

RCU2 — The ALICE TPC readout electronics consolidation for Run2

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2013 JINST 8 C12032

(<http://iopscience.iop.org/1748-0221/8/12/C12032>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 137.138.125.163

This content was downloaded on 08/07/2014 at 07:20

Please note that [terms and conditions apply](#).

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS 2013,
23–27 SEPTEMBER 2013,
PERUGIA, ITALY

RCU2 — The ALICE TPC readout electronics consolidation for Run2

J. Alme,^{a,1} T. Alt,^j L. Bratrud,^f P. Christiansen,ⁱ F. Costa,^e E. David,^h T. Gunji,^k
T. Kiss,^h R. Langøy,^f J. Lien,^f C. Lippmann,^d A. Oskarsson,ⁱ A. Ur Rehman,^g
K. Røed,^b D. Röhrich,^c A. Tarantola,^j C. Torgersen,^c I. Nikolai Torsvik,^c K. Ullaland,^c
A. Velure,^c S. Yang,^c C. Zhao,^b H. Appelshäuser^j and L. Österman^{i,2}

^aBergen University College, P.O. Box 7030, NO-5020 Bergen, Norway

^bUniversity of Oslo, P.O. Box 1048, Blindern, NO-0316 Oslo, Norway

^cUniversity of Bergen, P.O. Box 7800, NO-5020 Bergen, Norway

^dGSI Helmholtzzentrum für Schwerionenforschung,

Planckstr. 1, D-64291 Darmstadt, Germany

^eCERN, CH-1211, Genève 23, Switzerland

^fVestfold University College, Postboks 2243, NO-3103 Tønsberg, Norway

^gCOMSATS Institute of Information Technology,

Park Road, Chak Shahzad, Islamabad, Pakistan

^hCerntech, Petzvál J. u. 44, H-1119 Budapest, Hungary

ⁱUniversity of Lund, Box 117, 221 00 LUND, Sweden

^jGoethe University Frankfurt,

Senckenberganlage 31, 60325 Frankfurt am Main, Germany

^kUniversity of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

E-mail: johan.alme@hib.no

ABSTRACT: This paper presents the solution for optimization of the ALICE TPC readout for running at full energy in the Run2 period after 2014. For the data taking with heavy ion beams an event readout rate of 400 Hz with a low dead time is envisaged for the ALICE central barrel detectors during these three years. A new component, the Readout Control Unit 2 (RCU2), is being designed to increase the present readout rate by a factor of up to 2.6. The immunity to radiation induced errors will also be significantly improved by the new design.

KEYWORDS: Electronic detector readout concepts (gas, liquid); Front-end electronics for detector readout; Data acquisition concepts; Digital electronic circuits

¹For the ALICE TPC collaboration.

²Corresponding author.



Contents

1	Introduction	1
2	Present TPC readout electronics	2
2.1	Motivation for the RCU2 upgrade	3
3	Alternative LS1 upgrade proposals	4
4	RCU2	4
4.1	GTL backplane design	5
4.2	RCU2 hardware design	5
4.3	Radiation monitor	6
5	RCU2 FPGA design	6
6	RCU2 DCS software	8
7	Radiation tolerance measures and irradiation campaigns	8
8	Outlook and conclusions	9

1 Introduction

The Large Hadron Collider (LHC) accelerates protons and lead ions close to the speed of light and collides them inside the four experimental areas along the beamline. One of the LHC experiments is A Large Ion Collider Experiment (ALICE) [1].

The ALICE detector comprises several sub-detectors, of which one is the Time Projection Chamber (TPC) [2]. The TPC detector is the main tracking detector of ALICE. The end plates on both sides of the TPC barrel are populated with 557 568 readout pads. The TPC readout electronics is located directly behind the two detector end plates.

Following a successful running period from November 2009 to January 2013 (Run1) [3], the LHC is currently shut down for maintenance and preparation for even higher energies and luminosities. This period, named Long Shutdown 1 (LS1), lasts end of 2014.

The next running period (Run2) is between 2015 and 2018. In this period a peak luminosity up to $4 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ can be expected for Pb-Pb collisions with centre-of-mass energy of 5.5 TeV per nucleon pair. As a comparison, in the heavy ion run in 2011 the luminosity reached $10^{26} \text{ cm}^{-2} \text{ s}^{-1}$ [3]. The envisaged luminosities for Run2 correspond to interaction rates from 8 to 30 kHz. To get maximum benefit from the delivered luminosities, the most attractive data taking scenario for ALICE foresees data taking rates of at least a factor two faster than what can be achieved with the current TPC readout electronics. For a trigger mix containing central, semi-central, minimum bias,

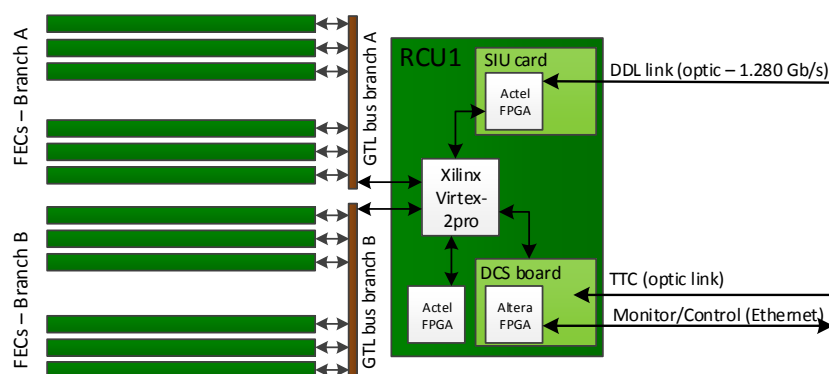


Figure 1. Sketch of the present TPC readout electronics with the RCU1.

calorimeter and Transition Radiation Detector (TRD) triggers, an event readout rate of at least 400 Hz is envisaged for the ALICE central barrel detectors, while keeping a reasonably low dead fraction (busy time). In comparison, the maximum achieved readout rate for the same trigger mix of the current TPC readout electronics has been measured to be about 320 Hz (in this case with high dead fraction). The higher luminosities for Run2 will also lead to a 40% increase in the event sizes since the number of tracks per interaction will be higher. The maximum event size for a central event in Run1 is 65 MB, for Run2 a corresponding event will produce 90 MB of data.

This paper presents the RCU2, which is the replacement of the present Readout Control Unit. It is currently being developed in order to provide a readout system capable to achieve the discussed performance requirements.

2 Present TPC readout electronics

A sketch of the present TPC readout electronics [2] is shown in figure 1, and consist of a Readout Control Unit (RCU) that connects to between 18 and 25 Front End Cards (FECs), depending on the radial position of the RCU in the TPC barrel. The radial direction is divided into 6 readout partitions with one RCU each. The connectivity between the RCU and the FECs is implemented using two branches of a multidrop, parallel Gunning Transistor Logic (GTL) bus with a bandwidth 1.6 Gbps. The FEC itself has 128 analog input channels, with each input channel corresponding to a single detector pad on the TPC detector.

The present RCU, from here on referred to as RCU1, consists of a motherboard with two attached mezzanine cards, the Source Interface Unit (SIU) and the Detector Control System (DCS) board. The SIU hosts one flash based FPGA and is the interface card implementing the Detector Data Link (DDL). The DDL is a custom protocol on a 1.280 Gbps bidirectional optical link connecting the RCU1 to the Data Acquisition (DAQ) system [2]. The TPC uses this interface for transmitting the event data for further online and offline analysis and for configuration of the Front End Electronics (FEE).

The DCS board host an SRAM based FPGA with an embedded ARM processor core [4]. It is configured with a small Embedded Linux platform, and communicates to the higher layer of the

DCS system [2] via a special transformerless Ethernet connection. In addition, it has an optical interface receiving the clock and trigger information from the Timing, Trigger and Control [5, 6].

The RCU1 motherboard hosts two FPGAs. The main FPGA, where the data readout functionality is implemented, is an SRAM based Xilinx Virtex2pro-vp7 [7], while the support FPGA is a Flash based Actel APA075 [8]. The purpose of the support FPGA is to program the main FPGA from an onboard Flash memory. It also implements Active Partial Reconfiguration to detect and correct Single Event Upsets (SEUs) [9] in the configuration memory of the main FPGA [10].

2.1 Motivation for the RCU2 upgrade

The present TPC readout electronics will be a limiting factor with the foreseen readout rate for Run2. In addition to this, stability issues have been observed during Run1, some of which can be traced to the SEUs in the SRAM based FPGAs. By itself, this is not necessarily a problem because of the implementation of Active Partial Reconfiguration. However, the FPGA resources are fully utilized and there is no room for design level mitigation. This implies that a single event upset (SEU) in a sensitive SRAM cell automatically leads to a functional failure.

Data rate limitations. The GTL bus between the RCU1 and the FECs is implemented as a 40 bit wide parallel bus where the RCU1 connects to two separate branches. Each of the two buses has a bandwidth of 1.6 Gbps and reads the data from up to 13 FECs, depending on the readout partition and branch. For high occupancy events like central Pb-Pb collisions the data readout through the GTL buses is the bottleneck of the readout system. The readout time is determined by the slowest readout partition, which is the second one from inside; Readout Partition 1 (RP1). The occupancy is highest on the inner partitions and RP1 has the largest number of FECs connected (25). Since the GTL bus is separated into two branches, here up to 13x128 channels must be read out sequentially. The readout time is characterized by an overhead needed for addressing of the channels and the 1.6 Gbps datarate of the GTL bus. It has been measured for pp and Pb-Pb collisions in 2010 [11]. For central Pb-Pb events the readout time reaches up to 4 ms depending on the number of tracks per event.

Stability issues. The current readout system is not radiation tolerant and suffers from radiation effects. In the 2013 p-Pb data taking at high interaction rates (up to 200 kHz) 9% of all data taking sessions (runs) had to be stopped (or were stopped automatically) due to errors that occurred in the TPC readout electronics. These errors, which lead to the readout getting stuck (busy) or to corrupted event headers, can be assigned to radiation effects where SEUs in the configuration memory of the FPGA is the dominant factor. The SEU sensitivity of the current RCU1 FPGA design has been characterized [12], and it was found that approximately 1% of the SEUs leads to a premature abort of a run. Based on the potential luminosity for Run2, and on the analysis of the SEU rate for the heavy ion runs in 2011, it can be expected that the ALICE data taking will stop every hour if no measures are taken [13].

Also, the DCS boards suffer from radiation related communication errors (DCS to RCU) and communication losses (Ethernet to DCS system). Even though uncritical for the data taking, the loss of monitoring capabilities has to be avoided.

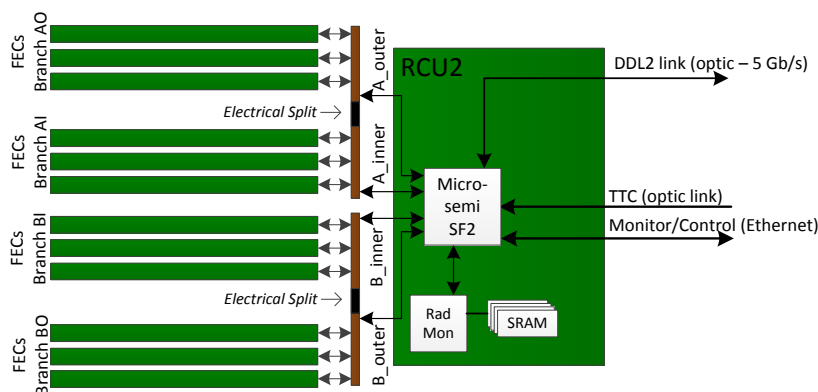


Figure 2. Sketch of RCU2 highlighting the main differences from RCU1.

Conclusion. The present TPC readout electronics will not be able to fulfill the ALICE ambitions for Run2 with the foreseen 40% increase in event size and the envisaged readout rate of 400 Hz. The stability problems are also expected to become worse with the present solution under the Run2 conditions (see section 7). With this in mind it is clear that an upgrade is needed.

3 Alternative LS1 upgrade proposals

Two upgrade options for LS1 have been discussed for the TPC electronics. The first option was to make a small add-on board, the Front End Card Interface (FECint) board, to be connected to the FEC with the purpose of translating the wide parallel bus into a high speed serial link. The initial idea of this design is given in [14]. A second prototype of the FECint board was designed in 2012, where two FECs are connected in parallel to one FECint [15].

This upgrade solution has two main advantages: (1) No bandwidth limitations are imposed by the readout electronics, and (2) it is relevant for the upgrades planned for Long Shutdown 2 (LS2) given the fact that it would use components and an infrastructure that would be reminiscent of the planned LS2 upgrade. However, this upgrade-plan was not prioritized as it was too ambitious given the time-budget available. Instead it was decided to develop the RCU2.

4 RCU2

The main motivation for the RCU2 is to develop a solution that gives the needed performance improvement, and at the same time is feasible within the limited time-frame of LS1. The RCU2 was initiated only in April 2013, giving less than 20 months to complete the project, including design, prototyping, mass-production and commissioning. This demanding time-schedule implies that there is no room for rearranging the infrastructure or architecture of the present TPC readout electronics. The cables for Ethernet, Trigger, DAQ and power must all be reused as is. Even if identified as the bottleneck of the present system, the GTL bus must also remain. However, each GTL bus is split into into four branches per RCU2 instead of two, see figure 2. This ensures at least a doubling of the data rate, which requires an upgraded DDL link as well. Using the same fiber, this is updated to the DDL2 protocol [16] with a bandwidth of 3.125 Gbps or 5 Gbps.

As shown in figure 2, the RCU2 hosts a Microsemi SmartFusion2 (SF2) System-on-Chip (SoC) FPGA M2S050-1FG896 [17, 18]. This is a state of the art flash-based FPGA that has SEU immune configuration memory, as well as several other radiation tolerance measures implemented [18] (see section 7). Since most of the the stability issues seen in Run1 can be traced back to SEUs in the configuration memory of the FPGA, these issues will be avoided just by the change of technology. It also comes with a Microcontroller Subsystem [19] which is based on a hardcore ARM Cortex-M3 microcontroller and many peripheral cores.

4.1 GTL backplane design

The shape of the TPC sector implies that different backplane designs are needed for the six readout partitions in one sector. For the RCU2, the existing two backplane branches A and B are further split forming four separate branches: A_inner, A_outer, B_inner and B_outer. This implies that each backplane is electrically split, giving two separate branches on the same PCB. It has been decided to keep the PCB formfactor of the original backplanes, as the position and angle of each FEC connector needs to be correct with sub millimeter precision. Even though the branches vary both in length and in the number of FECs connected, it is important that the connectors towards the RCU2 are located at fixed positions.

4.2 RCU2 hardware design

The formfactor and the placement of connectors and components on the RCU2 is largely decided by the wish to reuse the copper cooling envelopes from the RCU1 [2]. The water cooling provided by the cooling envelopes will be needed, since the GTL driver chips dissipate a lot of heat. A power consumption of about 22 W divided between the 4.3 V and 3.5 V power supplies have been estimated, where the equivalent number for the RCU1 is less than 10 W.

TTC interface. The data on the TTC link is transmitted using a 160 MHz biphas differential manchester encoded signal [20]. The data payload is split into channel A and B, where channel A is used for sending level 0 and level 1 trigger, while channel B is used for trigger information (e.g. level 2 triggers). On the RCU1, the clock and data recovery (CDR) and the channel splitting is done by the radiation tolerant TTC receiver chip (TTCrx) [20]. One important point of the TTCrx is that the clock is recovered with a predefined phase-offset to the LHC input clock, since it is used as the base of the sampling clock, and the phase of the sampling clock must be aligned for all of the 216 RCUs.

One of the main challenges during the design of the RCU2 is that the TTCrx is out of stock and no new production runs are foreseen. Solutions for replacing the TTCrx have been suggested [21], but none of them are intended for use in a radiation environment. Our solution is to use the Avago HFBR-2316TZ [22], of which an earlier version of the device was qualified by the TTC group at CERN [23] in 2003. A standard limiting amplifier follows the optical receiver to convert the signal to a digital LVDS signal, which is connected directly to user IOs on the FPGA where the CDR and channel splitting is implemented. Since no data exist on the limiting amplifier regarding radiation tolerance, and the data on the optical receiver is from a previous version, the components in the TTC chain will be tested for radiation tolerance early 2014 as part of the full system test mentioned in section 7.

DCS interface. The DCS interface is designed using a Marvell 88E1111 Ethernet PHY [24] connected to the Serial Deserialiser (SERDES) of the SF2. The advantages of this PHY are that there are no magnetic components inside, that it uses few IOs on the SF2, and that it gives a link speed of 100 Mbps opposed to 10 Mbps as on the DCS board. A design without transformers is needed because of the ALICE magnetic field. The solution for the analog part is taken over from the present DCS board design [25], which has proven to be quite reliable. There have been problems with communication losses on the Ethernet link, but these are not related to the analog design. Tests have been performed that shows that the Marvell PHY without transformer satisfies the requirements. Irradiation tests of the Marvell PHY will be performed early 2014 as part of the full system test mentioned in section 7. However, it is not regarded as a high risk component since a loss of communication is not critical to the data taking, and it is possible to reset the PHY from the SF2 if communication losses are detected.

DAQ interface. The DAQ interface is designed using an small formfactor pluggable (SFP) optical transceiver that is connected to the onchip SERDES on the SF2. The DDL2 protocol module is delivered by the ALICE DAQ group as a SF2 compatible IP core with the same front-end interface as on the RCU1. This makes the integration of the module easy.

4.3 Radiation monitor

The RCU1 includes a reconfiguration network comprised of a flash-based FPGA and a flash memory device, which is used for active partial reconfiguration of the main FPGA design [26]. Additionally, this is a powerful radiation monitor that has given valuable results concerning the impact of the radiation environment on the RCU1 [12, 27].

The reconfiguration network is made obsolete on the RCU2 with the introduction of the SF2 FPGA. However, the radiation monitoring functionality should be kept. A Microsemi ProASIC3 A3P250 [28] and four Cypress 8 Mbit SRAM memories [29] have been included on the board to implement this. The particular type of SRAMs are equal to those used on the latest LHC RadMon devices, meaning that they are extensively characterised [30]. The additional FPGA is needed since the SF2 lacks the pin resources for the SRAM interfaces. It should be emphasized that the RadMon design has been done as a standalone project earlier [31], so the extra functionality comes with a very low risk.

5 RCU2 FPGA design

The strict time constraint of the upgrade project implies that the design is largely based on the present RCU1 FPGA design. An important point is that the RCU1 FPGA design has proven to work very well, if the radiation related functional failures are disregarded.

The proposed design is given in figure 3. The overall structure of the design is inherited from the RCU1. The inclusion of the embedded ARM Cortex M3 core on the FPGA is a major simplification, regarding the DCS bus interface, compared to the present design. Additionally, the rich selection of hardcore peripherals provided by the Microcontroller Subsystem is utilized where appropriate, for instance the I²C interface to the ADCs and EEPROM, the SPI to the RadMon and external flash device, and the DDR interface to the external memory.

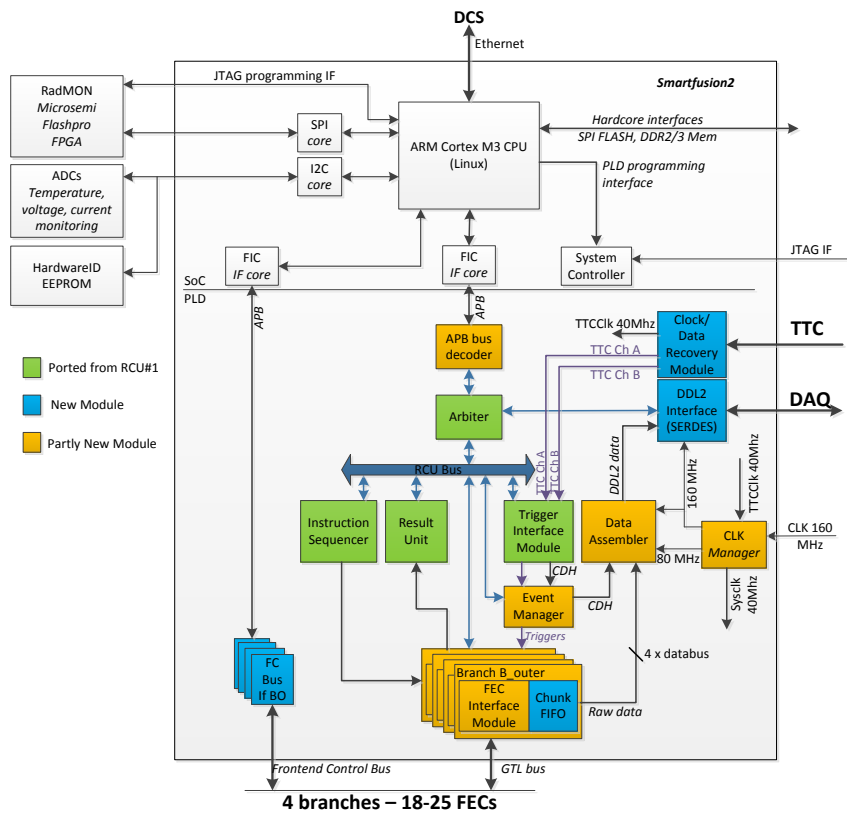


Figure 3. Simplified block schematic of the proposed RCU2 FPGA design.

Control system. Apart from some of the hardcore IP cores in the SF2 that are used for monitoring purposes, the Frontend Control Bus Interface is the only dedicated control system component in the SF2. This module implements the custom I²C interface to the FECs, used for monitoring of currents, voltages and temperatures on the FECs [32]. To improve the partitioning of the design, this module is connected to a separate fabric interface core (FIC) of the Microcontroller Subsystem than the rest of the logic in the RCU2 FPGA.

Readout scheme. A new readout scheme has been developed in order to utilize the improved parallelism of the RCU2 and so that the data analysis on the receiver side can be done as efficiently as possible. The receiver side expects the data to be ordered pad by pad and padrow by padrow. As seen in figure 3 there is one readout module per branch. The data from the channels (pads) for one single padrow belonging to one single branch is called a *chunk* of data, and the readout module stores this data into one large *chunk FIFO*. The channels in the padrows are divided between the branches so that one can read out the chunk FIFOs consecutively in a round robin fashion, and automatically keep the data in the correct order.

Except for the new structure with chunk FIFOs, most of the readout module is kept as explained in [11]. However, the increase in bandwidth given by the DDL2 implies that the readout of the data is separated in different clock domains. The data is read from the FECs with a 40 MHz clock (constrained by the FEC) into a buffer memory, then stored into the chunk FIFO at 80 MHz,

and finally read by the Data Assembler from the chunk FIFOs at 160 MHz. This ensures that no bandwidth limitations are imposed by the readout logic. The different clocks are generated from the recovered TTC clock using an onchip Phase Locked Loop (PLL).

Other new functionality. To further improve the efficiency of the data readout, the RCU2 FPGA will implement a new feature which discards junk data on the fly. Junk data typically comes from instabilities in the High Voltage system of the TPC and is recognized by several channels that have far more data than expected. This can be detected already at the RCU2 level by setting threshold levels regarding expected data size, and the RCU2 will simply not send these data to DAQ system at all. Simulations have shown that no good physics data is discarded by this algorithm.

Readout performance simulations. A simulation of the readout architecture has been developed in SystemC [33]. The focus of these simulations is the verification of the performance gain due to the new readout scheme and the confirmation that the resources offered by the SF2 device, in particular the available fabric memory blocks to be used for the chunk FIFOs, are sufficient. The simulation is based on real data from heavy ion collisions at different centralities (varying data sizes) recorded with the TPC in 2010. Different readout modes and hardware parameters (e.g. FIFO depth) are simulated. The readout time contains also the trigger delays and other external parameters.

Compared to the measurements of the present system, where the readout time for central events is about 4 ms [11], the first simulation results using data from 170 events indicates a readout time of approximately 1.5 ms that is an improvement factor of about 2.6 for the RCU2.

6 RCU2 DCS software

The inclusion of an Embedded Linux platform [34] on the RCU2 is permitting to reuse most of the existing low level DCS software, more specifically the FeeServer and the intercom layer [35]. The exception is the drivers for the RCU2 DCS which must be redesigned since the hardware interface of the RCU2 is different compared to the DCS board.

It was considered to discard the DCS implementation and instead send monitoring data in parallel to the physics data stream through the optical readout link. The latter might be reconsidered in the future. However, having a secondary interface that is capable of configuration and monitoring increases the flexibility of the RCU2 by giving redundant access to the complete RCU2 address space. Regardless of not playing any role in the event readout itself, the DCS can still be used to configure the FEE to perform readout. In addition, upgrading the designs in the RadMon FPGA and the SF2 fabric is much easier when the Linux is installed on the Microcontroller Subsystem.

7 Radiation tolerance measures and irradiation campaigns

Due to the increased luminosity the radiation load to the electronic components will also increase with respect to Run1. Radiation calculations for the present ALICE detector have previously been presented in [36] where the expected high energy hadron fluence rate ($>20\text{MeV}$) was estimated to 0.8 kHz/cm^2 for the worst case location of the TPC electronics, and for an interaction rate of 8 kHz. Scaling this fluence rate to an interaction rate of 30 kHz, the expected value for Run2

becomes 3 kHz/cm², which is a significant rate. The choice of the flash based Microsemi SF2 FPGA is therefore essential as the configuration memory of these devices are stated to be immune to SEUs [18]. This device also offers SEU protected memories, DDR bridges, Ethernet cores etc, and the DDR2/3 controllers are optionally secured with SECDED [18]. The latter is expected to improve the DCS stability since the Linux is uploaded to this memory on bootup. For the user memory and registers some mitigation means are needed and will be implemented.

For the total dose and 1 MeV neutron-equivalent fluence the initial numbers from [36] were 1.6 kRad and 4.5×10^{10} cm⁻² respectively, and this for a total operation of 10 ALICE years. Assuming a similar running program of pp, p-Pb and Pb-Pb interactions, and then scaling for a 3 year running period of Run2 including an increased interaction rate of 30 kHz for Pb-Pb, the expected numbers for the total dose and 1 MeV neutron-equivalent fluence will be roughly similar to those reported in [36]. A dose of less than a few kRads, and similar a 1 MeV neutron-equivalent fluence in the order of 10^{10} cm⁻², are not very significant as the onset for failure typically occurs for dose value above 10 kRad and 1 MeV neutron-equivalent fluence values of 10^{11} cm⁻² [37]. It should also be noted that the radiation calculations in [36] already contains a safety factor (2 to 3) because the multiplicity was overestimated for these calculations.

On a PCB level, there has been emphasis on the use of components already proven to function well in a radiation environment. For components selected where no data regarding radiation tolerance can be found, an irradiation campaign at the Oslo Cyclotron is presently being prepared. In the beginning of 2014 an irradiation campaign of the full readout chain will be done either at The Svedberg Laboratory (TSL) in Uppsala or at the Paul Scherrer Institute (PSI) in Zürich. At both these locations higher beam energies are available (for instance 180 MeV protons and neutrons at TSL).

8 Outlook and conclusions

The aim of the RCU2 is to do a small-scale upgrade project that enables ALICE to collect a significantly larger amount of events in the central barrel at a moderate cost. The cost is estimated to be about 455 kCHF. The readout time for TPC events is simulated to be improved by a factor of 2.6 for central heavy ion events, which enables the TPC readout to conform to the running scenario that is envisaged for Run2 of ALICE. The stability problems seen in Run1 will most likely be reduced to a negligible level, but this is yet to be confirmed by the irradiation campaigns.

The RCU2 is a project that relies quite heavily on the reuse of the current design and ideas. This reduces the risk of the project substantially. The biggest risk has been identified to be the aggressive time-schedule. The first prototype of the RCU2 is expected to be produced by the end of the year, and the mass production will be done from April 2014. The installation on the detector is planned in October 2014, after which the commissioning period starts. This will last until March 2015. At the time of writing the progress of the project is approximately as originally planned.

References

- [1] ALICE collaboration, *ALICE Experiment at the CERN LHC*, 2008 *JINST* **3** S08002.
- [2] J. Alme et al., *The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events*, *Nucl. Instrum. Meth. A* **622** (2010) 316.

- [3] K. Šafařík, *Overview of recent ALICE results*, *Nucl Phys. A* **904** (2013) 27C.
- [4] Altera, *Excalibur devices, hardware reference manual*, v3.1 (Nov. 2002).
- [5] D. Evans et al., *The ALICE Central Trigger System*, *Real Time Conference, 2005. 14th IEEE-NPSS* (2005) 5.
- [6] J. Alme et al., *Radiation-tolerant, SRAM-FPGA based trigger and readout electronics for the ALICE experiment*, *IEEE Trans. Nucl. Sci.* **55** (2008) 76.
- [7] Xilinx Inc., *Virtex-II Pro and Virtex-II Pro X platform FPGAs: complete data sheet*, DS083 v4.7 (Nov. 2007).
- [8] Actel Corporation. *Actel ProASIC^{plus} flash family FPGAs datasheet*, v5.7 (Sept. 2008).
- [9] F. Lima Karstensmidt et al., *Fault-tolerance techniques for SRAM-based FPGAs*, Springer, Germany (2006).
- [10] J. Alme, *Firmware development and integration for ALICE TPC and PHOS front-end electronics*, Ph.D. thesis, University of Bergen, Bergen, Norway (2008).
- [11] A. Ur Rehman, *The ALICE TPC readout electronics — Design, performance optimization and verification of the DAQ circuit*, Ph.D. thesis, University of Bergen, Bergen, Norway (2012).
- [12] J. Alme et al., *Radiation tolerance studies using fault injection on the readout control FPGA design of the ALICE TPC detector* [2013 JINST 8 C01053](#).
- [13] K. Røed and J. Alme, *Single event upsets in the readout control FPGA of the ALICE TPC detector during the first LHC running period*, poster session at *Topical Workshop on Electronics for Particle Physics 2013*, September 23–27, Perugia, Italy (2013).
- [14] A. Junique et al., *Upgrade of the ALICE-TPC read-out electronics*, [2010 JINST 5 C21026](#).
- [15] V. Holsen, *Upgrade of the readout electronics for the ALICE TPC*, Master thesis, University of Bergen, Bergen, Norway (2013).
- [16] F. Costa, *DAQ readout after LS1 and LS2*, at *ALICE week presentation* (2012).
- [17] Microsemi Inc., *SmartFusion2 system-on-chip FPGAs datasheet*, rev. 4 (Jun. 2013).
- [18] Microsemi Inc., *SmartFusion2 system-on-chip FPGAs product brief*, rev. 12 (Oct. 2013).
- [19] Microsemi Inc., *SmartFusion2 microcontroller subsystem user guide*, rev. 4 (Sept. 2013).
- [20] J. Christiansen et al., *TTCrx reference manual — A timing, trigger and control receiver ASIC for LHC detectors*, v. 3.11 (Dec. 2005).
- [21] M. Dimitriyev et al., *Development of a microTCA carrier hub for CMS at HL-LHC*, [2010 JINST 5 C12042](#).
- [22] Avago, *HFBR-1312TZ transmitter HFBR-2316TZ receiver 1300 nm fiber optic transmitter and receiver datasheet*, av02-1500en edition (2012).
- [23] M. Gastal and P. Moreira, *Radiation qualification of commercial off-the-shelf P-I-N receivers for the TTC system*, Technical Report, CERN-EP/MIC, Geneva Switzerland (2003).
- [24] Marvell, *Marvell 88E1111 product brief, integrated 10/100/1000 Ultra Gigabit Ethernet Transceiver*, mv-s105540 rev. a (Oct. 2013).
- [25] DCS Board Schematic, board v. 1.64 (Dec. 2005).
- [26] J. Alme et al., *Case study of a solution for active partial reconfiguration of a Xilinx Virtex-II Pro*, in the proceedings of *3rd FPGAWorld conference*, Stockholm, Sweden (2006).

- [27] K. Røed et al., *First measurement of single event upsets in the readout control fpga of the ALICE TPC detector*, [2011 JINST 6 C12022](#).
- [28] Microsemi, *ProASIC3 flash family FPGAs with optional soft ARM support datasheet*, rev. 13 (Jan. 2013).
- [29] Cypress Perform, *CY62157EV30 MoBL 8 Mbit (512K x 16) Static RAM datasheet*, rev. f (Jun. 2009).
- [30] A. Masi et al., *The new generation of the CERN accelerator radiation monitoring system for electronics*, [IEEE Trans. Nucl. Sci. 60 \(2013\) 3475](#).
- [31] A. Velure, *Design, implementation and testing of SRAM based neutron detectors*, Master thesis, University of Bergen, Bergen, Norway (2011).
- [32] C. G. Gutiérrez, *Readout and control system for the ALICE TPC electronics*, Ph.D. thesis, Escuela Técnica Superior de Ingenieros Industriales y de Telecomunicación University of Cantabria, Spain (2007).
- [33] Accellera Systems Initiative, *SystemC webpage*, <http://www.systemC.org>.
- [34] EmCraft Systems, *Linux Cortex-M user's manual*, rel. 1.11.0 (Jul. 2013).
- [35] D.T. Larsen, *Monitoring and calibration of the ALICE Time Projection Chamber*, Ph.D. thesis, University of Bergen, Bergen, Norway (2010).
- [36] A. Morsch and B. Pastirčák, *Radiation in ALICE detectors and electronic racks*, ALICE-INT-2002-28 (2004).
- [37] C.C. Foster, *Total ionizing dose and displacement-damage effects in microelectronics*, [MRS Bulletin 28 \(2003\) 136](#).