Understanding Single Event Effects (SEEs) in FPGAs

A Backgrounder

August 2011

Microsemi

Table of Contents

Overview 3

Sources and Effects of Ionizing Radiation 3

Galactic Cosmic Rays 3
Atmospheric Interactions 5
Packaging 6
Silicon Substrate 7

Single Event Effects 7

Effects of GCR on Silicon 7
Single Event Upsets 8
Single Event Functional Interrupt 10
Single Event Transients 11

SEE Mitigation 11

Block Memory 11
Logic 11
Configuration Memory 13

Summary 13

References 14

Overview

With the increasing popularity of programmable logic, FPGAs are finding their way into many applications that were once the territory of ASICs and ASSPs. At the same time, process nodes are shrinking and logic density is increasing, meaning that more of the system can be implemented in a single device. As programmable logic finds its way into avionics, communications and medical applications, designers face demands for increased reliability and safety over many of the traditional markets for FPGAs.

In these high-reliability markets, there has long been concern over the effect of ionizing radiation on memory circuits; however, its impact on programmable logic is not widely understood by the engineering community—nor is it broadly known that not all FPGA technologies share the same risks. With the focus on reliability and safety in these markets, designers must quantify these risks and understand how differing FPGA technologies react in this environment.

Sources and Effects of Ionizing Radiation

Galactic Cosmic Rays

Galactic cosmic rays (GCR), comprised of high-energy particles, overwhelmingly protons, impact the Earth's atmosphere constantly. These particles, originating in space, have sufficient energy to liberate nuclei when they collide with molecules in the Earth's atmosphere. The result of this collision is referred to as an air shower, where a wide range (and high number) of particles are generated. The primary spallation products of concern to terrestrial and avionics designers are neutrons and protons (in addition to remanent cosmic rays).

The flux of cosmic rays impacting the Earth's atmosphere is modulated by both the solar wind and the Earth's magnetic field. As a result, the greatest modulation occurs at the equator and when the solar wind is most active (i.e., when solar flare activity is high). The flux resulting from the air shower is modulated by the density of the atmosphere (expressed as depth). Combining all of these factors results in the particle flux being a function of latitude, longitude, altitude and solar activity, with the greatest flux occurring at high altitudes over the poles during quiet periods of solar activity.

Figure 1 and Figure 2 show the atmospheric flux for neutrons compared to altitude and latitude.

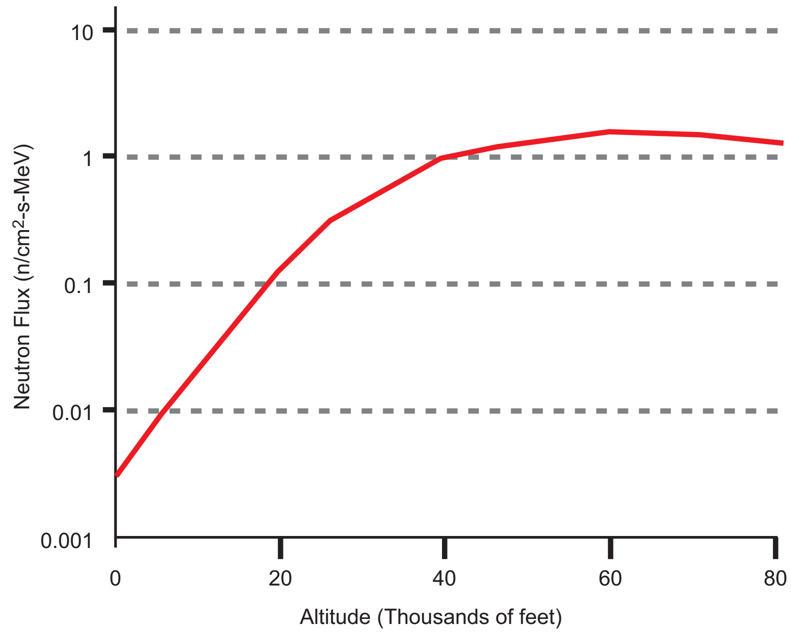

Figure 1: **Atmospheric Flux for Neutrons in the Range of 1–10 MeV over Altitude**

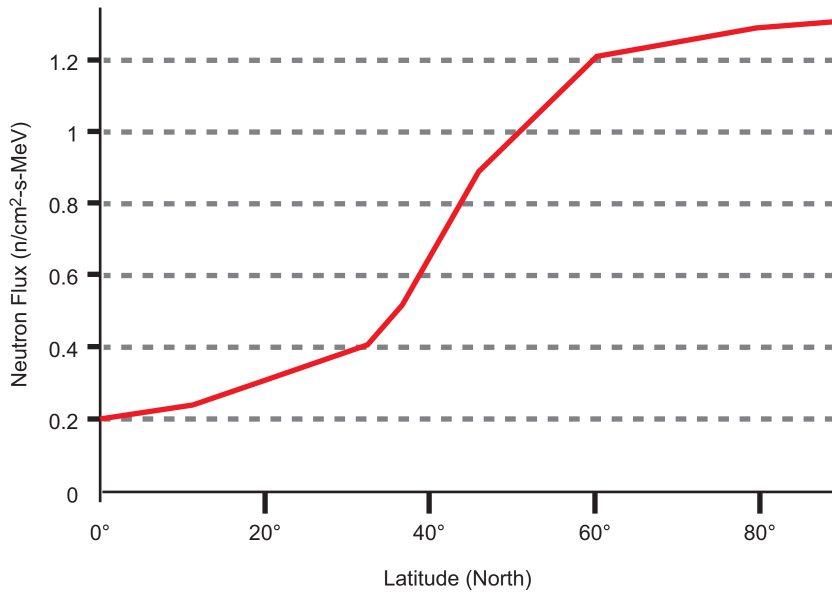

Figure 2: **Atmospheric Flux for Neutrons in the Range of 1–10 MeV over Latitude**

Figure 3 shows the atmospheric neutron energy spectrum at sea level in New York (the standard reference point). The curve illustrates that the vast majority of the atmospheric neutrons hitting Earth are under 10 MeV.

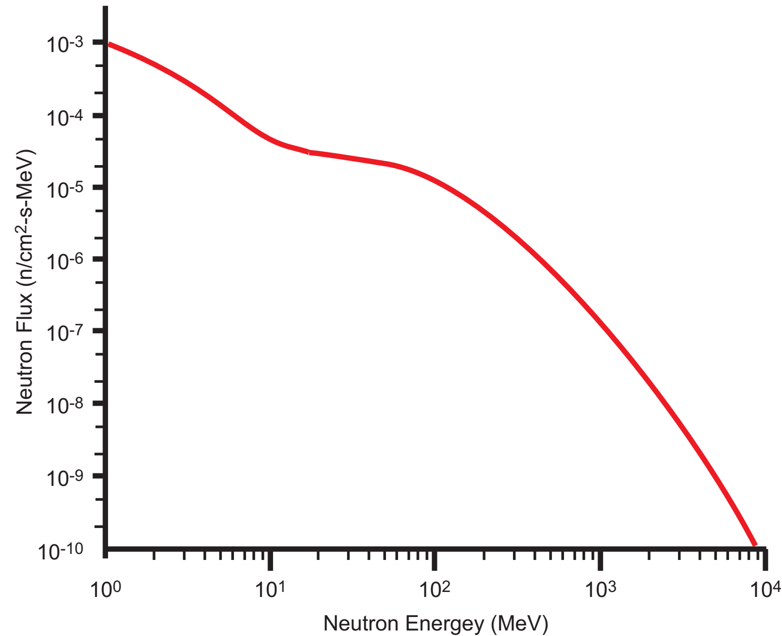

Figure 3: **Differential Neutron Flux at New York CIty (from JESD89A)**

JEDEC Standard 89A, *Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices*, provides a means for estimating neutron flux for a given latitude, longitude, altitude and level of solar activity, relative to a reference point set to the actual flux occurring at sea level in New York City. Based on these equations, an observer in an aircraft flying at 40,000 feet over the poles during a period of moderate solar activity will experience more than 500 times the neutron flux as a terrestrial observer in New York City.

JEDEC maintains a neutron flux calculator at www.seutest.com.

Atmospheric Interactions

Secondary particles are created when GCR collide with the upper atmosphere. These secondary particles also collide with molecules in the atmosphere. For example, when liberated high-energy neutrons collide with a nitrogen atom, which makes up more than 78% of the atmosphere by volume, a proton plus a carbon-14 atom are produced (EQ 1).

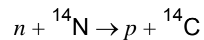

$$n + {}^{14}\mathrm{N} \rightarrow p + {}^{14}\mathrm{C}$$

EQ 1

These additional protons (along with remanent cosmic rays) are about 7 to 32% of the neutron flux at the Earth's surface.

Packaging

Additional radiation sources can be found in packaging itself. Packaging materials used for integrated circuits contain trace amounts of uranium and thorium. The elements $^{238/235}U$ and ^{232}Th naturally emit alpha particles as they decay (Figure 4). Although alpha particles that result from decay have low penetration depth—a few centimeters of air can act as sufficient shielding—the proximity of packaging material to the silicon substrate make them an issue for electronic circuits.

Manufacturers can combat this effect through the use of ultra-low-alpha (ULA) packaging materials. Materials qualify as ULA if they have emissions below 0.002 count/hour-cm^2.

In addition, the eutectic lead solders used for the solder bumps in flip-chip packaging are a source for alpha particles. Even in a solder purified of other radioactive impurities, it is not possible to separate out ^{210}Pb from the rest of the lead isotopes. At first glance, since the primary decay mode of ^{210}Pb is electron transmission (99.99999998%), alpha particle emission appears not to be a concern. However, the decay chain for lead-210 results in polonium-210, a strong alpha emitter (the heat generated by the decay of 1 gram of ^{210}Po is 140 W). Given the relatively short half-life of ^{210}Pb, the rate of alpha emission can increase by an order of magnitude in a matter of months. Because chemical separation of ^{210}Pb is not possible, ensuring that a eutectic solder is ULA requires direct alpha testing over a period of months.

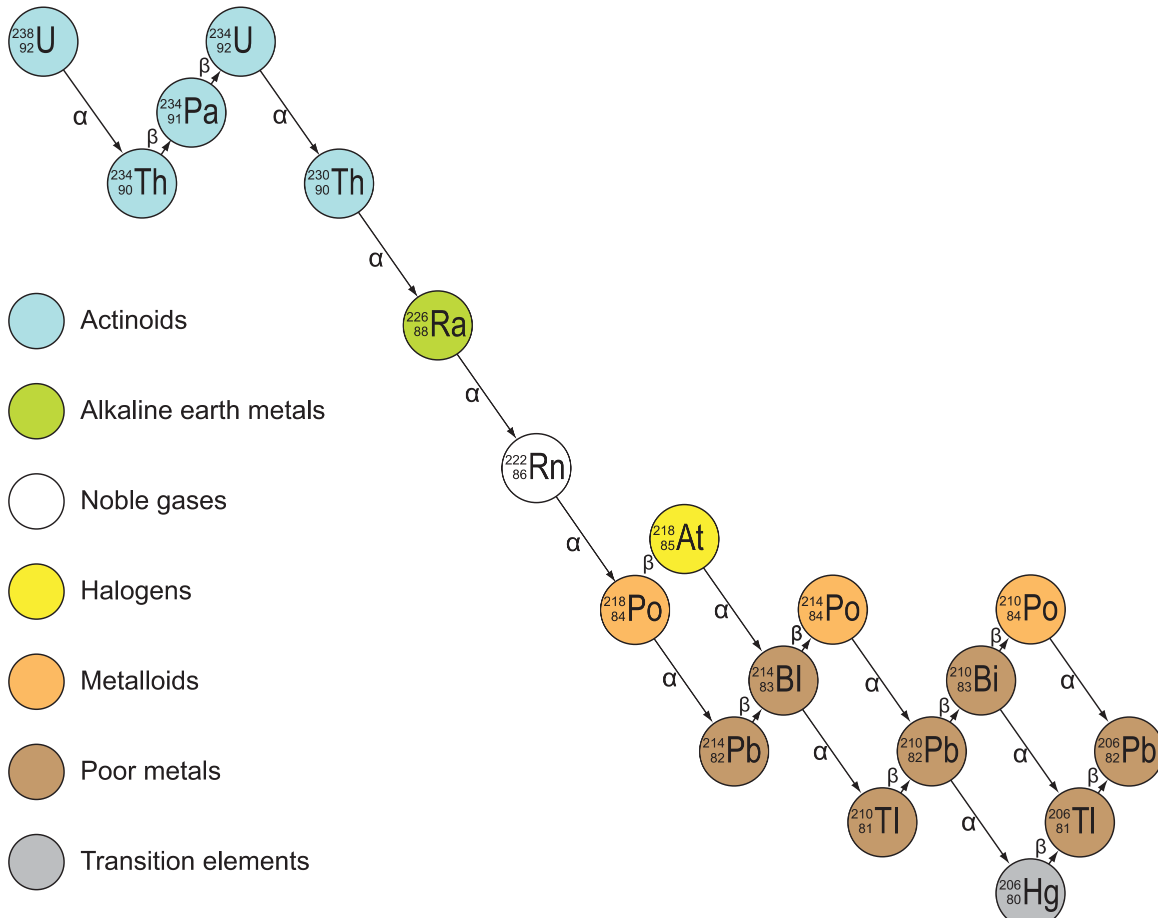

Figure 4: **Uranium Decay Chain**

Silicon Substrate

Another source of ionizing radiation is the element boron, used in polysilicon doping, substrate doping, or borophospho-silicate glass (BPSG) in large amounts. When one of the commonly occurring boron isotopes, ^{10}B, (about 20% of total boron in nature) is struck by low-energy (thermal) neutrons (referred to as neutron capture), both a lithium ion and an alpha particle are created. This spectrum can be significant given both the amount of boron present in substrates plus the number of low-energy neutrons present in the cosmic shower. Because both of these sources are in the device itself, no amount of outside shielding can protect against these particles.

Manufacturers can combat this issue by using depleted boron, a by-product of nuclear reactors. Depleted boron, composed almost entirely of ^{11}B, has a neutron capture cross-section six times less than ^{10}B.

Single Event Effects

Generally, any effect induced by a single radiation event on electronic circuit (as opposed to effects due to collective dosage), whether transient or damaging, is referred to as single event effect (SEE). There are three subclasses of SEEs that are the focus of this paper: single event upsets (SEUs), single event functional interrupt (SEFIs), and single event transients (SETs).

Effects of GCR on Silicon

When charged particles strike the silicon substrate of an IC, they leave an ionization trail (Figure 5).

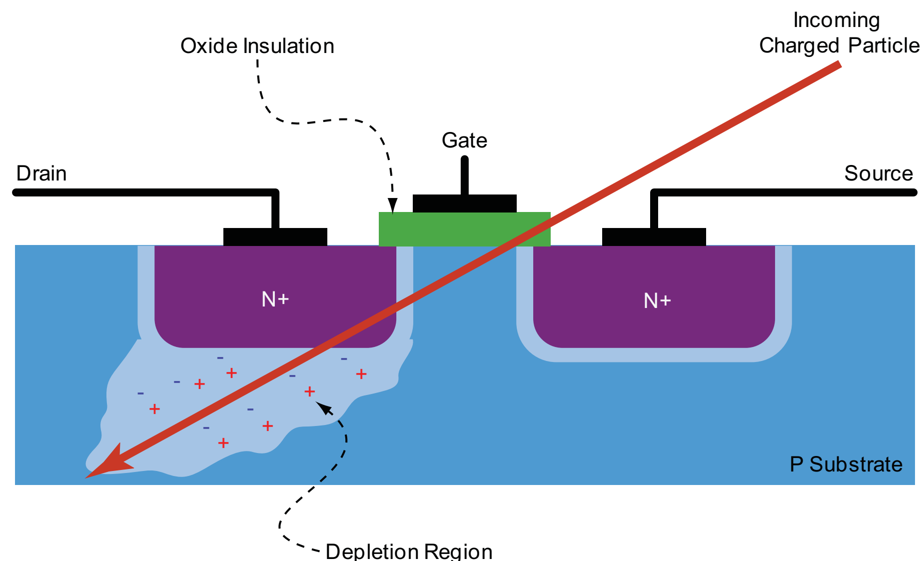

Figure 5: **Impact of a High-Energy Particle**

Similarly, when a high-energy particle such as a neutron strikes the substrate, it collides with atoms in the substrate, liberating a shower of charged particles which then leave an ionization trail. For example, a neutron striking a silicon atom can release energy through elastic and inelastic scattering events or via spallation events that release magnesium and aluminum ions along with alpha particles and protons.

The effect of impacting ionizing particles is typically quantified with respect to the amount of energy transferred to the material along its track, referred to as the linear energy transfer (LET). The amount of energy is often referenced to the density of the material in units of MeV-cm²/mg.

Table 1: **Common Reaction Products for Neutron Interactions in Silicon**

Reaction Products	Threshold Energy (MeV)
$^{25}Mg + \alpha$	2.75
$^{28}Al + p$	4.00
$^{27}Al + d$	9.70
$^{24}Mg + n + \alpha$	10.34
$^{27}Al + n + p$	12.00
$^{26}Mg + {}^{3}He$	12.58
$^{21}Ne + 2\alpha$	12.99

Neutrons are the most abundant component of the GCR spectrum; however, because they are not directly ionizing, predicting their impact is not straightforward—it is the energy of the reaction products that matter more than the energy of the impacting neutron. Looking at the likely reaction products from 14 MeV (a common neutron energy used in SEE testing), Mg and Al ions, the maximum recoil energies are 3.6 MeV and 2 MeV respectively. Taking the maximum recoil energy (for the Mg), the resultant LET in silicon is 7.8 MeV-cm^2/mg.

Single Event Upsets

When a high-energy particle or ion impacts at the depletion region of an N-P junction, charges in the range of femtocoloumbs to picocoloumbs can collect in the region, creating voltage and current transients. The resulting LET can be sufficient to overpower the junction and cause a change in state (bit flip) of the memory element (SRAM cell, register, latch, or flip-flop). This change in state is referred to as a single event upset. Because the effect is temporary, these errors are often referred as being soft—only the data stored in the element is corrupted.

The collected charge needed to change the bias (and thus the state of the transistor) is referred to as the Q_{CRIT} of that element. The amount of transferred energy needed to cause this change of state is referred to as the $LET_{THRESHOLD}$, often used as a measure of susceptibility of a structure to upset. Any circuit with a $LET_{THRESHOLD}$ in excess of 120 MeV-cm^2/mg is considered immune to upset.

Another measure of circuit sensitivity is cross-section. Cross-section is a derived (i.e., fictitious) area that represents the probability of particle interaction with a material and is dependent upon the energy of the particle. The cross-section for a circuit is different for the atmospheric GCR spectrum than it is for the high-energy heavy ions found in space environments). The smaller the number, the lower the probability, usually expressed in units of cm^2.

Effect on SRAM Cells

Of special concern for FPGAs is the effect of this radiation on SRAM cells. In a typical six-transistor SRAM cell (Figure 6), four of the transistors basically form cross-coupled inverters to store the bit value. If an ion strike of sufficient energy occurs near one of these transistors, the bit value stored in the cell can change or flip.

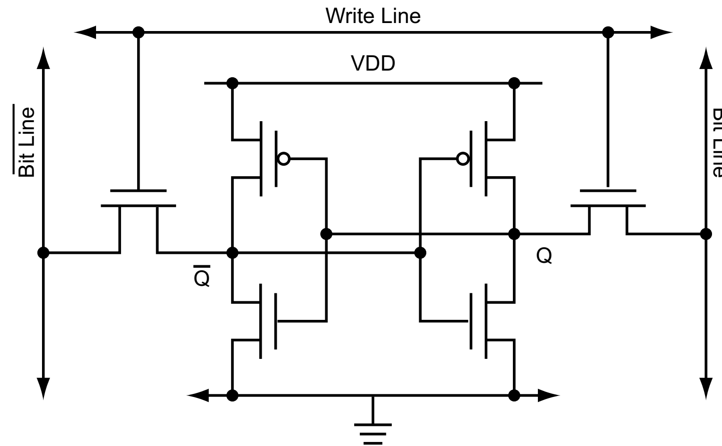

Figure 6: **Six Transistor SRAM Cell**

Due to active feedback of the cross-coupled transistors, the Q_{CRIT} for changing the bit value is also dependent upon the switching speed of the cell (the slower the cell, the higher the Q_{CRIT}). Generally, Q_{CRIT} has decreased as process nodes have shrunk; however, changes in process technology that have lowered junction collection efficiency have acted to ameliorate (but not cancel out) the effect of decreasing transistor size and falling operating voltages.

Simulation results indicate that at 65 nm, SRAM operating at a nominal supply of 0.8 V to 1.2 V is in the range of 2 fF to 3 fF. The decrease in Q_{CRIT} between a 65 nm and 45 nm process is estimated to be on the order of 30% less energy, further increasing the susceptibility of SRAM cells to soft errors.

The upset cross-section for the block RAM and configuration cells in Xilinx Virtex®-5 FPGAs (taken from UG116, *Device Reliability Report*) are 3.96×10^{-14} cm^2 and 6.70×10^{-15} cm^2 per bit, respectively. The resulting FIT rates per megabit from GCR normalized to sea level at NYC are 691 and 161.

Effect on Flash Cells

Flash, an alternative to SRAM for FPGA interconnect technology, is a nonvolatile storage structure. At the heart of a flash memory cell is a floating gate, located between a control gate and the MOSFET structure below, encased in good dialectic. The bit value is stored as a charge on the floating gate (with a charged gate representing a zero value for NOR flash cells). Writing or erasing the cell requires a high voltage (±17.5 V for a ProASIC®3 FPGA) and milliseconds of time to either add or dissipate the charge on the floating gate.

Similar to an SRAM cell, an ion strike in or near the depletion region will deposit a charge. However, the Q_{CRIT} of a flash cell is significantly larger than that of an SRAM cell. Moreover, the flash cells used for configuration feature a far more robust construction than those used in bulk flash memory, which has a focus on speed and size. For example, the amount of charge generated by an ion with a linear energy transfer (LET) of 37 MeV-cm^2/mg—an energy at the very high end of the atmospheric spectrum—is less than 1% of the Q_{CRIT} on a programmed floating gate in the 250 nm ProASIC (flash) configuration switch.

Testing of Microsemi's 130 nm flash process shows that the embedded nonvolatile memory (eNVM) has a $LET_{THRESHOLD}$ of 60 MeV-cm^2/mg. The flash cell used for configuring the FPGA is more robust in nature; about 7.6 times larger and more than double the $V_{THRESHOLD}$, giving an estimated $LET_{THRESHOLD}$ far in

excess of the 120 MeV-cm^2/mg needed to ensure immunity. As a result, flash cells used for FPGA configuration are immune to GCR-induced SEU.

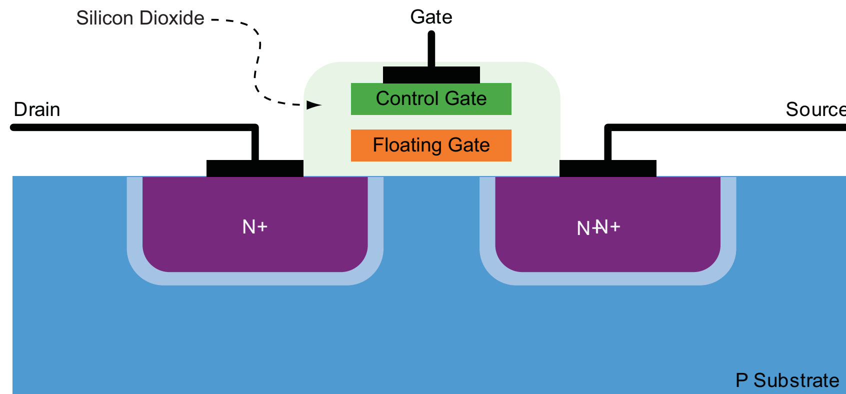

Figure 7: **Flash Memory Cell**

As experimental confirmation, iRoC Technologies conducted a set of experiments designed to test the effects of atmospheric neutrons on ProASIC3 FPGAs in December 2005. Tests were performed in compliance with JEDEC specification JESD-89. Each device tested was exposed to a cumulative neutron fluence of at least 2.42×10^{10} n/cm^2, which is equivalent to exposure to natural background neutron radiation at sea level for >300,000 years (calculated for New York city). Zero flash cell upsets were detected.

Single Event Functional Interrupt

All FPGAs have an array of logic modules (referred to as the fabric), embedded memory blocks, possibly some specialized blocks such as multipliers or DSP, clock management circuits such as PLLs, surrounded by a ring of programmable I/Os. One key area of differentiation between various families of FPGAs is their fabric. Families from different suppliers often differ in the exact structure of the logic modules and how these modules are interconnected or wired together. It is this interconnect that poses the greatest concern from an SEFI perspective.

There are two aspects to FPGA routing: the wire (or metal trace) and the interconnecting via. The vias used in all FPGAs are programmable—the basis of the entire technology. In SRAM-based FPGAs, the basic programmable via is a single-bit SRAM cell. This via is programmed and erased the same way as any other SRAM memory cell. Although more robust than block SRAM, the SRAM via is still susceptible to upset.

These programmable vias are also used to set the configuration of each logic module, I/O and embedded block. An SEU that does affect the functionality of a device is referred to a single event functional interrupt (SEFI).

Possible SEFIs in an SRAM-based FPGA are the following:

- Breaking of a routing connection
- Bridging of two signals
- Shorting of a signal to power or ground
- Changing the functionality of a logic module
- Changing the functionality of embedded block
- Changing the direction or standard of an I/O

In the case of the configuration memory for a typical SRAM-based FPGA, only a small fraction of the millions of configuration bits are used for a given design. As a consequence, not all SEUs to a

configuration bit result in a change to a customer design. Which configuration bits actually impact the configuration of a given design and therefore, if upset, would result in an SEFI is highly design-dependent and hard to predict. Estimates of how many configuration bits in an SRAM-based FPGA are actually used to configure a design ranges from 10% to 30%.

Single Event Transients

When impacting ions induce voltage pulses on combinatorial circuitry in a device, these effects are known as single event transients (SETs). If the induced voltage level exceeds that of the switching threshold and is of sufficient pulse-width, erroneous data values can be propagated through the circuit. As the name implies, these errors are temporary in nature, with pulse-widths on the order of 25 ps. SETs can impact combinatorial logic as well as other elements such as PLLs and charge pumps.

SEE Mitigation

Block Memory

Regardless of the base technology of the FPGA, SEUs in block memory must be mitigated. Fortunately, there are longstanding techniques that can be used to correct for memory errors. These techniques are very effective against SEUs. Error correction and detection depends on additional bits being added to every word in memory. The most common technique, error detection and correction (EDAC), uses Hamming codes that can detect up to two simultaneous bit errors and correct single-bit errors. As data is read from the memory, the EDAC function will correct any errors.

The mitigation performance can be improved by implementing scrubber functionality. A scrubber circuit reads the contents of a target memory constantly to check for and correct for errors. This scrubbing of the data takes place in background during memory idle times. For improved reliability, the EDAC and scrubber can be hardened through the implementation of triple-module redundancy (TMR). See the "Logic" section for more details.

Microsemi has developed a customizable core, CoreEDAC, for use with Microsemi FPGAs to automate the implementation of EDAC circuitry with support for scrubber functionality and TMR hardening. See the *CoreEDAC Handbook* for more details.

Logic

Although the core logic is less susceptible to upset when compared to SRAM, depending upon the criticality of the circuitry and the likelihood of upset, mitigation may be called for. The most common technique for hardening logic is to use triple voting, or triple module redundancy (TMR). With TMR, the logic path is triplicated and after each set of triplicated flops, all three paths feed into a majority gate voting circuit (Figure 8 on page 12). With this technique, if an SET occurs on one of the logic paths or if one of the flops is upset, the results of the other two paths override the error and the correct value is propagated to the rest of the circuit. TMR prevents SETs and SEUs from propagating through the circuit.

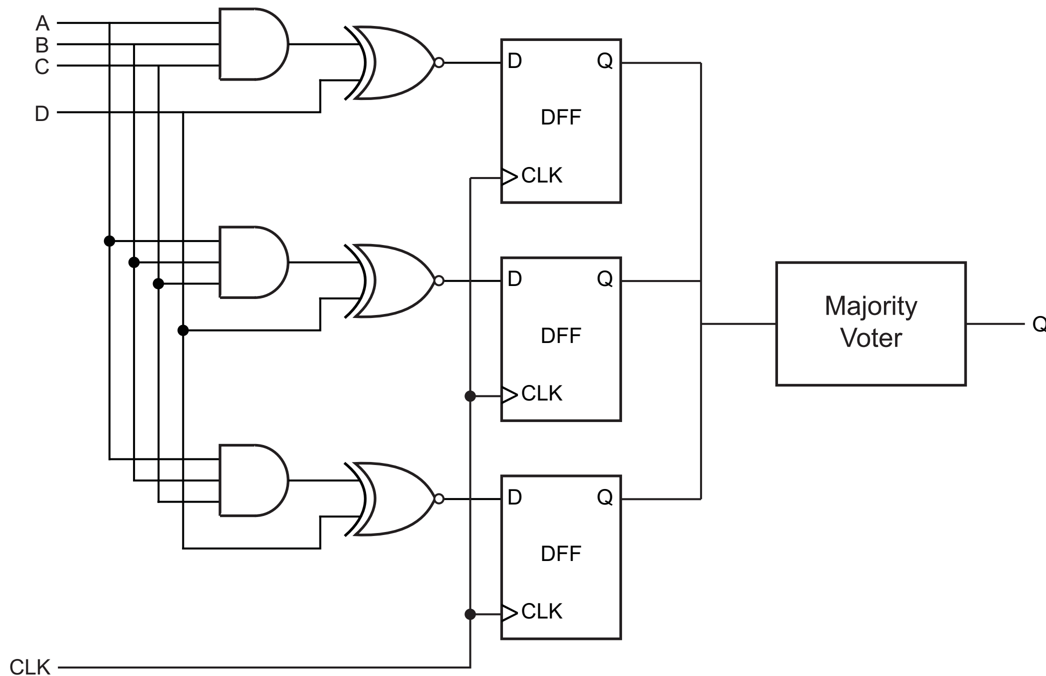

Figure 8: **Triple Module Redundancy Example**

While full TMR provides a high-level of protections against upsets, it is costly in terms of resources, requiring three times the logic to implement plus the logic for the voting circuit. In addition, the voting circuit adds another level of delay. These costs may be too high for many applications. An alternative is to use TMR at a higher level. For example, TMR can be implemented in the following ways:

- At the flop level only – If the FFP analysis determines that the risk presented by SETs to combinatorial paths is low enough that it does not require redundancy mitigation, TMR can be applied only to the flops in the data path. While this technique still requires triple the amount of flops versus the original design, it does lessen the impact on combinatorial resources and is often easier to implement.
- Selective mitigation – TMR can be applied only to the critical paths identified via the FFP analysis. Selective mitigation can greatly reduce the resource impact of TMR.
- At the block level – Rather than mitigate at the data-path level, individual blocks can be implemented normally and their outputs fed to voting circuits. The designer can either triplicate the blocks and allow the voter to propagate the correct result, or simply duplicate the blocks and allow the voter to flag an error condition on mismatch. This approach provides for greater separation between redundant circuits. In the case of simple duplication, it eases the impact on resources.

The exact approach taken to mitigate errors depends both on the FFP analysis and the required design assurance level. Depending upon the approach, TMR can be implemented manually or via synthesis. For example, Synplicity's Synplify® provides TMR support for designs targeted at Microsemi FPGAs.

Configuration Memory

Although the flash cells used for interconnect in FPGAs are immune to upset, the configuration memory of SRAM-based FPGAs require mitigation. Because of the growing awareness of SEUs, manufacturers of SRAM-based FPGAs recommend various mitigation techniques, ranging from simplistic to the more complex.

Mitigating Configuration Memory in SRAM-Based FPGAs

Whole Device Reconfiguration

The simplest method is to reconfigure the SRAM-based FPGA at regular intervals, clearing any SEUs that have accumulated. For this method to be successful, the designer must determine the impact of potential errors, plus the length of time it takes for these errors to propagate through the design. The goal is to reconfigure the FPGA inside the mean time between failures. The errors still propagate, but the potential damage is limited by the reconfiguration. However, the time scale for reconfiguration is on the order of hundreds of milliseconds, during which, the function hosted in the FPGA is unavailable. This downtime may not be acceptable in many applications.

Configuration Memory Scrubbing

With more recent generations of SRAM-based devices, the user can use the built-in error detection scheme in the configuration engine. Using a configuration memory readback feature, the CRC for each configuration frame is calculated and compared to a golden CRC. If a mismatch is detected, then an SEU has occurred, and the application can reconfigure the entire FPGA. Alternately, the application can attempt to correct the error and rewrite the frame in background.

However, scrubbing the configuration memory does not work the same as EDAC implemented on a memory block (with or without background scrubbing). With a mitigated memory block, data errors do not propagate beyond the memory. In other words, the error correction circuitry *intercepts* any errors coming from the block memory. WIth configuration memory scrubbing, an upset configuration bit still controls FPGA circuitry until it is corrected.

Again, the errors still propagate; only the time before they are corrected is reduced when compared to periodic whole-device reconfiguration. Moreover, the detection time is still on the order of milliseconds, equating to millions of clock cycles before an upset can be corrected—certainly enough time for an error to propagate through even the most complex systems. The only alternative to mitigating error propagation is to fully TMR the design, *including* the I/O.

Mitigation Does Not Equal Immunity

Regardless of the methodology, mitigation is used to correct errors *after* the fact, in other words, it attempts to lessen their impact. In all cases, the correction schemes are only able to handle single-bit errors within a configuration memory frame. Any multi-bit errors require full device reconfiguration. In addition, mitigation schemes require additional reliability analysis and engineering time to implement and fully assess the impact of errors that still propagate. Mitigation should not be confused with immunity.

Summary

Various memory elements within electronic devices are suspectable to being upset when impacted by high-energy particles within the Earth’s atmosphere. In addition, other elements of a device may propagate induced pulses or transients that can result in errors in function. Only one supplier of FPGAs offers devices with a base technology that is fundamentally immune to upset. Building on a 20-year history of delivering high-reliability products to commercial avionics, military, and space applications, Microsemi is uniquely positioned to help designers understand the impact of SEUs and SETs and mitigate their effect.

References

1. *Investigation and Characterization of SEU Effects and Hardening Strategies in Avionics*, IBM Report 92-L75-020-2, Aug. 1992, republished as DNA-Report DNA-TR-94-123, Defense Nuclear Agency, Feb. 1995.
2. Bauman, Richard, "Effects of Terrestrial Radiation on Integrated Circuits," Chapter 31, *Handbook of Semiconductor Manufacturing Technology,* Nishi, Yoshio, and Robert Doering, Ed. R Doering and Y Nishi, CRC Press, 2008.
3. Baumann, R.C. "Radiation-Induced Soft Errors in Advanced Semiconductor Technologies," *IEEE Transactions on Device and Materials Reliability,* Vol.5, No.3, pp. 305- 316, Sept. 2005.
4. Chandra, V., Aitken, R.; "Impact of Technology and Voltage Scaling on the Soft Error Susceptibility in Nanoscale CMOS," from *Defect and Fault Tolerance of VLSI Systems*, DFTVS '08, pp.114–122, IEEE International Symposium on 1-3 Oct. 2008.
5. J. Keane et al, “Method for Qcrit Measurement in Bulk CMOS Using a Switched Capacitor Circuit,” NASA Symposium on VLSI Design, June 2007.
6. Fogle, A.D., Don Darling, Blish, R.C. II, Daszko, E., "Flash Memory under Cosmic and Alpha Irradiation," Device and Materials Reliability, IEEE Transactions on, Vol.4, No.3, pp. 371- 376, Sept. 2004.
7. Friedberg W., Copeland K., *What Aircrews Should Know About Their Occupational Exposure to Ionizing Radiation*, FAA Report No. DOT/FAA/AM-03/16.
8. *Radiation Results of the SER Test of Actel FPGA in December 2005,* iRoC Technologies, March 2006.
9. Meyerhof, Walter E., *Elements of Nuclear Physics,* McGraw-Hill, New York, 1967.

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

Microsemi Corporation (NASDAQ: MSCC) offers a comprehensive portfolio of semiconductor solutions for: aerospace, defense and security; enterprise and communications; and industrial and alternative energy markets. Products include high-performance, high-reliability analog and RF devices, mixed signal and RF integrated circuits, customizable SoCs, FPGAs, and complete subsystems. Microsemi is headquartered in Aliso Viejo, Calif. Learn more at **www.microsemi.com**.

© 2011 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.

55900134-0/8.11