

HARDENED BY DESIGN APPROACHES FOR MITIGATING TRANSIENT
FAULTS IN MEMORY-BASED SYSTEMS

by

DANIEL RYAN BLUM

A dissertation submitted in partial fulfillment of
the requirements for the degree of

DOCTOR OF PHILOSOPHY

WASHINGTON STATE UNIVERSITY
School of Electrical Engineering and Computer Science

MAY 2007

To the faculty of Washington State University:

The members of the Committee appointed to examine the dissertation of DANIEL RYAN BLUM find it satisfactory and recommend that it be accepted.

Chair

ACKNOWLEDGEMENT

The author would like to extend appreciation to all of the sources of financial support that have contributed to the work contained in this dissertation. These sources include a Teaching Assistantship from the WSU School of EECS, a Research Assistantship from the WSU School of EECS Boeing Centennial Endowed Chair, a Harold P. Curtis fellowship, and a WSU Graduate School Scholars Fellowship. Without this support, I wouldn't have had the opportunity to pursue studies at WSU.

Appreciation is due to my advisor Dr. José Delgado-Frias for his guidance during my Ph.D. studies. Dr. Delgado has gone well beyond the call of duty in his advising, and this has largely influenced the work contained in this document. I also appreciate the help of the other members of my dissertation committee, namely Dr. Jabulani Nyathi and Dr. Partha Pande. Their assistance with my education and this dissertation has greatly contributed to my experiences at WSU.

I would like to thank my family for providing support and encouragement throughout the years. My mother, father and sister have provided assistance in all aspects of my life, and this has helped me greatly throughout my doctoral studies. In addition, I would like to thank my fiancée Rebecca Van Wyck for her love and patience. I look forward to our lives together, and I am excited to think of starting a family with her in the future.

Finally, I would like to acknowledge the camaraderie and dedication of my fellow students in the HiPerCopS research group. These students include Fredrick Anderson, Katie Blomster, Seon Kwang Jeon, Jonathan Larson, Mitchell Myjak, and Andy Widjaja.

HARDENED BY DESIGN APPROACHES FOR MITIGATING TRANSIENT
FAULTS IN MEMORY-BASED SYSTEMS

Abstract

by Daniel Ryan Blum, Ph.D.
Washington State University
May 2007

Chair: José G. Delgado-Frias

In radioactive environments, particle strikes can induce transient errors in integrated circuits (ICs). Strikes directly disrupting memory are known as Single-Event Upsets (SEUs), while strikes initially disrupting logic are called Single-Event Transients (SETs). Chips manufactured in aggressive technologies may also experience Multiple-Bit Upsets (MBUs). This research focuses on novel hardened by design circuit-level approaches to protecting integrated circuits against SEUs, SETs and MBUs. A number of system-level designs have been developed utilizing these approaches to demonstrate their capabilities.

Many of the design-hardened memory circuits considered in this study share a common theme, which is the ability to bypass transient faults. This is critical for performance, as it allows a system to proceed with subsequent operations while recovering from a disruption. Among the considered approaches, the novel Triple Path DICE (TPDICE) structure is the most balanced. This structure requires only two of its three inputs to be resolved during a write operation to ensure recovery. The recovery process requires approximately 50-100ps in 0.18 μ m CMOS.

A number of digital system applications for these approaches have been designed as part of this research. These applications include an SET-tolerant reconfigurable digital signal processing (DSP) architecture, SET-tolerant pipeline memory circuits, and an MBU-tolerant memory design. The radiation-tolerant DSP architecture is a memory-based system relying on TPDICE circuits to provide fault tolerance. Redundancy and additional circuit-level techniques are adopted to improve system-level reliability.

SET-tolerant pipeline memory structures are introduced to provide a means for protecting computational datapaths against faults. A number of single-ended and differential structures are presented and evaluated with respect to performance, energy consumption, and timing. Latches, master-slave flip-flops, and pulse-triggered flip-flops each offer a distinct balance of these comparison attributes.

Finally, a novel MBU-tolerant design is depicted, which utilizes layout-based interleaving and multiple-node disruption tolerant memory latches. This approach protects against non-grazing as well as grazing incidence particle strikes. Advantages with respect to size, complexity, and MBU tolerance are realized when compared to existing solutions.

Contents

Acknowledgement	iii
Abstract	iv
List of Tables	x
List of Figures	xi
1 Introduction	1
1.1 Single-Event Upsets and Single-Event Transients	5
1.2 SEU and SET Probability Calculations	8
1.3 SEU and SET Rate Trends with Respect to Feature Size	13
1.4 Overview of Fault-Tolerant Schemes	16
1.4.1 Error Correcting Codes and Hamming Code	17
1.4.2 Cross-Parity Approach	18
1.4.3 Triple Modular Redundancy	19
1.4.4 Circuit-Level Design-Hardened Approaches and the DICE Latch	21
1.5 Outline	22
2 SET-Tolerant Approaches	24
2.1 Triple Modular Redundancy	24
2.2 Basic Interlocked Design-Hardened Latches	26
2.2.1 SET-Tolerant Design-Hardened Latches	27
2.2.2 Read-Induced Upsets in Design-Hardened Latches	29
2.3 Modified Dice Approaches	30

2.4	Delay-Based Approaches.....	32
2.5	Fully-Differential DICE.....	34
2.6	Triple Path DICE Design.....	37
2.6.1	TPDICE Architecture.....	38
2.6.2	TPDICE Simulation Results.....	42
2.7	Analysis of Results	46
2.8	Summary	50
3	SEU and SET-Tolerant Reconfigurable DSP System	52
3.1	Reconfigurable Architecture.....	52
3.2	SEU and SET-Tolerant Design Approach.....	52
3.2.1	Integration of Design-Hardened Memory into the Proc. Architecture.....	56
3.2.2	Top-Level CMOS Implementation of Cell-Switch-Cell Design.....	58
3.2.3	Element Architecture.....	59
3.2.4	Switch Architecture.....	62
3.3	Analysis and Simulation	64
3.3.1	Element Simulation	64
3.3.2	System-Level Data Transfer Simulation.....	65
3.3.3	System-Level SET Response	66
3.3	Summary	68
4	Fault-Tolerant Pipeline Latches and Flip-Flops	69
4.1	Fault-Tolerant Pipelined Systems	70
4.1.1	Fault-Tolerant Structures.....	70
4.1.2	Pipeline Latches vs. Edge-Triggered Flip-Flops.....	73
4.2	Single-Ended Fault-Tolerant Pipeline Memories	77
4.2.1	TPDICE Latch.....	77

4.2.2	TPDICE Master-Slave Flip-Flop.....	78
4.2.3	TPDICE Pulse-Triggered Static Flip-Flop.....	78
4.2.4	TPDICE Dual-Pulsed Semidynamic Flip-Flop.....	84
4.3	Differential Fault-Tolerant Pipeline Memories	86
4.3.1	Barry-Dooley Differential Latch.....	86
4.3.2	Barry-Dooley Differential Master-Slave Flip-Flop.....	87
4.3.3	TPDICE Differential Latch	89
4.3.4	TPDICE Differential Master-Slave Flip-Flop.....	90
4.3.5	TPDICE Differential Pulse-Triggered Static Flip-Flop	92
4.3.6	TPDICE Differential Dual-Pulsed Semidynamic Flip-Flop	93
4.3.7	TPDICE Differential Sense Amplifier Flip-Flop.....	95
4.4	Comparison of Fault-Tolerant Pipeline Memories	97
4.5	Summary	106
5	MBU-Tolerant Approaches	107
5.1	Background.....	107
5.2	Ion Trajectories and Layout Interleaving.....	108
5.3	Multiple-Node Disruption Tolerant Latches.....	110
5.4	Layout Interleaving of MNDT Latches	113
5.5	Comparison to Existing Strategies.....	115
5.6	Summary	118
6	Conclusion	119
6.1	Contributions	122
6.2	Future Work.....	125
	References	127
A	SEU and SET-Tolerant Reconfigurable DSP Architecture Layout	133

B Publications	135
B.1 Journal	135
B.2 Conference	135
B.3 M.S. Thesis.....	136

List of Tables

2.1 DICE transient pulse-induced states and recovery states	36
2.2 Minimum clock periods of design-hardened memory circuits	47
2.3 Energy consumption figures of design-hardened memory circuits	47
2.4 Circuit complexity quantization of design-hardened memory circuits.....	47
4.1a Single-ended pipeline memory acronym definitions	100
4.1b Differential pipeline memory acronym definitions.....	100
4.2a Single-ended pipeline memory performance figures	101
4.2b Differential pipeline memory performance figures	101
4.3a Single-ended pipeline memory energy consumption figures.....	102
4.3b Differential pipeline memory energy consumption figures.....	103
5.1 Comparison of MBU-tolerant approaches.....	117

List of Figures

1.1	SER trends of single-bit SRAM and logic circuits with respect to feature size ..	15
1.2	Hamming parity bits inserted into an 8-bit data word	18
1.3	Overview of the cross-parity scheme.....	19
1.4	Block diagram of a TMR system.....	20
1.5	Basic DICE memory cell	22
2.1	Block diagram of a TMR scheme with delays.....	25
2.2	Block diagram of a bypass-capable TMR scheme with redundant voting	26
2.3	DICE design-hardened memory latch, modified to tolerate SETs	27
2.4	Enhanced DICE memory cell 1	30
2.5	Enhanced DICE memory cell 2	31
2.6	Delay-based tri-state buffer scheme for incorporating SET resistance	32
2.7	Delay-based DICE scheme for incorporating SET resistance	33
2.8	Timing diagram for bypassing transient disturbances	34
2.9	Fully-differential DICE circuit utilizing a shared data bus.....	35
2.10	TPDICE circuit utilizing a shared data bus.....	39
2.11	TPDICE schematic with separate read and write ports	39
2.12	TPDICE SET recovery process	41
2.13	Simulation of TPDICE internal nodes during SET recovery (H->L).....	43
2.14	Simulation of TPDICE internal nodes during SET recovery (L->H).....	43

2.15	Simulation of TPDICE internal nodes during SEU recovery (H->L)	45
2.16	Simulation of TPDICE internal nodes during SEU recovery (L->H)	45
2.17	Simulation setup used to compare the SET-tolerant latches.....	46
2.18	Plot of clock period results vs. SET width.....	48
2.19	Plot of energy consumption results.....	49
3.1	High-level view of the reconfigurable DSP architecture.....	53
3.2	Array of elements in the memory mode configuration.....	54
3.3	Array of elements in the mathematics mode configuration.....	54
3.4	8-bit multiply-accumulate function implemented with four cells.....	55
3.5	Cell and switch architecture used in the fault-tolerant DSP processor.....	56
3.6	Block diagram of cell-switch architecture with I/O interface labels	58
3.7	Block diagram of element architecture with I/O specifications.....	59
3.8	Precharge circuitry used to pull element data lines to Vdd/2	60
3.9	2:4 decoder circuit used in the design of the element.....	61
3.10	Selector used to enable column data in the element design.....	61
3.11	2:4 decoder used in the selector with clock and read/write inputs	62
3.12	Schematic of a switch node used in the cell-switch architecture.....	63
3.13	Latch used to store switch configuration data.....	63
3.14	Element simulation depicting clock, row address, data and output signals.....	65
3.15	Illustration of destination memory cell voltage levels during data transfer.....	66
3.16	Simulation of an SET affecting a switch during an H -> L transition	67
3.17	Simulation of an SET affecting a switch during an L -> H transition	67
4.1	System pipelined with a TPDICE latch	72

4.2 Barry/Dooley differential SEU and SET-tolerant memory cell.....	72
4.3 Basic flip-flop timing diagram.....	74
4.4 Pulse-triggered flip-flop timing diagram.....	76
4.5 Fault-tolerant single-ended TPDICE latch (SETL).....	77
4.6 Single-ended TPDICE master-slave flip-flop (SETMSFF).....	79
4.7 Single-ended TPDICE alternative master-slave flip-flop 1 (SETAMSFF1).....	80
4.8 Single-ended TPDICE alternative master-slave flip-flop 2 (SETAMSFF2).....	82
4.9 TPDICE-based flip-flop architecture block diagram.....	83
4.10 Single-ended TPDICE explicit pulse-triggered flip (SETPFF).....	83
4.11 Single-ended TPDICE dual-pulsed flip-flop (SETDPFF).....	85
4.12 Differential Barry-Dooley pipeline latch (DBDL).....	87
4.13 Differential Barry-Dooley master-slave flip-flop (DBDMSFF).....	88
4.14 Differential TPDICE latch (DTL).....	90
4.15 Differential TPDICE master-slave flip-flop (DTMSFF).....	91
4.16 Differential TPDICE alternative master-slave flip-flop 1 (DTAMSFF1).....	92
4.17 Differential TPDICE alternative master-slave flip-flop 2 (DTAMSFF2).....	93
4.18 Differential TPDICE pulse-triggered flip-flop (DTPFF).....	94
4.19 Differential TPDICE dual-pulsed flip-flop (DTDPPF).....	95
4.20 Differential TPDICE sense amplifier flip-flop (DTSAFF).....	96
4.21a Simulation setup for single-ended pipeline memories.....	99
4.21b Simulation setup for differential pipeline memories.....	99
4.22a Graph of delay vs. energy values for single-ended structures.....	104
4.22b Graph of delay vs. energy values for differential structures.....	105
5.1 MBU upset patterns resulting from grazing and non-grazing particle strikes ...	109

5.2	Six node latch possessing the ability to tolerate two simultaneous upsets	111
5.3	Eight node latch with dual-node and limited triple-node disruption tolerance..	112
5.4	Diagram of interleaved layout with selected grazing particle strikes	114
5.5	Interleaved layout of nine MNDT memory latches	115
5.6	Simulation of recovery from three disruptions affecting the MNDT approach.	116
5.7	Simulation of recovery from two SETs affecting the MNDT approach.....	116
A.1	VLSI Layout of the Radiation-Tolerant Reconfigurable DSP Circuitry	134

Chapter 1

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

Radioactive particles may pass through an IC, leaving behind a trail of charge. Faults induced by such occurrences are referred to as Single-Event Effects (SEEs). If 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]. In some situations, one particle strike may disrupt multiple nodes or bits in a circuit. Single particle strikes that upset multiple memory bits are known as Multiple-Bit Upsets (MBUs)

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 report.

One focus of this project is 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 100-200ps 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-cm²/mg often produce 100-200ps pulses, cosmic ray strikes with LET of 100 MeV-cm²/mg can induce SET pulses up to 2ns 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, 2ns SETs impose a ceiling of 500Mhz 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 report, including Triple Modular Redundancy (TMR), a basic SET-tolerant version of the Dual-Interlocked Storage Cell (DICE), buffered SET-tolerant DICE approaches, a delay-based DICE design, 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 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 200ps and 500ps SETs in the system. Energy consumption and circuit complexity figures are also considered.

To illustrate the system level benefits of bypass-capable structures, circuitry from a reconfigurable DSP processor has been implemented with the TPDICE design. The reconfigurable architecture consists of an array of four-bit LUTs that can be programmed and connected to perform various DSP algorithms. The LUTs are known as cells, and they are connected through reconfigurable switches. In this particular implementation, two fault-tolerant cells have been connected through one switch. The switch allows input data to be written to either cell or output data to be read from either cell. In addition, data can be transferred from one cell to the other. The transient bypass capability allows the system clock to be set independent of the maximum length SEU/SET.

Many high performance systems achieve timing benefits when flip-flops are utilized instead of basic latches. Pipelined systems can benefit from the clock skew and race condition tolerance provided by flip-flops, as their transparency window is much smaller than that of latches. However, the incorporation of fault-tolerance into pipeline flip-flops has not been explored in great detail. This report presents a number of pipeline memories that bypass transient faults, allowing them to achieve high performance.

Standard SEU and SET-tolerant designs may not function in situations where single particle strikes can upset multiple nodes or multiple bits in an IC. In these situations, MBU-tolerant structures must be adopted. The most basic way to mitigate multiple node disruptions is to increase the layout spacing between mutually vulnerable nodes. This spacing can be increased without increasing overall system area by utilizing layout interleaving techniques. These techniques can effectively mitigate the probability of

multiple node upset due to non-grazing particle strikes. However, grazing particle strikes can disrupt nodes separated by substantial distances along a line tracing the ion path. Fortunately, it is possible to protect against grazing strikes by combining layout interleaving with multiple node disruption tolerant memory.

1.1 Single-Event Upsets and Single-Event Transients

SEUs 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 [16]. 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 pulse 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.

Not all particle strikes result in upset. Protons, neutrons, heavy-ions, and alpha particles are all capable of causing SEU, and they each have unique energy transfer characteristics. Heavy-ions (cosmic rays) are prevalent in space, and they carry the highest energy of the group. On the other end of the spectrum, alpha particles possess the lowest energy, and are more likely to be found on earth [17].

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 and the circuitry surrounding the struck node [16]. SEU can occur in SRAM when the voltage of the struck node flips and is fed back through other inverter in the cell. The feedback process must occur before the cell can recover, or else upset will not occur [18]. Also, the load circuitry of the struck node directly impacts the recovery and feedback times of the cell. This significantly affects the probability of upset [16].

High concentrations of charged particles are often present in radioactive environments. Particles in space often possess very high energies, which makes SEU more likely [19]. Additionally, space shielding is ineffective protection against high-energy ions. Reducing the feature size of an IC results in smaller node capacitances, which are more susceptible to injected charge [17]. This increased sensitivity has resulted in SEU observation in terrestrial environments [20]. 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. Reduced feature sizes possess smaller node capacitances, which allow

larger voltage spikes. Smaller clock periods increase the probability that spurious data will propagate through combinational logic and be latched by a memory cell [23].

In the past, the fault-tolerance of combinational logic has been much less of a concern than that of memory. Masking effects have prevented almost all transient pulses from propagating through logic chains and into memory. However, the reduced feature sizes and supply voltages found in advancing technologies are diminishing 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 (which directly impact memory) can be controlled with relatively inexpensive schemes [23].

Since both logic and memory are constructed from transistors, the physics of SETs are similar to those of SEUs. Particle strikes deposit a trail of electron-hole pairs that can collect at reverse-biased junctions. If enough charge collects at a junction of a 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 [23].

As technology improves, smaller transistors and increased levels of superpipelining result in an overall reduction of these masking effects. Higher clock rates produce 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]. On the other hand, electrical and logical masking are not necessarily affected by technological improvements. 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].

1.2 SEU and SET Probability Calculations

SEU and SET probability is based on a number of factors, including IC fabrication technology, transistor drive strengths, node capacitances, IC operation environment, and circuit design techniques. As fabrication technologies improve, devices become smaller, making them more susceptible to upset. By the same token, circuits with reduced drive strength and smaller node capacitances are more likely to be affected by a given quantity of collected charge. The magnitude of radioactive strikes in the operating environment is determined by the particle flux and LET spectrum. Finally, the overall fault rate of an IC can be reduced by the use of circuit design techniques such as redundancy, interlocked feedback, and layout spacing.

Vulnerability of an IC to radioactive effects may be measured in terms of cross-section (σ) and sensitive volume. The cross-section is equal to the number of upsets

experienced by a chip divided by the particle fluence. Units for this measurement are often given in SEEs / Particle / cm². IC sensitive volume is defined as the portion of the chip that, if struck by a particle, may lead to upset. The expected error rate of an IC may be calculated by combining the cross-section and sensitive volume information with the LET spectrum of the operating environment [25].

A first-order IC cross-section analysis can be performed by considering the critical charge (Q_{crit}) or threshold LET (LET_{th}) of the design. Q_{crit} is defined as the minimum amount of collected charge ($Q_{collected}$) required at a sensitive node for upset to occur, with units of pC. In other words,

$$\sigma = 0 \text{ if } Q_{collected} < Q_{crit}, \quad (1.1)$$

$$\sigma = f(Q_{collected}) \text{ if } Q_{collected} \geq Q_{crit}, \quad (1.2)$$

where $f(Q_{collected})$ is a function with output values greater than zero.

As was mentioned above, SEE probability increases as feature sizes decrease. Generally speaking, a decrease in feature size is accompanied by a square increase in SEU and SET probability (assuming all other factors are held constant). More specifically, the general trend of IC critical charge has been modeled as

$$Q_{crit} = 0.023 t^2, \quad (1.3)$$

where t is the feature size in μm [26]. The critical charge of a device directly affects its probability of upset. However, it is important to keep in mind the ratio of collected charge leading to upset versus the total charge deposited, which is defined as the efficacy in [27]. $Q_{collected}$ is not equal to $Q_{deposited}$ in all circumstances. Efficacy differences exist

between various fabrication technologies and feature sizes. Because of this, efficacy factors into SEU probability with an importance similar to that of Q_{crit} .

Q_{crit} values can be converted to LET_{th} , which is useful when comparing individual particle LET values with the IC upset sensitivity threshold. LET_{th} is the minimum particle LET required to produce an upset, with units of MeV-cm²/mg. To perform this conversion, the angle of incidence of the particle strike and its path through the sensitive volume must be considered. LET_{th} values are calculated assuming the particle takes the longest possible path through the sensitive volume. Using a rectangular parallelepiped model for sensitive volume, the maximum possible chord length (s_{max}) of an ion track through this volume is equal to

$$s_{max}^2 = a^2 + b^2 + c^2, \quad (1.4)$$

with a , b , and c representing the volume width, length, and depth. Taking this into account, LET_{th} can be calculated as

$$LET_{th} = \frac{Q_{crit} w_{ehp}}{q\rho s_{max}}. \quad (1.5)$$

For this equation, w_{ehp} (energy required to generate an electron-hole pair) = 3.6 eV in Si, $q = 1.6022 * 10^{-19}$ C/e, and ρ (material density) = 2.33 g/cm³ for Si [25]. This equation assumes that $Q_{collected} = Q_{deposited}$, which is a simplification that removes efficacy effects from consideration.

If we consider a particle strike of a given $LET \geq LET_{th}$, it is possible to calculate the minimum distance it must pass through the sensitive volume to deposit at least Q_{crit} . This value, referred to as s_{min} , is equal to

$$s_{min} = \frac{Q_{crit} W_{ehp}}{q\rho LET} . \quad (1.6)$$

Additionally, it is possible to simplify the calculation process by assuming the ion strike cord length is equivalent to the depth of the sensitive region and adjusting the effective LET to compensate for the actual cord length. Given a particle strike of incident angle θ ,

$$LET_{eff} = \frac{LET_{inc}}{\cos(\theta)} , \quad (1.7)$$

where LET_{inc} is the actual LET of the incident particle strike, and LET_{eff} is the effective LET to be used in upset rate calculations.

The upset rate of a particular IC can be estimated through the use of the Figure of Merit (*FOM*) approach introduced in [28]. This model was designed to approximate IC sensitivity to cosmic rays at geosynchronous orbit. An updated version of this approach states that

$$FOM = \frac{\sigma_{HL}}{LET_{0.25}^2} \frac{(MeV / mg / cm^2)^2}{cm^2} , \quad (1.8)$$

with σ_{HL} defined as the limiting heavy ion cross section per bit at high LET and $LET_{0.25}^2$ equal to the LET at 25% of the limiting cross section [29]. $LET_{0.25}^2$ and σ_{HL} curves are obtained from experimental cross-section measurements of the IC and environment in question. If $LET_{0.25}^2$ values are not readily available, they may be estimated from LET_{th} as follows:

$$LET_{0.25}^2 = 2.77LET_{th}^{0.88}, \quad (1.9)$$

which is convenient due to the fact that LET_{th} is a basic first-order IC parameter [30, 31]. Additionally, if σ_{HL} measurement data is not available, it can be useful to simplify the FOM calculation through the use of sensitive volume dimensions and Q_{crit} estimations. This calculation is defined as

$$FOM = \frac{a * b * c^2}{Q_{crit}^2} \frac{pC^2}{\mu m^4}, \quad (1.10)$$

with a , b , and c representing the sensitive volume length, width, and depth [30]. Note that $1 \text{ pC}/\mu\text{m} = 98 \text{ MeV}/\text{mg}/\text{cm}^2$. From the figure of merit, the estimated soft error rate (SER) can be calculated as

$$SER = C * FOM, \quad (1.11)$$

where C is equal to a constant that is unique to the chip and its environment [29]. Average error rates of hardened devices are around 10^{-8} errors per bit-day. Unhardened devices possess error rates that are significantly higher than this [32].

It is important to note that the above analysis is based off of the critical charge model. While this offers an effective overview of the situation, it does not take the time dependency of charge collection into account. Generally speaking, the longer it takes for charge to collect at a reverse-biased junction, the less likely upset is to occur. This is due to the fact that the drive transistors controlling the affected node dissipate collected charge as it accumulates [16]. Models taking this factor into account require an input function relating the particle LET spectrum with charge collection duration values. This function would be dependent upon the characteristics of the IC as well as the target environment.

In practice, multiple factors affect the SER of fabricated digital systems. The above calculations provide a basic estimation. However, extensive simulations and/or actual data measurements are required to obtain accurate results. The following section presents SER trends obtained from these in-depth sources.

1.3 SEU and SET Rate Trends with Respect to Feature Size

A number of factors influence the SER of digital microelectronic circuits. These factors fall into two major categories: The attributes of a given IC and the environment it operates in. This section focuses on IC attributes, as this is the category that circuit designers can control. In total, four general IC attributes contribute to error rate trends, assuming constant sensitive volume. These factors include critical charge, efficacy, transistor drain area, and number of transistors per chip. As technology improves, critical charge decreases by a power of two and efficacy tends to decrease as well. The drain area of each transistor also decreases by a power of two, but this is canceled out by an equivalent increase in the number of transistors per chip. Therefore, the relationship

between Q_{crit} and efficacy is the major factor contributing to the SER of a given design [23].

Previous sections in this chapter introduced the concept that reductions in IC feature size are accompanied by a squared decrease in Q_{crit} on average. This attribute depends primarily upon node capacitances and supply voltage. Additionally, individual circuit structures possess differing Q_{crit} values. SRAM arrays, latches, and combinational logic have unique characteristics that affect their susceptibility to upset. Also discussed previously was the topic of charge collection efficiency, or efficacy. Efficacy, otherwise known as the difference between $Q_{deposited}$ and $Q_{collected}$, varies significantly between different technologies [23].

The trend of the critical charge vs. efficacy ratio is decreasing with feature size. This ratio is approaching a small constant as technology continues to improve. As this occurs, this ratio becomes less of a factor, and other considerations such as the size of the sensitive volume become more important [23].

The study featured in [23] presents error rates for SRAM, latch, and logic circuits across 0.6 μ m to 50nm feature sizes. The results of this study imply that the SER of logic is increasing at a much greater rate than that of memory structures. Single SRAM memory cell error rates decrease slightly as feature sizes decrease, as the critical charge vs. efficacy does not change significantly, whereas the area decreases. Single latch error rates decrease even more slowly, which is due to higher Q_{crit} at larger feature sizes. SERs for constant area SRAM arrays and pipeline latches increase at a marginal rate due to chip vulnerability factors that stay mostly constant with respect to each other. Finally, the SER for individual logic chains increases by over five orders of magnitude from 0.6 μ m to




Figure 1.1: Plot showing SER trends of single-bit SRAM and logic circuits with respect to feature size.

50nm, and constant area logic increases by nine orders of magnitude over the same technology interval. This can be attributed to the fact that improvements in technology reduce latching window masking effects. Generally speaking, the logic SER has been shown to increase at least linearly with clock frequency [23-24].

The graph in Figure 1.1 depicts the SER trends of single-bit SRAM and logic circuits as IC feature sizes decrease. Data for this graph was obtained by averaging results presented in [23, 33-35]. The SER values illustrated here result from neutron particle strikes, which are the most common form of SEE-inducing radiation on the surface of the earth. This radiation is the result of interactions between cosmic rays and our planet's

atmosphere [36]. SER trends for SRAM and logic in a space environment are not included due to the fact that they are not readily available. In any case, it is expected that this terrestrial data is analogous to the SER trends of space-based systems.

It may be observed that the SER for SRAM is staying relatively flat, whereas the logic SER is increasing at an average rate of approximately one order of magnitude per technology generation. By the 65nm generation, it is predicted that the SER of a single logic chain will exceed that of an SRAM cell. Note that ICs often possess significantly more SRAM cells than logic circuits, and so the chip-level SER of memory is likely to remain the major contributing factor for at least a handful of future generations. In any case, these trends reinforce the argument that SETs affecting logic will continue to grow in frequency at an alarming rate. ICs designed for operation in space must provide protection against SETs to maximize reliability. Additionally, consumer electronics targeted toward applications on the surface of the earth will likely require SET-tolerant circuitry in the near future.

1.4 Overview of Fault-Tolerant Schemes

A number of design-hardened approaches exist for the purpose of implementing fault-tolerant digital systems. These schemes rely on circuit design techniques such as temporal and spatial redundancy, feedback, and voting to detect and recover from transient disruptions. The approaches considered here are divided roughly into two categories, namely system-level and circuit-level designs.

System-level schemes rely on components outside of the memory structure to perform calculations and correct errors that may be present in the memory. Circuit-level schemes exist entirely inside of the memory structure, meaning that fault-tolerance is

incorporated directly into the design of the RAM latches [37]. Presented in this section are four existing schemes designed to implement fault-tolerance in hardware. Error correcting codes (ECCs) and cross-parity are provided as examples of system-level schemes. TMR may be considered a system-level or circuit-level approach based on the specific implementation. Finally, the DICE approach is examined as a representative circuit-level scheme.

1.4.1 Error Correcting Codes and Hamming Code

ECCs rely on the use of code or parity bits to protect data, and they are often used to provide fault-tolerance in digital systems. One popular ECC that can be used to efficiently protect individual data words is known as Hamming Code. This approach provides the ability to correct up to one error at a time in each data word. Encoded parity bits are inserted into each word. An example of an 8-bit data word with four inserted Hamming parity bits is shown in Figure 1.2. During a read operation, this encoded parity is decoded to give the position of an erroneous bit in the data word, if one exists. Accomplishing this requires an extensive XOR network to calculate the encoded parity, and a dedicated decoder to determine the position of a potential incorrect bit. This overhead would consume only a small percentage of the clock cycle in large SRAM or DRAM memory arrays. However, the overhead may not be acceptable in small memory arrays and pipelined datapaths that require very high levels of performance. Adding a multi-level XOR network and a decoder in series with the critical path of such a system would introduce a sizeable delay that would substantially increase the clock period.




Figure 1.2: Hamming parity bits P1-P4 inserted into an 8-bit data word D1-D8.

The suitability of ECCs as the only method of fault-tolerance in an IC may decrease as SET rates increase. ECCs can effectively prevent particle strikes from directly disrupting memory, but they are generally not useful at preventing SETs originating in logic. This is due to the fact that it is very difficult to maintain data coding through multiple logic layers. Additionally, encoding and decoding logic itself is susceptible to SETs. Generally speaking, redundant or delay-filtered data paths are the most robust methods for hardening logic against transient disruptions [34].

1.4.2 Cross-Parity Approach

As mentioned in the previous section, ECCs such as Hamming Code do not necessarily protect small memory arrays with great efficiency. One alternative to ECCs for small memories is known as the cross-parity approach. Error detection and correction in the cross-parity scheme are made possible by the storage of a parity bit for every row and column in a memory unit [38]. These parity bits are generated during writes to the memory. When a write occurs to the location corresponding to row i and column j of the memory, parity bits for i and j must be updated. If the same memory location is read at a later time, then the current parity of row i and column j is compared to the stored parity of i and j . Figure 1.3 illustrates the relationship between the parity bits and the rows and columns of the data memory. If the current parity is not consistent with the stored parity, then the memory location in question contains an incorrect bit (assuming there is no more

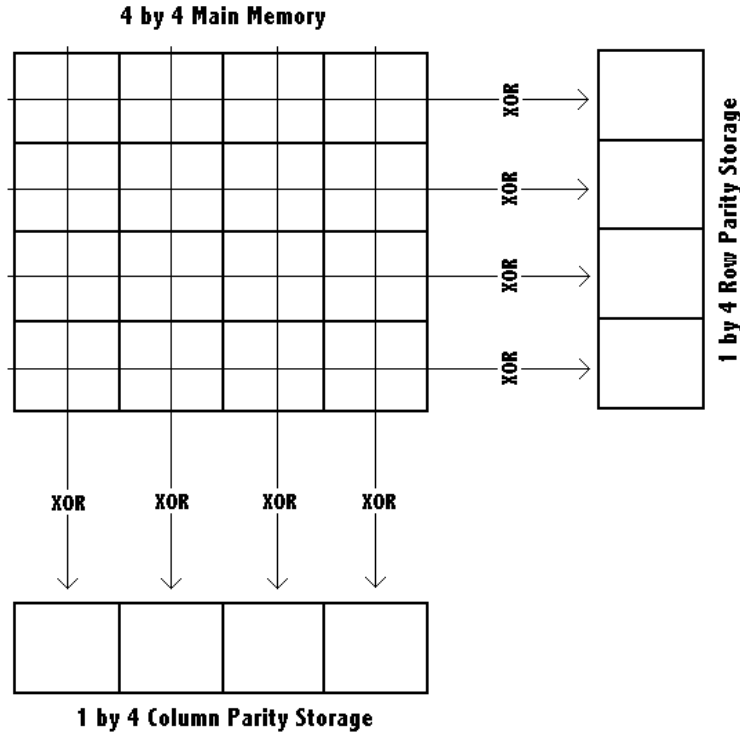


Figure 1.3: Overview of the cross-parity scheme.

than one error in the system). If this is not the case, then the bit is determined to be correct. A VLSI implementation of this structure is detailed in [39].

For small memory arrays, cross-parity configurations are more efficient than comparable Hamming Code designs. This is due to the smaller XOR network and lack of decoding required by cross-parity. Because of this, cross-parity is more optimal in terms of delay, size, and computational complexity.

1.4.3 Triple Modular Redundancy

In a TMR-based system, three copies of all memory cells are used along with voting to filter out one fault per memory triplet during data transfer. When a read operation is performed, the three memory cells each send their version of the data. The voting circuit then passes on the data that was sent by the majority of the memory cells. A graphical




Figure 1.4: Block diagram of a TMR system.

depiction of this is shown in Figure 1.4. Such a system continues to function correctly when multiple errors are present in distinct triplet locations of a memory array, which is not possible in a cross-parity system. However, this comes at the cost of 200% memory overhead, in addition to voting logic and optional control circuitry. In contrast, cross-parity only requires 50% memory overhead, in addition to XOR logic and control signal generation. Also, updating an incorrect bit in memory is more complicated in TMR, as the module containing the error must be identified and coordinated with the appropriate memory control signals. With these factors in mind, it can be reasoned that cross-parity would provide size and energy consumption advantages over TMR for small memory arrays. On the other hand, TMR would be more efficient in larger memory arrays and bit-slice pipelined datapaths. Additionally, TMR has the potential to provide greater performance, assuming its voting circuitry could be made faster than the XOR network required by cross-parity.

Permanent damage to memory that is caused by faults must be detected and bypassed to insure the robust operation of a system. The software in a reconfigurable architecture can be designed to detect this damage and reconfigure the system to avoid compromised circuitry. In addition, a number of hardware-based methods can be used to accomplish this, including TMR and the system-level approaches described above. System-level approaches require significant area and delay overhead outside of the memory itself. This overhead may not be necessary in all circumstances.

1.4.4 Circuit-Level Design-Hardened Approaches and the DICE Latch

In a number of situations, the only faults that will affect an IC are SEUs and other transient errors. Decreasing the feature size of a circuit increases the impact of an SEU by a power of two, which means that SEUs are becoming more of a problem as IC technology improves [1]. The system-level and block-level schemes listed above will protect against SEUs, but they are not optimal. First off, it would be preferable if the chosen solution reacted immediately to an SEU instead of only when read operations are performed. Secondly, the capability to bypass affected hardware is unnecessary if transient errors are the only concern. Both of these objectives can be accomplished by adopting a circuit-level approach, where SEUs are prevented by modifying the latches themselves, instead of adding additional components outside of the memory cells. This approach can significantly reduce area and delay of the system, as the error-correcting scheme is integrated directly into the memory it is protecting [40].

The Dual Interlocked storage Cell (DICE) shown in Figure 1.5 is a circuit-level design that has the capability to recover from transient faults at any of its feedback nodes [37]. The DICE memory cell is based off of four inverters which are connected in an

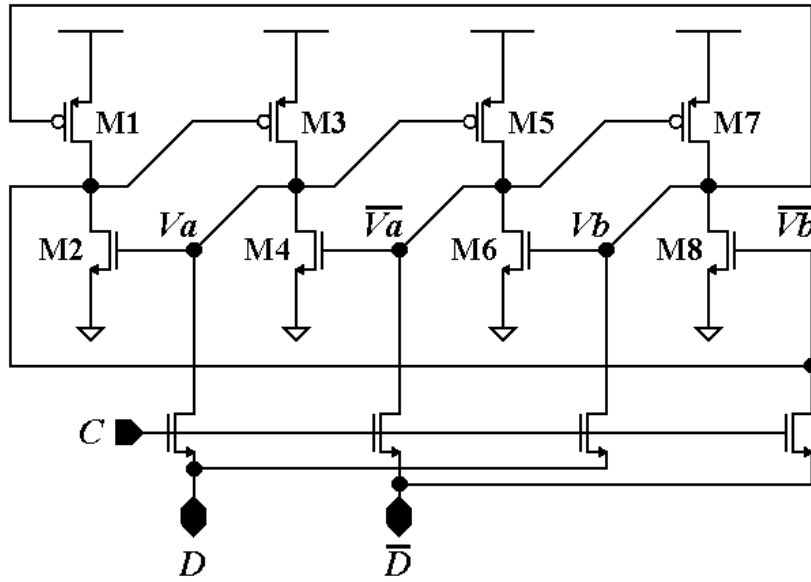


Figure 1.5: Basic DICE memory cell.

interlocked fashion. Arranging the latch interconnectivity this way assures that two separate inverters will restore an upset 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 report will present solutions to this and other weaknesses of the basic DICE latch.

1.5 Outline

The remainder of this document is organized as follows. Chapter 2 introduces and evaluates SEU and SET-tolerant circuit-level memory designs. The designs are compared with respect to minimum clock period, energy consumption, circuit complexity, and transient bypass capability. Chapter 3 presents an application of these fault-tolerant approaches, namely an SEU and SET-tolerant reconfigurable DSP processor. This architecture is a memory-based design that relies upon design-hardened memory cells to achieve fault-tolerance. Chapter 4 discusses the subject of SEU and SET-tolerant pipeline

memory circuits. A number of structures are introduced and evaluated based on performance, energy consumption, and timing considerations. Chapter 5 looks into an approach for tolerating Multiple-Bit Upsets (MBUs), which can be a problem with advanced technologies used in extremely radioactive environments. The presented solution combines a multiple-node disruption tolerant latch with layout interleaving to withstand grazing and non-grazing MBU-inducing strikes. Finally, Chapter 6 concludes the document by highlighting the contributions of this research and presenting closing remarks.

Chapter 2

SET-Tolerant Approaches

Hardened by design approaches have been used in a number of fault-tolerant digital system architectures. These approaches are an alternative to process hardening in cases where process-level techniques are not available or adequate for the desired level of fault-tolerance. Considering specific hardened by design approaches, circuit-level hardening can be used to achieve increased performance over ECCs, system-level schemes, and instruction-level approaches. Traditional circuit-level hardened by design approaches include TMR and memory cells relying on redundant transistors and interlocked feedback.

2.1 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




Figure 2.1: 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.

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 Figure 2.1 [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 combinational logic, the transient pulse can affect only one destination cell at any time during a write operation. Unfortunately,




Figure 2.2: Block diagram of bypass-capable TMR scheme with redundant voting.

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 Figure 2.2 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. 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.

2.2 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 [37]. It

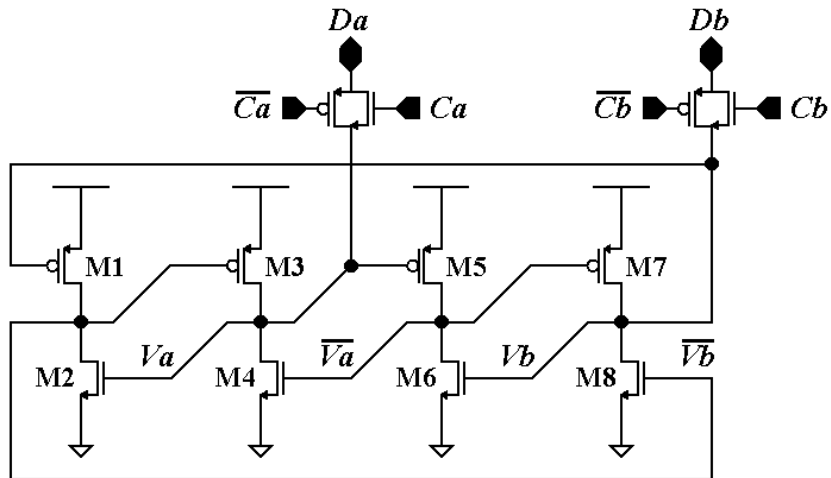


Figure 2.3: DICE design-hardened memory latch, modified to tolerate SETs.

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.

2.2.1 SET-Tolerant Design-Hardened Latches

Although the original DICE design is susceptible to SETs, the circuit can be modified to tolerate these disturbances. Figure 2.3 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 in terms of area and energy consumption.

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-cm²/mg alpha particles are a concern, SEU/SET pulses are generally limited to only 100-200ps. However, they can be as large as 300ps to 2ns when 100 MeV-cm²/mg cosmic rays are brought into consideration in heavily radioactive 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.

2.2.2 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 vs. the enable transistors. The bus capacitance isn't 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 tri-state 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 vs. node capacitance, and also precharging the data bus to an intermediate value such as $V_{dd}/2$ can help to avoid

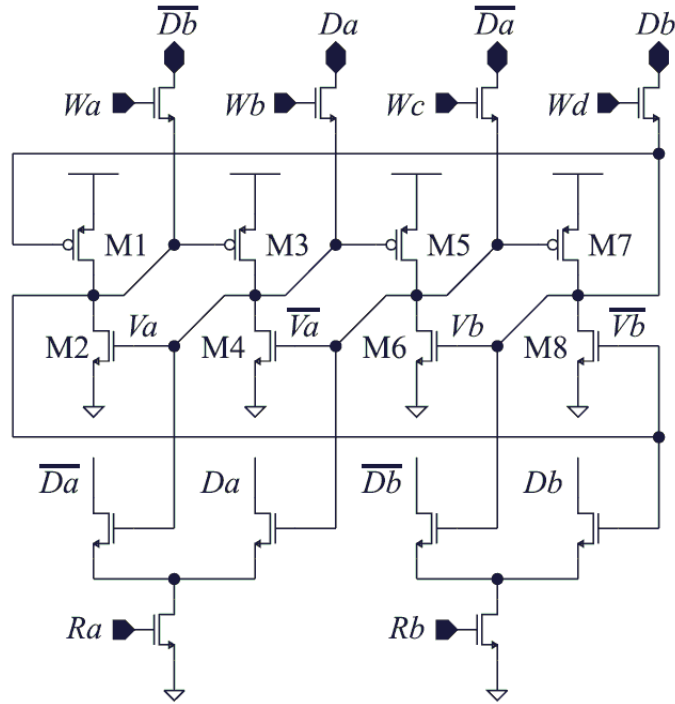


Figure 2.4: Enhanced DICE memory cell 1.

this issue. This is not practical in large memories, necessitating the adoption of output buffers for each cell.

2.3 Modified Dice Approaches

This section will describe additional 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. None of the approaches in this section have the ability to bypass transient pulses, which limits performance.

Four write transistors are used in Figure 2.4, each of which is enabled by an independent write line [5]. If an SET were to disrupt a write line, only one internal node

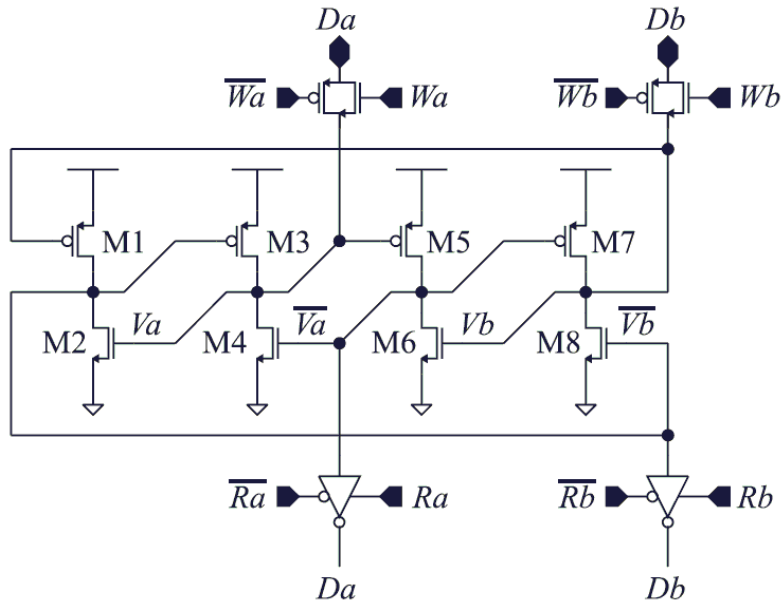


Figure 2.5: Enhanced DICE memory cell 2.

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.

Another important property of this cell 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 V_{dd} . 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 2.5 uses different approaches to protect the data and write enable lines and to provide buffering during reads [5]. Write operations drive two internal nodes

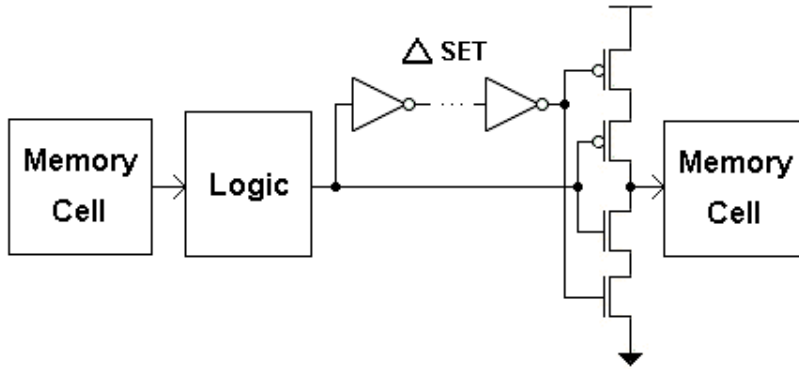


Figure 2.6: Delay-based tri-state buffer scheme for incorporating SET resistance.

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.

2.4 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. Obviously, incorporating this delay prohibits these approaches from bypassing transient disruptions. Figure 2.6 shows one implementation of this strategy [8]. A tri-state buffer (also known as a Muller C-Element) 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




Figure 2.7: Delay-based DICE scheme for incorporating SET resistance.

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 2.7 illustrates a temporally-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




Figure 2.8: Timing diagram for bypassing transient disturbances.

of the receiver cell are driven to opposite values for a greater amount of time, which degrades performance.

2.5 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. Figure 2.8 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

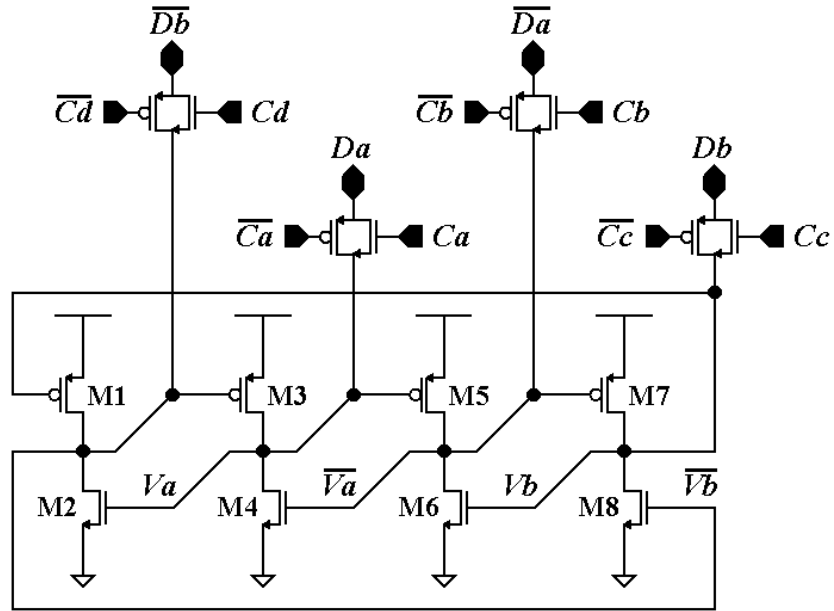


Figure 2.9: Fully-differential DICE circuit utilizing a shared data bus.

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 Figure 2.9. 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

Table 2.1: DICE transient pulse-induced states and recovery states.

Initial State (Va /Va Vb /Vb)	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

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 Figure 2.9. Table 2.1 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 reasonably balanced). The two 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.

2.6 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 Triple Path DICE (TPDICE) approach presented in this section addresses these issues.

2.6.1 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 Figures 2.10 and 2.11, 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 [41].

TPDICE latches possesses 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 Figure 2.10 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

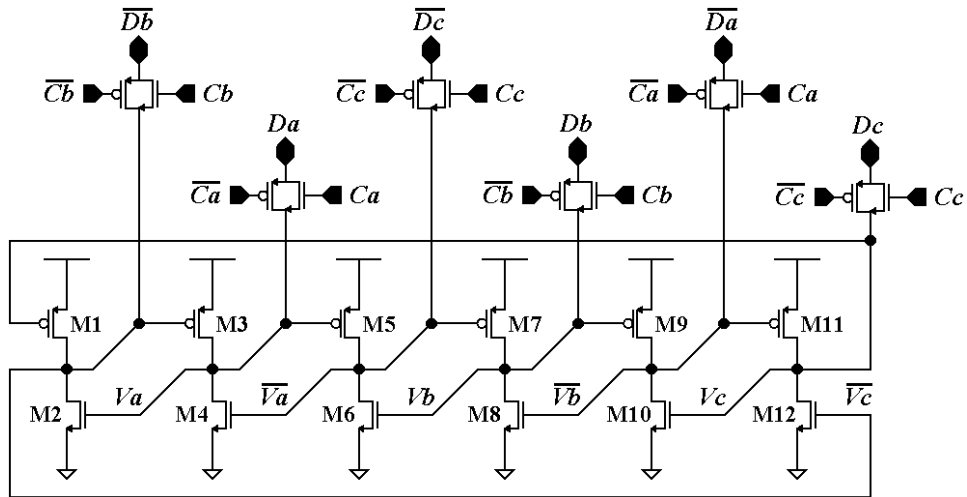


Figure 2.10: TPDICE circuit utilizing a shared data bus.

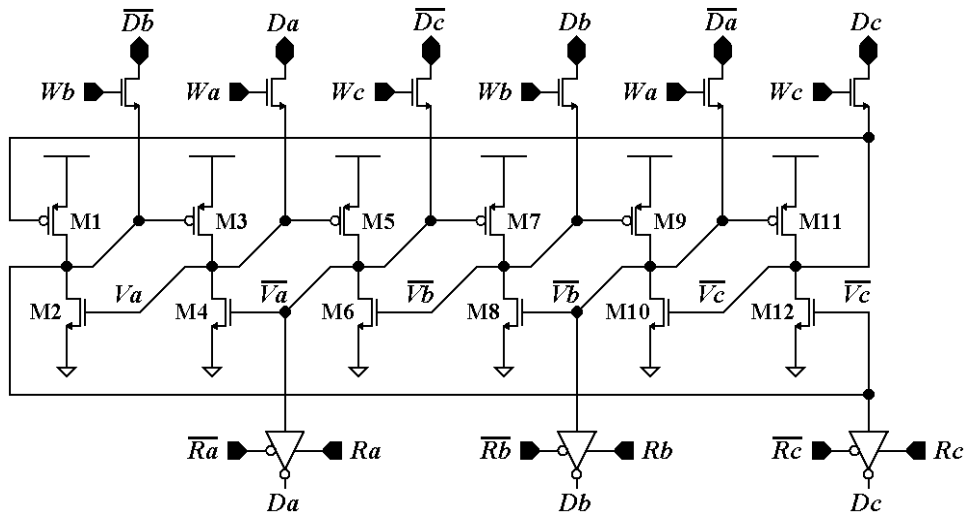


Figure 2.11: TPDICE schematic with separate read and write ports.

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.

Figure 2.11 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 insures that the state of the latch can not 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. Figure 2.12 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




Figure 2.12: TPDICE SET recovery process.

happens, two factors may combine to resist the write operation. First, the destination memory may initially be at the opposite state of the 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 Va , Vb , and Vc will be affected (similarly for $/Va$, $/Vb$, and $/Vc$). Assuming the outputs are Va , Vb , and Vc , 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 Figures 2.10 and 2.11.

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

2.6.2 TPDICE Simulation Results

Figure 2.13 depicts an SET affecting the Da data path with buffers driving data and enable signals, replacing the ideal voltage sources used for Figure 2.12. SETs affecting Db or Dc are not illustrated in this paper, as their effect is analogous to SETs on Da due




Figure 2.13: Simulation of TPDICE internal nodes during SET recovery process (H->L transition).




Figure 2.14: Simulation of TPDICE internal nodes during SET recovery process (L->H transition).

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 (V_b , V_c) 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 D_a during a 010101 to 101010 write transition is shown in Figure 2.14. This case is the inverse of the previous write simulation. Again, V_a and $\overline{V_b}$ are directly affected by the SET. V_a is held low and $\overline{V_b}$ 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, V_a and $\overline{V_b}$ are driven to their proper values, V_c and $\overline{V_c}$ are tri-stated, and $\overline{V_a}$ and V_b 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 $\overline{V_a}$ is directly affected during a 101010 to 010101 write operation. V_a is held high during the write, $\overline{V_c}$ experiences conflict, and all other nodes are driven to proper values. After the write operation is completed, V_a continues to be affected by the particle strike, and $\overline{V_c}$ 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




Figure 2.15: Simulation of TPDICE internal nodes during SEU recovery process (H->L transition).




Figure 2.16: Simulation of TPDICE internal nodes during SEU recovery process (L->H transition).

this is shown in Figure 2.15. 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. V_a is held low by the particle strike, and $\wedge V_a$ experiences conflict from this. The other four nodes are driven to desired values by the write buffers. At the end of the write cycle, V_a continues to be affected by the transient pulse, and $\wedge V_a$ loses its write buffer support. Fortunately, this cell state is

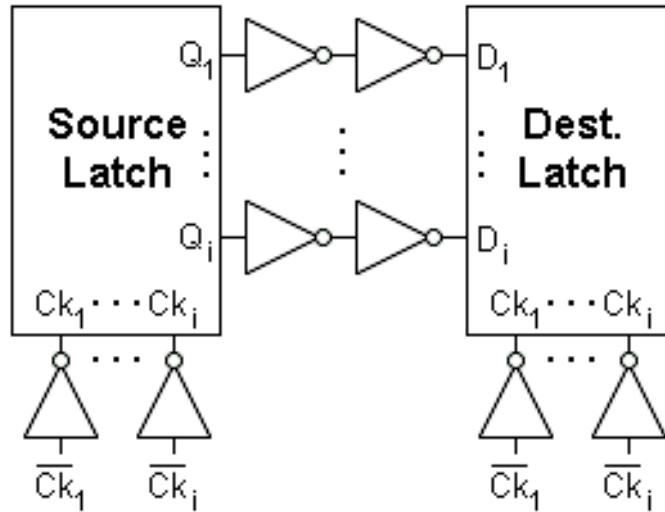


Figure 2.17: Simulation setup used to compare the SET-tolerant latches.

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. Figure 2.16 is a simulation of the internal node status during this scenario.

2.7 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 on the diagram shown in Figure 2.17. 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 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

Table 2.2: Minimum clock periods.

	No SET	200ps SET	500ps SET	Bypass
TMR	628ps	749ps	749ps	Data, Logic
SET-Tolerant DICE	515ps	1010ps	1310ps	none
Enhanced DICE 1	258ps	668ps	968ps	none
Enhanced DICE 2	587ps	988ps	1341ps	none
Delay DICE	863ps	1449ps	1753ps	none
Differential DICE	348ps	473ps	480ps	Data Only
TPDICE	434ps	545ps	552ps	Data, Logic

Table 2.3: Energy consumption figures.

	Read	Write	Buffers	Voting	Total
TMR	127.1fJ	180.5fJ	560.4fJ	264.1fJ	1132.1fJ
SET-Tolerant DICE	35.62fJ	84.69fJ	424.5fJ	n/a	544.81fJ
Enhanced DICE 1	7.801fJ	182.4fJ	404.6fJ	n/a	594.80fJ
Enhanced DICE 2	111.0fJ	170.2fJ	467.4fJ	n/a	748.60fJ
Delay DICE	102.4fJ	234.7fJ	506.9fJ	n/a	844.00fJ
Differential DICE	29.86fJ	44.75fJ	414.4fJ	n/a	489.01fJ
TPDICE	33.85fJ	84.00fJ	459.7fJ	n/a	577.55fJ

Table 2.4: Circuit complexity quantization.

	Memory	Vote/Delay	Logic	I/O
TMR	18	36	3x	3 data, 3 ck
SET-Tolerant DICE	12	n/a	2x	2 data, 2 ck
Enhanced DICE 1	18	n/a	4x	4 data, 4 w, 2 r
Enhanced DICE 2	20	n/a	2x	2 data, 2 w, 2 r
Delay DICE	20	4	1x	1 data, 2 w, 1 r
Differential DICE	16	n/a	4x	4 data, 4 ck
TPDICE	24	n/a	3x	3 data, 3 ck

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 2.2 and shown graphically in Figure 2.18. The left column of the table shows the minimum clock period attained for each approach with no transient pulses,

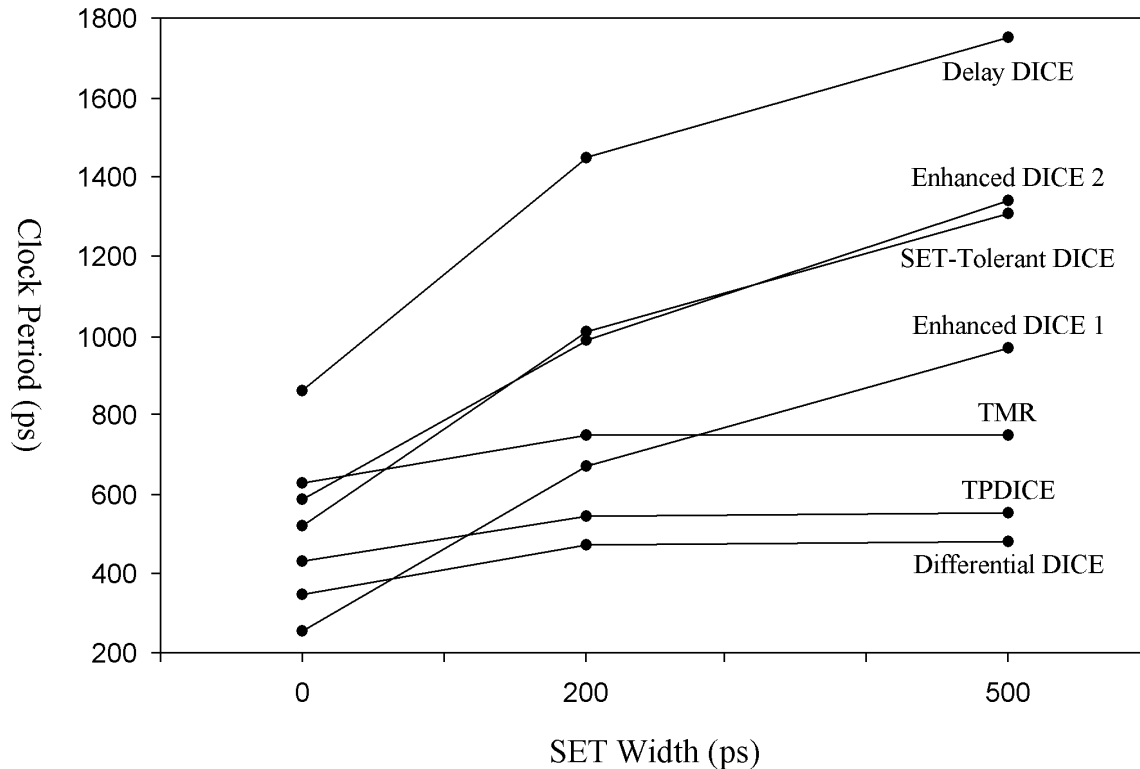


Figure 2.18: Plot of clock period vs. SET width. Note the quasi-linear dependence on SET width is only present in the SET-Tolerant DICE circuit, which is the only approach that can not bypass transient pulses.

while the other columns show the clock periods needed to tolerate 200ps and 500ps SETs. Table 2.3 presents energy consumption figures that are broken down into read memory, write memory, buffer, voting, and total contributions. Figure 2.19 is a graphic depiction of the energy consumption results. Finally, Table 2.4 depicts the circuit complexity of each approach, divided up into memory, voting/delay, logic and I/O categories. All simulations were performed in 0.18 μ m CMOS technology.

These results show that while no approach dominates all comparison categories, TPDICE is the most balanced. TMR achieves reasonable performance (749ps with a 500ps SET), but has substantial costs in terms of energy consumption (1132.1fJ) and

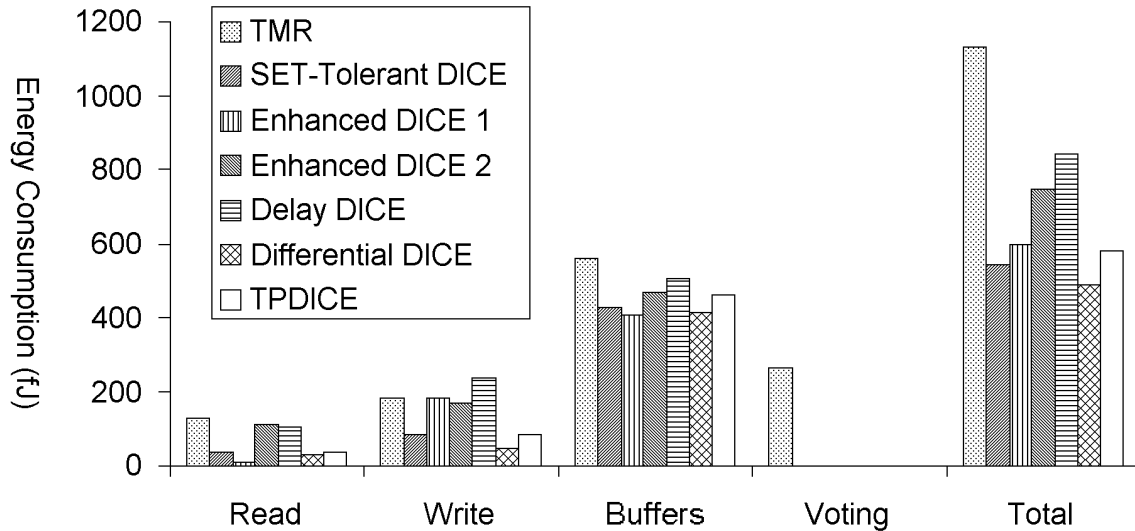


Figure 2.19: Plot of energy consumption results. The total consumption of TMR is nearly double that of each DICE approach. Note that the energy consumption of each approach is dependent on the logic in the data paths (if any is used).

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 (544.81fJ), but depends on a clock period that is proportional to the maximum SET width (515ps with no upset, 1010ps with 200ps SET, 1310ps with 500ps 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 (473ps clock period with a 200ps SET, 480ps with a 500ps SET) and low energy consumption (489.01fJ) 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 (545ps clock period with a 200ps SET, 552ps with a 500ps

SET) while using three data paths. It requires 577.55fJ 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 ΔSET , 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 ΔSET , making the system scalable to situations that must tolerate wide SETs.

2.8 Summary

This chapter discussed the concept of circuit-level hardened by design SEU and SET-tolerant approaches. Such approaches rely on redundancy, voting and/or interlocked feedback at the memory cell level. TMR designs utilize three copies of all memory and voting to filter out potential upsets. DICE-based structures rely on a four-node interlocked SEU-tolerant memory cell supported by SET-tolerant read/write enable

circuitry. The novel TPDICE approach presented above is based off a six-node memory cell that possesses the ability to bypass transient pulses. It is more efficient than bypass-capable TMR in terms of speed, energy consumption, and circuit complexity. This is due to the fact that area-redundant TMR requires triplicated external voting, whereas TPDICE does not.

Chapter 3

SEU and SET-Tolerant Reconfigurable DSP System

A reconfigurable DSP architecture has been designed to balance performance, power consumption, and versatility [42-43]. This design relies on a two-level memory-based configuration. The top level consists of a reconfigurable array of four-bit cells that can be programmed to perform memory and DSP-optimized arithmetic operations. At the lower level, a 4x4 matrix of LUT-based elements forms the processing core of each cell. An SEU and SET-tolerant implementation of this architecture is detailed below. TPDICE memory cells have been adopted to provide fault-tolerance and transient pulse bypass capability.

3.1 Reconfigurable Architecture

The reconfigurable DSP processor architecture is made up of medium-grain cells, as opposed to the fine-grain components that make up field programmable gate arrays (FPGAs). For this application, FPGAs require excessive area and power to achieve an unnecessary level of flexibility. Due to the regularity of most DSP algorithms, a medium-grain reconfigurable structure is sufficient. The structure featured in this section consists of an array of cells that perform 4-bit operations. Every cell is connected to its eight neighbors by sixteen 4-bit busses. Figure 3.1 is a high-level illustration of the reconfigurable DSP architecture.




Figure 3.1: High-level view of the reconfigurable DSP architecture.

The processing core of each cell is made up of a 4x4 array of elements. Each element is a 16x2-bit lookup table (LUT) that stores the truth table of a user-defined function. The array of elements can be arranged into a memory mode or mathematics mode configuration, which are displayed in Figures 3.2 and 3.3. The memory mode arrangement turns the cell into a 64x8 bit random access memory, providing storage capability for the processor. In mathematics mode, the structure of the array of elements is similar to that of a carry-save multiplier. This facilitates the efficient implementation of many arithmetic functions used in DSP, including addition and multiply-accumulate.

It is possible to implement functions of multiple word lengths with this reconfigurable architecture. Each cell manipulates 4-bit operands, and the cells can be cascaded to process longer data. Figure 3.4 shows four cells interconnected to implement an 8-bit multiply-accumulate function. Word lengths of 16, 32, 64 or even 128 bits can be achieved in this fashion. Many functions with long word lengths take multiple cycles to compute. To increase the throughput, pipeline latches are present in every cell, allowing a new operation to be initiated during every clock cycle [42-43].

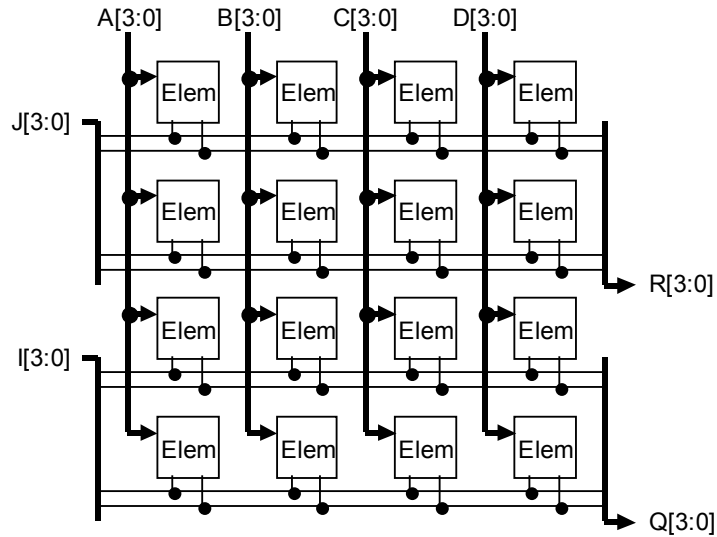


Figure 3.2: Array of elements in the memory mode configuration.

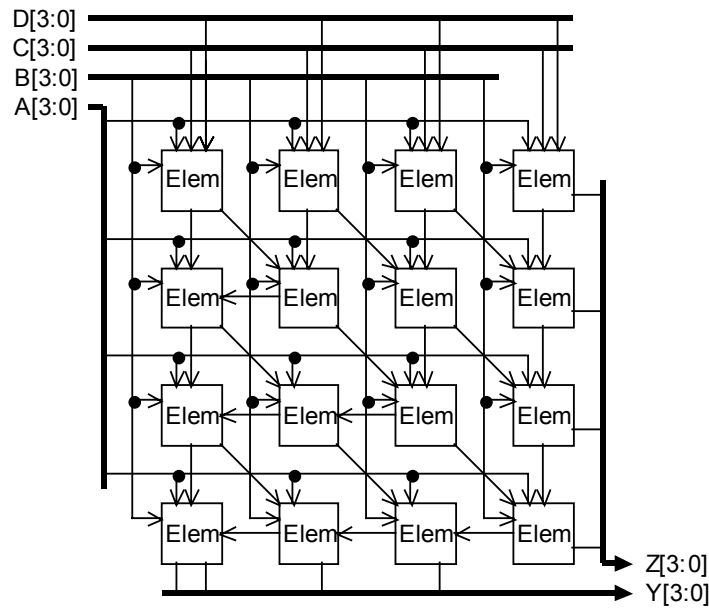


Figure 3.3: Array of elements in the mathematics mode configuration.

The reconfigurable DSP processor is constructed from LUTs and programmable switches. The LUTs are small SRAM memory arrays, relying on decoders and shared data busses. Because of the shared busses, the LUT memory cells require read enable capability. LUT performance depends on optimizing the decoder delays, buffer driving




Figure 3.4: 8-bit multiply-accumulate function implemented with four cells.

capabilities, and node/bus capacitances. The LUTs are programmed with truth tables of the algorithms implemented by the processing core. Faults that change values in the LUT could corrupt the functionality of the system until they are overwritten.

Switches provide reconfigurability in the DSP processor. They are constructed from transmission gates and memory cells that store switch configuration. The configuration memory turns on the necessary transmission gates to create the desired data paths. Switch performance is maximized by balancing impedance through the transmission gates with buffer strength. The switch architecture is vulnerable to SEUs affecting the switch configuration memory and SETs affecting the transmission gates. Adopting any of the SET-tolerant schemes presented in this report would protect against both possibilities.

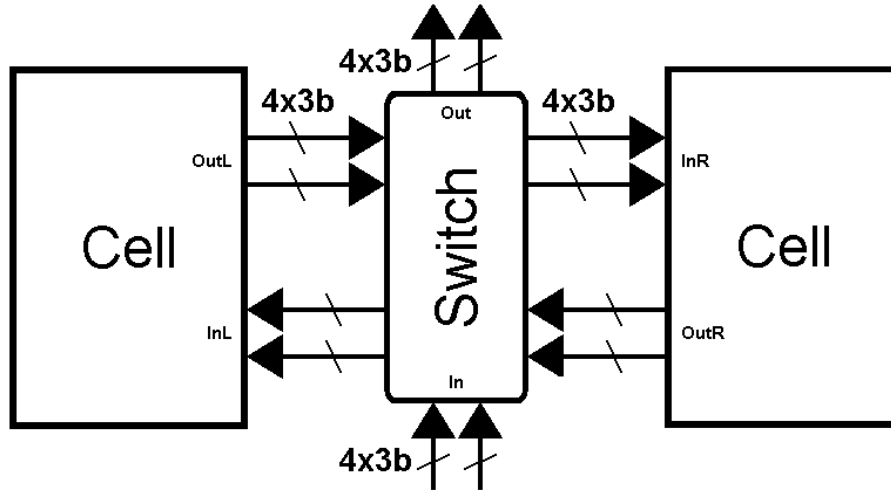


Figure 3.5: Cell and switch architecture used in the fault-tolerant reconfigurable DSP processor.

3.2 SEU and SET-Tolerant Design Approach

The fault-tolerant implementation of reconfigurable hardware considered in this chapter consists of two LUT-based cells that are connected by a reconfigurable switch, as shown in Figure 3.5. This setup is sufficient to demonstrate the fault-tolerance of all aspects of the reconfigurable architecture, including its resistance to faults affecting LUT memory, switch memory, switch logic, and write/read enable logic. The switch allows outside input data to be routed to either cell. Data can be transferred between cells in either direction. In addition, either cell can transmit data to outside outputs. Each data path consists of two triple redundant 4-bit paths in each direction.

3.2.1 Integration of Design-Hardened Memory into the Processor

Architecture

The Triple Path DICE (TPDICE) latch was chosen to provide fault-tolerance in this implementation of the DSP processor. This memory cell is shown in Figure 2.10. It

exhibits advantages in systems that utilize small memory arrays and combinational logic. Small memory arrays possess low bus capacitance, which limits the voltage ripple magnitude TPDICE latches experience during read operations. This reduces the probability of read-induced SEU. In addition, TPDICE possesses the capability to bypass transient faults. This bypass capability allows the system to withstand transient upsets without waiting for them to dissipate, thereby avoiding clock overhead proportional to the maximum transient pulse width. Studies have measured transient pulse widths of up to 2ns. Avoiding overhead of this magnitude is very desirable from a performance perspective [12-15].

Read-induced upsets are avoided in this implementation of the fault-tolerant reconfigurable DSP processor through capacitance management and precharging data busses to $V_{dd}/2$. Since only four latches are connected to each column-aligned data bus, the ratio of bus capacitance versus latch node capacitance can be maintained at a reasonable level.

Integration of the fault-tolerant hardware into the reconfigurable architecture creates a number of system-level concerns. It requires the duplication of logic, switches, data paths and read/write enable lines. This increases the complexity of the reconfigurable architecture and requires additional design considerations. Interconnect routing becomes more involved, and transistor sizing must be done in such a way as to avoid the magnification of SET pulses.

This configuration of reconfigurable DSP circuitry enables the testing of SRAM-to-SRAM memory transfer protection capability. The TPDICE latch and triplicated supporting circuitry provide the necessary fault-tolerance. Specifically, this setup




Figure 3.6: Block diagram of cell-switch architecture with I/O interface labels.

demonstrates that this approach possesses the ability to mitigate SETs that affect the read/write enable circuitry and switch logic.

3.2.2 Top-Level CMOS Implementation of Cell-Switch-Cell Design

An implementation of this design has been developed in 0.18 μ m CMOS. This implementation consists of a 15k transistor, five metal layer VLSI layout. Figure 3.6 shows the element and switch arrangement of this design. The I/O signals interfacing




Figure 3.7: Block diagram of element architecture with I/O specifications.

these components are labeled in this diagram. All signals are triplicated to ensure fault-tolerance and transient bypass capability.

3.2.3 Element Architecture

Figure 3.7 depicts the element architecture used in this implementation of the DSP processor. TPDICE LUT memory is arranged in a 4x4x2-bit array. Precharge circuitry drives the data lines to $V_{dd}/2$ during the decoding phase to help mitigate the possibility of read-induced SEU. 2:4 decoders receive the upper order address bits and use this data to enable a row of the LUT. This data is transmitted to the selectors, which enable a column of data based on the lower order address bits.




Figure 3.8: Precharge circuitry used to pull element data lines to $V_{dd}/2$ in between data operations.

The memory-based element relies on the precharge circuitry shown in Figure 3.8 to bring data lines to $V_{dd}/2$ in between all read and write operations. This is done to reduce the amount of charge sharing between the data busses and the memory cells during read operations, thereby protecting against read-induced SEU. The circuit uses NMOS transistors to drive the data lines to $V_{dd}/2$ during the low phase of the clock. NMOS transistors are also used to speed up the precharge process by short circuiting true and complement data lines, which redistributes charge to equal out the line voltages.

Row decoding operations in the element structure are performed by the circuit depicted in Figure 3.9. This circuit is a 2:4 decoder with enable capability. It takes the upper two element address bits, and uses this information to enable one row in the element memory. The decoder lines are disabled during the precharge cycle and enabled during data operations. When the decoder is disabled, the decoder logic is separated from the memory array and the decoder outputs are deasserted. In this state, all memory cells are disconnected from the data busses, allowing precharge operations to take place without conflict.

The selector circuit in Figure 3.10 enables columns in element memory for read and write operations. It operates on the lower two element address bits, as well as read and write input signals. A 2:4 decoder takes the input data and generates enable signals for

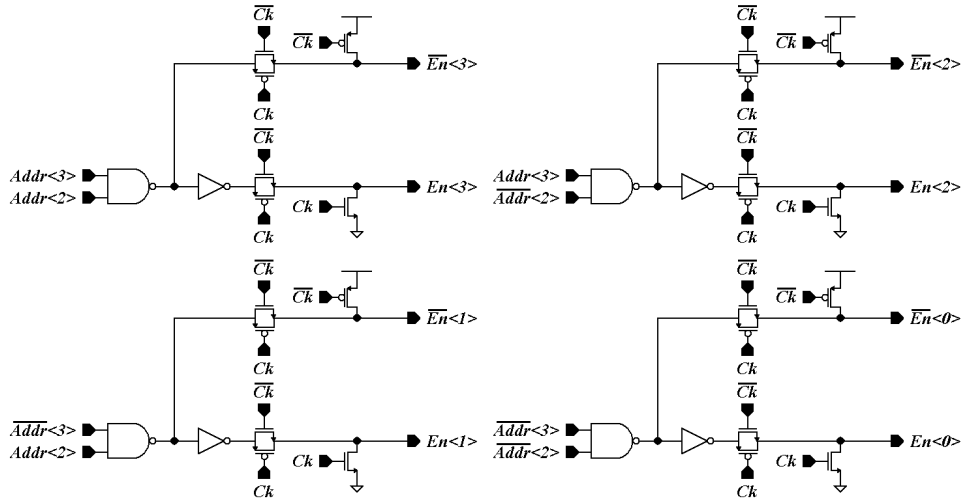


Figure 3.9: 2:4 decoder circuit used in the design of the element.

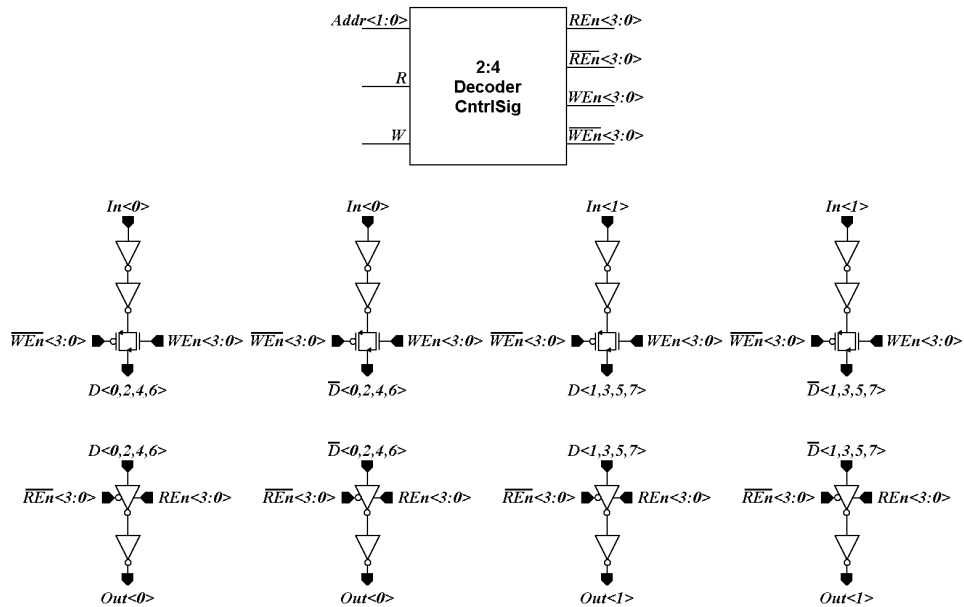


Figure 3.10: Selector used to enable column data in the element design.

the column selection operation. Separate data paths exist for read and write operations.

Columns are selected through the use of transmission gates and tri-state inverters.

The decoder circuit displayed in Figure 3.11 is used by the selector to enable the proper column of data during element read and write operations. This design is similar to that of the element row decoder, as it is a 2:4 structure with enable capability. It differs

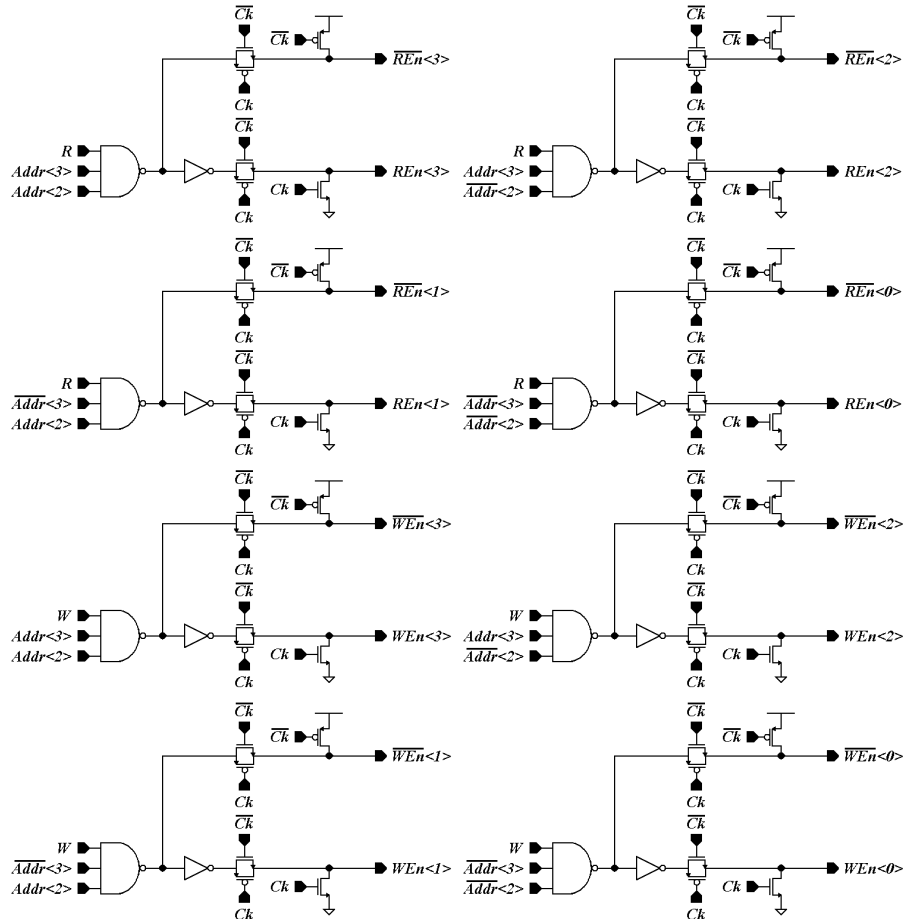


Figure 3.11: 2:4 decoder with clock and read/write inputs. This design is used in the column selector circuit.

from the row decoder in that it accepts the lower two address bits and read/write signals as inputs. The lower order address bits specify the desired column, while the read/write signals identify the direction of data transfer.

3.2.4 Switch Architecture

The circuitry used to implement a switch node in the cell-switch architecture is shown in Figure 3.12. Transmission gates enable the desired data path configuration, and a TPDICE latch is used to store the configuration data. The TPDICE configuration latch and triplicated switch data paths provide SEU and SET bypass capability.




Figure 3.12: Schematic of a switch node used in the cell-switch architecture.




Figure 3.13: Latch storing switch configuration data.

Figure 3.13 illustrates the design used to store configuration data in the switch nodes. A TPDICE latch is adopted to provide memory functionality and the ability to recover from transient disruptions. Write enable circuitry allows the configuration to be updated via external control signals. Read buffers control the transmission gates that are responsible for data path switching. These buffers are arranged to assure that any SEU or SET affecting the configuration memory will disrupt at most one data line per redundant triplet. With this attribute in place, the system can assure recovery from any single-event disruption affecting the configuration memory.

3.3 Analysis and Simulation

This section presents analysis and simulation data to demonstrate the functionality of the radiation-tolerant reconfigurable DSP processor. Specifically, the operational characteristics of the LUT-based element, reconfigurable switches, and overall system are examined. Additionally, the response of the system to a particle-induced transient disruption is described.

3.3.1 Element Simulation

As described in the previous section, elements are LUT-based structures that form the processing core of the reconfigurable DSP processor. Each element consists of a 4x4 array 2-bit TPDICE SRAM latches. Elements are arranged in a 4x4 matrix to perform arithmetic or memory operations inside of every cell. Only the memory mode configuration is considered here in order to simplify the radiation tolerance analysis of this architecture.

In memory mode, each element receives four bits of address information and outputs two bits of data. The data bus lines of the LUT are precharged to $V_{dd}/2$ in order to mitigate the effect of read-induced SEU. Figure 3.14 illustrates these functional characteristics of the element. This simulation depicts multiple clock cycles, during each of which a read operation is performed. Clock (Cka), row address (I0, I1), data bus (Da<0>, /Da<0>), and output signals (Outa<0>, Outa<1>) are displayed for each operation. The row address lines are incremented every cycle to select a new memory location from which to read (column address bits are held constant). Data lines are precharged while the clock is high, and are driven by the enabled memory cells when the




Figure 3.14: Element simulation depicting clock, row address, data and output signals.

clock is low. Finally, new output data becomes valid at the end of every low phase of the clock. The two output bits present data from the memory cells at the specified LUT address location.

3.3.2 System-Level Data Transfer Simulation

The memory mode reconfigurable DSP processor configuration considered in this chapter consists of two LUT-based cells connected by a reconfigurable switch. This switch allows a number of data transfer configurations, including transfer directly from one cell to another. Figure 3.15 is a simulation that demonstrates this capability. Data is read from one cell, passes through the switch, and is written to the second cell. The simulation illustrates internal node voltages of the destination memory. Memory locations I5 and I6 transition from low to high, whereas locations I4 and I7 transition from high to low. The low to high transition completes first due to datapath characteristics and initial conditions. All operations resolve by the end of the clock period.




Figure 3.15: Illustration of destination memory cell node voltage levels during a data transfer operation.

3.3.3 System-Level SET Response

The reconfigurable DSP processor discussed in this chapter offers protection against particle-induced transient faults affecting memory and logic. SEUs directly affecting latches in this architecture have an effect that is described in Chapter 2 of this dissertation. On the other hand, SETs affecting logic may propagate to latches and produce effects similar to SEUs affecting memory. The TPDICE memory cells utilized in this reconfigurable architecture tolerate such effects. Figure 3.16 is a simulation of an SET affecting switching logic in the reconfigurable DSP processor. The simulation shows signal values of the destination memory latch. One input signal (InRa) is held low for the entire write operation. Node Va of the destination latch follows this input signal. However, after the clock signal is deactivated, all destination latch nodes resolve to the desired output state. Thus, the SET is tolerated and bypassed by the TPDICE memory configuration.




Figure 3.16: Simulation of an SET affecting a switch in the reconfigurable DSP architecture during a L -> H data transition.




Figure 3.17: Simulation of an SET affecting a switch in the reconfigurable DSP architecture during a H -> L data transition.

Figure 3.17 depicts the converse of the operation shown in Figure 3.16. A high to low write operation is illustrated, with one input held high. All destination memory latch voltage levels resolve to the desired state shortly after the write operation is completed. Destination nodes Va and /Vc require the most time to resolve, however they are driven quickly to the proper levels after the write operation commences.

3.4 Summary

The proposed SEU and SET-tolerant reconfigurable DSP processor offers a unique combination of flexibility, high-performance and radiation tolerance for space applications. The medium-grain architecture featuring 4-bit LUT-based cells is particularly suited towards DSP algorithms. High-performance and radiation-tolerance are provided by the TPDICE memory cell, which provides SEU/SET bypass capability.

Chapter 4

Fault-Tolerant Pipeline Latches and Flip-Flops

This chapter presents an analysis of SEU and SET-tolerant approaches to constructing pipeline latches and flip-flops. TPDICE and Barry-Dooley basic cells are utilized to achieve fault-tolerance and transient bypass capability [41, 44-45]. A number of single-ended and differential structures are presented and evaluated with respect to performance and energy consumption. The general pipeline memory categories considered in this paper include level-sensitive latches, edge-triggered master-slave flip-flops, and pulse-triggered flip-flops. Performance and energy consumption evaluations are based off simulations performed in 90nm CMOS.

As of now, substantial research has focused on producing SEU and SET-tolerant SRAMs [5-11]. On the other hand, little research has focused on protecting pipelined systems. These architectures exist in mission critical designs, and so they must possess high reliability. When considering fault-tolerance, pipelined systems require protection for their pipeline latches/flip-flops. Additionally, the ability to bypass SEUs and SETs is critical for performance, as it allows a 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 2ns), and so high-performance systems cannot afford to pause operations while transient pulses are present [12-15].

4.1 Fault-Tolerant Pipelined Systems

Pipelined systems, such as processor data paths and ALU functional units, consist of uni-directional memory cells and combinational logic. The memory cells do not share a data bus, so they need no read enable capability. Modern systems are trending toward increasing levels of superpipelining, and so thinner logic stages are used between memory cells. Because of this, the clock period constraints for these systems are becoming dominated by the memory delays. Therefore, the performance aspects of these memories are extremely important. Additionally, improvements in manufacturing technology result in circuits that are more fragile, which increases the probability that faults may corrupt pipeline data. Fault-tolerant approaches, such as the schemes illustrated in this report, can mitigate this effect.

4.1.1 Fault-Tolerant Structures

The focus of this chapter is on designs that utilize SEU and SET-tolerant basic cells as building blocks to construct pipeline memory structures. This approach was chosen due to its modular architecture and circuit-level nature. The modular architecture allows pipeline memory to be constructed with any fault-tolerant SRAM cell placed in essentially any latch or flip-flop configuration (assuming the benefits of both the cell and the configuration). The circuit-level nature provides a number of benefits, including quick fault recovery, ease of integration into existing designs, and flexibility to adjust the tradeoff between performance, energy consumption, and size.

Specifically, TPDICE and Barry/Dooley SEU/SET-tolerant SRAM cells have been chosen for use in these pipeline memory designs. These cells possess transient bypass

capability, allowing them to tolerate SEUs and SETs without pausing until they are dissipated by the system. These cells transmit and receive sufficient information to allow operations to proceed even when one data signal is disrupted. Due to this capability, these designs do not require clock overhead equal to the minimum SEU/SET width. In contrast, designs without bypass capability do require this overhead. In space, cosmic ray strikes with LET of 100 MeV-cm²/mg can induce SET pulses up to 2ns long [12-14]. Transient bypass capability is extremely beneficial in these situations, as a clock overhead of 2ns would substantially hamper performance.

The TPDICE SEU and SET-tolerant SRAM cell is a novel design presented in Figure 2.11. It can be thought of as a regular DICE cell extended from four nodes to six, or SET-tolerant TMR with the triplicated voting integrated directly into the memory structure. Since the TPDICE design requires no external voting, it benefits from performance, energy consumption, and size advantages over TMR. Additionally, this design retains one of the most valuable TMR attributes, which is the ability to bypass transient disruptions. This attribute is extremely desirable in performance-oriented pipeline systems that must function in the presence of highly energetic radioactive particles.

Integration of TPDICE memory circuits into pipelined systems allows for the speed benefits of TMR without the need for majority voters. A graphical depiction of this concept is shown in Figure 4.1. TPDICE latches can be set up with one clock phase or in dual-phase master-slave configurations. Single-phase setups are less complex, while dual-phase designs place less restrictive timing constraints on combinational logic.

Additionally, each logic block can be pipelined through traditional unprotected means, as long as a maximum of one SET is expected within the latency of the pipelined

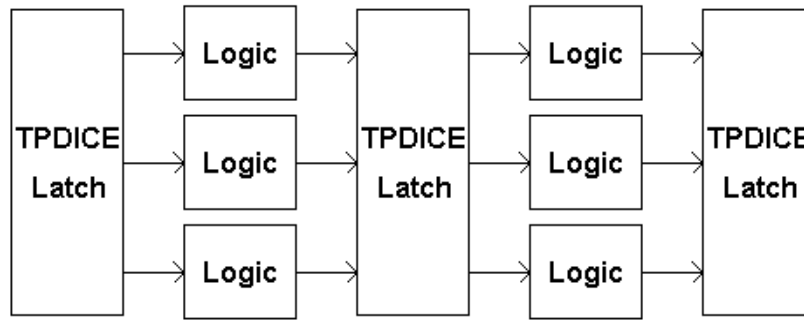


Figure 4.1: System pipelined with TPDICE latches.

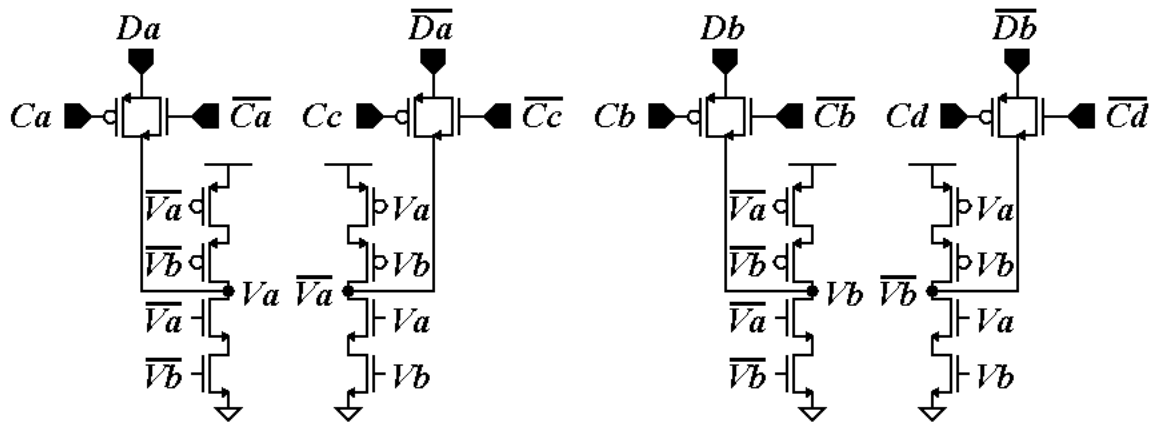


Figure 4.2: Barry/Dooley differential SEU and SET-tolerant memory cell.

block. A TPDICE cell at the end of the unprotected logic block performs a majority voting function to obtain the correct data output.

The Barry/Dooley basic cell can be substituted for its TPDICE counterpart in differential systems that do not require cross-coupled signal paths. Figure 4.2 is a schematic of the Barry/Dooley cell [44-45]. It consists of four nodes, each controlled by an NMOS and PMOS stack of two transistors each. A transient disruption affecting one node in the cell would tri-state the two other nodes of opposite polarity. The final node would not be affected. As a result, the voltage levels of the three nodes not directly

affected by the disruption are maintained at the correct levels. Because of this, sufficient information is preserved to bypass transient disruptions.

The Barry/Dooley design possesses less complexity when compared to the TPDICE approach, as it has only four internal nodes. However, it requires independent signal paths and combinational logic. This suits it towards designs that prioritize minimal complexity and do not use cross-coupling between true and complement data signals.

4.1.2 Pipeline Latches vs. Edge-Triggered Flip-Flops

The choice of pipeline memory architecture made by the designer is dependent upon the requirements of the system. Pipeline latches have the potential to achieve a lower power- delay product, while edge-triggered flip-flops provide greater protection against data race conditions. Latches work best in situations where logic min/max delays can be tightly constrained. On the other hand, flip-flops are the best choice when it is not practical to place tight timing constrictions on logic.

Latches are the most basic pipeline memory design, and they are advantageous in terms of performance, power consumption, and complexity. These structures are typically formed with a memory element that is fed by write enable circuitry. The strong attributes of this design are obtained via a one-stage architecture, simple clocking, and level-sensitive sampling. However, these benefits come at the cost of a large transparency window.

Pipeline flip-flops are desirable in situations where timing is critical and logic min/max delays invite the possibility of race conditions in basic latches. Latch/flip-flop transparency windows provide an opportunity for race conditions to occur, and so they should be minimized. During transparency, input data is allowed to pass through to the




Figure 4.3: Basic flip-flop timing diagram.

outputs of a memory cell. Timing faults may occur when this data has enough time to propagate through the subsequent logic stage and be latched by the next memory cell. Latches have significant transparency windows, which are often comparable in width to the clock period. On the other hand, flip-flops have much smaller (or non-existent) transparency windows. Thus, flip-flops are more successful at preventing race conditions and the timing faults associated with them.

Figure 4.3 is an illustration of basic flip-flop timing. Input sampling is performed at the rising edge of the clock. Data must be stable for a setup time before the active edge, and also for a hold period after this edge. The Clk-Q delay is defined as the period between the active edge and the point in time when output data achieves stability. The D-Q delay is the period between the beginning of the setup time and the occurrence of output data stability. When constructing timing budgets, the total flip-flop delay is calculated as the greater of the D-Q delay or the setup plus hold time. This is the value

that is subtracted from the clock period to account for data propagation through a flip-flop.

There are a number of ways to construct flip-flops, including master-slave and pulse-triggered flip-flop configurations. Master-slave designs utilize two latches in series in an attempt to reduce the transparency window. If two non-overlapping clocks are used to strobe the latches, then the transparency window can be reduced to zero. However, if clock skew induces overlap in the two clocks, then transparency can occur. Pulse-triggered flip-flops utilize a sampling pulse to reduce the transparency window of a standard latch. This sampling pulse enables write operations to the latch. Its width is usually set approximately equal to the sampling window of the flip-flop (setup time plus hold time). The sampling window is often substantially smaller than an active clock phase, which makes this technique effective at reducing the transparency window of a standard latch.

Pulse-triggered flip-flops possess timing characteristics that are logically analogous to those of standard edge-triggered flip-flops, although they are physically different [46]. Figure 4.4 depicts this. The rising edge of the clock initiates a sampling pulse, which enables the output latch. Data is physically captured at the falling edge of the sampling pulse as opposed to the rising edge of the main clock. Therefore, the timing parameters are shifted to center around this new active sampling edge. The setup and hold times form the sampling window around the new active edge. The Clk-Q delay begins at the active edge and ends at the point where the output data is stable. And as before, flip-flop performance is defined by the D-Q delay, which is the period of time between required input data stability and output data stability.




Figure 4.4: Pulse-triggered flip-flop timing diagram.

To conform to a timing methodology consistent with other flip-flops, pulse-triggered flip-flop timing can be referenced to the active edge of the main clock. This is the approach adopted by this paper. The resulting setup time is equal to the pulse-referenced setup time minus the time difference between the rising clock edge and the falling edge of the sampling pulse. The new hold time equals the pulse-referenced value plus the clock to pulse time differential. In many cases the new setup time will be negative, although the sampling window (the sum of the setup and hold times) will always remain the same. Clk-Q values increase, but they are generally not useful as a timing measurement [46]. On the other hand, the more revealing D-Q values remain the same.

Every pipeline memory design possesses strengths and weaknesses that suit it towards a particular application. Performance, power consumption, complexity, and




Figure 4.5: Fault-tolerant single-ended TPDICE latch (SETL).

transparency window size tradeoffs are made in all of these circuits. The subsequent section looks at specific fault-tolerant pipeline latches and flip-flops.

4.2 Single-Ended Fault-Tolerant Pipeline Memories

4.2.1 TPDICE Latch

A basic single-ended fault-tolerant latch can be formed using TPDICE memory as its core. This single-ended TPDICE latch (SETL) is shown in Figure 4.5. It relies on a transmission gate access stage that feeds a TPDICE memory cell. Output buffers are used to isolate the memory from subsequent pipeline circuitry. This design is simple, possessing low complexity, delay, and power consumption. However, it suffers from a large transparency window due to the fact that it samples data for an entire clock phase.

The purpose of the TPDICE latch circuitry is to recover from faults and save the state stored by the latch. Additional NMOS transistors are used in the output latch to facilitate write operations. These transistors are disabled when the clock is high, which allows TPDICE nodes to be pulled high without resistance. The output latch is in parallel with the output buffers in this architecture. Unlike in an SRAM design, the output latch is not responsible for driving the output buffers. Instead, the input buffers assume this responsibility. Because of this, it is optimal to minimize the TPDICE transistor widths, as this reduces the capacitance in the signal path.

4.2.2 TPDICE Master-Slave Flip-Flop

In many circumstances, it is impractical or impossible to tightly constrain the timing parameters of logic in a pipelined system. This may be due to the nature of the design, the timing variability of potential logic circuits, clock skew, etc. Race conditions and timing faults must be avoided at all costs, and so logic timing is a critical design parameter. In these cases, it is often advisable to adopt edge-triggered flip-flops in place of level-sensitive latches. Flip-flops maintain transparency windows that are only a fraction of a clock phase, as opposed to the full clock phase of transparency that latches exhibit. Smaller transparency windows provide less opportunity for data race conditions to occur.

Master-slave configurations are perhaps the most straightforward edge-triggered flip-flop architecture. These designs consist of a master latch that samples input data during the first clock phase, and a slave latch that releases the data during the second phase. Ideally, input data does not affect the outputs until after it is latched. This reduces the transparency window to zero. However, it should be noted that a significant transparency

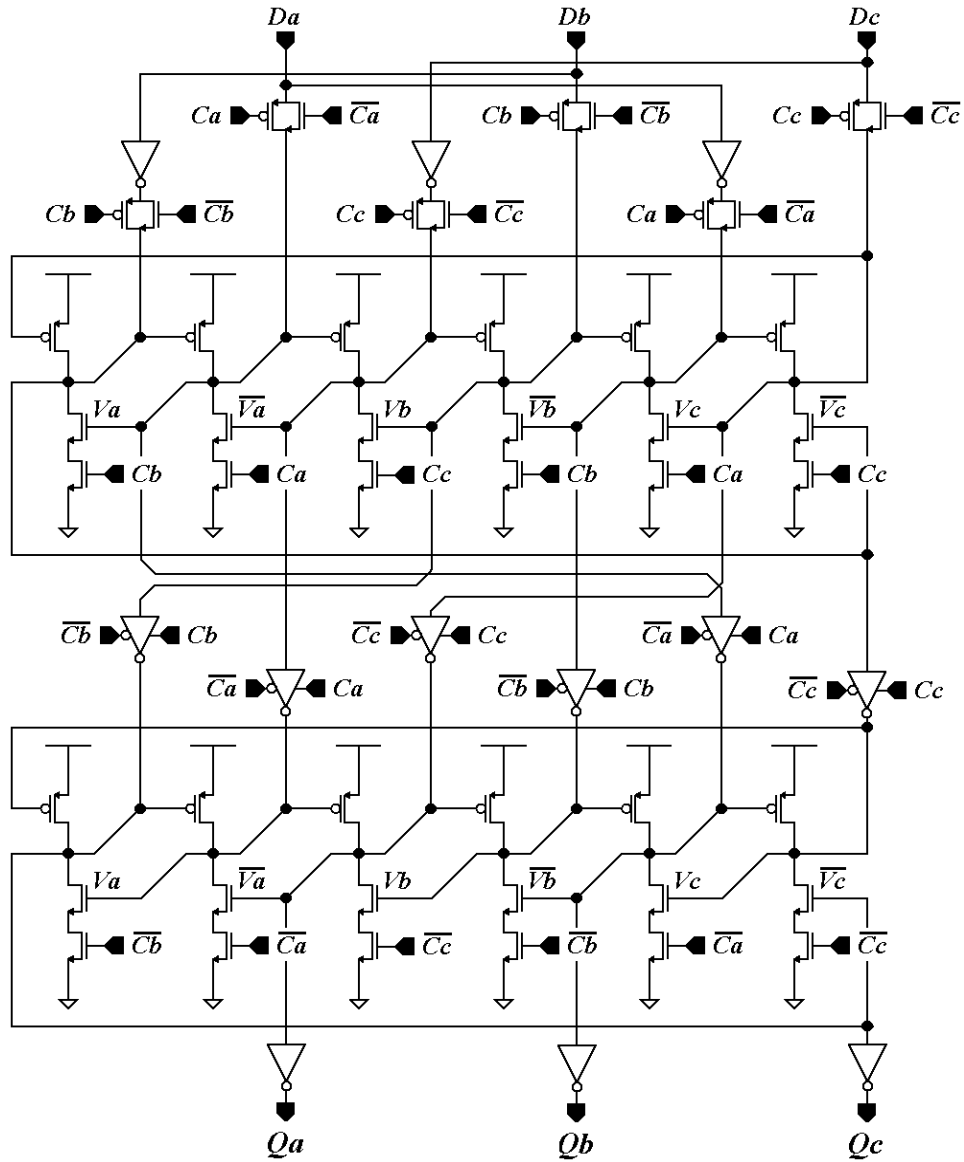


Figure 4.6: Single-ended TPDICE master-slave flip-flop (SETMSFF).

window may appear if clock skew becomes a factor. In any case, master-slave designs are generally very effective at protecting against data race conditions and timing faults.

A basic master-slave fault-tolerant flip-flop can be formed by utilizing a TPDICE latch in both the master and slave stages. Figure 4.6 illustrates this single-ended TPDICE master-slave flip-flop (SETMSFF). Transmission gates enable write operations, while buffers provide drive strength to the outputs. Tri-state buffers facilitate sample release

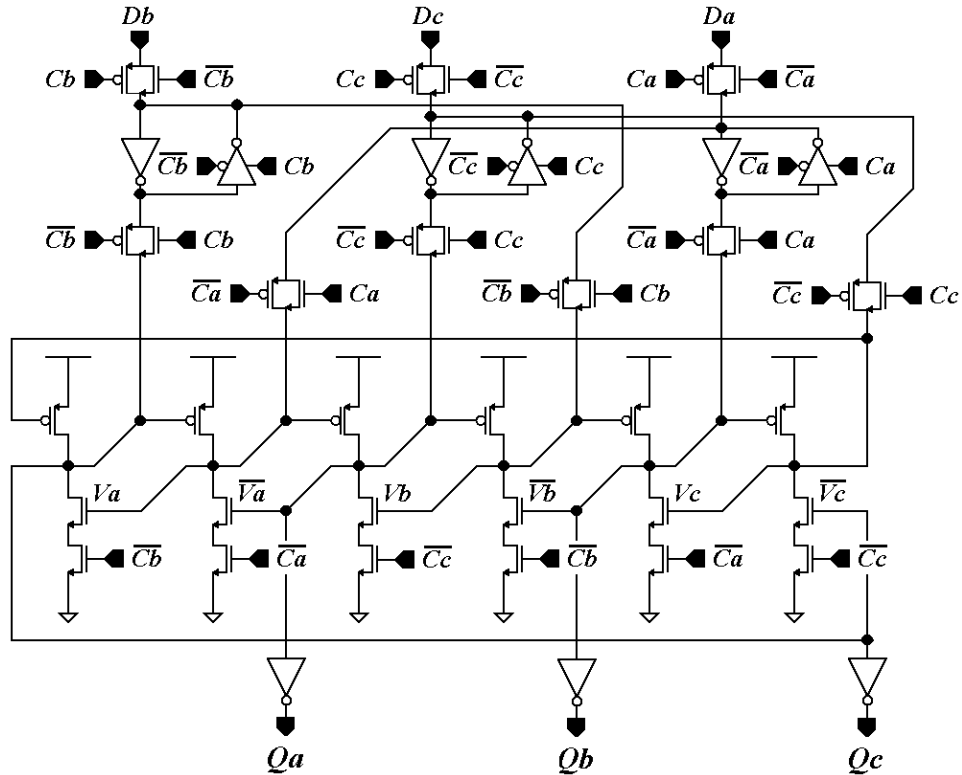


Figure 4.7: Single-ended TPDICE alternative master-slave flip-flop 1 (SETAMSFF1).

operations. Faults are corrected in both phases of the clock cycle because TPDICE latches are used in both stages of this flip-flop. This is an important property, as designs with one TPDICE latch in only the master or the slave section can correct faults in just one clock phase.

Figure 4.7 illustrates the case where a triplicated standard latch is used in the master section and a TPDICE latch in the slave. We refer to this configuration as the single-ended TPDICE alternative master-slave flip-flop 1 (SETAMSFF1). This master latch design is similar to that of the PowerPC flip-flop described in [47]. SEUs and SETs affecting the master section may compromise at most one of the three standard latches. The TPDICE slave latch would correct this during subsequent cycles. Due to this delay in correction, this configuration does not correct transient pulses during the sample release

phase, which is the initial phase during which output data becomes valid. This may result in faults that propagate indefinitely inside of systems that utilize logic with nearly the maximum possible delay in every pipe stage. However, this design would be suitable in systems with an appreciable mix of short logic delays, as data corrected after the sample release phase would have time to propagate to the following pipeline stages. The advantage of this configuration is performance. The simple master stage presents no conflict during write operations, and it can be sized to improve the efficiency of data transfer to the slave latch.

A compromise between the previous two configurations can be achieved by placing a TPDICE latch in the master stage and a standard latch in the slave stage. This provides an improvement over the performance of the SETMSFF, while maintaining the ability to correct faults during the sample release stage. However, the performance improvement is not as great as that of the SETAMSFF1. This is due to the fact that the TPDICE master latch presents some conflict during write operations, and it is not as effective as a standard latch at transferring data to the slave stage. Figure 4.8 illustrates this approach, which we refer to as the single-ended TPDICE alternative master-slave flip-flop 2 (SETAMSFF2).

4.2.3 TPDICE Pulse-Triggered Static Flip-Flop

While master-slave flip-flops are effective at preventing timing faults, they obtain this ability by sacrificing performance. By instead selecting pulse-triggered flip-flops, timing fault mitigation can be realized without this performance cost. Single-ended SEU and SET-tolerant pulse-triggered flip-flops can be constructed by adopting a triplicated sampling stage that feeds a TPDICE output latch. A block diagram of this is shown in

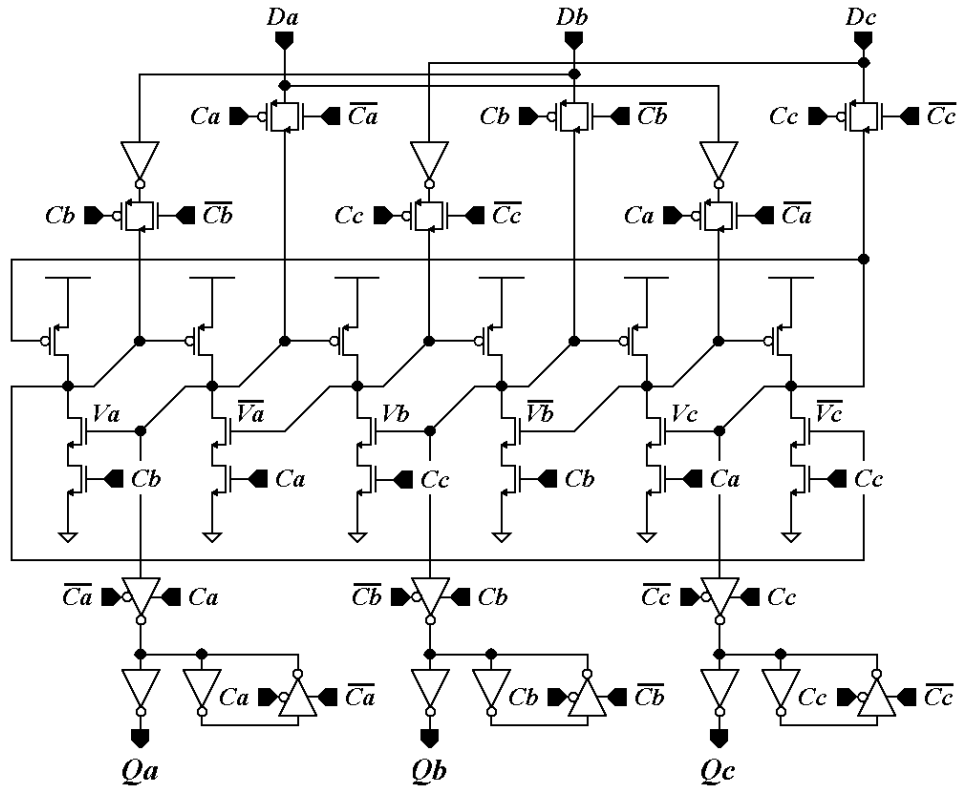


Figure 4.8: Single-ended TPDICE alternative master-slave flip-flop 2 (SETAMSFF2).




Figure 4.9: TPDICE-based flip-flop architecture block diagram.

Figure 4.9 [48]. Any single-ended sampling stage can be used in this configuration. Each possibility offers a balance between performance, sampling window size, power

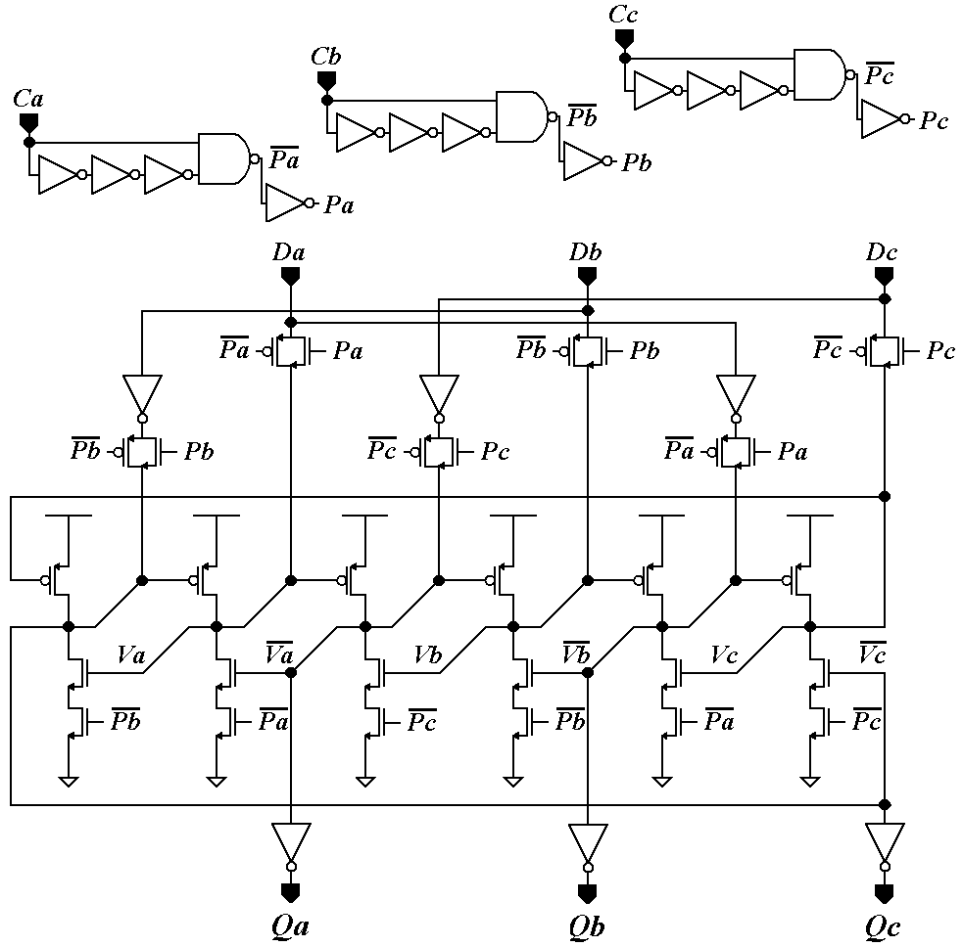


Figure 4.10: Single-ended TPDICE pulse-triggered flip-flop utilizing explicit sampling pulse generators (SETPFF).

consumption, and circuit complexity. It is up to the circuit designer to choose an approach that satisfies the demands of the situation.

A basic single-ended pulse-triggered flip-flop (SETPFF) design featuring a TPDICE output latch is shown in Figure 4.10. This design utilizes NAND gates and inverters to generate pulses that define the sampling window. This approach to pulse generation and sampling is analogous to the explicit pulsed static flip-flop described in [49]. The width of write operations decreases from an entire clock phase (for a basic latch) to the duration of the sampling pulse. This reduction in the width of write operations reduces the

transparency window, resulting in less opportunity for input data to pass through to the output of the circuit.

The explicit pulse generating circuitry can be shared amongst a group of flip-flops to reduce complexity and power consumption. The group must be small, or else device variations, noise, skew, and interconnect length would significantly degrade the quality of the generated pulse [49].

4.2.4 TPDICE Dual-Pulsed Semidynamic Flip-Flop

In the critical path of a digital system, additional speed is often desired at the expense of higher power consumption. By adopting semidynamic circuitry, this tradeoff can be realized. Performance levels even higher than that of standard latches can be achieved with this approach. Semidynamic flip-flops often possess a dynamic input stage that feeds a static memory cell. The dynamic input stage is precharged while the clocking signal is inactive. When the clock goes active, data is sampled and stored in the static output stage.

Figure 4.11 shows a semidynamic single-ended TPDICE dual-pulsed flip-flop (SETDPFF). The dual-pulsed input stage is a design that was presented in [50]. The first triggering pulse precharges the access stage, and input evaluation occurs during the second pulse. NMOS enable transistors allow the precharge to proceed without conflict. Overlap may occur between the two pulses without affecting functionality. Unlike in pure dynamic structures, data is never floating in this design. A transistor is driving all nodes in the circuit at all times, reducing the effect of noise and accumulated radiation.

The pulse generators described in [50] create the pulses needed for enabling the precharge and evaluation phase in the input stage. The generators are explicit, and they

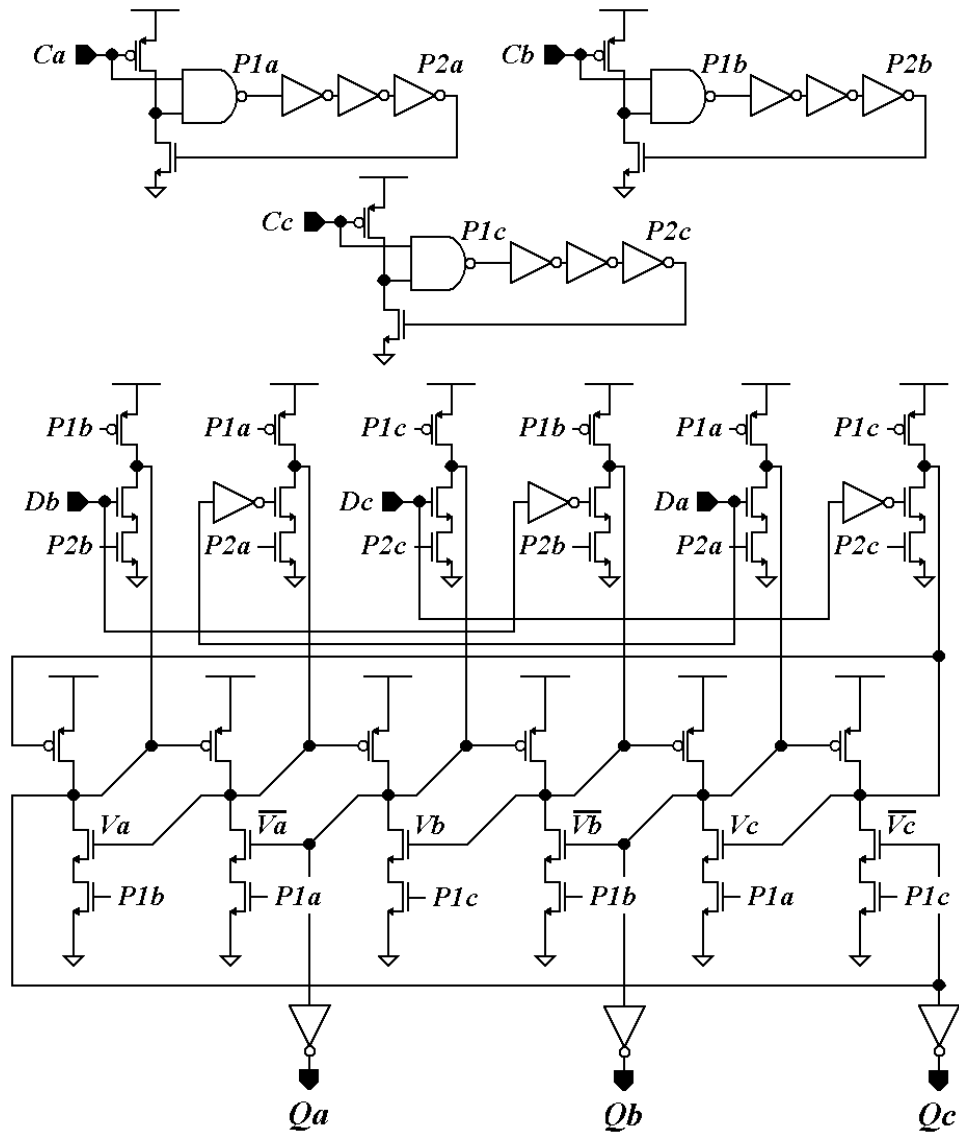


Figure 4.11: Single-ended TPDICE dual-pulsed flip-flop (SETDPFF).

create both the precharge and evaluate pulses, which reduces area and power consumption. Additionally, generators can be shared by a group of flip-flops, which further improves efficiency.

Speed in this design is achieved through low data and clock input capacitance, a thin path between the inputs and the outputs, and the fact that only one internal signal transition needs to be optimized. Each data and clock line is connected to only three transistor gates, resulting in a low load on these inputs. The data critical path consists of a

two-transistor NMOS stack that connects directly to the output buffer in parallel with the TPDICE output latch. This compact critical path allows for low D-Q delay. Finally, since all of the internal flip-flop nodes are precharged, only one internal signal transition needs to be optimized, and precharge transistors can be of minimum width.

This circuit is particularly useful when paired with dynamic domino logic-based systems. The precharge-evaluate architecture can be mated directly to domino logic on the inputs and outputs. Additionally, domino logic removes the need to deal with the output glitch created by the precharge phase of the flip-flop. If domino logic is not used, then inverting logic or additional setup/hold time requirements must be enforced to prevent unwanted timing faults.

4.3 Differential Fault-Tolerant Pipeline Memories

4.3.1 Barry-Dooley Differential Latch

A low-area differential pipeline latch can be created using the Barry-Dooley fault-tolerant cell as the memory stage [44-45]. This differential Barry-Dooley latch (DBDL), shown in Figure 4.12, requires only four internal memory nodes. Transmission gates enable write operations, while output buffers provide drive strength to the subsequent stage. If a fault affects one data path or internal node, the Barry-Dooley cell will contain this disruption to one node, and then correct it. Since at most only one node is affected at a time, sufficient information is transmitted to the subsequent stage to bypass transient disruptions.

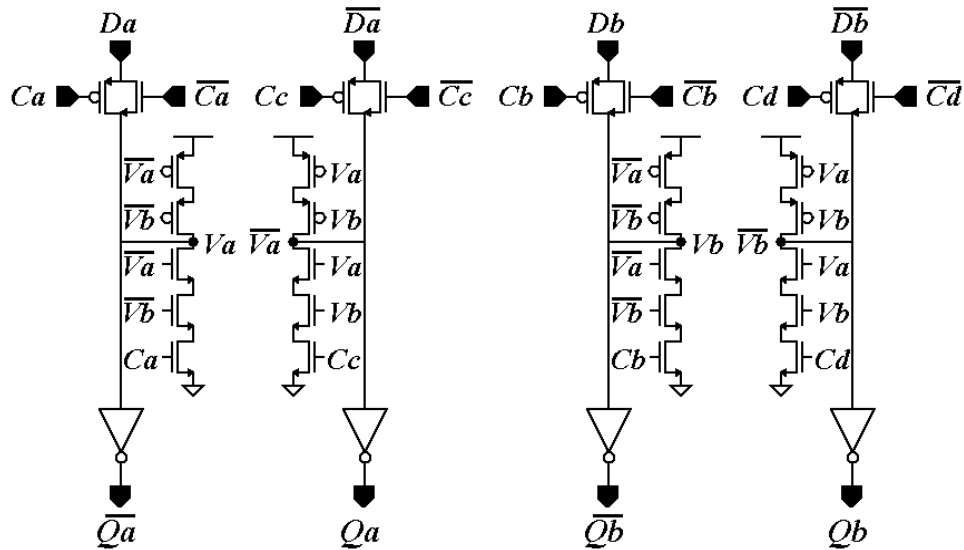


Figure 4.12: Differential Barry-Dooley pipeline latch (DBDL).

Power consumption and circuit complexity benefits can be realized with this circuit when compared to six-node designs. However, all four inputs must be completely independent. Thus, cross-coupled logic cannot be used with this design.

This latch is a desirable choice when simplicity is required, independent differential data paths are available, and logic timing can be tightly constrained. Independent differential data paths are available when logic is not shared between true and complement paths and cross-coupling is not implemented. Simplicity in this design is achieved through the adaptation of a level-sensitive latch structure and the use of four internal nodes, which is the minimum number necessary to construct balanced fault-tolerant SRAM.

4.3.2 Barry-Dooley Differential Master-Slave Flip-Flop

As with their single-ended counterparts, differential pipeline latches can be susceptible to timing faults if logic delays cannot be tightly constrained. One solution to this problem is the adoption of pipeline flip-flops instead of level-sensitive latches. The

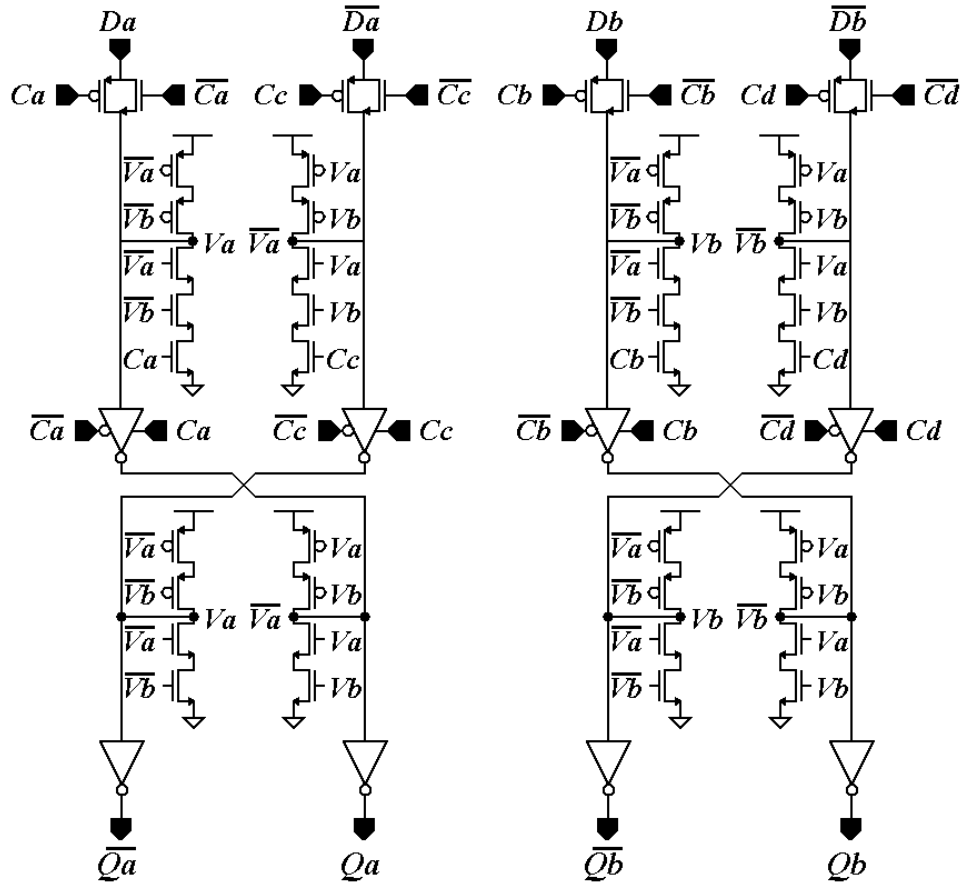


Figure 4.13: Differential Barry-Dooley master-slave flip-flop (DBDMSFF).

Barry-Dooley latch described above can be converted into a flip-flop by adding an additional latch stage to the input of the circuit. This differential Barry-Dooley master-slave flip-flop (DBDMSFF) is shown in Figure 4.13. Similar to the previous design, an SEU or SET affecting this flip-flop will disrupt only one node of its Barry/Dooley output latch. This includes faults captured by the input latch. In any case, the output latch will bypass and correct the transient disruption.

This flip-flop is best suited towards systems that would work with the Barry/Dooley latch, but do not feature logic with tightly bounded delays. These systems favor simplicity and possess independent data paths. The use of only four inputs, transmission gate enable circuitry, and a basic SRAM master latch provide the simplicity in this

design. However, shared logic and coupled data paths must be avoided to maintain independent data paths.

4.3.3 TPDICE Differential Latch

In differential structures, a requirement for independence between true and complement data paths often negatively affects performance and complexity in a design. Cross-coupled and sense-amplifier logic configurations can be adopted to amplify differential signals in such a way as to improve performance. Additionally, logic can be shared amongst true and complement paths, reducing complexity.

Barry-Dooley based differential pipeline memories require all inputs to be independent, which places a significant restriction on the logic style used in the pipelined system. By substituting the TPDICE design for its Barry-Dooley counterpart, this restriction on logic is removed. At first glance, simplicity is reduced due to the use of six data paths instead of four. However, this factor can be overcome by the advantages attainable in the logic and data transmission circuitry.

Figure 4.14 is a diagram of a differential TPDICE latch (DTL). This structure functions in a similar fashion as its single-ended counterpart, save for the differential input and output circuitry. All SEUs and SETs affecting this circuit can be tolerated. A transient pulse affecting a differential set of input signals is spread to a pair of offset nodes in the TPDICE cell, allowing straightforward recovery. Likewise, transient pulses directly affecting the latch disrupt one TPDICE node, which are tolerated by design. All things considered, this latch offers simplicity and high performance in a differential configuration, at the cost of a large transparency window.

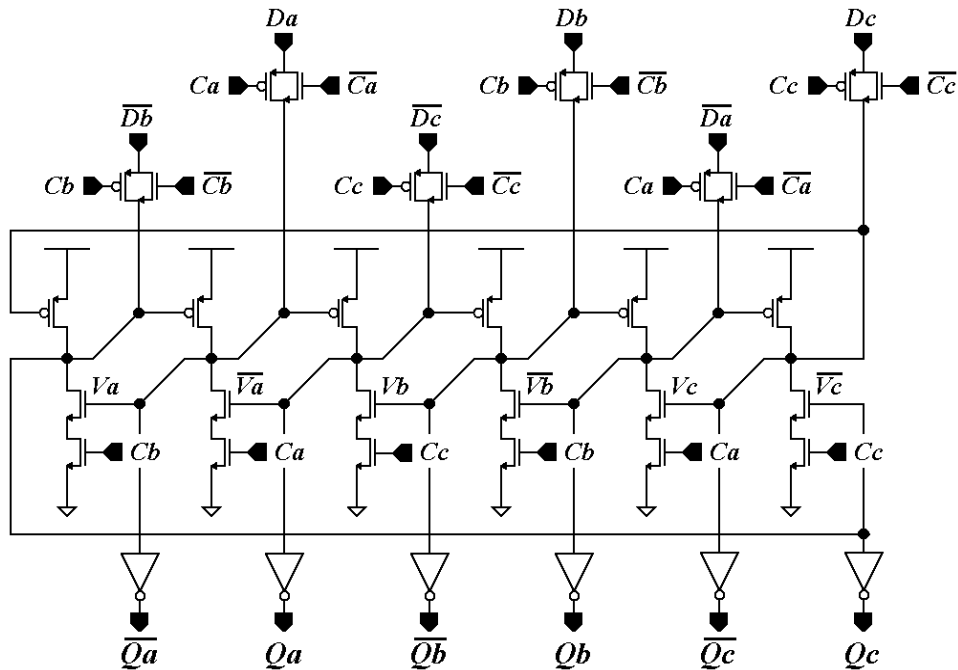


Figure 4.14: Differential TPDICE latch (DTL).

4.3.4 TPDICE Differential Master-Slave Flip-Flop

Figure 4.15 depicts a differential TPDICE master-slave flip-flop (DTMSFF) with TPDICE latches used in both the master and slave stages. This configuration corrects faults during both phases of the clock cycle. The core of this design is equivalent to its single-ended counterpart. Complementary inputs and outputs are incorporated to provide differential I/O functionality. Transmission gates serve as read enable circuitry, and inverters are used as output buffers. Tri-state inverters facilitate data transfer between the master and slave latches, providing enable capability and signal drive strength.

Performance can be improved by adopting unprotected latches in the master stage if transient upsets do not need to be corrected immediately in the sample release stage. This is also the case with the single-ended version of this flip-flop. One example of such a design is shown in Figure 4.16. We refer to this configuration as the differential TPDICE

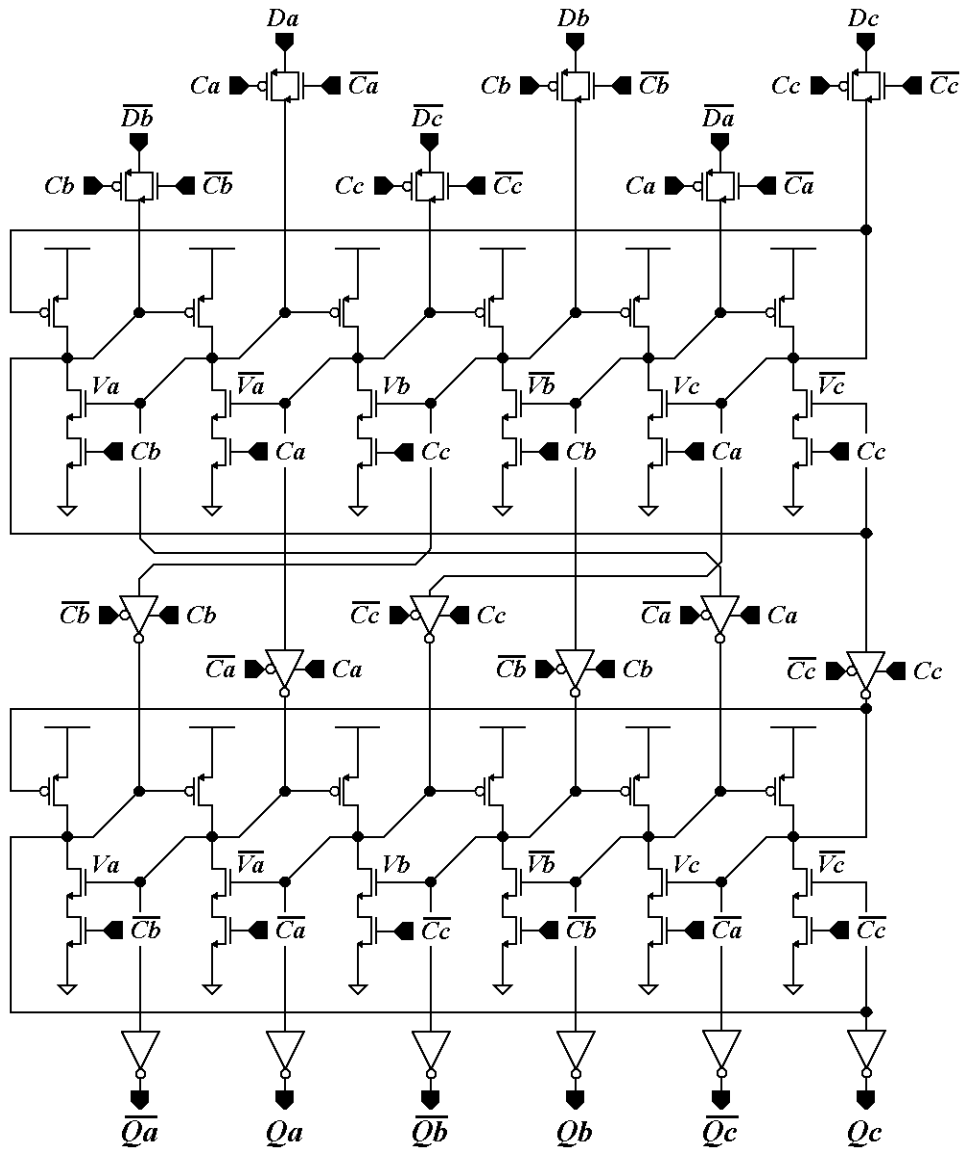


Figure 4.15: Differential TPDICE master-slave flip-flop (DTMSFF).

alternative master-slave flip-flop 1 (DTAMSFF1). The master stage utilizes the architecture from [51]. This master latch design possesses low complexity, as it uses NMOS pass transistors as enable circuitry and only one inverter per data path. Cross-coupled PMOS transistors provide keeper functionality and restore data to the positive supply rail. NMOS transistors disable keeper functionality to facilitate write operations.

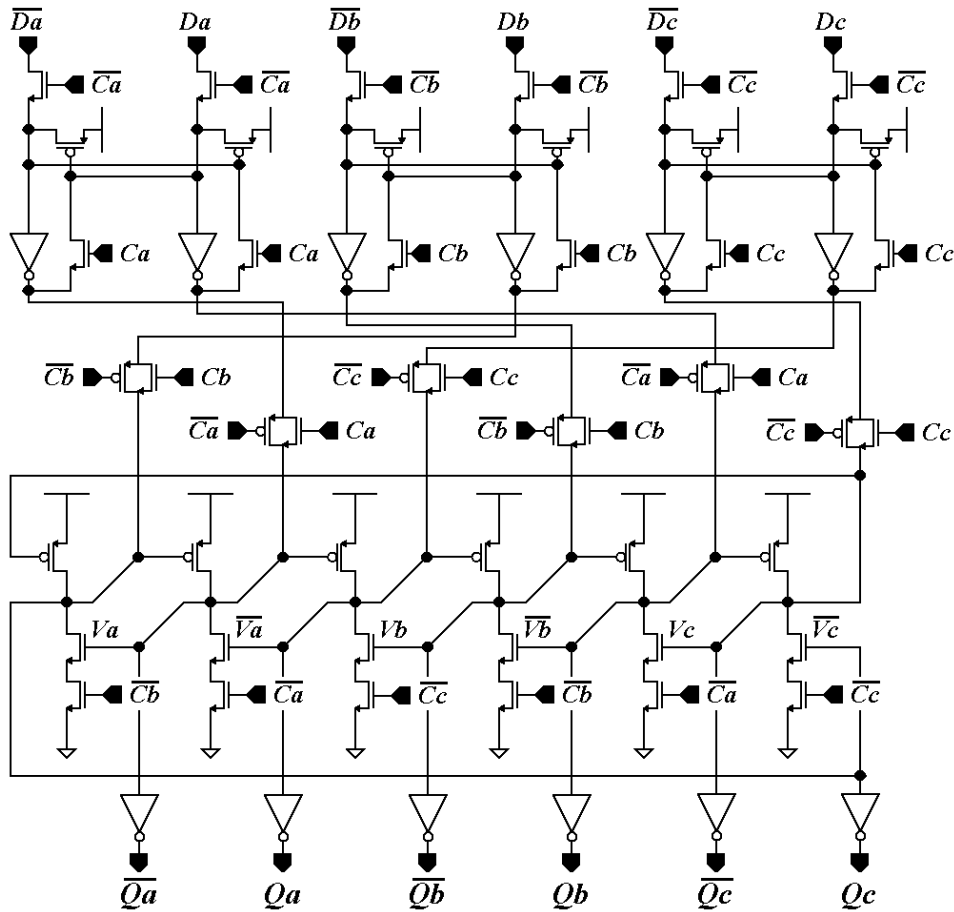


Figure 4.16: Differential TPDICE alternative master-slave flip-flop 1 (DTAMSFF1).

As with single-ended designs, placing a TPDICE latch in the master stage and an unprotected design in the slave stage results in a compromise between fault recovery time and performance. This setup achieves fault recovery in the sample release stage, as well as greater performance than using the TPDICE latch in both stages. Figure 4.17 depicts this approach, which we call the differential TPDICE alternative master-slave flip-flop 2 (DTAMSFF2).

4.3.5 TPDICE Differential Pulse-Triggered Static Flip-Flop

Figure 4.18 is a schematic for a differential version of the pulse-triggered flip-flop (DTPFF) introduced previously. Complement inputs and outputs are substituted for the

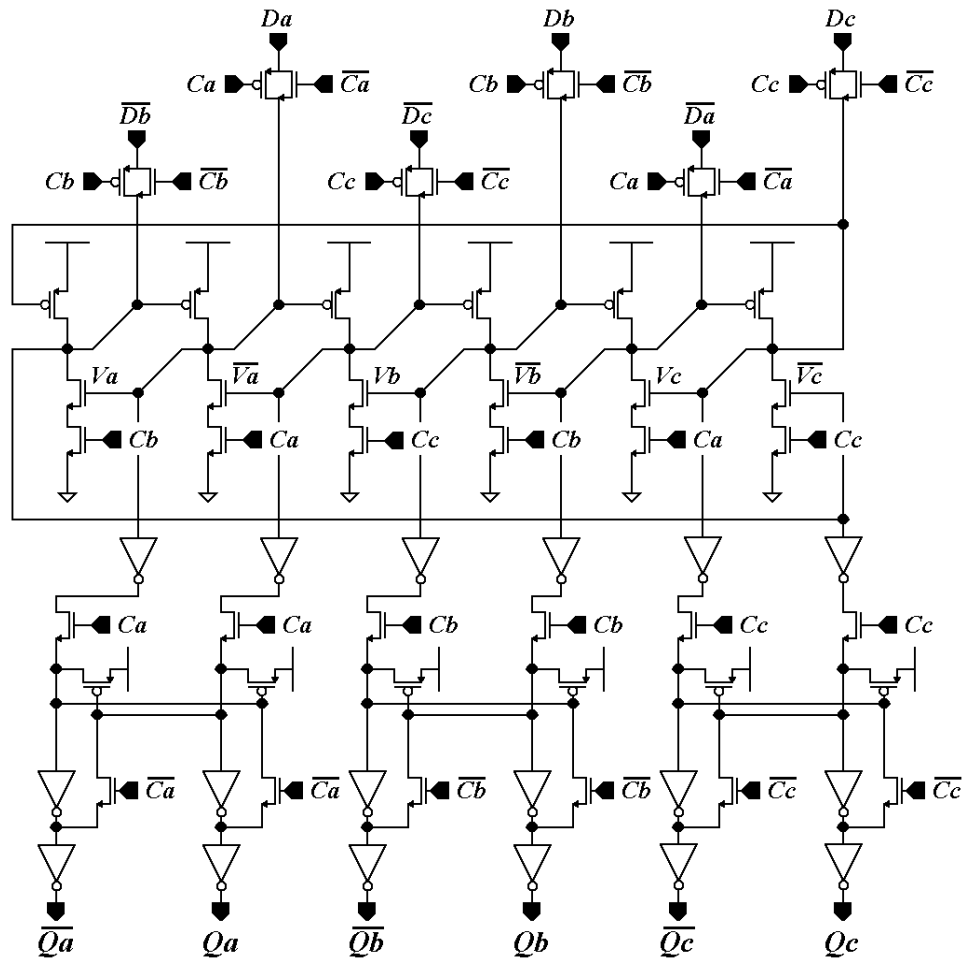


Figure 4.17: Differential TPDICE alternative master-slave flip-flop 2 (DTAMSFF2).

single-ended data paths used above. This configuration does not require inverters to generate complement inputs, thereby reducing the depth of the critical path. Performance is improved by this optimization.

4.3.6 TPDICE Differential Dual-Pulsed Semidynamic Flip-Flop

The single-ended dual-pulsed semidynamic flip-flop presented in the previous section can be extended to support differential data paths. This is accomplished by incorporating complement inputs and outputs to the circuit. A depiction of this differential TPDICE dual-pulsed flip-flop (DTDPFF) is shown in Figure 4.19.

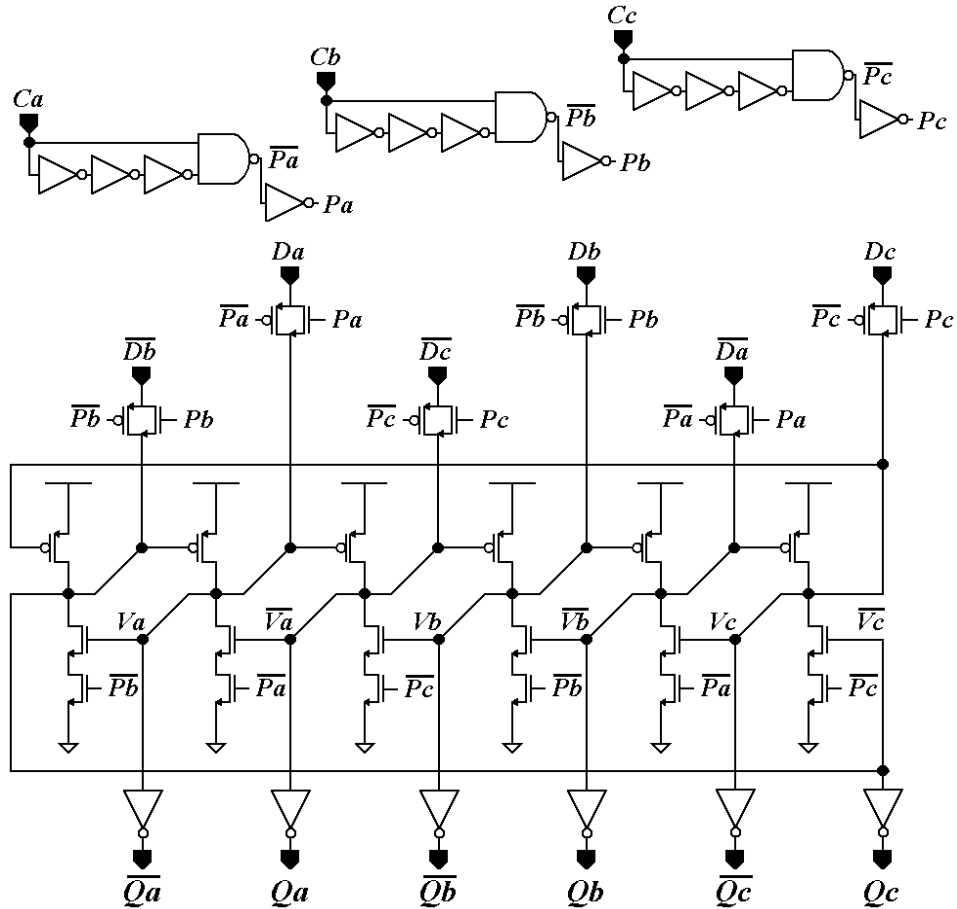


Figure 4.18: Differential TPDICE pulse-triggered flip-flop (DTPPF).

This version of the circuit benefits from balanced data transitions, allowing it to accept input data later in a write operation than its single-ended counterpart. Performance benefits are realized from this, as the reduced setup time decreases the overall delay between the arrival of input data and validity of output data.

On the other hand, the differential nature of this design results in increased average energy consumption. This is due to the fact that glitches appear on its outputs during every write operation. In contrast, glitches only affect the single-ended version during write operations that activate the circuit's enable transistors. Since the differential circuit

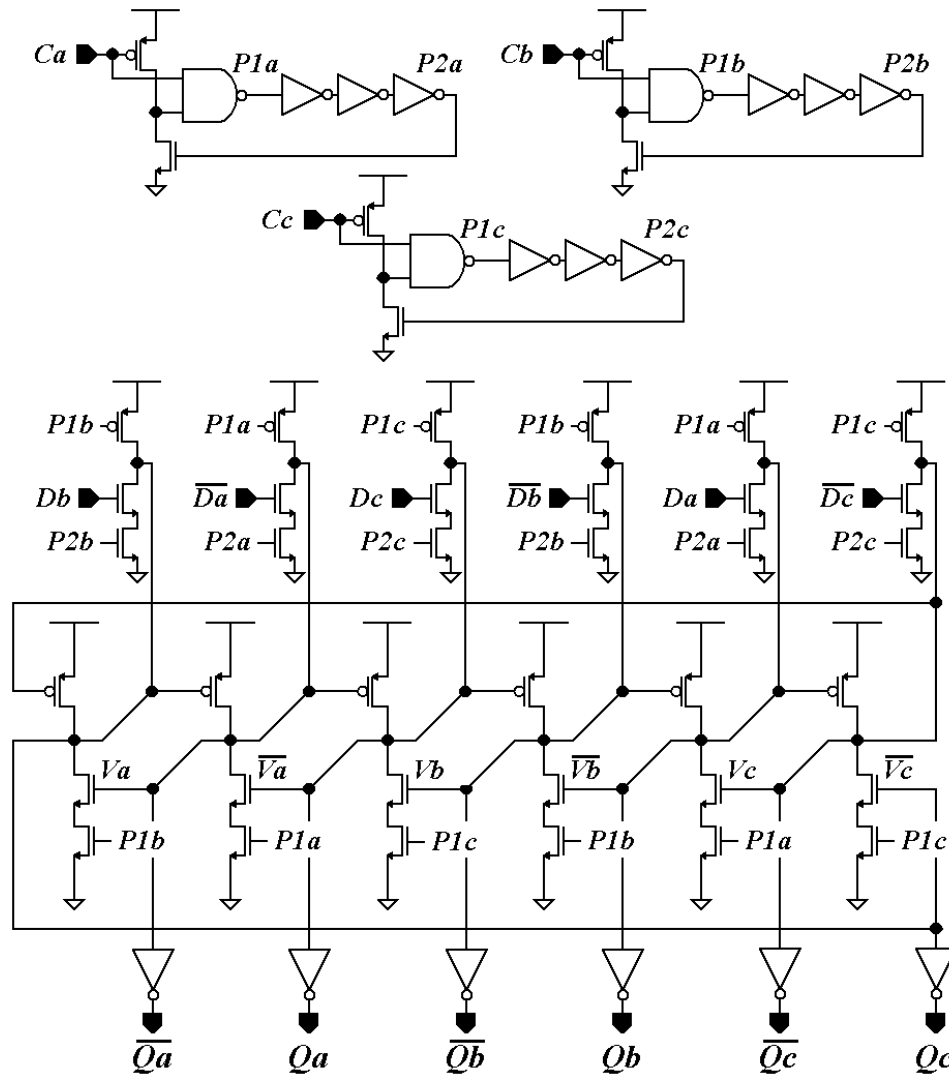


Figure 4.19: Differential TPDICE dual-pulsed flip-flop (DTDPFF).

experiences twice the number of glitches as its single-ended counterpart, it consumes additional energy.

4.3.7 TPDICE Differential Sense Amplifier Flip-Flop

Another robust class of fully-differential pipeline flip-flops can be formed by utilizing sense amplifiers as the front end to a TPDICE output latch. Such a design would exhibit a very short sampling window, as sense amplifiers can sample very small voltage deltas in differential data. Figure 4.20 illustrates this differential TPDICE sense-amplifier flip-flop

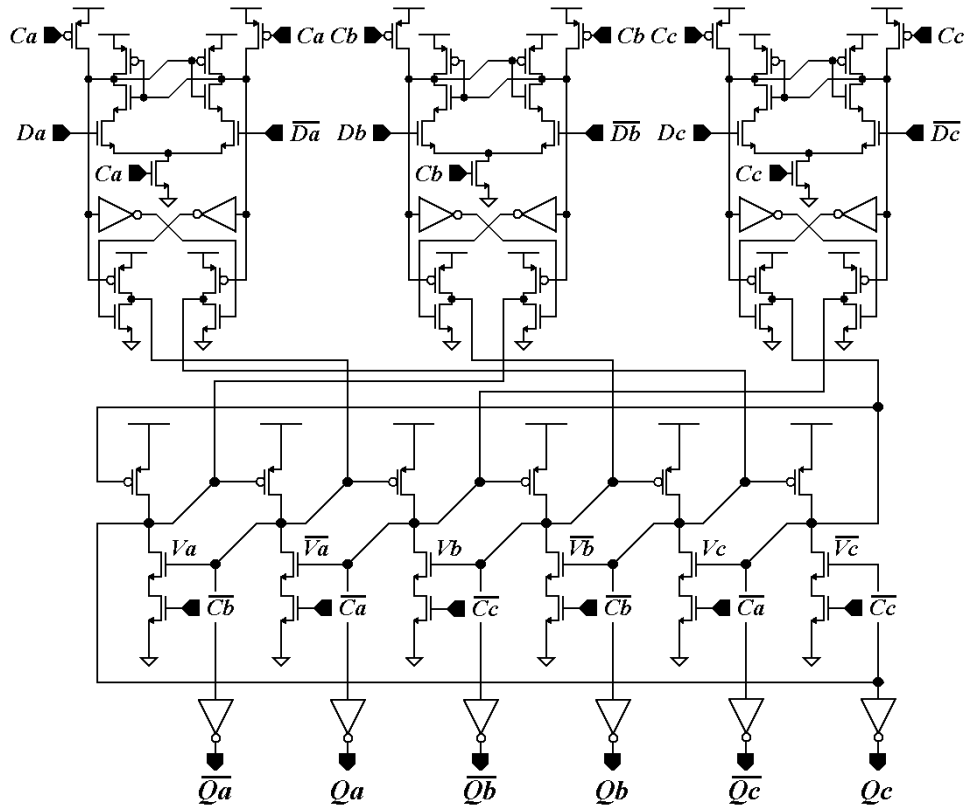


Figure 4.20: Differential TPDICE sense amplifier flip-flop (DTSAFF).

(DTSAFF), using the revised Sense Amplifier-Based Flip-Flop described in [52]. When the clock is low in this design, the sense amplifier stages are disabled, and the outputs of these stages are precharged to logic 1. During this phase, the output latch is in a “hold” state. When the clock transitions to a high value, the sense amplifier stages are enabled. Input values are sampled by the sense amps, which calculate their outputs based off the differences in input voltage pairs. After input values are sampled, the sense amps are disabled. The sampled input values are then sent to the output latch. The output latch saves the results and makes them available as the flip-flop outputs. Changes in the flip-flop inputs during a high clock phase do not affect the state of the flip-flop after data has already been stored. This mitigates timing faults caused by data race conditions.

The TPDICE output latch recovers from transient faults affecting any part of this circuit. If data in any of the sense amplifiers is corrupted, an offset complement pair of nodes will be disrupted in the output latch. The TPDICE circuit was designed to recover from this. However, the recovery process does not complete until the cycle following the sample release phase. Because of this factor, this design is best suited to systems featuring an appreciable percentage of logic stages with short delays.

4.4 Comparison of Fault-Tolerant Pipeline Memories

The fault-tolerant pipeline memory structures described in the previous sections have been evaluated with respect to performance and energy consumption. Each design possesses a distinct balance of attributes that suits it towards a target application. Single-ended and differential designs have been evaluated separately to allow for even comparisons. Evaluations have been performed in a 90nm CMOS process.

All designs analyzed here are fed by minimum-sized balanced data buffers and drive fan-out four (FO4) loads. Minimum-sized balanced buffers are also used as clock drivers in structures requiring clock and complement clock signals. Circuits that do not require complement clock signals are allowed double-sized clock buffers. Additionally, transistor sizes have been selected to optimize performance and energy efficiency.

The motivation behind this circuit configuration is based off a study performed in [53]. This investigation revealed that the average flip-flop in a pipelined data path utilized minimum-sized input data buffers and possessed a fan-out of four. Changes in this configuration alter the relative performance and energy consumption of a set of pipeline memory structures [54]. An analysis of this effect on the designs considered in this paper is planned for the future.

Performance attributes considered in this study include the setup time, hold time, Clk-Q delay, D-Q delay, and total delay. Setup time (also known as D-Clk delay) is defined as the amount of time input data must be stable before the arrival of the sampling clock edge. Hold time (Clk-D delay) is equal to the time between the sampling clock edge and the last instant that input data must be held stable. Clk-Q delay is the time between the clock edge and the instant output data becomes valid. D-Q delay is the time between the first instant that input data must be stable and the instant at which output data is valid. Finally, total delay is the total amount of time a flip-flop consumes from the clock period. It is equal to the greater of the D-Q delay or the setup plus hold time.

In the past, some reports have characterized pipeline memories based on Clk-Q delay. However, this attribute doesn't include all aspects of a circuit's performance, specifically the setup time. To be accurate, flip-flops should be compared through the use of total delay instead [46].

Energy consumption results have been obtained for 0%, 50%, and 100% activity factors. The 0% values represent consumption during cycles with no output transition, while the 100% values represent cycles with a transition. The average case is represented by the 50% values, which are calculated by averaging the 0% and 100% numbers. All measurements are broken down into values for each component of the memory structures. The components include input buffers, internal memory, clock buffers, output buffers, and pulse generation circuitry. Energy dissipated charging the output load capacitance is not included, as this is considered the responsibility of the subsequent stage [53]. However, short-circuit energy consumption in the output buffers is included.

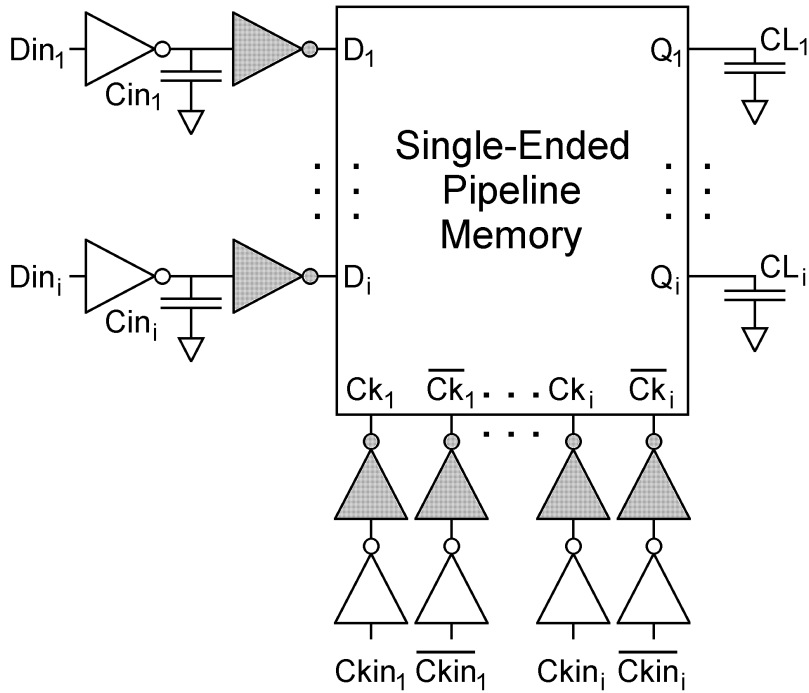


Figure 4.21a: Simulation setup for single-ended pipeline memories.

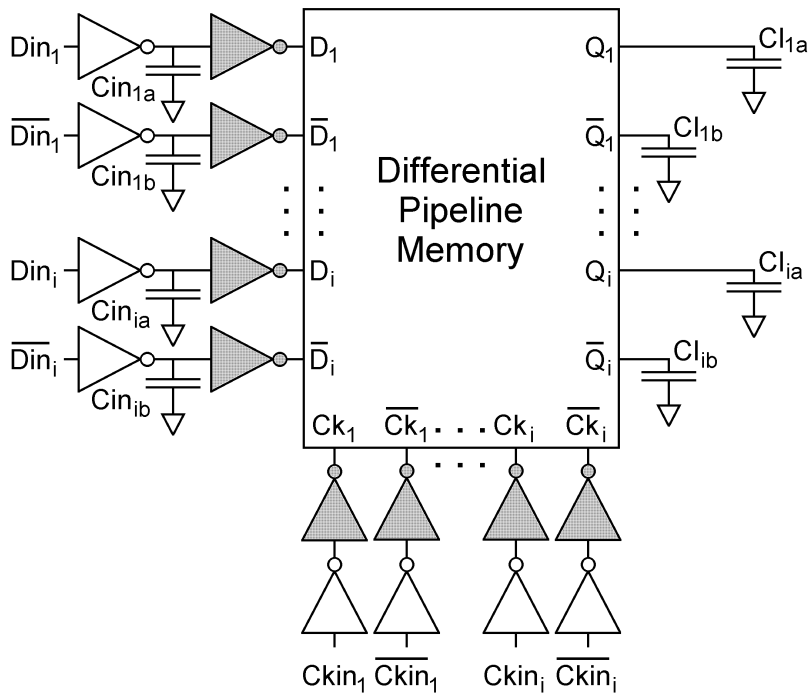


Figure 4.21b: Simulation setup for differential pipeline memories.

Figure 4.21a is a diagram depicting the simulation setup for single-ended architectures. The setup for differential structures is shown in Figure 4.21b. Inverters

TABLE 4.1a: Single-ended pipeline memory acronym definitions.

Acronym	Name
SETL	S-E TPDICE Latch
SETMSFF	S-E TPDICE Master-Slave Flip-Flop
SETAMSFF1	S-E TPDICE Alternative Master-Slave Flip-Flop 1
SETAMSFF2	S-E TPDICE Alternative Master-Slave Flip-Flop 2
SETPFF	S-E TPDICE Pulse-Triggered Flip-Flop
SETDPFF	S-E TPDICE Dual-Pulsed Flip-Flop

TABLE 4.1b: Differential pipeline memory acronym definitions.

Acronym	Name
DBDL	Diff. Barry-Dooley Latch
DBDMSFF	Diff. Barry-Dooley Master-Slave Flip-Flop
DTL	Diff. TPDICE Latch
DTMSFF	Diff. TPDICE Master-Slave Flip-Flop
DTAMSFF1	Diff. TPDICE Alternative Master-Slave Flip-Flop 1
DTAMSFF2	Diff. TPDICE Alternative Master-Slave Flip-Flop 2
DTPFF	Diff. TPDICE Pulse-Triggered Flip-Flop
DTDPFF	Diff. TPDICE Dual-Pulsed Flip-Flop
D TSAFF	Diff. Sense Amplifier Flip-Flop

buffer the clock and data inputs to the pipeline memory. Output buffers drive the capacitance of the subsequent stage. The shaded buffers are included in the delay and power consumption figures. Load capacitances are placed on the input and output lines to produce signals realistic to a FO4 data path.

Pipeline memory acronym definitions are shown in Tables 4.1a and 4.1b to serve as a reminder and increase the clarity of the simulation labels. Performance simulation results for the single-ended designs are displayed in Table 4.2a. Table 4.2b illustrates these results for the differential circuits. Single-ended energy consumption results are depicted in Table 4.3a, while differential values are shown in Table 4.3b.

Graphical depictions of the performance and energy consumption results are shown in Figure 4.22a for single-ended structures. The SETL offers the best balance of

TABLE 4.2a: Single-ended pipeline memory performance figures.

Delay (ps)	Setup	Hold	Clk-Q	D-Q	Total
SETL	119.9	-10.1	-4.4	115.6	115.6
SETMSFF	130.1	9.8	76.7	206.8	206.8
SETMSFFU1	85.3	14.1	55.9	141.2	141.2
SETMSFFU2	119.4	15.5	73.6	193	193
SETPFF	7	107.7	102.5	109.4	114.7
SETDPFF	-65.8	134.6	144.4	78.6	78.6

TABLE 4.2b: Differential pipeline memory performance figures.

Delay (ps)	Setup	Hold	Clk-Q	D-Q	Total
DBDL	97.2	2.3	-0.2	97	99.5
DBDMSFF	98.8	10.7	84.1	182.9	182.9
DTL	104.1	0.8	-8.7	95.4	104.9
DTMSFF	113.6	-3.7	105.6	219.2	219.2
DTAMSFF1	73.8	16.2	62.9	136.7	136.7
DTAMSFF2	93.5	11.3	82.9	176.4	176.4
DTPFF	-39.1	144.9	134.3	95.2	105.8
DTDPFF	-85.9	155.3	151.5	65.6	69.4
DTSAFF	7.3	42.7	107.5	114.8	114.8

performance and power consumption, although this comes at the cost of a large sampling window. The SETPFF offers almost identical performance with a significantly reduced transparency window, although energy consumption suffers. If performance is a lower priority, master-slave flip-flops can be adopted to avoid pulse generation circuitry and use less energy. The SETAMSFF1 is the fastest of this group, although it cannot recover from disruptions until after the initial sample release clock cycle. Faster recovery is provided by the SETAMSFF2, although with greater latency. The SETMSFF is the slowest of the master-slave structures, which is the price paid for the ability to recover from disruptions in all phases of the clock. Lastly, optimum performance with a small sampling window is obtained by the SETDPFF. The downside of this approach is that it

TABLE 4.3a: Single-ended pipeline memory energy consumption figures.

Energy (fJ)		Int	In	Clk	Out	Pulse	Total
SETL	100%	7.307	5.234	7.004	1.309	0	20.85
	50%	3.662	2.626	7.038	0.656	0	13.98
	0%	0.017	0.019	7.071	0.003	0	7.111
SETMSFF	100%	16.25	5.696	13.17	1.283	0	36.39
	50%	8.194	2.853	13.19	0.643	0	24.88
	0%	0.141	0.01	13.22	0.003	0	13.37
SETAMSFF1	100%	13.25	4.601	10.05	1.328	0	29.23
	50%	6.664	2.338	10.09	0.665	0	19.76
	0%	0.08	0.074	10.14	0.003	0	10.29
SETAMSFF2	100%	12.77	5.861	11.36	1.332	0	31.32
	50%	6.44	2.934	11.34	0.667	0	21.38
	0%	0.109	0.008	11.33	0.002	0	11.44
SETPFF	100%	7.258	4.898	7.576	1.317	25.6	46.64
	50%	3.646	2.458	7.575	0.66	25.63	39.97
	0%	0.034	0.018	7.575	0.003	25.66	33.29
SETDPFF	100%	12.93	2.694	5.546	1.319	25.82	48.3
	50%	12	1.348	5.545	1.128	25.82	45.84
	0%	11.07	0.003	5.545	0.937	25.82	43.37

consumes the most power and experiences output glitches during pulse overlap due to its semidynamic nature.

Figure 4.22b graphically presents the performance and energy consumption figures for the differential designs. The DBDL and DBDMSFF are faster with less energy consumption than their TPDICE counterparts. However, they require independent differential data paths, which can easily counterbalance these advantages. The relationships between the attributes possessed by the TPDICE differential designs are largely similar to that of their single-ended counterparts. The DTL offers the best combination of performance and energy consumption, although it does not offer much protection against data race conditions. Moving to structures that limit the transparency window, the DTPFF can offer the performance of the latch at the cost of higher energy

TABLE 4.3b: Differential pipeline memory energy consumption figures.

Energy (fJ)		Int	In	Clk	Out	Pulse	Total
DBDL	100%	2.194	7.107	6.836	0.186	0	16.32
	50%	1.097	3.612	6.921	0.095	0	11.73
	0%	0	0.117	7.006	0.004	0	7.127
DBDMSFF	100%	10.27	7.111	10.66	0.186	0	27.85
	50%	5.157	3.6	10.71	0.093	0	19.37
	0%	0.05	0.088	10.76	0	0	10.89
DTL	100%	2.587	8.858	6.999	2.841	0	21.29
	50%	1.303	4.453	7.008	1.424	0	14.19
	0%	0.018	0.047	7.017	0.006	0	7.088
DTMSFF	100%	10.34	7.355	11.13	1.595	0	30.42
	50%	5.202	3.702	11.09	0.8	0	20.79
	0%	0.063	0.049	11.05	0.005	0	11.17
DTAMSFF1	100%	18.59	5.425	10.84	2.645	0	37.5
	50%	9.309	2.715	10.89	1.326	0	24.23
	0%	0.028	0.005	10.93	0.006	0	10.97
DTAMSFF2	100%	15.14	8.061	9.257	2.565	0	35.02
	50%	7.57	4.051	9.37	1.285	0	22.28
	0%	0	0.041	9.483	0.005	0	9.529
DTPFF	100%	2.528	8.511	7.725	2.849	26.94	48.55
	50%	1.267	4.262	7.724	1.428	26.96	41.63
	0%	0.005	0.012	7.722	0.006	26.97	34.72
DTDPFF	100%	16.53	2.197	5.557	2.614	27.07	53.97
	50%	15.84	1.101	5.559	2.606	27.08	52.18
	0%	15.15	0.005	5.56	2.597	27.08	50.39
DTSAFF	100%	21.05	2.66	8.099	2.576	0	34.39
	50%	16.72	1.421	8.061	1.291	0	27.49
	0%	12.39	0.182	8.023	0.005	0	20.6

consumption. The Master-slave latches consume less power, but also offer less performance. Relatively speaking, the DTAMSFF1 and DTAMSFF2 are faster than their single-ended counterparts, and the DTMSFF is slower. The DTPFF is the fastest design, although it consumes even more energy than the single ended version. This is due to the fact that it induces output glitches during every write operation, and not just the ones that activate its data evaluation transistors. And finally, unique to a differential configuration,

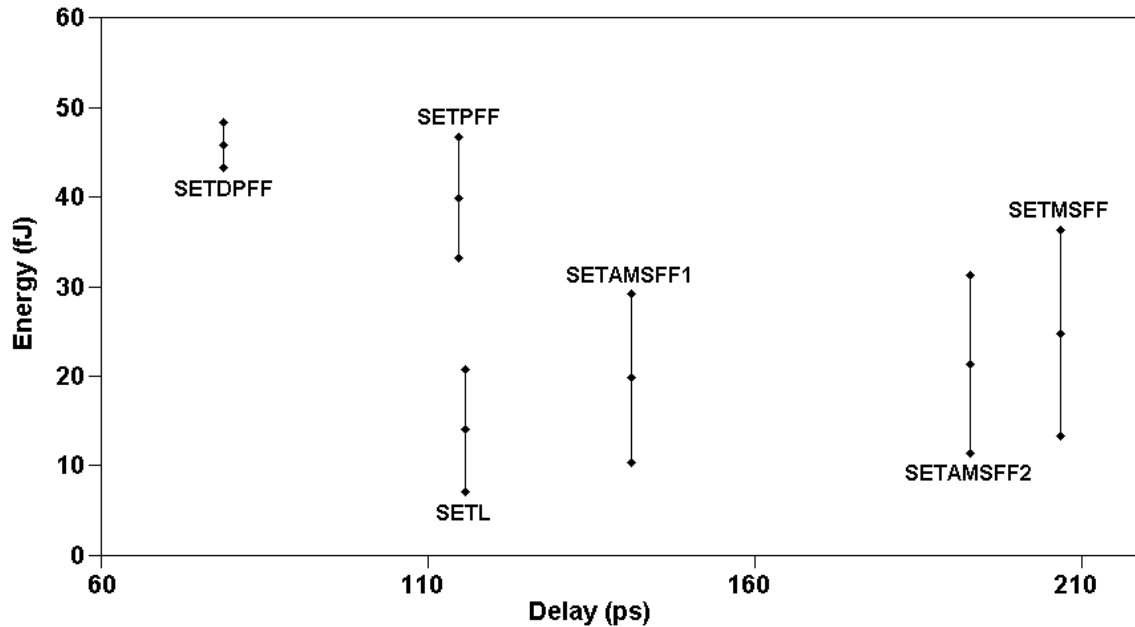


Figure 4.22a: Graph of delay vs. energy values for single-ended structures.

DTSAFF outperforms all master-slave flip-flops and consumes less energy than the pulsed structures. Unfortunately, it suffers from an inability to correct faults during the initial sample-release cycle, like the DTAMSFF1.

Designers selecting memory for a fault-tolerant pipelined system should first consider whether they would like a single-ended or differential data path. Single-ended designs require simpler logic, fewer data lines, and fewer output data buffers. Differential structures can take advantage of cross-coupled and sense amplifier logic (save for those using Barry-Dooley memory), as well as balanced data transition times, which can result in higher performance.

Latches are best in situations with tightly constrained logic delays, while flip-flops are a better choice when data race conditions are a possibility. Master-slave flip-flops are a good choice in designs that do not require top-level performance, but can benefit from a lack of pulse generation circuitry and moderate power consumption. A specific master-

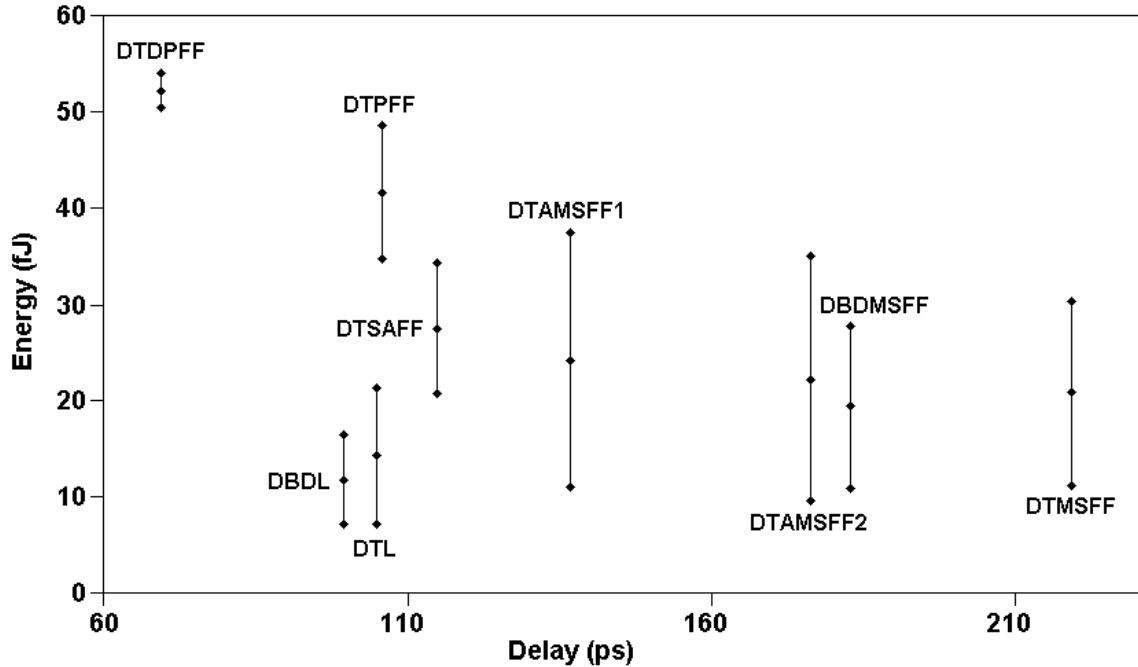


Figure 4.22b: Graph of delay vs. energy values for differential structures.

slave flip-flop should be selected based on performance and fault recovery response time considerations. Pulse-triggered flip-flops are ideal when performance is the most important consideration, and energy consumption is less important. The TPDICE pulse-triggered flip-flops offer performance at latch levels, while the semidynamic TPDICE dual-pulsed flip-flops can provide even greater performance if the output glitches and higher energy consumption can be tolerated.

The flip-flops specific to differential systems offer additional options to the system designer. Barry-Dooley approaches are beneficial when independent data paths are not needed. Additionally, the TPDICE sense amp flip-flop possesses better performance and energy consumption characteristics than master-slave designs. It is a good choice in systems that do not require immediate recovery from transient disruptions.

4.5 Summary

SEU and SET-tolerant pipeline memory structures were the focus of this chapter. Latches, master-slave flip-flops, and pulse-triggered flip-flops were presented, arranged in single-ended and differential configurations. TPDICE and Barry/Dooley memory cells were used in all pipeline memories to achieve SEU and SET-tolerance and bypass capability. All approaches were evaluated with respect to delay, energy consumption, and transparency window size. The transparency window is defined as the amount of time that input data may pass through to the outputs. This window could allow data to erroneously pass through multiple pipeline stages, resulting in a timing fault. Thus, it is desirable to minimize the transparency window.

Latches provide the best balance of performance and energy consumption, at the cost of a transparency window that is the width of a clock phase. Master-slave flip-flops eliminate the transparency window, but suffer in the performance and energy consumption categories. Pulse-triggered flip-flops offer a significantly reduced transparency window with latch-level performance, although they consume the most energy.

Chapter 5

MBU-Tolerant Approaches

This section describes a novel Multiple-Bit Upset (MBU)-tolerant design, which utilizes layout-based interleaving and multiple-node disruption tolerant memory latches. This approach protects against grazing incidence particle strikes, which produce disruptions with the widest possible spatial separation. Advantages with respect to size, complexity, and MBU tolerance are realized when this approach is compared to existing solutions.

5.1 Background

Transient errors may occur in ICs when radioactive particles impact them. Individual upsets that directly affect memory are known as Single-Event Upsets (SEUs), while upsets that originate in logic are known as Single-Event Transients (SETs). Many approaches have been designed to deal with SEUs and SETs [41]. However, most of these schemes cannot withstand multiple-node and Multiple-Bit Upsets (MBUs). MBUs are a growing concern in environments that possess substantial high-energy ion activity (e.g. in space) [17]. Because of this, additional techniques must be devised to protect against multiple disruptions.

Error Correcting Codes (ECCs), chip-level interleaving, and temporal redundancy are existing schemes that provide MBU tolerance. Multiple-bit correcting ECCs require significant overhead, and they may fail when faced with a large number of upsets. Chip-

level interleaving is effective, although the requirement of multiple chips per circuit is expensive. Finally, temporal redundancy does not possess robust MBU tolerance unless full system-level redundancy is implemented. This paper addresses this issue through the presentation of a layout interleaving scheme and multiple-node disruption tolerant (MNDT) memory latches. By combining these two techniques, protection from grazing and non-grazing MBU-inducing particle strikes can be achieved.

MBU probability is strongly dependent on node spacing, feature size, and supply voltage. As feature sizes shrink, MBUs are becoming more of an issue. The linear Energy Transfer (LET), range, track radius, and angle of incidence of the particle inducing upset are also important. In general, particles that deposit more energy, have a longer range, and have a larger radius are more likely to induce MBU. Additionally, grazing strikes have the capability to cause much greater problems than normal incidence strikes [17, 55].

5.2 Ion Trajectories and Layout Interleaving

Particle strikes that produce MBU fall in two general categories: Strikes virtually parallel to the IC surface that pass through sensitive regions of multiple nodes, and strikes that travel at an appreciable angle to the surface. Parallel strikes leave behind a linear trail of charge, and disrupt nodes along a line tracing the ion path. In general, these grazing strikes can affect the most nodes and nodes separated by the greatest distances. On the other hand, non-grazing particles may disrupt multiple nodes within a certain radius via diffusion, charge sharing or secondary particle scattering. If multiple nodes are disrupted by such a particle strike, the disruption pattern tends to form a cluster stretched along the




Figure 5.1: MBU upset patterns resulting from grazing incidence and non-grazing incidence particle strikes. The arrows represent grazing strikes, while the dots depict non-grazing strikes. The gray boxes symbolize typical disruption patterns initiated by these particle strikes.

angle of incidence [55-57]. Figure 5.1 depicts some typical upset patterns of grazing and non-grazing strikes [58-61].

It is possible to improve the MBU resistance of SEU/SET-tolerant latches, such as the DICE design in [37], by increasing the layout spacing between nodes in every cell. While it is impractical to use this strategy alone to prevent upset from grazing strikes, other strikes can be reasonably mitigated in this fashion. Grazing strikes have been shown to affect nodes separated by $300\mu\text{m}$ in [62], while non-grazing strikes typically affect a radius on the order of $10\mu\text{m}$ or less [17]. The layout of multiple cells could be interleaved in such a way as to increase the spacing between nodes in the same cell without substantially increasing the overall layout area. For example, consider interleaving the

layouts of four-node DICE latches. If four such latches are interleaved, each pair of nodes within the same latch can be separated by one node from a different latch. Spacing between nodes in the same latch is increased by a factor of n if n^2 memory latches are interleaved. Note that this specific example is vulnerable to grazing strikes. The following sections address this issue.

5.3 Multiple-Node Disruption Tolerant Latches

Most current SEU and SET-hardened memory structures can recover from only one disruption at a time. These hardened cells may fail in environments where MBUs occur. It is possible to decrease this probability of failure by adopting cells that tolerate multiple-node disruptions. Additionally, cells that can tolerate multiple disruptions have the potential to avoid upset due to grazing particle strikes that deposit a line of charge. If three or more nodes must be disrupted to flip a cell, it is possible to arrange the layout so that no line passes through a set of mutually vulnerable nodes.

At a minimum, five nodes are required to preserve sufficient information to tolerate dual faults in a latch. Six nodes (three true and three complement) are required to achieve this goal with balanced feedback. Unfortunately, the complexity required of such six node latches is very high. For example, one six node approach, illustrated in Figure 5.2, requires 84 transistors. Fortunately, this complexity can be significantly reduced by adopting structures with eight nodes. Adding additional nodes increases the amount of information stored by the cell, reducing the relative impact of a multiple-node disruption. This allows fewer transistors to be used to control the nodes in the latch.

Figure 5.3 illustrates an eight node MNDT latch [63]. The architecture is similar to that of a pair of connected Barry/Dooley cells [44-45]. The interconnections are arranged

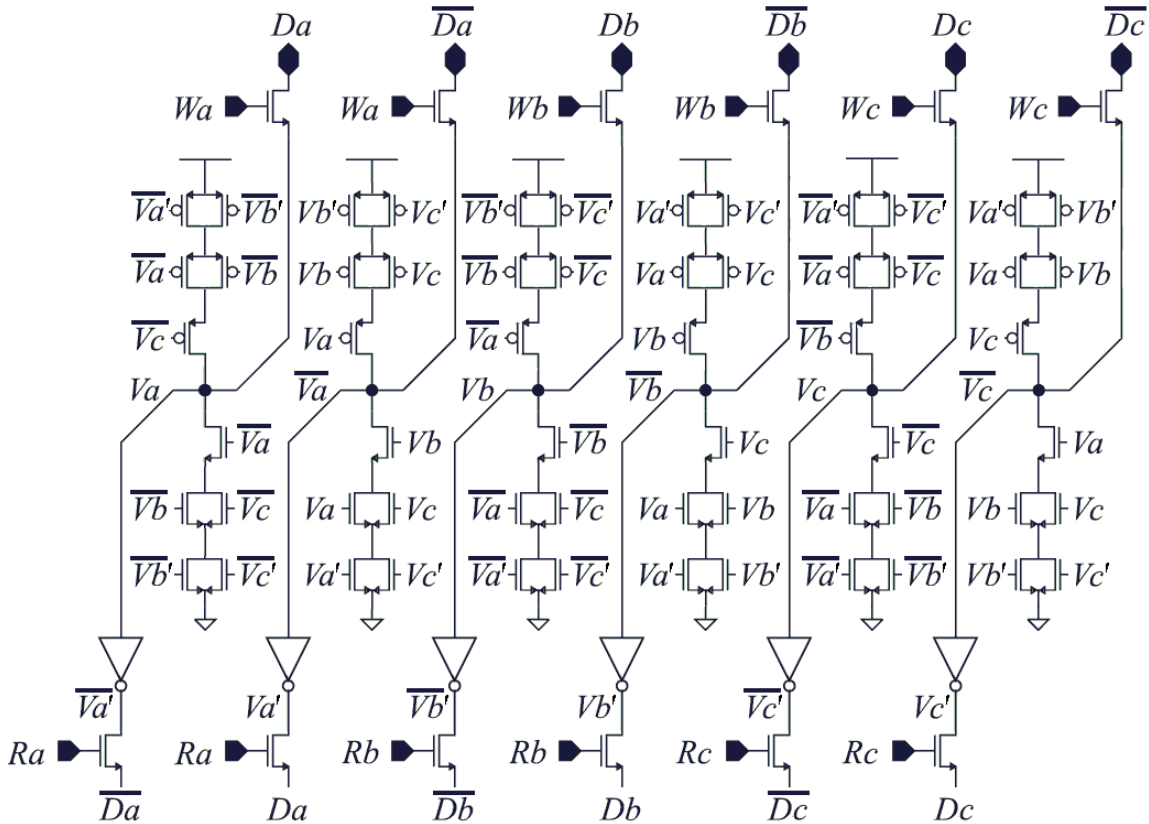


Figure 5.2: Six node latch possessing the ability to tolerate two simultaneous upsets. This design suffers from high complexity, as it requires 84 transistors.

in such a manner as to preserve sufficient information when faced with dual-node disruptions and allow for recovery after disruptions dissipate. Let sets $V = \{V_a, V_b, V_c, V_d\}$ and $\bar{V} = \{\bar{V}_a, \bar{V}_b, \bar{V}_c, \bar{V}_d\}$. If the two disrupted nodes are in different sets, then none of the other six nodes will be affected. During recovery, at least one of the upset nodes will be immediately pulled to the proper state, followed by the second node. Alternatively, if the two affected nodes are in the same set, then they could possibly flip the state of a node in the complement set. Fortunately, this does not prevent at least one of the disrupted nodes from recovering immediately after the effects of the particle strike dissipate. After this, the other nodes quickly follow suit.

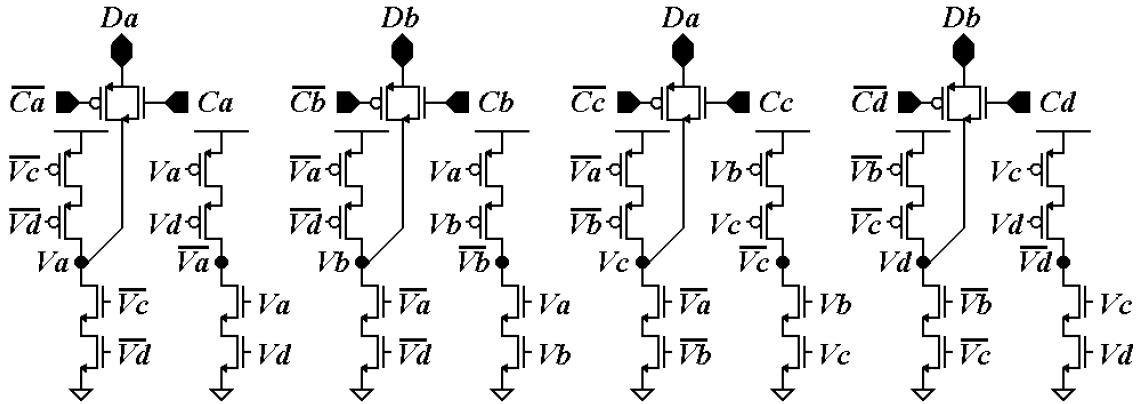


Figure 5.3: Eight node latch with dual-node and limited triple-node disruption tolerance.

One additional advantage of this circuit is a limited ability to tolerate three-node disruptions. A number of the possible disruptions affecting two nodes from one set and one node from the complement set can be withstood. This occurs when the two disrupted nodes from the same set do not control a node in the complement set that is distinct from the third disrupted node. This characteristic simplifies the task of creating an MBU-tolerant layout for this circuit.

As an example of this limited ability to tolerate three-node disruptions, assume the initial latch state is 01010101. Nodes Va and Vb are flipped from 0 to 1, and node /Vb is flipped from 1 to 0. The resulting latch state is 11100101. None of the other five nodes are affected, as they are not controlled by some combination of Va, Vb, and /Vb. After the transient disruption dissipates, Va and Vb are pulled back down to 0 by their controlling nodes (/Vc, /Vd and /Va, /Vd). Finally, /Vb is pulled back up to 1 by Va, which restores the proper state of the latch.

5.4 Layout Interleaving of MNDT Latches

In isolation, the usefulness of the multiple-node disruption tolerant latches described in the previous section is limited. The ability to tolerate multiple disruptions is mostly offset by the increase in area. However, if these latches are implemented with the layout interleaving technique described in Section II, a desirable characteristic can be achieved. The most hazardous MBU mechanism can be avoided by arranging the layout so no line passes through a set of nodes that, if disrupted, would cause a latch to flip state. Grazing particle strikes deposit charge along a linear path, and they can upset nodes spaced by large distances. Bits spaced by $300\mu\text{m}$ were shown to have an appreciable probability of mutual upset in [62]. However, if no line contains a set of mutually vulnerable nodes, then the probability of upset resulting from this mechanism can be mitigated.

The eight node circuit presented in the previous section can tolerate a number of three-node disruptions, including the following sets: $\{Va, /Va, Vd\}$, $\{Va, /Vb, Vb\}$, $\{Vb, /Vc, Vc\}$, $\{Vc, /Vd, Vd\}$. Using this information, the layout depicted in Figure 5.4 can be constructed with the purpose of avoiding upset due to multiple-node disruptions. Nine cells are interleaved, and so the spacing between nodes is increased by a factor of three. No line intersects a set of nodes that would, if disrupted, flip the state of the latch. This characteristic provides protection against grazing particle strikes, and it is illustrated in this figure. Additionally, since the spacing between each node is increased, tolerance to non-grazing particle strikes is provided. The level of interleaving can be adjusted to meet the demands of a particular situation. In particular, a worst-case grazing strike is represented by the dotted line in Figure 5.4 [63]. This strike has the potential to disrupt




Figure 5.4: Diagram of interleaved layout with selected grazing particle strikes. The solid lines represent strikes that pass through three nodes, while the dotted line represents a worst-case strike. All strikes can be tolerated with an appropriate level of interleaving.

four nodes via charge sharing [62]. The level of interleaving must be sufficient to avoid this possibility.

An implementation of the interleaved layout is shown in Figure 5.5. This design was performed in $0.18\mu\text{m}$ CMOS, and it utilizes four metal layers for interconnections. The layout requires an area of $36.4\mu\text{m} \times 49.2\mu\text{m}$. Fewer metal layers could be used at the cost of additional area. The minimum spacing between nodes in the same latch is $9.2\mu\text{m}$. Therefore, a worst-case particle strike must have an ion track radius of at least $5.6\mu\text{m}$ to upset four nodes in the same cell along the dotted trajectory shown in Figure 5.4. The vast majority of high-energy cosmic rays found in space have radii that are significantly less than this critical value [64].




Figure 5.5: Interleaved layout of nine hardened memory latches.

Figure 5.6 is a layout simulation illustrating recovery from three simultaneous disruptions affecting nodes Va , Vb and Vb in the same memory cell. Figure 5.7 depicts the recovery from simultaneous SETs affecting enable lines Ca and Cb .

5.5 Comparison to Existing Strategies

Existing MBU-tolerant schemes include ECCs, chip-level interleaving, and temporal redundancy. The ECCs can be broken down into schemes that operate on single data words and schemes that operate on blocks of data words. Hamming code is a single bit correcting code that functions by appending $\log_2(n+1)$ parity bits to each data word, where n is the total number of bits in the word. Multiple upsets can be corrected by adopting codes that use a greater number of parity bits. Block level ECCs, such as Reed-Solomon code, perform calculations on groups of data words to achieve greater




Figure 5.6: Simulation of recovery from three simultaneous disruptions affecting Va, Vb and /Vb.




Figure 5.7: Simulation of recovery from simultaneous SETs affecting Ca and Cb.

efficiency. Chip-level interleaving is implemented by distributing a single circuit across multiple chips. Finally, temporally redundant designs sample data at multiple instances in an attempt to detect and bypass faults.

Table 5.1 presents a comparison between various MBU-tolerant approaches. These include the MNDT latches with interleaving, a two-bit correcting ECC, Reed-Solomon block code, chip level interleaving, and temporal redundancy. The table columns compare the upset correction capabilities, redundancy, grazing strike protection, requirement of multiple chips, and latency. An analysis of this data follows:

Table 5.1: Comparison of MBU-tolerant approaches.

	Upset Correction	Redundancy (Cycle Time or Layout Area)	Grazing Protection	Multiple Chips	Latency
Proposed	3 per cell	4x (Area)	Yes	No	None
Two-Bit ECC	2 per word	$\log((n^2+n+2)/2)$ (Area)	No	No	Enc/Dec
Reed-Solomon	$(n-k) / 2$ symbols	$(n-k)$ symbols (Area)	No	No	Enc/Dec
Chip Interleaving	Any	$\log(n+1)$ (Area)	Yes	Yes	Enc/Dec
Temporal Redundancy	Any	3x (Time)	Yes	No	3x + Voting

The proposed approach is well balanced, as it protects against both grazing and non-grazing strikes, requires moderate redundancy, fits on one chip, and does not need latency inducing encoder, decoder, or voting circuitry.

- Multiple-bit correcting ECCs that operate on single words offer no protection against grazing strikes, require significant redundancy for moderate sized data words, and require encoding and decoding circuitry.

- Block level codes such as the Reed-Solomon approach are powerful, as they have correcting power equal to half the number of redundant symbols. However, they are only useful in large memories (such as those with 255 byte blocks), and they require encoders and decoders.

- Chip-level interleaving (placing each bit on a different chip) also provides robust MBU tolerance, although the use of multiple chips is not practical in many cases.

- Temporal redundancy is fully MBU-tolerant only when entire calculations are performed redundantly. In hybrid temporal-spatial redundant schemes, the spatially redundant circuitry is vulnerable to multiple upsets. With full temporal redundancy, performance is cut to one third of the original value, and is reduced further by voting.

5.6 Summary

MBUs are a growing concern in aggressive technologies that must operate reliably in harsh radioactive environments. Non-grazing particle strikes can upset a cluster of nodes within a moderate radius of the particle strike, whereas grazing strikes can disrupt a very long line (up to 300 μm) of nodes along the ion path. Grazing strikes are particularly dangerous due to the fact that it is very impractical to increase the spacing of mutually vulnerable nodes beyond the range of a grazing particle strike in a silicon IC. This chapter presented an MBU-tolerant approach that can withstand grazing and non-grazing particle strikes. It relies on a MNDT memory cell and layout interleaving to assure that a single grazing strike cannot pass through a set of memory nodes that, if disrupted, would cause a cell to flip state. This approach requires four times the memory complexity of a standard SRAM latch, as well as added interconnect complexity due to the layout interleaving. This level of complexity is less than that of other MBU-tolerant approaches that can tolerate grazing strikes, such as chip-level interleaving and full temporal redundancy.

Chapter 6

Conclusion

This dissertation has focused on hardened by design architectures for mitigating transient radiation-induced disruptions in digital integrated circuits. The disruptions considered here include SEUs, SETs, and MBUs. SEU/SET/MBU-tolerant SRAM designs form the foundation to our approach to diminishing the effects of radioactive particle strikes. An SEU and SET-tolerant reconfigurable DSP processor has been constructed to illustrate the capabilities of these memory cells. Radiation-hardened pipeline latches and flip-flops utilizing this fault-tolerant memory were also presented. Finally, a strategy for mitigating MBUs due to both grazing and non-grazing particle strikes was detailed.

From our evaluations we have observed that the ability to bypass SEU and SET transient pulses drastically improves performance. Some studies have measured SET width to be up to 2ns, a figure that is largely unaffected by reductions in feature size [12-15]. This in turn 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].

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, which can be accomplished by triplication of a basic voting circuit. In addition to this, feedback must be utilized to correct a corrupted source memory cell in TMR-based SRAM arrays.

Design-hardened circuit-level approaches were presented in this report with the goal of improving on the efficiency of TMR while retaining the capability to bypass transient pulses. These approaches include the fully-differential DICE and TPDICE memory latches. The fully-differential DICE approach utilizes a standard DICE core and dual-redundant complement data paths. It can only be used in pure data transfer systems that do not rely on combinational logic. In contrast, the TPDICE latch can be used with logic. It relies on six internal nodes, making it analogous to a DICE cell extended from two interlocked SRAM latches to three, or area-redundant TMR without external voting. This lack of voting allows the TPDICE approach to achieve superior performance, energy consumption, and circuit complexity when compared to TMR.

To achieve SET-tolerance, the use of one independent data path requires signal delays, which directly affect performance. Using two data paths allows transient pulse detection, but the system must pause until pulses dissipate. With three paths, majority voting (or the equivalent) allows the correct logic value to be selected and passed on before pulses dissipate. This pulse bypass capability removes clock period dependence on maximum SET width, drastically improving performance.

The proposed SEU and SET-tolerant reconfigurable DSP processor offers a unique combination of flexibility, high-performance and radiation tolerance for space

applications. The medium-grain architecture featuring 4-bit LUT-based cells is particularly suited towards DSP algorithms. High-performance and radiation-tolerance are provided by the TPDICE memory cell, which possesses SEU and SET bypass capability. Simulations have demonstrated the performance and SEU/SET mitigation advantages of this design.

This research addressed the concept of reliable pipelined datapaths through the examination of SEU and SET-tolerant pipeline memory structures. Single-ended and differential structures were characterized with respect to performance and energy consumption. All of the presented approaches possess SEU and SET bypass capability, which allows them to proceed with subsequent operations even when one of their nodes is affected by an upset. This characteristic is very important, as it removes the need to add clock cycle overhead proportional to the maximum upset width. Fault-tolerance and bypass capabilities are achieved by utilizing the TPDICE latch as the core memory in the majority of the proposed structures.

Pipeline memory circuits fall into three main categories that are considered by this report: Level-sensitive latches, edge-triggered master-slave flip-flops, and pulse-triggered flip-flops. Latches provide the best balance between performance and power consumption, but they possess large transparency windows. Master-slave flip-flops consume moderate power and have essentially no transparency window, although they offer the least performance. Pulse-triggered flip-flops possess high performance and reduced transparency windows, but at the cost of greater energy consumption.

Design-hardened techniques have been developed and presented herein with the purpose of tolerating MBUs and multiple-node upsets. Specifically, layout interleaving

was combined with MNDT memory latches. The increased spacing between nodes in a single latch and tolerance to multi-node disruptions provides substantial protection against non-grazing particle strikes that produce MBUs. More importantly, this scheme assures that grazing particles, which may disrupt distant nodes along their linear paths, cannot disrupt a set of nodes that would cause a latch to flip state. Few MBU-resistant schemes produced to date offer sufficient protection against grazing particle strikes. Those that do require substantial spacing between adjacent bits (even placing bits on separate chips), temporal redundancy, and/or the use of ECCs (expensive multiple-bit correcting codes are required in some schemes) [17, 56, 58, 60]. Due to this overhead, these schemes do not provide the high performance and balanced secondary characteristics of the proposed approach.

Section 6.1 of this chapter highlights the major contributions of this dissertation. Section 6.2 describes potential directions for future work.

6.1 Contributions

The research presented in this document has focused on augmenting the body of work produced by the radiation effects mitigation community. The novel designs and analysis presented in this paper provide the following contributions:

- ***Efficient non-bypass capable SET-tolerant SRAM latches:*** Structures without SET bypass capability are the best choices in situations where the maximum SET width is small or low complexity is more desirable than maximum performance. This research introduced and evaluated a number of such approaches, including basic SET-tolerant DICE, enhanced DICE 1, enhanced DICE 2, and delay-based DICE designs. Structures

relying on dual-rail logic or delay-filtered inputs possess relatively low complexity, but they must delay signals or pause write operations while a transient pulse exists on one of their inputs. This directly affects performance. Other SET-tolerant approaches considered in this report bypass transient pulses, so they do not need to pause when a pulse is detected.

- ***Fully-differential DICE structures with transient pulse bypass capability in SRAM systems:*** The novel fully-differential DICE approach utilizes the basic four-node DICE core. It relies on two independent differential data paths and four enable lines. This approach possesses the ability to bypass transient pulses, however this only applies in pure data transfer systems without combinational logic.

- ***TPDICE bypass-capable memory cell:*** The novel TPDICE structure is an original design that offers the ability to bypass transient faults while maintaining balance and efficiency. It 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 and energy consumption are 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.

- ***Techniques for mitigating read-induced upsets in design-hardened memory cells:*** Charge sharing affects many interlocked design-hardened memory cells during read operations, increasing susceptibility to SEU. This issue does not necessarily affect

pipeline latches, although it can potentially affect cells in memory arrays using a shared data bus. Tri-stated memory nodes preserving uncorrupted logic values have no drive strength to dissipate bus charge, making them sensitive to voltage drifting. This research alters the sizing of latch transistors, manages bus capacitance, and precharges the bus to $V_{dd}/2$ to diminish the probability of upset from this factor.

- ***Radiation-Tolerant Reconfigurable DSP Processor Architecture:*** SEU and SET-tolerance has been incorporated into the construction of reconfigurable DSP circuitry. Specifically, the design consists of two LUT-based cells connected by a reconfigurable switch. TPDICE latches possessing the capability to bypass SEUs and SETs provide radiation protection. Faults originating at any node in this architecture can be tolerated, including disruptions affecting LUT memory, switch memory, enable logic, and switching logic. This design features a number of unique attributes, including its memory-based reconfigurable architecture, design-hardened SRAM LUTs, and techniques utilized to ensure fault-tolerance exists in all aspects of its circuitry.

- ***Multiple pipeline memory options for use in fault-tolerant pipelined systems:*** A number of innovative SEU and SET-tolerant pipelined memory latches and flip-flops were described, offering a number of options to designers of radiation-tolerant datapaths. Single-ended and differential designs present a tradeoff between simplified logic and balanced data transitions. Basic latches the best choice when efficiency is desired and logic delays are tightly bounded, while flip-flops provide protection against potential timing faults. Master-slave designs are a good choice when modest performance and energy consumption is desired, while pulse-triggered designs provide high performance

at the cost of substantial energy consumption. Few studies have looked specifically at improving the radiation tolerance of high-performance datapaths. This research strives to serve as an example for future work in this area.

- ***MBU-tolerant SRAM-based approach utilizing layout interleaving combined with multiple node disruption tolerant memory:*** Design-hardened techniques have been developed and presented in this dissertation with the purpose of tolerating MBUs and multiple node upsets. Specifically, layout interleaving was combined with a novel MNMT memory latch. The increased spacing between nodes in a single latch and tolerance to multi-node disruptions provides substantial protection against non-grazing particle strikes that produce MBUs via diffusion or secondary particle scattering. More importantly, this scheme assures that grazing particles, which may affect distant nodes along their linear paths, cannot disrupt a set of nodes that would cause a latch to flip state.

6.2 Future Work

Future plans pertaining to this avenue of research include continued exploration of SEU/SET/MBU-tolerant approaches, fabrication of an IC containing fault-tolerant reconfigurable DSP circuitry, and construction of a high-performance computational datapath featuring SEU/SET-tolerant pipeline memories. The SEU/SET/MBU-tolerant approaches found in this report form a solid core, and work will be performed to expand on this foundation. Fabrication of the fault-tolerant DSP processor circuitry will enable verification of the design and characterization of its radiation tolerance threshold. Finally, construction of an SEU and SET-tolerant computational datapath will serve as an example for a class of circuits that has not seen significant radiation effects analysis. The

fault-tolerant pipeline latches and flip-flops proposed in this report will be utilized in this datapath design.

A number of factors should be taken into account when selecting an SEU/SET/MBU-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 Transactions on Nuclear Science*, vol. 29, no. 6, pp. 2055-2063, Dec. 1982.
- [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, 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 Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, 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," *Proceedings of the 2005 International Conference on Computer Design*, Las Vegas, NV, Jun. 2005.
- [6] D.R. Blum and J.G. Delgado-Frias, "Comparison of SET-resistant approaches for memory-based architectures," *Proceedings of the 12th NASA Symposium on VLSI Design*, Coeur d'Alene, ID, Oct. 2005.
- [7] D.G. Mavis and P.H. Eaton, "Soft error rate mitigation techniques for modern microcircuits," *Proceedings of the 40th Annual International Reliability Physics Symposium*, pp. 216-225, Dallas, TX, Apr. 2002.
- [8] P. Mongkolkachit and 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.
- [9] K.J. Hass, J.W. Gambles, B. Walker, and M. Zampaglione, "Mitigating single event upsets from combinational logic," *Proceedings of the 7th NASA Symposium on VLSI design*, Oct. 1998.
- [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 No. 6,326,809, Dec. 2001.

- [11] D.G. Mavis, "Single event transient phenomena -- challenges and solutions," *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 Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, vol. 48, no. 6, pp. 1931-1936, Dec. 2001.
- [14] L.M. Cohn, "Single-event effects in advanced digital and analog microelectronics," *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 Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, vol. 41, no. 6, pp. 2005-2017, Dec. 1994.
- [18] P.E. Dodd and F.W. Sexton, "Critical charge concepts for CMOS SRAMs," *IEEE Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, 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 Transactions on Nuclear Science*, vol. 43, no. 6, pp. 2797-2804, Dec. 1996.
- [21] K.J. Hass and J.W. Gambles, "Single event transients in deep submicron CMOS," *Proceedings of the 42nd Midwest Symposium On Circuits and Systems*, Las Cruces, NM, vol. 1, pp. 122-125, Aug. 1999.
- [22] M.P. Baze and S.P. Buchner, "Attenuation of single event induced pulses in CMOS combinational logic," *IEEE Transactions on Nuclear Science*, 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," *Proceedings of the 2002 International Conference on Dependable Systems and Networks*, Bethesda, Maryland, pp. 389–398, Jun. 2002.
- [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 Transactions on Nuclear Science*, vol. 51, no. 6, pp. 3480–3485, Dec. 2004.
- [25] K.E. Holbert. (2006, Jan. 18). *Single event effects* [Online]. Available: <http://www.eas.asu.edu/~holbert/eee460/see.html>.
- [26] Robinson, W. Lee, R. Aguero, S. Gabriel, "Anomalies due to single event upsets," *Journal of Spacecraft and Rockets*, vol. 31, no. 2, pp. 166-171, Apr. 1994.
- [27] E.L. Petersen, V. Pouget, L.W. Massengill, S.P. Buchner, and D. McMorrow, "Rate predictions for single-event effects—critique II," *IEEE Transactions on Nuclear Science*, vol. 52, no. 6, Dec. 2005.
- [28] E.L. Petersen, J.B. Longworthy, S.E. Diehl, "Suggested single event upset figure of merit," *IEEE Transactions on Nuclear Science*, vol. 30, no. 6, pp. 4533-4539, Dec. 1983.
- [29] E.L. Petersen, "SEE rate calculations using the effective flux approach and a generalized figure of merit approximation," *IEEE Transactions on Nuclear Science*, vol. 42, no. 6, Dec. 1995.
- [30] E.L. Petersen, "The SEU figure of merit and proton upset rate calculations," *IEEE Transactions on Nuclear Science*, vol. 45, no. 6, pp. 2550-2562, Dec. 1998.
- [31] N.J. Rudie, Unpublished, 1996.
- [32] A.J. Tylka, J.H. Adams, Jr., P.R. Boberg, B. Brownstein, W.F. Dietrich, E.O. Flueckiger, E.L. Petersen, M.A. Shea, D.F. Smart, and E.C. Smith, "CREME96: A revision of the cosmic ray effects on micro-electronics code," *IEEE Transactions on Nuclear Science*, vol. 44, no. 6, pp. 2150-2160, Dec. 1997.
- [33] P. Hazucha, T. Karnik, J. Maiz, S. Walstra, B. Bloechel, J. Tschanz, G. Dermer, S. Hareland, P. Armstrong, and S. Borkar, "Neutron soft error rate measurements in a 90-nm CMOS process and scaling trends in SRAM from 0.25 μ m to 90nm generation," *Proceedings of the IEEE International Electronic Development Meeting*, pp. 523-526, Dec. 2003.

- [34] R. Baumann, "The impact of technology scaling on soft error rate performance and limits to the efficacy of error correction," *Digest of International Electron Devices Meeting*, pp. 329-332, Dec. 2002.
- [35] R. Baumann, "Radiation-induced soft errors in advanced semiconductor technologies," *IEEE Transactions on Device and Materials Reliability*, vol. 5, no. 3, pp. 305-316, Sep. 2005.
- [36] R. Baumann, "Soft errors in advanced semiconductor devices-part I: the three radiation sources," *IEEE Transactions on Device and Materials Reliability*, vol. 1, no. 1, pp. 17-22, Mar. 2001.
- [37] T. Calin, M. Nicolaidis, and R. Velazco, "Upset hardened memory design for submicron CMOS technology," *IEEE Transactions on Nuclear Science*, vol. 43, no. 6, pp. 2874-2878, Dec. 1996.
- [38] M. Pflanz et al., "On-Line Error Detection and Correction in Storage Elements with Cross-Parity Check," *Proceedings of the Eighth IEEE International On-Line Testing Workshop*, pp. 69-73, 2002.
- [39] D.R. Blum, "VLSI implementation of cross-parity and modified DICE fault-tolerant schemes," M.S. Thesis, Washington State University, May 2004.
- [40] L.R. Rockett, Jr., "An SEU-Hardened CMOS Data Latch Design," *IEEE Transactions on Nuclear Science*, vol. 35, pp. 1682-1687, Dec. 1988.
- [41] D.R. Blum and J.G. Delgado-Frias, "Schemes for eliminating transient-width clock overhead from SET-tolerant memory-based systems," *IEEE Transactions on Nuclear Science*, vol. 53, no. 3, pp. 1564-1573, Jun. 2006.
- [42] M.J. Myjak and J.G. Delgado-Frias, "A two-level reconfigurable architecture for digital signal processing," *Proceedings of the 2003 International Conference on VLSI*, pp. 21-27, Las Vegas, NV, Jun 2003.
- [43] M.J. Myjak, "A two-level reconfigurable cell array for digital signal processing," M.S. Thesis, Washington State University, May 2004.
- [44] M.J. Barry, "Radiation Resistant SRAM Memory," U.S. Patent No. 5,157,625, Oct. 1992.
- [45] J.G. Dooley, "SEU-Immune Latch for Gate Array, Standard Cell, and Other ASIC Applications," U.S. Patent No. 5,311,070, May 1994.
- [46] V. Stojanovic and V.G. Oklobdzija, "Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems," *IEEE Journal of Solid-State Circuits*, vol. 34, no. 4, Apr. 1999.

- [47] G. Gerosa, S. Gary, C. Dietz, P. Dac, K. Hoover, J. Alvarez, H. Sanchez, P. Ippolito, N. Tai, S. Litch, J. Eno, J. Golab, N. Vanderschaaf, and J. Kahle, "A 2.2 W, 80 Mhz superscalar RISC microprocessor," *IEEE Journal of Solid-State Circuits*, vol. 29, pp. 1440-1452, Dec. 1994.
- [48] D.R. Blum and J.G. Delgado-Frias, "SEU and SET-tolerant pipelined systems," *Second IEEE International Conference on Space Mission Challenges for Information Technology*, MW-4 Mini-Workshop, Pasadena, CA, Jul. 2006.
- [49] J. Tschanz, S. Narendra, Z. Chen, S. Borkar, and M. Sachdev, "Comparative delay and energy of single edge-triggered & dual edge-triggered pulsed flip-flops for high-performance microprocessors," *Proceedings of the 2001 International Symposium on Low Power Electronics and Design*, pp. 147-152, Huntington Beach, CA, Aug. 2001.
- [50] A. Ma and K. Asanovic, "A double-pulsed set-conditional-reset flip-flop," *MIT Laboratory for Computer Science Technical Report 844*, May 2002.
- [51] M.J. Myjak, J.G. Delgado-Frias, and S.K. Jeon, "An energy-efficient differential flip-flop for deeply pipelined systems," *Proceedings of the 49th IEEE International Midwest Symposium on Circuits and Systems*, San Juan, Puerto Rico, Aug. 2006.
- [52] V. Stojanovic and V.G. Oklobdzija, FLIP-FLOP, U.S. Patent No. 6,232,810, May 2001.
- [53] S. Heo, R. Krashinsky, and K. Asanovic, "Activity-sensitive flip-flop and latch selection for reduced energy," *Proceedings of the 2001 Conference on Advanced Research in VLSI*, pp. 59-74, Salt Lake City, UT, Mar. 2001.
- [54] S. Heo and K. Asanovic, "Load-sensitive flip-flop characterization," *Proceedings of the IEEE Computer Society Workshop on VLSI 2001*, pp. 87-92, Orlando, FL, Apr. 2001.
- [55] O. Musseau, F. Gardic, P. Roche, T. Corbiere, R.A. Reed, S. Buchner, P. McDonald, J. Melinger, L. Tran, and A.B. Campbell, "Analysis of multiple bit upsets (MBU) in CMOS SRAM", *IEEE Transactions on Nuclear Science*, vol. 43, no. 6, pp. 2879-2888, Dec. 1996.
- [56] R. Koga, S.D. Pinkerton, T.J. Lie, and K.B. Crawford, "Single-word multiple-bit upsets in static random access devices," *IEEE Transactions on Nuclear Science*, vol. 40, no. 6, pp.1941-1946, Dec. 1993.
- [57] A. Makihara, H. Shindou, N. Nemoto, S. Kuboyama, S. Matsuda, T. Oshima, T. Hirao, H. Itoh, S.Buchner, and A.B. Campbell, "Analysis of single-ion multiple-bit

- upset in high-density DRAMs,” *IEEE Transactions on Nuclear Science*, vol. 47, no. 6, pp. 2400-2404, Dec. 2000.
- [58] T. Merelle, F. Saigne, B. Sagnes, G. Gasiot, Ph. Roche, T. Carriere, M.-C. Palau, F. Wrobel, and J.-M. Palau, “Monte-Carlo simulations to quantify neutron-induced multiple bit upsets in advanced SRAMs,” *IEEE Transactions on Nuclear Science*, vol. 52, no. 5, pp. 1538-1544, Oct. 2005.
- [59] S. Buchner, A. Campbell, T. Meehan, K. Clark, D. McMorro, C. Dyer, C. Sanderson, C. Comber, and S. Kuboyama, “Investigation of single-ion multiple-bit upsets in memories on board a space experiment,” *IEEE Transactions on Nuclear Science*, vol. 47, no. 3, pp. 705-711, Jun. 2000.
- [60] B. Gill, M. Nicolaidis, and C. Papachristou, “Radiation induced single-word multiple-bit upsets correction in SRAM,” *Proceedings of the 11th IEEE International On-Line Testing Symposium*, pp. 266-271, Jul. 2005.
- [61] G. Hubert, N. Buard, C. Weulersse, T. Carriere, M.-C. Palau, J.-M. Palau, D. Lambert, J. Baggio, F. Wrobel, F. Saigne, and R. Gaillard, “A review of DASIE code family: contribution to SEU/MBU understanding,” *Proceedings of the 11th IEEE International On-Line Testing Symposium*, pp. 87-94, Jul. 2005.
- [62] A.B. Campbell, O. Musseau, V. Ferlet-Cavrois, W.J. Stapor, and P.T. McDonald, “Analysis of single event effects at grazing angle,” *IEEE Transactions on Nuclear Science*, vol. 45, no. 3, pp. 1603-1611, Jun. 1998.
- [63] D.R. Blum and J.G. Delgado-Frias, “Hardened by design techniques for implementing multiple-bit upset tolerant static memories,” *Proceedings of the 2007 IEEE International Symposium on Circuits and Systems*, to be published, New Orleans, LA, May 2007.
- [64] W.J. Stapor and P.T. McDonald, “Practical approach to ion track energy distribution,” *Journal of Applied Physics*, vol. 64, no. 9, pp. 4430-4434, Nov. 1988.

Appendix A

SEU and SET-Tolerant Reconfigurable DSP

Architecture Layout

In this appendix, the CMOS VLSI layout of the SEU and SET-tolerant reconfigurable DSP architecture detailed in Chapter 3 is presented. This layout consists of over 15k transistors in a 0.18 μm , five metal layer technology. Figure A.1 is an illustration of the layout. This figure appears on the following page, and it has been rotated counter-clockwise by 90° to obtain a better fit within the page dimensions.




Figure A.1: VLSI Layout of the Radiation-Tolerant Reconfigurable DSP Cell-Switch-Cell Circuitry.

Appendix B

Publications

This appendix lists research papers produced during the course of the work presented in this dissertation. Included in this list are journal and conference papers that have been published, submitted for publication, or accepted for publication. The M.S. dissertation of the author is included below as well.

B.1 Journal

1. D.R. Blum and J.G. Delgado-Frias, "Schemes for Eliminating Transient-Width Clock Overhead from SET-Tolerant Memory-Based Systems," *IEEE Transactions on Nuclear Science*, vol. 53, no. 3, pp. 1564-1573, Jun 2006.
2. D.R. Blum and J.G. Delgado-Frias, "Delay and Energy Analysis of SEU and SET-Tolerant Pipeline Latches and Flip-Flops," submitted to the *IEEE Transactions on Nuclear Science*.

B.2 Conference

1. J.G. Delgado-Frias, M.J. Myjak, F.L. Anderson, and D.R. Blum, "A Medium-Grain Reconfigurable Cell Array for DSP Applications," *IASTED International Conference on Circuits, Signals and Systems*, pp. 231-236, Cancun, Mexico, May 2003.
2. D.R. Blum and J.G. Delgado-Frias, "A Fault-Tolerant Memory-Based Cell For a Reconfigurable DSP Processor," *2003 International Conference on VLSI*, pp. 21-27, Las Vegas, NV, Jun 2003.

3. M.J. Myjak, D.R. Blum, and J.G. Delgado-Frias, "Enhanced Fault-Tolerant CMOS Memory Elements," 2004 IEEE International Midwest Symposium on Circuits and Systems, pp. 453-456, Hiroshima, Japan, Jul 2004.
4. D.R. Blum, M.J. Myjak, and J.G. Delgado-Frias, "Enhanced Fault-Tolerant Data Latches for Deep Submicron CMOS," 2005 International Conference on Computer Design, pp. 28-34, Las Vegas, NV, Jun 2005.
5. D.R. Blum and J.G. Delgado-Frias, "Comparison of SET-Resistant Approaches for Memory-Based Architectures," 12th NASA Symposium on VLSI Design, Coeur d'Alene, ID, Oct 2005.
6. D.R. Blum and J.G. Delgado-Frias, "SEU and SET-Tolerant Pipelined Systems," Second IEEE International Conference on Space Mission Challenges for Information Technology, Mini-Workshop on Student Aerospace Research, Pasadena, CA, Jul 2006. (Best student paper award).
7. D.R. Blum and J.G. Delgado-Frias, "Hardened by Design Techniques for Implementing Multiple-Bit Upset Tolerant Static Memories," accepted for publication at the 2007 IEEE International Symposium on Circuits and Systems, New Orleans, LA, May 2007.
8. D.R. Blum, J.G. Delgado-Frias, and S.M. Ray, "SEU and SET-Tolerant Memory-Based Reconfigurable DSP Processor," accepted for publication at the 13th NASA Symposium on VLSI Design, Post Falls, ID, Jun 2007.

B.3 M.S. Thesis

1. D.R. Blum, "VLSI Implementation of Cross-Parity and Modified DICE Fault-Tolerant Schemes," M.S.E.E. Thesis, Washington State University, May 2004.