Enhanced Fault-Tolerant Data Latches for Deep Submicron CMOS

Daniel R. Blum, Mitchell J. Myjak, and José G. Delgado-Frias
School of Electrical Engineering and Computer Science
Washington State University
Pullman, WA, USA

Abstract

CMOS data latches used in critical applications must be immune to soft errors such as single event upsets. Existing designs protect the stored data against errors in the internal nodes, but may be vulnerable to transient faults in the control and data lines. The problem becomes more severe as feature sizes decrease. In this paper, we enhance the Dual Interlocked Storage Cell (DICE) to withstand soft errors at any node. We expand our scheme to encompass five fault-tolerant memory cells: one optimized for pipeline latches, and the others for SRAM. The designs have been verified through extensive layout simulations in 180-nm CMOS. Compared to the original DICE, the proposed cells can withstand a broader class of transient faults, but consume more energy during read and write operations. We explore the tradeoff between energy consumption and the number of redundant control lines required.

Keywords: Data latch, DICE, fault tolerant, radiation hardening, SEU, SRAM.

1. Introduction

Integrated circuits operating in the presence of radiation are susceptible to several types of errors, including latch-up, burn-out, oxide charging, reduction of carrier lifetimes, and single event upsets [1]. The first four types are permanent faults that must be addressed with appropriate process enhancements. However, single event upsets (SEU) cannot be eliminated at the process level. These errors are caused by the accumulation of charge from energized particle strikes. As a particle passes through an IC, it loses energy and frees electron-hole pairs along its path. This charge can be collected at a reverse-biased junction, such as the drain of a MOSFET in cutoff, via electrical drift and diffusion [2]. If enough charge is collected, the voltage at the corresponding node changes significantly. This effect does not cause any damage, and is usually temporary, as the charge dissipates through the surrounding circuitry. However, the upset can cause long-term disruption if the circuitry cannot recover its original state. In particular, the spurious charge may overwrite the logic value stored in a memory cell, and the resulting data corruption could cause the entire system to fail.

In deep submicron technologies, SEU 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 [3]. This effect is known as a single event transient (SET) [4]. If the disturbance reaches the write enable or data lines of a memory cell, the stored data could be overwritten. SET are becoming more of a problem as IC feature sizes and clock periods shrink. Reduced feature sizes create smaller node capacitances, which results 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 [5].

A variety of techniques have been proposed to mitigate SEU. Triple modular redundancy [6] and cross-parity checking [7] both offer system-level error correction but require substantial spatial and temporal overhead. To reduce the probability of SEU, process enhancements can hinder the charge collection process and resistors within memory cells introduce delays within feedback paths [2]. These approaches become less effective with smaller technologies. However, circuit-level schemes that utilize redundancy can provide immunity to transient errors regardless of feature sizes.

Researchers have designed several data latches that recover from transient faults by using redundant circuitry [8]–[11]. Of these designs, the Dual Interlocked Storage Cell (DICE) of Calin, Nicolaidis, and Velazco [11] is perhaps the most versatile. However, these designs have generally ignored the effects of SET on the control and data circuitry, relying instead on the capacitance of these lines to prevent large glitches. This assumption is no longer valid for deep submicron CMOS.

In previous work [12], we enhanced the DICE to tolerate faults in any internal or external node. We then optimized the design to create pipeline latches and SRAM cells. Layout simulations in 250-nm technology demonstrated that the designs achieve complete immunity to SEU and SET. In this paper, we repeat the verification of the data latches using 180-nm CMOS.

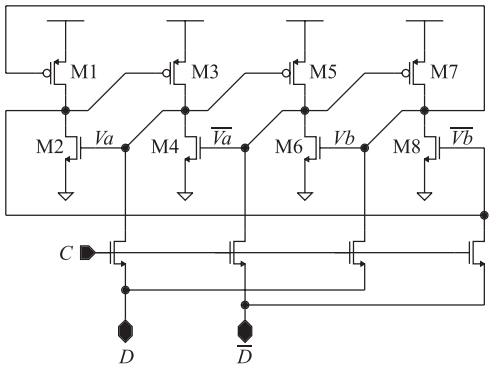

Figure 1. Dual Interlocked Storage Cell (DICE) [11].

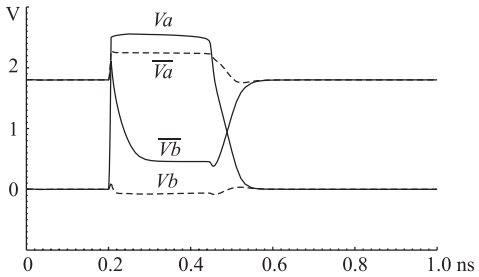

Figure 2. DICE recovers from injected charge on Va.

We then propose three additional SRAM cells that consume less energy for read and write operations, yet maintain the same degree of fault tolerance. We conclude with an analysis of the data latches in terms of area, number of control lines, and power consumption.

2. Dual Interlocked Storage Cell

Figure 1 illustrates the circuit schematic of the DICE [11]. The latch contains four inverters interlocked together so that the p-transistor and n-transistor are controlled by separate nodes. The four internal nodes Va, $\overline{Va}$, Vb, and $\overline{Vb}$ always contain alternating logic levels. Values of 0101 represent logic 0, and values of 1010 represent logic 1. Read and write operations are performed by asserting the enable input C and accessing the data on the bidirectional D and $\overline{D}$ lines.

Figure 2 depicts the simulated response of the latch to a transient fault at an internal node. The four internal nodes initially have logic levels 0101, where $V_{DD} = 1.8$ V. At $t = 0.2$ ns, charge is injected into Va. The error is modeled as a triangular current pulse with amplitude 50 mA, rise time 0.05 ns, and fall time 0.20 ns. The voltage at Va quickly rises above V_{DD}. Transistor M2 then turns on and discharges $\overline{Vb}$. However, $\overline{Va}$ and Vb remain unaffected. Thus, M1 and M4 can restore the latch to its original state by $t = 0.6$ ns. A positive glitch at Vb or a negative glitch at $\overline{Va}$ or $\overline{Vb}$ would produce similar results.

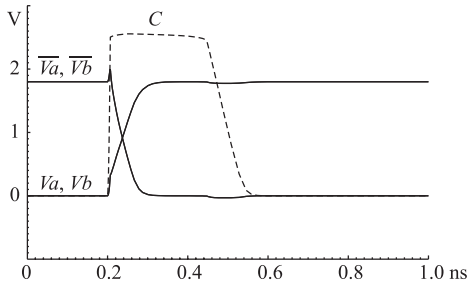

Figure 3. SET on C causes DICE to fail.

The advantages of the DICE include low transistor count, no static power consumption, and rapid recovery time. In addition, the latch does not depend on specific transistor ratios or process parameters to function properly. However, the latch is vulnerable to SET in the control and data lines. As shown in Figure 3, a positive glitch at C briefly enables the four pass transistors used for reading and writing. The latch changes state if the system is writing logic 1 to another latch sharing the same data lines. Likewise, a glitch in D or $\overline{D}$ at the end of a write operation may cause the wrong value to be stored in the latch.

3. Enhanced memory cell

Consider the circuit in Figure 4. The two bidirectional data lines of the original DICE are separated into two write lines, Da and Db, and two read lines, $\overline{Qa}$ and $\overline{Qb}$. These lines connect to the internal nodes of the cell via transmission gates. Each transmission gate is driven by two control signals, so the cell requires four write enable signals (Wa, $\overline{Wa}$, Wb, and $\overline{Wb}$) as well as four read enable signals (Ra, $/Ra$, Rb, and $/Rb$).

This enhanced memory cell requires many more control signals than the DICE, but can withstand a transient fault at any internal or external node. All the inputs and outputs have duplex redundancy, as denoted by a and b suffixes in the figure. Signals with opposite suffixes, such as Wa and Wb, should be generated by independent logic blocks, so that an isolated fault anywhere on the chip will not affect both signals simultaneously. The latch is designed to only change state when both sets of input signals are consistent.

For example, suppose an SET occurs in the logic that generates the write enable signals, so that Wa experiences a positive glitch while $\overline{Wa}$ experiences a negative glitch. If the cell is storing logic 0 but the current value on Da is logic 1, Va will undergo a positive glitch. However, the core of the DICE protects against these types of faults. Therefore, the latch will revert to its original state once the SET has run its course. The simulated results in Figure 5 demonstrate that the latch behaves as predicted. One can also show

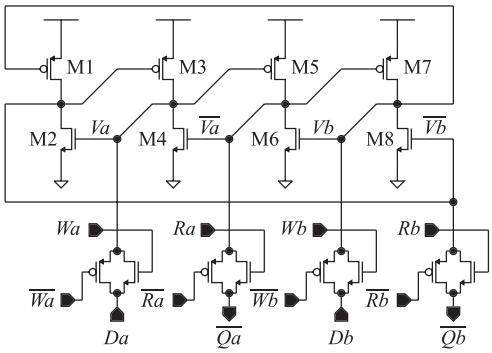

Figure 4. Enhanced memory cell.

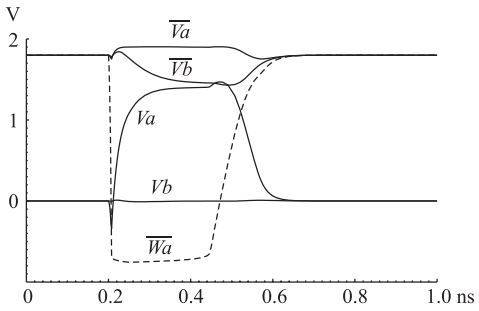

Figure 5. Enhanced design recovers from SET on Wa and $/Wa$.

that a transient fault in any enable signal or data line will affect only one of the internal nodes of the latch, and thus not cause an SEU.

Although the enhanced memory cell is immune to isolated transient faults, normal read and write operations may cause problems under certain conditions. The original DICE writes to all four internal nodes at once, whereas the enhanced memory cell only writes to nodes Va and Vb. Because of this, the four transistors driving Va and Vb (M3, M4, M7, and M8) must be made about one-half as strong as the pass transistors on the write lines; otherwise, normal write operations will not be able to change the state of the cell. This adjustment increases the susceptibility of the cell to injected charge at Va and Vb. If a glitch occurs at one of these nodes while no operation is being performed, the circuit will recover normally. However, if the glitch occurs at the beginning of a read operation, and the $\overline{Qa}$ and $\overline{Qb}$ lines are charged to the opposite logic levels as the internal nodes, the latch will change state.

Figure 6 depicts this SEU in the memory cell. The internal nodes initially have logic values 0101, but a transient spike pulls up Va at $t = 0.2$ ns. The $\overline{Qa}$ and $\overline{Qb}$ lines are both discharged, and pull down $\overline{Va}$ and $\overline{Vb}$ after the read enable signals are asserted at $t = 0.35$ ns. This combination of events causes the cell to flip state, as the internal nodes settle to logic values 1010 by the end of the simulation.

Rectifying this problem would require additional

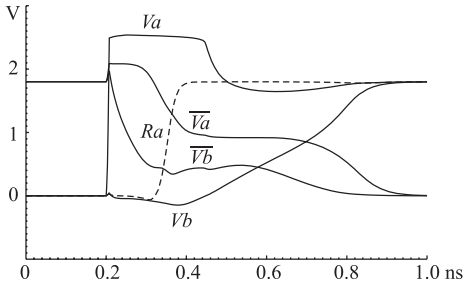

Figure 6. SEU caused by residual charge on data lines.

transistors in the latch. Rather than discuss possible solutions, we create several variants of the basic design that are optimized for particular applications. As demonstrated in the next section, these optimizations eliminate the likelihood of this type of SEU, while simultaneously removing excess input signals.

4. Optimized designs

In this section, we present five novel fault-tolerant memory cells: the first optimized for pipeline latches, and the others for SRAM. Layout simulations in 180-nm CMOS demonstrate that the designs achieve complete immunity to both SEU and SET in the control logic. The first two designs were proposed in previous work [12], whereas the last three are improvements on the SRAM cell that reduce the energy consumption of reads and writes.

4.1 Pipeline latch

The enhanced memory cell in Figure 4 fails when stored charge in $\overline{Qa}$ and $\overline{Qb}$ combines with injected charge during a read operation to affect the values of three internal nodes. One way to correct this problem is to replace the pass transistors on the read lines with inverters, as shown in Figure 7. This change prevents $\overline{Qa}$ and $\overline{Qb}$ from affecting the internal nodes during read operations. The resulting circuit is ideally suited for pipeline latches, since it continuously outputs the stored data on the read lines. One can also create an edge-triggered flip-flop by cascading two pipeline latches.

A write operation to the pipeline latch can only occur if both sets of write enable signals are asserted, and both data lines are set to the same logic value. As shown in Figure 8, a write operation completes in 0.3 ns, assuming the data line drivers are at least twice the size as the internal latch transistors. Figure 9 demonstrates that no write operation occurs if an SET spuriously activates Wa and $\overline{Wa}$. In this simulation, the latch is storing logic 0 but the data lines are being driven to the opposite logic level. The voltage at Va temporarily moves up to 1.4 V, but quickly returns to ground when the fault dissipates. Notice that $\overline{Va}$ and

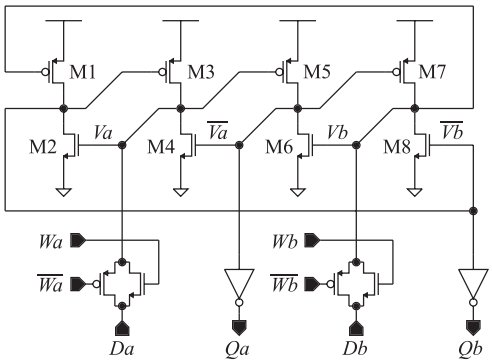

Figure 7. SEU-immune pipeline latch.

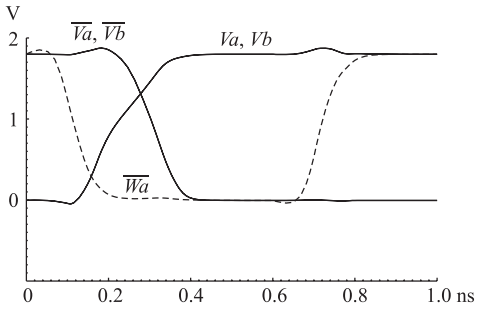

Figure 8. Normal write operation to pipeline latch.

$\overline{Vb}$ remain essentially unaffected, so neither Qa nor Qb changes during this event.

4.2 SRAM cell

Although the memory cell just described works well as a pipeline latch, it may not be appropriate for all situations. For example, SRAM typically uses common data lines for both read and write operations. For this reason, we have created another derivative of the enhanced memory cell in Figure 4 that is optimized for SRAM. As shown in Figure 10, the design uses twelve interface transistors connected to the bidirectional data lines Da, $\overline{Da}$, Db, and $\overline{Db}$. Inputs Ra, Rb, $\overline{Wa}$, and $\overline{Wb}$ enable the cell for reading and writing. To function correctly, the transistors that perform write operations must be made about twice as strong as

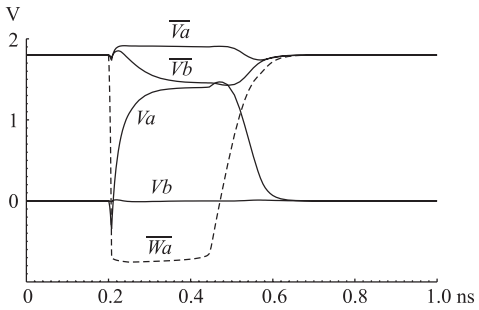

Figure 9. Pipeline latch recovers from SET on Wa and $\overline{Wa}$.

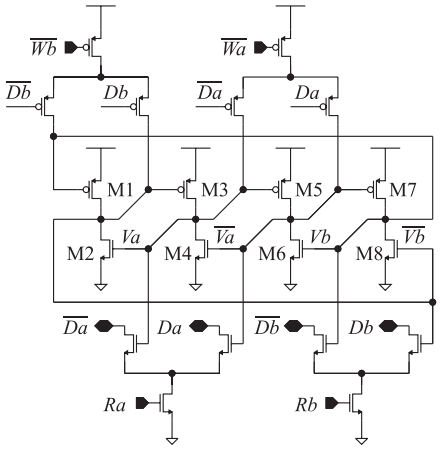

Figure 10. SEU-immune SRAM cell.

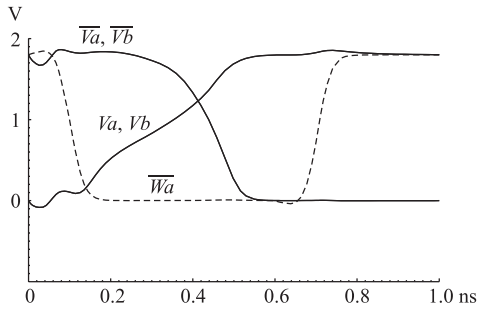

Figure 11. Write operation to SRAM cell.

internal transistors M1 through M8.

To read from the SRAM cell, the system first precharges all data lines to V_{DD}. The read enable signals Ra and Rb are then asserted. Assuming the cell is storing logic 0, $\overline{Va}$ and $\overline{Vb}$ will be at logic 1. The n-transistors on the bottom will then discharge Da and Db to ground. These transistors isolate the data lines from the internal nodes, preventing latent capacitance from feeding back into the cell during read operations. If the cell is storing logic 1, $\overline{Da}$ and $\overline{Db}$ will be discharged instead. Normally, the complementary D and $\overline{D}$ lines should never be low simultaneously.

Figure 11 depicts a normal write operation to the cell. Here, the cell initially stores logic 0 and the system writes logic 1. Thus, $\overline{Da}$ and $\overline{Db}$ will be discharged to ground. The write enable signals $\overline{Wa}$ and $\overline{Wb}$ are then asserted. The interface transistors on the top then pull up Va and Vb. Although only two internal nodes are driven, the cell changes state within 0.4 ns.

As with the pipeline latch, the SRAM cell does not complete a write operation if only one of the write enable signals is asserted. In Figure 12, an SET disrupts the value of $\overline{Wa}$ when no write operation should be in progress. The internal nodes have logic levels of 0101 but the data lines are set to 1010. When

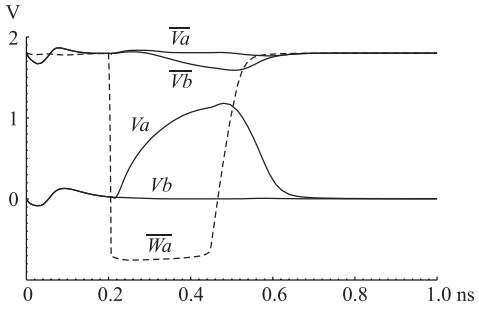

Figure 12. SRAM cell recovers from SET on $\overline{Wa}$.

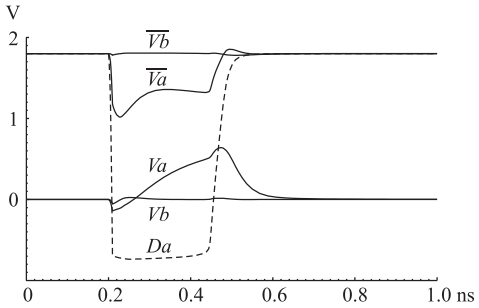

Figure 13. SRAM cell recovers from SET on Da.

the SET occurs, the p-transistors on top pull up Va. However, no other nodes are affected and the fault quickly dissipates.

An undesired situation can occur if two complementary data lines are both discharged to ground. For example, in Figure 13 the cell and data lines have the same initial logic levels as the previous simulation. The system performs a read operation on some other cell in the memory, so that $\overline{Da}$ and $\overline{Db}$ are pulled down to ground. A negative glitch then discharges Da as well. The p-transistors driven by these data lines turn on and provide a path between Va and $\overline{Va}$. However, the stored value will not be affected if the total upset time is fairly brief. Note that the SRAM cell by itself cannot restore Da to V_{DD} due to the single n-transistor driving the line. To correct this problem, weak cross-coupled p-transistors should be placed on each pair of data lines.

Although not detailed here, one can verify that a transient fault at any node will not cause the latch to malfunction. Faults can even occur during read and write operations, as demonstrated in the previous simulation, as long as the a and b signal paths are not affected simultaneously.

4.3 Reduced-energy SRAM cells

Power consumption is a growing concern in modern integrated circuit designs. Small embedded processors require low power consumption to maintain battery life. Large processors also need to minimize power dissipation to prevent overheating due to rapidly increasing transistor densities. Although the proposed SRAM cell achieves complete immunity to SEU and SET, write operations consume appreciable power since the write circuitry only drives two internal nodes at a time. Fortunately, the design of the SRAM cell has some flexibility. Buffering the internal nodes from the data lines is not necessary for write operations. However, buffering must be maintained for read operations to prevent faults caused by charge sharing coupled with particle strikes.

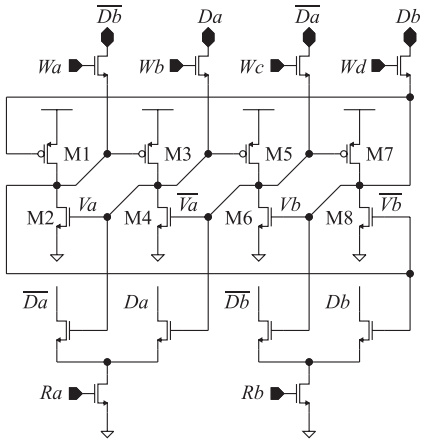

Figure 14. Energy-aware SRAM cell with four write transistors.

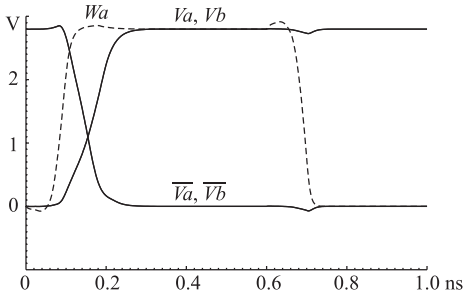

Figure 15. Write operation through four n-transistors.

Figure 14 shows a version of the SRAM cell that uses four n-transistors for the write circuitry. The four write enable lines (Wa, Wb, Wc, Wd) must be generated by independent blocks so that an SET in one path does not affect the other paths. Although this change increases the overhead of the control logic, write operations require less energy since all four internal nodes are driven simultaneously. Figure 15 demonstrates that a write operation completes in 0.2 ns. This write circuitry can tolerate an SET affecting any one of its data or write enable lines. Because of redundancy, such an SET would only impact one internal node. Figure 16 shows the SRAM cell recovering from an SET affecting Wa.

Another alternative is to replace the write circuitry with transmission gates. A schematic of this configura-

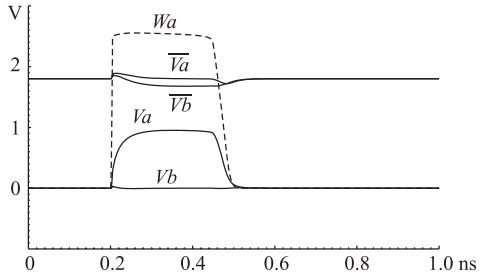

Figure 16. Energy-aware SRAM cell recovers from SET on Wa.

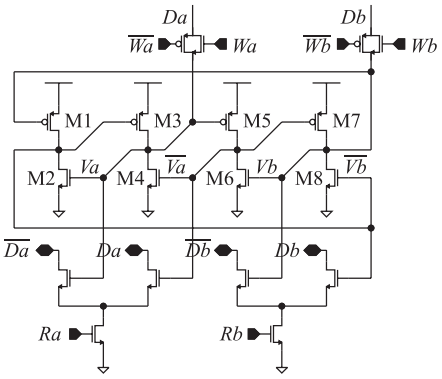

Figure 17. Energy-aware SRAM cell with transmission gates.

tion appears in Figure 17. This cell requires four write enable lines but only two redundant paths. Compared to the original SRAM cell, using transmission gates improves the speed of write operations, especially if transistors M3, M4, M7, and M8 are reduced in size. Figure 15 depicts a write operation. The response of the cell to an SET at Wa and $\overline{Wa}$ is similar to that of the pipeline latch described previously.

A further modification to the SRAM cell is replace the read circuitry with tri-state buffers, as shown in Figure 19. This change eliminates the need for charging the data lines to V_{DD} before read operations, but increases the number of control signals to eight. During read operations, the inverters drive the data lines while isolating the internal nodes from the latent capacitance. In the idle state, the inverters are disconnected, allowing other cells to drive the data lines.

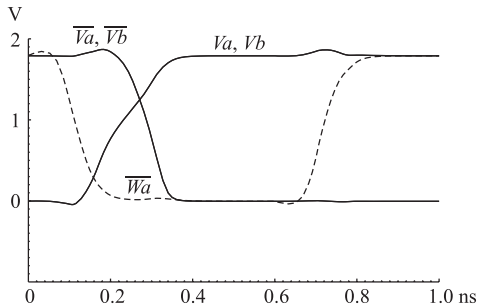

Figure 18. Write operation through transmission gates.

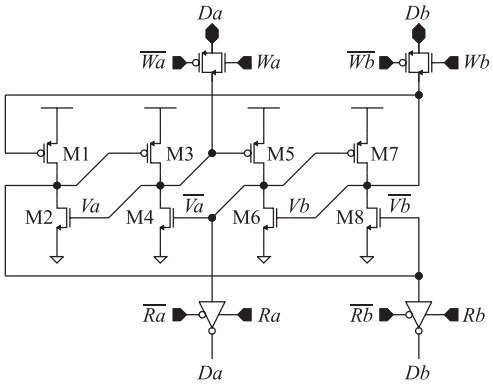

Figure 19. Energy-aware SRAM cell with tri-state buffers.

Table 1. Comparison of fault-tolerant data latches

Design	Ctrl	Data	Fault tolerance
DICE	1	2	Internal nodes
Enhanced	8	4	Problem with reads
Pipe	4	4	Complete SEU/SET
SRAM	4	4	Complete SEU/SET
(four n-transistors)	6	4	Complete SEU/SET
(transmission gates)	6	4	Complete SEU/SET
(tri-state buffers)	8	2	Complete SEU/SET

5. Analysis

Table 1 compares the fault-tolerant memory elements in terms of control lines, data lines, and degree of fault tolerance. All designs retain the minimal static power consumption and rapid recovery time of the original DICE. The pipeline latch and SRAM cells incur some overhead insulating the stored data against errors in the external circuitry. However, these designs are appropriate for applications that must balance high performance with high reliability. The necessity to protect memory elements against external errors will become more apparent as technologies scale down and capacitances continue to decrease.

Table 2 contains the power dissipation statistics for each cell, expressed as the energy consumption for read and write operations. The two numbers in each column represent the energy consumption of the memory cell and the external drivers, respectively. During read operations, the memory cell consumption results from sourcing current, while the driver consumption comes from switching and driving the data lines. During write operations, short circuit currents induced by transistor switching produce the majority of the internal cell consumption. The driver consumption is dominated by the sourcing of current to the memory cell load.

The energy dissipation for write operations varies widely between cells. The original DICE dissipates the least energy because it has the simplest design. All four internal nodes are driven differentially, reducing the load on the driver and the duration of short circuit

Table 2. Energy consumption of fault-tolerant data latches

Design	Writes (fJ)		Reads (fJ)	
	cell	drivers	cell	drivers
DICE	83.71 /	98.49	21.69 /	88.65
Enhanced	61.39 /	359.3	103.7 /	132.0
Pipeline	74.06 /	339.7	(n/a)	
SRAM	414.1 /	135.2	55.73 /	95.51
(four n-transistors)	74.60 /	152.4	15.19 /	86.50
(transmission gates)	49.10 /	341.7	21.73 /	95.24
(tri-state buffers)	77.90 /	349.9	87.02 /	143.9

currents inside of the cell. The next lowest dissipation comes from the SRAM cell with four n-transistors, which also has differential write capability. The original SRAM cell uses significant internal power because the driving capability is incorporated within the cell. In addition, the driving current passes through two p-transistors in series into only two of the internal nodes. Finally, the remaining cells use transmission gates for write operations, placing a considerable burden on the driving circuitry. The transmission gates require complementary enable inputs and only drive two of the internal nodes.

In general, the energy dissipation for read operations shows more similarity. The best designs are the original DICE, the SRAM cell with four n-transistors, and the SRAM cell with transmission gates. In these cells, the read enable circuitry consists of one n-transistor per line, resulting in low resistance. The original SRAM cell has higher internal consumption due to the additional loading of the write circuitry. The SRAM cell with tri-state buffers requires high energy during reads because of the high capacitive loading of the tri-state inverter inputs. Energy consumption is also an issue for the redundant cell because the capacitance of the transmission gates increases the load on the internal cell nodes. Finally, the pipeline latch does not have a specific read operation, as it always outputs the internal state of the cell.

6. Conclusion

In this paper, we have presented several CMOS memory elements that improve the fault tolerance of the DICE. The first design requires more control signals than the original cell and can experience an SEU from stored charge on the read lines. We then optimized the design to create a pipeline latch and an SRAM cell that can recover from transient faults at any node. In our scheme, all input and output signals are generated from two identical, independent paths so that faults cannot appear in both paths at once. Compared to triple modular redundancy, one less subsystem is required. Compared to cross-parity schemes, the system can recover from multiple independent errors without requiring periodic scans through the whole memory.

The energy consumption results showed that the original DICE and the SRAM cell with four n-transistors consumed the least energy during read and write operations. Both circuits achieved this result using differential write circuitry and n-transistor read circuitry. The SRAM cell has the advantage of SET immunity, along with data buffering during read operations. However, it does require more area and control lines.

Acknowledgment

M. Myjak is supported by a graduate fellowship from the U.S. Department of Homeland Security (DHS). The DHS Scholarship and Fellowship Program is administered by the Oak Ridge Institute for Science and Education (ORISE) for DHS through an interagency agreement with the U.S. Department of Energy (DOE). ORISE is managed by Oak Ridge Associated Universities under DOE contract number DE-AC05-00OR22750. All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORISE.

References

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

[2] P.E. Dodd, M.R. Shaneyfelt, and F.W. Sexon, "Charge collection and SEU from angled ion strikes," IEEE Trans. Nuclear Science, vol. 44, no. 6, pp. 2256-2265, Dec. 1997.

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

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

[5] 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.

[6] S. D'Angelo et al., "Fault-tolerant voting mechanism and recovery scheme for TMR FPGA-based systems," Proc. 1998 International Symp. on Defect and Fault Tolerance in VLSI Systems, pp. 233-240, Nov. 1998.

[7] M. Pflanz et al., "On-line error detection and correction in storage elements with cross-parity check," Proc. 8th IEEE International On-Line Testing Workshop, pp. 69–73, 2002.

[8] L.R. Rockett, Jr., "An SEU-hardened CMOS data latch design," IEEE Trans. Nuclear Science, vol. 35, pp. 1682–1687, Dec. 1988.

[9] N. Liu and S. Whitaker, "Low power SEU immune CMOS memory circuits," IEEE Trans. Nuclear Science, vol. 39, pp. 1679–1684, Dec. 1992.

[10] R. Velazco et al., "Two CMOS memory cells suitable for the design of SEU-tolerant VLSI circuits," IEEE Trans. Nuclear Science, vol. 41, pp. 2229–2234, Dec. 1994.

[11] T. Calin, M. Nicolaidis, and R. Velazco, "Upset hardened memory design for submicron CMOS technology," IEEE Trans. Nuclear Science, vol. 43, pp. 2874–2878, Dec. 1996.

[12] M.J. Myjak, D.R. Blum, and J.G. Delgado-Frias, "Enhanced fault-tolerant CMOS memory elements," Proc. 2004 IEEE International Midwest Symp. on Circuits and Systems, Hiroshima, Japan, Jul. 2004.