NASA/TM– 20250009692

EdgeCortix SAKURA-II Machine-Learning Accelerator SEE Heavy Ion Test Report

Seth S. Roffe
Goddard Space Flight Center, Greenbelt, MD

Scott D. Stansberry
Columbus Technologies and Services, Inc., Greenbelt, MD 20770

Edward J. Wyrwas
Columbus Technologies and Services, Inc., Greenbelt, MD 20770

Jeffrey Grosman
EdgeCortix, Kawasaki, Kanagawa, Japan

Stan Crow
EdgeCortix, Kawasaki, Kanagawa, Japan

Uzzal Podder
EdgeCortix, Kawasaki, Kanagawa, Japan

September 2025

NASA STI Program Report Series

The NASA STI Program collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/ and select the "General" help request type.

NASA/TM–20250009692

EdgeCortix SAKURA-II Machine-Learning Accelerator SEE Heavy Ion Test Report

Seth S. Roffe
Goddard Space Flight Center, Greenbelt, MD

Scott D. Stansberry
Columbus Technologies and Services, Inc., Greenbelt, MD 20770

Edward J. Wyrwas
Columbus Technologies and Services, Inc., Greenbelt, MD 20770

Jeffrey Grosman
EdgeCortix, Kawasaki, Kanagawa, Japan

Stan Crow
EdgeCortix, Kawasaki, Kanagawa, Japan

Uzzal Podder
EdgeCortix, Kawasaki, Kanagawa, Japan

Test Date: 9/18/2025
Report Date: 9/29/2025

National Aeronautics and
Space Administration

Goddard Space Flight Center
Greenbelt, MD 20771

September 2025

Acknowledgments

This work was sponsored by NASA Electronic Parts and Packaging (NEPP) Program and supported by the Defense Innovation Unit.

Available from

NASA STI Program
Mail Stop 148
NASA's Langley Research Center
Hampton, VA 23681-2199

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
703-605-6000

This report is available in electronic form at

https://nepp.nasa.gov/

1. Introduction and Purpose

To enable autonomy in space, machine-learning and computer vision applications become invaluable for sensor processing. However, these algorithms are computationally complex and unfeasible for many embedded central processing units (CPUs) and usually require external coprocessors, such as graphics processing units (GPUs) or accelerators specific to the application, including application specific integrated circuits (ASICs). In power-constrained systems, GPUs tend to consume more power than is acceptable (>40W), so lower-power accelerators have shown promise to provide the performance needed under spacecraft constraints. For radiation engineers, developing methodologies that can properly test CPUs, GPUs, and accelerators, and enable comparisons between them remains a necessary complication to solve as the devices become more complex. The methodology in this test aims to be a start in developing a baseline single-event effect (SEE) test for client-device machine learning accelerators. This category of devices does not host their own operating system.

This testing campaign is a continuation of a previous 200 MeV proton test performed in January 2024 [1] and a heavy ion test campaign performed in April and June of 2024 [2]. Specifically, this testing campaign evaluates the next generation card from EdgeCortix, the SAKURA-II machine-learning accelerator [3].

This experiment characterizes SEEs and data error susceptibility of the SAKURA-II accelerator under heavy ions. The device was monitored for single event upsets (SEUs), single event functional interrupts (SEFIs) and Single Event Latchup (SEL) at the Texas A&M K500 cyclotron. The SAKURA-II board accelerates machine-learning inference applications on a host computer through an m.2 x16 connector. For the purposes of devising an end-to-end automated analysis workflow for this experiment and a comparison to the SAKURA-I chip, the YOLO-V5 objection-detection model and the ResNet-50 classification models were used as a representative suite of analytical machine-learning models.

2. Test Result Summary

The SAKURA-II card was irradiated with the 25 $\frac{MeV}{amu}$ beam tune at the TAMU K500 cyclotron. Linear Energy Transfers (LETs) were tested up to 40.9 $MeV\ \frac{cm^2}{mg}$ at a fluence of >1E7 cm^2. A fitted Weibull curve for SEFIs is presented in this report, with an onset LET of 0.9 $MeV\ \frac{cm^2}{mg}$ and a limiting cross section of 1.00E-04 cm^2.

The SAKURA-II card did not experience any destructive SEEs nor any observed SEL during the heavy ion testing campaign. The most commonly observed effects were SEFIs within the PCIe interface. SEUs were observed in the form of changes in the output confidence scores of the models. Tolerable SEUs that were observed in the object-detection models were defined as inferences where the confidence scores were different from the expected value, but the objects in the image were correctly identified with only minor changes in the bounding box. For the image classification models, tolerable errors were described by changes in the output confidence scores that did not lead to a misprediction. Conversely, there were also several cases where the model

began consistently predicting incorrect objects from what was expected, requiring a system reset to fix the issue.

Similar to the results of the SAKURA-I chip, tolerable upsets were self-recovered on the next inference due to refreshing from off-chip DRAM that was not under irradiation. However, there were events in which the scores were permanently altered, and a system reboot was needed to recover.

3. Device Description

The device-under-test (DUT) was the SAKURA-II card, an m.2 2280 ASIC accelerator designed to accelerate inference on machine-learning applications on the edge. It contains 20 MB of on-chip memory and 16 GB of external LPDDR4. The interface for the card connects to a host PC via an m.2 2280 x16 slot. Further details can be seen in Table 1. A picture of the card used in the test can be seen in Figure 1

For test preparation, the ASIC SAKURA-II card, was decapsulated and thinned to 100 μm by Sage Analytical Labs in San Diego, CA. The card, after decapsulation and thinning, can be seen in Figure 2. The depth maps for each card from Sage Analytical Labs can be seen in Figure 3.

Table 1: SAKURA-II Card Details [1]

Part	SAKURA-II Edge Accelerator
REAG ID	25-026
Manufacturer	EdgeCortix
Cache	20MB
External Memory	16 GB LPDDR4
Reported Performance	60 TOPS (INT8)
Reported Power	8W
Interface	m.2 2280 form factor PCIe Gen 3.0 x16

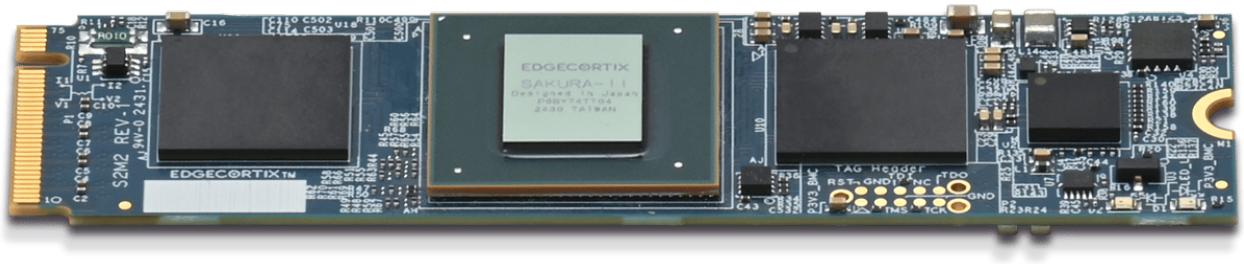

Figure 1: Image of the m.2 SAKURA-II Card [3]

Figure 2: De-lidded and thinned SAKURA-II card

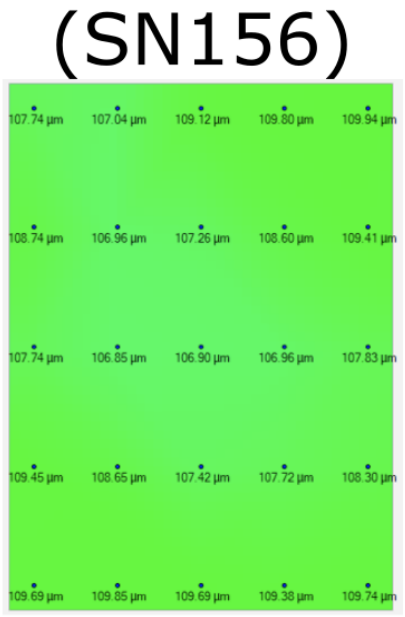

Figure 3: Depth map from Sage Analytical for the tested SAKURA-II device and its respective serial number.

4. Test Setup

The SAKURA-II card was mounted and clamped at normal incidence to the end of the beam line approximately 1 cm away from the beam port. and connected to a host PC running Ubuntu 20.04 via a 0.5 m PCIe extension cable. Additionally, a modified m.2 X16 bus breakout card with external power was used to isolate the current draw of the device from the rest of the test bench. A one-inch collimator was used to narrow the beam to roughly the size of the ASIC. Two images of the test setup can be seen in Figure 4. Compressed air was used to cool the ASIC in a laminar flow (noted by the orange nozzle in the picture on the right).

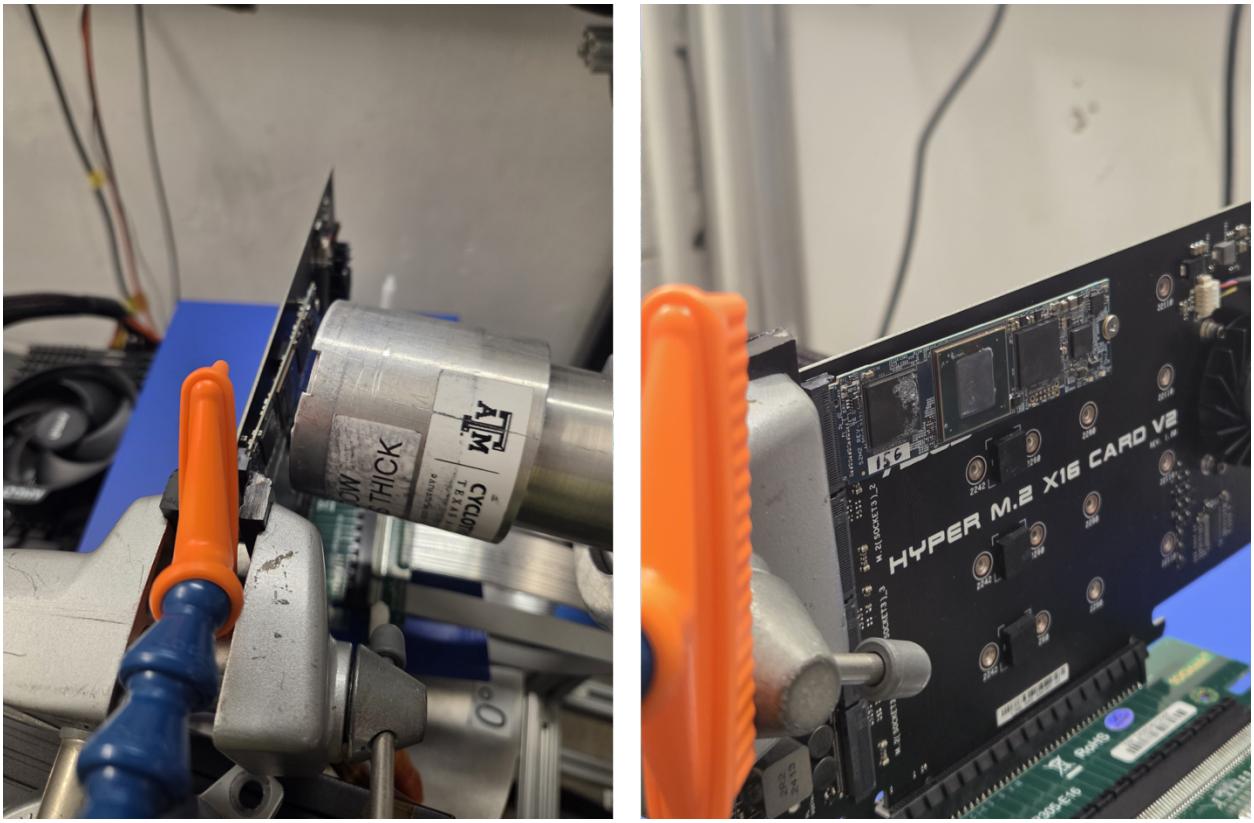

Figure 4. Test setup of the radiation test.

5. Test Facility

Facility:	Texas A&M K500 Cyclotron
Type of Radiation:	Heavy Ion
Energy:	25 $\frac{MeV}{amu}$

6. Test Conditions

Temperature:	Room Temperature
In-Air or Vacuum:	In-air
Supply Voltages:	12 V
Distance to beam port:	~1 cm

7. Test Methods and Procedures

This section covers the methodology used in this experiment. Details about the model choices and how runs were defined are discussed herein.

7.1. Model Selection

YOLOv5m and ResNet50 was chosen as the object detection and image classification models to focus on since it was the model that was used in the high energy proton test and the heavy ion tests done on the SAKURA-II card. The sample sizes of each model in this experiment can be seen in Table 2.

Table 2. Sample sizes for each model.

Model	Sample Size
ResNet-50	21
YOLOv5m	19

The YOLOv5 [4] object detection models were used with images selected from the Common Objects in Context (COCO) 80 dataset, which contains 80 classes [5] of objects. To control the input data to the model, only one image was cycled due to the number of objects that would be detected within it, after significant input dependence was not observed in [6] and [1]. By reducing the inputs to a small number, thereby limiting the overall quantity of objects that could be observed across all input images, the output vectors will be consistent between runs without introducing the additional data size of having too many input vectors. Restricting the experiment to only one image, keeps the test realistic to a real-world case, where there would likely be inferences on only one item at a time from a continuous feed, while still analyzing how different

confidences affect the results. Analyzing only one image also acts as an experimental control to make SEU analysis easier. The images used for this test can be seen in Figure 3. An example of a correctly classified outcome can be seen in Figure 5.

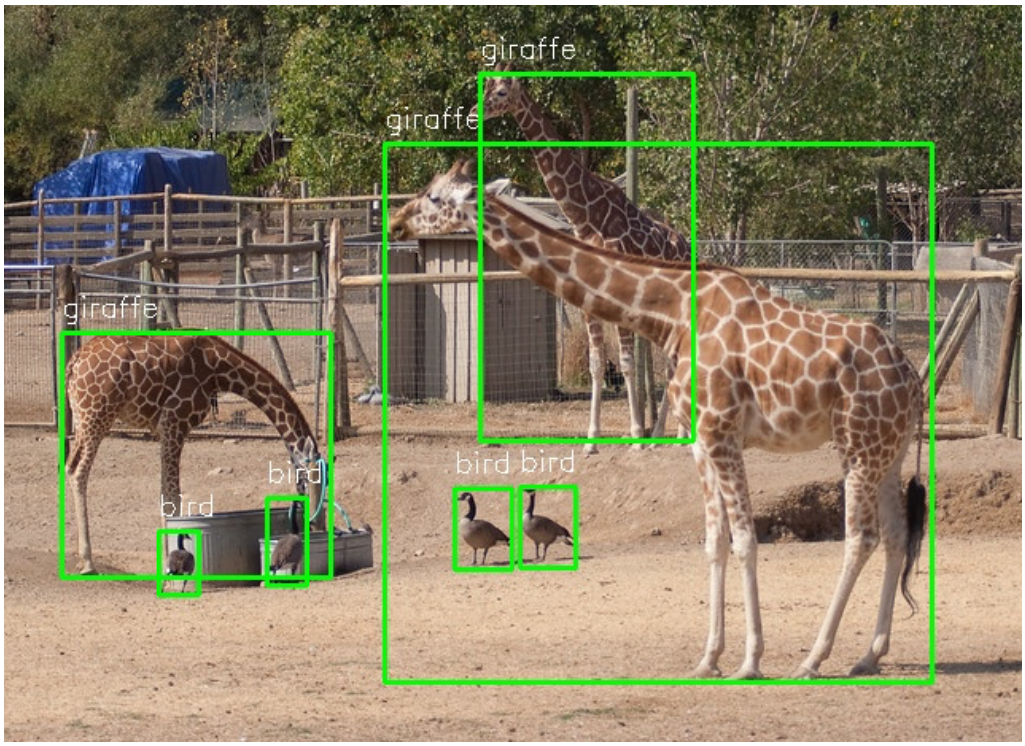

Figure 5. Output of object detection algorithm for a single image

For image classification, the ResNet50 model was used with images from the ImageNet dataset [7], which contains 1000 classes of objects. The same images were used for every run as an experimental control. Image Classification, as a contrast to object detection, which identifies several classes of objects within an image, assigns just one class to the object(s) within an image. An example of an image used in this test can be seen in Figure 7, with the class of “cat.”

Figure 6. ImageNet image of a cat used for classification [3].

7.2. Run Methodology

The test performed in this experiment was an extension of the high energy proton test performed in [1] and an equivalent of the heavy ion test performed in [2]. The device memory was programmed with a known pattern, irradiated, and then evaluated for single-event upsets. Two test conditions were used in this test for both image classification and object detection with varying experimental controls. These tests help understand the nature of any upsets seen within the SAKURA-II chip. These experiments were the following:

1. Running repeated inferences without any data transfers with the host beyond the initial setup with one input image. This allows us to observe any degradation in the output, if any, when inferences are run without constantly updating model parameters from the host.
2. Running repeated inferences with a reloading of the model parameters after a specified number of iterations. This allows us to observe if there is a recovery from recent upsets, if any, upon a refresh from the host PC.

Since the host PC was not in the beam path, and any critical or relevant data is passed to the DUT on a run start, the PC was only rebooted when there was a system hang, or PCIe communication with the card was upset. At the beginning of each new beam run, all model data and configuration were sent from the host PC to the DUT. The beam was powered on simultaneously with a run start.

For each run, the number of inferences was defined by a command line argument on start. Additionally, a command line argument was used to define the number of model parameter reloads if the run was following experiment #2. For both test conditions, the image with the giraffes and geese were used for object detection, shown in Figure 5, and the image of the cat was used for image classification, shown in Figure 6. After a run, the fluence was recorded along with the output confidence scores for all objects detected within the image. Additionally, the confidence score for all 80 classes within the dataset for each object were also recorded to observe any changes within any other class score, even if it was not the predicted output.

7.3. Data Analysis

To analyze the data, the confidence scores of the top scoring class for each detected object or image was plotted against the inference iteration number. This provides a timeline of how the confidence scores change over time. The model should be deterministic without any upsets. In other words, without radiation, the plot of confidence score vs inference iteration number should be a straight, horizontal line, both in object detection and in image classification. An example of this deterministic behavior in object detection and image classification can be seen in Figure 7 and Figure 8, respectively, which were run prior to the radiation experiment.

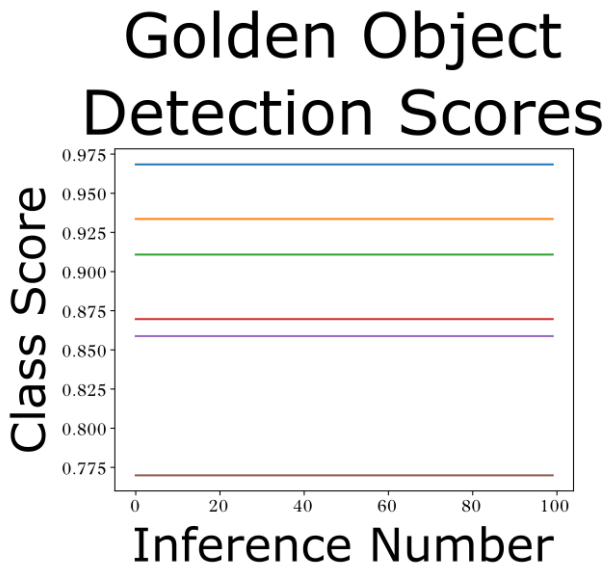

Figure 7. Golden confidence scores for one image of Object Detection

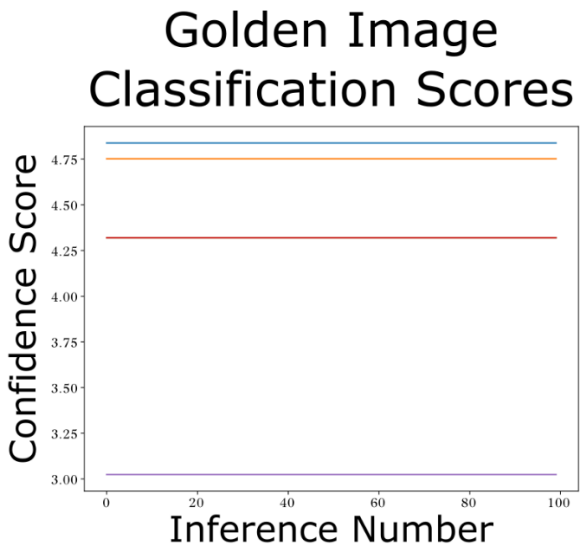

Figure 8. Top five golden confidence scores for Image Classification

The confidence scores were monitored for any upsets or deviations from the expected outcome. An error was considered tolerable if the expected output deviated but did not lead to a misprediction and recovered on the next inference. Any deviation from these golden scores is considered an upset, even if they were tolerable errors.

8. Test Results

The SAKURA-II card was irradiated with the 25 $\frac{MeV}{amu}$ beam tunes at Texas A&M's K500 cyclotron. Upsets were seen in the form of a drop or change in confidence score of the output classes. No destructive effects were observed during this heavy ion test, through a LET of 40.9 $MeV\frac{cm^2}{mg}$ at a fluence up to >1E7 cm^{-2}.

8.1. Object Detection Single-Event Effects

The single event effects observed were similar in type as those seen in [2]. For object detection, tolerable upsets were observed in the form of a small change in the output scores and bounding boxes for all objects within the image. Minor mispredictions were observed but were corrected upon the next inference, likely due to a cache refresh from the DDR outside the beam path. The small changes in the bounding boxes were likely due to the non-maximum suppression algorithm embedded in the model selecting a different, slightly less optimized bounding box for the detection. Persistent errors were observed where the upsets behaved in an uncontrolled fashion until a full reset of the system. These persistent errors were potentially tolerable and intolerable. Finally, there were catastrophic model failures where the outputs were entirely incorrect, no objects in the image were correctly detected or no objects were detected at all, and there were many false positives detected in the corners of the image. A more detailed description of the upsets observed, along with plots and image examples can be read about in the proton test report for the SAKURA-I card [1].

8.2. Image Classification Single-Event Effects

For many of the upsets seen, the error signature was a tolerable, small change in the confidence scores that were immediately recovered to its expected state on the next iteration. This is likely due to model data being transferred to the on-chip memory from off-chip DDR memory, which was not under irradiation, on the start of every iteration, even if there are no data transfers with the host PC. These temporary upsets seemed tolerable and did not impact the overall accuracy of the model across the inference iterations. This result is consistent with the results found during the proton test and the heavy ion tests [1][2] of the SAKURA-1 card.

Tolerable, non-persistent upsets were observed as small spikes in the output scores. An example of this kind of non-persistent error for the top five scoring classes of ResNet50 for the cat image can be seen in Figure 9.

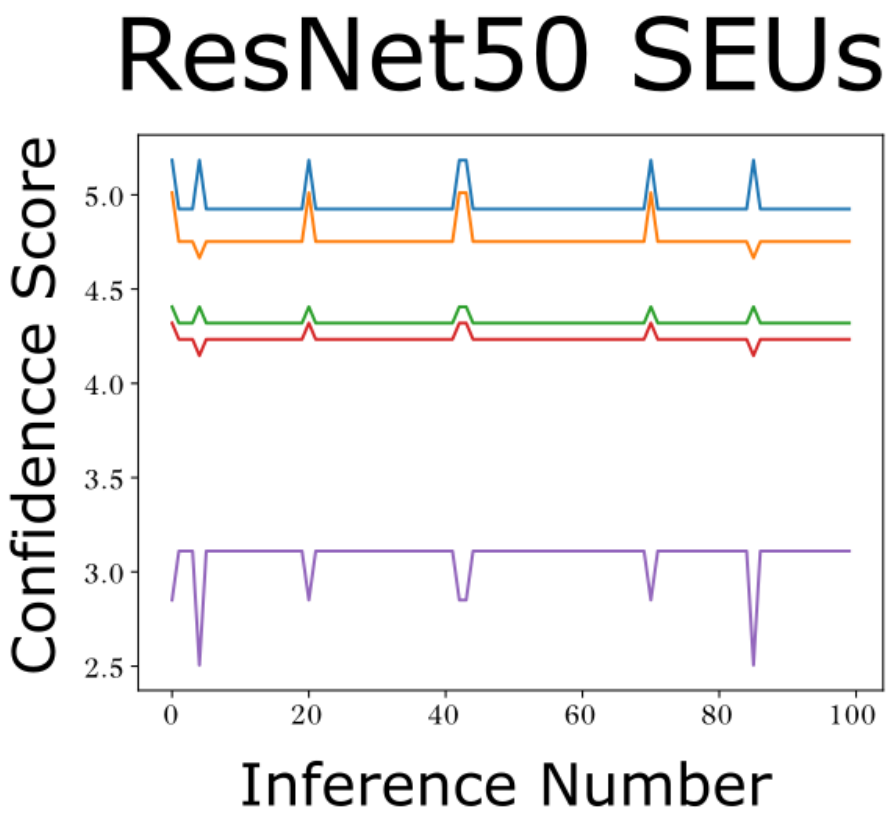

Figure 9. Example of image classification tolerable errors in the top five scoring classes which recover on the next inference.

Similar to object detection and the results seen in [2], the image classification models also saw persistent upsets. These persistent upsets required a full restart of the system to recover. Two examples of a persistent upset in the top scoring class of image classification can be seen in Figure 10 and Figure 11.

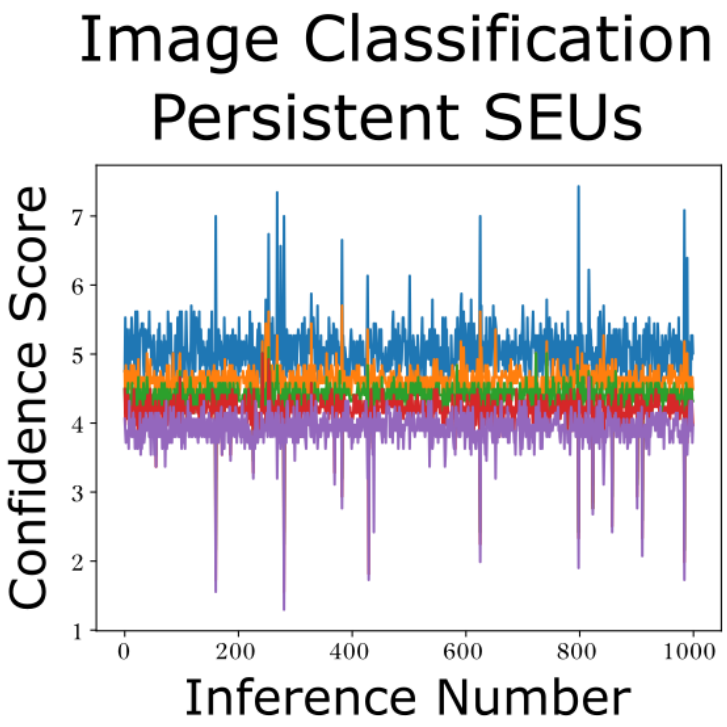

Figure 10. Persistent upset example in ResNet50.

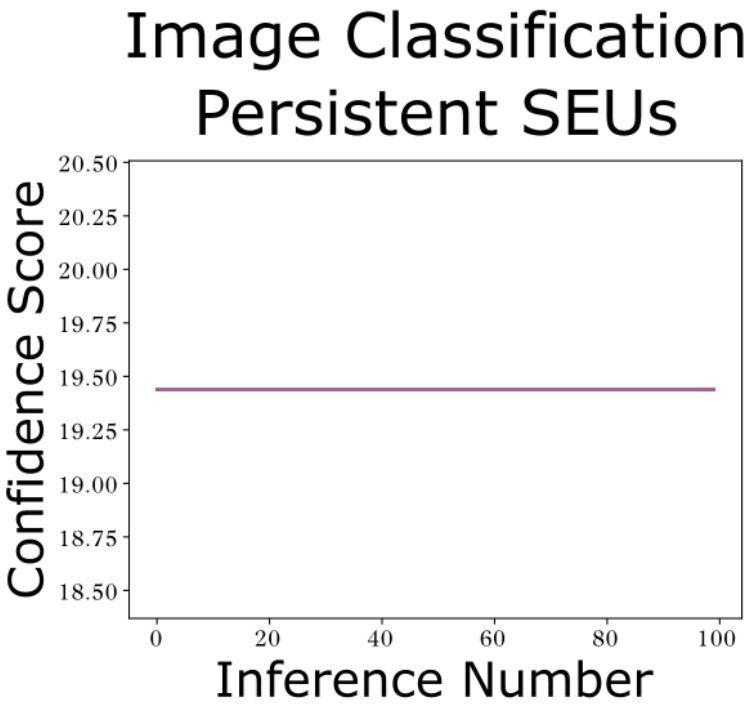

Figure 11. Persistent SEU in ResNet50 where the class scores all converge.

8.3. Fluence to Failure of Persistent Upsets

Both object detection and image classification saw persistent upsets that required a full system reset to recover. The fluence to persistent upset of the SAKURA-II card was measured. This measurement included the fluence up to the point of a persistent error, called fluence to failure (FTF). A persistent error for this measurement was any error that persisted for 5 or more iterations, any error that needed a reboot to recover, or any error that led to more than the expected number of detected objects in the case of object detection models.

The mean fluence to failure (MFTF) for each model was measured, seen in Figure 12 along with standard deviation. Both models showed very high errors due to low sample size. However, [2] found no significant differences between models and MFTF, and it is expected that this would be consistent here as well.

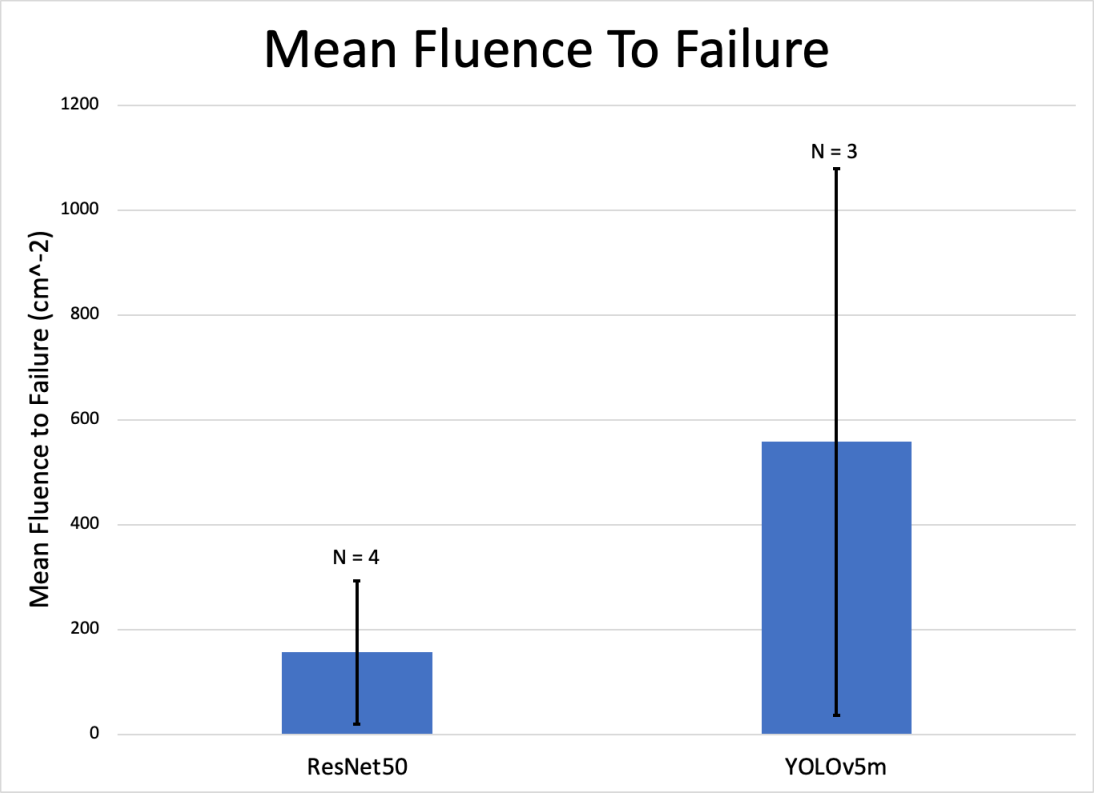

Figure 12. Mean Fluence to Failure for all models tested with standard error and sample size.

8.4. Single-Event Functional Interrupts

Most of the upsets that ended a run were categorized as single-event functional interrupts (SEFIs). These SEFIs were often seen and defined as losses in the communication with the device and required a restart of the system to recover, or a loss of ability to write data out. As expected, the cross section of the SEFIs increased with LET, as seen in the Weibull curve fit in Figure 13. These errors likely consisted of either an upset within the m.2-PCIe hardware in the card or a system hang which led to a loss in data telemetry. Additionally, the Weibull parameters can be seen in Table 3. Throughout the entire heavy ion test, the SAKURA-II never experienced

any permanent or destructive effects up to Xe with an LET of 40.9 $MeV\frac{cm^2}{mg}$. Inferences performed after the beam test performed as expected.

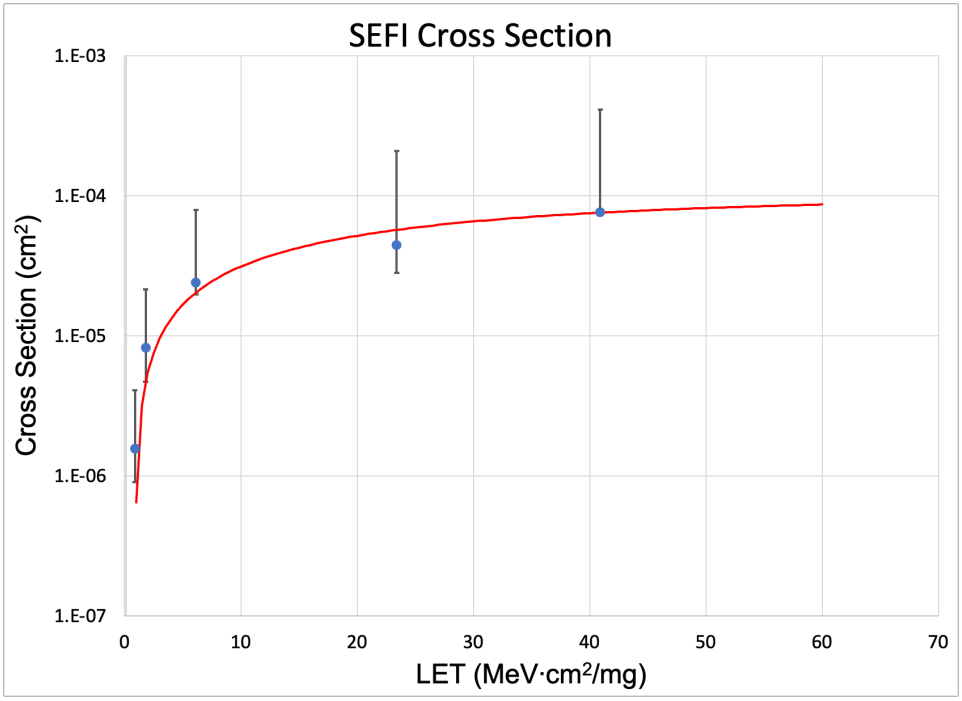

Figure 13. SEFI cross section and Weibull fit for the SAKURA-II.

Table 3. Fitted Weibull Parameters

Parameter	Fitted Value
Onset LET	0.9 $MeV\frac{cm^2}{mg}$
Limiting Cross Section (Max)	1.00E-04 cm^2
Shape	0.9
Width	27.0 $MeV\frac{cm^2}{mg}$

To compare the results of the SAKURA-II card to the SAKURA-I [2], the SEFI cross section curves were overlayed. This overlay shows that the SEFI results between the SAKURA-I and SAKURA-II are effectively identical. This is the expected result as the test models were the same, and the lithography between the two generations are identical. The difference in m.2 vs full-size PCIe as the connector shows no difference likely due to m.2 following the same control logic and drivers as PCIe. The comparison between the SEFI curves between SAKURA-I and SAKURA-II can be seen in Figure 14.

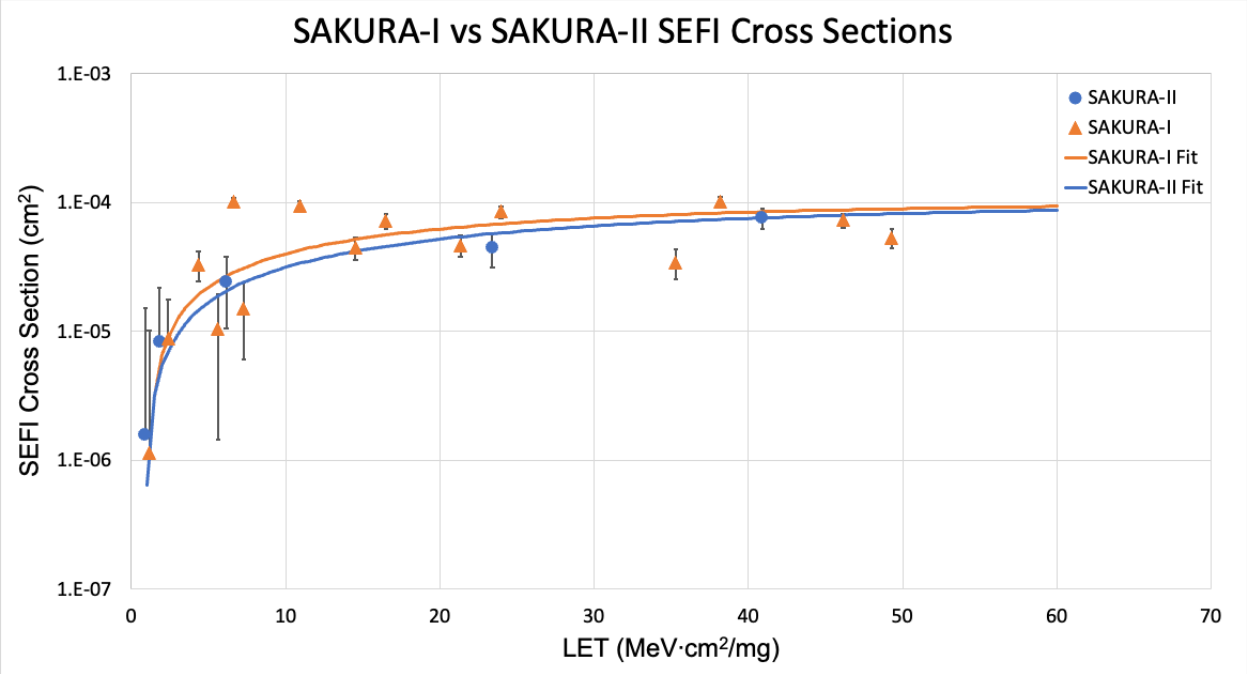

Figure 14. Comparison in SEFI cross sections between SAKURA-I and SAKURA-II.

An interesting implication of the results shown in Figure 14 is that the test methodologies between the heavy ion tests became more refined for the SAKURA-II test. There is less variation in the cross section results in the more recent test due to understanding the error signatures of the card prior to irradiation which enables more accurate fluence measurements when upsets occur. The methodologies used in this test will hopefully aid in the experimental design of future radiation tests on machine-learning accelerators.

9. Conclusion

This test report evaluates the second generation of the EdgeCortix machine-learning accelerator, the SAKURA-II. This work compares the results to the heavy ion and proton tests of the SAKURA-I card, evaluated in 2024 [1][2]. Most of the errors were hangs and crashes in the system that required a reboot to recover, classified as a SEFI. Of the data errors, most of the SEUs observed were temporary changes in the class confidence scores which recovered on the next inference iteration. However, some upsets saw persistent errors in the model which required a power cycle of the host PC. Throughout the entire test campaign, the SAKURA-II card did not experience any observe latchup events up to an LET of 40.9 $MeV\frac{cm^2}{mg}$ at a fluence up to 1E7 cm^{-2}. The results from the SAKURA-II card shows extremely similar results to that of the SAKURA-I card.

10. Appendix

This section summarizes the orbital correlation results pertaining to the Lunar space environment for the SAKURA-II device. This device was subjected to accelerated radiation testing using Heavy Ions. The statistical data from those tests was used with SPENVIS to calculate on-orbit single event upset rates. SPENVIS is ESA's SPace ENVironment Information System, an internet interface to models of the space environment and its effects, including galactic cosmic rays and Solar energetic particles. Because the failure metric in the radiation test data was a single event functional interrupt (SEFI), which requires a device power cycle, the resultant orbital rates from SPENVIS are for SEFIs. Mitigations such as redundancy or planned system resets may be necessary to assure availability of the device within a system.

Weibull Parameters used in the following rate calculations are those found in Table 3.

Table 4: Radiation Test Data Summary

Avg Cross Section	Degrees of Freedom	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
7.63609E-05	9	2.82171E-04	4.14289E-04
4.461E-05	8	1.63415E-05	2.09454E-04
2.41407E-05	6	4.40106E-06	7.93099E-05
8.246E-06	5	3.5232E-06	2.15469E-05
1.57591E-06	5	6.73328E-07	4.11788E-06

Lunar:

This environment is defined as a 385K km near-earth interplanetary orbit and referred to as Lunar Cruise. Guidance from NASA Ames Research Center's avionics trade study regarding Mission Radiation Environment Modeling and Analysis for their GCD Rad-Neuro Project (NASA/TM-20220011775, pages 10 and 11) was utilized to identify the orbital parameters. "The Moon has no atmosphere nor magnetosphere. At this distance from Earth, only the long tail of the magnetosphere proves any protection. Therefore, the Lunar mission environment is the same as the Cruise Phase radiation environment with one notable exception." The notable exception is the aggregation of shielding from the spacecraft, the Lunar body, and the Earth. The calculations for the Sakura-II are for worst case and do not include any consideration for shielding from these larger bodies. The critical time windows identified for this analysis were 2-hours and 15 days. The mission duration is 1 year. The results show that to assure system availability a system-level mitigation should verify functionality of the Sakura-II at least once every few days.

Table 5: Radiation Models Used for GEO

Trapped Radiation Models	AP-8, AE-8
Trapped proton anisotropy	Badhwar & Konradi 1990 MAX
Solar particle fluxes	CREME-96
Solar particle fluences	ESP-PSYCHIC
Galactic Cosmic Rays Spectra	ISO-15390
Shielded Flux	MFLUX

Table 6: Single Event Functional Interrupt Rates for Lunar Cruise Environment

Environment	SEFI/second
Worst Week - Solar Max	1.14E-07
Worst Day - Solar Max	3.23E-07
Worst 5 minutes - Solar Max	1.17E-06
Worst Week - Solar Min	1.14E-07
Worst Day - Solar Min	3.23E-07
Worst 5 minutes - Solar Min	1.17E-06

Table 7: Survivability Percentage for Lunar Cruise Environment

Survivability per Environment	120 minutes (2-hours)	1440 minutes (1-day)	21600 minutes (15-days)
Worst Week - Solar Max	99.918%	99.836%	86.258%
Worst Day - Solar Max	99.767%	99.535%	65.760%
Worst 5 minutes - Solar Max	99.161%	98.330%	21.960%
Worst Week - Solar Min	99.918%	99.836%	86.258%
Worst Day - Solar Min	99.767%	99.535%	65.760%
Worst 5 minutes - Solar Min	99.161%	98.330%	21.960%

Table 8: Total Mission SEFIs for 1-year Mission Duration in Lunar Cruise Environment

	Direct Ionization
Solar Minimum	2.4971E-01
Solar Maximum	2.4971E-01

LEO:

The environment for this study is defined as 800 km altitude with a 98.6° inclination. The mission duration is 7 years. The results show that to assure system availability a system-level mitigation should verify functionality of the Sakura-II at least once per week.

Table 9: Radiation Models Used for GEO

Trapped Radiation Models	AP-8, AE-8
Trapped proton anisotropy	Badhwar & Konradi 1990 MAX
Solar particle fluxes	CREME-96
Solar particle fluences	ESP-PSYCHIC
Galactic Cosmic Rays Spectra	ISO-15390
Shielded Flux	MFLUX

Table 10: Single Event Functional Interrupt Rates for LEO Environment

Environment	SEFI/second
Worst Week - Solar Max	2.95E-08
Worst Day - Solar Max	8.20E-08
Worst 5 minutes - Solar Max	2.95E-07
Worst Week - Solar Min	2.96E-08
Worst Day - Solar Min	8.20E-08
Worst 5 minutes - Solar Min	2.95E-07

Table 11: Survivability Percentage for LEO Environment

Survivability per Environment							
days	1	2	3	4	5	6	7
minutes	1440	2880	4320	5760	7200	8640	10080
Worst Week - Solar Max	99.745%	99.491%	99.237%	98.984%	98.732%	98.480%	98.229%
Worst Day - Solar Max	99.294%	98.593%	97.898%	97.207%	96.521%	95.839%	95.163%
Worst 5 minutes - Solar Max	97.480%	95.023%	92.628%	90.294%	88.018%	85.800%	83.637%
Worst Week - Solar Min	99.745%	99.491%	99.237%	98.984%	98.731%	98.480%	98.229%
Worst Day - Solar Min	99.294%	98.593%	97.897%	97.206%	96.520%	95.839%	95.162%
Worst 5 minutes - Solar Min	97.480%	95.023%	92.628%	90.293%	88.018%	85.799%	83.637%

Table 12: Total Mission SEFIs for 7-year Mission Duration in LEO Environment

	Direct Ionization	Proton Induced Ionization	Total
Solar Minimum	1.8208E-04	9.4402E-01	9.4421E-01
Solar Maximum	1.8208E-04	8.1138E-01	8.1156E-01

GEO:

The environment for this study is defined as 42K km with equatorial inclination (0°). The mission duration is 5 years. The results show that to assure system availability a system-level mitigation should verify functionality of the Sakura-II at least once per day.

Table 13: Radiation Models Used for GEO

Trapped Radiation Models	AP-8, AE-8
Trapped proton anisotropy	Watts et al. 1989 VF1-MAX
Solar particle fluxes	CREME-96
Solar particle fluences	ESP-PSYCHIC
Galactic Cosmic Rays Spectra	ISO-15390
Shielded Flux	MFLUX

Table 14: Single Event Functional Interrupt Rates for GEO Environment

Environment	SEFI/second
Worst Week - Solar Max	1.18E-07
Worst Day - Solar Max	3.33E-07
Worst 5 minutes - Solar Max	1.20E-06
Worst Week - Solar Min	1.18E-07
Worst Day - Solar Min	3.33E-07
Worst 5 minutes - Solar Min	1.20E-06

Table 15: Survivability Percentage for GEO Environment

Survivability per Environment							
day	1	2	3	4	5	6	7
minutes	1440	2880	4320	5760	7200	8640	10080
Worst Week - Solar Max	98.989%	97.989%	96.998%	96.018%	95.047%	94.087%	93.136%
Worst Day - Solar Max	97.168%	94.416%	91.743%	89.144%	86.620%	84.167%	81.783%
Worst 5 minutes - Solar Max	90.154%	81.277%	73.274%	66.059%	59.555%	53.691%	48.404%
Worst Week - Solar Min	98.989%	97.989%	96.998%	96.018%	95.047%	94.087%	93.136%
Worst Day - Solar Min	97.168%	94.416%	91.743%	89.144%	86.620%	84.167%	81.783%
Worst 5 minutes - Solar Min	90.154%	81.277%	73.274%	66.059%	59.555%	53.691%	48.404%

Table 16: Total Mission SEFIs for 5-year Mission Duration in GEO Environment

	Direct Ionization	Proton Induced Ionization	Total
Solar Minimum	7.3544E-01	1.7085E-02	7.5252E-01
Solar Maximum	7.3544E-01	1.7085E-02	7.5252E-01

11. References

[1] Roffe, Seth, Scott Stansberry, Jeffrey Grosman, Jeffry Milrod, Manish Sinha, Uzzal Podder, and Stan Crow. *EdgeCortix SAKURA-I Machine-Learning, PCIe Accelerator SEE Proton Test*. No. NASA/TM-20240006221. NASA Electronic Parts and Packaging (NEPP) Program, 2024.

[2] Roffe, Seth S., Scott D. Stansberry, Edward J. Wyrwas, Jeffrey Grosman, Jeffry Milrod, Uzzal Podder, and Stan Crow. EdgeCortix SAKURA-I Machine-Learning, PCIe Accelerator SEE Heavy Ion Test Report. No. NASA/TM-20240015800. National Aeronautics and Space Administration, 2024.

[3] EdgeCortix, “SAKURA-I Edge AI Accelerator,” April 2024.

[4] Jocher, Glenn, Ayush Chaurasia, Alex Stoken, Jirka Borovec, Yonghye Kwon, Kalen Michael, Jiacong Fang et al. "ultralytics/yolov5: v7. 0-yolov5 sota realtime instance segmentation." *Zenodo* (2022).

[5] Lin, Tsung-Yi, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. "Microsoft coco: Common objects in context." In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pp. 740-755. Springer International Publishing, 2014.

[6] Garrett, Tyler, Seth Roffe, and Alan George. "Soft-Error Characterization and Mitigation Strategies for Edge Tensor Processing Units in Space." *IEEE Transactions on Aerospace and Electronic Systems* (2024).

[7] Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. "Imagenet: A large-scale hierarchical image database." In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248-255. IEEE, 2009.