FPGA Qualification and Failure Rate Estimation Methodology for LHC Environments Using Benchmarks Test Circuits

FPGA Qualification and Failure Rate Estimation Methodology for LHC Environments Using Benchmarks Test Circuits

Antonio Scialdone, Rudy Ferraro, Rubén García Alía, Luca Sterpone, Salvatore Danzeca and Alessandro Masi

***Abstract*—When studying the behavior of a Field Programmable Gate Array (FPGA) under radiation, the most commonly used methodology consists in evaluating the SEE cross-section of its elements individually. However, this method does not allow the estimation of the device failure rate when using a custom design. An alternative approach based on benchmark circuits is presented in this article. It allows standardized application-level testing, which makes the comparison between different FPGAs easier. Moreover, it allows the evaluation of the FPGA failure rate independently from the application that will be implemented. The employed benchmark circuit belongs to the ITC'99 benchmark suite developed at Politecnico di Torino. Using the proposed methodology, the response of four FPGAs, the NG-Medium, the ProASIC3, the SmartFusion2 and the PolarFire, was evaluated under high-energy protons. Radiation tests with thermal neutrons were also conducted on the PolarFire to assess its potential sensitivity to them. Moreover, its performances in terms of TID effects have been evaluated by measuring the degradation of the propagation delay during irradiation.**

***Index Terms*—Field-programmable gate array (FPGA), Total Ionizing Dose (TID), Single Event Effects (SEE), Failure Estimation, Radiation Tests, Benchmark Tests, protons, thermal neutrons**

I. Introduction

AT CERN, many electronic systems are installed in the mixed-field radiation environment of the Large Hadron Collider (LHC). Thus, the Radiation Hardness Assurance (RHA) of electronic components is fundamental to ensure a reliable operation of accelerators and experiments. Because of their benefits in terms of costs, flexibility, and performances, Field Programmable Gate Arrays (FPGAs) are often at the core of several electronic systems. However, their lifetime and performances are affected by radiation-induced effects, such as Single-Event Effects (SEEs) and Total Ionizing Dose (TID). This sensitivity creates the necessity to perform many qualification tests in order to find a suitable FPGA to use in a specific system or experiment. In the past, many FPGAs, such as the ProASIC3, the SmartFusion2 [1], the Artix7 [2] and the NG-Medium [3], were qualified under radiation for CERN purposes. The ProASIC3 has been embedded in most of the accelerator systems in the past, whereas the SmartFusion2 and Igloo2 are used nowadays in the new developments. However, with the imminent approach of the High-Luminosity LHC (HL-LHC) and the consequent increase of the radiation levels, more robust FPGAs, able to withstand a higher fluence and ionizing dose, are necessary. Therefore, CERN is considering two new FPGAs as possible candidates for its applications: the NG-Medium and the PolarFire. The first is a Radiation Hardened by Design (RHBD) FPGA, manufactured using the STM C65 space process. The latter is the fifth generation of non-volatile FPGA device from Microsemi built on the state-of-the-art SONOS 28nm non-volatile process technology [4].

When performing a qualification test, the most commonly used procedure consists in evaluating the cross-section of each functional element (FE), i.e. DSPs, Flip-Flops, RAM, PLLs, separately [5]. Several radiation test data sets for this test topology are already available in the literature for many FPGAs, including those analyzed in this article, such as the PolarFire [4], [6]. Even though this approach yields a lot of information about the sensitivity of each FE, it does not give a realistic overview of how a custom application will work. Extrapolating the SEE susceptibility for an user design starting from the SEE response of its FEs is a difficult task. Other works, like [7] and [8], discussed the challenges and the consequences to face when performing such analyses on mission specific designs. This limitation makes the estimation of the device failure rate during LHC operation quite complex. For this reason, as mentioned also in [5], application-style tests are recommended to derive a realistic behavior of the system. Nonetheless, performing a test for each application is an expensive and time consuming task. In addition, interpreting the obtained results for other FPGAs belonging to different families and with different technology can be difficult. These limitations, together with the imminent approach of the HL-LHC, make the qualification process even more important and create the need to find a more reliable testing approach for the estimation of the device failure rate.

In this paper, a testing methodology to overcome the aforementioned problems is presented. The approach uses a benchmark application, which better reflects the workload of a real application and facilitates standardized application-level testing, making the comparison of different families of FPGAs easier. The benchmark belongs to the IT99 suite [9] developed by Politecnico di Torino. Such circuits have already been used in other reliability experiments, like [10] and [11], for evaluating the performances of different mitigation techniques

This work was partially funded thanks to the French national program "Programme d'Investissements d'Avenir, IRT Nanoelec" ANR-10-AIRT-05. The authors acknowledge the ILL for the beamtime allocated INDU-225.

Antonio Scialdone, Rudy Ferraro, Rubén García Alía, Salvatore Danzeca and Alessandro Masi are with the European Organization for Nuclear Research (CERN), CH-1211 Genève 23 Switzerland (e-mail: antonio.scialdone@cern.ch).

Antonio Scialdone and Luca Sterpone are with the Politecnico di Torino, Dip. Automatica e Informatica, Torino, Italy.

against SEEs. The advantages of this new methodology compared to the standard methodology will be proven by showing the results obtained for the SmartFusion2, the ProASIC3, the NG-Medium and the PolarFire under high-energy protons. Using the benchmark, a device reliability analysis has been performed to understand the failure modes of the different FPGAs, and compare the advantages and disadvantages of using one instead of another. For the PolarFire FPGA only, the results for the FEs under protons and thermal neutrons are presented, since an higher sensitivity is expected because of its newer technology node. Finally, its performances against TID effects are analyzed.

II. CERN RADIATION ENVIRONMENT

The radiation environment inside the CERN accelerators is composed of different particles over a large spectrum of energies, from meV up to GeV, whose distribution may vary significantly depending on the location. When exposed to this sort of environment, FGPAs' performances can be impacted by both TID and SEEs. The two main contributions to SEEs inside the LHC are coming from High Energy Hadrons (HEH) and Thermal neutrons (ThN). HEH are defined as all hadrons with kinetic energy above 20 MeV. The different areas inside the LHC and the corresponding radiation levels of the upcoming HL-LHC are analyzed in detail in [12]. The LHC areas where the COTS component can be installed can be divided into two main groups, according to their radiation levels: tunnel areas and shielded areas. The shielded areas, like UJs (junction chamber) and ULs (liaison gallery), located near the Interaction Points (IP), are heavily shielded, whereas RRs, located between the IRs (insertion regions), and Dispersion Suppressors (DS), are lightly shielded. In the heavily shielded areas, the thermal neutron contribution is much higher, because high-energy neutrons are thermalized by the shielding, whereas high-energy particles are attenuated. Different studies, like [2], [13], and [14], show that the thermal neutron contribution can significantly affect the sensitivity of recent technologies. Therefore, a Risk factor (R_{th}) [15] was introduced to identify locations with higher thermal neutrons contribution. The aforementioned risk factor is defined as

$$R_{th} = \frac{\Phi_{ThN}}{\Phi_{HEH}},$$

where Φ_{ThN} is the thermal neutron fluence and Φ_{HEH} is the high-energy hadron fluence. Table I summarizes the expected fluences with the HL-LHC upgrade for these areas obtained by FLUKA simulation [12], and the average R_{th} derived from measurements in [14]. In the tunnel the risk factor is around 4, but it can go up to 25 in the ULs, and up to 50 in the UJs. Therefore, thermal neutrons can have a non-negligible contribution to the total failure rate. This adds more complexity to the RHA procedure [16], making thermal neutron tests mandatory for FPGAs with smaller node technologies.

To address the devices' sensitivity to HEH and ThN individually, the FPGAs are tested in two different facilities. Concerning the HEH, according to the RHA procedure, they all have the same probability of inducing SEEs due to their similar nuclear interaction cross-section. Therefore, the devices are tested at the Paul Scherrer Institute (PSI) in Villigen, Switzerland, under a high-energy proton beam of 200 MeV. The ThN tests are carried out at the TENIS beam line, at the Institute Laue-Langevin (ILL) [17] in Grenoble, France.

TABLE I
EXPECTED ANNUAL RADIATION LEVELS IN THE LHC AREAS FOR THE HL-LHC UPGRADE AND AVERAGE MEASURED RISK-FACTOR

Location	TID (Gy)	HEH Fluence (p/cm^2)	R_{th}
Shielded areas			
UJ	10	$5 \cdot 10^9$	50
UL	0.2	10^8	25
RR	6	$3 \cdot 10^9$	6
Tunnel			
DS	100	$5 \cdot 10^{10}$	4
ARC	2	10^9	3

III. TEST SETUP AND METHODOLOGY

Different radiation campaigns were performed in order to investigate the robustness of different FPGAs. The NG-Medium and the PolarFire have been proposed as candidates for future CERN applications, considering the TID levels expected for the HL-LHC. The former, despite its cost not comparable to that of a fully commercial solution, is a RHBD FPGA, therefore it might be the only solution for systems in very harsh environments, especially in terms of TID. The PolarFire instead, belongs to the same family of the SmartFusion2 and ProASIC3 which had already been qualified for the LHC environments. Thus, similar/better performances are expected considering the addition of the SONOS technology.

To carry out the radiation campaigns, a setup comprising two different systems was adopted; these are the FPGA Under Test, acting as Device Under Test (DUT), and a second FPGA, acting as a tester. By means of an additional FPGA, the circuitry necessary to transfer the data from the DUT to a host computer is removed, e.g. logic necessary to implement a UART/Ethernet peripheral or a Built-in Self-Test (BIST) architecture. Thus, the circuits under test are monitored directly. This allows performing a better analysis of the results because the testing logic on the DUT is minimized, reducing the potential sources of errors. Fig. 1 illustrates a top-level view of the test setup. The Zynq-7000 SoC was adopted. It contains an ARM Cortex-A9 CPU connected to an Artix-7 FPGA. The FPGA implements the necessary test routine, sends the relevant data to the ARM CPU, and an external computer can be used to monitor the test. The tester and the DUT are connected through an Low Pin Count (LPC) FMC cable. Therefore, the development board hosting the DUT must be equipped with an FMC connector. This is the case for the NG-Medium and the PolarFire. However, for the ProASIC3 and SmartFusion2, whose development boards lack an FMC connector, a UART interface was integrated on the DUT to communicate with the external computer.

Implementing the benchmark structure on all the FPGAs allowed a better comparison between multiple devices belong-

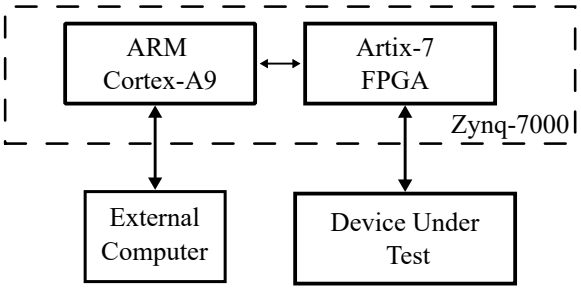

Fig. 1. Overview of the test setup adopted for the characterization of the Device Under Test.

ing to different families. Additionally, the SEE sensitivity of the PolarFire FEs was studied, to investigate the possibility of estimating the benchmark circuit sensitivity starting from these results. Propagation delay tests were also performed to quantify the design lifetime with the dose. These tests were conducted only on the PolarFire since they had already been performed on the other FPGAs in previous works.

A. Standard methodology

Following the already established guidelines detailed in [5], the sensitivity of the FEs was retrieved through dedicated circuits.

Concerning the Flip-Flops, multiple chains of Windowed-Shift-Register (WSR) were implemented. They are shift-registers with a serial-to-parallel output window, that captures the output of the last N Flip-Flops. This structure was chosen because it allows performing tests at higher speed, reducing signal integrity issues at the FPGA output. Thanks to this structure, in normal conditions the output window always contains the same values. Only in the case of an error, the content of the output window is different. The same structure was used to test the Single-Event Transient (SET) capture sensitivity, by adding combinatorial gates between each Flip-Flop of the chain (WSR-SET). Both topologies of chain were implemented with and without Triple-Module Redundancy (TMR) to test the efficiency of this mitigation technique. Moreover, as many chains as possible were placed on the DUT to retrieve a good amount of statistics. However, in order to keep the number of required I/Os low, a comparator mechanism for each chain, that checks the correctness of the data present in the window, was included. The output of this comparator was monitored by the tester architecture. Additionally, two test modes were included: static and dynamic. In static mode, the bit shifted inside the chain is set to either 0 or 1. In dynamic mode, the input alternates between the two. Fig. 2 shows the overview of the logic circuit used to retrieve the Flip-Flop sensitivity.

DSPs can be usually configured to implement various operations. In the case of the PolarFire, the possible configurations are: Multiplier only (MULT), Adder (ADD) and Multiply and Accumulate (MAC). To retrieve all of their sensitivities, a circuit containing all the configurations was implemented. In order to increase the amount of statistics measured, while keeping the number of I/Os within the limit of the FMC

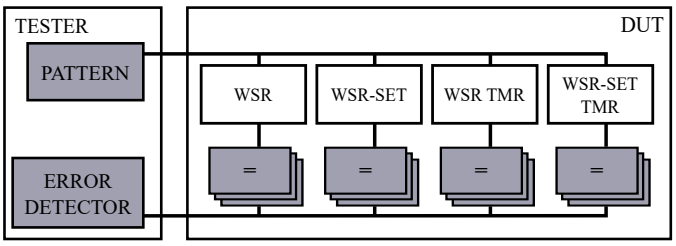

Fig. 2. Top-level architecture of the logic circuit used to retrieve the Flip-Flop sensitivity. The tester provides the inputs to the various structures on the DUT, and monitors their outputs.

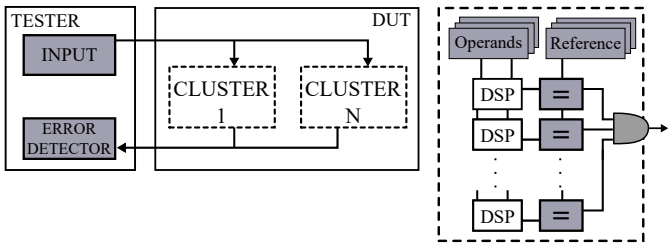

Fig. 3. Top-level architecture (left) of the circuit used to characterize the DSPs sensitivity. Internal structure (right) of a cluster of DSP.

connector, the DSPs are organized in clusters. Each cluster contains DSPs performing the same operation. All the clusters are fed with the same input, stored on the DUT in triplicated registers. Thus, all the DSPs in the same cluster are expected to produce the same output. A comparator, for each DSP, checks that the output is correct. Finally, all the comparators of a cluster are connected through an AND gate to the DUT output that is monitored by the tester. This way, an error is detected as soon as one DSP in the cluster is affected by a SEE. The test can be performed either in static or dynamic mode: in static mode, the inputs are fixed and the DSPs perform the same operation starting from the same input. In dynamic mode, the input alternates between two values. The overview of this architecture is shown in Fig. 3.

As far as clocking circuitry is concerned, the PolarFire comes with eight Phase-Locked Loops (PLL). Their sensitivity was analyzed by monitoring their LOCK signal using the tester. For each PLL, the tester registers an event every time the PLL is not locked.

B. Benchmark test methodology

Evaluating the radiation response of an FPGA is a challenging task because there are many aspects to take into consideration. Starting from a given RTL description of a circuit, its final implementation will depend on the target FPGA, the tool used, and the applied constraints. Moreover, starting from the RTL model, there are many steps automatically performed by the tool that are configurable by the user. All these aspects can lead to a different implementation of the same RTL circuit, with different area, power consumption, performance characteristics, and eventually a different radiation response because the resulting number of configuration bits is different, or because an SET may propagate differently depending on the resource used. However, when choosing the FPGA candidate

for a particular application at CERN, there is no defined implementation strategy, since it depends on each application. Most of the times, optimizations are carried out only at RTL level and through primitives available from the tool, i.e. to apply the TMR mitigation scheme. In addition to that, when using different tools, the available implementation strategies may not be the same, which would make the comparison harder. Moreover, if on one hand applying a different strategies can improve the circuit response, on the other hand it can mask some of the failure modes that could arise during operation when such techniques are not applied by the FPGA developers. Thus, since the goal is to perform an application-level test, testing all the FPGAs using the same implementation strategies was important. Because of all these reasons, it was decided to study the behaviour of each device running the same benchmark application constraining only the clock frequency, the IOs, and applying or not the TMR mitigation technique, leaving everything else to the vendor tools.

B.1 Benchmark presentation

The adopted benchmark belongs to the ITC'99 suite, in particular to the IT99 portion developed by Politecnico di Torino. The goal of the suite is to provide a set of RTL circuits with different test cases, different complexity, but uniform characteristics. They were built starting from public VHDL files, modified and combined to obtain larger circuits. Following this process, the circuits might have lost their original functionalities in favour of uniformity of description. The result of this process is a set of fully synthesizable circuits, without any hardware/compiler-specific directive that allows their implementation on different families of FPGAs. The VHDL descriptions range from tiny to larger circuits which implement a variety of functionalities, from Finite-State Machines (FSM) to soft-core microprocessors. Therefore, they are a good representation of how FPGAs are used in a real environment. In this work, the B13 was chosen. Its original function was to act as an interface with a weather sensor. The circuit occupies relatively few resources (339 gates, 53 FFs, 20 I/O). It is mainly based on FSMs, which are quite common in real applications. Even though it is quite small, it was selected because it had already been used in other reliability experiments on other FPGAs. Hence, it was possible to compare our results with the ones obtained in the other measurements. Moreover, tests were planned for FPGAs of different sizes, with the aim of collecting many statistics from the irradiation tests. Thus, the B13 represented a good choice given its small size, allowing one to replicate it as much as possible on the DUT.

B.2 Benchmark setup

When choosing the DUT implementation, a trade-off between error observability and complexity of test structure is necessary, depending on the target FPGA. Each B13 has 10 outputs, therefore it is not possible to directly monitor each instance using the tester because of the limited amount of pins. Thus, a comparison on the DUT was adopted, and additional test logic was implemented on the tester in order to keep a good level of observability. The input generator and the golden reference were instead moved to the tester. During the test, the input generator is used to feed the golden circuit on the tester and all the B13s on the DUT. The outputs of the B13s on the DUT are compared against the output of the golden B13. In case of an error, the test logic raises an error signal, generates the identifier of the B13 affected by the error and its output. All this information is sent to tester, to help understanding in which structure the error occurred. Fig. 4 shows such an architecture. A second version of this design with TMR, was also tested. In both versions, TMR is implemented on all the comparators and the failure detection logic to make them more robust against SEEs. This architecture was used for both the NG-Medium and the PolarFire. For the SmartFusion2 and ProASIC3 without the FMC connector, another solution was necessary because of the limited number of pins available on the test board. The error observability was sacrificed in favour of a less complex DUT design. The B13's were compared two-by-two, and an UART interface was used to notify the errors to the external computer for data analysis. In this case, the tester was not employed. Even for this design, TMR was applied on all the logic surrounding the B13s.

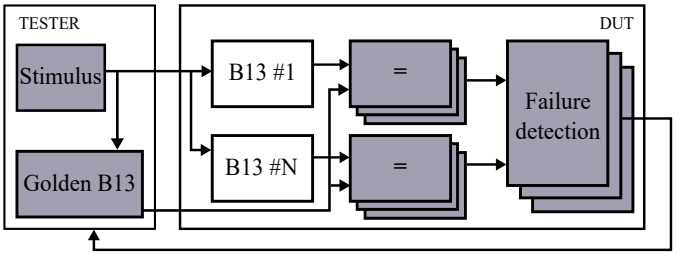

Fig. 4. Top-level architecture of the circuit used to study the response of the benchmark application.

C. Propagation delay monitoring

While exposed to radiations, the propagation delay of the FPGA elements can increase [18], [19]. Since every logic circuit works at a specific frequency, if the propagation delay changes while under radiation, the design could fail even before the device breakdown. As experimented in [20], the propagation delay degradation reached values of up to 1100% before the device failed. This kind of event, is the main source of failure related to TID for an user design. Therefore, the propagation delay increase was measured by monitoring the frequency of many ring oscillators. A ring oscillator is simply a loop of inverters. For an odd number of inverters N, each implemented with a LUT whose propagation delay is t_{pd}, the output of the last inverter will oscillate with a frequency f expressed by:

$$f = \frac{1}{t_{pd} \times N \times 2}.$$

Fig. 5 shows the circuit used to monitor the propagation delay. The TID effects on the DUT were analyzed by measuring the change of frequency of each ring oscillator during irradiation using the tester.

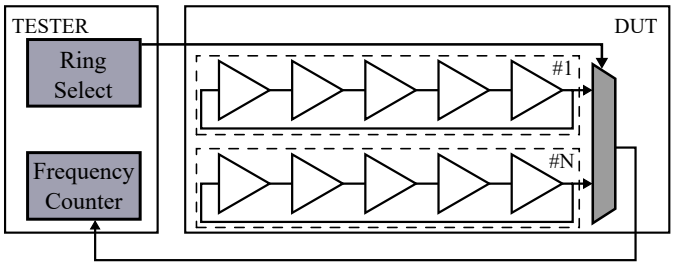

Fig. 5. Circuit for measuring the propagation delay change in the DUT. The tester selects the ring oscillators to observe, and a frequency counter on the tester FPGA measures its output frequency.

TABLE II
NUMBER OF CIRCUITS AND ELEMENTS USED FOR THE CHARACTERIZATION OF THE POLARFIRE MPF300

Structure	# Replica	# Element
WSR	4	8000
WSR-TMR	2	8000
WSR-SET	4	8000
WSR-SET-TMR	2	8000
DSPs (All configuration)	-	80
PLLs	-	8

IV. RADIATION TESTS AND RESULTS

This section analyzes the results obtained from the different radiation campaigns. For all the FPGAs, the benchmark circuit was used to retrieve the devices' response at application level and to estimate their failure rate during operation. For the PolarFire only, the functional elements and its performances against TID were evaluated too.

A. PolarFire functional elements

The sensitivity of Flip-Flops, DSPs, and PLLs of the PolarFire was studied for protons and thermal neutrons irradiations. The Flip-Flops were tested using the WSR and WSR-SET structure implemented with and without TMR in dynamic mode. Multiple tests at multiple frequencies were performed, with the aim of understanding the possible relation between the frequencies and SEE cross-sections. On the same DUT, also DSPs and PLLs were placed, together with the FFs. The DSPs were tested at different frequencies and in dynamic mode. The PLLs were all fed by the same clock, and they were all using the same configuration. During the irradiation, in addition to SEUs, global failures were observed. Most of the circuits on the DUT (FFs, DSPs and PLLs) were failing continuously at the same time. These events, considered to be Single Event Functional Interrupts (SEFI), were treated as a separate category. Table II describes the number of replicas used for each topology of element. Fig. 6 shows the cross-section for SEUs and SEFIs for the different elements. It must be noted that no SEU was observed on the PLLs. For each cross-section, the 95 % confidence interval calculated using the methodology presented in [21] is reported.

Several considerations can be made about these results. First, despite the smaller technology, the PolarFire shows an average FF cross-section 2.5x lower than its predecessor, the

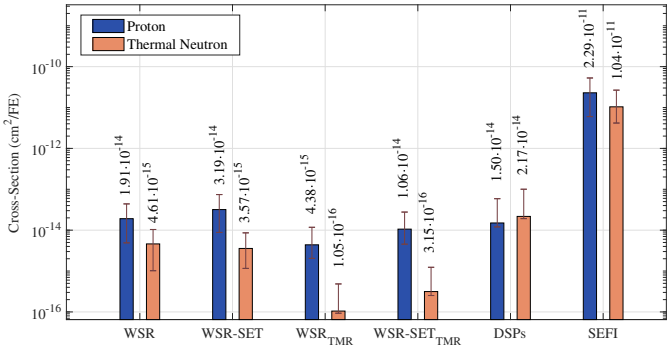

Fig. 6. SEUs and SEFIs cross-section under protons and thermal neutrons for the PolarFire MPF300 elements.

SmartFusion2 [1]. The difference is even higher for the DSPs, with an average cross-section of two orders of magnitude lower than its predecessor.

Then, a lower cross-section in the TMR version of the WSR chain can be observed, both for protons and thermal neutrons, which proves the efficiency of the mitigation technique.

However, global design failures (SEFIs) were observed and their cross-section is relatively high; they represent the dominant failure type. These SEFIs can originate because of SETs on a global route, which can be either the reset or the clock. However, not all the structures were failing, probably because the SETs were attenuated before reaching the structures placed far from the affected location.

Furthermore, a very interesting outcome from these tests is the cross-section measured for the thermal neutrons. For every metric, except for FFs with TMR, the protons and thermal neutrons cross-sections are relatively comparable. The cause for this high cross-section could be related to the presence of boron-10 in the device, since this element has a very high thermal neutron absorption. Considering the TMR version instead, the thermal neutron cross-section is lower. One possible reason could be the different SET duration. It is possible that the SETs induced by thermal neutrons are shorter than those induced by protons, thus they are better mitigated by TMR. However, the SET duration should be measured with techniques such as those presented in [22] to verify such assumption. This confirms that thermal neutrons tests are necessary because they can have a huge impact on the failure rate of these devices in operation. Fig. 7 instead, shows the FFs cross-section for SEUs as a function of the frequency. From the results it is visible that the cross-section is stable with the frequency, except for the WSR version in TMR mode.

B. PolarFire MPF300 propagation delay degradation

The degradation of the propagation delay was monitored using 1952 ring oscillators, each of them containing 47 inverters. With such a number, considering the LUT delay and the routing path, each of them generates a 100 MHz signal. The structures were spread through the entire FPGA to investigate if some areas could be affected more than others. Moreover, the ring oscillators were placed manually. For each of them, all

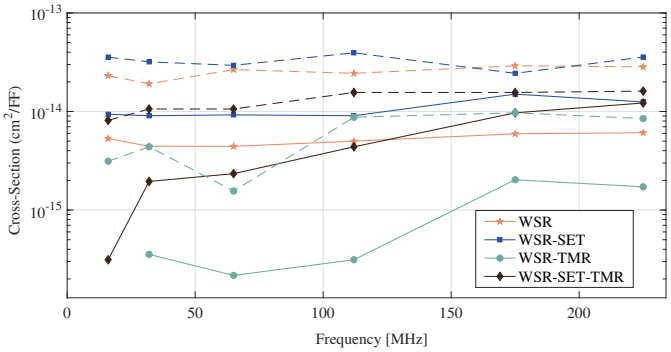

Fig. 7. PolarFire Flip-Flops SEE cross-section expressed as a function of the frequency. The filled lines refer to the static test, whereas the dashed lines to the dynamic test.

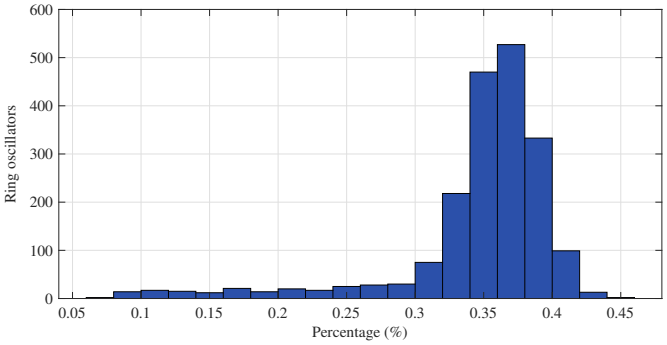

Fig. 8. Propagation delay difference before and after irradiation of the PolarFire MPF300 for a total cumulative dose of 5.2 kGy.

the inverters were placed next to each other, avoiding the extra logic that could be added by the placement tool, resulting in an identical structure for all the rings. The rings were monitored during the whole irradiation run using the tester. Fig. 8 shows the difference, in percentage, between the frequency measured before and after the irradiation. As it can be seen, the FPGA exhibited a very good behaviour. Most of the ring oscillators frequency changed of only 0.3%, and the maximum observed variation is of 0.45%, which is a very good result considering that the level of dose absorbed reached 5.2 kGy.

C. Benchmark application

Before presenting the results for the benchmark application, it is necessary to analyze the different events and failure types observed.

C.1 NG-Medium

It is an SRAM-based FPGA whose cells are hardened by design, therefore they are more resilient to single events compared to those of the other FPGAs. However, its configuration memory (CRAM) is based on an SRAM architecture which is more sensitive to SEU compared to Flash-based FPGAs. However, the NG-Medium is equipped with a Configuration Memory Integrity Check (CMIC). It is an embedded engine performing automatic verification and repair of single-bit error inside the CRAM. When the bitstream is generated, a CMIC reference is automatically added by the tool. During the

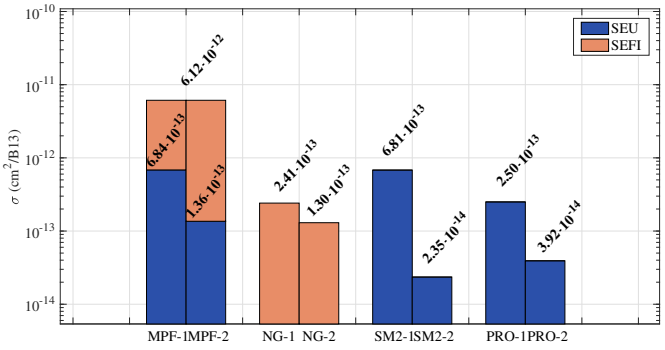

Fig. 9. SEU and SEFI cross-section of the benchmark application tested on NG-Medium, PolarFire, SmartFusion2 and ProASIC3, with version #2 implemented with TMR.

download process, the bitstream is loaded into the CRAM, while the CMIC reference is loaded into its own on-chip memory, protected by ECC. The engine periodically scans the CRAM, and in case of a single-bit mismatch in a word, it corrects the error. According to the datasheet, this process takes around 4 ms. In the case of a double-bit error, however, the engine stops working and further single bit errors will not be corrected anymore. Without scrubbing, an error in the CRAM would change the interconnect configuration of the FPGA, leading to a permanent failure of the design. Thanks to the CMIC instead, a SEU in the CRAM leads to a temporary failure only, until the repair is finished. While the CRAM is corrupted, the entire design is in a faulty state, and this event will be referred to as a Single Event Functional Interrupt (SEFI) from now on.

Another type of failure observed is related to the FPGA itself. During the irradiation, the NG-Medium suffered from radiation-induced resets causing a reload of the configuration memory. In our test setup, the configuration was stored inside an external memory. Hence, when a reset occurred, the FPGA could reload it and resume operation. However, this may not always be the case when in operation, and the CRAM could need to be manually reloaded again. For this reason, this kind of event is referred to as permanent failure.

C.2 PolarFire, SmartFusion2, ProASIC3

They are Flash-based FPGAs, therefore their CRAM is more resilient to SEUs, but since they are not radiation hardened, their cells are more sensitive. In this case, the source of failures for the application is due to SETs or SEUs inside the FPGA logic and the design functionalities are affected only temporarily, so these events are referred to as temporary failures.

C.3 Results analysis

Fig. 9 shows the response of the four FPGAs under high-energy protons. Table III reports the number of B13 replicas considered for each FPGA, the fluences and the errors detected. Results were gathered across multiple campaigns, thus the table contains cumulative values.

Starting with the NG-Medium, it was not planned to use any sort of mitigation technique in operation given the Rad-Hard nature of the FPGA. However, an unexpectedly high

TABLE III
CIRCUITS EMPLOYED, FLUENCES AND ERRORS DETECTED FOR EACH DESIGN ON THE DIFFERENT FPGAS

FPGA	Design Name	TMR	B13s instance	Fluence (p/cm^2)	Errors detected
PolarFire	MPF-1	N	2048	$3.00 \cdot 10^{10}$	42
	MPF-2	Y	2048	$2.3 \cdot 10^{11}$	64
NG-Medium	NG-1	N	184	$7.90 \cdot 10^{11}$	35
	NG-2	Y	32	$2.65 \cdot 10^{12}$	11
SmartFusion2	SM2-1	N	160	$1.10 \cdot 10^{11}$	12
	SM2-2	Y	50	$8.50 \cdot 10^{11}$	1
ProASIC3	PRO-1	N	180	$2.00 \cdot 10^{11}$	9
	PRO-2	Y	60	$8.50 \cdot 10^{11}$	2

TABLE IV
B13 SEU CROSS-SECTION FOR THE POLARFIRE MPF300 UNDER THERMAL NEUTRONS.

B13s instance	Fluence (p/cm^2)	SEU	
		# Events	σ (cm^2/B13)
Thermal Neutron			
2048	$5.49 \cdot 10^{11}$	82	$7.29 \cdot 10^{-14}$
1024	$1.21 \cdot 10^{12}$	7	$5.65 \cdot 10^{-15}$

number of CMIC corrections was observed. For this reason, the TMR-ed version of the benchmark was tested as well. Since the placement tool ('NanoXmap 2.7') provided by the manufacturer was not offering any automatic triplication routine, the design was triplicated manually at block level. The results show that Block-TMR reduced only slightly the problem. Further analysis demonstrated that this is due to the placing/routing tool, which is creating many long common paths in the design, i.e the one between the output of the circuits and their corresponding voters, also referred in the literature as Single Point of Failure (SPF) which reduce the effectiveness of the triplication technique. At the moment of the writing of this paper, no possibility was given to manually place the different elements of the design to apply known placement techniques [23] to mitigate this problem. Therefore, a quite high margin of improvement can be expected with the planned update of the FPGA development tools by NanoXplore. Concerning the permanent failure, the cross-section for this event has been measured as $1.10 \cdot 10^{-12}$ cm^2/device. Following a collaboration with the manufacturer, it emerged that one of the possible sources of this reset could be an error affecting the peripheral available in the FPGA as hard-coded block, i.e the SpaceWire interface. This peripheral, if affected by an error, triggers a reset and causes the reloading of the configuration memory even though it is not used by the user logic. Another reason could be the failure of the CMIC engine. When a double error is detected, the FPGA is reset and the CRAM is reloaded, so that the CMIC operation can restart. Thus, tests were performed again after these reset conditions were masked. Nonetheless, the errors still occurred, indicating that there might be other sources, such as a micro latch-up, that power resets the FPGA. Further investigation on this problem is necessary.

Concerning the PolarFire, some temporary failures due to SEUs inside the B13 circuits were observed, but also some SEFIs where many B13s were failing, as in the functional element tests. As visible in Fig. 9, even though the use of triplication reduced the SEU sensitivity by a factor of 5, the total failure rate is dominated in both cases (mitigated and non-mitigated), by the SEFIs. Further investigation is necessary to mitigate this effect for this FPGA.

On the other hand, no SEFI was observed on the SmartFusion2 and ProASIC3. The TMR version of the design also shows a lower sensitivity compared to the PolarFire, which is reasonable since they are based on a larger technology, and so they are less sensitive.

C.4 Comparison

Comparing the total cross-sections for all the FPGAs, it is clear that the NG-Medium exhibits a better behavior among them all. However, mitigating the SEFIs on the PolarFire would lower down the cross-section at the same level as the NG-Medium, which is remarkable when comparing a commercial FPGA with an RHBD FPGA. Moreover, the SEU cross-section for the Flash-based FPGAs are comparable for the non-TMR version, but the SmartFusion2 and ProASIC3 show better performances with the TMR version compared to the PolarFire.

It is important to note that this conclusion would have been really hard to draw by using only the FEs test, for different reasons. The PolarFire failure rate is a combination of SEUs and SEFIs, and it is different from the SEFI cross-section retrieved at FE level. Moreover, for the NG-Medium, it is impossible to estimate the impact of the CMIC operations at circuit level using the CRAM sensitivity, since the number of critical bits given by the tool is inaccurate as it was not possible to retrieve the same results based on their estimation. This demonstrates why the FEs test is not sufficient for the FPGA qualification, and it also shows how benchmarks allow for comparison of different types of technologies subject to different kinds of effects, where the sources of failures are different.

For the PolarFire, SmartFusion2, and ProASIC3, the benchmark application was tested also under thermal neutrons. As expected, no event was observed for the SmartFusion2 and ProASIC3, in contrast with what was observed on the PolarFire, which is based on a newer technology. Table IV reports its cross-section, that is 10x lower compared to the one observed under protons. Surprisingly, no SEFI was observed during the thermal neutron tests, even though the design was the same for both the tests. Moreover, two different PolarFire FPGAs were used for the functional element tests under thermal neutrons, and in both cases SEFIs were observed.

D. FPGA lifetime

The robustness of the FPGAs against TID effects was also investigated during the tests, in terms of lifetime and programmability. For the lifetime, three events are considered as the end of life of the device: an exponential propagation delay increase, a sudden failure of the device, or the corruption of the flash memory. However, only the first event was observed during the tests.

No degradation or failures were observed with the NanoXplore up to 3 kGy [3], but it was more surprising that

similar performances were reached with the PolarFire. The same evaluation board was used for three different campaigns separated by 6 and 5 months respectively, where the doses reached 500 Gy, 2 kGy, and 3 kGy respectively. Nevertheless, the FPGA did not show any significant sign of degradation, neither in terms of current consumption nor propagation delay. Also, no loss of re-programmability was observed after the second campaign (2.5 kGy), before being unable to program the board at the end of the third campaign (5.5 kGy). These are remarkable great results for a commercial FPGA, also when compared to the low lifetime and programmability threshold of its predecessors, the SmartFusion2 and the ProASIC3. The SmartFusion2 lost programmability features after only 70 Gy, and survived up to 650 Gy. The ProASIC3 instead lost programmability after only 20 Gy, whereas its lifetime was of 540 Gy. In these two cases, the lifetime corresponds to an exponential increase in the propagation delay.

The NG-Medium instead had a shorter effective lifetime because of two issues. The CMIC engine, as mentioned before, stops operating when it detects a double error. Thus, single-bit errors start accumulating until the design fails, without the possibility to recover without reloading the CRAM. Moreover, radiation-induced FPGA resets were observed, causing the loss of the CRAM content and thus, the permanent failure of the design. As described before, during the test the FPGA could reload the design from an external memory and resume operation. However, the radiation levels of the LHC areas where this FPGA will be used are high enough to cause the failure of flash memories. Therefore, in that situation, this kind of event would cause the failure of the system requiring a manual reload of the user design. To mitigate these problems, a solution could be the remote programming of the FPGA trough JTAG chains, just as it is already done for some FPGAs in the CERN detectors, for instance ATLAS. However, deploying the same solution in the whole accelerator part of the complex will significantly increase the cost of this solution. More investigations are necessary to understand the origins of these resets and possible mitigation techniques.

V. Failure Rate Estimation

By using the failure rate of the benchmark circuit, the failure rate of FPGAs belonging to different families can be compared. As previously done in [2], but considering two SEE contributions (thermal neutrons and high-energy hadrons), the Homogeneous Poisson process (HPP) is used to estimate the probability of having x failures over a period of time t. The process considers a constant failure rate and that the failures are independent from each other. It is expressed as:

$$f(x;\lambda_T;t) = \frac{(\lambda_T t)^x e^{-\lambda_T t}}{x!},$$

where λ_T is the combination of the thermal neutrons and high-energy hadrons contributions, and it is defined as

$$\lambda_T = \lambda_{HEH} + \lambda_{ThN} = \sigma_{HEH}\Phi_{HEH} + \sigma_{ThN}\Phi_{ThN}.$$

Φ_{HEH} and Φ_{ThN} are the HEH and ThN annual fluences respectively, whereas σ_{HEH} and σ_{ThN} are the HEH and ThN cross-sections. For the hadron, the proton cross-section was used as explained in section II.

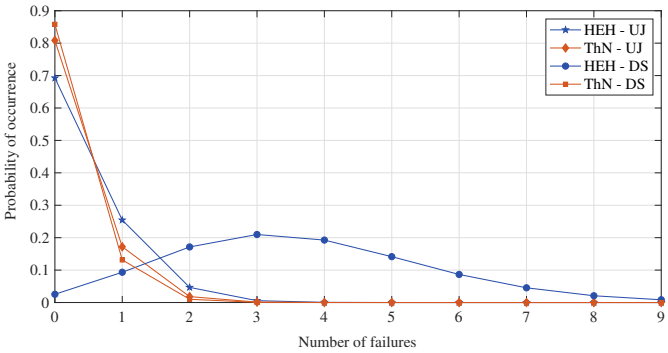

Fig. 10. System failures probability for the PolarFire FPGA considering 12 years of operation in the UJ and DS, calculated using the annual fluence level of the HL-LHC, for both HEH and ThN.

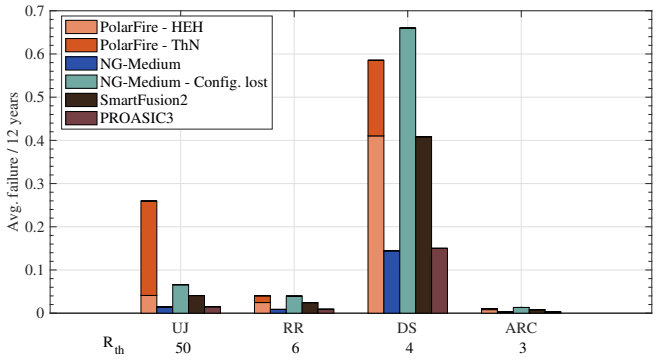

Fig. 11. Average failures for all the FPGAs in the four areas of the LHC, considering 12 years of operation

As an example, Fig. 10 shows the HPP curves calculated considering the PolarFire working for 12 years inside the DS and UJ, with two different risk-factors. The thermal neutrons and protons contributions in this case are separated. As it can be seen, in the DS where the HEH fluence is higher, the PolarFire has 70% chance of failing three or more times. On the other hand, there are less chances of failing because of ThN. In UJ instead, where the risk-factor is higher, there is a 20% probability of observing one or more failure induced by ThN. Fig. 11 shows the average number of failures, calculated from the HPP equation, for the non-TMR benchmark on all the FPGAs in four main LHC areas. The NG-Medium has very good performances compared to the others in terms of SEEs. However, the loss of configuration represents a problem, especially in the DS. Considering the PolarFire instead, the impact of ThN in the different areas can be appreciated. In the area with a low risk factor, such as the DS, 42% of the failures observed are induced by thermal neutrons. With the increase of the risk factor, the thermal neutron contribution increases, until it becomes the predominant failure source. In the UJ, 81% of the failures are induced by ThN. It must be noted that Fig. 10 and Fig. 11 show the failures for a single FPGA, thus the numbers seem low, but there will be hundreds of devices installed inside the LHC tunnel.

Thus, it is clear that the PolarFire represents a better

solution in areas with a low R_{th} factor, since its lifetime and programmability are much higher compared to those of the SmartFusion2 or ProASIC3. However, the situation is reversed when considering areas with a high R_{th} factor. Here, the TID levels are very low, therefore the higher sensitivity to thermal neutrons make the PolarFire a bad candidate compared to the SmartFusion2 and ProASIC3, since they are less sensitive to ThN and can still survive for 12 years. Nonetheless, the racks hosting the FPGAs could be shielded to reduce thermal neutron fluxes, mitigating this problem. Therefore, the correct FPGA must be chosen mainly depending on the LHC area, but also according to the failures modes tolerated for the application and to whether shielding is possible or not.

VI. Conclusions

These experiments allowed to obtain a suitable characterization of the PolarFire and the NG-Medium. For the PolarFire, a complete characterization has been performed in terms of SEE and TID. The work presented an approach for FPGA testing, based on benchmark circuits. The results obtained proved the efficiency and the advantages of this technique compared to standard SEE testing. A benchmark failure rate, different from the one that can be obtained by analyzing the FEs by themselves, was derived. Moreover, by testing the same application on four different FPGAs, it was possible to demonstrate how the methodology allows for comparison of different FPGAs belonging to different families and affected by different errors. The PolarFire showed a relatively high sensitivity to thermal neutrons, most probably due to the presence of boron-10 in the device, therefore tests under these particles are necessary for devices based on recent technologies. Using the experimental data, in conjunction with the classical reliability analysis, the failure rates of the various FPGAs in the different areas of the LHC were estimated. The results gathered from such estimation proved that the NG-Medium and the PolarFire are two interesting candidates for CERN applications. Results also confirmed the importance of assessing thermal neutrons sensitivity, since they represent the major cause of SEEs in the shielded areas. Future works will focus on a deeper study of the SEFIs occurred in the PolarFire, and will also expand the benchmark tests using other benchmark applications.

References

[1] G. Tsiligiannis, R. Ferraro, S. Danzeca, A. Masi, M. Brugger, and F. Saigné, "Investigation on the sensitivity of a 65nm Flash-based FPGA for CERN applications," in *16th European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Paper No. DW11, Sep. 2016.

[2] G. Tsiligiannis, S. Danzeca, R. García Alía, A. Infantino, A. Lesea, M. Brugger, A. Masi, S. Gilardoni, and F. Saigné, "Radiation Effects on Deep Submicrometer SRAM-Based FPGAs Under the CERN Mixed-Field Radiation Environment," *IEEE Transactions on Nuclear Science*, vol. 65, no. 8, pp. 1511–1518, Aug. 2018.

[3] G. Tsiligiannis, C. Debarge, J. Le Mauff, A. Masi, and S. Danzeca, "Reliability analysis of a 65nm Rad-Hard SRAM-Based FPGA for CERN applications," in *19th European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Sep. 2019.

[4] N. Rezzak, J.-J. Wang, S. Varela, G. Bakker, and A. N. Gu, "Neutron and Proton Characterization of Microsemi 28 nm PolarFire SONOS-Based FPGA," in *IEEE Radiation Effects Data Workshop (REDW)*, 2018, pp. 210–215.

[5] M. Berg, "Field programmable gate array (FPGA) single event effect (SEE) radiation testing," NASA/Goddard Space Flight Center, USA, Tech. Rep., 2012. [Online]. Available: https://nepp.nasa.gov/files/23779/FPGA_Radiation_Test_Guidelines_2012.pdf

[6] J. J. Wang, N. Rezzak, F. Hawley, G. Bakker, J. McCollum, and E. Hamdy, "Radiation characteristics of field programmable gate array using complementary-sonos configuration cell," Microchip, San Jose, CA, USA, Tech. Rep., 2019. [Online]. Available: https://www.microsemi.com/document-portal/doc_view/1244474-rt-polarfire-radiation-test-report

[7] M. Berg, "SEU System Analysis: Not Just the Sum of All Parts," presented at the Single Event Effects (SEE) Symp. and the Military and Aerospace Programmable Logic Devices (MAPLD) Workshop, La Jolla, CA, USA, 2014. [Online]. Available: https://ntrs.nasa.gov/citations/20140008977

[8] M. Berg, "SEE Test and Data Analysis for Complex FPGA Systems," presented at the Microelectronics Reliability & Qualification Working Meeting (MRQW), El Segundo, CA, USA, 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20200000820

[9] F. Corno, M. Reorda, and G. Squillero, "RT-level ITC'99 benchmarks and first ATPG results," *IEEE Design Test of Computers*, vol. 17, no. 3, pp. 44–53, 2000.

[10] H. Quinn, W. H. Robinson, P. Rech, M. Aguirre, A. Barnard, M. Desogus, L. Entrena, M. Garcia-Valderas, S. M. Guertin, D. Kaeli, F. L. Kastensmidt, B. T. Kiddie, A. Sanchez-Clemente, M. S. Reorda, L. Sterpone, and M. Wirthlin, "Using Benchmarks for Radiation Testing of Microprocessors and FPGAs," *IEEE Transactions on Nuclear Science*, vol. 62, no. 6, pp. 2547–2554, Dec. 2015.

[11] A. M. Keller, T. A. Whiting, K. B. Sawyer, and M. J. Wirthlin, "Dynamic SEU Sensitivity of Designs on Two 28-nm SRAM-Based FPGA Architectures," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 280–287, Jan. 2018.

[12] R. García Alía, M. Brugger, F. Cerutti, S. Danzeca, A. Ferrari, S. Gilardoni, Y. Kadi, M. Kastriotou, A. Lechner, C. Martinella, O. Stein, Y. Thurel, A. Tsinganis, and S. Uznanski, "LHC and HL-LHC: Present and Future Radiation Environment in the High-Luminosity Collision Points and RHA Implications," *IEEE Transactions on Nuclear Science*, vol. 65, no. 1, pp. 448–456, Jan 2018.

[13] S. Wen, R. Wong, M. Romain, and N. Tam, "Thermal neutron soft error rate for SRAMS in the 90NM–45NM technology range," in *IEEE International Reliability Physics Symposium*, 2010, pp. 1036–1039.

[14] M. Cecchetto, R. García Alía, F. Wrobel, M. Tali, O. Stein, G. Lerner, K. Bilko, L. Esposito, C. Bahamonde Castro, Y. Kadi, S. Danzeca, M. Brucoli, C. Cazzaniga, M. Bagatin, S. Gerardin, and A. Paccagnella, "Thermal Neutron-Induced SEUs in the LHC Accelerator Environment," *IEEE Transactions on Nuclear Science*, vol. 67, no. 7, pp. 1412–1420, Jul. 2020.

[15] D. Kramer, M. Brugger, V. Klupak, C. Pignard, K. Roeed, G. Spiezia, L. Viererbl, and T. Wijnands, "LHC RadMon SRAM Detectors Used at Different Voltages to Determine the Thermal Neutron to High Energy Hadron Fluence Ratio," *IEEE Transactions on Nuclear Science*, vol. 58, no. 3, pp. 1117–1122, Jun. 2011.

[16] K. Roeed, M. Brugger, and G. Spiezia, "An overview of the radiation environment at the LHC in light of R2E irradiation test activities," CERN, Tech. Rep., 2011. [Online]. Available: https://cds.cern.ch/record/1382083

[17] J. Beaucour, J. Segura-Ruiz, B. Giroud, E. Capria, E. Mitchell, C. Curfs, J. C. Royer, M. Baylac, F. Villa, and S. Rey, "Grenoble Large Scale Facilities for Advanced Characterisation of Microelectronics Devices," in *15th European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, Sep. 2015, pp. 312–315.

[18] T. Oldham and F. McLean, "Total ionizing dose effects in MOS oxides and devices," *IEEE Transactions on Nuclear Science*, vol. 50, no. 3, pp. 483–499, Jun. 2003.

[19] H. Hatano and M. Shibuya, "Total dose radiation effects on CMOS ring oscillators operating during irradiation," *IEEE Electron Device Letters*, vol. 4, no. 12, pp. 435–437, Dec. 1983.

[20] F. L. Kastensmidt, E. C. P. Fonseca, R. G. Vaz, O. L. Goncalez, R. Chipana, and G. I. Wirth, "TID in Flash-Based FPGA: Power Supply-Current Rise and Logic Function Mapping Effects in Propagation-Delay Degradation," *IEEE Transactions on Nuclear Science*, vol. 58, no. 4, pp. 1927–1934, Aug. 2011.

[21] "Single Event Effects test Method and Guidelines - Basic Specification," ESA, Tech. Rep., 1995. [Online]. Available: https://escies.org/download/webDocumentFile?id=62690

[22] N. Battezzati, S. Gerardin, A. Manuzzato, A. Paccagnella, S. Rezgui, L. Sterpone, and M. Violante, "On the Evaluation of Radiation-Induced

Transient Faults in Flash-Based FPGAs," in *14th IEEE International On-Line Testing Symposium*, Jul. 2008, pp. 135–140.

[23] M. J. Cannon, A. M. Keller, C. A. Thurlow, A. Pérez-Celis, and M. J. Wirthlin, "Improving the Reliability of TMR With Non-triplicated I/O on SRAM FPGAs," *IEEE Transactions on Nuclear Science*, vol. 67, no. 1, pp. 312–320, Jan. 2020.