Effects of Correlations on Accuracy of Power Analysis - An Experimental Study

Peter H. Schneider[1,2] Shankar Krishnamoorthy[3]

[1] Siemens AG, ZFE T SE 5, D-81730 Munich
[2] Inst. of EDA, TU-Munich, D-80290 Munich
[3] Synopsys, Inc., 700 E. Middlefield Road, Mountain View, CA, 94043-4033

Abstract – In this paper, we describe the correlation assumptions made by different power analysis methods and evaluate the impact on the accuracy of total power dissipation calculation as well as of the power dissipated by individual signals. Industrial circuits and applications are used. The results show that some assumptions cause inaccuracies of more than 100% for certain circuit types.

1 INTRODUCTION

1.1 Power Consumption

In CMOS circuits, power consumption is dominated by charging and discharging of capacitances, when a transition occurs at an internal or at an output signal [1]. Short circuit and leakage current dissipate additional power but can made small with proper design technique. The total capacitive power consumption in a circuit is given by

$$P_{circuit} = \frac{1}{2} V_{dd}^2 \sum_{\text{signal } i} C_i \alpha_i \quad (1)$$

where C_i is the total capacitance driven by signal *i*. The supply voltage is denoted by V_{dd}. Let α_i be the switching activity of signal *i*, i.e., the average number of $1\rightarrow0$ or $0\rightarrow1$ transitions on signal *i* per second. Since V_{dd} and C_i are known from the technology library parameters, the main problem of power analysis at logic level is the estimation of the switching activities α_i.

1.2 Switching Activity Analysis

In an early design stage, a designer wants to know the *total power consumption* of his design to explore architectural trade-offs and to decide whether a ceramic or a plastic packaging will be needed. However, in order to optimize a design for low power using automatic methods, *power consumption of each signal* must be known accurately since the optimization goal is to minimize $\sum C_i \alpha_i$. Then, α_i of each signal is part of the cost function.

In early work on power analysis, simulation of application vectors was used to estimate total power or transient behavior [2] - [6]. Event driven simulation was applied, which yields very high accuracy. The main problem of this strategy is the large CPU time requirement.

In order to overcome the large CPU time requirements but to preserve the accuracy, Najm et al. suggested to use the Monte Carlo technique in [7] - [11]. Primary input vectors are randomly generated according to statistical data at primary inputs. During the simulation of these pseudo random vectors, the Monte Carlo technique determines whether the switching activities have already converged and thus, whether the simulation can be terminated. With this technique, the number of simulated pseudo random vectors is minimum for a certain user-given accuracy value.

Two driving factors made most researchers work in the field of probabilistic computation. Firstly, the CPU time requirements of event driven simulation are too large to compute the cost function for optimization very often. Since at the logic-level, circuits are typically optimized step by step, a limit on the number of optimization steps deteriorates the optimization quality. Secondly, computing probabilities on BDDs is very efficient if all variables in the support are uncorrelated. Most probabilistic techniques use BDDs to propagate statistical data like signal probability and switching activity from primary inputs to each internal signal and primary output. With the number of correlations taken into account, these BDD based techniques require increasing computational costs. Several approaches [12] - [20] provide different trade-offs between accounting for correlations and computational costs.

ISLPED 1996 Monterey CA USA
0-7803-3571-8/96/$5.00 © 1996

1.3 Our Goal

Our goal is to understand the limitations of various analysis methods better so as to weigh these methods in terms of applicability for analysis and optimization. For this purpose, we describe the correlation assumptions and evaluate the inaccuracies caused by them.

Surprisingly, the inaccuracies caused by the correlation assumptions have hardly been examined yet, although they often dominate the accuracy of an analysis technique. In [16], the importance of sequential correlations is shown, and [19] considers the effect of spatial correlations at primary inputs.

In this paper, we define the different correlations and map each correlation assumption to the classes of analysis techniques that assume them. Then, we apply industrial circuits and applications to demonstrate which correlation assumption causes a major inaccuracy on what type of circuit. For the purpose of this evaluation, several experiments were performed. We describe in detail the obtained results. Inaccuracies are examined in terms of both total switching activity and switching activity of each signal. Some of the examined correlation assumptions cause inaccuracies of more than 100%. Such an inaccuracy is unacceptable for low power optimization. A cost function estimated with this level of inaccuracy will guide the optimization process to poor results.

The paper is organized as follows. In the next section, correlations and correlation assumptions are introduced. Section 3 describes the setup for experiments. In Section 4, the experiments are presented in detail. They are summarized and examined in Section 5. Conclusions are drawn in Section 6.

2 OVERVIEW OF CORRELATIONS

Ignorance of certain correlations is a common assumption that is made in power analysis research. In this section, we introduce different types of correlations and present a survey on which analysis techniques ignore what type of correlation.

2.1 Analysis Model

Let us first recall the definition of the terms *signal probability* and *switching activity*. Signal probability of a signal *i* denoted by $p(i)$ reflects the fraction of the time that signal *i* takes on value 1 as opposed to value 0. Switching activity of a signal *i* denoted by $E(i)$ is the average number of transitions on signal *i* (0 to 1 and 1 to 0) per second.

Most probabilistic approaches assume *zero gate delay*. Thus the power dissipation due to glitches is not reflected. For all presented evaluations, we also consider zero delay and thus neglect inaccuracies caused by this approximation because of the following reasons. Firstly, the contribution of glitch power has already been examined in several papers [13, 21, 22]. Secondly, the computation of the number of glitches needs accurate information on the delay caused by gates and by their interconnects. Logic-level power analysis was developed to estimate the power consumption at an early design stage. Typically at this stage, placement and routing has not yet been carried out and even gate delays are not known exactly until technology mapping has been finished. Thus, it is not possible to account for glitches correctly.

2.2 Correlations

Correlations were originally described for statistical processes [23]. Each signal in a logic circuit can be related to a statistical process [12]. Assuming zero-delay, correlations characterize signals as described in this section.

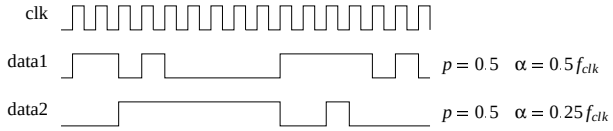

Figure 1: First order temporal correlation

Temporally uncorrelated: A signal i is said to be *temporally uncorrelated* if its next value is independent of its current and of all its previous values. Because of this independence, the probability of a 1 to 0 transition of signal i from one clock cycle to the next ($p(i\bar{i})$) is equal to $p(i)\,p(\bar{i})$. Similarly, $p(\bar{i}i)$ equals $p(\bar{i})\,p(i)$, and we obtain $\alpha_i = p(i\bar{i}+\bar{i}i)\,f_{clk} = (p(i)p(\bar{i}) + p(\bar{i})p(i))\,f_{clk}$. Thus, the signal probability of signal i describes its statistical behavior completely. An example for such a signal is signal data1 in Figure 1.

First order temporal correlation: A signal i is said to have first order temporal correlation if its next value depends on its current value but is independent of all previous values. Signal i can be described by a Markov chain. Its statistical behavior is completely determined by the signal probability and the switching activity. An example for such a signal is data2 in Figure 1.

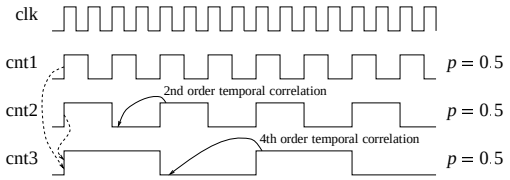

Figure 2: Higher order temporal and spatial correlation

Higher order temporal correlation: A signal is said to have higher order temporal correlation if its next value depends on its current value and on previous values. In Figure 2, bit 1, 2, and 3 of a binary counter are shown for illustration. Bit 2 (i.e., signal cnt2) has a second order temporal correlation. The next value of signal cnt2 will be 1 if its previous value was 0. Bit 3 has also an fourth order temporal correlation.

Spatial correlation: Two (or more) signals are said to be spatially correlated if each value that is taken on by one of these signals is correlated to the values the other signal(s) take on. In Figure 2, signals cnt1, cnt2, and cnt3 are spatially correlated. The signal cnt3, e.g., makes a transition only when cnt1 and cnt2 make a transition.

Spatial correlation of internal signals: Even if primary inputs are spatially uncorrelated, internal signals can be spatially correlated due to reconvergent fanouts. A reconvergent node has at least two fanins that depend on the reconvergent fanout stem. Thus, these fanins are spatially correlated. This correlation is referred to as spatial correlation of internal signals.

Sequential correlations: Spatial or temporal correlations at state lines that are induced by the cyclic structure of a sequential circuit are called sequential correlations. Even if primary inputs are uncorrelated or only first order temporally correlated, the states and state lines of an FSM can be higher order temporally correlated. An example is again the binary counter in Figure 2. A reset signal may be the only input to the counter. In the example, this reset signal is uncorrelated and has a very low signal probability. Nevertheless, the state lines (which are the signals in Figure 2) are higher order temporally and spatially correlated.

2.3 Survey of Current Analysis Methods

Table 1 presents a survey of the assumptions that are made by various analysis strategies. Symbol "x" (or "app") is used if a technique exactly (or approximately) accounts for a certain type of correlation. Symbol "ign" is used if a technique completely ignores a certain type of correlations. Obviously, simulating application vectors (column "sap") accounts for all correlations exactly.

The Monte Carlo simulation (column "mcs") requires pseudo random vectors, which are only first order temporally correlated, instead of application vectors as input to the simulation. Thus, spatial and higher order temporal correlations of primary inputs are ignored. First order temporally correlated pseudo random vectors will be defined in the next section.

	pcb	psq	mcs	sap
First order temp. corr. of PI	x	x	x	x
Sequential correlations	ign	app	x	x
Spatial corr. of internal signals	app	app	x	x
Spatial corr. of PI	app	ign	ign	x
Higher order temp. corr. of PI	ign	ign	ign	x

Table 1: Correlation assumptions vs. analysis techniques

In column "pcb", probabilistic techniques are considered that are developed for analysis of combinational circuits, while column "psq" considers probabilistic techniques developed for analysis of sequential circuits. We assume that the most advanced method is applied, i.e., the method that has the least number of entries "ign" in Table 1.

For combinational circuits, the most advanced technique is the technique of [19]. All other techniques do not account for spatial correlations at primary inputs. For sequential circuits, no probabilistic estimation technique considers spatial correlations at primary inputs. The only probabilistic technique that considers sequential correlations and temporal correlations at primary inputs is [20].

Even these most advanced probabilistic techniques can handle spatial and sequential correlations only approximately. Higher order temporal correlations are never taken into account by probabilistic techniques presented so far.

3 Experimental Setup

3.1 Metrics

In this section, we define the accuracy measurements that are used later in the paper. These measurements are the relative average error in terms of total or signal switching activity.

The *total switching activity* denoted by A is the sum of the switching activities of all signals:

$$A = \sum_{\text{signal } i} \alpha_i$$

The *average relative error* in terms of *total switching activity* is:

$$\mathbf{e}_{\text{total}} = \left| \frac{A_{\text{acc}} - A_{\text{ign}}}{A_{\text{acc}}} \right|$$

where subscript $_{\text{ign}}$ denotes values, which are measured while ignoring the correlation under evaluation, and subscript $_{\text{acc}}$ denotes values, which are measured while accounting for the correlation accurately.

The *average relative error* in terms of *signal switching activity* is:

$$\mathbf{e}_{\text{signals}} = \frac{1}{\#\text{signals}} \sum_{\text{signal } i} \left| \frac{(\alpha_i)_{\text{acc}} - (\alpha_i)_{\text{ign}}}{(\alpha_i)_{\text{acc}}} \right|$$

The measurement $\mathbf{e}_{\text{signals}}$ provides a much more refined assessment of the estimation quality than $\mathbf{e}_{\text{total}}$. During the computation of $\mathbf{e}_{\text{total}}$, over and under estimated switching activities compensate and thus, even if each signal is estimated very inaccurately, $\mathbf{e}_{\text{total}}$ can show reasonable accuracy. For optimization, however, $\mathbf{e}_{\text{signals}}$ must be small. Otherwise the optimization process will not be guided properly and will thus yield poor results. This will be shown in Section 4.2.

As suggested in [9], internal signals are divided into signals having high activity and signals having low activity. For the quantity $\mathbf{e}_{\text{signals}}$, only signals with switching activity $\alpha_i > \frac{\text{frequency}}{1000}$ are considered. The reason for disregarding low activity signals is that their contribution to the power dissipation is very small but they cause large relative errors: For example, a signal makes 2 transitions during an 10^5 vector simulation although it should only make 1 transition; the relative error is 100% but because of the signal's low α_i, its power dissipation can be neglected.

3.2 Circuits

The following circuits were used for the described experiments:

- **Controller**
  - Alarm clock: Alarms can be set and triggered.
  - 8-bit Microcontroller
  - Sort: Bubble sort of integer numbers

- **Datapath modules**
  - Two 8-bit multiplier: wallace tree, carry save
  - Two implementations of 8-bit carry look ahead adder
- **Mixed circuit** (i.e., circuit with control and datapath)
  - IIR: Infinite Impulse Response Digital Filter

3.3 Stimuli

For each circuit, a VHDL testbench was available representing either the application for which the circuit was designed or the functional test of the application. The vectors generated by these testbenches are referred to as *application vectors*. These vectors contain all correlations of the application.

For several experiments, we need vectors that are only first order temporally correlated but have the same signal probability and switching activity as the application vectors. For this purpose, a random generator was biased such that vectors are generated according to signal probabilities and switching activities of the application. The only correlation these vectors have is a first order temporal correlation. We denote these vectors as *pseudo random vectors*.

For each experiment that needs pseudo random vectors, 10^5 vectors were generated and simulated. Simulating 10^5 pseudo random vectors causes very low inaccuracies as shown in [11, 24].

4 Description of Experimental Procedures

In this section, we evaluate the inaccuracies caused by each correlation assumption. In the next section, we will summarize and comment the results of this section. Due to the lack of space, the experimental setups are not shown in detail. Upon request, the authors will be glad to provide an extended version of this paper.

4.1 Inaccuracies Due to Ignoring Sequential Correlations

To evaluate the inaccuracies due to ignoring sequential correlations, a simulation that accounts for sequential correlations is compared to a simulation that does not account for sequential correlations. Sequential correlations are ignored if pseudo random vectors are assumed at state lines instead of the correlated patters.

	$\mathbf{e}_{\text{total}}$	$\mathbf{e}_{\text{signals}}$
Datapath modules	No sequential elements	
Mixed circuit	3%	8%
Controller	25%	120%

Table 2: Inaccuracies due to ignoring sequential correlations

Table 2 gives the average relative error for this experiment. The small datapath modules do not contain any sequential element. The inaccuracies for the mixed circuit are very low, lying below 10% for both total and signal switching activity. For the controller type circuits, the error in terms of total switching activity is about 25%. Considering the signal switching activity, the error was typically higher than 100%.

4.2 Inaccuracies Due to Ignoring Spatial and Higher Order Temporal Correlations at Primary Inputs

The activities obtained by a simulation of application vectors are compared to a simulation of pseudo random vectors. The results are presented in Table 3. For the small datapath modules, ignoring all primary input correlations beside of first order temporal correlations causes an error smaller than 10%. Unfortunately, the error in terms of signal switching activities is 30% on average. For the wallace tree multiplier, for example, we obtained $\mathbf{e}_{\text{signals}} = 39\%$. This multiplier does not contain sequential elements.

	$\mathbf{e}_{\text{total}}$	$\mathbf{e}_{\text{signals}}$
Datapath modules	8%	30%
Mixed circuit	26%	56%
Controller	5% (. . . 50%)	10% (. . . 150%)

Table 3: Inaccuracies due to ignoring PI correlations

The scattered diagram for this multiplier is shown in Figure 3. The horizontal axis "accounting for PI correlations" shows the switching activity of a signal if all primary input correlations are taken into account, i.e., the switching activity obtained by simulating application vectors. The vertical axis "ignoring PI correlations" gives the switching activity if at primary inputs only first order temporal correlations are accounted for. Each diamond in the diagram corresponds to one signal of the considered 8-bit multiplier. If ignoring primary input correlations causes no inaccuracy for a certain signal then the diamond of this signal hits the dotted line.

Low power optimization techniques try to delete or "hide" signals with high activity or assign low capacitances to these signals, while signals with low switching activities are allowed to be additionally created or be assigned high capacitances. If power optimization is guided by the inaccurate analysis result in Figure 3, the highly overestimated signal marked with label "a)" is considered to have a higher switching activity than most signals actually having a higher switching activity. For example, the signal with label "b)" would be assumed to have a lower switching activity than the signal with label "a)", and thus, a power optimization technique may assign a higher capacitance to this signal than to the signal with label "a)". Therefore, the optimization process will result in a higher instead of a low power implementation.

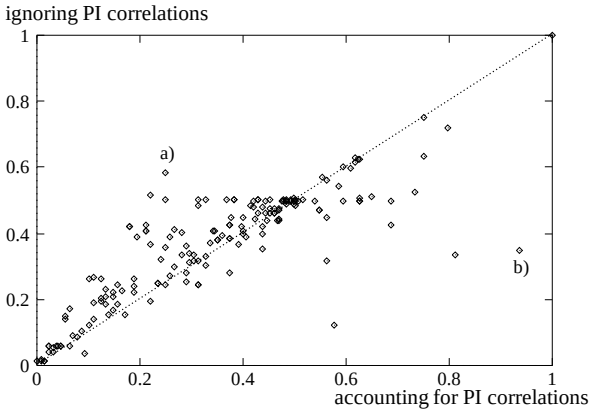

Figure 3: 8-bit multiplier w/ and w/o ignoring PI correlations

For the mixed circuit, the error in terms of total switching activity is 26%. The error in terms of signal switching activity is 56%, which is too high to guide the optimization.

For the controller type circuits, it is difficult to make a statement on the impact of the primary input correlations, because the impact heavily depends on whether the random vectors represent real operations. We will illustrate this with an example. In the application, the reset signal may initially be 1 for four clock cycles and then be 0 for all other vectors. So, there was one transition. Now, vectors are randomly generated. The signal probability and switching activity of the reset signal are preserved if this signal is 0 for all vectors and switches to 1 for the four last vectors. The controller, however, may start in a state, which is unreachable from the initial state of the application. If this happens, the error for pseudo random vectors can be very high while it is typically very low if the random vectors represent real operations.

4.3 Inaccuracy Due to Ignoring Spatial Correlations of Internal Signals and Simultaneous Switching

We computed the switching activities of the combinational logic with the probabilistic technique of [12]. This probabilistic technique accounts for first order temporal correlations accurately. Spatial correlations at primary inputs and at internal signals are ignored.

	$\mathbf{e}_{\text{total}}$	$\mathbf{e}_{\text{signals}}$
Datapath modules	10%	15%
Mixed circuit	25%	50%
Controller	10%	20%

Table 4: Inaccuracies due to ignoring internal spatial correlations and simultaneous switching

In Table 4, the activities obtained by using this probabilistic technique are compared to the activities obtaind by a simulation of pseudo random vectors. For the small datapath modules and for the controller type circuits, the accuracy in terms of total and the signal switching activity is sufficient. For the mixed circuit, the accuracy of the total switching activity may be acceptable, while the signal switching activity shows an average relative error of 50%.

4.4 *Inaccuracy Due to Ignoring Different Application Vector Sets*

We also performed an experiment to determine the sensitivity of circuits to different applications. This was done to determine whether it is possible to assign an invariant power number to a module (e.g., a multiplier) in order to characterize it in a module library.

We applied two different applications to the same circuit, simulated the vector sets of these two applications, and compared the determined switching activities of the internal signals. The vector sets of these two applications had different statistical behavior, i.e., a primary input or state line has different switching activities due to the different vector sets.

	$\mathbf{e}_{\text{total}}$	$\mathbf{e}_{\text{signals}}$
Datapath modules	35%	150%
Mixed circuit	30%	100%
Controller	5%	15%

Table 5: Different switching activities due to different applications

Switching activities of datapath modules and mixed circuits depend on the application. In our experiments, the total power differs by up to 50% and the average relative error in terms of signals was for all experiments larger than 100%. However, the examined control type circuits tend to be very insensitive to the input vectors as long as these vectors represent useful operations. Even the error for signal switching activity was only about 15%.

5 Overview And Implications

5.1 *Summary of All Experimental Results*

Table 6 summarizes the main source of inaccuracy for each circuit type. Let us first consider total power consumption. The accuracy appears to be sufficient for estimates to be used at an early design stage to evaluate different architectures for a module.

	Main Source of Inaccuracy	$\mathbf{e}_{\text{total}}$	$\mathbf{e}_{\text{signals}}$
Datapath m.	Ignore PI corr.	8%	30 %
Mixed circuit	Ignore PI corr.	26%	56%
	Ignore internal spatial corr. & simultaneous switching	25%	50%
Controller	Ignore sequential corr.	25%	120%

Table 6: Main sources of inaccuracy

Let us now consider the inaccuracies in terms of signal switching activities. Two observations strike out from Table 6. Firstly, even the smallest error in Table 6, which has appeared for small datapath modules, can produce a scattered diagram like the one in Figure 3. Using such analysis data will cause poor low power optimization results. This indicates that the correlation assumption is unacceptable. Secondly, each circuit type has a different main source of inaccuracy. Therefore, an estimation technique that is useful for all types of circuits must account for all these sources of inaccuracy.

Note that in the tables above, we have presented relative errors. In several papers about switching activity analysis, the absolute error instead of the relative error is given in terms of transition probabilities (i.e., switching activity per clock cycle). Obviously, probabilities are smaller than 1 and our observation is that transition probabilities are on average significantly smaller than 1. This is also the case for the example in Figure 3. The scattered diagram of this example illustrates an example of an average relative error $\mathbf{e}_{\text{signals}}$ equal to 39%. For this example, the average absolute error equals 0.067. Thus, it is about 6 times smaller than the relative error.

5.2 *Implications for Power Analysis and Optimization*

Table 1 shows that all analysis techniques beside of simulating application vectors cannot account for either sequential correlations or spatial correlations of primary inputs. Additionally, none of these techniques can handle higher order temporal correlations.

Probabilistic techniques handle spatial and sequential correlations only approximately as shown in Table 1. Furthermore, a sequential circuit is represented as a system of equations. The accuracy of the solution of a system of equations is unclear, given that the equations are inaccurate due to the correlation approximations. For large sequential ISCAS benchmarks, no results of probabilistic techniques have been presented so far.

Monte Carlo simulation can handle sequential correlations and spatial correlations of internal signals exactly. But spatial correlations at primary inputs and higher order temporal correlations are ignored. For datapath modules and mixed circuits, ignoring primary input correlations causes inaccuracies that are too large to guide low power optimization.

For power optimization using automatic methods, the only analysis technique that is sufficiently accurate for all circuit types is the simulation of application vectors. Unfortunately, this technique may require considerable CPU time.

6 Conclusion

In the recent years, many techniques have been presented to analyze switching activities at signals inside a circuit. To achieve high efficiency, all these techniques are based on correlation assumptions. We described which correlation assumptions are made by what kind of estimation techniques, and we examined the inaccuracies caused by making these assumptions.

For total power consumption, existing analysis techniques are reasonably accurate for most circuit types. However, the accuracy in terms of signal switching activity suffers significantly due to several sources of inaccuracy. The only existing strategy that covers all sources of inaccuracy for all circuit types is the simulation of application vectors.

Acknowledgement

The authors would like to thank Ashutosh Mauskar of Synopsys for providing circuits and applications and for his help.

References

[1] L. Glasser and D. Dobberpuhl, *The Design and Analysis of VLSI Circuits.* Addison-Wesley Publishing Company, 1985.

[2] A. Deng, Y. Shiau, and K. Loh, "Time Domain Current Waveform Simulation of CMOS Circuits," in *IEEE/ACM ICCAD*, pp. 208 – 211, Nov. 1988.

[3] S. Chowdhury and J. S. Barkatullah, "Estimation of Maximum Current in MOS IC Logic Circuits," *IEEE Transactions on CAD*, vol. 1, pp. 63 – 71, Jan. 1993.

[4] F. Rouatbi, B. Haroun, and A. J. Al-Khalili, "Power Estimation Tool for Sub-Micron CMOS VLSI Circuits," in *IEEE/ACM ICCAD*, pp. 204–209, Nov. 1992.

[5] T. H. Krodel, "PowerPlay - Fast Dynamic Power Estimation Based on Logic Simulation," in *IEEE ICCD*, pp. 96–100, October 14-16 1991.

[6] W. T. Eisenmann and H. E. Graeb, "Fast Transient Power and Noise Estimation for VLSI Circuits," in *IEEE/ACM ICCAD*, pp. 252–257, November 6–10 1994.

[7] C. M. Huizer, "Power dissipation analysis of CMOS VLSI circuits by means of switch-level simulation," in *Sixteenth European Solid-State Circuits Conference (ESSCIRC'90)*, pp. 61 – 64, Sept. 1990.

[8] R. Burch, F. Najm, P. Yang, and T. Trick, "A Monte Carlo Approach for Power Estimation," *IEEE Transactions on VLSI*, pp. 63 – 71, 1993.

[9] M. G. Xakellis and F. N. Najm, "Statistical Estimation of the Switching Activity in Digital Circuits," in *ACM/IEEE DAC*, pp. 728 – 733, June 1994.

[10] F. N. Najm, S. Goel, and I. N. Hajj, "Power Estimation in Sequential Circuits," in *ACM/IEEE DAC*, pp. 635 – 640, June 1995.

[11] T.-L. Chou and K. Roy, "Statistical Estimation of Sequential Circuit Activity," in *IEEE/ACM ICCAD*, pp. 34 – 37, Nov. 1995.

[12] F. N. Najm, "Transition Density: A Stochastic Measure of Activity in Digital Circuits," in *ACM/IEEE DAC*, pp. 644 – 649, June 1991.

[13] A. Ghosh, S. Devadas, K. Keutzer, and J. White, "Estimation of Average Switching Activity in Combinational and Sequential Circuits," in *ACM/IEEE DAC*, pp. 253 – 259, June 1992.

[14] B. Kapoor, "Improving the Accuracy of Circuit Activity Measurement," in *ACM/IEEE DAC*, pp. 734 – 739, June 1994.

[15] P. H. Schneider, U. Schlichtmann, and K. J. Antreich, "A new power estimation technique with application to decomposition of boolean functions for low power," in *EURO-DAC*, pp. 388 – 393, Sept. 1994.

[16] C.-Y. Tsui, M. Pedram, and A. M. Despain, "Power Efficient Technology Decomposition and Mapping Under an Extended Power Consumption Model," *IEEE Transactions on CAD*, vol. 13, pp. 1110–1122, Sept. 1994.

[17] G. D. Hachtel, E. Macii, A. Pardo, and F. Somenzi, "Probabilistic Analysis of Large Finite State Machines," in *ACM/IEEE DAC*, pp. 270 – 275, June 1994.

[18] T.-L. Chou, K. Roy, and S. Prasad, "Estimation of Circuit Activity Considering Signal Correlations and Simultaneous Switching," in *IEEE/ACM ICCAD*, 1994.

[19] R. Marculescu, D. Marculescu, and M. Pedram, "Efficient Power Estimation for Highly Correlated Input Streams," in *ACM/IEEE DAC*, pp. 635 – 634, June 1995.

[20] P. H. Schneider and B. Wurth, "Transition probability estimation for combinational and sequential circuits," in *International Workshop on Logic Synthesis (IWLS)*, pp. P4b1 – P4b10, May 1995.

[21] M. Favalli and L. Benini, "Analysis of Glitch Power Dissipation in CMOS IC's," in *ACM/IEEE International Symposium on Low Power Design*, Apr. 1995.

[22] M. Eisele and J. Berthold, "Dynamic Gate Delay Modeling for Accurate Estimation of Glitch Power at Logic Level," in *International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, Oct. 1995.

[23] A. Papoulis, *Probability, Random Variables and Stochastic Processes.* McGraw-Hill, 1991.

[24] P. H. Schneider, "PAPSAS: A Fast Switching Activity Simulator," in *International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, pp. 350 – 360, Oct. 1995.

[25] Synopsys, Inc., *VSS product reference manual*, 1995.