Gate Overdrive with Split-Circuit Biasing to Substitute for Body Biasing in FinFET and UTB FDSOI Circuits

A Thesis

Presented to the faculty of the School of Engineering and Applied Science University of Virginia

in partial fulfillment

of the requirements for the degree

Master of Science

by

Andrew Ross Whetzel

May

2015

APPROVAL SHEET

The thesis

is submitted in partial fulfillment of the requirements

for the degree of

Master of Science

Mary MUT

The thesis has been read and approved by the examining committee:

Mircea Stan

Advisor

Benton Calhoun (Chair)

Ronald Williams

Accepted for the School of Engineering and Applied Science:

James H. Ayl

Dean, School of Engineering and Applied Science

May

2015

Acknowledgements

This thesis is dedicated to my family for their love and support throughout my education. I wouldn't have had the opportunities I've had without them, and I am thankful every day. I would also like to thank my friends and colleagues from home, James Madison University, and the University of Virginia. You all made this journey unforgettable.

Abstract

Body Biasing (BB) in bulk CMOS is an important tool for circuit designers as it allows for dynamic modulation of device threshold voltage post-fabrication. The ability to modulate device threshold voltages has many advantages. A higher threshold voltage results in lower standby power, while a lower threshold voltage results in increased performance. Threshold modulation allows different power and performance modes to meet both energy and throughput constraints. BB may be used to reduce the impact of process variations by adjusting nMOS and pMOS thresholds independently to maximize performance given a power constraint.

Fully-depleted silicon-on-insulator (FDSOI) FETs such as ultrathin body (UTB) devices may benefit from a similar effect to BB when the buried oxide (BOX) is thin enough to allow the back plane potential to affect the accumulation or inversion in the channel. However, when the BOX is thick the back plane potential has very little effect on the channel, eliminating the ability to modulate threshold voltage via back plane biasing (BPB). Similarly, FinFETs benefit very little from controlled body effect because the gate has virtually full control over the channel while the body potential has none.

In this thesis a new circuit topology is presented which substitutes for body biasing without relying on the body effect. The inputs, outputs, and supply rails are split in such a way that the gates of some devices may be overdriven without increasing voltage swing, resulting in a higher I_{on} and reduced latency under forward bias. Under reverse bias this topology drives V_{GS} below 0 V, thereby reducing leakage current. Through SPICE simulations of a 28nm FDSOI technology

a speedup of up to 21% has been realized under forward bias with an increase in power of 27%, while static power can be reduced by up to 43% with a 19% decrease in performance.

Table of Contents

Acknowledgements	i
Abstract	ii
List of Figures and Tables	vi
Chapter 1: Introduction	1
Introduction	1
Related Work	3
Thick BOX vs. Thin BOX SOI	8
Chapter 2: Split-Circuit Biasing	
Intuition: Modulate V_{GS} instead of V_{th}	12
Split-Circuit Topology	
A. Split-Circuit Inputs	14
B. Split-Circuit Outputs	14
C. Split-Circuit Supply Rails	15
D. Splitting and Merging Signals	
Split-Circuit Simulation Results	21
A. Ring Oscillators	21
B. FinFET FFT Butterfly Module	29
Split-Circuit Sequential Analysis	32
Chapter 3: Implementation	
Split-Circuit Layout	34
Generating the Voltage Difference (ΔV)	
Chapter 4: Conclusion	41
Future Work	

Conclusions	
REFERENCES	

List of Figures and Tables

igure 1: TCMS Buffer	6
igure 2: Quadrail AND Gate	7
igure 3: UT2B SOI Transistor	9
igure 4: Static CMOS and Split-Circuit Inverter Schematic	14
igure 5: Static CMOS and Split-Circuit NAND Schematic	16
igure 6: Cascaded Split-Circuit Inverters	16
igure 7: Static Power of Split-Circuit Buffer	19
igure 8: Stacked Split-Circuit Buffer Inverter Schematic	20
igure 9: Frequency and Static Power of BB Ring Oscillators	22
igure 10: Frequency, Static Power, and Energy/Op of Split-Circuit Ring Oscillators	23
igure 11: Frequency of RVT and LVT Ring Oscillators	28
igure 12: Delay and Static Power of Split-Circuit FFT Butterfly Module	31
igure 13: Stick Diagram of CMOS and Split-Circuit Inverter Layout	34
igure 14: Charge Pump Schematic	37
igure 15: Charge Pump Power Requirement	37
igure 16: Switched Capacitor Ground Split	38
igure 17: Charge Pump Output Voltage vs. Clock Frequency	40
able I: V _{GS} Swing of Split-Circuit Devices	17

Chapter 1: Introduction

Introduction

As transistors scale into the deep sub-micron era, process variations play a larger role in their operation. Silicon-on-insulator (SOI) devices have shown promise over bulk devices in being more variation tolerant, but there is always something to be desired in the way of increasing this tolerance [1]. Adaptive Body Biasing (ABB) has been shown to reduce the impact of die-to-die and within-die variations in [5] by changing the nMOS and pMOS threshold voltages independently in order to maximize performance given a power constraint. In [11], Qi and Stan use ABB to limit the effects of Negative Bias Temperature Instability (NBTI) and increase the lifetime of pMOS'. ABB may be used alongside other techniques for allowing different power modes to further increase the number of power modes available to users. In [14], Yan, Jiong, and Niraj use ABB alongside dynamic voltage scaling (DVS) to both control the leakage in sleep modes and control the active power consumption in active modes depending on the performance requirement. Body biasing is a ubiquitous tool used at varying granularities which allows designers to meet both power and performance requirements. SOI devices fabricated on a thin buried oxide (BOX) may benefit from body biasing in a similar way to bulk CMOS devices. In SOI, the back plane (silicon under the BOX) is biased in order to adjust the device threshold voltages. The back plane effectively acts as a second gate which can control the accumulation or inversion at the BOX-channel interface. This process is called back plane biasing (BPB). The top gate potential continues to have significantly more control over the channel because the gate oxide is considerable thinner than the BOX, which separates the channel from the back plane (the top gate is about 1/10th as thick for a thin BOX process, or about 2 nm) [2]. The back plane

potential may bring the channel closer or further from inversion, which either decreases or increases the device threshold voltage relative to the top gate [6] [13]. However, when SOI devices are fabricated on a thick BOX, the back plane and channel are significantly decoupled and BPB has little to no effect. In thick BOX devices, a back plane bias on the order of 10 V may be required in order to see a significant change in threshold voltage [11]. Similarly, FinFETs fabricated in bulk processes receive little benefit from controlled body effect because the channel in the FinFET is mostly in the top of the fin, away from the body [4]. While these technologies offer advantages in other dimensions, it is a disadvantage to have little control of threshold voltages post-fabrication.

Dynamic voltage scaling (DVS) is another method for both reducing power when high performance is not required and increasing performance to meet the demands of large workloads. In DVS this is accomplished by either raising the supply voltage to enhance performance or lowering the supply voltage to save power [15]. The main disadvantage of DVS is that active power increases as the square of the voltage swing of a gate, so enhancing performance comes at a high cost.

In this thesis a circuit topology is presented which substitutes for body biasing without relying on the body effect. While two supply voltage domains are modulated with respect to each other, the voltage swing of a gate is never increased, so the quadratic increase in active power is not an issue as it is in DVS. We will continue this chapter by comparing and contrasting related works, and delve further into the advantages of both thin BOX and thick BOX. In Chapter 2, we will discuss the intuition behind the split-circuit biasing scheme and the equations which dictate the operation of the devices. Chapter 2 will also include the split-circuit topology, and the differences between that and traditional CMOS will be outlined. Results from SPICE simulations will be presented and analyzed from two different implementations; ring oscillators and a butterfly module used in an FFT. An analysis of sequential elements in the split-circuit topology will wrap up Chapter 2. In Chapter 3, the layout of the split-circuit topology will be discussed, along with some area and design complexity tradeoffs. The generation of the bias voltage will also be discussed in Chapter 3. The thesis will wrap up in Chapter 4 with a discussion of future work and a conclusion.

Related Work

Power and performance optimization is a well traversed field. There are countless techniques for reducing static and active power and increasing performance, and in this paper we will compare against 4: Dynamic voltage scaling (DVS), body biasing (BB), and two similar schemes called threshold voltage control through multiple supply voltages (TCMS) and mixed swing quadrail.

Dynamic voltage scaling (DVS) is a method for dynamically changing the supply voltage at varying granularity to meet performance and power demands [15]. Similar to split-circuit biasing, static power is reduced when the supply (or voltage swing) is reduced. However, under DVS performance is enhanced at the expense of active power by increasing the supply voltage. Equation (1) illustrates the relationship between power and Vdd, or voltage swing:

$$P_{active} \propto C V_{dd}^2$$
 (1)

Where P_{active} is active power, C is load capacitance, and Vdd is the voltage swing of the gate. While DVS may achieve similar benefits as BB or BPB without relying on the body effect, active power increases quadratically with increases in Vdd. The topology presented in this paper allows for dynamic increases in performance without the expensive overhead of a quadratic increase in active power. This is accomplished by overdriving device gates by strategically varying source voltages such that the gate-to-source voltage (V_{GS}) increases without increasing the voltage swing of any gate.

As discussed in the Introduction, body biasing is a ubiquitous technique for modulating threshold voltage post fabrication through changes in the potential of the body. In bulk processes, the body is biased such that device threshold voltage (V_{th}) is decreased under forward bias, creating a faster but leakier transistor, while V_{th} is increased under reverse bias, creating a slower but less leaky device. In this manner, performance can be increased under forward bias and static power can be reduced under reverse bias. This technique may also be used to counteract die-to-die and within-die Vth variations, thereby increasing yield. In [19], a dynamic fine-grain body biasing (D-FGBB) approach is proposed which both reduces the impact of V_{th} variations and decreases static power compared to a similar static approach. In SOI, the approach is similar but the mechanisms by which V_{th} is varied are slightly different. Under the buried oxide (BOX) is a Si back plane which can work as a second gate if the BOX is thin enough. This oxide is thicker than the gate oxide, and therefore does not have as much control over the channel as the top gate. However, changing the potential of the back plane (back plane biasing, or BPB) does change the threshold of the device relative to the top gate, thereby mimicking body biasing. When the BOX is thick, however, this effect is severely reduced. The split-circuit

topology presented in this paper mimics BB or BPB without relying on the body effect, so it may be employed in bulk, thin BOX, and thick BOX technologies.

The two most closely related works to this thesis are techniques known as threshold voltage control through multiple supply voltages (TCMS) and mixed swing quadrail [20]. The TCMS work presented in [20] is motivated by connected-gate FinFETs, where the top gate of a FinFET acts as one gate rather than two independent gates. When the gates are left independent, applying a small voltage to one effectively reduces the threshold voltage of the device with respect to the other gate. This technique is not available when the two gates are shorted, presenting a situation similar to that of thick BOX SOI. The main goal of TCMS is decreasing power and delay on repeaters situated on long interconnects where the output capacitance is rather high. They accomplish this by varying the supply voltages for individual inverters in buffers, such that the voltage swing for inverters driving a long interconnect is not affected, but the delay and static power consumption of these inverters are reduced via a widened supply range (ground below 0V and supply above Vdd) on their inputs. See Figure 1 for an illustration of the buffer design in TCMS.

There are two main differences between TCMS and the work presented in this thesis. The first is that TCMS does not propose a change in circuit topology, but rather a strategic allocation of supply voltages, while the work presented in this thesis changes both topology and supply voltages. The second difference is the motivation/focus of the paper. While TCMS focuses on reducing power and delay on long interconnects, the split-circuit topology in this paper is

designed to enhance power or performance in any implementation which may call for traditional CMOS.



(b)

Figure 1: TCMS buffer (a) and repeaters on long interconnects (b) based on [20]. The second inverter stage of each buffer has devices which are sized larger (hence the larger symbol) in order to drive the relatively high capacitance of the long interconnects. The authors in [20] propose a VddH, VddL, and VssH of 1.08 V, 1 V, and -0.1 V, respectively, while VssL is ground. These voltages are derived for a 32nm FinFET model.

Mixed swing quadrail is very similar to TCMS in that it involves varying supply voltages in consecutive gates to reduce voltage swing on high capacitive loads while increasing their gate drive with high swing inputs. Unlike TCMS, however, the topology is changed in the quadrail scheme. Refer to Figure 2 for an illustration of the quadrail topology.



Figure 2: Mixed Swing Quadrail 3-input AND gate based on [21]. The inputs A, B, and C swing from Vss2 to Vdd2, which is a small voltage swing compared to Vss1 to Vdd1. Similarly, the output swings from Vss2 to Vdd2 because the final stage inverter is supplied by these voltages. The logic gate in this case is the 3-input NAND which is supplied by the larger swing Vss1 and Vdd1, allowing the final stage inverter to be overdriven with 'on' gate voltages above its own supply.

Buffers are added in between logic gates. These buffers are made up of minimum sized balanced inverters, so their input capacitance is very low. While the buffers are supplied with a relatively small voltage swing, the logic is supplied with a much larger voltage swing. Because the input

capacitance on the buffers is small, the power consumed at the output of high swing logic gates is kept relatively small. The buffers are able to drive the larger load due to their overdriven gates, as well as provide only a small voltage swing to minimize active power. Mixed swing quadrail is more similar to TCMS than it is to the split-circuit topology because it does not alter any traditional CMOS design, but simply adds buffers at a fine granularity between logic gates.

Thick BOX vs. Thin BOX SOI

This work is meant to substitute for body biasing. Therefore, this topology is most useful in technologies which cannot benefit from body biasing. Silicon-on-insulator (SOI) processes which are fabricated on a thin buried oxide (BOX) may benefit from a technique known as backplane biasing (BPB), whereby the bulk substrate under the BOX is biased and a result similar to body biasing in bulk processes is achieved. Refer to Figure 3 for an illustration of a transistor fabricated on thin BOX SOI.

It should come as no surprise that when the BOX is thicker, the electric field from the substrate has a smaller effect on the channel of the SOI device, thereby decreasing the effect BPB can have. While a thin BOX may be around 20nm or smaller, what we consider a thick BOX may be as thick as 150nm [3]. The strength of an electric field decreases as the inverse square of the distance from the source, as seen in the equation (2).

$$E = \frac{k \cdot Q}{d^2} \tag{2}$$



Figure 3: Ultrathin Body and BOX (UT2B) SOI transistor [18]. The source and drain are raised to reduce resistance. The BOX thickness, t_b , of a thin BOX device is on the order of 20nm [3], while the channel thickness, t_{si} , may be as small as 10nm [11].

Therefore, if we increase the BOX thickness by 7x, we will see a response to BPB that is decreased by a factor of about 50. It is these thick BOX devices that will benefit most from a supply biasing topology.

The question remains, why would we choose a thick BOX over a thin BOX? Certainly both have their advantages. Thin BOXs have been shown to allow better control of short channel effects (SCEs) via the reduction of the field-fringing effect, whereby the BOX potential mildly inverts the channel directly above the BOX (away from the gate) and leakage similar to DIBL is experienced [2] [7]. However, in order to significantly reduce the field-fringing effect, the BOX must be thinned down below 25 nm, which is not a simple process [2]. Additionally, because of the strong charge coupling between the back plane and channel in thin BOX devices, a relatively

large parasitic capacitance is experienced in thin BOX devices resulting in longer gate delay and higher power. Equation (3) explains this phenomenon further:

$$t_d = 0.69RC \tag{3}$$

 t_d is the gate delay, R is the output resistance, and C is the output capacitance. The only variable that is affected by the BOX thickness is C, which decreases as the BOX thickness increases. Equation (4) illustrates how the active power is affected:

$$P = f \alpha C_L V_{DD}^2 \tag{4}$$

P is the active power, f is the clock frequency, α is the activity factor, C_L is the output, or load capacitance, and V_{DD} is the operating voltage. Again, the only variable which changes between a thick and thin BOX is C_L ; C_L is reduced with a thick BOX, thereby reducing active power. Thick BOX devices become a very attractive option as chips become denser, both from a power and performance perspective. Additionally, we can expect the yield from a chip fabricated on a thick BOX to be higher because of the ease of fabrication when compared to a thin BOX [2].

Bulk FinFETs were the first three dimensional (3D) structure to be implemented on a commercial chip (Intel Ivy Bridge CPU) [3]. While they have been shown to respond to reverse body-biasing (i.e., leakage is decreased when the body potential is lowered), they do not experience a significant change in threshold voltage under different body biases [8]. This means

that we cannot increase or decrease the gate delay post fabrication to meet a performance need. Because bulk FinFETs are the most ubiquitously manufactured structure which shows promise in scalability, it is desirable to find a way to mimic BB and give users post-fab control over gate delay.

Chapter 2: Split-Circuit Biasing

Intuition: Modulate V_{GS} instead of V_{th}

Body biasing is effective in either increasing performance or decreasing static power by changing the device threshold voltage. In equation (5), which illustrates the relationship between leakage current (I_{off}), gate voltage (V_G), and threshold voltage (V_{th}), we can see that V_{th} is in the exponent of the exponential [17].

$$I_{off} \propto exp(V_G - V_{th}) \tag{5}$$

For the case of an nMOS, V_G is less than V_{th} when the transistor is off, or leaking, so the exponent is negative. When V_{th} is increased, as is the case is reverse BB, the exponent becomes more negative and the leakage current is decreased. However, the same effect could be achieved by decreasing V_G to a value below 0 (for an nMOS).

Equation (6) illustrates the relationship between dynamic current (I_{on}), gate-to-source voltage (V_{GS}), and V_{th} for both the triode and saturation regions of a short channel device [16].

$$I_{on} \propto (V_{GS} - V_{th})^{\alpha} \tag{6}$$

If we increase I_{on} we will decrease the delay of a pull up or pull down network, thereby increasing performance. Under forward BB this effect is achieved by lowering V_{th} . Both modes

of operation have an equivalent dependence on the gate-to-source voltage (V_{GS}) as they do on V_{th} . Therefore, if we increase V_{GS} we will achieve the same effect as forward body biasing.

Split-Circuit Topology

As discussed in the Intuition section, increasing the gate voltage compared to the source voltage of an nMOS when receiving a logic '1' will increase the drive strength. Similarly, decreasing the relative gate voltage of a pMOS will increase its drive strength with receiving a logic '0'. We can apply an analogous argument for the off state of transistors. If we decrease the relative gate voltage for an nMOS receiving a logic '0', the leakage current will be reduced. The same goes for increasing the relative gate voltage of a pMOS receiving a logic '1'. These insights are the driving force for the circuit topology presented in this thesis.

In order to allow V_{GS} to be modulated, a circuit topology with two supply domains is presented. One domain will be the nominal domain, with voltage swings from 0 to Vdd. The second domain will have the same differential, but both the ground and supply rail will be shifted up by some bias voltage, ΔV , such that the voltage swings of gates receiving this supply domain will swing from $0 + \Delta V$ to Vdd $+ \Delta V$. The inputs, outputs, and supply rails of a traditional CMOS topology are split such that the number of each are doubled. The following sections explain each in more detail. Refer to figure 4 for an illustration of the split inputs, outputs, and supply rails.



Figure 4: Static CMOS inverter with 2 fingers per transistor (a) and split-circuit inverter with split inputs, outputs, and supply rails (b).

A. Split-Circuit Inputs

For every traditional input, there are two inputs under the circuit topology presented in this thesis. Any two corresponding inputs will carry the same logic, but one is shifted up by some bias voltage, ΔV . Under forward bias, the higher inputs will drive the gates of the nMOS' while the lower inputs will drive the pMOS'; this will result in a higher V_{GS} for half of the nMOS' and a higher absolute value of V_{GS} for half of the pMOS', and therefore a higher I_{on} for half of the devices (see equation (6)). Under reverse bias, ΔV will be negative. Therefore the nMOS' will receive the lower voltage while the pMOS' receive the higher voltage. This results in a lower leakage current in half of the 'off' devices (see equation (5)).

B. Split-Circuit Outputs

As with the inputs, there will be two outputs per traditional output. There are two pull up and pull down networks per gate in the split-circuit topology. Between a pull-up and pull-down

network there is an output, just as there is in a traditional CMOS implementation. One of these outputs will be shifted up by ΔV .

C. Split-Circuit Supply Rails

There are four supply rails, two sets of two. Each set of supply rails will have the same differential voltage, but one set will be shifted by ΔV . Each set will be the source for one of the outputs, thereby shifting one of the outputs up by ΔV . The power rail voltages will be obtained by two off-chip supplies which have the same differential (Vdd), and one on-chip charge pump to maintain the voltage separation between the two domains (ΔV).

Figure 4 shows a split circuit inverter next to a CMOS inverter with two fingers per transistor (illustrated as two transistors). Notice that there are two inputs entering and two outputs leaving the circuit in figure 4 (b). These inputs and outputs correspond to the inputs and outputs you would expect to see in a static CMOS inverter, but in this design there is one extra input/output which is shifted up by ΔV . For an inverter, each input carries the same bit (logic '1' or logic '0'), as does each output. In general, every gate will have two inputs/outputs, one high and one low, for each input/output you would expect to see in a traditional implementation. Refer to figure 5 for a comparison of schematics of a CMOS NAND and a split-circuit NAND.

Figure 6 shows two cascaded split-circuit inverters, with the outputs of one driving the inputs of the second. Let us focus on the devices in the right gate in figure 6. Table 1 summarized the V_{GS} swings of these devices. Notice only half of the transistors, one pMOS and one nMOS, have a V_{GS} which is affected by split-circuit biasing. That is, their V_{GS} does depend on ΔV . Let us examine those devices which are dependent on ΔV .



Figure 5: Static CMOS NAND with 2 inputs (A & B), 1 output (Z), and 2 fingers per transistor (a) and split-circuit NAND with 4 inputs (AH, AL, BH, & BL) and 2 outputs (ZH & ZL).



Fig. 6: Two split-circuit inverter gates. The outputs of the left gate (inside dotted box) drive the inputs of the right gate such that the high output drives the nMOS' and the low output drives the pMOS'. The capacitors illustrate the relationship between the supply rails; two off-chip domains will supply the two Vdd differentials (0 - Vdd and ΔV - Vdd + ΔV) while an on-chip charge pump will maintain the difference between the two domains, ΔV .

Device	VGS Swing
MP1	0 to Vdd
MN1	ΔV to Vdd + ΔV
MP2	- ΔV to Vdd - ΔV
MN2	0 to Vdd

Table I: V_{GS} Swings of Devices in Figure 6

 V_{GS} for MN1 swings from ΔV when the input is a logic '0' to Vdd + ΔV when the input is a logic '1'. Under forward bias (positive ΔV), MN1 will be overdriven when the input is a logic '1', increasing drive strength and decreasing the time the output takes to pull-down to '0'. When the input is a logic '0' under forward bias, the V_{GS} will be ΔV . The leakage current through MN1 will increase exponentially with increased ΔV , which limits how far we can apply a forward bias in the split-circuit topology (see equation (5)). Under reverse bias (negative ΔV), when MN1 receives a logic '0' the V_{GS} will be $-\Delta V$, while V_{GS} is Vdd - ΔV when MN1 receives a logic '1'. Therefore static power is reduced when the transistor is off at the expense of performance when the transistor is on. The same analysis can be applied to MP2, and we find that it is overdriven in forward bias mode and static power is reduced in reverse bias mode.

D. Splitting and Merging Signals

In order for a signal to come in as one signal, be processed in this topology, and emerge as one signal, it will need to be split into a high and low signal at the beginning and merged into one

signal at the end. Merging the signal is simple; since we have one signal which swings from 0 to Vdd, we can simply allow that signal to propagate out of the logic which implements the splitcircuit topology while terminating the high signal. To split the signal we may use a buffer comprised of two split-circuit inverters. The first inverter will experience high static current due to a gate voltage which does not depend on ΔV , but rather swings from 0 to Vdd. Sizing the first inverter minimally will reduce the static power to a minimum under higher forward and reverse biasing. Figure 7 shows how the static power changes in this buffer under different bias voltages.

The static power increases dramatically in both directions when the input is low. Under forward bias, the increase is due to the second inverter in the buffer; the leakage increases in the same manner which it would for any other combinational logic as discussed earlier. Under reverse bias, the leakage increases only in the first gate. This is due to the leakage in the nMOS device which is supplied by ΔV ; the input to the gate is 0 V, so the V_{GS} of that device is the absolute value of ΔV under reverse bias. This causes the leakage to increase exponentially, as we see in the curves in figure 7 (a), and as we will later see in figure 10 (b). When the input to the buffer is high, we only see an increase in static power under forward bias. This is because both stages are seeing an increase in static power, which is why there is a sharper increase in static power for the high input under forward bias than for the low input.







(b)

Figure 7: Static power response of split-circuit buffer vs. ΔV for both a low and high input (a) and static power response of stacked buffer (b). The static power at nominal bias ($\Delta V = 0 V$) is 1.8 nW for the standard buffer and 52.9 pW for the stacked buffer.



Figure 8: Stacked split-circuit inverter for signal-splitting buffer. AL and AH correspond to the low and high input, respectively, while ZL and ZH correspond to the low and high outputs. For the inverter at the first stage of the buffer, AL and AH would be shorted and would always swing from 0 to Vdd.

These static power numbers are rather high. Under forward bias, we may assume that the circuit is in active mode; when in active mode, the static power is less of a concern as the active power takes over, so the static power of the buffer may be acceptable when the circuit is under forward bias [29]. However, under reverse bias we may assume that the circuit is often in standby or hibernate mode. In this case, the static power of the entire circuit is reduced because little or no processing is required [30]. Although the static power of the buffer increases similarly in both forward and reverse bias (for a low input), this increase may not be tolerated in reverse bias. A simple stacked structure as illustrated in [31] would cut down on static power in the buffer. See Figure 8 for a schematic of a stacked split-circuit inverter used in the stacked buffer. The normalized static power of the stacked buffer is shown in figure 7 (b). Although the static power

continues to increase exponentially in both forward and reverse bias for some input, the nominal $(\Delta V = 0 \text{ V})$ static power is about 3% that of the nominal for the non-stacked buffer (52.9 pW vs. 1.8 nW). Even at its maximum, the static power of the stacked buffer is less than twice the static power of the non-stacked buffer at nominal bias. The static power of the stacked buffer is acceptable even in reverse bias.

Split-Circuit Simulation Results

A. Ring Oscillators

In order to test the hypothesis that this biasing scheme will act similarly to body biasing, six ring oscillators were constructed; three oscillators which are body biased, made up of inverters, NAND, and NOR gates, and three comparable oscillators which are split-circuit biased. Each ring oscillator has 49 stages, including one NAND gate to allow the oscillations to be stopped to perform an analysis on the static power of the circuit. Figure 9 shows the performance and static power response of the body-biased ring oscillators, while Figure 10 shows the response of the split-circuit biased ring oscillators.



(a)



(b)

Figure 9: Body biased ring oscillators normalized frequency (a) and normalized static power (b) vs. bias voltage. The technology used to simulate was a 28nm FDSOI process with a thin BOX and ultrathin body. The nominal frequency ($\Delta V = 0 V$) for the inverter, NOR, and NAND oscillators was 1.98, 0.93, and 1.19 GHz, respectively. The nominal Vdd was 1.0 V.



(a)



(b)



(c)

Figure 10: Split-circuit ring oscillator normalized frequency (a), normalized static power (b), and normalized energy per operation vs. ΔV . The simulations were performed with a 28nm FDSOI ultrathin body and thin BOX technology. The body of each device was tied to the device's supply to prevent the body effect from playing a role. The nominal frequency ($\Delta V = 0 V$) for the inverter, NOR, and NAND oscillators was 2.05, 1.04, and 1.26 GHz, respectively. The nominal Vdd was 1.0 V.

The graphs in figures 9 (a) and 10 (a) illustrate the similarities between the performance response of a body biased oscillator and an oscillator biased in the split-circuit topology. While the body biased topology can have biases ranging from -4 V to around 1 V, the split-circuit biased topology realizes similar performance responses in a range of -0.2 V to -0.2 V. There is about a 10:1 correlation between the performance responses of body and split-circuit biasing, meaning that a 1 V body bias is equal to a 0.1 V split-circuit bias. The split-circuit frequency reached a maximum of a 49% increase when ΔV was 0.2 V. Figures 9 (b) and 10 (b) show that both body and split-circuit biasing can achieve a reduction in static power under reverse bias. As mentioned in the topology section, half of the devices are affected under the split-circuit topology. This means that the theoretical limit of static power reduction is 50% that of nominal. In Figure 10 (b) we can see that the static power reaches an asymptote near the 50% mark. Under deeper reverse bias, the static power begins to increase again. This is likely due to gate induced drain leakage (GIDL), whereby significant bandbending in the drain caused by a large gate-to-drain voltage allows band-to-band tunneling. Gate oxide tunneling could be another of the reasons for the increase in static power in deep reverse bias [26]. Gate oxide tunneling is a similar phenomenon to GIDL, where a large gate-to-drain voltage causes tunneling through the gate oxide between the gate and the drain [27]. Given that the static power increases under deep reverse bias and reaches an asymptote near 50% of the nominal static power, there is little reason to bias ΔV below -0.10 V for this technology. This observation is important when considering what charge pump will be used to create the differential, and what properties we will need that charge pump to have.

Figure 10 (c) shows the energy per operation as ΔV is swept. Although the voltage swing of every gate does not increase under different bias voltages, there is another mechanism which causes active power to increase in forward bias and decrease in reverse bias: short circuit current. This short circuit current does not manifest in the same way which we may be used to seeing in traditional CMOS, however. In traditional CMOS, short circuit current arises when the input signal is both high enough to turn the nMOS devices on and low enough to turn the pMOS devices on. This happens during transitions of the input, where decreasing the time that the transition takes will mitigate short-circuit current. This may also arise in traditional CMOS

when the inputs do not swing fully, or when SRAM cells or registers are at their metastable point [32]. However, because the inputs to the nMOS' and pMOS' are disjoint in the split-circuit topology, the mechanism by which short-circuit current arises is different. Under no bias, the mechanism appears to be almost identical; the signals are in phase, so they can be modeled as though they are not disjoint. In this case, decreasing the time it takes to make the transition will indeed decrease the short circuit current. However, when the bias voltage is non-zero, the short circuit current is caused by the difference in delay between the 'low' and 'high' signal; these signals do arrive at the same time because one of the outputs is overdriven while the other is not, resulting in a longer transition time for the output which is not overdriven. Under forward bias $(\Delta V > 0V)$ the 'high' signal leads the 'low' signal during a low-to-high transition, while the 'low' signal leads the 'high' signal during a high-to-low transition (remember that 'low' and 'high' refer to the parts of a signal which swing from 0 to Vdd and from ΔV to Vdd+ ΔV , respectively; these signals carry the same logic). During a low-to-high transition, both signals start low, so the nMOS devices are off and the pMOS devices are on. The 'high' signal transitions faster than the 'low' signal. Recall that the 'high' signal drives only the nMOS' while the 'low' signal drives only the pMOS'; this means that as the 'high' signal transitions from logic '0' to logic '1', the nMOS' turn on. But because the 'low' signal is lagging, the pMOS' do not turn off as quickly, resulting in short circuit current. The larger ΔV is, the larger the discrepancy between the 'low' and 'high' signal. The same consequence is seen during a high-to-low transition; the 'low' signal leads the 'high' signal, the pMOS' turn on before the nMOS' turn off.

Under reverse bias, the opposite is true. The 'low' signal leads the 'high' signal during a low-tohigh transition and the 'high' signal leads the 'low' signal during a high-to-low transition (we call the signal which is shifted by ΔV the 'high' signal, even when ΔV is negative. So in the case of negative ΔV , the 'high' signal actually swings between two lower voltages than the 'low' signal). Therefore, during a low-to-high transition, the pMOS devices actually begin to turn off before the nMOS devices begin to turn on, resulting in significantly reduced short-circuit current. The same effect is seen during high-to-low transitions. These phenomena explain why we see increases in energy/operation under forward bias in figure 10 (c), and decreases in energy/operation under reverse bias. Although we still exhibit some increase in active energy under forward bias, it is less than quadratic as we would see in a DVS scheme. At a forward bias of 20% Vdd ($\Delta V = 0.2$ V), the energy per operation increases on average by 30%, while we would see a 44% increase for the same supply increase in DVS. Under reverse bias, we see a minimum energy/op that averages 3% less than the nominal case

While split-circuit biasing provides a similar effect to body biasing in both forward and reverse bias modes, the inherent limitation of static power reduction makes it more attractive as a knob which can increase performance rather than as a knob which reduces static power. The 28nm FDSOI PDK provides two transistor types, low (LVT) and high (or regular, RVT) threshold. The higher threshold devices have a lower static power and decreased performance, while the low threshold devices have higher performances at the expense of static power. Because the split-circuit topology performs best under forward bias, it may be beneficial to use the high threshold devices to reduce standby power and use split-circuit biasing to increase their performance in active mode. An experiment was conducted to evaluate an implementation with high threshold devices versus an implementation with low threshold devices. Figure 11 shows the results from this experiment.



(a)



Figure 11: Ring Oscillator Frequency vs. ΔV (a) and static power vs. ΔV (b) for low threshold (LVT) and regular threshold (RVT) devices in a 28nm FDSOI process. The circuit simulated was a 49-stage NAND ring

oscillator. The nominal frequency ($\Delta V = 0$) for the LVT and RVT implementations were 1.33 GHz and 1.26 GHz, respectively.

At the nominal point ($\Delta V = 0$) the LVT implementation has a frequency of 1.33 GHz and a static power of 0.37 μ W, while the RVT implementation has a frequency of 1.26 GHz and a static power of 0.051 μ W. With a ΔV of 0.04 V, the frequency of the RVT implementation is 1.40 GHz while the static power is 0.095 μ W. This results in a 5% increase in frequency over the nominal case for the LVT implementation. In fact, we can increase the frequency of the RVT oscillator to 1.52 GHz (14 % higher than the nominal case of the LVT implementation) with a ΔV of 0.08 V and continue to have a static power that is less than 66% of the nominal case of the LVT implementation. With split-circuit biasing we can achieve better performances with higher threshold devices without sacrificing static power in standby mode.

B. FinFET FFT Butterfly Module

While oscillators allow us to perform power and performance analyses they do not have much application. To further test the split-circuit topology, we implement a butterfly module for a pipelined 16-point FFT with 8 bits/sample. In [9] a pipelined FFT architecture is discussed which uses one butterfly module to calculate a 16-point FFT, which is then used to calculate a 1024-point FFT in hardware without the expensive area overhead. An FFT implemented in hardware processes signals considerably faster than an FFT implemented in software [22].

As discussed earlier, FinFETs suffer from a similar phenomenon as thick BOX FDSOI devices. Because the fin protrudes from the bulk, the gate is able to cover a much larger portion of the channel. This added gate control allows the gate potential to have nearly full control over the inversion in the channel, which leaves the bulk potential with very little to contribute. Therefore, the bulk potential has very little effect on the threshold voltage of a FinFET device, eliminating the possibility for body biasing. In order to test the split-circuit topology in a FinFET technology, the butterfly module for the FFT was simulated using 7nm and 20nm high performance FinFET models from http://ptm.asu.edu. These models are based off of the BSIM-CMG models constructed by the BSIM group at UC Berkeley [23], and do not include effects from the body potential because they are rather negligible [24]. We can therefore assume that all changes in power and performance are due to split-circuit biasing (although the same assumption was made previously, it is even stronger in the case where there is no body effect accounted for in the device models). Figure 12 (a) shows the performance response of the FinFET butterfly module for the 7nm and 20nm models, while figure 12 (b) shows the static power response to split-circuit biasing.

As discussed earlier, the theoretical limit for decrease in static power is 50%. In figure 12 (b), we see that the static power is reduced to about 52% that of nominal ($\Delta V = 0 V$) when ΔV is reverse biased at -0.20 V. We can again see that the static power reaches an asymptote at around -0.10 V.

At a forward bias of 0.2 V the delay of the butterfly module reduced to 58% of the nominal delay for the 7nm models. Beyond this bias the module failed (that is, the outputs were no longer valid) for both the 7nm and 20nm FinFET models.



(a)



(b)

Figure 12: Normalized delay from a change in inputs to a change in outputs (a) and normalized static power in standby mode (b) for FFT butterfly module simulated with FinFET models. The nominal Vdd for the 7nm and 20nm models are 0.7 V and 0.9 V, respectively.

Split-Circuit Sequential Analysis

The results shown up until this point have been using circuits which are combinational. As discussed earlier, the 'high' and 'low' signals do not arrive at the same time when ΔV is non-zero. This creates a few negative impacts on sequential elements. The first is that the clock frequency will be limited by the slowest signal. Because CMOS is inverting by nature, the signal differential (difference in timing between 'high' and 'low' signals) does not compound significantly. This means that the main contributor to the signal differential, as it pertains to a sequential element, is the final stage of the combinational logic before the sequential element. The clock frequency will be limited by the final signal to arrive. SPICE simulations in a 28 nm FDSOI commercial technology of an 8-bit ripple carry adder show that the differential between the 'high' and 'low' parts of the slowest signal from input to output differ by about 100 ps in maximum reverse bias ($\Delta V = -0.2 V$) and about 50 ps in maximum forward bias ($\Delta V = 0.2 V$). These differentials are enough to impact the energy/operation due to short circuit current, but do not add a significant constraint on the clock period for longer combinational circuits.

For the FFT butterfly module, the largest differential between a corresponding 'high' and 'low' signal was slightly higher than 1% of the delay of the faster signal; this means that for longer combinational circuits, the 'high' and 'low' differential does not add a significant limit on the clock period. However, as we decrease the length of the stages in between registers, the difference in the final 'high' and 'low' signals becomes more and more significant, and further limits the minimum clock period. This means that the more pipelining we do, the smaller the advantage we can achieve from split-circuit biasing. To illustrate this further, consider the case where pipelining is so severe that between two registers all we have is one gate, such as a

NAND. Because the minimum clock period is constrained by the slow signal, and the slower signal is impacted very little by split-circuit biasing, we will basically see no change in total gate delay with split-circuit biasing because the most constraining signal, the slower one, will not see a change in delay. Split-circuit biasing is most advantageous in longer combinational circuits.

Chapter 3: Implementation

Split-Circuit Layout

There is some added complexity to the layout of the split-circuit topology due to doubling the inputs, outputs, and supply rails. While the total transistor width is the same across both the split-circuit implementation and the traditional static CMOS implementation, the added complexity for the split-circuit implementation resulted in an average increase in area of 32% for a small standard cell library for which the layout was conducted in a 28 nm FDSOI technology. Refer to figure 13 for stick diagrams of a traditional CMOS inverter and a split-circuit inverter.



Figure 13: Stick diagram of a CMOS inverter with two fingers per transistor (a) and a split-circuit inverter (b). The green rectangles represent active area, the red dotted lines represent poly, the solid blue line represents metal 1, the double purple line represents metal 2, and a black 'X' represents a contact when it is between metal 1 and poly or active area, or a via when it is between two metal layers. In (b), the two power and the two ground rails are shown as non-overlapping in order to illustrate what metals are used. In practice both ground rails and both power rails overlap, reducing the area illustrated in (b).

The total number of vias can provide us with a method for roughly quantifying layout complexity. This is because vias cause more strict and limiting spacing rules due to the interface between two metal layers, both of which must adhere to the spacing rules relative to the via (not necessarily relative to each other) [25].

Additionally, vias must be enclosed by both metals layers which puts a lower limit on the size of the metal enclosing the via, as well as causing spacing problems around vias [25]. This all contributes to lower density and higher areas, as well as higher complexity. A traditional CMOS inverter stick diagram is shown in Figure 13 (a). There are no vias in the traditional inverter, and only one layer of metal is used. Figure 13 (b) shows the stick diagram for a split-circuit inverter. Although the overall transistor size is the same across both implementations, the splitcircuit implementation requires two metal layers and six metal-to-metal interfaces. There are many problems which can arise at metal-to-metal interfaces. Since the vias are smaller than their enclosing metals, resistance is higher at these interfaces [28]. This also leads to higher current densities, especially for vias connecting supply rails. High current densities in vias may lead to electromigration [12]. Electromigration is a phenomenon by which metal migrates due to high density unidirectional currents [12]. This migration leads to even higher resistances, and in severe cases the via metal may completely break, resulting in fewer effective vias and even higher current densities in those which remain. Including multiple vias per junction helps to mitigate these problems. Figure 13 (b) shows that 6 metal 1-to-metal 2 interfaces are required for the split-circuit topology, which equates to 12 vias at this level (2 vias per interface). At broader levels of abstraction (block-level, etc.) more vias are required to accommodate higher currents. These vias were the main contributing factor to the 32% increase in area.

Generating the Voltage Difference (ΔV)

One of the challenges in the transition from traditional CMOS to the split-circuit topology is generating the bias voltage, ΔV , dynamically. The 'high' power and ground rails have the same voltage differential as the 'low' power and ground rails; we just need to find some way of maintaining a relative differential between the two power domains. The simplest way to do this would be to generate the voltages off-chip, but this is not so elegant. A simple on-chip charge pump will suffice, provided that the differential can be maintained without requiring that the charge pump provides a significant amount of power. This is a reasonable assumption because the only inherent current paths are between the 'high' power and ground rails and the 'low' power and ground rails, never between domains. The most significant path between domains would be gate-to-drain leakage, which is a small percentage of the overall power in the system. Figure 14 shows the power that will be required of the charge pump as a percentage of the total power while the inverter ring oscillator, whose results are shown in figure 10, is oscillating. Because there is no voltage drop across the charge pump at nominal bias point (ΔV), the power provided by the charge pump is exactly 0 W. As discussed earlier, the main current path between the two power domains is through gate oxide leakage. The percentage of total power provided by the charge pump is higher in reverse bias because the total power of the circuit decreases while gate oxide leakage increases.

The on-chip charge pump would reduce clutter and bulky circuitry. In [10] a charge pump is proposed to generate bias voltages in a body biasing implementation. The charge pump proposed in [10] may be found in figure 15; the charge pump has been modified to fit the requirements of split-circuit biasing.



Figure 14: Percentage of total power supplied by charge pump in a 49 stage split-circuit inverter ring oscillator. The oscillator was simulated with models from a 28 nm FDSOI ultrathin body and buried oxide technology.



Figure 15: Charge pump for generating body bias voltages based on Dickson charge pump in [10]. Switch A is closed during forward bias mode ($\Delta V > 0V$) while B is open, and switch B is closed in reverse bias mode ($\Delta V < 0V$) while A is open. The clock frequency determines the output voltage; a simple programmable ring oscillator is proposed in [10] for changing the clock frequency dynamically. All diodes and caps are nMOS devices in practice.

The charge pump is designed to maintain a relative voltage difference between true ground and the ground which is shifted by ΔV , while two off chip supplies maintain the voltage differential within each supply domain. The three diodes on the left, which in practice are three nMOS devices with their gates tied to their source, will pump charge toward the node to their right, increasing the voltage at ΔV when the circuit is in forward bias mode. When the circuit is in reverse bias mode, the three diodes on the right, which are pMOS devices, will pump charge away from ΔV , bringing ΔV below 0 V.

The off chip supplies need to have disjoint grounds. This is not a trivial problem to tackle because any supplies to the same chip would presumably have a common ground. However, we can use a switch capacitor circuit to maintain two voltage domains without shorting the two grounds together. Figure 16 illustrates such a switch capacitor.



Figure 16: Switch capacitor circuit to maintain separate grounds while the two supply domains maintain the same voltage swing. While the circuit illustrated above is off chip, the difference between the two domains (ΔV) is maintained with the on chip charge pump shown in Figure 15.

Provided that the capacitors in Figure 16 are large enough to maintain the voltage differential between clock periods, the switch capacitor circuit will allow us to disconnect the grounds of the two supplies and maintain the ΔV differential on chip without requiring the on chip charge pump

to supply a significant amount of current. The size of the capacitors in Figure 16, as well as the clock period, will need to be tuned to meet the power demand of the ΔV to Vdd + ΔV domain.

While a body biasing scheme may have bias voltages in excess of 1 V, the split-circuit biasing scheme cannot have bias voltage above the device threshold voltage, or about 300 - 400 mV, because a forward bias larger than the device threshold voltage would fully invert the channel of transistors which are supposed to be off, resulting in high short-circuit current and poor output characteristics. Figures 9 (a) and 10 (a) compare the performance response of body biasing and split-circuit biasing in comparable circuits. The results show that split-circuit biasing is much more sensitive to the bias voltage than body biasing, and therefore a charge pump with higher granularity would be required to achieve the same granularity of effective control over V_{th} . This sensitivity is advantageous in a few ways: first, the charge pump will not have to supply as much power as it would if it needed to supply a rather high voltage. A charge pump that doesn't need to supply a significant amount of power can be smaller, saving both area and power to operate the charge pump. Second, the charge pump can be operated at a lower clock frequency. A lower clock frequency to operate the charge pump also means lower power to operate the charge pump. See equation (4) for the relationship between power and clock frequency. Finally, a smaller bias voltage means a shorter time for the charge pump to adjust and reach a steady voltage.

Figure 17 shows the output voltage of the charge pump shown in figure 15 as a function of the clock frequency. When the clock does not oscillate, the output voltage settles at a non-zero value

for both the forward and reverse bias mode; therefore the clock must always be active when the circuit is operating, even when the bias voltage is 0 V.



Figure 17: Output voltage of charge pump illustrated in figure 13 as a function of clock frequency. The nominal Vdd is 1 V, which means the clock swings from 0 V to 1 V.

Chapter 4: Conclusion

Future Work

Thus far, all layouts have been conducted in a custom fashion, although there has been some work in the way of synthesis. Characterizing the function of the inputs and outputs has proven to be a significant hurdle. The tools do not necessarily know that a 'high' output should go to the 'high' input (the input which drives the gates of the nMOS devices) of the following gate, because both the 'high' and the 'low' signals carry the same logic. Additionally, when characterizing the power and delay of the gates, there are forbidden cases which the tool must not consider. These are the cases where the 'high' and 'low' part of a signal carry different logic; because the tool treats inputs as independent, by default it will evaluate these cases unless we can tell it that these cases will not arise and need not be considered.

A tapeout is scheduled for May 6th, 2015. This chip will include ring oscillators, which will allow us to confirm the SPICE simulations in silicon. The next step will be to implement the split-circuit topology in a more holistic circuit, including both combination and sequential elements, in silicon. The synthesis research will aid in this; a flow from hardware description language to layout would significantly reduce design time and allow for larger circuits to be built.

Conclusions

Planar bulk MOSFETs are reaching an asymptote of their scalability, and designers are turning to Silicon-on-insulator devices to take us towards the single-digit nm nodes. SOI devices

fabricated on thick buried oxides may be desired for their ease of fabrication, reduced parasitic capacitance, and increased performance, but the thick BOX impedes the ability to modulate device threshold voltage post-fabrication. In this thesis, a circuit topology is proposed which substitutes for body biasing when the body effect is negligible and body biasing is not an option. Unlike DVS, the voltage swing is always the nominal supply voltage for all gates under all biases, which allows dynamic power to increase sub-quadratically while performance is enhanced. SPICE simulations have confirmed that, similarly to body biasing, we can increase performance or decrease static power with forward or reverse split-circuit biasing postfabrication. While the theoretical limit of static power reduction is 50% that of nominal (zero bias), simulations in a 28nm FDSOI technology have shown a maximum reduction in static power to 57% that of the nominal case. At the highest bias voltage before circuit failure, a performance increase of 92% is realized for a NOR ring oscillator in the 28nm FDSOI technology. A butterfly module for a pipelined FFT was simulated with 7nm and 20nm FinFET models from http://ptm.asu.edu. A propagation delay decrease of 66% was realized under maximum forward split-circuit bias, and static power was reduced to 52% of that of nominal static power consumption under maximum reverse bias. Simulations confirm that split-circuit biasing gives effective control over device threshold post-fabrication in technologies which cannot benefit significantly from controlled body effect.

REFERENCES

[1] Flandre, Denis, et al. "Fully-depleted SOI CMOS technology for low-voltage low-power mixed digital/analog/microwave circuits." *Analog Integrated Circuits and Signal Processing* 21.3 (1999): 213-228.

[2] Trivedi, Vishal P., and J. G. Fossum. "Nanoscale FD/SOI CMOS: Thick or thin box?." *Electron Device Letters, IEEE* 26.1 (2005): 26-28.

[3] James, Dick. "Intel Ivy Bridge unveiled—The first commercial tri-gate, high-k, metal-gate CPU." *Custom Integrated Circuits Conference (CICC), 2012 IEEE*. IEEE, 2012.

[4] Chiarella, Thomas, et al. "Benchmarking SOI and bulk FinFET alternatives for PLANAR CMOS scaling succession." *Solid-State Electronics* 54.9 (2010): 855-860.

[5] Tschanz, James W., et al. "Adaptive body bias for reducing impacts of die-to-die and withindie parameter variations on microprocessor frequency and leakage." *Solid-State Circuits, IEEE Journal of* 37.11 (2002): 1396-1402.

[6] Fenouillet-Beranger, C., et al. "FDSOI devices with thin BOX and ground plane integration for 32nm node and below." *Solid-State Electronics* 53.7 (2009): 730-734.

[7] Ernst, T., et al. "Fringing fields in sub-0.1 μm fully depleted SOI MOSFETs: optimization of the device architecture." *Solid-State Electronics* 46.3 (2002): 373-378.

[8] Park, Tai-su, Euijoon Yoon, and Jong-Ho Lee. "A 40nm body-tied FinFET (OMEGA MOSFET) using bulk Si wafer." *Physica E: Low-dimensional Systems and Nanostructures* 19.1 (2003): 6-12.

[9] Li, Weidong, and L. Wanhammar. "An FFT processor based on 16-point module." *Proc. of NorChip Conf.* 2001.

[10] Kim, Chris, and Kaushik Roy. "Dynamic VTH scaling scheme for active leakage power reduction." *Proceedings of the conference on Design, automation and test in Europe*. IEEE Computer Society, 2002.

[11] Eminente, S., et al. "Ultra-thin fully-depleted SOI MOSFETs: Special charge properties and coupling effects." *Solid-State Electronics* 51.2 (2007): 239-244.

[12] Zhang, Runjie, et al. "Some limits of power delivery in the multicore era."*Proceedings of WEED* (2012).

[13] Ohata, A., et al. "Mobility enhancement by back-gate biasing in ultrathin SOI MOSFETs with thin BOX." Electron Device Letters, IEEE 33.3 (2012): 348-350.

[14] Yan, Le, Jiong Luo, and Niraj K. Jha. "Joint dynamic voltage scaling and adaptive body biasing for heterogeneous distributed real-time embedded systems." *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 24.7 (2005): 1030-1041.

[15] Putic, Mateja, et al. "Panoptic DVS: A fine-grained dynamic voltage scaling framework for energy scalable CMOS design." *Computer Design*, 2009. ICCD 2009. IEEE International Conference on. IEEE, 2009.

[16] Sakurai, Takayasu, and A. Richard Newton. "Alpha-power law MOSFET model and its applications to CMOS inverter delay and other formulas." *Solid-State Circuits, IEEE Journal of* 25.2 (1990): 584-594.

[17] Roll, Guntrade. Leakage Current and Defect Characterization of Short Channel MOSFETs.Diss. Zugl.: Erlangen, Nürnberg, Univ., Diss, 2012. N.p.: n.p., n.d. Print.

[18] Tang, Stephen H., et al. "FinFET-a quasi-planar double-gate MOSFET." Solid-StateCircuits Conference, 2001. Digest of Technical Papers. ISSCC. 2001 IEEE International. IEEE, 2001.

[19] Teodorescu, Radu, et al. "Mitigating parameter variation with dynamic fine-grain body biasing." Microarchitecture, 2007. MICRO 2007. 40th Annual IEEE/ACM International Symposium on. IEEE, 2007.

[20] Muttreja, Anish, Prateek Mishra, and Niraj K. Jha. "Threshold voltage control through multiple supply voltages for power-efficient FinFET interconnects."VLSI Design, 2008. VLSID
2008. 21st International Conference on. IEEE, 2008.

[21] Krishnamurthy, Ram K., and L. Richard Carley. "Exploring the design space of mixed swing quadrail for low-power digital circuits." *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on* 5.4 (1997): 388-400.

[22] Chang, Yun-Nan, and Keshab K. Parhi. "An efficient pipelined FFT architecture." Circuits and Systems II: Analog and Digital Signal Processing, IEEE Transactions on 50.6 (2003): 322-325.

[23] Hu, Chenming et al. "BSIM-CMG." *BSIM Group*. UC Berkeley, 2012. Web. 16 March 2015.

[24] Hu, Chenming et al. "BSIM-MG F.A.Q." *BSIM Group*. UC Berkeley, 2012. Web. 16 March 2015.

[25] Cong, Jason, Jie Fang, and Kei-Yong Khoo. "Via design rule consideration in multilayer maze routing algorithms." Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on 19.2 (2000): 215-223.

[26] Choi, Yang-Kyu, et al. "Investigation of gate-induced drain leakage (GIDL) current in thin body devices: single-gate ultra-thin body, symmetrical double-gate, and asymmetrical doublegate MOSFETs." Japanese journal of applied physics 42.4S (2003): 2073.

[27] Lee, Dongwoo, David Blaauw, and Dennis Sylvester. "Gate oxide leakage current analysis and reduction for VLSI circuits." Very Large Scale Integration (VLSI) Systems, IEEE Transactions on 12.2 (2004): 155-166.

[28] Savidis, Ioannis, and Eby G. Friedman. "Closed-form expressions of 3-D via resistance, inductance, and capacitance." Electron Devices, IEEE Transactions on 56.9 (2009): 1873-1881.

[29] Lackey, David E., et al. "Managing power and performance for system-on-chip designs using voltage islands." Computer Aided Design, 2002. ICCAD 2002. IEEE/ACM International Conference on. IEEE, 2002.

[30] Lee, Y-H., Krishna P. Reddy, and C. Mani Krishna. "Scheduling techniques for reducing leakage power in hard real-time systems." Real-Time Systems, 2003. Proceedings. 15th Euromicro Conference on. IEEE, 2003.

[31] Rao, Rahul M., Jeffrey L. Burns, and Richard B. Brown. "Circuit techniques for gate and sub-threshold leakage minimization in future CMOS technologies."Solid-State Circuits Conference, 2003. ESSCIRC'03. Proceedings of the 29th European. IEEE, 2003. [32] Holcomb, Daniel E., Wayne P. Burleson, and Kevin Fu. "Power-up SRAM state as an identifying fingerprint and source of true random numbers." Computers, IEEE Transactions on 58.9 (2009): 1198-1210.