



**HAL**  
open science

## Characterization of a Fault-Tolerant RISC-V System-on-Chip for Space Environments

Douglas Santos, André Martins Pio de Mattos, Douglas Melo, Luigi Dilillo

► **To cite this version:**

Douglas Santos, André Martins Pio de Mattos, Douglas Melo, Luigi Dilillo. Characterization of a Fault-Tolerant RISC-V System-on-Chip for Space Environments. DFT 2023 - 36th IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems, Oct 2023, Juan-Les-Pins, France. pp.1-6, 10.1109/DFT59622.2023.10313549 . hal-04266887

**HAL Id: hal-04266887**

**<https://hal.science/hal-04266887v1>**

Submitted on 31 Oct 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

This is a self-archived version of an original article.  
This reprint may differ from the original in pagination and typographic detail.

**Title:** Characterization of a Fault-Tolerant RISC-V System-on-Chip for Space Environments

**Author(s):** Douglas A. Santos, André M. P. Mattos, Douglas R. Melo, and Luigi Dilillo

**Document version:** Post-print version (Final draft)

**Please cite the original version:**

D. A. Santos *et al.*, "Characterization of a Fault-Tolerant RISC-V System-on-Chip for Space Environments," 2023 36th IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), 2023.

*This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorized user.*

# Characterization of a Fault-Tolerant RISC-V System-on-Chip for Space Environments

Douglas A. Santos\*, André M. P. Mattos\*, Douglas R. Melo†, and Luigi Dilillo\*

\*IES, University of Montpellier, CNRS, Montpellier, France

†LEDS, University of Vale do Itajaí, Itajaí, Brazil

{douglas.santos, andre.martins-pio-de-mattos}@umontpellier.fr, drm@univali.br, luigi.dilillo@umontpellier.fr

## Abstract

Processing units are sensitive to the harsh radiation conditions present in space applications. Thus, radiation testing is a mandatory step toward high reliability for these systems. Notably, as an emerging processor architecture, RISC-V has attained interest for its utilization in these applications. For this reason, in previous work, we presented a fault-tolerant RISC-V System-on-Chip (SoC) with enhanced fault awareness, known as HARV-SoC, to attend to this increasing demand. The SoC includes structures to detect, correct, and report radiation-induced errors. In this work, we characterize HARV-SoC under proton irradiation, enabling the proposition of space applications based on a low-cost and open-source RISC-V solution. We observed and reported the effects of Single Event Effects (SEEs) and Total Ionizing Dose (TID) in the entire SoC, individually evaluating the HARV core, SoC elements, and external memories, with a detailed analysis of error propagation.

## Index Terms

RISC-V, System-on-Chip, Fault Tolerance, Radiation Effects, Single-Event Effects, Total Ionizing Dose, Protons

## I. INTRODUCTION

Designing systems for space applications is a challenging task due to various strict requirements and harsh operational conditions. In orbit, a serious threat to these systems is the degradation caused by radiation. Ionizing particles constantly interact with the system's components and materials, leading to a plethora of effects. In some cases, these particles are capable of inducing faults, either due to single interactions, a Single Event Effect (SEE), or caused by their accumulative effect, Total Ionizing Dose (TID) and Displacement Damage (DD). The induced faults can be transient, intermittent, or permanent [1].

In order to understand these effects and enable mitigation techniques to be applied, tests are prepared to characterize these systems under irradiation [2]. For that, radiation conditions are recreated in particle accelerators, where users have the ability to change important parameters in a controlled environment for accelerated experimental analysis [3]. These facilities usually provide specific particle species and spectra that allow reliability assessment of applications targeting a certain environment. For instance, for representative SEE characterization for space applications, proton and heavy-ion characterization is suggested. In space standards [4], [5], this recommendation is further defined, and its applicability is discussed.

An emerging trend in the last years has been the employment of Commercial Off-The-Shelf (COTS) components in modern space developments with higher risk acceptance [6]. This movement is mainly motivated by the high cost, long lead times, and limited performance of traditional radiation-hardened components. However, to mitigate the several reliability challenges raised by this approach, developers often exploit risk acceptance by applying fault tolerance techniques on their COTS-based systems. They apply strategies and techniques that usually rely on redundancy, using temporal, spatial, and informational characteristics of a system [7].

Processing units are a central element of complex space systems, for which COTS alternatives have been applied and investigated. In particular, RISC-V [8] is an emerging processor architecture that has attained an important role in various application domains, reaching even space systems as a promising alternative for future developments [9]. For this reason, we started the development of a RISC-V processor prospecting space applications in prior work [10]. Then, the processor was extended to become a robust SoC (known as HARV-SoC) with enhanced fault observability and monitoring for proper experimental validation, in which extensive tests were performed in various irradiation campaigns [11].

This study has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101008126, the Region d'Occitanie and the École Doctorale I2S from the University of Montpellier (contract no. 20007368/ALDOCT-000932), the Foundation for Support of Research and Innovation, Santa Catarina (FAPESC-2021TR001907), and the Brazilian National Council for Scientific and Technological Development (350794/2023-5).

In this work, we continue the validation effort by extending HARV-SoC characterization with high-energy protons. Thus, we present the experimental preparation, acquired results, and post-analysis, contributing to a more comprehensive understanding of the SoC under harsh radiation conditions. As aforementioned, the characterization of protons is an important step towards enabling the usage of systems in space applications. Therefore, we aim to present to the community the robustness of HARV-SoC as an alternative COTS solution targeting in-orbit applications.

This paper is organized as follows. The fault-tolerant RISC-V SoC is presented in Section II. Section III describes the irradiation facility and the proposed experiment. Section IV presents results for the proton irradiation. In Section V, we discuss these results. Finally, Section VI provides the conclusions.

## II. FAULT-TOLERANT RISC-V SYSTEM-ON-CHIP

The RISC-V System-on-Chip tested in this work (HARV-SoC) focuses on reliability for harsh environments [11]. The SoC is based on HARV (Hardened RISC-V) processor core [10], which applies hardening techniques on the architectural level to allow detecting, correcting, and reporting faults. HARV is a multi-cycle RISC-V-based processor core supporting the integer instruction set (RV32I), Control and Status Registers (CSRs), interrupts, and exceptions. The multi-cycle architecture is composed of five stages: instruction fetch, execution, memory access, write back, and trap handler. HARV applies techniques such as Error-Correcting Code (ECC) and Triple-Modular Redundancy (TMR) for the detection and correction of errors in its internal components. The register file, the program counter, and the instruction register are hardened using Single-Error Correction and Double-Error Detection (SECDED). The control, Arithmetic-Logic Unit (ALU), and critical CSRs use TMR for hardening. HARV also provides error reporting capabilities, which are provided through exceptions that synchronously interrupt the application flow and report detailed information on the detected error, such as the error type, application context, and corrupted data. These features improve the reliability of the SoC and support fault awareness at the application level.

HARV-SoC is composed of the HARV processor core and various peripherals (e.g., communication interfaces, memory controllers) integrated into an architecture targeting reliability with flexibility and low hardening overheads. This reliability is achieved by hardening critical internal components and providing recovery schemes, e.g., data memory with ECC, bus access timeout, reset controller, watchdog timer reset, and application checksum. For the data memory, the SoC uses an external SDRAM memory, accessed through a customized controller that enables SECDED for each 32-bit word. For access to peripherals, the bus controller provides a timeout flag that triggers an exception when a peripheral is not responding. With the bus timeout, the application receives information about the faulty peripheral. The application may use this information to perform actions to restore the functionality of a peripheral. For that, the reset controller is used, which enables resetting specific peripherals. Besides the timeout from the bus, the system also provides a watchdog timer that resets the SoC when a system timeout is reached, preventing hang processor failures.

Besides these hardening strategies, an application integrity verification is performed by the startup routine using a Cyclic Redundancy Check (CRC) algorithm. This enables the detection of errors while loading the program from external non-volatile memories. Furthermore, the reporting capabilities of the error monitor enable the application to perform actions according to the detected error. If the error is not correctable, the application uses the available information to decide, for example, to ignore it and continue the execution or even force a system reset to avoid faulty system behavior. Fig. 1 presents an overview of the SoC architecture, including the fault-tolerant HARV processor, its interconnection with peripherals, and the error monitor module, responsible for detection and reporting.

## III. RADIATION EXPERIMENT

To evaluate the SoC reliability, we performed a proton irradiation campaign. The following subsections provide a description of the irradiation facility, experimental setup, and test modes.




Fig. 1. HARV-SoC architecture [11].

### A. Irradiation Facility

PARTREC [12] is a dedicated research facility located in Groningen, Netherlands. It includes an accelerator that is used for nuclear physics research and radiation-hardness testing of electronic systems. This facility operates a large superconducting cyclotron capable of generating high-energy protons up to 190 MeV with fluxes up to  $1 \times 10^8$  p/cm<sup>2</sup>/s depending on the beam area. For this experiment, the beam configuration had a homogeneity uncertainty under 10 %. This type of irradiation provides the necessary conditions for triggering representative SEEs in the SUTs (Systems Under Test), the primary phenomena under investigation in this work. However, another effect is also observed: the accumulation of damage in the SUTs, known as TID. Both phenomena are presented and discussed in Section IV.

### B. Experimental Setup

The experimental setup consisted of various boards hosting the SoC in flash-based FPGA devices and the auxiliary circuitry for operation (e.g., power supply, logging interfaces). Only the FPGAs and external SDRAM memory were irradiated, leaving the remaining devices outside the direct beam sight. In particular, the FPGA is the Microchip's SmartFusion2 M2S010 [13], and the SDRAM memory is the Winbond W9864G6JT [14]. The FPGA configuration memory is based on flash cells, which are inherently more robust against SEEs than traditional FPGAs with SRAM-based cells, as shown by the manufacturer [15], [16]. Despite that, other internal structures of the chosen FPGA are still sensitive to SEEs, such as the Block RAMs (BRAMs), which are based on SRAM technology, and D-type Flip-Flops (DFFs), as shown in [16], [17]. For logging the experimental data, interfaces were implemented using a serial protocol available in the SoC architecture, in which each board had a direct and independent connection with a host computer. To achieve a representative characterization, i.e., a statistically significant number of radiation-induced effects, we used six boards for the irradiation. The boards were attached in frames without superposition and as close as possible to the center of the defined beam area. The supply current was individually monitored and limited in case of faulty behavior.

### C. Test Modes

Each SUT executed a benchmark: EEMBC's CoreMark [18] and the open-source Embench [19]. The usage of multiple benchmarks intends to improve the test stimuli, as discussed in [20]. CoreMark is a consolidated performance benchmark that includes four main workloads: list processing, matrix manipulation, state machines, and CRC. The CRC algorithm is used as a workload and provides a self-checking scheme for the inner steps of the benchmark execution. CoreMark is executed in iterations and provides a final result check. Embench is a free and open-source performance benchmark and intends to provide a broader set of algorithms, resulting in twenty-two different workloads. We implemented a combination of these algorithms in a single test application, which uses a custom set of these algorithms due to memory size constraints.

For the proposed radiation experiment, we assigned the first three boards with CoreMark and the other three with Embench. The execution is performed in a cycling test mode, in which the benchmark runs a short number of iterations and outputs the final results, then the SoC performs a soft reset before starting a new cycle. Every cycle runs the processors with a different hardening configuration, i.e., reporting and corrections enabled or just the reporting. In this experiment, we have two hardening configurations: baseline, with correction disabled, except the SECDED in the external memory, and hardened, with all the available protections enabled. We kept the SECDED enabled in the baseline to reduce the number of memory-related errors, as observed to be the most contributing error in all the previous characterizations, such as [11]. Moreover, these experiments consistently proved the importance of ECC in the system memory, leaving the opportunity for exploring other reliability aspects in this work.

It is worth mentioning that, for the first beam runs of this experiment, we experimented with different proton fluxes to define the most adequate for testing HARV-SoC. We performed the test in two stages: in the first, we irradiated four boards (SUTs #1, #2, #4, and #5) until unrecoverable failure, as deeply discussed in Section IV, with fluxes between  $1 \times 10^6$  and  $1 \times 10^8$  p/cm<sup>2</sup>/s, then we tested the remaining two boards (SUTs #3 and #6) following the same approach, but mostly with a flux of  $2 \times 10^8$  p/cm<sup>2</sup>/s. Both stages included the same setup and procedures, with the exception that the second run used a single flux for most of the experiment instead of different ones based on the initial verification.

## IV. RESULTS

This section reports the results and analysis of the test campaign, providing details on the observed proton-induced SEE and TID effects on the SUTs. We perform an initial investigation regarding the impact of TID on the SUTs and device failures to ensure that no bias is present in the SEE analysis due to accumulative effects. Then, we detail the detected errors and their cross sections, including an error propagation analysis.

## A. Device Failure Analysis

During the test campaign, the SUTs accumulated a significant dose, which led to unrecoverable device failures. For this reason, the first analysis that we performed was to understand the impact of the accumulative effect on each SUT. However, it is important to note that this work focuses on the RISC-V SoC analysis regarding SEE effects as the opposite of device (FPGA) characterization, in which the TID analysis only is performed to ensure proper filtering of possibly biased data. Besides that, to report a more insightful result, we investigated the most probable cause for the TID limitation of the device.

In Fig. 2, we show the accumulative number of HARV Fault Tolerance Exceptions (HFTE) over the received dose in the test campaign. HFTEs are exceptions triggered in the processor related to the fault-tolerant structures in the HARV implementation. The curves are plotted until the last reported HFTE and the dose values are estimations since only doses per beam run were reported by the facility. As expected, an abrupt increase is observed at the end of the irradiation for most SUTs since faulty and critical behavior was noticed during the test campaign. After 25 *krad*, SUT #2 presented a rapid increase in the number of HFTEs, followed by SUT #6 and SUT #5 at 29 *krad* and 31 *krad*, respectively. The remaining SUTs shortly presented a similar behavior. Hence, due to a clear dose dependency for the number of HFTEs, we decided to adopt a worst-case approach, in which we defined the threshold for SEE analysis at 25 *krad*, avoiding biases a

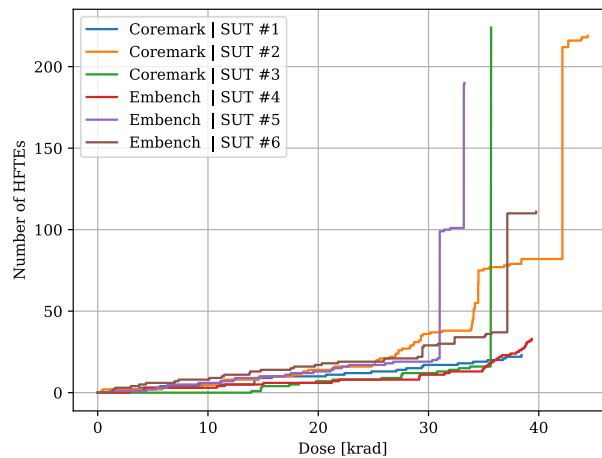


Fig. 2. Accumulated number of HFTEs vs TID of each board.

As mentioned, we further investigated an important behavior that may point to the root cause of the observed TID limitation. In most cases, moments before reaching an unrecoverable device failure, the SoC reported many CRC errors used to guarantee the integrity of the application code. As described, the integrity verification is executed during the startup routine, which is always called after a soft or hard reset, meaning a watchdog recovery attempt or a manual power cycle, respectively. Fig. 3 presents the accumulated number of CRC errors over the received dose in the test campaign. Similar to Fig. 2, the SUTs present a significant increase in CRC error at approximately 30 *krad*, corroborating the previous proposition of setting a conservative threshold for SEE analysis at 25 *krad*. It is worth mentioning that the curves were plotted until the last reported CRC error, which may differ from the last reported HFTE in Fig. 2. Since the application code is stored in an internal flash memory connected to the HARV through complex interfaces, our main hypothesis is that an element in this interface or in the controlling logic of this memory block failed with the accumulated doses.

Thus, to define the fluence and dose of the discussed unrecoverable device failure, we set the condition as the last execution of the SUT when even a power cycle was ineffective in restoring the operation. Table I presents these results for each SUT. Besides that, for comparison, we included the fluence used as a threshold for SEE analysis.

Finally, an evaluation of device failure throughout the test campaign was performed. The SUTs experienced failures that required external intervention with a complete power cycle. It is important to note that these failures are not recovered with soft reset attempts triggered by the SoC watchdog timer. Table II details the number of device failures per SUT. For the SUTs #1, #2, #4, and #5, we noticed a similar number of failures, while SUTs #3 and #6, which were tested in the second batch of the experiment, presented fewer failures.

## B. Error Distribution

Before investigating the reported errors, an analysis was performed to verify the distribution of errors for each SUT. Table III summarizes the reported errors per board considering both benchmarks used in the experiment. From this

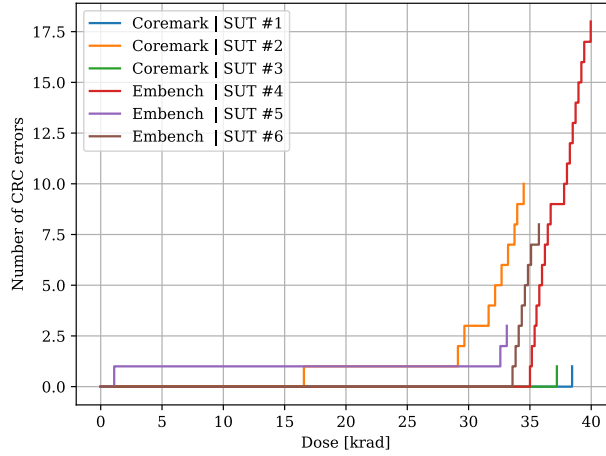


Fig. 3. Accumulated number of CRC errors per TID.

TABLE I  
ACCUMULATED FLUENCE AND DOSE UNTIL UNRECOVERABLE DEVICE FAILURE PER SUT.

SUT	TID [krad]	Fluence [ $\text{p}/\text{cm}^2$ ]	Fluence [ $\text{p}/\text{cm}^2$ ] (TID < 25 krad)
SUT #1	38.44	$5.94 \times 10^{11}$	$3.87 \times 10^{11}$
SUT #2	35.52	$6.17 \times 10^{11}$	$4.04 \times 10^{11}$
SUT #3	37.20	$4.97 \times 10^{11}$	$3.38 \times 10^{11}$
SUT #4	39.96	$5.89 \times 10^{11}$	$4.03 \times 10^{11}$
SUT #5	39.50	$5.66 \times 10^{11}$	$4.05 \times 10^{11}$
SUT #6	37.19	$5.92 \times 10^{11}$	$4.05 \times 10^{11}$

table, we did not identify a clear and consistent distinction between both benchmarks. Hence, for classifying and analyzing the propagation of the reported errors, we combined the results of both benchmarks. Some fluctuations in the number of errors per SUT were observed, mainly concerning SUT #3 and SUT #4, as expected of a stochastic phenomenon since SUTs had the same experimental conditions.

TABLE III  
ERROR DISTRIBUTION FOR TID LESS THAN 25 krad

Benchmark	SUT	# Errors <sup>1</sup>	Error XS <sup>1</sup> [ $\text{cm}^2/\text{device}$ ]
Coremark	SUT #1	13	$3.36 \times 10^{-11}$
	SUT #2	15	$3.71 \times 10^{-11}$
	SUT #3	7	$2.07 \times 10^{-11}$
Embench	SUT #4	7	$1.74 \times 10^{-11}$
	SUT #5	16	$3.95 \times 10^{-11}$
	SUT #6	18	$4.44 \times 10^{-11}$

<sup>1</sup> Analysis considering before device failure at 25 krad.

### C. Error Classification

We used HARV reporting features to classify the detected HFTEs. This classification is shown in Table IV, which reports the number of each type of HFTE occurrence (#HFTE) and its specific cross sections (HFTE XS). It is worth mentioning that the results presented in this table consider the experiment before the permanent damage caused by the TID, defined as 25 krad.

The most frequent HFTE is single-bit upset from the SDRAM memory, representing 57.9% of all HFTEs, while memory double-bit upsets account for 2.7% of the total. Although all single-bit upsets from the memory are corrected, these results show that the memory is the weakest structure of the SoC. Following that, several upsets were reported from the register file, mostly single-bit upsets, which account for 22.4% of all exceptions and are corrected only in the hardened HARV experiments.

TABLE II  
CROSS SECTION OF DEVICE FAILURES PER SUT.

SUT	# Device failures <sup>1</sup>	Failure XS <sup>1</sup> [cm <sup>2</sup> /device]
SUT #1	11	$1.85 \times 10^{-11}$
SUT #2	11	$1.78 \times 10^{-11}$
SUT #3	8	$1.61 \times 10^{-11}$
SUT #4	12	$2.04 \times 10^{-11}$
SUT #5	11	$1.94 \times 10^{-11}$
SUT #6	5	$8.44 \times 10^{-12}$

<sup>1</sup> Analysis considering accumulated fluence until unrecoverable failure.

Furthermore, the results show that the third most common HFTE is the timeout load access fault, which is detected by the bus master block. This HFTE is reported when the bus master attempts to access a peripheral, but the peripheral does not respond before a specified time limit, likely due to errors either in the bus interconnection structures or in the peripherals' controllers. When this HFTE is triggered, the HARV hardening structures attempt to recover access to the peripherals by resetting the entire bus, including the bus master, interconnects, and peripherals, through a specific reset signal.

Despite the most frequent HFTEs, some appeared only once during the experiment, namely the program counter single- and double-bit upset, instruction register double-bit upset, and load access faults. Out of those, single-bit upsets in the program counter are corrected for the hardened HARV configuration, and double-bit upsets and load access faults are the most critical because they are not correctable.

#### D. Error Propagation

In order to explore the effectiveness of the HARV-SoC's hardening structures, we analyze the propagation of each HFTE for each HARV-SoC configuration separately, as shown in Table V. This table presents the number of reports for each HFTE type (#HFTE) and the number of those that propagated to cause an execution error and did not finish correctly (#Propagated HFTEs). Although the HFTEs propagated to cause execution errors, these were still identified by the watchdog timer, which is always enabled and recovers the processor operation to continue the experiment. Furthermore, we present the cross section for HFTE propagations, which is normalized by the fluence during the execution of each HARV-SoC configuration.

Most of the propagated HFTEs were timeout load access faults, which were most likely non-recognized faults in unprotected structures that led to the peripherals being unresponsive. At times, the faulty state due to these HFTEs was actually recovered by the peripheral reset mechanism, which we left intentionally enabled for baseline configuration to improve its error reporting, which would simply leave the processor in a stuck state. Although these HFTEs did not always recover the processor operation, the application is still provided with valuable information that aids in recovering the processor operation through software strategies, such as rollbacks, task restarting, and soft resets.

The HFTEs from the register file were mostly single-bit upsets, which were corrected in the hardened HARV-SoC and ignored for the baseline. Out of those, only one propagated to cause an execution error, which in turn propagated to cause a timeout load access fault.

TABLE V  
ERROR PROPAGATION PER HFTEs FOR EACH HARV-SOC CONFIGURATION.

Configuration	HFTE Type	#HFTE <sup>1</sup>	#Propagated HFTEs <sup>1</sup>	HFTE Propagation XS <sup>1</sup> [cm <sup>2</sup> /device]
<b>Baseline</b>	Memory single-bit upset	21	0	-
	Register file single-bit upset	6	1	$6.99 \times 10^{-13}$
	Timeout load access fault	6	4	$2.80 \times 10^{-12}$
<b>Hardened</b>	Memory single-bit upset	23	0	-
	Register file single-bit upset	11	0	-
	Timeout load access fault	3	1	$6.28 \times 10^{-13}$
	Memory double-bit upset	2	0	-
	Program counter single-bit upset	1	0	-
	Instruction register double-bit upset	1	0	-
	Program counter double-bit upset	1	0	-
	Load access fault	1	0	-

<sup>1</sup> Analysis considering before device failure at 25 krad



TABLE IV  
REPORTED ERROR TYPES AND CROSS SECTIONS (ALL SUTs COMBINED).

HFTE Type	#HFTE <sup>1</sup>	HFTE XS <sup>1</sup> [cm <sup>2</sup> /device]
Memory single-bit upset	44	$1.88 \times 10^{-11}$
Register file single-bit upset	17	$7.25 \times 10^{-12}$
Timeout load access fault	9	$3.84 \times 10^{-12}$
Memory double-bit upset	2	$8.53 \times 10^{-13}$
Program counter double-bit upset	1	$4.27 \times 10^{-13}$
Instruction register double-bit upset	1	$4.27 \times 10^{-13}$
Program counter single-bit upset	1	$4.27 \times 10^{-13}$
Load access fault	1	$4.27 \times 10^{-13}$
Total	76	$3.24 \times 10^{-11}$

<sup>1</sup> Analysis considering before device failure at 25 krad.

Considering the HFTEs from the memory, all single-bit upsets were mitigated in both HARV-SoC configurations. The double-bit upsets, on the other hand, were not corrected, but they affected data that was not critical for the benchmark execution, which finished correctly. Similar behavior was observed for the least common HFTEs in the hardened configuration, in which the affected data was not critical for the benchmark execution.

As a result of the error propagation, the hardened HARV-SoC improved the reliability by reducing the number of propagations by approximately 16% based on the HFTE propagation cross section.

## V. DISCUSSION

In Subsection IV-A, we investigated device failures to enable proper SEE analysis since unrecoverable failures were observed at the end of the test campaign. Notably, we provided a hypothesis for the observed behavior, indicating the application code, stored in internal flash memory, as the root cause. The failure affected the read operation performed to fetch instructions for the processor in an intermittent manner since in some cases hard resets temporally restored the functionality. Despite that, this behavior was consistently observed for the SUTs after a certain dose, as shown in Fig. 3, leading to unpredictable execution as supported by Fig. 2. We suspect that the failure source is related to the CMOS logic, known to be sensitive to TID [21], in the path to access the memory cells and not the flash cells themselves, which are more robust to this effect, as shown by the manufacturer [15], [16]. To access this memory block, HARV uses an unprotected interface implemented in the device, leaving a margin for failures. Moreover, the controlling logic of the memory is not hardened by design, in which buffers, encoders/decoders, multiplexers, and many other structures are susceptible to dose effects.

We explored the execution of different benchmarks to extend the test stimuli and improve the characterization. Despite not finding a consistent distinction between both benchmarks, as discussed in Subsection IV-B, a difference is expected for distinct workloads. In this case, both CoreMark and Embench have a broad set of algorithms with different processor and memory utilization profiles, minimizing variations.

Finally, the error classification and propagation analysis provide insights into the improvements achieved with the fault tolerance techniques implemented in HARV-SoC. In Table V, we can observe that the hardened configuration is highly capable of avoiding the propagation of errors since the hardened version is able to reduce the error propagation by 16%.

## VI. CONCLUSION

This work presented the effects of high-energy proton irradiation in a fault-tolerant RISC-V SoC, reporting a detailed error analysis and providing insightful discussions on the system's reliability. We discussed error propagation on the SoC, compared the baseline and hardened implementations, and identified the impact of accumulative effects on the SEE analysis and the device failure. The fault tolerance techniques applied to the SoC showed significant improvement in reliability for high-energy proton irradiation. Moreover, the available reporting capabilities of HARV-SoC greatly supported our analysis. Finally, the proton characterization extends the evaluated environments of HARV-SoC, enabling the proposition of applications targeting space systems. In future work, we intend to further explore the impact of using different benchmarks as test stimuli and improve the protection for the application code.

## REFERENCES

- [1] M. Yang, G. Hua, Y. Feng, and J. Gong, *Fault-tolerance techniques for spacecraft control computers*, 1st ed. Wiley Publishing, 2017.
- [2] A. Coronetti *et al.*, "Radiation hardness assurance through system-level testing: Risk acceptance, facility requirements, test methodology, and data exploitation," *IEEE Transactions on Nuclear Science*, vol. 68, no. 5, pp. 958–969, 2021, doi: 10.1109/TNS.2021.3061197.

- [3] T. Rajkowski, F. Saigné, and P.-X. Wang, "Radiation qualification by means of the system-level testing: Opportunities and limitations," *Electronics*, vol. 11, no. 3, 2022, doi: 10.3390/electronics11030378.
- [4] *Single Event Effects Test Method and Guidelines*, European Space Components Coordination (ESCC) Std. ESCC Basic Specification No. 25 100, Issue 2, 2014.
- [5] *Space product assurance: Radiation hardness assurance - EEE components*, European Cooperation for Space Standardization (ECSS) Std. ECSS-Q-ST-60-15C, 2012.
- [6] J. Budroweit and H. Patscheider, "Risk assessment for the use of COTS devices in space systems under consideration of radiation effects," *Electronics*, vol. 10, no. 9, 2021, doi: 10.3390/electronics10091008.
- [7] D. J. Sorin, "Fault tolerant computer architecture," *Synthesis Lectures on Computer Architecture*, vol. 4, no. 1, pp. 1–104, 2009.
- [8] "RISC-V: The free and open RISC instruction set architecture," RISC-V Foundation, 2023. [Online]. Available: <https://riscv.org/>
- [9] G. Furano, S. Di Mascio, A. Menicucci, and C. Monteleone, "A european roadmap to leverage RISC-V in space applications," in *2022 IEEE Aerospace Conference (AERO)*, 2022, pp. 1–7, doi: 10.1109/AERO53065.2022.9843361.
- [10] D. A. Santos, L. M. Luza, C. A. Zeferino, L. Dilillo, and D. R. Melo, "A low-cost fault-tolerant RISC-V processor for space systems," in *2020 15th Design Technology of Integrated Systems in Nanoscale Era (DTIS)*, 2020, pp. 1–5, doi: 10.1109/DTIS48698.2020.9081185.
- [11] D. A. Santos, A. M. P. Mattos, D. R. Melo, and L. Dilillo, "Enhancing fault awareness and reliability of a fault-tolerant RISC-V system-on-chip," *Electronics*, vol. 12, no. 12, 2023, doi: 10.3390/electronics12122557.
- [12] A. Gerbershagen *et al.*, "Establishment of the new particle therapy research center (PARTREC) at UMCG groningen," in *Proc. 15th International Conference on Heavy Ion Accelerator Technology (HIAT'22)*, ser. International Conference on Heavy Ion Accelerator Technology, no. 15. JACoW Publishing, Geneva, Switzerland, 08 2022, pp. 20–23.
- [13] *IGLOO2 and SmartFusion2 Datasheet, DS0128*, Microsemi, 2018, rev. 12.
- [14] *W9864G6JT Datasheet*, Winbond, 2017, rev. A03.
- [15] *Single Event Effects - A Comparison of Configuration Upsets and Data Upsets, WP0203 White Paper*, Microsemi, 2015, rev. 1.
- [16] *IGLOO2 and SmartFusion2 65nm Commercial Flash FPGAs, Interim Summary of Radiation Test Results*, Microsemi, 2014, rev. 2.
- [17] D. Dsilva, J.-J. Wang, N. Rezzak, and N. Jat, "Neutron SEE testing of the 65nm SmartFusion2 flash-based FPGA," in *2015 IEEE Radiation Effects Data Workshop (REDW)*, 2015, pp. 1–5, doi: 10.1109/REDW.2015.7336722.
- [18] S. Gal-On and M. Levy, *Exploring CoreMark a benchmark maximizing simplicity and efficacy*, Embedded Microprocessor Benchmark Consortium, 2012.
- [19] J. Bennett, C. Garlati, G. S. Madhusudan, T. Mudge, and D. Patterson, "Embench: An evolving benchmark suite for embedded IoT computers from an academic-industrial cooperative," in *RISC-V Workshop Zurich Proceedings*, 2019.
- [20] H. Quinn *et al.*, "Using benchmarks for radiation testing of microprocessors and FPGAs," *IEEE Transactions on Nuclear Science*, vol. 62, no. 6, pp. 2547–2554, Dec 2015, doi: 10.1109/TNS.2015.2498313.
- [21] M. Sajid *et al.*, "Analysis of total ionizing dose effects for highly scaled CMOS devices in low earth orbit," *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 428, pp. 30–37, 2018, doi: 10.1016/j.nimb.2018.05.014.