Design and development of the ALICE common readout unit user-logic firmware for the Muon Identifier readout chain

D O Thys-dingou[1]*, A Raji[1], E Z Buthelezi[2,3] and S V Förtsch[2]

[1] Department of Electrical, Electronic and Computer Engineering, Cape Peninsula University of Technology, Bellville South Industrial, Cape Town, 7530, South Africa
[2] Department of SSC lab, iThemba LABS, Faure, Cape Town, 7100, South Africa
[3] School of Physics, University of the Witwatersrand, Johannesburg, 2000, South Africa

E-mail: dieuveil.orcel.thys-dingou@cern.ch

Abstract. A Large Ion Collider Experiment (ALICE) at the Large Hadron Collider (LHC) at CERN went through a major upgrade during which some of its sub-detectors were replaced with new ones, while others are equipped with new electronics to handle the expected higher collision rates in the current running period (Run 3), which started in 2022. As part of the upgrade, certain sub-detectors such as the Muon Trigger (MTR), renamed to Muon IDentifier (MID), now operate in a continuous, trigger-less readout mode, in addition to the previous triggered readout mode. Due to the increased quantity of data, typical methodologies are impossible to employ without massive efforts to expand the processing capacity. Since the new ALICE computing system cannot keep up with the increased data flow of the MID, a new processing algorithm has been established. This paper provides an insight to the new approach of processing the MID readout data based on a customized user-logic FPGA firmware.

1. Introduction

ALICE [1] is one of the four main experiments at the LHC. It is designed to study the strongly interacting matter, namely the quark gluon plasma (QGP) [2] and its properties. In order to unravel the enigma of the universe, the ALICE detector records data during lead-lead (Pb-Pb), proton-lead (p-Pb), and proton-proton (pp) collisions for a complete study. Based on data collected during Run 1 and 2, ALICE is the leading heavy-ion experiment in the world and is quickly expanding the knowledge gathered in previous experiments all over the world. The LHC completed the three-years planned second Long Shutdown, which started at the end of 2018 to prepare for Run 3. In line with the LHC upgrade, the ALICE detector had a major upgrade [3] . This upgrade addresses the challenge of reading out Pb-Pb collisions at a rate of 50 kHz, pp, and p-Pb at 200 kHz and higher. At the center of the ALICE upgrade strategy, is a high-speed readout approach based on a Common Readout Unit (CRU), which has been developed for detector data readout, reconstruction, multiplexing, and data decoding on the Online-Offline (O^2) computing system.

*Corresponding author

Many of the proposed physics observables require a change in the data-taking strategy, moving away from triggering a small subset of events to continuous online processing and recording of all events. To achieve these goals, ALICE has been upgraded in such a way that all interactions will be scrutinized with precision. The upgrade entailed the replacement of some sub-detectors with new ones, making use of new technologies, while others are now equipped with new front-end and readout electronic systems. Thus far, the selection of single muon and di-muon events with a maximum trigger rate of 1 kHz, limited by readout capabilities, was provided by the MTR, as well as muon identification. However, the upgrade strategy described in the letter of intent [3] does not require a muon trigger since all events of interest are now read out upon the interaction trigger before online selections. For this reason, as part of the upgrade, the new MID sub-detector plays the role of muon identifier.

2. Muon IDentifier

The MID [4] is based on 72 single-gap Resistive Plate Chamber (RPC) detectors, arranged in 2 stations of 2 chambers, each at a distance of about 16 m and 17 m from the interaction point, respectively. Its readout chain block diagram is shown in Fig. 1.

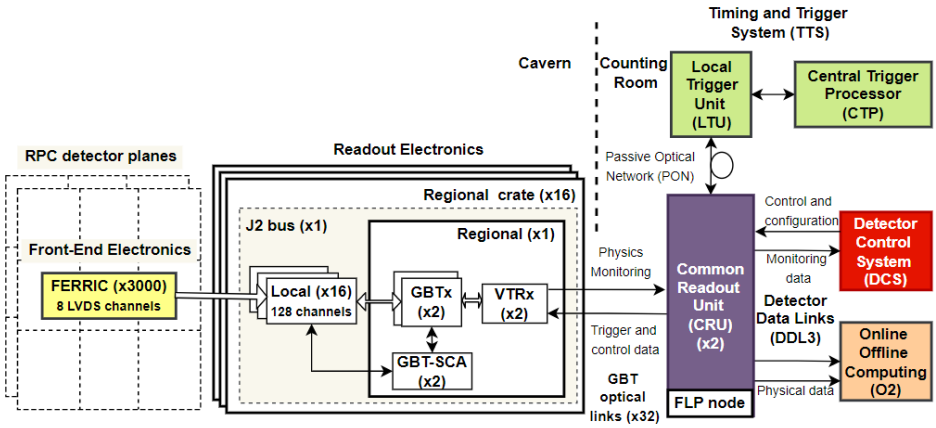

Figure 1. A schematic description of the MID readout chain architecture for LHC Run 3.

The MID readout chain consists of about 21 000 strips connected to the 72 RPC detectors spread over multiple Front-End Electronic Rapid Integrated Circuit (FEERIC) cards equipped with one or two customized Application-Specific Integrated Circuits (ASICs). The strip signals from the FEERICs are propagated to the readout electronics using high-speed Low-voltage Differential Signaling (LVDS) channels. The readout electronics (Local and Regional cards) act as interface between the on and off-detector electronics. They are mounted inside the cavern a little further away from the detector stations, where the radiation is lower. Since the colliding beams will produce a lot of radiation in the area around the ALICE detector in the cavern, the Regional cards are equipped with a radiation hardening Gigabit Transceiver set of chips (GBTx and GBT-SCA) [5, 6] used to facilitate the bidirectional connections between the readout electronics and CRUs through optical links, namely, GBT links. The CRUs are the key components of the chain. They combine and multiplex data from multiple readout electronic cards as well as timing and trigger information generated from the Central Trigger Processor (CTP) via the Local Trigger Unit (LTU) before transmitting the data to the O^2 computing facility for processing and storage. The CRUs are mounted on computers housed in the intermediary computer room, called the counting room, tens of meters above the ALICE

cavern and thus do not require radiation hardening, as is the case for the readout electronics. These computers can be reached over the network from the main Detector Control System (DCS). The DCS manages the readout chain by sending commands and monitoring the system. Experimental data are moved from the First Level Processor (FLP) to the Event Processing Node (EPN) for processing and storage. The EPN is an internal component of the O^2 computing system.

3. CRU firmware

The standard approach of delivering raw data to the O^2 system is no longer sufficient to meet the needs of the newly enhanced MID sub-detector. As a result, an alternative option was presented to the ALICE collaboration [7]. Since the release of its first official version in 2018, the CRU firmware provides basic functionalities allowing to interact with multiple systems via a variety of interfaces and read out any sub-detector without conducting any first stage data analysis. However, the CRU firmware can be customised to meet the demands of different sub-detectors in the ALICE experiment. In particular, a first stage data analysis can be implemented before online and offline reconstruction. This customisation is referred to as "User Logic". It is the responsibility of the MID specialists to decide how data should be handled in their user logic.

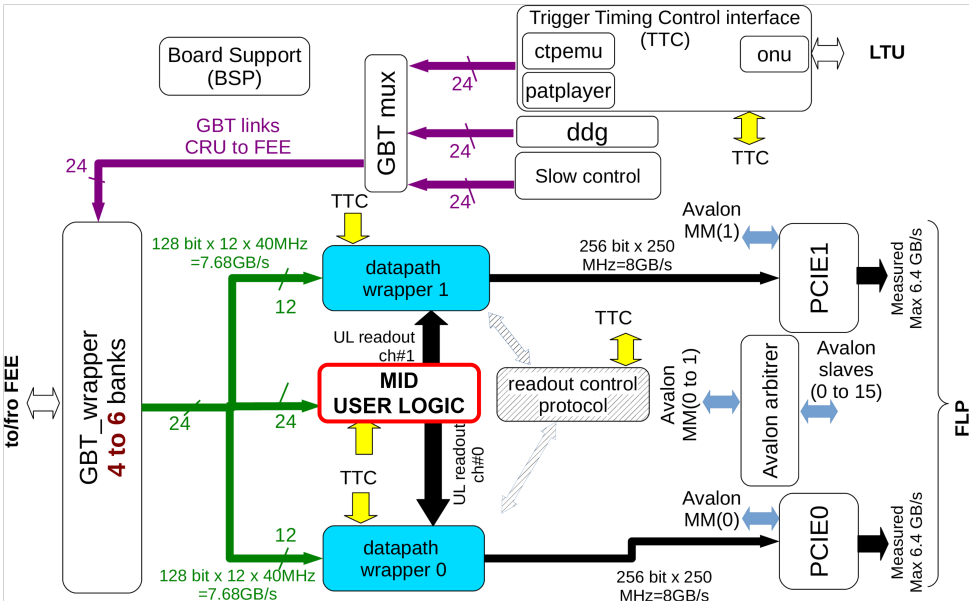

Figure 2. CRU firmware architecture. Adapted from [8].

The CRU firmware architecture is illustrated in Fig.2. It is composed of several modules interacting with several interfaces. From left to right, the main interfaces are the GBT wrappers, Board Support Package (BSP), Datapath Wrappers (DWs), Timing and Trigger Control (TTC), Dedicated Data Generator (DDG), slow control, and PCIe endpoints. All of these interfaces provide indispensable functionalities to the CRU firmware, and at the heart of it, is a tailored user logic component with functionalities that are unique to MID.

The key general requirements of the user logic component is derived from several documents and are cited in this paper, but not duplicated. These requirements are described in [8, 9] and [10]. One of the main functional requirements of the user-logic component is to reduce the data rates transmitted to the O^2 system by performing zero-suppression. It also has to carefully handle the readout electronics anomalies and errors, without stopping the data acquisition of the readout chain. Transient losses of payload (physics) data is not tolerated. This means that

at all cost, every level of the user logic hierarchy must be able to identify anomalies from the levels below and continue to transmit packets and follow the communication rules with the levels above. The violation of communication protocols, corruption of data, and unresponsiveness of the system in case of errors shall be prevented.

4. Architecture and design of the user-logic component

The user logic is designed in a sequential manner, with all processing occurring one after another from an input-output perspective. The objective of this approach is to facilitate error tracking. The user logic consists of three main segments, each of which is linked to a specific interface of the CRU firmware (See Fig. 2). Starting from the top is the TTC segment (Grey), which receives data from the timing and trigger system through the TTC interface. Next is the GBT segment (Blue), which receives data from the readout electronics via the GBT wrappers, analyses it then combines it with the Raw Data Header (RDH) extracted from timing and trigger information before transmitting them to the O^2 system via the datapath wrappers. The GBT segment is the only part of the design that can be duplicated through parameterization. Hence, enabling the possibility to process multiple GBT links, allows for improvement and adaption to diverse testing scenarios. The last segment is the Avalon (Orange), which provides configuration and monitoring through the PCIe interface. A representation of the user-logic block diagram and its interfaces is shown in Fig. 3.

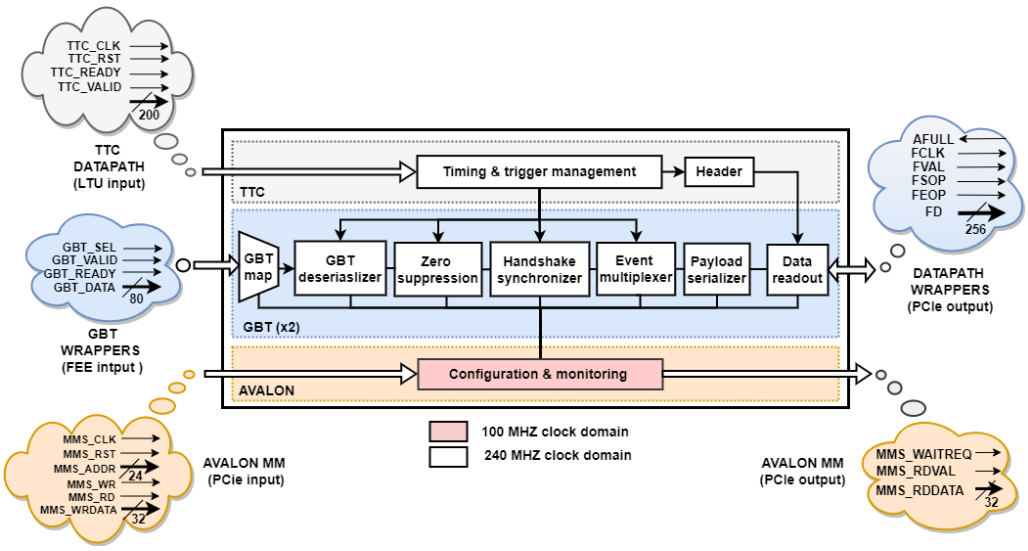

Figure 3. Structure of the user-logic design showing the three main segments and data flow.

5. Test-bench layout

A test bench at iThemba LABS was implemented to extend the capabilities of the user logic and match the test scenarios performed with readout electronics at SUBATECH, Nantes (where the readout electronics cards were designed [11]). The test bench is a scaled-down replica of the MID readout chain without the RPC detectors. It includes a fully-equipped VME crate (16 x Local, 1 x Regional, and 1 x J2 bus boards), the LTU, CRU, and FLP. Fig. 4 depicts the test bench setup and illustrates how various components are linked together. The full setup can be observed on the top left, while in the bottom left the fully-equipped VME crate is shown. The Local and Regional cards are plugged into the crate via the J2 bus card sitting at the back of the crate. Three cables are exiting the regional card, two of which are optical cables, and

connect the Regional card to the CRU. The latter is a USB (2.0) cable connected to a CentOS PC, which is used to configure and program the Local and Regional FPGAs. At the top-right is the LTU, which can be used to interact with multiple CRUs via a splitter. However, for this application only one CRU is needed. The connection between the LTU and the CRU is done via a single-mode SC to SC optical cable. The LTU uses an Ethernet cable to interact with the FLP software, which runs on CentOS 8. Finally, on the bottom-right is the CRU board housed by the FLP server. The CRU board is internally attached to the FLP via the PCIe connectors, its FPGA can be programmed using the PCIe interface or via its integrated USB blaster programmer, which connects to the FLP server using a micro-USB (2.0) cable.

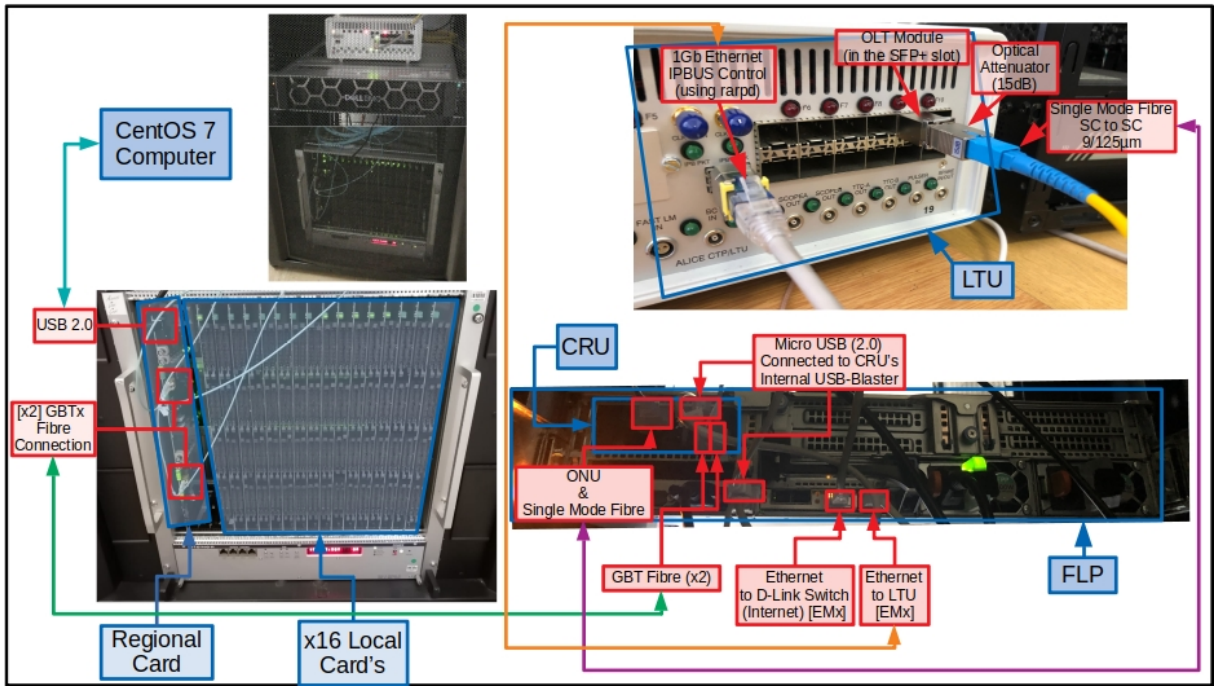

Figure 4. New test bench located in the laboratory S64 at iThemba LABS. The diagram shows the readout chain components and how they are connected.

6. Results

The simulated results of the working user-logic firmware were generated in ModelSim using a simulation framework provided by SUBATECH. A subset of the results is illustrated in Fig. 5. As can be observed, the user logic output packets fulfill the readout control protocol criteria specified in [10], which states that packets created by the user-logic must begin and terminate with the Start Of Packet (SOP) and End Of Packet (EOP) signals, and each packet must be enclosed by the RDH. The output data format transmitted by the user logic is based on the GBT raw data, pre-analysed and concatenated into multiple data blocks of 256-bit to form the payload included in the packets.

Figure 5 demonstrates that the user logic is operational and the simulation tests were successful. The hardware tests result did not reflect the simulation results at first, but were refined after each iteration until complete accuracy was achieved. The user logic helped decreasing the amount of data transmitted from the CRU to the FLP by roughly 80%, and improved the readability of the data at the O^2 level.

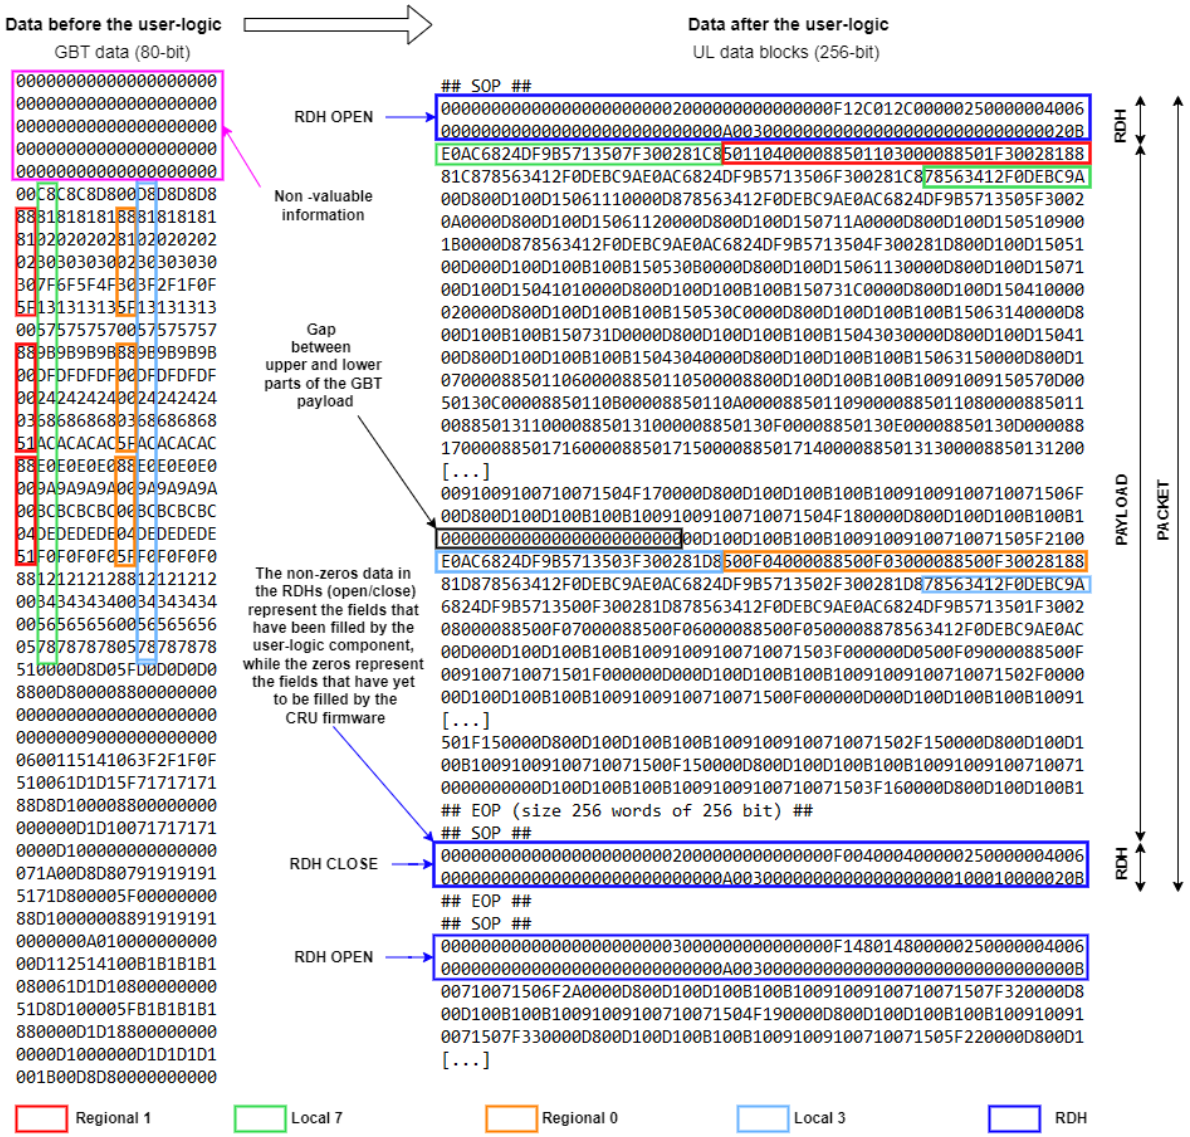

Figure 5. Simulation results extracted from Modelsim. This image compares data before and after the user logic processing. As it can be seen, the packets start and end with SOP and EOP signals. Each packet includes RDHs (open and close), that specify the borders of the packet.

7. FPGA resource usage

The CRU firmware combined with the user-logic component use about 160k (38%) Adaptive Logic Modules (ALMs) and 1355 (50%) RAM blocks of the available resources. These results were obtained after integrating and compiling the CRU firmware with the user logic component. Table 1 provides a summary of the total FPGA resource used. The user logic consumes around 37k (9%) ALMs and 271 (10%) RAM blocks of the overall resources. These findings meet the requirement of this study, but are good enough as the long-term aim is to process data from 16 GBT links while maintaining the overall RAM consumption below 75%. The user logic is currently being optimized and will soon be available before the start of LHC Run 3.

Table 1. FPGA resource usage of the CRU firmware after insertion of the user logic component.

Resource name	Total in ratio	Total in percentage
`Logic utilization (in ALMs)`	160,282 / 427,200	38%
`Pin`	369 / 960	38%
`Block memory bits`	19,982,660 / 55,562,240	36%
`RAM blocks`	1,355 / 2,713	50%
`RX channels`	41 / 72	57%
`TX channels`	41 / 72	57%
`Phase Locked Loops (PLLs)`	59 / 144	41%

8. Conclusion

The user-logic firmware performed reasonably well during the simulation and hardware tests. The findings shown in this paper demonstrate that the user logic is stable, reliable, and built to read out Pb-Pb collisions at a rate of 50 kHz, pp, and p-Pb at 200 kHz and higher without issues. The methodology implemented shows that it is feasible to considerably reduce the data bandwidth transmitted from the CRU to the FLP by roughly 80%. The user logic has passed the validation tests and so fulfils the MID requirements. The results obtained also indicate that with some optimizations, the user logic will be ready to process data from the complete MID system before the start of LHC Run 3.

References

[1] ALICE Collaboration 2008, The ALICE experiment at the CERN LHC, *JInst*, **3.08**: S08002."

[2] Graczykowski L. 2016, Soft QGP probes with ALICE, *Acta Phys.Polon.B* **vol 47**, pp. 1473-1487, 1608.00040.

[3] ALICE Collaboration 2014, Upgrade of the ALICE Experiment: Letter of Intent, *Journal of Physics G: Nuclear and Particle Physics*. DOI:10.1088/0954-3899/41/8/087001

[4] Terlizzi L. 2020, The ALICE Muon IDentifier (MID), *Journal of Instrumentation*, **15** C10031.

[5] Moreira P. et al. 2009, The GBT Project, *Proceedings of the TWEPP*. DOI:10.5170/CERN-2009-006.342

[6] Caratelli A. et al. 2015, The GBT-SCA: A radiation tolerant ASIC for detector control and monitoring applications in HEP experiment, *JInst*, **10** C03034.

[7] David E. et al. 2015, ALICE COMMON READ-OUT UNIT, *ALICE internal report*, (Geneva: CERN). https://www.phenix.bnl.gov/WWW/publish/mxliu/sPHENIX/Readout/CRU/CRU_Specification_v0.7.pdf

[8] Bourrion O., Bouvier J. et al. 2021, Versatile firmware for the Common Readout Unit (CRU) of the ALICE experiment at the LHC, *JInst*. DOI:10.1088/1748-0221/16/05/P05019

[9] D. Evans. et al. 2016, Trigger System Design Review, *ALICE internal report*, (University of Birmingham). https://indico.cern.ch/event/533412/contributions/2172791/attachments/1291065/1922927/CTPLTU12.pdf

[10] Bourrion O. et al., CRU firmware gitlab project. https://gitlab.cern.ch/alice-cru/cru-fw

[11] C. Renard, MID Readout electronics project. http://www-subatech.in2p3.fr/~electro/projets/alice/dimuon/trigger/upgrade/index.html