



WP395 (v1.1) May 19, 2015

Mitigating Single-Event Upsets

By: Jameel Hussein and Gary Swift

Occasionally, electronic devices exhibit erroneous behavior for no apparent reason. Through careful experimental design and statistical analysis, scientists and engineers discovered that normal background radiation is often the cause. These failures are generally rare and, thus, easily ignored for common applications. However, for high-assurance applications, it is important to consider the role of background radiation in system reliability. Reliability problems due to radiation most commonly fall into the category termed single-event effects (SEE) and show up as a type of soft errors called single-event upsets (SEUs).

Over the last decade, Xilinx has made significant investments in research and testing to offer customers a full range of SEU mitigation options as well as the lowest intrinsic FIT rates in the industry.

Introduction

SEUs are soft errors (non-permanent errors) caused by secondary particles liberated by the collision of a neutron with a silicon atom or from a contaminant emitting an alpha particle in an electronic device. The neutrons are generated when cosmic rays and protons from space interact with the atmosphere. The cosmic rays are from both inside (the sun) and outside (novas and supernovas) of the solar system. The neutrons range in energy from below 1 million electron volts (MeV) to more than 1,000 MeV. Electronic equipment could be built with shielding to protect against neutrons with this amount of energy—but this is not practical for most applications because the amount of material required to make the shield is prohibitive, e.g., as much as 30 meters of water for neutrons with high energy.

In addition to neutron effects, electronic devices are also vulnerable to alpha particles emitted by natural radioactive isotopes present in device materials and packaging.

The major circuit effects from neutrons and alpha particles include: transient current pulses, change in memory values (bit flips or SEUs), and latch-up. Of these effects, only latch-up is potentially destructive because the deposition of charge leads to a persistent short-circuit, which can result in severe overheating, melting, or vaporization. To prevent latch-up, Xilinx FPGAs follow proprietary design rules that specifically mitigate this effect and are tested for standard latch-up immunity [Ref 1] before they can be production qualified. The transient effects either dissipate quickly (~100 ps) in the circuit or get clocked into a storage element and show up as an indirect SEU. Indirect SEUs only become a concern in extreme radiation environments, such as inter-planetary space, and are essentially negligible in Xilinx FPGAs for commercial applications.

Regardless of origin, when a high energy neutron passes through the device and strikes an atom in the silicon substrate, a reaction occurs. Energetic and positively charged reaction products liberate a cloud of electron hole pairs. These create a single-event transient (SET) in the circuit, which can propagate through combinational logic. If the SET has enough amplitude and duration and coincides with a clock, then an incorrect value is registered and the SET becomes an SEU. Figure 1 shows an example of this; on the left, a neutron-induced reaction product liberating charge from the silicon nucleus that, on the right, disturbs the circuitry and changes the stored value. Only a fraction of SETs become SEUs, which can then manifest as system errors. The probability that an SET will become an SEU is lower in an FPGA than in an ASIC or ASSP because of the Xilinx foundry-independent design rules and certain inherent features related to FPGA configurability compared to standard devices like memory array.

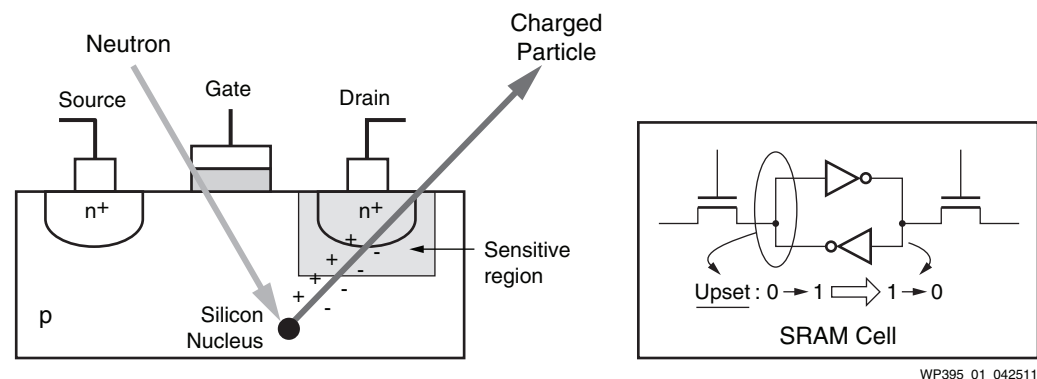


Figure 1: An SEU Occurs in Electronic Memory

In reconfigurable FPGAs, bit flips to the configuration memory and to the block RAM are rare but they occur often enough that they cannot be ignored in high reliability or high availability applications. Direct SEUs in the user registers are negligible.

Due to the programmability of FPGAs, the resources are loaded with capacitive structures, therefore, they have a large charge and tolerate a much larger percentage of particle interactions without upsetting. The memory elements and logic trees in ASICs and ASSPs lack the capacitive load provided by FPGA configuration memory; thus, combining the greater susceptibility of the memory elements and increased propagation of transients through the logic results in a far greater susceptibility to transient faults in ASICs and ASSPs. Moreover, this problem is exacerbated at higher clock speeds.

Any Failure Affects Reliability

Although Xilinx FPGAs are intrinsically reliable when it comes to hard and soft error effects, the probability of soft errors can exceed the probability of hard errors. At some point, the system can be unavailable or fail to function due to an SEU. Certain steps must be taken during the design implementation to mitigate soft errors. Xilinx FPGAs provide designers a number of unique solutions for soft error mitigation.

The designer must first consider the level of reliability and availability that the design requires. The range of reliability requirements covers many orders of magnitude depending on the class of product and its intended market. Availability or recovery from an SEU is also a factor. The highest levels of reliability and availability require redundancy and duplication of hardware functionality because the hard failure rate (wear-out, end of life) becomes significant in safety critical applications.

This white paper separates the levels of availability and levels of reliability:

- Availability: The industry-standard measure of up-time in terms of a percentage (e.g., 99.9995%)
- Reliability: Soft or hard failures in time (FIT) per billion hours, e.g., 1 FIT is equal to one failure in 10^9 hours

Different industries have varying requirements. For example, a communications hub can require:

- Less than 1,000 FIT (114 years mean time between operational interruption) and 99.9995% availability (less than 2.7 minutes a year) for a trunk or communications hub
- Less than 10,000 FIT (11.4 years mean time between failures (MTBF)), and 99% available (3.7 days total outage time per year)

Other reliability and availability requirements can be stated in terms of mean time to fail (in years), probability of failure, or consequences of a failure. Failure from a soft error typically only means a temporary interruption to operation.

For example, safety critical systems might be rated at $1e^{-5}$ probability of a dangerous failure (i.e., can affect human life or cause an environmental catastrophe) for a low-demand system (i.e., occasionally used) and $1e^{-9}$ probability of a dangerous failure for a high-demand system (continuously used). The $1e^{-9}$ probability of failure is a difficult specification to meet and probably requires hardware redundancy. The hardware itself can fail or wear out before a soft error effects an application. If a hard failure requires redundancy, then a soft failure is easier to mitigate because of the ability to switchover to a redundant unit if an SEU is detected.

After designers determine the acceptable failure rates and availability requirements, they must determine how reliable their system components are. Extensive research and testing is required to obtain this information. Xilinx has a long history of studying, testing, and mitigating SEUs to provide designers with accurate reliability and availability data.

History

In the early 1970s, it was discovered that alpha particles from radioactive decay from gross contamination in package materials can also cause errors in electronic devices [Ref 2]. At roughly the same time, on-orbit satellite anomalies were observed [Ref 3]. The same cosmic rays that interact with the atmosphere to create neutrons were studied. As a result, radiation-hardened space electronics became standard for space missions.

At first, the errors were so rare that the study of SEUs was an academic curiosity. Then as circuit dimensions got smaller and voltages were reduced with process shrinks, the stored charge at a node became smaller and smaller and the error rate became more significant.

Today, relative to the traditional failure mechanisms, soft error rates dominate, e.g., aging, electromigration, hot electron affects, and Negative-Bias Temperature Instability (NBTI). This is easily demonstrated when comparing the charge stored in a memory cell to the charge deposited by a neutron reaction product or an alpha particle. At 28 nm, the stored charge is less than 1 femto-coulomb ($1e^{-15}$ coulomb). Neutron reaction products can deposit up to 150 femto-coulomb, so upsets in the cell can be common if nothing is done in the design to protect the cell from upsetting.

Starting with Virtex®-II FPGAs, which had a soft error rate of 405 FIT/Mb (meaning 10^6 configuration bits), Xilinx embarked on a program to ensure that the device failure rate was kept low even in spite of the industry process moving toward higher failure rates. With Virtex-6 FPGAs, the failure rate is 160 FIT/Mb. The 7 series FPGAs have been tested and show a rate of about 100 FIT/Mb. See Figure 2. For the latest data, see UG116, *Device Reliability Report* [Ref 4].




Figure 2: Xilinx FPGA Failure Rate by Product

The failure rate improvements are a result of years of research and design work. One of the most demanding natural radiation environments for electronics is outer space. Lessons learned from Xilinx space products are used to improve commercial products as well. Recognizing Xilinx's success in space, the Air Force Research Labs funded an effort to produce an FPGA for space use that is more resistant to upsets by several orders of magnitude. Not only is that product shipping today, but the data and knowledge obtained is invaluable. This knowledge has been and continues to be leveraged in Xilinx's commercial products and translates into lower FIT rates generation after generation. In addition, Xilinx has more than 40 patents that focus on soft error mitigation and memory cell construction. For more information, visit: <http://www.xilinx.com/applications/aerospace-and-defense/space/index.htm>.

FPGAs vs. ASICs and ASSPs

At 65 nm and beyond, ASICs and ASSPs exhibit significant soft error rates. The SEU problem is worse with each process node; at 28 nm, the upset rate of an ASIC is dominated by single event transients that propagate through the logic. The soft error upset rate (bit flip) has also steadily increased as the voltage drops and dimensions shrink. ASIC manufacturers often require a signed non-disclosure agreement before sharing an estimate of the ASIC's error rate. They can only share an estimate because testing every ASIC is expensive and, therefore, simply not performed. It also takes a significant amount of time to perform real atmospheric testing as opposed to accelerated beam testing, which does not correlate to real life results without large margins of error. Xilinx proactively provides SEU-related information to customers and helps the customers to determine a failure rate (<10% margin of error) and to optimize the best mitigation solution for their application.

Testing

In addition to the design work and extensive research that has been done on SEUs, Xilinx also performs the testing outlined in JESD89A [Ref 1]. This testing includes neutron beam testing at Los Alamos National Labs and actual working arrays of 100 parts each, located on mountain tops as well as underground. This real world testing is called the Rosetta Program; see WP286, *Continuing Experiments of Atmospheric Neutron Effects on Deep Submicron Integrated Circuits* [Ref 5] for more details. Whether using accelerated beam testing or real-time atmospheric testing, the results accumulated over many process nodes are invaluable, and the knowledge gained is applied to the designing of each new generation of FPGAs.

The testing is also part of production qualification. Products cannot be production qualified until they meet their soft error goals for configuration and block RAM memories. Products must also be tested in proton and/or neutron beams to demonstrate latch-up immunity before they can be production qualified. A device that has any history or susceptibility to latch-up is a serious risk not only for the device, but for the entire system that incorporated it. In all the years and generations of the real world and beam testing, no Xilinx FPGA has ever experienced latch-up, and Xilinx does not qualify any FPGA for production until it is proven to not to exhibit latch-up in proton or neutron beam testing.

Construction

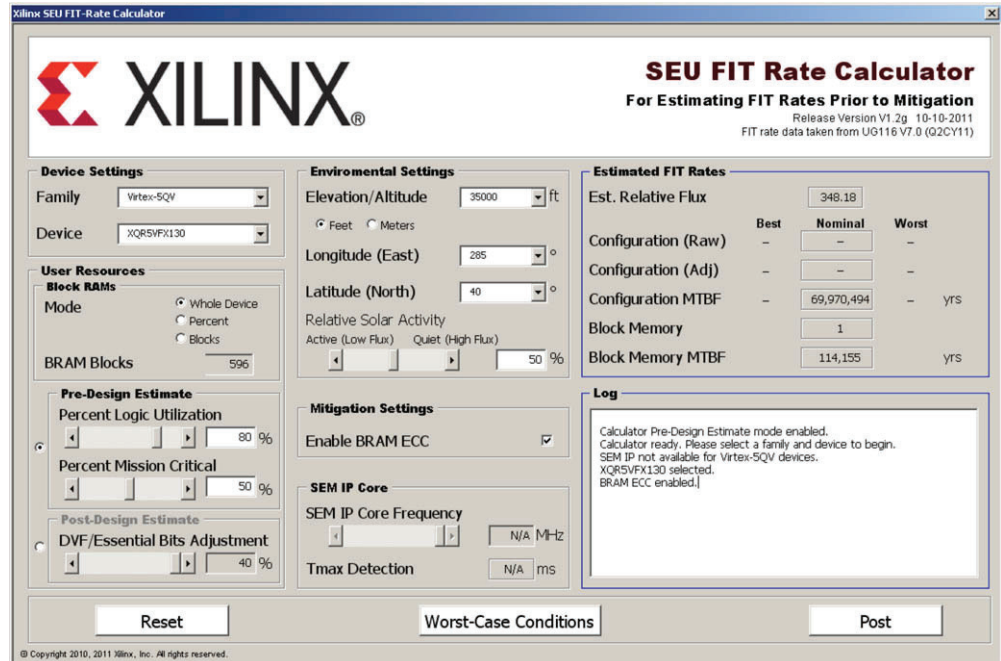
In addition to neutrons, soft errors can also be caused by alpha particles from impurities in the package material or the silicon device itself. Tiny amounts of contaminants in the package emit alpha particles that can have similar effects on static memory as neutron reaction products. These alpha particles can cause bits to flip in the device from 0 to 1 or from 1 to 0 by depositing charge in the active circuit nodes. To reduce the effects of alpha particles, Xilinx requires the use of ultra-low alpha package materials. If these materials are not used, the FIT rates can be hundreds, if not thousands, of times higher. Constant real-time monitoring of production FPGAs assures the ultra-low alpha specification is actually being met in practice.

The Data

Xilinx's SEU test data is updated and published quarterly in UG116, *Device Reliability Report* [Ref 4]. Currently, Xilinx is the *only* company in the industry to openly publish this type of reliability data. Xilinx has been committed to publishing SEU data and working with customers to solve SEU challenges for over 10 years.

Xilinx also provides a FIT Rate Calculator tool that helps customers to predict the FIT rate for their targeted device. See Figure 3. The tool also takes input to fine tune the rate based on a customer's application deployment; parameters such as longitude latitude and altitude have significant effects on the overall FIT rate. For more information or to download the calculator, go to:

<http://www.xilinx.com/member/avionics>.



WP396_03_021512

Figure 3: SEU Fit Rate Calculator

Mitigation Techniques

After customers establish their estimated FIT rate and their reliability requirements, they might need to determine their options for further SEU mitigation when using an FPGA. The first and most common mitigation "technique" is to do nothing.

As strange as it might seem, most designs require no mitigation. Xilinx FPGA intrinsic FIT rates are extremely low. Any Xilinx device that is power cycled daily, weekly, or even monthly, is very unlikely to have an upset that affects operation in the window between resets. Power cycling the device reconfigures it and essentially fixes any upsets that occurred. Many line-side communication applications fit into this category.

Log Errors

If choosing intrinsically low FIT devices does not quite achieve the required reliability and/or availability, the user needs to consider a "detect and log" approach. The idea is to log and fix the errors without any system level actions or resets. The design can easily keep a log of errors using the built-in cyclic redundancy check (CRC) and the error correction code (ECC) in the FPGA (described in UG470, *7 Series FPGAs Configuration User Guide* [Ref 6]). Once enabled, the FPGA continuously reads through the FPGA and checks for any bit flips without interrupting the running design. If a bit flip is detected, an error flag can be recorded with a time stamp to allow customers to check the log for an operational interrupt in the system and determine if it coincided with a logged soft error. Since the bit flips are then automatically corrected by the ECC, the residence time of an upset is short, and the system is more likely to continue and operate error free.

Soft IP and Error Classification

In reality, designs only use a small fraction of the total resources in an FPGA. No design ever uses every capability of every element in the entire device. Thus, an SEU probably will not affect the functionality of the design. Xilinx data, which has been confirmed by independent customer testing, shows that typically less than 30% of all configuration bits are used in any one design. However, even with high resource utilization, the actual rate of SEU impacting design operation is often less than 10% and is typically less than 5%.

That means that out of the total number of upsets, only a small percentage requires a system-level response.

Xilinx has software tools to help determine essential bits, the configuration bits that might cause a specific customer design to malfunction. In association with the LogiCORE™ IP Soft Error Mitigation (SEM) IP core (see [Figure 4](#)), Xilinx can provide a mask file of bits that are essential to the hardware function. When an upset occurs, the IP not only corrects the error but also determines if the upset can affect any resource that the design depends on. This process is called error classification. Based on the classification, the system can decide to keep operating normally with no reset if the upset bit is not essential. The IP performs the correction and classification in the background so that the system does not need to stop, increasing availability considerably. For more information on essential bits and the SEM IP core, see PG036, *LogiCORE IP Soft Error Mitigation Controller User Guide* [Ref 7].

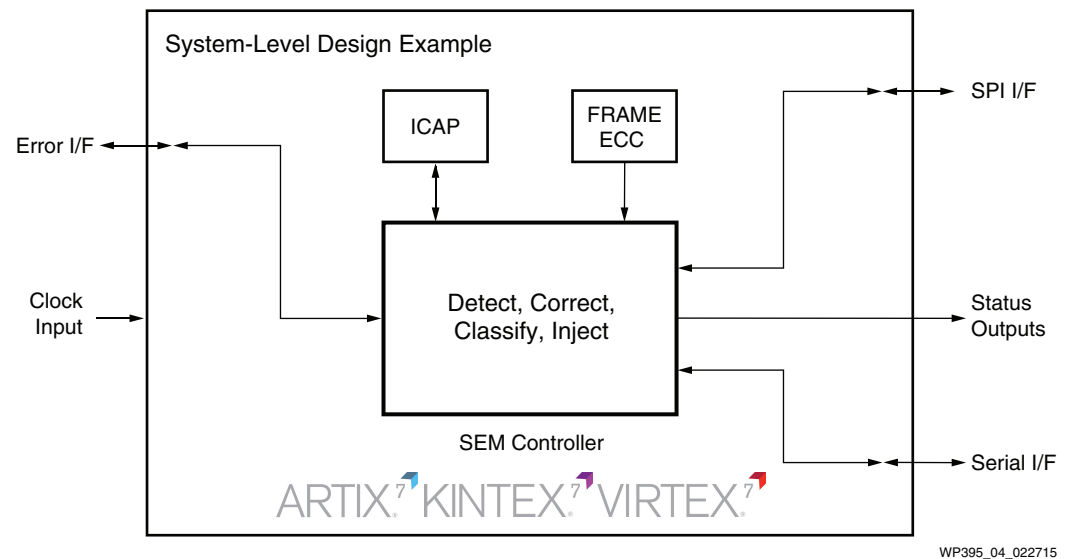


Figure 4: Soft Error Mitigation Controller

SEU Detection and Correction

The ability to quickly determine that a bit has flipped is key to high reliability and high availability applications. Starting with Virtex-4 devices, Xilinx FPGAs have had the built-in capability to detect errors. But correcting the errors is also very important; Virtex-6 and 7 series FPGAs have the built-in ability to do both.

As dimensions have shrunk, a single particle can now upset more than one bit; this is known as multiple bit upset (MBU). To reduce the effects of MBUs, configuration memory is interleaved such that physically adjacent errors are separated in the memory map. This makes them appear as separate single upsets thus enabling them to be repaired by the built-in error correction. The built-in correction feature in Virtex-6 and 7 series FPGAs can be easily activated in the constraints file (see UG360, *Virtex-6 FPGA Configuration User Guide* [Ref 8] and UG470, *7 Series FPGAs Configuration User Guide* [Ref 6]). The correction feature can fix all single-bit errors and almost all of the double-bit errors because MBUs span across frames. In the extremely rare event that an MBU is not correctable, the CRC checker, also built into the FPGA, detects 100% of MBUs. Combining these features with the SEM IP core offers customers complete detection and correction coverage.

Block RAM is also susceptible to SEUs. Bit flips in block RAM can be mitigated with built-in error correction code (ECC) that uses a Hamming code for a 64-bit encoder/decoder. The ECC combined with physical interleaving reduces the FIT rate for user data well below the hard failure rate that is essentially reduced to 0.

Typically, neither ASICs nor ASSPs have the ability to detect or correct errors. With no knowledge that a soft error occurred, an ASIC or ASSP design can be silently malfunctioning. The CRC check in Xilinx FPGAs confirms that the design is running without error. If the user must know that a system is error free (e.g., high assurance applications), then Xilinx FPGAs are a much better solution.

Xilinx can also significantly improve availability with error classification. This allows a system to reset only when an SEU might affect the design, and to ignore all other upsets, greatly reducing downtime and raising availability. Because correction and classification takes only a few hundred microseconds, many applications meet their requirements by taking advantage of these features.

For the highest levels of reliability, triplication and voting within the design can be required. This does not prevent upsets but allows the design to get through the period of detection and correction without interruption to error-free operation.

For some space applications, where failure is not an option, hardware triplication (three of everything) and external voters are required. Xilinx has many FPGA devices in space, including the Mars Rovers [Ref 9] and packages aboard the International Space Station [Ref 10]. Xilinx also has devices in other high reliability applications, such as in fighter jets, and in high-end luxury automobiles.

Summary

SEU mitigation can be very challenging if it has to be done without the benefit of mitigation IP and tools [Ref 11]. Xilinx provides engineering and tools that do most of the work so customers do not have to. Xilinx performs extensive testing and provides the results of these comprehensive atmospheric and beam tests. Customers can use these results to estimate the FIT rates for their design and accurately decide on the optimum mitigation solution. Starting with the FPGA silicon itself, which is designed for low SEU sensitivity, Xilinx offers tiered solutions to the SEU problem. These solutions include innovative IP and tools to predict rates and mitigate effects. The built-in features are simple and easy to use for both error detection and correction. In addition, Xilinx provides fully verified soft IP that can be used for further FIT rate reduction and system testing.

For more information, go to:

<http://www.xilinx.com/products/quality/single-event-upsets.htm>.

References

1. JESD89A, *Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices*
<http://www.jedec.org/sites/default/files/docs/jesd89a.pdf>
2. T. C. May and M. H. Woods, "Alpha-Particle-Induced Soft Errors in Dynamic Memories," *IEEE Trans. Electron Dev.*, ED-26, 2 (1979).
3. D. Binder, E. C. Smith, and A. B. Holman, "Satellite Anomalies from Galactic Cosmic Rays," *IEEE Trans. Nucl. Sci.*, NS-22, 2675 (1975).
4. [UG116](#), *Device Reliability Report*
5. [WP286](#), *Continuing Experiments of Atmospheric Neutron Effects on Deep Submicron Integrated Circuits*
6. [UG470](#), *7 Series FPGAs Configuration User Guide*
7. [PG036](#), *LogiCORE™ IP Soft Error Mitigation Controller User Guide*
8. [UG360](#), *Virtex-6 FPGA Configuration User Guide*
9. Xilinx Press Release
http://www.xilinx.com/prs_rls/design_win/0412_marsrover.htm
10. NASA Presentation
https://nepp.nasa.gov/mapld_2008/presentations/t/08%20-%20Blansett_Ethan_mapld_08_pres_1.pdf
11. [WP402](#), *Considerations Surrounding Single Event Effects in FPGAs, ASICs, and Processors*

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
05/19/2015	1.1	Updated Figure 4, SEU Detection and Correction , and References .
04/09/2012	1.0	Initial Xilinx release.

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 Xilinx’s limited warranty, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; 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 such critical applications, please refer to Xilinx’s Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS.