

# Multiple-Population Moment Estimation: Exploiting Interpopulation Correlation for Efficient Moment Estimation in Analog/Mixed-Signal Validation

Chenjie Gu, *Member, IEEE*, Manzil Zaheer, *Student Member, IEEE* and Xin Li, *Senior Member, IEEE*

**Abstract**—Moment estimation is an important problem during circuit validation, in both presilicon and postsilicon stages. From the estimated moments, the probability of failure and parametric yield can be estimated at each circuit configuration and corner, and these metrics are used for design optimization and making product qualification decisions. The problem is especially difficult if only a very small sample size is allowed for measurement or simulation, as is the case for complex analog/mixed-signal circuits. In this paper, we propose an efficient moment estimation method, called multiple-population moment estimation (MPME), that significantly improves estimation accuracy under small sample size. The key idea is to leverage the data collected under different corners/configurations to improve the accuracy of moment estimation at each individual corner/configuration. Mathematically, we employ the hierarchical Bayesian framework to exploit the underlying correlation in the data. We apply the proposed method to several datasets including postsilicon measurements of a commercial high-speed I/O link, and demonstrate an average error reduction of up to 2x, which can be equivalently translated to significant reduction of validation time and cost.

**Index Terms**—Analog/mixed-signal validation, Bayesian inference, extremely small sample size, moment estimation.

## I. INTRODUCTION

**D**URING circuit validation, it is crucial to make statistically valid predictions of the circuit performances of interest. The statistical nature of the problem comes from the fact that the latest process technology witnesses increasingly larger variability, and that systems are becoming so complex that effects from environment and surrounding circuits cannot be neglected, and as a result, they exhibit randomness in the circuit performance. Such statistical predictions are important since they are used to guide design optimization and to make key decisions such as whether the product is ready for high-volume manufacturing/shipping. A key problem in this process is the problem of estimating the probability distribution of

circuit performances, *a.k.a.*, density estimation. From this distribution, metrics such as the probability of failure (PoF) or yield can be derived for further analysis and optimization.

Traditional approaches for density estimation [1], [2] include parametric and non-parametric methods. While the existing techniques have obtained much success in various applications, they all require enough number of samples for the result to be accurate. That is, if the sample size is small, the result can be biased by the data, and may not be trusted.

This is the small-sample-size problem in circuit validation. It is further exacerbated for analog/mixed-signal applications, because both simulation and measurement of many analog/mixed-signal circuit performances are time and cost consuming [3]–[5]. For example, postlayout simulation can be slow, especially for circuits such as SRAM/PLL where extremely small time steps are required for high accuracy. As another example, during postsilicon validation, due to tight product release schedules, only a limited amount of measurement may be performed within the postsilicon time-frame. In addition, the measurement of performance metrics, such as bit-error-ratio (BER) and time/voltage margins of high-speed I/O links, takes a long time, and requires expensive equipment (such as BER testers) [6]–[8]. Taking into consideration all the practical issues, a very small number of samples are affordable within reasonable timeframe.

Unfortunately, there is few existing satisfying solution to get around this problem. To the best of our knowledge, the usual practice is to increase the sample size as much as possible to reach a certain confidence level, or to set an empirical guard-band on top of the estimation. There is a recent work [9] that considers a similar problem, but for performance modeling. Another recently published technique [10] solves a similar problem for postlayout performance distribution estimation, but with mildly small number of samples (50 or more).

Another problem that is sometimes ignored in circuit validation is that circuit performance distributions need to be estimated at various corners and configurations for various similar products at different steppings. For example, during I/O interface validation such as PCIE [11] and DDR [12], in addition to the traditional process, voltage and temperature (PVT) corners, we must also validate against different board/add-in card/Dual inline memory module (DIMM) configurations, input patterns, different equalization settings, *etc.*. In another word, the interface should meet the PoF specification for any

Manuscript received September 10, 2013; revised November 26, 2013 and January 16, 2014; accepted January 28, 2014. Date of current version June 16, 2014. This work was supported by Intel Corporation. This paper was recommended by Associate Editor H. E. Graeb.

C. Gu is with Intel Strategic Computer-Aided Design Laboratories, Hillsboro, OR 97124 USA (e-mail: chenjie.gu@gmail.com).

M. Zaheer and X. Li are with the Electrical and Computer Engineering Department, Carnegie Mellon University, Pittsburgh, PA 15213 USA (e-mail: manzil@cmu.edu; xinli@cmu.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCAD.2014.2304929

customer configuration of board and add-in cards. Therefore, it is inappropriate to mix the measurements under different configurations, because even with a low PoF across all configurations, we may obtain a very high PoF at a particular configuration. In this case, combining data from all configurations does not help us to increase the sample size. In fact, estimating the overall distribution can lead to misleading validation results.

In this paper, we present multiple-population moment estimation (MPME) which encapsulates a class of methods to efficiently estimate the moments of performance distributions at multiple corners and configurations. We try to solve the small sample size problem (*i.e.*, sample size ranging from 1 to 10) by exploiting the underlying correlation of data collected at multiple corners and configurations. In particular, we emphasize that data collected at different design stages, different configurations and different corners are not independent, but are correlated. Taking advantage of this non-intuitive fact leads to a theoretically guaranteed better estimator. While we focus on the moment estimation problem in this paper, it is possible to extend the idea to more general parametric and non-parametric density estimation problems.

Mathematically, MPME builds a generative graphical model to model the data obtained from simulation and measurement. Equivalently, the statistical graphical model defines a (parameterized) joint prior distribution of the moments at multiple populations. With the graphical model, MPME estimates the moments in two steps. First, the Maximum Likelihood Estimation (MLE) method is used to learn the prior distribution of moments. Second, the prior distribution learned in the first step is used to obtain the maximum *a posteriori* (MAP) estimation of moments at individual populations. Experimental results show that in comparison to traditional sample moment estimators, MPME reduces the average error by up to 2x in the best case for examples obtained from measurement of commercial designs.

The rest of paper is organized as follows. Sec. II formulates the problem, and explains why existing techniques can be problematic when a small number of samples are present. Sec. III describes rational and theory behind the MPME approach, and Sec. IV discusses advantages, potential limitations, and practical applications of the method. Sec. V presents experimental results on several datasets to demonstrate that MPME is consistently superior than traditional techniques in terms of estimation accuracy.

## II. BACKGROUND AND PROBLEM FORMULATION

In this paper, we consider the problem of estimating a circuit performance metric, denoted by  $x$ , which depends on many variables such as process parameters, voltage, temperature, board, add-in card, *etc.*. The performance metric  $x$  can also depend (indirectly) on time, because a subset of the parameters, such as process parameters, change over time.

As a concrete example application, we consider the problem of postsilicon validation of high speed I/O interfaces. In this application, a configuration is defined by fixing the values of a subset of the parameters. By considering variability of

all the other parameters,  $x$  exhibits a distribution at each configuration. For example, a configuration of an I/O link can be defined by the combination of a specific board and a specific add-in card. The variability of time/voltage margin (of the eye diagram) is caused by parameter variations such as PVT variations. Measurement of margins is repeated at each configuration for each silicon stepping, and the goal of validation is to ensure that PoF meets the specification at each stepping and at each configuration.

### A. Problem Formulation

To formalize the above description, we define a population to be a specific (corner, configuration, stepping) combination,<sup>1</sup> and denote  $P$  by the number of populations. For each population, we define a random variable  $x_i$ , ( $i = 1, \dots, P$ ) to model the variability of the performance metric at the corresponding population, and  $x_i$  satisfies a Gaussian distribution  $x_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$  where  $\mu_i$  is the mean and  $\sigma_i^2$  is the variance. For notational convenience, we define  $\boldsymbol{\mu} = [\mu_1, \dots, \mu_P]^T$  and  $\boldsymbol{\sigma}^2 = [\sigma_1^2, \dots, \sigma_P^2]^T$ .

In this formulation, the Gaussian distribution assumption is a simplification of the problem which is often used in practice. We discuss the potential extensions to non-Gaussian distributions and higher-order moments in Sec. IV-D.

For each population, we obtain a set of independent observations  $\mathcal{X}_i = \{x_{i,1}, \dots, x_{i,N_i}\}$ , where  $N_i$  is the sample size of the  $i$ th population. Each element in  $\mathcal{X}_i$  corresponds to one independent measurement at the  $i$ th population. The problem we aim to address is to estimate the moments  $(\mu_i, \sigma_i^2)$ ,  $i = 1, \dots, P$ , given the observations  $\{\mathcal{X}_1, \dots, \mathcal{X}_P\}$ .

For example, in Sec. V-C,  $\mathcal{X}_i$ ,  $i = 1, \dots, 8$  represent eight sets of observations at eight different link configurations, and  $x_{i,j}$ ,  $j = 1, 2, \dots$  represent the time margin measurement of the I/O link. We would like to estimate the time margin distributions at eight different configurations by estimating the first two moments.

The difficulty of this problem is that the sample sizes  $N_i$ 's can be extremely small. On the one hand, each individual sample can be very expensive to obtain due to long simulation/measurement time. On the other hand, since the validation must be performed at each configuration and corner, we have to obtain  $\sum_{i=1}^P N_i$  samples in total. With a large  $P$ , it might be impossible to obtain that many samples within a reasonable amount of time. This effectively results in even smaller  $N_i$ 's. With a very small sample size, the estimated moments could have a large error.

### B. Low Confidence under Small Sample Size

For a specific population, the most widely used estimator for mean and variance is the sample mean  $\bar{x}_i$  and sample variance  $S_i$ , respectively

$$\bar{x}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} x_{i,j}, \quad S_i = \frac{1}{N_i - 1} \sum_{j=1}^{N_i} (x_{i,j} - \bar{x}_i)^2. \quad (1)$$

<sup>1</sup>In this definition, (VT) corner refers to the assignment of supply voltage or temperature; configuration refers to the I/O link configurations such as data rate, board impedance, add-in card; stepping refers to a Silicon tape-out. Obviously, this definition is closely related to the postsilicon I/O validation problem. Readers can define the population that suits the application at hand.

Since  $\bar{x}_i \sim \mathcal{N}(\mu_i, \frac{\sigma_i^2}{N_i})$  and  $S_i \sim \frac{\sigma_i^2}{N_i-1} \chi_{N_i-1}^2$ , we obtain

$$\text{Std}(\bar{x}_i) = \frac{1}{\sqrt{N_i}}\sigma_i, \quad \text{Std}(S_i) = \frac{\sqrt{2}}{\sqrt{N_i-1}}\sigma_i^2. \quad (2)$$

If the standard deviation of an unbiased estimator is used as a measure of accuracy and confidence level, (2) shows that the accuracy of both sample mean and sample variance estimators depend on  $N_i$ . As  $N_i$  approaches infinity, the error converges to 0. However, when  $N_i$  is small, both estimators suffer from significant error.

### C. Handling Multiple Populations

One common way to handle multiple populations is to build a performance model. For example, consider the  $P, V, T$  variations, one might fit a response surface model (RSM) [13]

$$x = h(P, V, T). \quad (3)$$

Define the  $i$ th population by a specific  $(V, T)$  combination, denoted by  $(v_i, t_i)$ , we have

$$x_i = h(P, v_i, t_i) \quad (4)$$

from which the distribution of  $x_i$  can be derived given the distribution of  $P$ .

This is a viable solution, but its success is dependent on two critical assumptions. First, the configuration variables are continuous (not categorical). Second,  $x$  has a strong dependence on the configuration variables, and the underlying performance model template (such as RSM) is correct. These assumptions can often be broken in practice. Furthermore, a potential drawback for RSM technique is that the number of measurements must be at least as many as the number of underlying random variables. If there are too many parameters (*e.g.*, for characterizing process variability), we need many measurements which might not be affordable. Other techniques must be sought to handle multiple populations.

## III. MULTIPLE POPULATION MOMENT ESTIMATION

### A. Overview

As is evident in Sec. II-B, if each population is treated independently, there is little room for improvement. In contrast, MPME views data at different populations as correlated, and it tries to exploit such correlation to improve the accuracy of the estimator.

To model the correlation, MPME imposes a generative graphical model<sup>2</sup> which describes how the data are generated at multiple populations. Equivalently, it specifies a joint prior distribution on the moments  $\mu_i$ 's and  $\sigma_i^2$ 's. For example, the generative graphical model shown in Fig. 1(a) specifies a model where  $(\mu_i, \sigma_i^2)$  follow a distribution  $p(\mu_i, \sigma_i^2|\theta)$  parameterized by  $\theta$ , and the  $i$ th population  $x_i$  follows a Gaussian distribution with mean  $\mu_i$  and variance  $\sigma_i^2$ .

With the graphical model, MPME follows a two-step approach to estimate the moments.

<sup>2</sup>Graphical models [1], *a.k.a.*, probabilistic graphical models, provide a way to describe the probabilistic structure in a set of random variables. We provide a short introduction in Appendix A that is relevant for this paper.

- 1) First, a prior distribution of  $p(\mu_i, \sigma_i^2|\theta)$  is learned from data at all populations, using maximum likelihood estimation.
- 2) Second, maximum *a posteriori* (MAP) estimation is applied to each population using the prior distribution learned from the first step.

### B. Correlation Helps Improving Estimation Accuracy

Before we introduce MPME, it is instructive to look at two specific examples for which we can perform error analysis. The closed-form expressions intuitively explain why correlation can help improving estimation accuracy. It can also be shown that the estimators described in the two examples can be thought of as extreme cases of MPME. In these two examples, for simplicity, we consider the case where all populations have same number of independent samples, *i.e.*,  $N_1 = \dots = N_P = N$ .

*Example 1 (unequal mean, equal variance):* Assume that  $\mu_i$ 's are different, and  $\sigma_1^2 = \dots = \sigma_P^2 = \sigma^2$ , and consider the problem of estimating  $\sigma^2$ .

Since  $S_i \sim \frac{\sigma^2}{N-1} \chi_{N-1}^2$ , we obtain an unbiased estimator for  $\sigma^2$

$$\frac{1}{P}[S_1 + \dots + S_P] \sim \frac{1}{P} \frac{\sigma^2}{N-1} \chi_{NP-P}^2 \quad (5)$$

from which  $\text{Std}(\frac{1}{P}[S_1 + \dots + S_P]) = \sigma^2 \sqrt{\frac{2}{P(N-1)}}$ . Hence, the estimation error decreases as  $P$  increases, and is smaller than  $\text{Std}(S_i)$ .

*Example 2 (equal mean, unequal variance):* Assume that  $\mu_1 = \dots = \mu_P = \mu$ , and  $\sigma_i^2$ 's are different, and consider the problem of estimating  $\mu$ .

Since  $\bar{x}_i \sim \mathcal{N}(\mu, \frac{\sigma_i^2}{N})$ , we obtain an unbiased estimator for  $\mu$

$$\frac{1}{P}[\bar{x}_1 + \dots + \bar{x}_P] \sim \mathcal{N}(\mu, \frac{1}{P^2}[\frac{\sigma_1^2}{N} + \dots + \frac{\sigma_P^2}{N}]). \quad (6)$$

As  $P$  increases, the variance of  $\frac{1}{P}[\bar{x}_1 + \dots + \bar{x}_P]$  decreases. This shows that when there are many populations, (6) gives a very accurate estimate of  $\mu$ .

The above two examples show that with the extra (deterministic) information of equal variance or equal mean, we can reduce the estimation error roughly as  $1/\sqrt{P}$ . That is, the estimation error decreases as the number of population  $P$  increases. The reason for the error reduction is that the extra correlation information enables us to fuse the data from all populations, and it effectively increases the sample size.

In practice, however, it is too strong a statement to claim equal variance or equal mean. Rather, MPME imposes a soft correlation structure on the mean/variance. In particular, MPME imposes a joint prior distribution  $p(\mu, \sigma^2)$  to model the correlation.

### C. Modeling Correlation Among Multiple Populations

By imposing a joint prior distribution  $p(\mu, \sigma^2)$  on  $\mu$  and  $\sigma^2$ , MPME assumes an underlying generative graphical model which describes how the data  $\mathcal{X}_1, \dots, \mathcal{X}_P$  are generated. Assuming further that the prior distribution is parameterized




Fig. 1. Generative graphical models for multiple population Gaussian samples. (a) Probabilistic. (b) Deterministic.

by  $\theta$ ,<sup>3</sup> the graphical model is shown in Fig. 1(a). The graphical model describes that  $\mu_i$ 's and  $\sigma_i^2$ 's are independent samples from the distribution  $p(\mu, \sigma^2 | \theta)$ , and  $\mathcal{X}_i$ 's are conditionally independent samples from the corresponding Gaussian distributions  $\mathcal{N}(\mu_i, \sigma_i^2)$  given  $\mu_i$ 's and  $\sigma_i^2$ 's.

Compared to the traditional approach where  $\mathcal{X}_1, \dots, \mathcal{X}_P$  are independent from each other, the graphical model in Fig. 1(a) asserts that  $\mathcal{X}_1, \dots, \mathcal{X}_P$  are conditionally independent given  $\mu$  and  $\sigma^2$  and that  $\mu$  and  $\sigma^2$  are conditionally independent given  $\theta$ . Therefore, with  $\theta$  unobserved, the populations  $\mathcal{X}_1, \dots, \mathcal{X}_P$  are correlated.<sup>4</sup> This is a key difference between the traditional approach and MPME, it allows MPME to fuse the data from all populations, thus improving estimation accuracy.

It is important to note that in practice, the moments  $\mu$  and  $\sigma^2$  are deterministic fixed quantities given the circuit and the configuration, and are not random variables. For example, considering only V, T dependencies, the  $\mu_i$ 's and  $\sigma_i^2$ 's are deterministic functions of V, T, as shown in Fig. 1(b). The probabilistic generative model in Fig. 1(a) is simply a way to avoid estimating the potentially highly nonlinear functions  $f(\cdot)$  and  $g(\cdot)$ . It replaces the deterministic function of  $\mu_i$ 's and  $\sigma_i^2$ 's with a joint distribution that approximates the correlation defined by  $f(\cdot)$  and  $g(\cdot)$ . However, this is a very mild assumption. The probabilistic modeling not only boosts estimation accuracy, but also provides significant scalability/flexibility compared to direct performance modeling of  $\mu_i$ 's and  $\sigma_i^2$ 's.

The above generative graphical modeling idea can be extended to more general scenarios, including parametric and non-parametric multiple population density estimation problems. For example, consider the parametric density estimation problem where  $x_i$  satisfies the distribution  $p(x_i | \alpha_i)$  parameterized by  $\alpha_i$ . By imposing a joint distribution  $p(\alpha | \theta)$  over  $\alpha_i$ 's, we obtain the generative graphical model in Fig. 2. The 2-step approach in MPME can be similarly applied to this model

<sup>3</sup>Here,  $\theta$  is known as hyperparameters in statistical literatures [1].

<sup>4</sup>We elaborate in Appendix B the correlation induced by applying a (unobserved) prior distribution, and its relationship to traditional concept of the correlation coefficient.




Fig. 2. Generative graphical model for multiple population parametric density estimation problem.

for estimating  $\alpha_i$ 's. However, this is out of the scope of the paper, and we will only focus only on the moment estimation problem.

#### D. Choosing Prior Distributions

Intuitively, the prior distribution for  $\mu_i$ 's and  $\sigma_i^2$ 's, denoted by  $p(\mu_i, \sigma_i^2)$ , describes the belief about the correlation among  $\mu_i$ 's and  $\sigma_i^2$ 's. It is useful to note that the probabilistic models encompass deterministic relationships between parameters at different populations. For example, in Example 1,  $\sigma_1^2 = \dots = \sigma_P^2$  corresponds to a Dirac distribution  $p(\sigma_i^2) = \delta(\sigma_i^2 - \sigma^2)$ , and in Example 2,  $\mu_1 = \dots = \mu_P$  corresponds to a Dirac distribution  $p(\mu_i) = \delta(\mu_i - \mu)$ .

However, in real applications, it is too strong to claim a priori that  $\mu_i$ 's and  $\sigma_i^2$ 's at all populations are the same. Instead, it is often the case that  $\mu_i$ 's and  $\sigma_i^2$ 's at different populations are similar, but not equal, this is often observed in practical analog/mixed-signal circuits, especially in those carefully designed to account for variability. For example, many circuits have compensation loops and self-reconfigurable/self-healing features that cancel out the effects due to certain variability, which effectively pushes  $\mu_i$ 's toward each other. On the other hand, the variance in the circuit performance is usually caused by a small set of parameters (such as critical process parameters, temperature, voltage), and the dependency at different configurations tends to be similar, which effectively pushes  $\sigma_i^2$  toward each other.

Based on the above observation, we consider two candidates for the prior distribution.

1) *Independent Uniform Prior (UNI)*: The first candidate is the uniform prior distribution defined by

$$p(\mu_i, \sigma_i^2) = p(\mu_i | a, b) p(\sigma_i^2 | c, d) \quad (7)$$

where

$$p(\mu_i | a, b) = \begin{cases} \frac{1}{b-a} & \text{if } \mu_i \in [a, b] \\ 0 & \text{otherwise} \end{cases}, \quad (8)$$

$$p(\sigma_i^2 | c, d) = \begin{cases} \frac{1}{d-c} & \text{if } \sigma_i^2 \in [c, d] \\ 0 & \text{otherwise} \end{cases}$$

where  $a, b \in \mathbf{R}, c, d \in \mathbf{R}^+$  are hyperparameters that satisfy  $a \leq b, c \leq d$ .






Fig. 3. Generative graphical model corresponding to uniform prior (UNI).

As is evident from (7),  $\mu_i$  and  $\sigma_i^2$  are independent, and are parameterized by  $(a, b)$  and  $(c, d)$ , respectively. The corresponding generative graphical model is shown in Fig. 3.

The uniform prior is interesting because it has a straightforward interpretation when applied, the process of learning a uniform prior can be thought of as obtaining a bound on the quantities to be estimated, and the process of applying the uniform prior during estimation can be thought of as restricting the estimators to be within the bound defined by  $(a, b)$  and  $(c, d)$ . Details of the derivation are presented in Appendix C.

2) *Normal-Inverse-Chi-Squared Prior (NIX)*: The second candidate is known as the normal-inverse-chi-squared prior defined by

$$p(\mu_i, \sigma_i^2) = p(\mu|\sigma_i^2)p(\sigma_i^2) \quad (9)$$

where

$$\begin{aligned} p(\mu_i|\sigma_i^2) &= \mathcal{N}(\mu_i|\mu_0, \sigma_i^2/\kappa_0) \\ p(\sigma_i^2) &= \chi^{-2}(\sigma_i^2|\nu_0, \sigma_0^2) \end{aligned} \quad (10)$$

where  $\mu_0 \in \mathbf{R}$ ,  $\nu_0, \kappa_0, \sigma_0^2 \in \mathbf{R}^+$  are hyperparameters.

Unlike the independent uniform prior,  $\mu_i$  and  $\sigma_i^2$  are not independent in the normal-inverse-chi-squared prior. The corresponding generative graphical model is shown in Fig. 4.

The normal-inverse-chi-squared prior is particularly useful because it is a conjugate prior, *i.e.*, the posterior distribution  $p(\mu_i, \sigma_i^2|\mathcal{X}_i)$  is also a normal-inverse-chi-squared distribution. It allows for closed-form expressions of the posterior, leading to closed-form expressions of the MAP solution. Therefore, the MAP estimation using this prior is extremely computationally efficient. Details of the derivation are presented in Appendix D.

Similar to the UNI prior, the NIX prior also has a straightforward interpretation, it is equivalent to increasing the effective number of samples by adding fake data samples that reflect the prior. As is shown in Appendix D, the MAP mean estimation is equivalent to adding  $\kappa_0$  data samples with mean  $\mu_0$ , and the MAP variance estimation is equivalent to adding  $\nu_0$  data samples with variance  $\sigma_0^2$ . Therefore, if  $\kappa_0$  and  $\nu_0$  are large, we effectively have more samples, and that lead to

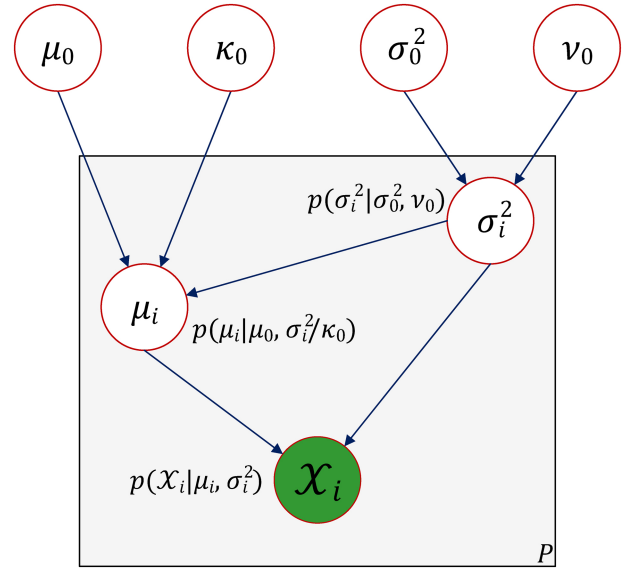


Fig. 4. Generative graphical model corresponding to normal-inverse-chi-squared prior (NIX).

more accurate estimation. As will be illustrated on a dataset in Sec. V-A, MPME can significantly increase the number of effective samples.

It is also interesting to note that both prior distributions can converge to the Dirac distribution  $p(\sigma_i^2) = \delta(\sigma_i^2 - \sigma^2)$  in Example 1 and  $p(\mu_i) = \delta(\mu_i - \mu)$  in Example 2. For the uniform prior, the Dirac prior may be obtained as  $|b - a| \rightarrow 0$  and  $|d - c| \rightarrow 0$ . For the normal-inverse-chi-squared prior, the Dirac prior may be obtained as  $\kappa_0 \rightarrow \infty$  and  $\nu_0 \rightarrow \infty$ .

### E. Learning the Prior Distribution

In MPME, the first step is to learn a prior distribution from data collected at all populations. We employ the maximum likelihood approach to learn the prior  $p(\mu_i, \sigma_i|\theta)$ , where  $\theta$  are hyper-parameters of the prior distribution. For example,  $\theta = [a, b, c, d]$  for the UNI prior, and  $\theta = [\kappa_0, \mu_0, \nu_0, \sigma_0^2]$  for the NIX prior.

The optimization problem can be formulated as

$$\underset{\theta}{\text{maximize}} \quad p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta) \quad (11)$$

where  $p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta)$  is the likelihood function. We may either use a nonlinear optimizer to solve for the optimal  $\theta$ , or we may derive closed-form solutions by solving

$$\frac{d}{d\theta} p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta) = 0. \quad (12)$$

To compute the likelihood function  $p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta)$ , we resort to the graphical model and integrate out  $\mu$  and  $\sigma$ , that is

$$\begin{aligned} & p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta) \\ &= \int_{\mu, \sigma} p(\mathcal{X}_1, \dots, \mathcal{X}_P|\mu, \sigma^2) p(\mu, \sigma^2|\theta) d\mu d\sigma. \end{aligned} \quad (13)$$

The integral (13) can be computed by numerical integration, or we may derive its closed-form expression for special prior distributions. The derivations of  $p(\mathcal{X}_1, \dots, \mathcal{X}_P|\theta)$  for the UNI prior and NIX prior are presented in Appendix C and Appendix D, respectively.

**Algorithm 1** Multiple Population Moment Estimation**Inputs:**  $\mathcal{X}_1, \dots, \mathcal{X}_P$ .**Outputs:**  $(\mu_i, \sigma_i^2), i = 1, \dots, P$ .

- 1: Solve maximize  $p(\mathcal{X}_1, \dots, \mathcal{X}_P | \theta)$  ((11)) for  $\theta$
- 2: **for**  $i = 1 \rightarrow P$  **do**
- 3: Solve maximize  $p(\mu_i, \sigma_i^2 | \mathcal{X}_i)$  ((14)) for  $(\mu_i, \sigma_i^2)$
- 4: **end for**

*F. Maximum a Posteriori Estimation of  $\mu$  and  $\sigma$* 

Once the prior  $p(\mu_i, \sigma_i^2 | \theta)$  is learned, MAP estimation can be applied to obtain a point estimate of  $\mu_i$ 's and  $\sigma_i^2$ 's. MAP formulation searches for the values of  $\mu_i$ 's and  $\sigma_i^2$ 's that maximize the posterior distribution, *i.e.*, it solves

$$\underset{\mu_i, \sigma_i^2}{\text{maximize}} \quad p(\mu_i, \sigma_i^2 | \mathcal{X}_i, \theta). \quad (14)$$

According to Bayes' rule

$$p(\mu_i, \sigma_i^2 | \mathcal{X}_i, \theta) \propto p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i, \sigma_i^2 | \theta) \quad (15)$$

where  $p(\mu_i, \sigma_i^2 | \theta)$  is learned as described in Sec. III-E, and

$$\begin{aligned} & p(\mathcal{X}_i | \mu_i, \sigma_i^2) \\ &= \prod_{j=1}^{N_i} \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \left\{ -\frac{(x_{i,j} - \mu_i)^2}{2\sigma_i^2} \right\} \\ &= \frac{\sigma_i^{-N_i}}{(2\pi)^{N_i/2}} \exp \left\{ -\frac{N_i(\bar{x}_i - \mu_i)^2 + (N_i - 1)S_i}{2\sigma_i^2} \right\} \end{aligned} \quad (16)$$

because  $x_{i,j}, j = 1, \dots, N_i$  are independent samples from the Gaussian distribution  $\mathcal{N}(\mu_i, \sigma_i^2)$ .

The details of the MAP estimation for the UNI prior and NIX prior can be found in Appendix C and Appendix D, respectively.

*G. MPME Algorithm*

Summarizing Sec. III-E and Sec. III-F, the MPME algorithm is shown in Algorithm 1.

## IV. REMARKS

*A. Practical Implementation*

It should be noted that the optimization problems in Algorithm 1 may not be convex, and may have multiple local optimal points. There is no guarantee that the numerical algorithm will find the global optima. However, since initial guesses can be estimated from the same data, the optimizer has a good guess to start with, and is less affected by local optimal points.

To alleviate the computational cost associated with solving the optimization problems, we may impose an empirical prior distribution, instead of learning one from data. For example, experienced designers may have a good idea of the range of  $\sigma_i^2$  at each population (*e.g.*, either from results of test chips or previous products), in this case, a uniform prior for  $\sigma_i^2$ 's can be asserted. However, empirical priors should be used with great caution, since it may incur unexpected bias. To be less biased, one may apply cross-validation [14] to check the validity of the empirical prior.

*B. Connections to Empirical Bayes Estimators*

The ideas presented in this paper are similar to the philosophy of a class of Bayesian estimators, called empirical Bayes estimators (EB) [15]. EB applies Bayes' rule to obtain either a point estimation or a posterior distribution of the parameters to be estimated. Unlike standard Bayesian methods that specify an arbitrary prior, EB learns the prior distribution from data. In particular, if a Gaussian prior is used for the mean, EB gives the so called James–Stein estimator [16] for the mean.

Particularly, a nice feature of the James–Stein estimator is that it is superior to the sample mean estimate, in the sense that the expected sum of mean square error of  $\mu_i$ 's at all populations is smaller than that of the sample mean estimator, that is

$$E \left\{ \sum_{i=1}^P (\mu_i - \mu_i^{JS})^2 \right\} < E \left\{ \sum_{i=1}^P (\mu_i - \bar{x}_i)^2 \right\} \quad (17)$$

where  $\mu_i$  is the actual mean,  $\mu_i^{JS}$  is the James–Stein estimator and  $\bar{x}_i$  is the sample mean. One can show that if the Gaussian prior on  $\mu_i$ 's is used in our method, we obtain an estimator very similar to the James–Stein estimator, and (17) still holds.

Unlike the James–Stein estimator, our method allows for more general prior distributions. In particular, we have derived the case for the UNI prior and the NIX prior. We will show in Sec. V that our method can significantly out-perform sample mean/variance estimators.

*C. Other Prior Distributions*

The choice of the prior distribution largely depends on its modeling capability as well as the computational tractability. In terms of the modeling capability, both UNI and NIX prior can model the closeness of mean/variance across populations pretty well. On the other hand, the likelihood functions using both priors have (semi-)analytical expressions, and the MAP estimation for both priors are extremely efficient due to the simplicity of the priors.

In addition to the UNI and NIX prior mentioned in Sec. III-D, one can apply other prior distributions and follow the same procedure of MPME. Different prior distributions encode different information and therefore encourage solutions of particular structures. For example, the Laplace distribution is a prior that encourages sparsity in the solution [1]. More generally, one can use a mixture of Gaussian to approximate any distribution to arbitrarily accurately. In this paper, we exploit the underlying structure that the mean/variance values cluster together, and find UNI and NIX prior are good enough for that purpose.

Using more complicated prior distributions also raises the question of computational tractability. For example, consider the problem of MAP estimation for the mean value of a Gaussian distribution. If we use a mixture of two Gaussians as the prior for the mean, then the posterior distribution for the mean is again a mixture of two Gaussians. The MAP estimation is in general no longer a convex optimization problem (as is the case for UNI and NIX priors), and therefore we lose the theoretical tractability for the MAP estimation. In addition, it is not hard to see that the parameter learning

problem will become more complicated and computationally more expensive.

#### D. Non-Gaussian Distributions and Higher-Order Moments

The discussion in Sec. III focuses on the case where the distribution at each population is Gaussian. This is an engineering assumption that is often used in practice. And with very few samples (*e.g.*, five), it is impossible to obtain an accurate estimation of the moments/distributions without extra knowledge about the problem.

For non-Gaussian distributions, the distributions of mean/variance are not Gaussian and  $\chi^2$ , and therefore the derivations need to be modified. The shape of the mean/variance distribution, however, may not have a closed form expression, and need to be treated on a case-by-case basis.

On the other hand, it is straightforward to extend MPME to non-Gaussian distributions if they have a limited number of parameters or sufficient statistic (*e.g.*, the exponential family). For many distributions in the exponential family, the sufficient statistic include the first two moments of  $x$  or  $\ln x$ . The adaption of MPME to these distributions include the choice of prior and the derivation of the posterior distribution. This is relatively straightforward because it is well established [1] that all members of the exponential family have conjugate priors (which lead to closed-form posterior distributions and make the MAP estimation procedure efficient).

In rare cases in circuit validation, one might also want to estimate higher-order moments (such as skewness and kurtosis). In theory, MPME may be applied to estimate higher-order moments, but with small sample sizes, the estimation error may still be too large for the method to be practical. Indeed, to apply MPME, one needs to further define  $p(x|m_i)$  where  $m_i$  is the  $i$ th moment. The rest of the algorithm can be derived by following the steps in Sec. III. This means that we need a way to convert a series of moments  $m_1, m_2, \dots$  to a probability distribution. This is a very hard problem, and deserves a paper by itself, Li *et al.* [17] proposed a solution which might be used together with the MPME algorithm.

An engineering solution to this problem is to assert  $p(x|m_i)$  is Gaussian with mean  $m_i$ . In this case, MPME can be readily applied by treating  $x^i$  as samples. We have used this method for the variance estimation problem and compared it to the rigorous treatment using  $\chi^2$  distribution. The empirical result shows that this method is not too much worse than the rigorous method.

#### E. Potential Limitations

Although our method may obtain a theoretically better overall estimate according to conclusions such as (17), it can be the case, theoretically, that for a specific population, our method introduces a large bias.

As an extreme example, consider 100 populations, each with one observation, and  $\mu_1 = \dots = \mu_{99} = 0, \mu_{100} = 1, \sigma_1^2 = \dots = \sigma_{100}^2 = 1$ . Effectively, our method will shrink the estimated mean toward 0. Therefore, for the 100th population, the bias can be large.

However, due to the reasons mentioned in Sec. III-D, such extremely pathological cases are unlikely to happen. Even if it happens, the outliers can be easily identified in a preprocessing step, and therefore, accuracy will not be compromised by outliers.

#### F. General Guideline of Applying MPME

There are two key questions that one may ask before applying MPME.

- 1) When is MPME (significantly) better than the sample estimators?
- 2) Which prior (NIX or UNI) should be used in MPME?

While it is hard to give a definite answer and rigorous theoretical analysis, we provide several general guidelines that help answering these questions.

First, MPME is significantly better than sample estimators only if the sample size is small. From (2), the error of sample estimators decreases as  $N$  increases. Therefore, if the sample size is large, sample estimators are good enough, and the benefit brought by MPME is negligible.

Second, MPME is significantly better than sample estimators only if the variance is large. Similarly, from (2), the error of sample estimators decreases as the variance  $\sigma_i^2$  decreases. Therefore, if  $\sigma_i$ 's are small, sample estimators also give very accurate results, and MPME estimation will be very similar to that of the sample estimators.

Third, obvious outliers need to be pruned in MPME. When using MPME, it is helpful to first inspect how the sample mean/variance spread. If there are obvious population outliers, they need to be removed. As explained in Sec. IV-E, the outliers are unlikely to be correlated to other populations, and therefore including them in MPME could lead to worse results.

Fourth, empirically, the NIX prior is usually better for the overall error than the UNI prior, and the UNI prior is more consistent across populations than the NIX prior. We can explain this empirical result by inspecting the MAP estimation equations of mean for the UNI prior and NIX prior.

Equation (42) shows that NIX prior pulls the mean estimate toward the prior mean  $\mu_0$ , which is likely to be close to the mean across all populations. Therefore, for a specific population, if  $\mu_i$  is close to the overall mean  $\mu_0$ , then NIX prior will give almost perfect estimation. However, the NIX prior can lead to large bias if  $\mu_i$  is far from  $\mu_0$ .

In contrast, equation (32) shows that applying the UNI prior is equivalent to applying a bound  $[a, b]$  on the sample estimator. Since  $a, b$  are learned from data, they usually cover the range of the mean values of every population. This means that no matter where the  $\mu_i$  is, the accuracy improvement tend to be similar because the probability that sample mean is out of the range of  $[a, b]$  is low.

## V. EXPERIMENTAL RESULTS

In this section, we illustrate the proposed method, MPME, on a few synthetic examples as well as an industrial example of a commercial high-speed I/O link. By the synthetic examples, we demonstrate that MPME can achieve much more accuracy







Fig. 5. Comparison of sample estimators and MPME ( $P = 20$ , Example 1). (a)  $\epsilon_{\mu}$  vs  $N$ . (b)  $\epsilon_{\sigma^2}$  vs  $N$ .

Fig. 6. Comparison of sample estimators and MPME ( $N = 5$ , Example 1). (a)  $\epsilon_{\mu}$  vs  $P$ . (b)  $\epsilon_{\sigma^2}$  vs  $P$ .

compared to traditional methods such as sample mean and sample variance estimator, and we conclude empirically the scenarios under which MPME may significantly outperform traditional methods. By the industrial example, we illustrate that MPME can increase validation quality and potentially reduce test time by more than 2X. All the numerical experiments are carried out using multiple threads on a Linux machine with Intel Xeon E5-2430 CPUs capable of running 24 threads in parallel and 64 GB of total physical memory.

#### A. Synthetic Examples 1

In this example, the data is generated as follows.

- 1) Determine  $P$  (the number of populations) and  $M$  (the number of independent trials).
- 2) Choose  $N_1 = \dots = N_P = N$  (i.e., all populations have same number of independent samples) and determine  $N$  (the number of samples at each population).
- 3) Choose  $\mu_i$ 's to be equally spaced over  $[9.5, 10.5]$ , i.e.,  $\mu_i = 9.5 + \frac{1}{P-1}(i-1)$ ,  $i = 1, \dots, P$ .
- 4) Choose  $\sigma_i$ 's to be equally spaced over  $[0.95, 1.05]$ , i.e.,  $\sigma_i = 0.95 + \frac{0.1}{P-1}(i-1)$ ,  $i = 1, \dots, P$ .

- 5) For  $i = 1, \dots, P$ , draw  $x_{i,j}$ ,  $j = 1, \dots, N$  from  $\mathcal{N}(\mu_i, \sigma_i^2)$ .

In our experiments, we choose  $M = 500$ , i.e., we generate 500 independent random trials from the same distribution. To compare MPME against sample estimators, we compute the average error across populations, defined by

$$\begin{aligned} \epsilon_{\mu} &= \frac{1}{P} \sum_{i=1}^P \sqrt{\frac{1}{M} \sum_{j=1}^M (\mu_i - \hat{\mu}_{i,j})^2} \\ \epsilon_{\sigma^2} &= \frac{1}{P} \sum_{i=1}^P \sqrt{\frac{1}{M} \sum_{j=1}^M (\sigma_i^2 - \hat{\sigma}_{i,j}^2)^2} \end{aligned} \quad (18)$$

where  $\hat{\mu}_{i,j}$  and  $\hat{\sigma}_{i,j}^2$  are the estimated mean/variance for the  $i$ th population in the  $j$ th trial.

We apply three methods (sample estimator, MPME with UNI prior and MPME with NIX prior) to this data set with varying  $P$  and  $N$  chosen from  $P \in \{5, 10, 15, 20, 30, 40, 50, 100\}$  and  $N \in \{5, 11, 15, 21, 31, 41, 51, 101\}$ . Under all combinations of  $P$  and  $N$ , we observe that MPME always out-performs the sample estimator in terms of accuracy.

Out of all combinations of  $P$  and  $N$ , we discuss the results of two special cases  $P = 20$  and  $N = 5$ , to illustrate how the





Fig. 7. Histogram of  $\kappa_0$  and  $\nu_0$ , ( $N = 5, P = 20$ , Example 1). (a)  $\kappa_0$ . (b)  $\nu_0$ .

accuracy of MPME estimation improves with the number of samples  $N$  and the number of populations  $P$ .

Fig. 5 shows the error of three methods for different values of  $N$  when  $P = 20$ . It can be observed that as  $N$  becomes large, the error of three methods all converge to a small value, and MPME does not present much advantage over the sample estimators. However, when  $N$  is extremely small, MPME obtains significantly better accuracy.

Fig. 6 shows the error of three methods for different values of  $P$  when  $N = 5$ . It can be observed that as  $P$  becomes large, the error of MPME decreases roughly as  $1/\sqrt{P}$ , while the error of the sample estimators stays the same. The reason is that sample estimators treat each population independently, while MPME exploits the joint information in the dataset to improve the estimation accuracy at individual populations.

As mentioned in Sec. III-D2, the application of the NIX prior can be interpreted as increasing the effective number of samples by  $\kappa_0$  (for mean estimation) and  $\nu_0$  (for variance estimation). Fig. 7 shows the histogram of  $\kappa_0$  and  $\nu_0$  over 500 trials for the case ( $N = 5, P = 20$ ). The mean values of  $\kappa_0$  and  $\nu_0$  are 32.9 and 79.4, respectively. This means, effectively, MPME increases the number of samples  $N = 5$  to around 40 and 80, this significantly improves the accuracy of the estimation.


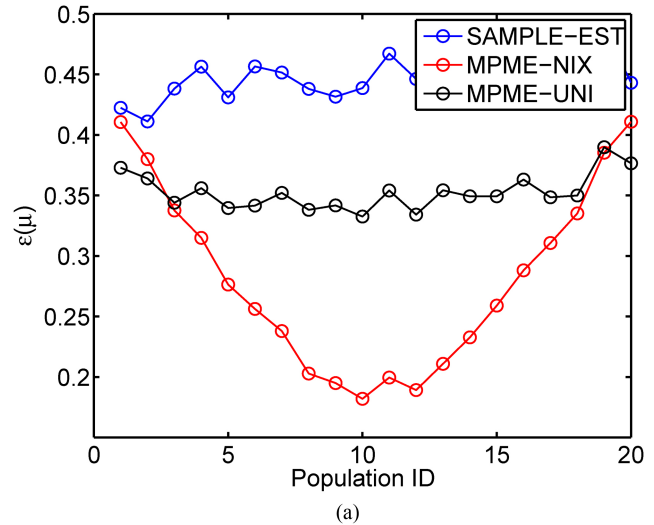


Fig. 8. Comparison of sample estimators and MPME ( $N = 5, P = 20$ , Example 1). (a)  $\epsilon_\mu$  across populations. (b)  $\epsilon_{\sigma^2}$  across populations.

From Figs. 5 and 6, we also find that for the MPME method, the NIX prior is usually better than the UNI prior, in terms of  $\epsilon_\mu$  and  $\epsilon_{\sigma^2}$ . While this shows that NIX prior might be preferred, we emphasize that the UNI prior could lead to a better accuracy for a particular population.

For example, Fig. 8 shows the average error for each of the 20 populations for the setting ( $N = 5, P = 20$ ). For the estimation of  $\sigma_i^2$ 's, MPME-NIX is consistently better than MPME-UNI. However, for the estimation of  $\mu_i$ 's, although the NIX prior leads to a smaller error for most of the populations, the UNI prior does better at the populations that have extreme  $\mu_i$  values. Intuitively, during the second step (MAP) in MPME, the UNI prior applies lower/upper bounds on the estimated mean, and the NIX prior pulls the estimated mean toward the joint mean (across populations). Therefore, if the population mean is close to the overall mean, NIX prior leads to a better estimation. On the other hand, if the population mean is far from the overall mean (*e.g.*, at extreme corners), UNI prior will be better. In both cases, however, MPME-UNI and MPME-NIX are always better than the sample estimators.



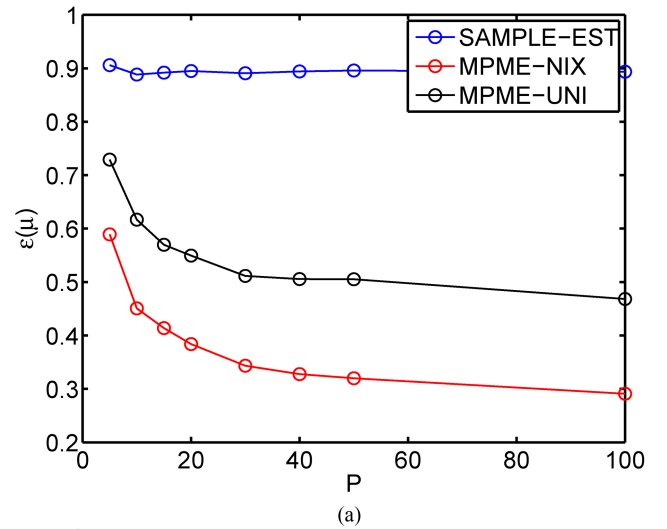



Fig. 9. Comparison of sample estimators and MPME ( $P = 20$ , Example 2). (a)  $\epsilon_\mu$  vs  $N$ . (b)  $\epsilon_{\sigma^2}$  vs  $N$ .

Fig. 10. Comparison of sample estimators and MPME ( $N = 5$ , Example 2). (a)  $\epsilon_\mu$  vs  $P$ . (b)  $\epsilon_{\sigma^2}$  vs  $P$ .

### B. Synthetic Examples 2

In this example, we use almost the same setting as the previous one, except for  $\sigma_i$  values for the  $i$ th population. We choose  $\sigma_i$  to be equally spaced over  $[1.9, 2.1]$ , *i.e.*,  $\sigma_i = 1.9 + \frac{0.2}{P-1}(i-1)$ , *i.e.*, twice of  $\sigma_i$  in the previous example.

Similar trends are observed as in the previous example, as shown in Figs. 9 and 10. However, compared to the previous example, it is worthwhile to note that MPME obtains relatively more reduction in error over sample estimators. The reason for this is that when the variance at each population is smaller, the data show less uncertainty/randomness. For example, the sample mean estimator has a confidence interval proportional to  $\sigma_i$  and the sample variance estimator has a confidence interval proportional to  $\sigma_i^2$ . Therefore, when  $\sigma_i$ 's are small, sample estimators achieve relatively good accuracy, and MPME provides less improvement. However, when  $\sigma_i$ 's are large, MPME beats sample estimators significantly, by exploiting the collective information gathered from multiple populations.

### C. Validation of High-Speed I/O Link

In I/O link validation, one critical performance metric is Bit-Error-Ratio (BER). For the state-of-the-art high-speed

links, the BER is extremely small. For example, in the latest PCIe specification [11],  $\text{BER}_{\text{spec}} = 10^{-12}$  with 8Gb/sec data rate. This makes BER measurement a very time-consuming process. An alternative is to measure the eye width and eye height [*a.k.a.*, time margin (TM) and voltage margin (VM), respectively] of the eye diagram at the receiver, which can be converted to BER under reasonable assumptions. Margin measurement, although much faster than direct BER measurement, is still expensive in terms of time and cost. For a limited time period, only a small number of data can be measured for each configuration.

In this example, we have measured the time margin of 50 dies (randomly sampled) for eight different configurations. (Note that we measured 50 dies simply for the purpose of validating our algorithm. In practice, only about five dies might be measured.) The mean and standard deviation at different configurations are shown in Fig. 11. We have also observed from the histogram that the distribution of time margin can be well approximated by Gaussian distributions.

To compare the results of MPME and sample estimators, we take  $N$  samples of data for each configuration from the






Fig. 11. Mean and standard deviation at eight configurations.

50 measurements, and apply both methods. We repeat this experiment for 500 times, and compare the statistics of  $\epsilon_\mu$  and  $\epsilon_{\sigma^2}$ . This is also known as bootstrap in the statistics literature [1].

The results for  $\epsilon_\mu$  and  $\epsilon_{\sigma^2}$  for different values of  $N$  are plotted in Fig. 12. Similar to the synthetic examples, it is observed that when the sample size is small, sample estimators are much less accurate than MPME, thus may lead to unreliable validation conclusions.

Besides accuracy improvement, Fig. 12 shows another practical implication of MPME, for the same overall accuracy, MPME requires much less samples than the sample estimators. In this particular example, MPME-NIX would need just about 50% samples than the sample estimators, in order to obtain the same accuracy. It directly implies less validation time, and thus faster product time-to-market.

Fig. 13 shows a detailed comparison of  $\epsilon_\mu$  and  $\epsilon_{\sigma^2}$  for all eight populations, and the results confirm our conclusion drawn from synthetic examples, *i.e.*, if the variance of the population is larger, MPME is more effective in reducing the error. In particular, the population #1 in this example has the largest variance, and MPME has the most significant error reduction over the sample estimators.

## VI. CONCLUSION

In this paper, we have proposed MPME, an efficient method for estimating moments (mean and variance) of multiple populations. The key difficulty we try to address is the problem of extremely small sample size, which is commonly seen in circuit validation. MPME alleviates this problem by considering samples obtained from many populations, which in practice can refer to different corners and configurations. MPME leverages data from all populations to improve the estimation accuracy for each population, and the method fits nicely under the hierarchical Bayesian framework. We validated MPME on several datasets, including measurement of a commercial I/O link. We show that MPME is consistently better than the sample mean/variance estimators, and can





Fig. 12. Comparison of sample estimators and MPME. (a)  $\epsilon_\mu$  vs  $N$ . (b)  $\epsilon_{\sigma^2}$  vs  $N$ .

achieve up to  $2\times$  average accuracy improvement. Furthermore, the accuracy improvement can also be equivalently translated to a potentially large test/validation time reduction.


## APPENDIX A PROBABILISTIC GRAPHICAL MODELS

Probabilistic graphical models use graphs (directed or undirected) to describe multivariate probability distributions and the probabilistic structures (*e.g.*, conditional independences). For the interest of this paper, we only discuss the concepts and notations relevant to the MPME method. For more details about graphical models, we refer the readers to two excellent books [1], [18].


In a graphical model, each node represents a random variable (or a set of random variables), and the edges represent the probabilistic relationships between these variables. In a directed graphical model, the edges can be interpreted as the dependency among variables. For example, a tree-like graphical model shown in Fig. 14(a) describes a joint probability distribution over  $\theta, \alpha_1, \dots, \alpha_P$  as

$$p(\theta, \alpha_1, \dots, \alpha_P) = p(\theta)p(\alpha_1|\theta) \cdots p(\alpha_P|\theta) \quad (19)$$





(a)



(b)

Fig. 13. Comparison of sample estimators and MPME. (a)  $\epsilon_\mu$  across populations. (b)  $\epsilon_{\sigma^2}$  across populations.

which encodes the conditional independence

$$\alpha_1 \perp\!\!\!\perp \alpha_2 \perp\!\!\!\perp \dots \perp\!\!\!\perp \alpha_P | \theta. \quad (20)$$

Here, the notation  $(A \perp\!\!\!\perp B|C)$  means that  $A$  and  $B$  are conditionally independent given  $C$ .

To simplify the graph notation, we use the plate notation to compactly represent multiple nodes. In the plate notation, we draw a single representative node and then surround it with a box labeled with  $P$  indicating there are  $P$  nodes of this kind. Using the plate notation, the graphical model in Fig. 14(a) can be compactly represented by Fig. 14(b).

## APPENDIX B

### CORRELATION INDUCED BY IMPOSING PRIOR DISTRIBUTION

In this section, we explain why random variables that are conditionally independent are correlated, and how that relates to the traditional concept of correlation coefficient.




Fig. 14. Simple tree-like graphical model. (a) Simple directed graphical model. (b) Plate notation.

For simplicity, we consider the graphical model in Fig. 14(a) where there are only two leaf nodes  $\alpha_1$  and  $\alpha_2$ . We further assume that

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} | \theta \sim \mathcal{N} \left( \begin{bmatrix} 1 \\ 1 \end{bmatrix} \theta, \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} \right) \quad (21)$$

where  $\sigma$  is known. (21) implies that  $\alpha_1$  and  $\alpha_2$  are conditionally independent given  $\theta$ .

However, since  $\theta$  is not observed, we need to study the marginal distribution of  $(\alpha_1, \alpha_2)$  to compute their correlation coefficient. To compute the marginal distribution, we assume that  $\theta$  also follows a Gaussian distribution:

$$\theta \sim \mathcal{N}(\mu_0, \sigma_0^2). \quad (22)$$

Then, we can compute

$$\begin{aligned} p(\alpha_1, \alpha_2) &= \int_{\theta} p(\alpha_1, \alpha_2, \theta) d\theta \\ &= \int_{\theta} p(\alpha_1, \alpha_2 | \theta) p(\theta) d\theta. \end{aligned} \quad (23)$$

With some algebraic manipulation, we obtain

$$\begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix} \sim \mathcal{N} \left( \begin{bmatrix} 1 \\ 1 \end{bmatrix} \mu_0, \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2 \end{bmatrix} + \sigma_0^2 \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \right). \quad (24)$$

Therefore, the correlation coefficient between  $\alpha_1$  and  $\alpha_2$  is

$$\rho = \frac{\sigma_0^2}{\sigma^2 + \sigma_0^2} = \frac{1}{1 + \frac{\sigma^2}{\sigma_0^2}}. \quad (25)$$

From (25), we can conclude that if  $\sigma_0$  is large and  $\sigma$  is small, strong correlation will exhibit between  $\alpha_1$  and  $\alpha_2$ . Furthermore, consider the case where  $\sigma$  is fixed. Hence, if we have a strong prior, *i.e.*,  $\sigma_0 \rightarrow 0$ , then  $\alpha_1$  and  $\alpha_2$  will show no correlation (and indeed are independent). If, however, we have a weak prior, *i.e.*,  $\sigma_0 \rightarrow \infty$ , then strong correlation exists between  $\alpha_1$  and  $\alpha_2$ . In MPME, we further exploit the fact that the mean values are close to each other, *i.e.*,  $\sigma$  is small. It is evident from (25) that this also advocates strong correlation across different populations. Similar arguments can be made for the case of multiple populations.

In general, for a graphical model shown in Fig. 1(a) with nontrivial conditional probability distributions,  $\alpha_i$ 's are correlated [18].

APPENDIX C  
LEARNING UNIFORM PRIOR AND MAP USING  
UNIFORM PRIOR

A. Learning Hyperparameters Using MLE

According to the graphical model for the UNI prior in Fig. 3, the following conditional independence relationships are satisfied

$$\begin{aligned} &(\mu_i \perp\!\!\!\perp \sigma_i^2 | a, b, c, d), \\ &(\mu_1 \perp\!\!\!\perp \dots \perp\!\!\!\perp \mu_P | a, b), \\ &(\sigma_1^2 \perp\!\!\!\perp \dots \perp\!\!\!\perp \sigma_P^2 | c, d), \\ &(\mathcal{X}_1 \perp\!\!\!\perp \dots \perp\!\!\!\perp \mathcal{X}_P | \boldsymbol{\mu}, \boldsymbol{\sigma}^2). \end{aligned} \quad (26)$$

Applying (26), (13) can be simplified

$$\begin{aligned} &p(\mathcal{X}_1, \dots, \mathcal{X}_P | \boldsymbol{\theta}) \\ &= \int_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} p(\mathcal{X}_1, \dots, \mathcal{X}_P | \boldsymbol{\mu}, \boldsymbol{\sigma}^2) p(\boldsymbol{\mu}, \boldsymbol{\sigma}^2 | \boldsymbol{\theta}) d\boldsymbol{\mu} d\boldsymbol{\sigma}^2 \\ &= \int_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} \left( \prod_{i=1}^P p(\mathcal{X}_i | \mu_i, \sigma_i^2) \right) \left( \prod_{i=1}^P p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) \right) d\boldsymbol{\mu} d\boldsymbol{\sigma}^2 \quad (27) \\ &= \prod_{i=1}^P \int_{\mu_i, \sigma_i^2} p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) d\mu_i d\sigma_i^2 \\ &= \prod_{i=1}^P \int_{\mu_i, \sigma_i^2} p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i | a, b) p(\sigma_i^2 | c, d) d\mu_i d\sigma_i^2. \end{aligned}$$

For the UNI prior, we have  $p(\mu_i, \sigma_i^2 | a, b, c, d) = p(\mu_i | a, b) p(\sigma_i^2 | c, d)$ , that is

$$p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) = \begin{cases} \frac{1}{b-a} \frac{1}{d-c}, & \text{if } a \leq \mu_i \leq b, c \leq \sigma_i^2 \leq d, \\ 0, & \text{otherwise.} \end{cases} \quad (28)$$

Inserting (28) and (16) into (27), we need to compute for each  $i$

$$\begin{aligned} &\int_0^\infty \int_{-\infty}^\infty p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i, \sigma_i^2 | a, b, c, d) d\mu_i d\sigma_i^2 \\ &\propto \frac{1}{b-a} \frac{1}{d-c} \times \\ &\left[ \begin{aligned} &Q_{N_i-3} \left( \frac{(b-\bar{x}_i)\sqrt{N_i(N_i-3)}}{\sqrt{(N_i-1)S_i}}, 0; 0, \sqrt{\frac{(N_i-1)S_i}{c}} \right) \\ &- Q_{N_i-3} \left( \frac{(a-\bar{x}_i)\sqrt{N_i(N_i-3)}}{\sqrt{(N_i-1)S_i}}, 0; 0, \sqrt{\frac{(N_i-1)S_i}{c}} \right) \\ &- Q_{N_i-3} \left( \frac{(b-\bar{x}_i)\sqrt{N_i(N_i-3)}}{\sqrt{(N_i-1)S_i}}, 0; 0, \sqrt{\frac{(N_i-1)S_i}{d}} \right) \\ &- Q_{N_i-3} \left( \frac{(a-\bar{x}_i)\sqrt{N_i(N_i-3)}}{\sqrt{(N_i-1)S_i}}, 0; 0, \sqrt{\frac{(N_i-1)S_i}{d}} \right) \end{aligned} \right]. \end{aligned} \quad (29)$$

If  $\phi(\cdot)$  and  $\Phi(\cdot)$  denotes the PDF and CDF of standard normal distribution, respectively, then  $Q_f(t, \delta; 0, R)$  is defined as

$$Q_f(t, \delta; 0, R) = \int_0^R \frac{\sqrt{2\pi} y^{f-1} \phi(y)}{\Gamma(\frac{f}{2}) 2^{\frac{f-2}{2}}} \Phi\left(\frac{ty}{\sqrt{f}} - \delta\right) dy \quad (30)$$

which can be solved by repeated integration by parts to yield closed form solutions [19].

B. MAP Estimation

For uniform priors of  $\mu_i$  and  $\sigma_i^2$ , the right-hand side of (15) is

$$\frac{1}{b-a} \frac{1}{d-c} p(\mathcal{X}_i | \mu_i, \sigma_i^2), \quad \text{if } \mu_i \in [a, b] \text{ and } \sigma_i^2 \in [c, d]. \quad (31)$$

Therefore, MAP is equivalent to maximum likelihood estimation on the support  $\mu_i \in [a, b]$  and  $\sigma_i^2 \in [c, d]$ . The solution is simply

$$\mu_{i,MAP} = \begin{cases} a & \text{if } \mu_{i,MLE} < a \\ \mu_{i,MLE} & \text{if } a \leq \mu_{i,MLE} \leq b \\ b & \text{if } \mu_{i,MLE} > b \end{cases} \quad (32)$$

$$\sigma_{i,MAP}^2 = \begin{cases} c & \text{if } \sigma_{i,MLE} < c \\ \sigma_{i,MLE}^2 & \text{if } c \leq \sigma_{i,MLE} \leq d \\ d & \text{if } \sigma_{i,MLE} > d \end{cases} \quad (33)$$

where  $\mu_{i,MLE}$  and  $\sigma_{i,MLE}^2$ <sup>5</sup> are equal to the sample mean and sample variance estimators, respectively, [20].

APPENDIX D

LEARNING NIX PRIOR AND MAP USING NIX PRIOR

A. Learning Hyperparameters Using MLE

According to the graphical model for the NIX prior in Fig. 4, the following conditional independence relationships are satisfied

$$\begin{aligned} &(\sigma_1^2 \perp\!\!\!\perp \dots \perp\!\!\!\perp \sigma_P^2 | \nu_0, \sigma_0^2), \\ &(\mu_1 \perp\!\!\!\perp \dots \perp\!\!\!\perp \mu_P | \kappa_0, \mu_0, \nu_0, \sigma_0^2), \\ &(\mathcal{X}_1 \perp\!\!\!\perp \dots \perp\!\!\!\perp \mathcal{X}_P | \boldsymbol{\mu}, \boldsymbol{\sigma}^2). \end{aligned} \quad (34)$$

Applying (34), (13) can be simplified

$$\begin{aligned} &p(\mathcal{X}_1, \dots, \mathcal{X}_P | \boldsymbol{\theta}) \\ &= \int_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} p(\mathcal{X}_1, \dots, \mathcal{X}_P | \boldsymbol{\mu}, \boldsymbol{\sigma}^2) p(\boldsymbol{\mu}, \boldsymbol{\sigma}^2 | \boldsymbol{\theta}) d\boldsymbol{\mu} d\boldsymbol{\sigma}^2 \\ &= \int_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} \left( \prod_{i=1}^P p(\mathcal{X}_i | \mu_i, \sigma_i^2) \right) \left( \prod_{i=1}^P p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) \right) d\boldsymbol{\mu} d\boldsymbol{\sigma}^2 \quad (35) \\ &= \prod_{i=1}^P \int_{\mu_i, \sigma_i^2} p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) d\mu_i d\sigma_i^2. \end{aligned}$$

For the NIX prior, we have  $p(\mu_i, \sigma_i^2 | \kappa_0, \mu_0, \nu_0, \sigma_0^2) = p(\sigma_i^2 | \nu_0, \sigma_0^2) p(\mu_i | \sigma_i^2, \mu_0, \kappa_0, \nu_0)$ , that is

$$\begin{aligned} &p(\mu_i, \sigma_i^2 | \boldsymbol{\theta}) \\ &= \mathcal{N}(\mu_i | \mu_0, \sigma_i^2 / \kappa_0) \chi^{-2}(\sigma_i^2 | \nu_0, \sigma_0^2) \\ &= \frac{\sigma_i^{-\nu_0-3}}{Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)} \exp\left\{-\frac{\nu_0 \sigma_0^2 + \kappa_0 (\mu_i - \mu_0)^2}{2\sigma_i^2}\right\} \end{aligned} \quad (36)$$

where  $Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)$  is a normalizing constant depending on the hyperparameters, explicitly

$$Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2) = \sqrt{\frac{2\pi}{\kappa_0}} \Gamma\left(\frac{\nu_0}{2}\right) \left(\frac{2}{\nu_0 \sigma_0^2}\right)^{\nu_0/2}. \quad (37)$$

Inserting (36) and (16) into (35), we need to compute for each  $i$  as given in (38) where  $\kappa_{N_i,i}$ ,  $\mu_{N_i,i}$ ,  $\nu_{N_i,i}$  and  $\sigma_{N_i,i}^2$  are defined as

$$\begin{aligned} \kappa_{N_i,i} &= \kappa_0 + N_i, \\ \mu_{N_i,i} &= \frac{\kappa_0 \mu_0 + N_i \bar{x}_i}{\kappa_{N_i,i}}, \\ \nu_{N_i,i} &= \nu_0 + N_i, \\ \sigma_{N_i,i}^2 &= \frac{\nu_0}{\nu_{N_i,i}} \sigma_0^2 + \frac{N_i - 1}{\nu_{N_i,i}} S_i + \frac{\kappa_0 N_i (\mu_0 - \bar{x}_i)^2}{\nu_{N_i,i} \kappa_{N_i,i}}. \end{aligned} \quad (39)$$

Substituting (37) into (38) and then back into (35), we obtain the likelihood in closed form as

$$\begin{aligned} &p(\mathcal{X}_1, \dots, \mathcal{X}_P | \mu_0, \kappa_0, \nu_0, \sigma_0^2) \\ &= \prod_{i=1}^P \frac{\Gamma(\nu_{N_i,i}/2)}{\Gamma(\nu_0/2)} \sqrt{\frac{\kappa_0}{\kappa_{N_i,i}}} \frac{(\nu_0 \sigma_0^2)^{\nu_0/2}}{(\nu_{N_i,i} \sigma_{N_i,i}^2)^{\nu_{N_i,i}/2}} \frac{1}{\pi^{N_i/2}}. \end{aligned} \quad (40)$$

<sup>5</sup>  $\sigma_{i,MLE}^2$  is a biased estimator. To eliminate the bias, we may replace  $\sigma_{i,MLE}^2$  in (33) by its unbiased estimator.

$$\begin{aligned}
& \int_0^\infty \int_{-\infty}^\infty p(\mathcal{X}_i | \mu_i, \sigma_i^2) p(\mu_i, \sigma_i | \kappa_0, \mu_0, \nu_0, \sigma_0^2) d\mu_i d\sigma_i^2 \\
&= \int_0^\infty \int_{-\infty}^\infty \frac{\sigma_i^{-N_i}}{(2\pi)^{N_i/2}} \exp\left\{-\frac{N_i(\bar{x}_i - \mu_i)^2 + (N_i - 1)S_i}{2\sigma_i^2}\right\} \frac{\sigma_i^{-\nu_0 - 3}}{Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)} \exp\left\{-\frac{\nu_0 \sigma_0^2 + \kappa_0(\mu - \mu_0)^2}{2\sigma_i^2}\right\} d\mu_i d\sigma_i^2 \\
&= \frac{1}{(2\pi)^{N_i/2}} \frac{1}{Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)} \int_0^\infty \int_{-\infty}^\infty \sigma_i^{-\nu_{N_i, i} - 3} \exp\left\{-\frac{\nu_{N_i, i} \sigma_{N_i, i}^2 + \kappa_{N_i, i}(\mu - \mu_{N_i, i})^2}{2\sigma_i^2}\right\} d\mu_i d\sigma_i^2 \\
&= \frac{1}{(2\pi)^{N_i/2}} \frac{Z(\kappa_{N_i, i}, \mu_{N_i, i}, \nu_{N_i, i}, \sigma_{N_i, i}^2)}{Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)} \int_0^\infty \int_{-\infty}^\infty \mathcal{N}(\mu_i | \mu_{N_i, i}, \sigma_i^2 / \kappa_{N_i, i}) \chi^{-2}(\sigma_i^2 | \nu_{N_i, i}, \sigma_{N_i, i}^2) d\mu_i d\sigma_i^2 \\
&= \frac{1}{(2\pi)^{N_i/2}} \frac{Z(\kappa_{N_i, i}, \mu_{N_i, i}, \nu_{N_i, i}, \sigma_{N_i, i}^2)}{Z(\kappa_0, \mu_0, \nu_0, \sigma_0^2)}. \tag{38}
\end{aligned}$$

## B. MAP Estimation

For NIX priors of  $\mu_i$  and  $\sigma_i$ , the posterior in (14) is

$$\mathcal{N}(\mu_i | \mu_{N_i, i}, \sigma_i^2 / \kappa_{N_i, i}) \chi^{-2}(\sigma_i^2 | \nu_{N_i, i}, \sigma_{N_i, i}^2). \tag{41}$$

Therefore, MAP estimates of  $\mu_i$  and  $\sigma_i$  are the modes of the posterior, which can be seen to be simply [21]

$$\begin{aligned}
\mu_{i, MAP} &= \mu_{N_i, i} = \frac{\kappa_0 \mu_0 + \sum_{j=1}^{N_i} x_{ij}}{\kappa_0 + N_i} \\
\sigma_{i, MAP}^2 &= \frac{\nu_{N_i, i} \sigma_{N_i, i}^2}{\nu_{N_i, i} + 3}. \tag{42}
\end{aligned}$$

The simplified expression for  $\mu_{i, MAP}$  can be interpreted as adding  $\kappa_0$  number of fake data samples with value (and so mean also)  $\mu_0$  to the measured data  $\mathcal{X}_i$  for population  $i$ . Similarly by expanding expression for  $\sigma_{i, MAP}^2$  one can obtain an analogous elucidation for the variance update. Basically this additional fake data samples incorporates the information present in the prior learnt using data from all the populations.

Furthermore, note that  $\sigma_{i, MAP}^2$  in (42) is a biased estimator. An unbiased estimation can be obtained by

$$\sigma_{i, MAP-UB}^2 = \frac{\nu_{N_i, i} \sigma_{N_i, i}^2}{\nu_{N_i, i} - 1}. \tag{43}$$


## ACKNOWLEDGMENT

The authors would like to thank E. Chiprout for various discussions and continuous support.

## REFERENCES

- [1] C. Bishop, *Pattern Recognition and Machine Learning*, vol. 4. New York, NY, USA: Springer, 2006.
- [2] J. Friedman, T. Hastie, and R. Tibshirani, *The Elements of Statistical Learning (Springer Series in Statistics)*, vol. 1. Berlin, Germany: Springer, 2001.
- [3] G. Balamurugan, B. Casper, J. Jaussi, M. Mansuri, F. O'Mahony, and J. Kennedy, "Modeling and analysis of high-speed I/O links," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 237–247, May 2009.
- [4] J. Keshava, N. Hakim, and C. Prudvi, "Post-silicon validation challenges: How EDA and academia can help," in *Proc. DAC*, 2010, pp. 3–7.
- [5] C. Gu, "Challenges in post-silicon validation of high-speed I/O links," in *Proc. ICCAD*, 2012, pp. 547–550.
- [6] Intel Corp. (2003). *Intel Platform and Component Validation* [Online]. Available: [http://download.intel.com/design/chipsets/labtour/PVPT\\_WhitePaper.pdf](http://download.intel.com/design/chipsets/labtour/PVPT_WhitePaper.pdf)
- [7] E. E. Lior Shkolnitsky. (2010, Jun.) *Electrical System-Validation Methodology for Embedded DisplayPort* [Online]. Available: <http://download.intel.com/design/intarch/PAPERS/323931.pdf>
- [8] K. Gambill. (2008, Dec.) *System Margin Validation* [Online]. Available: <http://download.intel.com/design/intarch/papers/321078.pdf>
- [9] W. Zhang, T. Chen, M. Ting, and X. Li, "Toward efficient large-scale performance modeling of integrated circuits via multi-mode/multi-corner sparse regression," in *Proc. DAC*, 2010, pp. 897–902.
- [10] X. Li, W. Zhang, F. Wang, S. Sun, and C. Gu, "efficient parametric yield estimation of analog/mixed-signal circuits via Bayesian model fusion," in *Proc. ICCAD*, 2012, pp. 627–634.

- [11] (2014, Jan.) [Online]. Available: <http://www.pcisig.com>
- [12] (2014, Jan.) [Online]. Available: <http://www.jedec.org>
- [13] G. E. Box and N. R. Draper, *Empirical Model Building and Response Surfaces*. New York, NY, USA: Wiley, 1987.
- [14] S. Arlot and A. Celisse, "A survey of cross-validation procedures for model selection," *Statist. Surveys*, vol. 4, pp. 40–79, 2010.
- [15] G. Casella, "An introduction to empirical Bayes data analysis," *Amer. Statist.*, vol. 39, no. 2, pp. 83–87, 1985.
- [16] B. Efron and C. Morris, "Data analysis using Stein's estimator and its generalizations," *J. Amer. Statist. Assoc.*, vol. 70, no. 350, pp. 311–319, 1975.
- [17] X. Li, J. Le, P. Gopalakrishnan, and L. T. Pileggi, "Asymptotic probability extraction for nonnormal performance distributions," *IEEE Trans. Comput. Aided Design Integr. Circuits Syst.*, vol. 26, no. 1, pp. 16–37, Jan. 2007.
- [18] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA, USA: MIT Press, 2009.
- [19] D. B. Owen, "A table of normal integrals: A table," *Commun. Statist. Simulation Comput.*, vol. 9, no. 4, pp. 389–419, 1980.
- [20] A. Papoulis and S. Pillai, *Probability, Random Variables and Stochastic Processes*. New York, NY, USA: McGraw-Hill, 2001.
- [21] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.



**Chenjie Gu** (M'11) received the Ph.D. degree in electrical engineering and computer science from University of California, Berkeley, CA, USA, in 2011.

He is currently a Research Scientist with Intel Strategic Computer-Aided Design Laboratories, Hillsboro, OR, USA. His current research interests include design, modeling, and validation of nano-scale systems, with a special interest in the development and application of computational tools and algorithmic techniques.

Dr. Gu received an ASP-DAC Best Paper Award in 2008 and an ICCAD Best Paper Nomination in 2010.



**Manzil Zaheer** (S'10) received the B.Tech degree in electronics and communication engineering from the Indian Institute of Technology Guwahati, Guwahati, India, in 2012, and is currently pursuing the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, USA.

His current research interests include machine learning with specific interests in nonparametric Bayesian methods, random graph generative methods, and their application to computer-aided design.



**Xin Li** (S'01–M'06–SM'10) received the Ph.D. degree in electrical and computer engineering (ECE) from Carnegie Mellon University (CMU), Pittsburgh, PA, USA, in 2005.

He is currently an Associate Professor with the ECE Department, CMU. His current research interests include integrated circuit and signal processing.

Dr. Li was a recipient of the NSF Faculty Early Career Development Award in 2012, the IEEE Donald O. Pederson Best Paper Award in 2013, a Best Paper Award from Design Automation Conference in 2010, and two IEEE/ACM William J. McCalla ICCAD Best Paper Awards in 2004 and 2011.