Analysis and Verification of Power Grids Considering Process-Induced Leakage Current Variations

Imad A. Ferzli, *Student Member, IEEE,* Farid N. Najm, *Fellow, IEEE*

***Abstract*— The ongoing trends in technology scaling imply a reduction in the transistor threshold voltage (V_{th}). With smaller feature lengths and smaller parameters, variability becomes increasingly important, for ignoring it may lead to chip failure and assuming worst-case renders almost any design non-achievable. This work presents a methodology for the analysis and verification of the power grid of integrated circuits considering variations in leakage currents. These variations are large due to the exponential relation between leakage current and transistor threshold voltage and appear as random background noise on the nodes of the grid. We propose a lognormal distribution to model the grid voltage drops, derive bounds on the voltage drop variances, and develop a numerical Monte Carlo method to estimate the variance of each node voltage on the grid. This model is used toward the solution of a statistical formulation of the power grid verification problem.**

***Index Terms*— Power grid, leakage current, process variations, Monte Carlo sampling, statistical analysis.**

I. Introduction

Managing leakage is among the most eminent challenges to be hurdled by the integrated circuits industry, with the advent of deep-submicron technologies in the nanometer regime. Several effects contribute to chip leakage currents; weak channel inversion leading to subthreshold leakage, band-to-band tunneling, and gate oxide tunneling are among the many factors which induce leakage at various levels [1]. Being the dominant leakage component in modern, high-performance designs operating at high temperatures [2], [3], subthreshold leakage plays a particularly critical role, especially given its high sensitivity to threshold voltage $V_{\rm th}$ variations and the scaling trends of $V_{\rm th}$.

High subthreshold leakage comes as a price tag on better performance and reduced active power, as a result of the need to scale the MOSFET threshold voltage, (V_{th}), to accompany the reduction in supply voltages (V_{dd}) and oxide thickness. This is because the decrease of V_{th} translates as an exponential increase in subthreshold leakage current, $I_{\rm sub}$, as can be seen from the following expression of $I_{\rm sub}$, based on the BSIM3 device model [4]:

$$I_{\rm sub} = I_0 \exp\left(\frac{V_{gs} - V_{th}}{nV_T}\right)\left(1 - \exp\left(-\frac{V_{ds}}{V_T}\right)\right), \quad (1)$$

This research was supported in part by Micronet, with funding from ATI and from Altera, and by the SRC under contract 2003-TJ-1070.

The authors are with the Department of Electrical and Computer Engineering at the University of Toronto, Ontario, Canada.

where V_T is the thermal voltage, and n and I_0 is are constant factors related to device characteristics. $I_{\rm sub}$ reportedly increases at a rate as high as about 5X per generation [5], or ten-fold for a 0.1V decrease of threshold voltage [6], with total subthreshold leakage current in dense high-performance chips forecast to be about half the total chip current [7].

In today's 1.2 V, 0.13 μm technologies, V_{th} is about 0.3 V (25% of V_{dd}). Compare these figures with older 1μm technologies, using a supply voltage of 5V and having a threshold voltage of about 0.8 V (16% of V_{dd}) [8]. Clearly, the scaling trends of threshold voltage do not keep up with those of supply voltage.

The disparity in scaling the supply and threshold voltages is specially consequential to pursuing aggressive designs, inducing tighter design margins, thus placing process variations under great scrutiny. In particular, since the gap between V_{dd} and V_{th} has narrowed, variations in the level of supply voltage become very significant when it comes to meeting timing closure. The network that delivers supply voltage to circuit devices is referred to as the *power grid*. The consequences of power grid voltage level variations in the estimation of a circuit timing vulnerability are being currently researched [9], [10], and it has been reported in an industrial survey that 50% of chips would fail to meet their timing requirements if no power grid verification were done prior to tapeout [11]. However, modern power grid verification tools fail to account for the impact of the statistical variations of leakage on the power grid voltage levels. Our present work addresses this issue, and we focus on subthreshold leakage, referred to thereafter simply as "leakage". Preliminary versions of this work appeared in [12], [13].

If leakage were uniform or predictable across the chip, then the contribution of leakage to the supply voltage drop can be simply accounted for by adding a deterministic leakage-induced component to the total voltage drop on the power grid. However, this is not the case and process variations in V_{th} and the transistor channel length cause exponentially larger variations in leakage currents, as a direct consequence of the exponential dependence of leakage currents on V_{th} or channel length [4], [14]. In 0.1 μm technology, it is possible to get 30 mV standard deviation in V_{th} [15]. Considering a supply voltage of, say, around 1 V, this means that a $\pm 3\sigma$ interval for V_{th} would span 18% of the supply! For whole chips, leakage variations have been *measured* at almost 20X, leading to a slowdown of the chip operating frequency by a factor of 1.4 [16], [17]. These variations are also known to have

a significant within-die, local component [18], [19], so that transistors in close proximity on the layout can have significant variations in their leakage currents. This is also referred to as mismatch and its effect on delay has been studied [20].

Recent publications have focused on the characterization of leakage currents subject to process variations. In [14], an approach was proposed to estimate the leakage current considering within-die variations. In [21], the sensitivity of leakage currents was studied with respect to physical parameters, including channel length, oxide thickness, and channel doping. This last approach was extended in [4] resulting in a technique for computing the probability distribution function of leakage currents, including die-to-die and intra-die variations. Iterative methods were suggested in [22] to estimate chip leakage considering variations in the supply voltage and temperature distribution across the chip, and a probabilistic approach was introduced in [23] to estimate statistical parameters of the chip leakage current.

We will assume statistical variability of leakage currents, *i.e.*, that the statistics of leakage currents are available, with the knowledge that both Monte Carlo and analytical techniques can be used to obtain this information [4]. While worst-case analysis is necessary when considering power grid voltage drop due to global, correlated die-to-die variations [18] in leakage currents, intra-die (or within-die) variations require statistical analysis in order to avoid overly pessimistic conclusions. There are no good tools today to estimate this statistical voltage drop on the power grid due to leakage current variations.

The paper is organized as follows: section II presents the equations that provide a model for the power grid and presents a breakdown of currents loading the grid in order, isolating the components that represent leakage current variability. Section III further discusses leakage current variations, where it is argued that corner-case analysis should be used to account for die-to-die variations, but that statistical analysis is necessary when dealing with within-die variations. Focusing on within-die variations, we first propose an analytical probability distribution model for the voltage drop at every node in section IV, and we derive a relation between the second-order statistics of leakage currents and grid voltage drops. Observing that the main computational difficulty lies in variance calculations, we propose a numerical Monte Carlo technique to estimate the variances and covariances in section V. In section VI, we suggest a statistical verification for power grids that is aware of leakage variability, and the lognormal voltage drop model is used to derive verification conditions. Section VII derives direct and iterative criteria to check these conditions, where the Monte Carlo technique proves very useful again. Results are shown in section VIII and we conclude in section IX.

II. System Models

A. Power Grid Equations

We consider an RC model of the power grid [24]. Let t be the total number of nodes on the grid. C4 sites make up for nodes with a voltage source of value V_{dd} to ground attached to them. Let p be the number of such sites, then $t = N + p$, where N is the number of nodes with no voltage source attached. These nodes are of interest since the voltage level of the p C4 sites is known. Let c_k ($k = 1, \ldots, N$) be the capacitance from every node to ground and $\mathbf{C}$ the diagonal matrix of all such capacitances. Let $i_k(t)$ be the value of the current source to ground attached to node k (note that if no current source is attached to some node, then $i_k(t) = 0$, $\forall t$). Let $\mathbf{i}(t)$ be the vector whose k^{th} component is $i_k(t)$. Let $\mathbf{G}$ be the conductance matrix [25] of the power grid, obtained from the resistive branches. Let $u_k(t)$ be the voltage at node k and $\mathbf{u}(t)$ be the vector of all such node voltages. Modified nodal analysis [25] applied to the power grid leads to the following system equation:

$$\mathbf{G}\mathbf{u}(t) + \mathbf{C}\dot{\mathbf{u}}(t) = -\mathbf{i}(t) + \mathbf{G}\mathbf{V}_{\mathbf{dd}}, \tag{2}$$

where every component of the vector $\mathbf{V}_{\mathbf{dd}}$ equals V_{dd}. Let $v_k(t) = V_{dd} - u_k(t)$ be the *voltage drop* at node k, and $V(t)$ the vector of voltage drops, then 2 can be written as:

$$\mathbf{G}\mathbf{v}(t) + \mathbf{C}\dot{\mathbf{v}}(t) = \mathbf{i}(t). \tag{3}$$

This is a *revised* equation for the node voltage drops, where the voltage sources have been set to zero and the direction of the currents reversed. In the following, we will be mainly concerned with the above revised system. When the circuit is in standby state, we solve (3) in DC state, discarding the capacitance matrix $\mathbf{C}$ and eliminating the time dependence from the voltage drops and currents, to yield:

$$\mathbf{G}\mathbf{V} = \mathbf{I}. \tag{4}$$

In the sequel, we use the static power grid model to derive the distribution of the node voltage drops on the power grid and their variances/covariances, due to within-die leakage current variations. The transient model is used to calculate the vector of voltage drop means (see section IV).

B. Breakdown of Current Components

Since the power grid has the property of being a linear system as shown in (2), the grid response, *i.e.*, its voltage drops due to circuit currents, can be obtained simply by adding up the grid response to individual current components. In particular, both active and leakage current effects ought to be considered when determining the grid response, which can be done by superposition of individual responses to both switching and leakage currents (static and dynamic).

Process-induced variations on the leakage component of the total current are relevant in the context of this work. It is standard practice to break down these process variations into *die-to-die* and *within-die* components [26], [27]. For the variations on a given parameter, the die-to-die component takes the same value for all instances of that parameter within a single chip, but differentiates among distinct chips. The within-die component causes variations within a given chip, depending on location.

We therefore model the total current as follows:

$$I = I_{\text{active}} + I_{\text{sub,nom}} + I_{\text{sub,dd}} + I_{\text{sub,wd}}, \tag{5}$$

where $I_{\rm active}$ refers to the active current, $I_{\rm sub,nom}$ is the *nominal* value of leakage current at the design point, $I_{\rm sub,dd}$ is the die-to-die component of leakage variations, and $I_{\rm sub,wd}$ is the within-die component of these variations. By linearity of the power grid, we can sum up the contributions of the four current components in (5) to get the total voltage drop on the power grid, as follows:

$$V = V_{I_{\rm active}} + V_{I_{\rm sub},\rm nom} + V_{I_{\rm sub},\rm dd} + V_{I_{\rm sub},\rm wd}. \quad (6)$$

Voltage drops on the power grid induced by active currents are calculated on either a transient basis or on a DC basis, using average or peaks of active currents, or combinations thereof. Typically, this is what power grid verification tools perform today. Clearly, active currents and nominal leakage constitute *deterministic* cases. This leaves inter- and intra-die variations to be incorporated in the analysis and verification of the power grid.

III. Dealing with Leakage Variability

A. Die-to-Die Variations

We recall the following *monotonicity* property of the power grid [28]:

Proposition. (monotonicity) *If $\mathbf{v}(t)$ is the voltage drop due to $\mathbf{i}(t)$ and $\mathbf{v}^*(t)$ is the voltage drop due to $\mathbf{i}^*(t)$, then the power grid has the following property:*

$$\text{if} \quad \mathbf{i}^*(t) \geq \mathbf{i}(t), \forall t \geq 0, \quad \text{then} \quad \mathbf{v}^*(t) \geq \mathbf{v}(t), \forall t \geq 0 \quad (7)$$

which we will express by saying that the grid is monotone.

Note that the monotonicity property applies to both DC and transient power grid analyses. A similar result was earlier proven [29] for the special case of an RC tree driven by a single voltage source. The monotonicity property allows one to handle die-to-die variations easily.

Since die-to-die variations are correlated across the chip, corner-case leakage currents on a given chip (due for example to largest variations in transistor threshold voltage) are not unlikely and are accounted for through case file analysis, since, by virtue of the monotonicity property of the grid, these corner leakage currents correspond to corner voltage drops on the grid. Therefore, it becomes very easy to calculate $V_{I_{\rm sub},\rm dd}$ for all power grid nodes, by setting the components of the die-to-die leakage variations at the top of their range (*e.g.*, at the $\rm mean + 3\sigma$ point) and solving the grid once (using (3) or (4)), then setting them at the bottom of their range (*e.g.*, at the $\rm mean - 3\sigma$ point) and solving the grid again. Cases other than the $\pm 3\sigma$ corners (*e.g.*, the median or $50^{\rm th}$ percentile of leakage) can be used to reflect more aggressive verifications.

B. Intra-Die Variations

Handling within-die variations is inherently more difficult than die-to-die variations, primarily because the currents are not all correlated, so that occurrence of a corner-case is highly unlikely.

Within-die variations have both *systematic* and *random* components. Systematic variations arise from the observed wafer-level variation trends reflected on a given die [18], [30] or the spatial locations of some features on the die and their context in terms of the neighboring layout patterns. They account for the local, layout-dependent *correlations* due to physical parameter variations across a chip, and the correlations that arise from "environmental conditions" [23], such as supply voltage and temperature. Random variations constitute the residual component of the total variations which, in essence, cannot be explained systematically [30], and are modeled as *statistically independent* random variables. Ideally, one wants to extract systematic variations and treat them as "deterministic" [31]. This approach, however, necessitates detailed variation models. and is hard to put to use pre-layout. It is standard practice today to ignore the systematic within-die component, at least in the early stages of the design, for several reasons: 1) systematic trends are layout-dependent, so that incorporating systematic variation models can occur only post-layout, 2) within-die correlation generally dies down quickly across short distances, and 3) detailed variation models are often not easily available or readily usable by the designers.

These considerations extend to the special case of within-die leakage current variability. Specifically, within-die leakage variations are induced by within-die variations in physical device parameters such as threshold voltage and channel length [14], and are also affected by supply voltage and temperature [22]. This work focuses on within-die leakage variations induced by threshold voltage tolerances within-chip. We will proceed under the assumption that leakage currents in different areas of a chip are statistically independent, which effectively amounts to considering all within-die $V_{\rm th}$ variations on a random basis. To emphasize this, we make explicit the following:

Assumption 1. *At any time point, we model all within-die leakage current variations as statistically independent.*

Since corner-case analysis would yield overly pessimistic predictions of the effect of intra-die variations because of their locality, *statistical analysis* needs to be effectively put to use. Analytical methods can be found in the literature to derive expressions for the statistics of witin-die leakage current variations (see [4], [14]), namely, their means and variances. Given such distributional information, our problem is to estimate the means and variances (and possibly covariances) of the corresponding grid voltages drops, $V_{I_{\rm sub},\rm wd}$. The rest of this paper develops the statistical analysis and verification of the grid in response to intra-die leakage variations.

IV. Statistics of Node Voltage Drops

A. Distributional Model for Node Voltage Drops

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.

In order to characterize the distribution of voltages at the nodes of the power grid, we have found it necessary to introduce the following *"pseudo-static" assumption*:

Assumption 2. *In order to obtain the distribution of the power grid voltage drops in response to within-die variations in leakage currents, 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 a precise characterization for the voltage drop distributions. Notice that this assumption is automatically true for the case of standby leakage. Thus, our analysis is exact for 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 we will include in the analysis *some* dynamics of the system through the computation of the mean response (see section IV-B below).

With this assumption and using (4), we have:

$$\mathbf{V}_{\mathbf{I}_{\mathrm{sub,wd}}} = \mathbf{G}^{-1}\mathbf{I}_{\mathrm{sub,wd}}, \tag{8}$$

where we consider $\mathbf{I}_{\mathrm{sub,wd}}$ to be the *random vector* representing within-die leakage variations, and $\mathbf{V}_{\mathbf{I}_{\mathrm{sub,wd}}}$ the corresponding grid response vector, with reference to (5) and (6). In the following, and for simplicity, we drop the subscripts indicating within-die leakage variations.

From (8), it is clear that the voltage drop at an arbitrary node i, V_i, can be expressed as:

$$V_i = q_{i1}I_1 + \cdots + q_{iN}I_N, \tag{9}$$

where q_{ij} is the (i,j)th entry of $\mathbf{G}^{-1}$. As (9) shows, the voltage drop at any node is a linear combination of the leakage currents loading the grid.

Since intra-die leakage variations arise from intra-die variations in transistor threshold voltage, modeled as Gaussian [14], then intra-die leakage variations are *lognormal* [4], by virtue of the exponential relation between transistor leakage and its threshold voltage shown in (1).

If the I_j are lognormal, then the $q_{ij}I_j$ are also lognormal. Hence, under *Assumption 1*, grid node voltages, given in (9), are each a summation of independent lognormal RVs. Sums of independent lognormal variables have been extensively studied and characterized in the literature pertaining to communications, and it was found that such sums can be accurately captured by another lognormal RV [32]. Therefore, we model the distribution of a node voltage drop as lognormal. We provide empirical data to corroborate this argument in section VIII-B.

Since V_i is lognormal, then $W_i = \ln(V_i)$ is Gaussian. The probability distribution function (pdf) of V_i, $f_{V_i}(.)$, at $v > 0$, is given by [33]:

$$f_{V_i}(v) = \exp\left(-\frac{[\ln(v) - \mu_{W_i}]^2}{2\sigma_{W_i}^2}\right) \Big/ \left(\sqrt{2\pi}\sigma_{W_i} v\right), \tag{10}$$

where μ_{W_i} and σ_{W_i} are respectively the mean and standard deviation of W_i. The cumulative distribution function (cdf) of V_i, $F_{V_i}(.)$, at $v > 0$, is given by [33]:

$$F_{V_i}(v) = \Phi\left(\frac{\ln(v) - \mu_{W_i}}{\sigma_{W_i}}\right), \tag{11}$$

where $\Phi(.)$ is the cdf of the standard normal distribution (Gaussian with 0 mean and unit variance). From (10) and (11), it is clear that the distribution of V_i has two parameters: the mean and standard deviation of the associated Gaussian RV W_i.

μ_{W_i} and σ_{W_i} are related to the mean μ_{V_i} and standard deviation σ_{V_i} of V_i as follows [33]:

$$\mu_{W_i} = \ln(\mu_{Vi}) - \frac{1}{2}\ln\left(1 + \frac{\sigma_{V_i}^2}{\mu_{V_i}^2}\right), \tag{12}$$

$$\sigma_{W_i} = \left(\ln\left[1 + \frac{\sigma_{V_i}^2}{\mu_{V_i}^2}\right]\right)^{1/2}. \tag{13}$$

From (12) and (13), we clearly see that computing the mean and variance of the voltage drop V_i completely specifies the distribution of V_i.

B. Distribution Parameters

Mean calculation is actually trivial. Since the system (3) is linear, then due to linearity of the mean ($E[\cdot]$) operator [33], one can write:

$$\mathbf{G}E[\mathbf{v}(t)] + \mathbf{C}\frac{d}{dt}E[\mathbf{v}(t)] = E[\mathbf{i}(t)]. \tag{14}$$

Thus, if we solve the system (3) once, using simply the means of leakage current variations as inputs, the solution gives the voltage means at all the nodes, obtained from the dynamic model of the grid. Note that mean calculation may be performed on a transient basis, and introduces dynamic properties of the power grid into our formulation.

We proceed to derive the covariance matrix of the node voltage drops induced by the random vector of within-die leakage current variations. Recall that the off-diagonal entries of covariance matrix of a random vector represent pairwise covariances, while the diagonal entries are the variances of the corresponding random variables [34]. Under static conditions (*Assumption 2*), we can write from (14):

$$\mathbf{G}E[\mathbf{V}] = E[\mathbf{I}]. \tag{15}$$

We now combine (4) and (15) to yield:

$$\mathbf{G}(\mathbf{V} - E[\mathbf{V}]) = \mathbf{I} - E[\mathbf{I}] \tag{16}$$

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

$$\mathbf{G}E\left[(\mathbf{V} - E[\mathbf{V}])(\mathbf{V} - E[\mathbf{V}])^T\right]\mathbf{G}^T = E\left[(\mathbf{I} - E[\mathbf{I}])(\mathbf{I} - E[\mathbf{I}])^T\right]. \tag{17}$$

We recognize the expectations on the left– and right–hand side of (17) as being simply covariance matrices of node

voltage drops $\Sigma_{\mathbf{V}}$ and within-die leakage current variations, $\Sigma_{\mathbf{I}}$, respectively. Thus, the above result can be rewritten as:

$$\mathbf{G}\Sigma_{\mathbf{V}}\mathbf{G}^T = \Sigma_{\mathbf{I}}. \tag{18}$$

Since $\mathbf{G}$ is symmetric, $\mathbf{G} = \mathbf{G}^T$. Therefore, (18) becomes:

$$\Sigma_{\mathbf{V}} = \mathbf{G}^{-1}\Sigma_{\mathbf{I}}\mathbf{G}^{-1}. \tag{19}$$

Under the assumption of statistical independence of within-die leakage current variations (*Assumption 1*), implying that $\Sigma_{\mathbf{I}}$ is diagonal, it can be seen from (19) that:

$$\begin{aligned}[\Sigma_{\mathbf{V}}]_{ij} &= \left[\mathbf{G}^{-1}\right]_{i1}\left[\mathbf{G}^{-1}\right]_{j1}[\Sigma_{\mathbf{I}}]_{11} + \cdots + \\ &\quad \left[\mathbf{G}^{-1}\right]_{iN}\left[\mathbf{G}^{-1}\right]_{jN}[\Sigma_{\mathbf{I}}]_{NN} \\ &= q_{i1}q_{j1}\sigma^2_{I_1} + \cdots + q_{iN}q_{jN}\sigma^2_{I_N}, \end{aligned} \tag{20}$$

where $[.]_{ij}$ is the (i,j)th entry of a matrix, q_{ij} is as defined in (9), and $\sigma^2_{I_i}$ is the variance of the i^{th} leakage current, that is, the (i,i)th entry of $\Sigma_{\mathbf{I}}$. Specifically, the variance of V_i, $\sigma^2_{V_i}$, is the (i,i)th entry of $\Sigma_{\mathbf{V}}$:

$$\sigma^2_{V_i} = [\Sigma_{\mathbf{V}}]_{ii} = q^2_{i1}\sigma^2_{I_1} + \cdots + q^2_{iN}\sigma^2_{I_N}. \tag{21}$$

Let $\mathbf{Var}(.)$ be the operator which takes a random vector and returns the *vector* of variances of each component of that random vector, respectively. Define a matrix $\mathbf{G}^{-1(2)}$ such that:

$$[\mathbf{G}^{-1(2)}]_{ij} = q^2_{ij}. \tag{22}$$

Now (21) can be written in matrix form as:

$$\mathbf{Var}(\mathbf{V}) = \mathbf{G}^{-1(2)}\mathbf{Var}(\mathbf{I}). \tag{23}$$

From the above discussion, we have established that calculating the voltage drop means can be easily done. However, we observe from (19) through (23) that computing the covariance matrix of the voltage drops or even only their variances requires full knowledge of the cumbersome $\mathbf{G}^{-1}$, the inverse of the grid conductance matrix. The next section discusses estimating these second-order statistics efficiently.

V. Variance/Covariance Estimation

A. A Current-Sampling Approach

One way of estimating the (output) variance/covariance of voltage given the (input) variance of current is to proceed by brute-force random sampling on the currents. Generate a randomly-chosen vector of current values, according to the lognormal distributions of the currents, and solve the system (4) for a corresponding sample of the voltages. Do this repeatedly and collect statistics on the voltages. Stop the sampling when the desired statistics have been estimated with sufficient accuracy. We have implemented this technique and found that it is viable in some cases, but that it suffers serious disadvantages regarding the accuracy of the variance/covariance estimates: as compared with the column-sampling approach, to be presented below, up to 10 times as many variance/covariance estimates may not reach convergence within the same execution time. Further, even in the suggested simple brute-force scheme, and since the sample variance can be related to a χ^2 distribution, one will be required to evaluate certain percentiles of the χ^2 distribution for every sample, in order to check whether an estimator of the variance/covariance has converged, This can get very expensive because it requires numerical integration of the χ^2 pdf - using a table is not a practical option because the number of samples can be huge. As such, this current sampling approach will not be discussed further; in contrast, we present a numerical Monte Carlo technique, based on a column-sampling method, which effectively transforms the variance/covariance estimation problem to a mean estimation problem, so that simple Monte Carlo methods, based on the central limit theorem, can be applied.

B. A Numerical Monte Carlo Approach

We propose a numerical Monte Carlo technique to estimate variances/covariances of node voltage drops [12], [35]. Observe that (20) is a weighted summation. Let $S = \sum_{i=1}^{N}\sigma^2_{I_i}$, $p_i = \sigma^2_{I_i}/S$, and $\sigma_{ij} = [\Sigma_{\mathbf{V}}]_{ij}$. Then (20) can be rewritten as:

$$\sigma_{ij} = S\sum_{l=1}^{N} p_l(q_{il}q_{jl}). \tag{24}$$

Since $\sum_{k=l}^{N} p_l = 1$ and $p_l \geq 0$, $l = 1 \cdots N$, then we can view the p_l weights as being *probability values* associated with the $q_{il}q_{jl}$ values, so that the summation in (24) becomes the mean (weighted average) of all the $q_{il}q_{jl}$ values in the i^{th} and j^{th} rows. If we define an RV $\mathbf{r_{ij}}$ as being a discrete RV that takes the values $r_{ijl} = q_{il}q_{jl}$ with probabilities p_l, $l = 1, 2, \ldots, N$, then we can write (24) as:

$$\sigma_{ij} = SE[\mathbf{r_{ij}}]. \tag{25}$$

Let the mean of $\mathbf{r_{ij}}$ be $\mu_{ij} = E[\mathbf{r_{ij}}]$ and its variance be v_{ij}. We can now use methods of *mean estimation* from statistics, basically Monte Carlo random sampling [36], in order to estimate the *population mean* μ_{ij} 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_l a sequence of indices of columns of $\mathbf{G}^{-1}$ to be included in the sample. From these, we form the following sample mean for any pair of rows i and j:

$$\bar{r}_{ij} = \frac{1}{n}\sum_{l\in\mathcal{L}} r_{ijl}, \tag{26}$$

where $\mathcal{L}$ is the set of indices included in the random sample.

We also compute the *sample standard deviation* of $\mathbf{r_{ij}}$, $s_{ij} \geq 0$ given by:

$$s^2_{ij} = \frac{1}{n-1}\sum_{l\in\mathcal{L}}(r_{ijl} - \bar{r}_{ij})^2 = \frac{n\left(\sum_{l\in\mathcal{L}} r^2_{ijl}\right) - \left(\sum_{l\in\mathcal{L}} r_{ijl}\right)^2}{n(n-1)}. \tag{27}$$

Note that s^2_{ij} is an estimator of v_{ij}. Now $\bar{r}_{ij}$ itself, being a sample mean, can be considered as an RV, with mean μ_{ij} (since the sample mean is an unbiased estimator of the mean

of a random variable [33]) and variance s_{ij}^2/n (for large n). Then, we have that

$$\hat{\sigma}_{ij} = S\bar{r}_{ij} \tag{28}$$

is an unbiased estimator of σ_{ij} (*i.e.*, $E[\hat{\sigma}_{ij}] = SE[\bar{r}_{ij}] = \sigma_{ij}$), with variance $S^2 s_{ij}^2/n$. Furthermore, by the *central limit theorem* [33], $\bar{r}_{ij}$ will be approximately normally distributed, so that the RV $(\sigma_{ij} - \hat{\sigma}_{ij})\sqrt{n}/(Ss_{ij})$ is normal with 0 mean and unit variance [36], for large n.

Fig. 1 illustrates the process when estimating the covariance between nodes 1 and 3. In the figure, columns 9, 5, 7, and 2 of $\mathbf{G}^{-1}$ are sampled. Notice that sampling $\mathbf{G}_i^{-1}$, the i^{th} column of $\mathbf{G}^{-1}$, involves solving $\mathbf{G}\mathbf{G}_i^{-1} = \mathbf{I_i}$, where $\mathbf{I_i}$ is the i^{th} column of the identity matrix, and that in fact, full knowledge of $\mathbf{G}^{-1}$ is circumvented. With reference to Fig. 1, the initial three column samples lead to values for r_{139}, r_{135}, and r_{137}, with corresponding sample mean $\bar{r}_{13}$ and sample standard deviation s_{13}. The sample mean and standard deviation are updated when column 2 is sampled. This process continues until convergence of $\bar{r}_{13}$, at which point $S\bar{r}_{13}$ is taken as an estimator of the covariance between the voltage at nodes 1 and 3. As will be discussed below, the convergence of the sample mean is directly related to how small the sample standard deviation becomes, and to the number n of samples taken. Observe that $\mathbf{G}^{-1}$ is a symmetric matrix, so that when columns i and j of $\mathbf{G}^{-1}$ are sampled, the covariance between V_i and V_j and the variances of V_i and V_j can be computed exactly, as per (20). With reference to Fig. 1, since columns 9, 5, 7, and 2 have been sampled, the variances of, and pairwise covariances between V_9, V_5, V_7, and V_2 can be computed, thus eliminating the need to bring estimates for these quantities to converge via the sampling process.

We can use tables of the standard normal to establish how large n should be in order for the sample standard deviation of $\bar{r}_{ij}$ to be small enough for it to be a viable estimator of μ_{ij}, with a certain confidence, and up to a predefined tolerance [36]. 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:

$$|\bar{r}_{ij} - \mu_{ij}| < \epsilon, \tag{29}$$

then it is known from sampling theory [36] that n should be larger than n_0 where:

$$n_0 = \left(\frac{z_{\alpha/2} s_{ij}}{\epsilon}\right)^2, \tag{30}$$

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 $\alpha/2$, *i.e.*, $\Phi(z_{\alpha/2}) = 1 - \alpha/2$. Thus, for instance, for 95% confidence, $\alpha = 0.05$ and $z_{\alpha/2} = 1.960$; for 99% confidence, $\alpha = 0.01$ and $z_{\alpha/2} = 2.576$. In practice, one samples until n is larger than 30 or so, then starts to use (30) to monitor convergence. Observe that tables are not needed to compute $z_{\alpha/2}$, a task that can be easily done by software (*e.g.*, using the erf() function) and that $z_{\alpha/2}$ needs to be evaluated only once.

The process described above can be used to estimate any entry in the covariance matrix $\Sigma_{\mathbf{V}}$ of the node voltage drops. In particular, the diagonal entries of $\Sigma_{\mathbf{V}}$ represent the variances of these voltage drops and can be estimated by setting $i = j$, for $i = 1, \ldots, N$ in the above calculations. In the following, we are mainly concerned with the described Monte Carlo method applied for variance calculations.

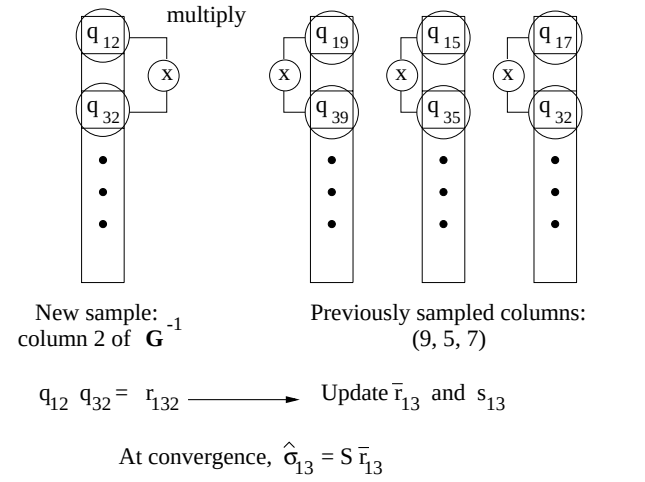

Fig. 1. Estimating the covariance between nodes 1 and 3 on the power grid.

C. Error Bound for Variance Estimation

Our intent here is to approximate the variance of every node voltage by choosing a meaningful error bound ϵ for use in (30), with $i = j$, to check for the convergence of the estimated variance value. Using the same notation as in the previous section, we have from (28) that:

$$\sigma_{V_i}^2 = \sigma_{ii} \approx S\bar{r}_{ii}. \tag{31}$$

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

$$\left|\sigma_{V_i} - \sqrt{S\bar{r}_{ii}}\right| \leq \delta V_{dd}. \tag{32}$$

In other words, we want to find ϵ (to be used in (30)) as a function of δ in order for the following to be true:

$$|\bar{r}_{ii} - \mu_{ii}| < \epsilon \quad \Longrightarrow \quad \left|\sqrt{S\bar{r}_{ii}} - \sqrt{S\mu_{ii}}\right| < \delta V_{dd}. \tag{33}$$

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

$$|x - x_0| < \epsilon \quad \Longrightarrow \quad |y - y_0| < \gamma. \tag{34}$$

There are two cases to consider, according to whether y_0 is small or not, as shown in Fig. 2. When y_0 is small, small enough so that $y_0 - \gamma < 0$, then (since $y > 0$ in all cases) in order to guarantee that $|y - y_0| < \gamma$, it is sufficient to impose an upper bound on y in the simple form $y - y_0 < \gamma$. This is achieved by imposing an upper bound on x in the simple form $x - x_0 < \epsilon_1$, where ϵ_1 is the corresponding upper bound on $(x - x_0)$, as shown in Fig. 2(a).

Let $\Delta x = x - x_0$ and $\Delta y = y - y_0$. Since $y^2 = x$, then $\Delta(y^2) = \Delta x$, and since $\Delta(y^2) = (y_0 + \Delta y)^2 - y_0^2 = (\Delta y)^2 + 2y_0\Delta y$, then $\Delta x = ((\Delta y)^2 + 2y_0\Delta y)$. For $\Delta y = \gamma$, $\epsilon_1 = \Delta x = (\gamma^2 + 2y_0\gamma) > 0$. In order to guarantee (34), in this case, we need to set $\epsilon = \epsilon_1 = (2y_0\gamma + \gamma^2)$.

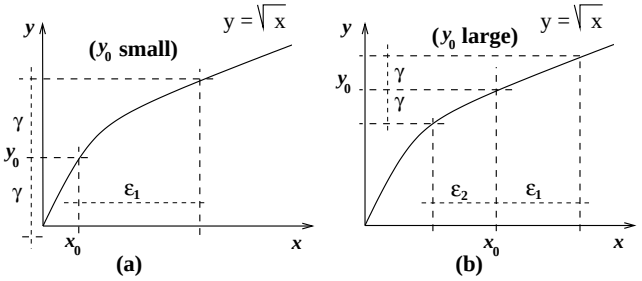

Fig. 2. Choice of error bound ϵ in the two cases when (a) y_0 is small and (b) y_0 is not so small.

When y_0 is not so small, i.e., when $y_0 - \gamma > 0$, then we need to consider both upper bounds and lower bounds, as shown in Fig. 2(b). If ϵ_2 is the lower bound on $(x - x_0)$, as shown in the figure, then we may compute ϵ_2 by setting $\Delta y = -\gamma$, which leads to $\Delta x = (\gamma^2 - 2y_0\gamma) < 0$, and we set $\epsilon_2 = -\Delta x = (2y_0\gamma - \gamma^2) > 0$. Since $\epsilon_2 < \epsilon_1$, as is obvious from the expressions found for each, then, in order to guarantee (34) in this case, we need to set $\epsilon = \epsilon_2 = (2y_0\gamma - \gamma^2)$. Notice that the condition $y_0 > \gamma$ translates to $\bar{r}_{ii} > \delta^2 V_{dd}^2/S$.

In summary, then:

$$\epsilon = \begin{cases} \frac{\delta V_{dd}}{S}\left(2\sqrt{S\bar{r}_{ii}} + \delta V_{dd}\right) & \text{, if } \bar{r}_{ii} < \delta^2 V_{dd}^2/S, \\ \frac{\delta V_{dd}}{S}\left(2\sqrt{S\bar{r}_{ii}} - \delta V_{dd}\right) & \text{, if } \bar{r}_{ii} > \delta^2 V_{dd}^2/S. \end{cases} \quad (35)$$

Notice also that, in either case, $\epsilon > \delta^2 V_{dd}^2/S$. Plugging (35) into (30) leads to:

$$n_0 = \left(\frac{z_{\alpha/2} s_{ii} S/\delta V_{dd}}{\delta V_{dd} \pm 2\sqrt{S\bar{r}_{ii}}}\right)^2, \quad (36)$$

where the $+$ or $-$ sign depends on whether $\bar{r}_{ii}$ is smaller or larger than $\delta^2 V_{dd}^2/S$, respectively. To summarize, if it is desired to estimate σ_{V_i} to within $\pm\delta V_{dd}$, with $(1-\alpha)\times 100\%$ confidence, then n must be larger than the threshold given by (36). This provides a useful trade-off between accuracy and speed, as more samples would be required for smaller δ.

By transforming the variance estimation problem into a mean estimation problem, we were effectively able to apply Monte Carlo methods to estimate the variance of each node voltage drop. Our sampling procedure, where the likelihood of a column to be sampled is driven by the current load at the corresponding node, intuitively amounts to giving a greater chance for a node with high-current variance and the corresponding column of $\mathbf{G}^{-1}$ to be sampled, which amounts to treating at as likely to have a greater impact on voltage drop variances on the grid than nodes with lower current variances. The error bound (δV_{dd}) and confidence level ($1-\alpha$) impose a statistical bound on the error. The Monte Carlo technique along with the stopping criteria derived above thus enable the estimation of the variances of grid voltage drops, while circumventing the problem of knowing the full inverse of the grid conductance matrix. Since information on the mean voltage drop is readily available (section IV-B), we have all parameters necessary to complete the description of the probability distribution of the node voltage drops in response to within-die leakage variations, given in (10) through (13).

VI. Power Grid Verification

A. Verification Methodology

Within-die variability of power grid node voltage drops appears as a process-induced background level of noise on the grid. In this section, we propose to verify the grid in presence of this noise [13]. The idea is to verify whether the "bulk of the distribution" at any/every every node falls below a user-specified "safety threshold voltage level". The user specifies what exactly is meant by the bulk of the distribution, by specifying a percentage of the probability distribution function to be below the threshold so that a node may be deemed "safe". The safety threshold voltage and the confidence level can be chosen to reflect aggressive designs or more conservative choices. For instance, when designers are able to make a statement such as: "Node i is safe if, with confidence $1-\beta_i$, the voltage drop at i is less than V_{Ti} volts", the approach described in this paper enables the verification of such a statement for all nodes, specifying parametrically the levels of "high confidence", $1-\beta_i$, and "safety threshold voltage", V_{Ti}. This leads to a statistical definition of safety, that we state as follows:

Definition 1. *Node i is said to be safe if $\mathcal{P}\{V_i < V_{Ti}\} > 1 - \beta_i$. Conversely, node i is said to be unsafe if $\mathcal{P}\{V_i < V_{Ti}\} < 1 - \beta_i$.*

In the above, $\mathcal{P}\{\}$ denotes a probability, V_i the voltage drop at node i (due to within-die leakage variations), V_{Ti} the safety threshold voltage at node i, and β_i a small positive number between 0 and 1. The term $1-\beta_i$ specifies the bulk of the voltage drop that is to be below the threshold V_{Ti} to assert that a node is safe, and we would refer to it as the *safety parameter*. The choice of the safety parameter and the safety threshold voltage at a given node takes into account how critical a high voltage drop at that node is, and we made explicit their dependence on the node. Given safety parameters and threshold voltages at all nodes, our objective is to *verify each node* in the grid, *i.e.*, tell which nodes/regions are safe with respect to the background noise, and which present a hazard.

Let $V_{1-\beta_i}$ be the $1-\beta_i$ percentile of the voltage drop at node i, i.e., $V_{1-\beta_i}$ is such that $F_{V_i}(V_{1-\beta_i}) = 1-\beta_i$, where $F_{V_i}(.)$ is the cdf of the distribution of the voltage drop at node i, defined in (11). Clearly, *Definition 1* can be restated in terms of $V_{1-\beta_i}$ as follows:

Definition 2. *Node i is said to be safe if $V_{1-\beta_i} < V_{Ti}$ and unsafe if $V_{1-\beta_i} > V_{Ti}$.*

Thus, $V_{1-\beta_i}$ becomes a *figure of merit* of the verification procedure. It can be viewed as a parametric measure of the voltage drop on the grid nodes taking into account process variations on the leakage currents: if the safety parameter is 50%, this percentile is the median voltage drop, 90%, 95% represent more conservative measures, and 100% represents an upper bound on the voltage drop. The verification problem reduces to determining whether the $1-\beta_i$ percentiles of voltage drops are greater or less than a given safety threshold for any node i.

Let z_{β_i} be such that $\Phi(z_{\beta_i}) = 1 - \beta_i$. From (11), we have:

$$\frac{\ln(V_{1-\beta_i}) - \mu_{W_i}}{\sigma_{W_i}} = z_{\beta_i}, \tag{37}$$

where $W_i = \ln(V_i)$ is a Gaussian RV with mean μ_{W_i} and standard deviation σ_{W_i}, as defined in section IV-A. From (37), and using (12) and (13), we can derive an expression for $V_{1-\beta_i}$ in terms of the mean and variance of V_i:

$$\begin{aligned} V_{1-\beta_i} &= \exp\left(z_{\beta_i}\sigma_{W_i} + \mu_{W_i}\right) \\ &= \mu_{V_i}\frac{\exp\left[z_{\beta_i}\left(\ln\left[1+\sigma_{V_i}^2/\mu_{V_i}^2\right]\right)^{1/2}\right]}{\left(1+\sigma_{V_i}^2/\mu_{V_i}^2\right)^{1/2}} \end{aligned} \tag{38}$$

The above equation relates $V_{1-\beta_i}$ to μ_{V_i} and $\sigma_{V_i}^2$. We argued in section IV-A that the mean can be easily computed. This leaves the variance of voltage drop as the only parameter on which $V_{1-\beta_i}$ depends, and we write $V_{1-\beta_i}(\sigma_{V_i}^2)$ to emphasize this dependence.

The following outlines our verification methodology. We first rewrite the safety/unsafety definition (*Definition* 2) in terms of voltage variances and derive critical variance values for verifying the node safety status (section VI-B). Then we derive quick direct checks to see where the actual variance falls with respect to those critical values (section VII-A). For nodes where direct checks do not provide conclusive information, we make use of the numerical Monte Carlo technique described in section V-B to estimate the variance of the voltage at these nodes (section VII-B), and hence the $1-\beta$ percentile voltage drops, and determine the node safety status.

B. Critical Variances

Writing $V_{1-\beta_i}(\sigma_{V_i}^2)$ as a function of $\sigma_{V_i}^2$, and using (38), one can easily show the following:

1) $V_{1-\beta_i}(0) = \mu_{V_i}$.
2) $V_{1-\beta_i}(.)$ admits one local maximum such that:

$$\begin{aligned} \max_{\sigma_{V_i}^2} V_{1-\beta_i}(\sigma_{V_i}^2) &= V_{1-\beta_i}\left(\mu_{V_i}^2\left[\exp(z_{\beta_i}^2) - 1\right]\right) \\ &= \mu_{V_i}\exp(z_{\beta_i}^2/2). \end{aligned} \tag{39}$$

3) $\lim_{\sigma_{V_i}^2 \to \infty} V_{1-\beta_i}(\sigma_{V_i}^2) = 0.$

Fig. 3 shows a typical plot of $V_{1-\beta_i}(\sigma_{V_i}^2)$.

In order to translate the safety/unsafety criteria on $V_{1-\beta_i}$ to conditions on the variances, nodes will be divided into three groups according to the values of V_{Ti}, μ_{V_i}, and z_{β_i}.

Group 1. *Includes all nodes i such that:*

$$V_{Ti} > \mu_{V_i}\exp(z_{\beta_i}^2/2).$$

These nodes will all satisfy $V_{1-\beta_i} < V_{Ti}$, for all possible values of the variance of their voltage drop $(0 < \sigma_{V_i}^2 < \infty)$, and therefore are safe irrespective of their variances.

The reason nodes satisfying the *Group 1* criterion are always safe can be seen graphically by drawing a horizontal line at $V_{Ti} > \mu_{Vi}\exp(z_{\beta_i}^2/2)$. Such a line is always above the curve representing $V_{1-\beta_i}$, as can be noted from Fig. 3. Therefore, $V_{1-\beta_i} < V_{Ti}$ and these nodes are safe.

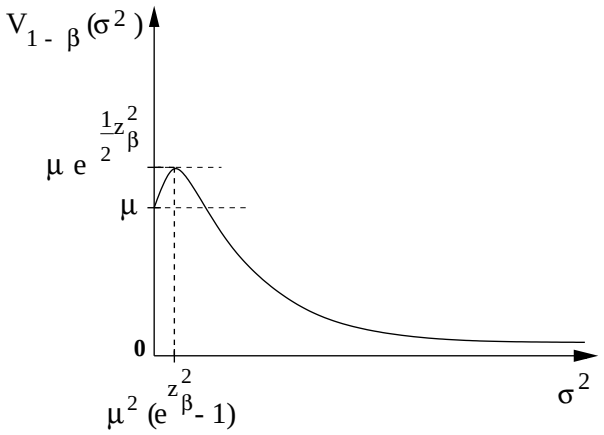

Fig. 3. A typical plot of $V_{1-\beta_i}$ versus $\sigma_{V_i}^2$.

Group 2. *Includes all nodes i such that:*

$$\mu_{V_i} < V_{Ti} < \mu_{V_i}\exp(z_{\beta_i}^2/2).$$

A node in this group is safe if and only if the variance of its voltage drop, $\sigma_{V_i}^2$, is such that: $\sigma_{V_i}^2 < \sigma_{1,i}^2$ or $\sigma_{V_i}^2 > \sigma_{2,i}^2$, where

$$\sigma_{1,i}^2 = \mu_{V_i}^2\exp\left[2z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}} - 2z_{\beta_i}\left(z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}}\right)^{1/2}\right] - \mu_{V_i}^2 \tag{40}$$

and

$$\sigma_{2,i}^2 = \mu_{V_i}^2\exp\left[2z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}} + 2z_{\beta_i}\left(z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}}\right)^{1/2}\right] - \mu_{V_i}^2 \tag{41}$$

Conversely, a node in this Group is unsafe iff $\sigma_{1,i}^2 < \sigma_{V_i}^2 < \sigma_{2,i}^2$ (see Fig. 4(a)).

Group 3. *Includes all nodes i such that: $V_{Ti} < \mu_{V_i}$. A node in this group is safe iff $\sigma_{V_i}^2 > \sigma_{3,i}^2$ and unsafe iff $\sigma_{V_i}^2 < \sigma_{3,i}^2$, where*

$$\sigma_{3,i}^2 = \mu_{V_i}^2\exp\left[2z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}} + 2z_{\beta_i}\left(z_{\beta_i}^2 - 2\ln\frac{V_{Ti}}{\mu_{V_i}}\right)^{1/2}\right] - \mu_{V_i}^2 \tag{42}$$

(see Fig. 4(b)).

With this, determining whether a node i is safe or unsafe reduces to knowing where the variance of the voltage at that node is located with respect to $\sigma_{1,i}^2$ and $\sigma_{2,i}^2$ for *Group 2* nodes and with respect to $\sigma_{3,i}^2$ for *Group 3* nodes. We refer to $\sigma_{1,i}^2$ and $\sigma_{2,i}^2$ as *critical variances* for a node i in *Group 2* and to $\sigma_{3,i}^2$ as the critical variance for a node i in *Group 3*. *Group 1* nodes are deemed safe, irrespective of the variance of their voltages. Observe that the group of each node as well as critical variances are known automatically, and require no a-priori knowledge of the variance of the node voltage drop, since z_{β_i} is directly obtained knowing β_i, V_{Ti} is a user-defined parameter, and the means of the voltage drops are easily calculated using (14).

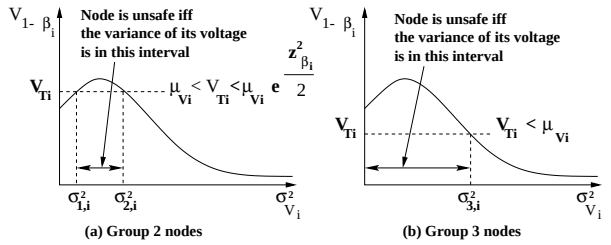

Fig. 4. Illustration of safety/unsafety conditions on the variances for *Group 2* and *Group 3* nodes.

VII. Checking Voltage Variances Against Critical Values

A. Direct Criteria

In this section, we make use of bounds that may directly reveal whether $V_{1-\beta_i}$ is greater or less than the safety threshold V_{Ti}, without having to compute the variances. From the previous section, if a node is in *Group 1*, then this node is safe for all possible values of the variance. The following provides similarly useful checks.

For convenience, we reproduce (21):

$$\sigma^2_{V_i} = [\Sigma_{\mathbf{V}}]_{ii} = \sum_{j=1}^{N} q_{ij}^2 \sigma^2_{I_j}, \tag{43}$$

where q_{ij} was defined to be the (i,j)th entry of $\mathbf{G}^{-1}$. Since $\sigma_{I_j} \geq 0$ and $q_{ij} \geq 0$, $\forall i,j$, then (43) yields:

$$\sigma^2_{V_i} = \sum_{j=1}^{N} (q_{ij}\sigma_{I_j})^2 \leq \left(\sum_j q_{ij}\sigma_{I_j} \right)^2. \tag{44}$$

Let $\mathbf{Std}(.)$ be the operator which takes a random vector and returns the corresponding vector of standard deviations. The above leads to the following upper bound on deviations of node voltage drops:

$$\mathbf{Std}(\mathbf{V}) \leq \mathbf{G}^{-1}\mathbf{Std}(\mathbf{I}), \tag{45}$$

where the vector inequality is defined component-wise. Given an LU-factorization of $\mathbf{G}$, and the variances of the within-die variability components of leakage currents, the cost of computing the upper bound given in (45) is only one forward/backward solve.

Now using the notation of section V-B, we can write from (24):

$$\sigma^2_{V_i} = \sigma_{ii} = S \sum_{j=1}^{N} p_j q_{ij}^2, \tag{46}$$

where $S = \sum_j \sigma^2_{I_j}$ and $p_j = \sigma^2_{I_j}/S$. Since $0 \leq p_j \leq 1$, then $p_j \geq p_j^2$, $\forall j$. Then we can write from (46):

$$\sigma^2_{V_i} = S \left(\sum_{j=1}^{N} p_j q_{ij}^2 \right) \geq \frac{1}{S} \left(\sum_{j=1}^{N} q_{ij}\sigma^2_{I_j} \right)^2. \tag{47}$$

This leads to a lower bound on $\mathbf{Std}(\mathbf{V})$, as follows:

$$\mathbf{Std}(\mathbf{V}) \geq \frac{1}{\sqrt{S}} \mathbf{G}^{-1}\mathbf{Var}(\mathbf{I}), \tag{48}$$

where $\mathbf{Var}(.)$ is the variance vector operator, defined in section IV-B. Given an LU-factorization of $\mathbf{G}$, the cost of computing this lower bound given in (48) is only one forward/backward solve.

Putting (45) and (48) together yields the following interval on $\mathbf{Std}(\mathbf{V})$:

$$\frac{1}{\sqrt{S}} \mathbf{G}^{-1}\mathbf{Var}(\mathbf{I}) \leq \mathbf{Std}(\mathbf{V}) \leq \mathbf{G}^{-1}\mathbf{Std}(\mathbf{I}). \tag{49}$$

The above inequality provides an upper and a lower bound on the standard deviations, and hence the variances, of the voltage drop at each node on the power grid, at the cost of 2 forward/backward solves.

To put these results to work, let u_i^2 and l_i^2 be respectively the upper and lower bounds on the *variance* of the voltage drop at node i. If node i is in *Group 2*, with corresponding critical variances $\sigma^2_{1,i}$ and $\sigma^2_{2,i}$, then this node can be deemed safe if ($u_i^2 < \sigma^2_{1,i}$ or $l_i^2 > \sigma^2_{2,i}$) and unsafe if ($l_i^2 > \sigma^2_{1,i}$ and $u_i^2 < \sigma^2_{2,i}$). Similarly, if node i is in *Group 3* with a corresponding critical variance $\sigma^2_{3,i}$, then it can be deemed safe if $l_i^2 > \sigma^2_{3,i}$ and unsafe if $u_i^2 < \sigma^2_{3,i}$.

B. Verification through column sampling

For nodes in *Group 2* or *Group 3* where the variance bounds given in (49) do not provide conclusive information regarding their safety status, we can estimate their voltage drop variances so as to determine the location of these variances with respect to critical values. For variance estimation, we make use of the numerical Monte Carlo method described in section V-B.

Observe that the variance estimates obtained by the Monte Carlo technique inherently carry a certain error, bounded by a given confidence level (as in section V). Some error will then be propagated to the estimates of the $1-\beta$ percentiles of voltage drops, as per (38), so that a certain confidence level will be required when asserting whether $V_{1-\beta_i}$ is greater or less than V_{Ti}. Accordingly, we extend *Definition* 2 to consider that a node is safe if $\mathcal{P}\{V_{1-\beta_i} < V_{Ti}\} \geq 1-\alpha$ and unsafe if $\mathcal{P}\{V_{1-\beta_i} > V_{Ti}\} \geq 1-\alpha$, where α is a small number between 0 and 1. In this perspective, for the nodes successfully checked in section VII-A, $\alpha = 0$.

Establishing the desired confidence level on the relative position of $V_{1-\beta_i}$ and V_{Ti} for node i translates to establishing that same confidence level on the relative position of the variance of the i^{th} node voltage and the corresponding critical variances ($\sigma^2_{1,i}$ and $\sigma^2_{2,i}$ if node i is in *Group 2* and $\sigma^2_{3,i}$ if node i is in *Group 3*). Therefore, when applying the Monte Carlo technique for variance estimation, the estimate of the variance per se matters less than where the variance is with respect to critical values. For example, suppose node i is in *Group 3*, variance estimation is only needed in order to obtain high confidence on whether $\sigma^2_{V_i}$ is less than (greater than) $\sigma^2_{3,i}$, so that the node can be marked as unsafe (safe). Therefore, while the sampling procedure as described in section V-B is unchanged, new error bounds, *i.e.*, stopping criteria, must be derived.

A few cases need to be considered. If node i is in *Group 2*, the estimate of the variance of the voltage drop at node i, $\hat{\sigma}_{ii}$ (following the notation of section V-B), may fall to the left of

$\sigma^2_{1,i}$, to the right of $\sigma^2_{2,i}$, or between the two, If node i is in *Group 3*, $\hat{\sigma}_{ii}$ may fall to the left or to the right of $\sigma^2_{3,i}$. We thus have variance intervals defined with respect to $\sigma^2_{1,i}$ and $\sigma^2_{2,i}$, or $\sigma^2_{3,i}$. We recognize that the probability of the true variance being outside the interval where the variance estimate lies is not very high, and indeed for α small enough, less than $1-\alpha$. Based on this, in the sampling process, we seek to find the smallest number of samples n that verifies, with confidence $1-\alpha$, whether the node is safe or unsafe, according to the interval where the variance estimate lies.

Assume node i is in *Group 2*. Then as was shown in section VI-B, $V_{1-\beta_i} < V_{Ti}$ is equivalent to $\sigma^2_{V_i} < \sigma^2_{1,i}$ or $\sigma^2_{V_i} > \sigma^2_{2,i}$, and $V_{1-\beta_i} > V_{Ti}$ is equivalent to $\sigma^2_{1,i} < \sigma^2_{V_i} < \sigma^2_{2,i}$. If $\hat{\sigma}_{ii} < \sigma^2_{1,i}$, $1-\alpha$ confidence that the node is safe is reached when:

$$\mathcal{P}\{V_{1-\beta_i} < V_{Ti}\} \geq 1-\alpha,$$

The above can be written as:

$$\mathcal{P}\{\sigma^2_{V_i} < \sigma^2_{1,i}\} + \mathcal{P}\{\sigma^2_{V_i} > \sigma^2_{2,i}\} \geq 1-\alpha,$$

leading to:

$$\mathcal{P}\left\{\frac{\sigma^2_{V_i}-\hat{\sigma}_{ii}}{\sqrt{\mathrm{Var}(\hat{\sigma}_{ii})}} < \frac{\sigma^2_{1,i}-\hat{\sigma}_{ii}}{\sqrt{\mathrm{Var}(\hat{\sigma}_{ii})}}\right\}$$
$$+\mathcal{P}\left\{\frac{\sigma^2_{V_i}-\hat{\sigma}_{ii}}{\sqrt{\mathrm{Var}(\hat{\sigma}_{ii})}} > \frac{\sigma^2_{2,i}-\hat{\sigma}_{ii}}{\sqrt{\mathrm{Var}(\hat{\sigma}_{ii})}}\right\} \quad \geq 1-\alpha,$$

where $\mathrm{Var}(.)$ represents the variance of a random variable. ¿From the arguments presented in section V-B, we know that $(\sigma^2_{V_i}-\hat{\sigma}_{ii})/\sqrt{\mathrm{Var}(\hat{\sigma}_{ii})}$ is standard normal, and that $\mathrm{Var}(\hat{\sigma}_{ii}) = S^2 s^2_{ii}/n$, where S and s_{ii} are as defined in section V-B and n is the number of samples. Therefore, the above condition reduces to:

$$\Phi\left(\frac{(\sigma^2_{2,i}-\hat{\sigma}_{ii})\sqrt{n}}{Ss_{ii}}\right) - \Phi\left(\frac{(\sigma^2_{1,i}-\hat{\sigma}_{ii})\sqrt{n}}{Ss_{ii}}\right) \leq \alpha. \quad (50)$$

That is, if $\hat{\sigma}_{ii} < \sigma^2_{1,i}$, then a $1-\alpha$ confidence level that node i is safe is attained if and only if (iff) n satisfies (50). Note that (50) does not give a closed-form solution for n, but it is easy to check it using the *erf*(.) function. Observing that in this case, $\hat{\sigma}_{ii} < \sigma^2_{1,i}$, we can obtain a closed-form *sufficient* condition for n to satisfy in order to verify safety, by neglecting $\mathcal{P}\{\sigma^2_{V_i} > \sigma^2_{2,i}\}$, yielding:

$$n \geq \left(\frac{Ss_{ii}z_\alpha}{\sigma^2_{1,i}-\hat{\sigma}_{ii}}\right)^2 = n_1, \quad (51)$$

where z_α is such that: $\Phi(z_\alpha) = 1-\alpha$.

Identical reasoning applies when $\hat{\sigma}_{ii} > \sigma^2_{2,i}$, and we obtain that the necessary and sufficient condition on n to verify safety is:

$$\Phi\left(\frac{(\hat{\sigma}_{ii}-\sigma^2_{1,i})\sqrt{n}}{Ss_{ii}}\right) - \Phi\left(\frac{(\hat{\sigma}_{ii}-\sigma^2_{2,i})\sqrt{n}}{Ss_{ii}},\right) \leq \alpha,$$

and that a sufficient condition that yields a closed-form for n is:

$$n \geq \left(\frac{Ss_{ii}z_\alpha}{\hat{\sigma}_{ii}-\sigma^2_{2,i}}\right)^2 = n_2. \quad (52)$$

Finally, if $\sigma^2_{1,i} < \hat{\sigma}_{ii} < \sigma^2_{2,i}$, we need to find the number of samples that will establish $1-\alpha$ confidence that node i is unsafe, *i.e.*, that $\sigma^2_{1,i} < \sigma^2_{V_i} < \sigma^2_{2,i}$. Extension of the above arguments leads to the following necessary and sufficient condition:

$$\Phi\left(\frac{(\hat{\sigma}_{ii}-\sigma^2_{1,i})\sqrt{n}}{Ss_{ii}}\right) + \Phi\left(\frac{(\sigma^2_{2,i}-\hat{\sigma}_{ii})\sqrt{n}}{Ss_{ii}}\right) \leq 2-\alpha.$$

In order to write a closed-form sufficient condition, we recall that a $1-\alpha$ confidence interval on $\sigma^2_{V_i}$ is given by [36]:

$$\hat{\sigma}_{ii} \pm \frac{Ss_{ii}z_{\alpha/2}}{\sqrt{n}}.$$

So, to obtain a $1-\alpha$ confidence level that $\sigma^2_{V_i}$ is between $\sigma^2_{1,i}$ and $\sigma^2_{2,i}$, it is sufficient that both extremes of the above interval lie within $[\sigma^2_{1,i}, \sigma^2_{2,i}]$. This leads to the following condition:

$$n \geq \left(\frac{Ss_{ii}z_{\alpha/2}}{\min(\hat{\sigma}_{ii}-\sigma^2_{1,i}, \sigma^2_{2,i}-\hat{\sigma}_{ii})}\right)^2 = n_3. \quad (53)$$

Obtaining bounds for nodes in *Group 3* is easier as we have only one critical variance. If $\hat{\sigma}_{ii} > \sigma^2_{3,i}$, then we need to check for safety, i.e.:

$$\mathcal{P}\{\sigma^2_{V_i} > \sigma^2_{3,i}\} \geq 1-\alpha.$$

This is equivalent to:

$$\Phi\left(\frac{(\hat{\sigma}_{ii}-\sigma^2_{3,i})\sqrt{n}}{Ss_{ii}}\right) \geq 1-\alpha,$$

which reduces to:

$$n \geq \left(\frac{Ss_{ii}z_\alpha}{\hat{\sigma}_{ii}-\sigma^2_{3,i}}\right)^2 = n_4. \quad (54)$$

Note that (54) is a necessary and sufficient condition on n to achieve $1-\alpha$ confidence that node i is safe. Similarly, if $\hat{\sigma}_{ii} < \sigma^2_{3,i}$, we obtain $1-\alpha$ confidence that node i is unsafe iff n satisfies:

$$n \geq \left(\frac{Ss_{ii}z_\alpha}{\sigma^2_{3,i}-\hat{\sigma}_{ii}}\right)^2 = n_5. \quad (55)$$

Note that $n_4 = n_5$, and that (54) and (55) are closed-form necessary and sufficient conditions on n, the required number of samples. The following summarizes convergence of node i after n samples.

```
If node i is in Group 2:
  If σ̂_ii < σ²_{1,i}:
    If n ≥ n_1: the node is done -safe.
  Else if σ̂_ii > σ²_{2,i}:
    If n ≥ n_2: the node is done -safe.
  Else if σ²_{1,i} < σ̂_ii < σ²_{2,i}:
    If n ≥ n_3: the node is done -unsafe.
Else if node i is in Group 3:
  If σ̂_ii > σ²_{3,i}:
    If n ≥ n_4: the node is done -safe.
  Else If σ̂_ii < σ²_{3,i}:
    If n ≥ n_5: the node is done -unsafe.
```

C. Residual Nodes

It can be seen from (51) – (55) that the number of samples required to establish the desired confidence level may be large if the estimated variance is very close to $\sigma^2_{1,i}$, $\sigma^2_{2,i}$, or $\sigma^2_{3,i}$. Suppose for example that the variance estimator of a node in *Group 3* is very close to $\sigma^2_{3,i}$. The variance itself may be either greater or less than $\sigma^2_{3,i}$, but the closer the estimator is to $\sigma^2_{3,i}$, the harder it is to establish a high confidence level on where the true variance actually lies.

With an estimate $\hat{\sigma}_{ii}$ of $\sigma^2_{V_i}$, we use (38) to define a natural estimator of $V_{1-\beta_i}$ given by:

$$\hat{V}_{1-\beta_i} = V_{1-\beta_i}(\hat{\sigma}_{ii}) = \mu_{V_i} \frac{\exp\left[z_{\beta_i}\left(\ln\left[1+\hat{\sigma}_{ii}/\mu^2_{V_i}\right]\right)^{1/2}\right]}{\left(1+\hat{\sigma}_{ii}/\mu^2_{V_i}\right)^{1/2}}. \tag{56}$$

Hence, if $\hat{\sigma}_{ii}$ is close to $\sigma^2_{1,i}$, $\sigma^2_{2,i}$, or $\sigma^2_{3,i}$, then $\hat{V}_{1-\beta_i}$ is correspondingly close to V_{Ti}. In this case, instead of seeking to establish a high confidence level on whether $V_{1-\beta_i}$ is greater or less than V_{Ti}, we estimate an upper bound (a conservative value) on $V_{1-\beta_i}$, with $1-\alpha$ confidence, that we denote $V_{ub,i}$.

This can be achieved in the following way. Let δV_{dd} be a user-defined resolution on the estimation of $V_{1-\beta_i}$. Let $\mathcal{D}$ denote the subset of nodes which have not converged and $\mathcal{R}$ the subset of $\mathcal{D}$ including all nodes i such that $|\hat{V}_{1-\beta_i} - V_{Ti}| \leq \delta V dd$. If at any time in the iteration process, $\mathcal{R} = \mathcal{D}$, and $\mathcal{R}$ is not empty, then we stop iterating and we call $\mathcal{R}$ the set of *residual nodes*.

We know from section VII-B that with $1-\alpha$ confidence, the variance of each residual node lies in the interval $[v_1, v_2]$, where:

$$v_1 = \hat{\sigma}_{ii} - (S s_{ii} z_{\alpha/2})/\sqrt{n}$$

and

$$v_2 = \hat{\sigma}_{ii} + (S s_{ii} z_{\alpha/2})/\sqrt{n}.$$

Knowing the variations of $V_{1-\beta_i}(\sigma^2_{V_i})$ versus $\sigma^2_{V_i}$ (see Fig. 3), it is easy to determine the point v_m in $[v_1, v_2]$ such that $V_{1-\beta_i}(v_m)$ is largest. Knowing the variations of $V_{1-\beta_i}(\sigma^2_{V_i})$ versus $\sigma^2_{V_i}$, described in section VI-B, and depending on the values of $\hat{\sigma}_{ii}$, v_1, and v_2, v_m can be equal to v_1, v_2, or $\mu^2_{V_i}[\exp(z^2_{\beta_i}) - 1]$, as follows:

$\texttt{If}\ \ v_1 < \mu^2_{V_i}\left[\exp(z^2_{\beta_i}) - 1\right] < v_2\texttt{:}$
$\quad v_m = \mu^2_{V_i}\left[\exp(z^2_{\beta_i}) - 1\right].$
$\texttt{If}\ \ v_1 < v_2 < \mu^2_{V_i}\left[\exp(z^2_{\beta_i}) - 1\right]\texttt{:}\ \ v_m = v_2.$
$\texttt{If}\ \ \mu^2_{V_i}\left[\exp(z^2_{\beta_i}) - 1\right] < v_1 < v_2\texttt{:}\ \ v_m = v_1.$

$V_{ub,i}$ can then be written as:

$$V_{ub,i} = V_{1-\beta_i}(v_m). \tag{57}$$

Fig. 5 illustrates a case for a residual *Group 3* node.

Residual nodes are simply nodes having their $V_{1-\beta}$ very close to the safety threshold voltage that it becomes difficult to tell whether they are safe or not. They are on the verge of safety or unsafety, as they were defined. Observe that $V_{ub,i}$ is always greater than V_{Ti} (otherwise node i would have converged), therefore, $V_{ub,i} - V_{Ti}$ can be viewed as the required increase in the safety threshold on node i to establish safety on that node.

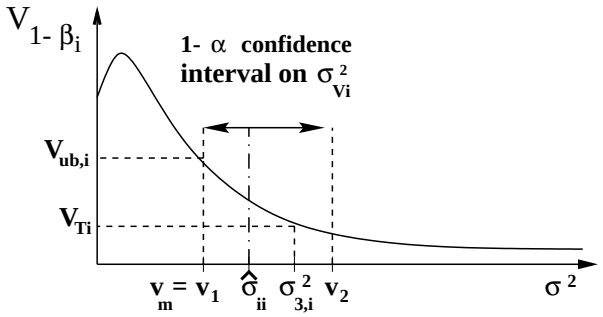

Fig. 5. Finding an upper bound with $1-\alpha$ confidence on $V_{1-\beta_i}$ for a residual, *Group 3* node.

VIII. Results

A. Test Grids

The proposed analysis and verification methods have been implemented and tested on a number of test-case grids. Not having access to power grids from industrial designs, and because we need a large and number of grids under different conditions, we have opted to *generate* a number of grids ourselves. In our test grids, technology and topology parameters are user-specified. The grid generation process is automatic and employs a pseudo-random number generator, based on the built-in function *random(.)* in C. Starting with a square uniform grid of a given size, we proceed to randomly delete a user-specified percentage of nodes (which we call the *degree of non-uniformity*), thus rendering the grid structurally non-uniform. Typical geometric and physical grid characteristics (e.g. grid dimensions) and characteristics of the fabrication process (e.g. sheet resistance of a particular level of metallization) are given by the user, leading to an initial value of the conductance of every branch. When a node is deleted, the conductances of the remaining surrounding edges (branches) are increased by a random amount around a user-specified percentage of their initial values. The rationale behind this is to allow the non-uniform grid to be loaded with currents comparable to its uniform predecessor while exhibiting comparable IR-drops. The number of V_{dd} (C4) sites is supplied by the user and C4s are placed randomly on the grid. The C4s and current sources are then distributed at random over the grid nodes. The user also specifies a typical value of leakage current variance and mean and the assigned values are generated randomly around these typical values; Current sources are then distributed at random over the grid. All results were obtained on a 1 GHz, dual-processor Sun Fire server with 4.0 GB of main memory. The implementation was carried out by writing C programs, and LU factorization and forward/backward solves done through the package UMFPACK [37], using Basic Linear Algebra Subprograms (BLAS) routines from the Sun Performance Library [38].

B. Lognormality of Voltage Drops

Fig. 6 corroborates the fact that voltage drops are well modeled by a lognormal distribution. We randomly generated $n_s = 10,000$ independent load vectors, with each individual current load following a lognormal distribution, and with different current loads in any given vector being statistically

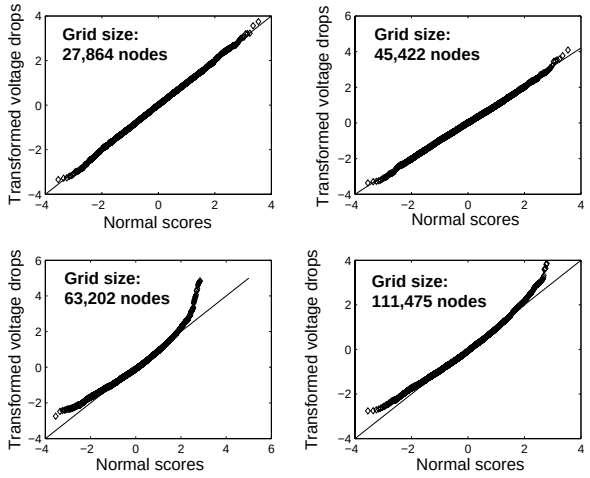

Fig. 6. Checking graphically the goodness-of-fit of the voltage drop data against a lognormal distribution using the method of normal scores.

independent from one another. We collected voltage drop data on all grid nodes, in response to each of the randomly generated load vectors, then computed the natural logarithms of each node voltage drop, *i.e.*, $\ln(V_{i,j})$, where $V_{i,j}$ is the voltage drop of node i and in response to current vector j, for $i = 1, \cdots, N$, and $j = 1, \cdots, n_s$. We then obtained the sample mean $\bar{\mu}_i$ and sample standard deviation s_i of each $\ln(V_i)$, for $i = 1, \cdots, N$. The y-axes in Fig. 6 represent voltage drop data transformed as $(\ln(V_{i,j}) - \bar{\mu}_i)/s_i$, $\forall i, j$. The point being to verify lognormality of all voltage drops, the task becomes to check these transformed voltage drops, which are statistically independent, against a standard normal distribution. This was done graphically, and the data was fit on normal scores plots [33], shown in Fig. 6. As can be seen from the figure, voltage drops showed good fits, except for certain outliers, validating the choice of lognormal distributions to model the voltage drops on the power grid, induced by independent, within-die leakage current variations.

C. Variance Estimation

Fig. 7 verifies the accuracy of the proposed column-sampling method for variance estimation. The correlation plot in Fig. 7(a), shows the estimated standard deviations versus exact ones, for a grid of 138,500 nodes, obtained after convergence of 100% of the grid nodes, reached after 14,589 samples. The estimated standard deviations are obtained via the column-sampling approach and the exact standard deviations obtained via the deterministic solution of (23). The plot confirms the accuracy of the results. The error distribution is shown in Fig. 7(b): only 1 among 138,500 data points (exact standard deviation = 1.25%V_{dd} and estimated standard deviation = 1.76%V_{dd}). carried an error larger than 0.5%V_{dd}, and this error is 0.51%V_{dd}. 70 data points carried errors larger than than 0.4%V_{dd}. The average error over all variance estimates is 0.06%V_{dd}. For reference, we note that the mean standard deviation of the voltage drop over all nodes is 0.99%V_{dd} and the maximum standard deviation of the voltage drop at any node is 3.77%V_{dd}. The mean and maximum standard deviation are obtained from the exact values, representing the

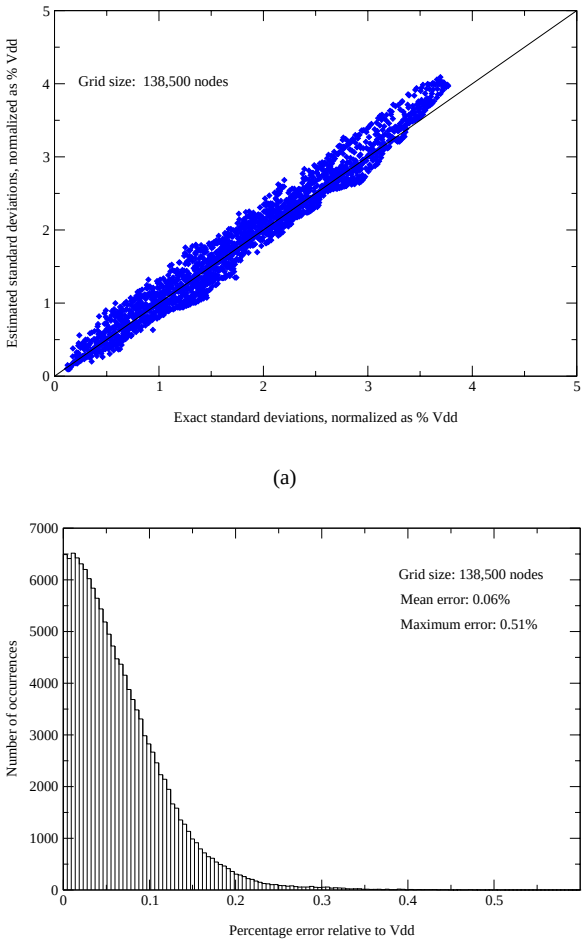

Fig. 7. Correlation plots (a) and error distribution (b) of variance estimates.

x-axis of Fig. 7(a). Clearly, the number of samples required for convergence depends on the error bound δ and the actual variance values to be estimated. For a given power grid and with a fixed error bound, it can be observed from (35) that as the standard deviations to be estimated increase, the number of required samples increases quadratically.

An important point, underscored in Fig. 8 is that a small fraction of nodes require a very large number of samples to reach convergence, as defined in (36): in the figure, and for a grid of 264,592 nodes, 95% of grid nodes converge after 1,566 samples, 99% converge after 3,304 samples, but it takes more than 13,000 samples for all nodes to reach convergence. In effect, for the "difficult" 1% of nodes, the $1-\alpha$ confidence interval after 3,304 samples extends beyond the error bound (δV_{dd}).

Fig. 9 shows the number of samples needed for convergence of 100% of the grid nodes as the error bound (δ) is varied, as a percentage of the supply voltage V_{dd}. The figure shows a sharp increase in the total number of samples for an error bound of 1% of V_{dd} or less. Indeed, if δ is greater than 5% of V_{dd}, convergence is attained immediately after 50 samples, while more than 5,000 samples would be needed to attain convergence when the error bound is fixed at less than 1% of V_{dd}.

TABLE I

VARIANCE ESTIMATION RESULTS ON GRIDS THAT ALLOW A DETERMINISTIC SOLUTION.

Size (#nodes)	# C4s	Degree (%) of non-uniformity	Std. Dev. Avg./Max. (% V_{dd})	Avg. Error (% V_{dd})	Max. Error (% V_{dd})	%nodes with error $> 1\%V_{dd}$	# samples	Runtime: LU	Runtime: column sampling	Runtime: deterministic solution
17,678	150	3	1.13/3.44	0.13	1.51	0.28	1,229	$<$ 1 sec.	45 sec.	9.5 min.
20,767	100	18	2.05/5.36	0.13	1.06	0.005	1,074	$<$1 sec.	46 sec.	12.8 min.
32,130	120	3	1.68/5.82	0.14	1.19	0.05	2,084	1.8 sec.	4.7 min.	1.1 hr.
45,980	400	5	1.04/3.11	0.14	1.11	0.004	2,699	2.5 sec.	8.9 min.	2.5 hrs.
50,327	100	18	2.02/5.97	0.15	1.88	0.35	1,851	2.2 sec.	6.2 min.	2.6 hrs.
59,850	250	5	0.94/2.30	0.13	0.90	0	1,410	3.3 sec.	7.1 min.	4.6 hrs.
76,484	260	19	2.52/7.86	0.20	2.04	0.35	4,861	3.5 sec.	29 min.	7.4 hrs.
89,302	300	10	1.12/3.14	0.13	1.63	0.01	2,435	5 sec.	20 min.	11 hrs.
100,796	200	8	1.56/4.07	0.15	1.33	0.18	2,390	5.8 sec.	23 min.	15.1 hrs.
142,956	400	1	1.04/3.61	0.14	1.22	0.003	2,743	10 sec.	43 min.	35.8 hrs.

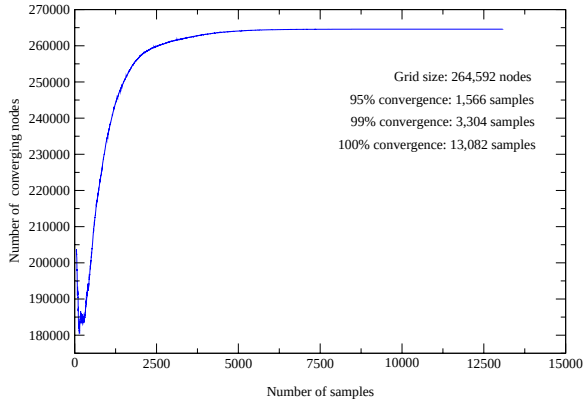

Fig. 8. Convergence progress.

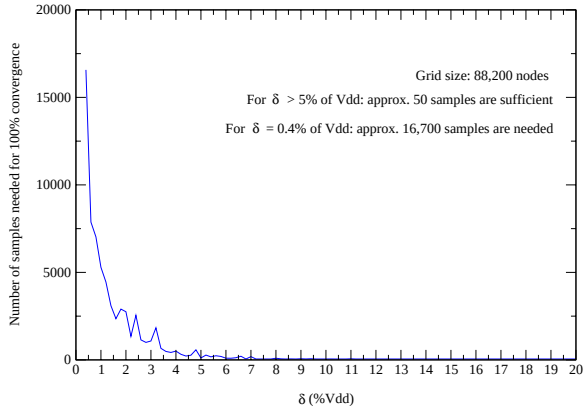

Fig. 9. Convergence versus error bound.

Table I shows characteristics and results for a number of power grids whose sizes were chosen small enough to allow one to compute the voltage variances by solving (23) deterministically, and compare results with the proposed Monte Carlo approach. The statistical estimation was conducted at 95% confidence ($\alpha = 0.05$) and with a 1% error bound ($\delta = 0.01$). The fourth column of the table shows the mean and maximum voltage drop standard deviation, taken over all the grid nodes. The values and spread of the voltage standard deviation continue the trend observed in Fig. 7. The column-sampling process was stopped when 99% of grid nodes reached convergence, as specified in (36). Since each column sample contributes to the estimation of every node voltage drop, each sample serves to update the variance estimate at all nodes, so that by the time the last column sample is taken, the confidence interval on the standard deviation estimate of the majority of nodes is smaller than the error bound (δV_{dd}). This explains the relatively small average errors in the node voltage standard deviation estimates, shown in the fourth column, which are found to be in the order of 0.1%-0.2% of V_{dd}. By the same token, we expect the number of nodes where the error in the standard deviation estimate is actually greater than δV_{dd} to be small relative to the grid size. Indeed, the sixth column shows the percentage of nodes (including non-converging ones) where the error in the voltage standard deviation estimate is in fact greater than the error bound of 1% of V_{dd}. In most cases, this percentage is effectively smaller than 1%, the percentage of nodes which did not converge. The fifth column reports the maximum error on any node: in each case, this error is quite comparable to the error bound, mostly in the range of 1%-2% of V_{dd}, and exceeding 2% of V_{dd} only in 1 in 10 cases. The last two columns of Table I compare the runtime of the column-sampling method with the deterministic solution of (23). The speedup ranges from about 12.7X, for the 17,678-node grid, to about 50X for the 142,956-node grid, and is, on average, about 26X.

Results on larger grids are shown in Table II. These grids are too large for the deterministic solution of node voltage variances to be carried out, therefore only the proposed column-sampling approach was applied to them, and only runtime and memory usage data are shown. It is clear that this technique can be applied to relatively large grids, with a tight (δ =1%, in this case) error bound.

D. Verification of Power Grids

Tables III and IV show the overall performance of the proposed verification approach on grids of various sizes. The grids in the first table are small enough so that we were also able to solve (23) deterministically to obtain the exact value of the variance at each node, and consequently, the exact value of the $1 - \beta$ percentile of the voltage drop. The tables show runtimes for LU factorization, application of the direct checks (see section VII-A), and verification through column sampling (see section VII-B). Table III further

TABLE II

VARIANCE ESTIMATION RESULTS ON LARGER GRIDS THAT DO NOT ALLOW A DETERMINISTIC SOLUTION.

Size (#nodes)	Number of C4s	Degree (%) of non-uniformity	Number of samples	Runtime: LU	Runtime: sampling	Memory usage
164,510	550	23	3,610	8 sec.	53 min.	125 MB
203,251	580	6	3,937	14 sec.	1.4 hrs.	186 MB
251,856	590	10	5,527	17 sec.	2.5 hrs.	212 MB
369,024	568	4	6,663	33 sec.	5.0 hrs.	375 MB

indicates the time required for carrying out grid verification by a deterministic solution of (23), allowing us to compute exact values for grid voltage variances and therefore, to determine the safety status of each node with 100% confidence, *i.e.*, without resorting to statistical variance estimation through column sampling. With this, we can measure the error of our verification approach when column sampling is put to use: the last column of Table III reports the resulting errors, defined as the ratio of the number of nodes which were deemed safe and are actually unsafe and vice versa, to the total number of nodes, excluding residual nodes.

Observe the difference in number of samples required for variance estimation (Tables I and II) and grid verification (Tables III and IV), which directly translates in significantly smaller runtimes for grid verification. This difference is due to the fact that *Group 1* nodes and upper and lower variance bounds rule out the need for column sampling to estimate voltage variances for a fraction of nodes. Also, the sampling process for grid verification does not seek an estimate of the node voltage variance per se, but rather to establish a given confidence level on whether this variance is greater or less than a critical value. This implies a difference in the stopping criteria used for column sampling between the variance estimation and grid verification problems, such that sampling for grid verification may stop when establishing the desired confidence level on the location of the variance estimate with respect to critical variances, without having an accurate estimate of the variance itself.

We observed in our experiments that some grids were fully verifiable by the direct criteria, described in section VII-A, in which case the run time is dramatically reduced. The first grid in Table III and the second grid in Table IV are examples of such grids. The errors, shown in Table III, were very much in accordance with the specified bounds, and in the case of grids there were verifiable using only direct criteria, the error is 0.

Fig. 10 shows distribution of the distance between the upper bounds on $V_{1-\beta_i}$ and the safety threshold voltages V_{T_i}, for the grid of 1.3 million nodes, featured in the last row of Table IV. The safety parameter $(1-\beta)$ was set at (90%) for all nodes and the confidence level for convergence $(1-\alpha)$ at 95%. The resolution was fixed at 1% of V_{dd} ($\delta = 0.01$). Fig. 10 shows that there were 14,362 residual nodes, corresponding to 1.1% of the grid size. The number of such nodes is primarily related to the user-specified resolution and how it compares with the actual variances of the node voltages. All experiments feature a resolution of 1% of V_{dd}, and the number of residual nodes was consistently small, compared to the grid size (see also Tables III and IV). It was stated in section VII-C that the distance between $V_{ub,i}$ and V_{Ti} is an absolute measure of how much the requirement on the safety threshold voltage at a residual node should be relaxed (*i.e.*, how much V_{Ti} should increase) in order to deem a node safe. Fig. 10 illustrates that this distance is small (the average being 0.96% of V_{dd}), implying that $V_{1-\beta_i}$ and V_{Ti} are indeed close (that closeness being controlled by the resolution).

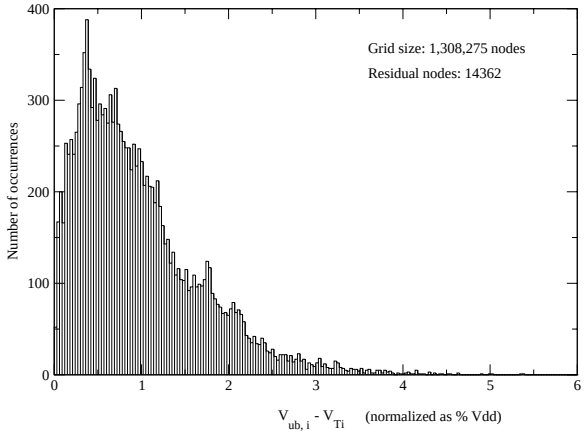

Fig. 10. Distribution of upper bounds on the $1-\beta$ percentiles.

IX. CONCLUSION

Technology scaling trends exacerbate the importance of proper leakage management and reinforce the need for robust power distribution across the chip. This work addressed the vulnerability of chip power distribution to process-induced leakage current variations. Leakage variations appear as random drops of the power grid voltage levels and may be thought of as a background level of noise on the grid. We suggested a lognormal statistical model for the voltage drops on the grid in presence of this noise and a numerical Monte Carlo technique to overcome the computational complexity associated with the grid size in order to determine the statistical parameters of this leakage-induced noise. With this ability, we were able to formulate a verification methodology in order to localize areas of the grid that are particularly susceptible to voltage violations due to leakage variations. Results showed the efficiency and applicability of the proposed approach.

X. ACKNOWLEDGEMENTS

The authors thank the anonymous reviewers of this paper for their comments and suggestions which helped improve this work.

TABLE III

VERIFICATION RESULTS ON SMALL GRIDS, WHERE CHECKING THE ACCURACY OF THE PROPOSED APPROACH IS POSSIBLE.

Size (#nodes)	Safety parameter $(1-\beta)$	Confidence level $(1-\alpha)$	Safety threshold (% of V_{dd})	% nodes safe	% nodes residual	Runtime: LU	Runtime: direct checks	Runtime: column sampling	Runtime: deterministic solution	% error
20,814	99%	95%	10%	100	0	1.1 sec.	0.00 sec	0	19 min.	0
40,804	92%	95%	10%	77.8	2.5	2.4 sec.	$<$ 1 sec.	6.2 sec.	1.9 hr.	1.3
51,711	99%	99%	10%	99.86	0.1	2.8 sec.	$<$ 1 sec.	7.8 sec.	3.2 hrs.	0.00
71,084	99%	99%	10%	95.8	0.3	3.7 sec.	$<$ 1 sec.	10.7 sec.	7 hrs.	0.00
79,925	95%	95%	10%	81.8	0.6	4.8 sec.	$<$ 1 sec.	14.9 sec.	9.3 hrs.	0.11
104,125	95%	95%	10%	91.9	0.6	5.4 sec.	1.2 sec.	18.9 sec.	13 hrs.	1.2

TABLE IV

PERFORMANCE OF THE PROPOSED VERIFICATION APPROACH ON LARGE GRIDS.

Size (#nodes)	Safety parameter $(1-\beta)$	Confidence level $(1-\alpha)$	Safety threshold (% of V_{dd})	% nodes safe	% nodes residual	Time (LU)	Time (direct)	# samples (iterative)	Time (iterative)	Memory usage
307,655	95%	92%	15%	65.4	5.4	25 sec.	5 sec.	433	15.2 min.	311 MB
415,410	95%	90%	10%	100	0	32 sec.	7 sec.	0	0	360 MB
691,850	95%	95%	10%	92.8	0.01	1.3 min.	16 sec.	2,484	4.1 hrs.	760 MB
811,912	95%	95%	15%	83.9	1.0	1.6 min.	18 sec.	2,440	5.7 hrs.	800 MB
1,008,899	99%	99%	10%	97.5	1.9	2.8 min.	23 sec.	131	21.7 min	1.3 GB
1,308,275	90%	95%	10%	95.3	1.1	4.7 min.	36 sec.	64	14 min.	1.8 GB

REFERENCES

[1] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," *Proc. IEEE*, vol. 91, no. 2, pp. 305–327, Feb. 2003.

[2] Y.-S. Lin, C.-C. Wu, C.-S. Chang, R.-P. Yang, W.-M. Chen, J.-J. Liaw, and C. H. Diaz, "Leakage scaling in deep submicron CMOS for SoC," *IEEE Trans. Electron Devices*, vol. 49, no. 6, pp. 1034–1041, June 2002.

[3] "The International Technology Roadmap for Semiconductors," 2003. [Online]. Available: http://public.itrs.net.

[4] R. Rao, A. Srivastava, D. Blaauw, and D. Sylvester, "Statistical analysis of subthreshold leakage current for VLSI circuits," *IEEE Trans. VLSI Syst.*, vol. 1, no. 2, pp. 131–139, Feb. 2004.

[5] V. De and S. Borkar, "Technology and design challenges for low power and high performance," in *ACM/IEEE International Symposium on Low Power Electronics and Design*, San Diego, CA, August 16-17 1999, pp. 163–168.

[6] Y. Taur, "CMOS design near the limit of scaling," *IBM J. Res. & Dev.*, vol. 46, no. 2/3, pp. 213–222, March-May 2002.

[7] T. Kam, S. Rawat, D. Kirkpatrick, R. Roy, G. S. Spirakis, N. Sherwani, and C. Peterson, "EDA challenges facing future microprocessor design," *IEEE Trans. Computer-Aided Design*, vol. 19, no. 12, pp. 1498–1506, Dec. 2000.

[8] Y. Taur, D. A. Buchanan, W. Chen, D. J. Frank, K. E. Ismail, S. Lo, G. A. Sai-Halasz, R. A. Viswanathan, H. C. Wann, S. J. Wind, and H. Wong, "CMOS scaling into the nanometer regime," *Proc. IEEE*, vol. 85, no. 4, pp. 486–504, Apr. 1997.

[9] R. Ahmadi and F. N. Najm, "Timing analysis in presence of power supply and ground voltage variations," in *IEEE/ACM International Conference on Computer-Aided Design*, San Jose, CA, November 9-13 2003, pp. 176–183.

[10] S. Pant, D. Blaauw, S. Sundareswaran, V. Zolotov, and R. Panda, "Vectorless analysis of supply noise induced delay variation," in *IEEE/ACM International Conference on Computer-Aided Design*, San Jose, CA, November 9-13 2003, pp. 184–191.

[11] "Power grid verification," White Paper, Cadence Design Systems, San Jose, CA, 2002. [Online]. Available: http://www.cadence.com/whitepapers/powerdistplan.html

[12] I. A. Ferzli and F. N. Najm, "Statistical estimation of leakage-induced power grid voltage drop considering within-die process variations," in *ACM/IEEE Design Automation Conference*, Anaheim, CA, June 2-6 2003, pp. 856–859.

[13] ——, "Statistical verification of power grids considering process-induced leakage current variations," in *IEEE/ACM International Conference on Computer-Aided Design*, San Jose, CA, November 9-13 2003, pp. 770–777.

[14] 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 *ACM/IEEE International Symposium on Low Power Electronics and Design*, Monterey, CA, August 12-14 2002, pp. 19–23.

[15] M. Kishor and J. P. de Gyvez, "Threshold voltage and power-supply tolerance of CMOS logic design families," in *IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, Yamanashi, Japan, October 25-27 2000, pp. 329–357.

[16] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De, "Parameter variations and impact on circuits and microarchitecture," in *ACM/IEEE Design Automation Conference*, Anaheim, CA, June 2-6 2003, pp. 338–342.

[17] T. Karnik, S. Borkar, and V. De, "Sub-90nm technologies - challenges and opportunities for CAD," in *IEEE/ACM International Conference on Computer-Aided Design*, San Jose, CA, November 10-14 2002, pp. 203–206.

[18] D. Boning and S. Nassif, *Models of Process Variations in Device and Interconnect*, A. Chandrakasan, W. J. Bowhill, and F. Fox, Eds. New York, NY: IEEE Press, 2001, ch. 6.

[19] 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 *IEEE International Solid State Circuits Conference*, San Francisco, CA, February 4-8 2001, pp. 278–279.

[20] M. Eisele, J. Berthold, D. Schmitt-Landsiedel, and R. Mahnkopf, "The impact of intra-die device parameter variations on path delays and on the design for yield of low voltage digital circuits," *IEEE Trans. VLSI Syst.*, vol. 5, no. 4, pp. 360–368, Dec. 1997.

[21] A. Srivastava, R. Bai, D. Blaauw, and D. Sylvester, "Modeling and analysis of leakage power considering within-die process variations," in *ACM/IEEE International Symposium on Low Power Electronics and Design*, Monterey, CA, August 12-14 2002, pp. 64–67.

[22] H. Su, F. Liu, A. Devgan, E. Acar, and S. Nassif, "Full chip leakage estimation considering power supply and temperature," in *ACM/IEEE International Symposium on Low Power Electronics and Design*, Seoul, Korea, August 25-27 2003, pp. 78–83.

[23] E. Acar, A. Devgan, R. Rao, F. Liu, H. Su, S. Nassif, and J. Burns, "Leakage and leakage sensitivity computation for combinational circuits," in *ACM/IEEE International Symposium on Low Power Electronics and Design*, Seoul, Korea, August 25-27 2003, pp. 96–99.

[24] J. N. Kozhaya, S. R. Nassif, and F. N. Najm, "A multigrid-like technique for power grid analysis," *IEEE Trans. Computer-Aided Design*, vol. 21, no. 10, pp. 1148–1160, Oct. 2002.

[25] L. T. Pillage, R. A. Rohrer, and C. Visweswariah, *Electronic Circuit and System Simulation Methods*. New York, NY: McGraw-Hill, 1995.
[26] S. R. Nassif, "Design for variability in {DSM} technologies," in *IEEE International Symposium on Quality Electronic Design*, San Jose, CA, March 20-22 2000, pp. 451–454.
[27] S. G. Duvall, "Statistical circuit modeling and optimization," in *IEEE International Workshop on Statistical Metrology*, Honolulu, HI, June 11 2000, pp. 56–63.
[28] 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 Trans. Computer-Aided Design*, vol. 14, no. 8, pp. 998–1012, Aug. 1995.
[29] J. Rubenstein, P. Penfield, and M. A. Horowitz, "Signal delay in RC tree networks," *IEEE Trans. Computer-Aided Design*, vol. 2, no. 3, pp. 202–211, July 1983.
[30] S. R. Nassif, "Within-chip variability analysis," in *IEEE International Electronic Devices Meeting*, San Francisco, CA, December 6-9 1998, pp. 283–286.
[31] V. Mehrotra, S. L. Sam, D. Boning, A. Chandrakasan, R. Vallishayee, and S. Nassif, "A methodology for modeling the effects of systematic within-die interconnect and device variations on circuit performance," in *ACM/IEEE Design Automation Conference*, Los Angeles, CA, June 5-9 2000, pp. 172–175.
[32] N. C. Beaulieu, A. A. Abu-Dayya, and P. J. McLane, "Estimating the distribution of a sum of independent lognormal random variables," *IEEE Trans. Commun.*, vol. 43, no. 12, pp. 2869–2873, Dec. 1995.
[33] R. E. Walpole, R. H. Myers, and S. L. Myers, *Probability and Statistics for Engineers and Scientists*, 6th ed. Upper Saddle River, NJ: Prentice Hall International, Inc., 1998.
[34] A. Papoulis and S. U. Pillai, *Probability, Random Variables, and Stochastic Processes*, 4th ed. Boston, MA: McGraw-Hill, 2002.
[35] I. A. Ferzli and F. N. Najm, "Statistical estimation of circuit timing vulnerability due to leakage-induced power grid voltage drop," in *IEEE International Conference on Integrated Circuit Design and Technology*, Austin, TX, May 17-20 2004, pp. 17–24.
[36] S. K. Thompson, *Sampling*, 2nd ed. New York, NY: John Wiley & Sons, Inc., 2002.
[37] T. A. Davis, *UMFPACK Version 4.3 User Guide*, University of Florida, Gainesville, FL, Jan. 2004. [Online]. Available: http://www.cise.ufl.edu/research/sparse/umfpack.
[38] *The Sun Performance Library User's Guide*, Sun Microsystems, Santa Clara, CA, 2003. [Online]. Available: http://docs.sun.com/source/817-0935/index.html.

Imad A. Ferzli received the B.E. degree in Computer and Communications Engineering from the American University of Beirut in 2001, and the M.A.Sc. degree in Electrical and Computer Engineering from the University of Toronto in 2004, where he is currently working toward the Ph.D. degree.

Mr. Ferzli is a Student Member of the IEEE and the Semiconductor Research Corporation (SRC). His research interests include CAD development for design for manufacturability and process variations, power grid verification and optimization, timing analysis, and yield modeling and prediction.

Farid N. Najm received the B.E. degree in Electrical Engineering from the American University of Beirut (AUB) in 1983, and the M.S. and Ph.D. degrees in Electrical and Computer Engineering (ECE) from the University of Illinois at Urbana-Champaign (UIUC) in 1986 and 1989, respectively. He worked as Electronics Engineer with AUB from 1983 to 1984. In 1989, he joined Texas Instruments in Dallas, TX, as Member of Technical Staff with the Semiconductor Process and Design Center. In 1992, he joined the ECE Department at UIUC as an Assistant Professor, and then became a Tenured Associate Professor at UIUC in 1997. In 1999, he joined the ECE Department at the University of Toronto, where he is currently Professor and Vice-Chair of ECE.

Dr. Najm is a Fellow of the IEEE, and is Associate Editor for the IEEE Transactions on CAD. He received the IEEE Transactions on CAD Best Paper Award in 1992, the NSF Research Initiation Award in 1993, the NSF CAREER Award in 1996, and was Associate Editor for the IEEE Transactions on VLSI 1997–2002. He served as General Chairman for the 1999 International Symposium on Low-Power Electronics and Design (ISLPED-99), and as Technical Program Co-Chairman for ISLPED-98. He has also served on the technical committees of ICCAD, DAC, CICC, ISQED, and ISLPED. Dr. Najm has co-authored the text "Failure Mechanisms in Semiconductor Devices," 2nd Ed., John Wiley & Sons, 1997. His research interests are in the general area of CAD tool development for low-power and reliable VLSI circuits, including power estimation and modeling, low-power design, power grid analysis and verification, and reliability analysis and prediction.