

Bridging the Accuracy of Functional and Machine-Learning-Based Mixed-Signal Testing

Haralampos-G. D. Stratigopoulos and Yiorgos Makris

Electrical Engineering Dept., Yale University
New Haven, CT 06520-8285, USA

E-mail: {haralampos-g.stratigopoulos, yiorgos.makris}@yale.edu

Abstract—Numerous machine-learning-based test methodologies have been proposed in recent years as a fast alternative to the standard functional testing of mixed-signal/RF integrated circuits. While the test error probability of these methods is rather low, it is still considered prohibitive for accurate production testing. In this paper, we demonstrate how to minimize this test error probability and, thus, how to bridge the accuracy of functional and machine-learning-based test methods. The underlying idea is to measure the confidence of the machine-learning-based test decision and retest the small fraction of circuits for which this confidence is low via standard functional test. Through this approach, the majority of circuits are tested using fast machine-learning-based tests, which, nevertheless, are equivalent to the standard functional ones with regards to test error probability. By varying the acceptable confidence level, the proposed method enables exploration of the trade-off between test time and test accuracy.

I. INTRODUCTION

While functional testing of mixed-signal/RF circuits is highly accurate, it involves elaborate measurement procedures that incur an often-prohibitive cost. As an alternative, machine learning inspired a new test paradigm, wherein the results of functional testing are inferred from a small number of measurements [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17]¹. This inference is based on learning the complex mapping between the few measurements and the functional test results, which is achieved by mining the knowledge contained in a characterized set of circuit instances (training set). In particular, test methods based on machine learning explore two directions. In the first direction [1], [2], [4], [5], [10], [16], [17], the nominal and faulty training circuit instances are projected on the space of measurements and a hypersurface is allocated to separate them. A new circuit is, subsequently, tested by simply examining the location of its measurement pattern with respect to this hypersurface. In the second direction [3], [6], [7], [8], [9], [11], [12], [13], [14], [15], multivariate adaptive regression splines (MARS) [24] model the functions that map a measurement pattern to the performance parameters of the circuit. A new circuit is, subsequently, tested by processing its measurement pattern through these functions, obtaining its performance parameters, and comparing to the circuit specifications.

¹Machine learning was first used for fault diagnosis purposes [18], [19], [20], [21], [22], [23]; however, this field of application suffers from the lack of representative and widely accepted analog fault models.

These machine-learning-based test methods offer a fast alternative to the standard functional test, yet their accuracy is not up to par due to the inherent discontinuity between the nominal and faulty distributions. In extensive experiments with various circuits and measurement spaces, the test error probability of the ontogenic neural classifier proposed in [1], the most advanced solution in the first direction, never drops below 2%. Similar percentages are reported for the MARS-based methods². Therefore, additional care is necessary in order to make machine-learning-based test competitive. In essence, what is needed is a method that identifies the circuit instances that have a high probability of misclassification, along with a less error-prone strategy for testing them.

In this paper, we propose such a method for the hypersurface-based test solution, which we demonstrate by extending the ontogenic neural classifier proposed in [1]. The underlying principle is to allocate two hypersurfaces (guard-bands) instead of one, thus creating a trichotomy of the measurement space, where an ambivalent region is interjected in between the nominal and faulty regions. In the outer sides of the guard-bands, either the nominal or the faulty circuit population dominates and, therefore, circuits falling therein are classified to the corresponding class with high confidence. In contrast, circuits falling in the guard-banded region are suspect to misclassification. These circuits are identified and retested through the standard functional test approach, in order to provide test decisions with high confidence throughout the entire circuit population. In short, the majority of circuits are tested through the fast machine-learning-based approach, while a small fraction is retested through the lengthier functional test method. Thus, the proposed method is overall faster than functional test, yet without sacrificing the accuracy of the test decisions. Furthermore, it provides the ability to explore the trade-off between test accuracy (i.e. test error probability) and test application time (i.e. number of retested circuits) by varying the area of the guard-banded region.

Guard-banding has been mentioned in [9] as an option for the MARS-based methods as well. Therein, guard-bands are defined as a percentile deviation from the circuit specification limits, but no experimental data is reported regarding the

²Often, instead of the test error probability, the performance parameter prediction error is reported. However, this error metric is only loosely related to the actual misclassification rate.




Fig. 1. Allocation of guard-bands in a two-dimensional measurement space extracted from the response of the switched-capacitor filter in Fig. 3 to band-limited white noise.

resulting number of retested circuits. Also, within the context of specification test compaction [25], guard-bands are allocated by perturbing the entire optimal decision hypersurface by a predefined distance, thus creating a guard-banded region of constant width. This rigidity of the reallocation method might inadvertently guard-band regions with non-overlapping populations, resulting in an unnecessarily large number of retested circuits. Instead, in the proposed method, the two guard-bands are viewed as independent decision boundaries and, thus, are allocated regardless of the position of the optimal decision hypersurface.

The rest of this paper is organized as follows. In section II, we briefly review the ontogenic neural classifier proposed in [1] and extend it to allocate guard-bands in a given measurement space. In section III, we discuss the extraction of an initial discriminative measurement space by applying a band-limited white noise test stimulus. In section IV, we present a genetic search method for selecting measurement subspaces wherein effective guard-bands can be allocated. In section V, the effectiveness of the proposed methodology is illustrated on a switched-capacitor ladder filter.

II. GUARD-BAND ALLOCATION

A. The two-class separation problem

The ontogenic neural classifier described in [1] allocates a hypersurface that separates the projections of the nominal and faulty instances of a representative set of characterized circuit instances (training set) in a d -dimensional measurement space. The key novelty of this classifier is that it possesses the ability to allocate arbitrarily non-linear hypersurfaces, thus reciprocating very well even in the presence of complex distributions. Fig. 1 illustrates an example of a separation hypersurface in a two-dimensional measurement space $x_1 - x_2$. The classifier is constructed progressively (ontogenically) by placing layers of linear perceptrons [26] above existing ones. The weights of the

synapses of each added layer are adjusted using the thermal perceptron learning rule [27]. This construction mechanism is known as the *I-Pyramid* algorithm [28]. The first layer receives as input the d -dimensional measurement vector, while successive layers receive as additional inputs the outputs of the preceding layers and a parabolic term, which is the sum of the squared d measurements. The hypersurface allocated by the output perceptron (top layer) is linear to its input vector, but it is non-linear to the original d -dimensional measurement space due to the parabolic term and the shared information from the preceding perceptrons. Thus, non-linear hypersurfaces can be obtained by training a sequence of linear perceptrons. As training progresses, the complexity of the hypersurface increases and the classification error in the training set decreases. However, the generalization of the trained classifier (i.e. its ability to correctly classify previously unseen circuit instances) decreases when the separation hypersurface is over-fitted to the training data. In fact, the generalization reaches a maximum and, subsequently, decreases monotonically. Therefore, the generalization is monitored on a second independent set of circuit instances (validation set) and after training is complete, the classifier is pruned down to the layer that achieved the best generalization. An unbiased estimate of the generalization is then computed on a third independent set of circuit instances (test set). The interested reader is referred to [1] for more details and relevant references.

B. Guard-banding

As shown in Fig. 1, there exist regions in the measurement space wherein nominal and faulty circuit populations overlap. Due to this overlap, the optimal test hypersurface is bound to have a non-zero error probability. To date, no extracted measurement space with zero test error has been reported. The lowest observed error has been around 2%, which is deemed prohibitive for accurate production testing. Thus, the hypersurface-based test cannot substitute the standard functional tests without sacrificing test accuracy. In this section, we propose a method for reducing the error through the use of guard-bands.

The underlying idea is to allocate two hypersurfaces (nominal and faulty guard-bands) that surround the overlapping populations. The guard-bands partition the measurement space into three regions: two regions of predominantly nominal or faulty circuits, respectively, and a zone in between that contains a mixed distribution. Fig. 1 illustrates a possible allocation of guard-bands, with the grey area representing the guard-banded zone. A new circuit is now tested by examining the location of its measurement pattern with respect to the two guard-bands instead of the single test hypersurface. If its measurement pattern falls in the predominantly nominal or predominantly faulty regions, it is assigned to the respective class with low error probability. Otherwise, if its measurement falls within the guard-banded zone, the circuit is retested via standard functional test to ensure a credible decision.

By varying the guard-banded area, the number of retested circuits and the test error probability can be traded off.

Specifically, as the guard-bands become more distant from the original test hypersurface, the test error probability is reduced at the expense of retesting more circuits. In the two limits, the guard-bands surround the entire distribution or merge onto the optimal test hypersurface. Thus, if h_r denotes the error of the test hypersurface, then the test error, g_r , associated with the guard-bands drops from h_r to zero as the number of the retested instances increases. In practice, in a discriminative measurement space, guard-bands can be allocated such that the test error approaches zero when a small fraction of circuit instances is retested.

Each guard-band is allocated separately to classify perfectly all the training patterns of one class and, under this constrain, to provide optimum classification for the training patterns of the opposite class. Without loss of generality, consider the allocation of the faulty guard-band. First, the overlapping regions are cleared out of the nominal training patterns. This is achieved by identifying the nominal patterns that lie within hyperspheres of radius D centered at faulty patterns. The radius D is defined as:

$$D = \frac{1}{N_f} \sum_{i \in C_f} \min_{j \in C_n} \|\bar{x}^i - \bar{x}^j\| \quad (1)$$

where \bar{x}^k is the measurement pattern of instance k , C_n and C_f denote the nominal and faulty classes, respectively, $\|\cdot\|$ is the Euclidian norm, and N_f is the number of faulty patterns in the training set. Nominal patterns j are successively paired with all faulty patterns i and if the inequality:

$$\|\bar{x}^i - \bar{x}^j\| < D \quad (2)$$

holds, then they are temporarily excluded from the training set. After the overlapping regions have been resolved, the ontogenic neural classifier described in [1] is employed to allocate the faulty guard-band. The dual procedure is followed to allocate the nominal guard-band.

Evidently, given a measurement space, the aforementioned trade-off between test time and test error probability can be explored by simply varying the parameter D around its value calculated in eq. (1).

III. MEASUREMENT EXTRACTION

The ability of the classifier to allocate effective separation hypersurfaces and, by extension, guardbands, is bounded by the choice of the measurement space wherein the problem is solved. Measurements should provide adequate discrimination and, preferably, should be extracted by interfacing the circuit under test (CUT) to a single test configuration, in order to maintain low test application time. To this end, we use white noise limited up to a frequency multiple of the bandwidth of the CUT [20] as a test stimulus in our method. Intuitively, this stimulus is promising, since it contains infinite tones that can generate distinct intricate response waveforms.

Band-limited white noise can be digitally synthesized on-chip using reliable low cost resources, as shown in Fig. 2. The pseudo-random bit sequence generated by a linear




Fig. 2. Pseudo-random white noise source.

feedback shift register (LFSR) has a spectrum whose envelope is proportional to the square of $(\sin x)/x$ [29]. The spectrum is flat within $\pm 0.1dB$ up to 12% of the clock frequency f_{clk} of the LFSR, dropping rapidly beyond its $-3dB$ point of $0.44 \cdot f_{clk}$. Thus, low-pass filtering with a high frequency cutoff of 5% – 10% of f_{clk} will convert the LFSR output to a band-limited white noise voltage. Since a sharp cutoff characteristic is not required, simple RC filtering suffices. The filtered bit pattern is applied to the CUT through a driving buffer. The response of the CUT is sampled³ at d equidistant points.

According to the above discussion, f_{clk} must be chosen such that $0.1 \cdot f_{clk} \geq \nu \cdot BW$, where BW is the bandwidth of the CUT and ν is a positive integer. Let t_r be the time resolution between consecutive measurements and t_o be the settling time of the CUT. In order to avoid repeating measurements, the length m of the LFSR must be chosen such that its period, $(2^m - 1) \cdot T_{clk}$ (assuming that the LFSR generates a maximal-length pseudorandom sequence), satisfies $(2^m - 1) \cdot T_{clk} \geq t_o + t_r \cdot d$.

IV. MEASUREMENT SUBSPACE SELECTION

The extraction method of the previous section provides an initial set of d candidate measurements that may be used to define the space wherein the classifier will allocate guard-bands. In order to increase the probability of extracting useful candidate measurements, d should be adequately high. However, due to a paradoxical phenomenon termed curse of dimensionality [26], increasing the dimensionality of the measurement space does not necessarily result in a better solution to the guard-band allocation problem. Due to the finite size of the training set, as the dimensionality increases, the distributions become increasingly sparse, the overlapping regions gradually disappear and the percentage of ambivalent circuits that are resolved through functional retesting diminishes. However, this incurs an adverse effect on the test error probability. Since the classifier has to assign a target value to each point in the measurement space, it ends up assigning random target values to large subspaces

³Direct sampling of the transient response is not suitable for multi-gigahertz RF devices due to expected timing inaccuracies. Thus, a new measurement extraction paradigm is required in order to apply machine-learning-based methodologies for RF devices. The research group at Georgia Institute of Technology has investigated this problem, which continues to be a topic of on-going research. Possible solutions include modulation of a baseband test stimulus and subsequent demodulation of the circuit response [8], power spectral density features [11] and the use of embedded sensors to extract low-frequency varying or dc responses [12], [15]. A comprehensive overview of the above techniques can be found in [14].

that are empty because of the sparse training distributions. Therefore, to avoid the curse of dimensionality, guard-bands should be allocated in a low-dimensional subspace of the initial d -dimensional measurement space, wherein the ratio of training instances to the number of coordinates is sufficiently large.

In order to select appropriate low-dimensional measurement subspaces, we employ an adapted version of the genetic algorithm (GA) developed in [30]. Recent comparative studies [31] show that GAs [32], [33] are the most suitable for large scale measurement selection problems. GAs start with a base population of chromosomes and generate successive populations through an intrinsically parallel search process that mimics the mechanics of natural selection and genetics [30]. In the search for subsets of measurements, a subset is encoded in a chromosome as a d -element bit-string, with the k -th bit denoting the presence or the absence of the k -th measurement. After each generation, a mating pool is selected to include copies of the fittest current chromosomes. Successive populations are reproduced by mating the chromosomes in the pool at random and allowing them to crossover and mutate. Crossover proceeds with high probability and is performed by crossing and separating the chromosome pair at a chosen crossover site. Mutation is the occasional alteration of a bit position in the chromosome string in order to reinstate potentially valuable measurements that have been discarded during the search. The GA evolves by the juxtaposition of schemata (bit templates), which results in rapid optimization of the target fitness function.

As pointed out in section II-B, our method enables exploration of the trade-off between test time and test accuracy by varying the radius D during guard-band allocation in a given measurement space. Within this trade-off, various optimization objectives can be pursued. For example, one may place more emphasis on test accuracy and, hence, pursue minimization of the percentage of retested circuits under the constraint that the test error probability does not exceed a certain threshold. Alternatively, one may place more emphasis on test time and, hence, pursue minimization of the test error probability, under the constraint that the percentage of retested circuits does not exceed a certain threshold. Mixed objectives may also be of interest. However, the measurement subspace wherein guard-banding better meets each of these objectives may be different. Thus, in order to obtain tight solutions along the trade-off curve, the GA is run multiple times to select subspaces that maximize fitness functions which reflect various competing objectives. The following three fitness functions are considered:

$$f_1 = h_r, \quad n_r = 0$$

$$f_2 = \frac{e(\frac{(100-n_r)-th}{m})}{e(1)} \cdot (1 + g^\theta \cdot v)$$

$$f_3 = \frac{e(\frac{(100-n_r)-th}{m})}{e(1) \cdot (g_r + \alpha)}$$

where h_r is the test error of the optimal test hypersurface, g_r is the test error of the guard-bands, n_r is the percentage of retested circuits and $e(\cdot)$ is the exponential function. Fitness f_1 is simply the test error of the ontogenic neural classifier [1] in the absence of guard-bands ($n_r = 0$). Fitness f_2 aims to identify the lowest n_r under the constrain $g_r < \epsilon_r$, where ϵ_r is an accepted error threshold. The parameter v is 1 when $g_r > \epsilon_r$ and 0 otherwise. g^θ , where $\theta = 0.5$, is a penalty factor that is kept low at the beginning of the GA in order to explore the schemata with larger flexibility, and it increases as the GA evolves in order to discard the invalid solutions that satisfy $g_r > \epsilon_r$. The exponential function scales n_r in order to offer higher fitness to solutions that have low n_r . The parameters th and m , termed feasibility threshold and tolerance margin respectively, define the scaling operation. Fitness f_3 is unconstrained and attempts to jointly optimize n_r and g_r . The parameter in the denominator is a small real number set to $\alpha < 1$.

V. RESULTS

The effectiveness of the proposed methodology in exploring the trade-off between test time (percentage of retested circuits) and test accuracy (test error probability), as well as in reducing the test application of standard functional test without sacrificing accuracy, is illustrated on the fifth-order elliptic switched-capacitor ladder filter shown in Fig. 3.

In a production environment, the training, validation and test sets would comprise chips across several lots, in order to capture the statistical impact of process drifts on the performance parameters. For the purpose of this experiment, we generate $N = 2000$ instances by Monte Carlo analysis, letting various design parameters follow a normal distribution centered at their nominal values with 3% variance. The parameters considered are the switched-capacitor values and the geometry, oxide thickness, threshold voltage, body effect coefficient and junction capacitances of the transistors in the op-amps. Catastrophic shorts and opens in the MOS switches are easily detected since they generate outlier points in the faulty distribution. $N/2$ instances are assigned to the training set and $N/4$ to each of the validation and test sets. The standard functional tests performed to characterize the N instances concern the ripples in the pass- and stop-bands, gain errors, group delay, phase response and total harmonic distortion.

The pass-band of the filter is in the range $0 - 1KHz$. The band-limited white noise stimulus is generated by a maximal length LFSR with characteristic polynomial $x^{10} + x^7 + 1$, clocked at $f_{clk} = 100KHz$, thus having a period of $10.23ms$. $d = 30$ measurements are extracted with a conservative resolution $t_r = 0.3ms$ (the settling time is approximately $0.5ms$).

The percentage of retested circuits, n_r , and the test error probability, g_r , in f_2 and f_3 are computed using the radius

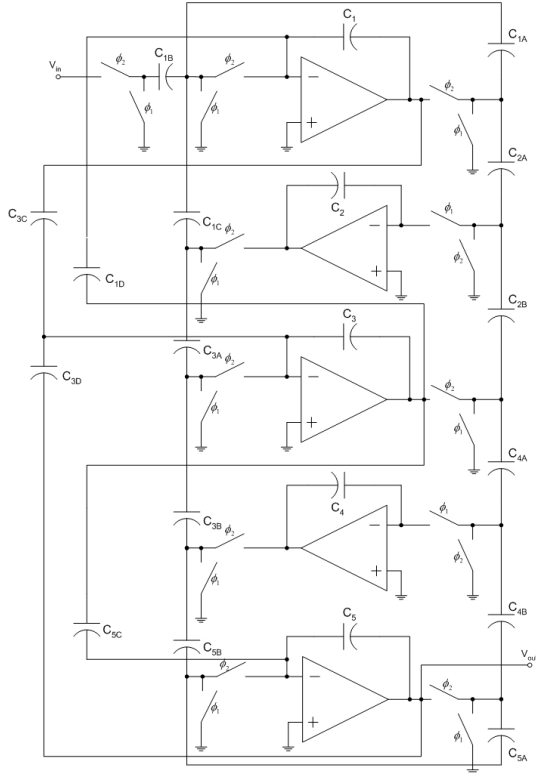


Fig. 3. LDI ladder realization of a fifth-order elliptic switched-capacitor filter [34].

D defined in eq. (1). The threshold value ϵ_r in f_2 is set to 1, i.e. only solutions with $g_r < 1$ are considered valid. The real parameter α in f_3 is set to 0.05, i.e. solutions with very low test error probability are given a very high fitness. The CPU elapsed times⁴ and the fittest identified measurement subsets for f_1 , f_2 , and f_3 are shown in Table I. Each bit in the strings corresponds to one of the d consecutive measurements performed after the response has settled.

For each fittest subspace, several guard-bands are allocated by varying the radius D ⁵. The corresponding trade-off between test time and test accuracy are illustrated in Fig. 4. Note that $g_r = 0$ is a reference point that equals the error probability of standard functional tests, that is, when $g_r = 0$, functional and machine-learning-based tests are equally accurate. The thick line runs along the dominant identified (n_r, g_r) points. These points are dominant in the sense that, for a specific n_r , they correspond to an optimal g_r and vice versa. Four interesting guard-bands can be observed along the dominant trade-off curve: the test error is 4% when $n_r = 0$,

⁴The GA and the allocation of guardbands in the fittest identified measurement subspaces are routines that are executed off-line and only once for any particular design or production run.

⁵The classifier is unbiased, i.e. $g_r > 0$ corresponds to both test escapes and yield loss. If test escape elimination is of higher importance, the faulty guard-band is allocated using a large radius D , such that it is adequately distanced from all faulty training patterns. Under this scenario, a trade-off can be obtained by varying the position of only the nominal guard-band.

TABLE I

CPU TIMES FOR THE GA AND FITTEST MEASUREMENT SUBSETS.

fitness function	CPU time	fittest measurement subset
f_1	5.78h	100111010000000001010000001000
f_2	10.28h	110010000100011110001000111000
f_3	9.97h	100000010011110110101000001100




Fig. 4. Trade-off between test time and test accuracy for the fittest measurement subspaces identified by running the GA with the three different fitness criteria defined in section IV. The thick piece-wise linear curve runs along the dominant points.

it decreases down to 2.0% when 8.1% of the circuits are retested, it drops further to 0.6% when the percentage of retested circuits increases to 15.6% and it reaches zero when this percentage is further increased to 20.7%.

With regards to test application time, let T and T' denote the time for standard functional and guard-band-based test application per CUT, respectively. The average guard-band-based test application time is $E[T'] = t + n_r \cdot T$, where t denotes the test application time for a circuit whose measurement pattern falls outside the guard-banded region. Thus, the average test time difference is $T - E[T'] = (1 - n_r) \cdot T - t$, which is positive when $T > t / (1 - n_r)$. Given that in this experiment $t < 10ms$ and that the standard functional test application time for the filter is typically well above 0.4s, i.e. $T > 0.4s$, a pessimistic estimate of the average test time using the guard-bands corresponding to the point (20.7%, 0%) is $E[T'] = 92.8ms$. This corresponds to an approximate 4X relative reduction in test time or, equivalently, a 0.3s average absolute reduction per CUT.

VI. CONCLUSIONS

The use of guard-bands in machine-learning-based testing of mixed-signal circuits enables exploration of the trade-off between test time and test accuracy. As demonstrated in this paper, efficient allocation of guard-bands in carefully selected

measurement subspaces allows the majority of circuits to be tested through faster yet equivalently accurate test criteria as standard functional tests. Additionally, it pinpoints the circuits that are suspect to misclassification through the machine-learning-based test method and should be retested via standard functional testing, in order to ensure the credibility of the test decision. Results obtained on a switch-capacitor ladder filter show that, by retesting 20.7% of the circuits, the accuracy of standard functional test is maintained while an estimated 4X relative or 0.3s absolute reduction in test time per CUT is achieved.

REFERENCES

- [1] H.-G. D. Stratigopoulos and Y. Makris, "Constructive derivation of analog specification test criteria," in *IEEE VLSI Test Symposium*, 2005, pp. 395–400.
- [2] C. Y. Pan and K. T. Cheng, "Pseudorandom testing for mixed-signal circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 16, no. 10, pp. 1173–1185, 1997.
- [3] P. N. Variyam and A. Chatterjee, "Enhancing test effectiveness for analog circuits using synthesized measurements," in *IEEE VLSI Test Symposium*, 1998, pp. 132–137.
- [4] C. Y. Pan and K. T. Cheng, "Test generation for linear time-invariant analog circuits," *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 46, no. 5, pp. 554–564, 1999.
- [5] W. M. Lindermeir, H. E. Graeb, and K. J. Antreich, "Analog testing by characteristic observation inference," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 18, no. 9, pp. 1353–1368, 1999.
- [6] P. N. Variyam and A. Chatterjee, "Specification-driven test generation for analog circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 10, pp. 1189–1201, 2000.
- [7] 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.
- [8] R. Voorakaranam, S. Cherubal, and A. Chatterjee, "A signature test framework for rapid production testing of RF circuits," in *Design, Automation and Test in Europe*, 2002, pp. 186–191.
- [9] R. Voorakaranam, R. Newby, S. Cherubal, B. Cometta, T. Kuehl, D. Majernik, and A. Chatterjee, "Production deployment of a fast transient testing methodology for analog circuits: Case study and results," in *IEEE International Test Conference*, 2003, pp. 1174–1181.
- [10] V. Stopjakova, P. Malosek, D. Micusik, M. Matej, and M. Margala, "Classification of defective analog integrated circuits using artificial neural networks," *Journal of Electronic Testing*, vol. 20, pp. 25–37, 2004.
- [11] S. S. Akbay and A. Chatterjee, "Feature extraction based built-in alternate test of RF components using a noise reference," in *IEEE VLSI Test Symposium*, 2004, pp. 273–278.
- [12] S. Bhattacharya and A. Chatterjee, "Use of embedded sensors for built-in-test of RF circuits," in *IEEE International Test Conference*, 2004, pp. 801–809.
- [13] A. Raghunathan, J. H. Chun, J. A. Abraham, and A. Chatterjee, "Quasi-oscillation based test for improved prediction of analog performance parameters," in *IEEE International Test Conference*, 2004, pp. 252–261.
- [14] S. S. Akbay, A. Halder, A. Chatterjee, and D. Keezer, "Low-cost test of embedded RF/Analog/Mixed-signal circuits in SOPs," *IEEE Transactions on Advanced Packaging*, vol. 27, no. 2, pp. 352–363, 2004.
- [15] S. S. Akbay and A. Chatterjee, "Built-in test of RF components using mapped feature extraction sensors," in *IEEE VLSI Test Symposium*, 2005, pp. 243–248.
- [16] H.-G. D. Stratigopoulos and Y. Makris, "Generating decision regions in analog measurement spaces," in *ACM Great Lakes Symposium in VLSI*, 2005, pp. 88–91.
- [17] H.-G. D. Stratigopoulos and Y. Makris, "Non-linear decision boundaries for testing analog circuits," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 24, no. 11, pp. 1760–1773, 2005.
- [18] P. Collins, S. Yu, K. R. Eckersall, B. W. Jervis, I. M. Bell, and G. E. Taylor, "Application of Kohonen and supervised forced organization maps to fault diagnosis in CMOS opamps," *Electronic Letters*, vol. 30, no. 22, pp. 1846–1847, 1994.
- [19] S. Yu, B. W. Jervis, K. R. Eckersall, I. M. Bell, A. G. Hall, and G. E. Taylor, "Neural network approach to fault diagnosis in CMOS opamps with gate oxide short faults," *Electronic Letters*, vol. 30, no. 9, pp. 695–696, 1994.
- [20] R. Spina and S. Upadhyaya, "Linear circuit fault diagnosis using neuromorphic analyzers," *IEEE Transactions on Circuits and Systems-II: Analog and Digital Signal Processing*, vol. 44, no. 3, pp. 188–196, 1997.
- [21] Y. Maidon, B. W. Jervis, N. Dutton, and S. Lesage, "Diagnosis of multifaults in analogue circuits using multilayer perceptrons," *IEE Proc.-Circuits Devices Syst.*, vol. 144, no. 3, pp. 149–154, 1997.
- [22] Z. R. Yang, M. Zwolinski, C. D. Chalk, and A. C. Williams, "Applying a robust heteroscedastic probabilistic neural network to analog fault detection and classification," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 1, pp. 142–151, 2000.
- [23] F. Aminian, M. Aminian, and Jr. H. W. Collins, "Analog fault diagnosis of actual circuits using neural networks," *IEEE Transactions on Instrumentation and Measurement*, vol. 51, no. 3, pp. 544–550, 2002.
- [24] J. H. Friedman, "Multivariate adaptive regression splines," *The Annals of Statistics*, vol. 19, no. 1, pp. 1–67, 1991.
- [25] S. Biswas, P. Li, R. D. (Shawn) Blanton, and L. Pileggi, "Specification test compaction for analog circuits and MEMS," in *Design, Automation and Test in Europe*, 2005, pp. 164–169.
- [26] C. M. Bishop, *Neural Networks for Pattern Recognition*, Oxford University Press, 1995.
- [27] M. Frean, "A "thermal" perceptron learning rule," *Neural Computation*, vol. 4, pp. 946–957, 1992.
- [28] R. Parekh, J. Yang, and V. Honavar, "Constructive neural-network learning algorithms for pattern classification," *IEEE Transactions on Neural Networks*, vol. 11, no. 2, pp. 436–451, 2000.
- [29] P. Horowitz and W. Hill, *The Art of Electronics*, Cambridge University Press, 1989.
- [30] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, 1989.
- [31] M. Kudo and J. Sklansky, "Comparison of algorithms that select features for pattern classifiers," *Pattern Recognition*, vol. 33, pp. 25–41, 2000.
- [32] W. Siedlecki and J. Sklansky, "A note on genetic algorithms for large-scale feature selection," *Pattern Recognition Letters*, vol. 10, pp. 335–347, 1989.
- [33] F. Z. Brill, D. E. Brown, and W. N. Martin, "Fast genetic selection of features for neural network classifiers," *IEEE Transactions on Neural Networks*, vol. 3, no. 2, pp. 324–328, 1992.
- [34] R. Gregorian and G. C. Temes, *Analog MOS Integrated Circuits for Signal Processing*, John Wiley and Sons, 1986.