

Evaluating Voltage Islands in CMPs under Process Variations

Abhishek Das, Serkan Ozdemir, Gokhan Memik, and Alok Choudhary

Electrical Engineering and Computer Science Department

Northwestern University, Evanston, IL USA

{ada829, soz463, memik, choudhar}@eecs.northwestern.edu

Abstract

Parameter variations are a major factor causing power-performance asymmetry in chip multiprocessors. In this paper, we analyze the effects of with-in-die (WID) process variations on chip multicore processors and then apply a variable voltage island scheme to minimize power dissipation. Our idea is based on the observation that due to process variations, the critical paths in each core are likely to have a different latencies resulting in core-to-core (C2C) variations. As a result, each core can operate correctly under different supply voltage levels, achieving an optimal power consumption level. Particularly, we analyze voltage islands at different granularities ranging from a single core to a group of cores. We show that the dynamic power consumption can be reduced by up to 36.2% when each core can set its individual supply voltage level. In addition, for most manufacturing technologies, significant power savings can be achieved with only a few voltage islands on the whole chip: a single customized voltage setting can reduce the power consumption by up to 31.5%. Since the nominal operating frequency remains unchanged after the modifications, our scheme incurs no performance overhead.

1. Introduction

As the silicon technology keeps on shrinking to 45nm and beyond, variations in manufacturing process parameters manifest themselves in all forms. These *process variations* can cause *inter-die* or *die-to-die* (D2D) variations, and *intra-die* or *with-in-die* (WID) variations. Moreover, these variations consist of both systematic (correlated) variations and non-systematic (random) variations. In general, process variations can have an unpredictable impact on the power, performance, and reliability of the systems.

Literature surveys [4, 11] reveal that the effect of process variations is profound on the power dissipation and performance of chips. Recent studies show that even in a relatively mature technology like 130nm, these variations are known to result in as much as a 30% decrease in max-

imum frequency and 500% increase in leakage power [9]. For newer technologies, these variations can be even higher: a 20-fold increase in leakage has been reported for 90nm technology [4]. A direct impact of this phenomenon is reduced chip yields. A chip may under-perform or dissipate power greater than a certain threshold and hence may eventually be dropped resulting in effective yield loss. Clearly, another important impact is the variation of power on the remaining chips. In this work, we try to optimize for dynamic power under process variations in chip multiprocessors (CMPs).

CMPs are the latest trend in chip industry. With process technology advancements, multiple cores are laid down on the same die to exploit parallelism in applications and provide higher performance. Just like their single processor counterparts, CMPs are no exceptions with respect to D2D and WID variations. In fact, in CMPs, the problem of parameter variability is more acute because rampant WID variations may result in *core-to-core* (C2C) variations [16]. As a result, the performance of certain cores drop beyond the expected level and a nominal frequency of operation is chosen to be equal to the frequency of the slowest core. Besides, D2D variations also cause chips to differ from each other. Under such circumstances, having a single V_{dd} level for all the manufactured chips is power-inefficient, since there are significant variations between chips manufactured in the same batch. Intuitively, power savings can be achieved by setting a customized V_{dd} for each chip, or a set of cores in a chip thus forming one or more *voltage islands*. In this paper, we try to investigate the impact of process variations on multicore chips and mitigate the power-inefficiencies by using single/multiple voltage islands.

Particularly, we make the following contributions:

- We develop an extensive model, which encompasses process variations for a CMP using statistical estimations and the detailed floorplan for Alpha EV7-like cores.
- We develop a variation-aware scheme for power optimization using single/multiple voltage islands across different cores in a CMP.

- We analyze varying voltage island granularities and show that depending on the technology, even a single voltage island can reduce the power consumption significantly.
- Finally, we formulate an analytical model that can be used to estimate the advantages of voltage islands for different manufacturing processes.

Overall, our results show that the multiple voltage island scheme results in up to 36.2% power reduction in our target architectures. A single voltage island, on the other hand, can save up to 31.5% of the dynamic power.

The remainder of the paper is organized as follow. Section 2 describes the models for power estimation under process variations. In Section 3, we present the multiple voltage island scheme aimed at dynamic power optimization. Section 4 illustrates the experimental results followed by the discussion of related work in Section 5. The paper is concluded with a summary in Section 6.

2. Architecture and Process Models

2.1. Architecture Modeling

Our goal in this study is to evaluate the impact voltage islands in a CMP system, where each core or a set of cores will be able to set its supply voltage such that it will continue correct operation with optimal power consumption. In this approach, we need to examine the impact of changing supply voltage levels on latency of the critical path of the processor. To model the critical path, we have taken into account the 7-stage pipeline of an Alpha-21364 (EV7) processor. The main components of our processor model are the issue queue, the register file, the integer execution units and the memory hierarchy consisting mainly of the L1 data cache. These structures are known to be the critical components in high-performance microprocessors and hence are included in our study, whereas the remaining structures are omitted to reduce the simulation time. The detailed modeling description of the individual microarchitectural units can be found in [8].

To understand the impact of variations on our processor architecture, we have also analyzed the latency distribution of different architectural components under process variations (the variation models are described in the next section). In Figure 1 we plot the latency distributions for different architectural units under process variations. The figure also shows the cumulative latency distribution (i.e., the latency of the entire chip) determined by the longest latency component for each modeled chip. We must highlight that the studied components have equal latency before the introduction of process variations. The results clearly reveal that the latency distribution of the cache dominates the




Figure 1. Latency distribution of each unit for a set of 2000 processors for $\phi = 0.5$. The cumulative distribution corresponds to the critical path of the chip.

cumulative latency distribution of the chips. For a set of 2000 chips we simulated, 58.9% of the critical paths were found to lie in the L1 cache.

The fact that caches are most vulnerable to unit-to-unit variations can be explained by several reasons. First, level 1 caches have a high frequency requirement and consequently tend to utilize low threshold voltages [3]. Second, according to the "FMAX" model introduced by Bowman et al. [5], the number of independent critical paths (N_{cp}) and critical path logic depth (L_{cp}) are two factors determining the criticality of a component. Thus a unit having a high N_{cp} to L_{cp} ratio, will have a larger variance to mean ratio in its delay distribution, leading to an increased susceptibility to process effects. SRAM structures have a high number of critical paths with low logic depths in those paths, making them highly susceptible to process variations. In fact, Humenay et al. [17] have demonstrated that unit-to-unit variations will be dominated by SRAM structures under process variations. Figure 2 depicts the frequency slowdown of several microarchitectural components having different N_{cp} due to random process variations. It is clear that SRAM structures with a high N_{cp} value have a significantly higher chance to contain the critical path of the processor. This analysis shows that the designers can employ voltage scaling techniques that treat the cache separately.

We must note that a change in the supply voltage will not affect the cache only. Regardless of the critical path selected, if the supply voltage of a core can be reduced, the power consumption of the whole core will be reduced.

2.2. Modeling Process Variations

Process variations can be defined as statistical variations in circuit parameters like gate-oxide thickness, channel length and Random Dopant Effects (RDE) due to the shrinking process geometries [4]. Process variations mainly




Figure 2. Frequency slowdown for L1 caches compared to other microarchitectural units (courtesy [17]).

consist of D2D and WID variations. D2D variation refers to the variation in process parameters across dies and wafers, whereas WID variation is the variation in device features within a single die, causing non-uniform characteristics inside a chip. Independent of their type, process variations generally fall into two categories: spatially-correlated variations where devices close to each other have a higher probability of observing a similar variation level, and random variations causing random differences between various devices within a die.

Table 1. Nominal and 3σ variation values for each process parameter modeled

	L_{gate}	V_{th}	W	T	ILD-T/H
Nominal Value	45nm	220mV	$0.25\mu m$	$0.55\mu m$	$0.15\mu m$
3σ -variation	± 10	± 18	± 33	± 33	± 35

To measure the impact of process variations on our processor model, we considered 5 different variation parameters. These are interconnect metal thickness (T), inter-layer dielectric thickness (ILD-T or H), line-width (W) on interconnects, gate length (L_{gate}), and threshold voltage (V_{th}) for the MOS devices. The variation limits for these parameters are given by Nassif [20]. The mean(μ) and 3σ values for each source of variation are listed in Table 1.

We model both spatially correlated and random process variations for our architecture. To take into account the spatial correlation we use a range factor (ϕ) in the 2D layout of the chip. Thus each process parameter x can be expressed as a function of its mean(μ) and variation(σ) and the range(ϕ) values as shown in Equation 1.

$$Param_x = f(\mu_x, \sigma_x, \phi) \quad (1)$$

$$C_i = \frac{\phi}{d_i} - 1 \quad \text{if } d_i \leq \phi \quad ; \quad \text{else } C_i = 0 \quad (2)$$

Equation 2 illustrates the role of the range parameter in determining the correlation coefficient. If in a 2D plane two points x_i and y_i are separated by a distance d_i , then the spatial correlation factor C_i between them can be thought of as an inverse linear function involving ϕ and d_i . Note that there is no correlation between two spatial points which are more than ϕ units apart. With this background, we have generated a spatial map of various parameter values using the R statistical tool [2]. To generate process parameters for a multicore chip, we have replicated the Alpha EV7 floor-plan described in the HotSpot tool [25] to form a 16 core CMP. Figure 3 depicts the procedure for threshold voltage generation under process variations.




Figure 3. V_{th} Map for 16 cores in each chip with (a) $\phi = 0.3$, (b) $\phi = 0.7$ and (c) $\phi = 0.5$. Inset in (c) shows the V_{th} map for a single core.

Random variations caused by RDE mainly manifest themselves as random changes in V_{th} under process variations. Hence to model random variations we use random values from a uniform random distribution to augment the spatially correlated values. The amount of random variations in each parameter can be different and is set according to the results presented in previous work [15]. Compared to spatially-correlated variations, the magnitude of random variations remain lower and generally does not exceed 30% of the overall variations. Note from Figures 3(a) and 3(b) that in our model the random variations depend on the ϕ value. For a small ϕ value (0.3), the parameters are highly random, whereas for a large value of ϕ like 0.7, the param-

eters are highly correlated. A ϕ value of zero would imply totally random variations.

A set of values are generated for each process parameter and are fed into the parameterized SPICE models described in Section 2.1. A batch of 2000 chips are simulated from these models. Thus once the simulation of a chip (precisely the 16 cores within a single chip) is completed, we generate a new set of parameter values corresponding to another chip. This automated procedure accurately simulates the WID variations for each chip. Since we pick the initial parameters from a normal distribution, the effects of D2D variations are also captured.

3. Power-Aware Multiple Voltage Islands

3.1. Methodology

This section presents an overview of the power-aware multiple voltage island scheme for CMPs. When processors are manufactured, they operate at a voltage level set during the design of the processor. This voltage level called the *nominal voltage* is usually chosen at the design time. However, under process variations setting a constant level for all the manufactured chips is considerably inefficient. First, different chips will have different latency slacks, which can be taken advantage of by customizing the voltage level for each chip. In addition, if we consider C2C variations, different cores will tend to have different latencies. In such cases, the operating frequency of the whole chip is determined by the maximum latency across all cores. It is known that the dependency of delay (D) or latency on the supply voltage is given by:

$$D \propto \frac{V_{dd}}{(V_{dd} - V_{th})^\alpha} \quad (3)$$

where V_{th} is the threshold voltage and α is technology constant varying between 1 and 2. Equation 3 implies that cores which have a latency lower than that of the slowest core (nominal delay), can increase their latencies by scaling down the V_{dd} in steps till they reach some minimum value. We refer to this voltage as the *minimum stable* supply voltage (V_{opt}). Beyond this point the circuit operation fails. On the other hand, nominal supply voltage can be defined as the voltage set during design time which gives the desired latency for the set of manufactured chips. Our experiments indicate that for most cases the supply of one or more cores can be reduced below the nominal V_{dd} value. This optimization can significantly cut down static and dynamic power dissipation, hence lowering the energy of the whole system. Thus, in a multicore system a single core or a group of cores can be clustered on the basis of critical latencies and assigned a custom supply voltage. Such clusters with a customized V_{dd} can be referred to as a *voltage islands*.

In general, if there are k voltage islands in a system having a nominal clock frequency of f_{clk} and a corresponding supply voltage V_{dd} , the dynamic and leakage power savings can be denoted by:

$$\Delta P_{dynamic} = C_{load} \cdot f_{clk} \cdot \sum_{k=1}^n (V_{dd}^2 - V_k^2) \quad (4)$$

$$\Delta P_{static} = \sum_{k=1}^n I_{leak_k} \cdot (V_{dd} - V_k) \quad (5)$$

where I_{leak_k} and C_{load} represents the average leakage current and load capacitance for each voltage island. Since our target CMP includes 16 cores, k can take values from 1 to 16. In the former case, the entire multicore system operates on a customized V_{dd} , while in the latter case each core has a different supply.




Figure 4. Various voltage island schemes.

Since voltage islands can contain a single core or a collection of cores, we divide the possible variable voltage islands into the following cases. On one extreme, each core can be allocated an individual V_{dd} and thus the chip will have 16 different voltage islands. On the other extreme, the chip can be assigned a single customized V_{dd} to optimize for power, thus having one voltage island. The other possibilities can be to divide the chip into 2, 4 and 8 islands. Figure 4 depicts some possible voltage island schemes for our 16-core CMP model. Note that for 2 and 8 voltage island configurations, we have two options. The cores can be selected horizontally or vertically to form the islands. For example, for the 2 islands case, the upper 8 and lower 8 cores can form the 2 islands (2-horizontal) or the left and right 8 cores can form the voltage island (2-vertical). In Section 4, we present results with both types of orientations.

One way of implementing the multiple voltage island scheme is to have a configurable DC-DC voltage converter in each voltage island. Once the chip has been tested the voltage levels for each island will be set once; in that way a dynamic adjustment can be avoided. Besides, several software tools allow user-level voltage control to change the supply particularly in mobile processors [1]. This concept can similarly be extended to multicore systems. The extra




Figure 5. Optimum voltage for different voltage islands given as a function (h) of latency.




Figure 6. An example latency distribution curve ($g(l)$).

overhead is going to be in the form of a supply voltage table keeping the voltage specifications for each island or core. The kernel is going to use this table during the boot operation. Alternatively, hardware mechanisms like [12–14] can be easily adapted for this purpose.

3.2. Modeling Power Optimization

In this section, we develop a model that can predict the amount of power savings for a given manufacturing technology. In the core of the model lies the observation that latency and voltage levels are correlated. For example, if a circuit operates at 8ns and our frequency requires 10ns operation, we can reduce the supply voltage until the latency is 10ns. Thus for a particular initial latency value (l), a corresponding minimum stable voltage (henceforth called optimal supply voltage V_{opt}) level exists that guarantees correct operation and results in the minimal power consumption. Since this metric depends on the latency, we first have to extract the relation between the latency and optimal supply voltage: $V_{opt} = h(l)$. Note that this function is circuit-specific. For our target architecture, we have first plotted the latency (l) versus the corresponding V_{opt} values

as shown in Figure 5. Then using curve fitting techniques the function h is found to be:

$$h(l) = \begin{cases} 615 & \text{if } l < 2.18ns \\ 66.8l^2 - 566.8l - 1202.2 & \text{if } 2.18ns \leq l < l_{cutoff} \\ 900 & \text{otherwise} \end{cases} \quad (6)$$

There are two important aspects of function h . First, our analysis of our circuit revealed that it does not work below 615mV (note that the nominal voltage level is 900mV). In addition, function h depends on the cutoff point set by the designer. This cutoff point corresponds to the frequency that the processor will run at and will be set by the designer. Thus using function h we can compute the value of V_{opt} . Since dynamic power is proportional to the square of the supply voltage, from Equation 4 we get dynamic power savings for a chip with latency l as:

$$\Delta P_{dynamic} = 1 - \left[\frac{h(l)}{V_{dd,nominal}} \right]^2 \quad (7)$$

We also need the latency distribution for the batch of chips manufactured to be able to understand the advantages of a voltage island scheme. Assuming that this distribution is Gaussian (Figure 6), we can formulate the probability of a chip having a certain latency by g such that:

$$g(l) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(l-\mu)^2/2\sigma^2} \quad (8)$$

where μ and σ are the mean and standard deviation of the latency(l) distribution. Note that these values can be estimated for a given manufacturing technology. Thus the average dynamic power dissipation (P) for a batch of chips with latency distribution $g(l)$ can be given as:

$$P \propto \int_{l=0}^{l=cutoff} g(l) \cdot [h(l)]^2 \cdot dl \quad (9)$$

Hence, given μ and σ of a distribution, Equation 9 can be used as an analytical model to compute dynamic power consumption with V_{opt} . In Section 4, we show that this model is highly accurate to estimate the optimized power consumption levels for our studied manufacturing technologies. Note that for a different technology, this model can be used by only providing the μ , σ , and cutoff values, which are easily available.

4. Results

In this section we present the power optimization results for different voltage island schemes and also analyze how accurately can the power consumption with voltage islands be predicted using our model. We conducted SPICE simulations on the circuit model described in Section 2. Since

the randomness of a parameter changes with ϕ , we evaluate 7 different voltage islands schemes on 1000 multicore chips each with ϕ values of 0.3, 0.5, and 0.7. The voltage islands can have a granularity of 1, 2, 4, 8, and 16 cores.




Figure 7. Power savings for 1, 2, 4, 8 and 16 voltage islands for (a) $\phi = 0.3$, (b) $\phi = 0.5$ and (c) $\phi = 0.7$.

Figure 7 illustrates the power savings for different voltage island schemes with different amounts of randomness in variation. It shows the percentage of the power saving compared to the processor that uses nominal voltage (900mV) in all its cores. For a highly random case ($\phi = 0.3$), the dynamic power improvement can lie between 13.5% and 35.1%. For the highest correlated variations ($\phi = 0.7$), on the other hand, the improvements range between 31.5% and 36.2%. For the $\phi = 0.5$ model, we see that the dynamic power reduction is between 30.5% and

36.2%. We can reach two important conclusions from these results. First, customized supply voltage levels can be an attractive mean to reduce the power consumption in CMPs under process variations. Particularly, we see that the dynamic power consumption of the chip can be reduced by as much as 36.2% on average, which is achieved when each core is individually controlled (16 voltage islands). Second, depending on the manufacturing technology, even a single customized voltage for the whole chip can reduce the power consumption significantly. Particularly, for the $\phi = 0.5$ and $\phi = 0.7$ models, we see that a single customized voltage level can reduce the power consumption by 30.5% and 31.5%, respectively. Only when the spatial correlation is diminishing ($\phi = 0.3$), we need individual control of the cores: for the $\phi = 0.3$ model, a scheme that uses 16 voltage islands can save 35.1% of the dynamic power while the single voltage island scheme reduces the power consumption by only 13.5%.

Another interesting trend we observe in the results is that voltage islands having the same number of cores have almost same energy savings. For example, 2-vertical and 2-horizontal voltage islands have similar power savings.

Accuracy of the model: We compare the results obtained from the analytical model (Equation 9) with the empirical data from our experiments. The average error in \mathbf{P} for ϕ values 0.3, 0.5, and 0.7 are found to be 0.01%, 0.30%, and 0.44%, respectively. Thus our model gives highly accurate estimations of the dynamic power consumption.

5. Related Work

Parameter variations have lately been a topic of interest in both industry and academia. Power optimizations under process variations have also been studied by several researchers. Previous works have proposed several circuit-level techniques to counter the negative effects of process variations [4, 6, 10]. The inter- and intra-die process variations and their effects on circuit leakage is studied in detail by Rao et al. [23]. In another work, Rao et al. [24] analyze the impact of process variations on circuit leakage and propose methods to reduce them. Most of these techniques focus on analyzing the design statistically or by using static timing analysis, and then modifying the parts of the circuits that are most susceptible to variations. Ozdemir et al. [22] have proposed architectural techniques to improve chip yield under process variation effects. Besides, many gate-sizing strategies have been used on the critical or near critical regions of the circuit in order to reduce the effective latency [9].

Variable Voltage/Frequency Islands (VFI's) have been previously used by other researchers [7, 18, 19]. Marculescu et al. [18] show that VFI-based latency-constrained sys-

tems are more likely to meet timing constraints than Single Clock, Single Frequency (SSV) based systems. In another work, Marculescu et al. [19] have suggested a GALS like architecture with multiple voltage islands for energy awareness under parameter variations. Dhar et al. [12] have designed a controller-based adaptive supply voltage scaling (AVS) mechanism for standard cell ASICs. Niyogi et al. [21] have addressed the issue of using multiple VFIs for energy optimization in media and signal processing applications. These works, although important and showing the advantages of customized voltage islands, do not study the CMPs but concentrate on application-specific processors. In a recent work, Humenay et al. [16] have studied the effects of core-to-core (C2C) variations on power dissipation and yield of chip multicore processors. The authors have investigated the effects of systematic variations on dense and distributed floorplans of a CMP, and used Adaptive Voltage Scaling (AVS) techniques to boost the performance of slow cores. Our work, on the other hand emphasizes on the importance of multiple voltage islands in a CMPs, to reduce power dissipation, and performs a detailed analysis of the advantages for various voltage island formations. To the best of our knowledge, there has been no previous work in analyzing the impact of multiple voltage islands on CMPs under process variation.

6. Conclusion

In this work we analyzed the effects of parameter variations on CMPs with an emphasis on the power dissipation. We presented a variation modeling technique which involves five different variation parameters affected by both systematic and random variations. We have first described an accurate model that can be used to estimate the advantages of forming voltage islands. Our simulations indicate that a custom supply voltage is more effective than a pre-determined nominal V_{dd} for the entire chip. Particularly, application of multiple voltage islands with a latency constraint cuts the power dissipation of CMPs by as much as 36.2%. We also show that for most manufacturing technologies, even a single customized supply voltage for the whole chip can reduce the power consumption substantially.

7. Acknowledgment

This work was supported in part by the National Science Foundation (NSF) under grants IIS-0613568, CCF-0541337, CNS-0551639, IIS-0536994, CNS-0406341, CCF-0621443, CCF-0546278, and NSF/CARP ST-HEC program under grant CCF-0444405, and in part by Intel Corporation.

References

- [1] Notebook Hardware Control, personal edition. <http://www.pbis-167.com/>.
- [2] The R project for statistical computing. <http://www.r-project.org>.
- [3] R. Bai, N.-S. Kim, D. Sylvester, and T. Mudge. Total leakage optimization strategies for multi-level caches. In *ACM Great Lakes symposium on VLSI*, Chicago, Illinois, 2005.
- [4] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De. Parameter variations and impact on circuits and microarchitectures. In *Proc. of the Design Automation Conference*, pages 338–342, Anaheim, CA, 2003.
- [5] K. A. Bowman and J. D. Meindl. Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration. *IEEE Journal of Solid State Electronics*, 37(2), Feb. 2002.
- [6] S. H. Choi, B. C. Paul, and K. Roy. Novel sizing algorithm for yield improvement under process variation in nanometer technology. In *Proc. of the Design Automation Conference*, pages 454–459. San Diego, CA, 2004.
- [7] P. Choudhary and D. Marculescu. Hardware based frequency/voltage control of voltage frequency island systems. In *CODES+ISSS '06: Proceedings of the 4th international conference on Hardware/software codesign and system synthesis*, pages 34–39, New York, NY, USA, 2006. ACM Press.
- [8] A. Das, S. Ozdemir, G. Memik, J. Zambreno, and A. Choudhary. Mitigating the effects of process variations: Architectural approaches for improving batch performance. In *Workshop on Architectural Support for Gigascale Integration (ASGI)*, June 2007.
- [9] A. Datta, S. Bhunia, J. H. Choi, S. Mukhopadhyay, and K. Roy. Speed binning aware design methodology to improve profit under parameter variations. In *Proc. of the Conf. on Asia South Pacific Design Automation*, pages 712–717, Yokohama, Japan, 2006.
- [10] A. Datta, S. Bhunia, S. Mukhopadhyay, and K. Roy. Delay modeling and statistical design of pipelined circuit under process variation. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 2006.
- [11] A. Devgan and S. Nassif. Power variability and its impact on design. In *VLSID '05: Proceedings of the 18th International Conference on VLSI Design held jointly with 4th International Conference on Embedded Systems Design (VLSID'05)*, pages 679–682, Washington, DC, USA, 2005. IEEE Computer Society.
- [12] S. Dhar, D. Maksimovi, and B. Kranzen. Closed-loop adaptive voltage scaling controller for standard-cell asics. In *ISLPED '02: Proceedings of the 2002 international symposium on Low power electronics and design*, pages 103–107, New York, NY, USA, 2002. ACM Press.
- [13] M. Elgebaly and M. Sachdev. Efficient adaptive voltage scaling system through on-chip critical path emulation. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2004.
- [14] D. Ernst, N. S. Kim, S. Das, S. Pant, R. Rao, T. Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner, and T. Mudge. Razor: A low-power pipeline based on circuit-level timing

speculation. In *MICRO 36: Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture*, page 7, Washington, DC, USA, 2003. IEEE Computer Society.

- [15] P. Friedberg, Y. Cao, J. Cain, R. Wang, J. Rabaey, and C. Spanos. Modeling within-die spatial correlation effects for process-design co-optimization. In *Proc. of the Intl. Symposium on Quality of Electronic Design*, pages 516–521, San Jose, CA, 2005.
- [16] E. Humenay, D. Tarjan, and K. Skadron. The impact of systematic process variations on symmetrical performance in chip multi-processors. In *Design, Automation and Test in Europe (DATE)*, April 2007.
- [17] E. Humenay, D. Tarjan, and K. Skadron. Impact of parameter variations on multi-core chips. In *Workshop on Architectural Support for Gigascale Integration*, June 2006.
- [18] D. Marculescu and S. Garg. System-level process-driven variability analysis for single and multiple voltage-frequency island systems. In *International Conference on Computer-Aided Design (ICCAD)*, Nov. 2006.
- [19] D. Marculescu and E. Talpes. Variability and energy awareness: a microarchitecture-level perspective. In *DAC '05: Proceedings of the 42nd annual conference on Design automation*, pages 11–16, New York, NY, USA, 2005. ACM Press.
- [20] S. Nassif. Modeling and analysis of manufacturing variations. In *IEEE Conference on Custom Integrated Circuits*, pages 223–228, San Diego, CA, May 2001.
- [21] K. Niyogi and D. Marculescu. Speed and voltage selection for gals systems based on voltage/frequency islands. In *ASP-DAC '05: Proceedings of the 2005 conference on Asia South Pacific design automation*, pages 292–297, New York, NY, USA, 2005. ACM Press.
- [22] S. Ozdemir, D. Sinha, G. Memik, J. Adams, and H. Zhou. Yield-aware cache architectures. In *International Symposium on Microarchitecture*, Orlando, FL, 2006.
- [23] R. Rao, A. Srivastava, D. Blaauw, and D. Sylvester. Statistical estimation of leakage current considering inter- and intra-die process variation. In *ISLPED '03*, Seoul, Korea, 2003.
- [24] R. R. Rao, D. Blaauw, D. Sylvester, and A. Devgan. Modeling and analysis of parametric yield under power and performance constraints. *IEEE Des. Test*, 22(4):376–385, 2005.
- [25] K. Skadron. HotSpot 3.1. <http://lava.cs.virginia.edu/HotSpot/>.