

Radiation-induced Effects on DMA Data Transfer in Reconfigurable Devices

Original

Radiation-induced Effects on DMA Data Transfer in Reconfigurable Devices / Portaluri, A., Azimi, S., De Sio, C., Sterpone, L., Merodio Codinachs, D.. - ELETTRONICO. - (2022), pp. 1-7. (IEEE The 28th International Symposium on On-Line Testing and Robust System Design (IOLTS 2022) Turin 12-14 September 2022) [10.1109/IOLTS56730.2022.9897262].

Availability:

This version is available at: 11583/2970901 since: 2023-01-13T10:04:33Z

Publisher:

IEEE

Published

DOI:10.1109/IOLTS56730.2022.9897262

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Radiation-induced Effects on DMA Data Transfer in Reconfigurable Devices

Andrea Portaluri, Sarah Azimi, Corrado De Sio, Luca Sterpone
Politecnico di Torino
Dipartimento di Automatica e Informatica
Turin, Italy
{andrea.portaluri, sarah.azimi, corrado.desio,
luca.sterpone}@polito.it

David Merodio Codinachs
European Space Agency (ESA)
European Space Research and Technology Centre
Noordwijk, The Netherland
david.merodio.codinachs@esa.int

Abstract—As the adoption of SRAM-based FPGAs and Reconfigurable SoCs for High-Performance Computing increased in the last years, the use of Direct Memory Access for data transfer becomes a key feature of many reconfigurable applications even in the space industry. For such kinds of applications, radiation-induced effects are a serious issue that mines the correctness and success of mission-critical tasks. In this paper, we evaluate the effects of proton-induced errors on a DMA-based application implemented on a Xilinx Zynq-7020 FPGA in order to quantify the robustness of this module in a typical hardware-accelerated configuration. The obtained results confirm the high criticality of the DMA module on programmable logic. Moreover, the Multiple Bits Upsets effect has been evaluated. The most recurring patterns have been reported in order to provide further tools to better characterize the behavior of these systems under future fault injection campaigns, as demonstrated in the experimental results.

Keywords—DMA, Fault Injection, MBU, FPGA, Radiation-induced Effects, Reconfigurable SoC, SEU.

I. INTRODUCTION

The usage of SRAM-based Field Programmable Gate Arrays (FPGA) has been tremendously increasing in the last decades, especially in avionics and aerospace [1]. This growing appreciation is mainly due to their reprogrammable capabilities and hardware-accelerating features. These devices support on-field reconfiguration i.e., a reprogramming of the implemented electronic circuit, either in the case of different tasks to perform or when nominal work parameters change. Moreover, major FPGA vendors are now capable of offering both programmable hardware and the computing efficiency of single or multi-core microprocessors on the same chip, making available real heterogeneous computers with the so-called Reprogrammable System-on-Chips [2]. All these features together make such devices pave their way through a very broad spectrum of uses like communications, satellites, automotive systems, and space data center applications [3]. The latter, in particular, is nowadays the ultimate frontier for data centers that are investigating the adoption of hardware-accelerated algorithms for enabling data-centric applications in space. Thus, the managing of data on the terascale is beginning to play a key role in FPGA devices for aerospace data processing. Among several alternatives, Direct Memory Access (DMA) offers a performant approach to it: allowing subsystems (e.g., peripherals) to access memory independently of the Central Processing Unit (CPU), it speeds up the data transfer to and from memories, that is one of the common bottlenecks of modern high-performant systems.

Moreover, recent DMA engines are capable of achieving a very high speed of transfer, making them suitable solutions for image processing and neural network applications [4][5]. However, the main features of the SRAM-based FPGAs come along with a major disadvantage: the vulnerability to radiation-induced errors. The current configuration of the device is stored in volatile memory, called Configuration Memory (CRAM), which is highly susceptible to the voltage spikes that high-energy particles (e.g., protons, neutrons, heavy ions) can create when striking the silicon surface of the CRAM. As a result of such an interaction, soft errors can eventually arise Single-Event Upsets (SEUs) effect. From the CRAM point of view, the occurrence of an SEU can be seen as a bit-flip in its content that can cause a change in a memory or a logic cell and, therefore, unpredictable behaviors of the device. Thus, in order to achieve robust and radiation tolerant designs [7] able to cope with data-intensive operations on reconfigurable devices, it is important to analyze and estimate the sensitivity of the FPGA and its components to radiation effects. No analysis focused on the DMA module yet, although being a key component for data transfer tasks and, to overcome this, our work aims to analyze and quantify the reliability of DMA architectures in proton-irradiated environments, like Low Earth Orbit (LEO), while performing high-intensity data transfer operations. A proton irradiation campaign and software fault injections have been carried out on a PYNQ-Z2 board implementing a benchmark circuit on the Xilinx Zynq-7020. The benchmark circuit exploits two DMAs for transferring data to and from two 16-core hardware accelerators. The radiation test brought out interesting results, highlighting the DMA system as the criticality of the design and how a single particle causing multiple bit upsets can eventually end up with different DMA transfers failing simultaneously, namely cross-domain errors. Finally, a fault injection campaign shows how using the analysis coming from the radiation test, it is possible to increase the fidelity of fault injection experiments predictions by exploiting the Multiple Bits Upset (MBU) patterns we found. All the processed data have been collected at Paul Scherrer Institute (PSI) in Villigen, Switzerland.

The paper is structured in the following way: section II gives a brief overview of DMA operations and radiation-induced errors, while section III presents the related works. Section IV and V describe the experimental workflow and results. Section VI discusses the conclusions.

II. TECHNICAL BACKGROUND

A DMA architecture allows hardware subsystems to directly access memory without the intervention of the CPU. This enables the processor to keep on working concurrently on other tasks while other memory operations take place. In order to extend the concept, with the term DMA, we refer to the controller (implemented in programmable hardware) that manages the data transfer rather than the actual hardware. DMA has become a key module in reconfigurable devices as it is also capable of performing data transfer and conversion among Advanced eXtensible Interface (AXI) Memory Mapped and AXI Stream interfaces [7]. The AXI is a standard communication interface supported by Xilinx and ARM [8] that allows designers to combine hardware and software solutions through a high-bandwidth infrastructure. In FPGAs, DMA data transfer is implemented through the instantiation of a controller in the programmable logic that exploits the hard-wired resources on the board. Thus, DMA implementation on reconfigurable SoC involves a high number of CRAM resources so highly prone to SEU-caused corruption. Moreover, this work focused on the Scatter-Gather (SG) engine, where the data transfer is initialized and defined in blocks of 32- or 64-bit (depending on the architecture) words, called Buffer Descriptors (BD). The BDs are responsible for describing read/write addresses, the direction of transfer, data length, number of words to transfer, and so on. As opposed to the Direct Register Mode approach (or Simple DMA), where every transfer cycle is called by the CPU, BDs are written only once and then pointed to the SG engine, allowing transfer loops with no involvement of the processor whatsoever. Figure 1 shows a conceptual scheme of DMA data transfer with the SG engine. Typically, FPGAs are programmed by loading configuration data, called bitstream in the configuration memory of the device. Each bit is responsible to configure logic elements such as Look-Up Tables (LUT), Flip-Flops (FF), Configuration Logic Block (CLB), and interconnections among them. Given correct timing and amplitude, voltage spikes caused by the interaction with an ionizing particle can often end up corrupting an electrical node corresponding to a bistable element in the CRAM and, so, modifying the implemented electrical circuit [9][10]. The occurrence of such a radiation-induced fault is known as SEU. Figure 2 shows the effects of SEU in the CRAM of reconfigurable devices.

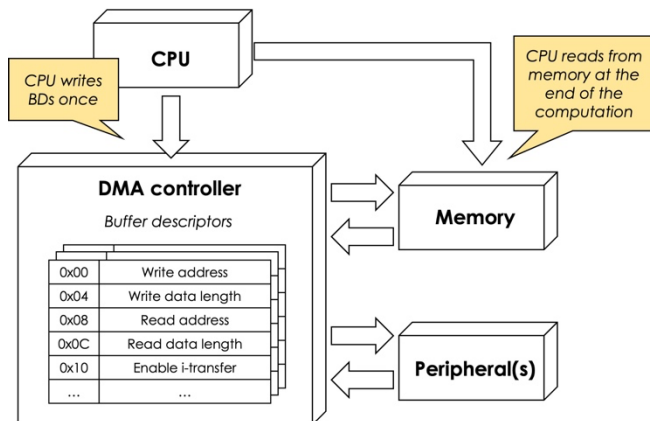


Fig. 1. A scheme of DMA data transfer architecture with SG engine.

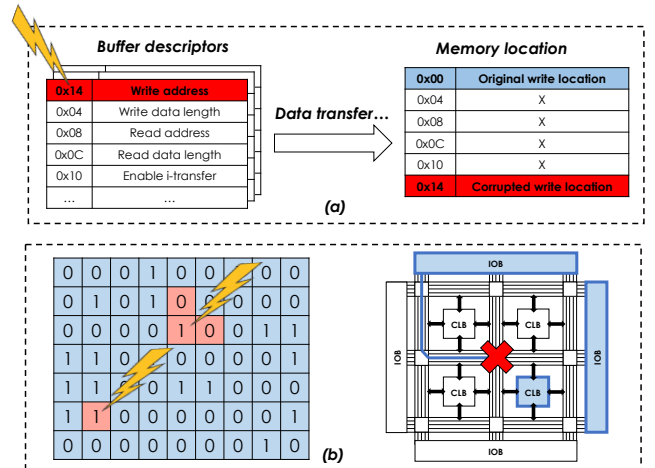


Fig. 2. Examples of errors in DMA due to radiation-induced bitflips. (a) Corruption of the BD causing wrong output products. (b) Corruption of routing bits causing output unavailability.

Although SEUs are classified as non-destructive events, such corruption may affect the device until a reconfiguration or a power cycle is performed, since CRAM content is not usually rewritten during the execution [11]. As it will be shown in this work, SEU effects can lead to severe effects such as the corruption of different modules at once, arising common-mode errors. In the DMA error scenarios, failures can mainly occur due to two separate factors: the corruption of the hardware in the programmable logic due to configuration memory faults (e.g., routing and PIP corruptions), or BDs modifications due to storage memory faults. In order to isolate the formers, in this experiment BDs have been re-written every application cycle and, in doing so, possible bitflips that corrupted them are continuously corrected.

III. RELATED WORKS

Different studies have described and analyzed data coming from radiation tests on reconfigurable devices, although the vast majority have focused on heavy-ions interaction. In [12], the authors characterize another Xilinx 28nm CMOS device, Kintex-7, against ultra-high energy heavy ions. The work presented in [13] describes the BRAM behaviors of a Zynq-7000 under heavy ions irradiation, an important and complementary result to our work since BRAM will contain the BDs as explained later. The authors in [14] present a characterization of embedded memories of Zynq-7000 to heavy ions and protons, while [15] and [16] focus on the embedded processor. Instead, several works have evaluated the reliability of the AXI interface within Xilinx Reconfigurable SoCs. Among the others, the authors in [17] quantified the errors and analyzed behaviors due to SEU effects in the AXI Interconnect Intellectual Property (IP) blocks, while in [18] the focus is on the interface between AXI Stream and Memory Mapped flows, also evaluating hardened circuits. However, although technically flawless, these works simulate the SEU by random bit-flipping the content of the CRAM in order to reproduce the soft error. As demonstrated in these references, heterogeneous circuits exploiting DMA are extremely susceptible to soft errors and, even if memories have been covered by previous investigations, DMA has been considered critical and not yet analyzed.

Moreover, to the best of our knowledge, no studies have focused on the correlation between common-mode failures between computing cores and DMA. Finally, additional data on MBU on Zynq-7000 technology coming from proton test are provided as further contribution.

IV. RADIATION ANALYSIS WORKFLOW

In order to evaluate the radiation-induced effects on reconfigurable DMA, a workflow has been followed. We first performed a proton radiation test and collected the test output coming from an FPGA implementing a DMA-based data transfer towards hardware accelerators. Data on cross-sections, failure rate, and patterns of MBU have been listed and fault injection campaigns have been carried out exploiting these a-posteriori data. In the following subsections, the analysis workflow is explained in detail.

A. Application Benchmark

Considering typical employment in hardware acceleration, we developed a benchmark that could manage a data transfer and computation of a workload, focusing on the heavy use of the DMA capabilities rather than the output itself. Thus, each DMA is fully using its multichannel capability (16 channels out of 16) and its burst length (maximum number of words transferrable per packet), while managing the data transfer to and from a matrix of hardware accelerators implemented on programmable logic. Figure 3 shows a conceptual scheme of the developed benchmark. The microprocessor is interfaced, through Master AXI High-Performance ports, with two different AXI DMA IP Blocks supporting multichannel mode and managing the data transfer to and from a 16-core hardware accelerator for each DMA. The accelerator is made up of multiple COordinate Rotation DIgital Computer (CORDIC) IP Cores [19]. The CORDIC IP cores implement the CORDIC algorithm, a widely adopted algorithm for computing transcendental functions, square roots, and vectors rotation without using a multiplier. In this particular instance, the cores were designed to compute the sine and cosine of a stream of input data. Memory Mapped-to-Stream (MM2S) and Stream-to-Memory Mapped (S2MM) transfer directions are addressed by different BDs allocated in a dedicated BRAM memory for each DMA. Interrupt signals coming from the DMAs have been concatenated and properly managed. A software routine runs on the processor in order to print the content of the output data buffers, write the BDs and handle the interrupts.

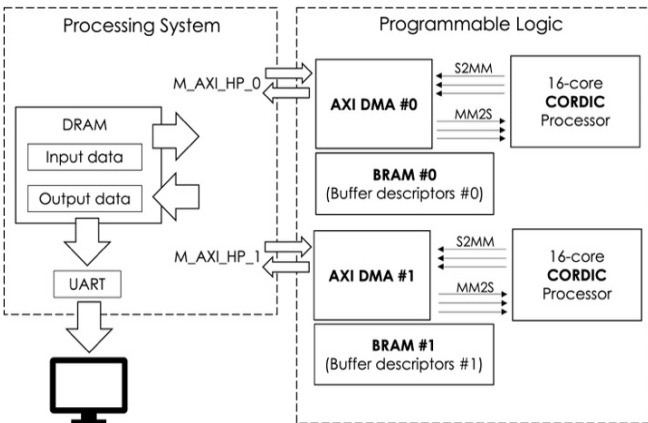


Fig. 3. Block scheme of the implemented benchmark.

B. Radiation test analysis

The proton irradiation campaign was performed at PSI Proton Irradiation Facility. For the later sections the following terminology has been used:

- *Energy* [MeV] determines the carried energy of a single proton.
- *Flux* [cm^2s^{-1}] describes the number of particles crossing a squared-centimeter section in a time unit.
- *Fluence* [cm^2] represents the actual number of particles that crossed a squared-centimeter section during the board exposure.

Finally, with the term *cross-section* [cm^2], we refer to the probability of a given unit (e.g., the DMA modules, CORDIC cores) to fail under the effect of a radiant excitation. Let us suppose that, for a given energy E , different N beam runs took place. Thus, the cross section per computational core is computed as the total number of occurred events during the N runs over the product of the sum of the N fluences (F) and a normalization constant k (e.g., number of cores), as shown in Equation 1.

$$\sigma_E = \frac{\sum_i^N \text{events}_i(E)}{\sum_i^N F_i(E) \cdot k} \quad (1)$$

Several energies and fluxes have been provided, allowing us to witness different behaviors of the devices while covering a quite broad spectrum of LEO scenarios [20], as reported in Section V.

The device has been mounted and centered with respect to the beam on an adjustable frame, aided by a cross laser sight. Figure 4 shows the experimental flow. In detail, the device has been monitored and configured by a dedicated host computer, connected through a serial port, that runs a Python script to manage output data and the periodic acquisition of the CRAM content, namely the readback procedure. The acquisition of the readback data has been performed with the maximum allowed frequency (about every 3 seconds) and in parallel to the execution of the application on the board. Detecting the differences between the bitstream and the readback file, we have been able to compute in real-time the number of SEU that occurred on the board.

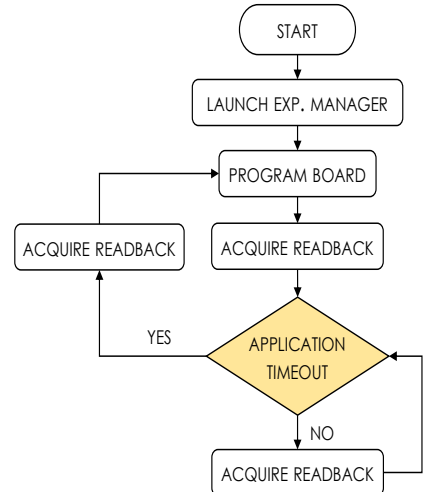


Fig. 4. Experimental flow diagram.

The application execution stopped only in case of application timeout in order to achieve accumulation of bitflips. A periodic check on the availability of output data was performed and, in case, the board re-programmed. From the control room of the facility, it has been possible to run beam sessions with arbitrary durations and to perform power cycles for the board whenever needed.

C. Fault injection analysis

After the analysis of radiation test data, the patterns of occurred MBU have been identified (whose more detailed information is presented in Section V). Two different fault injections with CRAM error accumulation, the former using the SEU fault model, and the latter using both SEU and MBU patterns fault models, have been performed and compared in order to discover if these data can be used to refine the prediction of faults injection campaigns. First, a fault injection platform has been developed in order to emulate the accumulation of faults in the configuration memory. The environment automatizes both fault generation and evaluation of output data. For this scope, we used the in-house developed PyXEL framework to perform bit-flips in the CRAM content [21]. PyXEL is a Python tool that eases the process of performing single or multiple bits fault injections by decoding the configuration memory of the device. A bitstream is corrupted one bit at a time and, then, downloaded into the board until it produces a fatal corruption causing a system crash. The bit-flips can involve randomly any configuration bits. A parallel thread receives the output data from the serial interface. Lastly, the collected output data are parsed and a report with the occurrences of faults is generated. In order to build a reliability curve and so compare the results of fault injections and radiation tests, the algorithm shown in Figure 5 has been developed. The fault injection algorithm has been developed in order to insert bit-flips randomly or on the basis of the patterns observed with the a-posteriori analysis of the radiation test experiments.

V. EXPERIMENTAL RESULTS

For the experiment, we tested an off-the-shelf Digilent PYNQ-Z2 Development Board, as shown in Figure 6. This device mounts a 28nm CMOS Z7020 Zynq 7-Series FPGA embedded with a dual-core ARM Cortex-A9 Processor. With reference to the detailed description of Section IV, the utilized resources for the implemented benchmark on PYNQ-Z2 are presented in Table I.

Algorithm 1 Algorithm for fault injection with accumulation with PyXEL

Input: *bitstream*, *no_failures*

Initialization:

- 1: Instantiate object of class FaultInjector with *bitstream* as an argument
- 2: Instantiate object of class SerialListener

Accumulation process:

- 3: **for** $i = 0$ to *no_failures* **do**
- 4: Duplicate and target *bitstream*
- 5: **while** True **do**
- 6: Perform random bit-flip (single or multiple) from patterns
- 7: Program board with *bitstream*
- 8: **if** state of serial listener = serial timeout **then**
- 9: break
- 10: **end if**
- 11: **end for**

Fig. 5. Pseudo-code for the fault injection analysis.




Fig. 6. Board exposed to the protons beam.

TABLE I
RESOURCES UTILIZATION FOR THE PYNQ-Z2

Resources	Used [#]	Available [#]	Usage [%]
LUTs	28,413	53,200	53.41
Flip-Flops	33,317	106,400	31.31
Memories	2,080	17,400	11.95
BRAM	14	140	10

Due to the characteristics of the programmable hardware, not all the occurred bit-flips have caused an error on the output. Indeed, some of them may target unused resources of the device that will not cause effects in the implemented design. First, we analyzed the radiation experiment data, computing the error rate, the particle cross-section of the peripheral, and the MBU patterns, then the comparison between fault injections and radiation test has been made.

a) Radiation test

We did not delid the chip on the device in order to maintain high fidelity between the experiment and the actual on-field application. Table II shows the values of energies and fluxes used during the radiation test experiment. The core errors were distinguished through a software voter running on the micro-processor. Golden values were re-written every cycle in order to avoid any possible corruption of them. DMA failures, instead, were counted through the interrupt signals. Table III shows the values of the cross-section for each energy, together with the CORDIC ones while Figure 7 plots while Figure 7 plots these values.

TABLE II
RADIATION TEST CONDITIONS: ENERGY, FLUX AND FLUENCE.

Energy [MeV]	Flux [$\text{cm}^{-2}\text{s}^{-1}$]	Fluence [cm^{-2}]
29.31	$4.124 \cdot 10^7$	$9.173 \cdot 10^{10}$
50.80	$4.024 \cdot 10^7$	$6.064 \cdot 10^{10}$
69.71	$4.110 \cdot 10^7$	$2.124 \cdot 10^{10}$
101.34	$4.319 \cdot 10^7$	$2.415 \cdot 10^{10}$
151.18	$4.094 \cdot 10^7$	$1.226 \cdot 10^{10}$
200	$4.144 \cdot 10^7$	$3.942 \cdot 10^{10}$

TABLE III
PARTICLE CROSS-SECTION PER DMA AND CORDIC

Energy [MeV]	Particle cross-section [cm ²]	
	DMA	CORDIC
29.31	2.725·10 ⁻¹¹	1.703·10 ⁻¹²
50.80	1.567·10 ⁻¹⁰	4.484·10 ⁻¹¹
69.71	2.354·10 ⁻¹⁰	5.592·10 ⁻¹¹
101.34	2.899·10 ⁻¹⁰	8.671·10 ⁻¹¹
151.34	3.263·10 ⁻¹⁰	6.627·10 ⁻¹¹
200	3.171·10 ⁻¹⁰	9.831·10 ⁻¹¹




Fig. 7. Particle cross-section per DMA and CORDIC.

Such a consideration can help the designer identify the most critical parts of the system. As matter of fact, the DMA data transfer on PYNQ-Z2 shows a higher sensitivity to protons-induced effects with respect to the computational cores since the cross-section is, on average, 20 times higher. Thus, the DMA modules must be counted as critical and they should be hardened accordingly. Further analysis has been carried out on the MBU events and their correlation with the DMA failure. Adopting the description given in [12], an MBU is intended as an event where a single particle, striking the silicon surface of the FPGA, causes multiple close bit-flips. Two bit-flips a and b are marked as close when their Euclidean distance is less or equal than $\sqrt{2}$, with the distance is computed as:

$$dist(a, b) = \sqrt{(bit(a) - bit(b))^2 + (frame(a) - frame(b))^2}$$

where $bit(\bullet)$ and $frame(\bullet)$ identify the coordinates of the bit-flip in the CRAM memory. Analyzing any two consecutive readbacks, we have been able to isolate several MBU patterns that occurred during the time frame of a readback acquisition. The few bitflips accumulating during each time window and the large size of the CRAM (more than 10^8 bits) allowed us to detect groups of SEUs with a strong correlation both in time and space. Figure 8 illustrates the detected patterns. They have been grouped in clusters for numbers of bitflips. The cross-section of the MBU clusters, shown in Figure 9, has been computed as the ratio between the occurrence of cluster size and the total number of particles that hit the board. Data show that the cross-section rate decreases as the cluster size increases. This may highlight the fact that the bigger sizes of MBU clusters require higher energy, thus turning out to be rarer events.




Fig. 8. MBU patterns per cluster size




Fig. 9. Cross-section of Different Clusters Pattern per device

Moreover, common-mode errors have been observed and analyzed. We refer to common-mode faults when two or more modules fail simultaneously due to the same SEU or MBU. Three kinds of common-mode errors are distinguished:

- *DMA-DMA (DD)* errors happen when the two DMA modules fail at the same time, due to the same SEU or MBU.
- *DMA-CORDIC (DC)* errors occur when DMA and computational core(s) are corrupted at the same time, due to the same SEU or MBU.
- *CORDIC-CORDIC (CC)* errors represent the failure of two or more cores simultaneously, due to the same SEU or MBU.

In Figure 10, the distribution of common-mode faults per MBU cluster size is presented.




Fig.10. Distribution of common-mode errors per MBU cluster. Common-mode errors have been registered only for some cluster sizes.

As can be noted from the data presented above, MBU occurrences can significantly impact the correctness of the application causing also common-mode errors. The vast majority of these happen simultaneously to a 2-cluster MBU, mainly due to a higher percentage of such a cluster size over the others. Moreover, CC common-mode errors are more likely to occur due to the elevated number of cores and their resources proximity in the programmable logic. Notice that such a corruption can also be due to single bit-flips. Previous works have demonstrated how physical isolation of resources can decrease these errors [22].

b) Fault injection

In order to verify whether the collected data can be used to refine the prediction of fault injection, two different campaigns have been carried out. Firstly, a total of 15,000 single bit-flips with accumulation have been performed collecting bitstreams that caused DMA corruption, following the injection algorithm shown in Fig. 5. Secondly, we performed another fault injection campaign, injecting MBU patterns on the basis of their rates of occurrences detected in the radiation test. The patterns shown in Figure 7 have been injected, together with single bit-flips, performing again 15,000 bit corruptions. Figure 11 presents the obtained results. The analysis shows how the single-bit with accumulation fault injection data can support the one coming from the radiation test, although small inaccuracies have been registered. These differences can be slightly accounted for by the fault injection with MBU fault models. Thus, it is reasonable to suppose that the main source of inequality is represented by the MBU effect. As we saw in the previous subsection, a striking particle may cause multiple bit-flips in a neighborhood of the impact that can surely affect the configuration in a different manner from the same amount of bit-flips injected randomly. Such a fault behavior is very complex to model in fault injection environments due to several variables such as the device and the chip technology that can affect the pattern of the multiple bits corrupted. Therefore, injecting SEU fault models only brings to an underestimation of the failure rate.

VI. CONCLUSION

This work focuses on the analysis of the sensitivity of reconfigurable DMA data transfer to radiation-induced faults on SRAM-based devices when adopted in typical hardware-accelerated applications.




Fig. 11. Failure rate with the accumulation of fault injections and radiation test data.

We tested a benchmark involving two multichannel DMA modules streaming data to arrays of processing elements implemented on a Zynq-7020 chip. Proton irradiation data, supplemented by an a posteriori fault injection campaign, have been collected and analyzed to quantify the criticality of DMA data transfer. Cross-section analysis has shown how the DMA, which is a central component for data transfer in the space data centers, is highlighted as one of the most critical modules in the programmable logic. DMA is resulted to be the cause of the majority of application errors, thus data transfer cannot rely exclusively on DMA especially when it represents a single point of failure for the system. Further fault injection analysis have shown that MBU patterns can be used to better predict failures of the systems when compared with single bit fault injection campaigns. As future works, we would like to deepen the MBU fault model in order to extend the capability of our fault injection platform, emulating such behavior. Moreover, we planned to focus on the effects that radiation-induced effects can have on the computational side, developing more coherent applications for high-performance computing.

ACKNOWLEDGMENT

The authors want to thank Dr. Wojciech Hajdas and all his staff at the Proton Irradiation Facility for the precious time and help using the proton beam.

REFERENCES

- [1] J. Romoth and M. Porrmann, "Survey of FPGA applications in the period 2000 – 2015", Technical Report, 2017.
- [2] Lee Hansen, "Unleash the Unparalleled Power and Flexibility of Zynq UltraScale+ MPSoCs", Xilinx WP470, 2016.
- [3] W. Anderson, et al., "The geostationary operational satellite R series SpaceWire based data system: Session: SpaceWire networks and protocols, long paper," *International SpaceWire Conference (SpaceWire)*, pp. 1-6, 2016.
- [4] Z. Wang et al., "CNN-DMA: A Predictable and Scalable Direct Memory Access Engine for Convolutional Neural Network with Sliding-window Filtering", *Proceedings on the Great Lakes Symposium on VLSI*, pp. 115-121, 2021.
- [5] M. Rhu, et al., "Compressing DMA Engine: Leveraging Activation Sparsity for Training Deep Neural Networks," *IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 78-91, 2018.
- [6] S. Azimi, et al., "A Radiation-Hardened CMOS Full-Adder Based on Layout Selective Transistor Duplication," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 29, no. 8, pp. 1596-1600, Aug. 2021
- [7] Xilinx, "AXI DMA v7.1", Xilinx PG021, 2019.
- [8] Xilinx, "AXI Reference Guide", Xilinx UG761, 2011.
- [9] S. Azimi, L. Sterpone B. Du and L. Boragno, "On the analysis of radiation-induced Single Event Transients on SRAM-based FPGAs," in *Microelectronics Reliability*, vol. 88-90, pp. 936-940, 2018.
- [10] S. Azimi and L. Sterpone, "Digital Design Techniques for Dependable High Performance Computing," *2020 IEEE International Test Conference (ITC)*, 2020, pp. 1-10.
- [11] C. De Sio, S. Azimi, L. Bozzoli, B. Du and L. Sterpone, "Radiation-induced Single Event Transient effects during the reconfiguration process of SRAM-based FPGAs," in *Microelectronics Reliability*, vol 100-101, 2019.
- [12] B. Du et al., "Ultrahigh Energy Heavy Ion Test Beam on Xilinx Kintex-7 SRAM-Based FPGA," in *IEEE Transactions on Nuclear Science*, vol. 66, no. 7, pp. 1813-1819, 2019.
- [13] V. Vlagkoulis et al., "Single Event Effects Characterization of the Programmable Logic of Xilinx Zynq-7000 FPGA Using Very/Ultra High-Energy Heavy Ions," in *IEEE Transactions on Nuclear Science*, vol. 68, no. 1, pp. 36-45, 2021.
- [14] L. A. Tambara, A. Akhmetov, D. V. Bobrovsky and F. L. Kastensmidt, "On the Characterization of Embedded Memories of Zynq-7000 All

Programmable SoC under Single Event Upsets Induced by Heavy Ions and Protons," *2015 15th European Conference on Radiation and Its Effects on Components and Systems (RADECS)*, 2015, pp. 1-4.

- [15] S. Azimi, C. De Sio, D. Rizzieri, and L. Sterpone, "Analysis of Single Event Effects on Embedded Processor," *Electronics*, vol. 10, no. 24, p. 3160, Dec. 2021
- [16] D. M. Hiemstra and V. Kirischian, "Single Event Upset Characterization of the Zynq-7000 ARM® Cortex™-A9 Processor Unit Using Proton Irradiation," *2015 IEEE Radiation Effects Data Workshop (REDW)*, 2015, pp. 1-3.
- [17] C. De Sio, et al., "On the analysis of radiation-induced failures in the AXI interconnect module," in *Microelectronics Reliability*, vol. 114, 2020.
- [18] F. Benevenuti and F. Kastensmidt, "Reliability evaluation oninterfacing with AXI and AXI-S," *IEEE 19th Latin-American Test Symposium*, 2018.
- [19] Xilinx, "CORDIC v6.0", Xilinx PG105, 2021.
- [20] G. D. Badhwar, "The Radiation Environment in Low-Earth Orbit," in *Radiation Research*, vol. 148, 1997.
- [21] L. Bozzoli, et al., "PyXEL: An Integrated Environment for the Analysis of Fault Effects in SRAM-Based FPGA Routing," *International Symposium on Rapid System Prototyping (RSP)*, pp. 70-75, 2018.
- [22] A. Portaluri, et al., "A New Domains-based Isolation Design Flow for Reconfigurable SoCs," *IEEE 27th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, pp. 1-7, 2021.