

COUNTERING AGING EFFECTS THROUGH FIELD GATE SIZING

A Thesis

by

TRENTON DEAN HENRICHSON

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

December 2008

Major Subject: Computer Engineering

COUNTERING AGING EFFECTS THROUGH FIELD GATE SIZING

A Thesis

by

TRENTON DEAN HENRICHSON

Submitted to the Office of Graduate Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

Approved by:

| | |
|---------------------|--------------------|
| Chair of Committee, | Jiang Hu |
| Committee Members, | Sunil Khatri |
| | Weiping Shi |
| | Duncan Walker |
| Head of Department, | Costas Georghiadis |

December 2008

Major Subject: Computer Engineering

ABSTRACT

Countering Aging Effects through Field Gate Sizing. (December 2008)

Trenton Dean Henrichson, B.S., Southern Methodist University

Chair of Advisory Committee: Dr. Jiang Hu

Transistor aging through negative bias temperature instability (NBTI) has become a major lifetime constraint in VLSI circuits. We propose a technique that uses antifuses to widen PMOS transistors later in a circuit's life cycle to combat aging. Using HSPICE and 70nm BPTM process numbers, we simulated the technique on four circuits (a ring oscillator, a fan-out four circuit, an ISCAS c432 and c2670). Over the lifetime of the circuit, our simulations predict a 8.89% and a 13% improvement in power in the c432 and c2670 circuits respectively when compared to similarly performing traditional circuits.

TABLE OF CONTENTS

| | Page |
|---|------|
| ABSTRACT..... | iii |
| TABLE OF CONTENTS..... | iv |
| LIST OF TABLES..... | vi |
| LIST OF FIGURES..... | vii |
| 1. INTRODUCTION..... | 1 |
| 2. TRANSISTOR AGING..... | 3 |
| 2.1 NBTI..... | 3 |
| 3. AGING COUNTERMEASURES..... | 6 |
| 3.1 Overdesign..... | 6 |
| 3.2 V_{dd} and V_{th} Tuning..... | 6 |
| 3.3 Forward Body Bias..... | 8 |
| 3.4 Stack Effect..... | 8 |
| 3.5 Duty Cycle..... | 10 |
| 3.6 Summary of Current Countermeasures..... | 10 |
| 4. FIELD TRANSISTOR SIZING..... | 12 |
| 4.1 The Circuit..... | 12 |
| 4.2 Simulation..... | 16 |
| 5. RESULTS AND ANALYSIS..... | 19 |
| 5.1 Simple Ring Oscillator..... | 19 |
| 5.2 Simulated Fan-out Test..... | 27 |
| 5.3 ISCAS Circuit Performance..... | 33 |
| 6. CONCLUSIONS..... | 47 |
| REFERENCES..... | 50 |

| | Page |
|---------------|------|
| APPENDIX..... | 52 |
| VITA..... | 57 |

LIST OF TABLES

| TABLE | Page |
|--|------|
| I Specifics of Antifuse Circuits..... | 14 |
| II Ring Oscillator Results..... | 20 |
| III Worst Case Ring Oscillator..... | 25 |
| IV Percentage Change Ring Oscillator | 26 |
| V Fan-out Four Results..... | 28 |
| VI Change in Power for Fan-out Four Circuits..... | 31 |
| VII Change in Period Fan-out Four circuits | 33 |
| VIII Predicted Gate Sizes for ISCAS Circuits..... | 36 |
| IX Change in Power for ISCAS Circuits | 38 |
| X Performance Difference for ISCAS Circuits..... | 40 |
| XI Power and Performance Change in Partial FTS ISCAS C432 Circuit..... | 44 |
| XII ISCAS C432 Fully Modified Results..... | 52 |
| XIII ISCAS C432 Partially Modified Results..... | 53 |
| XIV ISCAS C2670 Results..... | 54 |
| XV ISCAS CELL Libraries..... | 55 |
| XVI Antifuse Turn on Points..... | 56 |

LIST OF FIGURES

| FIGURE | Page |
|--|------|
| 1 Percentage Degradation of V_{th} Versus Time. © [2005] IEEE..... | 4 |
| 2 Absolute Change in ΔV_{th} Versus % Change in V_{dd} , V_{th} , and L_{eff} . © [2006] IEEE..... | 7 |
| 3 Static Conditions on 2-Input Nor. © [2006] IEEE..... | 9 |
| 4 FTS Inverter Circuit..... | 13 |
| 5 Behavior of New and Aged Ring Oscillator Circuits (a) Power Consumption (b) Propagation Delay..... | 22 |
| 6 D/SPMOS Ratio Versus Performance Ring Oscillator..... | 23 |
| 7 Behavior of New and Aged Fan-out Four Circuits (a) Power Consumption (b) Propagation Delay..... | 29 |
| 8 D/SPMOS Ratio Versus Performance Fan-out Four..... | 30 |
| 9 D/SPMOS Ratio Versus Performance in ISCAS Circuits (a) ISCAS c422 Partially Modified (b) ISCAS c432 Fully Modified (c) ISCAS c2670 | 35 |
| 10 Difference (%) in Power and Delay for Fully FTS Modified Circuits Compared to Similar Standard Circuits (a) ISCAS c432. (b) ISCAS c2670 | 42 |
| 11 Change in Power and Delay Compared to Similar Standard Circuits for the ISCAS c432 Partially Modified Circuits..... | 46 |

1. INTRODUCTION

Many of the established trends of VLSI circuits are contributing to an overall negative trend in product reliability. Transistors continue to scale rapidly. This leads to transistor gates that are thinner and narrower while at the same time the electrical current across the gates is increasing [1]-[3]. All these factors lead to an increase in negative bias temperature instability (NBTI) in PMOS transistors. NBTI is now recognized as a dominant aging factor in CMOS circuits [3], [4], and it will play a crucial role in the lifetime of VLSI technology in the future.

In order to keep reliability constant while continuing to increase performance and, in turn, scale transistors, many designers have resorted to oversizing PMOS transistors in anticipation of future degradation [1], [4], but this leads to an overall increase in circuit power. A second approach is the use of adaptive techniques, such as voltage scaling and body bias manipulation [5]-[7]. However, isolating voltage or body bias requires a large overhead in added nets and gates [5]. If the relative overhead is diminished by applying these techniques to a large portion of the circuit, this can have inverse effects on non-critical nets.

This paper investigates the use of antifuse technology to effectively resize PMOS transistors in later stages of a circuit's lifetime. Unlike previous techniques, field transistor sizing (FTS) is less wasteful of power than initial overdesign, and it is more easily targeted to fine grained critical paths than voltage scaling or body bias manipulation.

This thesis follows the style of *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*.

The technique was simulated using HSPICE and applied to ISCAS combinational circuits c432 and c2670. Both circuits showed modest improvements in performance and at least an 8% reduction in power when compared to similarly performing standard circuits (over the lifetime of the circuit

The rest of this paper is divided into sections. Section 2 further explains the mechanisms of transistor aging and the importance of NBTI countermeasures. Section 3 analyzes the application and the pitfalls of overdesign and adaptive age prevention techniques. Section 4 discusses the FTS technique and testing procedure in detail, including assumptions made and models used. Section 5 covers results and analyses of these tests. Finally, Section 6 consists of conclusions and suggestions for future research in this area.

2. TRANSISTOR AGING

2.1 NBTI

Negative bias temperature instability (NBTI) is the gradual increase in V_{th} after years of negative voltage are applied across the gate. It is the primary mechanism for aging of PMOS transistors [8]. Its effects on NMOS transistors are insignificant.

While there is still some controversy about the primary mechanism for NBTI, most researchers believe it is driven by the generation of positive ion traps in the Si-SiO₂ interface [8]-[10]. Hydrogen atoms are used in the Si-SiO₂ interface "tie off" loose silicon bonds [8], [9].

When a long term negative electric field or a very high short term field is applied across a gate, these bonds break apart and release H⁺ atoms. The H⁺ atoms then diffuse through the interface and form interface traps. Past research has shown that interface trap density N_{it} is related to the size of the electric field across the gate, and the age of the circuit in (1) [2]

$$\Delta N_{it}(E_{ox}, t) = t^{0.25} \chi (E_{ox} e^{E_{ox}/E_0})^{1/2}. \quad (1)$$

In Equation (1) χ represents the product of all field/time independent terms. The effect of the interface trap is to increase V_{th} and in turn reduce drive strength and performance. The relative degradation in voltage threshold due to trap degeneration can be related to the electric field and time in (2) [2]

$$\Delta V_{th}(E_{ox}, t) = (1 + m) [q t^{0.25} \chi (E_{ox} e^{E_{ox}/E_0})^{1/2}] / C_{ox}. \quad (2)$$

In (2) m accounts for V_{th} shift due to degradation in carrier mobility from other sources. Fig. 1 and charts the relative change in V_{th} over time (sec.) in a 70nm BPTM process [2].




Fig. 1. Percentage Degradation of V_{th} Versus Time. © [2005] IEEE

Using well known formulas for delay through a gate, in (3), we can approximate the change in delay through the gate relative to the change in V_{th} to be equal to (4) [2].

$$\tau = (C_L V_{dd}) / I_d = K_I / (V_g - V_{th})^\alpha \quad (3)$$

$$\Delta\tau / \tau = \alpha \Delta V_{th} = (\alpha \Delta V_{th}) / (V_g - V_{th}) \quad (4)$$

Since $(V_g - V_{th}) > V_{th}$ and α is close to 1, this means the degradation in performance due to NBTI is less than the degradation in V_{th} . Past research has shown that in an ISCAS c432 circuit with a BPTM 70nm process, the simulated performance degradation is 8.9%. Our own research used the 65 nm BPTM, in which we calculated a

10.8% change in performance. The discrepancy could be explained by the differences in process, choice of critical path, and P/N ratio.

As transistors continue to scale, E_{ox} will rise. This will cause the effects of NBTI to be increased with new technologies [9], [10]. Already E_{ox} is seen as a dominant mechanism in CMOS degradation [3], [4]. For this reason, it has been the focus of our research.

3. AGING COUNTERMEASURES

There has been little research into design specifically to address the effects of NBTI aging. However, many of the design techniques used to counteract process variation can be applied to combat NBTI aging.

3.1 Overdesign

The simplest countermeasure is overdesign, usually accomplished through gate sizing. Designers simply make transistors wider and therefore faster than they would normally be so after they age, they are still within standard operating margins [4]. However, by making transistors larger than they need to be, the designer is also making them consume more power than they would otherwise consume.

A slight variation on pure gate sizing is to oversize only the PMOS sections of the gate. Research has shown that modulating only the PMOS transistors is effective at minimizing the aging effect NBTI with a slightly lower cost in overall power. However, this is not feasible for large aging effects [5]. There is a limit to how far a P/N ratio can be skewed before the designer creates a performance cost in one transition in excess of any gains the designer would make later in the circuit's life. Even if PMOS sizing does work, this still causes an unnecessary penalty in power.

3.2 V_{dd} and V_{th} Tuning

Previous research has also found that aging effects are very sensitive to V_{dd} and V_{th} . These sensitivities are shown in Fig. 2 [5]. Therefore some designers attempt to minimize ΔV_{th} by altering a circuit's, V_{dd} or V_{th} .




Fig. 2. Absolute Change in ΔV_{th} Versus % Change in V_{dd} , V_{th} , and L_{eff} . © [2006] IEEE

However, if a designer lowers V_{dd} this has the effect of lowering performance. It is complicated to place multiple V_{dd} corners on a single circuit, and it becomes more burdensome as more distinct corners are added. Thus, it is impractical to apply V_{dd} tuning to targeted paths or devices, and the entire circuit would suffer to compensate for an effect that would only take place in a few critical areas.

Alternatively a designer could tune the V_{th} of a new circuit. By raising the V_{th} of a new cell a designer can reduce the change in V_{th} that will occur over the circuit's lifetime. However this is essentially pre-aging the circuit, which means the designers will once again suffer an upfront penalty in performance.

While voltage threshold tuning can be applied on a finer granularity in high performance circuits, its V_{th} is already carefully chosen in order to minimize power.

Increasing V_{th} in areas in which it has not already been set high for power reasons would have negative effects on performance.

3.3 Forward Body Bias

A more adaptive technique to countering NBTI effects is the use of forward body bias (FBB) to lower the V_{th} in later stages of the circuit's life cycle. This method is already being used to lower die-to-die process variations on large high performing circuits [6]. However, the circuitry for finding and locking in the right amount of FBB to create a balance of power and performance is nontrivial [6], and as the affected areas become smaller, the proportional overhead becomes larger. As a result, much like V_{dd} modulation, a FBB technique is not fine grained enough to be an ideal solution to the NBTI problem.

Furthermore, when FBB is used to combat process variation, the proper variation is calculated during burn in. This assures that the FBB will still be functionally correct under worst case working conditions [6]. Clearly, it is impractical to submit an IC to a second burn in halfway through its life cycle. Designers could get around this by providing more circuitry in the FBB driver to calculate the possible difference in worst case performance, but this would lead to an even larger amount of added circuitry. Designers could leave a substantial over performance margin in propagation delays, but this would defeat the purpose of an adaptive solution.

3.4 Stack Effect

A few more novel solutions for dealing with the aging problem have been suggested. One is to utilize signal scheduling to take advantage of the PMOS stack

effects [5]. Negative bias temperature instability, unlike HCI, occurs when there is a constant electric field even without a current. This means it is actually worse during static input. Fig. 3. [5] shows two static conditions for a simple 2-input nor.




Fig. 3. Static Conditions on 2-Input Nor. © [2006] IEEE

Note that they are logically equivalent only when the positions of signals A and B are switched, but they are not equivalent in the way they are affected by NBTI. In a PMOS stack the PMOS furthest from the power node (M_B) has a higher effective V_{th} due to body effect and is therefore less susceptible to NBTI. Because M_B has a higher effective V_{th} than M_A , a nor gate that spends most of its lifetime in ST1 wears out faster than a gate that spends most of its lifetime in ST2. If a circuit designer has adequate information about the behavior of a circuit during design, he can arrange signals to take advantage of stack effect to prevent aging due to NBTI.

However, stack effect is already used to improve performance of circuits. When stack effect is used to optimize performance, it is the gate's dynamic behavior and not its static behavior that a designer is interested in. Nonetheless, it is quite possible that a gate's behavior will result in one arrangement of signals being best for performance while the opposite arrangement is best to minimize aging. It is also possible that the designer will have insufficient information about the behavior of each signal.

3.5 Duty Cycle

Finally, if he knows enough about the behavior of a circuit, one can take advantage of the duty cycle to minimize the effect of NBTI [5]. The duty cycle of a circuit node is defined as the percentage of its lifetime that a state will remain high. NBTI is highest when a gate is stressed by a constant electric field. Therefore, the lower the duty cycle is on the gate of a transistor, the less degradation it will suffer from NBTI. In order to exploit this, not only will a lot of information about the behavior of the circuit need to be known at design time, but a serious burden will be put on the designer to reorganize the circuit. Ultimately, more serious study will be needed before a duty cycle can be utilized in this way.

3.6 Summary of Current Countermeasures

While there are many proposed techniques to reduce the effect of NBTI aging, most of them are in conflict with other features designers wish to optimize. Some of these techniques cannot be targeted to specific areas of interest in the circuit. Others are quite complex, and their overall effect on the power and performance of a circuit is not yet known.

Industry insiders have already begun to suggest that current design methods will not be adequate to overcome the aging effects predicted for processes below 32nm [1]. Future designers will have to seek out new ways to deal with the aging challenge. An ideal solution should have limited effect on the performance of a new circuit, have low overhead in terms of added area and engineering work, and be easily targeted to specific parts of the circuit without having unwanted effects on the behavior of the circuit as a whole.

4. FIELD TRANSISTOR SIZING

4.1 The Circuit

In order to avoid the pitfalls of overdesign and current adaptive techniques, we would like to design an approach to dealing with transistor aging that does not have a high upfront penalty and can be targeted to affect only the areas of the circuit where transistor aging is the most critical. We believe we can achieve both of these through the use of antifuses.

Antifuses are a mature technology primarily used in field programmable gate arrays. Modern antifuses are made with a layer of amorphous silicon and a dielectric sandwiched between two layers of metal [11], [12]. During normal operation an antifuse is highly resistant to current; ideally it acts as an open switch. When a series of high voltage pulses are applied across the two metal layers, the amorphous silicon aligns, and the antifuse forms a low impedance connection between its two metal layers. An antifuse can only be turned on once, and the antifuse is burned shut. It remains turned on for the lifetime of the circuit.

Traditionally, antifuses are used to make circuits that can be easily reprogrammed in the field. The field transistor sizing design uses this same technology to effectively resize PMOS transistors. The design takes a traditional CMOS gate and places additional PMOS transistors in parallel to any existing PMOS transistor in the traditional CMOS gate. [For the rest of this paper, we refer to those PMOS transistors bordered by antifuses as dynamic PMOS (DPMOS) and those not bordered by antifuses as static PMOS (SPMOS).]

From the time a circuit is manufactured until aging effects can be perceived by the user, all the antifuses are turned off. During this period of time, they act as very small capacitors. The numbers shown in Table I, based on technology from Actel [11], show the capacitance generated by an antifuse in the off state at various process sizes.

TABLE I
SPECIFICS OF ANTIFUSE CIRCUITS

| Process | Capacitance While Off | Resistance While On |
|------------------|-----------------------|---------------------|
| 0.6 μ m ONO: | 7.7 fF | 125 ohms |
| 0.6 μ m M/M: | 2.9 fF | 25 ohms |
| 0.35 μ m M/M | 1.6 fF | 25 ohms |
| 0.25 μ m M/M | 0.8 fF | 25 ohms |

During this stage of the FTS circuit's life, the antifuses at the bulk of each DPMOS transistor significantly lower the effective electric field at the gate and thus prevent significant power from being consumed. Antifuses at the source and gate of these transistors prevent current from going through the DPMOS transistor. This effectively prevents NBTI from taking place in the DPMOS transistor during the first stage of the circuit's life.

A power and performance penalty is paid for the additional parasitic capacitance of the antifuses, but the power penalty is small compared to overdesign, and the performance penalty is small compared to the penalty that would have been paid for aging. Effects of the added parasitics have been minimized by the location of the

antifuses. Because antifuses are placed on the bulk node but not the gate, they are not seen by transistors downstream and do not have a cumulative effect on power or delay through a path.

When the effects of aging can be felt by the customer, the antifuse is activated. While V_{th} cannot be directly measured delay can. So when the delay for a selected vector passes an arbitrary percent degradation point a user will know to turn on the antifuses.

In its on state the antifuse acts as a very small resistor. Again numbers provided by Actel in Table I show the amount of resistance provided by an antifuse in the on state at various process sizes [11].

At this stage in the circuit's life, each of the DPMOS-SPMOS transistor pairs effectively acts as a single larger transistor that has only undergone partial aging. This is the desired effect of overdesign without the additional power cost during the first part of the circuit's lifetime. The effect of such small amounts of resistance should be minimal; they were included in our simulation.

Metal to metal antifuses described above can fit on top of traditional metal vias so the proposed circuit changes should have minimal effect on circuit size [11]. The circuit only needs one additional net connecting antifuses and transistors along a given critical path. Dynamic transistors can be added or not added on a gate-by-gate basis. This makes the FTS technique much more fine grained than traditional adaptive techniques.

4.2 Simulation

HSPICE simulation was used to test the FTS technique. Four different circuits were simulated in a large variety of NMOS, SPMOS, and DPMOS size combinations. It was discovered that the technique varied widely depending on the circuit and gate sizes to which it was applied. All circuits were simulated using 65nm BPTM specifics with a 1.1V V_{dd} . BPTM numbers were also used in the simulation of wire delay (65nm local wires were used). Aging effects were simulated by increasing the V_{th} 33%. This increase was consistent with results in literature [2]. In their off states antifuses were simulated with a 0.8fF. In their on states these were substituted for 25 ohm resistors.

The first circuit was a simple 11 stage ring oscillator. This circuit was mainly used as a proof of concept before further investigation was warranted. Forty micrometer long wires were placed between each stage in the oscillator. Traditional PMOS widths from 1 to 3.5 were simulated. Dynamic PMOS transistors were simulated at 0.5 and 1.

The second circuit was a five stage fan-out four simulated circuit. This was used to gain further proof of concept before larger circuits were simulated. Along with the ring oscillator, it also provided us with the two extremes of critical path, one in which the critical path had no fan-out and one where the fan-out was rather high. This allowed us to gain insight into which circuits the FTS technique might perform best on without dealing with the multitude of factors that could cause differences in real world circuits. Like the ring oscillator, the fan-out circuit was simulated with SPMOS widths ranging from 1 to 3.5 and DPMOS widths of 0.5 and 1.

Finally, the technique was simulated on two combinational ISCAS circuits, c432 and c2670. The ISCAS c432 circuits had 244 gates including and-or-invert (AOI) and the or-and-invert (OAI) gates. The FTS technique was applied to a 16 gate critical path on the c432 circuits. A detailed gate list for both cells can be found in Table XV in the appendix. The ISCAS c2670 circuits had 1195 gates with a 22 gate critical path. The c2670 circuits did not have AOI or OAI gates but did have buffer's added, which the c432 did not. Eighty micrometer wires were added to the input and output of every gate using 65nm BPTM wire models.

Each of the ISCAS circuits was simulated in two separate flavors. In the basic FTS circuit, the FTS technique was applied to every gate in selected critical paths. The circuits were also simulated with partial FTS techniques. The same critical paths were chosen, but the technique was only applied to gates that would ordinarily have less than 4 PMOS transistors. The result was that the partial FTS circuits showed a smaller penalty in both the upfront power and performance as well as a smaller gain in the aged circuit. These two different variations of the FTS technique helped to demonstrate the technique's ability to be both fine grained and versatile. All ISCAS circuits were simulated with PMOS transistors, sizes ranging from 1 to 3.5, and dynamic transistor sizes 0.5 to 1.5.

For our simulations we assumed a user turned on the antifuse when V_{th} reached approximately 33%. The user would need to know the expected performance degradation at that point. This could be deduced from the standard circuit simulations. For instance for a s1d.5 circuit we would use the simulated degradation from a s1d0 cell.

Proposed turn on points for each simulated circuit are shown in table XVI in the appendix.

5. RESULTS AND ANALYSIS

5.1 Simple Ring Oscillator

The ring oscillator was used mainly as a proof of concept before more time intensive and demonstrative data was collected. However, the results of the ring oscillator proved to be echoed in later simulations. The ring oscillator was simulated with SPMOS widths at 2 - 3.5 and DPMOS widths 1 and 0.5. The results are shown below in Table II.

For this and all later charts, circuits were labeled according to the size of their static and dynamic PMOS transistors using the pattern s<static size>d<dynamic size>. For example, the ring oscillator with a 2 lambda SPMOS and a 1 lambda DPMOS was labeled s2d1. The control circuits were listed with a dynamic size of zero so a control circuit with size 3 PMOS gates was an s3d0 circuit.

There were eight numbers collected from every circuit, power (new and old), rise (new and old), fall (new and old), and period (new and old). Manipulating the widths of the transistors in the circuit affected all of these numbers, and they had to be analyzed together to see where the FTS technique was helpful and where it was not.


TABLE II
RING OSCILLATOR RESULTS

| | New | | | | Old | | | |
|-----------|------------------|----------------|----------------|------------|------------------|----------------|----------------|------------|
| | Power (μ w) | Rise Time (ps) | Fall Time (ps) | Delay (ps) | Power (μ w) | Rise Time (ps) | Fall Time (ps) | Delay (ps) |
| s.1d0 | 25.3 | 47.4 | 24.1 | 367 | 21.9 | 57.9 | 26.2 | 420 |
| s.1.5d0 | 33.3 | 33.9 | 22.4 | 311 | 29.0 | 40.8 | 23.9 | 353 |
| s.2d0 | 39.5 | 28.0 | 22.1 | 289 | 34.7 | 33.4 | 23.3 | 325 |
| s.2.5d0 | 44.7 | 25.1 | 22.6 | 280 | 39.5 | 29.3 | 23.5 | 313 |
| s.3d0 | 49.3 | 23.2 | 23.3 | 277 | 43.5 | 26.9 | 24.1 | 308 |
| s.3.5d0 | 53.1 | 22.1 | 24.4 | 278 | 47.0 | 25.2 | 24.9 | 307 |
| s.1d0.5 | 25.7 | 47.3 | 24.9 | 373 | 25.9 | 46.3 | 24.2 | 369 |
| s.1.5d0.5 | 33.5 | 34.0 | 22.9 | 318 | 32.2 | 35.8 | 23.1 | 329 |
| s.2d0.5 | 39.9 | 28.3 | 22.7 | 295 | 37.4 | 30.5 | 23.1 | 312 |
| s.2.5d0.5 | 45.1 | 25.3 | 22.9 | 286 | 41.7 | 27.7 | 23.4 | 304 |
| s.3d0.5 | 49.5 | 23.5 | 23.7 | 282 | 45.6 | 25.8 | 24.2 | 302 |
| s.3.5d0.5 | 53.2 | 22.4 | 24.8 | 283 | 48.9 | 24.4 | 25.1 | 304 |
| s.1d1 | 26.0 | 48.4 | 25.6 | 385 | 28.6 | 41.7 | 23.8 | 353 |
| s.1.5d1 | 33.7 | 34.9 | 23.5 | 327 | 34.4 | 33.8 | 23.2 | 324 |
| s.2d1 | 39.9 | 29.1 | 23.3 | 304 | 39.2 | 29.5 | 23.4 | 311 |
| s.2.5d1 | 45.2 | 25.8 | 23.7 | 294 | 43.2 | 27.1 | 23.9 | 306 |
| s.3d1 | 49.5 | 23.9 | 24.2 | 290 | 46.9 | 25.3 | 24.5 | 305 |
| s.3.5d1 | 53.4 | 22.8 | 25.1 | 290 | 50.0 | 24.3 | 25.7 | 307 |

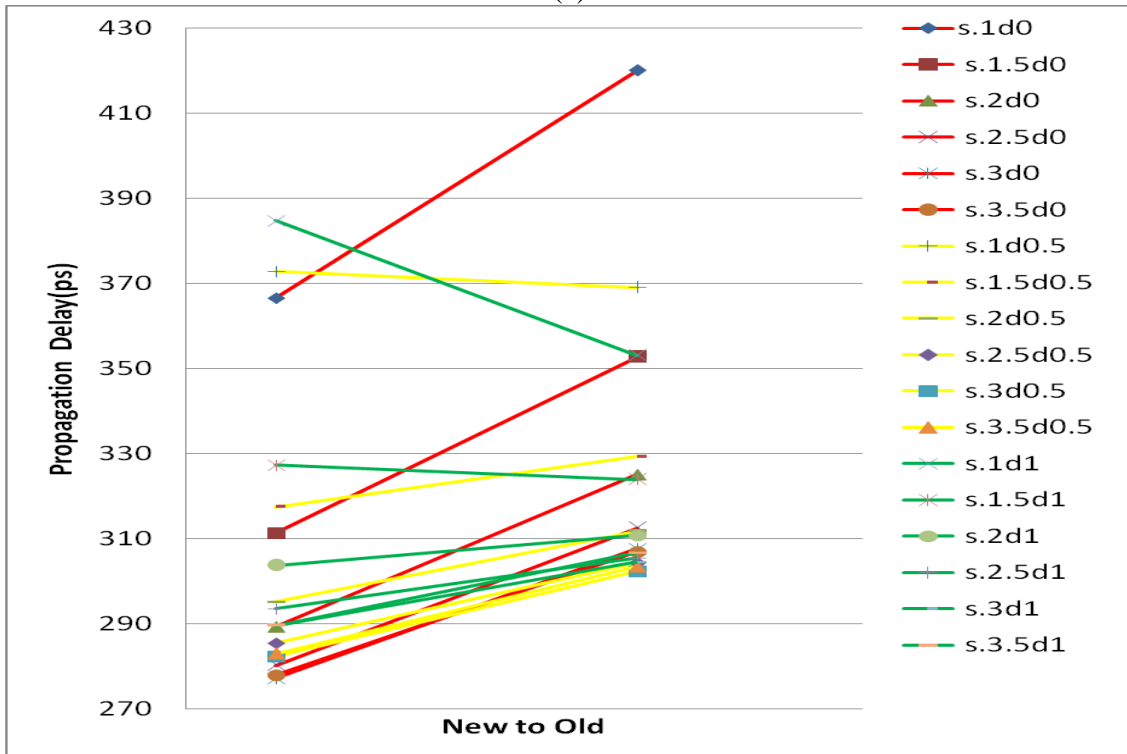
In all the simulations there was an upfront penalty in both power and performance for employing the FTS technique. This was to be expected; the antifuse circuits were not ideal, and they did introduce added parasitic capacitance into the circuit. What was slightly more surprising was that the penalty was higher when the DPMOS was wider. Even the small amount of current that was passing through the DPMOS in the off position was enough to add significant parasitics to the circuit.

Fig. 5 clearly shows how different gate sizes performed differently over their lifetime. Note that circuits with .5 lambda DPMOS transistors (yellow lines) came closest to having constant power and performance between their new and aged states. It was already clear from the oscillator that the key to making FTS pay off is properly sizing the DPMOS.

It was found that there were two key numbers that were most indicative of the FTS technique's overall effectiveness. The first was the sum of the DPMOS and SPMOS transistor sizes. It was not hard to see why this number was important. This number gave a quick and dirty estimation of the effective PMOS transistor size in an aged circuit. Of course, this did not take into account the change in V_{th} in the static PMOS, which lowered the effective combined strength of the aged transistors, but it provided a good starting place. It was found that a total PMOS length of between 3-3.5 was desirable in most simulated circuits.



(a)



(b)

Fig. 5. Behavior of New and Aged Ring Oscillator Circuits. (a)Power Consumption (b)Propagation Delay

The second key to transistor sizing was the ratio between the DPMOS and SPMOS widths. We saw a clear trend emerge when we graphed the DPMOS width / SPMOS width against the difference between the period in the old circuit and the period in the new circuit (Fig. 6).

As can be seen, there was an almost linear relationship between the SPMOS/DPMOS size ratio and the change in period due to aging. At the point where this line passes the Y intercept, the performance of the new circuit was expected to be almost equal to the performance of the old circuit. For the ring oscillator circuit, the magic ratio was about 0.53. One might suspect this would be an ideal ratio because it would not overuse power in either the old or the new state.

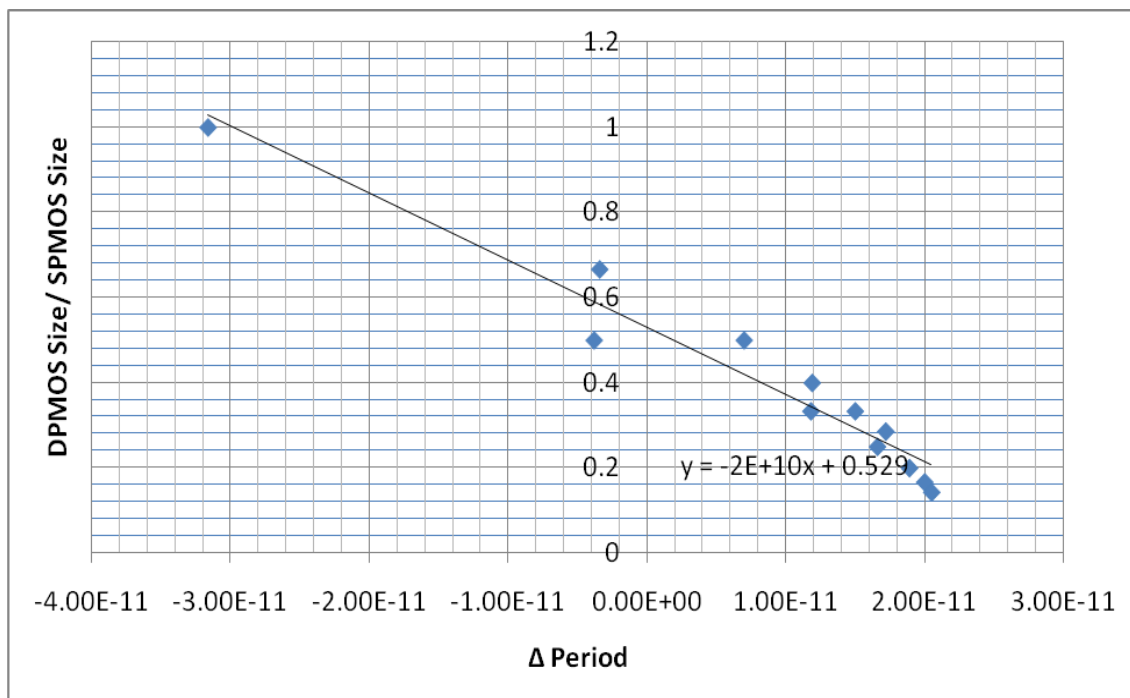


Fig. 6. D/SPMOS Ratio Versus Performance Ring Oscillator

In fact, our results showed that this number could be used as a guide to designers, but only when the sum of the two transistors was also accounted for. Using these numbers together, we could derive precise ideal SPMOS and DPMOS numbers. For instance, a sum of 3 and a ratio of 0.53 could be achieved at $s1.96d1.03$. A sum of 3.5 and a ratio of 0.53 yielded $s2.28d1.21$. We will justify the usefulness of these numbers later.

In order to establish the usefulness of the FTS technique, designers will have to consider many numbers in tandem. First, there is the problem of circuit behavior in its new state versus circuit behavior after it has been aged. The FTS technique does not present a benefit for people who are only interested in maximizing the usefulness of a circuit in its new condition. The FTS technique is only useful for designers who are interested in guaranteed behavior over the lifetime of their circuits.

For this reason, the real numbers of interest were worst case (over the lifetime of the circuit) numbers. For circuits with no DPMOS transistors, the worst case performance was that of the older circuit while the worst case power was that of a new circuit. In the circuits with DPMOS, the worst case performance/power could come at either stage, depending on the DPMOS/SPMOS ratio. For a general circuit we took the highest number (for power or performance) and used this as its worst case number (over the lifetime of the circuit). Table III shows the worst case numbers for all the circuits. We compared these numbers to determine which circuits performed best.

TABLE III
WORST CASE RING OSCILLATOR

| | Power (μw) | Delay (ps) |
|-----------|-------------------------|------------|
| s.1d0 | 25.3 | 420 |
| s.1.5d0 | 33.3 | 353 |
| s.2d0 | 39.5 | 325 |
| s.2.5d0 | 44.7 | 313 |
| s.3d0 | 49.3 | 308 |
| s.3.5d0 | 53.1 | 307 |
| s.1d0.5 | 25.9 | 373 |
| s.1.5d0.5 | 33.5 | 329 |
| s.2d0.5 | 39.9 | 312 |
| s.2.5d0.5 | 45.1 | 304 |
| s.3d0.5 | 49.5 | 302 |
| s.3.5d0.5 | 53.2 | 304 |
| s.1d1 | 26.0 | 385 |
| s.1.5d1 | 34.4 | 327 |
| s.2d1 | 39.9 | 311 |
| s.2.5d1 | 45.2 | 306 |
| s.3d1 | 49.5 | 305 |
| s.3.5d1 | 53.4 | 307 |

Secondly, for a fair comparison of the circuits, power and performance had to be considered together. First, we saw if power was improved. Ideally, period would be held constant while power was analyzed, but because we had static values, we could only compare the power of circuits with nearly the same period. Treating the numbers as individual points on a constant function would be convenient but unrealistic. VLSI circuits can only be manufactured at discrete values of lambda, so one can only realistically produce discrete performance levels. Instead, we compared each of the dynamic circuits to the static circuit that had the shortest period time in excess of the dynamic circuit. One can see the results of this comparison in Table IV:

TABLE IV
PERCENTAGE CHANGE RING OSCILLATOR

| Dynamic Circuit | Closest Power | Delta Period | Closest Period | Power Delta |
|-----------------|---------------|--------------|----------------|-------------|
| S1d0.5 | s1.5d0 | 5.70 | s1d0 | 2.21 |
| S1.5d0.5 | s2d0 | 1.29 | s1.5d0 | 0.75 |
| S2d0.5 | s2.5d0 | -0.26 | s2.5d0 | -10.77 |
| S2.5d0.5 | s3d0 | -1.07 | s3.5d0 | -15.21 |
| S3d0.5 | s3.5d0 | -1.50 | s3.5d0 | -6.85 |
| S3.5d0.5 | x | x | s3.5d0 | 0.19 |
| S1d1 | s1.5d0 | 9.07 | s1d0 | 12.71 |
| S1.5d1 | s2d0 | 0.68 | s1.5d0 | 3.24 |
| S2d1 | s2.5d0 | -0.58 | s2.5d0 | -10.75 |
| S2.5d1 | s3d0 | -0.71 | s3.5d0 | -14.96 |
| S3d1 | s3.5d0 | -0.78 | s3.5d0 | -6.76 |
| S3.5d1 | x | x | s3.5d0 | 0.51 |

As one can see, the power difference between FTS circuits and similarly performing standard circuits varied greatly from a decrease of 15% to as high as an increase of 12.71% , depending on the SPMOS and DPMOS sizes used. As already observed, proper sizing was key. Let's look at sizes close to numbers we proposed based on sum and D/S ratio. s2d1 was closest to s1.97d1.03 and had a 10.75% drop in power. s2.5d1 was closest to s2.3d1.2 and performed even better with a power drop of 14.96%.

Neither of these sizes was the most ideal; s2.5d0.5 had a power drop of 15.21% over similarly performing standard circuits. Our first circuit seemed to indicate that designers could yield significant improvements with the FTS technique through simple (back of the envelope) sizing formulas, or they could do even slightly better if they were willing to invest significant time in simulation.

If the goal of a designer is to increase performance while keeping power constant, the FTS technique is less impressive. Indeed, with a worst case lifetime period of $2.31E-11$, our best worst case performance did come from a FTS circuit (s2d0.5), but this was only a 0.26% improvement over similarly powered standard circuits. Our best performance improvement shift was the s3d0.5 circuit, which showed a 1.5% improvement in worst case performance and a 6.85% improvement in worst case power. Our proposed back of the envelope method would not have chosen this size because of its low D/S ratio.

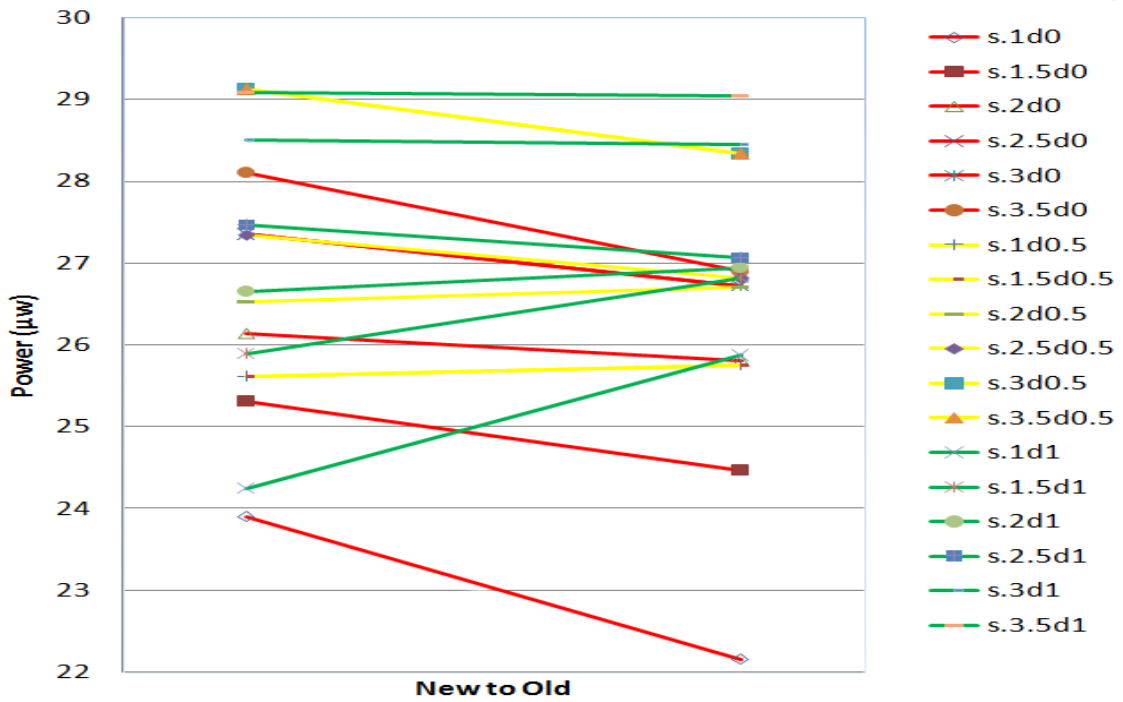
5.2 Simulated Fan-out Test

The second concept test was designed to simulate fan-out in more typical circuits. The circuit consisted of 5 inverters in a chain. Each inverter was four times larger than the inverter before it in order to simulate a load of four times as many inverters. Dynamic transistors were added to the third inverter, and measurements (power, rise time, and fall time) were made only on the third circuit. The results at various SPMOS (1, 1.5, 2, 2.5, 3, 3.5) and DPMOS (0.5, 1) were simulated. In general, the results of the fan-out four tests (shown in Table V) showed some of the same patterns as the oscillator ring but were much less positive.

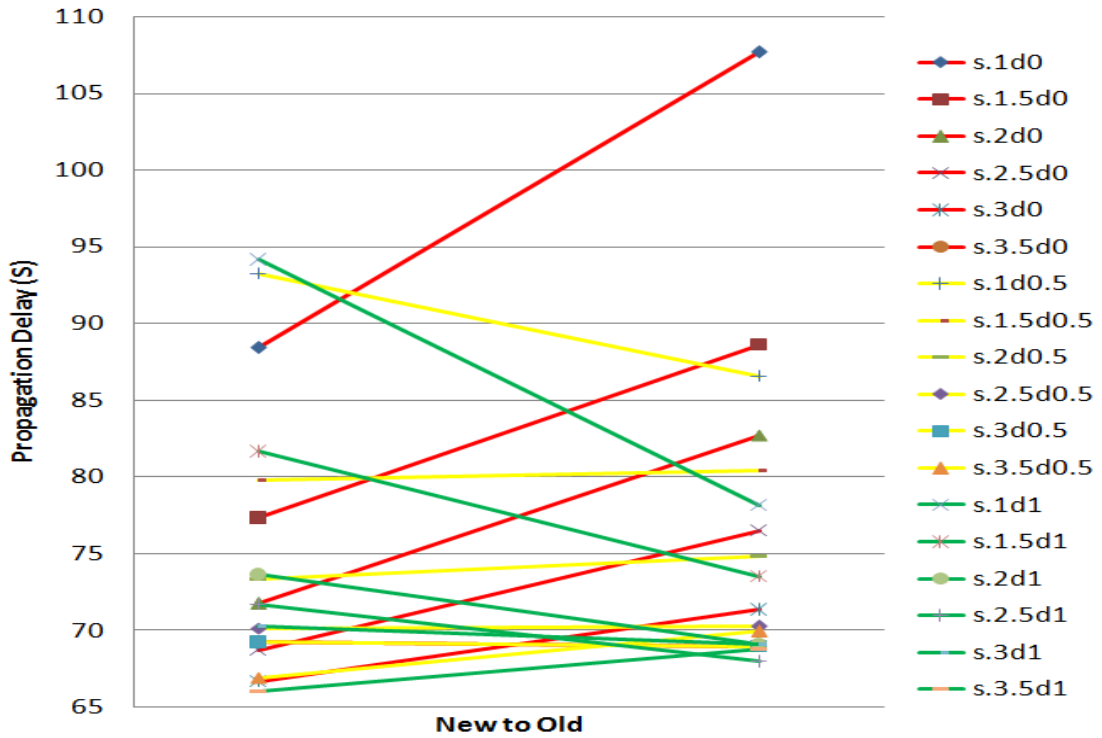
TABLE V
FAN-OUT FOUR RESULTS

| | New | | | | Old | | | |
|-----------|------------------|----------------|----------------|-------------|------------------|----------------|----------------|-------------|
| | Power (μ w) | Rise Time (ps) | Fall Time (ps) | Period (ps) | Power (μ w) | Rise Time (ps) | Fall Time (ps) | Period (ps) |
| s.1d0 | 23.9 | 64.9 | 23.5 | 88.4 | 22.2 | 80.6 | 27.0 | 108 |
| s.1.5d0 | 25.3 | 48.6 | 28.8 | 77.3 | 24.5 | 60.8 | 27.8 | 88.6 |
| s.2d0 | 26.1 | 41.7 | 30.0 | 71.7 | 25.8 | 52.6 | 30.2 | 82.7 |
| s.2.5d0 | 26.6 | 38.1 | 30.7 | 68.7 | 26.6 | 46.8 | 29.7 | 76.5 |
| s.3d0 | 27.4 | 35.4 | 31.3 | 66.7 | 26.7 | 39.5 | 31.9 | 71.3 |
| s.3.5d0 | 28.1 | 29.6 | 34.1 | 63.6 | 26.9 | 37.0 | 32.9 | 69.9 |
| s.1d0.5 | 24.3 | 65.6 | 27.7 | 93.3 | 24.7 | 58.7 | 27.9 | 86.6 |
| s.1.5d0.5 | 25.6 | 50.3 | 29.5 | 79.8 | 25.8 | 50.2 | 30.2 | 80.4 |
| s.2d0.5 | 26.5 | 43.0 | 30.4 | 73.3 | 26.7 | 45.1 | 29.8 | 74.9 |
| s.2.5d0.5 | 27.3 | 38.7 | 31.5 | 70.2 | 26.8 | 38.3 | 32.0 | 70.2 |
| s.3d0.5 | 28.4 | 35.9 | 33.4 | 69.2 | 26.9 | 36.0 | 32.9 | 69.0 |
| s.3.5d0.5 | 29.1 | 33.7 | 33.2 | 66.9 | 28.3 | 34.4 | 35.6 | 69.9 |
| s.1d1 | 24.2 | 65.7 | 28.5 | 94.2 | 25.9 | 47.8 | 30.4 | 78.1 |
| s.1.5d1 | 25.9 | 51.7 | 29.9 | 81.7 | 26.8 | 43.6 | 29.9 | 73.5 |
| s.2d1 | 26.7 | 43.7 | 29.9 | 73.6 | 26.9 | 37.0 | 32.1 | 69.1 |
| s.2.5d1 | 27.5 | 39.4 | 32.3 | 71.7 | 27.1 | 35.0 | 33.0 | 68.0 |
| s.3d1 | 28.5 | 36.4 | 33.9 | 70.3 | 28.4 | 33.4 | 35.7 | 69.1 |
| s.3.5d1 | 29.1 | 34.1 | 31.9 | 66.0 | 29.0 | 32.3 | 36.5 | 68.8 |

Another view of the fan-out four results can be seen in Fig. 7. As in the ring oscillator circuits, the power and performance over the lifetime of the fan-out four circuits was closer to constant when FTS gates (yellow and green lines) were used. However the range of power consumption was much more narrow in the fan-out four circuits than it was in the ring oscillator circuits. This is probably part of the reason the FTS technique was not as effective when applied to the fan-out four circuits.



(a)



(b)

Fig. 7. Behavior of New and Aged Fan-out Four Circuits. (a)Power Consumption (b)Propagation Delay

Once again there was a generally linear relationship between the ratio of dynamic PMOS and static PMOS transistors and the change in rise time over the life of the circuit. One can see this in Fig. 8. The ratio had no discernible relationship to the change in fall time over the lifetime of the circuit, but the change in overall performance (rise time + fall time) was dominated by rise time. Thus, overall performance had a linear relationship with the DPMOS/SPMOS size ratio.

The fan-out circuit had its Y intercept a little lower, at approximately 0.28. From this we could predict best results from sizes of approximately s2.34d.66 (sum 3) or s2.73d.77 (sum 3.5).

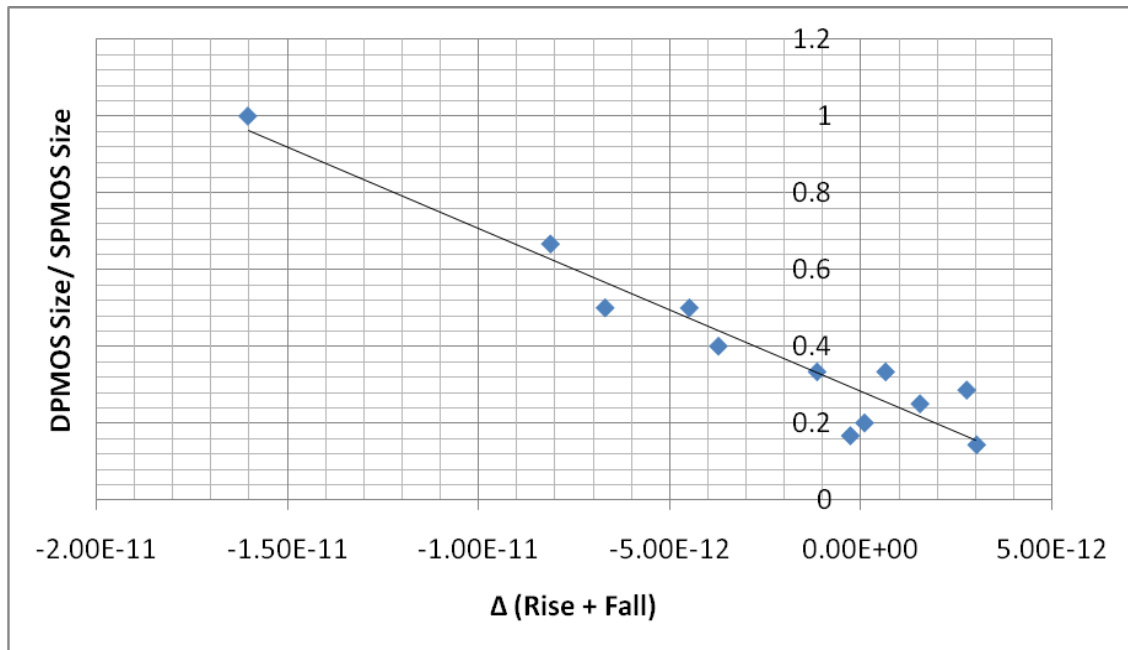


Fig. 8. D/SPMOS Ratio Versus Performance Fan-out Four

Again, in order to evaluate the effectiveness of the FTS technique, we had to do our best to hold one aspect of the circuit's behavior constant while we evaluated the other. First, we compared the worst case power of each FTS circuit with the standard circuit that had the shortest delay in excess of that circuit (Table VI). We could tell immediately that the FTS technique was not as effective in the simulated fan-out circuit as it was with the simple oscillator ring.

TABLE VI
CHANGE IN POWER FOR FAN-OUT FOUR CIRCUITS

| Dynamic Circuit | Closest Period | Power Delta (%) |
|-----------------|----------------|-----------------|
| s1d0.5 | s1d0 | 3.43 |
| s1.5d0.5 | s2d0 | -1.49 |
| s2d0.5 | S2.5d0 | 0.30 |
| s2.5d0.5 | s3d0 | -0.04 |
| s3d0.5 | S3.5d0 | 0.89 |
| s3.5d0.5 | S3.5d0 | 3.63 |
| s1d1 | s1d0 | 8.24 |
| s1.5d1 | s2d0 | 2.60 |
| s2d1 | S2.5d0 | 1.20 |
| s2.5d1 | S2.5d0 | 3.16 |
| s3d1 | s3d0 | 4.20 |
| s3.5d1 | S3.5d0 | 3.49 |

The best results came from the s1.5d0.5 circuit, but even this only showed a decrease in lifetime power of 1.49% . Our proposed circuits, rounded to manufacturable numbers, would be s2.5d0.5 and s2.5d1. The first of these would only provide a drop in power consumption of 0.04%, and the second would actually increase power consumption by 3.16%.

Most circuits resulted in a negative change in worst case power. One circuit had an 8.24% increase in power compared to similarly performing standard circuits. Clearly, the FTS technique could be quite costly if applied incorrectly.

The FTS technique was slightly more effective if we attempted to minimize delay. Table VII compares the performance of dynamic-sized circuits to static-sized circuits with the lowest power in excess of their dynamic counterparts. Performance is measured by the sum of both rise and fall times.

When we attempted to minimize delay, we were able to achieve an improvement of up to 2.76% in the sum of rise and fall times when compared to similarly powered circuits, but our proposed circuits would only provide an decrease in delay of 1.54% (s2.5d0.5) or worse yet, an increase in delay of 2.52% (s2.5d1). Again in most circuits the FTS technique actually made transition times worse for their power points. In fact, circuits could show increases in worst case rise time as high as 14.45%.

It is unclear why the FTS technique seems to perform better in long serial critical paths than shorter paths with high fan-out, but with this information one would expect the technique to work best in long repeated paths, which is a realistic situation in many real world circuits. Our results in ISCAS circuits reflected the results in the ring oscillator more closely than they reflected results in the fan-out four circuit.

TABLE VII
CHANGE IN PERIOD FAN-OUT FOUR CIRCUITS

| Dynamic Circuit | Closest Power | Delta Period |
|-----------------|---------------|--------------|
| s1d0.5 | s1.5d0 | 5.24 |
| s1.5d0.5 | s2d0 | -2.76 |
| s2d0.5 | s3d0 | 4.92 |
| s2.5d0.5 | s3d0 | -1.54 |
| s3d0.5 | s3.5d0 | -1.00 |
| s3.5d0.5 | x | x |
| s1d1 | S2d0 | 13.85 |
| s1.5d1 | S3d0 | 14.45 |
| s2d1 | S3d0 | 3.20 |
| s2.5d1 | S3.5d0 | 2.52 |
| s3d1 | x | x |
| s3.5d1 | x | x |

5.3 ISCAS Circuit Performance

In order to get a more definitive understanding of how useful field transistor sizing would be in real circuits, we tested it on two ISCAS combinational circuits, c432 and c2670. Dynamic transistors were applied to select critical paths of each, and the power consumed and propagation delay along each critical path were measured.


Every combination of transistor sizes was simulated twice. In one version the dynamic cells were used in every gate on the critical path. In the other version only gates with less than 4 PMOS transistors were replaced with FTS gates. The full amount of data collected from the ISCAS is too large to post here, so it can be found in the appendix.

The basic linear relationship between D/S size ratio and performance change continued to present itself in both ISCAS circuits as can be seen in Fig. 9. The trend


lines for the circuits which had only part of their critical paths modified kept very close to the trend lines of those that had their entire critical paths modified.

Of course, the most important thing to observe from these graphs was the point when performance difference approached zero. For both variations of the c2670 circuit, the difference in performance over their lifetimes tended toward zero as the SPMOS vs. DPMOS ratio tended toward 0.76. For the c432 circuits the difference in aged and new performance approached zero as the D/S ratio approached 0.39.


Without significantly more testing it is unclear what factors affected the zero points. The zero points ranged from 0.28 to 0.76, a difference of 170% . The fact that the D/S ratio had such a large variance clearly has negative implications on the usefulness of the FTS technique. Without more experimentation to establish what factors affect the D/S zero point, this cannot be estimated in advance of simulation. However, the linear nature of the D/S to performance difference graph does allow engineers to estimate the D/S zero point after simulation of only two gate sizes.



(a)



(b)



(c)

Fig. 9. D/SPMOS Ratio Versus Performance in ISCAS Circuits. (a) ISCAS c432 Partially Modified (b) ISCAS c432 Fully Modified (c) ISCAS c2670

Earlier findings have suggested that the sum of SPMOS and DPMOS transistors should be either 3.0 or 3.5. Using these numbers, we calculated several proposed FTS gate sizes [Table VIII]. These will be referred to later.

TABLE VIII
PREDICTED GATE SIZES FOR ISCAS CIRCUITS

| Circuit | DPMOS Size/ SPMOS Size | Sum 3.0 SPMOS | Sum 3.0 DPMOS | Sum 3.5 SPMOS | Sum 3.5 DPMOS |
|---------|---------------------------|---------------|---------------|---------------|---------------|
| c432 | 0.39 | 2.1582734 | 0.841727 | 2.517986 | 0.982014 |
| c2670 | 0.76 | 1.7045455 | 1.295455 | 1.988636 | 1.511364 |

Again in order to evaluate FTS usefulness in the ISCAS circuits, we attempted to analyze power while keeping delay as close as possible to constant. We compared each FTS circuit to the standard circuit with the smallest propagation delay in excess of the FTS circuit.

If there was a standard gate size (no DPMOS) that had worse performance and consumed more power than a smaller gate the smaller gate was not used it was never used for comparison because the smaller circuit would be used in any real world application. In the first set of circuits, we applied FTS modifications to every gate along the critical path (See the results in Table IX).

From the results we continued to see that proper gate sizing was the key. For poorly chosen FTS gate sizes, power increases of up to 9.53% (s1d0.5) could be seen. At this gate size poor performance was to be expected. The 1 lambda static finger was already too small in the new state to perform well. The addition of the dynamic finger added parasitic to an already underpowered new transistor. Since gate sizing is under the designer's control, as long as proper sizes can be found, this does not diminish the usefulness of the technique.

In the c2670 we were able to get a 13% improvement in (worst case) power without diminishing performance over the lifetime of the circuit. The best results came from the s1.5d1 circuit. Our proposed circuits also did rather well. Using the 3.0 target, we calculated the closest manufacturable circuit would be s1.5d1.5. The s1.5d1.5 showed a 12.1% improvement in power over similarly performing circuits. If we used 3.5 as a target, we would propose a s2d1.5 circuit. This showed a smaller but respectable power improvement of 7.89%.

TABLE IX
CHANGE IN POWER FOR ISCAS CIRCUITS

| FTS Circuit | C2670 | | C432 | |
|-------------|---------------|--------------|---------------|-------------|
| | Closest Delay | Power Delta | Closest Delay | Power Delta |
| S1d0.5 | S1.5d0 | 0.549176236 | s2d0 | 9.53809238 |
| S1.5d0.5 | S2d0 | -0.664767331 | s2d0 | 5.63887223 |
| S2d0.5 | S2.5d0 | -0.989653621 | s2.5d0 | 0.05803830 |
| S2.5d0.5 | S3.5d0 | -5.723630417 | s3.5d0 | -5.75079872 |
| S3d0.5 | S3.5d0 | -0.817661488 | s3.5d0 | -1.43769968 |
| S3.5d0.5 | S3.5d0 | 3.883892069 | s3.5d0 | 3.08839191 |
| S1d1 | S1.5d0 | 2.046929606 | s2d0 | 8.45830834 |
| S1.5d1 | S3.5d0 | -13.00081766 | s2d0 | 5.87882424 |
| S2d1 | S3.5d0 | -8.99427637 | s3.5d0 | -8.04046858 |
| S2.5d1 | S3.5d0 | -4.374488962 | s3.5d0 | -5.53780618 |
| S3d1 | S3.5d0 | 0.204415372 | s3.5d0 | -0.90521832 |
| S3.5d1 | S3.5d0 | 4.742436631 | s3.5d0 | 3.19488818 |
| S1d1.5 | S1d0 | 8.033385498 | s2d0 | 8.75824835 |
| S1.5d1.5 | S3.5d0 | -12.10139002 | s2d0 | 5.69886023 |
| S2d1.5 | S3.5d0 | -7.890433361 | s2.5d0 | 0.23215322 |
| S2.5d1.5 | S3.5d0 | -3.352412101 | s3.5d0 | -5.2715655 |
| S3d1.5 | S3.5d0 | 1.267375307 | s3.5d0 | -0.90521832 |
| S3.5d1.5 | S3.5d0 | 5.600981194 | s3.5d0 | 3.30138445 |

The same basic trend was followed in the c432 circuit. Our proposed sum 3.5 circuit (s2.5d1) showed a 5.53% improvement in power. Our proposed sum 3.0 circuit (s2d1) showed an even larger improvement of about 8.04%. The s2d1 circuit showed the largest power improvement of all the c432 circuits we tested. Based on the circuits we studied, a designer could expect the most reliably positive results from a sum 3 circuit closest to the D/S performance difference zero.

Next we evaluated FTS usefulness from the performance perspective. Again we used our tried and true method of attempting to keep power as close as possible to constant. We compared each FTS circuit to the standard circuit with the smallest power consumption in excess of the FTS circuit. If a standard gate size was larger than another standard gate size that achieved the same delay, it was never used for comparison because the smaller gate would be used in any real world application. If a FTS circuit consumed more power than all the standard circuits tested, its results were ignored for this comparison. The results for the fully modified circuits are shown in Table X:

TABLE X
PERFORMANCE DIFFERENCE FOR ISCAS CIRCUITS


| | C2670 | | c432 | |
|----------|---------------|--------------|---------------|-------------|
| | Closest Power | Delay Delta | Closest Power | Delay Delta |
| s1d0.5 | s2d0 | 4.249126892 | s3.5d0 | 29.34899 |
| s1.5d0.5 | s2d0 | -0.960419092 | S3d0 | 8.412698 |
| s2d0.5 | s2.5d0 | -0.658682635 | s2.5d0 | -0.258 |
| s2.5d0.5 | s3d0 | -0.697603882 | S3d0 | -1.0582 |
| s3d0.5 | s3.5d0 | -0.761962816 | s3.5d0 | -1.86766 |
| s3.5d0.5 | x | x | x | x |
| s1d1 | s2d0 | 3.667054715 | s3.5d0 | 30.84312 |
| s1.5d1 | s2.5d0 | -3.053892216 | S3d0 | 8.941799 |
| s2d1 | s3d0 | -2.820746133 | S3d0 | -1.0582 |
| s2.5d1 | s3.5d0 | -2.590673575 | S3d0 | -4.17989 |
| s3d1 | s3.5d0 | -2.712587626 | s3.5d0 | -4.16222 |
| s3.5d1 | x | x | x | x |
| s1d1.5 | s2d0 | 4.976717113 | s3.5d0 | 31.48346 |
| s1.5d1.5 | s2.5d0 | -3.293413174 | S3d0 | 9.047619 |
| s2d1.5 | s3d0 | -4.91355778 | S3d0 | 0.05291 |
| s2.5d1.5 | s3.5d0 | -4.327948796 | S3d0 | -3.80952 |
| s3d1.5 | x | x | s3.5d0 | -5.38954 |
| s3.5d1.5 | x | x | x | x |

The FTS technique was not as effective from the standpoint of increasing performance without affecting power. We did see that at least a 4.9% improvement in performance could be achieved by either ISCAS circuit if the sizing was done properly.


Our proposed FTS sizing did show improvement. At the target sum of 3.0, we saw a 3.29% improvement in worst case performance for the c2670 (s1.5d1.5) and a 1.05% improvement in the c432 (s2d1). The results were better with a target sum of 3.5. At the target sum of 3.5, we saw a 4.91% improvement in the c2670 (s2d1.5) and a 4.17% improvement in the c432 (s2.5d1). This was the best result for the c2670 but not the best result for the c432. In the c432 the s3d1.5 circuit showed a 5.38% improvement in worst case performance over similarly powered standard circuits. The results in both power and performance variation can be seen more easily in Fig. 10.

It is somewhat interesting to note that while targeting a sum of 3.0 works best when the designer is trying to optimize for power and SPMOS, a DPMOS sum of 3.5 seems to be the most reliable for obtaining high performance. This is somewhat intuitive because larger target sums yield larger gates that consume more power but have higher performance.

We should also note that some FTS circuits proved to do very poorly in the increased performance measure. Three of the c432 circuits had increased propagation delays of over 29%, but all three of these circuits had static PMOS sizes of lambda equal 1, grossly underpowered for the selected critical path. Furthermore, they had D/S ratios of at least 0.5. The c432's D/S zero point was as low as 0.39, so these FTS sizes would not have been considered by a knowledgeable designer.



(a)



(b)

Fig. 10. Difference (%) in Power and Delay for Fully FTS Modified Circuits Compared to Similar Standard Circuits. (a) ISCAS c432 (b) ISCAS c2670

The second variation of the ISCAS circuits attempted to minimize the upfront penalties in power and delay caused by the antifuse parasitics. In this variation dynamic transistors were only added to the smaller gates, those with fewer than 4 PMOS transistors in the standard circuit. Since the number of DPMOS transistors in a FTS circuit is equivalent to the number of PMOS transistors, this would keep the extra parasitic capacitance from growing very large.

In the c2670 circuit there were very few gates with more than 3 PMOS transistors in the standard gate. This was probably because the circuit was not optimized in the same way. In the c432 circuit multiple gate levels were compressed into single combined AOI and IOA type gates. In the c2670 circuit these gates were not utilized. The result was that only a single gate was different on the full FTS and partial FTS circuit critical paths, and no discernible difference was found in the power or performance of the c2670 circuit.

In the c432 circuit, however, significant differences in power and performance were seen when the FTS technique was only applied to the larger gates (at least 4 PMOS transistors). The power and performance comparison between each partial FTS c432 circuit and similar standard circuits can be seen in Table XI:

Once again our first calculation was the power consumed by each FTS circuit compared to the power consumed by the circuit that had the least propagation delay in excess of that circuit. The results in the partial FTS circuit were similar but slightly improved over the fully modified FTS circuit.

TABLE XI
POWER AND PERFORMANCE CHANGE IN PARTIAL FTS ISCAS C432 CIRCUIT

| FTS Circuit | Closest Power | Delay Delta | Closest Delay | Power Delta |
|-------------|---------------|-------------|---------------|-------------|
| s1d0.5 | s3d0 | 28.78 | S2d0 | 7.56 |
| s1.5d0.5 | s3d0 | 8.20 | S2d0 | 4.56 |
| s2d0.5 | s2.5d0 | -0.52 | S2.5d0 | -1.04 |
| s2.5d0.5 | s3d0 | -1.27 | S3.5d0 | -6.39 |
| s3d0.5 | s3.5d0 | -2.08 | S3.5d0 | -2.02 |
| s3.5d0.5 | x | x | S3.5d0 | 2.13 |
| s1d1 | s3d0 | 28.47 | S2d0 | 7.68 |
| s1.5d1 | s3d0 | 8.73 | S2d0 | 5.28 |
| s2d1 | s2.5d0 | -3.82 | S3.5d0 | -8.89 |
| s2.5d1 | s3d0 | -4.18 | S3.5d0 | -6.28 |
| s3d1 | s3.5d0 | -4.43 | S3.5d0 | -1.76 |
| FTS Circuit | Closest Power | Delay Delta | Closest Delay | Power Delta |
| s3.5d1 | x | x | S3.5d0 | 2.40 |
| s1d1.5 | s3d0 | 30.58 | S2d0 | 7.98 |
| s1.5d1.5 | s3d0 | 8.89 | S2d0 | 4.56 |
| s2d1.5 | s2.5d0 | -2.63 | S3d0 | -4.99 |
| s2.5d1.5 | s3d0 | -3.97 | S3.5d0 | -6.07 |
| s3d1.5 | s3.5d0 | -5.55 | S3.5d0 | -1.38 |
| s3.5d1.5 | x | x | S3.5d0 | 2.88 |

Our proposed sum 3.5 (s2.5d1) circuit had a 6.28% decrease in worst case power consumption compared to similarly performing standard circuits. This was slightly higher than the 5.53% decrease we saw in the fully modified FTS circuit. Our proposed sum 3.0 circuit (s2d1) showed an 8.89% decrease in worst case power consumption over similarly performing standard circuits. This was even higher than the 8.04% decrease we achieved with a fully modified FTS circuit. In both the full and partial c432 FTS circuits, the best results were achieved by the s2d1 (proposed sum 3) circuit.

Our second comparison was the worst case delay with constant power. The second part of Table XI compares the performance of each partial FTS circuit with the standard circuit that consumed the least amount of power in excess of that circuit. Those circuits that consumed more power than the largest standard circuit simulated were removed from the comparison.

Our proposed sum 3.5 circuit (s2.5d1) changed less than 0.01% when the larger gates were left unchanged. The partial FTS circuit continued to show a 4.18% improvement over a similarly performing standard circuit, but both the proposed 3.0 as well as the absolute best performing (over the lifetime of the circuit) circuits improved slightly. Our proposed sum 3.0 circuit (s2d1) showed a 3.82% increase in lifetime performance over a similarly powered standard circuit when only the smaller gates were changed. The same circuit with all the gates changed showed only a 1.05% improvement in performance. The results for the partially modified c432 circuit are summarized in Fig. 11.




Fig. 11. Change in Power and Delay Compared to Similar Standard Circuits for the ISCAS c432 Partially Modified Circuits

In both the full and partial FTS circuits, the best results for performance optimization were achieved by the s3d1.5 circuit. In the fully augmented FTS circuits the s3d1.5 circuit showed roughly 5.3% improvement in lifetime performance compared to the standard circuits with similar lifetime power consumption. When only gates with less than 4 transistors were changed, the results were slightly better. The partial FTS circuit achieved a 5.5% improvement in lifetime performance compared to the standard circuits with similar lifetime power consumption.

6. CONCLUSIONS

Transistor aging is a major problem that is only going to get worse with subsequent technology generations. In current technologies there is between a 10-15% degradation after ten years. If Intel's predictions are correct, aging in future processes will be much more significant after only 3-5 years. In these technologies overdesign will not be sufficient to makeup the power differences even if users are willing to pay the performance penalty [1].

Field transistor sizing presents an attractive solution to transistor aging for several reasons. First, while there is a small upfront cost in power and delay, it is much more efficient in power than overdesign. The FTS technique can be applied in a targeted manner to individual critical paths or even to specific gates within those critical paths with minimal routing overhead. The FTS technique is built to utilize inexpensive technology already in widespread use in FPGAs.

The only Area overhead would be that of the antifuses themselves. In the Actel [11] processes in which antifuses are used they are implemented between metal layers and have a small footprint $.40 \mu\text{m}$ in a $.15 \mu\text{m}$ process. Assuming the technology is scaled down to match its process total area overhead should be minimal.

Our results suggest that in real world circuits the FTS technique can produce equivalent lifetime performance as standard circuits with close to a ten percent reduction in power. The FTS technique can produce lifetime propagation delays that are 4-6% lower than those achievable from standard circuits of the same size. More impressively,

lifetime power consumption over a critical path can be reduced up to 13% in some circuits without affecting lifetime performance.

However, in order to benefit from its advantages, the FTS technique must be applied carefully. If the sizes of the DPMOS and SPMOS transistors are not chosen correctly, power and performance gains may be overtaken by the negative effects associated with the added parasitic capacitance in the antifuse circuits. Our tests have shown increases in power as high as 12.71% and increases in worst case delay as high as 31.48% over the lifetime of the circuit.

While we do not yet understand all the circuit factors that determine which transistor sizes work best, we have found several tools that will help the designer correctly apply the FTS technique. There is a clearly linear relationship between change in delay over the lifetime of the circuit and the DPMOS/SPMOS size ratio. The zero point for this line is a strong indicator for the sizes in which the FTS technique is most beneficial. This point, along with the general slope of this line, can be hard to predict from one circuit to the next. More research will be needed to determine what factors determine this critical number, and how, but the strong linear trend does allow designers to estimate these numbers by simulating relatively few gate sizes.

Furthermore, success of the FTS technique is highly dependent on the sum of the SPMOS and DPMOS transistor sizes. This in effect approximates the relative strength of the FTS circuit when the antifuse is in its on state. Unlike DPMOS/SPMOS ratio the best DPMOS, SPMOS sum seems to remain somewhat consistent for a given process. For the

Berkely predictive 65nm process we simulated, a sum of 3.0 was best when optimizing power while a sum of 3.5 was best when optimizing delay.

When both transistor size ratio and transistor size sums are taken into account, designers can reliably improve both lifetime power consumption and lifetime delay using FTS gates as compared to standard gates.

In some cases it appears FTS results can be further improved by selective application. In circuits such as the c432, where standard gates are combined into larger gates, with more than three traditional PMOS transistors, it may be advantageous not to apply the technique to these larger gates. At least for the c432, this technique did improve our results slightly. By modifying only the smaller gates, we were able to make power and performance scaling more gradual and thus achieve almost a 9%, as opposed to an 8%, improvement in power.

However, this result was not echoed in the c2670 circuit. This was not surprising since the c2670 circuit did not utilize complex gates. This is probably the more common real world case. In circuits where larger gates are utilized, designers may wish to investigate more selective application, but the relative difference this would achieve is expected to be small at best.

In summary, FTS is promising, but more research is advisable. FTS is promising because it provides substantial improvements in the aging of 65nm circuits, and current research shows aging will be critical in technologies to come. More research is suggested because we have not yet identified all the factors that dictate how FTS must be applied to achieve its best results.

REFERENCES

- [1] S. Borkar, "Designing reliable systems from unreliable components: the challenges of transistor variability and degradation," *IEEE Micro*, vol. 25, no. 6, pp. 10-16, Nov. 2005.
- [2] B. C. Paul, K. Kang, H. Kufluoglu, M. A. Alam, and K. Roy, "Impact of NBTI on the temporal performance degradation of digital circuits," *IEEE Electron Device Letters*, vol. 26, no. 8, pp. 560-562, Aug. 2005.
- [3] -----, "Negative bias temperature instability: estimation and design for improved reliability of nanoscale circuits," *IEEE Trans. Computer-Aided Design Integrated Circuits And Systems*, vol. 26, no. 4, pp. 743-751, Apr. 2007.
- [4] K. Kang, H. Kufluoglu, M. A. Alam, and K. Roy, "Efficient transistor-level sizing technique under temporal performance degradation due to NBTI," presented at the 2006 International Conference of Computer Design, San Jose, CA.
- [5] R. Vattikonda, W. Wang, and Y. Cao, "Modeling and minimization of PMOS NBTI effect for robust nanometer design," presented at the 2006 Design Automation Conference, San Francisco, CA, July 24-28, 2006.
- [6] J. W. Tschanz, J. T. Kao, S. G. Narendra, R. Nair, D. A. Antoniadis, A. P. Chandrakasan, and V. De, "Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage," *IEEE Journal Solid-State Circuits*, vol. 37, no. 11, pp. 1396-1402, Nov. 2002.
- [7] A. Keshavarzi, S. Ma, S. Narendra, B. Bloechel, K. Mistry, T. Ghani, S. Borkar, and V. De, "Effectiveness of reverse body bias for leakage control in scaled dual V_t CMOS ICs," presented at the 2001 International Symposium on Low Power Electronics and Design, Huntington Beach, CA.
- [8] S. Mahapatra, P. B. Kumar, and M. A. Alam, "Investigation and modeling of interface and bulk trap generation during negative bias temperature instability of p-MOSFETs," *IEEE Trans. on Electron Devices*, vol. 51, no. 9, pp. 1371-1379, Sept. 2004.
- [9] J. P. Campbell, P. M. Lenahan, A. T. Krishnan, and S. Krishnan, "Observations of NBTI-induced atomic-scale defects," *IEEE Trans. Device Materials*, vol. 6, no. 2, pp. 117-122, June 2006.
- [10] A. Shickova, B. Kaczer, P. Verheyen, G. Eneman, E. San Andres, M. Jurczak, P. Absil, H. Maes, and G. Groeseneken, "Negligible effect of process-induced strain

- on intrinsic NBTI behavior,” *IEEE Electron Device Letters*, vol. 28, no. 3, pp. 242-244, Mar. 2007.
- [11] “Quality and Reliability Guide,” *Actel*, Mountain View, CA, Mar. 2001.
- [12] J. Greene, E. Hamdy, and S. Beal, “Antifuse field programmable gate arrays,” *Proceedings of the IEEE*, vol. 81, no. 7, pp. 1042-1056, July 1993.

APPENDIX

TABLE XII
ISCAS C432 FULLY MODIFIED RESULTS

| Fully Modified ISCAS c432 Circuit | | | | |
|-----------------------------------|------------------------|------------------|------------------------|------------------|
| | New | | Aged | |
| | Propagation Delay (ps) | Power (μ w) | Propagation Delay (ps) | Power (μ w) |
| s.1d0 | 239 | 17.4 | 270 | 17.7 |
| s.1.5d0 | 201 | 17.0 | 223 | 16.8 |
| s.2d0 | 185 | 16.7 | 204 | 16.2 |
| s.2.5d0 | 176 | 17.2 | 194 | 17.0 |
| s.3d0 | 173 | 18.0 | 189 | 17.7 |
| s.3.5d0 | 171 | 18.8 | 187 | 18.5 |
| s.1d0.5 | 242 | 18.3 | 228 | 16.4 |
| s.1.5d0.5 | 205 | 17.6 | 204 | 15.9 |
| s.2d0.5 | 187 | 17.2 | 193 | 16.5 |
| s.2.5d0.5 | 181 | 17.7 | 187 | 17.4 |
| s.3d0.5 | 176 | 18.5 | 184 | 18.1E |
| s.3.5d0.5 | 175 | 19.4 | 184 | 18.8 |
| s.1d1 | 245 | 18.1 | 207 | 15.0 |
| s.1.5d1 | 206 | 17.7 | 191 | 15.7 |
| s.2d1 | 187 | 17.3 | 183 | 16.6 |
| s.2.5d1 | 181 | 17.7 | 180 | 17.3 |
| s.3d1 | 177 | 18.6 | 180 | 18.1 |
| s.3.5d1 | 176 | 19.4 | 179 | 18.7 |
| s.1d1.5 | 246 | 18.1 | 196 | 14.9 |
| s.1.5d1.5 | 206 | 17.6 | 182 | 15.7 |
| s.2d1.5 | 189 | 17.3 | 177 | 16.5 |
| s.2.5d1.5 | 182 | 17.8 | 176 | 16.8 |
| s.3d1.5 | 177 | 18.6 | 177 | 17.2 |
| s.3.5d1.5 | 176 | 19.4 | 177 | 17.6 |

TABLE XIII
ISCAS C432 PARTIALLY MODIFIED RESULTS

| Partially Modified c432 ISCAS Circuit | | | | |
|---------------------------------------|------------------------|-----------------|-----------------------|-----------------|
| | New | | Old | |
| | Propagation Delay (ps) | Power(μ w) | Propagation Delay(ps) | Power(μ w) |
| s.1d0 | 239 | 17.4 | 270 | 17.7 |
| s.1.5d0 | 201 | 17.0 | 223 | 16.8 |
| s.2d0 | 185 | 16.7 | 204 | 16.2 |
| s.2.5d0 | 176 | 17.2 | 194 | 17.0 |
| s.3d0 | 173 | 18.0 | 189 | 17.7 |
| s.3.5d0 | 171 | 18.8 | 187 | 18.5 |
| s.1d0.5 | 243 | 17.9 | 229 | 16.3 |
| s.1.5d0.5 | 205 | 17.4 | 204 | 15.8 |
| s.2d0.5 | 187 | 17.1 | 193 | 16.4 |
| s.2.5d0.5 | 180 | 17.6 | 187 | 17.2 |
| s.3d0.5 | 175 | 18.4 | 184 | 17.9 |
| s.3.5d0.5 | 175 | 19.2 | 183 | 18.8 |
| s.1d1 | 243 | 18.0 | 207 | 15.0 |
| s.1.5d1 | 206 | 17.6 | 190 | 15.6 |
| s.2d1 | 186 | 17.1 | 183 | 16.4 |
| s.2.5d1 | 181 | 17.6 | 180 | 17.2 |
| s.3d1 | 176 | 18.5 | 179 | 18.0 |
| s.3.5d1 | 175 | 19.2 | 180 | 18.6 |
| s.1d1.5 | 247 | 18.0 | 196 | 14.8 |
| s.1.5d1.5 | 206 | 17.4E | 182 | 15.6 |
| s.2d1.5 | 189 | 17.1 | 177 | 16.3 |
| s.2.5d1.5 | 182 | 17.6 | 175 | 16.4 |
| s.3d1.5 | 177 | 18.5 | 175 | 16.8 |
| s.3.5d1.5 | 175 | 19.3 | 177 | 17.4E |

TABLE XIV
ISCAS C2670 RESULTS

| ISCAS c2670 circuit | | | | |
|---------------------|-----------------------|------------------|-----------------------|------------------|
| | New | | Old | |
| | Propagation Delay(ns) | Power (μ w) | Propagation Delay(ns) | Power (μ w) |
| s.1d0 | 3.44 | 19.2 | 4.00 | 18.9 |
| s.1.5d0 | 3.11 | 20.0 | 3.59 | 19.8 |
| s.2d0 | 2.97 | 21.1 | 3.44 | 20.8 |
| s.2.5d0 | 2.91 | 22.2 | 3.34 | 21.8 |
| s.3d0 | 2.89 | 23.4 | 3.30 | 22.9 |
| s.3.5d0 | 2.87 | 24.5 | 3.28 | 24.0 |
| s.1d0.5 | 3.53 | 20.1 | 3.58 | 19.3 |
| s.1.5d0.5 | 3.17 | 20.9 | 3.40 | 20.3 |
| s.2d0.5 | 3.03 | 22.0 | 3.32 | 21.3 |
| s.2.5d0.5 | 2.96 | 23.1 | 3.27 | 22.4 |
| s.3d0.5 | 2.92 | 24.3 | 3.26 | 23.5 |
| s.3.5d0.5 | 2.91 | 25.4 | 3.26 | 24.5 |
| s.1d1 | 3.56 | 20.4 | 3.31 | 19.6 |
| s.1.5d1 | 3.21 | 21.3 | 3.24 | 20.6 |
| s.2d1 | 3.05 | 22.3 | 3.20 | 21.8 |
| s.2.5d1 | 2.97 | 23.4 | 3.20 | 22.9 |
| s.3d1 | 2.93 | 24.5 | 3.19 | 24.0 |
| s.3.5d1 | 2.92 | 25.6 | 3.20 | 25.0 |
| s.1d1.5 | 3.61 | 20.7 | 3.16 | 20.0 |
| s.1.5d1.5 | 3.23 | 21.5 | 3.14 | 21.0 |
| s.2d1.5 | 3.06 | 22.5 | 3.14 | 22.2 |
| s.2.5d1.5 | 2.99 | 23.6 | 3.14 | 23.3 |
| s.3d1.5 | 2.96 | 24.8 | 3.15 | 24.3 |
| s.3.5d1.5 | 2.93 | 25.8 | 3.17 | 25.4 |

TABLE XV
ISCAS CELL LIBRARIES

| Cell | C432 | C2670 |
|--------------|------|-------|
| 2 Input Nor | X | X |
| 3 Input Nor | | X |
| 4 Input Nor | X | X |
| 5 Input Nor | | X |
| 2 Input Nand | X | X |
| 3 Input Nand | X | X |
| 4 Input Nand | X | X |
| 5 Input Nand | | X |
| Inverter | X | X |
| Buffer | | X |
| AOI21 | X | |
| AOI22 | X | |
| AOI211 | X | |
| AOI221 | X | |
| OAI21 | X | |
| OAI22 | X | |
| OAI221 | X | |

TABLE XVI
ANTIFUSE TURN ON POINTS

| PMOS Size | Performance Degradation Ring Oscillator | Performance Degradation Fan Out Four | Performance Degradation C432 | Performance Degradation C2670 |
|-----------|---|--------------------------------------|------------------------------|-------------------------------|
| 1 | 14.57% | 21.77% | 13.20% | 16.25% |
| 1.5 | 13.26% | 14.59% | 10.64% | 15.39% |
| 2 | 12.34% | 15.28% | 10.22% | 15.73% |
| 2.5 | 11.52% | 11.26% | 9.86% | 14.74% |
| 3 | 10.96% | 7.02% | 9.50% | 14.28% |
| 3.5 | 10.43% | 9.88% | 9.78% | 14.40% |

VITA

Trenton Dean Henrichson was born in Edinburg, TX. He completed his Bachelor of Science Degree in computer engineering at Southern Methodist University, Dallas, Texas, in May 2004. He started his graduate studies as a computer engineering major at Texas A&M University in 2005.

He interned for Advanced Micro Devices in Austin, TX, from August 2006 through May 2007 as a Computer Aided Design engineer and continued to work at AMD from January 2008 until the present time in the same position. His research interests are in ASIC Design and Computer Aided Design for VLSI physical design.

Mr. Henrichson can be reached at the following email address:

henrichson.trenton@gmail.com. His permanent address is c/o Mr. Trenton Henrichson, 4700 Staggerbrush Rd., Austin, TX 78749.