UNIVERSITY OF SOUTHAMPTON

FACULTY OF PHYSICAL AND APPLIED SCIENCES

Electronics and Computer Science

Variation and Reliability in Digital CMOS Circuit Design

by

Massoud Mokhtarpour Ghahroodi

Thesis for the degree of Doctor of Philosophy

April 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

Doctor of Philosophy

VARIATION AND RELIABILITY IN DIGITAL CMOS CIRCUIT DESIGN

by Massoud Mokhtarpour Ghahroodi

The silicon chip industry continues to provide devices with feature sizes at Ultra-Deep-Sub-Micron (UDSM) dimensions. This results in higher device density and lower power and cost per function. While this trend is positive, there are a number of negative side effects, including the increased device parameter variation, increased sensitivity to soft errors, and lower device yields. The lifetime of next-generation devices is also decreasing due to lower reliability margins and shorter product lifetimes.

This thesis presents an investigation into the challenges of UDSM CMOS circuit design, with a review of the research conducted in this field. This investigation has led to the development of a methodology to determine the timing vulnerability factors of UDSM CMOS that leads to a more realistic definition of the Window of Vulnerability (WOV) for Soft-Error-Rate (SER) computation.

We present an implementation of a Radiation-Hardened 32-bit Pipe-lined Processor as well as two novel radiation-hardening techniques at Gate-level. We present a Single-Event-Upset (SEU) tolerant Flip-Flop design with 38% less power overhead and 25% less area overhead at 65nm technology, compared to the conventional Triple Modular Redundancy (TMR) technique for Flip-Flop design. We also propose an approach for in-field repair (IFR) by trading area for reliability. In the case of permanent faults, spare logic blocks will replace the faulty blocks on the fly. The simulation results show that by tolerating approximately 70% area overhead and less than 18% power overhead, the reliability is increased by a factor of x10 to x100 for various component failure rates.

Contents

List of Figures

List of Tables

Nomenclature

ABB	Adaptive Body Biasing
$ASET$	Analogue Single Event Transient
AVS	Adaptive Voltage Scaling
BTI	Bias Temperature Instability
CDF	Cumulative Distribution Function
CMP	Chemical Mechanical Polishing
$CWSP$	Code Word State Preserving
$DCLS$	Dual Core Lock Step
DMR	Dual Modular Redundancy
$DVFS$	Dynamic Voltage and Frequency Scaling
ECC	Error Correcting Code
EDA	Electronic Design Automation
ELT	Enclosed Layout Transistors
FIT	Failure In Time
H	Wire Height
HBD	Hard Break Down
HCI	Hot-Carrier Injection
HMR	Hybrid Modular Redundancy
$ITRS$	International Technology Roadmap for Semiconductors
L	Length
LET	Linear Energy transfer
MBU	Multi-Bit Upset
MCU	Multi-Cell Upset
MPU	Memory Protection Unit
$MTBF$	Mean Time Between Failures
$MTTF$	Mean Time To Failures
$NBTI$	Negative Bias Temperature Instability
NMR	N Modular Redundancy
PCA	Principal Component Analysis

PDF	Probability Density Function
PLL	Phase-locked Loop
RSM	Response Surface Methodology
RTL	Register Transfer Language
SBD	Soft Break Down
SEB	Single Event Burnout
SEE	Single Event Effect
$SEFI$	Single Event Functional Interrupt
SEL	Single Event Latch-up
SER	Soft Error Rate
SET	Single Event Transient
SEU	Single Event Upsets
$SILC$	Stress Induced Leakage Current
SNR	Signal-to-Noise Ratio
SOI	Silicon-on-Insulator
STA	Static Timing Analysis
$TDDB$	Time-Dependent Dielectric Breakdown
TCM	Tightly Coupled Memories
TMR	Triple Modular Redundancy
Tox	Oxide Thickness
TRC	Two-Rail Checkers
TSC	Totally Self Checking
UV	Ultraviolet
$VHDL$	VHSIC Hardware Description Language
VRM	Voltage Regulator Module
Vth	Threshold Voltage
WM	Wire Width
WOV	Window of Vulnerability
$UDSM$	Ultra Deep Sub Micron

Declaration of Authorship

I, Massoud Mokhtarpour Ghahroodi , declare that the thesis entitled *Variation and Reliability in Digital CMOS Circuit Design* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as papers in peer-reviewed international conferences;

Signed:.....

Date:.....

Acknowledgements

I owe my deepest gratitude to my supervisor Professor Mark Zwolinski for his continuous help and support in this work. It has been an honour to be his Ph.D. student. I appreciate all his contributions of time, ideas, and funding to make my Ph.D. experience productive and stimulating.

I would like to extend my appreciation to Cisco Systems Inc., ARM Ltd. and HiPEAC (European Network of Excellence on High Performance and Embedded Architecture and Compilation) for their support for the development of the work in this thesis.

To my dad.

Chapter 1

Introduction

1.1 Motivations For Research

This work is motivated by the challenges faced in designing reliable circuits in modern technology nodes. As the feature size of CMOS devices shrink to Ultra-Deep-Sub-Micron (UDSM) dimensions, intrinsic parameter fluctuations of atomic scale transistors drastically impact the power, performance and yield of manufactured chips and limit the scaling and integration of them. In other words, due to intrinsic variations, unavoidable in modern fabrication processes, the taped-out chips can vary radically as every single UDSM transistor out of the billions of transistors on a die have different characteristics.

Moreover the incredible growth and complexity of semiconductor fabrication facilities has resulted in the isolation of process/device engineers from circuit design engineers, leading to some lack of understanding of the impact of circuit designs upon manufacturability and testability due to the fundamental limitations of technology and device physics. Most of today's technologies are subject to very high defect density. Increasing defect density decreases yield and with such a high defect densities in UDSM chips, the manufacturing costs can be prohibitively high, making chip yield a critical metric for manufacturers. From a performance point of view, the circuit must meet its speed requirements over a range of voltages and temperatures that reflect the environment in which the circuit will operate; while the performance requirements must be met at a set of worst-case conditions for speed, the power requirements must be simultaneously met at another set of worst case conditions.

Moreover, from a reliability perspective, the incidence of transient errors increases at UDSM dimensions, and consequently, the dependability of the systems decreases. Reliability is therefore another required metric for every UDSM-based system in general and safety-critical systems in particular.

Therefore any design methodology must consider the Power, Performance, Yield and Reliability figures of merit in the design and manufacturing flow. At present, there are neither methodologies nor EDA tools that can capture the full complexity of these problems and be used successfully to predict both the characteristics and scale of the intrinsic fluctuations in UDSM transistors and interconnects, and their subsequent impacts on the power-performance-yield and reliability of circuits and systems.

1.2 Power-Performance-Yield-Reliability

1.2.1 Power

Power consumption which is the rate of energy dissipation in a system is always an important issue; but due to the technology scaling to UDSM dimensions with more and more transistors on a chip, power consumption, thermal and cooling issues are some of the major problems in chip design especially for battery operated devices. Although the dominant factor of power consumption is the dynamic power dissipation which is related to the operating voltage, frequency, capacitance and the required performance specifications, at UDSM, leakage power dissipation is not ignorable anymore, in such a way that as reported in the literature, leakage power will represent up to 40 percent of total power consumption in near-future devices [25].

1.2.2 Performance

As the name suggests, higher performance or faster computation is always desired. However applying today's design methodologies and traditional deterministic worst-case timing analysis at UDSM is too pessimistic [12]. Moreover, due to the unavoidable variations in global buffer loads and interconnect wire length, the global clock signal arrives at different components at different times. This phenomenon which is known as the clock skew makes the traditional globally

synchronous design methodology impractical; because the clock signal cannot be distributed to all of the millions of flip-flops and registers on the chip at the same time significantly degrading the achievable clock speed.

1.2.3 Yield

Generally, the key metric in determining the success of a technology in chip design (especially from the market perspective) is the number of devices on the die which are fully functional and applicable. This generally necessitates that every single transistor out of billions of transistors on a chip work properly. This issue is known as yield. Any malfunction in a single device can potentially lead to yield loss. This yield loss can be either catastrophic or parametric; catastrophic yield loss is caused by physical defects (such as stuck-at, opens, bridging faults) which typically manifest themselves as functional failures on the chip , leading to defective chips which must be thrown away. In parametric yield loss, manufactured devices do not perform according to the design specification i.e. chip functionality is correct but they may work slower or consume more power than expected in the design process.

For a long time the parametric yield was not considered serious and catastrophic yield loss was the main yield issue and various solutions such as adding redundancy and fault-defect tolerant methodologies were used to surmount this type of yield loss. But at UDSM dimensions, not only catastrophic yield loss is important but due to the variations, the parametric yield loss is rapidly increasing as well [26].

1.2.4 Reliability

Semiconductor manufacturing continues to provide smaller feature sizes resulting in lower power, higher density, and lower cost per function. While this trend is positive, there are a number of negative side effects, including increased semiconductor parameter variability, increased sensitivity to soft errors, and lower device yields. These issues become more and more important for the semiconductor industry and modelling is increasingly required to provide design tools not only to achieve better device performance but also more robust reliability margins. Dependability, or the ability of a system to function correctly under given operating conditions during a given period of time can be quantified using measures of Reliability or Failure In Time (FIT). The lifetime of next-generation devices is decreasing due

to lower reliability margins and shorter product lifetimes putting the reliability and dependability of such systems at stake [10].

Another major emerging reliability problem that cannot be furthermore ignored in UDSM technologies is the susceptibility to Soft Errors. As clearly stated in The International Technology Roadmap For Semiconductors(ITRS) 2009-2011 [24]:"The impact of Soft Error Rate (SER) over the years is almost constant or even increasing in spite of the reduced sensitivity for the single units due to device scaling and the use of countermeasures (e.g., SOI, redundancy, error detection and correction). This is because of the corresponding increase of the number of units in a system. Viable models and simulators are still lacking to extrapolate the SER from the cell up to the system level from accelerated tests, which are able to keep track of the error propagation and to provide enough statistical accuracy."

1.3 Objectives

The main objectives of this research are:

- To investigate the impacts of variation and reliability issues on UDSM CMOS circuits from a design perspective.
- To investigate timing vulnerability of UDSM combinational circuits and present a more realistic methodology to determine the vulnerability.
- To investigate timing vulnerability of UDSM sequential circuits and state-holding elements and propose a possible hardening-by-design solution.
- To investigate ageing and the reliability issues of UDSM circuits and processors and propose a repair mechanism to avoid fatal shut-downs.

1.4 Contributions

This thesis provides a survey of various UDSM impacts on circuits and devices, reviewing current research and providing a summary of the state-of-the-art techniques to mitigate the UDSM impacts. Moreover, techniques are introduced to deal with UDSM impacts in terms of performance and reliability. These include a novel radiation-hardened flip-flop design and an in-field logic repair mechanism for UDSM reliable circuit design.

1.5 Thesis Structure

In this chapter, the main problems, obstacles and motivations for this research have been briefly presented. The second chapter provides a survey of various UDSM impacts on circuits and devices in the literature and the existing techniques to tackle them. First, the pre-silicon or design-time techniques are discussed and then post-silicon or run-time solutions are provided, followed by a brief survey of the reliability issues. Chapter 3 addresses timing vulnerabilities mainly due to soft errors in combinational logic. Chapter 4 discusses timing vulnerability of sequential circuits and state-holding elements. Chapter 5 provides a discussion of other reliability issues such as ageing and possible in-field repairable architectures. The last chapter summarises the conclusions of the thesis.

Chapter 2

Literature Review

In this chapter, we take a brief survey of the major impacts of UDSM scaling on design. First we discuss process and intrinsic parameter variations, and their impacts on Performance and Power consumption. Then we take a brief survey of the effects of radiation and soft errors on UDSM CMOS circuits. Afterwards, the major reliability issues will be discussed and finally, state-of-the-art and the proposed solutions in the literature to tackle process variation and to mitigate soft errors and the reliability issues will be discussed.

2.1 Process Variation

One of the major impacts of UDSM on logic are increased levels of fabrication process variation and additional random uncertainties caused by the random placement of dopant atoms in the channel of each transistor at UDSM dimensions. This intrinsic randomness of the placement of atoms, along with the extrinsic limitations in controlling the manufacturing process and its precision, have impacted various parameters such as oxide thickness (T_{ox}), threshold voltage (V_{th}) and transistor channel length (L) directly. The impacts of process variation are not limited to the transistors. Interconnect parameter variation on wire height (H) and wire width (WM) are also increasing dramatically. Such phenomena will result in huge variation in gate delays and interconnect delays. Moreover due to nanometre scale geometries of devices with very thin (angstrom scale) gate oxide layers, the reliability issues of such UDSM designs have been increasing at very high rates [27] [28] [29].

This is mainly due to the inability of semiconductor manufacturing industries to improve tolerance levels in the fabrication-lithography stage and keep up with the technology scaling. For instance, the light source (with a wavelength of 193 nm) used in lithography in older technologies (above 130 nm) is still used in newer technologies [30]. Lithography tolerances are limited by the granularity of the resist materials. for UDSM process, various immersion and multiple exposures lithography techniques are used which will result more expensive manufacturing processes.

The doping density of the transistor channel is the major determinant of threshold voltage in bulk and polysilicon gate MOSFETs. To achieve the desired doping density and consequently the desired threshold voltage, certain number of dopant atoms are required. Due to the fact that the implantation of dopant atoms in devices is random, the eventual exact number of dopant atoms in the transistor channel is also random. Because of technology scaling, such number of atoms in the channel region is becoming smaller, as the channel volume decreases, and thus the relative effect of a single change in dopant number is increasing [30] [31]. There is a similar situation for interconnects, however the main factors which are responsible for variation in interconnects are the limitations in process control over the manufacturing process.

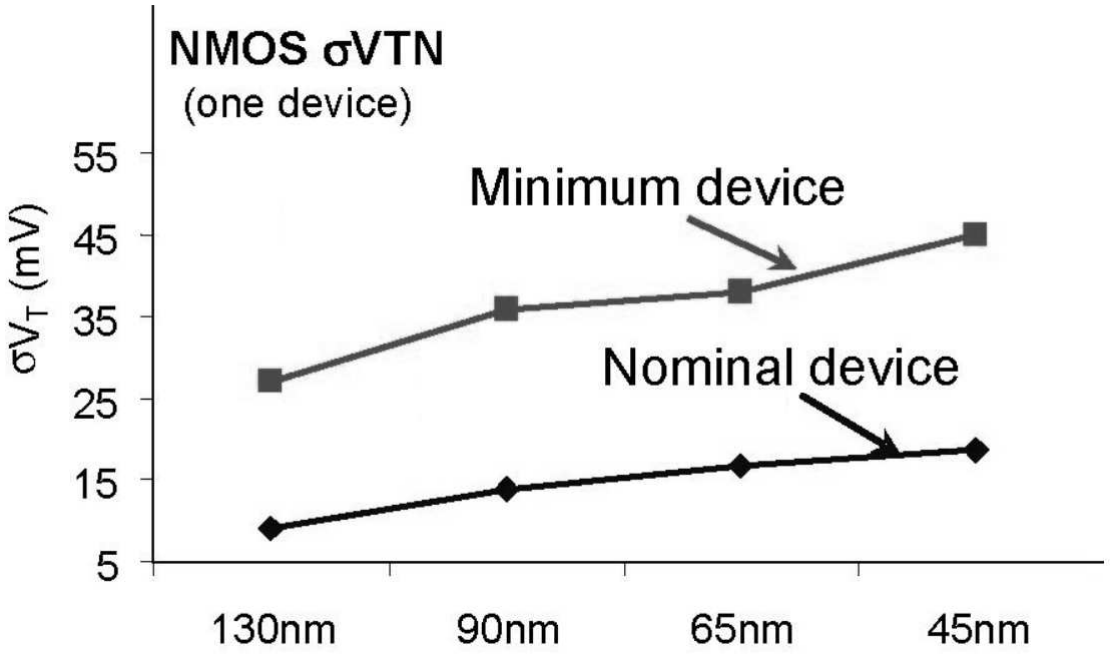

Figure 2.1: Variation in threshold voltage of devices [1]

Fig. 2.1, shows how the variance of the threshold voltage of an NMOS transistor has almost doubled between the 45nm and the130nm technology nodes. It can also be seen that the absolute value of the threshold voltage in 130nm technology is about 0.35V which is higher than 45nm technology (0.28V approx.). Hence the prediction of the device behaviors and eventually estimating the circuit performance metrics in the presence of device and interconnect variation is a major challenge for the chip industry.

From manufacturing perspective, one can classify the sources of variation in the transistor threshold voltage into two main categories: Global and Local. Global variation is caused by manufacturing process variations and local intrinsic variation is caused by local parameter fluctuations. For short channel transistors (channel length = 20nm in 45nm technology), the threshold voltage variation is caused by transistors geometries, and specifically the Leff parameter which can be classified as being local. In the case of long channel (channel length = 40nm in 45nm technology) transistors, the variation on V_{th} is mainly due to dopant diffusions, gate dielectric thickness, ion implantation and so forth that can be considered as being global [32]. For long channel transistors it is the average dopant diffusion that is more important. As shown in Fig. 2.2, the source and drain doping is very dense, but the channel doping is very vulnerable to variation.

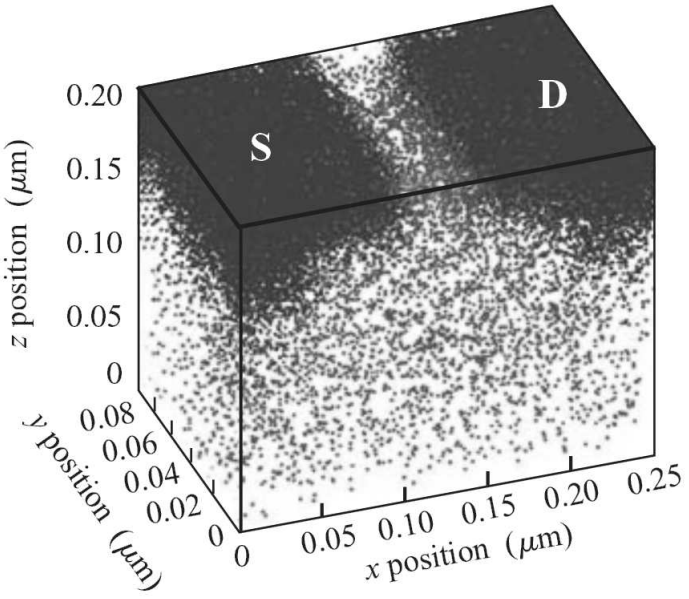

Figure 2.2: Random placement of dopant atoms a 50-nm channel-length MOSFET [2]

From another perspective, we can categorize variation as being either random or systematic [30] [33] [34]. Random physical effects such as poly-silicon gate line-edge roughness, random fluctuations of the number and the location of dopant atoms in the MOSFET channel and so forth [35] [33] [30]. Systematic variations are usually due to spatial dependencies of the manufacturing mechanisms for device processing, such as variation in chemical mechanical polishing (CMP) which produces predictable variation trends across the die [36], or predictable variations such as those caused by optical proximity effects [37].

It is also noteworthy to mention that when comparing the delay contribution of device variation to interconnect variation, delay variations due to interconnect are still less significant and device variations still hold the biggest share of the total delay variability in UDSM technologies [30]. The contribution of fluctuations in device parameters is about 90% of the total delay variation of a design in the real world [33].

2.1.1 Performance

Typically, the performance of a circuit is determined by the speed of the circuit which is rated by the operating clock frequency. Delays are generally the boundaries that determine this operating frequency. The rate at which information can propagate through the circuit depends on the longest path delay, and from a synchronous design perspective, the maximum clock frequency of a circuit is limited by the path with the maximum delay. Variation in process parameters will cause distributed small delay variations along any given path and when summed up, the path can become timing critical and even fail to meet the timing constraints, hence the chip can fail as a result. Random variations can cause a significant mismatch in the electrical performance of two identical devices placed next to each other. On the other hand, performance and timing verification in the presence of process variation is difficult because the critical path is no longer unique. This means different paths can become timing critical depending on the process-voltage corners that the manufactured chip is coming from. Therefore Critical Path selection and analysis cannot be deterministic anymore [38].

2.1.2 Power Consumption

Power consumption and thermal issues have always been important, but due to the high device density and integration of logic with UDSM dimensions, the problem has become even more serious. While dynamic power dissipation is the dominant component in the total power consumption, in UDSM, leakage power in forms of gate leakage or sub-threshold leakage is increasingly becoming a severe problem [39].

Thermal issues and generated heat in to days and future high performance and highly integrated devices is another significant factor. Dynamic power consumption activity can produce local hot spots on the die. These local hot spots can be several tens of degrees hotter than the rest of the die, even after application of the best cooling techniques to the package itself. As depicted in Fig. 2.3, a temperature difference of 40C to 50C corresponds to a 20 percent performance variation [3].

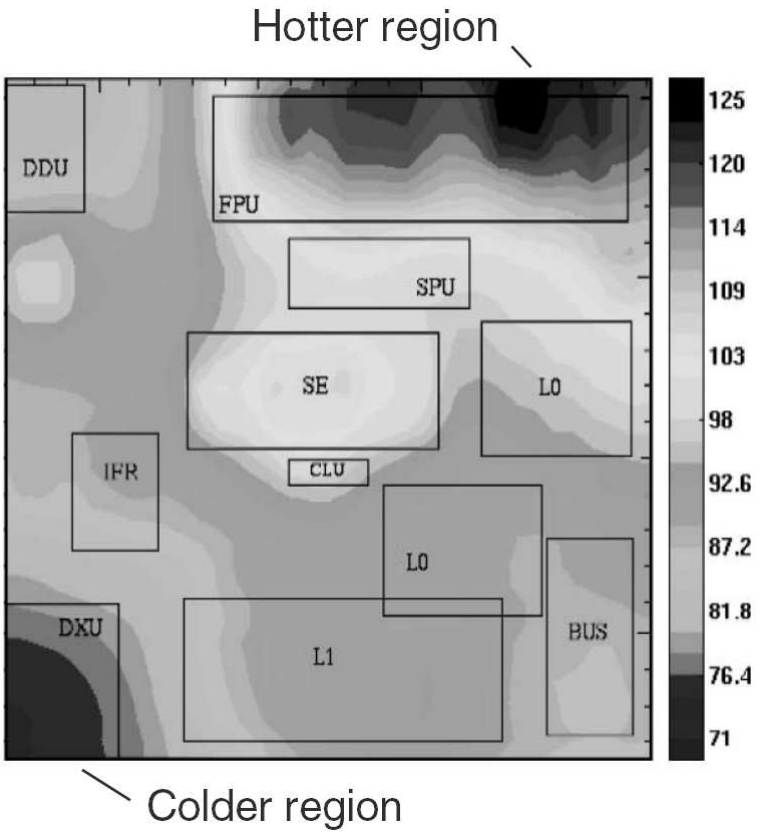

Figure 2.3: Temperature differences on a die: 40C to 50C temperature difference leads to 20% performance variation [3]

As it can be seen in Fig. 2.3, there are huge temperature differences between the cache area and the core area, with the core area being much hotter and the cache region is much colder also depicted in Fig. 2.4. This non-uniformity in thermal gradients is very prevalent in micro processors and it can be observed by doing local calculation (or local sensing) on the power densities at different locations on a chip. From a dynamic power dissipation point of view, slower devices are the result of higher temperatures. However according to the power consumption formula (i.e. C $\times$ V^2 $\times$ f), the total consumed power will remain the same. Nevertheless, the main issue will be because of the leakage power that grows exponentially that can potentially cause major IR drop issues [26]. In other words, dynamic power grows linearly with chip frequency (and since chip frequency used to be proportional to scaling, power draw would scale linearly with device shrink) but leakage power is increasing exponentially with device shrink.

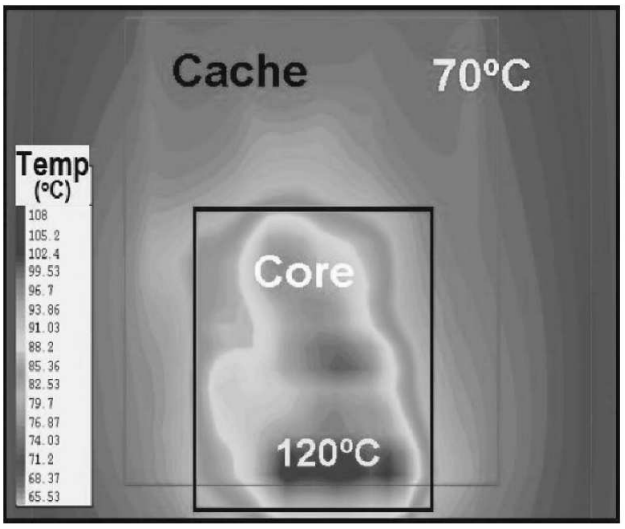

Figure 2.4: A Temperature Distribution Map of a Typical Chip with a Core and Cache [3]

IR drop is supply voltage drop across the chip. According to Ohms Law V=I $\times$ R, where R is the equivalent path DC resistance between the source location and the cell/macro location and I is the average current the chip draws from the supply down the paths. The power grid or the power mesh is comprised of multi-level metal structures. This includes planes, vias and tracks that feed the all of the standard cells and the memories and the macros across the chip. The supply voltage is produced by a voltage regulator module (VRM), which is usually a DC-DC converter that is connected to the power grid and distributes power across the chip.

Wire resistance can lead to excessive current draw from the power mesh that can cause significant performance degradations and signal integrity issues (an illustrated 2.5) . Due to smaller interconnect geometries the power mesh at UDSM scales, the level of vulnerability to power supply non-uniformity and IR drop is significant. The manifestation of such issues will be more sensitivity to noise and increased delay variation and eventually delay faults and timing errors. These impacts are aggravated when there are gates with different supply voltage levels that are connected across the chip communicating through level-shifters [4].

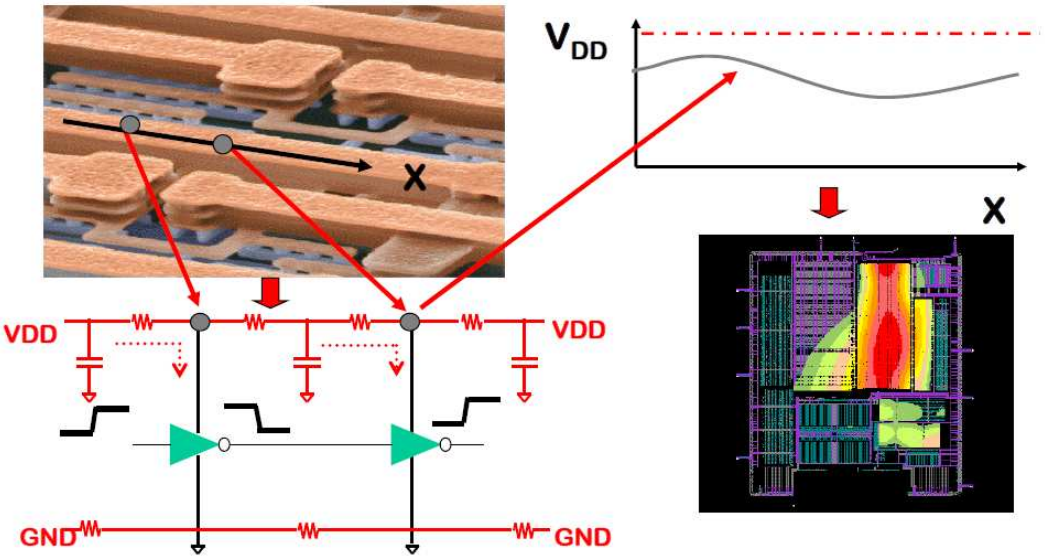

Figure 2.5: IR Drop in Power Distribution Network due to non-ideal components [4] [5]

The effective amount of resistance through the path between the voltage regulator module and the design blocks on the chip is the metric to estimate the IR drop. It should be mentioned that the topic of IR drop is not limited to the design of blocks on the chip. There are three main categories to be considered: On-chip IR drop, Package IR drop, and Board-level IR drop. Because of UDSM dimensions, accurate analysis of IR drop on the chip is critical as this is the most significant factor in determining if the chip is going to fail because of IR drop issues. However recently package IR drop and board level IR drop have become more important and their contribution to the over-all IR drop budget cannot be ignored [40]. This is mainly because of decreased supply voltage and increase vulnerability to noise issues that can have fatal impacts on the operation of high speed circuits. The increase in temperature and the creation of hot-spot across the chip will also add up to the IR drop that can potentially lead the chip towards failure [41].

2.2 Radiation and Soft Errors

The Earth and its surroundings are protected by the atmosphere, which acts as a filter, to let throughout visible light and heat, while stopping a significant amount of radiation and Ultraviolet (UV) light. Because of this natural protection , human beings and electronic devices are able to cope with solar flares, solar winds and cosmic rays. As reported in NASA reference publications [42] and also in [6], the two major sources of environment related spacecraft anomalies are, statistically, plasma and radiation effects, i.e. effects related to the charged particles from the space environment.

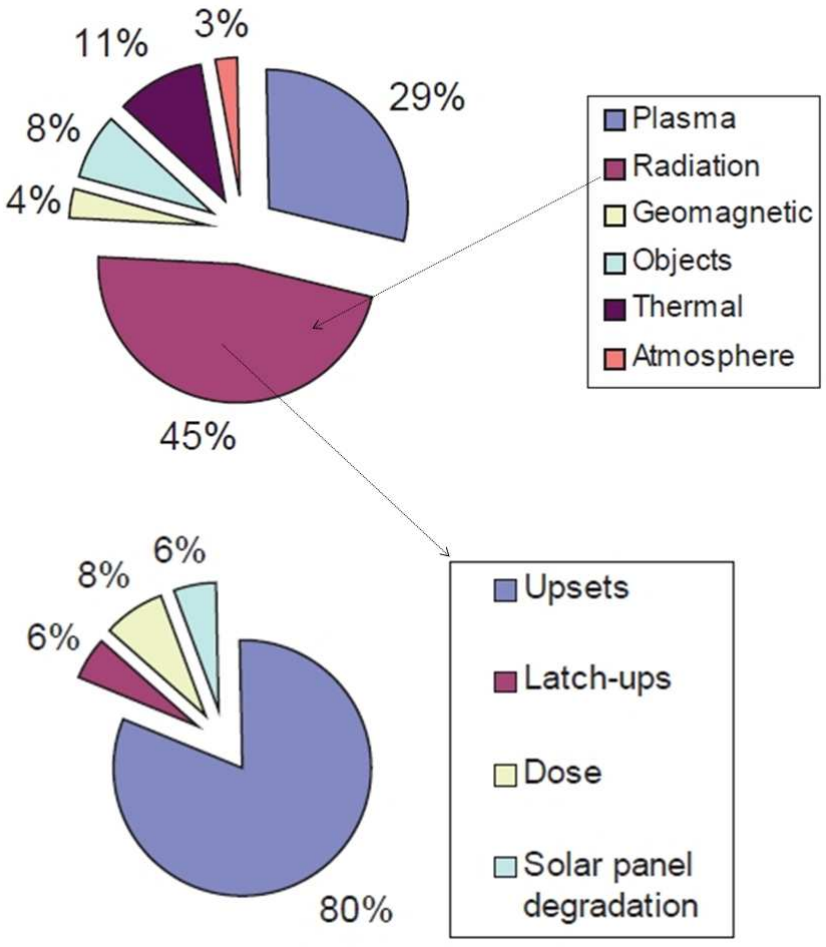

Figure 2.6: Spacecraft anomalies due to the space environment [6]

With technology scaling, radiation particle strikes are becoming increasingly problematic for both combinational circuits and memory elements even at sea level. The first report of serious industrial problem due to soft errors goes back to 1978 on the 2107-series 16-KB DRAMs by Intel. It was reported that the errors were

caused by the traces of radioactivity due to α particles in the package materials which led to radiation-induced Single-Event-Upsets (SEU) at sea level, referred to as "soft errors" [43]. From that era until now, radiation-induced problems have been some of the most challenging reliability issues in circuits and systems, not only in safety-critical applications and avionics, but also for Commercial, off-the-shelf (COTS) products. Therefore, the circuits used in these application must be tolerant to radiation particle strikes. In this section, we take a brief survey of radiation-induced errors on circuits and systems.

2.2.1 Single Event Effects Definition

Soft errors are a subset of non-destructive Single-Event Effects (SEEs) [44]. The interaction of nuclear particles with electronic components can create a series of SEEs. Such effects can be categorized as hard effects and soft effects. Hard effects or hard errors are permanent and non-recoverable. Soft errors are temporary and might be recoverable by applying power shut down, reset or rewriting the corrupted data. In CMOS-based circuits, the main hard error issues are Single Event Burnout (SEB) that can occur in power MOS devices, SEGR or die-electric breakdown caused by single event effects and micro-dose-induced threshold voltage variations due to SEEs in CMOS transistors. The PNPN parasitic structures can also be vulnerable. A Single Event Latch-up (SEL) can cause a strong current which can lead to overheating of the device, that if it is not stopped by a power cycle, it can have destructing impacts on the transistor. These hard errors are not discussed in this thesis, but they represent the most significant hard issues in the topic of SEEs.

Multi-Cell Upset (MCU) will occur when one high energy particle hits many state holding elements in a given clock cycle. On the other hand, Multi-Bit Upset (MBU) occurs when more than one bit of a word is struck by a single particle. There is another phenomenon known as Single Event Functional Interrupt (SEFI), that can happen in more sophisticated circuits and systems. SEFIs can cause loss of functionality because of perturbation of clocks or control registers that can lead to long periods of malfunctions in the system. Recovery might be obtained by switching off and back on, or rewriting configuration registers, or by applying a reset [45].

2.2.2 Major Soft Error Problems

Silicon devices have become more susceptible to radiation and energetic particle strikes. The energetic particle strikes can create localized ionization events in the silicon devices and if this happens in the sensitive region on the CMOS device, the resulting electron-hole pair can cause a transient current pulse that may alter the logic state of the struck node as depicted in Fig.2.7 and Fig. 2.8. This is known as a single event upset (SEU) on a memory element since it can upset the storage elements; if the particles strike any combinational node they can cause a transient current pulse that eventually becomes a voltage pulse at the output of the struck node. It is known as a single event transient (SET) on a combinational element [46] [47] [48] [49]. The transient pulse caused by a particle strike can be captured by the sequential elements depending on the existence of an active path from the struck node to the storage element, the arrival time and the width of the transient pulse at the storage element input as shown in Fig. 2.9 [50] [51]. The errors caused by SEUs or SETs are known as major soft error issues.

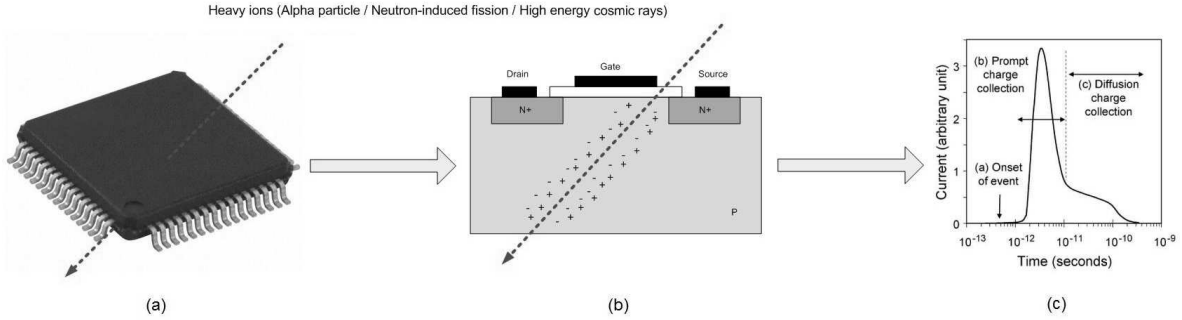

Figure 2.7: Illustration of single event transient pulse generation. Funnelling in an n+/p silicon junction following the ion strike and the resulting electrical transient current caused by the passage of a high-energy ion [7] [8].

As mentioned earlier, the transient pulse can be captured by a flip-flop and cause an error, provided that it is not masked by any of the following three derating factors or masking phenomena:

- Logic Masking happens when the particle strikes either a non-controlling input of a combinational logic gate, or the transient pulse is filtered out by other controlling nodes on the path to a sequential element.
- Electrical Masking occurs for transient pulses which will be attenuated due to characteristics of CMOS gates, such as size and load capacitance.

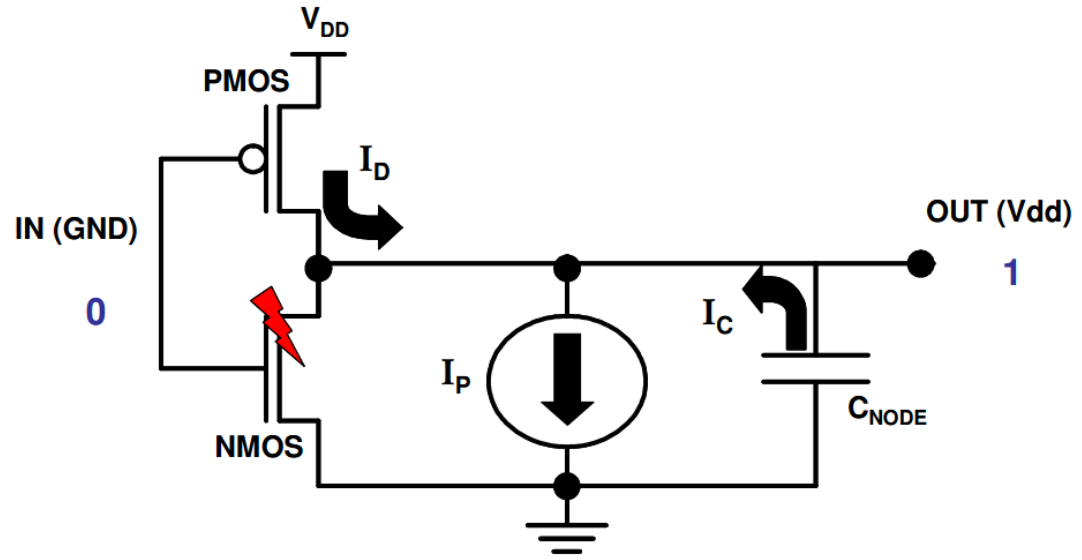

Figure 2.8: Particle strike on a sensitive node

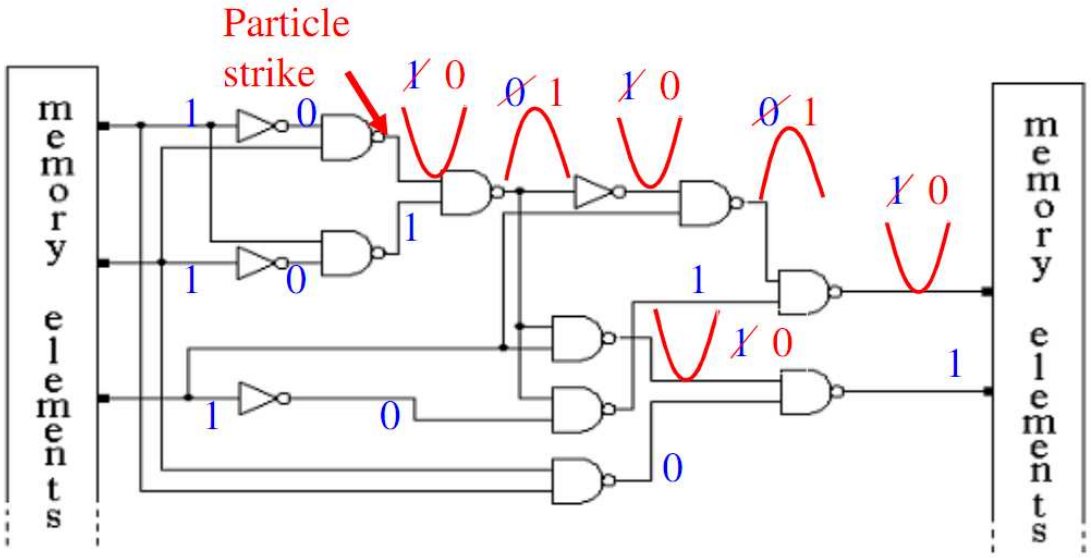

Figure 2.9: SET in Combinational Logic

- Temporal Masking occurs when the transient pulse is either narrower than the flip-flop window of vulnerability or the transient pulse is wide enough but reaches the memory element outside of the clock transition region and hence it is not sensed and captured.

The main sources of Soft Errors are reported to be [52] [47] :

- Alpha particles (caused solely by silicon packaging and radioactive impurities)

- Neutron-induced fission (interaction of neutrons from cosmic rays and boron in the silicon packaging)
- High-energy cosmic rays

The physical effects of radiation-induced charges on devices and circuits can be classified as: Direct and Indirect ionization. In direct ionization, as a high energy particle passes through a semiconductor material, it generates electron-hole pairs. As the radiation particle passes through the silicon, it loses its energy and after expending all its energy, the particle will come to rest. To define the transferred energy from the particle, the linear energy transfer (LET) value is used. LET is defined as the transferred energy that the radiation particle induced to generate the electronhole pair per unit length, normalized by the density of the target material (for VLSI designs, this is the density of Silicon) [31].

Indirect ionization consists of a light radiation particle with high energy that passes through the semiconductor material. Such particle can have a collision with the nucleus that can lead to a nuclear reaction. Protons and neutron particles are good examples indirect ionization. Such phenomenon can also create secondary particles such as heavy ions or alpha particles. Such secondary particle can then go through a direct ionization process and if the charge gets placed in different locations across the chip, multiple soft-errors can arise [31] [53].

Advances in packaging and fabrication have gradually reduced the effect of the alpha-particle induced soft errors and the neutron-induced soft errors dominate in most UDSM circuits [52]. Experimental results with heavy ions and alpha particles indicate that SET pulse widths can range from about 100 ps to over 1 ns for the 90-nm process. Such pulse widths are comparable to valid logic signals in 130nm and 90nm processes and indicate that as technology is scaled to lower operating voltages and higher operating frequencies, SETs may become a significant reliability problem [54] [55] [56].

The soft-error rate (SER) is measured in FIT units (failures in time), where 1 FIT denotes one failure per billion device hours (i.e., one failure per 114,077 years). For electronic systems, usually the SER values range between a few hundreds and around 100,000 FIT. This is about one soft error per year. The failure rate induced by soft errors can be relatively high in electronic devices, compared to other reliability issues that will be discussed later. The experimental results show that the failure rate for hard errors (for example latch-ups) is approximately equal to or lesser than 10 FIT. However, the soft-error rate is much higher. For instance,

the SER for an SRAM block with the size of 1 Mbit is usually of the order of 1,000 FIT in UDSM process technologies, which makes the memory blocks some of the most vulnerable parts of the chip [57]. The situation becomes even worse for systems with multiple memory blocks (which is the case for most of modern chips) in a way that it exceeds the cumulative failure rates because of other reliability issues. However, it should be noted the consequences of soft errors are totally different from hard errors. In the case of soft errors, the fault usually disappears when the system is reset or new data replaces the corrupted data, hence the damage is not permanent.

As depicted in Fig. 2.10, in previous technologies, memory elements such as SRAMs, flip-flops and latches contributed more to the overall SER of the chip and the contributions of of combinatorial logic to the overall SER was much lower. Therefore, memories (SRAMs, DRAMs, and latches) were mainly under consideration as they were more vulnerable to radiation particle hits. However, as the feature size of CMOS devices go below UDSM, the contribution of combinatorial logic gates to the overall SER has take a much bigger share, while the contribution of memory elements such as SRAMs to the overall SER has relatively remained constant. The is due to the fact that by using deeply pipe-lined circuits, the lengths of the combinatorial logic paths have been reduced dramatically which can result in the reduction of the masking or derating factors. In other words, because fewer number of SETs will get filtered out, the particle hits on combinatorial circuits can cause more faults and aggravate the overall SER of the chip.

As the clock frequency increases, the probability that transient pulses will be captured as valid data in combinational logic increases linearly. Particularly in the case of deeply pipe-lined processors, with an increase in circuit speeds, the chances of a given transient pulse propagating through the combinatorial circuit and getting latched increases, because the combination paths will become shorter and there will be pipeline registers at every pipeline stage latching the data, hence increasing the probability of an SET getting latched. However, we can also speculate that the duration of transient pulses decreases. Nevertheless due to both their higher chances to propagate in high-speed circuits and their higher probability of getting captured by the next stage state holding elements, such as flip-flops and latches, SETs have been predicted to become a very critical issue in deep Ultra-Deep Sub-Micron circuits [57].

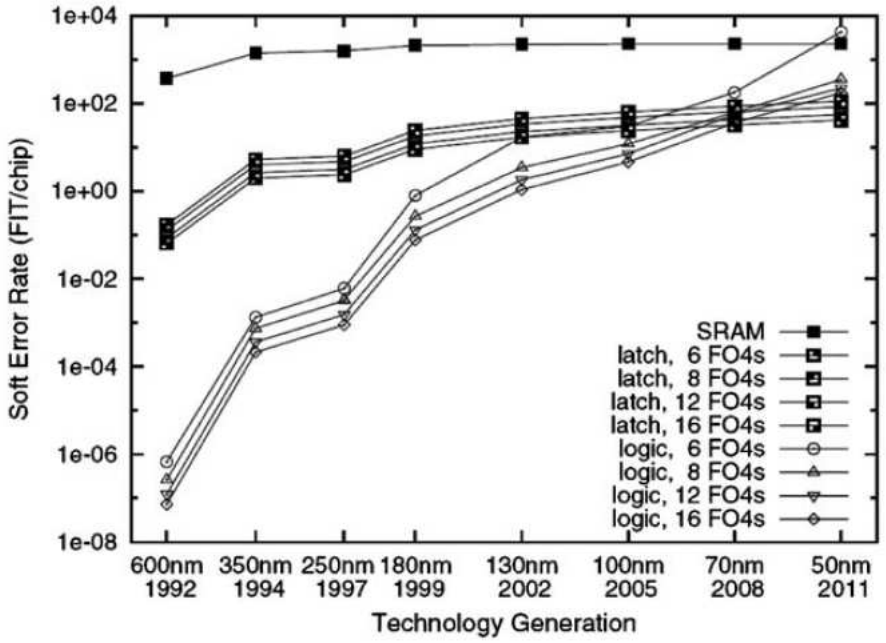

Figure 2.10: SER of an alpha processor for different technology nodes [9]

2.3 The Reliability Issues

Semiconductor manufacturing continues to provide smaller feature sizes, resulting in lower power, higher density, and lower cost per function. While this trend is positive, there are a number of negative side effects, including increased semiconductor parameter variability, increased sensitivity to soft errors, and lower device yields as mentioned before. The lifetime of the next-generation devices is also decreasing due to lower reliability margins and shorter product lifetimes. As demonstrated in [58], the average Mean Time To Failure (MTTF) of a modern-day super-scalar processor has dropped by approximately 4X between the 180nm to 65nm technology nodes. Design for reliability and resilience in the long term, is one of the major challenges flagged in the International Technology Roadmap for Semiconductors (ITRS) 2011 report [24].

Reliability is defined as the probability that a system or a device perform a specific function up to a specific time interval, in a pre-defined environment. Dependability can be defined as the ability of a system to deliver service at an acceptable level of confidence in either presence or absence of faults [20]. The metrics to calculate the dependability of a system usually consist of the assessments of availability and reliability, along with acceptable fault coverage and how the system meets the safety requirements [59].

2.3.1 Major Reliability Issues in Ultra Deep-Sub-Micron CMOS

Negative Bias Temperature Instability (NBTI), Hot-Carrier Injection (HCI) degradation and Time-Dependent Dielectric Breakdown (TDDB) or "Wear-Out" of MOS devices are some of the most important reliability concerns for UDSM designs.

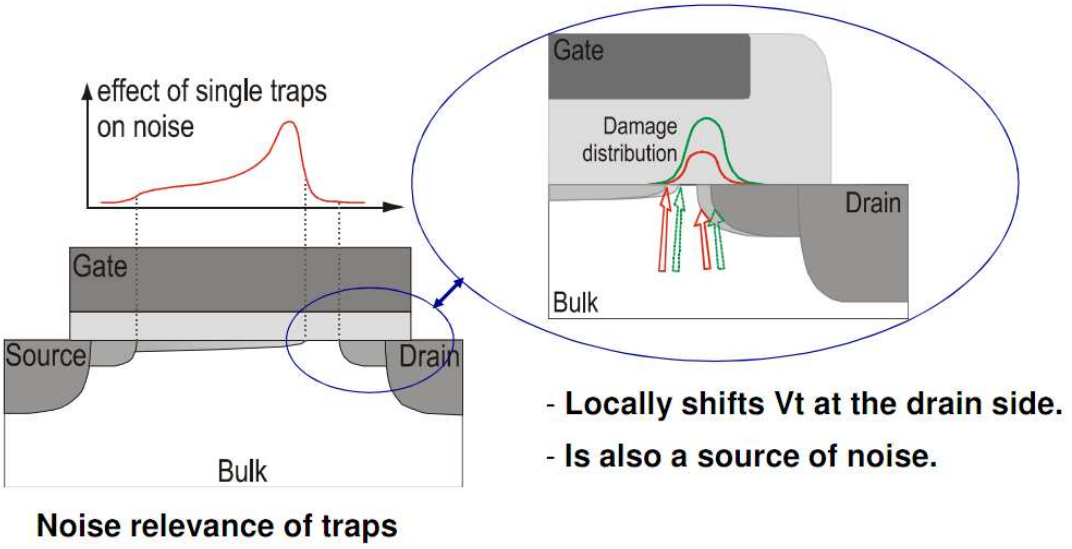

Figure 2.11: Hot carrier stress generates additional trap states near to the drain

2.3.2 Negative Bias Temperature Instabilities (NBTI) & Hot-Carrier Injection (HCI)

Negative bias temperature instability (NBTI) in pMOSFETs is considered a major reliability issue in Ultra Deep-Sub-Micron analogue and digital integrated circuits [60] [10]. This phenomenon occurs when a PMOS transistor is turned on at high temperatures (usually between 100 °C and 150 °C). When the gate of a PMOS transistor is negatively biased with respect to the substrate, defects are induced in the device, resulting in permanently reduced drive current and threshold voltage (Vth) shifts [10] [61].

In more details, the NBTI is caused the generation of traps at the $Si - SiO_2$ interfaces due to electrical stress on PMOS transistors. The manifestation of such electrical stress is reduction in channel mobility of the MOSFETs, through an increase in threshold voltage or the induction of parasitic capacitances which

degrade the performance. This challenging issue that the chip industry is facing can change the performance metrics of circuits and dramatically reduce the lifetime of a chip over time. An interesting phenomenon is that the threshold voltage can partially recover to its initial value when the gate bias is switched to 0V (interface traps can be alleviated partially when the electrical stress is reduced). However this Vth recovery is logarithmically time-dependent [10] [62]. As illustrated in [63] [64] [10], a substantial recovery in Vth is observed when the electrical stress is interrupted. This phenomenon is depicted in Fig 2.12 [65].

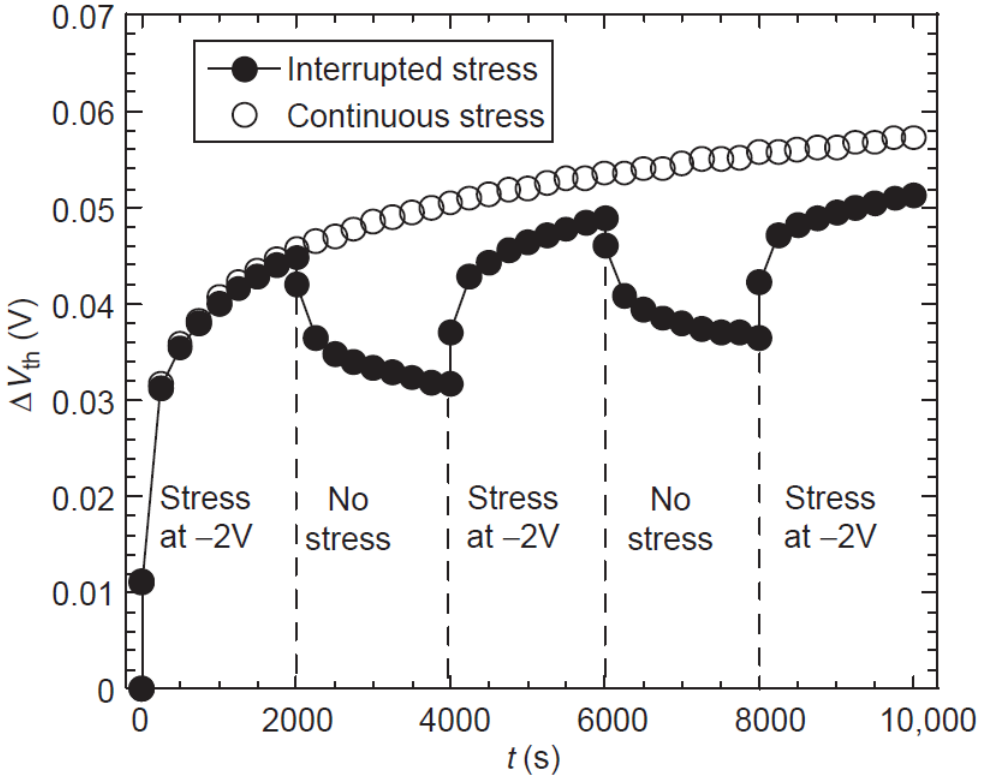

Figure 2.12: Vth differences as a function of stress time, showing the threshold voltage degradation during the stress and the partial-recovery when the gate bias is switched to 0V. From [10].

NBTI recovery can have advantages and drawbacks. The advantage would be in simpler circuits, in which the circuit can partially get closer to its nominal performance when the stress is removed. However for complex systems, this may backfire. For instance, in circuits with a lot of power saving features, such as clock gating, the clock-gated component do not age as much as the rest of the circuit, which means, after the removal of the clock-gating, those specific component will be faster than the rest of the circuit. This potentially can cause hold violations.

To explain NBTI, a hydrogen-release model is usually used. Under high temperature and applied voltage, the interface between the channel and the oxide will be

hit by high energy holes. The result of this will be the breaking of silicon-hydrogen bonds which will lead to the release of hydrogen atoms due to electrochemical reactions with the oxide interface. At the oxidechannel interface, positively charged traps will be created due to the combination of free hydrogen atoms with nitrogen or oxygen atoms. This leads to a shift in PMOS threshold voltage to become more negative as a result of the reduction in holes mobility. Moreover these effects impact the performance of the transistors by degrading the drive current of the devices [19] [66].

Another principal degradation issue of MOSFETs is hot-electron-induced depassivation, also known as HCI, of the $Si - SiO_2$ interface that limits the operating lifetime of the transistors [67]. Injection of hot-carriers can result in shifts in threshold voltages and trans-conductance degradation in CMOS devices. This is also caused by defects at the $Si - SiO_2$ which has been mitigated for the current generation of MOSFET devices. The manifestation of HCI is similar to NBTI i.e. reducing the transistor performance and shifting devices metrics. The damage is caused by hot carriers heating up in the high electric field near the drain side of the MOSFET that can lead to impact ionization and eventually degradation of device parameters.

The reason that such carriers are called Hot Carriers is that they are highly energetic. The process of ionization at the drain, produces electron-hole pairs. The substrate current I_{sub} will increase when some of these hot carriers enter the substrate region. Those carriers with high enough energy levels (i.e. 3.1 eV or higher for electrons and 4.6 eV or higher for holes) can potentially cross the oxide barrier and enter the oxide and consequently cause defects [19].

The conventional method to measure HCI degradation is by measuring the drain saturation current (I_{Dsat}) degradation. The reason is that I_{Dsat} is one of the key transistor parameters that can be used to determine the HCI-induced impacts on the circuit performance particularly because the HCI issues happen during the normal operation of the circuit (i.e. when the circuit is active) while the transistor is in saturation mode. A method to deal with HCI-induced issues is frequency guard-banding as HCI is directly related to the activity of the devices.

The expected lifetime of a silicon chip is often between 5 and 15 years [25]. Usually, the frequency degradation during the expected lifetime is between 1% and 10% [19]. Therefore, usually the manufactured chips are margined a few percent below the highest frequency at which they can actually operate. This frequency

marginalizing is known as the frequency guard-banding. Transistor lifetime degradation due to HCI (for example a 3% reduction in the threshold voltage) is typically speculated for the situation in which the chip is actually running (i.e. power on mode) [25].

HCI used to be more important in NMOS devices historically. This is because of the lower effective mass of electrons. Their mobilities are higher than holes and as a result, they can obtain higher levels of energy from the electric field in the channel of the transistors. Also NBTI has a slower rate of degradation comparing to HCI. It has been known that HCI usually happens in an NMOS device during the low to high transition at the gate input, This also means high switching activity or higher clock frequencies can increase HCI-induced ageing. Furthermore, the recovery in HCI is so small that it is negligible, which makes HCI the worst in stresses under AC conditions [68].

In any manufactured chip, the CMOS devices go through various stress conditions at different times and every stress condition will have its own degradation impacts on the devices. For instance, in a CMOS inverter, both the NMOS and the PMOS devices are connected to the same input voltage. So for example, when the input of the inverter gate is set to low (0V), the PMOS transistor experiences NBTI stress and therefore degrades while the NMOSFET is shut down. When the input transition from low (0v) to high (VDD), the NMOSFET goes through impact ionization condition and HCI degradation occurs. At the same time, the PMOS transistor is shut down and some of the NBTI-induced impacts can alleviate [68]. Due to the fact that each degradation mechanism (NBTI, HCI, also depends on signal transitions) generates defects either in the bulk oxide or at the interface, the overall MOSFET degradation can get very complex and modelling such phenomenon accurately is a challenge.

2.3.3 Time-Dependent Dielectric Breakdown (TDDB) or "Wear-Out"

The continued scaling of MOSFET devices requires ultra-thin gate dielectrics for controlling the short channel effect. This has reduced the reliability of the dielectric layer leading to dielectric breakdown over time due to the formation of a conductive path through the oxide to the substrate [69]. Even though TDDB has been studied for over three decades, the exact physical mechanism remains unclear. What is known is that the process is driven by voltage and temperature [10]. Major studies

have shown that electron fluence (current) and energy (voltage) are the driving factors for wearing out and eventual breakdown [10]. Oxide break down can be categorized into hard break down (HBD) and soft break down (SBD); HBD is considered as a catastrophic failure of the device and hence the entire circuit while SBD events do not cause immediate failure of the CMOS device but will affect the performance of the circuit [70]. Typically after soft breakdown the leakage current is only slightly larger than the pre-stress tunnelling characteristic. After some time the leakage current can continue to increase, finally resulting in a hard breakdown. To limit the thermal damages of TDDB, the power dissipation needs to be reduced. To do so, either the supply voltage needs to reduced or the percolation path current needs to be decreased by the application of resistance in series.

NMOS in inversion used to be the major factor of TDDB-induced lifetime degradations in previous technologies. But as reported by Intel in [11], in their 45nm CMOS technology, TDDB can occur on NMOS and PMOS under all operating bias conditions. During normal device operation, the electric field across the gate dielectric causes the generation of electrical defects which are known as "traps". The local electric field can then be impacted by such traps and leakage current can increase in the dielectric to a point where a conductive "chain" is formed between the cathode and the anode as depicted in Fig. 2.13.

The statistical theory that describes this process is called the Percolation Theory [71]. The percolation assumes that the traps are generated inside the oxide at random locations. At the vicinity of these traps, a sphere is considered with a constant radius 'r', and conduction happens when the sphere of two random traps overlap as depicted in Fig. 2.13. In the case of UDSM devices, the dielectric thickness is getting thinner (1.2nm in 65 nm technology consists of only a few monolayers of SI-O bonds), resulting in more susceptibility to gate current leakage and eventually leading to TDDB [72].

Although there is contradictory consensus in the literature on the exact physical mechanisms that lead to gate dielectric breakdown. It is generally accepted that a combination of several mechanisms such as trap-assisted conduction, charge injection, as well as bulk trap state generation contribute to TDDB. Through the constant presence of stress, more trap states are created and, eventually, there will be a gradual increase in the gate current. This phenomenon is known as Stress Induced Leakage Current (SILC) degradation [72] [11].

In summary, the impacts of technology scaling on CMOS circuits and devices manifest themselves as process variation leading to performance/power issues, more

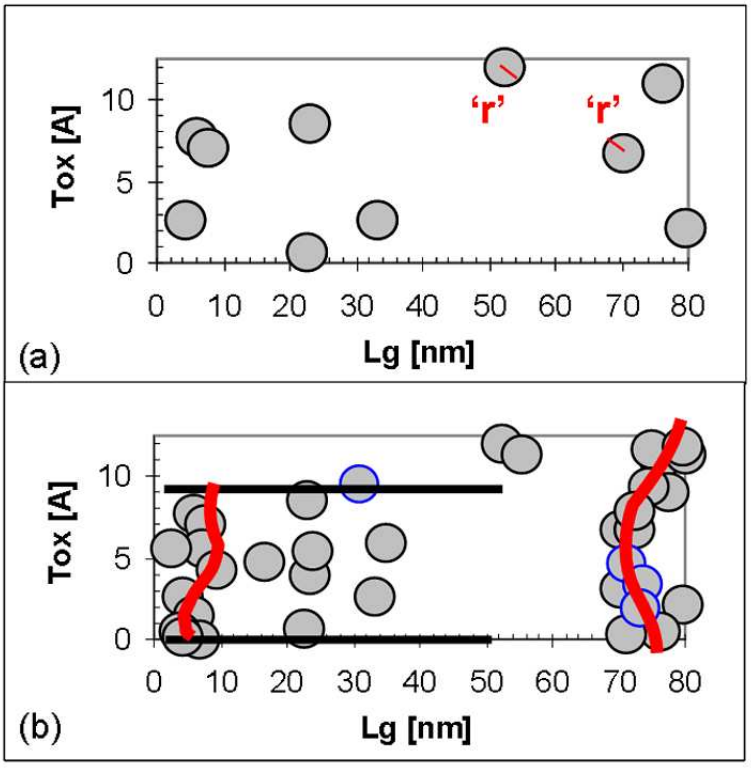

Figure 2.13: Percolation Theory describes traps as spheres of radius "r. When several of them form a complete chain from anode to cathode, breakdown (BD) occurs. The thinner the dielectric, the fewer the traps needed to cause Break down [11].

vulnerability to soft errors and causing severe reliability issues such as HCI/NBI and oxide break down. Simultaneous interactions of all of the aforementioned reliability issues, calls for a challenging comprehensive solution to deal with all of them. In the next section of this chapter, we take a brief survey of proposed techniques to cope with UDSM impacts on design.

2.4 Solutions and state-of-the-art

In the previous section, we took a brief survey of major challenges in UDSM CMOS devices. In this part we review the main techniques for the analysis and the mitigation of UDSM impacts on circuits. First, the Pre-Silicon techniques or static techniques that are applicable at design time are discussed. Next Post-Silicon or dynamic techniques that are based on adaptability and are applicable at Run-Time, are discussed.

2.4.1 Tackling Variations At Design-Time

2.4.1.1 Analysis Techniques

Traditionally, circuit performance is measured by deterministic timing analysis and by considering the path with the maximum delay or the worst-case or critical path. But in UDSM systems any path in a particular chip, can potentially become critical, depending on how variations manifest themselves on that particular chip. This phenomenon results in a circuit whose operation may be logically correct but does not perform at the required operating voltage and frequency [38] [73] [74].

The probabilistic nature of the timing behaviour of UDSM systems strongly suggests that statistical analysis and simulation should play a role in the selection and testing of critical paths. In statistical timing analysis, the propagation delays are modeled as random variables with given probability density functions (pdfs). By providing the gate-level netlist, design constraints, required clock frequency, the probability density functions of the cells from pin-to-pin and the interconnect delays, one would need to calculate the PDF of the actual signal arrival times, the required time and the slacks of primary outputs and the internal signals. Using all this data, we can compute the delay of the longest paths for setup timing checks, the shortest paths for hold timing tests and conclusively the determine the probable maximum speed of the design [75].

Statistical Timing Analysis Techniques

The field of Statistical Timing Analysis has been an active area of research and thus the literature is full of solid approaches that in many ways are built on each other; where each new approach vies to improve on a limitation or a short-coming of a previous approach.

At transistor level, there are diverse models and methods for statistical analysis and optimization Such as Monte Carlo, Response Surface Methodology (RSM), Principal Component Analysis (PCA), Projection-based performance modelling (PROBE) and Asymptotic Probability Extraction (APEX) methods [76].

Due to the huge size of the circuits and the complexity of computations in terms of number of transistors and random variables, it is not feasible to model and analyze variations for every single gate on a chip. Therefore analysis at higher abstraction levels is also performed.

From one point of view, statistical timing analysis can be divided in to two general categories:

- Path-Based Methods
- Block-Based Methods

Path-Based Methods

These methods are based on performing timing analysis on a selected set of critical paths in a circuit. Generally path-based methods are inefficient at UDSM scales, due to the uncertainty that exists in critical path selection.

Block-Based Methods

Block-Based Methods are similar to timing graph traversal which is preformed in traditional Static Timing Analysis (STA). But instead of pre-determined or nominal delays for each node, delay distributions are propagated through the timing graph. The advantage of this method in comparison with path-based methods is that there is no need for critical path selection. However due to the usage of delay distributions, the computational complexities of these methods are noteworthy.

Traditional Block-Based STA Methods are based on two atomic operations: SUM() and MAX() as depicted in Fig. 2.14. In a timing graph in which each node represents a logic gate, the value of SUM() for the node(j) is the sum of the Arrival Time(AT) from the previously traversed node (i) and the propagation delay from previous node (i) to current node(j). The Value of MAX() is the maximum Arrival Time(AT) of all the incoming paths to the current node(i). These two operations are repetitively executed to traverse the whole timing graph from the source node to the sink node.

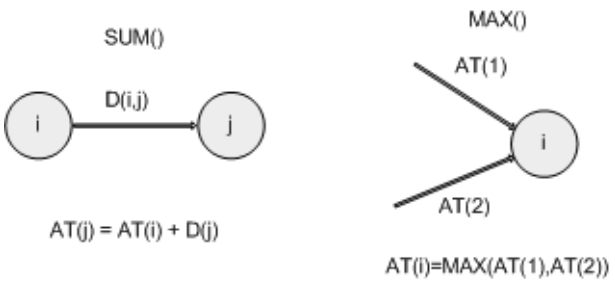

Figure 2.14: SUM and MAX Operations

Note that for statistical block-based timing analysis, SUM() and MAX() operations must calculate probability distributions instead of nominal deterministic values. Most of the techniques proposed in the literature are based on the assumption that variations are Normally Distributed. However, in practice, this is not proven.

In [77] and [78], algorithms are propose which are capable of calculating SUM() and MAX() by estimating PDF/CDF of either Normal or Non-Normal Distributions of arrival times, provided that all the distributions are mutually independent. But in practice, correlations exist; for instance, the arrival times can be correlated due to shared-paths or correlated process variations that results in joint or even multi-dimensional PDF/CDF which are computationally expensive to perform.

In [74] a variation-aware method based on statistical timing to select critical paths is introduced, in which node criticalities are computed to determine the probabilities of different circuit nodes being on the critical path across process variation. This methodology is aimed at uncovering performance violations in defect-free integrated circuits, however at nano-scale, physical defects are more likely to happen due to the fact that IC manufacturing process is inherently imperfect.

2.4.1.2 Implementation Techniques

At the architectural level, we can maintain the performance of the circuit and keep the supply voltage as low as possible (which is the ultimate goal in being power and energy efficient), by using parallel architectures also known as pipe-lining. As discussed in [79], this can be done by using parallelized circuits in a way that a bigger function will be broken into smaller functions, and each small function will be assigned to one of the parallelized circuits. By doing this the clock frequency requirements can be relaxed per parallelized circuits, provided that we can meet the target latency [80].

Gosh *et al.* [81] [82] proposed a pipeline-based design paradigm to achieve robustness with respect to timing failure and provide an opportunity for aggressive voltage scaling by critical path isolation. In their methodology called CRISTA, a set of possible paths that may become critical under process variations are predicted and isolated by increasing timing slack between critical and non-critical parts and their rare activation is ensured and afterwards any possible delay failure in the critical paths is avoided by dynamically changing to two-cycle operation using clock stretching (assuming all standard operations are single cycle). The drawback is that it is not generic and can be utilized only for pipeline designs considering single cycle operations.

Kourtev *et al.* [12] used clock skew scheduling techniques to decrease the number of paths with the maximum delay. They used clock scheduling (i.e. applying non-zero clock skews) to increase or decrease the amount of path delays. By applying this technique, a "shift" of the path delay distribution away from the maximum path delay can be achieved as depicted in Fig.2.15. There are two beneficial effects of that shift of delay which are either the circuit can be run at a lower clock period (or higher clock frequency) or the circuit can operate at the target clock period with a reduced probability of setup and hold time violations (improving the overall system reliability). However, this technique requires careful modification of clock distribution network to adjust desirable delays for each clocked-element at design time, but due to variations, the design-time properties and adjustments of the clock network may not be preserved after manufacturing.

Most of the techniques mentioned above are aimed at modelling, predicting, optimizing and accommodating power and performance issues caused by variations at pre-silicon stage or at design-time. On the other hand, there is another paradigm that addresses variation issues at post-silicon stage or run-time. These techniques are based on the idea of adaptive or tunable systems. In the next part, we will review these techniques.

2.4.2 Tackling Variations At Run-time

To facilitate low power, high performance, high yield products which are based on less reliable UDSM devices, post silicon or run-time techniques are introduced and applied for variety of applications. Here the main idea is instead of over-designing and over-calculation and simulations to cover all the variations, device parameters such as supply voltage or Body bias voltage are set based on the information such

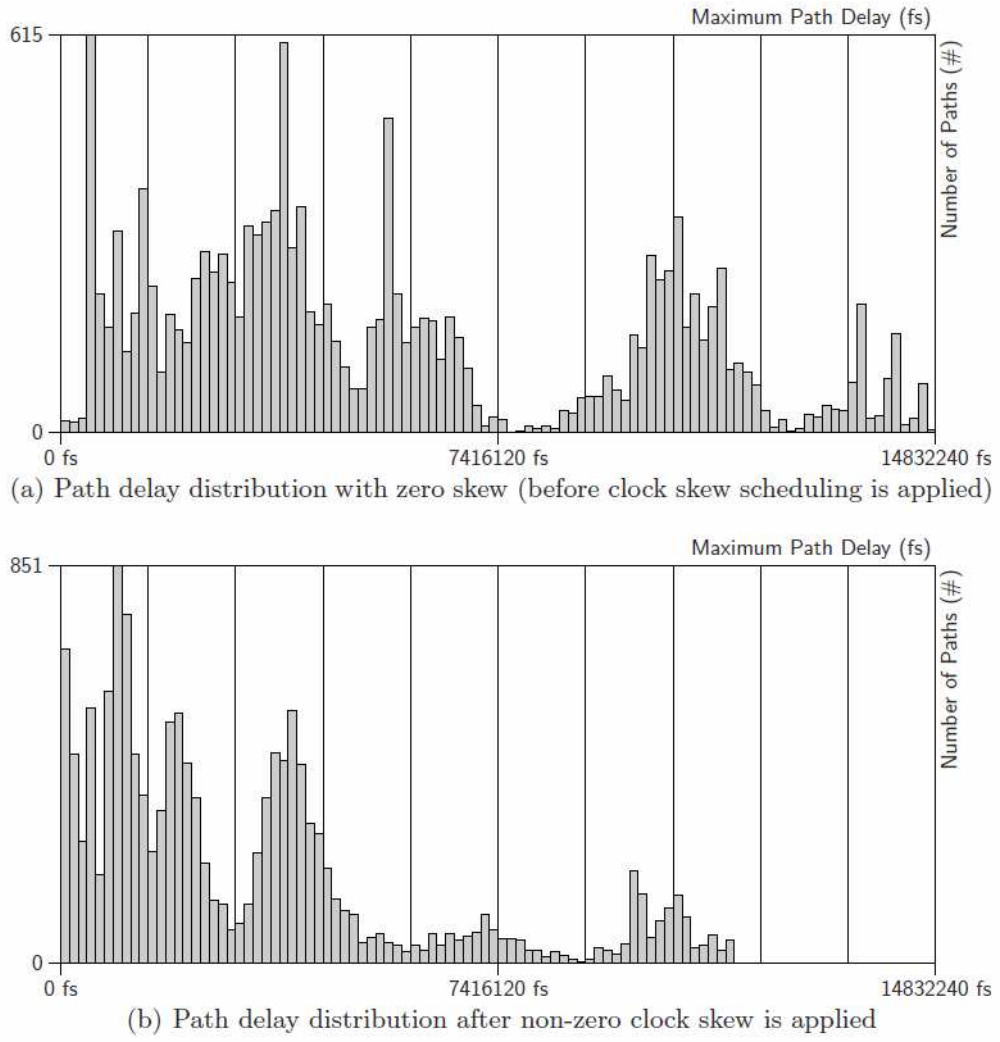

Figure 2.15: The application of clock skew scheduling to a commercial integrated circuit with 6,890 registers (note that the time scale is in femtoseconds) [12]

as voltage, leakage or delay measurements for each device. In more sophisticated techniques even logic functions can be moved to other processing elements on the device to meet the performance requirements.

2.4.2.1 Dynamic Voltage and Frequency Scaling (DVFS)

Recalling from the equations Eq. 2.1 and Eq. 2.2, it is obvious that to improve the delay of a given standard cell (i.e. reducing t_{pLH} and t_{pHL}), one could increase supply voltage (V_{dd}), reduce threshold voltage (V_{th}), increase transistor's gain factors ($\beta_n|p$), or reduce the load capacitance (C_I) [73].

$$t_{pLH} = \frac{C_I VDD}{I_p} = \frac{C_I VDD}{\beta_p (VDD - |V_{tp}|)^2} \tag{2.1}$$

$$t_{pHL} = \frac{C_I VDD}{I_n} = \frac{C_I VDD}{\beta_n (VDD - |V_{tn}|)^2} \tag{2.2}$$

Among these, the most feasible parameters for tuning are supply and threshold voltages and most of the adaptive/dynamic techniques are based on tuning these parameters. Threshold voltage can be changed by body biasing. If the body-source junction is reverse biased (V_{body}<0 for NMOS, V_{body}>VCC for PMOS), the magnitude of the threshold voltage increases. If the body-source junction is forward biased (V_{body}>0 for NMOS, V_{body}<VCC for PMOS), the magnitude of the threshold voltage reduces.

There is always a back-and-forth relationship between V_{dd} and V_{th} in DVFS. Low V_{th} leads to higher leakage power and lower dynamic power. With lower V_{th}, the target clock frequency can be met at lower V_{dd} while the leakage power will be higher. By [13] [83].

As the threshold voltage is increased, the supply voltage required to maintain the operating frequency is also increased and hence dynamic power increases. At the same time, the increasing threshold voltage results in a lower leakage power. For a given integrated circuit, there is an optimum point where the power is minimized as depicted in figure 2.16. This is the point where the increase in dynamic power is offset by the decrease in leakage power [13] [83].

An important voltage scaling technique is designing circuits with multiple supply voltages. Traditionally, synchronous chips are designed to work using a single voltage supply. Owing to the fact that the number of timing critical paths in a chip is usually a small portion of all of the paths, most of the paths can actually operate with lower voltages. In other words, most of the paths in a circuit, have wide positive setup timing slacks. So they usually arrive much earlier to the downstream logic comparing to critical paths and they have to wait until the data

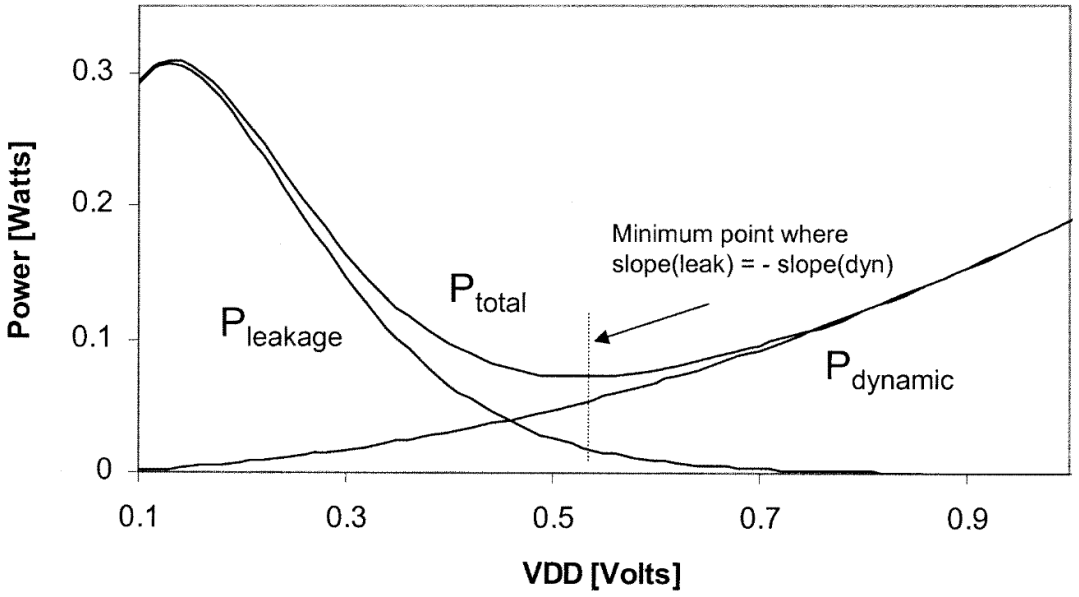

Figure 2.16: Dynamic and sub-threshold leakage power components for a fixed operating frequency in 140nm. As V_{DD} increases, V_{body} is adjusted to maintain the operating speed [13].

and the signals of the critical paths have also arrived and are valid. For such non-critical paths, although they are fast, but they cannot increase the performance of the circuit as the circuit speed is limited by the speed of the critical paths. Therefore, by operating such non-critical paths at lower voltages (up to the point that they will not become too slow to result in more critical paths), we can save energy.

By selectively decreasing the supply voltage for the gates or blocks which are not on recognised critical paths, and simultaneously maintaining or increasing the supply voltage for the gates on the critical paths we can meet the target clock frequency while optimizing power consumption, hence saving energy [80]. The problem is that scaling the supply voltage of all of the gates along a non-critical delay path, may not always be feasible due to local timing constraints. Moreover specialized voltage-level converter circuits are required to interface the circuits operating at different supply voltages in a multiple supply voltage circuit that will add overhead to the circuit.

Among the multiple supply voltage techniques, the clustered voltage scaling (CVS) technique, proposed in [80] [84], minimizes the number of voltage-level converters in a multiple supply voltage circuit. In the CVS technique, the supply voltages are assigned such that no low supply voltage gate drives a high supply voltage gate. This technique is applicable at gate-level; however impacts of DSM on critical

paths selection and delay uncertainties are not observable at gate-level, therefore the efficiency of this technique at UDSM is dubious.

Static leakage power can also be reduced using multi-threshold techniques along with power gating techniques and by scaling up the threshold voltage and scaling down the supply voltage. To disconnect power supply - V_{DD} or V_{SS} rails - from the blocks that are in idle mode, power-gating switches (transistors with high V_{th}) can be used [39]. The drawback here is that this technique is only applicable when the targeted circuit parts are idle for a considerable number of clock cycles. Also it should be noted that disconnecting power rails results in loss of data unless some memory elements are added to the circuits to save the states which adds to the overheads. Moreover, from a design perspective, this technique cannot be automatically added at gate-level and information and signals from the architectural-level are required to make the idle mode entrance or exiting decisions.

Another useful technique to deal with delays is frequency scaling (i.e. adjusting the global frequency and also adjusting the local clock frequency of the blocks on a chip). The optimal performance and delay tolerance might be achieved using a combination of supply/threshold voltage scaling and frequency scaling.

Tschanz *et al.* [65] explore schemes to dynamically adapt various combinations of frequency, supply and body bias to changes in temperature, supply noises, and transistor aging, to maximize average performance or improve energy efficiency. Their clocking scheme is comprised of three PLLs which operate on different frequencies and they are independent of one another. The scheme also includes a multiplexer to choose the appropriate clock source from these three PLLs. Various algorithms have been proposed on how and when to switch among the clock sources. In one simple method, the clock controller chooses one of the independently running PLLs and in a more complex method, PLL frequencies are chosen in a way that the circuit is running on one PLL, while one of the other two PLLs is locked to a lower frequency and the other one is locked to a higher frequency taking the currently under use PLL as the reference frequency. In the case of switching to one of the lower/higher frequency PLLs, the other two PLLs will re-lock to the 'new' lower and higher PLLs and this 3-stage procedure will be repeated at run-time continuously. However, the penalty with these on-line frequency scaling techniques is the PLL re-lock time due to PLL reconfiguration procedure, especially when PLL supply voltage and core circuit supply voltage are shared and changing dynamically.

Among various adaptive and dynamic voltage and frequency scaling techniques, Pass-Fail techniques are more pragmatic. Pass-Fail techniques scale the parameters until the system reaches the point of failure and then tune them to avoid timing errors. This significantly increases performance and reduce power consumption [13] but the problem is that adjusting these parameters is not instantaneous and may take multiple clock cycles.

On the other hand, although adaptive and dynamic techniques are beneficial in dealing with variations at run-time, yet their effectiveness is bounded by the well-known "Worst-Case" conditions or design margins. While these worst case conditions give a high level of confidence, in reality the worst case conditions seldom occur, and if a system is capable of detecting and correcting the errors on the occurrence of such worst case conditions, then more aggressive scaling can be applied. In other words, the design margins can be reduced and instead of worst case conditions, system can be design based on "Better-Than-Worst-Case (BTWC)" conditions rather than "Worst Case" with more relaxed design margins [85] [86].

The BTWC approach is generic and can be applied at different abstraction levels of the design. At circuit and architectural level an approach called Razor [14] has been proposed which is based on dynamic detection and correction of circuit timing errors. The key idea behind Razor is to automatically adapt the supply voltage based on the feedback we get from the timing error rate. This happens at run time, so there is a specific mechanism to monitor the timing errors. In theory, Razor-style architectures, can dramatically relax design-time margining and timing constraints.

In Razor flip-flop as shown in Fig. 2.17, the logic values at the down stream logic are sampled twice in every clock cycle. The first sample (main flip-flop) is taken using the very fast clock frequency (the normal operating speed) and the second sample (shadow flip-flip or latch) is taken using a delayed clock. A comparator compares the values of the main flip-flop and the shadow latch. When there is a timing error, the values of the main flip-flop and the shadow latch do not match. This will flag an error and consequently the pipeline will be flushed from that stage and the failed instruction will be redone [13] [14]. Here, the assumption is that during normal operation, the delay and power overhead caused by the error detection and correction phase is minimal, otherwise the power-performance efficiency of this method is not significant. Moreover, no shadow flip-flop should be placed on "short paths" as this may cause the shadow flip-flop to catch the

next data wave. In other words, checking the set-up and hold-time constraints becomes more complicated.

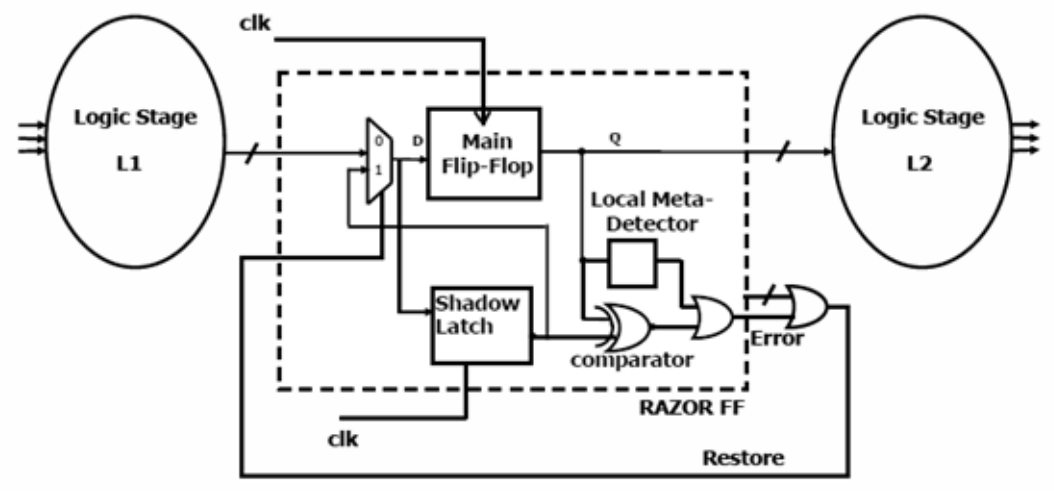

Figure 2.17: Razor Architecture [14]

At algorithmic level, Digital Signal Processors are good candidates for system design with reduced design margins. These techniques are known as Algorithmic Noise Tolerance (ANT). An example of ANT is shown in Fig. 2.18. In such systems an estimator (which is much simpler than the main processing block) that approximates the outcome of the complex computation can be added to the system and operates in parallel with the main computation block [15].

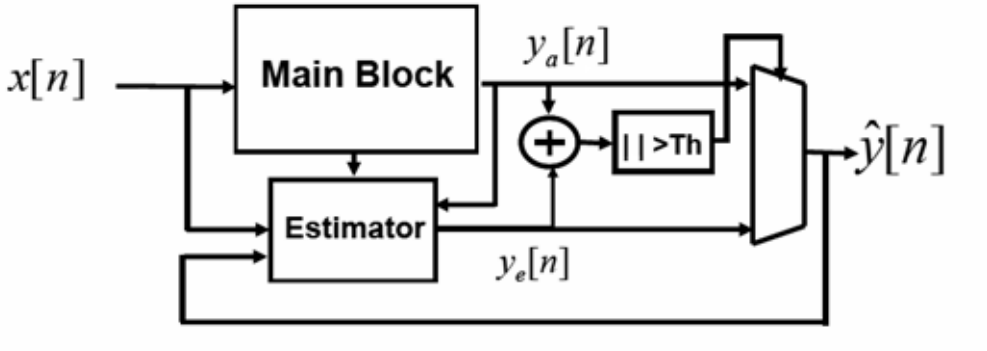

Figure 2.18: ANT Architecture [15]

In this case, assume that the main block parameters such as supply voltage have been aggressively scaled (with respect to the BTWC concept), so errors may start to occur and the main block faces values far from the predictions produced by the estimator. Therefore an error condition is flagged (detection), upon which the faulty outcome is replaced by the estimation (correction). This obviously deteriorates the quality of the processor - due to reduction in signal-to-noise ratio

(SNR), but if the estimator is good enough, the increase in the noise level is masked by the noise of the input signal or by the added noise of the signal-processing algorithm, and hence barely matters. Also it must be mentioned that "small errors" (errors that only effect the least-significant bits [LSBs]) may go undetected. For this scheme to work, clearly it is essential that the estimator does not make any errors itself. This requires that the "Estimate Module" be run at the nominal voltage. Since it is supposed to be a simple function, its energy overhead is small [15].

Based on this idea, an architecture for motion estimation is proposed in [87] which over scales the supply voltage at the expense of timing errors which are then corrected using the technique mentioned above. The main block in this architecture uses the MSAD (main sum of absolute differences) algorithm, whereas the estimator uses a simpler version called ISR-SAD (Input sub-sampled replica of sum of absolute differences) with reduced precision and reduced sampling rate compared to the MSAD. However, utilizing this technique, it is obvious that the estimator must not make any error by itself which necessitates that the estimator block must work at nominal voltage.

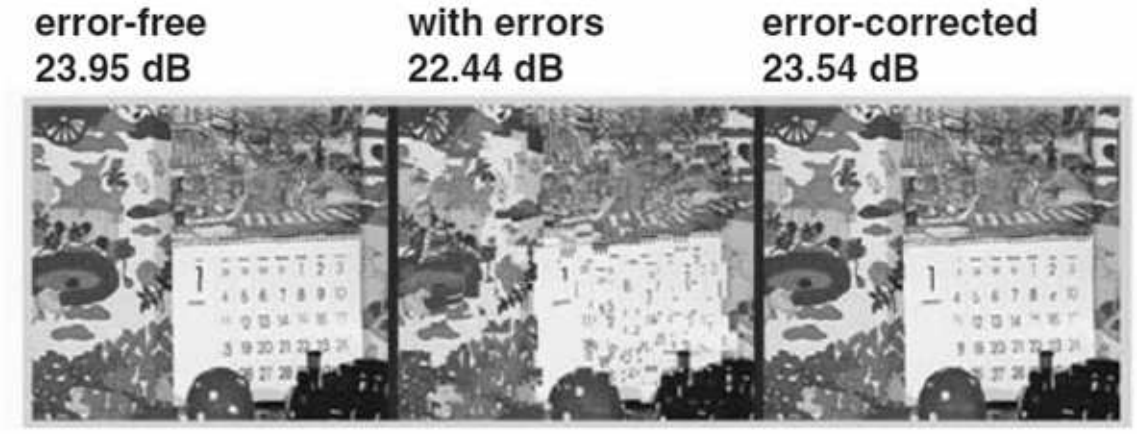

Figure 2.19: An example result of Motion Estimation with ANT error correction [15]

In [88], a Variation-Aware DVFS scheme is proposed for chip-multiprocessors. Chips are divided into Voltage/Frequency islands and two different hardware controllers considered for applying DVFS, the simple threshold-based controller and a greedy controller; The latter has higher overhead and higher power reduction capability. In this scheme, first the intra-die variation is calculated as a single effective parameter and feed to the system at the test time to determine the proper operating point to minimize power/throughput and choose the right voltage/frequency pairs for each island.

2.4.2.2 Considerations

Despite the fact that dynamic voltage and frequency scaling techniques generally offer quadratic reduction in power while maintaining the performance until recently, at UDSM scales, these techniques have not been as as fruitful and practical as before. For instance, The Adaptive Body Biasing (ABB) technique for below 65-nm technologies are not effective at all as illustrated in Fig. 2.20.

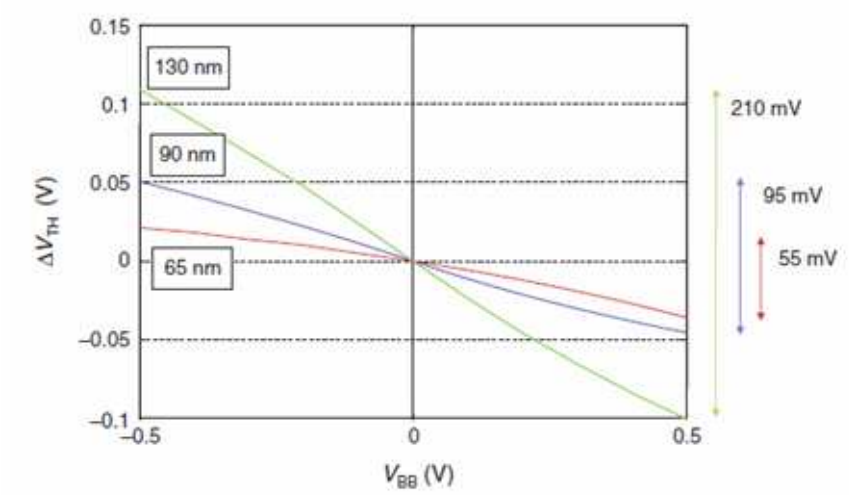

Figure 2.20: Comparison of body bias effectiveness in three technologies [15]

As it is shown, for 65 nm technologies, the broad body-biasing of 1-volt range, gives a narrow range of 55 mV for the threshold voltage and this will be even worse for 45nm or below.

Moreover, as predicted in ITRS reports, the next generation of circuits and power supplies (2007-2014) must operate at 0.9 V to 0.6 V with a dynamic range of 0.2 to 0.3 V. This level of voltage scaling denotes more susceptibility to noise and transient errors which result in more unreliability in the circuits. Therefore more power and area overhead may be required for error detection and correction which may overcome the power-performance savings using DVFS techniques [89] [90].

2.4.3 Soft Error Mitigation Techniques: Radiation Hardening By Design (RHBD)

Soft errors used to be primarily an issue for space and avionics applications but for UDSM technologies, it has also become a reliability problem at sea level. This has resulted in increased SEUs in state holding cells such as latches, and flip-flops caused by cosmic neutrons and alpha particles. Also, combinatorial logic is not more immune to radiation effects anymore and the number of particle-induced SETs in logic has been increased, leading to more SETS getting captured by the downstream flip-flops and latches.

The usual mechanism to overcome the soft error issues in memories has been the utilization of Error Correcting Codes (ECC) which imposes rather moderate overheads in terms of performance, area and power. This overhead penalty is tolerable is some designs and not acceptable in performance/area critical designs. Depending on the application of the design and also the design size, ECC techniques might be too costly and not feasible. In sequential cells and the logic, duplication, triplication, comparing and majority voting are some of the most well known techniques to overcome SEUs and SETs. However such techniques bring their own expensive performance, area and power overheads, making such techniques not suitable for all applications.

As stated earlier in this chapter, logic is affected by SEU and SET related soft errors. SEUs occur when an ionizing particle striking a sensitive node of a flip-flop or a latch cell flips the state of the cell. SET-related soft errors occur when a transient pulse, initiated by an ionizing particle striking a sensitive node of a logic gate, is propagated through the gates of the combinatorial logic and is captured by a sequential element such as a latch or a flip-flop.

Although SEUs are the most significant contributors to logic SER, SETs cannot be ignored as we move towards higher density chips. SETs will be discussed in chapter 3. If the logic SER of a given design exceeds the maximum allowable FIT due to SEUs, state-holding elements such as flip-flops or latches need to be protected to obtain an acceptable FIT figure. This can be achieved by replacing a selected group of conventional flip-flops and latches with hardened ones. However if such techniques are not sufficient to meet the required FIT figure (for example because of emerging SETs), then more comprehensive and perhaps sophisticated soft-error mitigation mechanisms must be utilized to make sure the circuit is resilient to SEUs and SETs. Such approach can be based on Hardware and Software using means

of redundancy. This can be time (temporal) redundancy or space redundancy or even a combination of both.

2.4.3.1 RHBD at device/Layout Level

At device/layout level, the simplest solution is increasing the charge needed for an SEU to occur which is known as the "critical charge". This can be achieved by increasing the capacitance in the sensitive nodes. The bigger the capacitance, the higher the immunity to SEUs with the drawback of imposing more power and area overhead [91]. In [16] [17], Enclosed Layout Transistors (ELT) has been proposed to eliminate the radiation-induced current between source and drain, hence avoiding the upset to happen as shows in Fig. 2.21 and Fig. 2.22. In these transistors the SEEs caused by radiation hits are prevented by cutting the current between the drain and the source of the transistor. This has been demonstrated to be very effective in CMOS processes of different technology nodes. However, due to challenges such as modelling the ELT transistors to compute W/L, the limitation in the W/L ratio that can be achieved and the lack of symmetry in the device, very few such radiation hardened cell libraries exist.

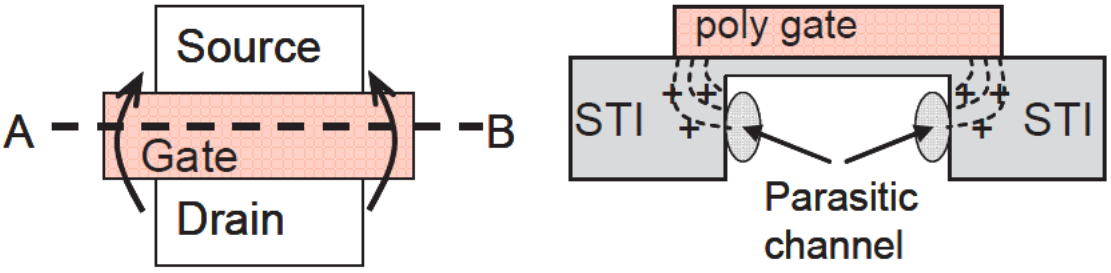

Figure 2.21: Top view of an open-layout NMOS transistor (left), and along its A-B line (right, view from the source or the drain electrode to the transistor channel) [16] [17]

In Fig. 2.21, the electric field is marked by the dashed line across the oxide of the STI at the transistor edge. This is the area where the STI and the polysilicon gate overlap each other. The + symbol, shows the positive charge that is caused by particle hits and is trapped in the STI. This trapped charge will improve the electric field up until the time that the P-doped inversion happens at the edges which will lead to opening two parasitic channels through which leakage current can flow from source to drain. In Fig. 2.22, the ends of the active areas and the beginning of the STI areas are shown by solid lines. Note that, n+ doping will

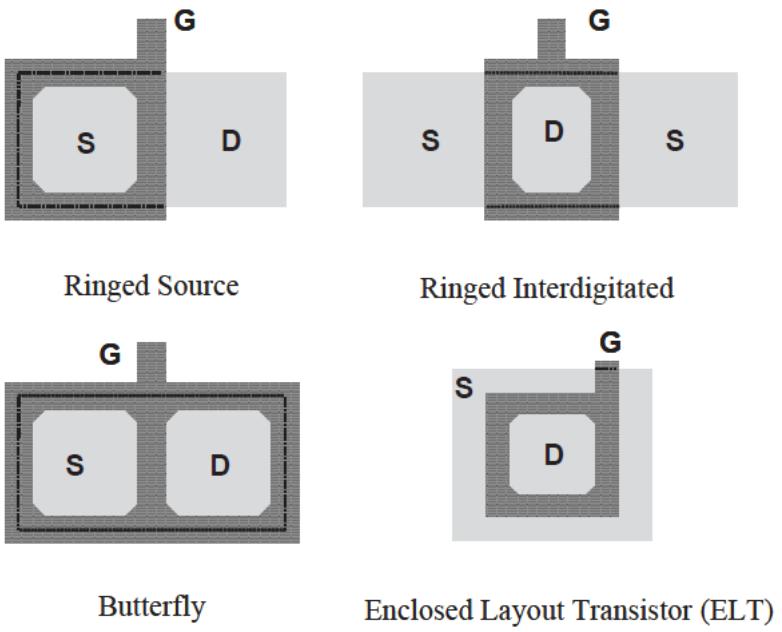

Figure 2.22: Transistor layout view for some of the possible NMOS designs eliminating the radiation-induced leakage current between source and drain [16] [17]

not be applied to the active area under the gate however this area is covered by a thin gate oxide that also covers the surroundings of both the drain and the source. This thin oxide layer is radiation tolerant [92].

This has been demonstrated to be very effective in CMOS processes of different technology nodes. However, there is a lot of challenges in the way of using such layouts. For instance, the limitations in modelling the ELT transistors and computing the W-L ratios and the inherent asymmetry in such devices, make it very difficult and expensive to be built for commercial applications. Hence there are not many of such hardened cells in existence.

2.4.3.2 RHBD at Transistor Level

Most of the proposed techniques at transistor level and above are based on various redundancies. The main feature of hardened storage cells (SRAM cells, latches, and flip-flops) is their capability in keeping their states when one of their internal nodes gets hit by a radiation particle that changes the state of that internal node. Various hardened storage cells have been proposed in the literature. We can categorize them in to three main types of hardened state holding cells [18].

The first type is based on increasing the critical charge by the addition of capacitors and/or resistors on the feedback loop of the state holding element. Passive elements such as poly-silicon can be utilized to create resistors in such cells. By using this technique, very strong resistors can be made at very low area overheads.

To increase the coupling capacitance can be done by using DRAM-style stacked capacitors which are placed on top of the state holding element. This will not incur any considerable area overheads [92]. The drawback of the above hardening approaches is that they require extra process steps that can have an impact on fabrication cost. In addition to the cost issue, such techniques may also have an impact on cell speed and power. These issues may reduce the interest of the above approaches for some commercial applications.

In the second category of designing hardened state holding, extra transistors are used. The radiation immunity of such cells are based on particular transistor sizing. The main challenge of using such cells is that while technology is going below UDSM scales, because of the the transistor sizing limits, the scaling of these do not track well with technology scaling. Also the addition of extra transistors will add to the area overheads.

The third and the last category of such hardened cells are Dual Interlock Cells also known as DICE as shown in Fig. 2.23. In such cells, radiation immunity to SEUs is obtained by duplication of the internal state holding nodes in the cell structure [93] [94]. The advantage of DICE is in its low performance penalty and the draw back is in its power and area overheads that can be twice the amount of a regular state holding cell. However since no specific scaling is needed for DICE cells, they are an attractive option for SEU hardening of UDSM technologies. Also Heavy Ion Tolerant (HIT) [95] cells fit in this category, in which the state-holding notes are duplicated to avoid the upsets. However for 90nm technologies and below, the SEU immunity achieved by these techniques is reported to be only 10 times better than standard cells. Moreover a particle strike on one of the state-holding nodes can cause the cell output to be wrong temporarily that can be fatal if it propagates to the next logic stage [18].

Also at transistor-level, techniques such as Code Word State Preserving (CWSP) have been proposed [18] [57]. Code Word State Preserving (CWSP) is based on replacing each transistor by a pair of transistors connected in series and driven by duplicated inputs. A CWSP cell, compares the values at its two inputs. When they are identical, the output value will be updated based on the input values.

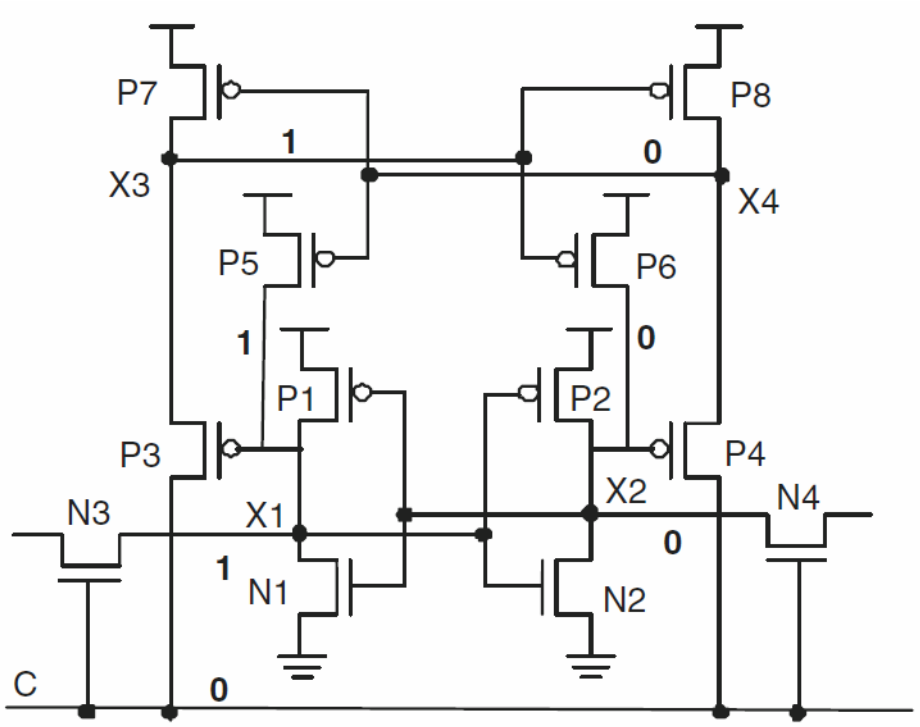

Figure 2.23: A hardened store holding cell - DICE [18]

Otherwise, when the inputs are not the same, the output value of the gate will be preserved. An example of such CWSP gates are depicted in Fig. 2.24.

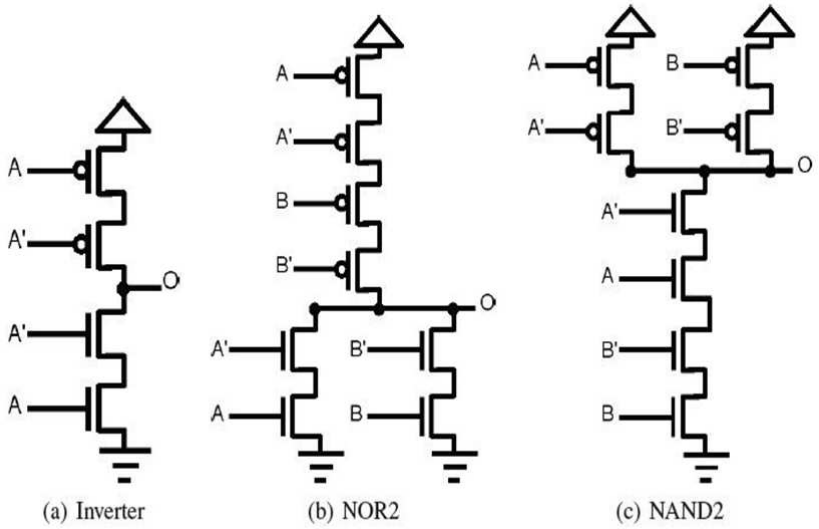

Figure 2.24: Code Word State Preserving (CWSP)

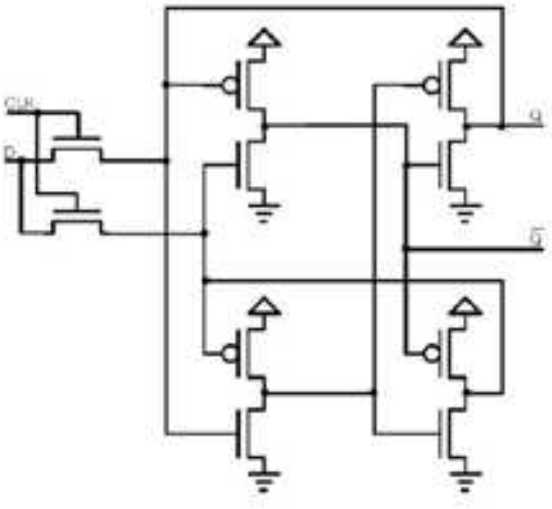

Figure 2.25: Dual Interlock Cell (DICE)

2.4.3.3 RHBD at Gate Level

Soft error mitigating techniques at gate level and higher levels of abstractions, are usually based on some sort of spatial redundancy, temporal redundancy or a combination of both. Among all of the proposed techniques at gate level, TMR is the most effective one and has been used extensively in the industry. At gate level, usually all the sequential elements in the design are triplicated with a majority voting circuit at the end. This imposes 3.2X overhead in terms of area and power compared to a non-TMR sequential cell. It is noteworthy to mention that the TMR concepts are also applicable at system level in which the whole core (sequential cells and combinational blocks) are triplicated as shown in Fig. 2.26; however again this adds more than 200% overhead to the whole area and power at system level. Fault tolerant techniques based on majority voting have fairly high fault coverage. In these systems, the output of the system is decided based on a voting mechanism among the sub-systems. However NMR (N Modular Redundancy) systems cannot necessarily handle all of the multiple-fault scenarios. For example a TMR system will fail if two sub-systems out of three are faulty at the same time.

2.4.3.4 RHBD at Register Transfer Level

The concept of TMR can be applied at Register Transfer Level (RTL) too. In [96], a method for an automatic insertion of radiation-hardened modules in designs at RTL is described. In their approach the VHDL RTL code is taken and the desired replicated blocks are added to design along with the required auxiliary

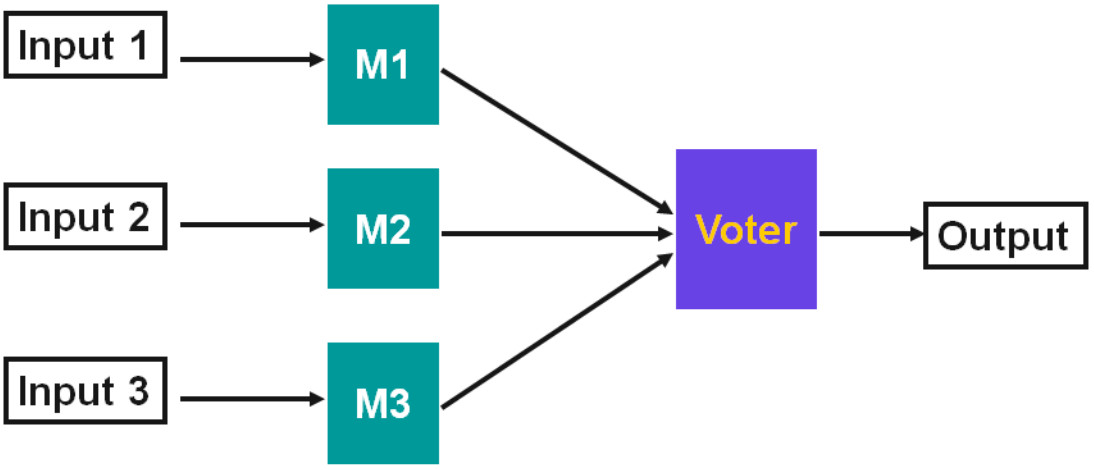

Figure 2.26: Triple modular redundancy

signals. This is done in two steps: 1) Target selection and replication, 2)Resolution function. However there is no commercial automatic RHBD at RTL tool available. In [97], an SEU error correction method is proposed in which the data-paths are duplicated and the outputs of every stage are monitored continuously. In the case of a mismatch at each stage, second computation is triggered on one of the two data path while the other data path continues processing the next input. Here the assumption is that neither of the computations requires error monitoring due to the probability of SEU occurrence on two consecutive iterations.

Another conventional technique is stand-by redundancy. In this technique, a diagnostic mechanism checks the outputs of the replicated sub-systems as shown in Fig.2.27. The fault coverage of this technique is the key element and the reliability of such a system is as good as its diagnostic mechanism. Any failure to detect faults, can lead to the system failure as the wrong output can be chosen. Stand-by redundancy systems are more suitable for environments in which permanent faults and multiple faults are the major concerns. Since detecting transient faults needs an on-the-fly and at-speed fault detection mechanism, stand-by redundancy systems are not very suitable for detecting and recovering from transient faults [98].

2.4.3.5 RHBD at Software Level

In the case that RHBD techniques are not applicable on hardware (because of architectural or technological limitations), software level is an interesting option. Various approaches have been proposed at software level like Computation Duplication [99], Procedure-level Duplication [100], Program-level Duplication [101] and

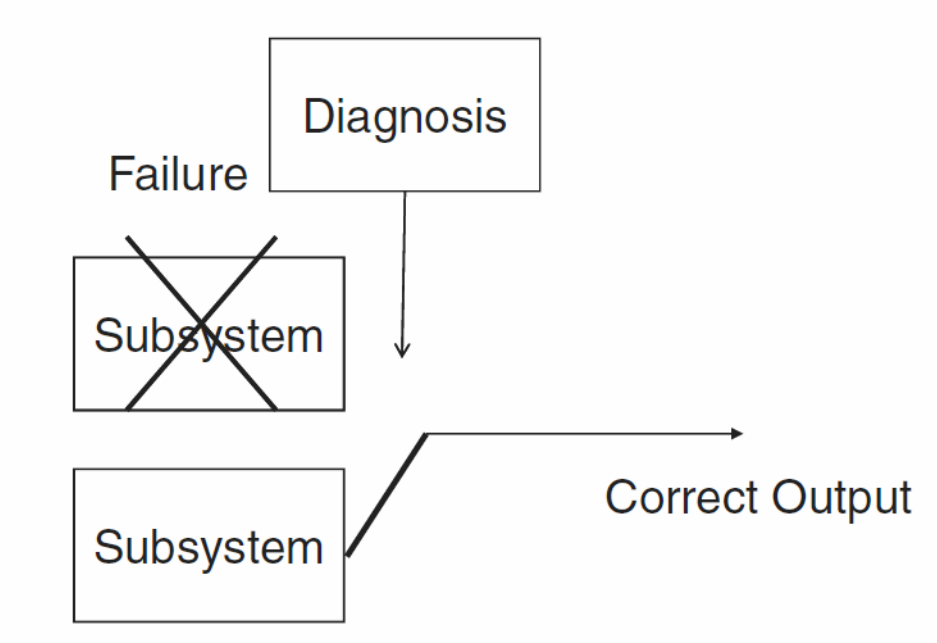

Figure 2.27: Stand-by redundancy

Redundant Multi-Threading (RMT) [102] [103]. In all of these approaches, the error detection & correction capabilities are obtained by virtually adding the Dual Modular Redundancy (DMR) or TMR schemes at different levels of granularity: instruction, instructions block, procedure, program, etc.

Applying RHBD techniques at each level of abstraction has its own advantages and drawbacks. There is a trade-off between the overhead and efficiency, and usually RHBD at higher levels of abstraction adds to the complexity of such techniques. Among all, Radiation Hardening at gate-level is the simplest and one of the most effective one, which is also supported by conventional EDA tools.

2.4.4 Dealing with the Reliability issues

A fundamental challenge in designing reliable systems is estimating whether a system will function properly in a predefined manner in a given environment for a given period of time. Providing this level of reliability of electronic systems out of intrinsically unreliable UDSM CMOS components is a major challenge [104]. For instance, Reliability requirements for computer systems that are used in military aircrafts, are typically in the range of 1 - 10^{-7} per mission, and the reliability requirements of 1 - 10^{-9} for a ten-hour flight are often expressed for mission-critical avionics systems [105].

Computer systems are designed to detect faults and be able to tolerate such faults by themselves. However this fault resilience is not 100% guaranteed and such systems are still vulnerable to failure. Therefore their reliability must be examined and we need to make sure that the fault tolerance requirement targets are met.

One issue is the complexity of analyzing and modelling the actual reliability of such fault-tolerant systems. Usually lifetime test is used as a measure of reliability. To determine the reliability of a highly-reliable design the following steps are taken:

1. Develop a mathematical model of the reliability of such system
2. Measure or approximate the parameters of the reliability model at elevated temperatures
3. Calculate the system reliability using the developed model and the specified model parameters

Obviously the precision of the estimated reliability of a given system solely depends on the accuracy of the model that has been used. Also due to the complexities of a highly fault-tolerant systems, deriving an accurate model that can comprehensively describe the behavior the system is a Herculean task. Such models should precisely consider all of the phases and the processes that lead to system failures along with the capabilities of the fault-tolerant system to operate in the presence of the faults and the broken parts.

Inherently the implementation of any fault-tolerance mechanism involves imposing additional overheads. In such mechanisms, redundancy has to be incorporated into the system with the aim of masking the faults. This will definitely increase both the cost and the development time. Moreover, any redundancy mechanism will impose some overheads in terms of power, performance, area on the system. Hence, there will always be a trade off between a suitable fault-tolerant technique and its inevitable overhead versus the fault coverage and the power, performance, area budget of the system. In other words, in the cost-benefit framework of a good fault-tolerant system, benefits - i.e. fault-tolerance and error recovery - should outweigh the costs which are the overheads and downtime of the system.

Hardware redundancy is perhaps the most commonly used method and can be employed in various forms. The two major forms are: Static redundancy in which fault-tolerance is achieved without actually detecting any faults. In Dynamic

redundancy such as "stand-by redundancy", a fault detection mechanism is built-in to the system which makes the system capable of recovering from the error. Although, in practice, methods known as Hybrid redundancy which exploit both static and dynamic techniques are typically used.

Also another sub-category of hardware redundancy is Reconfigurability to achieve higher levels of reliability which can enable real-time and compile-time reconfiguration with the aim of isolating faulty/defective units and reconfiguring at real-time to keep the system running. Such methods are very well known when it comes to using reconfigurable devices such as FPGAs. The effectiveness of any reconfiguration scheme is measured by two aspects:

- The probability that a redundant unit can replace a faulty unit.
- The amount of reconfiguration overhead involved.

Numerous schemes have been proposed for reconfiguration [106] [107] [108] [109]. However the difficulty and the complexity of using such reconfiguration methods by exploiting arrays of processing units in practice, is still a major challenge and it is out of the scope of this work.

A system with dynamic redundancy is comprised of several modules (usually identical) but only one of them is operating at a time. If a fault is detected in the current operating module, it will be cut out and one of the spare modules will replace it. Therefore in dynamic redundancy systems, continuously fault detection and recovery is taking place. The fault detection method can be based on periodic tests, self-checking circuits or watchdog timers.

Assuring that the design is evaluated properly and has met the defined dependability requirements is a significant challenge. Generally speaking, the evaluation methods for dependability can be classified into two main categories: quantitative methods and qualitative methods. As the name implies, the qualitative methods are usually subjective and such methods are used when certain factors and parameters related to the dependability of the system or the design cannot be quantified. As the quantitative methods deal with numerical analysis and they are extracted or represent certain dependability attributes of the system and each system can have different dependability parameters.

To model the reliability of a system Markov models are commonly used [110] [111]. For any given system, a Markov model is comprised of a list of the possible states

that system can go into plus the possible transition paths between those states, and the frequency of the parameters of those transitions. When analyzing the reliability of a system, the transitions are typically repairs of failures. Markov models can be represented graphically in a way that each state is shown as a bubble and the transitions can be depicted as arrows connecting the bubbles (i.e the states) as shown in Fig. 2.28. In Fig. 2.28, a single component that has just two states: healthy and failed.

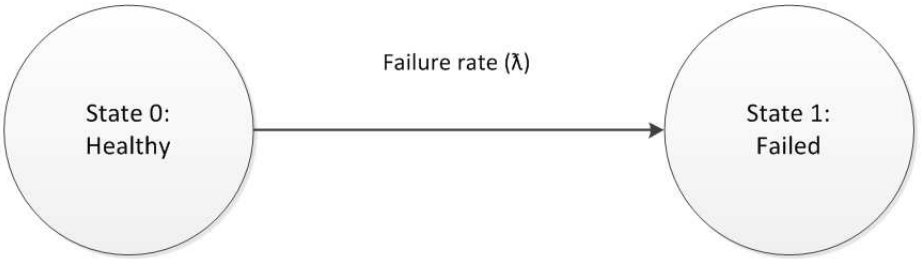

Figure 2.28: Markov model of a simple system

The Markov model representing such aforementioned systems are depicted in Fig. 2.29. State (1) represents the initial condition of working processors. The transition from state (1) to state (2) is labelled n*λ to represent the rate at which any one of the processing units fails (Initially in fault/defect free situations, n=3 for TMR and n=2 for Stand-by systems). In this model, the assumption is that, all of the processing units are identical, hence the failure rate λ is the same for every processing node. The system is in state (2) when one processor has failed. In TMR for example, the transition from state (2) to state (3) has the rate 2λ because only two working processors can fail and in Stand-by systems with one redundant component, the transition from state (2) to state (3) has the rate λ. For the Simplex system, state (2) is the death state while State (3) represents system failure for TMR due to the fact that in that state the majority of the processors in the system have failed. The same is true for Stand-by systems.

From the reliability analysis point of view, the failure distribution of electronics devices is considered to be exponential. This is specifically true for more mature products as it has been demonstrated that their failure rates follow the exponential distribution patten. Although for immature products and devices (i.e. new devices that have just been manufactured and have not been thoroughly tested before mass production) the failure rate is higher [110] [111]. The reliability issues can lead to transient faults, intermittent faults or permanent faults.

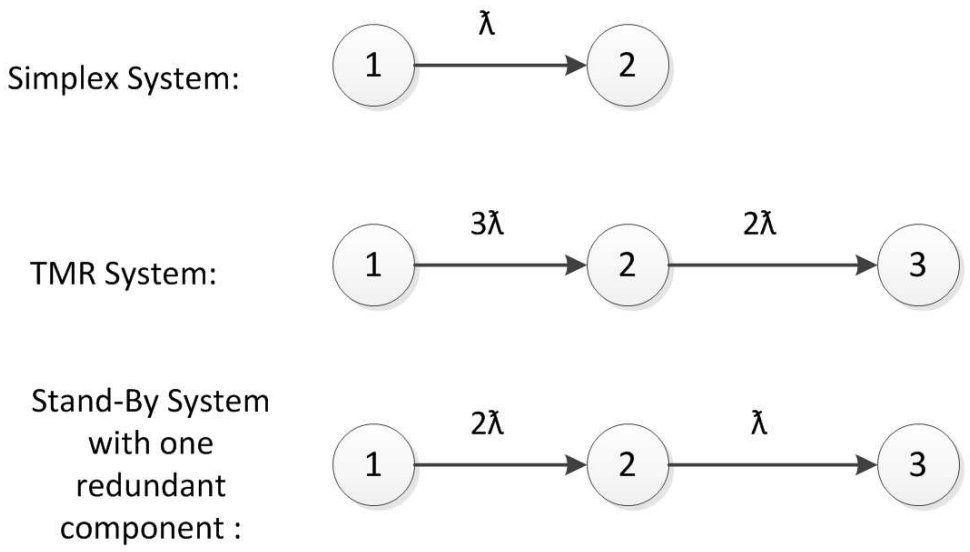

Figure 2.29: Markov models

It is generally accepted that under normal conditions, the failure rate of systems or individual components can be expressed as depicted in Fig. 2.30. A transient fault as the name implies is temporary (like soft errors) while a permanent fault is like a defect (hard errors). Intermittent faults can occur at regular intervals. The classic bathtub curve is usually used to demonstrate the potential permanent faults or hard errors. Hard errors can be the reason for both of the first phase or infant mortality and the second phase which is useful lifetime reliability as shown in Fig. 2.30 [20] [19].

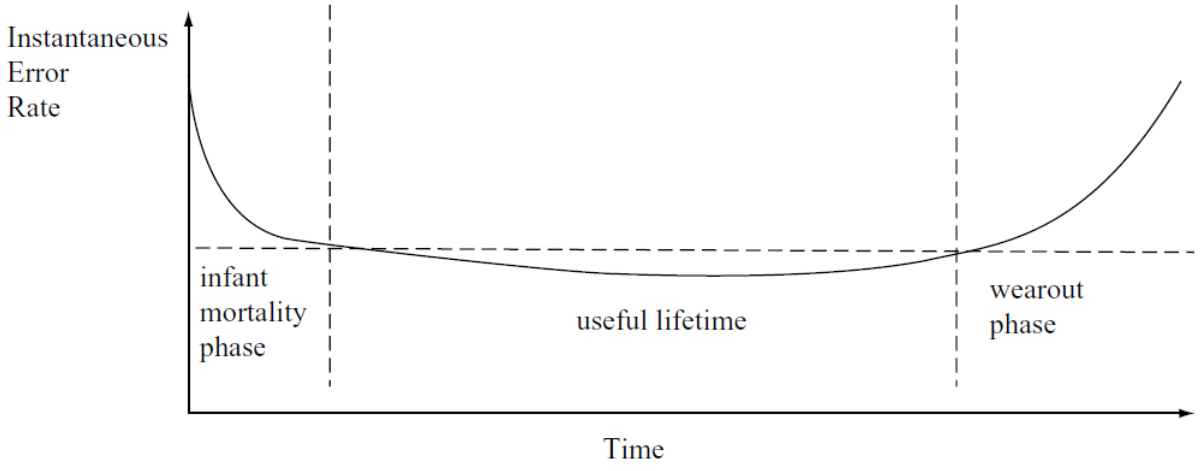

Figure 2.30: Bathtub curve showing the relationship between failure rate, infant mortality, useful lifetime, and wearout phase [19].

1. The first phase which starts at once at the beginning of the life span of a system or a device has a very high failure rate that decreases over time. This is the immature phase of a product which is usually known as Early failure phase or infant mortality. This phase is due to the existence of a small sample of the population that their defects cause very high failures in a short amount of time. It is also possible that the failure rate at this phase can fluctuate rather than having an continuous descending curve, the way it is depicted in the figure above.

2. In the second stage is known as the useful lifetime - the time interval between infant mortality and the wear out phase - in which the failure rate is usually almost constant. The failure rate in this phase usually follows Poisson distribution as the time interval is more or less fixed.

3. The third stage which is called the wear out phase is the period of time in which the failure rate increases rapidly and drastically. For our discussion, such failures are caused by ageing, BTI and other reliability issues.

In industry, a methodology knows as 'burn-in' is used to skip the infant mortality stage and reach the useful lifetime period faster. In burn-in tests, the chips are overclocked at elevated temperatures and the test vectors (which have very high activity rates) are applied to the chip. The aim of burn-in test is to make the weakest transistors fail rapidly [112]. By utilizing the burn-in technique, the chips that fail at the infant mortality stage (defective chips) can be spotted and removed from the product line, hence only the chips that reach the useful lifetime properly will become final products. Also by adding margins and setting technology parameters, the industry will try to increase the confidence level that the manufactures chips will at least survive for the defined minimum lifetime span [19].

Reliability is typically quantified as MTBF (Mean Time Between Failures) for repairable devices and MTTF (Mean Time To Failures) for non-repairable devices. In repairable systems, MTBF is the sum of the mean time of MTTFs of the device plus the MTTR (Mean time to repair/restore) as shown in Fig. 2.31 [20].

A system is assumed to function properly during most of its life-time. One way to determine if the level of faults or system malfunctions is within the acceptable range is by defining an availability factor. Availability (A) can be defined as [20] [21]:

$$A = \frac{Uptime}{Totaltime} = \frac{Uptime}{UpTime + DownTime} \tag{2.3}$$

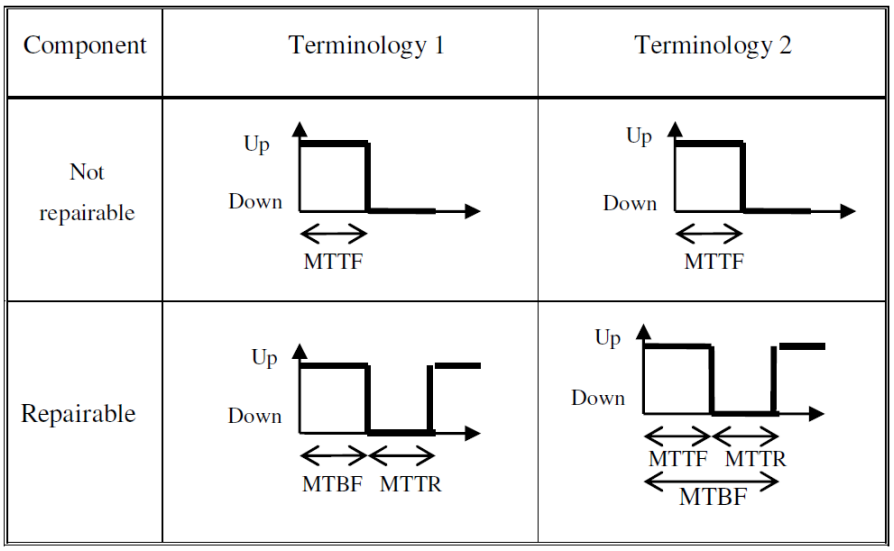

Figure 2.31: Definitions for MTBF [20] [21]

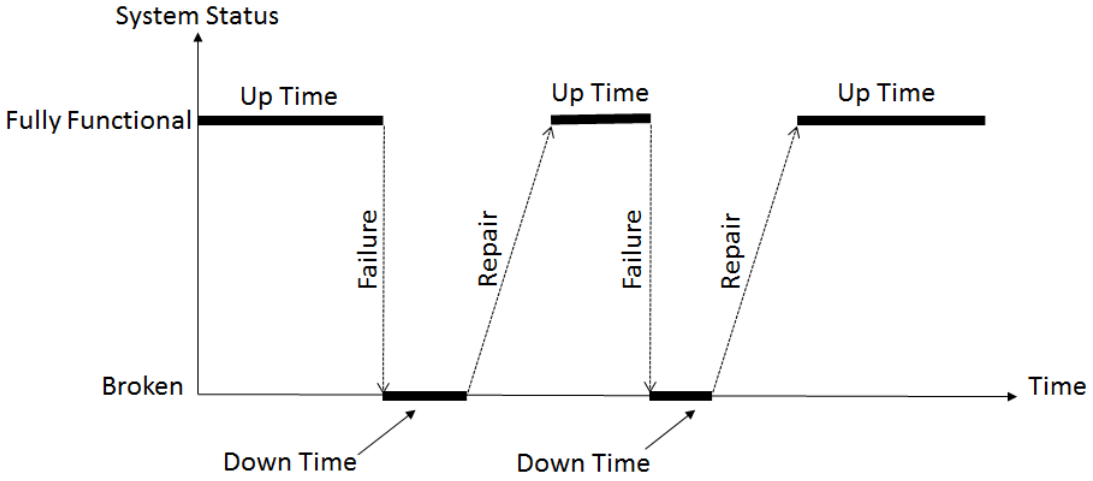

Figure 2.32: Different phases of a repairable system

To practically assess the availability factor of a system, the temporal elements should be replaced by other elements that represent the required functionality of the system. Depending on the situation and the desired purposes from the system, the availability factor should be defined with respect to effective 'Up time' (work time) and 'Down time' (repair/maintenance time) as shown in Fig. 2.32. This definition of availability is called 'inherent' availability [20] and is usually represented by:

$$A = \frac{MTTF}{MTTF + MTTR} = \frac{MTTF}{MTBF} \tag{2.4}$$

For simplification, two extra timing components are ignored here: waiting-for-maintenance time and recovery-time. This is because in an ideal world, these two timing components are zero. So to compute MTTR, only the maintenance time for correction is considered. But in practice maintenance and recovery time can be crucial. Particularly in safety-critical or medical applications such down-times can even lead to human casualties or even in non-stop computing systems in which losing even a few seconds of functionality or service can cause huge financial losses. Banks, reservation systems, servers or any infrastructures that deal with giving services to users are under this category. In such applications, occasional loss of services or disconnections is acceptable only if the system can restore quickly and provide the usual services to the end users with minimum downtime, rapid maintenance and low service delays.

Recently, the word "reliability" has been substituted by the term "dependability". Any methodology for designing, implementing and testing of the dependable systems must be able to identify the root causes of failures first. It should be able to predict the manifestations of such failures and eventually use the appropriate technology and techniques to deal with such failures at a releasable cost and tolerable overheads [20]

Preventing failures is the key factor in building dependable systems. To achieve this, is it crucial to understand the roots of the failures and the events that lead to such failures. A lot of failures can be temporarily inactive and latent for a specific period of time until they manifest themselves. In other words, a failure is in effect the external observation of an error inside the system. So errors are hidden until they become active which will lead to failures (externally observable). The failure themselves could also be present but not externally visible to the user. In other words, it also depends on the scope of observing the system. A failure in a sub-component might be totally hidden to the top-level without corrupting the desired system outputs. To make it even more complex, it is known that, similar failures can be rooted in different errors while the same errors can be responsible for different failures.

2.5 Concluding Remarks

In this chapter the first objective "To investigate the impacts of variation and reliability issues on UDSM CMOS circuits from a design perspective" has been addressed. This chapter provides a survey of various UDSM impacts on circuits and devices, reviewing the ongoing research and providing a summary of the state-of-the-art techniques to mitigate the UDSM impacts mainly the impacts of variation, soft errors and the reliability issues.

At 65nm, variation in transistor channel lengths and V_{th} in terms of standard deviation σ, has reached to 10% and the trend for deeper sub-micron technologies show that this is increasing. From a hardware designers' point of view, inherently the implementation of any of such fault-tolerance mechanism involves imposing additional overheads. Furthermore, redundancy will certainly have impacts on performance, power dissipation, weight, and size of the system. TMR as the most conventional fault tolerant technique imposes more than 200% overhead while its reliability is close to 100%, provided that only one out of three modules becomes faulty at any time, and of course if the majority voter is 100% robust. Thus a good fault tolerant design is a trade-off between the level of dependencies provided and the amount of redundancies used. or in other words, a good design is a trade-off between the cost of incorporating fault tolerance and the cost of errors that, includes losses due to downtime and the cost of erroneous results.

In the next chapter we investigate the timing vulnerability of UDSM circuits. The focus will be on the timing vulnerability to Soft errors and particularly Single-Event-Transients (SETs).

Chapter 3

Soft Errors and Timing Vulnerability

An mentioned, soft errors are a significant reliability issue for Ultra-Deep-Sub-Micron (UDSM) CMOS circuits. Therefore, an accurate assessment of the Soft-Error-Rate (SER) is crucial. In this part, we argue that the conventional definitions for the Window of Vulnerability (WOV) are too conservative and hence under-estimate the risk. We propose a new method for determining the timing factors and WOV for the sequential elements from the susceptibility perspective rather than the conventional performance perspective. Our methodology leads to a more realistic definition of the WOV for SER computation.

3.1 Introduction

As explained in chapter 2, section 2.2.2, because of decreasing circuit capacitance and increasing circuit speed, the SETs are becoming more important with the scaling of the technology [113] [55]. In recent years, despite reductions in the gate oxide thickness and increases in doping densities, which generally mitigate the susceptibility to soft errors, the reduced device dimensions and accompanying technological changes have resulted in increased sensitivity to transient radiation effects and particle hits [114] [115].

The impact of direct particle strikes on memory elements (SEUs) is well studied and various radiation-hardening techniques have been proposed to decrease the SER in memory blocks. However, the role of combinational elements along with

latches and flip-flops in determining the SER has not been investigated comprehensively. SER is the cumulative result of transient events on sequential elements and combinational parts in a circuit. This is quite different from SEUs on memory blocks such as SRAMs. The transient pulse caused by a particle strike can be captured by the sequential elements depending on the existence of an active path from the struck node to the storage element, the arrival time and the width of the transient pulse at the storage element input [50] [51].

The transient pulse can be captured by a flip-flop a t the down stream logic and cause an error if it is not masked by any of logic, electrical or temporal masking phenomena. Of these three masking effects, temporal masking is of greatest interest, since it plays a significant role in determining the SER [113]. In this chapter, we first take a survey of the conventional definitions for the Window of Vulnerability (WOV), then we propose our method in determining the WOV from the susceptibility perspective and at last we apply our method to determine the WOV of sequential cells for 130nm, 90nm and 45nm technologies.

3.2 Window of Vulnerability

A flip-flop is susceptible to capturing a spurious transient pulse, if it occurs inside the flip-flop's latching window, aperture window or window of vulnerability. The WOV is the basic factor in determining the effectiveness of temporal masking. To date, the WOV has been defined as the sum of the Setup time and Hold time constraints of the flip-flop [51] [116] [117].

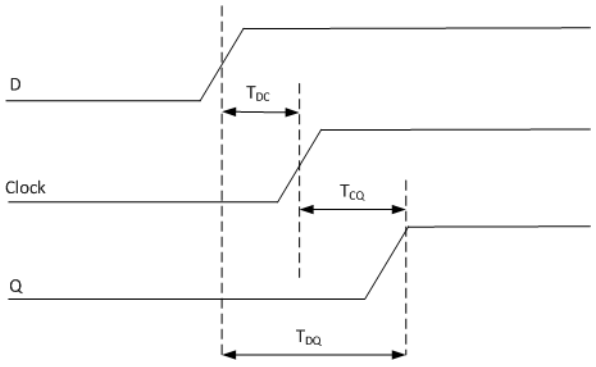

Figure 3.1: Flip-Flop Timing

As shown in Fig. 3.1, a common approach to characterizing the setup (hold) time is to consider the setup (hold) time with respect to the CLK-to-Q delay (T_{CQ}), while keeping a fixed value for the hold (setup) time [118] [119] [120]. According to

[120] [118], the output of the flip-flop falls into three regions: Stable, Metastable and Failure. In other words, depending on the size of the WOV and the width of the input pulse, the flip-flop can either latch the input data properly (stable region), become metastable, or fail to latch the input data as shown in Fig. 3.2.

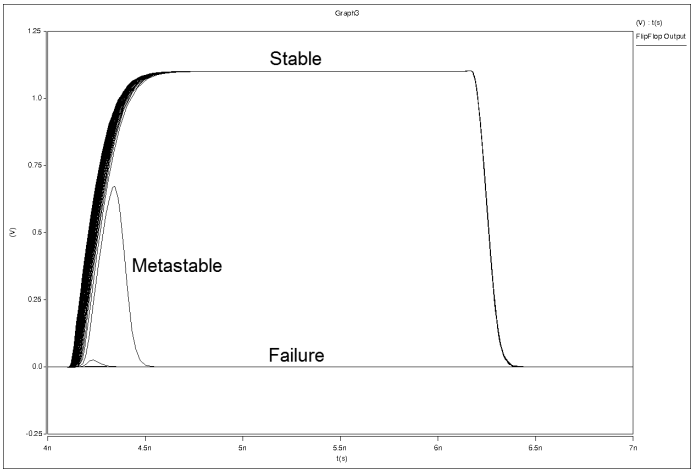

Figure 3.2: 45nm Technology - SPICE simulation of Flip-Flop output using Nangate 45nm SPICE models: When the input pulse is '1' for one clock cycle with varying input pulse width: Stable, Metastable and Failure regions.

Moreover, the chance of the flip-flop falling into the metastable region due to hold time violations is higher than for setup time violations, Fig. 3.3.

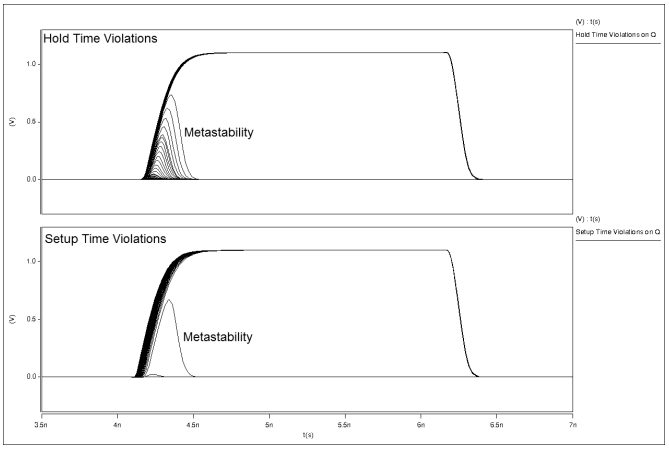

Figure 3.3: 45nm Technology - SPICE simulation of Flip-Flop output using Nangate 45nm SPICE models: Chances of metastability due to Hold time violations and Setup time violations.

The setup and hold times are traditionally calculated for the best performance. The setup time, T_{setup}, is usually defined as the D-to-Clock delay (T_{DC}) at which the *Minimum* D-to-Q delay (T_{DQ}) occurs, as depicted in Fig. 3.1, [119] [121]. The setup time, T_{setup}, is the time that input D must fall or rise before the clock

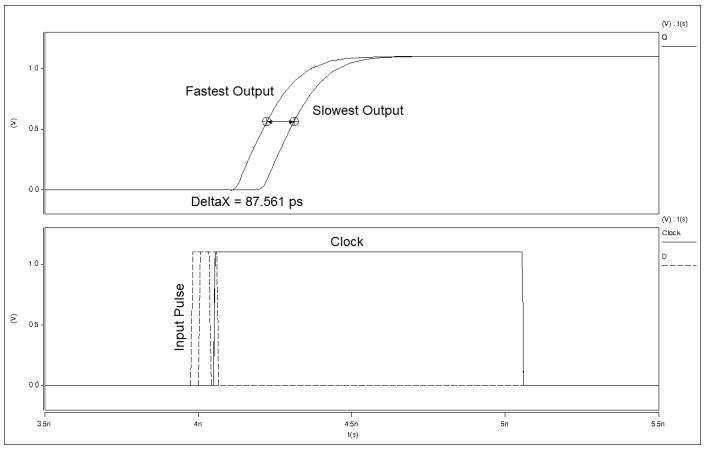

Figure 3.4: Fastest output vs Slowest output depending on the input pulse width and the pulse arrival time - 45nm technology

edge so that the data is properly captured with *the least possible* T_{DQ}. Industry standard EDA tools utilize the same procedure for sequential cell characterization. For instance, to perform cell characterization, a series of pass/fail simulations are run on the sequential cells to determine the setup and hold times for the minimum output delay [122]. The values of setup and hold times are chosen in a conservative manner to guarantee the best performance; the flip-flop might still capture its input properly with a longer Clock-to-Q delay (T_{CQ}) even if the data changes in less than the defined setup and hold times as shown in Fig. 3.4.

In [51], the setup time is defined as the D-to-Clock offset (T_{DC}) that corresponds to a 10% increase in the Clock-to-Q delay (T_{CQ}). Consequently, the conventional Window of Vulnerability (WOV) is defined as the sum of the setup and hold time windows around the clock edge during which the input data must not change. Using this definition, any pulse with a width equal to or greater than T_{DC} + $10\%T_{CQ}$ occurring around the clock edge, can be captured and any pulse narrower than this value will be masked and filtered out.

Our transistor-level simulations show that pulses which are much narrower than this conventional definition of WOV can, in fact, be captured and cause an error. In our SPICE simulation setup, we used a chain of inverters connected to a flop-flop and we injected SET pulses with various pulse widths and measured the flip-flop output. The transistor models were from Nangate 45nm cell library. For instance, considering the definition given in [51] and using the timing information for the 45nm cell library, the width of the WOV must be equal to 80 ps and any pulse narrower than 80 ps would not be captured, but for example a pulse with a width of 48 ps is captured by the flip-flop and will cause an error.

To the best of our knowledge, all the proposed definitions for the WOV and temporal masking effectiveness in the literature assume the above. In the next part, first we propose a new method to determine the WOV for the flip-flops from the susceptibility perspective and then we discuss the timing vulnerability and temporal masking effectiveness for Ultra-Deep-Sub-Micron CMOS technologies.

3.3 Methodology

From the susceptibility perspective, we define the WOV as the region around the clock edge where the narrowest input pulse can be properly captured and the flip-flop output stays in the stable region disregarding the minimum D-to-Output delay (T_{DQ}). To define this region, two points around the clock edge are determined:

1. The point (usually) before the clock edge, such that any input pulse starting after this time will not be captured properly by the flip-flop, no matter how wide the input pulse is. This is the point where any later-starting pulse will result in the flip-flop output falling into the metastable or failure regions.

2. The point (usually) after the clock edge, such that any input pulse ending before this time will not be captured properly by the flip-flop, no matter how wide the input pulse is. This is the point where any pulse ending earlier than this will result in the flip-flop output falling into the metastable or failure regions.

From simulations, a pulse with a width of the time between these two points will not be captured, but we observe that a pulse of approximately twice this width has sufficient energy to change the flip-flop state and hence will be captured, Fig. 3.5.

3.4 Results

Using Spice, we applied our methodology to find the narrowest capturable pulse width at 130nm, 90nm, 65nm and 45nm technologies. We used the fastest slew rate for the input pulse and the clock signal as specified in the timing library file ranging from 5ps to 450ps at 45nm. For the 130nm technology the minimum

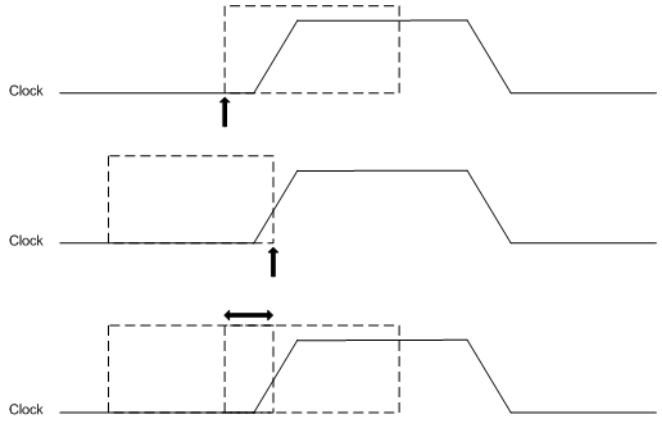

Figure 3.5: Defining the Window of Vulnerability

Table 3.1: An example of determining the minimum capturable pulse width - SPICE simulations using Nangate 45m technology library

Determining the minimum capturable pulse width		
Point 1	Input Pulse Start time(ns)	FF Output
	4.0213	Failure
	4.0212	Metastable
	4.0211	*Stable*
	4.0210	Stable
Point 2	Input Pulse End time(ns)	FF Output
	4.0376	Metastable
	4.0377	Metastable
	4.0378	*Stable*
	4.0379	Stable
Common Region	0.0167	-
Min Pulse Width	0.0334	Stable

captured pulse width is 65 ps. For the 90nm technology, the narrowest capturable pulse is observed to be 56 ps, and about 44 ps for 65nm and at last the narrowest capturable pulse width at 45nm is approximately 34 ps - Table 3.1. Note that to consider the worst case scenario, we used the smallest cells, so by up-sizing the cells, one could get better WOV immunity. This suggests that transient pulses with the widths equal to or greater than these values can be potentially captured by the flip-flop if they reach the flip-flop during the clock transition. These values are much less than the defined setup/hold time values in the cell library data-sheets, owing to the fact that the timing factors and minimum input pulse width for the sequential cells are characterized for the best performance rather than the susceptibility to SETs.

The results also suggest that the immunity against very narrow pulses and single-event-transients decreases with the technology scaling to UDSM as depicted in Fig. 3.6. It is noteworthy to mention the similarity between our method and previous methods for determining the metastability decay constant τ of the sequential cells [123] i.e. the amount of time that a flip-flop stays in the metastable region. The values of τ are calculated from experimental data obtained by uniformly varying the separation between the clock and event input timing. Since the obtained values are based on the experiments and observations and various factors such as process technology, temperature, supply voltage, clock rise time, etc influence the obtained values, thus formal definition of the WOV with agreeable precision is not achievable.

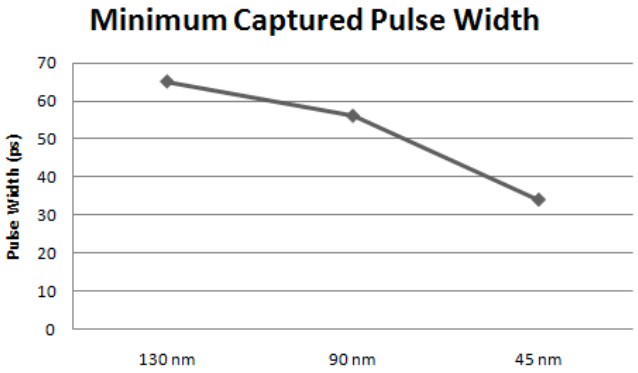

Figure 3.6: Minimum Captured Pulse Width by the Flip-Flops at three different technology nodes.

3.5 SET, WOV and Mitigation Factors

It is noteworthy to mention the importance of internal buffers inside the flip-flop cell in determining the minimum capturable pulse width. The most common approach for constructing an edge-triggered register is to use a Master-Slave configuration as shown in Fig. 3.7. The clock inverters and the buffers between the master and the slave latches inside the flip-flop cell can increase or decrease the minimum pulse width by adding or subtracting delays and clock skews. Since most of the sequential cells in UDSM cell libraries are based on master-slave latches, this factor can be used as a control knob to adjust the minimum capturable pulse width. By changing the size of the buffers inside the flip-flop, we can relatively achieve less susceptibility to narrow SET pulses by imposing some area and delay overheads.

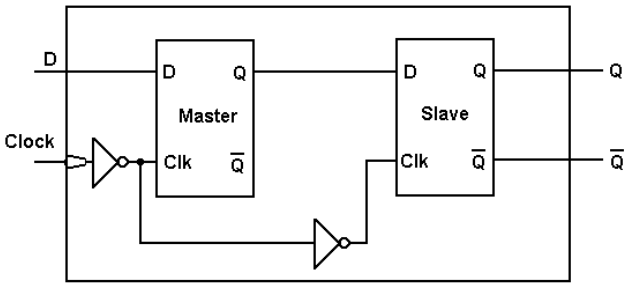

Figure 3.7: Master-Slave Flip-Flop

The minimum capturable pulse width for the 1X flip-flops in 45nm Nangate technology library has been observed to be 34 ps. By modifying the size of the clock inverter transistors between the master and the slave latches inside the flip-flop cell, the minimum pulse width can be increased from 34 ps to 100 ps and more. For instance, by increasing the transistor length, the minimum capturable pulse width of 109 ps can be achieved as depicted in Table 3.2. Consequently, for this particular internal inverter size, the flip-flop is immune to SET pulses with a width below 109 ps; which is approximately three times higher than the default value for the minimum capturable pulse width at 45 nm technology.

However the trade-off in the performance and the area should be considered. The rise or fall time depends on the channel resistance, which in-turn, depends on the device dimensions. The bigger the channel length, the larger the gate capacitance, hence the slower the circuit. With few exceptions, designers always use the smallest possible length available in a process to achieve the fastest speeds. There is more scope for varying the transistor width; for the case of an inverter, as the NMOS width increases, the fall time decreases but the rise time increases and as the PMOS width increases the rise time decreases but the fall time increases. The area taken up by the inverter must be also taken into account. The reduction in delay must be traded off against the increased area (and power) when the widths are increased. This is a cost-performance trade-off.

Moreover, interconnect capacitance plays an important role in masking the transient pulses and reducing the SER. Our simulations show that certain amount of capacitance at the output of the struck node flattens the transient pulse and reduces the pulse amplitude in such a way that the transient pulse cannot be sensed by the next combinational or sequential gate. For instance, we have run 10k Monte carlo simulations on 45nm chains of inverters and a flip-flop at the end with various output capacitances at the struck node and diverse SET pulse widths and amplitudes. The results are depicted in Fig. 3.8. It has been observed that

Table 3.2: Decreasing SET susceptibility using internal buffering re-sizing at 45 nm technology. Note that non-default values are non-physical as this is an experiment to look at operating boundaries.

Clock inverting transistors inside the flip-flop cell		
Transistor Width	Transistor Length	Min Capturable Pulse
0.27U (Default)	0.05U (Default)	34 ps
0.027U (Decreased)	0.05U (Default)	95 ps
0.27U (Default)	0.5U (Increased)	76 ps
0.27U (Default)	1.0U (Increased)	109 ps

the only key point from the SER perspective is the stuck node; the other nodes and their capacitance along the combinational path to the memory element almost have no effect on the generated SET, therefore the capacitances on the these nodes do not matter.

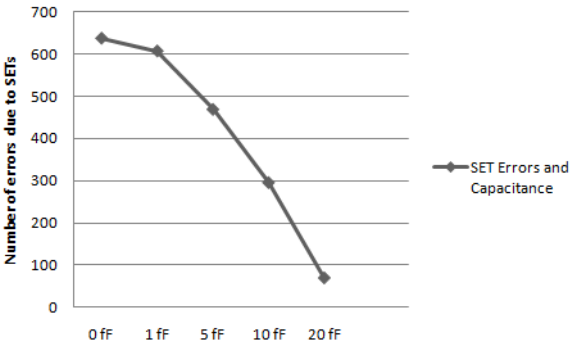

Figure 3.8: Capacitance at the struck node and SER

We also observed that the clock rise-time plays an important role in determining the location of the WOV. For very fast clock rise-times, the input pulse is captured right before the clock starts to rise as shown in Fig 3.9, but not with the minimum T_{DQ}. This is due to the internal state change of the Master-Slave latches inside the flip-flop.

3.6 Variation and the WOV

One of the major UDSM impacts on the devices are the variations discussed in Chapters 1 and 2 which cause significant unpredictability in the power and performance characteristics of integrated circuits. For a more accurate and realistic Failure In Time (FIT) and SER computation, the impacts of such variation must

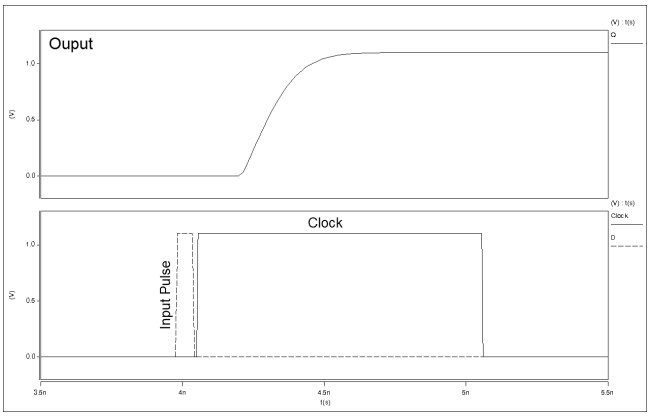

Figure 3.9: Narrow pulse properly captured right before the clock edge - 45nm technology

be taken into account. In this section, we investigate the impacts of variations on the WOV and the minimum pulse width.

Fig. 3.10 and Fig. 3.11 have been obtained using 10k Monte Carlo simulations to determine the vulnerability of 1X flip-flops (at 45nm Nangate technology) to SETs in the presence of process variations. Two parameters were chosen to mimic the impact of variation on transistors: Threshold voltage (V_{th}) and Oxide Thickness (T_{ox}). For each of the experiments resulting in figs 3.11 and 3.22, V_{th} and T_{ox} were varied fractionally to approximate the effect of process and intrinsic variations on both the sub-threshold and saturation regimes of device operation and the minimum captured pulse width has been determined.

The results show that even in presence of variations, the flip-flops are still very susceptible to narrow pulses and SETs. With 40% variation on Vth, the flip-flops could only cancel the 34 ps pulses in less than 60% of the cases and with 40% variation on Tox this immunity is less than 50% of the cases. The situation is aggravated when the pulses get wider in such a way that for pulses above 40 ps, the immunity and the probability to filter out the pulses are almost zero as depicted in Fig. 3.10 and Fig 3.11. In other words, variation does not have any significant impact on the minimum pulse width and it effectively neither improves nor reduces the vulnerability to very narrow pulses.

3.7 WOV and Soft Error Rate

The Soft Error Rate (SER) of any circuits can be defined as the average of all of the upset events for all of the particle strike times t_{strike}, with the collected charges

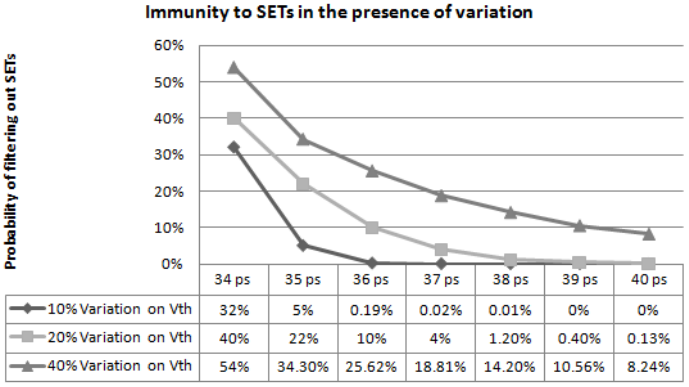

	34 ps	35 ps	36 ps	37 ps	38 ps	39 ps	40 ps
10% Variation on Vth	32%	5%	0.19%	0.02%	0.01%	0%	0%
20% Variation on Vth	40%	22%	10%	4%	1.20%	0.40%	0.13%
40% Variation on Vth	54%	34.30%	25.62%	18.81%	14.20%	10.56%	8.24%

Figure 3.10: 45nm technology

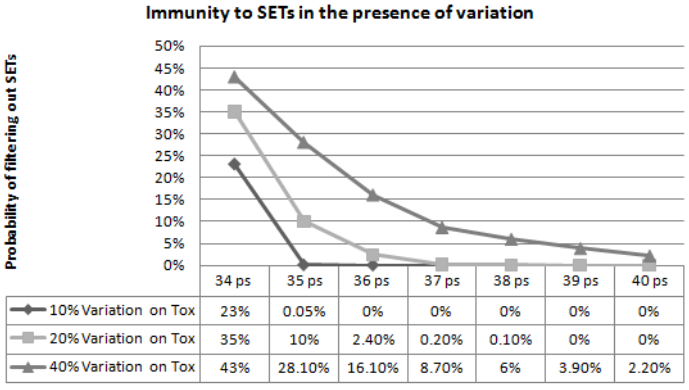

	34 ps	35 ps	36 ps	37 ps	38 ps	39 ps	40 ps
10% Variation on Tox	23%	0.05%	0%	0%	0%	0%	0%
20% Variation on Tox	35%	10%	2.40%	0.20%	0.10%	0%	0%
40% Variation on Tox	43%	28.10%	16.10%	8.70%	6%	3.90%	2.20%

Figure 3.11: 45nm technology

Q_i over all of the circuit nodes n. A_n is the impacted area of the struck node (drain area) and Prob($Q_{i,n}$) is the probability that charge Q_i is collected for every high energy particle at the struck node n. In this formula, the value of $Upset_{j,i,n}$ is 1 if and only if node n is upset by collected charge Q_i at strike time t_{strike} [124] [125]. We define an upset to be the condition in which the output of the struck node n changes more than half of the V_{dd}.

$$SER_{circuit} = (\sum_{n}^{nodes} A_n \sum_{i}^{Q} Prob(Q_{i,n})\Delta q \sum_{j=t_{strike}}^{T_{cycle}} Upset_{j,i,n}\Delta t) \times (\frac{Fluxof Particles}{T_{cycle}}) \tag{3.1}$$

The WOV can be defined as:

$$WOV = \Delta\tau_{n,i} = \sum_{j=t_{strike}}^{T_{cycle}} Upset_{j,i,n}\Delta t \tag{3.2}$$

Whereas the conventional WOV is equal to the sum of the setup and hold times.

Therefore the overall formula for the Soft Error Rate becomes [124] [125]:

$$SER_{circuit} = (\sum_{n}^{nodes} A_n \sum_{i}^{Q} Prob(Q_{i,n})\Delta q \Delta\tau_{n,i}) \times (\frac{Fluxof Particles}{T_{cycle}}) \tag{3.3}$$

As the WOV can directly affect the SER.

For memory arrays such as SRAM cells, it is generally safe to assume that $\Delta\tau_{\mathrm{n,i}}/T_{\mathrm{cycle}}$ is a constant value and equal to the duty cycle. For combinatorial circuits it depends on the location of the struck node in the downstream logic path and its distance to the destination flip-flop or latch, as the SETs can be masked logically or electrically (as discussed in section 3.1).

The SET induced SER of a circuit can be defined as:

$$SER_{SET} = \sum_{n}^{nodes} (Masking \times SET_{factor_n} \times \frac{PulseWidth}{T_{cycle}}) \tag{3.4}$$

In the formula above, the *Masking* factor is the canceling-out effect such as electrical masking or logical masking, and the SET_{factor} will depend on the flux energy, the impacted area, critical charge and charge collection efficiency of the devices for the nodes in the circuit.

Using the data for particle flux & critical charge per technology node data from [42] [57], and our calculated WOV, the SER of a chain of inverter circuit for 130nm, 90nm, 65nm and 45m can be calculated as shown in Table 3.3. The average SET PW of 45nm is extrapolated based on the values from 180nm down to 65nm technologies. Note that the logical masking and electrical masking factors have been ignored here as they would depend on the circuit topology, the layout and schematics, which would vary from circuit to circuit. We consider the worst case scenario.

From a reliability perspective, this means that to be able to account for SETs and to minimize SER, we need to characterize sequential cells beyond the conventional Setup and Hold time margins. We would need to make sure we have characterized the flip-flops and the latches up to the level that any input pulse that can be

Table 3.3: Average of SET-induced SER of different technology nodes - SER = FIT : Number of failures in 10^9 hours of operation

Calculated SER and WOV			
Technology	Min Capturable PW	Average SET PW	SET-induced SER (FIT)
130nm	65 ps	400 ps [126] [127]	1
90nm	56 ps	500 ps [128] [129]	10^1
65nm	44 ps	600 ps [130]	10^3
45nm	34 ps	700 ps (extrapolated)	10^4

captured by the flip-flop or the latch has been considered; even though it might result in a longer clk-to-q delay.

Moreover, we can characterize the SEE vulnerability of combinatorial cells in a look-up table fashion that can augment a standard cell library. Later, this data can be used to obtain more accurate circuit level SET-induced SER. To test the feasibility of this, we used our WOV methodology, SPICE and the formulas above to characterize and compute the results of FIT for an INVX1 cell connected to a DFFX1 Flip-Flop cell at 45nm for various SET pulse widths vs. different loads as shown in Table 3.4 The output loads are the values that are fed to the Flip-Flop input. The slew rate of 20ps (which is very fast for 45nm) has been used for the clock signal to account for the worst case scenario.

The steps in calculating (SET-pulse, output-load) pairs and creating the FIT tables are explained below:

1. Min_pair (Min_SET, min_output-load): Use SPICE to derive the 3-sigma probability distribution function (PDF) of the *minimum* capturable pulse width using our proposed WOV methodology with the *minimum* output load in the presence of process variation. Take the value of "$\mu - 3\sigma$" as the minimum point of SET-Load table.

2. Max_pair (Min_SET, Max_output-load): Use SPICE to derive the 3-sigma probability distribution function (PDF) of the *minimum* capturable pulse width using our proposed WOV methodology with the *maximum* output load in the presence of process variation. Take the value of "$\mu - 3\sigma$" as the maximum point of SET-Load table.

3. Middle_pairs: Divide the range between the Min_pair and the Max_pair into equal sections N. The value of N depends on the desired level of precision for the final SER-FIT table. Derive the N Middle_pairs.

4. Calculate the technology specific SER FIT of the Min_pair, Max_pair and the N Middle_pairs between, using the formulas above, based on the technology data.

In this example, the Threshold voltage (V_{th}) and Oxide Thickness (T_{ox}) were the chosen transistor parameters for variation in the 10k Monte Carlo runs. To find the min pulse width, in each scenario, the pulse height has been assumed to be constant Max i.e. V_{DD}.

Table 3.4: FIT-SET Char Table : INVX1 - Input = 0 - DFFX1

SET Pulse Width (ps)		200	300	400	500	600	700	800	900
Output Load (pF)									
	10	100	150	250	400	550	700	950	1000
	20	80	120	220	390	510	650	860	950
	30	65	100	200	350	470	600	770	900
	40	35	55	180	290	390	560	720	850
	50	20	40	120	220	330	510	650	800

The cell-specific SET-FIT data can be compiled and added to the design library along with the timing libraries. Such tables can be expanded to more dimensions to also include various slew rates and other significant factors such as supply voltage. In this case, by considering the output load, slew rate, the supply voltage as well as the SET pulse-width, we can obtain more accurate understanding of the SER of the circuit.

This is a first step towards generating SET FIT tables for WOV of flip-flops and latches. The results of such table are still pessimistic because of the following challenges:

- The first assumption is the minimum SET pulse width data is Gaussian. This is not necessary true. There are some scenarios where the data can fall into heavy-tailed distributions. Therefore the term "$\mu - 3\sigma$" might not represent the worst case. When dealing with non-Gaussian data, the term "$\mu - 3\sigma$" will not necessarily represent the the worst case corner. For instance, if the data distribution is Weibull or Log-normal (which is very common for reliability analysis [72]), then the distribution cannot be represented by μ

and σ parameters as Weibull distributions are parametrized by α (size) and K (shape), and depending on the values of α and K, the distribution can have zero tail on the left and very long tail on the right and "$\mu - 3\sigma$" become meaningless in this case. The same happens for Log-normal distributions which is defined by a shape and log-scale parameters and depending on the the values of these parameters, Log-normal distributions have very long tails on the right side.

- For simplicity and to save computation time and resources, we have assumed that the SET pulse height is fixed. In reality we will have to deal with bivariate distribution (Pulse Width and Pulse Height) or even multivariate distribution (Pulse Width, Pulse Height, Pulse Shape, Area under the pulse, etc) that again will not necessarily show Gaussian behaviors.

3.8 Discussion

The pulse width of transient glitches due to a particle hit is reported to be in the range of [78 ps, 206 ps] for 130 nm technology [126]. In [127] it is reported that, SET pulses range from about 400 ps to about 700 ps in a 130 nm process and this range increases to about 500 ps to 900 ps for a 90 nm process. Cannon *et al.* [128], measured heavy ion and proton-induced SETs in inverters, NAND and NOR gates for 90 nm technology. They observed SET pulses less than 400 ps wide in their library. In [129], a test circuit has been implemented to measure SET in IBM 130 nm and 90 nm processes. Test measurements with heavy ions and alpha particles show transient widths ranging from 100 ps to over 1 ns.

Although these empirical results reported in the literature slightly contradict each other, one point is evident: the minimum reported pulse width for the SETs due to particle hits in the literature is still much wider than the minimum captured pulse width as we observed in our simulations, making them susceptible to the SETs. This also suggests that SETs can lead to double-bit errors in UDSM circuits with clock frequencies in the GHz range, because pulses wider than one clock cycle can be captured by two consecutive clock cycles and eventually create double-bit errors instead of single-bit errors.

3.9 Concluding Remarks

This chapter addresses the second objective "To investigate timing vulnerability of UDSM combinational circuits and present a more realistic methodology to determine the vulnerability". We have presented an analysis of the conventional definitions for the WOV and have proposed a method to determine the minimum capturable pulse width for sequential cells which will lead to a more realistic SER computation. As suggested in the literature, the pulse-width of the most common SETs increases, for the same radiation environment, with technology scaling. This demonstrates the increasing importance of combinational logic soft errors. Considering this assumption for 45 nm technology and below, there is a high chance of transient pulses being captured by the flip-flops, because the WOV is very narrow at UDSM. Moreover, in circuits with GHz clock frequencies, this can even lead to double-bit errors rather than the conventional expectation of single-bit soft errors. Therefore conventional fault-tolerant techniques and single-bit error detection and correction methods may not be sufficient.

We also observed that interconnect capacitance plays an important role in masking the transient pulses and reducing the SER. Our simulations show that a certain amount of capacitance at the output of the struck node flattens the transient pulse and reduces the pulse amplitude in such a way that the transient pulse cannot be sensed by the next combinational or sequential gate. For instance, our simulations show that for typical 45 nm technology a load capacitance of greater than 20 fF at the output of the struck node flattens the SETs. Of course the exact amount of interconnect capacitance is not available before the place & route stage. For a more realistic calculation of SER, we should also consider interconnect capacitance.

Some of the results of the work in this chapter have been published as:

- M.M. Ghahroodi; M. Zwolinski; R. Wong; S.J. Wen, "Timing Vulnerability Factors of Ultra Deep-sub-micron CMOS," *European Test Symposium (ETS), 2011 16th IEEE , vol., no., pp.202,202, 23-27 May 2011*

Chapter 4

Soft Errors and Radiation Hardening By Design

Soft errors induced by radiation, causing malfunctions in electronic systems and circuits, have become one of the most challenging issues that impact the reliability of the modern processors even for sea-level applications. In this chapter we present an implementation of a radiation-hardened 32-bit pipe-lined Processor as well as two novel radiation-hardening techniques at gate-level. We present an SEU tolerant Flip-Flop design with 38% less power overhead and 25% less area overhead at 65nm technology compared to the conventional TMR Flip-Flop design. We also present an SEU-tolerant Clock-gating scheme with less than 50% area-power overheads and no performance penalty, compared to the conventional TMR for clock-gating. Our simulations show that the proposed schemes can recover from SEU errors in 99% of the cases.

4.1 Introduction

Since the main scope of this chapter is radiation hardening at gate-level, it is noteworthy to discuss Razor I and Razor II flip-flop architectures. Initially Razor flip-flop architecture has been introduced as a dynamic in-situ detection and correction of speed path failures. In the Razor methodology, the flip-flops in the critical paths are replaced by specific flop-flops called Razor (Fig. 4.1). A Razor flip-flop is comprised of a normal flip-flop and a shadow latch which receives the clock signal through a delay elements. Here the assumption is that, variability in

silicon, IR drop, noise, temperature and their manifestation as delay variation will show up on the critical paths first.

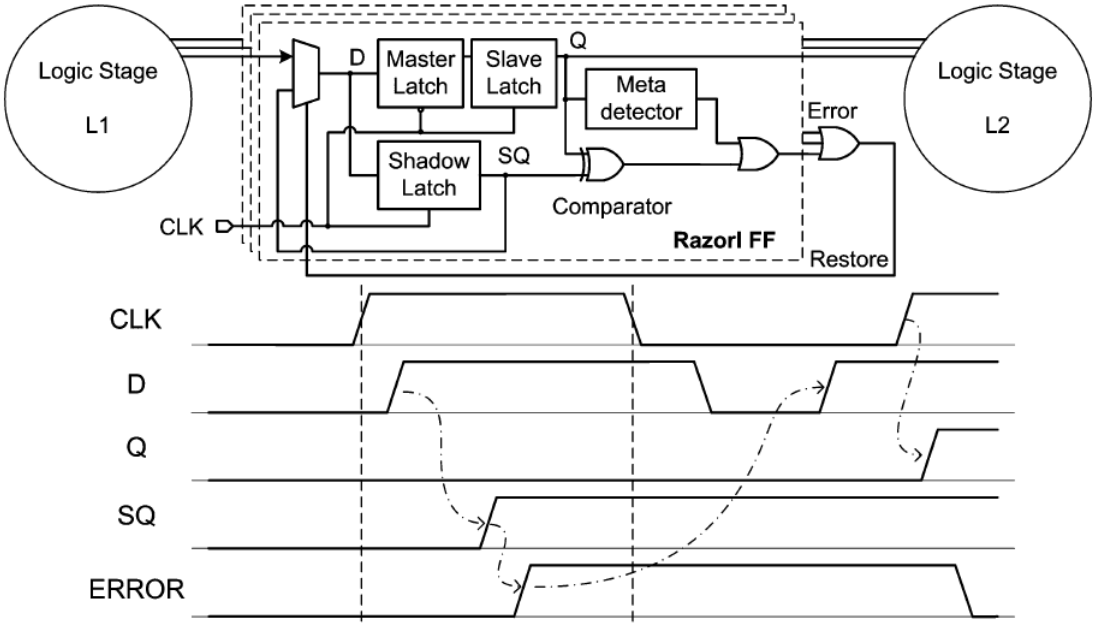

Figure 4.1: Razor I Flip-Flop [14]

In order to effectively address the design and timing issues in Razor I, Razor II was proposed.

Any delay increase that goes beyond the latching window of the downstream flip-flop in the critical paths will cause the flip-flops to fail in receiving the data. In this situation, the shadow latch will receive the correct data as the clock feeding the shadow latch is delayed by design. The comparator will compare the output of the flip-flop and the output of the shadow latch, hence the error can be detected. Here the problem is handling metastability conditions. This is because in the razor methodology, the system clock frequency is set based on the typical condition which does not unnecessarily cover the worst case scenario. This means the setup and hold time margins are calculated for the typical condition and in any chip that works in worst case scenario or slow corner or best case scenario or fast corner, there will be serious setup or hold time violations leading to metastability in the flip-flops. One of the features of Razor-style architectures is their ability to detect SEUs.

Razor I Features & Drawbacks:

- 75 transistors

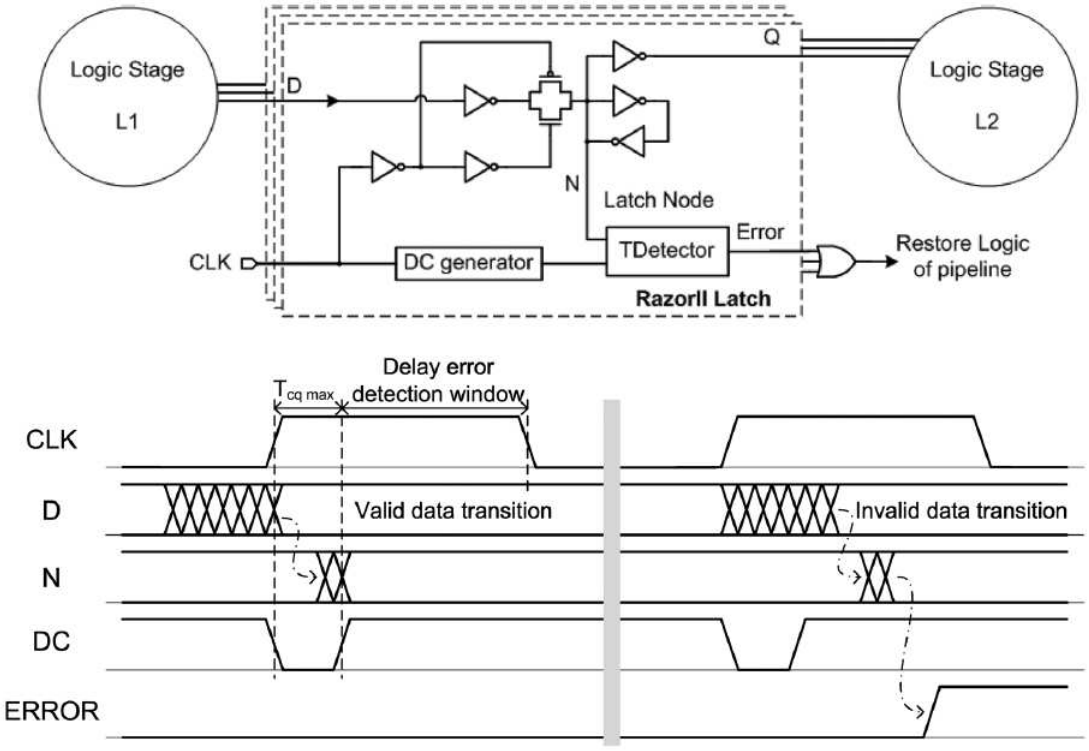

Figure 4.2: Razor II Flip-Flop [14]

- Only implemented on the Critical paths.
- Generation of a Global restore signal that goes to every pipeline flip-flop. This imposes significant timing constraints.
- Design of the meta-stability detector under the process variation is difficult. Because it needs to respond to meta-stable FF outputs across process corners.
- Additional risk of meta-stability on the restore signal itself that can potentially lead to system failures.
- Energy gain below the point-of-first-failure (POFF) is small (10%) compared to the energy gain from eliminating process-voltage-temperature (PVT) margins (35% to 45%). Due to the increase of the errors below POFF and the energy needed for recovery.
- Timing Error detection & Correction on the fly (at Flip-Flop level)

Razor II Features & Drawbacks:

- 39 transistors

- No need for a global pipeline restore signal, hence relaxing the timing constraints.
- Smaller in size, reduced clock pin capacitance (just one latch, unlike Razor I that is comprised of a Master-Slave FF plus one shadow Latch) hence less power, area overhead.
- Capable SEU detection.
- Error Detection only at Flip-Flop level (Recovery occurs at micro-architectural level)
- Increased cost of Instruction-per-cycle (IPC) penalty during recovering comparing to Razor I.
- Conventional Design flow has to be modified, because two different Clock trees are required: one for Critical Razor II flip-flops and one for non-Critical Razor II flip-flops. Otherwise excessive buffer insertion is required to balance the paths.

In the next part, we discuss the physical implementation of a rad-hard processor using the TMR scheme that we have done and then we propose two novel radiation-hardening techniques at gate-level; one for SEU-tolerant flip-flop design and the other for SEU-tolerant clock-gating scheme in a fully synchronous system.

4.2 Radiation Hardening of a 32 bit real-time processor at gate-level

We implement the TMR version of ARM Cortex-R4 [22] in 65nm general purpose technology node. The ARM Cortex-R Series embedded processors are fast, real-time and cost effective. They offer high performance, highly deterministic behavior, and built in safety features. Cortex-R4 is an implementation of the ARMv7-R architecture specifically designed for deeply embedded real-time applications such as HDD/SSD storage controllers, communications modems, and electronic control units (ECUs) for automotive and industrial systems. It offers significant energy efficiency, real-time response and predictable performance for real-time systems. The block diagram of Cortex-R4 is shown in Fig 4.3.

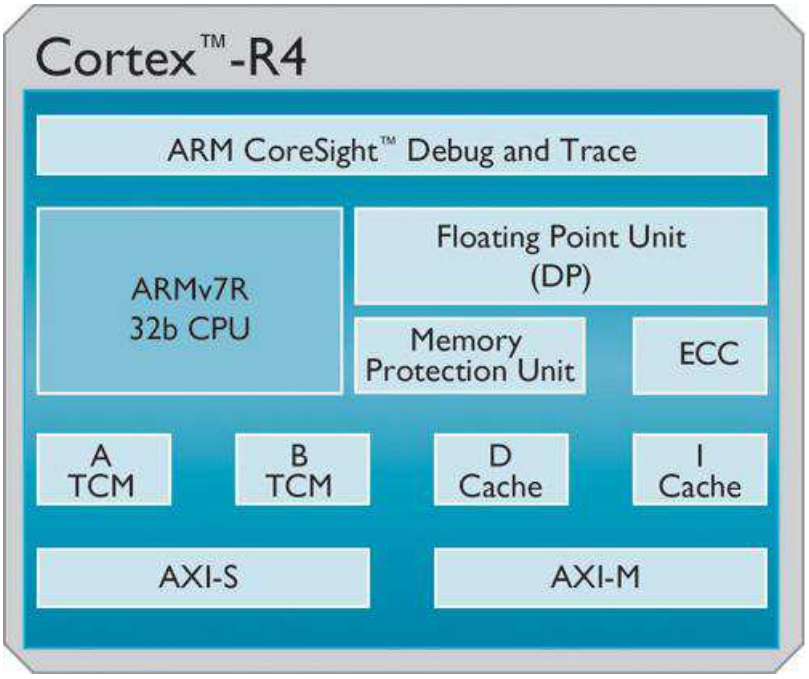

Figure 4.3: ARM Cortex-R4 [22]

Coretx-R4 has an 8-stage in-order dual-issue CPU pipeline, optional Memory Protection Unit (MPU) and CoreSight debug and trace units. It has optional tightly-coupled memories (TCMs), instruction/data caches and double-precision floating-point unit. Also, two 64-bit master and slave AXI ports are used to communicate with the external world. All caches and TCMs can be protected by ECC or parity against SEUs.

Cortex-R4 can also be used in a dual-core lock step (DCLS) configuration where a second redundant CPU can run in lock step with the first one. Both cores will share inputs and caches, and the outputs of the CPU are compared at every cycle to detect errors. This feature is used mainly in automotive systems but may also be useful for some space and avionics applications. The key difference between these two methods is that the majority voting of the TMR system enables the system to continue after a SEU or SET, providing tolerance and maintaining availability. With DCLS an SEU or SET can be detected but the correct result is unknown, as there is no majority voting, so the system must take an appropriate corrective action.

The Cortex-R4 CPU micro-architecture is mature and has undergone extensive verification and validation processes as an industrial product. Any change in the CPU micro-architecture will have repercussions in verification, validation. For these reasons, we opted to triplicate all the sequential cells at gate-level in a

product-quality CPU. All the flip-flops and latches are replaced by their TMR versions in the gate-level net-list. A majority voter is used to compare the output of the triplicated cell. Only the sequential cells are triplicated, so no redundancy is required for the combinational logic in the design.

The Synopsys IC compiler implementation flow is used to compile and place & route the entire design with the new TMR cell replacing the flip-flops of the default design. SPICE-level simulation is performed to validate the correct functionality of the design and the radiation immunity of the TMR cell. SET pulses of various widths are applied to the input of the TMR cell with various widths to find the point-of-failure in SET immunity. Different SET immunity levels can be achieved depending on the amount of input delay in the delay cells. For the case of our TMR cell, the immunity is for SET pulses with a width of 105ps. Any SET pulse wider than this can be potentially captured (depending on its time of arrival and its relation to the clock rising-edge) and cause an error. We have measured the target clock frequency, silicon area, and dynamic power. The results are presented normalized to the baseline Cortex-R4 CPU in Fig. 4.12, Fig. 4.13, Fig. 4.16 and Fig. 4.17. The target clock frequency of the TMR version is 35% slower than the baseline Cortex-R4. This is mainly because of the additional delays in the delay elements and majority voters in the TMR cells.

We use a baseline Cortex-R4 that does not have the optional floating-point unit and TCMs. However, it has instruction and data caches of 16KB each. Both caches are protected by ECC. So we take the TMR cell and apply it to every single flip-flop in the design excluding the caches, as they are protected by ECCs. Ideally, we should triplicate every single sequential cell in the cell library, characterize them and eventually add them to the cell library. During design implementation, the ECAD tools must be limited to choose only the TMR version of each flip-flop cell from the cell library. In this preliminary study, we pursue a non-optimal alternative where a single flip-flop cell that has all the functionalities (Set, Reset, etc) is selected from the standard cell library, and then this cell is triplicated to replace all the cells in the design.

To rad-hard a circuit, every single storage element (such as flip-flops or latches) should be hardened. We divide the core into two major parts: The core and the cache as depicted in Fig. 4.4. In our case, the memory and the data and instruction caches are protected with ECC, therefore our focus is on the core itself.

The rad-hard scheme we have used is based on spatial and temporal redundancy as shown in Fig. 4.5. In this case, all the flip-flops and latches will be immune to

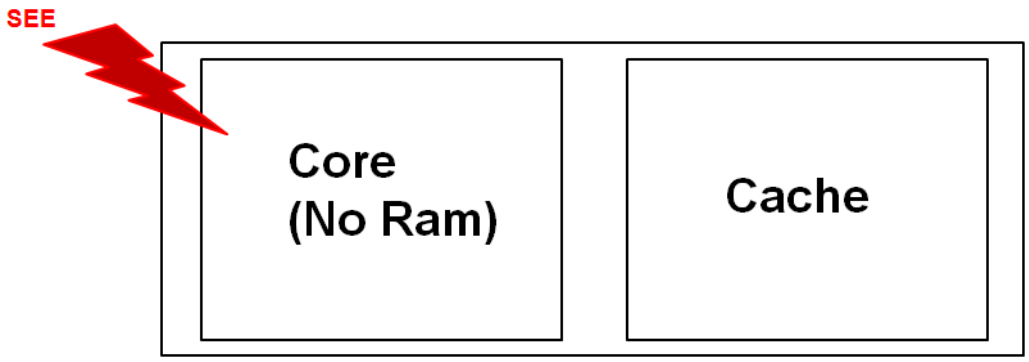

Figure 4.4: Processor Core

SEUs and SETs with certain pulse widths. It is noteworthy to mention that the immunity of this scheme to SETs depends on the amount of delays in the delays elements. By adding bigger delay elements, the SET immunity will increase at the cost of lost performance. Because the added delay elements will also add to the delay of all the path and especially in the case of the critical paths, this can be a huge performance drawback.

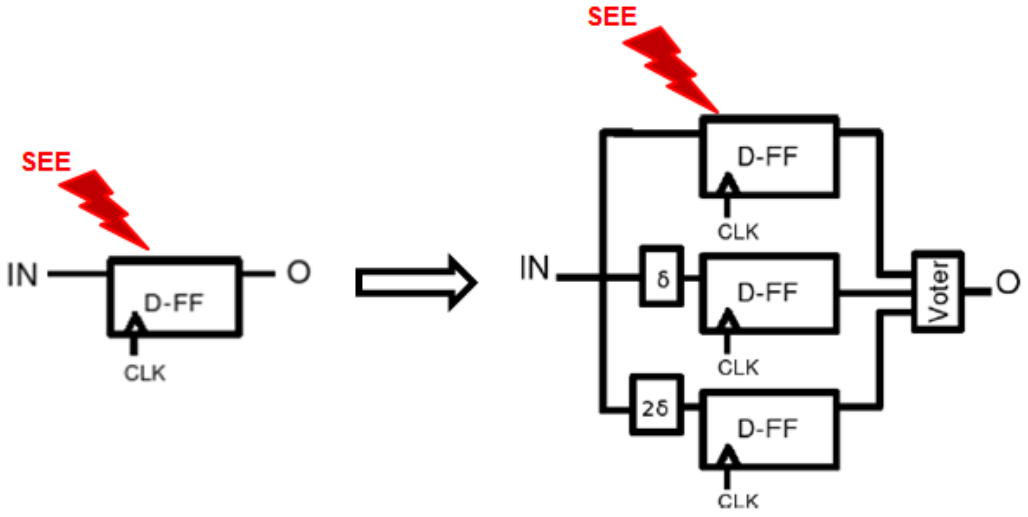

Figure 4.5: SEE Tolerant TMR Flip-Flop

For physical implementation, we used a 65nm general purpose TSMC standard cell library. A major bottle-neck is the lack of a TMR cell library. The cell library we have used contains 184 Sequential cells with various drive strengths. Implementing the SEE tolerant scheme on every one of these cells would be very time consuming and impractical, because it requires characterizing each new cell and then adding it to the library. A practical way of performing this task is to limit the synthesis tool in choosing only our design TMR cell. In this case, by just designing one TMR cell and forcing the EDA tools to use it we can implement the radiation

hardened version of the core much more quickly, but not optimally. Our proposed flow for choosing the right sequential cell for hardening is depicted in Fig.4.6.

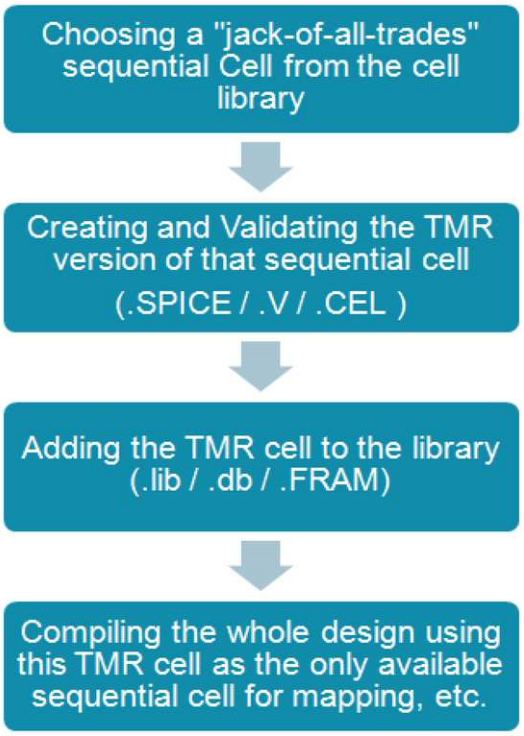

Figure 4.6: The proposed flow and the major steps in implementing a Rad-Hard core

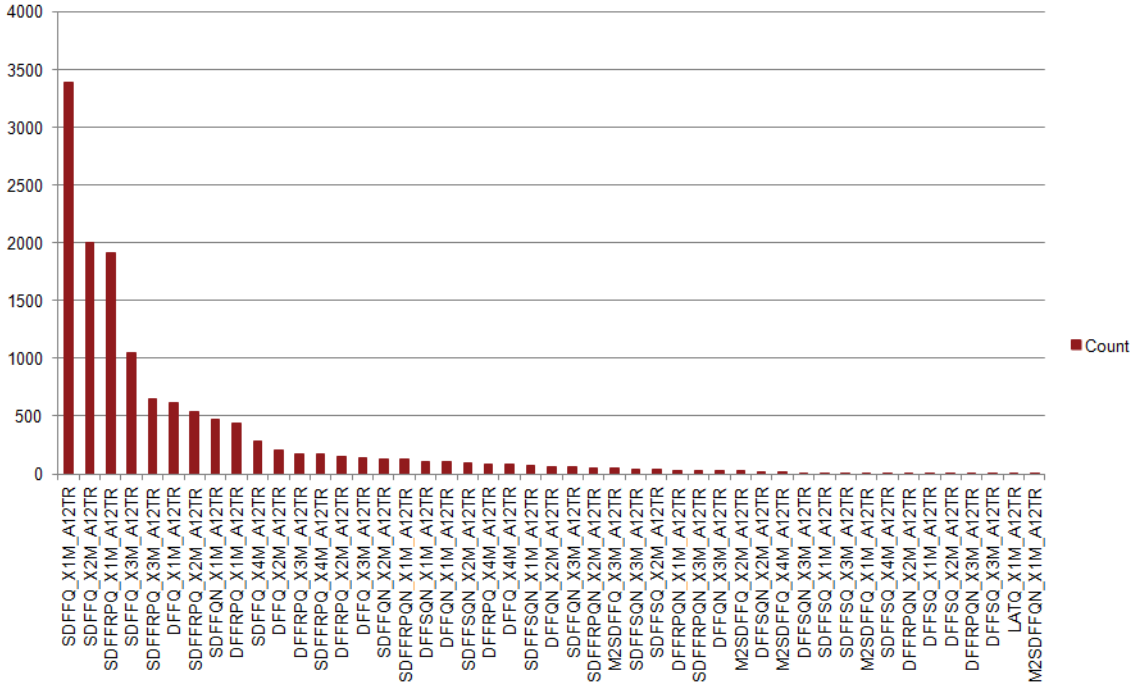

Figure 4.7: Default Sequential Cell usage by the core

As depicted in Fig. 4.7, the sequential cell SDFFQ-X1M-A12TR is the cell most used by the tools, but not all the sequential cells in the core core can be replaced by only this cell, because it does not have a Reset pin and some parts of the logic in the core need the Reset functionality. Therefore instead of that we picked the same cell but with a Reset input called SDFFSRPQ-X1M-A12TR which is a positive edge triggered static D-type flip-flop. This will be our jack-of-all-trades flip-flop. We have created a new cell named SDFFSRPQ-X1M-A12TR-TMR, which is a rad-hard version of that sequential cell. This new cell is comprised of 3 Flip-flops, delay elements and a majority voter as shown in Fig. 4.8 and Fig.4.9.

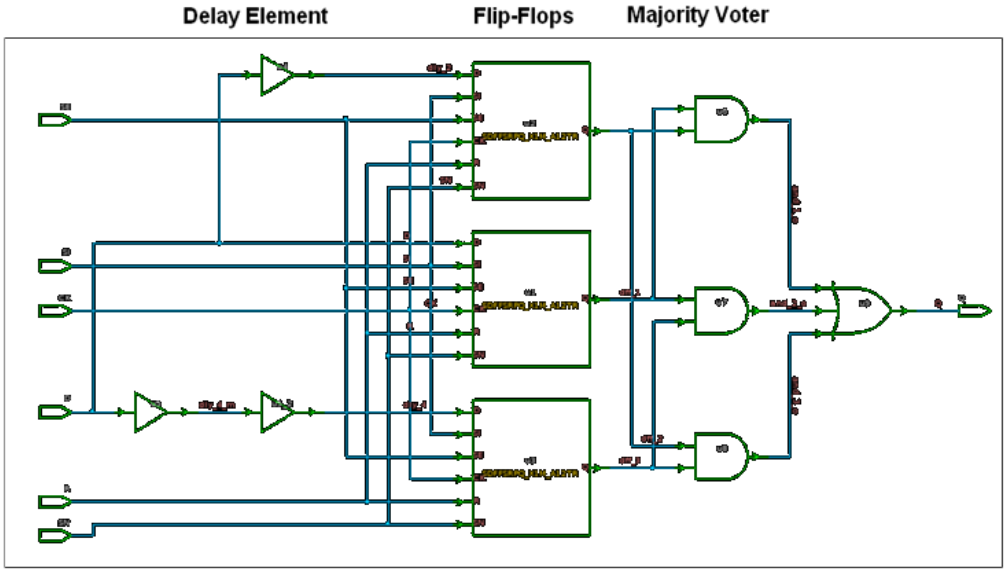

Figure 4.8: TMR cell schematic

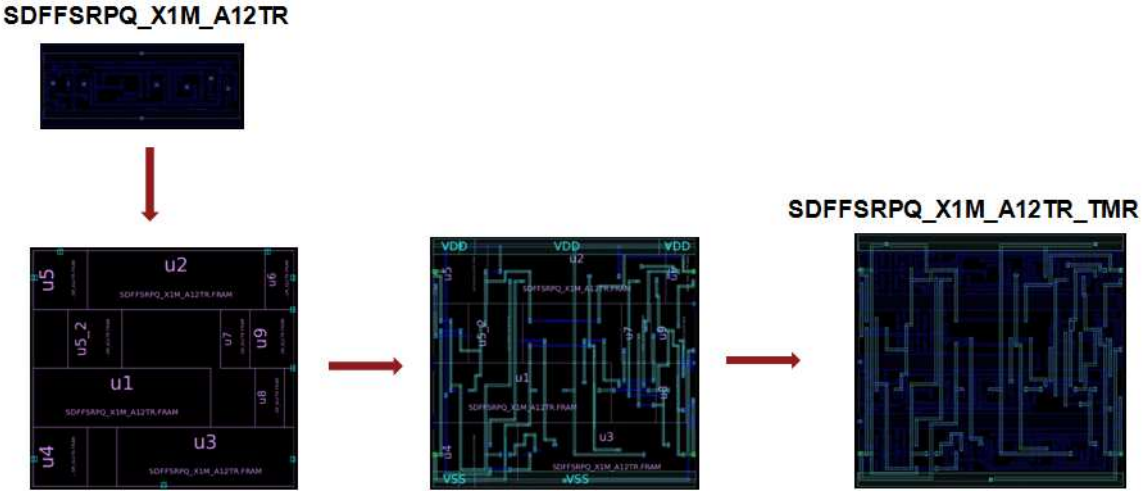

Figure 4.9: TMR cell layout

To validate the SET tolerance of the new cell, pulses have been injected into the flip-flop input in SPICE simulation. The simulation results show that this

architecture at 65nm is totally immune to the SETs with widths lower than 105 ps.

After adding the new cell to the library, the whole core has been synthesized and placed and routed using the new TMR cell as the only available cell to the tools. Since the new TMR cell is much bigger than a normal flip-flop, the floor plan must be expanded. Due to the fact that the cache memory consists of hard macro cells which are designed separately and floor-planned in the design at the bottom, it is easier to extend the floor plan from the top or at either side without moving the macro cells, as show in Fig. 4.10.

We have measured the total dynamic power of the design at post-layout. Total dynamic power increases by 41% in comparison to the baseline Cortex-R4 at its target CPU clock frequency. Because the clock frequency of the TMR design is 35% lower than the baseline Cortex-R4, then its total dynamic power is about 40% higher than the Cortex-R4 in spite of its doubled chip area. If the TMR version ran at the same clock frequency as the baseline, then we would expect the total dynamic power overhead to be closer to 100% higher based on 50-50 switching activity on average.

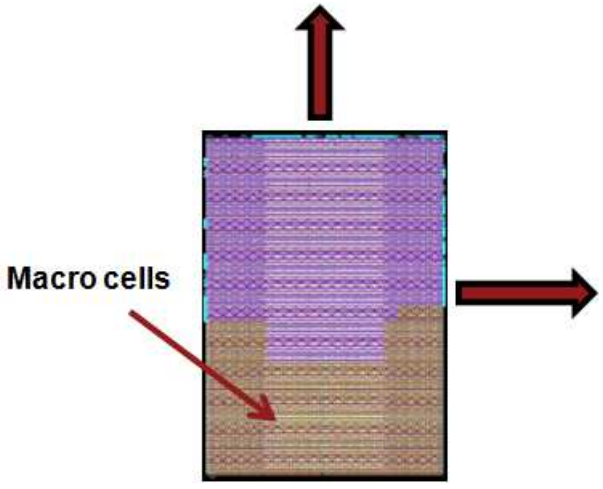

Figure 4.10: Extending Floor Plan

Post placed and route results are shown in the Fig. 4.11.

We have implemented four different versions of the core. 1) The default core, 2) The core with the jack-of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches. The power comparisons are shown in Fig. 4.12 and Fig. 4.13.

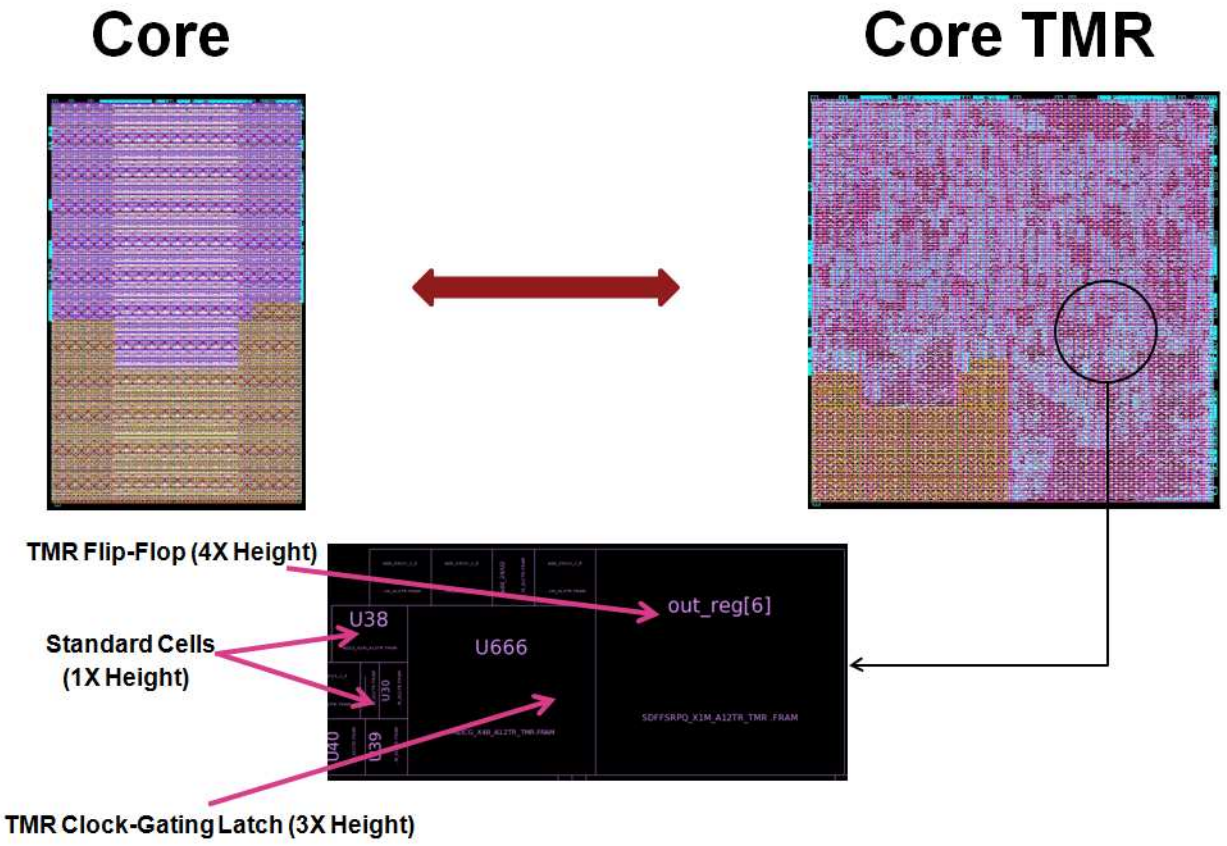

Figure 4.11: Core vs Core TMR

As shown in Fig. 4.14, Fig. 4.15 and Fig. 4.16, the chip area has been doubled because all flip-flops are triplicated. The area of an individual TMR cell is about 6 times larger than a regular flip-flop. In the original Cortex-R4, sequential cells occupy 20% of the CPU area, and the area occupied by TMR cells becomes 60% of the overall Cortex-R4 TMR chip area. So, the area occupied by TMR cells dominates the area of combinational and macro cells consisting of 16 KB instruction and data caches each. The overhead of the TMR cells also depends on the configuration options of the Cortex-R4 CPU. We have chosen 16KB instruction and data caches by default. If we had chosen 32KB instruction and data caches, the area overhead of the Cortex-R4 TMR would have been only 70% of the overall chip area rather than 100%.

The performance comparisons are shown in Fig. 4.17. The target clock frequency of the TMR core is about 35% less than the target CPU clock frequency of the original Cortex-R4.

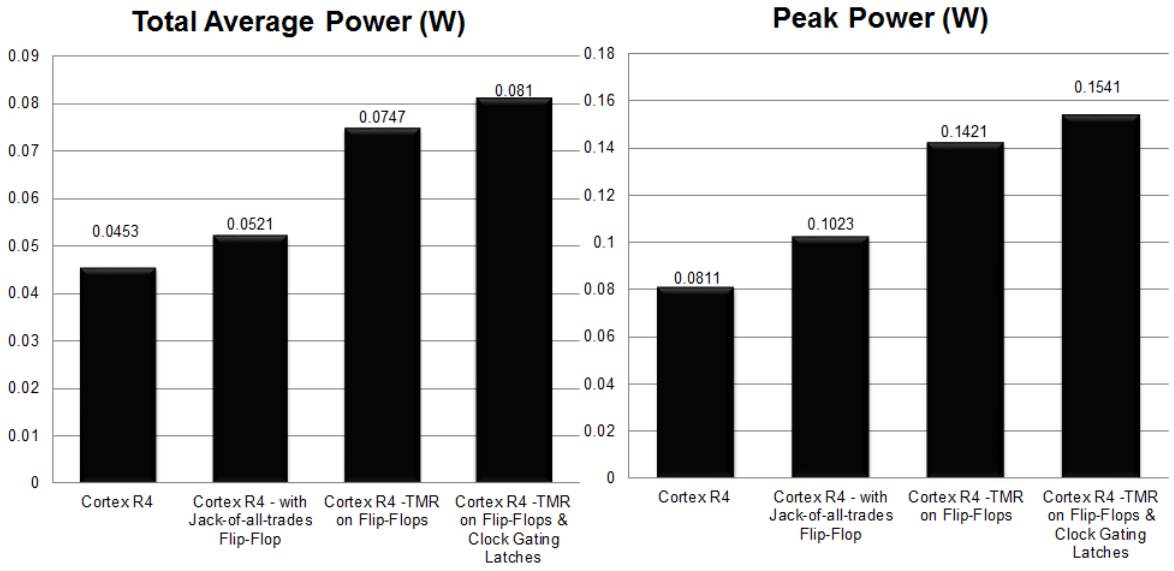

Figure 4.12: Total Average Power and Peak Power Consumptions - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

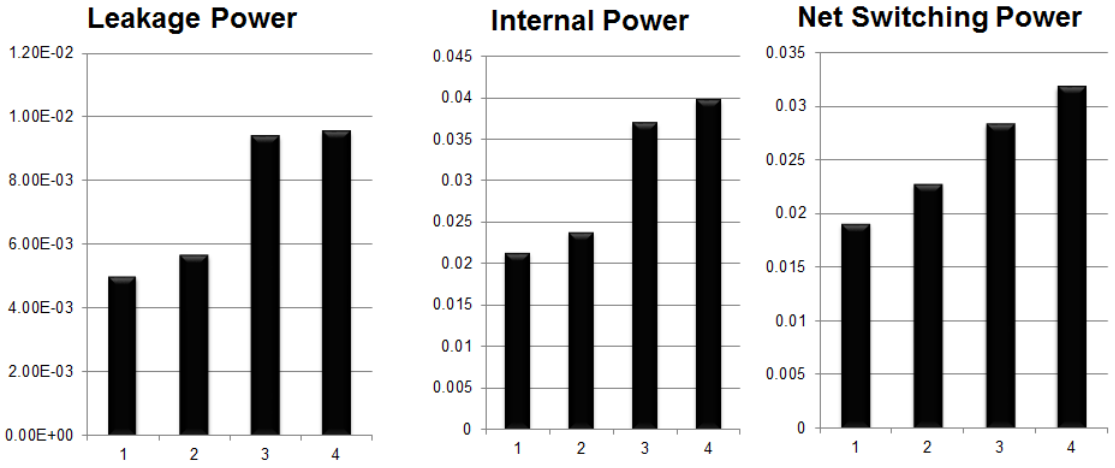

Figure 4.13: Leakage, Internal and Net Switching Power Consumptions - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

4.3 Discussion

This preliminary study investigates designing an SEU and SET tolerant ARM Cortex-R4 CPU targeting space and avionics applications. We design it by triplicating all flip-flops at the gate-level. In this way, the micro-architecture of the Cortex-R4 CPU is not modified and as this has the advantage of using a proven

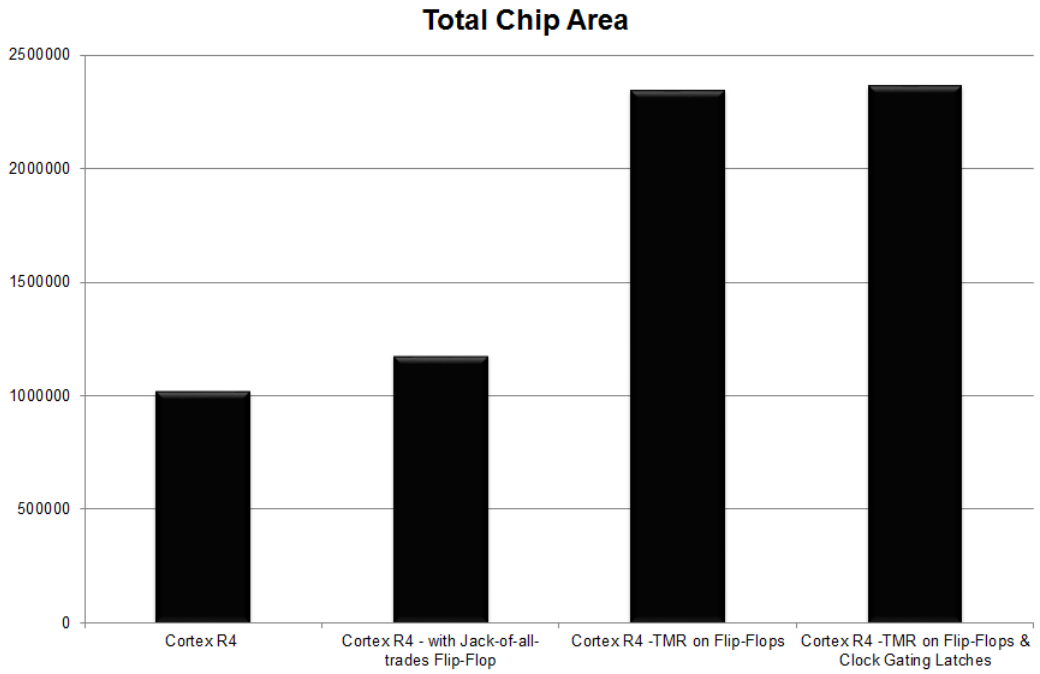

Figure 4.14: Area Comparisons in Total - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

CPU product with no software/tools change required. We have measured the overheads of the SEU and SET tolerant Cortex-R4 with respect to the original Cortex-R4. The SEU/SET tolerant Cortex-R4 occupies twice the chip area as the Cortex-R4 with 40% total dynamic power overhead, and its target clock frequency is about 35% less than the target CPU clock frequency of the original Cortex-R4.

The results of this study demonstrate that by triplicating all of the flip-flops at gate-level we can deliver radiation protection for mission-critical systems. There is always a trade-off between how much we would want the system to be SET-tolerant versus how much of performance degradation we can afford. Because the SET coverage of our scheme depends on temporal redundancy, the amount of the delay elements is the key factor in determining the SET-immunity. More immunity to SET-induced errors can be achieve at the cost of decreased performance.

Table 4.1 compares the availability and possible applications of our TMR scheme with the most discussed schemes in the literature. These comparisons are like-for-like.

There are various radiation-hardened processors in the literature. Table 4.2 compares our TMR core with some of the published radiation hardened processor

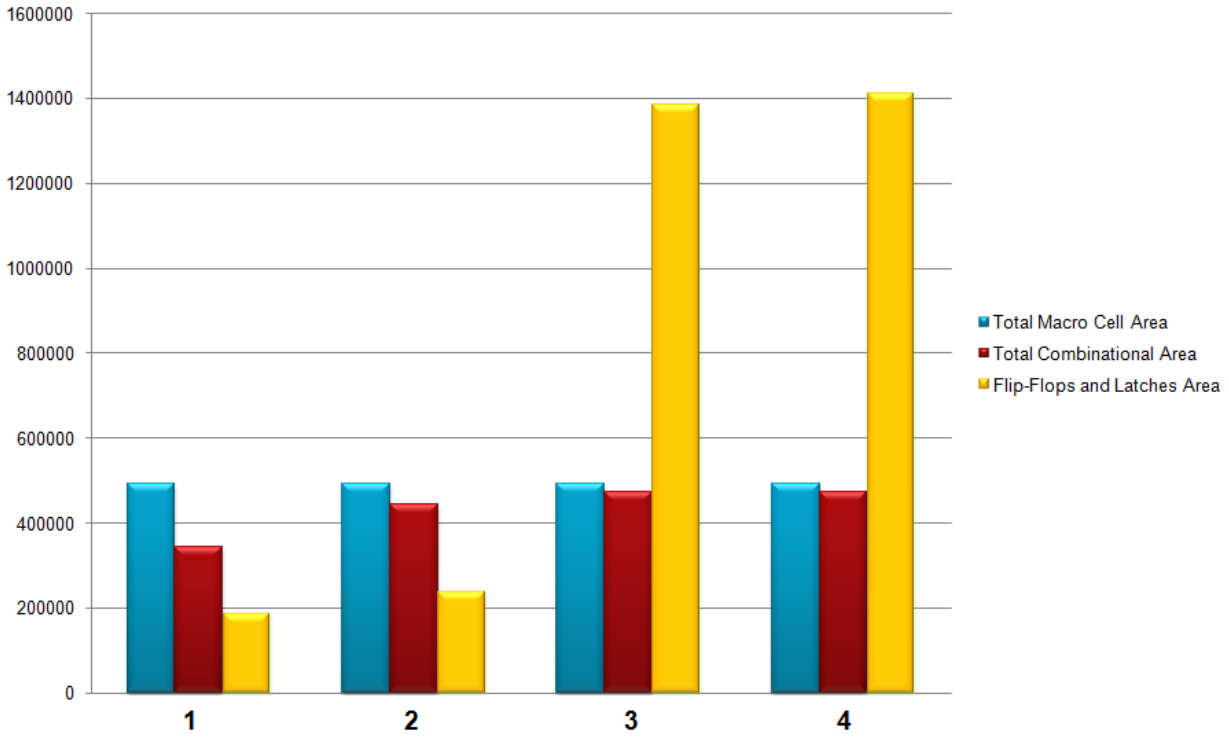

Figure 4.15: Area Comparisons in more Details I - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

Table 4.1: Availability and Applications of Conventional Radiation Hardening Schemes

Rad-Hard Scheme	Architectural changes	Availability	Application
Gate-level TMR	Not required	No down-time	mission-critical
System-level TMR	Minor	No down-time	mission-critical
Lockstepping	Moderate	low down-time	24x7 applications
RMT	Moderate	low down-time	24x7 applications
Pair-and-spare	High	very low down-time	24x7 applications

designs in the literature which they have also published their performance degradations. The performance degradations in Table 4.2 are as they reported in their published works.

One of the major benefits of gate-level TMR is the encapsulation of modifications at gate-level. This means that there are no architectural modifications and no changes in instructions and software programming. Such TMR system will always be in 'available' mode and the performance degradation will be due the TMR flip-flops and the majority voting circuits. System-level TMR will have approximately the same performance degradation, however the area and power overheads will

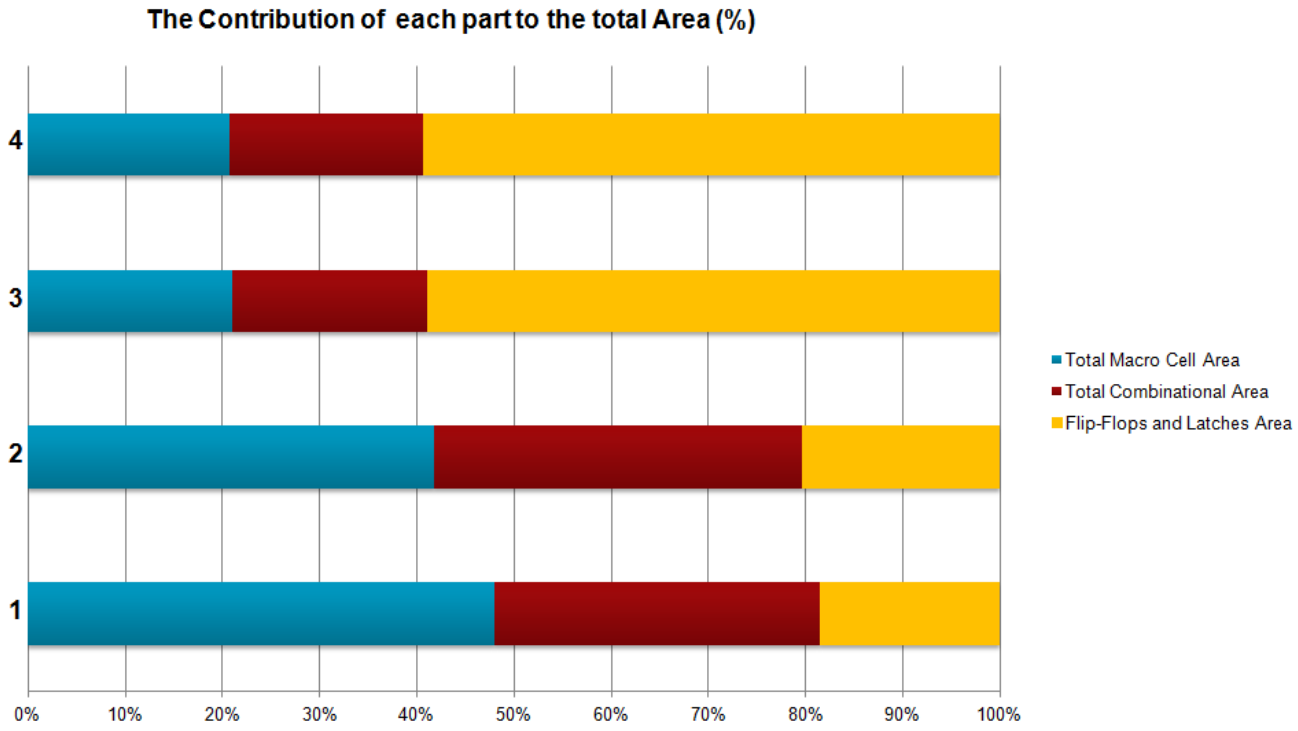

Figure 4.16: Area Comparisons in more Details II - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

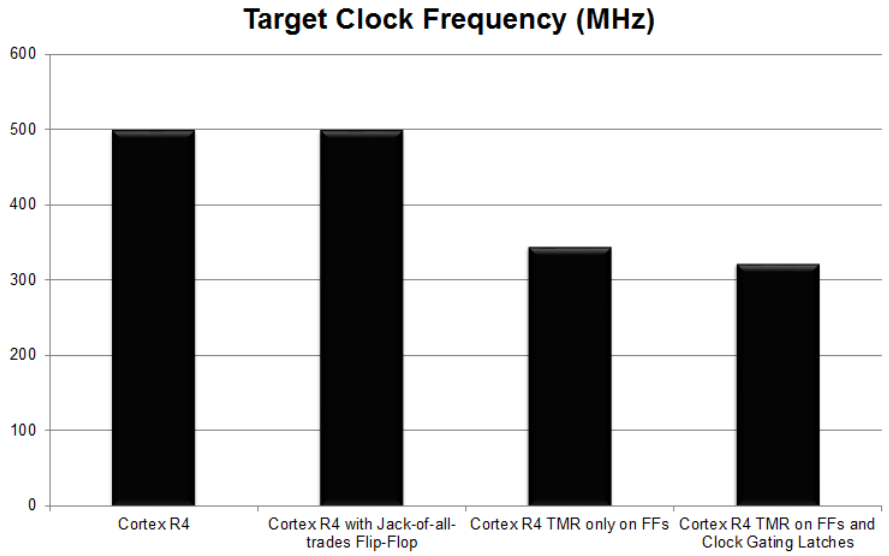

Figure 4.17: Performance Comparisons - 1) The default core, 2) The core with the jack- of-all-trades Flip-Flop, 3) The core with the TMR Flip-Flop, 4) The core with TMR scheme on Flip-Flops and the Clock gating latches.

be at least 200% comparing to a non-TMR system as every component will be triplicated. The benefit of system-level TMR is the simplicity of implementation

Table 4.2: Rad-Hard CPUs and their performance overheads

Rad-Hard Core Name	Architecture	Performance Degradation
ARM Cortex-R4 TMR (Ours)	32-bit ARMv7	35%
Phillips-Maxwell SBC-ATIM [131]	32-bit RISC	20%-40%
Honeywell RH-32 [132]	32-bit RISC	30%
Leon II Rad-hard [133]	32-bit RISC	40%
TI Radhard VC33 [134]	32-bit RISC	30%

as there will be three separate exact replicas of the system and there is no need to modify the implementation of the cores, although a wrapper circuit will be needed to connect the outputs of the replicated cores to the comparators.

In lock-stepping [135] [136], which is DMR by definition (with all the area-power overheads of a DMR system), redundant data are usually run on two identical but separate processors and the processor cores must have the exact same state in each cycle. These two cores will share the same inputs and the core are locked and synchronized cycle-by-cycle. In the event of a fault on one processor core, that core will be isolated and the other core will continue working while the faulty core is fixed or replaced. This decreases the availability of such scheme and makes it unsuitable for mission-critical applications.

Redundant Multi-threading (RMT) [102] [103], is similar to lock-stepping with the difference that in RMT, output comparison and input replication occur at the committed instruction stage, hence relaxing the cycle-by-cycle synchronization that is required in lock-stepping. Because outputs are only compared at the committing stage, the input replication mechanism of RMT becomes complex.

Pair-and-spare [137] [19], is basically a hybrid DMR system with RMT or lock-stepping mechanisms to detect faults. A pair-and-spare system is comprised of two pairs. The primary and secondary pair. Each pair is a DMR system by itself. The spare pair is always up-to-date since it receives continuous updates from the primary pair. This helps to make sure that the spare pair can replace the primary pair and resume the execution from any point where the primary pair failed. Obviously this can become very expensive in terms of overheads and architectural modifications.

Thus each radiation-hardening scheme has its own advantages, draw-backs and applications. In the next chapter we will discuss why TMR systems are suitable for mission-critical but short-term operations while other fault-tolerant mechanisms

that we discussed are more suitable for long-term operations with the penalty of probable sporadic down-times.

4.4 SEU-Tolerant Flip-Flop Design

In this section, we present a novel SEU-Tolerant Flip-Flop design. The main difference between our proposed design and other detection & recovery methods, which are typically based on the TMR concept is that our design is based on DMR. This obviously imposes less area and power overheads on the design. Conventional DMR methods can only detect errors with no recovery. However the presented method can detect and recover from SEU errors.

During any given clock cycle, the two flip-flops in a DMR scheme shown in Fig. 4.18 should hold the same value. If during any given clock cycle an SEU occurs on one of the flip-flops, the comparator compares the flip-flop outputs and detects the mismatch. But it cannot determine which one of the two flip-flops is hit by the particle. Hence error recovery is not possible. But the fact is that during any given clock cycle and right before the SEU occurrence and the mismatch between the outputs, both flip-flops were holding the correct value as depicted in Fig. 4.19. We exploit this fact and propose the SEU-Tolerant scheme depicted in Fig. 4.20. The timing diagram of the proposed scheme is shown in Fig. 4.21.

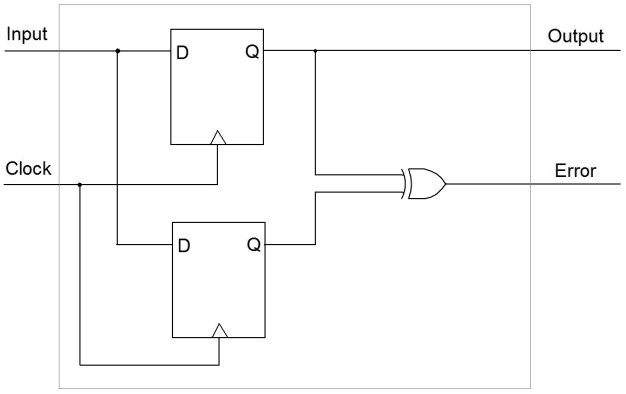

Figure 4.18: Dual-Module-Redundancy (DMR)

In SEU-free situations, the XOR output is always low and the active-low latch is transparent. The delayed version of the output from either of the flip-flops passes through the active-low transparent latch to the main output. By the time a particle hits one of the flip-flops and causes an SEU, the XOR goes High indicating the mismatch and it closes the latch. Since the latch is fed by the delayed version of one of the flip-flops (the amount of the delay is greater than the XOR propagation delay), the latch always closes on the correct value (the value before the SEU occurrence) and holds it. Therefore the main output remains unchanged and always correct.

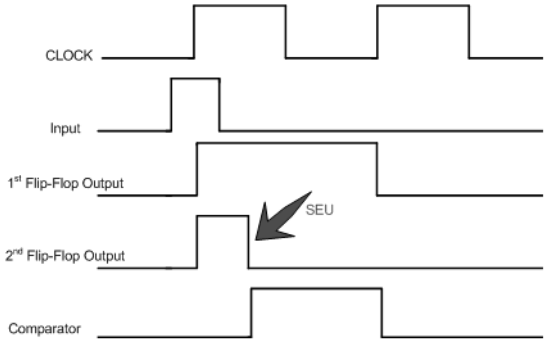

Figure 4.19: DMR Timing Diagram

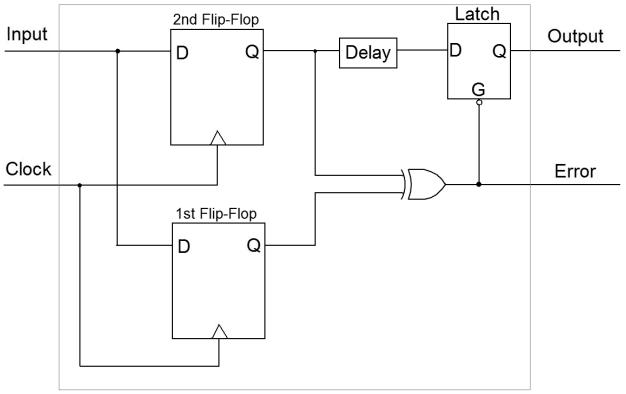

Figure 4.20: Proposed SEU-Tolerant Scheme - DMR with Error Recovery.

In other words, the latch is in transparent mode all the time behaving as a combinational gate and it is only in state-holding mode during an SEU occurrence. The advantage of such a circuit is that even if a particle hits the latch in any given clock cycle, it can only cause a glitch on the main output, because the latch is in transparent mode and not holding any state. This also means that, if in any give clock cycle, two particles strike the module, in such a way that the latch is hit first and one of the flip-flips is hit next, again the circuit can recover from the error, because the latch will close on the second particle hit and stores the correct value but with a glitch on the main output caused by the first particle hit.

The scheme has been implemented at transistor-level and gate-level for more accurate analysis. The proposed scheme can also be implemented at register-transfer level; however care should be taken at the place & route stage to reduce charge sharing and collecting between the sensitive nodes in a DMR/TMR sequential cell [138], [94]. It is also noteworthy to mention that the RTL implementation can complicate the timing issues by placing the storage elements of a DMR/TMR

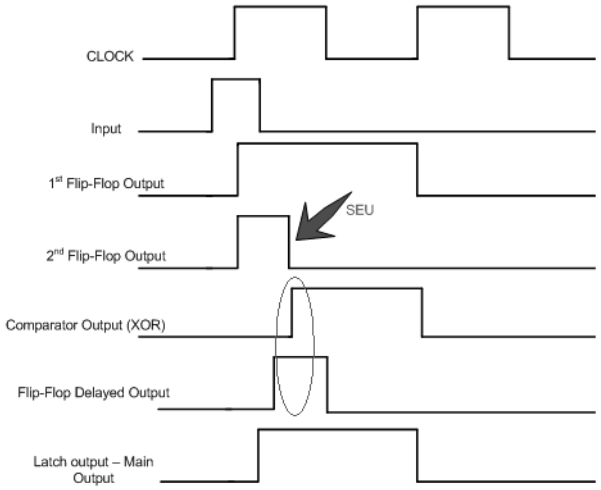

Figure 4.21: DMR with Error Recovery Timing Diagram - In the occurrence of an SEU, the latch closes on the correct value (the region under the oval), thus the main output is always correct.

sequential cell too far from each other, hence complicating the clock network synthesis in the place & route stage.

We have used 65nm general purpose TSMC standard cells, and assumed a 600 MHz clock frequency in the following results. The total number of transistors for the proposed flip-flop implementation is 70 compared to an equivalent TMR sequential cell (that is comprised of three flip-flops and the majority voting circuit implemented using the standard cells with the same cell size and driving strength) that contains 101 transistors. On average there is 38% less power overhead and 25% less area overhead because it can be implemented with fewer transistors and gates compared to a TMR sequential cell. The comparisons are depicted in Fig. 4.22.

The delay overhead in the TMR cell is due to the majority voting circuit which is comprised of three 2-input AND gates and one 3-input OR gate in our implementation. The propagation delay for the TMR flip-flop cell is the sum of $T_{\text{Clock-to-Q (of a none-TMR-Flip-Flop)}} + T_{\text{Majority-Voter}} + T_{\text{Interconnects}}$, while the propagation delay for the DMR with recovery cell is the sum of $T_{\text{Clock-to-Q (of a none-TMR-Flip-Flop)}} + T_{\text{(delay-element + latch(D-to-Q))}} + T_{\text{Interconnects}}$. The delay overhead in the DMR with recovery scheme is caused by the delay element and the latch. There is a 10% increase in the Clock-to-Q delay on average compared to the TMR cell, as shown in Fig. 4.22(b). This delay can be reduced by using smaller delay elements and faster latches or totally redesigning and characterizing the DMR cell as a new cell and adding it to the cell library.

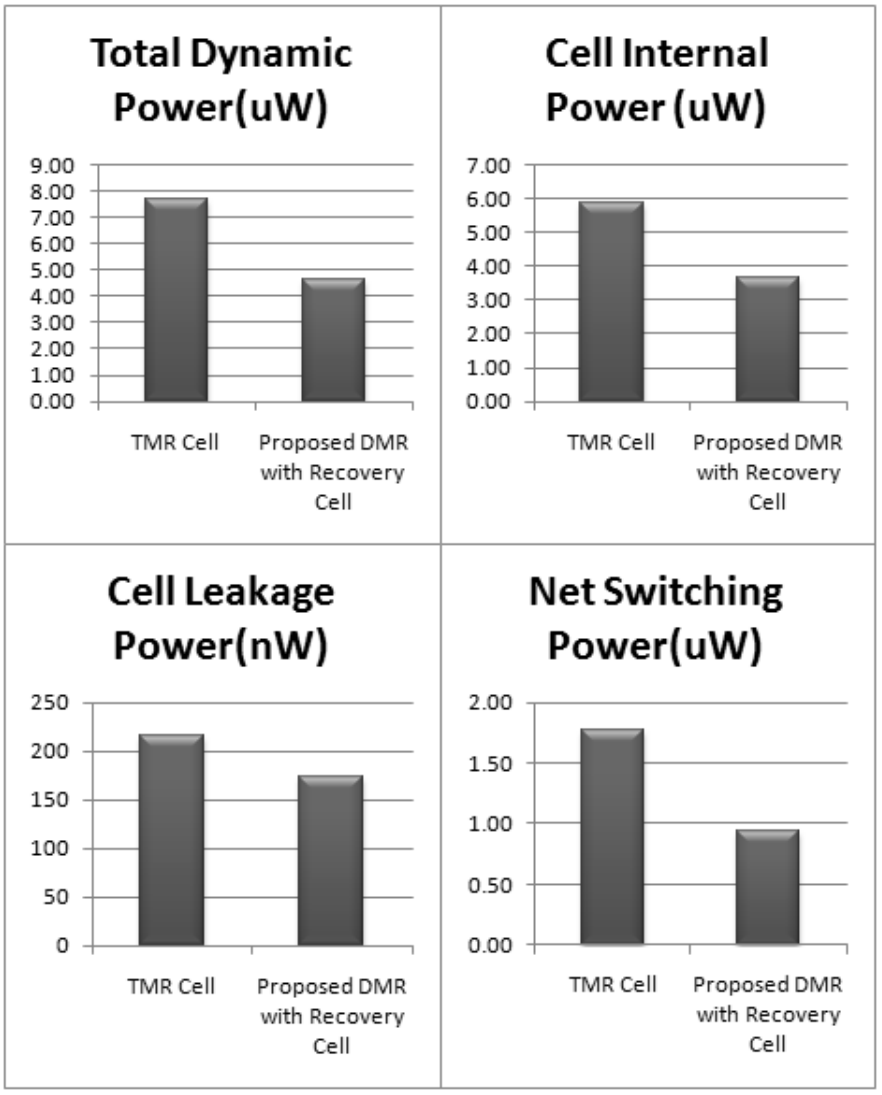

(a) Power Comparisons

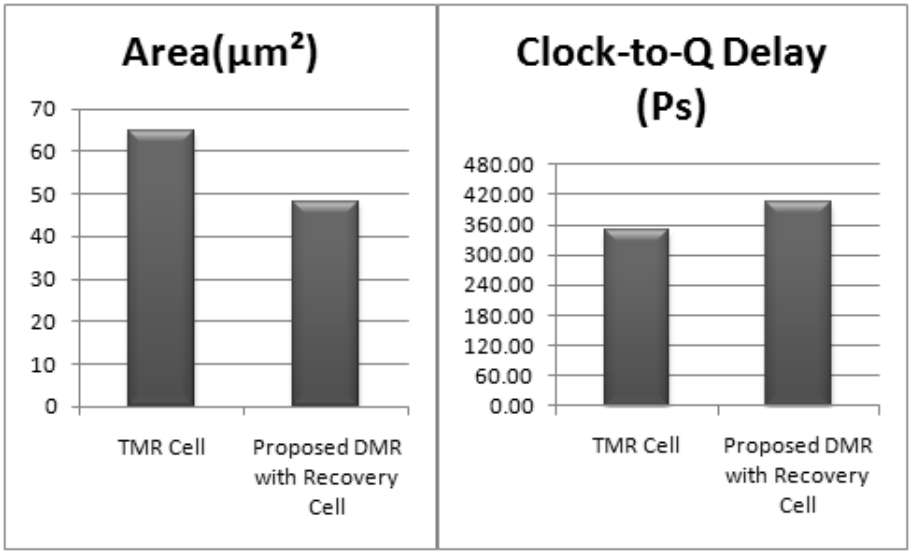

(b) Area & Delay Comparisons

Figure 4.22: Power, area & delay comparisons between two radiation-hardened sequential cells: a TMR cell *vs.* the proposed DMR with Recovery cell

To validate the SEU immunity of the proposed scheme, transistor-level simulations have been used for statistical SEU-fault injection. An example case is shown in Fig. 4.23. SEUs have been injected into either of the two flip-flops at different times during a given clock cycle in 10K Monte-Carlo runs to achieve a high level of confidence in the results. The results show that the proposed scheme can statistically detect 100% of SEU errors and recover from 99.1% of SEU errors. In less than 1% of cases, the SEU occurs right at the rising edge of the clock, in such a way that of one of the flip-flops does not have any chance to store the input value. In this case, the XOR gate goes high right on the rising edge of the clock indicating the error, but depending on which flip-flop the particle hits, the main output can be correct or incorrect. In these cases, if the flip-flop connected to the delay element is not the struck one, the main output is still correct, since the latch was fed by this flip-flop and closes on the occurrence of the SEU, but because of the mismatch in the XOR inputs, the error signal goes high and the output is considered faulty. We can also use the same methodology to design radiation hardened latches. We will present a specific technique for latches in the next section.

4.5 Radiation-Hardening and Clock-Gating Design

One of the most important issues that is usually ignored in radiation-hardening at gate-level is the radiation susceptibility of the low-power design techniques such as clock-gating. To save power, the clock signal is gated with an enable signal, in such a way that, when the flip-flop is holding its previous state and should not get updated, the clock will be disabled by the enable signal.

Conventional clock gating schemes use a latch to provide a glitch-free gated-clock to a number of flip-flops, as depicted in Fig. 4.24. This imposes more state-holding elements to the design with the same radiation susceptibility as the flip-flops. A particle hit on one of these clock-gating latches can create an SEU on the latch, that can eventually disregards the required enable signal status and update (or avoid updating) the stored values of the flip-flops during the clock cycle in which the flip-flops must hold their previous values (or get updated).

The conventional solution is to use a TMR scheme on the clock-gating latches. This imposes a 3.2x overhead in terms of area and power, plus the performance

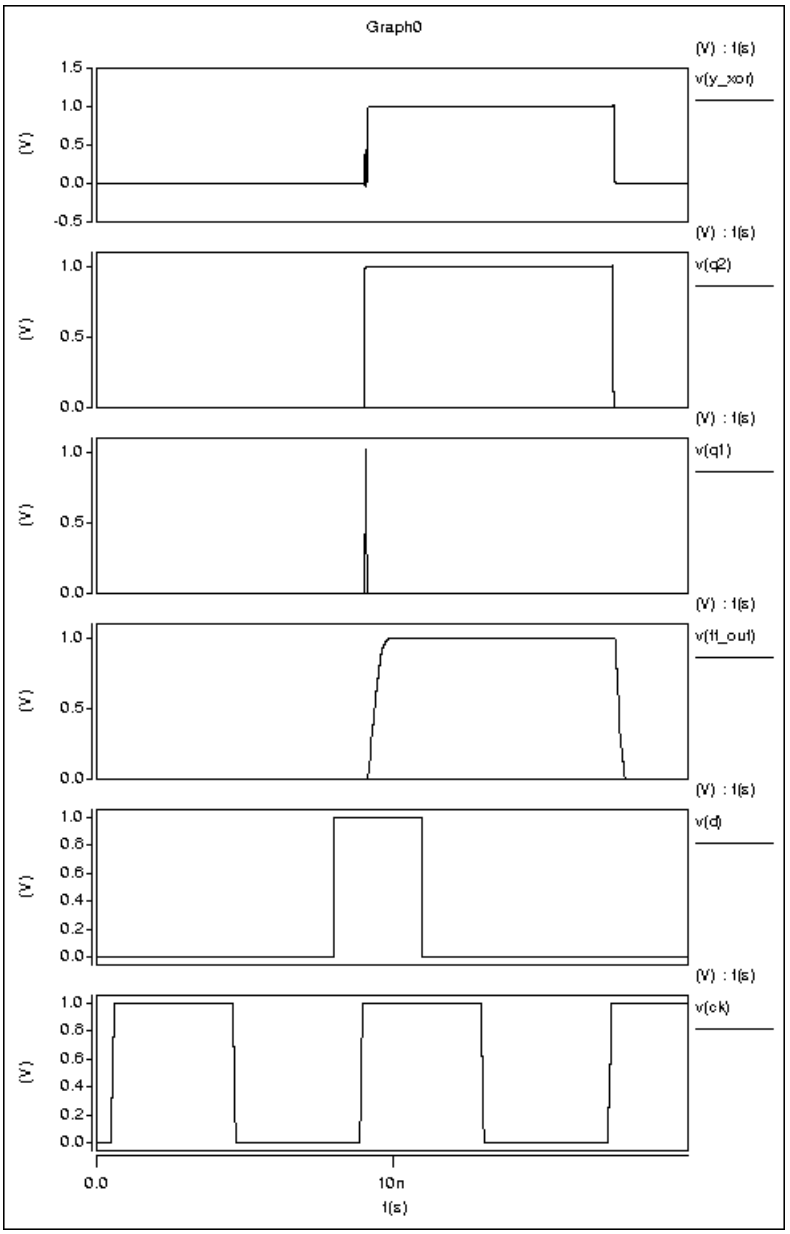

Figure 4.23: SPICE-level simulation. Despite one of the flip-flop outputs *q1* being almost destroyed due to an SEU, the main output *ff-out* is still correct.

overhead due to the existence of the majority voting circuit. Since the clock-gating is a special case, an alternative hardening technique is our proposed SEU-tolerant clock-gating scheme as shown in Fig. 4.25. A conventional TMR clock gating scheme uses three latches with the majority voting circuit. In our case, using the 65nm standard cell library, a TMR clock-gating latch contains 65 transistors. However the proposed scheme can be implemented using 27 transistors. This imposes less than 50% area-power overhead compared to the TMR version. Moreover there is no considerable delay overhead, because it does not have any majority voting

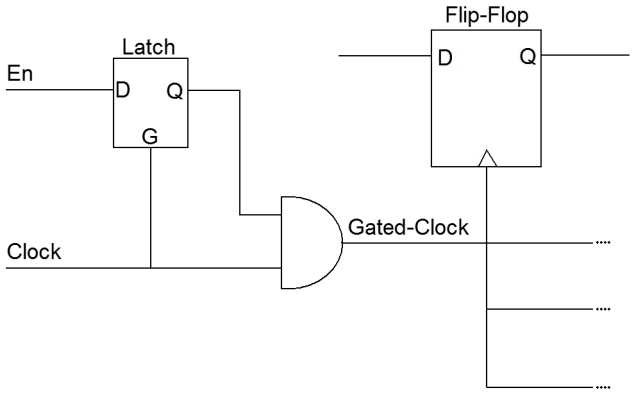

Figure 4.24: Conventional clock-gating scheme.

circuit.

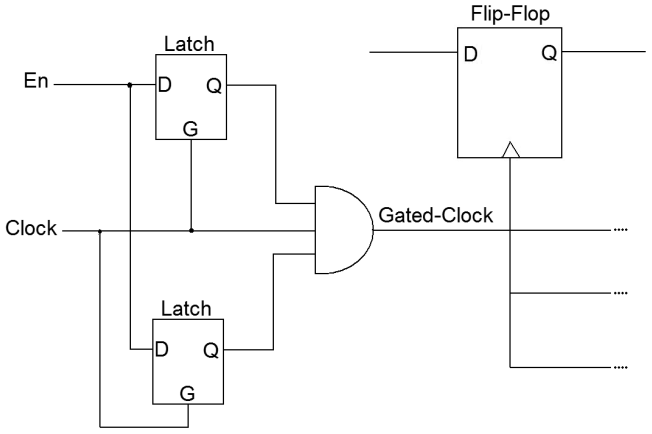

Figure 4.25: Proposed SEU-tolerant clock-gating scheme

The proposed clock-gating scheme is comprised of two active-high latches & one 3-input AND gate as depicted in Fig. 4.25. Two different scenarios exists:

- **Scenario 1**: *The SEU occurs when the Enable signal must be '0'*: Due to the fact that the controlling value on the AND gate is '0', therefore even an SEU on of the latches, changing '0' to '1' does not have any impact. This scheme guarantees that no SEU can activate the gated-clock signal and therefore in 100% of cases when the enable signal should be '0' it will remain '0', and an SEU on any of the latches cannot corrupt the flip-flop data by unwanted activation of the gated-clock signal "CLK-G" as shown in Fig. 4.26(a).

- **Scenario 2**: *The SEU occurs when the Enable signal must be '1'*: Since the controlling value on the AND gate is '0', any SEU on one of the latches can flip '1' to '0'. This causes the clock-gated signal "CLK-G" connected to the

flip-flops to have a narrower high phase, depending on the time that the SEU occurs during any given clock cycle Fig. 4.26(b). Our Spice-level simulations using 65nm technology show that only in less than 1% of cases this can lead to a data corruption on the flip-flop. For instance, in a worst case scenario, where the SEU occurs right at the rising edge of the clock signal in such a way that the gated-clock signal will be just a very narrow pulse looking like a glitch Fig. 4.27, but the flip-flop still gets updated properly.

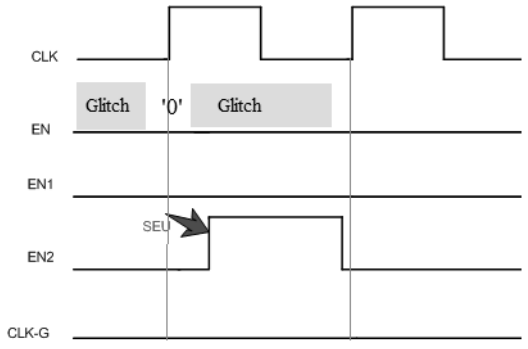

(a) Scenario 1: SEU when Enable signal must be 0

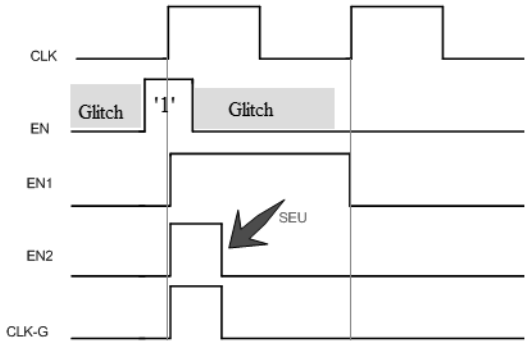

(b) Scenario 2: SEU when Enable signal must be 1

Figure 4.26: Timing Diagram for the proposed SEU-tolerant clock-gating scheme

Note that our scope in this section was focused on the RHBD clock gating. The flip-flops connected to this scheme need their own radiation hardening protection.

Table 4.3 summarizes the suitability of our proposed DMR-with-recovery scheme for different applications. As mentioned earlier, the notion of 99.1% SEU immunity is based on the statistical SEU injection during a given clock cycle. The Achilles' heel of this schemes falls into the 1% vulnerability that an SEU error cannot be corrected. Hence this scheme is not suitable for mission-critical or safety-critical applications in which data integrity, or system availability at 100% level

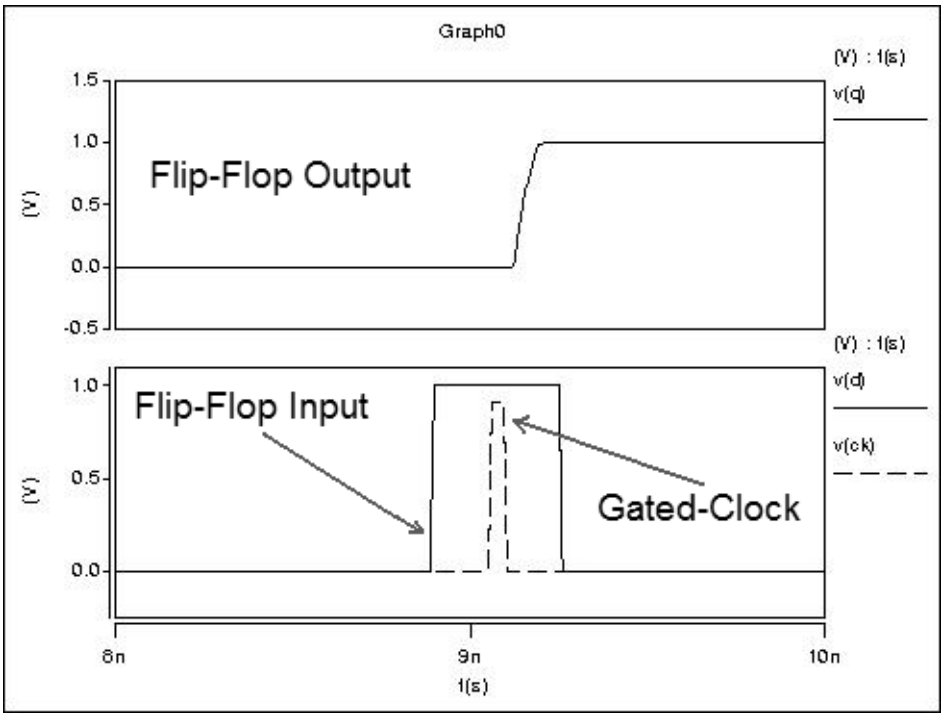

Figure 4.27: SEU-tolerant clock-gating scheme: A worst case scenario - The clock signal is almost destroyed by the SEU, but the flip-flop still gets updated properly but with a bit longer clock-to-q delay.

of confidence is a must (Although even with TMR, there can never be a 100% level of confidence as a simultaneous particle strike on two out of three of the flip-flops can corrupt the majority voter output). Apart from this, due to the performance overheads, our DMR-with-recovery scheme may not be suitable for non-critical but ultra-high-speed applications such as switches or web-servers. If resetting the system can be tolerated, then our scheme can be used in back-end databases systems and low-end web-servers due to the fact that our during that 1% timing window in which an SEU event cannot be correct, still the SEU event will be detected so by restarting the core or flushing the pipe-line, the system will be able to recover from such a fault.

Table 4.3: Applications

Applications	Data integrity	Availability	DMR-with-recovery scheme
Mission-critical	Critical	Critical	Not Suitable
Web-server	Moderate	High	Only Suitable for low-end market
Backend - database	High	Moderate	Only Suitable for low-end market
Desktop	low	low	Suitable
Mobile	low	low	Suitable

4.6 Concluding Remarks

This chapter addresses the third objective "To investigate timing vulnerability of UDSM sequential circuits and state-holding elements and propose a possible hardening-by-design solution". In this chapter we discussed the implementation of a radiation hardened processor core using our proposed TMR-based flow for the state holding elements such as flip-flops and clock-gating latches. The results show that the area-power overhead on average is less than 100% for such schemes. It is also noteworthy to mention that care should be taken when it comes to the layout design of such TMR cells and the placement of sensitive nodes should be done in such a way that redundant element (sensitive nodes in a flip-flop) should be placed far enough from one another to avoid an SEU on two store holding elements in one TMR flip-flop.

We have also presented a novel technique at gate level to design radiation-hardened sequential cells. The approach taken is based on DMR with error recovery, and results in 30% less area and power overhead compared to TMR sequential cells. We also presented a novel technique for a radiation-hardened clock gating scheme which results in less than 50% area and power overheads, plus no performance overhead comparing to the TMR version. Since we use conventional standard cell libraries and EDA tools to apply these techniques, no additional modification or custom libraries or tools are needed. Our SPICE-level simulations show that these methods are statistically able to recover from 99% of SEU errors.

Some of the results of the work in this chapter have been published as:

- M.M. Ghahroodi; M. Zwolinski; E. Ozer, "Radiation hardening by design: A novel gate level approach," *Adaptive Hardware and Systems (AHS), 2011 NASA/ESA Conference on , vol., no., pp.74,79, 6-9 June 2011*
- M.M. Ghahroodi; E. Ozer; D. Bull, "SEU and SET-tolerant ARM Cortex-R4 CPU for Space and Avionics Applications", *MEDIAN'13, The Second Workshop on Manufacturable and Dependable Multicore Architectures at Nanoscale, May 2013*

Chapter 5

Reliability and In-field Logic Repair

Ultra Deep-Sub-Micron CMOS chips have to function correctly and reliably not only during their early post-fabrication life, but also for their entire life span. In this chapter, we present an idea at architectural level to deal with this. In the case of any permanent faults, logic spare-blocks will replace the faulty blocks on the fly. Meanwhile by shutting down the main logic blocks partial threshold voltage recovery can be achieved which will alleviate the ageing-related delay impacts and timing issues. The proposed technique can avoid fatal shut-downs in the system and will decrease the down-time, hence the availability of such systems will be preserved. We have implemented the proposed idea on a pipe-lined processor core using conventional ASIC design flow. The simulation results show that by tolerating about 70% area overhead and less than 18% power overhead we can dramatically increase the reliability and decrease the downtime of the processor.

5.1 Introduction

In this chapter, we present a technique to retain the functionality of CMOS circuits and processors in the occurrence of any reliability issues that can lead to permanent faults, while keeping the overheads reasonable. We evaluate the reliability of the proposed techniques using Markov models and NASA SURE reliability tool. We also apply the technique to tackle timing and delay faults due to process variation or ageing impacts.

Various reliability factors such as safety and robustness as well as resilience to malfunctions need to be addressed in dependable systems [139] [140]. Infrastructure systems such as internet, banks, stock market, the electric power networks, etc. require dependability to ensure social stability and in many cases this must be 24 h, 365 days. Providing this level of dependability and availability and making reliable electronic systems out of unreliable CMOS components while keeping overheads as low as possible is a major challenge today.

To increase the reliability of systems, various fault-tolerant techniques based on redundancy in time/spatial domains are being used. As for hardware-based redundancy: majority voting redundancy (TMR), stand-by redundancy and hybrid modular redundancy (HMR) are the major techniques. It is generally believed that N-modular redundancy systems are more reliable than, say, stand-by redundancy systems. For instance, the reliability of an N-modular redundancy system such as TMR with three redundant components (with equal reliability R for each component and assuming that fault-detection coverage is 100%) can be formulated as:

$$R_{TMR} = R^3 + 3R^2(1-R) = -2R^3 + 3R^2 [21] \tag{5.1}$$

On the other hand, the reliability of stand-by redundancy systems with two redundant components (with equal reliability R for each component and assuming that fault-detection coverage is 100%) can generally be expressed as:

$$R_{stand-by} = R^2 + 2R(1-R) = -R^2 + 2R [21] \tag{5.2}$$

As depicted in Fig 5.1, mathematically speaking, the reliability of stand-by redundancy systems is higher that the reliability of TMR systems for all R. However this can be a tricky comparison. Because the TMR systems have higher fault-detection coverage with the feature of data comparison by nature, and higher reliability especially for shorter t. Therefore, the TMR systems are widely used for ultimate safety systems for life-critical applications with relatively shorter mission time. Also the overheads of TMR systems particularity in terms of power consumption are extremely high.

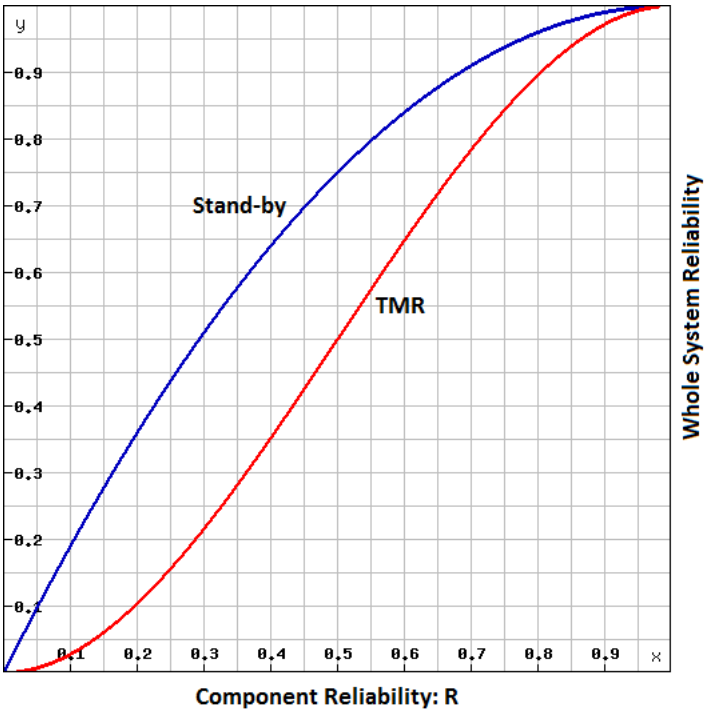

Figure 5.1: Reliability Comparisons

5.2 In-Field Repair in CMOS Circuit Design

5.2.1 Motivation

As mentioned in chapter 2, reliability issues such as Aging, Time-Dependent Dielectric Breakdown (TDDB), transistor, Hot-Carrier Injection (HCI) and Negative Bias Temperature Instability (NBTI) degradations are inevitable in the ultra deep-sub-micron era and each can significantly affect system MTBF. These phenomena can manifest themselves as performance degradation, timing errors or hard-errors in the chip, leading to a total failure of the processor permanently [141].

Here the challenge is increasing the maintainability of a system or a circuit, comprised of CMOS devices, once a malfunction is detected. Since these major reliability issues are inevitable by nature, therefore the goal is decreasing the downtime hence the MTTR of such systems. An overlooked fact is that the *"MTBF Countdown Clock"* does not start until the device is under stress i.e. the hardware is powered up. In other words, device degradation does not happen if the devices are permanently "off" because there is no electrical field in the transistor channels; hence electrons will not become energetic or "hot" enough to damage the channel oxide interface (causing HCI) and there will not be any high vertical electrical field in the channel at high temperatures (causing NBTI) for instance.

Moreover as reported in [10] [62] and as illustrated in [63] [64] and [10], some of the impacts of the aforementioned reliability issues such as the threshold voltage shifts in CMOS devices can be recovered to their original values by removing the stress from the devices and turning them off. For instance, in [142], the measurements of NBTI performance degradation of a processor core at 90nm have been demonstrated in which, due to NBTI stress for a duration of approximately 1000s, there is a frequency shift of about 0.2%. By removing the NBTI stress i.e. the voltage from the core ring oscillator, a recovery of about 0.1% which is about 50% of the frequency shift has been observed .

On the other hand, silicon is still the most commercially used material for integrated circuits because it is cheap and readily available. Taking all these facts into account, we propose an idea for logic in-field repair to avoid incidents such as fatal shut-down in processors and to increase the availability factor by providing logic spare-blocks using within-chip cold swapping. Our scope will be on the logic parts of a design rather than the memory blocks such as RAM cells, since we can protect the memory blocks using available ECC methods.

In the next section we present an architectural solution for increasing the reliability of processors using logic spare-blocks. Here the key point is trading area for reliability and trying to keep the power and the performance overheads as low as possible while keeping the processor running even in the occurrence of a permanent fault.

5.2.2 In-Field Logic Repair

Hard-errors and defects that disable the system from executing its applications can be fatal if the whole system shuts down. In case a set of failures disables the system from carrying out all applications, a subset of less important applications can be dropped while the more important applications can be kept alive. This concept is known as graceful degradation [143] [144]. The main idea is providing logic spare-blocks to the architecture in such a way that in the occurrence of any permanent faults or defects, the faulty logic block can be replaced by the spare-block as depicted in Fig. 5.2 to maintain the same functionality or with graceful degradation to preserve the vital functions of the faulty logic block.

5.2.3 Sphere of Replication and Levels of Granularity

The sphere of replication determines the logical boundary within which logic blocks are physically replicated. The size and the level of granularity of the sphere of replication can vary widely, and this level of granularity will lead to *implementation complexity* vs. *availability* trade-offs. A spare-block can vary from a simple logic gate to a whole processor core; a spare-block can be an exact replica of a logic block or a functionally equivalent structure of that logic block but with a different canonical form or physical implementation.

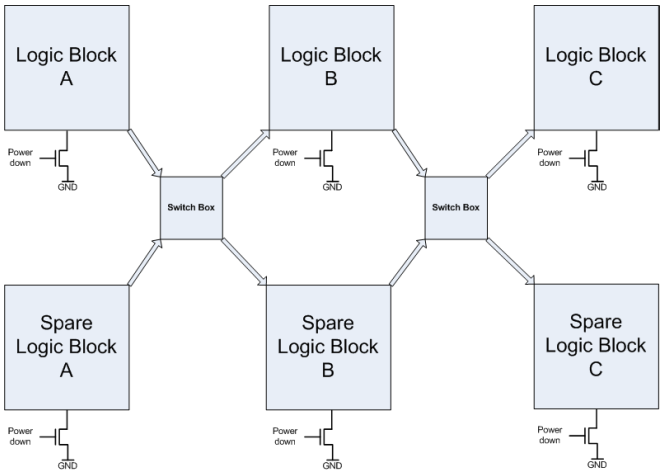

Figure 5.2: General idea of logic in-field repair: Spare logic blocks can be exact replicas of relevant logic blocks or functionally equivalent structures, or even a simplified version of the logic blocks with reduced functionalities.

It can be observed that about 80% of an ASIC chip uses less than 20% percent of the available cells in a given cell library. From a Boolean logic perspective, any arbitrary function can be implemented with 2-input NAND or NOR cells, but for power, performance and area optimization, a variety of cells for implementing the same logic functions are available in a given cell library. For example, any function implemented by NAND-NOR cells can be replaced the AOI(AND-OR-INVERT), OAI(OR-AND-INVERT) cells to save area and power, especially in CMOS technology. Because AOI/OAI cells consumes less transistors compared to NAND-NOR cells, this can result in lower fabrication costs. For instance, a 2-input CMOS AOI cell can be implemented by 6 transistors, comparing to the NAND-NOR equivalent which used 10 transistors. Therefore, any given logic block can be implemented using different cell varieties, all having the same functions.

A spare-block can also be a simplified version of a certain logic block that is smaller in size and provides a reduced level of functionality or service on the occurrence of

a permanent fault or defect on the main logic block rather than failing completely. In other words, the spare-blocks can also be designed to be used in the graceful degradation phase if any permanent fault or defect happens.

Obviously the easiest approach could be taking the whole processor core as a spare-block in the replication procedure. Such system would appear as a dual core processor, with one always-off core and one always-on core. In the occurrence of a permanent fault on such a dual core processor, the faulty core should be shut down and the other core should be powered up. However doing this at Core-level would cost many clock cycles due to initialization and re-execution of several instructions while switching between spare-blocks at pipeline level would cost just a few clock cycles hence the system would have a higher availability and lower downtime, particularly in case of safety-critical or non-stop computing applications. Therefore, in this work, we define the logic spare-blocks as pipeline stages in a processor core in favour of increased availability and decreased downtime.

It is noteworthy to mention that, there is a huge difference in power consumption between the proposed method and N-modular redundancy methods such as DMR or TMR. Here only the main core, the switches and the controller are always "on" and the spare-blocks are turned on only when they are needed. By the time a logic spare-block is turned on, its faulty/defective counterpart will be turned off to keep the power consumption overheads as low as possible. Moreover because the spare-blocks are always "off", they will be immune to the ageing, NBTI and HCI impacts. The down-time for the system will also be much lower than replacing or swapping a faulty chip with a new one, since all the replacements will happen within the chip.

Since the switching between the logic blocks is happening in a within-chip fashion, the delays and the downtime of such system is limited to a few clock cycles. Therefore the reliability of such system can be considered as equal to the reliability of stand-by redundancy systems with two redundant components. However the power overheads of such in-field repair system is significantly lower that any stand-by or TMR systems, thanks to the within-chip cold swapping mechanism.

The reliability of such in-field repair system R_{IFR}, with s number of spare-blocks and R_b as the reliability of each block, whether an active block or a spare-block -assuming that the fault detection and the switching mechanisms are flawless-, can be expressed as:

$$R_{IFR} = 1 - (1 - R_b)^{(s+1)} \tag{5.3}$$

R_{IFR} is an increasing function of the number of spare-blocks. However too many spare-blocks can have a detrimental effect on the reliability of such a system, accompanied by intolerable overheads. Table 5.1 and Table 5.2 show an example of the reliability of the IFR system, once assuming block reliability of 1, and increasing the number of spare blocks and another time assuming block reliability of 0.7, and increasing the number of spare blocks (assuming flawless fault detection and switching mechanisms). In an ideal world, that the spare blocks are 100% reliable, increasing the number of spare block only adds up to the overhead by 100%. In more practical cases, there is a tradeoff between the level of reliability and the amount of overhead that can be tolerated.

Table 5.1

Number of spare-blocks	Reliability of blocks	Reliability of IFR	Overhead
1	1	1	100%
2	1	1	200%
3	1	1	300%
4	1	1	400%
5	1	1	500%

Table 5.2

Number of spare-blocks	Reliability of blocks	Reliability of IFR	Overhead
1	0.7	0.91	100%
2	0.7	0.972	200%
3	0.7	0.9919	300%
4	0.7	0.99757	400%
5	0.7	0.999271	500%

The Markov model of such scheme can be expressed as Fig. 5.3 with F(t) as the distribution function of reconfiguration transitions from a working state (2) to a working state (3), that is after switching to spare-blocks. State (4) is the death state that represents the total failure of the system.

To evaluate the reliability of such system a program called SURE has been used which is released by NASA Langley Research Center. SURE is essentially a reliability analysis tool that can be used to compute and solve semi-Markov models based on SURE bounding theorem. This theorem states that: "The probability D(t) of a system with the mission time T to enter a particular death state

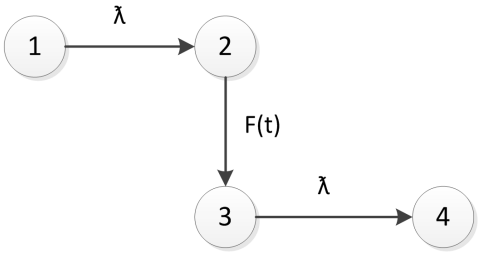

Figure 5.3: Markov model of In-Field Repair system

is bounded as: Lower-bound < D(t) < Upper-bound". The maths behind how SURE computes these, is discussed in [145]. For this work, we are interested in observing possibilities of the upper and the lower bounds of the main states: the operational state and the death state.

The SURE tool is particularly useful in analyzing the fault-tolerance features of reconfigurable systems. To investigate the reliability of mission-critical computer systems, the probabilistic boundaries are computed to be close enough (usually within 5% of each other). According the developers, even for enormous and complex systems, SURE bounding theorems have algebraic solutions which are consequently computationally efficient. Also, SURE can optionally take a specified parameter as a variable over a range of values, which can automate the process of sensitivity analysis of such systems [145].

To analyze the Markov reliability model of the system in SURE, all of the states and the transitions between them should be defined. Usually during early analysis of the processes of a system, there is no experimental reliability data available. So for simplicity, we have to start with some assumptions about the not-exactly-known issues in the process, components, mission-time and the failure rates. For instance, let's assume a TMR system in which, each processor has the failure rate λ over the specified range of 1E-6 to 1E-2 with the mission time equal to 1000 hours.

The SURE tools calculates an upper and a lower bound on the probability of system failure. As mentioned earlier, these bounds are usually within 5% of each other, and thus they usually provide an accurate estimate of system failure. Fig. 5.4 shows the probability of failure of a TMR system as opposed to an IFR system with the same failure rate λ over the same specified range of 1E-6 to 1E-2 with the same mission time equal to 1000 hours which is shown in Fig. 5.5. Because the upper and the lower bounds are very close, the bounds appear as one line in the

plot. It can be seen that the reliability of a TMR system is strongly dependent on the mission time and for longer mission times the IFR system is far more reliable.

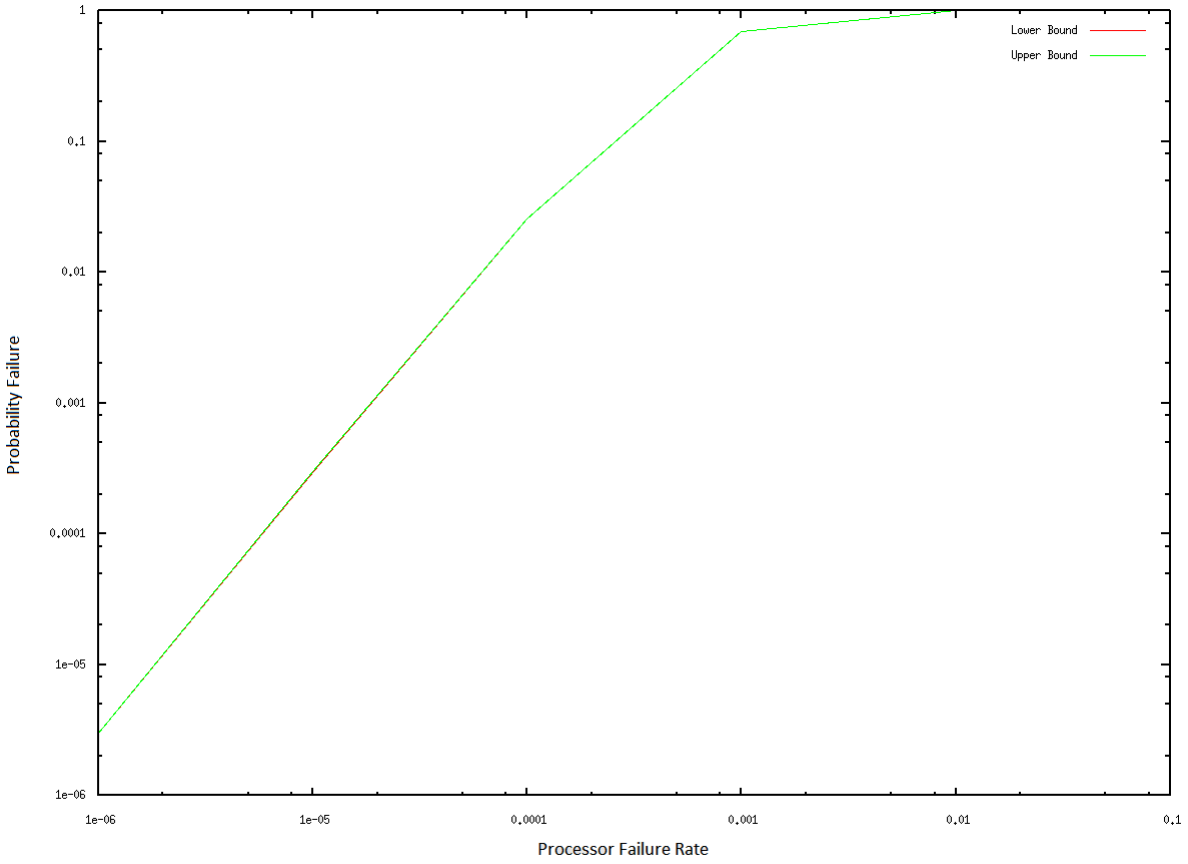

Figure 5.4: NASA SURE plot for TMR systems - Mission time: 1000 hours

In the next section, we discuss the implementation and the reliability gains of the general proposed general IFR architecture on a processor core at pipeline level.

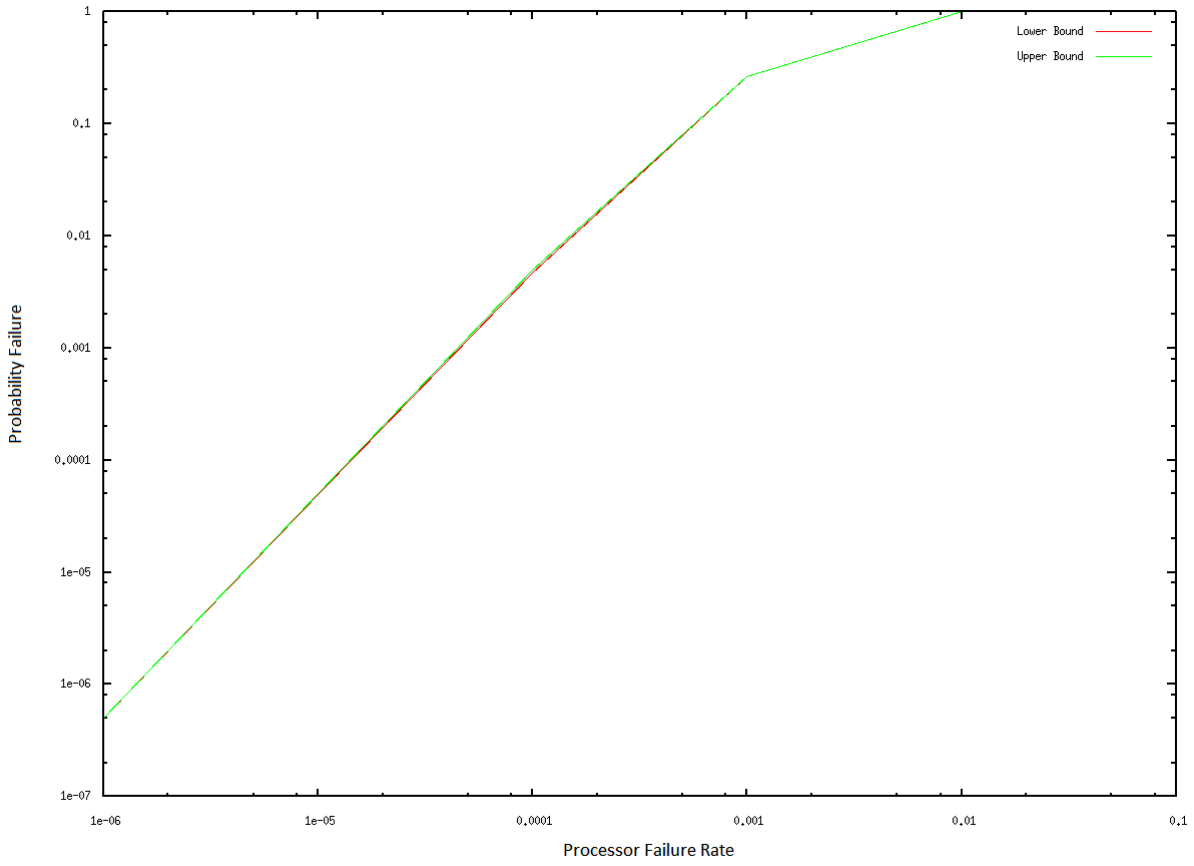

Figure 5.5: NASA SURE plot for the proposed general In-Field Repair system - Mission time: 1000 hours

5.2.4 In-Field Logic Repair at Pipeline Level

To minimize the reconfiguration process time F(t), we have implemented this idea at pipeline level on a processor core at RTL down to GDSII using a conventional ASIC design flow on 45nm Nanngate standard cell library. The RTL-to-GDSII flow is similar to what we used in Chapter 4, section 4.3.

The 32 bit processor core that has been used has a 3 stage pipeline and each pipeline stage is taken as one logic block. Pipelines stages are connected to one another through the switch boxes. The switch boxes have been implemented using multiplexor cells available in the standard cell library. Parity error detection method has been used with one parity bit for every eight bits of signals between the pipeline stages. The parity circuit has been added before the pipeline registers for every pipeline stage. The cost of the error coding and decoding logic is typically amortized over many bits.

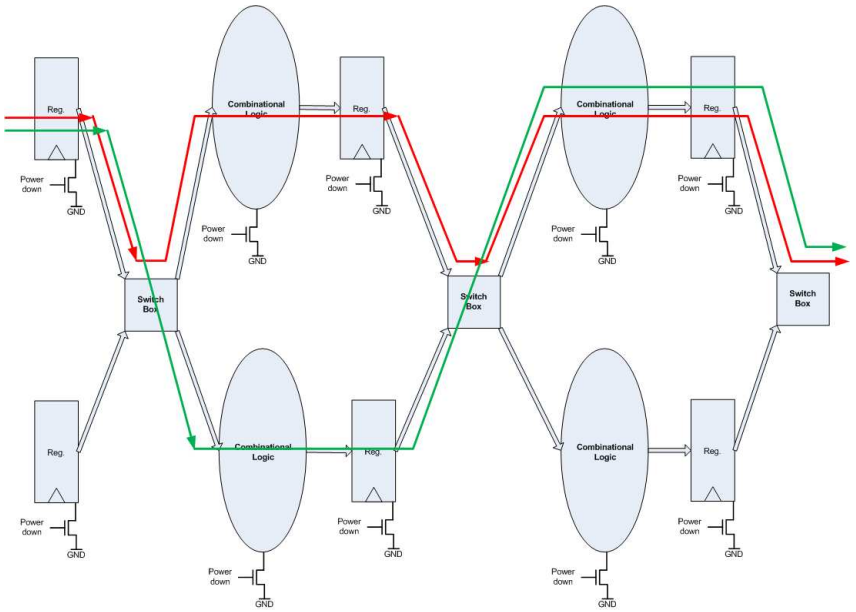

Figure 5.6: Proposed architecture

Note that the idea is to replicate the logic parts and not the register file or any other memory block. The memory blocks are protected by ECC which is the conventional method in such designs. The multiplexer-based switch boxes, the parity error detection circuits and the controller unit have been added to the core at RTL and the power-gating circuit has been added at the place & route stage. There is one parity bit for every eight bits of signals connecting the pipeline stages together. Any parity error lasting for more than a certain number of clock cycles is considered a permanent fault. Hence the faulty pipeline stage will be shut down, the pipeline will be flushed and the spare pipeline block will be turned on by the controller unit. The controller unit is also in charge of power management to avoid IR drop and simultaneous switching capacitance by turning on the blocks in a daisy chain style, hence avoiding any in-rush current. The controller has been implemented using the two-rail checker scheme.

To implement the proposed architecture, the duplication of the pipeline stages, switch-boxes, and the design of the controller have been done at Register Transfer level down to gate level. The core itself is designed in VHDL, the additional circuits such as parity circuits, the switch-boxes, the controller and the top level files are defined in Verilog and the power gating scheme and the power switches have been added at place & route stage.

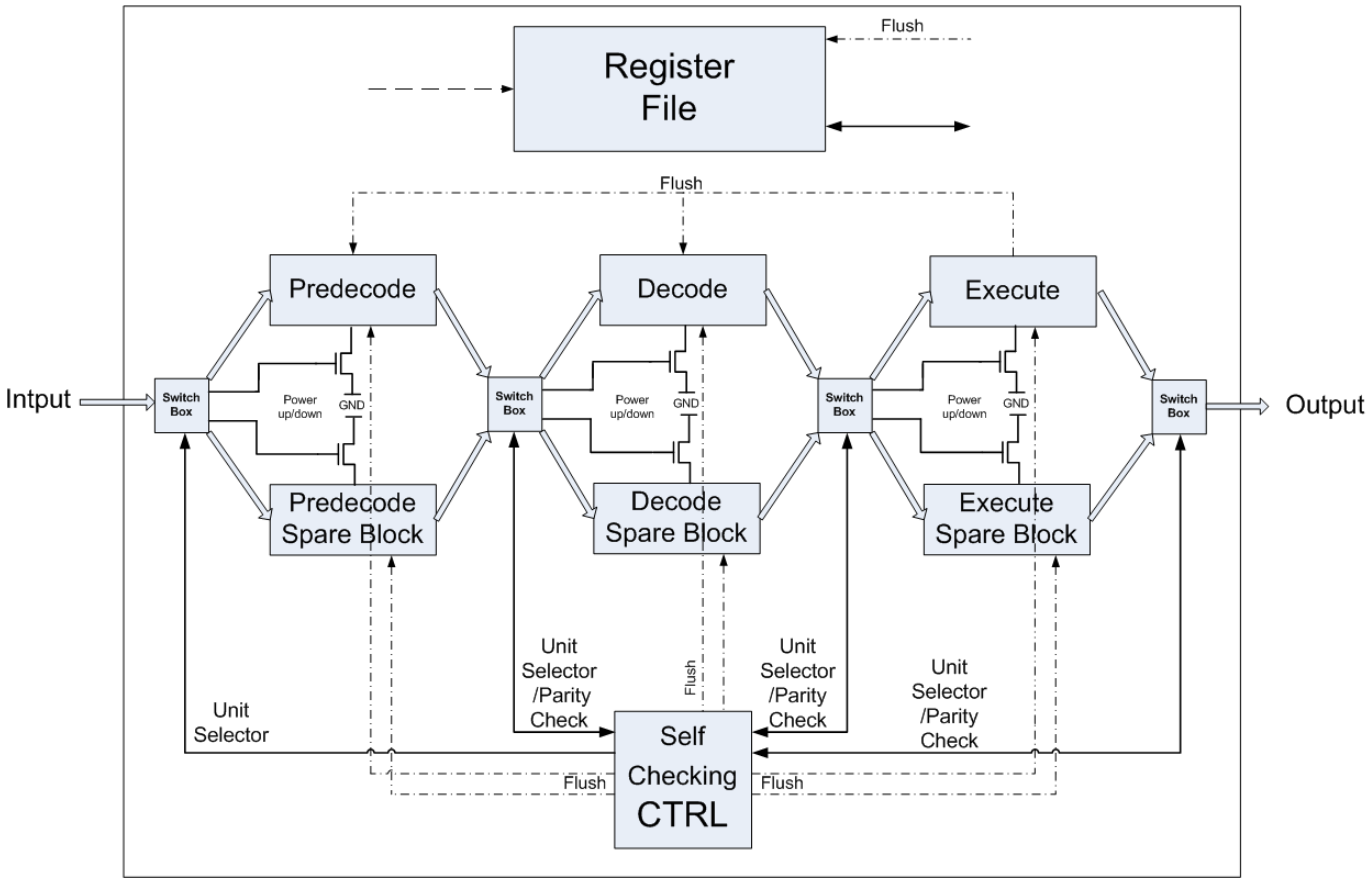

Figure 5.7: In-field repair architecture: Main pipeline blocks, spare-Blocks, the switch boxes and the controller.

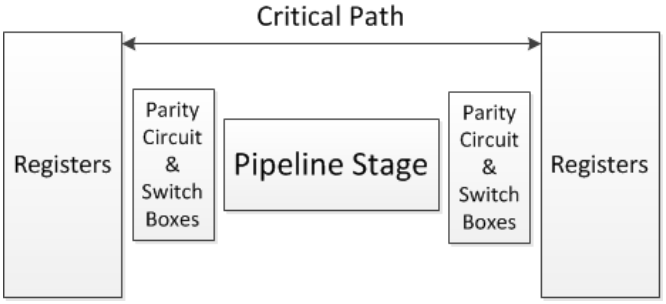

Figure 5.8: critical path

5.2.4.1 Switch Boxes

The switch boxes are essentially comprised of two 2×1 multiplexor cells for each bit, Fig 5.9. They are added as a VHDL structural block and instantiated in the RTL code connecting the pipeline blocks with one another.

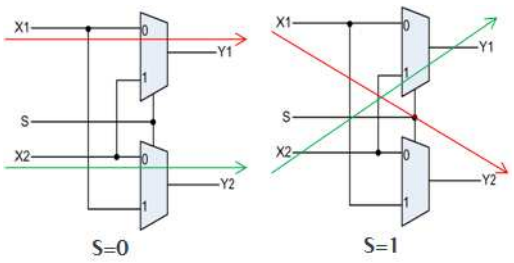

Figure 5.9: A 2-way switch for a single bit (costs: 20 transistors using 45nm cell library, almost equal to a D-flip-flop in size).

5.2.4.2 Error Detection Mechanism

Parity error detection has been added before and after the pipeline registers. Every eight bits of the registers is protected by one parity bit.

5.2.4.3 Self-Checking Controller

The controller is designed based on a duplication and comparison scheme as depicted in Fig 5.10. The duplicated copy of the controller has complemented output values. The comparison is done using totally self-checking (TSC) two-rail checkers (TRC). Both copies of the controller functions receive the same inputs.

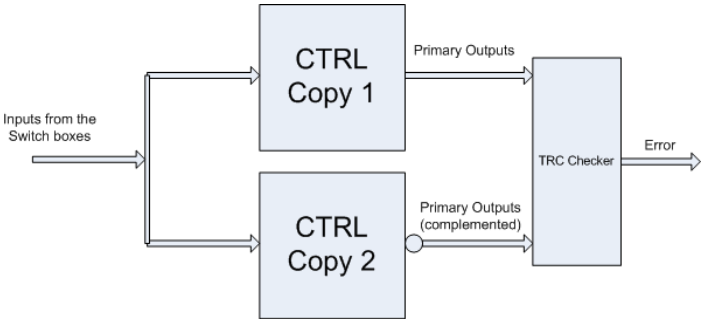

Figure 5.10: Self-checking controller

The problem with power gating schemes is that in the real world, power switches will not necessarily fully charge to supply power to the block or fully discharge to cut the power from the block. So there will be an equilibrium between the sub-threshold leakage of the power gating cell and the leakage current through the power gating cell. This is one of the main reasons we need isolation cells on the outputs of a block which is power gated [23]. This will help to avoid any crowbar-style currents flowing from the outputs of the power-gated block and the inputs of the the block that is connected to such outputs. Isolation cells are usually available

in low-power standard library cells. Each isolation cell has a control input signals that when engaged it will prevent any crowbar current on its output regardless of the fluctuations on its input. The controller provides this isolation control signal.

To differentiate transient faults from permanent faults, two different techniques are used. 1) Adjustable counters in the controller block. Whenever there is an error for more than certain number of clock cycles the error will be assumed to permanent. 2) Due to the fact that transient errors are typically last for less that a clock cycle, a cheap way of designing a fault differentiators circuit at gate-level is depicted in Fig. 5.11, that can easily differentiate permanent faults from transient faults.

For instance, four error/parity bits "ABCD" are fed to this circuit. In the occurrence of an error, if at least any of such bits goes high, the main error indicator "Error" will go high for one clock cycle. This can be due to an SET or an SEU. On the other hand, permanent faults will be more than one clock cycle, hence the resulting Error signal will go high, and stay high for many clock cycles, indicating that a permanent fault has occurred. This is shown in Fig. 5.12 and Fig. 5.13.

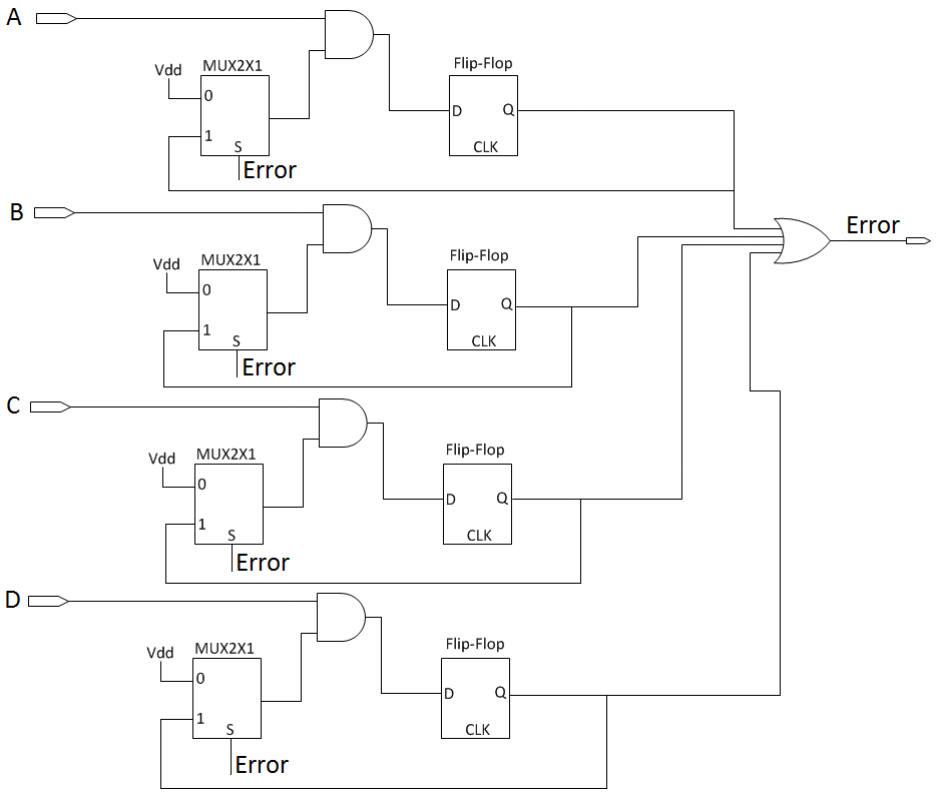

Figure 5.11: Differentiating permanent faults from transient faults

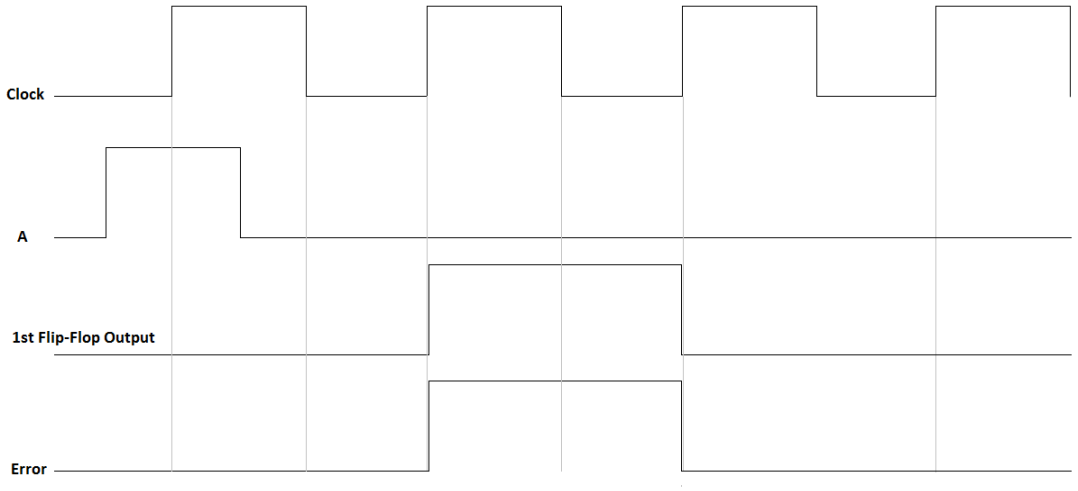

Figure 5.12: Transient fault

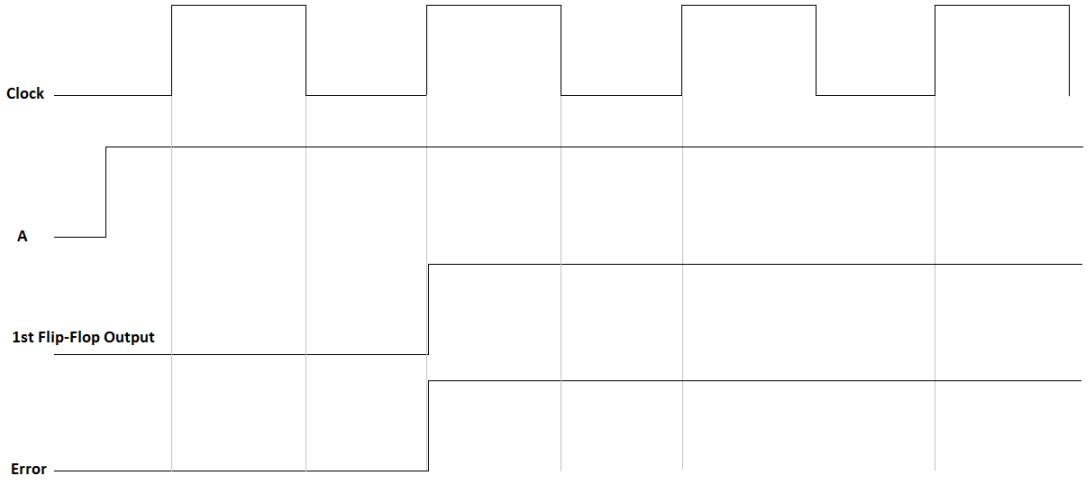

Figure 5.13: Permanent fault

In the occurrence of any permanent delay fault in any logic block (or any pipeline stage in our case), the controller flushes the pipeline, avoiding the ongoing instruction to commit, turns off the faulty block, turns on the spare-block and re-runs the instruction. In this case, the system only loses a few clock cycles rather than a longer period of time to swap the faulty chip with a new one, hence avoiding a total shut-down of the system. The timings of this scheme are depicted in 5.14.

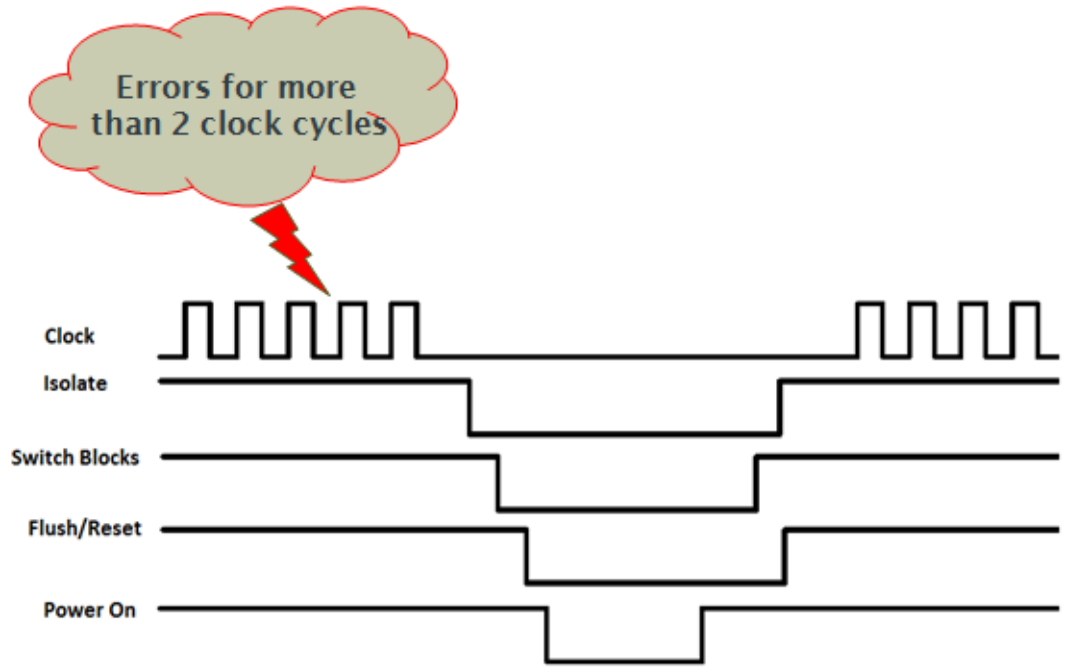

Figure 5.14: An example of the sequence of timings of power management unit based on the methodology in [23]

The reliability of such an architecture is equal to:

$$R_{IFR(pipeline)} = |{R_p}^2 + 2CR_p(1 - R_p)|R_{sw}R_{ctrl} \tag{5.4}$$

where R_p is the reliability of each logic block (pipeline stages in this case) with C as the fault coverage factor. R_{ctrl} is reliability of the controller and R_{sw} is the reliability of the switch boxes. Here the assumption is fault coverage of 100%. From this equation, it is obvious that R_{ctrl} and R_{sw} should not be ignored and they have key impacts on the $R_{IFR(pipeline)}$ as the switched and the controller will always be online, so their reliability factors are multiplied to the whole reliability of this system.

Fig. 5.15 shows the Markov model of this system. The system begins in state (1) where all components are operational. Either of two processors or the switch-boxes and the controller could fail. λ_p is the failure rate of the currently on-line pipeline stages that would make a whole functional core running, λ_p is also considered for the processor pipeline stages which are currently off-line, and λ_{sw} is the failure rate of the switch-boxes with λ_{ctrl} as the failure rate of the controller. Note that the sum of the rates of all failure transitions from state (1) add up to the sum of

the failure rates of all non-failed components ($\lambda_p + \lambda_{sw} + \lambda_{ctrl}$). This property should always be true for all operational states of a reliability model.

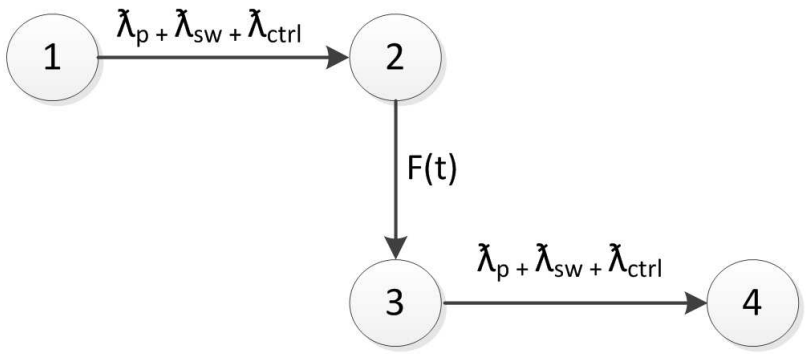

Figure 5.15: Markov model of IFR at pipeline level

5.2.5 Results

5.2.5.1 Overhead Comparisons

As shown in Fig. 5.16, the total area overhead is around 72% on this simple core, because all the logic blocks are duplicated and the controller unit and the switch boxes are also added. Because the spare-blocks are always kept off-line, the dynamic power overhead is less than 18% and the leakage power overhead is 14%. This is caused mainly by the controller unit and the switch boxes. The 9% delay overhead is mainly caused by the error detection mechanism and the switch boxes Fig. 5.8.

In the testing phase of the circuit, the reliability issues and ageing have been modelled as single stuck-at faults and delay faults injected in the main pipeline decode and execute units. The core clock frequency is 100 MHz. A pre-defined set of instructions has been run on both of the simple core and the IFR core. Having the same level of permanent faults in the simple core could result in the fatal shut-down of the core, while the IFR core can get back to its normal functioning status within approximately 1 micro second in this test as shown in Table 5.3.

Due to probable IR drop issues, and since we cannot account for dynamic IR drop accurately, we have to be conservative, so we allowed 50 clock cycles (0.5 us) for the power-on phase and to let the system switch to spare blocks. If the exact IR

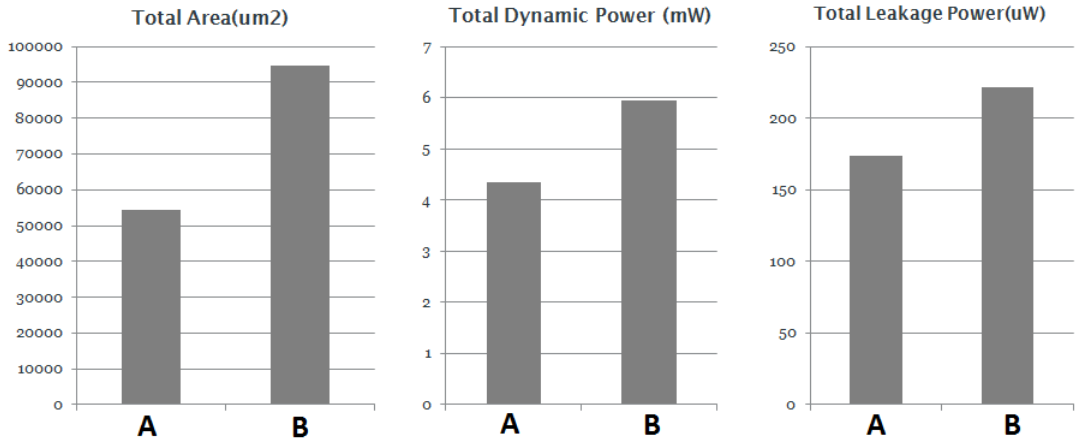

Figure 5.16: Area, power and performance comparisons: A)Simple core B)In-field repair core

Design	Faults	Total Recovery Time (us)
IFR Core	Stuck-at (Decode Unit)	0.82
IFR Core	Stuck-at (Execute Unit)	1.0
IFR Core	Delay (Decode Unit)	1.20
IFR Core	Delay (Execute Unit	1.51

Table 5.3: IFR Core Fault Test

drop behaviour of the system is known, the recovery times in Table 5.3 will be much quicker.

As shown in Fig. 5.17, the contribution of logic components (including all the pipeline stages: predecode unit, decode unit and execute unit) to the total area and power consumption is less than 40% for this specific simple core. It is noteworthy to mention that the relative contribution of logic to the total area is decreasing in modern processors since the sizes of on-chip memory blocks such caches are increasing rapidly, therefore applying this technique to more realistic and modern processors will result in a lower relative area overhead. As reported in ITRS 2011 roadmap report [24], the trend in processor chips shows that the contribution of memory in terms of chip area is predicted to be an order of magnitude higher than logic as shown is Fig. 5.18. This is particularly true for 90nm technology to 16nm and beyond, therefore the area overheads of the proposed technique will be justifiable.

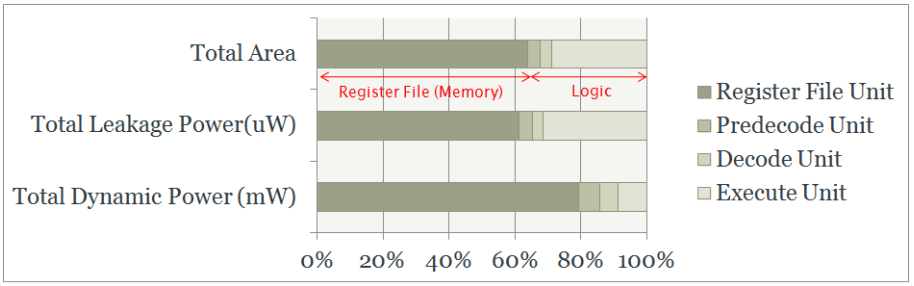

Figure 5.17: Contribution of each processor pipeline stage to the total area and power consumptions.

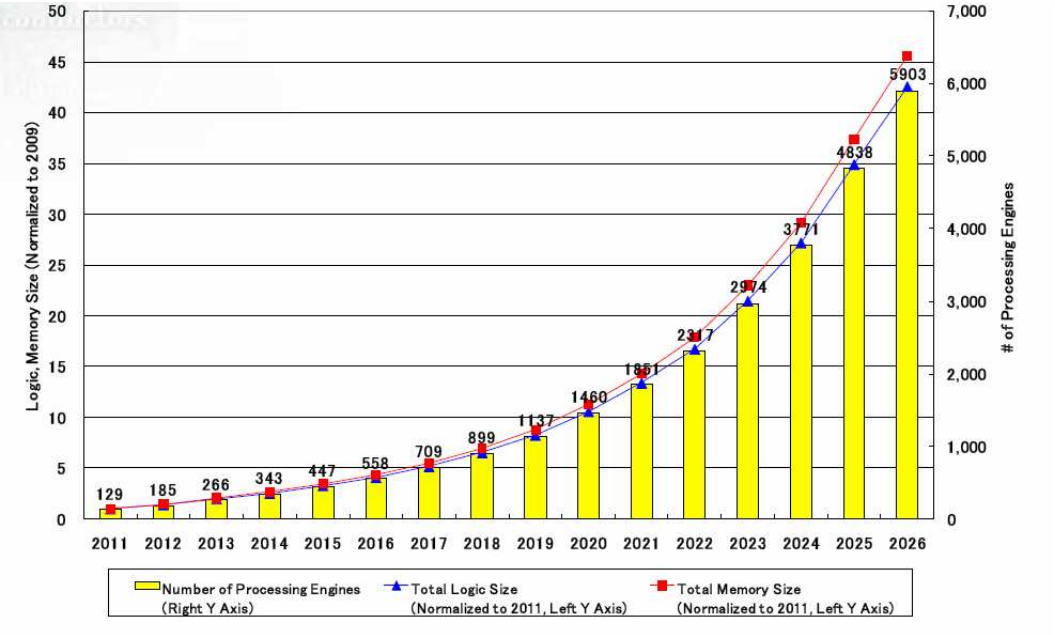

Figure 5.18: ITRS - Logic vs memory roadmap [24]

5.2.5.2 Reliability Comparisons

To compare the reliability of these systems, the SURE program has been used. The probability failure of a simple core with the failure rate of λ over the specified range of 1E-6 to 1E-2 with the mission time 1000 is shown in Fig. 5.19. Using the same failure rate λ and the same mission time, the graph in Fig. 5.20 shows that the probability failure of the IFR system is much lower than that of the simplex core particularly for components with higher failure rates λ's. In other words, for components/processors with higher failure rates, a simplex system has a higher probability to fail within the given mission time as opposed to an IFR-like system. Also considering the given range of 1E-6 to 1E-2 for the failure rates, the probability failure of a simple core starts with 0.001 while the probability failure of the IFR version of the core starts at 1E-6. For failure rates of around $\lambda = 0.001$,

the probability failure of the simple core is very close to 1, however for the IFR core it is approximately 0.3. It can be seen that the IFR core is far more reliable at any failure rates, of course at the cost of tolerating minor power-performance overheads and considerable area overheads.

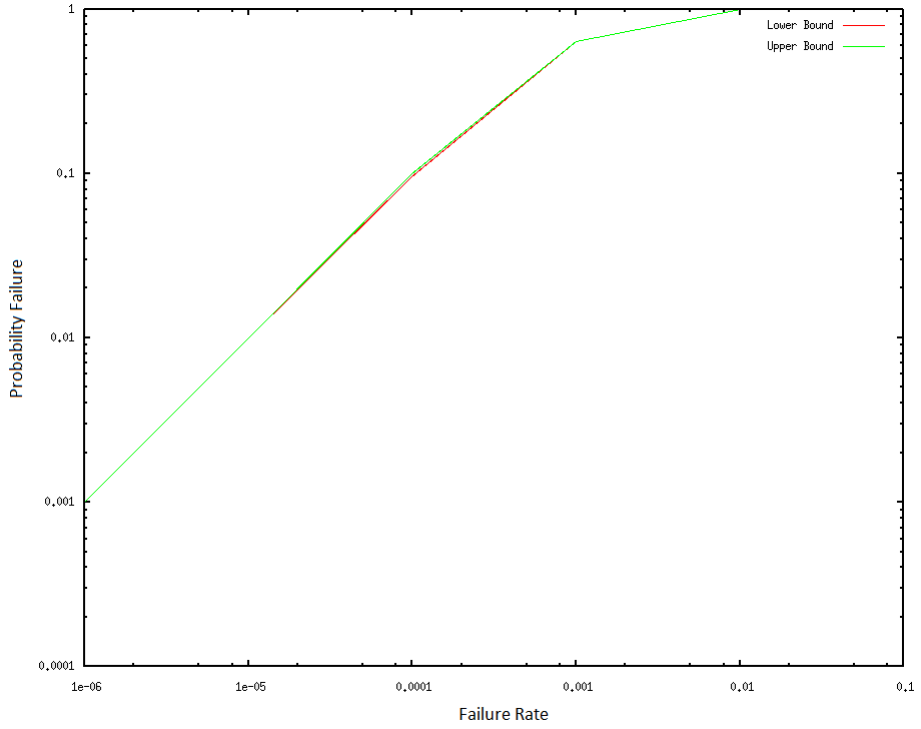

Figure 5.19: NASA SURE plot for probability failure of the simple core

5.3 Challenges in Static Timing Analysis of our proposed schemes

The work in this chapter was an attempt to assess the feasibility of the proposed in-field repair architectures. Care should be taken in Static Timing Analysis (STA) of such architectures as this can be challenging. This is due to the fact that the main blocks and the main timing paths and the spare blocks and spare timing paths do not necessarily age linearly or at exactly the same rates. This means that switching to spare-blocks can cause new hold time violations at run-time, as some paths in a spare-block can be faster than their counterparts in the main-block. These new hold time violations can be deadly to the system functionality just like any other hold time violations in a fully synchronous design. Therefore hold time robustness is a must in such scheme. To address this, two criteria should be specifically considered in functional STA:

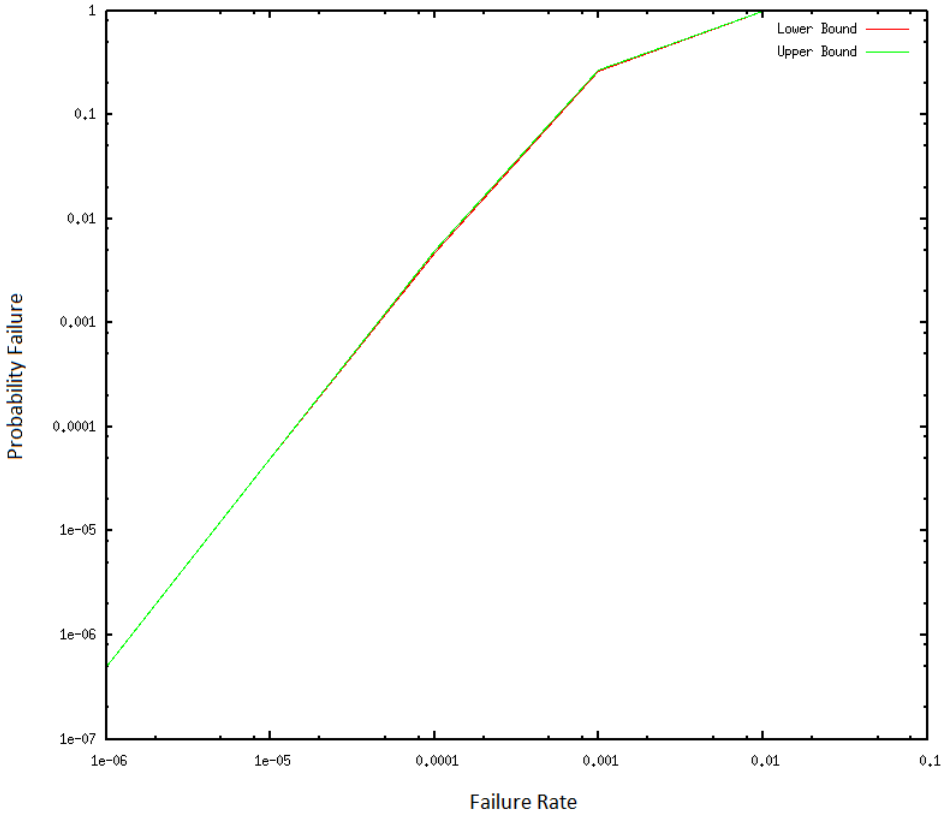

Figure 5.20: NASA SURE plot for probability failure of the IFR core

- Hold check of the design in all of the following modes: All of the main-blocks, all of the spare-blocks and all of the combinations of the main blocks and the spare blocks.
- Addition of extra hold margins on the top hold-time violating paths for all the modes mentioned above.

Therefore the drawback of such architecture is that it creates extra timing scenarios in the order of 2^n with n being the number of the blocks that have spare-block backups in the design. In other words, the functional mode in STA will not be only one mode anymore and it will be expanded to 2^n modes.

For instance, for our proposed in-field repair scheme on the 3-stage pipe-lined processor, n would be the number the pipe-line stages, hence all of the modes, should be analyzed in STA as shown in Table 5.4.

The deeper the pipe-line in a processor is, the more timing scenarios will be created which must be checked specially in terms of hold time violations, so this can be seen as an overhead at design time.

Table 5.4: STA Scenarios of the in-field repair scheme in functional modes with n=3 leading to 8 functional modes

Functional Timing Scenarios	Pre-decode	Decode	Execute
Scenario 1	Main block	Main block	Main block
Scenario 2	Main block	Spare block	Main block
Scenario 3	Main block	Main block	Spare block
Scenario 4	Main block	Spare block	Spare block
Scenario 5	Spare block	Main block	Main block
Scenario 6	Spare block	Spare block	Main block
Scenario 7	Spare block	Main block	Spare block
Scenario 8	Spare block	Spare block	Spare block

5.4 Variation-and-Ageing Resilient Design

It is also possible to take the idea of section 5.4, but instead of applying it to the whole core, the replication happens only on the critical paths in a design. In other words, besides the main group of the critical paths in a design, spare-paths also exist which are kept off-line. This is shown in Fig. 5.21. In the occurrence of any permanent delay-fault due to variation or ageing, instead of a total shut-down of the whole circuit or even slowing it down, we only swap the failed critical path with its spare-path.

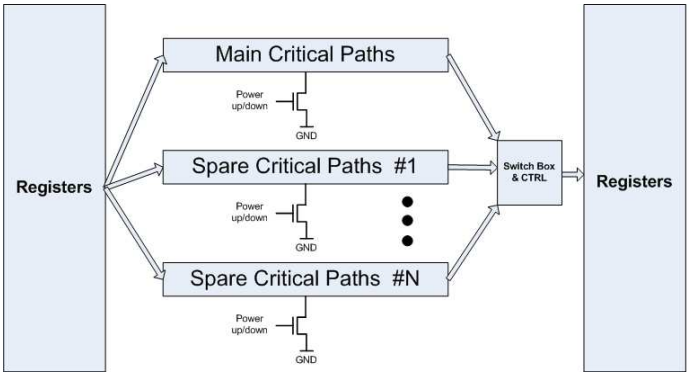

Figure 5.21: Variation-and-ageing resilient design

From the variation point of view, this technique can even increase the yield. At the post-fabrication testing stage, the fastest critical paths can be switched on as the main operating paths among the available spare-paths. The overhead of such scheme is minimal, since only the critical paths are replicated (depending on the number of replications). The additional area and power consumption overheads will be due to the power switches and a controller, very similar to the aforementioned in-field repair idea.

Also, as addressed in [10] and [62], some of the impacts of the aforementioned reliability issues such as delays and threshold voltage shifts can be recovered to their original values. Therefore, during normal operation of the circuit, the threshold voltages of shut-down devices in the turned-off blocks can partially recover to their initial values when the gate bias is switched to 0V. Therefore a partial V_{th} recovery can also be achieved that can alleviate the probability of any delay faults if the initially used critical paths are needed to be turned on again due to ageing-induced delay faults in the secondary group of critical paths.

Hence ideally, the Markov model of such a system can be expressed as Fig. 5.22. We can model the threshold voltage recovery rate with μ which is the repair rate. F(t) is the function of reconfiguration transitions from a working state (2) to a working state (3), that is after switching to spare-critical paths due to ageing-related delay faults in the main group of critical paths.

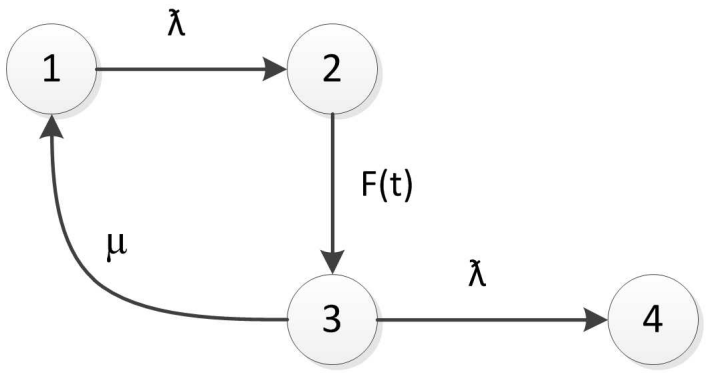

Figure 5.22: Markov model of Variation-and-Ageing Resilient Design

Note that the model in Fig. 5.22 contains a loop because of the repair rate, that is a path that returns to the first state. Loops can lead to infinitely long paths. In theory, evaluating such models will take forever unless a safety value is used. The SURE tools assesses such models using loop truncation automatically. The default truncation level is 25, which will not be reached in most models. The mathematical basis for loop truncation in SURE is given in [145]. The reliability plot of a simple design is the same as Fig. 5.19 and Fig. 5.23 shows the reliability plot of the proposed scheme.

To investigate the overheads of the proposed scheme, a group critical paths from the core has been taken. This group of critical paths has been replicated three times. In other words, each of these critical paths has two replicas, ready and kept off-line. The delay overhead of a such scheme is equal to the proposed in-field

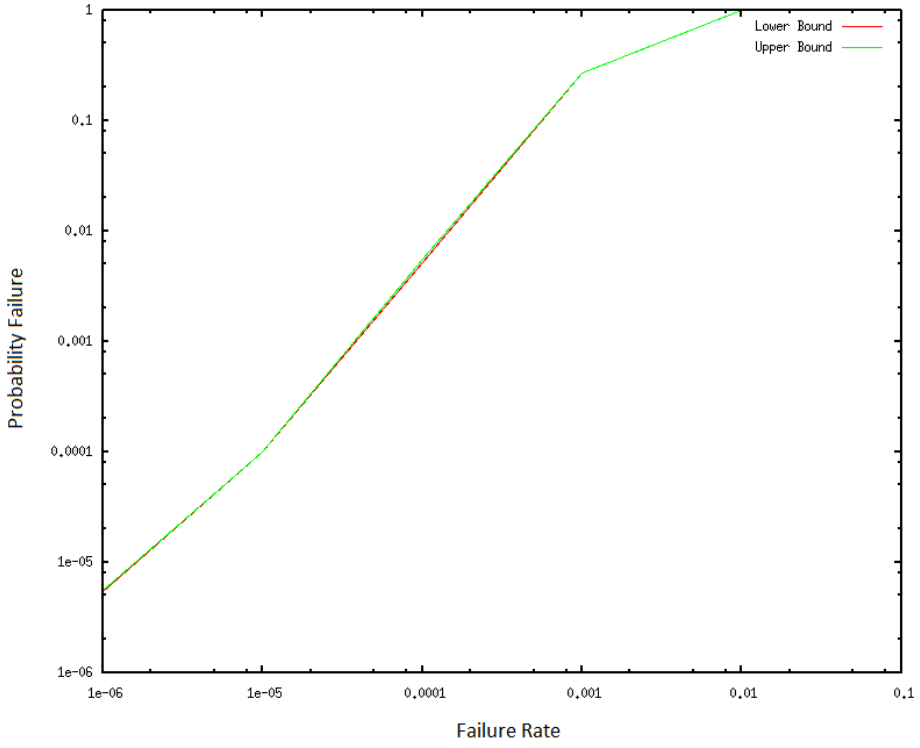

Figure 5.23: NASA SURE plot for probability failure of variation-and-ageing resilient design

repair structure however the area and power overheads are dramatically lower as depicted in Fig. 5.24.

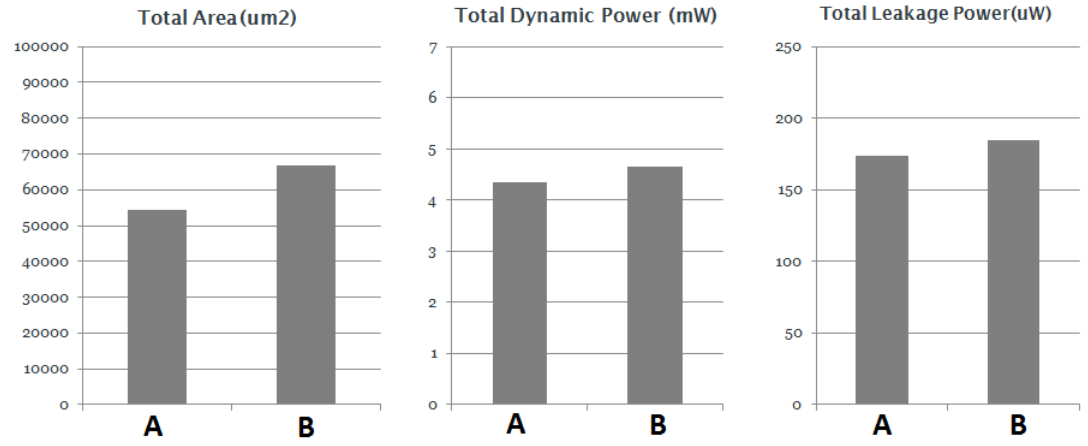

Figure 5.24: Area, power and performance comparisons: A)Simple core B)Variation-and-ageing resilient design core

It is also noteworthy to mention that having different implementations of the critical paths result in different delay variations. For instance, our Monte Carlo simulations at transistor level show that using AND-OR-Invert (AOI) & OR-AND-Invert (OAI) cells to implement the critical paths gives the lowest standard deviation for

delay variations but the highest mean delay as opposed to NAND-NOR implementations of the same paths. This is shown in Fig. 5.25. The delay distribution for AOI/OAI and NAND-NOR circuits is more normal than the default library versions as depicted in the QQ plots in Fig. 5.26. This was observed by running 10k transistor level Monte Carlo simulations adding 10% to 25% variation to the threshold voltages of the transistors in the critical paths. Therefore, by implementing the spare-critical paths using different cell types (and at the cost of greater mean delays), enhanced Normality, and therefore better variation predictability can also be achieved.

5.5 Concluding Remarks

This chapter addresses the last objective "To investigate aging and the reliability issues of UDSM circuits and processors and propose an in-field repair mechanism to avoid fatal shut-downs". In this chapter a logic in-field repair technique for UDSM CMOS processor design has been presented. Here the argument is that the *"MTBF Countdown Clock"* does not start until the device is under stress (i.e. the hardware is powered up) and device degradation does not happen if the devices are permanently "off". Therefore by trading area for increased reliability and by providing spare-logic blocks which are normally off (and will be turned on in case a logic block loses its functionality), any fatal shut-down will be prevented.

We have implemented the proposed technique on a pipe-lined processor core using the conventional ASIC design flow. The simulation results show that by tolerating about 70% area overhead and less than 18% power overhead we can dramatically increase the reliability and decrease the downtime of the processor. As this area overhead is purely logic based, and as trends in microprocessors indicate that logic will require less than one tenth of the silicon area of future chips the area overheads of the proposed technique will be justifiable. The reliability benefits of our proposed architecture have been assessed using Markov models and the NASA SURE reliability analysis tool. The results show that, using the proposed technique the reliability is increased by a factor of x10 to x100 for various component failure rates.

In addition, replication can be applied to only on a select group of critical paths to improve the system's resilience towards variation and ageing-induced delay faults. In turned-off blocks the threshold voltage can partially recover due to the 0V gate bias and this can alleviate the probability of any delay faults if the used logic-blocks or the group of critical-paths are needed to be turned on again due to the same ageing-induced delay faults in the secondary spare-blocks or the secondary group of the critical paths.

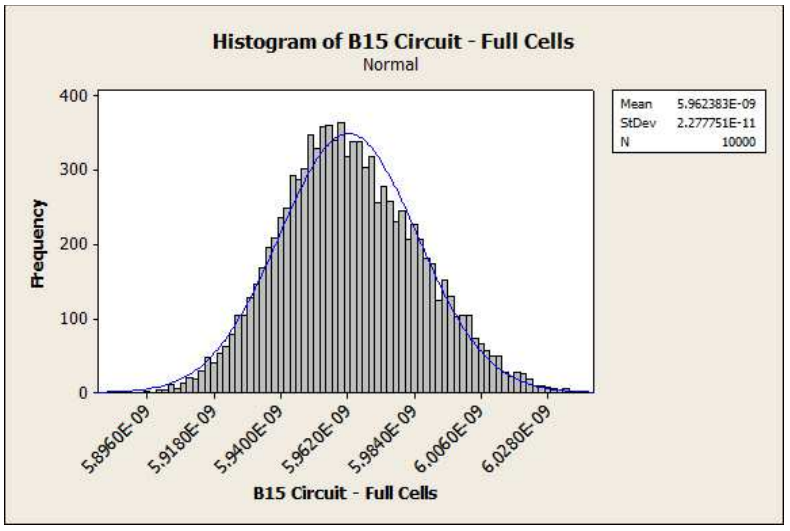

(a) Scenario I

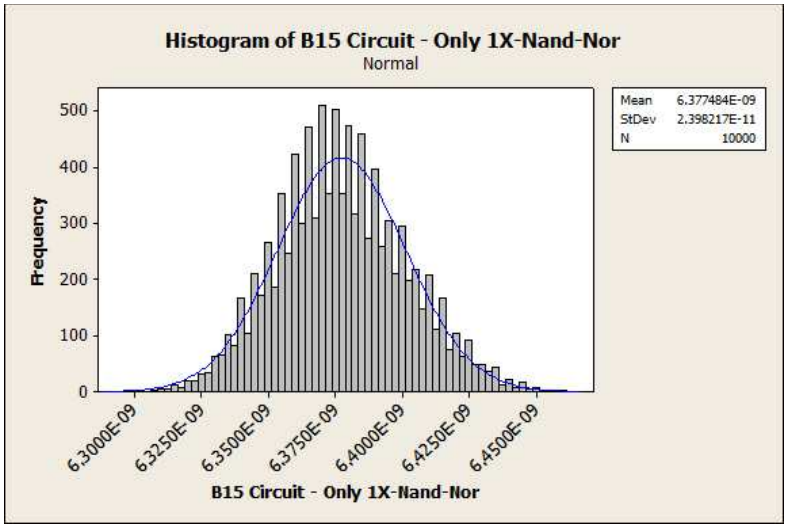

(b) Scenario II

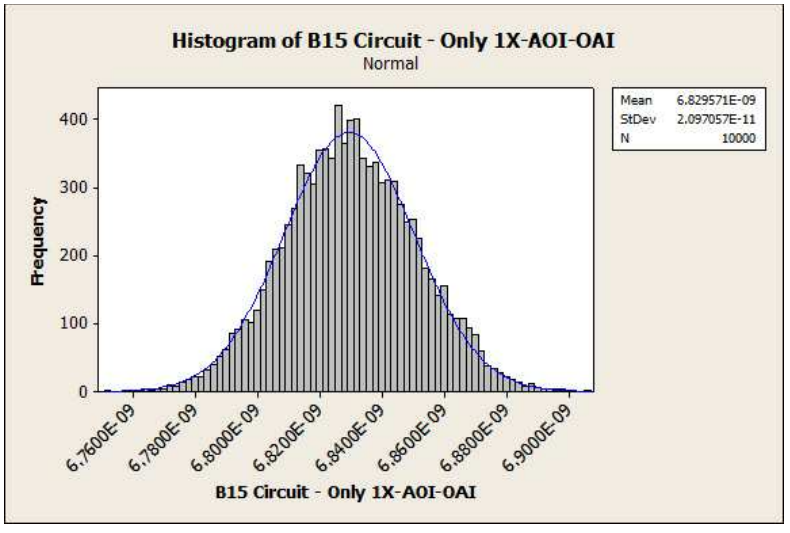

(c) Scenario III

Figure 5.25: Delay distribution of three different scenarios of implementing the B15 benchmark circuit: full Cells, only NAND-Nor, only AOI-OAI

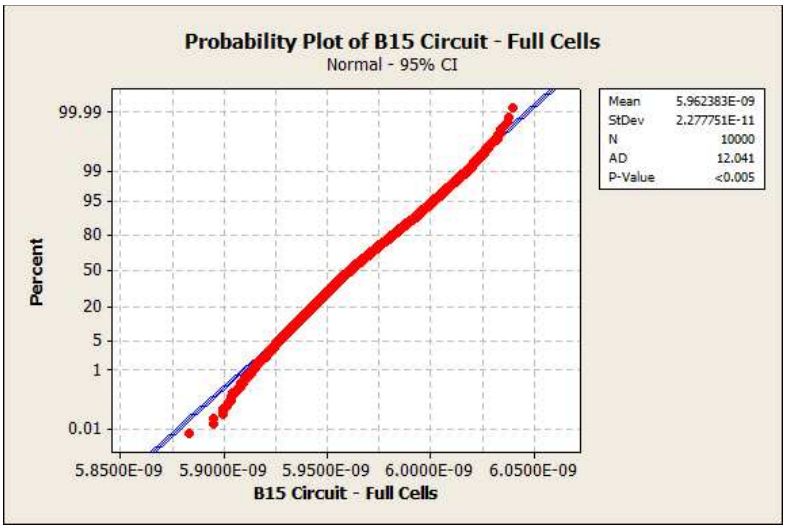

(a) Scenario I

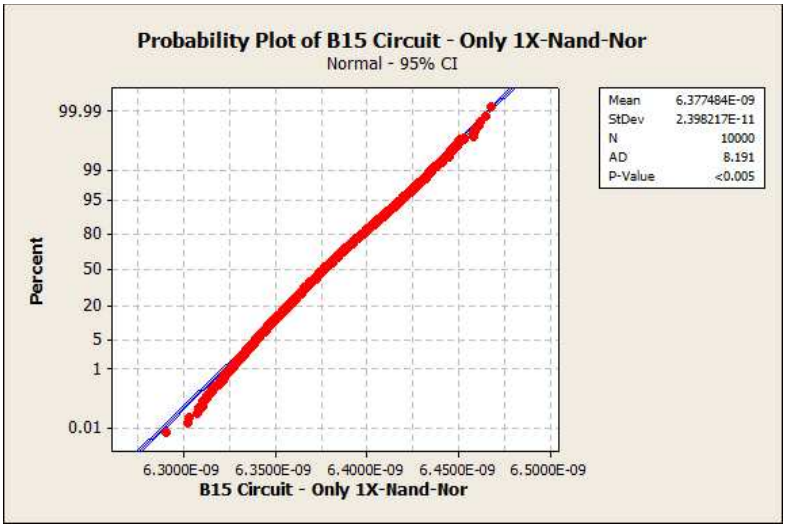

(b) Scenario II

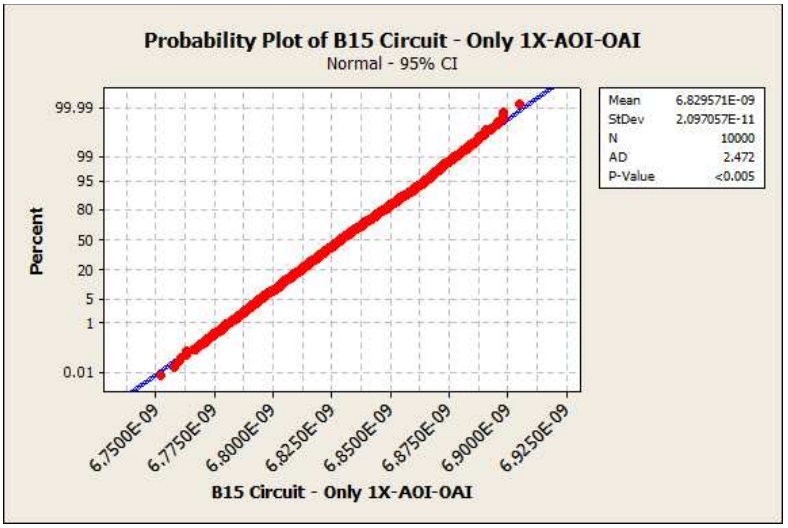

(c) Scenario III

Figure 5.26: Normal distribution predictability of three different scenarios of implementing the B15 benchmark circuit: full cells, only NAND-NOR, only AOI-OAI

Chapter 6

Conclusions

6.1 Conclusions

This thesis provides a survey of various UDSM impacts on circuits and devices, reviewing the ongoing research and providing a summary of the state-of-the-art techniques to mitigate the UDSM impacts.

We have proposed a method to determine the minimum capturable pulse width for the sequential cells which will lead to a more realistic SER computation. As suggested in the literature, the pulse-width of the most common SETs increases for the same radiation environment with technology scaling and demonstrate the increasing importance of combinational logic soft errors. Considering this assumption for 45nm technology and below, there is a high chance of transient pulses being captured by the flip-flops, because the WOV is very narrow at UDSM. For the 130nm technology the minimum captured pulse width is 65 ps. For the 90nm technology, the narrowest capturable pulse is observed to be 56 ps and the narrowest capturable pulse width at 45nm is approximately 34 ps. These values are much less than the defined setup/hold time values in the cell library data-sheets. The results also show that, at 45nm, even in the presence of process variation, the flip-flops are still very susceptible to narrow pulses. Moreover, in circuits with GHZ clock frequencies, this can even lead to multiple bit errors rather than the conventional expectation of single-bit soft errors.

We also observed that interconnect capacitance plays an important role in masking the transient pulses and reducing the SER. Our simulations show that certain amount of capacitance at the output of the struck node flattens the transient

pulse and reduces the pulse amplitude in such a way that the transient pulse cannot be sensed by the next combinational or sequential gate. For instance, our simulations show that for typical 45 nm technology a load capacitance of greater than 20 fF at the output of the struck node flattens the SET. Here the problem is that the exact amount of interconnect capacitance is not available before the place & route stage. For a more realistic calculation of SER, we should also consider interconnect capacitance.

Soft errors induced by radiation, causing malfunctions in electronic systems and circuits, have become one of the most challenging issues that impact the reliability of the modern processors even for sea-level applications. We have presented an implementation of a radiation-hardened 32-bit pipelined processor as well as two novel radiation-hardening techniques at gate-level. We obtained a single-event-upset (SEU) tolerant flip-flop design with 38% less power overhead and 25% less area overhead at 65nm technology comparing to conventional Triple Modular Redundancy (TMR).

To increase the reliability, availability, and lifetime of next-generation circuits, new methods for in-field repair or design for repair or self-healing should be investigated. In each of the cases, an in-field graceful repair action takes place after a mechanism for detecting faults. For general logic circuits, a variety of techniques may be used. In all cases, however, detection of an error requires that some defensive action needs to be taken to recover. This recovery could simply be a graceful shutdown. Hence, a checker and controller is needed.

We have proposed an approach for in-field logic repair. The key idea is trading area for reliability by adding logic in-field repair features to the system while keeping the power and the performance overheads as low as possible. In the case of any permanent faults, logic spare-blocks will replace the faulty blocks on the fly. We have implemented the proposed idea on a pipelined processor core using the conventional ASIC design flow. The simulation results show that by tolerating about 70% area overhead and less than 18% power overhead we can dramatically increase the reliability and decrease the downtime of the processor. The results show that, using the proposed technique the reliability can be increased by a factor of x10 to x100 for various component failure rates.

6.2 Future work

Technology scaling has resulted in significant variation between wafers and dies, as well as increasing levels of within-die variation. This makes the simulation and the analysis of temporal masking to be very difficult. More broadly, increasing variation has complicated accurate timing analysis. Any methodology that can include more accurate and realistic Failure In Time (FIT) analysis and utilizes the usefulness of temporal masking in the presence of variation can potentially reduce the hardware overheads.

An open question will be: Do all the SEUs (bit-flips) really matter? Speculative operations such as: Branch Prediction under Single-Event-Upsets, Memory disambiguation under Single-Event-Upsets or in certain certain blocks can only affect the performance without creating any functional error. Perhaps there is no need to protect those blocks with TMR which results in less overhead. In the field of billion-transistors chips, the winning might go with an 'ART' (Architecturally Radiation Tolerant). In other words, if in a processor with a certain radiation tolerant architecture, Architecturally Correct Execution in the presence of Single-Event-Upsets is achieved, then the savings especially in terms of area and power will be very high particularly in non-safety-critical applications.

Another question is what will happens if there is a fault in the checker or controller or 'who checks the checker?' Triplicating the controller would be absurd, because that would imply a further checker, another further checker to check the previous checker which can continue to infinity. Clearly, therefore, any self-checking or self-repairing system needs a reliable checker/controller that will be able to take appropriate action if a fault develops in the system at large or within the controller itself.

Simultaneous interactions of all of the aforementioned reliability issues, may deteriorate a single-ended mitigation technique. Such chaotic situations will even get worse by further scaling down, power-lowering, and speed-up of semiconductors considering the isolation of process/device engineers from circuit design engineers, leading to some lack of understanding of the impact of their designs upon manufacturability and testability due to the fundamental imitations of technology and device physics. This also necessitates the integration of novel methodologies and EDA tools that can capture the full complexity of these problems for the next generations of reliable circuits and systems.

Appendix A

45nm DFF Schematic

The Schematic of the D Flip-Flop Cell used to determine the WOV in Chapter 4 is provided overleaf.

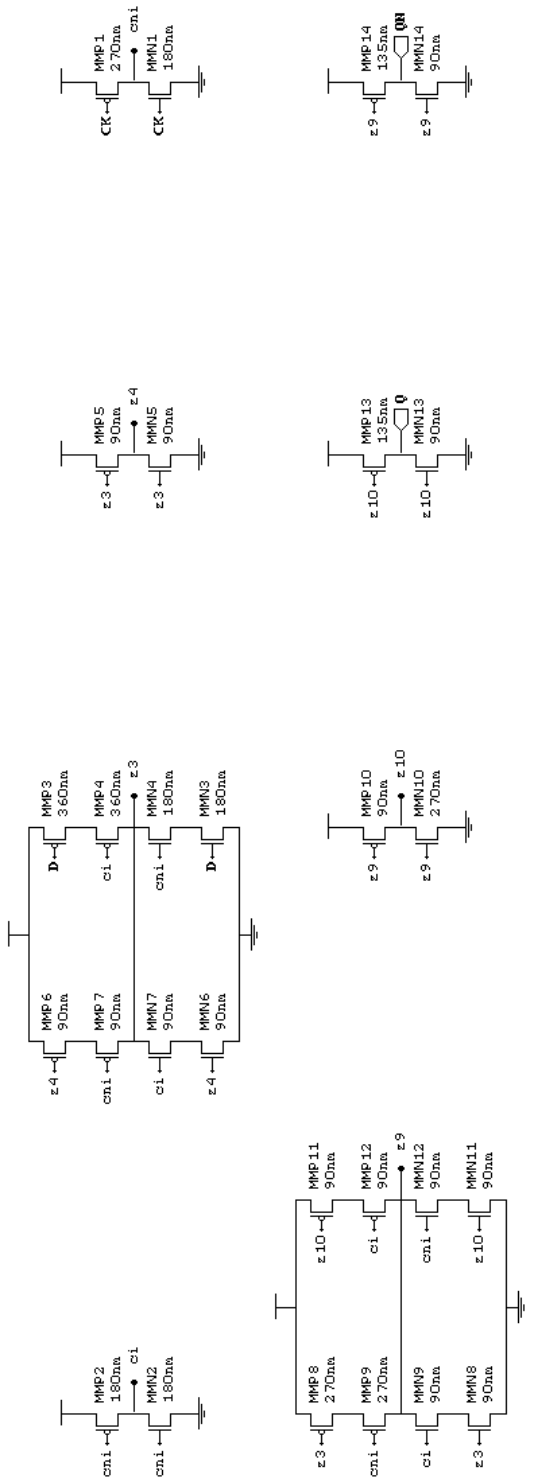

Figure A.1: 45nm DFF - Transistor Level

References

[1] K.J. Kuhn;. Reducing variation in advanced logic technologies: Approaches to process and design for manufacturability of nanoscale cmos. In *Electron Devices Meeting, 2007. IEDM 2007. IEEE International*, pages 471–474, Dec 2007.

[2] K. Bernstein; D.J. Frank; A.E Gattiker; W. Haensch; B.L. Ji; S.R. Nassif; E.J. Nowak; D.J. Pearson; N.J. Rohrer;. High-performance cmos variability in the 65-nm regime and beyond. *IBM Journal of Research and Development*, 50(4.5):433–449, July 2006.

[3] T. M. Mak; A. Krstic; K. T. Cheng; and Li. C. Wang;. New challenges in delay testing of nanometer, multigigahertz designs. volume 21, pages 241 – 248, may-june 2004.

[4] M. Tehranipoor; K. Peng; K. Chakrabarty;. *Test and Diagnosis for Small-Delay Defects*. Springer, 1st edition. edition, September 2011.

[5] P. Mc Crorie;. Ir drop analysis: It's not really necessary, is it? http://www.cadence.com/community/blogs/di/archive/2010/04/05/ir-drop-analysis, accessed on may 10 2012. *Cadence Systems*, 2010.

[6] K. L. Bedingfield; R. D. Leach; M. B. Alexander;. Spacecraft system failures and anomalies attributed to the natural space environment. *NASA reference publication 1390*, August 1996.

[7] R. C. Baumann;. Radiation-induced soft errors in advanced semiconductor technologies. *Device and Materials Reliability, IEEE Transactions on*, 5(3):305 – 316, sept. 2005.

[8] R. Baumann;. Soft errors in advanced computer systems. *IEEE Des. Test*, 22(3):258–266, May 2005.

[9] P. Shivakumar; M. Kistler; S. W. Keckler; D. Burger; L. Alvisi;. Modeling the effect of technology trends on the soft error rate of combinational logic. In *Dependable Systems and Networks, 2002. DSN 2002. Proceedings. International Conference on*, pages 389 – 398, 2002.

[10] D. M. Fleetwood; R. D. Schrimpf;, editor. *Defects in Microelectronic Materials and Devices*. CRC Press, November 2008.

[11] J. Hicks et al.;. 45nm transistor reliability. *Intel Technology Journal*, 12, 2012.

[12] E. G. Friedman I. S. Kourtev; B. Taskin. *Timing Optimization Through Clock Skew Scheduling*. Springer, nov 2008.

[13] A. Wang; S. Naffziger;. *Adaptive Techniques for Dynamic Processor Optimization: Theory and Practice*. Springer, 2008.

[14] D. Ernst; S. K. Nam; S. Das; S. Pant; R. Rao; P. Toan; C. Ziesler; D. Blaauw; T. Austin; K. Flautner; T. Mudge;. Razor: a low-power pipeline based on circuit-level timing speculation. In *Microarchitecture, 2003. MICRO-36. Proceedings. 36th Annual IEEE/ACM International Symposium on*, pages 7 – 18, dec. 2003.

[15] J. Rabaey;. *Low Power Design Essentials*. Springer, 2009.

[16] G. Anelli; M. Campbell; M. Delmastro; F. Faccio; S. Floria; A. Giraldo; E. Heijne; P. Jarron; K. Kloukinas; A. Marchioro; P. Moreira; W. Snoeys. Radiation tolerant vlsi circuits in standard deep submicron cmos technologies for the lhc experiments: Practical design aspects. *IEEE Trans. Nucl. Science*, 46(6):1690–1696, December 1999.

[17] K. Vleugels et al.;. Layout techniques to enhance the radiation tolerance of standard cmos technologies demonstrated on a pixel detector readout chip. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, pages 349–360, January 2000.

[18] R. Velazco; P. Fouillat; R. Reis;. *Radiation Effects on Embedded Systems*. Springer, 1 edition, may 2007.

[19] S. Mukherjee;. *Architecture Design for Soft Errors*. Morgan Kaufmann, 1 edition, February 2008.

[20] M. Lazzaroni;. *Reliability Engineering: Basic Concepts and Applications in ICT*. Springer, 1st edition. edition, September 2011.

[21] P. K. Lala;. *Self-Checking and Fault-Tolerant Digital Design*. Morgan Kaufmann, 2000.

[22] ARM Ltd;. Cortex-r4 and cortex-r4f technical reference manual. *ARM*, 2013.

[23] M. Keating; D. Flynn; R. Aitken; A. Gibbons; K. Shi;. *Low Power Methodology Manual For System-on-Chip Design*. Springer, Jan 2010.

[24] International Technology Roadmap for Semiconductors. International technology roadmap for semiconductors semiconductor industry association, www.itrs.net/links/2011itrs/home2011.htm (2011), accessed on may 1 2012. *Semiconductor Industry Association*, 2012, 2011.

[25] S. Borkar;. Designing reliable systems from unreliable components: the challenges of transistor variability and degradation. *Micro, IEEE*, 25(6):10 – 16, nov.-dec. 2005.

[26] C. Chiang; J. Kawa;. *Design for Manufacturability and Yield for Nano-Scale CMOS*. Springer, 2007.

[27] M. Orshansky; S. Nassif; and D. Boning;. Design for manufacturability and statistical design: A constructive approach. Springer, 1 edition, December 2007.

[28] K. Agarwal; S. Nassif ;. Characterizing process variation in nanometer cmos. In *Proc. of the Design Automation Conf*, pages 396–399, June 2007.

[29] K. Bernstein; D. Frank; J. Gattiker; A. E. Haensch; W. Ji; B. L. Ji; S. R. Nassif; E. J. Nowak; D. Pearson; N. J. Rohrer;. High-performance cmos variability in the 65-nm regime and beyond. *IBM Journal of Research and Development*, 50:433–449, 2006.

[30] M. Orshansky; S. R. Nassif; and D. Boning;. *Design for manufacturability and statistical design: A constructive approach*. Springer, US, 2008.

[31] R. Garg;. *Analysis and Design of Resilient VLSI Circuits: Mitigating Soft Errors and Process Variations*. Springer, 1 edition, nov 2009.

[32] S. B. Duane; S. Nassif;. Models of process variations in device and interconnect. In *Design of High Performance Microprocessor Circuits, chapter 6*. IEEE Press, 1999.

[33] K. Agarwal; S. Nassif;. Characterizing process variation in nanometer cmos. In *Proceedings of the 44th annual Design Automation Conference*, DAC '07, pages 396–399, New York, NY, USA, 2007. ACM.

[34] H. Chang; S. S. Sapatnekar;. Statistical timing analysis under spatial correlations. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 24(9):1467 – 1482, sept. 2005.

[35] K. Bernstein; D. J. Frank; A. E. Gattiker; W. Haensch; B. L. Ji; S. R. Nassif; E. J. Nowak; D. J. Pearson; N. J. Rohrer. High-performance cmos variability in the 65-nm regime and beyond. *IBM Journal of Research and Development*, 50(4.5):433 –449, july 2006.

[36] L. He; A. B. Kahng; K. H. Tam; J. Xiong;. Simultaneous buffer insertion and wire sizing considering systematic cmp variation and random leff variation. *Proc. of the Intl. Conf. on Computer-Aided Design*, 26(5):845–857, May 2007.

[37] K. Cao; S. Dobre; J. Hu;. Standard cell characterization considering lithography induced variations. In *Proc. of the Design Automation Conf*, pages 801–804, 2006.

[38] R. Datta et al.;. Test and debug in deep-submicron technologies. *http://www.cerc.utexas.edu*, May 2009.

[39] D. Chinnery; K. Keutzer;. *Closing the Power Gap between ASIC and Custom: Tools and Techniques for Low Power Design*. Springer, 2007.

[40] Q. K. ZHU;. *Power distribution network design for vlsi*. John Wiley and Sons, Inc, 1st edition. edition, 2004.

[41] E. E. Nigussie. *Variation Tolerant On-Chip Interconnects*. Springer, 2012 edition, December 2011.

[42] R. Velazco; P. Fouillat; R. Reis;. *Radiation Effects on Embedded Systems*. Springer, 1st edition, may 2007.

[43] T. C. May; M. H. Woods;. A new physical mechanism for soft errors in dynamic memories. In *Reliability Physics Symposium, 1978. 16th Annual*, pages 33 –40, april 1978.

[44] Jedec Standard. Measurement and reporting of alpha particle and terrestrial cosmic ray-induced soft errors in semiconductor devices. *Jedec Standard JESD89A*, October 2006.

[45] F. W. Sexton;. Destructive single-event effects in semiconductor devices and ics. *Nuclear Science, IEEE Transactions on*, 50(3):603 – 621, june 2003.

[46] I. W. Gilson; G. M. Vieira; H. N. Egas; F. L. Kastensmidt;. Modeling the sensitivity of cmos circuits to radiation induced single event transients. *Elsevier Microelectronics Reliability*, pages 29–36, January 2008.

[47] A. Pavlov; M. Sachdev;. *CMOS SRAM Circuit Design and Parametric Test in Nano-Scaled Technologies: Process-Aware SRAM Design and Test.* Springer, 1 edition, jun 2008.

[48] R. C. Baumann;. Soft errors in advanced semiconductor devices-part i: the three radiation sources. *Device and Materials Reliability, IEEE Transactions on*, 1(1):17 –22, mar 2001.

[49] R. C. Baumann; E. B. Smith;. Neutron-induced boron fission as a major source of soft errors in deep submicron sram devices. In *Reliability Physics Symposium, 2000. Proceedings. 38th Annual 2000 IEEE International*, pages 152 –157, 2000.

[50] L. W. Massengill; A. E. Baranski; D. O. Van Nort; J. Meng; B. L. Bhuva;. Analysis of single-event effects in combinational logic-simulation of the am2901 bitslice processor. *Nuclear Science, IEEE Transactions on*, 47(6):2609 –2615, dec 2000.

[51] V. Joshi; R. R. Rao; D. Blaauw; D. Sylvester;. Logic ser reduction through flip flop redesign. In *Quality Electronic Design, 2006. ISQED '06. 7th International Symposium on*, pages 6 pp. –616, march 2006.

[52] T. Karnik; P. Hazucha;. Characterization of soft errors caused by single event upsets in cmos processes. *Dependable and Secure Computing, IEEE Transactions on*, 1(2):128 – 143, april-june 2004.

[53] P. E. Dodd; L. W. Massengill;. Basic mechanisms and modeling of single-event upset in digital microelectronics. *Nuclear Science, IEEE Transactions on*, 50(3):583 – 602, june 2003.

[54] B. Narasimham; M. J. Gadlage; B. L. Bhuva; R. D. Schrimpf; L. W. Massengill; W. T. Holman; A. F. Witulski; K. F. Galloway;. Test circuit for measuring pulse widths of single-event transients causing soft errors. *Semiconductor Manufacturing, IEEE Transactions on*, 22(1):119 –125, feb. 2009.

[55] S. Mitra; T. Karnik; N. Seifert; M. Zhang;. Logic soft errors in sub-65nm technologies design and cad challenges. In *Design Automation Conference, 2005. Proceedings. 42nd*, pages 2 – 4, june 2005.

[56] R. Harada; Y. Mitsuyama; M. Hashimoto; T. Onoye;. Measurement circuits for acquiring set pulsewidth distribution with sub-fo1-inverter-delay resolution. In *Quality Electronic Design (ISQED), 2010 11th International Symposium on*, pages 839 –844, march 2010.

[57] M. Nicolaidis;. Soft errors in modern electronic systems. Springer, 1st edition. edition, September 2010.

[58] J. Srinivasan; S. V. Adve; P. Bose; J. A. Rivers;. Lifetime reliability: toward an architectural solution. *Micro, IEEE*, 25(3):70 – 80, may-june 2005.

[59] A. Avizienis; J. C. Laprie; B. Randell; C. Landwehr;. Basic concepts and taxonomy of dependable and secure computing. *Dependable and Secure Computing, IEEE Transactions on*, 1(1):11 – 33, jan.-march 2004.

[60] D. K. Schroder; J. A. Babcock;. Negative bias temperature instability: Road to cross in deep submicron silicon semiconductor manufacturing. *Journal of Applied Physics*, 94(1):1 –18, jul 2003.

[61] A. T. Krishnan; V. Reddy; S. Chakravarthi; J. Rodriguez; S. John; S. Krishnan;. Nbti impact on transistor and circuit: models, mechanisms and scaling effects [mosfets]. In *Electron Devices Meeting, 2003. IEDM '03 Technical Digest. IEEE International*, pages 14.5.1 – 14.5.4, dec. 2003.

[62] M. Denais; C. Parthasarathy; G. Ribes; Y. Rey-Tauriac; N. Revil; A. Bravaix; V. Huard; F. Perrier;. *On-the-fly characterization of NBTI in ultra-thin gate oxide PMOSFET's*. dec. 2004.

[63] M. Ershov; S. Saxena; H. Karbasi; S. Winters; S. Minehane; J. Babcock; R. Lindley; P. Clifton; M. Redford; A. Shibkov;. Dynamic recovery of negative bias temperature instability in p-type metal;semiconductor field-effect transistors. *Applied Physics Letters*, 83(8):1647 –1649, aug 2003.

[64] G. Chen; K. Y. Chuah; M. F. Li; D. S. H. Chan; C. H. Ang; J. Z. Zheng; Y. Jin; D. L. Kwong;. Dynamic nbti of pmos transistors and its impact on device lifetime. In *Reliability Physics Symposium Proceedings, 2003. 41st Annual. 2003 IEEE International*, pages 196 – 202, march-4 april 2003.

[65] J. Tschanz; N. S. Kim; S. Dighe; J. Howard; G. Ruhl; S. Vanga; S. Narendra; Y. Hoskote; H. Wilson; C. Lam; M. Shuman; C. Tokunaga; D. Somasekhar; S. Tang; D. Finan; T. Karnik; N. Borkar; N. Kurd; V De;. Adaptive frequency and biasing techniques for tolerance to dynamic temperature-voltage variations and aging. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pages 292 –604, feb. 2007.

[66] D. Bossen;. Ieee international reliability physics symposium (irps) tutorial notes. *IEEE*, apr 2002.

[67] K. Hess; L. F. Register; W. McMahon; B. Tuttle; O. Aktas; U. Ravaioli; J. W. Lyding; I. C. Kizilyalli;. Theory of channel hot-carrier degradation in mosfets. *Physica B: Condensed Matter*, 272(14):527 – 531, 1999.

[68] C.K. Maiti; T.K. Maiti;. *Strain-Engineered MOSFETs*. CRC Press, US, 2012.

[69] K. C. Wu; M. C. Lee; D. Marculescu; C. S. Chieh;. Mitigating lifetime underestimation: A system-level approach considering temperature variations and correlations between failure mechanisms. pages 1269 –1274, march 2012.

[70] H. Nan; L. Li; K. Choi;. Tddb-based performance variation of combinational logic in deeply scaled cmos technology. In *Quality Electronic Design (ISQED), 2012 13th International Symposium on*, pages 328 –333, march 2012.

[71] H. Miyazaki; D. Kodama;. Tddb lifetime of asymmetric patterns and its comprehension from percolation theory. In *Reliability Physics Symposium, 2009 IEEE International*, pages 814–818, April 2009.

[72] Z. Gan; W. Wong; J. J. Liou;. *Semiconductor Process Reliability in Practice*. McGraw-Hill, 2013.

[73] C. E. Stroud;. *A Designers' Guide to Built-in Self-test*. springer, 2002.

[74] V. Iyengar et al.;. Variation-aware performance verification using at-speed structural test and statistical timing. *Computer-Aided Design*, pages 4–8, December 2007.

[75] J. Liou; A. Krstic; Y. M. Jiang; and K. T. Cheng;. Modeling, testing, and analysis for delay defects and noise effects in deep submicron devices. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 22(6):756 – 769, june 2003.

[76] J. Le; X. Li and P. T. Lawrence;. *Statistical Performance Modeling and Optimization (Foundations and Trend*. Now Publishers Inc, aug 2007.

[77] A. Agarwal; V. Zolotov; and D. T. Blaauw;. Statistical timing analysis using bounds and selective enumeration. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 22(9):1243 – 1260, sept. 2003.

[78] A. Devgan; C. Kashyap;. Block-based static timing analysis with uncertainty. In *Computer Aided Design, 2003. ICCAD-2003. International Conference on*, pages 607 – 614, nov. 2003.

[79] A. P. Chandrakasan; R. W. Brodersen;. Minimizing power consumption in digital cmos circuits. volume 83, pages 498 –523, apr 1995.

[80] V. Kursun; E. G. Friedman;. *Multi-voltage CMOS Circuit Design*. Wiley-Blackwell, aug 2006.

[81] S. Ghosh; P. N. Dai; S. Bhunia; K. Roy;. Tolerance to small delay defects by adaptive clock stretching. In *On-Line Testing Symposium, 2007. IOLTS 07. 13th IEEE International*, pages 244 –252, july 2007.

[82] G. Swaroop; B. Swarup; K. Roy;. Crista: A new paradigm for low-power, variation-tolerant, and adaptive circuit synthesis using critical path isolation. *IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS*, 26(11), 2007.

[83] J. T. Kao; M. Miyazaki; A. P. Chandrakasan;. A 175-mv multiply-accumulate unit using an adaptive supply voltage and body bias architecture. *IEEE Journal of Solid-State Circuits*, 37(11):1545–1554, November 2002.

[84] K. Usami et al.;. Automated low-power technique exploiting multiple supply voltages applied to amedia processor. *IEEE Journal of Solid-State Circuits-March;*, 33(3):463–472, 1998.

[85] J. M. Rabaey;. *Better than worst case*. http://www.eecs.berkeley.edu/People/Faculty/jan/, Accessed, 2009.

[86] T. Austin;. *Diva*. http://www.eecs.umich.edu/taustin/, Accessed, 2009.

[87] G. V. Varatkar; N. R. Shanbhag;. Energy-efficient motion estimation using error-tolerance. In *Low Power Electronics and Design, 2006. ISLPED'06. Proceedings of the 2006 International Symposium on*, pages 113 –118, oct. 2006.

[88] S. Herbert; D. Marculescu;. Variation-aware dynamic voltage/frequency scaling. In *High Performance Computer Architecture, 2009. HPCA 2009. IEEE 15th International Symposium on*, pages 301 –312, feb. 2009.

[89] itrs. The international technology roadmap for semiconductors. *IEEE Trans*, jun 2012.

[90] S. Kaxiras; M. Martonosi;. *Computer Architecture Techniques for Power-Efficiency*. Morgan & Claypool Publishers, aug 2008.

[91] F. Faccio; K. Kloukinas; G. Magazzu; A. Marchioro;. Seu effects in registers and in a dual-ported static ram designed in a 0.25um cmos technology for applications in the lhc. In *in the proceedings of the Fifth Workshop on Electronics for LHC Experiments*. Cern, September 1999.

[92] H. Casier; M. Steyaert; A. V. Roermund;. *Analog Circuit Design, Robust Design, Sigma Delta Converters, RFID*. Springer, 1st edition edition, September 2011.

[93] T. Calin; M. Nicolaidis; R. Velazco;. Upset hardened memory design for submicron cmos technology. *Nuclear Science, IEEE Transactions on*, 43(6):2874 –2878, dec 1996.

[94] M. Haghi; J. Draper;. The 90 nm double-dice storage element to reduce single-event upsets. In *52nd IEEE International Midwest Symposium on Circuits and Systems*, pages 463–466, August 2009.

[95] R. Velazco; D. Bessot; S. Duzellier; R. Ecoffet; R. Koga;. Two cmos memory cells suitable for the design of seu-tolerant vlsi circuits. *Nuclear Science, IEEE Transactions on*, 41(6):2229 –2234, dec. 1994.

[96] L. Entrena; C. Lopez; E. Olias;. Automatic insertion of fault-tolerant structures at the rt level. pages 48 –50, 2001.

[97] H. Liang; P. Mishra; K. Wu;. Error correction on-demand: A low power register transfer level concurrent error correction technique. *IEEE Trans. Comput.*, 56(2):243–252, February 2007.

[98] N. Kanekawa;E. H. Ibe; T. Suga; Y. Uematsu;. Dependability in electronic systems: Mitigation of hardware failures, soft errors, and electro-magnetic disturbances. Springer, 1st edition. edition, November 2010.

[99] M. Rebaudengo; M. R. Reorda; M. Torchiano; M. Violante;. Soft-error detection through software fault-tolerance techniques. In *Defect and Fault Tolerance in VLSI Systems, 1999. DFT '99. International Symposium on*, pages 210 –218, nov 1999.

[100] N. Oh; E. J. McCluskey;. Error detection by selective procedure call duplication for low energy consumption. *Reliability, IEEE Transactions on*, 51(4):392 – 402, dec 2002.

[101] H. Engel;. Data flow transformations to detect results which are corrupted by hardware faults. In *High-Assurance Systems Engineering Workshop, 1996. Proceedings., IEEE*, pages 279 –285, oct 1996.

[102] S. S. Mukherjee; M. Kontz; S. K. Reinhardt;. Detailed design and evaluation of redundant multi-threading alternatives. In *Computer Architecture, 2002. Proceedings. 29th Annual International Symposium on*, pages 99 –110, 2002.

[103] C. Wang; H. Kim; Y. Wu; Y. Youfeng; V. Ying;. Compiler-managed software-based redundant multi-threading for transient fault detection. In *Code Generation and Optimization, 2007. CGO '07. International Symposium on*, pages 244 –258, march 2007.

[104] D. Bhaduri; S. Shukla; P. Graham; M. Gokhale;. Comparing reliability-redundancy tradeoffs for two von neumann multiplexing architectures. *IEEE Trans. Nanotechnol.*, 6(3):265–279, May 2007.

[105] R. W. Butler; S. C. Johnson;. Techniques for modeling the reliability of fault-tolerant systems with the markov state-space approach. *NASA Reference Publication 1348*, sep 1995.

[106] R. Kastner; A. Kaplan; M. Sarrafzadeh;. *Synthesis Techniques and Optimizations for Reconfigurable Systems*. Springer, illustrated edition edition, November 2003.

[107] R. Woods; K. Compton; C. Bourganis; P. C. Diniz;, editor. *Reconfigurable Computing: Architectures, Tools, and Applications: 4th International Workshop, ARC 2008, London, UK, March 26-28, 2008, Proceedings ... Computer Science and General Issues)*. Springer, March 2008.

[108] J. M. P. Cardoso; M. Hbner;, editor. *Reconfigurable Computing: From FPGAs to Hardware/Software Codesign*. Springer, 1st edition. edition, August 2011.

[109] M. Platzner; N. Wehn;, editor. *Dynamically Reconfigurable Systems: Architectures, Design Methods and Applications*. Springer, 1st edition. edition, March 2010.

[110] N. Limnios; G. Oprisan;. *Semi-Markov Processes and Reliability*. Birkhauser, March 2001.

[111] J. Pukite; P. Pukite;. *Markov Modeling for Reliability Analysis*. Wiley-Blackwell, June 1998.

[112] I. Wagner; V. Bertacco;. *Post-Silicon and Runtime Verification for Modern Processors*. Springer, 2011.

[113] L. Wissel; D. F. Heidel; M. S. Gordon; K. P. Rodbelland; K. Stawiasz; K. Cannon. Flip-flop upsets from single-event-transients in 65 nm clock circuits. *IEEE Transactions on Nuclear Science*, 56(6):3145–3151, December 2009.

[114] P. E. Dodd; L. W. Massengill;. Basic mechanisms and modeling of single-event upset in digital microelectronics. *Nuclear Science, IEEE Transactions on*, 50(3):583 – 602, june 2003.

[115] R. Velazco; P. Fouillat; R. Reis;. Radiation effects on embedded systems. Springer, 1 edition, may 2007.

[116] R. R. Rao; D. Blaauw; D. Sylvester;. Soft error reduction in combinational logic using gate resizing and flipflop selection. In *Proceedings of the 2006 IEEE/ACM international conference on Computer-aided design*, ICCAD '06, pages 502–509, New York, NY, USA, 2006. ACM.

[117] N. Seifert; N. Tam;. Timing vulnerability factors of sequentials. *Device and Materials Reliability, IEEE Transactions on*, 4(3):516 – 522, sept. 2004.

[118] E. Salman; E. G. Friedman; A. Dasdan; F. Taraporevala; K. Kucukcakar;. Pessimism reduction in static timing analysis using interdependent setup and hold times. In *Proceedings of the 7th International Symposium on Quality Electronic Design*, ISQED '06, pages 159–164, Washington, DC, USA, 2006. IEEE Computer Society.

[119] N. Weste; D. Harris;. *CMOS VLSI Design: A Circuits and Systems Perspective*. Addison Wesley, 3 edition, may 2004.

[120] V. Stojanovic; V. G. Oklobdzija;. Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. *Solid-State Circuits, IEEE Journal of*, 34(4):536 –548, apr 1999.

[121] J. Rabaey; A. Chandrakasan; B. Nikolic;. *Digital Integrated Circuits*. Prentice Hall, 2 edition, jan 2003.

[122] Synopsys. *Synopsys NCX Liberty User Guide*. Synopsys, 2011.

[123] T. J. Chaney;. Measured flip-flop responses to marginal triggering. *Computers, IEEE Transactions on*, C-32(12):1207 –1209, dec. 1983.

[124] N. Seifert; X. Zhu; D. Moyer; R. Mueller; R. Hokinson; N. Leland; M. Shade; L. Massengill;. Frequency dependence of soft error rates for sub-micron cmos technologies. In *Electron Devices Meeting, 2001. IEDM '01. Technical Digest. International*, pages 14.4.1–14.4.4, Dec 2001.

[125] L. Zeng; P. Beckett;. Soft error rate estimation in deep sub-micron cmos. In *Dependable Computing, 2007. PRDC 2007. 13th Pacific Rim International Symposium on*, pages 210–216, Dec 2007.

[126] R. R. Rao; K. Chopra; D. Blaauw; D. Sylvester;. An efficient static algorithm for computing the soft error rates of combinational circuits. In *Proceedings of the conference on Design, automation and test in Europe: Proceedings*, DATE '06, pages 164–169, 3001 Leuven, Belgium, Belgium, 2006. European Design and Automation Association.

[127] B. Narasimham; B.L. Bhuva; R. D. Schrimpf; L. W. Massengill; M. J. Gadlage; O. A. Amusan; W. T. Holman; A. F. Witulski; W. H. Robinson; J. D. Black; J. M. Benedetto; P. H. Eaton. Characterization of digital single event transient pulse-widths in 130-nm and 90-nm cmos technologies. *Nuclear Science, IEEE Transactions on*, 54(6):2506 –2511, dec. 2007.

[128] E. H. Cannon; M. Cabanas-Holmen;. Heavy ion and high energy proton-induced single event transients in 90 nm inverter, nand and nor gates. *Nuclear Science, IEEE Transactions on*, 56(6):3511 –3518, dec. 2009.

[129] B. Narasimham; M. J. Gadlage; B. L. Bhuva; R. D. Schrimpf; L. W. Massengill; W. T. Holman; A. F. Witulski; K. F. Galloway;. Test circuit for measuring pulse widths of single-event transients causing soft errors. *Semiconductor Manufacturing, IEEE Transactions on*, 22(1):119 –125, feb. 2009.

[130] S. Rezgui; R. Won; J. Tien;. Set characterization and mitigation in 65-nm cmos test structures. *IEEE Trans. Nuclear Science.*, 59(4), 2012.

[131] J. Nedeau; D. King; D. Lanza; K. Hunt; L. Byington;. 32-bit radiation-hardened computers for space. In *Aerospace Conference, 1998 IEEE*, volume 2, pages 241–253 vol.2, Mar 1998.

[132] G.R. Brown; L.F. Hoffmann; S.C. Leavy; J.A. Mogensen; J. Brichacek;. Honeywell radiation hardened 32-bit processor central processing unit, floating point processor, and cache memory dose rate and single event effects test results. In *Radiation Effects Data Workshop, 1997 IEEE*, pages 110–115, Jul 1997.

[133] M. Portolan; R. Leveugle;. A highly flexible hardened rtl processor core based on leon2. *Nuclear Science, IEEE Transactions on*, 53(4):2069–2075, Aug 2006.

[134] R. Fuller; W. Morris; D. Gifford; R. Lowther; J. Gwin; J. Salzman; D. Alexander; D. Hunt;. Hardening of texas instruments' vc33 dsp. In *Radiation Effects Data Workshop (REDW), 2010 IEEE*, pages 5–5, July 2010.

[135] T. J. Slegel; E Pfeffer; J. A. Magee;. The ibm eserver z990 microprocessor. *IBM Journal of Research and Development*, 48:295309, May/July 2004.

[136] C. Webb. z6the next-generation mainframe microprocessor. *Hot Chips*, August 2007.

[137] D. Bernick; B. Bruckert; P. D. Vigna; D. Garcia; R. Jardine; J. Klecka; J. Smullen;. Nonstop advanced architecture. In *in Proceedings of the International Conference on Dependable Systems and Networks (DSN*, page 1221, 2005.

[138] M. P. Baze; J. C. Killens; R. A. Paup; W. P. Snapp;. Seu hardening techniques for retargetable, scalable, sub-micron digital circuits and libraries. *Thirteenth Biennial Single Effects Symposium Manhattan Beach*, April 2011.

[139] J. C. Laprie;. Dependable computing and fault tolerance : Concepts and terminology. In *Fault-Tolerant Computing, 1995, ' Highlights from Twenty-Five Years'., Twenty-Fifth International Symposium on*, page 2, jun 1995.

[140] N. Kanekawa; E. H. Ibe; T. Suga; Y. Uematsu;. Dependability in electronic systems: Mitigation of hardware failures, soft errors, and electro-magnetic disturbances. Springer, 1st edition. edition, November 2010.

[141] R. Mishra, S. Mitra, R. Gauthier, D.E. Ioannou, D. Kontos, K. Chatty, C. Seguin, and R. Halbach. On the interaction of esd, nbti and hci in 65nm

technology. In *Reliability physics symposium, 2007. proceedings. 45th annual. ieee international*, pages 17 –22, april 2007.

[142] A. Bravaix; V. Huard; C. Parthasarathy; C. Guerin; G. Ribes; F. Perrier; M. Mairy; D. Roy; M. Denais. Paradigm shift for nbti characterization in ultra-scaled cmos technologies. In *Reliability Physics Symposium Proceedings, 2006. 44th Annual., IEEE International*, pages 735–736, March 2006.

[143] M. Glass; M. Lukasiewycz; C. Haubelt; J. Teich;. Incorporating graceful degradation into embedded system design. In *Design, Automation Test in Europe Conference Exhibition, 2009. DATE '09.*, pages 320 –323, april 2009.

[144] B. Randell; P Lee; P. C. Treleaven;. Reliability issues in computing system design. *ACM Comput. Surv.*, 10(2):123–165, June 1978.

[145] R. W. Butler; S. Johnson;. Techniques for modeling the reliability of fault-tolerant systems with the markov state-space approach. *NASA Reference Publication*, 1348, 1995.