Adaptive Power Gating of 32-bit Kogge Stone Adder

Alexander E. Shapiro*, Francois Atallah**, Kyugseok Kim**, Jihoon Jeong**, Jeff Fischer**,
and Eby G. Friedman*

* Department of Electrical Engineering, University of Rochester, Rochester, NY 14627

** Qualcomm Corporation, Raleigh, NC 27617

Abstract—Static power consumes a significant portion of the available power budget. Consequently, leakage current reduction techniques such as power gating have become necessary. Standard global power gating approaches are an effective method to reduce idle leakage current, however, global power gating does not consider partially idle circuits and imposes significant delay and routing constraints. An adaptive power gating technique is applied locally to a 32-bit Kogge Stone adder, and evaluated at the 16 nm FinFET technology node. This high granularity adaptive power gating approach employs a local controller to lower energy use and reduce circuit overhead. The controller conserves additional power when the circuit is partially idle (based on the inputs to the adder) by adaptively powering down inactive blocks. Moreover, the local controller reduces routing complexity since a global power gating signal is not required. The proposed adaptive power gating technique exhibits significant energy savings, ranging from 8% to 21%. This technique targets partially idle circuits, and therefore complements rather than replaces global power gating techniques. A 12% delay overhead results in a 5% area overhead. This delay overhead is reduced to 5% by increasing the area overhead to 16%, and can be further reduced by trading off additional area.

I. INTRODUCTION

Scaling of the minimum feature size has been applied to enhance the speed, area, and power of high complexity integrated circuits. Scaling, however, results in reduced threshold voltage, channel length, and gate oxide thickness. A nanometer scale gate oxide leaves insufficient material to prevent oxide tunneling. Additionally, the deeply scaled channel increases source-to-drain leakage current. Techniques to lower leakage power have therefore become a primary objective in modern microprocessors [1]. Power gating is a well known and efficient technique to reduce leakage current [2]. Microprocessor wide power gating has however significant overhead such as sleep and wake up delay, layout congestion, and physical area. These overheads as well as the complexity of system wide integration have limited industry wide adoption of global power gating [3].

Two decades have passed since the concept of power gating was first introduced [4]. In his work, S. Mutoh et al. utilized low threshold voltage transistors to enhance the speed of a power gated circuit, while high threshold voltage power gates reduced the leakage current during a sleep period. Consequently, a number of studies has been published to address the most common disadvantages of power gating. The

noise voltage induced by the rush current at wakeup has been addressed in [5] with skewed wakeup timing. Data retention mode operation is presented in [6] to eliminate information loss. The size of the power gate as a tradeoff between standby leakage current and speed degradation has been examined in [7]. An automated gate biasing technique has been developed to determine the gate bias voltage which minimizes leakage current [8]. Power gating characterization at early stages of the design process has been proposed [9]. Novel power gating approaches that utilize nano-electro-mechanical power switches with zero leakage current (off state) rather than MOS power switches have been examined [10]. Application guided power gating techniques that reduce leakage current in a register file based on the software state have also been developed [11].

This work proposes an adaptive and independent power gating technique. The proposed high granularity local power gating approach lowers overhead and improves efficiency as compared to global power gating. This technique is demonstrated here on a 32-bit Kogge Stone adder [12]. The adder is divided into an energy efficient number of clusters that can be independently powered down when inactive. The primary contribution of this work is the local controller that enables fine grain power gating of a clustered circuit by adapting to the current state of the adder. This technique saves additional power when a circuit is partially active. The nearby location of the controller also reduces system level complexity such as layout congestion and sleep/wake up delay, leading to additional power savings.

The paper is organized as follows. In Section II, the structure of the 32-bit Kogge Stone adder is reviewed. In Section III, the proposed adaptive power gating technique is described. Evaluation of the adaptive power gating technique is provided in Section IV. The paper is concluded in Section V.

II. 32-BIT KOGGE STONE ADDER STRUCTURE

The proposed adaptive power gating technique is evaluated on a 32-bit Kogge Stone adder [12]. The circuit consists of three major stages; the input, carry propagation network, and output stages. Two one-bit inputs are processed by the input stage, generating P_i (propagate) and G_i (generate) signals. The propagate signal indicates that the carry signal from the $i-1$ bit can propagate through bit i , and the generate signal indicates that the carry signal is generated at bit i . Both propagate and generate signals cannot be simultaneously high, and are described by, respectively, (1) and (2). The carry propagate network is responsible for propagating the carry signal from

This research is supported in part by the Binational Science Foundation under Grant No. 2012139, the National Science Foundation under Grant No. CCF-1329374, the IARPA under Grant No. W911NF-14-C-0089, and by grants from Qualcomm, Cisco Systems, and Intel.

Fig. 1: 8-bit Kogge Stone adder within a 32-bit Kogge Stone adder. The white blocks represent the bit propagate (BP) cells (input stage), solid gray blocks represent the group propagate (GP) cells, dotted gray blocks represent the group generate (GG) cells (carry propagation stage), and XOR blocks return the summation result (output stage). The critical delay path is highlighted by the bold line.

the previous bit lines. This carry propagate stage generates P_i^* and G_i^* , given by, respectively, (3) and (4), which passes the carry signal to the last output stage. The output stage performs the XOR operation between the generate signal of the previous bit (G_{i-1}) and the propagate signal of the current bit (P_i). The entire structure is illustrated with an 8-bit adder example, as shown in Figure 1.

$$(A_i)XOR(B_i) = P_i \quad (1)$$

$$(A_i)AND(B_i) = G_i \quad (2)$$

$$(P_i)AND(P_{i-1}) = P_i^* \quad (3)$$

$$(G_{i-1})AND(P_i)OR(G_i) = G_i^* \quad (4)$$

The adder is logically organized into horizontal clusters of one bit, each comprising the three major stages, the single bit input, carry propagate network, and output. Each cluster is independent except for the downstream carry network. Independent clustering enables the proposed fine grain, adaptive, and local power gating technique. Furthermore, the clusters can be combined to form up to 32-bit clusters. The most efficient cluster size is discussed in Section III-D.

III. ADAPTIVE POWER GATING OF 32-BIT KOGGE STONE ADDER

The proposed adaptive power gating methodology consists of three major components: power switches, isolation cells, and a controller, which are described in the following subsections. The optimal cluster size is examined in the last subsection.

A. Power switches

A power switch is a PMOS or NMOS transistor that disconnects the circuit from the power supply, ground, or both power and ground networks, when power gating is engaged. The addition of the power switches results in a secondary power network which can be disconnected from the primary power network. This secondary power network is referred to here as the virtual V_{DD}/Gnd network. Traditionally, the choice between a footer or header device [3] is based on the circuit structure and performance tradeoffs. Both footer and header power switches are rarely used in the same circuit due to the increased delay and area overhead. With FinFET technology, the performance gap between NMOS and PMOS is less [13], leaving the choice based solely on the circuit structure. For example, if the V_{DD} supply voltage is externally gated, a header switch powers down the internal circuit. Alternatively, if a high voltage is preferred at the outputs of the power gated circuit, disconnecting the ground with a footer switch is preferable.

In this work, the internal blocks benefit from a low voltage at the outputs of the power gated block. The power switches therefore use PMOS header transistors. The low voltage at the outputs, in this configuration, is provided by the isolation cells, as described in the next subsection.

The length of the power switch (18 nm) is based on the maximum I_{on}/I_{off} ratio as a function of gate length. The width of the power switch is a function of the size and peak current of the cluster as well as the number of distributed power switches. To limit the effect on circuit speed, the width is adjusted until a 10% constraint on delay is satisfied.

B. Isolation cells

Power gating is engaged by disconnecting a circuit from the power supply, resulting in a floating voltage at the circuit outputs. When this floating voltage is propagated to the inputs of the downstream logic, a short-circuit current is produced in the half open PMOS and NMOS transistors, as illustrated in Figure 2a. To lessen this effect, the active blocks are isolated from the power gated blocks by asserting high or low logic at the outputs of the power gated block.

A logic state is guaranteed by the isolation cells at the outputs of the power gated circuit. Two examples of an isolating cell are shown in Figure 3. These cells force either a high or low logic voltage at the output when *Isolate* is high, otherwise the cells are transparent to the downstream circuit. The AND/OR isolation cells, shown in Figure 3a, require more area than the single transistor isolation cells; however, the single transistor isolation cell (shown in Figure 3b) allows a direct current path between V_{DD} and ground.

To lower the area overhead, the isolation cells use an NMOS pull-down transistor. These isolation cells clamp the output to a low voltage when activated. The PMOS power switch network also complements the NMOS isolation cells by disconnecting the supply voltage when the power switches are engaged, eliminating the direct path between V_{DD} and ground. Using

Fig. 2: Power gated circuit with and without isolation cells. A short-circuit current is generated due to (a) floating output without isolation cell, as opposed to (b) constant output with isolation cell.

Fig. 3: Isolation cell structure; (a) single gate structure, and (b) single transistor structure.

high V_T transistors further lowers the leakage current caused by the inactive isolation cells.

C. Controller

Controllers are used in standard power gating applications to control and synchronize local power switches and isolation cells with clock gating or power gating signals. In the Kogge-Stone adder, the controller enables the adaptive power gating scheme. Two controller configurations are used to adapt to a specific input scenario of the 32-bit Kogge Stone adder. A simple controller and an enhanced controller are discussed in this section.

1) *Simple adaptive controller*: As described in Section II, the carry propagate stage of the adder is responsible for a significant fraction of the area and energy consumption of an adder. This carry propagation network, however, is redundant for bit i when both input bits, A_i and B_i , are zero. The simple adaptive controller recognizes this specific input pattern and applies power gating to the corresponding carry network. In the case of clusters larger than a single bit

size, as exemplified by the 4-bit clusters shown in Figure 4, the controller considers four bits of the corresponding inputs. In this case, the controller implements a boolean function, $(A_i + B_i + A_{(i+1)} + B_{(i+1)} + A_{(i+2)} + B_{(i+2)} + A_{(i+3)} + B_{(i+3)})'$ to recognize the four consecutive zero input bits. The controller dynamically power gates only those circuits that are part of the carry propagation network. The power gated circuits, as illustrated in Figure 4, are highlighted in gray and are isolated from the active output stage (highlighted by the black diagonal stripes) using isolation units (black/white diamonds). The input stage (highlighted in white) and the output stage are maintained active to guarantee correct operation of the adder.

Fig. 4: Simple adaptive controller applied to 8-bit Kogge Stone adder with four bit clusters. The power gated carry network of the adder is highlighted in gray. The isolation units are represented by black and white diamonds. The active (not power gated) input stage is highlighted in white, and the active output stage is highlighted in black diagonal stripes.

2) *Enhanced adaptive controller*: The enhanced controller includes the basic function of the simple adaptive controller. In addition to this basic function, this controller recognizes a second input pattern when the carry network is redundant and can be powered down. This case occurs when bit i does not generate a carry signal (i.e., both inputs A_i and B_i are not high), and no carry signal exists from the previous $i - 1$ bit. The proposed enhanced controller, as illustrated in Figure 5, monitors each input bit pair, A_i and B_i , for the occurrence of two input patterns (when the carry network is redundant) based on the Boolean functions described by (5) and (6),

$$G_0 + P_0 \times C_{in} = CR_0 \quad (5)$$

$$G_i + P_i \times CR_{i-1} = CR_i. \quad (6)$$

The primary benefit of this enhanced controller is the significantly expanded pool of input patterns that becomes available for power gating (not limited to a zero input as with the simple adaptive controller). A daisy chain connection between bit lines in the controller is however required, as highlighted by the dotted lines shown in Figure 5. Due to this long chain, the power gating signal from the controller to the power switches is delayed for most of the significant bits of the adder, leaving a smaller fraction of the clock period to save static power. Moreover, the daisy chain can lead to redundant

Fig. 5: Enhanced adaptive controller applied to 8-bit Kogge Stone adder with four bit clusters. The structure of the 8-bit Kogge Stone is the same as shown in Figure 4, and is therefore omitted except for the input stage.

logic switches inside the controller, further reducing energy efficiency.

D. Optimal cluster size with shared power switches

The proposed power gating circuit with the simple adaptive controller has been evaluated to determine the preferable clustering size. The gates within a single cluster share a common virtual power network with distributed power switches. The size of a cluster can be as small as a single bit line that includes the input stage of the relevant input bits, carry propagation network, and output XOR that returns the sum of the pair of input bits. Alternatively, the entire adder can be partitioned into a single 32-bit cluster sharing a single virtual power network.

The four major parts of the proposed power gating circuits, the controller, power switches, control wires, and isolation cells, have been evaluated to determine an efficient cluster size. The controller has a single logic stage (before the driver of the power gates) with a single bit cluster configuration. The logic depth of the controller increases by one stage with an increase in cluster size (decrease in cluster number). The controller is evaluated in different configurations that correspond to the partition of the adder. To determine the overhead of the power switch, the dynamic energy is evaluated for different cluster sizes. With a large number of clusters (producing smaller individual clusters), the total peak current sunk by the circuit is higher, requiring larger power switches to maintain the same performance. Similarly, the overhead of the isolation cell is dependent on the size and number of clusters. More isolation cells are needed with a larger number of clusters. Additionally, the distribution network of the control signal is adjusted for different cluster sizes.

The control network is split into L1 and L2 parts. The L1 section connects the controller output to L2 for each cluster, as illustrated in Figure 6a. The L2 "H" tree distributes the signal internally to the power switches, as shown in Figure 6b. Each wire segment of the control network is modeled with a corresponding parasitic capacitive and resistive T network. The

Fig. 6: Distribution network of the control signal; (a) L1 control lines from controller to the clusters, and (b) H-tree structured control lines inside the cluster from L1 to the distributed power switches and isolation units.

value of the parasitic capacitances and resistances are from ITRS [14], and scaled to the wire length of each segment. The wire lengths are approximated for a commercial 16 nm FinFET process, as shown in Figure 6. The energy overhead of the controller, control wires, and isolation cells is measured by dedicated power supplies. The L1 and L2 signal distribution network utilizes repeaters to drive the network impedance. The energy overhead of the control wires therefore includes the energy overhead of the inserted repeaters as well as the parasitic impedance of the interconnect. The power switch overhead includes the energy of the switching gate capacitance of the power gates.

The results of this analysis are presented in Figure 7. As illustrated in Figure 7a, the 4-, 8-, and 16-bit clusters exhibit the highest energy efficiency. A higher energy overhead is reported for both the smaller and larger clusters. As illustrated in Figure 7b, for the smaller clusters, the overhead is due to the inefficient control network and greater number of isolation units. For the larger clusters, the overhead is due to the larger controller as well as the sub-optimal control network. The smallest, yet energy efficient cluster size of four bits is therefore used to exploit greater granularity without compromising energy efficiency.

IV. EVALUATION AND RESULTS

Application of the proposed power gating technique to a 32-bit Kogge Stone adder with 4-bit clusters is evaluated

(a)

(b)

Fig. 7: Energy savings and overhead as a function of cluster size, (a) energy savings and overhead of the total recoverable energy (2 GHz cycle), and (b) distribution of the energy overhead. The diagonally striped bars represent the overhead of the power gating units, and the gray bars represent the savings in energy.

using 16 nm FinFET PTM models [15]. Both of the controller configurations, simple and enhanced controllers, are compared to a non-power gated version of the adder. The reported energy consumption of the power gated circuit includes the overhead of all of the power gating components described in Section III.

A. Simple adaptive controller

The circuits are injected with a randomly generated input pattern during the initial 2 GHz clock period followed by a new randomly generated input pattern during the following clock period. The second input pattern, however, is modified to include a different number of 4-bit zero groups starting from the most significant bit (MSB). The simulation is averaged over ten iterations for each number of 4-bit zero groups of the second input, as depicted in Figure 8. This analysis is used to evaluate the possible energy savings of the proposed fine grain power gating approach when the 32-bit input operands of the adder contain unused higher order bits. A sample analysis (which does not include the adder circuit) is performed to provide insight into the probability of the target input patterns. In this analysis, eight industry standard software benchmarks (eembc, coremark, linpack, octane, ragdoll, spec2006, spec2000, and sunspider) are executed on a commercial micro-

(a)

(b)

Fig. 8: Energy and delay of the power gating application with simple adaptive controller; (a) Energy consumption, and (b) delay overhead. The diagonally striped bars represent the standard circuit, gray bars represent the power gated circuit, and the black bars represent the difference in per cent.

processor to simulate general software activity. The number of leading zero bytes within the input variables to the arithmetic unit in the microprocessor is determined during the execution. This sample analysis demonstrates that for the majority of inputs to the adder (on average, 62%) at least one significant byte is equal to zero (two 4-bit clusters can be power gated, saving up to 12% energy). For 30% of the input patterns, the three most significant bytes are equal to zero (six 4-bit clusters can be power gated, saving up to 19% energy).

The energy is measured during the second cycle by evaluating the integral of the instantaneous power over the entire cycle. The delay is measured at the 50%-to-50% transitions at the output. An increase in energy savings of up to 21% is illustrated in Figure 8a. The power gated circuit is 8% more energy efficient when fully active due to the power switches that limit the dynamic and static power dissipation, trading off longer delay. The reduced power consumption is primarily due to the additional resistance (of the power gate) added between the power and ground networks, as well as the reduced voltage swing across the active circuit due to the voltage drop across the power gate. Alternatively, the power gated circuit exhibits up to 21% savings in energy when all of the eight clusters are powered down. Note that the energy consumption, both for the power gated and non-power gated circuits, declines with increasing number of powered down clusters. This behavior is

expected due to the lower activity of the adder with the larger number of 4-bit zero groups in the inputs.

The primary tradeoff of applying power gating is delay overhead, as illustrated in Figure 8b. As shown in Figure 8b, the delay overhead is, on average, 12% and does not vary significantly as a function of the number of powered off clusters. This result agrees with the expected 10% overhead, as described in Section III-A.

1) *Comparison of the simple adaptive power gating technique to standard power gating approach:* The simple adaptive power gating technique differs from standard power gating due to the dedicated controller and additional isolation units. Both techniques share the same underlying network of power gates and control lines. This network of power gates, as described in Section III-A, is the primary contributor to the delay overhead of the power gated circuit. The delay overhead of the power gated circuit is, however, a function of the size of the power gates which should be optimized for each power gated circuit based on local speed/area constraints. A dedicated controller, therefore, does not contribute additional delay since the controller operates parallel to the adder. Moreover, the additional isolation units contribute an insignificant delay with a similar contribution both to the standard power gating and simple adaptive power gating approaches.

The energy overhead, however, is higher in simple adaptive power gating but the contribution due to the controller and the additional isolation cells is low. The deleterious effect on the energy savings can be noted in the two extrema cases, as shown in Figure 9.

Fig. 9: Comparison of simple adaptive power gating technique to standard power gating approach.

In the first case, the entire adder is power gated (the right most column in Figure 9). The simple adaptive controller is redundant, wasting energy as compared to the standard power gating approach. In the second case, the adder is fully active (the left most column in Figure 9). The standard and simple adaptive power gating approaches are therefore both redundant. In the adaptive approach, additional energy is wasted due to the dedicated controller and the added isolation units. In these cases, the simple adaptive power gating technique exhibits a small energy overhead as compared to standard power gating. The contribution of the controller to the total energy overhead is insignificant since the controller is smaller

than the power gated circuit. 56% more isolation cells are required by the adaptive power gated circuit due to the fine grain clustering of the adder circuit. The contribution of the isolation units to the total overhead is however less than the controller, as depicted in Figure 7b. In other cases, however, when the power gated circuit is partially active, the additional savings in energy is significant, ranging between 3% to 13%. This comparison illustrates the significant benefit of the simple adaptive approach as compared to the standard power gating approach when the circuit is partially active. Although the adaptive technique has additional overhead when the circuit is either fully active or inactive, the overheads are, respectively, either small or insignificant due to the rare occurrence of these states.

B. Enhanced adaptive controller

The input pattern of the first cycle is identical to the simple adaptive controller. In this case, however, the second input pattern is entirely random. The results presented in this section are averaged over 1,000 randomly generated iterations. This analysis is, therefore, more conservative and overly pessimistic as compared to the analysis presented in section IV-A. The

Fig. 10: Probability distribution of inputs as a function of the number of powered off clusters.

randomly generated input exhibits the distribution depicted in Figure 10. The most common result of a random input pattern is one powered down cluster with a probability of 0.31. The next most common outcome is two powered down clusters with a probability of 0.24, and the third most common pattern is zero powered down clusters with a probability of 0.23. This distribution shows that in more than two thirds of the input scenarios at least one cluster can be powered down. In a practical application with correlated inputs, however, the number of clusters that can be powered down is likely to be greater.

The enhanced controller exhibits increased energy overhead resulting in insignificant energy savings at 2 GHz with a random input pattern, as shown in Figure 11a. Note that the savings in energy declines with increasing number of powered down clusters. The decline is due to the redundant transitions at the internal nodes which increase with larger number of powered down clusters due to the longer daisy chain in the controller, as discussed in Section III-C2. Nevertheless, during

(a)

(b)

Fig. 11: Energy consumption of power gating application with enhanced adaptive controller at (a) 2 GHz clock frequency, and (b) 1 GHz clock frequency. The diagonally striped bars represent the standard circuit, gray bars represent the power gated circuit, and the black bars represent the difference in per cent.

steady state operation, the enhanced controller exhibits a significant reduction in static power, as shown in Figure 12. A

Fig. 12: Static power savings of enhanced adaptive controller. The diagonally striped bars represent the standard circuit, gray bars represent the power gated circuit, and the black bars represent the difference in per cent.

longer clock period is, therefore, required to demonstrate these energy savings. As illustrated in Figure 11b, with a longer clock period (1 GHz frequency), the enhanced controller exhibits a significant savings in energy of up to 15%. Note

that, as opposed to the simple controller, the energy does not decline with an increasing number of powered down clusters. This result is expected due to averaging of the random input patterns.

Fig. 13: Delay overhead when power gating with enhanced adaptive controller. The diagonally striped bars represent the standard circuit, gray bars represent the power gated circuit, and the black bars represent the difference in per cent.

Similar to the previous analysis with the simple controller, the delay overhead is approximately constant over the range of powered down clusters. The delay overhead is a function of the number and size of the power switches which do not depend on the number of powered down clusters. The power switch network is the same for both controllers. The delay overhead is therefore similar. The average delay overhead, as shown in Figure 13, is close to 13% which is similar to the simple controller. This result agrees with the expected 10% overhead, as described in Section III-A.

V. CONCLUSIONS

Adaptive power gating of a 32-bit Kogge Stone adder is discussed in this paper. This technique is not limited to this adder and can be expanded to other major units within a modern microprocessor. The adaptive controller enables high granularity local power gating unavailable in global power gating. This high granularity provides additional power savings when the circuit is partially active (and cannot be globally power gated). Adaptive power gating exhibits significant energy savings, ranging from 8% to 21%, requiring a delay overhead of 12% that can be reduced to 5% by increasing the area overhead from 5% to 16%. Additional benefits such as lower layout complexity and reduced sleep/wake up delay overhead are also demonstrated.

REFERENCES

- [1] E. Salman and E. G. Friedman, *High Performance Integrated Circuit Design*. McGraw-Hill Publishers, 2012.
- [2] D. Flynn, R. Aitken, A. Gibbons, and K. Shi, *Low Power Methodology Manual For System-on-Chip Design*. Springer Publishers, 2007.
- [3] L. Benini and G. D. Micheli, *Dynamic Power Management: Design Techniques and CAD Tools*. Kluwer Academic Publishers, 1998.
- [4] S. Mutoh, T. Douseki, Y. Matsuya, T. Aoki, S. Shigematsu, and J. Yamada, "1-V Power Supply High-Speed Digital Circuit Technology with Multithreshold-Voltage CMOS," *IEEE Journal of Solid-State Circuits*, Vol. 30, No. 8, pp. 847–854, August 1995.

- [5] K. Usami, T. Shirai, T. Hashida, H. Masuda, S. Takeda, M. Nakata, N. Seki, H. Amano, M. Namiki, M. Imai, M. Kondo, and H. Nakamura, "Design and Implementation of Fine-Grain Power Gating with Ground Bounce Suppression," *Proceedings of the IEEE International Conference on VLSI Design*, pp. 381–386, January 2009.
- [6] S. Kim, S. Kosonocky, D. Knebel, K. Stawiasz, and M. Papaefthymiou, "A Multi-Mode Power Gating Structure for Low-Voltage Deep-Submicron CMOS ICs," *IEEE Transactions on Circuits and Systems II: Express Briefs*, Vol. 54, No. 7, pp. 586–590, July 2007.
- [7] A. Ramalingam, B. Zhang, D. Pan, and A. Devgan, "Sleep Transistor Sizing Using Timing Criticality and Temporal Currents," *Proceedings of the IEEE Asia and South Pacific Design Automation Conference*, Vol. 2, pp. 1094–1097, January 2005.
- [8] A. Valentian and E. Beigne, "Automatic Gate Biasing of an SCCMOS Power Switch Achieving Maximum Leakage Reduction and Lowering Leakage Current Variability," *IEEE Journal of Solid-State Circuits*, Vol. 43, No. 7, pp. 1688–1698, July 2008.
- [9] H. Xu, R. Vemuri, and W.-B. Jone, "Dynamic Characteristics of Power Gating During Mode Transition," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 19, No. 2, pp. 237–249, February 2011.
- [10] M. Henry and L. Nazhandali, "NEMS-Based Functional Unit Power-Gating: Design, Analysis, and Optimization," *IEEE Transactions on Circuits and Systems I: Regular Papers*, Vol. 60, No. 2, pp. 290–302, February 2013.
- [11] H. Tabkhi and G. Schirner, "Application-Guided Power Gating Reducing Register File Static Power," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Vol. 22, No. 12, pp. 2513–2526, December 2014.
- [12] P. M. Kogge and H. S. Stone, "A Parallel Algorithm for the Efficient Solution of a General Class of Recurrence Equations," *IEEE Transactions on Computers*, Vol. C-22, No. 8, pp. 786–793, August 1973.
- [13] S. Sinha, G. Yeric, V. Chandra, B. Cline, and Y. Cao, "Exploring Sub-20nm FinFET Design with Predictive Technology Models," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 283–288, June 2012.
- [14] "The International Technology Roadmap for Semiconductors," 2013.
- [15] Y. K. Cao, "Predictive Technology Models," June 2012. <http://ptm.asu.edu/>.