

Power Signal Processing: A New Perspective for Power Analysis and Optimization

Quming Zhou, Lin Zhong and Kartik Mohanram
Department of Electrical and Computer Engineering
Rice University, Houston, TX 77005
{quming, lzhong, kmram}@rice.edu

Abstract

To address the productivity bottlenecks in power analysis and optimization of modern systems, we propose to treat power as a signal and leverage the rich set of signal processing techniques. We first investigate the power signal properties of digital systems and analyze their limitations. We then study signal processing techniques to detect temporal and structural correlations of power signals. Finally, we employ these techniques to accelerate the simulation of an architecture-level power simulator. Our experiments with the SPEC2000 benchmark suite show that it is possible to accelerate power simulation by 100X without introducing significant errors at various resolution levels.

Categories and Subject Descriptors

J.6 [Computer-Aided Engineering]: Computer-aided design.

General Terms

Algorithms, Design.

Keywords

Power, Signal Processing, Power Simulation, Power Analysis.

1. Introduction

We have seen two designer productivity challenges to power optimization of a large electronic system, be it a system-on-a-chip (SoC), a system-in-a-package (SiP), or a complete computer system. First, techniques based on average power estimation are inadequate to identify and subsequently minimize system behavior that consumes high power. Moreover, a detailed and dynamic power trace covering a relatively long runtime is important to validate a system for performance and thermal management. Unfortunately, cycle-accurate power simulation of a large system for millions of cycles is notoriously slow [1]. For example, it takes about one hour to simulate only 4000 cycles for the SPE unit on the IBM CELL processor [2]. There is a need for techniques that improve the performance of power simulation tools and, ideally, minimally compromise accuracy.

This research was supported in part by a grant from Texas Instruments.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ISLPED'07, August 27–29, 2007, Portland, Oregon, USA.

Copyright 2007 ACM 978-1-59593-709-4/07/0008 ...\$5.00.

Second, power simulation or measurement of large electronic systems can produce a massive amount of data. Such data contains important information for design optimization and validation. Unfortunately, it is extremely hard and counter-productive for a designer to manually examine and analyze this data. Moreover, visual presentation and interactive manipulation of such massive data are also challenging. There is a need for tools to identify suspicious power behavior from massive power data and, ideally, suggest ways to improve it.

This paper describes a signal processing approach to address these two challenges. We treat power consumption of an electronic system as a digital signal and treat that of its components as a multi-dimensional signal or distributed signals. A component can be a gate, ALU, processor core, or even an entire chip on a printed-circuit board. Based on this, we explore advanced signal processing and pattern analysis techniques to study the power signal. We call this *power signal processing*. Whereas signal processing techniques, such as Fourier and Wavelet analysis, have been used for micro-architecture performance [3] and supply voltage analysis [4], they have not yet been applied to study and optimize power behavior. In this work, we

- study the properties of power signals,
- propose effective and efficient algorithms to detect *temporal* and *structural* correlations in power signals, and
- investigate the application of power signal processing to accelerate power simulation.

We believe that power signal processing introduces a new perspective into power analysis and optimization. Our experiments with the SPEC2000 benchmark suite show that it is possible to accelerate power simulation by 100X without introducing significant errors at various resolution levels. This work is thus an initial step toward utilizing the extremely rich collection of techniques from the signal processing and pattern analysis research communities to enhance power analysis and optimization of digital systems.

This paper is organized as follows. Section 2 describes preliminaries for power signal processing. Section 3 describes signal processing techniques that are relevant to power analysis and optimization, and techniques that make new discoveries regarding power behavior. Section 4 focuses on power simulation acceleration. Section 5 presents experimental results. Section 6 is a conclusion.

2. Power as a signal

We first provide necessary background and motivation for power signal processing and also address the unique properties of power signals.

2.1 Estimated and measured power signals

Dynamic power traces can be obtained by cycle-accurate power estimation or by direct power measurement. Techniques for cycle-accurate power estimation at various levels of abstraction have been widely used in industry [2]. The tradeoff in cycle-accurate power estimation is between speed and accuracy across abstraction levels. The most accurate estimation is running a SPICE-like simulator on a transistor-level netlist, which is too slow to be practical for large circuits. Register-transfer level power estimation can produce relatively accurate traces but still suffers from slow speed [1]. Many techniques to accelerate cycle-accurate power estimation have been studied in literature, e.g., [1, 5, 6]. On the other hand, however, architectural level power simulators for microprocessors [7–9] can be fast, but they lack the accuracy required to guide clock gating at the RTL level [2]. To achieve both high speed and high accuracy, power signal processing seeks to realize multi-resolution power estimation, i.e., to run architecture-level estimation as much as possible while selectively applying gate-level estimation only over interesting cycles.

Cycle-accurate power estimation is, however, limited in the accuracy to reflect the power dynamics of the real system. Most estimation technologies and simulators are memoryless, meaning that power consumption in each cycle and in each component is calculated independently. In a real system, decoupling capacitors, parasitic capacitance, and even by-pass capacitors make this untrue. Their net effect on the system power behavior is similar to a low-pass filter. As a result, truly cycle-accurate power estimation is indeed *not* an accurate reflection of system power behavior or necessary for power analysis unless the designer wants to study each and every cycle. Power signal analysis leverages this and enables the designer to examine power behavior at different resolutions rapidly.

The other way to obtain dynamic power traces is direct power measurement. While direct power measurement offers absolute accuracy, it is limited in both temporal and structural resolution. As mentioned above, due to the existence of decoupling capacitors, parasitic capacitance, and by-pass capacitance, the power consumption over a cycle or of a component is affected by its temporal or spatial neighbors. The power trace obtained through measurement, though accurate, is unable to offer the highest, i.e., cycle-by-cycle or component-by-component resolution.




Figure 1: Second-order RLC model for power-supply network

To further examine the inherent uncertainty in power signals introduced by decoupling capacitance and by-pass capacitors, we model the power-supply network of an electronic system with a second-order resistive, inductive, and capacitive (RLC) circuit as shown in Figure 1. In the model, the resistor represents the resistance of the power-supply network; the inductor represents parasitic inductance, e.g., that introduced by chip-die connectors [10]; and the capacitor represents parasitic capacitance and on-die decoupling capacitance to curb abnormalities in the power-supply network. The current drawn by the system can be represented by a current source, I_{chip} . Since I_{chip} is not directly observable to power

measurement, power measurement instead documents I_{measure} . However, the power-supply network suppresses most of the temporal dynamics in I_{chip} so that I_{measure} will be at most the low-pass filtered I_{chip} . When I_{chip} is spectrally steady, the RLC circuit behaves as a low-pass filter. For example, we use parameters for a high-performance processor with a 1GHz clock [10]: $R = 500\mu\Omega$, chip-die connector inductance $L = 0.005nH$, and on-die decoupling capacitor $C = 500nF$. The circuit model for the power-supply network has a resonant frequency of 100MHz given by $\frac{1}{2\pi\sqrt{LC}}$ [11].

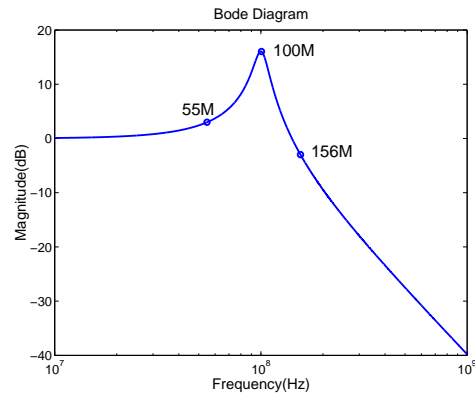


Figure 2: Bode diagram of the power-supply network

The frequency response of the chip is shown in Figure 2. The -3dB cutoff frequency is 156MHz, 56% higher than the resonant frequency. Any harmonic frequency of I_{chip} greater than 156MHz will be attenuated. As shown in Figure 2, the magnitude of 1GHz frequency will be reduced to 1% of the original value.

When I_{chip} is not spectrally steady, the power-supply network will further impact the accuracy of I_{measure} when the RLC circuit takes time to enter a new steady state. Therefore, the power-supply network will attenuate the frequency components in I_{chip} that are higher than the resonant frequency. Hence, the frequency components higher than the resonant frequency in I_{measure} will not accurately reflect those in I_{chip} . In other words, *a sampling rate much higher than the resonant frequency will not produce a power signal with more reliable temporal dynamics*. SPICE is used to simulate the circuit in Figure 1 with I_{chip} running at 1GHz with a triangular shape [12], which is higher than the resonant frequency of 100MHz. Figure 3 presents the plots for both I_{chip} and I_{measure} . The current I_{measure} is heavily modulated by the power supply circuit as shown by its fluctuating waveform. An error will occur if the current is directly measured to estimate the cycle-accurate power. The waveform stabilizes after 70 cycles in the figure, which implies that the measurable current is an average value for at least 70 cycles.

In summary, the power-supply network significantly limits the temporal dynamics that power measurement can capture. As a side effect, it also suppresses security attacks based on power analysis [13]. As long as security-sensitive behavior occurs at a higher frequency than the -3dB cutoff frequency or the resonant frequency, direct power measurement will be unlikely to uncover it.

2.2 Properties of power signals

The rationale behind our proposed approach is that power traces obtained through simulation and measurement can be naturally treated as time-discrete signals, or power signals. Moreover, power signals exhibit many properties that are amenable to digital signal processing.

To illustrate the properties of a power signal, we use a cycle-accurate power trace generated by an industry RTL power simula-




Figure 3: Cycle-accurate current (power) at 1GHz: the ringing of the measured current I_{measure} disallows a cycle-accurate measurement.

tion for an HDTV ASIC module [1] as an example. Part of the trace is shown in Figure 4. The figure also shows power contributed by three different types of data path units: functional units, multiplexers, and registers. Power traces typically have rich periodicity, as is apparent from Figure 4. Knowing the periodicity of a power trace, it is possible to recover or synthesize a power trace that approximates the original one, and potentially accelerate power simulation significantly. Figure 4 also shows that the power consumption by multiplexers and functional units are highly related. Knowing such structural relations among components, we can significantly speed up power simulation by skipping the simulation for either multiplexers or functional units.






Figure 4: Cycle-accurate power traces generated from RTL-level simulation exhibit periodicity and correlations.

Figure 5(a) is a power trace of a smart-phone measured at 10K samples/s, when it is playing a video clip using Windows Media Player. From the figure, it is clear that there is a repetitive pattern every 670 cycles. It corresponds to a frequency of 15Hz ($10K/670 = 15$), the number of video frames per second. The frame rate can also be visualized in the frequency domain. Figure 5(b) gives the time-frequency characteristics of the power trace, which reveals a strong frequency component at 15Hz. Additionally, that the dominant frequency of 15Hz is quite stable across the whole trace supports the periodicity of 670 cycles in the trace.

The highly predictable power trace is essentially correlated with the executed program. For example, loops in the algorithmic specification of a system create frequency components in the power trace. Nested loops create co-existing frequency components. Moreover, finer power behavior revealed under high temporal resolution is usually introduced by lower level design features. Through



(a) Power signal: the periodicity is 670 cycles



(b) Time-spectrum of power signal: prominent energy at 15Hz

Figure 5: Power signal of a smart-phone playing a video at 15 frames/s and its spectrum: the sampling rate is 10K/s.

power signal analysis and processing, we can relate power behavior with design features, and identify sources that introduce undesirable power behavior. Undesirable power behavior can include

- extremely high peak power,
- long-lasting high power period,
- repeated high-power patterns, and
- power behavior that reveals implementation information.

Whereas the first three are quite obvious for power and thermal management reasons, the last is related to system security. Differential power analysis [13] has been used to attack a system by comparing power traces generated by different inputs.

2.3 Resolution of power signals

We use “resolution” to refer to the level of detail of the temporal dynamics in a power signal. If a power signal can provide the average power over any m consecutive cycles, we say that it has a resolution level of m . Average power estimation for a whole simulation corresponds to the level ∞ ; cycle-accurate power traces are of the level 1, which is the highest level. The accuracy of a power trace can be measured at different resolution levels by using the following definition for error at the level m :

Definition: Error at the level m : Given a power trace sequence $S = [W_1, W_2, \dots, W_n]$, W_i being a sample window with m cycles, we

have measurement (or estimation) M_i for each window W_i . The error at the level m is defined as

$$\text{Error} = \frac{1}{n} \sum_{i=0}^n \left| \frac{\text{mean}(M_i) - \text{mean}(W_i)}{\text{mean}(W_i)} \right|, \quad (1)$$

where the absolute error is used to prevent positive and negative errors from canceling each other. The measurement M_i could be measured samples inside window W_i , or predicted values from adjacent windows if no simulation is carried out in window W_i .

Error at a resolution thus serves as a figure-of-merit to evaluate a power simulator or measurement. The error of measured current (power) consumption in Figure 3 is 79.2% at level 10, and reduces to 2.7% at level 70.

3. Correlation analysis

In this section, we discuss two types of correlations in power signals, temporal correlation and structural correlation.

3.1 Temporal correlation

Temporal correlation is the relation of a group of cycles with another group in the power signal. The most apparent temporal correlation is the periodicity. The periodicity of a trace will be revealed as peaks in the power signal spectrum. The spectrum gives the average energy of a signal at each frequency. A peak at frequency f_i is significant if

$$\text{Magnitude}(f_i) > u_p + k\sigma_p, \quad (2)$$

where u_p is the average magnitude over all frequencies, k is a threshold value (typically 3), and σ_p is the standard deviation in the magnitude over all frequencies. For an N -cycle power trace, we use the average power spectrum of L -point windows. A moving window of L -points with 50% overlap is applied to the N -cycle trace to form $2N/L - 1$ sections of length L . Then the spectrums of these sections are averaged. We refer to the largest significant frequency as the periodicity, p , of the trace.




Figure 6: Power spectrum of power signal in Figure 4: a significant peak is detected at 56 cycles, indicating a periodicity of 56 cycles.

The spectrum of the HDTV trace is shown in Figure 6. The significant periodicity is 56 cycles as denoted by the peak. It means that the trace repeats every 56 cycles.

3.2 Structural correlation

Structural correlation is the cross correlation between different components in a system. Figure 4 provides an example for the cor-

relation between the power consumption of different system components. Cross correlation is a standard method of estimating the degree to which two series are correlated. We use cross correlation analysis to explore the correlation between different power components. However, cross correlation cannot reveal the causal relationship between two components. When two power components are cross-correlated, we choose the one with larger power consumption as the dominant component for analysis.

Consider two power signals $x(i)$ and $y(i)$, where $i = 0, 1, 2, \dots, N-1$. The cross correlation r at delay d is defined as

$$r(d) = \frac{\sum_{i=0}^{N-1} [(x(i) - u_x)(y(i-d) - u_y)]}{\sqrt{\sum_{i=0}^{N-1} (x(i) - u_x)^2} \sqrt{\sum_{i=0}^{N-1} (y(i-d) - u_y)^2}}, \quad (3)$$

where u_x and u_y are the means of the corresponding series. When the index of the series is out of the range $[0, N-1]$, we use zero as the values. The denominator in the expression above serves to normalize the correlation coefficients such that $r(d) \in [-1, 1]$, the bounds indicating maximum correlation and 0 indicating no correlation. A high negative correlation indicates a high correlation but of the inverse of one of the series. The range of delay d is chosen between $[-p/2, p/2]$, where p is the detected periodicity. We use the maximum $|r(d)|$ among $d \in [-p/2, p/2]$ as the cross correlation of two series.

We employ t-test [14] to evaluate the statistical significance of r . The t-test evaluates whether the means of two groups are statistically different from each other. The hypotheses for the test are $H_0 : r = 0$ and $H_a : r \neq 0$. A low p-value for the test (less than 0.05 for example) indicates that there is evidence to reject the null hypothesis H_0 in favor of the alternative hypothesis H_a , or that there is a statistically significant relationship between the two series.




Figure 7: Power signal correlation matrix of components: components 1, 3, 6, and 8 are chosen as the major components in power simulation.

Figure 7 gives correlations of 13 components in an architectural power simulator, Sim-Panalyzer [8], distributed by the University of Michigan. If a significant correlation with p-value of 0.01 exists between two components i and j , we mark a star at position $[i, j]$. Since the correlation matrix is symmetric, only the upper portion is presented. The four components 1, 3, 6 and 8 that are highly correlated with all other components are chosen as the major components in power simulation. By tracking the power of these major components instead of all the components, it is possible to accelerate power simulation.

4. Adaptive acceleration of power simulation

To illustrate the applications of power signal processing, we next demonstrate how it can be applied to accelerate power simulation. We show that power traces can be obtained by selectively running the cycle-accurate power simulator without sacrificing accuracy significantly.

In Section 3, we showed that there are significant harmonic frequencies in power signals. This inspired us to employ the temporal relations for selective simulation. Similarly, the inspiration for structural selection comes from the high correlations between components in large systems. A power simulator usually divides the system into smaller functional components, each having its own power model. The total power is the sum of the power consumption of the individual components. In Section 3, our structural correlation analysis showed that a small number of dominating components are sufficient for total power estimation. As a result, speed-ups are achieved in power simulation if only the major components are simulated.

Based on the temporal and structural correlation detection, we devise an adaptive power simulation process presented in Algorithm 1. In the process, we start with extracting a vector \vec{T}^* for each simulated N -cycle trace, and compare it with vector \vec{T} . If vectors are matching, we double the skipped cycles and run another N -cycle simulation; otherwise, we simulate the successive N cycles. In step 2, a frequency of zero is used in case no significant frequency is detected as Eqn. 2. We use thresholding to determine the vector matching in step 7. Two vectors match if the absolute differences of all corresponding terms are less than the thresholds.

Algorithm 1 : Adaptive Sampling Power Simulation

- 1: Run a power trace Tr_0 for N cycles
 - 2: Calculate mean (μ), variance (σ), and periodicity (p)
 - 3: Initialize $\vec{T} = [u, \sigma, p]$ and index number $ind = 1$
 - 4: Skip $(ind - 1) * p$ cycles and simulate N -cycle power trace Tr
 - 5: Build $\vec{T}^* = [u_r, \sigma_r, p_r]$
 - 6: **if** $\vec{T} \approx \vec{T}^*$ **then**
 - 7: $ind = 2 * ind$
 - 8: **else**
 - 9: $ind = 1$
 - 10: **end if**
 - 11: Let $\vec{T} = \vec{T}^*$ and **goto** step 5
-

The power simulation employed in step 5 can use a full power simulator including all components, or use a partial simulation based on the correlation analysis of different components. The partial simulation reduces the simulation time and data by reducing the number of simulated components. Algorithm 2 presents the partial simulation technique to generate an N -cycle power trace in step 5 of Algorithm 1.

Algorithm 2 : Partial Simulation

- 1: Run L -cycle simulation fully
 - 2: Analyze structural correlation
 - 3: Determine major and non-major power components
 - 4: Simulate major components for $N - L$ cycles
 - 5: Add average power of non-major components from previous L cycles
-

The structural correlation analysis is used to identify the components that are highly correlated. For two highly correlated components, if one is much less than the other in average power, the power of the small one can be simplified into a constant value without

utilizing its detailed and time-consuming power model. This was described in Section 4.

5. Results for adaptive sampling

We used the SPEC2000 [15] benchmark suite to evaluate the effectiveness of adaptive acceleration based on power signal processing. We run Sim-Panalyzer [8] on SPEC2000 applications with the default inputs. Sim-Panalyzer models an ARM processor and performs cycle-accurate power simulation. Although the accuracy of most architectural power simulation is often disputable, we view Sim-Panalyzer as a system itself, instead of the ARM processor it attempts to model. The power traces of all 13 components for five million cycles was used as the baseline to validate the adaptive sampling and partial simulation techniques described in Section 4. The results of these experiments are presented in Table 1.

In the table, the first major column is the name of the benchmark. The second major column denotes the speed-up achieved using adaptive full simulation, which does not ignore any of the components. Speed-up is given by the ratio of the total cycles to the simulated cycles based on adaptive sampling and all power components. The third major column reports results for adaptive partial simulation. Num denotes the average number of major power components used for partial simulation, and speed-up is calculated as defined above. Under the column Error, we compare the results with the baseline full simulator at three different resolution levels: level ∞ (average power over the whole trace), level 1000, and level 100.

To validate the efficiency of power simulation based on adaptive sampling, we compare its results with two other sampling methods, periodic [16] and random (reported under the next major heading). In both cases, the whole trace is still divided into windows of m cycles each. Periodic sampling chooses the first cycle from every window; random sampling uniformly chooses a random cycle from every window. The error at level 100 for both periodic and random sampling is reported, which is significantly higher than that for adaptive sampling.

The table clearly demonstrates that the adaptive sampling algorithm is able to accelerate simulation by 96.7X on average across all the benchmarks with negligible error at a resolution level of 100. The performance of both periodic sampling and random sampling are comparable and both highly depend on the benchmark. The standard deviation of approximation errors across the eighteen benchmarks is 1.7% for adaptive sampling, much smaller than 9.9% for periodic or random sampling. It clearly shows that the adaptive sampling achieves much lower error over all cases, making it more suitable for simulation acceleration.

6. Conclusions

In this paper, we first investigated the power signal properties of digital systems and analyzed the limitations of power signal sources: cycle-accurate simulation and direct measurement. Next, we investigated signal processing techniques to discover temporal and structural relationships among power signals. To demonstrate the applications of power signal processing, we applied these techniques to accelerate an architecture-level processor power simulator. Experiments with the SPEC2000 benchmark suite demonstrate that power signal processing can accelerate its performance by 100X with negligible impact on power signal properties.

Our study shows that cycle accuracy at the system level is not necessary for many design tasks, such as power management and simulation. First, a well-designed power-supply network with decoupling capacitance will suppress cycle-accurate dynamics so that

Table 1: Simulation acceleration speed-up (X) and errors at different resolution levels (%)

Benchmark	Adaptive full simulation	Adaptive partial simulation					Traditional sampling	
	Speed-up	Num	Speed-up	Error (%)			Error at level 100 (%)	
				level ∞	level 100	level 1000	Periodic	Random
ammp	57.9	2.0	227.1	0.3	5.7	3.3	43.3	43.2
applu	22.9	2.0	102.7	0.5	6.4	2.9	19.7	19.7
apsi	88.2	2.0	324.9	0.1	4.0	4.2	2.2	2.2
art	42.8	2.8	113.8	0.3	4.7	0.9	6.6	6.7
bzip	39.4	3.0	110.0	0.1	5.4	4.1	5.5	5.6
craf	26.0	3.9	54.1	2.4	2.7	0.6	14.2	14.1
equa	31.2	3.9	62.7	1.8	5.0	2.6	13.3	13.3
gal	9.6	3.0	26.2	2.6	3.9	2.3	29.1	29.2
gap	24.3	3.0	59.8	0.2	1.1	1.9	3.0	3.1
gcc	9.9	3.4	24.7	0.4	6.1	3.1	12.9	12.9
gzip	17.3	3.2	42.8	1.2	4.2	2.5	13.3	13.2
luca	40.8	2.1	173.4	0.1	4.7	1.0	8.9	9.0
mcf	44.3	2.1	159.5	0.9	1.8	3.1	5.3	5.3
mesa	19.5	3.3	43.2	0.6	1.2	1.0	18.5	18.5
mgrid	21.4	4.3	49.9	1.2	2.2	1.3	13.4	13.6
swim	18.9	3.0	53.8	0.5	4.8	3.2	15.0	14.9
twolf	26.4	3.3	69.1	0.9	6.3	3.8	8.6	8.6
vpr	15.8	3.2	42.3	0.4	5.1	4.1	13.1	13.0
Average	30.9X	3.0	96.7X	0.8	4.2	2.5	13.7	13.7

it cannot be measured accurately. Second, simulation-based power traces are highly predictable. Our acceleration of 100X in simulation of the SPEC2000 benchmarks motivates that a power simulator should be able to support various tradeoffs between resolution and speed. Power signal processing readily supplies basic techniques for such a simulator.

We believe power signal processing provides a new perspective into automatic power analysis and optimization that will help address the two design productivity bottlenecks highlighted in the introduction. Beyond accelerating power simulation, future applications of power signal processing include tools that automatically analyze massive amounts of power data, detect undesirable power behavior for higher resolution simulation, and identify suspicious system components and behaviors.

7. References

- [1] L. Zhong *et al*, “RTL-aware cycle-accurate functional power estimation,” *IEEE Trans. Computer-aided Design*, vol. 25, pp. 2103–2117, Oct. 2006.
- [2] D. Stasiak *et al*, “Cell processor low-power design methodology,” *IEEE Micro*, vol. 25, pp. 71–78, Dec. 2005.
- [3] T. Sherwood, E. Perelman, and B. Calder, “Basic block distribution analysis to find periodic behavior and simulation points in applications,” in *Proc. Intl. Conf. Parallel Architectures and Compilation Techniques*, pp. 3–14, 2001.
- [4] R. Joseph, Z. Hu, and M. Martonosi, “Wavelet analysis for microprocessor design: Experiences with wavelet-based di/dt characterization,” in *Proc. Intl. Symposium High Performance Computer Architecture*, pp. 36–46, 2004.
- [5] N. R. Potlapally *et al*, “Accurate power macro-modeling techniques for complex RTL components,” in *Proc. Intl. Conference VLSI Design*, pp. 235–241, 2001.
- [6] S. Ravi, A. Raghunathan, and S. Chakradhar, “Efficient RTL power estimation for large designs,” in *Proc. Intl. Conference VLSI Design*, pp. 431–439, 2003.
- [7] D. Brooks, V. Tiwari, and M. Martonosi, “Wattch: A framework for architectural-level power analysis and optimizations,” in *Proc. Intl. Symposium Computer Architecture*, pp. 83–94, 2000.
- [8] Sim-Panalyzer: The SimpleScalar-Arm Power Modeling Project, <http://www.eecs.umich.edu/panalyzer/>.
- [9] W. Ye *et al*, “The design and use of simplepower: A cycle-accurate energy estimation tool,” in *Proc. Design Automation Conference*, pp. 340–345, 2000.
- [10] M. Powell and T. Vijaykumar, “Exploiting resonant behavior to reduce inductive noise,” in *Proc. Intl. Symposium Computer Architecture*, pp. 288–299, 2004.
- [11] R. A. DeCarlo and P. M. Lin, *Linear circuit analysis: Time domain, phasor, and Laplace transform approaches*. Oxford University Press, 2001.
- [12] J. Kozhaya, S. Nassif, and F. N. Najm, “A multigrid-like technique for power grid analysis,” *IEEE Trans. Computer-Aided Design*, vol. 21, pp. 1148–1160, Oct. 2002.
- [13] P. Kocher, J. Jaffe, and B. Jun, “Differential power analysis,” *Lecture Notes in Computer Science*, vol. 1666, pp. 388–397, 1999.
- [14] S. M. Ross, *Introduction to probability and statistics for engineers and scientists*. Elsevier Academic Press, 2004.
- [15] J. L. Henning, “SPEC CPU2000: Measuring CPU performance in the new millennium,” *Computer*, vol. 33, pp. 28–35, July 2000.
- [16] J. J. Yi and D. J. Lilja, “Simulation of computer architectures: Simulators, benchmarks, methodologies, and recommendations,” *IEEE Trans. Computers*, vol. 55, pp. 268–280, Mar. 2006.