Universitatea
Transilvania
din Brașov

ȘCOALA DOCTORALĂ INTERDISCIPLINARĂ

Facultatea: Inginerie Electrică și Știința Calculatoarelor

Ing. Ștefan POPA

Circuitul integrat ROC pentru experimentul ATLAS de la LHC

The Read Out Controller ASIC for the ATLAS Experiment at LHC

REZUMAT / SUMMARY

Conducător ştiinţific

Prof.dr.ing. Mihai IVANOVICI

BRAȘOV, 2021

Universitatea Transilvania din Brașov

D-lui (D-nei) .....

COMPONENȚA

Comisiei de doctorat

Numită prin ordinul Rectorului Universității Transilvania din Braşov
Nr. ..... din .....

PREŞEDINTE:	Conf. dr. ing. Carmen GERIGAN
CONDUCĂTOR ŞTIINȚIFIC:	Prof. dr. ing. Mihai IVANOVICI
REFERENȚI:	Senior research fellow Lorne LEVINSON
	Prof. dr. ing. Gheorghe ȘTEFAN
	Prof. dr. ing. Liviu GORAȘ

Data, ora şi locul susținerii publice a tezei de doctorat: 27 septembrie 2021, ora 10:00, sala NII1, Corp N, Strada Politehnicii 1, Brașov.

Eventualele aprecieri sau observații asupra conținutului lucrării vor fi transmise electronic, în timp util, pe adresa *stefan.popa@unitbv.ro* .

Totodată, vă invităm să luați parte la şedința publică de susținere a tezei de doctorat.

Vă mulțumim.

The Read Out Controller ASIC for the ATLAS Experiment at LHC

Summary of the Ph.D. Thesis

Transilvania University of Brașov, România

Ph.D. Candidate: Ștefan Popa
Ph.D. Supervisor: Prof. Mihai Ivanovici

2021

To _____

Ștefan Popa

is kindly inviting you to the public

Defence *of his* Doctoral Thesis, entitled

"The Read Out Controller ASIC for the ATLAS Experiment at LHC",

on the 27th of September 2021 at 10:00 AM,

at Transilvania University of Brașov, room NII1,

1, Politehnicii street, Brașov.

Examination Jury:

President:	Conf. dr. ing. Carmen GERIGAN
Coordinator:	Prof. dr. ing. Mihai IVANOVICI
Member:	Senior research fellow Lorne LEVINSON
Member:	Prof. dr. ing. Gheorghe ȘTEFAN
Member:	Prof. dr. ing. Liviu GORAȘ

Abstract

The ATLAS Experiment at LHC is used for fundamental research in particle physics. For its HL-LHC TDAQ system upgrade, new ASICs were developed. The Read-Out Controller (ROC) is such an on-detector radiation-tolerant ASIC that acts as a concentrator, buffer, filter and real-time data packet processor for the new end-cap muon detectors. The thesis presents its elaboration, implementation, quality assurance and control with emphasis on real-world experimental results. The IC is implemented in a 130 nm CMOS technology, resulted in a square die of 22.5 mm^2 with 232 pads and is packaged as 16×16 BGA. The design and its performance model were validated using custom analog and digital functional FPGA-based test setups. The digital test setup emulates the asynchronous chip context, employs optimizations and automatic clock and data synchronization and is used for mass-testing. The ROC's operation was tested while controlled ultrafast neutron beams were incident to its die. Its tolerance to the induced SEUs was evaluated and predictions for the operating environment were made. A proposed implementation of an FPGA Integrated Logic Analyzer that mitigates the observed limitations and constraints of the existing ones is included. The ROC design passed reviews within the ATLAS Collaboration and is included in the TDAQ system.

Rezumat

Experimentul ATLAS de la LHC este utilizat pentru cercetare fundamentală în fizica particulelor. Pentru modernizarea HL-LHC a sistemului său TDAQ, noi ASIC-uri au fost dezvoltate. Circuitul Read-Out Controller (ROC) este un astfel de ASIC tolerant la radiație cu rolul de concentrator, amortizor, filtru și procesor în timp real de pachete de date de la noile detectoare de miuoni. Această teză prezintă elaborarea, implementarea, asigurarea și controlul calității sale cu accent pe rezultate experimentale din lumea reală. Circuitul integrat este implementat într-o tehnologie CMOS de 130 nm, a rezultat într-o pastilă de siliciu pătrată de 22.5 mm^2 cu 232 de pini și este încapsulat ca BGA 16×16. Implementarea și modelul de performanță au fost validate utilizând sisteme de testare funcțională analogică și digitală personalizate și bazate pe FPGA. Sistemul de testare digitală emulează contextul asincron al cipului, conține optimizări și metode de sincronizare automată a semnalelor de date și ceas și este utilizat la testarea în masă. Funcționarea ROC-ului a fost testată sub incidența unor fascicule controlate de neutroni ultra-rapizi. Toleranța sa la efectele imediate de tip SEU ale radiației nucleare au fost evaluate și estimări pentru mediul de operare au fost realizate. O propunere de implementare pentru un Integrated Logic Analyzer pentru FPGA-uri care estompează limitările și constrângerile celor existente este inclusă. Circuitul ROC a trecut cu succes evaluările din comunitatea ATLAS și este inclus în sistemul TDAQ al experimentului.

Contents

1 Introduction **6**
1.1 The NSW-ATLAS-LHC-CERN context 6
1.2 Objectives 10

2 The Read-Out Controller (ROC) **11**
2.1 ROC's context and architecture 11
2.2 ROC's layout and package 14
2.3 ROC's steady-state model 15

3 ROC Testing **17**
3.1 The Quality-Control Digital ROC Test 18
3.2 Experimental testing results 21

4 Immunity to radiation-induced faults **24**
4.1 Test setup 25
4.2 Irradiation tests results 25

5 An application **29**

6 Conclusions **32**

References **36**

Cuprins

1 Introducere **6**
1.1 Contextul NSW-ATLAS-LHC-CERN 6
1.2 Obiective 10

2 Read-Out Controller (ROC) **11**
2.1 Context și arhitectură 11
2.2 Implementare fizică și încapsulare 14
2.3 Modelul de echilibru 15

3 Testarea ROC-ului **17**
3.1 Mediul de testare digitală 18
3.2 Rezultate experimentale 21

4 Imunitatea la efectele radiației nucleare **24**
4.1 Mediul de testare 25
4.2 Rezultatele testelor de iradiere 25

5 O aplicație **29**

6 Concluzii **32**

Referințe **36**

1 Introduction

The work presented in this thesis summary relates to the ATLAS (A Toroidal LHC Apparatus) Experiment [1] at the Large Hadron Collider (LHC) [2] particle accelerator, operated by the European Organization for Nuclear Research (CERN) which is situated near Geneva, Switzerland. CERN has in its repertory many important scientific achievements like the discovery of the Higgs boson [3], the production and maintenance of antihydrogen atoms [4] and the birth of the World Wide Web (WWW) information system [5]. In this introductory section, the thesis context is explained and its objectives are presented.

1.1 The NSW-ATLAS-LHC-CERN context

Particle physics is a branch of physics that studies the structures of matter and radiation and their interactions [6]. The notion of elementary particle denotes the subatomic particles with no substructure [6]. The Standard Model (SM) of particle physics is the theory that classifies all known elementary particles and describes three out of the four known fundamental forces [7], [6]. To validate the theoretical assumptions related to particle physics and the SM, particle accelerators [8] have been build and used since the 1930s [9]. A particle accelerator is a system that forms and transfers energy into well-defined beams of particles using electromagnetic fields. The electrical fields provide the acceleration while the magnetic fields concentrate and direct the beam. Thus, the particles contained in the beam reach very high speeds and are focused as much as possible. The target of the particle beam is a piece of material or another accelerated particle beam traveling in the opposite direction. Besides fundamental research, particle accelerators have many other uses: e.g. particle therapy (treatment of cancer), radiation sterilization of medical devices, ion implantation (semiconductor device fabrication), nuclear physics (production of isotopes), etc.

The performance of a particle accelerator is determined using two metrics: the energy transferred to the particle beam and the luminosity [10]. The first refers to the kinetic energy of each beam particle gained in the accelerator and is measured in eV (i.e. electronvolt; 1 eV = $1.602176634 \times 10^{-19}$ J). The higher the energy is, the higher the probability of generating a particle with a higher mass and the possibility of reaching further into the structure of matter [7]. The luminosity is a performance metric defined as the ratio between the number of particle interactions produced in a set time and the cross-section (σ) of the interaction [10]. In physics, σ represents the probability that a specific event will occur when a radiant phenomenon intersects a localized object or variation of density. It has the same unit of measure as the area of a surface since it represents the transverse size of the targeted object that the radiant phenomenon must hit for the process to take place. Thus, the unit of measure for the luminosity is $cm^{-2} \cdot s^{-1}$. Because the luminosity can vary in time, the final figure that reflects the number of observed events and as a result the quantity of produced data is the integrated luminosity. It represents the total data quantity over the sensitive time related to an experiment. As a reference,

the LHC reached a peak instantaneous luminosity of $L = 2 \times 10^{34}$ cm$^{-2} \cdot$ s^{-1} for the proton collisions in 2018 and accumulated 160 fb^{-1} (i.e. femtobarn, 1 fb $= 10^{-43}$ m^2) of proton collisions data between 2015 and 2018 [11].

The LHC is the largest and highest-energy synchrotron in the world [2]. A synchrotron is an electrodynamic (i.e. employing varying electromagnetic fields) particle accelerator in which the particle beam follows a closed-loop trajectory and the intensity of the magnetic field that bends the beam on its path is synchronous to the beam's energy. LHC is situated in an underground tunnel, at 170 m from the ground level and has a circumference of 27 km [2], beneath the border of Switzerland and France, near Geneva. Within the accelerator, proton (i.e. p^+) or ion beams travel in opposite directions. The p^+ beams are organized in up to 2808 bunches distributed along the circumference approx. every 7.5 m, each containing roughly 1.15×10^{11} particles [2] [12]. The energy reaches 7 TeV per proton (they are accelerated to 99.9999991 % of the speed of light in vacuum). Thus the bunches are spaced at approx. 25 ns and their crossing rate (i.e. Bunch Crossing - BC) is 40 MHz. Due to practical reasons, there are gaps in their pattern so, in one second, an average of 3×10^7 bunches cross [13] [12]. Up to 40 collisions are produced in each BC [13]. This amounts to 10^9 collisions every second [14] [13] from which several hundreds are of interest. The bunches travel like this within the tunnel between 10 and 24 hours [13]. New ones are then formed and accelerated.

The beams intersect each other in four collision points where detectors are installed, providing the necessary signals for the determination of the trajectory, energy and electrical charge of the resulted particles. The four experiments are: ALICE (A Large Ion Collider Experiment) [15], ATLAS [1], CMS (Compact Muon Solenoid) [16] and LHCb (LHC beauty) [17]. The ATLAS and CMS are general-purpose experiments and have high luminosity, while ALICE and LHCb are dedicated to the physics of heavy ions and the bottom quark, respectively.

The ATLAS detector, depicted in Figure 1, has a cylindrical shape around the interaction point such that the particle beams are perpendicular to the bases in their centers. The height (i.e. diameter of the base) is 25 m and the length (i.e. cylinder's height) is 44 m. It can characterize any particles produced from the beams collision (i.e. determine their masses, momentum, lifetimes, charges, spins and energies). Thus it contains different types of detectors, organized in layers of cylindrical shape, called *barrel* sections. The disk detector assemblies, parallel to cylinder bases, which close a barrel section are referred to as *end-cap* sections. The detector has no blind spots, being hermetic.

The main constituents of the ATLAS detector are the Inner Detector (IDET), the calorimeters, the Muon Spectrometer (MS) and the magnet system. All these sub-systems are divided into multiple layers and are complementary: the IDETs determine the trajectory while the calorimeters measure the energy of the particles that are stopped within them. The highly penetrating muons are measured by the muon system (i.e. tracking and energy). The magnet system bends the trajectories of the resulted charged particles so that their momentum can be determined.

Considering the 40 MHz BC rate and the approx. 100 million readout chan-

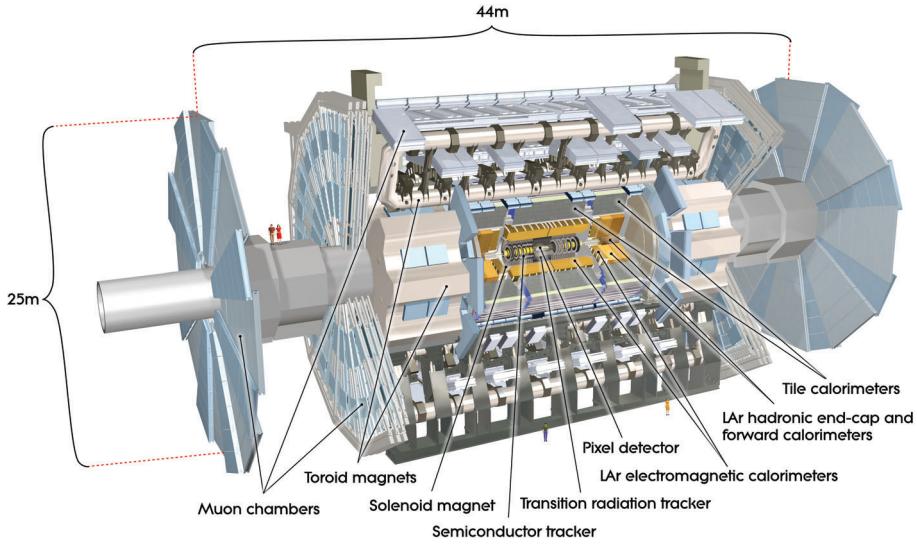

Figure 1: The ATLAS detector [18].

nels of the ATLAS detector, a throughput in the order of PB/s is produced [19]. Since this amount is currently impossible to transmit, process and store, the ATLAS Trigger and Data Acquisition (TDAQ) system interprets the detector signals and determines in real-time the Regions of Interest (ROIs) for each BC. Event filters reduce the rate further since only a few of the collisions are of interest. Based on these decisions, the TDAQ system converts the selected detector signals into a dataset, resulting in a data rate of approx. 1 GB/s [14].

The current LHC update and maintenance process aims at implementing the first step (Phase-I) of the High Luminosity LHC (HL-LHC) project [20]. The main objective of HL-LHC is the accumulation of at least 3000 fb^{-1} integrated luminosity over 10 years of operation. The modifications of the ATLAS detector related to HL-LHC are presented in [21]. Phase-I updates are scheduled to end in 2022, being followed by an operational run called Run 3. After that, the Phase-II update will be implemented with HL-LHC beginning operation in 2027. The luminosity increase means that the detectors and the associated on-detector electronics will be subjected to increased nuclear radiation. They must also cope with the higher collision rates. Thus, the detector technology, the TDAQ systems and the controlling and monitoring software tools must be optimized and adapted. Within ATLAS, the most changes will be implemented for the IDET, the calorimeters end-caps and the MS. In Phase-I, within the ATLAS MS, the *Small Wheel* inner end-cap regions will be replaced with the *New Small Wheel (NSW)* [22] containing approx. 2.45×10^6 detectors of Micro-Megas (MM) and small-strip Thin Gap Chamber (sTGC) types and new electronics. The first filtering element is the hardware-based Level-1 (L1) trigger while the

second filtering element will be the software-based High-Level Trigger (HLT). While both new muon detector types contribute to the trigger formation and the precision tracking, the MM detectors are used mainly for the track reconstruction due to their high spatial resolution (up to 100 μm) [23] and the sTGC for determining the trigger candidates due to their capability of identifying a singular BC [22]. In Phase-II, the L1 trigger will become Level-0 (L0) and a new L1, also hardware-based, will implement a more complex selection.

The ATLAS TDAQ system is paced by a 40 MHz clock signal called the BC or LHC clock synchronous to the collisions. The BCs are tagged with a BCID (BC IDentification - bunch count within an LHC orbit) and an orbit counter called OrbitID (Orbit IDentity). The ATLAS Central Trigger Processor determines the Regions of Interest ROIs within the BCs of interest, based upon the detector data fed through detector-specific electronics, processed by dedicated trigger processors and transmitted on dedicated links. Through the Time Trigger and Control (TTC) system the LHC clock is distributed; the BCIDs, OrbitIDs and trigger counters from all ATLAS systems and subsystems are synchronized and the BCs of interest are selected within the ROIs. The read-out system is responsible for managing the detector data fed by the detector-specific electronics and based upon the received trigger decisions, constructing, buffering and transmitting event fragments tagged with the above-mentioned identifiers. The time interval between the selected BC and the arrival of the corresponding selecting trigger command is called trigger latency.

The block diagram of the associated NSW electronics is depicted in Figure 2. The main radiation-tolerant Application Specific Integrated Circuit (ASIC) for reading and interpreting the signals from both types of muon detectors is the *VMM*[1], which represents an Amplifier Shaper Discriminator. One VMM3 (i.e. third version) chip contains 64 distinct detector read-out channels providing peak amplitude and time measurements via configurable charge amplification, discrimination and precise Analog-to-Digital Conversion. The VMM3's read-out digital logic buffers and aggregates the data from its channels based upon the received L0 trigger selection commands and transmits the resulting L0 events to the Read-Out Controller (ROC) ASIC. On the trigger path, the VMM3 can provide faster but coarser measurements for the detector signals. One ROC chip collects and buffers the L0 events from up to eight VMM3 ASICs (i.e. data from up to 512 NSW detectors), implements the L1 selection if used and forms more complex events with the aggregated data. Thus, it reduces the number of necessary data links on the read-out path. As it is highly configurable, it allows the optimization of bandwidth utilization. *All the contributions of this thesis are related to the ROC ASIC, therefore this chip is extensively described. At least 4875 ROC chips will be installed and will function concurrently within the NSW TDAQ system.*

The MM and sTGC readout path consists of the following components: VMM [25], ROC [26], Level-1 Data Driver Card (L1DDC) [27] [28] GigaBit Transciever (GBTx) [29] [30] and Front End LInk eXchange (FELIX) [31]. The

[1] not an acronym.

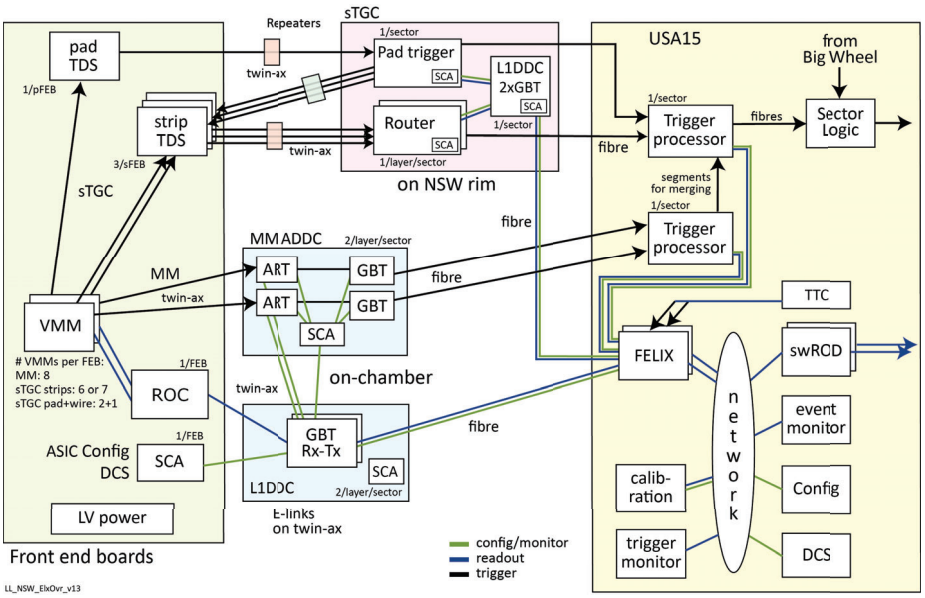

Figure 2: NSW TDAQ system overview [24].

MM trigger scheme consists of: VMM, Address in Real-Time (ART) [32], ART Data Driver Card (ADDC) [33] [34] GBTx, MM Trigger Processor, FELIX and sTGC Trigger Processor. The sTGC trigger path consists of VMM, pad Trigger Data Serializer (TDS) [35], Pad Trigger Extractor board (i.e. Pad trigger in Figure 2), strip TDS, Router board, sTGC Trigger Processor and FELIX.

1.2 Objectives

This Ph.D. thesis represents the culmination of the author's work within the CERN-LHC-ATLAS-NSW context. The objectives were: (i) designing and assuring the quality of the ROC's logic; (ii) correctly implementing it into the targeted 130 nm Complementary Metal Oxide Semiconductor (CMOS) technology (Global Foundries, previously IBM); (iii) functional validation of the manufactured design; (iv) performance measurements for the manufactured design; (v) quality control of the mass-fabricated ROCs; (vi) real-world nuclear radiation qualification for the design; (vii) integration support for the chip; (viii) identification and pursuit of new related research opportunities; (ix) dissemination of results and (x) participation at relevant courses, seminars, workshops and summer and winter schools.

2 The Read-Out Controller (ROC)

The ROC ASIC represents an on-detector custom real-time data packet processor which is a key part of the NSW TDAQ system, on the readout path. The latency requirements for the hardware triggers can be relaxed due to its relatively large buffer spaces, compared to other ASICs from the context. This, in conjunction with the implemented flow and congestion control mechanisms, minimizes data loss. Being a highly configurable concentrator, the ROC can be used to aggregate data from multiple channels with different amounts of throughput. Thus, the bandwidth utilization can be optimized since not all regions of the NSW will produce the same quantity of data. As two hardware-based levels of trigger will be employed within the ATLAS TDAQ system, the ROC will be responsible for the second level of event selection for NSW.

2.1 ROC's context and architecture

The ROC's context within the NSW TDAQ system is depicted in Figure 3. A top-level view of its architecture is included. Its logic is divided into two distinct parts with different purposes: the *analog* part (with a light red background) which is responsible for supplying the internal and external clock signals and for forwarding of the decoded TTC commands and the packet processing logic (depicted on a light yellow background) which is referred to as the *digital* part.

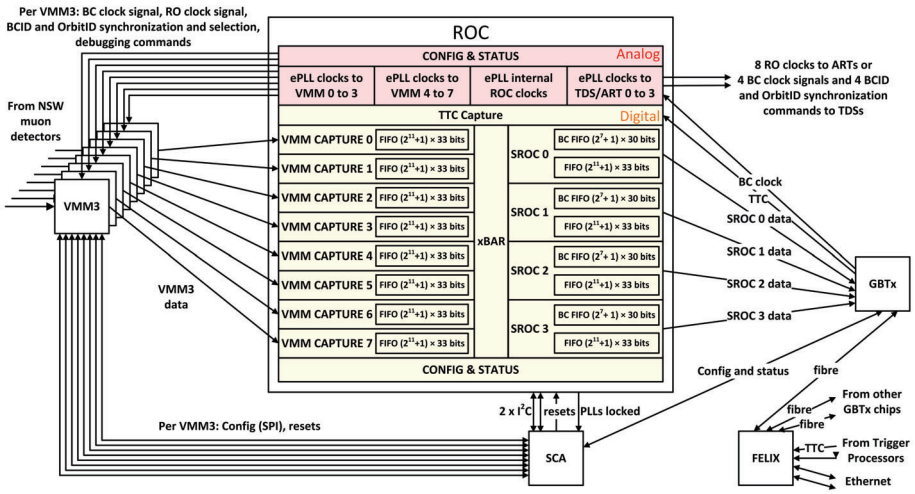

Figure 3: The ROC context within the NSW TDAQ system, its main interfaces and top-level architecture.

The ROC receives 8b10b [36] encoded data representing L0 packets from up to eight VMM3 ASICs on separate input channels called VMM Capture channels or modules. The data from one VMM3 ASIC are DDR (Double Data Rate),

serialized on two SLVS (Scalable Low-Voltage Signaling) [37] transmission lines using a 160 MHz Read-Out (RO) clock signal. The bits are alternatively sent on the two lines. The used electrical interface is called e-link, is radiation-hardened and can cope with data rates up to 320 Mbps (i.e. 320×10^6 bps) [38]. Thus, the resulting total encoded bandwidth is 640 Mbps for each channel. The receiving VMM Capture channels deserialize the incoming data, determine its alignment, decode it, determine its correctness and buffer the L0 packets into queues - First-In-First-Outs (FIFOs).

The ROC input packets that contain VMM3 channel data are called L0 hit packets while the ones without any data are called L0 null-events. Both types start with a 16-bit header word representing the L0 trigger information that caused the formation and transmission of the packet. An L0 null-event contains only the header word while in an L0 hit packet, the header is followed by at least one 32-bit hit word. Since one VMM3 chip is responsible for at most 64 detector channels and each hit word corresponds to one channel, the maximum number of hit words in an input packet is 64. The data fields of the hit word represent signal measurements and flags. The uncoded buffering format of the input packet words is detailed in Figure 4. Both types of words occupy one address in the VMM Capture FIFO. The Most Significant Bit (MSB) signals the End Of Packet (EOP) and is padded by the receive logic. The VMM3 packets are separated by at least two special *K.28.5* 8b10b symbols called *commas*. When data is not available the stream consists of commas. They are used at the receiving end to determine the start position of the 8b10b symbols within the bit stream (i.e. finding the alignment) and are not buffered.

	32	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
header	EOP	V	P	orbit(2)		BCID(12)												16'b0															
hit data	EOP	1	P	0	T	CHAN(6)						ADC (10)										TDC (8)								N	rel BCID(3)		

Figure 4: The buffering format of the input data packets [39].

The ROC receives the TTC stream and the LHC clock signal through the GBTx ASIC. The clock signal is used as reference in the analog part. The TTC stream is organized into bytes whose bits represent commands that synchronize the BC and triggers counters and select the BCs of interest. The bytes are not 8b10b encoded but are directly serialized on a single SLVS DDR lane also using a 160 MHz RO clock signal. The resulting TTC stream is 320 Mbps $= 320 \times 10^6$ bps and is interpreted by the TTC Capture module.

The ROC contains four Sub-ROC (SROC) modules that have distinct VMM Capture channels associated with them through a fully configurable cross-bar module. One SROC can have more than one VMM Capture FIFOs associated with it, but a VMM Capture FIFO cannot be linked to more than one SROCs. The cross-bar module is purely combinational and acts as a multiplexer-demultiplexer. For each L1 trigger an L1 event is formed in each enabled SROC, marked with the information of the selected BC and containing the aggregated and reformatted associated valid L0 data from the associated VMM Capture channels. The L1 events are buffered into an SROC FIFO with the format

depicted in Figure 5. The MSB signals the end of each packet and is not transmitted. As soon as they are ready to be sent (i.e. they are complete), the L1 packets are 8b10b encoded, serialized and transmitted to the GBTx. Each SROC has two SLVS DDR serial output lines that can operate at up to 320 Mbps (i.e. 320×10^6 bps) each. The following configurations are possible: i) 80, 160 or 320 Mbps with only the first line active or ii) 640 Mbps with both lines operating at 320 Mbps and the transmitted bits being interleaved, same as for the VMM3 data. The SROC can saturate the output bandwidth in all four cases. Each SROC contains a TTC FIFO for buffering the information BC identifiers of the L1 triggers. The buffering format is depicted in Figure 6.

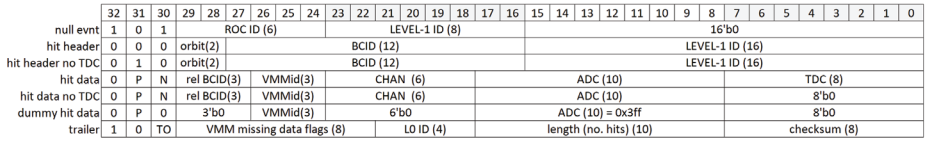

Figure 5: The buffering format of the L1 packets within the SROC FIFOs [39].

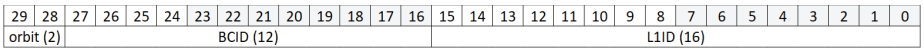

Figure 6: The buffering format of the L1 triggers within the TTC FIFOs [39].

The ROC's digital Configuration and Status module is a register bank interfaced with the digital logic on one side and with an Inter-Integrated Circuit (I^2C) slave on the other. Some of these registers drive signals within the digital part (i.e. configuration) while the other registers supervise relevant signals from the digital logic (i.e. monitoring).

The ROC's internal and externally supplied clock signals are generated by four Phase-Locked Loop (PLL)[2] blocks within the analog part. They all use the 40 MHz TTC BC clock signal as reference. The three ePLLs supplying the external clock signals are slightly modified versions of the design presented in [40]. They include phase-shifting circuits that receive the relevant TTC commands from the digital part (i.e. TTC Capture) and forward them with a configurable phase. All the ePLLs are configured and monitored through a separate register bank associated with an I^2C slave.

The two ROC I^2C slaves are interfaced to the master Slow Control Adapter (SCA) chip on dedicated buses as shown in Figure 3. Several SCA GPIOs (General Purpose Input/Outputs) drive the ROC reset signals and sample its ePLLs locking signals. The SCA ASIC also configures the eight VMM3s through dedicated SPIs (Serial Peripheral Interfaces) and drives their reset signals.

[2]called ePLL within the ATLAS Collaboration as shown in [40].

2.2 ROC's layout and package

The layout of the ROC silicon die in Electronic Design Automation (EDA) representation die is depicted in Figure 7, with highlighted main regions and macroblocks. The die is a square with a side length of 4.744 mm (i.e. 22.5 mm^2 area) and has a border formed by one row of pads. From the 232 pads, 187 are used for Input-Output (IO). The layout is slightly core limited, meaning that the size of the core determines the die size and not the pad number. The total width of the pad filler cells is less than the narrowest IO or power pad.

Similar to its architecture, the core is divided into two areas based on their function: the *analog* and the *digital* parts. The digital part has an area of 13.3 mm^2 from which 64 % represents SRAM - Static RAM (i.e. approx. 8.5 mm^2). Identical buffers are used for the eight VMM Capture and four SROC Packet FIFOs: for each a 0.689 mm^2 dual-port, dual clock domain SRAM with $2^{11} \times 33$ bits. Four dual-port, single clock domain $2^7 \times 30$ bits SRAMs represent the buffers of the four TTC FIFOs. Each has an area of 0.054 mm^2. The rest of the packet processing logic is contained in the remaining area of 4.8 mm^2 and consists of 83,079 combinational gates and 18,458 flip-flops. In Figure 7 the densities of sequential elements (i.e. flip-flops and latches) within the entire ROC can be visualized since they are highlighted with white border.

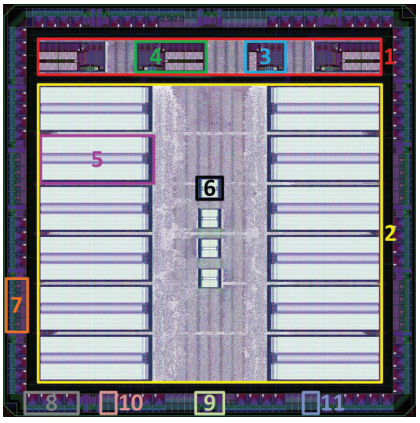

1. The *analog* part
2. The *digital* part
3. The ePLL for the internal clock signals
4. One ePLL for 4 BC (40 MHz) and 4 RO (160 MHz) output clock signals
5. One of the 12 dual-port, dual clock domain $2^{11} \times 33$ bits SRAMs used for the VMM Capture and SROC Packet FIFOs
6. One of the 4 dual-port, single clock domain $2^7 \times 30$ bits SRAMs used for the TTC FIFOs
7. Several input SLVS pads
8. Several output SLVS pads
9. Several power supply pads
10. Two output single-ended pads
11. Two input single-ended pads

Figure 7: ROC's layout with highlighted sequential cells and macroblocks [26].

Initially, the silicon die was wire-bonded on a prototype testing Printed Circuit Board (PCB) as depicted in Figure 8a. An intermediate Quad Flat Package (QFP) with 144 pins, illustrated in Figure 8b, was used for partial design validation. The final ROC package, depicted in Figures 8c and 8d, is a 16×16 Ball-Grid Array (BGA) package with a pitch of 1 mm and a ball diameter of 0.6 mm, resulting in a total footprint area of 289 mm^2.

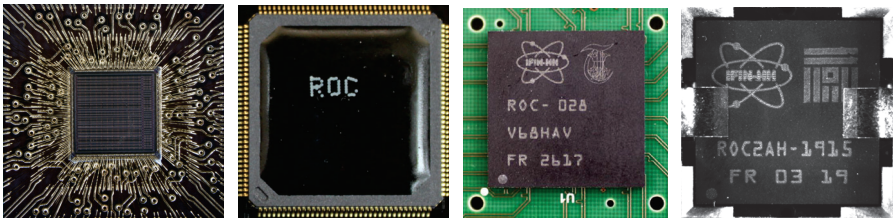

(a) Wire-bonded (b) QFP 144 (c) BGA on PCB (d) BGA in socket

Figure 8: A ROC die wire bonded on a prototype testing PCB (a), the intermediate QFP 144 package (b) and the final 16×16 BGA package on a testing PCB (c) and in an open-top socket (d).

2.3 ROC's steady-state model

From a queuing theory perspective, each SROC contains two servers: the Packet Builder Finite State Machine (FSM) that constructs the L1 events based upon the L0 events from the associated VMM Capture FIFOs and the received L1 triggers (i.e. the processing server) and the Streamer FSM that implements the sending protocol for the completed output events (i.e. the transmission server). The VMM Capture channels associated with an SROC are numbered 1 to m.

Since all the VMM3 ASICs associated with one ROC respond to the same L0 trigger commands and assuming that none of the triggers are dropped, the average input data packet rates are equal: $\lambda_{\text{VMM3 1}} = ... = \lambda_{\text{VMM3 i}} = ... = \lambda_{\text{VMM3 m}} = \lambda_{\text{VMM3}} = \lambda_{\text{L0 trigger}},\ \forall i \in \mathbb{N}, 1 \leqslant i \leqslant m$.

When only one trigger level is deployed within the ATLAS TDAQ system, the average arrival rate of L1 triggers is the same: $\lambda_{\text{L1 trigger}} = \lambda_{\text{L0 trigger}} = \lambda_{\text{VMM3}}$. When two stages are deployed for the hardware trigger, the ROC will be responsible for the second level matching within the NSW readout system. Thus, the condition $\lambda_{\text{L1 trigger}} \leqslant \lambda_{\text{VMM3}} = \lambda_{\text{L0 trigger}}$ covers both HL-LHC upgrade phases. The selection ratio is defined as $sel = \frac{\lambda_{\text{L1 trigger}}}{\lambda_{\text{VMM3}}} = \frac{\lambda_{\text{L1 trigger}}}{\lambda_{\text{L0 trigger}}}$, $0 \leqslant sel \leqslant 1$ and has the requirement that $\lambda_{\text{VMM3}} = \lambda_{\text{L0 trigger}} > 0$.

The condition for SROC FIFO to not fill and consequently throttle the processing server is $\rho_{\text{tx}} \leqslant 1$, where ρ_{tx} is the transmission server's utilization, defined as $\rho_{\text{tx}} = \frac{\lambda_{\text{tx}}}{\mu_{\text{tx}}}$. μ_{tx} is the average transmission rate when the server is busy while λ_{tx} is the average arrival rate of *customers* (i.e. the L1 packets from the processing server). But $\lambda_{\text{tx}} = \mu_{\text{SROC proc}}$, where $\mu_{\text{SROC proc}}$ is the average Packet Builder's service rate of L1 triggers.

Similarly, the TTC and VMM Capture FIFOs will not fill (and consequently data will not be lost) as long as the following condition is true: $\rho_{\text{proc}} \leqslant 1$. The processing server utilization is defined as $\rho_{\text{proc}} = \frac{\lambda_{\text{L1 trigger}}}{\mu_{\text{SROC proc}}}$. The conditions for not filing any ROC FIFOs translate to $\lambda_{\text{L1 trigger}} \leqslant \mu_{\text{SROC proc}} \leqslant \mu_{\text{tx}}$. Data loss does not occur within the SROC Packet FIFO if $\mu_{\text{SROC proc}} > \mu_{\text{tx}}$ but as soon as the FIFO is full $\mu_{\text{SROC proc}}$ will be decreased to μ_{tx}.

The maximum average lossless SROC transfer rate[3] (in pkt/s) is:

$$\mu_{SROC\_max\_tx} = \\ = \frac{kv}{10[k\sum_{i=1}^{m}\binom{m}{i}(1-p)^i p^{m-i}(9+e+i\bar{n}h) \ + kp^m(3+e)o+3-e]} \tag{1}$$

The variables used within the formula are detailed in Table 1. The selection ratio *sel* does not appear in the formula because only the output packets are considered. However, to reach and maintain the maximum SROC transfer rate, *sel* must be high enough so that the SROC Packet FIFO always contains at least one complete L1 packet.

Var.	Description
e	enable state for the use of EOP symbols between the back-to-back L1 packets, $e \in \{0,1\}$.
h	number of bytes of an L1 hit word, $h \in \{3,4\}$.
k	maximum number of back-to-back L1 packets, $1 \leqslant k \leqslant 255$.
m	number of associated VMM Capture channels, $1 \leqslant m \leqslant 8$.
$\bar{n}$	mean of the discrete random variable $n \in \mathbb{N},\ 1 \leqslant n \leqslant 64$ which represents the number of L0 hit words in the non-empty VMM3 packets.
o	enable state for the transmission of L1 null-event packets, $o \in \{0,1\}$.
p	probability that a VMM3 would output an L0 null-event.
v	transmission bit rate for the SROC, $v \in \{80,160,320,640\} \times 10^6$ bps.

Table 1: Description of the variables from the eq. 1, 2 and 3.

The maximum average lossless VMM3 output rate, in pkt/s and the maximum average lossless SROC processing rate, in trigg/s, are:

$$\lambda_{VMM3\_max\_tx} = \frac{1.6 \cdot 10^7}{(1-p)(\bar{n}+1)+p} \tag{2}$$

$$\mu_{SROC\_max\_proc} = \frac{40 \cdot 10^6}{2 + m\{\frac{1}{sel}[(1-p)(\bar{n}+1)+p] + 2(1-p^m) + op^m + 2\}} \tag{3}$$

In Figure 9 the theoretical maximum average lossless packet transmission and processing rates are plotted as functions of $\bar{n}$ considering $p = 20\%$. The maximum average processing rate when $sel = 1\%$ is so low because for each L1 trigger, 100 L0 packets are dequeued from each of the two associated VMM Capture channels on average. Only 1% of the input data is transmitted to the output, thus, a transmission bit rate $v = 80 \cdot 10^6$ bps is sufficient.

[3]The $\binom{n}{k}$ notation represents the combination of n taken as k: $\binom{n}{k} = \frac{n!}{k!(n-k)!}$.

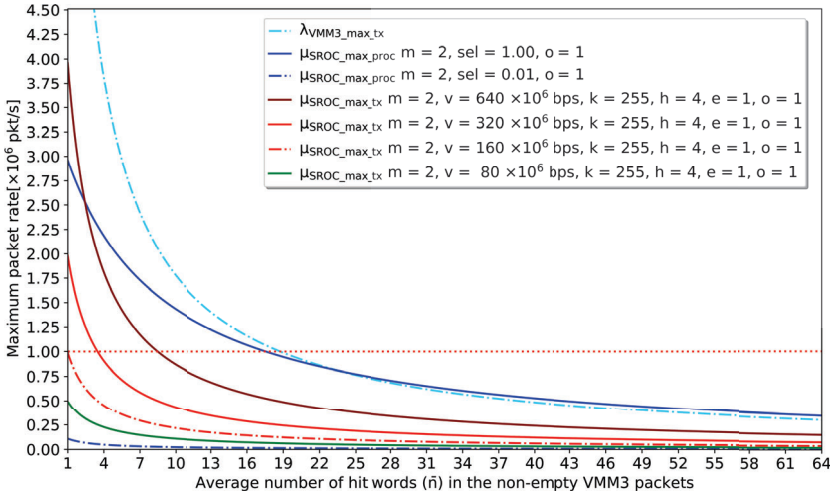

Figure 9: The maximum average lossless rates of packet processing and transmission as functions of $\bar{n}$.

3 ROC Testing

The ROC ASIC required real-world validation of its design, confirmation of its mathematical model, quality control of its mass-produced chip samples, quality assessment of its output clock and TTC signals and a demonstrator. Consequently, two separate FPGA-based (Field Programmable Gate Array) functional test environments were developed and used: a digital one for the packet processing logic and an analog one that assesses the forwarded clock signals and commands and their phase-programmability. Both of them were used for the quality control mass-testing of the fabricated ROC chips. A modified version of the digital test environment was used for the ROC radiation qualification tests which is the subject of Section 4. The emphasis is on the digital test environment because the author contributed exclusively to it.

The quality control of the mass-manufactured chip samples is motivated by the non-null probability of induced defects during fabrication and the process variability. The latter translates into the variability of the electrical characteristics. An example of the former is depicted in Figure 10 which contains two microscope photos of the upper left corner of distinct real-world ROC dies, with (right) and without (left) physical defects. Testing aims at discovering the flaws produced in the manufacturing process and separating the manufactured samples into categories based upon their performance and/or usability, process called binning, the simplest of which is the separation of the *good* and *bad* devices. In addition, functional testing also validates that the device meets its specifications, assuring its operation in the real world [41]. The thoroughness of the functional test determines the quality of the validated samples but implies more time and a lower yield. The ROC does not include scan chains.

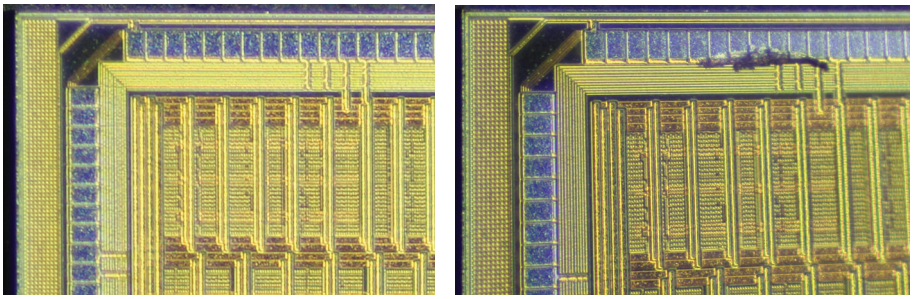

(a) Normal ROC die. (b) Defective ROC die.

Figure 10: Microscope photos (by Sorin Mărtoiu from Horia Hulubei National Institute for Research and Development in Physics and Nuclear Engineering - IFIN-HH, Măgurele) of normal and defective real-world ROC samples.

3.1 The Quality-Control Digital ROC Test

The digital test setup is based on the Xilinx Kintex Ultrascale KCU105 FPGA evaluation board [42] and custom accommodating PCBs for the ROC. It consists of input data generators, output monitor modules and a status and control unit, all implemented on the FPGA, as firmware. The firmware top-view architecture is detailed in Figure 11. The Device Under Test's (DUT's) 8b10b serial input channels are the eight VMM Capture channels. The ROC's TTC Capture module represents the DUT's control serial input channel, while the DUT's 8b10b serial output channels are the four SROCs. L0 packet generators reproduce the VMM3 data streams. Based upon the injected data and the test configuration, a separate generator outputs the TTC stream. The configuration and monitoring of the ROC's digital and analog parts are achieved through two I^2C master modules. The output data analyzers check the encoding and coherency of the ROC output data streams and the parity bits, checksum fields and content of the L1 events. A Xilinx MicroBlaze [43] 32-bit Reduced Instruction Set Computer (RISC) soft-core microprocessor is instantiated in the FPGA design and runs a custom software that controls and monitors the other modules within the FPGA design. Several peripherals are associated with the processor and are used for debugging (through Joint Test Action Group - JTAG), communication with the user (through Universal Asynchronous Receiver-Transmitter - UART), configuration of the custom ROC PCB (through I^2C) and precise timing. The clock signals that pace the FPGA design and the ones supplied to the ROC are generated by a PLL. The locking signal of this PLL is used for the generation of reset signals. A Xilinx Integrated Logic Analyzer (ILA) [44] increases the observability within the FPGA design by sampling relevant signals and transferring the resulted data to a computer through JTAG for analysis. The custom ROC PCB assures its power supply and the interconnection between it and the FPGA through FPGA Mezzanine Card (FMC) connectors [42].

Five versions of ROC testing PCBs were designed, implemented and used.

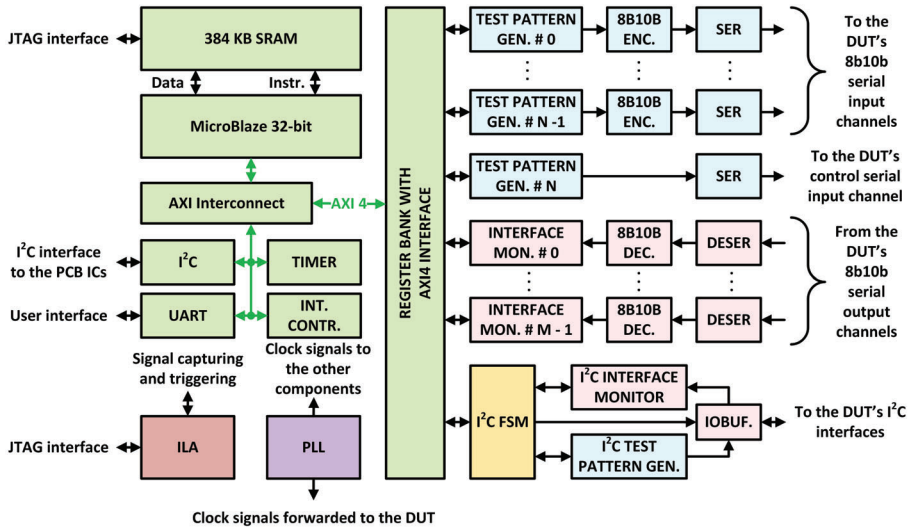

Figure 11: The top-view architecture of the ROC digital functional test setup.

The first four, illustrated in Figure 12, differ mainly by how the ROC package evolved: (i) wire-bonded directly to the PCB, (ii) in QFP 144 package mounted on the PCB and in BGA 16×16 package either (iii) surface mounted or (iv) placed within a closed-top socket. In all four figures the single High Pin Count (HPC) FMC connector is situated left on the other side of the board. The fifth PCB type (detailed in Figure 13) uses an open-top BGA socket for easy and rapid insertion/removal of the tested ASIC, different and supplementary Integrated Circuits (ICs) and the second FMC connector which is Low Pin Count (LPC). The signals mapped to the HPC FMC connector are kept on the same pins for all PCB versions to maintain firmware compatibility and correspond exclusively to the packet processing logic. In Figure 12a the area where the ROC die is placed and wire-bonded is covered with a low-viscosity polymer resin for protection. The same area was depicted uncovered in Figure 8a.

The FPGA design is highly configurable. The average rate of L0 packets is adjustable from 100 kHz to 1.4 MHz (above the requirements) in steps of 100 kHz. For each rate, predefined scenarios with various average percentages of empty packets and average sizes for the non-empty packets can be selected. In total there are 2^7 scenarios, each containing lists of 2^8 pseudo-random values for each descriptor. These lists are looped as the *VMM3 emulator* outputs packets. A distinct scenario with the worst-case theoretical burst in a loop is included and was used during mass testing and validation of the mathematical model. The queuing of the BC selection commands inside the VMM3 is emulated. Constant adjustable rate and size for the L0 packets can also be used. The fields of the hit words are filled with individual counter-generated values, incremented for each new hit word. For each *VMM3 emulator*, the starting and the increment values

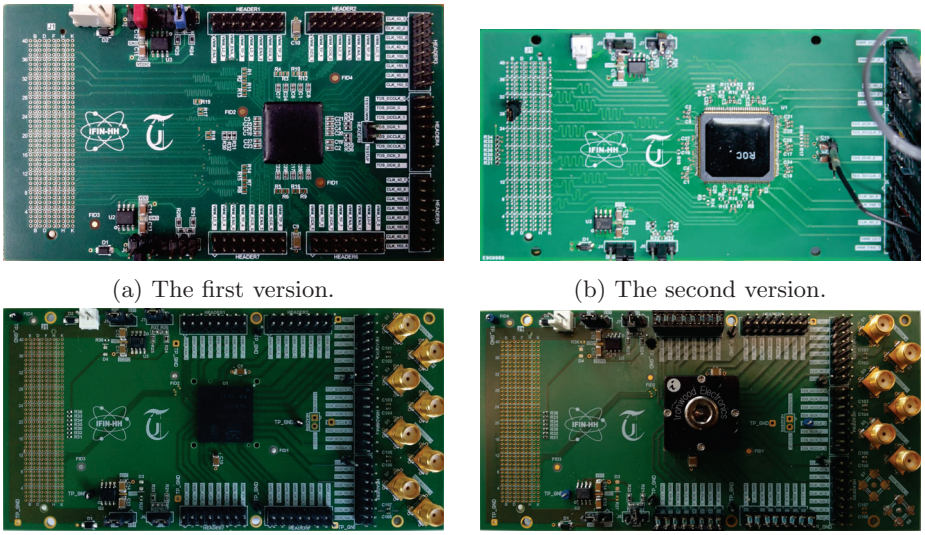

(a) The first version. (b) The second version.

(c) The third version. (d) The fourth version.

Figure 12: The first four versions of the ROC testing PCB.

of these counters are different. There is the option of using constant value hits (i.e. used for the radiation test setup from Section 4). The L1 trigger latency can be adjusted between 10 and 185 μs in 1 μs steps. The selection ratio can be set from 1 to 100% in steps of 1% and the selection pattern can be configured.

The ROC context is asynchronous because its components (e.g. the VMM3s associated with one ROC) have distinct clock domains and delay lines are used (e.g. within ROC and GBTx). Consequently the functional digital test setup is also asynchronous. The chip is mostly synchronous. It contains asynchronous elements: its FIFOs, the TTC Capture module, delay lines for the forwarded clock signals and TTC commands. It does not contain built-in automatic phase-detection circuits for the incoming data so it requires a more complex functional digital test setup and a more complicated set-up procedure once it is installed on the NSW Front-End Boards (FEBs).

The calibration of the phases for the high-speed serial data streams coming out of and into the FPGA, relative to the capturing clock signals within the ROC and the FPGA is achieved using configurable built-in FPGA delay lines. In a first approach, an automatic synchronization mechanism has been developed as a software component, running on the instantiated MicroBlaze microprocessor from the FPGA digital functional test setup, that sweeps the delay lines and validates the responses from the ASIC and FPGA. The middle of the largest interval of consecutive valid delay steps is the chosen delay value. At each step, the injected valid data is varied as much as possible to decrease the Mean Time Between Failures (MTBF) as defined in [45]. For the 8b10b channels, the status flags of the receiving end are checked for decoding, parity and coherency errors.

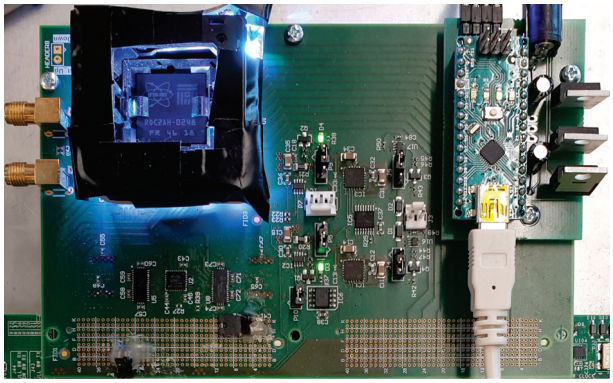

Figure 13: The final PCB version with an open-top BGA socket.

If no error occurred the delay is valid. The last channel that is calibrated is the TTC stream. The expected behavior is checked considering all the other data channels already calibrated in this case.

Because the initial approach is inefficient, a double symmetrical binary search inspired by the classical binary search algorithm was implemented. It searches the extremities of the valid intervals. The proposed approach is depicted in Figure 14 where the valid delay steps are represented with green and the invalid ones with red. The delay line extremities are d_{min} and d_{max}. The algorithm first samples the middle of the delay interval ($\Delta d/2$), where $\Delta d = d_{max} - d_{min}$, and then moves in both directions with $\Delta d/4$. Then, from each of the two resulting positions, it moves with $\Delta d/8$ in one or both directions based on the previous results and so on. Until the desired resolution is reached, the step size is halved at each iteration. The complexity of the proposed algorithm is $O(logN)$ compared to $O(N)$ for the classical sweep.

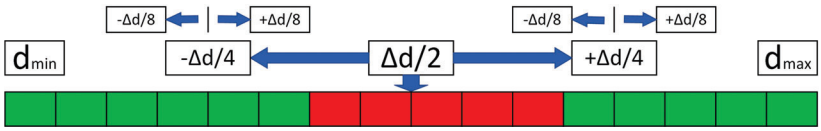

Figure 14: The implemented and used double binary search calibration.

3.2 Experimental testing results

The ROC's packet processing logic Register Transfer Level (RTL) code and netlists and the digital functional test setup RTL were validated in simulations in which they were interfaced. The implemented FPGA design met the timing, placing and routing constraints. The first real-world ROC chips were extensively tested in different scenarios for periods ranging from a couple of minutes up to 48 hours and ILA-captured data were analyzed. The ROC passed preliminary

integration tests with the other ASICs from the NSW context [46] [47]. The faster calibration method was successful (i.e. 70% faster in the worst case). The functional digital test setup was used for chip mass-testing. The performance of the ROC as theorized in Section 2.3 in terms of maximum data rates without loss, was assessed. The worst-case triggers and input packets burst loop was applied. The theoretical threshold limit is correct.

In Figure 15 the normalized histograms of the size of the largest valid delay interval (left) and the found optimal delay value (i.e. middle of the largest valid interval - right) for two VMM Capture channels, at nominal (1.2 V, top) and sub-nominal (1.1 V, bottom) voltages are depicted. The difference between the average sizes at nominal voltage is 137.55 delay steps (i.e. 572 ps) despite the identical time constraints and RTL code used during design. The spreading of the histograms at 1.1 V is caused mainly by the increased variability of the clock phase shift introduced by the ePLLs at lower than nominal voltages.

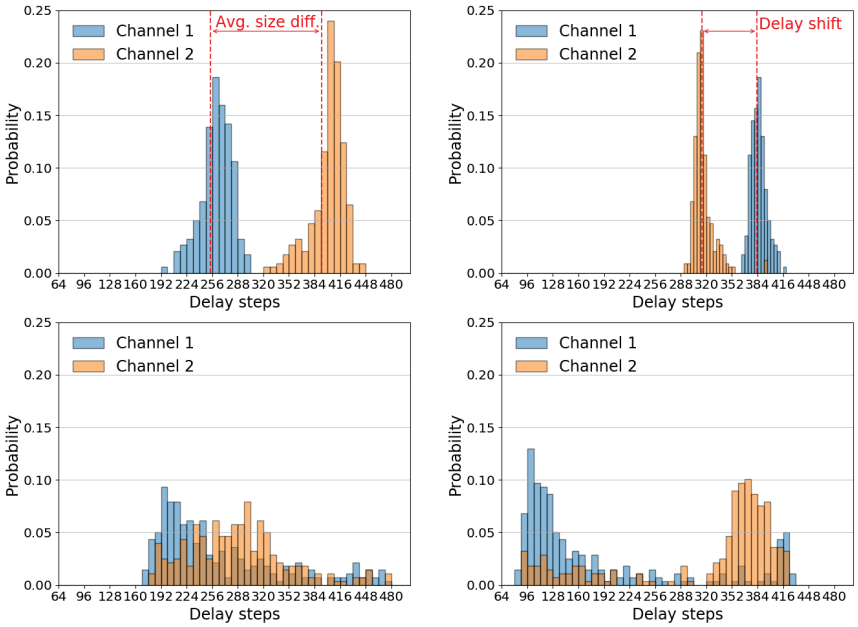

Figure 15: For two DUT input channels: histograms of the valid delay interval size (left) and the found optimal delay value (right) at nominal (top) and sub-nominal (bottom) voltages [48].

The presented synchronization methods proved to be useful for evaluating the size of the crossing region of the eye diagram (i.e. jitter estimation) for the DUT input channels. In Figure 16 the oscilloscope-measured eye diagrams on a sample chip for an input (top) and an output DUT channel (bottom) both at nominal (right) and at sub-nominal (left) voltages are shown. The estimated jitter values resulted from the waveforms are 670, 834, 1477 and 1664 ps from top to bottom and left to right showing degradation when lowering the supplied

voltage and passing through the DUT, as expected. For the same chip and channels, the sizes of the invalid delay interval are 665.6, 844.5, 561.1 and 773.76 ps in the same order. In this case, the jitter estimates for the output channel are lower than those for the input one, contrary to what was measured in the diagrams, due to the shorter observation duration of 20 ms/step compared to approx.10 s. For comparable results, the same observation time must be used.

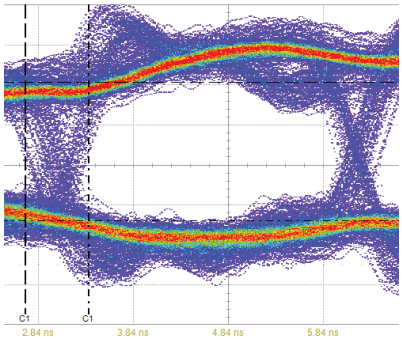

(a) Jitter 670 ps; FPGA 665.6 ps

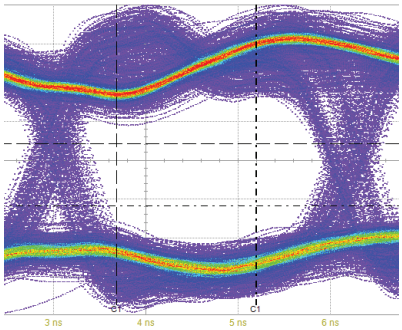

(b) Jitter 834 ps; FPGA 844.5 ps

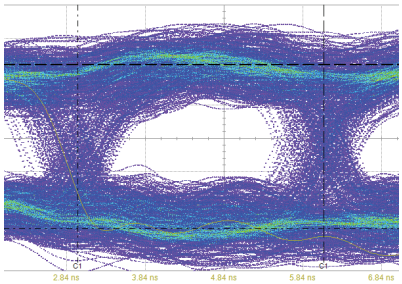

(c) Jitter 1477 ps; FPGA 561.1 ps

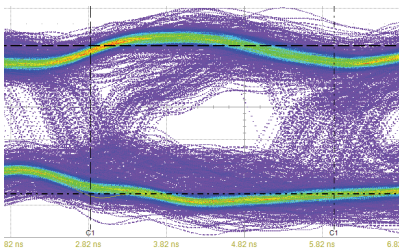

(d) Jitter 1664 ps; FPGA 773.76 ps

Figure 16: Measured eye diagrams for a DUT input channel at nominal (top-left) and sub-nominal (top-right) voltages and for a DUT output channel at nominal (bottom-left) and sub-nominal (bottom-right) voltages.

At Transilvania University of Brașov, 2,677 BGA ROC samples were functionally tested in both the digital and analog test setups. The following statements refer solely to the digital testing. Two supply voltages were used: (i) the nominal 1.2 V and (ii) the sub-nominal 1.1 V. Ten tests were designed to cover all the functional features. Only the chips that passed all the tests at both voltages were marked as *good*. A test is failed if any output bit is incorrect. 2,343 (87.52 %) chips passed the testing at 1.2 V and from them, 1,819 (67.95 % from the total) were declared *good*. An analysis of the causes of failure is included in the thesis.

4 Immunity to radiation-induced faults

Nuclear radiation disrupts the electronic devices and can even permanently damage them through two phenomena: *ionization*[4] and *displacement damage*[5]. The radiation-induced effects are classified into two types: cumulative and Single Event Effects (SEEs). The cumulative or total dose effects are proportional to the integrated flux of particles that were incident to the device. Eventually, they lead to the complete failure of the device. SEEs are immediate results of ionization. Two typical types of *soft* (i.e. temporary affecting a signal) SEEs are Single Event Transient (SET) and Single Event Upset (SEU). A SET is an induced rapid toggling (i.e. glitch) of the voltage level within a circuit net [51]. The effect on the system functionality depends on the net's function, e.g. the most disruptive SET happens on a clock tree branch. An SEU is an induced change of state within a sequential element [50] caused by charge deposition. The value of the stored bit is flipped and persists until a new value is written, the reset is enabled or the device is powered off.

The on-site ICs used for space applications, nuclear and particle physics experiments are exposed to larger quantities of nuclear radiation compared to the off-site or general purpose ones. In addition to the possibility of permanent damage, during operation the *soft* SEEs can induce errors and bring the circuits into states not considered in their design stage and thus cause them to not function properly. CERN defines two categories of electronic systems exposed to radiation: radiation-tolerant and radiation-hardened [50]. The radiation-tolerant devices are designed to operate within irradiated environments despite being vulnerable to radiation. They implement design mitigation techniques for SEEs. The radiation-hardened devices are immune to all radiation-induced faults in the specified limits of their operating environment. They employ both constructive and design mitigation measures.

Triple Modular Redundancy (TMR) represents a classical hardware redundancy technique [52] for mitigating SEEs. It consists of replicating the system three times, feeding the same input to all systems and determine the correct output by majority voting. If one system fails, the other two can determine the correct output. TMR was implemented for all ROC's Flip-Flops (FFs). All the majority voting modules from the digital part also signal if the value of one of the instances does not match with the other two. A large OR gate concentrates all these signals into one 1-bit wide signal that is used to increment an internal 8-bit SEU counter and is also fed to an output pad.

The ROC design in between the radiation-tolerant and radiation-hardened categories as it employs both design-based SEEs mitigation mechanisms and radiation-hardened technologies (i.e. the e-link pads and ePLL macroblocks) mixed with commercial off-the-shelf components (e.g. the SRAMs and the logic cells). Its SRAM FIFO buffers are the most susceptible to radiation.

[4]The process of acquiring positive or negative charge by losing or gaining electrons, respectively, through collisions of sufficiently energetic particles with the ionized material [49].

[5]The process of inducing defects in the lattice of semiconductor atoms [50].

4.1 Test setup

The Tandem accelerator from the National Centre for Scientific Research (NCSR) Demokritos, Athens, Greece, produces beams of neutrons with energies between 0.1 and 25.7 MeV [53] using various nuclear reactions. A total neutron fluence comparable to the maximum NSW annual one can be achieved in several tens of hours. Even if the energy spectrum of the produced beams does not completely match the spectrum within the NSW and the parasitic nuclear interactions are not very well studied or documented, the Demokritos facility was considered to be appropriate by the NSW Group for the radiation qualification process of the ROC and the other context ASICs.

The functional digital test setup used for design validation and chip mass-testing, presented in the previous section, was adapted for the irradiation tests. It was ideal since it emulates the ROC's context and covers all possible states and functional features for the packet processing logic being highly configurable. In addition, it can detect any incorrect bit of any output packet while the ROC functions in real-time at its nominal frequencies.

The mass-production testing PCB with the open-top socket (depicted in Figure 13) hosting the BGA packaged ASIC (as illustrated in Figure 8d) was positioned so that the silicon die was directly exposed to the neutron beam as close as possible to the source. The interaction area of the test setup is detailed in Figure 17. The testing board was interfaced with the FPGA evaluation board through a 1.8 m long FMC HPC cable. The FPGA evaluation board was positioned outside the beam trajectory, being shielded by lead bricks filled with paraffin and was interfaced through JTAG to Universal Serial Bus (USB) and UART to USB with a host computer with remote internet access. The UART was the main user interface to the test setup and its data was logged into text files. Time windows of relevant signals from the FPGA design, triggered by the encountered output packet error flags, were captured by a Xilinx ILA and transferred to the host computer via USB-JTAG.

The used configuration of the test setup increased the occupancy levels of all ROC FIFOs as much as possible without inducing data loss. Thus, the probability that the ROC's behavior and output data will be affected was higher.

4.2 Irradiation tests results

Two validated ROC chips were subjected to neutron beam runs with nominal energies of 20, 22 and 24 MeV. The neutron flux was estimated for each run by the accelerator operators, as a function of the distance and the interposed materials between the source and the ROC die. The stated uncertainty is 10%. The accumulated fluence is equivalent to 8.08 ± 0.81 months of operation in the most irradiated part of NSW when $L_{\mathrm{LHC}} = 10^{34}\ \mathrm{cm}^{-2} \cdot \mathrm{s}^{-1}$. A total of 69 FF SEUs were recorded from the 18458 FFs of the ROC digital part, the majority of which happened during the 24 MeV runs. The recorded log files and the captured ILA waveform windows were analyzed. No misalignment, encoding, decoding or protocol syntax errors were detected in the ROC's output

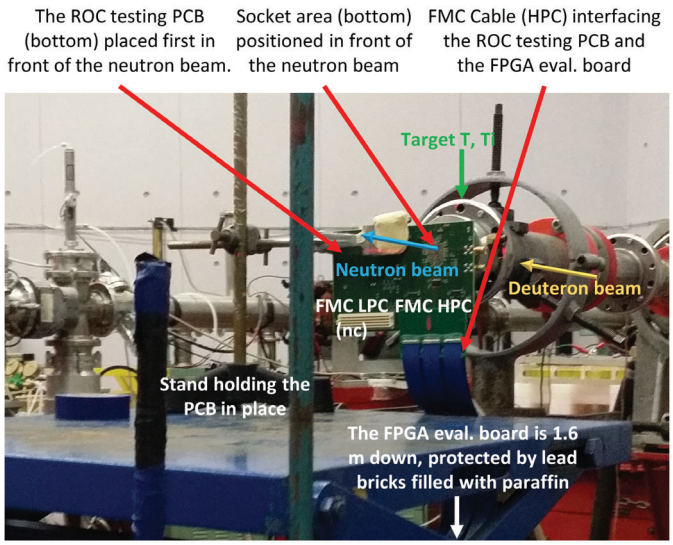

Figure 17: The interaction area within the ROC radiation test setup.

data streams proving that its logic was not driven into illegal states and recovered gracefully from the FF SEUs. In Table 2, the obtained FF cross-section averages over the runs with the same beam energy are listed, including their confidence intervals computed after the methodology presented in [54] considering a confidence level of 95%.

Nominal energy [MeV]	**FF SEUs**	$\sigma_{\mathbf{FF}}$ **[$\times 10^{-14}$ cm^2·bit^{-1}]**
20	22	$1.37^{+0.75}_{-0.58}$
22	10	$3.46^{+2.98}_{-1.93}$
24	37	$6.28^{+2.68}_{-2.23}$

Table 2: The recorded FF SEUs per beam energy, the confidence intervals and the associated cross-sections computed after the methodology from [54], considering a fluence uncertainty of 10% and a 95% confidence level.

Content errors in the output packets caused by SRAM SEUs occurred at a considerable rate, as depicted in Figure 18. Every 10 s (i.e. one iteration) the newly encountered content errors were logged. They are classified by the affected packet field. However, a single SRAM SEU can produce multiple errors.

In Figure 19 the normalized histograms of the interarrival times for the FF SEUs and the first three most frequent types of output packet errors are depicted for a 24 MeV run. The ideal exponential probability density function with the same mean arrival rate (i.e. λ) is plotted for reference in each case. To analyze the validity of the memoryless hypothesis, Chi-square and Kolmogorov-

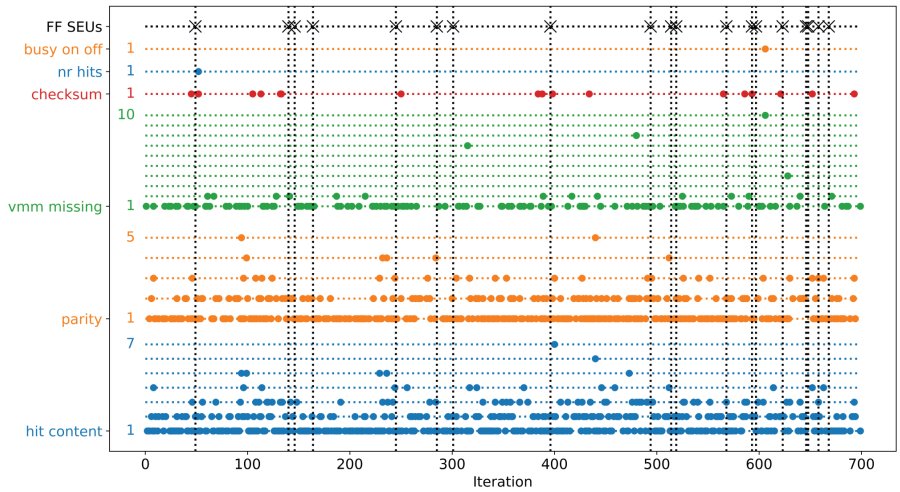

Figure 18: A time window of 700 consecutive iterations (i.e. 7000 s) from a 24 MeV run depicting with markers the moments of detection for 19 FF SEUs and various output packet errors separated by type.

Smirnov goodness of fit tests were performed in all cases. Their verdicts are listed in each histogram. This characteristic is confirmed only for the FF SEUs and the *VMM missing* errors. The used time resolution is not satisfactory for the *hit content* and *parity* errors. Since a single SRAM bit flip can cause the detection of multiple packet errors, one cannot assume exponential distributions for the interarrival times of the output packet errors even if the SRAM SEUs that caused them do follow such distributions.

In Figure 20 the cross-section confidence intervals for the FF SEUs and the four most frequent types of packet errors are depicted categorized by the energy of the neutron beam that produced them. For each packet content error type only the occupied SRAM latches that could cause it if altered are considered. The values assume a 10% fluence uncertainty and use a 95% confidence level. The probability for the SRAM latches to suffer SEUs is approximately equal between the error types. The rate of SEUs is significantly increased at 24 MeV compared to 20 MeV, translating into the relatively substantial cross-section values. The SRAM SEUs cross-sections were estimated at $5.82^{+1.18}_{-1.18} \times 10^{-14}$; $9.16^{+2.01}_{-1.99} \times 10^{-14}$ and $37.18^{+7.39}_{-7.39} \times 10^{-14}$ cm^2/bit at 20, 22 and 24 MeV energies, respectively. These values are more pessimistic than in reality because they are based upon the total output packet error counts, but a single bit flip may cause multiple errors to be reported.

SRAM bit flips can cause temporary or permanent (i.e. until reset) losses of synchronization between the VMM3 data and the L1 triggers. Two such scenarios were simulated and are presented in the full thesis.

Considering the 24 MeV cross-sections, the estimation is that each ROC will experience $150 < 232 < 331$ FF and $17593 < 21957 < 26323$ SRAM SEUs per

year in NSW at $L_{\text{LHC}} = 10^{34}$ cm$^{-2} \cdot$s^{-1}. These are equivalent to $1.39 < 2.15 < 3.06$ FF SEUs/min and $163.07 < 203.52 < 243.99$ SRAM SEUs/min, in all 4875 ROC ASICs which will be installed and function concurrently for NSW.

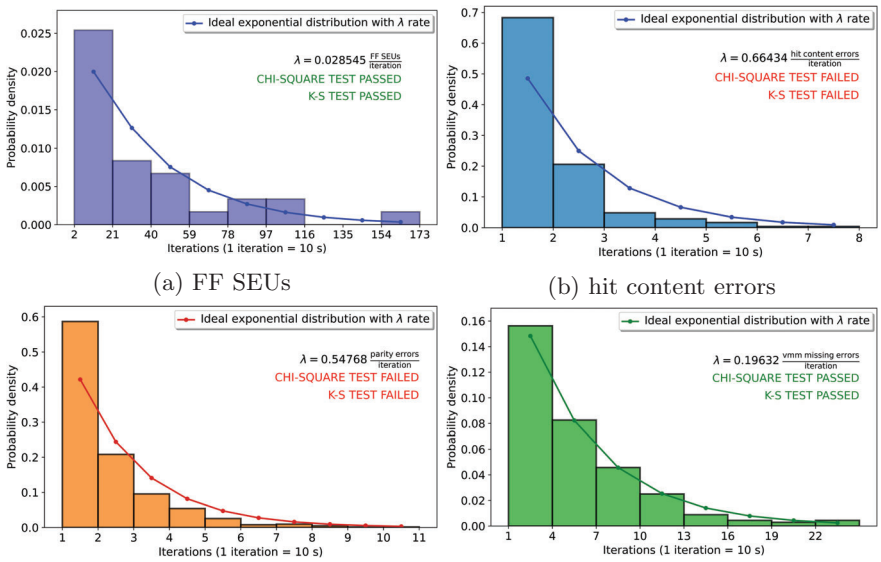

(a) FF SEUs (b) hit content errors

(c) parity errors (d) *VMM missing* errors

Figure 19: Normalized histograms and theoretical probability density functions of the interarrival times in the case of FF SEUs and the three most frequent types of output packet errors observed during a 24 MeV run.

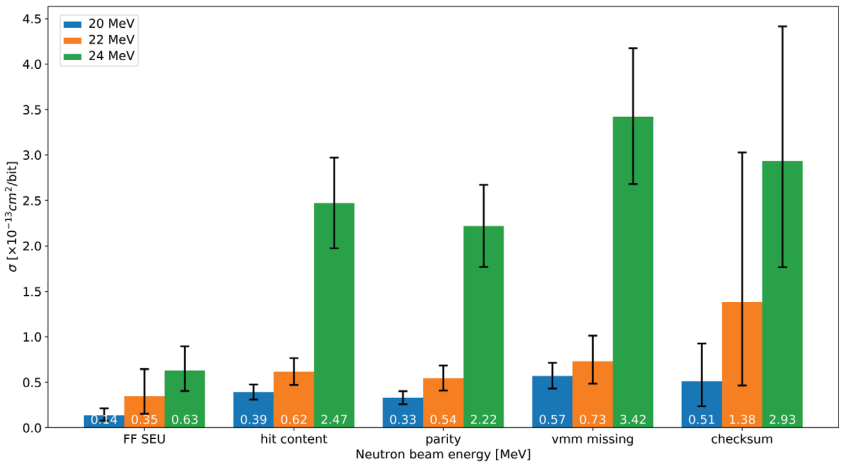

Figure 20: Cross-sections of the FF SEUs and the four most frequent types of packet errors with a fluence uncertainty of 10% and a 95% confidence level.

5 An application

During the development of the FPGA design for the validation of the ROC digital part, its digital functional mass-testing, performance assessment and irradiation qualification, one of the most useful deployed tools was the Xilinx ILA [44]. It increased considerably the observability within the FPGA design and therefore its debuggability and allowed the extensive analysis of the radiation-induced output packet errors. However, the tool has some constraints and limitations. Thus, the idea of a new FPGA ILA design that mitigates these limitations with changes to both the current method and apparatus was born. The improved performance and efficiency come at the cost of increased complexity.

The Logic Analyzer (LA) is an electronic instrument used to digitize, sample, store, display, decode, debug and analyze windows of multiple concurrent digital signals of interest from digital circuits and systems. In contrast to oscilloscopes, LAs are dedicated to exclusively probe digital signals at the data level on many more channels. Even if real-world digital signals are in fact analog, the probing LA channels transform them into digital representations. Oscilloscopes are used to analyze characteristics of signals considered digital (e.g. rise and fall times, fill factor, jitter and skew, etc.). LAs translate the same signals into logical 1's and 0's to determine the data and protocol correctness. An integrated LA (ILA) is an accessory that enhances the functionality of an electronic instrument or device, e.g. a Digital Storage Oscilloscope that incorporates an LA becomes a Mixed Signal Oscilloscope. An FPGA ILA is partially implemented into the FPGA's configurable logic. It acquires sample windows of selected signals from the implemented design and transfers them to a host computer where they can be interpreted, checked, stored and visualized as timing diagrams.

The identified limitations of the available FPGA ILAs (i.e. Xilinx ILA and Altera SignalTap II [55]) are the following: (i) the sampling depth (i.e. the acquisition window size) is constrained by the size of the available unused SRAM on the FPGA device; (ii) the sampling depth is constrained to be strictly a power of two; (iii) the available SRAM cannot be fully used; (iv) the SRAM cannot be combined with other sequential resources available on the device like dedicated FIFOs, distributed memory (from Look-Up Tables - LUTs) and FFs/registers; (v) the queuing algorithm and the relatively slow speed of the interface to the host computer cause increased and in some conditions unnecessary dead-time translating into missed triggers and signals of interest; (vi) re-implementation of the design is necessary when different signals must be probed which means lost time proportional with the complexity of the design and (vii) no available data about the moments of occurrence or the counts of triggers.

The proposed solutions are: (i) use a higher speed interface (e.g. 10 Gbps Ethernet) for sample data transfer; (ii) use a configurable scheduling algorithm that shares the available bandwidth with the user design and transmits data while being sampled instead of waiting for the buffer to fill; (iii) remove the constraint that the depth of the buffer must be strictly a power of two; (iv) extend the queuing depth by combining the available sequential resources of

the FPGA; (v) remove the need for re-implementation when different design signals must be sampled by using the partial reconfiguration feature; (vi) add timestamps and trigger counter values in each window of sampled data; (vii) maintain the rest of the current functional features of the available FPGA ILAs.

In Figure 21 the proposed top-level architecture of an FPGA ILA which uses an Institute of Electrical and Electronics Engineers (IEEE) 802.3 10 Gbps Ethernet high-speed interface is depicted. The user logic (Device Under Verification - DUV and Verification Environment - VE) has n different clock domains and for each one an FPGA ILA probing system is deployed. The Ethernet interface is shared in both directions with the DUV and the VE. Three variants for the implementation of the Media Access Controller (MAC) and Physical (PHY) Open Systems Interconnection (OSI) components are depicted: (i) both implemented inside the FPGA; (ii) the MAC sub-layer implemented in the FPGA logic and communicating with a PHY IC from the PCB hosting the FPGA and (iii) both implemented outside the FPGA as ICs on the accommodating PCB. The clock signal of the FPGA ILA processing system is shown as having the two possible frequencies of 156.25 and 312.5 MHz because these are the frequencies associated with most standard interfaces of communication between 10 Gbps Ethernet MAC and/or PHY ICs, e.g. Advanced eXtensible Interface 4 (AXI4) Stream [56] or 10 Gigabit Media-Independent Interface (XGMII) [57]. The proposed scheduling algorithm for the usage of the 10 Gbps Ethernet interface is Weighted Round Robin (WRR) implemented on a per-quanta basis. The size of the quanta is configurable. The weights allocated to the clients are determined by the throughput ratio of the FPGA ILA probing systems and the DUV and VE logic. The FPGA ILA's custom on-top of Ethernet protocol proposes an Ethernet payload format that allows the multiplexed segmented transmission of sampled data, data from VE, data from DUV, status notifications and the reception of configuration commands and VE and DUV data.

The main advantage of the proposed design is the ability to sample larger windows of signals than current JTAG-to-USB-based FPGA ILAs. The dead time after a trigger fires is reduced and in some conditions even eliminated. The design adds flexibility by using and combining built-in FIFOs, distributed memory, registers and Block BRAM (BRAM) rather than only BRAM for data buffering. It implies a more efficient use of the available FPGA resources (e.g. FIFO depths not constrained to be powers of two). The design is suitable both for close to real-time hardware-in-the-loop verification and reaching long DUV/DUT operational intervals faster than in simulations. By using partial FPGA reconfiguration, no re-implementation is required when other design signals must be sampled. Also, the designed ILA can transmit the data while it is being sampled instead of waiting for the buffer to fill, stopping the sampling and triggering and then flushing the buffer content to the host computer. In addition, the proposed ILA can share the used high-speed interface with the user design. Because most interfaces are full-duplex (e.g. the 10 Gbps Ethernet interface), the FPGA's receive link can be used for other communications with the host computer, e.g. for driving internal FPGA signals in close to real-time or receiving data to be processed by the DUV.

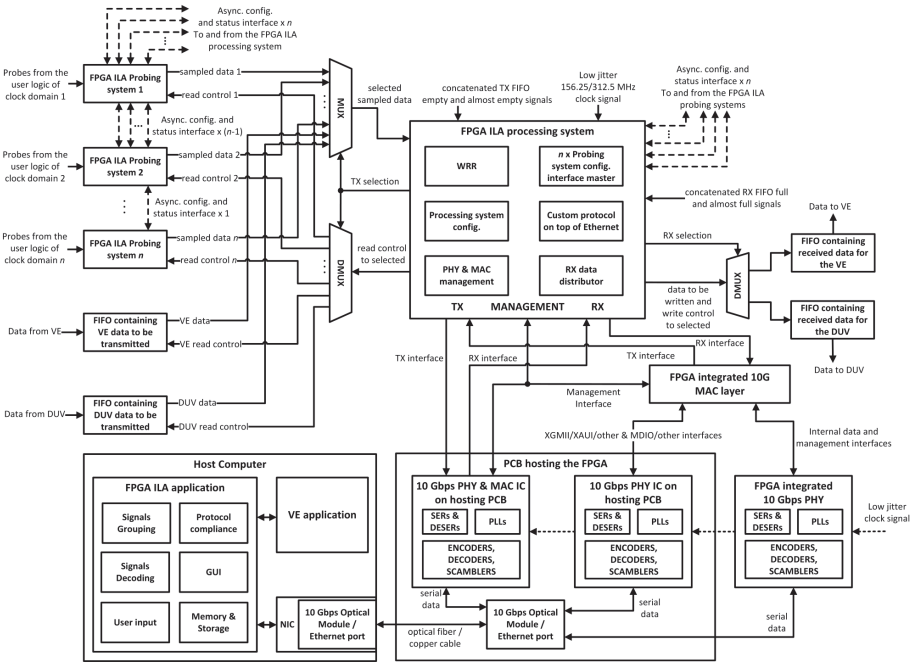

Figure 21: The architecture of the proposed FPGA Logic Analyzer.

The main disadvantage is that the proposed design requires an FPGA device capable of concurrently accommodating the user design, the ILA logic and the high-speed interface. Additional resources are required: low-jitter clock sources for pacing the high rate transfers and ICs implementing Data Link (DL) and/or PHY OSI layers. If the FPGA internal logic also implements the DL and PHY layers of the interface, the device requires high-speed transceivers, encoders, decoders, scramblers and PLLs. A relatively large amount of the available IOs must be dedicated when using the XGMII in the 10 Gbps Ethernet case. Other disadvantages include tighter constraints upon the user design and higher power consumption caused by the increased usage of resources and the higher operating frequencies. The proposed design has increased complexity and requirements. Last but not least, a host computer capable of accommodating the high-speed interface and handling the transferred data is required.

The proposed FPGA ILA would have been a very useful tool during the ROC's irradiation testing since it enhances the observability of the test setup. The results and analysis from Section 4 would have been consequently improved. Unfortunately, Xilinx and other companies directly interested in this subject protect their inventions with patents but are not bound to implement and provide them to the user of their tools and/or devices. Nevertheless, the proposed high-speed time-multiplexed FPGA ILA resulted as an application/consequence of the ROC's mass-testing and irradiation qualification.

6 Conclusions

All the subjects touched by this thesis relate to the high luminosity general-purpose ATLAS Experiment from LHC. The LHC, operated by CERN, Geneva, Switzerland, is the largest and highest-energy synchrotron in the world. The experiments at CERN (including ATLAS) are used for fundamental research in particle physics. To detect and characterize all the products resulted from the collisions of the particle beams within ATLAS, advanced detector systems and associated TDAQ systems were designed and successfully used, leading to the confirmation of the theorized Higgs boson in 2012.

The current upgrade process implements the first phase of the HL-LHC project which implies one order of magnitude increase in size for the collected data (3000 fb^{-1} integrated luminosity at 14 TeV in 10 years). To cope with the increased nuclear radiation and the necessity for more precise discrimination and characterization of the collision's effects taking place at higher rates, the ATLAS TDAQ system, the associated software tools and part of the detector system are being upgraded. Part of the upgraded NSW readout path, the ROC is an on-detector radiation-tolerant ASIC acting as a concentrator, buffer, filter and packet data processor. *The presented work represents the author's contribution to its design, verification, implementation, quality assurance and control (including radiation qualification), support for integration, documentation, mathematical and algorithmic descriptions.* The author's work was performed while being a research assistant in the "Experimentul ATLAS de la LHC" (i.e. The ATLAS Experiment from LHC) national research project as member of the team from Transilvania University of Brașov.

The performed work covers: digital electronics design, digital circuits simulation, digital circuit verification, programming, data structures and algorithms, data communications, embedded systems, EDA tools, operating systems, scripting, statistics, signal processing, packet processors, network-on-chip, measuring instruments and scientific and technical document preparation. The work implied extensive real-world debugging, equipment use and measurements. In addition, the author attended special courses presented by prestigious scientific research centers[6] [7]. Other covered topics include but are not limited to: modular redundancy, schematic PCB design, digital circuit synthesis, digital circuit constraining, static timing analysis, design for testability, logic equivalence checking for digital circuits, gate-level simulations of digital circuits, ASIC floor-planning, ASIC placement and routing, digital circuit testing, goodness-of-fit assessment, international scientific collaboration, research project management and dissemination of scientific results.

The ROC design was successfully taped out for fabrication in August 2016. Custom accommodating PCBs and functional FPGA-based test setups were

[6] "Comprehensive Digital IC Implementation & Sign-Off" held in November 2015 by the Microelectronics Support Centre Science and Technology Facilities Council from Rutherford Appleton Laboratory (RAL), Oxford, Didcot OX11 0QX, United Kingdom.

[7] "Verification for Digital Designs" held between 4 - 6 December 2019 by the Microelectronics Support Centre Science and Technology Facilities Council from RAL.

elaborated. The first ROC ASIC was powered up in the digital test setup in March 2017. The chip was active but error flags were raised. Since then, the test environment was massively improved. All the ROC's functional features were covered. The main causes for the initial errors were setup and hold violations inducing meta-stability in the capturing FFs from both the ASIC and the FPGA design. The high-speed serial lines between the ASIC and the FPGA required calibration and the initial issues were solved.

The ROC digital design was also tested and validated by a team from University of Michigan, part of the NSW collaboration, using the same design for the accommodating PCBs and the same type of FPGA evaluation boards but a different, separate, in-house FPGA design based upon the chip's specifications. The ROC's digital design was validated by both sides but some bugs that were not design-breaking were discovered in the digital part. The Michigan team also thoroughly analyzed the ROC's analog part. In parallel, at Horia Hulubei National Institute for Research and Development in Physics and Nuclear Engineering (IFIN-HH), Măgurele, Romania, the ROC's analog part was also tested. An uncertain phase after power-up for each output clock signal relative to the reference BC clock signal was discovered [47]. Thus the ROC's ePLLs macros were redesigned at IFIN-HH while the digital part remained the same.

The redesign solved the clock signals phase uncertainty and the new chip was named both ROC1A and ROC2 [58]. Its mass fabrication was achieved on an Multi-Project Wafer (MPW) alongside other NSW ASICs, as shown in Figure 22. Once it was confirmed that the analog redesign solved the phase uncertainty, the AF region was also occupied by the ROC2 design. It was decided by the NSW Group that the quality control of the fabricated samples will be carried out using the digital test setup presented in Section 3 and a separate automated analog test setup developed at IFIN-HH, Măgurele.

In the summer of 2018, the ROC was successfully tested while being hit by ultrafast controlled neutron beams produced by the TANDEM accelerator from NCSR Demokritos, Athens, Greece. The chip required a basic radiation qualification process since it was part of the on-detector ATLAS electronics. The ROC chip's operation was not directly disturbed by the incident radiation due to the implemented TMR but the data corruption in its SRAMs occurred at a considerable rate. Estimations as accurate as possible were made for the HL-LHC. The idea of a high-speed FPGA ILA with a different operation than the provided one was born from the lessons learned during these experiments.

At Transilvania University of Brașov, a total of 2677 ROC1A/ROC2 chips were tested using the two functional test FPGA designs (i.e. analog and digital) replicated in up to three test benches. At least 4875 ROC1A/ROC2 ASICs will function simultaneously within the NSW TDAQ system. The initial yields were considered unsatisfactory, as the majority of chips were failing the digital test. The calibration procedures for the high-speed transmission lines between the ASIC and the FPGA were further improved. This also motivated the development of the faster calibration algorithm from Section 3. The final digital testing yields (87.5 % at nominal voltage and 68 % at sub-nominal supply voltage) were considered satisfactory [58].

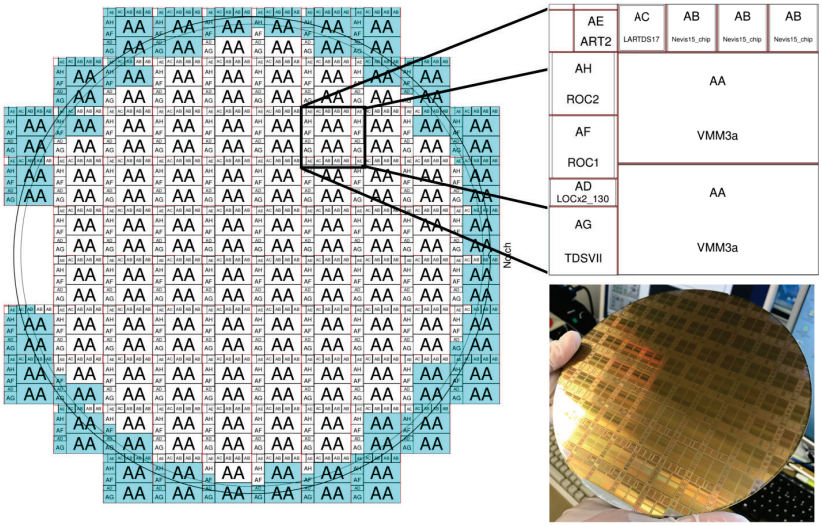

Figure 22: MPW containing NSW ASICs including the two ROC versions [59].

The author periodically reported the progress of the work and the testing results in NSW coordination meetings. As an achievement and recognition for his contribution to the Experiment, he completed the one-year ATLAS author[8] qualification procedure in November 2018.

The ROC1A/ROC2 design passed all the NSW internal reviews and is included in its readout system as shown in Figure 23. Published preliminary integration results are beginning to appear [46].

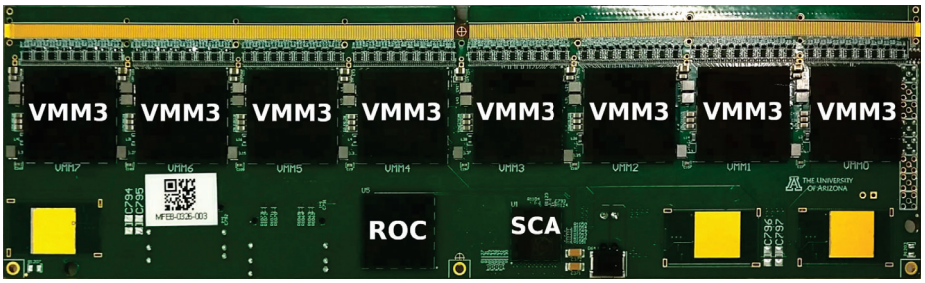

Figure 23: Front-end NSW board containing one ROC chip. Original photos by Anne Fortman, Harvard University ATLAS group.

As questions on the ROC's performance in specific scenarios often arise, the author proposed and validated a mathematical model for the maximum data rates that do not induce loss (presented in Section 2).

[8]Person who officially has made significant contributions to the ATLAS experiment [60]. All ATLAS General Publications are signed by all active ATLAS authors (e.g. [61]).

The scientific publications related to the presented work are:

1. **Ș. Popa**, M. Ivanovici, R.-M. Coliban, "Time-multiplexed 10Gbps Ethernet-based Integrated Logic Analyzer for FPGAs", *International Symposium on Electronics and Telecommunications, ISETC 2020*, Timișoara, 5-6 November 2020, https://doi.org/10.1109/ISETC50328.2020.9301115 [62]

2. **Ș. Popa**, S. Mărtoiu, M. Ivanovici, "Study of the ATLAS new small wheel read-out controller ASIC in a neutron irradiation environment", JOURNAL OF INSTRUMENTATION, Volume 15 P10023, October 2020, https://doi.org/10.1088/1748-0221/15/10/P10023 [39]

3. **Ș. Popa**, S. Mărtoiu, M. Ivanovici, "The quality-control test of the digital logic for the ATLAS new small wheel read-out controller ASIC", JOURNAL OF INSTRUMENTATION, Volume 15 P04023, April 2020, https://doi.org/10.1088/1748-0221/15/04/P04023 [63]

4. **Ș. Popa**, M. Luchian, M. Ivanovici, "Clock and data signals synchronization for an FPGA-based ASIC testing setup", *14th International Symposium on Signals Circuits and Systems, ISSCS 2019*, Iași, România, July 2019, https://doi.org/10.1109/ISSCS.2019.8801780 [48]

5. **Ș. Popa**, S. Mărtoiu, M. Luchian, R.-M. Coliban, M. Ivanovici, "The Quality-Assurance Test of the ATLAS New Small Wheel Read-Out Controller ASIC", *Topical Workshop on Electronics for Particle Physics, TWEPP 2018*, Antwerp, Belgium, 17-21 September 2018, https://doi.org/10.22323/1.343.0081 [64]

6. R.-M. Coliban, **Ș. Popa**, T. Tulbure, D. Nicula, M. Ivanovici, S. Mărtoiu, L. Levinson, J. Vermeulen, "The Read Out Controller for the ATLAS New Small Wheel", JOURNAL OF INSTRUMENTATION, Volume 11 C02069, February 2016, https://doi.org/10.1088/1748-0221/11/02/C02069 [26]

Papers 2, 3 and 6 are published in an international peer-reviewed scientific journal dedicated to the instrumentation for detector and accelerator science, indexed in the Web of Science database, having an impact factor of 1.454 in 2021. Papers 1 and 4 were published in the proceedings of international scientific conferences organized in Romania, conferences that are also indexed within the Web of Science database. Paper 5 is part of the proceedings of the Topical Workshop on Electronics for Particle Physics (TWEPP) 2018 where the author also participated with a poster (link[9]: https://indico.cern.ch/event/697988/contributions/3056039/).

Paper 6 presents the ROC's context, data formats and design. Papers 2 and 3 include such descriptions for completeness. Thus, the research presented in Section 2 is disseminated through these three papers. Paper 3 presents the ROC

[9]requires a CERN account.

digital functional test setup, its design and performance validation and the mass-testing results. Paper 4 describes the improved synchronization method for the clock and data signals within this test setup. Thus, the research presented in Section 3 is disseminated through these two papers. The ROC's radiation qualification presented in Section 4 is disseminated through paper 2. The research related to the FPGA ILA from Section 5 is disseminated through paper 1.

References

[1] The ATLAS Collaboration et al. The ATLAS experiment at the CERN large hadron collider. *Journal of Instrumentation*, 3(08):S08003–S08003, aug 2008. doi:10.1088/1748-0221/3/08/s08003.

[2] Lyndon Evans and Philip Bryant. LHC machine. *Journal of Instrumentation*, 3(08):S08001–S08001, aug 2008. doi:10.1088/1748-0221/3/08/s08001.

[3] ATLAS Collaboration et al. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Physics Letters B*, 716(1):1–29, 2012. doi:https://doi.org/10.1016/j.physletb.2012.08.020.

[4] The ALPHA Collaboration et al. Confinement of antihydrogen for 1,000 seconds. *Nature Physics*, 7(7):558–564, Jul 2011. doi:10.1038/nphys2025.

[5] Stephanie Sammartino McPherson. *Tim Berners-Lee: Inventor of the World Wide Web*. Twenty First Century Books, USA, 1st edition, 2009.

[6] Sylvie Braibant and Giorgio Giacomelli. *Particles and fundamental interactions: An introduction to particle physics*. Undergraduate lecture notes in physics. Springer, Dordrecht, Netherlands, 2012. doi:10.1007/978-94-007-2464-8.

[7] David J Griffiths. *Introduction to elementary particles; 2nd rev. version*. Physics textbook. Wiley, New York, NY, 2008. URL: https://cds.cern.ch/record/111880.

[8] *Particle Accelerators*, chapter 12, pages 363–393. Springer, Boston, MA, USA, 1972. doi:10.1007/978-1-4615-9701-8_12.

[9] John Douglas Cockcroft and E. T. S. Walton. Experiments with high velocity positive ions. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 129(811):477–489, 1930. doi:10.1098/rspa.1930.0169.

[10] Werner Herr and B Muratori. Concept of luminosity. 2006. URL: https://cds.cern.ch/record/941318, doi:10.5170/CERN-2006-002.361.

[11] Jorg Wenninger. Operation and Configuration of the LHC in Run 2. Mar 2019. URL: https://cds.cern.ch/record/2668326.

[12] ATLAS Electronics Coordination TDAQ Phase-II Upgrade Project. Atlas trigger & daq interfaces with detector front-end systems: Requirement document for hl-lhc. Atlas doc.: At2-di-es-0002, edms id: 1563801 v.1, CERN, Oct 2018. URL: https://edms.cern.ch/ui/file/1563801/1/RequirementsPhaseII_v1.1.0.pdf.

[13] CERN. Lhc the guide. Available at https://home.cern/sites/home.web.cern.ch/files/2018-07/CERN-Brochure-2017-002-Eng.pdf. Accessed: 05-01-2021.

[14] CERN. Lhc season 2. facts & figures. Available at https://home.cern/sites/home.web.cern.ch/files/2018-07/CERN-Brochure-2015-003-Eng.pdf. Accessed: 05-01-2021.

[15] The ALICE Collaboration et al. The ALICE experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08002–S08002, aug 2008. doi:10.1088/1748-0221/3/08/s08002.

[16] The CMS Collaboration et al. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004–S08004, aug 2008. doi:10.1088/1748-0221/3/08/s08004.

[17] The LHCb Collaboration et al. The LHCb detector at the LHC. *Journal of Instrumentation*, 3(08):S08005–S08005, aug 2008. doi:10.1088/1748-0221/3/08/s08005.

[18] Joao Pequenao. Computer generated image of the whole ATLAS detector. Mar 2008. URL: https://cds.cern.ch/record/1095924.

[19] Lyndon Evans and Lyn Evans. *The Large Hadron Collider: a marvel of technology; 2nd ed.* Physics (EPFL Press). EPFL Press, Lausanne, 2018. On the cover : Including the discovery of the higgs boson. URL: https://cds.cern.ch/record/2645935.

[20] I. Béjar Alonso, O. Brüning, P. Fessia, M. Lamont, L. Rossi, L. Tavian, and M. Zerlauth. Vol. 10 (2020): High-luminosity large hadron collider (hl-lhc): Technical design report. Cern yellow reports: Monographs, CERN, 2020. URL: https://e-publishing.cern.ch/index.php/CYRM/issue/view/127/95.

[21] Peter Vankov and ATLAS Collaboration. ATLAS Future Upgrade. Technical report, CERN, Geneva, Jun 2016. URL: https://cds.cern.ch/record/2195333, doi:10.22323/1.273.0061.

[22] T Kawamoto, S Vlachos, L Pontecorvo, J Dubbert, G Mikenberg, P Iengo, C Dallapiccola, C Amelung, L Levinson, R Richter, and D Lellouch. New Small Wheel Technical Design Report. Technical report, Jun 2013. ATLAS New Small Wheel Technical Design Report. URL: https://cds.cern.ch/record/1552862.

[23] Bernd Stelzer. The new small wheel upgrade project of the atlas experiment. *Nuclear and Particle Physics Proceedings*, 273-275:1160–1165, 2016. 37th International Conference on High Energy Physics (ICHEP). doi:10.1016/j.nuclphysbps.2015.09.182.

[24] Lorne Levinson. Nsw electronics overview, Oct 2020. Accessed: 5 January 2021. URL: https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/Overview%20and%20General/LL_NSW_ElxOvr_notitle_v13.png.

[25] G. De Geronimo, J. Fried, S. Li, J. Metcalfe, N. Nambiar, E. Vernon, and V. Polychronakos. Vmm1 - an asic for micropattern detectors. In *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)*, pages 633–639, Oct 2012. doi:10.1109/NSSMIC.2012.6551184.

[26] R.-M. Coliban, S. Popa, T. Tulbure, D. Nicula, M. Ivanovici, S. Martoiu, L. Levinson, and J. Vermeulen. The read out controller for the ATLAS new small wheel. *Journal of Instrumentation*, 11(02):C02069–C02069, feb 2016. doi:10.1088/1748-0221/11/02/c02069.

[27] P. Gkountoumis. Level-1 data driver card of the atlas new small wheel upgrade compatible with the phase ii 1 mhz readout scheme. In *2016 5th International Conference on Modern Circuits and Systems Technologies (MOCAST)*, pages 1–4, May 2016. doi:10.1109/MOCAST.2016.7495115.

[28] I. P. Mesolongitis, A. Gkountis, E. D. Kyriakis - Bitzaros, K. Zachariadou, P. Gkountoumis, and T. Alexopoulos. Testing the level-1 data driver card for the new small wheel of the atlas detector. In *2017 6th International Conference on Modern Circuits and Systems Technologies (MOCAST)*, pages 1–4, May 2017. doi:10.1109/MOCAST.2017.7937667.

[29] P Moreira, A Marchioro, and Kloukinas. The GBT: A proposed architecure for multi-Gb/s data transmission in high energy physics. 2007. URL: https://cds.cern.ch/record/1091474, doi:10.5170/CERN-2007-007.332.

[30] P Leitao, S Feger, D Porret, S Baron, K Wyllie, M Barros Marin, D Figueiredo, R Francisco, J C Da Silva, T Grassi, and P Moreira. Test bench development for the radiation hard GBTX ASIC. *Journal of Instrumentation*, 10(01):C01038–C01038, jan 2015. doi:10.1088/1748-0221/10/01/c01038.

[31] J Anderson, A Borga, H Boterenbrood, H Chen, K Chen, G Drake, D Francis, B Gorini, F Lanni, G Lehmann Miotto, L Levinson, J Narevicius, C Plessl, A Roich, S Ryu, F Schreuder, J Schumacher, W Vandelli, J Vermeulen, and J Zhang. FELIX: a high-throughput network approach for interfacing to front end electronics for ATLAS upgrades. *Journal of Physics: Conference Series*, 664(8):082050, dec 2015. doi:10.1088/1742-6596/664/8/082050.

[32] Address in real-time concentrator asic. Available at https://espace.cern.ch/ATLAS-NSW-ELX/Shared%20Documents/ADDC/ART%20ASIC/artasic_specs1.2_020316.pdf. Accessed: 02-04-2021.

[33] L. Yao, H. Chen, K. Chen, S. Tang, and V. Polychronakos. Design and Testing of the Address in Real-Time Data Driver Card for the Micromegas Detector of the ATLAS New Small Wheel Upgrade. *IEEE Trans. Nucl. Sci.*, 67(9):2155–2160, 2020. arXiv:1806.06912, doi:10.1109/TNS.2020.2986418.

[34] L. Yao, V. Polychronakos, H. Chen, K. Chen, H. Xu, S. Martoiu, N. Felt, and T. Lazovich. The address in real time data driver card for the MicroMegas detector of the ATLAS muon upgrade. *Journal of Instrumentation*, 12(01):C01047–C01047, jan 2017. doi:10.1088/1748-0221/12/01/c01047.

[35] Jinhong Wang, Liang Guan, J. W. Chapman, Bing Zhou, and Junjie Zhu. Design of a trigger data serializer asic for the upgrade of the atlas forward muon spectrometer. *IEEE Transactions on Nuclear Science*, 64(12):2958–2965, Dec 2017. doi:10.1109/TNS.2017.2771266.

[36] A. X. Widmer and P. A. Franaszek. A dc-balanced, partitioned-block, 8b/10b transmission code. *IBM Journal of Research and Development*, 27(5):440–451, Sep. 1983. doi:10.1147/rd.275.0440.

[37] Solid State Technology Association. Scalable low-voltage signaling for 400 mv (slvs-400). Standard jesd8-13, JEDEC, Oct 2001. URL: http://www.jedec.org/sites/default/files/docs/jesd8-13.pdf.

[38] S Bonacini, K Kloukinas, and P Moreira. E-link: A Radiation-Hard Low-Power Electrical Link for Chip-to-Chip Communication. 2009. URL: https://cds.cern.ch/record/1235849, doi:10.5170/CERN-2009-006.422.

[39] S. Popa, S. Mărtoiu, and M. Ivanovici. Study of the ATLAS new small wheel readout controller ASIC in a neutron irradiation environment. *Journal of Instrumentation*, 15(10):P10023–P10023, oct 2020. doi:10.1088/1748-0221/15/10/p10023.

[40] K Poltorak, F Tavernier, and P Moreira. A radiation-hard PLL for frequency multiplication with programmable input clock and phase-selectable output signals in 130 nm CMOS. *Journal of Instrumentation*, 7(12):C12014–C12014, dec 2012. doi:10.1088/1748-0221/7/12/c12014.

[41] M.E. Levitt. Asic testing upgraded. *IEEE Spectrum*, 29(5):26–29, May 1992. doi:10.1109/6.135405.

[42] Xilinx. Kcu105 board, ug917 (v1.10). Available at https://www.xilinx.com/support/documentation/boards_and_kits/kcu105/ug917-kcu105-eval-bd.pdf, February 2019. Accessed: 19-04-2021.

[43] Xilinx. Microblaze processor reference guide. ug984 (v2020.2). Available at https://www.xilinx.com/support/documentation/sw_manuals/xilinx2020_2/ug984-vivado-microblaze-ref.pdf, November 2020. Accessed: 19-02-2021.

[44] Xilinx. Integrated Logic Analyzer v6.2, PG172. Available at https://www.xilinx.com/support/documentation/ip_documentation/ila/v6_2/pg172-ila.pdf, 2016. Accessed 14-05-2020.

[45] R. Ginosar. Metastability and synchronizers: A tutorial. *IEEE Design Test of Computers*, 28(5):23–35, Sep. 2011. doi:10.1109/MDT.2011.113.

[46] Theodoros Vafeiadis. Integration and commissioning of ATLAS New Small Wheel Micromegas detectors with electronics at CERN. In *Proceedings of 40th International Conference on High Energy physics — PoS(ICHEP2020)*, volume 390, page 791, 2 2021. doi:10.22323/1.390.0791.

[47] ATLAS Resources Review Board. Atlas upgrade status report 2018 - 2019. Technical report, October 2018. Accessed: 29 October 2019. URL: https://cds.cern.ch/record/2638085/files/CERN-RRB-2018-079.pdf.

[48] S. Popa, M. Luchian, and M. Ivanovici. Clock and data signals synchronization for an fpga-based asic testing setup. In *2019 International Symposium on Signals, Circuits and Systems (ISSCS)*, pages 1–4, July 2019. doi:10.1109/ISSCS.2019.8801780.

[49] Omid Zeynali, Daryoush Masti, Maryam Nezafat, and Alireza Mallahzadeh. Study of "radiation effects of nuclear high energy particles" on electronic circuits and methods to reduce its destructive effects. *Journal of Modern Physics*, 2(12):1567–1573, 2011. doi:10.4236/jmp.2011.212191.

[50] R. Bailey, editor. *CAS - CERN Accelerator School: Power Converters - Radiation Risks and Mitigation in Electronic Systems*, Geneva, May 2014. CERN, CERN. URL: https://cds.cern.ch/record/2038628/files/245-263-Todd.pdf, doi:10.5170/CERN-2015-003.

[51] L. Artola, M. Gaillardin, G. Hubert, M. Raine, and P. Paillet. Modeling single event transients in advanced devices and ics. *IEEE Transactions on Nuclear Science*, 62(4):1528–1539, Aug 2015. doi:10.1109/TNS.2015.2432271.

[52] Tejinder Singh, Farzaneh Pashaie, and Rajat Kumar. Redundancy based design and analysis of alu circuit using cmos 180nm process technology for fault tolerant computing architectures. *International Journal of Computing and Digital Systems*, 4:53–62, 01 2015. doi:10.12785/ijcds/040106.

[53] T. Alexopoulos, G. Fanourakis, T. Geralis, M. Kokkoris, A. Kourkoumeli-Charalampidi, K. Papageorgiou, and G. Tsipolitis. Study of the VMM1 read-out chip in a neutron irradiation environment. *Journal of Instrumentation*, 11(05):P05015–P05015, may 2016. doi:10.1088/1748-0221/11/05/p05015.

[54] European Space Agency. Single event effects test method and guidelines., Oct 2014. ESCC Basic Specification No. 25100, Accessed: 15 June 2020. URL: https://escies.org/download/webDocumentFile?id=62690.

[55] Altera. Design debugging using the signaltap ii logic analyzer. Available at https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/qts/qts_qii53009.pdf, 2013. Accessed: 19-04-2021.

[56] ARM. Amba 4 axi4-stream protocol, version: 1.0, specification. Available at https://developer.arm.com/documentation/ihi0051/a/Introduction/About-the-AXI4-Stream-protocol, 2010. Accessed: 19-04-2021.

[57] Ieee standard for ethernet. *IEEE Std 802.3-2018 (Revision of IEEE Std 802.3-2015)*, pages 1–5600, Aug 2018. doi:10.1109/IEEESTD.2018.8457469.

[58] Philippe Farthouat. New small wheels roc production readiness review. Technical report, May 2018. Accessed: 19 May 2018. URL: https://edms.cern.ch/ui/file/1977683/1/ROC-PRR_docx_cpdf.pdf.

[59] Georgios Iakovidis. VMM3a, an ASIC for tracking detectors. May 2019. URL: http://cds.cern.ch/record/2675779/files/ATL-MUON-SLIDE-2019-205.pdf?version=1.

[60] Atlas authorship policy. version 7.4. Available at https://twiki.cern.ch/twiki/pub/AtlasProtected/AtlasPolicyDocuments/Authorship_Policy.pdf. Accessed: 26-05-2021.

[61] ATLAS Collaboration. Search for flavour-changing neutral currents in processes with one top quark and a photon using 81 fb(-1) of pp collisions at root s=13 TeV with the ATLAS experiment. *PHYSICS LETTERS B*, 800, JAN 10 2020.

[62] S. Popa, M. Ivanovici, and R. M. Coliban. Time-multiplexed 10gbps ethernet-based integrated logic analyzer for fpgas. In *2020 International Symposium on Electronics and Telecommunications (ISETC)*, pages 1–4, Nov 2020. doi:10.1109/ISETC50328.2020.9301115.

[63] S. Popa, S. Mărtoiu, and M. Ivanovici. The quality-control test of the digital logic for the ATLAS new small wheel read-out controller ASIC. *Journal of Instrumentation*, 15(04):P04023–P04023, apr 2020. doi:10.1088/1748-0221/15/04/p04023.

[64] Stefan Popa, Sorin Martoiu, Mihai Luchian, Radu Coliban, and Mihai Ivanovici. The Quality-Assurance Test of the ATLAS New Small Wheel Read-Out Controller ASIC. In *Proceedings of Topical Workshop on Electronics for Particle Physics — PoS(TWEPP2018)*, volume 343, page 081, 2019. doi:10.22323/1.343.0081.