A Stochastic Approach To Power Grid Analysis

Sanjay Pant, David Blaauw, Vladimir Zolotov*, Savithri Sundareswaran*, Rajendran Panda*

University of Michigan, Ann Arbor, MI
*Motorola Inc., Austin, Texas

ABSTRACT

Power supply integrity analysis is critical in modern high performance designs. In this paper, we propose a stochastic approach to obtain statistical information about the collective IR and LdI/dt drop in a power supply network. The currents drawn from the power grid by the blocks in a design are modelled as stochastic processes and their statistical information is extracted, including correlation information between blocks in both space and time. We then propose a method to propagate the statistical parameters of the block currents through the linear model of the power grid to obtain the mean and standard deviation of the voltage drops at any node in the grid. We show that the run time is linear with the length of the current waveforms allowing for extensive vectors, up to millions of cycles, to be analyzed. We implemented the approach on a number of grids, including a grid from an industrial microprocessor and demonstrate its accuracy and efficiency. The proposed statistical analysis can be use to determine which portions of the grid are most likely to fail as well as to provide information for other analyses, such as statistical timing analysis.

Categories and Subject Descriptors

B.8.2 [Performance and Reliability]: Performance Analysis and Design Aids.

General Terms

Algorithms, Reliability.

Keywords

IR drop, Ldi/dt, power supply networks.

1 INTRODUCTION

Power supply networks are essential in providing reliable and constant power/gnd supply to the onchip devices. Due to the parasitic resistance, capacitance and inductance of the interconnects, as well as the I/O packages, the power supply is non-ideal and exhibits fluctuations, both in space and time. These fluctuations can either increase the delay of gates in the circuit, reducing the operating frequency, or inject noise in the circuit, making the circuit prone to functional failures. The voltage drop occurring in a supply network can be broadly classified into IR-drop, which is due to the parasitic resistances of the interconnects and LdI/dt drop, which is due to the package and on-chip inductance of interconnects.

For modern high performance designs, it is not uncommon for the supply network to conduct as much as 50-100 Amperes of total current [1]. Furthermore, current consumption is expected to increase further, due to the increasing complexity of the designs and reduction in supply voltage, making it more difficult to meet stringent supply integrity constraints. In particular, the LdI/dt drop is expected to become more prominent with increasing operating frequency and current consumption [2].

Power supply analysis involves analysis of multimillion linear elements constituting the power grid and millions of non-linear devices on the chip. Modern microprocessors with on-chip memory can have more than 25 million nodes [3]. As a first step, the non-linear devices are replaced by current sources in order to make the whole network linear. Nevertheless, simulation of such a network is extremely challenging and significant research has been focussed on developing fast simulation techniques [4][5][6][7]. However, even with fast linear solvers, it is generally not possible to simulate the power grid for more than handful of cycles in reasonable time, making the selection of the input vectors which induce the largest voltage drop extremely crucial.

Power grid analysis can be classified into input vector based approaches and vectorless approaches. The vector based approaches employ search techniques to find a set of input patterns which cause the worst case drop in the grid. A number of methods have been proposed which use genetic algorithms or other search techniques to automatically find vectors that maximize the total current drawn from the supply network [8][9]. These approaches are computationally intensive and are limited to circuit blocks rather than full chip analysis. Furthermore, these approaches are inherently optimistic, underestimating the voltage drop and thus letting some of the supply noise integrity problems go unnoticed. The vectorless approaches aim to compute an upper bound on the worst-case drop in an efficient manner. A number of vectorless approaches for constructing worst-case currents have been proposed, using methods such as propagation of timing windows [10] or constraint graph formulations [11]. Vectorless approaches have the advantage of being fast and conservative. However, these methods address only the static IR-drop analysis and not the LdI/dt drop, which is increasingly becoming a key concern in supply integrity analysis.

In this paper, we propose an alternate approach for supply voltage drop analysis which is based on stochastic analysis. We determine the statistical parameters of supply voltage variation at any node in the circuit, which includes both IR and LdI/dt drops. The variability is defined over the input vector space where different vectors cause different currents to be drawn from the power supply network. We model the currents drawn by major blocks in the design as stochastic processes and extract their statistical information which includes correlations between different blocks both in space and in time. We present an approach to propagate this information through the linear model of a power grid and show how we can obtain the distribution of voltage drops at any node in the grid.

The statistical characteristics of supply variations can be useful in a number of ways. First, they enable the designer to identify the regions in the grid which are more likely to fail and should be given higher priority when the grid is corrected. The probability distributions of the voltage drops can also be used to obtain the distribution of the delay of gates in the critical paths of a circuit, which can be used in statistical timing analysis to compute the probability distribution of the circuit delay. In our analysis, we found that the occurrence of the worst-case drop is an extremely rare event. This demonstrates that the traditional worst-case analysis, where each

gate delay is characterized with the worst-case drop, can be very conservative and illustrates the need for a statistical power grid analysis approaches.

A particular advantage of the proposed approach is that the statistical information is obtained directly from block currents and allows for very large sets of input vectors to be incorporated in the analysis, which is not feasible in traditional full power grid simulation. We implemented the proposed approach on a number of grids, including a power grid extracted from an industrial processor design. We compared the results against SPICE simulation and demonstrate the efficiency and accuracy of the proposed method.

The remainder of the paper is arranged as follows. In section 2, we describe the behavior of a linear system, modelling the power supply network as a linear system with currents as random processes and derive the statistical parameters of the voltage drops. Section 3 presents the results obtained for different power grids. We draw our conclusions in section 4.

2 PROPOSED APPROACH

In this section, we present the proposed method for computing the statistical parameters of supply voltage variations at any node in the power supply network. Figure 1 shows the general flow of the analysis. A chip design can consist of millions of transistors forming a sea of gates. As a first step, this sea of gates is grouped into large blocks such that there is minimum correlation in the currents drawn by these blocks. Each block is simulated using a suitable power simulator, such as PowerMill to obtain the currents drawn by the blocks over time. Next, the mean, auto-correlation function, and cross-correlation functions are computed for each block current. Since the complexity of computing these correlation functions is linear with the vector length, very large vector sequences can be accommodated, consisting of millions of cycles or more. The power supply network is modelled as a linear system with block currents as stochastic processes characterized by the extracted statistical information. We then compute the impulse response for every node due to each block current by simulating the power grid in SPICE or a fast linear solver. The impulse response attains its steady state quickly and the grid needs to be simulated only for a short period of time. These impulse responses, along with the statistical parameters of block currents are then used to generate statistical parameters for the voltage drop. Initially, the block currents are assumed to be independent and later both spatial and temporal correlations in block currents are incorporated in the voltage drop statistical analysis. Finally, we show that the voltage drops closely approximate normal distributions, allowing arbitrary confidence points on the voltage distribution to be obtained. We now discuss each of the analysis steps in more detail.

2.1 Power grid as a linear system

The power supply network of a chip consists of the ideal supply voltage sources, power and ground wires modelled as a linear RLC
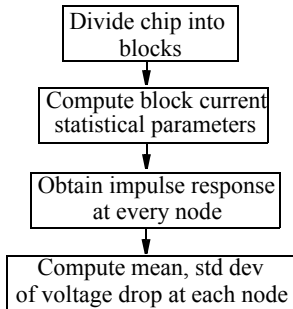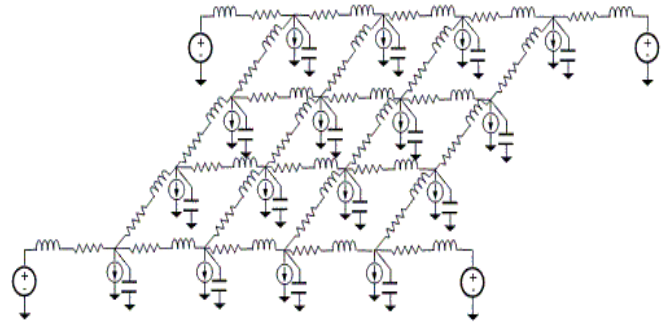

Figure 2. Two layers of a V_{dd} supply network

network, non-linear logic cells/blocks which draw current while switching and explicit decoupling capacitances, introduced for filtering out the high frequency power supply voltage variations. Ideal supply voltage sources are connected to the power and ground grids either by bond wires or by C4s in flip chips. The common approach to model the power grid is to substitute the non-linear logic cells/blocks with linear current sources and capacitances in parallel. This substitution of logic cells/blocks with their respective current sources transforms the entire nonlinear circuit into a linear network. The value of each current source is exactly equal to the current drawn by the corresponding logic cell/block. These currents are a function of time and can be obtained by simulating the circuit either at transistor or gate level. Thus, the computation of the currents is decoupled from the analysis of power grid which greatly simplifies the analysis. This is based on the implicit assumption that the values of the currents consumed by the cells do not depend on the voltage drop in the supply network which has been shown to introduce a small conservative error [7]. The capacitance in parallel with each current source models the capacitance of non-switching CMOS transistors in each cell/block. The non-switching capacitance in the design acts as decoupling capacitance and helps in filtering out the high frequency components of supply variations. Bond wires or C4s are modelled with their equivalent RL models. Another common simplification is to model the power and the ground grids separately. All the current sources and decoupling capacitances are connected to a common sink (ground) instead of being connected between power and ground grids. Figure 2 shows the linear model of a 2 metal layered power supply network. A multi-level power grid structure is similar. In this paper we will assume these simplifications to transform the power supply network into a linear system.

Using the above simplifications, the model of the power supply network consists of RLC elements, time varying current sources and constant voltage sources. We are interested in voltage variations at each supply point where the logic cells are connected. Thus, the simplified model of the power supply network can be considered as a linear system with time varying block currents as inputs and voltage variations at the supply connections as outputs (Figure 3).

We analyze this system using the field of linear system theory [12]. This system has impulse response function given by a matrix $H(t)$, whose element at row b and column n denotes the impulse response at node n due to current block b. Since the system is linear, the voltage response at node n due to any current waveform of block b is given as follows:

$$V_n(t) = \int_0^\infty i_b(t-\tau) \cdot h_{b,n}(\tau) \cdot d\tau \qquad \text{(EQ 1)}$$

currents of
logic blocks
$i_b(t)$

$H(t)$

voltage
variations at
logic blocks
$V_n(t)$

Figure 3. Linear system model of a power distribution network

Divide chip into
blocks

Compute block current
statistical parameters

Obtain impulse response
at every node

Compute mean, std dev
of voltage drop at each node

Figure 1. Flow diagram of proposed approach

where, $h_{b,n}(\tau)$ is the impulse response at node n due to the excitation at block b, $i_b(t)$ is the current waveform of block b and $V_n(t)$ is the voltage response at node n.

If the total number of blocks in the design is B, then the voltage response at node n due to all the current blocks acting together is the superposition of individual responses as shown below.

$$V_n(t) = \sum_{b=0}^{B-1} (\int_0^\infty i_b(t-\tau) \cdot h_{b,n}(\tau)d\tau) \qquad \text{(EQ 2)}$$

There are many ways of computing the impulse response $h_{b,n}(t)$ at a node n due to a block current i_b. In our approach, we first obtain the step response at node n by applying a unit step current at block b and simulating the grid. We then numerically differentiate the unit step response to obtain the impulse response. For typical grids, the unit pulse response dampens out quickly and we need to simulate the grid only for a small period of time. In the next section, we describe the representation of block currents, $i_b(t)$ and supply voltage variations $V_n(t)$ as stochastic processes.

2.2 Block currents and voltages as random processes

During the operation of a chip, the blocks consume currents that vary in time. The actual waveforms of these currents have rather complicated shapes. In order to completely analyze the behavior of a power supply network, it is required to simulate the power grid for all possible current waveform patterns, which is clearly not feasible. Usually, power supply networks are simulated using current waveforms for only a very limited period of time which does not guarantee that the voltage variation for all possible situations is analyzed. In this Section, we discuss how currents are modeled as random processes characterized by their statistical characteristics. This approach allows for very long waveforms to be analyzed, covering a large input vector space.

We represent each block current as a random process $i_b(t)$. This implies that the value of a block current at each time point is considered a random variable and the current over time is a random function of time. Note that the current values at different points in time are not independent. Their values depend on their history which means that we must account for possible correlations between current values over different time points. This random variable is described by a probability density function $p(i_b)$, which in the general case varies over time. Complete description of a random process requires specification of joint probability density function, $p(i_b(t_1), i_b(t_2),...)$ for all block currents at any time which very difficult to obtain and analyze. Our goal therefore is to compute the statistical characteristics of the voltage drop only, i.e. the mean and variance (standard deviation), which greatly simplifies the overall computation. Hence, we only need to determine the mean, auto-correlation function and cross-correlation functions of the block currents in order to compute the mean and standard deviation of the voltage response at any node in the power grid.

The mean and variance of a random process $X(t)$ are the expectation and the variance respectively of the random variable obtained by observing the process at some time t.

$$\mu_X(t) = E(X(t)) \qquad \text{(EQ 3)}$$

$$\sigma_X^2(t) = E((X(t)-\mu_X)^2) \qquad \text{(EQ 4)}$$

The auto-correlation function (or simply the correlation function) of a random process $X(t)$ is defined as the expectance of the product of two random variables obtained by observing the random process at times t_1 and t_2:

$$R_X(t_1, t_2) = E(X(t_1) \cdot X(t_2)) \qquad \text{(EQ 5)}$$

Figure 4. Variation of a block current with time

The cross-correlation function of two random processes $X(t)$ and $Y(t)$ is the expectance of the product of two random variables obtained by observing *two* processes at times t_1 and t_2, respectively:

$$R_{XY}(t_1, t_2) = E(X(t_1) \cdot Y(t_2)) \qquad \text{(EQ 6)}$$

We assume that each block current is a stationary random process. This implies that the statistical characteristics of block currents do not depend on any time shift or equivalently, the probability density functions of block currents do not depend on any time shift. This assumption is reasonable for currents and voltages because the only non-repeating part of current waveforms is chip initialization which is negligibly short as compared to the total time of operation of a chip. For stationary processes all the moments do not depend on time and the auto-correlation function depends only on time difference $\tau = t_2 - t_1$. Hence, the mean, auto-correlation and the cross-correlation functions are given as follows:

$$\mu_X = \mu_X(t) = E(X(t)) \qquad \text{(EQ 7)}$$

$$R_X(\tau) = E(X(t) \cdot X(t+\tau)) \qquad \text{(EQ 8)}$$

$$R_{XY}(\tau) = E(X(t) \cdot Y(t+\tau)) \qquad \text{(EQ 9)}$$

Another reasonable assumption which significantly simplifies the analysis is to assume the block currents to be ergodic processes. Ergodicity implies that averaging a random process over the *sample space* for each particular time (ensemble average) gives the same result as averaging it over the *time* of one implementation of the random process (time average). This assumption is justified by the fact that if a chip operates long enough it definitely exposes all possible operation modes and all its states. Therefore, investigating any single waveform long enough is sufficient for predicting the statistical behavior of a chip.

For stationary ergodic processes with an observation window $-T \le t \le T$, the mean and the auto-correlation functions can be expressed as stated below:

$$\mu_X = \lim_{T \to \infty} \frac{1}{2T} \cdot \int_{-T}^{T} X(t)d\tau \qquad \text{(EQ 10)}$$

$$R_X(\tau) = \lim_{T \to \infty} \frac{1}{2T} \cdot \int_{-T}^{T} X(t) \cdot X(t+\tau)dt \qquad \text{(EQ 11)}$$

and the cross-correlation of two random processes $X(t)$ and $Y(t)$ is given as:

$$R_{XY}(\tau) = \lim_{T \to \infty} \frac{1}{2T} \cdot \int_{-T}^{T} X(t) \cdot Y(t+\tau)dt \qquad \text{(EQ 12)}$$

The limits in these integrals indicate that these expressions are valid for a large value of T, i.e. the observation window for the random process $X(t)$ should be large. For our approach, this implies that the time interval of the block current waveforms should be sufficiently large.

Now we discretize EQ10 through EQ12 for practical computation of the statistical parameters of block currents. Figure 4 shows the

173

Figure 5. PDF (a) and auto-correlation (b) of a block current

variation in current for a short period of time for a block in an industrial microprocessor. The spikes correspond to periods of high activity in a clock cycle. As a first step, this waveform is divided into small time steps of width Δt such that there is no significant variation in block current within a time step. Then, we compute the mean of the block currents as follows:

$$\overline{I_b} = E(I_b) = \frac{1}{K}\sum_{k=0}^{K-1} I_{b_k} \qquad \text{(EQ 13)}$$

where K is the total number of time steps for which the waveform is observed. Similarly, the values of auto-correlation and cross-correlation functions in time step n are computed as follows:

$$R_{I_b}(n) = \frac{1}{K-n+1}\sum_{k=0}^{K-n} I_{b_k} \cdot I_{b_{k+n}} \qquad \text{(EQ 14)}$$

$$R_{I_i, I_j}(n) = \frac{1}{K-n+1}\sum_{k=0}^{K-n} I_{i_k} \cdot I_{j_{k+n}} \qquad -K \leq n \leq K \qquad \text{(EQ 15)}$$

where, I_{b_k} is the current of block b in time step k. Figure 5 shows the probability distribution and auto-correlation functions for the time varying block current shown in Figure 4. To compute the correlation functions for all possible values of n ($-K \leq n \leq K$) results in a complexity $O(K^2)$. However, we later show that we can restrict n to values less than the response time of the power grid without loss in accuracy, which is significantly smaller than K for large vector sequences and hence, the complexity is linear with K.

In the next subsection, we describe the computation of mean and standard deviation of voltage supply fluctuations, using the statistical parameters of block currents as derived in this section.

2.3 Linear system with stochastic excitations

An important property of a linear system is the fact that if we know the impulse response function and the expectance and auto-correlation functions of a signal at the input of linear system, we can compute the expectance and auto-correlation function at the output of the system. The output of a linear system, $Y(t)$ can be expressed as the convolution of the impulse response, $h(t)$ with its input random process $X(t)$ as given in EQ1. Using stationarity and ergodicity assumptions, we can average the expression over time to get the expectance of the signal at the output of the linear system.

$$\overline{Y} = E(Y(t)) = E\left(\int_0^{\infty} X(t-\tau) \cdot h(\tau) \cdot d\tau\right) \qquad \text{(EQ 16)}$$

$$= \int_0^{\infty} E(X(t-\tau)) \cdot h(\tau)d\tau = \overline{X} \cdot \int_0^{\infty} h(\tau)d\tau$$

This implies that the mean of the random process at the output of a linear time-invariant system in response to random process $X(t)$ is equal to the mean of $X(t)$ multiplied by the dc response of the system. This property is useful in power grid analysis because we can obtain the exact value of mean of the supply voltage variations without any complex computation if we know the mean of the input block currents and the dc response of the network. We simulate the

power grid in order to obtain the unit step response of the system and observe the response till it dies out and attains a steady state value. This steady state value is the dc response of the system and can be used in computing the mean of the supply voltage fluctuations.

Similarly, we can compute the second moment of the random signal at the output of the linear system using the auto-correlation function of the input signal as follows:

$$\overline{Y^2} = E(Y^2(t)) \qquad \text{(EQ 17)}$$

$$= E\left(\int_0^{\infty} X(t-\tau) \cdot h(\tau)d\tau \cdot \int_0^{\infty} X(t-\tau) \cdot h(\tau)d\tau\right)$$

$$= \int_0^{\infty}\int_0^{\infty} R_X(\tau_1 - \tau_2) \cdot h(\tau_1) \cdot h(\tau_2)d\tau_1 d\tau_2$$

Once we determine the second moment and the mean of the output process, we can obtain its variance using this well known formula:

$$\sigma_Y^2 = \overline{Y^2} - (\overline{Y})^2 \qquad \text{(EQ 18)}$$

EQ16 through EQ18 can be used to compute the mean and variance of the voltage drop at any node n. Let the number of blocks in the circuit be B. The voltage response at node n due to all the current blocks acting together is the superposition of individual responses and is given by EQ2. Thus the mean of the voltage response at node n is given by,

$$\overline{V_n} = \sum_{b=0}^{B-1}\left(E(I_b)\int_0^{\infty} h_{b,n}(\tau)d\tau\right) \qquad \text{(EQ 19)}$$

For practical implementation, we can discretize this expression as follows:

$$\overline{V_n} = \sum_{b=0}^{B-1} \overline{I_{b_k}} \sum_{i=0}^{M-1} h_{b,n}(i), \qquad \text{(EQ 20)}$$

where M is the total number of time steps after which the impulse response $h_{b,n}$ remains zero.

Now, we discuss the computation of variance of the voltage drop for the cases when the block currents are independent and when they are correlated. We also discuss the run time complexity in both cases.

Block currents considered independent

If the block currents are assumed to be independent, then the second moment (variance) of the voltage response due to all the blocks acting together, depends only on the auto-correlation functions of the individual block currents and is equal to the summation of the second moments (variances) of voltage responses due to individual block currents.

$$\overline{V_n^2} = \sum_{b=0}^{B-1}\int_0^{\infty}\int_0^{\infty} R_{I_b}(\tau_1 - \tau_2) \cdot h_{b,n}(\tau_1) \cdot h_{b,n}(\tau_2)d\tau_1 d\tau_2 \qquad \text{(EQ 21)}$$

We again discretize the above expression as follows:

$$\overline{V_n^2} = \sum_{b=0}^{B-1}\sum_{i=0}^{M-1}\sum_{j=0}^{M-1} R_{I_b}(i-j) \cdot h_{b,n}(i) \cdot h_{b,n}(j) \qquad \text{(EQ 22)}$$

where M is the number of time steps until the voltage response to the impulse current becomes zero or the step response reaches its steady state, K is the number of time steps used in computing the auto-correlation function of $I_b(t)$ (the number of time steps over which the block currents are observed) and B is the total number of blocks. It is important to note that the values of M and B are much less than K. In EQ22, we need to compute the auto-correlation function only for time steps n before which the impulse response has died down and is zero i.e. $0 \leq n \leq M$. Thus any temporal correlation in current waveforms, shifted in time beyond the time period taken by the impulse response to reach its steady state can be ignored. Hence, the com-

plexity of computing auto-correlation functions for B blocks (EQ14) is $O(KMB)$ which is linear in the length of a block current waveform K and, the complexity of EQ22 is $O(M^2B)$ which is independent of block current waveform length. Thus, it is possible to run millions of vectors and use the statistical information of corresponding current waveforms to obtain the statistical parameters of voltage drop.

Block currents considered to be correlated

If the block currents are correlated, the second moment and the variance of the voltage response due to all the blocks acting together depends not only on the auto-correlation functions, but also on the cross-correlations between different block currents. In general, when each block current is correlated both spatially and temporally to other block currents, the second moment of the voltage response at any node n is given by,

$$\overline{V_n^2} = \int_0^\infty \int_0^\infty \sum_{j=0}^{B-1} \sum_{k=0}^{B-1} R_{I_j, I_k}(\tau_1 - \tau_2) \cdot h_{j,n}(\tau_1) \cdot h_{k,n}(\tau_2) d\tau_1 d\tau_2 \quad \text{(EQ 23)}$$

or,

$$\overline{V_n^2} = \sum_{i=0}^{M-1} \sum_{j=0}^{M-1} \sum_{k=0}^{B-1} \sum_{l=0}^{B-1} R_{I_k, I_l}(i-j) \cdot h_{k,n}(i) \cdot h_{l,n}(j) \quad \text{(EQ 24)}$$

where R_{I_j, I_k} is the cross-correlation function between block currents $i_j(t)$ and $i_k(t)$. In general, not all the blocks are strongly correlated with each other and cross-correlation of only those blocks which have significant correlations can be considered with limited impact on the accuracy. If each block b in the design is correlated with C other blocks, then only BC correlations need to be considered instead of a total possible B^2, where $C \ll B$. Again, although the complexity of computing the cross-correlation function between two waveforms of length K is quadratic in K, we only need values of $R_{I_j, I_k}(n)$ such that $n < M$. Thus, the complexity of computing all the auto-correlations for all the blocks and cross-correlations for the possible BC combinations is $O(KMBC)$ while the complexity of computing the second moment from the correlation functions is $O(M^2BC)$. The overall complexity is therefore still linear with K.

Till now, we considered both spatial and temporal correlations in block currents. Typically in a design, blocks will have much larger correlation in space as compared to the correlation in time. If the block current correlations in time are ignored and each block is assumed to be spatially correlated to C other blocks, then as a special case, EQ24 reduces to:

$$\overline{V_n^2} = \sum_{i=0}^{M-1} \sum_{j=0}^{M-1} \sum_{k=0}^{B-1} \sum_{l=0}^{B-1} R_{I_k, I_l}(0) \cdot h_{k,n}(i) \cdot h_{l,n}(j) \quad \text{(EQ 25)}$$

The run time of the analysis improves in this case because now the complexity of computing cross-correlation coefficient for the block currents is $O(KBC)$ instead of a previous $O(KMBC)$ for computing the auto-correlation and cross-correlation functions.

2.4 Voltage drop probability distribution

The currents of a block vary in accordance to the input vectors applied to the circuit and generally show a large variation. The shape of the probability distribution of these block currents is not at all fixed and may vary depending on the size of the blocks, the functionality of the block and various other factors. Thus we cannot definitely state anything about the shape of the probability distribution of the current blocks.

However, the overall voltage drop is the sum of the voltage drops due to all the current sources acting together. The blocks have been formed such that there is minimum correlation between the current of the blocks. Thus, assuming that the number of the blocks is relatively large and most of the blocks are independent from each other,


Figure 6. Probability distribution of the overall drop

central limit central limit theorem can be applied which states that if $X_1, X_2...X_n$ are independent random variables, then the random variable formed by summing the variables $Y=X_1+X_2+...+X_n$ has a distribution which approaches a normal distribution for large values of n. Thus, if the design is divided into large number of blocks with most of the blocks being independent, the overall drop at any node can be approximated as having a Gaussian distribution function.

This means that while the proposed approach computes only the mean and variance of the voltage drop at a particular node, the voltage distribution can be approximated as a normal distribution. We show in Section 3 that the normal distribution closely fits the exact distribution. Using a Gaussian approximation of the voltage distribution, we can compute any desired confidence interval of the voltage drop, such as the 95% or 99% voltage confidence points.

3 RESULTS

The proposed approach for determining the mean and standard deviation of voltage drop was implemented and tested on a number of grids of different sizes for both flip-chip and wire bond package models. The block currents were generated by simulating an industrial microprocessor for thousands of cycles. The statistical parameters of the voltage drop were computed both assuming the block currents to be independent as well as by taking block current correlations (spatial and temporal) into account. We discussed in Section 2.4 that the probability distribution function of the voltage response was expected to be close to Gaussian if the number of blocks in the chip is large and most of the blocks are uncorrelated. We found that this assumption held for all the grids that we tested. Figure 6 shows the exact probability distribution function (PDF) of the voltage drop at a node in a grid and its Gaussian fit, which forms a close approximation. The exact PDF was obtained by simulating the grid with SPICE for the entire length of the current vectors and statistically analyzing the resulting voltage drops.

Table 1 shows the mean and standard deviation computed using the proposed approach, assuming block current independence. The 95% confident points are also given, obtained using a Gaussian fit of

Table 1. Mean, std dev and 95% interval of the voltage drops

Grid	Grid type	# of nodes	# of blocks	Mean (mV)	Std dev (mV)	95% conf int. (mV)	Maxdrop (mV)
Grid1	WB	3772	12	50.3	66.1	179.9,-79.3	327.6
Grid2	FC	3772	12	39.6	49.4	136.4,-57.2	281.3
Grid3	WB	7712	20	56.9	67.3	188.8,-75.0	297.1
Grid4	FC	7712	20	48.1	51.5	149.0,-52.8	169.1
Grid5	WB	17037	30	39.0	47.6	132.3,-54.3	231.7
Grid6	FC	17037	30	33.5	36.5	105.0,-38.0	186.0
Grid7	WB	32897	40	103.4	111.2	321.4,-114.6	376.9
Grid8	FC	32897	40	86.6	94.5	271.8,-98.6	292.0
Grid9	FC	157180	40	88.2	113.6	310.9,-134.5	-

Table 2. Comparison with HSPICE

Grid	Proposed Approach		HSPICE		% Error in Std Dev	Run Time
	Mean (mV)	Std Dev (mV)	Mean (mV)	Std Dev (mV)		
Grid-1	50.3	66.1	50.3	78.2	15.47%	48s
Grid-2	39.6	49.4	39.6	57.8	14.53%	48s
Grid-3	56.9	67.3	56.9	79.0	14.81%	1m40s
Grid-4	48.1	51.5	48.1	60.1	14.31%	1m29s
Grid-5	39.0	47.6	39.0	54.6	12.82%	2m37s
Grid-6	33.5	36.5	33.5	40.7	10.32%	2m41s
Grid-7	103.4	111.2	103.4	123.8	10.18%	3m14s
Grid-8	86.6	94.5	86.6	103.8	8.96%	3m25s
Grid-9	88.2	113.6	-	-	-	3m31s

the voltage drop PDF. The maximum voltage drop, as observed during SPICE simulation of the complete set of vectors is also shown. Grid-1 through Grid-8 are different size grids in 9 layers of metal, generated using pitches and widths of an industrial microprocessor design with a PEEC based extraction tool. The metal lines in the grid are modelled as an RLC network, which consists of the self inductance, capacitance and resistance of the wires. Mutual inductance of on-chip interconnects was ignored for the sake of simplicity. Grid-9 is the grid of an industrial processor, extracted using a commercial extraction tool and consists of over 1 million elements. I/O pads were modeled using an industrial package model. It is interesting to observe that the maximum voltage drop is between 2.17 and 4.89 times the standard deviation higher than the mean. This demonstrates that the occurrence of the worst-case drop is extremely rare.

Table 2 compares the mean and standard deviation of the voltage drops computed with the proposed approach against those obtained from SPICE simulation. Grid-9 could not be simulated in SPICE for the whole current waveforms because of its size. The mean of the voltage drop does not depend on correlations between block currents and is only affected by the error due to the discretization of current waveforms. Hence, the mean computed with the proposed method was found to be within 1% of those obtained with SPICE simulation. There is some error in the standard deviation of the voltage drops since the block currents are assumed to be independent. As a general trend, the error reduced with increasing number of nodes in the grid, which is a favorable characteristic since most industrial grids are very large. This results from the fact that with increasing size of the grid, the number of C4 bumps/wire bond pads in the grids increase, which attenuates the effect of current correlations on the voltage drop. Also, wire bond chips had a greater error due to correlation as compared to chips with flip-chip pads. The last column shows the run time for finding the mean and standard deviation at a single node. This run time includes the time to compute the mean, the auto-correlation functions of the currents and computation of the standard

Table 3. Effect of correlations on accuracy and runtime

Degree of Correlation	No. of Cross Correlations	% Error in Std Dev	Run Time
None	0	10.32%	2m41s
2	60	8.03%	14m53s
4	120	6.64%	24m48s
6	180	5.89%	36m36s
8	240	5.40%	51m03s
10	300	4.42%	1h9m24s
15	450	2.84%	1h37m41s

deviation. Note that the required time for computing the voltage statistics for additional nodes in the circuit would be substantially less since the auto-correlation and the cross-correlation functions of block currents need to be computed only once.

In Table 2, we show the effect of incorporating spatial and temporal correlations between the block currents on the run time and accuracy. We assume blocks that are close together to have a large correlation among them and assume distant blocks to be independent. Column 1 indicates the degree of correlation between adjacent blocks. A degree of correlation n implies that each block current is correlated with its n neighbors. Column 2 gives the total number of cross-correlations taken into account. We compute the cross-correlation functions of a particular block with all its neighboring blocks and choose n blocks with the largest cross-correlation coefficient. We observe that the accuracy of the computation can be improved with reasonable additional run-time.

4 CONCLUSION

In this paper, we presented a new approach for computing the statistical parameters of supply voltage fluctuations with variability in block currents. The analysis considers both IR-drop and LdI/dt drop in a power supply network and takes into account both spatial and temporal correlations in block currents. The analysis was implemented and tested on a number of grids, including the power grid of an industrial processor and we demonstrate the effectiveness of the proposed approach. The approach provides useful information about which parts of the grid which are most likely to fail and should be given priority during grid correction.

5 ACKNOWLEDGEMENTS

6 REFERENCES

[1] G. Steele, D. Overhauser, S. Rochel and Z, Hussain, "Full-chip verification methods for DSM power distribution systems," in *Proc. DAC*, pp. 744-749, 1998.

[2] S. Taylor, "The challenge of designing global systems in UDSM CMOS," in *Proc CICC*, pp.429-435, 1999

[3] A. Chandrakasan, W. J. Bowhill and F. Fox, *Design of high performance microprocessor circuits*, NY: IEEE Press, 2001.

[4] S. R. Nassif and J. N. Kozhaya, "Fast power grid simulation," in *Proc. DAC*, pp. 156-161, 2000.

[5] M. Zhao, R. V. Panda, S. S. Sapatnekar and D. Blaauw, "Hierarchical analysis of power distribution networks," in *IEEE Trans. on CAD*, pp. 159-168, 2002.

[6] R. Panda, D. Blaauw, R. Chaudhry, V. Zolotov, B. Young and R. Ramaraju, "Model and analysis for combined package and on-chip power grid simulation," in *Proc. ISLPED*, pp. 179-184, 2000.

[7] A. Dharchoudhury, R. Panda, D. Blaauw and R. Vaidyanathan, "Design and Analysis of Power Distribution Networks in PowerPC microprocessors," in *Proc. DAC,* 1998.

[8] A. Krstic and K. Cheng, "Vector generation for maximum instantaneous current through supply lines for CMOS circuits," in *Proc. DAC*, pp. 383-388, 1997.

[9] Y. M. Jiang, T. Young and K. Cheng, "VIP - an input pattern generator for identifying critical voltage drop for deep submicron designs," in *Proc. ISLPED*, pp. 156-161, 1999.

[10] H. Kriplani, F. Najm, I. Hajj, "Pattern independent maximum current estimation in power and ground buses of CMOS VLSI circuits," in *IEEE Trans. on CAD*, vol. 14. no. 8, pp. 998-1012, 1995.

[11] G. Bai, S. Bobba and I.N. Hajj, "RC power bus maximum voltage drop in digital VLSI circuits," in *Proc. ISQED*, pp. 205-210, 2001.

[12] George R. Cooper, Clare D. McGillem, *Probabilistic Methods of Signal and System Analysis*, Oxford Series, 1998.