Within-die Process Variations: How Accurately Can They Be Statistically Modeled?

Brendan Hargreaves
Division of Engineering
Brown University
Providence, RI 02912

Henrik Hult
Division of Applied Mathematics
Brown University
Providence, RI 02912

Sherief Reda
Division of Engineering
Brown University
Providence, RI 02912

Abstract— Within-die process variations arise during integrated circuit (IC) fabrication in the sub-100nm regime. These variations are of paramount concern as they deviate the performance of ICs from their designers' original intent. These deviations reduce the parametric yield and revenues from integrated circuit fabrication. In this paper we provide a complete treatment to the subject of within-die variations. We propose a scan-chain based system, vMeter, to extract within-die variations in an automated fashion. We implement our system in a sample of 90nm chips, and collect the within-die variations data. Then we propose a number of novel statistical analysis techniques that accurately model the within-die variation trends and capture the spatial correlations. We propose the use of maximum-likelihood techniques to find the required parameters to fit the model to the data. The accuracy of our models is statistically verified through residual analysis and variograms. Using our successful modeling technique, we propose a procedure to generate synthetic within-die variation patterns that mimic, or imitate, real silicon data.

I. Introduction

The advanced sub-wavelength semiconductor fabrication techniques have resulted in nanometer feature sizes with a substantial amount of process variations. These process variations are the result of the inability to robustly print geometric features [3, 11] and the inability to precisely control the diffusion of dopants [3, 13]. Process variations translate to variations in the key electrical parameters of circuit devices and interconnects, which increase the uncertainty in the outcome from the design process, and consequently jeopardize the parametric yield of the fabrication process. Process variations are typically divided into two components: inter-die and intra-die [3, 9]. *Inter-die* variations account for variations that arise between different chips in the same wafer or different wafers, while *intra-die*, or *within-die*, variations account for variations that arise between different devices and interconnects that reside within the same chip.

To cope with process variations, designers attempt to characterize the underlying sources of variations and then either apply statistical techniques [1, 5, 10] or design guard bands [16]. It is also possible to cope with these variations in post-manufacturing and during operational time [7]. Before developing or applying solutions to process variations, it is first fundamental to characterize, or to develop models, for process variation trends, and then use these models to derive design and manufacturing. While process variations are random in nature, within-die variations typically exhibit spatial correlations, i.e., devices that are spatially close to each other are likely to be more strongly correlated than devices that are spatially far from each other. This correlation has been the subject of a number of recent works [6, 2, 17, 8].

The objective of this paper is to develop a complete treatment to the subject of within-die process variations. We develop accurate statistical modeling techniques that fit realistic variability trends that we extract from 90nm chips. We also propose applications for our model. The contributions of this paper can be summarized as follows.

- We propose and implement a system, vMeter, to extract process variation data from sample 90nm chips. This extraction of within-die variations from actual silicon chips provides a solid basis for the development of realistic statistical modeling and estimation techniques.
- Based on the concepts of Gaussian random fields, we propose novel statistical analysis techniques that accurately fit a model to data. In particular, we propose a generic statistical model, the *Matérn* model, with enough flexibility to capture different within-die variation trends. We also propose the use of maximum-likelihood estimation techniques to calculate the required parameters of the proposed model. We thoroughly verify the accuracy of our models against the extracted data using statistical techniques such as the Kolmogorov-Smirnov test and statistical variograms.
- As an application of our model, we develop an algorithm that can generate synthetic variability trends that mimic within-die variations of real chips. Our algorithm is randomized and thus can be used to generate as many synthetic trends as required. The proposed algorithm provides a useful method to drive future research with realistic process variation models.

The organization of this paper is as follows. Section II overviews the previous research work that is related to this paper. Section III describes the components of our data extraction system (vMeter). Section IV develops the main concepts behind the statistical modeling techniques proposed by this work, and Section V shows how to use maximum-likelihood estimation techniques to find the best parameters for our models. Section VI proposes a number of techniques to thoroughly verify

978-1-4244-1922-7/08/$25.00 ©2008 IEEE

Technique	Ours	Friedberg *et al.* [6]	Bhardwaj *et al.* [2]	Xiong *et al.* [17]	Liu [8]
Correlation model	exponential & Matérn	piece-wise linear	linear & exponential	exponential & generic ("Matérn")	linear & exponential
Model parameter calculation method	maximum-likelihood estimation	linear fitting	Karhunen-Loève expansion	constrained nonlinear optimization	generalized least square fitting
Model tested on silicon measurements	yes (10 chips)	yes (35 chips)	no	no	yes (1 chip)
Statistical accuracy verficiation method	Kolmogorov-Smirnov test & SSE with variograms	SSE	SSE	SSE	difference between Kriging predictor & measurements

TABLE I
A SUMMARY COMPARING BETWEEN OUR APPROACH AND SOME OF THE RECENT PROPOSED APPROACHES. SSE STANDS FOR SUM OF SQUARED ERROR.

our model. In Section VII, we propose a method to generate synthetic within-die process variation trends that mimic realistic chips. Section VIII provides an extensive set of results from our model and ten sample chips. Finally, Section IX summarizes the main conclusions of this paper.

II. PREVIOUS WORK

A ring oscillator (RO) is a simple yet powerful device in measure process variations. It consists of an odd number of inverters cascaded sequentially in a feedback loop. For a given ambient temperature and a number of inverter stages in a RO, the frequency of oscillation of a RO depends on physical manufacturing processes [14, 11] at the location of the RO. The frequency of a RO represents a lumped value of all process variations regardless of their source. To measure the within-die process variations over a chip, it is necessary to position a number of ROs in a number of locations, and then use additional circuitry to connect the ROs in order to transfer the within-die variation data to where it is stored or processed [4, 12].

The measured variations from a test structure at a particular die location can be considered as a random variable that takes different values depending on the considered die. The collection of random variables that represent the results of all test structures form a *stochastic process* or a *random field* that is spatially indexed by locations of the test structures. A number of recent efforts investigate how to model within-die variations [6, 2, 17, 8]. Friedberg *et al.* [6] design critical dimension test structures to capture the variations in gate length, and then model the correlations between the variables of the resultant random field using piece-wise linear functions. Bhardwaj *et al.* [2] propose reducing the number of variables in the random field by using the Karhunen-Loève expansion to write the field as a series expansion of uncorrelated random variables. The uncorrelated random variables are fewer in number than the original correlated random variables, which reduces the complexity of the problem. Xiong *et al.* [17] propose modeling the correlations in the random field using exponential and "general" functions (that are reducible to the "Matérn" model). The parameters of the model are found through constrained nonlinear optimizations. As [8] notes, none of the models of [2] and [17] are applied to measured data, so further validation is still to be carried out. Recently, Liu [8] proposes the use of correlograms and variograms to model the spatial variations, and where the model parameters are determined through generalized least square fitting that is solved using Nelder-Mead simplex method. Table I gives a comparison between the previous approaches and the proposed approach in this paper.

III. VMETER: A SYSTEM TO EXTRACT WITHIN-DIE VARIATIONS

Realistic modeling of within-die process variations must start by first acquiring or extracting raw process variations data from silicon chips. In this section we briefly describe our process variation acquisition system vMeter.

Our main device in measuring the within-die variations is the *ring oscillator (RO)*. As the output frequency of a ring oscillator is sensitive to the inherent process variations of the chip, a RO frequency provides a succinct signature that determines the speed of a die at any desired location [4, 12, 11]. To measure the variations across all locations on a die, it is necessary to cover the entire die with ROs and connect them in a way that facilitates the automated extraction of their signatures. Towards that goal, our RO circuitry (Figure 1) is designed similar to a scan chain [4, 14], where the ROs are sequentially enabled one at a time for a *sample period*, during which the RO frequency is measured using a *frequency counter* and stored in a memory subsystem. Enabling one RO at a time ensures minimal current consumption, which reduces the runtime variations on the power supply network that can introduce noise in the measurements. Sequentially chaining all the ROs also reduces the bandwidth needed to transfer the signatures of all ROs to the analysis circuitry.

A block diagram of our overall extraction circuitry is given in Figure 1. The circuitry consists of $n_1 \times n_2$ ROs, where each RO occupies one of the numbered *tiles* and connects to its subsequent neighboring RO in the chain via two signals: the *scan-enable* signal which turns one RO at a time and the

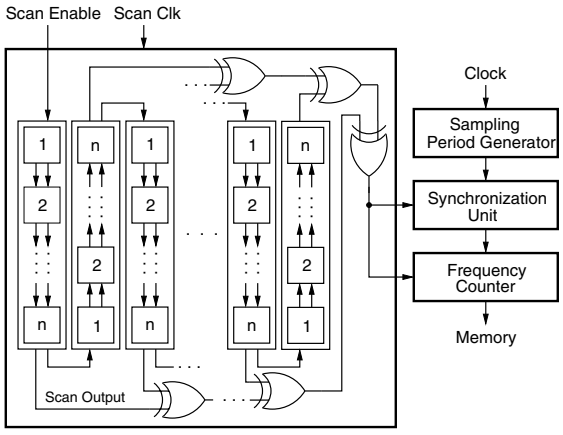

Fig. 1. Overall hardware organization of the proposed RO-based within-die process extraction system (vMeter).

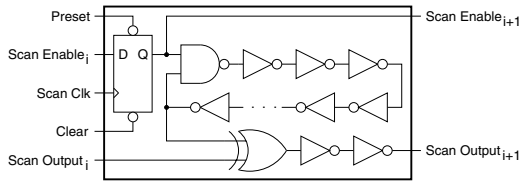

Fig. 2. Ring oscillator tile.

scan-output signal which carries the oscillatory output of each RO tile down the scan chain and ultimately to the frequency counter. The *scan clock* signal advances the enabled RO one at a time. The *frequency counter* counts the number of pulses it receives from the enabled RO tile for a *sample period* that is synchronized with the RO outputs via a *synchronization unit*.

The ring oscillator tile, given in Figure 2, forms the backbone for the within-die measurements. Each RO tile uses only local interconnects, and thus any variations in the RO frequency are mainly contributed by physical device variations. Each tile consists of a string of an odd number of inverters and the control circuitry. The flip-flop holds the value of the scan enable signal that controls the operation of the RO. The scan output and scan enable of the RO are connected to the scan input and scan enable signals of the subsequent RO tile as part of the scan chain.

To acquire accurate measurements and reduce the impact of switching noise on the power-ground network, we reduce the length of the RO scan chain using *interleaving*. A long scan chain implies that the output oscillation of each RO will have to travel down the chain, creating unnecessary switching noise on the power supply network. Thus we *interleave* the outputs of the RO tiles, by dividing the chain into *columns* as shown in Figure 1, where the scan output of each column is not chained to scan output of the next column, but rather to the outputs of other odd (or even) RO columns. More complex schemes can be used to interleave the odd and even rows together.

We implement our system into ten sample chips at 90nm technology. Figure 3 visually shows the within-die process variations for four chips, where it is clear that the variation trends exhibit systematic and random components that are

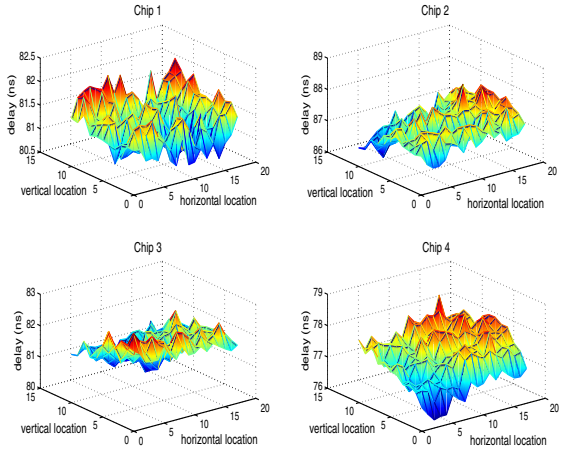

Fig. 3. Within-die process variations for four sample chips.

unique to each chip. Our goal in the subsequent sections is to develop a statistical analysis technique that can accurately model these trends.

IV. Modeling with Gaussian random fields

The die can be considered as consisting of a grid of n_1 by n_2 locations. A location on the chip will be denoted by $l = (x, y)$ where x is the horizontal coordinate and y is the vertical coordinate. The delay at a location l on chip i is denoted $D_i(l)$. Each $D_i(l)$ will be considered as a random variable with mean μ_i, where μ_i does not depend on the location l but is dictated by the inter-die variations and varies from chip to chip. The delays $D_i(l)$ and $D_i(l')$ at any two locations l and l' on the same chip i will be correlated. The correlation is typically strong at nearby locations and weak for locations far apart. To fully describe the intra-die variations we assume that the collection of random variables $D_i = \{D_i(x, y), 1 \le x \le n_1, 1 \le y \le n_2\}$, representing the delays at all different locations on chip i, form a random field. The random field D_i is assumed to be Gaussian with mean μ_i. That is, the delays at any vector of locations $(D_i(l^{(1)}), \ldots, D_i(l^{(m)}))'$ has a multivariate Gaussian distribution with mean μ_i.

To impose further structure the Gaussian random field is assumed to be *stationary* and *isotropic*. This implies that the variance σ_i^2 of the random variable $D_i(l)$ does not depend on the location l, and that the covariance between two locations l and l' only depends on the (Euclidean) distance $h = \|l - l'\|$ between l and l'. Then the distribution of D_i is completely determined by its covariance function, which can be written as

$$\mathrm{Cov}(D_i(l), D_i(l')) = \sigma_i^2 \varrho_i(\|l - l'\|).$$

The parameter $\sigma_i > 0$ is a *scale parameter* (σ_i^2 is the variance of $D_i(l)$) and the function ϱ_i is called the *correlation function*. Note that the mean μ_i, the scale parameter σ_i, and the correlation function ϱ_i all depend on i and may be different for different chips.

In many cases it is convenient to represent the random field D_i as a random vector of length $n = n_1 \times n_2$. The random vector, which also will be denoted D_i, is constructed such that the delay $D_i(l)$ at location $l = (x, y)$ is the $[n_2(x-1)+y]$'th entry in the vector. With this representation D_i is a Gaussian random vector with mean μ_i and covariance matrix $\sigma_i^2 \Omega_i$ where

$$[\Omega_i]_{n_2(x-1)+y, n_2(x'-1)+y'} = \varrho_i(\|l - l'\|).$$

To fully specify the model for intra-chip variations it only remains to specify a valid correlation function ϱ_i. Two models will be considered: the exponential model and the Matérn model, the first actually being a special case of the second.

A. The exponential model

A simple and natural model that allows for correlation between different locations is the exponential model. For this model the correlation function decays exponentially as a function of the distance $h = \|l - l'\|$, i.e.

$$\varrho_i(h) = e^{-\lambda_i h}, \quad \lambda_i > 0. \tag{1}$$

Note that as λ_i increases the correlation decays faster as a function of the distance. In this respect λ_i can be interpreted as the strength of correlation. Under this model, the random field D_i has three parameters; the mean level μ_i, the scale parameter σ_i, and the strength of correlation λ_i.

B. The Matérn model

The exponential model is attractive because of its simplicity but it is not very flexible in capturing a wide range of correlation structures. Another popular and more flexible class of correlation functions is the Matérn class [15]. In contrast to the exponential class the Matérn correlation function is parameterized by two parameters, $\theta_{1i} > 0$ and $\theta_{2i} > 0$, and has the functional form

$$\varrho_i(h) = \frac{1}{2^{\theta_{2i}-1}\Gamma(\theta_{2i})}\Big(\frac{2h\sqrt{\theta_{2i}}}{\theta_{1i}}\Big)^{\theta_{2i}}\mathcal{K}_{\theta_{2i}}\Big(\frac{2h\sqrt{\theta_{2i}}}{\theta_{1i}}\Big),$$

where $\mathcal{K}_\alpha(\cdot)$ denotes the modified Bessel function of the second kind of order α, and $\Gamma(\cdot)$ denotes the Gamma function. The parameter θ_{1i} can be interpreted as the rate of decay of the correlation as a function of distance. Large values of θ_{1i} lead to faster decay of correlations as distance increases. The parameter θ_{2i} controls the general shape of the correlation function and in particular the behavior of the correlation at small distances.

An attractive feature of the Matérn class is that it contains three important and widely used correlation functions as special cases; the linear model ($\theta_{2i} \to 0$), the exponential model ($\theta_{2i} = 0.5$), and the Gaussian model ($\theta_{2i} \to \infty$). Under the Matérn model the random field D_i has four parameters; the mean level μ_i, the scale parameter σ_i, and the correlation parameters θ_{1i} and θ_{2i}.

V. Parameter estimation using maximum likelihood

To fit the exponential model or the more general Matérn model to particular within-die measurements of a sample chip, one has to find the model parameters that provide the best fit. In statistical literature one of the most popular approaches to estimate unknown parameters from observed data is the maximum likelihood estimation method. This entails maximizing the likelihood function (or equivalently the log-likelihood function) over all possible parameter values. The likelihood function is the probability density evaluated at the observed values. Representing the Gaussian random field D_i as a random vector with mean μ_i and covariance matrix $\sigma_i^2\Omega_i$ the logarithm of the likelihood function is, up to an additive constant,

$$\ell(\mu_i, \sigma_i, \Omega_i) = -\frac{1}{2}\log\det(\sigma_i^2\Omega_i) - \frac{1}{2\sigma_i^2}(D_i-\mu_i)'\Omega_i^{-1}(D_i-\mu_i).$$

For the exponential model the correlation matrix Ω_i is completely determined by the unknown parameter λ_i. Then the following procedure leads to point estimates for μ_i, σ_i, and λ_i.

Maximizing with respect to μ_i and σ_i^2 yields the estimates

$$\begin{aligned}\hat{\mu}_i &= \frac{\mathbf{1}'\Omega_i^{-1}D_i}{\mathbf{1}'\Omega^{-1}\mathbf{1}}\\ \hat{\sigma}_i^2 &= (D_i-\hat{\mu}_i)'\Omega_i^{-1}(D_i-\hat{\mu}_i)/n,\end{aligned}$$

with $\mathbf{1} = (1,\ldots,1)'$. Plugging these expression back into the log-likelihood function we obtain the so-called profile log-likelihood

$$\ell_p(\lambda_i) = -\frac{1}{2}\log\det\Omega_i - \frac{n}{2}\log\hat{\sigma}_i^2 - \frac{n}{2},$$

which is to be maximized over $\lambda_i > 0$. This can be done using any standard numerical optimization algorithm.

For the Matérn model the only difference is that Ω_i depends on two parameters θ_{1i} and θ_{2i} instead of λ_i. Then the profile log-likelihood function has to be maximized over all positive values of θ_{1i} and θ_{2i}. This can also be done using standard numerical optimization. In our implementation, we use the MATLAB function *fminsearch* which is an implementation of the Nelder-Mead algorithm for unconstrained nonlinear optimization.

VI. Model verification

In this section two measures of model verification, or goodness-of-fit tests, are presented. The proposed procedures are intended to evaluate if the suggested model is a believable model for measurement data. The main objective is to evaluate the explanatory power of the model.

A. Verification by residual analysis

An effective measure of goodness-of-fit is to perform a statistical test to investigate if the measured data is consistent with the model or not. To this end we will analyze the residuals of the Gaussian random fields model D_i described in Section IV.

We will make use of the representation of D_i as a Gaussian random vector with mean μ_i and covariance matrix $\sigma_i^2\Omega_i$. Then D_i can be written in the form

$$D_i = \mu_i + \sigma_i A_i W_i, \tag{2}$$

where A_i is the Cholesky decomposition of Ω_i (the matrix A_i such that $A_i'A_i = \Omega_i$) and W_i is a vector of n independent $N(0,1)$ random variables. To see this note that with this representation D_i is a linear combination of independent $N(0,1)$ variables. Hence, it has a joint Gaussian distribution with mean μ_i and covariance matrix

$$\mathrm{Cov}(D_i) = E[(D_i-\mu_i)'(D_i-\mu_i)] = \sigma_i^2 A_i'A_i = \sigma_i^2\Omega_i.$$

Inverting the relation (2) yields the *residuals* W as

$$W_i = \frac{1}{\sigma_i}A_i^{-1}(D_i-\mu_i).$$

Using the observed values of D_i and the estimated parameter values, we use the last equation to obtain observed values of the residuals W_i. Then a statistical test can be performed to test if the observed residuals are consistent with observations from independent $N(0,1)$ random variables. To this end the Kolmogorov-Smirnov test will be applied. For this test the largest distance between the empirical distribution function of the observed residuals and the standard normal distribution function is measured. The Kolmogorov-Smirnov test statistic is

$$KS_i = \max_x |F_{in}(x) - \Phi(x)|,$$

where $F_{in}(x) = \sum_{j=1}^{n} I\{W_{ij} \leq x\}$ is the empirical distribution of $W_i = (W_{i1}, \ldots, W_{in})'$ and $\Phi(x)$ is the standard normal distribution function. If the null hypothesis, that the variables in W_i are independent standard normal, is rejected then there is evidence that the variables in W_i are not standard normal. Otherwise, if the null hypothesis is not rejected, the data is consistent with the standard normal assumption. This provides evidence in favor of the model. The hypothesis test will be performed at the 1% confidence level implying that there is only a 1% chance the null hypothesis will be (incorrectly) rejected when the null hypothesis is true.

In addition to the result of the hypothesis test it is also desirable to report the P-value. The P-value can be interpreted as the amount of support for the null hypothesis. It is computed as the probability (under the assumption that the null hypothesis is true) that the test statistic KS would take a value equal or larger than the observed test statistic. It is intimately connected to the result of the hypothesis test as the null hypothesis is rejected if the P-value is below 1%.

B. Verification using variograms

One of the main features of the suggested model is to allow for spatial correlation between different locations on the chip. In this respect it is desirable to investigate if the suggested correlation structure captures the main features of the data. An exploratory analysis of the correlation structure can be performed by studying the *variogram* and the *sample variogram*. The variogram is a popular tool in spatial statistics; in particular in geostatistical analysis. It has also been suggested in connection with spatial variations [8]. For a stationary isotropic random field D_i the variogram is defined by

$$\gamma(h) = \frac{1}{2}\mathrm{Var}(D(l) - D(l')), \quad \|l - l'\| = h.$$

The sample version of the variogram, called the sample variogram, is based on observations of a random field. It is computed by

$$\hat{\gamma}(h) = \frac{1}{N_h} \sum_{l,l':\|l-l'\|=h} (D(l) - D(l'))^2,$$

where N_h is the number of pairs l, l' such that $\|l - l'\| = h$.

If the model accurately captures the correlation structure, one expects the sample variogram for the measurement data to be similar to the theoretical variogram. However, one cannot expect a perfect correspondence. Because of the intrinsic randomness, the sample variogram will differ from the theoretical.

VII. Synthetic generation of within-die variations

A very attractive feature of Gaussian random fields is that it is easy to generate pseudo-random samples from them on a computer. That is, given the model and values of the model parameters it is easy to generate a large number of synthetic chips. In this section we describe a simple algorithm to generate samples from Gaussian fields. We present the algorithm for the isotropic field D_i in the previous section, but it is elementary to generalize to more complicated structures when the mean level and the scale parameter depend on the location on the chip as well as more complex correlation function.

Suppose the Gaussian random field D_i is represented as a Gaussian random vector of length $n = n_1 \times n_2$ with mean μ_i and covariance matrix $\sigma_i^2 \Omega_i$.

To generate a sample of D_i the following algorithm is implemented.

1. Generate a column vector $W_i = (W_{i1}, \ldots, W_{in})'$ of independent $N(0, 1)$ variables.
2. Compute A_i, the Cholesky decomposition of Ω_i, i.e. the matrix A_i such that $A_i' A_i = \Omega_i$.
3. Let $D_i = \mu_i + \sigma_i A_i W_i$.

Then D_i has the representation (2) which shows that it has the desired distribution.

VIII. Experimental Results

To evaluate the appropriateness and accuracy of the suggested statistical models, we first design the proposed vMeter system to extract within-die process variations. We then implement the system in ten sample chips (Altera's EP2C35 devices) manufactured in 90nm technology, with all chips belong to the same speed bin (C6). Each chip holds 198 ring oscillator tiles organized as a 18×11 lattice, i.e., $n_1 = 18$ and $n_2 = 11$. Each tile is composed of 135 inverters that are organized in a 3×3 logic blocks. As described earlier, Figure 3 gives the within-die variations from four chips of the ten sample chips. With the extracted data in hand, we carry out the following procedure based on the previous discussions.

- Depending on the modeling flexibility required, choose the appropriate model: exponential (Subsection IV.A) or *Matérn* (Subsection IV.B).
- Calculate the model parameters using maximum-likelihood estimation (Section V) and the provided extraction data.
- Verify the accuracy of the model using Kolmogorov-Smirnov test (Subsection VI.A) and/or using sum of squared errors from data variograms (Subsection VI. B).

A. Statistical analysis of measurement data

The first thing to explore is if the Gaussian random field model is reasonable for describing the measured intra-chip variations. We certainly require that the model is able to capture the essential structure of the spatial correlations. For this purpose, we first plot the variograms calculated for the silicon measurements of the ten chips in Figure 4 (left). It is natural to ask what sample variograms are likely to look like if, for example the exponential model with $\sigma = 0.05$ and $\lambda = 0.4$ or the Matérn model with $\sigma = 0.07$, $\theta_1 = 30$, and $\theta_2 = 0.1$ were true. By generating 100 synthetic chips for the exponential model and the Matérn model, using the algorithm described in Section VII, and computing the corresponding sample variograms, we may examine if the model is likely to produce sample variograms that are similar to the actual chips. The sample

variogram for synthetic chips together with the true variogram for the exponential and Matérn model are illustrated in the center and right plots of Figure 4, respectively. From these plots, we observe the following.

- The 100 random variograms generated from the model, whether using the exponential or the Matérn model, often deviate from the theoretical variogram (displayed by thick black line). In particular, we observe that the Matérn model is flexible enough to describe a wider range of spatial correlations than the exponential model.
- All of our silicon chips have variograms that are consistent with the Matérn model with 'roughly' the suggested parameter values. That is, the 10 variograms computed from the silicon chips appear as a subset from the sample space variograms generated from the Matérn model.

B. Fitting silicon data to statistical models

The maximum likelihood estimation method described in Section V is used to estimate the parameters μ, σ, and λ of the exponential model for each of the ten chips. To evaluate the goodness of fit, the Kolmogorov-Smirnov test of the residuals is performed to check whether they appear to have standard normal distribution. The results are summarized in Table II. For the Matérn model the corresponding maximum likelihood estimation and goodness-of-fit tests are also performed. The results are summarized in Table II and the resulting model variograms and sample variograms are illustrated in Figure 5. It is noteworthy that none of the Kolmogorov-Smirnov tests is rejected at the 1% level. This indicates a good fit. We also note that the mean value, as dictated by the inter-die variations, vary from chip to chip as expected. The maximum difference due to inter-die variations is $\frac{9.68-8.47}{8.47} = 14.2\%$.

To see how well the model captures the correlation structure the theoretical variogram based on the exponential model (blue line), the Matérn model (green) is illustrated together with the actual variograms of chips 1, 2, 3 and 4 (red line) in Figure 5.

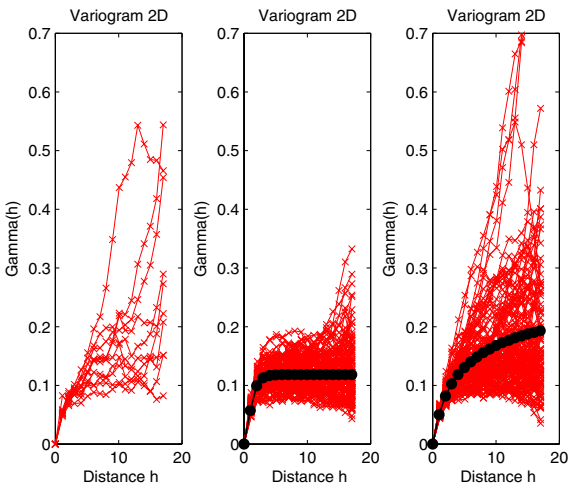

Fig. 4. *Left*: Sample variograms for ten chips. *Center*: Sample variograms for a hundred synthetically generated chips using the exponential model with $\sigma = 0.05$ and $\lambda = 0.4$. *Right*: Sample variograms for a hundred synthetically generated chips using a Matérn model with $\sigma = 0.07$, $\theta_1 = 30$, and $\theta_2 = 0.1$. The variogram for the model is given in (black o).

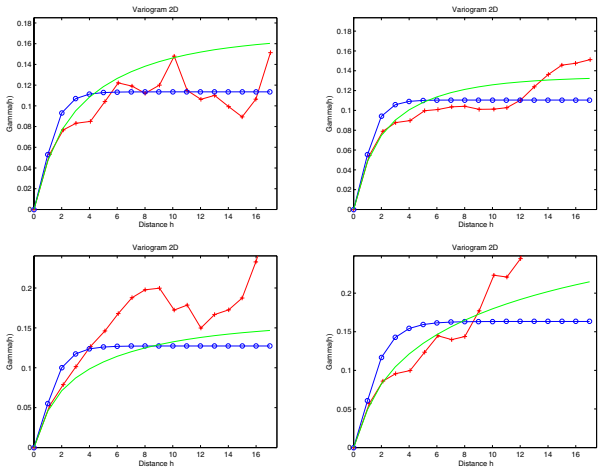

Fig. 5. Sample variogram (red x) and estimated variogram with exponential (blue o) and Matérn (green) for chip 1 (top left), chip 2 (top right), chip 3 (bottom left) and chip 4 (bottom right).

- From Figure 5, we observe that the Matérn model captures more of the correlation structure in comparison to the exponential model. For example, in chip 4 where the variogram deviates considerably from the exponential model, the Matérn model gives a better fit.

Besides generating synthetic variograms as was given in Figure 4, our method is capable of generating complete within-die variations patterns as outlined in Section VII. In Figure 6, we show four synthetic within-die variation trends ($\mu_i = 0$ and $\sigma_i = 1$). Besides the thorough statistical validation provided in the previous two subsections, we visually compare the synthetic trends in Figure 6 to the actual silicon trends of Figure 3. We find that our synthetic trends are visually similar to the actual trends. Our method for generating synthetic data can be of great value for researchers who would like to derive their process variation-based research with realistic within-die process variation models.

IX. Summary and Conclusions

In this paper we have developed a complete treatment for the subject of within-die process variations. We have designed and implemented a process variation extraction system in 90nm chips. To find a model that captures the extracted data, we have proposed a number of statistical models that accurately capture the correlation between the different spatial locations on the test chips. We have also described how to calculate the required parameters for our proposed models using maximum likelihood estimation, and thoroughly verified the correctness of our models and parameters. We have also proposed a procedure to generate synthetic variability trends that mimic realistic silicon chips. The procedure can be utilized by other researchers to generate accurate synthetic within-die variation trends for their experiments.

die	Exponential Model						Matérn Model						
	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\lambda}$	Rej.	P-value	KS	$\hat{\mu}$	$\hat{\sigma}$	$\hat{\theta}_1$	$\hat{\theta}_2$	Rej.	P-val	KS
1	9.04	0.048	0.38	N	0.97	0.034	9.05	0.064	16.3	0.10	N	0.98	0.034
2	9.68	0.049	0.42	N	0.91	0.039	9.67	0.064	9.76	0.080	N	0.34	0.066
3	9.08	0.050	0.34	N	0.03	0.101	9.08	0.065	18.6	0.08	N	0.20	0.079
4	9.20	0.054	0.27	N	0.76	0.048	9.22	0.075	59.3	0.11	N	0.16	0.079
5	9.07	0.074	1.32	N	0.31	0.068	9.07	0.094	8.36	0.0069	N	0.03	0.104
6	8.56	0.057	0.20	N	0.62	0.053	8.55	0.077	46.4	0.19	N	0.42	0.062
7	8.60	0.061	0.17	N	0.78	0.047	8.59	0.077	92.0	0.12	N	0.43	0.062
8	8.70	0.052	0.26	N	0.42	0.062	8.70	0.076	90.4	0.12	N	0.37	0.065
9	8.47	0.051	0.30	N	0.62	0.053	8.48	0.066	32.5	0.078	N	0.93	0.038
10	8.59	0.047	0.38	N	0.47	0.060	8.60	0.063	22.4	0.086	N	0.70	0.050

TABLE II

PARAMETER ESTIMATION FOR THE EXPONENTIAL AND THE MATÉRN MODELS USING MAXIMUM-LIKELIHOOD ESTIMATION. KOLMOGOROV-SMIRNOV TEST IS USED FOR RESIDUALS FOR EACH OF THE TEN CHIPS.

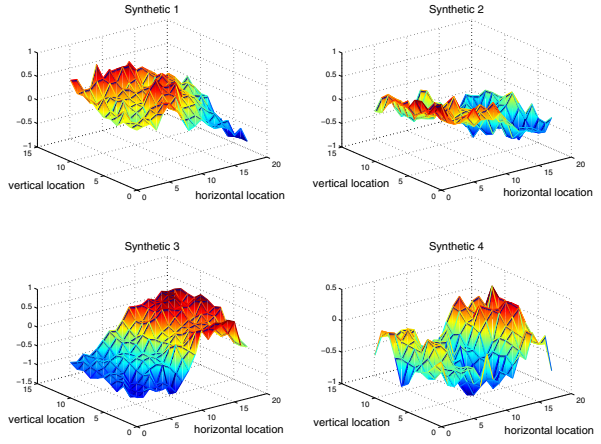

Fig. 6. Four random synthetic within-die process variations trends ($\mu_i = 0$ and $\sigma_i = 1$).

REFERENCES

[1] A. Agarwal, D. Blaauw, and V. Zolotov, "Satistical Timing Analysis for Intra-Die Process Variations For Spatial Correlations," in *Proc. IEEE International Conference on Computer Aided Design*, 2003, pp. 900–907.

[2] S. Bhardwaj, S. Vrudhuka, P. Ghanta, and Y. Cao, "Modeling of Intra-Die Process Variation for Accurate Analysis and Optimization of Nano-Scale Circuits," in *Proc. ACM/IEEE Design Automation Conference*, 2006, pp. 791–796.

[3] D. Boning and S. Nassif, "Models of Process Variations in Device and Interconnect," in *Design of High-Performance Microprocessor Circuits*, 1st ed., A. Chandrakasan, W. J. Bowhill, and F. Cox, Eds. IEEE Press, 2001, pp. 98–115.

[4] D. Boning, J. Panganiban, K. Gonzalez-Valentin, S. Nassif, C. McCowell, A. Gattiker, and F. Liu, "Test Structures for Delay Variability," in *ACM/IEEE International Workshop on Timing Issues in the Specification and Synthesis of Digital Systems*, 2002, p. 109.

[5] H. Chang and S. Sapatnekar, "Statistical Timing Analysis Considering Spatial Correlation Using a Single PERT-Like Traversal," in *Proc. IEEE International Conference on Computer Aided Design*, 2003, pp. 621–625.

[6] P. Friedberg, Y. Cao, J. Cain, R. Wang, J. Rabaey, and C. Spanos, "Modeling Within-Die Spatial Correlation Effects for Process-Design Co-Optimization," in *International Symposium on Quality Electronic Design Automation*, 2005, pp. 516–521.

[7] C. Kim, J.-J. Kim, I.-J. Chang, and K. Roy, "PVT-Aware Leakage Reduction for On-Die Caches with Improved Read Stability," *IEEE Journal of Solid-State Circuits*, vol. 41(1), pp. 170–178, 2006.

[8] F. Liu, "A General Framework for Spatial Correlation Modeling in VLSI Design," in *Proc. ACM/IEEE Design Automation Conference*, 2007, pp. 817–822.

[9] K. Okada and H. Onodera, "Statistical Parameter Extraction for Intra- and Inter-Chip Variabilites of Metal-Oxide-Semiconductor Field-Effect Transistor Characteristics," *Japanese Journal of Applied Physics*, vol. 44(1A), pp. 131–134, 2005.

[10] M. Orshansky and A. Bandyopadhyay, "Fast Statistical Timing Analysis Handling Arbitrary Delay Correlations," in *Proc. ACM/IEEE Design Automation Conference*, 2004.

[11] M. Orshansky, L. Milor, and C. Hu, "Characterization of Spatial Intrafield Gate CD Variability, Its Impact on Circuit Performance, and Spatial Mask-Level Correction," *IEEE Transactions on Semicondictor Manufacturing*, vol. 17(1), pp. 2–11, 2004.

[12] J. S. Panganiban, *A Ring Oscillaor Based Variation Test Chip*. Masters thesis. Massachusetts Institute of Technology, 2006.

[13] S. Roy and A. Asenov, "Where Do the Dopants Go?" *Science*, vol. 309, pp. 388–390, 2005.

[14] S. B. Samaan, "Parameter Variation Probing Technique," in *U.S. Patent Number 6,535,013*, 2003.

[15] M. L. Stein, *Interpolation of Spatial Data*, First ed. Springer, 1999.

[16] C. Visweswariah, "Death, Taxes and Failing Chips," in *Proc. ACM/IEEE Design Automation Conference*, 2003, pp. 343–347.

[17] J. Xiong, V. Zolotov, and L. He, "Robust Extraction of Spatial Correlation," in *Proc. ACM/IEEE International Symposium on Physical Design*, 2006, pp. 2–9.