

# Schemes for Eliminating Transient-Width Clock Overhead From SET-Tolerant Memory-Based Systems

Daniel R. Blum, *Student Member, IEEE*, and José G. Delgado-Frias, *Senior Member, IEEE*

**Abstract**—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 and spread to memory are called single event transients (SETs). In this paper, we propose two novel approaches to hardening integrated circuits against SEUs and SETs. The proposed approaches are fully-differential dual-interlocked storage cell (DICE) and triple path DICE (TPDICE).

The fully-differential DICE and TPDICE approaches are compared against two existing approaches, which are triple modular redundancy (TMR) and basic SET-tolerant DICE. All approaches except for the basic SET-tolerant DICE scheme share a common theme, which is the ability to bypass SEUs and SETs. This is critical for performance, as it allows the system to proceed with subsequent operations while a cell is recovering from the effects of a particle strike. SET pulse widths can be substantial (up to 2 ns), and so high-performance systems cannot afford to pause operations while these pulses are present. The minimum clock periods obtained for the basic SET-tolerant approach were 515 ps with no SET, and 1310 ps with a 500 ps SET (in 0.18  $\mu\text{m}$  CMOS). In contrast, the clock periods for the bypass-capable approaches with no SET/500 ps SET were 628/749 ps for TMR, 348/480 ps for fully-differential DICE, and 434/552 ps for TPDICE. Among the approaches that bypass transient pulses, TPDICE is the most balanced. TMR suffers from overhead due to its need for external voting circuitry. In addition to this, fully-differential DICE cannot be used with combinational logic, while TPDICE can.

**Index Terms**—Single-event upsets, single-event transients, soft errors, radiation effects, hardened by design.

## I. INTRODUCTION

**R**ELIABILITY 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 susceptible to upset. Decreased gate capacitances allow charged particles to exert greater influence on transistor operation [1]–[3].

Manuscript received February 9, 2006; revised March 22, 2006. This work was supported in part by the Boeing Centennial Endowed Chair, School of Electrical Engineering and Computer Science, Washington State University.

The authors are with the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99164-2752 USA (e-mail: dblum@eecs.wsu.edu; jdelgado@eecs.wsu.edu).

Digital Object Identifier 10.1109/TNS.2006.874496

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 a single event upset (SEU). Additionally, voltage spikes that are initiated in combinational logic can propagate 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 upsets originating in memory [4].

A number of approaches have been proposed and implemented 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 circuitry. If a transient pulse 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 potential transient errors are 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 paper.

In this paper, we focus on SETs, as their effect is expected to increase substantially in the near future. The majority of current SET-tolerant approaches halt the system while a transient event is present. For example, many of these schemes utilize memory cells with multiple inputs, fed by redundant or delay-separated logic paths [5]–[10]. The redundant or time-separated data paths assure that a transient pulse can be expected at only one of the inputs at any given time. If the inputs are not equal, the cell has detected a potential SET, and so it pauses the write operation until the voltage spike dissipates. This approach is effective, although it requires a write cycle overhead of at least the maximum expected SET duration. In contrast, this paper focuses on designs that bypass transient voltage pulses, allowing write operations to complete before these pulses dissipate. These approaches transmit sufficient data to the destination cell during write operations to ensure that the cell resolves to the desired state, even if a transient pulse affects an input throughout the entire operation. Because of this, clock period overhead dependent on SET width is avoided. This is beneficial from a performance perspective, especially in situations where the SET pulse widths are comparable to the clock period.

Many SET-tolerant circuits have been designed assuming 100–200 ps SET width [11]. Adding this overhead to the clock

period is tolerable, but it negatively impacts performance. Greater system frequency can be achieved by avoiding this overhead. However, recent in-depth studies have concluded that SET pulses can be much larger than this. While particle strikes with linear energy transfer (LET) of 5–10 MeV-cm<sup>2</sup>/mg often produce 100–200 ps pulses, cosmic ray strikes with LET of 100 MeV-cm<sup>2</sup>/mg can induce SET pulses up to 2 ns long [12]–[14]. In general, SET pulse widths increase linearly with increasing LET from particle strikes [15]. Additionally, pulse widths do not necessarily shrink with decreasing feature size. Obviously, wide SET pulses would seriously limit the performance of systems that halt while transients are present. Specifically, 2 ns SETs impose a ceiling of 500 MHz to the clock frequency. This is not acceptable in many high-performance designs. In future technologies, this constraint will be even more of an issue.

A number of SET-tolerant approaches are considered in this paper, including triple modular redundancy (TMR), a basic SET-tolerant version of the dual-interlocked storage cell (DICE), a fully-differential DICE circuit, and a triple path DICE (TPDICE) design. TMR, fully-differential DICE, and TPDICE have the capability to bypass SEUs and SETs, thereby avoiding clock period overhead dependent on the maximum transient pulse width. All four approaches are evaluated using a number of comparison metrics. Performance comparisons are made by considering minimum clock period figures taken with no disruptions, and also with 200 ps and 500 ps SETs in the system. Energy consumption and circuit complexity figures are also considered.

## II. SINGLE-EVENT UPSETS AND SINGLE-EVENT TRANSIENTS

SEUs are radiation-induced soft errors in ICs that are caused by the accumulation of charge from an energized particle strike. This charge can be collected at reverse-biased drain junctions via electrical drift and diffusion [16], [17]. The logic value stored in an SRAM cell will be lost if the circuit cannot dissipate the charge before feedback causes the cell to flip state [18]. The corruption of memory inside of such a circuit could cause an entire system to fail.

The amount of charge collected at a sensitive node in an SRAM cell is only one of the factors that affects SEU probability. Other important factors include the time over which charge is collected, the response time of the circuit, and the load circuitry [16]. The contribution of these variables implies that the widely used model of a constant critical charge threshold for upset may not be accurate enough for some situations.

High concentrations of charged particles are often present in environments possessing high levels of radiation. Particles in space often possess very high LET, which makes SEU more likely [19]. Feature sizes are shrinking as IC manufacturing technology improves, resulting in smaller node capacitances that are more susceptible to injected charge [17]. This increased sensitivity has resulted in SEU observation in terrestrial environments [20]. In addition, radioactive materials inside of the IC packaging can emit alpha particles. Even if there are no alpha particles produced by packaging, terrestrial cosmic rays are becoming more of a factor as cell capacitance decreases [17].

In modern technologies, upsets can occur from particle strikes that do not directly impact memory circuitry. Charge injected into combinational logic may initiate a voltage spike that propagates to memory [21]. This effect is known as an SET [22]. If the disturbance affects 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.

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 memory. Electrical masking occurs when the capacitance of gates in a logic chain dissipates a transient pulse before it reaches memory. Logical masking takes effect when a pulse occurs at a node that does not 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 [23].

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. A linear dependence has been experimentally demonstrated between system clock frequency and SET error rate [24]. Finally, 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 result in fewer logic gates between memory latches, which could reduce the effect of logical masking. All things considered, SETs are becoming a problem requiring substantial attention [23].

## III. TRIPLE MODULAR REDUNDANCY

TMR is a classic fault-tolerant approach that utilizes temporal and/or spatial redundancy and voting to filter out transient pulses. Temporally redundant TMR relies on an extended clock period to sample data at multiple different times. Spatially redundant TMR-based systems rely on three copies of all critical circuitry. 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 through combinational logic. 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 voting circuitry will not be able to withstand an SET unless signal delays are incorporated, as shown in Fig. 1 [7]. This circuit is an example of a temporally redundant TMR approach. The left path in this figure has no delay, the center path is delayed by the maximum SET width, and the right path is delayed by twice the maximum SET width. Thus, if an SET disrupts the voting or

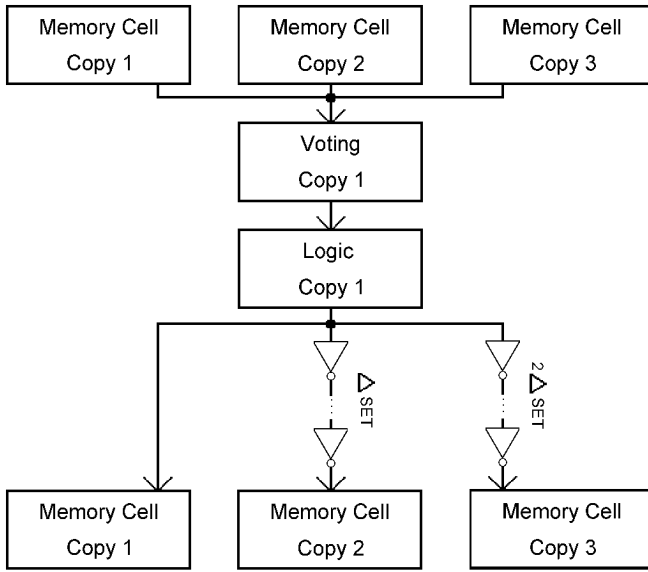


Fig. 1. Block diagram of TMR scheme with delays. This design is not bypass-capable, as overhead of at least twice the maximum SET width must be added to the clock cycle.

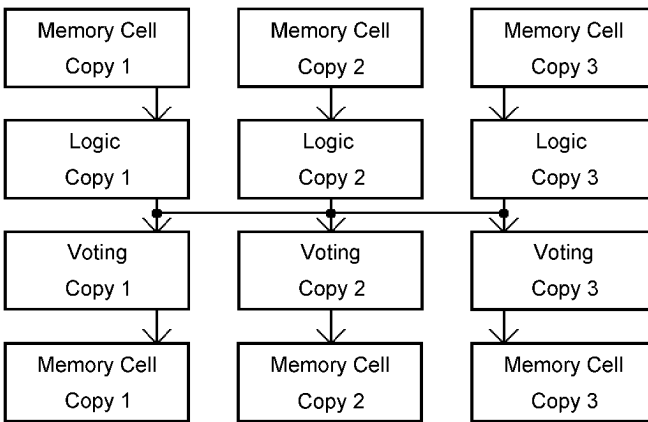


Fig. 2. Block diagram of bypass-capable TMR scheme with redundant voting.

combinational logic, a transient pulse can affect only one destination cell at any time during a write operation. Unfortunately however, the use of signal delays has a direct impact on performance. An overhead of at least two times the maximum width SET must be added to the clock cycle. Since this overhead is proportional to the SET width, this scheme does not bypass transient pulses.

The approach presented in Fig. 2 attempts to maximize performance by relying on three copies of logic and voting circuitry instead of any signal delays. This is an example of spatially redundant TMR. For this paper, a 1-bit data transmission system was simulated using the design in Fig. 2. The source cells read their data through logic, the voting gates select majorities, and the majorities are written to the receiver cells. Since signal delays are not used, this setup has the capability to bypass transient pulses. The voting logic can pass the correct value to the subsequent stage before a potential pulse dissipates. This is very beneficial from a performance perspective, as SET-width overhead does not have to be added to the clock cycle. High performance can be maintained even when faced with wide SET pulses.

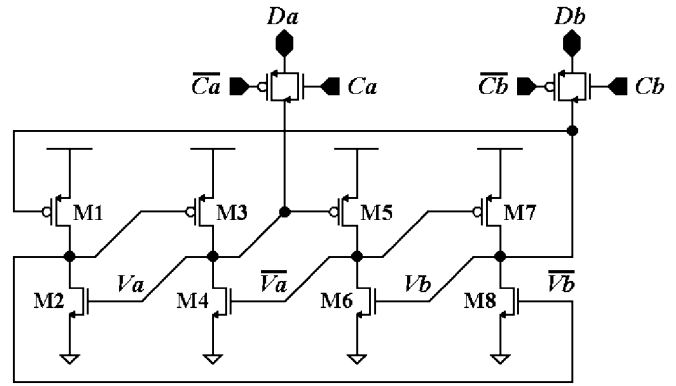


Fig. 3. DICE design-hardened memory latch, modified to tolerate SETs.

#### IV. BASIC INTERLOCKED DESIGN-HARDENED LATCHES

Significant effort has been put into the construction of design-hardened memory latches. The goal of this effort has been to efficiently protect digital systems from SEU and SET-initiated faults. The DICE circuit is a standard example of such a latch [25]. It relies on two inverter pairs and interlocked feedback to provide the capability to recover from a transient pulse at any one of its four internal nodes. However, the original DICE design did not provide protection against SETs affecting its data or enable paths.

##### A. SET-Tolerant Design-Hardened Latches

Although the original DICE design is susceptible to SETs, the circuit can be modified to tolerate these disturbances. Fig. 3 illustrates one such approach. This design uses two independent data paths, two copies of any combinational logic, and two independent access enable paths. The dual-rail logic approach presented in [9] and the two-input modified DICE cell in [6] are analogous to this design. If an SET is present in one data path during a write operation, the two paths will not be equal, and the cell state will not change. After the SET dissipates, the two data paths will be equal and correct, and the write operation will proceed as intended. It is important to note that this circuit does not bypass transient pulses because it pauses while SETs are present.

Generally speaking, design-hardened latches possess a number of advantages when compared to TMR, which uses voting to filter out erroneous pulses. Firstly, design-hardened latches often require significantly less complexity. For example, the basic DICE circuit requires approximately twice the complexity of a 6T SRAM cell, whereas TMR requires more than three times this complexity. Secondly, design-hardened cells correct transient pulses at the location of the occurrence, instead of relying on external voting to filter them out. This is particularly beneficial in memory arrays, as feedback or some similar technique would be needed to correct the cell directly affected by an SEU in a TMR system. Lastly, TMR is burdened by performance, power consumption and circuit complexity overhead introduced by voting circuitry, which is not present in design-hardened latches. The bypass-capable TMR approach presented in the previous section requires three independent voters, which is a significant expense.

The main disadvantage of the DICE latch above, when compared to bypass-capable TMR, is the fact that it tolerates transient pulses by pausing until they dissipate. Most other current design-hardened latches function in this manner as well [5]–[10]. This necessitates a clock period overhead of at least the width of the maximum possible SEU/SET pulse. In situations where only 5–10 MeV-cm<sup>2</sup>/mg alpha particles are a concern, SEU/SET pulses are generally limited to only 100–200 ps. However, they can be as large as 300 ps to 2 ns when 100 MeV-cm<sup>2</sup>/mg cosmic rays are brought into consideration in high radiation environments such as space [12]–[15]. This severely restricts performance in systems that do not bypass transient pulses, especially if the pulse duration is comparable to the clock period. With future technologies, this effect will increase, as SEU/SET pulse widths do not necessarily shrink as clock periods decrease.

### B. Read-Induced Upsets in Design-Hardened Latches

When design-hardened memory cells are used in SRAM arrays with shared data busses, read enable capability is needed. In any SRAM array, including non-hardened designs, voltage ripples caused by the load of the data bus on a cell during read operations are an important concern. If the magnitude of the ripple is large enough, the state of the cell reading data can be flipped unintentionally. This problem can be mitigated by increasing the sizing ratio of the core latch transistors versus the enable transistors. The bus capacitance is not involved in this calculation, as the ripple magnitude is dependent on transistor drive strength and the discharge current magnitude, not the total charge dissipated. If the core latch transistors are sufficiently strong in comparison to the enable transistors, they will drop an acceptably low drain-source voltage for the amount of discharge current and thus the voltage ripple magnitude will not flip the latch state.

Generally speaking, the voltage ripple issue is more of a concern in design-hardened memory cells. Most of these designs tristate one or more internal nodes during recovery from charged particle strikes. This provides a protected source of redundant data that is used to restore other nodes after the transient current is dissipated. However, tri-stated nodes have no capability to charge/discharge the bus during read operations. This changes the dynamics of the voltage ripple calculation. The bus capacitance is brought into play, as the ripple magnitude becomes dependent on charge sharing between the bus and the tri-stated node(s). Transistor drive strength is no longer the dominating factor, as all transistors controlling tri-stated nodes are off. Managing the bus versus node capacitance, and also precharging the data bus to an intermediate value such as VDD/2 can help to avoid this issue. This is not practical in large memories, necessitating the adoption of output buffers for each cell.

## V. FULLY-DIFFERENTIAL DICE

For maximum performance in SET-tolerant structures, transient pulse bypass capability must be incorporated. This bypassing attribute is desirable because it avoids pausing the system until pulses dissipate. Bypass-capable structures resolve to the desired state even if one of their inputs is affected by a particle-induced pulse for an entire write operation. Fig. 4

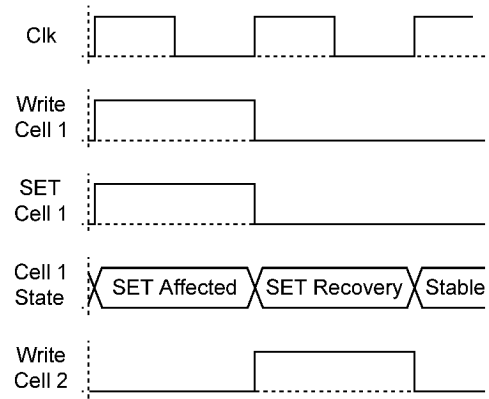


Fig. 4. Timing diagram for bypassing transient disturbances.

illustrates a basic timing scheme that facilitates the ability to bypass transient pulses. During the initial clock cycle, data is written to the first memory cell. An SET may disrupt one input during this operation. During the following cycle, the cell recovers from the transient pulse. Simultaneously, this first cell reads the recovering data, which is written to a second cell. By design, the recovering data transmitted to the second cell must contain enough information to guarantee that the second cell also resolves to the desired state. This scheme limits the propagation of transient pulses by holding them for one cycle between the time they are latched by a cell and the time they are transmitted to another cell. This provides an opportunity for latches to dissipate the pulses. Simulations on the structures presented in this paper have shown that when a transient pulse affects an input for an entire write cycle, the pulse disappears from the system within three cycles.

It is possible to create a bypass-capable hardened by design memory structure utilizing the basic DICE core. However, the bypass capability of this structure only exists in data transfer systems with no combinational logic. This design can be implemented by transmitting all four DICE node values independently. A depiction of this novel approach is shown in Fig. 5. This circuit uses four transmission gates to connect the cell nodes to the data lines during read and write operations. Alternatively, nmos devices could be used, at the cost of lower noise margins during writes while a transient pulse is affecting a data path. If a particle strike affects the source cell during data transfer, the transient pulse-induced state is transmitted to the destination cell if no combinational logic exists in the data path. Since the destination cell receives a standard transient pulse-induced pattern, the DICE core of the destination cell is able to resolve its node voltages to desired values during its recovery period.

A detailed analysis of this process can be performed by first considering the possible transient pulse-induced states for the cell in Fig. 5. Table I depicts the transient pulse-induced states of a standard DICE cell, which are applicable here. In general, if a particle strike affects one internal node, one adjacent node will also be affected. Even if both affected nodes are pushed to the respective opposite rails, the DICE design guarantees that the cell recovers to the desired state (assuming pmos and nmos transistor sizes are set appropriately, as described in [26]). The two

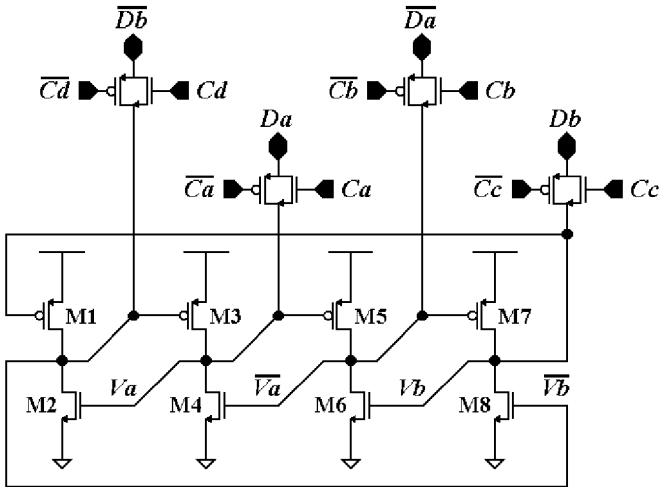


Fig. 5. Fully-differential DICE circuit utilizing a shared data bus.

TABLE I  
DICE TRANSIENT-INDUCED STATES AND RECOVERY STATES

| Initial State | Affected Nodes   | Induced State | Final State |
|---------------|------------------|---------------|-------------|
| 1010          | Va low, /Va high | 0110          | 1010        |
| 1010          | Vb low, /Vb high | 1001          | 1010        |
| 0101          | Va high, /Vb low | 1100          | 0101        |
| 0101          | Vb high, /Va low | 0011          | 0101        |

unaffected nodes are tri-stated, while the two corrupted nodes experience conflict. During recovery, the unaffected node voltages are preserved, while positive feedback pushes the corrupted nodes to desired values.

Taking this information into account, if a worst-case transient pulse-induced state is written to a destination cell, proper recovery is assured. This characteristic provides the capability to bypass transient pulses. A transient pulse affecting the source cell does not have to dissipate before the end of the write cycle to achieve proper functionality. Simply transmitting the transient pulse-induced state provides the destination cell with enough information to resolve to the desired state during its recovery period.

Four independent copies of enable lines are needed for this scheme to function as desired, assuming the enable lines do not possess the capacitance and buffer drive strength to be considered immune to SETs. Undesired writes could occur to memory cells if fewer independent lines are used. To illustrate this, consider the case where two independent enable lines are used. A transient pulse affects one enable line while the write buffers are on and the data bus is at the state opposite to that of the latch we are considering. This causes the values stored in the two cell nodes controlled by the affected enable line to be corrupted. Now two nodes in the cell are at incorrect values, and two nodes are still at proper values. From this initial condition, the cell may resolve to the opposite of the desired state, disrupting the functionality of the system.

In addition to the necessity of four independent data and enable paths, this approach is also limited by the requirement that no logic exists in the data paths for the bypass capability to function properly. These weaknesses are not present in the approach described in the following section.

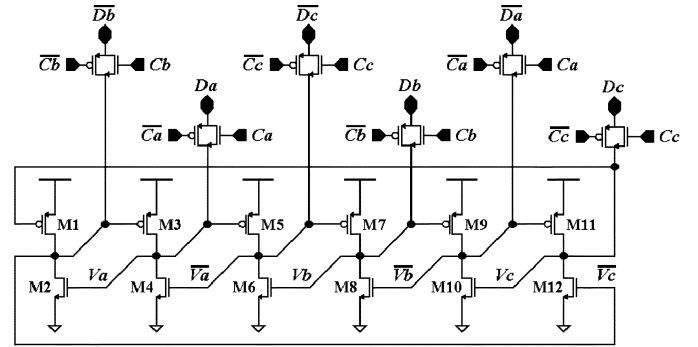


Fig. 6. TPDICE circuit utilizing a shared data bus.

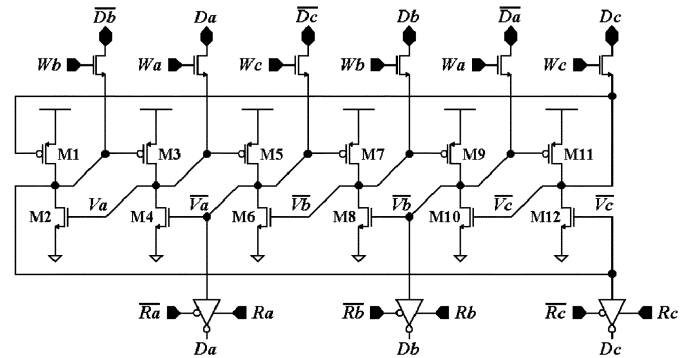


Fig. 7. TPDICE schematic with separate read and write ports.

## VI. TRIPLE PATH DICE DESIGN

All of the previous three SET-tolerant approaches presented in this paper have limitations that make them unsuitable for certain applications. TMR relies on voting circuitry that negatively impacts performance, energy consumption and circuit complexity. The basic SET-tolerant DICE approach does not have the capability to bypass transient pulses, which results in poor performance when faced with wide duration SEUs and SETs. Finally, the fully-differential DICE approach requires four independent data paths, and it cannot be used in systems with combinational logic. The TPDICE approach presented in this section addresses these issues.

### A. TPDICE Architecture

It is possible to integrate the majority voting directly into the TMR memory structure, reducing complexity and improving performance. This novel structure, shown in Figs. 6 and 7, is essentially a DICE cell extended from four internal nodes to six. It retains the main TMR advantage, which is the ability to bypass transient pulses. This ability allows the clock period to be set independent of the maximum length transient pulse. Pulses of greater width than the clock period can be tolerated. Additionally, this approach can be used with combinational logic and requires only three data/enable paths. Both of these properties cannot be matched by the fully-differential DICE design presented previously.

TPDICE latches possess a number of advantages over TMR. Majority voting in standard TMR approaches prevents the spread of faults to other cells, but it does not correct the actual cell affected by a particle strike. When TMR is used in SRAM

arrays, feedback is required to fix cells corrupted by SEUs. In contrast, the TPDICE approach needs no additional feedback to correct corrupted data. Additionally, TPDICE circuitry needs no majority voting to prevent the spread of faults to other memory locations. The absence of external feedback and voting circuitry in this approach provides increased efficiency when compared to TMR.

The TPDICE structure shown in Fig. 6 relies on transmission gates to connect the cell to a shared read/write bus during data transfer operations. All six data paths need to be driven during write operations, but only three paths need to be used during reads. If only three read paths are used, complementary write data can be generated locally with inverters for write operations. This structure can be used as a pipeline latch, and also in SRAM arrays. However, in SRAM arrays, bus capacitance must be minimized to prevent unwanted bit flips when the source cell is affected by a particle strike during a read operation. This is not always possible in larger arrays, necessitating the use of read buffers.

Fig. 7 depicts an alternate construction of the TPDICE cell with nmos write enable transistors and tri-state buffers providing read enable capability. Relying on nmos devices instead of transmission gates for write enable circuitry halves the transistor count of this portion of the circuit. However, noise margins are affected, especially when a transient pulse disrupts an input path. The use of tri-state read enable buffers ensures that the state of the latch cannot be erroneously flipped if a transient pulse affects a cell while it is reading data. This read enable approach is not as efficient as the previous design, but it works in all situations and is easy to implement.

During write operations, at least four nodes must be driven to proper values to achieve the desired functionality. If fewer nodes are driven, the write operation may not be able to overcome the initial state of the cell, as at least three of the nodes would resist the transition. For this reason, the approach adopted for the TPDICE design was to write to all six nodes through access transistors. NMOS transistors pass sufficient logic 1's for this setup, although the noise margins can be improved by using full transmission gates.

Particle strikes that do not originate in the destination memory (SETs) form one classification of transient pulse-related events that can affect TPDICE. If a transient pulse affects the source cell or combinational logic during a data transfer, the associated data path may be flipped to the incorrect value. This could affect two nodes of the destination cell during the write operation. However, these nodes are no longer affected after the write operation, allowing the other four nodes to restore the state of the latch during the recovery period. Fig. 8 depicts one possible SET-influenced write operation and the recovery process. The destination cell has an initial state of 101010, and the write value is 010101. The SET holds  $V_a$  high and  $V_b$  low during the write cycle. After the write enable signal is de-asserted, all nodes quickly resolve to desired values.

A second transient pulse-related event classification that can affect TPDICE designs is a particle strike to the destination cell (potential SEU) during data transfer. When this happens, two factors may combine to resist the write operation. First, the destination memory may initially be at the opposite state of the




Fig. 8. TPDICE SET recovery process.

write data. Second, the particle strike may hold two of the destination memory nodes at the incorrect value during the write operation. This effect may persist beyond the write operation, during the recovery process. The circuit can recover from this scenario if the write buffers have enough time to drive the four nodes that are not directly affected to their proper values. If this goal is achieved, the recovery process becomes analogous to a standard particle strike recovery that is not influenced by read or write operations.

The ability to bypass transient pulses comes from the use of three paths to transmit data and three inverter pairs in the core of the latch. An SEU impacting the destination cell during a write operation directly affects the logic value of the struck node, causes an adjacent node to become tri-stated, and induces conflict in the other adjacent node. The result is that the voltage of the struck node and one adjacent node is altered. Thus, at most one node out of  $V_a$ ,  $V_b$ , and  $V_c$  will be affected (similarly for  $/V_a$ ,  $/V_b$ , and  $/V_c$ ). Assuming the outputs are  $V_a$ ,  $V_b$ , and  $V_c$ , only one output (out of three) is altered during recovery. The case is similar for SETs, where nodes storing complementary values are directly affected, but only one output is altered. Because of this, subsequent cells can proceed with write operations before the affected cell has completed recovery. The cells can tolerate and bypass a transient pulse on one input, which is the case here, as only one output of the first cell was altered. This is extremely beneficial for performance, as it allows the TPDICE design to maintain a clock rate that is independent of the maximum SET width. In general, the recovery of an affected cell does not halt concurrent operations involving that cell.

Deadlock can occur in the TPDICE core if an SET affects two adjacent nodes and the destination cell begins at the logical state complementary to that of the write operation. Each node in the affected pair controls one transistor that drives the other node in the pair. If these drivers are responsible for restoring the nodes to their proper values but are switched off by the SET, then a lockout situation has occurred. Each node's recovery can begin only after the other node is restored to its proper state. Therefore, this situation must be avoided if proper functionality is to be achieved. This can be accomplished by offsetting each cell input from its complement input, as shown in Figs. 6 and 7.

With these SEU/SET tolerant characteristics in mind, the TPDICE circuit was designed with three independent data paths. Each path drives two non-adjacent nodes to prevent deadlock (i.e., the path passing through  $D_a$  and  $/D_a$  drives  $V_a$




Fig. 9. Simulation of TPDICE internal nodes during SET recovery process ( $H \rightarrow L$  transition).

and  $Vb$ , etc.). Inverters are used to produce the complementary signals at the input to each memory cell.

### B. TPDICE Simulation Results

Fig. 9 depicts an SET affecting the  $Da$  data path with buffers driving data and enable signals, replacing the ideal voltage sources used for Fig. 8. SETs affecting  $Db$  or  $Dc$  are not illustrated in this paper, as their effect is analogous to SETs on  $Da$  due to the symmetry of the TPDICE circuit. The initial state of the cell is 101010, and 010101 is the write value.  $Va$  and  $Vb$  are directly affected by the SET.  $Va$  is held high during the write operation, and  $Vb$  is held low. The other four nodes are pulled to the desired levels by the end of the write cycle. At the beginning of recovery,  $Va$  and  $Vb$  are driven strongly to appropriate values, conflict is induced in  $Vc$  and  $Vc$ , and  $Va$  and  $Vb$  are tri-stated. The restoration of  $Va$  and  $Vb$  removes the conflict from  $Vc$  and  $Vc$ , allowing the cell to quickly recover to the desired state. It is important to note that two of the three outputs ( $Vb$ ,  $Vc$ ) are at appropriate levels by the end of the write cycle. This allows subsequent operations to proceed in parallel with the recovery of the affected cell.

A simulation of an SET affecting  $Da$  during a 010101 to 101010 write transition is shown in Fig. 10. This case is the inverse of the previous write simulation. Again,  $Va$  and  $Vb$  are directly affected by the SET.  $Va$  is held low and  $Vb$  is held high during the write cycle. The other four nodes are driven to the desired values by the write buffers. After the write enable signal is deasserted,  $Va$  and  $Vb$  are driven to their proper values,  $Vc$  and  $Vc$  are tri-stated, and  $Va$  and  $Vb$  experience conflict. This initial condition easily recovers to the desired state. The nodes directly affected by the SET are strongly pulled towards proper values, quickly removing the side effects of the SET and allowing the cell to resolve properly. Again, two out of three outputs arrive at proper values before the end of the write operation. This allows the output data to be used immediately.

Another situation that must be considered occurs when a particle directly strikes the destination cell during a write operation (potential SEU). In this scenario, the effects of the particle strike may persist after the write operation is completed, as the disturbance is not cut off by write enable transistors. For example, consider the situation where  $Va$  is directly affected during a 101010 to 010101 write operation.  $Va$  is held high during the write,  $Vc$  experiences conflict, and all other nodes are driven to




Fig. 10. Simulation of TPDICE internal nodes during SET recovery process ( $L \rightarrow H$  transition).




Fig. 11. Simulation of TPDICE internal nodes during SEU recovery process ( $H \rightarrow L$  transition).

proper values. After the write operation is completed,  $Va$  continues to be affected by the particle strike, and  $Vc$  experiences greater conflict, as it is no longer driven by a write buffer. However, since all other nodes have been driven to desired values, the situation is essentially equivalent to that of a regular particle strike to memory (potential SEU). Because of this, the memory cell easily resolves to the desired state during the recovery period. A graphic depiction of this is shown in Fig. 11. Again, two of three outputs are driven to proper values before the end of the write procedure.

The complement to the previous simulation occurs when a potential SEU affects the destination cell during a 010101 to 101010 write procedure.  $Va$  is held low by the particle strike, and  $Va$  experiences conflict from this. The other four nodes are driven to desired values by the write buffers. At the end of the write cycle,  $Va$  continues to be affected by the transient pulse, and  $Va$  loses its write buffer support. Fortunately, this cell state is also equivalent to that of a basic particle strike to memory (potential SEU). Since this is true, the desired cell resolution is achieved by the end of the recovery period. Fig. 12 is a simulation of the internal node status during this scenario.

## VII. ANALYSIS OF RESULTS

In this section, simulation results are presented and analyzed for the TMR, basic SET-tolerant DICE, differential DICE, and TPDICE approaches described above. Each of these designs was simulated in an environment based off the diagram shown in Fig. 13. Inverters were used to generate realistic enable signals. Additionally, two inverters have been added to each data path to provide drive strength for write operations. These inverters account for buffers in SRAMs and logic in pipelines. However, the buffer/combinational logic configurations of real world systems




Fig. 12. Simulation of TPDICE internal nodes during SEU recovery process (L → H transition).

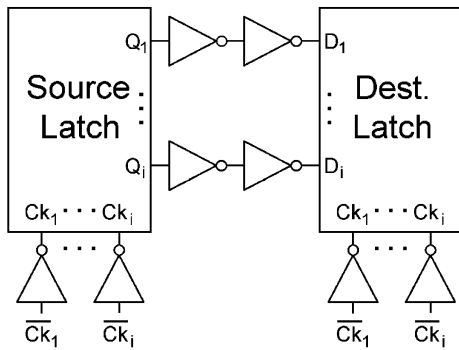


Fig. 13. Simulation setup used to compare the SET-tolerant latches.

TABLE II  
MINIMUM CLOCK PERIODS

|               | No SET | 200ps SET | 500ps SET | Bypass      |
|---------------|--------|-----------|-----------|-------------|
| TMR           | 628ps  | 749ps     | 749ps     | Data, Logic |
| SET-Tol. DICE | 515ps  | 1010ps    | 1310ps    | none        |
| Diff. DICE    | 348ps  | 473ps     | 480ps     | Data Only   |
| TPDICE        | 434ps  | 545ps     | 552ps     | Data, Logic |

are unique, and so delay, energy consumption, and complexity figures are ultimately dependent on the actual design chosen. For example, the power consumption penalty paid by a pipelined system using three independent data paths may be small if thin logic is required between pipeline registers. On the other hand, the penalty may be more substantial if deep logic stages are required.

Clock period, energy consumption, and circuit complexity statistics for all of the SET-tolerant approaches are presented below. The clock period simulation results are listed in Table II and shown graphically in Fig. 14. The left column of the table shows the minimum clock period attained for each approach with no transient pulses, while the other columns show the clock periods needed to tolerate 200 ps and 500 ps SETs. Table III presents energy consumption figures that are broken down into read memory, write memory, buffer, voting, and total contributions. Fig. 15 is a graphic depiction of the energy consumption results. Finally, Table IV depicts the circuit complexity of each approach, divided up into memory, voting/delay, logic and I/O

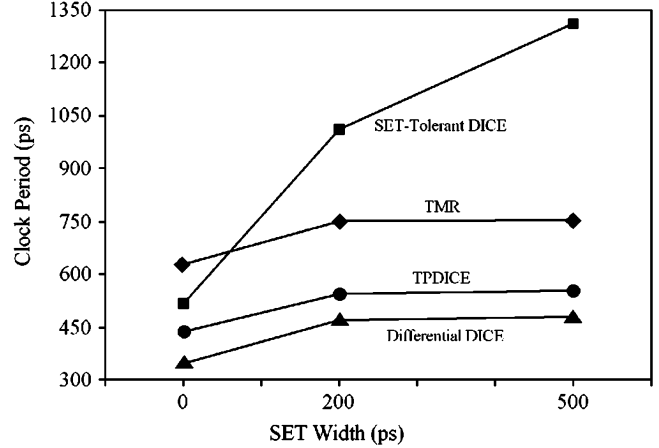


Fig. 14. Plot of clock period versus SET width. Note the quasi-linear dependence on SET width is only present in the SET-tolerant DICE circuit, which is the only approach that can not bypass transient pulses.

TABLE III  
ENERGY CONSUMPTION FIGURES

|               | Read    | Write   | Buffers | Voting  | Total    |
|---------------|---------|---------|---------|---------|----------|
| TMR           | 127.1fJ | 180.5fJ | 560.4fJ | 264.1fJ | 1132.1fJ |
| SET-Tol. DICE | 35.62fJ | 84.69fJ | 424.5fJ | n/a     | 544.81fJ |
| Diff. DICE    | 29.86fJ | 44.75fJ | 414.4fJ | n/a     | 489.01fJ |
| TPDICE        | 33.85fJ | 84.0fJ  | 459.7fJ | n/a     | 577.55fJ |




Fig. 15. Plot of energy consumption results. The total consumption of TMR is nearly double that of each DICE approach. Note that the energy consumption of each approach is dependent on the logic in the data paths (if any is used).

TABLE IV  
CIRCUIT COMPLEXITY QUANTIZATION

|               | Memory | Voting | Logic | I/O          |
|---------------|--------|--------|-------|--------------|
| TMR           | 18     | 36     | 3x    | 3 data, 3 ck |
| SET-Tol. DICE | 12     | n/a    | 2x    | 2 data, 2 ck |
| Diff. DICE    | 16     | n/a    | 4x    | 4 data, 4 ck |
| TPDICE        | 24     | n/a    | 3x    | 3 data, 3 ck |

categories. All simulations were performed in 0.18 μm CMOS technology.

These results show that while no approach dominates all comparison categories, TPDICE is the most balanced. TMR achieves reasonable performance (749 ps with a 500 ps SET), but has substantial costs in terms of energy consumption

(1132.1 fJ) and circuit complexity (18 memory and 36 voting transistors, as well as three independent data paths). The basic SET-tolerant DICE approach requires only two independent data paths and moderate energy consumption (544.81 fJ), but depends on a clock period that is proportional to the maximum SET width (515 ps with no upset, 1010 ps with 200 ps SET, 1310 ps with 500 ps SET). It is the only approach considered here that does not have the ability to bypass transient pulses. The fully-differential DICE design is well balanced, achieving high performance (473 ps clock period with a 200 ps SET, 480 ps with a 500 ps SET) and low energy consumption (489.01 fJ) at the cost of moderate size. However, it requires two independent differential data paths and four enable lines. Additionally, it loses bypass capability when logic is included in the data paths. Finally, TPDICE also provides good performance (545 ps clock period with a 200 ps SET, 552 ps with a 500 ps SET) while using three data paths. It requires 577.55 fJ of energy consumption. The TPDICE cell achieves full functionality in both data transfer and logic-based systems.

The approaches can be further classified based on the number of independent data paths (or copies of combinational logic) they require. The basic SET-tolerant DICE approach requires only two independent paths, the TMR and TPDICE approaches require three paths, and the fully-differential DICE approach requires two differential paths (four total transmitted signals). 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 transient pulses are 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  $\Delta\text{SET}$ , the length of the longest possible SET. On the other hand, schemes that rely on three or more paths can bypass transient pulses, and so they do not have to pause the system until these pulses dissipate. This means that the clock period is not dependent on  $\Delta\text{SET}$ , making the system scalable to situations that must tolerate wide SETs.

## VIII. CONCLUDING REMARKS

In this paper, we introduced and evaluated two novel SET-tolerant approaches, namely fully-differential DICE and TPDICE. Existing schemes (TMR and basic SET-tolerant DICE) have been compared to the new approaches. From our evaluations we have observed that the ability to bypass SEU/SET transient pulses drastically improves performance. Some studies have measured SET width to be up to 2 ns, a figure that is largely unaffected by reductions in feature size [12]–[15]. This puts a very restrictive cap on performance for approaches that do not bypass transient pulses. Most current approaches are not bypass-capable, including delay-based and dual-rail logic designs [5]–[10].

The following conclusions can be drawn from the comprehensive evaluation of the four SET-tolerant approaches presented in this paper:

- *TMR designs can bypass transient pulses, but they require substantial overhead due to the need for external SET-tolerant voting circuitry.* High-performance SET-tolerant TMR requires substantially 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. External voting circuitry must be hardened against SETs. In addition to this, feedback must be utilized to correct a corrupted source memory cell in TMR-based SRAM arrays.
- *Basic SET-tolerant DICE designs do not bypass transient pulses.* Basic DICE structures featuring delay-filtered inputs or dual-rail logic possess relatively low complexity, but they must pause write operations while a transient pulse is detected on one of their inputs. This directly affects performance. The other approaches considered in this paper bypass transient pulses, so they do not need to pause when a pulse is detected.
- *Fully-differential DICE structures can bypass transient pulses in SRAM systems.* Fully-differential DICE approaches utilize the basic four-node DICE core. They only have the ability to bypass pulses in pure data transfer systems without combinational logic. Additionally, they require two independent differential data paths and four enable lines.
- *TPDICE is the most balanced approach with SEU/SET bypass capability.* TPDICE is similar to TMR, in that it requires three data paths and three times the memory circuitry of a traditional SRAM cell. However, it possesses some very important advantages over TMR. First, size is reduced, as voting is directly integrated into the TPDICE latch structure. Second, delay is minimized due to the lack of external voting. Finally, transient pulses are corrected as they occur, removing the need for feedback circuitry to correct a corrupted source memory cell. Additionally, TPDICE can be used in systems with combinational logic, unlike fully-differential DICE.

A number of factors should be taken into account when selecting an SET-tolerant approach for a particular application. Clearly, each approach has its strengths and weaknesses, and so the optimum design depends on the demands of the application.

## REFERENCES

- [1] E. L. Petersen, P. Shapiro, J. H. Adams, and E. A. Burke, "Calculations of cosmic ray induced soft upsets and scaling in VLSI devices," *IEEE Trans. Nucl. Sci.*, vol. 29, no. 6, pp. 2055–2063, Dec. 1982.
- [2] A. M. Chugg, "Ionising radiation effects: A vital issue for semiconductor electronics," *Eng. Sci. Edu. J.*, vol. 3, pp. 123–130, Jun. 1994.
- [3] J. J. Wang, R. B. Katz, J. S. Sun, B. E. Cronquist, J. L. McCollum, T. M. Speers, and W. C. Plants, "SRAM based re-programmable FPGA for space applications," *IEEE Trans. Nucl. Sci.*, vol. 46, no. 6, pp. 1728–1735, Dec. 1999.
- [4] P. E. Dodd and L. W. Massengill, "Basic mechanism and modeling of SEU in digital microelectronics," *IEEE Trans. Nucl. Sci.*, vol. 50, no. 3, pp. 583–602, Jun. 2003.
- [5] D. R. Blum, M. J. Myjak, and J. G. Delgado-Frias, "Enhanced fault-tolerant data latches for deep submicron CMOS," in *Proc. 2005 Int. Conf. Computer Design*, Las Vegas, NV, Jun. 2005, pp. 28–34.
- [6] D. R. Blum and J. G. Delgado-Frias, "Comparison of SET-resistant approaches for memory-based architectures," in *Proc. 12th NASA Symp. VLSI Design*, Coeur d'Alene, ID, Oct. 2005.

- [7] D. G. Mavis and P. H. Eaton, "Soft error rate mitigation techniques for modern microcircuits," in *Proc. 40th Annu. Int. Reliability Physics Symp.*, Dallas, TX, Apr. 2002, pp. 216–225.
- [8] P. Mongkolkachit and B. Bhuvu, "Design technique for mitigation of alpha-particle-induced single-event transients in combinational logic," *IEEE Trans. Device Mater. Reliab.*, vol. 3, no. 3, pp. 89–92, Sep. 2003.
- [9] K. J. Hass, J. W. Gambles, B. Walker, and M. Zampaglione, "Mitigating single event upsets from combinational logic," in *Proc. 7th NASA Symp. VLSI Design*, Oct. 1998, pp. 4.1.1–4.1.10.
- [10] J. W. Gambles, K. J. Hass, and K. B. Cameron, "Apparatus for and method of eliminating single event upsets in combinational logic," U.S. patent 6,326,809, Dec. 4, 2001.
- [11] D. G. Mavis, "Single event transient phenomena—Challenges and solutions," presented at the Microelectronics Reliability and Qualification Workshop, Dec. 2002.
- [12] M. Gadlage, R. Schrimpf, J. Benedetto, P. Eaton, D. Mavis, M. Sibley, K. Avery, and T. Turflinger, "Single event transient pulse widths in digital microcircuits," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 6, pp. 3285–3290, Dec. 2004.
- [13] K. Castellani-Coulie, J. M. Palau, G. Hubert, M. C. Calvet, P. E. Dodd, and F. Sexton, "Various SEU conditions in SRAM studied by 3-D device simulation," *IEEE Trans. Nucl. Sci.*, vol. 48, no. 6, pp. 1931–1936, Dec. 2001.
- [14] L. M. Cohn, "Single-event effects in advanced digital and analog microelectronics," presented at the Microelectronics Reliability and Qualification Workshop, Dec. 2004.
- [15] P. Eaton, J. Benedetto, D. Mavis, K. Avery, M. Sibley, M. Gadlage, and T. Turflinger, "Single event transient pulsewidth measurements using a variable temporal latch technique," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 6, pp. 3365–3368, Dec. 2004.
- [16] P. E. Dodd, M. R. Shaneyfelt, and F. W. Sexton, "Charge collection and SEU from angled ion strikes," *IEEE Trans. Nucl. Sci.*, vol. 44, no. 6, pp. 2256–2265, Dec. 1997.
- [17] P. E. Dodd, F. W. Sexton, and P. S. Winokur, "Three-dimensional simulation of charge collection and multiple-bit upset in Si devices," *IEEE Trans. Nucl. Sci.*, vol. 41, no. 6, pp. 2005–2017, Dec. 1994.
- [18] P. E. Dodd and F. W. Sexton, "Critical charge concepts for CMOS SRAMs," *IEEE Trans. Nucl. Sci.*, vol. 42, no. 6, pp. 1764–1771, Dec. 1995.
- [19] P. E. Dodd, O. Musseau, M. R. Shaneyfelt, F. W. Sexton, C. D'hose, G. L. Hash, M. Martinez, R. A. Loemker, J.-L. Leray, and P. S. Winokur, "Impact of ion energy on single-event upset," *IEEE Trans. Nucl. Sci.*, vol. 45, no. 6, pp. 2483–2491, Dec. 1998.
- [20] P. E. Dodd, F. W. Sexton, G. L. Hash, M. R. Shaneyfelt, B. L. Draper, A. J. Farino, and R. S. Flores, "Impact of technology trends on SEU in CMOS SRAMs," *IEEE Trans. Nucl. Sci.*, vol. 43, no. 6, pp. 2797–2804, Dec. 1996.
- [21] K. J. Hass and J. W. Gambles, "Single event transients in deep submicron CMOS," in *Proc. 42nd Midwest Symp. Circuits and Systems*, Las Cruces, NM, Aug. 1999, vol. 1, pp. 122–125.
- [22] M. P. Baze and S. P. Buchner, "Attenuation of single event induced pulses in CMOS combinational logic," *IEEE Trans. Nucl. Sci.*, vol. 44, no. 6, pp. 2217–2223, Dec. 1997.
- [23] P. Shivakumar, M. Kistler, S. W. Keckler, D. Burger, and L. Alvisi, "Modeling the effect of technology trends on soft error rate of combinational logic," in *Proc. 2002 Int. Conf. Dependable Systems and Networks*, Bethesda, MD, Jun. 2002, pp. 389–398.
- [24] J. Benedetto, P. Eaton, K. Avery, D. Mavis, M. Gadlage, T. Turflinger, P. E. Dodd, and G. Vizkelethy, "Heavy ion-induced digital single-event transients in deep submicron processes," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 6, pp. 3480–3485, Dec. 2004.
- [25] T. Calin, M. Nicolaidis, and R. Velazco, "Upset hardened memory design for submicron CMOS technology," *IEEE Trans. Nucl. Sci.*, vol. 43, no. 6, pp. 2874–2878, Dec. 1996.
- [26] Q. Shi and G. Maki, "New design techniques for SEU immune circuits," in *Proc. 9th NASA Symp. VLSI Design*, Albuquerque, NM, Nov. 2000, pp. 7.4.1–7.4.16.