Microsemi®

FPGA Reliability and the Sunspot Cycle

September 2011

Table of Contents

Abstract/Executive Summary

Engineers design for system-level reliability when creating complex electronic systems. System reliability is dependent on a number of factors, including the reliability of the individual semiconductor devices selected for each design. Radiation effects impact device reliability in terrestrial applications, and by choosing FPGA components judiciously, designers can significantly enhance the operational reliability and stability of any system.

Natural background radiation has a number of sources, including cosmic rays originating outside our solar system, charged particles streaming out from our sun in the solar wind, and terrestrial radioactive decay of materials relatively abundant in our earthbound environment. As solar wind particles interact with Earth's atmospheric gases, they produce showers of neutrons that reach the Earth's surface. Neutron counts at sea level vary inversely with the 11-year sunspot cycle. In 2007 the sun was approaching a new sunspot minimum, with expectations that sunspot counts would begin to increase in 2008. This has not yet happened. As of August 2009, the sunspot count has been at or near zero for the past 12 months, and sea-level neutron flux measurements are now well past the previous maximum readings measured in 1965, with predictions of even higher peak readings.

Also of concern are alpha particles emitted by isotopes present even in the low-alpha plastic molding compounds used in semiconductor packaging. Because these materials are in close proximity to the semiconductor die, package-originated alpha particle effects are impossible to avoid.

With the ambient neutron flux rising well past previous maximum counts over the last 45 years, and with the unavoidable exposure to alpha particles from within semiconductor packaging itself, it is critical for engineers to have an understanding of the real-world effects that these particles have on common FPGA device types. This article examines these effects for SRAM FPGAs, explores the difference between soft and firm errors, looks at FPGA firm-error immunity, and details third-party alpha and neutron testing results for a wide range of FPGA architectures.

As reliability requirements for ground-based systems increase, and as the terrestrial operating environment becomes more challenging, these effects are much more relevant to system designers. By carefully choosing FPGA devices, designers can ensure they deliver highly reliable systems that can operate even in the most challenging extended availability ground-based applications.

Introduction

Engineers design for system-level reliability when creating complex electronic systems. System reliability is dependent on a number of factors, including the reliability of the individual semiconductor devices selected for each design. If a design includes field programmable gate arrays (FPGAs), there are environmental influences on reliability that directly impact the reliability of any system they populate, and choices in components can have a significant impact on system-level reliability. One aspect of the operating environment, often overlooked, is the terrestrial radiation environment. Radiation effects impact device reliability in terrestrial applications, and by selecting FPGA components judiciously, designers can significantly enhance the operational reliability and stability of any system.

Cosmic Rays and Neutrons and Alpha Particles, Oh My!

Natural background radiation has a number of sources, including cosmic rays from sources outside our solar system, charged particles and neutrons streaming out from our sun in the solar wind, and terrestrial radioactive decay of materials relatively abundant in our earthbound environment.

Alpha particles are helium nuclei—two protons and two neutrons—emitted by the radioactive decay of radioisotopes of elements such as radium, uranium and thorium. Cosmic rays are high-energy protons, electrons and alpha particles that originate in distant exotic astronomical phenomena such as supernovas, black holes and neutron stars within our galaxy, as well as in yet more distant quasars and radio galaxies. Solar wind is the outflow from the sun's corona of plasma, mostly composed of electrons and light ions of hydrogen and helium, plus small amounts of heavy ions. The lighter charged components of the solar wind are deflected by the Earth's magnetic field. Only the heavier components of solar wind and cosmic rays can break through to the Earth's atmosphere.

Cosmic rays and heavy ions in the solar wind react with gases in the upper atmosphere to produce high-energy neutrons. A significant number of the resulting neutrons penetrate the atmosphere and reach ground level. These neutrons interact with semiconductors to produce random failures in everyday ground-level applications.

The Earth's magnetic field traps and deflects incoming particles from the solar wind and other more distant sources. Particles that make it through the magnetic field are further attenuated by the atmosphere. At sea level this atmospheric blanket is thickest and the neutron environment is the most benign. Neutron flux densities increase with altitude to an atmospheric maximum at 60,000 feet (about 20 kilometers) above sea level.

Neutron flux densities also increase with latitude. Since magnetic field lines are closer together at the poles, more cosmic particles are able to penetrate closer to the Earth's surface than at the equator. The maximum neutron flux readings on the Earth's surface are seen at the poles.

Neutron flux density is still significant at sea level on the equator, and since neutrons can penetrate many feet of concrete, shielding integrated circuits (ICs) from neutrons is not practical. Any increase in neutron flux densities in the operating environment results in real reliability concerns for system designers using FPGAs.

Alpha Particles – My Plastic Package is Doing What?

Alpha particles can be easily blocked— a sheet of paper or a few centimeters of air provide sufficient shielding to protect against alpha emissions— so the issue of concern to designers is alpha emission from within the package or die. Radioactive isotopes of Uranium (U^{238}) and Thorium (Th^{232}) are present in plastic package molding compounds and assembly materials, and emit alpha particles as they decay. The radioactive isotope Boron-10 (B^{10}) may also be present in the die passivation layer. When Boron-10 interacts with thermal neutrons, it decays into a lithium ion, an alpha particle and a gamma ray. Semiconductor manufacturers are moving to low-alpha molding compounds, which contain smaller quantities of radioactive isotopes. While these low-alpha materials reduce the probability of an alpha event impacting the device, they do not eliminate these events.

The Neutron Environment and the Sunspot Cycle

Scientists around the world maintain a network of neutron monitoring stations. One such station is located at the University of Oulu in Finland, and historical Oulu neutron monitor data is available from 1964 through the present. A graph of the Oulu data is markedly cyclical (Figure 1).

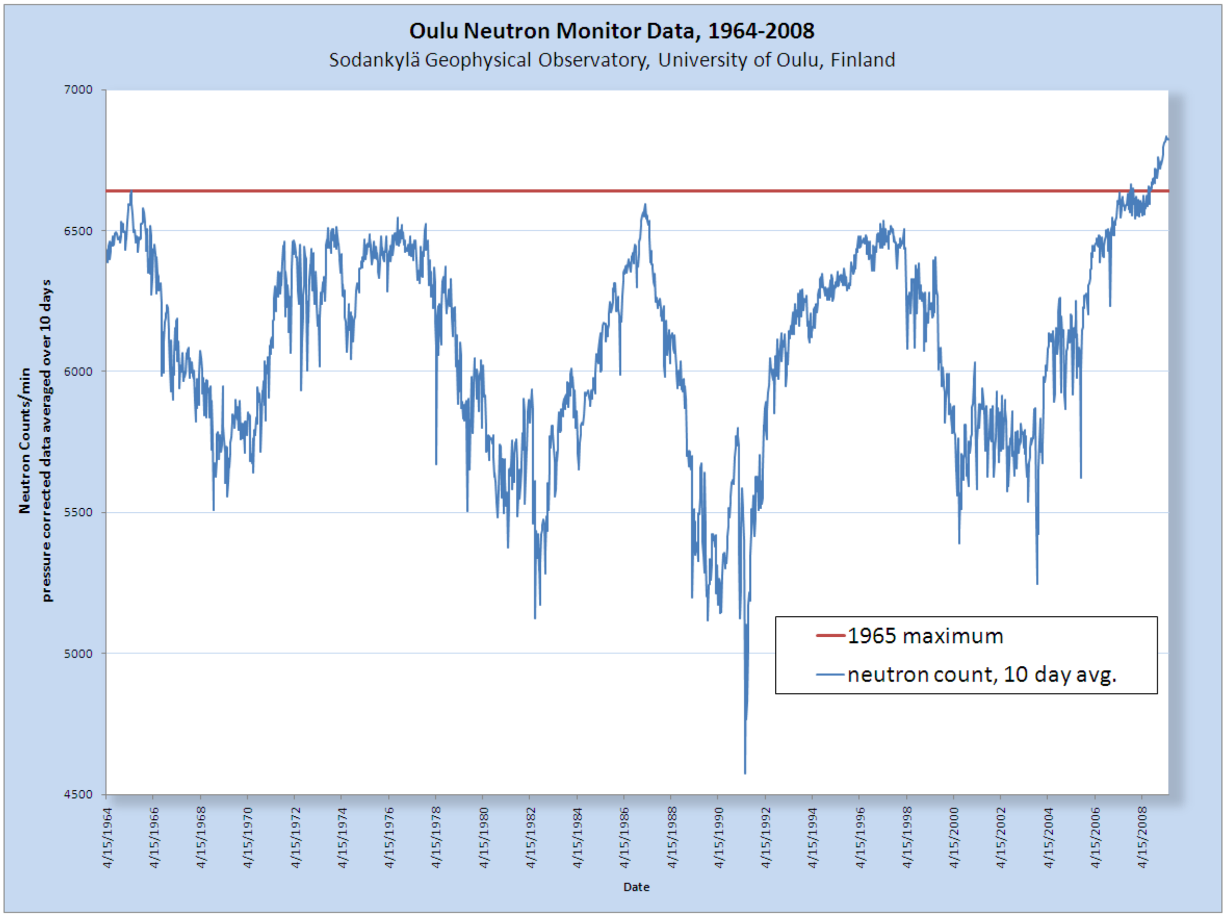

Figure 1: **Oulu Neutron Monitor Data, 1964-2008**

Since neutron flux varies with both altitude and latitude, measurements at the various neutron monitor stations around the world vary as well, but they all show the same periodic variation. When the mean for each site's dataset for 2008 data is subtracted, measurements from all the stations correlate (see Figure 2). The Oulu data is showing a worldwide variation in the Earth's neutron flux environment.

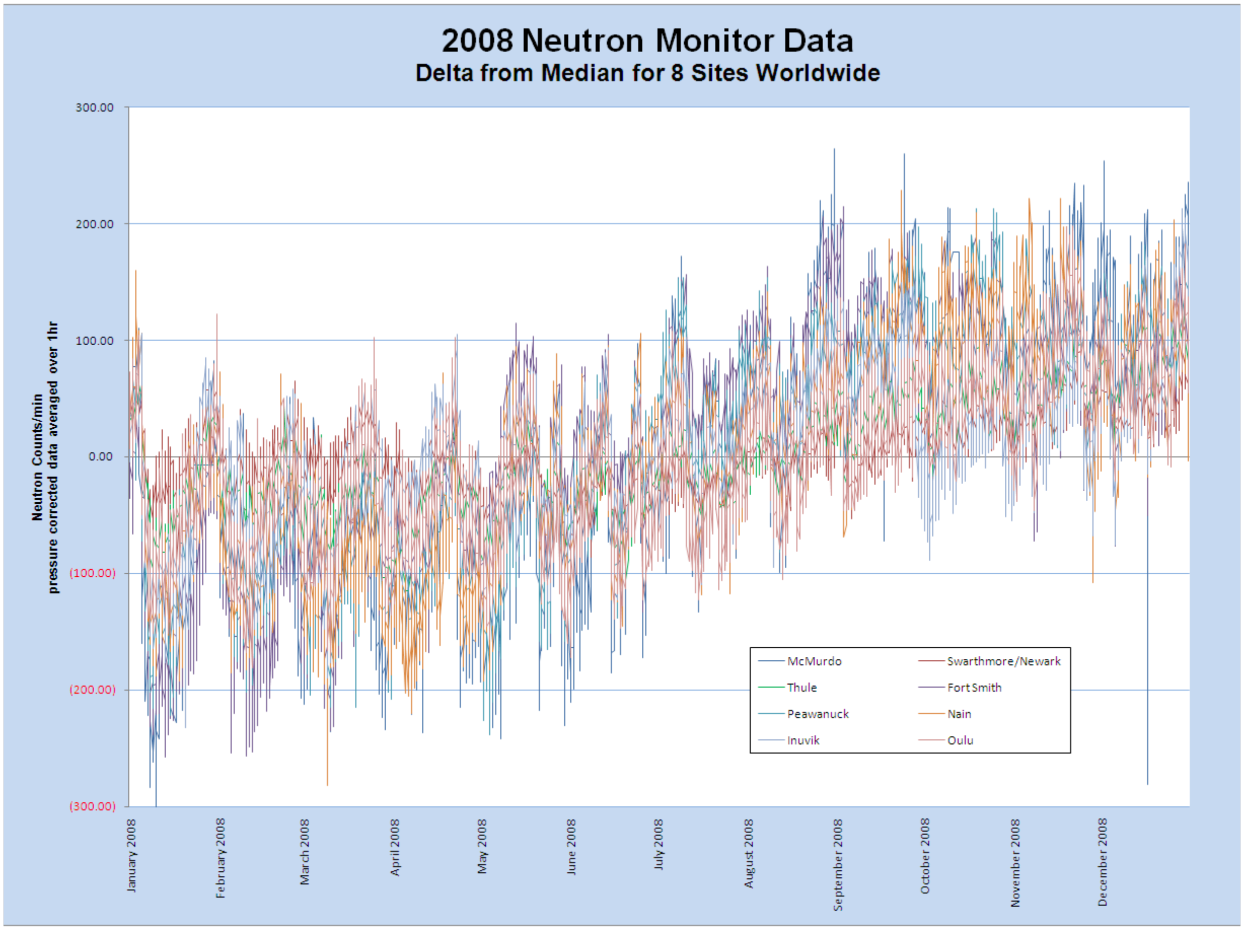

Figure 2: **2008 Neutron Data from Eight Sites in the World Neutron Monitor Network**

So, looking at the long term dataset, what explains the cyclical neutron flux variations? A look at variations in activity provides one correlation. Sunspot count data is available for an extensive period (back to the 17th century), with solid daily sunspot counts available for the last 200 years. These sunspot counts show a variation on a cycle of roughly 11 years. By overlaying sunspot count data since 1964 on the Oulu neutron measurements, we see that the neutron count measured at Oulu varies inversely with the sunspot count (see Figure 3). That is, when the sunspot count is high, the neutron count is low, and when the sunspot count is low, the neutron count is high.

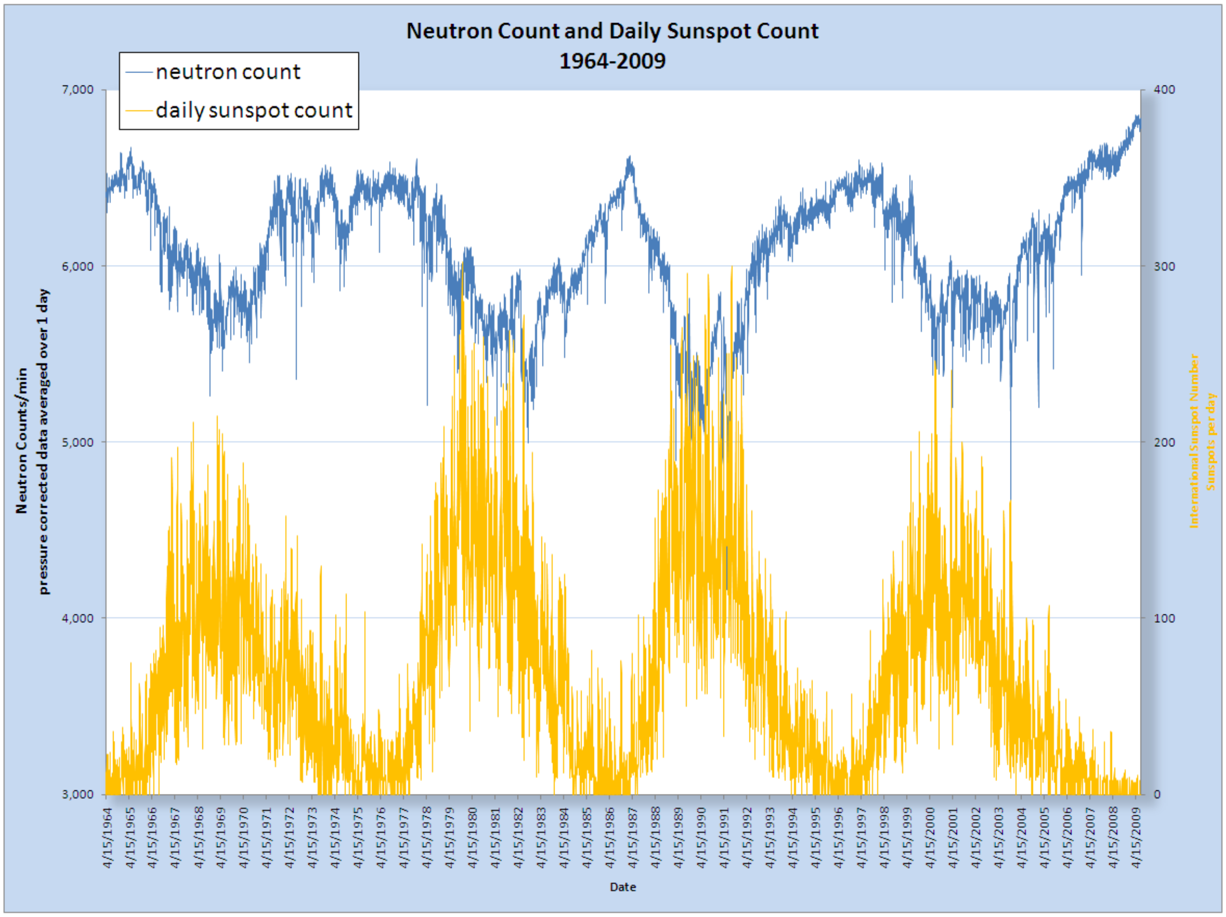

Figure 3: **Oulu Neutron Data and Sunspot Count, 1964-2009**

The historical maximum in the Oulu neutron data occurred in 1965. Since then the periodic neutron count maximum readings corresponding to each sunspot minimum have approached but never surpassed that 1965 maximum—until 2008. Since the last sunspot activity peak in 2000-2002, the count of sunspots has declined, and the measured neutron count has been increasing. By 2008 the sun was entering an unusually quiet period for sunspots, and by the beginning of 2008 the neutron count readings were consistently exceeding the previous 1965 maximum. By October 2008 the daily neutron count readings rarely fell below the 1965 maximum (see Figure 4).

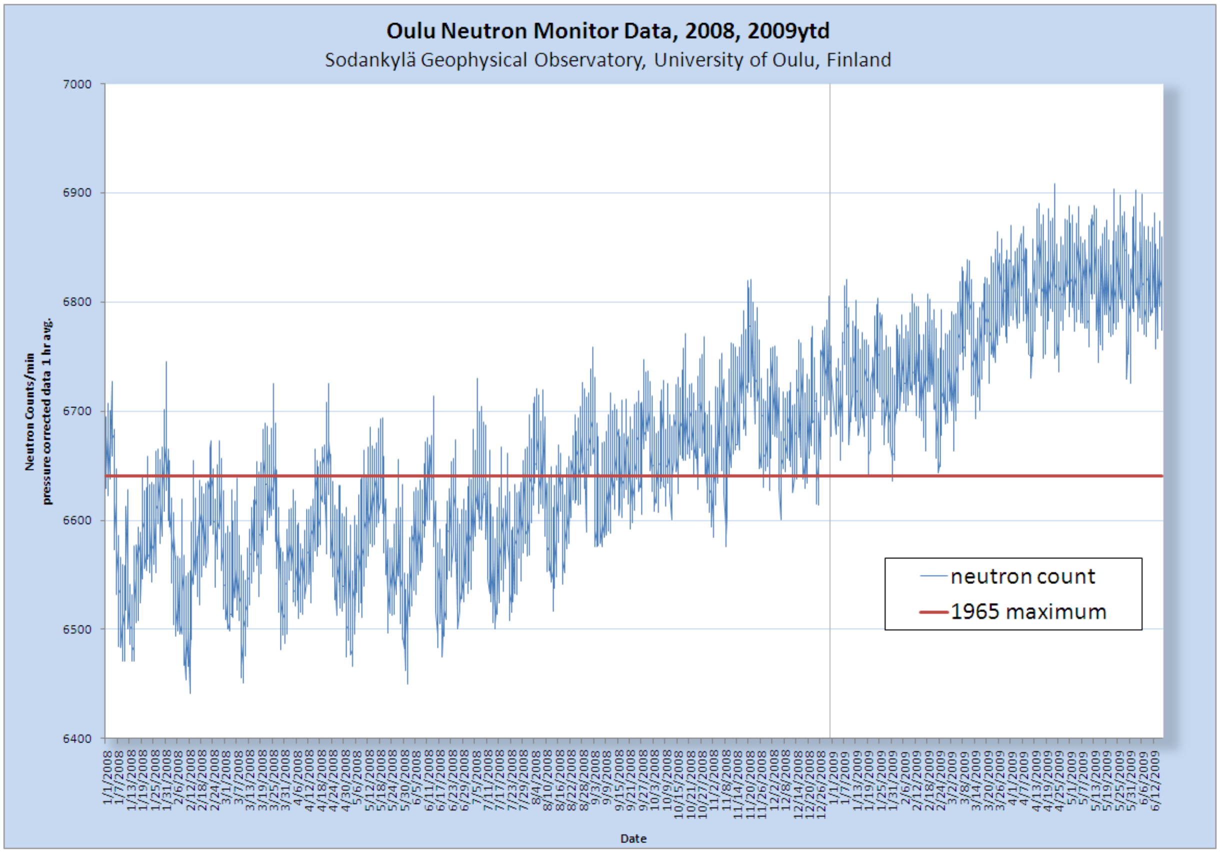

Figure 4: **Oulu Neutron Data 2008 and 2009 Year-to-Date**

An examination of the sunspot/neutron count cycles (Figure 3 on page 7) shows that the neutron count generally does not begin to fall off until after the sunspot count begins to pick up. With sunspot counts remaining effectively at zero from early 2008 through August 2009, the eventual timing of this turnaround is currently indeterminate. By applying statistical analysis techniques to previous solar cycles to build an estimate of when the sunspot count is likely to pick up again, some analysts have arrived at a maximum neutron flux reading for this cycle of well in excess of 7,000 neutron counts per minute. This level would be a significant new maximum, surpassing the previous 1965 peak readings by more than 20% compared to the 1964-2008 mean.

What Chaos Has This Neutron Wrought?

What does this enhanced neutron flux environment mean for system reliability? How are semiconductor devices, and specifically programmable FPGA devices, affected by these subatomic particle events?

When an incoming charged particle (an alpha particle or heavy ion, for example) impacts a semiconductor, the results are significant. Charged particles leave a trail of ionization through the substrate, causing a momentary current pulse in nearby transistors that is proportional to its incoming energy state. Data in memory cells and flip-flops can change (see Figure 5).

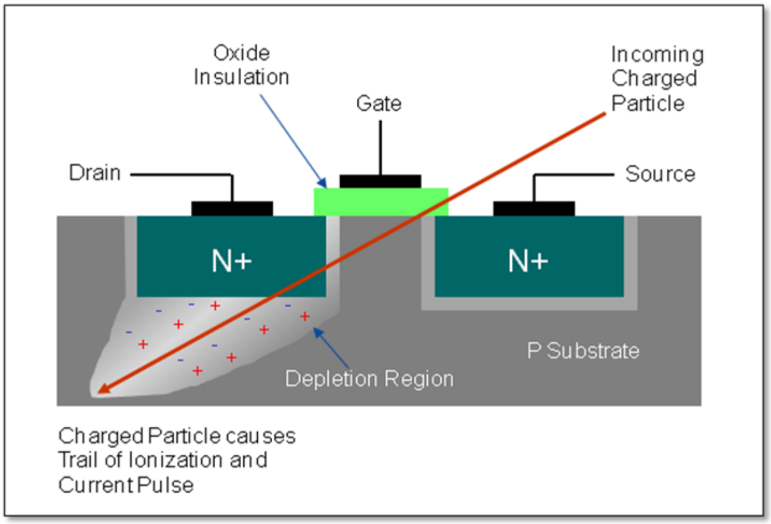
Oxide Insulation
Incoming Charged Particle
Gate
Drain
Source
N+
N+
P Substrate
Depletion Region
Charged Particle causes Trail of Ionization and Current Pulse

Figure 5: **Effect of a Charged Particle on a Semiconductor**

When a neutron impacts a device, the effects are also significant. An incoming neutron may collide with a silicon atom in the substrate. That impact can eject a spray of heavy ions that induce a current pulse in CMOS ICs (see Figure 6). This current pulse, which is proportional to the energy state of the incoming neutron, can cause data in memory cells and flip-flops to change.

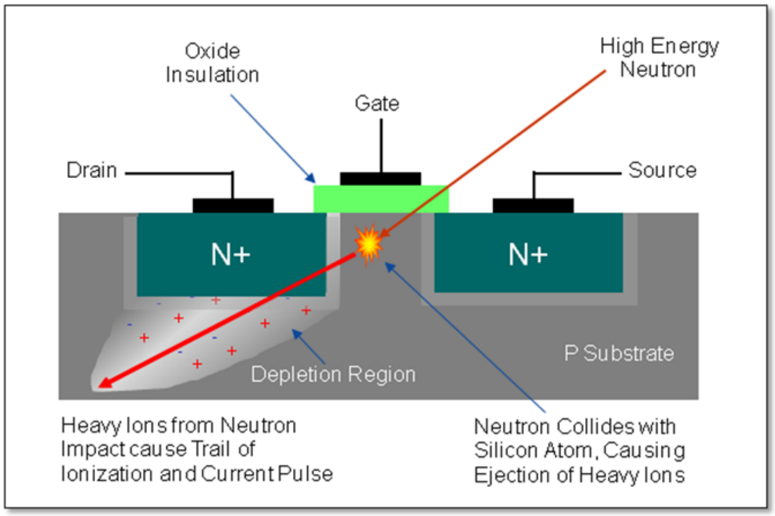
Oxide Insulation
High Energy Neutron
Gate
Drain
Source
N+
N+
P Substrate
Depletion Region
Heavy Ions from Neutron Impact cause Trail of Ionization and Current Pulse
Neutron Collides with Silicon Atom, Causing Ejection of Heavy Ions

Figure 6: **Effect of a Neutron on a Semiconductor**

FPGAs are semiconductor integrated circuit devices that can be reconfigured by the end customer. FPGA manufacturers have taken several approaches in creating the underlying architecture used in FPGA devices, including architectures based on antifuse technology, which is programmable only once; static random-access memory (SRAM) memory, which is reprogrammable multiple times but loses configuration when power is removed; and flash memory cells, which are also reprogrammable and retain their configuration when power is removed from the FPGA.

Both alpha and neutron-induced errors can destroy the integrity of data stored in SRAM FPGA cells. If a memory or flip-flop data change does not change the fit, form or function of the device, it is a “soft error.” Mitigation techniques for soft errors are relatively straightforward and include error detection and correction (EDAC) and error correcting code (ECC) techniques. These techniques consume FPGA circuitry and require a larger device—to perform the same functionality—than would be needed without these errors.

If the cell that changes due to the neutron or charged particle event is an SRAM configuration memory cell that controls logic functions or a routing matrix, the data error may change circuit functionality, internal cell connectivity or signal routing between cells within the SRAM FPGA. This is a “firm error,” and it persists until detected and cleared by rebooting or power cycling the volatile SRAM FPGA and reloading the FPGA with the correct configuration programming stream (see Figure 7).

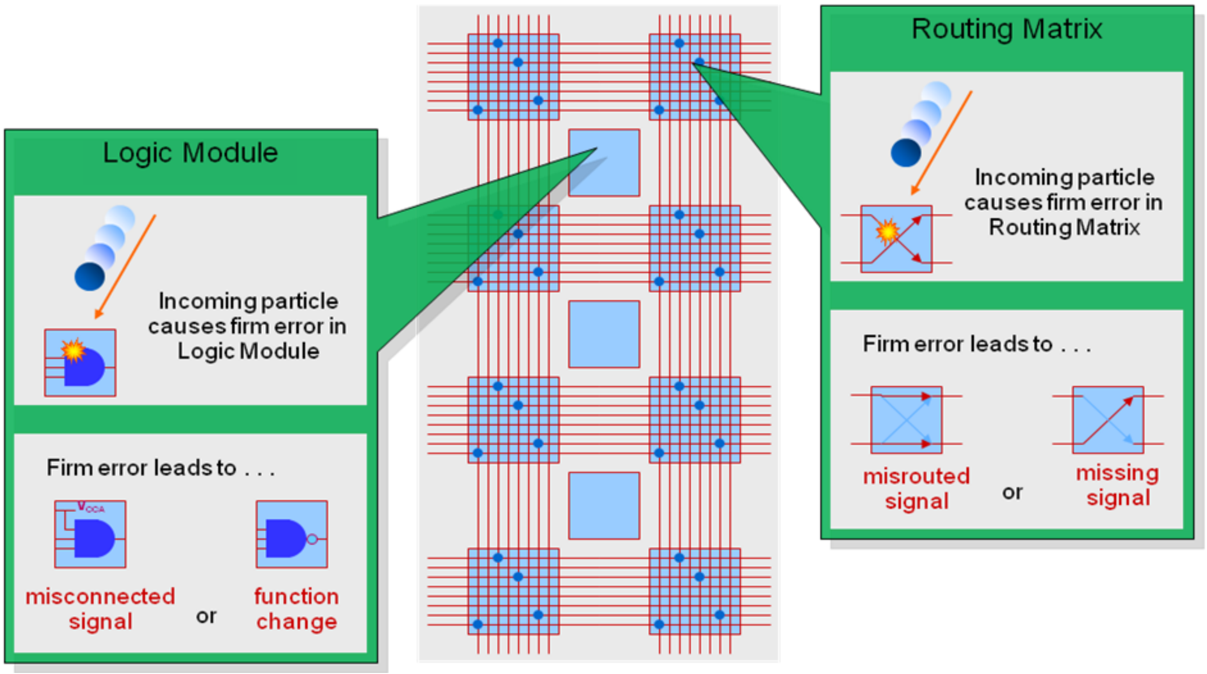
Logic Module
Incoming particle causes firm error in Logic Module
Firm error leads to . . .
VCCA
misconnected signal
or
function change
Routing Matrix
Incoming particle causes firm error in Routing Matrix
Firm error leads to . . .
misrouted signal
or
missing signal

Figure 7: **Firm-Error Effect on an SRAM FPGA**

Mitigation of firm errors is virtually impossible in SRAM FPGAs. System level triple module redundancy (TMR) and voting schemes, where three parallel FPGAs implement identical logic with one additional device monitoring the results and adjudicating any mismatches, are possible solutions to the susceptibility of SRAM FPGAs to these errors, but voting systems still require one firm-error immune FPGA to act as the voting judge.

Mitigation schemes based on cyclic redundancy check (CRC), in which a known good CRC checksum is loaded when the device is programmed and is also continuously recalculated during operation for the FPGA’s current state, are used by some SRAM FPGA vendors, but these are detection-only strategies. These strategies suffer from problems such as CRC-lag, where the CRC value register reflects a previous SRAM FPGA state during the CRC recalculation cycle. Firm-error susceptibility in the CRC register is another problem, where a soft error in the CRC register cells would yield a false-positive CRC mismatch flag. Depending on where in the FPGA a firm error occurs, it can sometimes be cleared with a system power cycle, or might require a full FPGA reset. In any case, these mitigation strategies only detect firm errors—they do nothing to prevent the firm error in the first place.

In the past few years SRAM FPGA vendors have developed hybrid SRAM plus NVM devices, offering persistent data retention with SRAM FPGA speeds. While these are an improvement over pure SRAM FPGAs in that they do not completely lose all configuration data when they lose power, they do not gain any advantages in reliability. These hybrid devices exhibit the same firm-error susceptibility and failure rates as pure SRAM FPGA devices.

Geometry Matters

As semiconductor manufacturers strive to meet the performance and cost improvements predicted by Moore's Law,[1] smaller process geometries have yielded improvements in speed, power consumption and cost in ICs. As high volume IC manufacturers lead the way to successively smaller process nodes, the rest of the semiconductor industry trails along behind, adopting smaller processes after manufacturing has been proven and most of the production bugs worked out. This race to smaller process nodes has brought successive challenges, including a net increase in vulnerability to single-event upset (SEU) errors.

Firm-error rate depends on a combination of particle density and various geometry-related factors. The critical charge required to cause a soft or firm error in an SRAM FPGA is proportional to device capacitance and voltage, and both of these decrease as SRAM FPGA devices migrate to smaller process geometries. Smaller SRAM cells are more easily upset by random low-energy particles, and a greater proportion of neutron and alpha events can transfer enough energy to cause an error. Supply voltage scaling also increases firm-error susceptibility, as lower voltages lower the upset threshold and make single-event errors more common. Decreases in gate area cause capacitance to decrease, which increases error susceptibility because the charge required to cause an upset increases exponentially as the required charge decreases. Smaller cell geometries also reduce cell area, effectively reducing the size of the target and the probability of a hit. In any given area, however, there will be more transistors, and the lower energy needed to cause an upset outweighs the slightly lower probability of an individual element being struck. The net effect is that smaller process geometries dramatically increase the probability of single-event upset soft or firm errors in FPGAs based on SRAM technologies.

[1] The observation made in 1965 by Gordon Moore, co-founder of Intel, that the number of transistors per square inch on integrated circuits had doubled every year since the integrated circuit was invented. In subsequent years, the pace slowed down a bit, but data density has doubled approximately every 18 months.

Testing SRAM and Flash FPGAs for Single-Event Upset Susceptibility

iRoC Technologies, an independent organization with expertise in soft error testing and regular testing slots at Los Alamos National Laboratory, conducted both neutron and alpha testing on FPGAs using three different programming technologies in five different architectures from three major FPGA vendors. The FPGAs were tested until a significant number of failures were observed, and based on these results, failures-in-time (FIT) rates were calculated. One FIT represents a single failure in one billion hours of operation; a system that experiences one failure in 13,158 hours has a failure rate of 1E9 / 13,158 = 76,000 FITs. The neutron test results are shown below in Table 1 and the alpha test results are shown in Table 2. Acceptable FIT rates for commercial applications are fewer than 100, while acceptable FIT rates for system-critical applications are fewer than 20.

Table 1: **JEDEC JESD-89 Compliant Neutron Test Results**

		Equivalent Functional Failure FIT Rates per Device			
		Ground-Level Applications		Commercial Aviation	Military Aviation
FPGA	**Technology**	**Sea Level**	**5,000 Feet**	**30,000 Feet**	**60,000 Feet**
Microsemi AX1000 1M gates	0.15 µm antifuse	No failures detected	No failures detected	No failures detected	No failures detected
Microsemi APA1000 1M gates	0.22 µm flash	No failures detected	No failures detected	No failures detected	No failures detected
Microsemi A3PE600 600K gates	0.13 µm flash	No failures detected	No failures detected	No failures detected	No failures detected
Vendor1 SRAM FPGA 3M gates	0.15 µm SRAM	1,150 FITs	3,900 FITs	170,000 FITs	540,000 FITs
Vendor1 SRAM FPGA 1M gates	90 nm SRAM	320 FITs	1,100 FITs	47,000 FITs	150,000 FITs
Vendor2 SRAM FPGA 1M gates	0.13 µm SRAM	460 FITs	1,600 FITs	67,000 FITs	220,000 FITs
Vendor2 SRAM FPGA 1M gates	90 nm SRAM	730 FITs	2,500 FITs	108,000 FITs	346,000 FITs
Vendor2 SRAM FPGA 2M gates	90 nm SRAM	1,600 FITs	5,500 FITs	236,000 FITs	751,000 FITs

Table 2: **JEDEC JESD-89 Compliant Alpha Particle Test Results**

		Equivalent Functional Failure FIT Rates per Device	
FPGA	**Technology**	**Low-Alpha Mold Compound (0.001 α/cm^2-hr)**	**Standard Mold Compound (0.04 α/cm^2-hr)**
Microsemi AX1000 1M gates	0.15 µm antifuse	No failures detected	No failures detected
Microsemi APA1000 1M gates	0.22 µm flash	No failures detected	No failures detected
Vendor1 SRAM FPGA 3M gates	0.15 µm SRAM	140 FITs	5,600 FITs
Vendor1 SRAM FPGA 1M gates	90 nm SRAM	260 FITs	10,400 FITs
Vendor2 SRAM FPGA 1M gates	0.13 µm SRAM	100 FITs	4,000 FITs

This shows that in neutron and alpha particle testing performed by third parties to date, SRAM FPGA devices in a variety of process geometries and cell technologies from a number of leading vendors all exhibit significant failure rates. Reprogrammable FPGA devices based on flash technology, however, experienced zero failures in both alpha and neutron testing. Antifuse FPGA devices also had zero failures, illustrating why these one-time-programmable devices dominate high-reliability spaceflight applications that do not require reprogrammability.

Conclusion: The Sky is Not Falling, but Immunity is a Good Thing

The bottom line is that system designers can no longer afford to ignore radiation effects. Systems in terrestrial internet infrastructure applications are already logging errors that cause system uptime issues that have been identified as having been caused by SEU radiation events. This applies not only at sea level—no terrestrial application system designer can ignore the fact that Denver, Colorado is more than 5,000 feet (1,600 meters) above mean sea level, and cell towers with their associated electronics are installed on mountaintops at significantly higher altitudes than that.

As network bandwidth demands increase, and as the move towards “cloud-computing” extended-uptime, high-reliability resource models gathers momentum, SEU firm-error elimination will become more critical in designing for reliability and achieving system-level FIT rate requirements.

Neutron monitor data from around the world shows that sea-level neutron flux rates are continuing to increase, which only makes matters worse for system designers. Radiation firm errors have now progressed from a nuisance to a significant problem.

Since SRAM FPGAs store configuration information that defines their circuit level functionality in SRAM cells that are very susceptible to firm errors, any upset could potentially result in system-level malfunction and functional failure. Designers using SRAM-based FPGAs in high-reliability applications can implement dedicated circuits to detect and correct configuration errors, adding to system cost and complexity. By implementing some variety of multiple-voting-logic or error-checking schemes such as EDAC or TMR, whether within individual devices or using three or more identical devices at the board level, designers can at least hope that when the inevitable firm errors do occur they will be detected, the system halted, reconfigured, rebooted and brought back online again as quickly as possible—unless critical circuitry implementing that checking is itself corrupted.

Radiation testing has shown that reprogrammable flash-based FPGAs are not subject to these loss-of-configuration firm errors. Device choice has a significant impact on system-level operational reliability, and considerable effort has been expended to build a solid understanding of semiconductor device reliability impact from these ambient environmental sources. By leveraging this information, system designers can make optimal design choices that reduce operational reliability impacts and deliver high-reliability systems to their customers.

If system designers want to avoid the complicated error detection and remediation schemes required when using SRAM FPGAs in system-critical applications, and the extra expense, increased board area, more complicated board routing and extended design times that are the result, they can instead choose firm-error immune FPGA devices such as those offered by Microsemi.

By Mike Brogley
IP and Solutions Product Marketing Manager, Microsemi SoC Products Group (formerly Actel)

Mike Brogley is a 20-year veteran of the semiconductor industry, having contributed in a variety of technical, operational, marketing and management roles at leading ASIC and FPGA manufacturers. Mike is a proud graduate of San José State University's School of Engineering with a Bachelor of Science degree in Aeronautics. He is a member of IEEE, AIAA, ASME, AAAS and SAE and is a life member of AFCEA and USNI.

Microsemi Corporate Headquarters
One Enterprise Drive, Aliso Viejo CA 92656
Within the USA: (800) 713-4113
Outside the USA: (949) 221-7100
Fax: (949) 756-0308 · www.microsemi.com