# A Comparative Study of One-Shot Statistical Calibration Methods for Analog / RF ICs

Yichuan Lu*, Kiruba S. Subramani*, He Huang†, Nathan Kupp†, Ke Huang‡, and Yiorgos Makris*

*Department of Electrical Engineering, The University of Texas at Dallas, Richardson, TX 75080
†Department of Electrical Engineering, Yale University, New Haven, CT 06520
‡Department of Electrical and Computer Engineering, San Diego State University, San Diego, CA 92115

***Abstract*—Growing demand for more powerful yet smaller devices has resulted in continuous scaling of fabrication technologies. While this approach supports aggressive design specifications, it has resulted in tighter constraints for circuit designers who face yield losses in analog / RF ICs due to process variation. Over the last few years, several statistical techniques have, therefore, been proposed to counter these losses and to recover yield through individual post-manufacturing calibration of each fabricated chip using tuning knobs. These techniques can be broadly classified as iterative or one-shot calibration methods, with the latter having the benefit of being faster and, therefore, more likely to be cost-effective in a high volume manufacturing (HVM) environment. In this paper, we first put three previously proposed one-shot statistical calibration methods to the test using a custom-designed tunable LNA, which was fabricated in IBM's 130nm RF CMOS process. We, then, introduce an improvement to the tuning knob selection criterion, which applies to all three methods, increasing their effectiveness. Finally, we demonstrate the efficacy of a previously proposed approach which uses simulation data and Bayesian model fusion in order to reduce the number of chips required for training the statistical models employed by the three one-shot calibration methods.**

## I. Introduction

As semiconductor manufacturing technologies continue to scale down transistor sizes in order to produce smaller and more powerful integrated circuits (ICs), the impact of process variations on the fabricated circuits becomes increasingly prominent. In analog/RF ICs, this impact is predominantly visible as a wider and less controlled distribution of circuit performances. This, in turn, results in increased probability that a device may be pushed outside its specifications by process variation, causing higher yield loss in newer technologies. Conservative design has traditionally been employed to mitigate this risk, yet at the cost of sacrificing performance. Alternatively, in an effort to support aggressive performance-driven design while ensuring satisfactory yield, post-manufacturing calibration has recently received extensive attention [1]–[8].

The underlying premise of such methods is that the impact of process variations may be counteracted through tuning knobs, which are included in the design and are individually tuned for every device after its fabrication, in order to bring its performances within the design specifications. The chosen knob settings are then permanently stored in some form of on-chip non-volatile memory (i.e., Flash, EEPROM, OTP, etc.) before the device is deployed. Thereby, aggressive performances and high yield can both be met.

Selecting an optimal knob setting for each device, however, is a challenging problem, due to the complex interplay between knob positions, process variations, and device performances. Compounding the problem, the realities of analog/RF IC test costs indicate that device performances can only be measured explicitly at most once, if profit is to be expected. As a result, post-production calibration needs to rely on alternate low-cost measurements in order to be viable.

Accordingly, following the general paradigm of *alternate test* [9], a significant number of statistical calibration methods have appeared in the literature, based on simple measurements from either intrusive or non-intrusive on-chip sensors. Generally, these methods can be divided into *iterative test-&-tune* [1]–[5] and *one-shot* [6]–[8] approaches. The former typically perform a directed search, each time selecting a new knob setting and repeating the low-cost measurements, seeking to optimize some criterion in the space of device performances. The latter aim at the same objective, yet rely on only one set of low-cost measurements for a single knob setting, based on which they select the final knob setting for a device. Since knob tuning is performed while the device is interfaced with expensive Automatic Test Equipment (ATE), one-shot statistical calibration approaches are more appealing in an HVM environment and are, hence, the focus of this study.

In an effort to better understand the strengths and shortcomings of one-shot statistical calibration methods, in this work we perform a comparative study of three such methods [6]–[8], which have recently appeared in the literature. This study is performed using a custom-designed test circuit which was fabricated in IBM's 130nm RF CMOS process. We, then, propose a modification in the criterion that these methods use for selecting the optimal knob setting and we demonstrate that the new criterion helps all three methods in achieving lower yield loss. Finally, we investigate the utility of a previously proposed method [10] in reducing the number of chips needed for training the statistical models employed by the three one-shot calibration methods. This is achieved through the use of simulation data, which is combined with measurements from a small number of chips using Bayesian model fusion.

The remainder of this paper is organized as follows: In Section II, we briefly review the three one-shot statistical calibration methods which are studied herein. In Section III, we introduce the experimental platform, including the fabricated IC, the evaluation board, and the obtained measurements. In


978-1-4673-6578-9/15/$31.00 

TABLE I: Summary of one-shot calibration methods

| Tuning method | Sensor type | Predicted parameter | Orthogonality assumption |
|---|---|---|---|
| [6] | Intrusive | Knob setting | Yes |
| [7] | Intrusive | Performances | Yes |
| [8] | Non-intrusive | Performances | No |

Section IV, we provide details of the comparative study setup, including the optimization criterion and the metric used for evaluation. Results are presented in Section V. A limitation of the optimization criterion employed by the three one-shot calibration methods is highlighted in Section VI, where an improved version is introduced and experimentally evaluated. Finally, the effectiveness of simulation data and Bayesian model fusion in reducing the number of chips required for training the statistical models is demonstrated in Section VII.

## II. One-shot calibration methods

One-shot statistical calibration methods are characterized by two main features: (i) they only rely on a set of low-cost measurements obtained for a single setting of the tuning knobs, and (ii) they employ statistical learning in order to train models which are used for selecting a final knob position based on the low-cost measurements. The key advantage of one-shot calibration methods is that they minimize the time an IC spends on the ATE for knob tuning, as only one set of low-cost measurements is needed[1]. Their primary objective is to recover yield, by bringing the performances of out-of-specification ICs back within their acceptable range through the use of tuning knobs. Additional objectives, such as power minimization, may also be considered during this post-manufacturing calibration stage.

Table I summarizes the features of three popular one-shot calibration methods, which are comparatively studied in this work. A brief description of these methods is provided in the following subsections, while more details can be found in the original publications introducing these methods. We note that these statistical approaches (i) may use either intrusive or non-intrusive sensor measurements, (ii) may predict either directly the best knob setting for a given optimization criterion or first predict performances for all knob settings and then select among them based on the optimization criterion, and (iii) may rely on various orthogonality assumptions regarding the impact of knobs and process variations, as we further explain below.

### A. DSP-driven self-tuning

Chronologically, the first one-shot calibration method was described in [6], with the objective of reducing test cost while recovering yield. This method relies on training a regression model which maps the low-cost alternate measurements to the optimum knob setting. In particular, a gradient-descent search algorithm is employed to find the optimum knob setting for a given set of low-cost alternate measurements such that a predefined cost function is minimized. In its simplest form, this cost function can be defined as the difference between the desired and the actual performances of the device. Once the model is trained, for new devices to be calibrated, the low-cost measurements are taken and the regression model is used to predict the optimum knob setting. The direct prediction of this optimum knob setting enables a one-shot calibration approach based on the low-cost alternate measurements. In [6], the output transient signal from an envelope detector at the output of a CMOS LNA was used as the low-cost measurement and bias currents/voltages were used as tuning knobs for calibrating third-order inter-modulation distortion, gain and supply current. A potential limitation of this method is that it assumes orthogonality between the impact of knobs on circuit performances, which is a desirable attribute but often difficult to achieve in practice when designing the knobs of a tunable analog/RF IC.

### B. Mid-point alternate test-based calibration

The knob orthogonality assumption was revoked in the method proposed in [7], which approached the problem from a different angle. This method, however, makes another assumption, namely that process variations and tuning knobs have an orthogonal impact on device performances [7]. Besides providing an additional degree of freedom to counter the impact of process variation, this orthogonal relationship offers another significant advantage: the regression model trained under this condition requires fewer measurements, thereby simplifying test setup and reducing test time and cost. The calibration process is again divided into a training phase and a tuning phase. During training, both performances and low-cost measurements are taken for a sample of devices. Regression models predicting performances for every knob setting based on low-cost measurements for a single knob setting (i.e., the "mid-point") are then trained using this data. During tuning, the low-cost measurements are taken for the mid-point knob position and the regression models are used to predict performances for all the knob combinations. Finally, the best knob setting is chosen for a given optimization criterion, which involves yield maximization and power minimization. In [7], this method was demonstrated using a tunable LNA circuit with bias voltage knobs and a peak detector with DC output as the low-cost measurement sensor, in order to calibrate input reflection, noise figure, gain, output reflection and power.

### C. Non-intrusive sensor-based calibration

The key novelty of the third one-shot method, which was introduced in [8], is the use of non-intrusive sensors (NISs) for predicting and calibrating the performances of an RF circuit. NISs are copies of circuit structures which are not electrically connected to the main circuit, but which are placed in very close proximity to the original structures. As a result, they are subjected to the same process variation as the original circuit and can, therefore, serve as a proxy for characterizing it. In all other respects, this method resembles the one described in [7]: a trained regression model is used to predict performances for

[1]Once the knob positions are chosen, standard analog/RF specification testing may be performed in order to confirm that the calibrated circuit meets its specifications, as discussed in [7], [8].

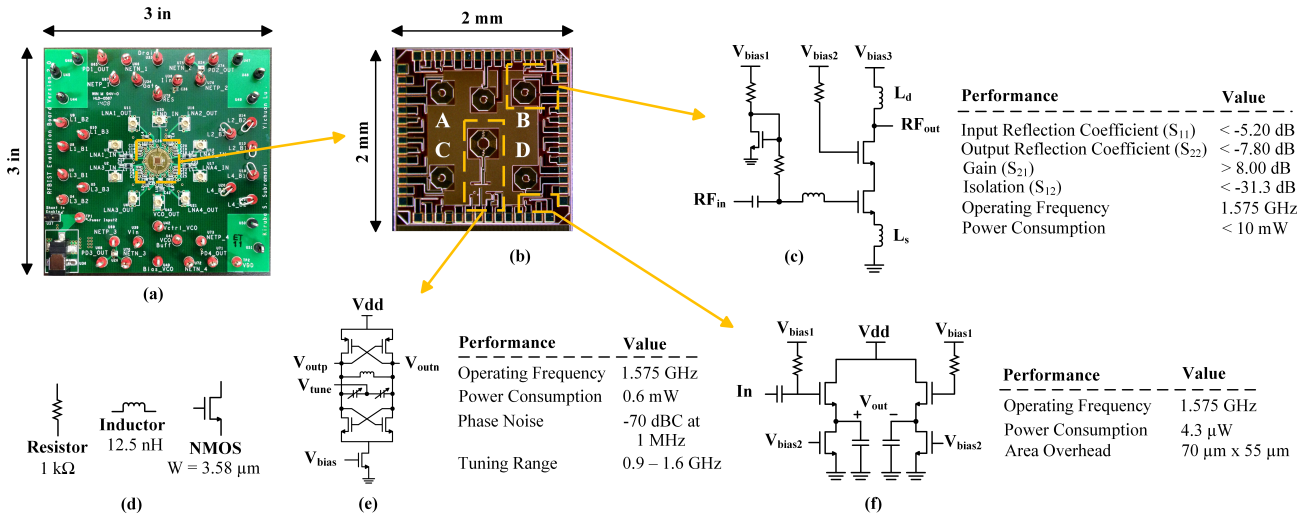


Fig. 1: Experimental Platform: (a) Evaluation board (b) Micro-photograph of fabricated die (c) LNA schematic (d) Non-intrusive sensors (e) VCO schematic (f) Peak detector schematic

all knob settings based on the NISs, and the optimum knob setting is chosen based on an optimization criterion. Evidently, since NISs are independent of the tuning knob position, they only need to be measured once. In [8], this method is demonstrated on an RF power amplifier (PA) designed in ST Microelectronics' 65nm CMOS technology. Power supply and bias voltages of the two PA stages are used as tuning knobs and two dummy resistors/capacitors along with the DC currents drawn by two dummy PA stages are used as non-intrusive sensors for calibrating 1-dB compression point, maximum power added efficiency, gain and power consumption.

## III. Experimental Platform

A comparative study of the aforementioned three methods requires a common evaluation platform. For this purpose, we designed and fabricated a test chip in IBM's 130nm RF CMOS process. The tunable RF circuit is a Low Noise Amplifier (LNA) with three bias voltage tuning knobs. The circuit also includes one intrusive sensor and three non-intrusive sensors, which are used for driving the knob selection procedure, as well as an on-chip Voltage Controlled Oscillator (VCO), which is used for stimulus generation. The fabricated IC is housed in a custom-designed evaluation board to interface with power supplies and external measurement instrumentation. Figure 1 shows the evaluation board, the micro-photograph of the fabricated IC, as well as the schematics and specifications of the various components included in the IC. Details about each of these components are provided in the following subsections.

### A. RF circuit

The RF circuit used in this platform is a Cascode LNA with inductive source degeneration, the schematic and specifications of which are shown in Figure 1(c). While this is a fairly straightforward design, it is an excellent candidate for our comparative study as it has the sophistication to demonstrate the impact of tuning knobs on performance characteristics. The input and output matching networks for the LNA are included on the evaluation board. Load inductance (Ld) is implemented using an on-chip spiral inductor. In this architecture, source degeneration inductance (Ls) allows the designer to match the input impedance to a $50\Omega$ load and is implemented by taking advantage of the parasitic bond-wire inductance.

### B. Tuning knobs

Three tuning knobs were carefully selected to establish full control on the LNA's performances. These knobs directly control the power supply and the bias voltages of the circuit. The fabrication technology used for this design features a supply voltage of 1.2V and a breakdown voltage of 1.6V. Evidently, flexibility in the range and granularity of the tuning knobs is highly desired for calibration purposes, since it allows exploration of a large area in the space of performances. From a practical implementation perspective, however, the lower limit for the knob values is directed by their ability to support acceptable performances, while the upper limit is governed by power consumption and breakdown voltage concerns. With these criteria in mind, we designed each knob in our circuit such that it can be swept from 0.8V to 1.4V in steps of 0.1V. This creates $7^3 = 343$ tuning knob combinations for the LNA circuit.

### C. Measurement sensors

As discussed earlier, relying on actual specification tests for post-manufacturing performance calibration of RF ICs is prohibitively expensive. Rather, measurements from low-cost on-chip sensors, either intrusive or non-intrusive, are used in all one-shot statistical calibration methods that can be found in the literature, in a process that is generally known as alternate test [9]. For this purpose, both intrusive and non-intrusive sensors were designed and integrated in our test chip. The first sensor is a peak detector (PD) circuit, which was designed and interfaced to the output of the LNA. The schematic and specifications are shown in Figure 1(f). This peak detector provides a DC output which is equivalent to

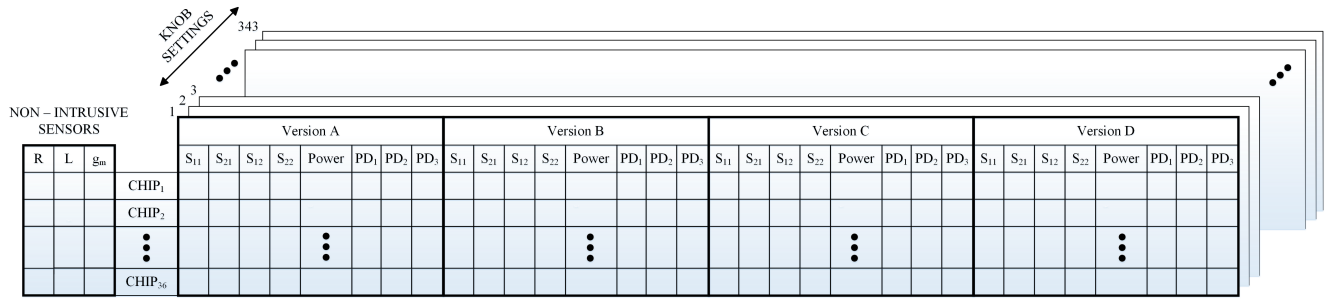


Fig. 2: Dataset Collection

the peak value of the transient signal. Since the PD has an electrical connection to the parent LNA circuit, it is considered an intrusive sensor. Simple non-intrusive sensors in the form of test coupons, including a resistor, an inductor and an NMOS transistor copied from the original LNA circuit are also integrated in our design, as shown in Figure 1(d). While these are not electrically connected to the LNA, they provide useful information regarding the impact of process variations on each fabricated IC. By placing these non-intrusive sensors nearby the LNA, it is likely that the process variations affecting the LNA will be adequately captured, i.e. the LNA performances will be strongly correlated to the measurements of these non-intrusive sensors [11]. We note that our non-intrusive sensors are not as sophisticated and are not as closely placed to the LNA as the ones described in [8]. Nevertheless, they reflect the general intention of the method described therein and end up performing very well as predictors for the knob calibration procedure.

### D. Stimulus generator

In order to excite the LNA and collect the response of the intrusive PD sensor, we also designed and integrated a symmetrical cross-coupled LC VCO on the chip. The schematic and specifications of the VCO, which serves as an on-die stimulus generator, is shown in Figure 1(e). The $V_{tune}$ pin provides flexibility in controlling the VCO's output frequency to a range of different values, thereby facilitating multiple observations through a single sensor.

### E. Fabricated chip

Statistical methods require a considerable number of devices for training and validation, in order to be fairly evaluated. Nevertheless, due to practical considerations regarding cost for fabricating and time for characterizing each circuit, our study is limited to a set of 40 chips, which were fabricated in IBM's 130nm RF CMOS process through MOSIS. However, in order to increase the number of samples, each chip was designed to accommodate 4 LNAs, each with its own intrusive PD sensor, as can be observed in the die micro-photograph of Figure 1(b). The layout of the four versions A, B, C, and D of the LNA and their PDs is identical, except for rotation across the X and/or Y axis, which was done to ensure the same routing distance to pins. In general, differences among these four replicas in terms of routing and placement are kept to a minimum, in order to avoid introducing additional variation through the design. Therefore, differences within the set of 160 LNAs can be largely attributed to process variations. We also note that there is only one VCO on each chip, which can drive the input of any of the four LNAs, and only one set of non-intrusive sensors, as these are not electrically connected to the LNA.

### F. Evaluation board

Interfacing the fabricated IC with external test equipment for carrying out our experiments is achieved through a custom-designed Printed Circuit Board (PCB) on which each IC is mounted. Figure 1(a) shows the evaluation board which we designed and fabricated for this purpose. Apart from housing the IC, the board also holds the input and output matching networks, ultra-miniature connectors for RF interfaces, as well as DC connectors for controlling the tuning knob values.

### G. Dataset

Each of the 40 fabricated chips was extensively characterized through an experimental setup which was used to measure the four performances ($S_{11}$, $S_{21}$, $S_{12}$, $S_{22}$) and the power consumption ($Power$) of each LNA for every one of the 343 knob settings. The intrusive sensor readings ($PD_1$, $PD_2$, $PD_3$) for 3 different stimulus frequencies generated by the VCO and for every one of the 343 knob settings were also recorded through a second experimental setup. Finally, the three non-intrusive sensor readings ($R$, $L$, $g_m$) were also collected for each chip. In our bench-top setup, tuning knobs were controlled by Agilent DC power supplies and all instruments were interfaced to National Instruments' LabVIEW [12], through which the entire measurement process was automated. By the end of this process, a single document containing all measurement results for all possible knob combinations for the 160 LNAs was created. Among the 40 chips, 4 were excluded from the study due to defective bond wires, leaving us with the dataset shown in Figure 2, which includes 36 chips with 4 LNAs each, for a total of 144 LNAs. This dataset constitutes the basis for comparing the three one-shot statistical calibration methods through a process that involves splitting it in training and validation subsets, as we explain in the following section.

## IV. Comparative Study Setup

In this section, we provide details regarding the setup used for this comparative study.

### A. Methods

The three one-shot statistical calibration methods described in Section II are considered. We note that methods [6] and [7] use the three intrusive sensor measurements $PD_1$, $PD_2$, and $PD_3$ as predictors for selecting a knob setting, while method [8] uses the three non-intrusive sensor measurements $R$, $L$, and $g_m$. We also note that methods [7] and [8] first predict the performances for all possible knob settings and then select among them the most appropriate for a given optimization criterion, while method [6] incorporates the optimization criterion in the prediction model and reports the knob setting that best approximates it.

### B. LNA set

The entire dataset shown in Figure 2 is used in this study. The 36 working chips (144 LNAs) are split into 20 chips (80 LNAs), which are used for training the statistical models in each of the three one-shot calibration methods and 16 chips (64 LNAs), which are used for evaluating the effectiveness of the trained models. All experiments are repeated 10 times, each time splitting randomly the available chips into training and validation sets, and all reported results are averaged over these 10 cross-validations.

### C. Optimization criterion

The optimization criterion used in this study reflects the main objective of the original works describing the three one-shot statistical calibration methods [6]–[8], i.e. yield optimization[2]. Specifically, we seek to maximize the probability that the chosen knob setting will result in a specification-compliant circuit for the four performances ($S_{11}$, $S_{21}$, $S_{12}$, and $S_{22}$) and the power consumption ($Power$). Among the –potentially many– knob settings that meet this objective, however, we seek to choose the one that brings the performances of the circuit as far from the specification planes as possible for one-sided specifications and as close to the middle of the specification range as possible for two-sided specifications, thereby **maximizing robustness**.

Formally, let $(sp_1, sp_2, \ldots, sp_k)$ be the design specifications of a circuit[3] and $(i_1, i_2, \ldots, i_k)$ be the performances of a fictitious, ideal device, which is the best we can hope to get for a given design and technology (in our case, we use the best performances across all 144 LNAs and 343 knob settings as a proxy for the ideal). Also, let $(p_1, p_2, \ldots, p_k)$ be the measured performances for a chosen knob setting, after normalization in the range $((sp_1, i_1), (sp_2, i_2), \ldots, (sp_k, i_k))$, respectively, so $0 \leq p_j \leq 1$, $\forall j = 1, \ldots, k$. We define robustness as:

$$Robustness = \sqrt{\sum_{j=1}^{k} (p_j)^2} \tag{1}$$

[2]We note that all three methods support additional optimization criteria, such as choosing among the knob settings that yield a specification-compliant circuit the one that minimizes power consumption. Our future work will extend this comparative study to encompass such advanced criteria.

[3]Since all performances considered in our circuit only have one bound, we limit the definition of this metric to one-sided specifications. Extending it to two-sided specifications is straightforward, as described in [7].

TABLE II: Yield Loss Results

| Method | Before Tuning | [6] | [7] | [8] |
|---|---|---|---|---|
| Yield Loss Average | 19.9/64 | 6.8/64 | 3.6/64 | 3.8/64 |
| Standard Deviation | 2.234 | 3.824 | 1.506 | 2.079 |

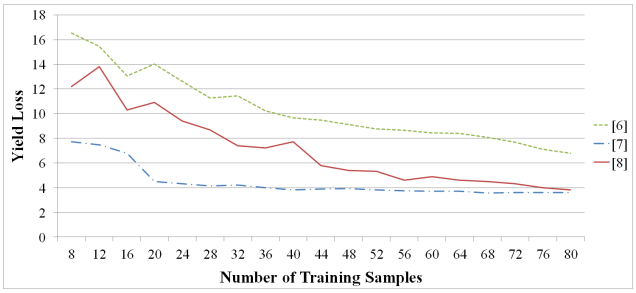


Fig. 3: Training Set Size Impact on Yield Loss

and seek to select the knob setting that minimizes this value. This metric is essentially the Mahalanobis distance between the device performances for the chosen knob setting and the ideal device performances, which is simply a covariance-scaled version of the Euclidean distance to ensure that each performance is weighted equally.

### D. Metric

The metric used for evaluating the three methods with respect to the above optimization criterion is defined below:
**Yield Loss:** Yield Loss reflects the limitations of a method in selecting a knob setting that results in a circuit which passes all specifications. Let $N_p$ be the number of circuits in a dataset under consideration, for which there exists at least one knob setting for which the circuit passes all specifications. Also, let $N_h$ be a subset of these $N_p$ circuits, for which the method under evaluation selects a knob setting that results in at least one failing specification. We define yield loss as:

$$Yield\ Loss = N_h / N_p \tag{2}$$

## V. Results

The first result that we report is the Yield Loss sustained by each of the three one-shot calibration methods, when training is performed on 20 chips (80 LNAs) and validation on 16 chips (64 LNAs). The average and the standard deviation over 10 cross validation runs are reported in Table II. The second column of the table reports the results before tuning, assuming that the knob setting is in the mid-point position, while the third through fifth columns report the result after the knobs are tuned to the positions selected by each of the three methods.

As may be observed, all three methods achieve a very significant improvement over the baseline, with [7] performing the best among the three. We should clarify, however, that the reason due to which the method described in [6] is less effective is because it assumes knob orthogonality, which is not obeyed by the design of knobs in this case study and is, in general, hard to satisfy. We would also like to point out again that the non-intrusive sensors used in this case study are less sophisticated than the ones described in [8]. Nevertheless, the ability of non-intrusive sensors to successfully tune the LNAs despite not being electrically connected to them is impressive.

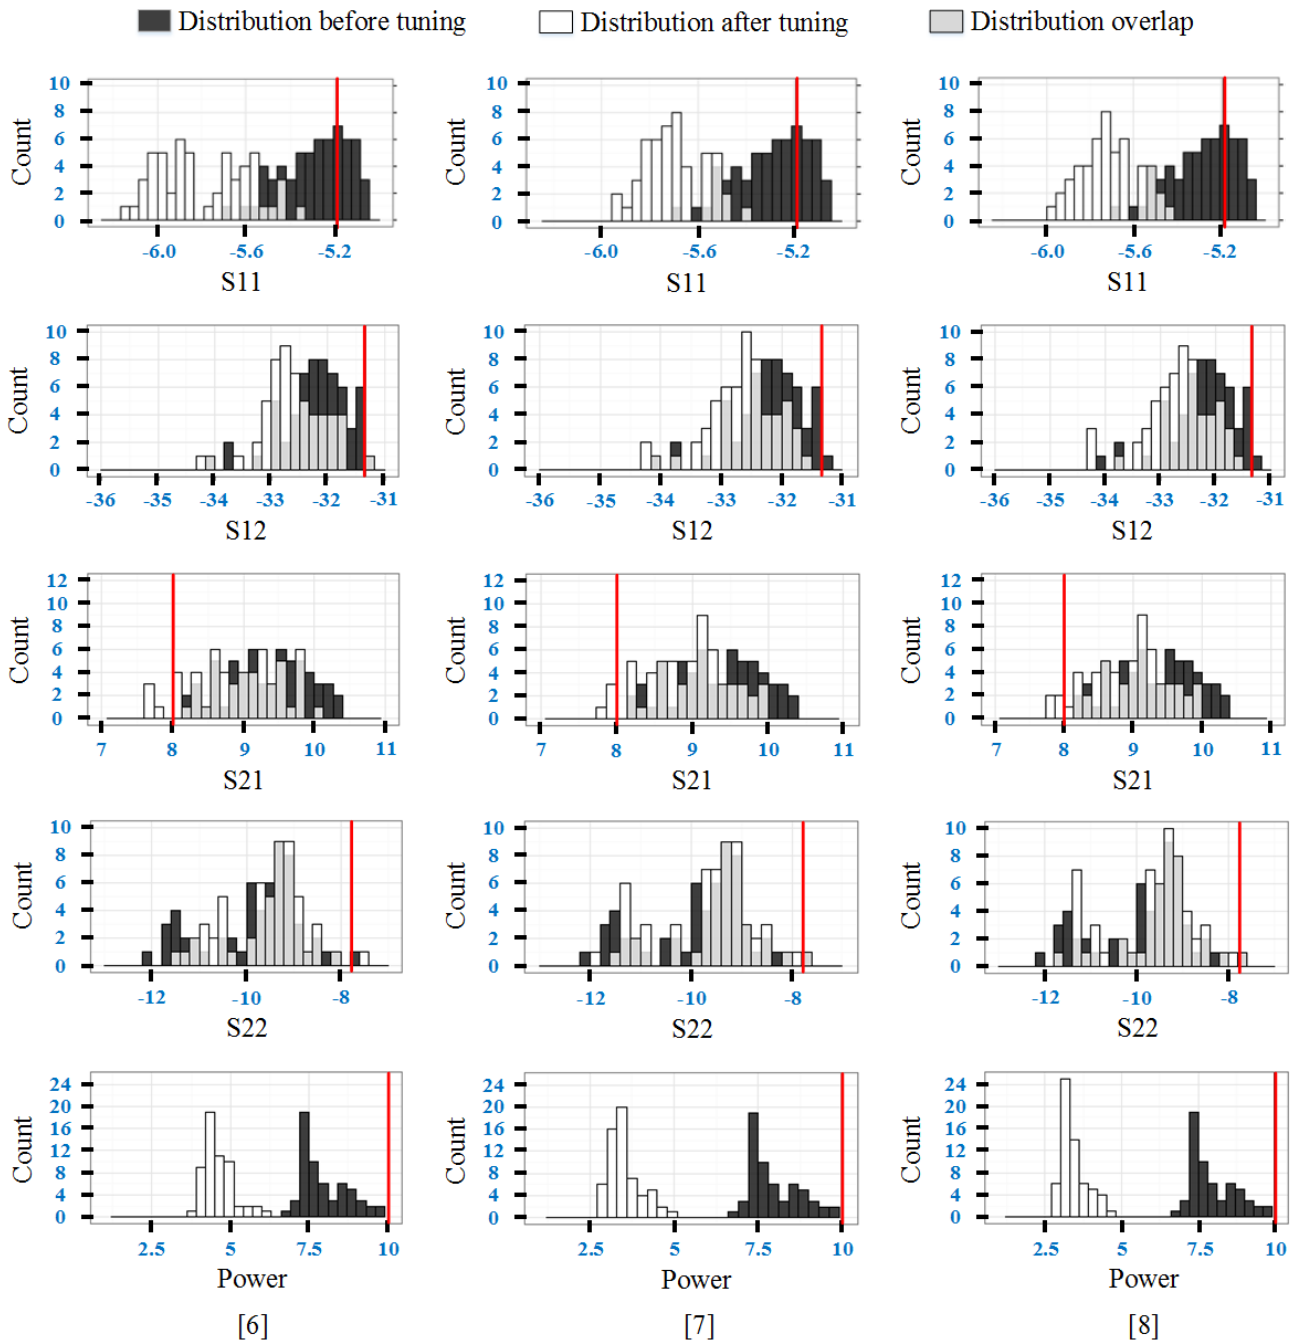


Fig. 4: Performance Distribution Before and After Calibration with Each of the Three Methods

The second result that we report is the impact of the training set size on the Yield Loss metric. Specifically, in Figure 3 we plot the Yield Loss of each of the three one-shot calibration methods as a function of the number of LNAs used for training, starting with 8 LNAs and increasing in steps of 4 until we reach the total of 80 devices in the training set. Once again, the results are averaged over 10 cross validations. As may be observed, the method described in [7] requires approximately 36 devices in the training set to achieve yield loss below 4/64 and its further improvement is very flat after that. In contrast, the methods described in [6] and [8] continue to improve as the training set size increases. For [6], this is attributed to the fact that the knobs in our design do not obey the orthogonality assumption, hence it takes more samples to capture the complex relation between the sensor measurements, the knob positions, and the LNA performances. For [8], this is attributed to the fact that the non-intrusive sensors do not directly reflect the impact of knobs, as they are not electrically connected to the LNA, hence it takes again more samples to capture the indirect relation between process variation and knob impact on performances.

The last result that we report is the impact of knob tuning on the distribution of performances for the entire population of LNAs that are subjected to calibration. Specifically, in Figure 4 we contrast the initial distribution of the four performances and power consumption prior to calibration, assuming that the

knobs are set to the mid-point position, to the distribution after calibration with each of the three methods. The red vertical line in each histogram represents the specification value which was listed in Figure 1. What is evident in the results is that all three methods are very effective in improving and bringing $S_{11}$ well within the acceptable range, even though a good percentage of LNAs fail this specification in the mid-point knob position. Improvement is also observed on $S_{12}$, where [7] and [8] bring all LNAs within the acceptable range, while [6] succeeds in all except one LNA. Gain in the calibrated circuits, however, exhibits an interesting behavior: all three methods select a knob position which results in some LNAs failing $S_{21}$, even though they were specification-compliant in the mid-point knob position. We attribute this effect to a limitation of the optimization criterion used for knob selection; indeed, as we will explain further in Section VI, this is a side-effect of the way robustness is quantified during knob selection. In this particular case, the benefit of pushing $S_{11}$ higher and $Power$ lower outweighs the risk of bringing $S_{21}$ closer to its limit, where statistical error of the regression models has a higher likelihood of pushing it in the failing zone. A solution that ameliorates this effect is also proposed in Section VI. Regarding $S_{22}$, all three methods only slightly distort the overall distribution, mildly pushing it closer to the specification limit due to the reduced $S_{21}$, which results in lower output impedance. Nevertheless, no more LNAs fail $S_{22}$ after tuning, as compared to the mid-point knob setting. Finally, $Power$ is drastically and almost equally effectively improved by all three methods, as seen in the last row of histograms.

## VI. Improved Knob Selection Criterion

As mentioned above, the optimization criterion of maximizing robustness, as defined in Equation 1, has an inherent limitation which results in unnecessary yield loss. To demonstrate this limitation, consider the simple example depicted in Figure 5, where two knob settings exhibit two pairs of normalized performances $p_1$ and $p_2$ for a circuit. In the first knob setting, the two performances are (0.25, 0.25), therefore robustness is calculated at $\sqrt{(}0.125)$. In the second knob setting, the two performances are (0.75, 0.05), therefore robustness is calculated at $\sqrt{(}0.565)$. Naturally, the given optimization criterion of maximizing robustness selects the second knob setting. However, such selection carries far more risk, because these performances are not measured but rather predicted through trained regression models, which are subject to statistical error. Such error is far more likely to result in a specification violation of a performance that is very close to the limit (as is 0.05 in the second knob setting in our example) than a performance that is further away. This explains why some LNAs fail $S_{21}$ after calibration with any of the three methods, even though they were passing this specification at the mid-point knob setting, as was shown in Figure 4.

To alleviate this problem, the contribution of a performance to the overall optimization criterion should be slowly tapered off above a certain value. This would, essentially, prevent

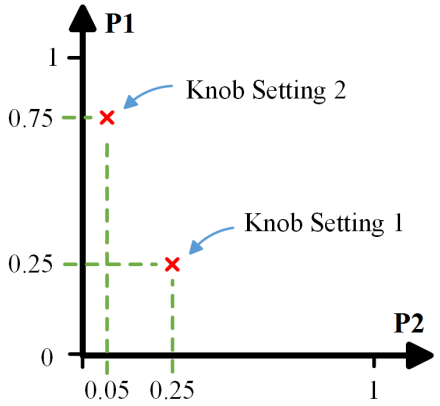


Fig. 5: Example Demonstrating Criterion Limitation

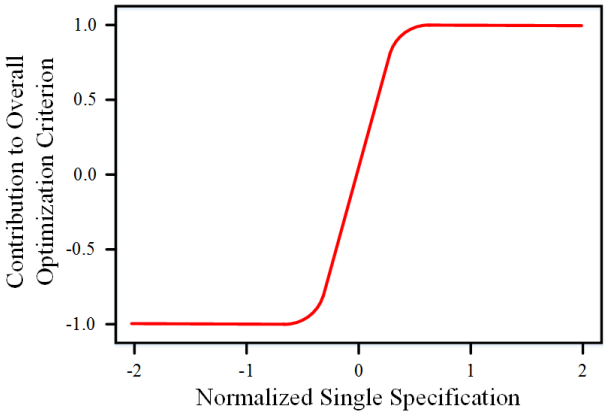


Fig. 6: Sigmoid Function for Tapering Off Contribution of Performances to Overall Optimization Criterion

overly pushing a strong performance at the expense of bringing other performances close to their limits. In other words, a knob setting chosen using this method would seek to tune a circuit in a position where all performances contribute to robustness in a more balanced fashion. This, in turn, reduces the risk that yield loss will ensue due to statistical error in the regression models used for predicting performances. To achieve this objective, we subject each performance to a cumulative distribution function $P_{norm}$ which has a sigmoid shape, as shown in Figure 6. The new optimization criterion is, then, defined as:

$$New\_Robustness = \sum_{j=1}^{k}(2 * P_{norm}(3.5 * \sqrt{2} * p_j) - 1) \quad (3)$$

which we, again, seek to maximize.

Unlike the Mahalanobis distance used in Equation 1, the $P_{norm}$ function used in the new criterion reduces the relative additional contribution of a performance as its value increases. In other words, the weight of a performance is reduced as its value increases. To understand this better, let us revisit the example of Figure 5. For the first knob setting whose performances are (0.25, 0.25), the cumulative distribution function gives an output of (0.9867, 0.9867) and the new distance value is 1.9734 which is the sum of the two numbers. Similarly, for the second knob setting whose performances are (0.05, 0.75), the corresponding output is (0.1955, 0.9998)

TABLE III: Yield Loss Results (New Criterion)

| Method | Before Tuning | [6] | [7] | [8] |
|---|---|---|---|---|
| Yield Loss Average | 19.9/64 | 5.7/64 | 2.9/64 | 3.4/64 |
| Standard Deviation | 2.234 | 2.921 | 1.287 | 1.647 |


Fig. 7: Training Set Size Impact on Yield Loss (New Criterion)

and its new distance value is 1.1953. Therefore, the first knob setting, which has more balanced performances, is selected.

Jointly considering magnitude and balance of performances improves the effectiveness of one-shot calibration methods in reducing Yield Loss by selecting appropriate knob settings. Table III reports the Yield Loss sustained by each of the three methods when the new criterion is used. As may be observed when comparing to Table II, all three methods benefit from this improved knob position selection criterion, enjoying both lower Yield Loss average and lower standard deviation across the 10 cross validation repetitions.

Figure 7 plots the Yield Loss of each of the three methods as a function of the number of LNAs used for training when the new criterion is used. Again, in comparison to the plot in Figure 3, all three methods benefit by the improved knob position selection criterion, converging to their best performance at a smaller training set size. The positive impact on the non-intrusive sensor-based method [8] is remarkable.

Finally, in Figure 8 we contrast the initial distribution for the four performances and power consumption prior to calibration, assuming that the knobs are set to the mid-point position, to the distribution after calibration with each of the three methods when the new criterion is used. A careful examination of these histograms vis-a-vis the ones in Figure 4 reveals that all three methods are able to reduce the number of devices that fail certain specifications, while not significantly changing the overall histogram distributions. Specifically, the method in [6] reduces the number of LNAs that fail $S_{21}$ from 4 to 2 and the number of LNAs that fail $S_{12}$ from 1 to 0 on average, the method in [7] reduces the number of LNAs that fail $S_{21}$ from 4 to 2 on average, and the method in [8] reduces the number of LNAs that fail $S_{21}$ from 4 to 3 on average.

In short, the new optimization criterion provides a more effective mechanism for selecting a knob position that not only maximizes robustness but also reduces the risk that a calibrated device will fail a specification due to statistical error in the regression models which predict the performances for the various knob settings.

## VII. Training Set Reduction via BMF

In an effort to reduce the size of the training set required for learning the regression models which predict the performances of a circuit for the various knob settings, an elegant method was proposed in [10]. By employing a Bayesian Model Fusion (BMF) technique, the coefficients of a regression model, which is learned at an earlier stage (i.e., from simulation or previous tape-out data), are used as prior knowledge to estimate the new regression model. Using BMF and prior knowledge of an early-stage model, it is shown therein that only a few samples are sufficient to achieve a reasonable accuracy level in the estimated regression model. We note that this method makes no orthogonality assumptions regarding the impact of knobs and/or process variations on circuit performances and, while it is more complex, it can significantly reduce the number of actual chips needed for training. In [10], this method was demonstrated using a voltage-controlled oscillator (VCO) designed in a commercial 32 nm CMOS process. Bias voltage was used as a tuning knob for calibrating phase noise, based on simple measurements such as oscillation frequency/amplitude and bias current.

To evaluate its effectiveness, we adapted this technique to each of the three one-shot calibration methods contrasted in this study. Specifically, we used Monte Carlo simulation of the post-layout extracted netlist of our circuit, thereby generating 100 synthetic chips, each including 4 LNAs, along with the VCO and the intrusive and non-intrusive sensors shown in Figure 1. Using these synthetic circuits, we then generated a measurement dataset similar to the one shown in Figure 2. This synthetic dataset was used to train the regression models for each of the three one-shot calibration methods and the trained models were, then, applied to 64 actual LNAs as the validation set. The sustained yield loss[4] is 14.2/64, 15.8/64, and 14.7/64, for each of the three methods [6], [7], and [8], respectively, and is shown in Figure 9, as the starting point of the three plots [6]+[10], [7]+[10], and [8]+[10]. Subsequently, we applied BMF, wherein we combined the synthetic data with an increasing number (from 4 to 32, in steps of 4) of actual measurements from the 80 LNAs in our training set, each time applying the trained regression models to the remaining 64 actual LNAs in order to evaluate efficiency. The entire experiment was repeated 10 times and the results were averaged and are reported in Figure 9.

Comparing to the plots shown in Figure 7, which we superimpose on Figure 9 to facilitate comparison, it is evident that the use of BMF, as proposed in [10], reduces the number of actual devices needed for training the regression models for all the one-shot calibration methods. Specifically, leveraging the prior knowledge obtained from Monte Carlo simulation data, even as small as a set of 4 actual LNA devices can be used to drastically reduce Yield Loss. As the number of actual devices from which measurements are fused in the model

[4]We note that in evaluating the positive effect of BMF in reducing training set size we used the new optimization criterion of Equation 3 for selecting a knob setting. However, the results are similar when the old optimization criterion of Equation 1 is used.

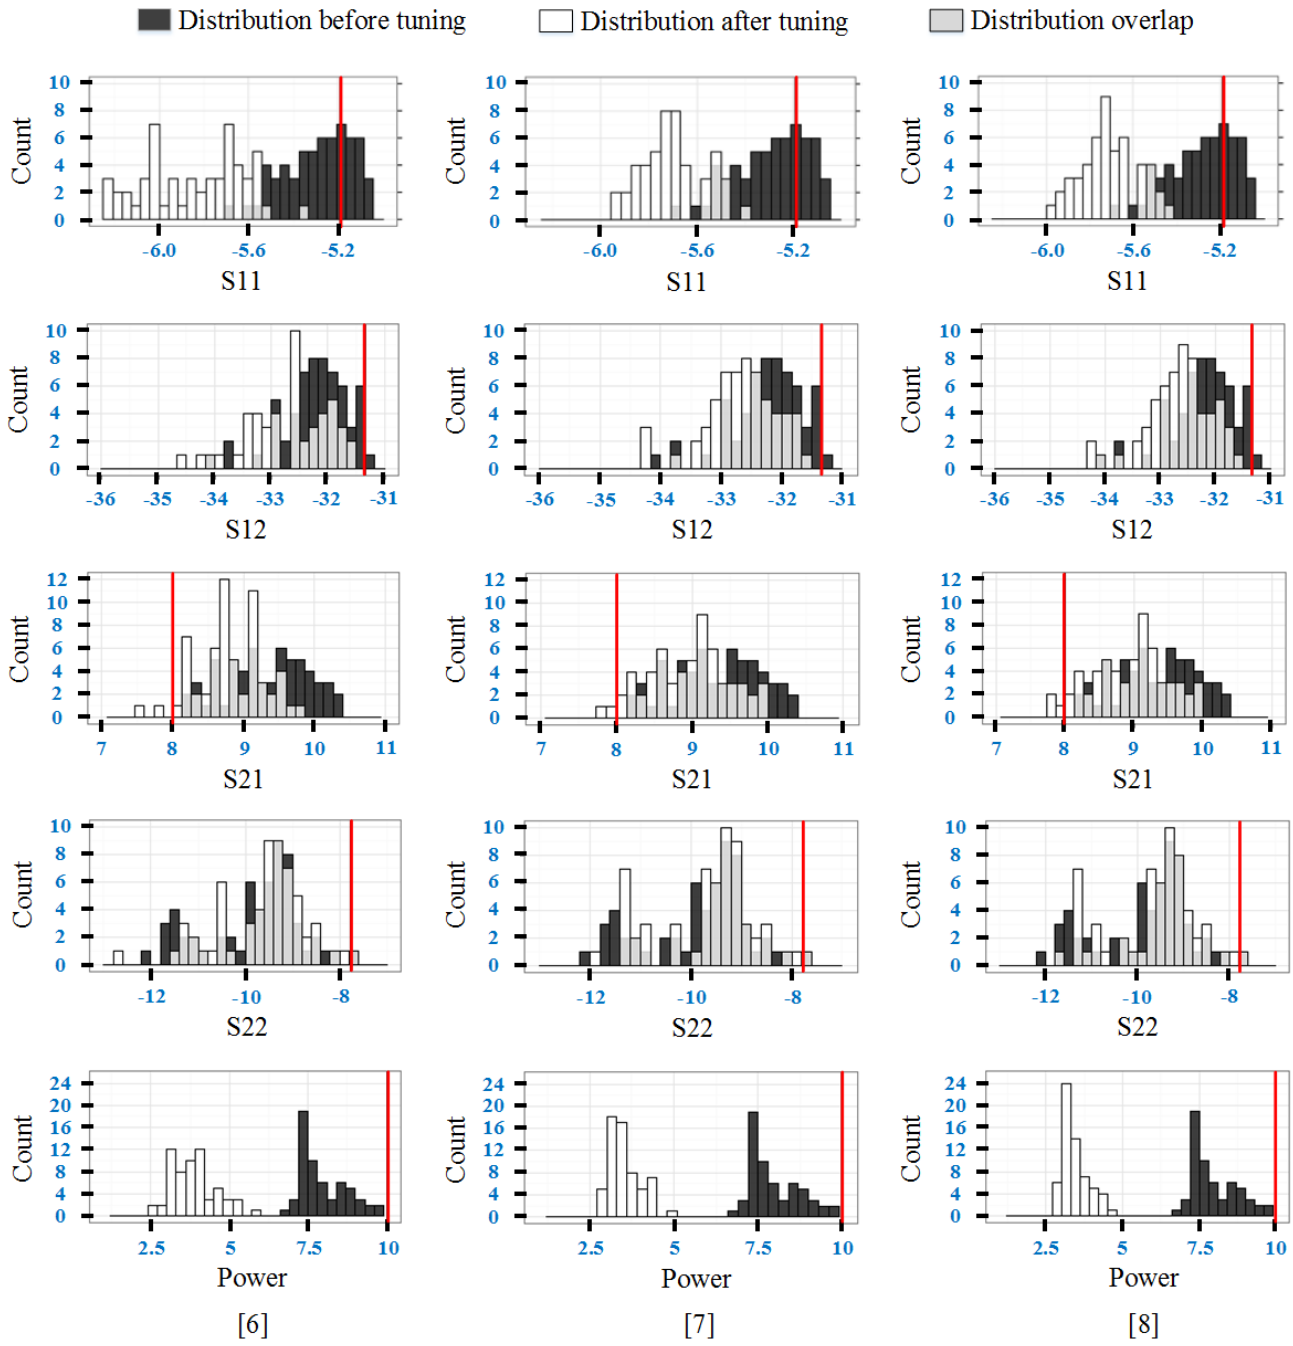


Fig. 8: Performance Distribution Before and After Calibration with Each of the Three Methods (New Criterion)

increases from 4 to 32, the utility of the synthetic data and the BMF method remains evident. Beyond that point, however, the plots converge, implying that the information provided by the actual measurements dominates, hence the pruned plot.

## VIII. Conclusions

One-shot methods for statistical post-manufacturing calibration of analog/RF circuits have emerged as an appealing option because they rely only on a single round of low-cost measurements, as opposed to the multiple rounds required by their iterative counterparts. Three such one-shot methods, namely [6]–[8], were comparatively studied in this work, using as a common evaluation platform a tunable test-chip, which was designed and fabricated in IBM's 130nm RF CMOS technology. The main conclusion reached by this comparative study is that one-shot statistical calibration is very effective in selecting a near optimal knob setting towards meeting the optimization criterion of yield loss minimization, while reducing the incurred calibration overhead to a single round of low-cost measurements. For this particular design and subject to the limitations discussed herein, the method in [7] exhibited a slightly superior ability in reducing Yield Loss. However, the method in [8] offers the competitive advantage of employing non-intrusive sensors while being only slightly less effective.

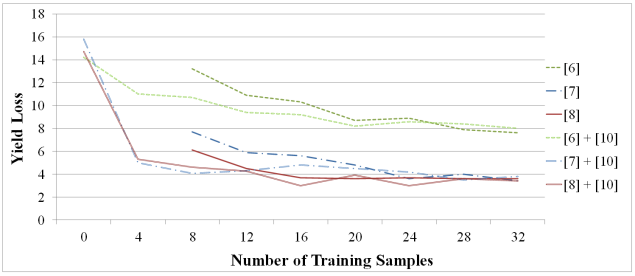


Fig. 9: Training Set Size Reduction via BMF (New Criterion)

Similarly, the method in [6] is faster and more effective in dealing with a large number of knobs, each having many settings, as it directly predicts the knob setting rather than first predicting the performances for every knob setting and then choosing among them. An additional conclusion of this study is that the optimization criterion used for selecting a knob setting needs to be carefully designed to accommodate the statistical error in predicting circuit performances from low-cost measurements, in order to avoid unnecessary yield loss. Lastly, this study also concluded that the Bayesian model fusion method proposed in [10] can effectively leverage simulation data in order to reduce the number of actual circuits required for training the regression models used by the three one-shot calibration methods compared herein.

## REFERENCES

[1] V. Natarajan et al., "Yield recovery of RF transceiver systems using iterative tuning-driven power conscious performance optimization," *IEEE Design & Test of Computers*, vol. 32, no. 1, pp. 61–69, 2015.

[2] S. Sun et al., "Indirect performance sensing for on-chip self-healing of analog and rf circuits," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 61, pp. 2243–2252, 2014.

[3] V. Natarajan et al., "Analog signature-driven postmanufacture multidimensional tuning of RF systems," *IEEE Design & Test of Computers*, vol. 27, no. 6, pp. 6–17, 2010.

[4] S. K. Devarakond et al, "Bist-assisted power aware self healing RF circuits," in *IEEE International Mixed-Signals, Sensors, and Systems Test Workshop*, 2009, pp. 1–4.

[5] V. Natarajan et al., "Act: Adaptive calibration test for performance enhancement and increased testability of wireless RF front-ends," in *IEEE VLSI Test Symposium*, 2008, pp. 215–220.

[6] D. Han, B. S. Kim, and A. Chatterjee, "DSP-driven self-tuning of RF circuits for process-induced performance variability," *IEEE Transactions on VLSI Systems*, vol. 18, no. 2, pp. 305–314, 2010.

[7] N. Kupp, H. Huang, P. Drineas, and Y. Makris, "Post-production performance calibration in analog/RF devices," in *IEEE International Test Conference*, 2010, pp. 245–254.

[8] M. Andraud, H.-G. D. Stratigopoulos, and E. Simeu, "One-shot calibration of RF circuits based on non-intrusive sensors," in *ACM Design Automation Conference*, 2014, pp. 1–6.

[9] P. N. 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] X. Li, F. Wang, S. Sun, and C. Gu, "Bayesian model fusion: a statistical framework for efficient pre-silicon validation and post-silicon tuning of complex analog and mixed-signal circuits.," in *ACM International Conference on Computer-Aided Design*, 2013, pp. 795–802.

[11] L. Abdallah, H.-G. D. Stratigopoulos, and S. Mir, "True non-intrusive sensors for RF built-in test.," in *IEEE International Test Conference*, 2013, pp. 1–10.

[12] National Instruments, "LabVIEW system design software," http://www.ni.com/labview/, 2014.