

Guido Alexander Willems

# Hardware Development for the Upgrade of the Transition Radiation Detector and Development of a Nuclei Trigger



# Hardware Development for the Upgrade of the Transition Radiation Detector and Development of a Nuclei Trigger

Inaugural-Dissertation  
zur Erlangung des Doktorgrades (Dr. rer. nat.)  
der Naturwissenschaften im Fachbereich Physik  
der Mathematisch-Naturwissenschaftlichen Fakultät  
der Universität Münster

vorgelegt von  
Guido Alexander Willems  
aus Münster  
– 2023 –





Dekan:

Prof. Dr. Rudolf Bratschitsch

Erster Gutachter:

Prof. Dr. Johannes Wessels

Zweiter Gutachter:

Prof. Dr. Christian Klein-Bösing

Tag der mündlichen Prüfung:

---

Tag der Promotion:

---



# Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. Physics Introduction</b>	<b>5</b>
2.1. The Standard Model . . . . .	5
2.2. Strong interaction . . . . .	6
2.3. Quark-Gluon Plasma (QGP) . . . . .	7
2.4. Heavy-Ion Collisions . . . . .	8
<b>3. Physics of Light Nuclei</b>	<b>11</b>
3.1. Production mechanisms of light nuclei . . . . .	11
3.1.1. Statistical thermal model . . . . .	12
3.1.2. Coalescence model . . . . .	13
3.2. Recent A Large Ion Collider Experiment (ALICE) results in Proton – Proton (pp) collisions . . . . .	15
<b>4. A Large Ion Collider Experiment</b>	<b>23</b>
4.1. ALICE particle detectors for track reconstruction and Particle Identifica- tion (PID) . . . . .	23
4.2. Central Trigger Processor . . . . .	29
4.2.1. Trigger hierarchy of ALICE . . . . .	29
4.2.2. Technical realisation of the communication between the Central Trigger Processor (CTP) and individual detectors . . . . .	31
4.2.3. Configuration of the CTP . . . . .	32

<b>5. TRD Operation and Developments in Run 2</b>	<b>33</b>
5.1. Transition Radiation Detector . . . . .	34
5.1.1. Detector Structure . . . . .	34
5.1.2. Principle of Particle Tracking and PID inside a ROC . . . . .	35
5.1.3. Gas system of the TRD (Service Task) . . . . .	38
5.1.4. Front End Electronics (FEE) and data flow in a ROC during readout . . . . .	40
5.1.5. Data Sampling and Processing in a Multi Chip Module . . . . .	42
5.1.6. Readout time windows . . . . .	43
5.1.7. Tracklet Data Format of the MCM . . . . .	45
5.1.8. Global Tracking Unit . . . . .	46
5.2. Transition Radiation Detector Nuclei Trigger . . . . .	51
5.2.1. Motivation of using the Transition Radiation Detector (TRD) for the nuclei trigger . . . . .	52
5.2.2. Feasibility studies for the nuclei trigger implementation . . . . .	53
5.2.3. Hardware implementation of the nuclei trigger . . . . .	75
5.2.4. TRD nuclei trigger in Proton – Lead (pPb) collisions . . . . .	76
5.2.5. TRD nuclei trigger in pp collisions . . . . .	84
<b>6. TRD Upgrade and Developments for Run 3</b>	<b>91</b>
6.1. Transition Radiation Detector Upgrade Overview . . . . .	92
6.1.1. Optimisation of the existing Front End Electronics . . . . .	93
6.1.2. Common Readout Unit . . . . .	95
6.1.3. O2 Data Acquisition System . . . . .	99
6.1.4. New Run 3 Readout Scheme for the TRD . . . . .	102
6.2. TRD laboratory setup at the University of Münster . . . . .	104
6.2.1. Location and available material . . . . .	104
6.2.2. New laboratory setup . . . . .	105
6.3. Consideration of the TRD FEE optical drive capabilities . . . . .	108
6.3.1. Theoretical estimation of the optical margin . . . . .	108
6.3.2. Investigation on the actual current setpoints in Run 2 . . . . .	113
6.3.3. Measurement of the Optical Path Attenuation . . . . .	114
6.3.4. Verification of data reception in CR1 . . . . .	116

6.3.5.	Data taking with the GTU and the optical fibre loop . . . . .	117
6.3.6.	Sensitivity test with the CRU in the lab at the Universität Münster	117
6.3.7.	Optical power measurements with the CRUs after their installation in CR1 . . . . .	118
6.4.	Readout Fibre Installations at Point 2 . . . . .	122
6.4.1.	Mapping of Common Readout Unit (CRU) Optical Readout Interface (ORI) link numbers to half chambers . . . . .	122
6.4.2.	Installation . . . . .	124
6.5.	Upgrade of the CTP . . . . .	127
6.5.1.	Fully revised Local Trigger Unit . . . . .	127
6.5.2.	Trigger and Timing Control over Passive Optical Network (TTC-PON) . . . . .	127
6.5.3.	Trigger and Timing Control (TTC)-Passive Optical Network (PON) Latency Considerations . . . . .	128
6.6.	The CRU Common Logic Firmware in the context of the TRD . . . . .	129
6.6.1.	CRU FPGA firmware . . . . .	130
6.6.2.	Common CRU Firmware . . . . .	131
6.6.3.	O2 Data Format . . . . .	135
6.7.	Prerequisites for TRD CRU Custom Firmware Development . . . . .	142
6.7.1.	TRD specific implementation of the O2 data format . . . . .	143
6.7.2.	Note on RDH fields for TRD . . . . .	144
6.7.3.	O2 payload packetisation scheme for TRD . . . . .	145
6.8.	Validation of the TRD CRU Custom Firmware and Tests . . . . .	150
<b>7.</b>	<b>Summary</b>	<b>155</b>
<b>8.</b>	<b>Zusammenfassung</b>	<b>159</b>
	<b>Appendix</b>	<b>163</b>
<b>A.</b>	<b>Theory of Operation of the TRD CRU Custom Firmware</b>	<b>165</b>
A.1.	trd_custom_module_wrapper . . . . .	166
A.2.	ori_wrapper . . . . .	166
A.3.	ori_link . . . . .	167

A.4. <code>trd_user_logic</code> . . . . .	169
A.5. <code>buffered_ttc_pon_handler</code> . . . . .	170
A.6. <code>cdc_ttc_pci</code> . . . . .	170
A.7. <code>trigger_buffer</code> . . . . .	172
A.8. <code>ttc_pon_handler</code> . . . . .	175
A.9. <code>trigger_word_fifo</code> . . . . .	179
A.10. <code>aggregate_single_link_cores</code> . . . . .	179
A.11. <code>single_link_core</code> . . . . .	181
A.12. <code>link_realigner</code> . . . . .	182
A.13. <code>ready_valid_checker</code> . . . . .	182
A.14. <code>debug_memory</code> . . . . .	183
A.15. <code>cdc_ori_pci</code> . . . . .	184
A.16. <code>trd_ddg</code> . . . . .	186
A.17. <code>busy_logic</code> . . . . .	187
A.17.1. Communication of the <code>busy_logic</code> instances with DCS . . . . .	193
A.18. <code>convert_16to256</code> . . . . .	194
A.19. <code>link_output_fifos</code> . . . . .	198
A.20. Buffering Stage . . . . .	201
A.21. <code>packetisers</code> . . . . .	203
A.21.1. State Machine <code>s_state: wait_for_available_trigger_word</code> . . . . .	205
A.21.2. State Machine <code>s_state: ship_data</code> . . . . .	207
<b>B. Physics background</b> . . . . .	<b>213</b>
B.1. Kinematic variables . . . . .	213
B.2. Bethe-Bloch formula . . . . .	214
<b>C. List of Acronyms</b> . . . . .	<b>217</b>
<b>List of figures</b> . . . . .	<b>227</b>
<b>List of tables</b> . . . . .	<b>237</b>
<b>Bibliography</b> . . . . .	<b>239</b>
<b>Acknowledgement</b> . . . . .	<b>247</b>







# Chapter 1

## Introduction

This thesis has been accomplished at the A Large Ion Collider Experiment (ALICE), which is the Large Hadron Collider (LHC) experiment dedicated to investigate heavy-ion collisions. The work for this thesis started during the Run 2 period and was concluded during the subsequent shutdown phase of the LHC, when most of the LHC experiments as well as the LHC itself have undergone major technical upgrades for new challenges in Run 3.

A major constraint which hinders new discoveries and particle physics research in general is posed by the limitations in the number of collision events and particles of interest collected by the experiment. During the Run 1 and Run 2 periods of the ALICE experiment, this aspect is mitigated to a certain extent by employing trigger mechanisms. Triggers inspect events during the process of data taking and only record those which exhibit specific signatures and are beneficial for further physics analysis. In Run 3, the mitigation by using triggers is superseded by a continuous readout concept at an increased collision rate. The improvements allow for a remarkably faster data collection along with the opportunity of investigating effects without specific trigger signatures.

Within the scope of this thesis, a hardware trigger capable of triggering on events containing light nuclei has been studied and implemented for the first time in Run 2. The investigation of light nuclei physics is of fundamental interest, as the production of such loosely bound compound states from a hot and dense environment, as it is created in

heavy-ion collisions, is not yet fully understood. Furthermore, the interaction of light nuclei with matter has raised major interest as it can be applied to the transport of light nuclei within our Galaxy and can strengthen our understanding of dark-matter annihilation, as shown in the ALICE paper *Measurement of  ${}^3_2\text{He}^{2+}$  nuclei absorption in matter and impact on their propagation in the Galaxy* published in the journal *Nature* [1].

This thesis presents the development of the new nuclei trigger and shows the achieved performance using the collected data. The trigger has been operating successfully in both Proton – Lead (pPb) at  $\sqrt{s} = 8$  TeV and Proton – Proton (pp) at  $\sqrt{s} = 13$  TeV collisions until the end of Run 2 in 2018. In both collision systems, significant enhancements of (anti-)nuclei have been demonstrated. The new trigger has collected a significant amount of events containing light nuclei in both collision systems and the acquired data are currently being analysed further by the ALICE collaboration.

A description of the nuclei trigger developed in the context of this thesis has been published in the paper *The ALICE Transition Radiation Detector: Construction, operation, and performance* in the journal *Nuclear Instruments and Methods* [2].

The main goals of the ALICE upgrade for Run 3 are to deeply investigate heavy-ion collisions, including measurements of new observables, and improve the precision of measurements such as the production of light nuclei. For the Run 3 upgrade, the detectors are equipped with the latest available technology. An additional achievement of this thesis is the readout development for the upgrade of the Transition Radiation Detector (TRD) of ALICE. The main focus is on the development of the TRD specific parts of the Common Readout Unit (CRU) firmware, which has been entirely and exclusively performed as part of this thesis. The CRUs are new hardware devices which enable the high performance data taking operation in Run 3. The CRUs have been deployed at the core of the TRD infrastructure and control the TRD readout in Run 3 in close coordination with the other subsystems of the ALICE experiment. Furthermore, the CRUs receive and handle the entire data which is generated by the TRD.

The *ALICE upgrades during the LHC Long Shutdown 2* paper, which includes the TRD readout upgrade and the work presented in this thesis, has been published [3].

This thesis describes the evolution from the TRD Run 2 infrastructure and TRD readout systems, treats the planning and preparation for the Run 3 upgrade and documents the details of the newly developed hardware design. Within the scope of this work, the new hardware design has also been functionally tested and the interactions with other subsystems have been verified. The stability of the new readout system has been proven in long term technical runs without beam within its final environment and the system has successfully recorded its very first physics data sets during a short pilot beam period. The TRD readout systems perform successfully in today's regular Run 3 data taking operations in Proton – Proton (pp) and Lead – Lead (Pb–Pb) collisions, which have started in mid 2022.

For PhD students in the ALICE collaboration, the successful completion of a service task is mandatory and shall support the detector operation. In the context of this thesis, a TRD gas system service task has been completed. This task corresponded to an equivalent of 6 months full time work and consisted of supervising and operating the TRD gas system which must run permanently on a 24 hours / 7 days in a week base and required immediate actions in case of issues. In addition, the gas had to be manually purified from nitrogen contamination using a cryogenic procedure.

The thesis is organised as follows. A brief insight into physics results on light nuclei in pp collisions recently published by the ALICE collaboration is given in the beginning of this thesis. Furthermore, the ALICE experiment and the TRD infrastructure in Run 2 are introduced. Then the work concerning the newly developed nuclei trigger is presented. Finally, the ALICE and TRD upgrades are introduced, followed by the CRU hardware development and the deployment of the new hardware for the TRD upgrade. The theory of operation of the CRU firmware developed in the context of this thesis is given in the appendix.

The work for this thesis has been carried out during a long term stay at the Conseil Européen pour la Recherche Nucléaire (CERN). This thesis has been supported by the Bundesministerium für Bildung und Forschung (BMBF) by a Wolfgang-Gentner-Stipendium in the context of the Deutsches Technisches Doktorandenprogramm at CERN. In this program it is not foreseen to work on experimental or theoretical physics problems and it is required that the doctoral student contributes with engineering tasks and topics to the fundamental physics research at highest energies at the leading edge of today's physics research [4].

# Chapter 2

## Physics Introduction

The ALICE experiment is designed for investigating the physics of heavy-ion collisions. This chapter gives a short introduction into the nature of heavy-ion collisions and the Quark-Gluon Plasma (QGP). For a theoretical background, this chapter starts with an overview over the fundamental particles and their interactions.

This chapter is based on the References [5–8], which can be consulted for further details.

### 2.1. The Standard Model

Today’s understanding of the composition and interactions of matter and anti-matter is based on the standard model of particle physics. It assumes all matter to be composed of certain fundamental particle types. A fundamental particle does not have further substructure and can interact with other particles through certain fundamental forces [5].

The elementary particles are divided into quarks (up, down, charm, strange, top, and bottom) and leptons (electron, muon, tau, and their corresponding neutrinos), which are fermions. In addition, each elementary particle has its corresponding anti-particle. The standard model describes the interaction by introducing force carriers, which are gauge

bosons. All interactions in nature can be divided into four categories: the electromagnetic interaction, which is mediated by photons; the weak interaction, which is mediated by W and Z bosons; the strong interaction, which is mediated by gluons; and the gravitation, which is not described within the scope of the current standard model [5].

In order to interact electromagnetically, a particle has to carry an electric charge. In analogy to this concept, it is possible to define a colour charge which is carried by the particles participating in the strong interaction. In the standard model, only quarks and gluons carry colour charge [5].

## 2.2. Strong interaction

Quarks have never been observed isolated in nature, but only as constituents of composite particles, which are called hadrons. The strong interaction is stronger than any electromagnetic repulsion and binds quarks within hadrons [5,6].

Hadrons are required to be neutral with respect to the colour charge. As opposed to the scalar electric charge, the colour charge is described by a three dimensional vector. Its components are referred to as red, green, and blue. This convention is chosen in analogy to optical colours as to how colour neutrality can be achieved [5,6].

Historically, two groups of hadrons have been defined: baryons, which are composed of three valence quarks; and mesons, which are composed of valence quark and anti-quark. For the baryons, colour neutrality is obtained by the three quarks carrying different colours. For the meson, colour neutrality is achieved by combining colour with anti-colour. In addition to the valence quarks, hadrons contain sea quarks and gluons as product of interactions [5,6].

Furthermore, additional exotic types of hadrons are being investigated. For example, tetraquarks are hadrons composed of four valence quarks and pentaquarks are hadrons com-


posed of five valence quarks. These two exotic states have been recently measured [9, 10].

Protons and neutrons, which are the constituents of atomic nuclei, are themselves baryons. Protons are composed of two up valence quarks and one down valence quark, while neutrons are composed of one up valence quark and two down valence quarks. To the outer environment, protons and neutrons in principle appear as colour neutral, i.e. they do not interact strongly with other hadrons. While this is true at large distances from the proton or neutron, residual effects of the strong interaction are present in close vicinity of the hadrons. These residual effects decrease with the distance and are significantly weaker than the strong interaction among the quarks within the hadron, but their strength is still enough to bind the nuclei of atoms with a sufficient binding energy, overcoming all electromagnetic repulsions. The average binding energy of atomic nuclei is in the order of 8 MeV per nucleon [11].

The theory which describes the strong interaction is the Quantum Chromodynamics (QCD). The strength of the strong interactions of quarks and gluons is determined by the QCD running coupling constant  $\alpha_s$ . As shown in Fig. 2.1,  $\alpha_s$  depends on the momentum transfer  $Q$  in a process. The coupling  $\alpha_s$  becomes small for large  $Q$  values, which correspond to short distances. At asymptotically large  $Q$  values, quarks are effectively free. This phenomenon is called asymptotic freedom. Contrarily,  $\alpha_s$  becomes large for low  $Q$  values, which correspond to large distances. In this situation, quarks are effectively confined in hadrons [5, 6, 12].

## 2.3. Quark-Gluon Plasma (QGP)

Under extreme conditions of high temperature or large chemical potential, the QCD predicts a phase transition into a state of matter, called Quark-Gluon Plasma (QGP), where quarks and gluons are deconfined from hadrons [6, 8]. Fig. 2.2 shows a schematic view of the QCD phase diagram. The hadronic phase is present at low chemical potential and temperature [13]. The phase transition from the QGP to the hadron gas at high temperature and low chemical potential is supposed to be similar to the conditions at the



**Figure 2.1.:** QCD running coupling constant  $\alpha_s$  as a function of the momentum transfer  $Q$ . The results are obtained by different analyses from various experiments. Figure from Ref. [12].


time of the primordial universe. The QGP and the phase transition can be reproduced in the laboratory by colliding heavy ions at ultra-relativistic energies [6, 8]. For large chemical potential and low temperature, a colour superconducting phase is expected [13].

## 2.4. Heavy-Ion Collisions


Particle accelerators are used to collide heavy ions in the laboratory. In a circular collider, two beams in opposite directions are accelerated independently and are brought to collide at specific interaction points, where suitable particle detectors are installed. Figure 2.3 illustrates the time evolution of a central high-energy heavy-ion collision with creation of the QGP.

Due to the high velocities, the accelerated heavy ions are Lorentz contracted in the laboratory frame. The heavy-ion collision in the collider undergoes several stages. In a





**Figure 2.2.:** Schematic view of the temperature  $T$  as a function of the chemical potential  $\mu_B$  of the QCD phase diagram. Figure from Ref. [13].



**Figure 2.3.:** Schematic time evolution of a central high-energy heavy-ion collision with QGP formation. Figure from Ref. [14].

first stage, parton interactions take place converting energy to form other particles. In a preequilibrium phase, the temperature is not uniform across the fireball. In this stage, the system is not in thermodynamic equilibrium such that a hydrodynamic description cannot yet be used. Once the system reaches its thermodynamic equilibrium, the QGP is established. The system expands and behaves as a medium which can be described macroscopically by hydrodynamic models under the assumption of local thermal equilibrium. The dynamics of the system evolution is described by thermodynamic variables, such as temperature, pressure and energy density together with a suitable equation of state. In addition, the velocity of the expansion is taken into account in hydrodynamic models [6, 8].

As the expansion continues, the temperature decreases. Once the system is cooled down sufficiently, the confinement is restored and the quarks form hadrons. This phase is called hadronisation. When the system reaches the chemical freeze-out temperature, the energy involved in the interactions does not suffice to generate new hadrons. At this point, the total yield of each particle type is fixed. After the chemical freeze-out, hadrons can still interact elastically, such that the momenta and thus the momentum differential yields may still be altered. Only after the kinetic freeze-out, the momenta do not change as no interactions, neither inelastic nor elastic, among the particles occur anymore. In this stage, the particles travel outbound and can be measured by detectors [6, 8].

Recently, signatures of the QGP have been observed in small collision systems. However, the nature of these systems is not fully understood [15]. Except for the QGP phase under investigation, the time evolution from Fig. 2.3 remains valid for describing small collision systems [6, 8].

# Chapter 3


## Physics of Light Nuclei

This thesis involves the development and commissioning of a hardware trigger on light nuclei. Before presenting details about the development of the trigger and its technical implementation, this chapter briefly presents the physics background and the motivation of the trigger. Analyses involving nuclei usually have to deal with a limited number of events, which leads to statistical limitations of the accuracy. The developed trigger has increased the number of events containing nuclei significantly. Some of the most recent results in nuclei physics obtained by the ALICE collaboration are summarised in this chapter. This chapter is based on the references [16–24]. For further details, these references can be consulted.

The development of the CRU firmware in the context of the upgrade of the TRD along with the upgrade of the entire ALICE experiment also aims at increasing the number of events in general, including those containing nuclei, by increasing the detector readout rate.

### 3.1. Production mechanisms of light nuclei

An important research goal concerning light nuclei observed in high energy collisions is to understand their production mechanisms. Light nuclei - as deuterons, tritons,  ${}^3_2\text{He}^{2+}$ ,  ${}^4_2\text{He}^{2+}$  and their antiparticles - are compound particles of protons and neutrons which are



**Figure 3.1.:** Nucleon composition of deuteron, triton and  ${}^3\text{He}^{2+}$ .

held together by a small binding energy (Fig. 3.1). For light nuclei, the binding energy is lower than the average binding energy per nucleon, which is approximately 8 MeV [11].

The theory predominantly favours two production mechanism models, the production in the thermal statistical model and the production by coalescence. When testing the hypotheses of these models experimentally [20, 25, 26], so far none of them can be ruled out and both models can describe the measurements within the current experimental uncertainties. The two models make different assumptions about the fact during which stage of a collision the nuclei are produced.

### 3.1.1. Statistical thermal model

In the statistical thermal model, nuclei are produced at chemical freeze-out. The model makes predictions on the expected particle yields and formulates a dependence between the differential yield of a certain particle species, the particle mass  $m$  and the chemical freeze-out temperature  $T_{\text{ch}}$  [26]:

$$\frac{dN}{dy} \propto \exp\left(-\frac{m}{T_{\text{ch}}}\right). \quad (3.1)$$

If the model is valid, the chemical freeze-out temperature shall be equal for the experimentally measured nuclei. In this case, nuclei can be used as a precise thermometer for the chemical freeze-out temperature due to their large masses.

It can be deduced that, when comparing the experimentally accessible density ratio between 2 particle species  $n_1/n_2$  with a mass difference of  $\Delta m$ , the ratio is directly connected to the chemical freeze-out temperature [26]:

$$\frac{n_1}{n_2} = \exp\left(-\frac{\Delta m}{T_{\text{ch}}}\right). \quad (3.2)$$

Measurements of particle ratios in heavy-ion collisions are in agreement with a chemical freeze-out temperature of approximately 160 MeV [26] and the predictions of particle yields by this model generally agree with measurements [20].

### 3.1.2. Coalescence model

The statistical thermal model assumes that nuclei with their low binding energy of a few MeV are generated at the chemical freeze-out, i.e. in an environment with a temperature of approximately 160 MeV [20, 26], although the low binding energy of nuclei may affect their stable existence in an environment at such high temperature. Assuming this, the production may actually not be possible at chemical freeze-out and happens in fact later. The coalescence model [16–21] describes the production of nuclei in this later stage.

Also, in the coalescence model, the notion of particle generation at chemical freeze-out is still valid, but only for the constituting nucleons and lower-mass particles. The nuclei are bound states of nucleons at low binding energy and they are generated at kinetic freeze-out by a coalescence of the available nucleons. At this stage, the temperatures are sufficiently low to allow for the stable existence of the bound states.

The coalescence model assumes that a coalescence of nucleons to form a nucleus is only possible if the nucleons are close enough in phase space and their difference in momentum is smaller than an empirical parameter  $p_0$ . Assuming that protons and neutrons have the same mass and  $p_T$  spectrum, the yield of any nucleus can be expressed using the proton yield [20]:

$$E_A \frac{d^3 N_A}{dp_A^3} = B_A \left( E_p \frac{d^3 N_p}{dp_p^3} \right)^A. \quad (3.3)$$

In this equation,  $A$  is the number of nucleons within the generated nucleus and  $p_A = Ap_p$  its momentum.  $\frac{d^3 N_A}{dp_A^3}$  and  $\frac{d^3 N_p}{dp_p^3}$  are the predicted nuclei yield and the proton yield, respectively.  $B_A$  is a factor, which can be measured experimentally.

The simplest available coalescence model does not consider any spatial properties of the production mechanism, but considers only momentum space. The coalescence probability is related to the probability of the nucleons to be within a sphere with the radius  $p_0$  in momentum space. Ref. [20] finds for  $B_A$

$$B_A = \left( \frac{4\pi}{3} p_0^3 \right)^{A-1} \frac{M}{m^A}, \quad (3.4)$$

where  $M$  is the mass of the generated nucleus and  $m$  is the mass of the proton. The term  $\frac{4\pi}{3} p_0^3$  is also called the coalescence volume in momentum space.

In the simple coalescence model, the parameter  $B_A$  is a constant and does not depend on variables such as the transverse momentum. The independence of these variables can be experimentally tested, in order to investigate the validity of the simple coalescence model.

Generally,  $B_2$  can be measured by comparing the spectra of deuterons with the ones of protons, while  $B_3$  can be measured by comparing the spectra of  ${}^3\text{He}^{2+}$  with the ones of protons.

The simple coalescence has later been extended, such that the neglected spatial aspects, such as the size of the emitting source and the size of the produced nucleus are taken into account. In this advanced coalescence model,  $B_A$  is no longer constant. For the refined model, Ref. [27, 28] finds for  $B_A$

$$B_A = \frac{2J_A + 1}{2^A} \frac{1}{\sqrt{A}} \frac{1}{m_T^{A-1}} \left( \frac{2\pi}{R^2 + \left(\frac{r_A}{2}\right)^2} \right)^{\frac{3}{2}(A-1)}, \quad (3.5)$$


where  $J_A$  is the spin of the resulting nucleus,  $m_T$  is the transverse mass of each nucleon participating in the coalescence,  $R$  is the radius of the emitting source and  $r_A$  is the radius of the resulting nucleus.

## 3.2. Recent ALICE results in pp collisions

This section shows some recent ALICE results on light nuclei in pp collisions, where the nuclei trigger is required to enhance statistics. The means of PID are not shown here, but they will be discussed in more details in Section 4.1. In all figures in this chapters, systematic uncertainties are shown as boxes and statistical uncertainties are shown as vertical bars.

Figure 3.2 shows the integrated yields of antiprotons, antideuterons, and  $\overline{{}_2^3\text{He}^{2+}}$  in pp collisions at  $\sqrt{s} = 7 \text{ TeV}$  [22]. A reduction of the  $p_T$  integrated yield ( $dN/dy$ ) is observed of about 1000 for each additional nucleon. The horizontal lines represent a fit with the function  $c_0 A c_1^A$  in accordance with the coalescence model.

The antideuteron to deuteron ratio  $\bar{d}/d$  is presented in Fig. 3.3 for different collision energies and in Fig. 3.4 for different multiplicities in pp collision at  $\sqrt{s}=7 \text{ TeV}$ . The ALICE collaboration expresses multiplicities as multiplicity classes, where class I+II corresponds to events with the largest multiplicity. The multiplicity decreases with the class number. Furthermore, Fig. 3.3 also reports the squared ratio between antiparticles



**Figure 3.2.:**  $p_T$  integrated yields ( $dN/dy$ ) of antiprotons, antideuterons and  ${}^3\bar{\text{He}}^{2+}$  nuclei as a function of the number of antinucleons in pp collisions at  $\sqrt{s} = 7$  TeV. Figure from Ref. [22].


and particles for protons,  $(\bar{p}/p)^2$  for comparison.

It can be observed that across all presented energies, multiplicity classes and  $p_T$  bins, antideuterons are produced with similar abundances as deuterons within the uncertainties.


Figure 3.5 shows the coalescence parameter  $B_2$  of deuterons and antideuterons as a function of the transverse momentum per nucleon in different multiplicity classes in pp collisions at  $\sqrt{s} = 7$  TeV [23]. Different scalings are applied for better readability.

Figure 3.6 shows the coalescence parameter  $B_2$  of deuterons and antideuterons, presented separately, as a function of the transverse momentum per nucleon  $p_T/A$  at mid-rapidity in pp collisions at  $\sqrt{s} = 0.9, 2.76$ , and 7 TeV [22]. For more details about the kinematic variables, refer to appendix B.1.







**Figure 3.3.:** Antideuteron to deuteron ratios ( $\bar{d}/d$ ) as a function of the transverse momentum per nucleon  $p_T/A$  in pp collisions in different collision energies. The  $(\bar{p}/p)^2$  ratios are shown for comparison. Figure from Ref. [22].





**Figure 3.4.:** Antideuteron to deuteron ratios as a function of  $p_T$  in different multiplicity classes in pp collisions at  $\sqrt{s} = 7$  TeV. Figure from Ref. [23].



**Figure 3.5.:** Coalescence parameter  $B_2$  of deuterons and antideuterons as a function of the transverse momentum per nucleon  $p_T/A$  in different multiplicity classes in pp collisions at  $\sqrt{s} = 7$  TeV. Figure from Ref. [23].



**Figure 3.6.:** Coalescence parameter  $B_2$  of deuterons and antideuterons, shown separately, as a function of the transverse momentum per nucleon  $p_T/A$  in pp for different collision energies. Figure from Ref. [22].



**Figure 3.7.:** Coalescence parameters as a function of the transverse momentum per nucleon  $p_T/A$  in pp collisions at  $\sqrt{s} = 13 \text{ TeV}$  for different multiplicity classes. Left hand panel: Deuterons and antideuterons (parameter  $B_2$ ).  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  (parameter  $B_3$ ). Figures from Ref. [24].

Figure 3.7 shows the coalescence parameter  $B_2$  for deuterons and antideuterons and the coalescence parameter  $B_3$  for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of  $p_T/A$  in pp collisions at  $\sqrt{s} = 13$  TeV [24]. The parameter  $B_2$  is obtained using High Multiplicity (HM) triggered data. The results are shown separately for different multiplicity classes, where class I corresponds to the largest multiplicity. Subsequent multiplicity classes correspond to lower multiplicities. The parameter  $B_3$  is obtained separately, using HM and Minimum Bias (MB) triggered data. In addition, the MB data sample is separated into 2 multiplicity classes I and II, where class I corresponds to the larger multiplicity.

It can be observed that in Fig. 3.5 the coalescence parameters  $B_2$  of deuterons and antideuterons show no significant dependence on the transverse momentum per nucleon within the uncertainties in all presented multiplicity classes [23]. This behaviour is in agreement with the simple coalescence model.

In Fig. 3.6 an increasing trend of the coalescence parameter  $B_2$  can be observed. The difference becomes visible when results are integrated over the multiplicity and the behaviour is observed across different collision energies. Simulations using an event generator together with an afterburner to simulate the coalescence of nucleons with similar momenta confirm this trend [24].

In Fig. 3.7 an increasing trend of the coalescence parameter  $B_3$  with respect to  $p_T/A$  is also observed for both integrated and differential multiplicities [24]. This observation is not expected in the simple coalescence model, but it is found to be in agreement with the advanced coalescence model described by equation 3.5 [28].

Differential results – in multiplicity,  $p_T$ , particle/antiparticle, etc – are important to investigate the production of light particles. In order to obtain conclusive results with sufficient accuracy, large statistics are required. In addition, results for light nuclei with  $A > 2$  are generally impacted by lower statistics due to lower production yields. The nuclei trigger greatly enhances the statistics, especially for the case of light nuclei with

$A > 2$ , and thus extends the physics prospects of the ALICE experiment.

This is not only relevant for the directly measured nuclei, but also for strengthening the understanding of the physics of hypernuclei, such as their production yields and their lifetimes. An example is the hypertriton, which is the lightest known hypernucleus and consists of a proton, a neutron and a  $\Lambda$  baryon, whose properties raise major interest in the research community [29].

Other measurements which can be extended with larger statistics concern the interactions of antinuclei with matter. Recent ALICE measurements allow for the conclusion that  $\overline{{}_2^3\text{He}^{2+}}$  can travel long distances within our Galaxy and can be used to study cosmic-ray interactions and dark-matter annihilation [1].



# Chapter 4

## A Large Ion Collider Experiment


The ALICE experiment is located at CERN and belongs to the four largest LHC experiments. The LHC beams collide within ALICE and the generated particles are measured by the particle detectors of ALICE. The experiment is operated in all proton–proton (pp), proton–lead (pPb) and lead–lead (Pb–Pb) collision systems. The detectors are optimised for the operation in heavy-ion collisions, which are the primary research field of ALICE. The detectors of the ALICE experiment provide excellent particle tracking capabilities with track reconstruction down to the level of single particles, even in high-multiplicity environments, at low to medium transverse momentum.

In this chapter, a summary about the most important subsystems of ALICE is given. It is based on the references [30, 31], which can be consulted for further details.

### 4.1. ALICE particle detectors for track reconstruction and PID

Figure 4.1 shows a sectional view of the setup of the ALICE experiment with its detectors. The collisions take place in the centre of the central barrel (in red). The particles emerge from the collision vertex into all directions. ALICE is capable of measuring the particles in the transverse region (perpendicular to the beam pipe). Many detectors in the central

## THE ALICE DETECTOR



**Figure 4.1.:** Sectional view of the ALICE experiment during Run 2. Figure from Ref. [32].

barrel provide full azimuthal coverage within the pseudo rapidity region  $-0.9 \leq \eta \leq 0.9$ . For more details about the kinematic variables, refer to appendix B.1. Furthermore, ALICE is equipped with specific forward detectors, including the muon spectrometer at the forward pseudo rapidity region  $-4.0 \leq \eta \leq -2.5$  [30, 31].

In the transverse region at mid rapidity, charged particles are deflected by a homogeneous magnetic field in beam direction. The magnetic field is generated by a large solenoid, also called the L3 magnet (10 in Fig. 4.1). During the reconstruction of tracks, it is possible to reconstruct the momentum of each particle by measuring the amount of track bending in the magnetic field. Low momentum particles travel on trajectories with smaller radii due to lower Lorentz forces [30, 31].

At mid rapidity, the Inner Tracking System (ITS) (1 in Fig. 4.1) surrounds the beam pipe. The ITS is a silicon based detector and is subdivided into three subdetectors: the Silicon Pixel Detector (SPD) (a in Fig. 4.1), the Silicon Drift Detector (SDD) (b in




Fig. 4.1) and the Silicon Strip Detector (SSD) (c in Fig. 4.1). The ITS is used as a tracking detector in order to locate the vertex position. In addition, it also provides energy loss measurements, such that it can be used for Bethe-Bloch based PID (for more details, refer to appendix B.2) together with the global particle transverse momentum reconstruction [30,31].

Apart from the ITS, there are the VZERO Detector (V0) (d in Fig. 4.1), the Forward Multiplicity Detector (FMD) (e in Fig. 4.1) and the Time 0 Detector (T0) (d in Fig. 4.1) in close vicinity to the vertex. The V0 is based on Photo Multipliers (PMs) and serves as a detector to provide minimum bias triggers when a collision occurs. Its acceptance in pseudo rapidity is  $2.8 \leq \eta \leq 5.1$  on the A-side and  $-3.7 \leq \eta \leq -1.7$  on the C-side. The FMD is used to measure the multiplicity in the forward region and to provide multiplicity based triggers. The T0 provides time measurements of the precise starting time of the particle trajectories. Together with other detectors, it allows for precise particle velocity measurements, which can be used for PID [30,31].

The TPC (3 in Fig. 4.1) is the main tracking detector of ALICE and consists of a large volume which contains a noble gas mixture. The TPC contains high voltage electrodes in order to generate a homogeneous electric field parallel to the beam direction inside the gas volume. Particles which enter the gas volume, ionise the gas and liberate charges which are subsequently drifting along the electric field lines towards the readout electrodes located at the end of the cylindrical TPC shape. The impact location of the drift charges together with the measured drift time, allow for a precise reconstruction of the locations of each track in 3-dimensional space and for transverse momentum reconstruction [30,31].

Furthermore, the TPC measures the energy loss of the particle. The mean energy loss of a charged particle can be predicted by the Bethe-Bloch formula and depends on the particle charge, rest mass and its momentum. For details, refer to appendix B.2. The energy loss within the TPC is characteristic for a particular particle species and a given momentum. However, ambiguities are present in certain momentum regions. When plotting the energy loss of a particle measured by the TPC against the momen-



**Figure 4.2.:** Top panel: Separation of different particle species for PID with the TPC. The figure shows for each particle the reconstructed particle momentum on the horizontal axis TPC energy loss on the vertical axis. Bottom panel: Separation of different particle species for PID with the TOF detector. The figure shows for each particle the reconstructed particle momentum on the horizontal axis and the particle velocity measured with TOF on the vertical axis. Both figures have been adapted from [33].

tum from the global reconstruction, particles of the same species gather around specific splines in the plot. The location of each particle within this plot can be used for particle identification purposes [30,31]. An example of the TPC PID plot is shown in Fig. 4.2.

The TRD (4 in Fig. 4.1) is another gas detectors which measures the energy loss and which in addition enhances the energy loss by a radiator material, especially for electrons and their antiparticles. Furthermore, the TRD also provides tracking and triggering functionalities [30,31]. Since the TRD is the central detector in this thesis, more details will be given in Section 5.1.

The TOF detector (5 in Fig. 4.1) measures for each particle the arrival time at the TOF detector. When combining this data with the starting time measured by T0, a time difference can be calculated for each particle. With the track length between the vertex and TOF, as it can be obtained from the tracking detectors via a global reconstruction, the velocity of the particle expressed as a fraction of the speed of light  $\beta = v/c$  can be obtained for every particle. Since different particles have different rest masses, the dependence between the velocity  $\beta$  and the momentum is different for each particle species. When plotting the velocity  $\beta$  against the momentum, particles with the same rest masses gather around specific splines in the plot. The location of each particle within this plot can be used for PID purposes [30,31]. An example for the TOF PID plot is shown in Fig. 4.2.

ALICE employs the Electromagnetic Calorimeter (EMCal) (7 in Fig. 4.1) and the Di-Jet Calorimeter (DCal) (8 in Fig. 4.1), which are used for electron identification by measuring the relation between energy and momentum  $E/p$ . They can also be used to identify photons and neutral pions. EMCal and DCal cover only a fraction of the full azimuth and they are located in opposite azimuth [30,31].

ALICE also contains the high-resolution Photon Spectrometer (PHOS) (9 in Fig. 4.1). In front of the PHOS, the Charged Particle Veto Detector (CPV) is located, which is used

to reject charged particles [30,31].

The Photon Multiplicity Detector (PMD) (16 in Fig. 4.1) is located in a different acceptance region and is used for measuring photon multiplicity. Internally, it also contains a veto detector for charged particles in its first layer [30,31].

The High Momentum Particle Identification Detector (HMPID) (6 in Fig. 4.1) is a Ring-imaging Cherenkov Detector (RICH) detector and allows to identify particles with high momentum, where the major detectors of ALICE do not provide enough separation between the particles. The HMPID covers only a fraction in rapidity and azimuth [30,31].

In forward direction, ALICE features an absorber cone made of tungsten (11 in Fig. 4.1) which absorbs mainly hadrons and muons at low momentum. The muons travelling in forward direction are deflected by a dipole magnet (15 in Fig. 4.1) and are then detected by using muon trigger (14 in Fig. 4.1) and tracked by the muon tracker (12 in Fig. 4.1). The muon wall (13 in Fig. 4.1) provides means of shielding in order to reduce the background of low energy particles in the muon trigger (14 in Fig. 4.1) [30,31].

In the very forward rapidity region, far away from central barrel in the LHC tunnel, there is a calorimeter, the Zero Degree Calorimeter (ZDC) (18 in Fig. 4.1), and the scintillator detector, the ALICE Diffractive Detector (AD) (17 in Fig. 4.1) [30,31].

On top of the L3 magnet, the ALICE Cosmic Ray Detector (ACORDE) is located. The ACORDE is a scintillator detector and it functions as a trigger detector on cosmic rays [30,31].

## 4.2. Central Trigger Processor

The typical time distance between bunches in the LHC is in the order of 25 ns and the crossing of bunches depends on the concrete LHC filling scheme. When bunches are forced to cross each other, new particles may be generated. However, this is a statistical process and implies that the time between interesting events (e.g. containing nuclei) is not constant. The timing of a particular event is impossible to predict.

Furthermore, ALICE consists of many detectors with very different properties with respect to their readout timing, some detectors provide a fast readout, while others require a longer time for the readout of a particular event. Typically, the latter ones provide in turn a more elaborate data set which may be required for meaningful offline analysis.

The raw data rate of all detectors in ALICE triggered at their maximum rates is in the order of several TB/s and it is not feasible to permanently record on permanent storage at this data rate with resources available in Run 2.

Fortunately, not all events are interesting for physics analysis, which allows to reduce the rate of events which need to be recorded on permanent storage significantly.

This section is based on information from the references [34–39], which may be consulted for further details.

### 4.2.1. Trigger hierarchy of ALICE

In order to technically address all these issues and to efficiently record data at a feasible rate while filtering systematically for interesting events, a sophisticated trigger system is in place which is centrally managed by the CTP. The idea is to work with a hierarchic system in which the fast detectors perform a pre-examination of the event directly after the collision. The readout of the slower detectors is only triggered in case the event was

**Table 4.1.:** Trigger levels used in ALICE during Run 2 [36].

Trigger level	Typical time after the collision
Collision (not detected yet)	0
Level Minus One (LM)	0.3 $\mu\text{s}$
Level 0 (L0)	1.2 $\mu\text{s}$
Level 1 (L1)	6.5 $\mu\text{s}$
Level 2 accept (L2a)/Level 2 reject (L2r)	$\approx 88 \mu\text{s}$

found worth recording. This concept can be generalised by introducing multiple stages, which are called trigger levels. Only events which pass the last trigger level qualify for permanent data storage. The slowest detector determines the maximum event recording rate of the experiment. Table 4.1 shows the hardware based trigger levels as used during Run 2.

For the LM, the L0 and the L1 the timing is fixed as with the due times given in the table. If the CTP does not send one of these fixed timing triggers in due time, the event is considered as globally discarded after the timeout. For the Level 2 (L2) trigger, the timing is more flexible and explicit L2 accept or L2 reject trigger messages are issued by the CTP [36].

In fact, there is another last instance in addition to the hardware based triggers which is able to discard an event entirely before it goes to the final storage. It is called the High Level Trigger (HLT). This trigger is not considered here because it is not hardware based and works independently from the CTP. Instead, it is implemented in software and therefore does not operate with strict real time requirements with respect to the data taking, i.e. it is also not synchronous to the LHC Bunch Crossing (BC) clock [36].

During the operation throughout Run 1 and Run 2, not all trigger levels, for which the system is designed for, have been permanently used. In practice, the experiment has been operating with two important strategies. The first one is minimum bias, where in principle every L0 trigger leads to the recording of the event. Of course, the rate has to

be limited artificially to account for the limits which slow readout detectors and storage rate bottlenecks impose. This down scaling has to select events randomly, i.e. without introducing any bias. The second strategy is the recording of rare triggers. In this mode, the trigger scheme works as described up to the L1 trigger level. In case that an event passes the L1 level, the data are determined to be recorded on final storage, i.e. after the L1 trigger, the CTP issues the L2a in all cases. With these strategies, also the HLT is not used for discarding entire events. Instead it has been compressing the data using lossless methods in order to exploit the available storage bandwidth to a better extent.

#### 4.2.2. Technical realisation of the communication between the CTP and individual detectors

In Run 2 the communication between the CTP is realised differently for the downstream data (CTP  $\rightarrow$  detector) and upstream data (detector  $\rightarrow$  CTP). The downstream data consist mainly of the different triggers which have to be sent to the detectors with the correct timing. These trigger data are broadcasted by the CTP via the Trigger and Timing Control (TTC) network. This is an unidirectional optical network designed with single mode fibres and driven by laser diodes on the CTP side using Trigger and Timing Control Emitter (TTCex) devices. To receive the triggers, all detectors are required to use a specific Integrated Circuit (IC), the Trigger and Timing Control Receiver chip (TTCrx), which is common to all LHC experiments. Since the trigger data are always synchronous to the LHC BC clock, the TTCrx IC is also able to recover the LHC BC clock from the established link, though it is not a dedicated clock transmission. The LHC orbit signal can also be recovered locally [35,37].

For the upstream communication, copper wires are used. The data to be sent from the detectors to the CTP are the detector specific trigger contributions. After the pre-examination of an event by a fast detector, the detector has to communicate to the CTP whether the detector considers the event worth to climb up the trigger hierarchy. This is typically the case if the detector finds a specific signature, like a particle with high transverse momentum or other possibly more complex criteria. Every detector can have several trigger contribution inputs to the CTP, one for every specific signature. This gives

full control of the experiment-wide readout policy to the CTP. In case that the readout policy defined by the Run Coordination (RC) changes, only the CTP configuration needs changing and individual detector settings can be left untouched [35,37].

### 4.2.3. Configuration of the CTP

The configuration of the CTP is very thorough and allows a fine grained control of all parameters in order to meet the requirements of RC. The trigger scheme which is applied by the CTP technically enforces the readout of those events which are considered to be interesting in view of the ALICE physics programme. It also allows to share in a fair manner the available number of events which can be read out among different trigger types and strategies as the total event statistics is limited by the running time and the slowest readout detectors. This can be achieved by applying down scaling to the trigger inputs in order to account for the requirements in statistics for different types of data. The requirements in statistics and the split factors between different triggers are defined by the Physics Board (PB) of ALICE depending on the priorities of different aspects of the ALICE physics programme.

In addition, the trigger hierarchy shown in Section 4.2.1 does not necessarily need to be applied to the entire experiment at once. ALICE consists of different detectors and not all detectors, which provide trigger inputs, have to trigger the readout of all detectors in the experiment. Instead, the detectors can be grouped for readout into readout clusters. Every readout cluster can be triggered by an individually configurable set of trigger inputs. This is of great advantage in order to reduce the impact that slow detectors have on the overall data taking performance.



## Chapter 5

# TRD Operation and Developments in Run 2

The achievements of this thesis during the LHC Run 2 period are presented in this chapter. For readability, the chapter starts with a theoretical description of the Transition Radiation Detector (TRD), including the detector structure and the principles of its operation. In Run 2, the data collection is governed by hardware triggers. These triggers are able to inspect each event for interesting signatures prior to data recording. The TRD Run 2 setup is able to perform a preliminary track reconstruction for this purpose and provide a trigger decision within microseconds after the collision. In the context of this thesis, this functionality has been extended in order to identify light nuclei on this timescale with excellent performance. As the data recording can be focused on events containing nuclei, a significant enhancement in available interesting events is achieved.

The development of the nuclei trigger with the TRD along with prior feasibility studies and measurements of the enhancement factors for light nuclei using the triggered data are part of the achievements of this thesis.

The trigger has been operated in both pPb and pp collision systems. An outstanding result is the discovery of two  $\overline{{}_2^4\text{He}^{2+}}$  candidates within the pPb data collected with the

nuclei trigger. This observation is unprecedented in the small pPb collision system.

The introductory overview of the TRD setup in Run 2 is given in section 5.1. The section 5.1.3 contains a brief description of the TRD gas system service task carried out as part of this thesis. The achievements of this work concerning the developed nuclei trigger and the corresponding results are presented in section 5.2.

## 5.1. Transition Radiation Detector

The TRD is a gaseous detector which provides charged particle tracking and particle identification (PID) capabilities. Due to its fast online data processing capabilities and its low latency optimised readout tree, it can furthermore be used as a trigger detector to provide other detectors with physics triggers based on specific event signatures. Its cylindrical shape around the collision vertex allows for full azimuthal coverage and  $-0.9 \leq \eta \leq 0.9$  coverage in pseudorapidity. The full inner barrel coverage around the Time Projection Chamber (TPC) makes the TRD suitable for the calibration of the TPC mitigating the negative impact of space-charge distortions within the TPC.

This section is based on information from the References [2, 34, 36–41], which may be consulted for further details.

### 5.1.1. Detector Structure

The detector structure is organised into 18 Super Modules (SMs) along the  $\varphi$ -direction and each SM is further subdivided into 5 stacks along the beam direction. The stacks are further subdivided into 6 Readout Chambers (ROCs) each, constituting 6 layers along the radial direction. The full detector has support for a total of 540 ROCs, although 18 of them have never been installed in consideration of concerns regarding the material budget in front of the PHOS. This missing part of the TRD is referred to as PHOS hole in this thesis. A ROC, which is also referred to as a module, is the smallest unit of the

detector which can operate in standalone mode as there are no inter dependencies among the ROCs [2, 34, 40]. The detector structure is illustrated in Fig. 5.1.




Figure 5.1.: Structure of the TRD [34, 40].

### 5.1.2. Principle of Particle Tracking and PID inside a ROC

The particle detection principle inside a ROC is based on the ionisation of gas molecules in combination with the generation of Transition Radiation (TR). A particle which hits the detector from inside the cylindrical volume enclosed by the TRD first encounters a radiator material on the bottom of the ROC (compare Fig. 5.2). The radiator has dielectric properties, such that the speed of light inside the radiator medium is significantly lower than outside. A particle travelling through the medium polarises the molecules of the material locally. In case the particle travels at a velocity higher than the speed of light in the medium, the locally produced field fluctuations of neighbouring molecules cannot interfere destructively and a cone shaped wave front emerges. This happens as soon as the particle passes the boundary air/medium. Therefore, the generated photons

are referred to as transition radiation (TR). They are emitted within a small cone around the direction of flight of the initial particle and the effect gets more intense, the more relativistic (larger  $\gamma$  factor) the initial particle is. In practice, the cone size can be considered negligible, such that the generated photons travel into the direction of flight of the initial particle and contribute to the measurable energy loss behind this point. This helps distinguishing light charged particles from heavy ones at equal transverse momenta, which otherwise would only have less significant differences solely in their own Bethe-Bloch energy losses. In particular, this can be exploited to generate pure electron data samples by separating the electrons from the background formed by the more abundant pions [2, 34, 40].





Figure 5.2.: Principle of particle detection within a ROC [34, 40].

Behind the radiator, the particle travels through a drift chamber. The drift region has a length of 3 cm, across which an electric field is generated via high voltage (HV) power supplies. The drift chamber is entirely filled with a gas mixture composed of 85 %

Xenon (Xe) and 15 % carbon dioxide ( $\text{CO}_2$ ) of high purity. Particles travelling through the drift region, ionise gas molecules. The deliberated charges cannot recombine as they are immediately separated in the electric field and drift towards the electrodes. They drift at a constant drift velocity which is limited due to the interactions with the gas molecules and thus depends primarily on the gas mixture. To exploit the TR effect to the maximum extent, a short photon absorption length of the detector gas within the energy range of several keV of TR photons is highly desirable [2]. This motivates the use of Xe as primary component over the other nobel gases (compare Fig. 5.3).




**Figure 5.3.:** Photon absorption length of different nobel gases vs. energy of the impinging photon [2].

After passing through the drift region, the negative drift electrons enter an amplification region. The amplification region contains in its centre plane thin anode wires at positive High Voltage (HV) potential, while the bottom plane of the amplification region contains thin wires at ground potential. Similarly, the top plane of the amplification region consists of small separated readout pads at ground potential. This geometry is also called Multi Wire Proportional Chamber (MWPC) and the static electric field is largest in the close

vicinity of the thin anode wires. The closer a drift electron approaches an anode wire, the stronger it gets accelerated by the electric field. In this region, the energy which the electron gains between interactions with the gas molecules suffices to ionise the gas molecules. This process generates additional electron/ion pairs which contribute to the total electric current. Since the generated electrons can ionise additional gas molecules, this effect is also referred to as the avalanche effect and effectively leads to a strong signal amplification. The generated ions drift away from the anode. Partially, they drift towards the cathode wires or fall back into the drift chamber, but a significant fraction drifts towards the cathode readout pads, where they contribute to a measurable signal. The readout pads are connected to different channels of the readout electronics and allow for reconstructing the particle trajectory coordinates in the  $r\varphi$ - $z$  plane. During the entire drift time, charge deposition values are sampled in a configurable number of time bins, which allows the reconstruction to resolve the trajectory coordinates in the radial direction. In order to improve the resolution beyond the resolution provided solely by the readout pad size, the Pad Response Function (PRF) can be used, which describes the fact that the deposited charge is typically not concentrated in a single pad, but shared between adjacent pads, such that an interpolation of a more precise position is possible [2]. In beam direction  $z$ , the readout pads are larger than in the  $r\varphi$  direction. In order to improve the position resolution also in beam direction, the elongated pads are tilted by a small  $2^\circ$  angle (compare Fig. 5.4) with respect to the beam direction  $z$ . The sign of the tilting in a particular stack alternates from layer to layer. This tilting is exploited during reconstruction to improve the resolution. [2, 34, 40].

### 5.1.3. Gas system of the TRD (Service Task)

The gas mixture, ideally composed of 85 % Xenon (Xe) and 15 % carbon dioxide ( $\text{CO}_2$ ), is recirculated permanently through the TRD SMs. The TRD ROCs expose large outer surfaces and their thin walls make them susceptible to damage from small pressure differences between the inside gas mixture and the outside atmospheric pressure. The pressure difference which can be tolerated is in the order of only 1 mbar. Since the variations of the atmospheric pressure due to weather changes generally exceed this tolerance, the gas system has to regulate the pressure inside the detector by storing excess gas in an external



**Figure 5.4.:** Illustration of tilted readout pads in adjacent layers. This is used to improve the tracking resolution beyond the limitations imposed by the finite number of readout pads and channels [34, 40].

high pressure buffer [2, 41].

Another important aspect of the TRD gas system is maintaining the purity of the gas. Although the gas system is an entirely closed system, neither continuous small gas losses nor the contamination of the gas with impurities from the ambient can be avoided. The gas system provides means of automatically refilling any amount of gas which gets lost with fresh gas and it provides means to clean the gas from atmospheric impurities [2, 41].

The most important contaminants from the atmosphere are  $\text{O}_2$ ,  $\text{H}_2\text{O}$  (humidity) and  $\text{N}_2$ .  $\text{O}_2$  and  $\text{H}_2\text{O}$  are eliminated by the gas system automatically. The elimination is based on 2 purifiers which are able to absorb the contaminants and can afterwards be regenerated at high temperatures. In a recurrent cycle of 4 days, the 2 purifiers are alternated between absorption of contaminants and regeneration. The contamination of the gas mixture with  $\text{N}_2$  is far more problematic, as no automatised decontamination procedure is available [2, 41].

The removal of  $\text{N}_2$  is however possible by a manual procedure, which uses cryogenics and takes advantage of the difference in the boiling point between Xe ( $-108^\circ\text{C}$ ) and  $\text{N}_2$

( $-196^{\circ}\text{C}$ ). The contaminated gas mixture is pumped into a dewar which cooled down using liquid nitrogen. The desirable Xe and  $\text{CO}_2$  condensate inside the dewar while the contaminating  $\text{N}_2$  does not. The purified mixture of Xe and  $\text{CO}_2$  can afterwards be heated up and compressed into gas bottles. The gas bottles can later be used to refill the gas system with purified gas. Since this procedure involves the removal of the gas mixture from the detector and takes many days to complete, it can only be performed during experiment shutdown phases.

As part of the work for this thesis, a gas system service task has been completed within 2 years at a workshare of 25%. This service task has consisted of supervising and operating the TRD gas system which runs permanently on a 24 hours / 7 days in a week base and has required immediate actions in case of issues. In addition, the cryogenic purification procedure has been carried out during a Year-End Technical Stop (YETS) shutdown period.

#### 5.1.4. Front End Electronics (FEE) and data flow in a ROC during readout

The readout electronics of a ROC is organised in 6 to 8 readout boards (ROBs). They are separate printed circuit boards (PCBs). The exact number of ROBs depends on the location of the ROC within the detector, as they vary in size. Every ROB contains 16 mutli chip modules (MCMs), which constitute the first data sampling, shaping and processing stage. The MCM is an application-specific integrated circuit (ASIC), which has been specifically developed and manufactured for TRD readout. Its hybrid design integrates both analogue and digital processing steps as well as suitable Analogue-to-Digital converters (ADCs) in between. Every MCM provides 18 separate channels and collects the data from 18 adjacent readout pads [2, 39]. For details on the data processing in an MCM, refer to the following Section 5.1.5.

In addition to the MCMs providing the readout channels, the ROBs contain additional MCMs which act as data concentrator nodes. Every ROB contains a Board Merger (BM)



collecting the data from the other MCMs of the same ROB. The data from all BMs are merged at the half chamber (HC) level by the half chamber merger (HCM). The HCM pushes the data to the optical readout interface (ORI) board, which is a dedicated PCB attached to the HCs. Every ROC is equipped with two ORI boards, one for each HC. They provide data serialisation and encoding capabilities to push out the data through an optical link using a laser diode [2,38,39]. More details on the ORI can be found in Section 6.3.1.

Besides the data push out infrastructure within the FEE, there are additional and independent infrastructures to distribute trigger information to the MCMs and to communicate with the MCMs via slow control. These networks are controlled by the detector control system (DCS) board. The DCS board is an additional dedicated PCB mounted on the ROCs [2,36,39].

It receives triggers via the trigger and timing control (TTC) system from the CTP using a specific integrated circuit (IC), the TTCrx chip, which is common to all LHC experiments and also recovers the bunch crossing (BC) clock as well as the orbit signal locally from the LHC Prévessin site. Triggers are always synchronous to these clocks. The triggers are redistributed to all the MCMs of the ROC. In particular, this allows the MCMs to wake up in time and start data sampling after a collision has taken place. The corresponding trigger is therefore also referred to as the wake-up trigger [2,35,39].

The Slow Control Serial Network (SCSN) is used to exchange configuration and debugging data with the MCMs using low voltage differential signalling (LVDS). Differential signalling with two signal paths always carrying strictly opposite signal levels is more robust against common mode interference from external sources. As the receiver is only sensitive to the difference, potentially harmful interference gets compensated. The SCSN is used in particular during detector configuration before the start of data taking [2,36,39].

For the interface with the DCS and for convenience during debugging, the DCS board implements a full Linux system. In view of the limited resources of the DCS board, the memory footprint of the Linux system is kept as small as possible by using the BusyBox

Linux distribution. The DCS board is equipped with an Ethernet interface and it is integrated into the DCS network of the experiment. Most features of the DCS board can be accessed by command line for debugging. During detector operation, the interaction between the DCS and the DCS board occurs via a locally running Distributed Information Management System (DIM) server [2, 36, 39].

### 5.1.5. Data Sampling and Processing in a Multi Chip Module

As every MCM collects the data from 18 adjacent readout pads, it can already perform data processing and compression for the corresponding local detector segment. The concept of early data reduction is an important aspect of the whole low latency optimised and fast readout detector design. However, there is still the possibility to circumvent processing stages and to read out the full and unmodified raw data if required [2, 36, 39].

The data path through the MCM with the processing steps is illustrated in Fig. 5.5. The 18 readout pads, to which the MCM is connected to, are connected to the Preamplifier and Shaper (PASA) unit which constitutes the analogue circuit of the MCM. It measures the charge which has been deposited on the readout pads and creates a charge dependent output signal prescaled for the 10 bit input range of the following Analogue-to-Digital converter (ADC). The ADC has a maximum conversion rate of 10 MHz, effectively imposing a limit on the number of time bins which can be sampled throughout the drift time. The digitised, but unprocessed data directly behind the ADC, is often referred to as raw data and allows for the largest flexibility in terms of offline data analysis. It is written into an event buffer for a potential later complete readout. This event buffer can hold an entire single event. At the same time the data undergoes several digital processing steps. They are carried out by specialised custom ASIC logic, but the logic is kept highly configurable through slow control. In the first step digital filters can be applied, e.g. to subtract background noise or to compensate for different chamber gain factors. All filter can be disabled entirely and filter parameters can be configured through slow control. The filtered data enters then the main processing stage, which consists of tracklet preprocessor and tracklet processor. Tracklets are small track segments measured and reconstructed in a ROC by a single MCM. They result as a fit through detected charge clusters (hits) within

the drift volume. This early reconstruction eliminates the need to send single hits through the readout network. Instead, every found tracklet is represented by only one data word as small as 32 bit, which encodes all geometric parameters of the tracklet as well as the energy loss. Details on the tracklet data format can be found in Section 5.1.7. The local geometric slope of a tracklet is approximative for the curvature and hence for the rigidity  $p_T/z$  of the global trajectory due to bending in the magnetic field. Since high  $p_T$  particles are often more interesting for physics, the MCM is capable of applying a configurable cut on the rigidity. Throughout Run 2, this threshold has been set to 2 GeV/c. Every MCM can send a maximum of 4 tracklets through the network interface (NI) over the readout tree [2,36,39].

The bifurcation behind the ADC with its parallel processing and buffering of data reflects the trigger scheme of ALICE: The TRD is designed to provide data after both the LM/L0 and the L1 trigger levels. If an LM trigger is received, the data sampling starts and undergoes all digital processing steps. All processing steps are completed before the global decision by the CTP whether to issue an L1 trigger is due. Therefore the TRD is able to contribute actively to the L1 trigger decision. In case that the L1 trigger decision gets confirmed, the TRD will ship the raw data from the event buffer over the readout tree. This takes more time, but it is only necessary in case that the event is already to confirmed to exhibit interesting signatures [2,36,39]. Details on the timing can be found in Section 5.1.6.

### 5.1.6. Readout time windows

The typical timing of reading out an event is given in Fig. 5.6. The figure shows a typical drift time duration of 2  $\mu$ s. For latency optimisation, the processing pipeline is started in parallel at the earliest possible time. The total duration of building tracklets is around 3.75  $\mu$ s. These 3.75  $\mu$ s are the absolute minimum dead time of the TRD when reading out tracklets [2,36].

After the tracklet processing is finished, data shipping via the optical readout fibres starts. Fig. 5.6 shows an exemplary short shipping time of less than 1  $\mu$ s. In general, the

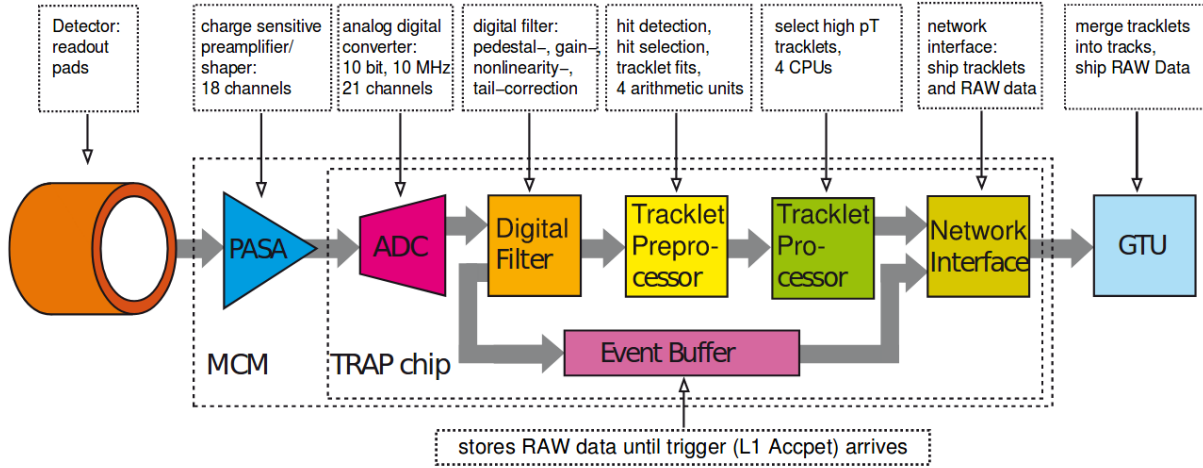


Figure 5.5.: Schematic processing steps in the FED [39].

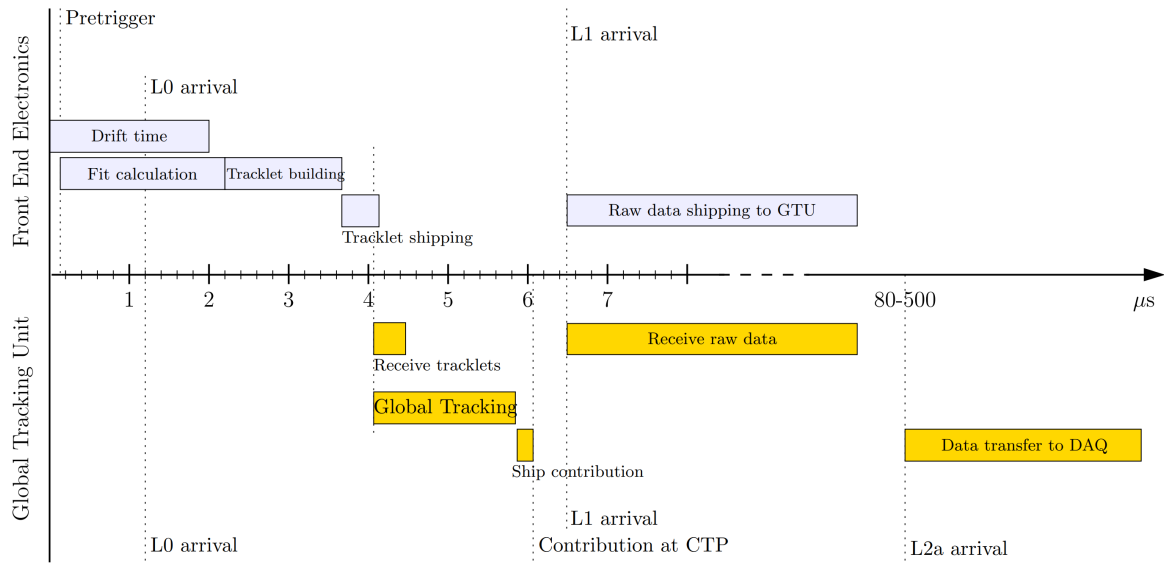


Figure 5.6.: Typical timing of the involved processing steps when reading out an event with the TRD [36].

shipping time however depends on the event multiplicity and the presented performance is typical in low multiplicity environments. For black events, a shipping time of several  $\mu\text{s}$  is necessary, the limited buffer sizes however guarantee that tracklet shipping completes

within the first 10  $\mu\text{s}$  after the collision [2, 36].

The shipped tracklets are buffered and further processed in the Global Tracking Unit (GTU). More details on the GTU are given in Section 5.1.8.

In case of a positive L1 trigger decision after about 6.5  $\mu\text{s}$ , the raw data is shipped via the optical links. The timing depends on the involved event sizes and may take several 100  $\mu\text{s}$  [2, 36].

### 5.1.7. Tracklet Data Format of the MCM

Since all data sampling and processing on the fast readout path has to fit into the short time window between the LM and the L1 trigger levels ( $\approx 6 \mu\text{s}$ ), the data format which represents a tracklet has to be carefully optimised. A compromise is needed between precision and the time needed for saving data, shipping and processing. This requirement excludes the possibility of relying on standardised data types like integer or float for the tracklet transmission. Instead a careful bit by bit optimisation is necessary [36, 42].

Therefore, the MCM ships a 32 bit wide word for every found tracklet which encodes the information shown in Table 5.1.

**Table 5.1.:** Format of the 32 bit wide word generated for every found tracklet by the MCM [42].

Name	Granularity	Range	Number of Bits
Axis intercept	160 $\mu\text{m}$	$-643.2 \text{ mm} - 643.2 \text{ mm}$	13
Deflection	140 $\mu\text{m}$	$-8.8 \text{ mm} - 8.8 \text{ mm}$	7
Pad row	1	0 – 15	4
Charge deposition	1/256	0 – 1	8

The axis intercept encodes the position of the tracklet in  $\varphi$  direction. The value is given as a distance between the ROC centre line at the outer surface in beam direction

and the crossing point of the particle trajectory with the outer ROC surface [42].

The deflection encodes the slope of the tracklet and is directly related to the rigidity of the particle. The value is given as the  $\varphi$  directional component of the distance between particle entry point at the ROC inner surface and the particle exit point at the ROC outer surface [42].

The pad row encodes the tracklet position in beam direction. The Multi Chip Modules (MCMs) do not communicate with neighbouring MCMs. The readout pads, to which a single MCM is connected to, are aligned in  $\varphi$  direction. This value thus depends solely on the mounting position of the transmitting MCM [42].

The charge deposition value corresponds to the charge deposited by the corresponding particle integrated over all time bins throughout the drift time. Although this field is named charge deposition in this thesis, there have been many attempts in the past to use it more efficiently in terms of electron identification. This has been motivated by the fact that only 8 bits, i.e. 256 distinct values, are available. The charge deposition of electrons is larger than the one of the more abundant pions due to TR, but the difference is small and might become invisible after scaling down the integrated charge deposition value (initially 16 bit) down to an 8 bit resolution. Therefore so called Lookup Tables (LUTs) have been introduced to map the initial 16 bit charge deposition to the 8 bit value for transmission using a configurable table. The LUT is configured during the MCM configuration by slow control and typically scales non linearly, trying to assign more resolution to the electron/pion overlap region, while reducing the resolution in other parts of the range. This can even be tuned such that the 8 bit value corresponds directly to an electron probability. Therefore, the entire field is named accordingly in some references [39, 42].

### 5.1.8. Global Tracking Unit

The GTU, which is based on Field Programmable Gate Arrays (FPGAs), is used in Run 2 as a global data concentrator and data processor node. It receives the data from all

ROCs of the TRD which are delivered via 1044 unidirectional optical links (1 link per HC). The processed data are forwarded to the Data Acquisition System (DAQ).

In addition, the GTU is capable of performing preliminary online reconstructions of the tracks in an event and generate TRD trigger contributions based on certain conditions. They are sent to the CTP via Low Voltage Differential Signalling (LVDS). The GTU can receive copies of the L0, L1, L2 accept and L2 reject triggers from the CTP for controlling the TRD readout. More details will be given in the following sections.

This section is based on information from the References [34, 37, 38, 43, 44], which may be consulted for further details.

## Structure of the GTU

The GTU consists of 90 Track Matching Units (TMUs), 18 Super Module Units (SMUs) and 1 Trigger Generation Unit (TGU). These 3 units have different Printed Circuit Board (PCB) designs, but are all based around a Xilinx Virtex-4 XC4VFX40 FF1152 FPGA. In addition, there are backplane PCBs for interconnecting all PCBs and power supplies. All components are housed in 3 dedicated racks. The GTU structure is hierarchical and reflects the hierarchic structure of the TRD: every TMU receives the data from 1 TRD stack and also houses the necessary optical receiver modules. 5 TMUs each form a group and serve 1 entire SM. Such 5 TMUs are connected to an SMU which concentrates their data and provides it to the Data Acquisition (DAQ) on a dedicated optical output. The TGU is responsible for concentrating the trigger contributions generated in the other units and sending them to the CTP. It also provides the necessary electrical outputs [34].

The hierarchic and parallel GTU structure also reflects the design idea of facilitating highly parallel data processing and thus enabling low latency trigger decisions.

## TMU

Each TMU is connected to a particular stack. After an L0 trigger, the TMU receives the tracklet data generated by the HCs of that stack and uses it to perform fast online tracking. The tracking algorithm starts the processing already at the arrival of the first tracklet in order to optimise latency [34, 43, 44].

The TMU groups geometrically matching tracklets from different layers together. The matches are identified by projecting all tracklets on a common plane which is perpendicular to the radial direction and identifying geometrical clusters. A minimum of 4 contributing layers in a cluster is required to use the cluster for online tracking and to reconstruct a GTU track from it. The threshold ensures a reasonable track quality. Tracklets which do not have matching counterparts in at least 3 other layers, are not considered for online tracking [34, 43, 44].

The online reconstruction is performed by fitting a straight line through the contributing tracklets. This allows for a significantly better estimation of the track inclination and position than it has been possible in the FEE using a single tracklet only. The approximation of the circular trajectory as a straight line within the space region covered by the TRD stack is sufficient for many purposes [34, 43, 44].

Besides the geometric GTU track position, the online reconstruction algorithm calculates [34, 43, 44]:

- The arithmetic average of the charge deposition values of the contributing tracklets of the GTU track.
- The online transverse momentum of the GTU track  $p_{T,GTU}$  under the assumption that the particle trajectory originates from the primary vertex. Note that for particle not originating from the primary vertex, the GTU online transverse momentum  $p_{T,GTU}$  may not be representative.



- The sagitta of the trajectory, which allows for a direct estimate of the actual track curvature, independent of the position of the vertex. This parameter can be used to reject tracks from secondary vertices (Late Conversion Rejection (LCR)).

These track parameters are vital for physics and their exceptional early availability before the L1 trigger decision allows the generation of TRD L1 physics trigger contributions for an efficient online event selection.

## SMU

The SMUs act as data concentrator nodes. Every SMU receives data from 5 associated TMUs through electrical back plane connectors. Both readout data and trigger contributions are concentrated by the SMUs.

The SMU is equipped with optical outputs which connect to the DAQ to ship out the data from the connected TMUs for further processing and final storage.

Furthermore the SMUs support multi-event buffering. For that, the SMU FPGA communicates with an external Static Random Access Memory (SRAM) on the PCB and buffers events in this memory until they can be shipped. This feature effectively decreases the detector dead time as the detector can start to sample new events, even though the previous event is still in the readout chain. The details can be found in [\[43\]](#).

## TGU

The TGU acts as a concentrator node to collect relevant trigger information from the 18 SMUs and generates the final L1 trigger contributions for the CTP. For every possible trigger contribution, a dedicated LVDS cable between the TGU and the CTP is in place.

For most trigger types, there is no additional processing required in the TGU and the L1 trigger contribution for the CTP is obtained from the SMU contributions by applying an OR-function. However, the TGU is equipped with the same FPGA model as used for

the TMUs and SMUs and provides resources to allow more complex trigger types, which correlate the SMU trigger contributions in a sophisticated manner. A possible application is a trigger on particles from the cosmic radiation, which impinge on a straight track and thus generate correlated signals in 2 opposed SMs on both detector sides [34, 43, 44].

## TRD Triggers

In Run 2, prior to the development of the nuclei trigger in the context of this thesis, the TRD has been already providing the following L1 trigger contributions [34, 43–45]:

- **TRD Single Electron Trigger (HSE)**

This trigger intends to fire when an electron with high transverse momentum is detected. The focus is on the transverse momentum, not charge deposition, when compared to the HQU trigger, however, there is an overlap region between both triggers. The condition for a firing the HSE trigger is the occurrence of at least 1 track in any stack fulfilling all of the following conditions:

- The track has been identified in at least 5 TRD layers.
- The track has been identified in the TRD layer 0 (vertex side of the detector).
- The PID value averaged across all participating layers is  $\geq 120$ .
- The GTU transverse momentum calculated by the TMU is  $\geq 3 \text{ GeV}/c$ .
- The LCR cut parameter is  $\leq 0.2 c/\text{GeV}$ .

- **TRD Quarkonia Trigger (HQU)**

This trigger intends to fire when an electron with high transverse momentum is detected. The focus is on the charge deposition, high transverse momentum is secondary, when compared to the HSE trigger. The trigger is supposed to enhanced electrons which result from heavy quarkonia decays. The condition for a firing of the HQU trigger is the occurrence of at least 1 track in any stack fulfilling all of the following conditions:

- The track has been identified in at least 5 TRD layers.

- The track has been identified in the TRD layer 0 (vertex side of the detector).
- The PID value averaged across all participating layers is  $\geq 130$ .
- The GTU transverse momentum calculated by the TMU is  $\geq 2 \text{ GeV}/c$ .
- The LCR cut parameter is  $\leq 0.2 c/\text{GeV}$ .

- **TRD Jet Trigger (HJT)**

This trigger intends to enhance statistics in jets, which are showers of particles collimated within a limited cone size. The jet trigger operates on a stack level and the stack size is connected to the considered jet cone size. The condition for a firing of the HJT trigger is that at least 1 stack exists which identified at least 3 particle tracks, where each of the tracks fulfils all of the following conditions:

- The track has been identified in at least 5 TRD layers.
- The track has been identified in the TRD layer 0 (vertex side of the detector).
- The GTU transverse momentum calculated by the TMU is  $\geq 2 \text{ GeV}/c$ .
- The LCR cut parameter is  $\leq 0.2 c/\text{GeV}$ .

Within the scope of this work, this list has been extended by the nuclei trigger. The details on the nuclei trigger are given in the next Section 5.2.

## 5.2. Transition Radiation Detector Nuclei Trigger

In the context of this thesis, the TRD has been extended with an additional trigger type, the nuclei trigger HNU. This trigger is capable of significantly enhancing the abundance of light nuclei, such as deuterons, tritons,  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  particles and their antiparticles. It has been widely used since its commissioning in late 2016 until the end of Run 2 data taking in 2018. Today, there are nuclei triggered data sets from a pPb at  $\sqrt{s_{\text{NN}}} = 8 \text{ TeV}$  period as well as pp at  $\sqrt{s} = 13 \text{ TeV}$  periods available for use by data analysers. For instance, the measurement of the production of hypertritons and antihypertritons in pp collisions at  $\sqrt{s} = 13 \text{ TeV}$  using events selected by the TRD nuclei trigger HNU is in the

process of publication. The latest results have been presented in Ref. [46].

The first development and implementation of the nuclei trigger based on the TRD has been done in the context of this thesis. The possibility of a nuclei trigger development had been already considered earlier and first feasibility studies are e.g. presented in Ref. [47]. However, the nuclei trigger had never been implemented previously. In this thesis, improved feasibility studies and detailed optimisations are presented. These results have been used for the hardware realisation of the nuclei trigger in the context of this thesis and have been further used during data taking.

This section outlines the development of the new TRD nuclei trigger from first feasibility studies to the hardware implementation and the commissioning in the experiment. Finally, the actual trigger performance is verified by measuring it based on the nuclei triggered data sets in pPb and pp collisions.

### 5.2.1. Motivation of using the TRD for the nuclei trigger

Due to their large masses, nuclei are rarely produced in a collision. Especially in collision systems with low average multiplicities like pp and pPb, physics analyses targeting nuclei, have to deal with very low available statistics from minimum bias data taking. Therefore, an enhancement of events containing these nuclei is desired.

Considering the available TRD hardware, the TRD appears suitable for the implementation of a nuclei trigger due to its following properties:

- The TRD is capable of measuring the particle energy loss within the detector gas. Nuclei are expected to undergo excessive energy loss compared to other more abundant particles at low transverse momentum  $p_T$ . Especially in case of the  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  species, there is in addition the doubled charge  $Z = 2$ , which is expected to quadruple the energy loss compared to the very abundant  $Z = 1$

particles and should thus produce a unique signature. For more details, refer to the Bethe-Bloch formula, described in appendix B.2.

- The TRD is equipped with powerful hardware which allows for a preliminary online reconstruction of the particle tracks. It has all the required data, including the energy loss signature, readily available prior to the L1 trigger decision. The hardware has proven its capabilities by providing the already existing TRD L1 trigger contributions and the necessary adaptations for an additional trigger contribution at the same latency might be feasible with specific strategic hardware modifications, but without a major upgrade or other disruption of the Run 2 data taking activities at that time.

### 5.2.2. Feasibility studies for the nuclei trigger implementation

Although the envisaged trigger targets the low-multiplicity pp collision system, most studies, which have been carried out before the actual implementation of the trigger, have been done using Pb–Pb data. Data sets from small collision systems did not provide enough statistics of nuclei tracks to estimate the trigger performance. The studies presented in this section mainly investigate the detector response at the traverse of given particles. The detector response for a given particle should in principle not depend on the collision system in which the particle has been originally created. Following this argument, an arbitrary choice of the collision system for the feasibility studies appears justified and should not affect the results significantly. There are limitations of this assumption though, as the angle of incidence of the particle on the detector does play a role for its response. The distribution of the angles of incidence depends on the collision system and becomes asymmetric for asymmetric collision systems.


#### Feasibility studies based on existing Monte Carlo data

**Data sets** For this study, the Monte Carlo (MC)-production LHC14a6 anchored to the the Pb–Pb at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV period LHC14a and the MC-production LHC14b3a anchored to the pPb at  $\sqrt{s_{\text{NN}}} = 5.02$  TeV period LHC14b are used. The LHC14a6 MC-production has been produced with artificial Hijing nuclei injection. Due to its large statistics, it thus proves ideal for detector response studies. Similarly, the LHC14b3a


MC-production contains an enhancement of electrons due to the use of both TRD and ECal electron triggers during the production of this MC. It thus proves ideal for the study of the detector response of electrons. Electrons may be a predominant source of signal contamination for the trigger, as the electron signal in the TRD is increased due to additional transition radiation. Although the use of radiator material to ease electron identification is a fundamental design principle of the TRD, it may negatively impact the performance of the nuclei trigger.

Figures 5.7 and 5.8 show the  $p_T$ -dependent abundance of different particles in the considered MC-productions. These results have been obtained based on the MC-truth labels. The artificial and almost uniform injection of deuterons, tritons,  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  as well as their respective antiparticles up to a maximum transverse momentum of  $p_T = 10 \text{ GeV}/c$  is clearly visible in Fig. 5.7. In contrast, the LHC14b3a MC-production shown in Fig. 5.8 without artificial nuclei injection from the pPb collision system does not provide sufficient statistics in nuclei. However, the much larger abundance of electrons and positrons in the triggered LHC14b3a data set is visible when comparing the two figures.


**Effects due to altered TRD installation** When using data sets from the year 2014, it has to be accounted for the fact that the TRD had not been fully installed at that time. As opposed to the time of the nuclei trigger development, the TRD did not cover full azimuth and only the sectors 0, 1, 7, 8, 9, 10, 11, 15, 16 and 17 had been available. The installation with 10 fully available super modules, had been established in 2011. Considering that even in the final TRD installation, the sectors 13, 14 and 15 miss their middle stack to avoid shielding the PHOS detector, in total 49 stacks have been available in the year 2014. At the time of nuclei trigger commissioning, the full TRD is available with 87 stacks. From this, a correction factor  $87/49$  can be derived by which all detector acceptance efficiencies have to be multiplied to reflect the actual expected detector performance. For illustration of this effect, Fig. 5.9 shows the distribution of TRD tracks across the detector super modules and detector stacks. The blind spots due to missing super modules are clearly visible.



**Figure 5.7.:** MC Production LHC14a6, Pb-Pb collisions at  $\sqrt{s_{NN}} = 2.76$  TeV Hijing nuclei injection for sufficient statistics, especially with respect to  ${}^3_2\text{He}^{2+}$ ,  ${}^4_2\text{He}^{2+}$  and their respective antiparticles.



**Figure 5.8.:** MC Production LHC14b3a, pPb collisions at  $\sqrt{s_{NN}} = 5.02$  TeV, triggered TRD and EMCal data for sufficient electron statistics.

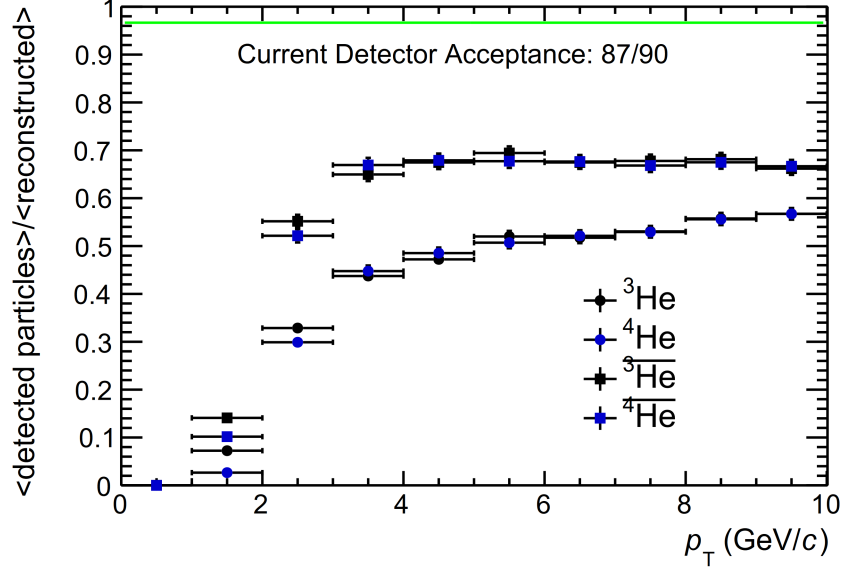


**Figure 5.9.:** Distribution of TRD tracks in the MC production LHC14a6 across sectors and stacks in the TRD.

**Acceptance efficiency for  $Z = 2$  particles** The trigger optimisation will be done for the most desirable  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  nuclei. For the evaluation of the TRD *acceptance efficiency* for these species, Event Storage Data (ESD) tracks of  ${}^3_2\text{He}^{2+}$ ,  ${}^4_2\text{He}^{2+}$  and their respective antiparticles have been selected based on the available MC truth information. A track matching between these ESD tracks and the available TRD tracks has been conducted by matching the MC labels. Due to detector acceptance limitations, not every ESD track can be attributed to a TRD track. The ratio of successful matches is the detector *acceptance efficiency* which generally depends on the transverse momentum  $p_T$ . Note that in this work the term *acceptance efficiency* is defined as the number of detected tracks in the TRD divided by the number of globally reconstructed tracks. This *acceptance efficiency* includes effects due to limitations in the geometric acceptance and also all other effects, as long as they are not related to triggering. The  $p_T$  dependent *acceptance efficiency* is shown in Fig. 5.10.

Figure 5.10 reveals that the efficiency for very low transverse momentum  $p_T \leq 0.5 \text{ GeV}/c$   $Z = 2$  particles is zero and rises up to 68% with increasing  $p_T$ . For very low transverse momentum with  $p_T \leq 0.5 \text{ GeV}/c$ , no TRD tracks are expected, as these particles are slow





**Figure 5.10.:** Fraction of  ${}^3_2\text{He}^{2+}$ ,  ${}^4_2\text{He}^{2+}$  and their respective antiparticles measured by TRD with respect to all reconstructed tracks of the corresponding particle species. Statistical uncertainties are shown as vertical bars. Small uncertainties are hidden by the point markers.

and the radii of their trajectories in the solenoid magnetic field are too small to reach the TRD. For increasing  $p_T$ , this effect decreases and the *acceptance efficiency* improves. It flattens for  $p_T \geq 6 \text{ GeV/c}$ , above this level the majority of particles reaches the TRD.

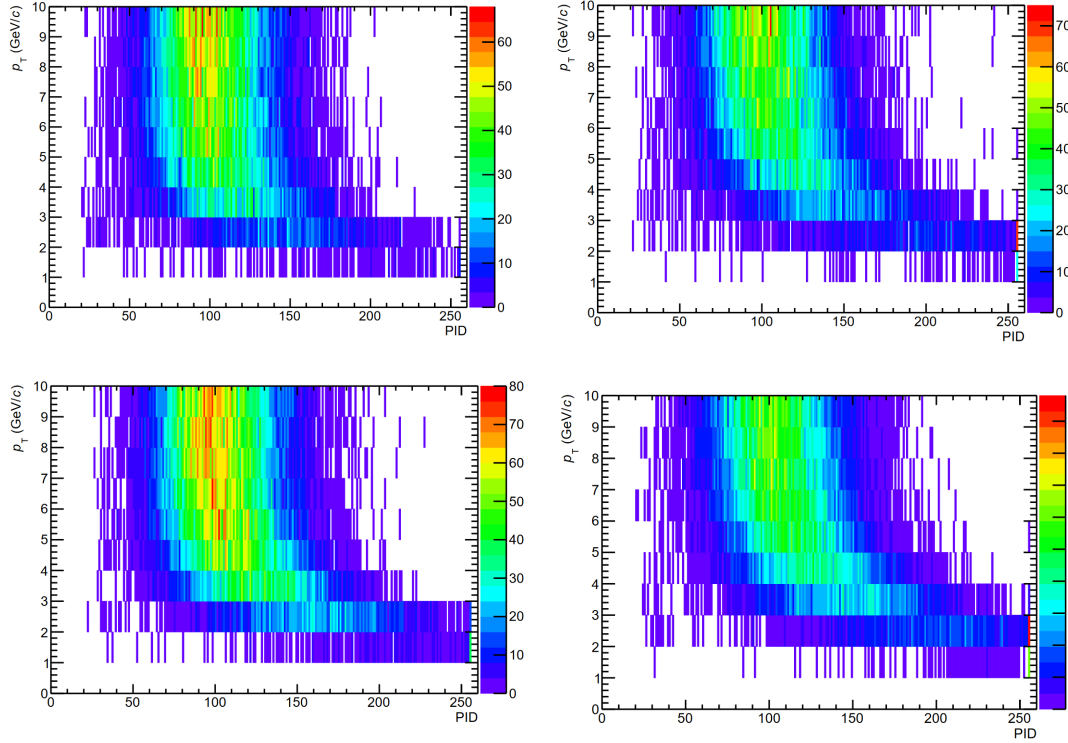
For  $p_T \leq 3 \text{ GeV/c}$ , a significantly larger *acceptance efficiency* for  ${}^3_2\text{He}^{2+}$  is observed when compared to the *acceptance efficiency* for  ${}^4_2\text{He}^{2+}$  within the same  $p_T$  bin. The effect gets stronger the lower the  $p_T$  and gets insignificant for  $p_T \geq 3 \text{ GeV/c}$ . The same observation can be made when comparing  $\overline{{}^3_2\text{He}^{2+}}$  with  $\overline{{}^4_2\text{He}^{2+}}$ .

Across all  $p_T$  bins it can be observed that the negatively charged anti-nuclei exhibit a higher *acceptance efficiency* than their positively charged nuclei counterparts. This is an effect which can be observed in the TRD for all pairs of particles with opposed charges. It is due to the fact that the drift charges which are liberated within the TRD gas are not only influenced by the electric field which makes them drift, but also by the outer solenoid magnetic field. The magnetic field is perpendicular to the electric field, thus resulting in

a Lorentz force on the drift charges. The drift current, the magnetic field and thus the resulting Lorentz force always show into the same direction, independent from whether the drift process has been initiated by a positive or negative traversing particle. Due to the modularity of the TRD structure and fact that distinct MCMs work independently and do not perform tracking together in the border regions, the TRD tracking can fail on tracks which are effectively too much inclined and localised in unfortunate positions, such that the same track traverses regions handled by different MCMs. The probability for the tracking to fail therefore increases with the track inclination corrected by the additional inclination introduced by the Lorentz force on the drift charges. The direction into which the charged particle trajectories are bent in the magnetic field depends on the sign of the charge. The same applies for the direction of the impinging angle of the particles onto the inner TRD surface. In case of negatively charged particles, the Lorentz force on the drift charges effectively makes the measured tracklets be more straight, such that they have a higher probability to be registered in the TRD, while for positively charged particles, the Lorentz force increases the effective inclination, worsening the performance. This effect is also referred to as the  $\vec{E} \times \vec{B}$  effect.

**Efficiency of the trigger condition for  $Z = 2$  particles** The layers of the TRD in which a traversing particle is registered measure the energy deposition within the layer and generate unsigned 8-bit numbers which is representative for the energy losses. The idea of the nuclei trigger operation is to apply an online cut on these values, such that only events which contain particles which energy deposition above the threshold are recorded. The 8-bit energy deposition values of the participating layers are gathered at the level of the GTU and an 8-bit arithmetic average is calculated for every particle track, in the following called the PID value.

For studying the properties of the trigger condition, data samples with  ${}^3_2\text{He}^{2+}$ ,  $\overline{{}^3_2\text{He}^{2+}}$ ,  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  have been selected from the MC production, requiring the condition that the particle track has been registered in both, the TPC and the TRD. The matching of global tracks to TRD tracks is here based on the MC truth labels.



**Figure 5.11.:** PID value as a function of  $p_T$ . Top left panel:  ${}^3_2\text{He}^{2+}$ . Top right panel:  ${}^4_2\text{He}^{2+}$ . Bottom left panel:  ${}^3_2\text{He}^{2+}$ . Bottom right panel:  ${}^4_2\text{He}^{2+}$ .

For the different particles species, 2-dimensional histograms illustrating the distribution of PID values along with the corresponding track  $p_T$  from the global reconstruction have been generated. The histograms are shown in Fig. 5.11.

Figure 5.11 shows that the energy loss generally increases as  $p_T$  decreases. The nuclei of interest exhibit a broad spectrum of possible energy losses in the region of  $p_T > 4 \text{ GeV/c}$  and the 8-bit value ranges between approximately 30 and 200. In this region, energy loss based triggering on nuclei is not possible, as the broad energy loss spectrum may overlap with almost any other particle. Note that in the  $p_T > 4 \text{ GeV/c}$  region, even an offline nuclei identification with the TPC is difficult and suffers from severe contamination with other particles.

In the region  $p_T < 4 \text{ GeV}/c$ , a strong shift of the PID distribution mean towards larger  $\text{PID} > 200$  can be observed. For the heavier  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$ , the effect is even stronger. In this region triggering on nuclei appears to be promising.

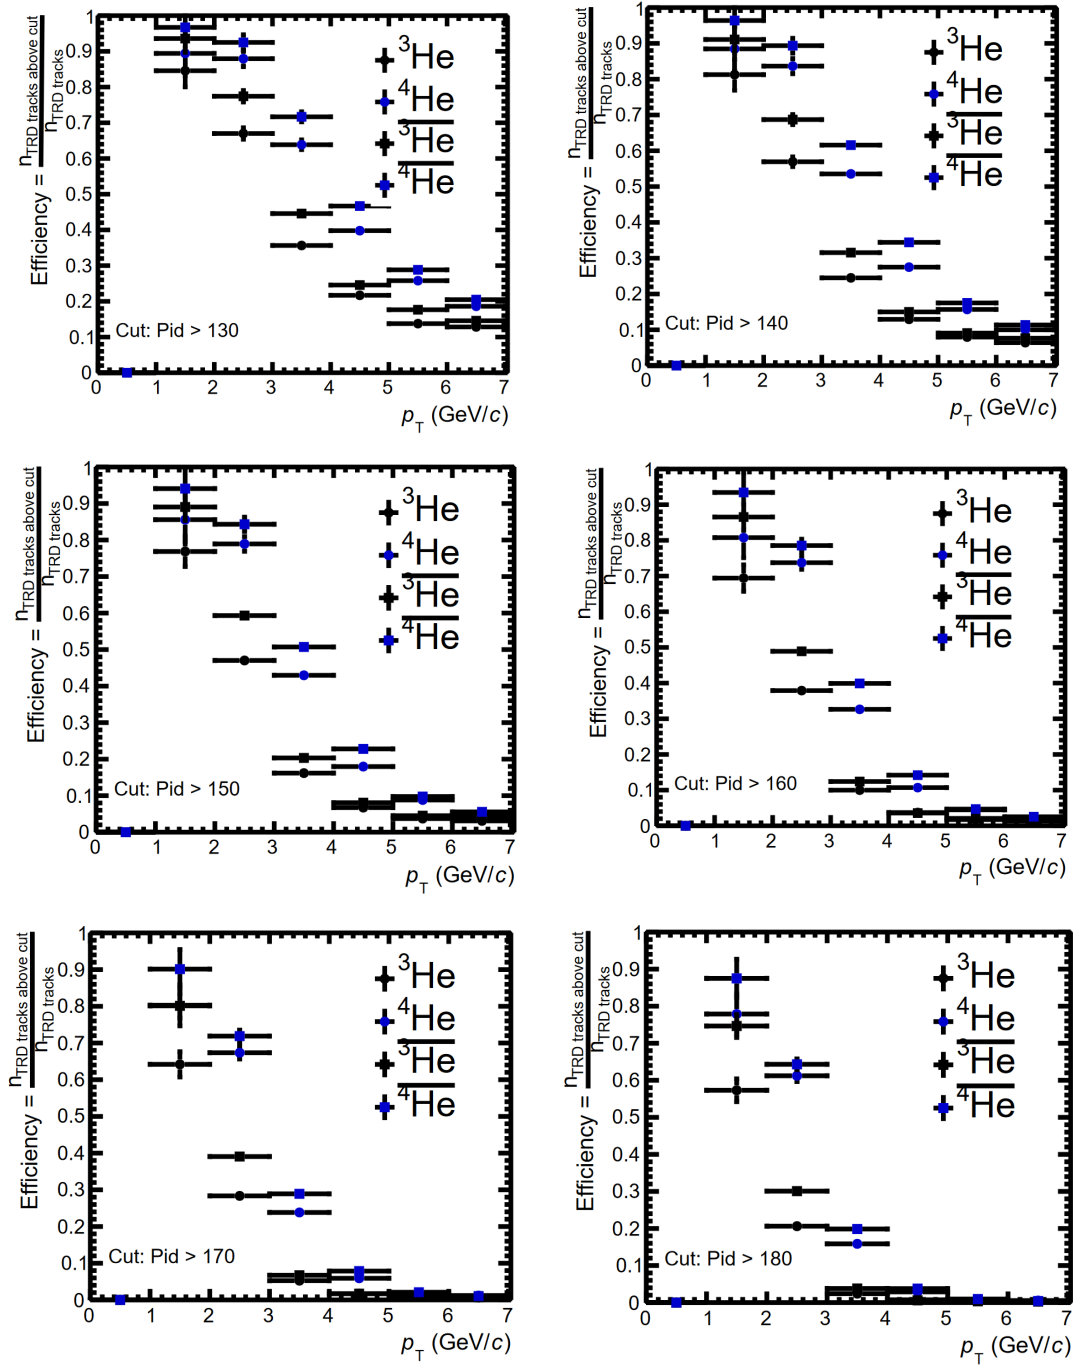
The fraction of nuclei with TRD tracks which get accepted by the trigger has been investigated separately for the different nuclei species and for different PID thresholds. The resulting fractions, in the following called trigger condition efficiencies, are plotted for exemplary thresholds in Fig. 5.12. Statistical uncertainties are shown as vertical bars.

Figure 5.12 shows quantitatively, that significant trigger condition efficiencies can be achieved in the region  $p_T < 4 \text{ GeV}/c$ . The efficiency of the trigger is in general larger for lower  $p_T$  and a lower PID threshold. In addition, the trigger is generally more efficient for  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  than for the lighter  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$ . For the same particle types, a small difference in the trigger condition efficiency can be observed between particles and their corresponding antiparticles. For the antiparticle, the trigger condition is slightly more efficient.

### Consideration of different LUTs

Since the PID values determine significantly the performance of the trigger, it is important to consider the details on how these values are generated by the TRD Front End Electronics (FEE).

When the MCM samples the amplified drift charges of a tracklet, the charge deposition raw value is converted to digital by an Analogue-to-Digital converter (ADC) with a 16-bit resolution. This corresponds to a numeric range from 0 to 65535 and the obtained value is truly proportional to the energy loss of the particle. However, during the optimisation of the TRD for fast triggering during the FEE design, it has been decided to not ship the 16-bit ADC raw value within the tracklet, but to shorten it and only ship 8 bits as a PID value, effectively limiting the range of possible numeric values from 0 to 255. Both the GTU, which is responsible for evaluating the trigger condition, and the tracklet data in



**Figure 5.12.:** Trigger condition efficiency as a function of  $p_T$ , shown separately for the different  $Z = 2$  nuclei for various PID cuts. Statistical uncertainties are shown as vertical bars.

the final storage are concerned by the limitation of having only a range from 0 to 255 available as PID value.

For translating the original 16-bit ADC value into the 8-bit number, different strategies have been used previously by the TRD team. Technically, the translation is performed on the TRD FEE using a LUT. The LUT is a static list of 256 16-bit values, mapping uniquely each 8-bit value to a certain subrange of the 16-bit value range. The LUT is fixed during the creation of the TRD FEE configuration and gets loaded into the MCMs via the Detector Control System (DCS) network at FEE configuration time.

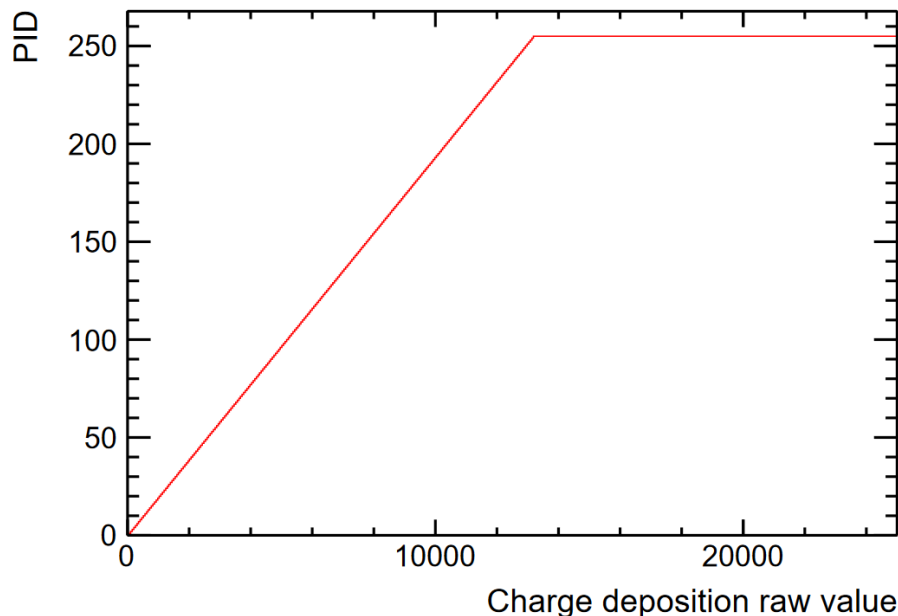
Since the LUTs may differ between different runs, different LHC periods and different MC productions, while only the 8-bit values are stored in the data, only trigger performance analyses which rely on data with the same LUT can be compared. During the previous data taking periods, 2 principle strategies have been widely used to determine the LUT:

- **Linear LUT**

The relation between the 16-bit value and the 8-bit value is linear up to a certain threshold of a reasonable maximum energy loss. Particles exceeding this threshold get an 8-bit value of 0xFF (overflow bin) assigned. An example of this type of LUT is given in Fig. 5.13.

- **Electron probability LUT**

This LUT aims at facilitating the identification of electrons by electron triggers and is designed such that the 8-bit represents no longer a direct measure for the energy loss, but it represents the probability that the tracklet is part of an electron or positron track. An example of this type of LUT is given in Fig. 5.14. Electrons and positrons have characteristic energy losses within a certain range, but electron triggers suffer from pion contamination within the same energy loss range. Pions have a smaller energy loss signature than electrons, but the separation is small and there is an overlap. In order to be able to distinguish electrons from pions at the level of the GTU, the LUT is designed such that a major subrange of the full 8-bit range maps to a comparatively small subrange of the original 16-bit range in the

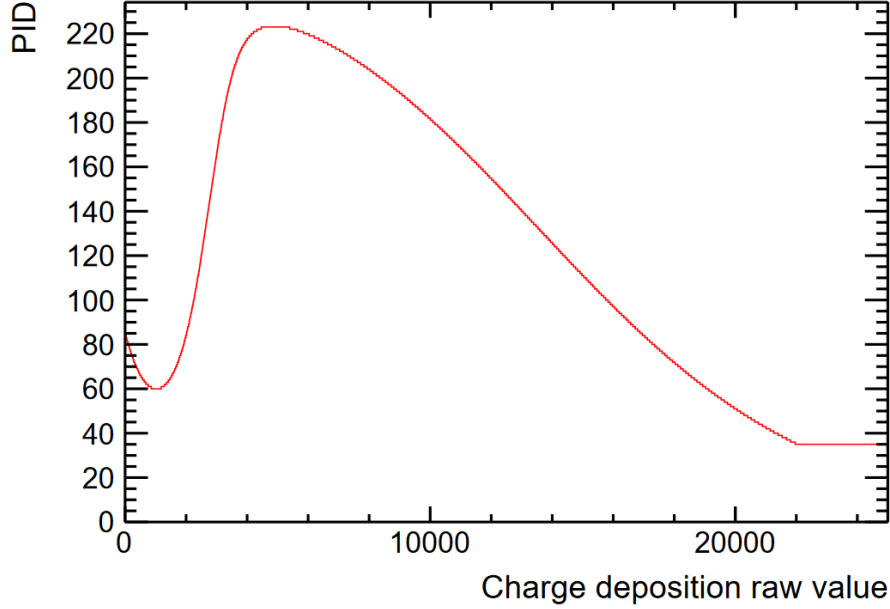


**Figure 5.13.:** Shipped 8-bit value as a function of the measured 16-bit raw ADC value (linear LUT).

electron/pion overlap region. This explains the steep rise of the curve in Fig. 5.14. After the maximum of the LUT, the probability drops again, as electrons - as opposed to nuclei - do not exhibit such large energy loss signatures.

The electron probability LUT is disadvantageous for triggering on nuclei because the mapping is not monotonous and after applying the LUT it is not possible to distinguish the rare particles with a large energy loss from the very abundant particles in the electron/pion overlap region. This is problematic for both, the actual online trigger and for offline studies of the nuclei trigger performance. Concerning the data recorded in the past, only data and MC production which have been recorded with the linear LUT can be used.

The nuclei-enhanced MC production LHC14a6 (compare Section 5.2.2) fulfils the condition of LUT linearity, while the electron-enhanced MC production LHC14b3a has been found to use the electron probability based LUT. An analysis of the trigger contamination with an MC production aimed at electrons would thus not be representative. Therefore



**Figure 5.14.:** Shipped 8-bit value as a function of the measured 16-bit raw ADC value (electron probability LUT).

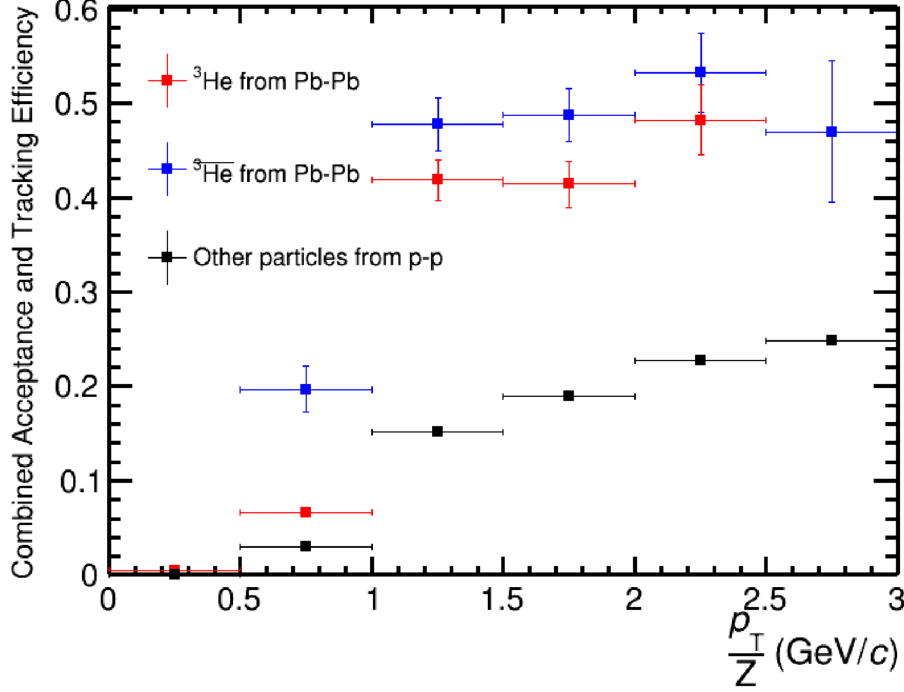
and for an even more realistic approach, the trigger contamination will be investigated later using real pp data within the scope of this thesis.

### Feasibility studies based on existing real data

For this study, the nuclei tracks have been collected from the Pb–Pb data taking in LHC11h, while contaminating tracks have been taken from pp data taking in LHC12a. Both data sets have been recorded with the linear LUT.

On the LHC11h data set, the nuclei of interest have been identified using the TPC. The Bethe-Bloch parametrisations for  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  are not available within the ALICE analysis framework AliPhysics and have been taken instead from Ref. [48]. Also, the track selection cuts have been adapted from Ref. [48]. The optimisations in this section are mostly based on  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  only, as for  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  the available statistics is often too low for conclusive results. However, as shown in Section 5.2.2, the expected





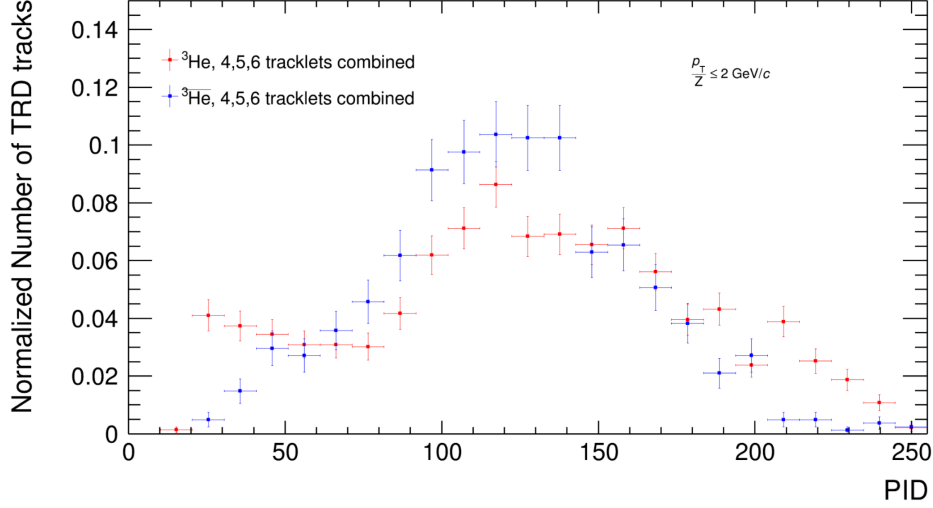
**Figure 5.15.:** *Acceptance efficiency* as a function of  $p_T/Z$  of nuclei in Pb–Pb collisions at  $\sqrt{s_{NN}} = 2.76$  TeV compared to the average *acceptance efficiency* in pp collision at  $\sqrt{s} = 8$  TeV. Statistical uncertainties are shown as vertical bars.

trigger performances for  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  are superior to those for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$ .

For the LHC12a pp data set, all particle tracks without any selection cuts are used in order to keep the contamination estimate as conservative as possible.

**Acceptance efficiency** Figure 5.15 shows the ratio between particle tracks with a TRD track and those identified by the TPC. The  $p_T$  dependent ratio is shown separately for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as well as for the pp particle mix.

For  ${}^3_2\text{He}^{2+}$ , the *acceptance efficiency* values are in reasonable agreement with the MC results in Fig. 5.10. For the same  $p_T$   ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  are about twice as efficiently accepted by the TRD than the particles from the pp particle mix.




**Figure 5.16.:** Normalised distribution of PID values for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV. Statistical uncertainties are shown as vertical bars.

**Distribution of PID values** The  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  particles which have been identified by the TPC have been matched to TRD tracks based on proximity in space. For this study only particle tracks in the target  $p_{\text{T}}$  range with  $p_{\text{T}}/Z \leq 2$  GeV/c are considered. Figure 5.16 shows a normalised distribution of the PID values measured by the TRD for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  particles. Statistical uncertainties are shown as vertical bars.

Figure 5.17 shows a reverse integrated version of the result in Fig. 5.16, i.e. for a given PID threshold, the y-axis indicates the fraction of particles which pass the threshold, also called the trigger condition efficiency. Statistical errors are shown as vertical bars. In Fig. 5.17 the statistical errors bars are small and hidden by the point markers.

When comparing the results for the trigger condition efficiency in real data (Fig. 5.17) with the trigger condition efficiencies obtained in MC (Fig. 5.12), it can be observed that the trigger performance in MC is slightly overestimated. The reason for the overestimation might be attributed to the particle transportation code in the GEANT framework, which is used in ALICE during the generation of the MC productions and is not well adapted to the




**Figure 5.17.:** Reverse integrated charge deposition distributions of  ${}^3\text{He}^{2+}$  and  $\overline{{}^3\text{He}^{2+}}$  in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV, which corresponds to the trigger condition efficiency as a function of the PID threshold. Statistical uncertainties are shown as vertical bars. Small error bars are hidden by the point markers.

simulation of the transition radiation detector response for rare nuclei with exceptionally large energy depositions.

**Consideration of the PID values of contaminating particles** Figure 5.18 shows the normalised distributions of the PID values of  ${}^3\text{He}^{2+}$  and  ${}^4\text{He}^{2+}$ , similar to Fig. 5.16, but on a logarithmic scale. In addition to the distributions for  ${}^3\text{He}^{2+}$  and  ${}^4\text{He}^{2+}$ , Fig. 5.18 also contains the normalised PID distribution for all particles from the pp data taking in LHC12a for a quantitative background estimate. The background distributions are further subdivided into 3 different cases: particle tracks which have been registered in all 6 layers of the TRD and tracks with only 5 or 4 tracklets in the TRD, respectively.


For the background distributions, it can be observed that the vast majority of background particles have very low PID values which are significantly lower than the typical



**Figure 5.18.:** Normalised PID distribution of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV compared to the normalised background spectrum in pp collision at  $\sqrt{s} = 8$  TeV. Statistical uncertainties are shown as vertical bars.

PID value range of the target nuclei. The trigger sample is only contaminated by the tail of the background particle distribution. Furthermore, it can be observed that the potentially problematic tail of the PID distribution of background particles depends strongly on how many tracklets the average PID is calculated. The less tracklets are available for calculating the average PID on the GTU, the more prominent is the tail of the background PID distribution and the more it interferes with the signal.

The strong dependence of the background PID distribution on the track quality (4, 5 or 6 tracklet cases) motivates a more sophisticated trigger condition. Instead of defining a single PID threshold for all cases, it is viable to define 3 different PID thresholds for the 3 different track qualities. The track qualities are known at the level of the GTU, so that the hardware implementation for this more complex trigger condition is feasible. Using this method, the increased background for low-quality tracks can be cut off with a stronger cut, while at the same time relaxing the cut for higher quality tracks and gaining additional trigger efficiency.



**Figure 5.19.:** Normalised distribution of 4, 5 and 6 tracklet cases, separately shown for  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV and background (other particles) in pp collision at  $\sqrt{s} = 8$  TeV. Statistical uncertainties are shown as vertical bars.

**Track qualities of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$**  It is important to understand how the track qualities are distributed for the target nuclei  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a preferred selection of certain track qualities may not only suppress the background, but also have an influence on the trigger condition efficiency. Figure 5.19 shows the relative abundances of the 3 different track qualities for the different particle species  ${}^3_2\text{He}^{2+}$ ,  $\overline{{}^3_2\text{He}^{2+}}$  and background particles. Statistical uncertainties are shown as vertical bars.

For  $\overline{{}^3_2\text{He}^{2+}}$ , the distribution is flat within the uncertainties and no significant dependence on the track quality can be observed. For  ${}^3_2\text{He}^{2+}$ , a significant dependence is observed, with a preference for lower quality tracks and lowering probability the more tracklets are required. The  $\vec{E} \times \vec{B}$  effect (also compare Section 5.2.2) explains the observed differences. For positively charged particles, the track bending due to the solenoid field is inclined into the same direction as the Lorentz angle of the drift charge path is inclined and thus the effective probability to not be measurable in all layers of the TRD is therefore larger.

While a strong cut on the 4-tracklet case is expected to reduce the background significantly, this has an undesirable impact on the efficiency for the positive  ${}^3_2\text{He}^{2+}$ , as a large fraction of  ${}^3_2\text{He}^{2+}$  particles is expected to leave only low quality tracks. Therefore, it is not possible to cut the 4-tracklet case too strong or even forbid 4-tracklet cases, a trade-off is required and thus a careful optimisation of the free threshold parameters promises best performance and will be presented in the next Section 5.2.2.

**Optimisation of the trigger condition** For the optimisation of the trigger cuts, a target trigger rejection factor of 1000 has been used. This number has been agreed upon with the ALICE trigger coordination team for the approval of the trigger during data taking for a fair use of ALICE readout resources while running the trigger at an envisaged 10 kHz inspection rate.

For finding the best cuts, an optimisation script has been written, which systematically loops over all free parameters and calculates the rejection factor based on LHC12a pp data. The results which match the target rejection factor of 1000 are further analysed regarding the resulting efficiencies. The final cuts are optimised for efficiency of  ${}^3_2\text{He}^{2+}$  in order to not cut into its signal (compare Fig. 5.19).

In addition to the PID values of all single tracklets, the GTU has even more track properties available within the online trigger decision logic. Among those, there are LCR flags, which have originally been developed for electron triggering. These are usually used to reject electron tracks which do not come from the primary vertex, but originate from later secondary decays. The flags are determined in the GTU based on the geometric track bending within the TRD. Though not originally developed for the nuclei trigger, the rejection of secondary electrons and positrons may be beneficial for triggering on nuclei too, as these background particles generate TR and may thus exhibit large PID. Therefore they may contribute to the trigger contamination.

The trigger cut optimisation has accordingly been adapted to support systematic variation of the following 6 free parameters:

**Table 5.2.:** Estimated required trigger cuts for a target trigger rejection factor of 1000 for a maximum  ${}^3\text{He}^{2+}$  trigger efficiency.


Parameter	Cut
PID cut on 4-tracklet tracks	126
PID cut on 5-tracklet tracks	118
PID cut on 6-tracklet tracks	112
LCR on/off for 4-tracklet tracks	on
LCR on/off for 5-tracklet tracks	off
LCR on/off for 6-tracklet tracks	off

- PID cut on 4-tracklet tracks
- PID cut on 5-tracklet tracks
- PID cut on 6-tracklet tracks
- LCR on/off for 4-tracklet tracks
- LCR on/off for 5-tracklet tracks
- LCR on/off for 6-tracklet tracks


The cuts have been found by the optimisation to allow for the largest trigger efficiency for  ${}^3\text{He}^{2+}$  at a constant trigger rejection factor of 1000 are shown in Table 5.2.

As the optimisation is carried out on a 6-dimensional space, only projections of the optimisation procedure can be plotted. For illustration, Fig. 5.20 and Fig. 5.21 show projections of how the results from Table 5.2 are obtained. In both figures all LCR cuts and the 4-tracklet PID cut are already fixed as in Table 5.2, leaving only 2 coordinates, the PID cut on 5-tracklet tracks (shown on x-axis) and the PID cut on 6-tracklet cuts (shown on y-axis).

The colours in Fig. 5.20 indicate the inverse trigger rejection factor and each contour line indicates an area of a constant inverse rejection factor. The contour line with the




**Figure 5.20.:** Contour lines of constant inverse trigger rejection factors as a function of the trigger cuts in pp collisions at  $\sqrt{s} = 8$  TeV. Exemplary projection from 6-dimensional optimisation. The marked contour line corresponds to the target inverse trigger rejection factor of 1 /1000.



**Figure 5.21.:** Average trigger efficiency for  ${}^3_2\text{He}^{2+}$  (arb. units) as a function of the trigger cuts in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV. Exemplary projection from 6-dimensional optimisation. The marked area corresponds to the target inverse trigger rejection factor 1/1000.





**Figure 5.22.:** Trigger condition efficiency on real data after optimisation and for a rejection factor of 1000 obtained in Pb–Pb collisions at  $\sqrt{s_{NN}} = 2.76$  TeV. Statistical uncertainties are shown as vertical bars.

target inverse rejection factor of 1/1000 is marked in the figure.

The colours in Fig. 5.21 indicate the resulting trigger efficiency for  ${}^3_2\text{He}^{2+}$ . The marking of the contour line from Fig. 5.20 along which the trigger runs with the target rejection factor of 1000 has been replicated in Fig. 5.21. The largest efficiency at the target rejection factor can be found along the contour line marking in the point with the darkest green colour. This point is at a 5-tracklet PID cut of 118 and a 6-tracklet PID cut of 112 (compare to Table 5.2).

**Trigger condition efficiency on real data** The resulting trigger condition efficiency on real data for the optimised cuts from Table 5.2 at a trigger rejection factor of 1000 for both  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  is shown in Fig. 5.22. Statistical uncertainties are shown as vertical bars.


## Improvements of the LUT

The electron probability LUTs presented in Fig. 5.14 is not usable with the nuclei trigger because it is not monotonous. The linear LUT presented in Fig. 5.13 works well for triggering on nuclei, but it comes with severe disadvantages for the existing electron triggers HQU and HSE. The resolution provided in the electron/pion overlap region is very small and can make it impossible to resolve the difference. However, data taking policies in the ALICE often require the operation of many different trigger types within the same run. This renders the option of switching between the 2 LUTs non-viable.

In order to overcome these issues and to commission the nuclei trigger along with existing triggers on electrons, a new type of LUT has been introduced and is shown in Fig. 5.23. The idea is to linearise the electron probability LUT in the electron/pion overlap region and in this way obtain a linear LUT which does not come with any disadvantages for electron triggering. The electron/pion overlap region profits from the same fine grained resolution as with the electron probability LUT. Also, the scaling is identical, so that no reoptimisation of the electron trigger thresholds is required.

Only the electron/pion overlap region is important for electron triggering, the other regions of the LUT can be freely optimised for triggering on nuclei. Therefore, the slope of LUT has been flattened towards the high energy deposition region, effectively resulting in a linear, monotonous LUT with 2 different slopes and a kink point between the region of interest for electron triggers and the region of interest for the nuclei trigger. The smaller slope in the high energy deposition region allows to distinguish high energy deposition particles with a very similar quality as with the previously used simple linear LUT. Without this adaption, the entire region of interest for nuclei triggering would become compressed into the overflow bin.

Note that after this optimisation, the interpretation of the LUT as an electron probability is no longer valid. However, due the monotone property of the new LUT, the original 16-bit energy loss value can always be at least approximately recovered on the GTU and



**Figure 5.23.:** Shipped 8-bit value as a function of the measured 16-bit raw ADC value (electron probability LUT from Fig. 5.14 in red color and newly introduced LUT in blue color). The newly introduced LUT allows for the coexistence between electron and nuclei triggers and it linearises the electron probability LUT in the electron/pion overlap region. After the kink point, the new LUT is optimised for triggering on nuclei.

during offline analysis and also the calculation of probabilities remains possible. The non-monotone electron probability LUT creates ambiguities and thus effectively discards information.

### 5.2.3. Hardware implementation of the nuclei trigger

A new FEE configuration has been generated which contains the proposed LUT presented in Section 5.2.2. The new LUT has been committed to the Subversion (SVN) repository with TRD FEE configurations to be available for building future configurations.

Furthermore, modifications for the GTU have been developed. The GTU is organised into TMUs, SMUs and the TGU. The tracks are reconstructed and track properties are evaluated within the TMUs, then the data is concentrated in the SMUs. The SMU is

also responsible for all trigger related processing which is possible at the level of each SM. The TGU is the trigger generation unit and has direct connectivity with the CTP. For details on the GTU, see Section 5.1.8.


A spare trigger connection to the CTP has been available with the label 1H14. This spare connection has not been used for triggering before and it has been agreed with the CTP team to be used for the new nuclei trigger. The CTP has introduced the name TRD Nuclei Trigger (HNU) for the new trigger on their side.

Since the nuclei trigger is a single track based trigger and no correlations of signals across different SMs are necessary, a distributed trigger processing is possible and no processing is required in the TGU. The spare connection to the CTP is connected to an OR-logic in the TGU. This logic picks up individual signals from the 18 SMUs and sends the trigger to the CTP when one of the SMUs delivers a trigger contribution. This logic can be used with the nuclei trigger, so that for the nuclei trigger, no modification of the TGU design is required.

The implementation of the nuclei trigger has been carried out by adapting the track based trigger decision logic in the 18 SMUs according to the trigger condition optimisation in Section 5.2.2. The thresholds found with the linear LUT have been scaled in order to work with the new LUT. The changes to the SMU design have been committed to GTU git repository for building future GTU firmware versions.

#### 5.2.4. TRD nuclei trigger in pPb collisions

The nuclei trigger has been optimised for pp collisions, however the data taking schedule of ALICE has allowed to operate the trigger in addition during the pPb collision at  $\sqrt{s_{NN}} = 8$  TeV periods LHC16r and LHC16s. The data from these periods are used in this analysis. Instead of MB, these periods have been exclusively dedicated to the operation of rare triggers, i.e. a certain fraction of the events is read out after firing the nuclei trigger. The other readout resources are dedicated to other rare triggers.



**Figure 5.24.:** TPC signal (arbitrary units) as a function of the rigidity  $p/Z$  in pPb collisions at  $\sqrt{s_{NN}} = 8$  TeV. Comparison of the HNU triggered data set (red points) with an equally sized reference data set (black points). The blue splines indicate the expected Bethe-Bloch energy loss within the TPC for, from bottom to top, protons, deuterons, tritons,  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  as well as their antiparticles.

### Comparison of the nuclei triggered data sample with the trigger mix

In order to check qualitatively, whether the nuclei trigger has worked as expected, the nuclei triggered data sample has been extracted from the collected data in LHC16r and LHC16s (in the following called HNU data set). In addition, the data set triggered by other triggers (in the following called trigger mix data set) has been extracted from the same periods. The trigger mix data set is larger than the HNU data set, therefore the trigger mix data set has been randomly downscaled such that the resulting number of events matches the one in the HNU data set and allow for comparison.

Figure 5.24 shows the TPC signal as a function of the rigidity  $p/Z$  for both data sets. Particles from the HNU data set are marked in red, particles from the trigger mix data set are marked in black. The blue lines show the TPC splines for the identification of

protons, deuterons, tritons,  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  as well as their respective antiparticles. The result does not contain any particle with a TPC energy loss lower than the proton spline. This is due to the use of a grid based pre-selection cut to reduce data sizes.

Since the 2 data sets have the same number of events, the ratio between the number of HNU points and the number trigger mix points in a particular region corresponds to the trigger enhancement. Especially for  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  as well as for their antiparticles, it can be observed that the majority of particles in the vicinity of the corresponding splines belong to the HNU data set. This indicates an outstanding enhancement performance of the nuclei trigger when compared to other triggers. As expected, an enhancement effect with a lower enhancement factor can also be observed for tritons, deuterons and their antiparticles.


Figure 5.24 clearly shows more nuclei than antinuclei. This unbalance is due to the fact that nuclei can not only be produced in the collision, but they can also be knocked-out by other particles from the detector material or from the LHC beam pipe. Antinuclei can only be produced during the collision.

### Selection of $\overline{{}^4_2\text{He}^{2+}}$ candidates


Figure 5.24 shows that the HNU data set contains an entry in the vicinity of the  $\overline{{}^4_2\text{He}^{2+}}$  spline at approximately  $p_{\text{T}}/Z \approx -1.6 \text{ GeV}/c$ . There is another entry at approximately  $p_{\text{T}}/Z \approx -2 \text{ GeV}/c$  which is covered by the  $\overline{{}^4_2\text{He}^{2+}}$  spline and therefore not visible in Fig. 5.24.

The production of these antiparticles during the collision is certain and a material knock-out from the beam pipe or detector can be excluded. The observation of 2  $\overline{{}^4_2\text{He}^{2+}}$  candidates is compelling as  $\overline{{}^4_2\text{He}^{2+}}$  has never been observed in neither the small pPb collision system nor in the small pp collision system before.

However, both  $\overline{{}^4_2\text{He}^{2+}}$  are in the region of  $p_{\text{T}}/Z \approx -2 \text{ GeV}/c$ , a point where the TPC splines of  ${}^3_2\text{He}^{2+}$  and  ${}^4_2\text{He}^{2+}$  start to approach each other, risking a contamination of the




**Figure 5.25.:** Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of the rigidity in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. The green lines indicate the PID relevant boundaries.



**Figure 5.26.:** Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  as a function of the rigidity in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. The green lines indicate the PID relevant boundaries.

$\overline{{}^4_2\text{He}^{2+}}$  spline with  $\overline{{}^3_2\text{He}^{2+}}$ . In order to increase the level of confidence that the observed  $\overline{{}^4_2\text{He}^{2+}}$  are not contaminating  $\overline{{}^3_2\text{He}^{2+}}$ , further investigations are important.

Figure 5.25 shows the relative deviation of the TPC signal from the expected  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  splines for the combined LHC16r and LHC16s data set. The green lines indicate the PID relevant boundaries.  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  candidates are selected from in between the




**Figure 5.27.:** TOF  $\beta$  as a function of the transverse rigidity for the TPC pre-selected  $\frac{3}{2}\text{He}^{2+}$ ,  $\frac{3}{2}\text{He}^{2+}$ ,  $\frac{4}{2}\text{He}^{2+}(\alpha)$  and  $\frac{4}{2}\text{He}^{2+}(\bar{\alpha})$  candidates in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Candidates pre-selected as  $\frac{3}{2}\text{He}^{2+}$  and  $\frac{3}{2}\text{He}^{2+}$  are marked as black points, while candidates pre-selected as  $\frac{4}{2}\text{He}^{2+}$  and  $\frac{4}{2}\text{He}^{2+}$  are marked as red points.

boundaries. For the positive values, the lower boundary has been increased in order to reduce the background from triton knock-out.

Figure 5.26 shows the relative deviation of the TPC signal from the expected  $\frac{4}{2}\text{He}^{2+}$  and  $\frac{4}{2}\text{He}^{2+}$  splines for the combined LHC16r and LHC16s data set. The green lines indicate the PID relevant boundaries.  $\frac{4}{2}\text{He}^{2+}$  and  $\frac{4}{2}\text{He}^{2+}$  candidates are selected from in between the boundaries. It can be observed that at least for  $|p_{\text{T}}/Z| > 2$  GeV/c that contamination from  $\frac{3}{2}\text{He}^{2+}$  and  $\frac{3}{2}\text{He}^{2+}$  is present in the band. In order to improve the purity of the PID, the TOF detector has been used.

For this study, the TOF has been used to calculate the particle velocity  $\beta$  (compare Chapter 4) for the TPC pre-selected  $\frac{3}{2}\text{He}^{2+}$ ,  $\frac{3}{2}\text{He}^{2+}$ ,  $\frac{4}{2}\text{He}^{2+}$  and  $\frac{4}{2}\text{He}^{2+}$  candidates. Figure 5.27 shows the TOF  $\beta$  as a function of the transverse rigidity for the TPC pre-selected  $\frac{3}{2}\text{He}^{2+}$ ,






**Figure 5.28.:** *Acceptance efficiency* as a function of the transverse momentum  $p_T$  for  ${}^3\text{He}^{2+}$  (red markers) and  $\overline{{}^3\text{He}}^{2+}$  (blue markers) candidates in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Statistical uncertainties are shown as vertical bars.

$\overline{{}^3\text{He}}^{2+}$ ,  ${}^4\text{He}^{2+}$  and  $\overline{{}^4\text{He}}^{2+}$  candidates. The candidates pre-selected as  ${}^3\text{He}^{2+}$  and  $\overline{{}^3\text{He}}^{2+}$  are marked as black points, while candidates pre-identified as  ${}^4\text{He}^{2+}$  and  $\overline{{}^4\text{He}}^{2+}$  are marked as red points.

It can be observed that for the  ${}^4\text{He}^{2+}$  candidates at approximately  $p_T/Z \approx -1.6$  GeV/c and approximately  $p_T/Z \approx -2$  GeV/c, the TOF detector shows a significant separation from the  $\overline{{}^3\text{He}}^{2+}$  band. Other particles pre-selected by the TPC as  $\overline{{}^4\text{He}}^{2+}$  are clearly identified as contamination from  $\overline{{}^3\text{He}}^{2+}$ . A similar behaviour can be observed with more statistics for positive  $p_T/Z$ .

### Measurement of the *acceptance efficiency*

Figure 5.28 shows the ratio of the number of tracks measured by the TRD with respect to the number of global tracks, defined as *acceptance efficiency*, for  ${}^3\text{He}^{2+}$  and  $\overline{{}^3\text{He}}^{2+}$  candidates




**Figure 5.29.:** Efficiency of the trigger condition as a function of the transverse momentum  $p_T$  in pPb collisions at  $\sqrt{s_{NN}} = 8$  TeV. Left hand panel:  ${}^3_2\text{He}^{2+}$  (red markers) and  ${}^3_2\overline{\text{He}}^{2+}$  (blue markers). Right hand panel:  ${}^4_2\text{He}^{2+}(\alpha)$ .  ${}^4_2\overline{\text{He}}^{2+}(\overline{\alpha})$  is not shown due to low statistics. Statistical uncertainties are shown as vertical bars.

on the LHC16r and LHC16s combined data sets.

It can be observed that the pattern in Fig. 5.28 is similar to the ones observed in the Fig. 5.10 and Fig. 5.15 for MC and the preliminary Pb–Pb data studies, respectively.

### Measurement of the trigger condition efficiency

Figure 5.29 shows the efficiency of the trigger condition on pPb data for  ${}^3_2\text{He}^{2+}$ ,  ${}^3_2\overline{\text{He}}^{2+}$ , and  ${}^4_2\text{He}^{2+}$ . The present statistics are not sufficient to calculate the  ${}^4_2\overline{\text{He}}^{2+}$  trigger condition efficiency. The results are obtained by dividing the number of tracks fulfilling the trigger condition by the number of tracks measured in the TRD.



**Figure 5.30.:** Left hand panel: Total efficiency of the HNU trigger on pPb data for  ${}^3\text{He}^{2+}$  (red markers) and  $\overline{{}^3\text{He}^{2+}}$  (blue markers). Right hand panel: Total efficiency of the HNU trigger on pPb data for  ${}^4\text{He}^{2+}$ .  $\overline{{}^4\text{He}^{2+}}$  is not shown due to low statistics.

The pattern of the trigger condition efficiency observed on pPb data in Fig. 5.29 seems to differ from the one observed on Pb–Pb data in Fig. 5.22. A possible interpretation of the discrepancies is due to the asymmetry of the collision system. It has been observed that the fraction of 4-tracklet TRD tracks with respect to all TRD tracks increases for high PID values on pPb data. The HNU trigger is sensitive to this effect as the trigger PID thresholds depend on the number of tracklets. However, the large statistical error bars do not allow for a final conclusion.

### Measurement of the total trigger efficiency

Figure 5.30 shows the total efficiency of the HNU trigger in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV for  ${}^3\text{He}^{2+}$ ,  $\overline{{}^3\text{He}^{2+}}$ , and  ${}^4\text{He}^{2+}$ . The present statistics are not sufficient to calculate the  $\overline{{}^4\text{He}^{2+}}$  total HNU trigger efficiency. The results are obtained by dividing the number

of tracks fulfilling the trigger condition by the number of measured global tracks.

The observed total efficiencies are in agreement with the estimates presented prior to the commissioning of the HNU trigger.

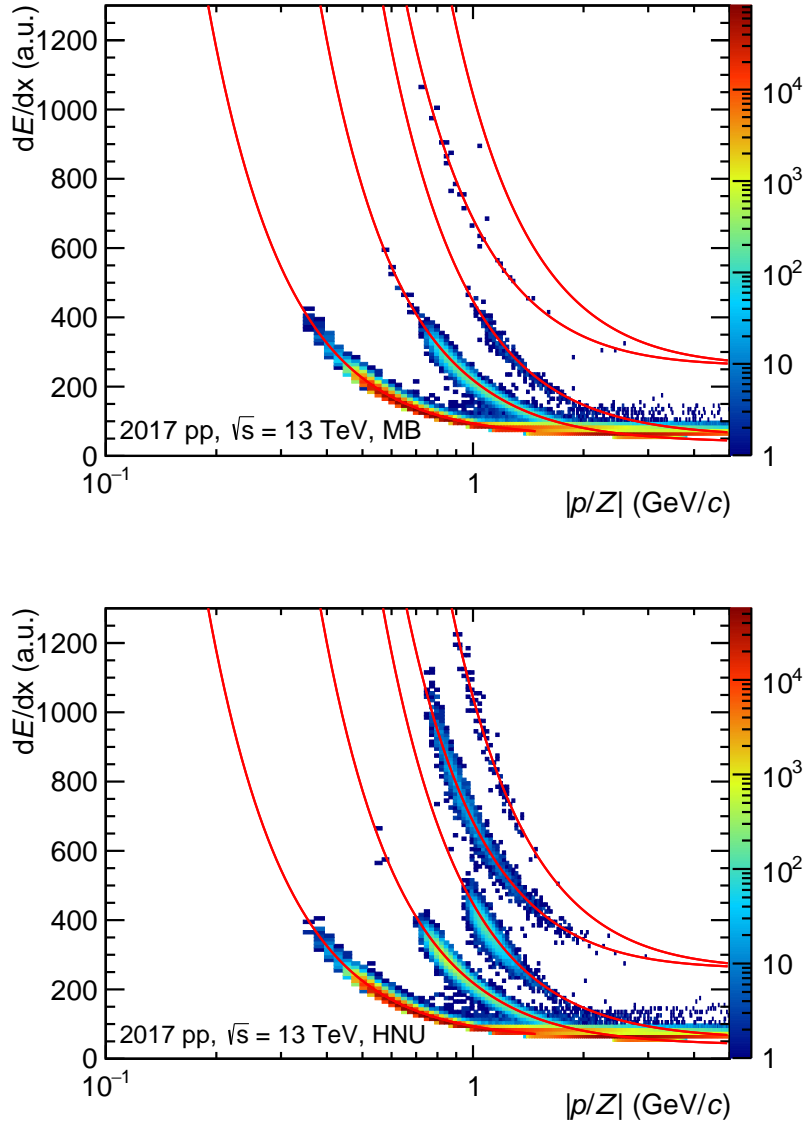
### 5.2.5. TRD nuclei trigger in pp collisions

The nuclei trigger has been operated in pp collisions during the data taking periods in 2017 and 2018. This section shows the performance of the nuclei trigger in pp collision during the triggered periods in 2017. In this section, the nuclei triggered data is compared to MB data. After quality selection cuts, the analysed MB data sample comprises approximately  $4.6 \cdot 10^7$  MB events and the analysed nuclei triggered data sample comprises approximately  $1.9 \cdot 10^7$  events.


#### Selection of nuclei candidates with the TPC

Figure 5.31 shows the energy loss in the TPC (arbitrary units) as a function of the rigidity  $|p/Z|$ . The red splines correspond to the expected Bethe-Bloch energy loss for each particle species (from top to bottom:  ${}^4_2\text{He}^{2+}$ ,  ${}^3_2\text{He}^{2+}$ , triton, deuteron, proton) for both data sets in comparison. A pre-selection cut in the level of the proton spline has been applied for early data reduction. Particles which exhibit a smaller energy loss in the TPC than the expected energy loss for protons are therefore not shown in Fig. 5.31.

A significant enhancement of  ${}^4_2\text{He}^{2+}$ ,  ${}^3_2\text{He}^{2+}$ , tritons and deuterons and their antiparticles can be observed in the triggered data sample with respect to the MB sample. As expected from previous studies, the enhancement increases with the mass of the particles and is most prominent for  ${}^4_2\text{He}^{2+}$  and  $\bar{{}^4_2\text{He}^{2+}}$ . Note that the MB data set in Fig. 5.31 contains approximately 2.4 times more events than the nuclei triggered sample, i.e. that the enhancement is 2.4 times larger than it appears. The enhancement will be measured using a more quantitative approach at the end of this section.



**Figure 5.31.:** TPC signal  $dE/dx$  in arbitrary units as a function of the rigidity  $|p/Z|$  in pp collisions at  $\sqrt{s} = 13$  TeV. The red splines correspond to the expected Bethe-Bloch energy loss for each particle species (from top to bottom:  ${}^4_2\text{He}^{2+}$ ,  ${}^3_2\text{He}^{2+}$ , triton, deuteron, proton). Particles and their respective antiparticles are combined. Top panel: MB data set with approximately  $4.6 \cdot 10^7$  MB events. Bottom panel: nuclei triggered data set with approximately  $1.9 \cdot 10^7$  events.



**Figure 5.32.:** Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of the transverse rigidity  $|p_T/Z|$  in pp collisions at  $\sqrt{s} = 13$  TeV. Top panel: MB data set. Bottom panel: nuclei triggered data set.

For the quantitative measurement, the particles have been identified using the TPC. For each expected Bethe-Bloch energy loss spline ( ${}^4_2\text{He}^{2+}$ ,  ${}^3_2\text{He}^{2+}$ , triton, deuteron), the relative deviation between each particle in the data sets and the expected splines has been calculated. Figure 5.32 shows an example for the case of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  for both data sets. The identified  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  gather around the horizontal at a deviation of 0. In order to obtain the PID relevant boundaries, the histogram has been projected on the y-axis and a Gauss function has been fitted around 0. For all particle species it has been observed that the  $3\sigma$  range of each fitted Gaussian is approximately  $\pm 0.2$  in accordance with the boundaries observed in the analysis of the pPb collisions at  $\sqrt{s_{\text{NN}}} = 8 \text{ TeV}$  in the Figures 5.25 and 5.26.

The  $p_{\text{T}}/Z$  range in the projections and for PID purposes has been limited to the values in Table 5.3 to limit contamination from overlapping bands.

**Table 5.3.:** Transverse rigidity ranges  $|p_{\text{T}}/Z|$  for the identification of light nuclei.

Particle	$ p_{\text{T}}/Z $ range (GeV/c)
Deuteron and Antideuteron	0.2 – 1
Triton and Antitriton	0.3 – 1.5
${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$	0.4 – 2.5
${}^4_2\text{He}^{2+}$ and $\overline{{}^4_2\text{He}^{2+}}$	0.5 – 2.5

### Further PID with TOF

For each particles species selected using the TPC, the TOF  $\beta$  and the rigidity have been used in order to calculate the parameter  $m^2/Z^2$

$$\frac{m^2}{Z^2} = \frac{1 - \beta^2}{\beta^2} \left( \frac{p}{Z} \right)^2, \quad (5.1)$$

where  $m$  corresponds to the particle mass.

Figure 5.33 shows an example of the parameter  $m^2/Z^2$  as a function of the transverse rigidity for the case of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$ . Each particles species has been selected around its expected of value  $m^2/Z^2$ . For  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$ , it is  $m^2/Z^2 = 1.97 \text{ GeV}^2/c^2$ , for  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  it is  $m^2/Z^2 = 3.47 \text{ GeV}^2/c^2$ . In order to identify the particles, the results in Fig. 5.33 have been projected onto the y-axis and a Gaussian function has been fitted for each particle species. The number of particles for each particle species and for each data set is obtained by calculating the integral below each Gaussian fit function.

### Enhancement factor

In order to calculate the quantitative nuclei trigger enhancement factor  $F_k$  for each particle species with respect to minimum bias data, the number of particle in each data set is normalised by the corresponding number of events in the data set and put into comparison. The enhancement factor for the particle species  $k$  is defined as


$$F_k = \frac{n_k^{\text{HNU}}/N^{\text{HNU}}}{n_k^{\text{MB}}/N^{\text{MB}}}, \quad (5.2)$$

where  $n_k^{\text{HNU}}$  and  $n_k^{\text{MB}}$  are the number of particles of the species  $k$  in the nuclei triggered data set and in the MB data set, respectively. The  $N^{\text{HNU}}$  and  $N^{\text{MB}}$  terms correspond to the number of events in the nuclei triggered data and in the MB data set, respectively.


Figure 5.34 shows the enhancement factor of the nuclei trigger in pp collisions at  $\sqrt{s} = 13 \text{ TeV}$  for each particle species and separately for particles and antiparticles with respect to MB.

Figure 5.34 shows that a significant enhancement factor is observed for all presented particles. For  ${}^3_2\text{He}^{2+}$  the enhancement is approximately 175 and for  $\overline{{}^3_2\text{He}^{2+}}$  the enhancement





**Figure 5.33.:** Parameter  $m^2/Z^2$  of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of the transverse rigidity  $p_T/Z$  in pp collisions at  $\sqrt{s} = 13$  TeV. Top panel: MB data set. Bottom panel: nuclei triggered data set.



**Figure 5.34.:** Enhancement factor of the nuclei trigger, shown separately for each particle species and antiparticle species, with respect to MB in pp collisions at  $\sqrt{s} = 13$  TeV. Statistical uncertainties are shown as vertical bars.

factor is approximately 328. For  $^4_2\text{He}^{2+}$  a clear enhancement with a factor of approximately 201 can be observed, however the lack of  $^4_2\text{He}^{2+}$  statistics in MB contributes a large error bar (also compare Fig. 5.31). The pattern of an increasing enhancement factor with increasing particle mass is reproduced in the figure.

## Chapter 6

# TRD Upgrade and Developments for Run 3

The achievements of this thesis for the LHC Run 3 period are presented in this chapter. For ALICE, a new readout concept is introduced in Run 3, which enables the permanent recording of data from all events, also referred to as continuous readout. This greatly improves the previous concept of employing online event selection mechanisms, as this method provides bias-free data for physics research and opens new opportunities in investigating effects without specific trigger signatures. Furthermore, ALICE is capable of recording remarkably faster collision rates of up to 50 kHz Pb–Pb collisions and up to 1 MHz pp collisions in Run 3.

These tremendous improvements are made possible by a large-scale technical upgrade of the majority of detector readout hardware in order to employ the latest available technologies in ALICE.

This thesis contributes significantly to the ALICE Run 3 upgrade by providing the Transition Radiation Detector (TRD) specific parts of the Common Readout Unit (CRU) firmware, which are developed entirely and exclusively within the scope of this work. The CRUs are Field Programmable Gate Array (FPGA) based readout equipment, which are used for the first time for the Run 3 upgrade. In case of the TRD, these devices are

placed at the core of the TRD readout infrastructure, such that they handle the entire data generated by the TRD and, in addition, provide interfaces for the other ALICE subsystems with the TRD. The CRUs control the TRD readout process during data taking.

The new readout system has been thoroughly tested in the laboratory and within ALICE. The developed CRU firmware is found to comply in all aspects with its prior specifications. The entire TRD readout chain is proven to perform stably and its performance exceeds previous estimations. After the beginning of Run 3 data taking in mid 2022, particle tracks can be successfully reconstructed based on the collected data with the CRUs.

The overview of the TRD upgrade for Run 3 is given in section 6.1, especially highlighting the differences to Run 2 (compare section 5.1). Sections 6.2 to 6.7 consider the aspects and prerequisites for the operation of the TRD in the new readout environment more closely. The validation of the TRD CRU custom firmware and performance evaluations are presented in section 6.8.

The Theory of Operation of the TRD CRU custom firmware along with its development is presented in details in the appendix A.

## 6.1. Transition Radiation Detector Upgrade Overview

The TRD readout tree in Run 2 has consisted of the TRD FEE, the GTU and the DAQ. The FEE and the GTU have received triggers from the CTP at LM and L1 trigger levels. Upon the reception of these triggers, the FEE has produced preprocessed tracklet data at the LM trigger level and raw data at the L1 trigger level, which it has transferred through optical links to the GTU. The GTU has relayed the data via additional optical links to the DAQ. In addition, the GTU has been able to provide TRD L1 trigger contributions for different types of triggers to the CTP based on the received data from the FEE after an LM trigger.

This entire readout chain has originally been designed for a Pb–Pb collision interaction rate of 8 kHz and for a significant event rejection at the L1 trigger stage. The chain has thus been heavily optimised in the past for a high event inspection rate at LM trigger level with a fast calculation of the L1 trigger contribution (LM tracklet data readout time  $< 10 \mu\text{s}$ , L1 decision time  $< 6.5 \mu\text{s}$ ), while transferring large and high-resolution raw data events for events accepted beyond the L1 trigger level. For further details, refer to Section 5.1.

In Run 3, the L1 trigger functionality is no longer required and the detector is required to provide readout rates as high as feasible while writing all events to permanent storage. No data shall be discarded in the readout chain and the fraction of recorded events in an 1 MHz pp collision interaction rate or 50 kHz Pb–Pb collision interaction rate running scenario shall be maximised. Since the FEE is based on Application-Specific Integrated Circuits (ASICs), which are specially designed and optimised for the Run 1 and Run 2 data taking scenarios, the reprogrammability is limited. Their replacement is not practical within the scope of the Run 3 upgrade and the adaption to the new requirements has to be made on a best-effort base.

The applied solution is presented in the following sections. Simulations confirm that it allows to collect around 70 % of events in a 50 kHz interaction rate Pb–Pb data taking scenario.

This section is based on information from the Ref. [3, 49–51], which may be consulted for further details.

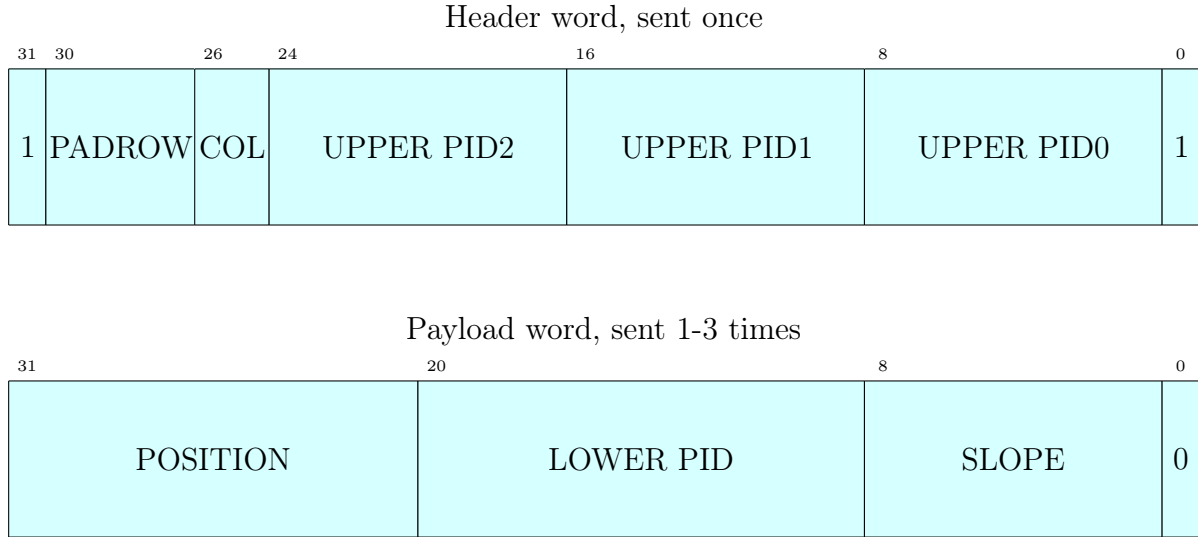
### 6.1.1. Optimisation of the existing Front End Electronics

In order to achieve a high event readout rate in Run 3, the tracklet readout mode which has been used to find fast L1 trigger contributions in Run 2, is leveraged for physics in Run 3. The data format, which had been optimised previously for the purpose of triggering in Run 2, has been carefully reconsidered and is enhanced with additional information valuable to

physics analysis while still fitting into the available data volume. The available data volume per LM trigger and per Multi Chip Module (MCM) is 4 tracklet words of 32 bits each.

The optimisations applied on the FEE are listed below [3, 49, 50]:

- In Run 2, all 4 words of 32 bits each have been used for actual tracklets, which has always limited the maximum number of tracklets per MCM and event to a maximum number of 4 tracklets. For the format used in Run 2, refer to Section 6.6.3. Data analysis however shows that even in the most central Pb–Pb collision events in the past Run 1 and Run 2, a track density of 4 tracklets per MCM has been rarely reached. Therefore only three tracklets per MCM are allowed in the Run 3 data format. The estimated fraction of tracklets lost by this measure in central Pb–Pb collision events is below 1%, while the freed-up space of 32 bit can be used to ship additional information.
- In Run 2, all information concerning a particular tracklet have been stored in a 32 bit word devoting only 8 bits to the PID purpose. In Run 3, the freed-up 32 bit word is used as a header to store position information about the MCM and 8 bit of PID information per tracklet. The header is followed by 1 to 3 32-bit tracklet words that store the position within the MCM, slope and additional 12 bits of PID information of the tracklet. In total, 20 bit PID resolution is available for every found tracklet, allowing the PID to profit from much higher resolution and/or dynamic range, respectively.
- In Run 2, most TRD triggers have involved an online cut on a minimum transverse momentum ( $p_T$ ). The online  $p_T$  has been calculated precisely on the Field Programmable Gate Arrays (FPGAs) of the GTU, which has applied the final trigger  $p_T$  cut. However, in order to shorten the readout time after an LM trigger, an additional preselection cut has been applied on the tracklets by the FEE. This cut has been based on the tracklet inclination. For the Run 3 upgrade, this cut is lifted in order to avoid a bias in the collected physics data.
- In Run 3, the TRD uses a physics trigger sent by the CTP at LM latency. Technically, it works identically to the LM trigger in Run 2. In addition, the TRD supports



**Figure 6.1.:** TRD tracklet data format. Each MCM that has reconstructed at least one tracklet will send a header with shared coordinate information and 8 bits of PID information per tracklet. For each reconstructed tracklet, one additional payload word with additional position and PID information as well as the reconstructed tracklet angle will be stored. Adapted from Ref. [3].

a new trigger type, called calibration trigger. The calibration trigger, also sent at LM latency, enables the shipping of physics data and, additionally, the full raw data which has been recorded by default in Run 2. This allows to trigger a full readout for a small fraction of events in Run 3, facilitating detector calibration. Apart from that, the calibration trigger is interpreted by the FEE as a command to reload its configuration parameters from hamming protected memory areas. This is a precaution measure and mitigates the impact of Single Event Upsets (SEUs) on data taking, sporadically observed on some isolated half chambers as Link Monitor Errors (LMEs) throughout Run 1 and Run 2.

The data format to be sent by the FEE in Run 3 is given in Fig. 6.1.

### 6.1.2. Common Readout Unit

The CRU is a new hardware equipment and the Run 3 upgrade is its very first application. It is designed as a PCIe40 card based on the Intel Altera Arria 10 FPGA. The board

also integrates power supply, the Peripheral Component Interconnect Express (PCIe) interface, 2 cages for Small Form Factor Pluggable (SFP) modules and Avago minipod optical components [51].

It has been developed in a joint venture between ALICE and Large Hadron Collider beauty (LHCb). During the development, different PCB versions of the board have been manufactured and used [51]:

- **Altera Arria 10 development kit**

Commercial development kit marketed by the FPGA manufacturer. It provides a PCIe interface and a maximum of 4 optical Receiver (RX)/Transmitter (TX) links. It has been used in the early stages of CRU firmware testing and development as long as no custom board version has been available yet.

- **CRUv1**

First version of fully customised CRU board for CERN.

- **CRUv2 pre-series version**


This version of the hardware comes with many improvements and it features a redesigned power supply unit. This version is almost identical to the final version, but PCB traces related to the TTC-PON upstream feature are connected to different FPGA pins. Therefore it is not firmware compatible with the final system version. The pre-series version is for example used in the TRD laboratory in Münster (compare Section 6.2).

- **CRUv2 final system version**

Final version for use at LHC Point 2 - ALICE (P2). At P2, only this type of CRU is used.

Figure 6.2 shows a photograph of the second version of CRU.





**Figure 6.2.:** Photograph of the second version of the CRU PCIe card. Major improvements with respect to first version are highlighted. Figure from Ref. [52].

The customised PCB versions can be further subdivided into different types depending on the components which are mounted. Especially the optical components for data transmission and reception (Avago Minipods) are easily replaceable as they are mounted through specific solder-less sockets. The custom CRU PCB provides 8 such sockets and every minipod can handle 12 RX or 12 TX links, depending on its type. There are 2 equipment levels of the CRU which are most widely used:

- The standard CRU version equipped with 24 optical RX links and 24 optical TX links. This version is used by all CRU based detectors in ALICE except by the TRD. Also, the CRU used in the laboratory in Münster is of this type.
- The special TRD CRU version equipped with 36 optical RX links and no TX links. This version is only used by the TRD at P2. This design is better adapted for the TRD with its large number of unidirectional optical links.

The minipods, which are mounted on the PCB, are optically connected through tiny optical fibres to Multi-fibre Push-On (MPO) receptacles in the card backplane

(PRIZM® (PRIZM) assembly). Each MPO connector houses 12 RX or TX links, respectively. The card version specialised for the TRD is equipped with 3 female MPO receptacles.

All versions of the CRU are equipped with 2 SFP cages in the card backplane to hold SFP modules. In the final system, one of these SFP cages is used to house an ONU SFP in order to interface with the trigger system via TTC-PON.

The CRU is equipped with 6 test Light Emitting Diodes (LEDs) in the card backplane for general debugging, therein 1 hard-wired power LED and 5 firmware controllable Red-Green-Blue (RGB) LEDs. In the final system, they are not used as they are hidden in the rack mounting positions of the CRUs and invisible to the operators.

Also, the CRUs provide a connector with spare electrical General Purpose Input Outputs (GPIOs) in the card backplane for eventual extensions. The TRD does not use them.

In order to satisfy the high power requirements which might arise under high processing load and to avoid overcurrents through the PCIe socket, the CRU provides an internal additional power connector which has to be attached directly to the +12V rail and ground of the host machine using an ATX PCIe connector.

The firmware on the FPGA is volatile and in principle needs reprogramming after every power cycle. To circumvent this limitation, there is a non-volatile flash memory chip on the board which provides a valid firmware to the FPGA at Power On Reset (POR). It has to be programmed independently of the FPGA with a suitable firmware.

The CRU features an on-board Universal Serial Bus (USB)-to-Joint Test Action Group (JTAG) adaptor to program both FPGA and serial flash memory through a mini

USB port in the backplane.

These information are based on the Ref. [\[51\]](#).

### 6.1.3. O2 Data Acquisition System

The Online Offline Processing System (O2) (pronounced: O-square) is a new central system of ALICE for data acquisition. It replaces the previous DAQ. In this section, 2 important components of the O2 system are introduced.

#### First Level Processor

The First Level Processor (FLP) constitutes the first O2 machine through which the data passes on its way from the latency critical Front End hardware systems at the detector to its final storage. There are no latency critical systems after the FLP. The FLP is based on a Dell PowerEdge R740 server machine and gets mounted in racks. This equipment is used for the first time in Run 3. It uses CentOS Linux as operation system [\[51\]](#).

The FLP serves as a host computer for the CRUs and each FLP houses up to 4 CRUs. All FLPs are installed in the Counting Room 1 (CR1). The CRUs are connected to the PCIe slots and power supply of the FLP. All other connections to the CRU are accessible through its backplane which integrates into the backplane of the FLP [\[51\]](#).

The FLP receives the data from the installed CRUs using Direct Memory Access (DMA) technology. It allows writing the data directly from the CRU into reserved memory pages in the Random Access Memory (RAM). This prevents the data from passing via the Central Processing Unit (CPU) and keeps the load on the CPU low even during data transfer [\[51\]](#).

The FLP supports PCIe bifurcation. This allows to split a PCIe interface logically into multiple interfaces with less lanes. The split PCIe interfaces appear as separate and independent devices to the operational system on the FLP. The CRU requires to split each PCIe slot into halves [51].

The FLP can access the CRUs via the PCIe interface through a dedicated slow control bus system (Avalon bus) which extends through all modules of the CRU firmware. This allows to run software on the FLP which enables configuration and monitoring of the CRUs by different means. During development, typically command line tools are used, but during productive detector operation, the slow control can be accessed using the DCS tools too [51].

It is furthermore possible to perform firmware updates of the CRUs through the PCIe interface. This eliminates the need to connect all CRUs to a USB host computer to reprogram them via JTAG. During productive use, the JTAG adaptor integrated into the CRU serves only for emergency purposes, in case that the firmware update via PCIe fails and the firmware has to be recovered during a physical access. Physical access to the FLP installation location in CR1 is possible at all times, even during beam operation [51].

Apart from the DMA software and CRU control software, there is additional software running on the FLP which take care of formatting the data for the downstream systems in the O2 chain. The Sub Time Frame Builder (STFB) takes care of repacking the data from the CRUs into so called sub time frames, where each sub time frame contains the data from a particular FLP and from a particular LHC orbit (time frame duration). The sub time frames are then passed on to the sub time frame sender, which takes care of shipping the data out of the FLP via a network interface. The FLPs are equipped with InfiniBand network cards which allow a throughput of up to 100 Gb/s per FLP [51].

It is possible to add additional software in between software components described above. This can be used in order to do detector specific data processing with the computing resources of the FLP. For the TRD, a data compressor will be used, which is in

development by the detector team.

These information are based on the Ref. [51].

## Event Processing Node

The data which has been passed on by the FLPs via the InfiniBand network, is received by the Event Processing Nodes (EPNs). The EPNs are new Run 3 server machines which assemble the sub time frames from different FLPs for a particular LHC orbit to generate so called Time Frames (TFs). A time frame in the final system contains the data from all detectors which are included in the run during data taking, generated within a particular LHC orbit. Finally, the data can also be reconstructed by the EPNs which are capable of performing this tasks online [51].

Similarly to the FLPs, also the EPNs can also be used for detector specific data processing tasks. This is especially interesting in cases where the FLPs do not provide enough resources for all detector specific data processing tasks and would otherwise risk to generate critical backpressure for the CRUs [51].

Apart from the online data processing and reconstruction features of the EPNs, the EPNs which are not allocated for data taking at a particular time, can be made available for offline data analysis tasks. In this way, available computing resources are used more efficiently. The nomenclature of the O2 as both the *Online* and *Offline* system is based on this feature [51].

These information are based on Ref. [51].

#### 6.1.4. New Run 3 Readout Scheme for the TRD

In Run 3, the GTU is replaced by CRUs. The CRUs receive the data directly from the FEE via the 1044 optical FEE links. Every CRU effectively provides 30 link inputs, implying that in total 36 CRUs are required. They are mounted in 12 TRD FLPs in total, each of them containing 3 CRUs.

A special FPGA firmware for the CRU has been developed within the scope of this thesis. For details, refer to Section 6.6. The firmware controls the readout process of the detector, receives, buffers and formats the data for the O2 system. All CRUs are connected to the Local Trigger Unit (LTU) on the CTP side to receive trigger information and to signal a detector busy status to the CTP. Each CRU determines an individual busy status contribution depending on the status of the readout of the connected FEE links. The CTP links the busy status contributions from all CRUs together to determine a global busy status of the TRD.

The O2 system performs additional processing steps in software, including a deeper level reformatting and data compression, before the data is written to a permanent storage.

The following points describe sequentially the process of acquiring an event explaining the role of the CRU and the interactions with other readout systems:

1. CTP sends a trigger at LM latency (physics trigger or calibration trigger). This trigger is sent via the LTU to the FEE and to the CRUs in parallel. The trigger to the FEE is shipped via the legacy TTC network which has already been used in Run 2. The trigger to the CRU is sent via 9 newly set up trigger distribution networks. They use the Trigger and Timing Control via TTC-PON technology, which is a new trigger distribution technology for Run 3. The large number of 9 networks is necessary to achieve minimum latency. For details, also refer to Section 6.5. The CRUs store all information from the received trigger message (e.g. orbit id, bunch crossing id, etc.) in internal buffers.

2. Upon the arrival of the trigger, the FEE begins recording the data while the drift charges are drifting to the readout pads. Each CRU receives the trigger at approximately the same time and internally opens a time window within which it is waiting for the input links to send all the acquired data. The timeout is programmable. In addition, each CRU generates its busy status contribution and sends it to the CTP. The TTC-PON networks support upstream communication through time multiplexing which is used to send the busy status signal. This indicates to the CTP to not send any other trigger as long any CRU provides an active busy contribution in order to avoid confusion of the FEE state machine.
3. When the FEE has acquired and processed the data, it starts shipping it via the TRD custom optical links. At the end of the transmission, the FEE appends specific end markers. The CRUs record the data received on all input links until it recognises the end markers. If for any reason, no data endmarker is recognised by the CRU within the programmable timeout or data words are received outside the data expectation window, the CRU marks the concerned link as erroneous and excludes it from data taking until a manual or automated recovery takes place. The CRU stores all received data in large internal data buffers whose sizes are sufficient to hold even entire calibration data black events. When the CRU has confirmed the reception of end makers on all active links or the timeouts have been reached, the CRU releases its busy status contribution. The CTP considers the detector as busy until all 36 CRUs have released their busy contribution.
4. Once the detector side of the event acquisition is finalised and the data is stored in internal buffers, the CRU is ready to acquire the next event, while at the same time starting to reformat the buffered data and shipping it to the O2 system. The CRU packs the data into packets of a maximum size of 8 kB and equips these packets with Raw Data Headers (RDHs). In addition, TRD customised headers are inserted into the data stream. The headers contain various information, most importantly the trigger timestamp information, in order to link the acquired data to other detector data during reconstruction. For details, refer to Sections 6.6.3 and 6.7.3.

Independent of the CRU processing physics and calibration triggers, the CRU processes Heartbeat (HB) triggers. These triggers synchronise the experiment every LHC

orbit and delimit the so called HB frames. The RDH in the output data stream of the CRU connects every piece of acquired data to a particular HB frame, allowing the data to be synchronised with continuous readout detectors. For details, refer to Section 6.6.

These information overlap with the revised TRD technical design report [49] and the ALICE Upgrade paper [3].

## 6.2. TRD laboratory setup at the University of Münster

Within the scope of this thesis, a TRD laboratory setup has been adapted to assist testing the readout chain and the development of the TRD CRU firmware. The laboratory setup is located at the Institut für Kernphysik at the Universität Münster and it has been extended with additional functionalities over time. Today, the lab setup provides a miniature version of the full Run 3 TRD readout chain up to the FLP. It has simplified the development and the testing during the TRD upgrade significantly and it has reduced the interference of the new developments with the production environment at P2 to a great extent. Tests at P2 are important too, but they have often been restricted, e.g. due to ongoing data taking activities in Run 2.

### 6.2.1. Location and available material

For the laboratory test setup, the TRD laboratory of the Institut für Kernphysik at the Universität Münster has been chosen. It had the advantage, that many important TRD components had been already available and could be repurposed for the upgrade related development alongside with new components.

The availability of material at the Münster laboratory is due to fact that many TRD SMs had been assembled at the Münster laboratory, when the TRD SMs had been commissioned in the past. Also, a large cosmic trigger setup, which is able to generate



cosmic triggers using coinciding photomultiplier signals, had been in place. It had been used for testing TRD SMs with cosmic rays. Details about this work can be found in Ref. [53]. After the early commissioning of the TRD SMs, the Münster laboratory had been predominantly used for developing a completely new TRD for the Compressed Baryonic Matter (CBM) experiment at the Facility for Antiproton and Ion Research (FAIR), Darmstadt.

For the new laboratory setup, the following items have been taken from the laboratory and have been repurposed:

- 1 wall mounted ROC, equipped with all Readout Boards (ROBs), the DCS board as well as the ORI boards. It had been used in the past to test TRD FEE configurations.
- 2 adjustable laboratory power supplies.
- 1 desktop computer.
- 1 Versa Module Eurocard (VME)-crate.
- 1 VME FPGA card for control purposes.
- 1 PCIe interface card.
- 1 TTCex transmitter module.
- various optical fibres and cables.
- 1 attenuator for legacy TTC.

### 6.2.2. New laboratory setup

For the new laboratory setup, the following items have been supplied in addition:


- 1 FLP server machine, Dell PowerEdge 730. This is the same machine as used at P2.
- 1 CRUv2 pre-series version, also compare Section 6.1.2.

- 1 MPO-to-12Lucent Connector (LC) fan-out fibre.
- 1 laboratory Run 3 LTU, bench-top version (compare Section 6.5.1).
- 1 additional ORI-board (compare Section 6.3.1).
- 8 additional ROBs as spare parts, courtesy of the Goethe-Universität, Frankfurt am Main.

Figure 6.3 shows the new TRD laboratory setup in its final version. On the top, the wall mounted and fully assembled ROC is visible. The 2 laboratory power supplies to aliment the ROC with its required supply voltages are below the ROC. On the lower left edge of the table, the VME-crate is visible, which houses trigger related components, most importantly, the TTCex module. On top of the VME-crate, there is the new Run 3 laboratory LTU in its bench-top version (compare Section 6.5.1) and the desktop computer. On the right side of the VME-crate, there is the FLP server machine which also houses the CRU. On the right side of the FLP, there is a gateway machine, configured as router. It provides remote access to the entire laboratory, e.g. for remote tests from CERN. This gateway machine is shared with other laboratory setups, e.g. those for the CBM experiment.

The components have been connected similarly to the production setup at P2 and electric installations have been renewed. In addition, all the software used at P2 has been installed in the laboratory too. The laboratory TRD ROC can be successfully configured, allowing for almost all features to be tested in this laboratory prior to their deployment at P2.

Note that Fig. 6.3 shows only the final version of the laboratory setup, but many intermediate versions had to be built due to the fact that some Run 3 components had still been in development and had not been readily available from the beginning of the work for this thesis. This concerned e.g. the FLP which has been a completely different machine in the beginning and therefore lacked the support for PCIe bifurcation. Also, the Run 3 LTU has not always been available. In the beginning, legacy trigger components from the Münster laboratory have been used to trigger the TRD ROC.



**Figure 6.3.:** New laboratory setup for testing the upgraded TRD readout chain at the Institut für Kernphysik, Universität Münster.

### 6.3. Consideration of the TRD FEE optical drive capabilities

During Run 1 and Run 2, the GTU has sent the TRD data from the cavern to CR1. Besides its data processing tasks, the GTU has thus also acted as an amplifying stage for the optical signals from the TRD FEE. Since the GTU will be obsolete in Run 3 and the FEEs are interfaced with the CRUs directly, major feasibility concerns existed questioning whether the optical signal from the TRD FEE is sufficient to drive a significantly prolonged optical path.

This section investigates the available optical power margins using theoretical calculations as well as estimations based on practical measurements. These investigations have been carried out before the actual decision of installing the CRUs in CR1. Based on these investigations, all concerns could be eliminated prior to completing the installation.

In the end of this section, measurements of the real conditions with the CRUs after their installation in CR1 are presented.

The information in this section which concern the already existing installation and the ORI boards are based on Ref. [38]. The specifications for the estimates in this section can be found in the datasheets of the involved components [54], [55] and [56].

#### 6.3.1. Theoretical estimation of the optical margin

##### Optical Readout Interface board


The TRD uses a custom optical link, the ORI, to connect the TRD FEE to the CRUs which runs at a gross optical serial link data rate of 2.5 Gb/s. Each ORI transmitter circuit board is equipped with a Vertical Cavity Surface Emitting Laser (VCSEL) diode, emitting light at a wavelength of 850 nm and an absolute maximum average output power

of up to 1000  $\mu\text{W}$ . The diodes have been manufactured by ULM Photonics and their type number is ULM850-02-TN-ULCBPP. These diodes are driven by an integrated driver IC from Texas Instruments (TI) with the type number LTC5100. The LTC5100 driver IC operates in Automatic Power Control (APC) mode, i.e. it regulates the optical power by monitoring a feedback current traversing a monitor diode that is integrated into the VCSEL transmitter diode. The setpoint for the regulation can be changed remotely, such that the maximum average power output of 1000  $\mu\text{W}$  of the diode can be achieved. All configuration parameters of the LTC5100 driver IC have been carefully optimised in the past minimising the signal eye closure. These optimisations have been carried out by measuring the electrical signal eye diagrams with a digital storage oscilloscope directly at the VCSEL diode and comparing it to the electrical signal eye diagram measured directly behind a test receiver. The datasheet of the LTC5100 driver IC provides an optical eye diagram, valid for operation at a link data rate of 2.5 Gb/s using a VCSEL diode with a wavelength of 850 nm and an average optical transmission power of 300  $\mu\text{W}$ . These conditions correspond to the situation of the TRD. At this working point, the datasheet reports an optical extinction ratio of 10 dB. The eye diagram from the datasheet is shown in Fig. 6.4.

## Optical Path

The VCSEL diodes are equipped with a support to allow plugging a LC directly into it. From the LC, an optical fibre is pulled to the GTU below the muon arm where it ends at an LC patch panel. For the TRD upgrade, the optical fibres have to be extended from here to CR1. For the extension, it is necessary to connect an LC-to-MPO patch cord to the patch panel, which extends the optical fibres up to a new MPO patch panel in the cavern. For details, also compare Section 6.4. From there, multi-mode optical fibres of Optical Multimode 3 (OM3) quality are pulled up to the counting room. Their total length is 130 m and the specified worst-case attenuation is 3.5 dB/km. In CR1, these optical fibres end at another MPO patch panel. From this patch panel, MPO-to-MPO patch cords are planned in order to connect the optical fibres to the CRUs. For LCs and MPO connectors, a worst-case attenuation of 0.5 dB each can be assumed. On the CRU PCB, there is another short optical fibre which connects the MPO connectors at the back plate

### Optical Eye Diagram at 2.5Gbps with 850nm VCSEL



**Figure 6.4.:** Eye diagram adapted from the LTC5100 datasheet: Emcore mode LC-TOSA VCSEL,  $2^{15}$  PRBS, 10 dB extinction ratio, 300  $\mu$ W average power [54].

internally to the minipods, which are directly mounted on the PCB through a PRIZM connector. This assembly has a specified worst-case attenuation of 2 dB. All attenuation estimates in this calculation are taken from the component datasheets wherever possible. In case no specifications could be found, they are generally accepted conservative estimates.

Unknown elements in this calculation are the exact length of the optical fibres from the detector to the GTU patch panel and the length of various patch cords, as well as the exact position of the TRD CRUs within CR1 which had not been chosen at the time of this investigation. As a very conservative assumption, every uncertain optical fibre length is assumed to be 50 m.

Table 6.1 summarises all conservatively estimated contributions to the total attenuation.

**Table 6.1.:** Contributions to the attenuation on the optical path.

Contribution	Conservatively estimated attenuation
LC at ORI board	0.5 dB
50 m optical fibre between detector and GTU patch panel	0.175 dB
LC at GTU patch panel	0.5 dB
50 m LC-to-MPO patch cord	0.175 dB
MPO connector at new upgrade patch panel in cavern	0.5 dB
130 m OM3 optical fibre from cavern to CR1	0.455 dB
MPO connector at new upgrade patch panel in CR1	0.5 dB
50 m MPO-to-MPO patch cord	0.175 dB
MPO connector at CRU	0.5 dB
CRU internal PRIZM connector to Minipod	2 dB
Total sum	5.48 dB

### Receiver characteristics and optical power margin

The receiver used on the CRU side is manufactured by Avago Technologies and its type number is AFBR-821vx3Z. The datasheet [56] states a receiver sensitivity of  $-12$  dBm expressed in the Optical modulation amplitude (OMA), which corresponds to the difference between the maximum power  $P_{\max}$  and the minimum power  $P_{\min}$  of the modulated optical signal.  $P_{\min}$  and  $P_{\max}$  are related to the optical average power  $P_{\text{avg}}$  by

$$P_{\text{avg}} = \frac{P_{\min} + P_{\max}}{2} \quad (6.1)$$

and the extinction ratio  $r$  (compare Section 6.3.1) is defined as

$$r = \frac{P_{\max}}{P_{\min}}. \quad (6.2)$$

It can be deduced that

$$\text{OMA} = 2P_{\text{avg}} \cdot \frac{r - 1}{r + 1}. \quad (6.3)$$

Considering a scenario of an average power of  $P_{\text{avg}} = 300 \mu\text{W}$ , which is a typical value for the TRD, emitted at the source and the given extinction ratio of 10 dB from the datasheet, we have  $r = 10$  and  $\text{OMA} = 490 \mu\text{W} = -3.1 \text{ dBm}$  at the source. After facing the attenuation of 5.48 dB in the worst case, the receiver still sees an OMA of  $-3.1 \text{ dBm} - 5.48 \text{ dB} = -8.58 \text{ dBm}$ . Considering the receiver sensitivity of  $-12 \text{ dBm}$ , an optical power margin of 3.42 dB is left. This margin is sufficient, especially if one considers that all assumptions are conservative. Apart from that, an average output power as low as  $300 \mu\text{W}$  has been assumed. This average power can still be adjusted up to  $1000 \mu\text{W}$ . The datasheets unfortunately do not provide information about the extinction ratio at different average powers, which makes it difficult to estimate the corresponding margin quantitatively. Assuming a negligible dependence of the extinction ratio on the average power, an accordingly even larger optical power margin is available.

Assuming that the extinction ratio does not change with the average power, one can crosscheck the result and roughly calculate the minimum average power needed at the source to have a margin of 0 dB at the receiver. This minimum average optical source power is  $P_{\text{avg}} = 136 \mu\text{W}$ . Chambers with an even lower average power setting, require an adaption of the laser diode current setpoint parameters.



### 6.3.2. Investigation on the actual current setpoints in Run 2


The optical power on the different TRD Half Chamber (HC) links varies. One reason for the variations is the fact that certain chambers have been calibrated in the past, aiming for a constant optical power level at the GTU, but others have been operated throughout Run 2 at default laser diode current setpoints, because no signal reception issues have been observed in Run 2.

Another reason for the fluctuations could be abnormally high attenuations in some optical components within the central barrel zone which is difficult to access for maintenance. Especially the connections between the laser diodes and the optical fibres use non-Commercial off-the-Shelf (COTS) components and might be susceptible to e.g. dust collection.

Considering these two influences on the optical power variations across the detector, it might have been possible that HCs with abnormally high attenuations existed, which could have been compensated for, in the calibration, by setting the laser diode current setpoint close to its absolute maximum rating of 12 mA. This possibility poses a major caveat against the argument of having available additional optical power margin gains by adjusting the setpoint, as outlined in Section 6.3.1.

For a better understanding of the situation, the setpoint data has been collected for almost all HCs along with the average optical power arriving at the GTU of the corresponding optical link. The collected data are shown in Fig. 6.5.

A few chambers have not been operational in Run 2 at the time of this measurement and an optical power of 0 is reported in plot. However, the plot still reports valid setpoint data for these chambers. Furthermore, the Fig. 6.5 shows a large number of HCs with an optical average power between 50  $\mu$ W and 250  $\mu$ W at the GTU. For these channels, the current setpoint is relatively low between 1 mA and 4 mA. These are the channels which have been calibrated in the past. Another large group of HCs exhibits large fluctuations



**Figure 6.5.:** Optical powers received at GTU vs. laser current setpoint.

in the measured GTU power, but a constant laser current setpoint of 4.5 mA. These are channels which have never been calibrated in the past and work at their default settings. No HC with a laser current setpoint close to its absolute maximum rating of 12 mA has been found in this investigation.

### 6.3.3. Measurement of the Optical Path Attenuation

Far in advance of the availability of the custom CRU board and its deployment in CR1, a measurement of the actual attenuation has been carried out at P2 as part of this work. The measurement has been conducted using the chamber in SM 07, stack 1, layer 2 (in the following named as 07\_1\_2). This choice has been motivated by the fact that this chamber exhibited issues with HV and a potential loss of the readout capabilities of this chamber would not have any impact on Run 2 data taking operations. The chamber did not have any issues concerning the readout tree. The optical powers have been measured

**Table 6.2.:** Composition of the prolonged optical test path.

Element	Length
TRD patch panel → Muon spectrometer patch panel	15 m
Muon spectrometer patch panel → CR1 patch panel	130 m
CR1 patch panel → CR1 patch panel (loop)	5 m
CR1 patch panel → Muon spectrometer patch panel	130 m
Muon spectrometer patch panel → GTU	15 m
Total length	295 m

using the GTU. It provides capabilities to read the average optical power arriving at a certain optical input using command line tools.

In order to measure the attenuation, the received optical power has been measured before and after prolonging the optical fibre. At the time of the measurement, no optical fibres for TRD between the cavern and CR1 had been available yet. Only a few test optical fibres dedicated for the muon spectrometer had been in place. These test optical fibres are identical to the ones planned for the TRD, but they do not start at the TRD. The starting point in the cavern is at the dipole magnet of the muon spectrometer. In order to conduct the measurement with the available installation, prolongation patch cords have been installed between the GTU and the muon spectrometer dipole magnet patch panel. These prolongation patch cords have been specifically manufactured for this measurement by a supplier contracted by CERN. Since the GTU could not be displaced to CR1 for practical reasons, an optical fibre loop has been installed in CR1 leading the signal back to the GTU through a separate test optical fibre. The total prolonged optical path is given in Table 6.2.

**Results** The optical power read by GTU without prolongation has been  $578\mu\text{W}$ . With the 295 m optical fibre loop in place, the reading dropped to  $289\mu\text{W}$ . This is a drop of 50 %. Considering that this value has been measured with signal travelling back and forth, the fraction of signal power lost on one way is

$$1 - \sqrt{0.5} \approx 30 \, \%.$$

This corresponds to a one-way attenuation of approximately 1.5 dB. This value compares very well to the relevant theoretical values given in Table 6.1, i.e. the theoretical worst case estimation is sufficiently reliable.

#### 6.3.4. Verification of data reception in CR1

Apart from the attenuation measurement, a first data reception attempt has been performed in CR1 at a very early stage, when neither the final optical installations nor the CRU had been available. For this test, the optical fibre loop in CR1 has been opened and connected to an Altera Arria 10 GX FPGA development board. This board features the same FPGA as the CRU and has been used for early developments before the actual CRU PCB production. The FPGA has been loaded with a very early development version of the TRD specific FPGA firmware, which has been already capable of establishing synchronisation with the TRD custom links.

The reception of control words from the optical link with the development board has worked correctly. While this has been an important result, the comparability of this test with the final CRU is limited, as the development uses different optical components. It uses an SFP RX module instead of the the Avago AFBR-821vx3Z receivers. Unlike for the Avago receivers, no comparable receiver sensitivity specification (compare Section 6.3.1) has been available in the datasheet of the SFP module.

This test has been repeated connecting the long test optical fibre temporarily to HCs with very low optical power (less than 100  $\mu$ W at the GTU). Links with these properties are rare and with such links, no data could be received in CR1. This supports the assumption that a modification of the laser diode current setpoint might be necessary in the future for some links.

### 6.3.5. Data taking with the GTU and the optical fibre loop


After the measurements described above, the optical loop via CR1 for chamber 07\_1\_2 has been left installed for the entire remaining time of Run 2. Even though the optical fibre has been even longer than what is expected for Run 3, no long term issues with data taking on this channel have been observed.

### 6.3.6. Sensitivity test with the CRU in the lab at the Universität Münster

Before the installation of the CRU, tests have been conducted with the CRU and a TRD chamber in the lab. A lab setup with a TRD ROC and a CRU has been set up in the TRD lab of the Institut für Kernphysik, Münster (compare Section 6.2). In these tests, the communication between a TRD chamber and an actual CRU with the final optical receiver component has been demonstrated for the very first time.

The CRU has been configured to keep the link synchronisation with an optical channel from the TRD chamber established and to reconnect automatically in case of a loss of synchronisation. In this situation the laser current setpoint of the corresponding ORI board of the TRD chamber has been varied, observing the optical power received by the CRU as well as the rate of synchronisation losses. This test has been conducted using a short optical link of approximately 5 m. The Figures 6.6 and 6.7 show the results.

The test shows that in a practical lab environment, the CRU allows for a reliable data reception at an optical average power as low as  $60\text{ }\mu\text{W}$ , achievable with a very low laser current setpoint as low as  $1.22\text{ mA}$ .

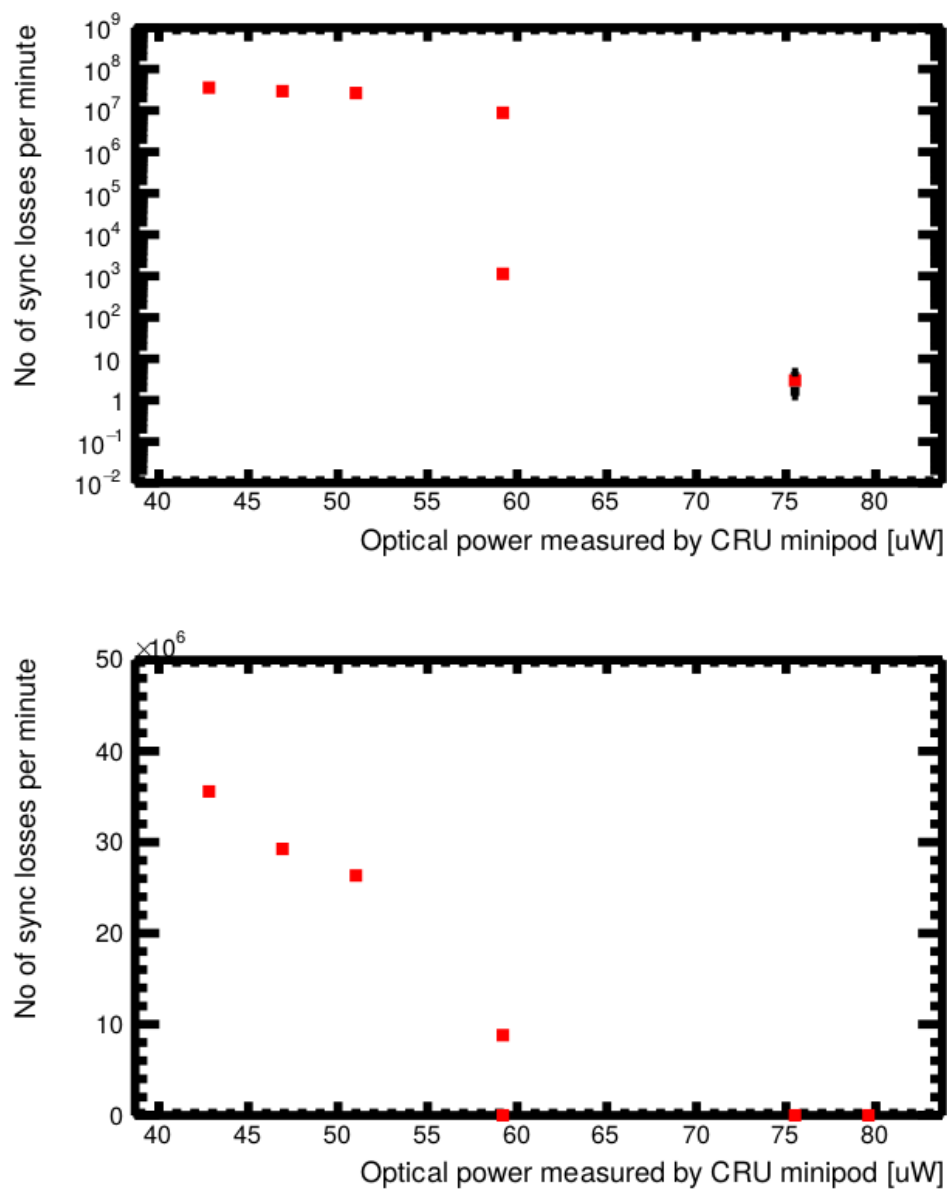


**Figure 6.6.:** Optical power measured by the CRU as a function of the laser current setpoint of a lab ORI board.


### 6.3.7. Optical power measurements with the CRUs after their installation in CR1

After the installation of the CRUs, it has been possible to use them for optical power measurements instead of the GTU and confirm the theoretical calculations and estimations in the previous sections at LHC P2. The optical power values are measured by the Avago Minipods, which feature Inter-Integrated Circuit (I2C) bus connections with the FPGA. The optical power can be read from the minipods with dedicated I2C commands.


The results presented here have been measured for SM05, which is connected to 2 CRUs, one for the A-side and another one for the C-side. Figures 6.8 and 6.9 show the optical power received by the GTU before the installation of the CRU together with the optical power received by the CRU after the installation. For some channels, a GTU power of 0 is reported because not all channels have been operational at the time of the measurement. In general, the CRU powers are lower than the GTU powers due to the additional attenuations. For very few channels, it is vice versa and the optical power



**Figure 6.7.:** Rate of synchronisation losses at logarithmic (top panel) and linear (bottom panel) scale. Vertical bars are statistical errors. They are mostly hidden by the points.




**Figure 6.8.:** Optical powers received at CRU vs. optical powers previously received at the GTU, SM05, A-side. The different channels are identified as SM\_stack\_layer\_HC.




**Figure 6.9.:** Optical powers received at CRU vs. optical powers previously received at the GTU, SM05, C-side. The different channels are identified as SM\_stack\_layer\_HC.





**Figure 6.10.:** Measured effective additional attenuations SM05, A-side. The different channels are identified as SM\_stack\_layer\_HC.



**Figure 6.11.:** Measured effective additional attenuations SM05, C-side. The different channels are identified as SM\_stack\_layer\_HC.

actually improved. This can be explained with rare defects of the previously used patch cords and connectors, e.g. caused by excessive contamination.

The Figures 6.10 and 6.11 show the ratio between CRU optical powers and GTU optical powers in SM05 expressed in dB. Implausible values are explained by inoperative channels during the measurement. The observed values are in line with the theoretical worst case estimations given in Table 6.1.

## 6.4. Readout Fibre Installations at Point 2

### 6.4.1. Mapping of CRU ORI link numbers to half chambers

The TRD sends its data via 1044 optical links. Most SMs have 60 optical links, except for the SMs 13,14,15 which only have 48 optical links due to the missing stack 2 in front of PHOS. Each CRU provides effectively 30 optical inputs and in total 36 CRUs are reserved for the TRD. A mapping scheme of HC-to-CRU optical link input has been developed, which accounts for all of the following aspects:

1. For simplicity, the mapping shall be such that no CRU serves distinct SMs.
2. For simplicity, the setup shall be symmetric and reflect the structure of the detector. No special attention shall be given to the PHOS hole, i.e. the concerned SMs are treated like all other SMs and a few available CRU inputs shall be left open. The possible economising of optical inputs is not worth in case it causes an increase in complexity or other downsides.
3. For simplicity, all CRUs shall use the same optical inputs and shall not require different firmware. This requirement is necessary because of the mechanical construction of the CRU and the fibres: The CRU has 3 MPO receptacles, where each one internally connects 12 optical links. An MPO-to-12 LC single fibre fan-out adaptor fibre is required in order to connect the TRD links. Mechanically, there are in total 36 optical inputs, though only 30 optical links can be used per CRU for other reasons. In order to be able to use the same firmware everywhere, it is important that the

mapping scheme maps the 6 unused optical inputs to the same mechanical position for every CRU.

4. Detector layers which are stacked on top of each other shall not be read out by different CRUs. This allows to implement tracking algorithms in the CRUs at a later time, as the data of every TRD track stays concentrated within the same CRU. Currently there is no such requirement, but in case it becomes necessary in the future, the mapping scheme shall already be compatible.

In order to fulfil all these requirements, for every SM, 2 CRUs have been attributed. The 60 links coming from one SM can be distributed exactly to the 2 CRUs with 30 links/CRU each and requirement (1) is easily fulfilled. For the PHOS hole inputs can be left unconnected, requirement (2) is thus easily fulfilled too.

The mapping of links within each SM to the 2 corresponding CRUs is more complex. The number of stacks in one SM is 5, which is uneven. In order to distribute the fibres equally to 2 CRUs, it cannot be avoided to distribute at least 1 stack across different CRUs. In order to nevertheless fulfil requirement (4) and not lose the feasibility of a future online tracking, the splitting between the 2 CRUs happens in the middle of the stack 2 at the HC boundary. HC 0 of the stack 2 is attributed to one CRU and HC 1 is attributed to the other CRU. In this way, each CRU still has the track information of all HCs which are stacked on top of each other locally available at all locations.

Every stack consists of 6 layers and thus 12 HCs in total. For simplicity by symmetry, it is a fortunate coincidence that this number 12 corresponds to the number of single channels within one MPO connector. The MPO connectors are connected by special fibres which also bundle 12 single fibres together. It has therefore been decided to use in principle one MPO input for one TRD stack. This is fully possible for the stacks 0,1,3 and 4. Stack 2 remains special, as it has to be split across 2 CRUs and thus 2 MPO fibres necessary, each one only half used.

For both CRUs, stack 2 is connected through a half-used MPO connector. In order to fulfil requirement (3), it is important that this half-used MPO connector plugs into each CRU at the same place. Therefore the mapping of the CRU serving the A-side of the SM has to be very different than the mapping of the CRU which is serving the C-side of the SM. In fact, the mapping scheme is mirrored at the HC boundary of the stack 2: For the A-side CRU the CRU ORI link inputs systematically numbered from 0 to 29 are populated starting from stack 0 up to the HC boundary of stack 2. For the C-side CRU, the CRU ORI link inputs systematically numbered from 0 to 29 are populated starting from stack 4 backwards up to the other side of the HC boundary of the stack 2.

For practical reference, the mapping of each CRU ORI link input to stacks, layers and HCs is given in Table 6.3 for the A-side CRUs and in Table 6.4 for the C-side CRUs.

### 6.4.2. Installation

The installation of the fibres according to the mapping scheme in the ALICE UX25 cavern has been carried out within the scope of the work for this thesis. In total, 108 patch cords of the MPO-to-12 LC single fibre type and a length of 3 m each have been purchased through a CERN contracted supplier. These patch cords have been installed in the GTU racks, such that the GTU is bypassed and the signal from the HCs is directly passed on to the CRUs in CR1. Prior to that, all patch cords have been labeled on both ends, clearly marking the FLP, the CRU and the CRU input to which they belong. The GTU has been left in place together with all required patch cords. In case that the GTU shall be recommissioned in the future, only fibre replugging is required. However, currently, no such planning exists.

**Table 6.3.:** Mapping of CRU ORI link inputs to stacks, layers and HCs for the A-side.

CRU ORI link number	Stack	Layer	Half chamber
0	0	5	0
1	0	4	0
2	0	3	0
3	0	2	0
4	0	1	0
5	0	0	0
6	0	5	1
7	0	4	1
8	0	3	1
9	0	2	1
10	0	1	1
11	0	0	1
12	1	5	0
13	1	4	0
14	1	3	0
15	1	2	0
16	1	1	0
17	1	0	0
18	1	5	1
19	1	4	1
20	1	3	1
21	1	2	1
22	1	1	1
23	1	0	1
24	2	5	0
25	2	4	0
26	2	3	0
27	2	2	0
28	2	1	0
29	2	0	0

**Table 6.4.:** Mapping of CRU ORI link inputs to stacks, layers and HCs for the C-side.

CRU ORI link number	Stack	Layer	Half chamber
0	4	5	0
1	4	4	0
2	4	3	0
3	4	2	0
4	4	1	0
5	4	0	0
6	4	5	1
7	4	4	1
8	4	3	1
9	4	2	1
10	4	1	1
11	4	0	1
12	3	5	0
13	3	4	0
14	3	3	0
15	3	2	0
16	3	1	0
17	3	0	0
18	3	5	1
19	3	4	1
20	3	3	1
21	3	2	1
22	3	1	1
23	3	0	1
24	2	5	1
25	2	4	1
26	2	3	1
27	2	2	1
28	2	1	1
29	2	0	1

## 6.5. Upgrade of the CTP

For Run 3, the CTP undergoes a major upgrade, which includes a redevelopment of many components. Only the changes which are most relevant for the TRD are discussed in this section.

### 6.5.1. Fully revised Local Trigger Unit

The CTP has exchanged the LTUs of all detectors by different and new models with new firmware. The new LTUs support the trigger distribution via the new TTC-PON (compare Section 6.5.2). Furthermore, the new LPU provides now the possibility of generating trigger signals which can be directly interpreted by the TRD FEE. In Run 2, a custom translator unit (LTU-T) had been necessary.

The new LPU is available in a rack mounted version for the ALICE UX25 cavern and in a lab bench-top version. The CTP provides special software tools, which allow operating the new LPU in standalone mode, facilitating the generation of triggers in the lab as well as during commissioning at P2.

For the TRD, 2 new LTUs are available. One of them is installed in the ALICE UX25 cavern and the other one is a bench-top version and is installed in the TRD lab of the Institut für Kernphysik, Münster, for development and testing purposes.

### 6.5.2. Trigger and Timing Control over Passive Optical Network (TTC-PON)

In Run 3, PON is used to distribute the trigger and timing signals from the CTP to the CRUs. Compared to the previous TTC system, TTC-PON offers the advantages that it adds an upstream channel and that it uses exclusively COTS components, simplifying

maintainability in the long run compared to specially developed prototypes.

In the case of a TTC-PON network, the Optical Line Terminal (OLT) is attached to the LTU at the CTP and the Optical Network Units (ONUs) are attached to the CRUs. The OLT and the ONUs are interconnected through a single mode fibre network with an optical splitter. The LTU with OLT is installed in the experimental cavern UX25 of ALICE, while optical splitters and CRUs are installed in CR1.

The LTU broadcasts the TTC signals to all CRUs simultaneously and the CRUs alternate in a round robin scheme sending the so called heartbeat acknowledge messages to the CTP. The arbitration time of each ONU is 125 ns.

The CTP uses TTC-PON internally, in order to connect the actual central trigger processor board with the detector LTUs. In these CTP-internal networks, the OLT is attached to the central trigger processor board and the ONUs are attached to the LTUs. The arbitration time of each ONU is 125 ns too.

### 6.5.3. TTC-PON Latency Considerations

The TRD has very challenging latency requirements. The TRD CRUs use a reserved bit within the heartbeat acknowledge message in order to signal their individual busy status to the CTP. Since the latency at which the busy signals are conveyed to the CTP has a direct impact on the maximum achievable event readout rate, the TTC-PON latency is absolutely critical.

Typically, a single optical network per detector is used which serves all CRUs of that detector. The time until a specific ONU has to wait before it falls into the arbitration window for sending upstream data, increases linearly with the number of CRUs connected to the same network. Given the critical latency requirements for the TRD, it has been



decided to exceptionally install the technical maximum of 9 optical networks for the TRD and group its 36 CRUs accordingly into groups of 4 CRUs each. The technical limit of 9 networks is imposed by the number of available SFP cages to insert OLTs on the LTU PCB design.

The total maximum latency of the upstream busy travelling from a CRU to the CTP is the sum of arbitration delays and propagation delays. If one assumes an effective propagation speed of  $1\text{ m}/5\mu\text{s}$  and a rough total optical fibre length of 150 m in order to estimate the propagation delay, the total expected latency is

$$3 \cdot 125\text{ ns} + 4 \cdot 125\text{ ns} + 150\text{ m} \cdot 1\text{ m}/5\mu\text{s} = 1.5\text{ }\mu\text{s}.$$

For the TRD busy, a latency as short as possible is desirable and the latency of  $1.5\text{ }\mu\text{s}$  appears reasonable considering the readout duration value range between  $4\text{ }\mu\text{s}$  and  $8\text{ }\mu\text{s}$  for tracklet only readout.

## 6.6. The CRU Common Logic Firmware in the context of the TRD

This section introduces the CRU common logic. It is a part of the CRU firmware whose design is partially shared across all CRU based detectors. However, the case of the TRD is special and this section discusses the usability of the common CRU firmware components for the TRD in details and presents ideas on how to adapt them.

The description of the common logic is based on Ref. [57].

### 6.6.1. CRU FPGA firmware

The CRU FPGA firmware consists of a common logic, which is similar for all CRU based detectors, and a detector specific logic, called custom user logic. The common logic is developed by a central team and is in principle able to receive data from a detector via the optical links and forward them to the FLP via the PCIe interface. In addition, the common logic supports sending data in the upstream direction of the optical links. This allows detector front ends to be configured and to receive trigger data via the CRUs and the optical links. For detectors which do not have any needs for customised FPGA data processing, it might be sufficient to use the common logic without adding custom user logic. This approach has for example been chosen by the ITS. Detectors with FPGA based processing requirements must implement their user logic and encapsulate it in a Very High Speed Integrated Circuit Hardware Description Language (VHDL) entity with predefined input and output ports to cooperate with the common logic. This approach is for example followed by the TPC and the Muon Chamber (MCH).

However, concerning the TRD, the situation cannot be compared to other detectors. An important aspect is that the TRD does not upgrade its Front End electronics in the course of the Run 3 upgrade. Therefore, the TRD has to deal with the specific properties of the Front End Electronics at the CRU level. The following issues have to be addressed:

1. The TRD FEE uses a custom link transport format, which is incompatible with the Giga-Bit Transceiver (GBT) link data format. In the context of the upgrade, GBT has been developed as a standard at CERN. It is used by the other CRU based detectors in ALICE in Run 3 and the CRU firmware developed by the common team only supports the link operation in GBT mode. This means that in addition to the actual TRD user logic, other components of the common firmware have to be rewritten for TRD compatibility too.
2. The links coming from the TRD FEE do not support receiving data and are strictly unidirectional. Therefore, the approach of configuring and triggering the detector FEE through the CRUs, is technically impossible for the TRD. The Front End needs

external channels via DCS and additional trigger fibres to operate correctly. The CRU firmware has to be tailored to this different approach.

3. The TRD FEE sends its data without any headers, i.e. the received tracking data cannot be attributed to a timestamp and geometrical location. This has been an important design decision for the past Runs where latency has been critical because the TRD has been working as a trigger detector and has provided fast trigger decisions. The CRU has to keep track of these meta data and add them to the data in order to make sure that the TRD data are compatible with other detectors and can be processed together with their data during reconstruction.
4. The TRD FEE is made of custom ASICs and they cannot deal with new triggers as long as data from a previous trigger is still being processed. The processing is not fast enough to cope with the envisaged high trigger rates expected in ALICE for the Run 3 readout (up to 1 MHz typically). Therefore, the FEE has to be protected from new triggers until the previous data has left the FEE. Since the CRU is the only real time capable system after the FEE, it has to implement this protection. This implementation requires an upstream data path from the detector specific user logic to the CTP. This path is unique for the TRD, as most detectors with only FPGA based data processing needs, do not need this direct signal path.

### 6.6.2. Common CRU Firmware

The structure of the common firmware developed by the central team is shown in Fig. 6.12. It shows that the firmware structure is designed in a modular manner. The shown modules are interconnected within a head entity (`top.vhd` and `core.vhd`).

The GBT wrapper bundles the functionality required for optical link data reception and transmission. In order to handle high speed serial data throughputs of multiple Gbit/s, which are common for optical interfaces, specialised circuitry, called Intellectual Property (IP) cores, inside the FPGA chip are required for data serialisation and deserialisation. This is mandatory because the programmable FPGA fabric is limited to typical frequencies in the MHz range for timing closure. The IP cores used for the implementation

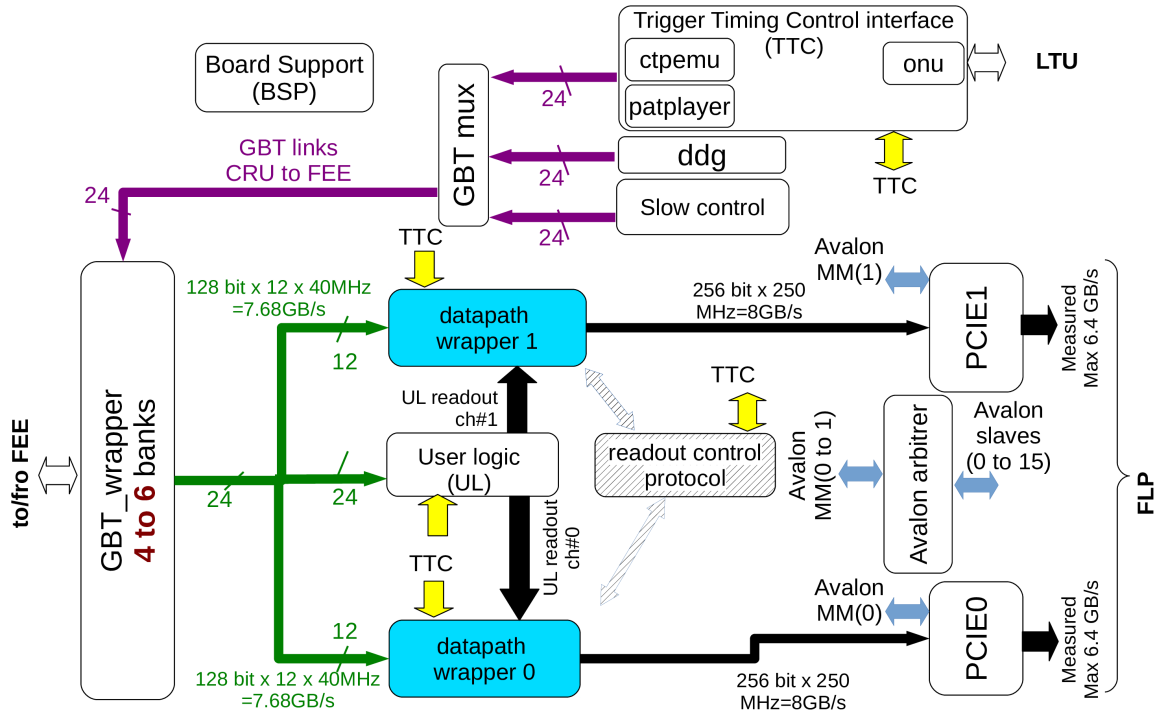


Figure 6.12.: Overview of the common CRU firmware [57].

of the GBT wrapper are located at the borders of the Arria 10 silicon die and are also referred to as transceiver or XCVR IP cores.

The TX functionality requires stable low jitter reference clocks for the data serialisers. These clocks are generated by Phase Locked Loop (PLL) circuits. The number of available PLLs on the FPGA is limited such that multiple serialisers have to share a single PLL.

The GBT wrapper can handle up to 24 channels, where every channel consists of 1 RX link and 1 TX link. In order to make the sharing of the PLLs technically possible, the channels are grouped into banks of 6 channels each.

The GBT wrapper is responsible for sending trigger data and detector configuration data upstream to the FEE. A multiplexer (GBT mux) is available to switch between the different data sources. For debugging purposes, it is possible to use a CTP emulator

(ctpemu) instead of real trigger data from the LTU. Also, a generator of artificial data (ddg) is available through the multiplexer.

The GBT wrapper receives downstream data from the detector FEE via optical links, decodes the GBT protocol and passes the decoded data on to downstream firmware components.

For debugging purposes, the GBT wrapper features a loopback mode which redirects the upstream data from the GBT mux to the downstream firmware components. This allows testing all CRU and firmware functionalities without actually attaching a detector through optical links.

For the TRD, the GBT wrapper is not usable as it supports only the GBT protocol. The TRD FEE uses a custom link protocol and does not support the reception of upstream data via optical links. Therefore, both the GBT wrapper and the debugging components connected through GBT mux need completely different implementations in order to support the TRD with similar functionality.

The data flowing downstream from the GBT wrapper are split into two halves, corresponding to 12 RX channels each. Both halves are processed independently which requires essentially a duplication of all generated data processing hardware up to the FLP. This is also the primary reason why it has to support PCIe bifurcation.

The downstream data takes multiple paths through the logic in parallel. It is received by 2 datapath wrappers and possibly a user logic in case the detector provides one. Assuming that the user logic is only performing an optional FPGA based data preprocessing, this design allows to switch between a raw data readout, entirely circumventing the user logic, and a readout of user logic processed data.

The datapath wrappers are responsible for basic data processing, filling raw data header fields, performing clock domain crossing and duty cycle adaption on the received data in order to provide the data in packages via 2 256bit@250MHz data push lanes to the DMA engines (PCIE0 and PCIE1). The data path wrapper has a large FIFO integrated at the end of its pipeline allowing to smoothen away short backpressure periods which can arise in case of short arbitration outages on the PCIe interfaces from the FLP. The DMA engines are responsible for the low level handling of the PCIe interface and the data transfer into the FLP memory via the FLP DMA controller. It uses XCVR IP cores too, in order to communicate via the PCIe interface.

In case that a user logic is implemented, it receives the data directly from the GBT wrapper and is responsible for the splitting into 2 endpoints. It injects its output data into the data path wrapper in parallel to the inputs for the GBT wrapper. For compatibility, the user logic has to follow a standardised format (compare Section 6.6.3).

The datapath wrappers do not work for the TRD because they expect the data in a format which the GBT wrapper produces. At this stage, the data are still similar to the GBT link data and data widths, clock speeds and duty cycles and these are unnatural to what a similar implementation for the custom TRD links requires. The direct data path through the data path wrappers is thus not usable. Only the direct user logic data inputs of the data path wrappers are used by TRD. It is worth noting that a direct data path without a correctly working TRD user logic would not work for the TRD readout also because of other reasons, e.g. because no busy protection would be present to steer the readout. Considering these facts, it has been decided to not implement an equivalent for the direct path feature in the CRU firmware for the TRD.

The firmware behind the user logic and/or data path wrappers features components controlling the data flow (readout control protocol). The readout control protocol allows for data throttling and for discarding erroneous data. The mechanism works as follows: for every LHC orbit, each CRU readout control protocol reports to the CTP whether the detector data have been gathered successfully (also called heartbeat acknowledge

message). The CTP generates a heartbeat map from these messages, which contains for a particular orbit the status report from all CRUs of the experiment. In case that an important subsystem reports an error, the CTP can decide to signal back to all CRUs of the experiment to discard the entire data for that orbit.

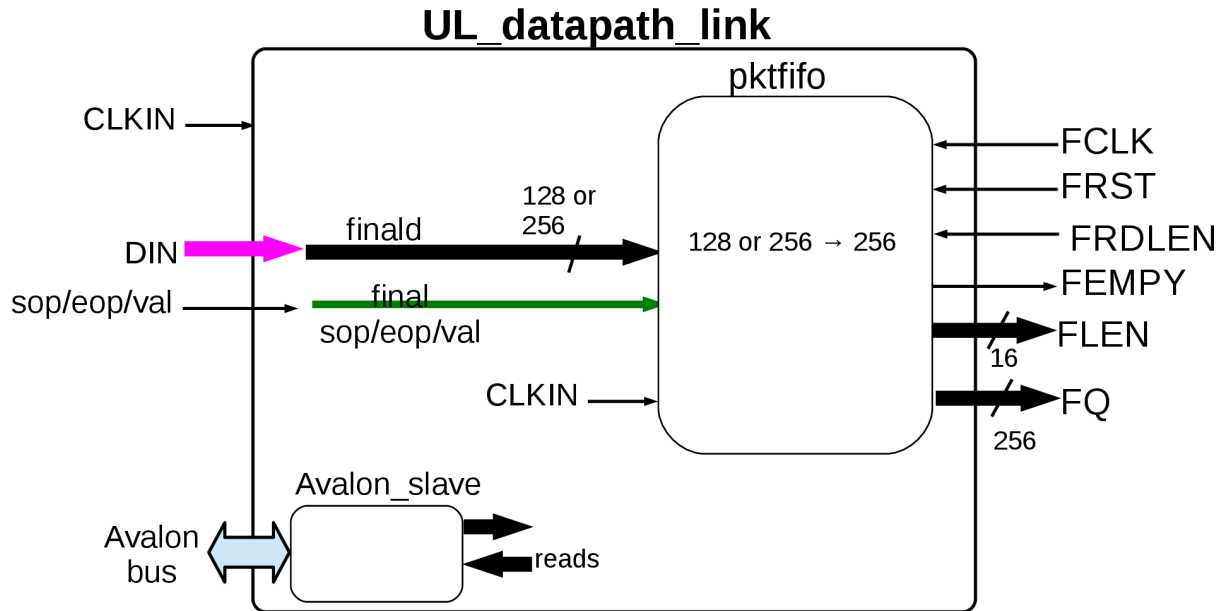
In case of the TRD, this feature works independent from the TRD user logic and thus does not take into account errors which happen at the level of the TRD FEE or in the TRD user logic. This behaviour is intended, as there will not be any cases of errors in the TRD which would justify discarding data from other detectors, especially considering the fact that the TRD FEE limits the fraction of events which can be recorded by its busy times.

The Avalon arbiter is responsible of managing the slow control network within the CRU firmware. The Avalon arbiter acts as a master on the firmware internal Avalon Bus. The master can be controlled by software through the PCIe interface. Every logical entity of the CRU can be equipped with corresponding Avalon slaves which must be connected to the Avalon Bus. They provide readable and/or writable 32 bit wide registers for slow control and monitoring purposes. The bus width is 32 bit at a clock frequency of 100 MHz. Only 1 register can be accessed at a time.

### 6.6.3. O2 Data Format

The CRU has to deliver the data to the O2 system in packetised format, which has to be compatible with the O2 system. Every packet must be equipped with a header, referred to as RDH, and can transport payload after the header. A maximum total packet size (header and payload) of 8 kiB must not be exceeded. No trailer is required after the payload.

The TRD adheres to the Raw Data Header, version 6 (RDHv6). The RDHv6 consists of 512 bits, transferred during 2 clock cycles. The definition of the 2 RDHv6 words is given in the Figures 6.15 and 6.16. The RDHv6 word 0 is sent in the first valid clock cycle of each packet, RDHv6 word 1 follows in the second valid clock cycle. The endianness of



**Figure 6.13.:** Block diagram of the data path link to connect a detector user logic [57].

the header words is big endian.

The main purpose of the RDH for the O2 is to attribute each packet to a particular timestamp and to a particular subsystem. The level of detail when filling the RDH fields is up to the detector, but basic information is mandatory in order for the O2 system to work properly. The relevant RDH fields are discussed below:

- **FEE ID**

This field provides 16 bit defined by the detector to identify the detector FEE which generated the data. Filling this field is not mandatory for the detector.

- **SYSTEM ID**

Unique number to identify the detector.

- **OFFSET NEW PACKET**



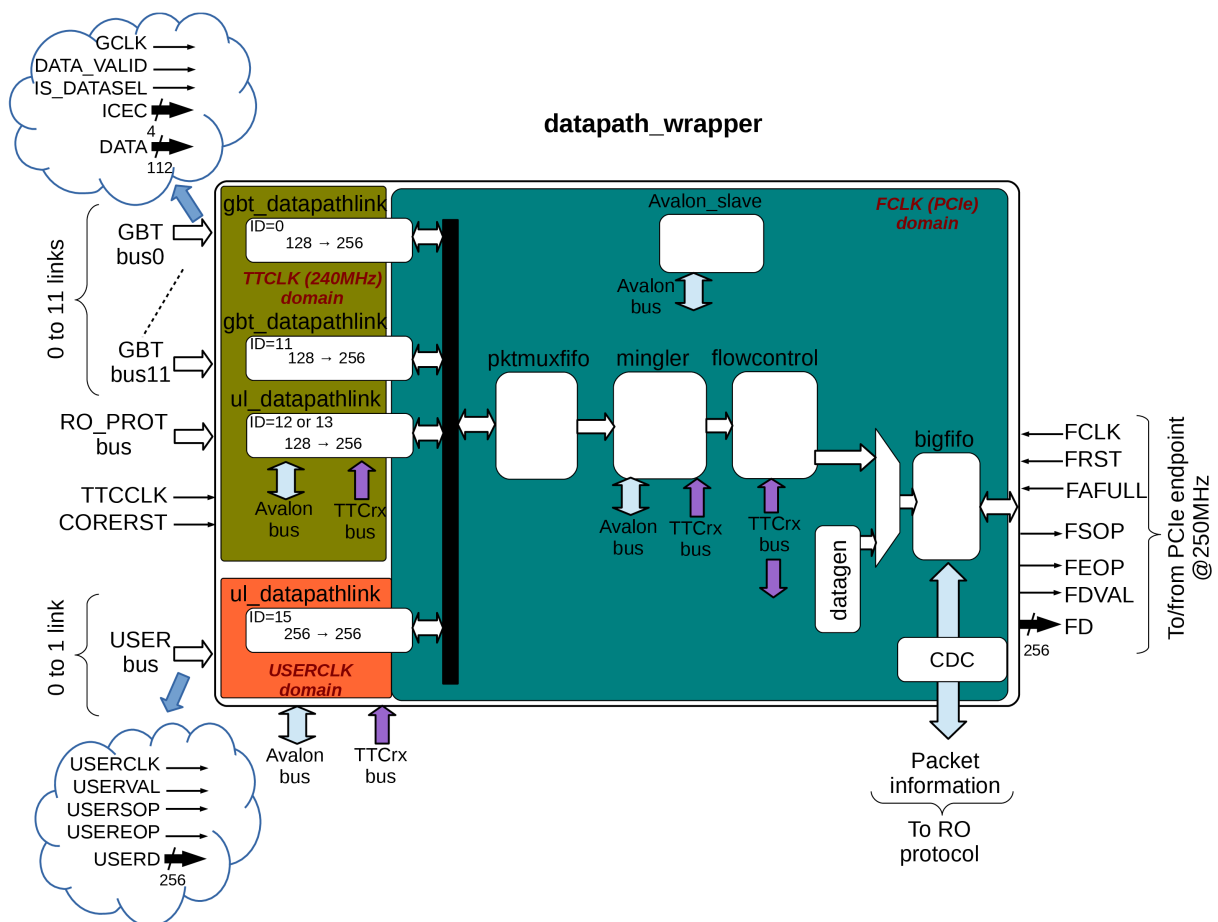
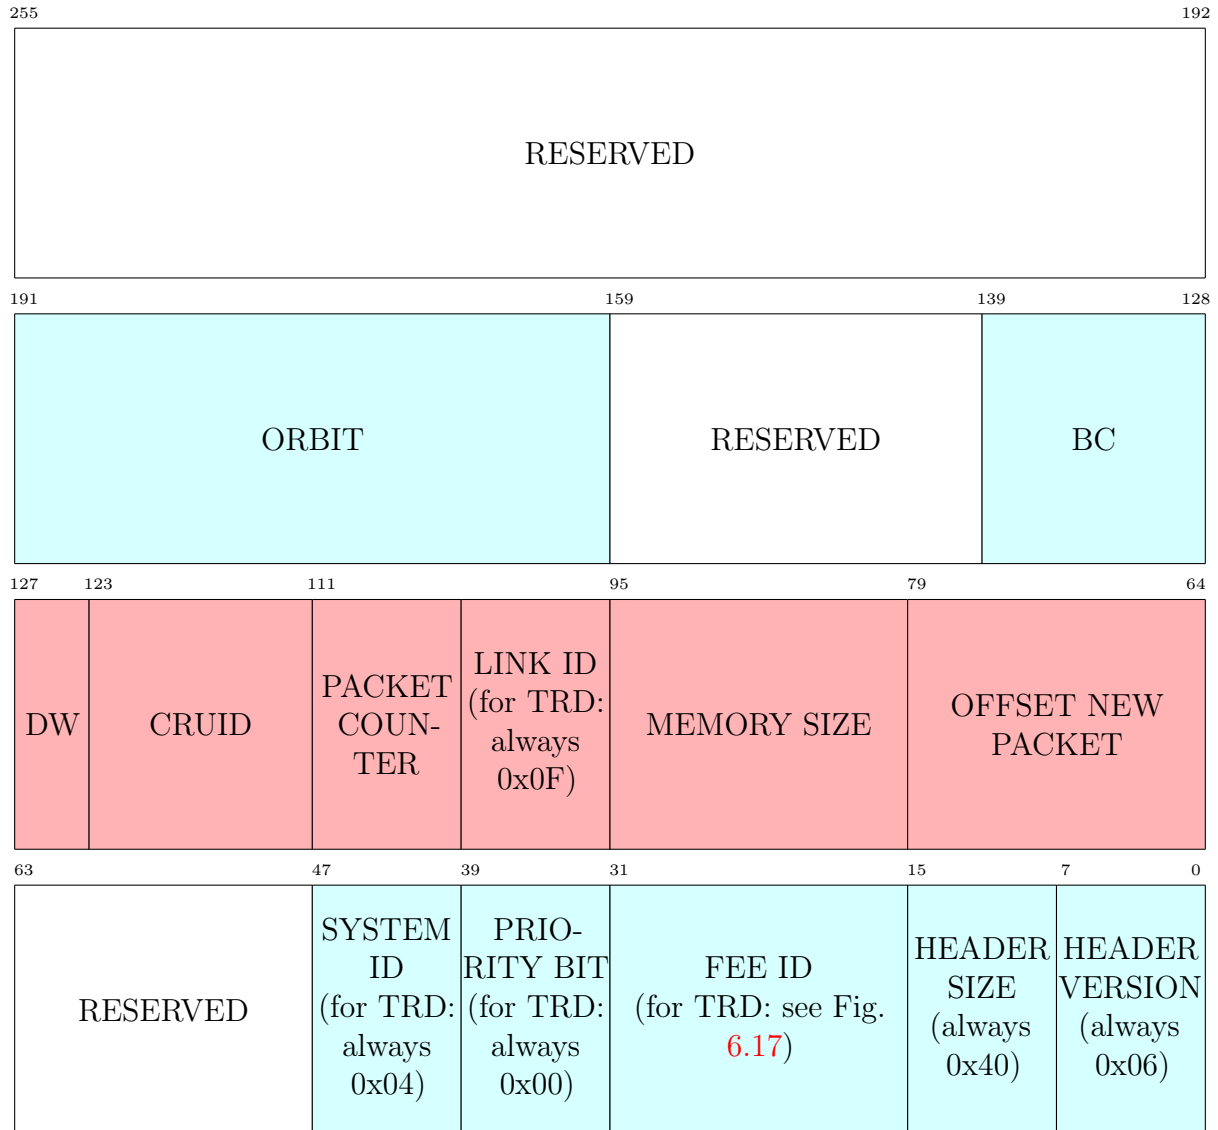
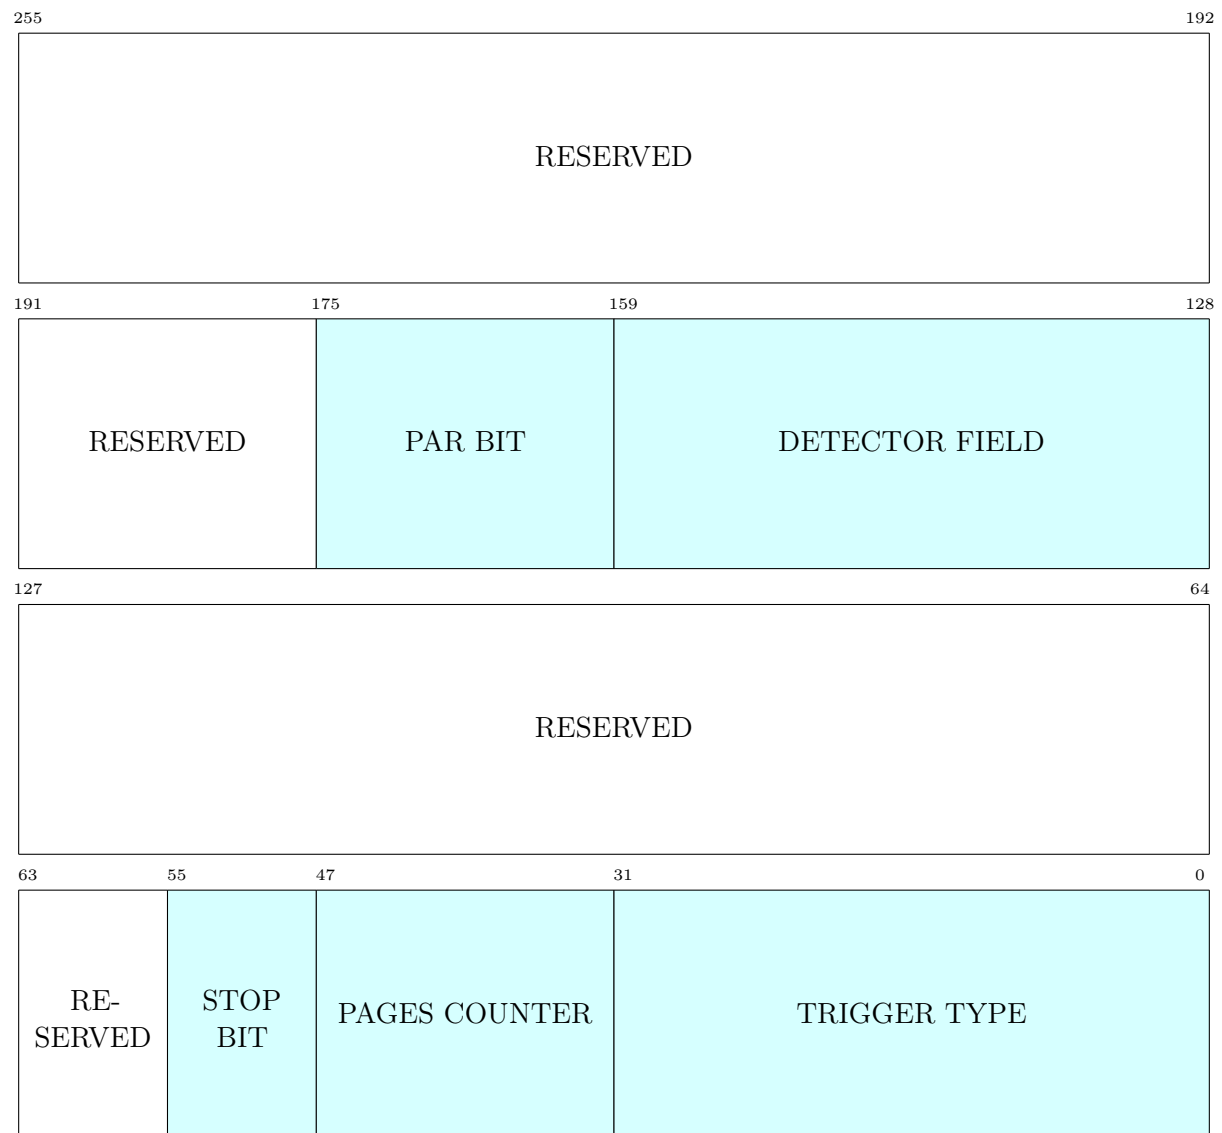


Figure 6.14.: Block diagram of the common logic datapath wrapper [57].



**Figure 6.15.:** Composition of the RDH Header word 0. Fields in red colour are only allocated by the user logic and later overwritten by the common logic with values generated within the common logic. All fields in blue colour are allocated and filled with values generated by the user logic. They are guaranteed to not be manipulated anymore by the common logic. Adapted from [58]



**Figure 6.16.:** Composition of the RDH Header word 1. Fields in red colour are only allocated by the user logic and later overwritten by the common logic with values generated within the common logic. All fields in blue colour are allocated and filled with values generated by the user logic. They are guaranteed to not be manipulated anymore by the common logic. Adapted from [58]

Common logic filled field informing the number of bytes to be counted in memory in order to find the next packet. This field simplifies reading the raw data by software once it is in the memory.

- **LINK ID**

Common logic filled field informing the link on which the data has been received. Note that this is mainly tailored to GBT detectors which use the GBT link wrapper provided by the common logic. In case the data are received via the direct path from the user logic to the datapath wrapper, this field indicates 0xF as a special link ID.

- **BC**

Bunch crossing ID to which the data in the packet belongs to. Filling this field is not mandatory for the detector.

- **ORBIT**

HB ID of the orbit to which the data belongs. This field is mandatory to be filled correctly and data from different orbits is not allowed to be mixed within the same packet. This is necessary for the O2 to work correctly because the experiment wide throughput throttling and data discarding in case of problems works by discarding data on a per-orbit base. O2 thus has to be able to identify which orbit certain data packets belong to.

Furthermore, it is not allowed to leave out orbits which do not have payload. Every non-rejected orbit announced by the CTP must generate at least 1 packet with the corresponding HB ID. This packet may be empty.

- **TRIGGER TYPE**

This field should carry the trigger type as received via the TTC-PON for the data in the packet. Filling this field correctly is not mandatory for the detector.

- **PAGES COUNTER**

This field must be filled with a counter which is incremented with every packet belonging to the same orbit. The counter is reset before the first data word of the

subsequent orbit. This field is mandatory to be filled correctly, as the O2 software would not be able to parse the data correctly otherwise.

- **STOP BIT**

This field has to be filled with 0x1 in case the packet is the last packet for the orbit and with 0x0 otherwise. This field is mandatory to be filled correctly, as the O2 software would not be able to parse the data correctly otherwise.

- **DETECTOR FIELD**

This field can be used for any purpose defined by the detector. Using this field is not mandatory for the detector.

The mandatory correct usage of the ORBIT field implies that at the orbit boundary, it is necessary to close a packet even it is not full and start the next packet.

For simplicity of the implementation, it is possible to generate an empty packet with the stop bit set with every subsequent heartbeat trigger (heartbeat frame closing packet). Additionally, one generates a new RDH without stop bit in case the heartbeat trigger type does not have the reject flag set (heartbeat frame opening packet).

## Datapath wrapper interface

The packets must be pushed to 2 independent datapath wrappers, serving the 2 PCIe endpoints. The datapath wrapper interface consists of the following signals:

- **CLOCK**

Must be driven with a clock from the user logic. The data transfer has to be synchronous to this clock. The maximum allowed clock frequency is 250 MHz.

- **DATA**

Must be driven with the data words formatted as 256 bit words and synchronous to the clock. Each clock cycle 1 of these words can be transferred.

- **VALID**

This valid flag has to be driven in order to mark the clock cycles where valid data has to be registered.

- **Start of Packet (SOP)**

SOP has to be driven high together with the first valid data word of a packet to indicate the packet start and 0 otherwise.

- **End of Packet (EOP)**

EOP has to be driven high together with the last valid data word of a packet to indicate the packet end and 0 otherwise.

- **Almost Full (AFULL)**

With this flag, the datapath wrapper indicates back to the user logic when its buffer is almost full. This indicates that the datapath is experiencing back pressure and cannot hand over the data to the subsequent logic. This is typically caused by congestion on the PCIe. The user logic shall stop pushing new data in this case, as it might be lost.

## 6.7. Prerequisites for TRD CRU Custom Firmware Development

This section presents thoughts and considerations in the context of this thesis, which are prerequisites for the implementation of the custom TRD CRU firmware. Most importantly, it also describes the idea on how to embed the custom TRD data into the global data format. This idea is the base for the actual implementation of the custom firmware, which is developed in the context of this thesis. The detailed Theory of Operation of the custom firmware is presented in Appendix [A](#).

### 6.7.1. TRD specific implementation of the O2 data format


The data sent by the TRD FEE in physics mode does not have any preceding header. The CRU must connect the data to the HC, which generated it, and to the BC by generating appropriate headers. While the official RDHv6 definition provides enough possibilities to store these meta data specifically required for the TRD inside the detector usable RDH, there are some important drawbacks to this approach.

Storing meta data with high level of detail in the RDH, namely including BC as a timestamp and the FEE HC link as subsystem information, would force the CRU to close the packet and open a new one, whenever any of the meta data changes. Depending on the actual event sizes, this could result in a huge number of packets with very small payload each. In this case, the fixed 512 bit header size per packet may become important with respect to the average payload size. The unavoidable overhead might have a severe impact on the overall performance.

In physics mode, the absolute maximum event size per link and per physics trigger is 1 kiB. This is only true for absolute black events. In the majority of cases, the HC has no tracks and the transmitted data consists of endmarkers only. A large number of comparatively much larger RDHs would need to be generated to process them.

Considering these drawbacks, the data format for the TRD has been based on a compromise: for internal TRD data processing, the RDH is used in order to identify the HB ID and the FEE ID only. For all other meta data which are important for the TRD, a custom packetisation scheme is used which is better tailored to the TRD specific requirements.

This scheme is compatible with the mandatory requirement of using the orbit field of the RDH correctly while at the same time systematically aiming for the generation of RDH preceded packets with the maximum size of 8 kiB. Only when the orbit changes and no sufficient data volume is available to fill the last RDH preceded packet of the



**Figure 6.17.:** Details on FEE ID field.

previous orbit entirely, the generation of a single smaller RDH preceded packet is necessary.

The custom TRD packetisation is performed on the payload of the RDH packets, i.e. it does not interfere in any way with the RDH packetisation for O2.

The TRD packetisation scheme is also independent from the O2 packetisation, i.e. a TRD packet does not need to necessarily fit into an O2 packet. It can spread across multiple O2 packets, without repeating the TRD header.

### 6.7.2. Note on RDH fields for TRD

The TRD uses the **FEE ID** field in order to distinguish different CRUs. The usage of the available 16 bits is shown in Fig. 6.17.

The endpoint can take the values 0x0 and 0x1. The CRU fills this field automatically as the endpoints are directly tied to the firmware. The field A/C side determines whether the CRU handles the A-side of a super module (0x0) or the C-side (0x1). The field SM is the super module index (values 0 - 17). For the latter 2 fields, there is no way for the CRU to determine these values automatically, as they solely depend on to which fibres the CRU is connected to. Therefore, they have to be configured via dedicated slow control registers prior to starting the data taking. There is a principle risk that this configuration is missed before taking data and that wrong fields are written into the final data. Therefore, the CRU starts up with the default values 0xF for the A/C side field and 0xFF for the SM field. These default values cannot occur in reality and therefore allow raw data reading software to recognise the fault afterwards.



The **PRIORITY BIT** field is not used.

The **SYSTEM ID** is always filled with 0x04, a unique identifier for the TRD defined by O2.

The **ORBIT** field is used to identify the orbit/HB ID, as it is mandatory.

The **BC** field of the RDH is not meaningful for the TRD as a custom scheme is used to identify single event timestamps. The CRU fills the RDH BC field with the BC value sent by the CTP together with the heartbeat trigger which started the heartbeat frame. The RDH BC is maintained for all subsequent packets belonging to the same heartbeat frame.

The **TRIGGER TYPE** field is filled according to the rules outlined above, as well as the pages counter and stop bit fields.

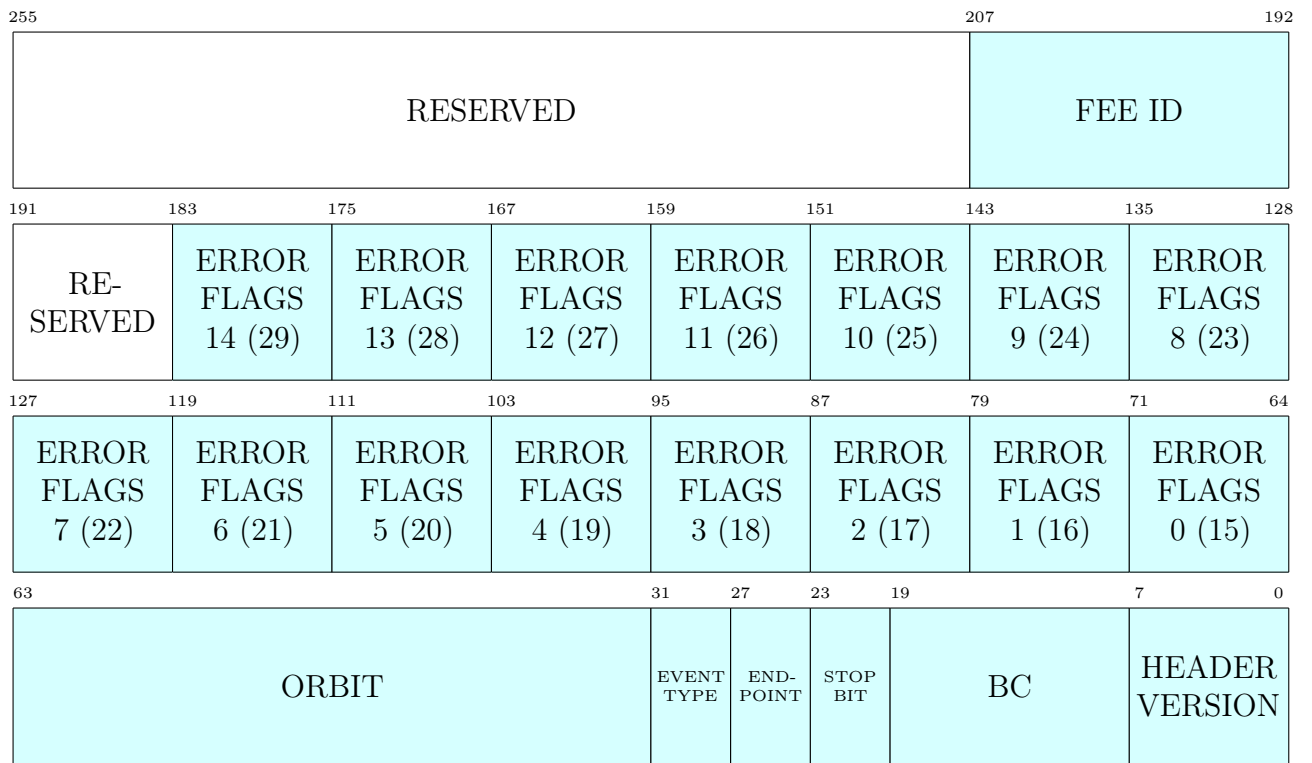
The TRD does not use the RDH **DETECTOR FIELD**.

The **PAR BIT** is not used either, as the TRD uses custom error recovery procedures which do not require any actions by O2.

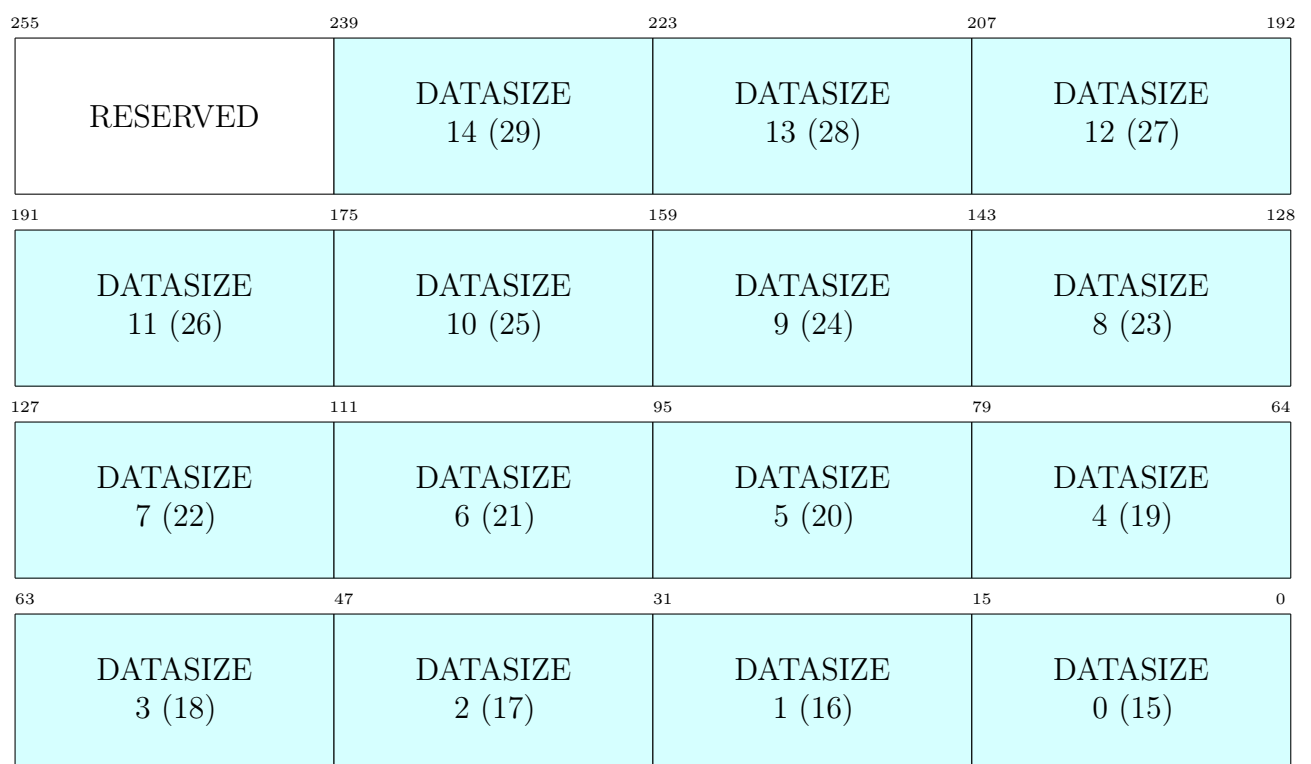
All unused fields are systematically filled with zeros by the CRU.

### 6.7.3. O2 payload packetisation scheme for TRD

The idea of the TRD custom packetisation is generating a single large packet for every physics or calibration event. The packetisation scheme does not impose a limit on the maximum packet size. This TRD packet gets equipped with a header indicating all required meta information, like bunch crossing time stamp as well as data sizes. The CRU



**Figure 6.18.:** Composition of the TRD Header word 0. Indices outside brackets are valid for endpoint 0, while indices in brackets are valid for endpoint 1.



**Figure 6.19.:** Composition of the TRD Header word 1. Indices outside brackets are valid for endpoint 0, while indices in brackets are valid for endpoint 1.

simplifies reading the data for the TRD offline by ensuring that all data from the same link is written contiguously into memory and that no data fragmentation occurs. After the TRD header, the non fragmented link data are embedded into the packet, starting from link 0 (15) up to link 14 (29) for endpoint 0 (1). No trailer is used at the end of the event data.

The composition of the TRD header is shown in the Figures 6.18 and 6.19. The TRD header has a total length of 512 bit. They are sent to O2 during 2 clock cycles, 256 bit each. The TRD header word 0 is sent in the first valid clock cycle of each packet, the TRD header word 1 follows in the second valid clock cycle. The endianness of the TRD header words is big endian.

The most important TRD header fields are discussed below:

- **HEADER VERSION**

Header version for distinguishing future modifications. The header version presented in this work is the first version 0x00.

- **BC**

The actual BC ID of the event in the payload.

- **STOP BIT**

Unused. It has been reserved in order to fragment data by spreading the same event across multiple TRD packets. This never happens as the CRU has more internal buffer resources available than the FEE can send.

- **ENDPOINT**

This field indicates 0x0 for endpoint 0 and 0x1 for endpoint 1.

- **EVENT TYPE**

This field indicates whether the event is of physics or calibration type.

- **ORBIT**



**Figure 6.20.:** Details on error flags field.

Copy of the RDH orbit field. This field exists for simpler offline reading of the orbit data after stripping off the RDHs.

- **ERROR FLAGS 0 (15) to 14 (29)**

This field indicates whether errors have occurred during the data acquisition on the corresponding link. 0x00 indicates that the data are trustworthy and no errors, neither during the current event acquisition, nor in the past have been detected on the corresponding link. A number different from 0x00 indicates problems. The composition is shown in Fig. 6.20.

A set timeout error bit indicates that after a trigger, no endmarker has been received within the configured timeout. A set data-outside-trigger-window bit indicates that data from the corresponding link has reached the CRU without a valid trigger context. Once an error bit is set, it remains also set for subsequent events. The error bits can only be reset manually or via DCS (also compare Table A.5).

- **FEE ID**

This field identifies the CRU and the endpoint. It is a copy of the RDH FEE ID and allows identifying the connected SM, the connected SM side, and the endpoint.

- **DATASIZE 0 (15) to 14 (29)**

These fields indicate how much data the following payload contains for the corresponding link. This field is crucial for offline processing of the data because it easily allows separating the data from different links. Note that there are no headers or spacers between the data from different links.

## 6.8. Validation of the TRD CRU Custom Firmware and Tests

While developing the custom CRU firmware, the Münster laboratory setup has been used for many verification tasks and for debugging.

In the long shutdown period after Run 2 and after the installation of all CRUs for the CRU based detectors of ALICE had been complete, an experiment wide commissioning period has been conducted. It has been composed of a standalone commissioning period, during which every detector has performed individual tests, and a global commissioning period, during which the ALICE experiment has been operated by central shifters, using the new Run 3 operational guidelines and software, to test the interplay of all ALICE systems together.

During standalone commissioning, many new features from other subsystems have become ready to be tested. These features could therefore not be tested in the laboratory earlier and their compatibility with the TRD CRU had to be verified. The new features are listed below.

- The calibration trigger. The trigger system started to support the sending of the required calibration trigger sequences with the possibility to alternately read Physics data and calibration data from the detector. Previously, only a verification of the TRD CRU with physics triggers had been possible.
- The new Run 3 TRD FEE configuration. This configuration started to support the reception of calibration triggers and the generation of the new Run 3 tracklet data format (compare Fig. 6.1). Previously, the TRD CRU firmware could only be verified based on Run 2 data.
- The busy protection. Previously, the TRD CRU has always generated its busy contribution, but this contribution could not be evaluated by the CTP earlier. Therefore, the readout performance could not be pushed to its limits and a conservative trigger

rate had to be used for testing. During commissioning, the busy has become fully operational.


- The raw data reading. For the TRD and its new Run 3 data format, dedicated raw data reading software has been developed by the TRD team. This software analyses the data packets generated by the CRU at the deepest possible level. In this way, it has become possible to verify the entire readout chain, from the FEE to the EPN down to the bit level.

While testing these new features, the TRD CRU firmware has proven excellent stability and no or only minor adjustments have been necessary to the CRU firmware after each test.

The availability of the full Run 3 trigger scheme and the entire busy signal processing chain have allowed for measuring the absolute maximum TRD readout rate as it is limited by the finite detector drift velocity as well as processing and propagation latencies. The result has been better than the estimates in the initial Technical Design Report [49] and yielded a maximum rate of 120 kHz.

Independent of testing the new features, also long term tests have been conducted. The TRD CRU has enabled the very first recording of particle tracks with the TRD in the new Run 3 setup. For that, the particle tracks have been generated by injecting instable Krypton isotopes into the TRD gas system, while the full detector readout chain has been operating. This test has proven stable running for multiple days. The data generated by the Krypton data taking have been verified for consistency and meanwhile also serve calibration purposes of the individual detector chambers.

In addition, global runs have been conducted by central shifters, which also have proven the stability and the excellent readout performance. Most importantly, all technical features of the CRU and the TRD readout chain have been verified in global runs during the very first Run 3 data recording with pp collisions during a pilot beam test period. The correctness and the usability of the recorded data has already been verified.




**Figure 6.21.:** Sampled ADC values during drift time in one of the first pp collisions in Run 3. The dark blue line shows the mean ADC value. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59].

Regular Run 3 data taking operations with beam have started in mid 2022. The collected data with the TRD and the CRUs can be successfully reconstructed.


Figure 6.21 shows the ADC raw values sampled during the drift time in one of the first pp collisions in Run 3. The dark blue line shows the mean ADC value. The amplification peak is clearly visible.

Figure 6.22 shows an exemplary event display of a Time Frame (TF) in pp collisions at  $\sqrt{s} = 13.6$  TeV extracted from run 520145 recorded on July 5th, 2022. The interaction rate is approximately 13 kHz [59]. The tracklets in the TRD are illustrated with blue markers. Figure 6.23 shows the 3-dimensional view of the the same TF, illustrating the reconstructed tracks within the TPC.






**Figure 6.22.:** Event display of an exemplary full Time Frame (TF) in pp collisions at  $\sqrt{s} = 13.6$  TeV from run 520145 recorded on July 5th, 2022. The tracklets in the TRD are illustrated with blue markers. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59].



**Figure 6.23.:** 3-dimensional view of the event display of an exemplary full Time Frame (TF) in pp collisions at  $\sqrt{s} = 13.6$  TeV from run 520145 recorded on July 5th, 2022. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59].



**Figure 6.24.:** 3-dimensional view of the event display of the first Pb–Pb collision at  $\sqrt{s_{\text{NN}}} = 5.36$  TeV in Run 3 recorded by ALICE on November 18th, 2022. The tracklets in the TRD are illustrated with cyan markers. Figure from Ref. [60].

In November 2022, ALICE has recorded Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 5.36$  TeV for the first time in Run 3. Figure 6.24 shows the 3-dimensional view of the event display of the first recorded collision. The tracklets in the TRD are illustrated with cyan markers.

# Chapter 7

## Summary

The main objective of this thesis was the development of the Common Readout Unit (CRU) firmware for the readout of the Transition Radiation Detector (TRD) in Run 3. The TRD is a detector of ALICE (A Large Ion Collider Experiment) at the Large Hadron Collider (LHC). An additional project within the scope of this thesis was the development of a TRD trigger on light nuclei for use in Run 2 as well as its performance analysis. Furthermore, a TRD gas system maintenance service task within the ALICE collaboration was part of this thesis.

The developed TRD trigger on light nuclei has been used to increase the yield of rare events containing light nuclei already in Run 2, prior to major technical upgrades for Run 3. This thesis covered the development of the trigger, starting with detailed feasibility studies, which were conducted by analysing data in Lead – Lead (Pb–Pb) collisions at  $\sqrt{s_{\text{NN}}} = 2.76 \text{ TeV}$  and Proton – Proton (pp) collisions at  $\sqrt{s} = 8 \text{ TeV}$  available from earlier data taking without trigger. Based on the obtained results, promising trigger performance estimates were presented and a set of best performing trigger parameters was determined.

The trigger was based on cuts on the particle energy loss within the TRD, which is expected to be exceptionally large for light nuclei. Technically, the cuts were applied by the Global Tracking Unit (GTU), which delivered the trigger decision to the Central

Trigger Processor (CTP) of ALICE as a Level 1 (L1) trigger contribution less than  $6.5\,\mu\text{s}$  after the collision. In the scope of this thesis, the necessary hardware modifications for the implementation of the trigger in the GTU were carried out and presented.

The developed nuclei trigger was operated successfully in Proton – Lead (pPb) collisions at  $\sqrt{s_{\text{NN}}} = 8\,\text{TeV}$  and Proton – Proton (pp) collisions at  $\sqrt{s} = 13\,\text{TeV}$  in Run 2. For both collision systems, the data were analysed with focus on the trigger performance. The results from the aforementioned feasibility studies were confirmed. Furthermore, significant enhancements of nuclei and antinuclei were observed.

For the first time, two  $\overline{{}_2^4\text{He}^{2+}}$  candidates were identified in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8\,\text{TeV}$ . These unprecedented candidates were observed in events collected by the developed nuclei trigger.

The trigger enhancement factors with respect to Minimum Bias (MB) data were measured in pp collisions at  $\sqrt{s} = 13\,\text{TeV}$ . Considerably large enhancement factors have been obtained for deuterons, tritons,  ${}_2^3\text{He}^{2+}$  and their antiparticles as well as for  ${}_2^4\text{He}^{2+}$ . The enhancement factors increase with the particle mass and their values are approximately 175 for  ${}_2^3\text{He}^{2+}$  and 328 for  $\overline{{}_2^3\text{He}^{2+}}$ .

The TRD gas system service task corresponded to an equivalent of 6 months full time work and consisted of supervising and operating the TRD gas system, which runs permanently on a 24 hours / 7 days in a week base and requires immediate actions in case of issues. In addition, the gas was manually purified from nitrogen contamination using a cryogenic procedure in the context of this service task.

After Run 2, major technical upgrades have been applied to ALICE. The aim of the upgrades was a significant increase in readout performance and data statistics. In Run 3, the CRUs have entirely replaced the GTU, which had handled the TRD data during LHC Run 1 and Run 2. All data collected by the TRD in Run 3 are received by the CRUs,

packetised and forwarded to reconstruction and permanent storage systems. The CRUs also perform a key role during data taking as they control and throttle the readout of the TRD Front End Electronics (FEE).

This thesis presented the evolution of the TRD readout upgrade, starting with a summary of the system as it had been operated in Run 2. The modifications for Run 3 were presented with focus on the development of the CRU firmware, which has been performed in the scope of this thesis. Details of the TRD specific components of the CRU firmware were outlined as a reference for possible future developments. In addition, this thesis contained detailed considerations on the optical drive capabilities of the TRD FEE as well as on the performance of the optical fibre installations for the new readout infrastructure. Both theoretical estimations and measurements were presented. In the scope of this thesis, a laboratory setup equivalent to the upgraded TRD readout systems in ALICE has been developed and installed at the Institut für Kernphysik at the Universität Münster. The laboratory setup and the obtained results were presented in this thesis.

The new readout system was thoroughly tested in the laboratory and within ALICE. The optical drive capabilities allow for a stable performance without limitations. The developed CRU firmware has been found to comply in all aspects with its prior specifications. The entire TRD readout chain has been proven to perform stably and its performance exceeds previous estimations.

Finally, after the beginning of Run 3 data taking in mid 2022, particle tracks were successfully reconstructed using the data collected by the CRUs.



# Chapter 8

## Zusammenfassung

Das Hauptziel dieser Doktorarbeit war die Entwicklung der Common Readout Unit (CRU)-Firmware für die Auslese des Transition Radiation Detector (TRD) in Run 3. Der TRD ist ein Detektor des ALICE-Experiments (A Large Ion Collider Experiment) am Large Hadron Collider (LHC). Ein zusätzliches Projekt im Rahmen dieser Doktorarbeit war die Entwicklung eines TRD-Triggers für leichte Atomkerne zur Verwendung in Run 2 sowie die Untersuchung der Leistungsfähigkeit des Triggers. Außerdem waren Wartungsaufgaben im Bereich des TRD Gas-Systems für die ALICE-Kollaboration Teil dieser Arbeit.

Die Entwicklung eines TRD-Triggers für leichte Atomkerne hatte den Zweck, die Anzahl der zu Analysenzwecken zur Verfügung stehenden Ereignisse bereits während des Run 2, noch vor einer umfangreichen technischen Aufrüstung, zu erhöhen. Die Doktorarbeit hat die Entwicklung des Triggers gezeigt, beginnend mit detaillierten Untersuchungen zur Vorhersage der Leistungsfähigkeit des Triggers. Für die Untersuchungen wurden Datensätze von früheren Aufnahmen ohne Trigger von Blei – Blei (Pb–Pb)-Kollisionen bei  $\sqrt{s_{NN}} = 2.76$  TeV und Proton – Proton (pp)-Kollisionen bei  $\sqrt{s} = 8$  TeV verwendet. Auf der Grundlage der so erhaltenen Ergebnisse wurde außerdem ein optimierter Satz von Triggerparametern bestimmt, für den die Leistungsfähigkeit des Triggers maximal war.

Der Trigger basierte auf Grenzwertschnitten für den Teilchenenergieverlust innerhalb des TRDs, der für leichte Atomkerne außergewöhnlich groß ist. Technisch sind

die Grenzwertschnitte von der Global Tracking Unit (GTU) durchgeführt worden, die die Triggerentscheidung als Level 1 (L1)-Triggerbeitrag innerhalb unter  $6.5\,\mu\text{s}$  nach der Kollision an den Central Trigger Processor (CTP) des ALICE-Experiments übertragen hat. Im Rahmen dieser Doktorarbeit sind die für die Umsetzung des Triggers notwendigen Veränderungen an der GTU-Hardware vorgestellt und durchgeführt worden.

Der Trigger für leichte Atomkerne ist für die Aufnahme von Proton – Blei (pPb)-Kollisionen bei  $\sqrt{s_{\text{NN}}} = 8\,\text{TeV}$  und für die Aufnahme von Proton – Proton (pp)-Kollisionen bei  $\sqrt{s} = 13\,\text{TeV}$  während Run 2 erfolgreich eingesetzt worden. Die oben genannten Vorhersagen über die Leistungsfähigkeit des Triggers sind anhand des getriggerten Datensatz mit pPb-Kollisionen bei  $\sqrt{s_{\text{NN}}} = 8\,\text{TeV}$  bestätigt worden. Außerdem sind signifikante Verstärkungsfaktor für  ${}^3_2\text{He}^{2+}$ ,  ${}^4_2\text{He}^{2+}$  und die zugehörigen Antiteilchen gefunden worden. Im getriggerten Datensatz mit pPb-Kollisionen bei  $\sqrt{s_{\text{NN}}} = 8\,\text{TeV}$  sind zwei  ${}^4_2\text{He}^{2+}$ -Kandidaten identifiziert worden. Es war dies das erste Mal, dass  ${}^4_2\text{He}^{2+}$ -Kandidaten in pPb-Kollisionen beobachtet wurden.

In dieser Doktorarbeit sind die Verstärkungsfaktoren des Triggers im Verhältnis zu Minimum Bias (MB)-Datensätzen für pp-Kollisionen bei  $\sqrt{s} = 13\,\text{TeV}$  gemessen worden. Für Deuteronen, Tritonen,  ${}^3_2\text{He}^{2+}$  und die zugehörigen Antiteilchen sowie für  ${}^4_2\text{He}^{2+}$  haben sich bemerkenswerte Verstärkungsfaktoren ergeben. Die Verstärkungsfaktoren steigen mit der Teilchenmasse, wobei sich für  ${}^3_2\text{He}^{2+}$  ein Wert von etwa 175 und für  ${}^4_2\text{He}^{2+}$  ein Wert von etwa 328 ergeben hat.

Die Wartungsaufgaben im Bereich des TRD Gas-Systems entsprachen dem Umfang einer 6-monatigen Vollzeitstelle und bestanden in der Überwachung und dem Betrieb des TRD Gas-Systems. Das TRD Gas-System ist 24 Stunden am Tag und 7 Tage jede Woche betrieben worden. In Störfällen war sofortiges Handeln erforderlich. Außerdem ist das Gas im Rahmen der Wartungsaufgaben manuell mit Hilfe eines Kryo-Prozesses von Stickstoff-Kontaminationen befreit worden.



Nach dem Ende des Run 2 ist eine umfangreiche technische Aufrüstung des ALICE-Experiments durchgeführt worden. Das Ziel der Aufrüstung war eine signifikante Steigerung der Leistungsfähigkeit der Datenauslese sowie eine Steigerung der Datenstatistik. Für Run 3 ist die GTU, die die Daten des TRDs in den Runs 1 und 2 verarbeitet hat, vollständig durch CRUs ersetzt worden. Die gesamten während des Runs 3 aufzunehmenden Daten des TRDs werden nun von den CRUs ausgelesen, in Datenpaketen formatiert und an die Systeme zur Rekonstruktion und Speicherung weitergeleitet. Den CRUs kommt eine Schlüsselrolle während der Datenaufnahme zu, da sie die Auslese der TRD Front End Electronics (FEE) steuern und den Datenfluss regeln.

Die Entwicklung der technischen Aufrüstung der TRD-Auslese ist in dieser Doktorarbeit dargestellt worden, beginnend mit einer Zusammenfassung des während Run 2 verwendeten Systems. Die Veränderungen für Run 3 wurden vorgestellt, wobei der Schwerpunkt auf der Entwicklung der CRU-Firmware gelegen hat, die im Rahmen dieser Dissertation durchgeführt worden ist. Die Details der TRD-spezifischen Teile der CRU-Firmware wurden ausführlich behandelt, sodass die vorliegende Doktorarbeit als Grundlage für mögliche Weiterentwicklungen in der Zukunft herangezogen werden kann. Darüber hinaus wurde auf die Details der Leistungsfähigkeit der optischen Laser-Sender der TRD FEE und auf die Leistungsfähigkeit der Glasfaserinstallation für die neue Ausleseinfrastruktur im Detail eingegangen. Sowohl theoretische Vorhersagen als auch Messergebnisse wurden gezeigt. Teil dieser Arbeit war auch die Entwicklung eines Laboraufbaus am Institut für Kernphysik an der Universität Münster, der in seiner Struktur den TRD-Auslesesystemen in ALICE entspricht und zum Testen der aufgerüsteten Auslesesysteme verwendet worden ist.

Das neue Datenauslesesystem ist im Rahmen dieser Doktorarbeit sowohl im Labor als auch innerhalb des ALICE-Experiments intensiv getestet worden. Die Leistungsfähigkeit der Laser-Sender ermöglicht einen stabilen Betrieb ohne Einschränkungen. Es hat sich erwiesen, dass die entwickelte CRU-Firmware alle zuvor spezifizierten Anforderungen erfüllt. Die gesamte TRD-Ausleseketten hat sich im stabilen Betrieb bewährt und die Leistungsfähigkeit der Auslese übertrifft vorherige Abschätzungen.

Abschließend ist festgestellt worden, dass seit dem Beginn der Datenaufzeichnungen in Run 3 ab Mitte des Jahres 2022 die Teilchenspuren basierend auf den über die CRUs ausgelesenen Daten erfolgreich rekonstruiert werden können.

# Appendix




# Appendix A

## Theory of Operation of the TRD CRU Custom Firmware

This appendix describes the implementation and the design of the custom CRU firmware for TRD, which has been carried out in the scope of this thesis.

The custom CRU firmware for the TRD is based on the common CRU firmware where the GBT wrapper and user logic are removed and replaced by custom logic. The GBT mux and data generator modules are also present in the TRD customised firmware for the sake of a simpler maintenance of the code base, but there is no practical use case for these modules for the TRD. Instead, the TRD user logic features a custom data generator taking into account the specific requirements of the TRD.

The TRD custom components of the logic are described in the following sections. The code is available at CERN GitLab <https://gitlab.cern.ch/alice-trd-upgrade/trd-cru-fw.git>. For better readability of the description along with the source code, the component and signal names in this text are kept identical to those used in the source code.

Figure A.1.: `trd_custom_module_wrapper` diagram.

## A.1. `trd_custom_module_wrapper`

For better code maintainability, all customised components of the TRD firmware have been wrapped into a large entity called `trd_custom_module_wrapper`. This entity replaces both the GBT wrapper and the user logic in the common firmware.

The `trd_custom_module_wrapper` features a PLL to generate the TRD user clock for the actual TRD specific user logic. It derives a clock with configurable frequency from a 240 MHz reference clock which the common logic provides. The TRD user logic has been functionally verified at a clock frequency of 160 MHz. Also, FPGA timing closure is typically achieved without any issues at this clock speed.

## A.2. `ori_wrapper`

The `ori_wrapppper` is responsible for receiving serial data streams from the optical input links into the CRU and converting them into parallel data streams, similar to the data format present in front of the ORI board on the ROC. It replaces the 2 GBT wrappers of




Figure A.2.: ori\_wrapper diagram.

the common firmware.

The `ori_wrappper` is primarily a wrapper entity for the `ori_link` entities (compare Section A.3). In addition, it employs a PLL in order to generate a common 125 MHz reference clock for the `ori_link` entities.

### A.3. ori\_link

Every `ori_link` entity handles a single unidirectional serial input link to the CRU. It receives the serial data stream at a rate of 2.5 Gbit/s and converts into a 16 bit wide parallel data stream at a frequency of 125 MHz with a valid bit. The core functionality is being carried out by an IP core from the FPGA manufacturer which has been configured to handle the following processing aspects on the incoming data stream:

- **Clock data recovery**

In this step, the clock of the transmitter is recovered using a PLL and a 125 MHz reference clock by synchronising the recovered clock on the edges of the incoming data.

- **Deserialisation**

This step converts the serial data stream of 1bit@2.5 Gbit/s into a parallel data stream of 20bit@125 MHz.

- **Decoding - Conversion of 10bits to 8bits**

The arriving data stream is 8b10b encoded. Each 8 bits of the original data has been replaced on the transmitter side by a 10 bit word according to a conversion table. The encoding guarantees that the serial data stream fulfils certain conditions during the transmission, especially it guarantees that the serial data stream contains enough edges, allowing the clock data recovery to work under all circumstances.

This step decodes the data by replacing each 10bits with the corresponding original 8bits. In order to decode the 20 bit wide stream, 2 decoders work in parallel to generate 16 bit words.

- **Decoding - Removal of control characters and generation of a valid bit**

The arriving serial data contain control characters which are important in the deserialisation step to deserialise the data with the original alignment of the data. The control characters are 10 bit characters which do not translate to any 8 bit word and can be uniquely recognised.


In this step the control characters are removed by generating a valid bit, which is asserted for actual data words and cleared for control characters. The net data rate of the 16bit@125MHz data stream with valid bit corresponds to the data rate present on the FEE in front of the ORI board.

The processing steps above essentially reverse the steps carried out on the transmitting ORI board by a transceiver chip TLK2051 from TI. For further details the datasheet in Ref. [55] can be consulted.


Figure A.3 shows a test of the IP core decoding capabilities in the simulation software ModelSim®. A test counter with control characters is encoded by an external encoder. The encoded data stream is fed into the IP core, which recovers the original data stream.

Apart from the `ori_link` core functionality, additional custom logic has been implemented in `ori_link` for debugging and error detection purposes. For example, it is possible to replace the parallel data stream from the IP core by fixed patterns. Further-





**Figure A.3.:** ModelSim® simulation to test the 8B10B decoder within the IP core. A counter with control characters is encoded by an external encoder and decoded by the IP core, which recovers the original data stream.




**Figure A.4.:** ori\_link diagram.

more, errors in the data stream are systematically counted and reported through slow control tools. In this way, unstable links or faulty configurations on the FEE side can be recognised early and narrowed down easily.

An overview of the input and output signals of the `ori_link` is given in Fig. A.4.

## A.4. trd\_user\_logic

The overview of the TRD user logic is given in Fig. A.6. Red arrows depict the main readout data flow, green arrows depict the trigger data flow, blue arrows depict control signals and the orange arrow is BUSY. The user logic receives the parallel data from the ORI wrapper from all optical input links. In addition, it receives the TTC-PON signal from the CTP via an interface to the common logic. The user logic processes the data and sends formatted output data to the datapath wrappers of the common logic. In addition,

Figure A.5.: `trd_user_logic` diagram.

it generates the busy signal for the CTP.

An overview of the input and output signals of the `trd_user_logic` is given in Fig. A.5.

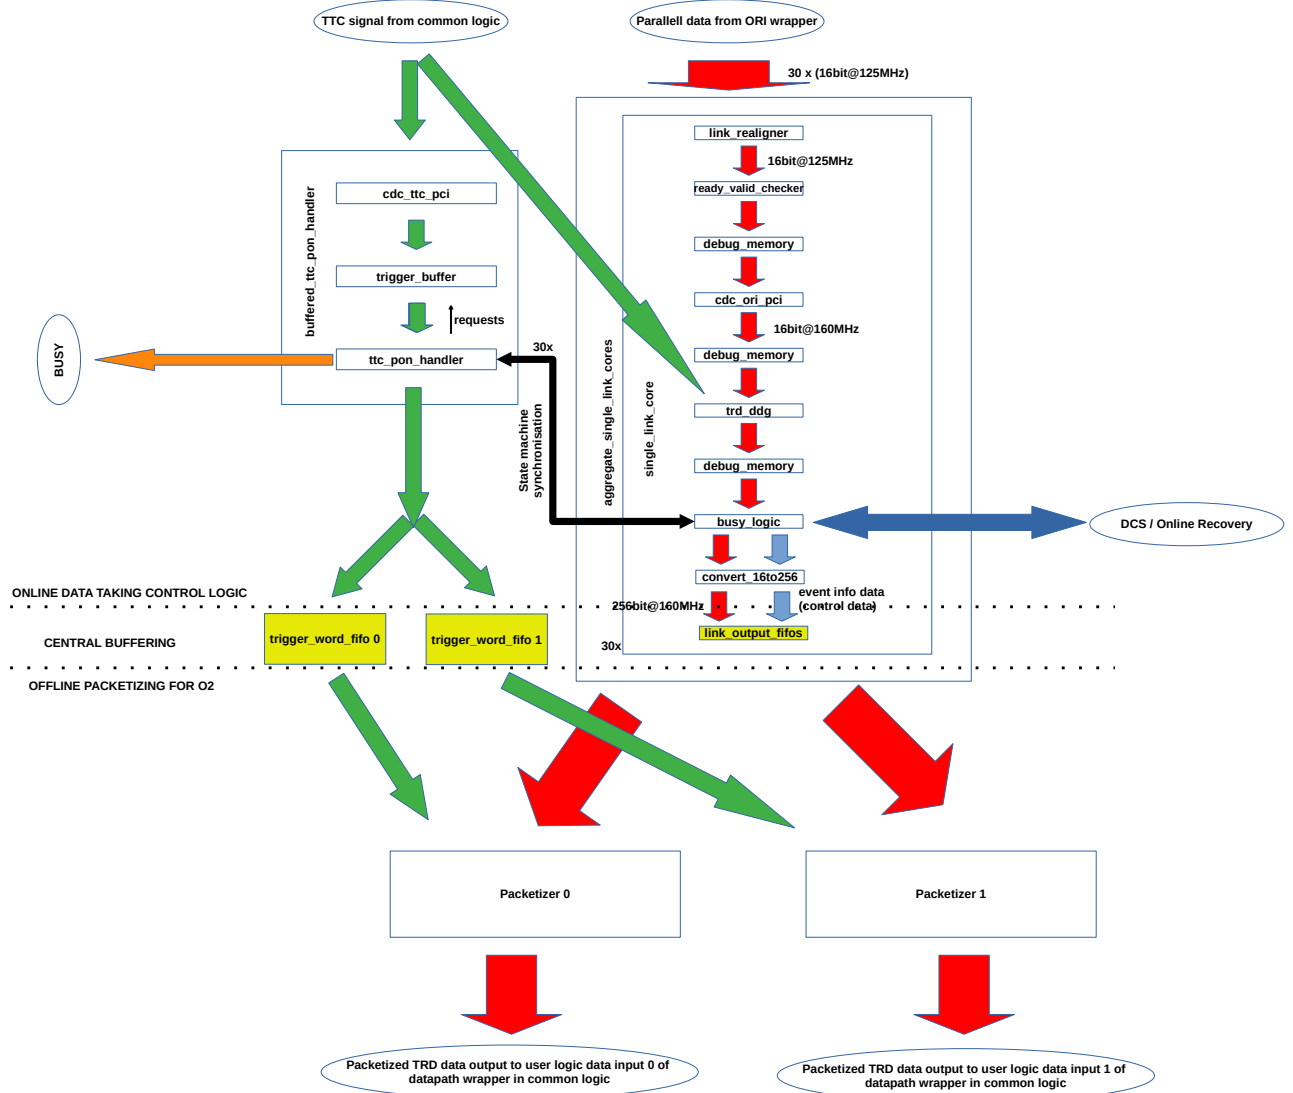
## A.5. `buffered_ttc_pon_handler`

This entity encapsulates the entire functionality around the trigger reception, readout steering and busy generation.

An overview of the input and output signals of the `buffered_ttc_pon_handler` is given in Fig. A.7.

## A.6. `cdc_ttc_pci`

The downstream trigger signals from the TTC-PON interface are delivered to the user logic via the common logic. Before being passed to the user logic, they are transferred into



**Figure A.6.:** Overview of the TRD CRU user logic. Red arrows depict the main readout data flow, green arrows depict the trigger data flow, blue arrows depict control signals and the orange arrow is BUSY.




Figure A.7.: buffered\_ttc\_pon\_handler diagram.




Figure A.8.: cdc\_ttc\_pci diagram.

a 240 MHz clock domain and a typical duty cycle of 1 valid word in 6 clock cycles. For compatibility with the TRD custom user logic clock of 160 MHz, the `cdc_ttc_pci` entity provides clock domain crossing of the TTC-PON signal into the TRD custom user logic clock domain. This is accomplished using a dual clock First In First Out buffer (FIFO) of low depth which is read whenever it is not empty.

An overview of the input and output signals of the `cdc_ttc_pci` is given in Fig. A.8.

## A.7. trigger\_buffer

The TRD custom user logic has to handle physics triggers, calibration triggers and heart-beat triggers. For what concerns physics triggers and calibration triggers, the readout scheme built around busy protection, guarantees that no new physics triggers or calibration triggers can be sent by the CTP while the data from a previous physics or calibration




Figure A.9.: trigger\_buffer diagram.

trigger is still being received by the CRU. These trigger types thus do not require buffering inside the CRU.

However, the same is not true for heartbeat triggers which are issued every orbit, regardless whether the TRD and its CRUs are busy recording data. Considering the long TRD readout times especially for calibration triggers, which can extend across multiple orbits, the heartbeat triggers which occur in the mean time have to be held in a buffer. This ensures that the CRU can systematically open and close every heartbeat frame for the O2 system, even if no TRD data is available in these heartbeat frames as the TRD has been busy processing the data from an earlier heartbeat frame which contained actual physics or calibration triggers.


The trigger\_buffer identifies the triggers which are relevant for the TRD in the TTC-PON data stream and converts them into a specially optimised data format, in the following named the TRD\_trig\_word format. The details of the TRD\_trig\_word format are given in Table A.1. For details on the format of the original TTC-PON data stream, refer to Ref. [61]. The selected and optimised trigger messages are written into an internal FIFO. This FIFO also performs the transition from a trigger push architecture to a trigger pull architecture, i.e. the downstream logic can request the next trigger whenever it is ready to handle it and does not need to be ready for new trigger words at all times.

**Table A.1.:** TRD\_trig\_word format. It is guaranteed that for every trigger (heartbeat, physics, calibration) separate TRD\_trig\_words are generated, i.e. only one of the 3 last bits can be high in the same TRD\_trig\_word. This is ensured by trigger\_buffer, even in case that CTP sends a heartbeat accept together with a physics or calibration trigger within the same clock cycle.

Bit	Number of Bits	Description
255..96	160	unused
95..64	32	trigger type field from TTC-PON message
63..47	17	unused
46..15	32	HB ID
14..3	12	BC ID
2	1	1: calibration trigger, 0: not calibration trigger
1	1	1: physics trigger, 0: not physics trigger (might still be calibration, this is different from bit 4 of TTC-PON RX message)
0	1	1: heartbeat trigger (might be both accepted or rejected), 0: not heartbeat trigger

The conversion of the trigger data format from the native TTC-PON data format to the TRD\_trig\_word format yields 2 advantages:

- Only relevant trigger words have to be stored, and therein only the relevant data have to be stored
- It significantly simplifies the downstream logic. The CTP definition of the native TTC-PON [61] allows sending multiple triggers within the same trigger word and clock cycle. In particular, a heartbeat trigger might coincide with a physics and/or calibration trigger. Even in case of such a coincidence, the CRU must ensure that in the final output data stream, the new heartbeat frame is opened first and only after that the physics data gets inserted. The TRD\_trig\_word format guarantees, that every TRD\_trig\_word in the FIFO corresponds to a particular action which the CRU must take and the order of the TRD\_trig\_words in the FIFO corresponds to the correct order in which these actions must be taken.

Figure A.10.: `ttc_pon_handler` diagram.

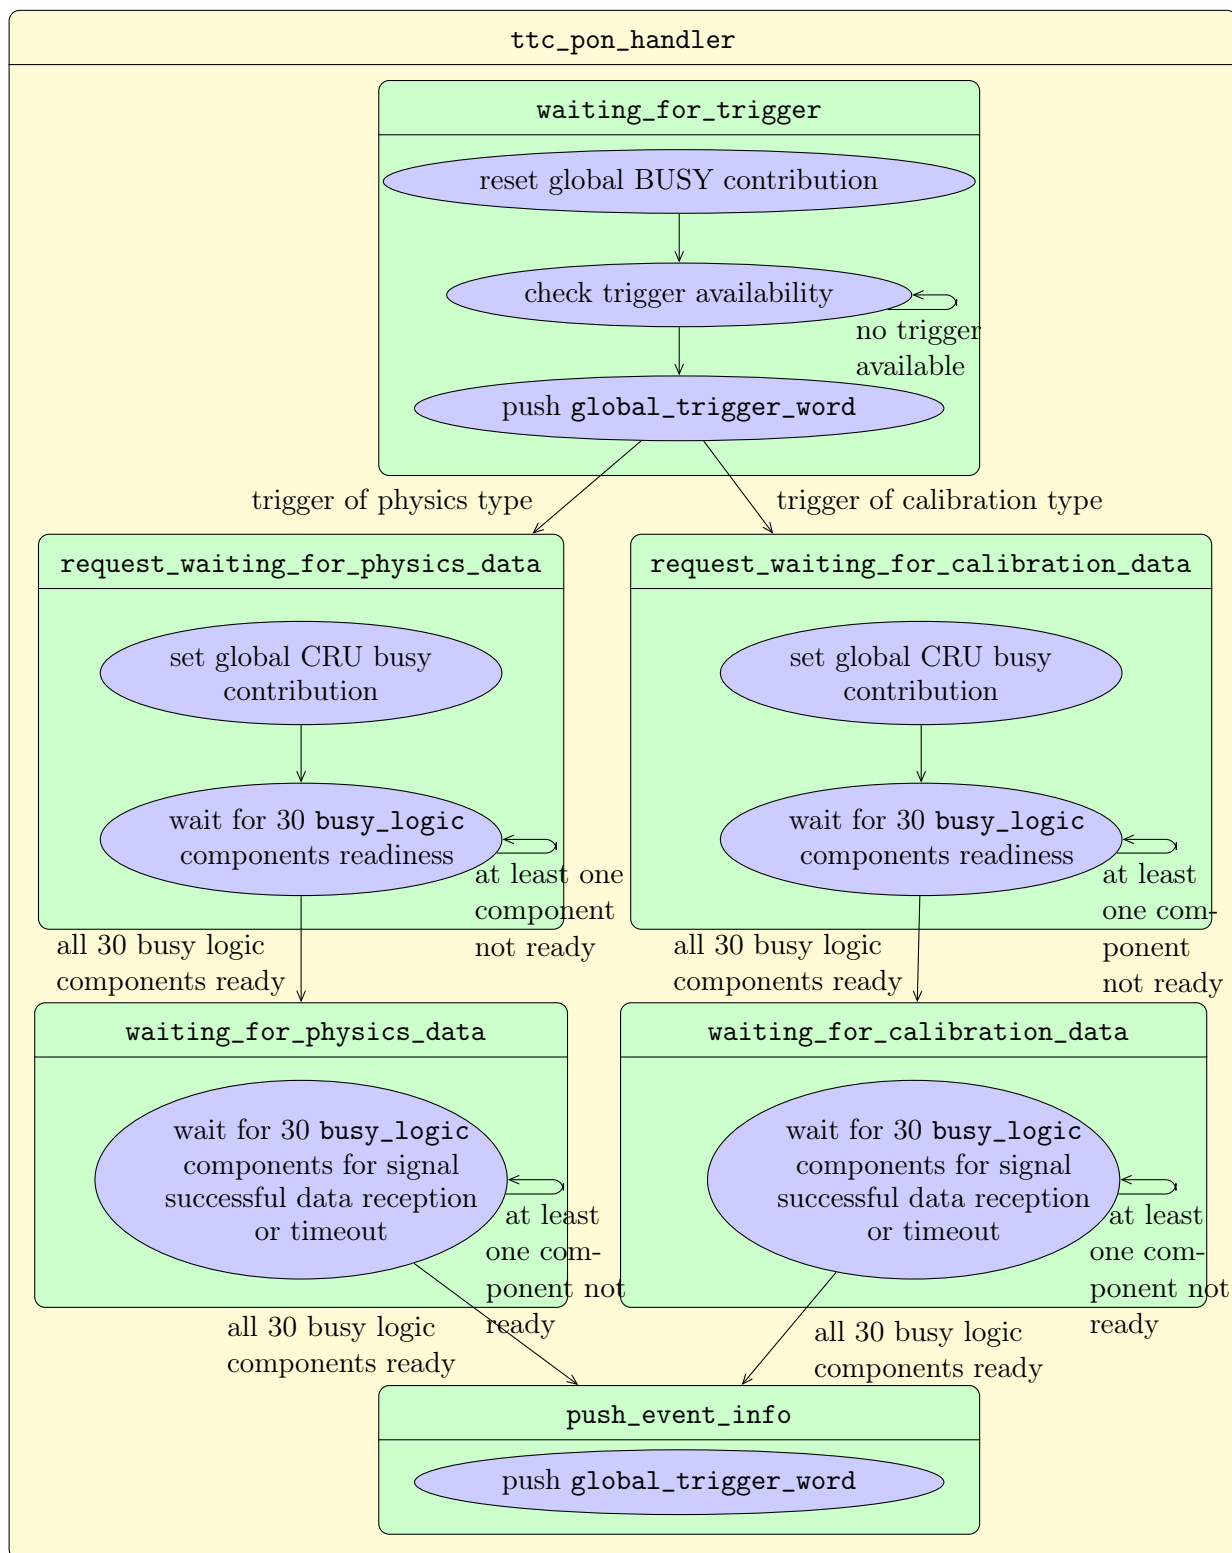
An overview of the input and output signals of the `trigger_buffer` is given in Fig. A.9.

## A.8. `ttc_pon_handler`

The `ttc_pon_handler` is the global core component which controls the TRD readout on the CRU global level depending on the received trigger. The triggers are sequentially picked up from the trigger buffer. A new trigger is fetched only once all data from the previous trigger have reached the CRU on all input links. The `ttc_pon_handler` generates the global busy contribution of the CRU for the CTP.

Whenever certain operational conditions occur (e.g. certain trigger received, event data acquisition complete), the `ttc_pon_handler` protocols the condition by generating a customised `global_trigger_word` to keep track of it. This tracking is important for the downstream logic.

The `ttc_pon_handler` communicates and synchronises itself directly with the `busy_logics` of each of the 30 input links. This close communication allows the `ttc_pon_handler` to establish a global state of the CRU, summarising the local states of all 30 input links as a global state. This global state is also representative for the readout state of the fraction of the TRD, which is handled by the concerned CRU.

Figure A.11.: `ttc_pon_handler` state machine diagram.



Internally, the `ttc_pon_handler` is based on a Finite State Machine (FSM) which involves several states and which are looped through sequentially, depending on the readout status. Figure A.11 illustrates the execution of a single cycle of the internal `ttc_pon_handler` state machine. In addition, the states and their transitions are described below:

- `waiting_for_trigger`

This is the idle state of the `ttc_pon_handler`. This state is maintained as long as no trigger is received and no data supposed to be taken. In this state, the `waiting_for_trigger` signals to not be busy and waits until a relevant trigger arrives. For this purpose, the `trigger_buffer` is checked in a fast loop for newly arrived triggers. In case that a new trigger is available, a `global_trigger_word` is generated and pushed out by the `ttc_pon_handler` in order to signalise this fact to downstream logic. In case that the trigger is of heartbeat type, no further action is taken and the `ttc_pon_handler` remains in the state `waiting_for_trigger`. In case the trigger is of physics or calibration type, the CRU must be prepared for data arrival and the FSM changes its state to `request_waiting_for_physics_data` or `request_waiting_for_calibration_data`.

- `request_waiting_for_physics_data` or `request_waiting_for_calibration_data`

In these states the CRU busy contribution is raised. This ensures that no further physics or calibration triggers can be received until the `ttc_pon_handler` clears busy. The `ttc_pon_handler` requests all `busy_logic` instances of all 30 input links to prepare themselves for data reception and listens back whether they are ready for data reception. Once all `busy_logic` instances report their readiness, the `ttc_pon_handler` FSM state changes to `waiting_for_physics_data` or `waiting_for_calibration_data`.

- `waiting_for_physics_data` or `waiting_for_calibration_data`

In these states the CRU busy contribution remains active. The CRU waits for the FEE on all input links for data to arrive. The `ttc_pon_handler` gets the status of data reception reported by the 30 `busy_logic` instances connected to the input links and remains in this global state until all expected data of the current event

**Table A.2.:** Composition of the `global_trigger_word`.

Bit	Number of Bits	Description
127..96	32	reserved
95..64	32	trigger type field from TTC-PON message
63	1	<code>global_event_info_is_preceding_word</code> (0: trigger word after TRD data acquisition, 1: trigger word before TRD data acquisition)
62..31	31	HB ID
30..19	12	BC ID
18..17	2	<code>global_event_info_event_type</code> (bit 1: calibration, bit 0: physics, bits 00: heartbeat (might be both accepted or rejected))
16..0	17	reserved

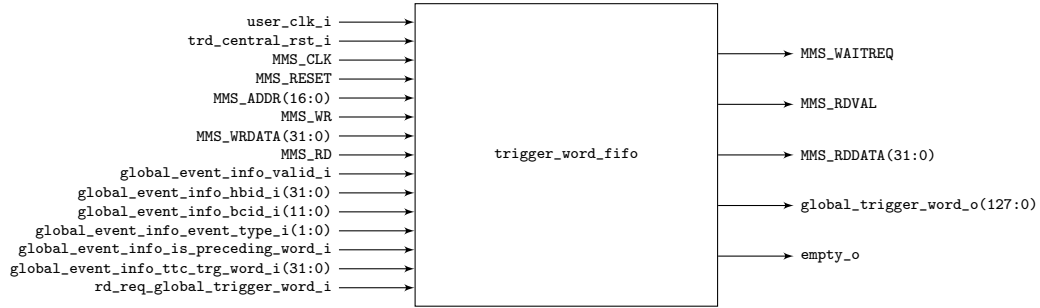
have been collected from all 30 input links. Once all `busy_logic` instances report that they do not expect more data, the `ttc_pon_handler` FSM enters the state `push_event_info`.

- `push_event_info`

This state serves the generation of a concluding `global_trigger_word`, indicating to downstream logic that the collection of the ongoing event has been completed and all the data are buffered in CRU buffers. After pushing out the `global_trigger_word`, the `ttc_pon_handler` FSM returns to the state `waiting_for_trigger`, which clears the busy contribution and checks for the arrival of a subsequent trigger to be processed.

The composition of a `global_trigger_word` is shown in Table A.2.

An overview of the input and output signals of the `ttc_pon_handler` is given in Fig. A.10.

Figure A.12.: `trigger_word_fifo` diagram.

## A.9. `trigger_word_fifo`

The `trigger_word_fifos` are single clock FIFOs which hold the `global_trigger_words` pushed by the `ttc_pon_handler`. They change from push architecture to pull architecture and allow the downstream logic to pick up words at any convenient time.

There are 2 `trigger_word_fifo` instances in the logic which receive the `global_trigger_word` from the unique `ttc_pon_handler` in parallel to thus store the same data, duplicated. This allows the subsequent 2 `packetiser` instances to work independently and serve both PCIe endpoints separately.

An overview of the input and output signals of the `trigger_word_fifos` is given in Fig. A.12.

## A.10. `aggregate_single_link_cores`





Figure A.13.: aggregate\_single\_link\_cores diagram.

Figure A.14.: `single_link_core` diagram.

The `aggregate_single_link_cores` entity serves as a container housing 30 `single_link_core` instances and its purpose is improving code readability. It allows the implementation of the link handling core entity to focus on a single link.

An overview of the input and output signals of the `aggregate_single_link_cores` is given in Fig. A.13.

## A.11. `single_link_core`

The `single_link_core` entity implements the handling of link data from the `ori_wrapper` for a single optical input link in the context of the global readout scheme.

It implements a central Avalon slave which is shared among its other sub components for slow control access.

An overview of the input and output signals of the `single_link_core` is given in Fig. A.14.




Figure A.15.: link\_realigner diagram.

## A.12. link\_realigner

The data decoded by the `ori_wrapper` for a single link is delivered as a 16bit@125 MHz data stream which has been generated by 2 parallel 8B10B decoders. Because of technical limitations in the configurability of the 8B10B decoder IP cores used in the `ori_wrapper`, the data stream unfortunately does not correspond exactly to the 16bit@125 MHz data stream prior to encoding by the Texas Instruments TLK2051 IC on the ORI board. Effectively, the data appears shifted by 8 bits, i.e. the 16 bit word transferred within the same clock cycle prior to the Texas Instruments TLK2051 encoded IC appears distributed across 2 clock cycles after decoding.

The `link_realigner` fixes this shortcoming by reshuffling the data bytes accordingly and restoring the original byte formatting prior to encoding.

An overview of the input and output signals of the `link_realigner` is given in Fig. A.15.

## A.13. ready\_valid\_checker

The `ori_wrapper` signals with a valid flag whether a valid data word or a control character is being transferred. In addition to the valid flag, there is a ready flag which signals whether the link synchronisation is established. In order to spare the downstream from dealing




Figure A.16.: ready\_valid\_checker diagram.




Figure A.17.: debug\_memory diagram.

with both flags, the 2 flags are combined using an AND operation resulting in a new valid flag incorporating the ready signal. The downstream logic only checks for the new valid flag.

An overview of the input and output signals of the `ready_valid_checker` is given in Fig. A.16.

## A.14. debug\_memory

A `debug_memory` has been implemented and inserted behind the `ready_valid_checker` and in other strategic places. The `debug_memory` internally implements a circular buffer consisting of 256 16bit registers. The `debug_memory` scrutinises the data which passes it and writes every valid data word into the circular buffer. The circular buffer can be inspected using a slow control interface.




Figure A.18.: `cdc_ori_pci` diagram.

This functionality is not strictly required during data taking operation but has proven to be useful in case of problems with the readout tree. In these cases, the readout can be halted and the latest data words in the readout tree can be inspected manually even if subsequent parts of the logic (e.g. packetising, disk dumping) are inoperative.

The `debug_memory` does not have its own connection to the Avalon bus for slow control due to missing address space. However, it uses some registers of the shared Avalon slave implemented in `single_link_core`. The registers concerning all implemented instances of `debug_memory` are documented in Table A.3.

An overview of the input and output signals of the `debug_memory` is given in Fig. A.17.

## A.15. `cdc_ori_pci`

This entity performs clock domain crossing and transfers the data from the recovered 125 MHz ORI board clock into the TRD custom user logic clock. In this way, it is guaranteed that the data from all links share the same clock domain, simplifying the design of the downstream logic.



**Table A.3.:** Slow control registers provided by the shared avalon slave in `single_link_core` for `debug_memory`.

Register address	r(eadable) / w(ritable)	Description
0x28	w	Word address pointer for <code>debug_memory</code> after <code>ready_valid_checker</code> . The number written to this register must be between 0 and 254 and selects which word of the stored words in the ring buffer appears at the output. 0 corresponds to the latest stored word, while 254 corresponds to the earliest.
0x2c	w	Word address pointer for <code>debug_memory</code> after <code>cdc_ori_pci</code> . The number written to this register must be between 0 and 254 and selects which word of the stored words in the ring buffer appears at the output. 0 corresponds to the latest stored word, while 254 corresponds to the earliest.
0x30	w	Word address pointer for <code>debug_memory</code> after <code>trd_ddg</code> . The number written to this register must be between 0 and 254 and selects which word of the stored words in the ring buffer appears at the output. 0 corresponds to the latest stored word, while 254 corresponds to the earliest.
0x1c	r	Selected word data output for <code>debug_memory</code> after <code>ready_valid_checker</code> .
0x20	r	Selected word data output for <code>debug_memory</code> after <code>cdc_ori_pci</code> .
0x24	r	Selected word data output for <code>debug_memory</code> after <code>trd_ddg</code> .




Figure A.19.: `trd_ddg` diagram.

Internally, `cdc_ori_pci` is based on a dual clock FIFO, which is read whenever it is not empty.

An overview of the input and output signals of the `cdc_ori_pci` is given in Fig. A.18.

## A.16. `trd_ddg`

The entity `trd_ddg` implements a TRD specific data generator which aims imitating the data as it is generated by the TRD FEE. The `trd_ddg` features an internal multiplexer which can be operated by slow control and determines whether the subsequent logic is alimented with real data or with generated data.

In order to imitate the TRD FEE, the `trd_ddg` is connected directly to the TTC-PON signal and it starts the data generation automatically when it receives a physics trigger or a calibration trigger. After such a trigger, a slow control configurable amount of data is sent. Slow control configurable end markers are inserted automatically at the end of the transmission. There are separate configuration parameters for the physics and calibration cases.




Figure A.20.: busy\_logic diagram.

Currently, only trivial data can be generated, which get assembled by repeating constant 16 bit template pattern. This pattern is configurable through slow control.

The `trd_ddg` is useful to debug, test and stress the readout chain with different event sizes and trigger rates. In this way, the system can also be operated without the requirement of having an operative TRD FEE connected.

The `trd_ddg` does not have its own connection to the Avalon bus for slow control due to missing address space. However, it uses configuration registers of the shared Avalon slave implemented in `single_link_core`. The registers concerning the `trd_ddg` configuration are documented in Table A.4.

An overview of the input and output signals of the `trd_ddg` is given in Fig. A.19.

## A.17. busy\_logic

The `busy_logic` has a similar internal structure as the `ttc_pon_handler` and communicates with it directly. While the `ttc_pon_handler` handles the readout at the CRU global level, the `busy_logic` handles the readout at the level of a single link. The `busy_logic` scrutinises the data as it flows through it and tries to identify the end of transmission by the recognition of specific end markers. It collects information about the data passing

**Table A.4.:** Slow control registers provided by the shared Avalon slave in `single_link_core` for `trd_ddg`.

Register address	r(eadable) / w(ritable)	Description
0x04	w	Enable. Writing 0x1 to this registers enables the artificial data generator and blocks the real data coming in via the optical link. Writing 0x0 disables the data generator and passes the real data from the optical link through.
0x08	w	Number of TRD user logic clock cycles to wait after a trigger before transmitting the generated data. This register can be used to simulate the drift and processing time of the FEE.
0x0c	w	Number of generated 16 bit words after a calibration trigger. This number excludes the 4 calibration endmarkers, which are always send, even if this register is 0x0.
0x10	w	Number of generated 16 bit words after a physics trigger. This number excludes the 4 physics endmarkers, which are always send, even if this register is 0x0.
0x14	w	Configures the endmarker patterns that are sent after the generated data. The upper 16 bit configure the calibration endmarker pattern (default: 0x0000) and the lower 16 bit configure the physics endmarker pattern (default: 0x1000). The <code>trd_ddg</code> appends the applicable pattern in the end of the transmission 4 consecutive times.
0x18	w	Configures the constant data patterns used to generate the artificial data. The upper 16 bit configure the calibration data pattern (default: 0x1234) and the lower 16 bit configure the physics data pattern (default: 0xABCD).

through it and is able to detect basic errors in the data stream. In case that errors are identified, the `busy_logic` can signalise an LME for the corresponding link and block further recording of erroneous data.

Internally, the `busy_logic` contains 2 FSMs. The main FSM controls the readout status and is synchronised with the `ttc_pon_handler` FSM, while the other FSM stores the error status of the link. By default, the error FSM state is `no_error` which corresponds to a flawlessly operating link. Depending on the encountered error, it can be switched anytime to an error state, effectively excluding the link from data taking until a manual or automatic recovery through slow control interface takes place.

The main FSM involves several states which are looped through sequentially, depending on the readout status. The stages are described below:

- `waiting_for_trigger`

This is the idle state of the `busy_logic`. This state is maintained as long as no instruction from the `ttc_pon_handler` is received. The instructions are identified by directly monitoring the `global_triggerstate` signal which represents the status of the `ttc_pon_handler` FSM. In case that the `ttc_pon_handler` FSM enters one of the states `request_waiting_for_physics_data` or `request_waiting_for_calibration_data`, the `busy_logic` FSM switches into one of the states `waiting_for_physics_data` or `waiting_for_calibration_data` accordingly.

In the `waiting_for_trigger` state, the `busy_logic` scrutinises the incoming link data and it expects not to see any incoming valid data words. In case that a valid word is identified though, the link is considered faulty and the error FSM is switched to the state `data_outside_triggerwindow`.

The data from the link is not forwarded to the data output in the `waiting_for_trigger` state, effectively discarding any illegal words sent by the FEE without prior trigger.

- `waiting_for_physics_data` and `waiting_for_calibration_data`

The behaviour of these modes depends on the status of the error FSM.

In case that an error is present, no data from the link is forwarded to the output and the main state machine leaves these modes already after a single clock cycle. The main FSM transitions to the state `push_event_info` in these cases.

In case that the error FSM reports no error, the `busy_logic` remains in waiting mode and scrutinises the incoming link data which are at the same time forwarded to the data output. The logic checks for the occurrence of a configurable end marker pattern (typically 4 times 0x1000 for physics data and 4 times 0x0000 for calibration data) in the data stream. The incoming 16 bit data words are counted for the ongoing event and the end marker patterns are required to be aligned with even counter values. This requirement is related to the fact that the FEE sends the words in multiples of 32 bit (tracklet word width), while the CRU processes the data as 16 bit words. If a valid end marker pattern is identified, the data transmission for the ongoing event is complete and the main FSM transitions to the state `push_event_info`.

During the entire data collection process for a particular event, a timer is used to measure the time passed since the `busy_logic` has entered the `waiting_for_physics_data` or `waiting_for_calibration_data` state, respectively. In case that a configurable timeout is reached without that an endmarker has been identified in the data stream, the error FSM is set to the state `timeout_reached` and the main FSM transitions to the state `push_event_info`.

- `push_event_info`

This state is reached after the collection of a particular event is considered finalised, either because appropriate endmarkers have been detected or because the link is not being waited for due to an error state. The objective of this state is to push out all collected meta data about the data of the handled link for the current event to downstream logic. These information are important for the functioning of downstream logic entities and for the building of meta data headers in the final data. Explicit output ports of the `busy_logic` and a dedicated valid signal are implemented for this meta data transfer:

- `event_info_valid_o`

All of the fields below are valid in case the `event_info_valid_o` bit is set.

- `event_info_wordnum_o`

The total number of 16 bit words effectively forwarded for the ongoing event. In case there is no error at all, `event_info_wordnum_o` is equal to the number of 16 bit words sent by the link. In case a timeout error has occurred during the current event, `event_info_wordnum_o` contains only the number of words forwarded until the timeout triggered. If data forwarding is locked due an error during previous events which has not been recovered yet, `event_info_wordnum_o` is 0.

- `event_info_error_o`

The state of the error FSM.

After a single clock cycle, the `busy_logic` main FSM transitions to the state `request_waiting_for_trigger`.

- `request_waiting_for_trigger`

This state is reached after finishing all operations for the handled link and current trigger. The purpose is to request the `ttc_pon_handler`, which sees this state, to transition the global state to `waiting_for_trigger`. Once this has happened, the `busy_logic` main FSM transitions to the idle state `waiting_for_trigger`, too.

Note that the `ttc_pon_handler` only fulfils the request, once all 30 `busy_logic` instances request the transition to the global idle state. The time by which an individual `busy_logic` finishes processing all operations for the current trigger can largely differ, as it depends on the amount of data which the connected link sends and which might vary in size, depending on the local multiplicity in the corresponding detector region. In addition, the `busy_logic` instances which are in an error state, reach the state `request_waiting_for_trigger` already very few clock cycles after receiving the trigger as they do not wait for incoming data. In this way, links with previously erroneous link behaviour get effectively excluded from the busy generation and do not slow down data taking by waiting for the timeout after every trigger.

**Table A.5.:** Slow control registers provided by `busy_logic`.

Register address	r(eadable) / w(ritable)	Description
0x00	r	Always 0x459814ee. Serves as a <code>busy_logic</code> type identifier allowing the slow control software to verify that <code>busy_logic</code> is present and answering.
0x04	r	Optical link number (0..29) which the <code>busy_logic</code> handles.
0x08	r	Status of the internal error state machine. The bit 0 indicates whether an error is present which has been caused by the reception of any data without a valid trigger context. The bit 1 indicates whether an error is present which has been caused by a timeout condition. The register value 0x0 indicates that no error is present.
0x0c	r	reserved.
0x10	w	Error state machine reset register. Writing 0x1 to this register resets the error state machine without resetting any other parts of the logic and effectively re-includes the link into the data taking. It is safe to write to this register during an event is read out. In this case, the re-inclusion of the link into data taking is delayed until the next trigger.
0x14	w	Configures the endmarker patterns that have to be identified in the data. The upper 16 bit configure the calibration endmarker pattern (default: 0x0000) and the lower 16 bit configure the physics endmarker pattern (default: 0x1000). The <code>busy_logic</code> waits until the configured 16 bit pattern appeared 4 consecutive times.
0x18	w	Configures the physics timeout in user logic clock cycles (clock period: 1/160MHz). If in physics mode no endmarker is identified after the configured time, a timeout error occurs.
0x1c	w	Configures the calibration timeout in user logic clock cycles (clock period: 1/160MHz). If in calibration mode no endmarker is identified after the configured time, a timeout error occurs.
0x20	w	<code>busy_logic</code> reset register. Writing 0x1 to this register resets the entire <code>busy_logic</code> .
0x24	r	Status of the main state machine of the <code>busy_logic</code> . The register stores an unsigned numeric value which corresponds to a state machine status according to Table A.6.
0x28	r	Copy of register at 0x08 for software compatibility reasons.
0x2c	w	Total number of incoming 16 bit words since the last reset of the <code>busy_logic</code> .
0x30	r	Total number of incoming 16 bit words since the exclusion of the link due to erroneous behaviour.
0x34	r	Total number of accepted incoming 16 bit words since the last trigger. This counter is reset automatically with every new trigger and stopped in case an erroneous link behaviour is detected.



**Table A.6.:** Legend for state machine states in numerical representation.

Numeric state machine status	State machine status name
0	busy_forced_by_slowcontrol
1	waiting_for_trigger
2	waiting_for_physics_data
3	waiting_for_calibration_data
4	prepare_request_waiting_for_trigger
5	last_counts
6	push_event_info
7	request_waiting_for_trigger
8	request_waiting_for_physics_data
9	request_waiting_for_calibration_data
10	error_illegal_trigger_received

A full list of slow control registers provided by each `busy_logic` instance is given in Table A.5. Apart from configuration and debugging registers, every `busy_logic` provides 2 slow control registers which are vital for the TRD online error recovery strategy. These registers are accessed in the final system by the DCS via the Alice Low-level Frontend (ALF) and Front End Device (FRED) software and the DIM protocol. More details are given in Section A.17.1.

An overview of the input and output signals of the `busy_logic` is given in Fig. A.20.

### A.17.1. Communication of the `busy_logic` instances with DCS

Reading the slow control register 0x28 of the `busy_logic` can be used to determine whether the associated link is working as expected or not (compare Table A.5). Any possible malfunction is also called LME and is most often caused by SEUs or similar error

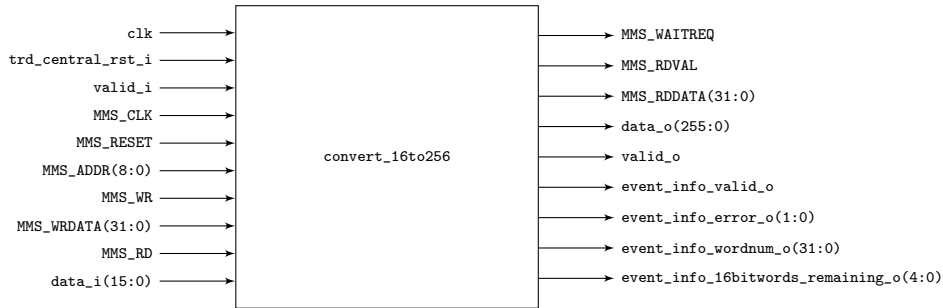


Figure A.21.: convert\_16to256 diagram.

states of the associated TRD FEE. The TRD DCS uses this functionality of the CRU to supervise the readout tree on the HC level. It reads all **busy\_logics** registers 0x28 of all TRD CRUs regularly and maps possible LMEs internally.

The DCS takes appropriate action and recovers the LMEs by power cycling and re-configuring the affected ROCs at a convenient time. Once the issues are solved, the DCS writes 0x1 to the registers 0x10 of the concerned **busy\_logics** (compare Table A.5) to reconsider the recovered ROCs in data taking.

This recovery procedure can be performed automatically through a DCS script, or manually, using the DCS GUI.

For convenience, the full addresses of the relevant registers 0x28 and 0x10 is given for all **busy\_logics** within a single CRU is given in Table A.7 and Table A.8, respectively.

## A.18. convert\_16to256


The data transfer up this stage has been based on transferring data words of 16 bits width per clock cycle as this is the natural width of the ORI link data. However, the O2 expects the detectors to send 256 bit per clock cycle. Therefore, data reformatting is

**Table A.7.:** Full addresses of the `busy_logic` error state registers 0x28 for all `busy_logic` present within a single CRU.

CRU ORI link number	Full register address address
0	0x00ca2b28
1	0x00ca3b28
2	0x00ca4b28
3	0x00ca5b28
4	0x00ca6b28
5	0x00ca7b28
6	0x00ca8b28
7	0x00ca9b28
8	0x00caab28
9	0x00cabb28
10	0x00cacb28
11	0x00cadb28
12	0x00caeb28
13	0x00cafb28
14	0x00cb0b28
15	0x00cb1b28
16	0x00cb2b28
17	0x00cb3b28
18	0x00cb4b28
19	0x00cb5b28
20	0x00cb6b28
21	0x00cb7b28
22	0x00cb8b28
23	0x00cb9b28
24	0x00cbab28
25	0x00cbbb28
26	0x00cbcb28
27	0x00cbdb28
28	0x00cbeb28
29	0x00cbfb28

**Table A.8.:** Full addresses of the `busy_logic` error state reset registers 0x10 for all `busy_logic` present within a single CRU.

CRU ORI link number	Full register address address
0	0x00ca2b10
1	0x00ca3b10
2	0x00ca4b10
3	0x00ca5b10
4	0x00ca6b10
5	0x00ca7b10
6	0x00ca8b10
7	0x00ca9b10
8	0x00caab10
9	0x00cabb10
10	0x00cacb10
11	0x00cadb10
12	0x00caeb10
13	0x00cafb10
14	0x00cb0b10
15	0x00cb1b10
16	0x00cb2b10
17	0x00cb3b10
18	0x00cb4b10
19	0x00cb5b10
20	0x00cb6b10
21	0x00cb7b10
22	0x00cb8b10
23	0x00cb9b10
24	0x00cbab10
25	0x00cbbb10
26	0x00cbcb10
27	0x00cbdb10
28	0x00cbeb10
29	0x00cbfb10



**Figure A.22.:** Embedding of 16 bit word into 256 bit words. The hexadecimal numbers within in the fields correspond to the order of words arriving from the optical link. They are embedded into the 256 bit in a pairwise swapped fashion.

required in order to fit the data into the new scheme. The `convert_16to256` implements the conversion of the data stream, as well as a reformatting of the event info data to work more conveniently with the new data format. In addition to that, the `convert_16to256` implements word order swapping, simplifying the offline raw data reading in software.

When the recording of a particular trigger begins, `convert_16to256` enumerates the incoming 16 bit words starting with 0x0 using a 4 bit counter. The counter overflows every 16 counts (maximum value: 0xF). A 256 bit output word can hold 16 incoming 16 bit words. The incoming 16 bit words get embedded into the 256 bit output according according to Fig. A.22.

Note that 2 consecutive incoming 16 bit words are not inserted strictly sequentially into the output word, but pairwise swapped. The FEE formats the data in tracklets of 32 bit words each, while the CRU logic operates with the base word width of 16 bit of the optical links. The natural order of 16 bit words flowing through the CRU logic corresponds to the ordering in which the words sequentially arrive via the optical link. The pairwise word swapping operation performed at this stage ensures that the offline software can operate with a 32 bit base width. It can later chop the data received from the CRU into chunks of 32 bit each and it is guaranteed that every obtained 32 bit chunk will correspond to a tracklet in its original format.

In case the number of received 16 bit words for a particular event is not a divisible by 16, every unpopulated fields in the last 256 bit output word are padded with the default pattern 0xEEEE.

At the end of every event, the `convert_16to256` receives meta information from the `busy_logic`, namely the signals `event_info_error_i` and `event_info_wordnum_i`. The `convert_16to256` stores these values. The value `event_info_wordnum_i`, indicating the number of 16 bit words the event consists of, gets converted into the number of 256 bit words for downstream logic. As soon as all the data of an event has been pushed to downstream logic, the converted meta data are pushed to downstream logic, using the signals `event_info_error_o`, `event_info_wordnum_o`, `event_info_16bitwords_remaining_o` and `event_info_valid_o`. `event_info_error_o` is the unmodified error state of the `busy_logic` error state machine, `event_info_16bitwords_remaining_o` is the converted event size, `event_info_16bitwords_remaining_o` is the number of 16 bit padding words used, i.e. the division rest and `event_info_valid_o` is a valid signal for the meta data. It is guaranteed that all meta data are pushed within one clock cycle, that every event (even empty ones) produce meta information and that the meta information is pushed to downstream logic only after the event link data payload is complete and has been pushed out entirely.

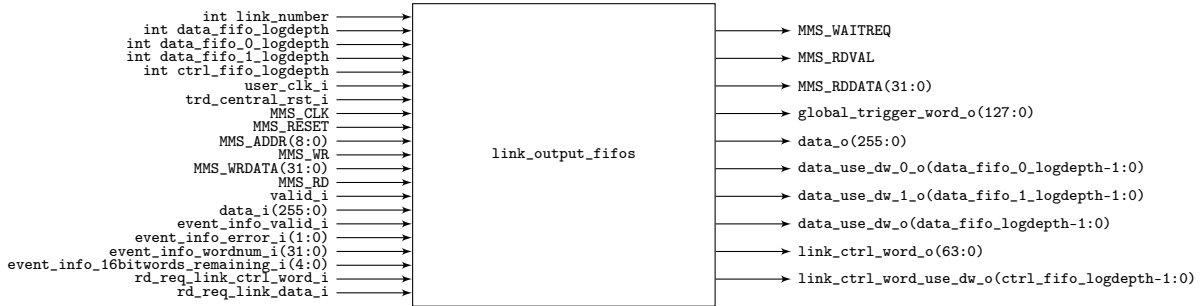
There are no slow control accessible registers provided by `convert_16to256`.

An overview of the input and output signals of the `convert_16to256` is given in Fig. [A.21](#).

## A.19. link\_output\_fifos

The `link_output_fifos` logic receives the data and the meta data for an event from the `convert_16to256` and buffers them. This allows changing from push architecture to pull architecture such that the downstream logic can pick up data and meta data words at any convenient time.

The logic is capable of handling the buffering of multiple events simultaneously without corrupting any data. The available storage can hold up to 96 kB of data, which corresponds

Figure A.23.: `link_output_fifos` diagram.

to 1.5 calibration events or more than 96 physics events assuming the largest possible multiplicity. A maximum of 32 meta data sets for events can be stored, effectively limiting the maximum number of events which can be buffered to 32.

Internally, the `link_output_fifos` logic contains of 3 FIFOs, 2 cascaded FIFOs buffering the data and 1 FIFO buffering all meta data.

The data FIFO cascade behaves from the outside perspective the same way as a larger FIFO with higher depth would behave. The reason for implementing the more complex cascade is a more efficient FPGA RAM resource usage: the data storage capacity of 96 kiB and the given data width of 256 translate to a required FIFO of 256 bit width x 3072 depth. However, the FPGA allows natively only the implementation of FIFOs where both data width and FIFO depth are powers of 2, but  $\log_2(3072) \approx 11, 5896....$  Rounding this up to 12 would result in a depth of 4096 and an unnecessarily high data storage capacity of 128 kiB. Given that a single FPGA has to handle 30 of such links, such wasteful resource utilisation would otherwise bring the FPGA to its limits. Therefore, 2 cascaded FIFOs of 256 bit width each are used, the first one with a depth of 2048 and the second one with a depth of 1024.

The data coming from the `convert_16to256` are immediately fed into the first FIFO with a depth of 2048. This FIFO is configured in show-ahead mode (also known as First Word Fall Through (FWFT) mode). The first word written to this FIFO appears at the

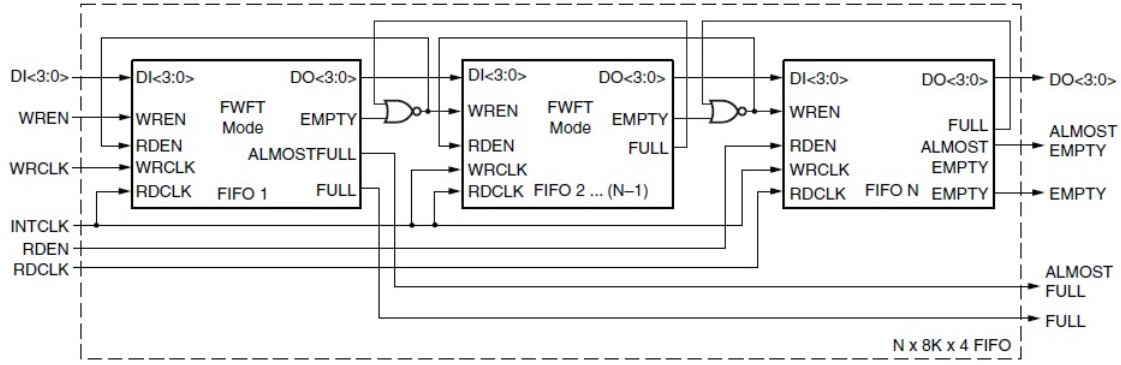


Figure -1 : Example: Cascading Multiple FIFOs by Depth

**Figure A.24.:** Principle of cascading multiple FIFOs to increase FIFO depth [62].

read side of the FIFO without that a prior read request is necessary, i.e. it falls through. In fact, the FIFO still has a read request input signal, but this signals does not act as a read request, but as a read acknowledge and switches the output to show the next data word. The FWFT mode of this first FIFO simplifies the implementation of the FIFO cascade. The second FIFO with a depth of 1024 is connected to the first FIFO following the cascading principle shown in Fig. A.24.

For buffering the event meta data, a separate single FIFO of 64 bit width and a depth of 32 is implemented. In order to buffer all metadata received via the signals `event_info_error_i`, `event_info_wordnum_i` and `event_info_16bitwords_remaining_i` from the `convert_16to256` efficiently, they are embedded into a 64 bit word according to Table A.9 such that the word width corresponds to the FIFO width. These 64 bit container words for meta data are called `link_ctrl_words` in the following.

Each `link_output_fifos` instance provides debug registers readable by slow control in order to supervise and verify the readout chain. An overview with addresses and functional descriptions is given in Table A.10.



**Table A.9.:** Definition of the `link_ctrl_word`. These words contain all event meta data related to a particular link and are used by the downstream logic when processing the data.

Bit	Number of Bits	Description
63..32	31	<code>event_info_wordnum_i</code>
31..26	5	<code>event_info_16bitwords_remaining_i</code>
25..2	25	unused
1..0	2	<code>event_info_error_i</code>

## A.20. Buffering Stage

The in total 30 `link_output_fifos` instances together with the 2 instances of `trigger_word_fifo` form a central buffering stage within the entire TRD custom CRU logic (compare Fig. A.6). For all trigger and data paths this buffering stage can be considered as a boundary between 2 functionally independent blocks of the firmware:

The firmware section upstream with respect to the buffering stage works strictly synchronised together with the trigger system and the state of the TRD FEE. It stores all required information in the buffering stage at a time convenient for the smooth online operation of the readout tree.

The firmware section downstream with the respect to the buffering stage works independently from the time critical readout operation. Its tasks is to read data from the buffering stage and to process it further, but these operations do not require to be strictly synchronous with the readout tree and are thus less time critical. However, the throughput has to be sufficient in order to guarantee that no overflow conditions can occur in any buffer within the buffering stage, as this inevitably would result in corruption and data loss.

**Table A.10.:** Slow control registers provided by `link_output_fifos`.

Register address	r(eadable) / w(ritable)	Description
0x00	r	Always 0x92471934. Serves as a <code>link_output_fifos</code> type identifier allowing the slow control software to verify that <code>link_output_fifos</code> is present and answering.
0x04	r	Number of <code>link_ctrl_words</code> currently stored in the meta data FIFO.
0x08	r	Number of data words currently stored in the first FWFT data FIFO.
0x0c	r	Sticky maximum number of <code>link_ctrl_words</code> stored at anytime since the last logic reset in the meta data FIFO.
0x10	r	Sticky maximum number of data words stored at anytime since the last logic reset in the first FWFT data FIFO.
0x14	r	Meta data FIFO full signal.
0x18	r	First FWFT data FIFO full signal.
0x1c	r	Meta data FIFO sticky full signal. This signal indicates whether the meta data FIFO has experienced an overflow condition at any time since the last logic reset.
0x20	r	First FWFT data FIFO sticky full signal. This signal indicates whether the first FWFT data FIFO has experienced an overflow condition at any time since the last logic reset.
0x24	r	Total implemented depth of the meta data FIFO (always 32).
0x28	r	Total implemented depth of the first FWFT data FIFO (always 2048).
0x2c	w	Sticky signals reset register. Writing 0x1 to this registers resets the values stored in 0x0c, 0x10, 0x1c, 0x20, 0x3c and 0x44 without resetting any other logic.
0x30	r	Total implemented depth of the second data FIFO (always 1024).
0x34	r	Total implemented depth of the entire data FIFO cascade (always 3072).
0x38	r	Second data FIFO full signal.
0x3c	r	Second data FIFO sticky full signal. This signal indicates whether the second data FIFO has experienced an overflow condition at any time since the last logic reset.
0x40	r	Number of data words currently stored in the second data FIFO.
0x44	r	Sticky maximum number of data words stored at anytime since the last logic reset in the second data FIFO.
0x48	r	Total number of data words currently stored in the entire data FIFO cascade. This value is equal to the sum of the values stored in the registers 0x08 and 0x40.

## A.21. packetisers

The **packetisers** are responsible for reading data from the buffering stage, formatting and packetising them according to the requirements described in Section 6.6.3. Furthermore, they push the data to the data path wrappers, using the interface introduced in Section 6.6.3. This makes the data usable for the O2 system.

There are 2 **packetisers** implemented. **Packetiser 0** handles the link data coming from the links 0 - 14 and **Packetiser 1** handles the link data coming from the links 15 - 29. The 2 **packetisers** work fully independently from each other by accessing different FIFOs in the buffering stage and ensuring that no data get mixed up. The **packetisers** push the packetised data to the user logic inputs of the **datapath\_wrappers**, which forward them via 2 separate PCIe endpoints (taking advantage of the PCIe bifurcation) to the FLP.

For trigger related data, **Packetiser 0** reads these information from **trigger\_word\_fifo 0** and **Packetiser 1** reads them from **trigger\_word\_fifo 1**. The trigger data provided by these FIFOs are strict duplicates of each other. The duplication of the trigger data with 2 **trigger\_word\_fifos** ensures that both **packetisers** can work independently and that they can pick up the next word from their respective **trigger\_word\_fifo** at possibly different times.




Figure A.25.: packetiser diagram.

Internally, the packetiser is based on multiple interleaved state machines. The outermost state machine `s_state` switches between the 2 states `wait_for_available_trigger_word` and `ship_data`. In this way, it alters between the most fundamental tasks of the packetiser: surveying the `trigger_word_fifo` for new `global_trigger_words` (state `wait_for_available_trigger_word`) and the actual shipping of packets (state `ship_data`).

### A.21.1. State Machine `s_state: wait_for_available_trigger_word`

A `global_trigger_word` acts as an instruction of what to do next, depending on whether it indicates a heartbeat, physics or calibration trigger and whether it marks the early stage of the trigger reception or the completion of the corresponding data acquisition (also compare the preceding-property, bit 63, Table A.2).

In case that the `global_trigger_word` is non preceding and it represents a physics or calibration trigger, all relevant information from the `global_trigger_word` are stored with internal signals (orbit, BC, event type) in order to reuse these information later for building the TRD header. The `s_state` state machine switches to the state `ship_data` to carry out the actual packet building and packet transmission. Note that the previous logic stages guarantee that in case a non-preceding `global_trigger_word` is available in the `global_trigger_fifo`, all the corresponding link data to form a packet from it, is already available in the `link_output_fifos`. This ensures that no packets can be generated without the necessary data available in the CRU, while at the same time enabling the packetiser to work as independent as possible from the previous logic stages.

Currently, no action is taken for preceding `global_trigger_word` representing a physics or calibration trigger.

Heartbeat triggers are special in the sense that there is no link data directly associated to them and that the previous logic only creates preceding `global_trigger_words`, but no concluding ones. Note however that it is still possible to identify the endings of heartbeat frames implicitly, because any `global_trigger_word` indicating a new heartbeat trigger

also implies that the previous heartbeat frame has ended. This implicit method allows the recognition of finalised heartbeat frames in all situations, including at End of Run (EOR) and when rate throttling is enforced by CTP, because in both situations, it is guaranteed that the CRU continues to receive new heartbeat triggers at the same rate. These heartbeat triggers have the rejected-flag (bit 2) in the TTC-PON message set. Rejected heartbeat triggers also generate preceding `global_trigger_words` and the full TTC-PON message is part of the `global_trigger_word` (compare Table A.2).

**Instruction State Machine: `s_02Instruction`** The packetiser implements the method of heart beat frame opening and closing according to the description in Section 6.6.3. In order to decide implicitly whether a heartbeat frame has to be opened or closed, the packetiser internally keeps track of the fact whether a heartbeat frame is currently opened or closed. Depending on this internal heartbeat frame state and the nature (rejected or not rejected) of the currently processed `global_trigger_word` representing the heartbeat trigger, an instruction `s_02Instruction` is generated, i.e. the state is changed accordingly. This instruction is important for the `ship_data` state to define what to do with respect to heartbeat frames. After that the state is switched to `ship_data`. The possible instruction values are:

- **`s_02Instruction: open_new_hb`**

No heartbeat frame is currently open and now a new heartbeat trigger without reject flag enforces the opening of a new heartbeat frame.

- **`s_02Instruction: close_current_hb_and_open_new_hb`**

A heartbeat frame is currently open and now a new heartbeat trigger without reject flag enforces the closing of the current heartbeat frame and the opening of a new one.

- **`s_02Instruction: close_current_hb`**

A heartbeat frame is currently open and now a heartbeat trigger with the reject flag set enforces the closing of the current heartbeat frame without further actions. This occurs after End of Continuous or Trigger (EOX) or when rate throttling is enforced by CTP by rejecting heartbeat frames.

### A.21.2. State Machine `s_state: ship_data`

The state `ship_data` of the `s_state` state machine is entered whenever a physics or calibration data are ready to be packetised and shipped or when a new heartbeat trigger requires writes on the data stream to open or close heartbeat frames.

The state `ship_data` handles both the RDH packetisation and the TRD packetisation. For these two aspects, 2 interleaved state machines are used. `s_02State` handles the RDH packetisation and `s_TRDState` handles the TRD packetisation.

**State Machine `s_02State`** The `s_02State` state machine organises the generation and the shipping of RDH packets. This process happens while the `s_state` state machine is in `ship_data` mode.

`s_02State` takes the following states:

- `s_02State: sendRDH0`

In this state the RDH header word 0 is pushed to the datapath wrapper. The fields are filled accordingly based on the information gathered previously from the `global_trigger_word`.

The next `s_state` state is not changed. The next `s_02State` state is in all cases `sendRDH1`.

- `s_02State: sendRDH1`

In this state the RDH header word 1 is pushed to the datapath wrapper. The fields are filled accordingly based on the information gathered previously from the `global_trigger_word`.

The next states are dependent on the circumstances under which the RDH header has been sent.

If it is a subsequent RDH that has been sent because the previous packet has reached its maximum size of 8 kB and more data has to be transferred, the next `s_state` state is left unaffected and the next `s_02State` state is `sendPayload`.

If it is an RDH which is sent to open a heartbeat frame as the result of a heartbeat trigger, the next `s_02State` state is `sendPayload`. The `s_state` state is changed to `wait_for_available_trigger_word`, effectively halting the `ship_data` mode until a new `global_trigger_word` is available.

If it is an RDH which is sent to close a heartbeat frame as the result of a heartbeat trigger, the next `s_02State` state is `copy_next_rdh_data`. The `s_state` state is changed to `wait_for_available_trigger_word`, effectively halting the `ship_data` mode until a new `global_trigger_word` is available.

- `s_02State: sendPayload`

In this state the payload of the RDH packets is generated and shipped. This is the only mode in which the `s_TRDState` state machine operates. The `s_TRDState` state machine is not altered in any way outside the `s_02State sendPayload`. This ensures that TRD packetising is entirely independent from RDH packetising.

However, the `s_TRDState` state machine has to count the volume it pushes and make sure not to exceed the maximum of 253 256 bit payload words to stay below 8 kiB in total, RDH included. It must switch `s_02State` to `copy_next_rdh_data` when the limit is reached in order handle the generation of a new RDH packet. Once all TRD payload is shipped, `s_state` is changed to `wait_for_available_trigger_word`, effectively halting the `ship_data` mode until a new `global_trigger_word` is available.

- `s_02State: copy_next_rdh_data`

This state lasts only 1 clock cycles and serves copying information from the `global_trigger_word` into local signals. They are used to build the RDH header subsequently.

The next `s_state` state is left unaffected. The next `s_02State` state is in all cases `wait_safety_time`.



**s\_02State:** wait\_safety\_time

The specifications of the datapath wrapper require some clock cycles of spacing between subsequent packets. This state ensures that by waiting for a configurable number of clock cycles. The number of clock cycles waited is configurable by slow control and defaults to 7 which is a sufficiently conservative choice.

The next **s\_state** state is left unaffected. Once the programmed time limit is reached, the next **s\_02State** state is in all cases **sendRDH0**.

- **s\_02State:** close\_packet\_and\_hb

This state is used in order to close a currently open packet and then initiate the closure of the heartbeat frame in absence of actual TRD link data. A valid 256 bit padding word containing only 16 times 0xEEEE is sent. This is necessary because the EOP flag has to be sent in order to close the currently open packet alongside with a valid data word.

The RDH field STOP BIT is set to 0x1 and **s\_02State** is altered to **wait\_safety\_time**. This ensures that a heartbeat closing packet is generated after the safety delay at the next turnaround of the **s\_02State** state machine.

**State Machine s\_TRDState** The **s\_TRDState** state machine organises the generation and the shipping of the actual TRD link data packets. This process happens while the **s\_state** state machine is in **ship\_data** mode and the **s\_02State** state machine is in **sendPayload** mode.

In all **s\_TRDState** states which can push data to the datapath wrapper interface, a counter keeps track of the total number of pushed words, in order to ensure not exceeding the 8 kiB O2 packet size limit. In case the limit is reached, the control is handed over to the **s\_02State** by changing its state to **wait\_safety\_time**. This leaves the **s\_TRDState** frozen, until a new RDH packet is generated and **s\_02State** is back in **sendPayload**, effectively having given the control back to **s\_TRDState** state machine.

`s_TRDState` takes the following states:

- `s_TRDState: readwordnum`

The state `readwordnum` is responsible to read a word from every meta data FIFO in the corresponding `link_output_fifos`. The relevant information therein are for all handled links stored in internal signals. Among those are the number of data words available to read for each link in the current event as well as the current link error flags for each link in the current event. These meta data are used to build the TRD header.

After having gathered these information, `s_TRDState` is changed to `sendTRD0`.

- `s_TRDState: sendTRD0`

In this state the TRD header word 0 is assembled based on the information gathered in the previous state. It is at the same time pushed to the datapath wrapper interface.

The data push is only executed in case that enough space is available within the current RDH packet, otherwise the control is handed over to the `s_02State` state machine in order to start a new RDH packet.

`s_TRDState` is changed to `sendTRD1` once the TRD header word 0 is pushed out.

- `s_TRDState: sendTRD1`

In this state the TRD header word 1 is assembled based on the information gathered in the `readwordnum` state. It is at the same time pushed to the datapath wrapper interface.

The data push is only executed in case that enough space is available within the current RDH packet, otherwise the control is handed over to the `s_02State` state machine in order to start a new RDH packet.

`s_TRDState` is changed to `sendPayload` once the TRD header word 1 is pushed out.

- `s_TRDState: sendPayload`

This state fills the TRD packet with payload. For that, all data FIFOs from the handled link are looped through and their data are pushed out sequentially into the

packet. The data from a particular link is kept contiguous in the packet, i.e. all data from the particular FIFO has to be processed before the next FIFO can be emptied. The data sizes, which have been gathered in the `readwordnum` state are used in order to only read the data words from each FIFO which belong to the event. This allows the logic in front of the buffering stage to work independently from the packetiser and to already write new data from a subsequent event into the link data FIFOs, even though the current event has not been entirely picked up yet.

Before pushing any data word to the datapath wrapper interface, it is checked whether enough space is available within the current RDH packet, otherwise the control is handed over to the `s_02State` state machine in order to start a new RDH packet.

Once all payload from the current event has left the user logic, the `s_TRDState` is changed to `readwordnum` and `s_state` is changed to `wait_for_available_trigger_word`. This prepares the `s_TRDState` to handle the next event when available and the `s_state` state machine gains back the control. It checks then for the availability of a new `global_trigger_word`, restarting the entire process.



# Appendix B

## Physics background

### B.1. Kinematic variables

The rapidity of a particle is defined as [\[6\]](#)

$$y = \frac{1}{2} \ln \left( \frac{E + p_z c}{E - p_z c} \right), \quad (\text{B.1})$$

where  $E$  and  $p_z$  are the energy and the momentum in the beam direction of the particle, respectively.

The pseudorapidity of a particle is defined as [\[6\]](#)

$$\eta = -\ln \left( \tan \frac{\theta}{2} \right), \quad (\text{B.2})$$

where  $\theta$  is the angle of the particle momentum with respect to the beam axis.


## B.2. Bethe-Bloch formula

When particles pass through a detector, they lose energy by interacting with the detector material. The mean energy loss of the particle per path length is described by the Bethe-Bloch formula [63]

$$-\left\langle \frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[ \frac{1}{2} \ln \left( \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} \right) - \beta^2 - \frac{\delta}{2} \right], \quad (\text{B.3})$$

where  $K = 4\pi N_A r_e^2 m_e c^2$  with  $r_e$  being the classical electron radius,  $m_e$  is the electron mass,  $\beta = v/c$  is the velocity of the particle with respect to the speed of light,  $\gamma$  is the Lorentz factor,  $z$  is the charge number of the particle,  $Z$  is the atomic number of the detector material,  $A$  is the atomic mass of the detector material,  $I$  is the mean excitation energy of the detector material,  $\delta$  is the density-effect correction to the ionisation energy loss and  $T_{max}$  is the maximum kinetic energy transferred in a single collision [63].

Figure B.1 shows the mean energy loss for muons, pions, and protons. At low values of  $\beta\gamma$ , the energy loss is dominated by the factor  $1/\beta^2$  and reaches a broad minimum around  $\beta\gamma \approx 4$ . Particles which have a mean energy loss close to the minimum are called Minimum-ionising Particles (MIPs). For higher values of  $\beta\gamma$ , the mean energy loss rises as the logarithmic factor increases [63].



**Figure B.1.:** Mean energy loss rate in liquid (bubble chamber) hydrogen, gaseous helium, carbon, aluminium, iron, tin, and lead. Figure from Ref. [63].





# Appendix C

## List of Acronyms

<b>ACORDE</b>	ALICE Cosmic Ray Detector
<b>AD</b>	ALICE Diffractive Detector
<b>ADC</b>	Analogue-to-Digital converter
<b>AFULL</b>	Almost Full
<b>ALF</b>	Alice Low-level Frontend
<b>ALICE</b>	A Large Ion Collider Experiment
<b>APC</b>	Automatic Power Control
<b>ASIC</b>	Application-Specific Integrated Circuit
<b>ATX</b>	Advanced Technology EXtended
<b>BC</b>	Bunch Crossing
<b>BM</b>	Board Merger
<b>BMBF</b>	Bundesministerium für Bildung und Forschung
<b>CBM</b>	Compressed Baryonic Matter
<b>CERN</b>	Conseil Européen pour la Recherche Nucléaire
<b>COTS</b>	Commercial off-the-Shelf

<b>CPU</b>	Central Processing Unit
<b>CPV</b>	Charged Particle Veto Detector
<b>CR1</b>	Counting Room 1
<b>CR2</b>	Counting Room 2
<b>CR3</b>	Counting Room 3
<b>CR4</b>	Counting Room 4
<b>CR5</b>	Counting Room 5
<b>CRU</b>	Common Readout Unit
<b>CTP</b>	Central Trigger Processor
<b>DAQ</b>	Data Acquisition
<b>DCal</b>	Di-Jet Calorimeter
<b>DCS</b>	Detector Control System
<b>DIM</b>	Distributed Information Management System
<b>DMA</b>	Direct Memory Access
<b>DQM</b>	Data Quality Monitoring
<b>ECS</b>	Experiment Control System
<b>EDR</b>	Engineering Design Review
<b>EMCal</b>	Electromagnetic Calorimeter
<b>EOC</b>	End of Continuous
<b>EOP</b>	End of Packet
<b>EOR</b>	End of Run
<b>EOT</b>	End of Trigger

---

<b>EOX</b>	End of Continuous or Trigger
<b>EPN</b>	Event Processing Node
<b>ESD</b>	Event Storage Data
<b>YETS</b>	Year-End Technical Stop
<b>FAIR</b>	Facility for Antiproton and Ion Research
<b>FEE</b>	Front End Electronics
<b>FIFO</b>	First In First Out buffer
<b>FLP</b>	First Level Processor
<b>FMD</b>	Forward Multiplicity Detector
<b>FPGA</b>	Field Programmable Gate Array
<b>FRED</b>	Front End Device
<b>FSM</b>	Finite State Machine
<b>FWFT</b>	First Word Fall Through
<b>GBT</b>	Giga-Bit Transceiver
<b>GPIO</b>	General Purpose Input Output
<b>GPU</b>	Graphics Processing Unit
<b>GTU</b>	Global Tracking Unit
<b>GUI</b>	Graphical User Interface
<b>HB</b>	Heartbeat
<b>HC</b>	Half Chamber
<b>HCM</b>	Half Chamber Merger
<b>HJT</b>	TRD Jet Trigger

<b>HLT</b>	High Level Trigger
<b>HNU</b>	TRD Nuclei Trigger
<b>HQU</b>	TRD Quarkonia Trigger
<b>HSE</b>	TRD Single Electron Trigger
<b>HV</b>	High Voltage
<b>I2C</b>	Inter-Integrated Circuit
<b>IC</b>	Integrated Circuit
<b>IP</b>	Intellectual Property
<b>ISP</b>	Internet Service Provider
<b>ITS</b>	Inner Tracking System
<b>JTAG</b>	Joint Test Action Group
<b>L0</b>	Level 0
<b>L1</b>	Level 1
<b>L2</b>	Level 2
<b>L2a</b>	Level 2 accept
<b>L2r</b>	Level 2 reject
<b>LC</b>	Lucent Connector
<b>LCR</b>	Late Conversion Rejection
<b>LED</b>	Light Emitting Diode
<b>LHC</b>	Large Hadron Collider
<b>LHCb</b>	Large Hadron Collider beauty
<b>LM</b>	Level Minus One

---

<b>LME</b>	Link Monitor Error
<b>LTU</b>	Local Trigger Unit
<b>LUT</b>	Lookup Table
<b>LVDS</b>	Low Voltage Differential Signalling
<b>HM</b>	High Multiplicity
<b>HMPID</b>	High Momentum Particle Identification Detector
<b>MB</b>	Minimum Bias
<b>MC</b>	Monte Carlo
<b>MCH</b>	Muon Chamber
<b>MCM</b>	Multi Chip Module
<b>MIP</b>	Minimum-ionising Particle
<b>MPO</b>	Multi-fibre Push-On
<b>MWPC</b>	Multi Wire Proportional Chamber
<b>NI</b>	Network Interface
<b>O2</b>	Online Offline Processing System
<b>OLT</b>	Optical Line Terminal
<b>OM3</b>	Optical Multimode 3
<b>OMA</b>	Optical modulation amplitude
<b>ONU</b>	Optical Network Unit
<b>ORI</b>	Optical Readout Interface
<b>P2</b>	LHC Point 2 - ALICE
<b>PASA</b>	Preamplifier and Shaper

<b>Pb–Pb</b>	Lead – Lead
<b>PB</b>	Physics Board
<b>PCB</b>	Printed Circuit Board
<b>PCIe</b>	Peripheral Component Interconnect Express
<b>PHOS</b>	Photon Spectrometer
<b>PID</b>	Particle Identification
<b>PLL</b>	Phase Locked Loop
<b>PM</b>	Photo Multiplier
<b>PMD</b>	Photon Multiplicity Detector
<b>PON</b>	Passive Optical Network
<b>POR</b>	Power On Reset
<b>pp</b>	Proton – Proton
<b>pPb</b>	Proton – Lead
<b>PRBS</b>	Pseudo Random Bit Stream
<b>PRF</b>	Pad Response Function
<b>PRIZM</b>	PRIZM®
<b>PT</b>	Pretrigger
<b>QCD</b>	Quantum Chromodynamics
<b>QGP</b>	Quark-Gluon Plasma
<b>RAM</b>	Random Access Memory
<b>RC</b>	Run Coordination
<b>RDH</b>	Raw Data Header

---

<b>RDHv6</b>	Raw Data Header, version 6
<b>RGB</b>	Red–Green–Blue
<b>RICH</b>	Ring-imaging Cherenkov Detector
<b>ROB</b>	Readout Board
<b>ROC</b>	Readout Chamber
<b>RX</b>	Receiver
<b>SCSN</b>	Slow Control Serial Network
<b>SDD</b>	Silicon Drift Detector
<b>SEU</b>	Single Event Upset
<b>SFP</b>	Small Form Factor Pluggable
<b>SM</b>	Super Module
<b>SMU</b>	Super Module Unit
<b>SOC</b>	Start of Continuous
<b>SOP</b>	Start of Packet
<b>SOR</b>	Start of Run
<b>SOT</b>	Start of Trigger
<b>SOX</b>	Start of Continuous or Trigger
<b>SPD</b>	Silicon Pixel Detector
<b>SRAM</b>	Static Random Access Memory
<b>SSD</b>	Silicon Strip Detector
<b>STF</b>	Sub Time Frame
<b>STFB</b>	Sub Time Frame Builder

<b>SVN</b>	Subversion
<b>TDM</b>	Time Division Multiplexing
<b>TF</b>	Time Frame
<b>TGU</b>	Trigger Generation Unit
<b>TI</b>	Texas Instruments
<b>TMU</b>	Track Matching Unit
<b>TOF</b>	Time of Flight
<b>TPC</b>	Time Projection Chamber
<b>TR</b>	Transition Radiation
<b>TRAP</b>	Tracklet Processor
<b>TRD</b>	Transition Radiation Detector
<b>TTC</b>	Trigger and Timing Control
<b>TTCex</b>	Trigger and Timing Control Emitter
<b>TTCrX</b>	Trigger and Timing Control Receiver chip
<b>TX</b>	Transmitter
<b>T0</b>	Time 0 Detector
<b>USB</b>	Universal Serial Bus
<b>VCSEL</b>	Vertical Cavity Surface Emitting Laser
<b>VME</b>	Versa Module Eurocard
<b>VHDL</b>	Very High Speed Integrated Circuit Hardware Description Language
<b>V0</b>	VZERO Detector
<b>WDM</b>	Wavelength Division Multiplexing



---

<b>Xe</b>	Xenon
<b>ZDC</b>	Zero Degree Calorimeter



# List of figures

2.1.	QCD running coupling constant $\alpha_s$ as a function of the momentum transfer $Q$ . The results are obtained by different analyses from various experiments. Figure from Ref. [12]. . . . .	8
2.2.	Schematic view of the temperature $T$ as a function of the chemical potential $\mu_B$ of the QCD phase diagram. Figure from Ref. [13]. . . . .	9
2.3.	Schematic time evolution of a central high-energy heavy-ion collision with QGP formation. Figure from Ref. [14]. . . . .	9
3.1.	Nucleon composition of deuteron, triton and ${}^3\text{He}^{2+}$ . . . . .	12
3.2.	$p_T$ integrated yields ( $dN/dy$ ) of antiprotons, antideuterons and $\overline{{}^3\text{He}^{2+}}$ nuclei as a function of the number of antinucleons in pp collisions at $\sqrt{s} = 7\text{ TeV}$ . Figure from Ref. [22]. . . . .	16
3.3.	Antideuteron to deuteron ratios ( $\bar{d}/d$ ) as a function of the transverse momentum per nucleon $p_T/A$ in pp collisions in different collision energies. The $(\bar{p}/p)^2$ ratios are shown for comparison. Figure from Ref. [22]. . . .	17
3.4.	Antideuteron to deuteron ratios as a function of $p_T$ in different multiplicity classes in pp collisions at $\sqrt{s} = 7\text{ TeV}$ . Figure from Ref. [23]. . . . .	17
3.5.	Coalescence parameter $B_2$ of deuterons and antideuterons as a function of the transverse momentum per nucleon $p_T/A$ in different multiplicity classes in pp collisions at $\sqrt{s} = 7\text{ TeV}$ . Figure from Ref. [23]. . . . .	18
3.6.	Coalescence parameter $B_2$ of deuterons and antideuterons, shown separately, as a function of the transverse momentum per nucleon $p_T/A$ in pp for different collision energies. Figure from Ref. [22]. . . . .	18

3.7.	Coalescence parameters as a function of the transverse momentum per nucleon $p_T/A$ in pp collisions at $\sqrt{s} = 13$ TeV for different multiplicity classes. Left hand panel: Deuterons and antideuterons (parameter $B_2$ ). ${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$ (parameter $B_3$ ). Figures from Ref. [24]. . . . .	19
4.1.	Sectional view of the ALICE experiment during Run 2. Figure from Ref. [32]. . . . .	24
4.2.	Top panel: Separation of different particle species for PID with the TPC. The figure shows for each particle the reconstructed particle momentum on the horizontal axis TPC energy loss on the vertical axis. Bottom panel: Separation of different particle species for PID with the TOF detector. The figure shows for each particle the reconstructed particle momentum on the horizontal axis and the particle velocity measured with TOF on the vertical axis. Both figures have been adapted from [33]. . . . .	26
5.1.	Structure of the TRD [34,40]. . . . .	35
5.2.	Principle of particle detection within a ROC [34,40]. . . . .	36
5.3.	Photon absorption length of different noble gases vs. energy of the impinging photon [2]. . . . .	37
5.4.	Illustration of tilted readout pads in adjacent layers. This is used to improve the tracking resolution beyond the limitations imposed by the finite number of readout pads and channels [34,40]. . . . .	39
5.5.	Schematic processing steps in the FED [39]. . . . .	44
5.6.	Typical timing of the involved processing steps when reading out an event with the TRD [36]. . . . .	44
5.7.	MC Production LHC14a6, Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV Hijing nuclei injection for sufficient statistics, especially with respect to ${}^3_2\text{He}^{2+}$ , ${}^4_2\text{He}^{2+}$ and their respective antiparticles. . . . .	55
5.8.	MC Production LHC14b3a, pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV, triggered TRD and EMCal data for sufficient electron statistics. . . . .	55
5.9.	Distribution of TRD tracks in the MC production LHC14a6 across sectors and stacks in the TRD. . . . .	56

5.10.	Fraction of ${}^3_2\text{He}^{2+}$ , ${}^4_2\text{He}^{2+}$ and their respective antiparticles measured by TRD with respect to all reconstructed tracks of the corresponding particle species. Statistical uncertainties are shown as vertical bars. Small uncertainties are hidden by the point markers. . . . .	57
5.11.	PID value as a function of $p_T$ . Top left panel: ${}^3_2\text{He}^{2+}$ . Top right panel: ${}^4_2\text{He}^{2+}$ . Bottom left panel: $\overline{{}^3_2\text{He}^{2+}}$ . Bottom right panel: $\overline{{}^4_2\text{He}^{2+}}$ . . . . .	59
5.12.	Trigger condition efficiency as a function of $p_T$ , shown separately for the different $Z = 2$ nuclei for various PID cuts. Statistical uncertainties are shown as vertical bars. . . . .	61
5.13.	Shipped 8-bit value as a function of the measured 16-bit raw ADC value (linear LUT). . . . .	63
5.14.	Shipped 8-bit value as a function of the measured 16-bit raw ADC value (electron probability LUT). . . . .	64
5.15.	<i>Acceptance efficiency</i> as a function of $p_T/Z$ of nuclei in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV compared to the average <i>acceptance efficiency</i> in pp collision at $\sqrt{s} = 8$ TeV. Statistical uncertainties are shown as vertical bars. . . . .	65
5.16.	Normalised distribution of PID values for ${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$ in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. Statistical uncertainties are shown as vertical bars. . . . .	66
5.17.	Reverse integrated charge deposition distributions of ${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$ in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV, which corresponds to the trigger condition efficiency as a function of the PID threshold. Statistical uncertainties are shown as vertical bars. Small error bars are hidden by the point markers. . . . .	67
5.18.	Normalised PID distribution of ${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$ in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV compared to the normalised background spectrum in pp collision at $\sqrt{s} = 8$ TeV. Statistical uncertainties are shown as vertical bars. . . . .	68
5.19.	Normalised distribution of 4, 5 and 6 tracklet cases, separately shown for ${}^3_2\text{He}^{2+}$ and $\overline{{}^3_2\text{He}^{2+}}$ in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV and background (other particles) in pp collision at $\sqrt{s} = 8$ TeV. Statistical uncertainties are shown as vertical bars. . . . .	69

- 5.20. Contour lines of constant inverse trigger rejection factors as a function of the trigger cuts in pp collisions at  $\sqrt{s} = 8$  TeV. Exemplary projection from 6-dimensional optimisation. The marked contour line corresponds to the target inverse trigger rejection factor of 1 /1000. . . . . 72
- 5.21. Average trigger efficiency for  ${}^3\text{He}^{2+}$  (arb. units) as a function of the trigger cuts in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV. Exemplary projection from 6-dimensional optimisation. The marked area corresponds to the target inverse trigger rejection factor 1/1000. . . . . 72
- 5.22. Trigger condition efficiency on real data after optimisation and for a rejection factor of 1000 obtained in Pb–Pb collisions at  $\sqrt{s_{\text{NN}}} = 2.76$  TeV. Statistical uncertainties are shown as vertical bars. . . . . 73
- 5.23. Shipped 8-bit value as a function of the measured 16-bit raw ADC value (electron probability LUT from Fig. 5.14 in red color and newly introduced LUT in blue color). The newly introduced LUT allows for the coexistence between electron and nuclei triggers and it linearises the electron probability LUT in the electron/pion overlap region. After the kink point, the new LUT is optimised for triggering on nuclei. . . . . 75
- 5.24. TPC signal (arbitrary units) as a function of the rigidity  $p/Z$  in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Comparison of the HNU triggered data set (red points) with an equally sized reference data set (black points). The blue splines indicate the expected Bethe-Bloch energy loss within the TPC for, from bottom to top, protons, deuterons, tritons,  ${}^3\text{He}^{2+}$  and  ${}^4\text{He}^{2+}$  as well as their antiparticles. . . . . 77
- 5.25. Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^3\text{He}^{2+}$  and  $\overline{{}^3\text{He}^{2+}}$  as a function of the rigidity in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. The green lines indicate the PID relevant boundaries. . . . . 79
- 5.26. Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^4\text{He}^{2+}$  and  $\overline{{}^4\text{He}^{2+}}$  as a function of the rigidity in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. The green lines indicate the PID relevant boundaries. . . . . 79

- 5.27. TOF  $\beta$  as a function of the transverse rigidity for the TPC pre-selected  ${}^3_2\text{He}^{2+}$ ,  $\overline{{}^3_2\text{He}^{2+}}$ ,  ${}^4_2\text{He}^{2+}(\alpha)$  and  $\overline{{}^4_2\text{He}^{2+}}(\bar{\alpha})$  candidates in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Candidates pre-selected as  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  are marked as black points, while candidates pre-selected as  ${}^4_2\text{He}^{2+}$  and  $\overline{{}^4_2\text{He}^{2+}}$  are marked as red points. . . . . 80
- 5.28. *Acceptance efficiency* as a function of the transverse momentum  $p_{\text{T}}$  for  ${}^3_2\text{He}^{2+}$  (red markers) and  $\overline{{}^3_2\text{He}^{2+}}$  (blue markers) candidates in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Statistical uncertainties are shown as vertical bars. . . 81
- 5.29. Efficiency of the trigger condition as a function of the transverse momentum  $p_{\text{T}}$  in pPb collisions at  $\sqrt{s_{\text{NN}}} = 8$  TeV. Left hand panel:  ${}^3_2\text{He}^{2+}$  (red markers) and  $\overline{{}^3_2\text{He}^{2+}}$  (blue markers). Right hand panel:  ${}^4_2\text{He}^{2+}(\alpha)$ .  $\overline{{}^4_2\text{He}^{2+}}(\bar{\alpha})$  is not shown due to low statistics. Statistical uncertainties are shown as vertical bars. . . . . 82
- 5.30. Left hand panel: Total efficiency of the HNU trigger on pPb data for  ${}^3_2\text{He}^{2+}$  (red markers) and  $\overline{{}^3_2\text{He}^{2+}}$  (blue markers). Right hand panel: Total efficiency of the HNU trigger on pPb data for  ${}^4_2\text{He}^{2+}$ .  $\overline{{}^4_2\text{He}^{2+}}$  is not shown due to low statistics. . . . . 83
- 5.31. TPC signal  $dE/dx$  in arbitrary units as a function of the rigidity  $|p/Z|$  in pp collisions at  $\sqrt{s} = 13$  TeV. The red splines correspond to the expected Bethe-Bloch energy loss for each particle species (from top to bottom:  ${}^4_2\text{He}^{2+}$ ,  ${}^3_2\text{He}^{2+}$ , triton, deuteron, proton). Particles and their respective antiparticles are combined. Top panel: MB data set with approximately  $4.6 \cdot 10^7$  MB events. Bottom panel: nuclei triggered data set with approximately  $1.9 \cdot 10^7$  events. . . . . 85
- 5.32. Relative deviation of the TPC signal from the expected Bethe-Bloch energy loss of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of the transverse rigidity  $|p_{\text{T}}/Z|$  in pp collisions at  $\sqrt{s} = 13$  TeV. Top panel: MB data set. Bottom panel: nuclei triggered data set. . . . . 86
- 5.33. Parameter  $m^2/Z^2$  of  ${}^3_2\text{He}^{2+}$  and  $\overline{{}^3_2\text{He}^{2+}}$  as a function of the transverse rigidity  $p_{\text{T}}/Z$  in pp collisions at  $\sqrt{s} = 13$  TeV. Top panel: MB data set. Bottom panel: nuclei triggered data set. . . . . 89

5.34.	Enhancement factor of the nuclei trigger, shown separately for each particle species and antiparticle species, with respect to MB in pp collisions at $\sqrt{s} = 13$ TeV. Statistical uncertainties are shown as vertical bars. . . .	90
6.1.	Tracklet data format . . . . .	95
6.2.	Photograph of the second version of the CRU PCIe card. Major improvements with respect to first version are highlighted. Figure from Ref. [52]. . . . .	97
6.3.	New laboratory setup for testing the upgraded TRD readout chain at the Institut für Kernphysik, Universität Münster. . . . .	107
6.4.	Eye diagram adapted from the LTC5100 datasheet: Emcore mode LC-TOSA VCSEL, $2^{15}$ PRBS, 10 dB extinction ratio, 300 $\mu$ W average power [54].	110
6.5.	Optical powers received at GTU vs. laser current setpoint. . . . .	114
6.6.	Optical power measured by the CRU as a function of the laser current setpoint of a lab ORI board. . . . .	118
6.7.	Rate of synchronisation losses at logarithmic (top panel) and linear (bottom panel) scale. Vertical bars are statistical errors. They are mostly hidden by the points. . . . .	119
6.8.	Optical powers received at CRU vs. optical powers previously received at the GTU, SM05, A-side. The different channels are identified as SM_stack_layer_HC. . . . .	120
6.9.	Optical powers received at CRU vs. optical powers previously received at the GTU, SM05, C-side. The different channels are identified as SM_stack_layer_HC. . . . .	120
6.10.	Measured effective additional attenuations SM05, A-side. The different channels are identified as SM_stack_layer_HC. . . . .	121
6.11.	Measured effective additional attenuations SM05, C-side. The different channels are identified as SM_stack_layer_HC. . . . .	121
6.12.	Overview of the common CRU firmware [57]. . . . .	132
6.13.	Block diagram of the data path link to connect a detector user logic [57].	136
6.14.	Block diagram of the common logic datapath wrapper [57]. . . . .	137



6.15. Composition of the RDH Header word 0. Fields in red colour are only allocated by the user logic and later overwritten by the common logic with values generated within the common logic. All fields in blue colour are allocated and filled with values generated by the user logic. They are guaranteed to not be manipulated anymore by the common logic. Adapted from [58] . . . . .	138
6.16. Composition of the RDH Header word 1. Fields in red colour are only allocated by the user logic and later overwritten by the common logic with values generated within the common logic. All fields in blue colour are allocated and filled with values generated by the user logic. They are guaranteed to not be manipulated anymore by the common logic. Adapted from [58] . . . . .	139
6.17. Details on FEE ID field. . . . .	144
6.18. Composition of the TRD Header word 0. Indices outside brackets are valid for endpoint 0, while indices in brackets are valid for endpoint 1. .	146
6.19. Composition of the TRD Header word 1. Indices outside brackets are valid for endpoint 0, while indices in brackets are valid for endpoint 1. .	147
6.20. Details on error flags field. . . . .	149
6.21. Sampled ADC values during drift time in one of the first pp collisions in Run 3. The dark blue line shows the mean ADC value. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59]. . . . .	152
6.22. Event display of an exemplary full Time Frame (TF) in pp collisions at $\sqrt{s} = 13.6$ TeV from run 520145 recorded on July 5th, 2022. The tracklets in the TRD are illustrated with blue markers. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59]. . . . .	153
6.23. 3-dimensional view of the event display of an exemplary full Time Frame (TF) in pp collisions at $\sqrt{s} = 13.6$ TeV from run 520145 recorded on July 5th, 2022. ALICE internal figure, kindly provided by Dr. Ole Schmidt [59].	153
6.24. 3-dimensional view of the event display of the first Pb–Pb collision at $\sqrt{s_{NN}} = 5.36$ TeV in Run 3 recorded by ALICE on November 18th, 2022. The tracklets in the TRD are illustrated with cyan markers. Figure from Ref. [60]. . . . .	154

A.1.	trd_custom_module_wrapper diagram. . . . .	166
A.2.	ori_wrapper diagram. . . . .	167
A.3.	ModelSim® simulation to test the 8B10B decoder within the IP core. A counter with control characters is encoded by an external encoder and decoded by the IP core, which recovers the original data stream. . . . .	169
A.4.	ori_link diagram. . . . .	169
A.5.	trd_user_logic diagram. . . . .	170
A.6.	Overview of the TRD CRU user logic. Red arrows depict the main readout data flow, green arrows depict the trigger data flow, blue arrows depict control signals and the orange arrow is BUSY. . . . .	171
A.7.	buffered_ttc_pon_handler diagram. . . . .	172
A.8.	cdc_ttc_pci diagram. . . . .	172
A.9.	trigger_buffer diagram. . . . .	173
A.10.	ttc_pon_handler diagram. . . . .	175
A.11.	ttc_pon_handler state machine diagram. . . . .	176
A.12.	trigger_word_fifo diagram. . . . .	179
A.13.	aggregate_single_link_cores diagram. . . . .	180
A.14.	single_link_core diagram. . . . .	181
A.15.	link_realigner diagram. . . . .	182
A.16.	ready_valid_checker diagram. . . . .	183
A.17.	debug_memory diagram. . . . .	183
A.18.	cdc_ori_pci diagram. . . . .	184
A.19.	trd_ddg diagram. . . . .	186
A.20.	busy_logic diagram. . . . .	187
A.21.	convert_16to256 diagram. . . . .	194
A.22.	Embedding of 16 bit word into 256 bit words. The hexadecimal numbers within in the fields correspond to the order of words arriving from the optical link. They are embedded into the 256 bit in a pairwise swapped fashion. . . . .	197
A.23.	link_output_fifos diagram. . . . .	199
A.24.	Principle of cascading multiple FIFOs to increase FIFO depth [62]. . . . .	200
A.25.	packetiser diagram. . . . .	204

- B.1. Mean energy loss rate in liquid (bubble chamber) hydrogen, gaseous helium, carbon, aluminium, iron, tin, and lead. Figure from Ref. [63]. . 215



# List of tables

4.1.	Trigger levels used in ALICE during Run 2 [36]. . . . .	30
5.1.	Format of the 32 bit wide word generated for every found tracklet by the MCM [42]. . . . .	45
5.2.	Estimated required trigger cuts for a target trigger rejection factor of 1000 for a maximum ${}^3_2\text{He}^{2+}$ trigger efficiency. . . . .	71
5.3.	Transverse rigidity ranges $ p_T/Z $ for the identification of light nuclei. . .	87
6.1.	Contributions to the attenuation on the optical path. . . . .	111
6.2.	Composition of the prolonged optical test path. . . . .	115
6.3.	Mapping of CRU ORI link inputs to stacks, layers and HCs for the A-side.125	
6.4.	Mapping of CRU ORI link inputs to stacks, layers and HCs for the C-side.126	
A.1.	TRD_trig_word format. It is guaranteed that for every trigger (heartbeat, physics, calibration) separate TRD_trig_words are generated, i.e. only one of the 3 last bits can be high in the same TRD_trig_word. This is ensured by trigger_buffer, even in case that CTP sends a heartbeat accept together with a physics or calibration trigger within the same clock cycle. . . . .	174
A.2.	Composition of the global_trigger_word. . . . .	178
A.3.	Slow control registers provided by the shared avalon slave in single_link_core for debug_memory. . . . .	185
A.4.	Slow control registers provided by the shared Avalon slave in single_link_core for trd_ddg. . . . .	188
A.5.	Slow control registers provided by busy_logic. . . . .	192
A.6.	Legend for state machine states in numerical representation. . . . .	193

A.7.	Full addresses of the <code>busy_logic</code> error state registers 0x28 for all <code>busy_logic</code> present within a single CRU. . . . .	195
A.8.	Full addresses of the <code>busy_logic</code> error state reset registers 0x10 for all <code>busy_logic</code> present within a single CRU. . . . .	196
A.9.	Definition of the <code>link_ctrl_word</code> . These words contain all event meta data related to a particular link and are used by the downstream logic when processing the data. . . . .	201
A.10.	Slow control registers provided by <code>link_output_fifos</code> . . . . .	202

# Bibliography

- [1] **ALICE** Collaboration, “Measurement of anti- $^3\text{He}$  nuclei absorption in matter and impact on their propagation in the Galaxy,” *Nature Phys.* **19** no. 1, (2023) 61–71, [arXiv:2202.01549 \[nucl-ex\]](#).
- [2] **ALICE** Collaboration, “The ALICE Transition Radiation Detector: Construction, operation, and performance,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **881** (2018) 88–127.  
<https://www.sciencedirect.com/science/article/pii/S0168900217309981>.
- [3] **ALICE** Collaboration, “ALICE upgrades during the LHC Long Shutdown 2,” [arXiv:2302.01238 \[physics.ins-det\]](#).
- [4] “Deutsches Technisches Doktorandenprogramm am CERN (Wolfgang-Gentner-Stipendien).”  
<https://wolfgang-gentner-stipendien.web.cern.ch>, 2016.
- [5] F. Halzen and A. Martin, *Quarks and Leptons: An introductory course in modern particle physics*. John Wiley & Sons, New York, USA, 1984.
- [6] C. Y. Wong, *Introduction to high-energy heavy ion collisions*. 1995.
- [7] J. D. Bjorken, “Highly relativistic nucleus-nucleus collisions: The central rapidity region,” *Phys. Rev. D* **27** (Jan, 1983) 140–151.  
<https://link.aps.org/doi/10.1103/PhysRevD.27.140>.
- [8] **STAR** Collaboration, “Experimental and theoretical challenges in the search for the quark–gluon plasma: The STAR Collaboration’s critical assessment of the evidence from RHIC collisions,” *Nuclear Physics A* **757** no. 1, (2005) 102–183.

- <https://www.sciencedirect.com/science/article/pii/S0375947405005294>.
- [9] **LHCb** Collaboration, “Observation of  $J/\psi p$  Resonances Consistent with Pentaquark States in  $\Lambda_b^0 \rightarrow J/\psi K^- p$  Decays,” *Phys. Rev. Lett.* **115** (Aug, 2015) 072001. <https://link.aps.org/doi/10.1103/PhysRevLett.115.072001>.
- [10] **LHCb** Collaboration, “Observation of an exotic narrow doubly charmed tetraquark,” *Nature Phys.* **18** no. 7, (2022) 751–754, [arXiv:2109.01038](https://arxiv.org/abs/2109.01038) [hep-ex].
- [11] J. Ongena, “Fusion: a true challenge for an enormous reward,” *EPJ Web Conf.* **189** (2018) 00015. <https://doi.org/10.1051/epjconf/201818900015>.
- [12] **ATLAS** Collaboration, “Measurement of dijet azimuthal decorrelations in  $pp$  collisions at  $\sqrt{s} = 8$  TeV with the ATLAS detector and determination of the strong coupling,” *Phys. Rev. D* **98** (Nov, 2018) 092004. <https://link.aps.org/doi/10.1103/PhysRevD.98.092004>.
- [13] J. N. Guenther, “Overview of the QCD phase diagram,” *Eur. Phys. J. A* **57** (2021) 136. <https://doi.org/10.1140/epja/s10050-021-00354-6>.
- [14] A. K. Chaudhuri, “Modeling relativistic heavy ion collisions,” in *A Short Course on Relativistic Heavy Ion Collisions*, 2053-2563, pp. 7–1 to 7–60. IOP Publishing, 2014. <https://dx.doi.org/10.1088/bk978-0-750-31060-4ch7>.
- [15] **ALICE** Collaboration, “Two-particle transverse momentum correlations in  $pp$  and  $p$ -Pb collisions at energies available at the CERN Large Hadron Collider,” *Phys. Rev. C* **107** (May, 2023) 054617. <https://link.aps.org/doi/10.1103/PhysRevC.107.054617>.
- [16] B. Dönigus, “Selected highlights of the production of light (anti-)(hyper-)nuclei in ultra-relativistic heavy-ion collisions,” *Eur. Phys. J. A* **56** no. 11, (2020) 280.
- [17] S. T. Butler and C. A. Pearson, “Deuterons from High-Energy Proton Bombardment of Matter,” *Phys. Rev. Lett.* **7** (1961) 69–71.
- [18] S. T. Butler and C. A. Pearson, “Deuterons from High-Energy Proton Bombardment of Matter,” *Phys. Rev.* **129** (1963) 836–842.
- [19] A. Schwarzschild and C. Zupancic, “Production of Tritons, Deuterons, Nucleons, and



- Mesons by 30-GeV Protons on A-1, Be, and Fe Targets,” *Phys. Rev.* **129** (1963) 854–862.
- [20] P. Braun-Munzinger and B. Dönigus, “Loosely-bound objects produced in nuclear collisions at the LHC,” *Nucl. Phys. A* **987** (2019) 144–201, [arXiv:1809.04681 \[nucl-ex\]](#).
- [21] L. P. Csernai and J. I. Kapusta, “Entropy and Cluster Production in Nuclear Collisions,” *Phys. Rept.* **131** (1986) 223–318.
- [22] **ALICE** Collaboration, “Production of deuterons, tritons,  $^3\text{He}$  nuclei and their antinuclei in pp collisions at  $\sqrt{s} = 0.9, 2.76$  and 7 TeV,” *Phys. Rev. C* **97** no. 2, (2018) 024615, [arXiv:1709.08522 \[nucl-ex\]](#).
- [23] **ALICE** Collaboration, “Multiplicity dependence of (anti-)deuteron production in pp collisions at  $\sqrt{s} = 7$  TeV,” *Phys. Lett. B* **794** (2019) 50–63, [arXiv:1902.09290 \[nucl-ex\]](#).
- [24] **ALICE** Collaboration, “Production of light (anti)nuclei in pp collisions at  $\sqrt{s} = 13$  TeV,” *JHEP* **01** (2022) 106, [arXiv:2109.13026 \[nucl-ex\]](#).
- [25] A. Andronic, P. Braun-Munzinger, J. Stachel, and H. Stöcker, “Production of light nuclei, hypernuclei and their antiparticles in relativistic nuclear collisions,” *Physics Letters B* **697** no. 3, (2011) 203–207.  
<https://www.sciencedirect.com/science/article/pii/S0370269311001006>.
- [26] **ALICE** Collaboration, “Production of light nuclei and anti-nuclei in  $pp$  and Pb-Pb collisions at energies available at the CERN Large Hadron Collider,” *Phys. Rev. C* **93** (Feb, 2016) 024917. <https://link.aps.org/doi/10.1103/PhysRevC.93.024917>.
- [27] R. Scheibl and U. Heinz, “Coalescence and flow in ultrarelativistic heavy ion collisions,” *Phys. Rev. C* **59** (Mar, 1999) 1585–1602.  
<https://link.aps.org/doi/10.1103/PhysRevC.59.1585>.
- [28] S. Hornung, “Production of (anti-) $^3\text{He}$  and (anti-) $^3\text{H}$  in p-Pb collisions at  $\sqrt{s_{\text{NN}}} = 5.02$  TeV measured with ALICE at the LHC.” [https://www.physi.uni-heidelberg.de/Publications/PhDThesis\\_SebastianHornung.pdf](https://www.physi.uni-heidelberg.de/Publications/PhDThesis_SebastianHornung.pdf).

- Dissertation, Universität Heidelberg (2020).
- [29] **ALICE** Collaboration, “Hypertriton Production in p-Pb Collisions at  $\sqrt{s_{NN}}=5.02$  TeV,” *Phys. Rev. Lett.* **128** no. 25, (2022) 252003, [arXiv:2107.10627 \[nucl-ex\]](#).
- [30] **ALICE** Collaboration, “The ALICE experiment at the CERN LHC,” *JINST* **3** (2008) S08002.
- [31] **ALICE** Collaboration, “Performance of the ALICE Experiment at the CERN LHC,” *Int. J. Mod. Phys. A* **29** (2014) 1430044, [arXiv:1402.4476 \[nucl-ex\]](#).
- [32] **ALICE** Collaboration, “Internal ALICE figure repository.”  
<https://alice-collaboration.web.cern.ch>.
- [33] **ALICE** Collaboration, S. Acharya *et al.*, “Production of light-flavor hadrons in pp collisions at  $\sqrt{s} = 7$  and  $\sqrt{s} = 13$  TeV,” *Eur. Phys. J. C* **81** no. 3, (2021) 256, [arXiv:2005.11120 \[nucl-ex\]](#).
- [34] J. de Cuveland, “Entwicklung der globalen Spurrekonstruktionseinheit für den ALICE-Übergangsstrahlungsdetektor am LHC (CERN).”  
<http://www.kip.uni-heidelberg.de/Veroeffentlichungen/download.php/4391/ps/cuveland-dipl.pdf>. Diplomarbeit, Universität Heidelberg (2003).
- [35] “Trigger and timing control.” <http://ttc.web.cern.ch/TTC/intro.html>.
- [36] J. Klein, “Commissioning of and Preparations for Physics with the Transition Radiation Detector in A Large Ion Collider Experiment at CERN.”  
[https://www.physi.uni-heidelberg.de/Publications/dipl\\_klein.pdf](https://www.physi.uni-heidelberg.de/Publications/dipl_klein.pdf).  
Diploma thesis, Universität Heidelberg (2008).
- [37] S. Kirsch, “Development of the Supermodule Unit for the ALICE Transition Radiation Detector at the LHC (CERN).” <http://www.kip.uni-heidelberg.de/Veroeffentlichungen/download.php/4654/ps/1818.pdf>. Diploma thesis, Universität Heidelberg (2007).
- [38] F. Rettig, “Entwicklung der optischen Ausleseketten für den ALICE-Übergangsstrahlungsdetektor am LHC (CERN).”

- <http://www-alice.gsi.de/trd/papers/dipl/rettig.pdf>. Diploma thesis, Universität Heidelberg (2007).
- [39] U. Westerhoff, “Development of the Electron Trigger System for the Transition Radiation Detector of the ALICE Experiment.”. Dissertation, Universität Münster (in progress).
- [40] J. de Cuveland, “A Track Reconstructing Low-latency Trigger Processor for High-energy Physics. Ein Triggerprozessor zur Online-Spurrekonstruktion in der Hochenergiephysik,” 2009. <https://cds.cern.ch/record/1295509>. Dissertation, Universität Heidelberg (2009).
- [41] M. Winn, “Inclusive  $J/\Psi$  production at mid-rapidity in p–Pb collisions at  $\sqrt{s_{NN}} = 5.02$  TeV.” [https://archiv.ub.uni-heidelberg.de/volltextserver/20656/1/pPb\\_jpsitoe\\_master\\_doctorthesis.pdf](https://archiv.ub.uni-heidelberg.de/volltextserver/20656/1/pPb_jpsitoe_master_doctorthesis.pdf). Dissertation, Universität Heidelberg (2016).
- [42] ALICE Collaboration, “The ALICE Offline Bible.” <https://alice-offline.web.cern.ch/sites/default/files/uploads/OfflineBible.pdf>.
- [43] S. Kirsch, “Multi-Event Buffering and Local Event Building for the ALICE Transition Radiation Detector.” [https://publikationen.ub.uni-frankfurt.de/opus4/frontdoor/deliver/index/docId/46472/file/Dissertation\\_Kirsch.pdf](https://publikationen.ub.uni-frankfurt.de/opus4/frontdoor/deliver/index/docId/46472/file/Dissertation_Kirsch.pdf). Dissertation, Universität Frankfurt (2017).
- [44] J. d. C. Prof. Dr. Volker Lindenstruth, “ALICE TRD Global Tracking Unit (GTU).” [https://www.physi.uni-heidelberg.de/~minjung/alicegirl4safe/alice\\_trd/gtu.pdf](https://www.physi.uni-heidelberg.de/~minjung/alicegirl4safe/alice_trd/gtu.pdf). EDR Documentation, Universität Heidelberg (2004).
- [45] J. Klein, “Jet Physics with A Large Ion Collider Experiment at the Large Hadron Collider.” <https://cds.cern.ch/record/1973326/files/CERN-THESIS-2014-186.pdf>. Dissertation, Universität Heidelberg (2014).
- [46] M. Hartung, “Pining down the (anti-)hypertriton production with ALICE at the LHC.” <https://www.dpg-verhandlungen.de/year/2023/conference/sruk/>

- part/hk/session/16/contribution/2, 2023.
- [47] J. F. Mücke, “Triggering on light nuclei with the ALICE Transition Radiation Detector.” [https://www.physi.uni-heidelberg.de/Publications/BscThesis\\_Felix\\_Muecke.pdf](https://www.physi.uni-heidelberg.de/Publications/BscThesis_Felix_Muecke.pdf). Bachelor thesis, Universität Heidelberg (2012).
- [48] N. A. Löher, “Search for exotic bound states and measurement of the (anti-)alpha production yield in Pb–Pb collisions with ALICE at the LHC (CERN).” [https://tuprints.ulb.tu-darmstadt.de/8222/7/Dissertation\\_Nicole\\_Loeher\\_genehmigt.pdf](https://tuprints.ulb.tu-darmstadt.de/8222/7/Dissertation_Nicole_Loeher_genehmigt.pdf). Dissertation, Universität Darmstadt (2018).
- [49] **ALICE** Collaboration, G. Willems, “Upgrade of the ALICE Readout and Trigger System.” [https://indico.cern.ch/event/735597/attachments/1669129/2676937/TDR\\_TDR\\_section\\_updated\\_June\\_2017.pdf](https://indico.cern.ch/event/735597/attachments/1669129/2676937/TDR_TDR_section_updated_June_2017.pdf). CERN-LHCC-2013-019, ALICE-TDR-015 (Internally revised version).
- [50] **ALICE** Collaboration, “Upgrade of the Readout and Trigger System, Technical Design Report.” <http://cds.cern.ch/record/1603472/files/ALICE-TDR-015.pdf>. CERN-LHCC-2013-019 ; ALICE-TDR-015 (2014).
- [51] P. Buncic, M. Krzewicki, and P. Vande Vyvre, “Technical Design Report for the Upgrade of the Online-Offline Computing System,” tech. rep., 2015. <https://cds.cern.ch/record/2011297>.
- [52] T. Kiss, “CRU HW update, Production Status.” [https://indico.cern.ch/event/829366/contributions/3472378/attachments/1869562/3076024/CRU\\_HFW\\_Status\\_SUF\\_26062019.pdf](https://indico.cern.ch/event/829366/contributions/3472378/attachments/1869562/3076024/CRU_HFW_Status_SUF_26062019.pdf). ALICE internal presentation (July 26th, 2019).
- [53] B. Bathen, “Aufbau eines Triggers für Tests der ALICE-TRD-Supermodule mit kosmischer Strahlung.” [https://www.uni-muenster.de/imperia/md/content/physik\\_kp/agwessels/thesis\\_db/ag\\_wessels/bathen\\_2007\\_diplom.pdf](https://www.uni-muenster.de/imperia/md/content/physik_kp/agwessels/thesis_db/ag_wessels/bathen_2007_diplom.pdf). Diploma thesis, Universität Münster (2007).
- [54] Analog Devices (formerly Linear Technology), “LTC5100 3.3V, 3.2Gbps VCSEL

- Driver, datasheet.” <https://www.analog.com/media/en/technical-documentation/data-sheets/5100fs.pdf>.
- [55] Texas Instruments, “Tlk2501, 1.5 to 2.5 gbps transceiver, datasheet.” [https://www.ti.com/lit/ds/symlink/tlk2501.pdf?ts=1662477964449&ref\\_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FTLK2501](https://www.ti.com/lit/ds/symlink/tlk2501.pdf?ts=1662477964449&ref_url=https%253A%252F%252Fwww.ti.com%252Fproduct%252FTLK2501).
- [56] Broadcom (formerly Avago Technologies), “AFBR-821vx3Z, Minipod, datasheet.” [https://docs.broadcom.com/doc/AV02-4467EN\\_PB\\_AFBR-811Xx3Z\\_2014-03-06](https://docs.broadcom.com/doc/AV02-4467EN_PB_AFBR-811Xx3Z_2014-03-06).
- [57] Olivier Bourrion et al., “cru-fw.” <https://gitlab.cern.ch/alice-cru/cru-fw>. Gitlab repository, CERN Gitlab (2021).
- [58] Filippo Costa et al., “wp6-doc.” <https://gitlab.cern.ch/Alice02Group/wp6-doc>. Gitlab repository, CERN Gitlab (2021).
- [59] O. Schmidt, “Private communication with Ole Schmidt.”
- [60] D. Dobrigkeit Chinellato, “ALICE event displays and geometry rendering of lead-lead ion collisions during Run 3,”. <https://cds.cern.ch/record/2841865>. ALICE-PHO-GEN-2022-009 (2022).
- [61] O. Bourrion, D. Evans, J. Imrek, A. Jusko, A. Kluge, M. Krivda, J. Kvapil, R. Lietava, L. A. Pérez Moreno, O. Villalobos-Baillie, E. Willsher, “Interface between CTS-CRU and CTS-Detector Front Ends, Trigger Notes for Developers.” [https://twiki.cern.ch/twiki/pub/ALICE/NoteForFEDevelopers/CTS\\_CRU\\_FE\\_interface.pdf](https://twiki.cern.ch/twiki/pub/ALICE/NoteForFEDevelopers/CTS_CRU_FE_interface.pdf), 2019.
- [62] “58928 - 2013.3 FIFO Generator v11.0 - Create deeper or wider FIFOs using smaller size FIFOs).” [https://support.xilinx.com/s/article/58928?language=en\\_US](https://support.xilinx.com/s/article/58928?language=en_US).
- [63] **Particle Data Group** Collaboration, “Review of Particle Physics,” *PTEP* **2022** (2022) 083C01.



# Acknowledgement

I am very grateful to everybody who has supported me accomplishing this thesis.

Especially, I would like to thank Prof. Dr. Johannes Wessels for the opportunity to work on this dissertation and for his enduring support. I also thank him for the opportunity of working at CERN. In addition, I thank Prof. Dr. Christian Klein-Bösing and Prof. Dr. Anton Andronic for their great support in all matters and under all circumstances. I thank Prof. Dr. Johanna Stachel for her trust and continuous support of my project. Furthermore, I am very grateful for the financial support throughout this work.

My CERN supervisor Dr. Filippo Costa has supported me greatly and I thank him for many intensive discussions, which have made this work a success.

I would like to thank Dr. Leticia Cunqueiro Mendez for all encouragements and support. I am grateful to Dr. Jorge Mercado Perez for his great help and encouragements. I would like to thank Dr. Benjamin Dönigus for his intensive help, his great support and encouragements. I thank Dr. Jochen Klein and Dr. Stefan Kirsch for their frequent help, their support and many productive discussions.

It has been a pleasure to work with the entire TRD team and I thank for the help and many important technical discussions. Especially, I would like to mention Dr. Venelin Angelov, Dr. Anastasia Berdnikova, Dr. Oliver Busch, Prof. Dr. Tom Dietel, Sean Murray, Dr. Yvonne Pachmayer and Dr. Ole Schmidt.

The CTP team has provided a lot of support and technical help and I would like to thank especially Dr. Anton Jusko, Dr. Marian Krivda, Dr. Jakub Kvapil and Dr. Roman Lietava.

I thank Dr. Chilo Garabatos Cuadrado for his intensive help and support with the TRD gas system.

Furthermore, I would like to thank the professors of the Institut für Kernphysik at the Universität Münster for their support and the interesting discussions, especially I would like to mention Prof. Dr. Dieter Frekers, Prof. Dr. Alexander Kappes, Prof. Dr. Alfons Khoukaz and Prof. Dr. Christian Weinheimer.

I am grateful for the many productive discussions with members of the ALICE collaboration and their help. In particular, I would like to mention Klaus Barth, Prof. Dr. Alexander Borissov, Prof. Dr. Peter Braun-Munzinger, Benjamin Brudnyj, Roberto Divia, Ömür Erkiner, Dr. Siegfried Förtsch, Dr. Martino Gagliardi, Dr. Alexander Kalweit, Tivadar Kiss, Dr. Alex Kluge, Dr. Robert Munzer, Dr. Maximiliano Puccio, Dr. Giuseppe Simonetti, Csaba Soos and Arturo Tauro.

I would like to thank the entire working group at the Institut für Kernphysik at the Universität Münster for the support, great help, interesting discussions and the fantastic work atmosphere. Especially, I would like to mention Dr. Bastian Bathen, Dr. Cruz de Jesús García Chávez, Dr. Katharina Demmich, Tabea Eder, Dr. Linus Feldkamp, Felix Fidorra, Dr. Rüdiger Haake, Ailec de la Caridad Bell Hechavarria, Florian Hermann, Fabian Hordt, Helena Horstmann, Florian Jonas, Philipp Kähler, Martin Kohn, Carolina Lopes, Dr. Jens Lühder, Adrian Meyer-Ahrens, Dr. Daniel Mühlheim, Phillip Munkes, Dr. Annika Passfeld, Dr. Hendrik Poppenborg, Dr. Lisa Schlüter, Dr. Lucia Tarasovicova, Ruben Weber and Uwe Westerhoff.



I thank Mrs. Uta Mayer for her competent help in all administrative matters. In addition, I thank Dr. Michael Hauschild for his assistance in the context of the Wolfgang-Gentner stipend.

Finally, I would like to thank my parents Ute Willems and Prof. Dr. Herbert Willems as well as my family for their support. I am very grateful to my wife Dr. Denise Moreira de Godoy Willems for her great support during the work for this thesis and far beyond.



# Guido Alexander Willems

Universität Münster  
Institut für Kernphysik  
Wilhelm-Klemm-Str. 9, 48149 Münster  
E-mail: guidowillems@uni-muenster.de

## Education

Since September 2015: **PhD Student in Physics**, Universität Münster, Germany  
Thesis: “Hardware Development for the Upgrade of the Transition Radiation Detector and Development of a Nuclei Trigger”  
Advisors: Prof. Dr. Johannes Wessels, Prof. Dr. Christian Klein-Bösing

August 2012 – January 2015: **Master of Science**, Universität Münster, Germany  
Thesis: “Software development for an automatic scanning microscope and its application for the MOEDAL experiment at LHC”  
Advisor: Prof. Dr. Dieter Frekers  
Overall average grade of the master degree: 1.0 with distinction

August 2012 – April 2013: **Erasmus Student**, Université Paris XI, France

October 2009 – August 2012: **Bachelor in Physics**, Universität Münster, Germany  
Thesis: “PLC logic control of an automatic scanning microscope - Hardware and basic programming concept”  
Advisor: Prof. Dr. Dieter Frekers  
Overall average grade of the bachelor degree: 1.3

June 2009: **Abitur**, Friedensschule Münster

Overall average grade of the Abitur: 1.0

## Languages

German (native language), English (fluent), French (fluent), Mandarin (elementary), Portuguese (elementary)

## Awards

July 2009 – December 2014: e-fellows.net stipend because of excellent achievements

September 2009: Distinction by the prime minister of North Rhine-Westphalia as an excellent high-school graduate

June 2009: Best high-school graduate of the year in chemistry by the Gesellschaft Deutscher Chemiker

June 2009: Excellent achievements by the Deutsche Physikalische Gesellschaft

February 2007: 3rd prize in the competition “Jugend forscht” with the topic “Power output measurements of a biological fuel cell”

## Conference talks

CRU status

TRD Comprehensive Status and Planning Meeting

September 2020, online

Trigger scheme, Readout scheme, Run 3 data format CRU status

TRD commissioning meeting

November 2019, Geneva, Switzerland

CRU status

TRD Comprehensive Status and Planning Meeting

July 2019, Frankfurt, Germany

Production of light (anti-)nuclei in pp collisions with ALICE at the LHC

Arbeitsreffen Kernphysik 2019

February 2019, Schleching, Germany

Status report of the ALICE Transition Radiation Detector (TRD)

ALICE week

July 2018, Geneva, Switzerland

Upgrade of the ALICE Transition Radiation Detector for Run 3

13th Gentner Day

April 2018, Geneva, Switzerland

Implementation of the ALICE TRD triggers

Deutsche Physikalische Gesellschaft (DPG) spring conference, Ruhr-Universität Bochum

February 2018, Bochum, Germany

Upgrade of the ALICE Transition Radiation Detector for Run 3

Hausseminar of the Institut für Kernphysik of the Universität Münster

January 2018, Münster, Germany

Triggering on  $Z = 2$  particles with the ALICE TRD

ALICE FSP meeting

September 2016, Altensteig-Wart, Germany

The Münster scanner and etch-pit recognition

5th MoEDAL Collaboration meeting

June 2016, Valencia, Spain

## Publications with significant contribution

ALICE Collaboration, “ALICE upgrades during the LHC Long Shutdown 2”,  
arXiv:2302.01238 [physics.ins-det].

ALICE Collaboration, “The ALICE Transition Radiation Detector: Construction,  
operation, and performance”, Nuclear Instruments and Methods in Physics Research  
Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 881 (2018)  
88–127.

ALICE Collaboration, G. Willems, “Upgrade of the ALICE Readout and Trigger System”,  
CERN-LHCC-2013-019, ALICE-TDR-015 (Internally revised version).