
Development of Time-to-Digital Converter based on Field-Programmable Gate Arrays

Maximilian Büchele



Fakultät für Mathematik und Physik
Albert-Ludwigs-Universität Freiburg

Development of Time-to-Digital Converter based on Field-Programmable Gate Arrays

Dissertation

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

vorgelegt von
Maximilian Büchele
aus
Rötenbach/Baden

Freiburg, Oktober 2017

Dekan:	Prof. Dr. Gregor Herten
Leiter der Arbeit	apl. Prof. Dr. Horst Fischer
Referent:	apl. Prof. Dr. Horst Fischer
Korreferent:	apl. Prof. Dr. Ulrich Landgraf
Tag der Verkündung des Prüfungsergebnisses:	20.12.2017

Teile dieser Arbeit wurden in folgenden Fachzeitschriften veröffentlicht:

M. Büchele *et al.*, “The ARAGORN front-end – An FPGA based implementation of a Time-to-Digital Converter”, *Conference proceeding to the 2016 Topical Workshop on Electronics for Particle Physics*, arXiv:1702.06713

Contents

List of Abbreviations	VI
1 Introduction	1
2 Theoretical Motivation	3
2.1 Deep-Inelastic Scattering	3
2.2 Inclusive DIS	5
2.2.1 Unpolarized Parton Distributions	5
2.2.2 Polarized Parton Distributions	7
2.3 Semi-Inclusive DIS	9
2.3.1 Flavour Decomposition of Helicity Distributions	10
2.3.2 Kaon Multiplicities and Fragmentation Functions	11
3 The COMPASS-II Experiment	15
3.1 The Beam Line	15
3.2 The Target Region	16
3.3 The Spectrometer	18
3.3.1 Tracking	19
3.3.2 Particle Identification	20
3.3.3 Trigger	21
3.4 The RICH-1 Detector	22
3.4.1 Hybrid Photon Detector	24
3.5 The Data Acquisition System	25
3.6 The GANDALF Framework	26
3.6.1 Mainboard	26
3.6.2 Mezzanine Cards	28

4	The ARAGORN Front-end	31
4.1	Design Overview	32
4.2	Field-Programmable Gate Array	33
4.2.1	XILINX Artix-7 FPGA	34
4.2.1.1	Configurable Logic Blocks	35
4.2.1.2	Embedded Memory	36
4.2.1.3	Digital Signal Processing	37
4.2.1.4	Clocking Resources	38
4.2.1.5	SelectIO Resources	39
4.2.1.6	GTP Transceiver Tiles	40
4.2.1.7	I/O Pin Planning	41
4.3	Power Supply	45
4.3.1	Voltage Regulators	45
4.3.2	Power Management	46
4.4	Board Configuration	48
4.5	Clocking Networks	53
4.5.1	Clock Sources	53
4.5.2	Jitter Attenuator	54
4.6	Fiber Optics	57
4.7	Interfaces	59
4.7.1	Extension Board Connectors	59
4.7.2	I ² C and PMBus	61
4.7.3	SPI and Microwire	63
4.7.4	Board LEDs and DIP Switches	64
4.7.5	JTAG Configuration and Debugging	65
4.8	PCB Design	66
4.8.1	Schematics	66
4.8.2	Layout	66
4.8.3	Simulations	69
5	Time-to-Digital Converter	73
5.1	Random Error	74
5.2	Counter-based TDC	76
5.3	Interpolating TDC	76

6	Firmware	79
6.1	TDC-FPGA Design	80
6.1.1	Multiphase Clock	80
6.1.2	TDC Core	81
6.1.3	Event Builder	88
6.1.4	Implementation	90
6.2	MERGER-FPGA Design	93
6.2.1	Constant-Latency Link	93
6.2.2	De-formatter	96
6.2.3	Data Concentrator	97
6.2.4	Embedded Processor	98
6.2.5	Analog Front-end Interface	100
6.3	Configuration Bus	102
6.4	TCS Interface	105
6.5	Project Management	105
6.6	Software Tools	107
7	Verification	109
7.1	Test Equipment	109
7.2	Test Setup	111
7.3	Test Results	113
7.3.1	Star Topology Network	113
7.3.1.1	Bit Error Ratio	113
7.3.1.2	Receiver Margin	115
7.3.2	TDC Characterization	116
7.3.2.1	Differential and Integral Non-linearity	118
7.3.2.2	Time Resolution	121
7.3.2.3	Time Interval Averaging	123
7.3.2.4	Rate Capability	125
8	Summary	127
A	PCB Layout	129
B	Connector Pinouts	133

C TDC Input Mapping **143**

Bibliography **145**

List of Abbreviations

BER	Bit Error Ratio
BRAM	Block RAM
CLB	Configurable Logic Block
CPLD	Complex Programmable Logic Device
CRC	Cyclic Redundancy Check
DAC	Digital-to-Analog Converter
DDR	Double Data Rate
DIS	Deep-Inelastic Scattering
DLL	Delay-Locked Loop
DNL	Differential Non-Linearity
DSP	Digital Signal Processing
EMCCLK	External Master Configuration Clock
FIFO	First-In First-Out
FPGA	Field-Programmable Gate Array
FSM	Finite-State Machine
HDL	Hardware Description Language
I/O	Input/Output
I²C	Inter-Integrated Circuit
IDE	Integrated Design Environment
INL	Integral Non-Linearity
IP	Intellectual Property
IPROG	Internal PROGRAM

LSB	Least Significant Bit
LUT	Look-Up Table
LVDS	Low Voltage Differential Signaling
MMCM	Mixed-Mode Clock Manager
PCB	Printed Circuit Board
PCS	Physical Coding Sublayer
PDF	Parton Distribution Function
PLL	Phase-Locked Loop
PMA	Physical Medium Attachment
PMBus	Power Management Bus
PWM	Pulse-Width Modulation
RAM	Random-Access Memory
RICH	Ring Imaging Cherenkov
RTL	Register Transfer Level
SDR	Single Data Rate
SFP	Small Form-factor Pluggable
SIDIS	Semi-Inclusive Deep-Inelastic Scattering
SPI	Serial Peripheral Interface
Tcl	Tool command language
TCS	Trigger and Control System
TDC	Time-to-Digital Converter
V_{CCO}	Output Buffer Supply Voltage
VCO	Voltage-Controlled Oscillator
VCXO	Voltage-Controlled Crystal Oscillator
VHDL	Very High Speed Integrated Circuit Hardware Description Language
XDC	XILINX Design Constraint

1. Introduction

Scattering experiments have always been an important tool in particle and nuclear physics aiming to explore the fundamental structure of matter. In the early years of the last century, H. Geiger and E. Marsden [1] observed when α particles pass through thin leafs of gold that a small fraction of the incident particles were diffusely reflected. With this discovery, the picture of the atom consisting of negatively charged corpuscles, to be identified with the electrons, within an electrically positive sphere was no longer viable. It was E. Rutherford [2] in 1911 who came up with a theory that suggested a heavy positive charge confined at the central spot in the atom to explain the observations of the large scattering angles.

Today we know that atomic nuclei consist of protons and neutrons. With the discovery that the values for the magnetic moment of nucleons deviate from the prediction for spin- $\frac{1}{2}$ Dirac particles, it became evident that nucleons cannot be considered as elementary constituents of matter [3, 4]. The form factors determined in elastic electron-nucleon scattering experiments indeed show that nucleons are rather extended objects. The challenge to resolve the nucleon substructure motivated a number of scattering experiments in the 1960s that extended further into the deep-inelastic regime. The scaling behaviour observed for the measured deep-inelastic lepton-nucleon cross-sections was considered as the first experimental evidence of point-like scattering centers within the nucleon. These findings established the parton model of the nucleon introduced by R. Feynman [5] and J. Bjorken [6] who postulated that the nucleon is constructed from point-like objects, the so-called partons, which were later identified with the quarks, anti-quarks and gluons, the mediators of the strong force. As far as we know today, the nucleon is composed of three valence quarks and gluons, accompanied by a cloud of quark-antiquark pairs, the sea quarks.

However, it is unclear to the present day how the nucleon spin is constructed from the spins and angular momenta of its constituents. Relativistic parton models predict that about 60% of the nucleon's spin is carried by the valence quarks. However, results from spin structure studies carried out for instance at CERN¹ [7, 8] and

¹Conseil Européen pour la Recherche Nucléaire

DESY¹ [9] suggest that the quark and antiquark contributions are actually only in the order of 30%. The attempt to explain the latter observation with a large gluon polarization could be excluded [10, 11]. Individual quark and antiquark helicity distributions have been extracted from semi-inclusive measurements [12]. However, these results are afflicted with large systematic uncertainties due to unsatisfactory knowledge of fragmentation functions that describe the fragmentation of quarks into hadrons. Hadron multiplicities to be extracted from semi-inclusive data recorded in 2016 and 2017 by the COMPASS-II² experiment at CERN using a liquid hydrogen target will constitute an important input to future parametrisations of quark fragmentation functions [13]. The advanced capabilities for hadron identification at COMPASS-II using a ring imaging Cherenkov detector provide sensitivity to individual quark and antiquark flavours which will allow for the extraction of some fairly known quark parton distribution and fragmentation functions. Quark parton distribution functions are key ingredients for the flavour decomposition of helicity distributions and for physics at hadron colliders [14]. In particular, there is a great interest to shed more light on the unpolarized strange quark and antiquark distribution functions at small values of the Bjorken scaling variable x_{Bj} . Upcoming analyses of the new proton data will extend to the region $0.001 < x_{\text{Bj}} < 0.2$ where they are little known [13].

With regards to the COMPASS-II programme [13], the spectrometer has undergone important hardware upgrades. Among other things, the ring imaging Cherenkov detector has been instrumented with a set of new photon detectors based on micro-pattern gas detector technologies. In the scope of this thesis, a compact hardware platform, which will be referred to as ARAGORN front-end, has been developed to address the increasing demand for high-performance time digitizers in high-energy physics experiments. This hardware is based on common field-programmable gate arrays. In the first place, the ARAGORN front-end has been designed in view of the new photon detectors installed at the ring imaging Cherenkov detector in the COMPASS-II spectrometer.

The thesis at hand is organized as follows. Chapter 2 is supposed to give a theoretical introduction to the structure of the nucleon. Here, fundamental experimental tools and important results of inclusive and semi-inclusive measurements, which motivated the upgrade of the COMPASS experiment, are summarized. Chapter 3 describes the setup of the COMPASS-II apparatus. Next the main part of this thesis starts with Chapter 4 that thoroughly describes the hardware aspects of this project. Details about the operation principle of time-to-digital converter and their characteristic parameters are given in Chapter 5 before the developed firmware designs are presented in Chapter 6. The results of this work on the project are summarized in Chapter 7.

¹Deutsches Elektronen Synchrotron

²Common Muon and Proton Apparatus for Structure and Spectroscopy

2. Theoretical Motivation

This chapter gives an introduction to the theoretical concepts and experimental techniques of spin structure studies. After a brief description of Deep-Inelastic Scattering (DIS) and Parton Distribution Functions (PDFs), important results of longitudinally polarized inclusive and semi-inclusive lepton-nucleon scattering measurements are summarized. Although the discovery that the quark and antiquark contributions to the spin of the nucleon are small dates back to the 1980s, the evaluation of the helicity distributions for the individual quark flavours is still ongoing. Knowledge of the quark fragmentation functions and efficient hadron identification, performed at COMPASS using the RICH-1 detector (see Sec. 3.4), is required for these analyses. In particular, the determination of strange quark helicity distributions is fairly challenging. Different values dependent on the choice of fragmentation functions have been observed for the strange quark polarization. The chapter concludes with a review of recent COMPASS results for charged kaon multiplicities from semi-inclusive measurements, aiming to improve the uncertainties on the quark-to-kaon fragmentation functions.

2.1 Deep-Inelastic Scattering

DIS is a fundamental tool of high-energy physics used to probe the parton structure of the nucleon. In the following, a lepton l with four-momentum k being scattered off the target nucleon N with four-momentum p is considered. At the typical center of mass energy of the lepton-nucleon system observed at the COMPASS experiment, the mediator of this reaction is a virtual photon. The Feynman diagram of this process is shown schematically in Fig. 2.1. As opposed to the elastic case, in which only the scattered lepton l' with four-momentum k' and the recoiled nucleon with four-momentum p' are observed in the final state, the invariant mass M_X of the final hadronic system X is greater than the mass M of the nucleon:

$$M_X^2 c^2 = (q + p)^2 = p^2 + 2pq + q^2 = M^2 c^2 + 2M\nu - Q^2 > M^2 c^2. \quad (2.1)$$

Due to the inequality in Eq. (2.1), the nucleon might get excited and create extra particles, for instance a π -meson. If the energy transfer is further increased, the

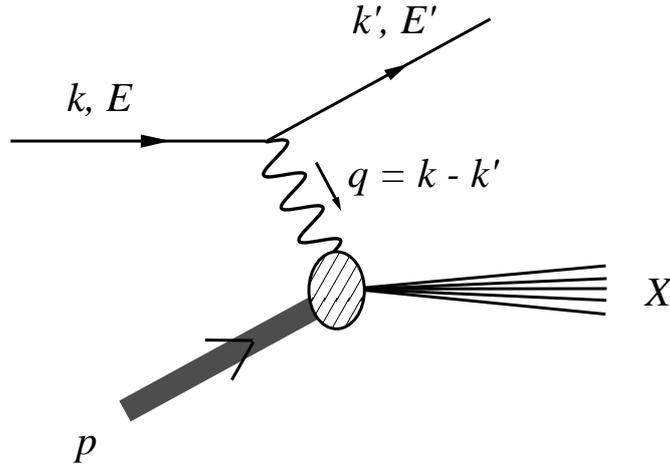


Figure 2.1: Feynman diagram of deep-inelastic lepton-nucleon scattering [15].

nucleon breaks up and fragments into hadrons. The DIS process can be measured in terms of the scattering angle θ between \vec{k} and \vec{k}' and the energy E and E' of the incoming and scattered lepton. However, the differential DIS cross-sections are usually expressed as functions of two of the Lorentz invariant variables defined as follows.

- The squared four-momentum transfer Q^2 given by the negative square of the four-momentum carried by the virtual photon:

$$Q^2 = -q^2 = -(k - k')^2 \stackrel{lab}{\approx} \frac{4EE'}{c^2} \sin^2\left(\frac{\theta}{2}\right).$$

- The energy loss of the lepton:

$$\nu := \frac{pq}{M} \stackrel{lab}{=} E - E'.$$

In this context, the scaling variable y describes the fractional energy loss:

$$y = \frac{pq}{pk} \stackrel{lab}{=} \frac{\nu}{E}.$$

- The dimensionless Bjorken scaling variable:

$$x_{\text{Bj}} := \frac{Q^2}{2pq} = \frac{Q^2}{2M\nu}.$$

The elastic case ($M_X = M$) corresponds to $x_{\text{Bj}} = 1$, while for inelastic processes ($M_X > M$) one finds $0 < x_{\text{Bj}} < 1$.

The *lab*-notation indicates the laboratory frame in which the target nucleon is at rest prior to the interaction with the incoming lepton. These equations are derived using the definitions of the four-momenta of the target nucleon and the virtual photon in the laboratory system, given by $p = (Mc, \vec{0})$ and $q = ((E - E')/c, \vec{q})$, respectively.

With regards to the kinematic variables defined above, DIS can also be referred to lepton-nucleon scattering in the limit:

$$Q^2, \nu \rightarrow \infty, x_{\text{Bj}} = \text{const.} < 1.$$

In the infinite momentum frame of the nucleon, the transverse momenta of the partons is neglected. Thus, x_{Bj} can be interpreted as the fraction of the longitudinal four-momentum of the nucleon carried by the struck quark.

2.2 Inclusive DIS

In inclusive DIS, only the scattered lepton is detected. Therefore, the inclusive cross-section accounts for all accessible hadronic final states. The differential cross-section for inclusive lepton-nucleon scattering can be written as [16, p. 12]:

$$\frac{d^2\sigma}{d\Omega dE'} = \frac{\alpha^2}{Q^4} \frac{E'}{E} L_{\mu\nu} W^{\mu\nu}, \quad (2.2)$$

where α is the electromagnetic coupling constant and $L_{\mu\nu}$, $W_{\mu\nu}$ are the leptonic and hadronic tensors associated with the vertices of the Feynman diagram shown in Fig. 2.1. While the leptonic tensor $L_{\mu\nu}$ is known from QED, structure functions are introduced to parametrize the spin-independent and spin-dependent part of the hadronic tensor $W_{\mu\nu}$.

2.2.1 Unpolarized Parton Distributions

In case of an unpolarized target nucleon, the inclusive cross-section can be written in terms of the structure functions $F_1(x_{\text{Bj}}, Q^2)$ and $F_2(x_{\text{Bj}}, Q^2)$ [17]:

$$\frac{d^2\sigma}{dx_{\text{Bj}} dy} = \frac{4\pi\alpha^2}{x_{\text{Bj}} y Q^2} \left[x_{\text{Bj}} y^2 F_1(x_{\text{Bj}}, Q^2) + \left(1 - y - \frac{x_{\text{Bj}}^2 y^2 M^2}{Q^2} \right) F_2(x_{\text{Bj}}, Q^2) \right]. \quad (2.3)$$

It was first reported by SLAC in 1968 that the cross-sections for inelastic electron and muon scattering only slightly depend on Q^2 (see Ref. [18]). Since the Fourier transform of a constant charge distribution is a δ -function, this *scaling* behaviour gave rise to the assumption that the virtual photon interacts with point-like, charged structures in the nucleon. In the parton model introduced by Feynman [5] and Bjorken [6], the cross-sections are deduced by the incoherent sum over elastic interactions between the lepton and all types of quarks and antiquarks. In first approximation, the structure function F_2 is given by:

$$F_2(x_{\text{Bj}}) = x_{\text{Bj}} \sum_f e_f^2 (q_f(x_{\text{Bj}}) + \bar{q}_f(x_{\text{Bj}})), \quad (2.4)$$

where for instance $q_f(x_{\text{Bj}})$ is the PDF for quarks of flavour f with fractional electric charge e_f . The quantity $q_f(x_{\text{Bj}}) dx_{\text{Bj}}$ yields the probability that the momentum fraction carried by the struck quark is within the interval $[x_{\text{Bj}}, x_{\text{Bj}} + dx_{\text{Bj}}]$. Another fundamental finding of the parton model is that F_1 and F_2 are connected by the Callan-Gross relation [19, p. 196]:

$$2x_{\text{Bj}} F_1(x_{\text{Bj}}) = F_2(x_{\text{Bj}}). \quad (2.5)$$

With the experimental confirmation of Eq. (2.5), the quarks were indeed identified as spin- $\frac{1}{2}$ particles.

Precision measurements of the structure function F_2 reveal small deviations from the naive parton model dependent on $\ln Q^2$, a characteristic increase for small values of x_{Bj} and a drop in the higher x_{Bj} region. Figure 2.2 exemplarily shows the combined results for the proton structure function $F_2^p(x_{Bj}, Q^2)$ from various collider and fixed target experiments as function of Q^2 for fixed values of x_{Bj} . The scaling violations can be explained as follows. The reduced wavelength of the virtual photon scales with $1/Q^2$. With increasing resolution, at high x_{Bj} , the quark distribution is shifted due to gluon radiation toward smaller values of x_{Bj} , whereas for small values of x_{Bj} the structure function rises since the virtual photon resolves, besides the valence quarks, more and more $q\bar{q}$ -pairs from gluon conversion. Scaling violations are also predicted by QCD through the Altarelli-Parisi equations [20] that describe the derivative of the quark and gluon distribution functions with respect to $\ln Q^2$. The QCD fits on the Q^2 -evolution of F_2 allow for a determination of the strong coupling constant $\alpha_s(Q^2)$ and provide an estimate of the gluon distribution function.

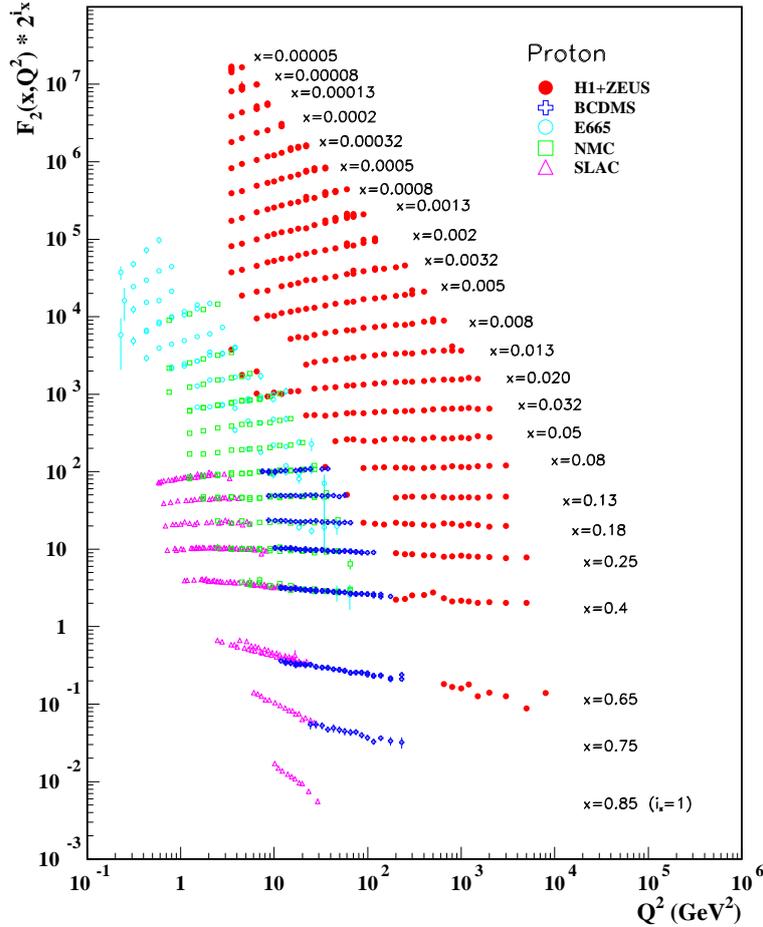


Figure 2.2: The proton structure function F_2^p as function of Q^2 shown in bins of fixed x_{Bj} obtained from various experiments [17]. For better comprehensibility, the results for F_2^p have been multiplied by 2^{i_x} , where i_x denotes the number of the x -bin, ranging from $i_x = 1$ ($x = 0.85$) to $i_x = 24$ ($x = 0.00005$).

2.2.2 Polarized Parton Distributions

In analogy to the unpolarized case, the asymmetric part of the hadronic tensor $W_{\mu\nu}$ in Eq. (2.2) is parametrized by the spin-dependent structure functions $g_1(x_{\text{Bj}}, Q^2)$ and $g_2(x_{\text{Bj}}, Q^2)$. The spin-dependent structure functions are measured using a longitudinally polarized lepton beam scattered off a longitudinally polarized nucleon target. The difference of the differential cross-sections for incoming leptons polarized anti-parallel (-) and target polarizations parallel (+) or anti-parallel (-) with respect to the beam direction is given by [21]:

$$\frac{d^3\sigma^{-+}}{dx_{\text{Bj}}dyd\phi} - \frac{d^3\sigma^{--}}{dx_{\text{Bj}}dyd\phi} = \frac{4\alpha^2}{Q^2} \left[\left(2 - y - \frac{\gamma^2 y^2}{2} \right) g_1(x_{\text{Bj}}, Q^2) - \gamma^2 y^2 g_2(x_{\text{Bj}}, Q^2) \right]. \quad (2.6)$$

In the quark parton model, the spin-dependent structure function $g_1(x_{\text{Bj}})$ is written as [22]:

$$g_1(x_{\text{Bj}}) = \frac{1}{2} \sum_f e_f^2 (\Delta q_f(x_{\text{Bj}}) + \Delta \bar{q}_f(x_{\text{Bj}})), \quad (2.7)$$

where e_f denotes the fractional electric charge carried by the struck quark of flavour f and

$$\Delta q_f(x_{\text{Bj}}) = q_f^{\rightarrow}(x_{\text{Bj}}) - q_f^{\leftarrow}(x_{\text{Bj}})$$

are the helicity distributions, the difference of the number densities of quarks with helicity parallel $q_f^{\rightarrow}(x_{\text{Bj}})$ and anti-parallel $q_f^{\leftarrow}(x_{\text{Bj}})$ to the helicity of the target nucleon. The helicity distributions for antiquarks are defined accordingly. The first moment of the helicity distributions:

$$\Delta q_f = \int_0^1 \Delta q_f(x_{\text{Bj}}) dx_{\text{Bj}}, \quad (2.8)$$

is linked to the sum rule for the spin structure of the nucleon introduced by Jaffe and Manohar [23]:

$$\frac{1}{2} = \frac{1}{2} \sum_f (\Delta q_f + \Delta \bar{q}_f) + \Delta g + L_q + L_g, \quad (2.9)$$

who proposed that the nucleon spin can be decomposed into contributions from all quark Δq_f and antiquark $\Delta \bar{q}_f$ flavours, the spin content of the gluons Δg and their respective orbital momentum contributions L_q and L_g . Just as for the unpolarized case, the spin-dependent structure functions depend logarithmically on Q^2 . Hence, the same conclusions may be drawn with regard to scaling violations. The spin-dependent structure function g_1 has been measured by various experiments at CERN, DESY, JLab and SLAC. The CERN experiments use a longitudinally polarized muon beam with momenta up to 200 GeV/c scattered off a solid-state target segmented into two or three cells with opposite polarisation. The analysis is based on the polarized DIS cross-section asymmetry rather than the difference given by Eq. (2.6), since important quantities like the unpolarized cross-section, the beam

flux, the number of target nuclei and the spectrometer acceptance cancel out in the asymmetry. The virtual photon asymmetry [24]:

$$A_1 = \frac{d\sigma_{1/2} - d\sigma_{3/2}}{d\sigma_{1/2} + d\sigma_{3/2}} = \frac{g_1 - \gamma^2 g_2}{F_1} \rightarrow \frac{g_1}{F_1}, \quad (2.10)$$

is of particular interest in spin physics. Here, $d\sigma_{1/2}$ and $d\sigma_{3/2}$ are the virtual photoabsorption cross-sections provided that the projection of the total angular momentum of the photon-nucleon system along the direction of the incoming lepton is $1/2$ and $3/2$, respectively. Taking into account that g_2 is suppressed by $\gamma^2 = Q^2/\nu^2$, the asymmetry A_1 corresponds to the ratio of the polarized and unpolarized structure functions g_1 and F_1 . The structure function g_2 is only accessible with a transversely polarized target. The virtual photon asymmetry A_1 can be related via the optical theorem to the longitudinal asymmetry $A_{||}$, which is the experimental observable, by the relationship [24]:

$$A_{||} = \frac{d\sigma^{-+} - d\sigma^{--}}{d\sigma^{-+} + d\sigma^{--}} \approx DA_1,$$

where $d\sigma$ is short for $\frac{d^3\sigma}{dx_{Bj}dyd\phi}$ and D is referred to the depolarization factor of the virtual photon, to be found for instance in Ref. [24]. The world data of the spin-dependent structure function $g_1(x_{Bj}, Q^2)$ as extracted from asymmetry measurements according to Eq. (2.10) is illustrated in Figs. 2.3 and 2.4.

The parton contributions to the nucleon spin defined in Eq. (2.8) can be derived from the first moment of g_1 . In leading order QCD, the first moment $\Gamma_1^p(Q^2)$ of $g_1(x_{Bj}, Q^2)$ for the proton can be written as [22]:

$$\Gamma_1^p(Q^2) = \int_0^1 g_1(x_{Bj}, Q^2) dx_{Bj} = \frac{1}{12} \left(a_3 + \frac{1}{3} a_8 \right) + \frac{1}{9} a_0, \quad (2.11)$$

with the three axial charges:

$$\begin{aligned} a_3 &= \Delta u + \Delta \bar{u} - \Delta d - \Delta \bar{d}, \\ a_8 &= \Delta u + \Delta \bar{u} + \Delta d + \Delta \bar{d} - 2(\Delta s + \Delta \bar{s}), \\ a_0 &\equiv \Delta \Sigma = \sum_f (\Delta q_f + \Delta \bar{q}_f) = a_8 + 3(\Delta s + \Delta \bar{s}). \end{aligned}$$

The isovector charge a_3 corresponds to the weak coupling constant $|g_A/g_V|$ obtained from the neutron β -decay. The octet charge a_8 can be extracted from hyperon β -decays assuming SU(3) flavour symmetry. The flavour-singlet charge a_0 is identical with the sum $\Delta \Sigma$ of the quark and antiquark contributions. In contrast to a_3 and a_8 , the singlet a_0 becomes Q^2 -dependent in higher orders of QCD. Under the assumption of a vanishing strange sea polarization ($\Delta s = \Delta \bar{s} = 0$), the octet and singlet charge are identical and the constraints on a_3 and a_8 predict [22]:

$$\Gamma_{1,EJ}^p \simeq 0.185, \quad (2.12)$$

which is known as the Ellis-Jaffe sum rule [25]. The EMC experiment at CERN obtained a somewhat smaller value for Γ_1^p than the naive prediction in Eq. (2.12).

From the EMC result, it was concluded that $\Delta\Sigma$ is significantly reduced due to a negative polarization of the strange sea. The finding that the quarks marginally account for the spin of the proton is referred to as the *spin puzzle* that gave birth to a number of polarized DIS experiments. Recent COMPASS results for the first moment of the deuteron spin-dependent structure function g_1^d point to a negative strange quark contribution of $\Delta s + \Delta\bar{s} = -0.08 \pm 0.01 \pm 0.02$ [26] and thus confirm the EMC conclusions.

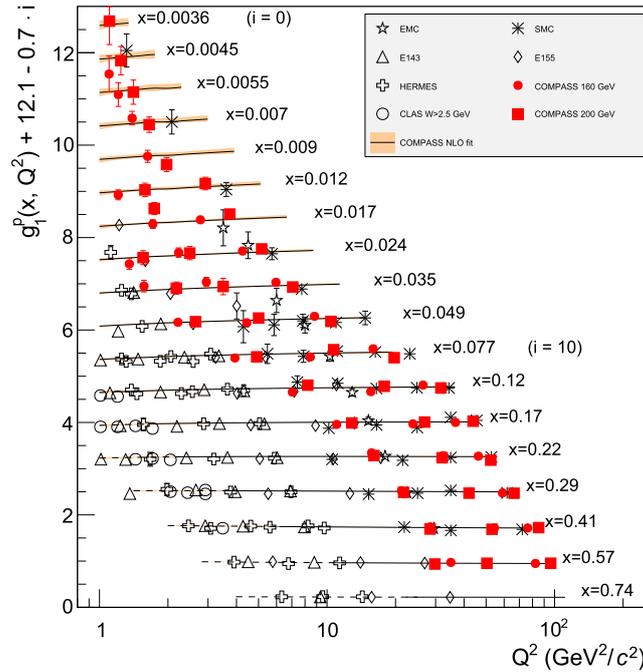


Figure 2.3: The spin-dependent structure function g_1^p of the proton as function of x_{Bj} and Q^2 . For clarity, the data of the different x_{Bj} -bins are offset. The lines correspond to a NLO-QCD fit [7].

2.3 Semi-Inclusive DIS

Polarized Semi-Inclusive Deep-Inelastic Scattering (SIDIS) measurements gain access to the individual quark, antiquark and gluon helicity distributions. Contrary to inclusive DIS, which tags only the scattered lepton, SIDIS measures additional final state hadrons. For the flavour decomposition of the helicity distributions mainly high-energy charged pions and kaons are of interest. For instance, an up quark most likely fragments into a π^+ and a down quark into a π^- . Analogically, a kaon in the final state preferentially originates from a strange quark. Similar to inclusive DIS, cross-section asymmetries for the production of hadrons h can be defined, which are written in LO-QCD as [12]:

$$A_1^h(x_{Bj}, Q^2, z) = \frac{\sum_f e_f^2 (\Delta q_f(x_{Bj}, Q^2) D_f^h(z, Q^2) + \Delta \bar{q}_f(x_{Bj}, Q^2) D_{\bar{f}}^h(z, Q^2))}{\sum_f e_f^2 (q_f(x_{Bj}, Q^2) D_f^h(z, Q^2) + \bar{q}_f(x_{Bj}, Q^2) D_{\bar{f}}^h(z, Q^2))}, \quad (2.13)$$

where $\Delta q(x_{Bj}, Q^2)$ and $q(x_{Bj}, Q^2)$ are the polarized and unpolarized PDFs. The probability for quarks and antiquarks of flavour f to produce a hadron h that car-

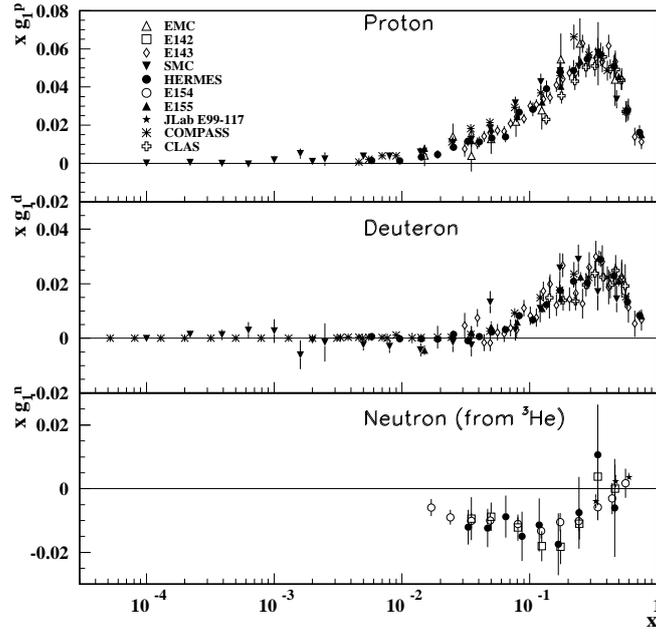


Figure 2.4: World data of the spin-dependent structure function $x_{Bj}g_1(x_{Bj})$ of the proton, deuteron, and neutron as obtained by different experiments in polarized deep-inelastic scattering [17].

ries the fractional energy $z = E_h/\nu$ in the target rest frame is described by the fragmentation functions D_f^h and $D_{\bar{f}}^h$.

2.3.1 Flavour Decomposition of Helicity Distributions

The fragmentation functions cancel out in the *difference asymmetry* that is derived from the difference of cross-sections for positive and negative hadrons [16, p. 140]:

$$A^{h^+-h^-} = \frac{(\sigma_{\uparrow\downarrow}^{h^+} - \sigma_{\uparrow\downarrow}^{h^-}) - (\sigma_{\uparrow\uparrow}^{h^+} - \sigma_{\uparrow\uparrow}^{h^-})}{(\sigma_{\uparrow\downarrow}^{h^+} - \sigma_{\uparrow\downarrow}^{h^-}) + (\sigma_{\uparrow\uparrow}^{h^+} - \sigma_{\uparrow\uparrow}^{h^-})}. \quad (2.14)$$

Under the assumption $\Delta s = \Delta \bar{s}$, the difference asymmetries for the deuteron target $A_d^{h^+-h^-}$ allow for a direct determination of the polarized valence distributions [16, p. 140]:

$$A_d^{h^+-h^-} \equiv A_d^{\pi^+-\pi^-} = A_d^{K^+-K^-} = \frac{\Delta u_v + \Delta d_v}{u_v + d_v}, \quad (2.15)$$

where $\Delta q_v \equiv \Delta q - \Delta \bar{q}$. The unpolarized PDFs $u_v + d_v$ are extracted from the structure function F_2 . Hadron identification is not required since pions and kaons equally contribute to the difference asymmetry. Semi-inclusive asymmetries were measured by the SMC and HERMES experiments and at JLab and BNL, and are until today studied by the COMPASS collaboration. Figure 2.5 shows the first moment of the polarised valence distribution:

$$\Gamma_v(x_{\min}) = \int_{x_{\min}}^{0.7} (\Delta u_v(x_{Bj}) + \Delta d_v(x_{Bj})) dx_{Bj}, \quad (2.16)$$

as function of the lower integration limit x_{\min} , extracted from COMPASS difference asymmetry data. The final COMPASS result at $Q^2 = 10 \text{ (GeV/c)}^2$ [27]:

$$\Gamma_v(0.006 < x_{Bj} < 0.7) = 0.40 \pm 0.07 \pm 0.06,$$

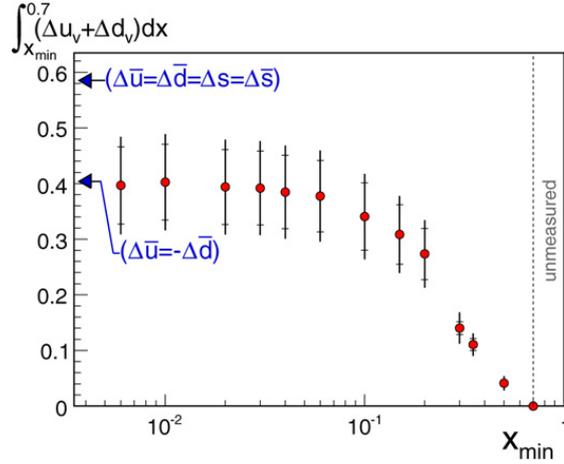


Figure 2.5: The first moment of $\Delta u_v(x_{Bj}) + \Delta d_v(x_{Bj})$ derived from COMPASS asymmetry data for the deuteron target as function of the lower integration limit x_{\min} . The arrows indicate the theoretical expectations for a flavour symmetric polarised sea and a non-symmetric polarization of the light sea quarks, respectively [27].

is two standard deviations below the value of the octet charge $a_8 = 0.58 \pm 0.03$. The assumption of a symmetric polarized sea ($\Delta \bar{u} = \Delta \bar{d} = \Delta s = \Delta \bar{s}$), which requires Γ_v to be equal to a_8 , is thus disfavoured.

A full flavour decomposition of the helicity distributions for the lightest quarks and antiquarks has also been performed by the COMPASS collaboration. Hadron identification using the RICH-1 detector plays a key role in this analysis. The quark helicity distributions extracted from double-spin asymmetries according to Eq. (2.13) for the production of identified charged pions and kaons are shown in Fig. 2.6. The data points of the up quark distribution $\Delta u(x_{Bj})$ are positive and the values of the down quark distribution $\Delta d(x_{Bj})$ are negative over the measured x_{Bj} range. The polarization of the sea quarks $\Delta \bar{u}(x_{Bj})$ and $\Delta \bar{d}(x_{Bj})$ is found to be small. While $\Delta \bar{u}(x_{Bj})$ is compatible with zero, $\Delta \bar{d}(x_{Bj})$ tends to be negative, which leads to a slightly positive flavour asymmetry of the sea $\Delta \bar{u}(x_{Bj}) - \Delta \bar{d}(x_{Bj})$. Surprisingly, the strange quark and antiquark helicity distributions $\Delta s(x_{Bj})$ and $\Delta \bar{s}(x_{Bj})$ are flat and both consistent with zero, contrary to the negative value of the strange quark polarization derived from the first moment of the structure function g_1^d . As will be pointed out below, the cause for this issue may be linked to the uncertainties on the strange quark-to-kaon FFs.

2.3.2 Kaon Multiplicities and Fragmentation Functions

A significant input to the analyses of the quark-to-kaon fragmentation functions comes from measurements of kaon multiplicities. In general, the differential multiplicity for charged hadrons h observed in unpolarized SIDIS measurements can be written as the semi-inclusive cross-section for charged hadron production normalized to the inclusive DIS cross-section [28]:

$$\frac{dM^h(x_{Bj}, z, Q^2)}{dz} = \frac{d^3\sigma^h(x_{Bj}, z, Q^2)/dx_{Bj}dQ^2dz}{d^2\sigma^{\text{DIS}}/dx_{Bj}dQ^2}. \quad (2.17)$$

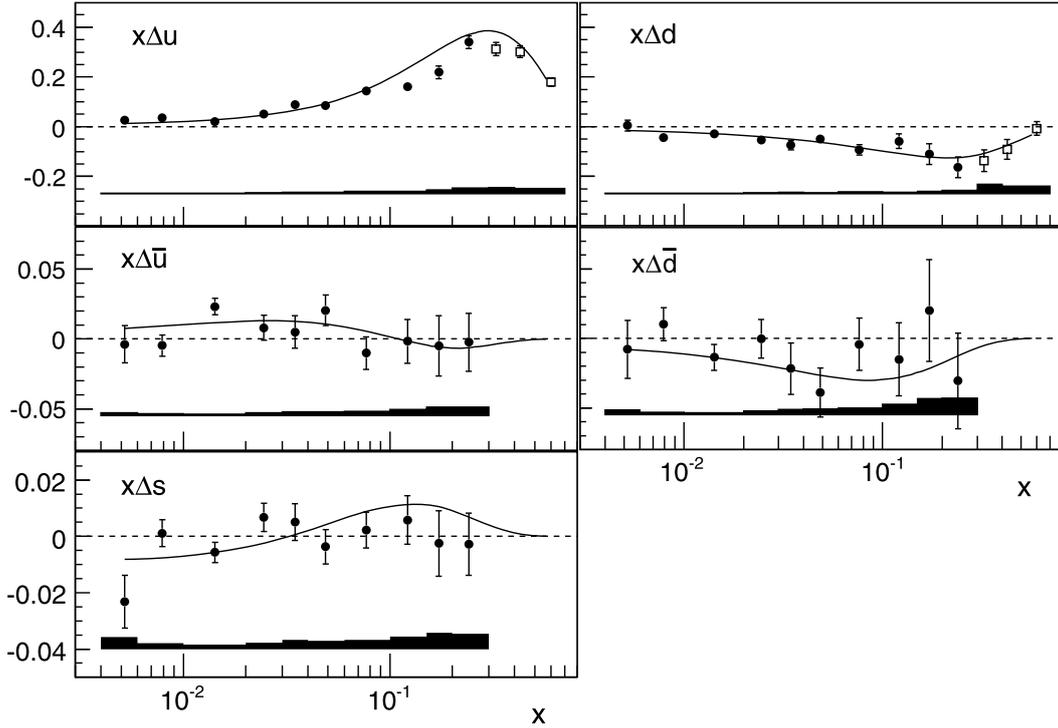


Figure 2.6: The quark helicity distributions at $Q^2 = 3 \text{ (GeV}/c)^2$ as function of x_{Bj} extracted from COMPASS asymmetry data for the proton and deuteron target [12]. The data points were derived from a LO analysis using the DSS fragmentation functions [29]. The curves represent the results of the NLO-DSSV fit [30].

When integrated over z , one obtains the mean number of hadrons h per DIS event. In LO-QCD, Eq. (2.17) takes the form:

$$\frac{dM^h(x_{Bj}, z, Q^2)}{dz} = \frac{\sum_f e_f^2 (q_f(x_{Bj}, Q^2) D_f^h(z, Q^2) + \bar{q}_f(x_{Bj}, Q^2) D_{\bar{f}}^h(z, Q^2))}{\sum_f e_f^2 (q_f(x_{Bj}, Q^2) + \bar{q}_f(x_{Bj}, Q^2))}, \quad (2.18)$$

where $q_f(x_{Bj}, Q^2)$ and $\bar{q}_f(x_{Bj}, Q^2)$ are the unpolarized PDFs for quarks and anti-quarks of flavour f . The fragmentation functions $D_f^h(z, Q^2)$ and $D_{\bar{f}}^h(z, Q^2)$ are defined in analogy to Eq. (2.13) since pions and kaons are both spin-0 particles.

The COMPASS data used for the extraction of the kaon multiplicities come from SIDIS measurements with a longitudinally polarized muon beam scattered off a longitudinally polarized ${}^6\text{LiD}$ target. The experimental setup covered the same kinematic range as used for the measurement of the strange quark polarisation (see Ref. [12]). The analysis is based on DIS events selected with inclusive triggers tagging the scattered muons only.

Kaon identification was performed using the RICH-1 detector. The particle identification is based on the extended likelihood method (see Ref. [31]). Knowing the momentum of the detected particle provided by the tracking system, likelihood values can be calculated for different mass hypotheses (π , K , p) dependent on the distribution of the measured Cherenkov photons. It is assumed that the detected particle

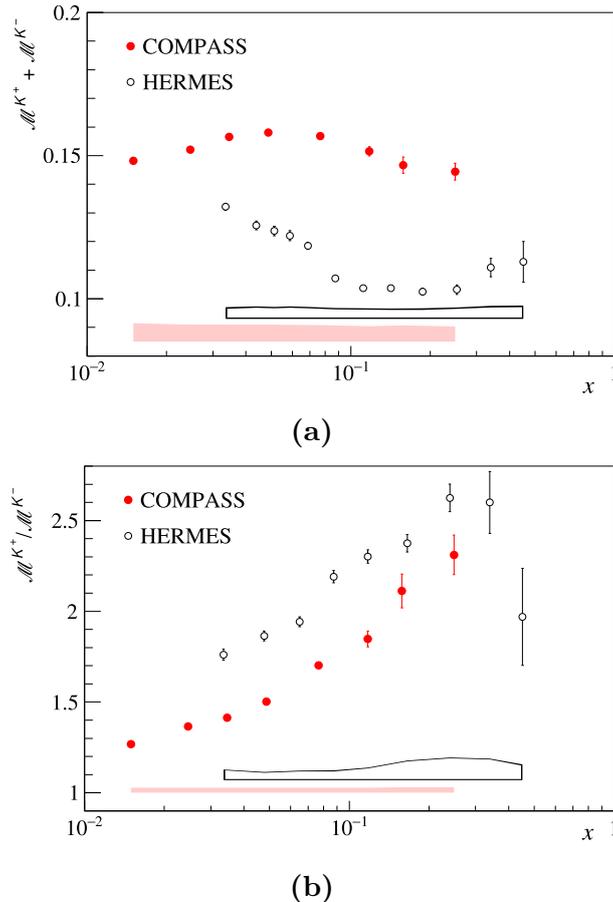


Figure 2.7: Sum $\mathcal{M}^{K^+} + \mathcal{M}^{K^-}$ (a) and ratio $\mathcal{M}^{K^+}/\mathcal{M}^{K^-}$ (b) of kaon multiplicities from the deuteron target as function of x_{Bj} extracted from COMPASS data in comparison with HERMES results [28].

corresponds to the hypothesis with the largest likelihood value. To ensure kaon identification at a confidence level of 95%, only particles with momenta between 12–40 GeV/c were accepted in the present analysis. Concurrently, the probability for falsely identified kaons was below 3% (see Ref. [28]).

The kaon multiplicities were obtained from the kaon yields summed over the different target polarizations, normalized to the number of DIS events. Corrections were applied considering QED radiative effects, the efficiency of hadron identification, the acceptance of the spectrometer and diffractive contributions from vector-meson production (see Ref. [28]). For an isoscalar target, the z -integrated sum of K^+ and K^- multiplicities can be written as [28]:

$$\mathcal{M}^{K^+} + \mathcal{M}^{K^-} = \frac{U\mathcal{D}_U^K + S\mathcal{D}_S^K}{5U + 2S}, \quad (2.19)$$

where $U = u + \bar{u} + d + \bar{d}$ and $S = s + \bar{s}$ are combinations of unpolarized PDFs and $\mathcal{D}^K(Q^2) = \int D^K(z, Q^2) dz$ denotes combinations of fragmentation functions integrated over the measured range of z , with $\mathcal{D}_U^K = 4\mathcal{D}_u^{K^+} + 4\mathcal{D}_{\bar{u}}^{K^+} + \mathcal{D}_d^{K^+} + \mathcal{D}_{\bar{d}}^{K^+}$ and $\mathcal{D}_S^K = 2\mathcal{D}_s^{K^+} + 2\mathcal{D}_{\bar{s}}^{K^+}$. Figure 2.7a shows the COMPASS result for $\mathcal{M}^{K^+} + \mathcal{M}^{K^-}$ as function of x_{Bj} in comparison with the HERMES results that were used at LO-QCD

to extract the product $S\mathcal{D}_S^K$ of the strange quark PDF and the strange quark-to-kaon fragmentation function (see Ref. [32]). The COMPASS data points show a rather flat and apparently larger distribution over the measured range of x_{Bj} than the HERMES results. It is assumed that the strange quark PDF $S(x_{Bj})$ approaches zero for high values of x_{Bj} . Hence, the sum $\mathcal{M}^{K^+} + \mathcal{M}^{K^-}$ can be equated with $\mathcal{D}_U^K/5$ in this region. At $x_{Bj} = 0.25$, $\mathcal{D}_U^K \approx 0.65 - 0.70$ was estimated from the COMPASS result, which is at variance with the value $\mathcal{D}_U^K = 0.43 \pm 0.04$ obtained from a global analysis of the fragmentation functions (see Ref. [33]). Most systematic effects are considered to cancel out in the ratio $\mathcal{M}^{K^+}/\mathcal{M}^{K^-}$ of the kaon multiplicities. Figure 2.7b shows the multiplicity ratio in dependency of x_{Bj} together with the HERMES results. Here, the COMPASS data points are systematically offset to smaller values from the HERMES results.

In conclusion, the COMPASS deuteron data together with the proton data of 2016 and 2017 on kaon multiplicities will provide a significant contribution to further constrain quark-to-kaon fragmentation functions. In particular, it can be expected that a better knowledge of the strange quark-to-kaon fragmentation function will shed more light on the contradictory results for the first moment of the strange quark helicity distribution as obtained from SIDIS measurements and the spin structure function g_1 , respectively.

3. The COMPASS-II Experiment

The COMPASS experiment is a fixed target experiment located at the M2 beam line of the Super Proton Synchrotron at the CERN accelerator complex. Since its founding, the COMPASS collaboration made important contributions in the fields of hadron structure and spectroscopy. This chapter gives an overview of the current experimental setup with a focus on physics with muon beams. Further details can be found in Ref. [13, 34].

3.1 The Beam Line

The COMPASS-II experiment can switch on-demand between high-intensity muon or hadron beams that originate from a primary proton beam extracted from the Super Proton Synchrotron at momenta up to $400 \text{ GeV}/c$. Protons are delivered to COMPASS at regular intervals during so-called *spills* with variable duration, dependent on the needs of other facilities. When injected to the M2 beam line, the protons are scattered off a Beryllium target. To control the flux of the pions and kaons created in the collision, production targets of different length are available.

During their transit through a 600 m long tunnel, the pions and kaons partially decay into muons and neutrinos. The remaining hadrons are finally stopped in a hadron absorber. Due to parity violation, the muon beam is naturally longitudinally polarized with respect to the beam axis. Since the degree of polarization depends on the ratio between the momenta of the muons and mesons, bending magnets for momentum separation are installed. The nominal momentum of the μ^+ beam selected for the COMPASS-II experiment is $160 \text{ GeV}/c$. At an intensity of $4.6 \times 10^8 \mu^+$ per spill, a polarization of $(-80 \pm 4) \%$ is achieved. Likewise, a longitudinally polarized μ^- beam with significantly lower intensity is available.

Due to the trade-off between beam flux and polarization, a momentum spread of around 5% is tolerated. Hence, accurate determination of the kinematics for each incident beam particle is required. This is done using the beam momentum station shown schematically in Fig. 3.1. The momentum measurement makes use of the fact that the beam is brought up to the surface level and subsequently bend back to the

horizontal axis before entering the experimental hall. The beam momentum station is composed of four scintillator hodoscopes (BM01–BM04) and two scintillating fibre detectors (BM05, BM06) positioned before and after a bending magnet (B6). From the curvature of the reconstructed particle tracks in the magnetic field, the momenta are derived at a precision of $\leq 1\%$ with an efficiency of $\approx 93\%$ [34, p. 15].

In the final section of the beam line, the incoming particles are focused on the target using an arrangement of additional bending and quadrupole magnets. However, the central beam spot is surrounded by a halo of particles, most of which are muons, that were not properly deflected.

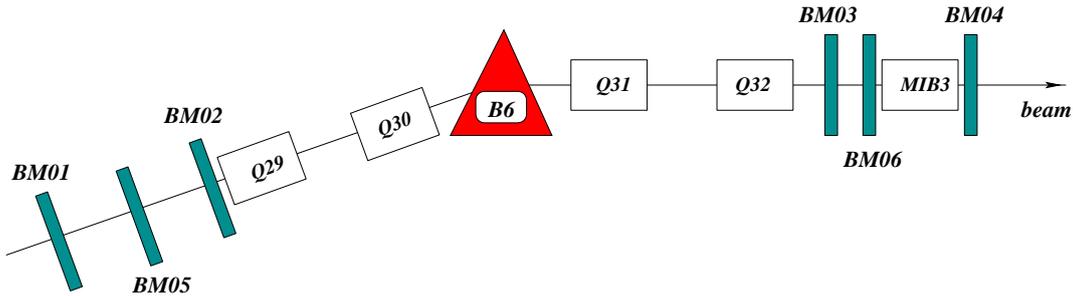


Figure 3.1: Schematic diagram of the beam momentum station [34].

3.2 The Target Region

Measurements at COMPASS using muon or hadron beams have mainly been performed with polarized solid-state targets. The muon programme measures cross-section asymmetries, the cross-section difference for different spin configurations of polarized muons and target nucleons divided by the corresponding spin-averaged cross-section. In order to extract the cross-section asymmetries from the measured raw asymmetries, the polarizations of the beam and target and the dilution factor, which accounts for the fraction of events from target nuclei other than polarized protons or deuterons, must be known. The target volumes are divided into two or three cells with opposite polarizations. Though the different target cells are exposed to the same beam flux, the polarization is inverted at regular intervals to cancel out acceptance effects. Deuterated lithium (${}^6\text{LiD}$) was selected as deuteron target and irradiated ammonia (NH_3) as proton target, respectively. The degree of polarization achieved by dynamic nuclear polarisation is $> 40\%$ for the deuteron target, while the proton target can be polarized to a degree of $> 80\%$ [34, p. 17]. However, the fraction of polarisable target nucleons in the ammonia target is only about one half of the value of deuterated lithium. Once the targeted polarization is achieved, the target material is kept in frozen spin mode at a temperature of around 60 mK.

In the frame of the COMPASS-II programme, a liquid hydrogen target was installed in 2012 to study quark generalized parton distributions using a $160\text{ GeV}/c$ muon beam. Simultaneously, the liquid hydrogen target is used to record semi-inclusive proton data that will allow for measurements of unpolarized quark parton distribution functions and to constrain quark fragmentation functions [13, p. 37]. The target installation is composed of a 2.5 m long cylindrical target cell with a

diameter of 40 mm that resides in a vacuum chamber with an outer diameter of 80 mm. To cool down the target material, a helium cryocooler with a cooling power of 30 W at 20 K is used. A detailed description of the target system is given in Ref. [35].

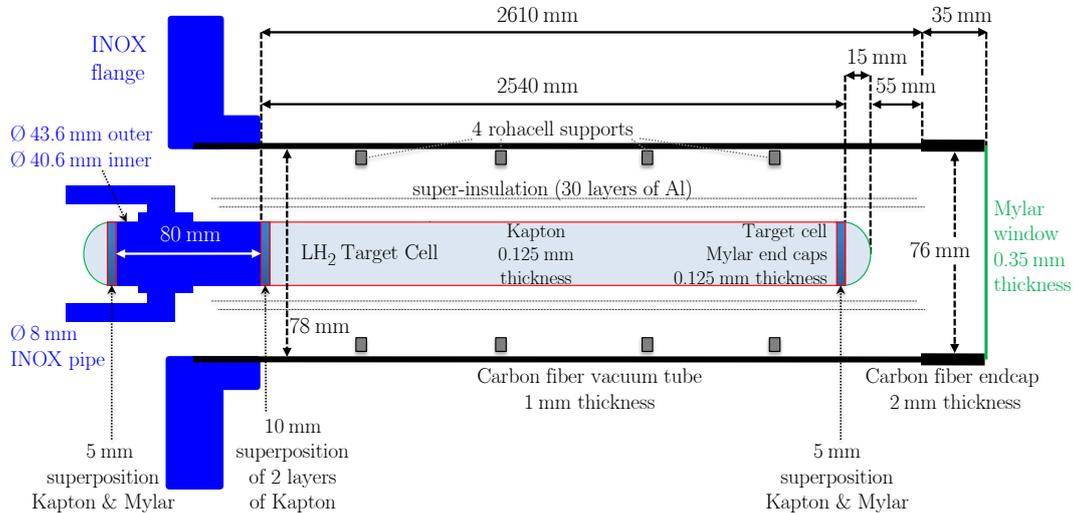


Figure 3.2: Sketch of the liquid hydrogen target system [36].

Certain generalized parton distributions can be constrained by measuring deeply virtual compton scattering cross-sections [37]. As pointed out in Sec. 3.1, the COMPASS apparatus has the unique possibility to switch simultaneously between μ^+ and μ^- beams, polarized along opposite directions. This feature allows to measure the deeply virtual compton scattering reactions $\overleftarrow{\mu}^+ p \rightarrow \mu^+ p \gamma$ and $\overrightarrow{\mu}^- p \rightarrow \mu^- p \gamma$ for lepton production of real photons from an unpolarized proton target. The length of 2.5 m of the target cell is based on the specification that a luminosity of about $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ shall be achieved for the μ^- beam. The exclusivity of the measurement requires that photons and recoiled protons with momenta down to 260 MeV/c can escape the target volume. In order to reduce the material budget of the target installation, the target cell is made of 125 μm thick Kapton sheet and Mylar end caps, whereas the vacuum chamber consists of carbon fiber with a total thickness of 1 mm and a 0.35 mm thick Mylar window (see Fig. 3.2).

Along with the liquid hydrogen target, the so-called CAMERA¹ detector was installed in 2012 to measure the momenta of the recoiled protons in exclusive scattering processes at full azimuthal and large polar angular acceptance. The CAMERA detector is composed of scintillator slats forming two concentric rings (inner and outer) surrounding the liquid hydrogen target (see Fig. 3.3). The inner and outer ring consists of 24 scintillator slats each which are equipped on both ends with light guides and a photomultiplier readout. The analog detector signals are digitized by analog-to-digital converter modules provided by the GANDALF² framework that will be described in Sec. 3.6. The measurement principle relies on the time and distance-of-flight between the intersection points of the particle tracks with the inner and outer ring. This constitutes a constraint for the thickness of the inner ring

¹COMPASS Apparatus for Measurements of Exclusive ReActions

²Generic Advanced Numerical Device for Analog and Logic Functions

since low-momentum protons must create sufficiently large signals in the outer ring and must not be stopped in the inner ring in order to be detected. The velocity and hence the momenta of the protons between the intersection points can be derived from the timing of the photomultiplier signals received on either side of the corresponding scintillator slats. It must be kept in mind that due to energy loss in the target and scintillators the measurement result can not be equated with the momenta of the recoiled protons at the interaction vertices. Therefore, the CAMERA detector has been calibrated using COMPASS data for exclusive ρ^0 muoproduction [38].

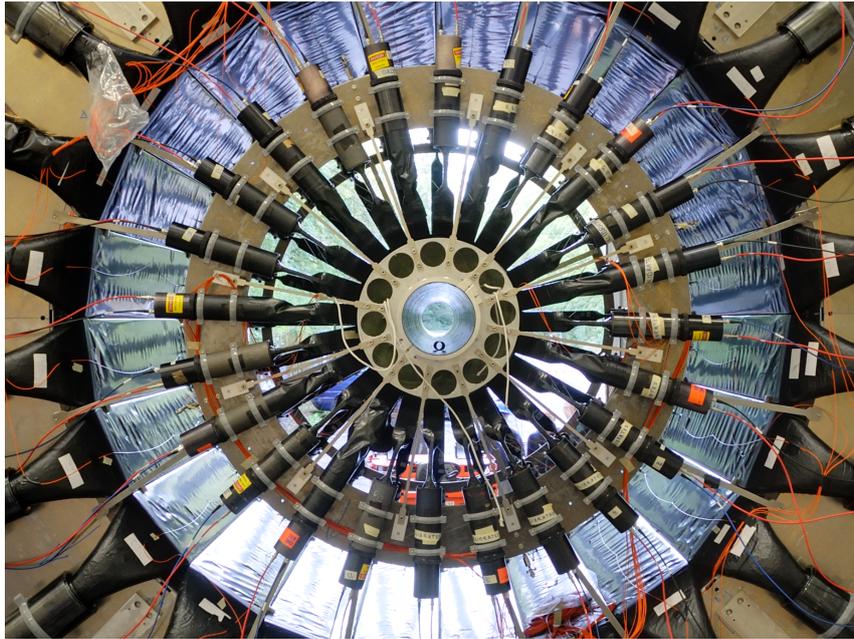


Figure 3.3: Photograph of the CAMERA detector looking downstream the beam axis. The photomultipliers and light guides of the inner ring and the scintillators and light guides of the outer ring are visible [38].

3.3 The Spectrometer

The large luminosity required for the COMPASS-II programmes demands for high-rate capability, efficient particle identification, large angular and momentum acceptance and the possibility to detect at the same time particles with very small scattering angles. Figure 3.4 shows a schematic view of the current experimental setup.

The COMPASS-II spectrometer extends over a length of about 60 m and is divided into two stages that are built around dipole magnets (SM1 and SM2). With precise knowledge of the magnetic fields, the momenta of the outgoing particles can be derived from the curvature of the reconstructed tracks. Low-momentum particles are detected in the Large Angle Spectrometer (LAS) that was designed for a polar acceptance of 180 mrad. Located downstream of SM1, a Ring Imaging Cherenkov (RICH) detector for charged hadron identification was installed. Particles with larger momenta scattered at small angles up to ± 30 mrad are covered by the Small Angle Spectrometer (SAS). Each spectrometer stage comprises a selection of tracking telescopes, electromagnetic and hadron calorimeters and a muon filter station.

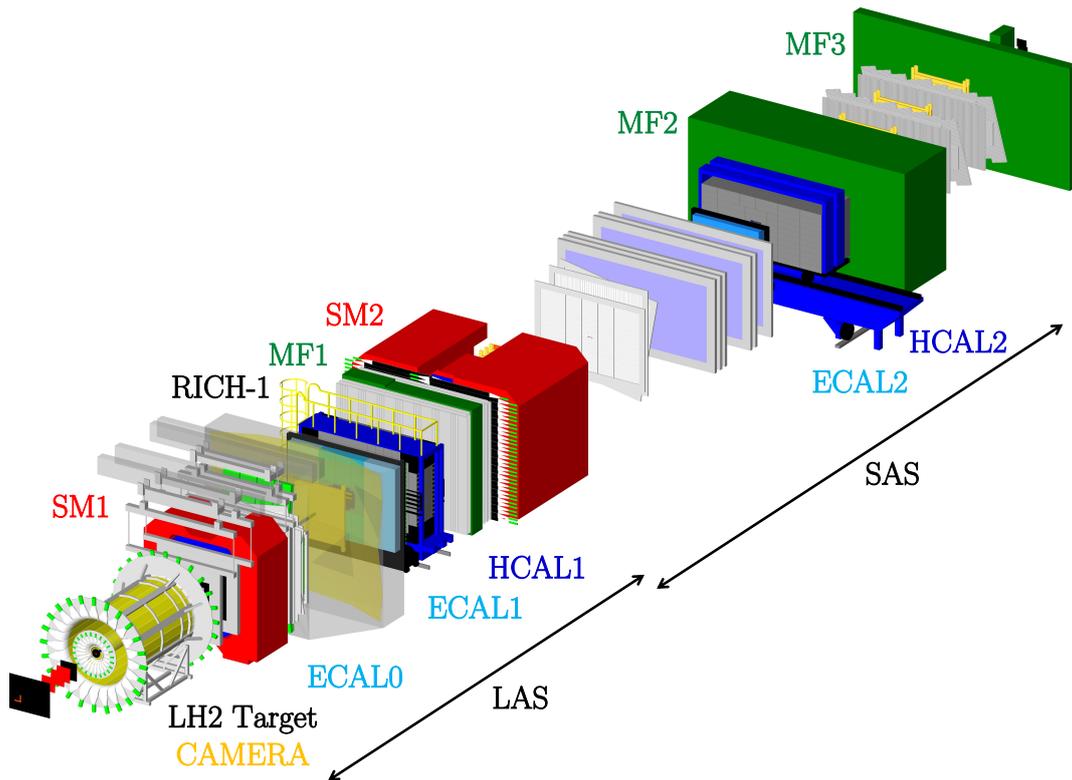


Figure 3.4: Schematic view of the COMPASS-II spectrometer [36].

3.3.1 Tracking

Depending on the distance to the target and the radial position with respect to the beam axis, different tracking detectors are installed for event reconstruction that can be categorized as following [34, p. 10].

- **Very Small Area Trackers:** The beam region up to a radial distance of 2.5–3 cm is the domain of the very small area trackers. The high rates of up to $10^5 \text{ s}^{-1} \text{ cm}^{-2}$ observed in the central beam spot demand for excellent time and spatial resolution and require robust detector designs that can withstand a high-intensity particle flow. These requirements are fulfilled by scintillating fiber and silicon microstrip detectors.
- **Small Area Trackers:** In the intermediate distance of 2.5–40 cm to the beam axis the requirement for high-rate stability is relaxed. This allows to employ micropattern gaseous detectors providing sufficient spatial resolution and, with regard to the covered area, a low material budget. In between the target and SM1, Micromegas¹ stations are in use, while the region downstream of SM1 is covered by GEM² stations. Extending to smaller scattering angles, the central part of both GEMs and Micromegas has been partially equipped with a pixelised readout that can cope with higher beam intensities. These so-called Pixel-GEM and Pixel-Micromegas can be attributed to the very small area trackers.

¹Micromesh gaseous structure

²Gas Electron Multiplier

- **Large Area Trackers:** Detector assemblies suitable to instrument large areas are installed to cover the outer regions defined by the acceptance specification of the spectrometer. These are drift chamber, straw drift tube and multi-wire proportional chamber stations.

3.3.2 Particle Identification

Efficient particle identification is a key requirement to study semi-inclusive and exclusive reactions. While hadron calorimeters measure the energy of hadrons produced at the interaction vertices, hadron identification is performed by the RICH-1 detector. Hadron calorimeters also serve as hadron absorbers for following muon detectors and constitute a decisive input for the trigger system, for instance triggering on scattered muons at low Q^2 . The energy of photons and electrons is measured using electromagnetic calorimeters. As opposed to momentum measurements, the performance of calorimeters improves with increasing energy E of the incoming particle according to the relationship $\sigma(E)/E \sim 1/\sqrt{E}$ [39, p. 602]. Therefore, calorimetry is essential at high energies. The critical energy for muons, defined as the energy where the energy loss due to ionization and radiation is equal, is several orders of magnitude higher than for electrons. Consequently, muons are able to penetrate thick absorbers like calorimeters. It is exactly this property the so-called muon filters take advantage of in order to distinguish muons from other particles.

- **RICH-1:** The angle under which Cherenkov light is emitted by particles traversing the large gas volume of the RICH-1 detector can be related to the particle velocity. In combination with the momentum of the incoming particle measured by the tracking system, the RICH-1 detector can separate protons, pions and kaons with momenta up to 50 GeV/ c . The RICH-1 principle and its associated photon detectors will be described in detail in Sec. 3.4.
- **Calorimeters:** Originally, two hadron calorimeters (HCAL1 and HCAL2) and two electromagnetic calorimeters (ECAL1 and ECAL2) were installed, one in each spectrometer stage. In the frame of the COMPASS-II programme, a third electromagnetic calorimeter (ECAL0) has been mounted after the target to cover larger photon angles. Lead glass modules are used to instrument the outer regions of ECAL1 and ECAL2. Incoming high-energy photons or electrons deposit their energy via pair production and Bremsstrahlung, which triggers electromagnetic showers in the absorber material. The secondary electrons and positrons in the shower emit Cherenkov light crossing the lead glass. Finally, the energy deposited in the calorimeter can be derived from the intensity of the Cherenkov light detected at the end of each module using photomultipliers. The ECAL0 and the central part of ECAL1 and ECAL2 are equipped with so-called Shashlyk modules. These sampling modules are made of alternating lead and plastic scintillator layers. The readout is performed using multi-pixel avalanche photo diodes that are less sensitive to magnetic fields.

The hadron calorimeters consist of sampling modules only. Each module is composed of alternating layers of iron and scintillators combined with a photomultiplier readout. Here, hadronic showers develop from inelastic interactions

of the incident hadron with the absorber material. The dimension of hadronic showers is described by the nuclear absorption length. This quantity is significantly larger than the radiation length, the counterpart to electromagnetic calorimeters. Hence, hadron absorbers are much thicker than the electromagnetic ones. As a consequence, hadron calorimeters are always located behind electromagnetic calorimeters in the spectrometer. Due to large fluctuations of the hadronic showers from one event to another, the relative energy resolution of electromagnetic calorimeters is about ten times better than for hadron calorimeters.

- **Muon Filters:** The COMPASS-II spectrometer is equipped with three Muon Filter stations (MF1, MF2, MF3) that consist of hadron absorbers preceded and followed by tracking telescopes. Since all particles other than muons are stopped in the absorbers, particle tracks continuing after the absorbers can be assigned to muons.

3.3.3 Trigger

Due to the luminosity at the COMPASS-II experiment, continuous data recording is not feasible. Hence, the purpose of the trigger system is to select events related to reactions that are in the focus of the interest. The limited buffer depth of the front-end electronics demand for a short decision time below 500 ns.

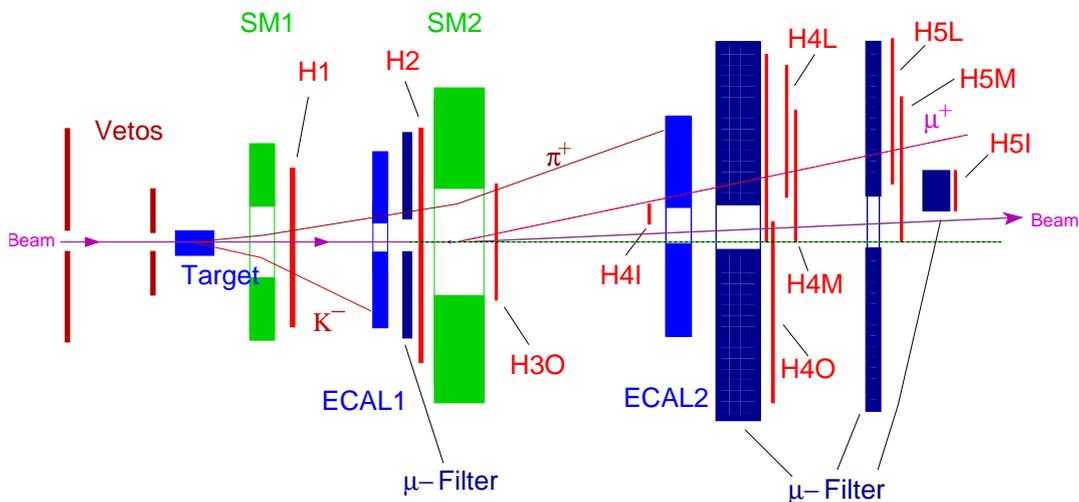


Figure 3.5: Placement of the hodoscope subsystems for the muon trigger [13].

The trigger system for the muon beam [40] is based on fast scintillator hodoscopes (see Fig. 3.5). Optionally, the hodoscope trigger can be combined with the fast response of the hadron calorimeters to select events with final state hadrons. In addition, a veto system is installed in front of the target to reject events triggered by halo muons. The trigger hodoscopes form subsystems, consisting of two stations each, that cover different kinematic regions of Q^2 .

For medium and large values of Q^2 , it is sufficient to measure the vertical position of the muon track at different distances from the target using two horizontal hodoscope

planes. This information is used to determine the projection of the muon trajectory in the non-bending plane of the spectrometer to be compared with the target position. At low Q^2 , vertical target pointing fails due to the very small scattering angles. In this case, the trigger accounts for the fractional energy loss y of the muon. Therefore, two vertical hodoscope planes measure the horizontal deflection of the muon trajectory at different distances downstream of the spectrometer magnets. In order to suppress background processes such as elastic scattering off target electrons, the energy loss trigger can be combined with a calorimeter trigger to request a certain energy deposit in the hadron calorimeters. The different hodoscope triggers are:

- **Inner and Ladder Trigger:** The inner (H4I, H5I) and ladder (H4L, H5L) hodoscopes consist of vertical scintillator slats. The inner trigger is sensitive to very small scattering angles, while the ladder trigger selects muons with large fractional energy loss.
- **Middle Trigger:** The middle (H4M, H5M) hodoscopes combine both horizontal and vertical planes to select deep-inelastic scattering events and to apply a coarse energy cut.
- **Outer Trigger:** The outer (H3O, H4O) hodoscopes use horizontal scintillator slats for vertical target pointing up to $Q^2 = 10 \text{ (GeV}/c)^2$ [13, p. 85].
- **LAS Trigger:** The LAS (H1, H2) hodoscopes are similar in structure to the outer hodoscopes. They cover the highest region of Q^2 up to $20 \text{ (GeV}/c)^2$ [40].

The hodoscope signals of each trigger subsystem enter coincidence matrices to select hit combinations that can either be related to extrapolated muon tracks matching the target position or, in case of vertical planes, to a certain range of y . Subsequent coincidence units also account for the calorimeter trigger and the veto signal [40].

3.4 The RICH-1 Detector

Particle identification using the RICH-1 detector (see Fig. 3.6a) relies on the detection of Cherenkov light emitted by particles crossing a 3 m long C_4F_{10} gas-filled vessel. The Cherenkov light cones are focused by a spherical UV mirror surface on two arrays of photon detectors mounted outside the spectrometer acceptance. A steel pipe with a diameter of 10 cm retains Cherenkov photons produced by beam particles from entering the gas volume [31].

The fundamental relationship between the photon emission angle θ_c and the particle velocity v is given by [39, p. 439]:

$$\cos \theta_c = \frac{1}{\beta n}, \quad (3.1)$$

with $\beta = v/c$ and the refractive index n . Since $\cos \theta \leq 1$, the threshold velocity for Cherenkov emission is given by:

$$\beta_{th} = \frac{1}{n}. \quad (3.2)$$

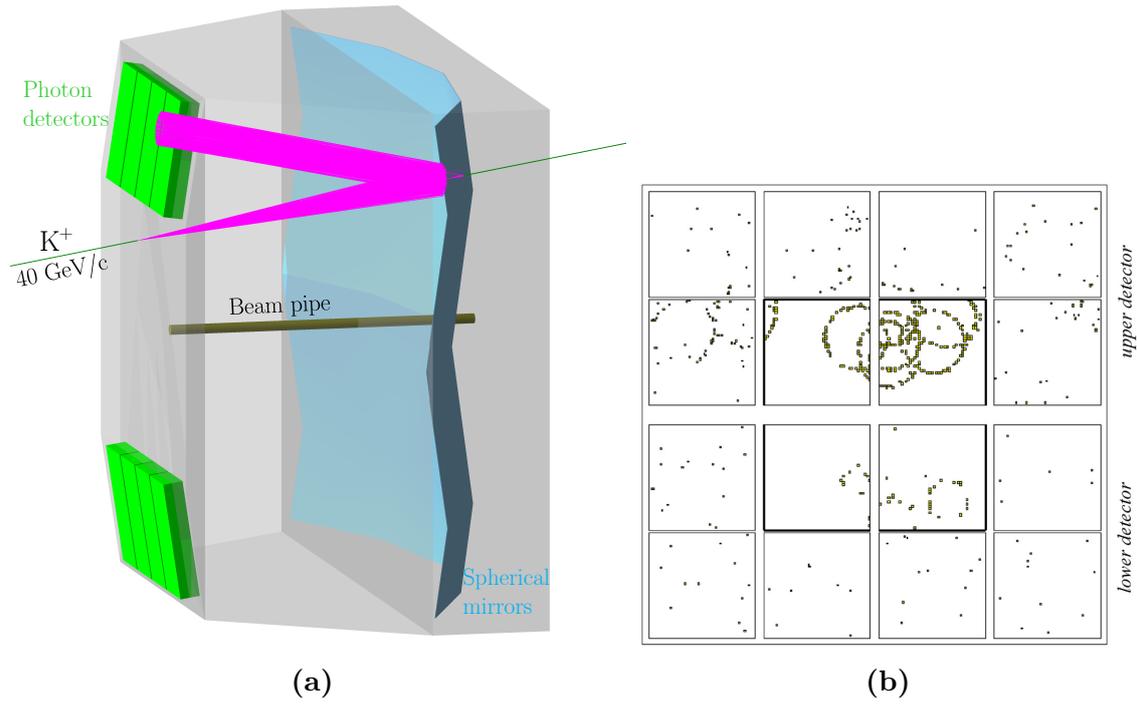


Figure 3.6: (a) Visualization of the RICH-1 detector showing the response to an incident K^+ with a momentum of $40 \text{ GeV}/c$ [36]. (b) Typical event display of the photon detectors located in the focal plane of the mirror system. Several photon rings are visible [31].

In other words, the particle velocity must be larger than the phase velocity in the radiator medium ($v > c/n$). Knowing the particle momentum p from the deflection in the spectrometer magnets, the mass m can be derived from the measurement of the Cherenkov angle θ_c [39, p. 452]:

$$m = \frac{p}{c} \sqrt{n^2 \cos^2 \theta_c - 1}. \quad (3.3)$$

At the RICH-1, the Cherenkov thresholds for pions, kaons and protons are found at 2.5, 9 and $17 \text{ GeV}/c$, respectively. It must be kept in mind that the refractive index $n = n(w)$ is actually a function of frequency. The radiated energy per unit frequency interval per unit path length is given by [39, p.442]:

$$\frac{d^2 E}{dw dx} = \frac{z^2 e^2}{4\pi\epsilon_0 c^2} w \left(1 - \frac{1}{\beta^2 n^2(w)}\right). \quad (3.4)$$

From Eq. (3.4), the number of Cherenkov photons emitted per unit wavelength interval over the radiator length L can be calculated [39, p.443]:

$$\frac{dN}{d\lambda} = \frac{2\pi z^2 \alpha}{\lambda^2} L \sin^2 \theta_c, \quad (3.5)$$

where α is the fine structure constant. The actual photon yields are less than Eq. (3.5) predicts since, in order to be detected, the emitted photons are converted into electrons. In view of the overall efficiency for photon collection and detection, the formula for the number of photoelectrons reads [39, p.445]:

$$N_{pe} = N_0 z^2 L \sin^2 \theta_c. \quad (3.6)$$

The quality factor N_0 is considered as the figure of merit of the given Cherenkov detector. Since the azimuthal angle of emission is uniformly distributed, the detected Cherenkov photons form a ring image in the focal plane of the RICH-1 mirror system (see Fig. 3.6b). In this light, the Cherenkov angle θ_c can be reconstructed from the radius of the photon ring and the known particle trajectory. Using the maximum likelihood method, the particle mass is assigned to the mass hypothesis that best fits the observed photon ring [31].

The central part of the RICH-1 photon detector, receiving the highest rates due to uncorrelated background events, is based on multi-anode photomultiplier tubes. The outer region, covering an active area of around 4 m^2 , was instrumented with multi-wire proportional chambers coupled to solid-state caesium iodide photocathodes [13, p. 112]. However, the multi-wire proportional chamber design suffers from performance limitations related to ion backflow from the multiplication process, which leads to a certain decrease of the quantum efficiency after a collected charge of a few mC/cm^2 [41]. Long recovery times ($\sim 1 \text{ d}$) after detector discharges prevent the operation at high gain ($> 10^4$) [42], thus limiting the efficiency for single photon detection. Therefore, an active area of 1.4 m^2 was upgraded in 2016 with new hybrid photon detectors based on micro-pattern gas detector technologies.

3.4.1 Hybrid Photon Detector

In selection of a large area photon detector providing higher robustness against electrical discharges, larger operational gain and reduced ion feedback, the choice was made for a hybrid design combining Thick Gas Electron Multiplier (THGEM) structures as photo-sensitive pre-amplification stage with a Micromegas.

The THGEM consists of a Printed Circuit Board (PCB) with a mechanically drilled hole structure similar to the GEM design but with much larger scale. Typical geometrical parameters are [43]: thickness of $0.2 - 1 \text{ mm}$ and hole diameter of $0.2 - 1 \text{ mm}$ with a spacing about twice the hole diameter. The rim, a small clearance ring around the holes, has a width $< 0.2 \text{ mm}$. Due to its robust design, THGEMs can be industrially manufactured at affordable costs. The operation principle is the same as for the GEM detector. A potential difference of a few kV between the copper layers creates a strong electric field within the holes. The electric field also reaches out to the surrounding attracting electrons, induced by gas ionization, into the holes where their multiplication is obtained by a gas avalanche process. In most applications, multiple THGEM layers are cascaded to achieve higher gain.

Toward the development of large area THGEM-based detectors, the geometrical parameters have been extensively studied and optimized [45]. Two effects that scale with the size of the active area have been identified. First, local thickness variations have a strong effect on the gain uniformity. By selecting the most uniform areas from the raw sheets, a thickness tolerance of 2% was achieved, corresponding to a gain deviation factor of ≤ 1.5 over a detector surface of $300 \times 300 \text{ mm}^2$. The second issue concerns electrical instabilities at maximum gain due to microscopic defects and copper residuals on the hole edges. Therefore, the production process includes a polishing treatment to improve the surface quality, which leads to a significant improvement of the operation stability.

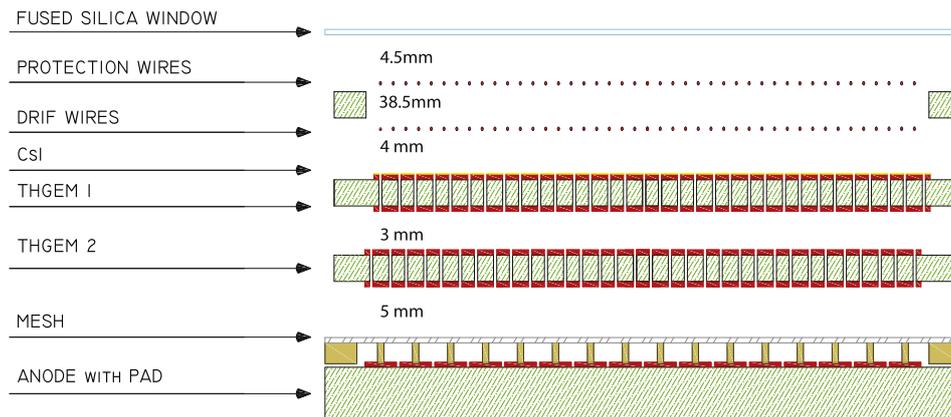


Figure 3.7: Schematic view of the hybrid photon detector [44].

The final hybrid hybrid photon detector has an active area of $600 \times 600 \text{ mm}^2$ and consists of a dual-layer THGEM with staggered hole alignment and a bulk Micromegas (see Fig. 3.7). The upper THGEM layer is coated with caesium iodide to act as reflective photocathode. The intrinsic parameters of the THGEMs are [44]: thickness of 0.47 mm, hole diameter of 0.4 mm with a spacing of 0.8 mm. The holes have no clearance rings. The micromesh and the anode plane, segmented into pads of $7.5 \times 7.5 \text{ mm}^2$, form a small ($128 \mu\text{m}$) amplification gap. The electrons from the avalanche process are collected on the anode pads, while the positive ions drift toward the mesh. The short drift times and hence fast signals induced by the positive ions and their effective removal at the mesh are intrinsic benefits of the Micromegas design. Together with the staggered design of the hole structures of the THGEMs, the ion backflow rate to the photocathode is effectively suppressed [46]. The final $600 \times 600 \text{ mm}^2$ hybrid photon detector is operated with a gain in the range between 1.2×10^4 and 2.5×10^4 [44].

3.5 The Data Acquisition System

The data acquisition system [47] accomplishes the readout of about 300 000 detector channels. The typical data rate is in the order of 1.6 GB/s during a spill. Figure 3.8 shows a flow diagram of the data acquisition system.

The detector signals received from pre-amplifier and discriminator modules are continuously digitized by analog-to-digital and time-to-digital converter front-end modules, mounted as close as possible to the detector stations. Upon reception of the trigger signal, the front-ends pass on the data, which coincide in time with the trigger event, to subsequent VME readout modules. At this stage, the event data is combined with identifier labels received from the Trigger and Control System (TCS). The collected data is finally transferred via fiber-optic links to 15:1 multiplexer cards following the S-LINK protocol [48]. The first data concentrator level of the data acquisition system can receive up to 120 incoming fiber-optic links. At the next level, a single switch card is employed to merge the sub-events of all detector channels before being sent to readout computers for temporary storage. Finally, the data files

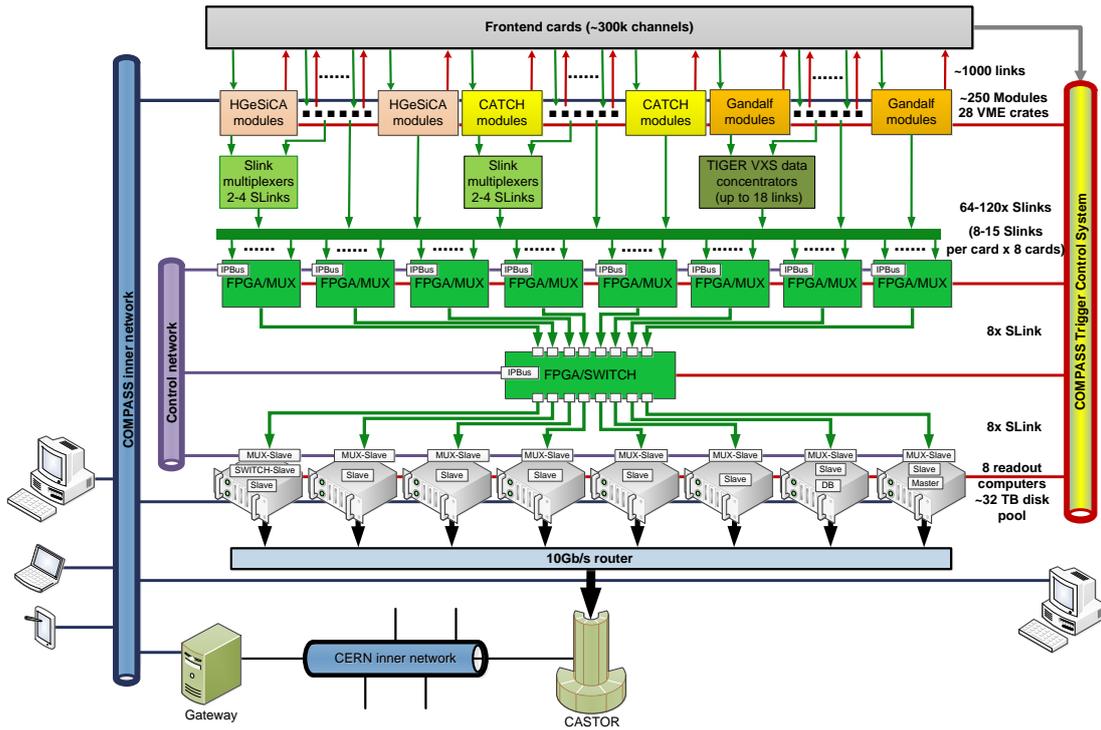


Figure 3.8: Structure of the data acquisition system [47].

are transferred via a 10 Gbit/s Ethernet network to the CERN storage management system CASTOR¹ that copies them to magnetic tapes.

3.6 The GANDALF Framework

Primarily developed for the readout of the CAMERA detector, the GANDALF framework is now employed for different applications at the COMPASS-II experiment. Detailed information about this project can be found in Ref. [49]. The GANDALF module was designed as 6U/VME64x [50] carrier board to employ the widely-used VMEbus interface [51]. Moreover, it complies with the VITA 41.0 VXS² standard [52] that defines a high-speed serial backplane link to connect each payload board on a VXS backplane to a central switch card. The GANDALF mainboard is combined with exchangeable, application specific mezzanine cards to adapt the system to the required readout task. Figure 3.9 shows a picture of the GANDALF module illustrating the main hardware components and interfaces.

3.6.1 Mainboard

Digital Signal Processing (DSP) is performed using a Field-Programmable Gate Array (FPGA) device³ that has been chosen owing to its large amount of fast logic units, allowing for a wide range of complex arithmetic operations up to 550 MHz [53]. Extending the integrated memory of the FPGA, the mainboard comprises both a

¹CERN Advanced STORAGE manager

²VME Switched Serial

³Xilinx Virtex-5 XC5VSX95T-2FFG1136C

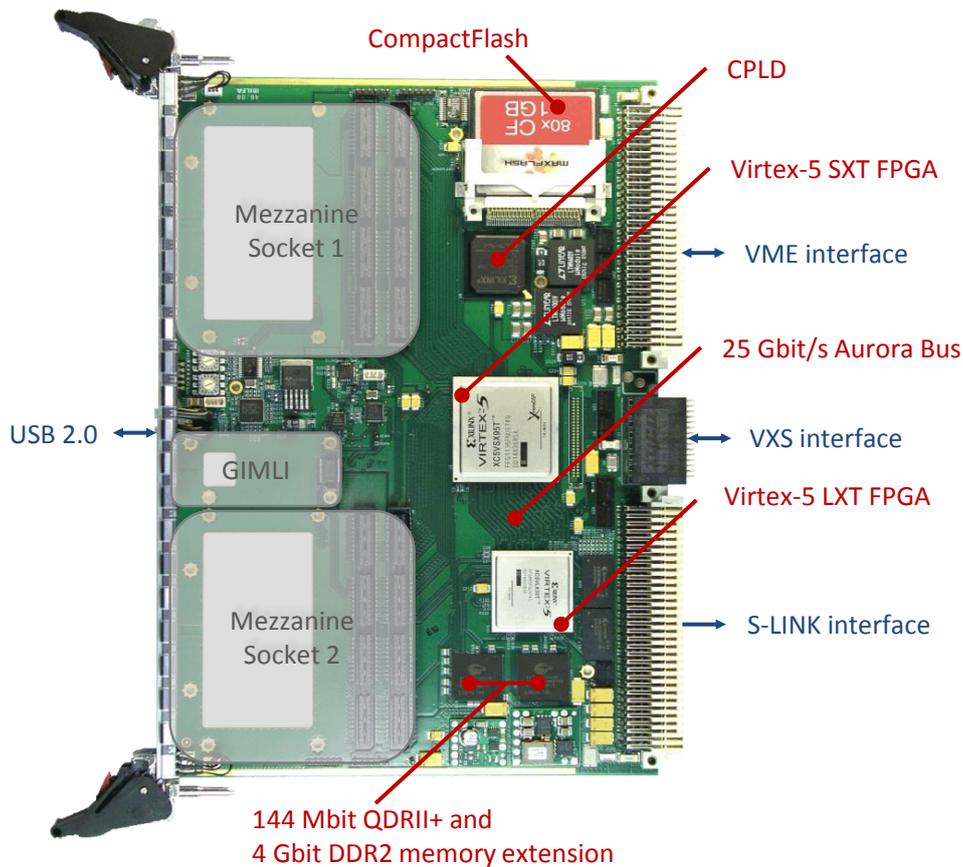


Figure 3.9: Overview of the GANDALF mainboard. External interfaces are depicted by blue labels. The main hardware devices allocating the GANDALF module as digital signal processor are labeled in red. The grey-colored boxes indicate the mezzanine sockets for the add-on cards.

144 Mbit QDRII+¹ and a 4 Gbit DDR2² memory. These devices are connected to a second FPGA³ that complements the system with a memory controller. In either direction, the FPGAs are connected by eight differential serial links facilitating the integrated high-speed transceiver tiles. The implemented Aurora bus protocol [54] is capable of exchanging data at a bandwidth of 25 Gbit/s.

The FPGA configuration is stored inside the device using volatile static Random-Access Memory (RAM) cells. Upon power-up, each FPGA device must therefore be programmed with a bitstream file to implement the intended circuit. Configuration bitstreams may be retrieved directly from an on-board Flash memory⁴. Though, the preferred way to configure the GANDALF module is through the VMEbus or the USB2.0 interface. These interfaces are accessed by a Complex Programmable Logic Device (CPLD)⁵ that transfers the received data on 8-bit wide SelectMAP buses to the FPGAs. The CPLDs are limited in terms of performance and logic resources,

¹Cypress Semiconductor CY7C1515V18

²Quimonda HYB18T2Gx02BF

³Xilinx Virtex-5 XC5VLX30T-2FFG665C

⁴System ACE CompactFlash

⁵Xilinx CoolRunner-II CPLD XC2C512-FG324

but are advantageous over other programmable logic devices as they are operational directly after power-up, thanks to non-volatile internal memory. Likewise, the VME64x backplane bus is used for various board monitoring and slow-control tasks.

Data readout is performed using either the high-speed VXS interface or dedicated backplane transition cards. The bus lanes of the transition card interface are directly connected to both FPGAs. This gives the user the flexibility to operate the protocol of choice, for instance, the 160 MB/s S-LINK interface [48]. Data acquisitions with limited bandwidth may also be carried out using the USB2.0 or the VMEbus interface.

3.6.2 Mezzanine Cards

The GANDALF mainboard allocates mezzanine sockets to extend the scope of readout tasks. Different kinds of add-on cards are available, consistent with particular input signal standards and operational modes.

GIMLI

At the COMPASS-II experiment, the TCS distributes a reference clock and the trigger signal together with corresponding event labels via an optical fiber network to every readout module employed in the spectrometer. A dedicated optical receiver and a clock recovery circuit are located on an add-on card mounted to the so-called GIMLI socket (cf. Fig. 3.9), in order not to restrain the range of application of the GANDALF module to a single clock distribution system. The recovered clock signal and the TCS data stream are routed separately from the GIMLI socket to the central FPGA where the final decoding of the event labels takes place. In addition, the clock signal is provided to the mezzanine sockets and a clock synthesizer chip¹ on the mainboard using differential buffers. Two more versions of the GIMLI card are presently available. One of these replaces the optical interface with LEMO input connectors for both clock and trigger signal. Optionally, the card hosts a 20 MHz oven-controlled crystal oscillator to substitute the external clock input. In a different scenario, the TCS is distributed via the VXS interface to the FPGA. Here, a dummy add-on card was designed to take the clock signal from the FPGA and redirect it to the targets mentioned before.

Analog Mezzanine Card

The analog mezzanine card is assembled with eight 12-bit analog-to-digital converter chips² running at approximately 500 MSample/s (see Ref. [55]). Using a time-interleaved method, the sampling rate is doubled (to 1 GSample/s) combining neighbouring channels. Both analog-to-digital converters are operated at the same frequency but receive clock signals with a constant 180° phase offset to one another. As a consequence, the number of input channels per mezzanine card is reduced by half. Real-time pulse shape analysis and feature extraction on the continuous data stream is carried out inside the FPGA on the mainboard.

¹Silicon Labs Si5326

²Texas Instruments ADS5463

The photomultiplier signals of the CAMERA detector are digitized by GANDALF modules equipped with analog mezzanine cards as transient analyzers. Pulse shape parameters like amplitude, integral and time stamp are extracted from the incoming data stream with a sophisticated digital constant fraction discrimination algorithm (see Ref. [56]). Its precision is inversely proportional to the signal amplitude in terms of the dynamic range. For pulses exceeding 25% of the dynamic range, the time resolution obtained at 1 GSample/s is better than 7 ps (see Ref. [56–58]).

Digital Mezzanine Card

The digital mezzanine card may be employed in systems that do not require the front-end electronics to be attached directly to the detector. Comprising two female VHDCI¹ connectors, the digital mezzanine card interconnects 64 differential channels with the FPGA on the mainboard. Differential buffers² accepting LVDS³ and LVPECL⁴ input signals, most commonly used by discriminator fabrics, are incorporated to protect the FPGA I/Os. With different placement of the buffer components during PCB assembly, the digital mezzanine card is turned into a LVDS output card. In this configuration, the GANDALF framework is utilized as versatile pattern generator that has extensively been used in the characterization of the ARAGORN front-end (see Sec. 7.2).

The scintillating fiber stations in front of the target region in the COMPASS-II spectrometer receive the highest input rates of several MHz per channel and exhibit very long trigger latency in the microsecond range. Reading these detectors, a 128-channel time-to-digital converter providing sufficiently deep hit buffer memory has been implemented using the GANDALF module in this configuration (see Ref. [59, 60]). The design could be extended by a scaler featuring online beam monitoring functionality (see Ref. [61, 62]). Further applications include a mean-timer with subsequent coincidence logic providing fast trigger decisions (see Ref. [63]).

Optical Mezzanine Card

The ARWEN⁵ add-on card complements the scope of applications as generic readout engine for front-end modules attached by optical links and shall hence be utilized for the ARAGORN front-end. The ARWEN card hosts a FPGA device⁶ to establish the electrical interface to four Small Form-factor Pluggable (SFP) transceiver sockets. The entire optical network has been verified at a bandwidth of 3.1104 Gbit/s in either direction using standard 850 nm optical SFP transceiver modules⁷. Toward the FPGA on the mainboard, a bidirectional interface is provided consisting of 32 differential channels per direction. Eighteen out of these signals are allocated to parallel data transfer whereas the remaining interconnects are assigned to miscellaneous functions. Furthermore, the mezzanine socket includes a SelectMAP interface to configure the FPGA with bitstream files received from external storage via the

¹Honda HDRA_ED136

²ON-Semi NB4N855S

³Low Voltage Differential Signaling

⁴Low Voltage Positive Emitter Coupled Logic

⁵Acronym for a mezzanine card providing four fiber optic transceiver modules

⁶Xilinx Spartan-6 XC6SLX45T-3CSG324C

⁷Finisar FTLF8524P2xNy and FTLX8571D3BCV

VMEbus interface. Figure 3.10 shows a simplified block diagram of the ARWEN mezzanine card.

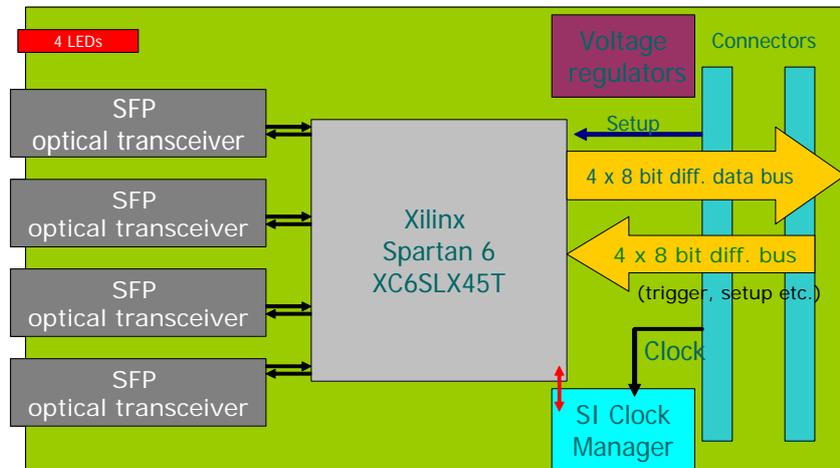


Figure 3.10: Functional block diagram of the ARWEN mezzanine card [64].

As mentioned earlier, a reference clock enters the GANDALF mainboard through the GIMLI add-on card. This clock signal is forwarded to a jitter attenuator chip¹ on the ARWEN card, providing the embedded transceiver tiles and the FPGA fabric with a clean reference. In the up-link direction, the design aims to distribute the trigger items along with various control signals with fixed latency to each payload board. This is achieved through bypassing the transmit buffers and disabling certain features like the integrated 8B/10B encoder inside the transceiver primitives. The same design constraints apply to the receiver side in order to maintain fixed-latency data transfer and recovery of a phase-synchronous reference clock upon the connected front-end boards. Hence, the up-link data stream regularly includes a recognizable bit sequence, referred to as comma symbol, to align the parallel data to the word boundary. The actual alignment process is device specific and needs to be tailored to the hardware realization of the payload board (see Sec. 6.2.1). More than that, the front-end modules can be remotely controlled via the VMEbus interface. For this purpose, a custom bus protocol has been introduced transferring configuration data to the slave boards and retrieving status information in return. Details about the configuration bus interface will be given in Sec. 6.3.

Down-link specifications are less restrictive as the constant-latency feature is not required for data readout. The FPGA on the ARWEN card not only merges the data received from the optical links but also acts as local event builder compliant with the data format of the COMPASS-II experiment (see Ref. [65]). This allows to perform certain consistency checks before passing on the received data to the subsequent data acquisition system. Following merger stages are also implemented inside the FPGA on the mainboard to process the output of both mezzanine cards. Detailed information about the employed FPGA designs can be found in Ref. [66].

¹Silicon Labs Si5326

4. The ARAGORN Front-end

The ARAGORN front-end constitutes a compact¹ hardware platform for FPGA based Time-to-Digital Converter (TDC) processing 384 input channels. Drawings of the PCB top and bottom layer are provided in Figs. A.1 and A.2 to locate the board components described in this chapter by means of their reference designators. The readout of the ARAGORN front-end is performed using high-speed optical links (see Fig. 4.1). This chapter focuses on the description of the hardware layout of the ARAGORN front-end.

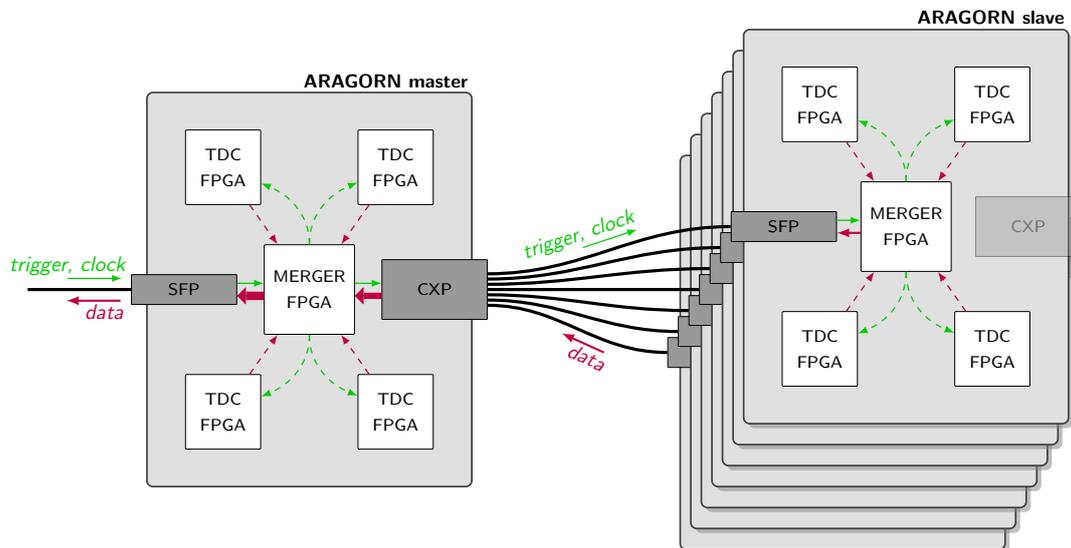


Figure 4.1: Schematic illustration of the optical readout network that combines eight ARAGORN front-ends through a star topology. Thereof seven slave boards are connected to the master switch card as satellites. A single fiber optic cable transfers the collected data from the master front-end to the subsequent readout engine.

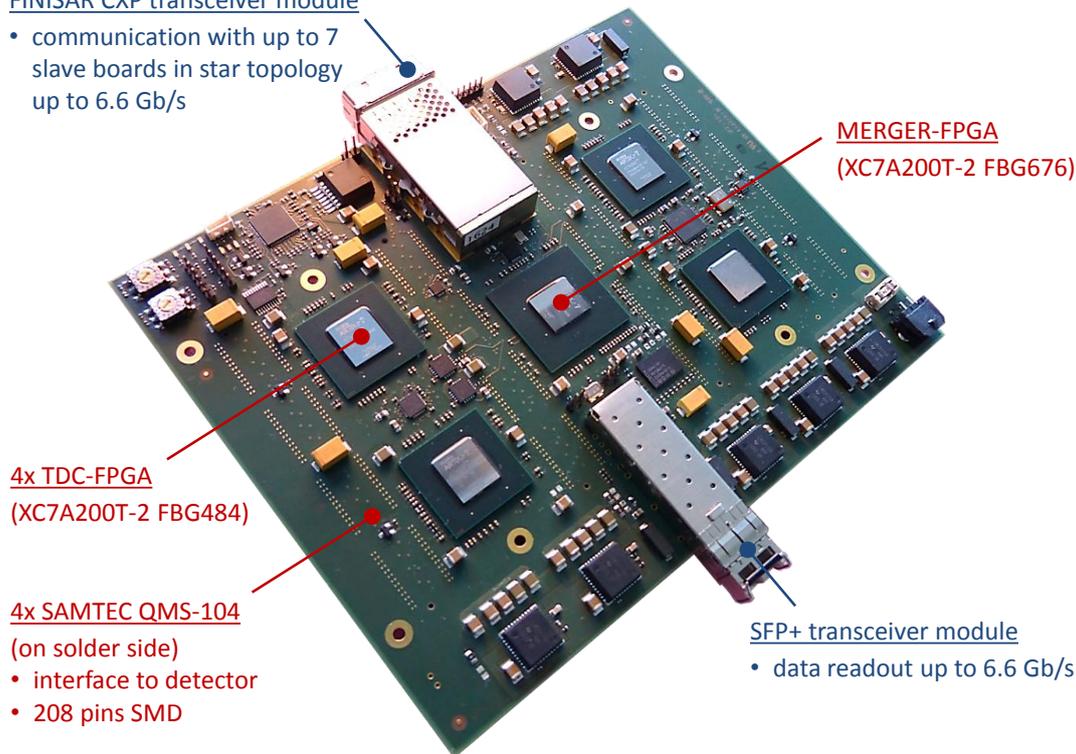
¹Board dimension: $140 \times 172 \text{ mm}^2$

4.1 Design Overview

Four FPGA devices¹ on the ARAGORN front-end, referred to as TDC-FPGAs, implement time-to-digital converter that accept 96 input channels each. An interface to the associated detector, or more specific analog readout modules, is provided by four high-density board-to-board connectors² located on the PCB bottom side. The TDC inputs make use of the LVDS standard. A fifth FPGA device³, briefly MERGER-FPGA in the following, provides differential point-to-point connections to each TDC-FPGA for data output. Consequently, the MERGER-FPGA acts as data concentrator and local event builder to present the collected data in an appropriate format to the data acquisition system.

FINISAR CXP transceiver module

- communication with up to 7 slave boards in star topology up to 6.6 Gb/s



MERGER-FPGA (XC7A200T-2 FBG676)

4x TDC-FPGA (XC7A200T-2 FBG484)

4x SAMTEC QMS-104

- (on solder side)
- interface to detector
- 208 pins SMD

SFP+ transceiver module

- data readout up to 6.6 Gb/s

Figure 4.2: This picture shows a photograph of the ARAGORN front-end. The main functional blocks are highlighted. The board comprises four FPGAs that implement time-to-digital converter in favour of 384 input channels. A fifth FPGA functions as generic data processor. A SFP transceiver socket and a CXP compliant receptacle are located on either side of the central FPGA. Four extension board connectors providing an external interface to the associated detector are located on the solder side of the PCB (not visible).

Furthermore, the MERGER-FPGA provides the infrastructure to control miscellaneous on-board devices that will be described within this chapter. The associated low-speed bus interfaces are accessed by a Microblaze embedded processor. This soft processor core is implemented with the programmable logic primitives and integrated memory cells provided by the MERGER-FPGA fabric. Including only

¹Xilinx Artix-7 XC7A200T-2 FBG484

²Samtec QMS-104-06.75-L-D-A-GP

³Xilinx Artix-7 XC7A200T-2 FBG676

the relevant features, the highly-configurable embedded design could be optimized to reduce the device utilization to a minimum. The Microblaze processor supports a rich set of peripherals which makes it a perfect match for slow-control tasks that can be efficiently processed by a high-level software application.

As already mentioned, the readout of the ARAGORN front-end is accomplished with fiber optic cables attached to a SFP+ transceiver module that in turn connects to integrated high-speed transceiver tiles of the MERGER-FPGA. The transceiver primitives provide a bandwidth of up to 6.6 Gbit/s. Though the transceiver links have been configured to adapt the system to the GANDALF module equipped with ARWEN add-on cards (cf. Sec. 3.6.2), the programmable logic is feasible of operating various industry standard protocols to fit the system in alternative readout configurations. Optical data communications is favored over standard copper cabling not only because it works over longer distances at higher bandwidths. Moreover, fiber optics overcome electromagnetic interference issues induced by high-speed board-to-board applications. Because the intrinsic signal-to-noise ratio of many particle detectors is small, this property is extremely valuable to keep the noise contribution from the readout electronics to a minimum. Toward the evolution of a high-speed optical readout network, the layout of the ARAGORN front-end permits to connect up subsequent *slave* boards through a star topology. This design objective is achieved with an optional CXP transceiver module, interconnecting with up to seven ARAGORN cards as satellites using an optical fanout cable. Consequently, eight boards thus 3072 detector channels can be read out via a single fiber optic cable attached to the SFP slot of the *master* board that hosts the CXP module (see Fig. 4.1). Thanks to the hot-pluggable realization of the costly CXP transceiver module, the PCB layout of the ARAGORN front-end is maintained independent of its final application.

The transceiver tiles inside the FPGA fabric integrate a clock recovery circuit to synthesize a clock signal from the incoming data stream that is phase-synchronous with the sender clock. This is of particular importance because the design intends to apply the recovered clock in the TDC application. It is thus essential that the transceiver channels linked to the CXP socket are driven with this particular clock alike, in order to forward the data stream to the connected slave boards. Unfortunately, the output of the clock and data recovery block does not comply with the jitter specifications of the transceiver tiles. To select a jitter attenuator circuit that reliably removes jitter from the recovered clock and at the same time maintains system synchronization upon repeated device initializations, extensive design studies were necessary before PCB production. These exercises were subject to a recent thesis conducted within the framework of this project (see Ref. [67]). The specifications of the selected device are described in detail in Sec. 4.5.2.

4.2 Field-Programmable Gate Array

Most digital components including application-specific integrated circuits are designed for particular functions that cannot be changed after manufacturing. In contrast, programmable logic devices perform user-defined operations on the targeted system. Programmable array logic components are considered as early precursors of today's programmable logic devices. These devices consist of a matrix

forming the logical AND, also called product term, of any combination of the optionally inverted inputs. Together with a set of fixed OR-gates, each output represents a sum-of-product term, implementing a variety of boolean functions. Further programmable array logic devices extended the combinatorial circuit with register components providing feedback to the programmable input array. In the delivery state, the AND-matrix is either completely connected or unconnected depending on the manufacturing technology. During the programming process, the unwanted interconnections are removed or as opposed to this, particular connections are made conductive. This procedure had to be applied beforehand PCB assembly. Later versions could be erased using ultraviolet light and electrically reprogrammed in the field (see Ref. [68, p. 270 ff.]). The implementation of more elaborate designs became feasible with the evolution of the CPLD family that extends the programmable array with embedded macrocells. These macrocells contain sophisticated logic gates like multiplexers and sequential components providing feedback and interconnections to other functional blocks via a central interconnection network. Compared to programmable array logic components, CPLDs provide significantly more logic resources along with an increased number of I/Os. In addition, the device configuration is stored in SRAM cells that are initialized from on-chip nonvolatile memory when the device is powered up (see Ref. [68, p. 274 ff.]).

The demand for high-performance programmable logic devices eventually led to the evolution of FPGAs. The embedded logic cells, arranged in a regular array, are linked to a programmable interconnection network. Compared to previous implementation fabrics, combinatorial circuitry is provided by so-called look-up tables, consisting of asynchronous RAM cells that store the truth table of the boolean function. Because the configuration resides in volatile memory, FPGA devices have to be programmed from on-board Flash memory or via external interfaces when power is applied. The I/O blocks are likewise programmable to support different signal standards to interact with other devices in a system. Modern FPGA devices include rich quantities of embedded RAM blocks to buffer the data to be processed or to store the executable image of embedded processor systems. Various portions of a FPGA design usually require to operate on different clock frequencies. Therefore, dedicated clock synthesizer modules and low-skew clock interconnection networks are provided on-chip. Some models also incorporate integrated transceiver tiles for high-speed serial I/O applications and dedicated primitives to perform complex arithmetic operations. As for today, FPGAs are in widespread use in many digital systems, allowing for design upgrades without the need for hardware modifications. The low-volume requirements to instrument facilities in the field of high-energy physics experiments can often not justify the tremendous engineering costs associated with application-specific integrated circuit designs. By contrast, FPGA based applications are operational within short development time and may offer a more cost-efficient alternative.

4.2.1 XILINX Artix-7 FPGA

The XILINX Artix-7 FPGA has been selected owing to its low power and cost specifications that perfectly meet the design constraints. The device is available in different versions that differ mainly in the quantity of the programmable logic resources.

Careful evaluation of the expected device utilization for the applied firmware designs indicated that the XC7A200T model provides a perfect trade-off between cost and processing capacity. Table 4.1 gives an overview of the available device features.

Another benchmark is the rich amount of embedded RAM blocks. The readout systems are often faced with high input rates from the detectors. However, the requirement for local buffering also scales with trigger latency. The latency parameter determines the duration of time the data must be kept in on-chip memory before a pre-selection algorithm distinguishes the physics events from background noise. The capacity for temporary data storage is crucial as well for the capability to combine the output of the different nodes into standardized chunks of data to be subsequently passed on to the following processing units.

Each FPGA model is available in a selection of packages that diverge in the number of user I/O and transceiver tiles. Different packages were selected for the TDC application and the inter-card pipelining feature, respectively. These ball grid array packages with a pitch of 1 mm are manufactured using flip-chip technology. The production process directly bonds the pads of the die, which are made conductive and coated with solder bumps, to the package substrate. The substrate again interconnects the bare chip with the external pins of the package.

Table 4.1: Summary of the Artix-7 XC7A200T-2 device features.

^a for the TDC-FPGAs

^b for the MERGER-FPGA

ⁱ see Sec. 4.2.1.1

ⁱⁱ see Sec. 4.2.1.2

ⁱⁱⁱ see Sec. 4.2.1.3

	CLB Slice ⁱ	33650
	CLB Flip-Flop ⁱ	269200
	Block RAM ⁱⁱ	365
	DSP Slice ⁱⁱⁱ	740
	f_{MAX} (BRAM / DSP)	461 MHz / 550 MHz
	LVDS I/O bandwidth	1.25 Gbit/s
	Package	FBG484^a FBG676^b
	Dimensions (mm)	23 × 23 27 × 27
	User I/O	285 400
	GTP Transceiver	4 8

4.2.1.1 Configurable Logic Blocks

The programmable logic resources of the Artix-7 FPGA are organized in Configurable Logic Blocks (CLBs). Each CLB contains a pair of slice elements accessing the switch

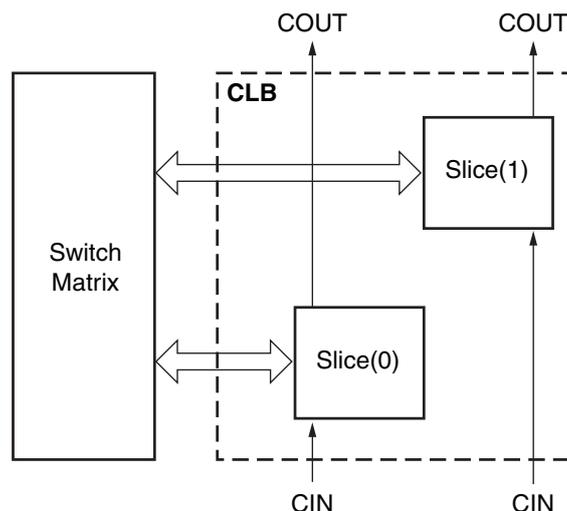


Figure 4.3: A CLB comprises two slice components, which are linked to a switch matrix accessing the routing network, to implement combinatorial and sequential logic functions. The fast look-ahead carry chain indicated with the CIN and COUT labels runs vertically up the slice columns [69, p. 9].

matrix that interconnects the CLB with the global routing network (see Fig. 4.3). A slice includes four 6-input look-up tables, eight storage elements as well as multiplexer and carry logic for arithmetic functions. No direct connection exists between the slice pair inside a CLB. However, the carry chain is cascadable in upward direction in each slice column to form wider arithmetic functions. The look-up tables are the intrinsic boolean function generators of the Artix-7 FPGA that implement arbitrary combinatorial circuitry. Each look-up table either supports a single 6-input function or two independent 5-input functions. Wider circuits are realized by combining multiple look-up tables using multiplexer interconnect logic. The storage elements register the look-up table outputs or directly sample the input signals by-passing the function generators. In general, the storage elements are configured as edge-triggered flip-flops, but four of these may alternatively be used as level-sensitive latch. In that case, the other elements remain unused. The various active-high control signals — clock enable, set/reset and write enable in particular — are shared across the sequential primitives of a slice. This property also applies to the optionally inverted clock input. The choice whether the set/reset inputs have synchronous or asynchronous behaviour is naturally common to all slice registers as well. A subset of the available slice components is capable of implementing shift registers or distributed RAM.

4.2.1.2 Embedded Memory

Apart from distributed RAM provided by the look-up tables, the Artix-7 FPGA comprises abundant embedded memory, organized in so-called Block RAM (BRAM) primitives. Each memory element can either be configured as single 36 kbit RAM or divided into two independent 18 kbit storage elements. The true dual-port configuration (see Fig. 4.4a) provides two independent access ports for concurrent synchronous read and write operations. The actual depth of the memory scales with

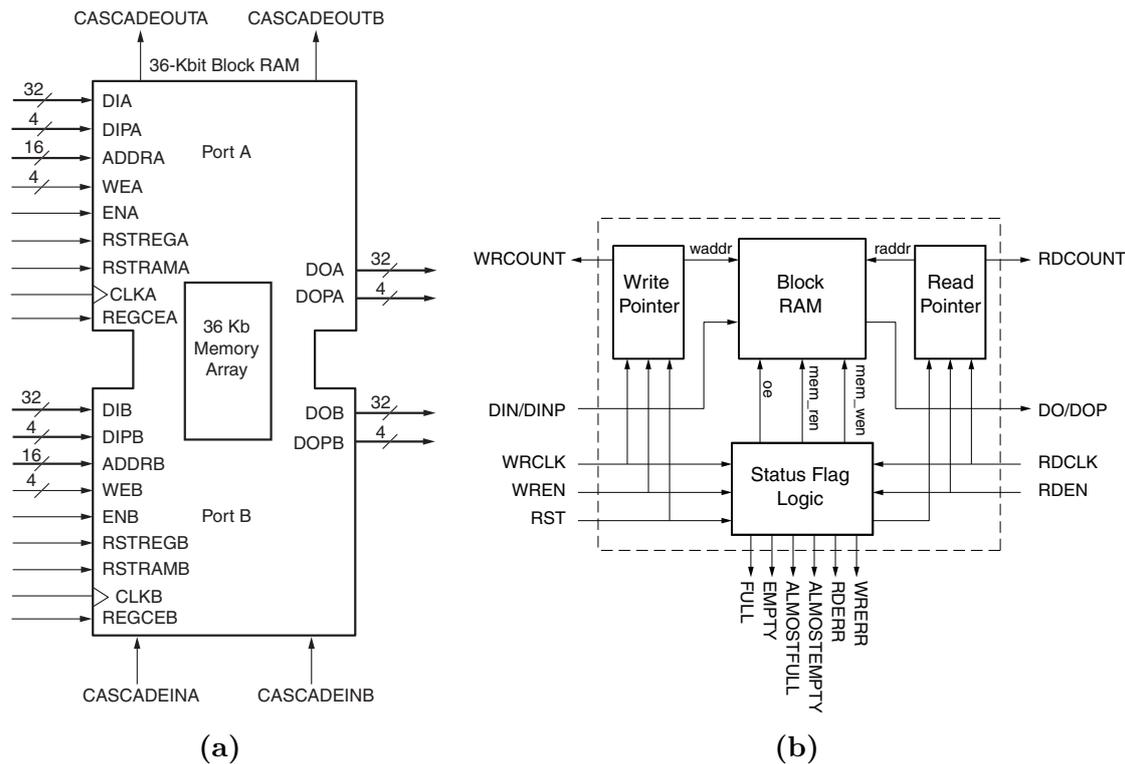


Figure 4.4: (a) The true dual-port RAM configuration features two independent access ports (Port A and Port B). Each port comes with individual address bus, data inputs/outputs and control signals [70, p. 16]. (b) The built-in FIFO support facilitates integrated logic for read/write pointer and status flag generation [70, p. 48].

the selected data width¹. Using the simple dual-port feature, the data width can be doubled to 36 bit for 18 kbit BRAM and 72 bit for 36 kbit BRAM primitives. In this mode, one port corresponds to the read-only port and accordingly the other port designates the write-only port. Wider and deeper storage areas are implemented by cascading multiple memory elements.

The TDC-FPGA design requires loads of memory to implement hit buffers storing the digitized timestamps from the individual input channels. It is furthermore necessary throughout this project to buffer the received data between independent clock domains and subsequent processing steps. This requirement is addressed by the dual-clock First-In First-Out (FIFO) memory support. For this purpose, dedicated counters, comparators and integrated logic for status flag generation are included (see Fig. 4.4b), thus saving programmable logic resources. Further integrated features like independent read/write port width selection and byte-write enable or optional output registers to increase performance are also available.

4.2.1.3 Digital Signal Processing

The Artix-7 FPGA provides dedicated primitives for DSP applications. A simplified diagram of the available operations is provided in Fig. 4.5. Each DSP48E1 slice consists of a 25-bit \times 18-bit multiplier with optional pre-adder followed by multiplexers

¹selectable data width: 1, 2, 4, 9, 18 for 18 kbit BRAM and 1, 2, 4, 9, 18, 36 for 36 kbit BRAM

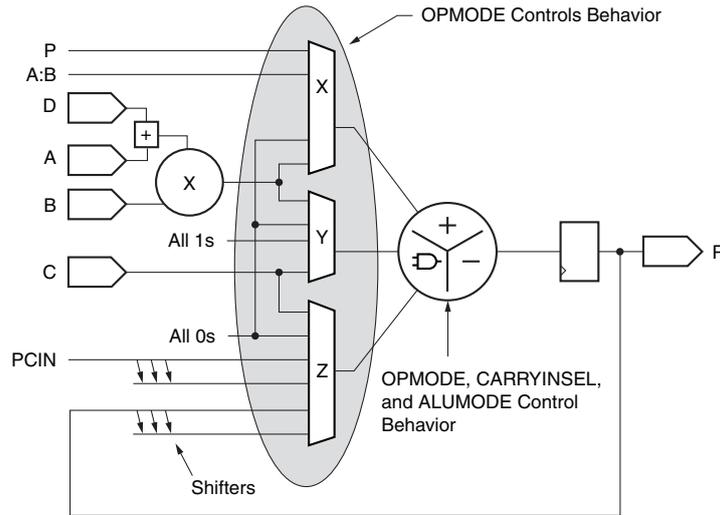


Figure 4.5: Functional diagram of the basic operations provided by the DSP48E1 slice [71, p. 25].

(X,Y,Z) selecting the inputs to an adder/subtractor or logic unit. The multiplexers are controlled by the OPMODE register, while the ALUMODE register specifies the particular add/subtract/multiply operation. When the multiplier is bypassed, the logic unit may implement bitwise logic functions on two 48-bit binary numbers. A pattern detection logic supports advanced features like the recognition of accumulator overflows or underflows and the implementation of counter resets on certain count values. Optional registers can be enabled for maximum speed. To realize more complex applications like DSP filters, the DSP48E1 slices located in the same device column can be cascaded using dedicated interconnections.

4.2.1.4 Clocking Resources

Providing the Artix-7 FPGA with high-precision, external clock signals through an optimal combination of dedicated package pins, clock buffers and clock routing resources is mandatory to achieve maximum design performance. The Artix-7 FPGA comprises global, regional and I/O clocking resources to manage different clock distribution strategies. A so-called clock region spans 50 CLB rows vertically and half the device in horizontal direction. Each clock region covers the sequential logic resources and an I/O bank consisting of 50 user I/O pins. The global clock tree is designed for low-skew and low-power applications and reaches every sequential component across the device. Independent of the global clock tree, the regional clock networks cover applications that can be restricted to a single clock region. The I/O clocking resources are tailored to implement source-synchronous interfaces.

Capabilities for clock frequency synthesis, clock deskew and jitter filtering are provided by embedded Mixed-Mode Clock Manager (MMCM) primitives. Figure 4.6 shows a functional diagram. In order to lock the circuit to the external reference clock, the frequency dividers D and M in the reference and in the feedback path must be selected such that:

$$\frac{f_{IN}}{D} = \frac{f_{VCO}}{M}, \quad (4.1)$$

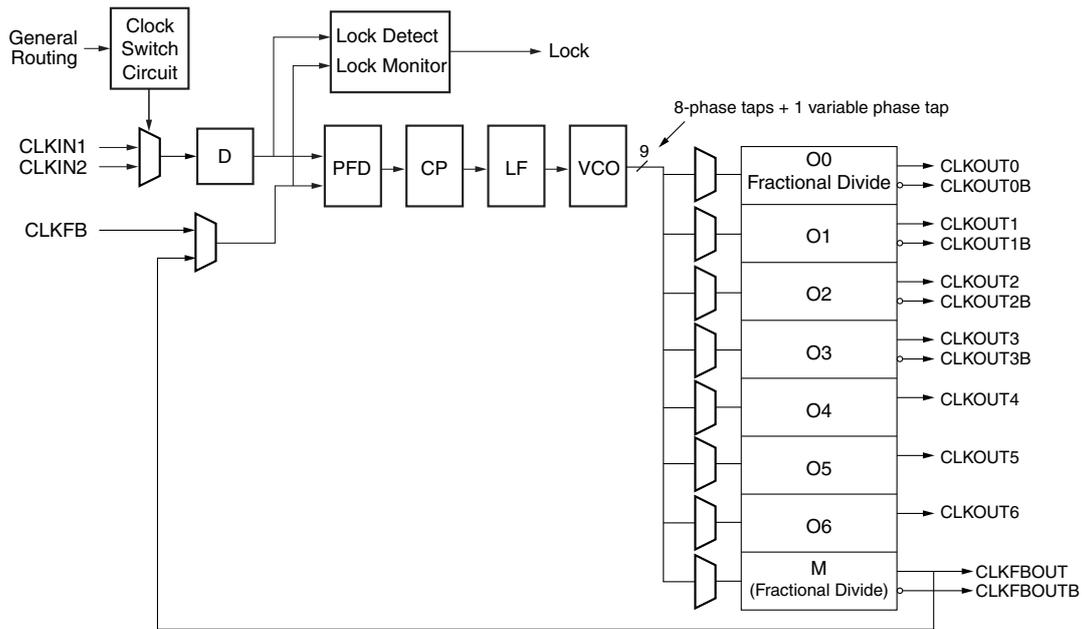


Figure 4.6: Block diagram of the MMCM circuit [72, p. 67]. The phase frequency detector (PFD) in combination with the charge pump (CP) and the subsequent loop filter (LF) synchronize the voltage-controlled oscillator (VCO) to the input clock.

where f_{IN} and f_{VCO} are the frequencies of the input clock and the Voltage-Controlled Oscillator (VCO), respectively. The clock outputs include programmable frequency divider (O_{0-6}), allowing for a wide range of clock frequencies ($f_{OUT,i}$), determined by the equation:

$$f_{OUT,i} = f_{IN} \frac{M}{D \cdot O_i}, \quad (4.2)$$

where O_i denotes the frequency divider of the corresponding clock output. The VCO provides eight output phases, each shifted by 45° , to be selected by the clock outputs and the feedback clock. Any delay added to the feedback path, which closes the loop to the phase frequency detector, results in a negative phase-shift of all clock outputs. Increasing system performance, the MMCM is commonly used in this way to compensate for clock buffer insertion delays. A comprehensive description of the clocking resources and the MMCM features is given in Ref. [72].

4.2.1.5 SelectIO Resources

The Artix-7 SelectIO drivers support a wide range of single-ended and differential signalling standards up to 3.3 V including programmable drive strength, output slew rate and integrated differential terminations. The user selectable inputs and outputs are organized in I/O banks containing 50 pins each. The dedicated Output Buffer Supply Voltage (V_{CCO}) pins of an I/O bank connect to the same external voltage rail. Thus, care must be taken when combining different signal standards within an I/O bank. The following signalling standards are employed by the ARAGORN front-end:

- 3.3 V and 2.5 V LVCMOS — miscellaneous control signals and low-speed bus interfaces.

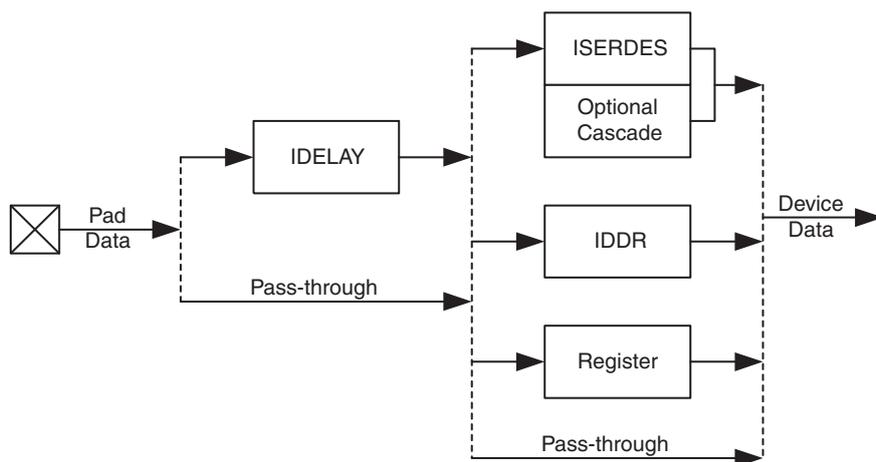


Figure 4.7: Sketch of the input datapath depicting the available primitives for data capture inside the I/O tile [73, p. 23].

- 1.8 V LVCMOS — memory interface between the on-board Flash and the MERGER-FPGA.
- 2.5 V LVDS — high-precision clock signals, TDC inputs, high-speed transceivers and source-synchronous interfaces.

Figure 4.7 shows an example of the data flow from the input pad through the I/O tile to the FPGA logic. Likewise, combinatorial interconnections are provided to directly interface the I/O drivers with the programmable logic inside the device. The I/O tiles include 31-tap, programmable input delay elements along with registers and dedicated Double Data Rate (DDR) primitives for data capture and corresponding features for serial data output and clock forwarding. High-speed source-synchronous applications are preferably implemented using the serial-to-parallel and parallel-to-serial converters, respectively. A so-called bitslip feature can be used to shift the parallel data to the correct word boundary. The advanced features provided by the I/O tiles greatly simplify timing issues encountered from signal routing delays inside the FPGA logic. Leveraging the dedicated I/O clock nets, the SelectIO resources are capable of processing data rates up to 1.25 Gbit/s.

4.2.1.6 GTP Transceiver Tiles

The Artix-7 FPGA incorporates configurable GTP transceiver tiles for high-speed serial I/O applications supporting a wide range of protocols up to 6.6 Gbit/s. The transceivers are organized in so-called quads. Each quad consists of four independent transmitter and receiver modules and is provided with two differential reference clock inputs along with dedicated routing to feed two Phase-Locked Loop (PLL) circuits (see Fig. 4.8). The clock pin pairs are internally terminated and biased to the analog power supply of the quad. The PLLs synthesize the high-speed clocks for the serial transmitter and receiver logic. Dedicated multiplexer interconnect logic gives the user high flexibility to implement different clocking schemes.

Both the transmitter and the receiver module are divided into a Physical Coding Sublayer (PCS) and a Physical Medium Attachment (PMA) sublayer. The PCS pro-

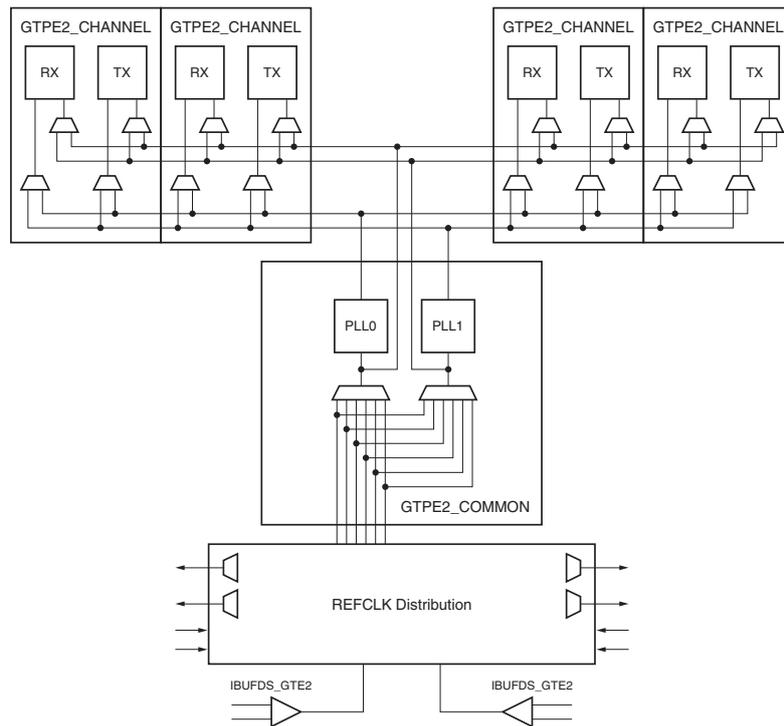


Figure 4.8: Internal structure of the GTP transceiver quad [74, p. 14]. The `GTPE2_CHANNEL` primitives comprise the transmitter (TX) and receiver (RX) modules. The PLLs (PLL0, PLL1) together with the multiplexers reside in the `GTPE2_COMMON` primitive.

vides the interface to the FPGA fabric, support for 8b/10b encoding/decoding and buffers to handle the clock transition between the parallel clock of the PMA sublayer and the FPGA fabric clock domain. The serializer/deserializer logic and the drivers of the serial transmitter and receiver reside inside the PMA sublayer. Moreover, the PMA sublayer features transmitter pre-emphasis and receiver equalization capabilities, which compensate for frequency-dependent loss in the PCB interconnects, to improve the data-eye opening at the receiver. Following the linear equalizer, the receiver PMA incorporates a clock and data recovery circuit to extract a clock from the incoming data stream. Designs that do not implement clock correction use the recovered clock as parallel receiver clock to drive the downstream FPGA logic. In this project, the GTP transceiver primitives of the MERGER-FPGA are used to handle the data traffic of the on-board optical networks.

4.2.1.7 I/O Pin Planning

I/O pin planning denotes the process of mapping logical design ports to physical device pins. This design step poses a trade-off between the requirements of the PCB design and the claim to achieve optimal data flow in and out the device. The interface pins providing the Artix-7 FPGA with configuration bitstreams are located in I/O bank 0, 14 and 15 (see Figs. 4.9 and 4.10). The configuration pins contained in bank 14 and bank 15 turn into standard user I/Os after the configuration process is completed. However, the V_{CCO} supplies are restricted to the voltage required for the selected configuration mode.

The pin planning for both the TDC-FPGA and the MERGER-FPGA was conducted using the XILINX *Vivado Design Suite* that provides a graphical interface to interactively control I/O pin placement in compliance with assigned signal standards, V_{CCO} voltages and distinct configuration modes. The elaborated pin plans have been validated during the early phase of the PCB design on the basis of provisional firmware designs, which covered the main functional blocks and the clocking networks, allowing for conclusive design rule checks.

TDC-FPGA Pin Plan

The FBG484 package comprises 484 pins out of which 285 are user I/Os. These pins are spread over I/O banks with 50 single-ended or 24 differential I/Os, respectively. Each TDC-FPGA implements 96 TDC input channels. The corresponding pin pairs are assigned to bank 15, 16, 34 and 35, which best support the routing process of the differential signal lines to the laterally located extension board connectors. A SelectMap x16 interface, which is located in bank 0 and 14, receives the configuration bitstream. The remaining multi-function pins of bank 14 and the partially bonded bank 13 are assigned to external clock inputs and residual board interfaces. Figure 4.9 shows the package pin plan of the TDC-FPGA.

The employed signal standards, LVDS for differential interconnects and LVCMOS for single-ended I/Os, are based on a V_{CCO} voltage of 2.5 V. As a result, all I/O banks of the TDC-FPGA could be connected to the same voltage rail. Dealing with a single I/O voltage greatly simplified the layout of the power planes and the placement of the decoupling capacitors in the PCB design.

MERGER-FPGA Pin Plan

The FBG676 package has been chosen for the MERGER-FPGA owing to the availability of two GTP transceiver quads. The high-speed interconnects of the GTP channels are subject to transmission line effects. Maintaining signal integrity, the FPGA device is mounted clockwise rotated by 90° on the PCB in order to minimize the length of the signal traces to the optical transceiver (SFP and CXP) sockets.

The 676 pads of the FPGA fabric are bonded out to package pins providing 400 user I/Os. Bank 13 and 34 are supplied with 3.3 V to control the analog readout modules attached to the ARAGORN front-end through the extension board connectors. These banks also include the low-speed bus interfaces of the peripheral on-board devices. When power is applied, the MERGER-FPGA retrieves the configuration bitstream from a parallel NOR Flash that requires an operating voltage of 1.8 V. This consequently defines the V_{CCO} supplies of bank 14 and 15, which comprise the memory interface to the Flash, and the configuration bank 0. The remaining four I/O banks accommodate the interconnections with the TDC-FPGAs, for instance the SelectMap x16 configuration or the source-synchronous data readout interfaces, that specify a V_{CCO} voltage of 2.5 V. Furthermore, the MERGER-FPGA receives and distributes a number of differential clock signals that are assigned to these banks alike. An overview of the package pin plan is given in Fig. 4.10.

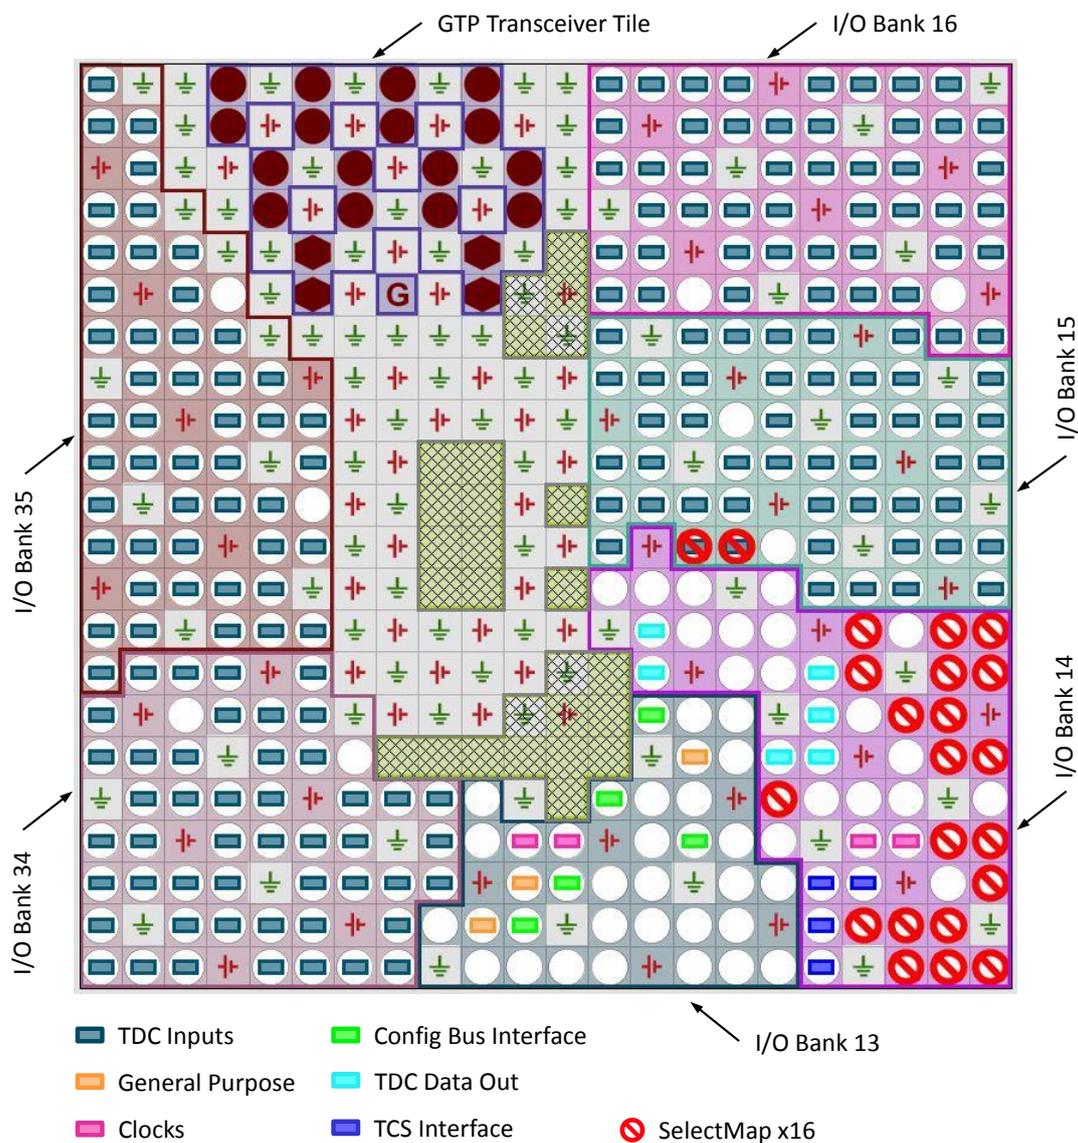


Figure 4.9: View of the I/O interfaces assigned to the package pins of the TDC-FPGA (exported from the *Vivado Design Suite*). The multi-function pins in bank 14 implement the SelectMap x16 configuration interface. All V_{CCO} supplies are connected to a common 2.5 V voltage rail.

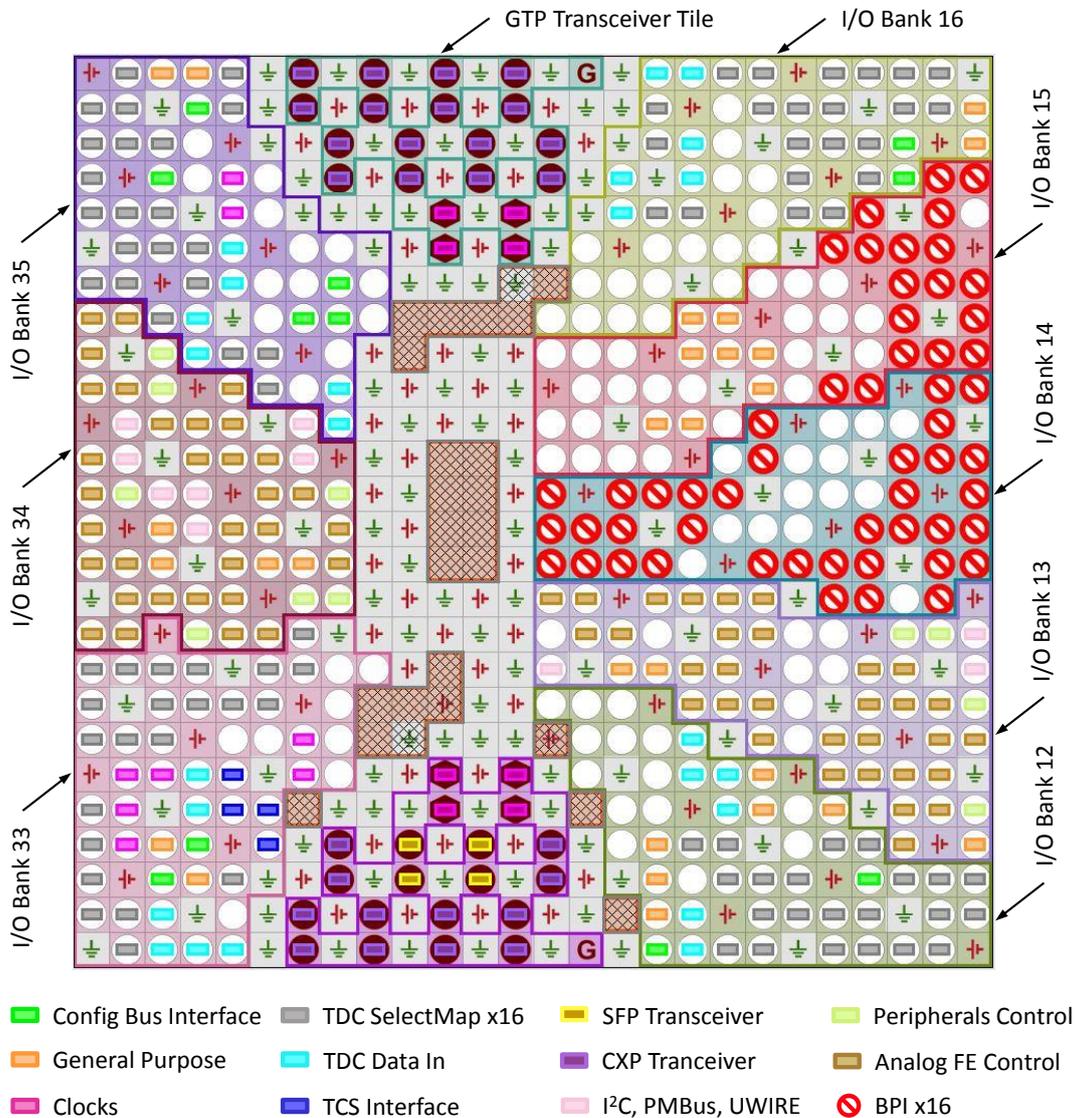


Figure 4.10: Package pin plan of the MERGER-FPGA (exported from the *Vivado Design Suite*). I/O interfaces related to the TDC-FPGAs are assigned to bank 12, 16, 33 and 35 using a V_{CC0} voltage of 2.5 V. Bank 13 and 34 provide 3.3 V-based signal standards for miscellaneous low-speed bus interfaces. Beside the transceiver channels, the transceiver quads receive two external reference clock signals each. Bank 14 and 15 host the memory interface of the configuration Flash (1.8 V).

4.3 Power Supply

The purpose of the power distribution network is to provide the on-board devices with stable supply voltages under all operating conditions. The power consumption of the ARAGORN front-end is dominated by the employed FPGA devices. The power dissipation of the optical transceiver modules is rather low¹. The power demand of an FPGA varies with programmable logic utilization, design speed and applied I/O interfaces. Nevertheless, the Xilinx *Power Estimator* spreadsheet tool [75] was consulted to evaluate the anticipated power supply requirements. A reliable assessment of the expected current draw on individual voltage rails could be retrieved for the TDC-FPGA firmware parameters that have been largely known at an early project stage. For the MERGER-FPGA design, a considerable device usage of 70 % CLB logic, 70 % BRAM primitives, 50 % DSP slices running at 400 MHz and 100 % GTP transceiver channel utilization was assumed. The simulation results are 6.6 W per TDC-FPGA and 9.8 W for the MERGER-FPGA. These specifications will provide sufficient capacity allowing for design upgrades and the implementation of future applications.

4.3.1 Voltage Regulators

The ARAGORN front-end facilitates a 4-pin power connector header² (J9), which is capable of carrying 5 A per contact, to interconnect the on-board power converters with an external power supply (12 V). Texas Instruments LMZ3 power modules [76] have been selected to generate the secondary core, auxiliary and V_{CCO} voltage rails required by the Artix-7 FPGA fabrics. These integrated DC/DC power converters achieve an efficiency of greater than 95 % and are available in pin compatible versions of 4, 7 and 10 A output current. A single external resistor is required to adjust the output voltage in the range of 0.6–5.5 V.

Under maximum load, the power consumption is dominated by the core supplies. The power distribution network incorporates four LMZ31710 devices providing up to 10 A each, which are configured for parallel operation in groups of pairs, to cope with any potential requirement emerging. The current sharing capability is enabled by connecting up dedicated pins of every device pair and synchronizing the switching frequencies to an external clock signal. As a consequence, the core voltage rail is distributed among two independent power split planes. The auxiliary voltage rail and the different I/O bank supplies are generated with LMZ31704 modules, providing up to 4 A output current. Furthermore, these modules power the miscellaneous on-board devices like optical transceiver modules, Flash memory, clock multiplier chips and fanout buffers.

The GTP transceiver quads feature separate power supplies for their internal analog constituents and integrated termination circuits. The power consumption of these primitives is rather low, but the specification requests very low-noise supply voltages. Linear regulator modules, which provide very clean output voltages, comply best with these demands. Compared to DC/DC converters, switching noise is non-existent but the power dissipation scales with the ratio between the input and

¹CXP transceiver < 3.5 W, SFP+ transceiver < 1 W

²Molex Micro-Fit

output voltage. To get around this issue, the transceiver supplies are generated from the auxiliary voltage rail using TPS74201 linear regulators [77] that can operate on very low dropout voltages. A schematic overview of the power distribution network is given in Fig. 4.11.

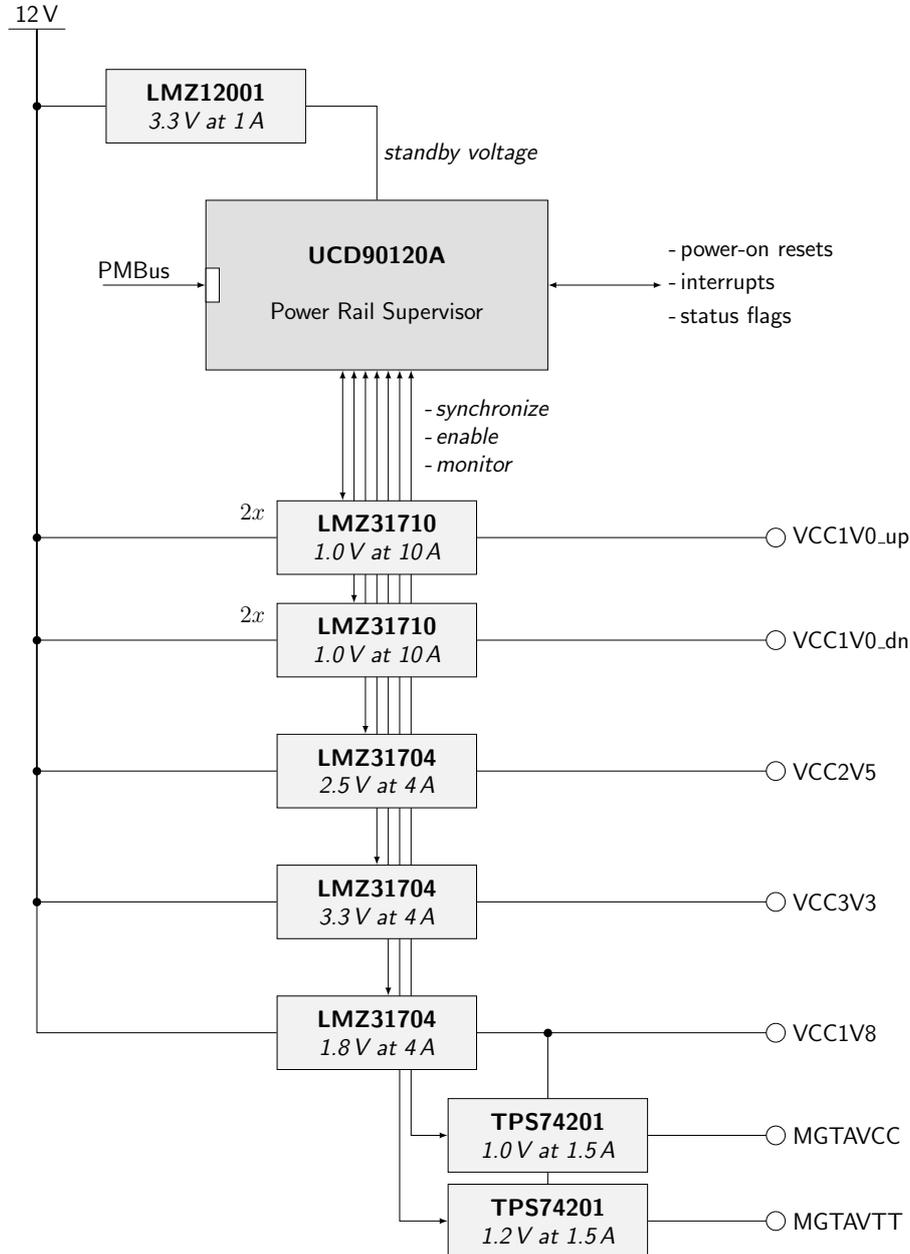


Figure 4.11: Overview of the power distribution network. The secondary voltage rails are generated from an external 12 V power supply using LMZ3 power modules and TPS74201 LDO linear regulators.

4.3.2 Power Management

The ARAGORN front-end incorporates an UCD90120A power rail supervisor [78] to ensure a reliable system start-up. The device continuously monitors the secondary voltage rails with an integrated 12-bit analog-to-digital converter. Using the configurable outputs of the supervisor chip to enable the individual voltage regulators, the

power rails can be sequenced-on and sequenced-off according to the specification. At power-up, the start-up currents of the FPGAs and the capacitive load currents, which charge the decoupling capacitors, may add up to a high surge current depending on the ramp-up rate of the voltage regulators. Avoiding such inrush currents, the power modules incorporate a built-in soft start feature to ensure a monotonic rise of the output voltages. Integrated overcurrent protection circuitry is likewise provided to protect the system from load faults. In addition, the entire power distribution network is safeguarded by a fast-acting fuse with an ampere rating of 7 A, assembled with a fuse holder (FH1) for easy replacement.

The *Fusion Digital Power Designer* software [79] provides a graphical user interface to setup the operating parameters of the power rail supervisor. Each rail definition includes the nominal output voltage, under/over voltage warnings, fault limits and power-good on/off settings. A power rail is turned on when a programmable time delay has passed and other selectable rails reach the power-good state. The same applies in reverse order for the sequence-off procedure. If a given rail does not achieve regulation within a certain time window, the device has been configured to shut down all rails and try once to re-sequence them back on. In response to a fault, which is still present after a programmable glitch time has expired, all rails are shut down immediately. Potential power failures are indicated by a red board LED (LD13). An overview of the secondary voltage rail definitions is given in Tab. 4.2.

Table 4.2: Overview of the power rail definitions.

Rail	V _{out} (V)	I _{out} (A)	Power-good	Power-good
	nominal	limit	On (V)	Off (V)
VCC1V0_up	1.0	20	0.95	0.9
VCC1V0_dn	1.0	20	0.95	0.9
VCC1V8	1.8	4	1.71	1.62
MGTAVCC	1.0	1.5	0.97	0.95
MGTAVTT	1.2	1.5	1.16	1.14
VCC2V5	2.5	4	2.38	2.25
VCC3V3	3.3	4	3.14	2.97

The general-purpose outputs provided by the power rail supervisor, of which some offer Pulse-Width Modulation (PWM) capabilities, implement simple boolean functions including optional assertion or deassertion delays on rail status flags and other general-purpose pin states. This feature is used to stall the configuration of the MERGER-FPGA and to hold the Flash in reset until all voltage rails are sequenced-on successfully in order to avoid abortive system initialization. Any erroneous completion of the configuration sequence is indicated with a red LED (LD1). Other system interrupts are asserted in response to status outputs of auxiliary on-board components. The integrated PWM functionality is used to provide a system clock and to synchronize the DC/DC converters generating the core supplies of the FPGAs.

The power supervisor chip operates on a detached standby voltage which is active as soon as the board is powered on. This standby supply implements undervoltage lookout circuitry adjusted to disable the device if the input voltage inadvertently drops below 7.5 V. Thus, early power rail switch-ons and system malfunctions are prevented. A green LED (LD2) indicates that the 12 V power is present and the standby voltage is up. The ARAGORN front-end hosts a Power Management Bus (PMBus) compliant receptacle to access the power rail supervisor for monitoring purposes and to store the user-defined device settings to on-chip nonvolatile memory. The secondary voltage rails are enabled according to these settings directly after the device comes out of reset. A diagram of the implemented ramp-up and ramp-down sequence is shown in Fig. 4.12.

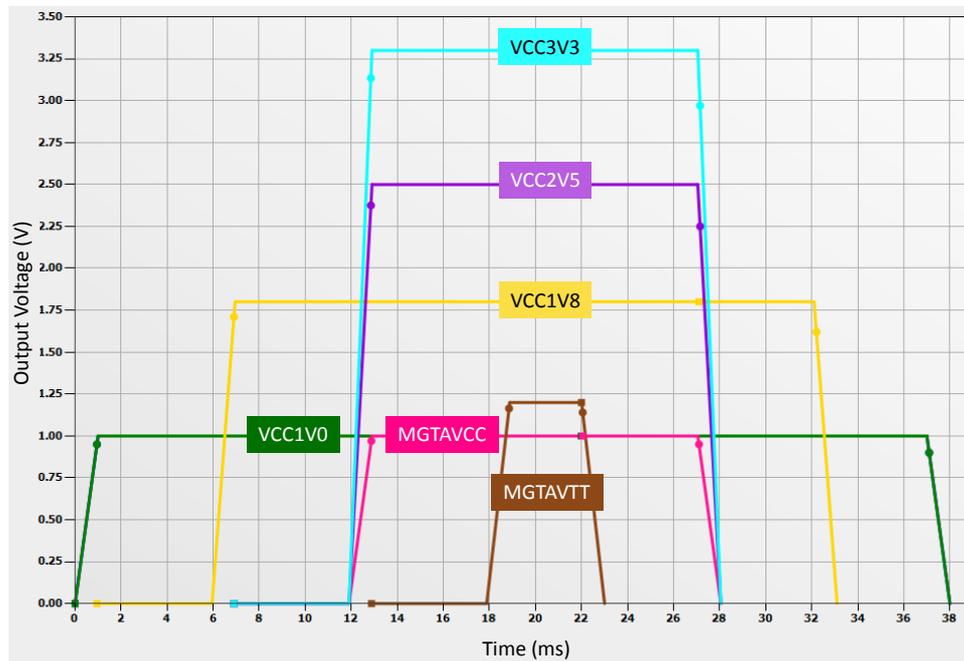


Figure 4.12: Power sequence-on and sequence-off operation for the secondary voltage rails. The auxiliary voltage rail (VCC1V8) is ramped up after the core supply (VCC1V0) is within regulation. As soon as the auxiliary rail reaches the power-good on state, the I/O bank rails (VCC2V5 and VCC3V3) are ramped up together with the MGTAVCC transceiver supply followed by the MGTAVTT transceiver supply.

4.4 Board Configuration

When the ARAGORN front-end is successfully powered up, configuration bitstreams are delivered to the FPGAs to implement the application-specific firmwares. During the startup process, the Artix-7 FPGA samples the state of dedicated mode pins to select the desired configuration mode. In the so-called master mode, the FPGA itself initiates the programming sequence and outputs a clock signal to read the bitstream data from nonvolatile memory that resides on the same host board. However, using an on-board microprocessor or CPLD to control the configuration process, which is referred to as slave mode, is preferable. It gives the user more flexibility in selection of the data source providing the bitstream image (see Ref. [80]).

Data communication between the ARAGORN front-end and a remote system can only be established by means of the optical transceiver modules via an intermediate host board, like the GANDALF module, connected to the VMEbus or similar interfaces. The optical transceiver sockets in turn directly interconnect with the GTP transceiver channels of the MERGER-FPGA. It is therefore required to boot the MERGER-FPGA from on-board nonvolatile memory first. After the optical network is established, the MERGER-FPGA provides other programmable devices like the TDC-FPGAs with configuration data received from network storage. The ARAGORN front-end comprises a 1 Gbit Micron parallel NOR Flash [81] (U7) that provides a 16-bit wide data bus. Dedicated configuration pins of the MERGER-FPGA connect to the address bus of the Flash. The main storage area of the Flash is subdivided into eight 128 Mbit partitions that can store an uncompressed bitstream each. This enables the MERGER-FPGA to load different images without losing previous versions. By default, the bitstream is automatically retrieved after power-up in asynchronous read mode. Thereby, the FPGA successively increments the Flash address and latches the received data using an internal oscillator. In asynchronous read mode, the clock speed is limited to 6 MHz, but the configuration process can be speed up enabling the synchronous read mode support of the Flash. An External Master Configuration Clock (EMCCLK), provided by a 80 MHz precision clock oscillator (U8), can be optionally enabled to substitute the internal oscillator. The configuration time is then given according to the relationship:

$$\text{configuration time} = \frac{\text{bitstream size}}{f_{\text{EMCCLK}} \times 16}.$$

Assuming a nominal bitstream length of 128 Mbit, the programming sequence is finished within 100 ms. The configuration Flash can be indirectly programmed via the JTAG interface in combination with an external programming cable (see Ref. [82]). This process is referred to as indirect, because the FPGA is in advance programmed with a bitstream that bridges the connection between the JTAG and the memory interface to the Flash. The intended use case of this approach is for debug purposes only and to provide the Flash with an initial bitstream image. The ARAGORN front-end is generally operated in radiation areas with limited access and besides that, the readout electronics are mounted to locations of the detector assembly that are accessible only with considerable effort. Hence, the capability for upgrading the bitstream image in the field is absolutely required for this project. On this account, the MERGER-FPGA design implements an embedded Microblaze processor (see Sec. 6.2.4) that enables in-system Flash programming from remote data sources. The different upgrade flows are illustrated in Fig. 4.13a.

The BPI configuration interface facilitates two revision select pins (RS[1:0]) that are wired to the most significant address bits of the Flash. These address lines are set externally to $\text{RS}[1:0] = \text{b}'01'$ using pull-up and pull-down resistors, respectively. With this setting, the FPGA initially boots the default image from an upper address space. The configuration process runs a Cyclic Redundancy Check (CRC) on received data packages. Any CRC error enables a fallback feature that ensures that the system can be recovered without manual user intervention. The fallback process actively drives the revision select pins low to load a so-called golden image from the base address of the Flash. The golden image implements a functioning design that

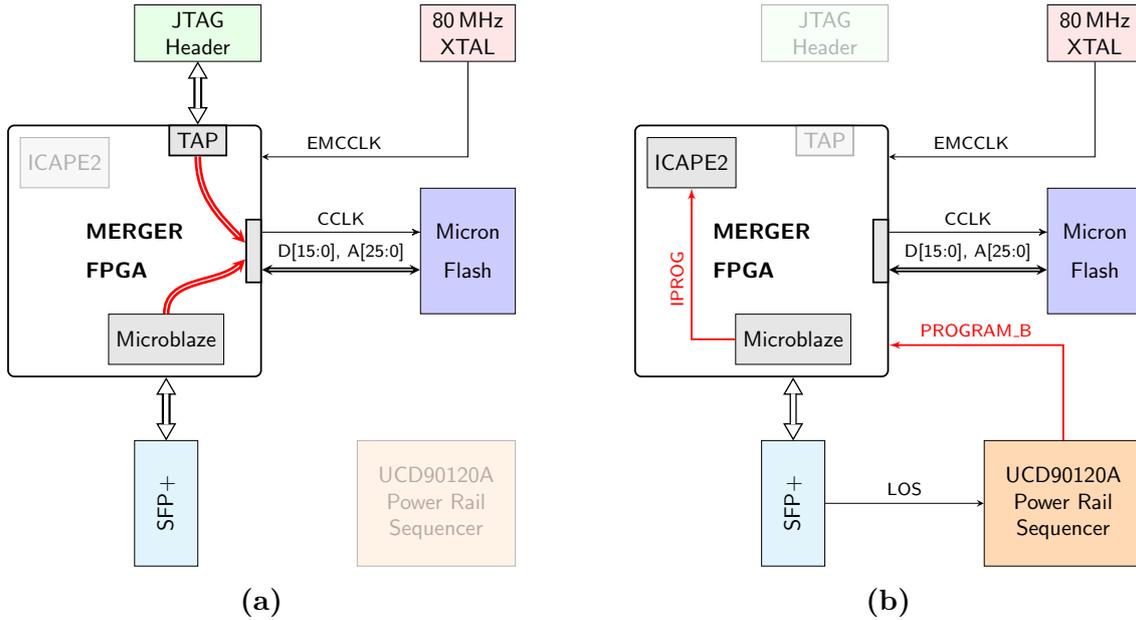


Figure 4.13: (a) Bitstream images can be delivered to the Flash either indirectly via the JTAG interface or in-system using the Microblaze processor. (b) A hard reset (PROGRAM_B) triggered by the SFP+ loss-of-signal (LOS) flag and a software interrupt (IPROG) can be used to reprogram the MERGER-FPGA during operation.

can fix corrupted memory contents and dynamically upgrade the default bitstream in the field. Apart from CRC errors, fallback is triggered as well if the configuration process does not complete and the address counter wraps around or an optionally enabled watchdog timer expires.

The design foresees two different strategies to reprogram the MERGER-FPGA during operation (see Fig. 4.13b). First, in the event that the remote connection to the ARAGORN front-end has been lost, a hard reset is executed from the power rail supervisor chip (cf. Sec. 4.3.2) whenever the loss-of-signal flag of the SFP+ module is cleared. In order to trigger this event, the user simply needs to toggle the laser on the sender side. For a complete reinitialisation of the readout system, this procedure is first carried out on the host board, e.g. the GANDALF module, to reprogram the MERGER-FPGA on the ARAGORN master. Once the master board is operational again, the connected slave boards are recovered by toggling the transmitters of the CXP module. The second option issues an Internal PROGRAM (IPROG) command through the internal configuration access port (ICAPE2) primitive [83] which provides an interface to the configuration registers of the Artix-7 FPGA. A dedicated interrupt to the Microblaze processor calls a service routine performing this task. Optionally, the warm boot start address register, which specifies the Flash address the bitstream is retrieved, can be modified. Thus, this multiboot function allows the user to switch between different bitstream images in order to enable in-system debugging of new firmware versions. If the upgraded firmware has been approved, the warm boot start address register and the IPROG command can be embedded inside the default image to automatically load the upgraded bitstream at power-up. Table 4.3 shows the relevant section of an exemplary bitstream that must be modified to enable the IPROG option. A flow diagram of the fallback and

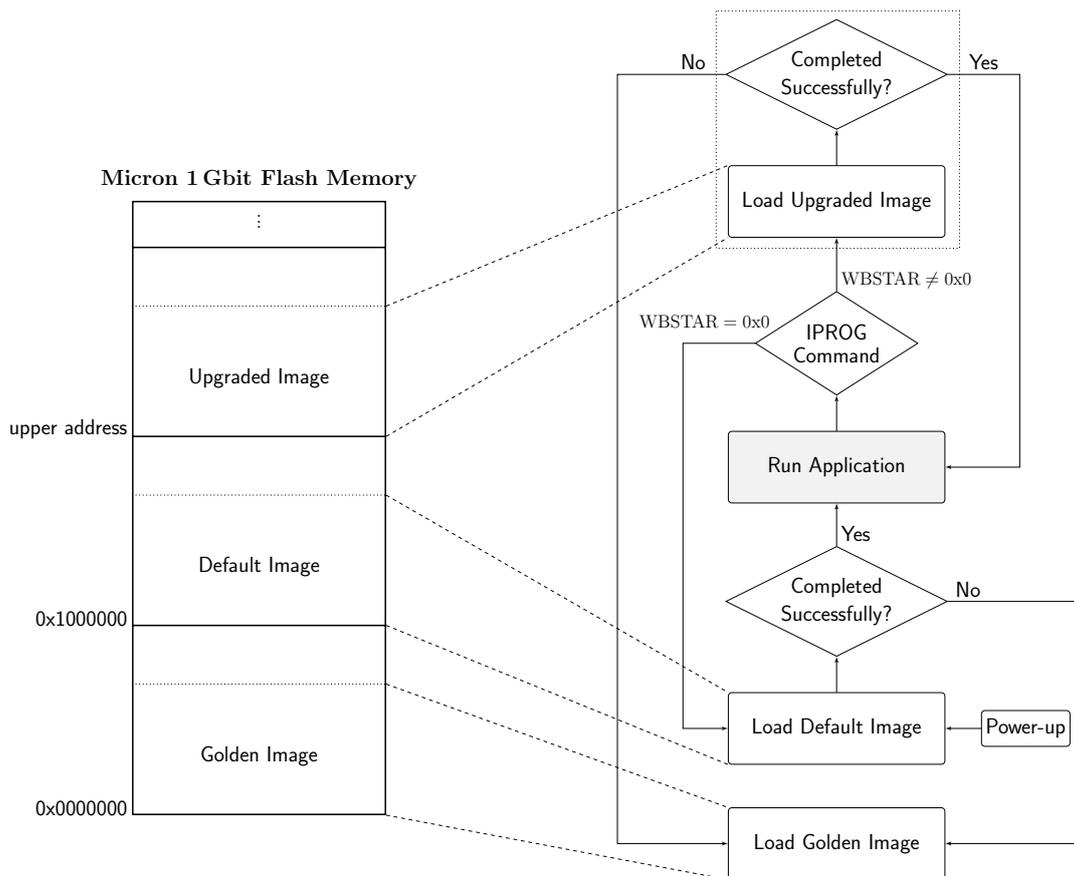


Figure 4.14: Flash memory allocation (left) and flow chart of the fallback and multiboot features (right). Following the IPROG command, the MERGER-FPGA is reprogrammed from the Flash address specified in the warm boot start address register (WBSTAR).

multiboot features is shown in Fig. 4.14. The configuration of the TDC-FPGAs is controlled by the Microblaze processor. The MERGER-FPGA provides individual SelectMap x16 interfaces to program the TDC-FPGAs in parallel. Bitstream images are delivered from remote storage in the same way as the Flash programming files. However, the data is not buffered in the configuration memory. Instead, the Microblaze processor directly outputs the data packages on four SelectMap x16 busses to the TDC-FPGAs. Just as the MERGER-FPGA, the TDC-FPGAs can be reprogrammed at any time of operation using a dedicated interrupt to the embedded processor.

Table 4.3: IPROG command (red color) embedded in an exemplary bitstream that triggers a reboot from the Flash address specified in the Warm Boot Start Address Register (red color), adapted from Ref. [80, p. 142].

Bitstream Data 32-bit hex	Explanation
FFFFFFFF	Dummy Word
AA995566	Sync Word
20000000	Type 1 NO OP
30020001	Type 1 Write 1 Words to WBSTAR
00000000	Warm Boot Start Address
30008001	Type 1 Write 1 Words to CMD
0000000F	IPROG Command
20000000	Type 1 NO OP

4.5 Clocking Networks

The ARAGORN front-end recovers the reference clock of the COMPASS-II experiment from the SFP+ transceiver link. This reference clock is employed as system-wide sampling clock in the TDC application and must therefore be used to drive the CXP channels that redirect the up-link data stream to the slave boards. The strict jitter specifications demand for an on-board jitter attenuator device that delivers a clean, phase-synchronous copy of the reference clock. Free-running clocks of different frequencies are likewise required for miscellaneous board functions. A schematic overview of the implemented clock structure is shown in Fig. 4.15.

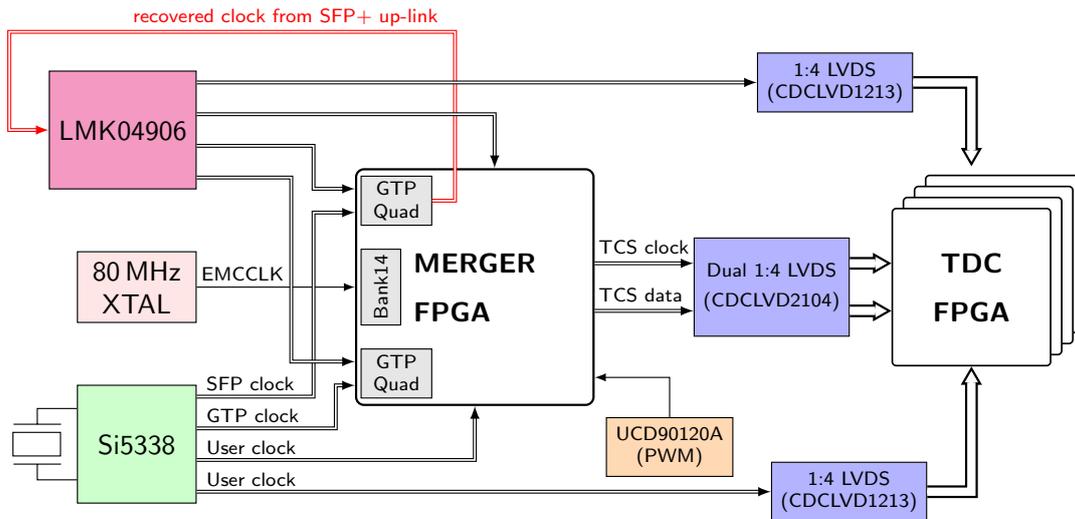


Figure 4.15: Clock structure of the ARAGORN front-end. Relevant devices are detailed in sections 4.5.1 and 4.5.2. The recovered clock (in red color) extracted from the SFP+ up-link inside the GTP receiver is forwarded to the LMK04906 jitter attenuator device. The jitter attenuator delivers a clean copy of the recovered clock to the GTP transmitter channels and the TDC-FPGAs. Moreover, this clock is redirected to the MERGER-FPGA driving the interface between the GTP receiver and the FPGA logic. The recovered clock is equivalent to the reference clock of the TCS. The TCS information received from the up-link is forwarded along with the TCS clock to the TDC-FPGAs.

4.5.1 Clock Sources

The Silicon Labs Si5338 [84] (U11) clock generator functions as programmable clock source providing four independent, differential output clocks. The Si5338 device is operated in free-running mode using an external 25 MHz crystal. Frequency synthesis is performed with an integrated PLL providing a VCO operating range of 2.2–2.84 GHz, followed by a fractional divider stage to produce arbitrary output frequencies up to 350 MHz. In this application, all clock outputs are configured as LVDS drivers.

To adapt the system to distinct readout tasks, the design of the ARAGORN front-end aims to support any desired transceiver line rate up to 6.6 Gbit/s. Therefore, two Si5338 clock outputs connect to the reference clock inputs of both GTP transceiver

tiles. The two remaining clock outputs provide user clocks to the FPGA fabrics. In case of the TDC-FPGAs, the corresponding clock output feeds a 1:4 fanout buffer that delivers a copy to each TDC-FPGA. The elaborated frequency plan of the Si5338 device and other on-board clocking devices is listed in Tab. 4.4. The device configuration in terms of a register map is created leveraging the *ClockBuilder* software tool [85]. After power-up, the register map is stored in-system to internal RAM of the Si5338 device via an I²C interface (see Sec. 4.7.2).

Table 4.4: Frequency plan of the on-board clocking devices.

Device	Receiver	Objective	Frequency
Si5338	GTP transceiver	Reference clock	155.52 MHz
	GTP transceiver	Reference clock	155.52 MHz
	MERGER-FPGA	User clock	200 MHz
	TDC-FPGAs	User clock	200 MHz
ECS-3518	MERGER-FPGA	EMCCLK	80 MHz
UCD90120A (PWM)	MERGER-FPGA	System clock	40 MHz

4.5.2 Jitter Attenuator

As outlined in the introduction to this chapter, the strict jitter specifications forbid to directly employ the unfiltered recovered clock as reference for the GTP channels that set up the star topology interconnections to the slave boards. Besides the ability to sufficiently remove jitter from the recovered clock, the designated circuit must maintain a deterministic phase to the input clock after a power cycle or loss-of-lock. These prerequisites are fulfilled by the Texas Instruments LMK04906 [86] (U31) jitter attenuator operated in zero-delay mode.

The device incorporates two PLLs that consist of a phase frequency detector, which compares the frequency of a tunable oscillator to the reference input, followed by a charge pump. Both frequencies are divided by programmable counters which are selected such that the divided frequencies coincide at the input to the phase frequency detector. The charge pump outputs a current, which is proportional to the phase offset, to a loop filter. Passing the loop filter, the current from the charge pump is converted into a voltage steering the tunable oscillator. Thereby, the loop filter acts as low-pass, aiming to remove jitter from the reference clock (see Ref. [87]).

The dual PLL architecture of the jitter attenuator is depicted in Fig. 4.16. The device provides superior jitter cleaning capabilities and facilitates six independent clock outputs. The output buffers can be programmed to support several single-ended or differential formats. In this project, all clock outputs are configured to use LVDS. The first PLL (PLL1) receiving the external reference clock is operated with an external Voltage-Controlled Crystal Oscillator (VCXO). Suppressing any high-offset frequency phase noise contained in the reference clock, the loop filter of PLL1 implements a narrow loop bandwidth using external components. The VCXO then outputs a low-jitter version of the reference clock to the second PLL (PLL2)

for frequency multiplication using an internal VCO. The partially integrated loop filter of PLL2 has been designed for a wide loop bandwidth to benefit from the low high-offset phase noise of the VCO.

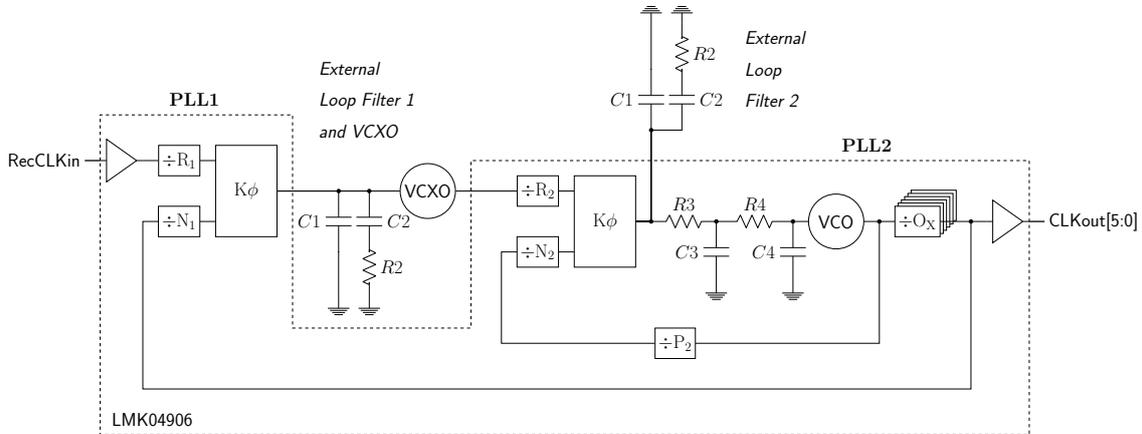


Figure 4.16: Simplified block diagram of the LMK04906 jitter attenuator.

The internal VCO of PLL2, which has a tuning range of 2370–2600 MHz, feeds the output buffers that include individual output dividers to produce the desired output frequencies up to 2.6 GHz. Depending on the selected divider values, there are multiple possible phases the output clocks can be aligned to the reference input. However, an internal feedback loop from the clock outputs back to the phase frequency detector of PLL1 can be enabled in the device to assure a fixed phase relationship to the reference input which is fundamental for the constant-latency feature of the star topology network. In this zero-delay mode, the output clock with the lowest frequency is connected to the feedback input of PLL1, satisfying the relationship:

$$\frac{f_{RecCLKin}}{R_1} = \frac{f_{VCO}}{N_1 \cdot O_x}, \quad (4.3)$$

where f_{VCO} is the frequency of the VCO in PLL2 and R_1 , N_1 and O_x denote the frequency divider at the phase frequency detector of PLL1 and the corresponding clock output. The feedback frequency of PLL2 must coincide with the frequency f_{VCXO} of the VCXO, according to:

$$\frac{f_{VCXO}}{R_2} = \frac{f_{VCO}}{N_2 \cdot P_2}, \quad (4.4)$$

where R_2 , N_2 and P_2 denote the frequency divider at the phase frequency detector and in the feedback path of PLL2. All external components of the LMK04906 device, among other things, the loop filters and the VCXO, have been particularly optimized for this project. The frequency of the recovered clock extracted from the SFP+ transceiver link is 155.52 MHz, four times the reference clock frequency of the COMPASS-II experiment. In selection of reasonable divider settings satisfying Eqs. (4.3) and (4.4), a 19.44 MHz VCXO component has been selected. The loop filter of PLL1 is designed as a 2nd order loop filter that is solely implemented with external resistor and capacitor components. In contrast, the external loop filter of PLL2 can be extended with integrated components to a higher order loop

filter. The particular values of the integrated resistors and capacitors are selected during device configuration. To achieve the best possible jitter performance, both loop filter designs have been simulated with the *Clock Design Tool* software [88] in consideration of the phase noise inherent with the recovered clock. Based on these studies, the optimized loop filter components, which correspond to loop bandwidths of 100 Hz for PLL1 and 111 kHz for PLL2, are listed in Tab. 4.5. Figure 4.17 shows a diagram of the simulated phase noise profile.

Table 4.5: Selected loop filter resistors and capacitors.

PLL #	external			internal			
	C1	C2	R2	C3	R3	C4	R4
1	330 nF	14.7 μ F	270 Ω	–	–	–	–
2	39 pF	3.3 nF	1.5 k Ω	10 pF	200 Ω	10 pF	200 Ω

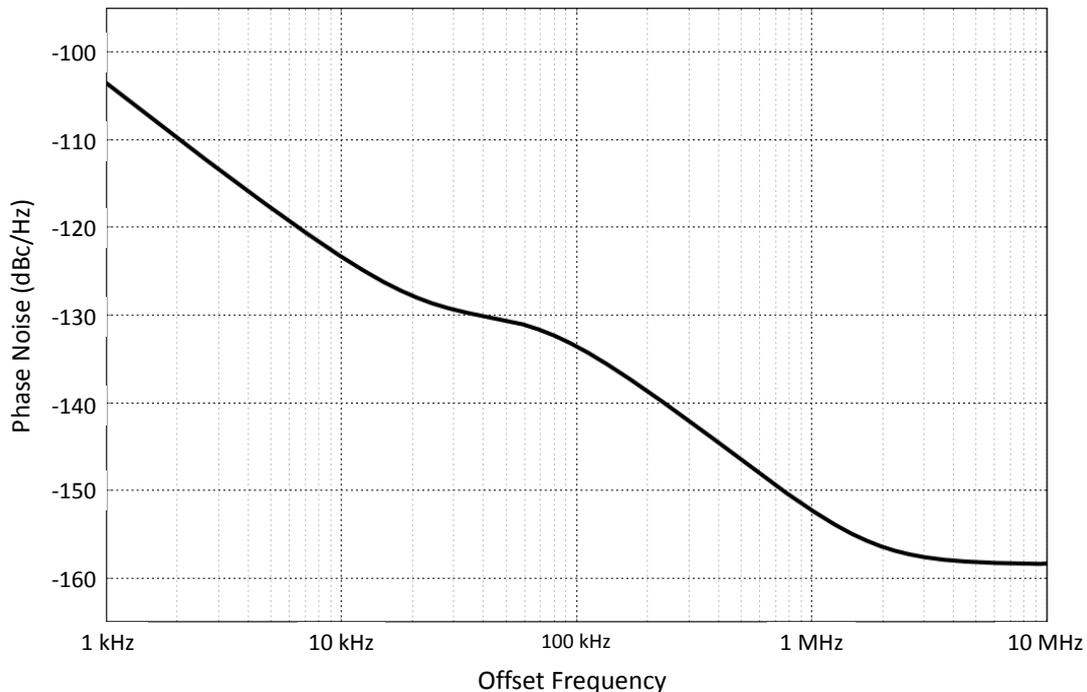


Figure 4.17: Simulated phase noise performance of the LMK04906 device (LVDS output at 155.52 MHz) in consideration of the jitter inherent with the recovered clock.

A comparison between the resulting phase noise performance of the LMK04906 device, the measured phase noise of the recovered clock and the recommended limits of the GTP transceiver reference at different offset frequencies is shown in Tab. 4.6. These numbers demonstrate that a jitter attenuator circuit is absolutely necessary for this application and that the implemented design complies with the specified requirements. A summary of the underlying device settings is provided in Tab. 4.7. The desired configuration can be selected with the *CodeLoader* software tool [89] that produces a register map to program the device through a Microwire interface (see Sec. 4.7.3).

Table 4.6: Phase noise of the recovered clock, the recommended phase noise limits of the GTP reference clock and the phase noise of the jitter attenuator clock outputs in dBc/Hz at different offset frequencies.

Offset	Recovered Clock (155.52 MHz)	GTP Ref. Clock [90] (156.25 MHz)	LMK04906 Clock Out (155.52 MHz)
10 kHz	-105	-121	-123
100 kHz	-115	-129	-133
1 MHz	-122	-133	-152
RMS Jitter 1 kHz to 10 MHz	8 ps	N/A	337 fs

Table 4.7: Outline of the device settings.

VCXO	Frequency	19.44 MHz
	Gain	1 kHz/V
VCO	Frequency	2488.32 MHz
	Gain	18 MHz/V
PLL1	PDF	19.44 MHz
	$K\phi$	1.6 mA
	Phase Margin	70°
	Bandwidth	100 Hz
PLL2	PDF	19.44 MHz
	$K\phi$	3.2 mA
	Phase Margin	70°
	Bandwidth	111 kHz

4.6 Fiber Optics

Fiber optic communication uses modulated light to transmit digital information along strands of glass fiber. A fiber optic cable consists of a core made of silica glass surrounded by a cladding and a protective jacket. According to the principle of total internal reflection, the index of refraction of the core must be greater than that of the cladding to confine the light within the core on its path through the cable.

There are two primary types of optical fibers – multi-mode and single-mode fibers – that are classified by their core diameter. Multi-mode cables have a core diameter of 50–62.5 μm . The thickness of the core allows the light rays to travel various paths. This results in varying propagation velocity for the different modes of light that propagate through the fiber. The spread of the arrival time, also known as modal dispersion, causes a distortion of the light pulse at the receiving node. Therefore,

multi-mode fibers are typically employed in short distance applications ranging up to 500 m. The international standard ISO/IEC 11801 defines three categories – OM1, OM2, OM3 and OM4 – of multi-mode cables. For instance, laser-optimized OM3 class fibers have a core-to-cladding ratio of 50/125 μm and are designed for 850 nm VCSEL¹ technology, supporting 10-Gigabit Ethernet (10GBASE-SR) links on distances up to 300 m. In contrast, single-mode fibers have a smaller core diameter of about 9 μm , allowing only a single mode of light to propagate. This has the benefit that signal attenuation is reduced due to the decreased number of reflections which permits to instrument long-distance links. For this type of fiber, the bandwidth is limited by chromatic dispersion that contributes to the deterministic jitter portion obtained at the receiver node. However, single-mode fiber products are costly as the interface to the optical transceiver is rather complex to implement.

Compared to copper cabling solutions, optical networking meets higher bandwidth applications over longer distances and offers a number of advantages like low loss, small size, light weight equipment. The front-end electronics employed in high-energy physics experiments are usually very noise-sensitive due to the small response of the particle detectors. A key design principle therefore is to minimize the impact of electromagnetic radiation that mainly originates from the digital part of the readout chain. Using fiber optic communication electrically decouples the sending and receiving nodes. This feature is highly desirable for the multi-tiered arrangement of the ARAGORN front-ends to avoid ground loops that can increase the level of background noise in the system by picking up electromagnetic energy through inductive coupling.

SFP Transceiver Socket

The Small Form-factor Pluggable (SFP) transceiver [91] is by now the standard interface for many applications using fiber optic data communication. The full-duplex, hot-pluggable SFP module provides an electrical interface to the host board and an optical interface that can be attached to different types of fiber optic cables to support the line rate and link distance required for the desired application. The ARAGORN front-end comprises a SFP/SFP+ compliant transceiver socket that connects to the GTP transceiver tile of the MERGER-FPGA. In the down-link direction, the SFP module allows for high-speed data readout up to 6.6 Gbit/s. The up-link instead provides an interface for slow control tasks and distributes the reference clock and the trigger primitives along with the corresponding event labels.

CXP Transceiver Socket

The CXP specification is part of the Infiniband architecture [92] that describes an interface for pluggable transceivers with 12 full-duplex lanes supporting data transmission rates in excess of 10 Gbit/s per lane. The ARAGORN front-end hosts a CXP compliant receptacle (see Fig. 4.18a) to enable the communication with up to seven boards through a star topology. For this purpose, seven transmit and seven receive lanes of the CXP socket are linked to the remaining GTP transceiver channels of the MERGER-FPGA. The connector layout permits to retrospectively mount

¹Vertical-Cavity Surface-Emitting Laser

the receptacle to the host board, enabling the CXP interface as required by the final application.

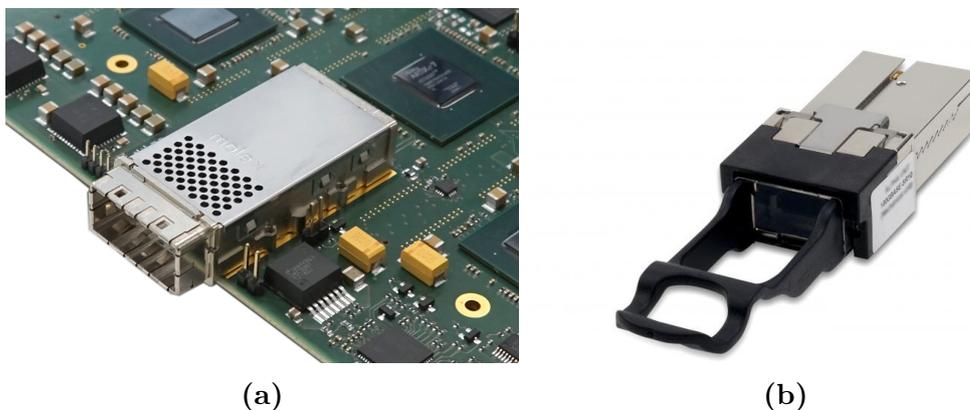


Figure 4.18: (a) CXP transceiver socket on the ARAGORN front-end. (b) Finisar FTLD10CE1C transceiver module.

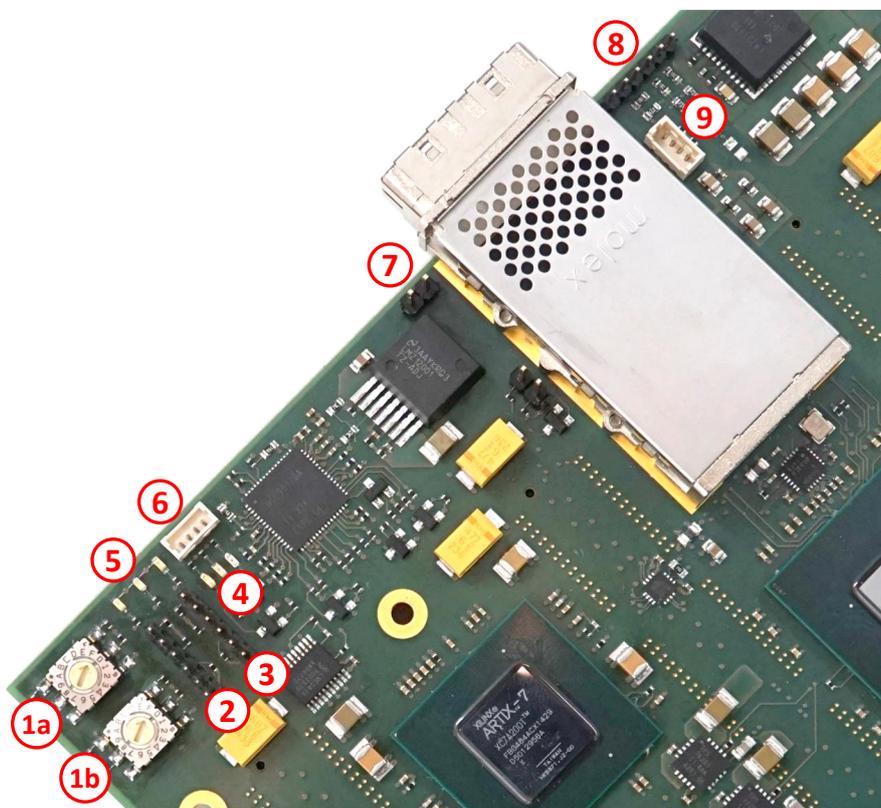
The Finisar FTLD10CE1C transceiver module [93], which is shown in Fig. 4.18b, has been selected to operate the star topology network. The device provides a 24-fiber, high-density receptacle for parallel optical communication over multi-mode fiber supporting a variety of industry formats. A fiber optic breakout cable interconnects the CXP module attached to the master front-end with the SFP slots on the satellite boards. The hot-pluggable CXP form factor has the advantage that this costly module is only required once to operate a set of eight boards.

4.7 Interfaces

The ARAGORN front-end comprises four high-speed connectors, providing an external interface for extension boards. Different industry-standard interfaces for communication with miscellaneous on-board components are furthermore available. Some of these are mastered solely by the MERGER-FPGA, while others are also accessible for debugging or monitoring purposes via external programming cables. Moreover, six LEDs are provided on the ARAGORN front-end to indicate the board status or system faults. Figure 4.19 shows a detailed view of the ARAGORN front-end for locating the different board interfaces.

4.7.1 Extension Board Connectors

The extension board connector provides 208 pins that are divided into four individual pin banks (see Fig. 4.20). The pins located on either side of each bank are separated by integral ground planes. The device has been selected due to its increased contact wipe, large number of mating cycles (up to 1000) and optional guide posts to provide a durable connection to extension boards comprising the analog front-end electronics like preamplifier and discriminator modules. Each TDC-FPGA allocates 192 pins corresponding to 96 differential input pin pairs. The differential input buffers of the TDC-FPGAs are configured for the LVDS format. Allowing for a space-optimized PCB layout, the TDC-FPGA design makes use of built-in, differential $100\ \Omega$ terminations that save loads of discrete board components.



-
- 1a Board ID (bits 4 to 7)
 - 1b Board ID (bits 0 to 3)
 - 2 JTAG: MERGER-FPGA
 - 3 JTAG: TDC-FPGAs
 - 4 LEDs (from left to right): LD2, LD1, LD13
 - 5 LEDs (from left to right): LD15, LD16, LD17
 - 6 PMBus header
 - 7 Power inhibit (PMBus_CTRL)
 - 8 Microwire header
 - 9 I²C header
-

Figure 4.19: The external interfaces are located alongside the board edge to maintain accessibility when the ARAGORN front-end is equipped with an optional heat sink.

The remaining 16 pins of each connector are attached to 3.3V I/O banks of the MERGER-FPGA to support the widely used LVTTTL or LVCMOS signal standard, implementing a general purpose interface to the extension boards. The connector pin-outs are listed in Appendix B.

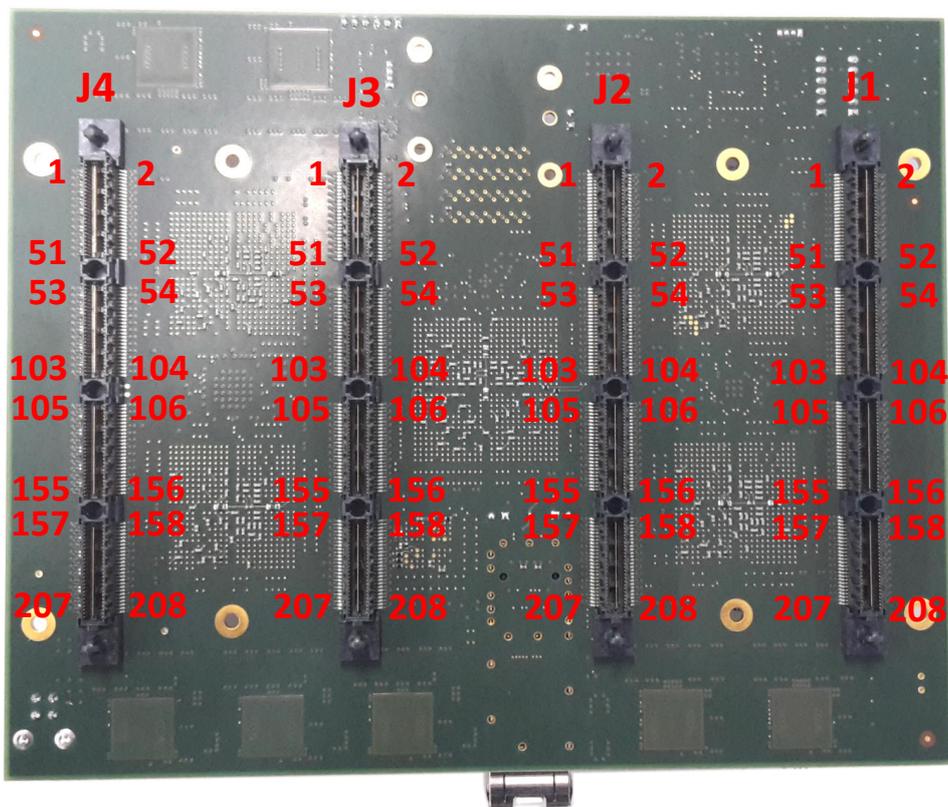


Figure 4.20: The Samtec QMS-104 connectors provide the interface for extension boards. The pin numbers of each pin bank are partially labeled in red color to enhance the understanding of the connectivity tables listed in Appendix B.

4.7.2 I²C and PMBus

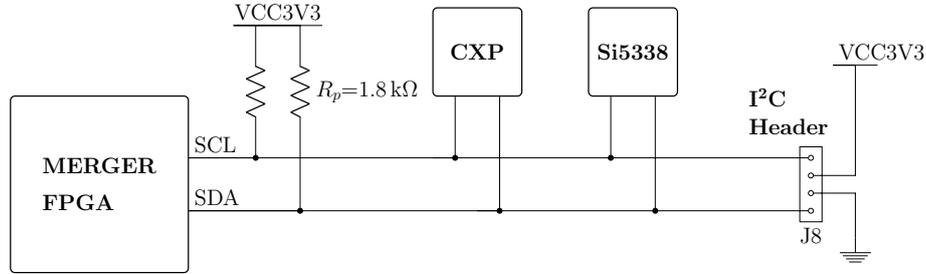
The Inter-Integrated Circuit (I²C) bus interface [94] is a popular serial bus specification for inter-IC control. The interface consists of two bidirectional lines – serial data (SDA) and serial clock (SCL) – for communication between a master (optionally multiple masters) and one or more slave ICs. Each slave device is identified by a unique, 7-bit software address. The slave addresses assigned to the different devices on the ARAGORN front-end are listed in Tab. 4.8. To avoid an address conflict between the SFP and CXP transceiver, the SFP module is connected to a separate bus that is driven by the MERGER-FPGA only. Furthermore, the main I²C bus is accessible through a pin header in combination with an external programming cable. A schematic diagram of the implemented bus architecture is shown in Fig. 4.21.

Devices are attached to the I²C bus using open-drain transistors. An open-drain output either actively drives the bus line low or disconnects the device from the bus when it is switched off. To make sure the devices on the bus see a valid logic high when the bus is left floating, the bus lines are connected via pull-up resistors

Table 4.8: Listing of I²C slave addresses.

Bus #	Address (7-bit)	Device	Description
1	50h	CXP transceiver	Transmitter Functions
	54h	CXP transceiver	Receiver Functions
	70h	Si5338	Configuration RAM
2	50h	SFP transceiver	Serial ID Memory Map
	51h	SFP transceiver	Digital Diagnostic Interface

to a positive supply voltage. This wired-AND structure permits multiple devices to share the same interconnect lines without the risk that bus conflicts end up in a short circuit when more than one device accidentally initiates a bus transfer at the same time.

**Figure 4.21:** Topology of the main I²C bus.

The I²C specification defines different operating speeds, ranging from bit rates of 100 kbit/s in Standard-mode, through 400 kbit/s in Fast-mode and up to 3.4 Mbit/s in High-speed mode. Each mode specifies a maximum rise time that is related to the time constant given by the capacitive load of the bus lines and the value of the pull-up resistors. In consideration of the voltage waveform for a charging capacitor $V(t) = V_{DD}(1 - e^{-t/RC})$, the rise time is defined as the period of time between $V_{IL} = 0.3 \cdot V_{DD}$ and $V_{IH} = 0.7 \cdot V_{DD}$. The maximum value of the pull-up resistor $R_{p(max)}$ can then be expressed as a function of the rise time t_r and the bus capacitance C_b [94, p. 55]:

$$R_{p(max)} = \frac{t_r}{0.8473 \cdot C_b}. \quad (4.5)$$

Conversely, the minimum specified sink current I_{OL} defines the minimum value of the pull-up resistor $R_{p(min)}$ that has to be large enough to drive the bus lines to a valid low-level voltage V_{OL} [94, p. 55]:

$$R_{p(min)} = \frac{V_{DD} - V_{OL}}{I_{OL}}. \quad (4.6)$$

Regarding the contributions from the device pins, the PCB traces and the optional off-board cabling, the parasitic capacitance of the main I²C bus is roughly 100 pF.

The slave peripherals attached to the bus are operated in Fast-mode using a supply voltage of $V_{DD} = 3.3\text{ V}$. Taking into account the parameters $t_r = 300\text{ ns}$, $V_{OL} = 0.4\text{ V}$ and $I_{OL} = 3\text{ mA}$ from the specification, the value of the pull-up resistors should be in the range of $0.97\text{--}3.5\text{ k}\Omega$. The selected pull-up resistance of $1.8\text{ k}\Omega$ is a trade-off between operating frequency and power consumption. A similar consideration yields a value in the range of $0.97\text{--}11.8\text{ k}\Omega$ for the I²C bus accessing the SFP transceiver module.

The Power Management Bus (PMBus) specification [95], which is an extension to the SMBus standard, defines a protocol for managing power modules via a serial interface based on I²C. The power rail supervisor on the ARAGORN front-end, which is mastered by the MERGER-FPGA, complies with the PMBus specification and can be operated at either 100 kHz or 400 kHz . Likewise, external host controllers can access the PMBus interface for configuration file download and monitoring of the power supplies via a connector header. The implemented PMBus topology is shown in Fig. 4.22.

Besides the standard PMBus commands transmitted through the serial interface, the PMBus specification permits to add pin-based functions, for instance with the control (CTRL) input that can be used to sequence-on and sequence-off the power rails. In this project, the CTRL signal is asserted with a jumper (J15) to inhibit the on-board voltage regulators. This is a useful feature to prevent unsupervised switch-ons of the voltage rails before the power rail supervisor has been programmed with a valid configuration at initial board operation. The PMBus address of the power rail supervisor is decoded from dedicated device pins. However, some addresses of the 7-bit address space are reserved for specific devices by the SMBus specification [96]. Again others can be restricted by the device manufacturer. The available address ranges for the power rail supervisor are 1–10 and 13–125.

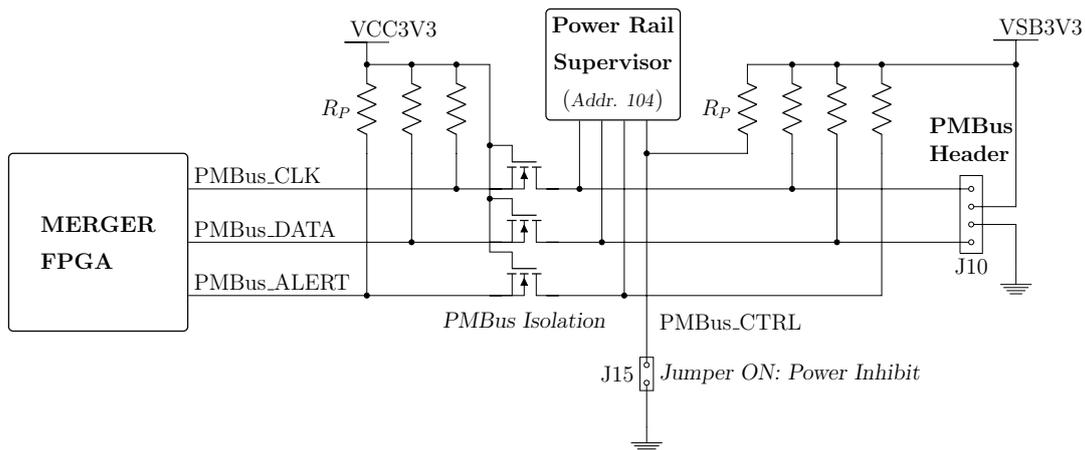


Figure 4.22: PMBus interface accessing the power rail supervisor.

4.7.3 SPI and Microwire

In selection of a communication interface for slow control tasks between the on-board FPGAs, the Serial Peripheral Interface (SPI) bus was chosen (see Fig. 4.23). Its interface is designed as a full-duplex serial link for connecting peripheral devices to

a single master. The master initiates the bus transfers on the master output (MOSI) and provides the slave peripherals with a clock signal (SCK). When multiple devices are attached to the bus in parallel, individual slave select (SS) lines are required to address the different devices. The master output signal is shared among the slaves. Reversely, the individual slave outputs can be tied together to the master input (MISO) alike, provided that tri-state drivers are available. In contrast to open-drain interfaces like I²C, push-pull drivers can be employed, allowing for higher data throughput and lower power consumption. Altogether, the SPI architecture provides a simple hardware interface and flexibility in the implementation of the software protocol applied. A comprehensive description of the inter-FPGA link is provided in Sec. 6.3.

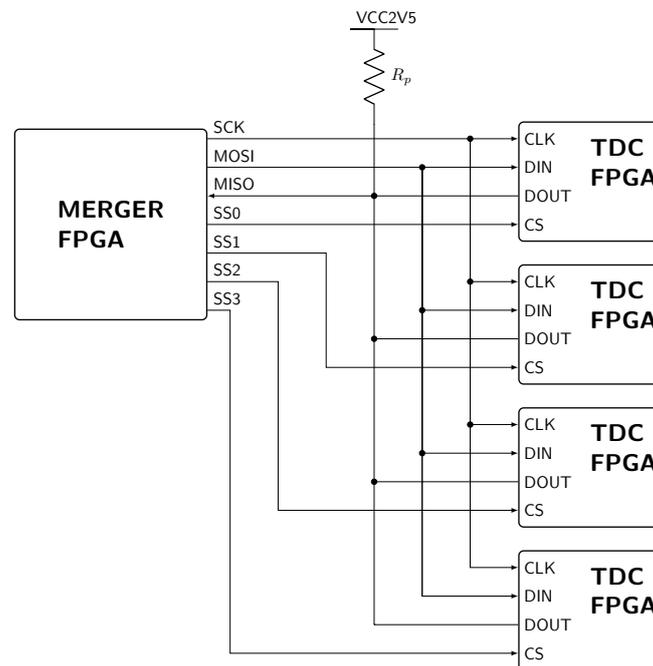


Figure 4.23: SPI bus architecture of the inter-FPGA link.

The Microwire interface [97] is a variant of the SPI bus developed by National Semiconductor. It is utilized by the jitter attenuator for register programming. The serial data is clocked into a shift register. The content of the shift register is subsequently transferred on the rising edge of the latch enable signal to the internal register addressed. The latch enable port corresponds to the slave select input of the SPI bus. Microwire peripherals might be operated in parallel with SPI devices. However, the Microwire interface is not shared with the inter-FPGA link due to its limited clock speed. It is also recommended by Texas Instruments not to share the bus wires as it can cause increased phase noise on the output clocks in case these signals are toggled while the jitter attenuator is in operation (see Ref. [86]).

4.7.4 Board LEDs and DIP Switches

Altogether six board LEDs are located on the ARAGORN front-end. Table 4.9 lists their assigned purposes. In order to address a particular ARAGORN board, two rotary-code DIP switches (U9, U10) are provided that form a 8-bit identifier

number. Thereby, the upper bits of the identifier are selected with DIP switch U10 and the lower bits with DIP switch U9. Both LEDs and DIP switches are highlighted in Fig. 4.19.

Table 4.9: Listing of the board LEDs.

RefDes	Colour	Meaning
LD1	red	MERGER-FPGA configuration failure
LD2	green	3.3 V standby power on
LD13	red	Secondary voltage rail failure
LD15	green	TDC-FPGAs configured
LD16	green	Jitter attenuator locked
LD17	green	SFP link aligned to valid word boundaries

4.7.5 JTAG Configuration and Debugging

In-system debugging tools interact with the FPGA fabrics through the JTAG interface that is compliant with the IEEE Standard 1149.1 [98]. This standard defines the boundary-scan architecture, a method for testing integrated circuits and interconnects on assembled PCBs. Targeted devices communicate through the test access port that contains the following connections:

TCK test clock input.

TMS next state select input to the test access port controller.

TDI serial data input to the instruction and data registers.

TDO serial data output of the JTAG registers.

The built-in test logic includes the test access port controller, an instruction register and several data registers accessing the test features of the design. Multiple devices can form a daisy chain path, interconnecting the TDO pin of one device with the TDI pin of the following device. The test access port of the Artix-7 FPGA is located in I/O bank 0. Thus, the on-board FPGAs cannot share the same JTAG chain because the Flash and consequently the MERGER-FPGA require a different supply voltage than the TDC-FPGAs.

Besides the standard boundary-scan registers, Xilinx FPGAs facilitate a dedicated register for bitstream download leveraging the configure instruction. In addition, the JTAG interface can be employed for indirect programming of Flash memories. The programming flow is referred to as indirect because the FPGA is preloaded with a bitstream that converts the JTAG instructions and subsequently outputs the received configuration data through the existing memory interface to the Flash (cf. Sec. 4.4). The configuration options described and the capability for debugging logic designs in hardware are integrated features of the XILINX *Vivado Design Suite* [99].

Programming cables to connect to the device under test for prototyping purposes are available from different vendors.

Debugging of logic designs is performed with special debug Intellectual Property (IP) cores¹ that can be customized to probe selected internal signals in the design. As described above, the logic analyzer software interacts with the implemented debug cores through the JTAG interface to examine the probed nets and to stimulate logic in the design. However, the sampling frequency is limited to system speeds and the length of the captured data frame depends on the integrated memory resources available in the target device. Despite these limitations, in-system debugging flows are advantageous over external logic analyzers and de facto standard in complex FPGA designs.

4.8 PCB Design

The PCB design process implies several steps, namely the design authoring, the placement of the board components and the routing of the physical interconnections between them. In high-speed digital designs, certain constraints need to be applied, for instance to match the targeted impedance of the transmission lines. Concomitantly with the layout process, intensive simulations have constantly been performed to ensure signal integrity and minimum noise coupling of the power distribution network. These challenging tasks were carried out using the Cadence *Allegro 16.6* software toolkit [100].

4.8.1 Schematics

In this project, the *Allegro System Architect* software was used for design authoring. This tool allows the user to accomplish the design entry either using the conventional schematics or a spreadsheet-based flow. With the schematic editor, the designer gathers the desired connectivity by drawing wire connections between component symbols representing the physical devices. This approach is most suitable for analog designs like power supplies. The spreadsheet editor in turn comes in handy dealing with high pin-count devices like FPGAs that involve only few discrete components like termination resistors or decoupling capacitors. The spreadsheet-based method was highly appreciated for capturing the connectivity of the FPGA devices and the extension board connectors on the ARAGORN front-end. For documentation purposes, schematic diagrams can be exported from spreadsheet-based designs as well.

4.8.2 Layout

In order to meet the targeted performance of digital systems, transmission line effects like reflections, cross-talk or ground bounce and their possible impact on the quality of the transmitted signals must be considered in the layout of the PCB. These requirements are then turned into design rules or constraints that can be captured with the *Allegro Constraint Manager* spreadsheet tool. A constraint set gathers multiple constraint parameters of a given domain.

¹Intellectual Property cores are reusable, configurable logic blocks

Table 4.10: Listing of the electrical constraint sets.

Electrical CSet	Static phase	Dynamic phase	
		max length	tolerance
DIFF_ECS	2.5 mm	15 mm	5 mm
HS_DIFF_ECS	0.15 mm	15 mm	0.15 mm

Any asymmetry in the length of a differential pair causes mode conversion, transforming the differential signal fractionally into the common signal. If the common signal again is not properly terminated, it may reconvert causing increased differential noise (see Ref. [101, p. 606 ff.]). Taking this into account, electrical constraint sets controlling the skew of the differential signal pairs in the design were created (see Tab. 4.10). The HS_DIFF_ECS constraint set applies strict phase tolerance to the high-speed differential lines interfacing the optical transceiver sockets with the MERGER-FPGA, while the noncritical nets are assigned to the DIFF_ECS constraint set. The static phase tolerance reflects the total difference in length, from the driver to the receiver, between the traces of a differential pair. The dynamic phase specifies a maximum tolerated etch length exceeding the phase tolerance along the signal path. The *Allegro PCB Designer* tool displays the affected nets and interactively assists the designer in identifying the source of the asymmetry, and hence to balance the differential lines by adding compensation length (see Fig. 4.24a).

Table 4.11: Listing of the physical constraint sets.

Physical CSet	Line width		Separation gap		Via	
	default	neck	default	neck	hole \varnothing	pad \varnothing
Single-ended	0.13 mm	0.10 mm	–	–	0.20 mm	0.50 mm
Differential	0.10 mm	0.09 mm	0.20 mm	0.14 mm	0.20 mm	0.50 mm
Power/GND	0.14 mm	0.10 mm	–	–	0.30 mm	0.60 mm

A major signal integrity issue in digital designs are signal reflections that occur whenever the instantaneous impedance of transmission lines change. Engineering a flat impedance profile, the impedance of the signal traces must be matched to the interface technology employed. The characteristic impedance depends on the line width, the separation gap in case of differential pairs, the distance to the reference plane and the value of the dielectric constant of the substrate material. Using the *Polar Si8000 Impedance Field Solver* [102], these parameters were selected such that the design constraints of $(50 \pm 5) \Omega$ for single-ended signals and $(100 \pm 5) \Omega$ for differential pairs are met on all signal layers. The physical dimensions of the signal traces and the board vias are specified in the physical constraint sets listed in Tab. 4.11. The BGA packages of the FPGAs have a pitch of 1 mm. With the default physical constraint set, it is impossible to route the differential signals in between the via rows of the BGA escapes. Thus, a so-called neck mode was introduced that squeezes down the traces in order to pass the differential signal in a tightly coupled

manner through the via mesh (see Fig. 4.24b). Spacing constraints are the third type of design constraints. For low-voltage circuit boards like the ARAGORN front-end, the minimum spacing between conductive elements is basically limited only by the manufacturing process of the PCB. However, sufficient edge-to-edge spacing must be observed to neighbouring traces to avoid cross-talk within the design.

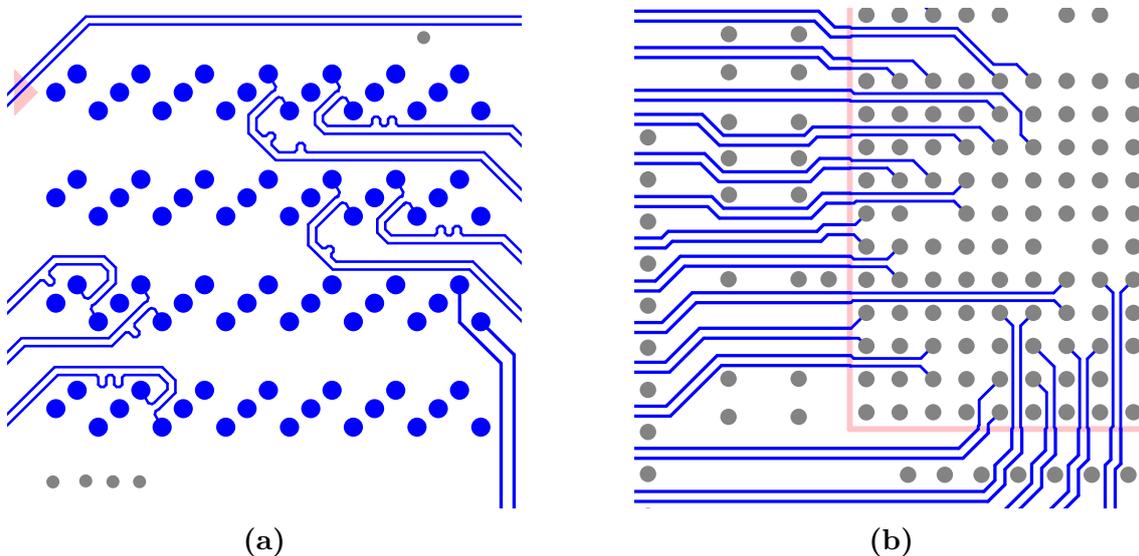


Figure 4.24: (a) The high-speed differential signals are phase-tuned by adding compensation length to the shorter trace close to the origin of the length mismatch. (b) The BGA escapes of the FPGAs are routed in neck mode passing the differential signal traces tightly-coupled through the dense via mesh.

Once completed, the logic design is exported to the *Allegro PCB Editor*. The PCB design process usually involves subsequent modifications such as pin swaps made in the physical layout and the logic design, accordingly. The first step in the layout process is to define the cross section of the PCB. The board areas facing the highest routing density determine the number of required signal layers. These regions have been identified on the ARAGORN front-end as the via fields underneath the FPGAs that require four layers to route all endpoints. The cross section of the 14-layer PCB is shown in Fig. A.3.

Special attention has been paid to the routing of the 384 differential pin pairs of the extension board connectors, because any noise picked up along these interconnect lines to the TDC-FPGAs directly affect the accuracy of the TDC application. Due to the dense layout, cross-talk induced noise must be considered that can be decreased by moving neighbouring traces farther apart, not less than twice the line width. However, enhanced near-end cross-talk is obtained between adjacent signal layers in dual-stripline configurations due to inadvertent overlap of signal traces (see Ref. [101, p. 405 ff.]). This effect could be reduced with orthogonal routing, which is often not applicable in dense designs like the ARAGORN front-end. Taking this into account, the signal layers have been separated from one another by return layers to eliminate noise contributions from broadside coupling.

Engineering the power distribution network, the basic design principle is to keep the supply voltages at the device pads within the specified noise limits under all

operating conditions. Fluctuating currents induced by switching events inside the on-board devices cause ripple voltages proportional to the impedance of the interconnects. Hence, the impedance spectrum of the power distribution network must not exceed the target impedance up to the bandwidth of the switching currents. Based on the worst case transient current I_{trans} , an estimate of the target impedance Z_{target} is given by the relationship (see Ref. [101, p. 624])

$$Z_{target} < \frac{VCC \cdot \text{ripple}(\%)}{I_{trans}}, \quad (4.7)$$

where VCC denotes the power supply voltage. To keep the impedance profile below the target impedance and thus the supply voltage within the noise limits, the peak impedance is damped out by a network of decoupling capacitors placed in close proximity to the current drawing devices, here the on-board FPGAs. However, the impedance of a real capacitor rises up for higher frequencies due to its parasitic inductance. To get around this issue, multiple capacitors are connected in parallel lowering their inductive contributions. The capacitor and FPGA mounting inductance together with the cavity spreading inductance, associated with the PCB power and ground planes, likewise contribute to the total effective inductance as the current flows from the capacitor to the power supply pins of the FPGA and back again from the ground pins to the capacitor. In order to reduce the current loop area, the design of the cross section minimizes the spacing between the power and adjacent ground planes. Apart from the capacitor quantities and values, their placement on the PCB is equally important. In order to be effective, the distance of a capacitor to the power supply pins must be matched to the transient frequency. Wherever possible, the high-frequency capacitors were directly attached to the BGA vias using a via-in-pad process. In practice, determining the target impedance is not an obvious task because it is very difficult to acquire precise knowledge about the current spectrum in the design that moreover varies with CLB and I/O utilization. Therefore, this design employs a decoupling network¹ recommended by the FPGA vendor, which surely exceeds the requirements for the applications performed with the ARAGORN front-end.

4.8.3 Simulations

Power and signal integrity analyses were performed using the *Allegro Sigrity* simulation environment. The DC currents delivered from the voltage regulator modules through the power distribution network to the device pads induce static voltage drops at the sinks. These drops are proportional to the series resistance observed in the PCB interconnects. Designing the power distribution network, the first step is to keep the DC voltage drop to a minimum as any deviation from the nominal voltage at the loads restrains the tolerable noise limits. The principle solution is to use wider and thicker copper traces. However, to fit the different voltage rails on the limited design space, certain trade-offs must be made in the layout of the PCB between what is desirable and what is feasible. Following this exercise, the simulation tools can be consulted to optimize the placement of the decoupling capacitors.

¹Decoupling capacitor quantities per FPGA device can be found in Ref. [103]

The core voltage supply (cf. Sec. 4.3) of the FPGA fabrics, which is distributed on two independent voltage rails, has the highest power demand, the lowest operating voltage (1.0 V) and the most stringent tolerance (5%). Identifying excessive voltage drop and thermal hotspots in the layout of these planes, the design was examined in consideration of the specified sink currents. The simulation is furthermore very helpful to evaluate the optimal locations for the sense lines of the voltage regulators that partially compensate the resistive loss to keep the voltage as close as possible to the nominal value. In addition, visualizations of the current flows through the vias connecting the voltage regulator outputs to the internal board layers were used to explore their ideal quantities and positioning at the output pads of the voltage regulators. Figure 4.25 shows the voltage drop across the core supply planes that could be optimized to deviate not more than 2.6% from the specification. The temperature distribution in Fig. 4.26 shows the simulated joule heating, more precisely the temperature increase due to the resistive power loss in the conductors. Although the effective board temperature is dominated by component heating, adequate and uniform heat transfer is important in order to avoid thermal stress to the board.



Figure 4.25: Simulation result for the voltage drop across the core supply planes. The deviation from the nominal voltage is below 2.6%.



Figure 4.26: Simulation result for the temperature rise of the core supply planes due to joule heating.

5. Time-to-Digital Converter

High-precision time measurements are frequently required for many applications in the field of high-energy physics, allowing for accurate time-of-flight or drift-time measurements. A Time-to-Digital Converter (TDC) translates the timing of signal changes into digital time values. Depending on the granularity of the converter circuit, a loss of information is associated with the digitization process. The TDC circuits commonly receive the digital input signals from preamplifier following discriminator modules that process the detector response. The timing of the physical event in the detector is determined by the rising or respectively falling-edge of the incoming pulse, referred to as a hit. While some architectures provide distinct inputs to start and stop the measurement process, in this project the incoming hits are processed on the fly using a reference clock as common time base. With this method, the dead time is virtually zero. Time intervals are later calculated off-chip from the recorded timestamps. In the COMPASS-II experiment, the reference clock is distributed globally to all readout modules employed in the spectrometer, allowing for time measurements between channels from different host boards. The performance characteristics of TDCs are defined as following:

- **time bin size** — intrinsic digitization step or Least Significant Bit (LSB).
- **time resolution** — random error or standard deviation of a time interval measurement.
- **dynamic range** — number of bits of the digital value representing the measurement result in units of LSB. If the measured time interval exceeds the measurement range, the result is equivalent to the time interval modulo the dynamic range.
- **double hit resolution** — minimum delay between consecutive hits received on the same input channel. Pulses with shorter spacing cannot be resolved by the TDC.
- **hit rate** — in our case, the input pulses to the TDC occur randomly in time. Not to be confused with the double hit resolution, the maximum hit

rate specifies the average input rate that can be processed on a certain time interval. Exceeding this parameter results in buffer overflow conditions and consequently loss of data.

- **trigger rate** — on-chip pre-processing of the raw data frames is essential for many different applications. In this design, only such hits which are time-correlated to the trigger primitives are selected for data output (see Sec. 6.1.2). The maximum trigger rate depends on various factors like the hit rate, the length of the trigger gate, the internal buffer depth and the bandwidth of the data acquisition system.
- **dead time** — period of time in which hit encoding tasks prevent the TDC from processing new hits. The dead time can be reduced by the use of pipelining registers and integrated memory for data buffering.

In real-world applications, precise time measurements are affected by jitter contributions from the input signals, the sampling clock and inherent noise in the system. The inhomogeneity of the conversion characteristic, which is described with the following metrics, likewise contributes to the accuracy of the measurements:

- **Differential Non-Linearity (DNL)** — the deviation from the nominal time bin size can be determined with a statistical code-density test. The DNL is usually illustrated with a graph showing the deviations of the individual bins normalized to LSB. Alternatively, the maximum value can be specified.
- **Integral Non-Linearity (INL)** — the deviation $t_{\text{INL},i}$ from the ideal transfer function after a given bin number i :

$$t_{\text{INL},i} = t_i - i \cdot \text{LSB}. \quad (5.1)$$

The graph of the INL can be obtained from summation of the differential non-linearities. The standard deviation of the INL is a good estimate for the deviation from the precision of an ideal TDC.

- **gain error** — the gain k denotes the slope of the regression line of the transfer function:

$$k = \frac{1}{\text{LSB}}. \quad (5.2)$$

The TDC application developed in this work implements an interpolating method. This technique combines a precision measurement of the timing of the hit within a reference clock cycle with the value of a coarse counter extending the measurement range. Due to the periodicity of the transfer characteristic it makes more sense to regard the INL rather than the gain error.

5.1 Random Error

Evaluating the precision of time interval measurements, it is assumed that the incoming hits are uncorrelated in time with the sampling clock of the TDC. The measured interval T can be decomposed in an integral part Q and a fractional part

F with $(0 \leq F \leq 1)$, so that $T = \text{LSB} \cdot (Q + F)$. Taking a repeated measurement of a constant interval T , the output of an ideal TDC is either $T_1 = Q$ or $T_2 = Q + 1$, provided that $F \neq 1$. The ratio of T_1 and T_2 follows a binomial distribution with

$$p(T_1) = 1 - F \quad \text{and} \quad q(T_2) = F. \quad (5.3)$$

The standard deviation σ of the binomial distribution is considered as an estimate of the random error or single-shot precision (see Ref. [104]):

$$\sigma = \text{LSB} \sqrt{F(1 - F)}. \quad (5.4)$$

The random error strongly depends on the fractional part F of the measured interval, reaching a maximum of $\sigma_{max} = \text{LSB}/2$ for $F = 0.5$ when both measurement results are evenly observed. In the event of $F = 0$ or $F = 1$, the TDC exclusively outputs either T_1 or T_2 and the random error becomes practically zero. The half-circle graph in Fig. 5.1 depicts the normalized standard deviation σ/LSB . The average standard deviation σ_{avg} can be derived by integration of Eq. (5.4) within the limits $0 \leq F \leq 1$ (see Ref. [104]):

$$\sigma_{avg} = \frac{\pi}{8} \text{LSB} \cong 0.39 \text{LSB}. \quad (5.5)$$

The INL of the converter circuit and the jitter contributions from different constituents prevent that the theoretical resolution from Eq. (5.4) is observed in real world applications. Determining the single-shot precision, the length of the measured interval is incremented in small steps in the range of T to $T + \text{LSB}$ to cover the periodic pattern of the converter characteristic. The worst case random error or the average single-shot precision obtained in this exercise is an important measure of the effective resolution.

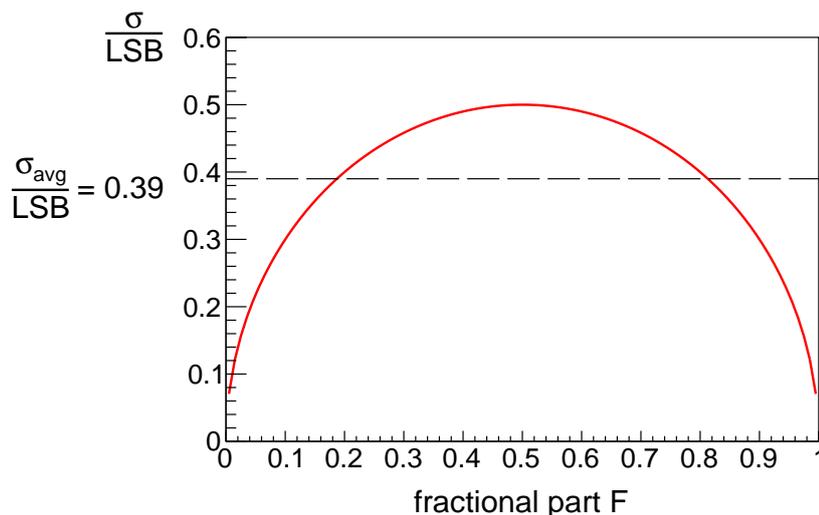


Figure 5.1: Normalized standard deviation σ/LSB dependent on the fractional part F of the measured time interval.

In some applications, where it is possible to take a series of N measurements of a constant interval T , the accuracy can be improved beyond the precision of a single-shot measurement by averaging the results (see Ref. [105]). An important finding

is that, if N is sufficiently large, the best estimate of the measured interval is the averaged result \hat{T} of the measurements, so that $\hat{T} \approx T$. It is demonstrated that the random error of the averaged reading is reduced by a factor of $1/\sqrt{N}$. The worst case and average uncertainty are then given by:

$$\sigma_{max,\hat{T}} = \frac{\text{LSB}}{2\sqrt{N}} \quad (5.6)$$

and

$$\sigma_{avg,\hat{T}} \cong 0.39 \frac{\text{LSB}}{\sqrt{N}}. \quad (5.7)$$

5.2 Counter-based TDC

The simplest implementation of a TDC is that of a coarse counter. The time base equals the clock period of the reference clock, thus $\text{LSB} = 1/f_{ref}$. Time interval measurements are performed by activating a preferably synchronized gate input to the counter. Alternatively, the current counter value is latched at the beginning and once again at the end of the interval. The measurement result calculated from the two counter values must take into account possible counter rollovers. This mode of operation is favoured if the incoming hits are received on different channels. The coarse counter approach has the advantage that the measurement range scales with the width of the counter primitives. However, an obvious drawback of this method is the limited precision because high-frequency oscillators are expensive devices and the large supply currents cause enhanced power consumption and often require efficient cooling. In the event of the Artix-7 FPGA, assuming the counter primitives are implemented with the integrated DSP components (cf. Sec. 4.2.1.3) providing a maximum frequency of 550 MHz and consequently a digitization step of 1.8 ns, this counter-based approach would be insufficient for most applications in modern high-energy physics experiments.

5.3 Interpolating TDC

Higher accuracy is achieved with an interpolating method (see Ref. [106, p. 13 ff.]). The interpolators resolve the timing of the hit signal with respect to the clock edge. The subdivision of the reference clock cycle can be achieved with a tapped delay line. Given that the individual delay cells introduce equidistant delay steps τ , the number of delay elements K are selected such that the total propagation delay $K\tau$ equals the clock period T_{ref} . Compared to the coarse counter approach, interpolating TDCs provide an improved time bin size of $\text{LSB} = T_{ref}/K$. When the hit signal propagates through the delay line, the state of the delay line is captured by sampling the delayed versions of the hit signal using latches or flip-flops. The measurement result is encoded from the thermometer code output of the associated register. The flip-flops reading the delay cells, which have been passed by the hit signal, output the 'high' state while the residuals in the delay line still remain in the 'low' state (see Fig. 5.2). Hence, the timing of the hit is given by the initial bit-flip position in the register output. In a subsequent encoding step, the measurement result is converted to a binary code word.

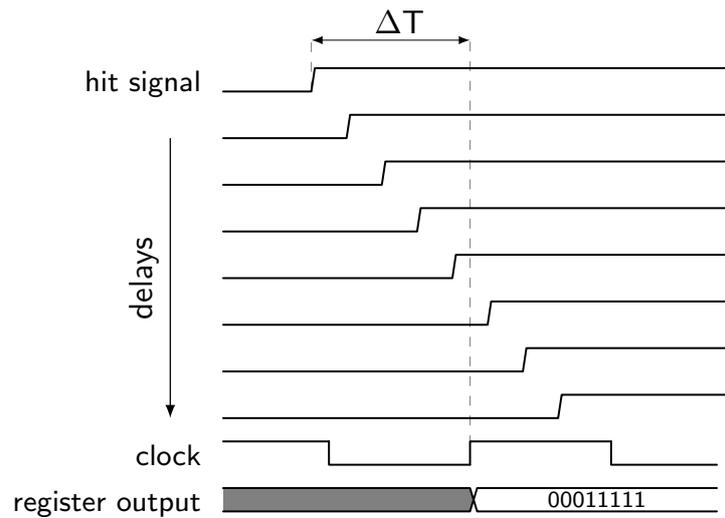


Figure 5.2: Operational principle of the interpolating TDC. The delayed versions of the incoming hit signal are sampled in parallel to store the state of the delay line in a register. The timing of the hit signal is encoded from the resulting thermometer code. Adapted from [106, p. 14].

An improvement of the interpolating method is achieved by the use of an integrated Delay-Locked Loop (DLL) or Phase-Locked Loop (PLL) circuit that automatically stabilizes the delay line against ambient temperature and supply voltage variations. The DLL facilitates a phase detector that compares the phase of the reference clock before and after it has passed the voltage-controlled delay line. A charge pump in combination with a loop filter following the phase detector tunes the delay line to match exactly the reference clock cycle. The PLL basically consists of a voltage-controlled oscillator that is synchronized to the reference clock. Depending on the bandwidth of the loop filter and the amount of phase noise inherent with the voltage-controlled oscillator, the PLL removes jitter from the reference clock, thereby improving the accuracy of the TDC. If the voltage-controlled oscillator is designed as a ring oscillator, this circuitry provides a delay line similar to the DLL approach.

Either way, the equidistant phase-shifted clock signals connect to the data inputs of the flip-flops and the incoming hit signal is used to sample the associated register. Alternatively, the interconnections are reversed so that the delay line acts as a multiphase clock source, activating the flip-flops in sequential order. In consideration of the available clocking resources provided by the Artix-7 FPGA (cf. Sec. 4.2.1.4), the latter was favoured for the implementation of the TDC application. The relatively short measurement range provided by the interpolators is easily extended with a synchronous counter counting up the reference clock cycles. The measurement result combines both the binary representation of the bit-flip position in the register n_i and the current coarse counter value m_i . The measured time interval T calculated from consecutive samples is (see Ref. [104])

$$T = (m_2 - m_1) \cdot T_{ref} + (n_2 - n_1) \cdot \text{LSB}. \quad (5.8)$$

6. Firmware

This chapter outlines the development of the firmware designs for the on-board FPGAs. The structure and the behaviour of digital circuits can be described in a text-based manner using a Hardware Description Language (HDL). This work makes use of VHDL¹, which was originally developed for behavioral simulations. However, a subset of the language constructs can be used to create Register Transfer Level (RTL) models of the design. The RTL description is then translated with a synthesis tool to a gate level netlist inferring the logic gates required to implement the desired circuit.

The Vivado Design Suite provides the toolset for the implementation of logic designs tailored to the Xilinx FPGA architecture. The implementation flow includes the synthesis of the RTL code and the mapping of the circuit to the logic resources provided by the FPGA fabric. Subsequent design steps conduct the placement of the logic primitives onto the hardware platform and the routing of the interconnections between them. Finally, a bitstream file for device programming is generated. The implementation process is carried out either in a graphical user interface for the Vivado Design Suite or with Tool command language (Tcl) commands in a shell environment. Likewise, the Tcl commands involved in the entire implementation flow can be embedded in a script file.

In order to obtain implementation results that meet timing closure, various design directives, so-called constraints must be applied during physical synthesis. The design constraints used with the Vivado Integrated Design Environment (IDE) are specified in the XILINX Design Constraint (XDC) format, a selection of industry standard and XILINX proprietary constraints that are collected in one or more (.xdc) files (see Ref. [107]). Two major categories can be distinguished, timing and physical constraints. Timing constraints refer to the specified input and output delays and the timing path requirements associated with the targeted clock frequencies. Physical constraints define pin locations and guide the place and route tools to use certain cell locations and routing resources. For instance, fixed-routing constraints play an

¹Very High Speed Integrated Circuit Hardware Description Language

important role in the implementation of the TDC-FPGA design in order to obtain uniform routing delays of the incoming hit signals to the sampling registers. Finally, the design tools support the use of RTL attributes, which are directly specified in the RTL code, to control the treatment of certain nets or registers.

6.1 TDC-FPGA Design

The TDC firmware implements a 96-channel TDC inside a single Artix-7 FPGA. A top-level diagram of the design is shown in Fig. 6.1. The main building blocks are detailed in the section at hand. The TDC architecture is based on the interpolating method outlined in Sec. 5.3. Intrinsic delay cells with similar propagation delays that can be linked together to form a tapped delay line are not eligible constituents of FPGAs. Commonly, delay lines are implemented based on the gate delays inherent with the carry chain structure propagating upward in the CLB slice columns (cf. Sec. 4.2.1.1). A drawback of this approach is that the total delay of the line must be adjusted to match the sampling clock period and that the cell delays are subjected to ambient temperature variations so that additional calibration steps are necessary. These design obstacles can be avoided when the interpolators are operated with a multiphase clock generated from an external reference clock input. This concept is beneficial in the sense that the design can make use of the integrated clocking resources to control the phase alignment.

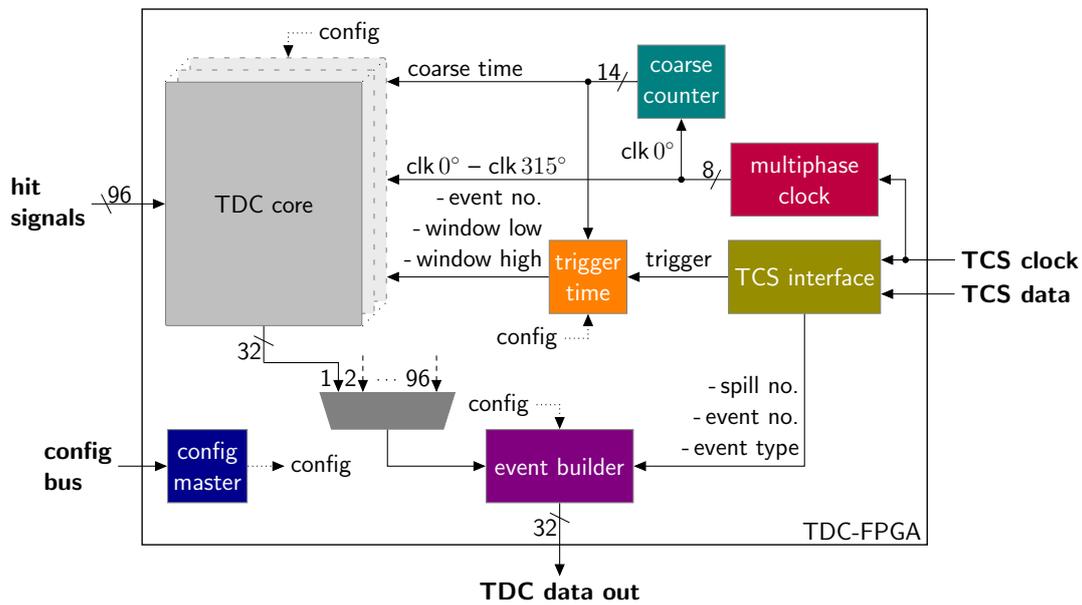


Figure 6.1: Top-level diagram of the TDC-FPGA design depicting the main functional blocks and interconnections.

6.1.1 Multiphase Clock

The TCS reference clock is recovered inside the MERGER-FPGA and distributed in parallel with the TCS data stream to the on-board TDC-FPGAs where it is fed into two MMCM primitives (see Fig. 6.2). The synthesized clock frequency ($f_{TDC} = 311.04 \text{ MHz}$) equals twice the reference clock frequency. The MMCM

basically consists of a PLL with some enhanced functionalities (cf. Sec. 4.2.1.4). The clock outputs of the MMCM can select from eight different phases of the voltage-controlled oscillator, shifted by 45° each. Depending on the selected output counter values, even higher phase-shift resolution is achieved.

The first MMCM is programmed to produce four phase-shifted clocks of 0° , 45° , 90° and 135° with respect to the reference clock input. Four more clocks shifted by 180° , 225° , 270° and 315° are delivered by the second MMCM. As a result, the dual MMCM configuration reveals a stabilized multiphase clock with eight evenly aligned phases. The phase-shifted clocks access the global clock tree of the FPGA that distributes the clock signals with minimal skew to any sequential element in the device. This ensures that the uniform delay steps introduced by the multiphase clock are retained at the clock inputs to the sampling registers.

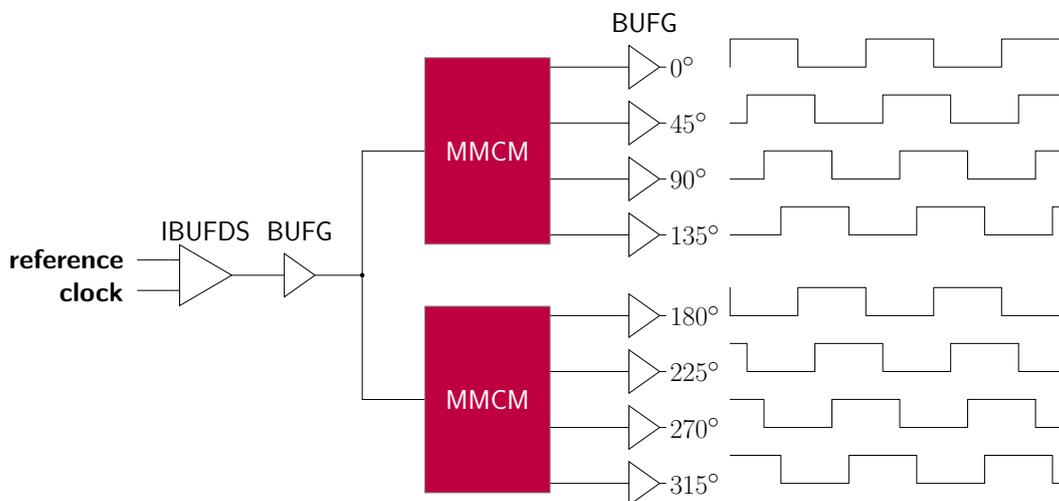


Figure 6.2: Two MMCMs synthesize eight equidistant phase-shifted clocks from the external reference clock that enters the FPGA through a differential input buffer (IBUFDS) followed by a global clock buffer (BUFG) primitive.

6.1.2 TDC Core

The TDC core instance covers the logic design of a single channel following the interpolation principle. Figure 6.6 shows a schematic diagram of the implemented circuit.

Fine Interpolators

Each input signal is connected to a set of eight edge-triggered D flip-flops. The flip-flops are driven in successive order by the multiphase clock (see Fig. 6.3). This approach results in a digitization step of $\text{LSB} = 1/(8 \cdot f_{TDC}) \cong 402 \text{ ps}$. Just as for the clock buffer tree, any routing skew in the input signal to the flip-flops has a direct impact on the uniformity of the transfer characteristic. The automatic router algorithms optimize the data paths between registers until the timing requirements are fulfilled. However, the results obtained with the standard implementation flow are inadequate for this application. It was therefore necessary to work out adequate implementation directives balancing the propagation delays to all net endpoints (see Sec. 6.1.4).

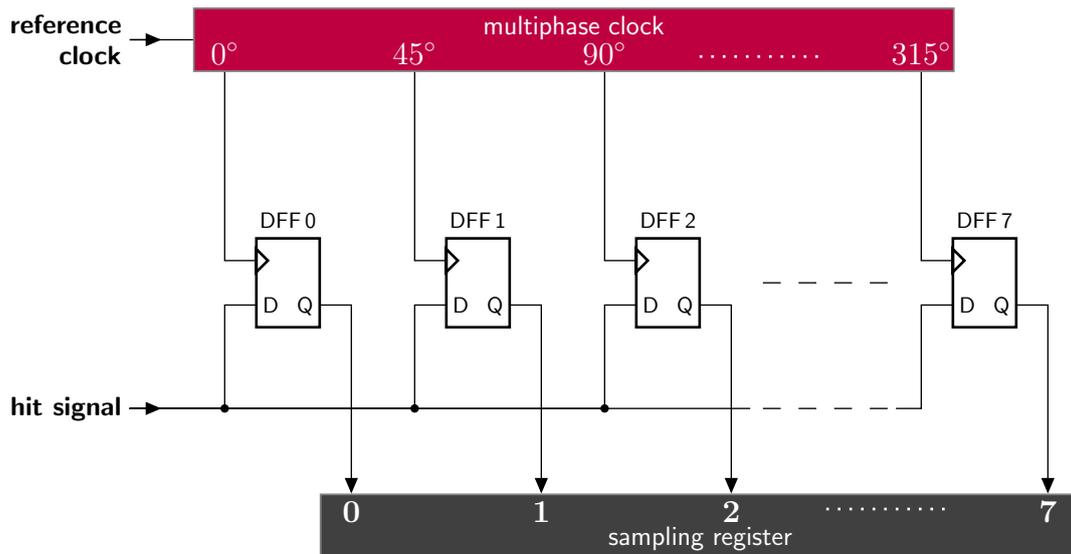


Figure 6.3: The multiphase clock activates the flip-flops (DFF0 – DFF7) receiving the incoming hit signal in consecutive order. The measurement result is encoded from the readings of the sampling register.

Before the fine time can be resolved from the readings of the sampling register, the output is synchronized with the first sampling clock. This is in fact the most timing critical part of the design, because the worst case margin equals the time delay between the phases of the multiphase clock. Relaxing the timing requirements of the design, the synchronization is performed by a three-stage pipelining register (see Fig. 6.4). A beneficial side effect is that the synchronizer protects the subsequent logic from unlikely but nevertheless potential effects of metastability that could be observed when the asynchronous hit signal violates the setup and hold time requirements of the flip-flops in the interpolators.

Hits are detected in the synchronized output of the sampling register using an 8-input OR gate. An edge detection circuit following the OR gate flags if a bit-flip occurred in the thermometer code. The OR gate remains in the high state for multiple clock cycles depending on the width of the input pulse. With this circuit, consecutive hits are only recognized if the OR gate returns to the low state for at least one clock cycle, determining the double hit resolution of the TDC. Analogously, the falling-edge of the input pulse is detected with an 8-input NAND gate which is beneficial for time-over-threshold measurements or when the polarity of the input signals is reversed. Covering those aspects, the detection logic was designed for user-selectable rising, falling or both-edge sensitivity.

The initial bit-flip position in the thermometer code yields the fine time of the measurement result. The low-to-high transition corresponds to the rising-edge of the hit while the falling-edge is identified with the lowest bit position storing the low state. The truth tables of the implemented binary encoders are listed in Tabs. 6.1 and 6.2. The binary value T_{hit} of the final measurement result¹ is composed of the

¹the binary output of the TDC is the measured timestamp in units of LSB

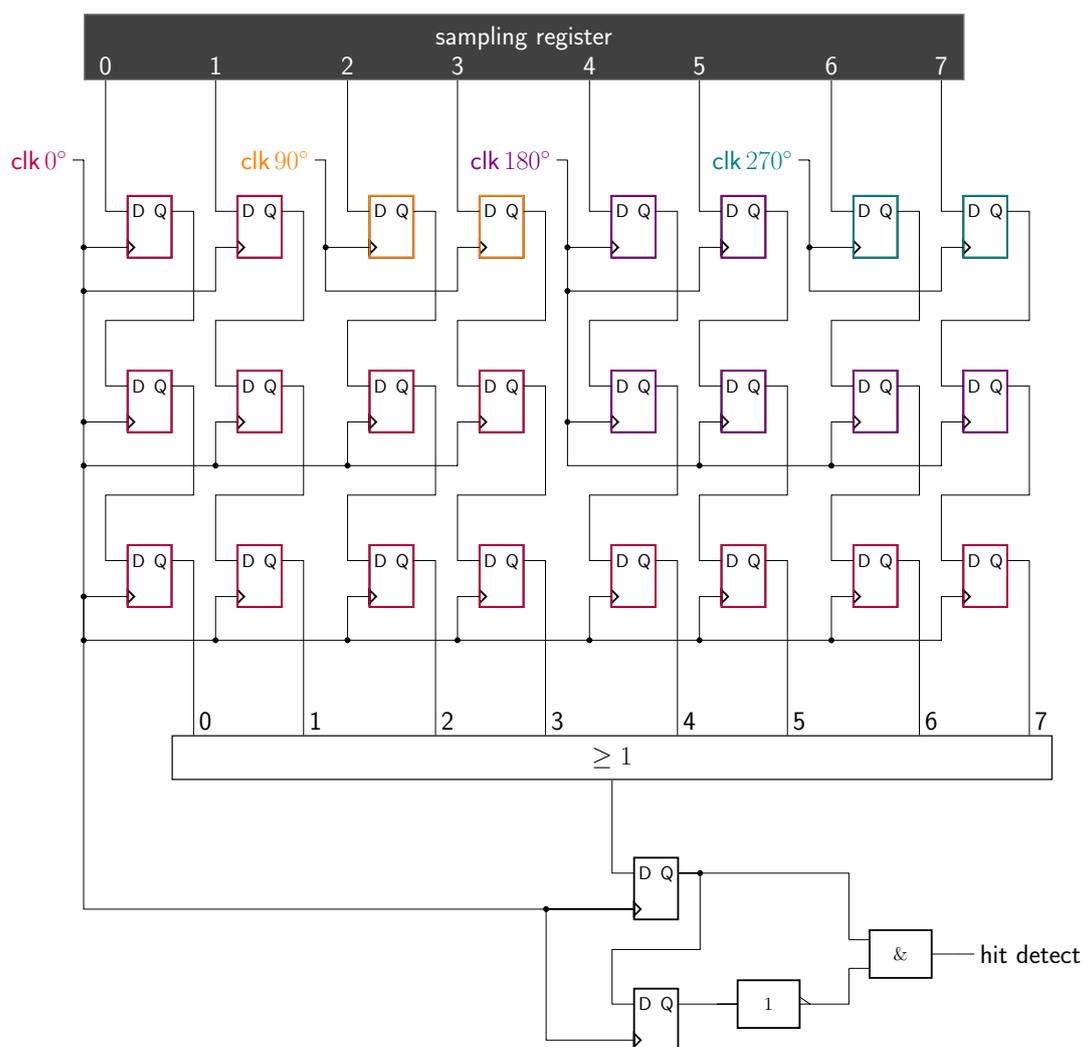


Figure 6.4: The output of the hit register is successively synchronized to the clock domain of $\text{clk } 0^\circ$ using a three-stage pipelining register. A bit-flip in the synchronized register is detected with an 8-input OR gate. The subsequent edge detection circuit ensures that only the initial low-to-high transition is recognized.

binary-encoded fine time T_f and the current value T_c of a 14-bit coarse counter that keeps track of the clock cycles elapsed to extend the measurement range:

$$T_{hit} = 8 \cdot T_c + T_f. \quad (6.1)$$

Multiplying binary numbers by a power of two is equivalent to a repeated shift left operation. Hence, the fine time bits are simply concatenated with the output of the coarse counter to obtain the measurement result according to Eq. (6.1). The computed measurement values are stored in 2k deep dual-port hit buffers implemented with the embedded 36 kbit Block RAM cells (cf. Sec. 4.2.1.2). The most significant bit of the timestamps flags potential counter rollovers and is omitted in a later processing step. Whether the detected hits are actually written to the hit buffer primitives finally depends on the selected edge-sensitivity and additional control flags avoiding for instance buffer overflow conditions. It is furthermore possible to completely shut down particular channels in the event of unconnected thus floating inputs.

Table 6.1: Truth table of the rising-edge encoder. The bits following the initial low-to-high transition are treated as don't-care ('-').

Thermometer code								Binary code		
t_7	t_6	t_5	t_4	t_3	t_2	t_1	t_0	b_2	b_1	b_0
-	-	-	-	-	-	-	1	0	0	0
-	-	-	-	-	-	1	0	0	0	1
-	-	-	-	-	1	0	0	0	1	0
-	-	-	-	1	0	0	0	0	1	1
-	-	-	1	0	0	0	0	1	0	0
-	-	1	0	0	0	0	0	1	0	1
-	1	0	0	0	0	0	0	1	1	0
1	0	0	0	0	0	0	0	1	1	1

Trigger Matching

The trigger matching unit selects only such entries from the hit buffers for subsequent data readout which coincide in time with the trigger events. At the time of trigger arrival, the current value of the coarse counter is latched and a programmable latency time is subtracted from the trigger timestamp to account for the trigger generation and distribution delay of the TCS. The latency-corrected trigger time (`window_low`) defines the lower limit of the acceptance window. Adding up a configurable gate parameter, corresponding to the drift-time or time-of-flight in the detector, determines the upper limit (`window_high`). A locally generated identifier number is incremented for every received trigger to be later compared in the event building process with the corresponding event number provided by the TCS. The trigger information computed is buffered in 1k deep FIFO primitives using the format outlined in Tab. 6.3 before the trigger matching units are available to process

Table 6.2: Truth table of the falling-edge encoder. The bits following the initial high-to-low transition are treated as don't-care ('-').

Thermometer Code								Binary Code		
t_7	t_6	t_5	t_4	t_3	t_2	t_1	t_0	b_2	b_1	b_0
-	-	-	-	-	-	-	0	0	0	0
-	-	-	-	-	-	0	1	0	0	1
-	-	-	-	-	0	1	1	0	1	0
-	-	-	-	0	1	1	1	0	1	1
-	-	-	0	1	1	1	1	1	0	0
-	-	0	1	1	1	1	1	1	0	1
-	0	1	1	1	1	1	1	1	1	0
0	1	1	1	1	1	1	1	1	1	1

the events in parallel. The simple dual-port hit buffer primitives have one write-only and one read-only port accessing the storage area. New hits are written to the memory address specified by the write pointer register (`write_ptr`). Accessing a given memory entry, the trigger matching process facilitates a read pointer register (`read_ptr`). Another register, the start search pointer (`read_ptr_s`), stores the memory address the pre-selection process is expected to start off for the next trigger event scanning the hit buffer entries. The selected timestamps are transferred to a 1k deep output FIFO. A hit is considered to match the selective time window when the coarse time fraction of the timestamp (`ram_dout(15 downto 3)`) is greater than the lower and smaller than the upper acceptance limit. This condition is expressed with the following VHDL statements:

```
gt_wl_i <= '1' when (ram_dout(15 downto 3) >=
                    window_low(12 downto 0)) else '0';
st_wh_i <= '1' when (ram_dout(15 downto 3) <=
                    window_high(12 downto 0)) else '0';
```

However, overflows in the buffer entries that occur when the coarse counter wraps around within the bounds of the acceptance limits invert the comparisons listed above. The same applies to the comparison of the address pointers indicating buffer overflow conditions. Such events are detected with a XOR gate that outputs 'high' if the most significant bit of the corresponding binary words differ. The relevant inputs to the trigger matching process therefore include reserved rollover bits. The altered statements regarding overflow conditions are then as follows:

```
gt_wl <= gt_wl_i xor window_low'HIGH xor ram_dout'HIGH;
st_wh <= st_wh_i xor window_high'HIGH xor ram_dout'HIGH;
```

The trigger matching process is implemented with a Finite-State Machine (FSM). Figure 6.5 shows a flow chart to illustrate the sequence of states. After initialization, the FSM remains in the IDLE state until a trigger is received and the empty flag of the trigger FIFO is deasserted (`tf_empty = '0'`). Then the next entry is loaded from

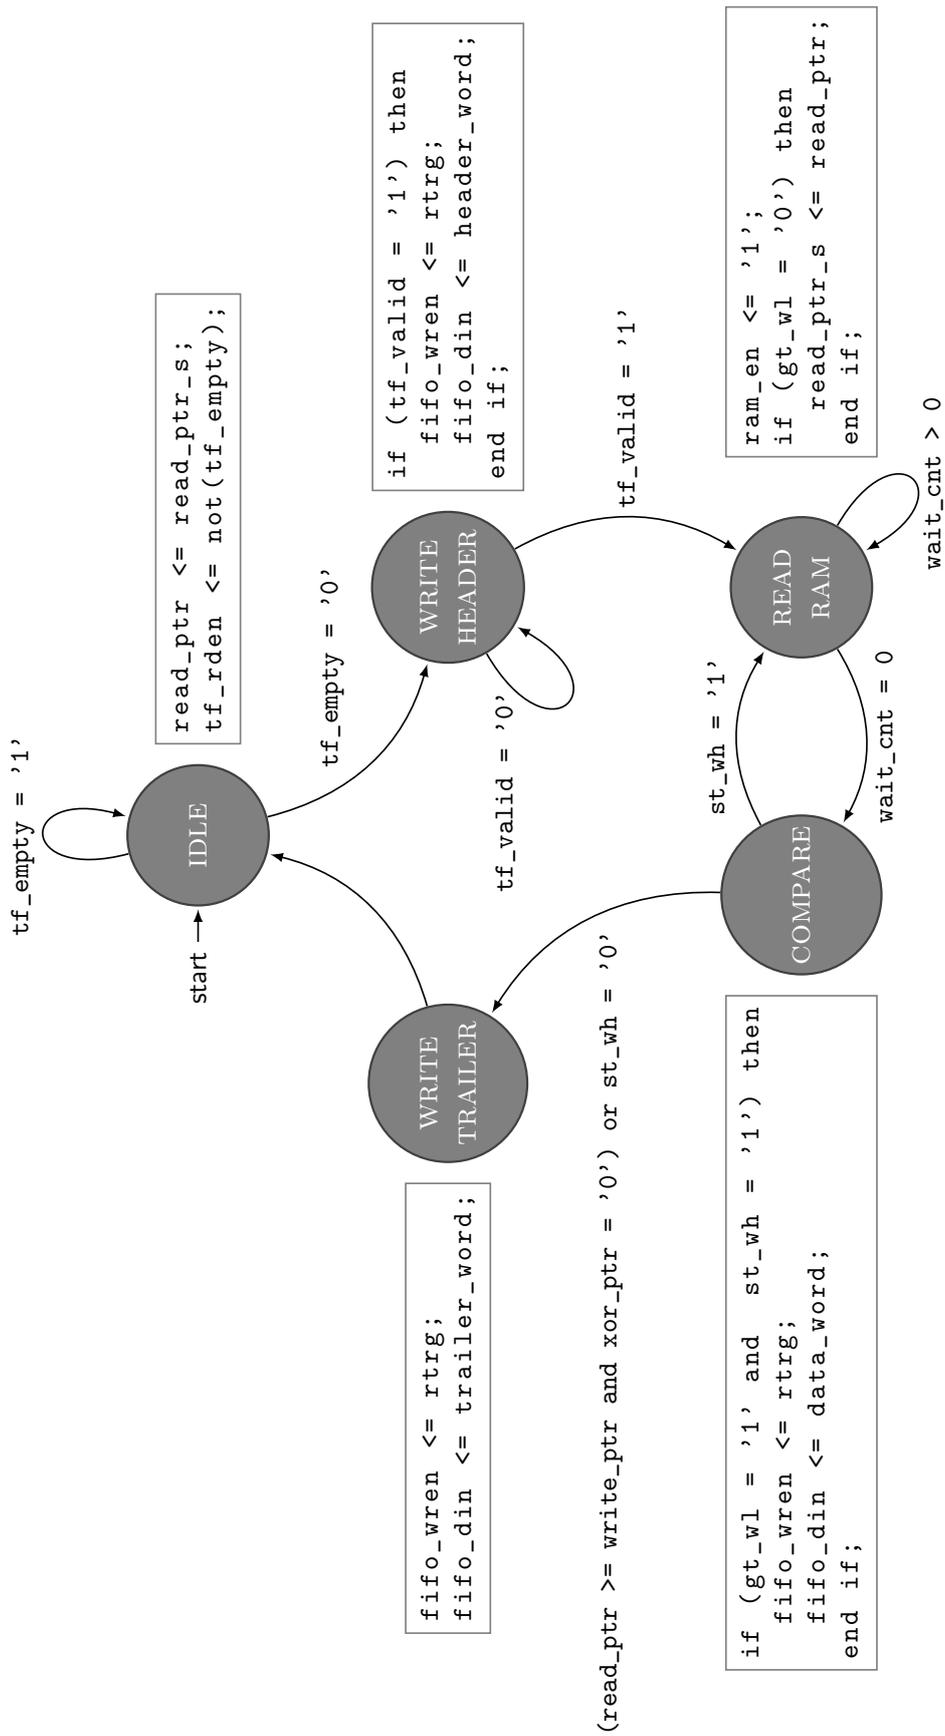


Figure 6.5: Flow chart of the trigger matching FSM.

the trigger FIFO (`tf_rden <= '1'`) and the FSM transitions to the WRITE HEADER state. As soon as the trigger information is available (`tf_valid = '1'`), a header word is written to the output FIFO to flag the beginning of the event. In the following, the FSM transitions between the READ RAM and the COMPARE state processing the hit buffer entries. The READ RAM state enables the read port of the hit buffer (`ram_en <= '1'`) and holds the process in an idle loop until the requested entry is available. The COMPARE state increments the read pointer and takes the decision whether the hits are discarded or written to the output FIFO. The search process terminates if a hit younger than the region of interest is received or no more hits are available in the hit buffer. The start search register is updated with the address of the first hit that was found in the acceptance window as the region of interest of consecutive triggers arriving at short intervals may overlap in the event of detectors using long trigger gates. Before the next trigger event is processed, a trailer word identical with the header word is written to the output FIFO to indicate the end of the event.

If the trigger matching FSM is idle for longer periods, the write pointer may catch up with the read pointer and consequently new hits would be lost. In order to prevent such overflow conditions in the hit buffers, so-called artificial triggers are generated internally at regular intervals whenever the trigger FIFO is empty and the previous event has been processed. Artificial triggers are distinguished from physics triggers by the most significant bit (`rtrg`) of the entries in the trigger FIFO and are treated equally in the trigger matching process except that no data is written to the output FIFO. This practice updates the read pointer position in memory, thereby discarding old hits that are no longer of interest for upcoming trigger events.

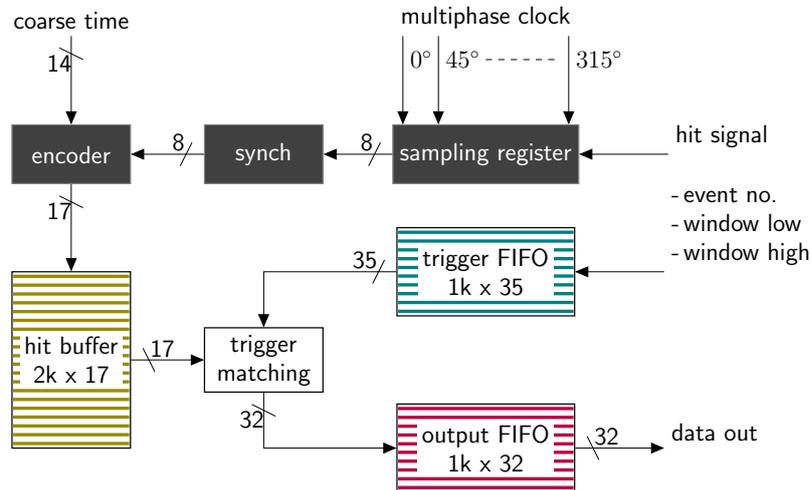


Figure 6.6: Block diagram of a TDC channel.

The format used for the header and data words (see Tab. 6.4) complies with the COMPASS online data format (see Ref. [65]). Headers are distinguished from data words by the most significant bit of the entries in the output FIFOs. The TDC timestamps (data) are included in the data words. The header words contain the internal trigger number (`event_no`) and the lowest bits of the lower acceptance limit

(trg_time) provided by the trigger FIFO entries in consistency with the existing F1-TDC employed at the COMPASS-II experiment (see Ref. [108]). The four lowest bits (lock) of both header and data words indicate if the MMCMs generating the multiphase clock were locked to the external reference clock. Each TDC input of the ARAGORN front-end is assigned to the corresponding channel of the detector assembly according to the port and channel identifiers for later analyses of the recorded data frames. Thereby, the port identifier allocates a number of 32 TDC inputs that are individually addressed by the given channel number. The mapping between the extension board connectors and the channels in the logic design can be specified after FPGA initialization as desired. The default configuration is listed in Appendix C.

Table 6.3: Format definition of the trigger FIFO entry.

34 0

rtrg	window_low (14)	window_high (14)	event_no (6)
------	-----------------	------------------	--------------

Table 6.4: Format definition of the header and data words.

31 0

header word

0	0	event_no (6)	trg_time (9)	0	channel (6)	port (4)	lock (4)
---	---	--------------	--------------	---	-------------	----------	----------

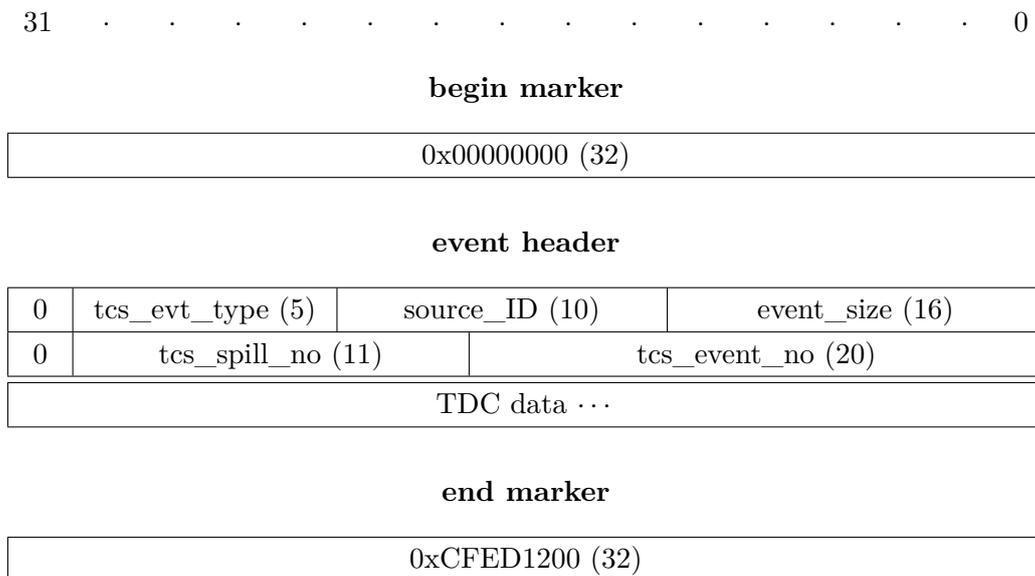
data word

1	0	channel (6)	data (16)	port (4)	lock (4)
---	---	-------------	-----------	----------	----------

6.1.3 Event Builder

Prior to data output, the data sets belonging to the same trigger event have to be collected from the individual output FIFOs. This is done using a three-stage scheme. In the first stage, the output FIFOs are processed in parallel by a number of twelve 8:1 data concentrator units. The second and final stage consist of three 4:1 data concentrator and a single 3:1 merger unit, respectively. After each stage, the collected data is buffered in 1k deep FIFO primitives. The merger algorithm implies that the event data is delimited by header respectively trailer words. To reduce the data volume transferred, these control words are successively dropped except from the header of the first channel and the trailer of the last channel with the same port identifier.

The data packets are finally combined with the event labels distributed by the TCS. The data format is outlined in Tab. 6.5. Every data packet is prepended with two event header and encapsulated within control words, the begin and end marker. An additional bit (cw) indicating control words in the data stream is appended to every buffer entry. For further reading please refer to the COMPASS online data format (see Ref. [65]).

Table 6.5: Format of a data packet.

The event_size bits in the first event header state the number of words in a data packet including the event headers. This value has to be determined beforehand data readout. Therefore, another FIFO memory stores the accumulated number of words in each data packet until the subsequent logic is available to compose the data packet in the final format. Concurrently, the internal event number contained in the TDC header (cf. Tab. 6.4) is cross-checked with the corresponding event label (tcs_event_no) to discard any incoherently received data in case of a fault.

Each TDC-FPGA provides a source-synchronous interface for data transfer to the MERGER-FPGA. The event data is loaded into a shift register for parallel-to-serial conversion and subsequently fed into an output DDR register provided by the I/O tiles of the Artix-7 FPGA (cf. Sec. 4.2.1.5). In DDR mode, two bits are presented on the same clock edge to the inputs of the register that outputs the data on both clock edges. Additional DDR primitives are used to forward a copy of the serial clock and a control flag indicating valid data along with the serial data stream (see Fig. 6.7).

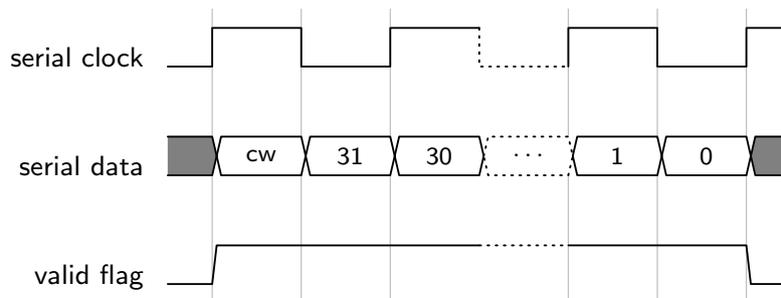


Figure 6.7: The source-synchronous interface is composed of three differential lines, serial data, serial clock and valid flag, using LVDS signalling. The figure depicts the structure of the DDR transfer.

On the receiver side inside the MERGER-FPGA, the incoming data stream and the valid flag pass through variable 31-tap delay primitives (cf. Sec. 4.2.1.5) to adjust the input timing so that the clock edge is centered in the middle of the data eye. Once at initial board operation, the link is configured to send a recognizable bit pattern, here '1's and '0's in alternating order, corresponding to a clock signal. Then the delays are stepped through until a bit-flip in the captured data is observed. At this point, the clock transitions at the edge of the data eye. With this information, the optimal delay values to be loaded at subsequent initializations are calculated.

The data transfer rate of the serial links must be adjusted to the bandwidth of the optical transceiver network. For a single board readout, the available bandwidth may be evenly partitioned among the four on-board TDC-FPGAs. With the maximum transceiver speed of 6.6 Gbit/s, the line rate equals:

$$R_{b,board} = \frac{6.6 \text{ Gbit/s} \cdot 0.8}{4} = 1.32 \text{ Gbit/s.} \quad (6.2)$$

The factor of 0.8 regards for the overhead due to 8b/10b encoding. For comparison, the maximum speed of the LVDS transmitter provided by the Artix-7 FPGA is 1.25 Gbit/s in DDR mode (see Ref. [109]). Provided that the star topology readout is applied, the line rate decreases by a factor of eight:

$$R_{b,star} = \frac{6.6 \text{ Gbit/s} \cdot 0.8}{4 \cdot 8} = 0.165 \text{ Gbit/s.} \quad (6.3)$$

The current design is operated in Single Data Rate (SDR) mode using the reference clock with a frequency of 155.52 MHz, satisfying the specification given in Eq.(6.3).

6.1.4 Implementation

As already mentioned in previous sections of the thesis at hand, potential skew in the signal path to the interpolating flip-flops has a direct impact on the uniformity of the transfer characteristic and hence on the feasible resolution. However, adequate directives are not included in the tool chain such that the automatic place and route algorithms produce satisfactory results with regard to the design requirements. In order to minimize the linearity error, the manual routing capability of the Vivado IDE, which permits the designer to interactively assign routing to specific nets, was applied (see Ref. [110]). Comprehensive studies conducted in a recent thesis [111] within the frame of this project revealed that the best results are obtained when the primary net from the input pad to the sampling flip-flops is in advance divided into multiple segments by introducing Look-Up Table (LUT) primitives as branching points (see Fig. 6.8).

The advantage of this approach is that the routing of the resulting subnets is generally less complex due to the smaller number of loads. In addition, the routing assigned to the first subgroup of flip-flops can simply be adopted by the second one. In order to keep the routes as short as possible, the flip-flops are arranged one below the other in the same CLB column with their related branching LUTs placed in the next column to the left (see Fig. 6.9). The manually assigned routes induce a skew of only 16 ps. The routing path connecting the two symmetric subgroups to the initial

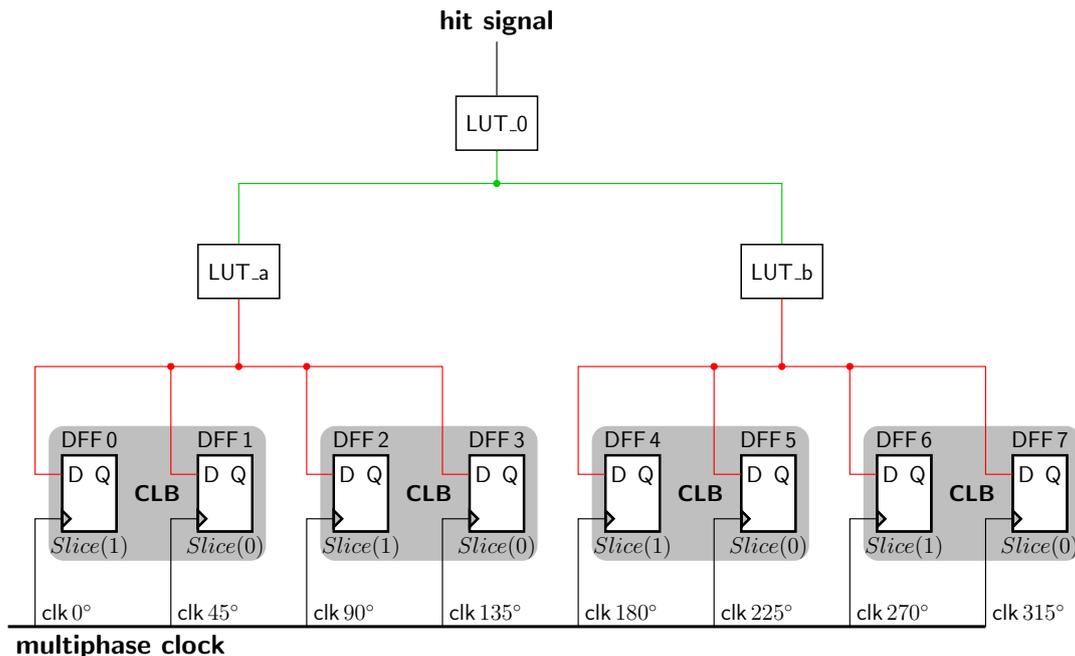


Figure 6.8: Schematic drawing of the signal path to the sampling flip-flops (DFF0 – DFF7) split into subnets using LUT primitives (LUT_0, LUT_a, LUT_b). Each sampling flip-flop occupies a different slice because the sequential primitives of a slice share the same clock input. As a reminder: a CLB comprises two slice primitives, depicted by *Slice(0)* and *Slice(1)* in this figure.

branching point marginally increases the net skew by 3 ps. A static timing analysis of the net delays (see Tab. 6.6 and 6.7) revealed that the routing skew is primarily caused by a static delay offset between the slice primitives located inside the same CLB, whereas the delays to the same slice of neighbouring CLBs is negligible. An evaluation of the associated clock nets showed an analogous distribution with an arrival time offset of around 35 ps between the slices of a CLB. This finding predicts a total effective skew of 54 ps on average, corresponding to a non-linearity of 13.4 % of the interpolators, disregarding potential phase errors of the multiphase clock.

The final routing is exported to be reused for the implementation of the remaining channels in the design. To lock down the routing in future implementation runs, the placement of the driver and receiver cells must be preserved as well. This is achieved by specifying the site locations these elements occupy on the device and their position within a slice. Furthermore, the mapping between the logical and physical inputs to the LUT components must be specified in advance to prevent pin swapping during implementation. In a script-based flow, the constraint set was replicated to implement all input channels. The site locations are defined such that the logic cells of the sampling registers are located in two columns on either side of the die close to the I/O banks. Once the routing of the hit signals is preserved, re-synthesis does not change the converter characteristic, allowing for design updates adding new functionality. No further floorplanning of the design was required to achieve timing closure. The entire TDC-FPGA design utilizes only 16 % of the flip-flop registers, 22 % of the look-up tables and 86 % of the BRAM resources available in the Artix-7 FPGA.

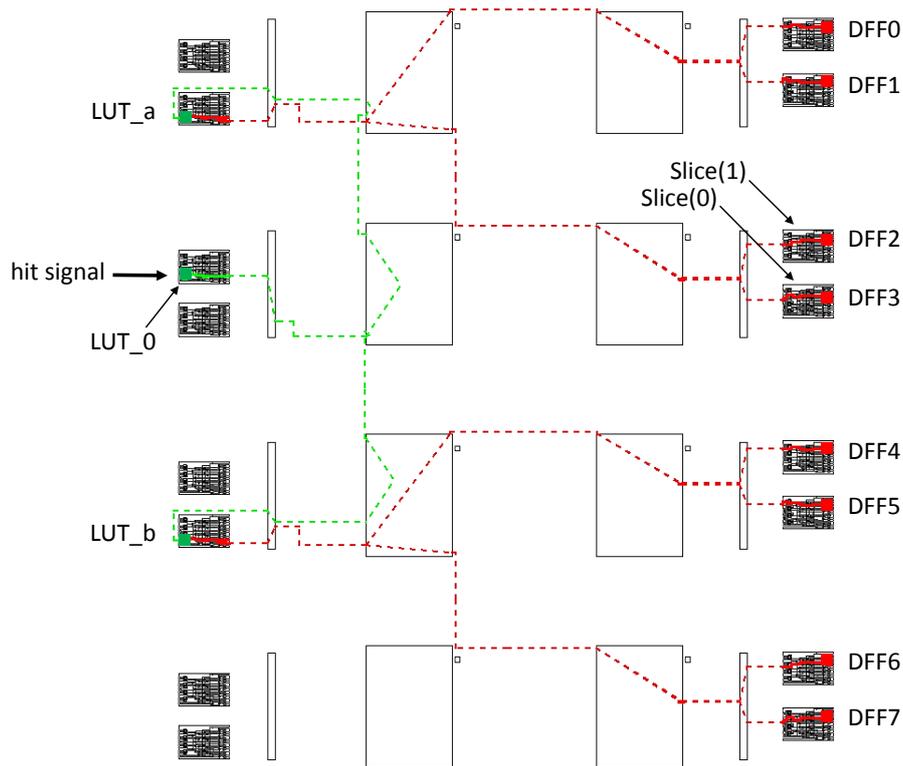


Figure 6.9: Device view of the manually assigned signal path to the interpolators of a single TDC channel.

Table 6.6: The net delays from the first branching point (LUT_0) to the next branching level (LUT_a/LUT_b) add a skew of 3 ps.

Cell	CLB slice	BEL	Delay (ps)
LUT_0	Slice(1)	A6LUT	0
LUT_a	Slice(0)	C6LUT	620
LUT_b	Slice(0)	C6LUT	623

Table 6.7: The net delays within a subgroup of sampling flip-flops cause a maximum skew of 16 ps.

Cell	CLB slice	BEL	Delay (ps)
LUT_a	Slice(0)	C6LUT	0
DFF0	Slice(1)	DFF	607
DFF1	Slice(0)	DFF	623
DFF2	Slice(1)	C5FF	607
DFF3	Slice(0)	C5FF	622

6.2 MERGER-FPGA Design

The MERGER-FPGA has the purpose to collect the output of the on-board TDC-FPGAs and, provided that the board is operated as a master in the multi-tiered front-end arrangement, to act as an inter-card pipelining device combining the data packets received from the connected slave boards via the CXP transceiver slot. Either way, the collected data is transmitted to the SFP transceiver slot to be transferred to the subsequent readout engine. Besides, the MERGER-FPGA is the centerpiece of the constant-latency up-link that aims to distribute the trigger primitives and the reference clock among the front-end boards linked together in the star topology network.

6.2.1 Constant-Latency Link

Physical events detected in the COMPASS-II spectrometer are recorded with a substantial number of different readout modules. Evaluating time intervals between the measured timestamps thus demands for a global time base. To allow for synchronous time measurements, the reference clock and control signals are expected to approach the receiver nodes with predictable latency and deterministic phase with respect to subsequent initializations and after a reset or a loss-of-lock.

At the COMPASS-II experiment, this is accomplished by the Trigger and Control System (TCS). The GANDALF module receives the TCS information via an experiment-wide passive optical network that is subsequently forwarded on fiber optic links to the ARAGORN master front-end using dedicated mezzanine cards (cf. Sec. 3.6.2). The CXP module on the master front-end again interconnects with the SFP slots of up to seven ARAGORN cards as satellites via a fiber optic breakout cable. Handling the communication with the on-board optical transceiver modules, the MERGER-FPGA is supposed to provide a constant-latency link that meets the above prospective. The high-speed links are operated using integrated transceiver primitives. A brief description of the GTP transceiver tiles is provided in Sec. 4.2.1.6. Basically, the independent transmitter and receiver channels act as parallel-to-serial and serial-to-parallel converter between the FPGA logic and the serial interface of the optical modules, respectively. To enhance the understanding of the following explanations, the link topology and the general clock structure is shown in Fig. 6.10.

Inevitable clock domain crossing between the serial and parallel side of the transceiver primitives and the programmable logic of the FPGA fabric constitutes a critical design issue. The transceiver tiles feature integrated PLLs that feed the internal clocking resources. One PLL receives a reference clock from the Si5338 clock multiplier to synthesize a clock seed for the clock data recovery circuit that extracts the high-speed clock and a divided version from the up-link data stream. This recovered clock could be facilitated as the fabric clock forwarding the up-link to the CXP module. However, this is not directly feasible due to the strict jitter requirements. Instead, the recovered clock is delivered to the on-board jitter attenuator that returns a clean copy to the FPGA fabric and to another shared PLL providing the high-speed transmitter clock.

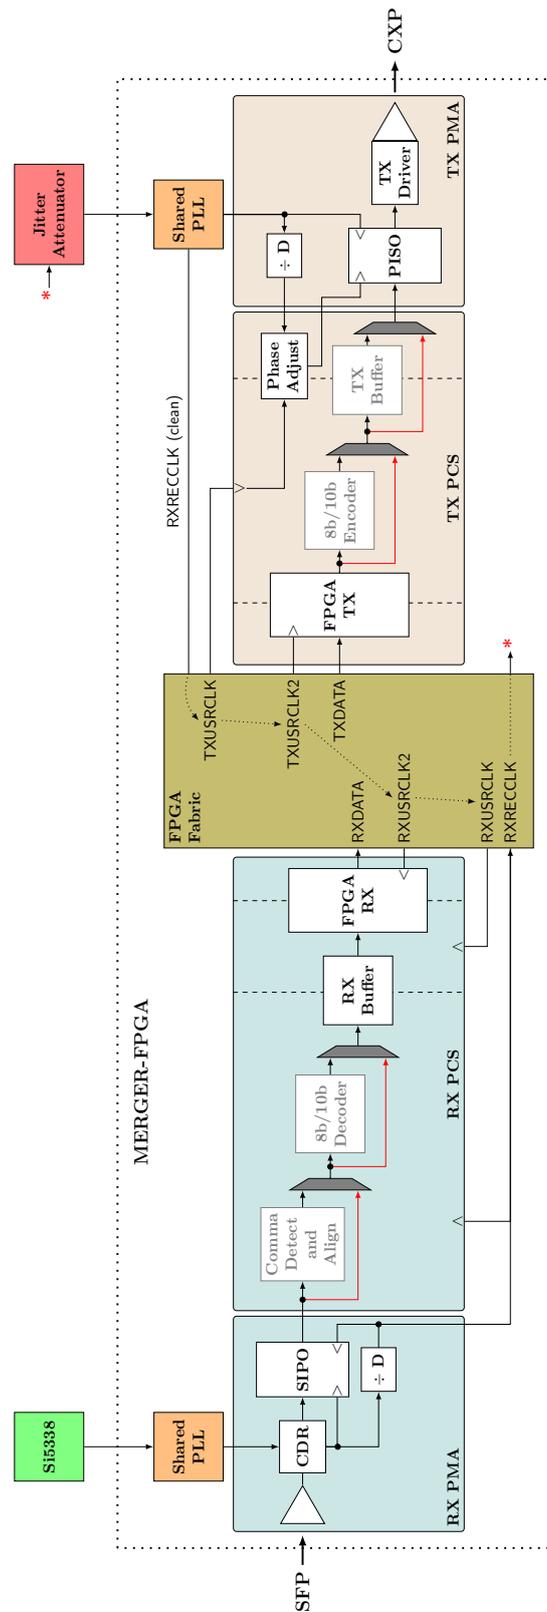


Figure 6.10: Topology of the constant-latency link showing the receiver channel attached to the SPF slot and an exemplary transmitter channel forwarding the up-link data stream to the CXP slot. The building blocks of the transceiver primitives are detailed in Ref. [74].

With the standard settings, the transceivers do not maintain the same latency through the link after a reset or loss-of-lock. The latency variations arise from the multiplication of the parallel clock at the sender side and the subsequent division of the recovered high-speed clock at the receiver side. As the divided clock can be aligned to different clock edges of the high-speed clock, the phase between the parallel clocks and thus the latency of the data transferred through the link varies alike (see Ref. [112]). Phase variations are observed in many clocking devices including PLL circuits. That is why the jitter attenuator device has particularly been selected for this application (cf. Sec. 4.5.2).

The latency variations demand for an alignment mechanism so that the received data matches the word boundaries. This project facilitates the 8b/10b encoding scheme that translates 8-bit data words into 10-bit symbols. On the one hand, the 10-bit symbols ensure a proper number of transitions to allow for a safe recovery of the clock signal. On the other hand, 8b/10b specifies 12 symbols as comma characters, commonly used to give specific meaning to the data transferred and to define an alignment sequence. The alignment sequence must not be contained within any combination of data symbols or comma characters. Only a subset of the available commas fulfills this requirement (see Ref. [113]). In fact, the alignment sequence selected for this design is a combination of two commas¹ because the internal data path of the transceivers is 20 bits wide, receiving two symbols at a time. The receiver primitives contain dedicated circuitry for automatic word alignment. When enabled, the alignment feature shifts the parallel data logically to the word boundary without changing the phase of the parallel clock. Alternatively, the alignment can be changed manually from the FPGA fabric. A dedicated control signal, when asserted, initiates a bit-shift in the parallel data. If the total number of bit-shifts is even, the alignment adjusts the phase of the recovered clock, but in event of an odd number of bit-shifts, the data word is shifted logically. As pointed out before, it is essential that the parallel clock is always aligned to the same clock edge of the high-speed clock. Hence, the integrated alignment features do not achieve the desired result.

To work around this issue, the comma detection unit and the 8b/10b decoder in the receiver is bypassed and implemented externally inside the FPGA fabric. The customized comma detection logic scans the incoming data for the alignment sequence and resets the receiver whenever the link is found incorrectly locked to the word boundaries. This approach completes quickly, as in the link idle state every fourth transmitted 20-bit symbol contains the alignment comma combination. Thereby, the process guarantees that the link is constantly monitored and automatically realigned after a loss-of-lock. Each time the up-link receiver comes out of reset and the jitter attenuator gets locked to the recovered clock, it is necessary to reset the transmitters connected to the CXP module including their associated PLLs as well. The latency through the transmitter datapath is effected by the variable phase relationship between the fabric clock and the clock domain of the parallel-to-serial converter. The phase offset between these clocks is resolved using the integrated phase-alignment capability. The fixed-latency feature is not required for the down-link which is therefore operated in the default configuration.

¹K28.1 + 28.5 – a complete listing of data symbols and comma characters can be found in Ref. [74]

6.2.2 De-formatter

The data transferred in both directions through the fiber optic network is delimited into packets to recognize related data sets at the receiver side. A packet starts with a begin marker, followed by a number of data words containing the information transferred and concludes with the corresponding end marker including an optional check sequence. Since the internal datapath of the transceivers is 20 bits wide, the markers can be defined from combinations of comma characters or commas with data symbols. A listing of control symbols that have been specified as markers for the different kind of data packets is provided in Tab. 6.8. However, specific controls like the trigger signal constitute a special case. These information must not be embedded into packets which are thus indicated with dedicated control symbols. Another example is the idle sequence that is sent to keep the link aligned whenever there is no pending data to be transmitted.

Table 6.8: Listing of control symbols.

Name	Comma	Description
command_align_comma	K28.1 + K28.5	alignment comma
command_idle	K28.2 + K28.2	idle comma
command_ftt	K28.0 + D0.0	first-level trigger
command_bos	K28.3 + D0.0	begin-of-spill
command_tcs_reset	K28.4 + D0.0	TCS reset
command_tcs_data	K23.7 + D0.0	TCS data
command_tcs_stream	D0.0 + K23.7	TCS stream data
command_config_data	K27.7 + D0.0	configuration data
command_bitstream_data	K27.7 + D1.0	TDC-FPGA bitstreams
command_start_of_data	K28.6 + D0.0	TDC data
command_eod	K29.7 + D0.0	packet end marker

The de-formatter unit de-multiplexes and buffers the received packets in FIFO primitives until the corresponding logic is available to process the data. The de-formatter process distinguishes between bitstream data to be received by the embedded processor programming the on-board TDC-FPGAs (see Sec. 6.2.4), configuration data to be stored in control registers by the configuration bus interface (see Sec. 6.3) and decoded TCS data to be later crosschecked in the data concentrator units (see Sec. 6.2.3) for consistency with the received TDC data frames. Concurrently, the up-link distributes the encoded TCS data stream so that it can be serially transferred together with the reference clock to the TDC-FPGAs using differential fanout buffers. Except for the TCS packets that are received by all front-end modules, the first word after the begin marker in a packet contains the destination address. The packet is dropped by the de-formatter logic if the destination address does not coincide with the 8-bit identifier number selected with the on-board rotary-code DIP switches.

6.2.3 Data Concentrator

The data concentrator units are designed to combine up to four input ports that can be individually disabled during operation if unused. A flow diagram of the process is shown in Fig. 6.11. The data packets are expected in the format outlined in Tab. 6.5 that permits the data concentrator unit not only to merge the received data but also to act as a local event builder rejecting incoherently received or corrupted data packets.

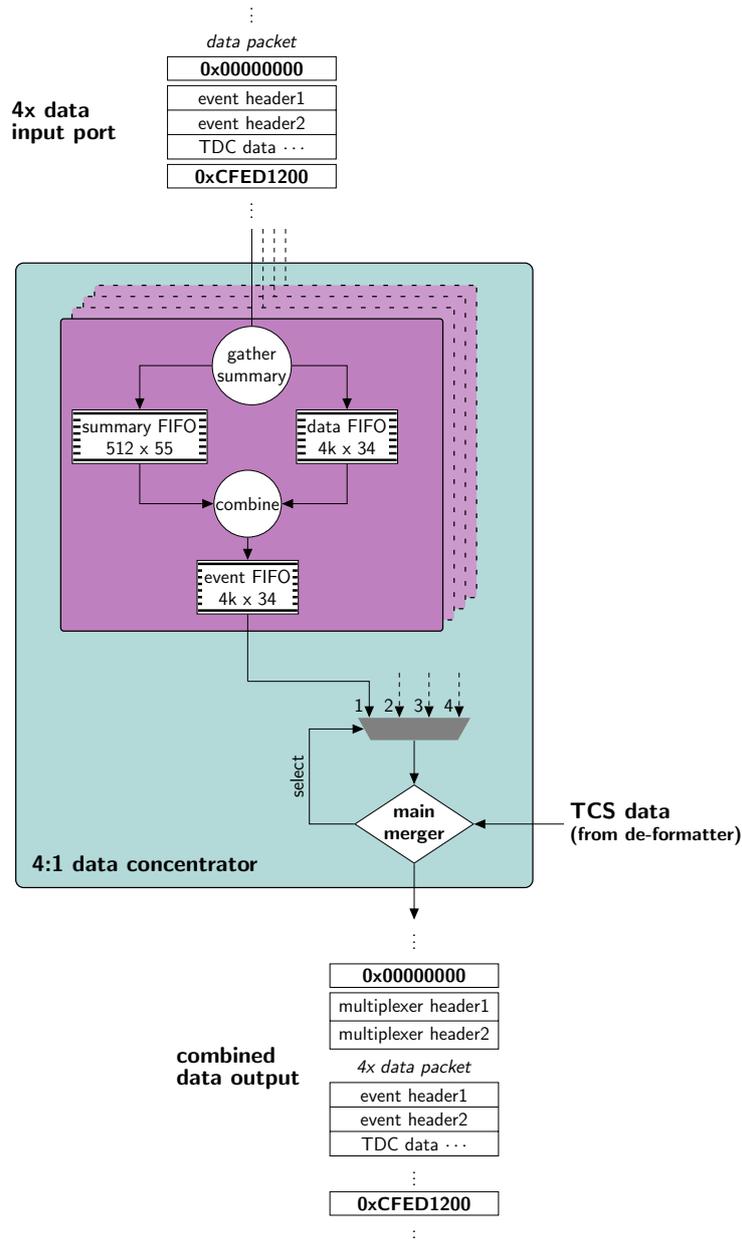


Figure 6.11: Flow diagram of the data concentrator process.

Besides the data inputs, each port receives two control bits indicating valid data and the begin and end markers in the data packets, respectively. Upon reception of the begin marker, the event labels are extracted from the subsequently received header words and the event data is temporarily buffered in a FIFO memory to calculate

the number of words contained in the packet. In case of a buffer overflow or if the received packet does not comply with the expected data structure, the residual data is rejected and the packet is properly concluded with an end marker. An internal status register flags the occurrence of such errors. The gathered information about every received packet including the error flags, the extracted event labels and the calculated event size is collected in another FIFO primitive. The data packets are subsequently prepended with the event summaries so that the subsequent logic can take appropriate measures. The combined packets are again buffered before being collected by the main process of the data concentrator unit.

The main merger process is paused until the TCS information is available and data packets are present at the input ports. If no data is received until a programmable timeout counter expires, the affected inputs are ignored in the following processing steps. Then, the event summaries are checked for active error flags and potential mismatches with the TCS event labels received from the de-formatter unit. When any of these situations occurs, the reset of the input ports concerned is asserted and the data is discarded for the duration of the current beam extraction period. Before the data packets from the different input ports are successively read out, two more header words are generated to comply with a special format that permits to join data packets from different readout modules (see Ref. [65, p. 12]). These so-called multiplexer headers use the same format like the event headers listed in Tab. 6.5. At the COMPASS-II experiment, multiplexer events are allocated to a certain address range, in the following referred to as multiplexer range, that covers all identifiers larger equal 896. By default, multiplexer headers received at the input ports are again removed which makes it simple to cascade the data concentrator units as required for the design. In fact, the MERGER-FPGA not only receives the output of the TDC-FPGAs but also the packets from the connected slave boards. Therefore, the MERGER-FPGA design implements a two-stage scheme consisting of three initial data concentrator units in the first stage and a single merger component in the second stage. The multiplexed packets in the final merger are buffered, subsequently split into 2-byte words and combined with control symbols (cf. Tab. 6.8) before the data stream is transmitted to the SFP transceiver slot for data output. Preferably, the TDC data related to a given ARAGORN board are associated with a single data packet, not only because the different TDC-FPGAs can already be distinguished by their port identifier (cf. Tab. 6.4). Moreover, it is desirable to reduce the amount of occupied addresses, which are limited in number, as they are shared among all readout modules employed in the COMPASS-II experiment. In order to gather the output into a single packet, the TDC-FPGAs are assigned to identifier numbers from the multiplexer range so that the corresponding headers are removed by the first-stage data concentrator unit. The event headers enclosing the assembled data packet, which in turn are generated by the main merger process, use an address beyond the multiplexer range to be retained in the second-stage merger.

6.2.4 Embedded Processor

The Microblaze [114] is a highly configurable, 32-bit reduced instruction set computer soft processor core, optimized for the implementation in Xilinx FPGAs. Within the Vivado IDE, the embedded subsystem is assembled using a schematic editor to

a so-called block design. Connections between the Microblaze and memory-mapped peripherals are based on the AXI interface and protocol specification [115]. Apart from the processor core itself and the general clocking, processor reset and interrupt logic, the block design contains I²C and SPI cores programming the on-board clocking devices, an external memory controller accessing the configuration Flash and a Block RAM controller to exchange data with the programmable logic. Custom interfaces like the SelectMap x16 configuration bus are implemented with dedicated general-purpose I/O cores. Once the block design is completed, the embedded system is integrated as a VHDL entity into the top-level design. Despite the rich diversity of applications, the entire embedded design consumes only 2.3% of the flip-flop registers, 4.5% of the look-up tables, 10% of the BRAM resources and 0.4% of the DSP slice elements available in the MERGER-FPGA fabric.

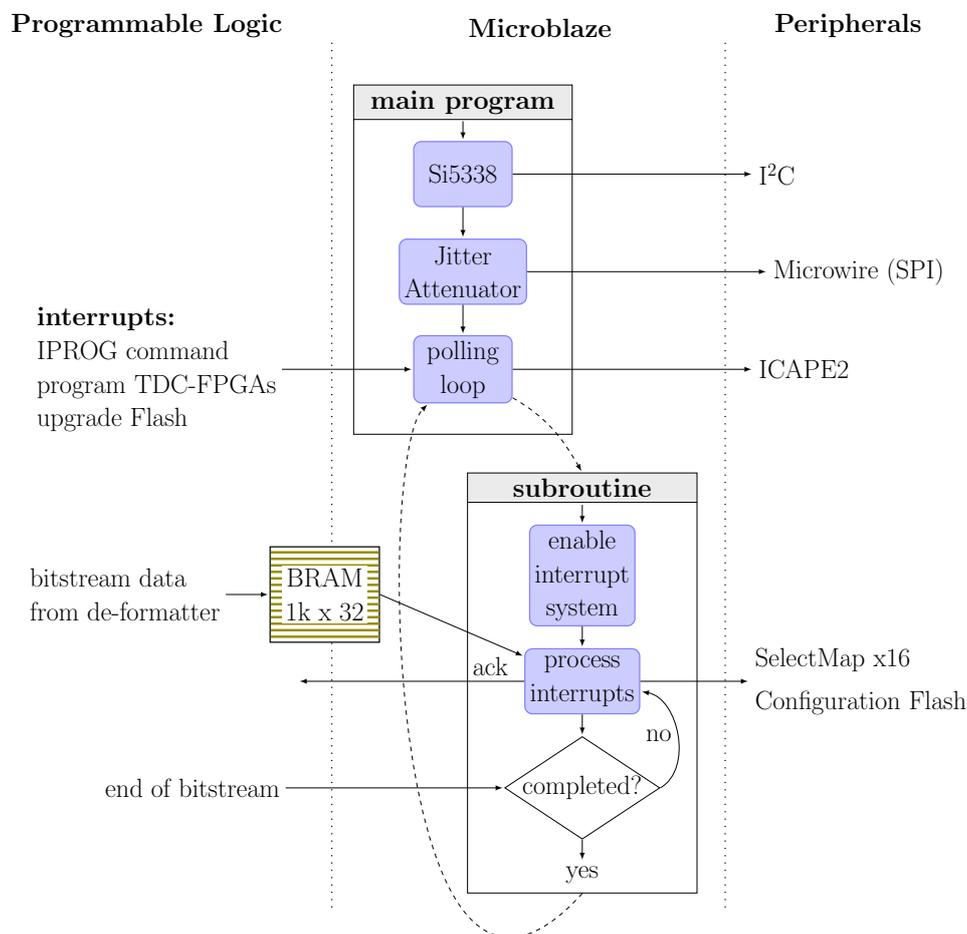


Figure 6.12: Flow chart of the embedded software application. After programming the on-board clocking devices, the software application enters a polling loop. Dedicated interrupts to the processor start a subroutine to configure the TDC-FPGAs and to upgrade the Flash content with bitstream files from remote storage.

For the development of the software application, the hardware description is exported to the Xilinx *Software Development Kit*. At first, a board support package is created providing the processor and peripherals with the required device driver libraries. Then, the build tool chain is executed on the software project to produce the executable image that is annotated back to the Vivado IDE to be included in the

bitstream, so that the Microblaze runs the embedded software application directly after the FPGA configuration is completed. The flow diagram of the embedded software application is shown in Fig. 6.12.

The strategies for programming the TDC-FPGAs and upgrading the configuration Flash by means of the embedded design using bitstream images from remote sources are briefly describes in Sec. 4.4. Because the size of the bitstream files is too large to be stored inside the MERGER-FPGA as a whole, the file content is successively exchanged between the programmable logic and the processor. Initially, the bitstream image is sent via the VMEBus interface to the GANDALF board where it is stored in on-board memory. The up-link controller then divides the file into packets by periodically sending a sequence of idle symbols after a given number of data words to give sufficient time for processing the bitstream data on the ARAGORN front-end. An interrupt informs the processor about pending data and calls a subroutine handling the data output to the peripherals. Before the next block of data is transferred, the process in the programmable logic is paused until an acknowledge flag is received from the embedded software. The repetitive process of downloading bitstream data from the programmable logic to the processor is carried out independent of the final destination interface, specified only by the interrupt service routine enabled in the interrupt controller. However, the length of the idle period and the size of the data packets must be adjusted with regard to the bandwidth of the peripherals.

At system startup, the processor programs the Si5338 clock multiplier and the jitter attenuator via the I²C respectively Microwire interface. The register maps are embedded in the data section of the executable image. These steps are obligatory for the initialization of the constant-latency link. Afterwards, the program enters a polling loop that permits the user to trigger a number of subroutines, namely the SelectMap x16 configuration mode for the TDC-FPGAs or the in-system Flash programming capability. Another interrupt asserts the advanced IPROG command in order to boot the MERGER-FPGA from a user-selectable address space in the Flash. Upon completion, the exit code of the subroutines is stored in internal status registers and is visualized using the board LEDs.

6.2.5 Analog Front-end Interface

The commissioning and testing of the ARAGORN front-end has been conducted using the analog readout electronics of the RICH-1 detector system (cf. Sec. 3.4). The readout of the photon detectors is partially performed with the 8-channel CMAD amplifier/discriminator chip [116]. Each channel consists of a low-noise amplifier followed by a shaper, a discriminator, a one shot and a LVDS driver. The gain of the amplifier can be changed from 0.4 mV/fC to 1.2 mV/fC in steps of 0.1 mV/fC and from 1.6 mV/fC to 4.8 mV/fC in steps of 0.4 mV/fC by acting on its feedback resistors and capacitors to compensate for channel-to-channel gain variations. Likewise, the threshold of the comparator and the baseline of the amplifier output can be individually adjusted by means of integrated Digital-to-Analog Converter (DAC) modules. The gain settings and the binary-encoded inputs of the 10-bit DACs are programmed using a serial transmission interface. The task of the analog front-end interface is now to distribute those settings to the CMAD channels attached via

dedicated interconnect boards to the I/O connectors of the ARAGORN front-end. Apart from providing the physical connections, the interconnect board serves as an address switch to select beforehand programming the control interface of the CMAD chip addressed.

The MERGER-FPGA design incorporates four instances of the analog front-end interface, one for each extension board connector, that can be addressed with the configuration bus interface (see Sec. 6.3). The channel settings are provided as 16-bit words. The decoding of the serial data is listed in Tab. 6.9. The upper four bits are used as the binary address to switch the de-multiplexer devices on the interconnect board (cf. Tab. 6.10), whereas the following four bits determine the DAC address (cf. Tab. 6.11). The remaining eight bits make up the opcodes specified in Tab. 6.12 together with the DAC register data. The CMAD chips are calibrated to have the baseline positioned at 650 digits so that the same threshold can be applied to all channels. The effective pedestals retrieved from calibration files must be regarded during DAC programming. A rich software toolset (see Sec. 6.6) has been developed within the frame of this thesis to load the entire settings for every detector channel connected to the ARAGORN front-ends and, inter alia, to update the configuration files with different calibrations or to change the thresholds of all channels in one go, which is particularly useful for performing fully-automated threshold scans.

Table 6.9: Decoding of the serial transmission interface.

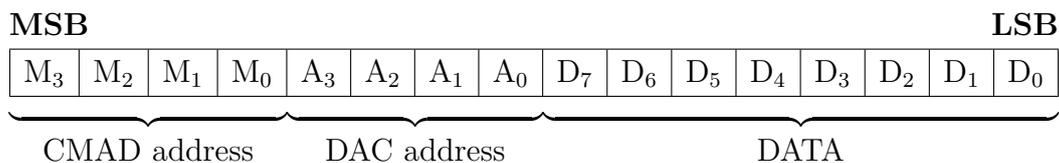


Table 6.10: CMAD address decoding.

M ₃	M ₂	M ₁	M ₀	CMAD
0	0	0	0	CMAD 1
0	0	0	1	CMAD 2
0	0	1	0	CMAD 3
0	0	1	1	CMAD 4
0	1	0	0	CMAD 5
0	1	0	1	CMAD 6
0	1	1	0	CMAD 7
0	1	1	1	CMAD 8
1	0	0	0	CMAD 9
1	0	0	1	CMAD 10
1	0	1	0	CMAD 11
1	0	1	1	CMAD 12
1	1	0	0	No Operation
		⋮		⋮
1	1	1	1	No Operation

Table 6.11: DAC address decoding.

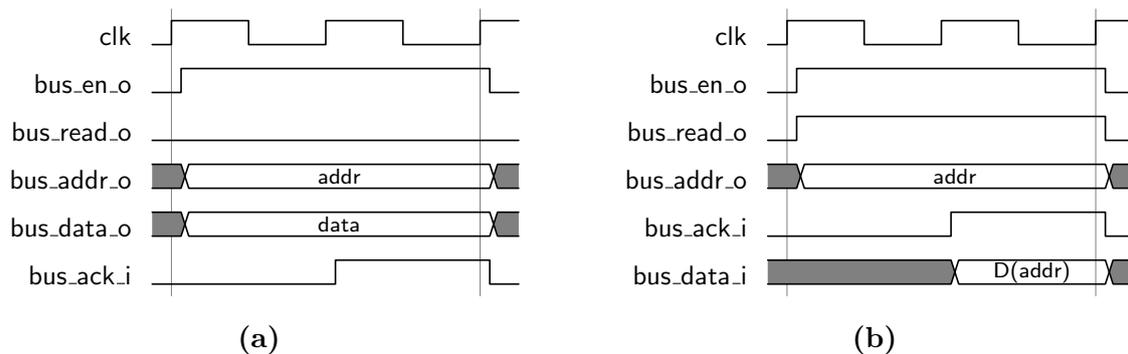
A ₃	A ₂	A ₁	A ₀	DAC
0	0	0	0	No Operation
0	0	0	1	DAC A
0	0	1	0	DAC B
0	0	1	1	DAC C
0	1	0	0	DAC D
0	1	0	1	DAC E
0	1	1	0	DAC F
0	1	1	1	DAC G
1	0	0	0	DAC H
1	0	0	1	No Operation
		⋮		⋮
1	1	1	1	No Operation

Table 6.12: Opcodes for the DAC registers.

D ₇	D ₆	D ₅	Operation
0	0	0	D ₄ D ₃ D ₂ D ₁ D ₀ → b ₉ b ₈ b ₇ b ₆ b ₅ threshold DAC
0	1	0	D ₄ D ₃ D ₂ D ₁ D ₀ → b ₄ b ₃ b ₂ b ₁ b ₀ threshold DAC
0	0	1	D ₄ D ₃ D ₂ D ₁ D ₀ → b ₉ b ₈ b ₇ b ₆ b ₅ baseline DAC
0	1	1	D ₄ D ₃ D ₂ D ₁ D ₀ → b ₄ b ₃ b ₂ b ₁ b ₀ baseline DAC
1	0	0	D ₃ D ₂ D ₁ D ₀ → gain CAP control
1	1	0	D ₃ D ₂ D ₁ D ₀ → gain RES control
1	-	1	Invalid Opcode

6.3 Configuration Bus

The configuration bus interface is an altered form of the Wishbone architecture specification [117], providing a custom configuration and monitoring interface for the on-board FPGAs. The shared bus topology consists of a single master that initiates the bus transactions accessing multiple register-based slave cores. All bus operations are synchronous to the master clock. Figure 6.13 shows timing diagrams for the bus read/write cycles.

**Figure 6.13:** Timing for bus operations: write cycle (a) and read cycle (b).

The bus architecture provides a 16-bit address space that is divided among the 32-bit wide slave registers as required in the design. The address mapping is defined in a global parameter file. Any data exchange on the bus follows the handshaking protocol that permits the selected slave to stall the operation until the request from the master is processed. The different bus signals are defined as shown in Tab. 6.13. The configuration bus interface is implemented in hardware as a multiplexer-based bus system. Figure 6.14 depicts the overall bus topology.

The configuration bus seeks not only to act as an intra-FPGA bus, but also to provide remote communication with the VME CPU. Thereby, the GANDALF module implements dedicated circuitry to bridge the gap between the VMEBus interface and the ARAGORN front-ends in order to read and modify slave register contents from the command shell (see Sec. 6.6). On this account, the bus master connects to the de-formatter unit processing the bus requests decoded from the up-link data stream.

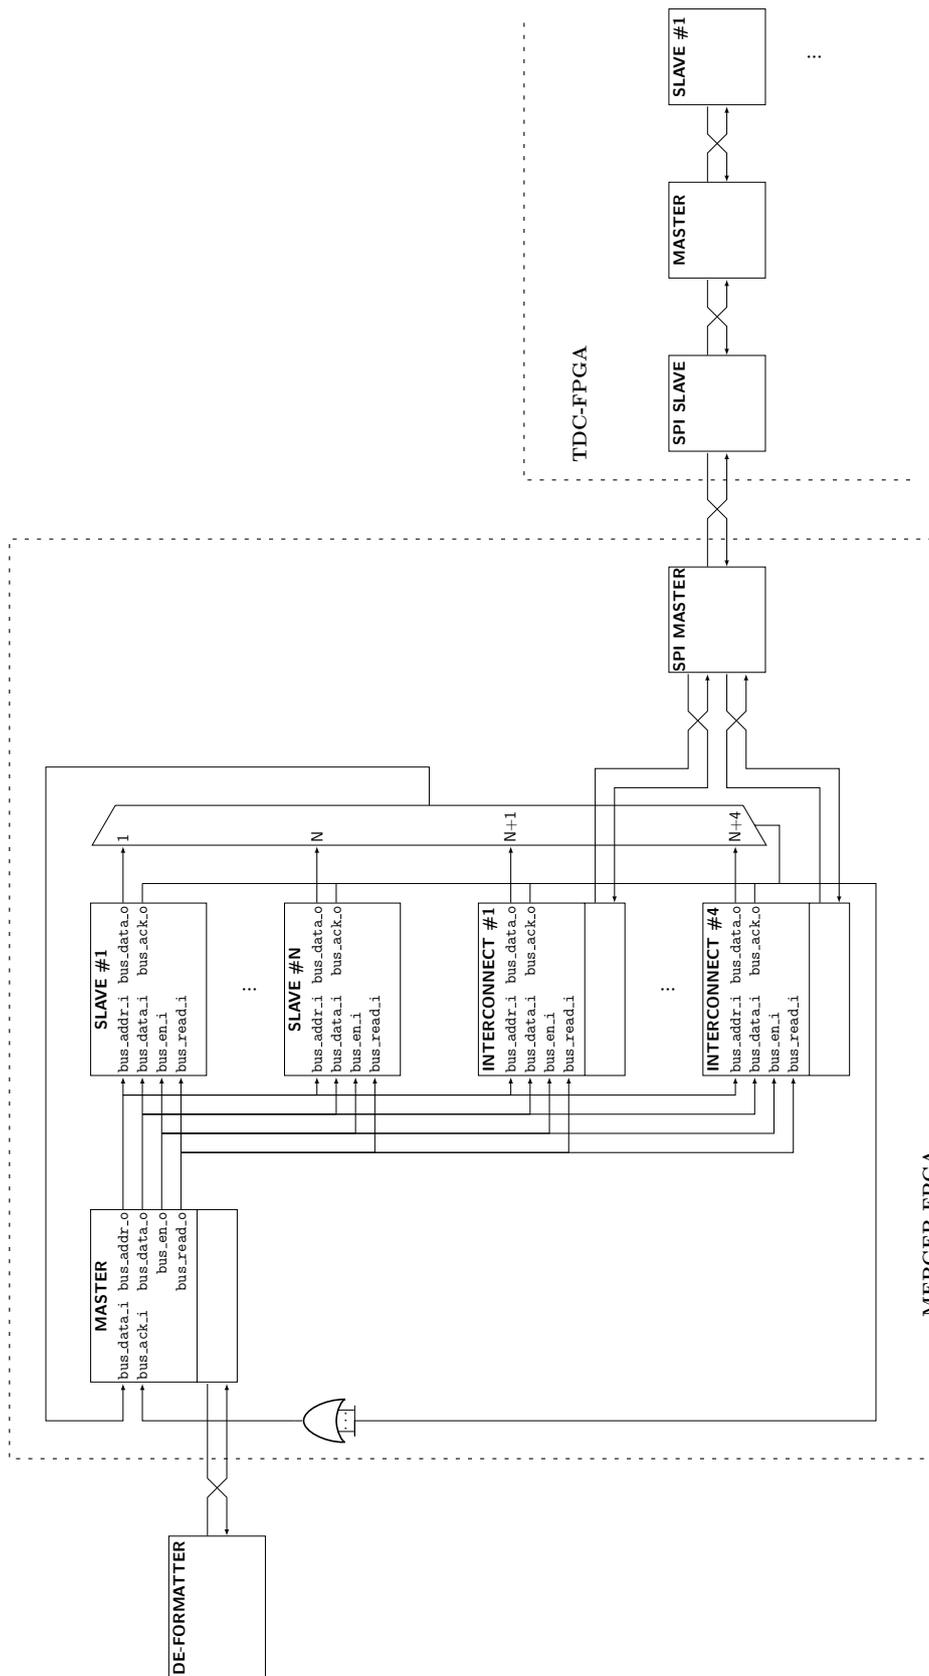


Figure 6.14: Topology of the configuration bus. Multiplexer interconnect logic connects the slave `bus_data_o` outputs to the master `bus_data_i` input. The multiplexer select lines are driven by the slave `bus_ack_o` outputs which are again collected with an OR gate to interconnect with the master `bus_ack_i` input.

Table 6.13: I/O port specification of the configuration bus interface.

Master	Slave	Description
bus_en_o	bus_en_i	indicates active bus cycle
bus_ack_i	bus_ack_o	acknowledge completion of operation
bus_read_o	bus_read_i	control signal indicating a read/write cycle
bus_addr_o	bus_addr_i	16-bit slave address
bus_data_o	bus_data_i	32-bit data to addressed slave
bus_data_i	bus_data_o	32-bit data from addressed slave

Even though the circuit area is sufficiently large for parallel transmissions within the FPGA fabrics, it is obviously clear that the parallel configuration bus interface would consume significant design space of the already very dense PCB layout if it would be directly extended to the TDC-FPGAs. To work around this issue, the connections between the MERGER-FPGA and the TDC-FPGAs are established using dedicated interconnect cores. Thereby, each of the interconnect slaves incorporates an interface accessing the SPI bus that provides a full-duplex serial link to the TDC-FPGAs. Resuming the bus architecture, each SPI slave in the TDC-FPGAs is again attached to another bus master instance.

However, this pipelining approach demands for a revision of the bus protocol. A write cycle, for instance, addressing a slave of the TDC-FPGAs, requires to store the data together with the slave address first to dedicated registers of the interconnect core, which subsequently conducts the SPI transaction. Accordingly, a read cycle involves a write operation to provide in advance the slave address before the SPI request can be processed. A subsequent read operation finally retrieves the data from a dedicated register of the interconnect core as shown in Figure 6.15.

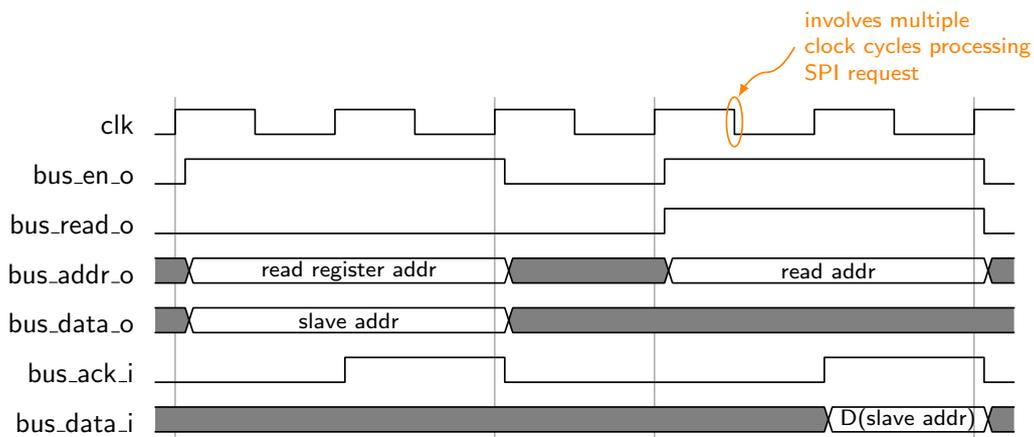


Figure 6.15: Bus read cycle addressing a slave of the TDC-FPGAs. First, *slave addr* is written to *read register addr* of the interconnect core. Second, a read operation from *read addr* retrieves the register content via the SPI interface.

6.4 TCS Interface

As mentioned earlier, the Trigger and Control System (TCS) broadcasts the first-level trigger items together with the corresponding event labels via optical links to the readout modules employed in the COMPASS-II experiment. The TCS makes use of bi-phase mark encoding and the time-division multiplexing method to transmit two independent channels simultaneously (see Ref. [118]). The first channel is designed for minimum latency and carries the first-level trigger signals, whereas the second channel transmits the TCS commands. One of these commands distributes the begin-of-spill and end-of-spill signals together with a spill number identifier that indicate the time slot in which protons are extracted from the Super Proton Synchrotron. The begin-of-spill signal for instance synchronizes the coarse counter primitives in the TDC-FPGA design across all ARAGORN front-end boards. Another command contains the event number and event type labels to the corresponding first-level trigger. The TCS receiver of the GANDALF module, which is located on a separate add-on card, extracts the global reference clock and transmits both the clock and the data stream to the FPGA on the mainboard, where the individual channels are de-multiplexed and the TCS commands are decoded. The ARAGORN front-ends receive the TCS information via the constant-latency link. Inside the MERGER-FPGA, the de-formatter unit parses the decoded data to the user logic. Reducing the limited number of FPGA I/Os, the constant-latency link likewise distributes the encoded data stream so that the TCS information can be serially transmitted to the TDC-FPGAs, which incorporate another instance of the TCS interface module each.

6.5 Project Management

The source files of the FPGA design projects are managed using the Git¹ version control system that has been chosen due to a number of advantages over similar tools. As an example, each working copy of a Git repository includes the full development history which lets designers work independently of the project's evolution in the central repository. The Git version control system also allows for the development of new features in isolated branches which has the benefit of a 'clean' master branch comprising only stable commits. At a later time, when the work on the design feature is completed and successfully verified, the feature branch can be merged into the master branch. In the scope of this thesis, it turned out that it is desirable to reuse specific modules of the code base, for instance the configuration bus feature that is implemented in both the TDC-FPGA and the MERGER-FPGA design. Likewise, the TCS interface module developed for the GANDALF framework can be used as is for this project. The Git version control system addresses this issue with the possibility to include other repositories as subdirectories, referred to as submodules. Thereby, the parent repository points to a particular commit of the submodule's repository to keep track of a certain development stage.

The Tcl support of the Vivado Design Suite permits to run the implementation flow in a script-based manner. Though by default the synthesis and implementation steps were carried out in the Vivado IDE, a Tcl script can be used to recreate and

¹<http://git-scm.com/>

recover the entire project and tool settings so that only a subset of the design files must be put under version control. Figure 6.16 depicts the directory structure of the repositories.

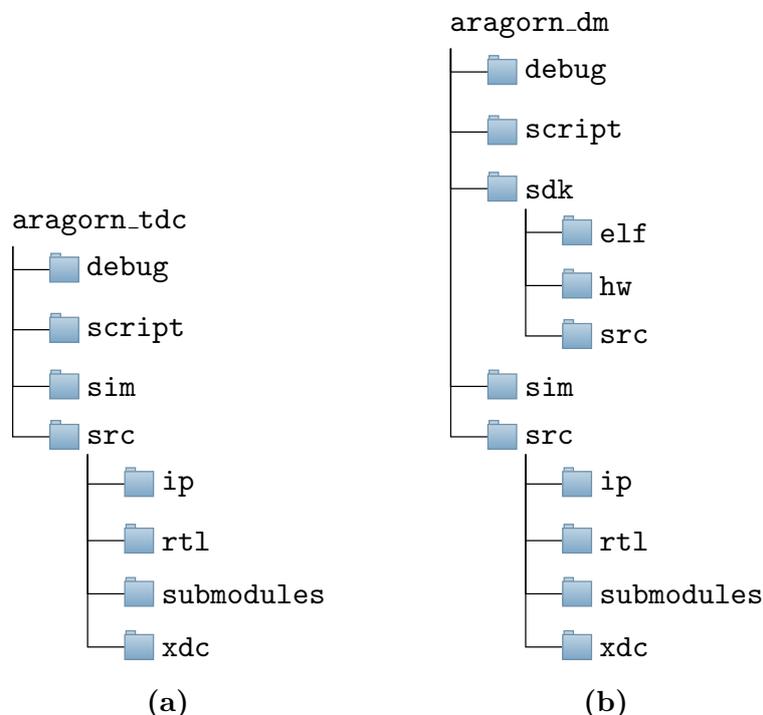


Figure 6.16: Directory structure of the TDC-FPGA design repository (a) and the MERGER-FPGA design repository (b).

The intended purpose of the different folders is as follows:

debug files related to the Vivado Logic Analyzer debug cores.

script Tcl scripts for project creation and compilation.

sim VHDL testbench modules for behavioral simulations and scripts related to the simulator software.

src the source files of the design grouped in following subdirectories:

ip IP catalog file (.xml), IP configuration file (.xci) and the VHDL (.vho) IP instantiation file for each IP core.

rtl VHDL (.vhd) modules.

submodules imported Git repositories.

xdc constraint (.xdc) files.

In addition to the folders listed above, the repository for the MERGER-FPGA design includes the subtree of the embedded software project under the **sdk** directory:

elf the executable image to be embedded in the bitstream.

hw the exported hardware description.

src the C source code files.

To recreate the Vivado IDE project, for example of the MERGER-FPGA design, the user executes:

```
vivado -source script/aragorn_dm_prj.tcl -notrace
```

in a command prompt. Accordingly, the embedded design can be restored, without the need to include any IP core files in the repository, by running:

```
source ./script/emb_dsgn_prj.tcl
```

in the integrated Tcl shell.

Any changes to the project settings or the block design must be annotated back to the scripts using appropriate Tcl commands that are detailed in Ref. [119]. After downloading the repository, the implementation results are reproduced unambiguously using this workflow. Thus, there is no need to put the configuration bitstreams under version control. However, the relevant commits have been tagged to clearly associate the bitstream files on the server with the development history in the repositories.

6.6 Software Tools

This section describes a selection of command line tools that have been developed for the configuration and monitoring of the ARAGORN front-end.

send_frontend

This console command addresses slave registers of the configuration bus interface. Though the capability to process bus read requests is provided by the ARAGORN front-end, the corresponding transactions have not yet been implemented in the GANDALF framework at the time of completion of the thesis at hand.

```
usage: send_frontend hexid addr data
arguments:
  hexid : address of the GANDALF module
  addr   : 16-bit slave address; the upper 8 bits determine the
           identifier number of the ARAGORN front-end,
           the lower 8 bits specify the register address.
  data  : register data
"arguments given as hexadecimal numbers"
```

aral.py

This application conducts the programming of the TDC-FPGAs, provides the system settings to distinct slave registers of the configuration bus and configures the analog readout modules attached to the ARAGORN front-end. The main program leverages a number of Python tools that can also be used autonomously as console commands. Due to the large number of system parameters, the application retrieves the settings from a user-editable configuration file. Similar to the format used by Windows INI files, the file content is categorised by so-called sections that contain the entries as 'name/value' pairs, for instance slave address and related register data passed to the configuration bus interface. The successful completion of the program can be verified with an application log that gathers the exit codes and status messages received from the different subprograms. The log messages are classified according to different levels of importance. This lets the user control the logging system, for instance, to record only error messages or to increase the amount of information when debugging.

```
usage: aral.py [options] brdIDs
arguments:
  brdIDs: list of IDs of the ARAGORN front-ends to be configured
options:
  -c cfg, --config=cfg      : configuration file with all settings
  -g hexid, --hex_id=hexid  : address of the GANDALF module
  -t, --tdc                    : program TDC-FPGAs
  -v, --verbose                : show info messages
  -h, --help                    : show help message and exit
```

cmad.py

This tool is used by the previous application to program the gain, baseline and threshold settings of the CMAD channels. Those parameters are retrieved from a separate configuration file. Furthermore, the program loads the baseline calibrations that are required once after initial system installation or when an analog readout module has been exchanged. Thereby, the tool browses the directory given as an optional argument for a file that contains the list with calibration files to be applied. Another useful feature is the capability to absolutely or incrementally modify the threshold values in the settings file.

```
usage: cmad.py [options] brdID
arguments:
  brdID: ID of the ARAGORN front-end
options:
  -g hexid      : address of the GANDALF module
  -s sett_dir   : the directory containing the settings file
  -c calib_dir : load calibrations from this directory
  -t thres     : set threshold values in settings file
  -i            : increment threshold values in settings file by thres
  -p           : program CMAD chips
  -h           : show help message and exit
```

7. Verification

7.1 Test Equipment

A fanout buffer extension board (see Fig. 7.1) has been developed in collaboration with the local electronic workshop to provide test signals to the TDC inputs of the ARAGORN front-end. The centerpiece of the fanout board is a 4x4 crosspoint switch [120] that receives two differential input signals via an RJ45 Ethernet plug. The remaining inputs of the crosspoint switch are constantly driven to a logic low level. The differential output drivers of the crosspoint switch incorporate individual 4:1 multiplexers¹ that can select any of the four input signals. Four out of eight multiplexer select pins are driven by the two remaining outputs² of the RJ45 plug, converted on-board to a single-ended standard (3 V CMOS). Table 7.1 lists the input/output combinations that can be selected using this design.

Table 7.1: Crosspoint switch input/output select table. Even more combinations are possible using on-board jumpers that access the remaining control pins.

S20	S21	S40	S41	Input/output combination
0	0	0	0	1st test signal to all outputs
0	1	0	1	} 1st test signal to output 1,3 and output 2,4 logic low
1	0	1	0	
1	1	1	1	1st and 2nd test signal interleaved

The fanout board incorporates a right-angle connector³ mating with the extension board connectors on the ARAGORN front-end. To provide copies of the test inputs to 96 differential pin pairs of the output connector, the crosspoint switch

¹The crosspoint switch provides eight control pins: S10/S11 select the input to the first output, S20/S21 second output, S30/S31 third output and S40/S41 fourth output.

²To select the input to the second and fourth output of the crosspoint switch, the first signal drives S20 and S40, the second one controls S21 and S41.

³Samtec QFS-104-01-L-D-RA

is followed by a two-stage buffer scheme. The first stage includes four 1:4 LVDS buffers¹ that again connect to a set of sixteen 1:6 LVDS buffers² in the second stage. Corresponding outputs of device pairs from the first buffer stage are routed in alternating order to the inputs of the second buffer stage. The same approach was followed in the layout of the interconnects between the second stage and the output connector. The idea behind this is to pursue the input/output configuration of the crosspoint switch after each buffer stage. As a result, the board not only acts as a 1:96 fanout, but also can provide the ARAGORN front-end with copies of two independent test signals received on neighbouring input channels. Furthermore, every second channel can be disabled to study the impact of cross-talk induced noise. The measured time jitter between the test signal outputs is in the order of 40 ps. Figure 7.2 shows a photo of the ARAGORN front-end mated with fanout buffer extension boards.

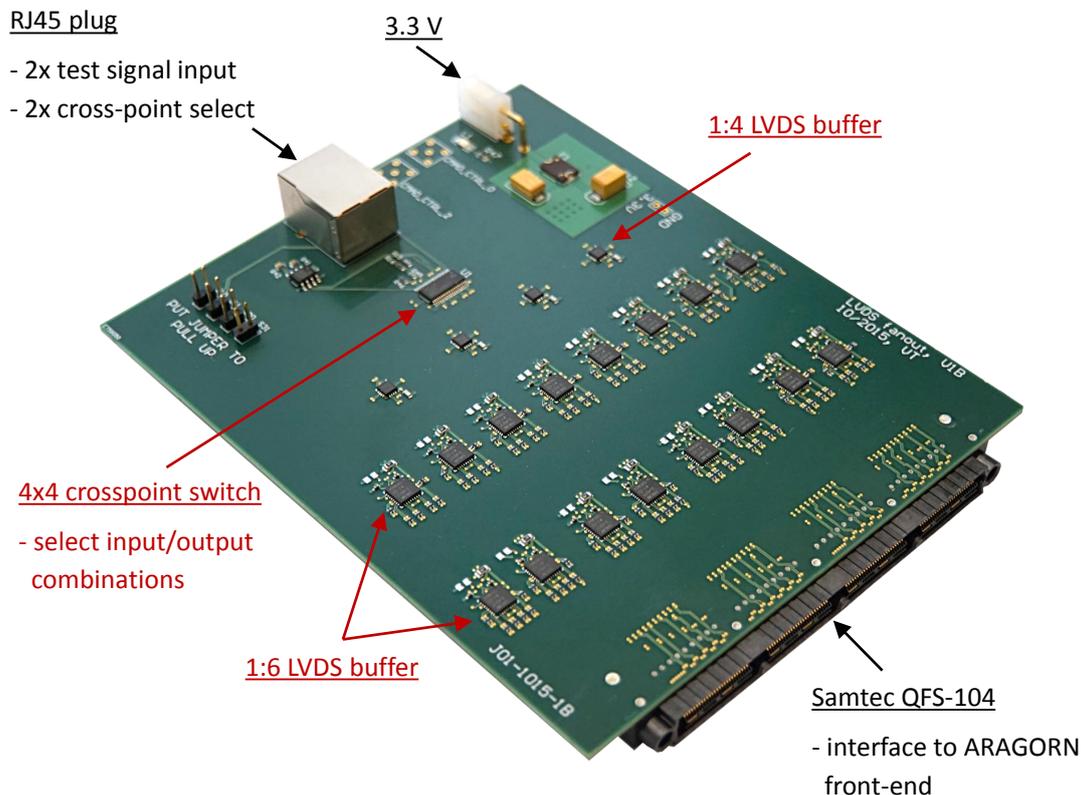


Figure 7.1: The fanout board delivers copies of up to two independent test signals to the TDC inputs of the ARAGORN front-end.

¹Microsemi ZL40215

²Microsemi ZL40217

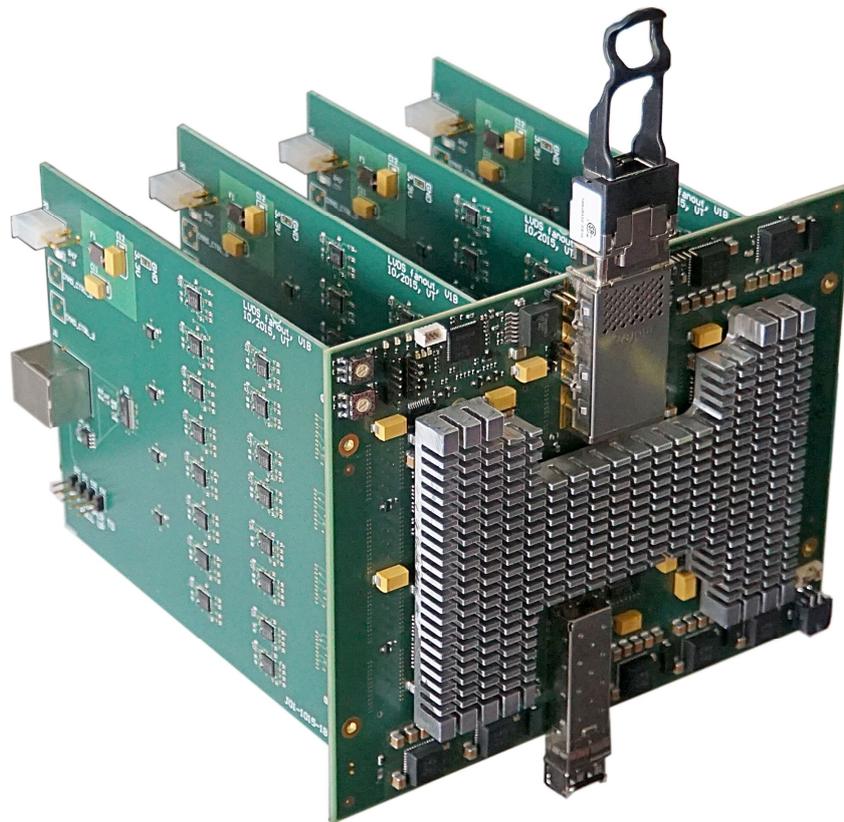


Figure 7.2: Picture of the ARAGORN front-end equipped with four fanout buffer extension boards.

7.2 Test Setup

The measurement setup used for the TDC characterization involved an ARAGORN master card and subsequent slave boards connected up to the CXP transceiver module using a fiber optic breakout cable. The master front-end was attached via the SFP slot to a GANDALF module equipped with ARWEN mezzanine cards. Data readout was accomplished using a dedicated backplane link card following the S-LINK specification [48]. All measurements were conducted using test signals generated with a Tektronix AFG3252 [121] dual-channel pulse generator.

A VHDCI-to-RJ45 breakout cable was employed to interconnect the fanout boards with a GANDALF module equipped with digital mezzanine cards (cf. Sec. 3.6.2), that again received the test signals from the pulse generator. This had the advantage that the multiplexer select signals were accessible via the VMEBus interface in order to change the configuration of the fanout boards during operation. Providing the ARAGORN front-ends under test with a reference clock and the trigger signals, another GANDALF module with ARWEN add-on cards was employed as TCS controller. The reference clock was generated using the 20 MHz oven controlled-crystal oscillator located on the GIMLI add-on card (cf. Sec. 3.6.2). The trigger signals entered through a LEMO input connector. Figure 7.3 shows a schematic overview of the test setup. A photograph of the test setup is shown in Fig. 7.4.

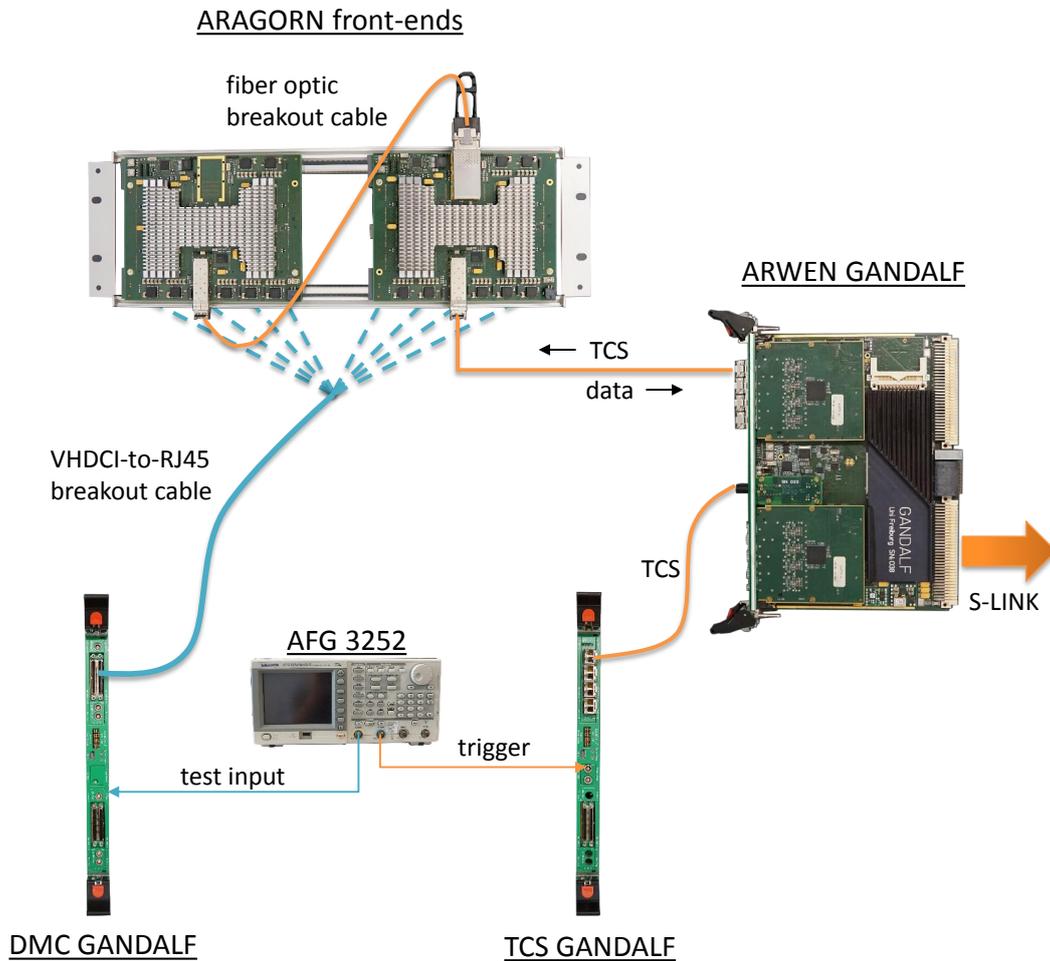


Figure 7.3: Overview of the measurement setup. The ARAGORN front-ends – the master front-end and a single slave card in this picture – were combined with fanout buffer extension boards. The test and trigger inputs originated from a AFG3252 pulse generator. A GANDALF module mated with a digital mezzanine card (DMC) was employed to pass on the test signals to the fanout boards. Data readout was performed with a GANDALF module equipped with ARWEN mezzanine cards. In the up-link direction this module distributed the TCS information received from another GANDALF acting as TCS controller.

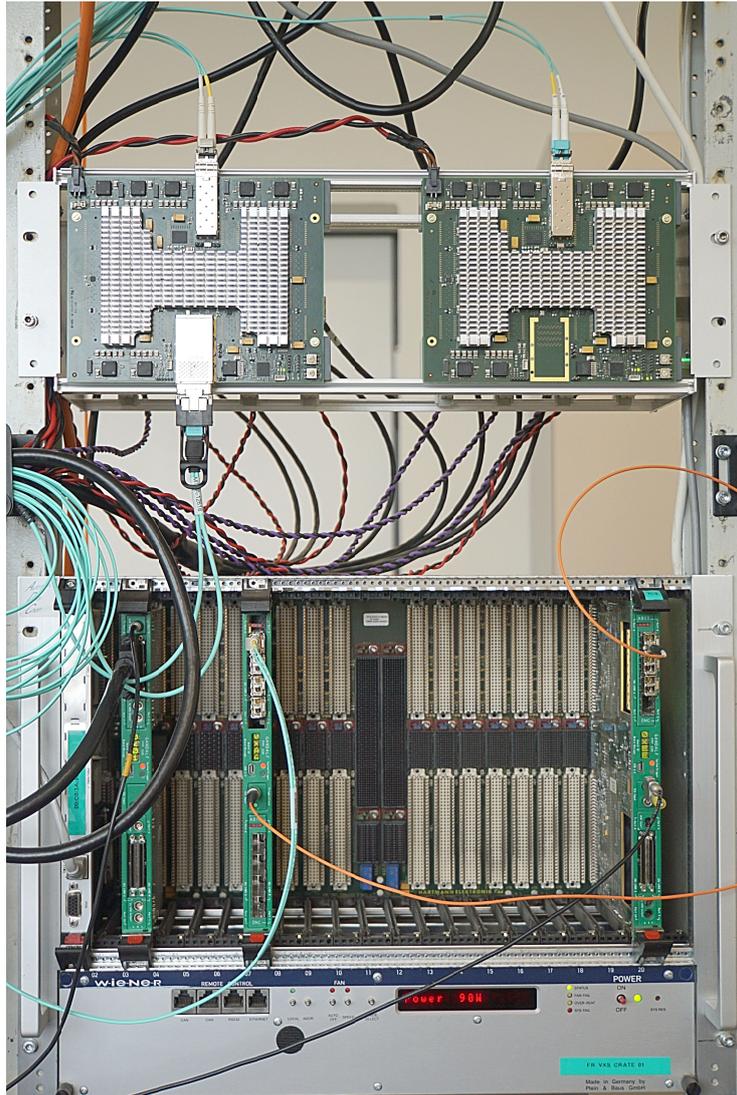


Figure 7.4: Picture of the test setup used to verify the performance metrics of the TDC application. The AFG3252 pulse generator and the data acquisition system are not visible.

7.3 Test Results

7.3.1 Star Topology Network

The design of the transceiver network has been validated using the integrated Bit Error Ratio (BER) tester [122]. The BER tester implements data pattern generators and checkers and provides access to the dynamic reconfiguration port of the transceiver primitives. This allows designers to examine the performance of the system under test and to tune various link attributes like pre-emphasis settings at run time.

7.3.1.1 Bit Error Ratio

The BER tester offers different pseudo-random bit sequence patterns for serial link validation. A pseudo-random bit sequence generator can be implemented in hard-

ware using a linear-feedback shift register. Unlike basic shift registers, the linear-feedback shift register provides a feedback path from the output of the last flip-flop to its input and, using XOR connections, to selected bits in the register. The principle underlying the shift operation is based on the idea that bit strings can be interpreted as polynomials with coefficients either 1 or 0. A n -bit linear-feedback shift register thus defines a characteristic polynomial of degree n . Its non-zero coefficients correspond to the bits with feedback connection. The different states represent polynomials of degree less than n . Sequential shifts perform arithmetic operations of the state polynomial with the characteristic polynomial as follows. At first, the present polynomial is multiply by x . Then, the intermediate polynomial is divided by the characteristic polynomial and the remainder gives the polynomial in the next state. The length of the binary sequence produced, before it repeats, depends on the initial state and the characteristic polynomial used. The maximum length is $2^n - 1$, which means that the linear-feedback shift register cycles through all possible binary values except zero. Such maximum length generators are referred to the degree of the underlying polynomial. For instance, the transceiver network on the ARAGORN front-end was tested using PRBS-7 and PRBS-31, which resemble 8b/10b and 64b/66b encoded data streams, respectively (see Ref. [123]). Compared to PRBS-7, the larger number of consecutive identical digits observed in PRBS-31 causes increased deterministic jitter and therefore applies a more stringent test to high-speed transceiver links.

The bit error ratio is a measure for the probability that a bit is incorrectly transmitted through a device under test. In practice, a predefined bit sequence is provided to the input of the system that can be checked against the output data stream. The ratio of bit errors detected to the total number of bits transmitted can then be used as an estimate of the actual bit error ratio. As bit errors can occur at random times, the accuracy of this approach depends on the data volume transmitted. In order to quantify the amount of time required to confirm that the true bit error ratio is better than an upper limit, a confidence level must be specified.

The probability $P_n(k)$, that k bit errors are detected within n transmitted bits can be described with the binomial distribution function. For most digital systems, the assumption can be made that p , the probability that a bit error occurs for each transmitted bit, tends toward zero. If n is at least BER^{-1} , so that $np > 1$, the Poisson theorem can be used to approximate the binomial distribution:

$$P_n(k) = \frac{(np)^k}{k!} e^{-np}. \quad (7.1)$$

The probability that more than N errors are detected can be written as:

$$P_n(\epsilon > N) = 1 - \sum_{k=0}^N \frac{(np)^k}{k!} e^{-np}. \quad (7.2)$$

If for an actual measurement less than N bit errors are observed, this probability can be treated as the confidence level in percent that the true bit error ratio is

smaller than a specified p (see Ref. [124]). Followed from Eq. (7.2), the number of bits required to validate a bit error ratio limit for a specified confidence level CL is:

$$n = -\frac{\ln(1 - CL)}{p} + \frac{\ln\left(\sum_{k=0}^N \frac{(np)^k}{k!}\right)}{p}. \quad (7.3)$$

If no errors are detected during the test ($N = 0$), the equation is simplified to:

$$n = -\frac{\ln(1 - CL)}{p}. \quad (7.4)$$

To test the performance of the optical transceiver network from the transmitter on the master front-end to the receiver on the slave boards and in the reverse direction, the BER tester feature was implemented inside the MERGER-FPGA of two ARAGORN boards. The BER tester core on the transmitter side was configured to repeatedly transmit a test pattern through the link to be checked against an internally generated pattern by the BER tester core on the receiver side.

The objective was to verify a bit error ratio better than 10^{-14} for a confidence level of 99%. For comparison, the IEEE standard for 10-Gigabit Ethernet [125] specifies $BER < 10^{-12}$. Solving Eq. (7.4) for $p = 10^{-14}$, $CL = 99\%$ yields 4.61×10^{14} bit to be transmitted. For the default data rate of 3.1104 Gbit/s, assuming an error-free measurement, the amount of test time required is 41 h. Indeed, no errors were observed neither for PRBS-7 nor PRBS-31 test patterns. Without modifying the reference clock structure on the ARAGORN front-end, it was possible to operate the system at 6.2208 Gbit/s. Again, the tests were completed with zero errors. These results confirm the prior adoption that the design can operate with $BER < 10^{-14}$ even at the bandwidth limit (6.6 Gbit/s) of the transceivers.

7.3.1.2 Receiver Margin

Although the performance of the design has successfully been verified, it is interesting to examine the signal quality to determine the available margin at the receiver. This can be done on the basis of an eye diagram. The eye diagram is commonly derived from samples of the data stream recorded with a high-speed sampling oscilloscope by superimposing each bit from the digital waveform.

The continuous time linear equalizer embedded in the receiver primitives compensates for high-frequency attenuation in the conductors so that any eye diagram measured with external instruments will differ from the data eye as seen internal to the receiver. Therefore, the transceiver provides built-in eye scan circuitry to examine the eye opening after the equalizer. The eye scan circuitry operates an additional sampler that can be programmed with different horizontal and vertical offsets from the sampling point determined by the clock and data recovery circuit. The horizontal offset adjusts the sampling time of the offset samples, while the vertical offset raises or lowers the differential voltage threshold. For each offset point, the eye scan measurement compares (bit by bit) the offset samples with the data samples. The calculated bit error ratio is the ratio of the error count to the specified number of data samples transmitted. Scanning the full range of horizontal and vertical offsets, a BER map is created to visualize the eye margin at the receiver.

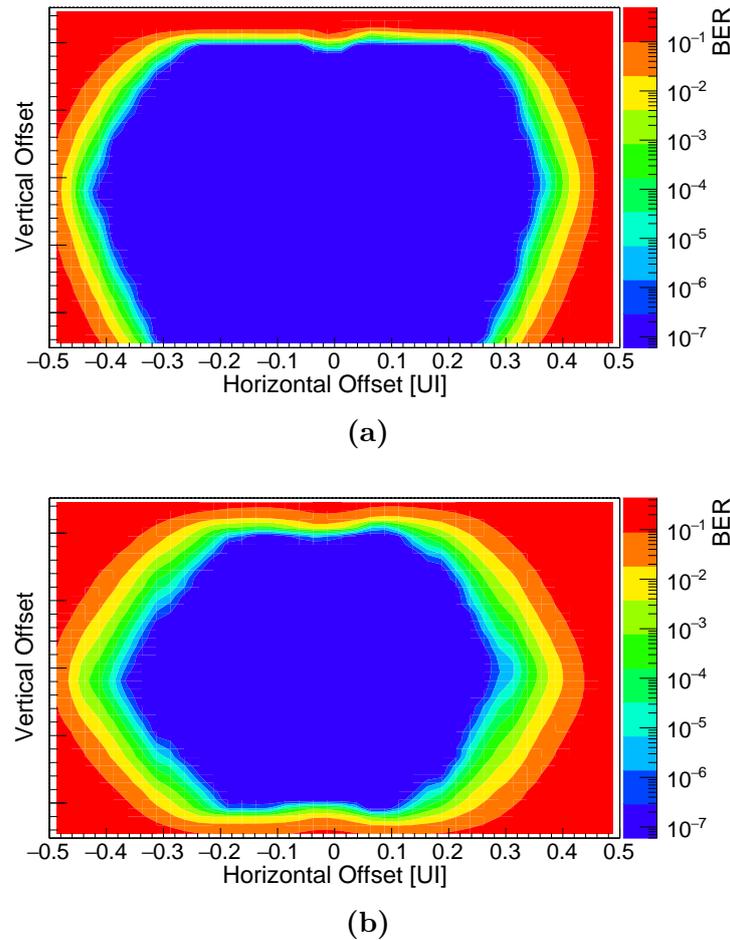


Figure 7.5: Image of the statistical eye at 3.1104 Gbit/s (a) and 6.2208 Gbit/s (b) generated from 2D eye scans. The horizontal offset denotes the sampling time with respect to the data samples. The vertical offset describes the differential voltage threshold to which the equalized waveform is compared.

At 3.1104 Gbit/s, the data eye appears almost completely open (see Fig. 7.5a), while at 6.2208 Gbit/s the eye opening is apparently smaller (see Fig. 7.5b). Unfortunately, margins cannot be quantified as the FPGA vendor does not provide information about actual compliance masks. However, from a qualitative view of the eye diagrams it can be concluded that the link does not suffer from significant impairments that would demand for further tuning.

7.3.2 TDC Characterization

This section describes in detail the measurements that have been performed to fully characterize the performance metrics of the TDC-FPGA design. The measurements cover inter alia the linearity errors of the transfer function, the time resolution under different operating conditions, the rate capability and the long-term stability of the system. The measurement results and the benchmarks of the TDC application are summarized in Tab. 7.2.

Table 7.2: Characteristic parameters of the TDC application.

Channels TDC-FPGA/ board	96 / 384
Reference clock frequency	311.04 MHz
Time bin size (LSB)	402 ps
Dynamic range	211 μ s (16 bit)
Double hit resolution	3.2 ns
Rate capability	34 MHz
Dead time	none
Differential non-linearity	min. 15.0 %, max. 35.7 %, avg. 23.5 %
Integral non-linearity	min. 11.5 %, max. 35.7 %, avg. 22.8 %
Time resolution ($\Delta t < T_{\text{clk}}$):	
Single board w/o cross-talk	163.6 ps
Single board w/ cross-talk	164.0 ps
Time resolution ($\Delta t > T_{\text{clk}}$):	
Single board	164.4 ps
Multiple boards	167.2 ps
Systematic effects:	
Repeated system initializations	16.4 ps peak-to-peak
Long-term stability	10.5 ps RMS
Temperature stability	0.21 ps/ $^{\circ}$ C

7.3.2.1 Differential and Integral Non-linearity

The differential non-linearity was determined using a statistical code density test. In this exercise, the TDC inputs received a train of pulses that originated randomly in time from the function generator. After a substantially large number of measurements N , assuming the TDC bins are identical in size and the input pulses appear asynchronous to the sampling clock, a histogram of the number of entries by bin number would show a uniform distribution $n = N/8$. In real applications, inevitable skew in both the signal path and the clock lines to the interpolating flip-flops causes imperfections in the transfer characteristic that lead to differential non-linearities:

$$\text{DNL}_i = \frac{n_i - n}{n}, \quad i = 0, 1, \dots, 7, \quad (7.5)$$

where i denotes the bin number. For the measurement of the differential non-linearity, a sample of $N = 53 \times 10^3$ hits per channel was recorded. The binary-code output of the TDC gives the measurement result in units of LSB. Hence, the bin number in which a hit was observed can be derived from a modulo eight operation. Despite the fact that the routing of the input signals to the interpolators was fixed in the implementation flow and a uniform arrival time offset between the multi-phase clock nets was presumed throughout the device, the differential non-linearity slightly varies from channel to channel. Histograms of the transfer characteristic for the channel with the smallest and largest differential non-linearity are shown in Fig. 7.6. As can be seen from Fig. 7.7, the differential non-linearity of corresponding channels from different TDC-FPGAs is not identical either. This behaviour might be explained by local process variations inside the FPGA fabrics.

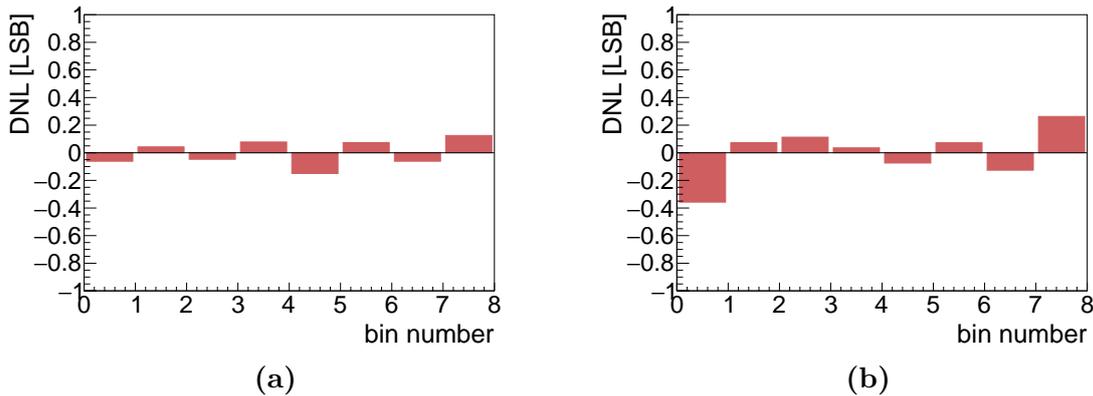


Figure 7.6: Channel with the smallest (a) and largest (b) differential non-linearity.

Examining the transfer characteristic in relation to the routing skew of both the input signals and the clock lines, the median of the differential non-linearities was calculated for every channel. The distribution of the median values is illustrated in Fig. 7.7. The channel average is 11.6 % LSB, corresponding to 47 ps. This result is in good agreement with the effective arrival time skew of 54 ps retrieved from a static timing analysis (cf. Sec. 6.1.4).

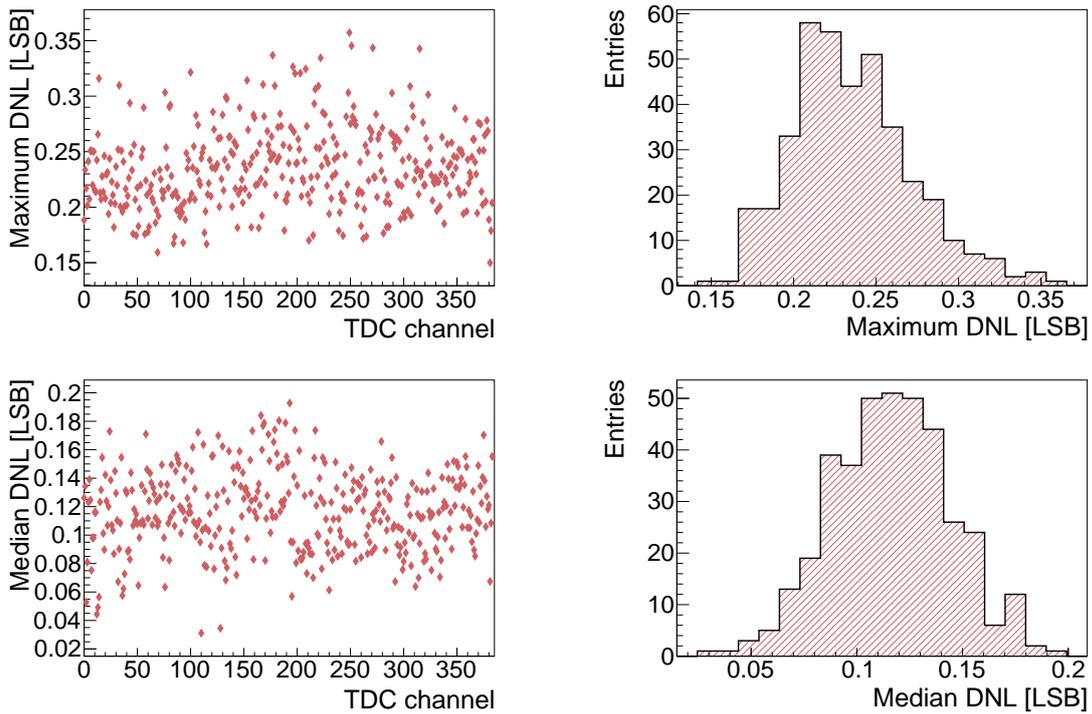


Figure 7.7: Upper left: Maximum differential non-linearity by channel. Upper right: Projection of the upper left distribution. Lower left: Median of the differential non-linearities by channel. Lower right: Projection of the lower left distribution.

Once the differential non-linearity is known, the integral non-linearity after the j th bin can be derived from the summation of the differential non-linearities:

$$\text{INL}_j = \sum_{i=0}^j \text{DNL}_i, \quad j = 0, 1, \dots, 7. \quad (7.6)$$

Figure 7.8 shows histograms of the integral non-linearity calculated according to Eq. (7.6) for the channel with the smallest and largest deviation from the ideal transfer function.

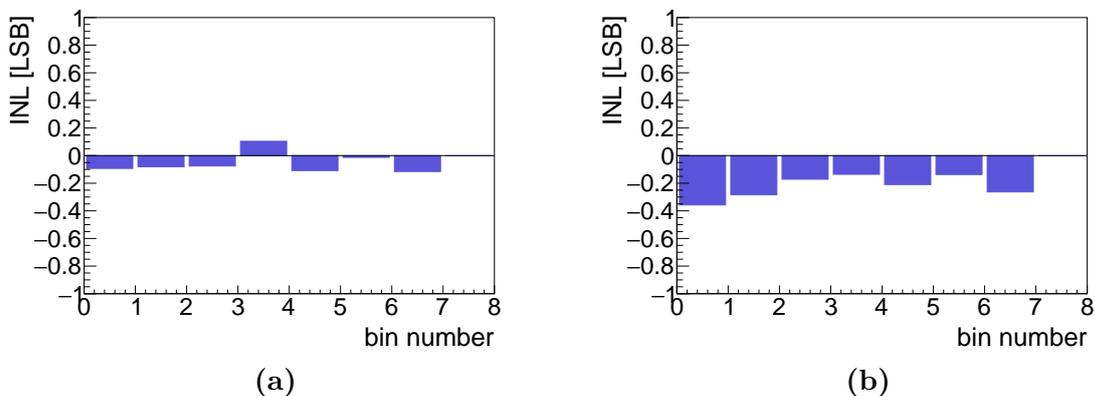


Figure 7.8: Channel with the smallest (a) and largest (b) integral non-linearity.

The distribution of the maximum values is illustrated by the upper plots in Fig. 7.9. The lower plots in Fig. 7.9 depict the standard deviation of the integral non-linearities. This characteristic quantity allows for an estimate of the expected time resolution (see Sec. 7.3.2.2). The results of the measurements discussed in this section are outlined in Tab. 7.3.

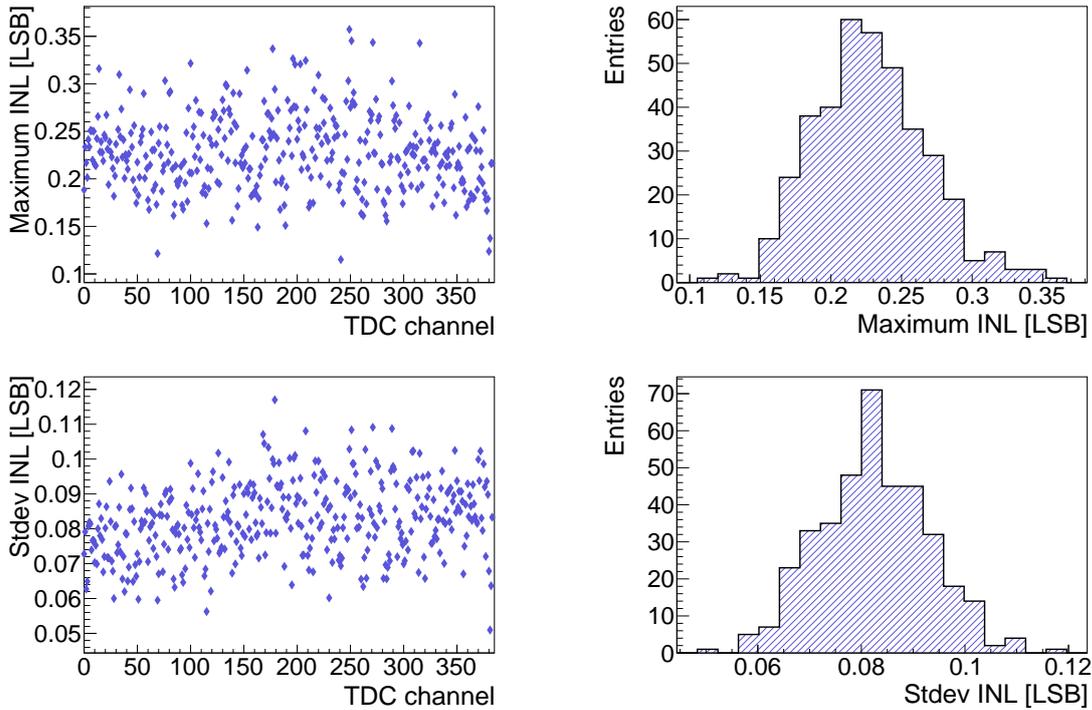


Figure 7.9: Upper left: Maximum integral non-linearity by channel. Upper right: Projection of the upper left distribution. Lower left: Standard deviation of the integral non-linearities by channel. Lower right: Projection of the lower left distribution.

The previous measurements were repeated several times in order to investigate the variance of the transfer characteristic under repeated system initializations. After each iteration, the differential non-linearity was determined and the TDC-FPGAs were reprogrammed to detect for instance potential phase variations of the MMCM clock outputs. It was found that the standard deviation of the sample, calculated for every bin of specific channels is negligibly small, not larger than 0.8% LSB. With the knowledge that the transfer characteristic is virtually constant, the integral non-linearities shown in Fig. 7.8 can be used to improve the accuracy of the TDC measurements (see e.g. Ref. [126]).

To correct the measurement values, correction charts may be stored for each channel in on-chip memory. However, depending on the resolution of the correction values, a substantial reduction of the effective dynamic range is provoked by this linearization technique. In consideration of the comparably small deviations from the ideal transfer function, it was not expected to gain a significant improvement of the single-shot precision.

Table 7.3: Results of the non-linearity measurements calculated from a sample of 53×10^3 hits per channel. The nominal size of LSB is 402 ps.

Linearity error	LSB	ps
max (DNL _i) channel minimum	15.0 %	60.3
max (DNL _i) channel maximum	35.7 %	143.5
max (DNL _i) channel average	23.5 %	94.5
median (DNL _i) channel average	11.6 %	46.6
max (INL _j) channel minimum	11.5 %	46.2
max (INL _j) channel maximum	35.7 %	143.5
max (INL _j) channel average	22.8 %	91.7
σ_{INL} channel average	8.2 %	33.0

7.3.2.2 Time Resolution

The studies presented in the following assess the time resolution, also known as single-shot precision, of time interval measurements. Furthermore, it was investigated how cross-talk induced noise and clock jitter affect the time resolution. The long-term stability and the consistency of the measurements with respect to repeated system startups were additional factors in the characterization of the TDC platform.

Time intervals refer to the time difference between hits and are calculated off-chip from the recorded timestamps. The applications performed with the ARAGORN front-end at the COMPASS-II experiment demand for time interval measurements between channels of the same front-end or even different host boards. It must be considered that the single-shot precision depends on the measured interval Δt , or more precisely, on the fractional part of the quotient $\Delta t/\text{LSB}$. This fundamental relationship can be evaluated by taking a series of measurements of a constant time interval. The length of the interval is then iteratively swept over a range equal to the period of the reference clock.

In consideration of the large amount of channels processed by the ARAGORN front-end, it is not necessary to follow this approach. Instead, the input channels were provided with copies of the same test signal from the pulse generator. Time intervals were analyzed between neighbouring channels and for each channel, the standard deviation and the mean was extracted from the sample. Since the length of the signal lines on the PCB and inside the FPGA fabric are not matched, the measured intervals slightly vary dependent on propagation delays. In Fig. 7.10, the

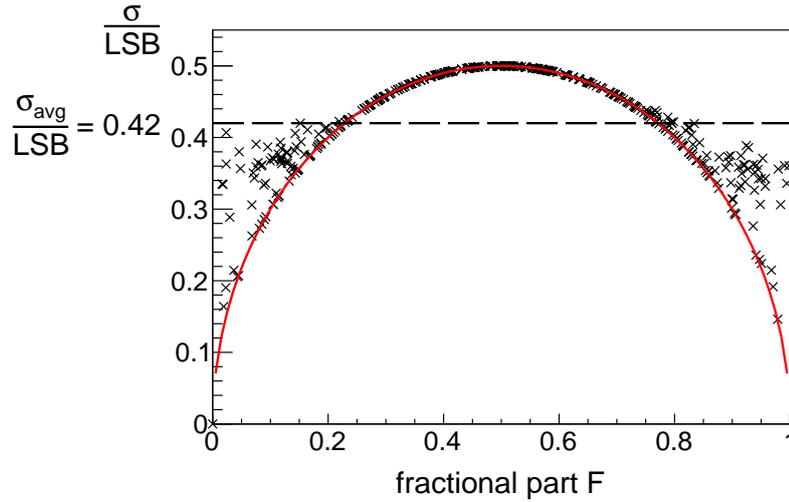


Figure 7.10: Normalized standard deviation dependent on the fractional part of the mean for time intervals measured between neighbouring channels of the ARAGORN front-end. The dashed line corresponds to the average time resolution. For comparison, the ideal case according to Eq. (5.4) is indicated by the red line.

standard deviation is plotted against the fractional part F of the mean value for all channels of the ARAGORN front-end. Around the maximum at $F = 0.5$, the results are in good agreement with the ideal case. Larger deviations are found in the outer regions where the imperfections of the transfer characteristic take effect. This finding has also been reported for other TDC designs (see e.g. Ref. [59, 127]). In fact, the values of the relative minima, observed for time intervals equal to integer multiples of LSB, are periodic with the period of the reference clock, as is the case for the non-linearity of the transfer function (see Ref. [127]). The average time resolution obtained using this method is 0.418 LSB, close to the theoretical resolution of 0.39 LSB of an ideal TDC given by Eq. (5.5).

To study the impact of cross-talk induced noise on the TDC application, the previous measurement was repeated with every second channel switched off. Thereafter, all outputs were enabled and every second channel was driven with an asynchronous aggressor signal. From a comparison of the average time resolution retrieved for the different configurations, it can be concluded that the jitter contribution from cross-talk is not larger than 0.022 LSB, corresponding to 8.8 ps.

Any measurement of the time resolution inevitably includes the time jitter inherent with the test signals received from the fanout boards. Bearing in mind that the integral non-linearity constitutes the deviation from the ideal transfer function, the time resolution of the TDC can be estimated with:

$$\sigma = \sqrt{\sigma_{\text{ideal}}^2 + \sigma_{\text{xtk}}^2 + \sigma_{\text{inp}}^2 + 2 \cdot \sigma_{\text{INL}}^2}, \quad (7.7)$$

where σ_{ideal} is the single-shot precision of the ideal TDC, σ_{xtk} is the cross-talk induced noise, σ_{inp} is the time jitter between the test inputs and σ_{INL} is the standard deviation of the integral non-linearities from Tab. 7.3. This quantity enters twice in Eq. (7.7) because time intervals are calculated from two measurement values. With these values and Eq. (7.7), σ_{inp} is determined to be 37.4 ps. Using external instruments,

the typical time jitter of the test inputs was measured to be 40 ps, which is in good agreement with the previous result taking into account the additive jitter from the measuring devices.

The previous investigations did not consider the jitter of the reference clock because the measured time intervals were shorter than the period of the reference clock. Examining the single-shot precision for time intervals exceeding the clock period, the fanout boards were programmed to provide neighbouring channels with independent test signals. These test inputs originated from the function generator with a delay of 2 μ s. The average time resolution obtained in this configuration is 0.419 LSB. Another measurement was conducted to study the correlations between multiple front-ends. Therefore, time intervals were analyzed between channels from different boards. With this setup, the average time resolution was measured to be 0.426 LSB. This result deviates by only 2.8 ps from the time resolution obtained for a single front-end.

The results of the measurements discussed in the section at hand are summarized in Tab. 7.4. The values presented take into consideration the time jitter inherent with the test signals.

Table 7.4: Average time resolution of all channels under different operating conditions.

Configuration	Time resolution	
	LSB	ps
Single board w/o cross-talk ($\Delta t < T_{\text{clk}}$)	40.7 %	163.6
Single board w/ cross-talk ($\Delta t < T_{\text{clk}}$)	40.8 %	164.0
Single board ($\Delta t > T_{\text{clk}}$)	40.9 %	164.4
Multiple boards ($\Delta t > T_{\text{clk}}$)	41.6 %	167.2

7.3.2.3 Time Interval Averaging

In Sec. 5.1 it was demonstrated that, if the number of measurements of a constant time interval is sufficiently large, the intrinsic precision of the TDC can be significantly improved by averaging the results. In regards to the accuracy of time interval averaging, the long-term and temperature stability and potential effects from repeated system startups have been examined.

Clock Phase Variations

Potential phase variations between the reference clocks of different ARAGORN front-ends have been thoroughly investigated with respect to subsequent initializations of the fiber optic network as follows. In a series of measurements of a constant time interval, the fiber optic cable interconnecting the master board with the ARWEN mezzanine card was iteratively unplugged and subsequently reconnected to force a reset of the GTP transceiver primitives. Following a reset, on every ARAGORN front-end in the star topology network, the jitter attenuator and as

well the MMCM primitives inside the TDC-FPGAs generating the multiphase clock re-lock. Time intervals were evaluated between corresponding channels of the master and a slave board. For each iteration, the mean of the sample was calculated. The change in the sample mean for the different iterations is 16.4 ps peak-to-peak. This measurement error can be attributed to phase variations between the reference clocks.

The constant-latency link, which has been described in Sec. 6.2.1, aims to transmit control signals with fixed latency to the receiver nodes. One of these signals synchronously resets all coarse counter primitives inside the TDC-FPGA design. Assuming that the latency through the link does not remain the same with respect to subsequent initializations, time jumps by integer multiples of the clock period would have been observed. This error mode was indeed never observed, neither in this exercise nor in a long-term measurement that will be discussed in the following paragraph.

Long-term Stability

To study the long-term stability, a constant time interval was measured for 74 hours between two ARAGORN front-ends. Evaluating changes in the TDC readings over time, the recorded hits data was subdivided into smaller samples comprising approximately 48×10^3 hits per channel, corresponding to periods of about 22 minutes. The standard deviation of the mean values observed for the different samples is 10.5 ps on average for all channels. Figure 7.11 shows the change in the mean over the measurement period for an exemplary channel. The rather systematic distribution gives rise to the assumption that the variations can be attributed to the phase alignment capability of the transmitter primitive that must be used if the transmit buffer is bypassed (cf. Sec. 6.2.1). This feature not only resolves the phase difference between the fabric clock and the parallel-to-serial converter block, but also continuously adjusts the phase of the parallel clock with respect to the serial transmitter clock in order to compensate for temperature and voltage variations (see Ref.[74]).

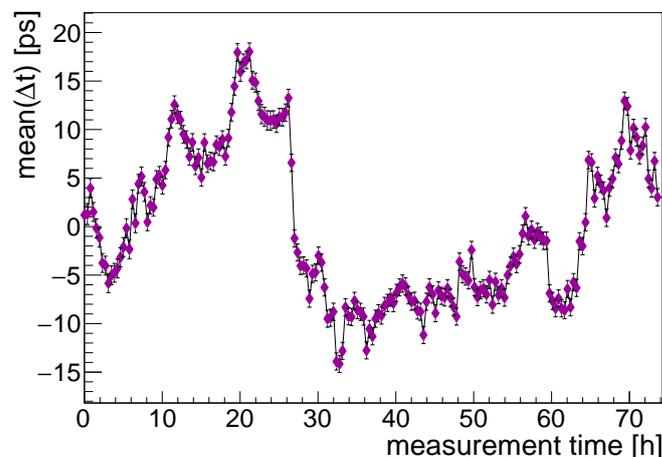


Figure 7.11: Relative change in the measured time interval mean over 74 h (exemplary channel).

Temperature Stability

The integrated MMCM primitives generating the multiphase clock are basically PLLs circuits providing some enhanced features. It is therefore expected that the design shows low dependence on ambient temperature or supply voltage variations. For this test, a constant time interval was measured repeatedly while the board temperature was increased in three steps from 35 °C to 50 °C by controlling the air flow around the board. The temperature was measured using the temperature sensor provided by the power rail supervisor chip on the ARAGORN front-end. Over the measured temperature range, a relative change of 3.2 ps was observed in the sample mean. From this result the estimated temperature drift is only 0.21 ps/°C.

7.3.2.4 Rate Capability

Every channel of the TDC-FPGA design comprises a 2k deep hit buffer RAM to withstand excessive input rates that may occur occasionally and for short periods of time. In this scenario, the input rate is limited only by the double-hit resolution of the design, that is in the order of the period of the reference clock. However, the average input rate that can be processed without running into buffer overflow conditions has a certain limit, briefly rate capability. The rate capability is determined by the amount of time required by the trigger matching algorithm to process entries in the hit buffer (cf. Sec. 6.1.2). If this limit is exceeded, at some point, the memory address supposed to store the next hit reaches the memory address the process is expected to start off for the next trigger event. In such a case, new hits are discarded until free storage is available again.

Integrated logic analyzer cores are provided by the Vivado Design Suite to probe selected internal signals of an FPGA design in real time. Evaluating the rate capability, a logic analyzer core was implemented to monitor the status flag indicating buffer overflows in the hit buffers. The input rate was then gradually increased until the assertion of the buffer full flag. This condition was observed for an input rate of 34 MHz per channel.

8. Summary

For recent measurements of the COMPASS-II physics programme, the RICH-1 detector was instrumented in 2016 with a set of new photon detectors based on a hybrid design combining micro-pattern gas detector architectures. The initial objective of this thesis was to develop the associated digital front-end electronics. It was foreseen to perform the analog readout with existing preamplifier/discriminator modules [116] of the COMPASS-II RICH-1 that are particularly customized to accommodate the demands for fast photon detection. However, during the work on this project, it turned out that the characteristic properties of the device pose specific problems concerning this application. With regards to the rather long current signals from the hybrid detector, it was observed that the device shows a ballistic deficit, the loss of pulse height at the output of the shaper, causing a poor signal-to-noise ratio. Here, the crux of matter lies in the 10 ns peaking time of the shaper. Therefore, it was decided to provisionally use an analog readout based on the APV25 chip [128] with a peaking time of 50 ns.

Nonetheless, the FPGA-based ARAGORN front-end provides the hardware platform for the implementation of high-performance time digitizers. Since its key interfaces solely connect to programmable logic devices, the board constitutes a versatile readout engine. The author of this thesis has conducted all steps of the development process, from conceptual design studies to system verification. In detail, this implies the following steps: the schematic diagram along with component selection, the PCB layout, the hardware commissioning and the development of the firmware designs.

Processing 384 input channels, four FPGAs on the ARAGORN front-end are operated with a time-to-digital converter firmware. Within the frame of this thesis, the performance metrics of the time-to-digital-converter have been thoroughly characterized. The principle results are summarized below. The time digitizers are sampled with a 311.04 MHz reference clock. This yields a time bin size of $\text{LSB} = 402 \text{ ps}$ and a dynamic range of $211 \text{ } \mu\text{s}$. With that in mind, the nominal time resolution is 156.8 ps . However, any real time digitizer reveals deviations from the ideal transfer function. For this design, the differential and integral non-linearities are both smaller than

36 % of LSB, compatible with the evaluated skew in the input and clock signal path to the interpolators. The time resolution was determined in distinct configurations. Without cross-talk and clock jitter, the single-shot precision is 163.6 ps. Considering the latter effects, the precision declines marginally to 164.0 ps and 164.4 ps, respectively. Evaluating time intervals between different ARAGORN front-ends, a time resolution of 167.2 ps is obtained. In view of these findings, it can be concluded that the time resolution is limited only by the integral non-linearity of the interpolators. Systematic effects on the precision of time interval averaging have also been examined. Long-term variations turned out to be in the order of 10.5 ps RMS. Owing to the fact that the multiphase sampling clock is obtained from a PLL-based design, the observed temperature drift is only about 0.21 ps/°C. Phase variations encountered in repeated system initializations between the reference clocks of multiple ARAGORN front-ends account for a measurement error of 16.4 ps peak-to-peak. The rate capability constitutes another benchmark parameter. Short-term peaks in the input rate are handled by the embedded memory resources of the FPGA fabrics, but the average processible input rate has of course a certain limit which is observed at 34 MHz.

The hardware features advanced optical interfaces to interconnect with up to seven boards as satellites. On this account, the central FPGA implements a constant-latency up-link for clock and trigger distribution which has successfully been verified concurrently with the commissioning of the time-to-digital converter. In the reverse direction, the firmware implements an event builder processing the received data for output via a single optical transceiver module. In bidirectional data transfers up to 6.2208 Gbit/s through the optical network, a bit error ratio better than 10^{-14} has been verified for a confidence level of 99 %.

Special attention has been paid to the development process that the ARAGORN front-end can be remotely controlled and configured without manual user intervention. A rich software toolset is provided for configuration bitstream and settings file download.

In light of the final results it can be claimed that this work has resulted in a fully operational time digitizer and readout engine covering a wide spectrum of applications in high-energy physics.

A. PCB Layout

Figures A.1 and A.2 show the top and bottom layers of the ARAGORN front-end for locating the components detailed in the thesis at hand by their reference designators. Figure A.3 depicts the cross section of the multilayer PCB.

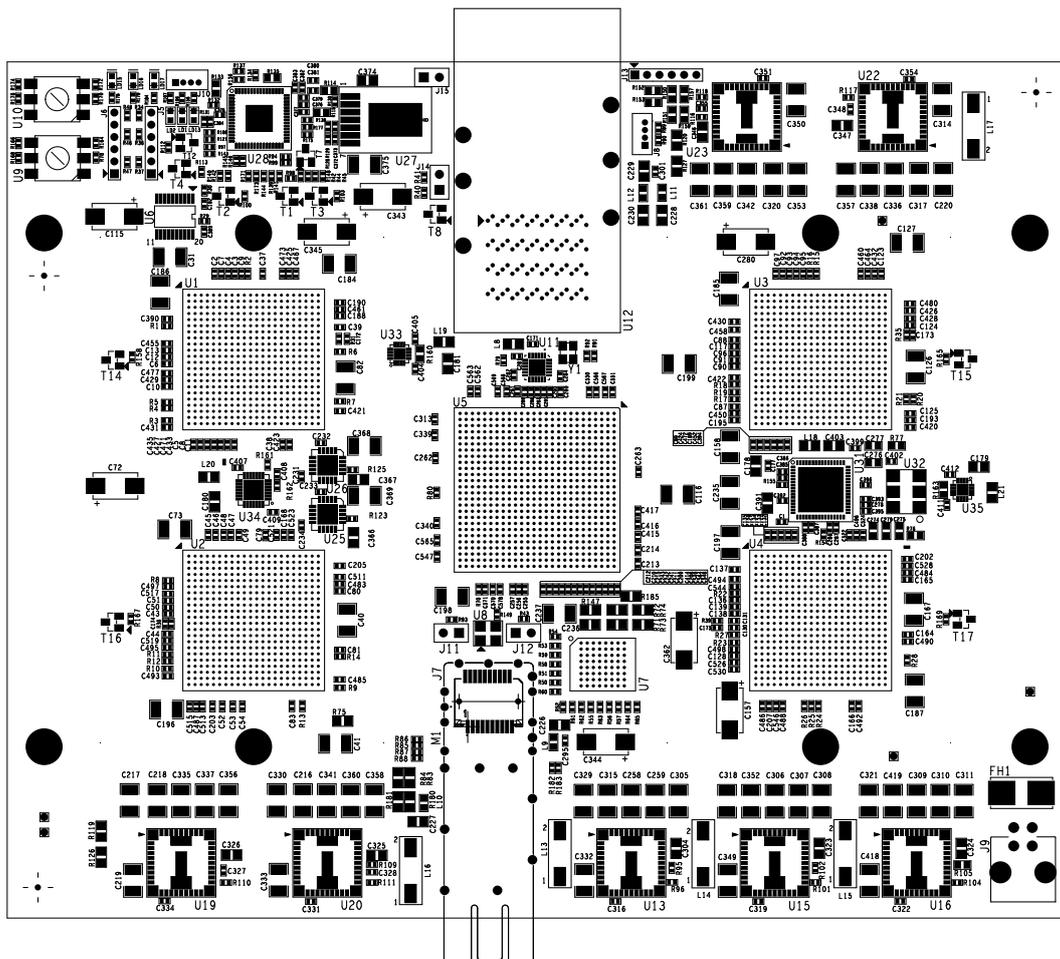


Figure A.1: Top view of the ARAGORN front-end.

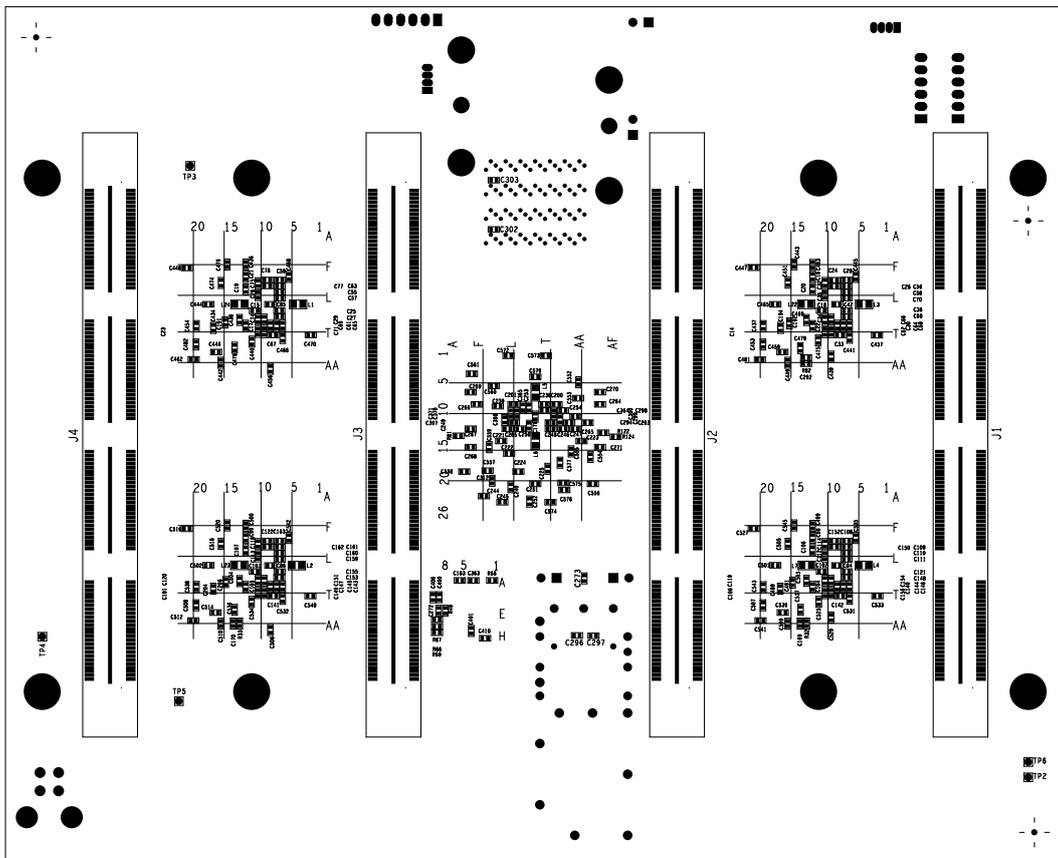


Figure A.2: Bottom view of the ARAGORN front-end.

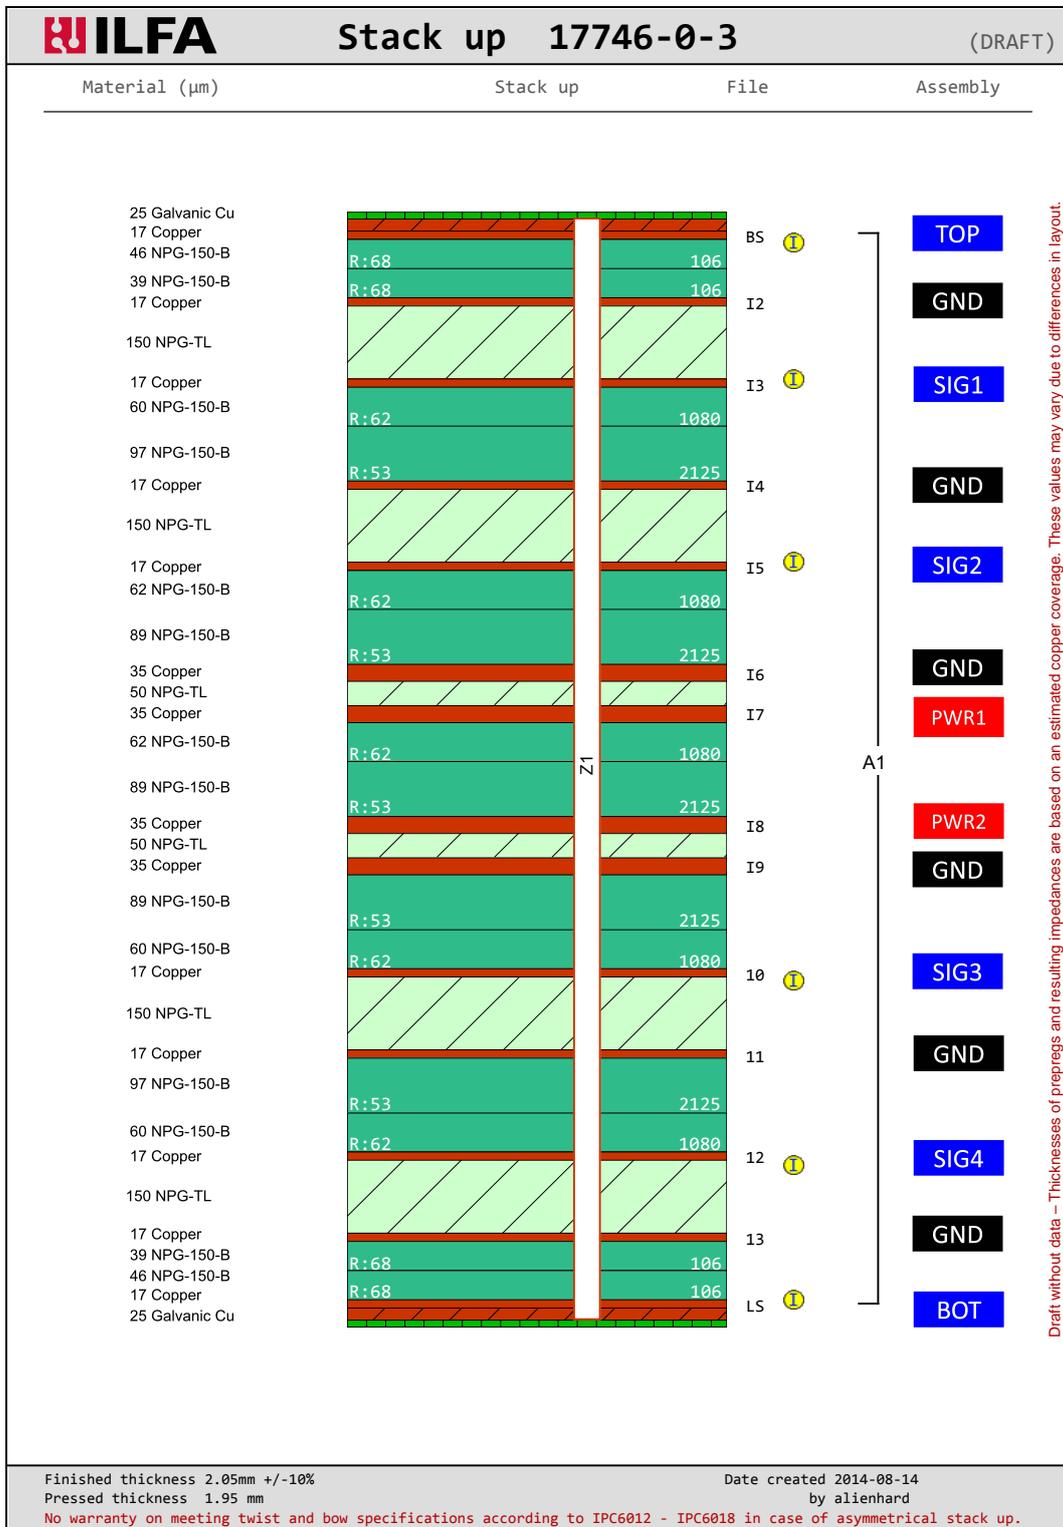


Figure A.3: Multi-layer PCB stack-up of the ARAGORN front-end.

B. Connector Pinouts

The following tables show the connectivity of the extension board connectors (J1 – J4). The differential TDC input signals are indicated with capital letters. The remaining pins connect to the MERGER-FPGA that applies 3.3V single-ended formats. Each of the four connector pin banks features an integral ground plane (pins 209 – 212).

Table B.1: Extension connector (J1) pinout.

Pin		Signal Name	FPGA	Pin	
+	-			+	-
5	7	TDC0_BANK0<0>	U1	F3	E3
9	11	TDC0_BANK0<1>	U1	C2	B2
13	15	TDC0_BANK0<2>	U1	B1	A1
17	19	TDC0_BANK0<3>	U1	E1	D1
21	23	TDC0_BANK0<4>	U1	H2	G2
25	27	TDC0_BANK0<5>	U1	J5	H5
29	31	TDC0_BANK0<6>	U1	K1	J1
33	35	TDC0_BANK0<7>	U1	L5	L4
37	39	TDC0_BANK0<8>	U1	M1	L1
41	43	TDC0_BANK0<9>	U1	N4	N3
45	47	TDC0_BANK0<10>	U1	P2	N2
49	51	TDC0_BANK0<11>	U1	R1	P1
6	8	TDC0_BANK0<12>	U1	K6	J6
10	12	TDC0_BANK0<13>	U1	H4	G4
14	16	TDC0_BANK0<14>	U1	E2	D2
18	20	TDC0_BANK0<15>	U1	G1	F1
22	24	TDC0_BANK0<16>	U1	H3	G3
26	28	TDC0_BANK0<17>	U1	K4	J4
30	32	TDC0_BANK0<18>	U1	K2	J2
34	36	TDC0_BANK0<19>	U1	L3	K3
38	40	TDC0_BANK0<20>	U1	M6	M5
42	44	TDC0_BANK0<21>	U1	M3	M2
46	48	TDC0_BANK0<22>	U1	P6	N5

continued on next page

50	52	TDC0_BANK0<23>	U1	P5	P4
53	55	TDC0_BANK1<0>	U1	R6	T6
57	59	TDC0_BANK1<1>	U1	T5	U5
61	63	TDC0_BANK1<2>	U1	U3	V3
65	67	TDC0_BANK1<3>	U1	V4	W4
69	71	TDC0_BANK1<4>	U1	W2	Y2
73	75	TDC0_BANK1<5>	U1	Y4	AA4
77	79	TDC0_BANK1<6>	U1	AA1	AB1
81	83	TDC0_BANK1<7>	U1	AB3	AB2
85	87	TDC0_BANK1<8>	U1	U6	V5
89	91	TDC0_BANK1<9>	U1	AB7	AB6
93	95	TDC0_BANK1<10>	U1	AA8	AB8
97	99	TDC0_BANK1<11>	U1	W9	Y9
54	56	TDC0_BANK1<12>	U1	R3	R2
58	60	TDC0_BANK1<13>	U1	R4	T4
62	64	TDC0_BANK1<14>	U1	T1	U1
66	68	TDC0_BANK1<15>	U1	U2	V2
70	72	TDC0_BANK1<16>	U1	V7	W7
74	76	TDC0_BANK1<17>	U1	W1	Y1
78	80	TDC0_BANK1<18>	U1	Y3	AA3
82	84	TDC0_BANK1<19>	U1	AA5	AB5
86	88	TDC0_BANK1<20>	U1	W6	W5
90	92	TDC0_BANK1<21>	U1	Y6	AA6
94	96	TDC0_BANK1<22>	U1	Y8	Y7
98	100	TDC0_BANK1<23>	U1	V9	V8
109	111	TDC1_BANK0<0>	U2	F3	E3
113	115	TDC1_BANK0<1>	U2	C2	B2
117	119	TDC1_BANK0<2>	U2	B1	A1
121	123	TDC1_BANK0<3>	U2	E1	D1
125	127	TDC1_BANK0<4>	U2	H2	G2
129	131	TDC1_BANK0<5>	U2	J5	H5
133	135	TDC1_BANK0<6>	U2	K1	J1
137	139	TDC1_BANK0<7>	U2	L5	L4
141	143	TDC1_BANK0<8>	U2	M1	L1
145	147	TDC1_BANK0<9>	U2	N4	N3
149	151	TDC1_BANK0<10>	U2	P2	N2
153	155	TDC1_BANK0<11>	U2	R1	P1
110	112	TDC1_BANK0<12>	U2	K6	J6
114	116	TDC1_BANK0<13>	U2	H4	G4
118	120	TDC1_BANK0<14>	U2	E2	D2
122	124	TDC1_BANK0<15>	U2	G1	F1
126	128	TDC1_BANK0<16>	U2	H3	G3
130	132	TDC1_BANK0<17>	U2	K4	J4
134	136	TDC1_BANK0<18>	U2	K2	J2
138	140	TDC1_BANK0<19>	U2	L3	K3
142	144	TDC1_BANK0<20>	U2	M6	M5
146	148	TDC1_BANK0<21>	U2	M3	M2
150	152	TDC1_BANK0<22>	U2	P6	N5
154	156	TDC1_BANK0<23>	U2	P5	P4
157	159	TDC1_BANK1<0>	U2	R6	T6

continued on next page

161	163	TDC1_BANK1<1>	U2	T5	U5
165	167	TDC1_BANK1<2>	U2	U3	V3
169	171	TDC1_BANK1<3>	U2	V4	W4
173	175	TDC1_BANK1<4>	U2	W2	Y2
177	179	TDC1_BANK1<5>	U2	Y4	AA4
181	183	TDC1_BANK1<6>	U2	AA1	AB1
185	187	TDC1_BANK1<7>	U2	AB3	AB2
189	191	TDC1_BANK1<8>	U2	U6	V5
193	195	TDC1_BANK1<9>	U2	AB7	AB6
197	199	TDC1_BANK1<10>	U2	AA8	AB8
201	203	TDC1_BANK1<11>	U2	W9	Y9
158	160	TDC1_BANK1<12>	U2	R3	R2
162	164	TDC1_BANK1<13>	U2	R4	T4
166	168	TDC1_BANK1<14>	U2	T1	U1
170	172	TDC1_BANK1<15>	U2	U2	V2
174	176	TDC1_BANK1<16>	U2	V7	W7
178	180	TDC1_BANK1<17>	U2	W1	Y1
182	184	TDC1_BANK1<18>	U2	Y3	AA3
186	188	TDC1_BANK1<19>	U2	AA5	AB5
190	192	TDC1_BANK1<20>	U2	W6	W5
194	196	TDC1_BANK1<21>	U2	Y6	AA6
198	200	TDC1_BANK1<22>	U2	Y8	Y7
202	204	TDC1_BANK1<23>	U2	V9	V8
<i>single-ended section</i>					
1		mezz0_ctrl<0>	U5		R2
2		mezz0_ctrl<1>	U5		T3
3		mezz0_ctrl<2>	U5		P1
4		mezz0_ctrl<3>	U5		T4
101		mezz0_ctrl<4>	U5		U2
102		mezz0_ctrl<5>	U5		T2
103		mezz0_ctrl<6>	U5		U1
104		mezz0_ctrl<7>	U5		R1
105		mezz0_ctrl<8>	U5		AA25
106		mezz0_ctrl<9>	U5		W25
107		mezz0_ctrl<10>	U5		AB25
108		mezz0_ctrl<11>	U5		W24
205		mezz0_ctrl<12>	U5		Y26
206		mezz0_ctrl<13>	U5		V24
207		mezz0_ctrl<14>	U5		Y25
208		mezz0_ctrl<15>	U5		Y22
209		GND			
210		GND			
211		GND			
212		GND			

Table B.2: Extension connector (J2) pinout.

Pin		Signal Name	FPGA Pin	
+	-		+	-
5	7	TDC0_BANK2<0>	U1	C13 B13
9	11	TDC0_BANK2<1>	U1	E13 E14
13	15	TDC0_BANK2<2>	U1	D14 D15
17	19	TDC0_BANK2<3>	U1	B15 B16
21	23	TDC0_BANK2<4>	U1	E16 D16
25	27	TDC0_BANK2<5>	U1	B17 B18

continued on next page

29	31	TDC0_BANK2<6>	U1	B20	A20
33	35	TDC0_BANK2<7>	U1	C18	C19
37	39	TDC0_BANK2<8>	U1	E19	D19
41	43	TDC0_BANK2<9>	U1	E21	D21
45	47	TDC0_BANK2<10>	U1	F16	E17
49	51	TDC0_BANK2<11>	U1	F13	F14
6	8	TDC0_BANK2<12>	U1	A13	A14
10	12	TDC0_BANK2<13>	U1	C14	C15
14	16	TDC0_BANK2<14>	U1	A15	A16
18	20	TDC0_BANK2<15>	U1	D17	C17
22	24	TDC0_BANK2<16>	U1	A18	A19
26	28	TDC0_BANK2<17>	U1	B21	A21
30	32	TDC0_BANK2<18>	U1	C22	B22
34	36	TDC0_BANK2<19>	U1	D20	C20
38	40	TDC0_BANK2<20>	U1	E22	D22
42	44	TDC0_BANK2<21>	U1	F18	E18
46	48	TDC0_BANK2<22>	U1	F19	F20
50	52	TDC0_BANK2<23>	U1	G21	G22
53	55	TDC0_BANK3<0>	U1	H20	G20
57	59	TDC0_BANK3<1>	U1	J22	H22
61	63	TDC0_BANK3<2>	U1	J20	J21
65	67	TDC0_BANK3<3>	U1	K21	K22
69	71	TDC0_BANK3<4>	U1	L19	L20
73	75	TDC0_BANK3<5>	U1	M18	L18
77	79	TDC0_BANK3<6>	U1	N22	M22
81	83	TDC0_BANK3<7>	U1	G17	G18
85	87	TDC0_BANK3<8>	U1	G15	G16
89	91	TDC0_BANK3<9>	U1	M15	M16
93	95	TDC0_BANK3<10>	U1	L14	L15
97	99	TDC0_BANK3<11>	U1	M13	L13
54	56	TDC0_BANK3<12>	U1	J19	H19
58	60	TDC0_BANK3<13>	U1	H17	H18
62	64	TDC0_BANK3<14>	U1	K18	K19
66	68	TDC0_BANK3<15>	U1	M21	L21
70	72	TDC0_BANK3<16>	U1	N20	M20
74	76	TDC0_BANK3<17>	U1	N18	N19
78	80	TDC0_BANK3<18>	U1	K17	J17
82	84	TDC0_BANK3<19>	U1	L16	K16
86	88	TDC0_BANK3<20>	U1	J15	H15
90	92	TDC0_BANK3<21>	U1	J14	H14
94	96	TDC0_BANK3<22>	U1	K13	K14
98	100	TDC0_BANK3<23>	U1	H13	G13
109	111	TDC1_BANK2<0>	U2	C13	B13
113	115	TDC1_BANK2<1>	U2	E13	E14
117	119	TDC1_BANK2<2>	U2	D14	D15
121	123	TDC1_BANK2<3>	U2	B15	B16
125	127	TDC1_BANK2<4>	U2	E16	D16
129	131	TDC1_BANK2<5>	U2	B17	B18
133	135	TDC1_BANK2<6>	U2	B20	A20
137	139	TDC1_BANK2<7>	U2	C18	C19

continued on next page

141	143	TDC1_BANK2<8>	U2	E19	D19
145	147	TDC1_BANK2<9>	U2	E21	D21
149	151	TDC1_BANK2<10>	U2	F16	E17
153	155	TDC1_BANK2<11>	U2	F13	F14
110	112	TDC1_BANK2<12>	U2	A13	A14
114	116	TDC1_BANK2<13>	U2	C14	C15
118	120	TDC1_BANK2<14>	U2	A15	A16
122	124	TDC1_BANK2<15>	U2	D17	C17
126	128	TDC1_BANK2<16>	U2	A18	A19
130	132	TDC1_BANK2<17>	U2	B21	A21
134	136	TDC1_BANK2<18>	U2	C22	B22
138	140	TDC1_BANK2<19>	U2	D20	C20
142	144	TDC1_BANK2<20>	U2	E22	D22
146	148	TDC1_BANK2<21>	U2	F18	E18
150	152	TDC1_BANK2<22>	U2	F19	F20
154	156	TDC1_BANK2<23>	U2	G21	G22
157	159	TDC1_BANK3<0>	U2	H20	G20
161	163	TDC1_BANK3<1>	U2	J22	H22
165	167	TDC1_BANK3<2>	U2	J20	J21
169	171	TDC1_BANK3<3>	U2	K21	K22
173	175	TDC1_BANK3<4>	U2	L19	L20
177	179	TDC1_BANK3<5>	U2	M18	L18
181	183	TDC1_BANK3<6>	U2	N22	M22
185	187	TDC1_BANK3<7>	U2	G17	G18
189	191	TDC1_BANK3<8>	U2	G15	G16
193	195	TDC1_BANK3<9>	U2	M15	M16
197	199	TDC1_BANK3<10>	U2	L14	L15
201	203	TDC1_BANK3<11>	U2	M13	L13
158	160	TDC1_BANK3<12>	U2	J19	H19
162	164	TDC1_BANK3<13>	U2	H17	H18
166	168	TDC1_BANK3<14>	U2	K18	K19
170	172	TDC1_BANK3<15>	U2	M21	L21
174	176	TDC1_BANK3<16>	U2	N20	M20
178	180	TDC1_BANK3<17>	U2	N18	N19
182	184	TDC1_BANK3<18>	U2	K17	J17
186	188	TDC1_BANK3<19>	U2	L16	K16
190	192	TDC1_BANK3<20>	U2	J15	H15
194	196	TDC1_BANK3<21>	U2	J14	H14
198	200	TDC1_BANK3<22>	U2	K13	K14
202	204	TDC1_BANK3<23>	U2	H13	G13
<i>single-ended section</i>					
1		mezz1_ctrl1<0>	U5	P5	
2		mezz1_ctrl1<1>	U5	R5	
3		mezz1_ctrl1<2>	U5	P6	
4		mezz1_ctrl1<3>	U5	T5	
101		mezz1_ctrl1<4>	U5	U6	
102		mezz1_ctrl1<5>	U5	U5	
103		mezz1_ctrl1<6>	U5	R8	
104		mezz1_ctrl1<7>	U5	P8	
105		mezz1_ctrl1<0>	U5	AA22	
<i>continued on next page</i>					

106	mezz1_ctrl1<1>	U5	AA23
107	mezz1_ctrl1<2>	U5	AB24
108	mezz1_ctrl1<3>	U5	AC24
205	mezz1_ctrl1<4>	U5	W23
206	mezz1_ctrl1<5>	U5	Y23
207	mezz1_ctrl1<6>	U5	V23
208	mezz1_ctrl1<7>	U5	AA24
209	GND		
210	GND		
211	GND		
212	GND		

Table B.3: Extension connector (J3) pinout.

Pin		Signal Name	FPGA Pin	
+	-		+	-
5	7	TDC2_BANK0<0>	U3	F3 E3
9	11	TDC2_BANK0<1>	U3	C2 B2
13	15	TDC2_BANK0<2>	U3	B1 A1
17	19	TDC2_BANK0<3>	U3	E1 D1
21	23	TDC2_BANK0<4>	U3	H2 G2
25	27	TDC2_BANK0<5>	U3	J5 H5
29	31	TDC2_BANK0<6>	U3	K1 J1
33	35	TDC2_BANK0<7>	U3	L5 L4
37	39	TDC2_BANK0<8>	U3	M1 L1
41	43	TDC2_BANK0<9>	U3	N4 N3
45	47	TDC2_BANK0<10>	U3	P2 N2
49	51	TDC2_BANK0<11>	U3	R1 P1
6	8	TDC2_BANK0<12>	U3	K6 J6
10	12	TDC2_BANK0<13>	U3	H4 G4
14	16	TDC2_BANK0<14>	U3	E2 D2
18	20	TDC2_BANK0<15>	U3	G1 F1
22	24	TDC2_BANK0<16>	U3	H3 G3
26	28	TDC2_BANK0<17>	U3	K4 J4
30	32	TDC2_BANK0<18>	U3	K2 J2
34	36	TDC2_BANK0<19>	U3	L3 K3
38	40	TDC2_BANK0<20>	U3	M6 M5
42	44	TDC2_BANK0<21>	U3	M3 M2
46	48	TDC2_BANK0<22>	U3	P6 N5
50	52	TDC2_BANK0<23>	U3	P5 P4
53	55	TDC2_BANK1<0>	U3	R6 T6
57	59	TDC2_BANK1<1>	U3	T5 U5
61	63	TDC2_BANK1<2>	U3	U3 V3
65	67	TDC2_BANK1<3>	U3	V4 W4
69	71	TDC2_BANK1<4>	U3	W2 Y2
73	75	TDC2_BANK1<5>	U3	Y4 AA4
77	79	TDC2_BANK1<6>	U3	AA1 AB1
81	83	TDC2_BANK1<7>	U3	AB3 AB2
85	87	TDC2_BANK1<8>	U3	U6 V5
89	91	TDC2_BANK1<9>	U3	AB7 AB6
93	95	TDC2_BANK1<10>	U3	AA8 AB8
97	99	TDC2_BANK1<11>	U3	W9 Y9
54	56	TDC2_BANK1<12>	U3	R3 R2

continued on next page

58	60	TDC2_BANK1<13>	U3	R4	T4
62	64	TDC2_BANK1<14>	U3	T1	U1
66	68	TDC2_BANK1<15>	U3	U2	V2
70	72	TDC2_BANK1<16>	U3	V7	W7
74	76	TDC2_BANK1<17>	U3	W1	Y1
78	80	TDC2_BANK1<18>	U3	Y3	AA3
82	84	TDC2_BANK1<19>	U3	AA5	AB5
86	88	TDC2_BANK1<20>	U3	W6	W5
90	92	TDC2_BANK1<21>	U3	Y6	AA6
94	96	TDC2_BANK1<22>	U3	Y8	Y7
98	100	TDC2_BANK1<23>	U3	V9	V8
109	111	TDC3_BANK0<0>	U4	F3	E3
113	115	TDC3_BANK0<1>	U4	C2	B2
117	119	TDC3_BANK0<2>	U4	B1	A1
121	123	TDC3_BANK0<3>	U4	E1	D1
125	127	TDC3_BANK0<4>	U4	H2	G2
129	131	TDC3_BANK0<5>	U4	J5	H5
133	135	TDC3_BANK0<6>	U4	K1	J1
137	139	TDC3_BANK0<7>	U4	L5	L4
141	143	TDC3_BANK0<8>	U4	M1	L1
145	147	TDC3_BANK0<9>	U4	N4	N3
149	151	TDC3_BANK0<10>	U4	P2	N2
153	155	TDC3_BANK0<11>	U4	R1	P1
110	112	TDC3_BANK0<12>	U4	K6	J6
114	116	TDC3_BANK0<13>	U4	H4	G4
118	120	TDC3_BANK0<14>	U4	E2	D2
122	124	TDC3_BANK0<15>	U4	G1	F1
126	128	TDC3_BANK0<16>	U4	H3	G3
130	132	TDC3_BANK0<17>	U4	K4	J4
134	136	TDC3_BANK0<18>	U4	K2	J2
138	140	TDC3_BANK0<19>	U4	L3	K3
142	144	TDC3_BANK0<20>	U4	M6	M5
146	148	TDC3_BANK0<21>	U4	M3	M2
150	152	TDC3_BANK0<22>	U4	P6	N5
154	156	TDC3_BANK0<23>	U4	P5	P4
157	159	TDC3_BANK1<0>	U4	R6	T6
161	163	TDC3_BANK1<1>	U4	T5	U5
165	167	TDC3_BANK1<2>	U4	U3	V3
169	171	TDC3_BANK1<3>	U4	V4	W4
173	175	TDC3_BANK1<4>	U4	W2	Y2
177	179	TDC3_BANK1<5>	U4	Y4	AA4
181	183	TDC3_BANK1<6>	U4	AA1	AB1
185	187	TDC3_BANK1<7>	U4	AB3	AB2
189	191	TDC3_BANK1<8>	U4	U6	V5
193	195	TDC3_BANK1<9>	U4	AB7	AB6
197	199	TDC3_BANK1<10>	U4	AA8	AB8
201	203	TDC3_BANK1<11>	U4	W9	Y9
158	160	TDC3_BANK1<12>	U4	R3	R2
162	164	TDC3_BANK1<13>	U4	R4	T4
166	168	TDC3_BANK1<14>	U4	T1	U1

continued on next page

170	172	TDC3_BANK1<15>	U4	U2	V2
174	176	TDC3_BANK1<16>	U4	V7	W7
178	180	TDC3_BANK1<17>	U4	W1	Y1
182	184	TDC3_BANK1<18>	U4	Y3	AA3
186	188	TDC3_BANK1<19>	U4	AA5	AB5
190	192	TDC3_BANK1<20>	U4	W6	W5
194	196	TDC3_BANK1<21>	U4	Y6	AA6
198	200	TDC3_BANK1<22>	U4	Y8	Y7
202	204	TDC3_BANK1<23>	U4	V9	V8
<i>single-ended section</i>					
1		mezz2_ctrl0<0>	U5		H2
2		mezz2_ctrl0<1>	U5		K5
3		mezz2_ctrl0<2>	U5		H1
4		mezz2_ctrl0<3>	U5		L4
101		mezz2_ctrl0<4>	U5		N7
102		mezz2_ctrl0<5>	U5		M6
103		mezz2_ctrl0<6>	U5		N6
104		mezz2_ctrl0<7>	U5		M5
105		mezz2_ctrl0<0>	U5		T14
106		mezz2_ctrl0<1>	U5		T15
107		mezz2_ctrl0<2>	U5		U15
108		mezz2_ctrl0<3>	U5		U16
205		mezz2_ctrl0<4>	U5		V18
206		mezz2_ctrl0<5>	U5		T18
207		mezz2_ctrl0<6>	U5		W18
208		mezz2_ctrl0<7>	U5		T17
209		GND			
210		GND			
211		GND			
212		GND			

Table B.4: Extension connector (J4) pinout.

Pin		Signal Name	FPGA Pin	
+	-		+	-
5	7	TDC2_BANK2<0>	U3	C13 B13
9	11	TDC2_BANK2<1>	U3	E13 E14
13	15	TDC2_BANK2<2>	U3	D14 D15
17	19	TDC2_BANK2<3>	U3	B15 B16
21	23	TDC2_BANK2<4>	U3	E16 D16
25	27	TDC2_BANK2<5>	U3	B17 B18
29	31	TDC2_BANK2<6>	U3	B20 A20
33	35	TDC2_BANK2<7>	U3	C18 C19
37	39	TDC2_BANK2<8>	U3	E19 D19
41	43	TDC2_BANK2<9>	U3	E21 D21
45	47	TDC2_BANK2<10>	U3	F16 E17
49	51	TDC2_BANK2<11>	U3	F13 F14
6	8	TDC2_BANK2<12>	U3	A13 A14
10	12	TDC2_BANK2<13>	U3	C14 C15
14	16	TDC2_BANK2<14>	U3	A15 A16
18	20	TDC2_BANK2<15>	U3	D17 C17
22	24	TDC2_BANK2<16>	U3	A18 A19
26	28	TDC2_BANK2<17>	U3	B21 A21
30	32	TDC2_BANK2<18>	U3	C22 B22
<i>continued on next page</i>				

34	36	TDC2_BANK2<19>	U3	D20	C20
38	40	TDC2_BANK2<20>	U3	E22	D22
42	44	TDC2_BANK2<21>	U3	F18	E18
46	48	TDC2_BANK2<22>	U3	F19	F20
50	52	TDC2_BANK2<23>	U3	G21	G22
53	55	TDC2_BANK3<0>	U3	H20	G20
57	59	TDC2_BANK3<1>	U3	J22	H22
61	63	TDC2_BANK3<2>	U3	J20	J21
65	67	TDC2_BANK3<3>	U3	K21	K22
69	71	TDC2_BANK3<4>	U3	L19	L20
73	75	TDC2_BANK3<5>	U3	M18	L18
77	79	TDC2_BANK3<6>	U3	N22	M22
81	83	TDC2_BANK3<7>	U3	G17	G18
85	87	TDC2_BANK3<8>	U3	G15	G16
89	91	TDC2_BANK3<9>	U3	M15	M16
93	95	TDC2_BANK3<10>	U3	L14	L15
97	99	TDC2_BANK3<11>	U3	M13	L13
54	56	TDC2_BANK3<12>	U3	J19	H19
58	60	TDC2_BANK3<13>	U3	H17	H18
62	64	TDC2_BANK3<14>	U3	K18	K19
66	68	TDC2_BANK3<15>	U3	M21	L21
70	72	TDC2_BANK3<16>	U3	N20	M20
74	76	TDC2_BANK3<17>	U3	N18	N19
78	80	TDC2_BANK3<18>	U3	K17	J17
82	84	TDC2_BANK3<19>	U3	L16	K16
86	88	TDC2_BANK3<20>	U3	J15	H15
90	92	TDC2_BANK3<21>	U3	J14	H14
94	96	TDC2_BANK3<22>	U3	K13	K14
98	100	TDC2_BANK3<23>	U3	H13	G13
109	111	TDC3_BANK2<0>	U4	C13	B13
113	115	TDC3_BANK2<1>	U4	E13	E14
117	119	TDC3_BANK2<2>	U4	D14	D15
121	123	TDC3_BANK2<3>	U4	B15	B16
125	127	TDC3_BANK2<4>	U4	E16	D16
129	131	TDC3_BANK2<5>	U4	B17	B18
133	135	TDC3_BANK2<6>	U4	B20	A20
137	139	TDC3_BANK2<7>	U4	C18	C19
141	143	TDC3_BANK2<8>	U4	E19	D19
145	147	TDC3_BANK2<9>	U4	E21	D21
149	151	TDC3_BANK2<10>	U4	F16	E17
153	155	TDC3_BANK2<11>	U4	F13	F14
110	112	TDC3_BANK2<12>	U4	A13	A14
114	116	TDC3_BANK2<13>	U4	C14	C15
118	120	TDC3_BANK2<14>	U4	A15	A16
122	124	TDC3_BANK2<15>	U4	D17	C17
126	128	TDC3_BANK2<16>	U4	A18	A19
130	132	TDC3_BANK2<17>	U4	B21	A21
134	136	TDC3_BANK2<18>	U4	C22	B22
138	140	TDC3_BANK2<19>	U4	D20	C20
142	144	TDC3_BANK2<20>	U4	E22	D22

continued on next page

146	148	TDC3_BANK2<21>	U4	F18	E18
150	152	TDC3_BANK2<22>	U4	F19	F20
154	156	TDC3_BANK2<23>	U4	G21	G22
157	159	TDC3_BANK3<0>	U4	H20	G20
161	163	TDC3_BANK3<1>	U4	J22	H22
165	167	TDC3_BANK3<2>	U4	J20	J21
169	171	TDC3_BANK3<3>	U4	K21	K22
173	175	TDC3_BANK3<4>	U4	L19	L20
177	179	TDC3_BANK3<5>	U4	M18	L18
181	183	TDC3_BANK3<6>	U4	N22	M22
185	187	TDC3_BANK3<7>	U4	G17	G18
189	191	TDC3_BANK3<8>	U4	G15	G16
193	195	TDC3_BANK3<9>	U4	M15	M16
197	199	TDC3_BANK3<10>	U4	L14	L15
201	203	TDC3_BANK3<11>	U4	M13	L13
158	160	TDC3_BANK3<12>	U4	J19	H19
162	164	TDC3_BANK3<13>	U4	H17	H18
166	168	TDC3_BANK3<14>	U4	K18	K19
170	172	TDC3_BANK3<15>	U4	M21	L21
174	176	TDC3_BANK3<16>	U4	N20	M20
178	180	TDC3_BANK3<17>	U4	N18	N19
182	184	TDC3_BANK3<18>	U4	K17	J17
186	188	TDC3_BANK3<19>	U4	L16	K16
190	192	TDC3_BANK3<20>	U4	J15	H15
194	196	TDC3_BANK3<21>	U4	J14	H14
198	200	TDC3_BANK3<22>	U4	K13	K14
202	204	TDC3_BANK3<23>	U4	H13	G13
<i>single-ended section</i>					
1		mezz3_ctrl1<0>	U5		N1
2		mezz3_ctrl1<1>	U5		M1
3		mezz3_ctrl1<2>	U5		L5
4		mezz3_ctrl1<3>	U5		M4
101		mezz3_ctrl1<4>	U5		K2
102		mezz3_ctrl1<5>	U5		K1
103		mezz3_ctrl1<6>	U5		L3
104		mezz3_ctrl1<7>	U5		J1
105		mezz3_ctrl1<0>	U5		W20
106		mezz3_ctrl1<1>	U5		T20
107		mezz3_ctrl1<2>	U5		Y20
108		mezz3_ctrl1<3>	U5		U20
205		mezz3_ctrl1<4>	U5		V19
206		mezz3_ctrl1<5>	U5		U19
207		mezz3_ctrl1<6>	U5		W19
208		mezz3_ctrl1<7>	U5		T19
209		GND			
210		GND			
211		GND			
212		GND			

C. TDC Input Mapping

In compliance with the COMPASS online data format [65], the TDC inputs of the ARAGORN front-end are distinguished by port and channel identifiers. The default mapping of the extension board connectors is listed in Tab. C.1.

Table C.1: Mapping of the TDC inputs in the firmware design to the extension board connectors (J1 – J4).

Port	Channel		Signal Name		FPGA	Conn.
	from	to	from	to		
0	0	23	TDC0_BANK0<0>	<23>	U1	J1
0	24	31	TDC0_BANK1<0>	<7>	U1	J1
1	0	15	TDC0_BANK1<8>	<23>	U1	J1
1	16	31	TDC0_BANK2<0>	<15>	U1	J2
2	0	7	TDC0_BANK2<16>	<23>	U1	J2
2	8	31	TDC0_BANK3<0>	<23>	U1	J2
3	0	23	TDC1_BANK0<0>	<23>	U2	J1
3	24	31	TDC1_BANK1<0>	<7>	U2	J1
4	0	15	TDC1_BANK1<8>	<23>	U2	J1
4	16	31	TDC1_BANK2<0>	<15>	U2	J2
5	0	7	TDC1_BANK2<16>	<23>	U2	J2
5	8	31	TDC1_BANK3<0>	<23>	U2	J2
6	0	23	TDC2_BANK0<0>	<23>	U3	J3
6	24	31	TDC2_BANK1<0>	<7>	U3	J3
7	0	15	TDC2_BANK1<8>	<23>	U3	J3
7	16	31	TDC2_BANK2<0>	<15>	U3	J4
8	0	7	TDC2_BANK2<16>	<23>	U3	J4
8	8	31	TDC2_BANK3<0>	<23>	U3	J4
9	0	23	TDC3_BANK0<0>	<23>	U4	J3
9	24	31	TDC3_BANK1<0>	<7>	U4	J3
10	0	15	TDC3_BANK1<8>	<23>	U4	J3
10	16	31	TDC3_BANK2<0>	<15>	U4	J4
11	0	7	TDC3_BANK2<16>	<23>	U4	J4
11	8	31	TDC3_BANK3<0>	<23>	U4	J4

Bibliography

- [1] H. Geiger and E. Marsden, “On a Diffuse Reflection of the α -Particles”, *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **82** 495–500 (1909), [doi:10.1098/rspa.1909.0054](https://doi.org/10.1098/rspa.1909.0054)
- [2] E. Rutherford, “LXXIX. The scattering of α and β particles by matter and the structure of the atom”, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **21** 669–688 (1911), [doi:10.1080/14786440508637080](https://doi.org/10.1080/14786440508637080)
- [3] R. Frisch and O. Stern, “Über die magnetische Ablenkung von Wasserstoffmolekülen und das magnetische Moment des Protons. I”, *Zeitschrift für Physik* **85** 4–16 (1933), [doi:10.1007/BF01330773](https://doi.org/10.1007/BF01330773)
- [4] L. W. Alvarez and F. Bloch, “A Quantitative Determination of the Neutron Moment in Absolute Nuclear Magnetons”, *Phys. Rev.* **57** 111–122 (1940), [doi:10.1103/PhysRev.57.111](https://doi.org/10.1103/PhysRev.57.111)
- [5] R. P. Feynman, “Very High-Energy Collisions of Hadrons”, *Phys. Rev. Lett.* **23** 1415–1417 (1969), [doi:10.1103/PhysRevLett.23.1415](https://doi.org/10.1103/PhysRevLett.23.1415)
- [6] J. D. Bjorken and E. A. Paschos, “Inelastic Electron-Proton and γ -Proton Scattering and the Structure of the Nucleon”, *Phys. Rev.* **185** 1975–1982 (1969), [doi:10.1103/PhysRev.185.1975](https://doi.org/10.1103/PhysRev.185.1975)
- [7] **COMPASS** Collaboration, C. Adolph *et al.*, “The spin structure function g_1^p of the proton and a test of the Bjorken sum rule”, *Phys. Lett. B* **753** 18 – 28 (2016), [doi:10.1016/j.physletb.2015.11.064](https://doi.org/10.1016/j.physletb.2015.11.064)
- [8] **COMPASS** Collaboration, C. Adolph *et al.*, “Final COMPASS results on the deuteron spin-dependent structure function g_1^d and the Bjorken sum rule”, (2016), [arXiv:1612.00620](https://arxiv.org/abs/1612.00620)
- [9] **HERMES** Collaboration, A. Airapetian *et al.*, “Precise determination of the spin structure function g_1 of the proton, deuteron, and neutron”, *Phys. Rev. D* **75** 012007 (2007), [doi:10.1103/PhysRevD.75.012007](https://doi.org/10.1103/PhysRevD.75.012007)
- [10] D. de Florian *et al.*, “Evidence for Polarization of Gluons in the Proton”, *Phys. Rev. Lett.* **113** 012001 (2014), [doi:10.1103/PhysRevLett.113.012001](https://doi.org/10.1103/PhysRevLett.113.012001)
- [11] **COMPASS** Collaboration, C. Adolph *et al.*, “Leading-order determination of the gluon polarisation using a novel method”, (2015), [arXiv:1512.05053](https://arxiv.org/abs/1512.05053)

- [12] **COMPASS** Collaboration, M. Alekseev *et al.*, “Quark helicity distributions from longitudinal spin asymmetries in muon–proton and muon–deuteron scattering”, *Phys. Lett. B* **693** 227 – 235 (2010), [doi:10.1016/j.physletb.2010.08.034](https://doi.org/10.1016/j.physletb.2010.08.034)
- [13] **COMPASS** Collaboration, F. Gautheron *et al.*, “COMPASS-II Proposal”, *CERN-SPSC-2010-014*, *SPSC-P-340* (2010)
- [14] A. Kusina *et al.*, “Strange quark parton distribution functions and implications for Drell-Yan boson production at the LHC”, *Phys. Rev. D* **85** 094028 (2012), [doi:10.1103/PhysRevD.85.094028](https://doi.org/10.1103/PhysRevD.85.094028)
- [15] A. Manohar, “An Introduction to spin dependent deep inelastic scattering”, *7th Lake Louise Winter Institute: Symmetry and Spin in the Standard Model* 1–46 (1992), [arXiv:hep-ph/9204208](https://arxiv.org/abs/hep-ph/9204208)
- [16] S. D. Bass, *The Spin Structure of the Proton*, World Scientific Publishing (2008)
- [17] C. Patrignani *et al.*, “Review of Particle Physics”, *Chin. Phys. C* **40** 100001 (2016), [doi:10.1088/1674-1137/40/10/100001](https://doi.org/10.1088/1674-1137/40/10/100001)
- [18] W. Panofsky, “Electromagnetic Interactions: Low q^2 Electrodynamics, Elastic and Inelastic Electron (and Muon) Scattering”, *Proceedings of the 14th Int. Conf. on High-Energy Physics, Vienna* (1968)
- [19] F. Halzen and A. D. Martin, *Quarks & Leptons: An Introductory Course in Modern Particle Physics*, John Wiley & Sons (1984)
- [20] G. Altarelli and G. Parisi, “Asymptotic freedom in parton language”, *Nuclear Physics B* **126** 298 – 318 (1977), [doi:10.1016/0550-3213\(77\)90384-4](https://doi.org/10.1016/0550-3213(77)90384-4)
- [21] V. Barone and P. G. Ratcliffe, *Transverse Spin Physics*, World Scientific, Singapore (2003)
- [22] B. Lampe and E. Reya, “Spin physics and polarized structure functions”, *Phys. Rept.* **332** 1 – 163 (2000), [doi:10.1016/S0370-1573\(99\)00100-3](https://doi.org/10.1016/S0370-1573(99)00100-3)
- [23] R. L. Jaffe and A. Manohar, “The g_1 problem: Deep inelastic electron scattering and the spin of the proton”, *Nucl. Phys. B* **337** 509 – 546 (1990), [doi:10.1016/0550-3213\(90\)90506-9](https://doi.org/10.1016/0550-3213(90)90506-9)
- [24] M. Anselmino *et al.*, “The Theory and phenomenology of polarized deep inelastic scattering”, *Phys. Rept.* **261** 1–124 (1995), [doi:10.1016/0370-1573\(95\)00011-5](https://doi.org/10.1016/0370-1573(95)00011-5), [Erratum: *Phys. Rept.* **281** 399(1997)]
- [25] J. R. Ellis and R. L. Jaffe, “A Sum Rule for Deep Inelastic Electroproduction from Polarized Protons”, *Phys. Rev. D* **9** 1444 (1974), [doi:10.1103/PhysRevD.10.1669.2](https://doi.org/10.1103/PhysRevD.10.1669.2), [10.1103/PhysRevD.9.1444](https://doi.org/10.1103/PhysRevD.9.1444), [Erratum: *Phys. Rev. D* **10** 1669 (1974)]

- [26] **COMPASS** Collaboration, V. Alexakhin *et al.*, “The deuteron spin-dependent structure function g_1^d and its first moment”, *Phys. Lett. B* **647** 8 – 17 (2007), [doi:10.1016/j.physletb.2006.12.076](https://doi.org/10.1016/j.physletb.2006.12.076)
- [27] **COMPASS** Collaboration, M. Alekseev *et al.*, “The polarised valence quark distribution from semi-inclusive DIS”, *Phys. Lett. B* **660** 458 – 465 (2008), [doi:10.1016/j.physletb.2007.12.056](https://doi.org/10.1016/j.physletb.2007.12.056)
- [28] **COMPASS** Collaboration, C. Adolph *et al.*, “Multiplicities of charged kaons from deep-inelastic muon scattering off an isoscalar target”, *Phys. Lett. B* **767** 133 – 141 (2017), [doi:10.1016/j.physletb.2017.01.053](https://doi.org/10.1016/j.physletb.2017.01.053)
- [29] D. de Florian *et al.*, “Global Analysis of Helicity Parton Densities and their Uncertainties”, *Phys. Rev. Lett.* **101** 072001 (2008), [doi:10.1103/PhysRevLett.101.072001](https://doi.org/10.1103/PhysRevLett.101.072001)
- [30] D. de Florian *et al.*, “Extraction of spin-dependent parton densities and their uncertainties”, *Phys. Rev. D* **80** 034030 (2009), [doi:10.1103/PhysRevD.80.034030](https://doi.org/10.1103/PhysRevD.80.034030)
- [31] **COMPASS** Collaboration, P. Abbon *et al.*, “Particle identification with COMPASS RICH-1”, *Nucl. Instrum. Meth. A* **631** 26 – 39 (2011), [doi:10.1016/j.nima.2010.11.106](https://doi.org/10.1016/j.nima.2010.11.106)
- [32] **HERMES** Collaboration, A. Airapetian *et al.*, “Measurement of parton distributions of strange quarks in the nucleon from charged-kaon production in deep-inelastic scattering on the deuteron”, *Phys. Lett. B* **666** 446 – 450 (2008), [doi:10.1016/j.physletb.2008.07.090](https://doi.org/10.1016/j.physletb.2008.07.090)
- [33] D. de Florian *et al.*, “Global analysis of fragmentation functions for pions and kaons and their uncertainties”, *Phys. Rev. D* **75** 114010 (2007), [doi:10.1103/PhysRevD.75.114010](https://doi.org/10.1103/PhysRevD.75.114010)
- [34] **COMPASS** Collaboration, P. Abbon *et al.*, “The COMPASS experiment at CERN”, *Nucl. Instrum. Meth. A* **577** 455–518 (2007), [doi:10.1016/j.nima.2007.03.026](https://doi.org/10.1016/j.nima.2007.03.026)
- [35] E. Bielert *et al.*, “A 2.5 m long liquid hydrogen target for COMPASS”, *Nucl. Instrum. Meth. A* **746** 20 – 25 (2014), [doi:10.1016/j.nima.2014.01.067](https://doi.org/10.1016/j.nima.2014.01.067)
- [36] T. Szameitat, “New Geant4-based Monte Carlo Software for the COMPASS-II Experiment at CERN”, Dissertation, Albert-Ludwigs-Universität Freiburg (2017), [doi:10.6094/UNIFR/11686](https://doi.org/10.6094/UNIFR/11686)
- [37] A. V. Belitsky *et al.*, “Theory of deeply virtual Compton scattering on the nucleon”, *Nucl. Phys. B* **629** 323–392 (2002), [doi:10.1016/S0550-3213\(02\)00144-X](https://doi.org/10.1016/S0550-3213(02)00144-X)
- [38] P. Jörg, “Deeply Virtual Compton Scattering at CERN - What is the Size of the Proton?”, Dissertation, Albert-Ludwigs-Universität Freiburg (2017), [doi:10.6094/UNIFR/12397](https://doi.org/10.6094/UNIFR/12397)

-
- [39] H. Kolanoski and N. Wermes, *Teilchendetektoren*, Springer Spektrum (2016)
- [40] C. Bernet *et al.*, “The COMPASS trigger system for muon scattering”, *Nucl. Instrum. Meth. A* **550** 217 – 240 (2005), [doi:10.1016/j.nima.2005.05.043](https://doi.org/10.1016/j.nima.2005.05.043)
- [41] H. Hoedlmoser *et al.*, “Long term performance and ageing of CsI photocathodes for the ALICE/HMPID detector”, *Nucl. Instr. and Meth. A* **574** 28 – 38 (2007), [doi:10.1016/j.nima.2007.01.101](https://doi.org/10.1016/j.nima.2007.01.101)
- [42] E. Albrecht *et al.*, “Status and characterisation of COMPASS RICH-1”, *Nucl. Instr. and Meth. A* **553** 215 – 219 (2005), [doi:10.1016/j.nima.2005.08.036](https://doi.org/10.1016/j.nima.2005.08.036)
- [43] M. Alexeev *et al.*, “THGEM-based photon detectors for the upgrade of COMPASS RICH-1”, *Nucl. Instr. and Meth. A* **732** 264 – 268 (2013), [doi:10.1016/j.nima.2013.08.020](https://doi.org/10.1016/j.nima.2013.08.020)
- [44] M. Alexeev *et al.*, “The MPGD-based photon detectors for the upgrade of COMPASS RICH-1”, *Nucl. Instr. and Meth. A* (2017), [doi:10.1016/j.nima.2017.02.013](https://doi.org/10.1016/j.nima.2017.02.013), In press
- [45] M. Alexeev *et al.*, “Status and progress of the novel photon detectors based on THGEM and hybrid MPGD architectures”, *Nucl. Instr. and Meth. A* **766** 133 – 137 (2014), [doi:10.1016/j.nima.2014.07.030](https://doi.org/10.1016/j.nima.2014.07.030)
- [46] M. Alexeev *et al.*, “Ion backflow in thick GEM-based detectors of single photons”, *Journal of Instrumentation* **8** P01021 (2013), [doi:10.1088/1748-0221/8/01/P01021](https://doi.org/10.1088/1748-0221/8/01/P01021)
- [47] M. Bodlak *et al.*, “FPGA based data acquisition system for COMPASS experiment”, *J. Phys. Conf. Ser.* **513** 012029 (2014), [doi:10.1088/1742-6596/513/1/012029](https://doi.org/10.1088/1742-6596/513/1/012029)
- [48] H. C. van der Bij *et al.*, “S-LINK, a data link interface specification for the LHC era”, *IEEE Transactions on Nuclear Science* **44** 398–402 (1997), [doi:10.1109/23.603679](https://doi.org/10.1109/23.603679)
- [49] F. Herrmann, “Development and Verification of a High Performance Electronic Readout Framework for High Energy Physics”, Dissertation, Albert-Ludwigs-Universität Freiburg (2011)
- [50] VMEbus International Trade Association (VITA), VME64 Standard (1994), ANSI/VITA 1.0-1994
- [51] W. D. Peterson, *The VMEbus Handbook*, VITA, 4th edition (1997)
- [52] VMEbus International Trade Association (VITA), VXS VMEbus Serial Standard (2006), ANSI/VITA 41.0-2006
- [53] Xilinx, Inc., Virtex-5 Family Overview, DS100
- [54] Xilinx, Inc., Aurora 8B/10B Protocol Specification, SP002

-
- [55] S. Schopferer, “Entwicklung eines hochauflösenden Transientenrekorders”, Diploma thesis, Albert-Ludwigs-Universität Freiburg (2009)
- [56] P. Jörg, “Untersuchung von Algorithmen zur Charakterisierung von Photomultiplierpulsen in Echtzeit”, Diploma thesis, Albert-Ludwigs-Universität Freiburg (2013)
- [57] S. Bartknecht *et al.*, “Development of a 1 GS/s high-resolution sampling ADC system”, *Nucl. Instrum. Meth. A* **623** 507 – 509 (2010), [doi:10.1016/j.nima.2010.03.052](https://doi.org/10.1016/j.nima.2010.03.052)
- [58] S. Bartknecht *et al.*, “Development and Performance Verification of the GANDALF High-Resolution Transient Recorder System”, *IEEE Transactions on Nuclear Science* **58** 1456–1459 (2011), [doi:10.1109/TNS.2011.2142195](https://doi.org/10.1109/TNS.2011.2142195)
- [59] M. Büchele, “Entwicklung eines FPGA-basierten 128-Kanal Time-to-Digital Converter für Teilchenphysik-Experimente”, Diploma thesis, Albert-Ludwigs-Universität Freiburg (2012)
- [60] M. Büchele *et al.*, “A 128-channel Time-to-Digital Converter (TDC) inside a Virtex-5 FPGA on the GANDALF module”, *Journal of Instrumentation* **7** C03008 (2012)
- [61] C. Michalski, “Entwicklung eines Echtzeit-Strahlprofil-Monitoring-Systems für das COMPASS-II Experiment”, Diploma thesis, Albert-Ludwigs-Universität Freiburg (2013)
- [62] T. Baumann *et al.*, “The GANDALF 128-channel Time-to-Digital Converter”, *Journal of Instrumentation* **8** C01016 (2013)
- [63] J. Bieling *et al.*, “Implementation of mean-timing and subsequent logic functions on an FPGA”, *Nucl. Instrum. Meth. A* **672** 13 – 20 (2012), [doi:10.1016/j.nima.2011.12.104](https://doi.org/10.1016/j.nima.2011.12.104)
- [64] C. Schill, Private communications
- [65] H. Fischer *et al.*, The COMPASS Online Data Format - Version 4, COMPASS note 2002-8 (2003)
- [66] T. Grussenmeyer, “Entwicklung eines modularen und verteilten Datenaufnahmesystems für Testexperimente”, Diploma thesis, Albert-Ludwigs-Universität Freiburg (2013)
- [67] M. Becker, “Serielle Takt- und Datenübertragung am COMPASS-Experiment”, Scientific paper, Albert-Ludwigs-Universität Freiburg (2014)
- [68] P. Ashenden, *Digital Design: An Embedded Systems Approach Using VHDL*, Morgan Kaufmann (2007)
- [69] Xilinx, Inc., 7 Series FPGAs Configurable Logic Block, UG474
- [70] Xilinx, Inc., 7 Series FPGAs Memory Resources, UG473

- [71] Xilinx, Inc., 7 Series DSP48E1 Slice, UG479
- [72] Xilinx, Inc., 7 Series FPGAs Clocking Resources, UG472
- [73] Xilinx, Inc., SelectIO Interface Wizard v5.1, PG070
- [74] Xilinx, Inc., 7 Series FPGAs GTP Transceivers, UG482
- [75] Xilinx, Inc., Xilinx Power Estimator (XPE), http://www.xilinx.com/products/design_tools/logic_design/xpe.htm
- [76] Texas Instruments, LMZ3-series Power Modules, Data sheets
- [77] Texas Instruments, TPS74201 Single Output LDO, Data sheet
- [78] Texas Instruments, UCD90120A Power Supply Sequencer and Monitor, Data sheet
- [79] Texas Instruments, Fusion Digital Power Designer Software, http://www.ti.com/tool/fusion_digital_power_designer
- [80] Xilinx, Inc., 7 Series FPGAs Configuration, UG470
- [81] Micron Technology, Inc., Micron StrataFlash Embedded Memory, Data sheet
- [82] Xilinx, Inc., BPI Fast Configuration and iMPACT Flash Programming with 7 Series FPGAs, XAPP587
- [83] Xilinx, Inc., Vivado Design Suite 7 Series FPGA and Zynq-7000 All Programmable SoC Libraries Guide, UG953
- [84] Silicon Laboratories, Si5338 I²C-Programmable Any-Frequency, Any-Output Quad Clock Generator, Data sheet
- [85] Silicon Laboratories, Si5335/38/51/56 ClockBuilder Desktop Software
- [86] Texas Instruments, LMK04906 Ultralow Noise Clock Jitter Cleaner and Multiplier With 6 Programmable Outputs, Data sheet
- [87] Texas Instruments, Clock Conditioner Owner's Manual, SNAA103
- [88] Texas Instruments, Clock Design Tool – Loop Filter & Device Configuration + Simulation, <http://www.ti.com/tool/clockdesigntool>
- [89] Texas Instruments, CodeLoader Software for Device Register Programming, <http://www.ti.com/tool/codeloader>
- [90] Xilinx, Inc., 7 Series FPGA GTX/GTH/GTP Transceivers – Reference clock phase noise masks, Answer record
- [91] SFF Committee, Specification for SFP (Small Formfactor Pluggable) Transceiver (2001), INF-8074i
- [92] InfiniBand, Supplement to Infiniband Architecture Specification (2009), Volume 2, Release 1.2.1

-
- [93] Finisar, 100GBASE-SR10 100 m CXP Optical Transceiver Module (2012), Product Specification
- [94] NXP Semiconductors, UM10204 – I²C-bus specification and user manual (2014), Rev. 6
- [95] System Management Interface Forum, PMBus Power System Management Protocol Specification (2010), Rev. 1.2
- [96] System Management Interface Forum, System Management Bus (SMBus) Specification (2014), Version 3.0
- [97] National Semiconductor, MICROWIRE Serial Interface (1992), AN-452
- [98] IEEE Computer Society, IEEE Standard for Test Access Port and Boundary-Scan Architecture (2013), IEEE Std 1149.1-2013
- [99] Xilinx, Inc., Vivado Design Suite User Guide – Programming and Debugging, UG908
- [100] Cadence Design Systems, Inc., Allegro PCB Design and Analysis, Release 16.6
- [101] E. Bogatin, Signal and Power Integrity – Simplified, Prentice Hall, 2nd edition (2010)
- [102] Polar Instruments, Si8000 PCB Controlled Impedance Field Solver, <http://www.polarinstruments.com/products/cits/Si8000.html>
- [103] Xilinx, Inc., 7 Series FPGAs PCB Design Guide, UG483
- [104] J. Kalisz, “Review of methods for time interval measurements with picosecond resolution”, *Metrologia* **41** 17–32 (2004)
- [105] Hewlett-Packard, Time interval averaging (1970), Application Note 162-1
- [106] S. Henzler, Time-to-Digital Converters, Springer Series in Advanced Microelectronics 29 (2010)
- [107] Xilinx, Inc., Vivado Design Suite User Guide – Using Constraints, UG903
- [108] G. Braun *et al.*, “TDC Chip and Readout Driver Developments for COMPASS and LHC-Experiments”, (1998), [arXiv:hep-ex/9810048](https://arxiv.org/abs/hep-ex/9810048)
- [109] Xilinx, Inc., Artix-7 FPGAs Data Sheet: DC and AC Switching Characteristics, DS181
- [110] Xilinx, Inc., Vivado Design Suite User Guide – Implementation, UG904
- [111] M. Bickel, “Charakterisierung einer Zeitmesseinheit für das Compass-Experiment am Cern”, Scientific paper, Albert-Ludwigs-Universität Freiburg (2015)

-
- [112] R. Giordano and A. Aloisio, “Fixed-Latency, Multi-Gigabit Serial Links With Xilinx FPGAs”, *IEEE Transactions on Nuclear Science* **58** 194–201 (2011), [doi:10.1109/TNS.2010.2101083](https://doi.org/10.1109/TNS.2010.2101083)
- [113] A. Athavale and C. Christensen, High-Speed Serial I/O Made Simple, Xilinx Connectivity Solutions, 1st edition (2005)
- [114] Xilinx, Inc., MicroBlaze Processor Reference Guide, UG984
- [115] Xilinx, Inc., AXI Referenc Guide, UG1037
- [116] O. Cobanoglu *et al.*, “"CMAD", a Full Custom ASIC for the Upgrade of COMPASS RICH-1”, “12th Workshop on Electronics for LHC and Future Experiments”, 434–437 (2006), [doi:10.5170/CERN-2007-001.434](https://doi.org/10.5170/CERN-2007-001.434)
- [117] OpenCores Organization, WISHBONE System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores (2010), Rev. B.4
- [118] B. Grube, I. Konorov and L. Schmitt, “COMPASS TCS documentation”, (2001), COMPASS note 2001-9
- [119] Xilinx, Inc., Vivado Design Suite Tcl Command Reference Guide, UG835
- [120] Texas Instruments, LVDS 4x4 Crosspoint Switch, Data sheet
- [121] Tektronix, Inc., Arbitrary/Function Generators AFG 3011 / 3021B / 3022B / 3101 / 3102 / 3251 / 3252 (2012), Data sheet
- [122] Xilinx, Inc., LogiCORE IP Integrated Bit Error Ratio Tester 7 Series GTP Transceivers v3.0 Product Guide, PG133
- [123] Xilinx, Inc., Transceiver Bit Error Isolation Methodology, WP425
- [124] J. Redd, “Calculating Statistical Confidence Levels for Error-Probability Estimates”, *Lightwave* 110–114 (2000)
- [125] IEEE Computer Society, IEEE Standard for Ethernet (2005), IEEE Std 802.3-2005
- [126] J. Kalisz *et al.*, “Single-chip interpolating time counter with 200-ps resolution and 43-s range”, *IEEE Transactions on Instrumentation and Measurement* **46** 851–856 (1997), [doi:10.1109/19.650787](https://doi.org/10.1109/19.650787)
- [127] F. Baronti *et al.*, “On the differential nonlinearity of time-to-digital converters based on delay-locked-loop delay lines”, *IEEE Transactions on Nuclear Science* **48** 2424–2431 (2001), [doi:10.1109/23.983253](https://doi.org/10.1109/23.983253)
- [128] M. J. French *et al.*, “Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker”, *Nucl. Instrum. Meth. A* **466** 359 – 365 (2001), [doi:10.1016/S0168-9002\(01\)00589-7](https://doi.org/10.1016/S0168-9002(01)00589-7)

Ich erkläre hiermit, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt und die wörtlich oder inhaltlich übernommenen Stellen als solche kenntlich gemacht habe.

Freiburg, den 25. Oktober 2017