

Comparison of SET-Resistant Approaches for Memory-Based Architectures

Daniel R. Blum and José G. Delgado-Frias
School of Electrical Engineering and Computer Science
Washington State University
Pullman, WA 99164-2752
Email: {dblum, jdelgado}@eecs.wsu.edu

Abstract- Modern mission critical digital systems often require some form of fault tolerance to achieve adequate functionality in their operating environments. In the presence of radiation, particle strikes can cause temporary signal errors in ICs. Particle strikes that directly affect memory are known as Single Event Upsets (SEUs), while strikes that affect combinational logic are called Single Event Transients (SETs). SETs are becoming more of an issue as technology improves, as the properties that masked these faults in the past are decreasing in influence. In this paper we compare multiple approaches to hardening integrated circuits against SETs. These approaches include modified DICE, TMR, and delay-based schemes. The metrics for comparison are clock period (with worst-case SEU/SET), energy consumption, and circuit complexity (area) requirements. Each approach possesses unique strengths and weaknesses, and so the best choice for a given system is dependent on the circumstances of the situation.

I. INTRODUCTION

Reliability and fault-tolerance are primary concerns in the design of digital systems, particularly when considering mission critical systems such as space communications. The failure of such a system to function correctly may result in undesirable consequences. In many cases, people's lives depend on their functionality. Research in the area of fault-tolerance is becoming increasingly important because the reduction of integrated circuit (IC) feature size is resulting in circuits that are more fragile. Decreased gate capacitances allow charged particles to exert greater influence on transistor operation [1].

Latch-up, burn-out, oxide charging, reduction of carrier lifetimes, and single-event upsets (SEUs), are some of the faults that affect ICs [2]. The first four faults cause permanent damage to a chip, and must be addressed through process enhancements or a software/hardware scheme that can bypass the damaged circuitry. On the other hand, the final fault is temporary, and it can be mitigated with schemes that are not as involved. Different types of faults occur in different environments, so it is important to tailor individual fault-tolerant schemes to demands of the situation.

Environments that are not adequately shielded against radiation, such as in space or at high altitude, exhibit high transient fault rates [3]. Energetic particles may pass through an IC, leaving behind a trail of charge. If the charge collects at a reverse-biased junction in the circuit, the resulting voltage spike could alter the state of a memory cell, disrupting the

system for a long period of time. This disruption is known as an SEU. SEUs are becoming more of a problem as IC technology improves [1].

Integrated circuits rely on combinational logic to perform computations and memory to store the results of these computations as well as the state of the system. Memory cells inside of digital systems are especially susceptible to error, as a voltage spike on a feedback line would be amplified by the latch inverters, which could easily change the state of the cell, resulting in SEU. A single erroneous value stored in this memory can continually disrupt the functionality of the entire system. Additionally, transient faults that originate in combinational logic can prorogate through multiple gates and be stored in memory. This phenomenon is known as a Single Event Transient (SET). SETs can cause the same level of disruption as faults that strike memory directly.

A number of approaches have been designed to mitigate the effect of SEUs and SETs. At the circuit-design level, area and/or time can be traded for increased fault-tolerance. Area-redundant techniques utilize multiple copies of a circuit. If a fault affects one copy, the other copies take over and correct the error. In contrast, time-redundant techniques rely on the fact that transient errors must dissipate eventually. These schemes use an extended clock period, allowing them to sample the data multiple times or halt write operations until the transient error is gone. Many approaches rely on both area and time redundancy. Software-level and process-level schemes also exist, but they are outside of the scope of this report.

In this paper, we focus on SETs, as their effect is expected to increase substantially in the near future. A number of SET-resistant approaches are considered, including modified versions of the Dual-Interlocked Storage Cell (DICE), Triple-Modular Redundancy (TMR), and delay-based schemes. These approaches are evaluated using a number of comparison metrics, including energy consumption, clock period (with worst-case SEU/SET), and circuit complexity/interconnect requirements.

Energized particle-induced current is often modeled as a current source at the struck node in Spice. Schemes in this paper focus on SEU/SET immunity, so the level of detail incorporated into the particle strike model is not extremely important. Specifically, a 700 μ A triangular current pulse is used with rise time 50ps, and fall time 150ps [4].

In section II of this paper, the characteristics of SEUs and SETs are explained in greater detail. Section III looks into modified DICE designs that have the ability to tolerate SETs. A TMR approach is described in section IV. Delay-based memories are discussed in section V. Section VI contains an analysis of the simulation results presented in this paper. Finally, a conclusion is offered in section VII.

II. SINGLE EVENT UPSETS AND SINGLE EVENT TRANSIENTS

Single event upsets (SEU) are radiation-induced soft errors in microelectronic circuits that are caused by the accumulation of charge from an energized particle strike. As a particle passes through an IC, it loses energy and frees electron-hole pairs along its path. This charge can be collected at reverse-biased junctions via electrical drift and diffusion [5]. The drain of a transistor in an “off” state is reverse-biased, and so it is a collector of charge from ion strikes. If enough charge is collected at the drain of an “off” transistor, the voltage of the corresponding node will change significantly. This effect does not damage the circuit, and it is usually temporary, as the charge can be dissipated by the surrounding circuitry. However, the upset can cause long-term disruption to the operation of the system if the circuit is not able to recover from the change in voltage. In particular, the logic value stored in an SRAM cell may be altered if the circuit cannot dissipate the collected charge before feedback causes the cell to flip state. The corruption of memory inside of such a circuit could cause an entire system to fail.

In modern technologies, upsets can occur from particle strikes that do not directly impact memory circuitry. Injected charge within combinational logic blocks may initiate a voltage spike that propagates to other circuitry [10]. This effect is known as a single event transient (SET) [11]. If the disturbance reaches the write enable or data lines of a memory cell, the stored data could be overwritten. SETs are becoming more of a problem as IC feature sizes and clock periods shrink. Reduced feature sizes create smaller node capacitances, which result in larger voltage spikes. Smaller clock periods increase the probability that spurious data will propagate through combinational logic and be latched by a memory cell [12].

In the past, the fault-tolerance of combinational logic has been much less of a concern than that of memory. Logical, electrical, and latching-window masking have prevented almost all upsets from propagating through logic chains and into memory. However, the reduced feature sizes and supply voltages found in advancing technologies are reducing the strength of these masking effects. It has been predicted that by 2011, the rate of SET occurrence will finally approach the SEU rate. This is important because many of the current techniques available for SET mitigation require significant cost in terms of circuit complexity, delay, and/or energy consumption. On the other hand, the effect of SEUs that directly affect memory can be controlled with relatively inexpensive schemes [12].

Since both combinational logic and memory cells are made up of transistors, the physics of SETs are similar to those of SEUs. Particle strikes leave behind a trail of electron-hole pairs that can collect at reverse-biased junctions. If enough

charge collects at a junction of a combinational logic transistor, a significant voltage spike will occur. This spike could produce a temporary pulse in the output of the associated logic gate. In turn, this pulse could propagate through the logic chain and be latched by memory, resulting in an upset.

Assuming a large enough pulse occurs at the output of a logic gate, only electrical, logical, and latching-window masking can prevent the pulse from being latched by a memory cell. Electrical masking occurs when the capacitance of gates in a logic chain dissipates a transient pulse before it reaches a memory cell. Logical masking takes effect when a pulse occurs at a node that does not logically affect the boolean equation implemented by the chain. Finally, latching-window masking is responsible when a pulse reaches the input to a memory cell while a write operation is not enabled [12].

As technology improves, smaller transistors result in diminished electrical masking, as they carry less capacitance. Additionally, higher clock rates initiate a greater number of write opportunities per unit time, which decreases the effect of latching-window masking. Logical masking is not necessarily affected by technological improvements, as it is dependent on the function implemented by the associated circuitry. However, increasing pipeline depths results in fewer logic gates between memory latches, which could reduce the effect of logical masking. All things considered, SETs will become a problem requiring substantial attention in the near future [12].

III. MODIFIED DICE APPROACHES

The Dual Interlocked storage Cell (DICE) is a circuit-level design that has the capability to recover from transient faults at any of its feedback nodes [4]. The DICE memory cell is based off of four inverters which are connected in an unorthodox fashion. Arranging the latch interconnectivity this way assures that two separate inverters will restore the node to its original state via feedback. However, the unmodified DICE cell does not have the capability to recover from transient errors on its data or feedback lines. This section will describe schemes that add redundant data and control signals to the DICE cell, providing it with the ability to resist a single transient error at any location in the system. The transient error may be an SEU or SET, which is a particle-induced upset that originates in combinational logic and propagates to a memory cell.

DICE approaches rely on redundancy built into the memory cells to provide the capability to recover from faults [4]. The core of the DICE cells considered in this paper consists of four inverters that are connected in an interlocked fashion (as transistors M1-M8 in figure 1). This interconnectivity allows the cells to tolerate a single SEU at any internal location in the memory core. However, the core does not guarantee protection against SETs that propagate to the write enable or data lines. The circuits in figures 1 and 2 were designed to add SET protection to the DICE core.

Four write transistors are used in figure 1, each of which is enabled by an independent write line [13]. If an SET were to disrupt a write line, only one internal node would be directly affected. The DICE core can tolerate this, and so the state of

the memory cell would be preserved. Redundancy is also present in the data lines. One SET could affect a pair of data lines, but this would not be sufficient to cause an error during a write operation.




Figure 1. Enhanced DICE memory cell 1.




Figure 2. Enhanced DICE memory cell 2.

Another important property of the DICE 1 design is the fact that its read circuitry does not affect its ability to tolerate upsets. This is due to buffering between the internal cell nodes and the data lines while the read transistors are enabled. Before a read operation, all of the data lines are precharged to VDD. During a read operation, the two appropriate data lines are pulled down to GND. Cells without buffering have exhibited increased probability of failure resulting when SEU strikes occur at the beginning of a read operation. The combined influence of a disrupted internal node and charge sharing from the data lines can be sufficient to flip the state of the memory.

The cell in figure 2 uses different approaches to protect the data and write enable lines and to provide buffering during reads [13]. Write operations drive two internal nodes through transmission gates. An SET affecting one of these two paths would not be sufficient to cause an upset. During reads, tri-

state inverters are enabled, which would drive the data lines to the desired values while buffering the internal nodes from the data lines. This approach requires only two independent write enable paths instead of four, and uses fewer data lines than the previous circuit. However, these benefits come at the cost of additional energy consumption and additional write delay.

Simulations have been performed on both of the above DICE Cell designs for this report. Specifically, a 1-bit data transfer system was set up for each cell design. These systems are relatively generic, in that they can be thought of as a stage in a pipeline or an SRAM-to-SRAM memory transfer. For each cell design, one copy of the cell functions as the source, and a second copy functions as the destination. The source reads out its data, which passes through any necessary buffers or combinational logic. The output of the buffers or logic is then written to the destination cell. Figure 3 is a depiction of this design.




Figure 3. Enhanced DICE memory cell data transmission.

Each DICE design was evaluated with respect to delay, energy consumption, and circuit complexity. For the first DICE cell, a minimum clock period of 350ps was obtained in the 0.18μm process utilized for this report. The clock period expanded to 550ps after a 200ps SET was inserted into the simulation. The second DICE design required a clock period of 650ps, which had to be increased to 850ps to maintain functionality after a 200ps SET was added.

Energy consumption for the first DICE design totaled to 712.18fJ per clock cycle in the worst case without upset. The read cell contributed 8.88fJ, the write cell 215.2fJ, and the buffers 488.1fJ. In comparison, the second DICE cell used 848.26fJ of energy. This figure consists of 89.46fJ from the read cell, 347.1fJ from the write cell, and 411.7fJ from the buffers.

It is possible to quantify the circuit complexity of the DICE designs by the number of transistors used in memory, the number of combinational logic copies needed, and the I/O requirements. The first DICE cell requires 18 memory transistors, two differential copies of combinational logic, two differential data signals, four write enable lines, and two read enable lines. On the other hand, the second DICE cell requires 20 memory cells, two copies of combinational logic, two data lines, two write enable lines, and two read enable lines.

IV. TRIPLE MODULAR REDUNDANCY

TMR designs utilize area redundancy and voting to bypass faults. In a TMR-based system, three copies of all critical circuitry are used. If one copy is compromised, the other two will dominate the voting process and the output data will still be valid. When a read operation is performed in a TMR design, three memory cells each send their version of a data bit. Voting circuitry then passes on the data that was sent by

the majority of the memory cells. The drawback of this approach is that it requires 200% additional memory cells, in addition to voting logic and other circuitry. Also, updating an incorrect bit in memory requires additional complexity, as feedback of the voting results to the source memory cell inputs must be utilized.

Many different TMR configurations exist, and some of them do not provide sufficient protection against SETs. Schemes relying on fewer than three copies of combinational logic or the voting circuitry will not be able to withstand an SET unless signal delays are incorporated, as shown in figure 4 [14]. Unfortunately, the use of signal delays has a direct impact on performance.

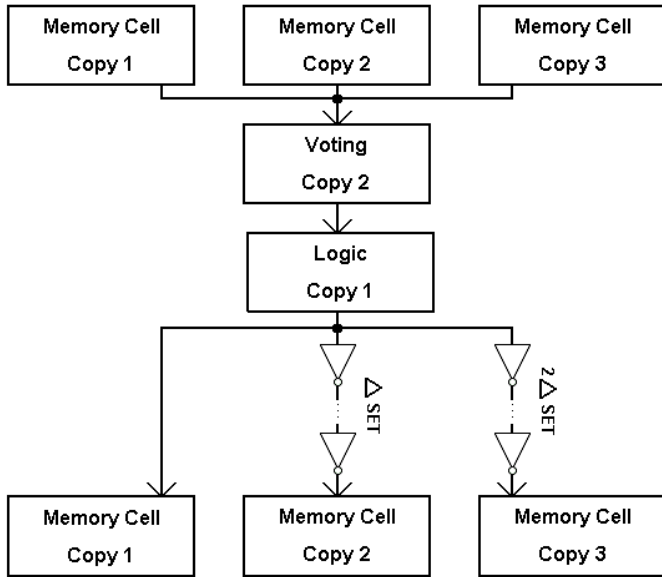


Figure 4. Block diagram of TMR scheme with delays.

The approach presented in figure 5 attempts to maximize performance by relying on three copies of logic and voting circuitry instead of any signal delays. For this report, a TMR 1-bit data transmission system was simulated with three source SRAM cells, three receiver cells, and triple-redundant CMOS voting circuitry. The source cells read out their data, the voting gates select majorities, and the majorities are written to the receiver cells.

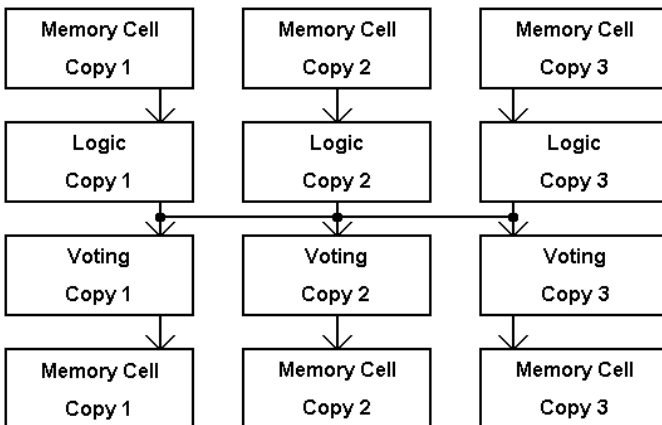


Figure 5. Block diagram of TMR scheme with redundant voting.

The total energy consumption of the TMR setup was 1150.38fJ per clock cycle in the worst case without upset. The read cells consumed 327fJ, the write cells 151.38fJ, the buffers 208.4fJ, and the majority gates 463.6fJ.

This report measures circuit complexity in the TMR design by the number memory transistors, the number of voting transistors, the number of combinational logic copies, and the I/O configuration. Memory in this TMR design consists of 18 transistors, while the voting circuitry requires 36 transistors (other majority gate designs can be selected to reduce transistor count). Three copies of combinational logic are required, as well as three data lines and three read/write enable lines.

V. DELAY-BASED APPROACHES

Delay approaches halt the system until a delayed version of each data signal agrees with the primary data path. If the length of the delay is equal to the maximum SET duration (ΔSET), the primary and delayed signals will not be equivalent while an SET is present. Figure 6 shows one implementation of this strategy [15]. A tri-state buffer is used to hold off writes to the memory until the primary and delayed signals are equivalent. The clock period must be expanded by ΔSET for this approach to work. Low circuit complexity requirements are a major benefit to this approach. The delayed data path is generated after the combinational logic, so logic duplication is not needed. Additionally, this design is compatible with any SEU-redundant memory cell, including DICE, enhanced capacitance/resistance cells, process enhanced cells, etc. Unfortunately, this scheme does not protect against an SET occurring at the output of the tri-state buffer. However, if a memory cell with two or more write paths is used (such as a DICE cell), the tri-state buffer could be eliminated, and greater immunity could be achieved.

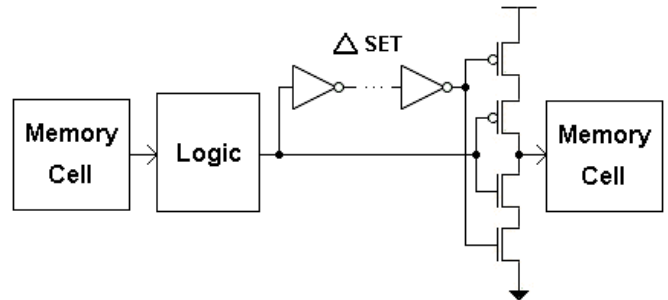


Figure 6. Delay-based tri-state buffer scheme for incorporating SET resistance [15].

Figure 7 illustrates a temporal-redundant design that does not rely on a tri-state buffer, eliminating the SET vulnerability described above. The DICE cell is designed in such a way that both inputs must be equivalent for writes to occur. Because of this, the presence of an SET in one of the write paths halts the system. As with the previous temporal-redundant approach, this setup also benefits from relatively low circuit complexity, as only one copy of the combinational logic is needed. However, the lack of complexity takes a toll on both the power consumption and the performance. Only one input of the receiver cell is driven during the first

portion of a write operation, which causes conflict between the internal nodes that consumes additional power. Additionally, the clock period must be wide enough to tolerate the worst-case SET. For this setup, a worst-case SET occurs in the middle of a clock cycle. When this happens, the two inputs of the receiver cell are driven to opposite values for a greater amount of time, which degrades performance.




Figure 7. Delay-based DICE scheme for incorporating SET resistance.

Delay, energy consumption, and circuit complexity statistics have been recorded for the delay-based DICE approach. With no upsets, a clock period of 950ps was achieved. After incorporating a 200ps SET, the clock period had to be stretched to 1.1ns.

In the worst case without an upset, 1028.49fJ of energy was required in a clock cycle. This breaks down to 91.19fJ from the read cell, 503.9fJ from the write cell, and 433.4fJ from buffers.

The circuit complexity of this scheme can be divided into the number of memory transistors, the number of delay transistors, the number of combinational logic copies, and the necessary I/O. Each memory unit requires 20 transistors, while the delay circuitry consists of four transistors. One copy of the combinational logic is needed, along with one data line, two write enable lines, and one read enable line.

VI. ANALYSIS OF RESULTS

In the previous sections, clock period, energy consumption, and circuit complexity statistics have been presented for four different SET-resistant approaches. The clock period results are shown in table 1. The left column shows the minimum clock period attained for each approach with no upsets, while the right column shows the clock periods needed to tolerate a 200ps SET. Table 2 presents energy consumption figures that are broken down into read memory, write memory, buffer, voting, and total contributions. Finally, table 3 depicts the circuit complexity of each approach, divided up into memory, voting/delay, logic and I/O categories.

These results show that no single SET-resistant approach dominates all of these comparison categories. The first DICE design is well balanced, achieving high performance (350-550ps clock period) and low power consumption (712.18fJ) at the cost of moderate size and the requirement of two independent differential data paths. The second DICE approach removes the requirement for differential data paths, but this results in an increased clock period (650-850ps) and energy consumption (848.26fJ). The delay-based DICE approach requires only one non-differential data path, but this degrades the clock period (950-1100ps) and energy consumption (1028.49fJ) even further.

The clock period for all approaches except TMR must be increased to tolerate SET pulses of greater than 200ps width.

This robustness is the main advantage of the TMR approach. Unfortunately, this benefit has substantial costs in terms of energy consumption (1150.38fJ) and circuit complexity (18 memory and 36 voting transistors, as well as three independent data paths).

TABLE 1. MINIMUM CLOCK PERIODS

	With No Upset	With 200ps SET
DICE 1	350ps	550ps
DICE 2	650ps	850ps
TMR	550ps	600ps
Delay DICE	950ps	1.1ns

TABLE 2. ENERGY CONSUMPTION FIGURES

	Read	Write	Buffers	Voting	Total
DICE 1	8.88fJ	215.2fJ	488.1fJ	n/a	712.18fJ
DICE 2	89.46fJ	347.1fJ	411.7fJ	n/a	848.26fJ
TMR	327fJ	151.38fJ	208.4fJ	463.6fJ	1150.38fJ
Delay DICE	91.19fJ	503.9fJ	433.4fJ	n/a	1028.49fJ

TABLE 3. CIRCUIT COMPLEXITY QUANTIZATION

	Memory	Vote/Delay	Logic	I/O
DICE 1	18	n/a	2x (diff)	2 data (diff), 4 write, 2 read
DICE 2	20	n/a	2x	2 data, 2 write, 2 read
TMR	18	36	3x	3 data, 3 read/write
Delay DICE	20	4	1x	1 data, 2 write, 1 read

From a performance perspective, the DICE 1 design comes out on top. Figure 8 is a plot of the clock period with a 200ps SET vs. the energy consumption of each design. The lower left corner of the graph is optimal, and the DICE 1 cell is closest to this position.

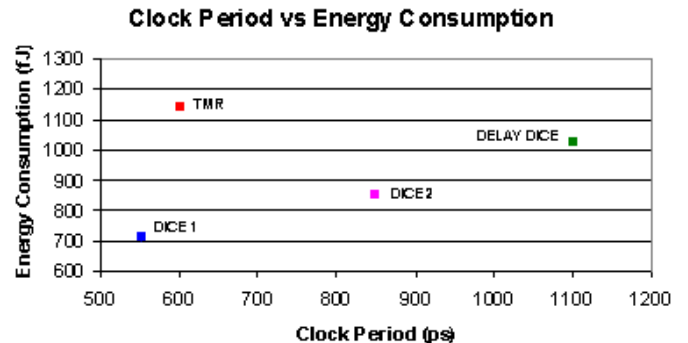


Figure 8. Plot of clock period vs. energy consumption for all of the SET-resistant schemes.

The four approaches can be further classified based on the number of independent data paths (or copies of combinational logic) they require. The delay-based DICE approach requires only one independent path, while the DICE 1 and DICE 2 designs require two paths, and the TMR setup requires three. Using a greater number of independent paths obviously increases the complexity of the design, but it can also improve performance. Approaches that rely on one independent path must incorporate some type of delay to filter out SETs from memory inputs. This directly impacts speed and power consumption, even when no upset is present. Designs that utilize two paths do not need to incorporate delays, but they

must pause the system while under the influence of an SET. Because of this, the clock cycle must be increased by ΔSET , the length of the longest possible SET. On the other hand, schemes that rely on three paths do not have to pause the system when under the influence of an upset. This means that the clock period is not dependent on ΔSET , making the system scalable to situations that must tolerate wide SETs.

VII. CONCLUSION

In this paper, the modified DICE, TMR and delay-based approaches were evaluated with respect to energy consumption, clock period, and circuit complexity/interconnect requirements. The first DICE cell exhibits low energy consumption, has low write delay, requires approximately three times the complexity and twice the interconnect of a regular SRAM cell. The second DICE cell requires additional energy consumption and write delay when compared to the first cell, but does not require differential data signals. TMR requires more than three times the energy consumption, complexity, and interconnect of a traditional SRAM cell, although it maintains a reasonable clock period when affected by SET pulses with long duration. Finally, the delay-based DICE approach only requires one data path and one copy of combinational logic, but this comes at the cost of a clock period increase and additional energy consumption when compared to the other DICE designs.

A number of factors should be taken into account while selecting an SET-resistant approach. The first two modified DICE designs are the best choices in situations where the maximum SET width is small and performance and power consumption are the most important factors. This is the case in a number of modern custom IC designs. The DICE 1 cell is suited for differential data paths, while DICE 2 was created for non-differential systems. The TMR design is a good choice in situations where wide SET pulses are expected and circuit complexity is not a major concern. Wide SET pulses are an issue in heavily radioactive environments. Finally, in many designs such as mobile applications, a major concern is circuit complexity. In these situations, the delay-based DICE design may be the best candidate presented in this paper. Clearly, each approach has its strengths and weaknesses, and so the optimum design depends on the demands of the application.

REFERENCES

[1] E.L. Peterson et al, "Calculations of cosmic ray induced soft upsets and scaling in VLSI devices," *IEEE Transactions on Nuclear Science*, vol. NS-29, no. 6, Dec 1982, pp. 2055-2063.

[2] A.M. Chugg, "Ionising radiation effects: A vital issue for semiconductor electronics," *Engineering Science and Education Journal*, vol. 3, pp. 123-130, Jun 1994.

[3] J.J. Wang et al, "SRAM based re-programmable FPGA for space applications," *IEEE Transactions on Nuclear Science*, vol. 46, no. 6, Dec 1999, pp. 1728-1735.

[4] T. Calin, M. Nicolaidis, and R. Velazco, "Upset hardened memory design for submicron CMOS technology," *IEEE Transactions on Nuclear Science*, vol. 43, pp. 2874-2878, Dec. 1996.

[5] P.E. Dodd, M.R. Shaneyfelt and F.W. Sexton, "Charge collection and SEU from angled ion strikes," *IEEE Transactions on Nuclear Science*, Vol. 44, no. 6, Dec 1997.

[6] P.E. Dodd, F.W. Sexton and P.S. Winokur, "Three-dimensional simulation of charge collection and multiple-bit upset in Si devices," *IEEE Transactions on Nuclear Science*, Vol. 41, no. 6, Dec 1994.

[7] P.E. Dodd and F.W. Sexton, "Critical charge concepts for CMOS SRAMs," *IEEE Transactions on Nuclear Science*, Vol. 42, no. 6, Dec 1995.

[8] P.E. Dodd et al, "Impact of ion energy on single-event upset," *IEEE Transactions on Nuclear Science*, Vol. 45, no. 6, Dec 1998.

[9] P.E. Dodd et al, "Impact of technology trends on SEU in CMOS SRAMs," *IEEE Transactions on Nuclear Science*, Vol. 43, no. 6, Dec 1996.

[10] K.J. Hass and J.W. Ambles, "Single event transients in deep submicron CMOS," *Proc. 42nd Midwest Symp. On Circuits and Systems*, Las Cruces, NM, vol. 1, pp. 122-125, Aug 1999.

[11] M.P. Baze and S.P. Buchner, "Attenuation of single event induced pulses in CMOS combinational logic," *IEEE Transactions on Nuclear Science*, vol. 44, pp. 2217-2223, Dec. 1997.

[12] P. Shivakumar et al., "Modeling the effect of technology trends on soft error rate of combinational logic," *Proc. 2002 International Conf. on Dependable Systems and Networks*, Bethesda, Maryland, pp. 389-398, Jun. 2002.

[13] D.R. Blum, M.J. Myjak, and J.G. Delgado-Frias, "Enhanced fault-tolerant data latches for deep submicron CMOS," *Proc. 2005 International Conference on Computer Design*, Las Vegas, NV, June 2005.

[14] D.G. Mavis, P.H. Eaton, "Soft error rate mitigation techniques for modern microcircuits," *Reliability Physics Symposium Proceedings*, 40th Annual, pp. 216-225, Apr. 2002.

[15] P. Mongkolkachit, B. Bhuvu, "Design technique for mitigation of alpha-particle-induced single-event transients in combinational logic," *IEEE Transactions on Device and Materials Reliability*, vol. 3, no. 3, pp. 89-92, Sep. 2003.