Case Study: Soft Error Rate Analysis in Storage Systems

Brian Mullins, Hossein Asadi, Mehdi B. Tahoori, David Kaeli
Dept. of Electrical & Computer Engineering
Northeastern University
360 Huntington Ave., Boston, MA 02115
{bmullins,gasadi,mtahoori,kaeli} @ece.neu.edu

Kevin Granlund, Rudy Bauer, Scott Romano
EMC Corp., RAS Dept.
Hopkinton, MA 01748
{granlund_kevin,bauer_rudy,romano_scott}@emc.com

Abstract

Soft errors due to cosmic particles are a growing reliability threat for VLSI systems. In this paper we analyze the soft error vulnerability of FPGAs used in storage systems. Since the reliability requirements of these high performance storage subsystems are very stringent, the reliability of the FPGA chips used in the design of such systems plays a critical role in the overall system reliability. We validate the projections produced by our analytical model by using field error rates obtained from actual field failure data of a large FPGA-based design used in the Logical Unit Module board of a commercial storage system. This comparison confirms that the projections obtained from our analytical tool are accurate (there is an 81% overlap in FIT rate range obtained with our analytical modeling framework and the field failure data studied).

1 Introduction

Soft errors are intermittent malfunctions of the hardware that are not reproducible [8, 10]. These errors, also called transient errors, occur more often than permanent errors [6, 11]. *Single Event Upsets* (SEUs) that cause soft errors are generated by cosmic particles, energetic neutrons, and alpha particles hitting the surface of silicon devices [13].

Device scaling significantly affects the susceptibility of integrated circuits to soft errors [12]. As the feature size shrinks, the amount of charge per device decreases thereby enabling a particle strike to be much more likely to cause an error. As a result, particles of lower energy, which are far more plentiful, can generate sufficient charge to cause a soft error. Hence, in the absence of error correction schemes, the system error rate will grow in direct proportion to the number of bits on the chip. Thus, while Moore's Law predicts an exponential increase in the transistor count, this growth comes at the cost of an exponential increase in the error rates for unprotected chips [4, 9].

One of the key design points in storage systems is availability. From the users' stand point, they want to be sure that data is never lost (*reliability*). Also, the application which heavily depends upon that data is required to remain available to access (*availability*). Therefore, a considerable amount of redundancy, error checking (parity), and error correction has been integrated into the design of these systems [5].

Field Programmable Gate Arrays (FPGAs) are commonly used in the design of state-of-the-art storage systems. FPGAs are also frequently used in the implementation of the adapters that interface to either hosts or disk arrays. Typically, FPGA-based designs are more vulnerable to soft errors than *Application Specific Integrated Circuit* (ASIC) implementations [3, 7]. Hence, the soft error reliability of these FPGAs is critical in the overall dependability of storage systems.

To achieve a reasonable balance between reliability and performance, the effect of soft errors at the system level and the contribution of each component to the overall soft error rate of the system need to be precisely analyzed. Tools are needed to be able to identify the most vulnerable components in the system. Given the availability of such a tool and redundancy budget, the most vulnerable components can be protected in the most effective manner (using hardware or software redundancy).

In this paper, we focus on soft error rate estimation of FPGAs used in storage systems designs. We present a case study on failure rate data for a particular FPGA-based con-

25th IEEE VLSI Test Symmposium (VTS'07)
0-7695-2812-0/07 $20.00 © 2007 IEEE

troller. We compare two different SEU rates: one resulting from a comprehensive field failure data analysis and one obtained from our analytical tool. To the best of our knowledge, this is the first attempt to validate a system-level soft error modeling tool using real and comprehensive field failure data.

The rest of the paper is organized as follows: in Section 2, an overview of the FPGA-mapped controller unit used in the design of storage system is presented. In Section 3, the failure field data analysis methodology is described. The results obtained from the analytical model are presented in Section 4. Finally, Section 5 concludes the paper.

2 Embedded Control Unit in Storage Systems

In this study we consider FPGAs that are commonly used in the design of high performance storage systems. These systems typically hold several hundreds disks, which can be protected via RAID protocols (RAID-1, RAID-5). The internal architecture provides for a high degree of redundancy so that a failure in any bus component does not disconnect any component from the system. One of the components that connects to the buses is the *Logical Unit Module* (LUM). Multiple LUMs connect the server host to the internal buses of the disk arrays. These LUMs run the algorithms to manage the memory caches.

Each LUM has several *Embedded Control Units* (ECUs); each ECU controls multiple microprocessors, DRAMs, and L2 caches. All ECUs have been implemented on an SRAM-based FPGAs, and each ECU acts independently. This FPGA component is the target of our study.

The ECU design has been implemented using a very popular and very large commercial FPGA device. Table 1 depicts the FPGA utilization of this device for the ECU design. As seen in this table, the ECU design uses 99.9% of the available slices and 73.4% of the available *Look-up Tables* (LUTs).

3 Field Data Analysis

The goal of this field data analysis is to calculate the FIT rate for one Logical Unit Module (LUM) based solely upon SEUs observed in the field. Specifically, SEUs that occurred on distinct FPGA components (we will refer to these as FSEUs) contained within the specified LUM were investigated. It should be noted that this field analysis was not limited in scope to one particular geographic location; data was collected from *all* functioning systems distributed in more than 35 countries, spanning regions from all across the globe. Thus, the analyzed field data represents a *global* distribution. In total, approximately 12,000 systems and twenty-eight months of field data was analyzed, including more than 750,000 FPGAs. The information needed for this analysis was readily made available via the technology and error reporting systems implemented in the field.

Initially, all errors in the LUM that were flagged in the field as *No Evidence of Failure* (NEOF) were collected and analyzed. NEOFs can describe a range of errors including:

- state-dependent logic or timing errors,
- software-based errors,
- signal crosstalk, and
- particle-induced soft errors (SEUs).

In total, $\sim$ 7900 NEOF events were analyzed.

The next step in this process was to analyze the error code associated with each NEOF event. The storage manufacturer has identified six specific error codes that have been verified in the field to be associated with FSEUs in the investigated FPGA devices. The way that these error codes were identified is as follows: once a trend of errors has been seen in the field, the storage manufacturer actually goes out into the field and scans out the FPGA to compare the bit pattern to the original configuration pattern. When a bit difference is found, the conclusion is that this bit flip was definitively caused by an FSEU. Note that even though an error event may possess one of the six error codes, it still may not have been caused by an FSEU. We will refer to errors possessing these six error codes as *Probable FSEUs*.

An additional three error codes have been identified as being potentially associated with FSEUs in the LUM, though field studies have not confirmed this yet (we will refer to error events possessing these three error codes as *Potential FSEUs*). We will label any NEOF event as a *Possible FSEU* if they possess one of these nine error codes.

Furthermore, when an FSEU is confirmed to have occurred in the field, there is a unique relationship between the observed error and the affected logic within the FPGA design. This is defined as the FSEU's failure signature. Over time, these signatures are documented, detailing both the specific field errors that represent possible FSEUs and their distinct failure signatures.

In general, FSEUs in the LUM have been observed to manifest themselves in several ways such as CRC errors, parity errors, timeout errors, and data miscompares. A specific example of an FSEU is a parity bit being set during an interrupt operation (e.g. a situation that can only arise by the internal logic of the FPGA being incorrect).

After all LUMs that were categorized as NEOF were collected and analyzed, all potential FSEUs are identified according to their specific field error code as described above. This reduces the analysis space by an order of magnitude.

Table 1. FPGA resource utilization information for ECU chip.

Parameter	Slices	IOBs	BRAM bits	Look-up Tables	FFs	Configuration bits
Number Used	12286	631	290816	18033	9725	1.7M
Total	12288	728	393216	24576	24576	6.5M
Usage	99.9%	86.7%	73.9%	73.4%	40.0%	26.2%

After completing this step, the focus of our field study was then reduced in scope to two specific FPGAs within the LUM. This decision was motivated by a high degree of understanding of the effects of FSEUs on these components and the significant role that these FPGAs play within the LUM.

The final step in identifying all FSEUs that occurred within the two FPGAs of interest was to inspect the log file provided with each *Possible FSEU*. Each error's log file contained all of the system information at the time when the error occurred, in addition to information that described how the system behaved both before, and after, the error was observed. In order for an error to be classified as a FSEU, it was required that the error's log file contain the failure signature that specifically pointed to a bit flip in the configuration bits in the FPGA (bit flips can also occur in the non-configuration bits of the FPGA, but in these designs the physical space in both FPGAs is heavily dominated by configuration bits [3]). There are also cases where the signature confirms that the error was not caused by a bit flip in an FPGA, and thus is not an FSEU (i.e., a non-FSEU).

Analyzing the collected data, it was found that a significant percentage of the total cases studied could not be discerned as either an FSEU or non-FSEU. For the majority of the indeterminate cases, our analysis is incomplete due to the lack of an error log file being available, and thus these errors could not be properly analyzed to the detail needed. However, by using the trends observed in the FSEU/non-FSEU errors, we can approximate what portion of the indeterminate cases are FSEUs and non-FSEUs. We multiply the ratio of the FSEUs/(FSEUs + non-FSEUs) to provide us with an estimate of the number of indeterminate cases that are FSEUs.

Table 2 represents all errors that were found to have a *Probable FSEUs* error code. The FSEU category represents the errors whose log file had an FSEU failure signature; the non-FSEU category represents the errors whose log file did not have an FSEU failure signature. The Indeterminate category represents the errors whose log file was not available, and thus whose FSEU error status could not be determined. Table 3 shows this same information for all errors that contained *Potential FSEU* error codes.

The next step in calculating the FIT rate of the targeted FPGA was to determine exactly where each FSEU occurred relative to the LUM. By again referencing each FSEU's field returned error code, information detailing the location of each FSEU could be extracted. Here, the FSEUs were categorized as having occurred on FPGA-1, on FPGA-2, or on either FPGA-1 or FPGA-2. The breakdown is shown in Table 4.

Table 2. Probable FSEU Error Code Statistics

Probable FSEU Error Code Stats	Percentage
FSEU	29.35
Non-FSEU	33.20
Indeterminate	37.45

Table 3. Potential FSEU Error Code Statistics

Potential FSEU Error Code Stats	Percentage
FSEU	14.45
Non-FSEU	40.46
Indeterminate	45.09

Assuming that FPGA-1 is the component of interest, and also that the distribution of FPGA-1 FSEUs within the category of FPGA-1 or FPGA-2 ranges from 20% to 80%, the following calculations result:

If 20% of these errors were to occur on FPGA-1, then the total number of FSEUs on FPGA-1 was calculated by:

$$(0.10 + (0.20 \times 0.31)) \times \text{(The total number of FSEUs)}$$

Similarly, if 80% of these errors were to occur on FPGA-1, then the total number of FSEUs on FPGA-1 was calculated by:

$$(0.10 + (0.80 \times 0.31)) \times \text{(The total number of FSEUs)}$$

The final step of this project was to then calculate an overall FIT rate for the FPGA of interest, in this case FPGA-1. By analyzing the number of FSEUs that occurred on this FPGA in combination with the total number of field hours that this component accrued, a device field FIT rate was obtained. Figure 1 summarizes all of the steps taken in computing this value.

Table 4. FSEU distribution in FPGA modules

FSEU Location	Percentage
FPGA-1	∼10
FPGA-2	∼59
FPGA-1 or FPGA-2	∼31

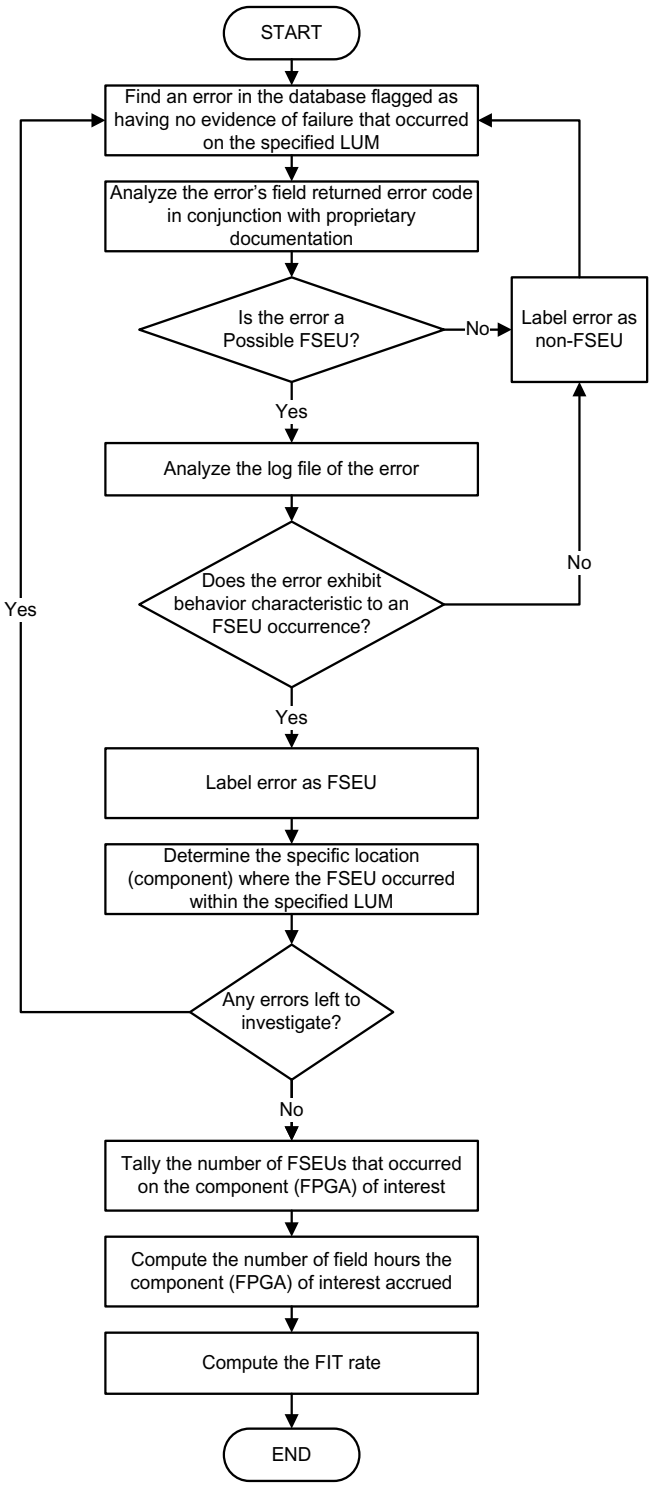

Figure 1. Field data analysis flow.

4 Analytical SER Projections

4.1 Overview of SER Modeling Framework

In order to estimate the soft error rate of a design mapped into an SRAM-based FPGA, we compute the probability of system failure due to bit-flips of FPGA memory bits. Memory elements in an SRAM-based FPGA device can be divided into two categories: 1) *configuration* and 2) *user* bits. Configuration bits are used for specifying the particular circuit mapped into the FPGA, whereas the user bits, such as flip-flops (FFs) or on-chip memory arrays, hold the current state of the circuit. The majority (more than 99%) of the memory bits in an FPGA are configuration bits and therefore, the probability of soft errors in the configuration bits is much greater than that in user bits [3].

The configuration memory bits are further categorized into *sensitive* and *non-sensitive* bits according to their vulnerabilities to SEUs. An SEU in a sensitive configuration bit affects the functionality of the particular circuit mapped into the FPGA; non-sensitive bits act as "don't care" configuration bits for that particular mapped design. Different nodes in the mapped design may consist of different quantities of sensitive bits.

In order to measure the vulnerability of each node of the circuit to SEUs we use two parameters: *Netlist Impact Probability (NIP)* and *Node Error Rate (NER)*. NIP is the probability that the error site is activated by the inputs and then propagated to the outputs. NIP depends on both the error model and the circuit topology (netlist) [1, 2, 3]. NER_i captures the effects of the utilization of specific FPGA resources via the placement and routing information of the FPGA design. NER_i is calculated based on the raw error rate of the device, the error model being used, and the number of sensitive SRAM configuration bits used to implement node i within the FPGA. The raw error rate of an SRAM cell depends on the device characteristics and the flux encountered by the device. Computation of NER_i involves extracting all sensitive configuration bits used in node i.The failure rate due to node i is then computed as the product of NER_i and NIP_i.

The system failure rate for the entire circuit can be computed by summing the calculated failure rates for all nodes. A detailed description and formulation of NER_i and NIP_i has been presented in [3].

4.2 Tool Specification

Figure 2 shows the overall flow of our SER estimation methodology. The inputs to this software tool are the FPGA physical device information, the FPGA interconnect and architecture information, and the raw error rates for different

Table 5. Execution time for ECU SER estimation.

Parameter	Read	Arrange	SP Computation	Extract error sites	SFR calculation	Total Time
Time (second)	108.74	3.85	437.30	50.99	11780.00	**12380.88**
Percentage	0.88%	0.03%	3.53%	0.41%	95.15%	**100.00%**

types of FPGA cells. The program outputs are the overall system derating FIT rate and the FIT rate of all internal nodes. There are four categories of modules in the program including: Physical Device Info, Design Info, Used Resource Info, and System Failure Computation. The only modules that need to be updated for each device family are the Physical Device Info Modules.

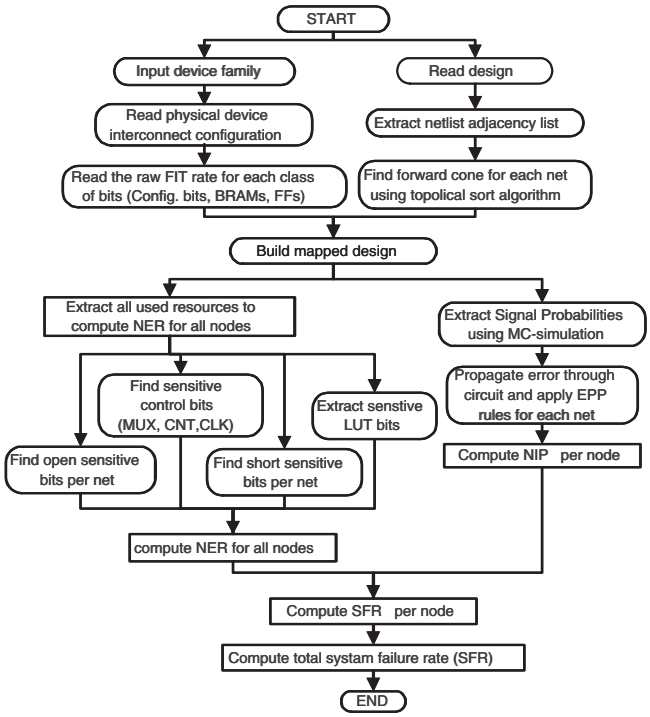

Figure 2. SER estimation flow.

4.3 Results

We have obtained SER projections for the ECU in the storage system using the described analytical framework. Our software tool extracts netlist information from the input design file including lists of used resources, sensitive bits, and error models of the design. The failure rate of all circuit nodes are computed based on the above information. This software tool has been executed on a Sun Blade 900 © workstation equipped with 4GB main memory and running Solaris 9 © operating system.

According to the outputs of our software tool, the configuration routing bits constitute almost 85% of the total sensitive configuration bits; LUT bits and control/clocking bits constitute 11% and 4% of the total configuration bits, respectively. The software tool also shows that the number of FFs is less than 0.6% of the total number of sensitive configuration bits.

The detailed execution time of our SER estimation method is listed in Table 5. The total run time of this SER estimation method includes the time required to read the netlist, arrange the netlist, compute signal probabilities, extract error sites, and compute the system failure rate using the error propagation probabilities for all nodes.

The system failure rate of the ECU has been computed using both our software tool and field data; Figure 3 shows a normalized representation comparing the two. Due to the fact that the field data was extracted from devices from different locations, the raw FIT rate which is used as an input to the predictive tool ranges between two values. Based on the raw FIT rates reported by vendors, our predictive tool reports a FIT rate ranging from 3.8 to 5.9. Furthermore, due to the range in the number of FSEUs in FPGA-1 (considering FSEUs that could have occurred on FPGA-1 or FPGA-2), the computed FIT rate also ranges between two values: from 2.5 to 5.5.

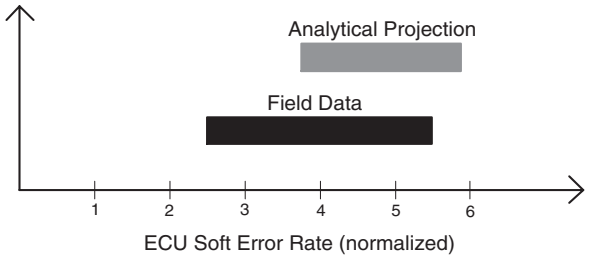

Figure 3. Comparison of predictive tool results with field data.

The slight mismatch between the field data and our analytical projection can be attributed to the fact that it was not possible to catch all FSEUs that occurred on the investigated device in the field. This is primarily due to two reasons: first, some FSEUs do not propagate to system outputs and remain undetected. Second, our list of *Possible FSEUs* does not cover all FSEUs and is continually being updated to try and do so. Additionally, this discrepancy can be partially attributed to some potential inaccuracies in the raw fit rates as published by FPGA vendors.

We have also analyzed the susceptibility of each architectural module within the ECU. Using this analysis tool, it

is possible to investigate the soft error vulnerability of each component in the hierarchical architecture, from a top-level module down to a net or cell within a sub-module. Such information is very useful for soft error debug, diagnosis, and re-design. The FIT contribution and the percentage of used SRAM cells of each top-level module, as an indication of relative size of each module, have been reported in Table 6. As can be seen in this table, the soft error contributions of different architectural modules are not uniform. Furthermore, some smaller modules (i.e., mapped using a small fraction of the FPGA resources) contribute to a considerable percentage of the overall FIT rate. Such information can be used by architects and designers to include protection (in terms of architectural or device redundancy) for those modules.

Table 6. Modular susceptibility analysis.

Module name	FIT(%)	% used SRAM cells
module_1	18.51	25.26
module_2	9.73	3.25
module_3	9.32	13.73
module_4	8.87	11.30
module_5	5.85	1.50
module_6	5.73	6.78
module_7	3.47	3.06
module_8	2.33	3.88
module_9	2.32	3.86
module_10	2.30	3.93
the rest	**31.57**	**23.45**

5 Conclusions

Dependability is one of the most critical factors for storage systems. Since FPGAs are commonly used in the implementation of these systems (particularly in the design of LUMs) it is important to be able to estimate the soft error reliability of FPGA-based designs. Designs mapped into FPGAs are more susceptible to soft errors than ASIC implementations since the majority of an FPGA chip area is dedicated to memory elements storing the configuration of the FPGA or circuit state. Moreover, soft errors in configuration memory cause permanent errors in the mapped design which cannot be corrected by traditional retry mechanisms.

We have presented a case study on an FPGA-based controller used in the design of commercial storage systems. We have developed a methodology to analyze the field failure (soft error) data gathered from more than 12,000 manufactured working machines and localized SEUs within the FPGA-based controller modules in the system. We have also compared FIT rates obtained from this comprehensive field failure analysis versus results obtained from our analytical framework. These results show that our analytical framework can accurately predict FIT rates (there is an 81% overlap in FIT rates obtained between the analytical mode and the field failure data) while the runtime of our framework is completely tractable (only 3.5 hours for a very large design mapped into one of the largest commercial FPGA devices).

Future work includes similar analysis for the additional FPGA-based modules contained within the LUM that was investigated in this case study as well as the effect of selective architectural protection on the overall soft error rate.

References

[1] G. Asadi and M. B. Tahoori, "An Accurate SER Estimation Method Based on Propagation Probability," In the IEEE/ACM Intl. Conference on Design, Automation and Test in Europe (DATE), pp. 306-307, Munich, Germany, March 2005.

[2] G. Asadi and M. B. Tahoori, "An Analytical Approach for Soft Error Rate Estimation in Digital Circuits," Proc. of the IEEE Intl. Symp. on Circuits and Systems (ISCAS), VOl. 3, pp. 2991-2994, Kobe, Japan, May 2005.

[3] G. Asadi and M. B. Tahoori, "Soft Error Rate Estimation and Mitigation for SRAM-Based FPGAs," Proc. of the 13th ACM Intl. Symp. on Field-Programmable Gate Arrays (FPGA-2005), pp. 149-160, Monterey, CA, Feb. 2005.

[4] R.C. Baumann, "Radiation-Induced Soft Errors in Advanced Semiconductor Technologies," IEEE Trans. on Device and Materials Reliability, pp. 305-316, Vol. 5, Issue 3, Sept. 2005.

[5] Hennessy and Patternson, "Computer Architecture: A Quantitative Approach," Third Edition, Morgan Kaufmann publishers, 2003.

[6] J. Karlsson, P. Ledan, P. Dahlgren, and R. Johansson, "Using Heavy-Ion Radiation to Validate Fault Handling Mechanisms," IEEE Micro, 14(1), pp. 8-23, Feb. 1994.

[7] A. Lesea, S. Drimer, J.J. Fabula, C. Carmichael, and P. Alfke, "The Rosetta Experiment: Atmospheric Soft Error Rate Testing in Differing Technology FPGAs," IEEE Transactions on Device and Materials Reliability, Volume 5, Issue 3, pp. 317-328, Sept. 2005.

[8] S. Mitra, N. Seifert, M. Zhang, Q. Shi and K. Kim, "Robust System Design with Built-In Soft-Error Resilience", IEEE Computer, vol. 38, pp. 43-52, Feb. 2005.

[9] S. S. Mukherjee, C. Weaver, J. Emer, S. K. Reinhardt, and T. Austin, "A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-Performance Microprocessor," Proc. of the 36^{th} Intl. Symp. on Micro-architecture (MICRO-36), pp. 29-40, 2003.

[10] H. T. Nguyen and Y. Yagil, "A Systematic Approach to SER Estimation and Solutions," Proc. of the 41^{st} Intl. Reliability Physical Symp., pp. 60-70, Dallas, Texas, 2003.

[11] E. Normand, "Single Event Upset at Ground Level," IEEE Transactions on Nuclear Science, vol. 43, No. 6, December 1996.

[12] P. Shivakumar, M. Kistler, S. W. Keckler, D. Burger, L. Alvisi, "Modeling the Effect of Technology Trends on the Soft Error Rate of Combinational Logic," Proc. of the International Conference on Dependable Systems and Networks (DSN'02), Washington D.C., June 2002.

[13] J. F. Ziegler, "Terrestrial Cosmic Rays," IBM Journal of Research and Development, pp. 19-39, Vol. 40, No. 1, January 1996.