Full-Chip Model for Leakage Current Estimation Considering Within-Die Correlation

Khaled R. Heloue, *Student Member, IEEE,* Navid Azizi, and Farid N. Najm, *Fellow, IEEE*

***Abstract*—We present an efficient technique for finding the mean and variance of the full-chip leakage of a candidate design, while considering logic-structures and both die-to-die and within-die process variations, and taking into account the spatial correlation due to within-die variations. Our model uses a "random gate" concept to capture high-level characteristics of a candidate chip design, which are sufficient to determine its leakage. These high-level characteristics include information about the process, the standard cell library, and expected design characteristics. We show empirically that, for large gate count, the set of all chip designs that share the same high level characteristics have approximately the same leakage, with very small error. Therefore, our model can be used as either an *early* or a *late* estimator of leakage, with high accuracy. In its simplest form, we show that full-chip leakage estimation reduces to finding the area under a scaled version of the within-die channel length auto-correlation function, which can be done in constant time.**

***Index Terms*—Statistical Analysis, Leakage Power, Process Variations, Within-Die Correlation.**

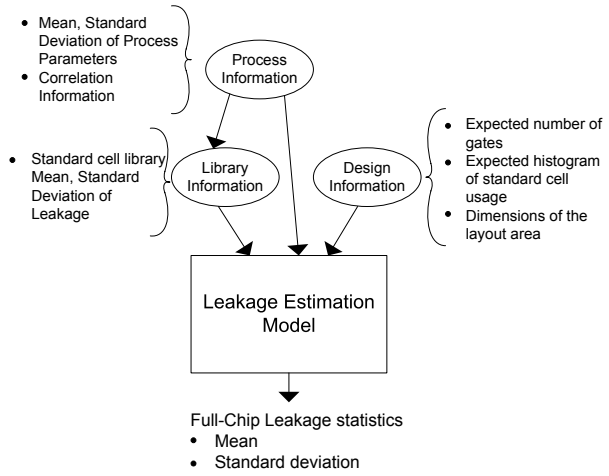

Fig. 1. Leakage Estimation Model and the High-level characteristics required

I. Introduction

As a result of technology scaling, leakage current is becoming a major design challenge, affecting both circuit performance and power. Leakage power is expected to continue to increase and due to limited power budgets, it will affect the feasibility of future microprocessor and ASIC designs [1]. Thus, estimating full-chip leakage becomes increasingly important. The leakage current of a circuit is not, however, simply the sum of the leakages of the devices in the circuit. Not only do logic-gate structures, such as stacking, affect the device leakage, but process variations make leakage estimation statistical in nature. Leakage current can be classified into two main components, namely sub-threshold leakage and gate tunneling leakage [2], [3]. Over the past few years, both types of leakage currents have been extensively studied, particularly sub-threshold leakage. Nevertheless, gate leakage can be important as noted in [4], although recent advances in the process, such as the introduction of new high-k materials and metal gates claim to have reduced gate leakage substantially.

Full-chip leakage estimation is useful at different points in the design flow. Towards the end of the design flow (late mode estimation), leakage estimation can be used as a final sign-off tool, and requires a complete netlist with possibly a circuit placement. On the other hand, early estimation of leakage (early mode estimation) provides the full-chip leakage given limited information about the design, which is very useful to allow for design planning.

Earlier work on leakage estimation [5], [6], [7], [8] concentrated on early mode estimators. However, they either did not consider logic-gate structures and other transistor topologies, and/or did not consider the effect of correlation between the variations on the total leakage, which is important to model. Narendra et al. [5] estimate the mean of full-chip sub-threshold leakage; they consider within-die variations, but ignore within-die correlations and do not take into account the effect of gate topologies. Furthermore, they do not estimate the standard deviation of full-chip leakage. Rao et al. [6] estimate sub-threshold leakage by first finding fitting parameters for the leakage current for individual gates in the library, and use the parameters to map the leakage distribution of the gate, due to within-die variations, to a log-normal distribution. They compute the total leakage of a circuit block using an approximation for sums of independent lognormal distributions. The authors, however, ignore within-die spatial correlations. Rao et al. have also tackled the problem of estimating full-chip leakage in another way [7]. They model different types of leakage separately as a product of a nominal value and a multiplicative function that represents the deviation from the nominal value due to variations. While, the authors separate the variations into global and local variations, the local varia-

K. R. Heloue and F. N. Najm are with the Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M5S3G4, Canada.

N. Azizi is with the Altera Toronto Technology Center, Altera Corporation, Toronto, ON M5S1S4, Canada.

This work was supported in part by Intel Corporation and in part by Altera Corporation.

tions are considered to be independent, and thus the effect of correlation is not factored into the final result; also, they do not provide an estimate of the standard deviation of full-chip leakage. Zhang et al. [8] in addition to considering process variations, also consider temperature and voltage variations. Instead of fitting the effect of process variations on leakage into an analytical equation, they use the BSIM model equations directly. However, just as the other early-estimators, they do not consider spatial correlations; they also do not consider gate topologies in their work.

More recent work [9], [3] has taken into consideration both the effects of gate topologies and within-die spatial correlation. Chang et al. [9] first precharacterize their library by fitting the different types of leakage currents to analytical forms. To model spatial correlation, they use the grid model [10] and determine the leakage in a grid by summing a set of correlated lognormal distributions, and then find the full-chip leakage distribution by summing the leakage distributions for each grid. Agarwal et al. [3] modeled spatial correlations differently using a quad-tree die partitioning method. To determine the final leakage distribution they sum the correlated lognormals using Wilkinson's method. Both these methods are late mode estimators of leakage, requiring minimally the circuit netlist and possibly a circuit placement to provide a leakage estimate. Also, since they operate at the level of the netlist, they can be expensive on large circuits, with a complexity of $\mathcal{O}(n^2)$ (some refinements are possible to reduce this cost, but with some loss of accuracy [9]).

Given the need to budget for power constraints, there is a need for accurate early mode estimators that take into consideration both correlation and gate topologies. As for late mode estimators, more efficient techniques are required. In this work, we present a new model and methodology for full-chip leakage estimation, in which certain high-level characteristics of a candidate chip design are used to determine its leakage statistics with high accuracy. For late mode estimation, these characteristics can be *extracted* from the netlist and/or placement. For early mode estimation, these characteristics can be simply specified as *expected* values based on previous design experience or on decisions made in the floorplanning stage. Our methodology uses a concept of a "random gate" to capture these characteristics and considers both correlations and gate topologies. We show that these high-level characteristics are sufficient to determine the leakage statistics of a design. We restrict our analysis to sub-threshold leakage estimation, although our mathematical framework can be easily extended to handle gate leakage.

A block diagram of the system is shown in Fig. 1. Given information about (1) the process, (2) the standard cell library, and (3) certain high-level design characteristics, we predict the mean and standard deviation of full-chip leakage. The process information includes the mean and standard deviation of the underlying process variations, such as the variations in transistor length or threshold voltage, and information regarding the within-die spatial correlation. The standard cell library information includes the leakage characteristics of the cell library under process variations; this information can be obtained by pre-characterizing the cells in the library. Finally, some information on the candidate design is needed, including the (extracted or expected) cell usage histogram (*i.e.*, frequency of use distribution) for cells in the library, the (extracted or expected) number of cells in the design, and the dimensions of the layout area. With this, we determine the full-chip leakage statistics (mean and variance) for the design.

To carry out the estimation, we propose a model which is *generic*, in the sense that it is a *template* for all designs that share the same values for these high-level characteristics. We use probability theory as the vehicle to implement this template, so that all designs that share the same values of these high-level characteristics will be *members* or *instances* of this probabilistic template model. We introduce the concept of the Random Gate (RG) which allows us to capture the characteristics of a candidate design. This allows the leakage statistics to be obtained in $\mathcal{O}(n)$ time, where n is the number of cells in the design, but we then also show that, for large gate counts, the statistics of the full-chip leakage can be written in integral form, allowing for the computational complexity of our estimator to become $\mathcal{O}(1)$ time[1]. The key point, the thesis of this work, is that *large* designs that share the same high-level characteristics will have approximately the same leakage statistics and, by leveraging this property, our estimation engine provides accurate and efficient estimation, either early or late in the design flow.

The rest of this paper is organized as follows: In Section II, we show how process parameters are modeled and propose to use a correlation function to model spatial correlations. In Section III, we present two flows, namely a Monte Carlo flow and an analytical flow, to characterize the different gates for leakage, and we show how to obtain the leakage correlation from the model. We also discuss some options in the face of uncertain signal probabilities. Section IV presents the random gate and the full-chip model, which are used in Section V to determine the full-chip leakage efficiently. We conclude in Section VI.

II. Modeling Process Variations

A. Parameter Model

Variations normally have two components: a Die-to-Die (D2D) component, and a Within-Die (WID) component. The D2D component is a variation between different instances of the die and is shared by all devices on the same die. The WID component of variation, however, causes different devices on the same die to have different process parameters; the WID variations have some correlation across the die. D2D and WID variations are considered to be (statistically) independent [11] and thus the total variance of a process parameter, such as the channel length (L), when both sources of variation are considered can be written as:

$$\sigma^2 = \sigma_{dd}^2 + \sigma_{wd}^2 \tag{1}$$

where σ_{dd}^2 is the variance of the D2D variation and σ_{wd}^2 is the variance of the WID variation. We will assume that all

[1]When used as a late mode estimator, there will be some additional cost to extract the cell usage histogram from the netlist, but that also can be constant-time, or linear-time in the worst case.

process parameters follow a Gaussian distribution, which is in line with the literature on leakage estimation. The resulting Random Variable (RV) for channel length (or any process parameter) can also be written with respect to their D2D and WID components as:

$$\mathbf{L}(i) = \mu + \sigma_{dd}\mathbf{Z}_0 + \sigma_{wd}\mathbf{Z}(i) \quad (2)$$

where i refers to an arbitrary device, μ is the mean of $\mathbf{L}$, $\mathbf{Z}_0$ is a zero mean standard normal RV with unit variance representing the D2D component which is shared by all devices on the die, and where the WID component is represented by a zero mean unit variance standard normal RV $\mathbf{Z}(i)$, a notation that emphasizes that it may be different for different devices on the die. For example, a second device j on the die will have a channel length with the following RV:

$$\mathbf{L}(j) = \mu + \sigma_{dd}\mathbf{Z}_0 + \sigma_{wd}\mathbf{Z}(j) \quad (3)$$

Given that $\mathbf{Z}(i)$ and $\mathbf{Z}(j)$ may be spatially correlated with a correlation of $\rho_{L_{wd}}(i,j)$, and that $\mathbf{Z}_0$ is shared by $\mathbf{L}(i)$ and $\mathbf{L}(j)$, then $\mathbf{L}(i)$ and $\mathbf{L}(j)$ will be correlated as well. We now show how we can express the total length correlation in terms of the D2D and WID breakdown, and the WID spatial correlation. Let us first write (2) and (3) as follows:

$$\mathbf{L}(i) - \mu = \sigma_{dd}\mathbf{Z}_0 + \sigma_{wd}\mathbf{Z}(i) \quad (4)$$
$$\mathbf{L}(j) - \mu = \sigma_{dd}\mathbf{Z}_0 + \sigma_{wd}\mathbf{Z}(j) \quad (5)$$

Assuming that the correlation between the WID variations $\rho_{L_{wd}}(i,j)$ is available as will be discussed in the next section, the covariance between the process parameters can then be written as:

$$E\left[(\mathbf{L}(i)-\mu)(\mathbf{L}(j)-\mu)\right] = \sigma_{dd}^2 + \sigma_{wd}^2\,\rho_{L_{wd}}(i,j) \quad (6)$$
$$= \sigma^2\rho_{L_{total}}(i,j) \quad (7)$$

where $\rho_{L_{total}}(i,j)$ is the total channel length correlation between $\mathbf{L}(i)$ and $\mathbf{L}(j)$ due to both D2D and WID components. We can now solve for the correlation between the total variation that composes both the D2D and WID variation to be:

$$\rho_{L_{total}}(i,j) = \frac{\sigma_{dd}^2 + \sigma_{wd}^2\,\rho_{L_{wd}}(i,j)}{\sigma^2} \quad (8)$$

The above equation can be simplified if we define α to be a ratio of the D2D variance to the total variance as:

$$\alpha = \frac{\sigma_{dd}^2}{\sigma^2} \quad (9)$$

which allows the total correlation to be written as a function of the correlation of the WID variation and the ratio of the D2D variance to the total variance:

$$\rho_{L_{total}}(i,j) = \alpha + (1-\alpha)\,\rho_{L_{wd}}(i,j) \quad (10)$$

The next section will describe the model that we used to capture the correlation in process parameters.

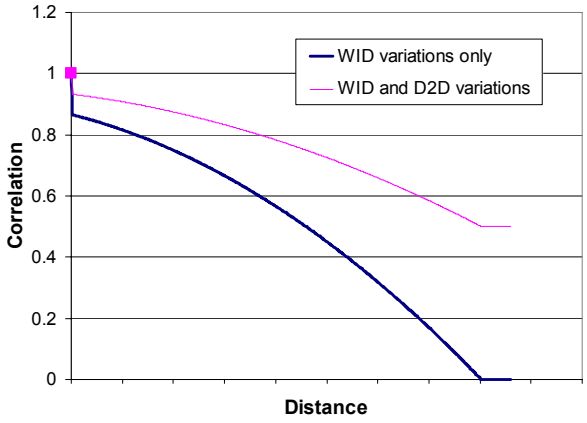

Fig. 2. Possible Correlation Model considering both WID and D2D variations

B. Correlation Model

Previous work on early-estimators of leakage did not take into account the spatial correlation that exists between the WID variations in the process parameters of different cells. However, in order to accurately estimate leakage, spatial correlation between variations must be taken into account [9], [3].

To model the WID spatial correlation between variations in transistor characteristics, we assume the existence of a spatial correlation function [12] that depends on the *distance* between the two transistors. Given the D2D and WID parameter variances, and the WID correlation, one can easily determine the total correlation between parameter variations (due to D2D and WID effects) by a simple normalization as was shown in (10). Not all functions, however, can be used as a spatial correlation function [12]; specifically spatial correlation functions are a family of monotonically decreasing non-negative functions [12]. One example of a spatial correlation function for WID variations is shown in bold line in Fig. 2; the function has a correlation of one at a distance of zero since the devices are in fact the same device [12]; furthermore the sudden drop from one at distance zero is the cause of a uncorrelated random variation in the parameter, which exists even in devices that are very close together [12]. The above spatial correlation function considers only the correlation in WID variations and therefore it dies down to zero after a certain distance. If we consider the D2D variation using (10), the total correlation will decrease down to α as shown in Fig. 2, where α is set to 0.5 [12].

III. Modeling at the Cell Level

While the statistics of the underlying process parameters can be obtained from the foundry, the leakage statistics of each cell can not be immediately obtained. Since each cell has a different topology, with different transistor stacks, the leakage in each cell is affected differently by the underlying variations in the transistor length and threshold voltage. Furthermore the cell's inputs also affect the leakage distribution of each cell.

A. Cell Leakage

Leakage current is determined primarily by transistor, not interconnect, parameters. Of the many transistor parameters that affect sub-threshold leakage, the truly relevant ones are channel-length (L) and threshold voltage (V_t), as shown in [13], due to the exponential dependence of sub-threshold leakage current on these two parameters. Threshold voltage variations are mainly due to two effects: random dopant fluctuations in the channel and the V_t roll-off effect whereby V_t varies in response to variations in L. For this work, when we refer to V_t variations, we specifically refer to the effect of random dopant fluctuations. We lump the effect of V_t roll-off on leakage into the L variations, because the two are directly related. This allows us to make the simple statement that V_t variations are purely random (independent) across the die [14], while L variations are not [9] (they include some within-die correlation). This approach is in line with the modern treatment of leakage in published work [6].

Since V_t variations are independent, while L variations are not, it follows immediately that, for full-chip leakage estimation, while V_t variations may be relevant for finding the *mean* of the total leakage, they are definitely not relevant for finding the *variance* of the total leakage. The reason for this is simple: the variance of the sum of n independent random variables is $\sim n\sigma^2$, while the variance of the sum of n highly correlated random variables is $\sim n^2\sigma^2$. Thus, for large chips (large n), the variance of chip leakage due to V_t variations is negligible compared to that due to L variations. This too is in line with the modern published work on leakage [6]. Thus, for leakage *variance* estimation, we can focus on L alone. As for the effect of V_t variations on the *mean* leakage, that can be easily determined through a multiplicative term that depends on the variance of V_t, which is derived from the mean of the log-normal distribution, similar to [15]. As this is standard textbook material, it will not be covered here.

To model the distribution of the leakage of each cell, we use two methods which have different levels of computational complexity and accuracy. The first method uses a Monte-Carlo (MC) analysis to obtain the leakage statistics of each cell. While this technique needs extensive simulations, it does give us some confidence in the resulting distributions. The second method, an analytical approach, uses a limited sampling of the leakage of the cell, and then fits the leakage of the cell into a functional form, from which we easily compute the mean and variance of the distribution. These two methods are discussed below, and we then discuss correlation and circuit state dependency.

1) Monte-Carlo Technique: We use a commercial 90nm CMOS technology, along with its associated standard cell library of which we use 62 cells which include the Static Random Access Memory (SRAM) cell, various flip flops and a range of different logic cells. For each cell and input combination, we perform a MC analysis to determine the mean and standard deviation of the cell's leakage. The MC analysis is done assuming all the variations in the transistor channel length within the cell are completely correlated, which is reasonable in practice given that the transistors in each cell are very close together. This is in line with previous work [9], where all cells within a grid are assumed to be completely correlated.

2) Analytical Technique: Rao et al. [6] introduced a mathematical model to express the leakage current, X, of a given cell as a function of channel length, L as:

$$X = ae^{bL+cL^2} \tag{11}$$

and showed that the analytical BSIM3 models vastly overestimated the leakage of devices that had gate lengths that deviated by more than 5% from their nominal, and that the fitted model above with the triplet (a, b, c) can accurately model the leakage of different topologies including individual transistors and transistor stacks [6].

In our work we first fit each cell's leakage into (11), and then use the triplet (a, b, c) to determine analytically the mean and variance of the underlying leakage distribution. To determine the triplet for each cell, we first perform a series of seven SPICE simulations, where the length of the transistors in the cell are modified from -3σ to 3σ in intervals of σ (where σ is the standard deviation in the transistor length) and measure the leakage. We then perform the Levenberg-Marquardt [16] method to fit the data into the above functional form and obtain (a, b, c). The model in (11) can fit the leakage of most cells quite well as can be seen in Fig. 3 where the analytical model is compared to SPICE simulations of a four-input-AND-into-OR cell. For some cells, however, the analytical model does not fit quite as well, as can be seen in Fig. 4, for a double-two-input-AND-into-two-input-NOR cell.

Note that, unlike [6] where numerical integration is used to approximate the leakage mean and variance, we use the fitted model with the triplet (a, b, c) to determine analytically and exactly the mean and variance of the underlying leakage distribution. The complete derivation, which was moved to Appendix A, results in the following:

$$\mu_{\mathbf{X}} = M_{\mathbf{Y}}(1) \tag{12}$$

$$\sigma^2_{\mathbf{X}} = M_{\mathbf{Y}}(2) - \mu^2_{\mathbf{X}} \tag{13}$$

where $M_{\mathbf{Y}}(t)$ is the moment-generating function of $\mathbf{Y} = \ln \mathbf{X}$ which can be shown to be:

$$M_{\mathbf{Y}}(t) = (1 - 2K_1 t)^{-\frac{1}{2}} e^{\left[\frac{K_2^2 K_1 t}{1-2K_1 t} + K_3 t\right]} \tag{14}$$

by using the moment generating function of the "Non-Central Chi-square" distribution where K_1, K_2 and K_3 are simple functions of the regression parameters (a, b, c) and the mean μ and standard deviation σ of the channel length, as follows:

$$K_1 = c\,\sigma^2 \qquad K_2 = \frac{1}{\sigma}\left(\frac{b}{2c} + \mu\right) \tag{15}$$

$$K_3 = \ln a + b\mu + c\,\mu^2 - c\left(\frac{b}{2c} + \mu\right)^2 \tag{16}$$

To check the accuracy of the analytical model in determining the mean and standard deviation of cell's leakage, we compare the results obtained from the fitted model to the results obtained through MC analysis for all 62 cells with all input combinations. For the mean, the analytical method is

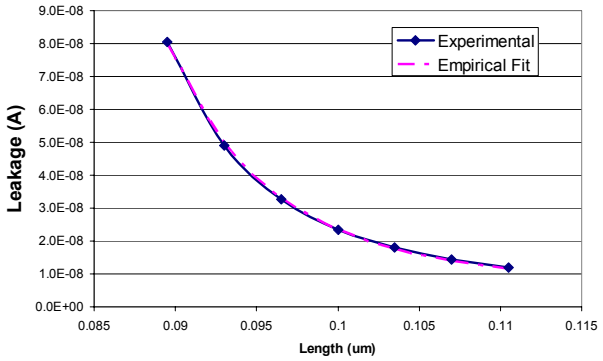

Fig. 3. Comparison of analytical fit with results from SPICE of an AO cell

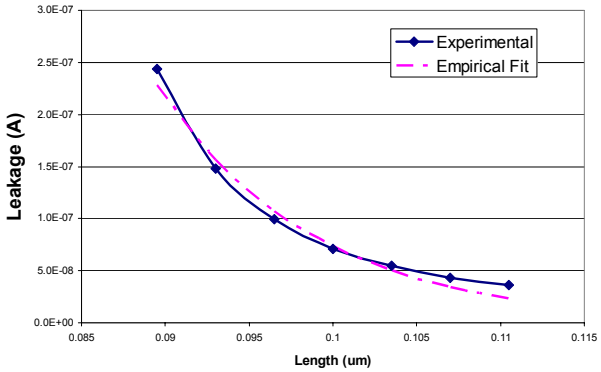

Fig. 4. Comparison of analytical fit with results from SPICE of an double-two-input-AND-into-two-input-NOR

quite close to the MC results; there is less than a 2% error for all gates, and the average absolute error is 0.44%. For the standard deviation, the average absolute error is 3.1%, and the maximum error is about 10%. The histogram of error in the mean and standard deviation for all cells and all input combinations is shown in Appendix A. Note that the error in the mean and standard deviation is not a result of the mathematical derivation, but due to the leakage curve not being exactly mapped to the functional form ae^{bL+cL^2}. Thus, there is a trade-off between computational complexity and accuracy; if MC analysis is performed on all gates, then the distribution models for all gates will have high accuracy; on the other hand, using the functional form requires minimal simulation time.

B. Leakage Correlation

As mentioned earlier, we assume the existence of a spatial correlation function which gives the correlation between *process parameters* as a function of the distance separating two locations, but which does not provide the correlation between the *leakages* of two cells at these locations. Using the regressed triplets, (a, b, c), we have developed an analytical method that determines the leakage correlation between any pair of gates placed at two arbitrary locations on the die given the correlation in their channel lengths. In other words, we

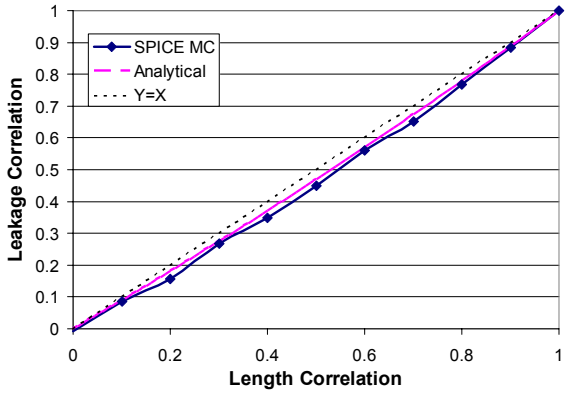

Fig. 5. Correlation in leakage vs correlation in channel length for a pair of gates

have determined a *mapping* $\rho_{m,n}(l_i, l_j) = f_{m,n}(\rho_L(l_i, l_j))$ where $\rho_L(l_i, l_j)$ is the channel length correlation between two locations l_i and l_j, $f_{m,n}(\cdot)$ is the derived mapping for gates m and n and $\rho_{m,n}(l_i, l_j)$ is the leakage correlation for gates m and n placed at locations l_i and l_j respectively. Note that the mapping depends on the types of gates m and n since a triplet (a, b, c) is associated with every gate type.

The details of the derivation leading to this mapping were omitted from this section and moved to Appendix B. Fig. 5 shows the results of the leakage correlation for a pair of gates given channel length correlation, as determined by the analytical mapping $f_{m,n}(\cdot)$, compared with the leakage correlation from MC analysis; note that the analytical technique shows a good match to the MC results. Also the leakage correlation is near the $y = x$ line, at which leakage correlation equals channel length correlation. We have performed the analysis for all pairs of gates, and shown that the analytical mapping provides accurate results in all cases. The set of mappings $f_{m,n}(\cdot)$ for different pairs of gates are slightly different but they all closely follow the $y = x$ line (refer to Appendix B). We will use this observation that the leakage correlation is close to the length correlation in the case where MC analysis is used to obtain the cell leakage statistics since we do not have the (a, b, c) triplet to obtain the leakage correlation exactly. We will discuss this in more details in Section V-A2.

C. Input Combinations

The signal probability (probability that a logic signal is 1) certainly has an effect on leakage. This effect is quite strong for single logic gates, causing a spread of 10X in some cases. However, for large circuits, the impact of signal probability is significantly diminished due to averaging of their effects (law of large numbers). To study this effect, we have swept the signal probabilities from 0 to 1 and have found, as shown in Fig. 6, that the effect on large circuit leakage is not pronounced and is also dependent on the frequency by which various cells are employed in the design. The figure shows the leakage mean, and similar behavior has been found for the leakage variance. For a practical solution approach, one has the option of simply setting the signal probabilities at

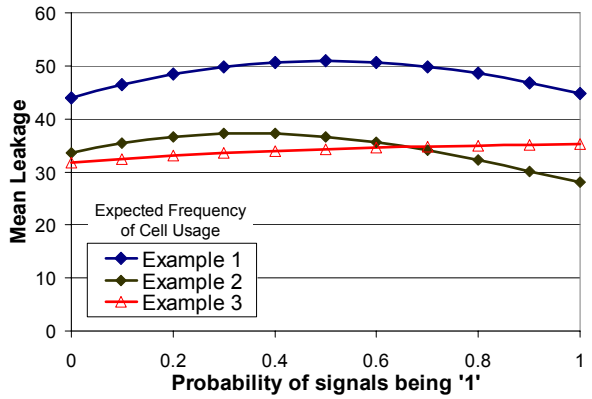

Fig. 6. Effects of signal probability on chip leakage

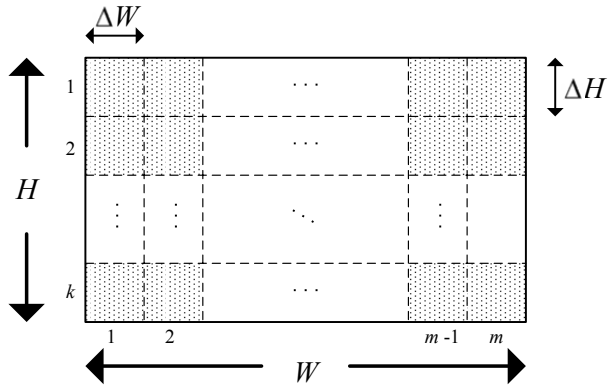

Fig. 7. Abstract organization of die

some ball-park mid-level value, such as 0.5. A better approach, which we employ, is to first characterize every cell for all its input states; then, based on this pre-characterized data, and for the given frequency of use distribution for cells, find the signal probability setting which maximizes the mean leakage, effectively finding the maximum of a plot such as Fig. 6. Empirically, we find that this setting turns to be very good for finding the maximum leakage mean for the candidate design, as well as its maximum leakage variance. This approach gives a conservative estimate, in the face of uncertainty about eventual signal probabilities.

IV. Full-Chip Model

What determines the leakage of a large circuit? We will demonstrate empirically that certain high-level characteristics of a candidate design are sufficient to determine its leakage. In a library-based standard-cell design environment, these characteristics are:

1) The cell library (characterized for leakage)
2) The (actual or expected) frequency of usage for cells in the library
3) The (actual or expected) number of cells in the design
4) The dimensions of the layout area

In order to carry out the leakage estimation, we propose a model for the candidate chip design which is *generic*, in the sense that it is a *template* for all designs that share the same values for these high-level characteristics. We use probability theory as the vehicle to construct this template, so that all designs that share the same values of these high-level characteristics will be *members* or *instances* of this probabilistic template model. After developing our leakage predictor based on this model, we will then show that the leakages of all instances of specific designs which are members of this model converge towards the predicted leakage value as the circuit size increases; Fig. 9 offers a a "sneak preview" of this convergence.

A. Model Definition and Suitability

Formally, our full-chip model is a rectangular array of a number (n) of identical *sites*, as shown in Fig. 7, where every site is occupied by a probabilistic abstraction which we call a *random gate* (RG), and such that the dimensions of the array are equal to the dimensions of the layout area of the candidate design, and that the number of sites n is equal to the number of cells in the design. But what is a RG? Simply put, a RG is similar to a Random Variable (RV); however, unlike a RV which assumes real numbers as *outcomes* or *instances*, the instances of a RG are gates from the standard-cell library, with probabilities identical to those in the frequency of use distribution. In other words, the RG discrete probability distribution is identical to the frequency of cell usage of the design.

This full-chip array model is a suitable probabilistic representation of all designs having the high-level characteristics highlighted earlier. On one hand, its dimensions and gate count match the dimensions of the layout and the number of cells in the candidate design. On the other hand, the frequency of cell usage of the design is also matched by the way the RG discrete probability distribution is defined. Hence, if an *instance* of the full-chip model is defined to be n RG instances at every site in the array, then the frequency of cell usage for that full-chip model *instance* will be identical to the frequency of cell usage of the candidate design, for large n. Therefore, the full-chip model is a probabilistic representation of a set of designs with the same high-level characteristics, and those designs are in fact *instances* of our model. Using this fact, we will use the full-chip model to estimate the leakage of the candidate design.

One possible reaction to this proposal is that all sites in the full-chip model are of identical size while obviously cells in the library are of different sizes. Another comment is that the array seems to leave no room for interconnect routing. Both these issues do not present a problem. In fact, the size of a site is really the size of the layout area, divided by the number of cells, thus it is the average size of a cell and the interconnect that may be associated with it. Thus, all that is captured by the notion of a RG site is the idea that the leakage due to one cell would on average be spread out or "allocated" to the layout area of a single site.

B. Leakage Statistics of a Random Gate

As stated earlier, the RG is simply a gate picked at random from the library, according to a discrete probability distribution which is identical to the frequency of gate usage. In order to perform full-chip leakage estimation based on our model, we need to construct and mathematically define the leakage statistics of the RG.

Let $\mathbf{I}$ be an RV that takes as values the *type* of a gate picked from the library at random to be used in the design. This means that $\mathbf{I} \in \{1, 2, \ldots, p\}$, where p is the total number of gates in the library, and that the distribution of $\mathbf{I}$ is identical to the frequency of gate usage. Let α_i be the frequency of usage of gate i. Then:

$$\mathcal{P}\{\mathbf{I} = i\} = \alpha_i \quad \forall i = 1, 2, \ldots, p \quad \text{and} \quad \sum_{i=1}^{p} \alpha_i = 1 \tag{17}$$

Let $\mathbf{X_I}$ be an RV that represents the leakage of a gate picked according to the distribution of $\mathbf{I}$. Then by definition, $\mathbf{X_I}$ is the leakage of the RG. Consequently, $\mathbf{X_I}$ is defined on two probability spaces; the space of $\mathbf{X}$ due to channel length variations, and the space of $\mathbf{I}$ due to the choice of gate type. Note that for an arbitrary realization of say $\mathbf{I} = i$, $\mathbf{X_I}$ will be equal to $\mathbf{X}_i$, that is the RV that represents the leakage of gate of type i. Recall that the statistics of $\mathbf{X}_i$, *i.e.*, its mean μ_i and standard deviation σ_i, have already been determined during pre-characterization for all gates i in the library, using either the MC or the analytical techniques. We can determine the mean leakage $\mu_{\mathbf{X_I}}$ of the RG as follows:

$$\mu_{\mathbf{X_I}} = E[\mathbf{X_I}] = E_I[E_X[\mathbf{X_I} \mid \mathbf{I} = i]] \tag{18}$$

$$= E_I[E_X[\mathbf{X}_i]] = \sum_{i=1}^{p} \alpha_i \mu_i \tag{19}$$

where $E_X[\cdot]$ and $E_I[\cdot]$ are the expected values over the spaces of $\mathbf{X}$ and $\mathbf{I}$, respectively. To determine the variance $\sigma^2_{\mathbf{X_I}}$ of $\mathbf{X_I}$, we start by determining its second moment $E\left[{\mathbf{X_I}}^2\right]$ as:

$$E\left[{\mathbf{X_I}}^2\right] = E_I\left[E_X\left[{\mathbf{X_I}}^2 \mid \mathbf{I} = i\right]\right] \tag{20}$$

$$= E_I\left[E_X\left[\mathbf{X}_i^2\right]\right] = \sum_{i=1}^{p} \alpha_i (\sigma_i^2 + \mu_i^2) \tag{21}$$

Given the second moment and the mean, the variance can be determined as follows:

$$\sigma^2_{\mathbf{X_I}} = E\left[{\mathbf{X_I}}^2\right] - \mu^2_{\mathbf{X_I}} = \sum_{i=1}^{p} \alpha_i (\sigma_i^2 + \mu_i^2) - \left(\sum_{i=1}^{p} \alpha_i \mu_i\right)^2 \tag{22}$$

To account for different input states, the summation in the above equations for the mean and variance are updated to account for the different weights of each input state.

C. Random Gate Leakage Correlation

In addition to the RG leakage statistics defined in the previous section, we need to construct and define the RG leakage correlation.

Recall that $\mathbf{X_I}$ is defined as the leakage of a random gate picked from the library according to the distribution of $\mathbf{I}$, and placed at some location on the die. Let $\mathbf{X_I}(l_i)$ and $\mathbf{X_I}(l_j)$ be the leakages of the two RGs at two arbitrary locations l_i and l_j. It is important to understand that $\mathbf{X_I}(l_i)$ and $\mathbf{X_I}(l_j)$ are identically distributed, and any correlation among these RVs is only due to the correlation over the space of process variations and not over the space of gate selection.

Let $C_{\mathbf{X_I}}(l_i, l_j)$ be the covariance of $\mathbf{X_I}(l_i)$ and $\mathbf{X_I}(l_j)$, which is defined as $C_{\mathbf{X_I}}(l_i, l_j) = E[\mathbf{X_I}(l_i)\,\mathbf{X_I}(l_j)] - \mu^2_{\mathbf{X_I}}$. It can be shown, using conditional expectation, that this covariance is given by:

$$C_{\mathbf{X_I}}(l_i, l_j) = \sum_{m=1}^{p} \sum_{n=1}^{p} \alpha_m \alpha_n C_{m,n}(l_i, l_j) \tag{23}$$

where $C_{m,n}(l_i, l_j)$ is the covariance of the leakage of two gates of types m and n, when placed at locations l_i and l_j, respectively, *i.e.*, $\mathbf{X}_m(l_i)$ and $\mathbf{X}_n(l_j)$. Note that the covariance of the leakage of the random gate $\mathbf{X_I}$ is the expected value over $\mathbf{I}$ of the covariances of all pairs of gate types. This result is somewhat intuitive since the random gate is an abstraction that embodies all gates in the library. Starting from (23), we can normalize $C_{m,n}(l_i, l_j)$ by the standard deviations of gates m and n to get their leakage correlation $\rho_{m,n}$. Then, we use the analytical mapping $f_{m,n}(\cdot)$ from Section III-B to relate the leakage correlation $\rho_{m,n}$ to channel length correlation ρ_L, as follows:

$$C_{\mathbf{X_I}}(l_i, l_j) = \sum_{m=1}^{p} \sum_{n=1}^{p} \alpha_m \alpha_n [\rho_{m,n}(l_i, l_j) \sigma_m \sigma_n]$$

$$= \sum_{m=1}^{p} \sum_{n=1}^{p} \alpha_m \alpha_n \sigma_m \sigma_n f_{m,n}(\rho_L(l_i, l_j)) \tag{24}$$

Let $F(\rho_L(l_i, l_j))$ be equal to the final expression in (24) above, and notice that this equation assumes that l_i and l_j are different. When they are the same, $C_{\mathbf{X_I}}(l_i, l_j)$ is just the variance $\sigma^2_{\mathbf{X_I}}$. Thus:

$$C_{\mathbf{X_I}}(l_i, l_j) = \begin{cases} F(\rho_L(l_i, l_j)) & \text{for} \quad l_i \neq l_j \\ \sigma^2_{\mathbf{X_I}} & \text{for} \quad l_i = l_j \end{cases} \tag{25}$$

By enforcing this correlation structure on our RG array, we ensure that instances of this array have the same correlation structure as the candidate design.

V. Full-Chip Leakage Estimation

For a specific placed design, based on a pre-characterized cell library, one can determine the full-chip leakage statistics using techniques from standard probability theory [17] for finding the sum of a number of correlated RVs (each RV corresponds to the leakage of one cell instance). This would be an $\mathcal{O}(n^2)$ approach, which can be expensive for large circuits (some refinements are possible to reduce this cost, but with some loss of accuracy [9]). Throughout this paper, we will refer to the leakage obtained from such an $\mathcal{O}(n^2)$ approach as the *true leakage* of a given design.

Apart from the issue of computational cost, such an approach is available only later in the design flow once a netlist and placement are available; it is useful only as a final check, and not as a prelude to corrective action. In this section, we

will first show how we can determine the full-chip leakage statistics in linear time, $\mathcal{O}(n)$, and then show how this can be improved to obtain the statistics in constant time, $\mathcal{O}(1)$. Importantly, we will also show that, for large gate counts, the statistics of any specific design that shares the same high-level characteristics under consideration converge to the values predicted by our model.

A. Linear-time method

Let $\mathbf{I}_T$ be an RV that represents the leakage of our full-chip model, *i.e.*, of the array of n RGs. This means that:

$$\mathbf{I}_T = \sum_{i=1}^{n} \mathbf{X}_\mathbf{I}(l_i) \tag{26}$$

where l_i is the location of the i^{th} random gate. We are interested in determining the statistics of $\mathbf{I}_T$, namely its mean $\mu_{\mathbf{I}_T}$ and variance $\sigma^2_{\mathbf{I}_T}$. The mean of $\mathbf{I}_T$ is equal to:

$$\mu_{\mathbf{I}_T} = E\left[\mathbf{I}_T\right] = \sum_{i=1}^{n} E\left[\mathbf{X}_\mathbf{I}(l_i)\right] = \sum_{i=1}^{n} E\left[\mathbf{X}_\mathbf{I}\right] = n\,\mu_{\mathbf{X}_\mathbf{I}} \tag{27}$$

The variance of $\mathbf{I}_T$ can be easily determined using a result from probability theory that the variance of a sum of correlated RVs is equal to the sum of pairwise covariances [17]. In other words:

$$\sigma^2_{\mathbf{I}_T} = \sum_{a=1}^{n}\sum_{b=1}^{n} C_{\mathbf{X}_\mathbf{I}}(l_a, l_b) \tag{28}$$

Note that the above double summation accounts also for the cases where $l_a = l_b$, for which the covariance is essentially the variance. Using the fact that any covariance can be written in terms of the correlation, $C_{\mathbf{X}_\mathbf{I}}(l_a, l_b) = \rho_{\mathbf{X}_\mathbf{I}}(l_a, l_b)\sigma^2_{\mathbf{X}_\mathbf{I}}$, we can write the total leakage variance in its final form:

$$\sigma^2_{\mathbf{I}_T} = \sigma^2_{\mathbf{X}_\mathbf{I}} \sum_{a=1}^{n}\sum_{b=1}^{n} \rho_{\mathbf{X}_\mathbf{I}}(l_a, l_b) \tag{29}$$

where the variance of the full-chip leakage is a function of the variance of the random gate and the extent of leakage correlation across the chip.

At this point, we have determined the mean of the total leakage (in constant time), and have shown that the computation of the variance of the total leakage requires a double summation over the number of gates on the chip. This $\mathcal{O}(n^2)$ complexity is not practically acceptable, especially knowing that n can be extremely large, on the order of millions. By taking into account the shape of the die and the sole dependence of the leakage correlation on the distance between different locations, we are able to cut down the complexity of computing the total leakage variance to $\mathcal{O}(n)$, as follows.

Let the RG array consist of k rows and m columns, where the total number of gates, n, is equal to the product $k \times m$, as shown in Fig. 7. Each location or "site" on the grid can be represented by a pair (r, s) where r is the horizontal index taking values $r = 1, \ldots, m$ and s is the vertical index taking values $s = 1, \ldots, k$. Also, assume that the height H and width W of the array are known. Let ΔH and ΔW be the height and width of the site where every gate will be placed.

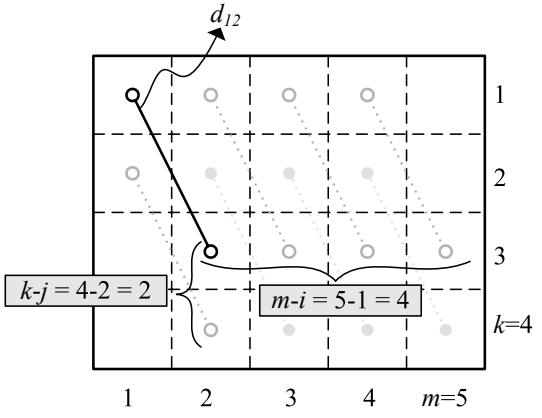

Fig. 8. Number of occurrences of a certain distance vector

Given the above parameters, the centre to centre distance d_{ij} between any two sites (r_1, s_1) and (r_2, s_2) can be easily determined to be $d_{ij}=\sqrt{(i\cdot\Delta W)^2+(j\cdot\Delta H)^2}$ where i is defined as the algebraic difference in horizontal indices, *i.e.*, $(r_2 - r_1)$, and j is defined as the algebraic difference in vertical indices, *i.e.*, $(s_2 - s_1)$. Note that $i = 0, \pm1, \ldots, \pm(m-1)$ and $j = 0, \pm1, \ldots, \pm(k-1)$.

Now recall the total leakage variance defined in (29) where the double summation covers all possible pairs of locations, and each location is a site on the grid defined by two indices. Since the correlation depends only on the distance d_{ij} between the pairs of locations, we can simplify the above expression greatly by performing the sum over the different distances rather than the pairs of locations. To do that, however, we need to determine the number of times each distance d_{ij} occurs. This is relatively easy for a rectangular $k \times m$ grid, as can be seen in Fig. 8, where the number of times a distance d_{ij} occurs along the width of the die is $m-i$ and along the height of the die is $k-j$. Using these two value, the number of occurrences n_{ij} of d_{ij} can be determined to be the following:

$$n_{ij} = (m - |i|) \cdot (k - |j|) \tag{30}$$

Since the leakage correlation between any two given locations depends only on the distance between these locations, we will explicitly highlight this fact, $\rho_{\mathbf{X}_\mathbf{I}}(l_a, l_b) = \rho_{\mathbf{X}_\mathbf{I}}(d_{ij})$ where i and j in the above equation are the algebraic differences in the horizontal and vertical indices of l_a and l_b.

Starting from (29), we will transform the quadratic summation that runs over all pairs of locations, into a summation that runs over the set of possible distances induced by the rectangular shape of the grid. This set will be covered if all the algebraic differences i and j are covered. After accounting for the number of times each algebraic difference occurs, n_{ij}, we get the following expression for the total leakage variance:

$$\sigma^2_{\mathbf{I}_T} = \sigma^2_{\mathbf{X}_\mathbf{I}} \sum_{i=-m}^{m}\sum_{j=-k}^{k} (m - |i|) \cdot (k - |j|)\ \rho_{\mathbf{X}_\mathbf{I}}(d_{ij}) \tag{31}$$

where the double summation runs at most $\mathcal{O}(k \times m) = \mathcal{O}(n)$ times, which is linear in circuit size. Note that the expres-

TABLE I
% ERROR IN FULL-CHIP STANDARD DEVIATION FOR ISCAS85 CIRCUITS COMPARED TO THE RG ESTIMATES

c499	c1355	c432	c1908	c880	c2670	c5315	c7552	c6288
1.04%	0.41%	1.14%	0.36%	0.74%	0.52%	0.23%	0.34%	1.38%

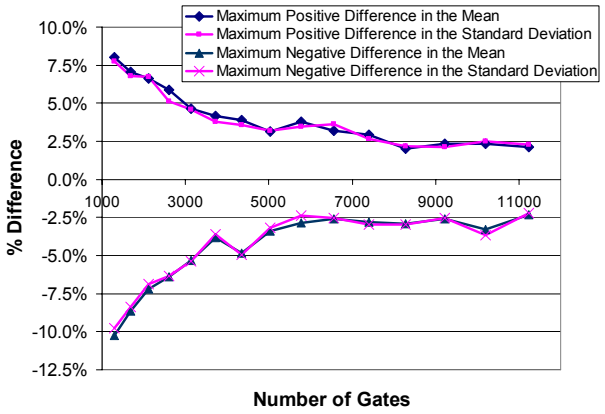

Fig. 9. Errors in the estimation of mean and standard deviation of full-chip leakage

sion in (31) is an exact transformation of (29) without any approximations, and was possible due to different factors:

1) The concept of random gate which allows us to express the total leakage as a sum of *identically* distributed RVs. This in turn made possible to extract the variance of the random gate outside the double summation in (29).
2) The sole dependence of the leakage correlation on the *distance* between the pair of locations rather on the location itself.
3) The rectangular shape of the grid which allows for closed form expression of the number of times each distance occurs.

Next, we validate our full-chip leakage model, both as an early and a late estimator of leakage.

1) Validation: Two types of validation tests were run, by first considering randomly generated circuits, as a way to make conclusions about the set of all circuits of a given size, and then by considering specific benchmark circuits.

In the first set of experiments, a large number of circuits were randomly generated so as to match a frequency of cell usage that was specified *a priori*. The circuits were then placed and routed, and their true leakage statistics (mean and variance) were found. Fig. 9 shows the maximum positive and negative difference between the means and standard deviations of the leakages of these circuits compared to the estimates provided by our model. It can be seen that as the number of gates in the circuits increases, the difference approaches zero; at a circuit size of 11,236 gates, the maximum difference is 2.2%. This small amount of error indicates that the set of all chip designs that share the same high level characteristics have approximately the same full chip leakage statistics and thus these high-level characteristics are sufficient to determine chip leakage. This first set of experiments serves to justify the statement that this approach is useful as an *early estimator* of full-chip leakage.

In the second set of experiments, we show how the model can be used as a *late estimator* of leakage for real (placed and routed) circuits. In this test, we have *extracted* the relevant high-level characteristics from each ISCAS85 circuit, namely the number of gates used, the histogram of cells used, and the dimensions of the layout; then with these values, we have used our model to estimate the leakage statistics of every circuit. Table I lists the errors in the full-chip leakage standard deviation, for all ISCAS85 circuits, between our model and the true leakage of these circuits. The errors are very small (notice, however, that these do not include any cell leakage modeling errors, which were discussed earlier in Section III). We do not show the errors in the mean leakage because they are truly negligible.

2) Simplified Correlation Assumption: In Section III we noted that the cell leakage statistics (*i.e.*, the mean and standard deviation of leakage) can be obtained in two ways; either (1) a MC analysis would be done or (2) the cell's leakage would be fitted into a functional form to get three fitting parameters (a, b, c). Using these parameters, the leakage mean and standard deviation were analytically obtained. The fitted parameters also allowed us to determine the leakage correlation between any pair of gates, $\rho_{m,n}$, given the channel length correlation ρ_L. Using the mapping, $f_{m,n}(\cdot)$, the RG leakage correlation was determined in (24).

If we, however, choose to obtain the leakage statistics of each cell through MC analysis, we would not be able to use $f_{m,n}(\cdot)$ to determine the leakage correlation between pairs of cells because the correlation mapping depends on the fitting parameters which are not available in MC mode. Without this mapping, the RG leakage correlation cannot be determined. The solution to this problem lies in Fig. 5, where we have noted that the leakage correlation of any pair of cells is approximately equal to the correlation in the channel length of these cells. In other words, $\rho_{m,n} \approx \rho_L, \ \forall m, n$. With this simplified correlation assumption, (24) can be used to determine the RG leakage correlation.

To determine the amount of error introduced by this assumption, we have compared the difference between the standard deviation when assuming $\rho_{m,n} = \rho_L$ compared to the analytical approach, *i.e.*, when using the true $f_{m,n}(\cdot)$ mapping. Regardless of whether we assume solely WID variations or have both WID and D2D variations, the percentage error is below 2.8%, as shown in Fig. 10.

B. Constant-time method

In this section, we show how, for large values of n, we can approximate the linear summation in (31) by an integral to obtain the statistics of full-chip leakage in constant time. This transformation is possible because the correlation function that shows up under the integral (as shown next) is a well-behaved monotonically decreasing function.

1) 2D Integration in Rectangular Coordinates: Starting from (31), let $x_i = i \cdot \Delta W$ and $y_j = j \cdot \Delta H$, and by multiplying

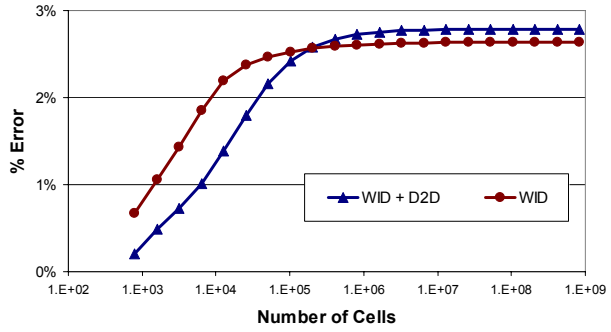

Fig. 10. % Error in leakage standard deviation for $\rho_{m,n} = \rho_L$ compared to $\rho_{m,n} = f_{m,n}(\rho_L)$

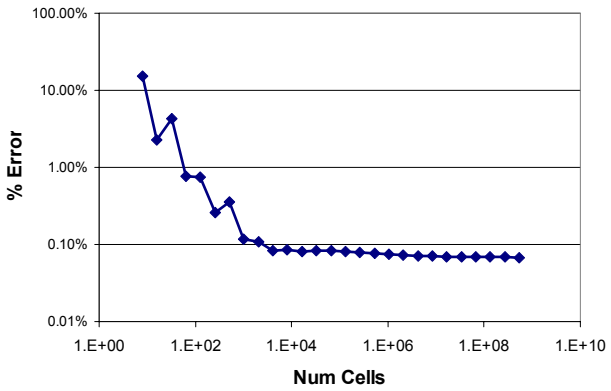

Fig. 11. % Error between numerical integration and linear time algorithm

out ΔW and ΔH we obtain:

$$\sigma^2_{\mathbf{I}_T} = \frac{\sigma^2_{\mathbf{X_I}}}{\Delta W \Delta H} \sum_{i=-m}^{m} \sum_{j=-k}^{k} (W - |x_i|) \cdot (H - |y_j|)\ \rho_{\mathbf{X_I}}(d_{ij}) \tag{32}$$

where $W = m \cdot \Delta W$, $H = k \cdot \Delta H$, and $d_{ij} = \sqrt{x_i^2 + y_j^2}$. By using a double integral to approximate the double summation over discrete values, we obtain:

$$\sigma^2_{\mathbf{I}_T} \approx \frac{\sigma^2_{\mathbf{X_I}}}{(\Delta W \Delta H)^2} \int_{-W}^{W} \int_{-H}^{H} (W-|x|)(H-|y|)\rho_{\mathbf{X_I}}\left(\sqrt{x^2+y^2}\right) \mathrm{d}y \mathrm{d}x \tag{33}$$

Let the area of a RG site be $\mathcal{A}_{\text{site}} = \Delta W \Delta H$ and the area of the die be $\mathcal{A} = n\mathcal{A}_{\text{site}}$. Note that the function being integrated is even, so that we can write:

$$\sigma^2_{\mathbf{I}_T} \approx 4\sigma^2_{\mathbf{X_I}} \frac{n^2}{\mathcal{A}^2} \int_0^W \int_0^H (W-x)(H-y)\rho_{\mathbf{X_I}}\left(\sqrt{x^2+y^2}\right) \mathrm{d}y \mathrm{d}x \tag{34}$$

The expression in (34) approximates the full-chip leakage variance for large values of n. Since the number of gates on the chip is typically in the order of millions, the approximation is valid in most cases. What is interesting about this expression is that it only requires the computation of an integral, which can be performed in constant-time using a good numerical integration routine; the leakage variance computation does not depend on the number of gates n, it is $\mathcal{O}(1)$.

2) 1D Integration in Polar Coordinates: To make our computation even more efficient, under certain conditions we can transform the double integral in (34) into a single integral in polar coordinates. First we write an exact mapping of (34) in double-integral form using polar coordinates:

$$\sigma^2_{\mathbf{I}_T} \approx \frac{4\sigma^2_{\mathbf{X_I}} n^2}{\mathcal{A}^2} \int_0^{\frac{\pi}{2}} \int_0^{D(\theta)} (W - r\cos\theta)(H - r\sin\theta)\rho_{\mathbf{X_I}}(r)\, r \mathrm{d}r \mathrm{d}\theta \tag{35}$$

where $D(\theta)$ is the distance from the origin to the boundary of the rectangular integration domain, which is less than the largest distance on the array. If the distance at which the WID correlation function reaches 0 is less than the minimum of the height or width of the array, then the double integral in (35) can be written as a single integral. To derive this single integral, let us for the moment assume that there are no D2D variations and that $\rho_{\mathbf{X_I}}$ becomes zero at a distance $D_{\max}$. If $D_{\max}$ is less than $\min(W, H)$ then (34) can be written as:

$$\sigma^2_{\mathbf{I}_T} \approx \frac{4\sigma^2_{\mathbf{X_I}} n^2}{\mathcal{A}^2} \int_0^{D_{\max}} \int_0^{\frac{\pi}{2}} (W - r\cos\theta)(H - r\sin\theta)\rho_{\mathbf{X_I}}(r)\, r \mathrm{d}\theta \mathrm{d}r \tag{36}$$

Since the correlation function does not depend on θ, we can further simplify the above expression by separating the integrals:

$$\sigma^2_{\mathbf{I}_T} \approx \frac{4\sigma^2_{\mathbf{X_I}} n^2}{\mathcal{A}^2} \int_0^{D_{\max}} \rho_{\mathbf{X_I}}(r) r \left[\int_0^{\frac{\pi}{2}} (W - r\cos\theta)(H - r\sin\theta) \mathrm{d}\theta \right] \mathrm{d}r \tag{37}$$

The expression in the brackets can be analytically integrated and results in the following expression:

$$g(r) \quad = \quad 0.5r^2 - (W+H)r + \frac{\pi}{2}WH \tag{38}$$

which leads to the final expression for full-chip leakage variance:

$$\sigma^2_{\mathbf{I}_T} \quad \approx \quad \frac{4\sigma^2_{\mathbf{X_I}} n^2}{\mathcal{A}^2} \int_0^{D_{\max}} r \cdot g(r) \cdot \rho_{\mathbf{X_I}}(r) \mathrm{d}r \tag{39}$$

When also considering D2D variations, recall from Section II that the correlation never reaches zero, and thus the single integral technique does not immediately apply. However, if we divide up the correlation function $\rho_{\mathbf{X_I}}(r)$ into a constant portion, ρ_c, and a portion that does go to 0 at $D_{\max}$, $\rho'_{\mathbf{X_I}}(r) = \rho_{\mathbf{X_I}}(r) - \rho_c$, then the single integral can be written as:

$$\sigma^2_{\mathbf{I}_T} \approx \left[\frac{4\sigma^2_{\mathbf{X_I}} n^2}{\mathcal{A}^2} \int_0^{D_{\max}} r \cdot g(r) \cdot \rho'_{\mathbf{X_I}}(r) \mathrm{d}r \right] + \rho_c \sigma^2_{\mathbf{X_I}} n^2 \tag{40}$$

3) Validation: The value of the standard deviation of the full-chip leakage obtained from the numerical integration (34) was compared to the value obtained from the $\mathcal{O}(n)$ approach presented in Section V-A.

As can be seen in Fig. 11, for circuits that have more than ten thousand gates there is less than 0.01% error between the numerical integration and that of the linear-time algorithm. For circuits with a small number of gates (<100) the % error is more than 1%; this is due to the granularity of the gates being a significant proportion of the total area of the design causing the integral to be less accurate than the true sum. For larger designs, the area of the logic gates compared to the area of the design approaches zero, allowing the numerical integration to provide good results, with less than 0.1% error.

Given that the $\mathcal{O}(n)$ time algorithm takes less than one second for circuits with less than 1000 gates, one can use the

$\mathcal{O}(n)$ time algorithm in those cases, and use the numerical integration for circuits with a much larger number of gates.

VI. Conclusion

We presented a probabilistic full-chip model that can be used to estimate, in constant-time, the leakage statistics of candidate designs either at an early or a late stage, while considering within-die correlations. We proposed and verified that certain high-level characteristics of a candidate chip design are sufficient to determine its leakage. These high-level characteristics, shown in Fig. 1, include information about the process, the standard-cell library, and the design in question. We showed that, for large gate count, the set of all chip designs that share the same high level characteristics have approximately the *same* full-chip leakage statistics, with very small error. We capture this set by a full-chip model based on Random Gates (RGs).

Appendix

A. Leakage Statistics - Analytical Method

In this section, we present the mathematical framework which allows us to analytically determine the mean and standard deviation of cell leakage, given the fitted functional form with the triplet (a, b, c). This framework can be applied for any leakage model that is quadratic exponential (which includes both sub-threshold and gate leakage). Recall that the leakage of each cell in the library is modeled as:

$$\mathbf{X} = a\, e^{b\mathbf{L}+c\mathbf{L}^2} \tag{41}$$

We are interested in determining the mean $\mu_{\mathbf{X}}$ and variance $\sigma^2_{\mathbf{X}}$ of the cell leakage, $\mathbf{X}$, given the mean and variance of the channel length, $\mathbf{L}$, and the regression parameters (a, b, c). Assume that $\mathbf{L}$ is a normally distributed RV with mean μ and standard deviation σ. Let $\mathbf{Y} = \ln \mathbf{X}$; then $\mathbf{X} = e^{\mathbf{Y}}$. The mean and variance of $\mathbf{X}$ can be written as follows:

$$\mu_{\mathbf{X}} = E[\mathbf{X}] = E\left[e^{\mathbf{Y}}\right] \tag{42}$$

$$\sigma^2_{\mathbf{X}} = E\left[\mathbf{X}^2\right] - \mu^2_{\mathbf{X}} = E\left[e^{2\mathbf{Y}}\right] - \mu^2_{\mathbf{X}} \tag{43}$$

Let $M_{\mathbf{Y}}(t)$ be the moment-generating function of $\mathbf{Y}$. By definition, this function is equal to:

$$M_{\mathbf{Y}}(t) = E\left[e^{t\mathbf{Y}}\right] \tag{44}$$

This function has been studied in the literature, and has a closed form expression for most known distributions [17]. Note that (42) and (43) can be written in terms of the moment-generating function of $\mathbf{Y}$:

$$\mu_{\mathbf{X}} = M_{\mathbf{Y}}(1) \tag{45}$$

$$\sigma^2_{\mathbf{X}} = M_{\mathbf{Y}}(2) - \mu^2_{\mathbf{X}} \tag{46}$$

The above result shows that the mean and variance of the cell leakage can be determined if $M_{\mathbf{Y}}(t)$ is known. To do that, we must determine the distribution of $\mathbf{Y}$. Since $\mathbf{Y} = \ln \mathbf{X}$, it follows from (41) that:

$$\mathbf{Y} = \ln a + b\mathbf{L} + c\mathbf{L}^2 \tag{47}$$

where $\mathbf{L}$ has a normal distribution. Let $\hat{\mathbf{L}}$ be a normalized version of $\mathbf{L}$:

$$\hat{\mathbf{L}} = \frac{\mathbf{L} - \mu}{\sigma} \tag{48}$$

This last equation shows that $\hat{\mathbf{L}}$ has a standard normal distribution with zero mean and unit variance. We can easily write (47) in the following form:

$$\mathbf{Y} = K_1 \left(\hat{\mathbf{L}} + K_2\right)^2 + K_3 \tag{49}$$

where:

$$K_1 = c\,\sigma^2 \tag{50}$$

$$K_2 = \frac{\left(\frac{b}{2c} + \mu\right)}{\sigma} \tag{51}$$

$$K_3 = \ln a + b\mu + c\,\mu^2 - c\left(\frac{b}{2c} + \mu\right)^2 \tag{52}$$

The motivation behind this transformation is to write $\mathbf{Y}$ in terms of an RV with a known distribution. Let $\mathbf{W} = \hat{\mathbf{L}} + K_2$; then $\mathbf{W}^2$ has a "Non-Central Chi-square" distribution with $\nu = 1$ degrees of freedom because $\mathbf{W}$ is normal with non-zero mean and unit variance. Therefore we can write (49) in terms of $\mathbf{W}^2$:

$$\mathbf{Y} = K_1\mathbf{W}^2 + K_3 \tag{53}$$

This allows to write the moment-generating function of $\mathbf{Y}$ in terms of the moment-generating function of $\mathbf{W}^2$ as follows:

$$M_{\mathbf{Y}}(t) = E\left[e^{t\mathbf{Y}}\right] \tag{54}$$

$$= E\left[e^{t\left(K_1\mathbf{W}^2+K_3\right)}\right] \tag{55}$$

$$= e^{tK_3} E\left[e^{tK_1\mathbf{W}^2}\right] \tag{56}$$

$$= e^{tK_3} M_{\mathbf{W}^2}(K_1 t) \tag{57}$$

Since $\mathbf{W}^2$ has a Non-Central Chi-square distribution with $\nu = 1$ degrees of freedom, then its moment-generating function is known [17]:

$$M_{\mathbf{W}^2}(t) = (1-2t)^{-\frac{1}{2}} \cdot e^{\frac{\lambda t}{1-2t}} \tag{58}$$

where $\lambda = (K_2)^2$ is the non-centrality parameter.

Now by using the above equation for the moment-generating function of $\mathbf{W}^2$, we can determine the moment-generating function of $\mathbf{Y}$ using (57) and get the final expression in (14), from which the mean and variance of the cell leakage can be determined as shown in (45) and (46).

To determine the accuracy of this analytical technique, we have compared its results to MC analysis. Histograms of the the percent error in the mean and standard deviation are shown in Fig. 12 and 13 respectively for all 62 cells with all input combinations. For the mean, the analytical method is quite

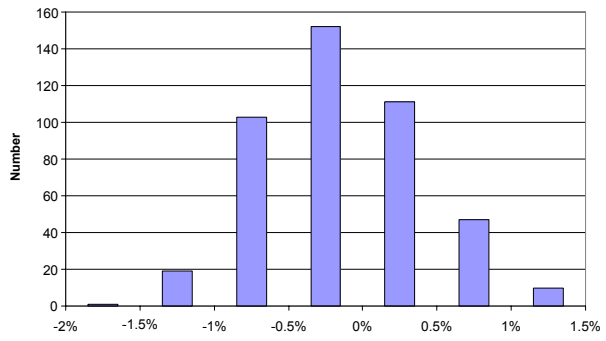

Fig. 12. Histogram of the % error in the mean of the analytical method compared to MC

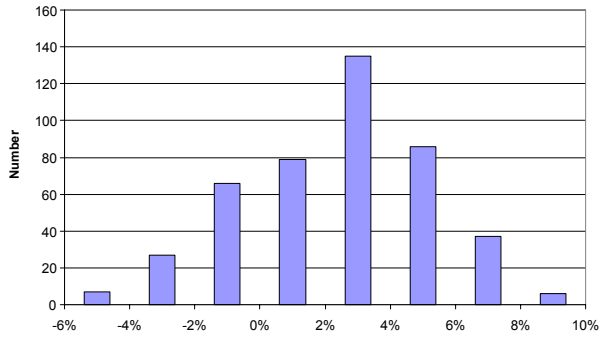

Fig. 13. Histogram % error in the standard deviation of the analytical method compared to MC

close to the MC results with errors less than a 2% for all gates. For the standard deviation the error is larger, with an average absolute error of 3.1% and a maximum error of about 10%. As mentioned in Section III-A2, the error in the mean and standard deviation is not a result of the mathematical derivation, but due to the leakage curve not being exactly mapped to (41).

B. Leakage Correlation - Analytical Mapping

In this section, we present the mathematical framework that allows us to analytically determine the correlation in the leakage currents of two cells given the correlation in their channel length. We also note that leakage correlation turns out to be very close to channel length correlation, in most cases.

Let $\mathbf{L}_1$ and $\mathbf{L}_2$ be two correlated RVs representing channel length at two arbitrary locations l_1 and l_2. We will assume that the correlation in channel length, $\rho_L(l_1, l_2)$, can be determined from the correlation model that we presented in Section II-B. Recall that $\mathbf{L}_1$ and $\mathbf{L}_2$ are normally distributed with mean μ and standard deviation σ.

We are interested in determining $\rho_{m,n}(l_1, l_2)$ defined as the correlation in the leakage of two gates m and n given the channel length correlation $\rho_L(l_1, l_2)$. Particularly, we will use the analysis that follows to find a mapping $f_{m,n}(\cdot)$ such that:

$$\rho_{m,n}(l_1, l_2) = f_{m,n}\left(\rho_L(l_1, l_2)\right) \tag{59}$$

Let $\mathbf{X}_m$ and $\mathbf{X}_n$ be the leakage of two gates of type m and n from the library; these RVs depend respectively on $\mathbf{L}_1$ and $\mathbf{L}_2$ in the following way:

$$\mathbf{X}_m = a_1 e^{b_1 \mathbf{L}_1 + c_1 \mathbf{L}_1^2} \tag{60}$$

$$\mathbf{X}_n = a_2 e^{b_2 \mathbf{L}_2 + c_2 \mathbf{L}_2^2} \tag{61}$$

It is important to understand that the channel length and leakage correlations are due to the spatial correlation between the locations l_1 and l_2, and depend particularly on the distance between the two locations. Let $C_{m,n}(l_1, l_2)$ be the covariance of $\mathbf{X}_m$ and $\mathbf{X}_n$ defined as follows:

$$C_{m,n}(l_1, l_2) = E\left[\mathbf{X}_m \mathbf{X}_n\right] - \mu_{\mathbf{X}_m} \mu_{\mathbf{X}_n} \tag{62}$$

The leakage correlation $\rho_{m,n}(l_1, l_2)$ can be expressed as a function of the covariance:

$$\rho_{m,n}(l_1, l_2) = \frac{C_{m,n}(l_1, l_2)}{\sigma_{\mathbf{X}_m} \cdot \sigma_{\mathbf{X}_n}} \tag{63}$$

where $\mu_{\mathbf{X}_m}$, $\mu_{\mathbf{X}_n}$, $\sigma_{\mathbf{X}_m}$, $\sigma_{\mathbf{X}_n}$ denote the means and standard deviations of $\mathbf{X}_m$ and $\mathbf{X}_n$ respectively, as determined in Section III-A2.

Examining (63) and (62), it is easy to see that the problem of finding $\rho_{m,n}(l_1, l_2)$ can be solved if $E\left[\mathbf{X}_m \mathbf{X}_n\right]$ is determined. By letting $\mathbf{Y} = \ln\left(\mathbf{X}_m \mathbf{X}_n\right)$, we can write $E\left[\mathbf{X}_m \mathbf{X}_n\right]$ as a function of the moment-generating function of $\mathbf{Y}$:

$$E\left[\mathbf{X}_m \mathbf{X}_n\right] = E\left[e^{\mathbf{Y}}\right] \tag{64}$$

$$= M_{\mathbf{Y}}(1) \tag{65}$$

since $M_{\mathbf{Y}}(t) = E\left[e^{t\mathbf{Y}}\right]$.

In this way, if we are able to determine the moment-generating function of $\mathbf{Y}$, we can evaluate it at 1 to determine $E\left[\mathbf{X}_m \mathbf{X}_n\right]$ using (65). Then, using (62) and (63), we can determine the leakage correlation $\rho_{m,n}(l_1, l_2)$. Using (60) and (61), we can write:

$$\begin{aligned} \mathbf{Y} &= \ln\left(\mathbf{X}_m \mathbf{X}_n\right) \\ &= \ln a_1 + \ln a_2 + b_1 \mathbf{L}_1 + b_2 \mathbf{L}_2 + c_1 \mathbf{L}_1^2 + c_2 \mathbf{L}_2^2 \end{aligned} \tag{66}$$

Assume that the correlation in the channel lengths $\mathbf{L}_1$ and $\mathbf{L}_2$ is $\rho_L(l_1, l_2) = \rho$. To model this correlation, we will use the following transformation; Let $\mathbf{Z}_1$ and $\mathbf{Z}_2$ be two RVs defined as follows:

$$\mathbf{Z}_1 = \frac{1}{2\alpha}\left(\frac{\mathbf{L}_1 - \mu}{\sigma} + \frac{\mathbf{L}_2 - \mu}{\sigma}\right) \tag{67}$$

$$\mathbf{Z}_2 = \frac{1}{2\beta}\left(\frac{\mathbf{L}_1 - \mu}{\sigma} - \frac{\mathbf{L}_2 - \mu}{\sigma}\right) \tag{68}$$

where:

$$\alpha = \sqrt{\frac{1+\rho}{2}} \tag{69}$$

$$\beta = \sqrt{1-\alpha^2} = \sqrt{\frac{1-\rho}{2}} \tag{70}$$

The way they are defined above, $\mathbf{Z}_1$ and $\mathbf{Z}_2$ are guaranteed to have certain properties. First, they are normally distributed

since $\mathbf{L}_1$ and $\mathbf{L}_2$ are jointly normally distributed. In addition, they are guaranteed to have zero-mean, unit variance, and zero correlation (or covariance). This can be easily shown:

$$E[\mathbf{Z}_1] = \frac{1}{2\alpha}\left(E\left[\frac{\mathbf{L}_1-\mu}{\sigma}\right] + E\left[\frac{\mathbf{L}_2-\mu}{\sigma}\right]\right) = 0 \tag{71}$$

$$\begin{aligned}\text{Var}(\mathbf{Z}_1) &= \frac{1}{4\alpha^2}\left(\frac{\text{Var}(\mathbf{L}_1)}{\sigma^2} + \frac{\text{Var}(\mathbf{L}_2)}{\sigma^2} + 2\frac{\text{Cov}(\mathbf{L}_1,\mathbf{L}_2)}{\sigma^2}\right) \\ &= \frac{2}{4(1+\rho)}(1+1+2\rho) = 1\end{aligned} \tag{72}$$

$$\begin{aligned}\text{Cov}(\mathbf{z}_1,\mathbf{z}_2) &= \frac{1}{4\alpha\beta}E\left[\left(\frac{\mathbf{L}_1-\mu}{\sigma}+\frac{\mathbf{L}_2-\mu}{\sigma}\right)\left(\frac{\mathbf{L}_1-\mu}{\sigma}-\frac{\mathbf{L}_2-\mu}{\sigma}\right)\right] \\ &= \frac{1}{4\alpha\beta}\left(E\left[\left(\frac{\mathbf{L}_1-\mu}{\sigma}\right)^2\right] - E\left[\left(\frac{\mathbf{L}_2-\mu}{\sigma}\right)^2\right]\right) = 0\end{aligned} \tag{73}$$

Being normally distributed with zero-mean and unit variance, and having zero correlation, imply that $\mathbf{Z}_1$ and $\mathbf{Z}_2$ are *independent standard normal* RVs. By reordering (67) and (68), we can express both $\mathbf{L}_1$ and $\mathbf{L}_2$ as a function of $\mathbf{Z}_1$ and $\mathbf{Z}_2$ as follows:

$$\mathbf{L}_1 = \sigma(\alpha\mathbf{Z}_1 + \beta\mathbf{Z}_2) + \mu \tag{74}$$

$$\mathbf{L}_2 = \sigma(\alpha\mathbf{Z}_1 - \beta\mathbf{Z}_2) + \mu \tag{75}$$

By substituting (74) and (75) in (66), we can write $\mathbf{Y}$ in the following matrix form:

$$\mathbf{Y} = K_1 + K_2 + \begin{bmatrix}K_3 & K_4\end{bmatrix}\begin{bmatrix}\mathbf{Z}_1\\ \mathbf{Z}_2\end{bmatrix} + \begin{bmatrix}\mathbf{Z}_1 & \mathbf{Z}_2\end{bmatrix}\begin{bmatrix}K_5 & K_6\\ K_7 & K_8\end{bmatrix}\begin{bmatrix}\mathbf{Z}_1\\ \mathbf{Z}_2\end{bmatrix} \tag{76}$$

where:

$$K_1 = \ln a_1 + b_1\mu + c_1\mu^2 \tag{77}$$

$$K_2 = \ln a_2 + b_2\mu + c_2\mu^2 \tag{78}$$

$$K_3 = \alpha\sigma[(b_1+b_2) + 2\mu(c_1+c_2)] \tag{79}$$

$$K_4 = \beta\sigma[(b_1-b_2) + 2\mu(c_1-c_2)] \tag{80}$$

$$K_5 = \alpha^2\sigma^2(c_1+c_2) \tag{81}$$

$$K_6 = \alpha\beta\sigma^2(c_1-c_2) \tag{82}$$

$$K_7 = K_6 \tag{83}$$

$$K_8 = \beta^2\sigma^2(c_1+c_2) \tag{84}$$

Generally, K_6 and K_7 are non-zero, which will lead to cross terms when performing the matrix multiplication in (76) (*i.e.,* terms in $\mathbf{Z}_1\mathbf{Z}_2$). These terms will complicate the determining of $\mathbf{Y}$ and we would ideally like to remove them from the expression. This can be achieved through matrix diagonalization; let's denote the 2×2 matrix in (76) by A. Because A is symmetric, we can diagonalize A in the following way:

$$A = PDP^T \tag{85}$$

where D is a diagonal matrix having as entries the eigenvalues of A, *i.e.,* $\begin{bmatrix}\lambda_1 & 0\\ 0 & \lambda_2\end{bmatrix}$ and P is a matrix having as columns the eigenvectors of A. Since A is symmetric, these eigenvectors are *orthogonal*. Moreover, we can choose P in such a way that its columns are also *orthonormal*. This decomposition is a standard practice and we use it here to transform $\mathbf{Z}_1$ and $\mathbf{Z}_2$ into a new set of RVs $\mathbf{V}_1$ and $\mathbf{V}_2$ that are also independent standard normals:

$$\begin{bmatrix}\mathbf{V}_1\\ \mathbf{V}_2\end{bmatrix} = P^T\begin{bmatrix}\mathbf{Z}_1\\ \mathbf{Z}_2\end{bmatrix} \tag{86}$$

Using the above transformation, (76) is written as follows:

$$\mathbf{Y} = K_1 + K_2 + \begin{bmatrix}\hat{K}_3 & \hat{K}_4\end{bmatrix}\begin{bmatrix}\mathbf{V}_1\\ \mathbf{V}_2\end{bmatrix} + \begin{bmatrix}\mathbf{V}_1 & \mathbf{V}_2\end{bmatrix}\begin{bmatrix}\lambda_1 & 0\\ 0 & \lambda_2\end{bmatrix}\begin{bmatrix}\mathbf{V}_1\\ \mathbf{V}_2\end{bmatrix} \tag{87}$$

where:

$$\begin{bmatrix}\hat{K}_3\\ \hat{K}_4\end{bmatrix} = P^T\begin{bmatrix}K_3\\ K_4\end{bmatrix} \tag{88}$$

Note that both D and P can be easily determined for a 2×2 symmetric matrix as there is a closed form expression for the eigenvalues and eigenvectors of A. Note also that since the off-diagonal entries of D are zero, $\mathbf{Y}$ will have no cross terms.

Now that we removed the cross terms, we can write $\mathbf{Y}$ into the following quadratic form:

$$\begin{aligned}\mathbf{Y} &= K_1 + K_2 + \hat{K}_3\mathbf{V}_1 + \hat{K}_4\mathbf{V}_2 + \lambda_1\mathbf{V}_1^2 + \lambda_2\mathbf{V}_2^2 \\ &= (K_1 + \hat{K}_3\mathbf{V}_1 + \lambda_1\mathbf{V}_1^2) + (K_2 + \hat{K}_4\mathbf{V}_2 + \lambda_2\mathbf{V}_2^2)\end{aligned} \tag{89}$$

where $\mathbf{V}_1$ and $\mathbf{V}_2$ are independent standard normal RVs.

Note that the above equation is essentially two instances of (47). Using exactly the same analysis that follows (47), we can write $\mathbf{Y}$ in terms of two independent RVs that have a non-central chi-square distribution; this allows us to determine the moment generating function of $\mathbf{Y}$, $M_{\mathbf{Y}}(t)$. Once this is done, we use (65) to find $E[\mathbf{X}_m\mathbf{X}_n]$, and consequently determine $\rho_{m,n}(l_1,l_2)$ from (62) and (63).

The above analysis, whereby the leakage correlation between any pair of gates placed at two arbitrary locations can be determined given the correlation in the channel length at these two locations, is referred to as the mapping $f_{m,n}(\cdot)$:

$$\rho_{m,n}(l_1,l_2) = f_{m,n}(\rho_L(l_1,l_2)) \tag{90}$$

We have used this mapping to determine the leakage correlation between the pairs of cells in our library. The results obtained for all pairs of gates, while being close to each other, are not exactly the same as can be seen in Fig. 14, where the correlation of 63 pairs of gates are plotted compared to the $y = x$ line. The difference can be better seen in a zoomed version of the plot in Fig. 15. The resulting curves are convex functions that pass through $(0,0)$ and $(1,1)$; they closely follow the $y = x$, deviating slightly at $\rho_L(l_1,l_2) = 0.5$.

REFERENCES

[1] S. Borkar. Design Challenges of Technology Scaling. *IEEE MICRO*, pages 23–29, 1999.

[2] A. Agarwal, C.H. Kim, S. Mukhopadhyay, and K. Roy. Leakage in nano-scale technologies: mechanisms, impact and design considerations. In *Proceedings of the 41st annual conference on Design automation*, pages 6–11. ACM New York, NY, USA, 2004.

[3] Amit Agarwal, Kunhyuk Kang, and Kaushik Roy. Accurate estimation and modeling of total chip leakage considering inter- & intra-die process variations. *IEEE International Conference on Computer-aided Design*, 2005.

[4] D. Lee, W. Kwong, D. Blaauw, and D. Sylvester. Simultaneous sub-threshold and gate-oxide tunneling leakage current analysis in nanometer CMOS design. *ISQED*, pages 287–292, 2003.

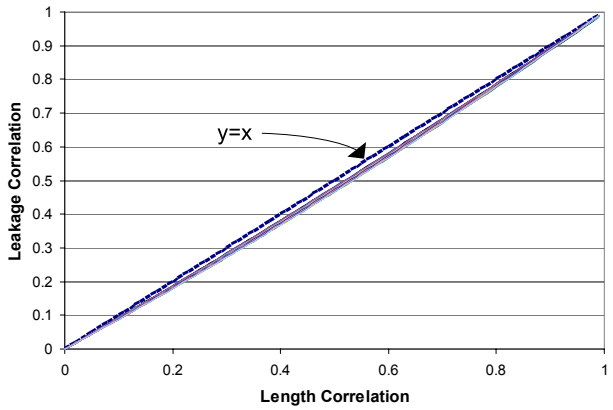

Fig. 14. Leakage Correlation of pairs of different gates

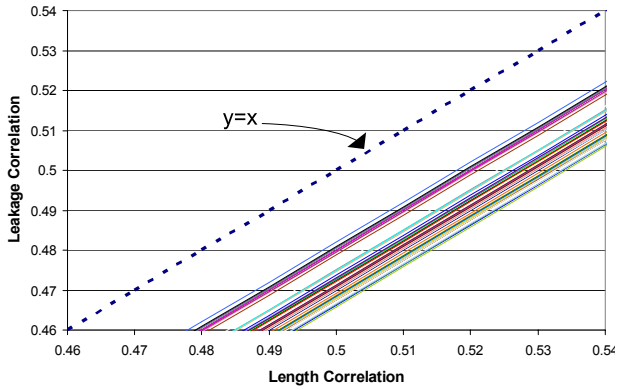

Fig. 15. Leakage Correlation of pairs of different gates (Zoomed)

[5] Siva Narendra, Vivek De, Dimitri Antoniadis, and Anantha Chandrakasan. Full-chip sub-threshold leakage power prediction model of sub-0.18μm CMOS. *IEEE/ACM International Symposium on Low Power Electronics and Design*, 2002.

[6] Rajeev Rao, Ashish Srivastava, David Blaauw, and Dennis Sylvester. Statistical analysis of subthreshold leakage current for VLSI circuits. *IEEE Transactions on VLSI Systems*, 12(2):131–139, February 2004.

[7] Rajeev R. Rao, David Blaauw, Dennis Sylvester, and Anirudh Devgan. Modeling and analysis of parametric yield under power and performance constraints. *IEEE Design & Test of Computers*, pages 376–385, July-August 2005.

[8] Songqing Zhang, Vineet Wason, and Kaustav Banerjee. A probabilisitic framework to estimate full-chip subthreshold leakage power distribution considering within-die and die-to-die P-T-V variations. *IEEE/ACM International Symposium on Low Power Electronics and Design*, 2004.

[9] Hongliang Chang and Sachin S. Sapatnekar. Full-chip analysis of leakage power under process variations, including spatial correlations. *IEEE Design Automation Conference*, 2005.

[10] H. Chang and S. S. Sapatnekar. Statistical timing analysis considering spatial correlations using a single PERT-like traversal. In *ICCAD*, pages 621–625, San Jose, CA, November 9-13 2003.

[11] S. G. Duvall. Statistical circuit modeling and optimization. In *Int'l Workshop on Statistical Metrology*, pages 56–63, June 2000.

[12] Jinjun Xiong, Vladimir Zolotov, and Lei He. Robust extraction of spatial correlation. *International Symposium on Physical Design*, 2006.

[13] A. Srivastava, R. Bai, D. Blaauw, and D. Sylvester. Modeling and analysis of leakage power considering within-die process variations. In *Proceedings of the 2002 international symposium on Low power electronics and design*, pages 64–67. ACM New York, NY, USA, 2002.

[14] Ali Keshavarzi, Gerhard Schrom, Stephen Tang, Sean Ma, Keith Bowman, Sunit Tyagi, Kevin Zhang, Tom Linton, Nagib Hakim, Steven Duvall, John Brews, and Vivek De. Measurements and modeling of intrinsic fluctuations in mosfet threshold voltage. *IEEE/ACM International Symposium on Low Power Electronics and Design*, 2005.

[15] D. Helms, G. Ehmen, and W. Nebel. Analysis and modeling of subthreshold leakage of RT-components under PTV and state variation. *Proceedings of the 2006 international symposium on Low power electronics and design*, pages 220–225, 2006.

[16] J.J. Mor. The Levenberg-Marquardt algorithm: implementation and theory.

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

PLACE PHOTO HERE

Khaled R. Heloue (S'01) received the B.E. degree in computer and communications engineering (with high distinction) from the American University of Beirut, Beirut, Lebanon, in 2003 and the M.A.Sc. degree in electrical and computer engineering from the University of Toronto, Toronto, ON, Canada, in 2005, where he is currently working toward the Ph.D. degree in the Department of Electrical and Computer Engineering.

His research includes computer-aided design (CAD) for integrated circuits, with a focus on timing analysis under variability, yield modeling and estimation, design for manufacturability (DFM), and full-chip leakage modeling and estimation. He interned with Intel's Strategic CAD Labs, Hillsboro, OR, in the summer of 2008.

PLACE PHOTO HERE

Navid Azizi received the B.A.Sc. degree, the M.A.Sc. degree, and Ph.D. degree in electrical and computer engineering from the University of Toronto, Toronto, ON, Canada in 2001, 2003, and 2007 respectively. He is currently working at the Altera Toronto Technology Centre in Toronto, ON, Canada as a Senior Software Engineer working on simultaneous switching noise (SSN) and memory interface modeling. He has previously worked with Intel Corporation in Hillsboro, OR.

His research interests include low-leakage and low-variability digital integrated circuit design and analysis. He has authored or co-authored 15 papers in refereed internal conferences and journals. Navid received the Postgraduate Scholarship Award from the National Sciences and Engineering Research Council (NSERC) of Canada for 2001 through 2003 and 2003 through 2005.

PLACE PHOTO HERE

Farid N. Najm (S'85-M'89-SM'96-F'03) is a Professor in the department of electrical and computer engineering (ECE) at the University of Toronto, Ontario, Canada. He received the B.E. degree in electrical engineering from the American University of Beirut in 1983 and the Ph.D. degrees in ECE from the University of Illinois at Urbana-Champaign (UIUC) in 1989. From 1989 to 1992, he worked with Texas Instruments, Dallas, Texas. He then joined the ECE department at UIUC as an Assistant Professor and became Associate Professor in 1997. In 1999, he joined the ECE department at the University of Toronto, where he is now Professor, and where he has served as department Vice-Chair from 2004 to 2007.

Dr. Najm is a Fellow of the IEEE, and is Associate Editor for the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS. He has received an IEEE Transactions on CAD Best Paper Award, an NSF Research Initiation Award, an NSF CAREER Award, and was an Associate Editor for the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS from 1997 to 2002. He serves on the executive committee of the International Symposium on Low-Power Electronics and Design (ISLPED), and has served on the technical committees of various conferences, including ICCAD, DAC, CICC, ISQED, and ISLPED. His research is on CAD for VLSI, with an emphasis on circuit level issues related to power, timing, variability, and reliability.