Feedback, Correlation, and Delay Concerns in the Power Estimation of VLSI Circuits

Farid N. Najm
ECE Dept. and Coordinated Science Lab.
University of Illinois at Urbana-Champaign
Urbana, IL 61801

Abstract – With the advent of portable and high-density microelectronic devices, the power dissipation of integrated circuits has become a critical concern. Accurate and efficient power estimation during the design phase is required in order to meet the power specifications without a costly redesign process. As an introduction to the other papers in this session, this paper gives a tutorial presentation of the issues involved in power estimation.

I. Introduction

Power dissipation of integrated circuits (ICs) is a major concern in VLSI circuits and systems design. One reason for this is the dramatic decrease in feature size and the corresponding increase in IC transistor count and clock frequency. The resulting high power dissipation elevates chip temperature and can cause performance degradation and decreased lifetime. A usually expensive solution is to use costly packaging and heat-sink technologies, which also complicates system design and increases system cost.

Another reason for the recent prominence of the power dissipation problem is the growing demand for portable communications and computing systems. These systems are only practical if they can be operated for extended periods of time without recharging or battery replacement. Given the relatively slow improvement in battery technology, designers have to use low-power ICs to increase the operating life of their portable products. Increasingly, it is the digital parts of these systems that are the high-power offenders, because of the extensive digital signal processing (DSP) that is usually performed.

Managing the power of a digital IC design adds to a growing list of problems that IC designers and design managers have to contend with. Computer Aided Design (CAD) tools are needed to help with the power management tasks. Specifically, there is a need for CAD tools to estimate power dissipation *during* the design phase in order to help meet the power specifications without a costly redesign process.

In the commonly used CMOS technology, the chip components (gates, cells) draw power supply current only during a logic transition (ignoring the small leakage current). While this is considered an attractive low-power feature of this technology, it makes the power-dissipation highly dependent on the *switching activity* inside these circuits. Simply put, a more active circuit will consume more power. This complicates the power estimation problem because the power becomes dependent on the circuit inputs, not only on the circuit structure. The power is said to be *input pattern-dependent.*

A simple and straight-forward method of power estimation is to *simulate* the circuit, using a circuit simulator, to obtain the power supply current, from which the average power can be computed. In order to use this method, *complete and specific* information about the input signals is required, in the form of voltage waveforms. Hence, we describe this simulation-based technique as being *strongly pattern-dependent*, which is a major problem with this approach. While it is not always the case, input signals may be unknown during the design phase because they depend on the system (or chip) in which the chip (or functional block) will eventually be used. Specifically, for a microprocessor or a DSP chip, the data inputs can not be determined apriori, because they depend on how the chip is deployed in the field. A simplistic method to get around this problem is to exhaustively simulate the circuit for all possible inputs, but this is obviously impractical.

Recently, several techniques have been proposed to overcome this problem [1] by using *probabilities* to describe the set of *all possible logic signals*, and then

ACM/IEEE Design Automation Conference, 1995.

studying the power resulting from the collective influence of all these signals. This formulation achieves a certain degree of *pattern-independence* that allows one to efficiently estimate and manipulate the power dissipation. Most of these techniques simplify the problem in three ways:

(1) It is assumed that the power supply and ground voltage levels throughout the chip are fixed, so that it becomes simpler to compute the power by estimating the current drawn by every sub-circuit assuming a given fixed power supply voltage.

(2) It is assumed that the circuit is built of logic gates and latches, and has the popular and well-structured design style of a *synchronous sequential circuit*, as shown in Fig. 1. In other words, it consists of latches driven by a common clock and combinational logic blocks whose inputs (outputs) are latch outputs (inputs). It is also assumed that the latches are edge-triggered and, with the use of a CMOS design technology, the circuit draws no steady-state supply current.

(3) Finally, it is commonly accepted that, in accordance with the results of [2], it is enough to consider only the charging/discharging current drawn by a logic gate, so that the short-circuit current during switching is neglected. This restriction is not absolutely required, and there are ways of avoiding it.

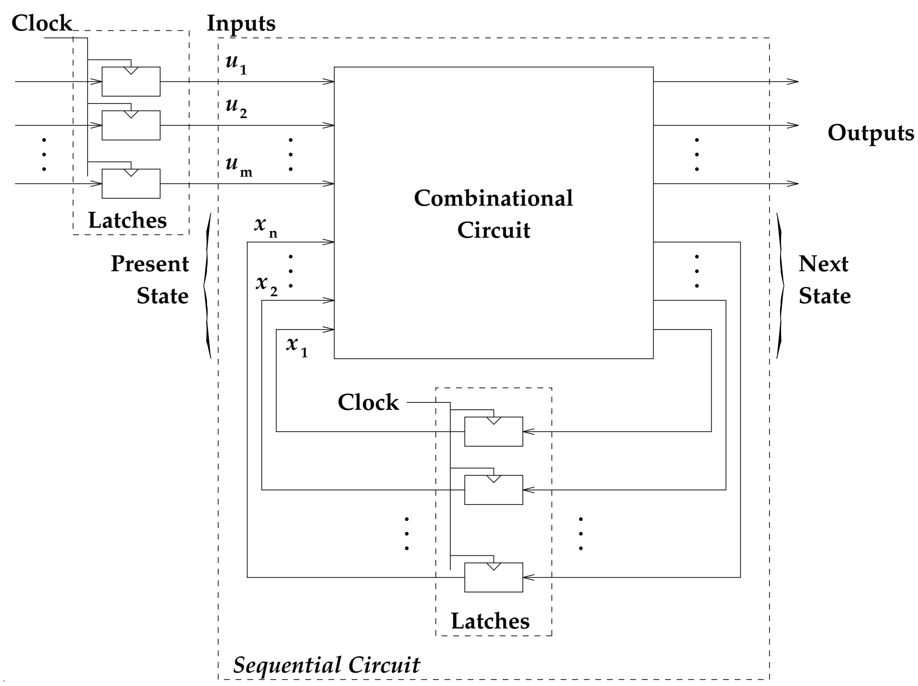

Figure 1. A combinational circuit embedded in a synchronous sequential design.

Therefore, the average power dissipation of a circuit can be broken down into (a) the power consumed by the latches and (b) that consumed by the combinational logic blocks. This provides a convenient way to decouple the problem and simplify the analysis. Correspondingly, we have found that an efficient way to estimate the power is the following two-step approach:

1. Solve for the latch power by examining the behavior of the whole circuit as a finite state machine (FSM).
2. Use the results of the FSM analysis to compute the power for the combinational circuit block.

This process is easily formulated using probabilities. In what follows, we will see how probabilities are relevant to power estimation, and then consider separately the computation of the latch power and the combinational circuit power. In doing so, we will uncover a number of challenges and difficulties, including the issues of feedback, correlation, and delay.

II. Using Probabilities

Probability has been used in order to solve the pattern-dependence problem, as follows. Instead of simulating the circuit for a large number of input patterns and then averaging the results, one can simply compute (from the input pattern set, for instance) the fraction of cycles in which an input signal makes a transition (a *probability* measure) and use that information to estimate (somehow) how often internal nodes transition and, consequently, the power drawn by the circuit. Conceptually, this idea is shown in Fig. 2, which depicts both the conventional path of using circuit simulation and the alternative path of using probabilities. In a sense, one performs the averaging *before*, instead of after, running the analysis. Thus, a single run of a probabilistic *analysis tool* replaces a large number of circuit simulation runs, provided some loss of accuracy can be tolerated. The issues are exactly what probabilities are required, how they are to be obtained and, most importantly, what sort of analysis should be performed.

In practice, a knowledgeable designer may be able to directly provide the required input probabilities, eliminating the need for a large set of specific input patterns. In any case, the results of the analysis will depend on the supplied probabilities. Thus, to some extent the process is still pattern-dependent and the user must supply information about the *typical* behavior at the circuit inputs, in terms of probabilities. However, since one is not required to provide complete and specific information about the input signals, we call these approaches *weakly pattern-dependent.*

There are many ways of defining probability measures associated with the transitions made by a logic signal, be it at the primary inputs or at an internal node. We start with the following two:

Definition 1. (signal probability): *The signal probability $P_s(x)$ at a node x is defined as the average fraction of clock cycles in which the steady state value of x is a logic high.*

Definition 2. (transition probability): *The transition probability $P_t(x)$ at a node x is defined as the average fraction of clock cycles in which the steady state value of x is different from its initial value.*

The signal probability is a relatively old concept that was first introduced to study circuit testability [3]. In the following sections, we will see how the signal and transition probabilities are related and how they can be used to compute circuit power.

III. Latch Power

Whenever the clock triggers the latches, some of them will make a transition and will draw power. Thus latch power is drawn in synchrony with the clock. If the transition probabilities $P_t(x)$ at the latch outputs are known, then the average power consumed by one latch is simply:

$$\frac{1}{2T_c} V_{dd}^2 C_x P_t(x)$$

where T_c is the clock period and C_x is the total capacitance at the latch output.

Thus the computation of the latch power reduces to finding the latch transition probabilities. However, computing the probabilities $P_t(x_i)$ from the input signal and/or transition probabilities is not trivial. In fact, it can be shown that finding these probabilities *exactly* is $\mathcal{NP}$-hard. Even finding them *approximately* is not easy, because the feedback creates the difficult situation where future signal values are related to their past and present values. Thus the signals in consecutive clock cycles may be correlated due to the feedback. This *feedback concern* is addressed in one of the papers [10] in this session. The other three papers address problems related to power computation in the combinational circuit.

IV. Combinational Circuit Power

Whereas latch power is drawn in synchrony with the clock, the same is not true for gates inside the combinational logic. Even though the inputs to a combinational logic block are updated by the latches (in synchrony with the clock), the internal gates of the block may make several transitions before settling to their steady state values for that clock period.

These additional transitions have been called *hazards* or *glitches.* Although unplanned for by the designer, they are not necessarily design errors. Only in the context of low-power design do they become a nuisance, because of the additional power that they dissipate. It has been observed [4] that this additional power dissipation is typically 20% of the total power, but can be as high as 70% of the total power in some cases such as combinational adders. We have observed that in a 16-bit multiplier circuit, some nodes make as many as 20 transitions before reaching steady state. This component of the power dissipation is computationally expensive to estimate, because it depends on the *timing relationships* between signals inside the circuit. Consequently, many proposed power estimation techniques have ignored this issue. We will refer to this component of power as the *toggle power.* Computing the toggle power is one main challenge in power estimation. How exactly do the circuit delays affect the switching activity? How sensitive is the activity (and, therefore, the power) to the exact delay values and to their relative magnitudes over reconvergent fanout paths? This *delay concern* is dealt with in one of the papers in this session [8].

Recall the signal and transition probabilities, defined above, and suppose they are computed for every gate output node in the combinational block. It is important to note that the resulting values are unaffected by the circuit internal delays. This is because, by definition, they depend only on steady state signal values in a clock cycle. Indeed, these values would remain the same even if a zero-delay timing model were used. If this is done, however, the toggle power would be automatically excluded from the analysis. This is a serious shortcoming of techniques that are based on these measures.

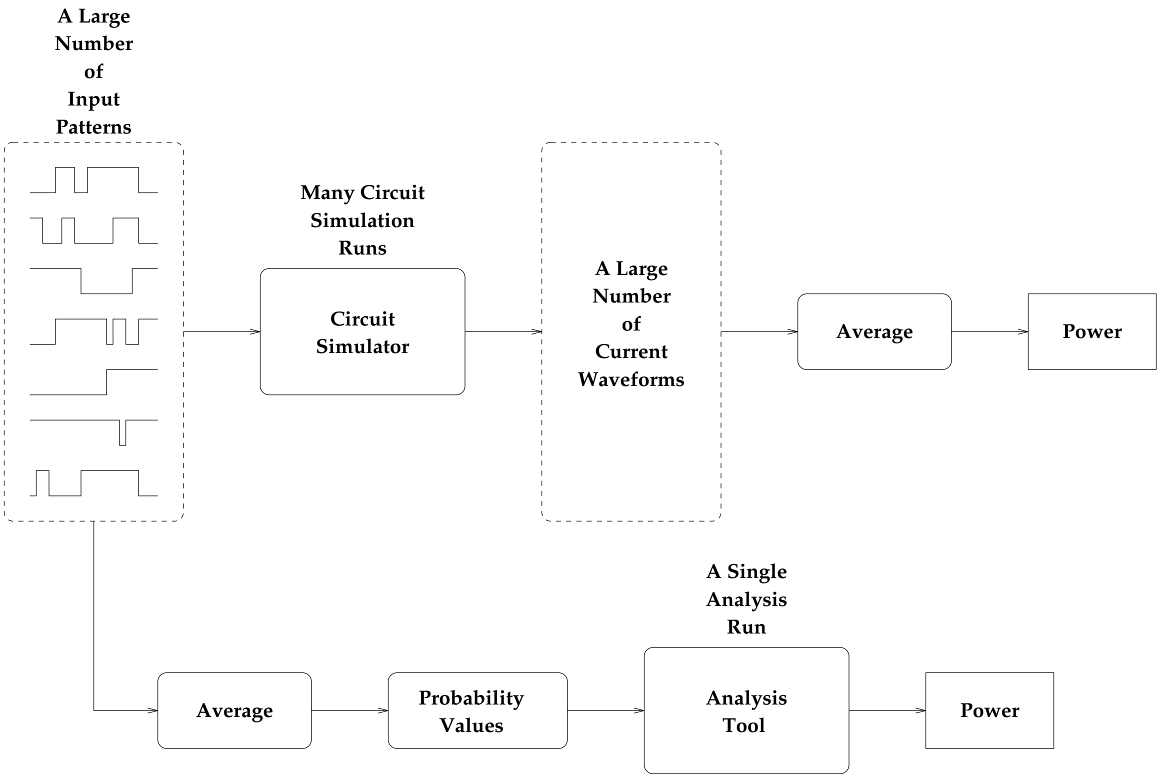

Figure 2. An alternative flow for power estimation.

A. Zero delay

If a zero-delay model is assumed and the transition probabilities are computed, then the combinational circuit power can be computed as:

$$P_{av} = \frac{1}{2T_c} V_{dd}^2 \sum_{i=1}^{n} C_i P_t(x_i) \tag{1}$$

where T_c is the clock period, C_i is the total capacitance at node x_i, and n is the total number of circuit nodes that are outputs of logic gates or cells. Since this assumes at most a single transition per clock cycle, then this is actually a *lower bound* on the true average power. Nevertheless, the results of a zero delay analysis may be useful as a rough technology-independent indication of the power requirements of a circuit.

In order to compute the internal transition probabilities, it is common to start by finding the signal probabilities. This, by itself, is not easy and can be shown to be $\mathcal{NP}$-hard. The problem has to do with whether the input signals to a logic gate (viewed as *random variables*) are independent or not. In practice, logic signals may be correlated so that, for instance, two of them may never be simultaneously high, or they may never (or always) switch together. Primary inputs to the combinational block may be correlated due to the feedback. And even if these inputs are assumed independent, other internal signals may be correlated due to reconvergent fanout (a gate fans out into two signals that eventually recombine as the inputs of some gate downstream). However, it is computationally too expensive to compute these correlations. This *correlation concern* is addressed by one of the papers [7] in this session.

Some have argued that the correlations do not significantly affect the final result, so that circuit input and internal nodes may be assumed to be *independent*. We refer to this as a *spatial independence* assumption. It leads to a significant simplification in computing the internal signal probabilities. If $y = ab$ is an AND gate output, and a and b are independent, then $P_s(y) = P_s(a)P_s(b)$. For an OR gate, we have $P_s(y) = P_s(a) + P_s(b)$. Thus the internal node probabilities are simply computed from those of the input nodes. The primary input node probabilities can be obtained as results of the analysis of the FSM, carried out previously.

To find the internal *transition* probabilities, we must deal with another independence issue of whether the values of the same signal in two consecutive clock cycles are independent or not. If assumed independent, then the transition probability can be easily obtained from the signal probability according to:

$$P_t(x) = 2P_s(x)P_s(\overline{x}) = 2P_s(x)\left[1 - P_s(x)\right] \qquad (2)$$

We refer to this as a *temporal independence* assumption. If this assumption is *not* made, then one must somehow represent the correlation between successive input vectors and internal signals. Given our formulation of power estimation as a two-step process, the correlation between two consecutive primary input bit values (on the same input line) can be obtained as transition probabilities computed during the FSM analysis. But that does not account for all input correlations. Correlations across more than one clock edge are not available, and correlations between one signal and previous values of other signals are also not available. Not only is computing these correlations too expensive, but making use of them during the computation of the combinational circuit power is also difficult. This is another aspect of the *correlation concern*, and is addressed by another paper [9] in this session.

B. Non-Zero delay

The problems described above become even worse in the case of non-zero delays. In this case, new probability measures are required to properly formulate the power dissipation problem. One such measure is the *transition density* [5, 6]. The transition density at node x is the average number of transitions per second at node x, denoted $D(x)$. Formally:

Definition 3. (transition density) *If a logic signal $x(t)$ makes $n_x(T)$ transitions in a time interval of length T, then the transition density of $x(t)$ is defined as:*

$$D(x) = \lim_{T\to\infty} \frac{n_x(T)}{T} \qquad (3)$$

The density provides an effective measure of the switching activity in logic circuits in the presence of any delay model. If the density at every circuit node is made available, the overall average power dissipation in the circuit can be computed as:

$$P_{av} = \frac{1}{2} V_{dd}^2 \sum_{i=1}^{n} C_i D(x_i) \qquad (4)$$

In a synchronous circuit, with a clock period T_c, the relationship between transition density and transition probability is:

$$D(x) \geq \frac{P_t(x)}{T_c} \qquad (5)$$

where equality occurs in the zero-delay case. Thus the transition probability gives a lower bound on the transition density.

In order to complete the density formulation, another measure is required: Let $P(x)$ denote the *equilibrium probability* [6] of a logic signal $x(t)$, defined as the *average fraction of time that the signal is high.* Formally:

Definition 4. (equilibrium probability) *If $x(t)$ is a logic signal (switching between 0 and 1), then its equilibrium probability is defined as:*

$$P(x) = \lim_{T\to\infty} \frac{1}{T} \int_{\frac{-T}{2}}^{\frac{+T}{2}} x(t)dt \qquad (6)$$

In contrast to the signal probability, the equilibrium probability depends on the circuit internal delays since it describes the signal behavior over time, not only its steady state behavior per clock cycle. In the zero-delay case, the equilibrium probability reduces to the signal probability.

The combination of correlation and delay compounds the problem. We now have the situation where even inside a single clock cycle, we must model the correlation between signals internal to a combinational circuit in both space *and* time. If these correlations are completely ignored, so that any two signals are completely independent both in space and time, we say that we have a *spatio-temporal independence assumption.* If this is assumed, then the transition density at the output y of a Boolean logic cell (gate) can be easily computed [6] from the density at its inputs, $x_1, \ldots, x_n$, according to:

$$D(y) = \sum_{i=1}^{n} P\left(\frac{\partial y}{\partial x_i}\right) D(x_i) \qquad (7)$$

where $\partial y/\partial x$ is the *Boolean difference* of y with respect to x, defined as:

$$\frac{\partial y}{\partial x} = y|_{x=1} \oplus y|_{x=0} \qquad (8)$$

where $\oplus$ denotes the exclusive-or operation. Due to the spatio-temporal independence assumption, it turns out that the probability of a simultaneous transition at two inputs x_i and x_j is 0. This feature of the density formulation may be acceptable inside the circuit, but not at the primary inputs. This issue of simultaneous switching is dealt with in one of the papers in this session [7], which considers both the zero and unit-delay cases.

V. Statistical Techniques

In the above description of combinational circuit power estimation, we have restricted our attention to the so-called *probabilistic techniques.* This was done in order to properly introduce the papers in this session. We call an approach *probabilistic* when it is based on propagating a probability measure directly through the logic. Alternative approaches are possible [1], that we call *statistical.* Statistical power estimation techniques are essentially Monte Carlo methods that are based on statistical sampling: Apply randomly generated input vectors to the circuit, and monitor the cumulative value of power dissipated, using a standard (logic or timing) simulator. Continue this until the monitored power converges. Even though it uses a standard simulator (strongly pattern-dependent), the required input information is only the signal and transition probabilities at the circuit inputs, so that the technique is effectively *weakly* pattern-dependent. The required probabilities can be obtained from the prior analysis step of the FSM.

These techniques have many advantages, such as predictable error, no internal independence assumptions, and ease of use. For more details the reader is referred to [1] and [11, 12]. At the level of the FSM, one of the papers in this session [10] uses a statistical technique to analyze the FSM and measure the latch output probabilities.

VI. Summary

In this paper, we have discussed power estimation techniques, and highlighted problems related to issues of feedback, signal correlation, and circuit delays. Logic signals are modeled using probabilities in order to allow one to efficiently model a large set of input vectors, leading to a weakly pattern-dependent approach. However, computing the probabilities of internal signals is not easy. The key issue is whether the signals are independent or not. An independence assumption is very attractive because it greatly simplifies the analysis. But it breaks down in practice for many reasons. One reason is feedback, which creates correlation between signal values in consecutive clock cycles. Another is reconvergent fanout, which creates correlation between internal signals even if the primary inputs are independent. A final complicating issue is circuit delay which leads to multiple transitions per clock cycle inside a combinational circuit. Other papers in this session will deal with all three issues.

References

[1] F. Najm, "A survey of power estimation techniques in VLSI circuits," *IEEE Transactions on VLSI Systems*, pp. 446–455, Dec. 1994.

[2] H. J. M. Veendrick, "Short-circuit dissipation of static CMOS circuitry and its impact on the design of buffer circuits," *IEEE Journal of Solid-State Circuits*, vol. SC-19, no. 4, pp. 468–473, Aug. 1984.

[3] K. P. Parker and E. J. McCluskey, "Probabilistic treatment of general combinational networks," *IEEE Transactions on Computers*, vol. C-24, pp. 668–670, June 1975.

[4] A. Shen, A. Ghosh, S. Devadas, and K. Keutzer, "On average power dissipation and random pattern testability of CMOS combinational logic networks," IEEE/ACM International Conference on Computer-Aided Design, pp. 402–407, Santa Clara, CA, November 8–12, 1992.

[5] F. Najm, "Transition density, a stochastic measure of activity in digital circuits," *28th ACM/IEEE Design Automation Conference*, San Francisco, CA, pp. 644–649, June 17–21, 1991.

[6] F. Najm, "Transition density : a new measure of activity in digital circuits," *IEEE Transactions on Computer-Aided Design*, vol. 12, no. 2, pp. 310–323, February 1993.

[7] H. Mehta, M. Borah, R. M. Owens, and M. J. Irwin, "Accurate estimation of combinational circuit activity," *ACM/IEEE 32nd Design Automation Conference*, 1995.

[8] F. Najm and M. Zhang, "Extreme delay sensitivity and the worst-case switching activity in VLSI circuits," *Proc. ACM/IEEE 32nd Design Automation Conference*, 1995.

[9] R. Marculescu, D. Marculescu, and M. Pedram, "Efficient power estimation for highly correlated input streams," *ACM/IEEE 32nd Design Automation Conference*, 1995.

[10] F. Najm, S. Goel, and I. Hajj, "Power estimation in sequential circuits," *ACM/IEEE 32nd Design Automation Conference*, 1995.

[11] R. Burch, F. Najm, P. Yang, and T. Trick, "A Monte Carlo approach for power estimation," *IEEE Transactions on VLSI Systems*, vol. 1, no. 1, pp. 63–71, March 1993.

[12] M. Xakellis and F. Najm, "Statistical Estimation of the Switching Activity in Digital Circuits," *31st ACM/IEEE Design Automation Conference*, San Diego, CA, pp. 728–733, 1994.