
A radiation tolerance study of the ALICE TPC Readout Control Unit 2

ZHAO Chengxin

Dissertation submitted for fulfilment of the degree of Doctor of Philosophy
at the University of Oslo, Norway

August 2017



© **Zhao Chengxin, 2017**

*Series of dissertations submitted to the
Faculty of Mathematics and Natural Sciences, University of Oslo
No. 1916*

ISSN 1501-7710

All rights reserved. No part of this publication may be
reproduced or transmitted, in any form or by any means, without permission.

Cover: Hanne Baadsgaard Utigard.
Print production: Reprosentralen, University of Oslo.

Acknowledgement

It has been almost five years since I started my PhD in August 2012. Lots of memories come to my mind while I am writing this page. I have met some gem-like personalities during my stay at Oslo, Bergen and Geneva. I am so grateful for your help and supports.

First of all, I want to thank my supervisors Ketil Røed and Johan Alme. If you hadn't supported and believed in me, I would not be able to break the darkness of my PhD and make it alright. Thank you for bringing me back to correct path while I was feeling hopeless in 2014. Ketil has helped in all the aspects from academic to personal life. He told me that the process of PhD life is to create the path in a wild forest, sometimes I might be lost but I need to keep on going. The valuable advices that I got from him in the tons of discussions and e-mail exchanges will benefit my whole life. Johan is always willing to give me the positive comments. This makes me confident and it is very important for me to persist on my research. He is always answering my questions in a patient manner, even for the most basic and boring ones.

The same thanks also go to my supervisor Helge Balk and former supervisor Kjetil Ullaland. Helge helped me a lot in the administrative stuffs and solved the practical issues of my first days in Oslo. Kjetil's guidance and advice during our occasional discussions in Bergen has been quite helpful in electronic design.

Dear Toralf Bernhard Skaali, I have to show my special gratitude to you. Thanks for picking me up and this changed the track of my life. Thanks for introducing me to the Norwegian group and the support you have given to me during the past years. Many thanks always go to Dieter Rörich for the advice on how to proceed my PhD and helping me to involve into the Norwegian community. I want to thank Attiq Ur Rehman for spending so many days with me in the lab in Bergen and at CERN. The hands-on experience in electronics design is quite important for my further career. Also thank you for having patience with my faults.

A very special thank you to Lars Bratrud. Our friendship started from debugging the PHOS electronics and then goes to all the parts in our life.

I am thankful to the whole Norwegian group and the RCU2 community. Thanks to all the guys that I met in Oslo, Bergen, Geneva, Uppsala, Stockholm, Beijing, Wuhan and Changsha.

The last but the biggest gratitude is to my family. Thanks for your supports during last three decades.
I am proud of you.

Chengxin

Oslo, May 2017

Abstract

ALICE is a general-purpose detector that is designed to study the physics of quark-gluon plasma. The Time Projection Chamber (TPC) is one of the major detectors of ALICE. The TPC electronics consists of 4356 Front-end cards (FECs), which are controlled by 216 Readout Control Units (RCU). Each RCU connects to between 18 and 25 FECs using a multi-drop bus. In LHC Run1, the Readout Control Unit 1 (RCU1) performed even better than specification. However, in Run2 the energy of colliding beams is increased from 8 TeV to 14 TeV (maximum value) and higher luminosity, which leads to larger event size and higher radiation load on the electronics. As a solution, the Readout Control Unit 2 (RCU2) is designed to provide faster readout speed and improved radiation tolerance with respect to the RCU1.

The RCU2 is conceptually similar to the RCU1 and it reuses the existing infrastructure and readout architecture of the TPC electronics. However, the multi-drop bus is split into four branches from the two branches and the bandwidth of the Detector Data Link (DDL) is increased from 1.60 Gbps to 3.125 Gbps. Correspondingly, the firmware is designed to utilize the improved parallelism. These actions ensure that the readout speed of the RCU2 can be improved by a factor of ~ 2 with respect to the RCU1. The flash-based Microsemi Smartfusion2 FPGA SOC is used as the main FPGA instead of the SRAM based Xilinx Virtex 2 Pro FPGA that was used on the RCU1. Because its configuration cells are immune to Single Event Effects, the radiation tolerance of the RCU2 was expected to be improved.

The primary objective of this thesis has been to study the radiation tolerance of the RCU2. This is done through several irradiation tests, which are divided into two steps. To start with, the radiation sensitivity of the Smartfusion2 FPGA and all the hardware interface are characterized. Afterwards, a system-level irradiation test is performed. Actions have been taken against all the radiation related problems that were revealed during the irradiation tests. Running experience shows that radiation tolerance of the readout system based on RCU2 is about 10 times better as compared to the RCU1 for p-Pb collisions at similar energy level.

The second objective of this thesis was to develop the firmware modules that realizes the readout algorithms. Development of the firmware has gone through three versions, the first prototype, the

second prototype and the commissioning version, and important contributions were made to the first two versions.

The integration and testing of the RCU2 is also an important task covered in this thesis. Functional tests were performed for the mass production, the irradiation tests at and the final installation at the TPC. Readout performance of the RCU2 has been characterized and the solutions aiming to further increase the readout speed have been proposed and verified.

Table of Contents

1	Introduction.....	1
1.1	The ALICE Experiment	2
1.1.1	ALICE sub-detectors	3
1.1.2	The Time Projection Chamber (TPC).....	4
1.2	The TPC readout electronics in Run1	5
1.2.1	The Front-end Card (FEC).....	6
1.2.2	The Readout Control Unit (RCU).....	6
1.3	TPC consolidation effort during LS1	8
1.3.1	Motivation.....	8
1.3.2	Solutions	9
1.4	Primary objective and main contribution	12
1.4.1	Radiation Tolerance	12
1.4.2	Design, Integration and Test of the RCU2.....	12
1.5	Outline of the thesis.....	13
2	Radiation effects on the RCU2	15
2.1	Interaction of particle with matters	15
2.1.1	Charged particles	15
2.1.2	Neutral particles	17

2.2	Radiation Effects related to the RCU2	18
2.2.1	Radiation environment of the RCU2	19
2.2.2	Single Event Effects (SEEs)	20
2.2.3	Total Ionizing Dose (TID) effect	23
2.2.4	Summary	24
2.3	Irradiation tests	24
2.3.1	Selection of test facilities	24
2.3.2	Dose calculation.....	27
2.3.3	Predicating the rate of SEEs induced errors	28
3	The RCU2.....	29
3.1	RCU2 overview.....	29
3.1.1	The RCU2 main FPGA.....	32
3.1.2	TTC Interface.....	34
3.1.3	DCS Interface.....	34
3.1.4	DAQ Interface.....	36
3.1.5	Radiation Monitor (RadMon)	36
3.1.6	ALTRO bus backplane	37
3.1.7	Software design.....	38
3.2	Smartfusion2 (SF2) firmware overview.....	41

3.2.1	Clocking and reset scheme.....	42
3.2.2	Firmware versions.....	43
3.2.3	Trigger Module	45
3.2.4	DDL2 Module	45
3.2.5	Monitoring and Safety Module.....	46
3.2.6	RCU2 Bus	47
3.3	Readout Module	48
3.3.1	First prototype.....	49
3.3.2	Second prototype	50
3.3.3	Commissioning version	59
3.4	Summary	63
4	Radiation Tolerance of the RCU2.....	65
4.1	Tested devices	66
4.2	Characterization of the Smartfusion2 (SF2).....	68
4.2.1	Single Event Latch-up (SEL) test	68
4.2.2	Fabric SRAM test	71
4.2.3	Embedded SRAM test.....	74
4.2.4	Flip-flop test.....	75
4.2.5	PLL test	77

4.2.6	Total Ionizing Dose (TID) effects test	80
4.3	Hardware Interface tests	81
4.3.1	TTC interface test	82
4.3.2	DAQ interface test	85
4.3.3	DCS interface test	87
4.4	System level irradiation test	89
4.4.1	Readout stability	90
4.4.2	DCS stability	92
4.5	Radiation Monitor test.....	94
4.6	Summary and Conclusion	96
5	Testing and integration.....	99
5.1	Validation of the ALTRO Interface	99
5.1.1	Test with simple ALTRO bus master	99
5.1.2	Test with ALTRO Bus Interface Module	101
5.2	Test with the second prototype of firmware.....	103
5.3	Test with the commissioning version of firmware	105
5.3.1	DDL2 link at 2.125 Gbps.....	105
5.3.2	DDL2 link at 4.25 Gbps.....	109
5.3.3	DDL2 link at 3.125 Gbps.....	110

5.3.4	Discussion on Radiation Mitigation.....	114
5.4	Summary	117
6	Summary and conclusion.....	119
6.1	Main Contribution.....	119
6.2	Running experience.....	120
6.3	Outlook.....	121
	Reference	123
	Appendix A. List of publications.....	129
A.1	As main contributor.....	129
A.2	As collaborator	129
	Appendix B. List of Abbreviations.....	130
	Appendix C. Fluence calculation.....	133
	Appendix D. RCU2 Data Format.....	134
	Appendix E. Screenshots and test results.....	136
E.1	SEU counts of SRAM tests	136
E.2	Screenshots of tests	137
E.3	Readout speed benchmark.....	138
E.4	Procedure of test prior to mass production.....	139
E.5	Commissioning of the RCU2	141

List of Figures

Figure 1-1 The LHC with four experiments	1
Figure 1-2 Roadmap of LHC to its full potential.....	1
Figure 1-3 The ALICE detector	2
Figure 1-4 Three-dimensional view of the TPC	4
Figure 1-5 Layout of the TPC readout electronics.....	5
Figure 1-6 Signal path in the TPC readout electronics	6
Figure 1-7 Front view of the RCU1	7
Figure 1-8 Schematic layout of the RCU1	7
Figure 1-9 Sketch of the RCU2 design [18]	10
Figure 1-10 Comparison of the readout time between RCU2 simulations and the measurement of RCU1 in LHC Run1	11
Figure 2-1 Mitigation of MBUs in memory cells in SF2	21
Figure 2-2 Structure of floating gate transistor in flash-based FPGA	23
Figure 2-3 Test facility of the Oslo Cyclotron	25
Figure 2-4 Test facility of the Svedberg Laboratory.	26
Figure 3-1 Overview of the RCU2.....	30
Figure 3-2 RCU2 Board (front side)	31
Figure 3-3 RCU2 Board (back side)	32

Figure 3-4 Schematic layout of the SF2 FPGA SoC.....	33
Figure 3-5 Digital part of the DCS Interface	35
Figure 3-6 The DAQ Interface.....	36
Figure 3-7 ALTRO bus backplane	38
Figure 3-8 Architecture of RCU2 Software Design.....	39
Figure 3-9 Flowchart of the software booting process	40
Figure 3-10 Overview of the RCU2 firmware.....	42
Figure 3-11 DDL2 protocol blocks	46
Figure 3-12 Monitoring and Safety module sub-modules	47
Figure 3-13 RCU2 bus structure topology.....	48
Figure 3-14 Second prototype of Readout Module.....	51
Figure 3-15 ALTRO Bus Interface sub-module.....	52
Figure 3-16 Chronogram of the CHRDO command.....	53
Figure 3-17 Screenshot of CHRDO operations	54
Figure 3-18 Branch Readout Unit sub-module.....	55
Figure 3-19 Flow chart of the Branch Readout Unit	56
Figure 3-20 RCU2 Data package structure	58
Figure 3-21 Commissioning version of Readout Module.....	60
Figure 3-22 Sub-module topology of the Channel Formatter.....	61

Figure 3-23 Sub-module topology of the Branch Readout Unit.....	62
Figure 4-1 Emcraft SF2 Starter-Kit	67
Figure 4-2 SEL test setup.....	68
Figure 4-3 Current consumption of the SF2 FPGA in first SEL test	69
Figure 4-4 Cross-section of current jumps vs. supply voltage in the first SEL test.....	70
Figure 4-5 Current consumption of the SF2 FPGA in second SEL test	71
Figure 4-6 SRAM irradiation test setup.....	72
Figure 4-7 SEUs and fluence for the SRAM test in campaign No.3	72
Figure 4-8 SEUs and fluence for the eSRAM test at campaign No.7.....	74
Figure 4-9 Flip-flop test setup.....	75
Figure 4-10 First PLL test setup in campaign No.3	77
Figure 4-11 Output clock of PLL with different configuration when it loses lock.....	78
Figure 4-12 Second PLL test setup in campaign No.6	80
Figure 4-13 TID effect on the SF2 chip.....	81
Figure 4-14 Setup of the TTC interface test	82
Figure 4-15 Example of radiation effect in an optical receiver	84
Figure 4-16 Setup of the DAQ interface irradiation test.....	86
Figure 4-17 Test setup of DCS Interface	86
Figure 4-18 Setup of system level irradiation test	88

Figure 4-19 Setup of the system-level irradiation test (without collimator).....	89
Figure 4-20 Setup of the system-level irradiation test (with collimator).....	90
Figure 4-21 Setup of the RadMon test.....	94
Figure 4-22 SEU counts as a function of fluence of the RadMon test	95
Figure 5-1 Test design with the simple ALTRO bus master	100
Figure 5-2 Observation of the write and read transaction.....	100
Figure 5-3 Test design with the ALTRO Bus Interface.....	101
Figure 5-4 Screenshot of CHRDO transaction	102
Figure 5-5 Test procedure for the RCU2 with the second prototype of firmware	103
Figure 5-6 The CERN setup - 1 RCU2 connects to 25 FECs	106
Figure 5-7 Test procedure for the RCU2 with the production of firmware	106
Figure 5-8 Benchmarking on the RCU2 with DDL2 at 2.125 Gbps	108
Figure 5-9 Benchmark on the RCU2 with DDL2 at 4.25 Gbps.....	109
Figure 5-10 Benchmarking of the RCU2 with DDL2 at 3.125 Gbps..	112
Figure 5-11 Measurement of reading single word from one channel..	114
Figure 5-12 Reconstructed data taken by TPC in the first p-Pb collision in Run2....	118
Figure D-1 CDH words of RCU2	134
Figure D-2 RCU2 payload words	134
Figure D-3 RCU2 Trailer words	135

Figure E.1-1 SEUs and fluence for the SRAM test in campaign No.2.....	136
Figure E.1-2 SEUs and fluence for the SRAM test in campaign No.3.....	136
Figure E.2-1 Measurement of RCU2 signals	137
Figure E.2-2 Measurement of CHRDO for an empty channel	137
Figure E.2-3 Measurement of CHRDO for the number of samples as 10.....	138
Figure E.4-1 Inspection of oscillator	140
Figure E.4-2 Screenshot of data-taking status	141
Figure E.5-1 The first 6 installed RCU2.....	142
Figure E.5-2 Data loop in sector (six readout partitions).....	143
Figure E.5-3 Radiation Monitor of the RCU2	143
Figure E.5-4 Check the DCS of installed partitions (colored blue).....	144
Figure E.5-5 Check the Status of FECs (Monitoring and Safety Module).....	144
Figure E.5-6 Check the power of installed partitions (colored purple)	145

List of Tables

Table 3-1 Resources comparison between the RCU1 main FPGA, the RCU2 main FPGA and the RCU1 firmware	33
Table 3-2 Execution time of each transaction for RCU1	63
Table 4-1 Overview of the irradiation campaigns (time-wise)	66
Table 4-2 SRAM test results	73
Table 4-3 Flip-flop test results	76
Table 4-4 PLL test results	79
Table 4-5 TTC interface test results (PLL lose lock)	83
Table 4-6 DAQ interface test results	85
Table 4-7 DCS interface irradiation test results	88
Table 4-8 Readout stability observations	91
Table 4-9 DCS stability observation	93
Table 4-10 Summary of the MTBF in Run2 of the RCU2	97
Table 5-1 Stress test of the ALTRO interface	102
Table 5-2 System level validation of the RCU2 (second prototype of firmware)	104
Table 5-3 System level validation of the RCU2 (commissioning version of the firmware with DDL2 bandwidth of 2.125 Gbps)	107
Table 5-4 Test results of 6 sample RCU2s	111
Table 5-5 Reliability estimation of the complete RCU2	116

Table 5-6 verification of firmware with mitigation actions	117
Table 6-1 Overview of End of Run (EoR) reasons for the ALICE experiment	120
Table E.3-1 Readout speed of single event (partition 1).....	138

1 Introduction

The Large Hadron Collider (LHC) is the largest particle accelerator in the world. It is hosted by CERN¹, the European Organization for Nuclear Research, which is located near Geneva on the border between Switzerland and France. The LHC lies about 100 meters beneath the ground in a tunnel with a circumference of 27 kilometers. Two particle beams accelerated close to the speed of light travel in opposite directions and collide at dedicated locations, where four major experiments, ALICE [1], ATLAS [2], CMS [3] and LHCb [4], are positioned, see Figure 1-1.



Figure 1-1 The LHC with four experiments [5]




Figure 1-2 Roadmap of LHC to its full potential (from [6] with the add-on of the Run1 scenario)

As shown in Figure 1-2, the roadmap of the LHC to achieve its full design energy has been divided

¹ CERN is the acronym of its French name Conseil Européen pour la Recherche Nucléaire

into several running periods and long shut-down periods. In the running periods, the LHC provides collisions for the experiments to take physics data. In the long shut-down periods, the LHC and the experiments are under maintenance and upgrade as a preparation for the next running period. During the first successful running period from November 2009 to February 2013 (Run1), the LHC ramped up its center-of-mass energy from the start-up 900 GeV to 7 ~ 8 TeV. In the second running period (Run2) which started in 2015 and will last until 2018, the center-of-mass energy of the collisions will reach up to 13 TeV for p-p collisions.

This thesis is part of the upgrade activities for the readout electronics of the ALICE Time Projection Chamber (TPC) [7] during the long shut-down 1 (LS1). Hence, this chapter will introduce the ALICE experiment and the TPC readout electronics used during Run1. As already mentioned, Run2 will introduce a higher energy in the collisions, and the practical implications of this will be discussed at the end of this chapter as part of the motivation for the upgrade.




Figure 1-3 The ALICE detector [1]

1.1 The ALICE Experiment

ALICE is a general-purpose detector designed to study the physics of quark-gluon plasma. In normal condition, quarks are bound into hadrons by the force carrier of the strong force (gluons). A Pb-Pb collision in the LHC will create an extremely high temperature and energy density so that the

hadrons undergo a phase transition into quark-gluon plasma, where the quarks and gluons are not in a bound state. According to the theory of the Big Bang, the universe was in a state of quark-gluon plasma up to a few milliseconds after the Big Bang. As the temperature and the density dropped, the quarks and gluons were bound into different kinds of hadrons, which constitute the basic building block of matters. Since the life-time of quark-gluon plasma is rather short, it cannot be observed directly. Hence, the ALICE detector is comprised of several sub-detectors that are designed to observe the signatures that indicate the existence of quark-gluon plasma and to study its properties. Details on the physics and experimental observables of the ALICE experiment can be found in [1].

1.1.1 ALICE sub-detectors

A collision at the LHC is called an event, and it produces a large number of secondary particles. The ALICE experiment is optimized to study Pb-Pb events, but pp and Pb-p events are recorded as well to provide reference data [1]. For each event, the momentum of the charged particles and the energy of the neutral particles are measured and in addition the types of particles (hadrons, electrons, photons and muons) are identified. Figure 1-3 gives the schematic layout of the ALICE detector. The size is $26 \times 16 \times 16 \text{ m}^3$ with a weight of 10,000 tons. To accomplish the above-mentioned tasks, a set of sub-detectors are placed in different layers, in some distance away from the central collision point. These sub-detectors can be sorted into three categories: the central detectors, the forward sub-detectors and the muon spectrometer. Details on these sub-detectors can be found in [1] and only a short summary is presented here.

The central detectors can be sorted into the central tracking detectors, the particle identification detectors and the calorimeters. The central tracking detectors include the Inner Tracking System and the cylindrical TPC. The Inner Tracking System is designed to localize the primary vertex and reconstruct the secondary vertices. It comprises six layers of silicon detectors. The innermost two layers are the Silicon Pixel Detector, the middle two layers are the Silicon Drift Detectors, and the outermost two layers are the Silicon Strip Detector. The TPC is the main tracking detector in the ALICE experiment. Together with other central detectors, it is optimized to provide the charged particle momentum measurement, the particle identification and the vertex determination. The Transition Radiation Detector, the Time of Flight detector and the High-Momentum Particle Identification Detector are particle identification detectors. The Transition Radiation Detector is designed to identify electrons with the momenta above 1 GeV/c. The Time of Flight detector and

the High-Momentum Particle Identification Detector identify the charged particles (protons, kaons and muons) having intermediate and large momentum, respectively. Two calorimeter detectors, the Photon Spectrometer and the electromagnetic calorimeter (EMCal, since 2015 DCal) are designed to detect photons and measure particle jets, respectively.

The forward sub-detectors include the Zero Degree Calorimeter, the Photon Multiplicity Detector, the Forward Multiplicity Detector, the Veto and the Time Zero. These forward sub-detectors measure the multiplicity and the spatial distortion of the non-interacting nucleons, which can be used to determine the geometry of the collision. In addition, the Veto and the Time Zero are also responsible for providing minimum biased triggers.

The muon spectrometer detects muons after all the other particles have been stopped by the absorbers in the forward region and provides fast trigger decisions.




Figure 1-4 Three-dimensional view of the TPC [8]

1.1.2 The Time Projection Chamber (TPC)

The TPC is the main tracking detector in the ALICE experiment. The layout of the TPC detector is shown in Figure 1-4. It is a cylindrical volume of 88 m^3 that is divided into two field cages by a high voltage electrode and filled with gas that can be ionized. The TPC has an inner radius of 0.85 m and an outer radius of 2.8 m . It spreads over 5.1 m along the path of the colliding beams. Charged particles created in the collisions ionize the gas. In the presence of the electric field, the

electrons will drift toward the end-plates, where multi-wire proportional chambers are used to multiply the electrons from primary ionization.

1.2 The TPC readout electronics in Run1

In the TPC there are in total 557 568 detector pads divided equally between the two end-plates, each of which is mapped to an individual channel in the readout electronics. The TPC readout electronics consists of 4356 Front-End Cards (FECs) and 216 Readout Control Units (RCUs), which are distributed into 36 trapezoidal sectors (18 in each end-plate). As shown in Figure 1-5, each sector covers six readout partitions along the radial direction in the TPC barrel. Each readout partition consists of from 18 to 25 FECs, depending on the readout partitions, which are connected to one RCU with a parallel multi-drop bus – the ALICE TPC Readout (ALTRO) [9] bus.

Figure 1-6 shows the signal path in the TPC readout electronics. The FEC processes the electric signals generated by the charges deposited on the detector pad and buffers the data. The RCU reads the data from the FECs, processes it and then transmits it to the Data Acquisition (DAQ) system [10].




Figure 1-5 Layout of the TPC readout electronics




Figure 1-6 Signal path in the TPC readout electronics

1.2.1 The Front-end Card (FEC)

As shown in Figure 1-6, each FEC contains 128 signal channels, which are realized by 8 Preamplifier Shaper (PASA) [11] chips and 8 ALTRO chips. The PASA amplifies and shapes the electric signals from the detector pads. The ALTRO chip does analog to digital conversion, digital signal processing and buffering of the acquired data.

In addition, the FEC holds one SRAM based FPGA, the Board Controller. The Board Controller does low level control system tasks like monitoring of current, voltages and temperatures on the FEC. In addition, it controls the direction of the ALTRO bus in the communication between the actual FEC and the RCU.

1.2.2 The Readout Control Unit (RCU)

A front view and schematic layout of the RCU used in LHC Run1 can be found in Figure 1-7 and Figure 1-8, respectively. From here on this RCU is named RCU1. It consists of a motherboard with two daughter boards: The Detector Control System (DCS) [12] board and the Source Interface Unit (SIU) card.

The DCS board hosts a TTCrx [13] chip that processes the trigger information coming from the Trigger, Time and Control (TTC) [7] via an optical link and provides the 40 MHz clock. In addition, a minimalistic Linux platform is running on the ARM processor embedded in a SRAM-based Altera FPGA [14]. Dedicated software operating on this Linux platform propagates the monitoring values to the DCS through an Ethernet link, so that any potential hazardous situation can be detected.

The SIU card ships the packaged data, coming from the motherboard, to the DAQ through an optical

link of 1.280 Gbps. The protocol of the Detector Data Link (DDL) [7] is implemented on a flash-based FPGA on the SIU card.




Figure 1-7 Front view of the RCU1




Figure 1-8 Schematic layout of the RCU1

The motherboard holds the RCU1 main FPGA (SRAM-based Virtex Pro2 [15]) and a supporting FPGA (flash-based Actel APA075 [16]). The main FPGA is in charge of the data readout algorithms. It moves the sampled data from the FECs to the RCU1, processes and packages the data, and then pushes the data to the SIU card. At the time when the choice of the main FPGA was made, no flash-based FPGAs with enough resources were available to implement the readout algorithm, thus the

SRAM-based FPGA was selected [17]. Because the configuration memory of the main FPGA is proven to suffer Single Event Upsets (SEUs), the flash-based supporting FPGA has been added to detect the SEUs and reconfigure the main FPGA.

1.3 TPC consolidation effort during LS1

1.3.1 Motivation

Due to the enlarged event size and the increased event rate, the RCU1 was expected to limit the readout rate for Run2 [18]. In addition, the radiation related issues are expected to be more critical in Run2 because of the higher radiation load [19]. The limitations on the RCU1 is discussed below.

Data rate limitations

All the FECs are connected to the RCU1 via the ALTRO bus, which is divided into two separate branches. The bandwidth of each branch is 1.60 Gbps and it serves from 9 to 13 FECs, depending on the readout partitions. The readout time of an event is defined as that of the slowest readout partition, which is the readout partition 1 with 25 FECs. In each branch, all the channels in all the FECs need to be read sequentially. Therefore, for high occupancy events like central Pb-Pb collisions, the readout through the ALTRO bus is the bottleneck of the readout system.

According to the measurement for Pb-Pb events in 2010, the readout time of the TPC reached up to 4 ms (250 Hz), depending on the number of tracks [20]. In Run2 a readout rate of 400 Hz is expected and the event size will increase by 25% [20]. To accomplish this performance, the readout speed of the RCU1 needs to be improved by a factor of at least 2.

Radiation Tolerance

One major challenge for the TPC electronics is the radiation created by the colliding beams in the LHC. There are two kinds of radiation effects that is of concern, the Single Event Effects (SEEs) and the Total Ionizing Dose (TID) effect. The SEE is a transient effect, which is induced by a single ionizing particle. The TID effect is cumulative effect which refers to the total dose received by the Front-End Electronics (FEE) during its life-time. Two quantities that are commonly used to describe the radiation environment are the flux and the dose. Details regarding the radiation environment of TPC and the radiation effects can be found in section 2.1 and section 2.2.

In general, the RCU1 was performing very well in Run1. However, it is not a radiation-tolerant system and the main reason is the SRAM-based main FPGA. The radiation effects, of which the dominant ones are the SEUs in the configuration cell of the main FPGA, leading to the readout getting stuck (busy) or corrupted event headers [21]. Consequently, these errors caused stops of the data-taking in Run1 [18]. The SEU sensitivity of the RCU1 FPGA design has been characterized in [19], and it was found that about 1% of the SEUs will lead to the abortion of a physics run. In addition, there is no radiation protection on the DCS board, whose main FPGA is also SRAM-based. In Run1, the DCS board has experienced frequent communication errors (DCS-RCU) and communication losses (Ethernet to the DCS) [18]. Although these scenarios are not critical for data-taking, the loss of monitoring should be avoided.

In Run1, the longest data-taking session in heavy-ion collisions is 8 hours and 4 mins (run 138275 in the logbook [22]). In Run2, the duration of each data-taking session should be similar to that in Run1, so the RCU2 should be capable of reading data continuously for at least ~8 hours.

As mentioned above, in Run2 the expected radiation load in terms of the flux of fast hadrons on the TPC electronics located in the innermost locations (worst-case) will increase to 3.0 kHz/cm² from the 0.8 kHz/cm² in Run1 [18]. With such a significant increase, radiation effects are therefore foreseen to occur more frequent on the RCU1 in Run 2 (discussed in section 2.2). Considering the study of the SEU rate for heavy ion runs in 2011 and the luminosity for Run2, the data-taking is expected to stop around every single hour, if no actions are taken on the RCU1 [18].

Conclusion

Based on the information presented above, it was concluded that the readout rate needed to be increased by a factor of at least 2. In addition, the radiation tolerance should also be improved to withstand the higher radiation load in Run2.

1.3.2 Solutions

To consolidate the readout system and improve its performance, two upgrade options have been discussed. The first one is the Front-End Card Interface solution [23] and the second one is the Readout Control Unit 2.

The Front-End Card Interface solution

In this solution, each FEC is connected to a Front-End Card Interface that translates the parallel ALTRO bus interface into a serial one, which could deal with the peak data rate of 1.60 Gbps [23]. Therefore, the readout speed of the upgraded system could be improved with a factor of 10 with respect to that of the RCU1 [23]. In addition, this solution is also 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 [18].

However, this solution was eventually dropped because a lot of new PCB boards and fibers needed to be produced and installed in the TPC, which was not suitable for the aggressive time scale of LS1.




Figure 1-9 Sketch of the RCU2 design [18]

The Readout Control Unit 2 (RCU2)

The RCU2 was then proposed to give the needed performance. A sketch of the RCU2 is shown in Figure 1-9. The RCU2 is conceptually similar to the RCU1 and it reuses the existing infrastructure and architecture of the TPC electronics, such as the cables for TTC, DCS, DAQ and power. However, the ALTRO bus has been split from the current two branches into four branches. Correspondingly, the DDL protocol has been upgraded to the DDL2 protocol [24], which uses the same fiber but has a higher theoretical bandwidth of 4.25 Gbps. To utilize the improved parallelism, a new readout algorithm has also been implemented (discussed in section 3.3.1). The RCU2 could not improve the readout speed by a factor of 10 as the Front-End Card Interface solution. However, assuming there is no other bottleneck in the system, it could ensure at least a doubling of the readout speed, which

still fulfills the requirements for Run2.

Before designing the RCU2, simulations were performed with a SystemC [25] model² of the new readout architecture. Actual data recorded in the heavy ion collisions in 2010 was used and bandwidth of the DDL2 link was set to be 4.25 Gbps. The simulation showed that the readout time of the largest event is ~ 1.6 ms (in subplot (a) in Figure 1-10), which is 2.5 times faster than the current speed of 4.0 ms (in subplot (b) of Figure 1-10). Taking the 25% increase on the event size into consideration, the readout rate of the RCU2 will reach ~ 500 Hz (2 ms), which doubles the readout speed of the RCU1.




Figure 1-10 Comparison of the readout time between RCU2 simulations [26] and the measurement of RCU1 in LHC Run1 [20]

The flash-based Microsemi Smartfusion2 (SF2) FPGA SoC [27][28], whose configuration memory is immune to SEU, was chosen as the main FPGA for the RCU2. Consequently, most of the stability issues seen in Run1, which can be traced back to the SEUs in the configuration cells of the RCU1 main FPGA [21], can be avoided in Run2. The PCB components of the RCU1 that has been proved functional in Run1 were considered to be reused. The components where no radiation tolerance related to the LHC environment was documented, including the SF2 FPGA, have been characterized and tested with several irradiation campaigns (discussed in Chapter 4).

² Developed and simulated by Christian Lippmann (christian.lippmann@cern.ch)

1.4 Primary objective and main contribution

The primary target of this Ph.D. project has been to study the radiation tolerance of the RCU2. This has involved irradiation testing of individual elements, work on improving the design, programming, connecting the elements and to run final tests.

1.4.1 Radiation Tolerance

Several irradiation campaigns have been performed to evaluate the radiation tolerance of the RCU2. In this thesis, these tests have been divided in two steps. In the first step, radiation sensitivity of different aspects in the SF2 FPGA has been characterized (section 4.2) and all the hardware interfaces (section 4.3) on the RCU2 have been tested. The tests revealed potential issues and appropriate actions were implemented afterwards.

As a second step, a full system-level test of the RCU2 including the hardware, the firmware and the software was done under radiation in a situation close to normal operation. Stability issues regarding data readout and status control were observed. Actions to minimize the impact of these issues were later taken (section 4.4).

Based on the results of these irradiation tests, the cross-sections for different failure types have been extracted. While taking all the 216 RCU2 plus 4356 FECs into consideration, Mean Time Between Failure (MTBF) numbers for those failure types have been estimated for Run2.

1.4.2 Design, Integration and Test of the RCU2

Firmware design. Originally, firmware modules that realize the readout algorithms (Readout Module) were planned to be inherited from the RCU1 FPGA design. Nevertheless, several engineering drawbacks were encountered while porting the firmware from the Xilinx FPGA to the SF2 FPGA (discussed in section 3.3.1). Therefore, the author of this thesis developed the first version of the Readout Module (section 3.3.2) for the RCU2. It is a new design but it inherits most of the concepts used in the RCU1 FPGA design. This module has been used in the system-level irradiation tests (section 4.4).

System integration and test. The integration and testing of the RCU2 is also an important task covered by this thesis. This task has been performed in two steps. Firstly, the functionality of the hardware interfaces has been verified (section 5.1). Secondly, all the firmware modules together with the Linux system have been integrated on the SF2 FPGA SoC. Then, the stability of the system has been tested and a benchmark of the readout speed has been performed (section 5.3). Two versions of the RCU2 system have been used in the tests: the second prototype (section 3.3.2) which has been used in the system-level irradiation test and the commissioning version³ (section 3.3.3) that has been commissioned at the TPC. In addition, several designs for dedicated tests have also been developed by the author of this thesis (Chapter 5).

1.5 Outline of the thesis

This thesis is structured into six chapters including this introduction chapter.

Chapter 2 describes the radiation environment of the TPC detector in Run2, the radiation effects on the RCU2 and the basics of the irradiation tests (selection of facility, dose calculation and SEU rate prediction).

Chapter 3 describes the RCU2 design. To start with, the hardware design is introduced. This includes the choice of the RCU2 main FPGA, the hardware interfaces, the Radiation Monitor and the ALTRO bus backplane. Furthermore, the firmware development on the RCU2 main FPGA has been discussed. This thesis focuses on the development of the modules that realizes the readout algorithms. In addition, functionality and structure of the other modules have been briefly discussed. And finally, the software design is presented.

Chapter 4 discusses the irradiation tests of the RCU2, which is the main contribution of this thesis. First of all, the test facilities are introduced and compared. Furthermore, characterization of the SF2 FPGA and test of the hardware interfaces are discussed. Finally, system-level irradiation tests of the RCU2, including the hardware, firmware and software, are presented. According to the test results, the expected error rate of various failures in the radiation environment of LHC Run2 is estimated and corresponded mitigation actions have been proposed. In addition, an evaluation of the Radiation

³ Author of this thesis was not involved in the development of this version of firmware but performed the integration and test of the whole system (section 5.3).

Monitor is also presented. The chapter ends with a discussion on the radiation mitigation techniques in the FPGA fabric SRAMs and registers.

Chapter 5 discusses integration, validation and commission of the RCU2. Firstly, tests performed on the RCU2 prototype are discussed. This includes the stress tests on the RCU2 hardware prototype before mass production and validation on the RCU2 with second prototype of firmware before the system irradiation campaign. Secondly, integration and verification of the final RCU2 design and preparation of the mass installation are presented. A special focus has been put on the stability and the readout rate. Finally, commission of the RCU2 is discussed.

Chapter 6 concludes the thesis and describes the ongoing and planned work.

2 Radiation effects on the RCU2

Electronics which are exposed to a radiation environment (e.g. space, LHC) will potentially be affected by radiation effects. This also applies to the RCU2 which is exposed to the environment of the TPC. Therefore, it is of significant importance to know how these radiation effects are induced, how they affect the RCU2, and how to predict the rate of the radiation induced errors for LHC Run2.

2.1 Interaction of particle with matters⁴

In ALICE, heavy (Pb-Pb) and lower mass (e.g. pp) particles are collided to produce primary particles in high density. These particles can be divided into two categories: charged particles and neutral particles. Many of these particles interact with the absorbers and structural elements of the experiment, which produce hadronic and electromagnetic showers. The cascade of secondaries poses a radiation load on the electronic devices and consequently causes potential damages.

2.1.1 Charged particles

Driven by the Coulomb force, which is the attraction or repulsion among particles due to the electric charge, charged particles interact with the atoms while passing through the material. In these interactions, the charged particles lose and transfer energy to the atoms through several processes: elastic or inelastic scattering with atomic electrons, elastic or inelastic scattering with nuclei, Bremsstrahlung, Cherenkov radiation, etc. Which of these processes dominate the energy loss depends on the energy, velocity, mass and charge of the particle as well as the properties of the material it collides with. For example, most of the energy loss in the interactions of heavy charged particles is through the non-elastic collisions with the atomic electrons in the material [29]. This process of energy loss is also called stopping power.

Stopping Power⁵: The stopping power (S) for a charged particle is defined as the differential energy loss ($-dE$) for this particle within the material divided by the corresponding differential length (d_x)

⁴ This section is based on reference [29] and [30] if not otherwise stated.

⁵ The unit of stopping power is keV/um or MeV/cm.

of the path:

$$S = -\frac{dE}{dx} \quad (2.1)$$

The stopping power is also called the rate of energy loss for a particle. It depends on the energy and type of the radiation as well as the property of the materials. The classical expression that describes the stopping power is the *Bethe-Bloch Formula* and is written as

$$-\frac{dE}{dx} = \frac{4\pi e^4 z^2}{m_0 v^2} NB \quad (2.2)$$

where

$$B \equiv Z \left[\ln \frac{2m_0 v^2}{I} - \ln \left(1 - \frac{v^2}{c^2} \right) - \frac{v^2}{c^2} \right] \quad (2.3)$$

with the following definitions:

v = velocity of the charged particle

z = charge of the particle in unit of e

N = number density of absorber atoms

Z = atomic number of absorber atoms

m = electron rest mass

e = electron charge

I = effective excitation and ionization potential of the absorber

B = stopping number (atomic number scaled for stopping)

The *Bethe-Bloch Formula* is valid for all types of charged particles provided their velocity remains large with respect to the velocity of the orbital electrons. Only the first item in the stopping number (B) is sufficient for the non-relativistic charged particles ($v \ll c$). The stopping number (B) varies slowly with particle energy and is proportional to the atomic number (Z) of the absorber. Thus, the general behavior of stopping power can be inferred from the residual multiplicative factor. For a given non-relativistic particle, $-dE/dx$ varies as $1/v^2$, or inversely with particle energy.

The stopping power consists of two components, the mass collisional stopping power and the mass radiative stopping power. The former is resulted from the interactions of particles with orbital electron (i.e. atomic ionizations and excitations) and the latter is resulted from the interactions of particles with nucleus (i.e. bremsstrahlung production).

Linear Energy Transfer (LET)⁶ is defined as the average energy (dE) locally deposited into the material by a charged particle of specific energy traversing a distance of d_x and it is written as dE/d_x . LET is closely related to stopping power except that it does not include radiative loss of energy (bremsstrahlung) and delta rays. For heavy charged particles, stopping power and LET are nearly equal; for beta particles, the delta-rays and the bremsstrahlung are not included in LET.

2.1.2 Neutral particles

Neutral particles are uncharged and therefore do not interact with matters by means of the Coulomb force. Neutrons and photons (gamma and X-rays) are typical neutral particles and the processes in their interactions are different.

Neutrons: Neutrons do not interact with atomic electrons, but interact with the nuclei of the atoms. Since the size of the nuclei is quite small compared to the whole atom, the probability of neutron interaction is rather low. Hence, neutrons could penetrate a long distance in the absorbing material before any interaction takes place. Processes in the nuclear interactions of neutrons highly depend on the available energy level. For example, the interactions of high energy neutrons will produce secondary radiation products (charged particles, neutrons, fission fragments, etc.), most of which transfer energy to the material through ionizing.

Photons: Photons are electromagnetic radiation with no rest mass, no charge, and travels at the speed of light. Energy of photon is in linear proportion to frequency (f) with the Plank's constant (h), and it is written as $E = hf$. All of the photon interactions lead to a partial or total transfer of the photon energy to the electron energy. There are three main processes in the energy transfer, which are photoelectric effects, Compton scattering and pair production.

⁶ LET is strictly defined in terms of energy divided by distance, e.g., MeV/cm. However, since the energy lost is directly proportional to the density of the material traversed, it is useful to divide the LET by the density of the material. Therefore, the units of LET are also typically expressed as MeV·cm²/mg.

2.2 Radiation Effects related to the RCU2⁷

Electronic devices which are exposed to a radiation environment are expected to experience two categories of radiation effects: SEEs and cumulative effects. The cumulative effects include TID effect and Displacement Damage.

Single Event Effects (SEEs): SEEs originate from the energy deposited by single particle through ionization in a given sensitive volume, occurring in a short time. It is a transient effect and occurs stochastically. For the LHC environment, the charged hadrons⁸ and the neutrons cannot deposit enough energy through direct ionization to induce a SEE. Instead, they generate a SEE through nuclear interaction with the material of the devices (section 2.1). Due to their statistical nature, SEEs are characterized with their probability of occurrence, which depends on the specifics of the radiation environment and the properties of the devices.

Total Ionizing Dose (TID) effect: TID is the progressive build-up of charges due to trapped holes in the insulating layers of MOSFET and BJT devices. Through ionization, electron-hole pairs are generated in the material along the particle track. Due to the high mobility, electrons can escape the oxide easily. In contrast, holes have a lower mobility, and can gradually be trapped in the dielectric. The TID effect may lead to parametric degradation (e.g. threshold voltage shift in MOSFET, current gain decreases in BJT) and eventually cause functional failure of the devices. TID is characterized by the maximum dose that a device can absorb before it no longer behaves within a given expected specification.

Displacement damage: Displacement damage is a non-ionizing effect and refers to the atomic displacement in the crystal lattice. If an incident particle can transfer enough energy to an atom in the crystal lattice by an elastic or inelastic collision, the atom can be knocked free from its lattice site and onto interstitial site. Displacement damage can change the electrical characteristics of certain components, e.g. reduced gain of bipolar transistors.

⁷ The basic principle of the radiation environment and the radiation effects on electronics is based on [31][32] and [33], if not otherwise stated.

⁸ In particle physics, a hadron is a composite particle made of quarks held together by the strong force in a similar way as molecules are held together by the electromagnetic force.

2.2.1 Radiation environment of the RCU2

Two quantities that are commonly used to describe a radiation environment are fluence rate and absorbed dose. Fluence rate⁹, also named flux, is defined as particles incident on a unit sphere or cross-sectional area per unit time. The time integrated flux is called fluence. Absorbed dose¹⁰, abbreviating as dose, is the mean energy imparted to per mass material minus the energy leaving the mass, either directly or through nuclear transformation. For a given number of particles, fluence and dose are correlated but not equivalent. In addition, another quantity named 1 MeV neutron-equivalent fluence¹¹ is normally used to express displacement damage.

Monte Carlo particle transport calculations shows that the radiation load in terms of the flux of fast hadrons on the TPC electronics locating in the innermost positions (worst-case) is estimated to be 0.8 kHz/cm², for the interaction rate of 8 kHz during Run1 [36]. Scaling the interaction rate to 30 kHz of Run2, the expected radiation load for Run2 will be 3.0 kHz/cm² [18]. This number is similar to what a satellite would encounter while traveling through the South Atlantic Anomaly [37] and ~ 0.6 million times of the radiation flux in ground level [38]. With such a significant number, SEEs are therefore expected to occur on the RCU2. For the 3 years running period of Run2, the total dose and the 1 MeV neutron-equivalent fluence are estimated to be less than a few krad and in the order of 10¹⁰ cm⁻², respectively [18]. These number are not significant as the onset for the typical failures occur when the dose is over 10 krad and the 1 MeV neutron-equivalent fluence is above 10¹¹ cm⁻² [39], so neither TID effect nor displacement damage is a big concern for the RCU2. However, TID effect on the SF2 FPGA still need to be considered, since it has previously been observed to lose its programmability at a low total dose level [40].

In the radiation environment of TPC, it is the high energy protons, neutron and pions (Energy > 10 to 20 MeV) that dominates the origin of the SEEs [36]. These high energy hadrons can be considered to be equally effective in their capability of producing SEEs [17]. In addition to these hadrons, there is also a considerable number of other particles (e.g. photons and electrons) that

⁹ Unit of flux is particles/cm²/s or a shorten version p/cm²/s [41].

¹⁰ Unit of dose is Gray and rad, where 1 rad = 0.01 Gray [41].

¹¹ Unit of 1 MeV neutron-equivalent fluence is cm⁻² [41].

contributes to the TID effects [36].

2.2.2 Single Event Effects (SEEs)

The family of SEEs is quite wide. The main members of SEEs that may occur on the SF2 and the hardware interfaces of the RCU2 are discussed in the following sub-sections.

Single Event Upset (SEU)

A SEU refers to a single bit-flip in the content stored in memory elements, which is induced by a single energetic particle strike [33]. A SEU will be provoked in the sensitive node if the energy deposited by single particle exceeds the critical charge of a storage element. SEUs are stochastic errors, which can happen in electronic devices at any time during their operation in radiation environment. SEUs are non-destructive and can be corrected by re-writing the memory elements.

The main FPGA of the RCU2 is the Microsemi SF2, which integrates a FPGA fabric, a Microcontroller Subsystem (MSS) [42] and several lanes of high speed serializer/deserializer (SERDES) interfaces [43]. Details regarding the SF2 are presented in section 3.1.1. The FPGA fabric of the SF2 is flash-based and its configuration cells are considered immune to SEU [28]. However, SEUs are still expected to occur in the SRAMs and in the flip-flops [44]. In addition, SEUs may also occur in the MSS of the SF2 and the hardware interfaces of the RCU2. If SEUs occur in critical bits, they may lead to Single Event Functional Interrupt, which will be discussed later in this section.

Multiple-Bits Upsets (MBU)

MBU refer to two or more bits in the same data word being flipped due to single radiation event. Because each bit-flip is actually a SEU, a MBU can be for simplicity treated as a set of SEUs.

In the SRAMs of the SF2, occurrence of MBUs is expected to be low for two reasons. Firstly, there is a physical distance between adjacent bits in the 65 nm manufacturing technology used for SF2 [45]. Secondly, as illustrated in Figure 2-1, logically adjacent data bits are physically separated in the memories. As a result, MBUs on physically adjacent bits can be divided into SEUs in several logical data words [45]. This dramatically reduces the probability of that an MBU will result in uncorrectable errors. In the irradiation tests performed by Microsemi in late 2014, the SF2 chips

were exposed to heavy ions at LET levels up to 90.3 MeV-cm²/mg and no MBUs were observed in the Large SRAMs, Micro SRAMs and flip-flops [45]. Therefore, no specific tests regarding MBUs will be discussed in this thesis.




Figure 2-1 Mitigation of MBUs in memory cells in SF2 [45].

Single Event Transient (SET)

If a single energetic particle hits the combinatorial logic in an integrated circuit, the deposited energy will give origin of a momentary pulse, which is defined as a SET [33]. In some cases, the transient pulse could propagate along the logic path until it is latched by some memory elements (e.g. SRAMs, flip-flop, latch), resulting in the changes on their output. As the clock frequency increases, the probability that SET will cause an upset in combinatorial logic increases. The ability of the SET to propagate and their probability of being captured by memory elements increases as well [34][35].

In this thesis, SET on the SF2 has been studied in terms of its probability of occurrence, varying the complexity of the combinatorial logics and the operating frequency of sequential logics (discussed in section 4.2.4).

Single Event Latch-up (SEL)

A spurious current pulse induced by a single highly energetic particle passing through the sensitive regions of electronic components could bias the parasitic PNP structure in the CMOS transistors and create a short between the power lines. In the JEDEC standard, this abnormal high-current state is defined as SEL [33]. SEL is potentially destructive and may cause permanent damage to electronic devices. If the device is not permanently damaged, a power cycle is required to the recover it back to normal operational situation. Several tests dedicated for SEL on the SF2 have therefore been

performed in this thesis (section 4.2.1). In the hardware interfaces of the RCU2, non-destructive SEL could induce Single Event Functional Interrupt and destructive SEL will lead to permanent damage.

Single Event Gate Rupture (SEGR)

Transient gate leakage current induced by a single particle strike can lead to a high electric field. In the presence of this electric field, a subsequent conducting path through the gate oxide of a MOSFET can be built. This phenomenon is defined as SEGR [33], to which the power MOSFET in OFF state is susceptible. The SEGR is a destructive effect and can cause permanent damage on the devices.

SEGR is expected to occur in MOSFETs operating with supply voltage higher than 100 V [46]. Therefore, it is not expected to take place during the normal operation of the RCU2, whose supply voltage is only 4.3 V and 3.3 V.

Single Event Functional Interrupt (SEFI)

Soft errors are non-destructive errors induced by a single energetic particle strike, which includes SEU, MBU, SET (if latched) and non-destructive SEL [33]. SEFI is defined as the reset, the lock-up, or the detectable malfunctions caused by a soft error on electronic components [33]. In case a SEFI occurs, the component will usually restore its operability automatically. Notably, a SEFI is usually related to the SEUs in the control elements of the components, and the underlying reasons for a SEFI can be difficult to find due to the complexity of the devices involved.

On the SF2, the components that have been tested in this thesis are the phase-locked loop (PLL) [49] and the MSS. For the PLL, lock signal was used as the monitor. For the MSS, SEFI was identified through observing the operating status of the software running on it. The tests for SEFI in the PLLs and the MSS are discussed in section 4.2.5 and section 4.4.2, respectively.

As discussed in the above sub-sections, hardware interfaces of the RCU2 are also expected to suffer SEFIs. Therefore, corresponding tests have been performed (section 4.3) in this thesis. In these tests, SEFIs were identified through observing the functionalities of these interfaces.

2.2.3 Total Ionizing Dose (TID) effect

All the electronic components in a radiation environment are expected to absorb a dose during their life-time. Any device that is sensitive to TID effect is expected to fail if it has been exposed to the maximum limitation. As a flash-based FPGA, the SF2 is potentially sensitive to TID effects [40] [47] , which may appear in two parts:




Figure 2-2 Structure of floating gate transistor in flash-based FPGA [47]

In the floating gate: Charge loss in the floating gate will lead to the shift of the threshold voltage and then the flips of stored bits [50]. Figure 2-2 demonstrates a typical flash structure, in which the bit value is stored as a charge on the floating gate. Electron-hole pairs are initiated while highly ionized particles passing through the transistors. There are three factors that reduces the threshold voltage of the floating gate [50] [51]: (1) injection of holes into the floating gate, (2) trapping of holes into the tunnel oxide and (3) emission of electrons over the poly-silicon/oxide barriers.

In the CMOS transistors: In the CMOS transistors, charges are progressively built up in the bulk of the oxides and the Si-SiO₂ interfaces due to the trapping of holes. Screening or enhancing of the charges in the gate electric field of the transistors lead to the shift of threshold voltage and the increase of leakage current. In the recent technology with thin oxide where transistors are isolated from each other, the trapped holes can invert the interface at the edges of the transistors, then create open leakage paths between the drain and the source or between adjacent devices [47]. Radiation sensitivity of the CMOS transistors is in positive proportion to the thickness of the oxide due the capability of trapping holes.

In flash-based FPGAs, the charge pump, which provides a higher programming voltage, is the block that is most vulnerable to TID effect, because it has high operation-voltage and uses the transistors with thick oxide [48]. Therefore, in this thesis TID effect on the SF2 has been characterized in terms of functionality and programmability (section 4.2.6).

TID effects of the hardware interfaces is not a concern, since the commercial CMOS components could stand a dose in the order of 10 krad [39], which is higher than the total dose (a few krads) that the RCU2 is expected to absorb in Run2.

2.2.4 Summary

Both the SF2 and the hardware interfaces on the RCU2 are expected to suffer radiation effects in the TPC. For the SF2, the following radiation effects should be considered: (1) SEU in the SRAMs and the flip-flops, (2) SEFI in the MSS and the PLLs, (3) SEL of the FPGA and (4) TID effects of the whole SF2 chip. For the hardware interfaces, the sensitivity of SEFI should be investigated.

2.3 Irradiation tests

The radiation tolerance of the RCU2 is evaluated through a set of irradiation tests. Selecting a proper test facility is the prerequisite to ensure the reliability of these tests. For SEEs, the rates of occurrence extracted from the tests are used to predict the corresponding error rate in Run2. For TID effects, the SF2 is exposed to a certain amount of dose and then checked in terms of functionality and re-programmability. This section discusses how to select the test facilities, how to calculate the dose and how to estimate the rate of SEE induced errors.

2.3.1 Selection of test facilities¹²

Different kinds of radiation effects should be tested with different radiation sources. For the SEEs testing, mono-energetic proton beams with energy over 60 MeV can be used. While testing for SEL, mono-energetic proton beams with the energy higher than 200 MeV is recommended. The proton

¹² Most of the recommendations regarding how to select test facility are from [54].

beams are good candidates because they are widely provided by many facilities. The reason for preferring mono-energetic beams is that the cross-section can be measured at a precise value of energy. For testing the TID effects, a ^{60}Co source is commonly used. Due to the limited number of test devices and time-slots, in our campaigns, the SEEs, the SEL and the TID effects need to be tested simultaneously. In this case, mono-energetic proton beams of 60 to 200 MeV can be used as the radiation source.

All the major tests were performed at the Svedberg Laboratory in Uppsala [52], with a mono-energetic proton beam of 180 MeV. In addition, several preliminary tests were carried out at the Oslo Cyclotron [53], with a mono-energetic proton beam of 25 MeV. The tests at the Oslo Cyclotron were intended to make a first screening of the candidate components. In addition, one supplementary test was performed at Nuclear Physics Institute [55] in Prague, with a mono-energetic proton beam of 35 MeV. The results of these test are discussed in detail.




Figure 2-3 Test facility of the Oslo Cyclotron. (a) Layout of the Oslo Cyclotron [53]. (b) Test setup and positioned beam center.

The Oslo Cyclotron

The Oslo Cyclotron is operated by the Department of Physics, University of Oslo. It is an accelerator in Norway that provides ionized particles for basic research. The Oslo Cyclotron can accelerate protons to the range from 2 MeV to 35 MeV. In our tests, a proton beam of ~ 25 MeV was used to irradiate the devices. Subfigure (a) of Figure 2-3 shows the layout of the Oslo Cyclotron, which includes the inner hall, where the MC-35 cyclotron is located, the outer hall, where the electronics are tested. Before performing the tests, the central position of the beam needs to be found in two

steps. Firstly, radiation films are exposed so that spread of the beam can be seen according to the area turned black. Afterwards, a radiation monitor¹³ connecting to a X-Y positioning system is moved within the area of the beam spot to find the position (beam center) where highest number of SEUs is produced (referred to the counts on the scintillator¹⁴ locating at fixed position). Subfigure (b) of Figure 2-3 demonstrates an example setup at the Oslo Cyclotron, in which the beam center is pointed by a laser on the reflection of the devices in a mirror.




Figure 2-4 Test facility of the Svedberg Laboratory. (a) Layout of the test area [52]. (b) Setup of our test.

The Svedberg Laboratory

The Svedberg Laboratory is operated by the Uppsala University in Sweden. With the Gustaf Werner cyclotron, it provides a proton beam ranging from 20 MeV to 180 MeV, with a beam spot diameter from 0.4 cm to 20 cm. The beam from the cyclotron is controlled to exit into the blue hall where the electronics are tested. The devices were exposed to the proton beam of ~ 180 MeV in our test. In contrast to the Oslo Cyclotron, beam dosimetry service, including calibration of the beam, is provided by the Svedberg Laboratory. Subfigure (a) of Figure 2-4 shows the layout of the test area. Subfigure (b) of Figure 2-4 shows the setup of the system level irradiation tests (discussed in section 4.4).

¹³ Details regarding the radiation monitor can be found in [56].

¹⁴ Details regarding scintillation counts can be found in [57].

2.3.2 Dose calculation

LET (dE/dx) is used to calculate the dose received by the tested devices. While passing through the devices, the proton beam transfers part of its energy to the material. The energy transfer (ΔE_{beam}) can be calculated with equation 2.4, where $\rho_{silicon}$ is the density of the silicon, $dE/dx_{silicon}$ is the LET of proton in silicon and $dx_{silicon}$ is length of particle path in the silicon.

$$\Delta E_{beam} = (\rho_{silicon} * \frac{dE}{dx_{silicon}} (E_{beam}) * dx_{silicon}) \text{ [MeV]} \quad (2.4)$$

Because the beam continuously loses energy along its path, the LET of the beam keeps increasing. However, for a short path as in the devices, these changes on the LET can be neglected. Therefore, equation 2.4 gives an approximation that is close enough to the real value.

As mentioned above, dose is defined as the total energy deposited on per unit mass of the device. Since fluence stands for the total number of particles hit per unit area of the device, the total energy transferred by the protons can be calculated as ΔE_{beam} multiplies the fluence; multiplies the area of the device ($A_{surface}$). Therefore, dose can be simply calculated with equation 2.5, where numerator is the total energy deposited by the beam, and denominator is the mass of the device. Assuming that all the protons pass through the device vertically, length of the path can be treated the same as the thickness of the silicon.

$$Dose (Si) = \frac{\frac{dE}{dx}(E_{beam}) * \rho_{silicon} * dx * fluence * A_{surface}}{A_{surface} * d_{thick} * \rho_{silicon}} \approx \frac{dE}{dx_{silicon}} (E_{beam}) * fluence \left[\frac{MeV}{mg} \right] \quad (2.5)$$

The SI-unit of dose is Gray (Gy) [41], which is 1 Joule of energy absorbed in a kilogram of matter (J/kg). When it comes to radiation of electronics, another unit that is often used is radiation absorbed dose (rad) [41]. The relation between Gy and rad can be seen in equation 2.6 [41].

$$1 \text{ Gy} = 1 \frac{J}{kg} = 100 \text{ rad} \quad (2.6)$$

$$1 \frac{MeV}{mg} = \frac{1.602E-13 J}{1E-6 kg} = 1.602E-7 \text{ Gy} = 1.602E-5 \text{ rad} \quad (2.7)$$

Equation 2.7 presents the conversion from $\frac{MeV}{mg}$ to rad. Hence, equation 2.5 can be transformed into

equation 2.8, which is used to calculate dose in this thesis.

$$Dose(Si) = 1.602 * 10^{-5} * \frac{dE}{dx_{silicon}}(E_{beam}) * fluence [rad] \quad (2.8)$$

2.3.3 Predicating the rate of SEEs induced errors

In the irradiation campaigns, the tested device is exposed to a certain fluence and the number of experienced SEE induced errors (N_{SEE}) are counted. Device sensitivity to SEEs is characterized as cross-section (CS), which is the probability that a particle strike will induce SEE. The cross-section is calculated by dividing N_{SEE} with fluence, as given by equation 2.9 [33],

$$CS_{device} = \frac{N_{SEE}}{fluence} = \frac{N_{SEE}}{flux*time} [cm^2 / device] \quad (2.9)$$

where the fluence can be simplified to time multiplied by flux for a stable radiation environment (e.g. TPC and test facility).

The uncertainty of the cross-section (σ_{CS}) is dependent on uncertainty of the N_{SEE} ($\sigma_{N_{SEE}}$) and the fluence ($\sigma_{fluence}$). Since the SEEs are random in time and linearly depends on the number of incoming particles, $\sigma_{N_{SEE}}$ can be given by Poisson distribution ($1/\sqrt{N_{SEE}}$). The $\sigma_{fluence}$ depends on the method of fluence calculation that is used in each irradiation campaign (Appendix C). Hence, the σ_{CS} can be calculated with equation 2.10 [29].

$$\sigma_{CS} = \sqrt{\left(\frac{1}{\sqrt{N_{SEE}}}\right)^2 + \sigma_{fluence}^2} \quad (2.10)$$

To characterize and evaluate the radiation tolerance of the tested devices, the MTBF in Run2 of different types of errors need to be calculated. This can be done with equation 2.11, which is derived from equation 2.10 by setting the N_{SEE} to be 1, applying the radiation flux in Run2 and using the cross-section for certain kind of error (failures).

$$MTBF = \frac{1}{flux_{Run2} * CS_{device}} \quad (2.11)$$

3 The RCU2

The RCU2 design consists of three layers: hardware, firmware and software. The firmware has gone through three versions: the first prototype, the second prototype and the commissioning version. The author has contributed to all these versions. For the first prototype, the author studied the feasibility of porting the RCU1 design into the RCU2. For the second prototype, the author was the main contributor to the Readout Module as well as the responsible for the system integration. For the commissioning version, the author tested the firmware in hardware. The author also proposed the optimizations to improve readout speed.

This chapter starts with an overview of the RCU2, in which the hardware design, including the main FPGA that host the firmware, and the software design are described. Afterwards, the overview of the firmware design is given, where the Readout Module is discussed in more detail.

3.1 RCU2 overview

The development of the RCU2 was initiated in April 2013, and 216 RCU2s with new backplanes were supposed to be installed at the end of 2014, just before the start of LHC Run2¹⁵. The restricted time-frame implied that the existing TPC readout electronics needed to be reused as much as possible. All the cabling for the hardware interfaces and power supply should remain as it was for the RCU1. Besides, the RCU2 still uses the ALTRO bus to communicate with the FECs, as changing this was considered to be too ambitious imposing a high risk given the available time to finish the project. Although being similar in appearance, the RCU2 has the following major improvements with respect to the RCU1¹⁶:

- (1) The ALTRO bus is split into four branches instead of the two branches for the RCU1, which ensures at least a doubling of the readout speed.
- (2) In the RCU1, the functionalities are distributed on the three PCBs: the RCU motherboard, the SIU card, and the DCS card. The RCU2 is one single PCB which improves the operational

¹⁵ Eventually the RCU2 was installed at the beginning of 2016 due to delays in the project.

¹⁶ The details of RCU1 can be found in [8] and [20].

robustness.

- (3) The flash-based Microsemi SF2 FPGA SoC [27][28] replaces the SRAM-based main FPGA in RCU1. Since the configuration cells of the SF2 are immune to SEU, radiation tolerance of the RCU2 is expected to be improved.
- (4) Bandwidth of the DDL link is increased from 1.280 Gbps to 3.125 Gbps. This is generally done by implementing a new DDL protocol and using the SERDES interface on the SF2.

Figure 3-1 shows the overview of the RCU2, which is generally divided into the readout path and the control path. The readout path starts with the TTC interface, which splits the TTC signal coming from the local trigger unit into the TTC clock and the TTC data. The firmware implemented in the SF2 decodes the TTC data, reads the event data from the FECs via the ALTRO Bus backplane and processes the captured event data. The DAQ interface ships the processed data to the ALICE DAQ.




Figure 3-1 Overview of the RCU2

The control path is centralized with the ARM processor in the SF2 and three off-chip DDR3 memories (discussed in section 3.1.7). The ARM processor has access to various PCB components

on the RCU2 as well as the firmware in the SF2. The TTCrx chip (in the TTC interface), the Small Form-factor Pluggable (SFP) transceiver (in the DAQ interface), the on-board ADCs (voltage, current and temperature) and the Hardware ID EEPROM¹⁷ are connected to the SF2 through the PC® bus [58]. The Radiation Monitor (RadMon) and the SPI Flash Memory [59] are connected to the SF2 through the SPI interface. The former monitors the radiation environment of the RCU2. The SPI Flash Memory [59] hosts the Linux embedded system and is used in In-System Programming [60]. The DCS interface, functioning as a bridge, interfaces the software to the ALICE DCS.

The front and the back of the RCU2 are shown in Figure 3-2 and Figure 3-3, respectively. Here the most important building blocks are highlighted. All these parts are commercial devices and they are not specially designed to be used in a radiation environment. This justifies the need for irradiation testing of these components. To better understand the impact and results of these tests (Chapter 4), these building blocks are discussed in detail in this section.




Figure 3-2 RCU2 Board (front side)

Clocking and reset scheme is essential for all digital designs, also for the RCU2. There are five clock sources on the RCU2: three on-board oscillators of 100 MHz, 156 MHz and 125 MHz, the 25/50 MHz oscillator in the SF2 [61] and the 40 MHz TTC clock. Both the 25/50 MHz oscillator

¹⁷ Each RCU2 has a unique serial number, which is stored in the Hardware ID EEPROM (Electrically Erasable Programmable Read-only Memory).

and the 40 MHz clock are used by the firmware in the SF2. The flash-based FPGA (ProASIC3 [62]) in the RadMon (refer to Figure 3-1) uses the 100 MHz clock as its system clock. The 156 MHz clock and the 125 MHz clock are dedicated for the DAQ interface and the DCS interface, respectively. At power-up, the power-on reset device gives a global reset to the SF2, where dedicated reset signals are generated for various PCB components and the firmware modules in the SF2. The clocking and reset scheme of the major building blocks is discussed in detail in their corresponding sub-sections.




Figure 3-3 RCU2 Board (back side)

3.1.1 The RCU2 main FPGA¹⁸

To withstand the significant radiation load for Run2, the flash-based Microsemi SF2 FPGA SoC has been chosen as the main FPGA of the RCU2. Schematic of the SF2 is shown in Figure 3-4. Besides its SEU immune configuration cells, the reason for preferring the SF2 is that it integrates a FPGA fabric, a microcontroller with various peripherals (the MSS) and several lanes of high speed SERDES interfaces.

The FPGA fabric of the SF2 is utilized to implement the RCU2 firmware. Due to the similar

¹⁸ Information about the SF2 is from reference [27] and [28] if not otherwise stated.

functionalities, the RCU2 firmware is estimated to use the same order of FPGA resources as the RCU1 firmware. According to Table 3-1, the SF2 on the RCU2, which is M2S050-FG896, should be able to provide enough design resources for the RCU2 firmware. The final resources count might end up differently since the logical resource maybe implemented and utilized differently for the SF2 and the Virtex2 Pro. Table 3-1 still gives a reasonable estimation on the feasibility of implementing the RCU2 firmware in the SF2.




Figure 3-4 Schematic layout of the SF2 FPGA SoC[63]

	Logic cells ¹⁹	RAMs	User IOs
RCU1- Virtex-II Pro XC2VP7	11,088	792 Kb	396
RCU2- M2S050-FG896	56,340	1314 Kb	377
RCU1 firmware	8,719	595 Kb	248

Table 3-1 Resources comparison between the RCU1 main FPGA, the RCU2 main FPGA and the RCU1 firmware [15] [20]

The MSS hosts the Linux system that replaces the functionalities of the DCS board on the RCU1. This ensures that the RCU2 can be backward compatible with regards to the layered structure,

¹⁹ Each logic cell contains 4 LTUs (look-up table) and 1 DFF (D flip-flop)

services and protocols used in the existing DCS. The Linux system runs on three off-chip DDR3 SRAMs, utilizing the hardcore DDR2/3 controllers of the MSS. The SERDES interfaces are used by the DAQ and the DCS interface, which are discussed in section 3.1.3 and section 3.1.4, respectively.

When the RCU2 project was initialized, the SF2 was available as engineering samples only and it was very promising on paper. Hence, there were lots of interest for this device at CERN, including for the RCU2 project. However, no available information existed on its radiation tolerance in the radiation environment of CERN. This made the irradiation testing of this device important beyond the scope of the RCU2 project. These irradiation tests are discussed in detail in section 4.2.

3.1.2 TTC Interface

The TTC signal contains two channels of TTC data, Channel A and Channel B, and the TTC clock. Channel A carries the L0 and the L1 triggers. Channel B conveys the serialized commands and trigger information. More information on the TTC signal can be found in [8].

The TTC interface receives and recovers the BiPhase Mark encoded TTC signal transmitted from the local trigger unit. The preferable device to recover the clock and the data was the TTCrx chip [13]. It implements the Clock and Data Recovery (CDR) algorithms for the TTC interface and is specially designed for the radiation environment of the LHC. In addition, the TTCrx chip has been tested in many irradiation campaigns [64] and, more importantly, it has been proved to be stable on the RCU1 in Run1. However, the TTCrx was out of production and just a limited amount was available, so it was important to consider alternative solutions. Two solutions were proposed and tested for the RCU2: (1) A commercial CDR IC - ADN2814 [65], and (2) a customized CDR module [66] that is implemented inside the FPGA fabric of the SF2. Both the TTCrx solution and the two alternative solutions have been tested in the irradiation campaigns. These tests are discussed in detail in section 4.3.1.

3.1.3 DCS Interface

The RCU2 communicates with the ALICE DCS through Ethernet. Due to the magnetic field in ALICE, transformers in the Ethernet Interface will not operate correctly. The design for the analog

part is adopted from the DCS board in the RCU1, because it contains no magnetic components and has been proved to be quite reliable in Run1. The digital part of the DCS interface is shown in Figure 3-5. It is constituted of the Marvell 88E1111 Ethernet PHY [68] and the Ethernet Module in the SF2. Data transmission between the Marvell PHY and the Ethernet Module is through a Serial Gigabit Media Independent Interface (SGMII), which aims at reducing the usage of the I/O pins of the SF2. The reasons why the Marvell PHY was chosen are that it supports SGMII protocol and uses only a few IOs on the SF2.




Figure 3-5 Digital part of the DCS Interface

The Ethernet Module comprises the MSS MAC Ethernet [69] and the SERDES Interface. The MAC Ethernet is a hardware peripheral provided by the MSS in SF2. The SERDES interface is a high-speed serial interface, which is used to serialize and de-serialize the data for high-speed serial transmission. It includes a serialize-deserializer and some peripheral modules to supports different communication protocols [43]. In the Ethernet Module, the SERDES interface is configured to interact with the MSS MAC Ethernet via ten-bit interface through Extra-long Physical Coding Sub-layer (EPCS) interface. As the only block that is driven by the external 125 MHz oscillator, the SERDES interface controls the clock distribution in the DCS interface. It provides the reference clock of the PLL that generates clocks for the MSS MAC Ethernet and the Marvell PHY. Additionally, the lock signal of this PLL is used as the reset of the PHY.

3.1.4 DAQ Interface

Data transmission between the RCU2 and the ALICE DAQ is through the DDL2 link, whose terminal on the RCU2 is named DAQ interface. As demonstrated in Figure 3-6, the DAQ interface is constituted of the SFP optical transceiver and the SERDES interface in the SF2 FPGA. The Readout Module communicates with the SERDES core via the custom VHDL modules that realize the DDL2 protocol and the EPCS interface. Firmware implementation of the DDL2 protocol is presented in section 3.2.4. The DDL2 link operates with a bandwidth of 3.125 Gbps, which meets the requirements of the upgrade proposal.




Figure 3-6 The DAQ Interface

3.1.5 Radiation Monitor (RadMon)

On the RCU1, the SEUs in the configuration memory of the SRAM-based main FPGA are continuously detected and corrected by the reconfiguration network during normal operation. This provided an intrinsic possibility for online monitoring of the SEUs. On the RCU2, the main FPGA is the flash-based SF2 FPGA and no SEUs are expected to occur in its configuration cells. Thus, a new radiation monitoring solution was needed to still provide this service.

The RadMon on the RCU2 consists of a flash-based FPGA (ProASIC3) and four 8 Mb Cypress SRAMs. The RadMon is based on the design from [56] and [72]. The FPGA writes a known pattern into each SRAM, reads it back, compares it and counts the number of differences, which are sent to the MSS through SPI interface, and then to the online monitoring system. The four SRAMs lead to some advantages compared to the solution of radiation monitoring on the RCU1. To start with, the sensitivity of the RadMon is expected to be increased by a factor of ~ 50 . This is because: (1) the number of sensitive bits is increased about 10 times from the 3 Mbits in RCU1 to the 32 Mbits in RCU2 and (2) the SEU cross-section of the sensitive bits of the RCU2 RadMon is about 5 times higher than that on the RCU1. Furthermore, variation among devices on each RCU2 can to some extent be evened out.

The new RadMon has been characterized at the Svedberg Laboratory in Uppsala (section 4.5), and it has behaved as expected. Details about its upgrade and the analysis of SEUs measurements in LHC Run1 can be found in [70] and [71].

3.1.6 ALTRO bus backplane²⁰

For RCU2, the backplane used for the RCU1 is electrically split from two branches into four branches. The two connectors for the RCU1²¹ were decided to be kept in the same position on the RCU2, which made the tests during the development easier. While considering the positions of the two new connectors (Branch AO and Branch BO), which need to fit to all the six readout partitions in each TPC sector, two solutions were proposed and prototyped (see Figure 3-7): (a) the all-in-one solution and (b) the adapter card solution.

The all-in-one solution was quite attractive because it is mechanically similar to the backplane used for the RCU1 and is easier to install on the detector. However, it was impossible to match the termination for both the outer branch and the inner branch, because the routing trace of the outer branch passes the termination of the inner branch. As the only working solution, the adapter solution was eventually selected, even if it required one extra board that made the assembling at the TPC

²⁰ The backplane was designed under the supervision of Anders Oskarson at University of Lund.

²¹ Branch AI and Branch BI on the RCU2.

slightly more challenging.




Figure 3-7 ALTRO bus backplane: (a) all-in-one solution, (b) adapter card solution [74]

3.1.7 Software design²²

Figure 3.20 shows the architecture of the RCU2 software design, which comprises the Linux system and the booting code (the bootstrap application and the Uboot). The booting code is stored in the embedded nonvolatile memory, while the Linux system in addition to the configuration bit-stream of the SF2 FPGA used for In-System Programming are stored in the external SPI flash memory.

Linux system

The 32-bit Linux system is running on the ARM Processor in the MSS and is uploaded from the flash memory to the three 16-bit DDR3 memories at power-on. Two of these DDR3 memories separately store the upper half and lower half or the 32-bit words of the Linux. The third one provides the parity bits for Single Error Correction and Double Error Detection (SECCDED) mechanism [75], aiming at improving the radiation tolerance. With several software programs running on it, the Linux platform bridges the hardware and the firmware to the ALICE DCS. The software program can be sorted into two logical entities, the FEE Server and the device drivers.

²² Taku Gunji (Taku.Gunji@cern.ch) is the main contributor to the software design on the RCU2.




Figure 3-8 Architecture of RCU2 Software Design

The FEE Server is a DIM²³ server implementation that runs on the RCU2. It provides the status of the RCU2 and FECs, such as the voltage, the current, the temperature and the number of SEUs. Additionally, it also receives commands to configure and control the FEE. The FEE Server comprises the FEE Server core and the Control Engine [76]. The FEE Server core is responsible for publishing services and receiving commands. This core itself is device-independent and it uses threads for device-dependent functions. It controls the executions of these threads. The FEE server was ported from the version running on the DCS board on the RCU1. The Control Engine is responsible for the device-dependent functions. It accesses the hardware through a set of Linux device drivers.

Device drives are running on the Linux platform to control and access the RCU2 hardware. Some of the drivers (e.g. SPI, I²C®) are provided by the Linux kernel and some are specially designed for the RCU2, including the ADC driver, the RCU2 Bus master and the RadMon driver. Particularly,

²³ DIM [73] is a client/server based inter-process communication system used in the ALICE DCS. The servers publish services, which are normally a set of data. The clients subscribe to these services and send commands to the serves. Once being subscribed, the services are subsequently updated by the servers at a fixed time interval or whenever the condition changes.

the Ethernet driver is customized from the driver provided by the Linux system.

Bootimg code

The bootimg code includes the Bootstrap application and the Uboot, and is stored in the embedded nonvolatile memory of the SF2 MSS. The bootimg process of the RCU2 software is shown in Figure 3-9. The Bootstrap application runs directly after power-on. It accomplishes the initialization of the SPI flash, the clocks, the DDR3 memories²⁴, the Ethernet, etc. The Uboot runs after the Bootstrap application. It reads the environment parameters (speed of the DDL2 link, configuration of the Linux, etc.) from the SPI flash memory and boots up the Linux system. Additionally, it is in charge of the In-System Programming of the SF2 FPGA.




Figure 3-9 Flowchart of the software booting process

In-System Programming (ISP)

ISP is needed since the RCU2 cannot be accessed during normal operation. As recommended in [60], ISP of the SF2 can be performed in three ways: (1) through JTAG, (2) through SPI port in slave mode and (3) through SPI port in master mode. Programming via JTAG was not considered since it was impractical given the physical location of the RCU2 in the ALICE cavern. Programming

²⁴ The Linux system is uploaded to the three DDR3 memories.

via the SPI port in slave mode needs an external device, usually a micro-controller, to act as a SPI Master. This was neither elegant nor practical on the RCU2.

Therefore, the ISP of the SF2 was decided to be implemented using the SPI port in master mode. As shown in Figure 3-9, the ISP of the SF2 are accomplished in the following steps:

- (1) The Uboot reads the bit-stream from the SPI flash memory
- (2) The Uboot calls the service in the system controller [77] to authenticate the bit-stream.
- (3) The Uboot calls the services in the system controller to program the device with the authenticated bit-stream.
- (4) The Uboot calls the services in the system controller to verify the device.

Noticeably, the prerequisite of executing ISP is that the firmware programmed to the SF2 is different from the bit-stream in the SPI flash memory.

3.2 Smartfusion2 (SF2) firmware overview

The top-level architecture of the FPGA design in the SF2 is shown in Figure 3-10. The firmware modules can be divided into the Readout Node and the DCS Node. In the Readout Node, the Trigger Receiver receives, decodes and processes the trigger sequence from the TTCrx chip. Afterwards, it generates triggers to the Readout Module and FECs. After receiving the triggers, the FECs starts to buffer the event data. Then the Readout Module reads data from the four branches of FECs in parallel, checks its integrity, processes and packages it. At the final stage, the packaged data is shipped to the DDL2 Module and then to the SFP transceiver.

The DCS Node includes the Monitoring and Safety Module and the Ethernet Module. The Monitoring and Safety Module is responsible for monitoring the status of the FECs through the Front-end bus (customized I²C® bus) [78]. It reads values when the FEE server orders it to. The Ethernet Module enables a communication up to the higher logical levels of the ALICE DCS. It is constituted of the IP cores of the SF2 and has already been described in section 3.1.3.

The firmware has been developed in stages and gone through several versions. In this section, a short overview of the FPGA fabric design is given, highlighting the main differences between the different versions. Then all the modules except the Readout Module are briefly discussed. The

Readout Module will be discussed in more details in section 3.3.




Figure 3-10 Overview of the RCU2 firmware

3.2.1 Clocking and reset scheme

As shown in Figure 3-10, there are four clock domains in the firmware: the 40 MHz TTC clock, the 80 MHz system clock, the 125 MHz clock and the 156 MHz clock. The 80 MHz system clock is generated by the system PLL (the MSS PLL in Figure 3-10) using the 25/50 MHz oscillator in the SF2. Processing data at 80 MHz, instead of 40 MHz as in RCU1, is done to utilize the doubled number of branches and the increased bandwidth of the DDL link. The reasons for a local clock instead of the 40 MHz global TTC clock to generate the 80 MHz system clock are as follow: First of all, the event readout in all the 216 readout partitions do not need to be synchronized, because the data from different readout partitions can be sorted offline. Furthermore, the local trigger unit performs a clock switchover at the beginning of each physics fill in LHC and the TTC clock is very unstable during this period. Even though no effects induced by the unstable clock has been seen on the RCU1, it cannot be disregarded as a potential pitfall. Using a local clock can avoid the potential problems caused by the clock switchover. In addition, the TTC clock is not always guaranteed during maintenance or development in the lab, so a local clock is anyhow needed.

In the Trigger Module, the sub-modules for decoding the TTC data work at 40 MHz and other modules for providing CDH words to the Readout Module work at 80 MHz. The 40 MHz TTC clock is synchronous with the bunch crossing (40 MHz) in the LHC and delivered to all the 216 RCU2s. By using this TTC clock to decode the TTC data and generate the sampling clock²⁵, the data sampling on the FECs can be done with equal reference to the time of interaction for all RCU2s in the TPC.

The clocking scheme of the Ethernet Module is presented in Figure 3-5. In the DDL2 Module, the sub-modules of the DDL2 protocol (section 3.2.4) work at 80 MHz, and the clocking scheme of the other parts is shown in Figure 3-6. All the other firmware modules, including the Readout Module, work at 80 MHz. The Readout Module uses the 80 MHz clock to generate the 40 MHz readout clock for the FECs.

In the second prototype of firmware, the lock signal of the system PLL is used as a reset to all the firmware modules, as shown in Figure 3-10. In the commissioning version of the firmware, a Reset Controller, as shown in Figure 3-13, is implemented. This is because the system PLL was observed to lose lock in the irradiation tests (discussed in section 4.4.1). In the Reset Controller, the PLL lock is used as a global reset only at power boot-up. Afterwards, the reset to dedicated modules can be given from DCS.

3.2.2 Firmware versions

The firmware has been developed through two prototype versions and then into the commissioning version used during the installation of the RCU2 in the TPC. The installed commissioning version has later been upgraded several times.

First prototype²⁶: The firmware was originally proposed to be ported from the RCU1 with some customized changes for the RCU2. However, porting this code from a Xilinx platform to a Microsemi platform soon proved to be far more challenging than what could be foreseen (discussed in section 3.3.1). Hence, it was decided to design a second prototype based on the ideas from the

²⁵ The sampling clock is used by the FECs to sample the event data. It can be configured to 2.5 MHz, 5 MHz, 10 MHz and 20 MHz. In this thesis, the sampling clock is 10 MHz if not otherwise stated.

²⁶ The author studied the feasibility of porting the firmware from a Xilinx platform into a Microsemi platform.

RCU1 FPGA design. The first prototype will be briefly discussed in 3.3.1. Before designing the second prototype, a simple ALTRO bus master was designed to test the basic write and read functions from the RCU2 to the FECs (discussed in section 5.1.1).

Second prototype²⁷: The second prototype of the RCU2 firmware was released for the system-level irradiation test at the Svedberg Laboratory in April 2015. The main change between the second prototype and the first prototype is the Readout Module, which performs essential readout functions. It handles trigger sequence from the Trigger Module, reads event data from the FECs and ships the formatted data packages to the DDL2 Module. All other modules have been developed according to the original plan for the first prototype. As of the irradiation campaign, the Trigger Module, the Ethernet Module and the Monitoring and Safety Module were fully functional. The DDL2 Module was realized with a data-rate of 2.215 Gbps. The author is the main contributor of the Readout Module, which is the largest part of the FPGA design. Therefore, the Readout Module in this version will be discussed in depth in section 3.3.2.

Commissioning version²⁸: The commissioning version of the RCU2 firmware was used in the system-level verification (discussed in section 5.3.3). It has been running in a stable manner since January 2016, only with minor bug fixes and some feature updates (discussed in section E.5).

This version of the Readout Module has been developed by team members at CERN after the system-level irradiation tests, since the focus of my work is mainly related to radiation testing and characterization of the RCU2. This Readout Module will be discussed in section 3.3.3, focusing on its difference and improvements with respect to the second prototype.

All other modules were inherited and upgraded from the second prototype. One major improvement is that the DDL2 Module is running at the speed of 3.125 Gbps.

²⁷ The author contributes to the development of the Readout Module and the system integration of the firmware. Responsibilities of the other team member: Torsten Alt (Goethe-Universität) - Readout Module. Attiq Ur Rehman (University of Bergen) - Readout Module. Ernő David (Cerntech, Budapest Hungary) - Ethernet Module. Fillipo Costa (CERN, Switzerland) - DDL2 Module. Johan Alme (University of Bergen) - Trigger Module and Monitoring and Safety Module.

²⁸ The author participated in the verification and optimization of the firmware. Responsibilities of the other team member: Alt Torsten (Goethe-Universität) - Readout Module and system integration. Ernő David (Cerntech, Budapest Hungary) - Ethernet Module. Fillipo Costa (CERN, Switzerland) - DDL2 Module. Johan Alme (University of Bergen) - Trigger Module, Monitoring and Safety Module and SIU Interface. Stefan Kirsch (Goethe-Universität) – Readout Module and system integration.

3.2.3 Trigger Module

The Trigger Module is generally ported from the RCU1 design. Its detailed functionalities can be found in [20] and [79] and a summary is given below.

It receives, processes and validates the two channels (Channel A and Channel B) of TTC data decoded by the TTC interface. Any trigger sequence is accepted only if the L1 trigger, the L1 message and the L2 trigger arrives within a certain timing region after the L0 trigger. In this case, a local L1 accept trigger and a local L2 accept trigger are issued to the FECs. The L1 accept trigger starts the data acquisition of an event in the ALTRO chips and the L2 accept trigger locks the sampled data in the multi-event buffers [9] in each ALTRO channel. If any trigger in the sequence violates the timing requirements, the data acquisition will be aborted.

The Trigger Module also generates the Common Data Header (CDH) words and passes them to the Readout Module. Some fields of the CDH words where the information should be filled by the Readout Module are padded with zeros.

Another important feature of the Trigger Module is to provide the 10 MHz sampling clock for the FECs. This sampling clock is used for the data acquisition in the ALTRO chips. As mentioned above, the sampling clock is derived from the 40 MHz TTC clock, which is synchronous with the LHC bunch crossing. This ensures the data sampling can be done with equal reference regarding the time of interaction for all the 216 RCU2s.

3.2.4 DDL2 Module

As shown in Figure 3-6, the DDL2 Module includes the SERDES interface in the SF2 MSS and the firmware modules of the DDL2 protocol²⁹. The SERDES has been discussed in section 3.1.4. This section focuses on how the DDL2 protocol is implemented.

As illustrated in Figure 3-11, the DDL2 protocol transmits data between the Readout Module and the SERDES Interface in both directions. The data path from the Readout Module to the SERDES

²⁹ Most of the information is based on [24].

interface starts with the FEE interface that admits 32-bit data words. Following which, the Data Classifier generates 2-bit code for each 32-bit word to indicate its type. At last, the Framing Module splits each 32-bit word into two 16-bit words, under the regulation of the 2-bit type code, because the SERDES Interface provides data bus up to only 20 bits. In the reverse direction, commands and data coming from the SERDES Interface are received by the Data Receiver and the Command Receiver, respectively. In the meanwhile, a Cyclic Redundancy Check block verifies the received information to detect potential errors. After that, data declared to be valid is pushed into a FIFO and then shipped to the Readout Module by the FEE interface.




Figure 3-11 DDL2 protocol blocks

The TLK_2051 emulators that stand in-between the SIU block and the SERDES interface are implemented to realize the functionalities of the TLK2051 transceivers [80] from Texas Instruments, which recover the DDL link and keep the DDL link aligned. These transceivers were employed on the RCU1 and they are replaced by the emulators on the RCU2.

3.2.5 Monitoring and Safety Module

The Monitoring and Safety Module is shown in Figure 3-12. It is designed to monitor the physical parameters of all the FECs in a readout partition. Each FEC contains a 10-bit, 5-channel ADC with an on-chip temperature sensor and an I²C® interface. One channel provides the data of the temperature sensor, while the other four provide the values of the analogue and digital voltages and currents measured at the input of the FEC [81]. The Board Controller can be configured to

continuously read the ADC and store the values in its register bank.

The Linux system controls the Monitoring and Safety Module to read or write to the Board Controller by setting a command register and read the result register. The Monitoring and Safety Module accesses the Board Controller via the front-end control bus. The front-end control bus interface is divided into four blocks, each of which communicates with its corresponded branch.




Figure 3-12 Monitoring and Safety module sub-modules

3.2.6 RCU2 Bus

The SF2 MSS and the DAQ Interface communicates with the fabric modules through the RCU2 Bus system. As presented in Figure 3-13, the RCU2 Bus system comprises one Bus Master, one Bus Arbiter and several Bus Slaves. The RCU2 Bus Master is a slave hooked on the Advanced Peripheral Bus (APB). It splits its assigned address span into several segments and allocates each of these segments to an individual RCU2 Bus Slaves. The Bus Arbiter is used to select whether the data goes into the slaves is from the RCU2 Bus Master or the Message Handler. The messages from the DAQ interface is used to configure the RCU2 or the FECs.

In addition, the RCU2 Bus master contains the RCU2 Commander and the Reset Controller. The RCU2 Commander is a simplified version of the Instruction Sequencer [20] in RCU1. It stores and executes a set of ALTRO commands and RCU1 commands. The DCS can control the Reset

Controller to provide reset to the whole firmware or dedicated module(s). As mentioned in section 3.2.1, the lock signal of the system PLL was used as a global reset in the second prototype of firmware. However, the PLL was found to be not reliable in radiation (discussed in section 4.4.1). Thus, this Reset Controller has been designed.




Figure 3-13 RCU2 bus structure topology

3.3 Readout Module

The Readout Module is the largest firmware module. The basic design requirements are as follow:

- The Readout Module needs to receive the L1 and L2 triggers from the Trigger Module and issue respective L1 and L2 triggers to the FECs. Additionally, it has to make the sampling clock from the Trigger Module differential and send it to the FECs.
- The Readout Module must read event data from all the four branches concurrently. The readout in each branch is done in dedicated order from ALTRO channel to ALTRO channel. There are two options for performing the readout: full readout and sparse readout. These two readout modes have been discussed in detail in [20]. In full readout, all the ALTRO channels are read. In sparse readout, only the channels that contain event data are read. The second prototype supports both these two readout modes. In the commissioning version, the sparse readout was not implemented at first but included later after the readout speed was benchmarked. This is discussed later in section 5.3.3. Additionally, the Readout

Module is also responsible for providing the readout clock for the FECs.

- The Readout Modules is required to encapsulate the data into packages in a dedicated format (see Appendix D) and send the packages to the DDL2 Module. It is beneficial for the efficiency of the data analysis that the data from one RCU is shipped sequentially pad by pad for each padrow, and that a complete padrow of data is received prior to shipping any data from the next padrow. The 128 channels of each FEC are not ordered by pads and padrow, but to match the physical constraints given by the electrical cable connection from the pad-plane to the FEC. However, it is always such that branches match areas of pads on the pad-plane from branch BO (leftmost) to branch AO (rightmost), and that no pads are ever connected in an interleaving mode between the branches. The chunk based readout algorithm makes use of this fact, and defines a chunk of data as the data ordered by pads for one padrow in one branch. To implement this scheme, one would need to control the order of which the channels are read out per branch, and to store these data in a FIFO including start and end markers for individual chunks. The next step is to ship the data to the DAQ system, which can be solved quite elegantly by reading a full chunk from the FIFOs in a round robin scheme one branch at the time.

As early mentioned, the development of the firmware has gone through three versions: the first prototype, the second prototype and the commissioning version. In this section, the important features of each version are discussed.

3.3.1 First prototype

Given the success of the RCU1 in Run1, the Readout Module was originally proposed to be ported directly from the RCU1, with some relatively small modifications to match the RCU2 hardware: (1) it needs to be expanded to four branches, (2) Xilinx IPs need to be replaced with Microsemi IPs and (3) the chunk-based readout algorithm should be implemented. However, porting the design from the RCU1 to the RCU2 was terminated by several engineering challenges. Firstly, the RCU1 design was dependent on Xilinx core components, which were not completely compatible with Microsemi core components. Furthermore, with the RCU1, it was the first time when a complete full-scale readout system and data acquisition system was integrated for ALICE TPC. The RCU1 was used to debug and verify many system level issues which unavoidably lead to an extensive patch work. Modularity and interfaces between sub-modules in firmware were affected adversely. Reuse and

cleanup of the VHDL code required more effort as compared to rewriting it from scratch with new Microsemi based core components. In addition, the ported Readout Module was verified in functional simulation but it never behaved as intended on the RCU2 hardware.

All the lab-tests with the first firmware prototype were performed on the first PCB prototype of the RCU2. Hence, it could not be excluded that the failures in these tests were caused by hardware issues. A single, simple ALTRO bus master was therefore designed from scratch to verify this (discussed in section 5.1.1). The results were that this bus master could both write and read to the registers in the ALTRO, and it was concluded that the hardware was behaving as intended.

Due to the reasons mentioned above, it was decided that the Readout Module should be fully redesigned while keeping the conceptual design structure of the RCU1.

3.3.2 Second prototype

The second prototype of the Readout Module was designed and released for the system-level irradiation test (section 4.4) at the Svedberg Laboratory in Uppsala in April 2015. It is similar to the commissioning version, both conceptually and on structural level, thus it could be used to evaluate the radiation tolerance of the final system. Different designs may have different radiation tolerance but at least the tests of the second prototype provided a good evaluation of our concept. As shown in Figure 3-14, the Readout Module consists of the ALTRO Interface Module, the Event Readout Manager and the Event Assembler. The Event Manager receives triggers from the Trigger Module and controls the readout, the ALTRO Interface Module interacts with FECs and reads the event data from the ALTRO channels, and the Data Assembler sorts the captured data and sends it to the DDL2 Module. In this prototype, the data transmission is done from channel to channel, but not from chunk to chunk as intended for RCU2. This is because there was no time available before the system-level irradiation campaign for implementing the chunk-based readout algorithm.

- **ALTRO Interface Module**

The ALTRO Interface Module communicates with the FECs and performs the readout of event data. Communication between the ALTRO Interface Module and the FECs includes write and read transactions to the FECs which are realized through a set of specified instructions. To accomplish the event readout process, channel readout (CHRDO) transaction and readout pointer increment

(RPINC) transaction need to be performed. In CHRDO transaction, the ALTRO Interface Module fetches data from each individual channel. In RPINC transaction, the readout pointer of the data buffers inside the ALTROs, where the event data is stored, is incremented.




Figure 3-14 Second prototype of Readout Module

As shown in Figure 3-14, the ALTRO Interface Module is split into four branches, each of which consists of an ALTRO Bus Interface, a Branch Readout Unit, a Memory Controller and two data memories. In addition, it also involves an Event Readout Manager, which acts as an abstraction layer between the Event Manager and the four Branch Readout Units. Further in this section, the functionality of each individual block is discussed in detail.

ALTRO Bus Interface: The ALTRO Bus Interface implements the protocol for the transactions on the ALTRO bus. Contrary to the RCU1 where three modules interact with the ALTRO bus, the ALTRO Bus Interface is the only module in the RCU2 FPGA design that does this. This improves the structure of the FPGA design. As shown in Figure 3-15, the ALTRO Bus Interface includes the ALTRO Interface Controller, the ALTRO Clock Generator, the ALTRO Data Synchronizer and the Trigger Generator.

A set of instructions are supported by the ALTRO chip and the Board Controller on the FECs. With these instructions, the RCU2 can (1) write to a register in a single FEC (normal mode) or the same registers on all the FECs (broadcast mode), (2) read from a single register in a single channel, and (3) control the execution of dedicated activities on the FECs. The instructions used in (3) are also called commands. The RCU2 sends these instructions to the FECs through the ALTRO bus, following the ALTRO bus protocol as described in [9]. The instructions are delivered by either the Branch Readout Unit (during normal operation) or the DCS via the RCU2 bus (during debug or configuration procedure). The task of the ALTRO Interface Controller is to handle the handshake protocol and deal with the potential erroneous situation during handshaking.




Figure 3-15 ALTRO Bus Interface sub-module

Detailed information regarding the instructions and the handshake protocols can be found in [9]. Only the handshake process of the CHRDO command, which will be repetitively used in this thesis, is mentioned below. The CHRDO command is used to read event data from a single channel in the ALTRO chip. It is the same as any Write command and only the address differs. As shown in Figure

3-16, the execution of a CHRDO command includes the execution of a regular Write command and a data dumping process. An ALTRO Write command is done by setting the appropriate word in the data and address field in the command, while the Command Strobe (CSTB) and the Write Enable (WRITE) line is pulled low. The ALTRO acknowledges by pulling the Acknowledge (ACK) line low. In the data dumping process, direction of the ALTRO bus is turned around and the data is transmitted from the FECs to the ALTRO Bus Interface. The assertion of Transfer Strobe (TRFS) indicates the start of the data dumping process and each 40-bit data word is valid on the falling edge of the Data Strobe (DSTB), which is synchronous with the readout clock (40 MHz).




Figure 3-16 Chronogram of the CHRDO command

The ALTRO Clock Generator is the source of the sampling clock and the readout clock for the FECs. It receives the 10 MHz sampling clock from the Trigger Module, and feeds it through a differential buffer to the output pins. It also divides the 80 MHz system clock to generate the 40 MHz readout clock.

The ALTRO Data Synchronizer is designed to capture the 40-bit data words from the FECs. Two options were proposed: (1) using a small dual clock FIFO, which are clocked by the 80 MHz system clock and the DSTB signal, (2) sampling the DSTB signal and the data word with the 80 MHz system clock. The option (1) was used on the RCU1. There is a concern regarding the quality of the DSTB signal since it is not a regular clock. The second option is potentially more stable but its prerequisite is that the timing of the DSTB signal is the same for all the FECs. Subfigure (a) of

Figure 3-17 shows the DSTB signals measured on the FECs locating at the far-end side in each branch. There are small variations in the range of 1 to 2 ns, which are the routing delay on the PCB backplane. Subfigure (b) of Figure 3-17 shows that the delay between the readout clock and the DSTB signal³⁰ is fixed to be around 21 to 22 ns. Depending on these values, the DSTB signal and the data words can be sampled at the falling edge of the readout clock. Considering that the second method is more reliable, this concept was used to implement the ALTRO Data Synchronizer. This sampling method has been verified to be stable through stress tests, which are discussed in section 5.1.2.




Figure 3-17 Screenshot of CHRDO operations. (a) DSTB signals measured on the FECs located at the far-end side of each branch³¹. (b) Timing of the Readout clock and the DSTB signal

The ALTRO Trigger Generator issues triggers to the FECs. Two schemes can be used in trigger generation: (1) bypassing the local triggers from the Event Manager and (2) generating L1 and L2 triggers whose parameters (length, delay, etc.) can be manually configured. The first scheme is default and applied in normal event readout procedure. The second one is designed for the purpose of system test.

Branch Readout Unit³²: Each Branch Readout Unit implements the event readout algorithm of its associated branch. It handles the orders from the Event Readout Manager and controls the ALTRO

³⁰ Both the readout clock and the DSTB signal are measured on the RCU2, so the delay of the complete signal looping path (RCU2-> backplane -> FECs -> backplane -> RCU2) is considered.

³¹ The screenshot was taken in persistent mode. The DSTB was asserted periodically and the signal line was high if the DSTB signal is not asserted.

³² All the instructions motioned in this sub-section are presented in detail in [9].

Bus Interface to read the event data from the FECs. Figure 3-18 shows the sub-modules of the Branch Readout Unit.

Subfigure (a) and Subfigure (b) of Figure 3-19 show the two procedures that operate in parallel in the Branch Readout Unit during a readout process: Data Readout and Address Scanning. In the procedure of Data Readout, the Branch Readout Controller is the main building block. It receives the orders from the Event Readout Manager to start the readout. Each readout process includes the several CHRDO transactions and one RPINC transaction. In each CHRDO, the Branch Readout Controller reads one channel address from the Channel Address FIFO and request the Command Encoder to send a CHRDO command to the ALTRO Bus Interface. Meanwhile, the Branch Readout Controller also request the ALTRO Bus Interface to perform the ALTRO Bus protocol.




Figure 3-18 Branch Readout Unit sub-module

After the CHRDO transactions of all the channels that should be read are completed, the Event Readout Manager requests the Branch Readout Unit to perform RPINC transaction. Empty channels are skipped in sparse readout mode. This will be discussed later in this section.

In the procedure of Address Scanning, two modes of readout are supported: full readout and sparse

readout. In full readout, all the channel addresses listed in the Readout List Memory (ROLM) are sequentially pushed into the Channel Address FIFO and then used in the CHRDO.

In sparse readout, only the channels that contain event data are read. Each bit in the Hit List Memory (HLM) indicates whether the corresponding ALTRO channel contains event data or not. These bits are set before each readout process by executing two dedicated command to the Board Controller: Scan Event Length (SCEVL) and Event Length Readout (EVL RDO). Execution of SCEVL and EVLRDO introduces an overhead of $\sim 91.3 \mu s$, so sparse readout is only beneficial if more than ~ 140 empty channels are skipped³³. During the readout process, the Address Sequencer reads a channel address from the ROLM and uses it as the index to access the HLM. If the channel is not empty, the address is pushed the Channel Address FIFO. Otherwise, the channel address is discarded.




Figure 3-19 Flow chart of the Branch Readout Unit. (a) Data Readout Procedure. (b) Address Scanning Procedure.

Different to the memories in Xilinx FPGA on RCU1, the memories in the SF2 do not support an asynchronous read operation. Thus, it takes two clock cycles to check whether a channel is empty,

³³ The readout time of each empty channel is measured to be ~ 650 ns.

which in total leads to an extra time of $\sim 22 \mu\text{s}$ ³⁴ for reading a single event. Considering that it takes only $\sim 106 \mu\text{s}$ to read an empty event (refer to section 3.3.3), this extra time is significant. To solve this problem, the Address Sequencer continuously reads channel addresses from the ROLM, check the corresponding bit in the HLM, and fill the addresses of the channels that are not empty to the Channel Address FIFO. Because this FIFO has asynchronous output, it can provide channel addresses to the Branch Readout Controller without any delay.

This firmware prototype supports sparse readout, since it uses the concept of the RCU1 FPGA design. Later in the commissioning version, the sparse readout was first removed and later included again. The reason for these decisions will be discussed in detail in section 5.3.3.

Event Readout Manager: The Event Readout Manager controls and monitors the event readout procedure. After being ordered by the Event Manager, it controls all the four Branch Readout Units to start the readout in parallel and wait for their completion. Afterwards, it controls all the Branch Readout Units to execute RPINC and then reports to the Event Manager after the RPINC is finished in all the branches.

Data Buffering: In this prototype, the data transmission is not in chunk-based mode but in linear mode, that is, data from a single channel is transferred round-robin from the four branches, i.e. one channel from Branch AI, then one channel from Branch AO, etc. Thus, data FIFOs are not implemented and two data memories are instead used to buffer the data. Each memory stores the event data that is read from a single ALTRO channel. Being controlled by the Memory Controller, the two memories work in ping-pong scheme. This avoids the ALTRO But Interface from being in idle state, and thereby increases the readout speed.

- **Event Manager**

The Event Manager consists of the Trigger Handler and the ALTRO Readout Controller. If the Trigger Handler receives a trigger sequence from the Trigger Receiver, it will issue L1 and L2 pulses to the FECs and increment the counter that records the number of pending events by one.

If pending event exists, the ALTRO Readout Controller initiate an event readout by commanding the

³⁴ $22.4 \mu\text{s} = 12.5 \text{ ns (clock cycle)} * 2 * 128 \text{ (channels in each FEC)} * 7 \text{ (number of FECs in the largest branch)}$

Event Readout Manager and the Data Assembler. Afterwards, it will inform the Trigger Handler to decrease the counter.

- **Data Assembler**

The Data Assembler reads event data from the data memories in the four branches in a round-robin fashion, converts the 40-bit words into 32-bit words³⁵, encapsulates them into packages, and then ships the packages to the DDL2 Module. Details regarding the RCU2 data format is shown in Appendix C and it is summarized in Figure 3-20.




Figure 3-20 RCU2 Data package structure

Each data package contains 12 CDH words, several segments of payload words, and 8 Trailer words. Each segment is the data from single channel. It contains the Channel Header and the Channel Payload. Correspondingly, the Data Assembler is partitioned in to the CDH constructor, the Payload constructor and the Trailer constructor, all of which are controlled by the Data Assembler Controller. The CDH constructor receives the CDH words from the Trigger Module and fill in the bits that stores the status of the readout. The Payload constructor reads data from the data memories in the four branches and then performs the 40-bit to 32-bit conversion. The Trailer constructor is implemented to generate the trailer words.

³⁵ The data words read from the ALTRO channels are 40-bit and the DDL2 protocol requires the data words to be 32-bit.

- **Discussion**

Compared to that in RCU1, this Readout Module has significant advantages on arranging the control flow of the event readout and modularity. Since the code was written from scratch, all the various patches are removed and the RCU2 design is closer to how the RCU1 design was envisaged originally. The following two features still need to be implemented to finalize the Readout Module:

- (1) The chunked-based readout algorithm, which is required by the data analyzing algorithms on the receiver side.
- (2) Improvement of the handling of the XOFF signal from the DDL2 link. Assertion of the XOFF signal means that the DDL2 link is saturated and the RCU2 must pause the data transmission immediately. In this version of Readout Module, the conversion and transmission are done in frames³⁶, each of which contains three sequential 40-bit words. The XOFF signal can be asserted at any given time, but the data conversion and transmission can only be suspended after a complete frame has been processed. Hence, there is a probability that the data is still being pushed into the DDL2 module even if the DDL2 link is saturated, and this will lead to the loss of data.

3.3.3 Commissioning version

Figure 3-21 shows the Readout Module in the commissioning version of the firmware. The design of this Readout Module is based on the second prototype, where many of the modules are similar. However, some parts are also completely redesigned by the design team in charge. At the time of writing, the design has been operating stably for several months on the TPC, even though a few features are still to be completed. This section discusses the commissioning version of the Readout Module highlighting on the following two major changes:

- (1) The chunk-based readout scheme has been implemented. The channel addresses in the ROLM are listed in dedicated order and a marker is inserted after the address of the last channel in each chunk. A data FIFO is instantiated in each branch to store the chunks of data and the Data Assembler has been modified to read to the four FIFOs in chunk-based round robin. In addition,

³⁶ The least common multiple of 40 and 30 is 120, so each frame contains 120 bits (3 x 40-bit words).

a Channel Formatter has been implemented in each branch to validate the event data that is read from the ALTRO channels, convert the data from 40-bit words into 32-bit words, and push the data into the data FIFO.

- (2) XOFF signal from the DDL2 link is handled in a proper manner.




Figure 3-21 Commissioning version of Readout Module

- **Channel Formatter**

As shown in Figure 3-22, the Channel Formatter is constituted of the Channel Trailer Checker, the Data Memory and the Channel Encoder.

The Channel Trailer Checker decodes the trailer word³⁷ of each data package from the ALTRO channels to fetch the signatures (e.g. number of data words, channel address, etc.), which are then checked against the same type of signatures recorded by the ALTRO Bus Interface. If these signatures are matched, the data is then pushed into the Data Memory and shipped to the Channel Encoder. In the Channel Encoder, the 40-bit words are converted into 32-bit words and then pushed into the data FIFO. For the empty channels that contain no event data, solely the trailer words are processed.

In the second prototype, data words are transmitted to the DDL2 module immediately after they are converted from 40-bit into 32-bit (30-bit data plus 2-bit symbol). As discussed in section 3.3.2, this scheme leads to the loss of data because the XOFF signal is not handled properly. In the commissioning version, the data conversion is done in the Channel Formatter and the data transmission is done in the Data Assembler. The Channel Formatter pushes the data into the dual-port FIFO. Data transmission from the dual-port FIFO to the Data Assembler can be suspended immediately if the XOFF signal is asserted.




Figure 3-22 Sub-module topology of the Channel Formatter

³⁷ Details can be found in reference [9].

- **Data Assembler**

As shown in Figure 3-21, a new module called Channel Data Conditioner has been implemented in the Data Assembler. It reads data from the four data FIFOs in chunk-based round-robin. Additionally, it stops reading data immediately if the XOFF signal is asserted.

- **Branch Readout Unit**

Figure 3-23 shows the Branch Readout Unit, which consists of the Readout Sequencer, the Branch Readout Controller and the Transaction Handler. The Readout Sequencer provides the addresses of the ALTRO channels to the Readout Controller, which controls the event readout process in a single branch. The Transaction Handler delivers instructions to the ALTRO Bus Interface and collects the above-mentioned signatures in each CHRDO process.




Figure 3-23 Sub-module topology of the Branch Readout Unit

The sparse readout functionality was not implemented at first, since it was assumed that the splitting of the backplanes would make it obsolete. The sparse readout is dedicated for the events that contain large number of empty channels, and the discussion below considers the extreme situation, that is, the empty events.

Readout time of a single event in full readout mode and sparse readout mode can be calculated with equation 3.1 and 3.2, respectively, where N_{FECs} means the maximum number of FECs in each branch.

$$RDO\ Time_{Full} = N_{FECs} * 128 * T_{CHRDO} + T_{RPINC} \quad (3.1)$$

$$RDO\ Time_{Sparse} = T_{SCEVL} + N_{FECs} * T_{EVL RDO} + T_{RPINC} + T_{ROLM\ scan} \quad (3.2)$$

	SCEVL	EVL RDO	CHRDO	RPINC	ROLM Scan
Time	~90.6 μ s	~0.775 μ s	~450 ns + 25 ns * N_{words}	~0.475 μ s	$N_{FECs} * 128 * 12.5$ ns

Table 3-2 Execution time of each transaction for RCU1 [20]

In the RCU1, the N_{FECs} equals 13. The $RDO\ Time_{Full}$ and the $RDO\ Time_{Sparse}$ for an empty event ($N_{words} = 0$) are ~750 μ s and ~146 μ s, respectively. Because the fixed busy time of the TPC in Run1 was 300 μ s³⁸, all the readout time that are shorter than this will still be counted as 300 μ s. Still, the full readout consumes a significant time of ~450 μ s more than the sparse readout. In the RCU2, the N_{FECs} is decreased to 7. The $RDO\ Time_{Full}$ and the $RDO\ Time_{Sparse}$ for an empty event drops to ~403 μ s and ~106 μ s, respectively. Both these two values are smaller than the fixed busy time of the TPC in Run2, which is 500 μ s. Therefore, it was concluded that the sparse readout would not be beneficial for the readout speed of the RCU2.

What needs to be emphasized is that the above calculations were based on the RCU1 firmware. However, for the RCU2, a more robust version of ALTRO bus protocol was implemented at the expenses of the time of each transaction. This has a direct implication on the readout speed, as is discussed in section 5.3.3. Thus, both the sparse readout and the ALTRO bus protocol with the same timing values as for RCU1 have later been implemented for the RCU2.

3.4 Summary

In this chapter, the RCU2 has been introduced and the hardware, software and firmware have been

³⁸ Values of busy time in TPC are from the TPC Run Coordinator Chilo Garabatos Cuadrado (chilo.garabatos.cuadrado@cern.ch).

discussed, where the main focus is on firmware. It has gone through two proper versions, the second prototype and the commissioning version. The commissioning version is the final firmware that has been installed and commissioned at TPC. However, the second prototype was the firmware used in the system level irradiation tests (discussed in section 4.4). Since the system-level irradiation test was very important for the development and enhancement of the commissioning version, both firmware versions are discussed in depth in this chapter.

4 Radiation Tolerance of the RCU2

The increased radiation level in LHC Run2 with respect to LHC Run1 requires an improved radiation tolerance of the TPC readout electronics. A flash-based Microsemi SF2 FPGA SoC was therefore chosen as the main FPGA of the RCU2. The SF2 can provide strong radiation tolerance mainly due to its SEU immune configuration cells. Nevertheless, SEUs or SEFIs may still occur in the SRAMs, the registers, the clocking elements (e.g. PLLs), the MSS as well as in the hardware interface of the RCU2. In addition, SEL and TID effects should also be considered. Therefore, irradiation campaigns were needed to investigate the radiation tolerance of the RCU2.

In total seven irradiation campaigns were performed between November 2013 and May 2015. The overview of each irradiation test is presented in Table 4-1. Campaign No.1 and No.2 were performed at the Oslo Cyclotron. Campaign No.3, No.4, No.6 and No.7 were performed at the Svedberg Laboratory. Campaign No.6 were performed at the Nuclear Physics Institute in Prague. Energy of the proton beams was 25 MeV, 180 MeV and 35 MeV, respectively. The RCU2 was tested in three stages. First, the PCB components, such as power regulators, bus transceivers and buffers were tested. Next, the SF2 FPGA, the hardware interfaces and the RadMon were tested separately. At last, a full system test of the RCU2 including the backplanes and FECs were performed.

The PCB level components were mainly tested for TID effects at the Oslo Cyclotron and no major problems were detected. These tests will not be discussed in this thesis but an overview can be found in [83]. This chapter focuses on the irradiation tests of the SF2 FPGA, the hardware interfaces, the RadMon and the RCU2 system test.

To characterize and evaluate the radiation tolerance of the tested devices, the MTBF in Run2 of the different types of errors has been calculated with the cross-section³⁹ extracted from the tests. The MTBF in Run2 presented is the worst-case estimation, that is, the radiation load in terms of the flux of fast hadrons on the RCU2s that populate in the innermost locations (3.0 kHz/cm²) has been used and all the 216 RCU2s plus 4356 Front-end Cards (FECs) have been counted in. In Run2, the RCU2 is expected to perform no worse than the RCU1 in Run1, in which the longest data-taking session

³⁹ If not otherwise stated, the uncertainty of the fluence and the number of SEE (N_{SEE}) is 15% and $\frac{1}{\sqrt{N_{\text{SEE}}}}$, respectively (discussed in section 2.3.3). The uncertainty of cross-section is calculated with equation 2.7.

in heavy-ion run lasted for ~8 hours. Hence, the MTBF in Run2 of all the RCU2s is compared with this number of 8 hours to compare with the radiation tolerance with the RCU1.

Campaign ID	Time	Test Devices	Test Objectives
No.1	Nov.2013	PCB components	Mainly for TID
No.2	April 2014	PCB components;	Mainly for TID
		ES-FG896	SRAM and TID
No.3	May 2014	FG484 x 2	SEL, flip-flop, SRAM, PLL and TID
		ES-FG896	TID of SF2
		RCU2 (FG896-v1)	TTC Interface with custom CDR; SEL and TID of SF2;
		RCU2 (ES-FG896)	DAQ interface TID of SF2
No.4	June 2014	RCU2 (FG896-v1)	TTC Interface with ADN2814 CDR; Signal quality of optical receivers
No.5	Sep. 2014	RCU2 (ES-FG896)	TTC Interface with TTCrx chip;
No.6	Nov. 2014	RCU2 (FG896-v2)	SEL, PLL and TID of SF2; TTC Interface with TTCrx chip
		RCU2 (FG896-v1) x 2	DCS Interface; TID of SF2
		Radiation Monitor	SEU sensitivity
No.7	April 2015	RCU2(FG896-v2) x2	System-level Test; TID of SF2
		Radiation Monitor	SEU sensitivity

Table 4-1 Overview of the irradiation campaigns (time-wise)

4.1 Tested devices

Four versions of SF2 chips were tested and all these chips are from the series of M2S050 [63]. As

listed in Table 4-1, they are shortened as their packages: the FG896-v1, the FG896-v2⁴⁰, the FG484 and the ES-FG896. FG896-v2 is the version that is used on the final RCU2 design. It is the version with SEL enhanced silicon. FG484 is the same as FG896 but with less IOs [63]. ES-FG896 is the engineering sample version of FG896. The FG484 and the ES-FG896 were available one year in advance to FG896, so they were good candidates to characterize the SF2 in the initial phase of the RCU2 project.




Figure 4-1 Emcraft SF2 Starter-Kit [84]

The SF2 chips were tested either on the Emcraft SF2 Starter-Kit [84] or the RCU2 board. The Starter-Kit is shown in Figure 4-1. It was small, user-friendly and with the correct FPGA (either FG484 or ES-FG896). It was used as the device under test in campaign No.2 and No.3. The Starter-Kits could be ordered and tested a year in advance of the first prototype of the RCU2, which enabled important irradiation tests to be performed prior to moving to hardware production. In addition, these Starter-Kits were also very easy to set up with Board Support Packages and Linux packages that could be downloaded from Emcraft. This meant that it was easy to utilize them as part of the test equipment as well, where they were set up to monitor the current consumption, read out test

⁴⁰ The subversions of FG896, i.e. FG896-v1 and FG896-v2 are defined by the author. It is used only in this thesis to identify the chips with SEL enhanced silicon.

data, etc.

The TTC interface, the DAQ interface, the DCS interface, the Radiation Monitor and the whole system were tested on the RCU2 board, which has been described in detail in Chapter 3.

4.2 Characterization of the Smartfusion2 (SF2)

Because the SF2 FPGA may suffer SEEs and TID effects in the radiation environment of LHC Run2, corresponding irradiation tests were therefore performed to characterize its vulnerability to these radiation effects. This section discusses the test procedure and gives an analysis of the test results. First of all, the tests for SEL are discussed. Furthermore, the tests for SEU in the fabric SRAMs, the embedded SRAMs and the flip-flops are discussed. At last, the tests for TID effects are discussed. Stability of the MSS, which was also predicted to be affected by SEEs, was tested as part of the system-level irradiation campaign, which will be discussed in section 4.4.

4.2.1 Single Event Latch-up (SEL) test

A sudden and large increase in the current consumption from the power supply of the SF2 FPGA can be interpreted as the occurrence of SEL. As discussed in section 2.2.2, SEL may occur on the SF2 and cause destructive damage. Therefore, irradiation tests for SEL were performed at the Svedberg Laboratory. The first SEL test was performed with the FG896-v1 in campaign No.3. After we got the FG896-v2 (with SEL enhanced silicon), the SEL test was performed again in campaign No.6. The setup and procedure of these two tests are the same.




Figure 4-2 SEL test setup

- **Test setup and procedure**

Figure 4-2 shows the setup for testing SEL. The power supply of the SF2 FPGA (1.2 V), the charge pump in the programming logic (2.5 V) and the DDR bank (0.5 V) are monitored in terms of voltage level and current consumption. These measurements are performed with three INA226 devices [82] on a monitoring board. The INA226 device is a current and power monitor with I²C® compatible interface from Texas Instruments. All the INA226 devices connect to the same I²C® bus and communicate with a SF2 Starter-Kit [84] through a I²C® Master Module. The Starter-Kit is set up with a Linux Platform and is controlled by the monitoring PC over an Ethernet link.




Figure 4-3 Current consumption of the SF2 FPGA in first SEL test

- **First SEL test**

The first SEL test was performed in campaign No.3, on a RCU2 prototype with FG896-v1. Current jumps were observed on the power supply of the SF2 FPGA. Figure 4-3 shows an example of the monitored current and the voltage level. To study how the SEL rate is effected by the supply voltage of the SF2, three voltage levels, 1.0 V, 1.1 V and 1.2 V, were used in the test. Figure 4-4 shows the cross-section of the current jumps with different amplitude. It was found that both the probability and amplitude of the current jumps can be reduced by lowering the supply voltage. As discussed in section 2.2.2, SEL is triggered by the formation of parasitic bipolar transistors in a CMOS circuit. At lower bias voltage, the gain of the parasitic transistors and the amount of charge collected from an impinging particle both decrease. These factors lead to the observed decrease in the susceptibility

of SEL. However, if the supply voltage is lower than 1.14 V⁴¹, the FPGA operates outside the timing models provided by the Libero SoC software [85] and the timing closure cannot be guaranteed. The measured SEL events were all non-destructive. A power cycle could always recover the SF2 back to normal operation. However, an SEL might cause errors in the PLL (section 4.2.5) and the hardware interfaces (section 3.1.3). Therefore, further study on the SEL was needed to confirm whether the SF2 was suitable for the RCU2.




Figure 4-4 Cross-section of current jumps vs. supply voltage in the first SEL test

- **Second SEL test**

By November 2014, Microsemi released the FG896-v2 with SEL enhanced silicon. The SEL test was performed again on one RCU2 with a FG896-v2 in campaign No.6. In the test, the SF2 chip was exposed to a fluence of 6.63×10^9 p/cm². As shown in Figure 4-5, the current consumption of the SF2 FPGA was stable, which means that no SEL was observed. The small variations on the current comes from normal operation of the firmware. Later in [44], Microsemi characterized that the SEL LET threshold of SF2 for maximum operating voltages at 100°C is determined to be higher than 22.5 MeV·cm²/mg. Devices having threshold LETs larger than 12 MeV·cm²/mg are often

⁴¹ The minimum requirement in the datasheet [27].

assumed to be immune to protons[87]. I.e. SEL in the SF2 was no longer a concern in the radiation environment of LHC Run2.




Figure 4-5 Current consumption of the SF2 FPGA in second SEL test

4.2.2 Fabric SRAM test

In the FPGA fabric of each SF2 there are 69 Large SRAMs (LSRAM) and 72 micro SRAMs (uSRAMs). Each LSRAM contains 1024 x 18 bits and each uSRAM contains 64 x 18 bits. The purpose of testing SRAM is to find the SEU cross-section. Measurements of SEUs in these fabric SRAMs were performed on one ES-FG896 chip in campaign No.2 at the Oslo Cyclotron and on two FG484 chips in campaign No.3 at the Svedberg Laboratory. The procedure of these two tests is the same.

- **Test setup and procedure**

The test setup is shown in Figure 4-6. It includes a SF2 Starter-Kit in radiation and a monitoring PC in the shielded area. On the tested Starter-Kit, all the 69 LSRAMs and 72 uSRAMs are configured to be 18-bit width. The SEU Monitor, which contains a Finite State Machine (FSM) and several registers, detects and counts the SEUs in the SRAMs and sends the numbers to the MSS, from where they are periodically transmitted to the Monitoring PC via an UART link.




Figure 4-6 SRAM irradiation test setup

Before the test, the FSM in the SEU Monitor fills the RAM blocks with a dedicated pattern of data. In the test, the FSM reads one 18-bit word at one time and all the locations in each RAM block and every block on the chip are read sequentially. If any upset is detected, the FSM records the time stamp, increases the corresponding counters, and then re-fills the location to correct the error. In case any SEU occurs in the state register of this FSM, it will be forced back to the state of detecting SEUs. This ensures that the test can continue. For each kind of the SRAMs, there is one address register, which stores the address to be checked, and two status registers, one counts the number of SEUs and the other one records the time stamp when the latest SEU is detected. Triple Modular Redundancy is used to protect these registers against SEUs.




Figure 4-7 SEUs and fluence for the SRAM test in campaign No.3

- **Test results**

In campaign No.2, the ES-FG896 chip was irradiated up to a fluence of 1.04×10^{11} p/cm². In total 2402 SEUs were detected in the LSRAMs and 72 SEUs were detected in the uSRAMs. In campaign No.3, the two FG484 chips were exposed to the fluence of 1.66×10^{11} p/cm² and 1.75×10^{11} p/cm², respectively. On the first FG484, 3362 SEUs were detected in the LSRAMs and 148 SEUs were detected in the uSRAMs. On the second FG484, 3779 SEUs were detected in the LSRAMs and 159 SEUs were detected in the uSRAMs.

The SEUs and the fluence of the first FG484 in campaign No.3 are plotted as a function of time in Figure 4-7. There is a linear dependence between the number of SEUs and the fluence. This is expected and proves the reliability of the test procedure. The plots for the second FG484 at campaign No.3 and the ES-FG896 at the campaign No.2 are shown in Appendix E.1, in which the linear dependency between fluence and the SEU counts also can be seen.

Memory (bits)	Campaign ID	package	Fluence (p/cm ²)	SEUs	cross-section (cm ² /bit)
LSRAM (1271808)	No.2	ES-FG896	1.04E+11	2402	$1.8\text{E-}14 \pm 0.3\text{E-}14$
	No.3	FG484 No.1	1.66E+11	3362	$1.6\text{E-}14 \pm 0.2\text{E-}14$
		FG484 No.2	1.75E+11	3779	$1.7\text{E-}14 \pm 0.3\text{E-}14$
	<i>Microsemi [44]</i>	<i>FG896</i>	<i>1.38E+11</i>	<i>4421</i>	<i>2.5E-14</i>
uSRAM (82944)	No.2	ES-FG896	1.04E+11	72	$8.4\text{E-}15 \pm 1.6\text{E-}15$
	No.3	FG484 No.1	1.66E+11	148	$1.1\text{E-}14 \pm 0.2\text{E-}14$
		FG484 No.2	1.75E+11	159	$1.1\text{E-}14 \pm 0.2\text{E-}14$
	<i>Microsemi [44]</i>	<i>FG896</i>	<i>1.38E+11</i>	<i>95</i>	<i>1.3E-14</i>
eSRAM (524288)	No.7	FG896-v2	4.99E+10	352	$1.4\text{E-}14 \pm 0.2\text{E-}14$

Table 4-2 SRAM test results⁴²

The results of both these tests are presented in Table 4-2. For the same kind of SRAM, the cross-section extracted from each test is similar. Later in August 2015, Microsemi published its test results of the FG896 in [44]. The cross-section of LSRAM and uSRAM reported by Microsemi is at the

⁴² Campaign Microsemi means the results are from Microsemi Inc. These are highlighted with an italic font.

same level as the one extracted from our tests. While estimating the reliability of the SRAMs on the RCU2 (discussed in section 5.3.4), the average cross-section extracted from our tests is used, that is, $(1.7 \pm 0.2) \times 10^{-14} \text{ cm}^2/\text{bit}$ for the LSRAM and $(1.0 \pm 0.2) \times 10^{-14} \text{ cm}^2/\text{bit}$ for the uSRAM.

4.2.3 Embedded SRAM test

Each SF2 also has two embedded SRAMs (eSRAMs) in the MSS, which can be protected by SECDED. The size of each SRAM is 64 KBytes (80 KBytes if SECDED is disabled). On the RCU2, the bootloader and the Uboot of the Linux system are uploaded in one of the eSRAMs on power-up. Thus, the cross-section for SEU in the eSRAM was characterized with one FG896-v2 chip (SEL enhanced silicon) in campaign No.7. The following procedure was used in the test⁴³:

- (1) Enable the SECDED protection on the eSRAM. The SECDED mechanism counts the SEUs and store the number into an internal error counter.
- (2) Write a dedicated pattern to all the addresses of the eSRAM.
- (3) Irradiate the SF2 and read the error counter periodically through a serial link (UART).




Figure 4-8 SEUs and fluence for the eSRAM test at campaign No.7

The eSRAM was exposed to a fluence of $4.99 \times 10^{10} \text{ p/cm}^2$ and 352 SEUs were detected. As shown

⁴³ The test program was written by Torsten Alt (torsten.alt@cern.ch) but the tests were performed by the author.

in Figure 4-8, the increase of the SEUs is linear proportional to the accumulated fluence. This proves the reliability of the test. As presented in Table 4-2, the cross-section for SEU in eSRAM is $(1.4 \pm 0.2) \times 10^{-14} \text{ cm}^2/\text{bit}$, which is similar to the fabric SRAMs.

4.2.4 Flip-flop test

The purpose of the flip-flop test is to estimate the cross-section of SET and SEU. The test was performed on a Starter-Kit with FG484 in campaign No.3 at the Svedberg Laboratory.

- **Test setup and procedure**




Figure 4-9 Flip-flop test setup

The test setup is shown in Figure 4-9. A SF2 Starter-Kit is positioned in the radiation area for testing. Another Starter-Kit and a PC are in the shielded area for monitoring. On the tested SF2, four chains of shift registers are instantiated. The input stream of each chain is alternating '0' and '1'. A window of 4 flip-flops is placed at the end of each chain to monitor the SEU and the burst errors⁴⁴. The

⁴⁴ Burst error refers to more than one bit get flipped, which is normally caused by an error on the clock. The idea of having an output window is from reference [86].

monitoring Starter-Kit reads the output window via 4 dedicated pins and counts the bit-flips.

To increase the sensitivity for SET, some tests include four inverters between subsequent flip-flops. It is expected that more bit-flips can be seen by adding combinational logic and increasing the clock frequency of the register chain [34][35].

- **Test results**

In the test, the length of each register chain was set to be 2500. The design was operated at 40 MHz, 80 MHz and 160 MHz. In addition, inverters were placed in the register chain in every second test of the same frequency⁴⁵. Several SEUs but no burst errors were observed. The test results are listed in Table 4-3. The results show no significant difference between the different test designs and operation frequencies. This could be due to low statistics combined with that the design does not sufficiently enhance the SET sensitivity, e.g. too few inverters. The cross-section extracted from all the tests is in the same level and similar to the number published by Microsemi. Since the RCU2 firmware operates at 80 MHz, the cross-section of the flip-flops on the RCU2 should be comparable to that of the test design with inverters operated at 80 MHz, that is, $(2.6 \pm 0.7) \times 10^{-14} \text{ cm}^2$.

Campaign ID (SF2 version)	flip-flops	Frequency (MHz)	Inverters	Fluence (p/cm ²)	Flips	cross-section (cm ² /flip-flop)
No.3 (FG484)	4 x 2500	40	0	6.21E+10	19	3.1E-14 ± 0.7E-14
	4 x 2500	40	4	1.42E+11	51	3.6E-14 ± 0.5E-14
	4 x 2500	80	0	7.22E+10	30	4.2E-14 ± 0.7E-14
	4 x 2500	80	4	9.91E+10	26	2.6E-14 ± 0.7E-14
	4 x 2500	160	0	1.61E+11	32	2.0E-14 ± 0.4E-14
<i>Microsemi [44] (FG896)</i>	<i>4 x 2000</i>	<i>10</i>	<i>0</i>	<i>1.38E+11</i>	<i>20</i>	<i>1.8E-14</i>

Table 4-3 Flip-flop test results⁴⁶

⁴⁵ The SF2 chip failed to be reprogrammed when we want to test the design with inverters at 160 MHz.

⁴⁶ Campaign Microsemi means the results are from Microsemi Inc. These are highlighted with an italic font.

4.2.5 PLL test

The use of PLLs is unavoidable in a modern digital design. Therefore, the radiation tolerance of the PLLs in the SF2 is critical and must be carefully evaluated. The PLLs were tested twice at the Svedberg Laboratory (campaign No.3 and No.6).

- **First PLL test**

The setup of the first PLL test is shown in Figure 4-10. One SF2 Starter-Kit is put into the radiation area. Another Starter-Kit and a monitoring PC are put in the shielded area. The lock signals of three fabric PLLs and one MSS PLL on the tested Starter-Kit are fed into the shielded Starter-Kit (Monitor Board) via General-purpose input/output. The Monitor Board counts the losses of the lock signal and sends the number to a monitoring PC through serial link. To study the stability of the output clock when a PLL loses its lock, the output clock of the PLL1 is used as the input clock of the PLL2. If the lock of PLL2 is lost every time when the lock of PLL1 is lost, it would be a clear indication of instabilities on the output clock of the PLL1.




Figure 4-10 First PLL test setup in campaign No.3

In this test, two SF2 Starter-Kits with FG484 were irradiated one by one. The first observation was that as long as the PLL1 lost its lock, the PLL2 lost its lock as well. This shows that the output clock of the PLL in SF2 is not reliable when it loses its lock. In the SF2, each PLL can be configured to three modes [49] :

- (1) It holds the output in reset (output low) until it is locked. After which, the output is released and synchronized with the reference clock.

-
- (2) It generates clock before being locked and resynchronizes with the reference clock after it is locked.
 - (3) It generates clock before it is locked and do not resynchronize with the reference clock after it is locked.




Figure 4-11 Output clock of PLL with different configuration when it loses lock.

As shown in Figure 4-11, no matter which mode the PLL is configured to, it cannot generate reliable clock output if it loses lock. If the PLL is configured to mode (1), the output will be low. If the PLL is configured to mode (2) or mode (3), the output clock will be unstable for a few clock cycles. The clocks must be stable a configurable amount of cycles before the lock is retained.

The two Starter-Kits were exposed to a total fluence of $5.38 \times 10^{11} \text{ p/cm}^2$. In total 381 losses were observed on the fabric PLLs and 143 losses were observed on the MSS PLLs. The losses of PLL2 lock caused by the PLL1 error has been excluded from the counts. As listed in Table 4-4, the cross-section of all the tested PLLs is similar. The worst cross-section of the fabric PLL and the MSS PLL

is $(2.6 \pm 0.5) \times 10^{-10} \text{ cm}^2/\text{PLL}$ and $(2.7 \pm 0.5) \times 10^{-10} \text{ cm}^2/\text{PLL}$, respectively. These cross-section numbers are about 100 times higher than the ones later reported by Microsemi. Since some SELs occurred in the test (section 4.2.1), it was considered that some of the observed PLL errors were induced by the SEL in the SF2.

Campaign ID	SF2 Version	PLL ID	Errors	Fluence (p/cm ²)	Cross-section (cm ² /PLL)
No.3	FG484 No.1 (SEL detected)	Fabric PLL1	52	2.42E+11	2.2E-10 ± 0.4E-10
		Fabric PLL2	62	2.42E+11	2.6E-10 ± 0.5E-10
		Fabric PLL3	60	2.42E+11	2.5E-10 ± 0.5E-10
		MSS PLL	65	2.42E+11	2.7E-10 ± 0.5E-10
	FG484 No.2 (SEL detected)	Fabric PLL1	65	2.96E+11	2.2E-10 ± 0.4E-10
		Fabric PLL2	73	2.96E+11	2.5E-10 ± 0.5E-10
		Fabric PLL3	69	2.96E+11	2.3E-10 ± 0.5E-10
		MSS PLL	78	2.96E+11	2.6E-10 ± 0.5E-10
No.6	FG896-v2 (SEL not detected)	Fabric PLL1	4	7.28E+11	5.5E-12 ± 2.9E-12
		Fabric PLL2	5	7.28E+11	6.9E-12 ± 3.4E-12
		Fabric PLL3	4	7.28E+11	5.5E-12 ± 2.9E-12
		MSS PLL	2	7.28E+11	2.8E-12 ± 2.0E-12
Microsemi [44]	M2S090	Fabric PLL	1	1.37E+11	7.3E-12

Table 4-4 PLL test results⁴⁷

- **Second PLL test**

After Microsemi corrected the SEL issue, the PLL test was performed again in campaign No.6 on a FG896-v2. The test setup is shown in Figure 4-10. A RCU2 prototype was put in the radiation area for testing. A Starter-Kit and a PC was placed in the shielded area for monitoring. The RCU2 communicates with the Starter-Kit through an SPI interface. Because the VHDL modules for detecting and counting the loss of the PLL lock signal is also in the radiation area, Triple Modular Redundancy are implemented to protect them against SEUs. The lock signals of three fabric PLLs plus one MSS PLL are monitored. No cascade of PLLs was implemented in the test design, since the aim of the test was to observe the lock signal of each PLL directly.

⁴⁷ Campaign Microsemi means the results are from Microsemi Inc.

The test card was exposed to a fluence of 7.28×10^{11} p/cm² and in total 15 losses were detected. During the test, no SEL was observed on the SF2. The test results are presented in Table 4-4. The average cross-section of the fabric PLL is $(6.0 \pm 3.8) \times 10^{-12}$ cm²/PLL and the cross-section of the MSS PLL is $(2.8 \pm 2.0) \times 10^{-12}$ cm²/PLL. These numbers are similar to the cross-section of the fabric PLL published by Microsemi. This implies that a considerable proportion of the PLL errors in the first test may be induced by the SEL.




Figure 4-12 Second PLL test setup in campaign No.6

4.2.6 Total Ionizing Dose (TID) effects test

As discussed in section 2.2.3, the SF2 FPGA, especially its charge pump, may be sensitive to TID effects [40][47]. During all the irradiation tests, in total 11 SF2 chips were exposed to different dose levels. Due to limited test time and available devices, the tests for TID effects were not very systematic. The programmability and functionality of each SF2 chip was checked immediately after being irradiated.

All the chips were still fully functional after being exposed to a dose from ~0.7 krad to ~48 krad. The observations on programmability are listed in Figure 4-13, in which No.1 to No.5 are ES-FG896 chips, No.6 to No.9 are FG896 chips, and No.10 and No.11 are FG484 chips. The chips that were exposed to more than ~2.5 krad could not be reprogrammed. The chips that received a dose less than ~2.4 krad could be programmed immediately after the tests. Interestingly, one FG896 (No.8 in

Figure 4-13) that was exposed to ~ 5.5 krad failed to be reprogrammed right after the test but retrieved its programmability after two weeks (room temperature). This indicates a possible annealing effect on the SF2. One of the ES-FG896 chips (No.5 in Figure 4-13) was dedicated to TID testing and was irradiated in steps of ~ 0.5 krad. It failed to be reprogrammed after receiving only 2.5 krad.




Figure 4-13 TID effect on the SF2 chip

For 10 years' operation in ALICE experiment, the TPC electronics located in the innermost partitions are estimated to receive a total dose of ~ 1.6 krad, this calculation is based on the radiation load in Run1 [1]. In Run2, the radiation load is about 3.75 times higher and the operation period is planned to be 3 years. Based on these numbers, the RCU2 is estimated to receive less than ~ 1.8 krad. Considering also the possible annealing effect, the TID effects should not be a concern for the SF2.

4.3 Hardware Interface tests

As discussed in section 2.2, the TTC interface, the DAQ interface, and the DCS interface on the RCU2 may experience SEFIs in the radiation environment of the TPC. Therefore, irradiation tests were performed to evaluate their radiation tolerance.

4.3.1 TTC interface test

The TTC interface includes an optical receiver and a CDR module. As discussed in section 3.1.2, the TTCrx chip, which was used as the CDR on the RCU1 and had been proven stable in Run1, was out of production and only a limited number was available when the RCU2 was designed. Therefore, a custom CDR design [66] in the SF2 FPGA and a commercial CDR - ADN2814 [65] were proposed as alternatives. The radiation tolerance of these new CDR solutions needed to be evaluated. In addition, an optical receiver whose radiation tolerance could satisfy the requirement of RCU2 was also missing.




Figure 4-14 Setup of the TTC interface test

Figure 4-14 shows the setup for testing the TTC interface. A local trigger unit feeds clock and triggers with a BiPhase Mark Encoded signal to the TTC interface. The triggers are issued at a configurable rate of 1 Hz, 10 Hz or 1 kHz. The TTC clock recovered by the CDR module serves as the reference clock of two PLLs (PLL1 and PLL2). The output clock of PLL1 drives the Trigger Decoder, which decodes the TTC data and records the errors in the data. The lock signals of both PLLs are monitored by a module using an independent clock. If only a single PLL loses lock, it should be caused by an error in the PLL itself. If both PLLs lose lock at the same time, the reason is most likely related to problems with the recovered TTC clock. A Linux system that operates in the MSS reads all the test results via an APB bus and passes the numbers to a monitoring PC via an UART interface.

The custom CDR and the ADN2814 CDR were tested in campaign No.3 and No.6, respectively. The optical receiver on the tested RCU2 is the Avago HFBR-2316TZ [88], which was chosen based on irradiation tests originally performed by the TTC group at CERN in 2003 [64].

In the test, the following three types of error were expected and monitored: (1) hamming errors in TTC data signal recovered by the CDR, (2) errors related to the decoding of the TTC data signal in the Trigger Decoder and (3) errors in the recovered TTC clock signal.

The number of error type (1) and type (2) were found to be highly dependent on the frequency of type (3) errors and on the rate of the input trigger. There was a fairly low number of hamming or decoding errors when the clock was stable. It is obvious that the increase of the trigger rate will enlarge the volume of the data received by the TTC interface, and consequently induce more errors. Therefore, focus of the test was put on observing the stability of the recovered TTC clock signal, i.e. the lock signals of the two PLLs.

Campaign ID (CDR solution)	Irradiated Device	Fluence (p/cm ²)	Cross-section (cm ²)	MTBF in Run2 (hours)
No.3 (Custom CDR)	SF2 chip	1.92E+11	1.0E-10 ± 0.2E-10 (20 errors ⁴⁸)	4.12 ± 0.64
	Avago HFBR-2316TZ	1.00E+11	6.1E-9 ± 0.9E-9 (607 errors)	0.07 ± 0.01
No.4 (ADN2814 CDR)	ADN2814	2.43E+10	3.7E-10 ± 0.6E-10 (9 errors)	1.16 ± 0.18
	Avago HFBR-2316TZ	2.53E+10	2.4E-9 ± 0.4E-9 (60 errors)	0.17 ± 0.12
	both ADN2814 and Avago HFBR-2316TZ	1.66E+11	3.2E-9 ± 0.5E-9 (538 errors)	0.18 ± 0.03
No.5 (TTCrx)	both TTCrx and PLD-2317TM	5.54E+10	< 1.8E-11 (0 errors)	> 23.7
No.6 (TTCrx)	both TTCrx and PLD-2317TM	4.49E+11	< 2.23E-12 (0 errors)	> 192

Table 4-5 TTC interface test results (PLL lose lock)

⁴⁸ In total 42 losses were detected. But 22 of them were observed when current jumps (SEL) occurred and a power cycled was needed to remove these errors. Hence, they are removed from the statistics.

In the test, the optical receiver and the CDR were irradiated either separately or together. The test results are listed in Table 4-5. In general, the error rate is higher when the optical receiver was irradiated than when the SF2 chip was irradiated. Assuming only the optical receiver is in radiation, the MTBF in Run2 of ADN2814 CDR solution and custom CDR solution is only ~ 0.17 hours and ~ 0.07 hours, respectively. Noticeably, the cross-section of the Avago HFBR-2316TZ extracted from the campaign No.3 and No.4 is different but still in the same order. This might be caused by the variations from device to device. To figure out the reason for these unacceptable error rates, several possible optical receivers for the RCU2 (Avago HFBR-2316TZ, Truelight TRR-1B43-000, Ficer FTPDA-R155-ST and PD/LD PLD-2317TM) were tested in campaign No.4. The low-voltage differential signaling output of each optical receiver was monitored with an oscilloscope. As can be seen in Figure 4-15, radiation effects in the optical receiver could cause glitches on its output. The TTCrx chip can compensate these glitches and ensures that the clock is recovered with a known and adjustable phase. However, neither the ADN2814 CDR nor the custom CDR can handle these glitches in a proper and stable manner.

After the test in campaign No.3, the TTC interface was redesigned with an existing batch of TTCrx chips with the PD/LD PLD-2317TM optical receivers, whose output showed the lowest rate of glitches in irradiation. The new TTC interface was tested both in campaign No.5 and No.6. No error was detected after the TTC interface has been irradiated up to a fluence of 5.54×10^{10} p/cm² and 4.49×10^{11} p/cm², respectively. This confirms the earlier irradiation tests results of the TTCrx chip in [64] and the experience from Run1.




Figure 4-15 Example of radiation effect in an optical receiver

4.3.2 DAQ interface test

The DAQ interface includes a SFP transceiver and a SERDES interface in the SF2. It transmits data from the RCU2 to the ALICE DAQ, so its radiation tolerance is critical to the whole readout system.

The irradiation test of the DAQ interface was performed in campaign No.3⁴⁹. Figure 4-16 demonstrates the test setup. One RCU2 prototype with ES-FG896 is put in the radiation area. In the shielded area, the Xilinx Virtex6 [89] generates 7-bit Pseudo Random Binary Sequence (PRBS) with its IBERT core [90] and sends the PRBS to the RCU2 at the speed of 2.125 Gbps. On the RCU2 side, the received PRBS stream is looped from the SF2 back to the Virtex6. The Virtex6 then compares the PRBS sent to and received from the RCU2, and counts the number of differences in the bit-stream. Before the irradiation campaign, this setup was verified in the lab (room temperature). No error was observed while ~100 TB of data was looped.

Irradiated Device	Total fluence (p/cm²)	Error type	Cross-section (cm²) Errors	MTBF in Run2 (hours)
SF2 chip	3.2E+11	bit-errors	9.4E-12 ± 5.6E-12 (3 errors)	45.7 ± 27.3
		link down	6.5E-12 ± 4.5E-12 (2 power cycle)	68.6 ± 49.6
AVAGO SFP	3.6E+11	bit-errors	1.1E-11 ± 0.6E-11 (4 errors)	38.6 ± 20.1
		link down	8.3E-12 ± 5.0E-12 (1 self-recover; 2 power cycle)	51.4 ± 30.7

Table 4-6 DAQ interface test results

The SFP transceiver on the tested RCU2 is AVAGO AFBR-57D7APZ [91]. It was chosen because it shows the best radiation hardness among all the possible transceivers in the earlier test, which is presented in [92]. The SF2 FPGA and the SFP transceiver were irradiated separately. They were exposed to a fluence of 3.2×10^{11} p/cm² and 3.6×10^{11} p/cm², respectively. Some bit-errors were detected and the data transmission link was observed to go down a few times. In most cases, a power

⁴⁹ The test was made by Filippo Costa (fillipo.costa@cern.ch). The results were analyzed by the author.

cycle was needed to re-establish the transmission link. At one time when the SFP transceiver was irradiated, the link recovered by itself in ~2 to 5 seconds. No errors that can be interpreted as the consequence of the SEL in the SF2 FPGA were observed in the test.




Figure 4-16 Setup of the DAQ interface irradiation test

The test results are shown in Table 4-6. For bit-errors, the MTBF in Run2 of the SF2 chip and the SFP transceiver is ~46 hours and ~39 hours, respectively. For the problems of transmission link, the MTBF in Run2 of the SF2 chip and the SFP transceiver is ~69 hours and ~51 hours, respectively. All these numbers are much larger than the duration of the longest data-taking session in Run1, which is ~8 hours.




Figure 4-17 Test setup of DCS Interface

4.3.3 DCS interface test

The major components of the DCS interface are the Marvell PHY and the SERDES interface in the SF2. The DCS interface is not critical for data-taking, still the probability of losing control to the RCU2 should be minimized.

Radiation tolerance of the DCS interface was tested in campaign No.6. Figure 4-17 shows a sketch of the test setup, which includes a RCU2 in the radiation area, a PC (PC1) located in the shielded area and another PC (PC2) standing in the control room. PC1 controls the RCU2 and PC2 monitors the status of the Ethernet link. PC2 exchanged data packages with the RCU2 at a rate of 1000 packages per second. Before the irradiation campaign, this setup was verified in the lab (room temperature). The Ethernet link was working in a stable manner in the test period of 48 hours and no package was lost⁵⁰.

In the irradiation campaign, the PHY and the SF2 chip were irradiated separately. The PHY was irradiated to a fluence of 1.78×10^{11} p/cm². The Ethernet link went down twice. When the Ethernet was working, no package was lost. As listed in Table 4-7, the estimated MTBF in Run2 of the Ethernet error is ~38 hours, which is quite promising for the RCU2.

However, when the SF2 FPGA was irradiated, the Ethernet link failed quite frequently. The SF2 was irradiated up to a fluence of 2.72×10^{11} p/cm² and the Ethernet link went down 18 times. The cross-section is ~60 times higher than when the PHY was irradiated. Nevertheless, current jumps on the supply voltage of the SF2 was observed in 16 of all these 18 failures. Thus, most likely these 16 failures were caused by SEL. By removing these 16 errors, the cross-section of the decreases to the same order as when the PHY was irradiated.

To further investigate the reason for the high frequency of failures, about 6.9 million data packages were looped inside the MAC of the MSS when only the SF2 was irradiated. In the test, current jumps on the supply voltage of the SF2 FPGA were observed to lead to the crash of the Ethernet link, which confirms our previous assumption.

⁵⁰ The test was designed by Ernő David and Tivadar Kiss. The test was conducted and analyzed by the author.

Irradiated Device	Total fluence (p/cm ²)	Ethernet down (times)	Cross-section (cm ²) Ethernet down	MTBF in Run2 (hours)
Marvel PHY	1.78E+11	2	1.1E-11 ± 0.8E-11	38.2 ± 27.6
SF2 chip	2.72E+10	18	0.6E-10 ± 0.2E-10	0.53 ± 0.14
		2 (removed SEL errors)	7.4E-11 ± 5.3E-11	5.38 ± 4.21

Table 4-7 DCS interface irradiation test results

After the test, SECDDED protection⁵¹ has been enabled on the corresponding eSRAMs in the MAC. In addition, Microsemi corrected the problem regarding the SELs. As a result, it is estimated that no more than one or two Ethernet errors will appear in a data-taking session of 8 hours in the heavy-ion runs in Run2. As a comparison, the requirement for the RCU1 was that no more than one or two errors should occur in a data-taking session of ~4 hours in the heavy-ion runs in Run1 [94].




Figure 4-18 Setup of system level irradiation test

⁵¹ The SECDDED is provided by Microsemi and can be enabled/disabled by the user.

4.4 System level irradiation test

To evaluate the radiation tolerance of the whole RCU2 system, a system-level irradiation test was performed at the Svedberg Laboratory in April 2015 (campaign No.7). Figure 3-10 shows the architecture of the RCU2 design used in the test, which includes the Readout Node and the DCS Node. The second prototype of firmware, which contains no radiation mitigations, was used for the test. The Readout Node and the DCS Node have been discussed in detail in Chapter 3.

As shown in Figure 4-18, the test setup consists of three parts:

- (1) In the radiation area, the RCU2 for testing is connected to four FECs. Besides, a SF2 Starter-Kit is used to monitor the current consumption of the SF2 FPGA on the RCU2. Note that the Stater-Kit is partially shielded.
- (2) In the shielded area, the trigger crate, the DAQ computer with a CRORC [24] and the monitoring PC are located. The trigger crate sends trigger sequences to the RCU2. The DAQ computer receives the data from the RCU2. The monitoring PC provides a serial link to the RCU2 to monitor the status of the Linux system.
- (3) In the control room, all the devices mentioned above are controlled and monitored by the three PCs via LAN.




Figure 4-19 Setup of the system-level irradiation test (without collimator)




Figure 4-20 Setup of the system-level irradiation test (with collimator)

In this test, the RCU2 was exposed to a proton beam of 84 mm x 72 mm (38 mm x 38 mm if the collimator is used) at a flux of $\sim 5.82 \times 10^6$ Hz/cm², which is ~ 190 times higher than the worst-case radiation load in Run2. The RCU2 was receiving and processing triggers, moving data from the FECs to the RCU2 and sending the packaged data to the DAQ computer. At the same time, all available registers in the RCU2 were read back periodically over the DCS interface. Bit-flips in the event data was unfortunately not checked, since there was no time available for implementing the related tools.

4.4.1 Readout stability

Data-taking of the RCU2 was monitored with the trigger rate set to be 10 Hz and two test cases were performed:

- (1) Irradiating both the whole RCU2 and all the four FECs (see Figure 4-19).
- (2) Irradiating only the SF2 chip on the RCU2. This was realized by shielding the other parts of the RCU2 board with a collimator (see Figure 4-20). In this case, the FECs were still irradiated but not completely shielded by the collimator.

The second case was used to observe potential problems internally in the SF2 FPGA and provide a reference to the first case.

The test results are presented in Table 4-8. During the tests, the readout was observed to stop a few times due to three categories of errors:

- (1) The PLL that provides system clock to the FPGA design loses its lock. At the time of testing, the lock signal of the system PLL was directly used as the reset signal of the FPGA design. In case the PLL loses its lock, the whole FPGA design was in reset and the readout stopped.
- (2) The Board Controller on the FEC, which control the driver of the ALTRO bus, is implemented in a commercial SRAM-based FPGA. No radiation mitigation has been implemented at design level, and SEUs occur in the configuration cells can cause the ALTRO bus on the FEC to be in erroneous state and will therefore lead to the stop of the readout.
- (3) The data transmission link goes down and a power cycle is needed to recover the link.

When the whole RCU2 was irradiated, the readout stopped 6 times in total. As presented in Table 4-8, 3 times were due to RCU2 failures and the other 3 times were caused by FEC errors. As to the RCU2 failures, the MTBF in Run2 is estimated to be 7.61 ± 4.53 hours. Since we are considering the worst-case scenario, the actual reliability of the RCU2 should be better than our estimation. Still, this is comparable to the performance of the RCU1 in Run1. Performance of the data transmission link is acceptable. However, cross-section of the PLL error is about 10 times higher than one from the previous PLL test (discussed in section 4.2.5). This observation was confirmed in the case when only the SF2 chip was irradiated.

The difference between the cross-section of PLL in this test and the previous test may be due to the uncertainty caused by low statistics, but it emphasizes the risk of using PLL. As discussed in section 4.2.5, the output clock of a PLL is not to be trusted when it loses lock. Hence, the 100 MHz on-board oscillator was proposed to provide system clock directly to the RCU2 FPGA design.

Irradiated Device	Fluence (p/cm²)	Error type	Errors	Cross-section (cm²/device)	MTBF in Run2 (hours)
RCU2	5.32E+10	PLL error	2	$3.8\text{E-}11 \pm 2.7\text{E-}11$	11.4 ± 8.24
		Transmission link	1	$1.9\text{E-}11 \pm 1.9\text{E-}11$	22.8 ± 22.8
		FEC error	3	$1.4\text{E-}11 \pm 0.8\text{E-}11$	1.13 ± 0.66
SF2	2.24E+10	PLL error	1	$4.5\text{E-}11 \pm 4.5\text{E-}11$	9.61 ± 9.61
		FEC error	2	$2.2\text{E-}11 \pm 1.6\text{E-}11$	0.95 ± 0.68

Table 4-8 Readout stability observations

Due to the large number (4356) of FECs in the TPC, in Run2 the FEC errors are estimated to cause readout error about every hour. The problems of FEC errors existed in Run1 as well. Hence, a

procedure called Pause and Recover⁵² has been implemented in ALICE since Run1 to reconfigure a complete single readout partition without terminating the data-taking session. In addition, the Pause and Recover scheme is also used for fixing the problems of the data transmission link. All the detectors support this Pause and Recover scheme and the following procedure is what the TPC does in a Pause and Recover:

- (1) If the readout is stuck due to FEC errors, the RCU2 will request the central DCS to pause the trigger.
- (2) When trigger is paused, the whole readout partition where the error occurs will be reconfigured.
- (3) Afterwards, the trigger is resumed and the data-taking session could continue without being terminated.

To detect the erroneous situations of the FECs in early stage, the following actions have been implemented in the RCU2. Firstly, the front-end control bus, indicating the status of the FECs, is continuously monitored. Secondly, the handshake procedure between the RCU2 and FECs is monitored. At last, the trailer word of each data package that contains the channel address, the length of the data, etc. is decoded and checked.

The Trigger Reception was working stable, that is, no error was observed on recovered TTC data and clock. This is consistent with the previous test in campaign No.5 and No.6. No stop of readout can be interpreted as being caused by the SEUs in the firmware. There was only limited time for preparing the test design, so the data error was not able to be checked. Anyway, according to the discussion in section 5.3.4, the data errors should not be a big concern for the RCU2.

4.4.2 DCS stability

The Linux system is the most important component for the DCS stability. As above-mentioned, the Linux system runs on the ARM processor of the SF2 MSS and three off-chip DDR3 SDRAMs.

Two test cases were carried out: irradiating the whole RCU2 board and irradiating solely the SF2

⁵² More detailed information about Pause and Recover can be found in section 5.2.7.2 of reference [20].

chip. Two kinds of errors were observed: sometimes the CPU rebooted and sometimes it was frozen. In both cases, the communication to the CPU was lost.

Irradiated Device	Total fluence (p/cm²)	Error type	Errors	Cross-section (cm²)	MTBF in Run2 (hours)
RCU2	4.22E+10	CPU reboot	20	$4.7\text{E-}10 \pm 1.3\text{E-}10$	0.91 ± 0.24
		CPU freeze	5	$1.2\text{E-}10 \pm 0.6\text{E-}10$	3.62 ± 1.71
SF2	2.69E+10	CPU reboot	7	$2.7\text{E-}10 \pm 1.1\text{E-}10$	1.60 ± 0.65
		CPU freeze	3	$1.2\text{E-}10 \pm 0.7\text{E-}10$	3.74 ± 2.23

Table 4-9 DCS stability observation

The test results are presented in Table 4-9. When the whole RCU2 is in radiation, the MTBF in Run2 of Linux reboot and Linux freeze is ~ 0.91 hours and ~ 3.62 hours, respectively. When only the SF2 is irradiated, the stability of the Linux system seems to be a little better because the DDR3 memories were shielded. However, it is hard to draw any conclusion due to the low statistics.

When CPU was frozen, the data taking could continue, but the control to the RCU2 and the FECs was lost. When CPU rebooted, the SERDES in the DAQ was re-initialized during the booting process. Hence, the readout was suspended until the re-initialization was completed.

The reason of these CPU-system errors is most likely related to MBUs in the DDR3 memories. Bit-flips in some places, most likely caused by MBUs, may lead to complete malfunction of the Linux System. In this case, the Linux is rebooted due to kernel panic⁵³. Bit-flips occurring in the other places may lead to the freezes or have no effects on the system until they are accessed.

Considering the high rate of CPU errors, one improvement has been implemented after the test⁵⁴, which separates the readout logic from the CPU. A stand-alone module for initializing the SERDES in the DAQ interface has been designed to replace the default initializing scheme, so that the SERDES will not be reconfigured in case the Linux reboots.

⁵³ A kernel panic occurs as a result of a hardware failure or a software bug in the operating system. The system is in an unstable state and rather than risking security breaches and data corruption, the operating system stops to prevent further damage and facilitate diagnosis of the error and, in usual cases, restart.

⁵⁴ Implemented by Fillipo Costa (fillipo.costa@cern.ch)

In addition, it is being investigated whether it is possible to have a Real-Time Operation System (RTOS) and all the needed software in the internal eSRAM of the MSS. A minimum RTOS uses fewer memory bits (~5 to 10 KBytes) than the Linux system (~4 MBytes) and MBUs are highly unlikely in the eSRAM (discussed in section 2.2.2). Therefore, the RTOS is expected to be less sensitive to radiation than the Linux system.

The DCS interface and the Monitoring and Safety Module were both irradiated up to the fluence of 5.32×10^{10} p/cm². The Ethernet link was observed to go down and recover by itself twice, during which the readout was not stopped. MTBF in Run2 of the DCS interface is estimated to be ~11.4 hours, which shows that its stability has been improved to an acceptable level. The Monitoring and Safety Module was also working stable, that is, no errors occurred on the RCU2 side.

4.5 Radiation Monitor test

The Radiation Monitor (RadMon) on the RCU2 was tested twice at the Svedberg Laboratory (campaign No.6 and No.7) to characterize the sensitivity for latch-ups and the cross-section of the SEUs in the SRAMs. These tests are discussed in detail in [93] and a summary is given below.




Figure 4-21 Setup of the RadMon test

Figure 4-21 shows the test setup, which includes a RCU2 board and a monitoring PC. In the RadMon, there are two SRAM interfaces, one ADC interface and one SRAM power control interface. Each of the two SRAM interfaces is connected to two SRAM ICs. It uploads a checkboard pattern to all

the addresses in the SRAMs, reads them back and compares them with the expected patterns. The ADC interface monitors the current, voltage and temperature of all the SRAM ICs. The SRAM power control interface controls the power regulators of the SRAMs. All these interfaces are accessed by the Register & Control Module in the RadMon FPGA. Dedicated software on the Linux system in the MSS of the RCU2 main FPGA reads the monitored values from the Control & Register Module via the SPI interface and sends these values to the monitoring PC periodically via a UART interface.




Figure 4-22 SEU counts as a function of fluence of the RadMon test: (a) in campaign No.6, (b) in campaign No.7 [93].

In the first campaign, two runs were made with only the RadMon irradiated. In these two runs, the RadMon was irradiated up to the fluence of 9.0×10^9 p/cm² and 1.2×10^{10} p/cm², respectively. The second irradiation campaign was carried out as part of the system-level test discussed in section 4.4. In this test, the whole RCU2 board was irradiated by a wide beam and the RadMon was exposed to a fluence of 1.5×10^{10} p/cm².

During all these tests, current consumption of the SRAMs was monitored by the RadMon FPGA. In the first campaign, the average current consumption was 9.4 mA and the peak value was 12.2 mA. In the second campaign, the average current consumption was 12.7 mA and the highest value was 14.6 mA. The variations in each test are due to that the SRAMs went into sleeping mode when they were not activated. The differences between the two tests can be explained by different RCU2 prototypes used. No current jumps were observed and the current level measured in these tests was in accordance with the one measured in the lab, implying that no latch-up was detected.

Both the SEU and MBU are counted by the RadMon. The RadMon treats a MBU as a set of SEUs. When MBU occurs, the SEU counter is increased as well. To distinguish different kinds of errors and compare with the values in [95], two kinds of counts were analyzed: (1) The raw SEU counts, including both MBUs and SEUs, (2) the pure SEU counts, including only the SEUs. These two counts are plotted as a function of fluence in Figure 4-22, in which the linear dependence proves the stability of the RadMon. Actually, most of the MBUs are contributed by the burst errors in the interface logic of the SRAMs while the chip select was toggled [93]. Since the chip selects of the SRAMs are continuously toggled in the RadMon, these burst errors must be considered while analyzing the data.

Taking all the 8 tested SRAM into account, the cross-section for raw SEU counts is $(2.6 \pm 0.5) \times 10^{-13} \text{ cm}^2/\text{bit}$ and for pure SEU counts is $(1.3 \pm 0.3) \times 10^{-13} \text{ cm}^2/\text{bit}$, which corresponds well to the numbers in [95].

4.6 Summary and Conclusion

In total seven irradiation campaigns were performed for the RCU2, in which the PCB components, the SF2 FPGA, the hardware interfaces, the RadMon and the whole RCU2 system were tested.

No major problems were observed on all the tested PCB components (e.g. the power regulator, the bus transceiver and the buffers, etc.).

In the first few tests, the SF2 showed some unexpected limitations: (1) SELs were experienced, (2) the PLL lost its lock at an unacceptable rate and (3) the FPGA failed to be reprogrammed at a surprisingly low dose. However, new SF2 chips were produced with SEL enhanced silicon, and the SELs are significantly reduced and even removed for the relevant radiation environment in the TPC. The stability of the PLL has also been improved at least by a factor of ~ 10 as a result of removing the SELs. However, special considerations still should be taken while using the PLLs. The problem of reprogramming still exists. However, since the onset for the reprogramming failures is higher than the expected radiation dose in Run2, it should not be a concern for the RCU2.

The test of the TTC interface revealed that both the custom CDR in SF2 FPGA and the commercial ADN2814 CDR are not suitable for use in the radiation environment of TPC. Therefore, the TTC interface has been redesigned with the already proven TTCrx.

The DAQ interface was observed to go down a few times and data errors were seen in the test. In general, the stability of the DAQ interface is acceptable.

The DCS interface went down at a high rate when the SF2 chip was irradiated. It was the SEL in the FPGA and the SEUs inside the eSRAMs of the MSS MAC that caused this problem. After solving these problems, the DCS interface has been proven to work in a stable manner in radiation.

In the tests for the Radiation Monitor, no latch-up was observed and the cross-section of the SRAM ICs was found to be in line with the previous characterized number in [95]. This proves that the Radiation Monitor can be used to monitor the radiation in the TPC.

The system-level irradiation test revealed some stability issues, especially regarding the readout and the DCS. All these radiation related problems have been dealt with or the mitigation actions for them have been planned.

Table 4-10 summarizes the MTBF in Run2 of the RCU2, including the hardware interfaces, the Monitoring and Safety Module, the readout stability and the DCS stability. In the test, the hardware version is the same as the ones that have been commissioned at the TPC (section 5.3.3), the firmware is the second prototype (section 3.3.2) and the software is as of April 2015 (section 3.1.7). The MTBF in Run2 of the readout stability is comparable with the duration of the longest data-taking session in heavy-ion runs in Run1 (~8 hours). The DCS stability also fulfills the requirement for the RCU1, that is, no more than one or two errors should occur in a data-taking session of ~4 hours in the heavy-ion runs in Run1 [94]. Therefore, the RCU2 is foreseen to work in a stable manner with the radiation load in Run2.

Components	TTC Interface	DAQ Interface	DCS Interface	MSM	DCS Stability	Readout Stability
MTBF in Run2 (hours)	>192	29.3 ± 4.5	11.4 ± 8.2	>22.8	3.6 ± 1.7	7.6 ± 4.5

Table 4-10 Summary of the MTBF in Run2 of the RCU2⁵⁵

⁵⁵ MSM stands for Monitoring and Safety Module.

5 Testing and integration

The test and integration work of the RCU2 was divided into three phases⁵⁶. The first phase dealt with the validation of the RCU2 hardware prototype for mass production. During the second phase, the main focus is the validation of the second prototype firmware. This is the preparation of the system-level irradiation test (section 4.4). The third phase refers to the tests performed with the commissioning version of the firmware, this is the prerequisite for the installation and commission. At the time of writing, all the 216 RCU2s have been installed and have been operational for about a year.

5.1 Validation of the ALTRO Interface

The functionality of the RCU2 hardware prototype needed to be verified before mass production. This section focuses on the validation of the ALTRO interface. Reading data without any error from the FECs through the ALTRO interface is the basis of developing the Readout Module in the firmware. The test procedure for the other hardware interfaces and the Radiation Monitor is the same as in the irradiation campaigns (discussed in chapter 4). These components were validated to be fully functional while developing the firmware and therefore not discussed further in this thesis.

5.1.1 Test with simple ALTRO bus master

As discussed in section 3.3.1, the Readout Module in the RCU2 firmware was planned to be ported from the RCU1, but this was not successful as discussed in section 3.3.1. Since all the lab-tests were performed on the first PCB prototype of the RCU2, it could not be excluded that the failures of communicating with FECs were caused by hardware issues. Therefore, a simple ALTRO bus master for a single branch was designed to verify the write and read transactions to the FECs.

Figure 5-1 shows the test design in the SF2, which includes the ALTRO bus master and some modules ported from the RCU1: the RCU decoder, the Instruction Sequencer and the Result Unit. The RCU Decoder receives and processes the orders to perform write and read transactions from

⁵⁶ All the tests were performed in room temperature if not otherwise stated.

the SF2 MSS. The Instruction Sequencer stores these orders and executes them to drive the ALTRO Bus master. The results of executing these orders are stored in the Result Unit. The tests are controlled by the shell script running on a PC.




Figure 5-1 Test design with the simple ALTRO bus master



Figure 5-2 Observation of the write and read transaction

The waveforms of a write transaction and a read transaction are shown in Figure 5-2, in which all the control signals and the readout clock are in good shape. In total, more than 5000 write and read transactions were tested. In these tests, the ALTRO bus was driven to the maximum load by changing the data pattern consecutively from all '0's to all '1's. No error was detected in the tests. This eliminated the concern of the hardware issues and laid the basis for the further development of




Figure 5-4 Screenshot of CHRDO transaction. (a) Handshake procedure. (b) DSTB on the four branches

Trigger rate (KHz)	Number of samples	Data pattern	Duration (hours)	Number of transaction (billion)	Error
10	10	Ramp	~2	~4	0
10	10	Zero - One	~2	~4	0
10	100	Ramp	~4	~15	0
10	100	Zero - One	~4	~15	0
1	1000	Ramp	~4	~14.5	0
1	1000	Zero - One	~4	~14.5	0
10	1000	Ramp	~8	~289	0
10	1000	Zero - One	~8	~289	0

Table 5-1 Stress test of the ALTRO interface

Two data patterns were used in the test: the consecutive all ‘0’s following all ‘1’s and the decimal ramp number from 0 to 1000. With the first pattern, the ALTRO bus can be driven to the maximum load. With the second pattern, errors in specific bits can be identified. The RCU2 was reading data from four channels (one in each FEC) concurrently. Various data patterns, trigger rates and number of samples (10-bit word) in each ALTRO channel were used. Subfigure (a) of Figure 5-4 shows the measurement of a CHRDO transaction, in which all the signals show good signal integrity and according to the ALTRO bus specification [9]. Subfigure (b) of Figure 5-4 shows the DSTB signals measured in persistent mode while the data-taking was performed at 10 kHz with the number of samples in each ALTRO channel set to 1000. More screenshots of the ALTRO bus signals are presented in Figure E.2-1. The complete test results are listed in Table 5-1. No error was observed

in a total of ~645 billion transactions. From this it was concluded that both the RCU2 and the new backplanes were working according to specification.

Considering the promising test results of the RCU2 hardware prototype in both room condition and irradiation campaigns (discussed in Chapter 4), a green light was given for the mass production at the beginning of 2015.

5.2 Test with the second prototype of firmware

As part of the preparation for the system-level irradiation campaign (discussed in section 4.4), the complete RCU2 system was validated in the lab. The test design and test setup were the same as in the irradiation campaign (see Figure 3-10 and Figure 4-18). In general, the test was performed with the second prototype of firmware on one RCU2 connecting to four FECs on separate branches. The DDL2 link was operating at 2.125 Gbps. A fixed pattern of data was filled into the on-board memories of the FECs, read back by the RCU2 and then sent to the DAQ computer, where the data errors in the data-stream were checked by the DAQ software.




Figure 5-5 Test procedure for the RCU2 with the second prototype of firmware

The test was divided into two steps. To begin with, the compatibility between the RCU2 and all the other devices (trigger crate, FECs and data computer) in the readout system was tested. Secondly, the stability of the data-taking and the errors in the data was checked. The test procedure shown in Figure 5-5 was used for both steps.

In the first step, the test procedure was executed for more than 100 times and it was observed that the data-taking could always start and stop in a stable manner. This implies that the RCU2 is

compatible with all the other devices. The data-taking in each test was running just for a short period (few seconds) to reduce the time it takes to perform the test.

In the second step, the test procedure was repeated a few times and the data-taking in each test lasted up to a few hours to check the stability of the readout and detect the data errors. At the time of testing, the DAQ software could not identify the exact position of the errors in each event. However, the number of the error words in each event can be counted.

No. of channels (AO/AI/BI/BO)	No. of samples (AO/AI/BI/BO)	Trigger Rate (Hz)	Event Rate (Hz)	Duration (hour)	Readout Stop	Data Error
128/128/128/128	998/998/998/998	10	10	2	No	No
128/128/128/128	998/998/998/998	100	100	2	No	No
128/128/128/128	998/998/998/998	500	300	2	No	Yes
128/128/128/128	998/998/998/998	1000	300	2	No	Yes
128/128/128/128	782/842/684/184	100	100	2	No	No
128/128/128/128	782/842/684/184	1000	370	2	No	Yes
55/80/75/50	483/483/483/483	500	500	1	No	No
55/80/75/50	483/483/483/483	1000	968	1	No	Yes
55/80/75/50	100/150/140/110	500	500	1	No	No
55/80/75/50	100/150/140/110	1000	1000	1	No	No
55/80/75/50	100/150/140/110	2000	1154	1	No	Yes

Table 5-2 System level validation of the RCU2 (second prototype of firmware)

The test scenarios and test results are listed in Table 5-2. The number of channels in each FEC, the number of samples in each ALTRO channel and the trigger rate were varied. The trigger crate was sending triggers with fixed time spacing⁵⁷. No stop of data-taking was observed in any of the test cycles, which indicates that the data-taking was working correctly. However, data errors were observed when the DDL2 link was saturated. After analyzing the recorded data, it was found that these errors were led by the flaws on handling the XOFF signal from the DDL2 link. This problem has been discussed in detail in section 3.3.2 and it has been solved in the commissioning version of firmware.

⁵⁷ The trigger crate can be configured to send triggers with fixed time spacing or random spacing. In this thesis, the triggers are always sent in fixed time spacing if not otherwise stated.

Ideally, this should have been corrected before the irradiation campaign, but due to time limitation it was not. However, the error was understood, and the rate of the error was known for given circumstances that was controllable. Hence, the second prototype could be used in the irradiation campaign.

5.3 Tests with the commissioning version of firmware

The commissioning version of the firmware can be divided into subversions, defined by the speed of the DDL2 link. To start with, the DDL2 link was working at 2.125 Gbps, which was the same as for the second prototype of firmware. Moreover, attempts were done to bring the speed of the DDL2 link to 4.25 Gbps but this failed. The reason was never fully understood. Finally, the speed of the DDL2 link was set to 3.125 Gbps, still within the requirements. With this speed, the readout was verified to be stable both in the lab and in the TPC. In this section, the tests that were performed on these subversions are discussed.

5.3.1 DDL2 link at 2.125 Gbps

Two setups were used to test the commissioning version of firmware with the DDL2 link at the speed of 2.125 Gbps. The first setup is the one used for testing the second prototype of the firmware, that is, 1 RCU2 connecting to 1 FEC in each branch. This setup is located at University of Bergen in Norway, from here on referred to as the Bergen setup. The second setup is located at CERN in Switzerland, from here on referred to as the CERN setup. It contains 1 RCU2 and 25 FECs⁵⁸, of which Branch AO contains 7 and each of the other three branches contains 6. A sketch of both setups is shown in Figure 4-18. Pictures of these two setups are shown in Figure 4-19 and Figure 5-6, respectively. Other than the number of FECs, all the devices in these two setups, including the hardware, the firmware, the software, etc., are the same. The reason for using two setups is to perform cross-checks to ensure that the tests are not corrupted by hardware issues. Noticeably, the busybox⁵⁹ was included in both the setups.

⁵⁸ This refers to readout partition 1, which is the partition with largest number of FECs.

⁵⁹ The busybox is discussed in detail in [96]. The purpose of busybox is to let the Central Trigger Processor [97] know when the data buffers on the FEE are full by asserting a busy signal which prevents further issuing of triggers.



Figure 5-6 The CERN setup - 1 RCU2 connects to 25 FECs

Hardware testing of the readout system




Figure 5-7 Test procedure for the RCU2 with the production of firmware

As discussed in section 5.2, these tests were sorted into two categories: (1) compatibility between the RCU2 and the other devices in the readout system and (2) stability of the data-taking and the errors in the data-stream.

Figure 5-7 presents the test procedure, which is similar to the one for the second prototype of the firmware, except for two major differences. First of all, the busybox is involved in the system, and it needs to be reset before data-taking. Furthermore, in some cases the whole RCU2 is power cycled (option (a)) and in the other cases only the Linux is rebooted (option (b)). Option (b) is added to verify that the data-taking is not dependent on the rebooting of the Linux OS.

To start with, the test procedure was executed 100 times with option (a) and 100 times with option (b). In the tests, it was observed that the data-taking could always start and stop in a stable manner. This proves that the RCU2 is compatible with the other devices in the readout system. In addition, when the RCU2 was tested with option (b), the data-taking was not stopped while Linux was rebooting. This proves that the standalone initializing scheme of the SERDES in DDL2 Module works as intended.

No. of channels (AO/AI/BI/BO)	No. of samples	Trigger Rate (Hz)	Duration (hours)	DDL2 link saturation
128/128/128/128	1000	100	4	No
128/128/128/128	1000	1000	4	Yes
128/128/128/128	100	100	4	No
128/128/128/128	100	5000	4	Yes
896/768/768/768	1000	10	8	No
896/768/768/768	1000	30	8	No
896/768/768/768	1000	100	8	Yes
896/768/768/768	10	10	4	No
896/768/768/768	10	1000	4	No
896/768/768/768	10	5000	4	Yes

Table 5-3 System level validation of the RCU2 (commissioning version of the firmware with DDL2 bandwidth of 2.125 Gbps)

Afterwards, all the tests of the second prototype of firmware (listed in Table 5-3) were repeated on the Bergen setup⁶⁰. No data error or stop of data-taking was observed. This proves that the XOFF signal (discussed in section 5.2) is now properly handled.

After the functional tests were done, stress tests with long period of data-taking were performed on both setups. In these tests, all the available ALTRO channels in each setup were used and configured with the same number of samples. The trigger rate was intentionally varied. As listed in **Table 5-3**, the test scenarios cover various combinations of the number of samples in each ALTRO channel and the trigger rate, e.g. high number of samples (1000) with low trigger rate (10 Hz), high number of samples (1000) with high trigger rate (5000 Hz), etc. In these tests, the data-taking was stable

⁶⁰ Busybox was not included in these tests.

and no data errors are detected. The results of these tests, especially the ones in which the DDL2 link was saturated, proves that the commissioning version of firmware could take data at 2.125 Gbps in a stable manner.

Readout speed of the RCU2

The readout speed of the TPC is counted as that of the slowest partition, which is the one with maximum number of FECs. This is readout partition 1 with 25 FECs. Therefore, the readout performance of the RCU2 was benchmarked on the CERN setup, which is a readout partition 1. The readout time of each event is measured on the RCU2 from when the L2 trigger is issued to the FECs until the data transmission over the DDL2 link is completed. In the test, the number of samples in each ALTRO channel was varied from 10 to 1000.




Figure 5-8 Benchmarking on the RCU2 with DDL2 at 2.125 Gbps

The subplot (a) of the Figure 5-8 compares the readout time⁶¹ of the RCU2 with that of the RCU1. The improvement factor is calculated as the ratio between the readout speed of the RCU2 and the RCU1 for the event with the same number of samples in each channel. For small events, the readout speed has been improved by a factor from 1.5 to 1.7. For large events, the improved factor decreases gradually to ~ 1.25 . As expected, the readout performance does not realize the planned improved factor of 2. According to the subplot (b) of the Figure 5-8, the DDL2 link starts to get saturated after the number of samples reaches ~ 30 to 40 and $\sim 33\%$ of the readout time was caused by this saturation.

⁶¹ In this thesis, the readout speed is measured in full readout mode if not otherwise stated. Readout speed equals readout time divides event size.

Therefore, it is concluded that the DDL2 link needs to work at a faster speed.

5.3.2 DDL2 link at 4.25 Gbps

The bandwidth of the DDL2 link was then increased to the intended value of 4.25 Gbps. This was accomplished by changing some configuration parameters of the SERDES interface in the firmware. Afterwards, the RCU2 was verified on the CERN setup. Unexpectedly, lots of errors were detected in the data, of which the data from Branch BI has the highest number.




Figure 5-9 Benchmark on the RCU2 with DDL2 at 4.25 Gbps

Theoretically, the bandwidth of the Readout Module is ~ 305 MBytes/s⁶² and that of the DDL2 link is ~ 425 MBytes/s⁶³. Hence, these data errors should not be caused by the saturation on the DDL2 link. The RCU2 was benchmarked on the CERN setup with 6 FECs in each branch. The number of samples per ALTRO channel was varied from 10 to 1000. Subplot (a) of the Figure 5-9 shows that the maximum throughput, i.e. the bandwidth, of the RCU2 is ~ 305 MBytes/s, which equals to the estimated value. Subplot (b) of the Figure 5-9 shows that the DDL2 link never was saturated. These observations confirm that these errors are not related to the saturation of the DDL2 link.

In order to figure out the origins of these data errors, the following three kinds of readout tests were performed on the CERN setup with 25 FECs. All the ALTRO channels in the setup were used and

⁶² 305 MBytes/s equals to 80 MHz multiplies 32-bit, where 80 MHz is working frequency of the Readout Module and 32-bit is the size of the data interface from the Readout Module to the DDL2 Module (link).

⁶³ 4.25 Gbps equals to 531 MBytes/s. In addition, the 10-bit to 8-bit conversion is performed on the DDL2 link. Hence, the DDL2 link should have a theoretical bandwidth of ~ 425 MBytes/s.

the number of samples in each channel was varied from 10 to 1000.

In the first test, the normal ALTRO readout was performed, i.e. the RCU2 reads data from the FECs, processes it and transmits it to the DAQ computer. If all the four branches were active, data errors started to occur when the number of samples became larger than ~ 20 . If the Branch BI was turned off, data errors started to appear only when the number of samples was larger than ~ 600 .

In the second test, the RCU2 was configured to not read data from the FECs. A Data Generator written in VHDL was used to generate a ramp pattern of data, which was then transmitted to the data computer. In this case, no data error was seen during the transmission of several TB of data.

In the last test, the RCU2 was configured to read data from the FECs but discard it. The data that was transmitted to the DAQ computer was generated by the Pattern Generator. In this case, the observations on the data errors were in line with that in the first tests. With all the four branches being active, data errors started to appear if the number of samples becomes larger than ~ 20 to 30. If the Branch BI was turned off, the onset of the number of samples for data errors was ~ 600 to 650.

The reason for these errors were never fully understood, but the most likely reason is that the switching noise of the ALTRO bus was disturbing the input clock to the SERDES used for the DDL2 link. Due to time constraints of the project, it was decided to rather go for a data rate of 3.125 Gbps, that was anyway within the specification.

5.3.3 DDL2 link at 3.125 Gbps

To bring the DDL2 link working at 3.125 Gbps, the on-board oscillator for the DAQ interface was changed from 106.25 MHz to 156 MHz. Correspondingly, the DDL2 module in the firmware was adapted to the new oscillator. On the first two modified RCU2s, the test procedure shown in Figure 5-7 was executed 100 times both with option (a) and option (b). No error was detected and this proved the functionality of the new design.

Stress test before installation

In total 240 RCU2s (216 for installation at TPC and 24 for backup) needed to be verified prior to installation. Due to the limited time available for the tests, 6 RCU2s were chosen as samples, each of which was tested for a long period of a tens of hours. Each of the other 234 RCU2s was tested

for a shorter period of ~2 hours. All these tests were performed on the CERN setup with 25 FECs. In addition, an attenuator was connected to the output of the DAQ interface to emulate the situation in TPC⁶⁴.

The tests performed on the 6 sample RCU2s are listed in Table 5-4. The number of samples in each ALTRO channel was varied between 10, 100 and 1000, which covers a wide range of data volume. Because these tests focused on the data errors, in each test the data rate was driven to be highest possible value so that the ALTRO bus was driven with maximum load.

Board number	Number of samples in each channel	Number of events (million)	Throughput (MByts/s)	Duration (hours)	Data volume (TB)
1	10	~77.7	78	17.8	4.98
1	100	~10.6	260	5.1	4.74
1	1000	~2.6	283	11.0	11.25
2	10	~127.0	78	29.0	8.14
2	100	~4.5	260	6.9	6.42
2	1000	~18.2	283	16.0	16.39
3	1000	~5.0	283	21.0	21.44
4	1000	~5.2	283	21.9	22.33
5	1000	~4.6	283	19.4	19.93
6	1000	~5.1	283	21.5	21.70

Table 5-4 Test results of 6 sample RCU2s

To start with, the normal ALTRO readout was performed on these 6 boards. The RCU2 read the pre-filled data from the FECs, processed it and transmitted to the DAQ computer, where the data was checked to detect errors. No data error or stop of readout was observed while reading ~110 TB of data over ~170 hours. In addition, the throughput and data rate of different RCU2s were the same as long as the number of samples was set to be the same. This proves the stability of the readout. Afterwards, the other 234 RCU2s were tested for ~2 hours⁶⁵. The number of samples was set to be

⁶⁴ In TPC, the data signal needs to transmit a long distance before it reaches the DAQ computer and the strength of the data signal will decrease along the path. In the lab, the transmitting path is short so that an attenuator was used to weaken the signal.

⁶⁵ These tests were performed by manpower from CERN, Goethe-Universität (Germany), Lund University (Sweden), University of Oslo (Norway), University of Bergen (Norway), Vestfold University College (Norway) and Bergen University College (Norway). Torsten Alt (Torsten.Alt@cern.ch) and the author were in charge of the test group.

1000 and the data rate was driven to be the highest possible value (~ 66 Hz). These tests lasted over 490 hours and no data error was seen in a total of ~ 500 TB of data.

All 240 RCU2s with the DDL2 link at 3.125 Gbps passed all the tests and were thereby ready for installation.

Discussion on Readout Speed⁶⁶

The bandwidth of the DDL2 link was measured with 6 FECs in each branch. As shown in subplot (a) of the Figure 5-10, the throughput of readout partition 1 reaches its peak value of ~ 295 Mbytes/s when the number of samples is ~ 110 and then decreases gradually to ~ 280 Mbytes/s with the increase in the number of samples.




Figure 5-10 Benchmarking of the RCU2 with DDL2 at 3.125 Gbps.

Subplot (b) of Figure 5-6 compares the readout speed of the RCU2 with that of the RCU1. For large events, the readout speed is increased with a factor of ~ 1.9 , which is slightly smaller than the design requirement of 2. For small events, the readout speed does not live up to requirements. The RCU2 needs ~ 670 μ s to read an empty event. Together with the round-trip time from the busybox to the local trigger unit of ~ 130 μ s, at least ~ 800 μ s is needed for reading an empty event. This is only ~ 1.6 times faster than the RCU1 and introduces ~ 300 μ s to the fixed busy time of the TPC at 500 μ s.

⁶⁶ The benchmark and optimization were performed by the author on the commissioning version as of January 2016.

There are three factors that affect the readout performance. Firstly, the Branch AO contains 1 FEC more than other branches. At the end of each readout, this FEC is the only one that provides data to the RCU2, at the maximum data rate of ~ 195 MBytes/s⁶⁷, which is far smaller than the bandwidth of the DDL2 link. Its impact on the readout speed becomes more significant with the increase in the data volume from a single FEC, i.e. the number of samples in each ALTRO channel. However, this issue is an intrinsic factor of the TPC structure and cannot be avoided.

Secondly, the firmware under test contained only one data buffer for each branch. Therefore, the next ALTRO channel cannot be read until the data from the previous channel has been processed. This single buffering structure affects the events with different sizes⁶⁸ to different extent. While reading small events, the DDL2 link is not saturated and always available for receiving data. So the processing time is therefore directly counted into the readout time. While reading large events, the DDL2 is saturated and it is the main factor that slows down the readout speed. However, while reading the channels in the last FEC in Branch AO, the DDL2 link is not saturated, so the data processing time is also directly counted into the readout time. The larger the volume of data in each ALTRO channel is, the larger this processing time is. This explains why the peak throughput appears at ~ 110 samples when the DDL2 link is just about to be saturated, and then decreases with the increase of the number of samples in each channel.

In addition to the single buffering, another issue that affects the readout speed, especially for the small events, is the execution time of the ALTRO protocol implementation. As discussed in 5.3.2, data errors were induced by the switching noise of the FPGA pins and the ALTRO bus pins when DDL2 link was working at 4.25 Gbps. A conservative scheme was therefore used to switch the direction of the ALTRO bus in the handshake protocol. Figure 5-11 compares the non-conservative switching scheme with the conservative switching scheme for reading a single word. Figure E.2-1 and Figure E.2-2 in appendix E.3 compares the situation of reading an empty channel and a channel with 10 samples, respectively. It is observed that 100 ns more is needed by the conservative scheme to performed a CHRDO transaction.

According to the previous discussions, three actions were proposed to improve the readout speed. Firstly, replacing the conservative bus switching scheme with the non-conservative switching

⁶⁷ 195 MBytes/s equals to 40 MHz (readout clock) multiplied by 40 bits (width of ALTRO data bus).

⁶⁸ The event size is defined as the number of samples in each ALTRO channel.

scheme. This can reduce the overhead of reading each ALTRO channel by at least 100 ns. Secondly, implementing the double buffering to compensate the data processing time. The efficiency of double buffering depends on the volume of the data in each ALTRO channel. At last, implementing the sparse readout to skip empty channels. As discussed in 3.3.2, in sparse readout an overhead of $\sim 106 \mu\text{s}$ is needed to build the “hit list”. This overhead can be compensated by skipping 157 empty channels (675 ns per channel), which occupy $\sim 17.5\%$ of all the 896 channels in the largest branch.




Figure 5-11 Measurement of reading single word from one channel. (a) RCU2 with non-conservative ALTRO bus switching. (b) RCU2 with conservative ALTRO bus switching.

The first two actions were implemented in an optimized firmware to study their efficiency. Throughput of the optimized design is shown in the subplot (a) of Figure 5-10. It is clear that the throughput of reading all the events, especially the small events, has been increased. With respect to that of the current RCU2, the time of reading an empty event and a full event has been increased by 9.2% and 3.4%, respectively. As shown in the subplot (b) of Figure 5-10, the readout speed has been improved by a factor of at least ~ 1.95 (in most cases above 2) with respect to that of the RCU1. All of the proposed optimization has later been included in the commissioning version of the firmware.

5.3.4 Discussion on Radiation Mitigation

When the RCU2 is in the radiation environment of Run2, the SEEs in the SRAMs, flip-flops and PLLs of the firmware may lead to two kinds of errors: reliability errors and data errors.

The reliability error is the most critical consequence of the SEEs. It will cause the TPC readout

system to stop and could lead to an unplanned termination of a data-taking session. As to the firmware, the reliability errors can be caused by the following two reasons:

- (1) There are ~1000 flip-flops in the critical path of the RCU2 firmware. SEUs in these flip-flops can cause a functional interrupt.
- (2) If the PLL that provides the system clock is hit, it will have fatal consequences for the firmware.

As discussed in section 4.2.4 and 4.2.5, the cross-section of SEE for each individual flip-flop and PLL is $(2.6 \pm 0.7) \times 10^{-14} \text{ cm}^2$ and $(2.8 \pm 2.0) \times 10^{-12} \text{ cm}^2$, respectively. Taking all the 216 RCU2s into account, the MTBF in Run2 of critical flip-flops and PLLs is calculated to be 16.3 ± 4.0 hours and 156 ± 113 hours, respectively. By multiplying the reliability of the flip-flops and the PLLs, the MTBF of reliability errors in Run2 is estimated to be 14.8 ± 4.3 hours, which is longer than the longest data taking session in heavy ion collisions in Run1 (~8 hours). However, considering also the failure rate of the TTC interface and the DAQ interface, the reliability of the whole RCU2 will be worse than this estimation. Failure rates of the components that can lead to an unplanned termination of a data-taking session have been calculated in the previous sections and they are listed in Table 5-5. By multiplying the reliability of these components, MTBF in Run2 of the RCU2 is estimated to be 9.4 ± 3.1 hours, which is comparable with the number estimated based on the system-level irradiation test in section 4.4.1 (7.6 ± 4.5 hours). This is also comparable with the longest data taking session in the heavy-ion collisions in Run1 (~8 hours). This implies that the reliability of the RCU2 in Run2 is expected to be in the same level as the RCU1 in Run1. Still, since only about 25% of the logic cells in the SF2 has been utilized by the firmware, some radiation mitigations can be implemented to protect the critical registers. As recommended by Microsemi in [99], state machine registers can be protected with hamming encoding and the other registers can be protected with local Triple Modular Redundancy.

Data errors are mainly caused by the SEUs that occur in the ROLM and data memories, both of which are instantiated with LSRAMs. SEUs in the ROLM cause bit-flips in the channel addresses. In case this happens, the RCU2 will read data from wrong channels and not from the correct channels. The channel addresses are static, so the SEUs in the ROLM are accumulated run-time. In total, there are 557568 channels spreading over 216 readout partitions. The address of each channel is 12 bits. It can be estimated that one channel address will turn wrong about every 0.84 hours in a data-taking session in Run2. At the end of a session that lasts ~8 hours, in each event the data of

~10 channels will be wrong. The erroneous channels occupy only 0.00179% of all the channels, so they have limited effects on the data quality.

Components	TTC Interface	SERDES (DAQ Interface)	AVAGO SFP (DAQ Interface)	Firmware
MTBF in Run2 (hours)	192 $\pm$ 192	68.6 $\pm$ 49.6	51.4 $\pm$ 30.7	14.8 $\pm$ 4.3

Table 5-5 Reliability estimation of the complete RCU2⁶⁹

Since the event data is buffered in the data memories before it is sent to the DAQ, the SEUs in the data memories will directly cause bit-flips in the event data. The cross-section of the bit-flips in the event data should be the same as the cross-section of the SEU in the LSRAMs, which is $(1.7 \pm 0.2) \times 10^{-14} \text{ cm}^2/\text{bit}$. For the LSRAMs with the total size of 81 MBytes, which is the same as the size of the estimated largest event in Run2 [26], the mean time between SEUs is estimated to be ~29.4 seconds. Given the event rate in Run2 is expected to be 400 Hz (discussed in section 1.3), one bit-flip is expected to occur during the readout of about every 11760 events, which corresponds to an error rate of $\sim 8.5 \times 10^{-5} \text{ bits/event}$.

Despite that the rate of data errors is low, it still has some effects on data quality and can be reduced by protecting the ROLM and the data memories with SECDED mechanism (hamming-3 encoding). In the mitigation scheme, parity bits are added to the channels addresses before they are written in to the ROLM. If SEU occurs in a channel address, it can be detected and corrected when the channel address is accessed. The corrected address is used in the readout and written back to the ROLM. All the event data words are encoded before they are pushed in the data memories. To keep the consistence in each data package, the CDH words that are stored in 4 uSRAMs need to be encoded as well. The decoding and correction are done on the RCU2 before the data is sending to the DAQ. If the decoding is performed on the DAQ side, the size of each event will be increased by ~20% due to the extra parity bit, which limits the bandwidth of the RCU2.

The mitigation techniques that are discussed above have been implemented in the commissioning version of the firmware. The state registers of the FSMs, the Readout List Memory and the data

⁶⁹ No failure of the TTC Interface was observed in the irradiation test, so the upper MTBF in Run2 is calculated based on 1 failure.

memories are hamming-3 protected. The other registers are protected with local Triple Modular Redundancy. While implementing these mitigation techniques, no timing issues were reported by the Libero SoC software. Afterwards, the firmware with these SEU mitigation techniques was verified on three RCU2 boards in the lab.

Board Number	Samples	Duration (hours)	Events (Million)	Data Volume (TB)	Data error	Readout stop
1	100	4	~8.45	~3.8	0	0
1	1000	4	~0.95	~4.1	0	0
2	100	4	~8.45	~3.8	0	0
2	1000	4	~0.95	~4.1	0	0
3	100	6	~12.7	~5.7	0	0
3	1000	6	~1.43	~6.1	0	0

Table 5-6 verification of firmware with mitigation actions

The test procedure is discussed in detail in section 5.2 and it is summarized below. Dedicated pattern of data is written into the memories on the FECs. The RCU2 reads the data from the FECs and sends it to a data computer through the DDL2 link. The data computer checks the received data to find errors. In each test, the trigger rate was driven to highest possible value. The number of samples in each ALTRO channel was set to be either a small value of 100 or a large value of 1000. The test results are presented in Table 5-6. No data error or stop of readout was observed while in total ~28 TB data was read during ~28 hours. This proves that the firmware with the proposed mitigation actions could perform data-taking in a stable manner. Unfortunately, it was not possible to test this version in radiation within the timeframe of this PhD thesis.

5.4 Summary

The commissioning of the RCU2 were carried out into two stages. In 2015, 6 RCU2s with the DDL2 link at 2.125 Gbps were commissioned in one of the 36 TPC sectors. In January 2016, all the 216 RCU2s with the DDL2 link at the speed of 3.125 Gbps were installed at the TPC. Several TB of data was looped on each sector and no data error or stop of readout was observed. At the time of writing (November 2016), the RCU2s have been taking data for about one year. Figure 5-12 shows the reconstructed data taken by the TPC in the first stable p-Pb collision of 2016.




Figure 5-12 Reconstructed data taken by TPC in the first p-Pb collision in Run2⁷⁰

⁷⁰ This picture is from Robert Helmut Munzer (robert.muenzer@cern.ch)

6 Summary and conclusion

In LHC Run1, the Readout Control Unit 1 (RCU1) performed even better than specification [20]. However, in Run2 the energy of colliding beams is increased from 8 TeV to up 13 TeV, so both the event size and the radiation levels are increased. Therefore, the RCU2 is designed to provide a faster readout speed and improved radiation tolerance compared to the RCU1.

The RCU2 has four major advantages over the RCU1 in hardware: (1) it has four branches instead of two branches, (2) the bandwidth of the DDL link is increased from 1.60 Gbps to 3.125 Gbps, (3) it uses a single PCB design that integrates all the functionalities of the three PCB board in the RCU1, and (4) the flash-based Microsemi SF2 FPGA SoC is used as the main FPGA instead of the SRAM based Xilinx Virtex 2 Pro FPGA that was used on the RCU1.

6.1 Main Contribution

Radiation tolerance of the RCU2 has been studied through several irradiation campaigns. Two kinds of radiation effects are the focus of these tests, SEEs and TID effects. To begin with, the SF2 chip, the TTC interface, the DCS interface and the DAQ interface were characterized in several irradiation campaigns. Afterwards, the whole RCU2 system was tested in a system-level irradiation campaign, in which the RCU2 was working similarly as in the TPC. Cross-sections for various kinds of failures (errors) were extracted and the corresponding MTBF in Run2 was calculated. The following actions have been taken or proposed to improve the readout stability: (1) The SF2 has been confirmed to be immune to SEL in the radiation environment of LHC. (2) For the TTC interface, the ideas of using alternative solutions other than the TTCrx chips have been confirmed to be infeasible. (2) Triple Modular Redundancy or hamming protection on vital modules of the firmware have been proposed. (3) The readout logic has been separated from the CPU with a stand-alone module for initializing the SERDES in the DAQ interface. Regarding the DCS stability, the following two actions have been taken or proposed: (1) SECDDED protection on the eSRAMs in the SERDES of the DCS interface have been enabled. (2) Investigation of having a RTOS and all the needed software in the internal eSRAM of the MSS has been proposed. In general, actions have been taken against all the radiation related problems that were revealed during the irradiation tests. In conclusion, despite the fact that the radiation level in Run2 is estimated to be ~ 3.75 times higher than that in Run1, the RCU2 is still expected to work as satisfactory as the RCU1 in Run1.

Development of the firmware has gone through three versions, the first prototype, the second prototype and the commissioning version. Many tests have been performed for the RCU2 and they are the prerequisite of the mass production of the hardware, the irradiation tests in different stages and the final installation. Readout performance of the RCU2 has been studied, based on which the solutions to further improve the readout speed have been proposed. Eventually, all 240 produced RCU2s have been verified to work in a stable manner with the DDL2 link at 3.125 Gbps.

6.2 Running experience

At the time of writing, the RCU2 has been recording data in p-p and p-Pb collisions without any major issues. The RCU2 has met its performance requirements; the current data acquisition system of the ALICE TPC is measured to record data at factor two higher rates than the readout rates during Run1.

System	Period	Beam	Energy (TeV)	Total EoR	EoR by TPC	Ratio (%)
RCU1	2010	p-p	7	754	31	~4.1
RCU2	2016	p-p	13	1246	36	~2.8
RCU1	2013	p-Pb	8	230	23	10.0
RCU2	2016	p-Pb	8	303	4	~1

Table 6-1 Overview of End of Run (EoR) reasons for the ALICE experiment⁷¹

Regarding the radiation tolerance of the RCU2, a qualitative measure is made in Table 6-1 by comparing End of Run (EoR) reasons for the ALICE experiment. An EoR is referred to as the reason which has caused to end a data-taking session during LHC operation. EoRs are mainly due to operational procedures and conditions of the LHC and the beam itself. Ideally, any detector of the ALICE experiment should not cause such a situation which can lead to end a data-taking session during normal operation of the LHC. An EoR reason caused by a detector may have been due to errors induced by radiation effects on readout electronics or malfunction of any other subsystem in the data acquisition system.

Table 6-1 shows that EoR reasons due to TPC readout electronics for p-p collisions at factor of two

⁷¹ Statistics are from the ALICE logbook.

higher beam energy has been reduced to half as compared to Run1. It is evident that radiation tolerance and stability of the system based on RCU2 has contributed significantly to detector uptime by reducing EoR reasons caused due to TPC readout electronics by approximately factor 10 as compared to the RCU1 for p-Pb collisions during Run1 at similar energy levels.

6.3 Outlook

The RCU2 will work at the TPC until the end of Run2. In 2018, it will process event data in Pb-Pb collisions. In Run3, the TPC will include a new ASIC chip, the SAMPA, for a new faster readout [98]. The Microsemi SF2 or IGLOO2 [100] is currently under consideration for several other projects at CERN, for instance the Beam Halo Monitor at CMS [101], the Muon Frontend control system at LHCb [102] and the TOF readout electronics at ALICE [103]. The experience gained from the RCU2 project is therefore valuable for the whole community and not only for the ALICE TPC detector.

Reference

- [1] The ALICE collaboration et al., *the ALICE Experiment at CERN LHC*, 2008 JINST 3 S08002.
- [2] The ATLAS collaboration et al., *the ATLAS Experiment at the CERN Large Hadron Collider*, 2008 JINST 3 S08003.
- [3] The CMS collaboration et al., *The CMS experiment at the CERN LHC*, 2008 JINST 3 S08004.
- [4] The LHCb collaboration et al., *The LHCb Detector at the LHC*, 2008 JINST 3 S08005.
- [5] <https://environmentalarmageddon.files.wordpress.com/2010/10/lhc-sim.jpg>
- [6] P La Rocca, *The upgrade program of the major experiments at the Large Hadron Collider*, 2014 J. Phys.: Conf. Ser. 515 012012.
- [7] J. Alme et al, *The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events*, Nuclear Ins. and Method. Section A: Volume 622, Issue 1, 1 October 2010, Pages 316–367.
- [8] J. Alme, *Firmware Development and Integration for ALICE TPC and PHOS Front-end Electronics*, Ph.D. dissertation, University of Bergen, 2008.
- [9] CERN, *ALICE TPC Readout Chip User Manual*, June 2002.
- [10] The ALICE collaboration et al., *The ALICE data acquisition system*, Nucl. Instrum. Meth. A 741 (2014) 130.
- [11] H. Soltveit et al, *The Preamplifier shaper for the ALICE TPC-Detector*, Nuclear Ins. and Method. Section A: Volume 676, 1 June 2012, Pages 106–119.
- [12] P. Chochula, *The ALICE Detector Control System*, IEEE Transactions on Nuclear Science, Volume: 57, Issue: 2, 2010, Pages: 472 - 478
- [13] J. Christiansen et al., *TTCrx reference manual - A timing, trigger and control receiver ASIC for LHC detectors*, v. 3.11 (Dec. 2005).
- [14] Altera, *Excalibur devices hardware reference manual*, v3.1 (Nov. 2002).
- [15] Xilinx Inc., *Virtex-II Platform FPGAs: Complete Data Sheet*, March 1, 2005
- [16] Actel Inc., *ProASIC PLUS® Flash Family FPGAs v5.8*, June 2009
- [17] K. Røed, *Single Event Upsets in SRAM FPGA based readout electronics for the Time Projection Chamber in the ALICE experiment*, PhD thesis, University of Bergen, 2009
- [18] J Alme et al., *RCU2 - The ALICE TPC readout electronics consolidation for Run2*, 2013 JINST 8 C12032

-
- [19] 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.
- [20] A. U. Rehman, *The ALICE TPC Readout Electronics*, PhD thesis, University of Bergen, 2012
- [21] J. Alme et al, *Proposal for an optimization of the ALICE TPC read-out for running at full energy*, ALICE TPC review meeting 6. June 2012, <https://indico.cern.ch/event/194489/>
- [22] The ALICE Electronic Logbook, <https://alice-logbook.cern.ch/logbook>.
- [23] A. Junique et al., *Upgrade of the ALICE-TPC read-out electronics*, 2010 JINST 5 C12026
- [24] F. Carena et al., *DDL, the ALICE Data Transmission Protocol and its Evolution from 2 to 6 Gb/s*, JINST 10 (2015) no.04, C04008
- [25] Accellera Systems Initiative, SystemC webpage, <http://www.systemC.org>.
- [26] A. Velure, *Upgrades of the ALICE TPC Front-End Electronics for Long Shutdown 1 and 2*, IEEE Trans. Nucl. Sci, Vol. 62, No. 3, June 2015.
- [27] Microsemi Inc., *SmartFusion2 system-on-chip FPGAs datasheet*, rev. 4 (June 2013).
- [28] Microsemi Inc., *SmartFusion2 system-on-chip FPGAs product brief*, rev. 12 (Oct. 2013).
- [29] Glenn F. Knoll, *Radiation Detection and Measurement*, 4th Edition
- [30] V. Balashov, *Interaction of particles and Radiation with Matter*, Springer-Verlag, 1997.
- [31] H. L. Olesen, *Radiation Effects on Electronic Systems*, 1993.
- [32] R D Schrimpf, *Radiation Effects and Soft Errors in Integrated Circuits and Electronic Devices*, 2004.
- [33] JEDEC STANDARD, *Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices*, Technical report, JEDEC Solid State Technology Association, Arlington, VA 22201-3834, Rev. JESD89 (August 2001).
- [34] R. A. Reed et al., *Single event upset cross sections at various data rates*, IEEE Trans. Nucl. Sci., vol. 43, pp. 2862–2866, Dec. 1996.
- [35] F. Irom et al., *Frequency Dependence of Single-Event Upset in Highly Advanced PowerPC Microprocessors*, IEEE Trans. Nucl. Sci., vol. 51, pp. 3505–3509, Dec. 2004.
- [36] A. Morsch et al., *Radiation in ALICE Detectors and Electronic Racks*, ALICE-INT-2002-28 version 1.0.
- [37] G. Spera, *A Space Oddity*, Crosslink summer 2003
- [38] R. G. Alía, *Radiation Fields in High Energy Accelerators and their impact on Single Event Effects*, PhD thesis, Montpellier University, 2014.
- [39] C.C. Foster et al., *Total ionizing dose and displacement-damage effects in microelectronics*, MRS Bulletin 28 (2003).

-
- [40] T. Vanat et al, *A System for Radiation Testing and Physical Fault Injection into the FPGAs and Other Electronics*, in the proceeding of 2015 Euromicro Conference on Digital System Design
- [41] K. Butcher et al, *The International System of Unit (SI) -Conversion Factor for general use*, NIST Special Publication 1038.
- [42] Microsemi Inc., *SmartFusion2 Microcontroller Subsystem User's Guide*, rev. 1 (April 2013).
- [43] Microsemi, *SmartFusion2 and IGLOO2 FPGA High-Speed Serial Interfaces UG0447 User Guide*, rev.5 (July 2015).
- [44] Microsemi Inc., *SmartFusion2 and IGLOO2 Neutron Single Event Effects (SEE)*, rev 2 (August 2015).
- [45] Microsemi Inc., *IGLOO2 and SmartFusion2 65nm Commercial Flash FPGAs, Interim Summary of Radiation Test Results*, Rev. 2 (Oct. 2014).
- [46] SMPS Technology, *Power Transistor Single Event Burnout*, Sep. 2005.
- [47] J. J. Wang et al., *Total Ionizing Dose Effects on Flash-Based Field Programmable Gate Array*, IEEE Trans. Nucl. Sci., vol. 51, no. 6, pp. 3759-3766, 2004.
- [48] J. J. Wang et al., *Investigating and modeling total ionizing dose and heavy ion effects in flash-based field programmable gate array*, Proceeding of Radiation Effects on Components and Systems Workshop, Athens, Greece, 2006.
- [49] Microsemi Inc., *SmartFusion2 and IGLOO2 Clocking Resources User Guide*, rev 4 (July 2015).
- [50] N. Rezzak et al., *Total Ionizing Dose Characterization of 65nm Flash-Based FPGA*, Proceeding of 2014 IEEE Radiation Effects Data Workshop (REDW), Paris, France, 14-18 July 2014.
- [51] E. S. Snyder et al., *Radiation response of floating gate EEPROM memory cells*, IEEE Trans. Nucl. Sci., vol. 36, pp. 2131–2139, Dec. 1989.
- [52] Proton beam facility at the Svedberg Laboratory, <http://www.tsl.uu.se/irradiation-facilities-tsl/PAULA-proton-beam-facility/>.
- [53] Oslo cyclotron laboratory, <http://www.mn.uio.no/fysikk/english/research/about/infrastructure/OCL/>.
- [54] M. Huhtinen et al., *E. Computational method to estimate single event upset rates in an accelerator environment*, Nucl.Instrum.Meth. A450 (2000) 155-172.
- [55] Nuclear Physics Institute, <http://neutron.ujf.cas.cz/>.

-
- [56] A. Velure, *Design, implementation and testing of SRAM based neutron detectors*, Master thesis, University of Bergen, Norway, 2011.
- [57] J. B. Birks, *The Theory and Practice of Scintillation Counting*, Pergamon Press, 1064.
- [58] I²C bus, <https://en.wikipedia.org/wiki/I%C2%B2C>.
- [59] MXIC Macronix International Co. Ltd., *MX29LV640B T/B 64M-BIT single voltage 3V only flash memory datasheet*, rev. 1.4 (Oct. 2009)
- [60] Microsemi Inc., *SmartFusion2 and IGLOO2 Programming User Guide*, rev 5.0 (May 2014).
- [61] Microsemi Inc., *SmartFusion2 Oscillators Configuration*, rev 1(March 2012)
- [62] Microsemi Inc., *DS0097: ProASIC3 Family Flash FPGAs Datasheet*, rev 18 (2016)
- [63] Microsemi SmartFusion2 SoC FPGAs, <https://www.actel.com/fpga/smartfusion2/>.
- [64] T. Toifl et al., *Measurements of Radiation Effects on the Timing, Trigger and Control Receiver (TTCrx) ASIC*, Proceeding of 6th Workshop on Electronics for LHC Experiments, 11-15 Sep 2000. Cracow, Poland
- [65] Analog Devices Inc., *Continuous Rate 10 Mb/s to 675 Mb/s Clock and Data Recovery IC with Integrated Limiting Amp*, revision March 2009.
- [66] C. Torgesen, *Clock and data recovery methods for the readout control unit 2 in ALICE TPC*, JINST 10 (2015) no.04, C04028.
- [67] PDL D Inc., *PLD-1315 and PLD-23XX*, Rev. D (May 2014).
- [68] Marvell, *8E1111 product brief*, Rev. A (Oct. 2013)
- [69] Microsemi, *SmartFusion2 MSS Ethernet MAC Configuration*, Rev. 5-02-00351-0 (Sep. 2012).
- [70] Ketil Røed, *Upgrade of the ALICE TPC FEE online radiation monitoring system*, 2015 JINST 10 P12019.
- [71] K. Røed et al., *First measurement of single event upsets in the readout control FPGA of the ALICE TPC detector*, 2011 JINST 6 C12022.
- [72] G. Spiezia et al., *A new Radmon version for the LHC and its injection lines*, IEEE Trans. Nucl. Sci. 61 (2014) 3424.
- [73] DIM - *A Distributed Information Management System for the Delphi experiment at CERN*, Presented at: IEEE Eight Conference REAL TIME '93 on Computer Applications in Nuclear, Particle and Plasma Physics (Vancouver, June 8-11 1993).
- [74] Anders Oskarsson, RCU2 weekly meeting.
- [75] Actel Inc., *APA750 and A54SX32A LANSCE neutron test report*, white paper edition (Dec.,2003).
- [76] RCU Control Engine: https://wikihost.uib.no/ift/index.php/RCU_ControlEngine.

-
- [77] Microsemi Inc., *RTG4 FPGA System Controller UG0576 User Guide*, rev. 2 (April 2015).
- [78] C. González Gutiérrez, *The ALICE TPC Readout Control Unit*, Proceeding of 2005 IEEE Nuclear Science Symposium Conference.
- [79] J. Alme, *A Trigger Based Readout and Control System operating in a Radiation Environment*. PhD thesis, University of Bergen, Bergen, 2008.
- [80] Texas Instruments Inc., *TLK2501 (ACTIVE) 1.5 to 2.5 Gbps Transceiver*, rev. 1 (June 2003).
- [81] CERN PH/ED, *ALICE TPC Board Controller*, rev 2.3 (Dec. 2004).
- [82] Texas Instrument Inc., *INA226 High-Side or Low-Side Measurement, Bi-Directional Current and Power Monitor with I²C Compatible Interface*, Rev. A (August 2015).
- [83] I. N. Torsvik, *Radiation Testing of RCU2 Components*, Master Thesis, University of Bergen, June 2014
- [84] Microsemi Inc., *Smartfusion2 Starter Kit Guide*, release 1.9.1 (Feb. 2012).
- [85] Microsemi Inc., *Using Identify with Libero SoC v11.7 TU0071 Tutorial*, Rev 2 (April 2016)
- [86] M. Berg, *Field Programmable Gate Array (FPGA) Single Event Effect (SEE) Radiation Testing*, NASA Electronic Parts and Packaging (NEPP); and Defense Threat Reduction Agency Under IACRO #11-4395 (Feb 2012).
- [87] L. D. Edmonds, *Proton SEU Cross Sections Derived from Heavy-Ion Test Data*, IEEE TRANSACTIONS ON NUCLEAR SCIENCE, VOL. 47, NO. 5, OCTOBER 2000
- [88] Avago Inc., *HFBR-1312TZ transmitter HFBR-2316TZ receiver 1300nm fiber optic transmitter and receiver datasheet*, AV02-1500EN edition, Jan. 2012.
- [89] Xilinx Inc., *Virtex-6 Family Overview*, rev 2.5 (August 2015).
- [90] Xilinx Inc., *Integrated Bit Error Ratio Tester 7 Series GTX Transceivers v3.0 Product Guide*, version 3.0 (June 2016).
- [91] AVAGO In., *AFBR-57D7APZ Digital Diagnostic SFP, 850 nm, 8.5/4.25/2.125 GBd Low Voltage (3.3 V) Fiber Channel RoHS Compliant Optical Transceiver*, AV02-1143EN, Jan. 2013.
- [92] T. Higuchi et al., *Study of Radiation Damage in Front-End Electronics Components*, Proceeding of 18th IEEE-NPSS Real Time Conference, 2012.
- [93] K. Røed et al., *Upgrade of the ALICE TPC FEE online radiation monitoring system*, 2015 JINST 10 P12019.
- [94] K. Røed et al., *Irradiation tests of the complete ALICE TPC Front-End Electronics chain*, Proceedings of the 11th Workshop on electronics for LHC and future experiments, Sept. 2005, Heidelberg, Germany, Page(s): 165-169, ISBN 9290832622

-
- [95] G. Spiezia et al., *A new Radmon version for the LHC and its injection lines*, IEEE Trans. Nucl. 61 (2014) 3424.
- [96] *BusyBox User Guide*, <https://wikihost.uib.no/ift/images/0/01/BusyBoxUserGuide.pdf>.
- [97] D. Evans et al., *The ALICE Central Trigger System*, Proceedings of the 14th IEEE-NPSS Real Time Conference, Stockholm, 4-10 June 2005.
- [98] S.H.I.Barboza et al, *SAMPA chip: a new ASIC for the ALICE TPC and MCH upgrades*, JINST C02008
- [99] F. Merkelov, *Design Techniques for Implementing Highly Reliable Designs using FPGAs*, Microsemi Space Forum Russia, Nov. 2013.
- [100] Microsemi Inc., *DS0128 Datasheet IGLOO2 FPGA and SmartFusion2 SoC FPGA*, rev 11 (October 2016)
- [101] N. Tosi et al, *The CMS Beam Halo Monitor electronics*, Journal of Instrumentation, Volume 11, February 2016
- [102] V.Bocci, *Architecture of the LHCb muon Frontend control system upgrade*, in the proceeding of 2015 IEEE Nuclear Science Symposium and Medical Imaging Conference.
- [103] P. Antonioli et al, *Design and test of a GBTx based board for the upgrade of the ALICE TOF readout electronics*, in the proceeding of IEEE 2016 Real Time Conference.

Appendix A. List of publications

A.1 As main contributor

C. Zhao et al., *First irradiation test results of the ALICE TPC Readout Control Unit 2*, Journal of Instrumentation, Volume 10, January 2015.

C. Zhao et al., *First performance results of the ALICE TPC Readout Control Unit 2*, Journal of Instrumentation, Volume 11, January 2016.

C. Zhao et al., *Performance of ALICE PHOS trigger and improvement for Run2*, Journal of Instrumentation, Volume 8, December 2013.

A.2 As collaborator

J. Alme et al., *RCU2 — The ALICE TPC readout electronics consolidation for Run2*, Journal of Instrumentation, Volume 8, December 2013.

K. Røed et al., *Upgrade of the ALICE TPC FEE online radiation monitoring system*, Journal of Instrumentation, Volume 10, December 2015

Additionally, a total number of 116 publications from February 2013 to present are listed where credited as part of the ALICE Collaboration or the ALICE TPC Collaboration (based on results from SPIRES-HEP Search).

Appendix B. List of Abbreviations

ADC	Analogue to Digital Converter
ALICE	A Large Ion Collider Experiment
AHB	Advanced High-performance Bus
ALTRO	ALICE TPC Readout
APB	Advanced Peripheral Bus
ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal LHC ApparatuS
CDH	Common Data Header
CDR	Clock and Data Recovery
CHRDO	Channel Readout
CERN	Conseil Européen pour la Recherche Nucléaire
CMS	Compact Muon Solenoid
CMOS	Complementary Metal-Oxide-Semiconductor
DAQ	Data Acquisition
DCS	Detector Control System
DDL	Detector Data Link
DDR	Double Data Rate
DSTB	Data Strobe
EEPROM	Electrically Erasable Programmable Read-only Memory
EPCS	Extra-long Physical Coding Sub-layer
EVLRO	Event Length Readout
FEC	Front-end Card
FEE	Front-end electronics

FIFO	First In First Out
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
GPIO	General Purpose Input Output
HLM	Hit List Memory
PC	Inter-Integrated Circuit
ISP	In-System Programming
JTAG	Joint Test Action Group
LAN	Local Area Network
LHC	Large Hadron Collider
LHCb	LHC beauty
MBU	Multiple-Bits Upsets
MOSFET	metal–oxide–semiconductor field-effect transistor
MSS	Microcontroller Subsystem
MTBF	mean time between failure
PC	Personal Computer
PHY	Physical Layer
PRBS	Pseudo Random Binary Sequence
PLL	Phase-locked loop
ROLM	Readout List Memory
RCU	Readout Control Unit
RPINC	Readout Pointer Increment
SCEVL	Scan Event Length
SECDED	Single Error Correction and Double Error Detection
SEE	Single Event Effect

SEGR	Single Event Gate Rupture
SEL	Single Event Latch-up
SERDES	Serializer/deserializer
SET	Single Event Transient
SEU	Single Event Upset
SFP	Small Form-factor Pluggable
SIGMII	Gigabit Media Independent Interface
SIU	Source Interface Unit
SoC	System on Chip
SPI	Serial Peripheral Interface Bus
SRAM	Static Random-Access Memory
TID	Total Ionizing Dose
TPC	Time Projection Chamber
TTC	Trigger, Timing and Control
UART	Universal Asynchronous Receiver/Transmitter
VHDL	VHSIC (Very High Speed Integrated Circuit) Hardware Description Language

Appendix C. Fluence calculation

At the Oslo Cyclotron, the correlating factor (f_{co}) between the scintillator counts (N_{scint}) and the number of the SEUs on the radiation monitor (N_{SEU}) were calculated during the calibration tests. Therefore, N_{SEU} in real tests can be estimated by multiplying this correlating factor with the scintillator counts. Cross-section of the radiation monitor (CS_{rad_mon}) is known as $1.14 \times 10^{-6} \text{ cm}^2/\text{device}$ [56]. So the fluence can be calculated with the equation C.1.

$$Fluence = \frac{N_{SEU}}{CS_{rad_mon}} = \frac{N_{scint} * f_{co}}{CS_{rad_mon}} \quad (C.1)$$

The uncertainty of the fluence can be calculated with

$$\sigma_{CS} = \sqrt{\sigma_{lin}^2 + \sigma_{fco}^2 + \sigma_p^2} \quad (C.2)$$

where σ_{lin} is the uncertainty introduced due to the nonlinear response of the scintillator to the intensity of the beam; σ_{fco} is the uncertainty resulting from finding the correlating factor between the number of scintillator counts and the number of SEUs in radiation monitor; σ_p is the uncertainty of positioning of the device from the exit of the beam. According to [56], these three uncertainties are 10%, 10% and 5%, respectively. Thus, the uncertainty of the fluence at the Oslo Cyclotron is calculated to be 15%.

At the Svedberg Laboratory, the exposure on a monitor in the control room was followed and total fluence can be calculated by using a conversion factor provided by the Svedberg Laboratory. According to the information from the Svedberg Laboratory, the uncertainty of the fluence is 15%.

Appendix D. RCU2 Data Format




Figure D-1 CDH words of RCU2




Figure D-2 RCU2 payload words

RCU2 Trailer Word 1:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	0	0	0	0	Payload Size																								

RCU2 Trailer Word 2:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0	
1	0	0	0	0	1	0	0	TRX Timeout		DSTB stuck	TRSF stuck	ACK stuck	No data	Altro Error	TRX Timeout	DSTB stuck	TRSF stuck	ACK stuck	No data	Altro Error	TRX Timeout	DSTB stuck	TRSF stuck	ACK stuck	No data	Altro Error	TRX Timeout	DSTB stuck	TRSF stuck	ACK stuck	No data	Altro Error
ID = 0x84								Branch BO				Branch BI				Branch AO				Branch AI												

RCU2 Trailer Word 3:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
100010						Address Err	Padding Err	Payload Err	Error Count										Address Err	Padding Err	Payload Err	Error Count									
ID						Branch AO										Branch AI															

RCU2 Trailer Word 4:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0																						
1					0					0					0					1					1					Address Err	Payload Padding Err	Error Count										Address Err	Payload Padding Err	Error Count									
ID						Branch BO															Address Err	Payload Padding Err	Branch BI																														

RCU2 Trailer Word 5:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	1	0	0	Unused														ACTIVE FEC A											

RCU2 Trailer Word 6:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	1	0	1	Unused														ACTIVE FEC B											

RCU2 Trailer Word 7:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	1	1	0	ALTRO DPCFG1																									

RCU2 Trailer Word 8:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	0	0	1	1	1	ALTRO DPCFG2																			SCLK-L1						

RCU2 Trailer Word 9:

31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
1	1	1	0	0	0	Firmware Version										RCU ID										Trailer Length					

Figure D-3 RCU2 Trailer words

Appendix E. Screenshots and test results

E.1 SEU counts of SRAM tests




Figure E.1-1 SEUs and fluence for the SRAM test in campaign No.2




Figure E.1-2 SEUs and fluence for the SRAM test in campaign No.3

E.2 Screenshots of tests




Figure E.2-1 Measurement of RCU2 signals. (a) Sampling clock and readout clock. (b) Quality of data lines. (c) L1 and L2 triggers. (d) Broadcast command




Figure E.2-2 Measurement of CHRDO for an empty channel. (a) Non-conservative bus switching (same as RUC1). (b) Conservative bus switching.




Figure E.2-3 Measurement of CHRDO for the number of samples as 10. (a) Non-conservative bus switching (same as RUC1). (b) Conservative bus switching.

E.3 Readout speed benchmark

Table E.3-1 Readout speed of single event (partition 1)

Number of samples	RCU1	RCU2 2.125 Gbps (conservative)	RCU2 4.25 Gbps (conservative)	RCU2 3.125 Gbps (conservative)	RCU2 3.125 Gbps (optimized)
10	1.1199	0.6727	0.6727	0.6727	0.5832
20	1.2399	0.8071	0.8071	0.8071	0.6504
30	1.4000	0.8743	0.8743	0.8743	0.6953
40	1.7199	0.9862	0.9863	0.9863	0.7625
50	1.9599	1.1822	1.0535	1.0535	0.8074
60	2.2799	1.3703	1.1655	1.1655	0.9216
70	2.4400	1.5583	1.2327	1.2327	1.0418
80	2.8398	1.8063	1.3671	1.3671	1.2017
90	3.0000	1.9935	1.4343	1.4343	1.3234
100	3.3200	2.1783	1.5463	1.5463	1.4484
110	3.5599	2.4251	1.6135	1.6135	1.6134
120	3.8798	2.5948	1.7255	1.7354	1.7385
130	4.0399	2.7799	1.8431	1.8604	1.8619

140	4.4397	3.0262	2.0176	2.0391	2.0286
150	4.5999	3.2116	2.1471	2.1737	2.1538
160	4.9198	3.3962	2.2782	2.3097	2.2790
170	5.1600	3.6490	2.4479	2.4828	2.4421
180	5.4796	3.8364	2.5784	2.6182	2.5702
190	5.6398	4.0269	2.7079	2.7494	2.6998
200	6.0397	4.2850	2.8802	2.9285	2.8710
250	7.2397	5.3021	3.5632	3.6314	3.5478
300	8.6796	6.3868	4.2904	4.3821	4.2725
350	9.9596	7.4706	5.0126	5.1243	4.9887
400	11.3192	8.4919	5.6981	5.8335	5.6691
450	12.5995	9.5698	6.4201	6.5749	6.3870
500	14.0392	10.6505	7.1421	7.3210	7.1045
550	15.2397	11.6597	7.8188	8.0177	7.7771
600	16.6794	12.7409	8.5434	8.7655	8.5026
650	17.9592	13.8172	9.2638	9.5069	9.2201
700	19.3193	14.8340	9.9472	10.2139	9.8964
750	20.5989	15.9127	10.6646	10.9543	10.6104
800	22.0386	16.9898	11.3880	11.6981	11.3252
850	23.2387	18.0074	12.0651	12.3887	12.0130
900	24.6793	19.0937	12.7898	13.1411	12.7234
950	25.9590	20.15659	13.5094	13.8845	13.4374
1000	27.3193	21.1931	14.1900	14.5921	14.1165

E.4 Procedure of test prior to mass production

Step 1. Make sure the oscillator is present

Do a visual inspection of the board to make sure the oscillator is actually present, and that it is soldered (red circle). The card number is indicated in the yellow rectangle below.




Figure E.4-1 Inspection of oscillator

Step 2. Update and configure the RCU2

- (1) Update Linux with the follow steps:

Login the RCU2 -> Upload bootloader -> Update Uboot -> Update Linux -> reboot the RCU2.

- (2) Configure the RCU2 and the FECs with the following steps:

Login the RCU2 -> configure RCU2 and FEC -> Initialize DDL to 3.125Gbps -> Check DDL status.

Step 3. Test the readout

Initialize Trigger Crate -> Start data checker -> start data-taking for a few hours -> stop data checker
-> stop trigger -> stop data-taking.

If the number of samples is set to 1000, the sub-event rate and the Byte recorded rate should be ~66 Hz and 283 MB/s, as shown in Figure E.4-2.

LDC status display	
LDC name	alone1dc
host	pcphaid001.cern.ch
Number of equipments	1
Number of triggers	3366
Current Trigger rate	66.0
Average Trigger rate	35.8
Number of sub-events	3366
Sub-event rate	66
Sub-events recorded	3368
Sub-event recorded rate	66
Bytes recorded	14434027504
Byte recorded rate	283.020 MB/s
Nb. evts w/o HLT decision	0

Figure E.4-2 Screenshot of data-taking status

E.5 Commissioning of the RCU2

From January to June in 2015, 6 RCU2s with DDL2 link at 2.125 Gbps were commissioned in one of the 36 TPC sectors. The appearance of these installed RCU2s is shown in the subplot (b) of the Figure E.5-1.

Readout was tested without the presence of magnetic field. Fixed pattern of data was written into the on-board memories of the FECs and read back by the RCU2. Several TB of data was looped and no error of the data or stop of the readout was detected.

The trigger reception, the Monitoring and Safety Module, the Ethernet link and the DDL2 link were working stable. The ISP programming was generally operational. It exited prematurely with a probability of ~10% to 15%, but a retry always worked.

The Linux system was observed in radiation. No reboots or freezes were seen on these RCU2s. It is hard to draw any conclusion on the stability of the Linux due to the low statistics. As a reference, about 10 reboots were observed on the other 210 RCU1s during the same running period of ~6 months.

As shown in the subplot (a) of the Figure E.5-1, operational temperature of the RCU2 with cooling system was around the normal range of 20 degrees.




Figure E.5-1 The first 6 installed RCU2⁷². (a) Temperature of the installed RCU2. (b) Appearance of the installed RCU2.

In early 2016, 6 RCU2s with DDL2 link at 3.125 Gbps were commissioned in one of the 36 TPC sectors⁷³. Fixed pattern of data was filled into the on-board memories of the FECs and then read back by the RCU2. ~19 TB of data was looped in the presence of magnetic field and ~954 PB of data was looped without the presence of magnetic field. No stop of readout was observed during the data-taking. ~10% of the captured data was checked. No error was detected in the data taken without radiation. Nevertheless, two errors were found in the data taken in radiation. Given that the max event size in Run2 is ~81 MBytes, this refers to an error rate of $\sim 8.1 \times 10^{-5}$ bits/event.

Afterwards, all the 216 RCU2s were installed at TPC. Several TB of data was looped on each sector without the presence of magnetic field (refer to Figure E.5-2). The TTC and DAQ interface worked in a stable manner, and no data error or stop of readout was observed. The DCS interface (Figure E.5-4), the Monitoring and Safety Module (Figure E.5-5), the RadMon (Figure E.5-3), and the power supply system (Figure E.5-6) were working stably. The ISP programming still failed in some cases but a retry always worked. At the time of writing (November 2016), the RCU2s are taking data in a stable manner in p-Pb collisions.

⁷² Pictures from Christian Lippmann (Christian.Lippmann@cern.ch)

⁷³ These commissioning action were mainly done by Christian Lippmann (Christian.Lippmann@cern.ch)




Figure E.5-2 Data loop in sector (six readout partitions)




Figure E.5-3 Radiation Monitor of the RCU2




Figure E.5-6 Check the power of installed partitions (colored purple)