

Post-Silicon Timing Characterization by Compressed Sensing

Farinaz Koushanfar
ECE and CS Depts
Rice University
Email: farinaz@rice.edu

Petros Boufounos
ECE Dept
Rice University
Email: petrosb@rice.edu

Davood Shamsi
ECE Dept
Rice University
Email: davood@rice.edu

Abstract—We address post-silicon characterization of the unique gate delays and their timing distributions on each manufactured IC. Our proposed approach is based upon the new theory of compressed sensing. The first step in performing timing measurements is to find the sensitizable paths by traditional testing methods. Next, we show that the timing variations are sparse in the wavelet domain. The sparsity is exploited for estimation of the gate delays using the compressed sensing theory. This estimation method requires significantly less number of timing measurements compared to the case where the dependence between the gate delays is not directly integrated within the estimation framework. We discuss a number of applications for the new post-silicon timing characterization method. Experimental results on benchmark circuits show that using compressed sensing theory can characterize the post-silicon variations with a mean accuracy of 95% in the pertinent sparse basis.

I. INTRODUCTION

Modern integrated circuits are variable and complex. Continuous CMOS scaling has made possible integration of billions of gates into a single multi-layer chip. Scaling to the physical device limitations and mask imprecisions have created nondeterminism in the chip’s characteristics. In the new regime, traditional models, design, and test methods have a limited effectiveness.

Furthermore, with miniaturization of devices beyond 65nm, the impact of intra-die variation and the spatial correlations are becoming more prominent [1]. Several key areas have been impacted. For example, the number of critical paths is increasing with variation, rendering the traditional test methodologies based on a few critical paths inexpressive.

In statistical static timing analysis (SSTA), instead of the single valued delays utilized in traditional models, the delay probability distributions and their correlations are used [2]. SSTA produces pre-silicon models and analysis. A post-silicon timing analysis of the chips was proposed in [3]. The method integrates the SSTA models with data collected from a few on-chip test points (e.g., via ring oscillators), to form the chip-specific distribution of the delays. Post-silicon gate-level leakage characterization by using noninvasive leakage measurements was recently proposed [4].

Our objective is to perform noninvasive post-silicon timing characterization of each chip. We exploit the theory of *compressed sensing* [5], [6] and the set of the sensitizable paths known from the testing phase to perform post-silicon delay modeling using very few measurements. We demonstrate how this method can be used for testing the chips and for efficient estimation of post-silicon the specific distribution of the individual chip’s timing. Compressed sensing exploits the sparsity of the delay distribution matrix [2], to reconstruct the timing information from a few analog timing measurements. Our contributions are as follows:

- We introduce the first post-silicon timing characterization method that is based on compressed sensing. Our method keeps the number of

measurements low without adding on-chip test structures or sensors. We only rely on the external nondestructive tests.

- We create a systematic method for exploiting the sparsity of the timing variation for post-silicon characterization.
- We present modifications to the original compressed sensing framework that is based upon regular grid-based sampling, so it can consider the irregularities of the placement in the spatial correlations of gate delays.
- We exploit spatial correlation to approximate the timing variation of the gates that are unobservable and uncontrollable because of their placement on unsensitizable paths. The key for compressed sensing-based gate characterization is the delay variation’s sparsity.

The remainder of the paper is as follows. Section II outlines the preliminaries. We introduce variation estimation by delay measurements in Section III. In Section IV, we use sparsity of the variation in the wavelet domain to recover variation with a small number of delay measurements. A number of applications of the proposed method are outlined in Section V. Evaluation results are presented in Section VI. We conclude in Section VII.

II. PRELIMINARIES

A. Variation Model and Delay Model

We adapt the Gaussian variation model by Liu [7] where the total variation, ψ_u^{total} , in a gate g_u is

$$\psi_u^{\text{total}} = \psi_u^{\text{inter}} + \psi_u^{\text{intra}} + F_u\beta \quad (1)$$

ψ_u^{inter} and ψ_u^{intra} represent inter-die and intra-die variation, respectively. ψ_u^{intra} is a multivariate Gaussian random vector. $F_u\beta$ models systematic variation. If (x_u, y_u) is the location of the gate g_u on the IC, then $F = [1, x_u, y_u]^T$ and β is a 3×1 constant vector.

Transition delay is usually modeled as a linear function of transistor feature size variation [8], [2], [9]. For example, consider a NAND2 gate that one of its inputs is 1 and its other input, at time $t = 0$, transits from 0 to 1. Because of propagation delay of the NAND2 gate, its output transits from 1 to 0 at time $t = d_r$. When there are variation in the transistor feature size, the rising-edge delay, denoted by d_r , varies among the NAND2 gates in the IC, modeled by [8]

$$d_r(\psi_u^{\text{total}}) = d_r^0 + \xi\psi_u^{\text{total}} \quad (2)$$

where ξ is a constant.

Note that even if the propagation delay is modeled as a quadratic (or higher order) polynomial [10], one can use a similar approach by assuming new variables for higher order parameters.

B. Sensitizable Paths

A path in an IC is defined as a sequence of logic gates from an input of the IC to one of its output pins. To determine the propagation delay of a path, one should find an appropriate input vector to the IC. If such an input vector exists, the path is called *sensitizable*; otherwise, it is called *unsensitizable*. For finding the sensitizable paths we use the path selection method introduced by Murakami et al. [11].

C. Compressed Sensing

Compressed Sensing is a recently emerging signal acquisition method that exploits sparse signal models to reduce the signal acquisition burden [12], [5]. Specifically, we assume that the signal of interest is a K -sparse vector \mathbf{x} in an N -dimensional space, i.e., that it only has K non-zero components. Using compressed sensing we can sample and reconstruct this vector by acquiring only $M = O(K \log(N/K))$ linear measurements:

$$\mathbf{p} = A\mathbf{x} + e, \quad (3)$$

where A denotes the measurement matrix of dimension $M \times N$, \mathbf{p} denotes the M -dimensional measurement vector, and e denotes the measurement noise.

Despite the dimensionality reduction and the rank deficiency of A , one can reconstruct the sparse vector of interest, \mathbf{x} from the measurement vector \mathbf{p} using the following convex optimization:

$$\min \|\mathbf{x}\|_1 + \lambda \|\mathbf{p} - A\mathbf{x}\|_2^2, \quad (4)$$

in which λ is a parameter chosen according to the noise variance and $\|\mathbf{x}\|_p = (\sum_{i=1}^N |x_i|^p)^{\frac{1}{p}}$. If the measurement matrix A satisfies certain conditions, it can be shown that the reconstruction using Equation 4 is exact [5].

The compressed sensing model is robust even when the acquired vector \mathbf{x} is approximately sparse, often referred to as compressible. A vector is compressible if it has very few (say K) coefficients with large magnitude and the remaining coefficients are approximately 0. Compressible vectors can be approximated very well using the best K -term approximation, i.e., using the K most significant coefficients and setting the remaining coefficients to 0.

In most practical applications, such as ours, a vector is not compressible in the canonical domain. In practice, a sparsity inducing basis W is typically necessary to expose the sparsity. The theory accommodates this case using the basis expansion

$$\mathbf{s} = W\mathbf{x}, \quad (5)$$

in which case W is the sparsity inducing transform, and the basis expansion vector \mathbf{s} is sparse instead of the vector of interest \mathbf{x} . In this case Equation 3 becomes

$$\mathbf{p} = AW^{-1}\mathbf{s} + e. \quad (6)$$

This is the same formulation as Equations 3 and 4, with only a change of variables. We now aim to recover a sparse representation \mathbf{s} from the measurements \mathbf{y} , which are acquired with a measurement matrix AW^{-1} . The signal is recovered from the transformation using Equation 5.

III. DELAY ESTIMATION BY ℓ_2 -NORM MINIMIZATION

In this section, we propose a method for post-silicon gate delay estimation by measuring the input/output path delays. First, we measure the signal propagation delays of a number of sensitizable paths. Then, based on the measured delays, we construct linear equations with the scaling factors of gate delays (defined in Section




Fig. 1. A sensitizable path from input to the output.

II-A) as the unknown parameters. Finally, we estimate the gate characteristics by solving those equations for the scaling factors. In Section IV, we use the variation in spatial correlations to improve the scaling factor estimation error.

The total path delay is an additive composition of the delays of its elements. For example, in Figure 1, the delay of path P_1 (bold path) can be written as the summation of the delays in wire w_1 , gate g_1 , wires w_5 and w_6 , gate g_3 , wire w_8 , and so on, more formally:

$$\begin{aligned} d_r(P_1) &= d(w_1) + d_r(g_1) + d(w_5) + d(w_6) + d_f(g_3) \\ &+ d(w_8) + d_r(g_4) + d(w_{10}) + d_f(g_6) + d(w_{12}) \\ &+ d_r(g_7) + d(w_{13}), \end{aligned} \quad (7)$$

where $d(w_i)$ is the delay of the wire w_i ; $d_r(g_i)$ and $d_f(g_i)$ are the rising and falling delays of the gate g_i , respectively.

For clarity of exposition, in this paper we assume interconnect delays (wire delays) are zero. The proposed method can be easily extended to accommodate non-zero interconnect delays. Note that variation in the interconnects may have a different statistical representation compared with the gates. In this case we may consider compressed sensing methods to address the sum of two distinct distributions in one framework [12]. Assuming zero interconnect delays, Equation 7 reduces to:

$$d_r(P_1) = d_r(g_1) + d_f(g_3) + d_r(g_4) + d_f(g_6) + d_r(g_7). \quad (8)$$

As discussed in Section II, because of process variation, the gate delays deviate from their nominal values [8], i.e.,

$$d_r(g_i) = d_r^{\text{nominal}}(g_i) + \xi_{r,g_i} l_{g_i}, \quad (9)$$

where $d_r^{\text{nominal}}(g_i)$ is the nominal delay for rising transition and l_{g_i} is the scaling factor of the variation for the gate g_i ; and ξ_{r,g_i} is a constant coefficient. Similarly for the falling transition,

$$d_f(g_i) = d_f^{\text{nominal}}(g_i) + \xi_{f,g_i} l_{g_i}. \quad (10)$$

Therefore, Equation 8 becomes

$$\begin{aligned} d_r(P_1) &= d_r^{\text{nominal}}(g_1) + \xi_{r,g_1} l_{g_1} \\ &+ d_f^{\text{nominal}}(g_3) + \xi_{f,g_3} l_{g_3} \\ &+ d_r^{\text{nominal}}(g_4) + \xi_{r,g_4} l_{g_4} \\ &+ d_f^{\text{nominal}}(g_6) + \xi_{f,g_6} l_{g_6} \\ &+ d_f^{\text{nominal}}(g_7) + \xi_{r,g_7} l_{g_7}, \end{aligned} \quad (11)$$

or

$$\xi_{r,g_1} l_{g_1} + \xi_{f,g_3} l_{g_3} + \xi_{r,g_4} l_{g_4} + \xi_{f,g_6} l_{g_6} + \xi_{r,g_7} l_{g_7} = b_{P_1}$$

$$\begin{aligned} b_{P_1} &= d_r(P_1) - d_r^{\text{nominal}}(g_1) - d_f^{\text{nominal}}(g_3) \\ &- d_r^{\text{nominal}}(g_4) - d_f^{\text{nominal}}(g_6) - d_f^{\text{nominal}}(g_7), \end{aligned}$$

where b_{P_1} is a constant. Thus, each sensitizable path in the circuit leads to a linear relation among the variation elements, l_{g_i} . The falling

and rising coefficients (ξ_{f,g_i} and ξ_{r,g_i}) are known and our goal is to estimate the variation l_{g_i} .

Assume that $P_1, P_2 \dots P_M$ are M sensitizable paths in a general combinational circuit C with N gates. For each path P_j , if it is stimulated by a rising transition,

$$\sum_{i=1}^N \alpha_{P_j}(i) \xi_{\lambda^r(P_j, g_i), g_i} l_{g_i} = b_j^r \quad (12)$$

where

$$\alpha_{P_j}(i) = \begin{cases} 1 & \text{if } g_i \text{ belongs to the path } P_j; \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\lambda^r(P_j, i) = \begin{cases} f & \text{if } g_i \text{ has a falling transition when path } P_j \\ & \text{is stimulated by a rising transition;} \\ r & \text{otherwise.} \end{cases}$$

Similarly for a falling transition,

$$\sum_{i=1}^N \alpha_{P_j}(i) \xi_{\lambda^f(P_j, g_i), g_i} l_{g_i} = b_j^f \quad (13)$$

where

$$\lambda^f(P_j, i) = \begin{cases} f & \text{if } g_i \text{ has a falling transition when path } P_j \\ & \text{is stimulated by a falling transition;} \\ r & \text{otherwise.} \end{cases}$$

To write Equations 12 and 13 compactly, we define the matrix A , the measurement vector \mathbf{b} and the variation vector \mathbf{l} as follows.

$$A = \begin{pmatrix} \alpha_{P_1}(1) \xi_{\lambda^r(P_1, g_1), g_1} & \dots & \alpha_{P_1}(N) \xi_{\lambda^r(P_1, g_N), g_N} \\ \alpha_{P_2}(1) \xi_{\lambda^r(P_2, g_1), g_1} & \dots & \alpha_{P_2}(N) \xi_{\lambda^r(P_2, g_N), g_N} \\ \vdots & & \vdots \\ \alpha_{P_M}(1) \xi_{\lambda^r(P_M, g_1), g_1} & \dots & \alpha_{P_M}(N) \xi_{\lambda^r(P_M, g_N), g_N} \\ \alpha_{P_1}(1) \xi_{\lambda^f(P_1, g_1), g_1} & \dots & \alpha_{P_1}(N) \xi_{\lambda^f(P_1, g_N), g_N} \\ \alpha_{P_2}(1) \xi_{\lambda^f(P_2, g_1), g_1} & \dots & \alpha_{P_2}(N) \xi_{\lambda^f(P_2, g_N), g_N} \\ \vdots & & \vdots \\ \alpha_{P_M}(1) \xi_{\lambda^f(P_M, g_1), g_1} & \dots & \alpha_{P_M}(N) \xi_{\lambda^f(P_M, g_N), g_N} \end{pmatrix},$$

$$\mathbf{b} = (b_1^r, b_2^r, \dots, b_M^r, b_1^f, b_2^f, \dots, b_M^f)^T,$$

and

$$\mathbf{l} = (l_1, l_2 \dots l_N)^T.$$

Finally, we estimate the variation in \mathbf{l} by solving the following least squares problem

$$\min \|\mathbf{A}\mathbf{l} - \mathbf{b}\|_2^2. \quad (14)$$

We call this method ℓ_2 minimization method.

IV. DELAY ESTIMATION USING COMPRESSED SENSING

This section incorporates sparsity in the wavelet domain as a model for the spatial correlation of the timing variation. Thus, one can use compressed sensing theory to measure and estimate the variation.




Fig. 2. Left: Spatial variation in a typical IC. Right: wavelet transform of the variation.

A. Sparse Representation of Variation

To capture the spatial correlation in the variation we use wavelet basis expansions. Wavelet basis expansions have two significant advantages that make them suitable for the problem at hand [13]. First, they can be computed efficiently using well-studied fast algorithms. Second, they are known to be good in sparsely describing smooth functions, such as in images because of the spatial correlations.

Figure 2 demonstrates the effectiveness of the wavelet transform in representing spatial variation. The figure on the left is the 2D plot of the variation in a typical IC, generated using the Gaussian model in [7]. The spatial correlation is evident in the figure. The figure on the right side represents the wavelet transform for the left hand side. Most of the transform coefficients are zero. Only the top-left part of the figure has a dense amount of significant non-zero elements.

B. Gates on Regular Grids

The derivations in this section assume that all the gates are located on a regular grid. Section IV-C considers the general case of nonuniform grids and relaxes this assumption.

For the gates that are located on a regular grid, the two-dimensional wavelet transform of the variation denoted by \mathbf{s} , can be expressed as the product of the variation vector, \mathbf{l} , with the wavelet transform matrix W :

$$\mathbf{s} = W\mathbf{l}, \quad (15)$$

where \mathbf{s} is assumed sparse because of the spatial correlation in the variation. We enforce the sparsity prior by regularizing Equation 14 using the ℓ_1 norm of \mathbf{s} , as described in Section II-C:

$$\min \|\mathbf{A}\mathbf{l} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{s}\|_1 \quad (16)$$

or, equivalently,

$$\min \|\mathbf{A}W^{-1}\mathbf{s} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{s}\|_1, \quad (17)$$

where λ is the regularization coefficient. The sparsity of the variations' wavelet transformation \mathbf{s} introduces a modeling prior that improves the reconstruction and resolves ambiguity. This prior is implemented using the regularization term $\lambda \|\mathbf{s}\|_1$, in Equations 16 and 17. We call this method the ℓ_1 regularization method.

C. Gates on Irregular Grids

In practice, gates are not placed on a regular layout grid. Thus, in this section, we extend our method to irregular grids.

Figure 3 shows an example of an IC in which gates are placed on an irregular grid. To address the irregular placement, we cover the IC with a finer regular grid. Then each gate is assigned to a point on the regular grid using Procedure 1 below. At the first step of Procedure 1, we label all the regular grid points as *unassigned*. This means that none of the regular grid points is assigned to any




Fig. 3. Gates on irregular grids.

gate. In the second step, for every gate, we find its closest regular point that is *unassigned*, assign the gate to this point, and label that point as *assigned* to prevent multiple selection.

Thus, after Procedure 1, each gate is assigned to its closest regular grid that is not assigned to any other gate.

PROCEDURE 1

Mapping from irregular gates to fine regular grids

- (1) Set all the regular grid points *unassigned*;
- (2) for all gates, g_i
 - a. p = the closest grid point to the gates that is *unassigned*;
 - b. assign gate g_i to p ;
 - c. label grid point p as *assigned*;

Finally, we assign auxiliary variables to all the unassigned points in the regular grid.

V. APPLICATIONS

The proposed timing characterization method is effective, inexpensive, and fast. A range of technical applications can profit from the extracted post-silicon delay characteristics, including:

- (1) *Post-silicon optimization*. Fast noninvasive IC characterization, enables applying chip-specific optimizations [14], [3].
- (2) *Improving simulations*. The post-silicon models can be integrated within the simulation platforms to enable more accurate simulations.
- (3) *Improving SSTA methods*. The aggregate statistics gathered from post-silicon characterization can also be used to enhance the quality of the pre-silicon models, such as SSTA.
- (4) *Manufacturing process characterization*. The processes and technologies of the state-of-the-art manufacturing are considered classified information that are not typically available to the users. The new method can make accurate post-silicon estimation for a number of important process parameters.
- (5) *IC identification*. Since the variation is unique and unclonable on each IC, it can be used as the chip's ID for security [15], [16], [17].

VI. EVALUATION RESULTS

In this section, we evaluate the performance of the proposed variation estimation methods on the MCNC benchmarks.

Not that it is not possible to find the exact delay characteristics of all gates. The estimation error is measured in the space of singular values. The estimation error is the least in the direction of the singular vector corresponding to the largest singular value and increases in the direction of the singular vectors corresponding to decreasing the singular values. We call the estimation subspace n_e ; we project estimation error to the space of the singular vectors corresponding to the largest singular values in the n_e subspace.




Fig. 4. Variation (delay) estimation error vs. measurement error.




Fig. 5. Variations (delay) estimation error vs. the number of measurements.

To evaluate the performance of the proposed methods, we simulated the variation model (Section II-A) on a number of MCNC benchmark circuits. A total of 12% random variations is assumed. Correlated intra-die variation is 60% of the total variation [18]; 20% of the total variation is uncorrelated intra-die variation and the remaining variation is allotted to the inter-die variation.

We used the ABC tool to map the benchmark circuits to NAND2, NAND3, NAND4, NOR2, NOR3, NOR4, and inverter gates. Then, the gate placement is done by the Dragon placement tool. The gates have different sizes and they are located on irregular grids.

Figure 4 shows variation estimation error for both the ℓ_2 minimization and the ℓ_1 regularization methods on two benchmarks, C432, and C880. The horizontal axis is delay measurement noise and the vertical axis is variation estimation error. In average, the ℓ_1 regularization improves the estimation error by a factor of 2 over the ℓ_2 minimization. The estimation subspace is 52 and 89 for the C432 and the C880 circuits respectively. When the measurement noise is small, delay measurements provide enough information to estimate the variations accurately. As measurement noise increases, sparsity provides a strong prior that effectively de-noises the measurements. Thus, the performance gap between the ℓ_1 regularization and the ℓ_2 minimization increases as the measurement noise increases.

The impact of the number of measurements is demonstrated in Figure 5. The x-axis is the number of delay measurements divided by the number of the gates. Again, ℓ_1 regularization exhibits a factor of 2 improvement compared to ℓ_2 minimization. The estimation subspace is the same as in Figure 4.

Table 1 summarizes the results of variation estimation on 12 benchmark circuits. After the benchmarks' name in the first column, the second, third and fourth columns are the number of gates, the number

Circuit properties					3% noise		6% noise		9% noise		
name	#gates	#inputs	#meas	$\frac{\sigma_{N/4}}{\sigma_1}$	subspace	ℓ_1 error	ℓ_2 error	ℓ_2 error	ℓ_2 error	ℓ_1 error	ℓ_2 error
C432	206	36	309	0.035	26	3.76	6.82	4.34	12.86	5.23	17.25
					52	6.57	12.58	7.75	21.22	9.5	30.846
C499	532	41	798	0.045	67	4.05	4.78	4.74	6.91	5.70	9.35
					135	11.52	12.28	12.48	15.11	13.80	18.60
C880	353	60	529	0.043	44	2.65	5.45	4.27	10.61	5.99	22.49
					89	5.34	11.56	7.93	21.71	10.9	36.5
C1355	517	41	775	0.038	65	2.55	4.11	4.17	7.87	5.90	11.69
					131	5.22	7.10	8.21	13.19	11.41	19.47
C1908	615	33	992	0.052	78	2.56	2.77	4.05	71.61	5.68	100
					156	4.78	5.25	7.57	70.94	10.58	97.21
C2670	900	233	1350	0.019	114	2.26	3.03	3.48	5.54	4.84	8.17
					229	5.22	7.27	7.66	13.29	10.51	19.60
alu2	360	10	540	0.0519	45	2.54	10.69	3.74	21.30	5.17	38.78
					91	4.88	25.70	7.89	51.28	11.28	78.55
alu4	733	14	1099	0.036	93	3.63	12.79	6.01	100	9.76	100
					186	6.42	20.41	10.22	102.93	15.76	102.93
comp	163	32	244	0.061	20	1.16	1.78	1.71	3.11	2.34	4.51
					41	2.63	4.43	3.81	8.05	5.19	11.87
cordic	102	23	153	0.099	13	3.37	5.11	5.04	9.41	6.93	13.90
					26	8.38	15.93	13.10	29.89	16.91	44.17
b9	113	41	169	0.15	14	1.62	11.19	2.13	22.34	2.75	33.50
					28	3.17	13.13	4.11	25.48	5.24	38.01
c8	165	28	247	0.22	20	2.32	9.43	4.12	18.72	5.85	28.03
					41	5.10	14.09	9.33	27.95	13.10	41.84

TABLE I

PERFORMANCE OF ℓ_2 -NORM MINIMIZATION AND ℓ_1 -NORM REGULARIZATION FOR A NUMBER OF MCNC BENCHMARK CIRCUITS.

of inputs to the circuit, and the number of delay measurements respectively. The fifth column is the ratio of the $N/4$ -th singular value to the first singular value in the measurement matrix (N is number of gates.) This column shows how fast the singular values decay; or how the measurement matrix is well conditioned. The sixth column is the estimation subspace. The rest of the columns represent the percent estimation error for ℓ_2 minimization and ℓ_1 regularization for 3%, 6%, and 9% measurement noise.

VII. CONCLUSION

We have introduced a novel approach for post-silicon gate-level timing characterization. The approach leverages the new theory of compressed sensing to accurately estimate the gate-level delays using only a few noninvasive measurements. To implement the approach, we employed the set of sensitizable paths, sparse representation of the delay variation, structural logic relations, and methods to account for gate layouts irregularities. Experimental results demonstrated that the post-silicon timing of the benchmark circuits could be characterized with an average accuracy of 95% in the pertinent subspace.

VIII. ACKNOWLEDGEMENT

This work is partly supported by the DARPA/MTO Young Faculty Awards W911NF-07-1-0198 and NSF CT-0716674.

REFERENCES

- [1] B. Cline, K. Chopra, D. Blaauw, and Y. Cao, "Analysis and modeling of CD variation for statistical static timing," in *ICCAD*, 2006, pp. 60–66.
- [2] A. Ramalingam, G. Nam, A. Singh, M. Orshansky, S. Nassif, and D. Pan, "An accurate sparse matrix based framework for statistical static timing analysis," in *ICCAD*, 2006, pp. 231–236.
- [3] Q. Liu and S. Sapatnekar, "Confidence scalable post-silicon statistical delay prediction under process variations," in *DAC*, 2007, pp. 497–502.
- [4] D. Shamsi, P. Boufounos, and F. Koushanfar, "Noninvasive leakage power tomography of integrated circuits by compressive sensing," in *ISLPED*, 2008.
- [5] E. Candes, "Compressive sampling," in *Int. Congress of Math*, 2006, pp. 1433–1452.
- [6] R. Baraniuk, "A lecture on compressive sensing," *IEEE Signal Processing Magazine*, vol. 24, no. 4, pp. 118–121, 2007.
- [7] F. Liu, "A general framework for spatial correlation modeling in VLSI design," in *DAC*, 2007, pp. 817–822.
- [8] X. Lu, Z. Li, W. Qiu, D. M. H. Walker, and W. Shi, "Longest path selection for delay test under process variation," in *ASP-DAC*, 2004, pp. 98–103.
- [9] V. Iyengar, J. Xiong, S. Venkatesan, V. Zolotov, D. Lackey, P. Habitz, and C. Visweswariah, "Variation-aware performance verification using at-speed structural test and statistical timing," in *ICCAD*, 2007, pp. 405–412.
- [10] Z. Feng, P. Li, and Y. Zhan, "Fast second-order statistical static timing analysis using parameter dimension reduction," in *DAC*, 2007, pp. 244–249.
- [11] A. Murakami, S. Kajihara, T. Sasao, I. Pomeranz, and S. M. Reddy, "A test structure for characterizing local device mismatches," in *ITC*, 2000, p. 376.
- [12] D. Donoho, "Compressed sensing," *IEEE Trans. on IT*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [13] S. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [14] Y. Alkabani, T. Massey, F. Koushanfar, and M. Potkonjak, "Input vector control for postsilicon leakage current minimization in the presence of manufacturing variability," in *DAC*, 2008.
- [15] Y. Alkabani and F. Koushanfar, "Active hardware metering for intellectual property protection and security," in *USENIX Security*, 2007, pp. 291–306.
- [16] Y. Alkabani, F. Koushanfar, N. Kiyavash, and M. Potkonjak, "Trusted integrated circuits: A nondestructive hidden characteristics extraction approach," in *IH*, 2008.
- [17] F. Koushanfar and M. Potkonjak, "CAD-based security, cryptography, and digital rights management," in *DAC*, 2007.
- [18] Y. Cao and L. T. Clark, "Mapping statistical process variations toward circuit performance variability: an analytical modeling approach," in *DAC*, 2005, pp. 658–663.