

Yield Prognosis for Fab-to-Fab Product Migration

Ali Ahmadi*, Ke Huang[†], Amit Nahar[‡], Bob Orr[‡], Michael Pas[‡], John M. Carulli Jr.[§] and Yiorgos Makris*

*Department of Electrical Engineering, The University of Texas at Dallas, Richardson, TX 75080

[†]Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA 92115

[‡]Texas Instruments Inc., 12500 TI Boulevard, MS 8741, Dallas, TX 75243

[§]GlobalFoundries, 400 Stone Break Road Extension, Malta, NY 12020

Abstract—We investigate the utility of correlations between e-test and probe test measurements in predicting yield. Specifically, we first examine whether statistical methods can accurately predict parametric probe test yield as a function of e-test measurements within the same fab. Then, we investigate whether the e-test profile of a destination fab, in conjunction with the e-test and probe test profiles of a source fab, suffice for accurate yield prognosis during fab-to-fab product migration. Results using an industrial dataset of ~3.5M devices from a 65nm Texas Instruments RF transceiver design fabricated in two different fabs reveal that (i) within-fab yield prediction error is in the range of a few tenths of a percentile point, and (ii) fab-to-fab yield prediction error is in the range of half a percentile point.

I. INTRODUCTION

The rapidly growing and dynamically changing consumer electronics market introduces interesting challenges to production planning of semiconductor manufacturing companies, calling for agility and flexibility in order to efficiently respond to fluctuating demand. Contingency plans for dealing with catastrophic events, such as earthquakes, floods, and hurricanes, which have severely hampered the market in the past, as well as political or sheer financial reasons, often place similar constraints in production planning as well. Migrating a product from one fab to another, however, is not a trivial endeavor. Different fabs, implementing the same technology node and even employing identical equipment and software suites, are bound to exhibit variations in the parametric profile of the silicon they produce, and by extension, the yield of a device fabricated therein. Accurate yield prognosis, however, is an indispensable piece of information during device migration and production planning.

Predicting parametric yield for a design produced by a specific fab is not a new problem, with solutions varying from pure simulation-based to silicon measurement-driven. Traditionally, Monte Carlo based approaches have been used to generate a large number of random samples based on expected process variations, in order to estimate the distribution of each performance of interest [1]. Alternatively, modeling techniques, which approximate a performance of interest as a linear or non-linear function of device-level parameters have also been employed. Subsequently, the distribution of a desired performance is estimated by numerical methods [2], [3]. Such simulation-based methods, however, are of limited accuracy. Along another direction, the authors of [4] introduced the use of Bayesian model fusion for yield estimation, wherein pre-silicon simulations are reinforced with a small set of post-silicon measurements to enhance model accuracy. Similarly, in [5], high volume manufacturing yield of a product is estimated

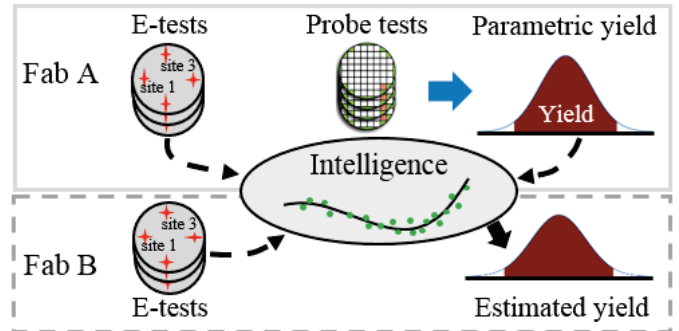


Fig. 1: Yield Prediction during fab-to-fab migration.

through spatio-temporal wafer correlation models learned from early silicon wafers. Such silicon measurement-based methods, however, assume access to probe measurements from a tangible number of wafers of the device, produced in the fab of interest, in order to estimate parametric production yield. Yet in the context of fab-to-fab product migration, this information is not available since the device has yet to be produced in the target fab.

To address this problem, in this work we develop a yield prognosis method which does not require target fab measurements from the device to be migrated. Instead, as shown in Figure 1, it relies on e-test and probe test measurements from the source fab, where the device is currently produced, as well as the e-test profile of the target fab, which can be obtained from other devices produced therein, as e-tests are typically common across devices in the same technology node. Toward this end, in Section II we first discuss a regression-based solution to the problem of correlating e-test measurements to parametric probe test yield within a single fab. Then, in Section III, we introduce three methods for extending this capability to the fab-to-fab product migration scenario, namely *model migration*, *importance sampling* and *predictor calibration*. Experimental results demonstrating the effectiveness of parametric yield prediction based on e-test measurements on actual production data for both the within-fab and the fab-to-fab migration scenarios are presented in Section IV and conclusions are drawn in Section V.

II. WITHIN-FAB CORRELATION

Before we address the problem of predicting parametric probe test yield from e-tests during fab-to-fab product migration, we discuss the simpler version of doing so within the same fab. Meaningful correlations among measurements from various stages of semiconductor manufacturing are known to




Fig. 2: E-test/yield correlation model.

exist and have been utilized in several tasks, such as test cost/time reduction and yield improvement in the past. In one such approach, described in [6], the authors proposed a statistical method for predicting probe test outcomes from e-test data of a wafer. This method employed a genetic algorithm-based approach to select a key subset of all e-test parameters and, then, build a multivariate nonlinear correlation model between selected e-tests and probe test outcomes. The identified correlations were iteratively refined through designer feedback, with the main objective of providing useful information for process health monitoring, rather than reducing test cost/time. Indeed, while the accuracy of these e-test to probe test prediction models is high, it would not suffice for omitting probe test of *individual dies* with acceptable test error. Nevertheless, if one is interested in predicting parametric probe test yield across an *entire wafer* from its e-tests, as is the case in our problem, these models are very accurate.

Figure 2 depicts how parametric probe test yield can be predicted through e-test measurements of wafers in a fab. First, for a training set of wafers, both e-tests and parametric probe test measurements are obtained. The device specifications are then used to compute parametric yield for each probe test across the wafer. Using the e-tests and the parametric probe test yield, a correlation model is trained. Finally, for a new wafer produced by this fab, its e-test measurements can be provided to the trained correlation model in order to predict parametric yield. The key component in this scheme is the construction of the correlation model, whereby the dependent variables (parametric probe test yield) are expressed as functions of predictors (e-test measurements).

Several methods exist in the literature for multivariate regression such as *Multivariate Adaptive Regression Splines (MARS)*, *Least-Angle Regression Splines (LARS)*, *Projection Pursuit Regression*, *Multi-Layer Perceptrons*, and *Radial Basis Function Networks* [7]. Among them, in this work, we use MARS [8], which was also used in [6] and several other test cost reduction methods in the past [9].

III. YIELD PROGNOSIS IN FAB-TO-FAB MIGRATION

We now turn our attention back to the fab-to-fab migration problem, wherein we seek a prognosis of the parametric yield of a product migrating from a source fab to a target fab. This prognosis may be based on the e-test and the parametric probe test of the source fab, where the device is currently produced so ample data is available. In addition, it may also be based on the e-test profile of the target fab, which can be obtained from other devices produced therein, as most e-tests

are typically shared across designs on the same technology node. The probe test profile for the target fab, however, is not available, since the device has yet to be produced therein. Our objective is to statistically predict parametric probe test yield in the target fab based on the above data. After introducing notation and formulating the problem, we describe three such methods, namely *model migration*, *importance sampling*, and *predictor calibration*.

A. Notation and Problem Formulation

Given a set of e-test measurements, \mathbf{eT}_S , and probe test measurements, \mathbf{PT}_S , from the source fab, we can use the device specification limits to compute the parametric probe test yield vector for every wafer in our dataset. Each wafer can, then, be represented by:

$$\text{wafer}_S^i = (\mathbf{eT}_S^i, \mathbf{y}_S^i) \quad (1)$$

where $\mathbf{eT}_S^i = [et_S^{i1}, et_S^{i2}, \dots, et_S^{im}]$ is the m -dimensional vector of e-test measurements for wafer i and $\mathbf{y}_S^i = [y_S^{i1}, y_S^{i2}, \dots, y_S^{ik}]$ is the k -dimensional parametric yield vector for the probe test measurements of wafer i . Let us also denote by $\mathbf{p}_S(\mathbf{eT}_S)$ the density function of e-tests over n_S wafers of the source fab. Similarly, given a set of e-test measurements, \mathbf{eT}_T , from the target fab, with their density function over n_T wafers denoted by $\mathbf{p}_T(\mathbf{eT}_T)$, a wafer can be represented by:

$$\text{wafer}_T^j = [\mathbf{eT}_T^j] = [et_T^{j1}, et_T^{j2}, \dots, et_T^{jm}] \quad (2)$$

Our objective is to predict the k -dimensional parametric yield vector for the probe tests of wafer j , $\hat{\mathbf{y}}_T^j = [\hat{y}_T^{j1}, \hat{y}_T^{j2}, \dots, \hat{y}_T^{jk}]$, for each of the n_T wafers of the target fab.

B. Model Migration

A straightforward approach is to use the method discussed in Section II to express parametric yield in the source fab as a function of its e-tests, $\mathbf{Y}_S = \mathbf{f}_S(\mathbf{eT}_S)$. Then, the trained regression function can be applied directly to the e-tests of the target fab, in order to predict its parametric yield, $\mathbf{Y}_T = \mathbf{f}_S(\mathbf{eT}_T)$. Model migration success relies on two assumptions:

1. Homogeneous distribution between training and testing data sets, i.e. e-tests in the source and target fabs must come from the same distribution, $\mathbf{p}_S(\mathbf{eT}_S) = \mathbf{p}_T(\mathbf{eT}_T)$.
2. Identical conditional distribution of yield values for training and testing data sets, $\mathbf{p}_S(\mathbf{Y}_S | \mathbf{eT}_S^i) = \mathbf{p}_T(\mathbf{Y}_T | \mathbf{eT}_T^j) \Rightarrow \mathbf{eT}_S^i \approx \mathbf{eT}_T^j$. In other words, if a wafer from the source fab and a wafer from the target fab have the same yield, they must also have similar e-test vectors.

As these assumptions do not necessarily hold true in a semiconductor manufacturing context, the accuracy of model migration is expected to be limited.

C. Importance Sampling

Another approach, which revokes the homogeneity assumption but retains the identical conditional distribution assumption discussed above, is importance sampling. In order to build a model using the training data, which will retain its accuracy




Fig. 3: Instance weighting example.

when used on the testing data, importance sampling biases the selection of the instances from which the model is built. Specifically, preference is given to the most relevant region of the training distribution (source fab e-tests), i.e. the region which overlaps the most with the testing distribution (target fab e-tests). To achieve this, higher weights are assigned to samples from the relevant region, hence the name importance sampling [10]. This method has been successful in various real-world applications [11]–[14].

In our context, we seek to assign weights (importances) to the instances of the source fab e-test distribution, favoring those which are frequently encountered in the target fab e-test distribution with higher weights and penalizing those that are rarely encountered with lower weights. To achieve this, we compute the weight of each instance in the source distribution as the ratio of the target density over the source density. Therefore, we first estimate the density of e-tests for the source and target fabs separately, and then we compute the weight vector through the following equation:

$$\mathbf{W} = \frac{\mathbf{p}_T(\mathbf{eT}_S)}{\mathbf{p}_S(\mathbf{eT}_S)} \quad (3)$$

Figure 3 shows an example of source and target fab densities for one e-test, along with the instance weights calculated using Equation 3. The dashed and solid lines represent the density of the training and testing data, while the dotted graph is the weight corresponding to each training instance.

To apply importance sampling to the fab-to-fab migration problem, we first allocate the instance weights for the source fab. Next, we train a regression model to express parametric yield in the source fab as a function of its e-tests, $\mathbf{Y}_S = \mathbf{f}_{wS}(\mathbf{eT}_S, \mathbf{W})$, but elements from the training set are selected with probability commensurate with their assigned weights, instead of equal probability, as in model migration. The trained regression function is, then, applied directly to the e-tests of the target fab to predict its parametric yield, $\mathbf{Y}_T = \mathbf{f}_{wS}(\mathbf{eT}_T)$.

Importance sampling is expected to perform better than model migration, as it does not rely on the homogeneity assumption. Nevertheless, it still assumes identical conditional distributions, hence there remains room for improvement.

D. Predictor Calibration

A third approach, which does not rely on any of these two assumptions, is predictor calibration. In this method, the distribution of e-tests in the target fab (i.e. predictors) is calibrated based on the distribution of e-tests in the source

fab, $\widehat{\mathbf{eT}}_T = \mathbf{h}(\mathbf{eT}_T, \mathbf{eT}_S)$, prior to being used for predicting parametric probe test yield in the target fab. A simple way of achieving this would be mean calibration, which removes the mean shift, $\Delta(\mu)$, from each instance of target distribution:

$$\widehat{\mathbf{eT}}_T^i = \mathbf{eT}_T^i - \Delta(\mu), \quad \Delta(\mu) = \mu(\mathbf{eT}_T) - \mu(\mathbf{eT}_S) \quad (4)$$

However, in order to achieve better precision, other parameters of the distribution, such as variance, skewness and kurtosis, also need to be calibrated. To accomplish this, we employ a two-step procedure. First, in the cumulative distribution function (CDF) of the target fab, we find the cumulative probability associated with each sample in the target distribution, $\mathbf{x}_i = \mathbf{F}_T(\mathbf{eT}_T^i)$. Then, using the inverse CDF of the source fab, we determine the e-test value associated with cumulative probability \mathbf{x}_i , $\widehat{\mathbf{eT}}_T^i = \mathbf{F}_S^{-1}(\mathbf{x}_i)$, where \mathbf{F}_S^{-1} is the inverse CDF of the source fab distribution. This procedure is applied to all n_T instances (i.e. wafers) of the target fab distribution. The predictor calibration algorithm is summarized below:

```

 $\widehat{\mathbf{eT}}_T = \emptyset$ 
for  $\mathbf{eT}_T^i \in \mathbf{eT}_T$  do
   $\mathbf{x}_i \leftarrow \mathbf{F}_T(\mathbf{eT}_T^i)$ 
   $\widehat{\mathbf{eT}}_T^i \leftarrow \mathbf{F}_S^{-1}(\mathbf{x}_i)$ 
 $\widehat{\mathbf{eT}}_T \leftarrow \widehat{\mathbf{eT}}_T \cup \widehat{\mathbf{eT}}_T^i$ 
end for

```

Using this method, the mapping function is defined as:

$$\widehat{\mathbf{eT}}_T = \mathbf{h}(\mathbf{eT}_T) = \mathbf{F}_S^{-1}(\mathbf{F}_T(\mathbf{eT}_T)) \quad (5)$$

In order to utilize predictor calibration in fab-to-fab migration, a regression function is first trained to express parametric yield in the source fab as a function of its e-tests, $\mathbf{Y}_S = \mathbf{f}_S(\mathbf{eT}_S)$. Then, the prediction calibration algorithm maps the distribution of e-tests in the target fab into the distribution of e-tests in the source fab, $\widehat{\mathbf{eT}}_T = \mathbf{h}(\mathbf{eT}_T)$. Eventually, the trained regression model is applied to the calibrated e-tests of the target fab, in order to predict parametric yield, $\mathbf{Y}_T = \mathbf{f}_S(\widehat{\mathbf{eT}}_T)$.

Since predictor calibration does not make any of the two assumptions stated earlier, it is expected to outperform both model migration and importance sampling.

IV. EXPERIMENTAL RESULTS

In order to experimentally evaluate the effectiveness of the proposed yield prognosis methods, we use actual production data from a 65nm analog/RF device currently in high volume manufacturing (HVM) production by Texas Instruments¹. This data, which is depicted in Figure 4, comprises devices from two geographically dispersed fabs wherein this device is fabricated, which we will refer to as fab A and fab B. The dataset for fab A includes 54 e-test and 168 parametric probe test measurements from a total of 1800 wafers, each of which has 9 e-test measurement sites and approximately 1500 die per wafer. The dataset for fab B includes the same e-test and parametric probe test measurements from a total of 500 wafers,

¹Details regarding the device cannot be released due to an NDA under which this data has been provided to us.

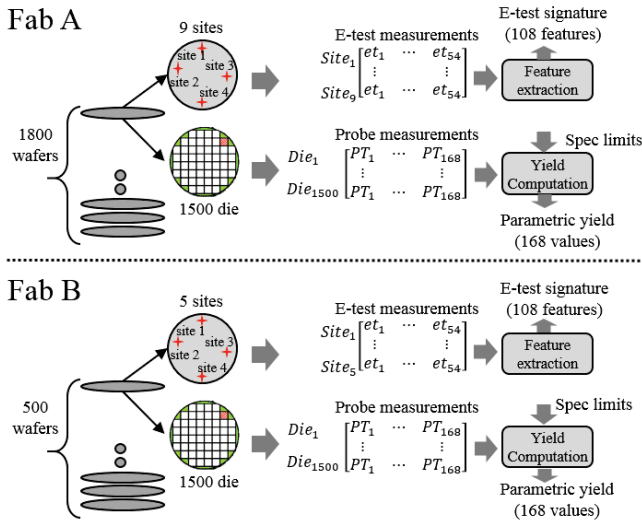


Fig. 4: Experimental dataset.

with the only difference being that e-tests are obtained on only 5 instead of 9 sites. These two datasets were obtained from the two fabs at approximately the same time period. Along with the data, we are also provided with the specification limits for each of the 168 parametric probe tests, hence we can compute the yield of each performance on every wafer for each of the two fabs. Additionally, for each of the 54 e-test measurements, we compute the mean and the standard deviation across the 9 sites on wafers produced in fab A (5 sites on wafers produced in fab B), hence the e-test signature of each wafer consists of 108 parameters. Using this dataset, we seek to:

- Quantify the accuracy of statistically predicting parametric yield from e-test measurements within a single fab.
- Quantify the accuracy of the described prognosis methods in statistically predicting yield during fab-to-fab product migration based on e-test and probe test profiles of the source fab and only an e-test profile of the target fab².

In both cases, we use two metrics to quantify prediction accuracy. The first metric is the average absolute difference, δ_i , between predicted and actual yield for the i -th probe test:

$$\delta_i = \frac{1}{n_T} \sum_{j=1}^{n_T} |\hat{y}_{ij} - y_{ij}| \quad (6)$$

where n_T is the number of wafers for which the prediction is applied, while \hat{y}_{ij} and y_{ij} are the predicted and the actual yield of the i -th probe test on the j -th wafer, respectively.

The second metric, ϵ_i , normalizes the average absolute difference to the yield range:

$$\epsilon_i = \frac{1}{n_T} \sum_{j=1}^{n_T} \frac{|\hat{y}_{ij} - y_{ij}|}{\max(\mathbf{y}_i) - \min(\mathbf{y}_i)} \quad (7)$$

where $\max(\mathbf{y}_i)$ and $\min(\mathbf{y}_i)$ are the highest and lowest yield values, respectively, of the i -th probe test across all wafers. Expressing prediction error as a percentage of this range is important towards gaging its significance.

²We note that the e-test profile of the target fab should be obtained from a different product fabricated therein. Since we only have data from one device, however, we use its e-test profile as a proxy.




Fig. 5: Within-fab yield prediction error.

A. Within-Fab Yield Prediction

In order to quantify the accuracy of statistically predicting parametric yield from e-test measurements within a single fab, we use 90% of the available wafers from a fab as the training set and the remaining 10% as the test set. Using the 108 e-test features and the 168 yield values reflecting each wafer in our training set, we train a separate regression model (i.e. MARS) for each of the 168 probe tests. The trained regression models are then applied to the 108 e-test features of each wafer in the test set, in order to predict the yield of each of the 168 probe tests across this wafer. The predicted results are, then, compared to the actual values, which are available in the dataset, in order to estimate prediction accuracy. To establish statistical significance, we apply a 10-fold cross validation approach where results are averaged over 10 repetitions, each time randomly splitting the dataset into training and test sets.

Figures 5(a) and 5(b) present the results for the datasets of fab A and fab B, respectively, using the first metric, δ_i , defined in Equation 6. The horizontal axis shows the 168 probe tests, sorted in increasing prediction error, while the vertical axis shows the corresponding average absolute difference between the predicted and actual yield³. As may be observed, this difference is in the order of a few tenths of a percentage point, corroborating that parametric probe test yield can be predicted very accurately from the e-test measurements of a wafer.

Figures 6(a) and 6(b) demonstrate the same results, this time using the second metric, ϵ_i , defined in Equation 7. In each histogram, the horizontal axis is the prediction error, while the vertical axis shows the percentage of probe tests that are predicted within a given error range. For example, the first bar shows the percentage of probe test measurements whose normalized average prediction error is below 2%, with the corresponding value being 60% and 21% for fab A and fab B, respectively. As may be observed, the yield of the vast majority of probe tests can be predicted using e-test measurements with an error which is well below 10% of their yield range.

³We note that since our test data is Continue on Fail (COF) and a device might fail multiple probe tests, the sum of the yield prediction errors over the 168 probe tests does not reflect the overall yield prediction error.




Fig. 6: Normalized within-fab yield prediction error.

TABLE I: Average within-fab yield prediction error.

Metric	Fab A	Fab B
δ_i	0.23%	0.16%
ϵ_i	3.2%	5.6%

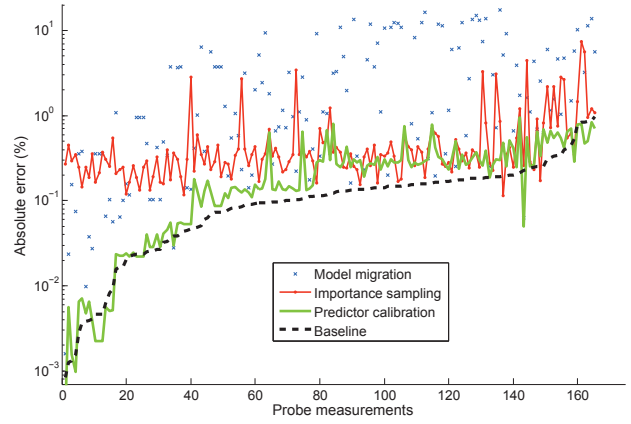
These results are summarized in Table I, where the average absolute prediction error is calculated at 0.23% for fab A and at 0.16% for fab B. Normalized to the yield range of each probe test, the results are 3.2% for fab A and 5.6% for fab B.

B. Yield Prognosis for Fab-to-Fab Migration


In order to quantify the accuracy of the described prognosis methods in predicting yield during fab-to-fab product migration, we performed the following experiment, first using fab A as the source and fab B as the target, and then reversing the roles: using the 108 e-test features and the 168 yield values from every wafer in the dataset of the source fab, as well as the 108 e-test features from every wafer in the dataset of the target fab, we apply the three methods described in Section III to predict the yield for the 168 probe tests in the wafers of the target fab. The predicted values are, then, compared to the actual yield values, which are available in our dataset, in order to estimate prediction accuracy. As a baseline for prediction accuracy, we use the within-fab yield prediction results for the target fab, which were presented in the previous subsection.

Figures 7(a) and 7(b) present the results for product migration from fab A to fab B and vice-versa, respectively, for each of the three methods of Section III, using the first metric, δ_i , defined in Equation 6. The within-fab baseline results are also shown as a point of reference. The horizontal axis shows the 168 probe tests, sorted in increasing prediction error, while the vertical axis shows the corresponding average absolute difference between the predicted and actual yield for each method. We note that, in this case, the vertical axis is in logarithmic scale and in the *model migration* plot connecting lines are omitted in order to enhance figure readability.

As may be observed, *model migration*, wherein the correlation models learned on the source fab are directly applied to the e-tests of the target fab, results in prediction error in the range of 10% and 5% for the two experiments, respectively. This is expected, since this approach assumes homogeneous e-test distributions and identical conditional yield distributions in the two fabs, something that is typically not the case. *Importance sampling*, on the other hand, reduces the yield prediction error to within a couple of percentage points. Evidently, the weighting policy used therein is effective in modeling the



(a) Migrating from fab A to fab B



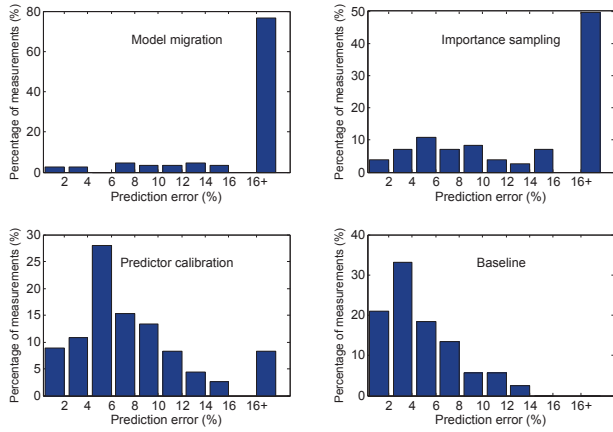
(b) Migrating from fab B to fab A

Fig. 7: Absolute fab-to-fab yield prediction error.

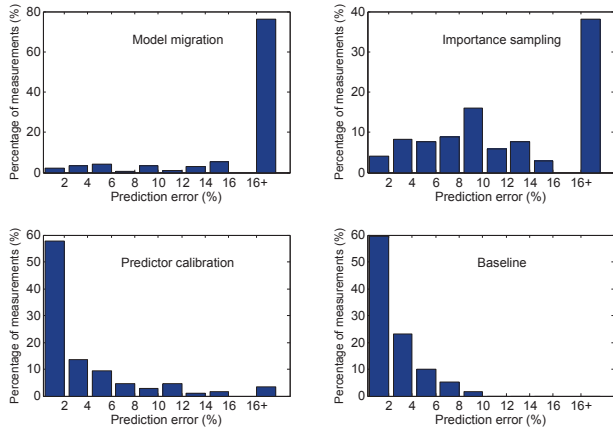
difference in the marginal distribution of e-tests in the source and target fabs. Finally, *predictor calibration* outperforms the other two methods due to the accurate mapping of the target into the source distribution. Indeed, the yield prediction error drops to below half of a percentage point and is very close to the baseline achieved by within-fab correlation models.

Figures 8(a) and 8(b) demonstrate the results of the same experiments, this time using the second metric, ϵ_i , defined in Equation 7. In each histogram, the horizontal axis is the prediction error, while the vertical axis shows the percentage of probe tests that are predicted within a given error range. Separate histograms are shown for each of the three fab-to-fab yield prognosis methods, as well as the within-fab baseline method. As may be observed, the results corroborate our previous observation that the *predictor calibration* method is almost as efficient as the baseline within-fab yield prediction method, achieving accuracy which is within a single-digit percentage of the yield range. In contrast, *model migration* and *importance sampling* are far less accurate, resulting in a normalized prediction error of more than 16% of the yield range for the majority of the probe test measurements.

These results are summarized in Tables II(a) and II(b) for the two experiments, respectively. When migrating from fab A to fab B, the average absolute prediction error over all probe tests and wafers is calculated at 5.52%, 0.98%, and 0.54% for



(a) Migrating from fab A to fab B



(b) Migrating from fab B to fab A

Fig. 8: Normalized fab-to-fab yield prediction error.

the three fab-to-fab migration methods, respectively, compared to 0.16% for the within-fab baseline yield prediction. Normalized to the yield range of each probe test, the results are 87.1%, 31.4%, 10.3%, and 5.6%, respectively. When migrating from fab B to fab A, the average absolute prediction error over all probe tests and wafers is 9.84%, 1.37%, and 0.35% for the three fab-to-fab migration methods, respectively, compared to 0.23% for the within-fab baseline yield prediction. Normalized to the yield range of each probe test, the results are 82.6%, 28.2%, 4.9%, and 3.2%, respectively.

V. CONCLUSION

E-test and probe test measurements exhibit strong correlation which can be statistically harnessed for yield learning purposes. As we demonstrated using a large dataset from a 65nm Texas Instruments RF transceiver produced in two different fabs, these correlations enable very accurate prediction of parametric probe test yield from e-test measurements within the same fab, with error ranging in the order of a few tenths of a percentage. Moreover, using the e-test and probe test profiles of a source fab and only the e-test profile of a target fab, which can be obtained from prior devices fabricated therein, these correlations facilitate highly accurate yield prognosis when migrating a product across these fabs, with error ranging in the order of half of a percentage.

TABLE II: Average error for fab-to-fab migration.

(a) Migrating from fab A to fab B				
Metric	Model migration	Importance sampling	Predictor calibration	Baseline
δ_i	5.52%	0.98%	0.54%	0.16%
ϵ_i	87.1%	31.47%	10.29%	5.6%
(b) Migrating from fab B to fab A				
Metric	Model migration	Importance sampling	Predictor calibration	Baseline
δ_i	9.84%	1.37%	0.35%	0.23%
ϵ_i	82.6%	28.2%	4.9%	3.2%

VI. ACKNOWLEDGEMENT

This research has been partially supported by the Semiconductor Research Corporation (SRC) Task 1836.131.

REFERENCES

- [1] F. Gong, H. Yu, and L. He, "Stochastic analog circuit behavior modeling by point estimation method," in *ACM International Symposium on Physical Design*, 2011, pp. 175–182.
- [2] X. Li, J. Le, P. Gopalakrishnan, and L. Pileggi, "Asymptotic probability extraction for non-normal distributions of circuit performance," in *IEEE/ACM International Conference on Computer-Aided Design*, 2004, pp. 2–9.
- [3] R. H. Myers, D. Montgomery, and C. Anderson-Cook, "Response surface methodology: process and product optimization using designed experiment," *John Wiley and Sons, New York*, pp. 343–350, 2002.
- [4] C. Fang, F. Yang, X. Zeng, and X. Li, "BMF-BD: Bayesian model fusion on Bernoulli distribution for efficient yield estimation of integrated circuits," in *ACM Design Automation Conference*, 2014, pp. 1–6.
- [5] A. Ahmadi, K. Huang, S. Natarajan, J. Carulli, and Y. Makris, "Spatio-temporal wafer-level correlation modeling with progressive sampling: A pathway to HVM yield estimation," in *IEEE International Test Conference*, 2014, pp. 1–10.
- [6] N. Kupp, M. Slamani, and Y. Makris, "Correlating inline data with final test outcomes in analog/RF devices," in *IEEE Design, Automation & Test in Europe Conference & Exhibition*, 2011, pp. 1–6.
- [7] V. Cherkassky and F. Mulier, *Learning from data: concepts, theory, and methods*, John Wiley & Sons, 2007.
- [8] J. H. Friedman, "Multivariate adaptive regression splines," *The annals of statistics*, pp. 1–67, 1991.
- [9] P. Variyam, S. Cherubal, and A. Chatterjee, "Prediction of analog performance parameters using fast transient testing," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 21, no. 3, pp. 349–361, 2002.
- [10] T. C. Hesterberg, *Advances in importance sampling*, Ph.D. thesis, Stanford University, 2003.
- [11] P. Baldi and S. Brunak, *Bioinformatics: the machine learning approach*, MIT press, 2001.
- [12] J. J. Heckman, "Sample selection bias as a specification error," *Journal of the Econometric Society*, pp. 153–161, 1979.
- [13] M. Sugiyama, M. Krauledat, and K. R. Müller, "Covariate shift adaptation by importance weighted cross validation," *Journal of Machine Learning Research*, vol. 8, pp. 985–1005, 2007.
- [14] H. Shimodaira, "Improving predictive inference under covariate shift by weighting the log-likelihood function," *Journal of Statistical Planning and Inference*, vol. 90, no. 2 pp. 227–244, 2000.