Worst-Case Performance Prediction Under Supply Voltage and Temperature Variation

Chung-Kuan Cheng, Amirali Shayan
CSE Department
UC San Diego
La Jolla, CA
ckchen@ucsd.edu, amirali@ucsd.edu

Andrew B. Kahng
CSE and ECE Departments
UC San Diego
La Jolla, CA
abk@ucsd.edu

Kambiz Samadi
ECE Department
UC San Diego
La Jolla, CA
ksamadi@ucsd.edu

ABSTRACT

The power delivery network (PDN) is a major consumer of interconnect resources in deep-submicron designs (i.e., more than 30% of the entire routing area) [18]. Hence, efficient early-stage PDN optimization enables the designers to ensure a desired power-performance envelope. On the other hand as technology scales, gate delays become more sensitive to power supply variation. In addition, emerging 3D designs are more prone to supply voltage and temperature variation due to increased power density. In this paper, we develop accurate inverter cell delay and output slew models under supply voltage and temperature variation. Our models are within 6% of SPICE simulations on average. We use our single-cell delay and output slew models to estimate the delay of a path (i.e., an inverter chain, etc.). We also present a methodology to find the worst-case input configuration (i.e., input slew, output load, cell size, noise magnitude, noise slew, noise offset and temperature) that causes the delay of the given path is maximized. We believe that our models can efficiently drive accurate worst-case performance-driven PDN optimization.

Categories and Subject Descriptors

B.7 [**INTEGRATED CIRCUITS**]: Performance Analysis and Design Aids

General Terms

Design, Performance

Keywords

Supply voltage noise, temperature variation, nonparametric regression modeling, worst-case delay variation

1. INTRODUCTION

In sub-65nm designs, power/ground voltage level fluctuations (PG noise) has become a primary concern for power integrity as circuit timing becomes more susceptible to supply voltage noise. Thus, designers must take into consideration the impact of supply voltage noise to ensure successful chip design [14]. Rising supply voltage variation has become a challenge for power distribution system (PDS) verification. Typically, PDS verification is based on simulation; however, all possible current waveforms and load circuits are not known early in the design cycle. Hence, it is important to develop methods of accurately predicting worst-case supply voltage noise to ensure that the design timing is met.

Existing works [8, 19, 20] on supply voltage noise and its implications on power distribution network (PDN) optimization or PDS verification are oblivious to the timing impacts of supply voltage noise. In this work, we develop early-stage closed-form performance models under supply voltage and temperature variations that aid designers to assess the impact of their PDN design choices on the performance of the design. Timing degradation due to PG noise is often estimated by considering voltage drops through static IR-drop analysis. However, these analyses fail to capture the dynamic behavior of the supply voltage noise.

On the other hand, temperature variation affects transistor characteristics including threshold voltage, drive current, drive resistance, and off-current. Hence, it is important to accurately model the impact of temperature on circuit performance. Exiting literature [1, 6] propose closed-form expressions that consider the impact of temperature on cell delay; however, in this work we consider the combined effect of supply voltage and temperature variation on circuit performance.

In addition, emerging 3D designs are more prone to supply voltage noise due to increase in power/current demand and variations among tiers. Compensation of the supply voltage variation requires a fair amount of the silicon real estate (e.g., decoupling capacitance allocation, etc.), routing resources, and increased packaging cost. Increased power density in 3D designs also requires close attention to the impact of temperature on circuit performance. Hence, to guarantee a given performance envelope, designers need to characterize the impact of supply voltage and temperature variation on circuit timing. Furthermore, [22] points out to a number of problems caused by dynamic effects of supply voltage noise. These effects include (1) change in maximum frequency of a critical path, (2) degradation of the clock network performance, etc. Thus, designers must consider the dynamic effect of supply voltage noise early in the design cycle.

Finally, the PDN is a major consumer of resources (e.g.,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
SLIP'10, June 13, 2010, Anaheim, California, USA.
Copyright 2010 ACM 978-1-4503-0037-7/10/06 ...$10.00.

more than 30% of the entire routing area) in wire-limited deep-submicron designs [18]. Conventionally, the PDN is designed to satisfy power integrity constraints, but without understanding the true implications of supply noise on delay, correct optimization of PDN is impossible. To close this gap, our present work gives a methodology for closed-form modeling of the delay impact of supply voltage noise (characterized by noise slew, offset, and magnitude). We believe our models can efficiently drive accurate worst-case performance-driven PDN optimization, as shown in Figure 1.

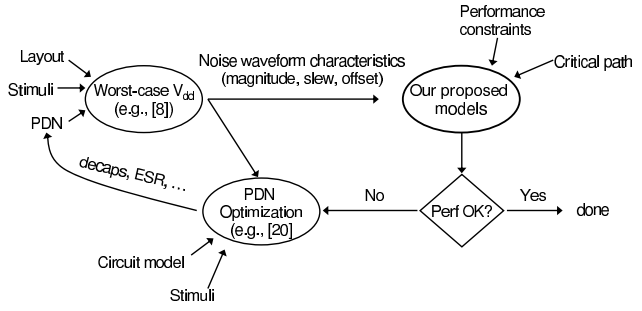

Figure 1: Accurate worst-case performance-driven power distribution network optimization flow.

In this paper, we propose a new modeling paradigm in which we use *machine learning-based nonparametric* regression techniques to develop accurate early-stage performance models under dynamic supply voltage and temperature variations. The contributions of our work are as follows.

- We propose a framework for gate delay modeling under supply voltage and temperature variations, using machine learning-based nonparametric regression methods. Further, we introduce a reproducible flow to aid *automatic* generation of accurate performance estimation models (e.g., using generic critical paths).
- We develop early-stage performance models using our basic gate delay models, to enable worst-case performance prediction that can efficiently drive PDN optimization.
- We validate our models against SPICE simulations using 65nm foundry SPICE models.

The remainder of this paper is organized as follows. In Section 2 we review and contrast prior related work. Section 3 describes our implementation flow and the scope of our study. Section 4 describes our modeling methodology using machine learning-based regression techniques. In Section 5 we describe the impact of different parameters on gate delay and output slew, and present our proposed worst-case performance model. In Section 6 we validate our proposed models against SPICE simulations. Finally, Section 7 concludes the paper.

2. RELATED WORK

Gate delay models under supply voltage noise can be classified as (1) static or (2) dynamic; with the former type, the dynamic behavior of the noise waveform is ignored. The majority of the existing literature focuses on the former type [3, 7, 11, 15, 17]. Hashimoto *et al.* [7] propose to replace supply voltage noise with an equivalent power/ground voltage. However, this method assigns static voltage value (time-invariant) during the static timing analysis (STA), and cannot appropriately capture the dynamic behavior of the noise waveform. Martorell *et al.* [11] present a probabilistic approach to estimate supply voltage noise bound given performance criteria. However, they assume that all gates in a combinational path have the same supply voltage value; this assumption is incorrect due to the presence of dynamic supply voltage noise.

Chen *et al.* [3] propose closed-form equations to estimate the change in delay of buffers in the presence of supply voltage noise. However, the authors do not consider specific noise waveform characteristics (magnitude, offset, and slew) in their analysis. In another effort, Weng *et al.* [17] propose a methodology to improve the accuracy of gate delay calculation under supply voltage noise by taking into account (time-varying) IR drop waveforms. To capture the dynamic impact of supply voltage noise, the authors of [17] discretize the noise waveform and assign an equivalent DC value across different time intervals. The DC values are calculated as the average supply voltage values over the entire interval. This method still does not capture the 'true' dynamic behavior of the supply noise waveform. To assess the impact of supply voltage noise on circuit performance, [15] suggests that using average supply voltage, rather than dynamic behavior, can be well-correlated with measurements; however, the authors fail to demonstrate the limitations of timing analysis using static IR-drop analysis as noted in [14].

Recently, Okumura *et al.* [14] have proposed a gate delay calculation approach which considers the dynamic behavior of the supply voltage noise by considering noise waveform slew and magnitude. However, in their characterization setup they do not allow all the relevant parameters (i.e., input slew, noise slew, noise magnitude, etc.) to change simultaneously; this limits the applicability of their proposed model. In our present work, we develop new gate delay and output slew models under supply voltage and temperature variations, where all the relevant parameters can interact with one another.

3. IMPLEMENTATION FLOW

Figure 2 shows our implementation flow, which beings with SPICE simulations using foundry SPICE models and extracted or CDL SPICE netlists for each gate type. We measure the 50% delay and output slew of each gate with respect to a number of different parameters. In our experiments we have three main axes: (1) cell delay parameters, (2) supply voltage noise parameters, and (3) temperature. These parameters, and the values that they take on in our experiments, are explained below. Cell delay parameters include (1) input slew $slew_{in}$, (2) output load $load_{out}$, and (3) cell size $cell_{size}$. For supply voltage we use 0.9V as the nominal value, with noise waveform superimposed on it.

Supply voltage noise parameters include (1) noise amplitude amp_{noise}, (2) noise slew $slew_{noise}$, and (3) noise offset $offset_{noise}$. Noise offset denotes the noise transition time with respect to that of the input signal transition. Finally, temperature denotes the operating temperature of the transistors. In our studies, we use two different cells (1) inverter, and (2) 2-input NAND to show the applicability of our modeling approach. For worst-case performance model we implement our basic cell delay and output slew models in C++. Using our basic delay and output slew models we construct path delay models with arbitrary number of stages and a mix of different cells. We run a total of 30720 SPICE simulations and gather delay and output slew values corresponding to different parameters (cf. Table 1).

We use *Synopsys HSPICE v.Y-2006.03* [24] for SPICE simulations using 65nm foundry SPICE models and netlists. We perform our experiment using typical corner and normal-

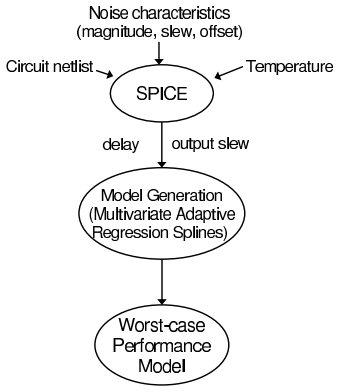

Figure 2: Implementation flow.

V_{th} (NVT) transistors. We also use *MARS3.0* [23] to implement nonparametric regression techniques.

Table 1: List of parameters used in our studies.

Parameter	Values
$slew_{in}$	{0.00056, 0.00112, 0.0392, 0.1728,0.56, 0.7088}ns
$load_{out}$	{0.0009, 0.0049, 0.0208, 0.0842}pF
$cell_{size}$	INV: {1, 4, 8, 20} ND2D: {1, 2, 4, 8}
amp_{noise}	{0, 0.054, 0.144, 0.27}V
$slew_{noise}$	{0.01, 0.04, 0.07, 0.09}ns
$offset_{noise}$	{-0.15, -0.05, 0, 0.05, 0.15}ns
$temp$	{-40, 25, 80, 125}°C

4. MODELING METHODOLOGY

4.1 Modeling Flow

Previous delay estimation techniques do not consider dynamic impact of supply voltage noise on cell delay [7, 11, 15]. By contrast, we propose to pursue a different modeling paradigm in which we use *machine learning-based nonparametric regression techniques* to capture the dynamic impact of supply voltage noise on cell delay. To illustrate the basic idea, consider the following baseline model generation flow:

- We begin with a parameterized SPICE netlist for a given inverter cell. We refer to this as a *configurable* inverter SPICE specification, which will be used to generate the representative inverter cell delay under different cell and supply voltage noise parameters. For example, a given SPICE simulation setup can be configured with respect to (1) input slew, (2) output load, (3) inverter size, (4) supply voltage noise magnitude, (5) supply voltage noise width (i.e., frequency), (6) voltage noise offset (i.e., with respect to the input transition), and (7) temperature.

- Using a small subset of selected configurations for *training*, we run through each configuration in this training set through SPICE simulations, to obtain an accurate cell delay for each instance.

- Finally, we apply machine learning-based nonparametric techniques on the training set of delay to derive the corresponding cell delay estimation models.

In general, the modeling problem aims to approximate a function of several to many variables using only the dependent variable space. This generic formulation has applications in many disciplines. The goal is to model the dependence of a target variable y on several predictor variables $x_1, \cdots, x_n$ given R realizations $\{y_i, x_{1i}, \cdots, x_{ni}\}_1^R$. The system that generates the data is presumed to be described by

$$y = f(x_1, \cdots, x_n) + \epsilon \tag{1}$$

over some domain $(x_1, \cdots, x_n) \in \mathcal{D} \subset \mathcal{R}^n$ containing the data [4]. Function f captures the joint predictive relationship of y on $x_1, \cdots, x_n$, and the additive stochastic noise component ϵ usually reflects the dependence of y on quantities other than $x_1, \cdots, x_n$ that are neither controlled nor observed. Hence, the aim of the regression analysis is to construct a function $\hat{f}(x_1, \cdots, x_n)$ that can accurately approximate $f(x_1, \cdots, x_n)$ over the domain $\mathcal{D}$ of interest. There are two main regression analysis methods: (1) global parametric, and (2) nonparametric. The former approach has limited flexibility, and can produce accurate approximations only if the assumed underlying function $\hat{f}$ is close to f. In the latter approach, $\hat{f}$ does not take a predetermined form, but is constructed according to information derived from the data. Multivariate adaptive regression splines (MARS) is a nonparametric regression technique which is an extension of linear models that automatically models nonlinearities and interactions, and is used in our methodology. In this paper, we use MARS-based approach to model the dynamic impact of supply voltage noise on cell delay.

4.2 Multivariate Adaptive Regression Splines

Given different cell and supply voltage noise parameters $\mathcal{X}$, we apply MARS to construct cell delay model, $d_{cell} = \hat{f}(x_1, \cdots, x_n)$. Variables $x_1, \cdots, x_n$ denote cell and supply voltage noise parameters. The general MARS model can be represented as [21]

$$\hat{y} = c_0 + \sum_{i=1}^{I} c_i \prod_{j=1}^{J} b_{ij}(x_{ij}) \tag{2}$$

where $\hat{y}$ is the target variable (i.e., inverter delay and output slew in our problem), c_0 is a constant, c_i are fitting coefficients, and $b_{ij}(x_{ij})$ is the truncated power basis function[1] with x_{ij} being the microarchitectural parameter used in the i^{th} term of the j^{th} product. I is the number of basis functions and J limits the order of interactions. In our experiments we set the number of basis functions to 100 and the order of interactions to 6, i.e., every parameter can interact with all the other parameters. The basis functions $b_{ij}(x_{ij})$ are defined as

$$b_{ij}^{-}(x^{\mu arch} - t_{ij}) = [-(x^{\mu arch} - t_{ij})]_+^q \tag{3}$$
$$= \begin{cases} (t_{ij} - x^{\mu arch})^q & x^{\mu arch} < t_{ij} \\ 0 & \text{otherwise} \end{cases}$$
$$b_{ij}^{-}(x^{\mu arch} - t_{ij}) = [+(x^{\mu arch} - t_{ij})]_+^q \tag{4}$$
$$= \begin{cases} (x^{\mu arch} - t_{ij})^q & x > t_{ij} \\ 0 & \text{otherwise} \end{cases}$$

where q (≥ 0) is the power to which the splines are raised to

[1] Each basis function can be a constant, a hinge function that is of form $max(0, c - x)$ or $max(0, x - c)$, or a product of two or more hinge functions.

adjust the degree of $\hat{y}$ smoothness, and t_{ij} is called a knot. When $q = 0$ simple linear splines are applied.

The optimal MARS model is built in two passes. (1) Forward pass: MARS starts with just an intercept, and then repeatedly adds basis function in pairs to the model. Total number of basis functions is an input to the modeling. Backward pass: during the forward pass MARS usually builds an overfit model; to build a model with better generalization ability, the backward pass prunes the model using a generalized cross-validation (GCV) scheme

$$GCV(K) = \frac{1}{n} \frac{\sum_{k=1}^{n}(y_k - \hat{y})^2}{[1 - \frac{C(M)}{n}]^2} \quad (5)$$

where n is the number of observations in the data set, K is the number of non-constant terms, and $C(M)$ is a complexity penalty function to avoid overfitting.

5. ACCURATE CELL DELAY MODELING

In this section, we discuss the impact of supply voltage noise and temperature variation on cell delay, and note that delay modeling under supply voltage and temperature variation is a nontrivial task. We show an example of our proposed delay and output slew models derived from machine learning-based nonparametric regression techniques. We also propose a methodology to find the worst-case input configuration that maximizes the delay of a given path.

5.1 Cell Delay and Output Slew Models

In the existing literature [7, 11], supply voltage variation is assumed to be constant (time-invariant). When the supply voltage varies slowly with respect to the clock period, this is reasonable. This assumption enables to predict the timing impact of the supply voltage noise: the worst-case delay corresponds to the worst-case noise that can occur when the target cell is switching. In other words, when the supply voltage varies slowly, the delay degradation is proportional to the peak of the noise [15]. However, to better capture the impact of time-varying supply voltage noise we must consider the noise waveform characteristics including (1) noise amplitude, (2) noise slew, and (3) noise offset. Figure 3 shows the impact of noise slew on cell inverter delay. We observe that noise slew affects cell delay only when it is comparable to input slew. Hence, we must take into consideration the specific noise waveform characteristics to ensure more accurate delay modeling.

Existing PDN optimization frameworks [19, 20] use *fluctuation area*, i.e., the area under the noise waveform, as the metric to represent the supply voltage noise. However, it is easy to see that such an approach can incur significant error in the delay estimation. Consider two scenarios: (1) $slew_{noise}$=0.2ns, amp_{noise}=0.2V and (2) $slew_{noise}$=0.4ns, amp_{noise}=0.1V. Using a triangular waveform to represent the supply noise, the two scenarios have different noise waveforms, yet have similar areas under the noise curve. When we evaluate gate delay under each of these scenarios, we observe 22% difference. (In this evaluation, we use a single inverter, with other parameters values being $slew_{in}$=0.4ns, $load_{out}$=0.002pF, $cell_{size}$=1X, $offset_{noise}$=0ns, and $temp$=25°C.) We conclude that to accurately model the impact of supply voltage noise on cell delay, we must consider both noise slew and noise magnitude parameters, and not simply the area under the noise waveform.

The other important supply voltage noise characteristic is *noise offset*, which denotes the time of the voltage noise transition relative to the time of the input signal transition.

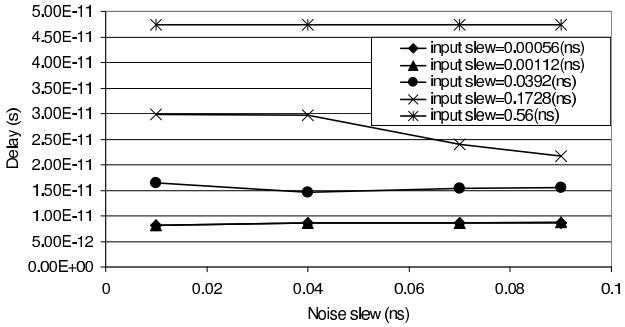

Figure 3: Delay of an inverter cell versus noise slew, for different input slew values.

We expect that as long as the supply voltage noise waveform is outside of the input signal transition window, it should not have any impact on cell delay. However, when the noise waveform overlaps with the input signal transition, there will be an effect on cell delay. Figure 4 shows the impact of noise offset on cell delay. In our experiment, input slew and noise slew are 0.09ns and 0.1ns, respectively. In our delay model, we explicitly consider noise offset as an input to the model.

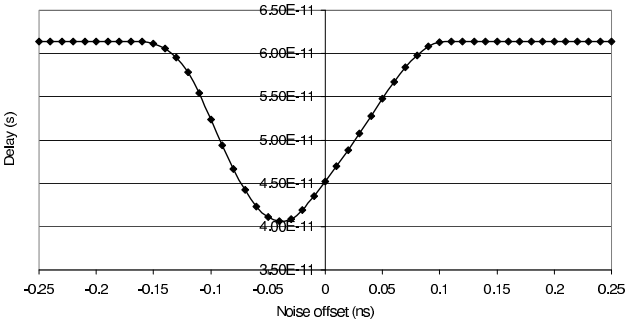

Figure 4: Impact of supply voltage noise offset on cell delay.

In addition, cell characteristics are influenced by temperature. Temperature impacts cell delay through voltage threshold, mobility, etc. parameters [6]. For example, as temperature decreases, both threshold voltage and mobility increase; the latter causes increased saturation current. However, the impact of temperature on cell delay depends on the gate voltage. The gate voltage at which the temperature shifts of threshold voltage and mobility exactly compensate each other's effects on delay is typically called zero-temperature-coefficient (ZTC) [10]. Hence, cell delay can increase or decrease with the increase in temperature. These complex relationships between cell delay and the aforementioned parameters make delay modeling a nontrivial task.

Finally, since our gate delay model depends on input slew, we must also model output slew of the previous stage of the critical path. Given the above discussion, we note that approximating CMOS gate delay is a nontrivial task with non-obvious implications, as seen from Figure 3. This has motivated us to explore machine learning-based nonparametric regression techniques to develop accurate cell delay and output slew models. Figure 5 illustrates the form of resulting inverter delay and output slew models using 65nm foundry SPICE models.[2]

[2]Note that our methodology can be straightforwardly ap-

Delay Model
$b_1 = \max(0, load_{out} - 0.0208)$;
$b_2 = \max(0, 0.0208 - load_{out})$; $\cdots$
$b_{98} = \max(0, offset_{noise} - 0.05) \times b_{92}$;
$b_{100} = \max(0, offset_{noise} + 2.4e\text{-}12) \times b_{37}$;
$d_{cell} = 1.018e-11 + 7.353e-10 \times b_1 - 5.890e-10 \times b_2$
$- 2.172e-11 \times b_3 + \cdots - 1.708e-7 \times b_{96} +$
$2.431e-7 \times b_{98} - 3.031e-8 \times b_{100}$
Output Slew Model
$b_1 = \max(0, load_{out} - 0.0009)$;
$b_2 = \max(0, cell_{size} - 4) \times b_1$; $\cdots$
$b_{99} = \max(0, 0.05 - slew_{noise}) \times b_{55}$;
$b_{100} = \max(0, offset_{noise} + 0.15) \times b_{94}$;
$slew_{out} = 1.227e-11 + 1.529 \times b_1 - 2.051e-10 \times b_2$
$+ 2.050e-9 \times b_3 + \cdots - 1.081e-8 \times b_{98}$
$- 4.327e-9 \times b_{99} - 7.422e-9 \times b_{100}$

Figure 5: Sample inverter delay and output slew models in 65nm.

5.2 Worst-case Performance Model

In this subsection we formalize the problem of finding the worst-case performance under dynamic supply voltage and temperature variations. We are interested in the specific configuration, i.e., set of seven parameters (7-tuple) described in Table 1, that causes the delay of a given path with arbitrary number of stages to be maximum.[3] Note that we construct our path delay model using our basic cell delay and output slew models. Our proposed delay and output slew models are essentially mappings f and g, respectively, from the set of all 7-tuples Q (cf. Table 1) to the positive reals, i.e., $f : Q \rightarrow \mathcal{R}^+$ and $g : Q \rightarrow \mathcal{R}^+$, where $Q = slew_{in} \times load_{out} \times cell_{size} \times amp_{noise} \times slew_{noise} \times offset_{noise} \times temp$.

For a single stage the problem of finding the worst-case configuration seeks $\vec{q}^* \in Q$ such that $f(\vec{q}^*)$ is maximized. With more than one stage in a path, i.e., $k > 1$, the output slew of the previous stage becomes the input slew to current stage, and the noise offset must be adjusted accordingly. Then, we seek $\vec{q_1}^*$ such that $f(\vec{q_1}^*) + \cdots f(\vec{q_k}^*)$ is maximized, where $\vec{q_m}^* = \vec{q_1}^*$ for all stages $1 < m < k$, except that the $slew_{in}$ component is replaced by $g(\vec{q_{m-1}}^*)$ and the $offset_{noise}$ component is adjusted at the beginning of each stage. Note that the worst-case configuration is always going to be an element of the cross-product of the various sets of parameter values. In other words, it is one of $|slew_{in}| \times |load_{out}| \times |cell_{size}| \times |amp_{noise}| \times |slew_{noise}| \times |slew_{noise}| \times |offset_{noise}| \times |temp|$ configurations. In our studies, the worst-case configuration is out of 30720 different configurations.

6. EXPERIMENTAL RESULTS AND VALIDATION

To generate our models, we randomly select 10% of our entire data set as training data; we then test the models on the other 90% of the data. To show that the selection of the training set does not substantially affect model accuracy, we randomly select 10% of the entire data set five times and show the corresponding models' maximum and average error values (Table 2).

To show the accuracy of our worst-case performance model, we compare our worst-case predictions with SPICE simulations. We construct three different paths with different number of stages, each consists of (1) only inverters, (2) only 2-input NAND, and (3) a mix of inverter and 2-input NAND gates. For (3), we construct the path starting with an inverter, and then alternating 2-input NAND gates with inverter gates. In our experiments, one of the NAND gate inputs is connected to supply voltage (v_{dd}). We evaluate our predictions using two metrics: (1) correlation of our predictions against SPICE results, and (2) relative (%) difference in delays between our proposed model and SPICE. For (1) we rank our model predictions (total of 30720 data points) in descending order with respect to the delay of the given path. Each delay value corresponds to a set of parameters (i.e., 7-tuple including all the parameters shown in Table 1). Next, we compare our predicted worst-case configuration with SPICE, and find the rank ($rank_{SPICE}$) of our predicted worst-case configuration within SPICE results. For multi-stage paths with $k > 1$ stages, we need to adjust the noise offset for each stage. To perform this we need to identify the time at which the input to stage i, where $i = 1 \cdots k$, makes the transition. This value can be estimated by calculating the delay up to stage $i-1$, and subtracting $\frac{slew_{in}^i}{1.6}$ from it, where $slew_{in}^i$ is the input slew to stage i, and $\frac{slew_{in}^i}{1.6}$ determines the 50% output slew transition.[4]

Tables 3, 4, and 5 show the comparison our our worst-case performance model with SPICE for a path consists of (1) only inverter, (2) only 2-input NAND, and (3) a mix of inverter and 2-input NAND gates, respectively. The second and third columns, represent our (2) and (1) comparison metrics, respectively. The fourth column shows where the SPICE worst-case configuration is ranked according to our proposed model ($rank_{MARS}$). We observe that our path delay models are within 4.3% of SPICE simulations. In addition, our predictions are always ranked in the top 3 (out of 30720 configurations) of the SPICE list ($rank_{SPICE}$). We note that the ability of our worst-case performance model to correctly predict worst-case configuration is beneficial for early-stage design and optimization of power distribution networks. Finally, the SPICE-computed worst-case performance value is always among top 5 predictions of our model.

Table 2: Model stability versus random selection of the training set.

Experiments	delay % diff		output slew % diff	
	max	avg	max	avg
Exp 1	56.993	5.660	55.117	6.012
Exp 2	53.342	5.458	56.896	5.976
Exp 3	53.661	5.401	56.237	5.526
Exp 4	55.419	5.552	54.883	5.311
Exp 5	55.015	5.609	55.614	5.672

Table 3: Comparison of our proposed worst-case performance model and SPICE for an inverter chain. Rank values are out of 30720 configurations.

#Stage	delay % diff	$rank_{SPICE}$	$rank_{MARS}$
1	1.08	1	1
3	3.54	3	2
5	4.29	1	1
10	3.26	2	4
20	2.42	1	1
30	2.88	1	1

plied to future technologies, as long as necessary SPICE models and device-level netlists are available.

[3]In our experiments a path consists of (1) only inverter, (2) only 2-input NAND, and (3) a mix of inverter and 2-input NAND.

[4]In our experiments, 10%-90% transition time is the slew value.

Table 4: Comparison of proposed worst-case performance model and SPICE for a 2-input NAND chain. Rank values are out of 30720 configurations.

#Stage	delay % diff	$rank_{SPICE}$	$rank_{MARS}$
1	1.34	1	1
3	3.21	1	1
5	3.69	2	3
10	3.11	1	1
20	3.43	2	3
30	2.37	2	2

Table 5: Comparison of proposed worst-case performance model and SPICE for a mixed inverter-NAND chain. Rank values are out of 30720 configurations.

#Stage	delay % diff	$rank_{SPICE}$	$rank_{MARS}$
1	1.08	1	1
3	2.73	2	4
5	3.24	3	5
10	3.36	1	1
20	3.93	2	4
30	2.85	1	1

7. CONCLUSIONS

In this paper, we have developed a methodology, based on nonparametric regression, to obtain accurate closed-form cell delay and output slew models under dynamic supply voltage and temperature variations. Our proposed models are within 6%, on average, of SPICE simulations. We show that our basic gate delay and output slew models can be used to construct delay estimates under supply noise for arbitrary critical paths. We also show that our models can accurately find the worst-case supply noise configuration that leads to worst-case delay performance. We believe that our proposed models can be beneficial in an accurate worst-case performance-driven power distribution network optimization, such as that shown in Figure 1.

8. REFERENCES

[1] K. Banerjee, S. J. Souri, P. Kapur and K. C. Saraswat, "3-D ICs: A Novel Chip Design for Improving Deep-Submicrometer Interconnect Performance and Systems-on-Chip Integration", *Proc. IEEE* 89(5) (2001), pp. 602–633.

[2] L. Carloni, A. B. Kahng, S. Muddu, A. Pinto, K. Samadi and P. Sharma, "Interconnect Modeling for Improved System-Level Design Optimization", *Proc. ASPDAC*, 2008, pp. 258–264.

[3] L. H. Chen, M. Marek-Sadowska and F. Brewer, "Buffer Delay Change in the Presence of Power and Ground Noise", *IEEE Trans. on VLSI Systems* 11(3) (2003), pp. 461–473.

[4] J. H. Friedman, "Multivariate Adaptive Regression Splines", *Annals of Statistics* 19(1) (1991), pp. 1–66.

[5] M. Fukazawa and M. Nagata, "Measurement-Based Analysis of Delay Variation Induced by Dynamic Power Supply Noise", *IEICE Trans. on Electronics* E89-C(11) (2006), pp. 1559-1566.

[6] M. Graziano, M. R. Casu, G. Masera, G. Piccinini and M. Zamboni, "Effects of Temperature in Deep-Submicron Global Interconnect Optimization in Future Technology Nodes", *Microelectronics Journal* 35 (2004), pp. 849–857.

[7] M. Hashimoto, J. Yamaguchi and H. Onodera, "Timing Analysis Considering Spatial Power/Ground Level Variation", *Proc. ICCAD*, 2004, pp. 814–820.

[8] X. Hu, D. Peng, A. Shayan and C.-K. Cheng, "Worst-Case Noise Prediction With Non-Zero Current Transition Times for Early Power Distribution System Verification", *Proc. ISQED*, 2010.

[9] Y.-M. Jiang and K.-T. Cheng, "Analysis of Performance Impact Caused by Power Supply Noise in Deep Submicron Devices", *Proc. DAC*, 1999, pp. 760–765.

[10] E. Long, W. R. Daasch, R. Madge and B. Benware, "Detection of Temperature Sensitive Defects Using ZTC", *Proc. VTS*, 2004, pp. 185–190.

[11] F. Martorell, M. Pons, A. Rubio and F. Moll, "Error Probability in Synchronous Digital Circuits Due to Power Supply Noise", *Proc. CDTISNE*, 2007, pp. 170–175.

[12] A. V. Mezhiba and E. Friedman, "Scaling Trends of On-Chip Power Distribution Noise", *IEEE Trans. on VLSI Systems* 12(4) (2004), pp. 386–394.

[13] Y. Ogasawara, T. Enami, M. Hashimoto, T. Sato and T. Onoye, "Validation of a Full-Chip Simulation Model for Supply Noise and Delay Dependence on Average Voltage Drop with On-Chip Delay Measurement", *IEEE Trans. on Circuit and Systems II* 54(10) (2007), pp. 868–872.

[14] T. Okumura, F. Minami, K. Shimazaki, K. Kuwada and M. Hashimoto, "Gate Delay Estimation in STA Under Dynamic Supply Voltage Noise", *Proc. ASPDAC*, 2010, pp. 775–780.

[15] M. Saint-Laurent and M. Swaminathan, "Impact of Power-Supply Noise on Timing in High-Frequency Microprocessors", *IEEE Trans. on Adv. Packaging* 27(1) (2004), pp. 135–144.

[16] N. H. E. Weste and D. M. Harris, *CMOS VLSI Design: A Circuits and Systems Perspective*, Addison Wesley, 2003.

[17] S.-H. Weng, Y.-M. Kuo, S.-C. Chang and M. Marek-Sadowska, "Timing Analysis Considering IR Drop Waveforms in Power Gating Designs", *Proc. ICCD*, 2008, pp. 532–537.

[18] J. Xiang and L. He, "Full-Chip Multilevel Routing for Power and Signal Integrity", *INTEGRATION, the VLSI Journal* 40 (2007), pp. 226–234.

[19] W. Zhang, Y. Zhu, W. Yu, A. Shayan, R. Wang, Z. Zhu and C.-K. Cheng, "Noise Minimization During Power-Up Stage for a Multi-Domain Power Network", *Proc. ASPDAC*, 2009, pp. 391–396.

[20] W. Zhang, L. Zhang, A. Shayan, W. Yu, X. Hu, Z, Zhu, E. Engin and C.-K. Cheng, "On-Chip Power Power Network Optimization with Decoupling Capacitors and Controlled-ESRs", *Proc. ASPDAC*, 2010, pp. 119–124.

[21] Y. Zhou and H. Leung, "Predicting Object-Oriented Software Maintainability Using Multivariate Adaptive Regression Splines", *Journal of Systems and Software* 80 (2007), pp. 1349–1361.

[22] "Power Noise Analysis for Next Generation ICs", *Apache Design Solutions*, 2009.

[23] *MARS User Guide*, http://www.salfordsystems.com/.

[24] *Synopsys HSPICE*, http://www.synopsys.com/.