Introduction to Single-Event Upsets

WP-01206-1.0 White Paper

This paper provides an overview of single event upsets (SEU), the capabilities provided in FPGAs to mitigate the effects of SEU, techniques that can be incorporated in user designs to mitigate the effects of SEU, and how tools and intellectual property (IP) can be used to validate a design to ensure high levels of tolerance for SEU.

Introduction

SEU is the change in state of a storage element inside a device or system. This state change is a soft error and can often be fixed by changing the state of the storage element back to its original value. The following sections discuss the taxonomy, sources of SEU, why SEU may or may not be important for your designs, and mitigation of SEU in Altera's FPGA devices. If you are already familiar with SEU, consider going directly to the overview section of Altera's SEU mitigation capabilities.

Cosmic Rays, Cascading Particle Showers, and Altitude

Cosmic rays, a term coined by Victor Hess in 1912, are high energy particles originating from outside of our atmosphere (1). When these high-speed particles collide with molecules in our atmosphere a cascading particle shower can be created. Figure 1 is a simplified view of a particle shower.

101 Innovation Drive
San Jose, CA 95134
www.altera.com

Figure 1. Simplified Particle Shower

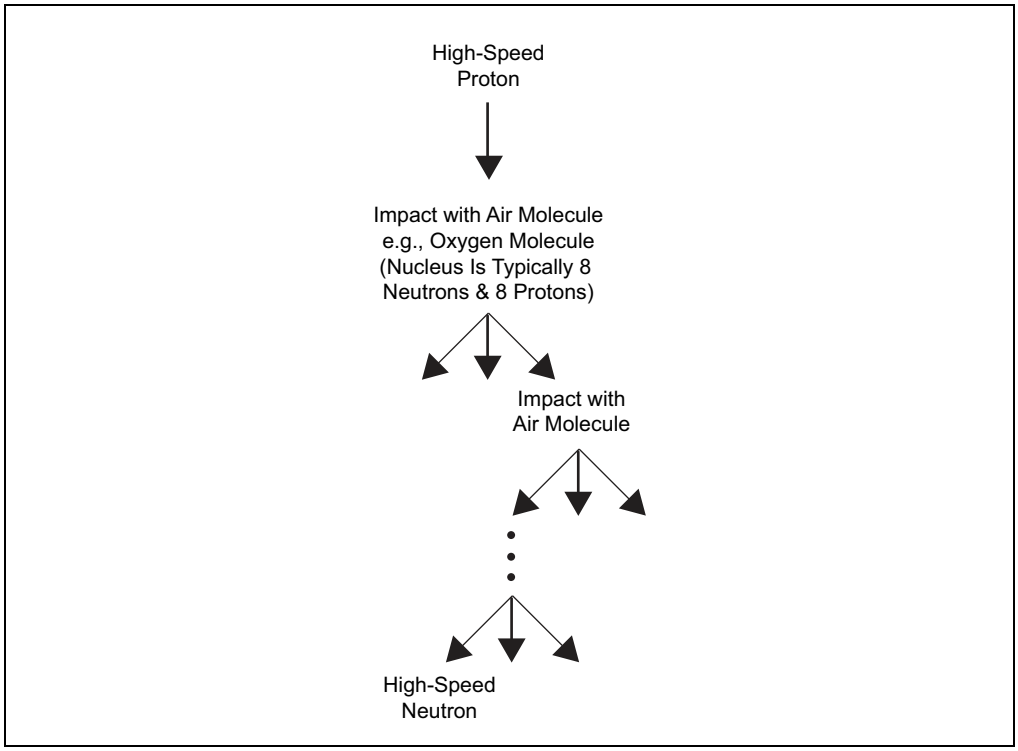

For additional information on particle showers, refer to the Air Shower (physics) page, available on Wikipedia (en.wikipedia.org/wiki/Air_shower_(physics)).

Cosmic rays are usually high-speed protons and when they collide with an air molecule, many nuclear interactions can occur including the generation of high-speed neutrons. These high-speed neutrons are one of the two main sources of SEU in semiconductor devices. The other main source is alpha particles, which will be discussed in the next section.

As the atmosphere gets denser closer to the surface of the earth, there are more air molecules for high-speed neutrons to collide with. Additionally, each collision results in lower energy neutrons. Conversely, at higher altitudes there are more high-speed neutrons because there is less likelihood of hitting another air molecule. Therefore, the neutron flux, or the neutron count per area per time, increases with altitude. Per the JEDEC specification (JESD89A), the typical neutron flux in New York City is 13 neutrons per square centimeter per hour. Using the calculator available at www.seutest.com, the flux at an altitude of 10,000 feet above New York City is 11.1 times greater, or 144 neutrons per sq. cm per hour. This altitude effect is the reason that one of the first groups interested in SEU mitigation was engineering teams focused on aeronautics designs using semiconductor devices.

What about shielding your design from neutrons? Since one foot of concrete only reduces the neutron flux by 30%, in most situations it is not possible to shield an electronic system from neutrons. Therefore, the focus is on understanding the failures in time (FIT rate) from high-speed neutrons and how that will affect the system uptime. One FIT is one failure in one billion hours of operation ($1x10^9$).

Alpha Particles and Ultra-Low Alpha Packaging Materials

What about the alpha particles mentioned earlier? Alpha particles are helium nuclei which are emitted as a result of radioactive decay. In semiconductor devices, the main source of alpha particles is from packaging materials. Alpha particles are charged particles so they can cause an ionization track in a semiconductor. An ionization track or path is shown in Figure 3. An ionization track is the path through the semiconductor material where the ion has created free electrons and holes.

Due to the low energy of alpha particles, only the packaging materials that are in very close proximity to the silicon are of consideration. This means that the die bumps and the underfill material for flip chip die are the two key sources of alpha particles in flip chip devices. Whereas the mold compound is a key source of alpha particles for wire bond devices. Since this is well known in the semiconductor industry, most semiconductor companies use ultra-low alpha materials for die bumps, underfills, and mold compounds.

A subsequent section will discuss the alpha particle interaction with silicon and provide the basics for how an SEU event happens.

Is SEU Important for Your Design?

All semiconductor devices are susceptible to SEU and for many years only designs that went to higher altitudes had SEU mitigation as one of their design requirements. However, as system uptime requirements continue to go up, SEU and related mitigation techniques are becoming more important in many applications. Examples are high-reliability computing and storage applications being driven by data centers and cloud computing. Still, the majority of designs have not reached the point where SEU mitigation is needed. If SEU requirements are not specified in your design specifications and product specifications, you will likely not need to use any mitigation techniques.

SEU Taxonomy

When it comes to formal terms used to describe SEU, there has been a lack of consistent industry usage, so some of the terms used in this paper may differ from other materials you may have seen. Figure 2 shows the terms and their relationships that this paper will follow.

Figure 2. SEU Acronyms

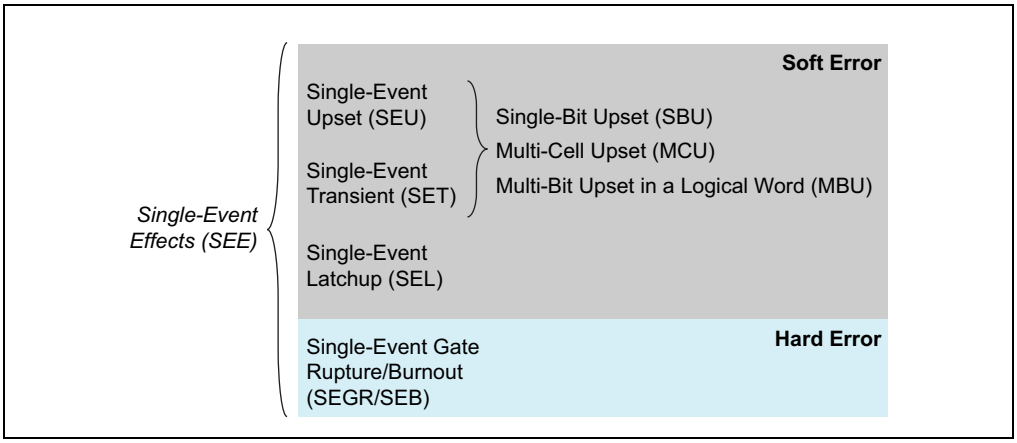

Definitions for the terms used in Figure 2:

Soft Error:	Storage element (memory cell, latch, or register) state change. No hardware damage and is correctable.
SET:	Single Event Transient. A glitch caused by single event effect, which travels through combinational logic and is captured into storage element.
SEU	Single Event Upset. Storage element state change – may affect a single bit or multiple bits.
SBU	Single Bit Upset. A single storage location upset from a single strike.
MCU	Multiple Cell Upset. Multiple storage locations upset from a single strike.
MBU	Multiple Bit Upset. Multiple upsets in a logical word from a single strike.
SEL	Single Event Latchup. The event creates an abnormal high-current state by triggering a parasitic dual bipolar circuit, which requires a power reset. It can possibly cause permanent damage to the device, in which case the result is a hard error. *(1)*

Additional definitions related to single event effects (SEE):

SER	Soft Error Rate. The statistical probability of a soft error over time typically specified in FITs
FIT	Failure in Time. One failure per 10^9 device operating hours.
SEFI	Single Event Functional Interrupt. A functional failure caused by an SEE. Refer to the *Understanding SEFI (Single Event Functional Interrupt) in FPGA Designs* White Paper for additional details.
SEFI ratio	The ratio of the number of SEE events divided by the number of functional interrupts or failures.
SEGR	Single Event Gate Rupture, damage of the gate oxide and the resulting current path.
SEB	Single Event Burnout, creation of a high-current state resulting in device damage.

Another resource is www.jedec.org, which has the official JEDEC definitions of these terms.

High-Speed Neutrons, Alpha Particles, and SEU

How do both a charge-neutral particle (neutron) and a charged particle (alpha) create a soft error? Let's start with the charged alpha particle and a graphic (Figure 3) that shows the effects:

Figure 3. Alpha Strike

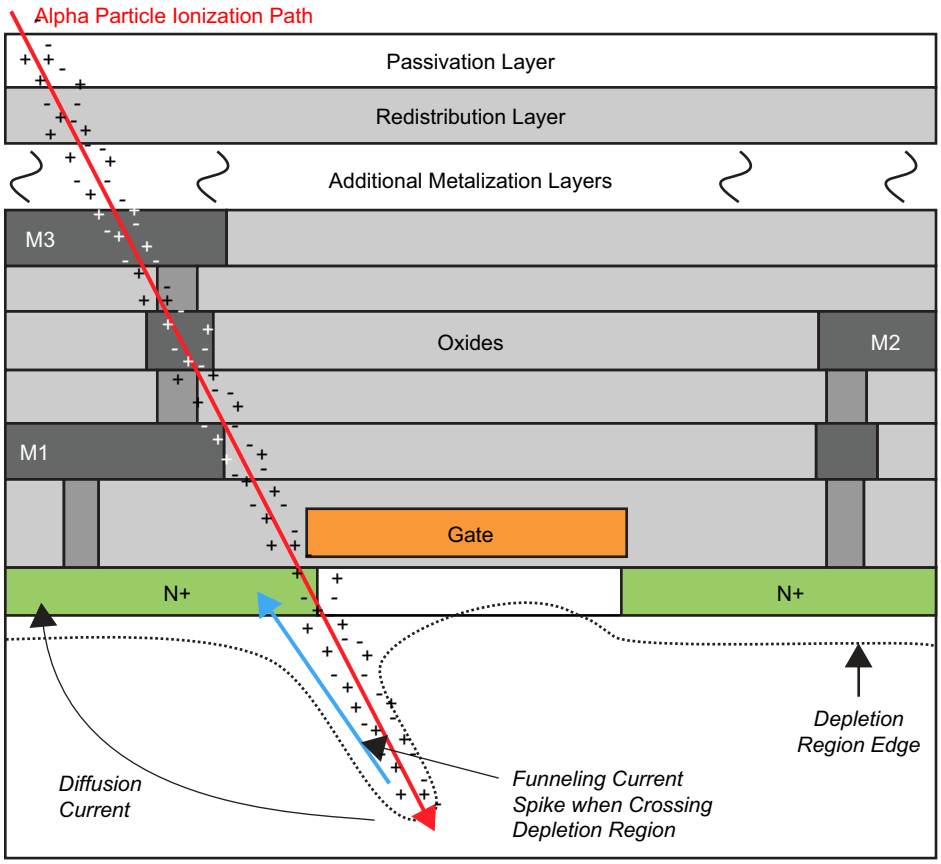

Alpha Particle: Helium Nucleus (2 Neutrons & 2 Protons)
Need a Charged Particle to Cause Ionization

As the charged particle goes through the various material layers in a semiconductor device, it creates an ionization path with free electrons and holes. In many layers, these electron hole pairs recombine back to a normal state. However, when the path crosses the depletion region underneath a drain-gate-source region the electrons it creates can be quickly attracted to a higher voltage NMOS drain diffusion, which sometimes results in the change of state of a storage element. Similarly, for PMOS transistors, the holes can be quickly attracted to a lower voltage PMOS diffusion. Even if the ionization path is nearby, the diffusion path of electrons or holes could also result in a storage element upset as well.

So if it takes a charged particle to create this ionization path, then how can a neutron also cause an upset? The answer is shown in Figure 4:

Figure 4. Neutron Strike

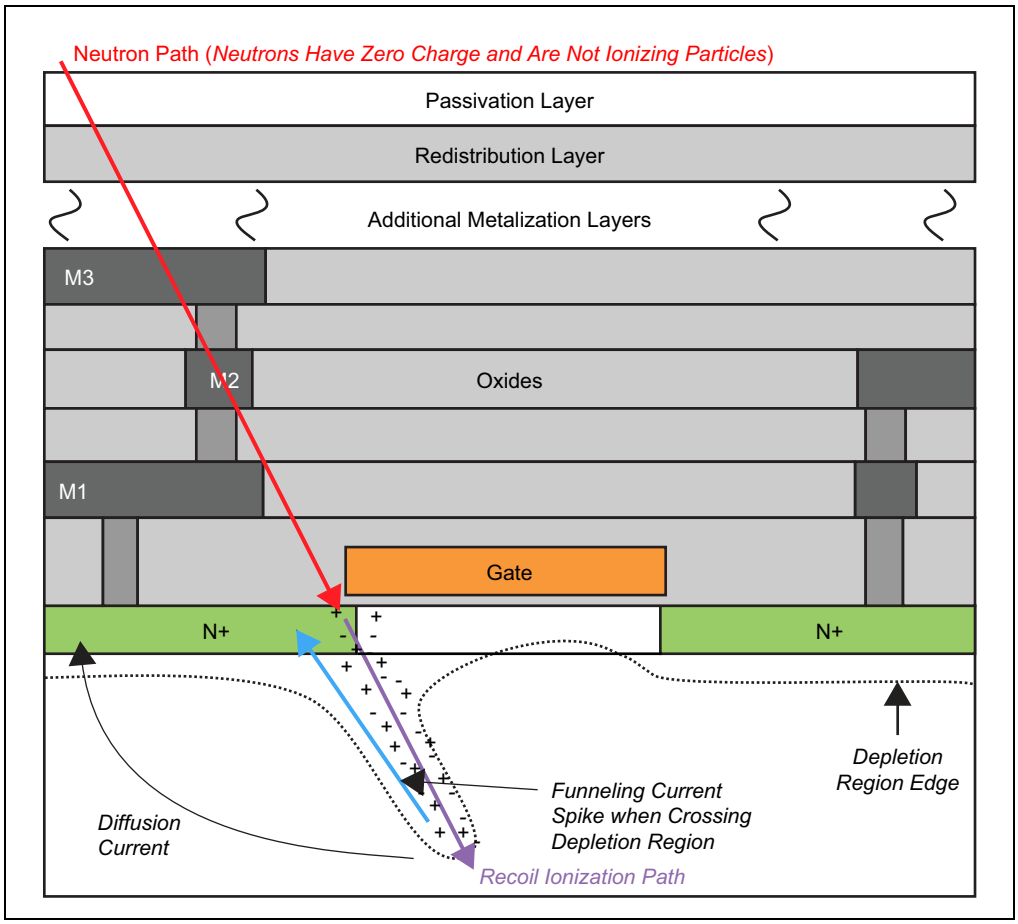

When the neutron collides with another molecule, it can create ions or charged particles with enough speed to create the same ionization path effect as shown in Figure 4. Again, when the ion path crosses the depletion region underneath a drain-gate-source region or is near a p-n junction, the electrons it creates can be attracted to a higher voltage NMOS drain diffusion sometimes resulting in the change of state of a storage element. In the case of PMOS transistors, holes can be attracted to a lower voltage also causing the change of state of a storage element.

FPGAs and SEU

The following sections provide information on one of the drivers for increased FPGA usage, FPGA building blocks and related SEU effects, and how SEU mitigation is a combination of the user designs usage of device capabilities as well as the mitigation techniques used in the design.

FPGAs, Applications, and SEU

The use of FPGAs in system design continues to increase driven by the dramatic cost increase of producing ASICs and ASSPs. Using an ASSP as an example, many companies spend approximately 20% of revenue on development of next-generation products. If the total cost of development of a deep sub-micron ASSP is $100M, then the expected revenue for this ASSP would be $500M. Assuming an average selling price of $25 for the ASSP, then the number of units would need to be 20 million units. For a market share of 50%, then the overall market size would have to be 40 million units. Given the ever increasing fragmentation of markets, outside of mobile devices, there are few other markets that ship at or above 40 million units. As FPGAs continue to take over more sockets from ASICs and ASSPs, there are additional markets and applications where SEU mitigation is important. As previously mentioned, the majority of markets and applications do not need to be concerned about SEU effects.

FPGA Building Blocks and Related SEU Effects

Here we will cover the basic building blocks of an FPGA and how they might be affected by an SEE. FPGAs are comprised of:

- I/O, transceivers, and associated configuration registers
- Hard IP—for example, PCI Express® (PCIe®) hard IP
- Embedded SRAMs—for example, M20K blocks
- Programmable logic fabric and registers
- Configuration RAM (CRAM)

The likelihood of SETs increases with frequency, where SETs don't become a significant contributor to the overall FIT rate until the design frequency reaches the multi-giga-hertz range. Given that the highest design frequencies of FPGAs are typically less than 2 GHz, the impact of SETs on the overall FIT rate is negligible. From a SEU standpoint, the key items to evaluate or mitigate are the configuration registers, the embedded SRAM, the core registers, and the configuration RAM (CRAM). Using the largest Stratix® V device as an example, the number of core registers is 1.4M, the embedded SRAM is 52 Mb, and the configuration RAM is around 250 Mb. As can be seen from these numbers, the first areas of focus are on the embedded SRAM and the CRAM. For designs where SEU mitigation is important, the SRAM FIT rate can be reduced to essentially zero using error correction codes (ECC) or sometimes called error checking and correcting. Suppliers typically have SEU test data using accelerated neutron testing to show how their SRAM failures, even including multi-bit errors, can be corrected with ECC. Similarly, the CRAM errors can be detected using cyclic redundancy check (CRC) and corrected using a process called scrubbing. There are several options and capabilities provided by different suppliers to determine if the CRAM upset is critical to your design – or not.

SEU Reduction by Design

The SEFI ratio for a programmable device is impacted by both the device and the user design implemented in the device. For designs which need to ensure a very low level of functional interrupts in the system, there are many device options and design techniques that can be used. The following sections provide an overview of the device related mitigation capability and some user design mitigation options.

Silicon Design Techniques to Reduce SEU FIT Rates

As mentioned in the previous paragraph, two key focus areas are SRAM and CRAM. For the SRAM, special interleaving and layout techniques are used to ensure that no more than double-bit errors are created from a single strike into a logical word. This is coupled with double error correction and triple error detection (DECTED) to ensure that all errors are corrected and if a case of a triple-bit error in a logical word was ever encountered, it would be detected. For Altera® devices, data is available to show that no errors larger than double-bit errors in a logical word have ever been seen and by using DECTED these soft errors are corrected.

The CRAM also uses proprietary layout techniques, coupled with increased voltage to increase the critical charge, to ensure a lower FIT rate by design. Error-detect cyclical redundancy checking (EDCRC) is used to continuously compare the contents of the CRAM to an expected syndrome result. In Altera devices, if any errors are detected, an EDCRC error is asserted and customers can control which actions are taken. Scrubbing is Altera's term for describing the capability to automatically correct the CRAM and determine the criticality of the SEU event.

Packaging Techniques to Reduce Alpha Particles

Products supplied today by Altera use ultra-low alpha packaging materials to ensure the lowest possible alpha flux. Most suppliers today use the same ultra-low alpha packaging materials.

SEU Mitigation by User Design

There are many techniques used by design teams today for SEU mitigation at the design level. These techniques include ECC, CRC, and Triple Modular Redundancy (TMR). ECC is typically used on user memories and is typically available as hard IP, soft IP, or both in today's FPGAs. As mentioned previously, the use of ECC on the M20K blocks reduces the FIT rate for SRAM effectively to zero. CRC is often used as well. One example is using end-to-end CRC checking for packet processing designs. This method helps ensure that if any SEU event causes a modification to the contents of the packet, it will be caught and a packet re-transmission is requested. Another option commonly used for critical logic is TMR. An example of TMR is triplicating state machines with voters to ensure that if an SEU changes state of one of the state machine flip flops then the design will not be affected. To ensure continued operation, the design needs to detect the inconsistency of one out of the three instances and get that design instance back into the same state as the other two.

IP Support

FPGA providers offer several features and different selections are made by different suppliers as to which are embedded hard IP features and which are soft IP features. These features included the EDCRC or scrubbing capability described previously as well as other features such as hierarchical tagging and sensitivity processing.

For additional information on these mitigation topics, refer to the *Enhancing Robust SEU Mitigation with 28-nm FPGAs* White Paper.

Supporting Quartus II Software Capabilities

The Quartus® II software supports fault injection using either a graphical user interface (GUI) or via scripting. The GUI is often useful to ensure that the initial design and testing is working properly. Converting to script-based testing then allows more exhaustive testing of the design. Because fault injection modifies the contents of the CRAM, it is used for emulating SEU events and subsequently testing the user design response to the SEU. If scrubbing is enabled, it will test the scrubbing capability as well as any associated user design options such as hierarchical tagging and/or sensitivity process.

Conclusion

A SEU is the change in state of a storage element inside of a device or system. This state change is a soft error and can often be fixed by changing the state of the storage element back to its original value. The main causes of SEE are from radioactive decay of the packaging materials (alpha particles) or high-speed neutrons colliding with silicon atoms creating secondary particles, which then create an ionization track where the electrons or holes can get collected on the source or drain of a transistor, causing the soft error. In FPGAs there is SEU mitigation through process selection, silicon design and layout techniques, specialized built-in IP, specialized soft IP, and user design techniques. Additionally, there are software capabilities to implement and test designs that include SEU mitigation. Altera's devices are built from the ground up to have robust SEU tolerance and the associated IP and software capability help designers create and test their design to ensure high levels of SEU tolerance.

For additional information on SEE and mitigation techniques, read the Soft Errors in Modern Electronic Systems book by Michael Nicolaidis.

Further Information

1. JEDEC89A: www.jedec.org
2. Soft Errors in Modern Electronic Systems by Michael Nicolaidis
3. White Paper: *Understanding SEFI (Single Event Functional Interrupt) in FPGA Design* http://www.altera.com/literature/wp/wp-01207-single-event-functional-interrupt.pdf
4. White Paper: *Enhancing Robust SEU Mitigation with 28 nm FPGAs*: http://www.altera.com/literature/wp/wp-01135-stxv-seu-mitigation.pdf

Acknowledgements

- Michael Sydow, Sr. Product Marketing Manager, High-End Products, Altera Corporation

Document Revision History

Table 1 shows the revision history for this document.

Table 1. Document Revision History

Date	Version	Changes
September 2013	1.0	Initial release.