The Final Barrier to High-Performance Computing in Space: Architectures for Mitigating Destructive Radiation Effects

Yoav Simhony
School of Physics and Astronomy
Tel Aviv University
Tel Aviv, Israel
yoavsimhony@mail.tau.ac.il

Alex Segal
Unit of Mathematics
Afeka College of Engineering
Tel Aviv, Israel
alexs@afeka.ac.il

Ofer Amrani
Faculty of Electrical Engineering
Tel Aviv University
Tel Aviv, Israel
ofera@tauex.tau.ac.il

Erez Etzion
School of Physics and Astronomy
Tel Aviv University
Tel Aviv, Israel
ereze@tau.ac.il

***Abstract*—The increasing ambition of commercial and institutional space missions is driving demand for onboard computing capabilities far beyond traditional spacecraft avionics. Applications such as orbital data centers, large satellite constellations, autonomous mission operations, and real-time data processing increasingly depend on high-performance GPUs, advanced processors, and high-density solid-state storage. While commercial off-the-shelf (COTS) electronics offer unmatched capability and development speed, their reliable operation in space remains a central engineering challenge.**

Many established environmental hazards, including thermal-vacuum behavior, cumulative radiation effects such as total ionizing dose (TID) and non-destructive single-event effects (SEUs), are routinely mitigated with well-developed design practices. In contrast, Destructive Single Event Effects (DSEE) remain a dominant unresolved reliability barrier for the deployment of modern, highly integrated commercial electronics in orbit. Traditional DSEE mitigation relies on radiation-hardened components or per-device current monitoring, both of which become impractical for GPU-class systems characterized by large and rapidly varying currents.

This paper introduces a board-level, non-intrusive mitigation architecture that addresses DSEE. By detecting and characterizing incident particles, rather than hardening or detecting the implications of the phenomenon, it is possible to mitigate DSEEs. This architecture is evaluated by simulating a highly sensitive, COTS-based board, empirically known to exhibit DSEEs on a weekly basis, encapsulated within particle detectors. Using a comprehensive GEANT4 simulation of the cosmic radiation flux and a hybrid analytic-machine learning algorithm, we demonstrate that the system successfully mitigates DSEEs with a twenty-fold increase in board lifetime in orbit. This approach is anchored by empirical validation from a dedicated experiment aboard the International Space Station (ISS), which recorded over 500 particle events. The results indicate that this architecture identifies and mitigates DSEEs with high statistical confidence, enabling the practical deployment of advanced commercial electronics for next-generation orbital infrastructure.

Keywords: space computing, radiation effects, destructive single-event effects, latchup, mitigation architecture, particle detectors, GEANT4, machine learning, ISS

I. Introduction

Space missions increasingly demand onboard processing capabilities that approach terrestrial high-performance computing. Emerging concepts, including orbital data centers, onboard AI inference, advanced Earth-observation analytics, and autonomous mission operations, depend on GPU-class accelerators, high-end CPUs, and high-density storage subsystems [1]–[3]. Commercial off-the-shelf (COTS) electronics provide a rapidly evolving performance roadmap and a mature ecosystem, but their reliability in the space radiation environment remains a central obstacle to widespread deployment.

The space environment introduces multiple hazards to electronic systems, including destructive and non-destructive effects caused by radiation, as well as by the thermal, mechanical, and vacuum environments. For most of these hazards, scalable and non-intrusive mitigation practices are well established. Among particle-radiation effects, non-destructive effects such as single-event upsets (SEUs) are routinely mitigated through a suite of software- or firmware-based error detection and correction (EDAC), triple modular redundancy (TMR), and related fault-tolerance techniques. Cumulative radiation effects, such as total ionizing dose (TID), are becoming less pronounced as semiconductor technology nodes continue to scale. These technologies enable high-density and higher-performance electronics while considerably reducing transistor gate volumes and charge-trapping volumes, thereby limiting the severity of these effects and in-

creasing electronic component lifetimes in the space environment. Furthermore, cumulative radiation effects can be effectively mitigated through passive shielding, a solution that is becoming increasingly viable as the cost per kilogram of launching mass into orbit decreases.

In contrast, destructive single-event effects (DSEE), including destructive latchup and related irreversible failure modes, remain a limiting factor for modern, highly integrated commercial electronics in orbit. Although DSEE susceptibility depends on device structure, operating voltage, thermal conditions, and process-specific parasitic mechanisms, the challenge is amplified by increasing transistor density. Traditional DSEE mitigation approaches depend either on (i) radiation-qualified components, when available and suitable; (ii) radiation-hardened components, which are often costly; or (iii) when required and possible, per-component current monitoring and power cycling. The third approach is feasible mainly for devices with low and relatively constant current draw. These approaches become increasingly challenging for advanced VLSI technologies, such as modern GPUs and memory boards, because the main components can draw a current of hundreds of amperes to the core while exhibiting rapid load transients and operating within a densely integrated power-delivery architecture.

In addition, established practices for reliable space systems rely on using only radiation-qualified or radiation-hardened components within each electronic board and subsystem. This approach requires redesigning commercial electronic boards and subsystems specifically for space applications, precluding the direct adoption of high-performance electronic boards that have already been designed, manufactured, and tested in large quantities. Rather than adopting existing board-level or system-level commercial designs, the space industry relies on component-level design and integration. As a result, space technology is often unable to utilize the highest-performance electronics, leading to high-cost space systems that require long development cycles and are inherently outdated relative to available terrestrial COTS electronics.

This work presents a system-level radiation-mitigation architecture that enables a non-intrusive, active protection scheme around a COTS electronic board or subsystem, providing a key missing capability for DSEE mitigation. The architecture detects and characterizes impinging particles in real time, assesses the DSEE risk associated with each particle event, and, when required, mitigates this risk by power-cycling the protected electronic board [4], [5]. The architecture is evaluated through simulation and validated using a dedicated spaceborne experiment that provided event-by-event particle measurements in low Earth orbit (LEO) [6], enabling a direct comparison between modeled and measured environments.

II. Spacecraft Design Practices to Withstand the Space Radiation Environment

The space environment, and particularly the necessity of radiation mitigation, forces space systems engineers to trade off electronic system reliability and risk against performance, development time, and cost.

A. Component-Level Space Hardened Design Methodology

When prioritizing high reliability and low risk in this trade-off, the prevailing design practice is to develop space systems from the bottom up, relying exclusively on space-hardened components. Many VLSI design techniques have been devised and implemented to allow these specialized components to withstand random particle interactions [7], [8]. This method affords the designer full control over the risk and reliability (average lifetime) of the subsystem. However, it incurs heavy penalties in cost, performance, and time-to-orbit due to prolonged development cycles, frequently resulting in satellites launching today with 1990s-era computing capabilities.

B. Component-Level Space Qualified Design Methodology

Driven by the commercial electronics industry's mass production of low-cost, non-radiation-hardened components, an alternative bottom-up design methodology has emerged based on radiation-qualified components. In this method, commercial terrestrial components are screened to withstand the hazards of the space environment by irradiating them in particle accelerators until failure. After mapping these components, electronic subsystem designers can tailor the risk and reliability of the spacecraft to accommodate specific mission durations and orbits [5]. Frequently, when commercial components fail to withstand the required radiation environment, designers are forced to revert to space-hardened components, or otherwise compromise on system performance, reliability, or risk. In some cases, additional mitigation techniques, such as error detection and correction (against SEU), cold or hot redundancy, passive shielding (against TID), and current monitoring and power-cycling (against DSEE), can be implemented to improve survivability and meet minimum environmental requirements. While this space-qualified approach benefits from components that are orders of magnitude cheaper than space-hardened equivalents, the requisite radiation testing and lot control add additional costs and limit component availability. Ultimately, the major constraint of this method is its continued reliance on bottom-up, component-level design.

C. Board-Level Space Qualified Design Methodology

The next significant evolution in space architecture is the application of commercial board-level designs

to spacecraft subsystems. For example, incorporating modern commercial GPU boards for orbital data processing has the potential to enable high-performance computing while massively reducing development time and costs. Recent testing on advanced commercial GPUs has categorized their vulnerabilities across TID, SEU, and DSEE mechanisms. Although their baseline TID tolerance is adequate for many mission lifetimes and SEUs can be addressed through software fault-tolerance, testing indicates that several of these boards are susceptible to DSEEs, making them a limiting factor for space deployment [9]–[13].

Furthermore, validating radiation resilience at the board level is inherently challenging due to the thermal dissipation requirements, particularly when the board is tested under vacuum, together with the limited penetration depth of accelerator-beam particles within the chip, which may require chip thinning. This complexity can obscure true component vulnerabilities; for example, recent findings [10] demonstrated, after repeated testing, that certain GPU boards previously considered immune to DSEE are in fact susceptible. In certain missions, it is possible to modify commercial boards by replacing unqualified components from the boards periphery with space-qualified or hardened alternatives; however, this is often problematic, as reaping the full benefits of this method precludes printed circuit board (PCB) redesigns.

This board-level methodology also allows for the implementation of the already mentioned additional mitigation techniques, although these are often less effective than at the component-level approaches. In particular, board-level current monitoring and power cycling are largely ineffective for high-performance boards and components such as GPUs. Owing to their high power consumption, often ranging from hundreds to thousands of watts, and highly dynamic load transients, the relatively small power increase triggered by a DSEE, often on the order of tens of milliwatts, is easily masked. Furthermore, the power-supply circuitry on these boards includes regulators and capacitance that further obscure the current spikes associated with DSEEs.

A notable mitigation strategy within this board-level methodology is the implementation of hot and cold redundancy. For example, if a high-performance commercial board has an expected average lifetime of six months in orbit before exhibiting a DSEE, mission planners could deploy twenty or more such boards in cold redundancy, switching to a new board upon each failure to achieve a system-level lifetime of ten years. However, although this approach may achieve an average lifetime of ten years, there remains a statistical probability that the system will fail before reaching the required mission duration. Since DSEEs follow a Poisson process, the lifetime of a system of twenty identical boards follows a Gamma distribution. Hence, there exists a probability of approximately 53% that the system will fail in less than ten years. Therefore, if the system requires high confidence that it will survive for ten years in orbit, this must be taken into account, and a larger number of redundant boards is required. While this consumes considerable volume and mass, both for the duplicated electronics and for the overhead required by the connectivity matrix, today it may be the only acceptable path to achieving the required computational performance.

III. Active Radiation Mitigation - Detecting the Particle Rather than the Consequence

The active radiation mitigation architecture introduces a non-intrusive approach to mitigating DSEEs, while also mitigating cumulative radiation effects and reducing SEU occurrence, which can be further addressed through error detection and correction software. This is achieved by detecting and characterizing each impinging particle in real time and power-cycling the protected commercial electronics only for particle events that may induce a DSEE. Since DSEE damage is caused by thermal runaway and local heating, if power cycling is induced within a limited time frame, e.g., 1 ms, from the moment the particle traverses the protected electronics, no damage will occur [5].

A. Concept and Hardware Framework

Physically, the active radiation mitigation architecture places the commercial terrestrial electronics enclosed within particle detectors (see Figure 1. As a particle traverses the particle detectors to reach the protected electronics, the particle is detected by the enclosing detectors. The signals from the detectors are sent to a simple space-qualified computer which analyzes the information, characterizes the particle event, evaluates the likelihood of a destructive SEE and if warranted, performs an immediate power-cycle before irreversible thermal damage develops.

B. Practical Properties

This architecture enables the implementation of high-performance electronics in space, benefiting from the abundance and performance of the commercial terrestrial electronics industry.

Because it is non-intrusive, it enables board-level and subsystem-level design, rather than component-level design. This allows the use of boards that have already been designed, manufactured, and qualified for terrestrial markets, thereby reducing time to orbit as well as development and manufacturing costs through mass production.

Furthermore, it is compatible with accelerator-derived device characterization: if per-device susceptibility information exists, the decision logic may be tuned to reduce false triggers; if not, conservative operating assumptions may be used.

The cost of such a non-intrusive system is that the power-cycle rate can be significantly higher than the rate of true DSEEs, which affects system performance.

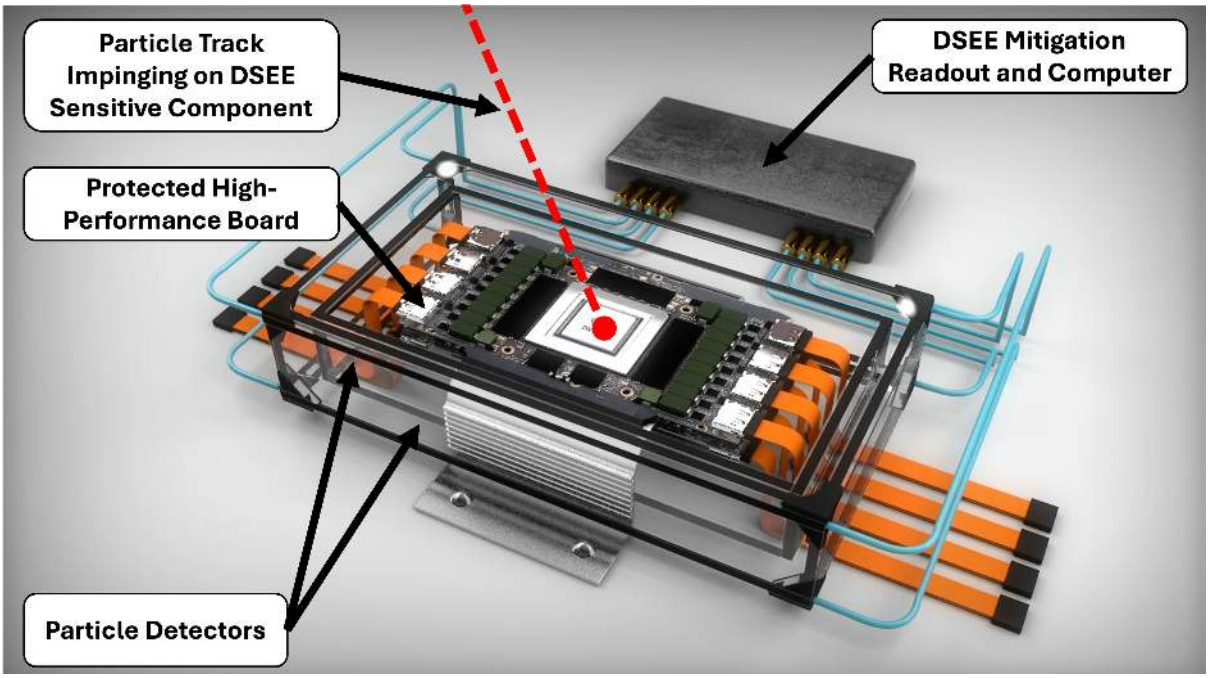


Fig. 1: Active radiation mitigation architecture. A high-performance commercial board is enclosed within a particle detector array. When a particle strikes a sensitive component on the board, the particle detectors register the event and transmit signals to a readout and computer board. The system then characterizes the event, determines whether it may induce a destructive SEE, and, if necessary, immediately power-cycles the protected board before thermal damage occurs.

Thus, a balance between the system sensitivity and the required uptime must be maintained.

IV. Evaluation and Validation: GEANT4 and Empirical ISS Measurements

To evaluate the mitigation architecture, a full-system GEANT4 particle transport simulation was constructed. We model a protected board composed of sensitive components that, under standard unmitigated orbital conditions, experience DSEEs on a weekly basis. This high-risk system is encapsulated within an array of particle detectors, mirroring the geometric configuration described in our prior physical designs [4].

A realistic cosmic radiation flux of ∼500,000 particles, with an spectral distribution derived using the CREME96 model [15], was simulated to interact with the encapsulated system. The transport of incident particles and secondary showers was governed by a standard hadronic physics list appropriate for space environments.

When an incident particle traverses the particle detectors, the resulting response is processed by a hybrid algorithm that combines analytical physical models with machine learning techniques such as gradient boosting [16]. This algorithm successfully recovers the incident particle's characteristics. A critical requirement of this hybrid approach is its execution speed; the inference pipeline is heavily optimized to ensure that the prediction latency remains well within the required temporal window, allowing the preemptive power cycle to initiate before a sustained high-current state induces permanent thermal damage to the component.

Based on these recovered particle event characteristics, the mitigation logic assesses the threat level and triggers a preemptive power cycle only when a threshold indicating potential DSEE risk is breached. The simulation shows that this approach successfully intercepts the particles responsible for destructive failures, effectively increasing the expected operational lifetime of the system by 7 to 20-fold, depending on the required mission parameters. This flexible lifetime improvement reflects the inherent trade-off between absolute system protection and continuous availability. By tuning the threshold of the classifier, the system can be biased toward extreme caution for critical operations—accepting a higher rate of false-positive power cycles, or biased toward maximum uptime in shorter lifetime or lower-risk orbital environments.

These simulations directly support our ongoing development of a full-scale orbital demonstrator system, scheduled for launch in March 2027. In this demonstrator, the protected payload will feature a sensitive array of ten ISSI IS61LV5128AL (IS61) SRAM components, enclosed within particle detectors.

The predicted mitigation performance is intrinsically linked to the physical susceptibility of these specific components. Ground characterization data, adapted from the PROBA-II Technology Demonstration Module campaign [14], establishes the heavy-ion DSEE

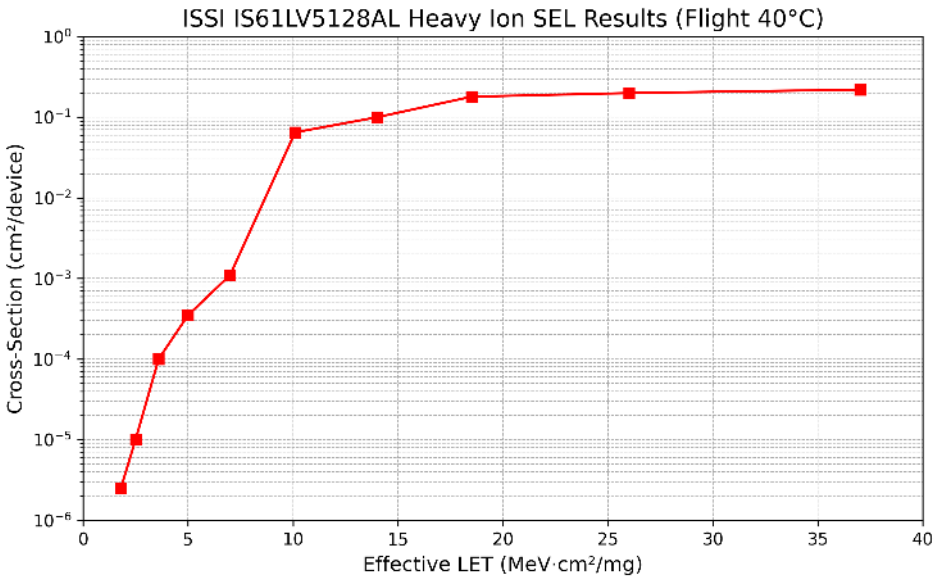


Fig. 2: Heavy ion SEL cross-section sensitivity for the ISSI IS61LV5128AL SRAM at a flight temperature of 40°C. Data adapted from [14].

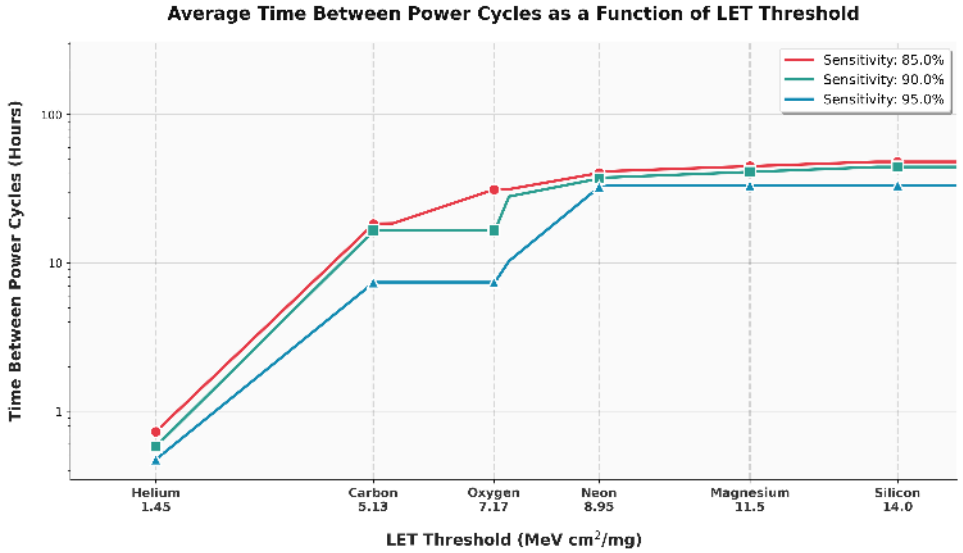


Fig. 3: Mean time between mitigation-induced power cycles as a function of the selected detection threshold. Tuning this threshold allows the system to balance a 7 to 20-fold increase in system lifetime against the operational overhead of periodic power cycling.

cross-section for the IS61 SRAM. As illustrated in Figure 2, this physical cross-section behaves as a continuous probability distribution, expanding as incident linear energy transfer (LET) increases until it approaches the total sensitive chip area. To maintain high system availability, the active mitigation architecture is designed to trigger mitigation only when the effective sensitive area expands into a statistically significant threat footprint.

Examining the characterization curve, the DSEE incidence transitions into a meaningful operational risk around an LET threshold of 5.0 MeV·cm^2/mg. Below this region, the cross-section is sufficiently low that the probability of a DSEE remains low. By anchoring our hybrid machine learning classification models near this threshold, the algorithm efficiently filters out low-risk interactions while reliably capturing critical threats.

Evaluated against our orbital simulation data at this specific operational threshold, our model reconciles this physical risk with strict availability constraints. Although the chosen IS61 components are extremely sensitive to DSEE, this threshold corresponds to a highly manageable mitigation overhead ranging from one to three preemptive power cycles per day (see Figure 3). Over a standard mission year, this approach is expected to successfully intercept hundreds of po-

tential DSEE failures while ensuring near-continuous operational throughout, affirming the architecture's capacity to protect commercial electronics in orbit.

Beyond these simulated projections, the efficacy of the active mitigation architecture is anchored by empirical validation from a dedicated experiment aboard the International Space Station (ISS). Capturing over 500 million independent particle events in LEO, the spaceborne hodoscope provided granular, event-by-event particle radiation characterization. As detailed in [5], this massive ground-truth dataset was utilized not only to strictly validate the GEANT4 orbital environment and simulated detector responses, but also to emulate real-time mitigation behavior. By passing the raw ISS event streams directly through the algorithmic decision layer, the system successfully demonstrated its ability to continuously isolate the rare heavy ions responsible for destructive events from the benign particles, establishing high statistical confidence in the architecture's deployability.

V. Conclusion & Discussion

High-performance onboard computing is essential for next-generation space systems. While cumulative radiation effects and non-destructive single-event upsets can be effectively mitigated through established engineering practices and fault-tolerant design, destructive single-event effects remain the primary barrier to deploying advanced commercial electronics in orbit.

This work presents a non-intrusive, system-level mitigation architecture that addresses this challenge by detecting and characterizing incident particles in real time, and preemptively power cycling protected electronics before the onset of localized thermal runaway. By shifting the mitigation paradigm from reacting to failure signatures to predicting high-risk particle events, the architecture enables effective protection without modifying the underlying hardware.

Evaluated using full-scale GEANT4 particle-transport simulations and a hybrid machine-learning decision framework, the proposed system demonstrates the ability to extend the operational lifetime of high-performance commercial electronics by a factor of twenty. These results are further supported by empirical validation using over 500 million particle events recorded by a dedicated ISS-based hodoscope, confirming the system's ability to reliably distinguish high-risk events from benign radiation.

Taken together, these results establish the feasibility of deploying high-performance commercial electronics in space through non-intrusive, system-level protection. This approach provides a scalable pathway toward bridging the performance gap between terrestrial and spaceborne computing, enabling a new class of data-intensive and autonomous orbital systems.

References

[1] S. Wang and Q. Li, "Satellite computing: Vision and challenges," vol. 10, no. 24, 2023, pp. 22 514–22 529.

[2] Y. Zengshan, W. Changhao, G. Chongbin, L. Yuanchun, X. Mengwei, G. Weiwei, and C. Chuanxiu, "A comprehensive survey of orbital edge computing: Systems, applications, and algorithms," *Chinese Journal of Aeronautics*, vol. 38, no. 7, p. 103316, 2025.

[3] Voyager Technologies, "The silent architecture of golden dome: Electronics, sensors and communications," Technical Insight Report, Tech. Rep., Jan 2026.

[4] Y. Simhony, A. Segal, Y. Orlov, O. Amrani, and E. Etzion, "Scintillator-sipm detector for tracking and energy deposition measurements," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1069, p. 169955, 2024.

[5] Y. Simhony, A. Segal, O. Amrani, and E. Etzion, "High-end space electronics: Active shielding to mitigate catastrophic single-event effects," *Aerospace Science and Technology*, p. 111344, 2025.

[6] Y. Simhony, A. Segal, Y. Orlov, D. Bashi, O. Amrani, and E. Etzion, "Spaceborne cots-capsule hodoscope: Detecting and characterizing particle radiation," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1070, p. 169996, 2025.

[7] T. V. Reddy and S. Nakhate, "Radiation hardened by design technique to mitigate single event transients in combinational logic circuits," in *2017 International Conference on Recent Innovations in Signal processing and Embedded Systems (RISE)*. IEEE, 2017, pp. 342–347.

[8] A. S. Keys and J. T. Howell, "Technology developments in radiation-hardened electronics for space environments," Tech. Rep., 2008.

[9] I. A. Troxel, J. J. Schaefer, M. Gruber, D. Sabogal, D. Ellis, J. Schaf, and G. West, "Heavy ion and proton test results for recent-generation gpus," in *2021 IEEE Radiation Effects Data Workshop (REDW)*. IEEE, 2021, pp. 1–6.

[10] M. Cannon, W. Evans, D. Lee, I. Troxel, M. Gruber, and D. Sabogal, "Proton and heavy ion see data on nvidia and amd gpus," in *2023 IEEE Radiation Effects Data Workshop (REDW)(in conjunction with 2023 NSREC)*. IEEE, 2023, pp. 1–8.

[11] F. C. Bruhn, N. Tsog, F. Kunkel, O. Flordal, and I. Troxel, "Enabling radiation tolerant heterogeneous gpu-based onboard data processing in space," *CEAS Space Journal*, vol. 12, no. 4, pp. 551–564, 2020.

[12] W. S. Slater, B. B. Rutherford, J. K. Mee, R. E. Pinson, M. Gruber, D. Sabogal, and I. A. Troxel, "Single event effects and total ionizing dose radiation testing of nvidia jetson orin agx system on module," in *2023 IEEE Radiation Effects Data Workshop (REDW)(in conjunction with 2023 NSREC)*. IEEE, 2023, pp. 1–6.

[13] W. S. Slater, N. P. Tiwari, T. M. Lovelly, and J. K. Mee, "Total ionizing dose radiation testing of nvidia jetson nano gpus," in *2020 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 2020, pp. 1–3.

[14] R. Harboe-Sorensen, C. Poivey, A. Zadeh, A. Keating, N. Fleurinck, K. Puimege, F.-X. Guerre, F. Lochon, M. Kaddour, L. Li *et al.*, "Proba-ii technology demonstration module in-flight data analysis," *IEEE Transactions on Nuclear Science*, vol. 59, no. 4, pp. 1086–1091, 2012.

[15] A. J. Tylka, J. H. Adams, P. R. Boberg, B. Brownstein, W. F. Dietrich, E. O. Flueckiger, E. L. Petersen, M. A. Shea, D. F. Smart, and E. C. Smith, "Creme96: A revision of the cosmic ray effects on micro-electronics code," *IEEE Transactions on Nuclear Science*, vol. 44, no. 6, pp. 2150–2160, 1997.

[16] Y. Simhony, A. Segal, O. Amrani, and E. Etzion, "From light to energy: Machine learning algorithms for position and energy deposition estimation in scintillator–sipm detectors," *Sensors*, vol. 26, no. 1, p. 101, 2025.