



BNL-112560-2016-CP

An Architecture for a Mitigated FPGA Multi-Gigabit Transceiver for High Energy Physics Environments

**M. Brusati, Camplani, M. Cannon, H. Chen,
M. Citterio, M. Lazzaroni, H. Takai, M. Wirthlin**

Presented at the 14th IMEKO TC10 Workshop on Technical Diagnostics
Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB)
Milano, Italy
June 27-28, 2016

June 2016

Physics Department

Brookhaven National Laboratory

**U.S. Department of Energy
USDOE Office of Science (SC),
High Energy Physics (HEP) (SC-25)**

Notice: This manuscript has been authored by employees of Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. The publisher by accepting the manuscript for publication acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or any third party's use or the results of such use of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof or its contractors or subcontractors. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

IMEKO

International Measurement Confederation

JUNE 27-28th
MILAN



Technical
Committee
TC10
Technical Diagnostic

An Architecture for a Mitigated FPGA Multi-Gigabit Transceiver for High Energy Physics Environments

M. Brusati¹, A. Camplani¹, M. Cannon³, H. Chen⁴,
 M. Citterio², M. Lazzaroni^{1,2,*}, H. Takai⁴, M. Wirthlin³

¹*Dipartimento di Fisica, Università degli Studi di Milano, via G. Celoria, 16, 20133 Milano, Italy, marta.brusati@studenti.unimi.it, alessandra.camplani@unimi.it, massimo.lazzaroni@unimi.it*

²*INFN Milano, via G. Celoria, 16, 20133 Milano, Italy, mauro.citterio@mi.infn.it*

³*NSF Center for High Performance Reconfigurable Computing, Department of Electrical and Computer Engineering, Brigham Young University, Provo, Utah, U.S.A., matthew.cannon@byu.edu, wirthlin@byu.edu*

⁴*Brookhaven National Laboratory, Upton, New York, U.S.A., chc@bnl.gov, takai@bnl.gov*

*Corresponding author

Abstract - SRAM-based Field Programmable Gate Array (FPGA) logic devices are very attractive in applications where high data throughput is needed, such as the latest generation of High Energy Physics (HEP) experiments. FPGAs have been rarely used in such experiments because of their sensitivity to radiation. The present paper proposes a mitigation approach applied to commercial FPGA devices to meet the reliability requirements for the front-end electronics of the Liquid Argon (LAR) electromagnetic calorimeter of the ATLAS experiment, located at CERN. Particular attention will be devoted to define a proper mitigation scheme of the multi-gigabit transceivers embedded in the FPGA, which is a critical part of the LAR data acquisition chain. A demonstrator board is being developed to validate the proposed methodology. Mitigation techniques such as Triple Modular Redundancy (TMR) and scrubbing will be used to increase the robustness of the design and to maximize the fault tolerance from Single-Event Upsets (SEUs).

I. INTRODUCTION

FPGAs are very sensitive to radiation-induced Single-Event Upsets (SEUs), which occur when a single charged particle passes through the device and transfers charge between nodes of the active regions. This sensitivity is intrinsic in FPGA architecture since they are composed of a large amount of memory cells in a small circuit area. Mitigation techniques have been widely investigated as their use is essential for hostile environments to ensure the correct operation [1, 2].

Even though the sensitivity of FPGAs and memories (such as SRAM or DRAM) are very similar, they have very different failure modes. Typical memory cells store data and any upset will only corrupt that data. In FPGAs, however, the largest component of memory is used as configuration memory, which define the operation of the configurable logic blocks, routing resources, input/output blocks and other programmable resources. An upset in a configuration memory cell may thus change the operation of the user defined circuit.

The failure modes of FPGAs are also different from custom circuit technology – such as Application-Specific Integrated Circuits (ASICs) – where the design has a fixed functionality and does not need a configuration memory. Since the custom circuit is hardwired, only latches need radiation protection techniques and the routing and logic are usually considered insensitive to soft errors [3]. FPGAs require SEU mitigation to ensure adequate reliability for many different resources (latches, logic, routing and I/O), since they are subject to alteration from SEUs.

The benefits of this technology are nonetheless very intriguing: the high throughput and parallel computing features are important elements in a high demanding application like the data processing of the ATLAS calorimeter.

The reprogrammability of the board is also an interesting feature as the circuit can be changed by simply loading a new bitfile: it allows firmware upgrades without the need to physically replace the electronics and also provides easy access for error correction.

The present work is part of an ongoing collaboration that is based on previous studies on the HEP environment, FPGA reliability, resistance to radiation and mitigation

techniques [4–11].

Serializers embedded in the latest generation of FPGAs present an interesting solution for high speed data transmission applications, but can operate unreliably in high radiation environments. In this paper we present an approach that allows the use of commercial off-the-shelf (COTS) FPGAs in a hostile environment (like the ATLAS LAr calorimeter front-end electronics), verified by creating a mitigated design of a multi-gigabit transceiver that will be installed on a validation board in the experiment.

The paper is structured as follows: the next section will describe the radiation environment expected for the LAr calorimeter front-end electronics, followed by a description of the damage caused to programmable logic by radiation and the error classification used. In Section III the techniques for mitigation used for this project will be introduced and in Section IV the project layout will be presented. In Section V some final considerations are drawn.

II. HIGH ENERGY ENVIRONMENT

The LAr calorimeter of the ATLAS experiment is located at the Large Hadron Collider (LHC) in Geneva. ATLAS is a very complex experiment, as can be seen in Fig. 1, and it is designed to precisely detect very high energetic events. It has observed a particle as massive as the Higgs boson, i.e. near 126 GeV [12]. The ATLAS detector consists of three main sub-systems: the inner tracking system, the calorimeter, and the muon spectrometer. The calorimeter system is composed of sampling detectors and is housed in one barrel and two endcap cryostats. The barrel is made of a LAr Electromagnetic Calorimeter (EC), surrounded by a hadronic calorimeter made of steel and scintillating tile (TileCal). In the endcap region of the detector all of the calorimetry uses LAr as the active material. The EC is made of alternating layers of accordion-shaped lead absorbers and electrodes. Liquid argon is placed between these layers as an active medium. Interactions in the absorbers transform the incident energy into a “shower” of particles that are detected by the sensing element. The energies of the particles generated from the collision event are reconstructed from a complex analysis of the signals coming from all detectors.

The present paper deals with the electronics for the barrel LAr electromagnetic calorimeter. The front-end electronics are placed near the EC to perform the first signal processing (such as preamplification, analog-to-digital conversion and serialization for transmission via optical fiber), and the crate is exposed to a high radiation level and magnetic fields that cause frequent failures in non-protected devices. The LAr calorimeter is described in more detail in [13] and [14].

The magnetic field that affects ATLAS front-end electronics goes up to 0.1 T, and some care in the choice of the components must be taken to avoid any unwanted effects. On the other hand the radiation exposure in the experiment is very consistent and is composed by ionizing and non-

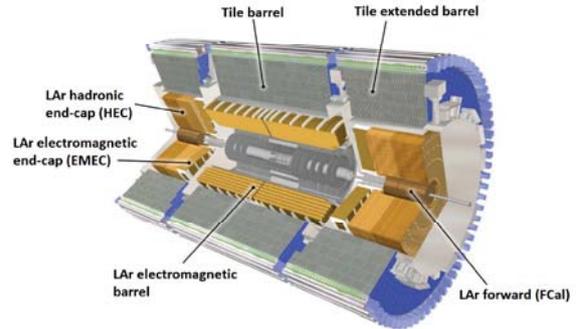


Fig. 1. Picture of the ATLAS experiment [14].

Table 1

	Simulation (one year)	Safety Factor	Test Target* (10years)
Ionizing Dose	3.0 rad	10	100 krad
1 MeV eq. Neutron	$6.0 \cdot 10^{11} \text{ cm}^{-2}$	2	$1.2 \cdot 10^{13} \text{ cm}^{-2}$
Hadrons ($>20 \text{ MeV}$)	$8.5 \cdot 10^{10} \text{ cm}^{-2}$	2	$2 \cdot 10^{12} \text{ cm}^{-2}$

* 1 LHC year = 10^7 s , $\sigma_{pp} = 80 \text{ mb}$,
 Luminosity = $5 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$;

ionizing particles, constituting a very hostile environment for programmable logic.

Some basic definitions for dealing with radiation damage are recalled in Appendix A.

The expected background for the LAr calorimeter electronics for Phase II runs are shown in Table 1, while Fig. 2 shows the background due to a mixed field of hadrons, electrons and photons. Phase II background will be ~ 10 times higher than the one shown in the figure.

Circuit damage caused by radiation due to Single-Event Effects (SEE) will be considered. SEEs are the radiation response of a semiconductor caused by a sin-

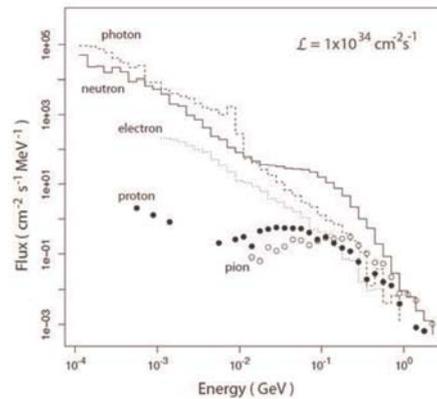


Fig. 2. Simulated spectrum of particles for the ATLAS LAr calorimeter electronics.

gle particle. SEEs include Single-Event Upset (SEU), Multiple-Bit Upset (MBU), Single-Event Functional Interrupt (SEFI), Single-Event Latch-up (SEL), Single-Event Transient (SET), Single-Event Burnout (SEB) and Single-Event Gate Rupture (SEGR) [15, 16].

There are two types of errors: soft and hard errors. Soft errors are nondestructive functional errors that do not cause a permanent damage to the material. They are a subset of SEE, and include SEU, MBU, SEFI, SET and SEL. These errors are usually associated with erroneous output signals from a latch or memory cell that can be corrected. On the other hand hard errors are irreversible: the device is permanently damaged, data is lost and its operation is no longer recovered even after power reset and re-initialization. Hard errors can be caused by SEGR and SEB. A description of each error is provided in Appendix A.

The present paper will focus on SEUs, in particular on sensitive bits, which are configuration bits that directly affect the behavior of the circuit when upset (as opposed to other configuration bits that do not affect the functionality of the circuit). Sensitivity depends on the specific design, as it is related to the location and number of configuration bits used for the design. There are a variety of ways to check if a bit is sensitive to SEUs, but the most common are radiation testing and fault injection. The latter is by far cheaper and faster, and has been shown to be a good solution. In [17] the fault injection tool predicted approximately 97% of the output errors found during radiation testing.

We are interested in measuring the failure occurrences of the system, by means of a diagnostic procedure that allows a further distinction into different modes of disruption. We can indeed classify errors into persistent and nonpersistent ones, as was done in [9].

The distinction is based on the duration of the failure: if the system restores its correct functionality after some time, the error is classified as nonpersistent and the upset bit is stored as a nonpersistent one. If the SEU causes a permanent error in the design performance which can be restored only through a global system reset, the faulty bit is treated as persistent instead.

A. Nonpersistent Errors

Nonpersistent errors are only temporary errors which are fixed after the scrubbing procedure. This technique rewrites the correct configuration memory on the device, as will be discussed in the next section. We can see the behavior of such an error in Fig. 3: the Device Under Test (DUT) calculates the square number of the input, but when an SEU occurs its output becomes incorrect. The third sequence of the figure shows the arithmetic difference between the expected result and the DUT output. As long as the faulty bit produces an incorrect output the difference is not zero, but after configuration scrubbing occurs correct operation mode of the upset circuit is restored and the sequence is null again. This example shows that nonpersis-

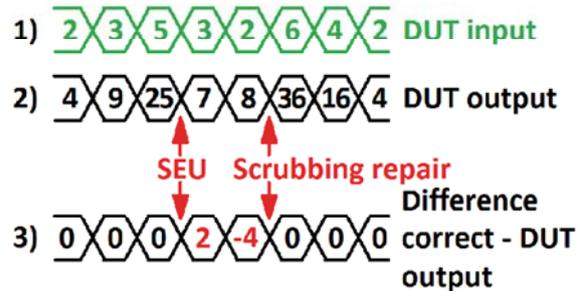


Fig. 3. Example of a nonpersistent error in a squarer.

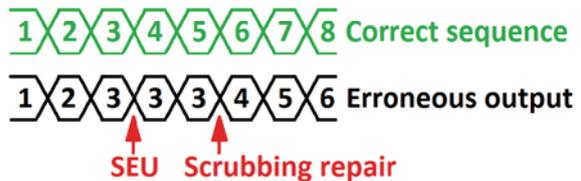


Fig. 4. Example of persistent error in a binary counter.

tent behavior appears in circuits where the output depends on the present value of the inputs, so that the correct functioning is restored when the configuration memory is repaired by scrubbing.

B. Persistent Errors

A bit is persistent if its alteration produces an error that will indefinitely propagate in the circuit, even after scrubbing has restored the original value. A correction of this type of error can be made only through a system reset.

These errors occur because the scrubbing procedure restores the circuit structure, but the temporary failure may have generated a faulty state which will not be corrected until a system reset. Persistent errors are thus caused by upsets in the configuration memory corresponding to feedback loops and internal state storage, which receive a faulty state and maintain it until a reset occurs.

An example of a circuit vulnerable to persistent errors is a binary counter: the next value is computed from the current one, so the output depends on the previous state in the counting sequence. If a SEU damages the counter and the output does not change (as shown in Fig. 4), this output will be wrong even after the scrubbing procedure. Correction of the faulty bit assures the restoration of correct operation but the counter will start from the wrong state, losing the correct sequence.

To conclude it is clear that the electronics used in the experiment must be protected in order to operate without performance and reliability degradation. Reliability is an important issue in experiments that can be physically accessed only during shutdown time, and remote maintenance must be provided in order to avoid complete data loss.

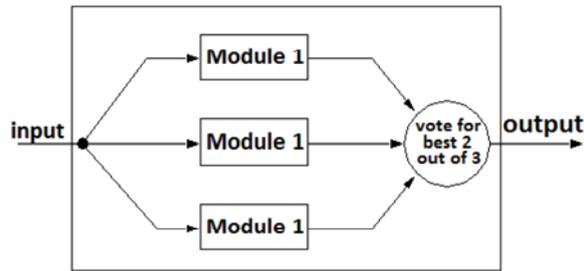


Fig. 5. Triple Modular Redundancy scheme.

III. MITIGATION TECHNIQUES

Since FPGAs can operate unreliably in a high radiation environment, mitigation techniques are needed to correct or mask the faults caused by SEUs. In this section the mitigation techniques used to improve the FPGAs robustness will be presented.

The mitigated circuitry itself is subject to SEUs, so it has to be carefully designed to minimize the potential sensitivity and avoid additional reliability losses.

A. TMR: Triple Modular Redundancy

TMR is a static hardware redundancy scheme for masking single faults in a digital circuit. The masking is performed by triplicating logic resources and inserting a majority voter, as in Fig. 5. In this way if a SEU changes the state of a module, the majority voter will see two equal results and a flawed one, and will choose to transmit only the good result, assuming that only one module failed. If a particle strike were to cause two errors or if errors are allowed to accumulate, two domains could fail and TMR would choose an incorrect output. To reduce the probability of fault accumulation, a scrubbing technique is usually implemented to prevent errors from accumulating and breaking TMR.

The TMR technique has a very high area cost, since it requires to triple the whole design and needs an additional circuitry for the majority voter implementation. To ensure protection against SEUs, the mitigated design will occupy an area that goes from a minimum of three up to six times the original design area [3]. There are also issues related to timing performance and power consumption of the design, since the additional resources needed for mitigation may worsen the design characteristics. In particular the voter circuitry is made of combinational logic which increases the critical path length and negatively affects timing [18].

Previous studies have investigated alternative mitigation techniques to TMR [3]. However, they have shown that common ASIC mitigation techniques are not viable for FPGA design, making TMR the best available choice for reliability and resource requirements.

TMR is therefore one of the best techniques for improving a FPGA's reliability under radiation: it has been widely

used and has shown significant enhancements on different designs. Some tools for automatic TMR implementation are available, such as the BLTMR, developed at BYU and LANL that will be used in the present work [2].

B. Configuration Scrubbing

Configuration scrubbing is a mitigation technique that provides error correction by repeatedly scanning and cleaning configuration upsets. It is usually composed of an external radiation hardened circuit that has access to the FPGA configuration memory and compares it against a golden copy. If a discrepancy is found, the fault is corrected and scanning resumes. Configuration scrubbing is necessary in circuits that use mitigation such as TMR, because it prevents the buildup of configuration faults. Without scrubbing TMR would become ineffective, since faults will accumulate and eventually overcome the voter output.

There are a variety of scrubbing modes that can be utilized for a system. Three of the most common scrubbing methods are blind, readback and hybrid. Blind scrubbing continuously writes the golden bitstream to the FPGA, even if there are no errors in the system. This method will correct errors the quickest, but is often unnecessary and utilizes a high amount of bandwidth. Readback scrubbing reads the configuration memory, compares it to the golden bitstream and then corrects any errors. While not very useful in a deployed system, this method provides useful data about upset rates and failure modes. The final method, hybrid scrubbing, utilizes the on chip scrubber to fix single bit errors (using the ECC codes) and only uses external circuitry and memory when a MBU occurs. This method is slower than blind scrubbing, but utilizes less system bandwidth.

IV. PROJECT SETUP

The board being developed hosts Xilinx 7-Series Kintex 325T FPGAs. It is a commercial off-the-shelf device, not Radiation Hardened By Design (RHBD), to prove that the needed reliability can be achieved with proper mitigation techniques. This Kintex device has also been chosen because its behavior under radiation has been extensively studied in previous tests that measured its static cross section with a wide-spectrum neutron beam, high energy hadrons, high energy protons and heavy ions [10]. Preliminary tests with proton irradiation of a GTX transceiver design have been carried out as well [11]. RHBD SRAM-based FPGAs are currently used in space related applications, but the cost of such devices is prohibitive for our purposes (the LAr front-end electronics has thousands of channel, needing ~ 1600 boards for the readout), not to mention that these devices have not been tested for the LHC radiation environment.

The ATLAS experiment is also very demanding in terms of throughput, since the bunch crossing of the LHC occur every 25 ns . The minimum data throughput of the

Table 2. Data from [19].

Kintex XC7K325T Important Features	
User FFs	407,600
Logic cells	326,080
DSP Slices	840
Block RAM Memory (36 Kb)	445
Internal memory	16.4 Mb
Transceivers	16

board using optical links must be 100 Gb/s. Communication from the front-end crate to the back-end readout electronics is provided via 12 optical links, so an acceptable solution is to provide 10 fibers running at 10 Gb/s, while using the extra links to provide redundancy or external data transfer. The device is provided with 16 GTX transceivers that can handle 12.5 Gb/s each, providing the needed performance. Other notable features of the used device are shown in Table 2. For a complete overview refer to the Xilinx documentation [19, 20].

ATLAS electronics must be designed to interface correctly with the existing crate, in particular with the power lines and cooling structure. Power distribution is bound by the preexisting experiment constraints and sets an upper limit of 100 W of power dissipation for the single board, one third of which is reserved for the FPGA (i.e. $\sim 33W$). The cooling structure consists of heat sinks and liquid cooling only, since the performance of a fan would be compromised by the strong magnetic fields of the area.

The final board will be configured as shown in Fig. 6. The FPGA will receive data from a radiation resistant SRAM filled with synthetic data that mimics the real calorimeter raw output and allows corrections from the outside through the optical links and the controller. The board is built to support many different kinds of scrubbing. All of the scrubbing methods previously explained will be utilized to determine which method is the best. In addition, these methods will be implemented in a variety of different ways. An external SRAM and flash memory will be provided as data storage for the golden bitstream. A optical link dedicated to the JTAG port will also be provided to allow support for remote scrubbing. We will investigate a variety of techniques including a new method called “Paired-Self Scrubbing”, where two FPGAs with the same bitstream will scrub each other. Each scrubbing configuration will be studied in detail before deciding upon the final method.

The external memories can be accessed by the controller to ensure that firmware upgrades in the FPGA are loaded into the golden copy as well: if only the FPGA is provided with a new bitfile, the scrubber will read a configuration memory content different from its own and will “correct” the discrepancies accordingly. The FPGA uses its GTX transceivers to serialize and transmit the synthetic “calorimeter” data to the controller which in turn passes it to the fiber optics transmitter that sends it to the back-end electronics. The board will be connected to the terminal of

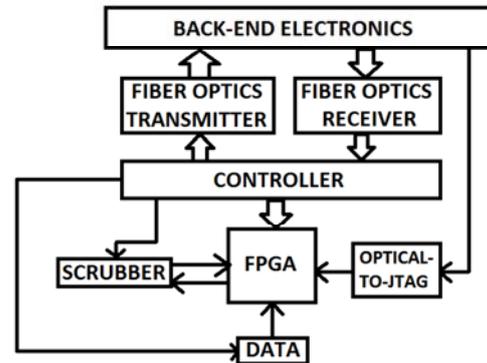


Fig. 6. Block diagram of the board.

an external computer that can be accessed remotely. The FPGA can send information on the line connection status and a message if an error has been detected. In this way a live check on the board’s behavior under radiation can be ensured.

In addition to scrubbing, each design will be mitigated with TMR using the latest tools available. Using these tools, we will be able to control the granularity of the domains and measure the trade-offs in the real system. The finer the granularity, the greater the reliability improvement, but the resource usage and power consumption greatly increase. We can also investigate the use of partial-TMR (pTMR) to protect the most sensitive and critical parts of the circuit to obtain the greatest reliability improvement while limiting the resource usage and power consumption.

V. CONCLUSION

In the current work we are presenting a new approach to radiation resistant front-end electronics for HEP experiments: a programmable-logic based serializer that meets the high performance requirements of the ATLAS experiment. The hardware and software development of the project and its hostile environment application are a challenging task that we are now confident to accomplish.

A. APPENDIX

The following definitions are used in the paper for radiation damage and error in semiconductor devices [16].

Dose is a measure of energy deposited per unit mass of medium and the unit of dose is rad or Gray. The SI unit for absorbed radiation is Gray (Gy) defined as the absorption of one Joule of radiation energy per one kilogram of matter.

Dose rate is the amount of ionizing radiation which an object would receive per unit of time.

Total dose is the total accumulated amount of absorbed ionizing radiation specified at a particular dose rate exposure at +25°C.

Fluence is the number of particles crossing a surface per unit area.

Single-Event Upset (SEU) is a soft error that occurs when a single particle strikes the sensitive volume of a memory cell, generating a charge which causes a change in the logic state of the cell. That node will remain in the upset state until new data is written into the memory element.

Single-Event Transient (SET) is a temporary voltage variation at a node in an integrated circuit caused by a single energetic particle strike. It is a glitch that propagates through the circuit, and if it results in a proper change of state it becomes an SEU.

Single-Event Functional Interrupt (SEFI) is a detectable functional failure, often associated with an upset in a control bit or register and can be corrected with a device reconfiguration or a reset.

Single-Event Latch-up (SEL) is an anomalous high-current state caused by the passage of a single particle in the device. If the current is above device specification, it must be immediately cleared with a power cycle to avoid permanent damage to the device. An example is in a CMOS device, where the PNP structure can be seen as a PNP and an NPN BJT, stacked next to each other. SEL occurs when the parasitic current energy is able to switch on one of the BJTs.

Single-Event Burnout (SEB) is a hard error, caused by a high-current state that (unlike SEL) results in device damage.

Single-Event Gate Rupture (SEGR) occurs when the gate oxide is damaged and a new current path is created. It is a hard error because it cannot be corrected.

Multiple-Bit Upset (MBU) takes place when two or more error bits occur in the same word, and cannot be corrected by a simple single-bit Error Correcting Code (ECC).

REFERENCES

- [1] L. Sterpone, M. Violante, "Analysis of the robustness of the TMR architecture in SRAM-based FPGAs", *IEEE Transactions on Nuclear Science*, Oct 2005, vol. 52, no. 5, pp. 1545–1549, ISSN 0018-9499, <http://dx.doi.org/10.1109/TNS.2005.856543>.
- [2] B. Pratt, M. Caffrey, P. Graham, K. Morgan, M. Wirthlin, "Improving FPGA design robustness with partial TMR", *IEEE, Reliability Physics Symposium Proceedings*, 2006. 44th Annual., IEEE International, pp. 226–232, <http://dx.doi.org/10.1109/RELPHY.2006.251221>.
- [3] K. S. Morgan, D. L. McMurtrey, B. H. Pratt, M. J. Wirthlin, "A comparison of TMR with alternative fault-tolerant design techniques for FPGAs", *IEEE Transactions on Nuclear Science*, 2007, vol. 54, no. 6, pp. 2065–2072, <http://dx.doi.org/10.1109/TNS.2007.910871>.
- [4] M. Lazzaroni, M. Citterio, S. Latorre, A. Lanza, P. Cova, N. Delmonte, F. Giuliani, "Features measurement and reliability considerations on a proposed main converter for LHC experiments", *Measurement*, 2015, vol. 71, pp. 73 – 88, ISSN 0263-2241, <http://dx.doi.org/10.1016/j.measurement.2015.04.002>.
- [5] S. Baccaro, G. Busatto, M. Citterio, P. Cova, N. Delmonte, F. Iannuzzo, A. Lanza, M. Riva, A. Sanseverino, G. Spiazzi, "Reliability oriented design of power supplies for high energy physics applications", *Microelectronics Reliability*, 2012, vol. 52, no. 9–10, pp. 2465 – 2470, ISSN 0026-2714, <http://dx.doi.org/10.1016/j.microrel.2012.06.088>.
- [6] P. Cova, N. Delmonte, F. Giuliani, M. Citterio, S. Latorre, M. Lazzaroni, A. Lanza, "Thermal optimization of water heat sink for power converters with tight thermal constraints", *Microelectronics Reliability*, 2013, vol. 53, no. 9–11, pp. 1760 – 1765, ISSN 0026-2714, URL <http://www.sciencedirect.com/science/article/pii/S0026271413002102>, <http://dx.doi.org/10.1016/j.microrel.2013.07.035>.
- [7] M. Lazzaroni, M. Citterio, S. Latorre, A. Lanza, "High B test of a commercial step-down Point of Load for LHC experiments", *Measurement*, 2015, vol. 60, pp. 194 – 206, ISSN 0263-2241, URL <http://www.sciencedirect.com/science/article/pii/S0263224114004850>, <http://dx.doi.org/10.1016/j.measurement.2014.10.018>.
- [8] J. Ameel, et al., "Radiation-hard power electronics for the ATLAS new small wheel", *Journal of Instrumentation C01009*, 2015, vol. 10, no. 01, pp. 1 – 8, ISSN 1478-0221, <http://dx.doi.org/10.1088/1748-0221/10/01/C01009>.
- [9] K. Morgan, M. Caffrey, P. Graham, E. Johnson, B. Pratt, M. Wirthlin, "SEU-induced persistent error propagation in FPGAs", *IEEE Transactions on Nuclear Science*, 2005, vol. 52, no. 6, pp. 2438–2445, <http://dx.doi.org/10.1109/TNS.2005.860674>.
- [10] M. J. Wirthlin, H. Takai, A. Harding, "Soft error rate estimations of the Kintex-7 FPGA within the ATLAS liquid argon (LAr) calorimeter", *Journal of Instrumentation (JINST)*, 2014, vol. 9, no. C01025, <http://dx.doi.org/10.1088/1748-0221/9/01/C01025>.
- [11] M. Cannon, M. Wirthlin, A. Camplani, M. Citterio, C. Meroni, "Evaluating Xilinx 7 Series GTX Transceivers for Use in High Energy Physics Experiments Through Proton Irradiation", *Nuclear Science, IEEE Transactions on*, 2015, vol. 62, no. 6, pp. 2695–2702, <http://dx.doi.org/10.1109/TNS.2015.2497216>.
- [12] G. Aad, et al., "Observation of a new particle in the search for the standard model higgs boson with the ATLAS detector at the LHC", *Phys.Lett. B716*, Jul 2012, pp. 1–29, <http://dx.doi.org/10.1016/j.physletb.2012.08.020>.
- [13] ATLAS LAr Collaboration, "LAr Technical Design Report", CERN/LHCC, 1996, URL <http://cds.cern>.