The Pennsylvania State University

The Graduate School

Department of Computer Science and Engineering

Towards Minimizing the Adverse Effects of Temperature on High Performance Digital Systems

A Dissertation in

Computer Science and Engineering

by

Andrew Jonathan Sylvester Ricketts

Submitted in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

May 2010

The dissertation of Andrew Jonathan Sylvester Ricketts was reviewed and approved* by the following:

Vijaykrishnan Narayanan

Professor of Computer Science and Engineering

Dissertation Advisor

Chair of Committee

Mary Jane Irwin

Robert E. Noll Chair of Engineering

Professor of Computer Science and Engineering

Yuan Xie

Associate Professor of Computer Science and Engineering

Kenan Unlu

Professor of Mechanical and Nuclear Engineering

Raj Acharya

Professor of Computer Science and Engineering

Head of the Department of Computer Science and Engineering

*Signatures are on file in the Graduate School

Abstract

There are wide power variations across a chip leading to temperature variations. High power density components such as register files and arithmetic logic units will trend towards elevated temperatures while other lower power density units will be cooler. Higher temperatures in general decrease the lifetime reliability of digital systems through a number of mechanisms. Furthermore, continued scaling leads to an increase in soft error occurrences. To this end in addition to increasing the error resiliency of register files a proposal that would also lower energy per access is presented. The power modes available to save power also have the effect of lowering the rate of certain degradation. The effect of negative bias temperature instability (NBTI) on SRAM cells under power saving modes is investigated for leakage current, read and hold margins. This is done for both symmetric and asymmetric cells. Noting that temperature variations occurs across a chip, we present techniques to lower its effects on synchronous digital circuits propagation delay variation in general and clock skew in particular.

TABLE OF CONTENTS

LIST OF FIGURES vi

LIST OF TABLES viii

LIST OF ABBREVIATIONS ix

ACKNOWLEDGEMENTS x

Chapter 1 Introduction 1

Chapter 2 Reliable Multiported Register Files 5

- 2.1 Introduction 5
- 2.2 Error types and recovery methods 8
  - 2.2.1 Danger of error propagation 13
- 2.3 Experimental description 14

Chapter 3 Impact of NBTI on Different Power Saving Cache Strategies 19

- 3.1 Introduction 19
- 3.2 Negative Bias Temperature Instability(NBTI)- Background and Modeling 21
- 3.3 SRAM Cells Under NBTI 25
- 3.4 Power Saving Strategies in Caches 29
- 3.5 SNM and WNM Recovery Under Different Cache Configurations 32
- 3.6 Effect of NBTI Under Process Variations 34

Chapter 4 Mitigating thermal effects on clock skew 37

- 4.1 2D Clock networks 37
  - 4.1.1 Introduction for 2D clock networks 37
  - 4.1.2 RLC model for including thermal effects 40
  - 4.1.3 Effect of temperature on clock skew 43
  - 4.1.4 Design of dynamically adaptive clock buffers 47
- 4.2 3D Clock networks 52
  - 4.2.1 Introduction for 3D clock networks 52
  - 4.2.2 Proposed clock tree topologies 55
  - 4.2.3 Clock tree modeling for including thermal effects 57
  - 4.2.4 Results 59

Chapter 5 Conclusion..62

References..65

LIST OF FIGURES

Figure 1. Benchmark-wise per cycle port utilization 6

Figure 2. Energy consumption and area of a register file with increasing access ports 7

Figure 3. Conceptual diagram showing three register file banked configurations 8

Figure 4. Interleaved parity generation 10

Figure 5. Average number of times a register is read with an error before correction 13

Figure 6. The number of bank saves 16

Figure 7. Number of stalls 17

Figure 8. PMOS transistors in the stress and recovery phases of NBTI 21

Figure 9. Shifted threshold voltage of a 32nm and 45nm technology node PMOS transistor versus time for different duty cycles (β) 23

Figure 10. Shifted threshold voltage of a 45nm technology node PMOS transistor versus time for different duty cycles 23

Figure 11. a) Symmetric 6T SRAM cell b) asymmetric 6T SRAM cell for reduced leakage current based on dual-V_{TH} transistors (shaded transistors are high $V_{TH)}$ 25

Figure 12. Different SRAM cells butterfly curves for read SNM measurements 26

Figure 13. Percentage of recovery of read static noise margin for different cache configurations based on 32nm node SRAM cells 32

Figure 14. Percentage of recovery of read static noise margin for different cache configurations based on 45nm node SRAM cells 33

Figure 15. Percentage of recovery of write noise margin for different cache configurations based on 32nm node SRAM cells 33

Figure 16. Percentage of recovery of write noise margin for different cache configurations based on 45nm node SRAM cells 34

Figure 17. Read SNM distribution of different SRAM cells for 32nm technology node....35

Figure 18. Read SNM distribution of different SRAM celllls for 45nm technology node....35

Figure 19. Leakage current distribution of different SRAM cells for 32nm technology node....36

Figure 20. Leakage current distribution of different SRAM cells for 45nm technology node....36

Figure 21. H-tree with node labels....38

Figure 22. Wire dimensions in the (a) first level H and (b) second level H's....41

Figure 23. Increasing skew as clock signal propagates through buffers at a temperature difference of 50^{o}C....44

Figure 24. Skew improvement by increasing width of buffers....45

Figure 25. Thermally adaptive buffer schematic....48

Figure 26. Skew improvement after using the thermally adaptive buffers....50

Figure 27. Example of a 3D integrated circuit....52

Figure 28. H-tree repeated in each layer, fed by a common clock....56

Figure 29. H-tree on single layer with clock propagating through vias....56

Figure 30. Dimensions of the wires and vias in the clock trees....58

Figure 31. Skew improvement for the split profile using adaptive buffers....59

LIST OF TABLES

Table 1. The evaluated recovery options and operations under error sizes 9

Table 2. Energy consumption, latency, and area overheads of SECDED and parity .. 14

Table 3. Register file configurations (without SECDED and parity codecs included but with the bits needed for these codecs included) 14

Table 4. Total area and energy requirements of register file configurations (with SECDED and parity encoders and decoders included) 15

Table 5. Summary of predictive model of |ΔVth| under NBTI 22

Table 6. SNM and WNM Changes due to NBTI for 32nm Technology at 125°C 27

Table 7. Recovery of SNM and WNM for 32nm Technology at 125°C and duty cycle (β)=0.25 27

Table 8 Duty cycles and threshold changes for the considered cases.......................... 29

Table 9. PTM Technology Model Parameters 31

Table 10. Possible skew improvement using buffer selection 46

Table 11. Results of skew improvement after using the thermally adaptive buffers... 51

Table 12. Results of skew improvement after using the thermally adaptive buffers... 61

LIST OF ABBREVIATIONS

ALU- Arithmetic Logic Unit

BSIM- Berkeley Short-Channel IGFET Model

CMOS- Complementary Metal Oxide Semiconductor

ECC- Error Correcting Code

IGFET- Insulated-Gate Field Effect Transistor

MOSFET- Metal Oxide Semiconductor Field Effect Transistor

NBTI- Negative Bias Temperature Instability

NMOS- n-type Metal Oxide Semiconductor

PEEC- Partial Element Equivalent Circuit

PMOS- p-type Metal Oxide Semiconductor

RAID- Redundant Arrays of Inexpensive Disks

RLC- Resistor Inductor Capacitor

SRAM- Static Random Access Memory

SPEC- Standard Performance Evaluation Corporation

SECDED- Single Error Correct Double Error Detect

SNM- Static Noise Margin

SoC- System on Chip

TSMC- Taiwan Semiconductor Manufacturing Company

VTC- Voltage Transfer Characteristics

VLSI- Very Large Scale Integration

WNM- Write Noise Margin

ZTC- Zero Temperature Coefficient

ACKNOWLEDGEMENTS

I would like to express my gratitude to my advisor Dr. Vijay Narayanan for his guidance and patience through these years. He has directed me well throughout this process being a source of guidance. I would also like to thank the members of my committee for their time and suggestions during the thesis process. Their feedback has been accepted and morphed my work into its current form. I would like to thank God for giving me the strength and ability to do this work. The invaluable support of my fellow students helped being a comfort during the late nights. In particular my friends Kevin Irick, Sri Hari Krishna Narayanan, and Charles Addo-Quaye have been with me through difficult times. Finally I would like to thank my family particularly my brother, Daniel Ricketts, and sister, Janique Ricketts. While we spent this time in different places the many phone calls were always appreciated.

Chapter 1

Introduction

Modern high performance microprocessors have been shown to have temperature profiles such that the difference between the coldest and hottest regions exceeds 50°C [1]. The coolest regions tended to be located within the level 2 caches while the hottest temperatures were within the ALU and register files. There is a positive relationship between dynamic power or activity level and temperature. This tends to cause the most active parts of a digital system to be the location of the highest temperature or hotspot. Elevated temperatures are associated with a number of negative effects on digital systems these include increased leakage current, shorter time for time dependent dielectric breakdown, and increased threshold voltage degradation due to NBTI [2, 3, 4].

Packaging solutions including heat sinks, heat spreaders, air cooled, and water cooled solutions have been used over the years to dissipate the heat generated by these systems. There are usually material and economic limitations of these solutions giving an overall thermal budget which chips must stay under to remain competitive and having a reasonable expectation of being reliable over their lifetime. Removing the heat generated is the responsibility of the packaging of the integrated circuit, but an alternative and synergistic way of addressing the temperature issue is to lower the power consumed by the digital system before it results in untenable temperatures.

Multiple circuit centric solutions have been proposed that are mostly classifiable as targeting power reductions from one source or other with the side effect of lower power leading to lower temperatures. These solutions can take the form of runtime active management or could be design time solutions. Within these solutions the target solution can usually be described as targeting either the dynamic and/or static leakage component of the power Equation 1. Some of the runtime solutions include dynamic voltage scaling, dynamic frequency scaling, and varying the transistors' threshold voltage, V_{th}. The design time solutions include the logic design style and choices, reducing the supply voltage, V_{DD}, appropriately sizing the transistors to match the power and timing budget, using multiple V_{DD}, multiple V_{th}, the stacking effect, and pin ordering [5]. Another major classification that combines both static design changes and runtime inactive times disables or otherwise lower the power used in non-active modules [6]. These include clock gating, use of sleep transistors, multiple V_{DD}, variable V_{th}, and input control. Clock gating in particular aims to lower the capacitive load, C_L, on the clock reducing the dynamic power. Sleep transistors disconnects the supply voltage from circuit components lowering the leakage power. Multiple V_{DD} is generally used in cases where timing can be sacrificed for power savings such as portions of memory that will not be accessed for multiple cycles. Alternatively certain signal paths have timing slack. That is multiple signals traveling arrive at differing times. The signal path that determines the overall clock period is termed the critical path. Other signals traveling on other paths can

travel slower by lowering their V_{DD}. Along these paths timing can be sacrificed for power

$$\mathrm{P} = \underbrace{\mathrm{C_L} V_{DD}^{\;2} f}_{\text{dynamic}} + \underbrace{\mathrm{V_{DD}} I_{leak}}_{\text{static}}$$

Equation 1. The power equation

savings without interfering with the overall performance of the circuit.

The dynamic power is proportional to V_{DD}^2 such that a quadratic decrease in dynamic power will occur for a decrease in V_{DD}. Both the capacitive component, C_L, and the frequency component, f, of the Equation 1 are linearly proportional to dynamic power and V_{DD} is linearly proportional to the leakage power. Finally, while it may seem that the leakage current, I_{leak}, is a linear contributor to the static power this current increases super linearly with temperature. According to Equation 2 [2], where $v_t = kT/q$, there is both a quadratic and exponential term of temperature that increases the leakage current. This could potentially set up feedback loop of the increased temperature increasing the leakage current causing increased power which further increases the temperature. This possibility is termed thermal runaway and must be prevented. Recent processors currently have

$$\mathrm{I_{Leak}} = \mu_0 \cdot C_{ox} \cdot \frac{W}{L} \cdot e^{b(Vdd - Vdd0)} \cdot v_t^{\;2} \cdot \left(1 - e^{\frac{-Vdd}{vt}}\right) \cdot e^{\frac{-|vth| - Voff}{n \cdot vt}}$$

Equation 2. The leakage current model

temperature sensors to detect and attempt to prevent thermal emergencies. General solutions are to reduce the operating voltage and/or frequency of the processor to reduce both the static and dynamic sources of power. In the event that no other solution is able to prevent the processor from exceeding the safe temperature range the system is shut down completely ceasing the operation of the processor rather than risking permanent damage [7].

Technology is scaling providing new opportunities and challenges. These include the continued reduction in supply voltage and transistor sizes. These offer reduced noise margins and increased opportunity for transient errors to occur. Low power modes enhance these effects and further decrease the reliability to soft errors of particular memory structures. The register file is one such structure particularly prone to these optimizations due to its SRAM core. The beginning offers a counter to the potential increase in soft errors in register files. Lower power mode optimizations in SRAMS enhance reliability particularly lifetime operation. Power is broken into active and passive components having a number of knobs to be turned based on Equation 1. Local power optimizations increase temperature variations. The clock needs to be able to handle these power variations and as such work towards lowering the temperature induced clock skew is presented. Finally section 5 concludes and offers future research directions.

Chapter 2

Reliable Multiported Register Files

2.1 Introduction

The register file is a critical component of high performance superscalar processors providing buffered communication of register values between producers and consumers of data. As the register file is accessed very frequently, faults occurring in the register file can propagate to the functional units and up the memory hierarchy rapidly which can lead to data corruption and system reliability problems. Transient faults are faults that are temporary and leave no permanent damage to the hardware. A number of sources such as power supply noise, crosstalk, and a radiation particle strike can cause them. The focus of this section is the on the detection and recovery from soft errors occurring in the stored values in register files but in many cases the same system will be able to detect and recover from permanent faults affecting the register file.

Many common instructions across varied instruction set architectures require two input register values and produce one output value. Each distinct input register value requires a register read similarly each output register value requires a register write. Across the SPEC benchmarks [9] the average read and write port utilization is fairly low for high issue machines requiring only a small subset of the available read and write ports for

most cycles as shown in Figure 1.

As the number of ports in a monolithic register file increases, the energy consumption and access latency increase drastically [10]. The size of the multiported memory cell increases approximately quadratically with of the number of access ports [11]. There have been many techniques to lower the energy, access latency, and area of multiported register files. Some of these techniques use a clustered architecture with each cluster containing a subset of the register file and functional units [12, 13, 14] while other techniques use a centralized architecture with banked register file having fewer ports per bank [15, 16]. Highly ported register files tend to consume less power and area when

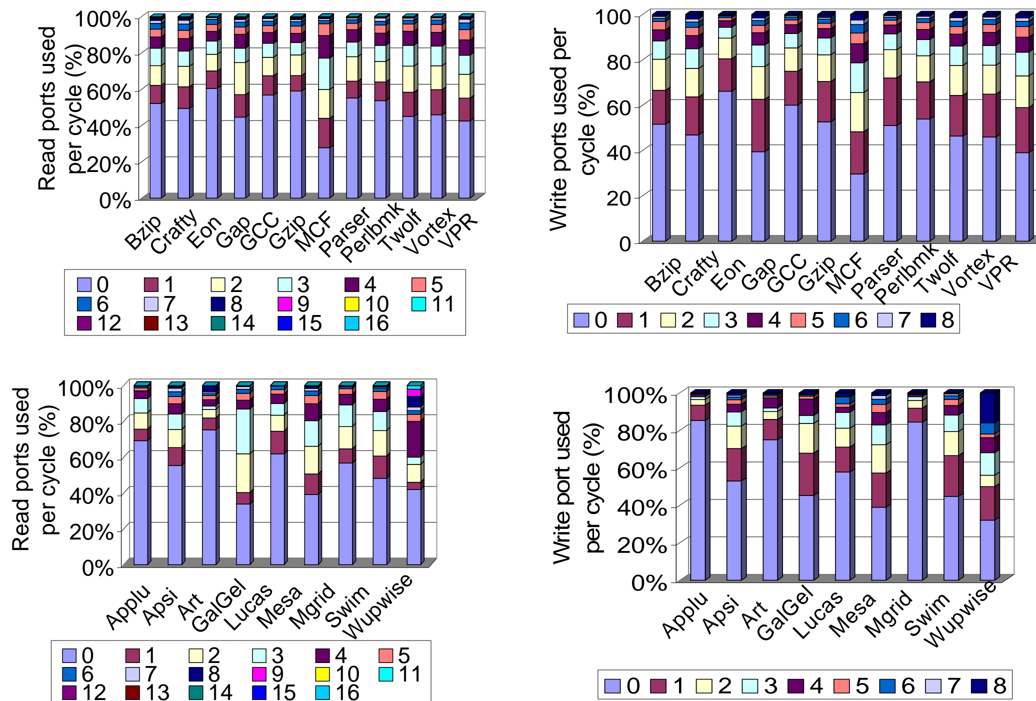

Figure 1. Benchmark-wise per cycle port utilization

designed in a distributed fashion rather than using a single monolithic register file. There could be multiple banks of register files with each bank consisting of a smaller number of ports than the original monolithic configuration. Figure 2 shows the increase in both energy and area required for larger number of ports on a register file with the same number of entries. Figure 3 shows conceptually a method of reducing the number of read

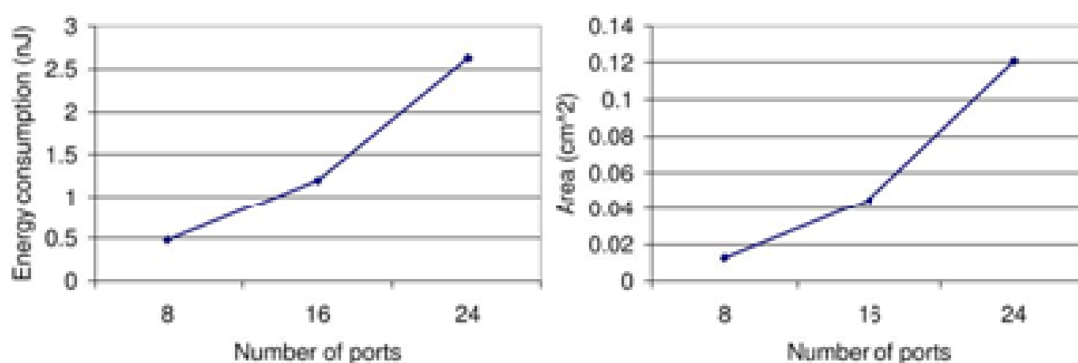

Figure 2. Energy consumption and area of a register file with increasing access ports

ports per bank as the number of banks increase. The number of write ports per bank must remain constant to replicate each write for each bank.

Multibanked register files contain inherent data redundancy. This fact enables the provision of low cost reliable register file operation in the presence of errors using detection and recovery. Recovery is possible because of the distributed register file redundant copies of data, which to our knowledge has gone untapped in previous works.

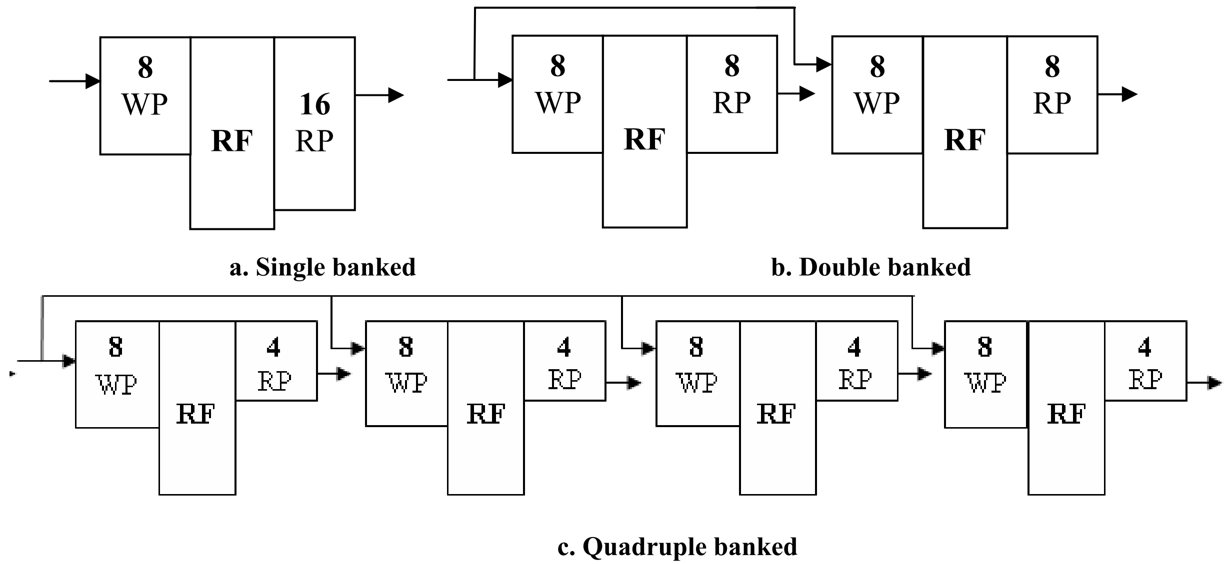

Figure 3. Conceptual diagram showing three register file banked configurations

2.2 Error types and recovery methods

The classification of processor errors in general is divided into two major classes termed as either hard or soft errors. These classifications are due to the longevity of the errors once manifested. Hard errors are permanent damage to the silicon or metal structures in a microprocessor. Soft errors are temporary errors that do not damage the microprocessor. Provided soft errors do not affect a memory element, they will have no lasting effects on operations and settle back to their correct value.

Hard errors could stem from a number of potential sources. Over time the use of the transistors changes their properties with the expectation that hot carrier effects and negative bias temperature instability will become a concern for lifetime reliability of

Table 1. The evaluated recovery options and operations under error sizes

<table>
<tr><td colspan="3"></td><td colspan="3">Size of Consecutive Error</td></tr>
<tr><td></td><td>Config</td><td>Number of Banks</td><td>1 bit</td><td>2 bits</td><td>3 bits</td></tr>
<tr><td rowspan="3">SECDED ECC</td><td>ECC1B</td><td>1</td><td rowspan="3">Detected/ Corrected</td><td>Detected/ Checkpoint</td><td rowspan="3">No guarantees</td></tr>
<tr><td>ECC2B</td><td>2</td><td rowspan="2">Detected/ Read Other Bank/ Write Correct Value</td></tr>
<tr><td>ECC4B</td><td>4</td></tr>
<tr><td rowspan="3">Parity</td><td>SB</td><td>1</td><td colspan="3">Detected/ Checkpoint</td></tr>
<tr><td>MB2</td><td>2</td><td colspan="3" rowspan="2">Detected/ Read Other Bank/ Write Correct Value</td></tr>
<tr><td>MB4</td><td>4</td></tr>
</table>

transistors in general and SRAMs in particular [17, 18]. Stuck at, stuck open, or stuck short faults are another source of hard errors [19]. Hard errors that are localized to a subset of the banks can be recovered completely.

Electrical noise sources such as crosstalk and power supply noises can cause temporary errors without causing lasting physical damage to the processor. A radiation particle strike can also cause localized electrical noise resulting in soft errors. These could have no effect on the system due to latching window masking, logical masking, or electrical masking. Soft errors could cause stored values in the register files to flip and this corrupted value to remain and propagate through the system if there are no corrective actions taken. A primary issue is that SRAMs are more susceptible to soft errors during read and write operations due to charge sharing [20]. A read operation could complete incorrectly leading to the supply of incorrect data. Detecting this and performing a corrective re-read would solve this issue. Writing the incorrect value poses a more difficult problem to solve if the same error is written to all banks in the register file. On a subsequent read of this register entry, this error can be detected and checkpointing would

be required to correct it.

Table 1 contains the different error recovery methods considered for the various register file configurations. For our purposes, SECDED is the protection system. The single bank configuration, ECC1B, degenerates to having a monolithic register file with ECC protection. If it detects an error but is unable to correct this error then checkpointing will be required to continue correct operation. The double bank configuration, ECC2B, is able to detect and correct single bit errors. For double bit errors, the alternate bank is read and provided this value is error free this value is then re-written to correct the error in the register file. The quadruple bank configuration, ECC4B, operates in similar fashion to ECC2B with the benefit of being able to maintain correct operation even if up to three banks are in error.

The implementation of parity is based on 16 bit interleaving. That is for the 64 bit register values bit *i* is covered by parity bit *i mod 4* as shown in Figure 4. This

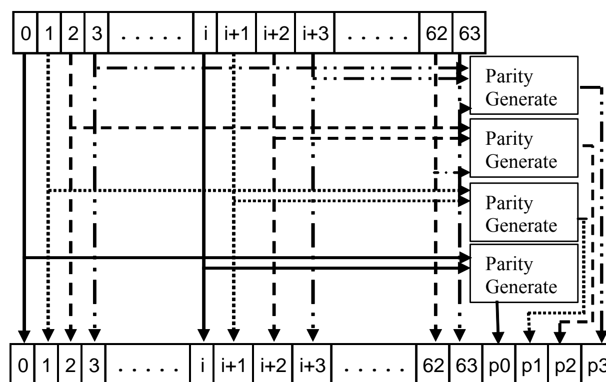

Figure 4. Interleaved parity generation

arrangement ensures that adjacent errors up to and including 4 erroneous bits can be detected. If multiple events cause distributed errors in a single register entry then parity will fail with below 4 erroneous bits if both bit i and bit $i + 4$ are flipped. Typically, physical effects such as radiation induced soft errors are confined to a specific region.

Single strikes are more likely than multiple strikes on different locations during the small duration of a register value's lifetime. Consequently, the parity design can protect against most 4-bit errors in practice. This technique also has similarities to the interleaved parity used for RAID [**21**] to combat disk errors.

The single banked error detection, SB, always requires a checkpointing recovery. For both the double, MB2, and quadruple, MB4, bank interleaved parity configurations recovery from errors proceeds through detection of the error and then reading from the other bank. Provided there is no error in the value read from the other bank then this value is re-written to correct the error in the register file with MB4 being able to maintain correct operation in the presence of up to three banks having this value in error.

Assuming there is no error in any of the input entries for an instruction then the register files of these recovery options will have no visible effect on normal operations. The actual effect on the pipeline during an error will depend on the pipeline organization, the type of error, and when it is detected. The main differences will depend on when the register file is read. In a reservation station styled architecture as each input becomes

ready the data is forwarded allowing for the possibility of error detection to proceed before instruction execution. An alternative would be to read the registers only upon confirmation that all the data is ready. These both allow, in the worst case, for error detection to be overlapped with instruction execution.

There are two main replay mechanisms used in contemporary processors namely, flush recovery for the Alpha 21264 [12] and selective replay as used in the Pentium 4 [22]. This work uses instruction based selective replay proposed in [23] for error recovery. In this mechanism each instruction carries a re-execute bit that is checked in the commit stage. This bit coupled with the bank and error information determines if this instruction needs to be re-executed and which bank(s) is known to be in error. Since the instruction is stopped in the commit stage this effectively isolates the error to within the executing functional units and the register file itself until it is corrected. To recover from an error firstly one of the alternate banks would be read and the instruction once again executed. Upon final confirmation that a bank has been located that is not in error then this bank's value will be rewritten to all the other banks to overwrite any errors. In all the multibank cases keeping track of which banks had erroneous values is a required addition and if all the available banks have been shown to be in error then an entire system rollback will be necessary.

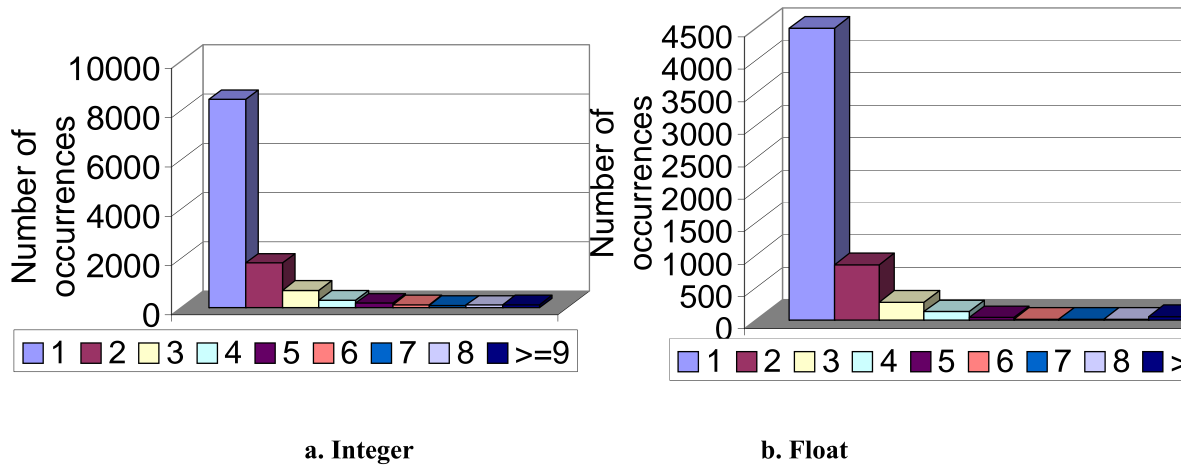

a. Integer **b. Float**

Figure 5. Average number of times a register is read with an error before correction

2.2.1 Danger of error propagation

Selective replay mechanism consumes limited processor resources to perform error correction. The farther an error spreads the more work is required to perform correction. To measure this, the number of accesses to an errant register was kept across the suit of SPEC benchmarks until one of the recovery measures corrected this register. It can be seen from Figure 5 that for MB4 after an error was injected these errors would have been consumed by multiple other instructions even within the short time it takes for recovery. This again highlights the importance of protecting register files from errors as these errors will rapidly spread to other instructions. While as expected the vast majority of these errors are only read once and then corrected, 29% of integer benchmarks and 24% of float benchmarks detected errors are read multiple times before correction. It has been

Table 2. Energy consumption, latency, and area overheads of SECDED and parity

Mechanism	Codec	Area (μm^2)	Energy per access (nJ)	Propagation Delay (ns)
SECDED (72,64)	Encoder	2170	9.40E-04	0.5
	Decoder	4067	3.12E-03	0.98
Parity (4 X 16 bit)	Encoder	2152	6.39E-04	0.3
	Decoder	755	3.47E-04	0.58

shown that certain high demand register entries are actually read over 15 times before the error is detected and corrected.

2.3 Experimental description

The Simplescalar 3.0 toolset [24] was modified to perform register renaming and multibanked register file on the Alpha SPEC2000 binaries [12]. Twelve SPEC2000 integer benchmarks and nine SPEC2000 floating-point benchmarks were considered. For

Table 3. Register file configurations (without SECDED and parity codecs included but with the bits needed for these codecs included)

Config	Banks	Register Width	Ports Per Bank		Area (μm^2)	Energy per access (nJ)	Access Time (ns)
			Read	Write			
Base	1	64	16	8	16128.800	3.935	1.269
ECC1B	1	72	16	8	16926.663	4.091	1.304
ECC2B	2	72	8	8	16531.463	3.862	0.971
ECC4B	4	72	4	8	20156.588	4.845	1.038
SB	1	68	16	8	16527.731	4.013	1.287
MB2	2	68	8	8	15876.731	3.762	0.960
MB4	4	68	4	8	19575.694	4.763	1.030

each benchmark, we fast-forwarded the first 100 million instructions and then simulated the next 300 million instructions.

Accelerated error injection was added with single bit errors, double bit errors, and triple bit errors occurring at the rate of 10^{-3}, 10^{-5}, and 10^{-7} respectively. Based on the existing work [**25**], the error probability considered in the register file for one-bit, two-bit, and three-bit errors occur with decreasing probability of approximately two orders of magnitude. The bank and location of the error was determined randomly throughout the simulation.

SECDED and parity codecs were designed in Verilog and synthesized using the Synopsys Design Compiler with 90nm TSMC technology library. The energy consumption per access, access latency, and area for different register file organizations were obtained using CACTI-3.0 [26] assuming a 90nm technology. From Table 2 and 3

Table 4. Total area and energy requirements of register file configurations (with SECDED and parity encoders and decoders included)

Config	Total Area (μm^2)	Normalized	Total Energy (nJ)	Normalized
ECC1B	99356.263	1.0000	98.2484	1.0000
ECC2B	98961.063	0.9960	92.7399	0.9439
ECC4B	102586.19	1.0325	116.3289	1.1840
SB	45814.131	0.4611	96.3244	0.9804
MB2	45163.131	0.4546	90.2872	0.9190
MB4	48862.094	0.4918	114.3125	1.1635

it is apparent that the energy per access required in reading from a register file is 4 orders of magnitude larger than the energy needed to do error detection either through parity or SECDED. From the viewpoint of energy efficiency, instead of performing duplicate reads and checks every cycle we use the redundant read only when an error is detected. The overall area and energy consumption of the different configurations including the register file and the error detection/correction codecs are shown Table 4 in for the different configurations. Each write port requires a SECDED encoder or 4 sets of parity encoders and each read port requires a SECDED decoder or 4 sets of parity decoders. In the monolithic case, a detectable error always manifests itself when identified by the decoder requiring corrective action. For both multibank sizes and error resilient strategies, there is the possibility of having an error residing in an unread bank. These occurrences are considered *bank saves*. Generally, the more banks the greater the opportunity for a bank save as shown in Figure 6 as a factor of the number of errors

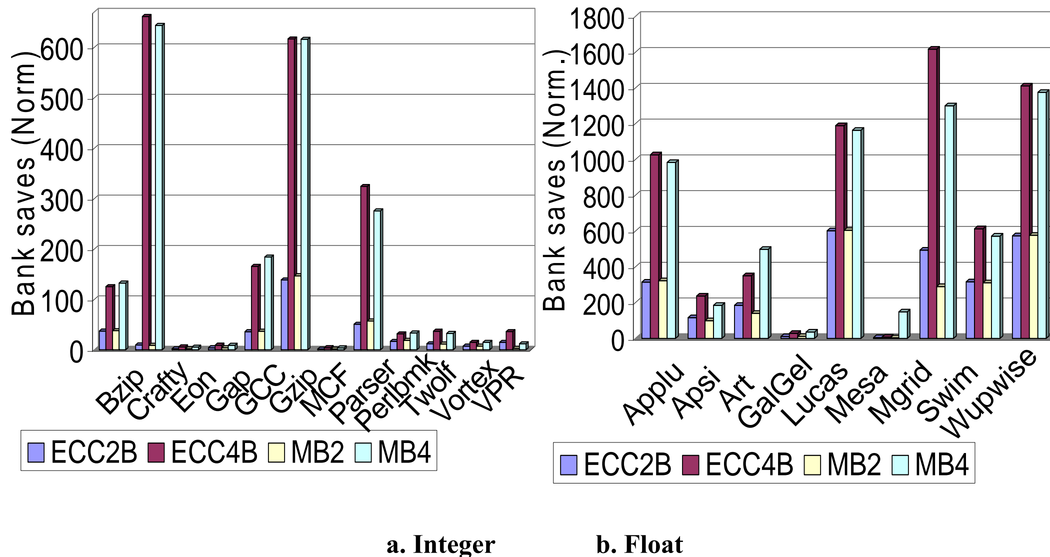

a. Integer **b. Float**

Figure 6. The number of bank saves

detected.

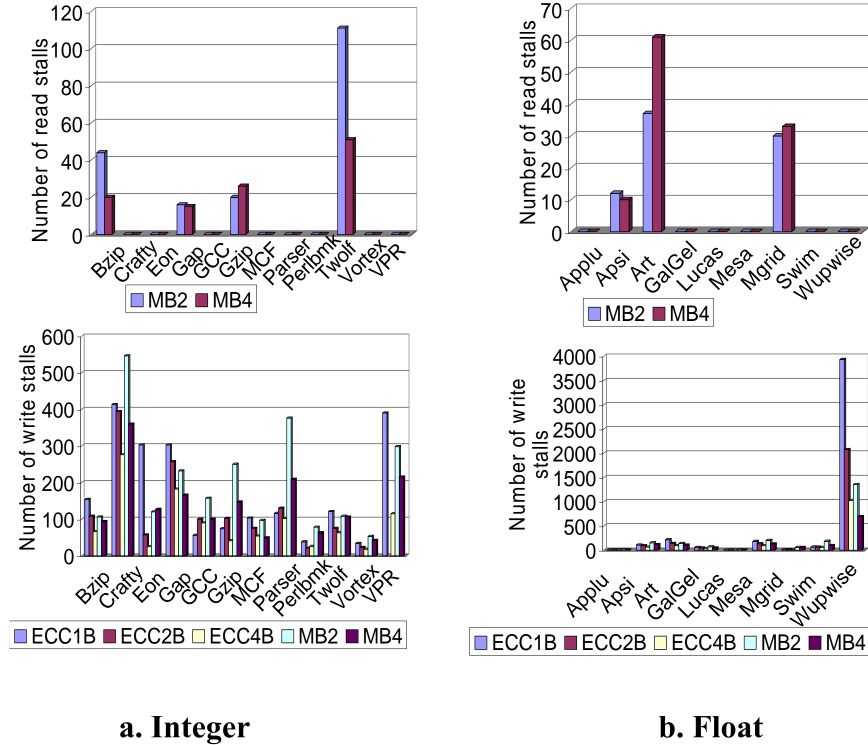

a. Integer **b. Float**

Figure 7. Number of stalls

The multibanked recovery methods require an additional read port to read the correct value and write port to overwrite the incorrect value. In the event that all read or write ports would have been needed due to ready instructions then preference is give to the error recovery request and the other port requests are delayed a single cycle. These are referred to as *read* and *write stalls*. Figure 7 shows the port stalls encountered. These numbers are appropriately placed in context when considered in conjunction with the average number of errors detected of over 53,000 across the integer benchmarks and over 62,000 across the float benchmarks. While the multibanked ECC recovery methods for

double bit errors could cause read stalls none were observed during simulation runs and so they are not included in Figure 7. The number of write stalls is significantly higher than that of read stalls because as a fraction of available resources register writes has a higher value. In particular Wupwise uses the total available 8 write ports approximately 15% of the clock cycles resulting in a much higher occurrence of write stalls than any other benchmark.

Triple bit error detection is not guaranteed with SECDED and as such the possibility of silent data corruption exists in ECC based systems faced with a triple bit error. While a scheme using an interleaved (39, 32) SECDED system could offer guaranteed protection against 4 consecutive erroneous bits and lower the latency of ECC protection it also increases the area and power overhead.

This section provided a detailed comparison of recovery options using SECDED and parity in multibanked register files. The parity scheme provides protection for up to 4 bit errors in a single register entry. Having multibanks for register files provides some inherent redundancy which to date has gone untapped. From an energy and area perspective the 2 banked configuration, MB2, is the best selection. Knowing that in a number of cases the 4 banked configuration, MB4, would allow an error to exist in one of the banks but not actually ever being read allows this to remain a feasible option even though its area and energy requirements are greater than that of the 2 bank configuration.

Chapter 3

Impact of NBTI on Different Power Saving Cache Strategies

3.1 Introduction

Systematic reduction in PMOS transistor parameters due to Negative Bias Temperature Instability (NBTI) over the lifetime of a system is becoming a significant reliability concern in nanometer regime [47,48]. Particularly, sub-threshold devices which demand a high drive current for operation are hugely affected by threshold shifts and drive current losses due to NBTI [49]. The NBTI induced shift in threshold voltage degrades the performance of CMOS digital circuits. However, degradation in performance can be offset by upsizing of the PMOS devices during the design phase. The area and performance tradeoffs due to NBTI effect have been widely studied in [50-53]. In SRAMs, these tradeoffs do not work effectively since the stability of these cells are sensitive to aging and degrade as time progresses. Also SRAM cells are particularly more susceptible to the NBTI effect because of their topologies. Since one of the PMOS is always under stress if the cell contents are not periodically flipped, it introduces an asymmetric threshold shifts in both the PMOS devices.

This work investigates the effect of NBTI independently and along with process variations on standard 6T SRAM cell and read static noise margin free 8T SRAM cell

based cache configurations. In standard 6T SRAM cell, symmetric and asymmetric (dual-V_{th}) SRAM cell topologies based caches are considered. Symmetric 6T SRAM cells are generally used for implementation of high performance caches. While asymmetric SRAM cells based caches, designed with dual-V_{th} technology have been recently studied for their strong potential of leakage power savings [54, 55]. In asymmetric SRAM cells, sub-threshold leakage current devices are made oh high V_{th}, while assuming a skewed distribution of storage bit ‘0’. The read static noise margin free SRAM cells consist of an isolated read-port and a cross coupled inverter pair have recently attracted much attention [56-58]. In these cells data storage nodes are isolated by providing a separate read-port (read current path), hence, there is no degradation in read Static Noise Margin (SNM) during read cycle, referred as read SNM free SRAM cells.

The effects of power saving strategies in 6T symmetric and asymmetric and 8T SRAM cells based caches under NBTI are investigated. It is determined that employing different power saving strategies to the caches can recover a substantial portion of the stability noise margins lost due to the predominant occurrence of logic ‘0’ being stored in caches. Based on different power saving strategies, six different cache configurations are formed and their duty cycles are derived. This leads to an additional design consideration while determining which of the power saving strategies should be applied in cache design, particularly if lifetime operation is a prime concern. Furthermore a study of the inter-die process variations in 6T and 8T SRAM cells based caches along with NBTI is done.

The following section gives an overview of the impact of NBTI on PMOS transistors and how duty cycle modulates the shift in threshold voltage. Then the experimental setups used for simulating the 6T symmetric and asymmetric and 8T SRAM cells are given. The different power saving strategies in caches and formulation duty cycles for different cache configurations are presented. The next section presents the recovery of noise margins under the different cache configurations. The effects of process variations on the noise margins and leakage currents are determined. Then the summary of key conclusions is made.

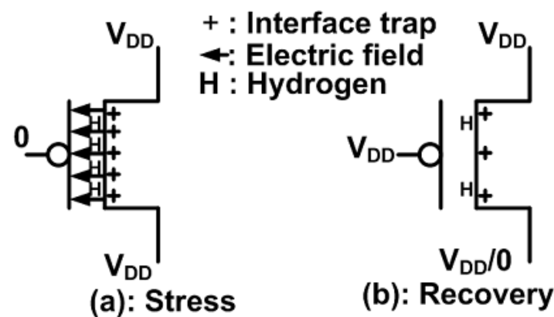

Figure 8. PMOS transistors in the stress and recovery phases of NBTI

3.2 Negative Bias Temperature Instability(NBTI)- Background and Modeling

Figure 8 shows the stress and recovery phases of NBTI in a PMOS transistor. The NBTI phenomenon that causes the gradual increase in threshold voltage (V_{th}) of a PMOS transistor due to the formation of interface traps over time. Under negative bias condition

(i.e. V_{GS} = -V_{DD}) of a PMOS transistor, interface traps are generated as hydrogen diffuses towards the gate. This phase of NBTI is called stress. When the voltage of the gate is set to V_{DD} then no new interface traps are generated and hydrogen diffuses back and anneals the broken bonds. However, full recovery becomes impossible as some of the hydrogen may no longer be available. Thanks to the annealing process dynamically recovering the threshold voltage a significant amount of performance and circuit stability can be recovered. A circuit designed for skewed activity while enabling the dynamic recovery can greatly reduce the aging effect due to NBTI. This is particularly true for SRAM cells where one of the PMOS transistors is always in stress mode.

Table 5. Summary of predictive model of |ΔVth| under NBTI

Stress	$\sqrt{K_v^2 \cdot (t-t_0)^{1/2} + \Delta V_{th0}^2} + \delta_v$
Recovery	$(\Delta V_{th0} - \delta_v) \cdot \left[1 - \sqrt{\eta (t-t_0)/t}\right]$
Kv	$A \cdot T_{ox} \cdot \sqrt{C_{ox} \cdot (V_{gs} - V_{th})} \cdot \left[1 - V_{ds} / \alpha (V_{gs} - V_{th})\right] \cdot e^{\frac{E_{ox}}{E_0}} \cdot e^{\left(-\frac{Ea}{kT}\right)}$

It is understood that NBTI follows a Reaction-Diffusion (R-D) process and that the stress and recovery phases are successfully analyzed using the R-D model as used in [59, 4].

As shown in Table 5, there are a number of factors that influence the shift in threshold voltage due to NBTI. There is a positive relationship among these parameters such as temperature, supply voltage, the duty cycle, and the magnitude of change in the threshold

voltage. By lowering any of these parameters, reduction in change in threshold voltage $\Delta|V_{TH}|$ will be observed. Temperature is a function of power density and the rate at

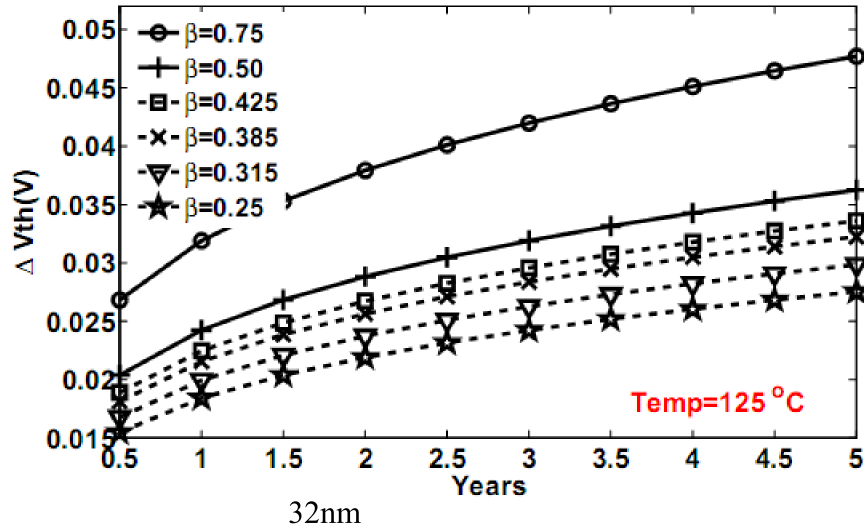

Figure 9. Shifted threshold voltage of a 32nm and 45nm technology node PMOS transistor versus time for different duty cycles (β)

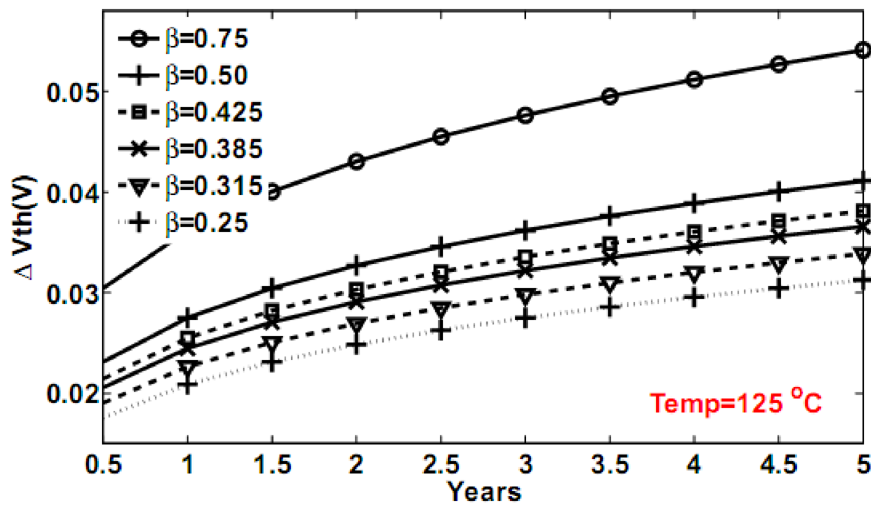

Figure 10. Shifted threshold voltage of a 45nm technology node PMOS transistor versus time for different duty cycles

which heat is being removed from the system. Supply voltage can also be modified in order to control the power density because of its quadratic relationship, depending upon the workload for increased or decreased performance. Under lighter workloads the

supply voltage can be lowered and conversely increased for demanding workloads. Since the PMOS only degrades when there is a gate to source (V_{GS}) voltage difference then the ration of this to no voltage difference, the duty cycle (β), heavily impacts the change in threshold.

Figure 9 and Figure 10 show the shift in threshold voltage $\Delta|V_{TH}|$ of a 32nm and 45nm PMOS transistor due to NBTI over a five years of time. The shifted threshold voltage is plotted for different duty cycles and $V_{GS} = -V_{DD} = -1V$ (-0.9V for 32nm) and T = 125^{o}C. Duty cycles for different cache configurations are explained and derived in a later section. It can be seen from these figures that the duty cycle (β) has significant role in modulating the threshold voltage of a PMOS transistor. Hence, it could be a good candidate to control the aging effects.

In caches, there is a strong bias towards logic bit value ‘0’ being stored on average 75% of bit value for most of the time in data or instruction caches [55, 60]. By periodically inverting the contents of the cache and marking that the data is inverted this occupancy can be further reduced to 50%. Even with the cache occupancy of logic bit value ‘0’ is 50%, then PMOSs are degrading but their degradation is occurring in a balanced fashion. Asymmetric SRAM cell based caches have been developed to take the advantage of skewed distribution of logic bit value ‘0’ in caches for leakage power reduction. However, the asymmetry introduced by the dual-V_{th} devices degrade the read SNM and makes the SRAM cells more vulnerable as shown in the next section.

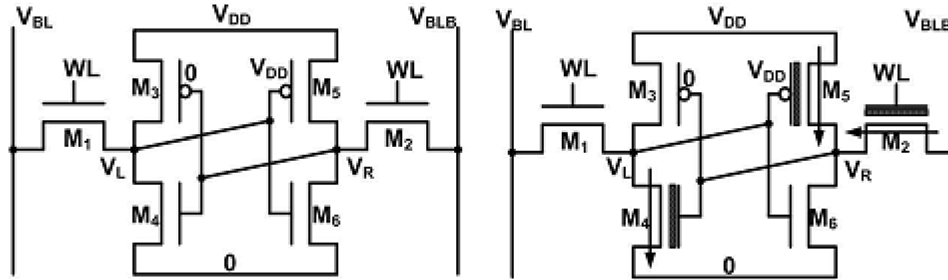

a. Symmetric standard 6T SRAM cell **b. Asymmetric 6T SRAM cell for reduced leakage**

Figure 11. a) Symmetric 6T SRAM cell b) asymmetric 6T SRAM cell for reduced leakage current based on dual-V_{TH} transistors (shaded transistors are high V_{TH})

3.3 SRAM Cells Under NBTI

Figure 11 shows the symmetric 6T SRAM cell and asymmetric 6T SRAM cell for reduced leakage current. In symmetric 6T SRAM cell simulation setup shown in Figure 11, transistor M_1, M_2, M_4, M_5, and M_6 have nominal V_{TH} values, while shifted $\Delta|V_{TH}|$ value due to NBTI is used for M_3 transistor model. In asymmetric 6T SRAM cell, shown in Figure 11 has three types of transistor models, transistor M_1 and M_6 have nominal value of V_{th} models, shifted $\Delta|V_{TH}|$ value due to NBTI is used for M_3, while transistors M_2, M_4, and M_5 have higher V_{TH} to reduce leakage current. All transistors in symmetric and asymmetric 6T SRAM cells are of minimum feature sized with cell ration = 2 and M_1 = M_2 = M_3 = M_5 = 45nm/45nm (32nm/32nm), and M_4 = M_6 = 45nm/90nm (32nm/64nm). In 8T SRAM cell [56], regular-V_{TH} and minimum feature sized transistors are used.

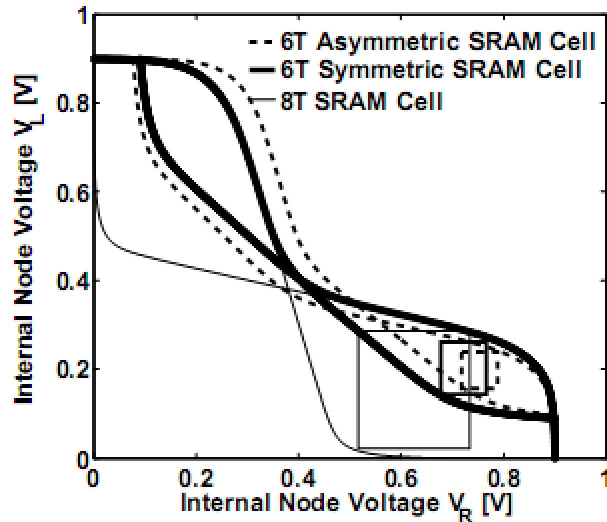

Figure 12. Different SRAM cells butterfly curves for read SNM measurements

Figure 12 shows the voltage transfer characteristics (VTCs) or butterfly curves of 6T symmetric and asymmetric and 8T SRAM cells for 32nm technology with NBTI effect and the duty cycle β = 0.25 for five years of time span. The butterfly curve of a symmetric 6T SRAM cell is almost symmetric and it has negligible effect of shifted threshold voltage due to NBTI. For an asymmetric 6T SRAM cell, butterfly curve is not symmetric because of dual-V_{TH} devices. However, shifted threshold voltage due to NBTI of transistor M_3 has less effect as compared to high-V_{TH} devices (M_2, M4, and M_5) used in the cell. For 8T SRAM cell, read SNM is equivalent to hold SNM or in other words isolated read port held the data storage nodes unchanged. Hence, 8T SRAM cell has better read SNM and 6T asymmetric SRAM cell has worst read SNM.

The stability parameters (SNM and WNM) of 6T symmetric and asymmetric and 8T SRAM cells are analyzed using HSPICE simulation results in order to investigate the NBTI effects. The static noise margin (SNM) obtained from the butterfly curve is used for measuring the read and hold stability. The SNM is estimated graphically as the length of a side of the largest square that can be embedded inside the lobes of the butterfly curve. The write stability (WNM) is measured using the write trip point, defined as the minimum amount of voltage needed on the bitline to flip the cell content [62].

Table 6. SNM and WNM Changes due to NBTI for 32nm Technology at 125°C

Cells	NBTI	Hold SNM [mV]	Read SNM [mV]	WNM [mV]
6T Sym SRAM	t=0	267.02	135.15	150
	t=5 yrs	247.16	120.29	179
6T Asym SRAM	t=0	245.36	101.02	140
	t=5 yrs	224.09	86.36	170
8T SRAM	t=0	267.02	267.01	150
	t=5 yrs	247.16	247.15	179

Table 7. Recovery of SNM and WNM for 32nm Technology at 125°C and duty cycle (β)=0.25

Cells	NBTI	Hold WNM [mV]	Read SNM [mV]	WNM [mV]
6T Sym SRAM	β=0.25	260.29	130.15	160.00
	Recovery(%)	66.11	66.35	65.52
6T Asym SRAM	β=0.25	238.07	96.00	151.00
	Recovery(%)	65.73	65.76	63.33
8T SRAM	β=0.25	260.19	260.19	160.00
	Recovery(%)	65.61	65.66	65.52

Table 6 and Table 7 show the degradation and recovery in stability parameters, respectively, for the 6T symmetric and asymmetric and 8T SRAM cell caches.

Degradation in read SNM and WNM are determined by simulating the SRAM cells for the shifted value of threshold voltage due to NBTI after five years of time for different duty cycles. The percentage of recovery in SNM and WNM are calculated by incorporating the dynamically recovered threshold voltage for $\beta = 0.25$ in the SRAM cells, simulation results for different SRAM cells are tabulated in Table 7.

The amount of degradation in read SNM and WNM is approximately 10% for both symmetric and asymmetric SRAM cells after five years time span. While there is a 28.3% reduction in read SNM for the asymmetric SRAM cell as compared to symmetric SRAM cell, as shown in Table 6. This drastic reduction in read SNM is mainly due to asymmetry introduced by the dual-V_{TH} devices used in the cell in order to minimize leakage current. Also an opposite trend of WNM is observed, WNM of the initial (t=0) SRAM cell is lower than the stressed (t=5 years) SRAM cell because of increase in trip-point of inverter M3 and M4. Therefore higher voltage is needed at the bitlines to write into the SRAM cells due to the aging effect.

Recovery in the read SNM cells is almost equal, as shown in Table 7. However, recovery in WNM for asymmetric SRAM cell is slightly higher than the symmetric SRAM cell. Higher recovery in WNM is observed in asymmetric SRAM cells due to high V_{TH} devices which have lesser effect due to NBTI as compared to the lower V_{TH} devices used in symmetric SRAM cells.

3.4 Power Saving Strategies in Caches

By lowering the supply voltage of the cache leakage energy can be saved. If the path to V_{DD} or ground is completely cutoff then the state then the state of the cache is lost and this is considered as a state destroying mode. Otherwise, if the supply voltage is lowered

Table 8 Duty cycles and threshold changes for the considered cases

Cases	β	Description	PMOS 32nm/45nm Δ|VTH|[mV]
1	0.75	Average occupancy of 0 in cache	47.7/85.12
2	0.5	Periodic bit flipping	36.2/41.1
3	0.425	Periodic bit flipping + Disable	33.6/38.13
4	0.385	Periodic bit flipping + Sleep mode + Speculative wake up	32.2/36.6
5	0.315	Periodic bit flipping + Sleep mode + On demand wake up	29.5/33.86
6	0.25	Periodic bit flipping + Sleep mode + Decay to state destroying	27.5/31.25

but the SRAM cells are still able to maintain their state then this is considered as a state preserving mode and the cache is termed to be in a sleep state. An earlier work [63] had considered a number of state preserving and state destroying strategies for the L2 cache. The five main strategies are as follows:

- CONSERVATIVE: When a block in L1 is written to, then the corresponding sub-block in L2 is fully turned off and put into a state destroying mode. Since the block in L1 is dirty then the block in L2 is dead and can be safely deactivated.

Since instructions are not written to, this strategy cannot be optimized for instructions.

- SPECULATIVE I: When a block is brought from L2 to L1, the block in L2 is put in a state-preserving mode immediately. It does not wait for the L1 block dirty nor does it lose data. If the evicted block had become dirty then the block in L2 is reactivated and written into.
- SPECULATIVE II: Similar to Speculative I but instead the L2 block is put into the state destroying mode instead of sleep mode
- SPECULATIVE III: This is similar to SPECULATIVE I, but the block in L2 is speculatively woken up when the L1 block is evicted.
- SPECULATIVE IV: This is similar to Speculative I except the L2 block is reactivated and written back whenever the corresponding L1 cache block needs to be replaced.

The leakage energy saved under these cache strategies in [63] was used to calculate the average inactive time experienced by the cache blocks. Based on the average inactive time, the duty cycles (β) are derived and the following cases are formed:

- CASE 1: Is the skewed distribution of ‘0’ with average occupancy of 75% of the time. Due to the symmetric nature of SRAMs if one of the PMOSs has a duty cycle of 0.75 then the other has a duty cycle of 0.25
- CASE 2: A cache with an additional value stating if the data is inverted or not. With periodic inversions the duty cycle will approach 0.5 equalizing the degradation on both the PMOS transistors.

- CASE 3: Extends CASE 2 and incorporates the Conservative strategy of disabling cache blocks.
- CASE 4: Employs sleep mode of cache blocks when the block is written into L1. It attempts to speculatively wake up the L2 block before it is needed.
- CASE 5: It is similar to CASE 4, but it only wakes up the block when needed
- CASE 6: It combines the sleep mode with the disabled mode using a timer that if it expires switches the cache block off instead of continued sleeping.

Table 8 summarizes the considered cases for different cache strategies and duty cycles. The NBTI induces shift in threshold voltages ($\Delta|V_{TH}|$) for 32nm and 45nm technology node PMOS transistors after 5 years is estimated using the model described in Table 8. The PTM 32nm and 45nm technology models were used with the parameters in Table 9.

The PMOS transistor models with shifted threshold voltage due to NBTI for the different duty cycles are incorporated in HSPICE net list to simulate different SRAM cells based cache configurations for their stability parameters and leakage current measurement.

Table 9. PTM Technology Model Parameters

Tech.	V_{DD} [V]	V_{TH} [V]	High V_{TH} [V]	T_{OX} [nm]
32nm	0.9	0.16	0.24	1
45nm	1	0.18	0.27	1.1

3.5 SNM and WNM Recovery Under Different Cache Configurations

This section presents the SNM and WNM recovery for different cache configurations (power saving cache strategies) formulated on the basis of duty cycles. Based on different cache strategies discussed in the above section, we calculated the recovery in read SNM and WNM. Figure 13 and Figure 14 show the recovery in read SNM for 32nm and 45 nm technology nodes respectively at 70°C and 125°C for different SRAM cells based cache configurations. Recovery in read SNM for cache configurations varies from 38% for CASE 1 to 66% for CASE 6. There is a slight increase in recovery of read SNM at higher temperature. As we have seen in Section II that duty cycle β has a significant role in modulating the shift in threshold voltage due to aging, which is also visible from the percentage of recovery of read SNM for different SRAM cell based cache configurations.

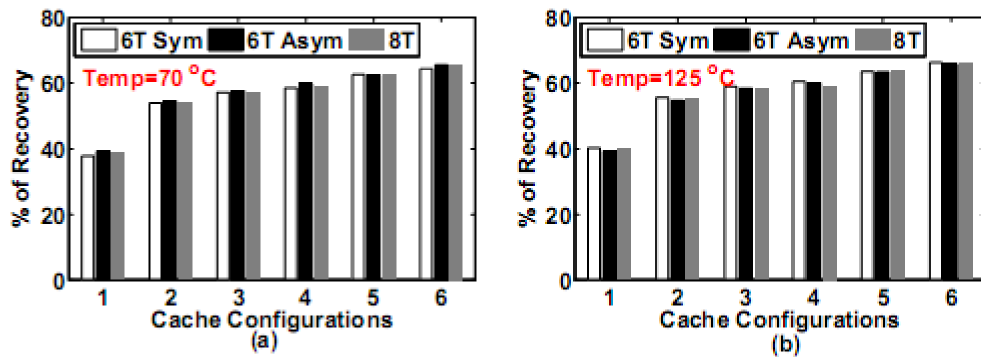

Figure 13. Percentage of recovery of read static noise margin for different cache configurations based on 32nm node SRAM cells

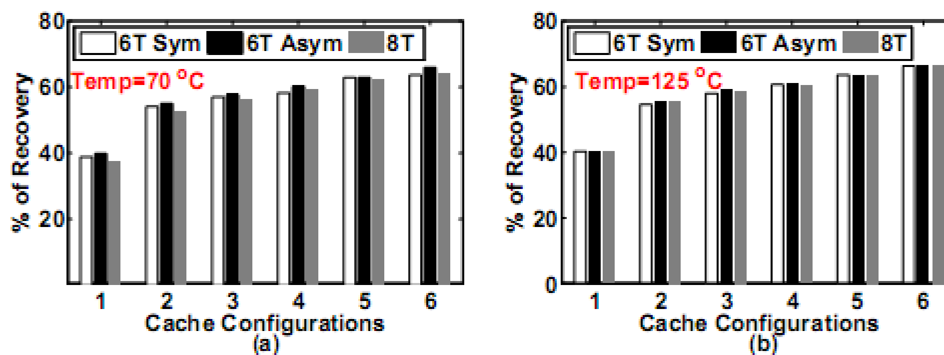

Figure 14. Percentage of recovery of read static noise margin for different cache configurations based on 45nm node SRAM cells

A similar trend has been observed in the recovery of WNM for different SRAM based cache configurations. However, rate of recovery of WNM for asymmetric SRAM cell caches is slightly lower than the symmetric SRAM cells based caches, as shown in Figure 15 and Figure 16 for 32nm and 45nm nodes respectively. While the rate of recovery of WNM of an 8T SRAM is almost equivalent to 6T SRAM since write operation in a 8T SRAM is done in similar fashion of 6T assuming that the regular V_{TH} devices are used.

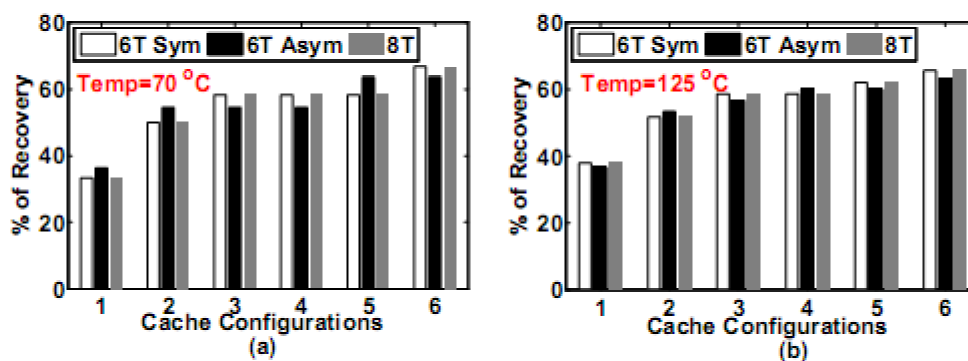

Figure 15. Percentage of recovery of write noise margin for different cache configurations based on 32nm node SRAM cells

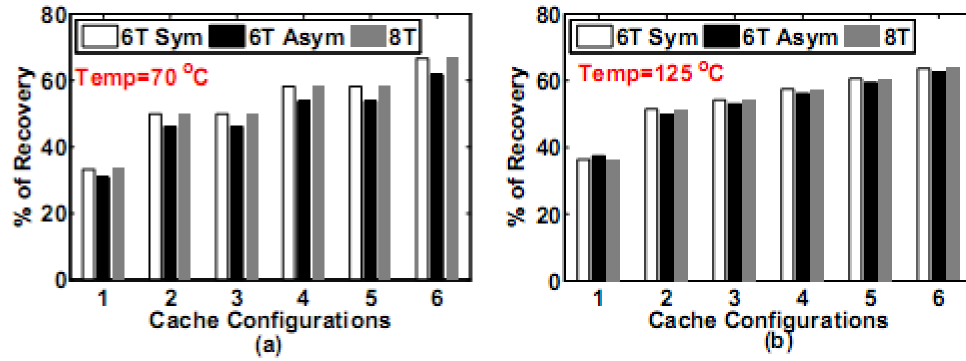

Figure 16. Percentage of recovery of write noise margin for different cache configurations based on 45nm node SRAM cells

Increased rate of recovery of WNM and SNM in different SRAM caches is specifically seen for lower duty cycles. It is mainly because of lesser impact of NBTI due to lower duty cycle. Since CASE 6 has the least impact of NBTI or has better capability of dynamically recovering the shifted threshold voltage due to NBTI. Hence, it can be a good candidate of cache configuration where reliability and life span is a major concern.

3.6 Effect of NBTI Under Process Variations

Variations in design and process parameters such as threshold voltage leads to a greater loss of parametric yield with respect to SRAM noise margins []. The effect of NBTI has a direct impact on the PMOS device threshold voltage as a result SRAM cells are more susceptible to parametric failure due to again effects. In order to investigate the effect of NBTI along with process variations 1000 Monte Carlo simulations were performed for each configuration assuming a 15% variation in V_{TH} with 3σ as an independent random variable for all the transistors in SRAM cells (6T and 8T) with Gaussian distribution.

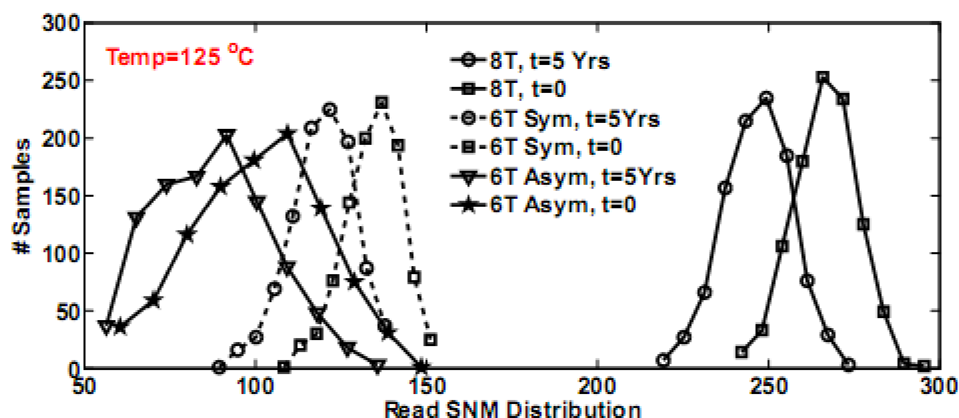

Figure 17. Read SNM distribution of different SRAM cells for 32nm technology node

Figure 17 and Figure 18 show the read SNM distribution of 32nm and 45nm technology nodes respectively at 125°C temperature for different SRAM cells. Degradation in mean read SNM due to NBTI after five years time is clearly visible. The effect of process variations along with NBTI is more dominant in 32nm node as compared to 45nm technology node. However, asymmetric 6T SRAM cell shows large standard deviations as compared to its counterpart 6T and 8T SRAM cells for both technology nodes.

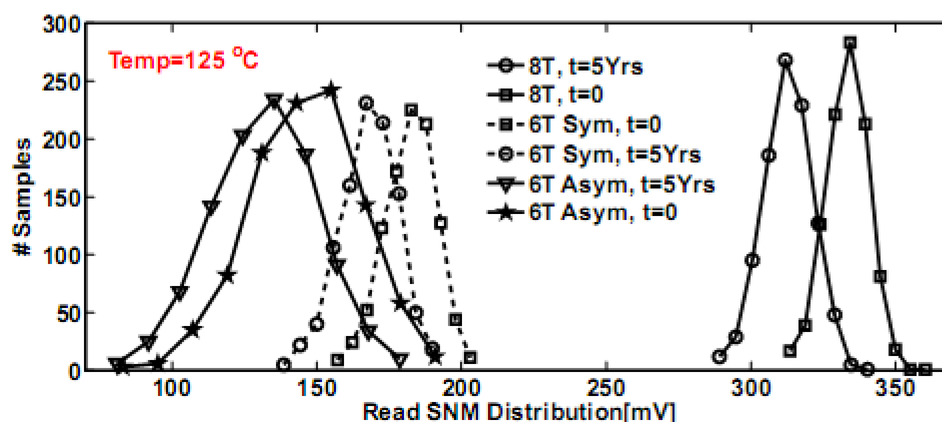

Figure 18. Read SNM distribution of different SRAM celllls for 45nm technology node

Figure 19 and Figure 20 show the leakage current distribution of 32nm and 45nm technology nodes respectively at 125°C temperature for different SRAM cell

configurations. It is seen that the leakage current follows the log-normal distribution because of its exponential dependence with threshold voltage. Also, there is significant reduction in mean leakage current for asymmetric 6T SRAM cell as compared to symmetric 6T and 8T cells, increased threshold voltage due to NBTI further reduces the leakage current.

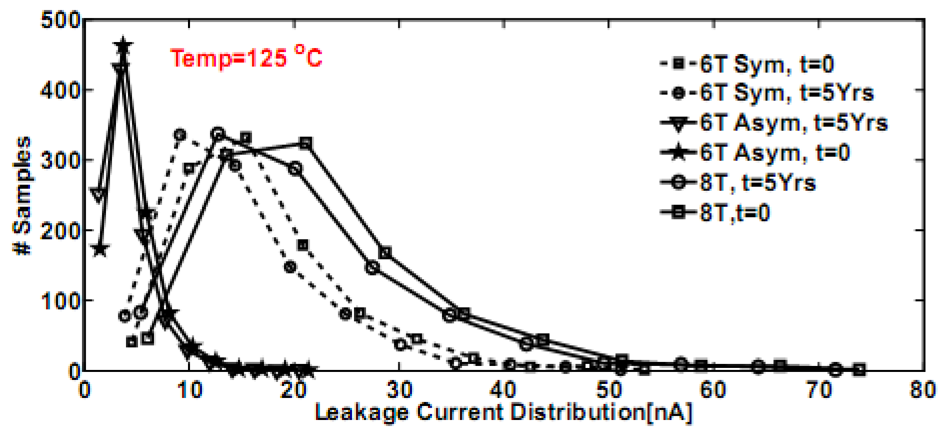

Figure 19. Leakage current distribution of different SRAM cells for 32nm technology node.

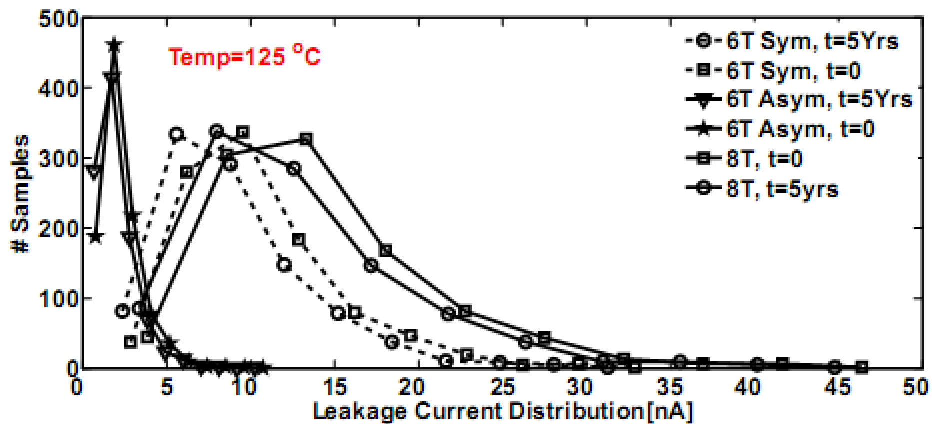

Figure 20. Leakage current distribution of different SRAM cells for 45nm technology node.

Chapter 4

Mitigating thermal effects on clock skew

4.1 2D Clock networks

4.1.1 Introduction for 2D clock networks

With the advancement of VLSI design and technology, temperature gradient across a single chip has become an important factor that significantly affects the performance of a chip. The difference in temperature between different parts of a chip arises because of the structural difference and diversity of computational activities across the chip. In modern processors, the temperature difference can be as high as 50°C [1] affecting the performance of different building blocks on the chip as well as the interconnects between them. Increasing the temperature changes the resistance of interconnect wires according to the relation $R = R_0[1 + \alpha(T - T_0)]$, where R_0 is the interconnect resistance at the nominal temperature T_0, and α is the temperature coefficient of resistance for the interconnect material. For common interconnect materials (Copper and Aluminum), the resistance increases with the increase in temperature. On the other hand, MOSFET devices behave differently than interconnects with the increase in temperature because of two contending effects on the drain current due to the decrease of both the carrier mobility (μ) and the threshold voltage (V_{th}). The square-law equation of the drain current of a MOSFET transistor $I_{DS} = \mu \cdot C_{ox} \cdot \left(\frac{W}{L}\right) \cdot \left(V_{GS} - V_{th}\right)^2$ suggests that the decrease in the mobility would decrease the current, and inversely, the decrease of the threshold

voltage would increase the current. However, the current-temperature relation is a function of the operating point of the transistor – the threshold voltage variation is dominant at relatively low bias voltages whereas the mobility variation with temperature is dominant at higher bias voltage. There exists a bias point, known as the zero

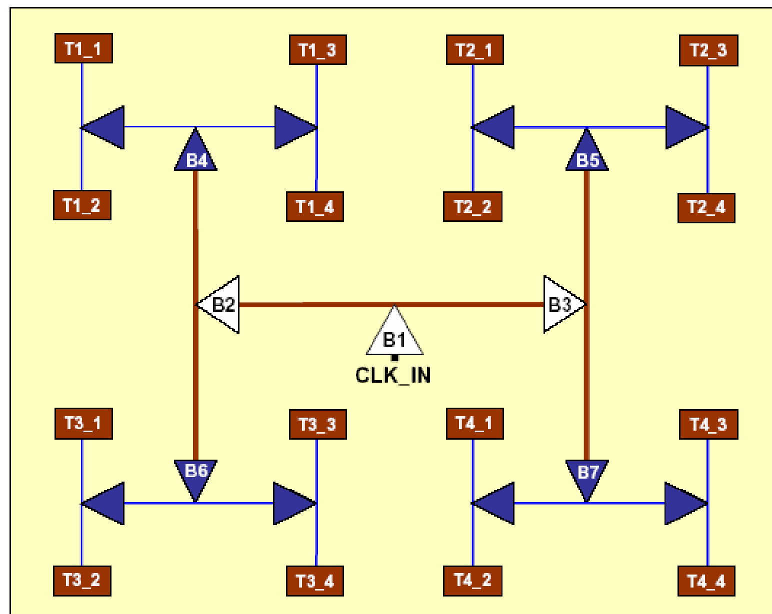

Figure 21. H-tree with node labels

temperature coefficient (ZTC) point, where the drain current of the transistor is insensitive to temperature variations.

In this work, we focus on the effect of temperature on the clock skew between physically close and functionally related terminals of a clocking network. In the H-tree shown in Figure 21, we can see that for a number of terminal locations, while physically close, the clock signals reach through completely different paths from the source. Consequently, temperature differences in the paths can lead to significant skews. Since the increase of

clock skew possesses a big performance threat to high performance digital integrated circuits, we need intelligent solutions to mitigate the effect of temperature on clock skew. Several techniques have been proposed in this context to compensate for the skew variation across the chip. One of the recent methods suggested by Shakeri and Meindl in [29] uses a temperature variable supply voltage to guarantee near constant delay across a temperature range. One issue associated with this method is that it needs different supply voltages for different sections and thus required a number of level converters. Moreover the ZTC point occurs at a voltage around 0.37V for the 45nm node, which will greatly reduce the speed of circuit operations. Moreover, in the context of the clock tree, this method requires fine-tuning of the supply voltage for multiple sections of the clock tree and level converters between each of these sections, which increases the complexity of the circuit. Other solutions in literature [30,31] assume some known fixed temperature profiles, which will be highly inaccurate particularly in the case of a processor that runs many different applications with varied characteristics, producing application specific temperature profiles. Recently, partially adaptive tunable delay buffers were proposed to compensate for clock skew due to thermal variations [32]. Tuning of delay buffers was computed off-line and stored in a lockup table, consequently, not allowing a true dynamic buffer tuning.

In this paper, we propose a new circuit technique that dynamically adjusts the driving strengths of the clock buffers according to the current temperature profile of the chip and reduces the overall clock skew between terminals, thus, allowing a truly dynamic tuning.

As opposed to the previously proposed methods, our method only uses modified clock buffers and, more importantly, does not assume any fixed temperature profile thereby maintaining uniform performance over time and for different applications. Applying the thermal profiles used in [30], we demonstrate that our buffering technique can reduce this thermally induced skew by 72.4% on the average with an insignificant area requirement.

4.1.2 RLC model for including thermal effects

The relative arrival times of the clock signal at different parts of a design is crucial for the proper functioning of synchronous digital circuits since synchronization of various modules ideally requires zero skew clock signals at all communicating clock sinks. Therefore, the clock distribution network of a circuit is carefully designed such that the clock signal arrives at each terminal at the same time. Among the different configurations of such ideally zero skew clocking schemes, H-tree is the most commonly used structure in regular array based designs [19]. Clock buffers are inserted at the clock source and at the fanouts to increase the strength of the clock signal. Due to different variations in the operating conditions as well as manufacturing variations, significant clock skew can be observed at the terminals of an H-tree. On-chip temperature variation happens to be a major source of clock skew in modern integrated circuits. The skew is the result of the uneven spatial variation of the wire resistance due to temperature variation as well as the temperature dependent variations of the clock buffers. For the analysis of the clock tree under temperature variation, we developed SPICE models for the H-tree where the

interconnect wires are modeled by their RLC parasitic components and the devices are modeled using the BSIM4 model.

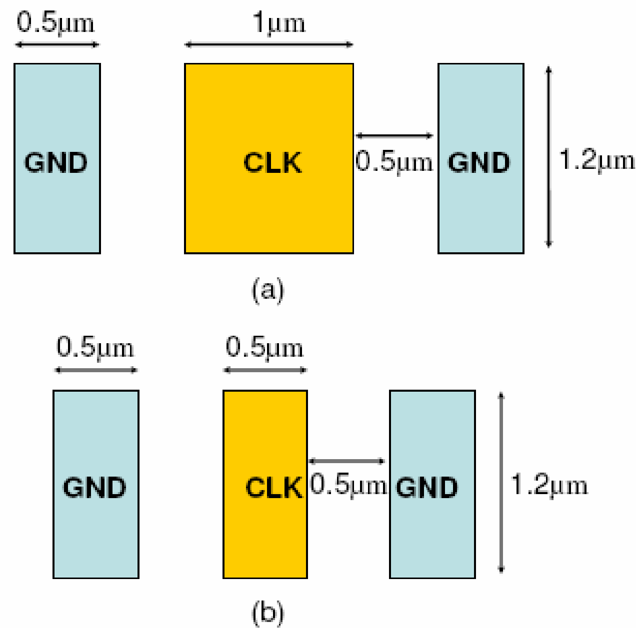

Figure 22. Wire dimensions in the (a) first level H and (b) second level H's

A typical shielded H-tree in a 45nm design is chosen for this work where the die length is 2cm. For the purpose of illustrating the effects of temperature variation on H-tree, we consider a 2 level H-tree, as depicted in Figure 21. The clock tree is shielded from interference on both sides by two ground lines. The dimensions of the wires in clock tree are shown in Figure 22 and clock buffers of appropriate size have been inserted in the H-tree.

Since the clock lines are extended over a long region in the global metal layers, they experience significant inductive effects. Therefore, when modeling the interconnect of

the H-tree, we use the RLC representation of the interconnect wires for accurate analysis of the clock tree. In this work we use a partial element equivalent circuit (PEEC) based interconnect model which is readily usable by circuit simulators like SPICE [33]. The values of the interconnect parasitics were computed using different parasitic extractors: the resistance and inductance values were extracted by FastHenry [34] and the capacitance was extracted using FastCap [35].

From the interconnect perspective, temperature mainly impacts the resistance of the interconnect wires since the resistivity of the interconnect material increases with temperature. As far as the capacitance is concerned, the effect of temperature can safely be ignored. On the other hand, since the partial inductance values used in our H-tree model depend solely on the geometry of the conductors [36], temperature does not have any direct impact on these values. However, temperature can affect the effective loop inductance because of the redistribution of the currents through different return paths due to uneven change in the resistance values due to varying temperature over the entire chip. Therefore, effectively, while performing temperature dependent simulation of the H-tree, we need to assign proper temperature dependent values to the resistances and the effect of temperature on the effective inductance will be automatically taken care of by the circuit simulator. The clock buffers are modeled using the predictive BSIM4 devices models for the 45nm CMOS technology, which proved to accurately capture process sensitivities in the nanometer regime [37].

4.1.3 Effect of temperature on clock skew

For investigating the impact of temperature variation across the chip on clock skew, we simulated several temperature profiles originally suggested by [30]. This earlier work focused on reducing the skew over a single insertion point while our work tackles the entirety of the tree and offers solutions for more generic scenarios. The profiles include a linear as well as an exponential temperature fall off across the chip with different temperature ranges. To illustrate the problem of temperature dependent skew variation, we start with a simple split profile where we divided the chip into two temperature zones, TL on the left half and TR on the right half with the left half being cooler than the right half. However, our evaluation considers more complex profiles representative of different on-chip temperature profiles. Figure 23 shows how the clock signal propagates from the input buffer across a temperature difference of 50°C through different paths towards two physically close terminals T3_4 and T4_2 (Figure 23). For this pair of paths, the skew between the outputs of the buffers equidistant from the clock source increases as the

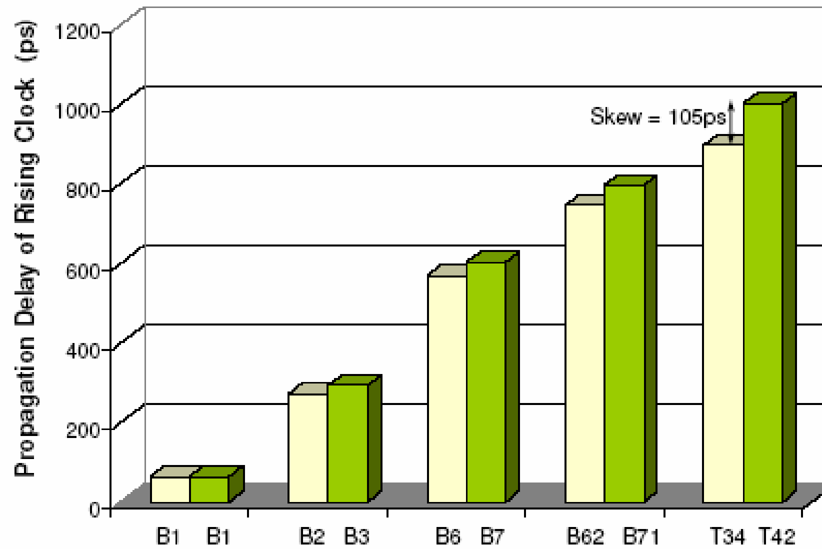

Figure 23. Increasing skew as clock signal propagates through buffers at a temperature difference of 50°C

distance from the source increases. However, it should be understood that the tendency of the skew propagation depends entirely on the temperature profile and the pair of terminals under consideration. After traversing the paths under the assigned temperature profile, the terminals T3_4 and T4_2 end up having a significant clock skew of 105ps between them.

Increasing the buffer size generally increases the speed of operation until the self loading or the interconnect delay stops any further improvement. Therefore, if the buffers in the high temperature region are made bigger than the buffers operating at a lower temperature, the buffers will speedup the signal in high temperature region, thereby reducing the skew. However, the sizes of the buffers are very specific to the design under consideration and will strongly depend on the temperature profile. Therefore, we cannot have any optimal buffer size that can compensate for every possible temperature profile

on the chip. Thus, exhaustive simulations on the fixed buffer sizes need to be performed for each temperature profile. Although this kind of exhaustive simulation based buffer sizing is expensive in terms of computational requirements and can address static temperature profiles only, we performed the simulations for two reasons: (i) this shows that modifying the buffer strength can reduce the clock skew and thus motivates the design of the adaptive strength buffers presented in the next section and (ii) this exhaustive buffer sizing provides an achievable upper bound that can be achieved by modifying the driver strength only.

The buffers are upsized in the slower path by a factor β compared to their counterparts in

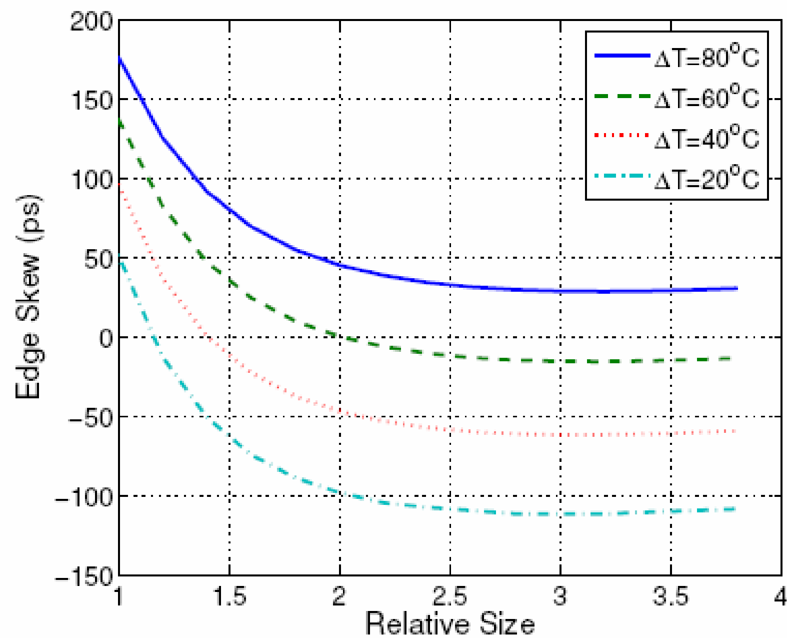

Figure 24. Skew improvement by increasing width of buffers

the faster path. For our simulations, we performed a sweep on β over the range 1 to 4 with an increment of 0.05 by fixing the temperature of one side at 170°C and varying the temperature of the other side to different temperatures, as shown in Figure 24. As mentioned before, for a given temperature difference, the slower path becomes faster with increasing β till the point when the interconnect delay or self loading does not allow any further reduction in the path delay. Therefore, the delay of the "slower" path first reduces to a minimum, then increases. However, as seen from Figure 24, there could be values of β that leads to over-compensation and the skew may become negative. This over-compensation may arise at a lower temperature difference where there is less need for compensation. For example, the value of β that gives zero skew for a temperature difference of 20°C is 1.16, whereas a value of 1.41 is required when the difference in temperature is 40°C. For even higher temperatures, the skew cannot be reduced to zero by buffer sizing alone, as found in the case of a temperature difference of 80°C. Table 10 provides results for the linear and exponential temperature profiles mentioned before. In the table, TH and TL represent the highest and the lowest temperatures, respectively. For

Table 10. Possible skew improvement using buffer selection

Thermal Profile	Parameters		Unmodified Skew (ps)	Driver Strength* Placement (ps)	Driver Strength Potential Improvement (%)
	T_H	T_L			
$T(x) = ax + b$ $a = \frac{T_H - T_L}{L}$ $b = T_L$	170	90	79.6	0.28	99.65
	170	110	61.4	0.81	98.68
	170	130	42.5	2.81	93.39
	170	150	22.5	1.76	92.18
$T(x) = a.e^{-bx}$ $a = T_H$ $b = \frac{1}{L} ln \left(\frac{T_H}{T_L}\right)$	170	90	77.6	1.47	98.11
	170	110	60.5	0.16	99.74
	170	130	42.2	2.47	94.15
	170	150	22.4	1.71	92.37

these profiles it is observed that on average skew can be potentially reduced, theoretically, by 96% with the worst case improvement of 92.18% over the unmodified H-tree in the presence of temperature variation. In this search, each individual buffer had two discrete size selections. This was used to determine the minimum skew possible from these available choices. In the next section, we describe an implementation of the adaptive clock tree to compare the effectiveness of our practical implementation.

4.1.4 Design of dynamically adaptive clock buffers

For reducing the variation of the clock skew with temperature gradient, we employ a dynamically adaptive circuit scheme where local temperature sensors sense the ambient temperatures and convert the temperatures to voltages that are used for dynamically changing the driving strength of the clock buffers to reduce the overall skew. We utilized the temperature sensor design in [38] on the chip to monitor the temperature variation.

The adaptive buffers use the combination of two techniques to compensate the temperature effect, which are buffer current control and body bias control. In the buffer current control technique, the currents of the NMOS and PMOS transistors are controlled by adaptively connecting an additional parallel branch which acts as a variable current source when controlled by a proper bias voltage. Adaptively controlling the buffer currents is used to compensate the speed degradation by increasing the driving current when the temperature increases. The second technique works on controlling the body-

bias of the transistors to control their threshold voltages. Reducing the threshold voltages of the transistors increases their speed, which can be used to compensate for the speed degradation at elevated temperatures.

Our first design technique for the adaptive-current buffer is illustrated by the circuit shown in Figure 25. In this technique, the effective transistor current in the pull-up and pulldown sections are adjusted dynamically to control the rise and fall times of the output, respectively. The control signals coming from the wave shaping circuits after the temperature sensor are connected to the PMOS and NMOS switches in the parallel branches. These switches are all off at the lower temperature and they begin to conduct current gradually as the temperature increases. At the maximum operational temperature, the two switches conduct the maximum current.

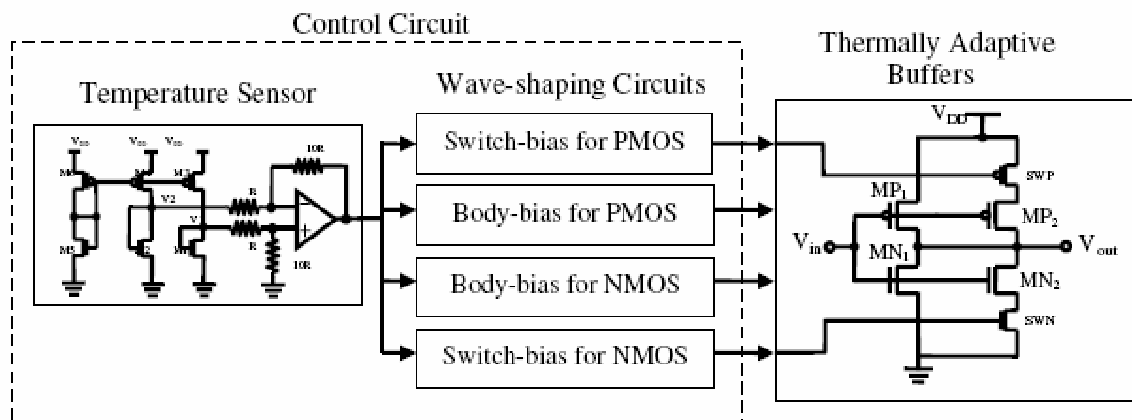

Figure 25. Thermally adaptive buffer schematic

Dynamic body bias, the other technique used in our temperature adaptive buffers, is an effective technique to achieve dynamic-V_{th} in bulk CMOS. Several circuit techniques that apply forward body bias to decrease V_{th} (and improve performance) have been proposed [39, 40]. In our design, we used two different sources to provide forward body bias to the PMOS and NMOS transistors. For the NMOS transistor, normally designed with bulk tied to the ground, we apply a positive voltage on its bulk. The applied voltage ranges from 0V to 0.5V as further increase in the body bias draws significant current in the junction, which increases both the power dissipation and the required size of the wave-shaping buffers. The inverse voltage, varying between VDD=1V and 0.5V, is applied to the PMOS transistor bulk.

The wave shaping circuits that follow the temperature sensor are responsible for generating the temperature dependent voltages that are used as the bias voltages of the switches (SWN and SWP) and the body bias of the transistors in the clock buffers (Figure 25). In order to generate the bias voltages of the switches (SWN and SWP), the wave shaping circuits amplify the signal coming from the temperature sensor and adjusts its voltage level for suitable switch bias for the temperature range of interest. The forward body bias is applied to all the transistors in the buffer to increase each individual current when the temperature increases and keep their function unaltered.

The area requirement of our adaptive scheme was found to be very small compared to the overall chip area. In our implementation, we found that the used silicon area was

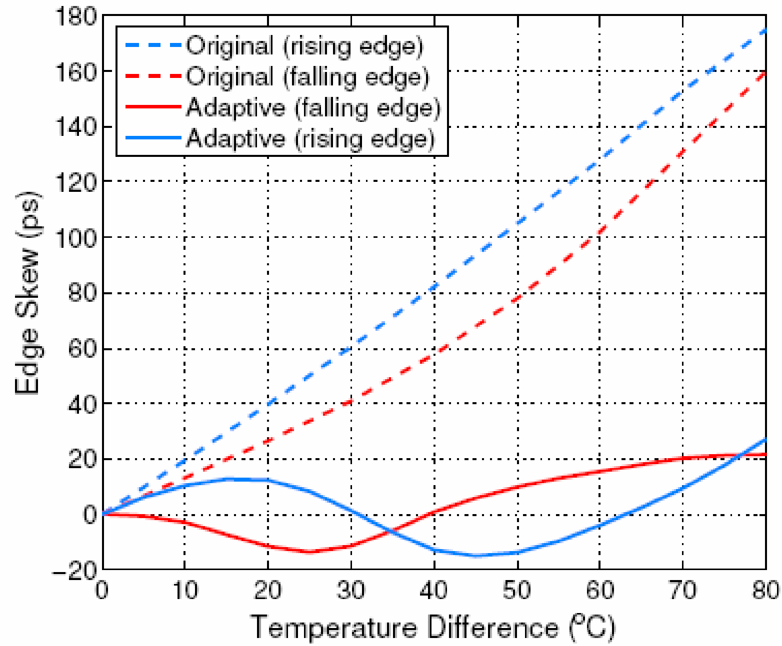

Figure 26. Skew improvement after using the thermally adaptive buffers

approximately 1mm^2. Since the chip area is 400mm^2, our clocking scheme used only 0.25% of the total chip area.

Spice simulations were performed to evaluate the performance of the proposed clock tree design technique. In the simulations, we used the BSIM4 predictive technology models of the 45nm CMOS technology [37] as well as the extracted temperature-dependent parameters of the H-tree discussed. The skew between the clock signals at different output terminals is evaluated for the rising and the falling edges. Figure 26 presents the edge skew variations between the terminals in the left and the right half of the chip (T3_4 and T4_2) versus their temperature difference. The edge skew equals zero when the temperature difference is zero and then it starts to increase with increased temperature difference. With a temperature difference of 80°C, the original buffers suffer edge skews

of 160ps and 175ps on the falling and rising edge respectively. The adaptive buffers reduce the values of the edge skews to 21ps and 27ps on the falling and rising edge respectively. In order to compare our adaptive technique with the theoretical optimal solution that is shown in Table 10 due to exhaustive optimal search, we re-simulated the same temperature profiles using our adaptive technique. The simulation results, shown in Table 11 demonstrate that our adaptive technique achieved an improvement in the clock skew ranging from 54.7% to 99.2%, with an average improvement of 72.4%. As mentioned before, this improvement in the clock skew using our adaptive technique was achieved with a minimal area overhead.

Table 11. Results of skew improvement after using the thermally adaptive buffers

Thermal Profile	Parameters		Unmodified Skew (ps)	Adaptive Buffers Skew (ps)	Adaptive Buffers Improvement (%)
	T_H	T_L			
$T(x) = ax + b$, $a = \frac{T_H - T_L}{L}$, $b = T_L$	170	90	79.6	0.64	99.2
	170	110	61.4	11.6	81.1
	170	130	42.5	15.2	60.0
	170	150	22.5	10.2	54.7
$T(x) = a.e^{-bx}$, $a = T_H$, $b = \frac{1}{L} ln\left(\frac{T_H}{T_L}\right)$	170	90	77.6	14.2	81.7
	170	110	60.5	8.11	86.6
	170	130	42.2	16.4	61.1
	170	150	22.4	10.1	54.9

4.2 3D Clock networks

4.2.1 Introduction for 3D clock networks

3D integration of chips is a promising approach to integrate large systems on a single chip where the problems associated with conventional interconnects are reduced since the average global wire length is reduced drastically [41, 42]. In the 3D integration technology, the overall two dimensional system is divided into a number of smaller size blocks that are fabricated at different layers, stacked on top of each other and connected through vertical interlayer vias [43, 44]. 3D integration offers the possibility of integrating heterogeneous technologies on the same system for high performance SoC, as shown in Figure 27, by fabricating optical, RF, analog and digital chips on different

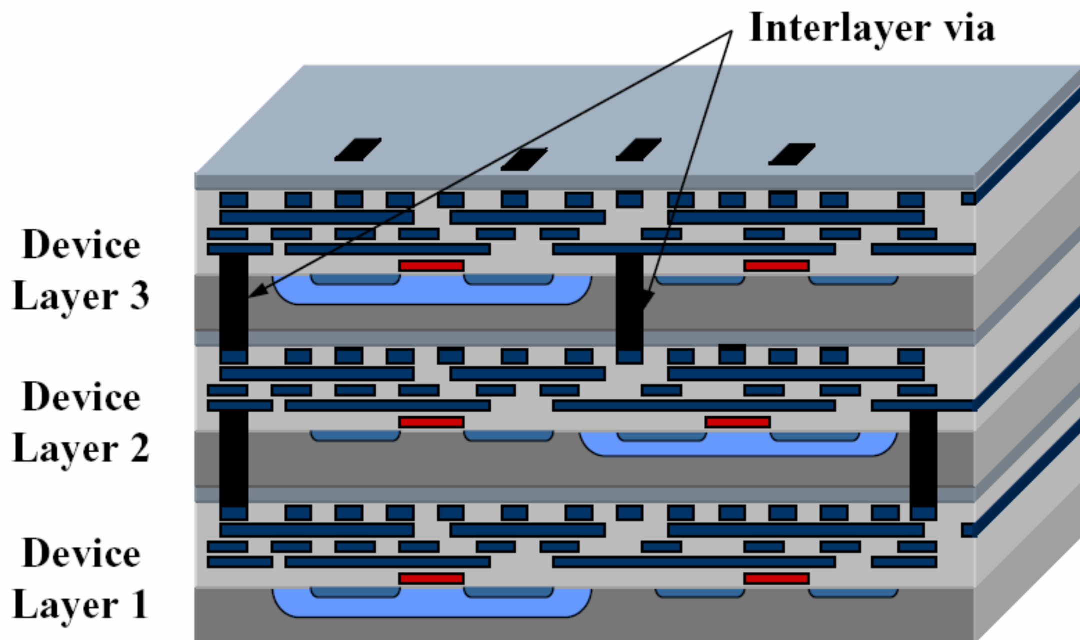

Figure 27. Example of a 3D integrated circuit

layers of the three dimensional chip. This offers much improved noise performance in mixed signal chips due to reduced electromagnetic interference between the digital and analog parts, implemented on different layers, as well as lower substrate noise. Additionally, a greater number of gates can be realized using 3D integration technology leading to larger system integration and integration of traditionally off-chip components on a single 3D chip. For example, the potential to fabricate bigger caches or integration of processor and memory on a single chip exists in the 3D realm. Overall, the main benefit of the 3D integration technique is the possibility of performance enhancement of the global interconnects and the prospect of integrating heterogeneous active layers to enable newer system architectures.

The realization of 3D integrated circuits faces complex challenges. Although there are issues associated with integrating different active layers on a single 3D chip, a major problem in the 3D chips is the inefficient heat dissipation from the chips that leads to thermally induced performance degradation and can reduce the lifetime of the fabricated chips [45]. In conventional 2D chips, the generated heat is dissipated through an external heat sink. In the 3D chips, all the layers will contribute to the generation of heat that has to be dissipated through one or two heat sink(s). Spreading the heat from the layers away from the heat sink is more challenging due to the use of dielectric materials, which have lower thermal conductivity, to isolate the different layers of the chip. The difference in temperature in a single layer and across layers arises because of the structural difference and diversity of computational activities. However, the profiles across the layers vary due

to the varying ease of heat dissipation from the layers as well. For example, the existence of large caches versus relatively small frequently accessed arithmetic units can cause significant temperature differences on a single layer. On the other hand, layers away from the heat sink and situated in between other layers are thermally more insulated than other layers leading to interlayer temperature variation.

As discussed above, the effect of spatial and temporal variation of temperature on the combined performance of the interconnect and devices is a complex phenomenon which becomes even more complicated in a 3D chip because of the uneven variation of temperature across layers. Advancement in cooling and packaging technology alone is not sufficient to fully manage the thermal effects in 3D chips, therefore, ensuring performance under thermal effects is the key to the success of the 3D integration technology.

The focus of this paper is the integrity of the clock signal in the context of 3D integrated circuits. For the synchronous part of a 3D chip, which may be distributed across layers, skewless clock signal is of utmost importance for the accuracy and speed of operation of a design. Since the clock network spans over most parts of the chip and thereby gets exposed to as diverse temperature range as occurs across the chip, the effect of temperature is very pronounced in clock trees. Figure 21 shows an H-tree for a single layer of a 3D chip where we can see that for a number of physically close terminals the clock signals traverse through entirely different temperature zones to reach the terminals

which can lead to significant skew between these terminals. For the clock distribution in the 3D chips, we investigate two different clock tree topologies and compare between them, as discussed in the next section. We demonstrate that our technique can reduce the thermally induced skew by 61.65% on the average.

4.2.2 Proposed clock tree topologies

We consider two alternative clocking schemes for 3D chips in this paper. The first scheme is shown in Figure 28 where the input clock signal, at the center of the clock tree, is fed to each layer through interlayer vias and each layer has its own clock tree with associated clock buffers implemented in the corresponding active layer. We refer to this scheme as the “replica topology”. The second scheme, as depicted in Figure 29, implements the clock tree with the clock buffers on a single layer and using interlayer vias the clock signals from the terminals of the clock tree are passed to all other layers. We call this scheme as the “via topology”.

In the replica topology, the clock tree for each layer can be customized according to the temperature profile of the layer. For example, the average temperature of a layer away from the heat sink is likely to be higher than the average temperature of the layer attached to the heat sink, which can be handled in the first scheme. However, the obvious disadvantage of this scheme is the design overhead, both in terms of resources and design efforts required for the layer wise customization. Moreover, because of the separate customization of the different layers, the skew between terminals in different layers may

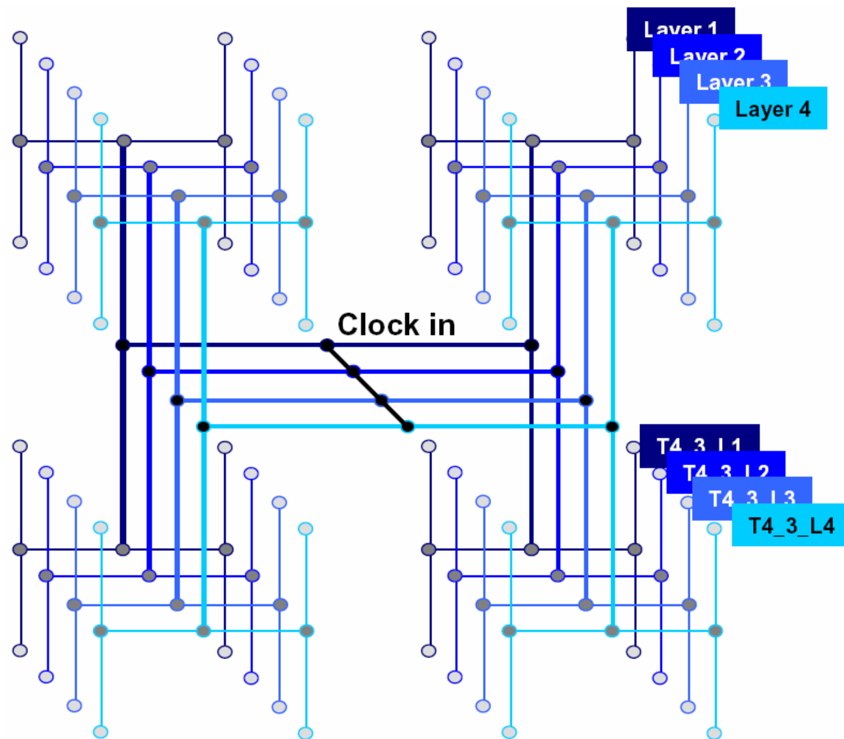

Figure 28. H-tree repeated in each layer, fed by a common clock

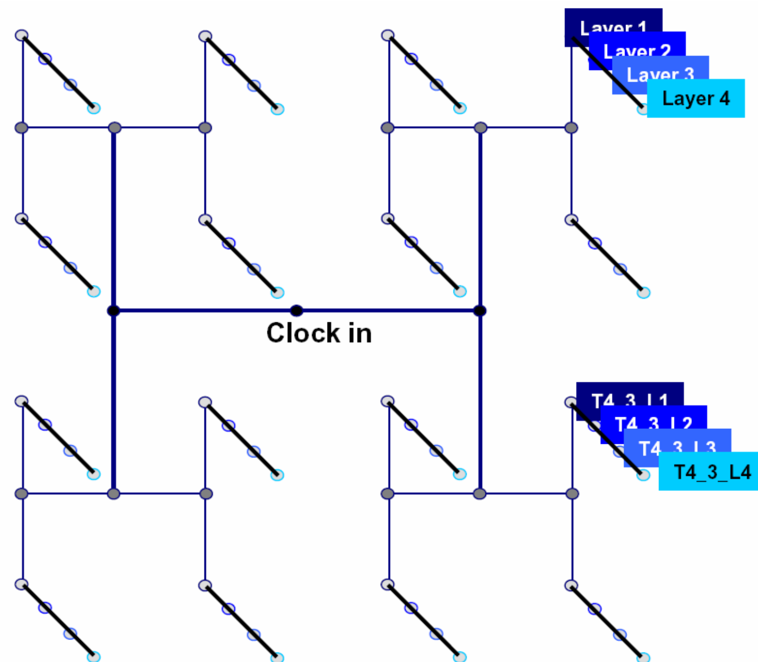

Figure 29. H-tree on single layer with clock propagating through vias

be high even if the skew is low in the same layer. On the other hand, the via topology will provide uniform skew compensation across layers since the same terminal clock signals are transmitted across layers. Additionally, the second scheme obviously has the advantage of less design overhead. In Section 3.2.4, we compare these two schemes. In particular, we compare the clock skew improvement achieved by the two schemes for a number of temperature profiles as well as the corresponding power requirements. It can be understood that the replica topology will require N times more area and wiring resources compared to the via topology, where N is the number of layers in the 3D clock tree.

4.2.3 Clock tree modeling for including thermal effects

For synchronous modules spread across multiple layers in a 3D chip, the arrival times of the clock signal at different communicating clock sinks should ideally have zero skew for the proper functioning of the chip. Therefore, the clock distribution network needs to be carefully designed so that the clock signal arrives at each terminal (ideally) at the same time. Designing the clock sinks at equal distances from the clock source is one of the ways of distributing the clock with low clock skew, as found in the H-trees most commonly used in regular array based designs [19]. On-chip temperature variation in a 3D chip happens to be a major source of clock skew since the uneven spatial variation of temperature leads to unequal change in the wire resistance and clock buffer characteristics along different paths. In this work, we consider two different schemes for clocking 3D synchronous modules and develop SPICE models for analyzing the clock

skew under temperature variation. The interconnect wires are modeled by their RLC

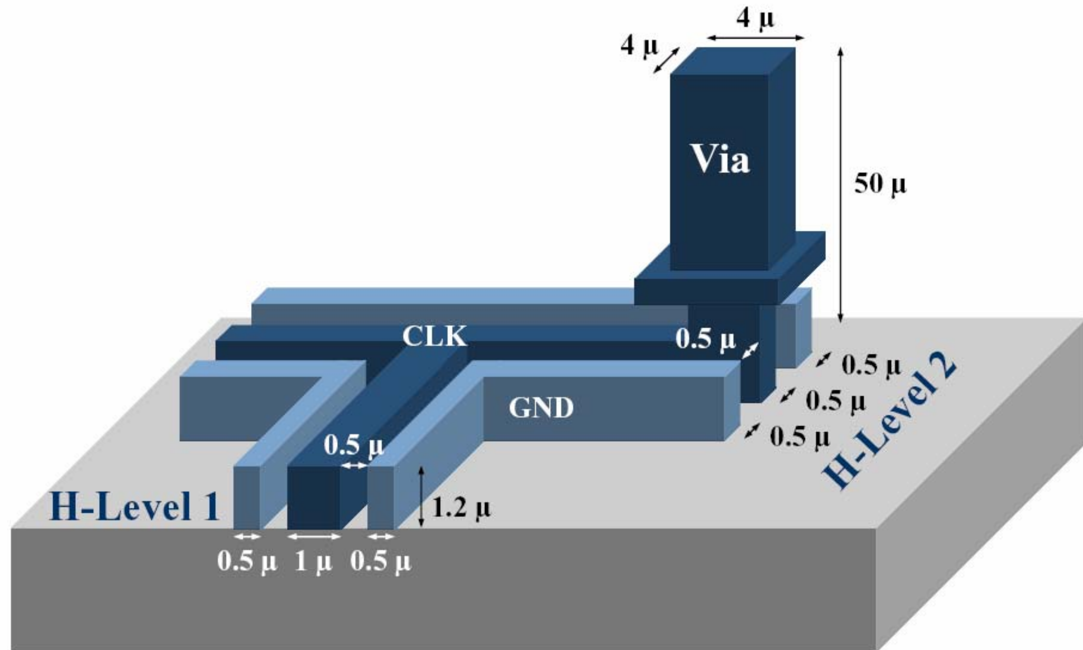

Figure 30. Dimensions of the wires and vias in the clock trees

parasitic components and the devices are modeled using the BSIM4 model. The details of the dynamically adaptive buffers were mentioned before.

In this work we consider a shielded H-tree in a 3D chip with four layers designed in the 45nm technology node having a die length of 2cm. For illustrating the effects of temperature variation on clock tree, we consider a 2 level H-tree, as depicted in Figure 21, where the first-level H has sides of length 10mm and the second-level H's have sides of length 5mm. The dimensions of the wires are shown in Figure 30. Clock buffers of sizes gradually reducing from the clock source toward the clock sink have been inserted in the clock tree, as depicted in Figure 21

4.2.4 Results

The performance of the proposed thermally robust clock trees was evaluated by SPICE simulations using the BSIM4 predictive technology models for the 45nm CMOS technology node [37]. The simulations were performed for both the topologies presented in Section 3.2.2. To demonstrate the efficacy of our method, we measure the skew between two physically close terminals, T1_3 in layer 1 and T2_1 in layer 4, in the three dimensional clock tree under different temperature profiles. [46] suggests that up to five

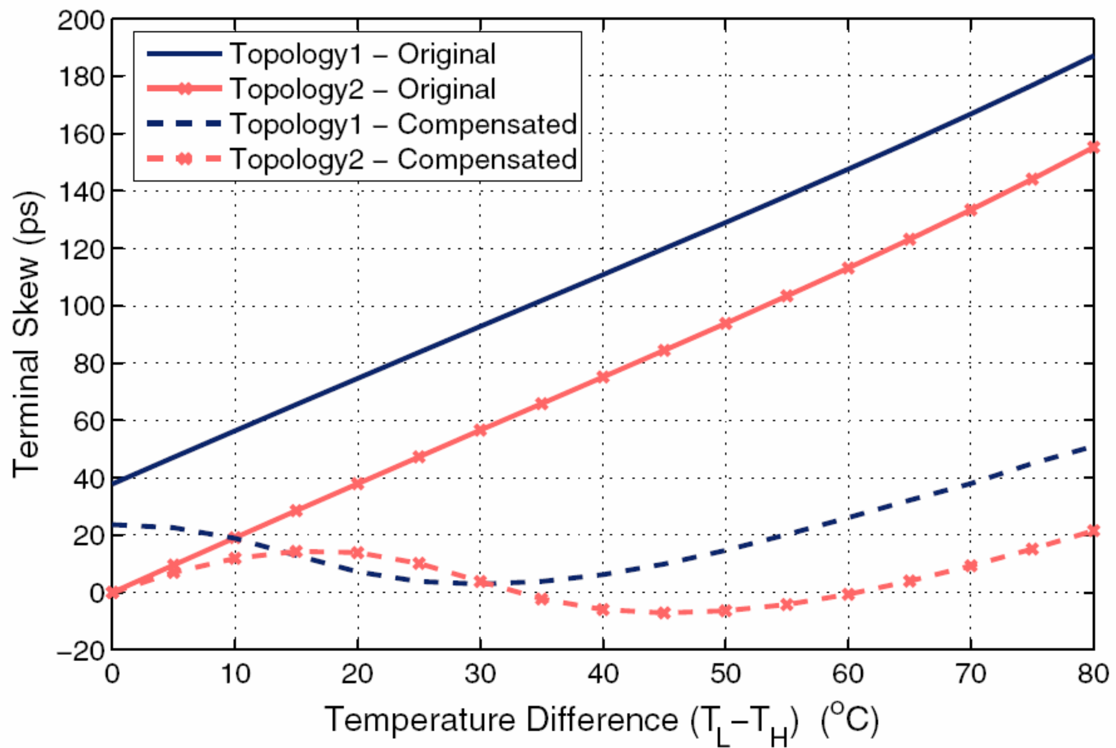

Figure 31. Skew improvement for the split profile using adaptive buffers

layers, the 3D temperature profile will be dominated by the heating of the first layer and the temperature rises almost linearly with the number of layers. In our experiments we assume that the temperature increases by 5°C for each layer away from the heat sink.

We start with the split thermal profile as shown in Figure 21 where the left and right halves of a layer are fixed at temperatures TL,i and TR,i, respectively. The skew variation with temperature difference (TL,i −TR,i) is illustrated in Figure 31. Note that the skew in the via topology is lower than the skew obtained in the replica topology since the vertical temperature gradient increases the skew in the replica topology. The adaptive buffers reduced the temperature dependence of the skew as shown in Figure 31. For the replica topology, the maximum compensated skew was 55ps as opposed to the original maximum skew of 188ps. On the other hand, the compensated skew was within ± 15ps over the 80°C temperature variation for the via topology, while the original maximum skew was 155ps over the same temperature range. The localization of the clock-tree in one layer makes the compensation in the via topology more effective than the replica topology. We performed additional simulations for two more generic thermal profiles. These profiles were extended from 2D generic cases modeled in [30] to the 3D case. The first profile models a linear temperature fall off across each layer of the 3D chip whereas the second profile is an exponential fall off of temperature across each layer. As assumed in case of the split thermal profile, the temperature difference between layers is 5°C. The simulation results for the two proposed topologies have been shown in Table 12. Note that the skew values reported in the table represents the maximum skew between any pair of terminals in the 3D clock tree within √2 times the distance between the nearest

Table 12. Results of skew improvement after using the thermally adaptive buffers

Thermal Profile	Parameters		Replica Topology			Via Topology		
	T_H (^oC)	T_L (^oC)	Unmodified Skew (ps)	Compensated Skew (ps)	% Improvement	Unmodified Skew (ps)	Compensated Skew (ps)	% Improvement
Linear:	100	20	103.50	29.26	71.73	61.08	23.83	60.99
$T(x,y)=a(x+y)+b$	100	40	92.32	37.10	59.81	47.58	11.06	76.75
$a=\frac{T_H-T_L}{2L}$	100	60	81.15	43.44	46.47	33.62	11.62	65.44
$b=T_L$	100	80	68.99	42.52	38.37	18.22	9.18	49.62
Exponential:	100	20	94.22	26.99	71.35	54.72	24.49	55.24
$T(x,y)=a.e^{-b(x+y)}$	100	40	88.09	29.88	66.08	45.18	15.77	65.10
$a=T_H$	100	60	79.51	40.92	48.53	32.82	9.87	69.93
$b=\frac{1}{2L}ln\left(\frac{T_H}{T_L}\right)$	100	80	68.63	42.28	38.39	18.07	9.01	50.14

terminals since the clock skew is more important for physically close terminals. Based on the results from the table, we find that the average unmodified skew is 2.17 times more in the replica topology than the via topology. It can be noticed that, on the average, our skew reduction technique reduces the skew by 55.09% for the replica topology and by 61.65%

for the via topology. Moreover, the average skew after the correction is approximately 2.5 times less in the via topology as compared to the average compensated skew obtained for the replica topology. We also compared the power consumption of the two clock tree topologies. Intuitively, the replica topology will consume approximately N times more power and area than the via topology, where N is the number of layers in the 3D clock tree. However, since the load driven by the single layer of clock tree in the via topology is equal to the combined load driven by all the layers of the replica topology, the ratio of power consumed by the first and second topologies will be less than N. The replica topology consumes 2.35 times more power, making the via topology the obvious choice for designing 3D clock trees.

Chapter 5

Conclusion and future work

This thesis presented a study of the reliability of certain key components of high performance digital systems. The focus was on register files, clock distribution, and caches. The reliability methods presented in this thesis can be used under a variety of circumstances improving the resiliency of these key components.

The redundancy and parity described increases the reliability of the register file. This minimizes the likelihood of any errors generated being propagated to the rest of the system thereby localizing and correcting the error. This is important due to the environmental nature and frequency of accesses of the register file.

Only a monolithic, dual banked, and quadruple banked register file was considered. The banked versions had a total of 16 read ports subdivided among the banks. The 8 write ports were duplicated for each bank to keep the contents of the banks the same. One potential aspect that could be explored is reducing the number of write ports and exploring the effect that has on performance. This could be coupled with a study of trading off reliability with redundancy.

The analysis and impact of NBTI along with process and temperature variations on different SRAM cells in low power cache configurations was performed. The ability to reliably recover a substantial portion of the lost noise margins is shown. The effect is

investigated for both symmetric and asymmetric 6-T SRAMS and 8-T SRAMS and similar trends of decreasing static noise margins are noticed for all three configurations. The 8-T SRAMS initially have the highest read noise margins and the asymmetric 6-T SRAMS having the lowest read and write noise margins. The relative order of the absolute values of the noise margins remain the same after degradation suggesting that the 8-T SRAMS are the best overall SRAM structure before and after NBTI degradation. Finally the process variation effects is investigated noticing a decrease in the read noise margin over time and a lowering of the leakage current due to the increased threshold voltages.

There are other types of SRAM cells than just the 6 and 8 transistor SRAMs. These range from a 1-T pseudo-static to 10-T SRAMs. These different configurations are targeted for a variety of operating environments such as low leakage, higher noise margins, faster access, and smaller sizes. An expanded study to include these varieties of SRAMS may yield an improved understanding of the tradeoffs that need to be considered before choosing a particular type of SRAM.

One of the most influential factors in determining NBTI is temperature. Temperature also has an effect on the propagation speeds of signals through interconnect wires. The arrival times of signals may jitter due to temperature variations across an integrated circuit. The clock signal can be heavily affected by temperature variation as it is distributed throughout the majority of the chip. Accurate timing is required for high

performance synchronous circuits and the ability to maintain low skew was demonstrated. This was furthermore extended to three dimensional clock distribution networks with an innately superior architecture proposed.

Combining a study of temperature induced clock skew and NBTI degradation would be one direction for future study. The portions of the clock tree that are experiencing the highest temperature will also be the slowest paths and the buffers suffering the most NBTI threshold changes. This will lead to further clock skew in the paths that are already the worst ones. Certain designs use a mesh structure to reduce the clock skew while other designs use delay locked loops. These were not investigated but will also be affected by temperature and NBTI. The negative effects of register file reliability, clock skew, SRAM noise margin, and NBTI have been studied and solutions to mitigate them have been proposed.

References

1. S. Borkar, et. al., "Parameter variation and impact on circuits and microarchitecture", in DAC, 2003.

2. Y. Zhang, et. al., "HotLeakage: An Architectural, Temperature-Aware Model of Subthreshold and Gate Leakage". University of Virginia Dept. of Computer Science Tech. Report CS, 2003.

3. K. P Cheung. "Temperature effect on ultra thin SiO_2 time-dependent-dielectric-breakdown", in PPID, 2003.

4. R. Vattikonda, et. al. "Modeling and minimization of PMOS NBTI effect for robust nanometer design", DAC, 2006.

5. G. Sery, et al., "Life is CMOS: why chase the life after?" DAC, 2003.

6. J. Tschanz, et al., "Dynamic-Sleep Transistor and Body Bias for Active Leakage Power Control of Microprocessors", ISSCC, 2003.

7. S. Gochman, et. al., "Introduction to Intel Core Duo Processor Architecture", ITC, 2006.

8. A. J. Ricketts, et. al., "Investigating Simple Low Latency Multiported Register Files", ISVLSI, 2007.

9. SPEC 2000 Benchmark, http://www.spec.org.

10. A. S. Palacharla, N. P., et. al., "Complexity-effective Superscalar Processors", ISCA, 1997.

11. M. Tremblay, B. Joy, and K. Shin, "A Three Dimensional Register File for Superscalar Processors", HICSS, 1995.

12. R.E. Kessler, "The Alpha 21264 Microprocessor", Micro, 1999.

13. A. Seznec, E. Toullec, and O. Rochecouste, "Register Write Specialization Register Read Specialization: a Path to Complexity Effective Wide-issue Superscalar Processors", Micro, 2002.

14. V.V. Zyuban and P.M. Kogge, "Inherently Lower-power High Performance Superscalar Architectures", IEEE Transactions on Computers, 2001.

15. R. Balasubramonian, S. Dwarakadas, and D.H. Albonesi, "Reducing the complexity of the register file in dynamic superscalar processors", Micro, 2001.

16. N.S. Kim and T. Mudge, "Reducing Register Ports Using Delayed Write-Back Queues and Operand Pre-Fetching", ICS, 2003.

17. A. Dasgupta and R. Karri, "Switch level hot-carrier reliability enhancement of VLSI circuits", DFT, 1995.

18. S. V. Kumar, C. H. Kim, and S. S. Sapatnekar, "Impact of NBTI on SRAM Read stability and Design for Reliability", ISQED, 2006.

19. J. M. Rabaey, A. Chandrakasan, and B. Nikolic, Digital Integrated Circuits, A Design Perspective, Second Edition, Prentice-Hall, Upper Saddle River, NJ, 2003.

20. N. Seifert et al., "Historical Trend in Alpha-particle Induced Soft Error Rates of the Alpha Microprocessor", IRPS, 2001.

21. D.A. Patterson, G. Gibson, and R.H. Katz, "A case for redundant arrays of inexpensive disks (RAID)", In Proc. of International Conference on Management of Data, 1988.

22. G. Hinton et al., "The Microarchitecture of the Pentium 4 Processor", ITJ, 2001.

23. G. Memik, G. Reinman, and W.H. Mangione-Smith, "Precise Instruction Scheduling", ILP, 2005.

24. D. Burger and T. Austin, "The Simplescalar Tool Set", Computer Science Department, University of Wisconsin-Madison, 1997.

25. L. Li et al., "Soft Error and Energy Consumption Interactions: a Data Cache Perspective", ISLPED, 2004.

26. P. Shivakumar and N.P. Jouppi, "CACTI 3.0: an Integrated Cache Timing, Power, and Area Model", Technical report TN-2001/2, Compaq Western Research Laboratory, 2001.

27. M. Mondal, A. J. Ricketts, et. al, "Mitigating Thermal Effects on Clock Skew with Dynamically Adaptive Drivers", ISQED, 2007.

28. M. Mondal, A. J. Ricketts, et. al, "Thermally Robust Clocking Schemes for 3D Integrated Circuits", DATE, 2007.

29. K. Shakeri and J. Meindl, "Temperature Variable Supply Voltage for Power Reduction", ISVLSI, 2002.

30. A. H. Ajami, K. Banerjee, and M. Pedram, "Modeling and Analysis of Non-Uniform Substrate Temperature Effects in High Performance VLSI," IEEE Transactions on Computeraided Design of Integrated Circuits and Systems, 2001.

31. M. Cho, S. Ahmed, and D. Z. Pan, "TACO: Temperature Aware Clock-Tree Optimization", ICCD, 2005.

32. A. Chakraborty et al., "Dynamic Thermal Clock Skew Compensation Using Tunable Delay Buffers", ISLPED, 2006.

33. H. Heeb and A. Ruehli, "Three-Dimensional Interconnect Analysis Using Partial Element Equivalent Circuits", IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications, 1992.

34. M. Kamon, M. J. Tsuk, and J. White, “Fasthenry: A Multipole-Accelerated 3-D Inductance Extraction Program”, IEEE Transactions on Microwave Theory and Techniques, 1994.

35. K. Nabors and J. White, “Fastcap: A Multipole Accelerated 3-D Capacitance Extraction Program”, IEEE Transactions on Computer-aided Design of Integrated Circuits and System, 1991.

36. A. Ruehli, “Inductance Calculations in a Complex Integrated Circuit Environment”, IBM Journal of Research and Development, 1972.

37. W. Zhao and Y. Cao, “New Generation of Predictive Technology Model for sub-45nm Design Exploration", ISQED , 2006.

38. Q. Chen, M. Meterelliyoz, and K. Roy, "A CMOS Thermal Sensor and Its Applications in Temperature Adaptive Design", ISQED, 2006.

39. K. Ishibashi, T. Yamashita, Y. Arima, I. Minematsu, and T. Fujimoto, “A 9µW 50MHz 32b Adder Using Self- Adjusted Forward Body Bias in SoCs", ISSCC, 2003.

40. H. Ananthan, C. H. Kim, and K. Roy, “Larger-than-Vdd Forward Body Bias in Sub-0.5V NanoscaleCMOS”, ISLPED, 2004.

41. H. Kurino et. al., "Intelligent Image Sensor Chip with Three Dimensional Structure", IEDM, 1999.

42. Y.-F. Tsai, Y. Xie, N. Vijaykrishnan, and M. J. Irwin, "Three-Dimensional Cache Design Exploration Using 3DCacti," ICCD, 2005.

43. A. Rahman, D. Antoniadis, and A. Agarwal, “System Level Performance Evaluation of Three-Dimentional Integrated Circuits", in IEEE Transactions on VLSI Systems, 2000.

44. R. H. Havemann and J. A. Hutchby, "High-Performance Interconnects: An Integration Overview," Proceedings of IEEE, 2001.

45. K. Banerjee, S. J. Souri, P. Kapur, and K. C. Saraswat, "3-D ICs: A Novel Chip Design for Improving Deep-Submicrometer Interconnect Performance and Systems-on-Chip Integration", Proceedings of the IEEE, 2001.

46. G. M. Link and N. Vijaykrishnan, "Thermal Trends in Emerging Technologies", ISQED, 2006.

47. “International technology road map for semiconductors, test and test equipment,” http://public.itrs.net/, 2006.

48.N. Kimizuka, et. al, “The impact of bias temperature instability for direct-tunneling ultra-thin gate oxide on mosfet scaling,” in VLSI Technology, 1999, Digest of Technical Papers. Symposium on, 1999.

49. D. K. Schroder and J. A. Babcock, “Negative bias temperature instability: Road to cross in deep submicron silicon semiconductor manufacturing,” Journal of Applied Physics, vol, 94, no. 1, pp. 1-18, 2003.

50. A. Krishnan, et. al., NBTI impact on transistor and circuit: models, mechanisms and scaling effects [mosfets]," in Electorn Devices Meeting, 2003. IEDM '03 Technical Digest. IEEE International, Dec. 2003.

51. S. Mahapatra, P. Kumar, and M. Alam, "Investigation and modeling of interfact and bulk trap generation during negative bias temperature instability of p-mosfets," Electorn Devices, IEEE Transactions on, September 2004.

52. B. Paul, et. al., "Impact of NBTI on temporal performance degradation of digital circuits," Electron Device Letters, IEEE, Aug 2005.

53. R. Wittmann, et. al., "Impact of random bit values on NBTI lifetime of an SRAM cell," in Physical and Failure Analysis of Integrated Circuits, 2006. International Symposium on, July 2006.

54. N. Azizi, F. Najm, and A. Moshovos, "Low-leakage asymmetric-cell SRAM," Very Large Scale Integration (VLSI) Systems, IEEE Transactions on, Aug. 2003.

55. A. Moshovos, B. Falsafi, F. Najm, and N. Azizi, "A case for asymmetric-cell cache memories," Very Large Scale Integration Systems, IEEE. Transactions on, July 2005.

56. N. Verma and A. P. Chandrakasan, "A 256kb 65nm 8T Subthreshold SRAM Employing Sense-Amplifier Redundancy," IEEE Journal of Solid State Circuits, Jan. 2008.

57. L. Chang, et. al., "An 8t-SRAM for variability tolerance and low voltage operation in high-performance caches," Solid State Circuits, IEEE Journal of, April 2008.

58. J. Singh et. al., "Single ended 6t SRAM with isolated read-port for low-power embedded systems," Design, Automation & Test in Europe Conference, April 2009.

59. M. A. Alam and S. Mahapatra, "A comprehensive model of PMOS NBTI degradation," Microelectronics Reliability, 2005.

60. Y. J. Chang and F. Lai, "Dynamic zero-sensitivity scheme for low-power cache memories," Micro, IEEE, July-Aug 2005.

61. E. Seevinick, F. List, and J. Lohstroh, "Static-noise margin analysis of MOS SRAM cells," Solid-State Circuits, IEEE Journal of, Oct. 1987.

62. E. Grossar, et. al., "Read stability and write-ability analysis of SRAM cells for nanometer rechnologies," Solid State Circuits, IEEE Journal of, Nov 2006.

63. L. Li, et. al, "Leakage energy management in cache hierarchies," in Parallel Architectures and Compilation Techniques, Proceedings, 2002.

64.PTM, "Predictive technology model," in Nanoscale Integration and Modeling (NIMO) Group at ASU, http://www.eas.asu.edu/ptm/, 2008.

65. A.J. Bhavnagarawala, et. al., "The impact of intrinsic device fluctutations on CMOS SRAM cell stability," IEEE Journal of Solid State Circuits, Apr. 2001.

VITA

Andrew Ricketts

Andrew Ricketts received a Capstone Scholarship to attend Howard University in 1999. He received his Bachelor of Science in Electrical Engineering from the Department of Electrical Engineering at Howard University in the Spring of 2003. He joined the PhD program in Computer Science and Engineering in the Fall of 2003. He was a Teaching Assistant for the Introductory Digital Design Lab for a number of semesters. He was also the Teaching Assistant for the Computer Architecture course and an instructor for the Field Programmable Gate Arrays course. His research interests revolve around the design of energy efficient and reliable architectures. He has given a research talk in Germany and has published papers in various conferences and journal. Additionally his research interests include alternate to CMOS technologies, clock distribution, and lifetime reliability.