

Process and Reliability Sensors for Nanoscale CMOS

John P. Keane*, Chris H. Kim, Qunzeng Liu**, and Sachin S. Sapatnekar

University of Minnesota, Minneapolis, MN 55455, USA

*now with Intel Corporation, Hillsboro, OR 97124, USA

** now with LinkedIn Corp., Mountain View, CA 94043, USA

Abstract

A key component for enabling smart silicon is the ability to create sense/respond feedback loops that allow adaptive circuit behavior as a response to variations. Such systems require the use of sensing techniques that describe the changes that the circuit experiences. In this paper, we describe the use of sensing schemes that drive on-chip process and aging variation measurements in manufactured silicon. We describe three case studies that show techniques for designing sensors: one of these is targeted to aging variations, while the other two capture process variation effects. The first case study, demonstrated in silicon, shows how beat frequencies can be used to build silicon odometers that capture how much a circuit has aged over its lifetime. The second and third case studies utilize a mix of presilicon process characterization data (as used in statistical timing/power analysis) and postsilicon measurements to predict the performance drift due to process variations.

Keywords

Reliability, Variations, Bias Temperature Instability, Sensors, Critical Paths, Silicon odometer

1. Introduction

In the nanometer regime, CMOS integrated circuits experience significant degradations in predictability and reliability. These variations in circuit behavior arise from a variety of sources whose underlying causes may be traced to process variations, environmental variations, or aging. Current solutions to overcoming variations include methods that overdesign systems, or that statistically design a population of circuits to ensure that a desired fraction meet specifications. However, such approaches can provide limited relief since the magnitude of variability and the number of failures, both along the spatial and temporal axes, are showing a markedly increasing trend. The net result of these effects is that variability is steadily eating into design margins and yield.

An alternative to static presilicon design margins is the use of postsilicon sensor-based adaptation schemes that capture multiple sources of variation, and use sensor data to mitigate the impact of these variations. Such “smart silicon” strategies are enabled by closed-loop control systems. Typically, these consist of a feedback loop that assimilates the information captured by various sensors, processes it, and triggers a response that suitably adapts the circuit to ensure that it meets its specifications. An example of such a scheme, using adaptive body biases and adaptive supply voltages, is shown in Figure 1. Data from a set of sensors is sent to a controller that generates body bias values, v_{bn} and v_{bp} , as well as the supply voltage, V_{dd} . An example of such an

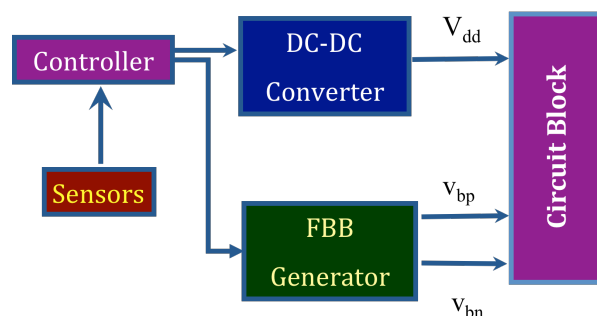


Figure 1: Schematic of a combined adaptive body bias and adaptive supply voltage adaptation scheme

optimization is a classical adaptive body bias scheme¹ where a sensing circuit is used to predict the performance drift, and this is used to feed back a signal that adapts the body bias over the die to return it to work within performance bounds.

The critical difference between presilicon and postsilicon optimizations can be summarized as follows: presilicon optimizations attempt to optimize over an entire *population* of manufactured parts, while postsilicon fixes are focused on die-specific solutions that improve the behavior of each specific manufactured part. In a large-volume manufacturing scenario, this implies that postsilicon optimizations must be simple, fast, and effective, and must be able to adaptively improve the behavior of the circuit over its entire lifetime.

The heart of any postsilicon detection scheme is a sensor structure that is used to determine the nature of the variations, and the focus of this paper is on sensing schemes for process and aging variations. The nature of these schemes can be different: process sensors detect the effect of one-time variations from manufacturing, which cause process parameters to shift from their nominal values on a specific part, while aging sensors detect run-time variations due to circuit aging. Process sensors can be used to enable compensations at “ $t=0$ ” (i.e., after manufacturing), while aging sensors supply different data over the life of the chip. Of course, even aging sensors must capture the effects of process variations on various aging and performance effects. Various schemes may be used to build on-chip sensors that enable smart silicon, and these may be categorized as follows:

Leveraging presilicon+postsilicon vs. postsilicon-only data: To enable the design cycle, the statistics of CMOS processes are typically characterized for presilicon analyses. While these statistics are true for the overall population of all chips and are not specific to a single manufactured part, they may be leveraged, using both this presilicon data and some postsilicon measurements, to build sensors that correlate well with design characteristics. However, process characteristics may shift over the lifetime of the process, possibly requiring recalibration of the sensors. An alternative class of sensors avoids this issue by using postsilicon measurement data alone.

Circuit-specific sensors vs. general sensors: Some sensing schemes are tailored to the specific circuit blocks that they are intended to measure, e.g., by setting their sensitivities to variations to be similar to those of the original circuit, often leveraging presilicon data (described above) to match the sensor to the circuit. Others capture general chip-wide trends and do not attempt to be circuit-specific.

In this paper, we overview some case studies describing the design of sensors to detect on-chip variations, bringing out the use of sensors under various combinations of the above three categories. Our first case study in Section 2 describes techniques for building general aging sensors based on postsilicon-only data. These sensors are designed to capture time-varying performance shifts due to aging phenomena based on ring oscillators (variously abbreviated as “ROSC” or “RO”), based on the concept of beat frequencies. Next, in Section 3 we describe two case studies for circuit-specific process sensors that leverage both presilicon and postsilicon data. Finally, we present concluding remarks in Section 4.

2. On-chip odometers for reliability sensing

Parametric shifts or circuit failures caused by Hot Carrier Injection (HCI), Bias Temperature Instability (BTI), and Time Dependent Dielectric Breakdown (TDDB) in CMOS technology have become more severe with shrinking device sizes and voltage margins. These phenomena may result in parametric variations, which cause circuit delays and leakage to shift as the circuit ages, or catastrophic failure. To enable adaptivity, or even to measure the performance degradation efficiently, it is essential to embed on-chip reliability monitors that aim at accurately measuring these aging effects.

A new class of on-chip reliability monitors has recently been demonstrated in silicon, targeted for a range of applications such as reliability characterization during process ramp up, chip lifetime projection, in-field product data collection, and real-time aging compensation. These monitors are referred to as *odometers*, capturing their ability to determine how far down the road to aging a chip has proceeded. Dedicated odometers provide several important advantages that are not achievable using traditional probing or standard ROSC structures. First, using odometers to control the aging measurements enables extremely high measurement resolution (e.g. sub-picosecond). This is particularly critical for accelerated stress experiments in which the voltage and temperature conditions are elevated beyond normal conditions for keeping the test time attainable. The odometer's high resolution allows voltage and temperature stress conditions closer to the actual product operating condition, significantly enhancing the confidence of the lifetime estimation results. Second, the measurement time can be made short (e.g. sub-microsecond) using a dedicated odometer which is essential when interrupting stress to record BTI measurements, as this mechanism is known to recover within microseconds. Finally, from a device characterization point of view, odometers are useful for monitoring statistical behaviors as a large number of devices can be stressed in parallel under various configurations, resulting in a large experiment time speedup.

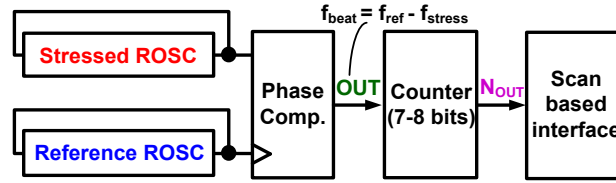


Figure 2: Simplified diagram of silicon odometer beat frequency detection system

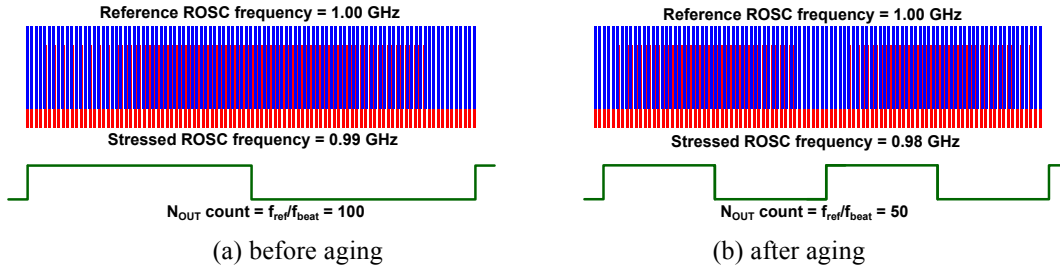


Figure 3: Illustration of the silicon odometer system. The beat frequency between the reference ROSC (blue bars) and the stress ROSC (red bars) is measured using simple digital logic gates. A 1% degradation in the stressed ROSC frequency (0.99 GHz \rightarrow 0.98 GHz) translates into a 50% change in the output count N_{OUT} (100 \rightarrow 50) achieving a high frequency resolution within a short measurement time.

We now describe the central idea behind the odometer designs that our group has demonstrated in several unique test chips^{2,3} in processes ranging from 65nm to 32nm. Figure 2 shows a simplified diagram comprising a pair of ROSCs, a phase comparator, a counter, and a scan based interface. During the short measurement periods, a phase comparator uses a fresh reference ROSC to sample the output of an identical stressed ROSC. The phase comparator can be built using a standard D-flip-flop circuit with outputs from the two ROSCs connected to the data and clock inputs, respectively⁴. In this configuration, the output of the stressed ROSC is sampled at every rising edge of the reference ROSC output producing a signal that exhibits the beat frequency as shown in Figure 3. ROSCs can be either put into stress mode or kept fresh by switching the local power supply using power gates. The output signal of this phase comparator exhibits the beat frequency, which is the difference between the frequencies of an unstressed reference ROSC and the stressed ROSC under test. A counter is used to measure the beat frequency by counting the number of reference ROSC periods during one period of the phase comparator output signal. This count is recorded after each stress period to calculate the shift down in the stressed ROSC frequency. Suppose the initial frequency of the reference ROSC is

called f_{ref} , that of the fresh ROSC to be stressed is f_{stress} , and the initial output count is N_1 . Also, without the loss of generality, we can assume that f_{ref} is slightly higher than f_{stress} . This condition can be achieved in a real chip using frequency trimming circuits and a simple scan-in test interface. The period of the beat frequency signal is the time it takes for the reference ROSC to accumulate N_1 periods or for the stressed ROSC, which has a slightly lower frequency, to accumulate (N_1-1) periods. This one clock period difference arises from the fact that the output of the stressed ROSC with a longer period will take one less period to cycle back and align with the output of the reference ROSC while the two ROSCs are free oscillating. Hence, the period of the beat frequency signal can be expressed as:

$$\frac{1}{f_{ref}} \cdot N_1 = \frac{1}{f_{stress}} \cdot (N_1 - 1) \quad (1)$$

At the end of a stress period, f_{ref} will remain unchanged, but f_{stress} will be decreased due to aging, and we call the new frequency f'_{stress} . We also have a new output count (N_2), so the resulting equation is:

$$\frac{1}{f_{ref}} \cdot N_2 = \frac{1}{f'_{stress}} \cdot (N_2 - 1) \quad (2)$$

Using these two equations, we can calculate a frequency shift during stress as follows:

$$\frac{f'_{stress} - f_{stress}}{f_{stress}} = \frac{N_1 \cdot (N_2 - 1)}{N_2 \cdot (N_1 - 1)} - 1 = \frac{(N_2 - N_1)}{N_2 \cdot (N_1 - 1)} \quad (3)$$

Those simple calculations show that if f_{ref} is only slightly higher than f_{stress} , the output count is high. For example, the count N_1 is 100 for a 1% frequency difference as shown in Figure 4(b). This slight difference can be ensured with trimming capacitors and calibration. The subsequent small decreases in f_{stress} due to aging cause a large change in this count. For instance, a 2% difference between the ROSC frequencies gives an N_2 of 50, so a 1% shift to that point is translated into a decreased count of 50 (Figure 4(b)). Consider another example in which the count value changes from 100 before stress, to 99 after stress. In this scenario, the corresponding frequency shift according to equation (3) can be calculated as

$$\frac{f'_{stress} - f_{stress}}{f_{stress}} = \frac{(99 - 100)}{99 \cdot (100 - 1)} \approx -0.010\%. \quad (4)$$

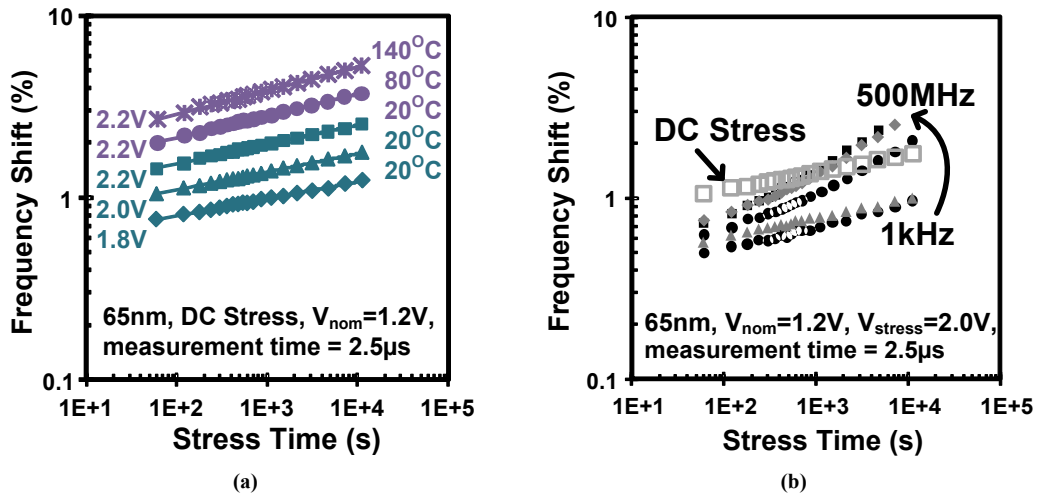
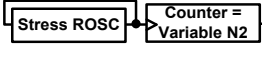
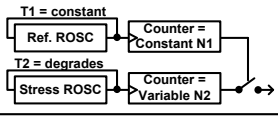
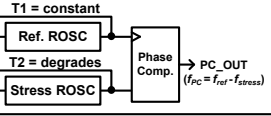


Figure 4: (a) DC stress results for different stress conditions. Voltage and temperature acceleration trends for BTI can be modeled for chip lifetime projection. (b) AC stress results compared with DC. We see a drop in degradation of roughly half at low frequencies, compared with DC stress, due to BTI recovery. As frequency is raised, HCI plays a larger role in the aging. This leads to a larger power law exponent, which is a signature of HCI⁵.

Therefore, with high frequency ROSCs, the beat frequency detection system can achieve extremely high measurement resolution on the sub-picosecond order. It is important to note that

the number of counts required to obtain such high resolution is very small (i.e. around 100 ROSC cycles). This is crucial for preventing the fast BTI recovery that occurs within microseconds after the stress conditions are removed for measurements. The proposed odometer technique can be used to measure aging in other circuit structures (for example, a critical path or a memory delay path) by simply replacing the inverter stages with other logic gates of interest. It's also worth mentioning that the accuracy of the beat frequency detection scheme becomes higher when the two ROSC frequencies are closer to each other.

Now, we present sample results collected from a previous test chip design to demonstrate the effectiveness and convenience of our odometer designs for reliability monitoring. Figure 4 (a) displays the frequency shifts measured under different voltage and temperature conditions. A DC stress was applied and the measurement time was limited to less than a microsecond. Results show that BTI induced degradation depends exponentially on the stress voltage and stress temperature. The acceleration factors with respect to the two stress parameters can be readily calculated based on the measured data. Figure 4 (b) shows the frequency shifts for both DC stress and AC stress conditions. At low AC stress frequencies, a drop in total degradation of $\sim 1/2$ compared with DC stress is observed, due to the BTI recovery that takes place during in each half cycle. As the frequency is raised, HCI plays a larger role in the aging due to the increased switching activity. This leads to a larger power law exponent, which is a signature of HCI.

System	Single ROSC	2 ROSC, simple	2 ROSC, beat freq.
Block Diagram			
Function	Count Stress ROSC periods during externally controlled meas. time	Count Stress ROSC periods during N1 periods of Ref. ROSC	Count Ref. ROSC periods during one period of PC_OUT
Features	Simple; compact	Simple; immune to common mode variations	High resolution w/ short meas. time; immune to common mode variations
Issues	Voltage and temp. variations; meas. time vs. resolution tradeoff; requires absolute timing reference (e.g. oscilloscope)	Meas. time vs. resolution tradeoff	Requires extra circuits (e.g., Phase Comp., edge detector, etc...)
Meas. time for 0.01% resolution *	30 μ s	30 μ s	0.3 μ s
Meas. error wrt. common mode variations **	+10.18% / -8.57%	+0.26% / -0.38%	+0.06% / -0.07%

*ROSC period = 3 ns ** simulated with $\pm 4\%$ ΔV_{CC}

Figure 5: Comparison of ROSC based aging monitors⁶.

Figure 5 compares various ROSC based systems for monitoring frequency shifts including the beat frequency detection method. The scheme denoted as “single ROSC” measures the output of a single ROSC whose frequency is divided down by a counter for easier off-chip frequency measurements using test equipment. The “2 ROSC, simple” scheme measures the degradation in one stressed ROSC by counting the number of periods it cycles through while a set number of periods in a fresh reference ROSC are counted¹. The resolution of the first two schemes in Figure 5 is simply the measurement time divided by the ROSC period, while that of the beat frequency technique can be derived from equation (3), we see that the latter technique reaches a maximum resolution of 0.01% within only 0.3 μ s with a single measurement recorded, while the other systems require 100X more time to reach the same accuracy. The longer measurement times in the standard period counter systems would result in unwanted BTI recovery which can occur within a few microseconds or less. In addition to the high frequency

resolution, the beat frequency technique benefits from a high immunity to voltage or temperature variations due to its differential nature. Simulation data provided in the bottom row of Figure 5 verifies that both the “2 ROSC, simple” and “2 ROSC, beat frequency” schemes have significantly less error compared to the “Single ROSC” scheme in the presence of common mode supply voltage fluctuations. Considering that these reliability sensors are simpler than the critical path monitors widely used in real products⁸, we do not anticipate the sensor area to be a major concern.

In summary, the beat frequency technique achieves a significantly higher frequency measurement resolution (sub-picosecond) in a shorter measurement time (sub-microsecond) than traditional measurement techniques with only a modest increase in the number of circuits, and is immune to common mode environmental variations. As proven through our previous test chip studies, the attractive features of this technique make it an indispensable apparatus for circuit aging research.

3. Process variability sensors for performance prediction

Process variations, which affect parameters such as dopant density, the effective channel length, and the oxide thickness, are widely acknowledged to significantly affect circuit performance. These may be classified into systematic or random variations, depending on whether they are deterministic or probabilistic. Random variations may be either die-to-die (D2D) variations from one chip to another, as well as within-die (WID) variations among different locations, which often show spatial correlations, within a single die.

At the presilicon stage, spatial correlations are typically precharacterized for a given process, and these parameters are fed into statistical static timing analysis (SSTA) computations for circuit analysis. Spatial correlations may be characterized using various models, based on square or hexagonal grids, and typically represent the delay as a function of a fundamental set of random variables, p_1, \dots, p_n . These variables are typically correlated Gaussians, with the correlation decaying with distance: it is common practice to use methods such as principal components analysis (PCA) to orthogonalize the variations into a set of uncorrelated normalized basis variables, v_1, \dots, v_n . The circuit delay D can then be represented in the following *canonical form*:

$$D = D_0 + \sum_{i=1 \text{ to } n} k_i v_i + k_{n+1} v_{n+1}$$

where k_i , $1 \leq i \leq n$ correspond to the coefficients associated with the n principal components, and k_{n+1} is the coefficient of all uncorrelated random sources, lumped into one variable v_{n+1} .

There is a well-established body of literature that efficiently conducts presilicon circuit analysis to determine the effects of random one-time variations on key circuit performance parameters such as delay and power. SSTA and statistical power analysis techniques⁹ to evaluate the effects of one-time process variations in performance have already found their way into industrial practice. The computational efficiency of these methods is made practical through a preprocessing step which has shown that Gaussian-distributed correlated variations can be orthogonalized using principal component analysis (PCA); similarly, non-Gaussian distributions can be processed using independent component analysis (ICA).

In our discussion, we refer to the circuit to be evaluated as the *original circuit*, and the detection mechanism as the *test circuit*. The parameters of interest for the original circuit are typically the delay and power dissipation. Without precisely measuring the original circuit, it is not possible to estimate its delay, and our goal is to use a limited set of measurements to substantially narrow down the range of the performance distribution. The designer has the freedom to design these test structures so that they provide the maximum amount of information.

Traditional sensing schemes are based on test structures in the form of ROs, inverter chains, “critical path replicas,” a set of simple delay elements, or using a rudimentary look-up table approach based on gross estimates of logic and interconnect speeds¹⁰ that aim to capture the

variations in the performance of the original circuit. However, such structures are based on simple assumptions, and may provide incorrect results. For example, measuring the RO can provide accurate information about its behavior – but there is a risk in using this information to draw an inference about the original circuit, which can have a very different structure, and very different sensitivities to variations. A critical path replica, on the other hand, can only capture the behavior of one path – but real circuits have a very large number of critical paths, each of which can have different sensitivities to variation. It is possible, and indeed likely, that after manufacturing, the actual critical path is different from the nominal critical path. Therefore, any inference drawn from a replica of a single critical path, such as the nominal one, is unreliable. While it is possible to account for these inaccuracies by “padding” the prediction with a conservative margin, such an approach must incorporate pessimism, which is liable to unnecessarily leave a considerable amount of performance on the table.

A fundamental limitation of on-chip surrogate test structures is that they cannot capture the effects of independent (spatially uncorrelated) variations. By definition, there is no relationship between uncorrelated parameter variations (aside from sharing gross statistics) of a test structure and those in the rest of the circuit, and the only way to determine these variations fully is through delay test of the original circuit. A second effect is that such variations add noise to the measurements from on-chip test structures: this can be overcome by “drowning out” uncorrelated random variations by increasing the number of stages in the test structure (noting that the sigma/mu of these variations reduces with the number of stages).

We describe two methods below: one based on building ring oscillators, and another on building a synthetic path that represents the variations in the entire circuit.

3.1 RO test structures

This framework is based on a set of distributed ROs, spread throughout the area of the chip. By measuring the speed of these ROs, it is possible to infer information about the speed of the original circuit, based on the use of conditional probabilities. Intuitively, each RO can capture the spatial variations in the region in its neighborhood; with a sufficient number of well spread-out ROs, it is possible to recapture the value of the deterministic variations and the spatially correlated random variations^{11,12}.

Assume that n ROs are placed on a chip, and define a delay vector

$$\mathbf{d}_i = [d_{i,1}, d_{i,1}, \dots, d_{i,n}]^T$$

for the test structures, where $d_{i,i}$ is the random variable (over all manufactured chips) corresponding to the delay of the i^{th} test structure.

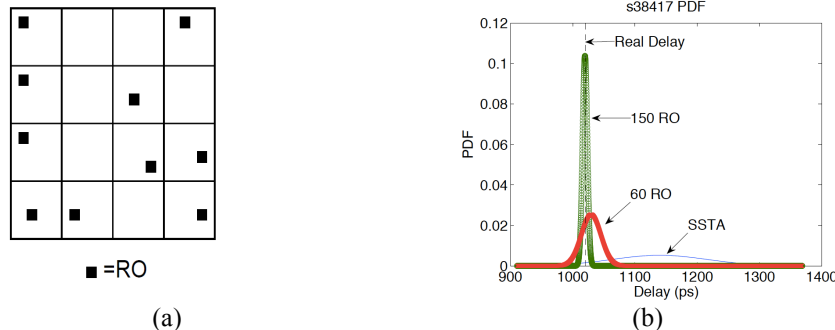


Figure 6: (a) Distribution of RO test structures over a die area (b) Prediction of circuit performance based on measurements from these structures.

For a particular fabricated die, the delay of the original circuit and the test structures correspond to one sample of the distribution that describes the underlying process parameters, and

this results in a specific value of $d = d_{real}$ and of $\mathbf{d}_t = \mathbf{d}_{real}$. Since the n -RO scheme is much smaller and less complex than the original circuit, these measurements can be performed rapidly. The problem of delay prediction can be stated as a conditional probability evaluation, corresponding to finding

$$d_{real} = f(d | \mathbf{d}_t = \mathbf{d}_r).$$

Based on the underlying statistics of the characterized process distribution, it is possible^{11,12} to compute the function f above. This computation is based on a presilicon analysis of the probability distribution of d and \mathbf{d}_t in canonical form, based on which the conditional distribution can be determined. With a sufficient number of test structures, the variance of the conditional distribution can be made sufficiently small, as illustrated in Figure 6. The widely-spread curve marked “SSTA” shows the overall distribution of delays based on presilicon SSTA for a placed instance of the ISCAS89 benchmark circuit, s38417. For a manufactured part whose real delay is shown by the dotted line marked “Real Delay,” a prediction based on 60 ROs is shown by red curve marked “60 RO”, and one using 150 ROs by the green curve marked “150 RO”: these can be seen to have progressively narrower variances.

3.2 Building a Representative Critical Path

The above RO-based approach consists of three steps: a) presilicon test structure insertion b) postsilicon test structure measurement c) postsilicon computations to infer the speed of the chip. The third step requires additional computation, which may not always be reasonably possible in the context of postsilicon adaptation. To bypass this step entirely, an alternative approach replaces the ensemble of ROs by a single synthesized test structure that replicates the delay characteristics of the original circuit¹³.

This structure, which we refer to as a representative critical path (RCP) for the circuit, leverages the property of spatial correlation between parameter variations and we optimally determine its structure and physical locations. Specifically, we synthesize the RCP and place it on the die so that its delay is highly correlated with that of the original circuit. Since the RCP is an on-chip test structure, it can easily be used within existing post-silicon tuning schemes, e.g., by replacing the nominal critical path in a typical sense/respond adaptive loop¹⁴. Two heuristic approaches for building the RCP were presented in our work¹³, and were demonstrated to work effectively.

This approach differs from the method based on a critical path replica in several ways. First, while the critical path replica attempts to exactly match the delay of the circuit by reproducing one or more critical paths, the delay of the RCP can be quite different from that of the circuit. Specifically, the criterion for building the RCP is to ensure that the sensitivity of its delay to variations is similar to that of the original circuit. In other words, if its delay changes by a certain value, then with a predictable level of confidence, we can forecast that the delay of the circuit changes by a specific value.

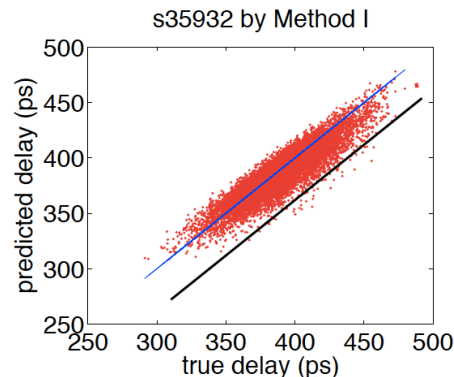


Figure 7: A comparison between the predicted and true delay for a synthesized RCP.

Figure 7 shows a quantitative comparison of our approach with a conventional method that replicates the nominal critical path for a benchmark circuit. The nominal critical path approach is always optimistic: to see this, note that the delay of any path (including the nominal critical path) in the circuit is a lower bound on the delay of the actual critical path in the manufactured circuit. Therefore, using the nominal critical path as a delay estimator is guaranteed to *underestimate* the actual delay. In contrast, the accepted practice in timing characterization is that delay estimates should be overestimated, or pessimistic.

In its purest form, our method may be either optimistic or pessimistic, and the scatter plot can be seen to lie on either side of the $x=y$ line in Figure 7. To overcome this, the mapping between the real and predicted delay is chosen to be on the black line, which captures a certain percentile (e.g., 99 percentile) of all points: the percentile can be chosen according to the required level of confidence. While such a pessimistic approach can also be applied to the nominal critical path, for the same confidence level, the inaccuracy is improved by 2x or more using the RCP. This result follows from the narrow variance of the predicted delay, similar to the sharp curve illustrated in Figure 6.

In this presentation, we have considered the use of an RCP to measure process variations. However, the same idea may also be used to measure aging variations: the RCP should be constructed so that its delay sensitivity to aging resembles that of the circuit it represents.

4. Conclusion

We have presented several case studies showing how sensor structures may be used to characterize process and aging variations in a chip. Various schemes have been presented, some of which rely purely on postsilicon measurements while others leverage presilicon characterizations of the process. Clearly, the latter class of structures are as accurate as the presilicon characterization, but if the process is seen to drift due to characterization, the updated parameters may be used, with occasional additional offline computations, to interpret the results of test structure measurements. Our case studies also show structures that are independent of the circuit structure and capture chip-wide variations, as well as others are more closely tied to the specific structure they aim to measure.

Acknowledgments

The authors would like to acknowledge the contributions of T. H. Kim, S. V. Kumar, R. Persaud, X. Wang, and W. Zhang to the work described in this paper.

Author Biographies

John P. Keane (S'06, M'11) received his Ph.D. in electrical engineering from the University of Minnesota, Minneapolis in 2010. He is currently with the Advanced Design organization at Intel Corporation where he is working on memory test circuits. John has co-authored over 20 peer-reviewed publications and is a co-inventor on five US patents.

Chris H. Kim is an Associate Professor of Electrical and Computer Engineering at the University of Minnesota. His current research interests include digital, mixed-signal, and memory circuit design for silicon and non-silicon technologies. He has a PhD in electrical and computer engineering from Purdue University. He is a senior member of IEEE.

Qunzeng Liu received his Ph.D. in Electrical Engineering from the University of Minnesota, Minneapolis, Minnesota. His research interests include statistical analysis and optimization of electronic circuits and software engineering. He currently works as a software engineer at LinkedIn Corp, Mountain View, California.

Sachin S. Sapatnekar received his Ph.D. from the University of Illinois at Urbana-Champaign and is currently a Professor of Electrical and Computer Engineering at the University of Minnesota. He is currently interested in research on circuit variability and reliability. He is a Fellow of the IEEE.

References

- ¹ J. W. Tschanz, J. T. Kao, S. G. Narendra, R. Nair, D. A. Antoniadis, A. P. Chandrakasan, and V. De, "Adaptive Body Bias for Reducing Impacts of Die-to-Die and Within-Die Parameter Variations on Microprocessor Frequency and Leakage," *IEEE Journal of Solid-State Circuits*, Vol. 37, No. 11, pp. 1396 – 1402, November 2002.
- ² J. Keane, W. Zhang, C.H. Kim, "An Array-Based Odometer System for Statistically Significant Circuit Aging Characterization," *IEEE Journal of Solid-State Circuits*, pp. 2374 – 2385, October 2011.
- ³ J. Keane, W. Zhang, C.H. Kim, "An On-Chip Monitor for Statistically Significant Circuit Aging Characterization", *Proceedings of the International Electron Devices Meeting*, pp. 4.2.1 – 4.2.4, December 2010.
- ⁴ T.H. Kim, R. Persaud, C.H. Kim, "Silicon Odometer: An On-Chip Reliability Monitor for Measuring Frequency Degradation of Digital Circuits," *IEEE Journal of Solid State Circuits*, Vol. 4, pp. 874 – 880, April 2008.
- ⁵ J. Keane, X. Wang, D. Persaud, and C.H. Kim, "An All-in-One Silicon Odometer for Separately Monitoring HCI, BTI, and TDDB," *IEEE Journal of Solid-State Circuits*, pp. 817 – 829, April 2010.
- ⁶ J. Keane, X. Wang, T. Kim, and C.H. Kim, "On-Chip Reliability Monitors for Measuring Circuit Degradation," *Microelectronics Reliability*, pp. 1039 – 1053, August 2010.
- ⁷ K. Stawiasz, K.A. Jenkins, L. Pong-Fei, "On-Chip Circuit for Monitoring Frequency Degradation Due to NBTI", *Proceedings of the IEEE International Reliability Physics Symposium*, pp. 532 – 535, April 2008.
- ⁸ A. Drake, R. Senger, H. Deogun, et al., "A Distributed Critical-Path Timing Monitor for a 65nm High-Performance Microprocessor," *Proceedings of the International Solid-State Circuits Conference*, pp. 398 – 399, February 2007.
- ⁹ H. Chang and S. S. Sapatnekar, "Statistical Timing Analysis Considering Spatial Correlations using a Single PERT-Like Traversal," *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 621 – 625, 2003.
- ¹⁰ A. Wang and S. Naffziger, eds., *Adaptive Techniques for Dynamic Processor Optimization: Theory and Practice*, Springer, New York, NY, 2008.
- ¹¹ Q. Liu and S. S. Sapatnekar, "Confidence Scalable Post-Silicon Statistical Delay Prediction under Process Variations," *Proceedings of the ACM/IEEE Design Automation Conference*, pp. 497 – 502, 2007.
- ¹² V. Zolotov and J. Xiong, "Optimal Statistical Chip Disposition," *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pp. 95 – 100, 2011.
- ¹³ Q. Liu, and S. S. Sapatnekar, "Synthesizing a Representative Critical Path for Post-Silicon Delay Prediction," *Proceedings of the International Symposium on Physical Design*, pp. 183 – 190, 2009.
- ¹⁴ J. W. Tschanz, J. T. Kao, S. G. Narendra, R. Nair, D. A. Antoniadis, A. P. Chandrakasan, and V. De, "Adaptive Body Bias for Reducing Impacts of Die-to-Die and Within-Die Parameter Variations on Microprocessor Frequency and Leakage," *IEEE Journal of Solid-State Circuits*, Vol. 37, No. 11, pp. 1396 – 1402, November 2002.