Understanding Soft and Firm Errors in Semiconductor Devices

Questions and Answers

December 2002

1) What are soft errors?

A soft error is a "glitch" in a semiconductor device. These glitches are random, usually not catastrophic, and normally do not destroy the device. They are caused by external elements outside of the designer's control. Many systems can tolerate some level of soft errors. For example, in a video application, soft errors can manifest themselves as missing or incorrectly colored bits on a display screen. These errors may or may not be noticeable or important to the user.

2) What causes soft errors?

Soft errors are caused by a charged particle striking a semiconductor memory or a memory-type element. Specifically, the charge (electron-hole pairs) generated by the interaction of an energetic charged particle with the semiconductor atoms corrupts the stored information in the memory cell. These charged particles can come directly from radioactive materials and cosmic rays or indirectly as a result of high-energy particle interaction with the semiconductor itself.

High-energy cosmic rays and solar particles react with the upper atmosphere generating high-energy protons and neutrons that shower to the ground. Neutrons are particularly troublesome as they can penetrate most man-made construction (a neutron can easily pass through five feet of concrete). This effect varies with both latitude and altitude. In London, the effect is two times worse than on the equator. In Denver, with its high altitude, the effect is three times worse than at sea-level San Francisco. In a commercial airplane, the effect can be 100-800 times worse than at sea-level.

Another common source of these errors are alpha particles, which are emitted by the trace amount of radioactive isotopes present in the packaging materials of integrated circuits. Bump materials used in the new flip-chip packaging technique have also been recently identified as containing significant alpha particle sources.

3) Is this a new problem?

Initially, the soft error problem gained widespread attention in the late 1970s as a memory data corruption issue, when DRAMs began to show signs of apparently random failures. Although the phenomenon was first noticed in DRAMs, SRAM memories and SRAM-based programmable logic devices are also subject to the same effects.

Unlike capacitor-based DRAMs, SRAMs are constructed of cross-coupled devices, which have far less capacitance in each cell. The lower the capacitance of a cell, the greater the chance of an upset. As both the voltage and cell size are reduced with each new process generation, the SRAM cell capacitance continues to decrease, making the cell even more vulnerable to more types of (lower energy) particles.

4) What are firm errors?

Although the physical phenomenon is often referred to as a soft error or as the soft error rate (SER), strictly speaking, this term only applies to memory elements used for data storage. An error in a memory element is considered soft because it corrupts the data. This same type of radiation induced error in an FPGA is a "firm" error, because it is not just a transient data error. When a firm error occurs, the data is not corrupted; it is the device's configuration or "personality" that is affected. The error changes the actual function of the device. There are no soft errors in an SRAM FPGA configuration memory; they are firm errors, and they can have serious system consequences.

5) Are radiation affects at ground-level just a theoretical problem?

No, in 2000, Sun's UltraSPARC II workstations were crashing at an alarming rate. The inability to initially locate the source of the problem created significant customer dissatisfaction issues for Sun. The root cause of the problem was finally traced to IBM supplied SRAMs that were experiencing high upset rates due to charged particles causing soft errors in the memory system. Ultimately, not only did Sun switch memory vendors, they also designed new error checking and correcting logic and implemented it across the entire cache architecture.

6) Have firm errors been reported in SRAM-based FPGAs?

Recent independent tests conducted by the European Space Agency (ESA) found two basic classes of firm errors in the Xilinx SRAM-based XQVR300 part:

1. Routing errors: Occur when the routing bits and configuration look-up tables are corrupted. This firm error results in a continuous functional error until new configuration data is reloaded. It can take a great number of clock cycles before this configuration loss is detected and recovery actions are initiated. During this time, the error can propagate to the rest of the system.

2. Persistent errors: Occur when the weak keeper circuits within the device are corrupted. This firm error cannot be corrected by a reconfiguration. The part must be completely reset and reinitialized. This may involve bringing down the entire system to do a complete power-on reset.

Firm errors can create "illegal" conditions within the FPGA. High current conditions are possible due to contentions in a misconfigured SRAM FPGA. This high current draw may damage the device or the board on which it is mounted. If not corrected, firm errors that result in simultaneously enabling pull-ups and pull-downs or serious bus contention may physically damage the FPGA creating a "hard" error.

7) Are there any other incidents with errors due to charged particles that have been widely reported?

This is a sensitive issue because most vendors do not like to publicly admit to latent design flaws in their equipment or admit their equipment has potential reliability issues. Due to concerns about these effects, Intel has announced plans to incorporate error-correction in its SRAM-intensive McKinley IA-64 processor, which is scheduled to be completed later this year. The April 2002 International Reliability Physics Symposium (IRPS), held in Dallas, TX, had a special focused session discussing, "Radiation Induced Soft Errors in Silicon Components and Computer Systems." The large semiconductor companies were well represented at this session.

8) Is this a significant concern for ground-based equipment?

Some time ago, IBM demonstrated that at an altitude as low as 10,000 feet, SER effects were already 14 times higher than at sea level because of the greater exposure to cosmic rays. Failures increase very rapidly at higher altitudes. However, today, even ground-based systems are not immune from these errors.

Due to process technology improvements, today's deep sub-micron SRAM-based devices have significantly increased sensitivity to radiation effects and are much more likely to be upset by a passing particle. A decade ago, it would have taken at least 50 femtocoulombs to change the state of a typical SRAM cell. Today, just 10 femtocoulombs are enough to upset an SRAM cell.

Leading telecommunications and networking companies are now specifying qualification tests designed to evaluate radiation resistance of the integrated circuits they are planning to use in their communications systems. JEDEC has published a specification (JESD89) that specifies SER testing methodologies for integrated circuits.

Radiation induced errors are a real problem today with both land-based and airborne equipment. This problem will continue to worsen as devices increase in density and geometries continue to shrink. 130 nanometer and smaller SRAM geometries are particularly sensitive to these problems.

9) How resistant to firm errors are Actel's antifuse-based products?

Firm errors are non-existent in Actel's antifuse-based products. The Actel antifuse has been shown to be immune to both ground and aero particle effects. For space applications, Actel offers the RT54SX-S product family. This family hardens the sequential logic flip-flops to provide a robust solution for use in the space market, where radiation effects are particularly severe.

10) How resistant to firm errors are Actel's Flash-based products?

Firm errors are non-existent in Actel's Flash-based products. The Actel Flash cell has been shown to be immune to ground and aero particle effects.

11) How common are soft and firm errors?

These errors are commonly expressed by "failure-in-time," or FIT rates. At the 0.13-micron process node, some memory technologies have error rates ranging from 10,000 to 100,000 FITs per megabit. This means that typically, in a single device, there will be a random data failure at a frequency ranging from one to ten years. On average, a bank of one hundred, 1-megabit memory devices would exhibit a failure every 3 to 30 days. For some applications, this would be considered noise-level and ignored. In other applications, designers would use error detection and correction techniques to mitigate these effects.

12) How do firm errors affect system reliability?

Firm Errors are a more serious issue to a system designer than data corruption. Once a firm error occurs, that functional error will remain (either 'routing' or 'persistent' in type) until the system is reconfigured or reinitialized, depending on the type of firm error. If this occurs, the system FIT rate essentially becomes infinite.

13) How can firm errors be prevented in SRAM-based programmable devices?

Firm errors are a fact of life in all SRAM-based PLDs; their occurrence cannot be prevented. Tripling each portion of the design and voting out the error may mitigate some of these effects. However, voting cuts the available gates by a factor of 4 or 5 and does not deal with circuits that are static (where voting does not fully protect against soft errors). A system designer can further mitigate firm error effects by adding detect and reconfiguring routines into the system, although this can significantly reduce system availability and adds cost and complexity.

14) What does this mean in the "real world?"

From published radiation testing data, a typical 1 megagate SRAM-based FPGA (XCV1000) would have a ground-level firm error rate of 1200 FITs. Using this FIT rate, it is easy to calculate several real world scenarios:

- A complex system that used 100 of these 1-megagate SRAM-based FPGA devices (or the equivalent number of FPGA gates) would on average exhibit a functional failure every 11 months. This is a best-case number. This same system, if deployed in London or Denver, **would fail every 3 to 7 months**.

- A system of similar complexity, if used in an aircraft or other high altitude environments, would have a firm error rate that is 100-800 times worse, and a functional failure **would occur every 12 to 36 hours**.

- This issue is not confined to just high altitude or complex systems. If a product contains just a single 1 megagate SRAM-based FPGA and has shipped 50,000 units, there is a significant risk of field failures due to firm errors. Even for such a simple system, the manufacturer can expect that within his customer base, there will be a field **failure due to a firm error every 17 hours. That is less than one day between failures!**

- OEMs should consider Sun's experience with errors and carefully evaluate the potential risks involved in ignoring the effects of radiation-induced errors.

Soft and Firm Error Rate Summary

- Designers cannot control the sources of soft or firm errors, but their effects can be mitigated through the use of careful design techniques and by utilizing error resistant programmable products.

- Although the underlying causes are the same, soft errors are transient errors in memories; firm errors are non-transient errors in SRAM based programmable logic devices. Unlike soft errors, firm errors remain until they are detected and corrected.

- FPGAs based on SRAM technology are inherently unsuitable for any avionics applications due to the significantly increased occurrence of firm errors at higher altitudes.

- Be aware of the "hidden costs" of firm errors, which include time lost in supporting and analyzing "random" field failures, reduced system availability and reliability, and reduced levels of customer satisfaction.

- Mission critical applications at ground level are subject to unplanned outages due to firm errors unless preventative measures are taken. The reliability of ground-based systems using SRAM-based FPGAs vary significantly with altitude, latitude, and design complexity.

- While an individual product with just a single FPGA may have a relatively low probability of failure, when this same product ships in large quantities, there is a greatly increased risk of system outages in the field due to firm errors.

- Today's deep sub-micron-based programmable devices are already very susceptible to neutron and alpha particle induced errors. Shrinking geometries are making the problem increasingly worse with each new generation of SRAM-based FPGAs.

- Previous generations of 5Volt CMOS technology had noise margins of a couple of volts, while newer nanometer technologies will have only a few tenths of a volt noise margin. The combination of neutrons from above and locally produced noise will continue to challenge designers in the quest to build reliable systems.

For more information, call **1.888.99.ACTEL** or visit our website at **http://www.actel.com**

www.actel.com

Actel Corporation

955 East Arques Avenue
Sunnyvale, CA USA 94086
Telephone 408.739.1010
Facsimile 408.739.1540

Actel Europe Ltd.

Maxfli Court, Riverside Way
Camberley, Surrey GU15 3YL
United Kingdom
Telephone +44 0 1276.401450
Facsimile +44 0 1276.401490

Actel Japan

EXOS Ebisu Building 4F
1-24-14 Ebisu Shibuya-ku
Tokyo 150, Japan
Telephone +81 0 3.3445.7671
Facsimile +81 0 3.3445.7668