Modeling and Characterizing Power Variability in Multicore Architectures

Ke Meng Frank Huebbers Russ Joseph Yehea Ismail

Department of Electrical Engineering and Computer Science
Northwestern University
Evanston, IL 60208

Abstract

Parameter variation due to manufacturing error will be an unavoidable consequence of technology scaling in future generations. The impact of random variation in physical factors such as gate length and interconnect spacing will have a profound impact on not only performance of chips, but also their power behavior. While circuit-level techniques such as adaptive body-biasing can help to mitigate mal-fabricated chips, they cannot completely alleviate severe within die variations forecasted for near future designs.

Despite the large impact that power variability will have on future designs, there is a lack of published work that examines architectural implications of this phenomenon. In this work, we develop architecture level models that model power variability due to manufacturing error and examine its influence on multicore designs. We introduce VariPower, a tool for modeling power variability based on an microarchitectural description and floorplan of a chip. In particular, our models are based on layout level SPICE simulations and project power variability for different microarchitectural blocks using statistical analysis. Using VariPower, (1) we characterize power variability for multicore processors, (2) explore application sensitivity to power variability, and (3) examine clustering techniques that can appropriately classify groups of processors and chips that have similar variability characteristics.

1 Introduction

In future technology generations, manufacturing variation will have a profound impact on the reliability, performance, and power consumption of microprocessor designs. Manufacturing deviations due to both systematic fabrication errors as well as random statistical variations affect gate size, dopant concentration, interconnect width, spacing, and thickness. This translates directly to chips that miss critical circuit design targets including latency, power, and resilience to noise. In current designs, foundry induced physical deviations already produce significant die-to-die variation. In particular, industry data for a high-performance processor in a 130nm technology shows that individual dies produced with the same fabrication equipment can have as much as a 30% die-to-die frequency variation and a $20\times$ leakage power variation [7]. ITRS predicts that manufacturing variability will have an increasing prominence in future designs.

The large magnitude of the power variability present in current chips is expected to worsen in future scaled process technologies. This is primarily due to the exponential relationship between transistor gate length and subthreshold leakage current [28] and increasingly intensified leakage power percentage in total power. Consequently, very small deviations in this critical parameter can have detrimental effects on the overall power profile of a chip. Statistical variations in other transistor parameters such as gate width can also have a significant impact on power consumption. Projective studies have shown that physical variations in interconnects will have an increasingly important influence on overall chip performance and will eventually overtake devices as a dominant source of performance variability [26].

The net effect of these manufacturing errors is that components and chips will be increasingly prone to *fabrication induced asymmetry* where physical instantiations of cores, interconnection components, and caches on the same chip may differ widely although they have identical schematic descriptions. In contrast, to *architected asymmetry* [21], which can be artfully constructed to balance power, throughput, latency, and area goals for a target workload, fabrication asymmetry is considerably more nettlesome. The major difficulty is that many of the fundamental characteristics such as circuit power and latency for various microarchitectural structures are no longer constant. They are subject to deviations due to imperfections in the materials and equipment used to fabricate the chip, as well as unavoidable, statistical variance. Furthermore, the Semiconductor Industry Association (SIA) whose forecast anticipates improvements from device/fabrication processes, still paints a grim picture for parameter control in deep submicron technology nodes [29].

Microarchitecture can have a significant impact on parameter variation. Pipeline depth and chip organization can influence the susceptibility of a design to parameter variation [7]. Furthermore, by choosing structures that can be configured on a per instance basis after fabrication and design styles that are more robust to the power, performance, and reliability consequences of parameter variation, designers can mitigate variability. In addition, cooperative strategies that consider both circuit-level implementation and architectural organization are promising because they allow for tradeoffs at many levels of the design. To properly understand these tradeoffs it is imperative that architects have access to models that consider physical structure and capture the relationship between early stage architectural organization decisions and the statistical profile of key design metrics.

1.1 Contributions

In this work, we describe VariPower, a microarchitectural tool for modeling statistical variability in the power consumption of a high-performance microprocessor design. We target power variability as an initial target for architectural manufacturability studies due to the emergence of power as a first class design constraint [29], and the large amount of power variation already seen in commercially available chips [7]. We focus our study on the impact of power variability on a chip multiprocessor (CMP) design which is composed of schematically homogenous cores and caches. Due to within-die parameter variation, these components may have fabrication induced asymmetry in power consumption. Furthermore, we argue that architects will also need high-level strategies for reasoning about statistical variation and classifying types of cores and chips with respect to their variation.

This makes the following principal contributions:

- We develop architectural models for studying probabilistic die-to-die and within-die manufacturing power variations

- From a power perspective, we explore application sensitivity to variability.

- We introduce an automatic approach for classifying representative groups of cores and chips that have considerable parameter variation.

Overall, this work is one of the first to consider architecture-level models for manufacturing variability. In addition, we offer approaches for characterizing power asymmetry due to process variation and illustrate the potential for variation aware management.

1.2 Organization

The remainder of this paper is organized as follows: In Section 2, we describe the chief sources of parameter variation in chip manufacture. We then introduce a model for projecting the statistical profile of a design under power variability in Section 3. In Section 4, we validate our model against detailed SPICE simulation and published data on commercially available chips. In Sections 5 and 6, we present our experimental methodology and a series of case studies that explore power variation in a multicore design. We offer a discussion and comparison to existing work in Section 7, and finally we conclude in Section 8.

2 Background

2.1 Physical Manufacturing Variations

In current fabrication technologies, the physical dimensions of circuit elements, such as gate length and wire width, commonly deviate from their nominal values. This is often a by-product of errors in the any one of many elaborate manufacturing steps used in modern VLSI manufacture. Recently the problem has received significant attention because technology scaling magnifies the slightest errors [6, 7].

Some manufacturing processes, such as lithography and chemical mechanic polishing, are fundamentally more difficult to control with current technology. Consequently, they introduce variation in the physical dimensions of devices and interconnects. Operational characteristics for MOS transistors are heavily determined by relevant physical parameters, such as gate length, gate oxide thickness, and dopant density. Manufacturing steps that influence these physical parameters have a larger bearing on the final result. As transition into the nanoscale era, the problem will worsen [25, 26]. This will have direct impact on both the yield and quality of the final products.

In general, manufacturing errors fall into two categories: systematic and random [15]. Systematic variations can affect the whole lot, wafer, die or portion of the die in a common and repeatable pattern. Most of these errors may be corrected or alleviated if the patterns are observed and corresponding measures, such as optical proximity correction, are taken. On the other hand, random variations are extremely difficult if not impossible to predict and generally require a statistical technique to analyze the problem. Random *die-to-die* variations will affect all the on-die devices in the same way, while random *within-die* variations will produce parameter differences that change on a device-to-device basis on a single die.

In some cases, within-die variations also have spatial correlation patterns [18]. For two transistors on the same die, it has been shown that gate lengths are linearly correlated with distance. This correlation has an important implication: neighboring devices are more likely to share common properties. Consequently, a leaky transistor is likely to be surrounded by other leaky transistors. As a result, regional clusters of leaky transistors can quickly make the microarchitectural unit that they appear in less attractive to use as whole. This is a concept that we explore in our case studies.

2.2 Static Leakage Current Under Variation and Correlation

Static leakage current is primarily composed of three components: subthreshold leakage, gate leakage and substrate leakage. Simulation on sample circuit structures with SPICE and PTM [13] 65nm technology predictive model card identifies subthreshold leakage as the dominate leakage source for near future technologies. The subthreshold leakage of a single

OFF transistor is determined by gate length, gate width, and threshold voltage. Subthreshold leakage has a linear relationship with gate length and has exponentially relationships with gate length and threshold voltage. Due to this exponential dependence on gate length, small deviations in this parameter can introduce leakage current that is significantly larger than the nominal case.

In general, leakage analysis is difficult in an arbitrary logic or memory cell due to the *stack effect* [28]. Under this phenomena, chaining of multiple OFF transistors reduces *both* the leakage variation and leakage current. Different circuit structures may consequently have very different leakage power [11]. Detailed leakage power modeling under the stack effect requires analysis to identify a set of prime inputs and their corresponding probability of appearing. However, one very common building block, the 6-T SRAM cell does not create a stack effect concern because it has no chained transistors.

2.3 Dynamic Circuit Power Under Variation

Dynamic power is consumed by the charging and discharging of internal capacitors contributed by transistors and wiring network. The capacitance of these parasitical capacitors mostly has a linear relationship with the structure dimension. Hence, the deviation in structural dimensions would only cause an approximately linear variation in dynamic power, instead of an exponential one as in the case of static power. As a result, dynamic power variation is significantly smaller than static power variation, a result confirmed by our SPICE simulations in the following sections.

Though the variation is limited, dynamic power is still a primary source of chip power dissipation. For the purposes of completeness and comparison, we continue to include dynamic power in the models that we propose.

3 VariPower: Statistical Power Modeling

The importance of early-stage power estimates to guide architectural design decisions will increase as manufacturing variations become more prominent. At first glance, it would appear that low-level fabrication errors are too far removed from the architectural domain, and those approaches targeted at the circuit level might offer more benefit.

However, architectural design options such as pipeline depth and cache sizing often dictate the overall performance and power characteristics of a processor. Furthermore, they frequently place bounds on which circuits can be used to implement functionality. In addition, pure circuit approaches such as adaptive body-biasing (ABB) [33] have non-negligible overheads and may not be suitable for fine-grain, localized adjustments necessary for within-die variation. Flexible approaches that simultaneously consider both architectural and circuit trade-offs [24], may offer a better chance of reaching an optimal design. Without high-level models that can capture the degree of variation, early stage studies will be difficult to produce.

In this work, we introduce VariPower an architecture-level power variability modeling tool based on circuit-level macros. At the heart of VariPower is a lookup table driven power density estimator that uses SPICE derived scaling factors to model the impact of physical parameter deviations on both dynamic and static power consumption. VariPower models within-die as well as die-to-die power variability via Monte Carlo simulation and can project the probability distributions for power consumption under parameterized architectural models and application usage profiles. This flexibility allows VariPower to predict the severity of fabrication induced power asymmetry for a design and its consequences on different classes of workloads and power management policies.

Figure 1 outlines the flow for the generation of a single sample. An architecture-level description for a chip is fed into VariPower. This description identifies the major components of the design, expresses their area and placement, and characterizes their basic circuit composition. VariPower uses this information to create a high-level description of the chip. Using a series of random numbers which represent parameter variation, and circuit-level models that estimate power, the simulator produces a power profile for the chip.

3.1 Hierarchical Representation of A Physical Design

In fabricated chips, parameter variation is partially dependent on spatial properties of the circuit blocks [18, 27]. Typically, nearby circuits tend to have strong parameter correlations. As the distance between circuits increase, the correlations decrease. Physical geometry of microarchitectural structures will have a significant impact on their statistical profiles. To capture these effects, VariPower models all components of a design in hierarchical floorplan.

In VariPower, a chip-level design is expressed in a *description file* using the `Python` scripting language. Individual cores, caches, and other architectural structures are represented as `Python` objects. The `Python` interface allows components to be represented in terms of physical dimensions: length, width, and placement on chip. To facilitate easy replication and placement of common superstructures, VariPower allows the user to define hierarchical components in a group. Using basic support routines, grouped items can be re-instantiated and stamped down anywhere on chip. This allows for easy representation of tiled architecture as depicted in Figure 2. In this example, a simple processor core is defined in terms of its major subcomponents: caches, router, and execution pipeline. The components are grouped to form a core. The core is then replicated many times to describe a whole chip layout.

In addition to describing physical placement of components, the description file also specifies circuit composition of architectural structures. Each component or physical re-

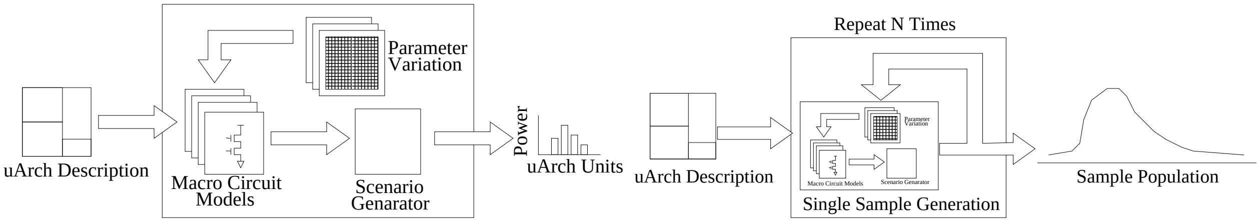

(a) Single Sample Generation

(b) Monte Carlo Approach

Figure 1: Modeling infrastructure used in this work.

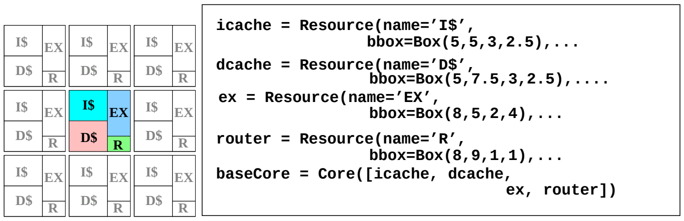

```
icache = Resource(name=’I$’,
                  bbox=Box(5,5,3,2.5),...
dcache = Resource(name=’D$’,
                  bbox=Box(5,7.5,3,2.5),....
 ex = Resource(name=’EX’,
                  bbox=Box(8,5,2,4),...
router = Resource(name=’R’,
                  bbox=Box(8,9,1,1),...
baseCore = Core([icache, dcache,
                                  ex, router])
```

(a)

(b)

Figure 2: Hierarchical models used in VariPower

source, can be comprised of one or more circuit macro blocks which together implement functionality. For example, a cache structure might be composed of a tag array, data array, sense amps, decoders, and drivers. Each of these functional subsections could be represented by its own circuit macro block within the larger cache structure.

3.2 Power Models for Circuit Macros

Circuit macros are the key to modeling statistical power profiles in VariPower. Each macro is a small circuit that is representative of a larger circuit structure. These macros can be thought of as basic building blocks. Just like the full circuits that they represent, these macros have their own corresponding power densities and sensitivity to parameter variations. To capture implementation specific details and the relationship between power and physical parameter deviation, VariPower uses layout-level circuit models to express variability for key classes of circuit structures that are used in a processor. Specifically, VariPower has representative circuit macros for regular array structures, such as cells from register files, cells from cache arrays, as well as slices of an ALU.

With the benefit of layout-level empirical models, VariPower can predict the impact physical variations will have on the power consumption of a design. The layout description for individual macro circuits includes nominal sizing and physical placement of devices and interconnect. Based on this geometric information, VariPower support utilities automatically extract resistances and capacitances for multiple metal layers and polysilicon, as well as gate sizing descriptions for devices. The resistive and capacitive values are determined via analytic models [35]. The electrical components are then output to a SPICE file. We then simulate the macro circuit under a number of different inputs and collect both static and dynamic power.

To model the impact of variation on the macro circuit, we vary physical dimensional parameters for the entire layout description, re-extract electrical component values and reevaluate the circuit under SPICE. For interconnect, VariPower layout level models allow us to directly model the impact that wire width, thickness, and inter-layer dielectrics have on power. For devices, we model gate length, which is known to have an exponential impact on leakage power, as well as gate width. We do not directly model the variational impact on dopant ion concentration or gate oxide thickness and plan to add extensions to model these parameters as part of future work.

In our circuit characterizations we assume that within a circuit, macro dimensional parameters of wires are perfectly correlated, and that dimensional parameters of devices are also perfectly correlated. This simplifying assumption is reasonable because physical parameters of circuit neighboring structures have been shown to be strongly correlated [18, 27]. This is a result of imperfections in a manufacturing step, for example etching, impacting neighboring structures in a similar manner. We construct tables for dynamic and leakage power by varying physical parameters for the macro circuit within the range +/-15% of the nominal value.

3.3 Modeling Parameter Variations

As described in section 2, integrated circuit fabrication processes introduce both systematic and unavoidable random variation in the physical features that comprise transistors and interconnects in high-performance processors. In particular, lithography, etching, and chemical mechanic polishing are all subject to error. Consequentially, physical dimensions such as gate length and wire width will vary for features in circuits, leading to electrical variations in their power behavior.

Each of the physical parameters can be modeled as the sum of die-to-die (D2D) and within-die (WID) variances as:

$$\sigma(x,y)^2 = \sigma_{D2D}^2 + \sigma_{WID}^2 \int_P \rho(x,y)dP \qquad (1)$$

where the overall variance of a parameter is a function of its location on the chip. The global variance is determined by σ^2_{D2D} and it relates the overall deviations that are present across all fabricated dies. The second term expresses the spatial correlation of the physical parameter. Empirical studies have shown that critical parameters can have strong positive correlation for two neighboring points [18, 27]. With this simple, yet flexible model, VariPower can effectively model common types of statistical parameter variation.

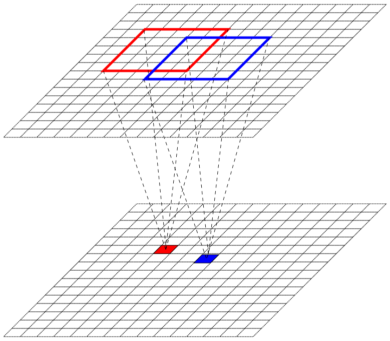

Figure 3: Computing correlated parameter variation using convolution sums.

VariPower introduces a novel scheme for generating virtually any kind of correlated parameter profile. For a given parameter that we wish to model, VariPower generates an $n \times n$ matrix P, which represents how the parameter varies locally on chip. The grid edge length, n, can be tuned for accuracy/speed tradeoffs. We note that by discretizing the chip into parameter domains, we are introducing parameter modeling error which is inversely proportional to the edge length n. This modeling error can manifest itself in one of two ways: (i) circuits in the same grid region may be separated from each other by a maximum distance of $ChipLength/n$ yet be modeled to have the same correlation and (ii) circuits in neighboring grid regions may have a separation distance of less than $ChipLength/n$ and hence have weaker correlations than their true separation would suggest. In our experience, modest values of n give good results. In the case of $n = 1024$, we note that the correlation step sizes are on the order of 0.001 for a linear spatial correlation model. We choose this edge length in our work.

To construct the final parameter matrix, P, VariPower first generates, G, an $n \times n$ matrix of independent Gaussian random numbers. To determine each element, P_{ij}, VariPower sums a subset of elements for G to create virtually any correlation pattern. Figure 3, shows an example where elements in the final matrix, P, will have a correlations in horizontal/vertical directions that decrease linearly and reach zero after a space of three elements. This works because elements in P that are near each other have a large number of items in common. Elements that are far away from each other have none in common. In essence, correlated parameter generation is very similar to 2D convolution kernels used in image manipulation. The procedure for generating these random, correlated values is simple yet flexible. We can essentially change the rate at which the correlation approaches zero by changing the size of the convolution sum. By changing the aspect ratio of the convolution box, we can also change the horizontal and vertical correlation factor. In addition, we can model non-monotonic correlations with more irregular shapes. Furthermore, by allowing the convolution sum to "wrap-around", we can model concave correlation patterns [18].

The last step in generating P is to add a single random number $(\mu = 0, \sigma = \sigma_{D2D})$ to all elements of the matrix. This represents the global sample-wide parameter deviation. The entire process is repeated to generate variation matrices for all physical parameters that VariPower models.

To compute the local dynamic and static power variation of each region of a chip, VariPower overlays the parameter matrix over the hierarchical layout of the design as shown in Figure 3. Within each region, VariPower uses the local parameter matrix values to index the lookup table for the corresponding circuit macro. This allows us to model spatially dependent power densities across the chip.

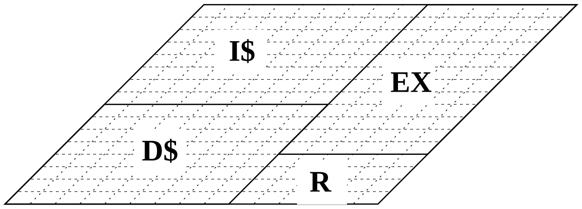

Figure 4: Division of architectural floorplan into regions subject to parameter variation.

3.4 Resource Utilization and Circuit Modes

While process variation affects the maximum power profile for a chip, application utilization patterns and runtime power saving strategies have a significant impact on the typical power consumption of the chip. For example, switching activity factors determine the dynamic power consumption of a pipeline. Dynamic and static power saving strategies such as clock-gating and power-gating transition unused portions of a processor into low-power states. The degree to which they can save power is highly dependent on the activity profiles and performance demands of running applications.

VariPower's scenario generator allows it to examine chip power and its variability under relevant workloads and policies. In particular, VariPower takes an activity profile which captures the hardware usage patterns of benchmark applications. The activity profile can be readily generated by cycle accurate power/performance simulators like Simplescalar/Wattch [10, 9] and expresses the number of cycles spent in active, idle, and low-power states for caches, register files, execution units, and other resources in the processor. VariPower applies application usage profiles on power macro models for each core to generate a cross product of usage patterns and cores. From this cross product it can select entries that represent interesting user-defined scenarios.

VariPower has built-in support for identifying the average power over all possible assignments (asymmetry agnostic), as well as worst-case power (pathologically bad assignment) and best-case power (optimal assignment strategy). Furthermore the Scenario Generation model can be extended to examine more complicated assignment patterns. This allows us to answer questions like, *What is the best assignment under limited knowledge of power asymmetry?* As architects begin to explore the effects of parameter variations, it will be increasingly important to answer these *what-if* questions.

4 Gaining Confidence in Variation Models

For architecture-level power models, the emphasis is traditionally placed on fidelity rather than absolute accuracy. In this way, architectural models can be used to help guide early stage design decisions without the complexity and detail that would be essential under an absolute accuracy requirement. VariPower is designed to produce high-fidelity projections on power variability.

At present, validating VariPower is difficult for two primary reasons. First, there is limited, detailed, published industrial data on parameter variation. In particular, we do not know of any comprehensive data on the power variation of microarchitectural structures. Second, VariPower aims to model the susceptibility of future architectures to variation trends in future technologies. This is a common challenge in the architecture community. In this section, we describe some of our low-level power building blocks and offer some partial validation for VariPower's projections.

4.1 Low Level Circuit Blocks

We choose several basic memory and logic cells as representative structures of the whole processor. At present, we have layout/circuit models for a simple dual bitline SRAM cell, a multi-ported (4r,2w) SRAM cell, a simple CAM cell, an ALU bit slice, and a pipeline latch comparable to one used in the PowerPC 603 [32]. As we continue the development of VariPower, we hope to extend this list to cover more circuits. We anticipate that the use of SPICE simulations with interconnect resistance/capacitance extracted from actual layouts would provide a sufficient accuracy for these blocks. Figure 5 shows the layout for two macro blocks used by VariPower. In the process of assembling these models, we sanity-checked for correct functional operation under the target clock frequency.

VariPower can generate power estimates under variation using two different mechanisms. Under the first mechanism, we directly apply the block level power estimates to calculate absolute power for a given processor model. The benefit of this approach is that the power variations and absolute power are tied to the same underlying circuit-level implementation. Under the second mechanism, we use existing power simulators to form a baseline power estimate and apply VariPower models to project the deviation under parameter variation. In the evaluations in this paper, we apply the later mechanism. At present, VariPower does not have enough representative circuit blocks to provide absolute, overall power projections for an entire processor. We therefore use a slightly modified version of `Wattch` [9] as our baseline and apply VariPower models to project variation.

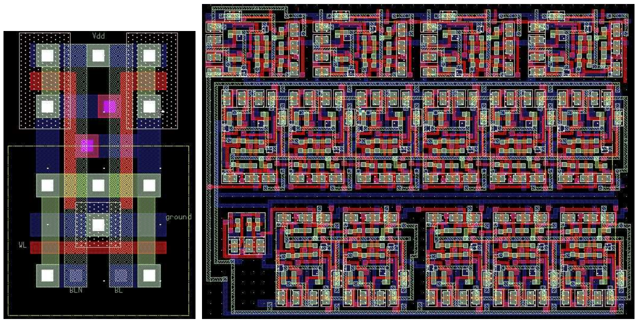

(a) (b)

Figure 5: Circuit layout for (a) SRAM cell used in caches and (b) portion of an adder implemented in dynamic logic.

4.2 Chip-Wide Parameter Deviation and Power Variation

Our first partial validation compares VariPower projections to published measurements from fabricated designs. We first compare our modeled on-chip gate length to those reported by Friedberg et al [18]. In their work, the authors used electronic linewidth metrology (ELM) to capture critical gate dimensions for a 200mm wafer fabricated in a 130nm process. ELM works by passing a known current through gates and measuring the voltage across a section of those gates. Friedberg et al found strong spatial gate length correlations between transistors on the same die. The correlation decreased roughly linearly with distance and leveled off at about half the chip length. We configured VariPower to model a similar correlation profile. Figure 7 shows the resulting gate length correlation. Note that our convolution based parameter generation is capable of producing a close facsimile of the empirical findings in [18]. In Figure 6, we present two samples of a four core CMP modeled using our correlation method. The two chips have very different gate length variation patterns. This underscores the impact that local paramater variation will have on multicore power.

Our second validation examines chip-wide power. VariPower allows us to model both dynamic and static power variation. In the literature, we could not find many reported figures for dynamic power variation. Nassif notes that the impact of manufacturing variation on this topic has not received much attention [15]. One of the benefits of VariPower is its ability to give a comprehensive power

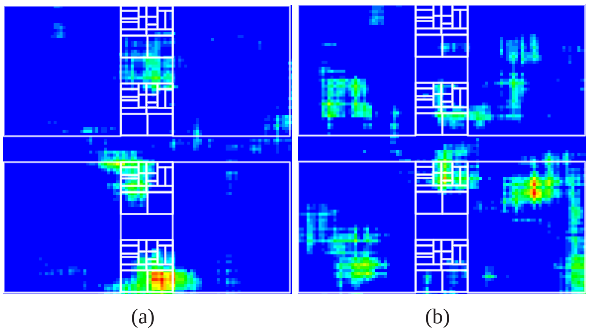

Figure 6: Gate length variation in two chip samples. The chips represent two physical instances of the four core CMP described in Section 5.

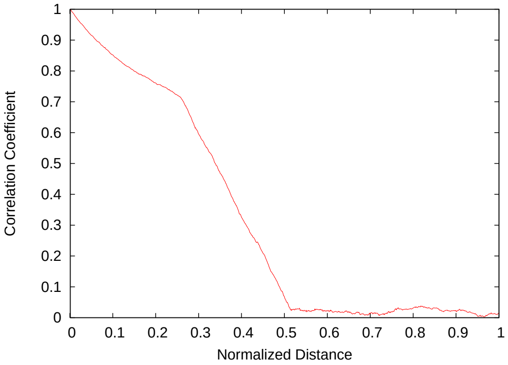

Figure 7: Modeled intra-die variation in gate length.

projection. In Figure 9, we present our estimates of dynamic and static power variation for a four core chip multiprocessor which we describe in detail in Section 5. Our results focus on within-chip parameter variation. We note first that the dynamic power variation is limited in comparison to the static power variation. In addition, the leakage distribution is skewed, with a small number of chips that have very large leakage factors. We also see approximately a $4\times$ variation in leakage power. These results are all comparable to those reported in [4]. However, the relative spread in leakage is much smaller than the $20\times$ variation described in [7]. We still believe that our projections are reasonable because they reflect only within-die parameter variation while the samples studied by Borkar also had substantial die-to-die parameter variation. Die-to-die variation is known to make a major contribution to total variation [8].

5 Experimental Methodology

In this section, we describe the processor model and workloads used in our case studies. While VariPower is capable of modeling the power variation of virtually any CMP configuration, we choose to show a number of different uses of VariPower on a single CMP design.

5.1 Processor Model

Our experiments model power variability and performance of 4-core homogenous chip multiprocessors for a 65nm process. Each core of the processor is comparable to an Alpha 21264 (EV6) scaled to current technology [19]. Under this simple technology scaling, we assume that the processor will not be able to reach the maximum frequency for 65nm, and instead operates at a 3.0GHz frequency. A similar scaling methodology was used by Kumar et al. in [21]. The cores in the processor have private L1 data and instruction caches and private L2 unified caches. Intercore communication and off-chip memory transfers travel across an on-chip bus network. Table 1 summarizes our base processor model.

Our simulation infrastructure is based on a heavily modified version of the M5 Stand-Alone Execution simulator [5] which includes detailed models of pipelines, caches, buses, and off-chip memory. We extend M5 by modeling nominal power under parameter variation as described in Section 3.

Single Core	
Clock Rate	3.0 GHz
Fetch/Decode Width	4 inst
Issue Width	6 inst, out-of-order
IQ/LSQ/ROB	32/40/80 entries
Functional Units	4 IntALU, 1 IntMult/Div 1 FPALU, 1 FPMul/Div 2 MemPorts
L1 Inst Cache	64KB 2-way 64B blocks
L1 Data Cache	64KB 2-way 64B blocks 3 cycle load hit
Chip Multiprocessor	
Cores	4
L2	2MB 16-way private 128B blocks
Off-chip memory latency	200 cycles
Power Parameters	
VDD	1.0V
Clock Rate	3.0GHz
Feature Size	65nm

Table 1: Processor Parameters

5.2 Workloads

To evaluate the efficacy of VariPower, we use several workloads that showcase a variety of hardware usage patterns. Individual applications are taken from the SPEC CPU2000 benchmark suite. To reduce the total number of simulations, we identify a subset of SPEC applications which exhibit a range of power and performance characteristics and then focus our case studies on these benchmarks. Table 2 lists all the benchmarks used in our experiments. To isolate representative simulation windows, we use SimPoint [30] to identify relevant instruction execution intervals for all benchmarks and save checkpoints. Using these checkpoints, we simulate until at least one thread has committed 200 million instructions.

INT benchmarks	FP benchmarks
crafty	swim
eon	mesa
bzip2	equake
twolf	ammp

Table 2: SPEC CPU2000 Benchmarks Used In This Work

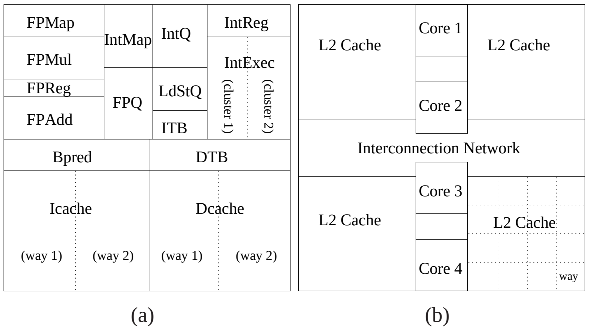

(a)

(b)

Figure 8: Floorplan of Simulated Processor

6 Results

In this section, we conduct a series of case studies using VariPower. These studies serve as examples of the kinds of early stage studies that VariPower can perform.

Figure 8 (a) shows the floorplan of a single core used in these case studies. The floorplan itself is borrowed from Skadron et al ([31]) and is a rough approximantion of an Alpha 21264 processor core. We also base our floorplan for our four-core CMP on work by Kumar ([22]) as shown in Figure 8(b).

In VariPower, Monte Carlo analysis is used to simulate the variations of five process parameters: gate width, gate length, wire length, wire height and inter-wire distance. In this study, we focus on within-die variation, and we include no additional die-to-die variation. For the 65nm predictive technology model [13] used in our SPICE simulations, we assume a 3σ variation of 9% deviation of nominal values for gate width and gate length, and a 3σ variation of 15% deviation of nominal values for the remaining process parameters. The whole chip is divided into a 1024×1024 grid. The devices in the same grid region are assumed to have perfect correlation. Furthermore, correlation between devices in different grid sections linearly drops as the separation increases as illustrated in Figure 7. The Monte Carlo simulations produce 10,000 samples.

6.1 Case Study 1: Core-To-Core Power Variations

As an example of how variability affects microarchitectural structures within a core, we compare the static power of floating point resources in each core of our CMP. We choose floating-point resources because they are not used by integer applications and hence are a likely candidate for leakage management techniques such as standby power modes or power-gating [16, 20]. The insight is that if there is significant variation in power across cores for a given functional unit, we may benefit from selecting specific applications to run on appropriate cores. For example, an application that does not utillize floating point resources could be run on a core with leaky FPUs because any reasonable leakage power management strategy would transition the FPUs into a low-power state. Note: This assumes that there is little or no difference in maximum operating frequency for the chosen core. Based on the effects that a large number of critical paths have circuit delay [8], this is a reasonable assumption. This assumption is confirmed by other high-level models [17].

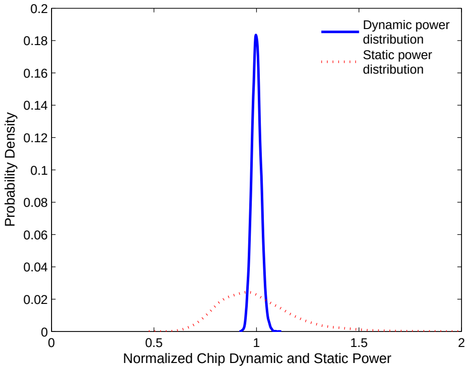

Figure 9: Normalized Chip Dynamic and Static Power Distribution

Floating Point Resources		Normalized Leakage Power Rank (Decreasing Power)			
		1	2	3	4
FPMap	mean	1.160	1.088	1.036	1.000
	stdev/mean	0.1116	0.0733	0.0548	0.0467
FPMul	mean	1.160	1.088	1.037	1.000
	stdev/mean	0.1105	0.0729	0.0552	0.0467
FPReg	mean	1.154	1.079	1.032	1.000
	stdev/mean	0.1258	0.0708	0.0497	0.0412
FPAdd	mean	1.161	1.088	1.037	1.000
	stdev/mean	0.113	0.072	0.0548	0.0469

Table 3: The power distribution of the same microarchitectural structures in different cores.

Table 3 presents the normalized mean leakage for floating point resources across all the cores in our CMP design. For each resource, we rank the structures by decreasing leakage power. We can see overall, that for a given resource type, the leakiest structure is considerably leakier than the least leaky. On average the most power-hungry resource uses 16% more power than the corresponding least power-hungry resource of the same type. This suggests that their might be some opportunity for assigning application threads to cores based on their resource usage and the chip leakage profile. We also note that in general, the ratio of leakage power for cores decreases in the same fashion for the all the functional resource types that we study. What is not evident in the table, is that when a given structure suffered from a higher leakage factor, other structures in the same core did as well. This can be expected due to the strong spatial correlation factors discussed in Section 3.

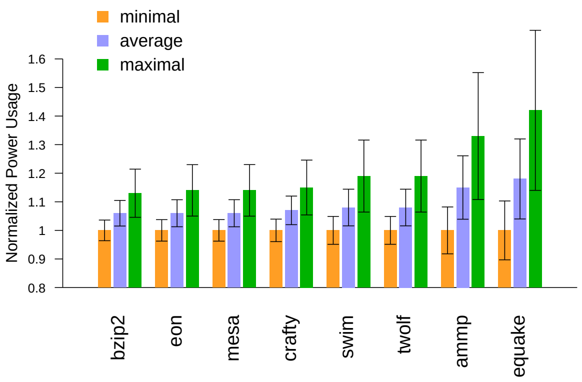

Figure 10: The average power usage of benchmark programs on best, randomly selected, and worst possible cores.

6.2 Case Study 2: Application Sensitivity to Within-Core and Core-To-Core Power Variation

General purpose applications normally have different utilization requirements on macro blocks. For example, some threads have larger memory footprints and require most of the available cache capacity. As mentioned in the previous case study, some applications may not require all of the available execution resources. Typically, these applications are candidates for power reduction strategies that may effectively resize processor components such as caches and queues [12, 3]. Under power variability, each core may have its own leakage profile for a given resource. Consequently, different assignments of threads to cores may yield different power savings when leakage management is applied. This constitutes an opportunity for *core-to-core* savings under leakage asymmetry.

For a given core, there are typically functionally identical structures that may be available to a thread. If the current application does not require all of those resources, there may be a choice of which resources to use and which to transition into a low power standby mode. Selective cache ways is an example of a power saving strategy to which this may apply [3]. Traditionally, structures are considered equivalent from a power savings standpoint. However, under parameter variation, there may be a considerable difference in leakage power for two structures that provide identical functionality. For example, one cache sub-array may be leakier than a neighboring sub-array. This is an example of a *within core* savings under leakage asymmetry.

We used VariPower to model the impact that within-core and core-to-core resource selection can have on power. Figure 10 shows the core leakage power for eight SPEC 2000 benchmarks under different application to core bindings. For each benchmark, the left bar corresponds to the best situation in which the application is assigned to the core that consumes minimal power and the right bar is the opposite scenario which exhibits the worst result. The central bar shows the average power usage when the application is randomly assigned to a core. All the results are normalized with respect to the minimal power of a program. Clearly, different applications can have widely different results. Over all the benchmarks, the minimal assignment saves 6% to 19% over the average (random) assignment. The minimal assignment saves from 13% to 41% over the worst case. The applications in Figure 10 are listed from left and right in increasing order of the static power percentage of the total power. With longer memory stall times, benchmarks like ammp and equake which have lower IPC have lower dynamic power when compared to other programs. As stated in Section 6.1, static power of a microarchitecture structure normally has much larger variation, and it is actually expected that a program using less dynamic power has a better chance to achieve arger power savings, as illustrated in Figure 10.

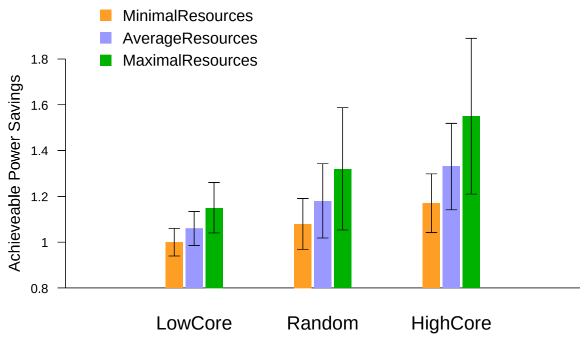

Figure 11: The possible power savings of twolf on different cores by disabling unneeded microarchitectural resources. When resizing caches, we tolerate at most a 2% performance loss.

When combined with other power management mechanisms, in-core power variations also provide power saving opportunities if proper assignment can be made. A focused power control strategy would choose to power-gate the right microarchitectural units to minimize performance loss, obtaining a better power-performance balance. We study such an example using the benchmark twolf in the remaining part of case study 2.

Detailed simulation shows that closing 2 of 16 ways of L2 Cache and half of the L1 Dcache would only cause a 2% performance loss for twolf. In Figure 11 the three color bars show the achieved power savings under three scenarios: (1) selecting the most power-efficient, (2) random and (3) most power-hungry blocks to close when resizing the caches. The three groups from left to right in the Figure correspond to the scenarios in which twolf runs on the most power-efficient core, a randomly selected core and the most power-hungry core. On average, the best selection achieves 12% more power savings over a random selection and 23% over the worst choice. Additionally, we see that the benefit is larger when the within core resource selection is used on a cores that have higher overall power.

6.3 Case Study 3: Automatic Clustering and Binning of Samples

As the impact of manufacturing variability increases, we expect that Monte Carlo simulation methods like VariPower, will become more common in architectural studies. Typically these types of simulations produce large amounts of data in the form of samples. While it is typically straightforward to collect important summary statistics for an individual variable (e.g. mean and standard deviation of leakage power) from the data, it can be considerably more difficult to perform multidimensional data analysis. Rather than organizing the data with respect to a single metric, it becomes necessary to simultaneously arrange the data with respect to many metrics. This is particularly true in the power variability analysis of a CMP because each core or processor component can essentially be thought of as its own dimension.

One likely application of multidimensional analysis is *multicore binning*, which can be thought of as a multicore extension of traditional uniprocessor binning. Under multicore binning, our goal is to identify a set of chip instances that have similar core-level profiles. In the case of power variability studies that we explore in this paper, we want to identify groups of chips that appear the same with respect to their core power consumption under variation. This knowledge could be used to partition a sample space for further study. In essence, multicore binning can be thought of as a way to apply some order to the mountain of data that emerges from Monte Carlo simulation.

In this case study, we explore the use of *clustering* a statistical data mining approach that groups and organizes multidimensional data. In particular, we apply the k-means algorithm to analyze the leakage profile of our four core CMP sample space. Previous work has examined clustering techniques to identify program phases [30]. To our knowledge, we are the first to propose a machine learning algorithm to analyze hardware projections.

For a given N, the k-means algorithm groups the given multidimensional data into N clusters. Each cluster contains a collection of data items that share some similarity, usually measured by a distance function (e.g. Manhattan distance or Euclidean distance). We applied k-means clustering to identify chips that have similar core leakage profiles, using Euclidean distance as a similarity criteria. For each chip in our sample population, we first sort the core leakage values in ascending order. This allows us to compare the cores from different chips using a consistent rank. We explored the benefit of clustering for three values of N: 3, 5, and 7.

Figure 12 summarizes the clustering results by presenting a centroid for each cluster. For N=3, there are two large clusters which comprise almost 90% of the population. The third cluster which represents 9.95% of the population is distinguished by a much higher overall leakage Figure (reaching 22W), and features a very large leakage value for one of its cores.

For N=5, the centroid with the largest total leakage values also represents the smallest fraction of the population just as before. But in this case, the representative 4.25% of the sample population are probably outliers that may be disregarded in analysis. Clusters 3 and 4 offer a particularly interesting view of the sample space. These two clusters have similar total leakage power values, but there is a noticeable difference in the way this power is divided. This highlights an advantage of clustering that could not be identified if we had opted to just characterize with respect to total leakage.

Finally, for N=7, the smallest bin, which again has the largest power total, only represents 1.75% of the population. The six remaining graphs show considerable absolute range (12.8W - 24.8W) and have very different core power distributions.

We can see from this small example, that clustering has the potential to help architects group multidimensional data. As part of our future work, we plan to investigate clustering techniques to further analyze data collected from architectural simulations.

7 Discussion and Related Work

Process variations and its impact on system performance and reliability have gained much attention in the research community in recent years. Borkar *et al.* in [7] discuss common parameter variations observed in today's industry and their impact on circuit and microarchitecture. This work also describes current challenges at the circuit level and offers opportunities for architects to help.

In an effort to better understand and describe the underlying physical mechanisms behind parameter variation, recent work has examined use of statistical models. These representations capture observable circuit characteristics (such as leakage power and maximum clock frequency) given the variation of the underlying technology parameters (such as transistor channel length and oxide thickness). In [28], Rao *et al.* developed a model to estimate the variation of chip leakage current due to gate length process variation. In [14], the authors established a similar model with additional considerations on oxide thickness variability and process parameter correlations. In [1], random dopant fluctuation is further included in estimating the leakage variation. Bowman *et al.* [8] developed a model describing the maximum clock frequency distribution of processors. This model was demonstrated to be extremely accurate when compared with wafer sort data.

Recent research [18, 34] has taken a closer look at the calibrating models against real, fabricated chips. The authors in [18] physically measured the critical dimensions on an industrial processed wafer using ELM and successfully observed the strong correlation of gate lengths. In [34] the authors implemented special testing structures and electronically measured leakage currents. As work in this area continues, we will benefit from higher fidelity parameter variation models.

While much progress has been made on modeling and ad-

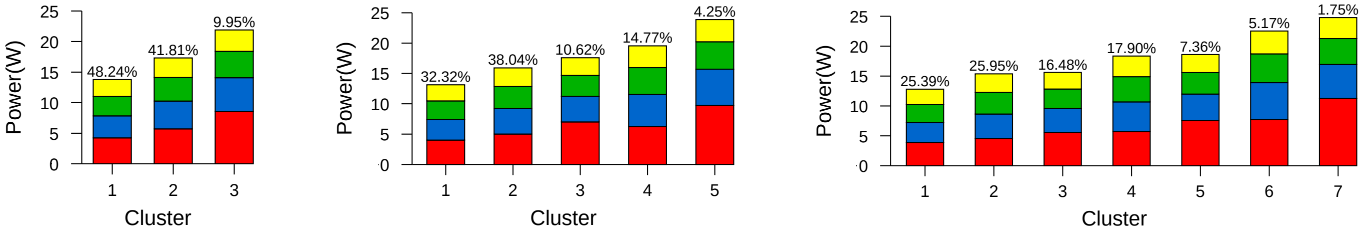

(a) N=3

(b) N=5

(c) N=7

Figure 12: Core leakage power binning using k-means clustering. N denotes the number of clusters. Bars represent core and chip leakage power for the centroid of each cluster. Percentages represent the size of the cluster relative to the entire sample population (10,000 chips).

dressing variation problems at both the device and circuit levels, microarchitects are only begining to examine the problem. Humenay *et al.* develop a model for power and performance variability for mulitcore chips [17]. The major differences between their power model and ours is that we build on SPICE level macro blocks, and we also model interconnect related variations and dynamic power. In addition, we have augmented VariPower with a very flexible model for modeling correlated parameter variation and a scenario generator that allows us to easily answer a number of “What if?” questions. Marculescu and Talpes in [23] propose a joint performance, power and variability metric design method considering the statistical uncertainty at microarchitecture level due to gate length and temperature variations. In [2] the authors establish a model describing the failure distribution of on-chip caches under severe process variations. They propose a fault-tolerant organization that improves the yield.

Our work in this paper is the one of the first to look at variation aware microarchitectural power modeling using physically grounded analysis. We have tried to include as many of the factors influencing manufacturing variation as we could for early stage power analysis. In particular, we have explored the impact that both gate sizing and interconnect variation have on dynamic and static power. In our future work, we plan to also examine the impact of variations in dopant concentration and gate oxide thickness. Low-level circuit macros form the basis for VariPower’s power estimates. As we develop this tool, we plan to add more layout level circuit macros to provide a better match for microarchitectural structures. Finally, we plan to continue our model validation and refine VariPower as industry data becomes publicly available. We anticipate that in the coming years, interest in parameter variaton will grow within the architecture community.

8 Conclusion

In this work, we propose VariPower, an architecture-level model and toolset for evaluating power variability due to manufacturing defects. VariPower uses empirical, layout-level models that capture the effects that physical parameter variation have on both dynamic and static power. We provide a partial validation of our model against published results. Finally, we provide a series of case studies that explore the potential for power variability analysis at the microarchitecture level.

Acknowledgments

We would like to thank the anonymous reviewers for their constructive feedback and helpful suggestions. This work was supported in part by NSF award CCF-0541337.

References

[1] A. Agarwal, K. Kang, and K. Roy. Accurate estimation and modeling of total chip leakage considering inter- and intra-die process variations. In *ICCAD*, 2005.

[2] A. Agarwal, B. C. Paul, S. Mukhopadhyay, and K. Roy. Process variation in embedded memories: Failure analysis and variation aware architecture. *IEEE Journal of Solid State Circuits*, 40:1804–1813, Sep. 2005.

[3] D. H. Albonesi. Selective cache ways: On-demand cache resource allocation. In *International Symposium on Microarchitecture*, pages 248–, 1999.

[4] M. Ashouei, A. Chatterjee, A. D. Singh, V. De, and T. Mak. Statistical estimation of correlated leakage power variation and its application to leakage-aware design. In *Proc. of the 19th VLSID*, 2006.

[5] N. L. Binkert, E. G. Hallnor, and S. K. Reinhardt. Network-oriented full-system simulation using m5. In *Sixth Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW)*, February 2003.

[6] S. Borkar. Design challenges of technology scaling. *IEEE Micro*, 19(4):23–29, July 1999.

[7] S. Borkar et al. Parameter variations and impact on circuits and microarchitecture. In *Proceedings of the 40th Design Automation Conference (DAC-40)*, 2003.

[8] K. A. Bowman, S. G. Duvall, 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 Circuits*, 37:183–190, Feb. 2002.

[9] D. Brooks, V. Tiwari, and M. Martonosi. Wattch: A framework for architectural-level power analysis and optimizations. In *Proceedings of the 27th International Symposium on Computer Architecture (ISCA-27)*, June 2000.

[10] D. Burger and T. M. Austin. The SimpleScalar Tool Set, Version 2.0. *Computer Architecture News*, pages 13–25, June 1997.

[11] J. A. Butts and G. S. Sohi. A static power model for architects. In *Proceedings of the 33rd International Symposium on Microarchitecture (MICRO-33)*, December 2000.

[12] A. Buyuktosunoglu, T. Karkhanis, D. H. Albonesi, and P. Bose. Energy efficient co-adaptive instruction fetch and issue. In *Proceedings of 30th International Symposium on Computer Architecture (ISCA-30)*, May 2003.

[13] Y. Cao, T. Sato, D. Sylvester, M. Orshansky, and C. Hu. New paradigm of predictive mosfet and interconnect modeling for early circuit design. In *Proc. of CICC*, pages 201–204, 2000. http://www.eas.asu.edu/p̃tm.

[14] H. Chang and S. S. Sapatnekar. Full-chip analysis of leakage power under process variations, including spatial correlations. In *Proceedings of DAC*, 2005.

[15] A. Devgan and S. Nassif. Power variability and its impact on design. In *Proceedings of the 18th International Conference on VLSI Design (VLSID-05)*, 2005.

[16] S. Dropsho, V. Kursun, D. Albonesi, S. Dwarkadas, and E. Friedman. Managing static leakage energy in microprocessor functional units. In *The 35th International Symposium on Microarchitecture (MICRO-35)*, 2002.

[17] K. S. Eric Humenay, David Tarjan. Impact of parameter variations on multicore chips. In *Workshop on Architectural Support for Gigascale Integration 2006 (ASGI in conjunction with ISCA-33)*.

[18] 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 6th Int. Symp. on Quality Electronic Design*, 2005.

[19] L. Gwennap. Digital 21264 sets new standard. *Microprocessor Report*, pages 11–16, Oct. 28, 1996.

[20] Z. Hu et al. Microarchitectural techniques for power gating of execution units. In *The International Symposium on Low Power Electronics and Design (ISLPED)*, July 2004.

[21] R. Kumar, K. Farkas, N. P. Jouppi, P. Ranganathan, and D. M. Tullsen. Single-isa heterogeneous multi-core architectures: The potential for processor power reduction. In *Proceedings of the 36th International Symposium on Microarchitecture (MICRO-36)*, December 2003.

[22] R. Kumar, V. Zyuban, and D. M. Tullsen. Interconnections in multi-core architectures: Understanding mechanisms, overheads and scaling. In *Proc. of the 32nd ISCA*, 2005.

[23] D. Marculescu and E. Talpes. Energy awareness and uncertainty in microarchitecture-level design. *IEEE Micro*, 25:64–76, Sept.-Oct. 2005.

[24] S. M. Martin, K. Flautner, T. Mudge, and D. Blaauw. Combined dynamic voltage scaling and adaptive body biasing for lower power microprocessors under dyanmic workloads. In *Proc. of ICCAD*, 2002.

[25] S. R. Nassif. Modeling and forecasting of manufacturing variations. In *Proceeding of the Fifth International Workshop Statistical Metrology*, 2000.

[26] S. R. Nassif. Modeling and analysis of manufacturing variations. In *Proceedings of the 18th International Conference on VLSI Design (VLSID-05)*, 2005.

[27] M. Orshansky, L. Milor, P. Chen, K. Keutzer, and C. Hu. Impact of spatial intrachip gate length variability on the performance of high-speed digital circuits. *IEEE Transactions on Computer Aided Design of Intergrated Circuits and Systems*, 21(5):544–553, May 2002.

[28] R. Rao, A. Srivastava, D. Blaauw, and D. Sylvester. Statistical analysis of subthreshold leakage current for vlsi circuits. *IEEE Trans. on VLSI Systems*, 12:131–139, Feb. 2004.

[29] Semiconductor Industry Association. International Technology Roadmap for Semiconductors, 2005. http://public.itrs.net/Files/2005ITRS/Home.htm.

[30] T. Sherwood, E. Perelman, G. Hamerly, and B. Calder. Automatically characterizing large scale program behavior. In *Proceedings Tenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Oct 2002.

[31] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-aware microarchitecture. In *Proceedings of the 30th International Symposium on Computer Architecture (ISCA-30)*, June 2003.

[32] V. Stojanovic and V. Oklobdzija. Comparative analysis of master-slave latches and flip-flops for high-performance and low-power systems. *IEEE Journal Solid-State Circuits*, 34(4):536–548, April 1999.

[33] J. W. Tschanz et al. Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage. *IEEE Journal of Solid-State Circuits*, 37(11), November 2002.

[34] R. Venkatraman, R. Castagnetti, and S. Ramesh. The statistics of device variations and its impact on sram bitcell performance, leakage and stability. In *ISQED*, 2006.

[35] S.-C. Wong, G.-Y. Lee, and D.-J. Ma. Modeling of interconnect capacitance, delay, and crosstalk in vlsi. *Transactions on Semiconductor Manufacturing*, 13(1), 2000.