

# VARIUS-TC: A Modular Architecture-Level Model of Parametric Variation for Thin-Channel Switches \*

S. Karen Khatamifard, Michael Resch, Nam Sung Kim<sup>†</sup>, Ulya R. Karpuzcu  
University of Minnesota <sup>†</sup>University of Illinois  
{khata006, resc0059, ukarpuzc}@umn.edu {nskim}@illinois.edu

## Abstract

*Under aggressive miniaturization, unconventional digital switches rapidly come to light, which introduce new sources of variation in design parameters, and hence challenge the manufacturing process further. As a result, performance and power of manufactured hardware becomes greatly unpredictable. Characterizing variation-incurred unpredictability at early stages of the design necessitates dependable architecture-level models of variation, which distill device- and circuit-level details to accurately evaluate system-level implications. In this paper, we introduce a modular architecture-level model of parametric variation to address this challenge. As a case study, we refine our discussion to a representative class of emerging thin-channel switches, FinFETs.*

## 1. Introduction

Both, aggressive miniaturization and rapid introduction of unconventional materials and device architectures (which in turn enables aggressive miniaturization) make the manufacturing process of contemporary digital switches less controllable. At the same time, device-level optimizations to facilitate robust operation under aggressive miniaturization often require modification of the manufacturing process itself, which usually gives rise to new sources of variation in process parameters. As a result, deviation of device parameters from their nominal design specification becomes more likely, rendering performance and power efficiency of manufactured hardware greatly unpredictable. Unlocking the performance and power efficiency benefits of emerging switches hence mandates characterization of variation-incurred unpredictability across the system stack. This necessitates dependable architecture-level models of variation to enable accurate evaluation of system-level implications at early stages of the design.

The speed at which unconventional materials, device architectures and optimizations come to light can easily prohibit the longevity of a tightly calibrated architecture-level model of parametric variation. To overcome this challenge, we introduce VARIUS-TC, a dependable model which decouples architecture-level analysis from circuit- and device-level characterization by careful abstraction. VARIUS-TC comprises three modules which span device, circuit, and architecture layers of the system stack under process variation, respectively:

\*This work was supported in part by NSF under grant CCF-1421988; DARPA under PERFECT Contract Number HR0011-12-2-0019; and Samsung Electronics under contract 2016-03835. Kim has a financial interest in AMD and Samsung Electronics.

- The *device module* encapsulates switch-level electrical characteristics such as current-voltage profiles. Closed-form formulae or look-up tables (LUT) can capture electrical characteristics. LUTs can be particularly useful in analyzing emerging devices where no known closed-form formulae exist.
- The *circuit module* encompasses performance and power characteristics at logic gate and memory cell granularity. The output of the device module represents the input of the circuit module.
- The *architecture module* derives the performance and power characteristics at pipeline and (on-chip) memory block granularity. The output of the circuit module represents the input of the architecture module.

By swapping corresponding variants of device and/or circuit modules, VARIUS-TC can keep track of architecture-level implications of process variation as novel switches get discovered. In this paper, we focus on a representative class of emerging thin-channel switches, FinFETs, as a case study. In the following, Section 2 provides the background; Sections 3 and 4 explain the parametric variation model; Sections 5 and 6 cover the evaluation and a representative use case of VARIUS-TC; Section 7 compares and contrasts VARIUS-TC to related work; and Section 8 concludes the paper.

## 2. Background

### 2.1. Thin-Channel Switches




Figure 1: Planar vs. thin-channel switches. Adapted from [1].

The G(ate) of a classic digital switch controls the formation of a conductive path – the channel – between the S(ource) and D(rain) terminals, to turn on the switch. An ideal switch prohibits current flow when turned off:  $I_{on}/I_{off} \rightarrow \infty$ , where  $I_{on}$  ( $I_{off}$ ) corresponds to the current flow between S and D when the switch is on (off). Under aggressive miniaturization

the distance between S and D becomes so small that G loses control over the channel, resulting in a non-negligible  $I_{off}$ . To prevent excessive growth of  $I_{off}$ , G needs stronger control over the channel. Two emerging classes of device architectures, ultra-thin body silicon on insulator (UTB-SOI) and fin field effect transistor (FinFET), achieve this by eliminating excess silicon material between S and D, i.e., by making the channel thinner [1]. Fig. 1 depicts the device architecture for a classic old-school switch, a *planar* metal-oxide field effect transistor (MOSFET) in (a); UTB-SOI in (b); and FinFET in (c), with the channel thickness explicitly marked. Without loss of generality, the rest of the paper is confined to FinFETs as representative thin channel devices. FinFET switches have their channel rotated by  $90^\circ$  with respect to the planar design, which results in a *fin* sticking out of the substrate plane (Fig. 1(c)). The gate surrounds the fin for enhanced channel control, which in turn prevents excessive  $I_{off}$ .

## 2.2. Process Variation

Due to the stronger control of the gate over the channel, FinFET features better performance and power characteristics than planar MOSFET [8], however, the 3D structure (Fig. 1(c)) complicates the manufacturing process, and introduces new sources of variation. The main challenge stems from the manufacture of thin and uniform fins [1].

Die-to-die (D2D) variation in process parameters mainly stems from *systematic*, i.e., strongly correlated, effects such as lithographic aberrations. D2D variation causes similar changes in the values of switch parameters across the die. Within die variation (WID) in process parameters, on the other hand, can be due both *systematic* and *random* (i.e., uncorrelated) effects such as random dopant fluctuation. Each switch in the die may have a different change in the value of each affected parameter due to WID variation. In the following, we focus on the system-level impact of WID variation in FinFET parameters. A global offset per die (for each affected parameter) can capture D2D variation on top.

All geometric dimensions shown in Fig. 1(c) – fin thickness, fin height, and channel length, i.e.,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$  – are subject to variation, along with the oxide thickness,  $T_{ox}$ , and metal gate work function,  $\Phi_g$  [13, 16, 10]. Variation in  $\Phi_g$  comes from the variation in the orientation of the gate metal grains. The threshold voltage,  $V_{th}$ , strongly depends on fin dimensions – all subject to variation [13], and governs both performance and power. Variation in  $V_{th}$  degrades operating speed by increasing the variance of the critical path distribution. Higher variance causes longer tails in the critical path distribution, where the fastest and the slowest paths reside. In the end, the slowest path determines the operating speed. Moreover, variation in  $V_{th}$  induces a higher static power consumption: for the same variation-induced  $\Delta$  change in  $V_{th}$ , switches with  $-\Delta$  leak more than switches with  $+\Delta$  save. The variation in  $L_{Fin}$ ,  $T_{Fin}$ , and  $\Phi_g$  has the highest impact on performance and (static) power [11, 10].

## 3. VARIUS-TC: Macroscopic View

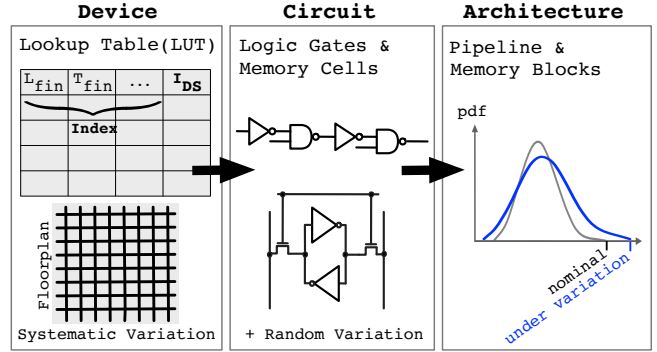


Figure 2: Overview of VARIUS-TC.

Parametric variation can easily degrade the operating speed and increase the static power consumption significantly, hence wipe out performance and power benefits of emerging switches such as FinFET. Accordingly, accurate characterization of variation-incurred unpredictability at early stages of the design is critical. Fig. 2 provides an overview of VARIUS-TC, a modular, dependable architecture-level variation model. VARIUS-TC features three basic modules which span device-, circuit-, and architecture layers of the system stack. Inputs to VARIUS-TC are the chip floorplan at functional block granularity and physical FinFET parameters for the respective technology node.

### 3.1. Device Module

The *device module* encapsulates closed-form formulae or lookup tables (LUT) to derive  $I_{on}$  and  $I_{off}$  from (variation-afflicted) physical FinFET parameters. As high-fidelity closed-form formulae can easily become computationally expensive for an architecture-level model, VARIUS-TC provides support for LUTs. The index (i.e., input query) to the LUT is a vector comprising  $V_{DS}$  (the drain-source voltage),  $V_{GS}$  (the gate-source voltage),  $T_{ox}$ ,  $\Phi_g$ ,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$ , respectively. The drain-source current  $I_{DS}$  represents the LUT output. Depending on the values of  $V_{GS}$  and  $V_{DS}$ ,  $I_{DS}$  may correspond to  $I_{on}$  or  $I_{off}$ . **LUT Generation:** Each row of the LUT corresponds to a specific index vector  $\langle V_{DS}, V_{GS}, T_{ox}, \Phi_g, T_{Fin}, H_{Fin}, L_{Fin} \rangle$ . We sweep the value of  $T_{ox}$ ,  $\Phi_g$ ,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$  over a specific range, at a specific resolution. Both the range and the resolution represent VARIUS-TC parameters which evolve as a function of the anticipated (maximum) deviation of physical FinFET parameters from their nominal values under variation.  $V_{DS}$ ,  $V_{GS}$  values reflect anticipated operating voltages. For each combination, a SPICE simulation generates the corresponding  $I_{DS}$  by utilizing a technology model file (such as Predictive Technology Model (PTM) [15, 18]) which tabulates low-level FinFET parameters. Each such index vector along with the corresponding  $I_{DS}$  represents a row of the LUT. A finer parametric resolution comes at the cost of increased storage overhead. The LUT also considers different temperature

values.

**Generation of Systematic Variation Maps:** In addition to LUT generation, the device module is in charge of the preparation of systematic variation maps. Following [13, 10] VARIUS-TC models the variation in  $T_{ox}$ ,  $\Phi_g$ ,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$  by separate Gaussian distributions, each characterized by a different  $\sigma/\mu$ . The nominal values from the technology model file represent the (expected value)  $\mu$ . The (standard deviation)  $\sigma$  represents yet another VARIUS-TC parameter. In order to capture within-die (WID) systematic variation, the device module super-imposes a grid on the chip floorplan and assigns to each grid point a sample from the Gaussian distributions of  $T_{ox}$ ,  $\Phi_g$ ,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$ , respectively. WID variation of each parameter exhibits spatial correlation, which VARIUS-TC captures by a spherical function: For each parameter subject to variation, the correlation between two switches on die only depends on their Euclidean distance. The correlation assumes its maximum value at distance zero, decreases with increasing distance, and vanishes once the distance exceeds the *correlation range*,  $\phi$ . Similar to [17, 9], VARIUS-TC captures random variation analytically: The circuit module (Section 3.2) combines the systematic and random variation in deriving gate- and (memory-)cell-level performance and power characteristics under variation.

### 3.2. Circuit Module

The *circuit module* uses the outcome of the device module – the LUT and systematic variation maps – in deriving gate- and (memory-)cell-level performance and power characteristics under variation. Systematic variation maps incorporate profiles for all,  $T_{ox}$ ,  $\Phi_g$ ,  $T_{Fin}$ ,  $H_{Fin}$ , and  $L_{Fin}$ . The circuit module extracts the vector  $\langle T_{ox}, \Phi_g, T_{Fin}, H_{Fin}, L_{Fin} \rangle$  for each grid point from the systematic variation map and uses this as the query vector to retrieve the corresponding  $I_{DS}$  from the LUT. Next, the module plugs  $I_{DS}$ , along with the query vector, into the performance model derived from [14]. VARIUS-TC mimics VARIUS(NTV) to factor in random variation analytically in deriving both performance and power, and in calculating the minimum safe operating voltage,  $V_{MIN}$ , per memory cell, under variation.

### 3.3. Architecture Module

Following VARIUS(NTV) methodology, the *architecture module* uses the outcome of the circuit module (1) to extract the critical path distribution within a pipeline stage or a memory block; (2) to determine the minimum safe operating voltage of each memory block,  $V_{MIN}$ ; (3) to calculate the minimum safe clock period,  $\tau_{MIN}$ , at a given supply voltage  $V_{dd}$ ; and (4) to report the probabilities of variation-induced logic and memory (timing or stability) errors as a function of the operating point (i.e., the clock period and supply voltage). A timing error in a logic block emerges if variation-incurred slowdown prevents safe operation at the designated clock period. Similarly, a timing error in a memory block emerges during a read or write, if

variation-incurred slowdown prevents completion of the read or write within the designated time window. A stability error is the case, if excessive leakage under variation corrupts the memory content even if the memory is not accessed.

## 4. VARIUS-TC: Microscopic View

### 4.1. Capturing Variation in Logic Blocks

A vector of variation-afflicted physical FinFET parameters,  $\langle T_{ox}, \Phi_g, T_{Fin}, H_{Fin}, L_{Fin} \rangle$ , represents the query index to the LUT (Sections 3.1, 3.2). VARIUS-TC always extracts the LUT entry with the minimum difference to the query index. After retrieving the respective (variation afflicted)  $I_{on}$  from the LUT for a given variation profile, *circuit module* determines the variation in gate delay from  $CI_{on}/V_{dd}$ , where  $C$  represents the equivalent load capacitance, and  $V_{dd}$ , the operating voltage. In this manner, VARIUS-TC implicitly considers  $I_{on}$ 's dependence on  $V_{th}$  and temperature.  $C$  is mainly a function of the floorplan and practically does not change with variation.

Using circuit module's gate delay distributions, VARIUS-TC's *architecture module* determines the path delay distribution of the slowest pipeline stage under variation following VARIUS(NTV) methodology [17, 9]. The path delay distribution under variation,  $D_{Var}$ , dictates the maximum path delay  $\max(D_{Var})$  which in turn determines the minimum possible value of the clock period  $t_{CLK}$ . The timing error rate per cycle while operating at a given clock period  $t_{CLK}$  becomes  $1 - cdf_{D_{Var}}(t_{CLK})$ , where  $cdf$  represents the cumulative distribution function.  $cdf_{D_{Var}}(t_{CLK})$  captures the cumulative probability of path delays (in the respective pipeline stage) to assume a lower value than  $t_{CLK}$ . A timing error occurs iff any path delay exceeds the designated clock period  $t_{CLK}$ .

A pipeline stage encompasses many paths of different delays. Even if there was no variation, there exist a specific distribution of path delays,  $D_{Logic}$ . Variation transforms this distribution to  $D_{VarLogic}$ :

$$\begin{aligned} D_{Logic} &= D_{Gates} + D_{Wire} \\ D_{Wire} &= k_W \times D_{Logic} \\ D_{Gates} &= (1 - k_W) \times D_{Logic} \end{aligned} \quad (1)$$

$D_{Gates}$  from Equation 1 corresponds to the delay of a sequence of gates along a path modulo wires;  $D_{Wire}$ , to the wire delay, were there no variation. VARIUS-TC neglects variation in wire delay. The multiplier  $k_W$  captures the share of wire delay in  $D_{Logic}$ .


Equation 2 gives the path delay distribution of a (purely logic) pipeline stage under variation.  $D_{VarGates}$  reflects changes in  $D_{Gates}$  under variation considering both the systematic ( $D_{VarSysGates}$ ) and random ( $D_{VarRandGates}$ ) components:

$$\begin{aligned} D_{VarLogic} &= D_{VarGates} + D_{Wire} \\ &= D_{SysGates} + D_{RandGates} + D_{Wire} \\ D_{SysGates} &= k_{Sys} \times D_{Gates} = k_{Sys} \times (1 - k_W) D_{Logic} \end{aligned} \quad (2)$$

Under aggressive miniaturization, the footprint of a pipeline stage becomes so small, that all of its paths fall within the correlation distance  $\phi$  of systematic variation. Therefore, being highly correlated, parameters of all gates along a path in a (purely logic) pipeline stage change practically in the same direction by the same quantity. The coefficient  $k_{Sys}$  captures this change due to systematic variation. The random component does not exhibit any correlation, hence a similar coefficient to  $k_{Sys}$  does not apply. VARIUS-TC derives the random component by composing independent and identically distributed gate delay distributions (under random variation) instead.

## 4.2. Capturing Variation in On-chip Memory Blocks

When compared to logic, on-chip memory blocks are more susceptible to variation due to both the higher density (i.e., higher number of transistors per area) and smaller size transistors. VARIUS-TC models the variation in a conventional 6T(ransistor) cell (Figure 3(a)) following VARIUS(NTV) methodology [17, 9]. The cell consists of two inverters, formed by  $PR-NR$  and  $PL-NL$  in a positive feedback loop, and two access transistors,  $AXR$  and  $AXL$ .  $V_R$  stores the cell's value and  $V_L$  its complement. To read from or write to the cell, word-line  $WL$  is driven high to connect the cell to the bit-lines  $BL$  and  $BR$ . To read, the bit-lines are pre-charged to logic high. To write,  $BR$  is pre-conditioned to the value to be written, and  $BL$ , to its complement. In the following, we assume that  $V_R=0$  without loss of generality. Since the cell is symmetric, the discussion applies directly to  $V_R=1$ .



**Figure 3: 6-T(ransistor) (a) and 8-T cell (b).  $V_R$  and  $V_L$  are the voltages at nodes  $R$  and  $L$ , respectively.**

If, during a read, the time needed to produce a voltage difference between the two bit-lines exceeds the period that  $WL$  stays high, a timing error occurs. Variation-induced increases in  $V_{th}$  of the discharge transistors,  $AXR$  and/or  $NR$  can trigger timing errors during read, because the transistors become slower. A timing error can also occur during a write, if the write is unable to change the logic state of the target cell by the end of the designated write duration. Variation-induced shifts in the switching threshold of the  $PR/NR$  inverter and/or  $PL$  becoming stronger than  $AXL$  under variation can trigger such errors. Variation can also distort the logic value stored in the cell due to excessive leakage of the transistors forming the inverters, even if the cell is not being accessed: Variation-induced shifts in  $V_{th}$  of the transistors forming the inverters, which increase the leakage, trigger such stability errors.

VARIUS-TC supports the 8-T(ransistor) cell of Figure 3(b) [5], as well, which is optimized for low power operation. This cell is easier to design reliably because it decouples the transistors used for reading from those used for writing such that they can be optimized separately. Specifically, the two additional transistors  $NRD$  and  $AXRD$  are only responsible for reads, while the rest coordinates the writes exactly as in the 6T cell.

For a cell storing 0 ( $V_R = 0$ ,  $V_L = 1$ ), a (non-redundant) write completes once  $V_L$  becomes logic 0.  $V_L$  represents the input to the inverter formed by transistors  $PR-NR$ , the output of which determines  $V_R$ . Therefore, as  $V_L$  becomes logic 0, this inverter forces  $V_R$  to logic 1. Accordingly, to model write timing errors, VARIUS-TC extracts  $D_{VarWriteCell}$ , the time that node  $L$  takes to reach the switching threshold ( $V_{SWITCH}$ ) of the  $PR-NR$  inverter –  $V_{SWITCH}$  in this case corresponds to the maximum voltage to be interpreted as logic 0 by the inverter. In doing this, the architecture module closely follows the methodology from [9]. The device and circuit modules, on the other hand, deviate from [9] due to the utilization of the LUT for the extraction of respective currents. After obtaining the probability distribution for  $D_{VarWriteCell}$ , VARIUS-TC computes the distribution of the maximum of  $D_{VarWriteCell}$  over all the cells in a line,  $D_{VarWriteLine}$ . The probability

$$P[D_{VarWriteLine} > t_{WRITE}]$$

gives the probability of a write timing error, where  $t_{WRITE}$  is the designated write duration.

Relaxing timing constraints (e.g., increasing the designated clock period,  $t_{CLK}$ ) can help mitigate timing errors, however, no such remedy applies to stability errors. For a cell storing 0 ( $V_R = 0$ ,  $V_L = 1$ ), at low  $V_{dd}$ , the voltage  $V_L$  reduces by construction. When the cell is not accessed, if  $V_L$  reduces enough – due to leakage through  $NL$  and  $AXL$ , to reach the  $V_{SWITCH}$  of the  $PR-NR$  inverter, the cell content can easily get distorted. A stability error occurs when the leakage current through the  $NL$  and  $AXL$  transistors in Figure 3(b) reduces  $V_L$  below the  $V_{SWITCH}$  of the  $PR-NR$  inverter while the cell is not being accessed. At that point, the cell's state gets lost.

VARIUS-TC's architecture module captures stability errors following [9]; the key difference from [9] is the adoption of the LUT by the device and circuit modules in calculating the respective currents. VARIUS-TC calculates the stability error probability per cell by

$$P_{Cell,Err} = P[V_L(V_{dd}) - V_{SWITCH}(V_{dd}) < 0].$$

If no redundant cells are available,

$$P_{Line,Err} = 1 - (1 - P_{Cell,Err})^{line\ size}$$

gives the corresponding error probability of a line, where  $line\ size$  denotes the number of cells per line; and  $1 - (1 - P_{Cell,Err})^{line\ size}$ , the probability that at least one cell fails. The

error probability per memory block becomes

$$P_{Mem,Err} = 1 - (1 - P_{Line,Err})^{number\ of\ lines}$$

in this case, with *number of lines* denoting the number of lines in the respective memory block. VARIUS-TC can also calculate the minimum allowable supply voltage to avoid such errors:  $V_{dd_{MIN,Cell}}$  by solving

$$V_L(V_{dd_{MIN,Cell}}) = V_{SWITCH}(V_{dd_{MIN,Cell}})$$

for voltage under variation, where  $V_L(V_{dd_{MIN,Cell}})$ , and  $V_{SWITCH}(V_{dd_{MIN,Cell}})$ , respectively, denote the values of  $V_L$  and  $V_{SWITCH}$  at  $V_{dd} = V_{dd_{MIN,Cell}}$ . VARIUS-TC calculates  $V_{dd_{MIN,Line}}$  by  $\max(V_{dd_{MIN,Cell}})$ .

### 4.3. Capturing Variation in Static Power

VARIUS-TC derives the static power distribution by integrating the static power distribution per switch under variation, over all switches within the encapsulating logic or memory block.  $V_{dd} \times I_{off}$  gives the per switch static power under variation, with  $V_{dd}$  being the operating voltage; and  $I_{off}$ , the leakage current under variation as generated/fetched from the LUT by the device/circuit modules.

## 5. Evaluation Setup

We evaluate VARIUS-TC for different FinFET technology nodes from PTM, as listed in Table 1. As the baseline for comparison, we deploy 16nm PTM high-performance (HP) for planar MOSFET, where the nominal  $V_{th}$ ,  $V_{thNOM}$  is 0.48V, and the nominal effective channel length, 16nm. We set  $\sigma/\mu$  for cumulative variation to 5% which results in  $\approx 3.5\%$  of systematic and random variation under equal share. We consider three different levels of variation – *low*, *medium*, and *high*, with the corresponding  $\sigma/\mu$  of FinFET parameters given in Table 2. To generate systematic variation maps, we set the correlation distance  $\phi$  to 0.1 [17, 9].

Table 1: FinFET technology nodes [15].

Parameter	Unit	PTM 10nm	PTM 16nm	PTM 20nm
$L_{Fin}$	nm	14	20	24
$T_{Fin}$	nm	8	12	15
$H_{Fin}$	nm	21	26	28
$T_{ox}$	nm	12	13.5	14
$\phi_g$	eV	4.42	4.40	4.37
$V_{NOM}$	V	0.75	0.85	0.9

Table 2: Different levels of variation used in evaluation.

Parameter	Low var.	Medium var.	High var.
$L_{Fin}$	3.5%	7%	10.5%
$T_{Fin}$	3.5%	7%	10.5%
$\phi_g$	0.16%	0.32%	0.48%

We repeat each experiment for 100 dies, and report the characteristics for the median die. The chip is organized

in  $4 \times 4$  clusters (Figure 4). Each cluster has 4 in-order Intel Xeon Phi [7] like cores (each of 1GHz and with 32KB L1 private instruction and data cache) and 2MB shared (4-bank) L2 cache. We deploy 8T memory cells. The intra-cluster interconnection network is a bus; the inter-cluster, a 2D torus. We deploy the microarchitectural simulator Sniper-6.0 [3] integrated with the architectural power model McPAT [12] (scaled to 16nm) for running PARSEC-3.0 [2] applications using simsmall input with 64 threads. We experiment with BS:*blackscholes*, BT:*bodytrack*, CN:*cannal*, FR:*ferret*, FA:*fluidanimate*, SC:*streamcluster*, SW:*swaptions*, and XX:*x264*.

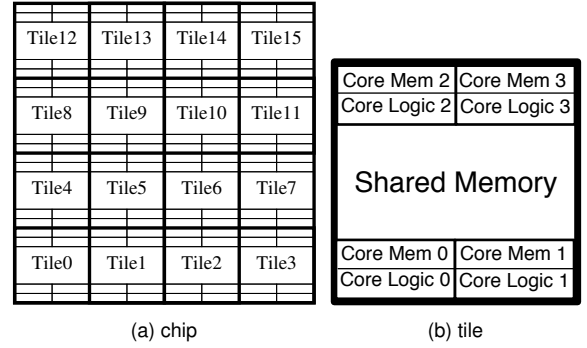


Figure 4: Floorplan of the evaluated manycore architecture.

## 6. Evaluation

### 6.1. FinFET vs. Planar MOSFET under Variation

We first characterize the WID variation in minimum safe operating voltage,  $V_{MIN}$ , and minimum safe clock period at the nominal voltage,  $\tau_{MIN}$  considering 16nm FinFET and 16nm planar MOSFET based designs.  $V_{MIN}$  represents the minimum safe operating voltage for each memory block, which practically prevents the onset of stability errors under variation.  $\tau_{MIN}$ , on the other hand, is the minimum clock period at which each core can safely operate (at the nominal supply voltage) under variation. Operation at  $\tau_{MIN}$  practically prevents the onset of timing errors.

Fig. 5(a) depicts the kernel density estimate for  $V_{MIN}$ ; Fig. 5(b), for  $\tau_{MIN}$ . Kernel density estimates correspond to continuous histograms, with the area under the curve = 1. The nominal supply voltage,  $V_{NOM}$  is 0.85V for FinFET, and 0.7V for planar MOSFET, respectively. We deploy the *high variation* profile from Table 2 for FinFET, and the *low variation* equivalent (where variation in  $L_{Fin}$  corresponds to the variation in effective channel length) for planar MOSFET, to favor the planar MOSFET based design.

Fig. 5(a) captures WID variation for  $V_{MIN}$  across all memory blocks on chip; Fig. 5(b), for  $\tau_{MIN}$  across all cores. The x-axis of Fig. 5(a) depicts the WID  $V_{MIN}$  spread in Volts. We observe that the FinFET based design outperforms its planar counterpart: The planar MOSFET based design renders a maximum  $V_{MIN}$  of 0.59V; the FinFET based, of only 0.46V. At the same time, the  $V_{MIN}$  spread across die remains notably






Figure 5: WID variation of FinFET vs. planar MOSFET.

larger for the planar case. This difference in the spread of the distributions is even more pronounced for  $\tau_{MIN}$ . In Fig. 5(b), the x-axis is normalized to the nominal critical path delay,  $\tau_{NOM}$ , when operating at the nominal operating voltage,  $V_{NOM}$  (were there no variation). The planar MOSFET based design renders a maximum normalized  $\tau_{MIN}$  of 1.412; the FinFET based, of only 1.035. Our results confirm previous studies which report enhanced resilience of FinFET based designs to variation [8].

## 6.2. Impact of Variation Level

We next analyze how WID variation in  $V_{MIN}$  and  $\tau_{MIN}$  changes as a function of the variation in each critical physical FinFET parameter (Section 2.2). As in Section 6.1, we use the 16nm PTM FinFET. Fig. 6 captures the trend for the three different variation profiles – *low*, *medium*, and *high* – from Table 2, following the same format as Fig. 5. We observe that the spread of both  $V_{MIN}$  and  $\tau_{MIN}$  distributions increase as the level of variation increases, to give rise to longer tails. When compared to  $V_{MIN}$ ,  $\tau_{MIN}$  exhibits higher sensitivity to the level of variation.




Figure 6: Impact of different levels of WID variation.

## 6.3. Sensitivity Analysis

**Sensitivity to Physical FinFET Parameters:** The variation in  $L_{Fin}$ ,  $T_{Fin}$ , and  $\Phi_g$  has the highest impact on power and performance [11, 10]. Accordingly, we confine our sensitivity analysis to these three parameters, and consider each in isolation in Fig. 7. We experiment with the 16nm PTM FinFET




Figure 7: Sensitivity to physical FinFET parameters.

under *high* variation (Table 2). In line with previous work [10], we observe that for both,  $V_{MIN}$  and  $\tau_{MIN}$  variation,  $\Phi_g$  represents the critical parameter.  $\Phi_g$  predominantly delivers the worst WID variation profiles, and its impact is higher on the distribution of  $V_{MIN}$ . The maximum value of  $V_{MIN}$  becomes 0.29V under *high* variation for  $L_{Fin}$  (in isolation); 0.25V for  $T_{Fin}$  (in isolation); and 0.42V for  $\Phi_g$  (in isolation), respectively. The maximum value of  $\tau_{MIN}$  becomes 1.019 for  $L_{Fin}$ ; 1.051 for  $T_{Fin}$ ; and 1.108 for  $\Phi_g$ , in isolation.





Figure 8: Sensitivity of WID variation in  $V_{MIN}$  (a) and  $\tau_{MIN}$  (b) to technology scaling, considering PTM model files optimized for high performance (HP).


**Sensitivity to Technology Scaling:** Fig. 8 captures the sensitivity to technology scaling, considering PTM FinFET at 20nm, 16nm, and 10nm under *high* variation (Table 2). We observe that the impact of WID variation increases with technology scaling, for both,  $V_{MIN}$  and  $\tau_{MIN}$ , to render longer tails. Our analysis so far covers PTM nodes optimized for H(igh) P(erformance). PTM L(ow) P(ower) nodes render similar trends, as shown in Fig. 9, when we repeat these experiments.

## 6.4. Impact on Static Power

Fig. 10 characterizes how chip-wide static power changes when compared to the no-variation case, under three different levels of variation (Table 2) for PTM HP FinFET at 16nm. We observe that the static power increases notably, by 2.04 $\times$ , 4.51 $\times$ , and 9.52 $\times$ , under *low*, *medium*, and *high* levels of variation.

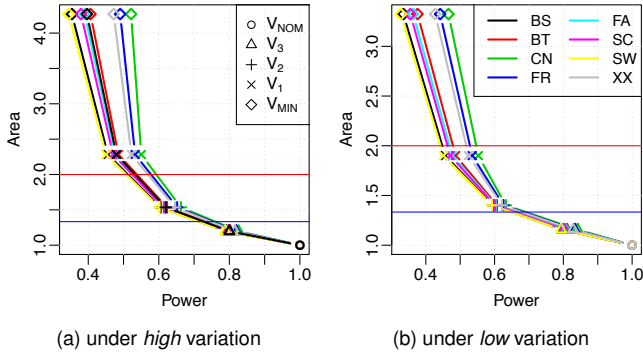


**Figure 9: Sensitivity of WID variation in  $V_{MIN}$  (a) and  $\tau_{MIN}$  (b) to technology scaling, considering PTM model files optimized for low power (LP).**



**Figure 10: Variation in (normalized) static power,  $P_{sta}$ .**

### 6.5. An Example Use Case



**Figure 11: Area vs. power trade-off as a function of the operating voltage for PTM HP FinFET at 16nm.**

Under contemporary technology scaling, one promising way to cram more cores into the available power budget is reducing the operating voltage  $V_{dd}$  aggressively. If  $V_{dd}$  remains slightly above the threshold voltage  $V_{th}$ , power consumption can decrease by more than an order of magnitude [4, 6]. This unconventional regime of operation, Near-Threshold Voltage (NTV) Computing, enables more cores to operate simultaneously. Power savings increase with the proximity of the near-threshold  $V_{dd}$  to  $V_{th}$ . Unfortunately, as  $V_{dd}$  reaches  $V_{th}$ , not only degrades the (minimum) clock period  $\tau_{MIN}$ , hence, the (maximum) operating frequency  $f$  ( $\propto 1/\tau_{MIN}$ ), but also, resilience to variation weakens.

For embarrassingly parallel applications, we can avoid performance degradation by distributing (the same) computation<sup>1</sup> to more cores. While each core operates at the degraded  $f$  at NTV, the execution time can still reduce due to the reduced amount of work per core. At the same time, power savings due to NTV operation exceed the power cost of more cores participating in computation. This solution, however, is only viable if the available chip area can accommodate the higher number of active cores required to mask the  $f$  degradation at NTV.

Fig. 11 captures the area overhead (y-axis) along with the corresponding power consumption (x-axis) as a function of the operating voltage, for the *high* (a) and *low* (b) variation profiles from Table 2, in compensating the  $f$  degradation at NTV by distributing computation to more cores, for the benchmark applications from Section 5. The axes are normalized, respectively, to the area and power of the non-variation-afflicted baseline operating at  $V_{NOM}$ . For both profiles, the slowest of the variation-afflicted cores determines the degraded operating frequency  $f$  at any given operating voltage. Each shape captures a different voltage value, and each trendline, a different benchmark application.

*High* variation renders a higher slowdown than *low* variation, hence demands more active cores (and incurs a higher area overhead) to compensate for the degraded  $f$ . Accordingly, the span of the y-axis is larger in Fig. 11(a). For each profile, we consider three voltage points, uniformly sampled from  $[V_{MIN}, V_{NOM}]$  range. Due to the difference in the variation profiles, (chip-wide)  $V_{MIN}$  values also differ between Fig. 11(a) and (b), and so do the values of the voltage samples,  $V_1$ ,  $V_2$ , and  $V_3$  (in ascending order)<sup>2</sup>. We deploy VARIUS-TC (i) to extract  $V_{MIN}$ ; (ii) to calculate the safe  $f$  and power consumption at each voltage level.

Featuring the maximum power savings at minimum area overhead, the bottom-left corner in Fig. 11 demarcates the desirable operating region. The two horizontal lines show boundaries for area overheads of  $2\times$  and  $1.33\times$ , respectively. For example, if the area overhead has to remain below  $1.33\times$ , we cannot operate at  $V_{MIN}$  – only  $V_3$  becomes feasible for both of the variation profiles. On the other hand, if an area overhead of up to  $2\times$  is acceptable, we can lower the operating voltage to  $V_2$  under *high* variation (Fig. 11(a)), and to  $V_1$  under *low* variation (Fig. 11(b)). Even under *low* variation, operating at  $V_{MIN}$  incurs an area overhead of more than  $3\times$ . Similarly, we can extract the maximum possible power savings at a given area budget from Fig. 11.

## 7. Related Work

Most of the proposed variation models for FinFET do not reach the architecture level [10, 13, 16]. Existing architecture-level variation models such as VARIUS(NTV) [17, 9], on

<sup>1</sup>i.e., by keeping the problem size constant

<sup>2</sup> $V_{NOM} = 0.85V$ ,  $V_3 = 0.75V$ ,  $V_2 = 0.65V$ ,  $V_1 = 0.55V$ , and  $V_{MIN} = 0.44V$  in Fig. 11(a).

the other hand, are tailored for planar MOSFET only. That said, VARIUS-TC represents an extension of VARIUS(NTV) to FinFET. A recent FinFET-based model, FinCANON [11] reaches the architecture level, but focuses on the network on chip and memory, rather than the processor logic. At the same time, FinCANON does not feature any probabilistic model to analyze major variation-triggered error modes, as opposed to VARIUS-TC. Another model that reaches the architecture-level is McPAT-PVT [19]. McPAT-PVT, as well, fails short of providing statistical reliability analysis under variation, as opposed to VARIUS-TC. On the other hand, both FinCANON and McPAT-PVT feature modular macro-models derived from TCAD-based device-level simulations, similar to VARIUS-TC's LUT.

## 8. Conclusion

Due to the enhanced control of the channel, emerging switches such as FinFET can operate more power-efficiently than their classic, planar counterparts. At the same time, they introduce new sources of variation in process parameters, and hence challenge the manufacturing process further. As a result, performance and power of manufactured hardware becomes greatly unpredictable, which can easily impair the power efficiency potential. Unlocking the power efficiency benefits of emerging switches hence mandates characterization of variation-incurred unpredictability across the system stack at early stages of the design. This paper introduces VARIUS-TC, a highly modular architecture-level model, which serves the purpose. VARIUS-TC makes different types of system-level studies under variation possible, including, but not limited to: (1) extraction of a safe (i.e., practically error-free) operating frequency and voltage; (2) generation of critical path delay, power, and  $V_{MIN}$  distributions; (3) calculation of error probabilities for critical logic (timing) and memory (timing and stability) error modes; and (4) design space exploration. VARIUS-TC is available for download at <http://altai.ece.umn.edu/variuss>.

## References

- [1] K. Ahmed and K. Schuegraf, "Transistor Wars," *IEEE Spectrum*, vol. 48, no. 11, 2011.
- [2] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," Princeton University, Tech. Rep. TR-811-08, January 2008.
- [3] T. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the Level of Abstraction for Scalable and Accurate Parallel Multi-core Simulation," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Nov 2011.
- [4] L. Chang, D. J. Frank, R. K. Montoye, S. J. Koester, B. L. Ji, P. W. Coteus, R. H. Dennard, and W. Haensch, "Practical Strategies for Power-Efficient Computing Technologies," *Proceedings of the IEEE*, vol. 98, no. 2, February 2010.
- [5] L. Chang, R. Montoye, Y. Nakamura, K. Batson, R. Eickemeyer, R. Dennard, W. Haensch, and D. Jamsek, "An 8T-SRAM for Variability Tolerance and Low-Voltage Operation in High-Performance Caches," *IEEE Journal of Solid-State Circuits*, no. 4, April 2008.
- [6] R. G. Dreslinski, M. Wiecekowsky, D. Blaauw, D. Sylvester, and T. Mudge, "Near-Threshold Computing: Reclaiming Moore's Law Through Energy Efficient Integrated Circuits," *Proceedings of the IEEE*, vol. 98, no. 2, February 2010.
- [7] J. Jeffers and J. Reinders, *Intel Xeon Phi Coprocessor High-Performance Programming*. Morgan Kaufmann, 2013.
- [8] E. Karl, Y. Wang, Y.-G. Ng, Z. Guo, F. Hamzaoglu, M. Meterelliyo, J. Keane, U. Bhattacharya, K. Zhang, K. Mistry, and M. Bohr, "A 4.6 GHz 162 Mb SRAM Design in 22 nm Tri-Gate CMOS Technology With Integrated Read and Write Assist Circuitry," *IEEE Journal of Solid-State Circuits*, vol. 48, no. 1, Jan 2013.
- [9] U. R. Karpuzcu, K. B. Kolluru, N. S. Kim, and J. Torrellas, "VARIUS-NTV: A Microarchitectural Model to Capture the Increased Sensitivity of Manycores to Process Variations at Near-Threshold Voltages," in *International Conference on Dependable Systems and Networks (DSN)*, June 2012.
- [10] V. B. Kleeberger, H. Graeb, and U. Schlichtmann, "Predicting Future Product Performance: Modeling and Evaluation of Standard Cells in FinFET Technologies," in *Design Automation Conference (DAC)*, June 2013.
- [11] C. Y. Lee and N. K. Jha, "FinCANON: A PVT-Aware Integrated Delay and Power Modeling Framework for FinFET-Based Caches and On-Chip Networks," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, no. 99, 2013.
- [12] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, in *International Symposium on Microarchitecture (MICRO)*, 2009.
- [13] D. Lu, C.-H. Lin, A. Niknejad, and C. Hu, "Compact Modeling of Variation in FinFET SRAM cells," *IEEE Design & Test of Computers*, vol. 27, no. 2, 2010.
- [14] D. Markovic, C. C. Wang, L. P. Alarcon, T.-T. Liu, and J. M. Rabaey, "Ultralow-Power Design in Near-Threshold Region," *Proceedings of the IEEE*, vol. 98, no. 2, February 2010.
- [15] Predictive Technology Model (PTM), <http://ptm.asu.edu/>.
- [16] S. Rasouli, K. Endo, and K. Banerjee, "Variability Analysis of FinFET-based Devices and Circuits Considering Electrical Confinement and Width Quantization," in *International Conference on Computer-Aided Design (ICCAD)*, Nov 2009.
- [17] S. Sarangi, B. Greskamp, R. Teodorescu, J. Nakano, A. Tiwari, and J. Torrellas, "VARIUS: A Model of Process Variation and Resulting Timing Errors for Microarchitects," *IEEE Transactions on Semiconductor Manufacturing*, vol. 21, no. 1, February 2008.
- [18] S. Sinha, G. Yeric, V. Chandra, B. Cline, and Y. Cao, "Exploring sub-20nm FinFET Design with Predictive Technology Models," in *Design Automation Conference (DAC)*, June 2012.
- [19] A. Tang, Y. Yang, C. Y. Lee, and N. K. Jha, "McPAT-PVT: Delay and Power Modeling Framework for FinFET Processor Architectures Under PVT Variations," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 23, no. 9, Sept 2015.