



Single Event Effects - A Comparison of Configuration Upsets and Data Upsets

WP0203 White Paper

Single Event Effects in Ground-Based and Airborne Systems

Single event effects (SEE) include instantaneous upsets, transients, and latch-ups due to particle radiation. Historically, SEEs were of interest only to design teams working on systems destined for high-radiation environments such as space. However, advances in modern semiconductor manufacturing technologies have rendered modern ICs susceptible to radiation effects at ground level due to atmospheric neutrons and other sources of terrestrial background radiation. SEEs in field programmable gate arrays (FPGAs) are particularly important for designers and reliability engineers working on high-availability or high-reliability systems, as FPGAs are increasingly relied upon to perform system-critical functions. This white paper addresses:

- The difference between configuration upsets and data upsets
- Why SRAM FPGAs are more vulnerable to SEE
- A comparison of their respective risks to system reliability

Testing for Neutron SEE

High energy neutrons form the most damaging naturally-occurring terrestrial background radiation. Neutrons arise in the upper layers of the atmosphere, created by the interaction of energetic sub-atomic particles in space (galactic cosmic rays) and gases in the atmosphere. Since neutrons do not possess electrical charge, they do not recombine, and instead persist in the atmosphere. At lower elevations, neutrons are attenuated by atmospheric gases. However, they still represent a significant threat to the reliability of electronics even at sea-level.

To gather significant statistical data on the effects of atmospheric neutron radiation on modern electronics, most semiconductor vendors and system manufacturers use accelerated neutron testing. Several test facilities provide neutron beams, which closely resemble the energy spectrum of the naturally-occurring background neutron radiation. However, these neutron beams are many orders of magnitude greater in **neutron flux**, measured in neutrons per cm² per second. The benchmark standard test facility is Los Alamos Neutron Sciences Center, LANSCE, at Los Alamos National Laboratory in New Mexico. The LANSCE neutron energy spectrum closely matches the background neutron energy spectrum from 1 MeV to 200 MeV at sea-level, but with flux 10⁶ higher than the background flux.

Configuration Upsets

Configuration upsets are most crucial to the operation of FPGAs. High-availability or mission-critical applications require FPGAs to retain their configuration in radiation environments. The consequences of an upset in the FPGA configuration memory can be as severe as the FPGA failing to operate as intended, potentially releasing millions of bits of corrupted data into the system, or even failing to respond in critical control situations.

Multiple generations of Microsemi flash-based FPGAs have been tested in the high-flux neutron environment at LANSCE, and have exhibited a complete absence of radiation induced upsets in the flash cells, which control the configuration of the FPGA, as shown in [Table 1 on page 3](#). This is in direct contrast to SRAM FPGAs, which have demonstrated very high levels of configuration upsets, over multiple generations from multiple vendors.

Table 1: Configuration Upsets Test Results for IGLOO2 and SmartFusion2 Devices


Environment	Configuration SEU
Ground Level (Sea-Level, New York City)	Immune
Aviation (40,000 feet, New York City)	Immune

Table 2 shows typical SRAM FPGA configuration upsets at sea-level in FIT/Mb, which is the standard unit of measurement. One FIT/Mb denotes a failure every billion device hours per Megabit of configuration memory.

Table 2: Typical Popular SRAM FPGA Configuration Upsets at Sea-Level in FIT/Mb (Failures in Time per Billion hours per Megabit of Configuration Memory)

Technology Node	FIT/Mb
250 nm	160
180 nm	180
150 nm	400
130 nm	400
90 nm	100
65 nm	160
45 nm	180
40 nm	100
28 nm	80

Configuration upsets change the function of an FPGA, and cause it to behave unpredictably. Normally, configuration upsets can be cleared by reconfiguring or power cycling the FPGA and have no lasting effect. However, configuration upsets can create illegal conditions within the FPGA. For example, by creating high-current conditions due to contentions as a result of the misconfiguration. The high-current draw may damage the device or the board on which it is integrated. If not corrected, configuration upsets can result in simultaneously enabling pull-ups and pull-downs or they can result in serious bus contention, both of which may physically damage the FPGA.


Figure 1: Configuration Errors in SRAM FPGAs

Mitigation of configuration upsets is very complex and cumbersome and not wholly effective. Mitigation schemes rely on reloading the correct configuration file. They can also rely on tripling each portion of the design and voting out the error, triple module redundancy (TMR), to mitigate the effects. But this reduces the available gates by a factor of four or five and does not deal with static circuits where voting does not protect against the accumulation of upsets, defeating the TMR. The corrupted configuration bit causes device malfunction from the occurrence of the single event until it is corrected, a period of at least several tens or hundreds of milliseconds. During this time, the FPGA behaves erroneously and pours corrupted data into the system.

When a configuration error occurs, it usually results in a large amount of corrupted data in the system as well as possibly requiring the FPGA to be reset, which causes downtime in operations. If the SRAM-based FPGA is used in a mission-critical or life-critical function, this introduces the risk of catastrophic failure.

Data Upsets

A data single event upset (SEU) is a soft error where radiation has changed a data bit to an incorrect value. These can be broadly categorized into single-bit and multi-bit data upsets.

Single-bit Data Upsets in Flip-Flops

An SEU in a flip-flop manifests as a single-bit error. The error will be overwritten when the flip-flop is clocked the next time. In many systems, a single-bit error does not have any consequences and mitigation is not necessary. However, for control applications in mission-critical or life-critical systems, single-bit errors can be mitigated by forward error correction codes, TMR, or fail-safe state machine design techniques. Data upset rates for Microsemi SmartFusion[®]2 system-on-chip (SoC) FPGA and IGLOO[®]2 FPGAs are shown in Table 3. As background neutron radiation varies with latitude and longitude, upset rates are usually quoted for sea-level, New York City (NYC). The SRAM data upset rates are comparable between SRAM FPGAs and flash-based FPGAs.

Table 3: Data Upset Test Results for IGLOO2 and SmartFusion2 Devices

Feature	Test Fluence (Neutrons/cm ²)	Error Rate Ground-Level (Sea-Level, NYC, FIT)	Error Rate Aviation (40,000', NYC, FIT)
Flip-flop	4.35×10^{11}	218.3 FIT / million flip-flops	1.13×10^5 FIT / million flip-flops
LSRAM	1.7×10^{11}	340.6 FIT / million bits	1.75×10^5 FIT / million bits
uSRAM	1.7×10^{11}	175.3 FIT / millions bits	9.04×10^4 FIT / million bits

Single-bit Data Upsets in the SRAM Memory

It is required to make the distinction between memory structures on the FPGA used to store data and configuration memory structures used to store the function of the FPGA. A single-bit data upset in the configuration memory of the FPGA fabric results in a stream of corrupted data propagated into the system. However, when a single-bit error occurs in data memory, it is just a single data bit among many, and can effectively be mitigated by error detection and correction encoding schemes such as shortened Hamming codes which provide single error correction, double error detection (SECDED). If the dwell time of data is long, background scrubbing may be required, so that multiple radiation events do not cause multiple errors to occur in the same data word. The decision on whether or not to deploy scrubbing can be reached by considering the worst-case dwell-time for data in the memory, combined with the upset rate for the memory structure in the intended environment.

Multi-bit Data Upsets in Flip-Flops

Multi-bit data upsets are problematic when they occur, but SEEs are usually localized to a very small area of the chip. By interleaving the memory, logically adjacent bits are physically separated. Physically adjacent bits belong to different logical words, rendering multi-bit upsets correctable.

Figure 2 shows the effect of physically interleaving in an SRAM memory.




Figure 2: Interleaving Eliminates Uncorrectable Multi-Bit Upsets

Single-bit data upsets in memory structures or in flip-flops can easily be mitigated by error correction and detection encoding (EDAC), or by redundancy with parity checking. In many cases, single-bit data upsets in the data stream are insignificant. This is a lower threat to the system reliability with only single bits of corrupted data, in contrast to streams of corrupted data which can arise from Configuration Upsets.

Microsemi's largest SmartFusion2 SoC FPGA is the M2S150 device with 146,124 flip-flops. If operated at commercial aviation altitudes, using the flip-flop upset rate of 1.13×10^5 FIT/million flip-flops, a single M2S150 device experiences 16,512 bit flips per billion hours, or 1-bit flip every 7 years.

Note: Single-bit errors are generally inconsequential.

Comparison Examples

To compare an SRAM FPGA with configuration upsets and a flash-based FPGA immune to configuration upsets an example can be drawn using published failure data of a popular manufacturer's 65 nm SRAM FPGA. With ~50 Mb of configuration memory, the FPGA experiences configuration failures at a rate over 8,000 FIT (8,000 failures per billion hours) at sea-level NYC, or one error in every 14 years. If there are 100 parts in a system, then there will be one error in the system for every 52 days on average. A modern commercial airliner with 20 or more line-replaceable units (LRU) composed of multiple boards with multiple FPGAs can easily contain a population of 100 FPGAs. The commercial airliner flies at 40,000 ft or higher altitudes, and using a general long-haul flight that crosses the Arctic Circle, the FPGA experiences 515x more configuration upsets than at sea-level NYC, as described in JEDEC industry document JESD89A. Therefore the airliner experiences one FPGA functional failure for every 2.5 hours of flight on average. Assuming a fleet of 200 airplanes with this equipment, the failure rate exceeds one per minute on average. If configuration SEU mitigation is deployed, it might take the system at least 10 msec to detect and correct each configuration SEU. Assuming that the FPGA has multiple outputs transmitting data at an aggregate rate of 2 Gb/sec, then each configuration failure causes the corruption of at least 20 million data bits. This occurs at a rate of once every 2.5 hours on each aircraft, or once every minute across the fleet. Worse than this, a configuration upset in an SRAM FPGA can cause not only corruption of large amounts of data, it can also cause the malfunction of a flight-critical system. In comparison, if the commercial airliner design uses flash-based FPGAs there will not be configuration upsets at all.

In contrast to configuration upsets, radiation upsets in flip-flops or data memory cause single-bit errors. Flash and SRAM FPGAs experience upsets at approximately the same rates, roughly 100,000 FIT per million flip-flops or memory bits at 40,000 ft altitude. For the airborne system example, each FPGA has 146,124 flip-flops. Each FPGA experiences 16,405 bit flips per billion hours of operation. On the aircraft using 100 FPGAs, this results in one single-bit flip every 600 hours of operation on average. Across the fleet of 200 aircraft, this will be a single-bit flip once every 3 hours.

It is clear that the consequences of configuration upsets are much more severe than the consequences of data upsets in flip-flops and data memory.

Table 4: Mean Time Between Failures for Configuration Upsets and Flip-Flop Data Upsets in a Commercial Aviation Application

Upsets	Mean Time	Description
Configuration Upset - Aviation		
Configuration Upset Rate - Sea-Level, per Mb	160	Per Billion hours, NYC sea-level, per Mbit of config memory
Amount of Config Memory (Mb)	50	65 nm 80KLE SRAM FPGA
Configuration Upset Rate - Sea Level, per FPGA	8,000	Per Billion hours, NYC sea-level, 65 nm 80KLE SRAM FPGA
Multiplier Sea-Level to 40K'	515	From seutest.com, JESD-89A
Configuration Upset Rate - 40K'	4,120,000	Per Billion hours, NYC 40K', for FPGA using 50 Mbits of configuration memory
FPGAs per airplane	100	Ten FPGAs per system, ten systems per airplane
Configuration upsets per airplane	412,000,000	Per billion hours, NYC 40K'
Mean Time Between Failures per airplane	2.43	Hours between configuration upsets
Airplanes in fleet	200	–
Configuration upsets per fleet	82,400,000,000	Per Billion hours, NYC 40K'
Configuration upsets per fleet per day	1,978	Each event can result in system malfunction
Mean Time Between Failures for fleet	0.73	Minutes between configuration upsets
Data Upset - Aviation		
Flip-flop Upset Rate	112,270	Per Million flip-flops per Billion hours, 40,000'
Number of flip-flops per FPGA	146,124	In M2S150
Flip-flop Upset Rate	16,405	Per M2S150 per Billion hours, 40,000'
FPGAs per plane	100	Ten FPGAs per system, ten systems per airplane
Bit-flips per plane	1,640,534	Per billion hours, NYC 40K'
Mean Time Between Failures per airplane	610	Hours between bit-flips
Planes in fleet	200	–
Bit-flips per fleet	328,106,830	Per Billion hours, NYC 40K'
Bit-flips per fleet per day	7.9	Single-bit errors per day
Mean Time Between Failures for fleet	3.0	Hours between bit-flips

A second example uses a wide-area communications network to illustrate the difference between configuration upsets and data upsets in a different context. In the network, there are 100 routers, each with 25 line cards, and each line card uses four FPGAs. The entire system comprises a total population of 10,000 FPGAs.

If the system uses a popular SRAM FPGA of 65 nm technology node with 80,000 flip-flops and 50 Mb of configuration bits, it results in 8,000 FIT per FPGA at NYC sea-level. Configuration SEU mitigation is deployed and the system detects and corrects each configuration SEU in 10 msec. Assuming that the FPGA has multiple outputs transmitting data at an aggregate rate of 4 Gb/sec, then each configuration failure causes the corruption of 40 million data bits. It happens at a rate of 8,000 times per billion hours per FPGA, or once every 5,210 days per FPGA. With a system population of 10,000 FPGAs, it means an error

causing the corruption of 40 million bits of data occurs once in every 12.5 hours. Worse than this, each configuration upset in an SRAM FPGA has the potential to cause a functional failure in the system.

In contrast, Microsemi flash-based FPGAs are immune to configuration upsets. The data upset rate for Microsemi SmartFusion2 and IGLOO2 FPGAs is determined through accelerated neutron testing to be around 218 FIT per million flip-flops at sea -level. In the communication system example, above, if the 146,124 flip-flop Microsemi M2S150 SoC FPGA is used instead of the 65 nm SRAM FPGA, each FPGA experiences a single-bit error on average 32 times per billion hours. This is once every 31 million hours, or once every 3,600 years. With 10,000 FPGAs in the network, roughly three single-bit data errors occur in the communications network each year.

Neutrons cause data corruption in the flash-based FPGA at the rate of three bits per year, whereas in the SRAM-based FPGA neutrons cause data corruption at the rate of 40 million bits every 12.5 hours.

Table 5: Mean Time Between Failures for Configuration Upsets and Flip-Flop Data Upsets in a Wide-area Network Application

Upsets	Mean Time	Description
Configuration Upset - Network		
Configuration Upset Rate - Sea-Level, per Mb	160	Per Billion hours, NYC sea-level, per Mbit of configuration memory
Amount of Configuration Memory (Mb)	50	65nm 80KLE SRAM FPGA
Configuration Upset Rate	8,000	Per Billion hours, NYC sea-level, 65 nm 80KLE SRAM FPGA
FPGAs per network	10,000	4 FPGAs per board, 25 boards per router, 100 routers in network
Configuration upsets per network	80,000,000	Per Billion hours, NYC sea-level
Configuration upsets per network per day	1.9	Each event can result in system malfunction
Mean Time Between Failures for network	12.5	Hours between configuration upsets
Data Upset - Network		
Flip-flop Upset Rate	218	Per Million Flops per Billion hours, sea-level
Number of flip-flops per FPGA	146,124	In M2S150
Flip-flop Upset Rate per FPGA	32	Per M2S150 per Billion hours, sea-level
FPGAs network	10,000	4 FPGAs per board, 25 boards per router, 100 routers in network
Bit-flips per network	318,550	Per billion hours, NYC sea-level
Bit-flips per network per year	2.8	Per year, NYC sea-level
Mean Time Between Failures for network	131	Days between bit flips

Summary

Upsets in configuration memory in SRAM FPGAs cause corruption of data many orders of magnitude worse than the corruption of data in flip-flops and data memory structures. Thus, they are significantly more dangerous in mission-critical and safety critical applications.

When the FPGA population in an entire system, fleet or network is considered, the risk of system outages or catastrophic failures in the field due to configuration upsets is severe. Users must additionally be aware of the hidden costs of configuration upsets, which include time lost in supporting and analyzing random field failures, reduced system availability, and reduced levels of customer satisfaction. Mitigation techniques to protect SRAM FPGAs against configuration errors cause additional design complexity and are ineffective at preventing errors - they can only correct the errors after they are detected. Flash-based FPGAs are immune to configuration upsets, and therefore present a much more effective solution for logic integration in high-reliability or safety-critical systems.

References

Microsemi Reliability Webpage

TR0020: SmartFusion2 and IGLOO2 Neutron Single Event Effects (SEE) Test Report

JEDEC Standard JESD89A

List of Changes

The following table shows important changes made in this document for each revision.

Date	Changes	Page
Revision 1 (November 2015)	Initial release.	NA



Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

E-mail: sales.support@microsemi.com

© 2015 Microsemi Corporation. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi Corporation (Nasdaq: MSCC) offers a comprehensive portfolio of semiconductor and system solutions for communications, defense & security, aerospace and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, Calif., and has approximately 3,600 employees globally. Learn more at www.microsemi.com.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.