The Certainty of Uncertainty: Randomness in Nanometer Design

Hongliang Chang†, Haifeng Qian‡, and Sachin S. Sapatnekar‡

† CSE Department ‡ ECE Department
University of Minnesota, Minneapolis, MN 55455, USA

Abstract. Randomness and uncertainty are rearing their heads in surprising and contradictory ways in nanometer technologies. On the one hand, uncertainty and variability is becoming a dominant factor in the design of integrated circuits, and on the other hand, algorithms based on randomness are beginning to show great promise in solving large scale problems. This paper overviews both aspects of this issue.

1 Introduction

A historical look at integrated circuit technologies shows several inflection points that have characterized the 250nm, 180nm and 130nm nodes. The move to the sub-100nm regime is projected to bring about the most revolutionary of these changes, in terms of how it impacts the way in which design is carried out. Most notably, randomness will become a fact of life that designers will be forced to confront, and perhaps, paradoxically, the only certainty in nanometer designs will be the presence of uncertainty. Several issues related to uncertainty and randomness will be discussed in this paper.

We will begin, in Section 2, by exploring the origins of randomness in nanometer circuits, and will then discuss methods that must be used in next-generation designs to handle such variations in Section 3. This first aspect of randomness, caused by process and environmental variations, is "problematic" and requires new solutions to overcome its effects, since such variations manifest themselves as changes in the delay and power dissipation of a circuit. As a consequence, the analysis of timing will move from a purely deterministic setting to a statistical analysis, as will the analysis of leakage power, which is becoming a major component of the total power dissipation. This has already lead to intense efforts in statistical static timing analysis (SSTA) and statistical power analysis in recent years. Finding efficient solutions to these problems presents numerous new challenges, and while some first steps have been taken, many problems remain unsolved.

Amid all these problems also lies an opportunity: there is a second facet of randomness that is likely to have very positive consequences in the future, as discussed in Section 4. As the electronic design automation world becomes more educated in the use of stochastic techniques, new opportunities will arise on the algorithmic side, as novel statistical approaches will be developed for solving

design problems. This has already been set into motion: problems as diverse as capacitance extraction, power estimation, Vdd net analysis, crosstalk analysis, placement, and ESD analysis are seeing viable stochastic solution techniques. An attractive feature of the random techniques is that when used in appropriate settings, they can scale extremely well with increasing problem sizes, and for several problems, they have the potential for localized computation. This paper will overview such algorithms and raise the challenge of harnessing the power of such methods for solving the problems of tomorrow.

2 Sources of Uncertainty

Current-day integrated circuits are afflicted with a wide variety of variations that affect their performance. Essentially, under true operating conditions, the parameters chosen by the circuit designer are perturbed from their nominal values due to various types of variations. As a consequence, a single SPICE-level transistor or interconnect model (or an abstraction thereof) is seldom an adequate predictor of the exact behavior of a circuit. These sources of variation can broadly be categorized into two classes

Process variations result from perturbations in the fabrication process, due to which the nominal values of parameters such as the effective channel length (L_{eff}), the oxide thickness (t_{ox}), the dopant concentration (N_a), the transistor width (w), the interlayer dielectric (ILD) thickness (t_{ILD}), and the interconnect height and width (h_{int} and w_{int}, respectively).

Environmental variations arise due to changes in the operating environment of the circuit, such as the temperature or variations in the supply voltage (V_{dd} and ground) levels. There is a wide body of work on analysis techniques to determine environmental variations, both for thermal issues [8, 7, 20, 10], and for supply net analysis [18].

Both of these types of variations can result in changes in the timing and power characteristics of a circuit.

Process variations can also be classified into the following categories:

Inter-die variations are the variations from die to die, and affect all the devices on same chip in the same way, e.g., they may cause all of the transistor gate lengths of devices on the same chip to be larger or all of them to be smaller.

Intra-die variations correspond to variability within a single chip, and may affect different devices differently on the same chip, e.g., they may result in some devices having smaller oxide thicknesses than the nominal, while others may have larger oxide thicknesses.

Inter-die variations have been a longstanding design issue, and for several decades, designers have striven to make their circuits robust under the unpredictability of such variations. This has typically been achieved by simulating the design at not just one design point, but at multiple "corners." These corners are

chosen to encapsulate the behavior of the circuit under worst-case variations, and have served designers well in the past. In nanometer technologies, designs are increasingly subjected to numerous sources of variation, and these variations are too complex to capture within a small set of process corners.

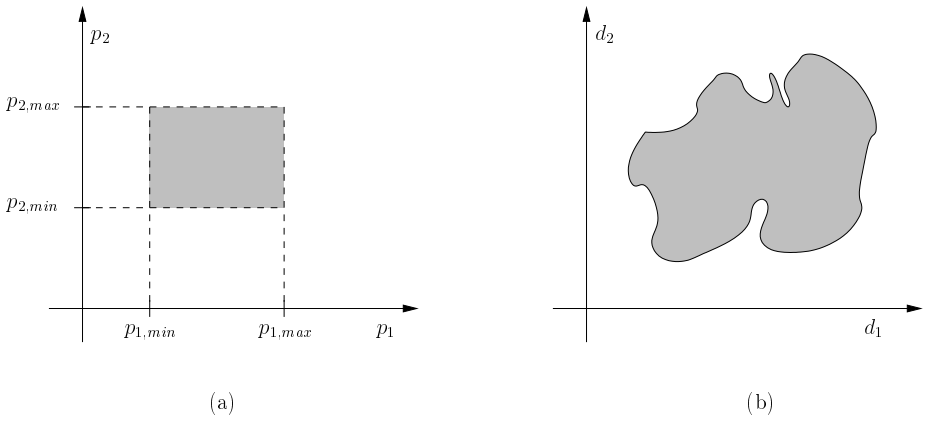

Fig. 1. The feasible region in (a) the performance parameter space and (b) the design/process parameter space.

To illustrate this, consider the design of a typical circuit. The specifications on the circuit are in the form of limits on performance parameters, p_i, such as the delay or the static or dynamic power dissipation, which are dependent on a set of design or process parameters, d_i, such as the transistor width or the oxide thickness. In Figure 1(a), we show the behavior of a representative circuit in the performance space of parameters, p_i, whose permissible range of variations lies within a range of $[p_{i,min}, p_{i,max}]$ for each parameter, p_i, which corresponds to a rectangular region. However, in the original space of design parameters, d_i, this may translate into a much more complex geometry, as shown in Figure 1(b). This may conservatively be captured in the form of process corners at which the circuit is simulated.

In nanometer technologies, intra-die variations have become significant and can no longer be ignored. As a result, a process corner based methodology, which would simulate the entire chip at a small number of design corners, is no longer sustainable. A true picture of the variations would use one process corner in each region of the chip, but it is clear that the number of simulations would increase exponentially with the number of such regions. If a small number of process corners are to be chosen, they must be very conservative and pessimistic. For true accuracy, a larger number of process corners may be used, but this number may be too large to permit computational efficiency.

The sources of these variations may be used to create another taxonomy:

Random variations (as the name implies) depict random behavior that can be characterized in terms of a distribution. This distribution may either be explicit, in terms of a large number of samples provided from fabrication line measurements, or implicit, in terms of a known probability density function (such as a Gaussian or a lognormal distribution) that has been fitted to the measurements. Random variations in some process or environmental parameters (such as those in the temperature, supply voltage, or L_{eff}) can often show some degree of local *spatial correlation*, whereby variations in one transistor in a chip are remarkably similar in nature to those in spatially neighboring transistors, but may differ significantly from those that are far away. Other process parameters (such as t_{ox} and N_a) do not show much spatial correlation at all, so that for all practical purposes, variations in neighboring transistors are uncorrelated.

Systematic variations show predictable variational trends across a chip, and are caused by known physical phenomena during manufacturing. Strictly speaking, environmental changes are entirely predictable, but practically, due to the fact that these may change under a large number (potentially exponential in the number of inputs and internal states) of operating modes of a circuit, it is easier to capture them in terms of random variations. Examples of systematic variations include those due to (i) spatial intra-chip gate length variability, also known as across-chip linewidth variation (ACLV), which observes systematic changes in the value of L_{eff} across a reticle due to effects such as changes in the stepper-induced illumination and imaging nonuniformity due to lens aberrations [15], and (ii) ILD variations due to the effects of chemical-mechanical polishing (CMP) on metal density patterns: regions that have uniform metal densities tend to have more uniform ILD thicknesses than regions that have nonuniformities.

The existence of *correlations* between intra-die variations complicates the task of statistical analysis. These correlations are of two types:

Spatial correlations To model the intra-die spatial correlations of parameters, the die region may be tesselated into n grids. Since devices or wires close to each other are more likely to have similar characteristics than those placed far away, it is reasonable to assume perfect correlations among the devices [wires] in the same grid, high correlations among those in close grids and low or zero correlations in far-away grids. Under this model, a parameter variation in a single grid at location (x, y) can be modeled using a single random variable $p(x, y)$. For each type of parameter, n random variables are needed, each representing the value of a parameter in one of the n grids.

Structural correlations The structure of the circuit can also lead to correlations that must be incorporated in SSTA. Consider the reconvergent fanout structure shown in Figure 2. The circuit has two paths, a-b-d and a-c-d. If, for example, we assume that each gate delay is a Gaussian random variable, then the PDF of the delay of each path is easy to compute, since it is the sum of Gaussians, which admits a closed form. However, the circuit delay is the maximum of the delays of these two paths, and these are correlated since

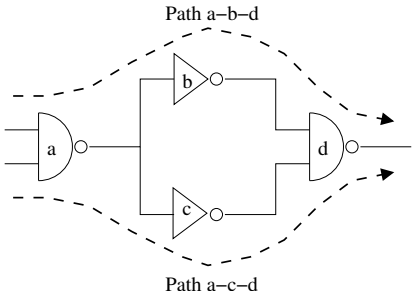

Fig. 2. An example to illustrate structural correlations in a circuit.

the delays of a and d contribute to both paths. It is important to take such structural correlations, which arise due to reconvergences in the circuit, into account while performing SSTA.

3 Analysis of Uncertainty

As an example, we will now illustrate the concepts involved in the statistical analysis of timing; similar techniques are being developed for power analysis.

The geometrical parameters associated with the gate and interconnect can reasonably be modeled as normally distributed random variables. Before we introduce how the distributions of gate and interconnect delays will be modeled, let us first consider an arbitrary function $d = f(\mathbf{P})$ that is assumed to be a function on a set of parameters $\mathbf{P}$, where each $p_i \in \mathbf{P}$ is a random variable with a normal distribution given by $p_i \sim N(\mu_{p_i}, \sigma_{p_i})$. We can approximate d linearly using a first order Taylor expansion:

$$d = d_0 + \sum_{\forall \text{ parameters } p_i} \left[\frac{\partial f}{\partial p_i}\right]_0 \Delta p_i \tag{1}$$

where d_0 is the nominal value of d, calculated at the nominal values of parameters in the set $\mathbf{P}$, $\left[\frac{\partial f}{\partial p_i}\right]_0$ is computed at the nominal values of p_i, $\Delta p_i = p_i - \mu_{p_i}$ is a normally distributed random variable and $\Delta p_i \sim N(0, \sigma_{p_i})$. The delay function here is arbitrary, and may include, for example, the effects of the input transition time on the gate/wire delay.

If all of the parameter variations can be modeled by Gaussian distributions, this approximation implies that d is a linear combination of Gaussians, which is therefore Gaussian. Its mean μ_d, and variance σ_d^2 are:

$$\mu_d = d_0 \tag{2}$$

$$\sigma_d^2 = \sum_{\forall i} \left[\frac{\partial f}{\partial p_i}\right]_0^2 \sigma_{p_i}^2 + 2 \sum_{\forall i \neq j} \left[\frac{\partial f}{\partial p_i}\right]_0 \left[\frac{\partial f}{\partial p_j}\right]_0 cov(p_i, p_j) \tag{3}$$

where $cov(p_i, p_j)$ is the covariance of p_i and p_j.

This approximation is valid when Δp_i has relatively small variations, in which domain the first order Taylor expansion is adequate and the approximation is acceptable with little loss of accuracy. This is generally true of the impact of intra-chip variations on delay, where the process parameter variations are relatively small in comparison with the nominal values, and the function changes by a small amount under this perturbation. For this reason, the gate and interconnect delays, as functions of the process parameters, can be approximated as a normal distributions when the parameter variations are assumed to be normal.

The existence of on-chip variations requires an extension of traditional STA techniques to move beyond their deterministic nature. The SSTA approach, which overcomes these problems, treats delays not as fixed numbers, but as probability density functions (PDF's), taking the statistical distribution of parametric variations into consideration while analyzing the circuit. The simplest way to achieve this, in terms of the complexity of implementation, may be through Monte Carlo analysis. While such an analysis can handle arbitrarily complex variations, its major disadvantage is in its extremely large run-times. Therefore, more efficient methods are called for.

The task of static timing analysis can be distilled into two types of operations:

- A gate is being processed in STA when the arrival times of all inputs are known, at which time the candidate delay values at the output are computed using the "sum" operation that adds the delay at each input with the input-to-output pin delay.
- Once these candidate delays have been found, the "max" operation is applied to determine the maximum arrival time at the output.

In SSTA, the operations are identical to STA; the difference is that the pin-to-pin delays and the arrival times are PDFs instead of single numbers.

The first method for statistical static timing analysis to successfully process large benchmarks under probabilistic delay models was proposed by Berkelaar in [4]. In the spirit of static timing analysis, this approach was purely topological, and ignored the Boolean structure of the circuit. It assumed that each gate in the circuit has a delay distribution that is described by a Gaussian PDF, and assumed that all process variations were uncorrelated.

The approach maintains an invariant that expresses all arrival times as Gaussians. As a consequence, since the gate delays are Gaussian, the "sum" operation is merely an addition of Gaussians, which is well known to be a Gaussian. The computation of the max function, however, poses greater problems. The set of candidate delays are all Gaussian, so that this function must find the maximum of Gaussians. In general, the maximum of two Gaussians is *not* a Gaussian. However, given the intuition that if a and b are Gaussian random variables, if $a \gg b$, $\max(a, b) = a$, a Gaussian; if $a = b$, $\max(a, b) = a = b$, a Guassian, it may be reasonable to approximate this maximum using a Gaussian. In later work in [11], a precise closed-form approximation for the mean and variance was utilized.

Another class of methods includes the work in [3], which uses bounding techniques to arrive at the delay distribution of a circuit, but again, these ignore any spatial correlation effects, and it is important to take these into consideration.

Figure 3 shows a comparison of the PDF yielded by an SSTA technique that is unaware of spatial correlations, as compared with a Monte Carlo simulation that incorporates these spatial correlations. The clear difference between the curves demonstrates the need for developing methods that can handle these dependencies.

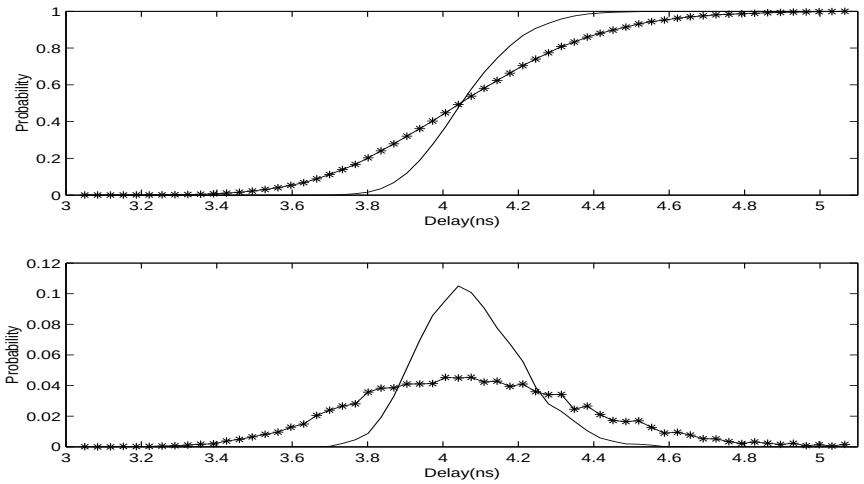

Fig. 3. A comparison of the results of SSTA when the random variables are spatially correlated. The line on which points are marked with stars represents the accurate results obtained by a lengthy Monte Carlo simulation, and the the solid curve shows the results when spatial correlations are entirely ignored. The upper plot shows the CDFs, and the lower plot, the PDFs [6].

The approach in [6] presents a novel and simple method based on the application of principal component analysis (PCA) techniques [13] to convert a set of correlated random variables into a set of uncorrelated variables in a transformed space; the PCA step can be performed as a preprocessing step for a design. The overall idea is similar to that of Berkelaar's, but the use of PCA permits rapid and efficient processing of spatial correlations. In reality, some parameters may be spatially correlated and others (such as T_{ox} and N_d) may be uncorrelated: this method is easily extended to handle these issues.

The overall flow of the algorithm is shown in Figure 4. The complexity of the method is $p \cdot n$ times the complexity of CPM, where n is the number of squares in the grid and p is the number of correlated parameters, plus the complexity of finding the principal components, which requires very low runtimes in practice.

Input: Process parameter variations
Output: Distribution of circuit delay

1. Partition the chip into $n = nrow \times ncol$ grids, each modeled by spatially correlated variables.
2. For each type of parameter, determine the n jointly normally distributed random variables and the corresponding covariance matrix.
3. Perform an orthogonal transformation to represent each random variable with a set of principal components.
4. For each gate and net connection, model their delays as linear combinations of the principal components generated in step 3.
5. Using "sum" and "max" functions on Gaussian random variables, perform a CPM-like traversal on the graph to find the distribution of the statistical longest path. This distribution achieved is the circuit delay distribution.

Fig. 4. Overall flow of the PCA-based statistical timing analysis method.

The overall CPU times for this method have been shown to be low, and the method yields high accuracy results.

4 Uncertainty as a virtue

4.1 Introduction

The concept of uncertainty can also be harnessed to advantage in providing efficient solutions to many difficult problems. Examples of such problems are as follows:

Randomized algorithms have been proposed in [14] for the solution of many combinatorial problems, including problems such as partitioning that arise in CAD. However, these have not been significantly developed in EDA.

Monte Carlo methods have been used very successfully to compute the average power dissipation of a circuit by applying a small fraction of the exponentially large space of possible input vectors to a circuit [5]. Such methods have also been employed for SSTA, as described earlier.

Random walk methods have been used to analyze large systems with localized behavior, such as in capacitance extraction [9], power grid analysis [16], and we are currently investigating their application to the analysis of electrostatic discharge (ESD) networks and to the problem of placement in physical design.

Other miscellaneous applications of random methods include techniques for crosstalk analysis [19] and in the probabilistic analysis of routing congestion [12, 21].

All of these point to the fact that the use of statistical methods in design is a vibrant and growing field with many upcoming challenges, particularly as, when used in the right contexts (e.g., when the computation is localized), these

methods can scale extremely well. We will illustrate one such method in the following section.

4.2 Case Study: Power grid analysis using random walks

On-chip power grids play an important role in determining circuit performance, and it is critical to analyze them accurately and efficiently to check for signal integrity, increasingly so in nanometer technologies.

A typical power grid consists of wire resistances, wire inductances, wire capacitances, decoupling capacitors, VDD pads, and current sources that correspond to the currents drawn by logic gates or functional blocks. There are two subproblems to power grid analysis: *DC analysis* to find steady-state node voltages, and *transient analysis* which is concerned with finding voltage waveforms considering the effects of capacitors, inductors and time-varying current waveform patterns.

The DC analysis of a power grid is formulated as a problem of solving a system of linear equations:

$$G\mathbf{X} = \mathbf{E} \tag{4}$$

where G is the conductance matrix for the interconnected resistors, $\mathbf{X}$ is the vector of node voltages, and $\mathbf{E}$ is a vector of independent sources. Traditional approaches exploit the sparse and positive definite nature of G to solve this system of linear equations for $\mathbf{X}$. However, the cost of doing so can become prohibitive for a modern-day power grid with hundreds of millions of nodes, and this will only become worse as the circuit size is ever growing from one technology generation to the next.

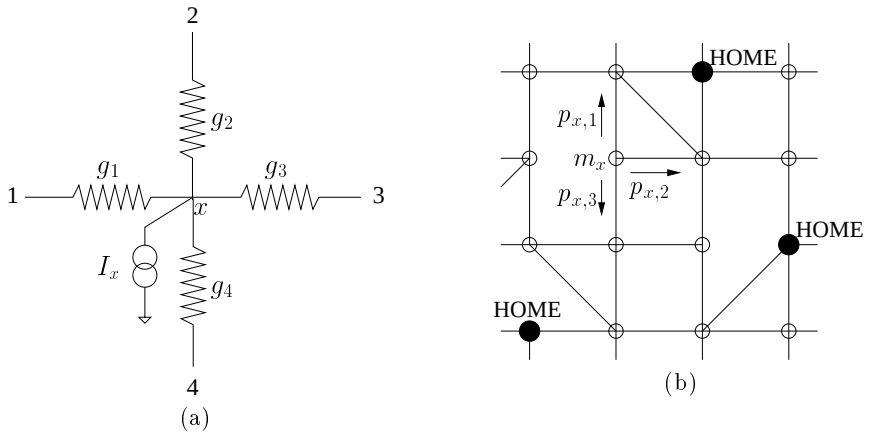

Fig. 5. (a)A representative power grid node. (b) An instance of a random walk "game."

For the DC analysis of a VDD grid, let us look at a single node x in the circuit, as illustrated in Figure 5(a). Applying Kirchoff's Current Law, Kirchoff's

Voltage Law and the device equations for the conductances, we can write down the following equation:

$$\sum_{i=1}^{\text{degree}(x)} g_i(V_i - V_x) = I_x \tag{5}$$

where the nodes adjacent to x are labeled $1, 2, \cdots$, degree(x), V_x is the voltage at node x, V_i is the voltage at node i, g_i is the conductance between node i and node x, and I_x is the current load connected to node x. Equation (5) can be reformulated as follows:

$$V_x = \sum_{i=1}^{\text{degree}(x)} \frac{g_i}{\sum_{j=1}^{\text{degree}(x)} g_j} V_i - \frac{I_x}{\sum_{j=1}^{\text{degree}(x)} g_j} \tag{6}$$

We can see that this implies that the voltage at any node is a linear function of the voltages at its neighbors. We also observe that the sum of the linear coefficients associated with the V_i's is 1. For a power grid problem with N non-VDD nodes, we have N linear equations similar to the one above, one for each node. Solving this set of equations will give us the exact solution.

We will equivalence this problem to a random walk "game," for a given finite undirected connected graph (for example, Figure 5(b)) representing a street map. A walker starts from one of the nodes, and goes to an adjacent node i every day with probability $p_{x,i}$ for $i = 1, 2, \cdots,$ degree(x), where x is the current node, and degree(x) is the number of edges connected to node x. These probabilities satisfy the following relationship:

$$\sum_{i=1}^{\text{degree}(x)} p_{x,i} = 1 \tag{7}$$

The walker pays an amount m_x to a motel for lodging everyday, until he/she reaches one of the homes, which are a subset of the nodes. If the walker reaches home, he/she will stay there and be awarded a certain amount of money, m_0. We will consider the problem of calculating the expected amount of money that the walker has accumulated at the end of the walk, as a function of the starting node, assuming he/she starts with nothing.

The gain function for the walk is therefore defined as

$$f(x) = E[\text{total money earned } | \text{walk starts at node } x] \tag{8}$$

It is obvious that

$$f(\text{one of the homes}) = m_0 \tag{9}$$

For a non-home node x, assuming that the nodes adjacent to x are labeled $1, 2, \cdots,$ degree(x), the f variables satisfy

$$f(x) = \sum_{i=1}^{\text{degree}(x)} p_{x,i} f(i) - m_x \tag{10}$$

For a random-walk problem with N non-home nodes, there are N linear equations similar to the one above, and the solution to this set of equations will give the exact values of f at all nodes.

It is easy to draw a parallel between this problem and that of power grid analysis. Equation (10) becomes identical to (6), and equation (9) reduces to the condition of perfect VDD nodes if

$$p_{x,i} = \frac{g_i}{\sum_{j=1}^{\text{degree}(x)} g_j} \qquad i = 1, 2, \cdots, \text{degree}(x)$$
$$m_x = \frac{I_x}{\sum_{j=1}^{\text{degree}(x)} g_j} \qquad m_0 = V_{DD} \qquad f(x) = V_x \tag{11}$$

A natural way to approach the random walk problem is to perform a certain number of experiments and use the average money left in those experiments as the approximated solution. If this amount is averaged over a sufficiently large number of walks by playing the "game" a sufficiently large number of times, then by the law of large numbers, an acceptably accurate solution can be obtained.

This is the idea behind the proposed *generic algorithm* that forms the most basic implementation. Numerous efficiency-enhancing techniques are employed in the implementation, and these have been described in [16, 17]. The DC solution above has also been extended to solve the transient analysis problem, which can be handled similarly, and with greater efficiency.

5 Conclusion

The effects of variability and uncertainty are here to stay in nanometer VLSI designs, and CAD techniques must be found to overcome them. This paper has outlined the basics of how a CAD engineer will have to deal with randomness in the future: not only in terms of dealing with it during design, but also in the sense of exploiting it by using algorithms that exploit randomness. This paper only skims the very surface of this issue, and there is likely to be considerable work in this field in the future.

References

1. A. Agarwal, D. Blaauw, V. Zolotov, and S. Vrudhula. Computation and refinement of statistical bounds on circuit delay. In *Proceedings of the ACM/IEEE Design Automation Conference*, pages 348–353, June 2003.
2. A. Agarwal, D. Blaauw, V. Zolotov, and S. Vrudhula. Statistical timing analysis using bounds. In *Proceedings of Design and Test in Europe*, pages 62–67, February 2003.
3. A. Agarwal, V. Zolotov, and D. T. Blaauw. Statistical timing analysis using bounds and selective enumeration. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 22(9):1243–1260, September 2003.

4. M. Berkelaar. Statistical delay calculation, a linear time method. In *Proceedings of ACM/IEEE International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems (TAU)*, pages 15–24, December 1997.
5. R. Burch, F. Najm, P. Yang, and T. N. Trick. A Monte Carlo approach for power estimation. *IEEE Transactions on VLSI Systems*, 1(1):63–71, March 1993.
6. H. Chang and S. S. Sapatnekar. Statistical timing analysis considering spatial correlations using a single PERT-like traversal. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pages 621–625, November 2003.
7. G. Chen and S. S. Sapatnekar. Partition-driven standard cell thermal placement. In *Proceedings of the ACM International Symposium on Physical Design*, pages 75–80, 2003.
8. Y. Cheng and S. M. Kang. A temperature-aware simulation environment for reliable ULSI chip design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(10):1211–1220, October 2000.
9. Y. L. Le Coz and R. B. Iverson. A stochastic algorithm for high-speed capacitance extraction in integrated circuits. *Solid-State Electronics*, 35:1005–1012, 1992.
10. B. Goplen and S. S. Sapatnekar. Efficient thermal placement of standard cells in 3D IC's using a force directed approach. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pages 86–89, 2003.
11. E. Jacobs and M. R. C. M. Berkelaar. Gate sizing using a statistical delay model. In *Proceedings of Design and Test in Europe*, pages 283–290, 2000.
12. J. Lou, S. Thakur, S. Krishnamoorthy, and H. S. Sheng. Estimating routing congestion using probabilistic analysis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 21(1):32–41, January 2002.
13. D. F. Morrison. *Multivariate Statistical Methods.* McGraw-Hill, New York, NY, 1976.
14. R. Motwani and P. Raghavan. *Randomized algorithms.* Cambridge University Press, Cambridge, UK, 1995.
15. M. Orshansky, L. Milor, P. Chen, K. Keutzer, and C. Hu. Impact of spatial intrachip gate length variability on the performance of high-speed digital circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 21(5):544–553, May 2002.
16. H. Qian, S. R. Nassif, and S. S. Sapatnekar. Random walks in a supply network. In *Proceedings of the ACM/IEEE Design Automation Conference*, pages 93–98, 2003.
17. H. Qian and S. S. Sapatnekar. Hierarchical random-walk algorithms for power grid analysis. In *Proceedings of the Asia/South Pacific Design Automation Conference*, pages 499–504, 2004.
18. S. S. Sapatnekar and H. Su. Analysis and optimization of power grids. *IEEE Design and Test*, 20(3):7–15, May-June 2002.
19. S. B. K. Vrudhula, D. Blaauw, and S. Sirichotiyakul. Estimation of the likelihood of capacitive coupling noise. In *Proceedings of the ACM/IEEE Design Automation Conference*, pages 653–658, 2002.
20. T.-Y. Wang and C. C.-P. Chen. 3-D thermal-ADI: A linear-time chip level transient thermal simulator. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 21(12):1434–1445, December 2002.
21. J. Westra, C. Bartels, and P. Groeneveld. Probabilistic congestion prediction. In *Proceedings of the ACM International Symposium on Physical Design*, pages 204–209, 2004.