been recorded in the DVCS test run 2012. [185]

8. Summary

In late 2012, the DVCS measurement campaign was started at the COMPASS-II experiment using 160 GeV/c high-intensity muon beams and a 2.5 m long liquid hydrogen target. The aim of these studies is to constrain better the GPDs, in order to gain more information about the transverse spatial parton distributions and to help with the disentanglement of the total angular momentum of the partons.

For the DVCS measurements the CAMERA recoil proton detector has been built and was commissioned successfully during a test run in 2012. It is designed to cover the liquid hydrogen target with full azimuthal angular acceptance and a polar angular acceptance from 45° to 90° . The detector comprises two ring-shaped layers, each with 24 scintillating counters, which are read out on both ends. This adds up to 96 channels in total. The PMT signals are digitized by 12 GANDALF transient analyzer modules at a sampling rate of 1 GS/s with an effective resolution of better than 10 bits. The GANDALF modules perform an online pulse-feature extraction based on a digital constant fraction discrimination, which results in a representation of the PMT pulses by arrays of time stamp, amplitude and integral values. The calculated pulse features are available within 130 ns after the signal detection in the PMT. Apart from being buffered for subsequent transmission to the DAQ, the pulse information also serve as trigger primitives for the so-called CAMERA trigger.

The concept of the CAMERA trigger system has been created within the frame of this thesis. Based on the hits in the recoil proton detector, a first-level trigger signal is generated, selecting events with a proton in the final state. At the same time, background signals from δ -ray electrons or pions are distinguished to suppress false trigger signals. For this particle identification a real-time track reconstruction and a combined analysis of the time of flight and the specific energy loss in the detector material is required. Furthermore, latency constraints due to the existing front-end electronics limit the available processing time to roughly 500 ns. In order to implement the sophisticated trigger processing algorithms with low latency, the development of a high-performance programmable logic module was mandatory.

For this reason, the TIGER module has been developed in the scope of this thesis. It constitutes the underlying hardware for the CAMERA trigger and for the connection of the readout modules to the DAQ system. It has been designed for integration into the GANDALF framework, which is based on the VME64x/VXS standard. The TIGER module takes on the role of a VXS switch board, having access to high-bandwidth backplane links from each GANDALF transient analyzer module in the crate. A high-performance FPGA device and a large on-board memory are provided for fast execution of complex algorithms and for extensive buffering capabilities. An embedded CPU is available for slow-control tasks like configuration and monitoring of the system and a GPU can support the FPGA with floating-point calculations. Apart from this, the board features interfaces for the TCS and S-LINK connection to the DAQ.

The author of this thesis has conducted all steps of the development process of the CAMERA trigger system, from concept creation to commissioning. This includes the schematic design, the component selection and the PCB layout of the TIGER board. A fundamental firmware design for the Virtex-6 FPGA has been developed, which implements the relevant logic to operate the board and access its interfaces. Finally, the hardware has been successfully put into operation and all features have been thoroughly tested and characterized.

The TIGER modules were first employed at the COMPASS-II experiment during the DVCS test run in 2012, providing key functionalities for the trigger and data acquisition systems. They were operated in two distinct firmware configurations:

- The TIGER *trigger processor* implemented the CAMERA trigger, based on the trigger primitives received from the GANDALF transient analyzer modules. It was utilized to generate first-level triggers during the pion beam runs and the laser runs for calibration purposes. During the muon runs this trigger was used in "tagging" mode to prove the trigger concept in a dedicated offline analysis.
- The TIGER *readout concentrator* was employed to multiplex the event data of all 12 GANDALF modules into a single S-LINK data stream for transmission to the DAQ system. The TCS signal was also received and distributed within the VXS crate by this module.

Based on the TIGER first-level trigger an online calibration of the CAMERA detector was performed with a laser pulser system, in order to obtain calibration constants for on-the-fly applications. The derived calibration constants are in good agreement with the values obtained from a precise offline calibration using physics events. The TIGER trigger processor made use of the online calibration constants during the treatment of the incoming trigger primitives.

The analysis of the data that has been acquired in the 2012 DVCS test run is currently in full swing and first results may be expected soon. The performance of the CAMERA detector and its readout chain is evaluated thoroughly and improvements will be made where necessary. The goal is to be well prepared for the forthcoming DVCS measurement phase after the restart of the CERN accelerator complex in 2015.

A. TIGER Board Layout

The figures A.1 to A.3 in this appendix can be used to locate components on the TIGER board by means of their reference designators:

- Fig. A.1 shows the debug LEDs on the back side of the PCB.
- Figs. A.2 and A.3 show a top and bottom view of the PCB with only the component pins and the silkscreen visible.

Fig. A.4 shows the cross section of the multi-layer PCB.




Figure A.1: Photo of the back side of the TIGER module. The debug LEDs are marked to ease locating.






Figure A.3: TIGER PCB bottom view.




Figure A.4: Cross section of the TIGER multi-layer PCB.

B. FPGA and CPLD Pin Plan

The FPGA I/O signals of the TIGER module are listed in tables B.1 to B.3, together with the pin numbers and signal standards.

Signals of the LVDS VXS backplane interface are listed in appendix C.1 in conjunction with the connector pin-out.

Signals of the LVDS VHDCI connector are listed in appendix C.2 in conjunction with the connector pin-out.

The CPLD I/O signals of the TIGER module are listed in table B.4, together with the pin numbers and signal standards.

FPGA Pin Plan

Signal	Pin	Standard
aurora_mgtclk_n	E9	MGTCLK
aurora_mgtclk_p	E10	MGTCLK
aurora_rx_n[0]	H8	MGT
aurora_rx_n[1]	G6	MGT
aurora_rx_n[2]	F8	MGT
aurora_rx_n[3]	E6	MGT
aurora_rx_p[0]	H7	MGT
aurora_rx_p[1]	G5	MGT
aurora_rx_p[2]	F7	MGT
aurora_rx_p[3]	E5	MGT
aurora_tx_n[0]	J2	MGT
aurora_tx_n[1]	H4	MGT
aurora_tx_n[2]	G2	MGT
aurora_tx_n[3]	F4	MGT
aurora_tx_p[0]	J1	MGT
aurora_tx_p[1]	H3	MGT
aurora_tx_p[2]	G1	MGT
aurora_tx_p[3]	F3	MGT
come_pcie_rx_n	AG2	MGT
come_pcie_rx_p	AG1	MGT
come_pcie_tx_n	AD4	MGT
come_pcie_tx_p	AD3	MGT
pcie_sys_clk_n	Y7	MGTCLK
pcie_sys_clk_p	Y8	MGTCLK
pcie_x8_rx_n[0]	R6	MGT
pcie_x8_rx_n[1]	U6	MGT
pcie_x8_rx_n[2]	V4	MGT
pcie_x8_rx_n[3]	W6	MGT
pcie_x8_rx_n[4]	Y4	MGT
pcie_x8_rx_n[5]	AA6	MGT
pcie_x8_rx_n[6]	AB4	MGT
pcie_x8_rx_n[7]	AC6	MGT
pcie_x8_rx_p[0]	R5	MGT
pcie_x8_rx_p[1]	U5	MGT
pcie_x8_rx_p[2]	V3	MGT
pcie_x8_rx_p[3]	W5	MGT
pcie_x8_rx_p[4]	Y3	MGT
pcie_x8_rx_p[5]	AA5	MGT
pcie_x8_rx_p[6]	AB3	MGT
pcie_x8_rx_p[7]	AC5	MGT

Signal	Pin	Standard
pcie_x8_tx_n[0]	P4	MGT
pcie_x8_tx_n[1]	R2	MGT
pcie_x8_tx_n[2]	T4	MGT
pcie_x8_tx_n[3]	U2	MGT
pcie_x8_tx_n[4]	W2	MGT
pcie_x8_tx_n[5]	AA2	MGT
pcie_x8_tx_n[6]	AC2	MGT
pcie_x8_tx_n[7]	AE2	MGT
pcie_x8_tx_p[0]	P3	MGT
pcie_x8_tx_p[1]	R1	MGT
pcie_x8_tx_p[2]	T3	MGT
pcie_x8_tx_p[3]	U1	MGT
pcie_x8_tx_p[4]	W1	MGT
pcie_x8_tx_p[5]	AA1	MGT
pcie_x8_tx_p[6]	AC1	MGT
pcie_x8_tx_p[7]	AE1	MGT
sfp_mgtclk_n	M7	MGTCLK
sfp_mgtclk_p	M8	MGTCLK
sfp1_abs	AL14	LVC MOS25
sfp1_los	AL15	LVC MOS25
sfp1_rx_n	N6	MGT
sfp1_rx_p	N5	MGT
sfp1_tx_en	AR18	LVC MOS25
sfp1_tx_fault	AR17	LVC MOS25
sfp1_tx_n	M4	MGT
sfp1_tx_p	M3	MGT
sfp2_abs	AM14	LVC MOS25
sfp2_los	AN15	LVC MOS25
sfp2_rx_n	P8	MGT
sfp2_rx_p	P7	MGT
sfp2_tx_en	AU17	LVC MOS25
sfp2_tx_fault	AT16	LVC MOS25
sfp2_tx_n	N2	MGT
sfp2_tx_p	N1	MGT

Table B.1: Multi Gigabit Transceiver (MGT) Signals

Signal	Pin	Standard
clk_ddr3_n	M12	LVDS_25
clk_ddr3_p	L12	LVDS_25
ddr3_addr[0]	G36	SSTL15
ddr3_addr[1]	H36	SSTL15
ddr3_addr[2]	A37	SSTL15
ddr3_addr[3]	D38	SSTL15
ddr3_addr[4]	C38	SSTL15
ddr3_addr[5]	B37	SSTL15
ddr3_addr[6]	B38	SSTL15
ddr3_addr[7]	A39	SSTL15
ddr3_addr[8]	A41	SSTL15
ddr3_addr[9]	C39	SSTL15
ddr3_addr[10]	H35	SSTL15
ddr3_addr[11]	A40	SSTL15
ddr3_addr[12]	F37	SSTL15
ddr3_addr[13]	J35	SSTL15
ddr3_addr[14]	B39	SSTL15
ddr3_addr[15]	E38	SSTL15
ddr3_ba[0]	E42	SSTL15
ddr3_ba[1]	B41	SSTL15
ddr3_ba[2]	F41	SSTL15
ddr3_cas_n	B42	SSTL15
ddr3_ck_n[0]	H34	DIFF_SSTL15
ddr3_ck_n[1]	F36	DIFF_SSTL15
ddr3_ck_p[0]	G34	DIFF_SSTL15
ddr3_ck_p[1]	F35	DIFF_SSTL15
ddr3_cke[0]	G41	SSTL15
ddr3_cke[1]	F42	SSTL15
ddr3_cs_n[0]	C40	SSTL15
ddr3_cs_n[1]	C41	SSTL15
ddr3_dm[0]	H39	SSTL15
ddr3_dm[1]	J40	SSTL15
ddr3_dm[2]	H30	SSTL15
ddr3_dm[3]	A32	SSTL15
ddr3_dm[4]	G33	SSTL15
ddr3_dm[5]	N29	SSTL15
ddr3_dm[6]	B22	SSTL15
ddr3_dm[7]	J21	SSTL15
ddr3_dq[0]	L37	SSTL15_T_DCI
ddr3_dq[1]	J38	SSTL15_T_DCI
ddr3_dq[2]	H41	SSTL15_T_DCI
ddr3_dq[3]	H38	SSTL15_T_DCI
ddr3_dq[4]	R29	SSTL15_T_DCI
ddr3_dq[5]	R27	SSTL15_T_DCI

Signal	Pin	Standard
ddr3_dq[6]	K38	SSTL15_T_DCI
ddr3_dq[7]	H40	SSTL15_T_DCI
ddr3_dq[8]	M34	SSTL15_T_DCI
ddr3_dq[9]	L34	SSTL15_T_DCI
ddr3_dq[10]	P31	SSTL15_T_DCI
ddr3_dq[11]	M32	SSTL15_T_DCI
ddr3_dq[12]	K42	SSTL15_T_DCI
ddr3_dq[13]	J42	SSTL15_T_DCI
ddr3_dq[14]	M33	SSTL15_T_DCI
ddr3_dq[15]	P30	SSTL15_T_DCI
ddr3_dq[16]	K29	SSTL15_T_DCI
ddr3_dq[17]	J32	SSTL15_T_DCI
ddr3_dq[18]	D37	SSTL15_T_DCI
ddr3_dq[19]	C35	SSTL15_T_DCI
ddr3_dq[20]	K30	SSTL15_T_DCI
ddr3_dq[21]	J30	SSTL15_T_DCI
ddr3_dq[22]	D36	SSTL15_T_DCI
ddr3_dq[23]	C36	SSTL15_T_DCI
ddr3_dq[24]	A34	SSTL15_T_DCI
ddr3_dq[25]	B33	SSTL15_T_DCI
ddr3_dq[26]	B32	SSTL15_T_DCI
ddr3_dq[27]	D32	SSTL15_T_DCI
ddr3_dq[28]	E35	SSTL15_T_DCI
ddr3_dq[29]	A35	SSTL15_T_DCI
ddr3_dq[30]	C33	SSTL15_T_DCI
ddr3_dq[31]	E32	SSTL15_T_DCI
ddr3_dq[32]	A36	SSTL15_T_DCI
ddr3_dq[33]	H31	SSTL15_T_DCI
ddr3_dq[34]	L30	SSTL15_T_DCI
ddr3_dq[35]	M29	SSTL15_T_DCI
ddr3_dq[36]	B36	SSTL15_T_DCI
ddr3_dq[37]	E33	SSTL15_T_DCI
ddr3_dq[38]	G32	SSTL15_T_DCI
ddr3_dq[39]	G31	SSTL15_T_DCI
ddr3_dq[40]	K32	SSTL15_T_DCI
ddr3_dq[41]	L31	SSTL15_T_DCI
ddr3_dq[42]	J37	SSTL15_T_DCI
ddr3_dq[43]	N28	SSTL15_T_DCI
ddr3_dq[44]	L32	SSTL15_T_DCI
ddr3_dq[45]	N30	SSTL15_T_DCI
ddr3_dq[46]	K33	SSTL15_T_DCI
ddr3_dq[47]	P28	SSTL15_T_DCI
ddr3_dq[48]	G23	SSTL15_T_DCI
ddr3_dq[49]	B24	SSTL15_T_DCI

Signal	Pin	Standard
ddr3_dq[50]	C23	SSTL15_T_DCI
ddr3_dq[51]	H23	SSTL15_T_DCI
ddr3_dq[52]	C24	SSTL15_T_DCI
ddr3_dq[53]	A24	SSTL15_T_DCI
ddr3_dq[54]	B23	SSTL15_T_DCI
ddr3_dq[55]	F21	SSTL15_T_DCI
ddr3_dq[56]	K22	SSTL15_T_DCI
ddr3_dq[57]	A22	SSTL15_T_DCI
ddr3_dq[58]	G21	SSTL15_T_DCI
ddr3_dq[59]	J20	SSTL15_T_DCI
ddr3_dq[60]	D22	SSTL15_T_DCI
ddr3_dq[61]	L21	SSTL15_T_DCI
ddr3_dq[62]	A21	SSTL15_T_DCI
ddr3_dq[63]	H21	SSTL15_T_DCI
ddr3_dqs_n[0]	L36	DIFF_SSTL15_T_DCI
ddr3_dqs_n[1]	N31	DIFF_SSTL15_T_DCI
ddr3_dqs_n[2]	F34	DIFF_SSTL15_T_DCI
ddr3_dqs_n[3]	F31	DIFF_SSTL15_T_DCI
ddr3_dqs_n[4]	C34	DIFF_SSTL15_T_DCI
ddr3_dqs_n[5]	K34	DIFF_SSTL15_T_DCI
ddr3_dqs_n[6]	F22	DIFF_SSTL15_T_DCI
ddr3_dqs_n[7]	E23	DIFF_SSTL15_T_DCI
ddr3_dqs_p[0]	L35	DIFF_SSTL15_T_DCI
ddr3_dqs_p[1]	M31	DIFF_SSTL15_T_DCI
ddr3_dqs_p[2]	E34	DIFF_SSTL15_T_DCI
ddr3_dqs_p[3]	F32	DIFF_SSTL15_T_DCI
ddr3_dqs_p[4]	B34	DIFF_SSTL15_T_DCI
ddr3_dqs_p[5]	K35	DIFF_SSTL15_T_DCI
ddr3_dqs_p[6]	G22	DIFF_SSTL15_T_DCI
ddr3_dqs_p[7]	E24	DIFF_SSTL15_T_DCI
ddr3_odt[0]	G37	SSTL15
ddr3_odt[1]	F40	SSTL15
ddr3_ras_n	D41	SSTL15
ddr3_reset_n	G42	LVC MOS15
ddr3_we_n	D42	SSTL15

Table B.2: DDR3 SDRAM Signals

Signal	Pin	Standard
ca[0]	AW13	LVC MOS25
ca[1]	AW12	LVC MOS25
ca[2]	BB14	LVC MOS25
ca[3]	BB13	LVC MOS25
ca[4]	AU13	LVC MOS25
ca[5]	AU12	LVC MOS25
ca[6]	AW15	LVC MOS25
ca[7]	AY15	LVC MOS25
ca[8]	AR13	LVC MOS25
ca[9]	AP13	LVC MOS25
ca[10]	AU14	LVC MOS25
ca[11]	AV15	LVC MOS25
ca[12]	AT12	LVC MOS25
ca[13]	AR12	LVC MOS25
ca[14]	BA14	LVC MOS25
ca[15]	BA15	LVC MOS25
ca[16]	AW17	LVC MOS25
ca[17]	AY17	LVC MOS25
ca[18]	AR15	LVC MOS25
ca[19]	AP15	LVC MOS25
ca[20]	BB17	LVC MOS25
ca[21]	BB16	LVC MOS25
ca[22]	AT14	LVC MOS25
cd[0]	U31	LVC MOS25
cd[1]	T31	LVC MOS25
cd[2]	AL32	LVC MOS25
cd[3]	AK32	LVC MOS25
cd[4]	R33	LVC MOS25
cd[5]	P32	LVC MOS25
cd[6]	AH33	LVC MOS25
cd[7]	AJ33	LVC MOS25
cd[8]	P33	LVC MOS25
cd[9]	N33	LVC MOS25
cd[10]	AG31	LVC MOS25
cd[11]	AH31	LVC MOS25
cd[12]	R30	LVC MOS25
cd[13]	T30	LVC MOS25
cd[14]	AF31	LVC MOS25
cd[15]	AG32	LVC MOS25
cd[16]	AD30	LVC MOS25
cd[17]	AD31	LVC MOS25
clk_ref_300_n	F14	LVDS_25
clk_ref_300_p	E14	LVDS_25
clk_si1_n	V30	LVDS_25

Signal	Pin	Standard
clk_si1_p	W30	LVDS_25
clk_si2_n	AF30	LVDS_25
clk_si2_p	AE30	LVDS_25
fcs_b	AH30	LVC MOS25
fle_b	AH29	LVC MOS25
foe_b	AJ30	LVC MOS25
fwe_b	V31	LVC MOS25
lvttlout1_fpga	AU18	LVC MOS25
lvttlout2_fpga	AT17	LVC MOS25
tmc_clk_v6_n	AY13	LVDS_25
tmc_clk_v6_p	AY14	LVDS_25
tmc_data_n	AW16	LVDS_25
tmc_data_p	AV16	LVDS_25

Table B.3: Misc. Signals

CPLD Pin Plan

Signal	Pin	Standard
BUSY	R14	LVC MOS25
CA<0>	N11	LVC MOS25
CA<1>	R16	LVC MOS25
CA<2>	P16	LVC MOS25
CA<3>	M15	LVC MOS25
CA<4>	T15	LVC MOS25
CA<5>	R12	LVC MOS25
CA<6>	R13	LVC MOS25
CA<7>	T10	LVC MOS25
CA<8>	P11	LVC MOS25
CA<9>	T14	LVC MOS25
CA<10>	R9	LVC MOS25
CA<11>	L12	LVC MOS25
CA<12>	P9	LVC MOS25
CA<13>	P10	LVC MOS25
CA<14>	R10	LVC MOS25
CA<15>	T13	LVC MOS25
CA<16>	T16	LVC MOS25
CA<17>	N8	LVC MOS25
CA<18>	T9	LVC MOS25
CA<19>	P13	LVC MOS25
CA<20>	N10	LVC MOS25
CA<21>	M9	LVC MOS25
CA<22>	L15	LVC MOS25
CCLK	N15	LVC MOS25
CD<0>	M11	LVC MOS25
CD<1>	L14	LVC MOS25
CD<2>	K14	LVC MOS25
CD<3>	T11	LVC MOS25
CD<4>	P14	LVC MOS25
CD<5>	L13	LVC MOS25
CD<6>	M13	LVC MOS25
CD<7>	J13	LVC MOS25
CD<8>	T12	LVC MOS25
CD<9>	K15	LVC MOS25
CD<10>	N14	LVC MOS25
CD<11>	M14	LVC MOS25
CD<12>	K16	LVC MOS25
CD<13>	J15	LVC MOS25
CD<14>	J14	LVC MOS25
CD<15>	R15	LVC MOS25
CD<16>	N13	LVC MOS25
CD<17>	M16	LVC MOS25
CLK_40M	P5	LVC MOS33
COMe_cb_reset_B	C4	LVC MOS33
COMe_pwr_ok	H2	LVC MOS33
COMe_pwrbtn_B	D7	LVC MOS33
COMe_sleep_B	G2	LVC MOS33
COMe_smb_alert_B	A3	LVC MOS33
COMe_smb_ck	M2	LVC MOS33
COMe_smb_dat	M1	LVC MOS33
COMe_sys_reset_B	F3	LVC MOS33
COMe_thrm_B	H1	LVC MOS33
COMe_thrmtrip_B	G1	LVC MOS33
COMe_wake0_B	F5	LVC MOS33
COMe_wake1_B	G4	LVC MOS33
COMe_wdt	E4	LVC MOS33
cpld_pg	E5	LVC MOS33
CSI_B	T7	LVC MOS25
DIN	P8	LVC MOS25
DONE	R8	LVC MOS25
FCS_B	M10	LVC MOS25
FLE_B	N16	LVC MOS25
FOE_B	M12	LVC MOS25
FWE_B	P15	LVC MOS25
gbe0_act_B	M7	LVC MOS33
gbe0_led1a	L4	LVC MOS33
gbe0_led1b	R6	LVC MOS33
gbe0_led2a	L5	LVC MOS33
gbe0_led2b	K4	LVC MOS33
gbe0_link1000_B	R5	LVC MOS33
gbe0_link100_B	P6	LVC MOS33
gbe0_link_B	N6	LVC MOS33
GPU_gpio0	J1	LVC MOS33
GPU_gpio1	K1	LVC MOS33
GPU_gpio2	L1	LVC MOS33
GPU_prsnt_l_B	E8	LVC MOS33
GPU_prsnt_r_B	B5	LVC MOS33
GPU_pwr_en	C3	LVC MOS33
GPU_pwr_good	C1	LVC MOS33
GPU_pwr_level	C2	LVC MOS33
GPU_th_alert_B	D2	LVC MOS33
GPU_th_overt_B	D1	LVC MOS33

Signal	Pin	Standard	Signal	Pin	Standard
GPU_vga_disable_B	A1	LVC MOS33	SI_LOS	N1	LVC MOS33
GPU_wake_B	B3	LVC MOS33	SI_RST_B	N2	LVC MOS33
hot1_en	A8	LVC MOS33	sus_s3_B	C7	LVC MOS33
HSWAPEN	T8	LVC MOS25	sus_s5_B	G5	LVC MOS33
i2c_ck_s0	R3	LVC MOS33	sus_stat_B	H4	LVC MOS33
i2c_dat_s0	P4	LVC MOS33	switch1	C5	LVC MOS33
ics_clkreq0_B	E6	LVC MOS33	switch2	A7	LVC MOS33
ics_clkreq1_B	D6	LVC MOS33	TMC_LOCK	P7	LVC MOS33
INIT_B	M8	LVC MOS25	TMC_RATE	M6	LVC MOS33
ldo_en	E10	LVC MOS33	ucd_gpi1	E7	LVC MOS33
ldo_pg	E9	LVC MOS33	ucd_gpi2	B1	LVC MOS33
led_cp1d1G_B	H3	LVC MOS33	ucd_gpo1	E3	LVC MOS33
led_cp1d1R_B	H5	LVC MOS33	ucd_gpo2	C6	LVC MOS33
led_cp1d2G_B	B7	LVC MOS33	ucd_gpo3	A2	LVC MOS33
led_cp1d2R_B	B8	LVC MOS33	ucd_gpo4	B4	LVC MOS33
led_gp1_B	K5	LVC MOS33	ucd_seq_2	B6	LVC MOS33
led_gp2_B	K2	LVC MOS33	ucd_tms_pgood	D4	LVC MOS33
led_gp3_B	L3	LVC MOS33	VXS_ga0_B	D3	LVC MOS33
led_gp4_B	M5	LVC MOS33	VXS_pen_B	H14	LVC MOS33
lpc_ad<0>	R4	LVC MOS33	VXS_pp_scl<1>	F12	LVC MOS33
lpc_ad<1>	N4	LVC MOS33	VXS_pp_scl<2>	B15	LVC MOS33
lpc_ad<2>	T2	LVC MOS33	VXS_pp_scl<3>	E12	LVC MOS33
lpc_ad<3>	T1	LVC MOS33	VXS_pp_scl<4>	A15	LVC MOS33
lpc_clk	M3	LVC MOS33	VXS_pp_scl<5>	H13	LVC MOS33
lpc_drq0_B	N3	LVC MOS33	VXS_pp_scl<6>	C15	LVC MOS33
lpc_frame_B	M4	LVC MOS33	VXS_pp_scl<7>	C14	LVC MOS33
lpc_serirq	P1	LVC MOS33	VXS_pp_scl<8>	B14	LVC MOS33
ltm_pg	E2	LVC MOS33	VXS_pp_scl<9>	D14	LVC MOS33
ltm_run	E1	LVC MOS33	VXS_pp_scl<10>	C13	LVC MOS33
M0	T5	LVC MOS33	VXS_pp_scl<11>	B12	LVC MOS33
M1	T4	LVC MOS33	VXS_pp_scl<12>	C12	LVC MOS33
M2	T3	LVC MOS33	VXS_pp_scl<13>	C10	LVC MOS33
main_en	A6	LVC MOS33	VXS_pp_scl<14>	D9	LVC MOS33
main_pg	A5	LVC MOS33	VXS_pp_scl<15>	A12	LVC MOS33
pex_clk_req_B	F1	LVC MOS33	VXS_pp_scl<16>	C9	LVC MOS33
pex_rst_B	F2	LVC MOS33	VXS_pp_scl<17>	B10	LVC MOS33
pex_std_sw_B	B2	LVC MOS33	VXS_pp_scl<18>	A9	LVC MOS33
pmbus_cntl	A4	LVC MOS33	VXS_pp_sda<1>	E15	LVC MOS33
PROGRAM_B	L16	LVC MOS25	VXS_pp_sda<2>	D16	LVC MOS33
RDWR_B	N9	LVC MOS25	VXS_pp_sda<3>	F14	LVC MOS33
SI2_INTR	N5	LVC MOS33	VXS_pp_sda<4>	C16	LVC MOS33
SI_DEC	R1	LVC MOS33	VXS_pp_sda<5>	E14	LVC MOS33
SI_INC	R2	LVC MOS33	VXS_pp_sda<6>	D15	LVC MOS33
SI_LOL	P2	LVC MOS33	VXS_pp_sda<7>	F13	LVC MOS33

Signal	Pin	Standard
VXS_pp_sda<8>	E13	LVC MOS33
VXS_pp_sda<9>	A14	LVC MOS33
VXS_pp_sda<10>	D13	LVC MOS33
VXS_pp_sda<11>	B13	LVC MOS33
VXS_pp_sda<12>	A13	LVC MOS33
VXS_pp_sda<13>	C11	LVC MOS33
VXS_pp_sda<14>	B9	LVC MOS33
VXS_pp_sda<15>	D11	LVC MOS33
VXS_pp_sda<16>	B11	LVC MOS33
VXS_pp_sda<17>	D10	LVC MOS33
VXS_pp_sda<18>	A11	LVC MOS33
VXS_sera	G16	LVC MOS33
VXS_serb	G15	LVC MOS33
VXS_sw_se<0>	H12	LVC MOS33
VXS_sw_se<1>	G13	LVC MOS33
VXS_sw_se<2>	G14	LVC MOS33
VXS_sw_se<3>	H16	LVC MOS33
VXS_sw_se<4>	G12	LVC MOS33
VXS_sw_se<5>	F16	LVC MOS33
VXS_sw_se<6>	E16	LVC MOS33
VXS_sw_se<7>	F15	LVC MOS33
VXS_sysfail_B	B16	LVC MOS33
VXS_sysrst_B	A16	LVC MOS33

Table B.4: CPLD I/O Signals

C. Connector Pin-Out

C.1 VXS Backplane Connectors

FPGA Pins		Link Name	Backplane Pins		
+	–		Conn.	+	–
AH39	AJ40	VXSPORT_1<0>	P4	A13	B13
AG42	AH41	VXSPORT_1<1>	P4	C14	D14
AF32	AG33	VXSPORT_1<2>	P4	A15	B15
AH34	AJ35	VXSPORT_1<3>	P4	C16	D16
AF37	AG37	VXSPORT_1<4>	P4	E13	F13
AG38	AH38	VXSPORT_1<5>	P4	G14	H14
AG36	AH36	VXSPORT_1<6>	P4	E15	F15
AJ36	AH35	VXSPORT_1<7>	P4	G16	H16
AK38	AJ38	VXSPORT_2<0>	P3	A1	B1
AK40	AL40	VXSPORT_2<1>	P3	C2	D2
AJ42	AK42	VXSPORT_2<2>	P3	A3	B3
AL42	AM42	VXSPORT_2<3>	P3	C4	D4
AH40	AJ41	VXSPORT_2<4>	P3	E1	F1
AK39	AL39	VXSPORT_2<5>	P3	G2	H2
AJ37	AK37	VXSPORT_2<6>	P3	E3	F3
AL41	AM41	VXSPORT_2<7>	P3	G4	H4
AD32	AE32	VXSPORT_3<0>	P4	A9	B9
AE33	AD33	VXSPORT_3<1>	P4	C10	D10
AE34	AE35	VXSPORT_3<2>	P4	A11	B11
AF42	AF41	VXSPORT_3<3>	P4	C12	D12
AG34	AF34	VXSPORT_3<4>	P4	E9	F9
AF40	AG41	VXSPORT_3<5>	P4	G10	H10
AF39	AG39	VXSPORT_3<6>	P4	E11	F11
AF35	AF36	VXSPORT_3<7>	P4	G12	H12

FPGA Pins		Link Name	Backplane Pins		
+	–		Conn.	+	–
AL34	AK34	VXSPORT_4<0>	P3	A5	B5
AK35	AL36	VXSPORT_4<1>	P3	C6	D6
AM34	AL35	VXSPORT_4<2>	P3	A7	B7
AN39	AM39	VXSPORT_4<3>	P3	C8	D8
AN35	AN36	VXSPORT_4<4>	P3	E5	F5
AP37	AR37	VXSPORT_4<5>	P3	G6	H6
AP36	AP35	VXSPORT_4<6>	P3	E7	F7
AR35	AT35	VXSPORT_4<7>	P3	G8	H8
AC41	AD41	VXSPORT_5<0>	P4	I5	J5
AB37	AB38	VXSPORT_5<1>	P4	K6	L6
AD36	AD35	VXSPORT_5<2>	P4	I7	J7
AE37	AD37	VXSPORT_5<3>	P4	K8	L8
AB39	AA40	VXSPORT_5<4>	P4	I1	J1
AB32	AB33	VXSPORT_5<5>	P4	K2	L2
AA42	AB42	VXSPORT_5<6>	P4	I3	J3
AA41	AB41	VXSPORT_5<7>	P4	K4	L4
AV41	AU41	VXSPORT_6<0>	P3	I13	J13
AY42	BA42	VXSPORT_6<1>	P3	K14	L14
BA40	AY40	VXSPORT_6<2>	P3	I15	J15
BA41	BB41	VXSPORT_6<3>	P3	K16	L16
AP42	AR42	VXSPORT_6<4>	P3	I9	J9
AV40	AW40	VXSPORT_6<5>	P3	K10	L10
AT42	AU42	VXSPORT_6<6>	P3	I11	J11
AT40	AU39	VXSPORT_6<7>	P3	K12	L12
AC36	AB36	VXSPORT_7<0>	P4	A5	B5
AC40	AD40	VXSPORT_7<1>	P4	C6	D6
AE38	AD38	VXSPORT_7<2>	P4	A7	B7
AD42	AE42	VXSPORT_7<3>	P4	C8	D8
AC35	AB34	VXSPORT_7<4>	P4	E5	F5
AC38	AC39	VXSPORT_7<5>	P4	G6	H6
AC34	AC33	VXSPORT_7<6>	P4	E7	F7
AE40	AE39	VXSPORT_7<7>	P4	G8	H8
AR40	AT41	VXSPORT_8<0>	P3	A9	B9
AN41	AP41	VXSPORT_8<1>	P3	C10	D10
AN38	AP38	VXSPORT_8<2>	P3	A11	B11
AR39	AT39	VXSPORT_8<3>	P3	C12	D12
AM37	AM36	VXSPORT_8<4>	P3	E9	F9
AL37	AM38	VXSPORT_8<5>	P3	G10	H10
AN40	AP40	VXSPORT_8<6>	P3	E11	F11
AW42	AW41	VXSPORT_8<7>	P3	G12	H12
AA35	Y35	VXSPORT_9<0>	P4	A1	B1
W37	Y37	VXSPORT_9<1>	P4	C2	D2
Y40	Y39	VXSPORT_9<2>	P4	A3	B3
AA36	AA37	VXSPORT_9<3>	P4	C4	D4

FPGA Pins		Link Name	Backplane Pins		
+	–		Conn.	+	–
W32	Y33	VXSPORT_9<4>	P4	E1	F1
V33	W33	VXSPORT_9<5>	P4	G2	H2
AA34	Y34	VXSPORT_9<6>	P4	E3	F3
AA32	Y32	VXSPORT_9<7>	P4	G4	H4
AW37	AW38	VXSPORT_10<0>	P3	A13	B13
AT37	AR38	VXSPORT_10<1>	P3	C14	D14
AV39	AV38	VXSPORT_10<2>	P3	A15	B15
AU34	AT34	VXSPORT_10<3>	P3	C16	D16
AU36	AT36	VXSPORT_10<4>	P3	E13	F13
AU37	AU38	VXSPORT_10<5>	P3	G14	H14
AY39	BA39	VXSPORT_10<6>	P3	E15	F15
BB39	BB38	VXSPORT_10<7>	P3	G16	H16
V40	W40	VXSPORT_11<0>	P5	A13	B13
U39	V39	VXSPORT_11<1>	P5	C14	D14
W42	Y42	VXSPORT_11<2>	P5	A15	B15
V41	W41	VXSPORT_11<3>	P5	C16	D16
W35	V35	VXSPORT_11<4>	P5	E13	F13
U42	U41	VXSPORT_11<5>	P5	G14	H14
V38	W38	VXSPORT_11<6>	P5	E15	F15
Y38	AA39	VXSPORT_11<7>	P5	G16	H16
BA37	BB37	VXSPORT_12<0>	P2	A1	B1
AY38	AY37	VXSPORT_12<1>	P2	C2	D2
AW36	AV36	VXSPORT_12<2>	P2	A3	B3
BA35	AY35	VXSPORT_12<3>	P2	C4	D4
AY34	AW35	VXSPORT_12<4>	P2	E1	F1
BB36	BA36	VXSPORT_12<5>	P2	G2	H2
BB34	BA34	VXSPORT_12<6>	P2	E3	F3
AV34	AV35	VXSPORT_12<7>	P2	G4	H4
P36	P35	VXSPORT_13<0>	P5	A9	B9
R40	T40	VXSPORT_13<1>	P5	C10	D10
U36	T36	VXSPORT_13<2>	P5	A11	B11
T41	T42	VXSPORT_13<3>	P5	C12	D12
U32	U33	VXSPORT_13<4>	P5	E9	F9
U37	U38	VXSPORT_13<5>	P5	G10	H10
W36	V36	VXSPORT_13<6>	P5	E11	F11
V34	U34	VXSPORT_13<7>	P5	G12	H12
AM26	AL26	VXSPORT_14<0>	P2	A5	B5
AT31	AU31	VXSPORT_14<1>	P2	C6	D6
BA32	AY33	VXSPORT_14<2>	P2	A7	B7
BB33	BB32	VXSPORT_14<3>	P2	C8	D8
AR33	AT32	VXSPORT_14<4>	P2	E5	F5
AV33	AW33	VXSPORT_14<5>	P2	G6	H6
AP30	AN30	VXSPORT_14<6>	P2	E7	F7
AN29	AM29	VXSPORT_14<7>	P2	G8	H8

FPGA Pins		Link Name	Backplane Pins		
+	–		Conn.	+	–
R39	P38	VXSPORT_15<0>	P5	A5	B5
M41	M42	VXSPORT_15<1>	P5	C6	D6
R35	R34	VXSPORT_15<2>	P5	A7	B7
T34	T35	VXSPORT_15<3>	P5	C8	D8
P40	P41	VXSPORT_15<4>	P5	E5	F5
R37	T37	VXSPORT_15<5>	P5	G6	H6
P42	R42	VXSPORT_15<6>	P5	E7	F7
T39	R38	VXSPORT_15<7>	P5	G8	H8
BA30	AY30	VXSPORT_16<0>	P2	A9	B9
AW31	AV31	VXSPORT_16<1>	P2	C10	D10
BA31	BB31	VXSPORT_16<2>	P2	A11	B11
AT29	AR29	VXSPORT_16<3>	P2	C12	D12
AY32	AW32	VXSPORT_16<4>	P2	E9	F9
AT30	AR30	VXSPORT_16<5>	P2	G10	H10
AW30	AV30	VXSPORT_16<6>	P2	E11	F11
AY29	BA29	VXSPORT_16<7>	P2	G12	H12
N40	N41	VXSPORT_17<0>	P5	A1	B1
M36	M37	VXSPORT_17<1>	P5	C2	D2
N36	P37	VXSPORT_17<2>	P5	A3	B3
N38	N39	VXSPORT_17<3>	P5	C4	D4
L39	L40	VXSPORT_17<4>	P5	E1	F1
L41	L42	VXSPORT_17<5>	P5	G2	H2
N35	N34	VXSPORT_17<6>	P5	E3	F3
M38	M39	VXSPORT_17<7>	P5	G4	H4
AR27	AT27	VXSPORT_18<0>	P2	A13	B13
AR28	AP28	VXSPORT_18<1>	P2	C14	D14
AN28	AM27	VXSPORT_18<2>	P2	A15	B15
AU28	AV28	VXSPORT_18<3>	P2	C16	D16
AU29	AV29	VXSPORT_18<4>	P2	E13	F13
BB29	BB28	VXSPORT_18<5>	P2	G14	H14
AW28	AY28	VXSPORT_18<6>	P2	E15	F15
AN26	AP27	VXSPORT_18<7>	P2	G16	H16
AK28	AK29	VXS_sp_rx<0>	P2	I5	J5
AU33	AU32	VXS_sp_rx<1>	P2	K6	L6
AP31	AN31	VXS_sp_rx<2>	P2	I7	J7
AL27	AM28	VXS_sp_rx<3>	P2	K8	L8
AR34	AP33	VXS_sp_tx<0>	P2	I1	J1
AN33	AN34	VXS_sp_tx<1>	P2	K2	L2
AM33	AM32	VXS_sp_tx<2>	P2	I3	J3
AP32	AR32	VXS_sp_tx<3>	P2	K4	L4

Table C.1: LVDS VXS Backplane Interface Signals

C.2 VHDCI Connector

The pin assignment of the LVDS connectors on the TIGER front panel is listed in the following table. The letter "L" or "R" in the column "Conn." denotes the connector on the left-hand side or on the right-hand side respectively.

FPGA Pins		Signal Name	Connector Pins		
+	–		Conn.	+	–
P18	P17	LVDS_IN<0>	L	1	35
G16	F16	LVDS_IN<1>	L	2	36
C15	D15	LVDS_IN<2>	L	3	37
D18	C18	LVDS_IN<3>	L	4	38
D13	E13	LVDS_IN<4>	L	5	39
M14	N14	LVDS_IN<5>	L	6	40
H14	G13	LVDS_IN<6>	L	7	41
J16	H16	LVDS_IN<7>	L	8	42
D17	E17	LVDS_IN<8>	L	9	43
E15	F15	LVDS_IN<9>	L	10	44
J12	J11	LVDS_IN<10>	L	11	45
K14	L14	LVDS_IN<11>	L	12	46
K12	L11	LVDS_IN<12>	L	13	47
J13	K13	LVDS_IN<13>	L	14	48
M16	N15	LVDS_IN<14>	L	15	49
L17	M17	LVDS_IN<15>	L	16	50
B14	C14	LVDS_IN<16>	L	17	51
C13	D12	LVDS_IN<17>	L	18	52
B18	A19	LVDS_IN<18>	L	19	53
A17	B17	LVDS_IN<19>	L	20	54
H18	G18	LVDS_IN<20>	L	21	55
F17	G17	LVDS_IN<21>	L	22	56
K18	J18	LVDS_IN<22>	L	23	57
M18	N18	LVDS_IN<23>	L	24	58
M13	N13	LVDS_IN<24>	L	25	59
J15	K15	LVDS_IN<25>	L	26	60
G19	F19	LVDS_IN<26>	L	27	61
C19	B19	LVDS_IN<27>	L	28	62
L16	L15	LVDS_IN<28>	L	29	63
K17	J17	LVDS_IN<29>	L	30	64
N16	P16	LVDS_IN<30>	L	31	65
E19	E18	LVDS_IN<31>	L	32	66
		GND	L	33	67
		GND	L	34	68

FPGA Pins		Signal Name	Connector Pins		
+	–		Conn.	+	–
AJ21	AJ20	LVDS_OUT<0>	R	1	35
AL20	AL21	LVDS_OUT<1>	R	2	36
AJ17	AK17	LVDS_OUT<2>	R	3	37
AK20	AK19	LVDS_OUT<3>	R	4	38
AL19	AM19	LVDS_OUT<4>	R	5	39
AN16	AM16	LVDS_OUT<5>	R	6	40
AK15	AK14	LVDS_OUT<6>	R	7	41
AJ16	AJ15	LVDS_OUT<7>	R	8	42
AN20	AP20	LVDS_OUT<8>	R	9	43
AU21	AT21	LVDS_OUT<9>	R	10	44
AM17	AM18	LVDS_OUT<10>	R	11	45
AP21	AP22	LVDS_OUT<11>	R	12	46
AP16	AP17	LVDS_OUT<12>	R	13	47
AN18	AN19	LVDS_OUT<13>	R	14	48
AT22	AR22	LVDS_OUT<14>	R	15	49
AM21	AN21	LVDS_OUT<15>	R	16	50
AV23	AU22	LVDS_OUT<16>	R	17	51
AV18	AV19	LVDS_OUT<17>	R	18	52
AP18	AR19	LVDS_OUT<18>	R	19	53
AY18	AW18	LVDS_OUT<19>	R	20	54
AT20	AR20	LVDS_OUT<20>	R	21	55
AV21	AW21	LVDS_OUT<21>	R	22	56
BB18	BB19	LVDS_OUT<22>	R	23	57
AY20	BA20	LVDS_OUT<23>	R	24	58
BA19	AY19	LVDS_OUT<24>	R	25	59
BA22	BA21	LVDS_OUT<25>	R	26	60
AV20	AW20	LVDS_OUT<26>	R	27	61
AW22	AY22	LVDS_OUT<27>	R	28	62
BB24	BB23	LVDS_OUT<28>	R	29	63
AW23	AY23	LVDS_OUT<29>	R	30	64
BB22	BB21	LVDS_OUT<30>	R	31	65
AY24	BA24	LVDS_OUT<31>	R	32	66
		GND	R	33	67
		GND	R	34	68

Table C.2: LVDS VHDCI Connector Signals

D. CAMERA Photos




Figure D.1: The CAMERA detector has arrived at its final position in the COMPASS experiment on September 27th, 2012. Two physicists are shown for size comparison.




Figure D.2: The CAMERA detector was assembled in the "clean area" next to the COMPASS hall.



Figure D.3: Transport of the CAMERA detector to the COMPASS hall.




Figure D.4: Left: VXS crate for the CAMERA readout with 12 GANDALF modules and 2 TIGER modules. Right: patch panel for the connection of the CAMERA PMT signals to the GANDALF modules.



Figure D.5: View into the space between *Ring A* and *Ring B* of the CAMERA detector. The orange fibers are used for the laser calibration. *Ring A* is held in place by the white strings.

Bibliography

- [1] J. D. BJORKEN, E. A. PASCHOS, 1969; *Inelastic Electron-Proton and γ -Proton Scattering and the Structure of the Nucleon*. Phys. Rev., 185:1975
- [2] R. P. FEYNMAN, 1969; *Very High-Energy Collisions of Hadrons*. Phys. Rev. Lett., 23(24):1415
- [3] EUROPEAN MUON COLLABORATION, J. ASHMAN, ET AL., 1988; *A measurement of the spin asymmetry and determination of the structure function g_1 in deep inelastic muon-proton scattering*. Physics Letters B, 206(2):364
- [4] SPIN MUON COLLABORATION, B. ADEVA, ET AL., 1998; *Spin asymmetries A_1 and structure functions g_1 of the proton and the deuteron from polarized high energy muon scattering*. Phys. Rev. D, 58:112001
- [5] P. ANTHONY, R. ARNOLD, ET AL., 1999; *Measurement of the deuteron spin structure function $g_1^d(x)$ for $1(\text{GeV}/c)^2 < Q^2 < 40(\text{GeV}/c)^2$* . Physics Letters B, 463(2–4):339
- [6] HERMES COLLABORATION, A. AIRAPETIAN, ET AL., 2007; *Precise determination of the spin structure function g_1 of the proton, deuteron, and neutron*. Phys. Rev. D, 75:012007
- [7] COMPASS COLLABORATION, C. ADOLPH, ET AL., 2013; *Leading and next-to-leading order gluon polarization in the nucleon and longitudinal double spin asymmetries from open charm muoproduction*. Phys. Rev. D, 87:052018
- [8] M. GELL-MANN, 1964; *A schematic model of baryons and mesons*. Physics Letters, 8(3):214
- [9] R. L. JAFFE, A. MANOHAR, 1990; *The g_1 problem: Deep inelastic electron scattering and the spin of the proton*. Nuclear Physics B, 337(3):509
- [10] COMPASS COLLABORATION, V. Y. ALEXAKHIN, ET AL., 2007; *The deuteron spin-dependent structure function g_1^d and its first moment*. Physics Letters B, 647(1):8
- [11] COMPASS COLLABORATION, M. ALEKSEEV, ET AL., 2010; *Quark helicity distributions from longitudinal spin asymmetries in muon-proton and muon-deuteron scattering*. Physics Letters B, 693(3):227
- [12] M. BURKARDT, C. A. MILLER, W.-D. NOWAK, 2010; *Spin-polarized high-energy scattering of charged leptons on nucleons*. Reports on Progress in Physics, 73(1):016201

- [13] A. W. THOMAS, W. WEISE, 2001; *The Structure of the Nucleon*, chap. “Deep-Inelastic Lepton Scattering”. Wiley-VCH, 1st edn.
- [14] J. C. COLLINS, D. E. SOPER, G. STERMAN, 1989; *Factorization of Hard Processes in QCD*. In *Perturbative QCD, Adv. Ser. Direct. High Energy Phys.*, vol. 5 (edited by A. H. MUELLER). World Scientific Publishing
- [15] F. HALZEN, A. D. MARTIN, 1984; *Quarks & Leptons: An Introductory Course in Modern Particle Physics*. John Wiley & Sons
- [16] C. G. CALLAN, D. J. GROSS, 1969; *High-Energy Electroproduction and the Constitution of the Electric Current*. Phys. Rev. Lett., 22:156
- [17] J. BERINGER, ET AL., 2012; *Review of Particle Physics*. Phys. Rev. D, 86:010001
- [18] U. ELSCHENBROICH. *Diagram 01-105*
<http://www-hermes.desy.de/notes/pub/trans-public-subject.html>
- [19] D. MÜLLER, D. ROBASCHIK, B. GEYER, 1994; *Wave Functions, Evolution Equations and Evolution Kernels from Light-Ray Operators of QCD*. Fortschritte der Physik, 42(2):101
- [20] A. V. RADYUSHKIN, 1997; *Nonforward parton distributions*. Phys. Rev. D, 56(9):5524
- [21] X. JI, 1997; *Gauge-Invariant Decomposition of Nucleon Spin*. Phys. Rev. Lett., 78(4):610
- [22] M. DIEHL, 2003; *Generalized parton distributions*. Physics Reports, 388(2–4):41
- [23] J. C. COLLINS, A. FREUND, 1999; *Proof of factorization for deeply virtual Compton scattering in QCD*. Phys. Rev. D, 59:074009
- [24] X. JI, J. OSBORNE, 1998; *One-loop corrections and all order factorization in deeply virtual Compton scattering*. Phys. Rev. D, 58:094018
- [25] K. GOEKE, M. V. POLYAKOV, M. VANDERHAEGHEN, 2001; *Hard exclusive reactions and the structure of hadrons*. Progress in Particle and Nuclear Physics, 47(2):401
- [26] X. JI, 1997; *Deeply virtual Compton scattering*. Phys. Rev. D, 55(11):7114
- [27] COMPASS COLLABORATION, 2010; *COMPASS-II Proposal*. Proposal, CERN, Geneva, SPS and PS Experiments Committee. CERN-SPSC-2010-014 / SPSC-P-340
- [28] M. BURKARDT, 2005; *Transverse deformation of parton distributions and transversity decomposition of angular momentum*. Phys. Rev. D, 72:094020
- [29] A. V. BELITSKY, D. MÜLLER, A. KIRCHNER, 2002; *Theory of deeply virtual Compton scattering on the nucleon*. Nuclear Physics B, 629(1–3):323
- [30] H. FISCHER. *DVCS Program at COMPASS-II*. Presentation at the ECT Trento GPD2010 Workshop, Trento, Italy, October 11-15, 2010
<http://wwwcompass.cern.ch/compass/gpd/index.html>

- [31] COMPASS COLLABORATION, 2009; *COMPASS Medium and Long Term Plans*. Letter of intent, CERN, Geneva, SPS and PS Experiments Committee. CERN-SPSC-2009-003 / SPSC-I-238
- [32] M. VANDERHAEGHEN, P. A. M. GUICHON, M. GUIDAL, 1998; *Hard Electroproduction of Photons and Mesons on the Nucleon*. Phys. Rev. Lett., 80:5064
- [33] K. KUMERIČKI, D. MÜLLER, 2010; *Deeply virtual Compton scattering at small x_B and the access to the GPD H*. Nuclear Physics B, 841(1–2):1
- [34] HERMES COLLABORATION, A. AIRAPETIAN, ET AL., 2012; *Beam-helicity and beam-charge asymmetries associated with deeply virtual Compton scattering on the unpolarised proton*. Journal of High Energy Physics, 2012(7):1
- [35] COMPASS COLLABORATION, P. ABBON, ET AL., 2007; *The COMPASS experiment at CERN*. Nucl. Instr. and Meth. A, 577(3):455
- [36] T. KUNZ, T. SZAMEITAT. *COMPASS-II Spectrometer Model*. 3D image rendered with TGEANT. Private communications
- [37] Sumitomo Cryogenics Group, Inc.;
<http://www.shicryogenics.com/>
- [38] N. DOSHITA. *LH2 target in the 2012 DVCS run*. Presentation at the COMPASS Collaboration Meeting, November 15th, 2012 at CERN
- [39] T. SZAMEITAT, A. FERRERO. Private communications
- [40] CERN CRYOGENICS GROUP. Private communications
- [41] E. BURTIN, N. D’HOSE, A. FERRERO. Private communications
- [42] H. BETHE, 1930; *Zur Theorie des Durchgangs schneller Korpuskularstrahlen durch Materie*. Annalen der Physik, 397(3):325
- [43] Saint-Gobain Crystals, 2011; *Organic Scintillation Materials Brochure*
- [44] MACHINE SHOP OF THE FREIBURG INSTITUTE OF PHYSICS. Private communications
- [45] Hamamatsu Photonics, 2012; *Photomultiplier Tube R10533 data sheet*
- [46] ET Enterprises Ltd, 2012; *9823B series data sheet*
- [47] P. JÖRG. *Time resolutions from A and B tests*. Presentation at the COMPASS GPD Meeting, July 10th, 2012 at CERN
- [48] A. GUSKOV. *ECALO in the 2012 DVCS run*. Presentation at the COMPASS Collaboration Meeting, November 15th, 2012 at CERN

- [49] N. ANFIMOV, V. ANOSOV, ET AL., 2013; *Shashlyk EM calorimeter prototype readout by MAPD with superhigh pixel density for COMPASS II*. Nucl. Inst. & Meth. A, 718(0):75 . Proceedings of the 12th Pisa Meeting on Advanced Detectors
- [50] Zecotek Photonics Inc., 2011; *Multi-pixel Avalanche Photo Diode - Enabling the future of imaging and detection*. Zecotek MAPD White Paper
- [51] N. ANFIMOV, I. CHIRIKOV-ZORIN, ET AL., 2010; *Beam test of Shashlyk EM calorimeter prototypes readout by novel MAPD with super high linearity*. Nucl. Inst. & Meth. A, 617(1–3):78
- [52] B. KETZER, A. AUSTREGESILO, ET AL., 2007; *A Triple-GEM Detector with Pixel Readout for High-Rate Beam Tracking in COMPASS*. IEEE Nuclear Science Symposium Conference Record, NSS '07, p. 242
- [53] M. ALEXEEV, F. BARBOSA, ET AL., 2012; *Detection of single photons with THickGEM-based counters*. Nucl. Inst. & Meth. A, 695:159
- [54] M. ALEXEEV, R. BIRSA, ET AL., 2013; *Development of a Thick GEM-based detector for single photon detection*. Nucl. Instr. and Meth. A. (to be published)
- [55] M. ALEXEEV, R. BIRSA, ET AL., 2013; *THGEM-based photon detectors for the upgrade of COMPASS RICH-1*. Nucl. Instr. and Meth. A, 732(0):264
- [56] B. GRUBE, 2006; *A Trigger Control System for COMPASS and a Measurement of the Transverse Polarization of Λ and Ξ Hyperons from Quasi-Real Photo-Production*. Ph.D. thesis, Technische Universität München
- [57] C. BERNET, A. BRAVAR, ET AL., 2005; *The COMPASS trigger system for muon scattering*. Nucl. Inst. & Meth. A, 550(1–2):217
- [58] L. SCHMITT, H. ANGERER, ET AL., 2004; *The DAQ of the COMPASS experiment*. IEEE Trans. Nucl. Sci., 51(3):439
- [59] Cypress Semiconductor, 2011; *HOTLink[®] Design Considerations*. AN1162
- [60] H. C. VAN DER BIJ, R. A. MCLAREN, ET AL., 1997; *S-LINK, a data link interface specification for the LHC era*. IEEE Trans. Nucl. Sci., 44(3):398
- [61] F. CARENA, W. CARENA, ET AL., 2006; *The ALICE Data-Acquisition Software Framework DATE V5*. 15th International Conference on Computing In High Energy and Nuclear Physics, Mumbai, India, p. 25
- [62] M. BODLAK, V. FROLOV, ET AL., 2013; *New data acquisition system for the COMPASS experiment*. Journal of Instrumentation, 8(02):C02009
- [63] F. HERRMANN, 2011; *Development and Verification of a High Performance Electronic Readout Framework for High Energy Physics*. Ph.D. thesis, ALU Freiburg¹, Physikalisches Institut

¹Albert-Ludwigs-Universität Freiburg

- [64] S. SCHOPFERER, 2009; *Entwicklung eines hochauflösenden Transientenrekorders*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [65] S. BARTKNECHT, H. FISCHER, ET AL., 2010; *Development of a 1GS/s high-resolution sampling ADC system*. Nucl. Inst. & Meth. A, 623(1):507
- [66] S. BARTKNECHT, H. FISCHER, ET AL., 2011; *Development and Performance Verification of the GANDALF High-Resolution Transient Recorder System*. IEEE Trans. Nucl. Sci., 58(4):1456
- [67] M. BÜCHELE, H. FISCHER, ET AL., 2012; *The GANDALF 128-Channel Time-to-Digital Converter*. Physics Procedia, 37(0):1827
- [68] M. BÜCHELE, 2012; *Entwicklung eines FPGA-basierten 128-Kanal Time-to-Digital Converter für Teilchenphysik-Experimente*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [69] M. BÜCHELE, H. FISCHER, ET AL., 2012; *A 128-channel Time-to-Digital Converter (TDC) inside a Virtex-5 FPGA on the GANDALF module*. Journal of Instrumentation, 7(03):C03008
- [70] C. MICHALSKI, 2013; *Entwicklung eines Echtzeit-Strahlprofil-Monitoring-Systems für das COMPASS-II Experiment*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [71] T. BAUMANN, M. BÜCHELE, ET AL., 2013; *The GANDALF 128-channel Time-to-Digital Converter*. Journal of Instrumentation, 8(01):C01016
- [72] J. BIELING, G. AHLUWALIA, ET AL., 2012; *Implementation of mean-timing and subsequent logic functions on an FPGA*. Nucl. Inst. & Meth. A, 672(0):13
- [73] C. SCHILL, 2012; *The ARWEN Optical Mezzanine Card for GANDALF*
<http://wwwhad.physik.uni-freiburg.de/gandalf/pages/download.php>
- [74] P. JÖRG, 2013; *Untersuchung von Algorithmen zur Charakterisierung von Photomultiplierpulsen in Echtzeit*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [75] NIST, PHYSICAL MEASUREMENT LABORATORY. *PSTAR Database for Protons*
<http://physics.nist.gov/PhysRefData/Star/Text/PSTAR.html>
- [76] T. SZAMEITAT, 2012; *Entwicklung einer Monte-Carlo-Simulation für das COMPASS-II-Experiment*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [77] M. GORZELLIK, 2013; *Entwicklung eines digitalen Triggersystems für Rückstoßprotonendetektoren*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [78] T. KUNZ, 2012; *Entwicklung einer Simulationsumgebung für das COMPASS-II-Experiment mit Geant4*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [79] R. SCHÄFER, 2013; *Charakterisierung eines Detektors zum Nachweis von Rückstoßprotonen am COMPASS Experiment*. Diploma thesis, ALU Freiburg, Physikalisches Institut

- [80] VMEbus International Trade Association (VITA), 2006; *VXS VMEbus Switched Serial Standard*. ANSI/VITA 41.0-2006
- [81] W. D. PETERSON, 1997; *The VMEbus Handbook*. VITA, 4th edn.
- [82] W-IE-NE-R Plein & Baus GmbH;
<http://www.wiener-d.com/sc/powered-crates/vxs/6u-vxs-6021.html>
- [83] Hartmann Electronic GmbH;
<http://www.hartmann-elektronik.de/en/products/vme-vme64-and-vme64x/vxs.html>
- [84] MEN Mikro Elektronik GmbH;
<http://www.men.de/products/vmebus,2,01A020-.html>
- [85] W-IE-NE-R Plein & Baus Elektronik, 2013; *Product Catalog 2013*
- [86] T. BAUMANN, 2013; *Entwicklung einer Schnittstelle zur Übertragung von Pulsinformationen eines Rückstoßdetektors an ein digitales Triggersystem*. Diploma thesis, ALU Freiburg, Physikalisches Institut
- [87] IEEE Computer Society, 1996; *IEEE Standard for Additional Mechanical Specifications for Microcomputers Using the IEEE Std 1101.1-1991 Equipment Practice*. IEEE Std 1101.10-1996
- [88] Xilinx, Inc., 2012; *Virtex-6 Family Overview*. DS150
- [89] K. T. POZNIAK, 2010; *FPGA-based, specialized trigger and data acquisition systems for high-energy physics experiments*. Measurement Science and Technology, 21(6):062002
- [90] Xilinx, Inc., 2013; *Zynq-7000 All Programmable SoC Overview*. DS190
- [91] Xilinx, Inc., 2012; *Virtex-6 FPGA Configurable Logic Block User Guide*. UG364
- [92] Xilinx, Inc., 2013; *Virtex-6 FPGA Clocking Resources User Guide*. UG362
- [93] Xilinx, Inc., 2011; *Virtex-6 FPGA Memory Resources User Guide*. UG363
- [94] Xilinx, Inc., 2011; *Virtex-6 FPGA DSP48E1 Slice User Guide*. UG369
- [95] Xilinx, Inc., 2012; *Virtex-6 FPGA Memory Interface Solutions User Guide*. UG406
- [96] SK Hynix Inc., 2010; *DDR3 SDRAM Unbuffered SODIMMs: HMT325S6BFR8C*. Data sheet
- [97] SK Hynix Inc., 2010; *DDR3 SDRAM Device Operation*
- [98] Integrated Device Technology, Inc., 2011; *ICS9DB102 Two Output Differential Buffer for PCIe Gen1 & Gen2*. Data sheet
- [99] ON Semiconductor, 2011; *NB6N11S 3.3 V 1:2 AnyLevel Input to LVDS Fanout Buffer / Translator*. Data sheet

- [100] Silicon Laboratories, 2010; *Si5326 Any Frequency Precision Clock Multiplier / Jitter Attenuator*. Data sheet
- [101] Silicon Laboratories, 2011; *Any-Frequency Precision Clocks Family Reference Manual*. Si53xx-RM
- [102] Silicon Laboratories, 2011; *Si5338 $\dot{P}$ C-Programmable Any-Frequency, Any-Output Quad Clock Generator*. Data sheet
- [103] Silicon Laboratories, 2012; *Si5335/38/51/56 ClockBuilder Desktop Software*
- [104] Xilinx, Inc., 2008; *CoolRunner-II CPLD Family*. DS090
- [105] Xilinx, Inc., 2007; *XC2C512 CoolRunner-II CPLD*. DS096
- [106] Xilinx, Inc.; *ISE Design Suite*
<http://www.xilinx.com/products/design-tools/ise-design-suite/>
- [107] Xilinx, Inc., 2010; *Platform Flash XL High-Density Configuration and Storage Device*. DS617
- [108] Xilinx, Inc., 2012; *Virtex-6 FPGA Configuration User Guide*. UG360
- [109] Xilinx, Inc., 2012; *Partial Reconfiguration User Guide*. UG702
- [110] IEEE Computer Society, 2008; *IEEE Standard Test Access Port and Boundary-Scan Architecture*. IEEE Std 1149.1-2001
- [111] International Rectifier, 2011; *IRF6718 DirectFET Power MOSFET*. Data sheet
- [112] Texas Instruments, 2010; *UCD9246 Digital PWM System Controller*. Data sheet
- [113] Texas Instruments, 2010; *PTD08A010W and PTD08A020W Digital PowerTrain Module*. Data sheets
- [114] Xilinx, Inc.; *Xilinx Power Estimator (XPE)*
http://www.xilinx.com/products/design_tools/logic_design/xpe.htm
- [115] Linear Technology, 2007; *LTM4605 High Efficiency Buck-Boost DC/DC μ Module Regulator*. Data sheet
- [116] Cadence Design Systems, Inc.; *Allegro PCB Designer*. Release 16.5
http://www.cadence.com/products/pcb/pcb_design/
- [117] M. MONTROSE, 2000; *Printed Circuit Board Design Techniques for EMC Compliance: A Handbook for Designers*. John Wiley & Sons, 2nd edn.
- [118] J. KNIGHTEN, B. ARCHAMBEAULT, ET AL., Summer 2006; *PDN Design Strategies: III. Planes and Materials – Are They Important Factors in Power Bus Design?* IEEE EMC Society Newsletter, 210:58

-
- [119] Xilinx, Inc.; *Device Models*
<http://www.xilinx.com/support/download/>
- [120] Signal Integrity Software, Inc.; *What is IBIS-AMI?*
<http://www.sissoft.com/elearning/ibis-ami.html>
- [121] PCI Industrial Computer Manufacturers Group, 2010; *COM Express Module Base Specification*. PICMG COM.0 R2.0
- [122] PCI Industrial Computer Manufacturers Group, 2009; *COM Express Carrier Design Guide*
- [123] Kontron Europe GmbH, 2012; *COMe-mTT10 User's Guide*. Rev. 120
- [124] SD Card Association, 2006; *SD Specifications Part 1: Physical Layer Simplified Specification*. Version 2.00
- [125] Intel Corp., 2012; *Intel Platform Controller Hub EG20T Datasheet*. Order Number: 324211-009US
- [126] Intel Corp., 2013; *Intel Atom Processor E6xx Series Datasheet*. Document Number: 324208-005US
- [127] Hirose Electric Co., Ltd., 2010; *FH12 Series 0.5mm and 1mm Pitch Connectors For FPC/FFC*. Part No. FH12-40S-0.5SV(55)
- [128] Kontron Europe GmbH, 2012; *ADA-LVDS-DVI 24bit User's Guide*. Rev. 123
- [129] NXP Semiconductors, 2012; *UM10204 - I²C-bus specification and user manual*. Rev. 5
- [130] Texas Instruments, 2008; *4-channel I²C and SMBus switch with reset function*. PCA9546A data sheet
- [131] SFF Committee, 2013; *Specification for Diagnostic Monitoring Interface for Optical Transceivers*. SFF-8472, Rev 11.3
- [132] Intel Corp., 2002; *Intel Low Pin Count (LPC) Interface Specification*. Document Number: 251289-001
- [133] Artec Design, 2012; *Artec Dongle II OpenCores project*. Entity: lpc_iow
http://opencores.org/project,artec_dongle_ii_fpga
- [134] MXM-SIG, 2012; *Mobile PCI Express Module Electromechanical Specification*. Version 3.1
- [135] MXM-SIG, 2009; *MXM Graphics Module Software Specification*. Version 3.0
- [136] Mercury Computer Systems, 2008; *Sensor Stream Computing Platform*. Data sheet
- [137] D. FRANKLIN. *Sensor Processing with Rugged Kepler GPUs*. Presentation at the GPU Technology Conference, May 16th, 2012 in San Jose

- [138] I. STRAZNICKY, 2012; *Ruggedization of MXM Graphics Modules*. In *High Performance Extreme Computing (HPEC), 2012 IEEE Conference on*, p. 1
- [139] MXM-SIG, 2008; *MXM Version 3.0 System Design Guide*. Revision 0.3
- [140] WOLF Industrial Systems; *MXM Products*
<http://wolf.ca/products/mxm/>
- [141] Khronos Group, 2012; *The OpenCL Specification*. Version 1.2
<http://www.khronos.org/>
- [142] PCI-SIG, 2003; *PCI Express Base Specification*. Revision 1.0a
- [143] R. BUDRUK, D. ANDERSON, T. SHANLEY, 2004; *PCI Express System Architecture*. Mind-Share, Inc, 1st edn.
- [144] PCI-SIG, 2006; *PCI Express Base Specification*. Revision 2.0
- [145] Intel Corp., 2011; *PHY Interface for the PCI Express, SATA and USB 3.0 Architectures*. Version 4.0
- [146] PLDA, 2013; *PCI Express XpressRICH3 Reference Manual*. Version 1.5.2
- [147] Mentor Graphics; *ModelSim SE-64 10.0b*
<http://www.mentor.com/products/fpga/simulation/modelsim>
- [148] PLDA, 2012; *PCI Express Bus Functional Model 3 Reference Manual*. Version 1.4.0
- [149] Xilinx, Inc., 2007; *Hot-Swapping Virtex-II, Virtex-II Pro, Virtex-4, and Virtex-5 Devices*. XAPP251
- [150] Xilinx, Inc., 2012; *Virtex-6 FPGA Data Sheet: DC and Switching Characteristics*. DS152
- [151] Xilinx, Inc.; *Will driving the I/Os of an unpowered bank cause damage to the part?*
<http://www.xilinx.com/support/answers/13428.htm>
- [152] F. HERRMANN, 2011; *GANDALF Framework User Guide*. Version 1.1
<http://wwwhad.physik.uni-freiburg.de/gandalf/pages/download.php>
- [153] Xilinx, Inc., 2012; *LogiCORE IP Aurora 64B/66B Product Guide*. PG074
- [154] SFF Committee, 2001; *Specification for SFP (Small Formfactor Pluggable) Transceiver*. INF-8074i
- [155] C. DECUSATIS, 2008; *Handbook of Fiber Optic Data Communication: A Practical Guide to Optical Networking*. Elsevier Academic Press, 3rd edn.
- [156] LEMO, 2011; *00 NIM-CAMAC & 01 Coaxial Connectors*
- [157] Honda Tsushin Kogyo, 2008; *HDRA series 0.8 mm Pitch Connectors*
- [158] Multi-Inno Technology, 2009; *OLED Module Specification – Model: MI12832DO*

- [159] P. ASHENDEN, 2008; *The Designer's Guide to VHDL*. Morgan Kaufmann, 3rd edn.
- [160] P. ASHENDEN, 2007; *Digital Design: An Embedded Systems Approach Using VHDL*. Morgan Kaufmann
- [161] Xilinx, Inc., 2010; *Virtex-6 FPGA SelectIO Resources User Guide*. UG361
- [162] B. GRUBE, I. KONOROV, L. SCHMITT, 2001. *COMPASS TCS documentation*. COMPASS note 2001-9
- [163] P. KREMSER; Diploma thesis in preparation, ALU Freiburg, Physikalisches Institut. Private communications
- [164] S. SABATINI, G. WENXUE, 2012; *PCIe SG DMA controller*
http://opencores.org/project,pcie_sg_dma
- [165] V. L. LORENZO, 2009; *Triple Port WISHBONE SPRAM Wrapper*
http://opencores.org/project,wb_3p_spram_wrapper
- [166] OpenCores, 2010; *WISHBONE System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores*
- [167] Xilinx, Inc., 2013; *ChipScope Pro Software and Cores User Guide*. UG029
- [168] Xilinx, Inc., 2013; *Partial Reconfiguration Tutorial*. UG743
- [169] T. GRUSSENMEYER; Diploma thesis in preparation, ALU Freiburg, Physikalisches Institut. Private communications
- [170] CERN HSI; *S-LINK Products*
<http://hsi.web.cern.ch/hsi/s-link/products.html>
- [171] S. HAAS, 2012; *HOLA LSC implementation with on-chip SERDES for Virtex-6*. CERN PH-ESE
- [172] Kontron Europe GmbH, 2011; *Linux BSP & KEAPI for COMe-mTT10*
<http://emdcustomersection.kontron.com/>
- [173] JLAB Data Acquisition Group, 2012; *CentOS6 Linux Diskless Setup*
https://coda.jlab.org/wiki/index.php/CentOS6_Linux_Diskless_Setup
- [174] J. CORBET, A. RUBINI, G. KROAH-HARTMAN, 2005; *Linux Device Drivers*. O'Reilly, 3rd edn.
- [175] A. K. PUGALIA, 2011; *Device Drivers, Part 9: I/O Control in Linux*
<http://www.linuxforu.com/2011/08/io-control-in-linux/>
- [176] G. MARCUS, 2006; *Linux pciDriver*. Lehrstuhl für Informatik V - Universität Mannheim

-
- [177] G. MARCUS, 2011; *MPRACE Software Library*. Institut für Technische Informatik Heidelberg
<http://li5.ziti.uni-heidelberg.de/mprace/>
- [178] J. PECH, 2011; *Re: Can't get the Xilinx cable drivers installed on SL6.1 (RHEL 6.1)*
<http://newsgroups.derkeiler.com/Archive/Comp/comp.arch.fpga/2011-09/msg00047.html>
- [179] AMD Staff, 2011; *OpenCL and the AMD APP SDK*
<http://developer.amd.com/resources/documentation-articles/articles-whitepapers/>
- [180] Solomon Systech, 2010; *SSD1306 OLED/PLED Segment/Common Driver with Controller*
- [181] L. FRIED, 2012; *SSD1306 OLED Driver Library*. Adafruit Industries
https://github.com/adafruit/Adafruit_SSD1306
- [182] HUBER+SUHNER, 2012; *RF cables catalog*
<http://www.hubersuhner.com>
- [183] COMPASS COLLABORATION; *PHysics Analysis Software Tools*
<http://ges.home.cern.ch/ges/phast/index.html>
- [184] A. FERRERO. *Exclusivity Selection with CAMERA*. Presentation at the COMPASS GPD Meeting, June 25th, 2013 at CERN
- [185] P. MALM, 2013; *Myoproduktion exklusiver ρ^0 -Mesonen an Protonen*. Master's thesis in preparation, ALU Freiburg, Physikalisches Institut. Private communications

