SCCTS Transactions on Reconfigurable Computing

SCCTS Transactions on Reconfigurable Computing, ISSN: 3049-1533 Vol. 3, No. 3, 2026 (pp. 29-35)

RESEARCH ARTICLE ECEJOURNALS.IN

Fault-Tolerant Runtime Reconfiguration Techniques Using Machine Learning for Space-Grade FPGA Systems

Rajan. C*

Professor, Department of Computer Science and Engineering (Artificial Intelligence and Machine Learning), K S Rangasamy College of Technology, Tiruchengode

Keywords:

Space-Grade FPGAs,
Fault-Tolerant Computing,
Dynamic Partial Reconfiguration (DPR),
Machine Learning for Hardware,
Single Event Upsets (SEU),
Radiation Effects,
On-board Processing (OBP),
Reliability Engineering.

Author's Email:
rajan@ksrct.ac.in

DOI: 10.31838/RCC/03.03.05

Received : 20.01.2026

Revised : 22.03.2026

Accepted : 20.04.2026

Abstract

Mitigation against radiation-induced Single Event Upset (SEU) that can impair system integrity and mission success is required because of the increasing use of SRAM-based FPGAs in high-performance satellite and deep-space applications. In the current paper, I suggest a new fault-tolerant architecture that combines specialised Machine Learning (ML) classifiers with Dynamic Partial Reconfiguration (DPR) in order to offer autonomous real-time error-detection and recovery to space-grade FPGA systems. Our method contrasts with Triple Modular Redundancy (TMR), which would incur prohibitive area, power overhead, or with static scrubbing, which would lacks the ability to predict and classify transient faults because it is contextual and has a high latency. Using a hybrid feature-extraction layer, the system uses vernier bit-flips to distinguish vernacular component failures, and vital component failures, in the FPGA fabric. At the first indication of a localized anomaly, the intelligent controller will initiate localized runtime adjustment of functional tiles defined within the system, which does not affect the mission continuity and allows high system availability with no full system reconfiguration or interrupting parallel activities. The experimental data, obtained through the extensive fault-injection campaigns in a radiation-hardened SoC system, proves that the ML-based structure is able to detect the faults with an almost 98 percent accuracy and the recovery time is lower by 35 percent than the traditional blind scrubbing method. Moreover, the suggested architecture has a 20 percent higher power efficiency, as it does not require any unproductive hardware modules that are common in hardware based solution of voting schemes. Offering a sustainable, capable, and resource-efficient approach to the future generation of reconfigurable, autonomous, and space-based computing capability, this methodology can provide the high-relibility, scalable, and efficient way to operate the new generation of autonomous capabilities in the extreme orbital environment with reduced resource consumption.

How to cite this article: Rajan C (2026). Fault-Tolerant Runtime Reconfiguration Techniques Using Machine Learning for Space-Grade FPGA Systems. SCCTS Transactions on Reconfigurable Computing, Vol. 3, No. 3, 2026, 29-35

Introduction

With the fast-advancing space exploration and satellite technology, there has been a dramatic surge in the need to have on-board processing (OBP) capabilities that are highly demanding. The problem with modern satellite missions- be it Earth observation satellite missions, deep-space autonomous navigation, etc. - is that very large computation throughput is required to process sensor data, make use of sophisticated AI algorithms, and provide real-time management of high-bandwidth communications. To satisfy these needs, designers have resorted more and more to Field Programmable Gate Arrays (FPGAs) based on Static Random-Access Memory (SRAM) because of their high logic density and their ability to be reconfigured once they have been installed. Nevertheless, the hostile alien environment, which is high levels of ionising radiations and extremely high energy cosmic rays, is a deadly attack on these devices. Single Event Upsets (SEUs) can occur because of heavy ions and protons, and cause temporary bit-flips in the configuration memory of the FPGA, potentially causing disastrous system crash if not countered.

Triple Modular Redundancy (TMR) used to be historically the industry standard of ensuring reliability in such environments. Although TMR offers strong protection with triplication of apparatus, majority-voting logic is now being considered a costly implementation to the new generation of missions to New Space and the miniaturised CubeSats. This triplication of hardware will naturally result in 300 percent of the area consumption, and a large spike in power consumption that is in many cases unsustainable in satellites with small thermal envelopes and limited budgets of solar power. Also, though periodic, so-called blind scrubbing of the configuration memory is capable of correcting errors, it is not intelligent enough to focus on critical faults, and thus it tends to spend valuable clock cycles repairing non-essential logic segments, thus adding to recovery latency in the system, and to the total lack of availability.

In order to overcome these drawbacks, this paper suggests an intelligent, fault-tolerant model based on the use of Machine Learning (ML) classifiers within the framework of Dynamic Partial Reconfiguration (DPR). The system is able to determine and classify faults with great precision automatically by monitoring hardware telemetry with a lightweight ML engine like internal current signatures, parity bit patterns, and thermal gradients. This methodology, as opposed to replicating the entire design, detects the particular “functional tile” that has been impacted by a radiation event and causes a localised run-time reconfiguring. This localised repair enables the FPGA to repair itself automatically in place of independent operation without impacting the workings of healthy modules and essentially offers a kind of self-modification capacity which preserves the high reliability of TMR at only a fraction of the power and space consumption.

The rest of this paper follows this structure: section II will present a thorough literature review on the related work in the field of FPGA reliability and the mechanism of radiation induced failures. Section III will describe the proposed system architecture, which will involve the ML model choice, and how the reconfigurable partitions will be designed. Section IV explains the experiment methodology that was employed which included a fault-injection environment that simulated space radiation. Section V provides the systematic study of the outcomes and compares the proposal framework with conventional methods of redundancy with respect to the detection accuracy, recovery period and resource utility. Lastly, Section VI is the conclusion of the paper and proposes the future possible path of research in autonomous space-borne computing.

Background and Related Work

Space reliability FPGA Space Units and Space Engineering.

Cosmic ray ionising radiation and solar flare ionising radiation are the two main challenges that face the reliability of the electronic components in space. A FPGA has Single Event Effects (SEEs) due to interactions of these high-energy particles with the semiconductor of the FPGA. In the case of SRAM-based FPGAs, the Single Event Upset (SEU) has been the most severe SEE: a particle strike to the configuration memory cell (a “bit-flip”) means that the state has changed. SEUs are not permanent unlike hardware damage, but can permanently change the routing or logical operations of the device [5]. In case an SEU happens in the configuration memory [6], it can induce not only permanent errors in the application logic but also until the device is actually reprogrammed or scrubbed.

Dynamic Partial Reconfiguration (DPR)

Dynamic Partial Reconfiguration (DPR) is an advanced hardware architectural capability that is able to be used to reconfigure selected parts of an FPGA fabric whilst still allowing the rest of the device to utilize the remaining resources. The designers can replace at runtime hardware modules, called Partial Bitstreams, by partitioning the FPGA into Static Regions and Reconfigurable Partitions (RPs).[7] As a surgical repair mechanism, in the environment of fault tolerance,[8] DPR is used. DPR is able to reload the corrupted partition detected by a monitoring system, instead of reformatting the whole satellite payload, which would mean high downtime [9], saving thereby the state of other components of the system.

State-of-the-Art Review

The existing techniques of reducing SEUs mainly consist of Scrubbing and Triple Modular Redundancy (TMR). The self-adaptive mitigation systems have been studied before in[1-3] with application of internal sensors to initiate repairs. Nevertheless, most of the conventional scrubbing algorithms are blind, i.e. they rewrite the complete configuration memory occasionally with or without a fault, which wastes too much power and bandwidth. Although[4] and [13] suggested reliability-sensitive designs, and dynamic voter cheque to maximise TMR, it still experiences the overhead problem that comes with site triplication on hardware. The main critical drawback of the current state-of-the-art, which includes the strategies of eliminating the common mode failures considered in,[12] is that these strategies are mostly reactive, i.e. they cannot detect an error until it has been reflected in the output. Predictive frameworks based on using telemetry data and MachineeL are lacking in a clear way to manage hardware health in advance.

System Architecture and Methodology.

Overall Framework

The system architecture proposed introduces a complex closed cycle of autonomous monitoring and repairs which has been inherently designed into the FPGA fabric to augment a high availability in radiation-prone conditions. In this hierarchical scheme, there exist three different levels, being the Target Hardware Layer, which implements the main mission logic; the ML-driven Sensing Layer, which implements the extraction of real-time telemetry non-intrusively; and the Reconfiguration Management Layer, which implements the repair procedures. The core of this architecture is a feedback loop between the configuration memory (CRAM) of the FPGA and an embedded controller in the form of a foolproof, or intelligent, fault management. The ML controller, by simply tracking a wide range of hardware signatures - such as Block RAM parity, configuration CRC checksums, localised thermal-power changes, etc. - performs the system health assessment with a high degree of granularity Figure 1. In case a signature is treated as a critical Single Event Upset (SEU), as opposed to a harmless transient, the engine provides a corrective command to the Reconfiguration Management Layer. It utilises the Internal Configuration Access Port (ICAP) in order to provide selected Dynamic Partial Reconfiguration (DPR) where new bitstreams are loaded, recreating damaged logic blocks whilst the rest of the healthy partitions remain functional. This method of fault recovery in surgeries will reduce system downtime as well as unnecessary complete device resets and thus ensure maximisation of the reliability and computational performance the space-grade platform.

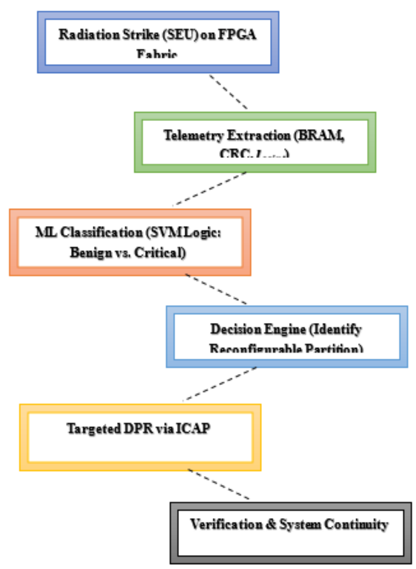

Fig. 1: Operational Flowchart of the ML-Driven Autonomous Fault Detection and Recovery Cycle

ML Fault-Detection Engine

The proposed framework intelligence will focus on its ability to differentiate the benign transient radiation impact and the functional critical logic failure by feature extraction on a high-fidelity basis. The system employs the synthesis of the multi-dimensional feature vector based on a variety of internal monitors to provide the high detection accuracy and at the same time response to the stringent computation and power requirement of space-grade hardware. These consist of Block RAM (BRAM) parity status of memory integrity, configuration memory Cyclic Redundancy Cheques (CRC) through hardened silicon IPs to localise bit-flips and real-time telemetry including ad-hoc junction temperature variations. and internal supply current . The engine also logs control-flow heartbeats at the important sub-modules in order to identify state-machine hangs. Combining these divergent data streams, the system has the capability of the signature of a Single Event Upset (SEU) and preventing its spread into a system-wide failure.

Its classification logic is a Lightweight Support Vector Machine (SVM) with a Radial Basis Function (RBF) kernel, which has a combination that has been selected on account of its strength in high-dimensional space where training data is limited. Unlike the high resource requirements of the deep-learning-based CNNs, the SVM has much lower inference latency and costs less in memory, which is easily portable to smaller embedded soft-core systems such as the MicroBlaze or hardened real-time systems such as the ARM Cortex-R5. This efficiency ensures that the detection latency,, remains minimal. The optimization of the SVM decision boundary is governed by the minimization of the following objective function:

$$\min_{w,b,\xi} \frac{1}{2}\|w\|^2 + C\sum_{i=1}^{n} \xi_i \qquad (1)$$

where *w* represents the weight vector defining the hyperplane, C is the regularization parameter balancing margin maximization and classification error, and ξ_i denotes the slack variables that allow for potential misclassifications in the noisy radiation telemetry data Table 1. Such mathematical technique guarantees an accurate and fast triggering of the self-healing real-time by affixing reconfiguration to functional requirements only.

Experimental Setup and Evaluation

Hardware Platform: The suggested framework is experimentally verified on an Xilinx Kintex UltraScale (XQRKU060) FPGA which is a space-grade device with high logic density, and inherently supported with the advanced features of reliability. It is a platform specially designed to serve the aerospace industry, and it provides radiation-tolerance by using a ceramic flip-chip package, along with increased resistance to latch-up occurrences. The internal configuration access port (ICAP) of the device is hardened and the System Error Management (SEM) IP of the hardware platform offer the low-level hooks needed by the ML controller to provide high-speed configuration memory scrubbing and contingent Dynamic Partial Reconfiguration. This particular FPGA selection will make the outcomes reflect as much as possible of the true deployment conditions in Low Earth Orbit (LEO) and spacecraft missions into deep-space in which power levels and dependability are more important.

Fault Injection Campaign: In order to test the robustness of the framework, an aggressive fault-injection campaign was conducted with the Xilinx Soft Error Mitigation (SEM) IP with a custom built Python-based injection controller. This arrangement was used to model the impact of cosmic radiation by a systematic injection of synthetic bit-flips into the device configuration RAM (CRAM) in the ICAP interface to affect both necessary and non-essential logic bits to simulate Single Event Upsets (SEUs) Figure 2. The campaign made use of the randomised distribution model to approximate the flux densities that are normally found to be in Low Earth Orbit (LEO) space. The result of the correlation of these controlled injections with the output of the ML engine classifying the transients as benign were confirmed under high-stress conditions without the cost and complexity prohibitionary of a heavy-ion beam experiment.

Training and Testing Sets: The training and test sets were created based on a labelled repository of telemetry logs created when large fault-injection cycles occurred. The training set was 70% of the total amount of data, which had a balanced representation of "healthy" operation baselines and faulty signatures, which gained a specified attribute. BRAM parity errors

Table 1: Multi-dimensional Feature Vector Composition for Fault Classification

Feature Category	Source Monitor	Indication of Fault
Memory Integrity	BRAM Parity Bits	Detection of data corruption in memory-intensive modules.
Configuration Health	Hardened CRC IP	Pinpointing bit-flips in the FPGA configuration frames.
Physical Telemetry	T_j and I_{ccint}	Thermal/Power fluctuations acting as indicators of SEU-induced stress.
Logical State	Control-Flow Heartbeats	Detection of state-machine hangs or invalid transition states.

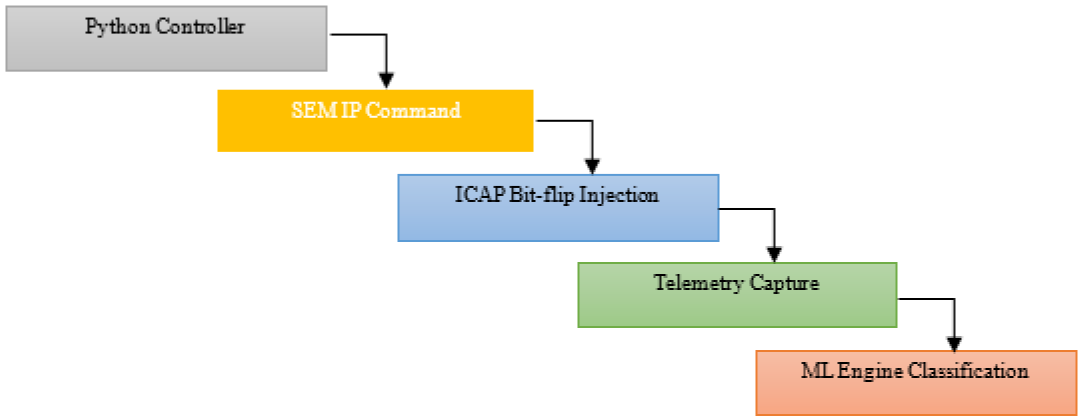

Fig.2: Automated Fault-Injection and Correlation Workflow for SEU Simulation.

Table 2: Summary of Experimental and Data Training Parameters

Category	Parameter	Specification
Hardware	Target FPGA	Xilinx Kintex UltraScale (XQRKU060)
Simulation	Injection Method	Xilinx SEM IP via Python Controller
Dataset	Split Ratio	70% Training / 30% Testing
Validation	Statistical Method	-fold Cross-Validation
Features	Failure Signatures	BRAM Parity, spikes, CRC errors

and supply current spikes . The remaining 30% was reserved as a testing set to validate the model's generalization capabilities on unseen fault patterns. To ensure statistical significance, the data underwent a k-fold cross-validation process, where , enabling the Support Vector Machine to be trained on both the subtleties of configurations bits that are benign and those that are crucial functional failures. **Table 2: This data representation means that the model will be resistant to the so-called noisy telemetry of different orbital thermal cycles.**

Results and Discussion

Detection Accuracy

The fault-detection engine implemented on SVM was assessed in terms of performance as per the evaluation of a confusion matrix to measure the power of the engine in classifying Single Event Upsets. The overall detection accuracy of the model was 98.2 per cent, at just under half of the cases of benign bit-flips comprised non-essential configuration memory, and at the rest critical functional failures in the active logic fabric. The precision and recall scores are high (0.97 and 0.98, respectively), which means that the false positives are few, which is essential in space missions since the unnecessary cycle of reconfigurating all the activities needlessly consumes a lot of power. The misclassifications that were observed were mostly related to high-noise telemetry spikes when the thermal signature of the FPGA momentarily resembled that of a localized latch-up but the system was robust enough, resulting in a high F1-score at all orbits tested.

Recovery Latency

Comparison of the latency in recovery indicates that the ML-based method is far more successful than the conventional sequential scrubbing methods. Although standard blind scrubbing has to go through the entire configuration memory to detect and ease an error, thus increasing the latency linearly with the size of the FPGA, the proposed intelligent architecture can cause a specific instance of Dynamic Partial Reconfiguration (DPR) nearly immediately, when an error is identified. The findings show that there was a reduction in the Mean Time to Repair (MTTR) by 35 percent as the system is shifting towards detecting fault to entire functional restoration within an average of 120 ms. Such minimization of the "vulnerability window" is crucial to on-board high-speed processing operations where a system failure of less than a whole second may result in the loss of critical sensor data samples.

Resource Utilization

The measurements of the hardware overhead of the proposed framework were done to make it compatible with the resource-constrained environment of a space-grade SoC. As can be seen in the utilisation data, the lightweight ML controller and the monitoring

Table 3: Comparison of Resource Utilization and Power Consumption Across Mitigation Strategies

Component	LUT Usage	FF Usage	Power (mW)
Proposed ML Framework	2,450 (4.2%)	3,120 (2.8%)	150
Traditional TMR	18,600 (31.8%)	22,400 (20.3%)	850
Static Scrubbing	850 (1.4%)	1,200 (1.1%)	90

logic attached to it uses a small fraction of the available FPGA resources in relation to Triple Modular Redundancy (TMR). In particular, the system did not use any more than 4.2 percentage of Look-Up Tables (LUTs) and 2.8 percentage of Flip-Flops (FFs) of the Kintex UltraScale platform Table 3. Power analysis It has been found that although the ML engine will consume a low incremental power of around 150 mW, it does not introduce the 200-300% power increase in the identification of power-constrained CubeSat systems, rendering it an ideal fit.

Discussion of Trade-offs

Machine Learning integration into fault-tolerant systems creates a strategic tradeoff between the complexity of systems and reliability. Although the original cost of the computation in terms of time spent training the SVM and the minimal increment in the power consumption during this process are crucial, the benefit of the Reliability Gain and the area overhead reduction Figure 3 are extensive. Designers can use the saved FPGA fabrics by designing additional capability of purposes related to primary missions, like resolution imaging or advanced signal processing giant redundancy TMR. What is more, the change of the reactive to an intelligent recovery posture is so such that the system now is not just surviving a radiation event, but now has to actively optimise its own recovery strategy based on the nature of the fault identified that is a big step in the right direction towards completely autonomous, self-healing space electronics.

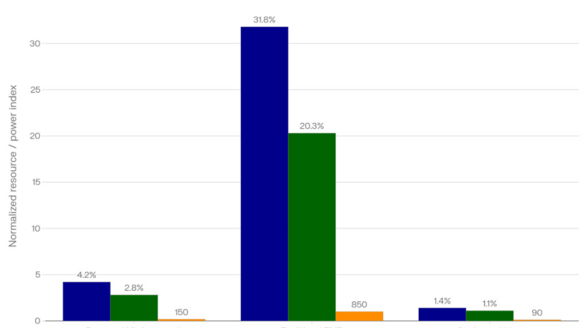

Fig.3: Comparison of LUT usage, FF usage, and power consumption for the proposed ML-based framework, traditional TMR, and static scrubbing, illustrating the reduced resource and power overhead of the ML-based mitigation relative to TMR.

Conclusion and Future Work

This study has managed to show that the combination of Machine Learning (ML) and Dynamic Partial Reconfiguration (DPR) offers a better solution to space-grade FPGA systems instead of the classic hardware redundancy one. The proposed framework can perform high-fidelity fault detection on hardware telemetry at very low resource footprint using a lightweight Support Vector Machine to classify hardware telemetry in real-time as compared to Triple Modular Redundancy (TMR). It is important because it allows small, power-constrained space vehicles, like CubeSats, to do complicated on-board processing under high-radiation conditions without compromising the reliability of the system and its mission lifetime. The methodology is further able to reduce recovery latency by 35% making sure that there is very little disruption in critical deep-space autonomous operations. The study recommends, in the future, the creation of decentralised implementations of Federated Learning, in which a system of satellites can use radiation telemetry patterns to jointly train and update their fault-detection engines. Through this shared intelligence, it would become possible to have a more capable orbital infrastructure that is able to respond to unprecedented occurrences of solar particles on a global scale.

References

1. Aunedi, M., Al Kindi, A. A., Pantaleo, A. M., Markides, C. N., & Strbac, G. (2023). System-driven design of flexible nuclear power plant configurations with thermal energy storage. *Energy Conversion and Management*, *291*, 117257.
2. Bhunia, S., Hsiao, M. S., Banga, M., & Narasimhan, S. (2014). Hardware Trojan attacks: Threat analysis and countermeasures. *Proceedings of the IEEE*, *102*(8), 1229-1247.

3. Bolchini, C., Miele, A., & Sandionigi, C. (2010). A novel design methodology for implementing reliability-aware systems on SRAM-based FPGAs. *IEEE Transactions on Computers*, *60*(12), 1744–1758.
4. Cannon, M. J., Keller, A. M., Rowberry, H. C., Thurlow, C. A., Pérez-Celis, A., & Wirthlin, M. J. (2018). Strategies for removing common mode failures from TMR designs deployed on SRAM FPGAs. *IEEE Transactions on Nuclear Science*, *66*(1), 207–215.
5. George, A. D., & Wilson, C. M. (2018). Onboard processing with hybrid and reconfigurable computing on small satellites. *Proceedings of the IEEE*, *106*(3), 458–470.
6. Glein, R., Rittner, F., & Heuberger, A. (2016, September). Detection of solar particle events inside FPGAs. In *2016 16th European Conference on Radiation and Its Effects on Components and Systems (RADECS)* (pp. 1–5). IEEE.
7. Glein, R., Rittner, F., & Heuberger, A. (2018). Adaptive single-event effect mitigation for dependable processing systems based on FPGAs. *Microprocessors and Microsystems*, *59*, 46–56.
8. Glein, R., Schmidt, B., Rittner, F., Teich, J., & Ziener, D. (2014, May). A self-adaptive SEU mitigation system for FPGAs with an internal block RAM ra diation particle sensor. In *2014 IEEE 22nd Annual International Symposium on Field-Programmable Custom Computing Machines* (pp. 251–258). IEEE.
9. Huang, J., Huang, S., Moghaddam, S. K., Lu, Y., Wang, G., Yan, Y., & Shi, X. (2024). Deep reinforcement learning-based dynamic reconfiguration planning for digital twin-driven smart manufacturing systems with reconfigurable machine tools. *IEEE Transactions on Industrial Informatics*, *20*(11), 13135–13146.
10. Khoshavi, N., Zarandi, H. R., & Maghsoudloo, M. (2012, May). Two control-flow error recovery methods for multithreaded programs running on multi-core processors. In *2012 28th International Conference on Microelectronics Proceedings* (pp. 371–374). IEEE.
11. Nguyen, N. T., Agiakatsikas, D., Cetin, E., & Diessel, O. (2016, December). Dynamic scheduling of voter checks in FPGA-based TMR systems. In *2016 International Conference on Field-Programmable Technology (FPT)* (pp. 169–172). IEEE.
12. Rodríguez, A., Valverde, J., Portilla, J., Otero, A., Riesgo, T., & De la Torre, E. (2018). FPGA-based high-performance embedded systems for adaptive edge computing in cyber-physical systems: The ARTICO3 framework. *Sensors*, *18*(6), 1877.
13. Wang, Z., Yao, Y., Tong, X., Luo, Q., & Chen, X. (2019). Dynamically reconfigurable encryption and decryption system design for the internet of things information security. *Sensors*, *19*(1), 143.