



Delft University of Technology

Proton Single Event Upsets Characterization of the NOEL-V Processor on the Xilinx Kintex Ultrascale FPGA

Hendrix, Tom; Mascio, Stefano Di; Menicucci, Alessandra

DOI

[10.1109/RADECS59798.2023.10752848](https://doi.org/10.1109/RADECS59798.2023.10752848)

Publication date

2023

Document Version

Final published version

Published in

2023 RADECS Data Workshop, RADECS 2023

Citation (APA)

Hendrix, T., Mascio, S. D., & Menicucci, A. (2023). Proton Single Event Upsets Characterization of the NOEL-V Processor on the Xilinx Kintex Ultrascale FPGA. In *2023 RADECS Data Workshop, RADECS 2023* (2023 RADECS Data Workshop, RADECS 2023). IEEE.
<https://doi.org/10.1109/RADECS59798.2023.10752848>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Proton Single Event Upsets Characterization of the NOEL-V Processor on the Xilinx Kintex Ultrascale FPGA

Tom Hendrix, Stefano Di Mascio, and Alessandra Menicucci

Abstract—This paper evaluates the Single Event Upset (SEU) susceptibility of the NOEL-V processor, a novel and highly modular Intellectual Property (IP) Core by Cobham Gaisler on the Xilinx Kintex Ultrascale SRAM FPGA. The processor is based on the promising RISC-V architecture, an open source Instruction Set Architecture (ISA) that is quickly rising in popularity.

In order to characterize the performance of the NOEL-V IP Core in the space radiation environment, the KCU105 development board is used as Device Under Test (DUT) and irradiated with medium and high energy protons. Thanks to the NOEL-V configurability, several versions of the NOEL-V were tested and microarchitectural differences could be exposed.

The biggest influence on user logic upsets is observed to be related to the use of an operating system. For a single-core, high-performance configuration, a foreseen in-orbit failure rate of one failure every 395 days is found for for a 51.6° circular orbit at 420 km altitude.

Findings indicate that the NOEL-V processor, with the implementation of targeted fault tolerant measures, can be a viable choice for space missions even as soft-core in SRAM FPGA. Due to its modularity, the processor can be used for a multitude of mission types ranging from high performance general-purpose to low-end microcontroller applications. Error Detection And Correction, which is not available in open source versions, will be needed to protect user memory and make sure upsets in caches and Configuration RAM (CRAM) do not lead to a failure of the processor.

Index Terms—FPGA, Protons, RISC-V, Single Event Effects

I. INTRODUCTION

Nominal operation of spacecraft is ensured by numerous integrated circuits (IC). These ICs have similar functions when used in terrestrial applications, but with the downside of having to deal with the hostile environment in space. Typical satellites employ between 47 and 400 ICs, with missions like Inmarsat2 having in excess of 1700 ICs [3]. These ICs are a mix of Application Specific Integrated Circuits (ASIC) and Field Programmable Gate Arrays (FPGA), with the trend in the past years seeing an increase of the latter over the first.

Among FPGA, SRAM-based FPGAs are very attractive because of their flexibility, performances and wide adoption in terrestrial applications. Examples of IP Cores already popular in space applications and that could be used on a similar FPGAs include the LEON3 and LEON4 cores by Cobham Gaisler, based on the SPARC ISA. The SPARC ISA has been

the most used ISA by the European Space Agency in recent years, but due to lost momentum of this ISA in terrestrial applications the way for a new ISA to be adopted by ESA is paved [5]. RISC-V offers the same openness as SPARC, with added benefits of modularity, compact code and larger address size (e.g. 64-bit) [4]. As opposed to SPARC, RISC-V shows momentum by wide adoption by academia and backing of big commercial companies [5, 6]. A new RISC-V based IP Core, NOEL-V, was recently released by Cobham Gaisler.

This paper evaluates the NOEL-V IP core implemented on a Xilinx Kintex FPGA. In order to characterize the performance of a NOEL-V processor in the harsh space radiation environment, the KCU105 development board is used as Device Under Test (DUT) and tested with a high energy proton beam. This high-performance FPGA also exist in a space grade version, with the possibility to have computing performances, in the near future, orders of magnitude better than any space grade processor.

II. METHODOLOGY

The objective of this research is to get a general overview of the susceptibility to radiation of the the NOEL-V processor, by means of using the modularity of the processor and software layers. The NOEL-V configurations employed are described in Tab. I.

TABLE I: Configurations employed during tests [7]. "L2C" stands for L2 Cache, "w" for ways.

Configuration	RISC-V subset	L1 Cache	MMU	PMP	L2C	FPU
Tiny	IM	1 KiB, 1 w	No	No	No	No
Minimal (EX1)	IMA	8 KiB, 2 w	No	Yes	Yes	No
GPP (Single)	IMAFD	16 KiB, 4 w	Yes	Yes	No	Yes
GPP (Dual)	IMAFD	16 KiB, 4 w	Yes	Yes	No	Yes
HPP (EX2)	IMAFD	16 KiB, 4 w	Yes	Yes	Yes	Yes

A. Software

The processor has been be tested while running benchmarks representative of the final application. Most space systems perform multiple tasks and are therefore usually running a Real-Time Operating System (RTOS). For this reason, the selected benchmark runs both on a RTOS and bare-metal. RTEMS, an RTOS provided with the NOEL-V processor, has been employed.

Tom Hendrix, Stefano Di Mascio and Alessandra Menicucci are with Department of Space Engineering, Delft University of Technology. Emails: T.J.B.Hendrix@student.tudelft.nl, s.dimascio@tudelft.nl and a.menicucci@tudelft.nl

B. Failure Classification

A number of tests have been performed and many distinct failure modes have been observed. To distinguish these failure modes, a grouping and naming is made, dividing the failures into Silent Data Corruption (SDC), Safe Failure (SF), Unsafe Failure (UF) and Fatal Failure (FF).

C. Test setup

The proton tests are conducted at HollandPTC¹ in Delft. As the development board also provides cooling for the FPGA, a fan is mounted on top of the FPGA. The influence of this fan on the proton beam is unknown and would have to be simulated. Therefore, it has been chosen to irradiate the backside of the FPGA. This is possible as the FPGA used is interconnected using the flip-chip method. Taking off the fan was no viable option as the temperature of the FPGA would become excessive for the maximum rated temperature. The FPGA is irradiated at a normal angle of incidence with the room kept at room temperature. Collimators are used to bring the beam area back to the size of the FPGA, an area of $4 \times 4 \text{ cm}^2$. This also means that the DDR4 memory is not in the beam path, avoiding upsets in main memory. Mounting of the development board at the right position of the beam is ensured by vertical and horizontal lasers. These lasers show the exact middle of the proton beam, while also ensuring the FPGA is horizontal and not at an angle. The test setup is shown in Fig. 1 and 2.




Fig. 1: KCU105 board clamped in place for the test.

The DUT is connected to a laptop in the radiation room. This laptop in turn is connected to a laptop in the control room via TeamViewer. Using this set-up, the KCU105 board could be controlled by the researcher at all times.

This tests were executed according to the test procedure shown in Fig. 3.

D. Choosing fluence

At the facility three beam energies were available, namely 100, 150 and 250 MeV. This translated to 70, 120 and 220 MeV at the DUT due to the effect of the air path to the device

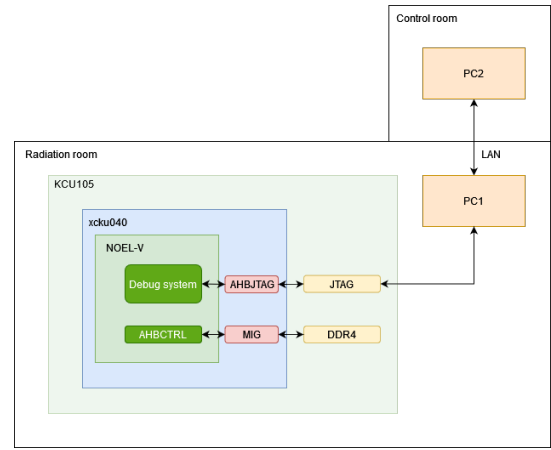


Fig. 2: Schematic representation of the setup in the radiation room and how two PC's are used for communication.

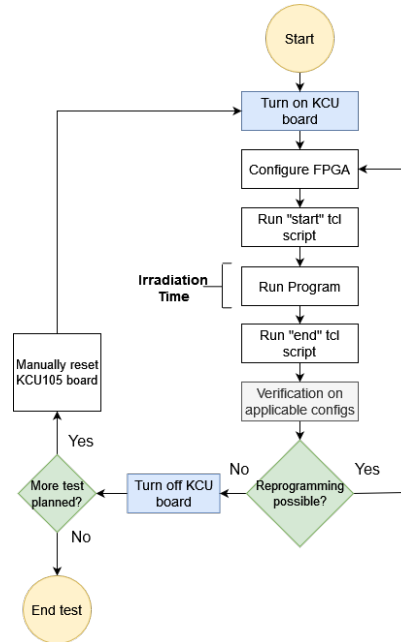


Fig. 3: General test flow for all tests.

and collimators. Having only 3 (high) energies available means that it might not be possible to construct a energy-cross section graph if all three energies are above the 'knee' region.

The beam area used is $4 \times 4 \text{ cm}^2$. This area was chosen as it coincides almost perfectly with the FPGA area. As the beam area is not easy to change within tests, the beam area was kept constant throughout all tests.

The beam flux, being one of three possible values, is chosen empirically. Error rates were observed for multiple levels of the flux. First starting off at the lowest flux available at $2 \cdot 10^6 \text{ p/cm}^2 \cdot \text{s}$, after which the flux is increased until a desired error rate was found. The highest energy available is used to perform these flux calibration tests to limit the accumulation of radiation damage [2].

¹<https://www.hollandptc.nl/>

E. Error metric definition

Vivado functionality is utilized to extract the amount of errors in the CRAM memory.² Specifically this is the Vivado readback functionality that compares a bitstream file that is readback from the FPGA with a mask file, containing an exact copy of the generated bitfile before it was send to the FPGA. In order to obtain the FEC (Functional Error Cross section) operations of a certain program are monitored in real-time and the beam is turned off. The FEC is thus the inverse of the fluence to error.

III. TEST RESULTS

During the three tests a combined test time of 174 minutes is reached, in which $2.14 \cdot 10^{11} \text{ p/cm}^2$ were fired at the DUT.

A. Test 1

The FPGA was irradiated over a total time of 59 minutes (beam on time). During the test this meant a total fluence of $7.93 \cdot 10^{10} \text{ p/cm}^2$ was reached. Test one started off with an initial empirical test to determine the optimal flux level to use. This level was found to be $5 \cdot 10^7 \text{ p/cm}^2 \cdot \text{s}$.

B. Test 2

The total time tested during Test 2 was 58.5 minutes, for a total fluence of $8.93 \cdot 10^{10} \text{ p/cm}^2$, at a flux of around $2 \cdot 10^7 \text{ p/cm}^2 \cdot \text{s}$. Set lower than the optimal flux due to beam limitations. The main results of these tests are described below.

1) *Total CRAM error rates:* As for the first test, CRAM error rates are reported for all examples applicable. In this test this was the case for all but the first 10 runs.

By dividing the number of errors reported in the CRAM memory by the fluence for the tests, an average cross section of $2.56 \cdot 10^{-16} \text{ cm}^2/\text{bit}$ is found, the individual cross sections for the runs are shown in Fig. 4, where it can be seen that the cross section is relatively constant between runs.




Fig. 4: Graphical depiction of CRAM cross section during test 2.

During the first test, the Dhrystone benchmark was used on each of the 19 runs. During the second test for 14 runs.

Comparison of FEC for runs running different benchmarks is plotted in Fig. 5. For initial comparison the runs of the first test and second test are kept separated, to keep into account different energies.




Fig. 5: FEC for different benchmarks.

Absolute values are shown in Tab. II. For the absolute value comparison the Dhrystone runs are joined together, warranted by the small differences observed during the energy and mitigation tests. Dhrystone runs are taken as the reference for ratio calculation.

TABLE II: Absolute values of the FEC and CRAM cross sections for different benchmarks.

	FEC (cm^2)	CRAM (cm^2/bit)
Dhrystone	$6.37 \cdot 10^{-10}$	$2.72 \cdot 10^{-16}$
CoreMark	$3.92 \cdot 10^{-10}$	$2.82 \cdot 10^{-16}$
Ratio	1.63x	0.96x

C. Test 3

Effective proton beam time of test 3 was 56.5 minutes, accounting for for a total fluence of $4.51 \cdot 10^{10} \text{ p/cm}^2$, at the same flux level as set for test 2.

1) *Cache:* 459 s of irradiation was performed on the processor when all caches were disabled, at the respective flux levels this means a subjected fluence of $6.11 \cdot 10^9 \text{ p/cm}^2$. During this time, no upsets occurred within the iCache and L2Cache. One bit flip occurred in the dCache, where a 1 was flipped to a 0.

Taking into account the sizes of the caches, being 32 KiB (instruction and data) and 256 KiB for the L1cache and L2cache respectively, the error rates can be estimated to be $2.50 \cdot 10^{-15} \text{ err/bit} \cdot \text{s}$ and $< 2.50 \cdot 10^{-15} \text{ err/bit} \cdot \text{s}$ respectively.

2) *Influence of operating system:* The influence of the employment of RTEMS is shown in Fig. 6.

In absolute values the FEC running bare-metal CoreMark was $1.18 \cdot 10^{-9} \text{ p/cm}^2$ and when running CoreMark in RTEMS was $0.39 \cdot 10^{-9} \text{ p/cm}^2$, showing a 3x improvement in cross section by not using an operating system.

3) *Impact of the software:* Fig. 7 shows the different FEC for an integer benchmark and a benchmark operating on floating point elements.

4) *Failure modes:* A visual depiction of the failure modes occurring during the third test are depicted in Fig. 8.

When RTEMS is not employed, none of the failures are deemed safe. In the case of the OS, failure information is

²<https://docs.xilinx.com/r/2021.1-English/ug908-vivado-programming-debugging/Bitstream-Verify-and-Readback>




Fig. 6: Influence of the use of RTEMS.

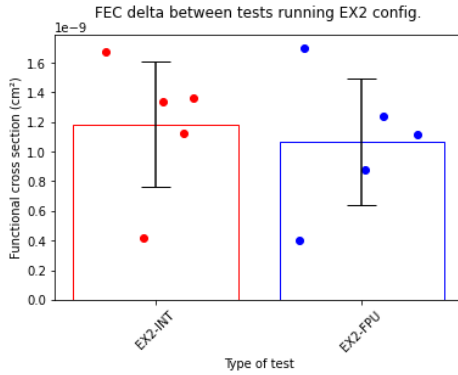


Fig. 7: Functional error cross section for INT and FPU software running on the EX2 configuration

shown and the execution is halted purposefully by the processor itself, this is not the case for the bare metal implementation. This is a big advantage when using RTEMS.

For Test 3, as the software was now executing as part of an OS, different failure types are identified from the earlier tests.

5) *Cross sections of different configurations*: Different configurations have different functional cross section, as shown in Fig. 9.

The FEC has been scaled by the area of the respective configuration, taking into account the expected effect increased resource usage would have, as shown in Fig. 11 & 10.

IV. CONCLUSIONS

The reduced susceptibility of the Floating Point Unit (FPU) compared to the integer unit states the findings in other research and has expanded the conclusion that not only is the Floating Point Register less susceptible than the General Purpose Register (because of a lower utilization), but also the FPU as a whole shows decreased functional error cross section.

The inclusion of an L2Cache lead to higher susceptibility due to the increase in vulnerable memory. The best configuration depends on mission requirements, but this research has proven that radiation susceptibility only increases slightly with processor resource usage, or not at all for some configurations like EX2. In general, the radiation susceptibility between configurations does not differ much with microarchitectural differences, all being in the same order of magnitude.




Fig. 8: Failure modes grouped.

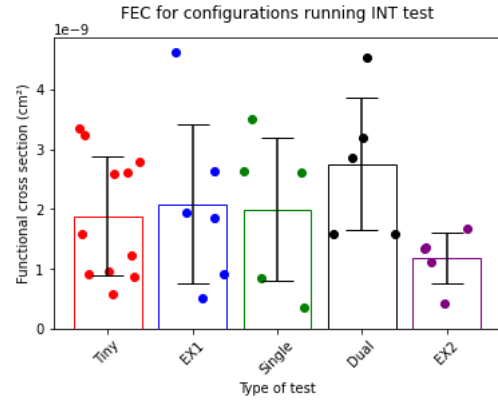


Fig. 9: Functional error cross section for all configurations running INT software

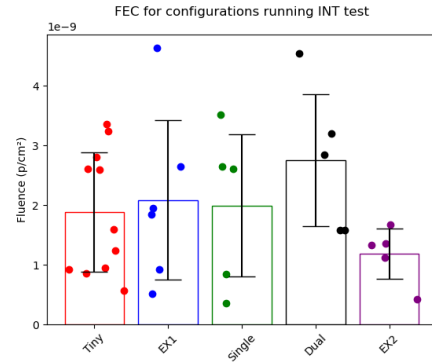


Fig. 10: FEC for all configurations running INT software, absolute.

Using the earlier described FOM method [1] it can be shown that the NOEL-V processor can be made suitable for in-space operations. Cache upsets are shown to occur but with the use of Error Detection And Correction (EDAC) methods such upsets can be avoided. Since no Multiple Bit Upsets (MBUs) were observed, EDAC methods with small overhead can be used.

The processor is promising for usage in space, however some mitigation measures should be taken to ensure this. The use of fault tolerant techniques will have to be investigated further.

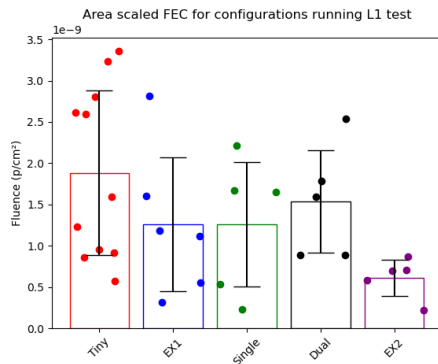


Fig. 11: FEC for all configurations running INT software, scaled by area overhead.

REFERENCES

- [1] E. L. Petersen. “The SEU figure of merit and proton upset rate calculations”. In: *IEEE Transactions on Nuclear Science* 45.6 (1998), pp. 2550–2562.
- [2] JEDEC Government Liaison Committee et al. *TEST STANDARD FOR THE MEASUREMENT OF PROTON RADIATION SINGLE EVENT EFFECTS IN ELECTRONIC DEVICES*. JESD234. 2013.
- [3] A Fernández León. “Trends and patterns of ASIC and FPGA use in European space missions”. PhD thesis. MS thesis, Delft Univ. of Technol., Delft, Netherlands, 2013.[Online ... , 2013.
- [4] Krste Asanović and David A Patterson. “Instruction sets should be free: The case for risc-v”. In: *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2014-146* (2014).
- [5] Stefano Di Mascio et al. “The case for RISC-V in space”. In: *International Conference on Applications in Electronics Pervading Industry, Environment and Society*. Springer. 2018, pp. 319–325.
- [6] Stefano Di Mascio et al. “Leveraging the Openness and Modularity of RISC-V in Space”. In: *Journal of Aerospace Information Systems* 16.11 (2019), pp. 454–472.
- [7] Cobham Gailser AB. *GRLIB IP Core User’s Manual*. 2021.