



Original Article

Using machine learning for anomaly detection on a system-on-chip under gamma radiation[☆]Eduardo Weber Wächter^{a, b}, Server Kasap^c, Şefki Kolozali^a, Xiaojun Zhai^{a, *}, Shoaib Ehsan^a, Klaus D. McDonald-Maier^a^a School of Computer Science and Electronic Engineering, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK^b School of Engineering, University of Warwick, Coventry, CV4 7AL, UK^c School of Computing, Electronics and Mathematics, Coventry University, Coventry, CV1 5FB, UK

ARTICLE INFO

Article history:

Received 22 March 2022

Received in revised form

21 June 2022

Accepted 26 June 2022

Available online 30 June 2022

Keywords:

Gamma radiation

Machine learning

Anomaly detection

Field programmable gate arrays. TID

ABSTRACT

The emergence of new nanoscale technologies has imposed significant challenges to designing reliable electronic systems in radiation environments. A few types of radiation like Total Ionizing Dose (TID) can cause permanent damages on such nanoscale electronic devices, and current state-of-the-art technologies to tackle TID make use of expensive radiation-hardened devices. This paper focuses on a novel and different approach: using machine learning algorithms on consumer electronic level Field Programmable Gate Arrays (FPGAs) to tackle TID effects and monitor them to replace before they stop working. This condition has a research challenge to anticipate when the board results in a total failure due to TID effects. We observed internal measurements of FPGA boards under gamma radiation and used three different anomaly detection machine learning (ML) algorithms to detect anomalies in the sensor measurements in a gamma-radiated environment. The statistical results show a highly significant relationship between the gamma radiation exposure levels and the board measurements. Moreover, our anomaly detection results have shown that a One-Class SVM with Radial Basis Function Kernel has an average recall score of 0.95. Also, all anomalies can be detected before the boards are entirely inoperative, i.e. voltages drop to zero and confirmed with a sanity check.

© 2022 Korean Nuclear Society, Published by Elsevier Korea LLC. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

One of the biggest challenges in the European Union is the cleaning of nuclear waste [1]. This task involves handling and moving extreme toxic material contaminated with different types of ionizing radiation. Unfortunately, higher doses of radiation harm the human body; therefore, the cleaning process should be taken with much precaution. Luckily, most of this work can be done by robots, but sadly the electronic devices controlling them are also susceptible to radiation.

[☆] This work is supported by the U.K. Engineering and Physical Sciences Research Council through grants EP/R02572X/1, EP/P017487/1, EP/V000462/1, and EP/V034111/1.

* Corresponding author. School of Computer Science and Electronic Engineering, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK.

E-mail addresses: eduardo.wachter@warwick.ac.uk (E.W. Wächter), server.kasap@coventry.ac.uk (S. Kasap), sefki.kolozali@essex.ac.uk (Kolozali), xzhai@essex.ac.uk (X. Zhai), sehsan@essex.ac.uk (S. Ehsan), kdm@essex.ac.uk (K.D. McDonald-Maier).

The radiation consequences on contaminated sites can be generally categorized into two broad types: permanent and transient effects. One example of the permanent damage is caused by Total Ionizing Dose (TID), while Single Event Upset (SEU) is a transient example. TID is a phenomenon that causes permanent damage, which inevitably, at some point, is going to result in total failure of the electronics. TID effects can only be minimized to extend the system life. The literature has come up with special hardened devices specifically targeting this effect [2,3]. Typically, these devices are much more expensive than unhardened ones. On the other hand, SEUs causes **transient** bit-flips on memory elements. For example, one of the most employed solutions for radiation environments, especially for space applications, is Field Programmable Gate Arrays (FPGAs) due to their reconfiguration capability and performance aspects. FPGAs are predominantly used for these applications [4–6] as they allow reconfiguration and thus hardware adaptability to deal with most transient effects.

The underlying idea of this paper is that instead of using radiation-hardened devices, we adopt consumer electronic level

COTS (commercial off-the-shelf) devices and then monitor, manage and replace them with healthy ones before they stop working. In [7,8], authors attempted to use commercial off the shelf micro-controller in gamma radiation environment, and a number of testings and results were reported. From these results, on-board voltage regulators were very sensitive to the effect of gamma radiation, this condition brings us a research challenge to anticipate when will the board results in a total failure due to radiation, specifically TID effects.

In this paper, we first employ a statistical analysis of the collected data to prove that the measured values of the board (e.g. voltage, temperature and current) under radiation were significantly different compared to a scenario compared to a scenario without radiation. Later, the paper tries to answer whether it is possible to predict when the board will stop working due to radiation with reasonable accuracy. For this purpose, we tested a low-cost COTS unhardened consumer electronic level 28 nm FPGA used in many fault-tolerant techniques for transient radiation effects. We employ this board under gamma (γ) radiation and log its behaviour.

The monitoring of these boards needs to be executed with an enhanced approach. The paper shows that using simple techniques to observe voltage and temperatures may not indicate that the board will stop working. For example, one widely used data analysis tool is R control charts. This type of chart, popularly known as a control chart, monitors the mean and range of normally distributed variables simultaneously when samples are collected at regular intervals. Such a technique uses upper and lower control limits (UCL and LCL) to monitor the behaviour of the variables. Still, it may not be sufficient to indicate that a given board is behaving abnormally. A few data points outside the operational limits caused by radiation do not suggest that the board will stop operating. This paper shows that this might take a few minutes or more than 1 h. For this reason, we employ state-of-the-art machine learning (ML) techniques to understand and try to predict when the board is behaving abnormally. The advantage of ML algorithms is that it can fine tune and predict failures with no need of expert knowledge (i.e. use of UCL/LCL). However, the disadvantage of using such approach is it requires data, meaning, a few boards will be totally damaged because of radiation. Authors in [9] reviews works that used ML for failure prediction in industrial mechanical systems for the last decade and identify opportunities for future research, although there is no focus on electronics.

The novelty of this works is three-folded:

- The first study to monitor and measure voltages and temperatures on a consumer electronic level SRAM-based FPGA SoC (System on Chip) under γ radiation.
- This is the first study performing a quantitative/statistical analysis of the effects of gamma radiation on voltages and temperatures of an FPGA and then compare to an environment without radiation.
- First work using machine learning algorithms trying to predict when the board will be rendered in-operational by gamma radiation through the observation of temperature and voltage values only.

The paper is organized as follows. Section 2 reviews the state-of-the-art of related techniques. Section 3 details the hardware setup to monitor and test the boards under experiment, later presenting an example where the proposed technique might be used. Section 4 explains the statistical analysis while Section 5 details the employed machine learning techniques and how to measure its accuracy. Section 6 details the experiments carried out under γ radiation and how the data was organized to feed the machine learning algorithms. Finally, Section 7 evaluates, compares and

discuss the results, and Section 8 draws conclusions and future plans.

2. Related work

As the sophistication of embedded systems grows, their vulnerability to errors is adversely affected due to an increase in critical points of failure. Adopting fault mitigation or fault tolerance techniques is vital if FPGAs are used in radiation environments. Fault tolerance techniques that enhance embedded processor reliability can be categorized as hardware-, software- and hybrid-based techniques [10,11].

The hardware-based techniques, which mainly rely on spatial redundancy, provide two or more instances of a hardware component, such as processors, memories, buses or power supplies, for protection against soft errors. This class of techniques include Triple Modular Redundancy (TMR) [12], Duplication with Comparison (DWC) [13] and hardware monitors [14] which incorporate watchdog or checker modules to monitor the system and detect errors by verifying the control-flow related memory accesses of the target processor. These techniques can protect the system from errors in the computation outputs, i.e. silent data corruption (SDC), as exemplified in [15].

Software-implemented hardware fault tolerance (SIHFT) approaches handle hardware malfunctions by merely shielding the software without any hardware alteration. These techniques rely on adding redundant software code for comparison to detect errors. However, they exhibit a high-performance overhead, which may not be viable for some real-time systems. These kinds of techniques, such as ABFT [16], HETA [17] and S-SETA [18], detect control-flow faults leading to FIs, which manifest themselves as hangs or crashes, and then place the system into a fail-safe state. Note that both hardware- and software-based techniques are not capable of correcting 100% of the errors, but rather detecting them to avoid a failure that would have adverse effects on the entire mission. Furthermore, they protect the system either from SDCs or functional interrupts (FI), but not both.

The hybrid techniques are the ones that use a SIHFT method combined with a hardware intellectual property (IP), which performs consistency checks in the processor, making them effective against both SDCs and FIs. For instance, the lockstep technique is a hybrid fault-tolerance technique based on software and hardware redundancy. It employs the concepts of checkpointing and recovery mechanisms (e.g. roll-back recovery, roll-forward recovery) at the software level, and processor replication and checker circuits at the hardware level, as explained in the following sections. Therefore, it is capable of both error detection and correction. The lockstep technique's most significant merit is its ability to detect and correct both SDCs and FIs, unlike many other fault tolerance techniques. Several researchers have developed and implemented their lockstep technique version, such as those in [19–23], to make a range of processors resistant to radiation-induced soft errors, extensively analyzed and compared in [24].

The authors in [25] evaluated a (^{60}C) gamma-ray radiation testing of a space application FPGA, namely the RT4G150 from Microsemi. Microsemi is a manufacturer of high-reliability FPGAs for space applications, while RTG4 is the 4th generation family of radiation-tolerant flash-based FPGAs. The work assesses the degradation of the flash cell through its threshold-voltage (V_T) shift. For space applications, dynamic burn-in (DBI) testing is used to evaluate the long term reliability of the device. Among all product screening tests employed by many business categories, including automotive, aerospace, and defence, the burn-in (BI) test is one of the most effective tests for early failure detection. The work indicates that RTG4 shows a shift of the programmed Pflash cell V_T

post-DBI is observed. The programmed Pflash VTshift is due to voltage degradation, resulting from approximately 1.75% degradation of the DAC's output. It is essential to emphasize that [25] does not monitor the temperature of the FPGA, this being out of the scope of the project.

In [26] Authors evaluated a COTS FPGA, namely Microsemi ProASIC3E A3PE1500. Despite their low reliability, the authors state that this FPGA has been considered a promising alternative to replace radiation-hardened ones. The study analyses the Single-Event Upset (SEU) sensitivity of the FPGA for a combined set of Electromagnetic Interference (EMI) and TID tests. This component was under the pre-qualification process for use in some satellites of the Brazilian Space Program. The TID test was performed by exposing the FPGA to a 10-keV effective energy X-ray beam. The device was roughly exposed to the TID expected to be cumulated on satellite electronics after operation for a period of 4 or 5 years in a given orbit, as specified by the Brazilian National Institute for Space Research. The conclusion is that flip-flops (FFs) present a lower SEU-immunity degradation when exposed to conducted-EMI (5.6%) than SRAM cells (resp. 6.3%); the latter memory elements are intrinsically more robust to EMI since they present a much lower cross-section than FFs. Note that this work does not monitor the temperature of the FPGA as well.

Tarrilho et al. [27] analyzed the behaviour of flash-based FPGA from Actel under TID. In this work, a design with an embedded system composed of a MIPS microprocessor hardened with fault-tolerant techniques is employed on a COTS flash-based FPGA ProASIC3E family from Actel. The TID experiment with no reconfiguration monitored the power supply current during radiation and the FPGA temperature. They reported that the power supply current started to change after 45 krad(Si) when some modules stopped working. The current increases promptly and reaches 1.5 times the original current just before 65 krad(Si). The temperature and current drop abruptly when most modules fail around 65 krad(Si). The study primarily showed the failure dose for some internal modules and did not present when the current starts behaving abnormally.

In [28] evaluate TID effects on an SRAM-based COTS FPGA. They combine hardware and software techniques to perform on-chip irradiation via a $^{90}\text{Sr}/^{90}\text{Y}$ electron source and assess the degradation of the system. The experiment consists of two executions of Zynq XC7Z020T chips hosted by two distinct Zedboards. The analysis focuses on ring oscillators (ROs) implemented on the programmable logic of the FPGA for estimating/predicting the performance degradation due to the TID effects. The authors show specifically the RO frequency according to temperature, current and accumulated dose. The authors indicate that specifically for COTS 28-nm Zynq-7000 chips, the results show increased TID tolerance beyond the Mrad level. The tests were conducted with the FPGAs surviving up to 2 Mrad(Si). In summary, this work shows that COTS FPGAs can survive significant amounts of ionizing radiation, although not displaying the limit, since it was not the objective of their work.

These papers [27,26,25] summarize most related work for TID effects in FPGAs. They either focus on flash-based devices [27,25] or point out [26] that they could be a promising alternative for radiation-hardened ones. Also, only one paper monitors current and temperature, although it is not the main objective of that paper to correlate or predict the FPGA stop working.

3. Experimental design

This section explains the setup, shown in Fig. 1, built to measure and log all the tested experiments, later detailed in Section 6. The experimental setup is composed of (i) a laptop to control, collect

and log all data; (ii) a *Design Under Test (DUT) board* and (iii) a special *monitoring board* to collect the data from it. The only electronic device exposed to radiation is the DUT, that being the reason for the monitoring board, which is outside the radiation cavity and connected to the DUT with voltage and temperature probes.

The DUT is a MiniZed board, depicted in Fig. 2, a widely available and, most used, equipped with a low-cost ($\approx \text{€}100$) Xilinx 28 nm Zynq FPGA. The monitoring board is an Arduino board coupled with temperature and voltage sensors, which during the experiments are located outside the radiation chamber; therefore, it can be reused (see Fig. 3).

The MiniZed board [29] (Fig. 2) is a development board containing a Xilinx Zynq single-core SoC XC7Z007S-1CLG225C [30] with 512 MB DDR3L micron storage, a 128MB QSPI flash and 8 GB eMMC. The FPGA contains a programmable logic and a Programmable Logic (PL) (FPGA) and Processing Subsystem (PS) (ARM Cortex-A9). The FPGA chip is fed by a power management integrated circuit (PMIC) Dialog DA9062 [31] which controls five different voltages [32]: 0.675, 1.0, 1.35, 1.8 and 3.3 V, detailed in Table 1.

For example, V_{DDR3} supplies the voltage to the DDR3 memory of the board. The monitoring board is organized to monitor these voltages. We also use a thermocouple wire to monitor the temperature on the surface of both FPGA and PMIC. This setup is arranged so that only the DUT and wires connecting to the monitoring board are inside the radiation chamber; therefore, the radiation does not interfere with the monitoring electronics.

The monitoring board is an Elegoo R3 board [33], which is completely compatible with the official Arduino R3 version. The board is composed of an Atmel ATMEGA328P chip [34] and coupled with five DC 0–25 V voltage sensors and two Digilent Pmod TC1 K-Type thermocouple modules [35]. The thermocouple sensor module works with a temperature range of -73 °C – 482 °C with a 0.25 °C resolution. The voltage sensor works with a range of DC 0–25 V with a resolution of 0.01 V. The board collects one reading each second, operating in a 1 Hz frequency. The data is sent to a laptop and stored for later computation. In parallel, the MiniZed board runs a compute an intense set of operations whose outputs are sent to the host laptop through a serial interface for a sanity check aiming to determine if the board is still operational. When the board stops sending data, it is considered to be dead. Later, all the boards used on experiments were tested individually outside the radiation chamber to ensure they were inoperational. We tried to download the provided sanity check, and none of the boards responded to the cable connection attempts, confirming that the boards were dead.

3.1. Motivational example

Fig. 4 shows the measurements of the Minized board under normal operation. We monitor the power controlling IC and the FPGA temperature (T_{PMIC} and T_{FPGA} respectively) and five voltages (V_{ddr3} , V_{aux} , V_{core} , V_{tt} and V_{cco}) supplied to the FPGA. One can see a normal fluctuation in the temperature and voltages within some bounds. On the other hand, Fig. 5 shows the same type of board under gamma radiation, where the black lines on voltages are the average of the initial readings. The temperature starts to rise faster, and voltages operate outside their normal bounds and eventually, the board stops working. These two simple examples clearly show that the board is not behaving normally under radiation, as expected. Later sections try to correlate these behaviours to predict when the board will stop working. Furthermore, we use state-of-the-art machine learning techniques to evaluate whether it is possible to indicate when a board will fail based on voltages and temperature sensor readings.




Fig. 1. Block diagram of the experimental setup. Only the DUT is under radiation.




Fig. 2. Employed a Minized board with connectors attached to monitor voltage and two thermocouple pairs to measure temperatures.




Fig. 3. The monitoring board composed of an Elegoo R3 board (1) coupled with voltage (2) and temperature (3) sensors.

Table 1
Temperature and voltage monitored values.

Symbol	Expected value	Description
T_{FPGA}	-	FPGA temperature
T_{PMIC}	-	PMIC temperature
V_{core}	1 V	Core supply
V_{aux}	1.8 V	Auxiliary supply
V_{ddr3}	1.35 V	DDR3
V_{tt}	0.675 V	Vtt
V_{cco}	3.3 V	Vcco / board voltage




Fig. 4. Voltage and temperature sensor reading in an environment without radiation. Temperature and voltage values operate inside certain bounds, and there are few deviations from the mean.

4. Understanding the statistical effect of gamma radiation on Minized boards

To examine the effects of different levels of γ radiation on our boards, we conducted a statistical analysis using measurements for boards that stopped working after being exposed to γ radiation and

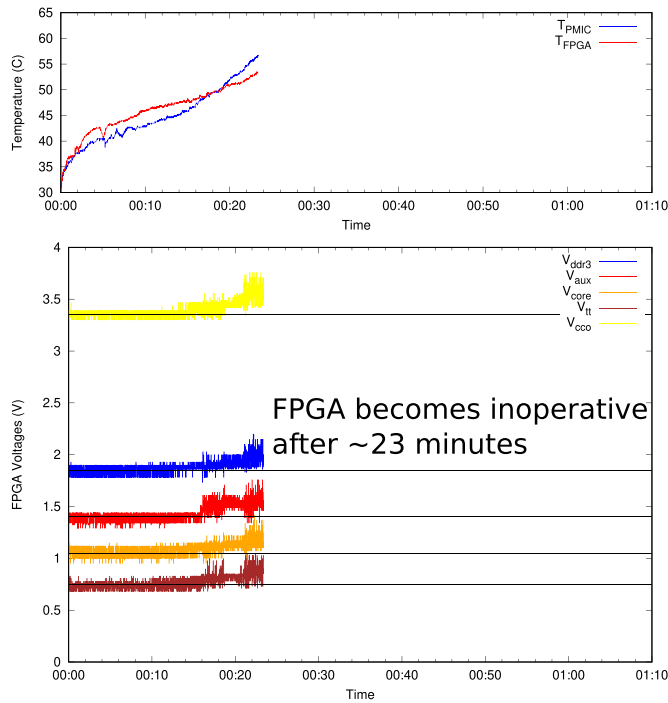


Fig. 5. Voltage and temperature sensor reading on an environment with γ radiation. Temperature and voltage values operate outside bounds compared to a mean without radiation.

for boards used in a non-radiation environment. To test whether any significant differences occurred in these measurements, we conducted one-way Multivariate Analyses of Variance (MANOVA). The first MANOVA test compared the γ radiation effects on the boards (i.e. 0: the measurements collected from boards working under γ radiation; 1: measurements obtained from boards on a non-radiation site) as independent variable, and sensor measurements, two temperatures (T_{PMIC} and T_{FPGA}) and five voltages (V_{ddr3} , V_{aux} , V_{core} , V_{tt} and V_{cco}) as dependent variables. The second MANOVA test involved different γ radiation levels (to which boards were exposed) as an independent variable (i.e. γ radiation levels detailed in Table 4) and temperature and voltage as dependent variables.

The statistical results are presented in Table 2 and Table 3. The partial eta squared (η^2) represents the effect size, determining how much the relationship will affect the values. On the other hand, the F-value is the test statistic used to determine how much one variable is associated with the response. The factors and interaction

Table 2

Results of One-Way Multivariate Analyses of Variance to discover the sensory observation difference between functioning and non-functioning sensor boards used in radiation and non-radiation environments. η^2 is the partial eta squared measure of effect size. * $p < .05$, ** $p < .01$, *** $p < .001$. The table demonstrates the statistical effect of the main factor. The error degrees of freedom was the same for each dependent variable.

Source	Dependent variable	df	F	η^2
γ Radiation Functioning	T_{PMIC}	1	33051.493***	.214
	T_{FPGA}	1	18569.785***	.133
	V_{aux}	1	1040594.28***	.896
	V_{ddr3}	1	463097.020***	.793
	V_{core}	1	30362.611***	.200
	V_{tt}	1	19258.336***	.137
	V_{cco}	1	87822.015***	.420
Error		121226		

Table 3

Results of One-Way Multivariate Analyses of Variance to discover the sensory observation difference for sensor boards at different gamma radiation levels. * $p < .05$, ** $p < .01$, *** $p < .001$. The table demonstrates the statistical effect of the main factor. The error degrees of freedom was the same for each dependent variable.

Source	Dependent variable	df	F	η^2
γ radiation	T_{PMIC}	6	9721.465***	.325
	T_{FPGA}	6	4066.381***	.168
	V_{aux}	6	197999.152***	.907
	V_{ddr3}	6	94615.951***	.824
	V_{core}	6	6617.718***	.247
	V_{tt}	6	4006.448***	.165
	V_{cco}	6	148.701***	.577
Error		121221		

Table 4

Radiation experiment for the DUT.

Exp. #	Radiation rate (Gy/h)	DUT time	Monitoring board time
0	1209	1:56:00	2:00:11
1	2469	0:39:14	0:41:31
2	5137	0:23:16	0:27:28
3	5871	0:19:49	0:23:23
4	7707	0:12:39	0:13:39
5	16966	0:06:58	0:07:38

effects were analyzed with one-way analysis using the partial eta squared index of effect size. The Bonferroni procedure was used here. The definitions in [36] have been adopted to discuss the effect sizes: small effect size ($\eta^2 \leq .01$), medium effect size ($.01 \leq \eta^2 \leq .06$) and large effect size ($.06 \leq \eta^2 \leq .14$). The MANOVA levels of significance are reported using the F-statistics and probability p . A risk of α of .05 was used in all statistical tests.

4.1. Comparison of boards behaviour on radiation and non-radiation sites

There was a highly significant effect of the functioning (*Functioning*) of the boards deployed in radiation site when compared to a non-radiation sites in Table 2. Overall, there was highly significant effect on the functioning of the sensor boards with very large effect sizes when the measurements collected from radiation and non-radiation sites: $F(1, 121226) = 33051.493, p < .001, \eta^2 = .214$ for T_{PMIC} ; $F(1, 121226) = 18569.785, p < .001, \eta^2 = .133$ for T_{FPGA} ; $F(1, 121226) = 1040594.28, p < .001, \eta^2 = .896$ for V_{aux} ; $F(1, 121226) = 463097.020, p < .001, \eta^2 = .793$ for V_{ddr3} ; $F(1, 121226) = 30362.611, p < .001, \eta^2 = .200$ for V_{core} ; $F(1, 121226) = 19258.336, p < .001, \eta^2 = .137$ for V_{tt} ; $F(1, 121226) = 87822.015, p < .001, \eta^2 = .420$ for V_{cco} .

Our results indicate a significant difference in the obtained measurements between the radiation and non-radiation sites. This was an expected result, but this is the first work reporting these findings to the best of the authors' knowledge. It is important to emphasize that since it justifies the underlying assumption of the proposed analysis of this paper.

4.2. The effect of different γ radiation levels

There was a highly significant effect of γ radiation levels (γ Radiation) with a very large effect sizes on the sensor board measurements in Table 3: $F(6, 121221) = 9721.465, p < .001, \eta^2 = .325$ for T_{PMIC} ; $F(6, 121221) = 4066.381, p < .001, \eta^2 = .168$ for T_{FPGA} ; $F(6, 121221) = 197999.152, p < .001, \eta^2 = .907$ for V_{aux} ; $F(6, 121226) = 94615.951, p < .001, \eta^2 = .824$ for V_{ddr3} ; $F(6, 121221) = 6617.718, p < .001, \eta^2 = .247$ for V_{core} ; $F(6,$

121221) = 4006.448, $p < .001$, $\eta^2 = .165$ for V_{tt} ; $F(6, 121221) = 27547.584$, $p < .001$, $\eta^2 = .577$ for V_{cco} .

The posthoc analysis showed highly significant differences between most sensors when exposed to different γ radiation levels $p < .001$. However, there were no significant differences in the following interactions. For T_{PMIC} measurements, there was no significant difference between 2469 and 7707 γ radiation levels. For T_{FPGA} , there was no significant difference between 0 and 5871 γ radiation levels, and there was only a significant difference between 7707 and 16966 γ radiation levels $p < .05$. For voltage 1, there was no significant difference between 2469 and 5137 and 5871 γ radiation levels. Similar to T_{FPGA} results, there was only a significant difference between 7707 and 16966 γ radiation levels for V_{aux} measurements. For V_{ddr3} measurements, there was a significant difference between 7707 and 16966 γ radiation levels ($p < .05$). For V_{core} measurements, there was no significant difference between 1209 and 7707 γ radiation levels and between 5137 and 5871 γ radiation levels. For V_{tt} measurements, there was no significant difference between 1209 and 7707 and 16966 γ radiation levels. For V_{cco} measurements, there was no significant difference between 2469 and 5971 γ radiation levels and between 7707 and 16966 γ radiation levels.

These results indicate that it is impossible to correlate a given voltage to a radiation level and create a relationship between them. Different sensor inputs would have different weights depending on the radiation rate level. For this reason, more elaborate approaches, such as machine learning algorithms, would have better results since there is a tuning of inputs through experience, in this case, the historical readings of temperature and voltage.

5. Anomaly detection with machine learning models

Typically, anomalous data in this study are connected to problems or rare events such as abnormal temperatures or voltages or malfunctioning components. This connection may imply which data points can be considered anomalies to identify these events that are typically useful for predicting the early failure of the system. This section explores three machine learning models: 1) Elliptical Envelope, 2) Local Outlier Factor, 3) One-Class Support Vector Machine as our anomaly detectors.

5.1. Machine learning models

5.1.1. Elliptical Envelope

Elliptical Envelope is a Gaussian distribution-based method that forms the key data parameters into an underlying multivariate Gaussian distribution expression. In short, it attempts to identify a boundary ellipse that covers most of the data. Therefore, any data not within the ellipse can be classified as an anomaly. The FAST-minimum covariance determinant is used to estimate the size of the ellipse, which selects non-overlapping samples of data and computes the mean u and covariance matrix C . Therefore, Mahalanobis distance d_{MH} for the input data vector x can be calculated using the following equation, and the data are then ordered ascendingly by d_{MH} [37].

$$d_{MH} = \sqrt{(x - \mu)^T C^{-1} (x - \mu)}$$

5.1.2. Local Outlier Factor

Local Outlier Factor is one of the Nearest Neighbour based methods for anomaly detection. In general, normal data are usually grouped in a neighbourhood that seems dense compared to the

abnormal data, which are far from their close neighbours. To quantify this neighbourhood, these types of approaches typically use distance-based or density-based methods, where both ways require a similarity or a distance calculation to determine whether the data are on the degree of abnormality or not. We use the Local Outlier Factor (LOF) abnormal detector in this study [38].

5.1.3. One-Class Support Vector Machine

One-Class Support Vector Machine (OCSVM) [39] is a classification-based anomaly detection method. Depending on the availability of labels, it can be divided into one-class and multi-class classification models. This approach is similar to all other supervised learning techniques, has two phases: 1) Training phase and 2) Testing phase. In the training phase, the classifier is trained using the labelled data, and then the data are classified as normal or abnormal using the trained model in the testing phase. In OCSVM, the classification rule for the linear decision boundary is given as follows:

$$f(x) = \langle \vec{w}, \vec{x} \rangle + b$$

where \vec{w} and b are the normal vector and bias, respectively. The algorithm is trying to find the rule f within the maximal geometric margin and then assign a label to a test example \vec{x} . For example, if $f(x) > 0$, the label of \vec{x} will be marked as normal; otherwise, it is labelled anomaly. This optimization problem can be solved by

$$\min_{\alpha} \frac{1}{2} \sum_{ij} \alpha_i \alpha_j K(x_i, x_j)$$

where α_i is the i^{th} weight, $0 \leq \alpha_i \leq \frac{1}{v}$ and $\sum_i \alpha_i = 1$. v is a variable to

control the maximizing the distance between the origin and the number of data points contained in the boundary. l is the number of points in the training dataset. $K(x_i, x_j)$ is the kernel function, and it is given as follows.

$$K(x, y) = \langle \phi(x), \phi(y) \rangle$$

where ϕ maps the training vectors from input space X to a high dimensional feature space.

A series of mathematical functions, known as the kernel, is used by SVM algorithms. The kernel function takes data as input and translates it into the appropriate form. Different SVM algorithms use various kernel function types. The adopted version, OCSVM, uses the Radial Basis Function (RBF) kernel [40].

5.2. Evaluation metrics: precision, recall, F1 score

This paper employs three well-known metrics for pattern recognition/classification to evaluate the machine learning models: Precision, Recall and F1 score [41]. Precision is the fraction of relevance among the retrieved instances, while Recall is the fraction of the total amount of relevant instances. All these metrics would signify better results as they approach the value of 1.

In the context of this paper, there are two kinds of data: annotated and not annotated – data with an anomaly and without it. The employed models should recognize if the data are an anomaly or not. Suppose we have a given dataset with ten anomalies and the remaining data are not an anomaly. A given ML model identifies eight data points as anomalies, of which five are anomalies (true positive) while the rest are not (false positives). The ML model's precision is 5/8, while the Recall is 5/10. In this example, precision means how valid the results are, while Recall shows how complete the results are. Therefore, in our study Recall is more important

since high values refer to early detection of anomalies rather than precision referring to the detection of anomaly at the same time point as anomaly is annotated.

F1 score is a metric for accuracy. It considers both the precision and the Recall to compute the score. The F1 score is then calculated using the harmonic mean of the Precision and Recall, where an F1 score has its best value at 1 (perfect precision and Recall).

6. Radiation experiments

The experiments took place on the Dalton Cumbrian Facility (DCF) laboratory, where a γ radiation source is available [42] [43]. The source is composed of a Cobalt (^{60}C) self-shielded irradiator that can provide absorbed dose rates of up to 20 kGy/h depending on the distance from the source to the DUT. The radiation cavity contains three rods, as shown in Fig. 6, with different dose rates for each one. Thus, different radiation rates are achieved using various configurations of the available rods, including lead obstacles to absorb the radiation and/or positioning the samples at different distances from the sources.

We employ six different experiments, each one using a separate board (DUT) under radiation – each board will be referred by its experiment number. Table 4 shows each feature of the experiments where different radiation rates were applied to see how the boards would behave. The DUT time is the time between the start of the experiment until the FPGA board stops sending data through the serial connection. This time should be considered the operational time of the DUT. On the other hand, the Monitoring board time is the time from the start of the experiment until the monitored voltages drop to zero, which means that the board itself, including the PMIC, is entirely inoperative. In each experiment, the board is under a constant radiation rate and stays the same distance from the rods. Before the actual radiation experiment begins, the radiation rate is measured using a probe removed afterwards. For this reason, the probe is inserted, the radiation is released and measured for 1 min, then the experiment is stopped, and the probe is removed. The radiation rate measured during that period is assumed to be constant for the whole experiment.

After the experiments, all the boards stopped working. During the experiments, one can see an increase in the voltage bounds on all experiments. The interesting point is that the deviation from the normal bounds does not indicate that the board will stop working right away. If we take, for example, the two extremes for the




Fig. 6. Minized board inside radiation chamber. The three pipes on the back contain three radiation rods that are lifted from the ground when the chamber is closed.

radiation rates, i.e. experiments 0 and 5, one took almost 2 h while the other took less than 10 min to stop working. Fig. 7 shows that although experiment 5 exhibits an early change in the bounds of the voltages, the board can take more than 1 h to stop working (experiment 0). On the other hand, the board might take a few minutes to stop working with a higher radiation rate after the first values outside the normal bounds are observed. This observation leads to a search for a more elaborate way of predicting when the board will stop working than just watching voltages outside the normal bounds of operation.

7. Evaluations

This Section compares the OCSVM method against the Local Outlier Factor and Elliptical Envelope methods. OCSVM copes well with non-linear functions and might be a more suitable approach for this problem. The section is divided into four subsections. The first one shows how the data was preprocessed and explains the methodology used in the second subsection, where the training set is built. Then, the third subsection compares the OCSVM against the two other methods using the training set. Finally, the last subsection further explores how early we can detect anomalies with OCSVM.

7.1. Training and testing methodology

The objective here is to compare how the training data would affect the results rather than comparing different ML algorithms. Hence, we show the comparison using the OCSVM algorithm.

The first step is to remove as least as possible from the collected data, all of which are time-stamped. Unfortunately, a few points show the temperature as zero or undefined (e.g. NaN). Therefore, all data points for that period were removed, even if another sensor showed consistent data. These removals do not interfere with the evaluation since all the data points are time-stamped.

After the data trimming step, we compare two approaches: one using seven features (five voltages and two temperatures) and one using the same features plus a constant value for the radiation rate – the rates shown in Table 4. Fig. 8 shows the precision of the two models with the six experiments. Using the constant radiation rate as input weakens the results significantly. Therefore, for further experiments, only voltage and temperature values are considered.

7.2. Exploring the best training and testing strategy in a radiation site

At this point, the number of features is defined, we came up with three training and testing strategies to evaluate, compute or organize the measured data. The target is to find the right balance of data points in the training set. The following strategies evaluate trade-offs between a set of experiments employed and the number of data points. Each strategy represents a data processing approach, and it is detailed next as follows:

- Strategy 1: Train the model with the first minutes of a given board and then test on the remaining measurements of the same board.
- Strategy 2: Train the model with the first minutes of a set of different experiments and test on the remaining measurements of all boards.
- Strategy 3: Find the optimum number of training samples for the ML system.

The first strategy was employed using the first few minutes of each board and then test on the remaining measurements on the same board. Fig. 9 shows the results for this strategy. Even using the




Fig. 7. Voltages behaviour of experiments 0 and 5, respectively.

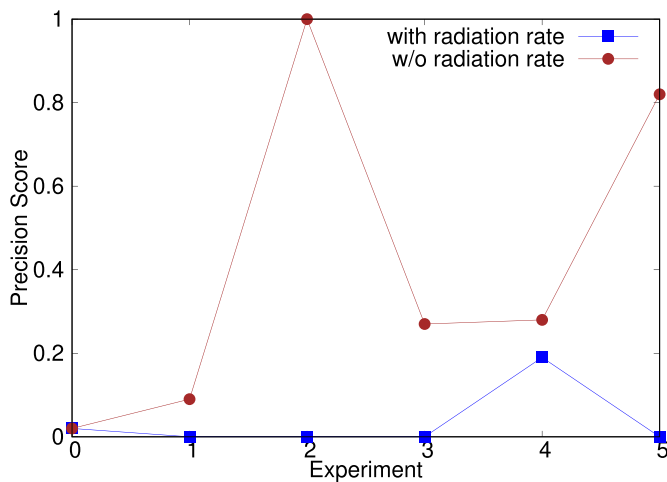


Fig. 8. Comparison using different features as input for the training data. Respectively lines show the results using i) seven features (five voltages and two temperatures) and ii) the same seven features plus the constant radiation rate.




Fig. 9. Results using each experiment separately as training data.

same board data and comparing it with the remaining data does not provide good results. Only board from experiment 5 shows good scores for F1 and precision. Nevertheless, the results for Recall are excellent showing 1 for all the scenarios, which means the model is capable of identifying all annotated anomalies.

Because of the weaker results of strategy 1, we decided to include more boards on the training data. Strategy 2 evaluates different sets of boards as training data to check how many we need to include to

get the best results. For this strategy, we trained all combinations sets of two up to all six boards, i.e. sets of two boards $\{\{0,1\}, \{0,2\}, \dots, \{4,5\}\}$, sets of three boards $\{\{0,1,2\}, \{0,1,3\}, \dots, \{3,4,5\}\}$ and so on. Later we compared the F1 score for each board separately. Results shown in Fig. 10 include the worst and best set of boards only for the sake of space. One can assume that we get better results as we feed more information to the model (different boards). For example, the worst results for sets with two, three and four boards are much lower than results using five or all six boards.

On the third strategy, we evaluate the effects of the amount of training data fed into the model. Fig. 11 shows the comparison of using the first 300, 360, 420, 480, 520 and all data points as input for the model. Interestingly, as we feed more data to the training set, it does not mean we would get better results. Adding more than 420 points does not increase the F1 score, and as it shows better average results, we keep this number of data points as a training set for all experiments. This means that only 7 min of observation can be enough to save all the DUT boards under radiation.

7.3. Comparison of OCSVM with other anomaly detection techniques

This subsection compares three state-of-the-art machine learning techniques: i) Elliptical Envelope, ii) Local Outlier Factor and iii) One Class Support Vector Machine (OCSVM). These techniques are trained using the proposed dataset, discussed previously, as inputs.

To compare the anomaly detection results, we first need to annotate the data showing where the anomalies happened and then compare with the ML algorithms outputs. We have calculated the size of the time window where each board exhibits values outside usual bounds during radiation experiments, visually observing the measured voltage (we could not observe such values for temperature). Table 5 shows the size of this window, and associated calculations are displayed for each board. As can be seen, the minimum window size observed is approximately 3 min. That is to say that the last 3 min of each experiment is the time where each board exhibits values outside normal bounds. Therefore, after adding a safety margin of 2 min, we annotate the last 5 min of each experiment as anomalies for the use in the training dataset fed into the machine learning algorithms. Five minutes is roughly equivalent to 300 data points; therefore, we annotate this number of points as an anomaly. Besides the observed window size, 300 data points also give the developer a reasonable amount of time to take action before the board stops working, e.g., move the computed data to safe storage, to another computing node in the system or even move it away from the radiation environment. We can also consider a different amount of time as annotation, but it would result in the retraining of the model. Then finally, we can summarize the training set with the annotation as follows:




Fig. 10. Using sets of 2, 3, 4, 5 and all boards as training data. Figure showing only best and worst results for each set.




Fig. 11. Comparison of training data for input data. Using the first 420 data points from all experiments show better results. 420 data points represent roughly 7 min of data.

Table 5
Window times where voltages exhibit values outside usual bounds for each experiment.

Exp. #	Window time
0	74 min
1	18 min
2	07 min
3	07 min
4	06 min
5	03 min

- Remove all the inconsistent sensor readings;
- Collect the first 420 data points from all boards (as justified in Fig. 11);
- We annotate the last 300 data points as an anomaly.

Using that training set, we then employ a multivariate analysis with seven variables (including five voltage and two temperature values) to feed the ML algorithms. Fig. 12 summarizes the results for Elliptical Envelope, LOF and OCSVM algorithms where the F1, Precision, and Recall scores are shown for each model. OCSVM demonstrates better outcomes for all scores. It is essential to highlight that the Recall score for OCSVM is the one for all experiments except number two, which is a remarkable result, showing an average Recall score of 0.95 and strong evidence that the model can

detect the relevant results. The recall score in experiment 2 is not the maximum one but shows a high value (0.842) and will be discussed next.

7.4. Exploring how early OCSVM can detect anomalies

As the OCSVM showed the best F1, Precision and Recall results, this subsection details each experiment separately, as illustrated in Fig. 13. Each experiment shows in the rows: (i) the five observed voltages; (ii) two temperatures; (iii) anomaly annotation; and (iv) the OCSVM model output (i.e. 0 represents normal data and 1 an anomaly both for annotation and model output).

Table 6 shows that in all experiments except one, the model marks the abnormal behaviour of the board before the annotation. This is an exceptional result, indicating that the model can be used as a suitable indicator/warning before the board stops working. Note that the ML output for experiment 2 is delayed by only 12 s (negative number in Table 6); nevertheless, there still would be enough time for the board to be saved before it is permanently damaged.

Experiment 0 is the longest in time under radiation mainly because it has the lowest radiation rate among the experiments. The board took almost 2 h to stop working, and it is the worst result for the model output. Although the model could point to the anomaly, it has done it earlier, more than 1 h before the board stops working. Experiment 1 has a shorter execution time, and the model was also capable of predicting before the anomaly, i.e. 25 min before the annotation. Experiment 2 has one particular behaviour different from the others. The model output shows one unique point marking as an anomaly in the first minutes. As this was the only point and not a continuous trend, we can disregard this result - one should always observe the trend rather than particular points. The model then marks only points 12 s after the anomaly annotation. We still consider this a good result since a few seconds later would still allow time to take action.

In experiments 3, 4 and 5, the model behaves similarly, marking the anomaly 7, 6 and 0.5 min before the annotation, which is a remarkable result. To summarize, all boards would have been saved since the model can point out before a board dies completely, thus allowing a few minutes for the designer to save or transfer the processed data to a safe environment.

It is essential to point out that the model has a trade-off. As we gathered more experiments with high radiation rates, five of the experiments have more than 2000 Gy/h; the model has been trained with more data for these environments. Therefore it also performs better, i.e. predicts near the annotation, on experiments with high radiation rates. As we feed more information, that is, more experiments in different radiation rates, the model should perform better.




Fig. 12. Comparison of machine learning models for the dataset.

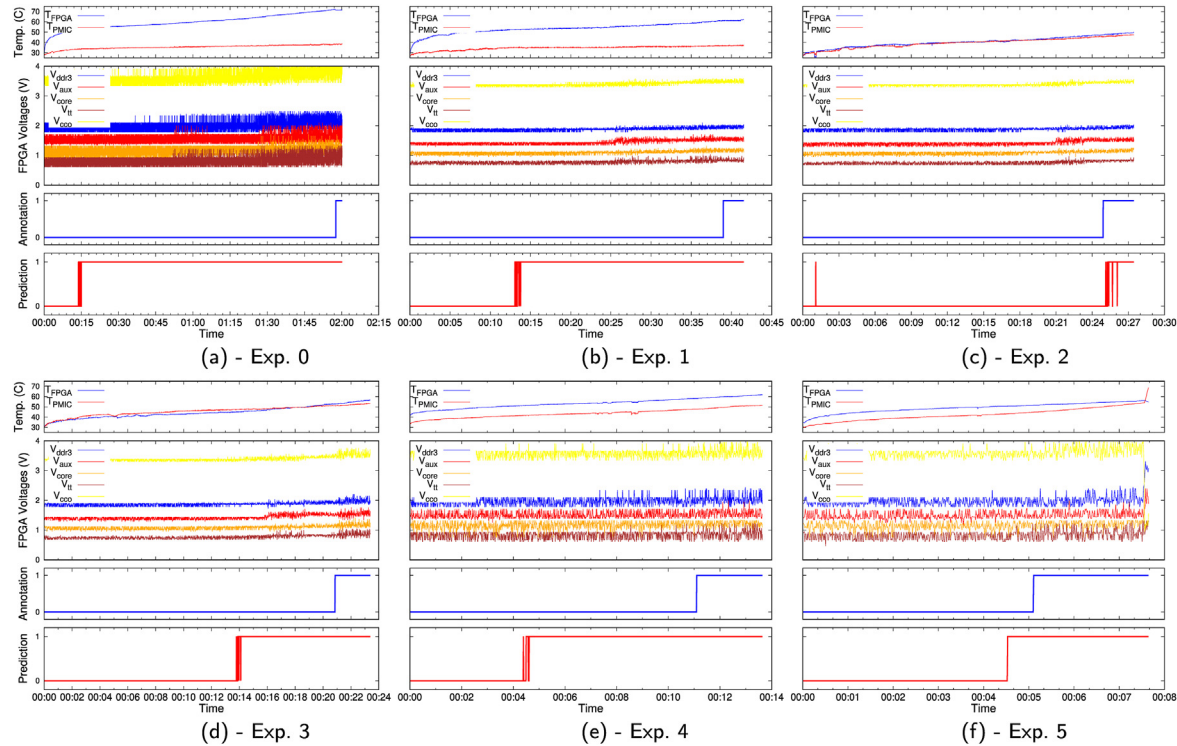


Fig. 13. Results for each experiment using One-Class Support Vector Machine to predict the anomaly before the board stops working.

Table 6

Time showing how earlier the OCSVM marks the annotation (around 5 min before the board stops working). The model prediction is calculated with (Annotation - Prediction) from Fig. 13.

Exp. #	Experiment duration	Model prediction before manually annotated anomaly
0	2:00:11	1:43:53
1	0:41:31	0:25:56
2	0:27:28	-0:00:12
3	0:23:23	0:07:04
4	0:13:39	0:06:42
5	0:07:38	0:00:34

8. Conclusion and future works

This paper proposed an anomaly detection machine-learning algorithm to predict when a COTS FPGA would stop working due to gamma radiation. The OCSVM algorithm showed the best results in terms of Recall score and was capable of pointing out 100% of the anomalies before the board stopped working. Annotating the anomaly before it stops working and giving time to the designer to take actions allowed to detect the anomaly before the board dies,

considered an extraordinary achievement since it will enable the designer to use this as an assumption for future works. This work employed six boards that were inoperable after the experiments. Using more boards with different radiation rates can improve the model results.

With such indicator before the board being inoperative, different design decisions might be taken to provide a further level of reliability to the overall system. For example, if this is a data collection application, move the computed data to a safe storage. If

the equipment used for prolonged radiation exposure, can be used as indicator to replace the equipment before it is inoperative. Finally, in the case of a moving robot, this indicator can be used to inform the robot to move away from the radiation environment or going back to the base.

Future works include using the OCSVM algorithm at run-time experiments. The proposed approach with DUT and a monitoring board can also be modified to execute on a self-contained solution. The Arduino board of the DUT was used to remove the monitoring board and sensors from the radiation environment. In a real-case scenario, this would not be possible. To tackle that approach, the Minized board (DUT) could use the sensor boards directly and execute the OCSVM algorithm itself.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

We acknowledge the support of The University of Manchester's Dalton Cumbrian Facility (DCF), a partner in the National Nuclear User Facility, the EPSRC UK National Ion Beam Centre and the Henry Royce Institute. We recognise Kevin Warren for their assistance during the experiments.

References

- [1] Nuclear Provision: explaining the cost of cleaning up Britain's nuclear legacy. <https://www.gov.uk/government/>.
- [2] O. Gonçalves, G. Prenat, B. Dieny, Radiation hardened mram-based fpga, *IEEE Transactions on Magnetics* 49 (7) (2013) 4355–4358.
- [3] S. Ramaswamy, L. Rockett, D. Patel, S. Danziger, R. Manohar, C.W. Kelly, J.L. Holt, V. Ekanayake, D. Eiftmann, A radiation hardened reconfigurable fpga, in: 2009 IEEE Aerospace Conference, 2009, pp. 1–10.
- [4] A. Lindoso, L. Entrena, M. García-Valderas, L. Parra, A hybrid fault-tolerant LEON3 soft core processor implemented in low-end SRAM FPGA, *IEEE Transactions on Nuclear Science* 64 (1) (Jan 2017) 374–381.
- [5] 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, M. Wirthlin, Using benchmarks for radiation testing of microprocessors and FPGAs, *IEEE Transactions on Nuclear Science* 62 (6) (2015) 2547–2554.
- [6] S.M. Trimmerger, Three ages of fpgas: a retrospective on the first thirty years of fpga technology, *Proceedings of the IEEE* 103 (3) (2015) 318–331.
- [7] Tomas Fried, Antonio Di Buono, David Cheneler, Neil Cockbain, Jonathan M. Dadds, Peter R. Green, Lennox Barry, C. James Taylor, Stephen D. Monk, Radiation testing of low cost, commercial off the shelf microcontroller board, *Nuclear Engineering and Technology* 53 (10) (2021) 3335–3343.
- [8] Andrew West, Jordan Knapp, Lennox Barry, Steve Walters, Stephen Watts, Radiation Tolerance of a Small Cots Single Board Computer for Mobile Robots, *Nuclear Engineering and Technology*, 2021.
- [9] Joerg Leukel, Julian González, Martin Rieker, Adoption of machine learning technology for failure prediction in industrial maintenance: a systematic review, *Journal of Manufacturing Systems* 61 (2021) 87–96.
- [10] S. Kasap, E. Weber Wächter, X. Zhai, S. Ehsan, K. McDonald-Maier, Survey of soft error mitigation techniques applied to LEON3 soft processors on SRAM-based FPGAs, *IEEE Access* 8 (2020) 28646–28658.
- [11] Yufan Lu, Xiaojun Zhai, Sangeet Saha, Shoaib Ehsan, D. Klaus, McDonald-Maier, A self-adaptive seu mitigation scheme for embedded systems in extreme radiation environments, *IEEE Systems Journal* (2022) 1–12.
- [12] Carl Carmichael, Triple Module Redundancy Design Techniques for Virtex FPGAs, Xilinx Inc., San Jose, CA, USA, XAPP197 Application Note, 2006.
- [13] D.K. Pradhan, N.H. Vaidya, Roll-forward and rollback recovery: performance-reliability trade-off, *IEEE Transactions on Computers* 46 (3) (March 1997) 372–378.
- [14] L. Parra, A. Lindoso, M. Portela-García, L. Entrena, B. Du, M.S. Reorda, L. Sterpone, A new hybrid nonintrusive error-detection technique using dual control-flow monitoring, *IEEE Transactions on Nuclear Science* 61 (6) (Dec 2014) 3236–3243.
- [15] H. Quinn, Z. Baker, T. Fairbanks, J.L. Tripp, G. Duran, Robust duplication with comparison methods in microcontrollers, *IEEE Transactions on Nuclear Science* 64 (1) (Jan 2017) 338–345.
- [16] Kuang-Hua Huang, J.A. Abraham, Algorithm-based fault tolerance for matrix operations, *IEEE Transactions on Computers*, C- 33 (6) (June 1984) 518–528.
- [17] J.R. Azambuja, M. Altieri, J. Becker, F.L. Kastensmidt, HETA: hybrid error-detection technique using assertions, *IEEE Transactions on Nuclear Science* 60 (4) (Aug 2013) 2805–2812.
- [18] E. Chielle, G.S. Rodrigues, F.L. Kastensmidt, S. Cuenca-Asensi, L.A. Tambara, P. Rech, H. Quinn, S-SETA: selective software-only error-detection technique using assertions, *IEEE Transactions on Nuclear Science* 62 (6) (Dec 2015) 3088–3095.
- [19] Harn Hua Ng, PPC405 Lockstep System on ML310, Xilinx Inc., San Jose, CA, USA, XAPP564 Application Note, 2007.
- [20] F. Abate, L. Sterpone, C.A. Lisboa, L. Carro, M. Violante, New techniques for improving the performance of the lockstep architecture for SEEs mitigation in FPGA embedded processors, *IEEE Transactions on Nuclear Science* 56 (4) (Aug 2009) 1992–2000.
- [21] M. Violante, C. Meinhardt, R. Reis, M. Sonza Reorda, A low-cost solution for deploying processor cores in harsh environments, *IEEE Transactions on Industrial Electronics* 58 (7) (July 2011) 2617–2626.
- [22] H. Pham, S. Pillement, S.J. Piestrak, Low-overhead fault-tolerance technique for a dynamically reconfigurable software processor, *IEEE Transactions on Computers* 62 (6) (June 2013) 1179–1192.
- [23] A.B. de Oliveira, G.S. Rodrigues, F.L. Kastensmidt, N. Added, E.L.A. Macchione, V.A.P. Aguiar, N.H. Medina, M.A.G. Silveira, Lockstep dual-core ARM A9: implementation and resilience analysis under heavy ion-induced soft errors, *IEEE Transactions on Nuclear Science* 65 (8) (Aug 2018) 1783–1790.
- [24] E.W. Wächter, S. Kasap, X. Zhai, S. Ehsan, K. McDonald-Maier, Survey of lockstep based mitigation techniques for soft errors in embedded systems, in: *Computer Science and Electronic Engineering Conference (CEEC 2019)*, 2019, pp. 124–127.
- [25] N. Rezzak, J. Wang, M. Traas, A. Zerrouki, G. Bakker, F. Xue, A. Cai, F. Hawley, J. McCollum, E. Hamdy, Investigation of tid and dynamic burn-in-induced v_T shift on rtg4 flash-based fpga, *IEEE Transactions on Nuclear Science* 65 (1) (2018) 64–70.
- [26] P. Villa, E. Bezerra, R. Goerl, L. Poehls, F. Vargas, N. Medina, N. Added, V. de Aguiar, E. Macchione, F. Aguirre, M. da Silveira, Analysis of cots fpga sensitivity to combined effects of conducted-emi and tid, in: 2017 11th International Workshop on the Electromagnetic Compatibility of Integrated Circuits (EMCCompo), 2017, pp. 27–32.
- [27] J. Tarrillo, J.R. Azambuja, F.L. Kastensmidt, E.C.P. Fonseca, R. Galhardo, O. Gonzalez, Analyzing the effects of tid in an embedded system running in a flash-based fpga, *IEEE Transactions on Nuclear Science* 58 (6) (2011) 2855–2862.
- [28] G. Lentaris, K. Maragos, D. Soudris, F. Di Capua, L. Campajola, M. Campajola, A. Costantino, G. Furano, A. Tavoularis, L. Santos, Tid evaluation system with on-chip electron source and programmable sensing mechanisms on fpga, *IEEE Transactions on Nuclear Science* 66 (1) (2019) 312–319.
- [29] Minized board product brief. <http://zedboard.org/>.
- [30] Zynq-7000 SoC data sheet. <https://www.xilinx.com/>.
- [31] DA9062 data sheet. <https://www.dialog-semiconductor.com/>.
- [32] MiniZed power architecture. zedboard.org.
- [33] Elegoo R3 board. <https://www.elegoo.com/>.
- [34] ATmega data sheet. <http://ww1.microchip.com/>.
- [35] PmodTC1 board reference manual. <https://reference.digilentinc.com/>.
- [36] J. Cohen, *Statistical Power Analysis for the Behavioral Sciences*, second ed., Lawrence Erlbaum Associates., Hillsdale, NJ, 1977.
- [37] Ben Hoyle, Markus Michael Rau, Kerstin Paech, Christopher Bonnett, Stella Seitz, Jochen Weller, Anomaly detection for machine learning redshifts applied to SDSS galaxies, *Monthly Notices of the Royal Astronomical Society* 452 (4) (2015) 4183–4194.
- [38] Malak Alshawabkeh, Byunghyun Jang, David Kaeli, Accelerating the local outlier factor Algorithm on a GPU for intrusion detection systems, in: *Proceedings of the 3rd Workshop on General-Purpose Computation on Graphics Processing Units*, Association for Computing Machinery, New York, NY, USA, 2010. GPGPU-3, page 104–110.
- [39] Katherine A. Heller, Krysta Marie Svore, Angelos D. Keromytis, salvatore J. Stolfo, One class support vector machines for detecting anomalous windows registry accesses, in: *Proceedings of the ICDM Workshop on Data Mining for Computer Security*, 2003.
- [40] David S. Broomhead, David Lowe, Radial Basis Functions, Multi-Variable Functional Interpolation and Adaptive Networks, Royal Signals and Radar Establishment Malvern, United Kingdom), 1988. Technical report.
- [41] D.M.W. Powers, Evaluation: From precision, recall and f-measure toroc, informedness, markedness and correlation, *Journal of Machine Learning Technologies* 2 (1) (2011) 37–63.
- [42] Dalton Cumbrian facility. <https://www.dalton.manchester.ac.uk/research/facilities/cumbria-facilities/>.
- [43] Zeba Khanam, Bilal Aslam, Sangeet Saha, Xiaojun Zhai, Shoaib Ehsan, Rustam Stolkin, Klaus McDonald-Maier, Gamma-induced image degradation analysis of robot vision sensor for autonomous inspection of nuclear sites, *IEEE Sensors Journal* (2021), <https://doi.org/10.1109/JSEN.2021.3050168>. In press.