

Improving Analog and RF Device Yield through Performance Calibration

Nathan Kupp, He Huang, and Yiorgos Makris

Yale University

Petros Drineas

Rensselaer Polytechnic Institute

As the semiconductor industry continues scaling devices toward smaller process nodes, maintaining acceptable yields despite process variations has become increasingly challenging. Analog and RF circuits are particularly sensitive to process variations. This article discusses the challenges of cost-effective postfabrication performance calibration in such analog and RF devices and introduces a single-test, single-tuning-step method to constrain cost and complexity while reaping the benefits of a tunable design.

■ **ALTHOUGH TECHNOLOGY SCALING** has been consistently favorable for digital devices, enabling higher performance per area per watt, analog and RF devices have not necessarily benefited at the same pace. Analog and RF circuit design requires careful balancing of many parameters and is particularly sensitive to even the slightest perturbation. Moreover, such devices are typically employed in wireless mobile products, where low power consumption is critical, thus compounding the problem. As a result, ensuring that the performances of an analog device meet the design specifications has become increasingly challenging, particularly in light of the increasing process variation of smaller process nodes. To handle these challenges, analog and RF designers have traditionally resorted to conservative circuit design approaches, trading off some performance for higher yields and better variation tolerance.

Recently, postproduction performance calibration has emerged as a new defense for combating the increasing challenges of analog and RF design. The key idea is the addition of postproduction tunable components (or *knobs*) in the design, to support individual calibration of the performances of each fabricated device. With an appropriately selected set of knobs and performance calibration

algorithm, failing devices can be fine-tuned until their performances fall within the design specifications. Thus, by adjusting the knobs, some devices that would simply be discarded under the traditional analog test regime can now be salvaged, thereby recovering yield. In short, the benefit offered to the existing design and

test flow by this performance calibration approach is that it lets analog designers aggressively optimize high-performance ICs, while maintaining expectations of high yield.

In this article, we discuss the challenges of cost-effective postfabrication performance calibration in analog and RF devices and introduce a novel single-test, single-tuning-step method that substantially constrains cost and complexity while reaping the benefits of a tunable design. Furthermore, we describe a cost-benefit model to facilitate comparison with respect to current industry practice, and we discuss the method's potential as demonstrated on a tunable RF low-noise amplifier device designed and simulated in 0.18- μm RF CMOS.

Postproduction performance calibration

Despite its potential, postproduction performance calibration has not yet achieved widespread use, mainly because of the perceived implementation cost and complexity. Interestingly, it is not the knobs themselves that cause the slow adoption of calibration methods. The chosen knob settings can be easily and inexpensively stored on chip using nonvolatile memory trimming,^{1,2} as is commonly practiced in industry, thus making the calibration process transparent to users. Rather, it is

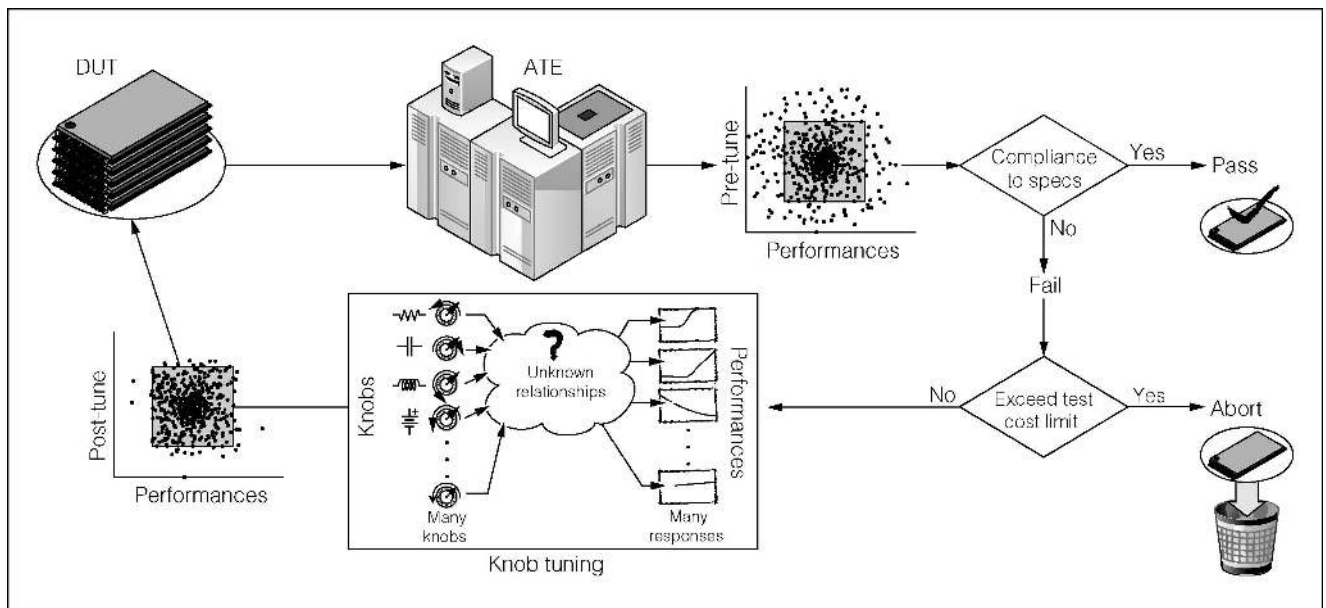


Figure 1. Iterative performance calibration in which knobs are tuned until a device under test (DUT) is “healed” or the number of iterations is exceeded such that the benefit of healing is surpassed by the corresponding cost.

the complex relationship between knobs and performances that engenders this perception of cost: this relationship is not yet well-characterized in the literature, and many of the performance calibration methods thus far proposed rely on iterative test-and-tune cycles to explore the large space of knob settings, which is cost prohibitive.

Consider, for example, the generic iterative performance calibration method shown in Figure 1. The performances of a fabricated device under test (DUT) are first measured using expensive analog and RF ATE. In an aggressively optimized design, process variation results in many devices falling outside the specification limits, as illustrated by the post-ATE scatter plot of performances. We can avoid having to simply discard these devices because of the presence of knobs in the circuit: the knobs enable a performance calibration loop wherein the knobs are tuned to a new setting and the process is repeated until either the device is “healed” or a threshold (i.e., number of iterations) is exceeded, beyond which the benefit from healing the device is surpassed by the corresponding cost. If this is implemented properly, the expectation is that such tuning will help moderate the impact of process variation and will result in tighter performance distributions and, by extension, a much larger percentage of devices that fall within the design specifications.

Tuning would be straightforward if cost were not a consideration: for every device, we could exhaustively iterate through test-and-tune cycles until a knob setting is found that enables the device to meet specification limits. However, two key challenges can be quickly recognized that could jeopardize the viability and cost-effectiveness of this iterative performance calibration framework. First, the standard industry practice for analog and RF devices, specification testing, is already very expensive, often accounting for more than 30% of the total cost of a device. Hence, multiple iterations in which specification testing is performed each time will quickly result in an economically unviable solution. Second, a tunable design could include many knobs, each with multiple positions and capable of impacting multiple design performances in complex ways, as alluded to by the “unknown relationships” cloud of Figure 1. Accordingly, blindly searching the space of knob settings will most likely result in a losing proposition. Addressing these two challenges lies, therefore, at the core of developing a cost-effective performance calibration method.

Midpoint alternate-test-based performance calibration

We propose a novel performance calibration method called *midpoint alternate-test-based performance calibration*, as Figure 2 shows. This method

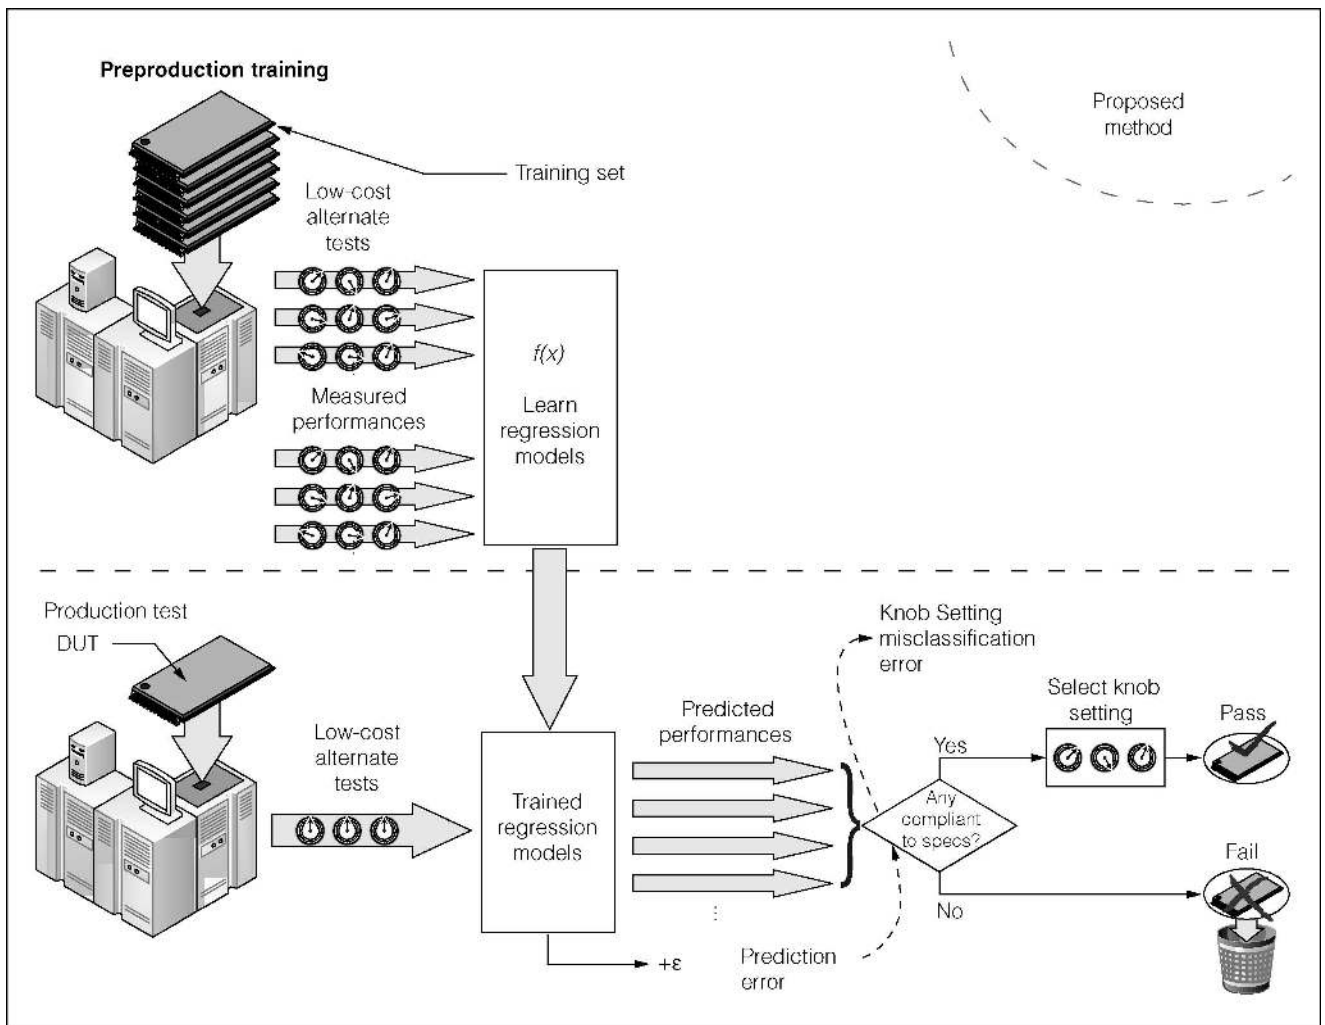


Figure 2. The proposed method of midpoint alternate-test-based performance calibration. Alternate test requires a preproduction training stage in which a small training set of devices is set aside, and on which both alternate tests and specification performances are measured. These measurements are used to construct regression models. In production, only the alternate tests are explicitly measured on every device, and then used in conjunction with the trained regression models to predict performances.

addresses the aforementioned challenges by eliminating both the need for expensive specification testing (by turning to low-cost alternatives instead) and the need for iterations, employing a single tuning cycle instead. By using low-cost test alternatives and a single tuning cycle, we gain the ability to statistically learn the impact of process variations and knob tuning on circuit performances.

Moderating test cost

Analog and RF device performances are often complex derivations that require expensive ATE to be obtained at operating frequencies. Thus,

repeatedly performing test-and-tune cycles to evaluate knob settings would quickly become economically infeasible. To manage this cost, we replace specification testing with alternate test,³ which substitutes low-cost measurements in lieu of performance measurements. These alternate tests are carefully designed to be well correlated with the specification tests, while consuming significantly fewer test resources to collect. To leverage these correlations, alternate test requires a preproduction training stage in which a small training set of devices is set aside, on which regression models are constructed. In production, only the alternate tests must be explicitly measured on every device, and used in conjunction

with the trained regression models to predict performances.

In a nutshell, the proposed performance calibration methodology uses the results of alternate test, which is performed once for the precalibration mid-point setting of the knobs, along with the learned models, to select the most appropriate knob setting. We do this by collecting alternate tests and using the alternate test regression models to predict performance across knob settings, and then by comparing the prediction results to specification limits to assign each knob setting a pass/fail label.

As Figure 2 shows, predicting pass/fail labels in this way introduces possible errors in the form of knob-setting label misclassifications. In our experiments, we've observed that such errors are substantially lower with our performance calibration methodology than with the traditional alternate test, even without the use of sophisticated error-moderation techniques such as guard-banding.⁴ Nevertheless, such techniques can also be applied, as necessary, to further reduce misclassification. Furthermore, our method—as with all alternate test-based methods—requires some mechanism to filter out devices with catastrophic faults, whose performances do not belong to the statistical distributions from which the regression models should be built. Details of various sophisticated defect-filtering methods that have been proposed are available elsewhere.^{5,6}

Eliminating iterative search

Even with alternate test-based performance calibration, iteratively performing test-tune cycles can incur excessive ATE time and cost for production ICs. Our performance calibration method eliminates the need for iterative search, via assertions about the properties of knob variation in tunable devices. First, we have empirically found that knob variation and process variation orthogonally act on device performance. This enables us to separately model each axis of variation and build a composite model that accounts for both. We have already stated that alternate tests are designed to correlate well with device performance. Implicitly, this means that we can model process variation from the alternate tests. To model knob variation, we employ a process-variation-free, simulated “ideal” device.

Because we address knob effects in the context of ideal device performances, our method overcomes the high cost of iterative test-tune cycles by

requiring only a single alternate test at a single knob setting to predict device performances across all knob settings. As Figure 2 shows, this provides a one-step solution for evaluating knob settings during production test.

Knob-setting selection

An important consideration for our performance calibration methodology is knob-setting selection. The system of Figure 2 provides knob pass/fail labels for every knob setting. When we encounter the (frequently occurring) case in which more than one knob setting heals the device, we must perform knob-setting selection (depicted by the “Select knob setting” block in Figure 2). To do this, we require a knob-setting selection metric to differentiate the optimal setting among the group of passing knob settings. For our work, we implement two approaches to knob-setting selection.

Distance from specification planes. The most conservative approach is to order potential knob settings on the basis of maximum distance from specification planes in a normalized performance space. (We use the normalized Mahalanobis distance instead of Euclidean distance to ensure that each specification is uniformly weighted.) Optimality is contingent on the type of specification limits provided: for single-sided specification limits, the maximum distance is simply the maximum distance from the specification plane itself, whereas for double-sided limits the maximum distance is the midpoint of the limits. Using this approach reduces the probability of a mistake due to marginal prediction error at the specification limit boundaries, at the expense of tending toward larger power consumption.

Power. Given a set of predicted-to-heal knob settings for a device, power is a natural optimizer for selection. To make this available as a ranking metric, we add power to the list of predicted device performances during the model-construction stage of our midpoint alternate-test-based performance calibration. This lets us predict device power consumption for every knob setting of every device in the test set. Significantly, we found that the prediction error for power was very low, letting us use predicted power to rank knob settings.

Once we've used our trained regression models to predict power values, we employ two power

rankings: minimum power and median power. Minimizing power while meeting specification limits would appear to be the global optimum; indeed, this would be the case were we to have performed an exhaustive specification test, establishing ground truth for every knob setting and then determining if the device passed or failed for each setting. However, using statistical models introduces slight errors in the pass/fail boundary. In some cases, minimizing power pushes the performances closer to their specification limits, thereby increasing the apparent misclassification error. Median power mitigates some of this error while avoiding the high power consumption of the distance metric presented previously.

Modeling

As noted previously, we believe it's possible to capture the knob effects by studying the "ideal" device, or the simulated performances of the circuit at each knob setting, without process variation. Because this simulated device does not contain process variation, it provides us the necessary information to model how the device responds to knob variation in isolation.

Knob and process variation modeling

Analog design closely approximates a zero-sum game, and is a careful balance of various trade-offs. Adding postproduction tunable elements to a circuit simply postpones a portion of this trade-off optimization process until after device fabrication. Thus, any nontrivial knob circuit element will affect more than a single specification performance—some positively, others negatively. Ideally, we would like to design knobs which are almost completely independent,^{7,8} so that a simple linear model will effectively approximate knob effects on performance. However, the nonidealities of analog design make complete independence impossible to achieve. More importantly, this is an unnecessary constraint. Although seeking knob independence remains a laudable objective, we can better model knob effects on performance by acknowledging and accommodating for knob interdependence through the inclusion of second-order knob interaction terms along with knob main effects in our model.

Thus, we model the performance responses of the ideal device as functionally dependent on the knob settings via a model that is linear in the parameters

but includes the pairwise quadratic interaction terms of explanatory variables:

$$\hat{P} = \hat{\beta}_0 + K^T \hat{\beta}_K + \varepsilon \quad (1)$$

where $\hat{\beta}_0$ is an intercept term representing the variation-free performance of the device, and K is the vector of knob settings and all pairwise interaction terms:

$$K = \underbrace{(K_1, K_2, \dots, K_p)}_{\text{main effects}} \underbrace{(K_1 K_2, K_1 K_3, \dots, K_{p-1} K_p)}_{\text{interaction terms}}^T \quad (2)$$

Finally, $\hat{\beta}_K$ is the knob effect parameter vector estimated by our model:

$$\hat{\beta}_K = \underbrace{(\beta_1, \beta_2, \dots, \beta_m)}_{\text{main effects}} \underbrace{(\beta_{1:2}, \beta_{1:3}, \dots, \beta_{(p-1):p})}_{\text{interaction terms}}^T \quad (3)$$

Were we to apply this knob-effect model to data from a real device, the prediction error would be large, as the model does not account for process variation at all. However, we posit a surprising result: given the orthogonality of process variation and knob variation, process variation is a constant offset from the presented knob-effect model. That is, we can jointly model knob and process variation effects by adding a single term to our ideal device model, which accounts for process variation. To obtain an estimate for this term, we look to alternate tests. Each alternate test gives a direct measure of the magnitude of process variation effects. Of course, each performance measure shows high correlation with different subsets of the alternate test set. Thus, we include all of the alternate tests A collected to improve our estimate, resulting in the following complete model of a performance measure as a function of alternate tests (process variation) and the knobs:

$$\hat{P} = f(A, K) = \hat{\beta}_0 + A^T \hat{\beta}_a + K^T \hat{\beta}_k + \varepsilon \quad (4)$$

We can simplify this model by concatenating the vectors A and K as X (following convention and prepending a unity constant term), and concatenating $\hat{\beta}_0$, $\hat{\beta}_a$, and $\hat{\beta}_k$ as β to arrive at the linear regression model:

$$\hat{P} = f(X) = X^T \hat{\beta} + \varepsilon \quad (5)$$

This equation provides a complete joint model for knob and process variation effects on a single performance measure. We follow this approach to

generate individual models for each of the device performance measures.

Cost model

As we've observed, one of the most significant roadblocks to the adoption of performance calibration is cost. Unless the cost-benefit ratio of deploying tunable architectures is comparable to existing design and test methods, it will not be implemented. Here, we develop an inclusive cost model that enables comparison of our midpoint alternate-test-based performance calibration method to specification test and alternate test. We use the notation of Table 1 for the duration of this discussion.

Table 2 presents a complete list of the cost models. The reference case for cost is specification testing, which includes only the baseline design cost C_0 and the cost of performing specification test once on every device $N_T P$.

As we just discussed earlier, alternate test replaces expensive specification tests with a set of low-cost alternate tests. Thus, our cost model for alternate test substitutes the $N_T P$ term with the cost of running alternate tests on every device $N_T A$. Because the models to predict performances from alternate tests must be learned, we also require a small training set in which both alternate and specification tests are performed, $N'_T(A + P)$. Note that, typically, $N_T \gg N'_T$.

We also include a cost model for our midpoint alternate-test performance calibration methodology. This adds a knob design cost term, C_D , and maintains a test set cost of $N_T A$. A key advantage of our midpoint alternate-test approach to performance calibration is that test set cost is independent of N_K . (Our approach maintains a training set cost that is proportional to N_K which we discuss later.)

Experimental validation

To validate our proposed performance calibration method, we designed a cascode low-noise amplifier (LNA) in TSMC 0.18- μm RF CMOS technology. Here, we document our design choices and show experimental results for the proposed midpoint alternate-test-based performance calibration method. (For a brief discussion of related work in this area, see the "Prior Work in Analog and RF Performance Calibration" sidebar.)

Table 1. Cost model notation.

Variable	Definition
C_0	Baseline cost of device development and production
C_D	Design cost to add knobs and implement device as a tunable architecture
N'_T	No. of devices in the training set
N_T	No. of devices in the test set
N_K	No. of knob settings
T_P	Relative cost for measuring all types of performances
T_A	Relative cost for measuring all alternate tests

Table 2. Cost models.

Configuration	Cost model
Specification testing	$C = C_0 + N_T T_P$
Alternate test	$C = C_0 + N_T T_A + N'_T (T_A + T_P)$
Midpoint alternate-test-based performance calibration	$C = \underbrace{C_0 + C_D}_{\text{baseline term}} + \underbrace{N_T T_A}_{\text{test set term}} + \underbrace{N'_T N_K (T_A + T_P)}_{\text{training set term}}$

Performance-calibration-enabled low-noise amplifier

The device we selected for experimental validation was an RF LNA, simulated using Cadence Design Systems' Spectre. We selected the LNA because it is one of the most common components in commercial transceiver RFICs. To perform postproduction performance calibration, we used three key bias voltages as our circuit knobs because these provided maximal control over performances.

Naturally, adding voltage knobs (or any knobs for that matter) to a design incurs additional cost overhead, which is accounted for by the term C_D of our cost model, and should be a consideration when selecting the type of knob to implement. Given the expanding adoption of SoC devices, which frequently integrate many DC-DC converters, we expect that integrating the three voltage regulators necessary for the knobs in this LNA will be feasible in SoCs. However, users should carefully weigh the cost-benefit trade-offs of different knob implementations; various postproduction tunable components have already been proposed in the literature.

Prior Work in Analog and RF Performance Calibration

Several researchers have attacked the problem of performance calibration in analog and RF devices.¹⁻⁶ This research falls into two categories: optimization and prediction. Several investigations involved optimization,³⁻⁶ which uses gradient descent-based methods for knob-setting selection by iteratively performing test-tune cycles to heal devices. This approach assumes that knob effects cannot be characterized in closed form, requiring use of iterative optimization methods. As we demonstrate in the main text of this article, we can make much stronger assertions about how knobs interact with device performances. Moreover, using an iterative approach is too expensive, requiring multiple test-and-tune cycles, whereas our proposed midpoint alternate-test-based performance calibration methodology requires only a single test-tune step.

Other researchers have employed prediction,^{1,2} which eliminates iteration by recognizing that first-order linear models can approximately characterize knob effects, and builds a series of such models to perturb baseline alternate test Multivariate Adaptive Regression Splines (MARS) model predictions. The effective cost of such methods is equivalent to our proposed method. However, these models are built on the assumption that designers can effectively build knobs that are approximately independent, to enable linear modeling. Because complete independence is not achievable, we avoid the error introduced by this oversimplification and include knob interaction effects in our model.

Moreover, rather than implementing a two-model approach (MARS and linear regression), we handle knob and process variation jointly in a single model.

References

1. A. Goyal, M. Swaminathan, and A. Chatterjee, "A Novel Self-Healing Methodology for RF Amplifier Circuits Based on Oscillation Principles," *Proc. Design, Automation and Test in Europe Conf. (DATE 09)*, European Design Automation Assoc., 2009, pp. 1656-1661.
2. A. Goyal, M. Swaminathan, and A. Chatterjee, "Self-Calibrating Embedded RF Down-Conversion Mixers," *Proc. Asian Test Symp. (ATS 09)*, IEEE CS Press, 2009, pp. 249-254.
3. V. Natarajan et al., "ACT: Adaptive Calibration Test for Performance Enhancement and Increased Testability of Wireless RF Front-Ends," *Proc. VLSI Test Symp. (VTS 08)*, IEEE CS Press, 2008, pp. 215-220.
4. S. Devarakond et al., "BIST-Assisted Power Aware Self Healing RF Circuits," *Proc. IEEE 15th Int'l Mixed-Signals, Sensors, and Systems Test Workshop*, IEEE CS Press, 2009, doi:10.1109/IMS3TW.2009.5158691.
5. V. Natarajan et al., "BIST Driven Power Conscious Post-Manufacture Tuning of Wireless Transceiver Systems Using Hardware-Iterated Gradient Search," *Proc. Asian Test Symp. (ATS 09)*, IEEE CS Press, 2009, pp. 243-248.
6. D. Han, B.S. Kim, and A. Chatterjee, "DSP-Driven Self-Tuning of RF Circuits for Process-Induced Performance Variability," *IEEE Trans. VLSI Systems*, vol. 18, no. 2, 2010, pp. 305-314.

Figure 3 shows the LNA schematic and specifications. Along with the LNA, this figure also shows the schematic and the specifications of an on-chip signal generator and an on-chip amplitude sensor that we designed and implemented for collecting alternate test data.

The signal generator and peak detector combination has repeatedly been demonstrated in the alternate test literature to be a highly successful means of capturing process variation impact in analog and RF devices when placed on the die with the device being tested. Accordingly, we found the signal generator and peak detector combination a natural choice for the alternate test implementation in our LNA, because the sole objective of alternate test as applied in the context of performance calibration is to quantify process variation effects on devices.

We placed amplitude sensors at both the input and output of the LNA, and we collected two measurements on each amplitude sensor, corresponding to two input frequencies of the signal generator, for a total of four alternate tests. With an appropriate choice of stimuli from the signal generator, the alternate test measurements produced by the amplitude sensor have been demonstrated to be well correlated with LNA performances.

Data set

For our experiments, we created 1,000 instances of the LNA with process variation effects included to simulate a production environment. The three knobs in the LNA designed for our experiment were assigned three discrete settings (1.6 V, 1.8 V, and 2.0 V) for a total of $3^3 = 27$ possible knob positions.

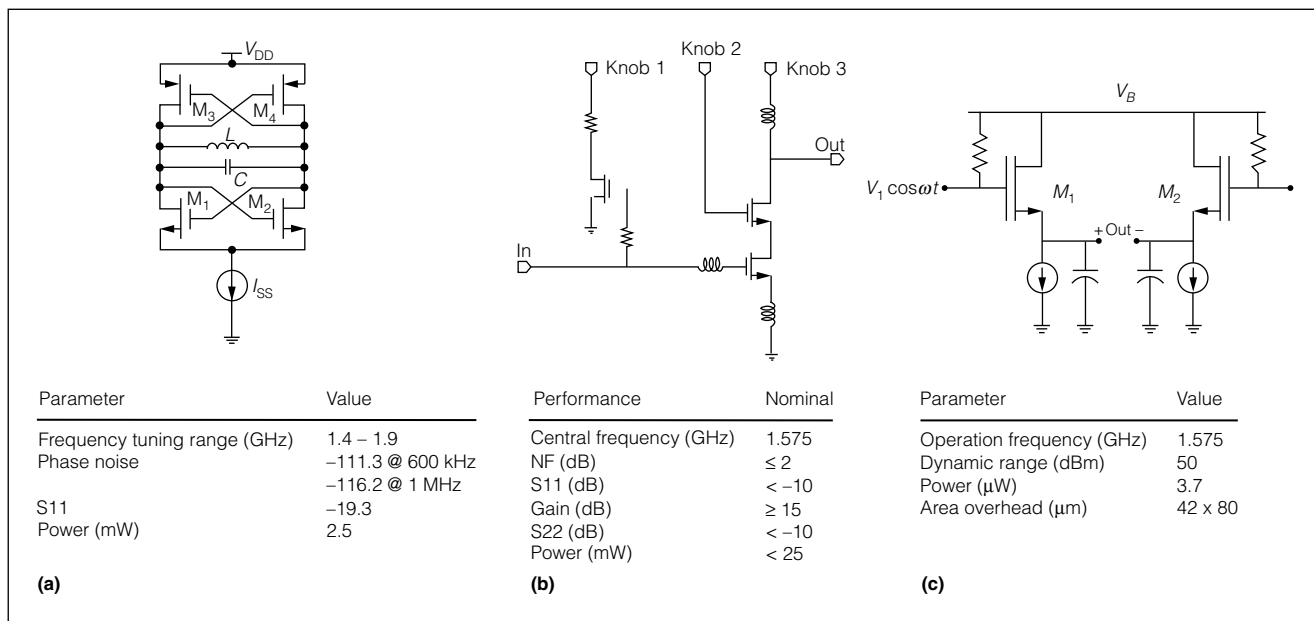


Figure 3. Signal generator (a), LNA circuit (b), and amplitude sensor (c), used for alternate test.

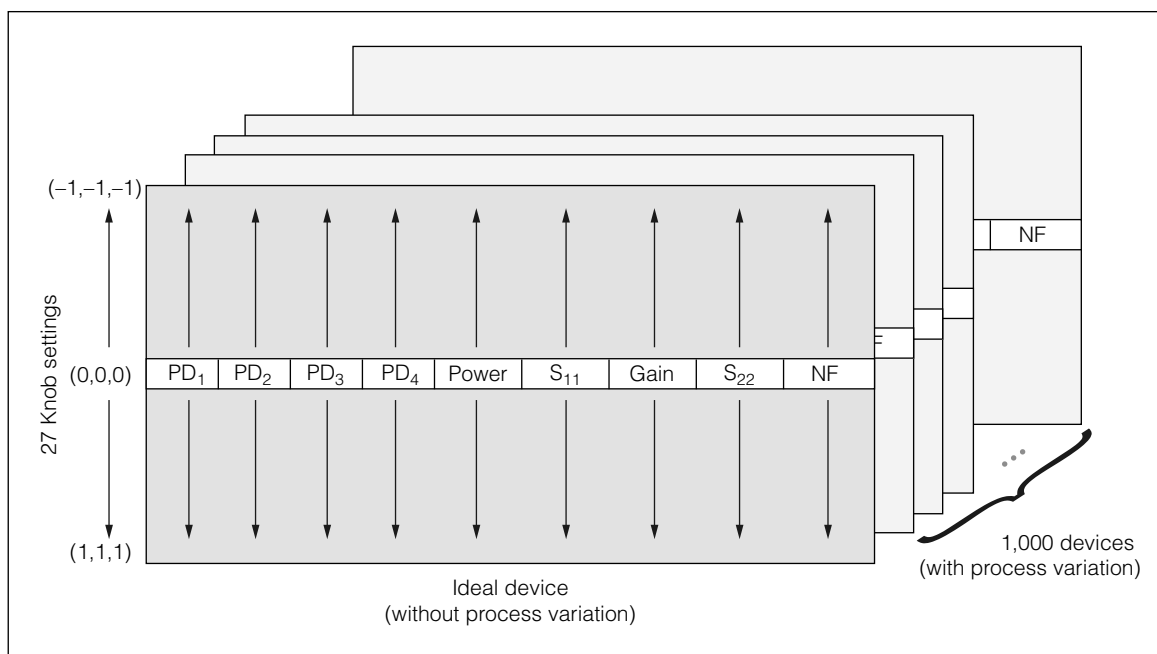


Figure 4. Graphical depiction of data set.

On every device in our data set, we collected a total of nine measurements: four performances (S11, noise figure [NF], gain, and S22), a power measurement, and the four low-cost alternate test measurements.

Thus, the entire data set is a $1,000 \times 27 \times 9$ matrix, as Figure 4 shows. Essentially, all of the performance calibration methods proposed to date can be

reduced to methodologies for systematically slicing away pieces of this 3D matrix.

If we are to model the circuit response to knob and process variation, an initial training set must be generated that includes the relationships we wish to model. For example, if we want to predict circuit performances at every knob setting, these performances must be explicitly assessed for a small training set to




Figure 5. Specification test. The test results in 85.1% of devices passing and 14.9% of devices failing.

Table 3. Alternate test results.

		Actual	
		Fail (%)	Pass (%)
Predicted	Fail	10.86	1.52
	Pass	3.54	84.08

construct our models. The training set need only be large enough to adequately capture normal variation and noise effects; as we will explain later, the number of training points required is actually relatively small. (A more general discussion of learning from circuit instances in the presence of noise is available elsewhere.⁹) Once these models are constructed, we can use them to predict circuit performances for the remaining circuits. For the experiments that required training statistical models, we split the data set 50/50, training on data from 500 devices and predicting on the remaining 500. We also performed 10 cross-validations to ensure statistical stability of the reported results.

Specification test

As Figure 5 shows, we used the center knob position to emulate a knob-free device, and we compared the performances at the center knob position to specification limits in order to obtain a pass/fail value for every device in the data set. Of the 1,000 devices,

851 passed specification testing and 149 devices failed, translating to 85.1% yield.

Alternate test

We also performed simple alternate test (without guard-banding or any other derivative performance improvement method) by only considering data from the midpoint knob setting, emulating a knob-free device. We constructed prediction models correlating each of the four device performances with peak detector measurements. The confusion matrix in Table 3 shows the results of this experiment.

Thus, employing standard alternate test results in a 3.54% test escape rate and a 1.52% yield loss rate. This is consistent with state-of-the-art alternate-test literature, excluding sophisticated error compensation techniques such as guard-banding.

Performance calibration

Exhaustive specification testing provides a useful reference point for the absolute ceiling on yield improvement possible by using performance calibration techniques. As Figure 6 shows, we exhaustively assess all circuit performances to determine a ground truth pass/fail label—that is, the pass/fail status—for every knob setting for every device.

Rather than simply looking at pass/fail labels for devices, using performance calibration let us extend the simple paradigm of pass/fail and label devices as healable or unhealable, whereby a healable device is defined as one with at least one knob setting that produces passing performances. For our data, 973 of the devices were healable, and 27 were unhealable. Recall that when the tuning was not used, 851 of the devices met specification limits and passed. Therefore, the maximum possible benefit from performance calibration methods was 122 devices, or a 12.2% yield improvement. Also, in all, approximately two-thirds (18,092) of the $1,000 \times 27 = 27,000$ total number of knob settings produced passing performances, which indicates that random knob-setting selection would introduce an unacceptably high error.

A second reference case we performed was an exhaustive alternate test by collecting alternate tests at all knob settings. Because this was a performance calibration method, we again labeled devices as healable or unhealable. The confusion matrix in Table 4 shows the error for unhealable and healable classification using exhaustive alternate test. Thus, alternate test introduced an approximately 1.04% test escape rate and a 0.34% yield loss rate, for a total error rate slightly greater than 1%.

Next, we demonstrate the performance of our proposed midpoint alternate-test-based performance calibration. Using our methodology, we classified devices as healable or unhealable, with a success rate as the confusion matrix of Table 5 shows.

Thus, with the use of alternate test, an approximately 0.62% test escape rate and a 0.48% yield loss rate were introduced, resulting in a total error rate slightly greater than 1%.

Knob-setting selection

As we explained earlier, once performances have been predicted using midpoint alternate test, knob-setting selection is performed via the specification plane distance or the predicted power knob-setting selection metric. Figure 7 presents the trade-off between power and the percentage of correct healings for the knob-setting selection optimality metrics: minimum power, median power, and maximum specified plane distance. As can be observed, the distance metric achieved a near-perfect 99.2% correct-healing rate, at the expense of high power consumption, whereas minimizing power substantially improved power consumption, as expected, but at the expense of increased error.

Training-set cost reduction

As we've already noted, the proposed midpoint alternate-test-based performance calibration method incurs an initial training-set cost $N'_T N_K (A + P)$ proportional to the number of knob settings N_K .




Figure 6. Exhaustive specification testing results in 18,092 knob settings passing, and 8,908 devices failing.

Table 4. Exhaustive alternate test results.

		Actual	
		Unhealable (%)	Healable (%)
Predicted	Unhealable	1.56	0.34
	Healable	1.04	97.06

Table 5. Midpoint alternate-test results.

		Actual	
		Unhealable (%)	Healable (%)
Predicted	Unhealable	1.98	0.48
	Healable	0.62	96.92

We found that this cost is far too pessimistic, and for real devices, the number of training instances required to adequately learn the statistics of knob and process variation is actually far smaller.

To demonstrate this finding, we used uniform sampling to reduce the size of the training set from the initial 13,500 observations (500 devices \times 27 knob settings) to 25, 50, 100, 250, 500, 1,000, and 10,000 observations. Figure 8 shows the percentage of correct healings versus the number of training set observations for the knob-setting selection methods. Error bars are displayed for the 10 cross-validations.

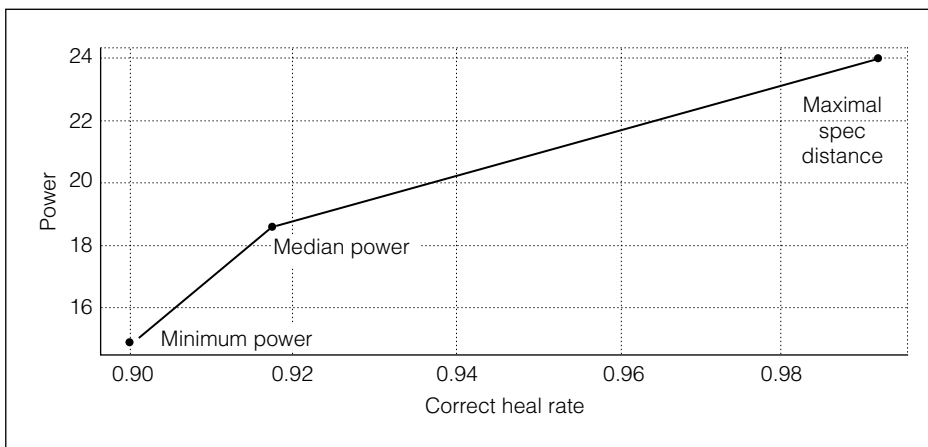


Figure 7. Trade-off between power and prediction quality.




Figure 8. Percentage of correct healings vs. training-set size.

The horizontal dashed lines present the baseline values obtained by building models from the complete training set. Once we have accounted for process variation, knob effects are relatively simple to model. Therefore, a large training set is not required to adequately capture knob variation in the performance space.

Observe that training on just 500 observations (3.7% of the original 13,500 observations) provides prediction quality on par with models constructed from the full training set. Thus, for our midpoint alternate-test-based performance calibration method, we can decouple the training cost from the number of knob settings. Therefore, our midpoint alternate-test method results in a total cost (preproduction training cost and production test cost) and total error on par with traditional alternate test, while gaining the benefits of postproduction performance calibration.

WE'VE DEMONSTRATED THAT appropriate modeling of knob variation and process variation enables highly successful performance calibration. The proposed midpoint alternate test is a cost-effective way to introduce performance calibration methodologies into an analog and RF device test flow. Indeed, it overcomes the limitations of both iterative approaches and two-model approaches by implementing a single model requiring a single alternate-test measurement step to perform tuning. Our next steps will be to investigate application of the proposed methodology to a fabricated device, and ultimately validate our tuning methodology on an industrial performance calibration-enabled device. ■

■ **References**

1. L.R. Carley, "Trimming Analog Circuits Using Floating-Gate Analog MOS Memory," *IEEE J. Solid-State Circuits*, vol. 24, no. 6, 1989, pp. 1569-1575.

2. T. Das and P.R. Mukund, "Self-Calibration of Gain and Output Match in LNAs," *Proc. IEEE Int'l Symp. Circuits and Systems (ISCAS 06)*, IEEE Press, 2006, pp. 4983-4986.
3. P.N. Variyam, S. Cherubal, and A. Chatterjee, "Prediction of Analog Performance Parameters Using Fast Transient Testing," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 21, no. 3, 2002, pp. 349-361.
4. H.-G.D. Stratigopoulos and Y. Makris, "Error Moderation in Low-Cost Machine Learning-Based Analog/RF Testing," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, no. 2, 2008, pp. 339-351.
5. H.-G. Stratigopoulos et al., "Defect Filter for Alternate RF Test," *Proc. 14th European Test Symp. (ETS 09)*, IEEE CS Press, 2009, pp. 101-106.
6. N. Kupp et al., "Confidence Estimation in Non-RF to RF Correlation-Based Specification Test Compaction," *Proc. 13th European Test Symp. (ETS 08)*, IEEE CS Press, 2008, pp. 35-40.
7. A. Goyal, M. Swaminathan, and A. Chatterjee, "A Novel Self-Healing Methodology for RF Amplifier Circuits Based on Oscillation Principles," *Proc. Design, Automation and Test in Europe (DATE 09)*, IEEE CS Press, 2009, pp. 1656-1661.
8. A. Goyal, M. Swaminathan, and A. Chatterjee, "Self-Calibrating Embedded RF Down-Conversion Mixers," *Proc. Asian Test Symp. (ATS 09)*, IEEE CS Press, 2009, pp. 249-254.
9. H.-G.D. Stratigopoulos and Y. Makris, "Non-linear Decision Boundaries for Testing Analog Circuits," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 11, 2005, pp. 1760-1773.

Nathan Kupp is pursuing a PhD in electrical engineering at Yale University. His research interests include applying machine learning and statistical learning theory to problems in analog and RF test. He has an MS in electrical engineering from Yale University. He is a student member of IEEE.

He Huang is an RF design engineer at Broadcom. His research interests include RFIC design, BIST, and self-healing RFICs. He has an MS in electrical engineering from Yale University.

Yiorgos Makris is an associate professor of electrical engineering at Yale University. His research interests include test and reliability of analog, digital, and asynchronous circuits and systems. He has a PhD in computer science and engineering from the University of California, San Diego. He is a senior member of IEEE.

Petros Drineas is an associate professor of computer science at Rensselaer Polytechnic Institute. His research interests include design and analysis of randomized and approximation algorithms for linear algebraic computations. He has a PhD in computer science from Yale University. He is a senior member of the ACM.

■ Direct questions and comments about this article to Yiorgos Makris, Yale University, Dunham Laboratory, 10 Hillhouse Ave., New Haven, CT 06520; yiorgos.makris@yale.edu.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.