Leveraging Existing Power Control Circuits and Power Delivery Architecture for Variability Measurement

Dhruva Acharyya+, Kanak Agarwal*, Jim Plusquellic
+Verigy Ltd. (dhruva.acharyya@verigy.com)
*IBM Corp. (kba@us.ibm.com)
Dept. of Electrical and Computer Engineering, Univ. of New Mexico (jimp@ece.unm.edu)

Abstract

Embedded test structures are increasingly being used to measure and analyze performance and power variations in product chips to better understand the impact of process variations. In this work, we propose a minimally-invasive, low-overhead technique for characterizing within-die and die-to-die leakage variation. The proposed technique leverages existing power control circuitry added by designers to reduce the power consumption of inactive functional units. We manipulate these 'sleep islands' to isolate and measure their leakage current contribution to the chip-wide leakage current. The measured set of sleep island leakage currents reflect the leakage current variation across the chip at a coarse level of resolution. In order to improve resolution, we propose a multiple power supply port (MSP) measurement technique to provide 'within-island' leakage current measurements. A calibration technique is described that corrects for differences between the sleep island and MSP approaches, effectively enabling the same information to be obtained using either technique. We demonstrate the techniques on a set of test chips fabricated in 65-nm SOI technology. The results show that leakage current variations across a small test structure array have both a locally random and globally deterministic component that can be accurately mapped using the sleep island or MSP approaches.

1 Introduction

Controlling process variations is becoming more difficult with scaling. This is particularly apparent at the lowest level of measurement granularity, i.e., across-field (or within die), in contrast to the lot-to-lot, wafer-to-wafer and across wafer levels. The main sources of across-field variation are related to 1) the limitations over the precise control of optical sources, resulting in across-field focus and dose variation [1] as well as mask errors [2] and 2), layout dependent systematic effects such as pitch and density dependent line-width variability [3, 4] and microscopic etch loading [5]. In contrast, the sources of across-wafer variation are related to wafer-level non-uniformities such as post exposure bake (PEB) temperature gradient [6] and resist thickness variation [7].

Scribe line test structures, i.e., those placed in the physical space between dies that is later destroyed when the wafer is diced, are routinely used to monitor process variations at the lot-to-lot, wafer-to-wafer and across-wafer levels [8]. Unfortunately, this strategy is not effective for measuring within-die variation, and instead, test structures within the product itself (embedded test structures) are needed. Moreover, these embedded test structures need to be distributed across the die in order to determine the spatial characteristics of variation on the entire die. The challenges to adding such structures are similar to those that the manufacturing test community deal with in regard to design-for-testability, i.e., they must minimize area overhead, yield loss, performance/power impact, as well as have a small I/O interface, test cost, etc.

Test structure design continues to be an active area of research [9,10,11]. Embedded ring oscillators (EROs) have been successfully used to characterize within-die performance variations, and due to their simple design and I/O interface, they are the preferred embedded test structure [12,13,14]. Unfortunately, the analysis of ERO frequency that is used to measure performance variations provides only limited information regarding leakage current variations, and therefore alternative methods are needed.

Process variation can cause an order of magnitude variation in chip leakage [15]. Leakage current variations can have a significant impact on variations in product power consumption. Large levels of within-die variations in leakage can lead to undesirable scenarios in which a low leakage, low performance region of a chip limits the maximum speed but a high leakage region(s) on the same chip increases its power consumption. A slow, high-power chip is not a desirable combination. Therefore, it is important to develop techniques that can track leakage variations so that these types of problems can be identified and mitigated.

In this paper, we propose a product-oriented method for measuring within-die leakage variation that utilizes both global and local I_{DDQ} measurements. Global I_{DDQ} refers to the leakage current of the entire chip. Local I_{DDQ}, on the other hand, refers to the leakage currents measured from the multiple supply ports (MSP) of the chip. An important distinguishing characteristic of the proposed global I_{DDQ} measurement process is that it leverages existing power control circuitry inserted for implementing sleep islands. By manipulating the state of the sleep island control circuitry and measuring the resulting global leakage current, the leakage characteristics of each island can be isolated.

The core logic embedded within a sleep island may span a large region of the chip, and therefore, the sleep island approach may be able to provide only a coarse level of leakage current characterization. Resolution can be

978-1-4244-7207-9/10/$26.00

improved by combining the sleep island method with the MSP technique. The regular, spatial distribution of the supply ports enables the measurement of 'within-sleep-island' leakage currents. Therefore, by combining the sleep island and MSP techniques, a higher resolution leakage current variation profile can be obtained for the entire chip.

A major focus of this work is to demonstrate the sleep island approach and to validate a method that calibrates MSP data to sleep island data. In an actual application in which a high resolution map of leakage current variations is needed, the process described in this paper is useful for achieving this objective. However, other applications may not need such detailed leakage variation information. The following outlines the basic flow of our proposed methods, with various 'exit' points identifying the level of information obtained.

- Sleep Island by itself: If the chip is designed with a large number of sleep islands and high level of control over sleep states, it will be possible to obtain an accurate map of within-die leakage variations with sleep island alone, i.e., MSP is not needed.
- Sleep Island and MSP: If the chip is designed with a limited number of sleep islands or a limited level of control, MSP can be used directly to obtain an approximation of 'within-island' leakage variations.
- Sleep Island and MSP with calibration: Non-uniform power grid architectures, 'bleeding' effects and metal resistance variations (to be discussed) 'distort' the MSP values. If an accurate mapping of the chip's leakage characteristics is needed, calibration can be used to correct for these detractors. Calibration makes MSP equivalent to the sleep island approach but requires additional effort, i.e., on-chip calibration circuits and simulation experiments.
- In any of these scenarios, it is also possible to obtain some defectivity information.

We validate the sleep island and MSP techniques on a set of test chips fabricated in a 65-nm SOI technology. The following identifies the main contributions:

- Using sleep island, we determined that the within-die leakage current variations of our test chips span a range of approximately 90%. In comparison, the die-to-die variation analysis yields a change of approximately 300%.
- Using MSP and calibration, we were able to show a high level of correlation (> 80%) between the sleep island and MSP techniques for chips with overall leakage currents above the noise floor of our measurement setup.
- We were also able to identify two defective chips using either the sleep island or MSP techniques. The current behavior of these chips is easily distinguishable from the variations that occur in leakage current.

The remainder of the paper is organized as follows. In the next section, we describe the test chip architecture and hardware experiments. In Section 3, we analyze leakage variations using the sleep island approach. In Section 4, the multiple supply port (MSP) technique is described. Calibration methods are discussed in Section 5 that equate the MSP and sleep island approaches. Section 6 presents the correlation analysis of the sleep island and MSP techniques. The results of analyzing two defective chips is described in Section 7, and we present our conclusions in Section 8.

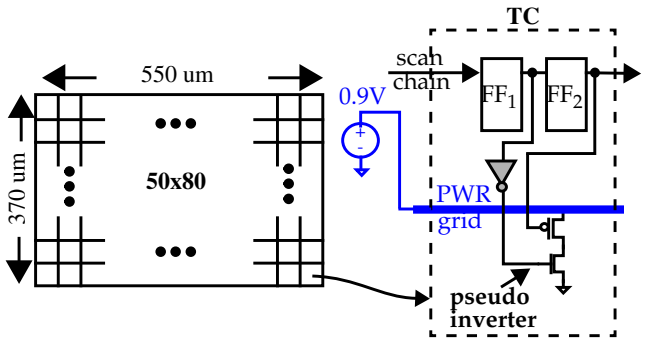

Fig. 1. Block diagram of the test macro. Schematic of one test circuit (TC) is also shown. The other TCs are identical.

2 Test Chip Architecture and Experiments

A block diagram of the test structure fabricated on our 65 nm SOI chips is shown in Figure 1 [18]. It consists of a 50x80 array of test circuits (TCs). The contents of a TC are shown on the right side of the figure. Each TC consists of two scan chain FFs whose outputs connect to individual gates of a pseudo-inverter. The scan chain allows control over the leakage state of each of the 4,000 pseudo-inverters. In this paper, we analyze the leakage characteristics of different regions or **islands** of the array by placing the TCs in that island into a high leakage state while placing all other TCs into a low leakage state. The low leakage state is realized by setting both FF_1 and FF_2 of a TC to logic '1', which disables both the nMOS and pMOS devices in the pseudo-inverter. The high leakage state is realized by placing the pattern '01' in FF_1 and FF_2. This pattern keeps the pMOS disabled but enables the nMOS device. By subtracting the current measured under a high leakage configuration from a base configuration, which places **all** TCs into low leakage, it is possible to isolate the leakage current contribution of the pMOS devices in the high leakage island.

Figure 2 shows a block diagram view of the TC array and the set of high leakage islands investigated in our experiments. The column labeled '4 island' shows 4 high leakage configurations of the array that provide a coarse level of leakage current characterization. This is true because the 4 region partitioning strategy places 1,000 of the TCs into the high leakage (HL) state and 3,000 TCs into the low leakage (LL) state. For example, the configuration labeled with **p1** (for leakage pattern 1) places the lower left island in HL (shown shaded in the figure) and the other three regions in LL. Similarly, pattern **p2** places the upper left island in HL, and so on. The configurations shown in columns labeled '16 island' and '64 island' incrementally improve on the leakage 'resolution' by partitioning the

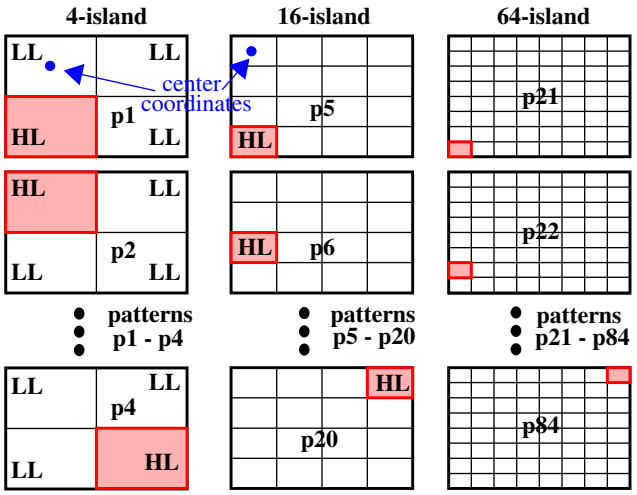

Fig. 2. Sleep island configurations of TC array. Left-most column labeled '4-island' provides low resolution, 16-island provides medium resolution and 64-island provides high resolution analysis of leakage current variations in the array.

array into a larger set of smaller islands.

In order to analyze leakage variations in the pMOS devices across the HL islands of the array, we first measure the global current with the array configured into a 'low leakage everywhere' pattern **p0** (not shown). We then subtract the p0 current from the global current measured under each of the 84 patterns shown in Figure 2. The resulting sets of 4, 16 and 64 'normalized' currents obtained for the 4, 16 and 64-island experiments, resp., define the **within-die** variation profile of the chip.

Given the propriety nature of the data, a **percentage change** metric is used to depict the variations in the normalized currents described above. The percentage change is computed by subtracting the normalized leakage current measured under a **reference pattern**, e.g. p1, from each of the currents measured under the remaining experiments in that group, e.g., patterns p2, p3 and p4 from Figure 2. A percentage is computed by dividing this difference by the reference current and multiplying by 100. Patterns p5 and p21 are used as the reference patterns for the 16 and 64 island experiments, respectively. A similar metric is defined for **die-to-die** variation. In this case, a **reference chip** is used instead of a pattern. The reference chip is constructed from a model in which 250 uAs of leakage current is uniformly distributed across the entire array. We collected data and processed it in this fashion for 33 chips.

3 Sleep Island Results

Figures 3 to 5 display the results in contour plots for three chips, resp., under each of the 4, 16 and 64 island experiments (rows). The column of contour plots on the left side in each figure give the with-in die results while the right column of contour plots give the die-to-die results. The (x,y) plane in the figures represent the (x,y) plane of the array as shown in Figure 2. The coordinates assigned to each island and plotted in Figures 3 to 5 are given as the center point of each region (see Figure 2 for some examples). The contours depict the leakage current characteris-

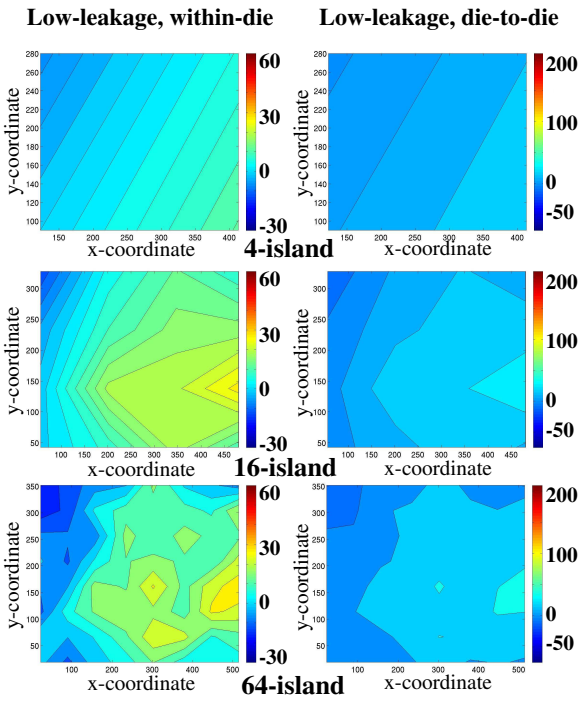

Fig. 3. Contour plots of low leakage chip of with-in die and die-to-die leakage current variations at 3 levels of resolution.

tics as a percentage change. For example, the contours for the 4 island experiments (top-most row) contain only 4 data points, one from each of the 4 island experiments shown in Figure 2. The current profiles for the 16 and 64 island experiments include 16 and 64 data points, resp. In order to facilitate comparisons between the three levels of resolutions and across chips, the range within all plots showing within-die variation results is set to -30% to +60%, and the range within all die-to-die variation plots is set to -75% to +210%.

The larger number of pMOS devices in high leakage, e.g., the 4 island experiments, tends to average out local variations, and reduces the overall range of variation (as expected). This is reflected in the contours by a smaller difference in the minimum and maximum values in comparison with the 16- and 64- island contour plots. The die-to-die analyses incorporate a DC component, with the entire contour skewed to one particular region of the range[1]. Figure 3 shows the results from a low leakage chip, Figure 4 gives a medium leakage chip and Figure 5 shows a high leakage chip from the population of 33 chips that we analyzed. Although the variation profiles exhibits a large degree of randomness, their also appears to be a systematic component where leakage magnitudes are smaller along the array's edges, particularly along the left edge. This is most easily observed in the 16- and 64-island analysis.

1. Note that the variation profile is identical for both the within-die and die-to-die plots for each chip. Only the DC component and scaling are different.

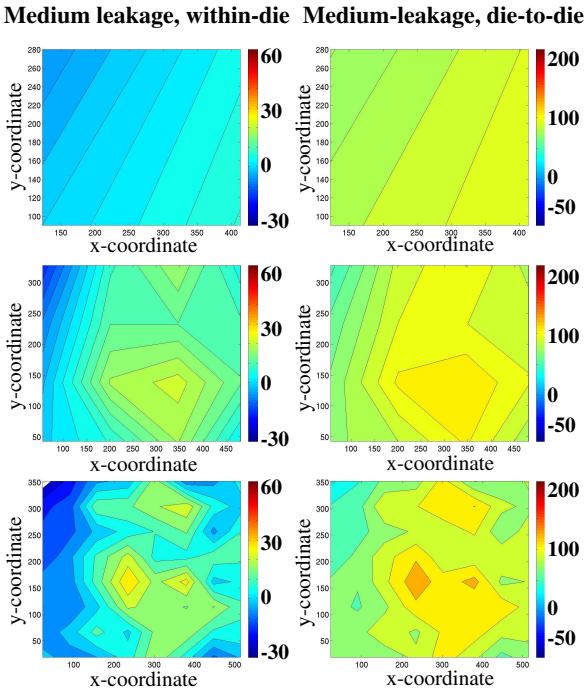

Fig. 4. Contour plots of medium leakage chip of with-in die and die-to-die leakage current variations at 3 levels of resolution.

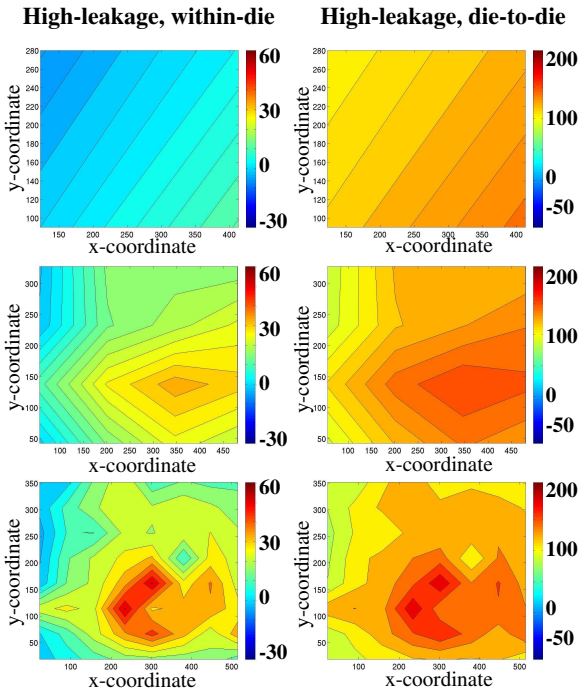

Fig. 5. Contour plots of high leakage chip of with-in die and die-to-die leakage current variations at 3 levels of resolution.

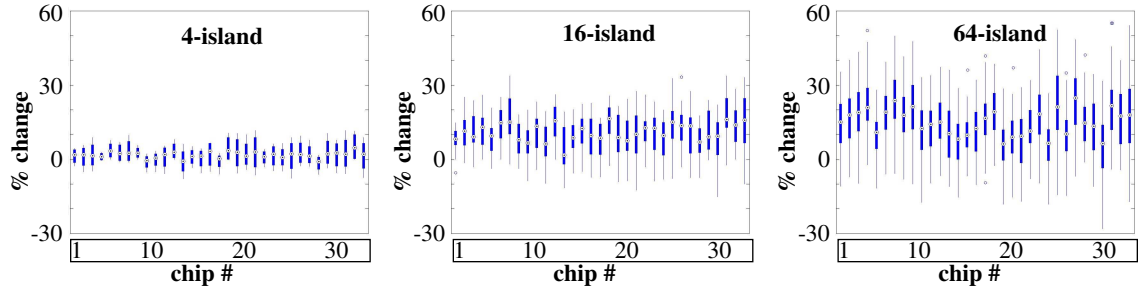

Fig. 6. Box plots of % change of within-die variation under 4-, 16- and 64-island experiments (left to right). Chips along x-axis are sorted such that overall leakage increases left to right in each plot.

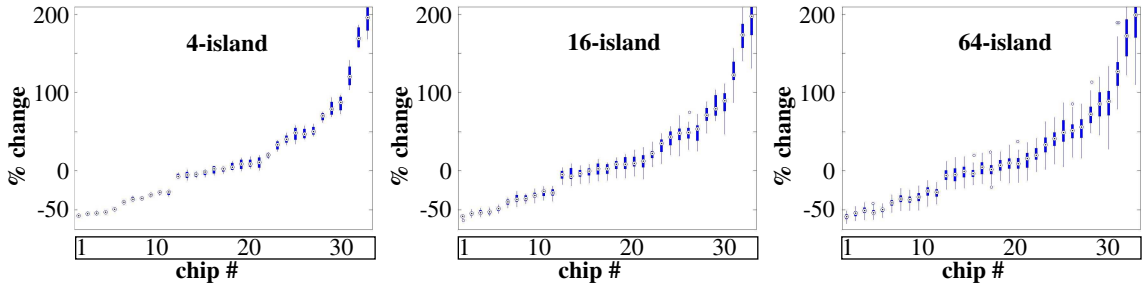

Fig. 7. Box plots of % change of die-to-die variation under 4-, 16- and 64-island experiments (left to right). Chips along x-axis are sorted such that overall leakage increases left to right in each plot.

The box plots in Figures 6 and 7 give a statistical view, i.e., medium value, extreme values, etc., of the within-die and die-to-die leakage variations, resp., of the 33 chips. The chips are sorted along the x-axis according to their overall leakage current magnitude characteristics, from low (chip 1) to high (chip 33). The y-axis plots the percentage change metric described above. The most notable characteristics in the box plots of Figure 6 is the increasing level of variation from left to right across them, i.e., the results for the 4-island experiment depicts small with-in die variations, while the 16- and 64-island results depict increasingly larger variations. Moreover, leakage current variation appears to track overall leakage current magnitude, as the range of percentage change within each box plot is relatively constant across the chips. In contrast, the die-to-die variation box plots in Figure 7 show a more pronounced

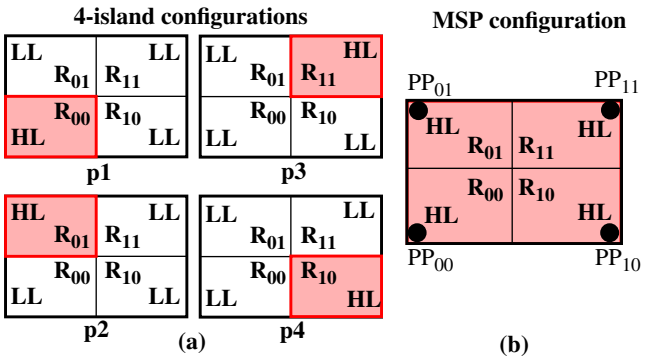

Fig. 8. Sleep Island leakage configurations (right) and MSP leakage configuration (left) for correlation analysis.

correlation in this regard, i.e., higher leakage chips have larger levels of variation. Moreover, the range of variation is approx. 3x larger in the die-to-die variation analysis. For example, the range of within-die variation is approx. 90% while the range of die-to-die variation is 300%.

4 Multiple Supply Port Analysis

A product design that incorporates large-area sleep island structures, and/or provides only limited control over the individual states of the sleep islands, may only allow a coarse level of resolution to be obtained of within-die leakage variations. For such designs, a better solution is to combine the sleep island approach with a multiple supply port (MSP) technique, that is capable of providing 'within-sleep-island' variation information [16]. Our proposed MSP technique obtains 'regional' leakage current information by measuring the leakage currents individually through each of the power ports (PPs). Since the power grids on nearly every product chip incorporate multiple PPs to mitigate IR and Ldi/dt voltage droop problems, our method simply leverages these existing architectural features. The set of leakage currents measured under MSP define a leakage current vector Q given by Eq. 1. Here, I_1 to I_N represent

$$Q = [I_1\ I_2\ \ldots\ I_N] \qquad \textbf{Eq. 1.}$$

the individual components of a chip's leakage current as measured through the MSPs.

It should be noted that access to the individual power ports is only possible during wafer probe, i.e., before the chips are packaged. In order to get around this limitation in our hardware experiments, we designed our test chips such that each of the power ports is connected to dedicated package pin. We then designed a test board that allowed the individual PP currents to be measured using an ammeter. Details of the test setup can be found in [18].

In previous work, we demonstrated that a normalized version of Q is well-correlated with the leakage characteristics measured through the PPs of our test chips [18]. In those experiments, logic states in the chips were manipulated to vary the leakage characteristics artificially, i.e., we carried out a vector-to-vector analysis of leakage current. In this work, we measure and report on *actual* within-die variations (as we did in the previous section) with the chips configured into a specific logic state. The goal of this work is to show that MSP is highly correlated to the sleep island results, and therefore, can serve as a high resolution alternative to the sleep island approach in cases where the sleep islands are too large to provide adequate resolution.

5 Calibrating MSP to Sleep Island

In order to determine the degree of correlation between the sleep island and MSP approaches, we focus this analysis on a subset of the leakage configurations described in Section 3, as shown in Figure 8(a). Each of the configurations labeled **p1** through **p4** place one region of the array in high leakage (HL) and the remaining regions in low leakage (LL). By measuring the global current in each of these configurations and subtracting a base leakage measurement, we were able to precisely measure the variations that occur in each of the HL islands. This is possible because of the high degree of control we have in the test structure. In actual product chips, such precise control over the leakage state of large number of localized devices is not possible.

An alternative approach is to measure the leakage currents through the MSPs on the chip, as shown in Figure 8(b)[1]. The power grid over the test structure on our chips is connected to an external power supply using 4 supply ports, shown as filled circles along the edge of the array and labeled PP_{00} through PP_{11}. The goal is to correlate the branch currents measured through the four supply ports, i.e., the current vector Q defined earlier, with the four global currents that we measure in each of the four sleep island experiments of Figure 8(a). In order to keep the leakage characteristics of the array the same in both scenarios, we configure the array with a 'high leakage everywhere' pattern for the MSP experiments. If the MSP approach correlates to the sleep island approach, then MSP can be used to obtain the same information as the 4-island approach described in Section 3.

Several issues associated with MSP need to be dealt with in order to achieve high correlation with the sleep island approach. First, from Figure 8(a) and (b), it is clear that MSP uses multiple measurements from 1 pattern to obtain leakage information that the sleep island approach obtains with 4 patterns. Since the power grid over the 4 leakage islands is continuous, the single pattern approach results in a 'bleeding' effect, in which the leakage currents from non-local islands 'bleed' into the local power port current measurements. For example, a portion of the leakage current drawn by the test structure in islands, R_{01}, R_{10} and R_{11} of Figure 8(b) will be sourced from PP_{00}. Other PPs will also be impacted by this bleeding effect. The mathematical model of this bleeding effect is given by Eq. 2, which defines a measured PP current, I_{xy}, as a fractional

1. Note that access to the individual power ports is only possible during wafer probe testing, and not after the dice have been packaged.

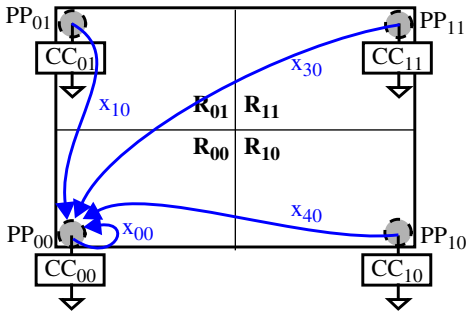

Fig. 9. Calibration circuit (CC) tests to obtain constants for estimating island leakage currents using MSP.

sum of the island leakage currents, L_{xy}, from islands R_{xy} in Figure 8(b).

$$\begin{aligned} I_{00} &= x_{00}L_{00} + x_{10}L_{01} + x_{20}L_{10} + x_{30}L_{11} \\ I_{01} &= x_{01}L_{00} + x_{11}L_{01} + x_{21}L_{10} + x_{31}L_{11} \\ I_{10} &= x_{02}L_{00} + x_{12}L_{01} + x_{22}L_{10} + x_{32}L_{11} \\ I_{11} &= x_{03}L_{00} + x_{13}L_{01} + x_{23}L_{10} + x_{33}L_{11} \end{aligned} \qquad \textbf{Eq. 2.}$$

The objective is to solve these equations for L_{xy}, the leakage current corresponding to each island. In order to do so, we need to determine the value of the constants x_{ab} which define the nature of the bleeding from the various islands to the PPs. One way to estimate these constants is to insert a special calibration circuit (CC) into the layout at positions underneath the PPs. In previous work, we proposed a calibration circuit similar to the TC shown on the right side of Figure 1 as a means of correcting for impedance variations in the power grid and probe card for application to manufacturing test [17]. This circuit can also be used here to derive an estimate of the constants x_{ab} in Eq 2. Figure 9 illustrates the process, where one copy of the TC shown in Figure 1 is placed underneath each PP (labeled CC in the figure). The calibration process involves creating a sequence of shorts between power and ground at these positions in the power grid and measuring the corresponding PP currents under each test. A short is created by setting the state of FF_1 and FF_2 in Figure 1 to '00', which enables both the NMOS and PMOS transistors in the pseudo-inverter.

The PP currents measured under these four calibration tests are then converted into the constants x_{ab} from Eq. 2. For example, the constants for the first equation of Eq. 2 are obtained by enabling each of the four CCs in Figure 9, one at a time, and measuring I_{00}, the current in PP_{00}. The four currents are converted to fractions by dividing them by the total current drawn by the corresponding CC pseudo-inverter. The total current is determined by summing the currents at all of the PPs under each CC test. The fractions effectively express the amount of current sourced through PP_{00} from these four points in the layout, and can serve as estimates for the x_{ab} constants directly. The x_{ab} constants for the other 3 equations in Eq. 2 are obtained in an analogous fashion. Once the matrix of 16 x_{ab} constants are computed, the leakage characteristics of each region, L_{xy}, are determined by inverting the matrix and multiplying it by the leakage current vector Q, represented by I_{xy} on the left side of Eq. 2, as given by Eq. 3.

$$[L] = [X]^{-1}[Q] \qquad \textbf{Eq. 3.}$$

Unfortunately, the island leakage estimates obtained from this procedure are not useful because of several other differences that exist between the sleep island and MSP approaches. In addition to the distortion introduced by 'bleeding', the MSP approach must also deal with current distribution distortions introduced by the power grid architecture. The power grid architecture does not affect the sleep island approach because it analyzes total (global) currents. Second, the 'point source' estimation of bleeding provided by the CC tests misrepresents what actually occurs from the distributed leakage sources in each island[1].

Both of these problems can be dealt with using a small set of simulation experiments, namely, a simulation of the 'leak everywhere' scenario depicted in Figure 8(b) and a set of simulations modeling the four CC tests depicted in Figure 9. The 'leak everywhere' simulation is carried out on a power grid model of the chip in which a set of uniform-valued current sources are distributed at regular intervals in regions that correspond to the core logic area(s) of the chip. This simulation enables the impact of the power grid architecture to be determined. The CC simulations enable a set of correction factors to be computed between 'point source' calibration and an ideal form of calibration that we will call 'leakage calibration', as explained below.

As we did with the calibration test data collected from the chips, we build a 2-D matrix of current fractions, i.e., PP currents divided by the total CC current, using the currents produced under the CC simulation experiments. This calibration matrix is inverted and used to transform the data measured under the 'leak everywhere' simulation experiment using a vector-matrix multiplication operation (Eq. 3). The calibrated 'leak everywhere' PP currents represent a nominal, zero-leakage-variation scenario.

Figure 10 shows the data from simulation experiments carried out on a power grid model of our chips. For the 'leak everywhere' simulation, we distributed 200 nA current sources uniformly across the layout. The total current introduced in each of the islands is 192.5 uA, labeled as LI_{xy} in the figure[2]. The currents drawn through the PPs

1. It is important to note that the 'point source' calibration serves an important additional role of significantly reducing current distortion caused by resistance variations in the power distribution system, including those introduced by process variations in the power grid metal wires.

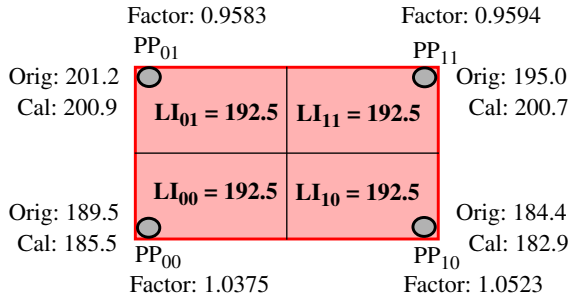

Fig. 10. Simulation derived 'leak everywhere' data before and after calibration with computed correction factors.

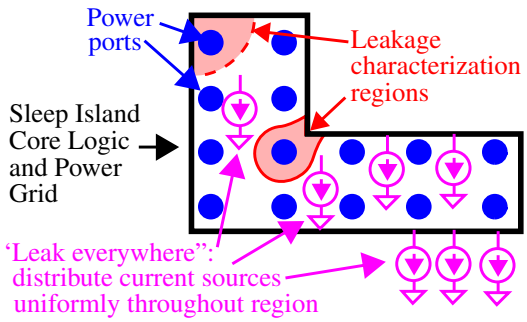

Fig. 11. Simulation experiment for correcting for misalignment in core logic and power ports.

under this simulation are labeled in the figure as "Orig:", and the calibrated values are labeled "Cal:". From the displayed values, it is evident that subtle variations in the power grid architecture 'distort' the individual sleep island current distribution to their respective PPs. For example, the calibrated PP_{00} current is 185.5 uA, which is 7 uAs smaller than the corresponding island value of 192.5 uA. The values labeled as "Factor:" in the figure give the constants needed to scale the PP values to make them equal to the island values. These factors, F_{xy}, are computed as the ratio of the island current, LI_{xy}, to the calibrated PP current, C_{xy}, as given by the left side of Eq. 4. Applying these factors to the calibrated chip data, L_{xy}, as given by Eq. 3, yields a 'corrected leakage' value for the chip, CL_{xy}. This process makes the MSP chip currents (CL_{xy}) nearly equivalent to the sleep island chip currents.

$$F_{xy} = \frac{LI_{xy}}{C_{xy}} \longrightarrow CL_{xy} = L_{xy} \times F_{xy} \qquad \textbf{Eq. 4.}$$

Although we developed this procedure using a regular, rectangular shaped sleep island architecture, it can be applied to any arbitrary shaped island in practice. The same process is carried out as shown for the "L-shaped" sleep island in Figure 11, i.e., calibration tests are simulated for each PP and a 'leak everywhere' simulation is run. For irregularly shaped sleep islands, the shape of the leakage characterization region around each power port may vary as shown by the two examples in Figure 11. This poses no problem because the objective of the technique is to analyze leakage variation across the sleep island and therefore, controlling the precise shape of the characterization region is not important. In fact, the leakage characterization regions in our rectangular shaped test structure are not discrete rectangles as shown in Figure 10, but rather are elliptical in shape. This shape mismatch in combination with measurement noise are the primary detractors to the correlation analysis of the MSP and sleep island methods presented in the next section.

2. The size of the simulated array is slightly smaller than the chip test structure and is 50x77 instead of 50X80.

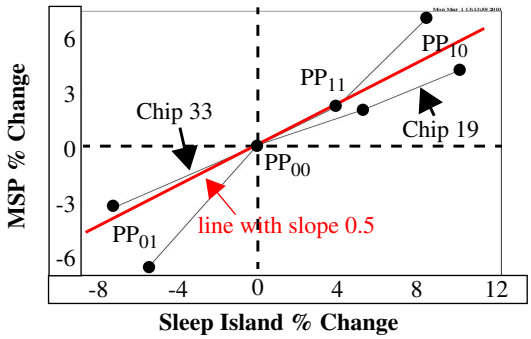

Fig. 12. Sleep Island vs. MSP % change for 2 chips. Correlation coefficients computed for each curve.

6 MSP and Sleep Island Correlation Analysis

Our correlation analysis focuses on determining the degree of correlation between the '**% change**' values computed under the sleep island and MSP analyses. Figure 12 plots the sleep island % change values (x-axis) against the MSP % change values (y-axis) for two chips that exhibit a high level of correlation (>90%). Each curve consists of 4 data points, one for each of the 4 islands/power ports. The % change values are computed with respect to R_{00} for the sleep island analysis and PP_{00} for the MSP analysis (see Figure 8(a) and (b)). Therefore, one of the data points in each curve is (0,0). The global currents are used as the sleep island data while the corrected PP currents, i.e., CL_{xy} from the previous section, are used as the MSP data. The % change values under each analysis are paired to preserve their spatial relationship, i.e., the % change value for sleep island R_{01} is paired with the % change for MSP PP_{01}, etc. The data points in each curve are sorted in ascending order according to their sleep-island values (x-axis) to enhance clarity in the figure. PP_{xy} labels are given in the figure to identify the correspondence of the data points to the islands/PPs shown in Figure 8(a) and (b).

The first important characteristic of the plot is the difference in the range of the sleep island values (approx. 20%) and the MSP values (approx. 12%). This range difference is reflected in the slope of the regression line[1] of 0.5, also shown in the figure. From this subset of data, there appears to be a factor of 2 reduction in sensitivity of the MSP approach in comparison to the sleep island approach.

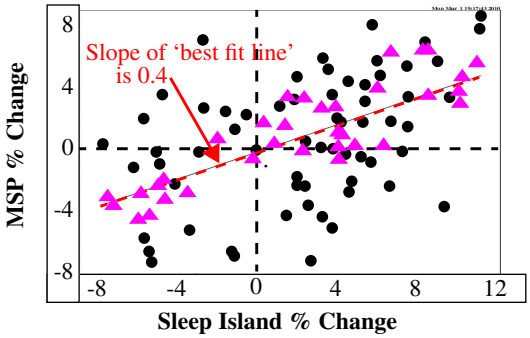

Fig. 13. Sleep Island vs. MSP % change analysis of 12 chips for scaling factor.

We suspect this is caused by the difference in the shape of the regions that are characterized under each of the methods (as mentioned earlier), and differences between the chips and the simulation model in the overall resistance characteristics of the power grid. The most important feature of the plot is the high degree of correlation that exists between the % change values from each analysis, which is reflected as the degree of 'straightness' in the curves. This suggests that MSP is capable of delivering similar results and is suitable as an alternative to the sleep island approach.

Figure 13 gives a scatter plot of the data points from the % change analysis for all 33 chips. The filled triangles correspond to the a 12-chip subset whose correlations (discussed below) are large, i.e., >+90%, while the filled circles correspond to chips with smaller correlations. The slope of the regression line using all data points is 0.4, but remains at 0.5 if the highly-correlated triangle-shaped data points are used.

We computed Pearson correlation coefficients (PCCs) for each of the % change curves of the 33 chips [19]. The range of PCC values is -100% (perfect negative correlation) through 0% (no correlation) to 100% (perfect positive correlation). A curve with data points that are perfectly aligned with positive slope receives a score of 100%, the best score possible. A bar graph depicting the PCCs from our analysis is given by Figure 14, with chip number plotted on the x-axis against the PCC on the y-axis. The chips are ordered according to the overall leakage value, with low leakage chips on the left and high leakage chips on the right.

The PCCs of 22 of the 33 chips are >70%, 17 chips are >80% and 12 chips are >90%. All chips to the right of chip #16 (high leakage chips) have PCCs >70%. The lower values of the PCCs for the chips on the left is attributed to a low signal-to-noise ratio. A combination of factors including the small size of the array (approx. 380 um by 560 um), the large series external resistances in series with the supply ports (approx. 3 Ohms) and the 0.5 uA noise floor of our setup, made it difficult to distinguish between the PP

1. The regression line is the best fit line through the data points from the curves.

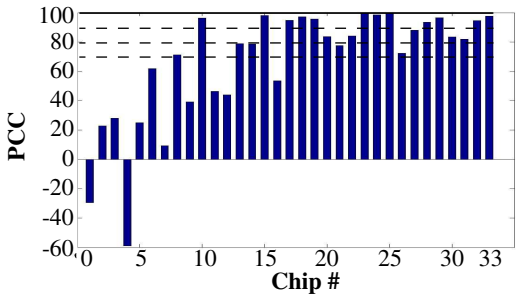

Fig. 14. Correlation coefficients associated with % change curves from sleep island analysis vs. MSP. Overall leakage for chips along x-axis increases from left to right.

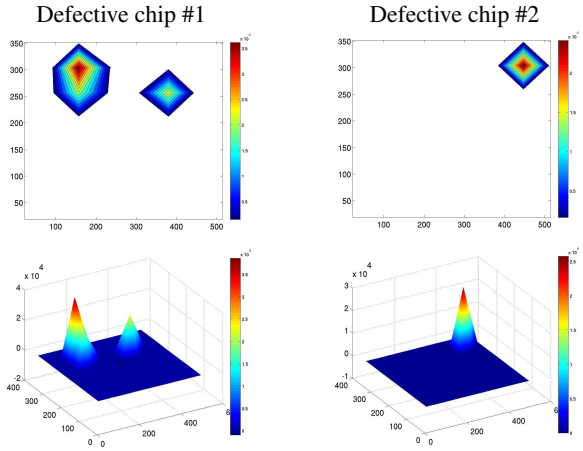

Fig. 15. 64-island within-die contour plots (top row) and 3-D surface plots (bottom row) of 2 chips with naturally occurring defects.

currents in the MSP experiment for these chips. For example, the measured differences in the PP currents for these chips are on order of a couple uAs. The higher leakage chips produce differences in the individual PP currents in the range of 5-10 uAs, which increased the stability and reproducibility of the measurements. Note that in product chips, the number and position of the PPs can be chosen such that the signal-to-noise ratio is much larger, simply by choosing PPs that have larger spatial separation. A trade-off needs to be considered here because widely separated PPs also reduces the measurable resolution of leakage current variations in the substrate.

7 Defect Analysis

Two additional chips that we analyzed were found to have naturally occurring defects. The defects are not provoked until a high leakage state is configured into the portion of the array that possess the defect. Figure 15 shows contour plots (top) and 3-D surface plots (bottom) from the 64-island experiments for the two chips. Defective chip #1 has multiple defects as is depicted by the multiple peaks. The current profiles of both chips show that the defect currents are significantly larger than the leakage current. The % change values shown in Table 1 under the sleep island and MSP techniques both indicate the presence of an anomaly (see bold entries). The point source nature of defects, as

opposed the distributed nature of leakage current variations, causes unique behavior in the % change metric that is easily identified using either the sleep island or MSP approach. Interestingly, for these defect types, the sleep island approach continues to provide valid data for the other regions because the defect is essentially de-activated when islands other than the defective island(s) are tested. Aside from applying these methods to product chips, these experiments demonstrate the utility of the test structure described in this work in providing valuable DFM and defect data for the purpose of yield learning and process bring-up.

	% Change			
Chip and Experiment	**PP_{00}**	**PP_{01}**	**PP_{10}**	**PP_{11}**
Def. chip #1 Sleep Island	0.00	**3716.13**	5.78	**1092.56**
Def. chip #1 MSP	0.00	**77.64**	-11.83	**25.76**
Def. chip #2 Sleep Island	0.00	1.24	6.13	**1576.07**
Def. chip #2 MSP	0.00	23.66	57.78	**264.38**

Table 1: % change values for defective chips under sleep island and MSP approaches.

8 Conclusions

We describe two approaches for measuring within-die and die-to-die leakage current variations in test chips and actual product chips, and demonstrate the methods on a set of test chips fabricated in IBM's 65 nm SOI process. The sleep island approach provides high resolution of within-die leakage variations, but only limited resolution on product chips because of the large area and/or limited control of typical sleep island structures. A multiple supply port technique is proposed to provide 'within-sleep-island' observability of leakage current variations. Experimental results from the fabricated test chips demonstrate that MSP correlates well with the sleep island data, and therefore can serve as an alternative for application to product chips.

References

[1] Y. Borodovsky, "Impact of Local Partial Coherence Variations on Exposure Tool Performance", *SPIE*, vol. 2440, pp. 750-770, 1995.

[2] A.K. Wong, R. Ferguson and S. Mansfield, "Mask Error Factor in Optical Lithography", *Trans. on Semi. Manuf.*, vol. 13, pp. 235â242, May 2000.

[3] D.G. Chesebro et al., "Overview of Gate Linewidth Control in the Manufacture of CMOS Logic Chips", *IBM J. of Res. and Dev.*, vol. 39, pp. 189-200, July 1995.

[4] J.-Y. Lai, N. Saka and J.-H. Chun, "Evolution of Copper-Oxide Damascene Structures in Chemical Mechanical Polishing", *J. of Electrochem. Soc.*, pp. G31-G40, 2002.

[5] C. Hedlund, H. Blom and S. Berg, "Microloading Effect in Reactive Ion Etching", *J. of Vacuum Science and Tech.*, vol. 12, pp. 1962â1965, 1994.

[6] Y. Lee, M. Sung, E. Lee, Y. Sohn, H. Bak and H. Oh, "Temperature Rising Effect of 193nm Chemically Amplified Resist during Post-Exposure Bake", *SPIE*, vol. 3999, pp. 1000-1008, 2000.

[7] C. Berger et al., "Critical Dimension Variations of I-line Processes due to Swing Effects", *SPIE*, vol. 6153, pp. 61523T, 2006.

[8] C. Hess et al., "Scribe Characterization Vehicle Test Chip for Ultra Fast Product Wafer Yield Monitoring", *Intl. Conf. on Microelectronic Test Structures*, pp. 110-115, 2006.

[9] K. Agarwal and S. Nassif, "Characterizing Process Variation in Nanometer CMOS", *Design Automation Conf.*, pp. 396-399, 2007.

[10] M. Ketchen and M. Bhushan, "Product-Representative At Speed Test Structures for CMOS Characterization", *IBM J. of Res. and Dev.*, vol. 50, pp. 451-468, July 2006.

[11] D. Boning et al., "Test Structures for Delay Variability", TAU Workshop, pp. 109, 2002.

[12] M. Bhushan, A. Gattiker, M. B. Ketchen and K. K. Das, "Ring Oscillators for CMOS Process Tuning and Variability Control", *Trans. on Semi. Manuf.*, vol. 19, pp. 10-18, Feb. 2006.

[13] L. Wang, W. Pope, L-T. Pang, A. Neureuther, E. Alon and B. Nikolic, "Hypersensitive Parameter-Identifying Ring Oscillators for Lithography Process Monitoring", *SPIE*, pp. 750-770, 2008.

[14] A. Gattiker, M. Bhushan, and M. Ketchen, "Data Analysis Techniques for CMOS Technology Characterization and Product Impact Assessment", *Intl. Test Conf.*, pp. 1-10, 2006.

[15] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi and V. De, "Parameter Variations and Impact on Circuits and Microarchitecture", *Design Automation Conf.*, pp. 338-342, 2003.

[16] J. Plusquellic, "IC Diagnosis Using Multiple Supply Port IDDQs", *IEEE Design & Test*, vol. 18, 2001, pp. 50-61.

[17] J. Plusquellic, D. Acharyya, A. Singh, M. Tehranipoor, C. Patel, "Quiescent-Signal Analysis: A Multiple Supply Pad I_{DDQ} Method", *Design and Test of Computers*, Volume 23, Issue 4, April 2006, pp. 278-293.

[18] K. Agarwal, D. Acharyya, J. Plusquellic, "Characterizing Within-Die Variation from Multiple Supply Port IDDQ Measurements, *Intl. Conf on Computer-Aided Design*, 2009, pp. 418-424.

[19] http://en.wikipedia.org/wiki/Pearson_correlation