Statistical Verification of Power Grids Considering Process-Induced Leakage Current Variations*

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

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

ABSTRACT

Transistor threshold voltages (V_{th}) have been reduced as part of on-going technology scaling. The smaller V_{th} values feature increased variations due to underlying process variations, with a strong within-die component. Correspondingly, given the exponential dependence of leakage on V_{th}, circuit leakage currents are increasing significantly and have strong within-die statistical variations. With these leakage currents loading the power grid, the grid develops correspondingly large statistical voltage drops. This leakage-induced voltage drop is an unavoidable background level of noise on the grid. Any additional non-leakage currents due to circuit activity will lead to voltage drop which is to be added to this background noise. We propose a technique for checking whether the statistical voltage drop on every node is within user-specified bounds, given user-specified statistics of the leakage currents.

1. INTRODUCTION

Technology scaling requires reduction of the MOSFET threshold voltage (V_{th}), to accompany the reduction in supply voltages (V_{dd}) and oxide thickness. Over the years, V_{th} has been reduced from about 1 V in 5 V technology ($V_{dd} = 5$ V) to about 0.3 V in today's 1.2 V, 0.13 μm, technology, and is forecast to be further reduced at rate of 15% per generation in future [5]. Due to the exponential dependence of transistor sub-threshold leakage current on V_{th}, this leads to a much higher rate of increase of I_{off}, reportedly as high as about 5X per generation [5]. As a result, it is forecast [7] that, in the 0.1 μm generation, total chip leakage current (for dense, high-performance chips) would be about half of the total chip current.

Another consequence of technology scaling is that the reduced V_{th} values exhibit increased statistical variations due to underlying process variations [17, 12]. In 0.1 μm technology, it is possible to get 30 mV standard deviation in V_{th} [3, 9]. Considering a supply voltage of, say, around 1 Volt, this means that a $\pm 3\sigma$ interval for V_{th} would span 18% of the supply! Due to the exponential dependence of leakage current on threshold voltage, these variations lead to much larger variations in leakage current. For individual transistors, and individual logic gates, close to 3X variations have been observed in leakage [17]. These variations are also known to have a significant within-die component [4, 11, 2, 20, 12], so that transistors in close proximity on the layout can have significant variations in their leakage currents.

For whole chips, leakage variations have been *measured* at almost 20X [8]. If half the total current will, in future, be due to leakage, and if that total leakage current is going to vary by 20X, then it is clear that this will cause strong statistical variations in the total chip current. This has many consequences for chip design, especially for design of the power grid. In this work, we consider the fact that, in response to these statistical leakage currents, the grid will develop voltage drop at all the nodes that is correspondingly significant, and statistically variable with a strong within-die component. This voltage drop is unavoidable and manifests itself as a *background level of noise* on the grid which will have an effect on circuit delay and operation. This must be considered during circuit design. Any further (non-leakage) current variations that may be due to circuit dynamic operation will lead to voltage drop on the grid which is *on top* of this background noise level. Thus, there is a need to *verify the grid* by quantifying this background noise and making sure it does not exceed certain user-specified values.

Being statistical, there are no good tools today to estimate this voltage drop. Simply setting all leakage currents to their maximum possible values is overly pessimistic because local variations in V_{th} will make this case extremely improbable. Instead, one must consider all the leakage currents to be random and must proceed to analyze the grid on that basis. In this paper, we focus on within-die variations, which have been receiving a great deal of attention due to the impact they are forecast to have in future technologies. For within-die variations, a technique was proposed in [6] for estimating the mean and variance of the (random) node voltage drop at every grid node, using a statistical estimation method. While that technique is efficient, it is not fast enough as it stands to handle the multi-million node power grids that are common today.

In this paper, we extend the approach presented in [6] to propose a way to verify the grid which can handle much larger grid sizes. This is achieved by "changing the problem" being solved. Rather than doing a variance analysis on the node voltages, we will seek to characterize the distribution of the voltage drops, in response to leakage current variations, as users may rather be interested in verifying that each node voltage drop is within bounds. This is partly motivated by the fact that, pure variance analysis needs to be placed in a general framework where variances are put to practical use in order to draw informative conclusions. Therefore, one way to verify the grid is to check that the voltage drops at all the nodes are less than certain user-supplied thresholds (for node i, we will denote this threshold by V_{Ti}). The voltage drop being statistical, we actually will check whether the "bulk of the distribution" of the voltage drop at every node is below threshold. The user can specify what exactly is meant by "the bulk of the distribution," by specifying a certain percentage of the distribution which must be below the voltage threshold. We denote this percentage by $(1 - \beta_i)$. Thus, if one specifies $1 - \beta_i = 90\%$, we will check whether 90% of the distribution of the voltage drop at node i is less than V_{Ti}. Such a node is declared *safe*, otherwise *unsafe*. We will present an efficient statistical technique for checking which nodes are safe and which are not.

Except for [17] and [12], there is not much prior work on the study of statistical leakage currents due to within-die process variations. Other than [6], there is no work that the authors are aware of on the study of the statistical grid voltages in response to the statistical leakage currents.

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

2. PROBLEM STATEMENT

The background noise on the nodes of the power grid appears as random voltage drops on the grid nodes. We will say that a node is safe with regard to this noise if it is (highly) likely that the voltage drop on this node is less than some threshold voltage, and that a node is unsafe otherwise. This leads to a statistical definition of safety, that we state as:

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

In the above, $P\{\}$ denotes a probability, V_i the voltage drop at node i, V_{Ti} the safety threshold voltage at node i, and β_i a small positive number between 0 and 1. The quantity $1-\beta$ specifies the bulk of the voltage drop that is to be below a certain threshold to assert that a node is safe, and we would refer to it as the *safety parameter*. The choice of the safety parameter as well as 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 that must be accounted for during the design/layout stages. 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 $P\{V_i < V_{1-\beta_i}\} = 1-\beta_i$. Clearly, *Definition 1* can be restated in terms of $V_{1-\beta_i}$:

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 at the nodes of the power grid 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. Characterizing the voltage drop by a percentile accounts for the randomness in the leakage currents.

The verification problem reduces to determining whether the $1-\beta$ percentiles of voltage drops are greater or less than a given safety threshold. The difficulty lies in the fact that both the type and the parameters of the distribution of voltage drops are unknown.

Our methodology is as follows. We will first argue that the power grid node voltage drops due to within-die variations of leakage currents are lognormally distributed (section 3.2). Thus, we will be able to derive an expression for $V_{1-\beta_i}$ in terms of the parameters of the lognormal distribution, namely, mean and variance. We will show that *computation* of the mean is trivial (section 3.3) which allows us to state the verification problem in terms of the variances (section 3.4), for which we will derive an analytical expression (section 4). Then, we will introduce upper and lower bounds on the variances that allow us to to have 100% confidence whether $V_{1-\beta_i}$ is greater or less than V_{Ti} for some subset of the nodes (section 5). For the remaining nodes, we will *estimate* the variances so as to achieve a high confidence level on whether the percentile voltage drop is greater or less than the safety threshold voltage (section 6). This step is a tradeoff between accuracy and speed, in the sense that when it is computationally prohibitive to ascertain (with confidence 100%) the location of $V_{1-\beta_i}$ relative to V_{Ti}, we seek to know with a high (but less than 100%) confidence what this relative position is. Finally, in the case where some residual nodes have their $V_{1-\beta_i}$ so close to V_{Ti} that obtaining a high confidence on the relative position of the two gets too expensive, we report an estimate of $V_{1-\beta_i}$, the figure of merit of safety/unsafety, to within a user-defined resolution (section 6.3). The following sections elaborate on the different steps involved in implementing this methodology.

3. STATISTICS OF 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.

3.1 System Equations

We consider an RC model of the power grid, where each branch of the grid is represented by a resistor and where there exists a capacitor from every grid node to ground. In addition, some nodes have ideal current sources (to ground) representing the current drawn by the circuit tied to the grid at that point, and some nodes have ideal voltage sources (to ground) representing the connections to the external voltage supply. Let the power grid consist of $N+p$ nodes, where nodes $1,2,\ldots,N$ have no voltage sources attached, and nodes $(N+1),(N+2),\ldots,(N+p)$ are the nodes with the voltage sources. Let c_k be the capacitance from every node k to ground. Let $i_k(t)$ be the current source connected to node k, where the direction of positive current is from the node to ground. We assume that $i_k(t) \geq 0$ and that $i_k(t)$ is defined for every node $k=1,\ldots,N$ so that nodes with no current source attached have $i_k(t)=0, \forall t$. Let $\mathbf{i}(t)$ be the vector of all $i_k(t)$ sources, $u_k(t)$ be the voltage at node k, and $\mathbf{u}(t)$ be the vector of all $u_k(t)$ signals. Applying Modified Nodal Analysis (MNA) [16] leads to:

$$\mathbf{G}\mathbf{u}(t) + \mathbf{C}\dot{\mathbf{u}}(t) = -\mathbf{i}(t) + \mathbf{G}\mathbf{V_{dd}} \tag{1}$$

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

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

This is a *revised* system equation which one can solve directly for the voltage drop values. Notice that the circuit described by (2) consists of the original power grid, but with all the voltage sources set to zero and all the current source directions reversed. In the following, we will mainly be concerned with this *modified power grid* and the revised system of equations (2). In cases when the circuit is in a standby state, where all the currents are constant, the circuit response is obtained using a DC analysis. The DC equivalent of (2) is readily seen as:

$$\mathbf{G}\mathbf{V} = \mathbf{I} \tag{3}$$

3.2 Probability Distribution of Node Voltages

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 1. In response to within-die variations on the 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 (section 3.3 below).

With this assumption, the vector of node voltage drops $\mathbf{V}$ can be written as a function of the vector of leakage currents $\mathbf{I}$:

$$\mathbf{V} = \mathbf{G}^{-1}\mathbf{I} \tag{4}$$

From (4), it is clear that the voltage drop at an arbitrary node i can be expressed as:

$$V_i = q_{i1}I_1 + \cdots + q_{iN}I_N, \tag{5}$$

where q_{ij} is the (i,j)th entry of $\mathbf{G}^{-1}$. As (5) shows, the voltage drop at any node is a linear combination of the leakage currents loading the grid.

Process variations have both *systematic* and *random* components [4, 18], and as is typically done with other parameters, we will break down variations on the leakage currents as occurring both on the wafer level and on the die level, plus a residual component [18]:

$$\mathbf{I} = \mathbf{I_{dd}} + \mathbf{I_{wd}} + \mathbf{r}, \tag{6}$$

where $\mathbf{I_{dd}}$ represents the die-to-die (or inter-die) variations, $\mathbf{I_{wd}}$ the within-die (or intra-die) variations, and $\mathbf{r}$ the residual component. Note that (6) confines the randomness to the residual component, extracting the inter-die and the intra-die variations in a systematic, deterministic way, according to a methodology explained in [18].

Inter-die variations are modeled as a shift in the mean of a given parameter and can usually be dealt with using conventional statistical analysis techniques, such as Monte-Carlo simulations, or worst-case analysis [4]. Within-die variations, on the other hand, arise due to a variety of *independent* factors related to the environment, process, and layout [14]. These factors may not be fully comprehended, especially in the early or pre-layout stages of the design. Hence, it is argued in [13] that characterization of these variations as random is necessary, so as to include in the residual term of (6) all variations that cannot be accounted for on a systematic basis. Within-die variations cause device and interconnect mismatches on the chip, described in [4] as "unintentional". They are known [18] to have high spatial frequency trends across the surface of the chip, reflecting strong local variations. In this work, we will neglect spatial correlations between within-die leakage current variations and lump these variations into the residual "noise" term in (6) to lay the following assumption on the statistical properties of intra-die leakage currents, which will subsequently prove very useful in simplifying the problem:

Assumption 2. All intra-die leakage currents may be modeled as statistically independent random variables.

We now consider intra-die variations in leakage currents, due to variations in the transistor threshold voltages. The latter variations are modeled as Gaussian [12] leading to *lognormal* [10] variations on the leakage currents [17], by virtue of the exponential relation between the transistor leakage current and its threshold voltage. Trivially, if the I_j are independent and lognormal, then so are the $q_{ij}I_j$. Hence, the (random) node voltages (5) are summations of independent lognormal RVs. Sums of independent lognormal variables have been extensively studied and characterized in the literature pertaining to communications engineering, and it was found that such sums can be accurately captured by another lognormal RV [1], hence the basis for characterizing the distribution of the node voltage drops as lognormal. We shall provide empirical data to corroborate this argument in section 7.

Since V_i is modeled as a lognormal RV, the distribution of V_i has two parameters, mean and variance. By virtue of lognormality of V_i, $\ln(Vi)$ is a Gaussian RV, and the cumulative density function (cdf) of V_i is given by [10]:

$$F_{V_i}(V) = P\{V_i < V\} = \Phi\left(\frac{\ln(V) - \mu_{\ln(V_i)}}{\sigma_{\ln(V_i)}}\right), \tag{7}$$

where $\Phi(.)$ is the cdf of the Gaussian distribution with 0 mean and unit variance, and $\mu_{ln(V_i)}$ and $\sigma_{\ln(V_i)}$ are respectively the expected value and the standard deviation of $\ln(V_i)$. It can be shown that the parameters of $\ln(V_i)$ can be expressed in terms of the mean and variance of V_i, as follows:

$$\mu_{\ln(V_i)} = \ln(\mu_{Vi}) - \frac{1}{2}\ln\left(1 + \frac{\sigma_{V_i}^2}{\mu_{V_i}^2}\right), \tag{8}$$

and

$$\sigma_{\ln(V_i)} = \sqrt{\ln\left(1 + \frac{\sigma_{V_i}^2}{\mu_{V_i}^2}\right)}, \tag{9}$$

where μ_{V_i} and σ_{V_i} are the expected value and standard deviation of the voltage drop at node i, respectively.

We can now derive an expression for the $1-\beta_i$ percentile of the voltage drop. Notice that $V_{1-\beta_i}$ is such that $F_{V_i}(V_{1-\beta_i}) = 1-\beta_i$, so we can write:

$$\Phi\left(\frac{\ln(V_{1-\beta_i}) - \mu_{\ln(V_i)}}{\sigma_{\ln(V_i)}}\right) = 1 - \beta_i, \tag{10}$$

so that:

$$\frac{\ln(V_{1-\beta_i}) - \mu_{\ln(V_i)}}{\sigma_{\ln(V_i)}} = z_{1-\beta_i}, \tag{11}$$

where $z_{1-\beta_i}$ is such that $\Phi(z_{1-\beta_i}) = 1 - \beta_i$, and can be easily calculated given β_i. Therefore, we have:

$$\begin{aligned} V_{1-\beta_i} &= e^{z_{1-\beta_i}\sigma_{\ln(V_i)} + \mu_{\ln(Vi)}} \\ &= \mu_{V_i}\frac{e^{z_{1-\beta_i}\sqrt{\ln\left(1+\frac{\sigma_{V_i}^2}{\mu_{V_i}^2}\right)}}}{\sqrt{1 + \frac{\sigma_{V_i}^2}{\mu_{V_i}^2}}} \end{aligned} \tag{12}$$

Equation (12) provides an expression for the $1-\beta_i$ percentile of the voltage drop at any node, in terms of the mean and variance of the voltage drop at that node.

3.3 Mean Estimation

Since the system (2) is linear, then due to linearity of the mean ($E[\cdot]$) operator [15], one can write:

$$\mathbf{G}E[\mathbf{v}(t)] + \mathbf{C}\frac{d}{dt}E[\mathbf{v}(t)] = E[\mathbf{i}(t)] \tag{13}$$

Thus, if we solve the system (2) once, using simply the current means as inputs, the solution gives the voltage means at all the nodes, obtained from the dynamic model of the grid. This leaves the variances as the only unknowns in determining the percentile voltage drops as per (12).

3.4 Verification Equations

From *Definition* 2, safety/unsafety of node i is determined by the relative position of $V_{1-\beta_i}$ and V_{Ti}. Note that (12) can be viewed as expressing $V_{1-\beta_i}$ as a function of $\sigma_{V_i}^2$, *i.e.*, $V_{1-\beta_i} = f(\sigma_{V_i}^2)$. One can easily show the following:

1. $f(0) = \mu_{V_i}$
2. $f(.)$ admits one local maximum for $\sigma_{V_i}^2 = \mu_{V_i}^2(e^{z_{1-\beta_i}^2} - 1)$ and the maximum value of $f(\sigma_{V_i}^2)$ is $\mu_{V_i}\sqrt{e^{z_{1-\beta_i}^2}}$
3. $\lim_{\sigma_{V_i}^2 \to \infty} f(\sigma_{V_i}^2) = 0$.

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

Group 1. Includes all nodes i such that:

$$V_{Ti} > \mu_{V_i}\sqrt{e^{z_{1-\beta}^2}}.$$

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.

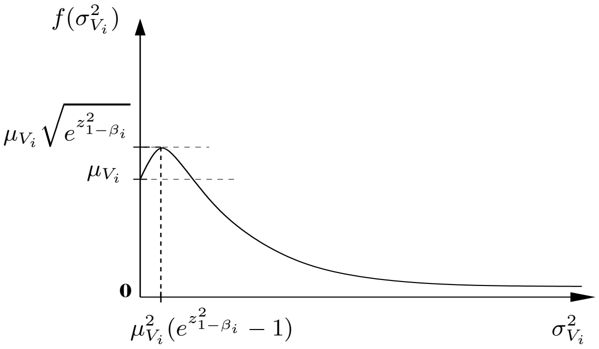

Figure 1: A typical plot of $f(\sigma^2_{V_i})$.

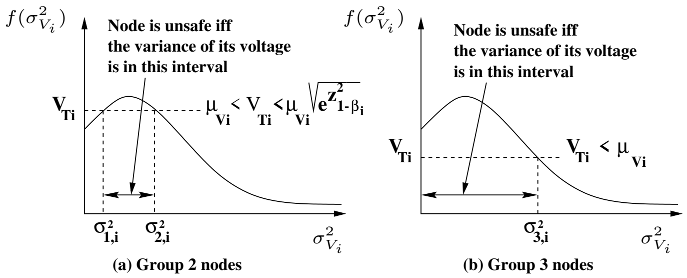

Figure 2: Illustration of safety/unsafety conditions on the variances for Group 2 and Group 3 nodes.

Group 2. Includes all nodes i such that:

$$\mu_{V_i} < V_{Ti} < \mu_{V_i}\sqrt{e^{z^2_{1-\beta_i}}}.$$

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

$$\sigma^2_{1,i} = \mu^2_{V_i}\left(e^{2\left(z^2_{1-\beta_i} - \ln\frac{V_{Ti}}{\mu_{V_i}} - z_{1-\beta_i}\sqrt{z^2_{1-\beta_i} - 2\ln\left(\frac{V_{Ti}}{\mu_{V_i}}\right)}\right)} - 1\right) \tag{14}$$

and

$$\sigma^2_{2,i} = \mu^2_{V_i}\left(e^{2\left(z^2_{1-\beta_i} - \ln\frac{V_{Ti}}{\mu_{V_i}} + z_{1-\beta_i}\sqrt{z^2_{1-\beta_i} - 2\ln\left(\frac{V_{Ti}}{\mu_{V_i}}\right)}\right)} - 1\right). \tag{15}$$

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

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

$$\sigma^2_{3,i} = \mu^2_{V_i}\left(e^{2\left(z^2_{1-\beta_i} - \ln\frac{V_{Ti}}{\mu_{V_i}} + z_{1-\beta_i}\sqrt{z^2_{1-\beta_i} - 2\ln\left(\frac{V_{Ti}}{\mu_{V_i}}\right)}\right)} - 1\right). \tag{16}$$

(see Fig. 2b).

With this, determining whether a node is safe or unsafe reduces to knowing where the variance of the voltage at that node is located with respect to σ^2_1 and σ^2_2 for *Group* 2 nodes and with respect to σ^2_3 for *Group* 3 nodes; *Group* 1 nodes are deemed safe, irrespective of the variance of their voltages. Observe that the group of each node is known automatically, and requires no a-priori knowledge of the variance of the node voltage drop, since $z_{1-\beta_i}$ is directly obtained knowing β_i, V_{Ti} is a user-defined parameter, and the means of the voltage drops are easily calculated using (13).

4. VARIANCE COMPUTATION

In this section, we derive an analytical expression for the variance of the voltage at every node. Under the pseudo-static assumption, the system (13) is simplified to its DC version:

$$\mathbf{G}E[\mathbf{V}] = E[\mathbf{I}] \tag{17}$$

We now combine (3) and (17) to yield:

$$\mathbf{G}(\mathbf{V} - E[\mathbf{V}]) = \mathbf{I} - E[\mathbf{I}] \tag{18}$$

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

$$\mathbf{G}E\left[(\mathbf{V} - E[\mathbf{V}])(\mathbf{V} - E[\mathbf{V}])^T\right]\mathbf{G}^T = E\left[(\mathbf{I} - E[\mathbf{I}])(\mathbf{I} - E[\mathbf{I}])^T\right] \tag{19}$$

We recognize the expectations as being simply covariance matrices [15], so that the above result can be rewritten as:

$$\mathbf{G}\mathrm{Cov}(\mathbf{V})\mathbf{G}^T = \mathrm{Cov}(\mathbf{I}) \tag{20}$$

Since $\mathbf{G}$ is symmetric, so is $\mathbf{G}^{-1}$. Therefore, (20) becomes:

$$\mathrm{Cov}(\mathbf{V}) = \mathbf{G}^{-1}\mathrm{Cov}(\mathbf{I})\mathbf{G}^{-1} \tag{21}$$

The assumption of statistical independence of leakage currents implies that $\mathrm{Cov}(\mathbf{I})$ is diagonal, and since $\mathbf{G}^{-1}$ is symmetric, it can be seen from (21) that:

$$\begin{aligned}[\mathrm{Cov}(\mathbf{V})]_{ii} &= \left([\mathbf{G}^{-1}]_{i1}\right)^2[\mathrm{Cov}(\mathbf{I})]_{11} \\ &\quad + \cdots + \left([\mathbf{G}^{-1}]_{iN}\right)^2[\mathrm{Cov}(\mathbf{I})]_{NN}, \end{aligned} \tag{22}$$

where $[.]_{ij}$ is the (i,j)th entry of the corresponding matrix. Let Σ_V and Σ_I represent the *vectors* of standard deviations of voltages drops and currents, respectively, and Σ^2_V and Σ^2_I the corresponding variance vectors, and define a matrix $\mathbf{G}^{-1(2)}$ such that:

$$[\mathbf{G}^{-1(2)}]_{ij} = \left([\mathbf{G}^{-1}]_{ij}\right)^2, \tag{23}$$

so that (22) can be written as:

$$\boxed{\Sigma^2_V = \mathbf{G}^{-1(2)}\Sigma^2_I} \tag{24}$$

The solution of (24) requires full knowledge of the inverse of $\mathbf{G}$, which is impractical for large grids, hence the difficulty in evaluating the variances of node voltages.

5. DIRECT CRITERIA

In this section, we will make use of bounds that may directly determine 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.

5.1 Bounds on the Variances

As in section 3.2, let q_{ij} denote the (i,j)th element of $\mathbf{G}^{-1}$, σ_{I_j} the standard deviation of the current source at node j and σ_{V_j} the standard deviation of the voltage drop at node j. Since $\sigma_{I_j} \geq 0$ and $q_{ij} \geq 0$, $\forall i, j$, then from (24):

$$\sigma^2_{V_i} = \sum_j q^2_{ij}\sigma^2_{I_j} = \sum_j (q_{ij}\sigma_{I_j})^2 \leq \left(\sum_j q_{ij}\sigma_{I_j}\right)^2. \tag{25}$$

This leads to an upper bound on Σ_V, as follows:

$$\Sigma_V \leq \mathbf{G}^{-1}\Sigma_I. \tag{26}$$

Given an LU-factorization of $\mathbf{G}$, the cost of computing this upper bound is only one forward/backward solve.

Now let $S = \sum_j \sigma_{I_j}^2$ and let $p_j = \sigma_{I_j}^2/S$. Since $0 \leq p_j \leq 1$, then $p_j < p_j^2$, $\forall j$. Then we can write:

$$\sigma_{V_i}^2 = S\sum_j p_j q_{ij}^2 \geq S\left(\sum_j p_j q_{ij}^2\right) = \frac{1}{S}\left(\sum_j q_{ij}\sigma_{I_j}^2\right)^2. \tag{27}$$

This leads to a lower bound on Σ_V, as follows:

$$\Sigma_V \geq \frac{1}{\sqrt{S}}\mathbf{G}^{-1}\Sigma_I^2. \tag{28}$$

Also, given an LU-factorization of $\mathbf{G}$, the cost of computing this lower bound is only one forward/backward solve.

Putting (26) and (28) together yields the following interval on Σ_V:

$$\frac{1}{\sqrt{S}}\mathbf{G}^{-1}\Sigma_I^2 \leq \Sigma_V \leq \mathbf{G}^{-1}\Sigma_I. \tag{29}$$

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 $\sigma_{1,i}^2$ and $\sigma_{2,i}^2$, then this node can be deemed safe if ($u_i^2 < \sigma_{1,i}^2$ or $l_i^2 > \sigma_{2,i}^2$) and unsafe if ($l_i^2 > \sigma_{1,i}^2$ and $u_i^2 < \sigma_{2,i}^2$). Similarly, if node i is in *Group* 3 with a corresponding $\sigma_{3,i}^2$, then it can be deemed safe if $l_i^2 > \sigma_{3,i}^2$ and unsafe if $u_i^2 < \sigma_{3,i}^2$.

6. ITERATIVE CRITERIA

If none of the checks proposed in section 5 succeeds in establishing whether some nodes are safe or unsafe, and since solving the system (24) to get the variances is computationally expensive, we will trade off some accuracy in the knowledge of the voltage variances in order to arrive to an efficient way of verifying these remaining nodes. Short of knowing with certainty whether $V_{1-\beta_i}$ is less than or greater than V_{Ti}, we will seek to know this information with high confidence. Specifically, we will extend *Definition* 2 to consider that a node is safe if $P\{V_{1-\beta_i} < V_{Ti}\} \geq 1-\alpha$ and unsafe if $P\{V_{1-\beta_i} > V_{Ti}\} \geq 1-\alpha$, where α is a small number between 0 and 1. In this perspective, for the nodes successfully checked in section 5, $\alpha = 0$.

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 ith node voltage and some known constants ($\sigma_{1,i}^2$ and $\sigma_{2,i}^2$ if i is in *Group* 2 and $\sigma_{3,i}^2$ if i is in *Group* 3).

6.1 Variance Estimation

In the following, we make use of a technique introduced in [6] to estimate the variances of node voltages. In order to simplify the notation, let r_{ij} denote the (i,j)th entry of the matrix $\mathbf{G}^{-1(2)}$ (*i.e.*, $r_{ij} = q_{ij}^2$, where q_{ij} is defined in section 3.2), and define $\sigma_{I_i}^2$, $\sigma_{V_i}^2$, S, and p_i as in section 5.1 . As in (27), we have:

$$\sigma_{V_i}^2 = S\sum_{j=1}^{N} p_j r_{ij} \tag{30}$$

Since $\sum_{i=1}^{N} p_j = 1$, then we can view the p_j weights as being *probability values* associated with the r_{ij} values, so that the summation above becomes the mean (weighted average) of all the r_{ij} elements in the ith row. If we define an RV $\mathbf{r_i}$ as being a discrete RV that takes the values r_{ij} with probabilities p_j, $j = 1, 2, \ldots, N$, then we can write (30) as:

$$\sigma_{V_i}^2 = SE[\mathbf{r_i}] \tag{31}$$

Let the mean of $\mathbf{r_i}$ be $\mu_i = E[\mathbf{r_i}]$ and its variance be σ_i^2. We can now use methods of *mean estimation* from statistics, basically Monte Carlo random sampling [10, 19], in order to estimate the *population mean* μ_i using the mean of a much smaller sample (say, of size $n \ll N$) from the population, *i.e.*, using the *sample mean*.

The process is simple. Using a weighted random number generator, we generate according to the probabilities p_j a sequence of indices of columns of $\mathbf{G}^{-1(2)}$ to be included in the sample. From these, we form the following sample mean for every row i:

$$\bar{r}_i = \frac{1}{n}\sum_{j\in\mathcal{J}} r_{ij} \tag{32}$$

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

$$s_i^2 = \frac{1}{n-1}\sum_{j\in\mathcal{J}}(r_{ij}-\bar{r}_i)^2 = \frac{n\left(\sum_{j\in\mathcal{J}} r_{ij}^2\right) - \left(\sum_{j\in\mathcal{J}} r_{ij}\right)^2}{n(n-1)} \tag{33}$$

Note that $\bar{r}_i$ itself can be considered as an RV, with mean μ_i and variance s_i^2/n (for large n). Since $\bar{r}_i$, the sample mean, is an *unbiased* estimator of $E[r_i]$ [10], then

$$\hat{\sigma}_{V_i}^2 = S\bar{r}_i \tag{34}$$

is an unbiased estimator of $\sigma_{V_i}^2$, with variance $S^2 s_i^2/n$. Furthermore, by the *central limit theorem* [10], $\bar{r}_i$ will be normally distributed, so that the RV:

$$\frac{\sigma_{V_i}^2 - \hat{\sigma}_{V_i}^2}{\frac{Ss_i}{\sqrt{n}}}$$

is normal with 0 mean and unit variance [19], for large n.

6.2 Error Bounds

As we are sampling columns of $\mathbf{G}^{-1(2)}$, $\hat{\sigma}_{V_i}^2$ may fall to the left of $\sigma_{1,i}^2$, to the right of $\sigma_{2,i}^2$, or between the two, if the node is in Group 2, or to the left or to the right of $\sigma_{3_i}^2$ if the node is in Group 3. We thus have intervals defined with respect to $\sigma_{1,i}^2$ and $\sigma_{2,i}^2$, or $\sigma_{3,i}^2$. 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 $1-\alpha$ large enough, less than $1-\alpha$. Based on this, in the sampling process, we shall seek to find the smallest number of samples n that verifies, with confidence $1-\alpha$, whether the node is safe, according to the interval where the variance estimate lies.

Assume node i is in *Group* 2, then as was shown in section 3.4, $V_{1-\beta_i} < V_{Ti}$ is equivalent to $\sigma_i^2 < \sigma_{1,i}^2$ or $\sigma_i^2 > \sigma_{2,i}^2$, and $V_{1-\beta_i} > V_{Ti}$ is equivalent to $\sigma_{1,i}^2 < \sigma_i^2 < \sigma_{2,i}^2$.

If $\hat{\sigma}_{V_i}^2 < \sigma_{1,i}^2$, $1-\alpha$ confidence on the safety is reached when:

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

The above can be written as:

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

leading to:

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

Knowing that $(\sigma^2_{V_i} - \hat{\sigma}^2_{V_i})/\sqrt{\mathrm{Var}(\hat{\sigma}^2_{V_i})}$ is standard normal, this condition reduces to:

$$\Phi\left(\frac{(\sigma^2_{2,i} - \hat{\sigma}^2_{V_i})\sqrt{n}}{Ss_i}\right) - \Phi\left(\frac{(\sigma^2_{1,i} - \hat{\sigma}^2_{V_i})\sqrt{n}}{Ss_i}\right) \leq \alpha. \quad (35)$$

That is, if $\hat{\sigma}^2_{V_i} < \sigma^2_{1,i}$, then a $1-\alpha$ confidence level that node i is safe is attained iff n satisfies (35). Note that (35) 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, $\hat{\sigma}^2_{V_i} < \sigma^2_{1,i}$, we can obtain a *sufficient* condition for n to satisfy in order to verify safety, by neglecting $P\{\sigma^2_{V_i} > \sigma^2_{2,i}\}$, yielding:

$$n \geq \left(\frac{Ss_n z_{1-\alpha}}{\sigma^2_{1,i} - \hat{\sigma}^2_{V_i}}\right)^2 = n_1, \quad (36)$$

where $z_{1-\alpha}$ is such that: $\Phi(z_{1-\alpha}) = 1-\alpha$. Thus, (36) provides a closed-form for n that checks the safety at node i.

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

$$\Phi\left(\frac{(\hat{\sigma}^2_{V_i} - \sigma^2_{1,i})\sqrt{n}}{Ss_i}\right) - \Phi\left(\frac{(\hat{\sigma}^2_{V_i} - \sigma^2_{2,i})\sqrt{n}}{Ss_i}\right) \leq \alpha, \quad (37)$$

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

$$n \geq \left(\frac{Ss_n z_{1-\alpha}}{\hat{\sigma}^2_{V_i} - \sigma^2_{2,i}}\right)^2 = n_2. \quad (38)$$

Finally, if $\sigma^2_{1,i} < \hat{\sigma}^2_{V_i} < \sigma^2_{2,i}$, 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:

$$\Phi\left(\frac{(\hat{\sigma}^2_{V_i} - \sigma^2_{1,i})\sqrt{n}}{Ss_i}\right) + \Phi\left(\frac{(\sigma^2_{2,i} - \hat{\sigma}^2_{V_i})\sqrt{n}}{Ss_i}\right) \leq 2-\alpha, \quad (39)$$

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 [19]:

$$\hat{\sigma}^2_{V_i} \pm \frac{Ss_n z_{1-\alpha/2}}{\sqrt{n}}, \quad (40)$$

where $z_{1-\alpha/2}$ is such that $\Phi(z_{1-\alpha/2}) = 1-\alpha/2$. So, to obtain a $1-\alpha$ confidence level that $\sigma^2_{V_i}$ is between $\sigma^2_{i,1}$ and $\sigma^2_{i,2}$, it is sufficient to have both extremes of the interval given in (40) lie within $[\sigma^2_{1,i}, \sigma^2_{2,i}]$. This leads to the following condition:

$$n \geq \left(\frac{Ss_n z_{1-\alpha/2}}{\min(\hat{\sigma}^2_{V_i} - \sigma^2_{1,i}, \sigma^2_{2,i} - \hat{\sigma}^2_{V_i})}\right)^2 = n_3. \quad (41)$$

Obtaining bounds for nodes in Group 3 is easier. We have two intervals where the variance estimate may fall: if $\hat{\sigma}^2_{V_i} > \sigma^2_{3,i}$, then we need to check for safety, i.e.:

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

This is equivalent to:

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

which reduces to:

$$n \geq \left(\frac{Ss_n z_{1-\alpha}}{\hat{\sigma}^2_{V_i} - \sigma^2_{3,i}}\right)^2 = n_4, \quad (42)$$

Note that (42) is a necessary and sufficient condition on n to achieve $1-\alpha$ confidence that node i is safe. Similarly, if $\hat{\sigma}^2_{V_i} < \sigma^2_{3,i}$, we obtain $1-\alpha$ confidence that node i is safe iff n satisfies:

$$n \geq \left(\frac{Ss_n z_{1-\alpha}}{\sigma^2_{3,i} - \hat{\sigma}^2_{V_i}}\right)^2 = n_5, \quad (43)$$

Note that $n_4 = n_5$, and that (42) and (43) are closed-form necessary and sufficient conditions on n, the required number of samples.

In summary, convergence of node i is achieved as follows:

```
If node i is in Group 2:
    If σ̂²_Vi < σ²_1,i:
        If n ≥ n1:  the node is done - safe.
    Else if σ̂²_Vi > σ²_2,i:
        If n ≥ n2:  the node is done - safe.
    Else if σ²_1,i < σ̂²_Vi < σ²_2,i:
        If n ≥ n3:  the node is done - unsafe.
Else if node i is in Group 3:
    If σ̂²_Vi > σ²_3,i:
        If n ≥ n4:  the node is done - safe.
    Else If σ̂²_Vi < σ²_3,i:
        If n ≥ n5:  the node is done - unsafe.
```

6.3 Residual Nodes

It can be seen from (36) – (43) 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.

In section 6.1, we used $\hat{\sigma}^2_{V_i}$ as an estimator of $\sigma^2_{V_i}$. Given (12), we will then use $\hat{V}_{1-\beta_i}$ as an estimator of the $1-\beta_i$ percentile of the voltage drop at node i, which can be written as:

$$\hat{V}_{1-\beta_i} = f(\hat{\sigma}^2_{V_i}), \quad (44)$$

where $f(.)$ is defined in section 3.4. Hence, if the variance estimator is close to $\sigma^2_{1,i}$, $\sigma^2_{2,i}$, or $\sigma^2_{3,i}$, then $V_{1-\beta_i}$ 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 $|\hat{V}_{1-\beta_i} - V_{Ti}| \leq \delta Vdd$. 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 (40) that with $1-\alpha$ confidence, the variance of each residual node lies in the interval $[v_1, v_2]$, where

$$v_1 = \hat{\sigma}^2_{V_i} - (Ss_n z_{1-\alpha/2})/\sqrt{n}$$

and

$$v_2 = \hat{\sigma}^2_{V_i} + (Ss_n z_{1-\alpha/2})/\sqrt{n}.$$

Knowing the variations of $f(.)$ (as per section 3.4), it is easy to determine the point v_m in $[v_1, v_2]$ where $f(.)$ is largest. Depending on the values of $\hat{\sigma}^2_{V_i}$, v_1, and v_2, v_m can be equal to v_1, v_2, or $\mu^2_{V_i}(e^{z^2_{1-\beta_i}} - 1)$. Clearly, then, $V_{ub,i}$ can be written as:

$$V_{ub,i} = f(v_m). \quad (45)$$

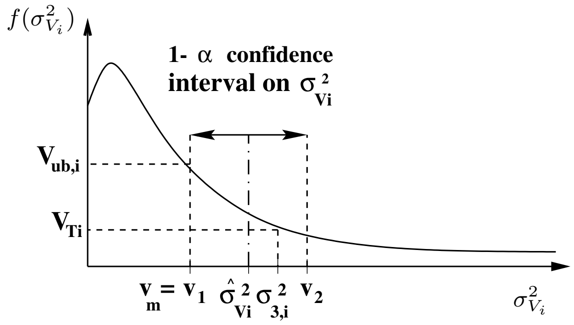

Figure 3: Finding an upper bound with $1-\alpha$ confidence on $V_{1-\beta_i}$ for a residual, Group 3 node.

Fig. 3 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.

7. EXPERIMENTAL RESULTS

This method has 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 number of grids to test our approach under different conditions, we have opted to *generate* a number of grids ourselves. The grid generation process is automatic, and employs a random number generator, as well as user-specified technology and topology parameters. Starting with a square uniform grid of a given size, we proceed to randomly delete a user-specified percentage of nodes, thus rendering the grid structurally non-uniform. Typical geometric and physical grid characteristics (e.g. grid dimensions) as well as 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 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 and leakage current sources are supplied by the user; the C4s and current sources are then distributed at random over the grid nodes.

Fig. 4 corroborates the fact that voltage drops are well modeled by a lognormal distribution. If indeed voltage drops are lognormally distributed, then their logarithms are normal. We generated several grids and collected voltage drop data, then verified graphically the goodness-of-fit of their logarithms on normal scores plots [10]. As can be seen from the figure, voltage drops showed good fits, except for certain outlying points, validating the choice of lognormal distributions to model the voltage drops on the power grid, induced by independent, within-die leakage current variations.

Tables 1 and 2 show the overall performance of the proposed approach on grids of various sizes. The experiments were run on a 1.5 GHz Sun Fire server with 4.0 GB of main memory, and we report CPU times. The grids in the first table are small enough so that we were able also to solve (24) 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. We observed in our experiments that some grids were fully verifiable by the direct criteria (section 5) in which case the run time is dramatically reduced. The first grid in Table 1 and the second grid in Table 2 are examples of such grids. Having the exact values of the $1-\beta$ percentiles of the voltage drops for the small grids, we define the percentage error 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. We report these

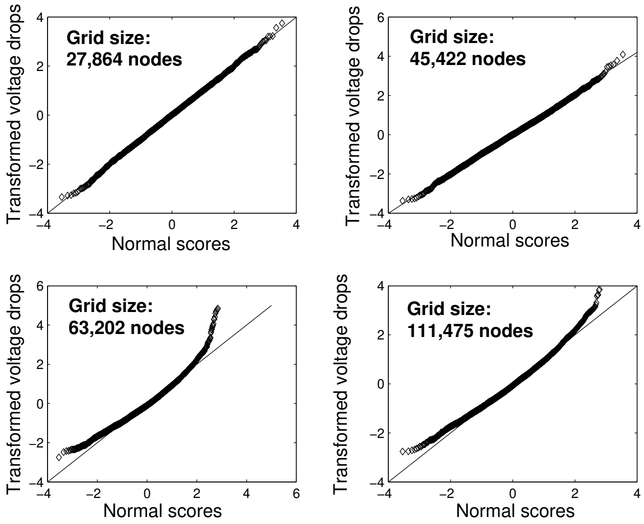

Figure 4: Checking graphically the goodness-of-fit of the voltage drop data against a lognormal distribution using the method of normal scores.

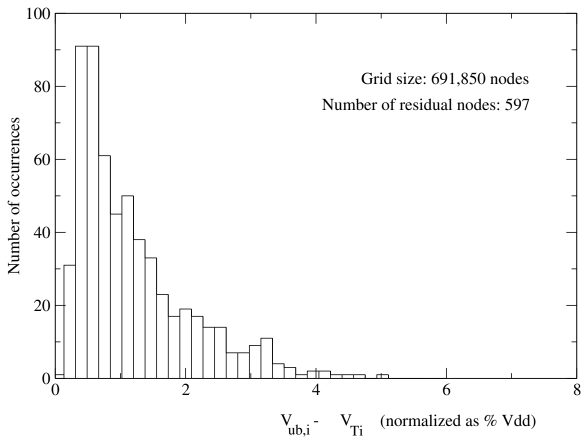

Figure 5: Distribution of upper bounds on the $1-\beta$ percentiles for residual nodes.

errors in the last column of Table 1. In every case, the error was very much in accordance with the specified bounds. The exact solution comes with only modest error penalty, and in the case of grids there were verifiable using only direct criteria, the error is 0%.

The histogram in Fig. 5 shows, for all residual nodes in a grid of roughly 700K nodes, the distribution of the distance between the upper bounds on the $V_{1-\beta_i}$ and the safety threshold voltages V_{Ti}. The safety parameter $(1-\beta)$ was set at (90%) for all nodes and the confidence level for convergence $(1-\alpha)$ at 90%. The resolution was fixed at 1% of V_{dd} ($\delta = 0.01$). The figure shows that there were only 597 residual nodes. The number of such nodes is primarily related to the user-specified resolution. The experiments conducted all 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 1 and 2). It was stated in section 6.3 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. 5 illustrates that this distance is small (the average is about 1% of V_{dd} in Fig. 5), implying that the $V_{1-\beta}$ and V_{Ti} are indeed close (that closeness being controlled by the resolution).

Our methodology assumes that an LU-factorization of the conductance matrix is available, and proceeds to apply verification checks, given this factorization. An important observation, that

Table 1: Results on small grids, where verifying 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	Time (LU)	Time (direct)	Time (iterative)	Time (exact)	% error
20,704	99%	95%	10%	100	0	32 sec.	0.00 sec	0	17 min	0
29,132	80%	95%	5%	85.4	2.0	23 sec.	0.26 sec.	36 sec.	37.5 min.	1.3
40,604	92%	95%	10%	70.6	2.7	3 min.	0.66 sec.	11 sec	2.4 hrs.	0.4
51,711	99%	99%	10%	99.94	0.06	2.3 min.	0.66 sec.	11 sec.	3.1 hrs.	0.2
72,085	90%	95%	10%	88.4	6.8	106 sec.	0.65 sec.	133 sec.	3.9 hr.	3.9
103,775	95%	92%	15%	91.5	4.8	5.6 min.	1.28 sec.	20 sec.	10.7 hrs.	0.95

Table 2: Performance of the proposed 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%	64	3.7	1.6 hr.	7 sec.	957	37 min.	552 MB
415,410	90%	90%	10%	100	0	2.1 hrs.	0.02 sec.	0	0	637 MB
691,850	90%	90%	10%	98.1	0.08	8.3 hrs.	20 sec.	1327	2.4 hrs.	1.3 GB
811,912	90%	90%	15%	91.5	0.4	8.65 hrs	22 sec.	1677	3.2 hrs.	1.4 GB
1,008,899	80%	95%	10%	99.8	0.10	30.3 hrs.	44.2 sec.	101	25.2 min	2.8 GB

appears more and more crucial for the larger grids, is that the LU becomes the bottleneck in terms of run time (besides being the bottleneck for memory usage). Despite the fact that the LU is both time and memory consuming, we were able to run grids of more than 800K nodes in about eleven hours in total, more than 8.5 of which are due to the LU, which is still a reasonable execution time. For the grid of 1M nodes, notice that the LU takes more than 30 hours, but the verification itself less than 30 minutes. As part of future work, we will investigate ways to get around this LU bottleneck.

8. CONCLUSION

Due to reduced threshold voltages (V_{th}), circuit leakage currents are much higher than before, and are projected to become even larger. Due to increasing V_{th} variations, which exhibit a strong within-die component, leakage currents are statistically variable, with a strong within-die component. The effect of these currents on the power grid is to generate a statistical background-noise voltage drop level on the grid. We have presented an efficient analytical methodology to verify every node in the grid in the presence of this noise, by checking whether the bulk of the distribution of the node voltage drop at any node falls below a user defined voltage level, with high confidence. We have derived bounds on the variances of the voltage drops, and direct and iterative criteria to estimate a given percentile of the voltage drop. We checked the accuracy of the proposed technique on small power grids. For large grids, our experiments showed that when applied on top of the LU factorization, this technique is impeded by the cost of this factorization. As such, future work may involve ways to alleviate the problem associated with the cost of the LU.

9. REFERENCES

[1] N. C. Beaulieu, A. A. Abu-Dayya, and P. J. McLane. Estimating the distribution of a sum of independent lognormal random variables. *IEEE Transactions on Communications*, 43(12):2869–2873, Dec. 1995.

[2] 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 *International Solid State Circuits Conference*, pages 278–279, 2001.

[3] D. Burnett, K. Eringtron, C. Subramanian, and K. Baker. Implications of fundamental threshold voltage variations for high-density SRAM and logic circuits. In *Symposium on VLSI Technology*, pages 15–16, 1994.

[4] A. Chandrakasan, W. J. Bowhill, and F. Fox, editors. *Design of High-Performance Microprocessor Circuits*, chapter 6, by D. Boning and S. Nassif. IEEE Press, 2001.

[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*, pages 163–168, San Diego, CA, August 16-17 1999.

[6] I. A. Ferzli and F. N. Najm. Statistical estimation of leakage-induced power grid voltage drop considering within-die process variations. In *ACM/IEEE 40th Design Automation Conference*, pages 856–859, Anaheim, CA, June 2-6 2003.

[7] T. Kam, S. Rawat, D. Kirkpatrick, R. Roy, G. S. Spirakis, N. Sherwani, and C. Peterson. EDA challenges facing future microprocessor design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 19(12):1498–1506, Dec. 2000.

[8] T. Karnik, S. Borkar, and V. De. Sub-90nm technologies - challenges and opportunities for CAD. In *International Conference on Computer-Aided Design*, pages 203–206, San Jose, CA, Nov. 10-14 2002.

[9] 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*, pages 329–357, Oct. 2000.

[10] I. R. Miller, J. E. Freund, and R. Johnson. *Probability and Statistics for Engineers*. Prentice-Hall, Inc., Englewood Cliffs, NJ, 4th edition, 1990.

[11] S. Narendra, D. Antoniadis, and V. De. Impact of using adaptive body bias to compensate die-to-die vt variation on within-die vt variation. In *International Symposium on Low-Power Electronics and Design*, pages 229–232, 1999.

[12] 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*, pages 19–23, Monterey, CA, August 12-14 2002.

[13] S. R. Nassif. Within-chip variability analysis. In *International Electronic Devices Meeting*, pages 283–286, 1998.

[14] S. R. Nassif. Design for variability in dsm technologies. In *First International Symposium on Quality Electronic Design*, pages 451–454, 2000.

[15] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, NY, 2nd edition, 1984.

[16] L. T. Pillage, R. A. Rohrer, and C. Visweswariah. *Electronic Circuit and System Simulation Methods*. McGraw-Hill, New York, NY, 1995.

[17] 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*, pages 64–67, Monterey, CA, August 12-14 2002.

[18] B. E. Stine, D. S. Boning, and J. E. Chung. Analysis and decomposition of spatial variation in integrated circuit processes and devices. *IEEE Transactions on Semiconductor Manufacturing*, 10(1):24–41, Feb. 1997.

[19] S. K. Thompson. *Sampling*. John Wiley & Sons, Inc., New York, NY, 2nd edition, 2002.

[20] J. Tschanz, J. Kao, S. Narendra, R. Nair, D. Antoniadis, A. Chandrakasan, and V. De. Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage. In *International Solid State Circuits Conference*, pages 344–345, 2002.