



WP402 (v1.0.1) March 7, 2012

Considerations Surrounding Single Event Effects in FPGAs, ASICs, and Processors

By: Dagan White

Single event effects (SEEs) are of a growing concern in high-reliability system development, yet there is much disparity among users of ASICs and FPGAs with regard to understanding how susceptible their designs might be. The avionics and industrial system development guidance that currently exists is only broadly beginning to consider SEEs and their impact on system reliability. Unfortunately, standards such as DO-254, DO-178, ARP 4754, ARP 4761, and IEC 61508 provide little or no guidance on how to handle SEEs. This white paper highlights concerns regarding effects of SEEs on ASICs and FPGAs and points to analysis and mitigation techniques for handling SEEs.

Introduction

All sub-micron integrated electronics devices are susceptible to SEEs to some degree. The effects can range from transients causing logical errors, to upsets changing data, to destructive soft error latch-up (SEL). Traditionally, FPGAs were targeted as being more sensitive due to their use of SRAM for the configuration storage. As dimensions shrink to below 90 nm, SEEs in all devices (ASICs, ASSPs, and FPGAs) must be considered.

While targeted to an avionics audience, this white paper has broad applicability to any industry where safety and reliability are of critical importance. It should be useful to a wide audience comprised of system architects, engineering and program managers, and certification authorities. Some knowledge of programmable or custom devices, with or without microprocessors, and associated design methodologies is assumed.

Understanding the SEE Phenomenon

A significant amount of literature currently exists regarding the physics of SEEs inside integrated circuits, but a brief introductory overview is given below to assist readers less aware of the phenomenon.

Causes of SEEs

SEEs result from interaction of high-energy particles with circuit elements in integrated circuits. When a high-energy particle passes through the silicon substrate of a device, charged particles are created as the result of sub-atomic particle collisions. These particles are generated by an ionization trail along the path of the incoming particle.

For example, if a charged particle impacts at or near a transistor junction, the collected charge can induce an upset to the state of that transistor. If the collected charge is larger than the critical charge of the element, the element changes state. This change in state (or bit flip in the case of a memory cell) is referred to as an SEU. Similarly, the charged particles can induce a current and voltage spike on a metal interconnect, which is referred to as a single-event transient (SET). If the pulse width of the spike is wide enough, the spike can propagate through the circuit (see [Types of Single Event Effects](#)).

Sources of Charged Particles

Two sources of charge particles are of concern to designers of high-reliability systems: cosmic ray interactions with the atmosphere, and impurities in packaging materials and the silicon substrate.

Atmospheric Sources

Galactic cosmic rays (GCR) originate in outer space, are primarily comprised of subatomic particles and light ions, travel at nearly the speed of light, and strike Earth from all directions. As high-energy cosmic rays enter the atmosphere and react with atoms, through a process known as direct nuclear spallation, neutrons are generated in the atmosphere. The result of this phenomenon is often referred to as an air shower. Neutrons with energy greater than 10 MeV carry sufficient energy to cause SEEs in integrated circuits.

Atmospheric depth (density) also plays a significant role in causing neutron-generating reactions and in transporting neutrons to ground level. An intense neutron environment exists at higher altitudes in the atmosphere, 10 km to 40 km and more above the surface. In addition, Earth's magnetic field causes the flux to vary from the equator to the poles, with the equator having the least flux and the poles having the greatest flux. The magnetic field of the sun as it varies during the sunspot cycle also influences the flux of cosmic rays [Ref 1]. For example, maximum flux occurs at a solar minimum.

Packaging Material Impurities

Packaging materials used for integrated circuits often contain impurities. Among these are trace amounts of uranium and thorium isotopes, which emit alpha particles as they decay. Although these particles are low in energy and have limited penetration depth, they are a concern for integrated circuits due to their close proximity to the silicon substrates.

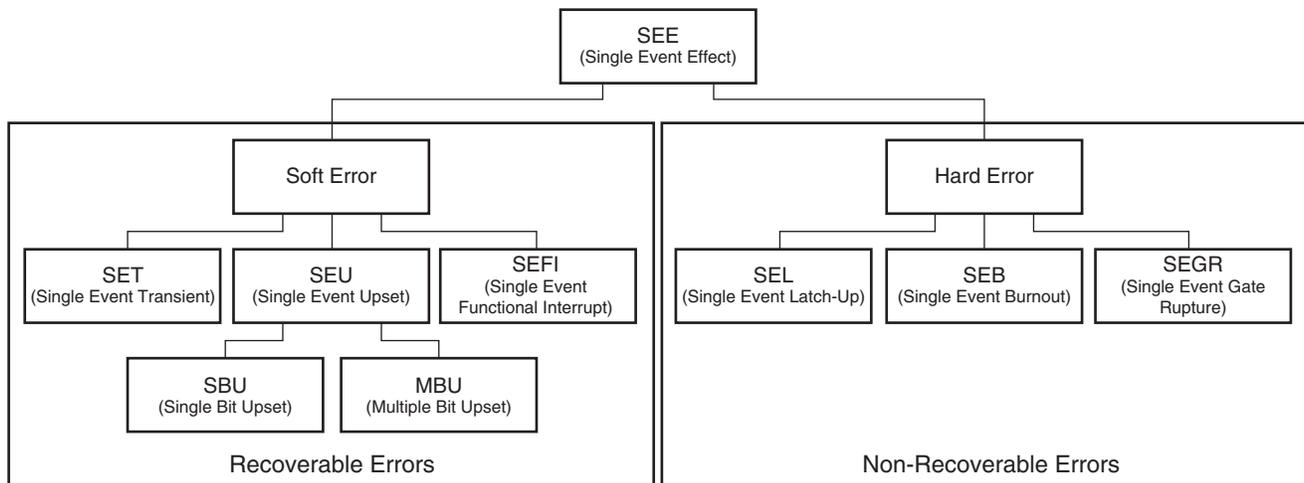
Another source of alpha particles in packaging is the eutectic lead solders used for the solder bumps in flip-chip packaging. Even if the solder is purified of other radioactive impurities, it is impossible to remove the lead isotope ^{210}Pb . Although ^{210}Pb is not an alpha emitter per se, its decay chain contains the strong alpha emitter ^{210}Po .

Substrate Impurities

The element boron used in borophospho-silicate glass (BPSG) is another source of ionizing radiation. When one of the common boron isotopes, ^{10}B , is struck by low-energy neutrons, an alpha particle and a lithium ion are generated. Given the significant amount of boron present in substrates plus the number of low-energy neutrons in the air shower, the effect is significant.

Types of Single Event Effects

A number of events fall under the general category of SEEs. These events or errors can be divided into two broad categories: soft versus hard errors. Soft errors are those events that have no damaging effects and are cleared by normal device operation. Hard errors are events that generally result in lasting damage to the circuitry. See [Figure 1](#).



WP402_01_082311

Figure 1: Types of Single Event Effects

Soft (Recoverable) Errors

Soft errors are upsets to the device operation and are self-correcting in time or are correctable by rewriting a memory element. The three subclasses of soft errors are:

- Single-event transients (SETs) result when a high-energy particle impacts a combinatorial path of a device and can induce a voltage/current spike. If the pulse-width of this spike is sufficient and at the right time, it can propagate through the circuit.
- Single-event upsets (SEUs) are the result of high-energy particles causing a change in the state of a memory element (SRAM, flash, flop, or latch). SEUs can be categorized as single-bit or multi-bit upsets (SBUs or MBUs). SBUs are by far the most common SEE seen in avionics applications.
- Single-event function interrupts (SEFIs) are disruptions to normal device operation (beyond a simple corruption of user data). These types of effects alter the functionality of the circuit and typically require reconfiguration/reset or power cycling for recovery.

Note: Failures-in-time (FIT) rates are commonly discussed in relation to SEUs, SETs, and SEFIs, but these are soft errors that affect the functionality and not permanent failures of the device.

Hard (Non-Recoverable) Errors

Errors that cause lasting damage to the device are classified as hard errors. The three subclasses of hard errors are:

- Single-event latch-up (SEL) is a circuit latch-up induced by radiation. This latch-up can be either permanent or clearable with power cycling.
- Single-event burnout (SEB) is a short-circuiting caused when a high-energy ion impacts a transistor source, causing forward biasing. SEBs are typically a threat to power MOSFETs but are also seen in IGBTs, high-voltage diodes, and similar circuits.
- Single-event gate rupture (SEGR) is a plasma spiked caused by a high-energy ion impact, resulting in rupture of the gate oxide insulation.

Xilinx® FPGAs are not susceptible to latch-up and gate rupture caused by neutron radiation, as demonstrated through both internal and external Xilinx device testing. The same might not be said for other FPGA or ASIC vendors. Xilinx space-grade parts are immune to latch-up from heavy ions as well, but this type of radiation is not an issue inside of earth's atmosphere.

ASICs, FPGAs, and SEUs

DRAM was the first technology where terrestrial SEU became a concern, but these devices are now fairly robust. SRAM soft error rates (SERs) then became a concern and are still a concern today because even though the per-bit SER has held steady despite the decreasing feature size, the total amount of SRAM bits per system/device has increased greatly. SRAM is used inside stand-alone memory devices as well as FPGAs. Concern over FPGAs arises from their use of SRAM for user block memory as well as device configuration memory. With the latest sub-90 nm technology nodes, concern over ASIC upset rates is rising.

SRAM-based FPGAs hold the device routing in a configuration memory, and they use block RAMs for user memory. Both of these memory structures, along with flip-flops, can be upset by radiation, although at different rates. User block RAM can be protected with error-correcting code (ECC) and parity schemes, as can external memory devices. FPGA configuration memory, however, cannot be directly protected in the same manner as block memory via ECC or parity checks. SEU mitigation techniques that monitor device configuration memories are recommended for FPGA designs. Xilinx devices have built-in configuration memory error detection capabilities (using ECC), and SEU mitigation IP is available to monitor and repair configuration memory. Other FPGA structures are upsettable as well but at an insignificant rate.

Note: The SRAM cells used for the configuration memory of Xilinx FPGAs are larger and more robust than the SRAM cells used for general-purpose memory, which are optimized for speed and cost. Moreover, the Xilinx configuration memory cells have been optimized for SEU resistance. Xilinx has been actively improving the SEE resilience of its configuration memory over the last 12 years, whereas no such action has been taken in the ASIC industry.

SEE concerns in ASICs have risen because of the decreased operating voltages and element capacitance combined with increased clock speeds. These factors mean that transient upsets are more likely and can easily translate to clocked functional errors. Soft error rates can now easily exceed 50,000 FIT per processor, including logic gates and on-chip memory. System-level consideration and mitigation techniques are necessary for ASICs [Ref 2]. Other data shows that ASIC designs below 90 nm have exhibited 1,000 FIT per million gates, and 1,000 FIT per million memory bits [Ref 3]. User memory can be protected—but logic upsets, which can account for a substantial portion of the upset rate, cannot be easily protected. Logical SETs, when latched, can lead to logic errors and consequently are no longer negligible in processors manufactured on deep sub-micron processes. System-level solutions are required [Ref 4].

At the same time that ASICs have become more susceptible to upset, Xilinx FPGAs have been designed for improved immunity and lower soft error FIT rates. In fact, Xilinx devices at 65 nm and below have shown improved immunity, with nominal rates on the order of below 100 FIT/Mb for configuration memories and below 500 FIT/Mb for user memories (see [Ref 5] for exact rates per device family). This improvement is the result of Xilinx's work and study of SEU phenomenon since 2002 [Ref 3].

For both ASICs and FPGAs, there are non-zero error rates, non-zero detection times, and non-zero correction times. It is imperative to consider SEEs both when using ASICs and FPGAs in any high-reliability application. Some vendors provide information to assist in analysis of system FIT rates. For example, Xilinx provides public information regarding SEU FIT rates for its devices via publications [Ref 5] [Ref 6] that helps customers estimate the FIT rate for their targeted device. Few, if any, ASIC vendors actually publish their FIT rate data, which only serves to mask an increasingly concerning issue. Exact processor FIT rates are also tricky to determine, requiring a combination of analysis, simulation, and beam testing.

Beyond vendor data, some airframe manufacturers have SEE models or estimations, which they apply and levy broadly across multiple vendors' technologies. This approach provides a rough estimate for those vendors that do not supply data, but this approach is risky. In contrast, Xilinx has significant history regarding radiation testing and characterization as evidenced through the Xilinx Radiation Test Consortium—a collaborative effort across many companies in the aerospace industry. See Xilinx's Space website for more information [Ref 7].

SEE rates are probabilistic and vary with geographic location, environmental conditions, and altitude. All FIT rates are estimates based on modeling, analysis, and/or testing, but all published FIT rates are not necessarily created equal. Xilinx data complies with the JEDEC Standard 89A (JESD89A). In fact, Xilinx played a role in updating this specification as a result of its expertise and leadership at the forefront of neutron radiation research [Ref 6]. The Xilinx FIT rate calculator applies the models from JESD89A with FPGA FIT rate data to yield an adjusted, application-specific FIT rate.

All radiation testing is not created equal either. For example, particle test beams can vary in their energy and particle distribution. To counteract this variability, Xilinx employs control devices when conducting beam testing to adjust for test setup and beam variation from run to run. Flight tests can capture real-world data regarding FIT rates, but geographic location and timing with the solar activity can cause variability in the data. Ultimately, soft error FIT rates are estimates that enable the developer to assess the probability of fielded system upset rates.

Mitigation Approaches

To understand the various mitigation approaches, several scenarios can be examined. For example, one scenario includes a processor having a FIT rate of 600 at sea level in New York City (USA) corresponding to a mean time between failures (MTBF) of roughly 190 years. While an MTBF of this magnitude can seem insignificant, if 1,000 systems are fielded, then the combined MTBF of all systems drops to 70 days—one upset every 70 days on average. This rate might not be tolerable for high-reliability system such as industrial applications or networking routers.

In a different scenario, altitude is examined. A FIT rate of 600 at sea level at NYC corresponds to a rate of 367,200 at 40k feet elevation over the poles, representing a MTBF of 110 days for a single fielded unit. Flying a hundred units results in roughly one upset per day. One system in the air has the nearly same magnitude of upset as 1,000 systems on the ground.

Both the memory and logical structures in ASICs are susceptible to SEEs, especially at sub-90 nm technology nodes. Similarly, FPGA configuration memory and user block memory are upsettable. This susceptibility does not mean that these technologies are

unsuitable for avionics and high-reliability systems; it means that SEEs should be considered and mitigation tactics must be employed.

Designers should assess the following before making a final ASIC or FPGA selection:

- Frequency of events - FIT rate and MTBF
- Detection time of events, and means of detecting the event
- Recovery time after event detection
- Performance, area, and monetary cost of the mitigation solutions
- System performance and system design implications

These fault detection and mitigation techniques should be considered when designing with both ASIC and FPGA solutions:

- Soft error mitigation IP (SEM IP)—good for FPGAs and soft processor only
- ECC or parity checks for user memories in both ASICs and FPGAs
- Software-implemented fault tolerance (SWIFT) for both soft and hard processor solutions
- Hardware mitigation solutions—lockstep operation, dual and triple module redundancy (DMR and TMR) for FPGA solutions or ASIC designs
- Watchdog timers

All mitigation approaches should consider area, performance, detection time, and correction time balanced against fixed and variable costs as well as system safety and reliability costs.

Available Xilinx FPGA SEE mitigation methods include:

- External watchdog timer with external handling control (lacks full device check)
- Full-device cyclic redundancy check (CRC) with external reset of FPGA (might upset operation when unnecessary)
- Full-device CRC with bit correction and flag to design (design can decide on further actions)
- Full-device CRC with correction and non-essential bit classification (ignores 66% of false positives). See [Architectures and Refinement of FIT Rates](#) for a description of essential bits.
- DMR and TMR design techniques, or lockstep operation (area hit)
- Additional built-in fault tolerance checks (custom generated)
- Safe state machines - “safe_implementation” and “when others” statement with recovery state
- SWIFT techniques (for processors)
- Memory protections using ECC or parity checks
- Flow checks, range checks, signatures, CRCs, parity, etc.

ASIC SEE mitigation methods include:

- External watchdog timers (can catch every time-dependent behavior)
- Architectural mitigation (costly solutions on top of increasingly costly technology nodes)
- SWIFT techniques (for processors)
- Memory protections using ECC or parity checks

ASIC and Processor Robustness

With each successive process node, the cost of ASIC NRE increases by \$5M or more. At the same time, the ASIC susceptibility to SEEs is increasing as the operating voltage and elemental capacitance decreases. These smaller technology nodes are the critical enabler of power reduction and increased performance with higher clocking speeds. All of these aspects drive greater design density. Larger and larger end-markets are now necessary to support the non-recurring cost of developing modern ASICs. TMRing, lockstepping, or whatever silicon mitigation techniques might be employed to enhance ASIC immunity to SEE are contradictory to the natural evolution of commercial-grade ASICs. While these reliability enhancing features are desirable for high-reliability markets, they are not necessarily desirable for mainstream COTS markets [Ref 8]. Commercial markets might not care about the SEE frequency.

Boeing conducted a detailed study of the affects of SEE on the clock, flip-flop, and logic structures inside of a commercial-grade 90 nm standard-cell ASIC, with the conclusion that hardening techniques must be considered and applied differently across all circuit structures in the device to achieve an appropriately hardened ASIC suitable for avionics applications. This study identifies some of the complex considerations that go into hardening different elements of ASIC structure for the ultimate goal of building a SEE-robust ASIC [Ref 9].

In lieu of hardware solutions or in combination with such solutions, software-implemented fault tolerance is one means of enhancing SEE handling in processor hardware. Much research has been conducted on SWIFT techniques, but more innovation might be required to turn the research into viable market solutions—and this burden will likely fall on the high-reliability market. Many of the techniques for handling SEUs in processors is applicable to both soft and hard processor solutions, a benefit for both ASIC and FPGA-based processors.

Many possible software methods are available to address processor upsets. Software techniques can include data-flow error monitoring and control-flow monitoring, but these techniques have not reached 100% coverage. Hardware techniques might include memory access checks, consistency checks, control-flow checks, watch dogs, and dynamic verification. Soft and hard processors may require different strategies in some cases. One study [Ref 10] has shown that for a soft processor, a hybrid hardware and software approach can yield 100% fault detection with processing time overhead around 150% of the non-mitigated design.

Research work from Brigham Young University (BYU) assesses SWIFT techniques versus DMR/TMR techniques in terms of performance and area solution costs. While this research is geared for space-based applications and is focused on soft processors, the same concepts can be applied to terrestrial and airborne systems that use both soft and hard processors. This work shows that software implemented techniques can achieve decent detection and correction rates versus DMR and TMR, with all solutions capturing greater than 90% of the errors. SWIFT techniques do, however, lead to a performance hit nearing 2X. On the other hand, DMR and TMR are costly in terms of area, with 2.5X and 3.7X area hit respectively, but they do achieve greater detection rates with only minimal performance hits [Ref 11]. The designer needs to review the trade-offs when selecting mitigation methods.

Architectures and Refinement of FIT Rates

It is a tricky proposition to assess the effect of an SEU to an FPGA configuration bit for any specific end design. First, which bits are truly critical to a user design? And second, if a bit is critical to the design, is it critical to the function at the time that the bit upset occurs and prior to its correction?

Joint research work by BYU and Los Alamos National Laboratory (LANL) [Ref 12] assesses the vulnerability of FPGA designs to configuration bit upsets and examines the bits that are critical to a design. For those configuration bits that are critical, the research explores which ones, even if functionally corrected, might not correct a disturbance of the processing state. The study classifies bits as either persistent or non-persistent, referring to state-machine control bits or feedback bits that when upset can corrupt processing versus corruption of passing information in the datapath (such as corruption of a video display data bit). The results demonstrate that the proportion of persistent bits in a design depends on the design architecture [Ref 12].

Similar questions arise when assessing SEEs in the logical structure that controls an ASIC. Intel Corporation and others have recognized this issue and have conducted research in an attempt to quantify an architectural vulnerability factor [Ref 13] [Ref 14]. The theory of the work applies to any soft or hard processor. Xilinx recognizes similar ideologies and has carried out similar research focused on FPGAs.

Generally, the FIT rate for the configuration memory of an FPGA is calculated simply by multiplying the FIT/Mb by the configuration memory size (after subtracting overhead bits and block RAM content). However, the results are overly pessimistic as only a maximum of 10% of the configuration bit upsets actually result in a functional failure in the design. Similarly, an SEU mitigation strategy that flags every upset to configuration memory as being critical results in many false positives.

However, determining which bits are critical to a design is a time-consuming project that requires injecting faults into every configuration bit of an end design. To simplify the process, Xilinx developed *essential bits* technology. The essential bits output produces a list of bits that affect functionality of the design. In contrast, critical bits represent a subset of the essential bits that results in a functional failure in the design if upset. For example, an essential bit upset in a non-active area of the design (in higher order bits of a counter, a rarely used state, or test circuitry) does not result in a functional failure. The essential bits output is conservative but can still allow the user to rule out 66% or more of the configuration bits for a given design.

Using the essential bits output with SEM IP, which detects and corrects upsets, allows the system to ignore non-essential bit upsets. Non-essential bits are still corrected to prevent accumulation of errors, but the design can continue to operate without further intervention. If an essential bit is upset, then that bit is corrected, and the user design can determine whether or not a device reset is prudent (depending on architectural knowledge of the design and the effects of persistent and non-persistent errors). Using this technology, the effective FIT rate of a full device is greatly reduced—to 33% or less.

Even if an essential and critical bit upset is corrected, an error can still propagate. DMR/TMR and other architectural techniques are required to guarantee uninterrupted operation. An upset that affects a feedback or decision path could propagate or place the design in an unintended mode prior to correction of the upset configuration bit. For this reason, short of robust architectural mitigations, it is prudent to correct all upset bits, and then, if it is an essential bit upset, internally reset the device. Xilinx is continuing to develop technologies that can enhance the fidelity of SEU responses.

Xilinx enables users to employ various levels of SEU protection (see [Mitigation Approaches](#)) and recommends that designers:

1. Assess the soft error data for device families [\[Ref 5\]](#).
2. Select a device family that supports SEM IP (Virtex®-5 FPGAs and later) [\[Ref 15\]](#).
3. Employ the SEU FIT Rate Calculator (available from Xilinx) to assess the soft error FIT rate and MTBF for the design and target device with the level of device utilization and environmental conditions that are expected. This is a preliminary assessment tool.
4. Complete the normal design process incorporating the SEM IP.
5. Simulate the design and use the SEU fault-injection simulation capability to verify the design. Additionally, simulate forced invalid states in state machines.
6. Use the ISE® Design Suite 13.2 (or later) essential bits output data to assess the estimated SEU rate for the design. These are refinements that can be fed back into the FIT Rate Calculator to yield a more accurate estimate of the design FIT rate. The Essential Bits outputs can be used with certain versions of the SEM IP and target devices to reduce unnecessary handling of false-positive SEU hits. See the Xilinx Avionics website [\[Ref 15\]](#) for further information on these refined SEU features.

Conclusion

Systems that utilize sub-90 nm geometries, products like ASICs and FPGAs, in any avionics or high-reliability application must adopt proper techniques to mitigate the susceptibility of such technologies to SEEs. FIT rate estimates can be used to assess the MTBF of these technologies for the proper mitigation at the device and system level. Any mitigation strategy ultimately needs to address trade-offs that include area, performance, detection time, and correction time. These factors need to be balanced against fixed and variable costs as well as system safety and reliability costs.

More Resources

- Xilinx SEU website
www.xilinx.com/products/quality/single-event-upsets.htm
- Xilinx Processor Solutions (Zynq™-7000 EPP, MicroBlaze™ processor, LEON, and miniMIPS)
<http://www.xilinx.com/products/technology/embedded-processing>
- Soft Error Mitigation (SEM) Core
<http://www.xilinx.com/products/intellectual-property/SEM.htm>
- DO-254 compliant version of MicroBlaze processor
<http://www.xilinx.com/applications/aerospace-and-defense/avionics>
- BYU Open Source TMR
<http://reliability.ee.byu.edu/edif/>
- Mentor Graphics solutions (Precision HiRel and TMR)
<http://www.mentor.com/products/fpga/synthesis/precision-hi-rel/>

References

1. [XAPP1073](#), *NSEU Mitigation in Avionics Applications*
2. *Soft Errors in Advanced Computer Systems*, Robert Baumann, Texas Instruments, 2005 IEEE Copublished by the IEEE CS and the IEEE CASS
3. [WP256](#), Xilinx FPGAs Overcome the Side Effects of Sub-90 nm Technology
4. *Coping With SEUs/SETs in Microprocessors by Means of Low-Cost Solutions: A Comparison Study*, M. Rebaudengo, M. Sonza Reorda, M. Violante, B. Nicolescu, and R. Velazco, Italian Space Agency and MIUR CERCOM, 2002, IEEE Transactions on Nuclear Science
5. [UG116](#), *Device Reliability Report*
6. [WP286](#), *Continuing Experiments of Atmospheric Neutron Effects on Deep Submicron Integrated Circuits*
7. Xilinx Space website
<http://www.xilinx.com/applications/aerospace-and-defense/space/index.htm>
8. [ReStore: Symptom-Based Soft Error Detection in Microprocessors](#), Nicholas J. Wang, Sanjay J. Patel, University of Illinois
9. [Clock, Flip-Flop, and Combinatorial Logic Contributions to the SEU Cross Section in 90 nm ASIC Technology](#), David L. Hansen, Eric J. Miller, Aj Kleinosowski, Kirk Kohnen, Anthony Le, Dick Wong, Karina Amador, 2009, IEEE Transactions on Nuclear Science
10. [Detecting SEEs in Microprocessors Through a Non-Intrusive Hybrid Technique](#), José Rodrigo Azambuja, Ângelo Lapolli, Lucas Rosa, and Fernanda Lima Kastensmidt, IEEE, Federal University of Rio Grande do Sul, 2011, IEEE Transactions on Nuclear Science
11. [Software Fault-Tolerant Techniques for Softcore Processors in Commercial SRAM-Based FPGAs](#) (and other related work based on discussions), Nathaniel H. Rollins and Michael J. Wirthlin, NSF Center for High-Performance Reconfigurable Computing (CHREC), Brigham Young University, 2011
12. [Detection of Configuration Memory Upsets Causing Persistent Errors in SRAM-based FPGAs](#), D. Eric Johnson, Keith S. Morgan, Michael J. Wirthlin, Michael P. Caffrey, Paul S. Graham, Brigham Young University and ISR-3 Los Alamos National Laboratory, 2004, 7th Annual Military and Aerospace Programmable Logic Devices International Conference
13. [A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-Performance Microprocessor](#), Shubhendu S. Mukherjee, Christopher Weaver, Joel Emer, Steven K. Reinhardt, and Todd Austin, Proceedings of the 36th International Symposium on Microarchitecture, 2003
14. [Versatile Prediction and Fast Estimation of Architectural Vulnerability Factor from Processor Performance Metrics](#), L. Duan, B. Li and L. Peng, In Proceedings of the 15th IEEE International Symposium on High-Performance Computer Architecture (HPCA-15), Raleigh, NC, Feb. 2009
15. Xilinx Avionics website
<http://www.xilinx.com/applications/aerospace-and-defense/avionics>

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
09/12/11	1.0	Initial Xilinx release.
03/02/12	1.0.1	Minor typographical edits.

Notice of Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.