

Statistical Estimation of Leakage-Induced Power Grid Voltage Drop Considering Within-Die Process Variations*

Imad A. Ferzli
 Department of ECE
 University of Toronto
 Toronto, Ontario, Canada
 ferzli@eec.utoronto.ca

Farid N. Najm
 Department of ECE
 University of Toronto
 Toronto, Ontario, Canada
 f.najm@utoronto.ca

ABSTRACT

Transistor threshold voltages (V_{th}) have been reduced as part of on-going technology scaling. The smaller V_{th} values feature increased fluctuations due to process variations, with a strong within-die component. Correspondingly, given the exponential dependence of leakage on V_{th} , circuit leakage currents are increasing significantly and have strong within-die statistical variations. With these currents loading the power grid, the grid develops large voltage drops, which is an unavoidable background level of noise on the grid. We develop techniques for estimation of the statistics of the leakage-induced power grid voltage drop based on given statistics of the circuit leakage currents.

Categories and Subject Descriptors

B.7 [Integrated Circuits]: Design Aids

General Terms

Design, Algorithms

Keywords

Power grid, leakage current, voltage drop, statistical analysis

1. INTRODUCTION

Technology scaling requires reduction of the MOSFET threshold voltage (V_{th}), to accompany the reduction in supply voltages (V_{dd}) and oxide thickness. In today's 1.2 V, 0.13 μm technology, V_{th} is about 0.3 V and is forecast to be further reduced at the rate of 15% per generation in future [2]. This reduced V_{th} comes with increased variations, due to underlying process variations [11, 7]. Considering sub-threshold leakage currents, and due to the exponential dependence of these currents on threshold voltage, we also get a high rate of increase (reportedly as high as about 5X per generation [2]) along with large variability (close to 3X for individual gates [11] and 20X for whole chips [4]). Moreover, these variations are known to have a significant within-die component [1, 7], so that transistors in close proximity on the layout can have significant variations in their leakage currents.

In this work, we consider the fact that, in response to statistical leakage currents, the power grid of VLSI chips exhibits a

*This research was supported by Micronet, with funding from ATI Technologies Inc. and from Altera Corp.

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.

DAC 2003, June 2-6, 2003, Anaheim, California, USA.
 Copyright 2003 ACM 1-58113-688-9/03/0006 ...\$5.00.

voltage drop at all the nodes that is statistically variable with a strong within-die component. This voltage drop is unavoidable and manifests itself as a *background level of noise* on the grid which will have an effect on circuit delay and operation and must be considered during circuit design. However, in order to account for this noise, simply setting all leakage currents to their maximum values is too pessimistic because strong local variations in V_{th} make this case highly improbable. Instead, one must consider all leakage currents to be random and focus on their independent components, consistent with the fact that within-die variations are significant. We will show that die-to-die variations are easily dealt with, and will focus on within-die variations, for which we propose an efficient approach to directly estimate the statistics of the voltages from the given statistics of the currents.

Except for [11] and [7], there is not much prior work on the study of statistical leakage currents due to within-die process variations. There is no work that the authors are aware of on the study of the statistical grid voltages in response to the statistical leakage currents.

2. PROBLEM FORMULATION

It is helpful to distinguish between two types of leakage in integrated circuits. A circuit certainly draws leakage current when it is in standby or sleep mode, what may be referred to as the *standby leakage*. The circuit also draws leakage current when it is active. Indeed, a logic gate draws leakage current any time that its supply is "on." Even inside a switching window, part of the current drawn from the supply may be attributed to leakage. The leakage drawn by the circuit during its active (non-standby) states, may be referred to as the *dynamic leakage*. The grid response to standby leakage may be obtained by a DC analysis of the grid, using only a resistive model, whereas response to dynamic leakage requires a transient analysis, using an RC or RLC model of the grid.

2.1 System Equations

We consider an RC model of the power grid, where each branch of the grid is represented by a resistor and where there exists a capacitor from every grid node to ground. In addition, some nodes have ideal current sources (to ground) representing the current drawn by the circuit tied to the grid at that point, and some nodes have ideal voltage sources (to ground) representing the connections to the external voltage supply. Let the power grid consist of $N + p$ nodes, where nodes $1, 2, \dots, N$ have no voltage sources attached, and nodes $(N + 1), (N + 2), \dots, (N + p)$ are the nodes with the voltage sources. Let c_k be the capacitance from every node k to ground. Let $i_k(t)$ be the current source connected to node k , where the direction of positive current is from the node to ground. We assume that $i_k(t) \geq 0$ and that $i_k(t)$ is defined for every node $k = 1, \dots, N$ so that nodes with no current source attached have $i_k(t) = 0, \forall t$. Let $\mathbf{i}(t)$ be the vector of all $i_k(t)$ sources, $u_k(t)$ be the voltage at node k , and $\mathbf{u}(t)$ be the vector of all $u_k(t)$ signals. Applying Modified Nodal Analysis (MNA) [10]

leads to:

$$\mathbf{G}\mathbf{u}(t) + \mathbf{C}\dot{\mathbf{u}}(t) = -\mathbf{i}(t) + \mathbf{G}\mathbf{V}_{dd} \quad (1)$$

where \mathbf{G} is an $N \times N$ conductance matrix, \mathbf{C} is an $N \times N$ diagonal matrix of node capacitances, and \mathbf{V}_{dd} is a constant vector each entry of which is equal to V_{dd} . Let $v_k(t) = V_{dd} - u_k(t)$ be the voltage drop at node k , and let $\mathbf{v}(t)$ be the vector of voltage drops, then (1) can be written as:

$$\mathbf{G}\mathbf{v}(t) + \mathbf{C}\dot{\mathbf{v}}(t) = \mathbf{i}(t) \quad (2)$$

This is a *revised* system equation which one can solve directly for the voltage drop values. Notice that the circuit described by (2) consists of the original power grid, but with all the voltage sources set to zero and all the current source directions reversed. In the following, we will mainly be concerned with this *modified power grid* and the revised system of equations (2).

In cases when the circuit is in a standby state, where all the currents are constant, the circuit response is obtained using a DC analysis. The DC equivalent of (2) is readily seen as:

$$\mathbf{G}\mathbf{V} = \mathbf{I} \quad (3)$$

2.2 Current Statistics

By virtue of the underlying process variations, the leakage currents exhibit two kinds of variations: within-die variations (also referred to as intra-chip or intra-die variations) and die-to-die variations (also referred to as inter-chip or inter-die variations). As is typically done with other types of parameter variations [12, 8], we can express each (random) current vector as the sum of two (random) current vectors that separately capture the die-to-die (i_{dd}) and within-die (i_{wd}) variations:

$$\mathbf{i}(t) = \mathbf{i}_{dd}(t) + \mathbf{i}_{wd}(t) \quad (4)$$

At each time point, the components of $\mathbf{i}_{dd}(t)$ are all correlated, so that they are all functions of a single underlying random variable that takes the same value all over the die. However, the components of $\mathbf{i}_{wd}(t)$ are all *independent* and have zero-mean. The same breakdown of currents can be done in the DC currents case.

Since the grid is a linear system, then by superposition, the voltage response of the power grid to the total current can be composed as the sum of its responses to the two types of random currents in (4). Note that the voltage drop on each node increases when the currents loading the grid increase [5], so that it becomes very easy to study the response due to $\mathbf{i}_{dd}(t)$, by setting all the components at the top of their range (at the mean + 3σ point) and solving the grid once, then setting them at the bottom of their range and solving the grid again. This gives the total spread in voltage values due to die-to-die variations.

Handling within-die variations is more difficult, primarily because the currents are all independent random variables. The rest of this paper will be focused on within-die variations. Precisely, given the means and variances of the individual independent leakage currents, our problem is to estimate the means and variances (and possibly covariances) of the grid voltages.

2.3 Mean Estimation

Due to linearity of the mean ($E[\cdot]$) operator, one can write:

$$\mathbf{G}E[\mathbf{v}(t)] + \mathbf{C}\frac{d}{dt}E[\mathbf{v}(t)] = E[\mathbf{i}(t)] \quad (5)$$

Thus, if we solve the system (2) once, using simply the current means as inputs, the solution gives the voltage means at all the nodes. The rest of this paper is focused on estimation of the node voltage variances.

3. VARIANCE ESTIMATION

Variance is inherently harder to estimate than the mean. We have no good way to formulate the system equations in as-simple a form as (5) that maintains the full dynamic model of the grid. Instead, we have found it necessary to introduce the following “*pseudo-static*” assumption:

Assumption 1. For variance estimation, we assume that the grid may be solved as a DC system at every/any time point.

This is purely a simplifying assumption which helps arrive at an efficient formulation for the variance computation. Notice that this assumption is automatically true for the case of standby leakage. For dynamic leakage, since the leakage current of a logic gate is constant when it is not switching, then this assumption may be acceptable in practice, especially since *some* dynamics of the system have been included in the analysis (through the computation of the mean response as per (5)).

One way of estimating the (output) distributions of voltage given the (input) distributions of current is to simply do random sampling on the currents. Generate a randomly-chosen vector of current values, according to the distributions of the currents, and solve the system (3) repeatedly to get a corresponding sample of the voltages, until the desired statistics have been estimated with sufficient accuracy. This technique is viable, and we have actually implemented it and found that, in some cases, it requires fewer samples than the column-sampling approach to be proposed below. However, it has a serious disadvantage in that when checking whether a certain estimator of the variance has converged, one is required to evaluate certain percentiles of the χ^2 distribution for every sample. This can get very expensive because it requires numerical integration of the χ^2 pdf (probability density function) - using a table is not a practical option because the number of samples can be huge. In contrast, the column-sampling method (below) transforms the variance estimation problem to a mean estimation problem, which requires estimating (only once) a percentile of the standard normal pdf. Thus, we have found that the current-sampling technique takes much longer time to converge (about 10-20X more) than the column-sampling method presented below, which is our main result for this paper.

3.1 Variance Computation

Under DC conditions, the system (5) is simplified to:

$$\mathbf{G}E[\mathbf{V}] = E[\mathbf{I}] \quad (6)$$

We now combine (3) and (6) to yield:

$$\mathbf{G}(\mathbf{V} - E[\mathbf{V}]) = \mathbf{I} - E[\mathbf{I}] \quad (7)$$

Multiplying each side by its transpose and applying the expected value operator to each side, leads to:

$$\begin{aligned} \mathbf{G}E[(\mathbf{V} - E[\mathbf{V}])(\mathbf{V} - E[\mathbf{V}])^T]\mathbf{G}^T = \\ E[(\mathbf{I} - E[\mathbf{I}])(\mathbf{I} - E[\mathbf{I}])^T] \end{aligned} \quad (8)$$

We recognize the expectations as being simply covariance matrices [9], so that the above result can be rewritten as:

$$\mathbf{G}\text{Cov}(\mathbf{V})\mathbf{G}^T = \text{Cov}(\mathbf{I}) \quad (9)$$

Since \mathbf{G} is symmetric, so is \mathbf{G}^{-1} . Therefore, (9) becomes:

$$\text{Cov}(\mathbf{V}) = \mathbf{G}^{-1}\text{Cov}(\mathbf{I})\mathbf{G}^{-1} \quad (10)$$

Generating and storing $\text{Cov}(\mathbf{V})$ is computationally very expensive for large grids. Furthermore, one can envision the need for the covariance of a few nodes in a small vicinity on the layout, but it is not clear how covariances among all pairs of nodes would be useful. Indeed, it may be argued that the only information of practical value today is the variances of the node voltages, *i.e.*, the diagonal of $\text{Cov}(\mathbf{V})$. Therefore, we will derive a simpler system based on the variances of the voltage drops, rather than their entire covariance matrix. Noting that the independence of the currents means $\text{Cov}(\mathbf{I})$ is diagonal, and since \mathbf{G}^{-1} is symmetric, it can be seen from (10) that:

$$\begin{aligned} [\text{Cov}(\mathbf{V})]_{ii} = & ([\mathbf{G}^{-1}]_{i1})^2 [\text{Cov}(\mathbf{I})]_{11} \\ & + \dots + ([\mathbf{G}^{-1}]_{iN})^2 [\text{Cov}(\mathbf{I})]_{NN}, \end{aligned} \quad (11)$$

where $[\cdot]_{ij}$ is the (i, j) th entry of the corresponding matrix. Let $\text{Var}(\mathbf{V})$ and $\text{Var}(\mathbf{I})$ be respectively the *vectors* of variances of all voltages and currents, and define a matrix $\mathbf{G}^{-1(2)}$ such that:

$$[\mathbf{G}^{-1(2)}]_{ij} = ([\mathbf{G}^{-1}]_{ij})^2, \quad (12)$$

so that (11) can be written as:

$$\boxed{\text{Var}(\mathbf{V}) = \mathbf{G}^{-1(2)} \text{Var}(\mathbf{I})} \quad (13)$$

The solution of (13) requires full knowledge of the inverse of \mathbf{G} , which is impractical for large grids. In the following, we propose an efficient random-sampling-based technique which uses (13) to estimate the variances of the node voltage drops.

3.2 A Column-Sampling Approach

In order to simplify the notation, let $r_{ij} = [\mathbf{G}^{-1(2)}]_{ij}$. Also, let $\sigma_{I_i}^2$ be the variance of the current source at node i and $\sigma_{V_i}^2$ be the variance of the voltage drop at node i . Let $S = \sum_{i=1}^N \sigma_{I_i}^2$ and let $p_i = \sigma_{I_i}^2/S$. Then, using (13) yields:

$$\sigma_{V_i}^2 = S \sum_{j=1}^N p_j r_{ij} \quad (14)$$

Since $\sum_{i=1}^N p_j = 1$, then we can view the p_j weights as being *probability values* that are associated with the r_{ij} values, so that the summation above becomes the mean (weighted average) of all the r_{ij} elements in the i th row. If we define a random variable (RV) \mathbf{r}_i as being a discrete RV that takes the values r_{ij} with probabilities p_j , $j = 1, 2, \dots, N$, then we can write (14) as:

$$\sigma_{V_i}^2 = SE[\mathbf{r}_i] \quad (15)$$

Let the mean of \mathbf{r}_i be $\mu_i = E[\mathbf{r}_i]$ and its variance be σ_i^2 . We can now use methods of *mean estimation* from statistics, basically Monte Carlo random sampling [6, 13], to estimate the *population mean* μ_i using the mean of a much smaller sample (say, of size $n \ll N$) from the population, *i.e.*, using the *sample mean*.

Using a weighted random number generator, we generate according to the probabilities p_j a sequence of indices of columns of $\mathbf{G}^{-1(2)}$ to form the following sample mean for every row i :

$$\bar{r}_i = \frac{1}{n} \sum_{j \in \mathcal{J}} r_{ij} \quad (16)$$

where \mathcal{J} is the set of indices included in the random sample. We also compute the *sample standard deviation*, $s_i \geq 0$ given by:

$$s_i^2 = \frac{1}{n-1} \sum_{j \in \mathcal{J}} (r_{ij} - \bar{r}_i)^2 = \frac{n \left(\sum_{j \in \mathcal{J}} r_{ij}^2 \right) - \left(\sum_{j \in \mathcal{J}} r_{ij} \right)^2}{n(n-1)} \quad (17)$$

It can be shown that, for large n (as $n \rightarrow \infty$), \bar{r}_i converges to μ_i . Furthermore, for finite n , it is possible to determine [6, 13] how large n should be in order for \bar{r}_i to be a good approximation (to within arbitrary accuracy) of μ_i , as follows.

Considering all possible samples of size n from a given population, one can view \bar{r}_i itself as a random variable, with mean μ_i and variance σ_i^2/n . For moderately large values of n (larger than 30 [6] or 50 [13]), the distribution of \bar{r}_i becomes close to a normal distribution. This means that the distribution of $(\bar{r}_i - \mu_i)/(\sigma_i/\sqrt{n})$ becomes close to a standard normal distribution. Thus, it becomes possible to use tables of the standard normal to establish how large n should be in order for the variance of \bar{r}_i to be small enough for it to be a viable estimator of μ_i , with a certain level of confidence [6, 13]. For example, if it is desired to have $(1 - \alpha) \times 100\%$ confidence (where α is a small positive number, $0 < \alpha < 1$) that the following is true:

$$|\bar{r}_i - \mu_i| < \epsilon \quad (18)$$

Then n should be larger than n_0 where:

$$n_0 = \left(\frac{z_{\alpha/2} \sigma_i}{\epsilon} \right)^2 \quad (19)$$

where $z_{\alpha/2}$ is such that the area to the right of it under the pdf of the standard normal curve is equal to $\alpha/2$ (*e.g.*, for 90% confidence, $\alpha = 0.05$ and $z_{\alpha/2} = 1.96$). Note that the strict use of (19) requires knowledge of σ_i . In practice, however, one samples until n is larger than 30 or 50 or so, then starts to use (19), substituting s_i for σ_i [6, 13], to monitor convergence, leading to:

$$n_0 = \left(\frac{z_{\alpha/2} s_i}{\epsilon} \right)^2 \quad (20)$$

3.3 Error Bound

In order to choose a meaningful error bound ϵ for use in (20), recall that our intent is to approximate the variance of every node voltage as follows:

$$\sigma_{V_i}^2 \approx S \bar{r}_i \quad (21)$$

Let δ be a small positive number, $0 < \delta < 1$. It makes sense to aim to achieve a user-defined error bound on σ_{V_i} defined relative to the supply voltage, V_{dd} , as follows:

$$\left| \sigma_{V_i} - \sqrt{S \bar{r}_i} \right| \leq \delta V_{dd} \quad (22)$$

In other words, we want to find ϵ (to be used in (20)) as a function of δ in order for the following to be true:

$$|\bar{r}_i - \mu_i| < \epsilon \implies \left| \sqrt{S \bar{r}_i} - \sqrt{S \mu_i} \right| < \delta V_{dd} \quad (23)$$

To simplify the notation, let $x = \mu_i$ and $x_0 = \bar{r}_i$. Also, let $y = \sqrt{\mu_i} = \sqrt{x}$ and $y_0 = \sqrt{\bar{r}_i} = \sqrt{x_0}$, and let $\gamma = \delta V_{dd}/\sqrt{S}$. Notice that $\gamma > 0$. We want to find ϵ in terms of δ so that:

$$|x - x_0| < \epsilon \implies |y - y_0| < \gamma \quad (24)$$

There are two cases to consider, according to whether $y_0 < \gamma$ or $y_0 > \gamma$. When $y_0 < \gamma$, then (since $y > 0$ in all cases) in order to guarantee that $|y - y_0| < \gamma$, it is sufficient to impose that $y < y_0 + \gamma$. Since y is an increasing function of x , then we require that: $|x - x_0| = x - x_0 < (y_0 + \gamma)^2 - y_0^2 = 2y_0\gamma + \gamma^2$, so $\epsilon = 2y_0\gamma + \gamma^2$.

When $y_0 > \gamma$, then we need to have $y_0 - \gamma < y < y_0 + \gamma$, which implies that $(y_0 - \gamma)^2 < x < (y_0 + \gamma)^2$. This translates to requiring that: $|x - x_0| < 2y_0\gamma - \gamma^2$, so $\epsilon = 2y_0\gamma - \gamma^2$.

Notice that the condition $y_0 > \gamma$ translates to $\bar{r}_i > \delta^2 V_{dd}^2/S$. In summary, then:

$$\epsilon = \begin{cases} \frac{\delta V_{dd}}{S} (2\sqrt{S \bar{r}_i} + \delta V_{dd}) & \text{if } \bar{r}_i < \delta^2 V_{dd}^2/S, \\ \frac{\delta V_{dd}}{S} (2\sqrt{S \bar{r}_i} - \delta V_{dd}) & \text{if } \bar{r}_i > \delta^2 V_{dd}^2/S. \end{cases} \quad (25)$$

Plugging (25) into (20) leads to:

$$n_0 = \left(\frac{z_{\alpha/2} s_i S / \delta V_{dd}}{\delta V_{dd} \pm 2\sqrt{S \bar{r}_i}} \right)^2 \quad (26)$$

where the $+$ or $-$ sign depends on whether \bar{r}_i is smaller or larger than $\delta^2 V_{dd}^2/S$, respectively.

4. EXPERIMENTAL RESULTS

This method has been implemented on a number of test-case grids. Not having access to power grids from industrial designs, we have opted to *generate* a number of grids ourselves. Starting with a square uniform grid of a given size, we randomly delete a user-specified percentage of nodes, which we call the *degree of non-uniformity*. Geometric and physical grid characteristics (*e.g.* grid dimensions) as well as characteristics of the fabrication process (*e.g.* sheet resistances) lead to initial values of branch conductances. When a node is deleted, the conductances of the remaining surrounding branches are increased by a random amount in order to allow the non-uniform grid to be loaded with currents

Table 1: Results on grids that allow an exact solution.

Size (#nodes)	Number of C4s	Degree (%) of non-uniformity	Avg. Error (% of V_{dd})	Max. Error (% of V_{dd})	Number of samples	CPU time (LU)	CPU time (sampling)	CPU time (exact)
17,678	150	3	0.07	0.73	5,559	21 sec.	6 min.	17 min.
38,800	170	3	0.09	0.65	7,381	2 min.	24 min	2 hrs.
41,058	170	19	0.08	0.63	20,523	45 sec.	41 min.	1.2 hrs
41,095	170	19	0.09	0.84	17,725	51 sec.	37 min.	1.3 hr.
41,147	170	19	0.09	0.78	20,068	46 sec.	41 min.	1.2 hr.
64,067	260	3	0.08	0.73	15,092	6 min.	1.6 hr.	6.5 hrs.
76,503	260	19	0.09	0.92	28,829	3 min	2.3 hrs.	5.6 hrs.
76,522	260	19	0.09	1.07	28,526	3 min.	2.2 hrs.	5.4 hrs.
81,000	350	10	0.07	0.60	11,236	5 min.	1.2 hrs.	8.1 hrs.
110,776	400	8	0.05	0.43	18,501	13 min.	3.4 hrs.	18.8 hrs.
138,500	450	8	0.06	0.51	14,589	18 min.	3.4 hrs.	29.6 hrs.

Table 2: Results on larger grids that do not allow an exact solution.

Size (#nodes)	Number of C4s	Degree (%) of non-uniformity	Number of samples	CPU time (LU)	CPU time (sampling)	Memory usage
164,510	550	23	24,314	11 min.	3.1 hrs.	163 MB
203,251	580	6	22,124	35 min.	5.6 hrs.	302 MB
251,856	590	10	31,052	39 min.	8.6 hrs.	335 MB
309,719	620	4	34,346	1.5 hrs.	16.6 hrs.	558 MB




Figure 1: Correlation plot of results for a grid with 138,500 nodes

that are comparable to those applied to its uniform predecessor, while exhibiting comparable IR-drops. The number of V_{dd} (C4) sites and leakage current sources are supplied by the user; the C4s and current sources are then distributed at random. Finally, the statistical estimation was conducted at 90% confidence ($\alpha = 0.1$) and with a 1% error bound ($\delta = 0.01$).

Table 1 shows results for a number of power grids that are small enough to allow for the exact solution of (13), in order to assess the accuracy of our column-sampling approach. These results were obtained on a Sun Blade 1000 with a 750MHz CPU and 2.5 GB of RAM. Clearly, column-sampling does a very good job; the average error is well below the specified 1% threshold, and only 1 of 11 cases has a maximum error above 1%. With an error threshold of 1% and a confidence of 90%, one would expect the error to exceed 1% with only 10% probability, which is exactly what we find. For the grid with 138,500 nodes, the correlation plot of standard deviations is given in Fig. 1, showing very good accuracy. Results on larger grids are shown in Table 2, obtained on a Dell machine with a 1.6GHz Intel Pentium-4 CPU and 2.0GB of RAM. Only the proposed column-sampling approach was applied to these grids, and run-time and memory usage data are shown. Clearly, this technique is applicable to relatively large grids, with a tight ($\delta = 1\%$, in this case) error bound.

5. CONCLUSION

Due to reduced threshold voltages (V_{th}), leakage currents are much higher than before, and due to increasing V_{th} variations, which exhibit a strong within-die component, leakage currents are

statistically variable, with a strong within-die component. The effect of these currents on the power grid is to generate a statistical background-noise voltage drop on the grid which needs to be accounted for before a power grid can be deemed safe. This voltage drop reduces the allowable dynamic voltage drop that may be tolerated on a power grid due to normal circuit operation. We have presented an efficient approach by which given statistics of the leakage currents can be used to derive statistics (means and variances) of the grid voltages.

6. REFERENCES

- [1] 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. In *ISSCC*, pages 278–279, 2001.
- [2] V. De and S. Borkar. Technology and design challenges for low power and high performance. In *ISLPED*, pages 163–168, 1999.
- [3] T. Kam, S. Rawat, D. Kirkpatrick, R. Roy, G. S. Spirakis, N. Sherwani, and C. Peterson. EDA challenges facing future microprocessor design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(12):1498–1506, Dec. 2000.
- [4] T. Karnik, S. Borkar, and V. De. Sub-90nm technologies - challenges and opportunities for CAD. In *ICCAD*, pages 203–206, 2002.
- [5] H. Kriplani, F. N. Najm, and I. Hajj. Pattern independent maximum current estimation in power and ground buses of CMOS VLSI circuits: algorithms, signal correlations, and their resolution. *IEEE Transactions on Computer-Aided Design*, 14(8):998–1012, August 1995.
- [6] I. R. Miller, J. E. Freund, and R. Johnson. *Probability and Statistics for Engineers*. Prentice-Hall, Inc., 4th edition, 1990.
- [7] S. Narendra, V. De, S. Borkar, D. Antoniadis, and A. Chandrakasan. Full-chip sub-threshold leakage power prediction model for sub-0.18um CMOS. In *ISLPED*, pages 19–23, 2002.
- [8] S. R. Nassif. Within-chip variability analysis. In *IEDM*, pages 283–286, 1998.
- [9] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 2nd edition, 1984.
- [10] L. T. Pillage, R. A. Rohrer, and C. Visweswariah. *Electronic Circuit and System Simulation Methods*. McGraw-Hill, 1995.
- [11] A. Srivastava, R. Bai, D. Blaauw, and D. Sylvester. Modeling and analysis of leakage power considering within-die process variations. In *ISLPED*, pages 64–67, 2002.
- [12] B. E. Stine, D. S. Boning, and J. E. Chung. Analysis and decomposition of spatial variation in integrated circuit processes and devices *IEEE Transactions on Semiconductor Manufacturing*, 10(1):24–41, Feb. 1997.
- [13] S. K. Thompson. *Sampling*. John Wiley & Sons, Inc., 2nd edition, 2002.