Neighbor Current Ratio (NCR): A New Metric for I_{DDQ} Data Analysis

Sagar S. Sabade D. M. H. Walker
Department of Computer Science
Texas A&M University
College Station, TX 77843-3112
Phone: (979) 862-4387
Fax: (979) 847-8578
E-mail: {sagars, walker} @cs.tamu.edu

Abstract

I_{DDQ} test loses its effectiveness for deep sub-micron chips since it cannot distinguish between faulty and fault-free currents. The concept of current ratios, in which the ratio of maximum to minimum I_{DDQ} is used to screen faulty chips, has been previously proposed. At the wafer level neighboring chips have similar fault-free properties and are correlated. In this paper, use of spatial correlation in combination with current ratios is investigated. By differentiating chips based on their nonconformance to local I_{DDQ} variation, outliers are identified. The analysis of SEMATECH data is presented.

***Keywords**:* I_{DDQ} testing, current ratios, spatial correlation

1. Introduction

Testing deep sub-micron (DSM) chips is a complex task. I_{DDQ} testing loses its effectiveness as leakage current increases exponentially with shrinking transistor geometries [1]. The International Technology Roadmap for Semiconductors (ITRS) projection for leakage current is 8-20 A when DRAM half pitch is in the range of 32-22 nm by year 2014 [2]. It is already difficult to distinguish defective currents from the background leakage. Several solutions have been reported in the literature to solve this problem. These solutions can be classified under three different categories. The first category of solutions relies on circuit/technology changes to reduce the background leakage current. This includes techniques like reverse body bias [3], multiple threshold transistor design, etc. A second category of solutions relies on exploiting the dependence of I_{DDQ} on other parameters like temperature [4, 5]. The third category of solutions uses various means to reduce fault-free I_{DDQ} variation by data analysis so as to make faulty I_{DDQ} values distinguishable. This involves use of statistical means [6, 7, 8], graphical display of data [9] or use of correlation between I_{DDQ} and other parameters [10, 11, 12]. Due to increasing spread in I_{DDQ} with each technology node, it is difficult to spot outliers in the data set [13]. The current ratios approach has shown some promising results [7]. In this work, we evaluate the capabilities of a combination of the current ratios technique and wafer-level spatial correlation.

This paper is organized as follows. In the next section we review the current ratios concept and describe the motivation behind this study. In Section 3 we describe our analysis methodology. Section 4 describes the experimental results for SEMATECH data[1]. Section 5 presents discussion of the results and Section 6 concludes the paper.

2. Current Ratios Concept

The main concern about I_{DDQ} testing is increasing background current that makes it impossible to distinguish between faulty and fault-free values. This is worsened due to

[1] This data comes from the Test thrust at SEMATECH, Project S-121 on Test Methods Evaluation. The conclusions drawn are our own and do not necessarily represent views of SEMATECH or its member companies.

increased process variations that cause wide variation in fault-free I_{DDQ}. A fault-free chip that leaks more should consume high current for *all* input patterns[2]. On the other hand, a chip having an active defect consumes high current only when the defect is excited. In this case the leakage current depends on the nature and resistance of the defect, among other parameters. Gattiker and Maly showed that the general shape of sorted I_{DDQ} readings (called a *current signature*) can reveal whether a chip is faulty or not [9]. The number of measurements limits the resolution of this approach. Due to the slow speed of I_{DDQ} test and cost concerns, it is usually not feasible to take many I_{DDQ} measurements in production.

Maxwell et al observed that in spite of an order of magnitude difference in I_{DDQ} values, two dice had similar signatures, as shown in Figure 1 [7]. It was proposed that ratios of maximum I_{DDQ} to minimum I_{DDQ} for fault-free chips would have small variation and can be used as a pass/fail criterion. The signature was described using an equation of the form:

$$Max\ I_{DDQ} = Slope \bullet Min\ I_{DDQ} + Intercept$$

Thus, the maximum I_{DDQ} was predicted based on the minimum I_{DDQ} value. The authors characterized a sample of chips and determined current ratios for several fault-free chips. Through linear regression they determined the parameters (the slope and the intercept) of the equation shown above. To account for unmodeled process variations, a guard band was added. In production, I_{DDQ} was measured for the minimum I_{DDQ} vector and the pass/fail limit on the maximum I_{DDQ} for all other vectors was set.

To examine the applicability of this technique to the SEMATECH data we computed current ratios for SEMATECH test chips that passed all wafer tests (11263 chips termed “all pass”), or at wafer level failed only I_{DDQ} test (1689 chips termed “I_{DDQ}-only fail”) [14]. For each chip 195 I_{DDQ} readings are available. These readings were sorted to obtain the current ratio of maximum I_{DDQ} to minimum I_{DDQ}. Figure 2 shows the current ratios for the all pass and I_{DDQ} only chips sorted in ascending order. Since the SEMATECH experiment used an I_{DDQ} pass/fail limit of 5 μA, the majority of chips that passed all tests naturally have small ratios (<3). Some I_{DDQ}-only failed chips (788) have current ratios less than 10, comparable to those of the all-pass chips. Many of these chips either pass all tests or fail only I_{DDQ} test after burn-in. Clearly not all of these chips are so flawed to be rejected. This indicates that the SEMATECH I_{DDQ} test limit of 5 μA resulted in considerable yield loss. Figure 3 shows the distribution of the current ratios for chips that passed all tests or failed only I_{DDQ} test (chip from the burn-in sample are shown). The current ratio bin size is 0.5 and some bins are offset for clarity. It can be seen that the number of I_{DDQ}-only failed chips in most of the bins is proportionately comparable to the number of all pass chips in the corresponding or the closest bin.

Spatial Correlation

The total leakage current contains two components: an intrinsic leakage component and a defective leakage component. The minimum I_{DDQ} is mostly due to intrinsic leakage (assuming that at least one vector does not excite the defect or the chip is fault-free). As neighboring chips on a wafer undergo similar processing, their fault-free device parameters are correlated [11]. Therefore, neighboring chips are expected to have similar current ratios. To verify this assumption, we obtained current ratios of all-pass and I_{DDQ}-only failed dice on a wafer. A surface plot with the projection of its contour is shown in Figure 4. Some of the contours are marked with the corresponding ratios. The dice at XY locations (8,10) and (12,10) failed only I_{DDQ} test at the wafer level and had maximum

[2] The converse is not true. A chip with a passive defect (e.g. V_{DD} to ground resistive short) also has elevated leakage for all input patterns.

leakage currents of 67 μA and 61 μA, respectively. It can be observed from Figure 4 that current ratios across the wafer vary relatively smoothly considering the logarithmic scale on the Z-axis. The *spatial outlier* chips are likely to be defective and are at greater risk of failing during burn-in. Unfortunately these two chips were not subjected to burn-in, so post burn-in results are not available.

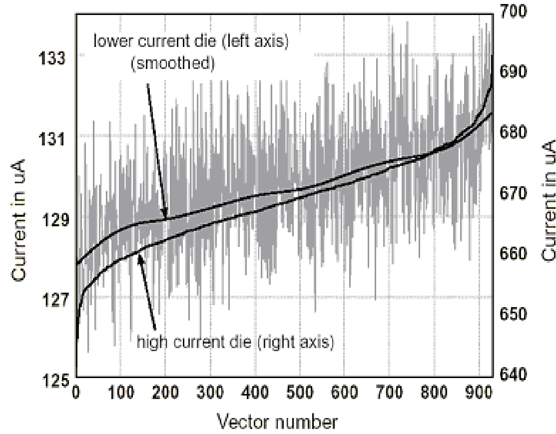

Figure 1: Current signatures of two dice [7]

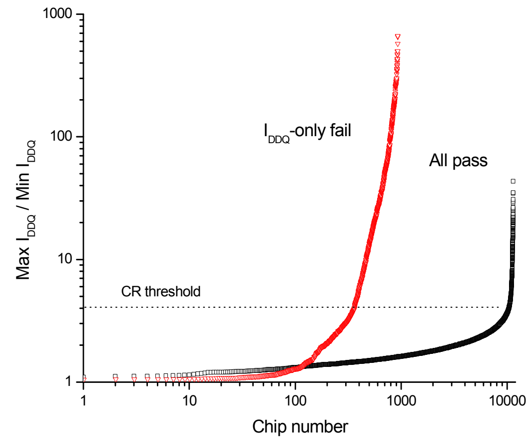

Figure 2: Current ratios for SEMATECH chips

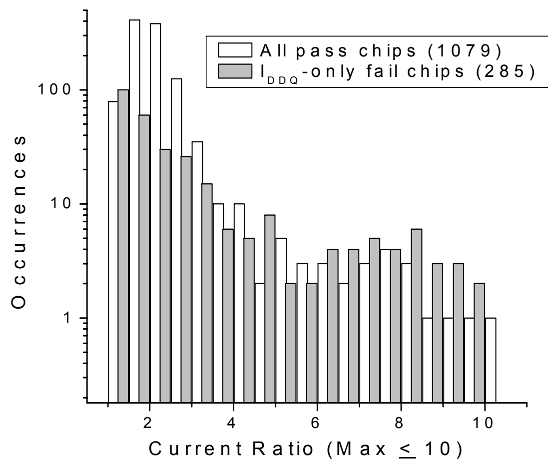

Figure 3: Histogram of current ratios for all-pass and I_{DDQ}-only fail chips (max current ratio = 10)

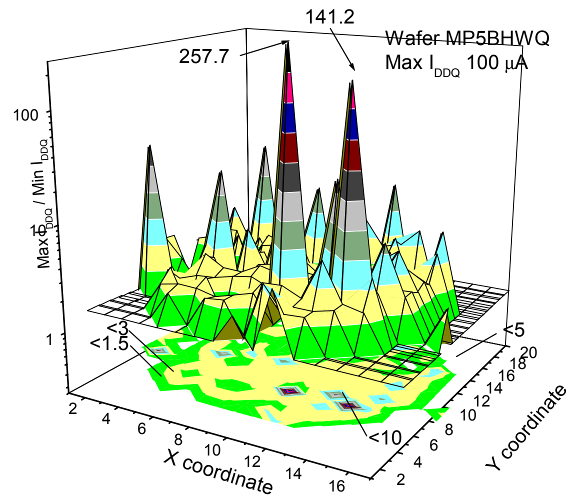

Figure 4: Current ratio surface and contour plot

Neighbor Current Ratios (NCR)

If we obtain ratios of I_{DDQ} readings of a die and its neighboring dice for each vector, they should exhibit a small variation. To distinguish them from current ratios, we denote these ratios as *neighbor current ratios* (NCR). Ideally, NCRs would all be equal to 1. Let I_1, $I_2 \ldots I_n$ denote I_{DDQ} readings for a fault-free die (die A) and I'_1, $I'_2 \ldots I'_n$ denote I_{DDQ} readings for a fault-free neighboring die (die B). Under ideal conditions (no process variations) we obtain:

$$NCR = \frac{I_1}{I'_1} = \frac{I_2}{I'_2} = \ldots = \frac{I_n}{I'_n} = 1$$

Due to process variations NCRs would exhibit small variation around the mean value of unity. If both the chips are fault-free, but one leaks more than the other, all NCRs

would be consistently more or less than unity. In this case, the mean value would be different but the variance would be small. Assuming process variations do not cause as much variation as caused by defects, different cases arise as shown in Table 1.

When either chip contains a passive defect, since defective leakage components are not known, the mean NCR cannot be predicted. In this case, however, NCRs would have small variance.

Figure 5 shows the histograms of NCRs for five different cases. In each case two neighboring dice from a SEMATECH wafer are used. Figure 5(a) corresponds to the case where both dice have active defects. Two groups of NCRs clustered near 1 and 4 are noticeable. When the input pattern excites the defect in one chip (but not the other), high NCR (~4) is obtained. The other cluster (near 1) occurs when both defects are excited and the chips have similar currents or when neither defect is excited. Figure 5(b) depicts the case where both dice have passive defects. Although the mean value is near 1, there is more variation than in the fault-free case. Understandably, it is difficult to distinguish it from the case when both the chips are fault-free, as shown in Figure 5(e). Figure 5(c) illustrates a case when one die has a passive defect and another die has an active defect. The mean and standard deviation depend on which and how many vectors excite the defects as well as on the nature of defects. Figure 5(d) shows the case when one die is leakier than the other. In this case all NCRs are less than 1 and the histogram has a very small standard deviation.

Table 1: Different cases for NCR (die A I_{DDQ} /die B I_{DDQ})

Die A	Die B	NCR[§]
Fault-free	Fault-free	~1
Fault-free[3]	Passive Defect	<1
Fault-free	Active Defect	<1
Passive Defect	Fault-free	>1
Active Defect	Fault-free	>1
Active Defect	Passive Defect	?
Passive Defect	Active Defect	?
Passive Defect	Passive Defect	?
Active Defect	Active Defect	?

3. Analysis Methodology

In this section we describe the analysis methodology. The analysis is performed using the wafer-level information from SEMATECH data. We consider four adjacent dice in N, S, W, E directions and 4 dice in SE, SW, NW, NE directions as the neighboring dice for the center die. Chips that fail other than I_{DDQ} test are not considered. Also, chips with very high I_{DDQ} are not considered in the analysis since such chips are definitely defective. NCRs are computed for each vector by considering the center die and each neighboring dice as $NCR_{ij} = \frac{I_{cj}}{I_{ij}}$, where I_{cj} is I_{DDQ} for the center die for j^{th} vector and I_{ij} is I_{DDQ} of the i^{th} neighbor ($1 \le i \le 8$) for the j^{th} vector. For each chip we obtain the maximum NCR for each available neighbor. The maximum of all NCR values is used for the pass/fail criterion (referred to as the NCR for the chip). The NCR value essentially indicates the degree of conformity of a chip to its neighbors. Thus, a chip having very high NCR is more likely to contain a defect and fail than a chip having NCR close to 1. Note that even the chips on a leaky wafer will show conformance with their neighboring chips. This

---

[§] Cases where NCR is dependent on the nature of defect and cannot be predicted are indicated by a question mark.

[3] For a die with a subtle defect the same relation also holds true.

would reduce the yield loss caused by a single pass/fail threshold. NCR-based rejection is expected to reduce the yield loss by rejecting only gross spatial outliers.

Figure 6 shows the NCR wafer surface plot for the same wafer shown in Figure 4. Notice that several chips having small current ratios in Figure 4 (likely passive defects) exhibit much higher maximum NCR in Figure 6. Thus use of NCR can improve the confidence in outlier detection. Figure 7 shows maximum NCRs for all chips that passed all wafer level tests or failed only I_{DDQ} test. All pass chips show a long tail (outlier chips) indicating not all chips are fault-free. On the other hand, many I_{DDQ}-only failed chips have NCR values (<10) comparable to that of all pass chips.

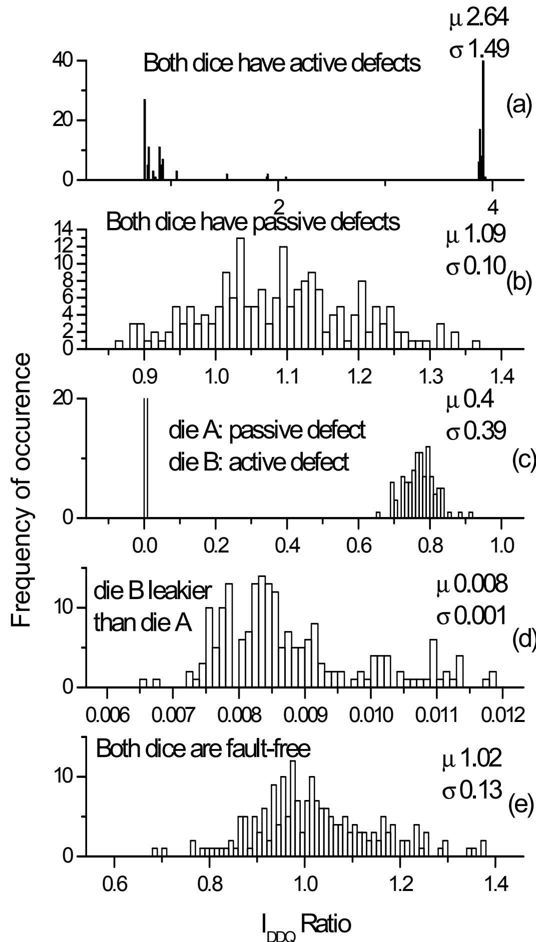

Figure 5: Histograms of NCRs for different cases

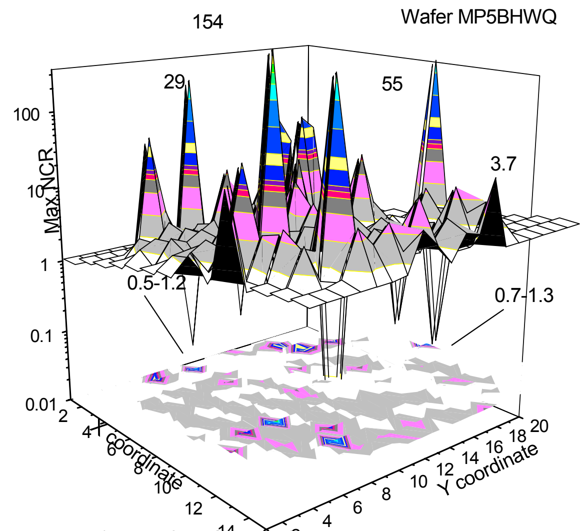

Figure 6: NCR surface and contour plot

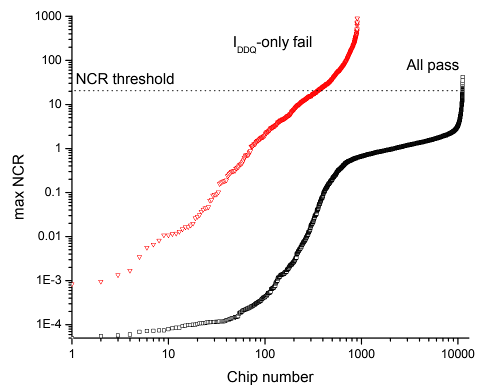

Figure 7: NCRs for SEMATECH chips

It is difficult, if not impossible, to distinguish NCRs of two fault-free chips and two chips with passive defects. However, it is unlikely that all the neighbors have identical passive defects with similar currents. One possible solution is to compare the difference between the minimum and maximum I_{DDQ} with the minimum I_{DDQ}. For a chip with a passive defect, the difference between the minimum and maximum I_{DDQ} values will be relatively small. The vector-to-vector within chip variation of I_{DDQ} for a fault-free chip is muffled by the I_{DDQ} due to a passive defect. Such chips would exhibit small current ratios.

4. Experimental Results

To evaluate the effectiveness of NCR in screening wafer-level spatial outliers, we applied this method to SEMATECH data. We considered only those chips that passed all tests (1102) or failed only I_{DDQ} test (1558) at the wafer level and underwent six hours of burn-in. If any I_{DDQ} reading exceeded 100 μA the chip was assumed to contain a gross defect [16]. Therefore all chips for which I_{DDQ} exceeded 100 μA were ignored from the analysis. This reduced the dataset from 1558 I_{DDQ}-only fails to 858 I_{DDQ}-only failed chips. The distribution of the entire dataset according to the number of available neighbors is shown in Table 2. Even though the BI sample was not random, it is interesting to note that a large number of dice that passed all wafer and post-BI tests or failed only I_{DDQ} test at both levels have 3 or more neighbors. Although NCRs can be obtained even if a single neighboring die is available, confidence in prediction is improved if multiple dice are present.

Table 2: Distribution of dice in the original dataset

Wafer Probe	Post-BI	Number of available neighbors								
		0	1	2	3	4	5	6	7	8
All- pass	All pass	4	16	39	82	120	185	233	225	152
	I_{DDQ}-fail	0	0	3	3	4	2	7	7	1
	Other	0	1	1	0	3	1	6	5	2
I_{DDQ}-fail	All pass	7	9	27	47	48	37	25	23	9
	I_{DDQ}-fail	7	19	51	99	95	105	107	82	40
	Other	1	0	1	5	4	4	5	0	1

Since we considered only immediate neighboring die positions, the dice having zero adjacent neighbors (19) could not be considered for analysis. Thus the total dice in the dataset were reduced to 1941 from the 1960 chips shown in Table 2 (1098 all pass, 843 I_{DDQ} fail). For each available neighbor a total of 195 ratios were obtained and the maximum NCR value (across all neighbors) was used for pass/fail decision.

To compare the effectiveness of NCR with current ratios, we considered the same dataset and used current ratios for the pass/fail decision. An important criterion was to select threshold values for a fair comparison. It can be observed from Figure 2 that chips that passed all tests but have a current ratio of more than 4 seem to be "outliers". We therefore considered a current ratio of 4 to be an appropriate pass/fail threshold. Then the NCR threshold was adjusted so as to match the same defect level obtained by the current ratio method. While computing the defect level, all post-BI I_{DDQ}-only failed chips were considered fault-free. For the same defect level as obtained by the CR threshold of 4, the NCR threshold value was 21.

Table 3 shows the distribution of chips in different categories. All chips are divided into two main categories: Chips accepted by current ratio (CR accept) and chips rejected by current ratio (CR reject). They are further subdivided into two categories: chips accepted by neighbor current ratio (NCR accept) and those rejected by neighbor current ratio (NCR reject). These four categories are divided depending on their wafer probe result: (a) chips that passed all SEMATECH tests ("All pass") and (b) chips that failed the 5 μA threshold I_{DDQ} test but passed other tests ("I_{DDQ}-only fail"). Each category is subdivided based on post burn-in SEMATECH test result. This distinction is made to understand the distribution of NCRs and CRs in different categories and understand if certain chips get detected by one method but not by the other and rejection rate of healer chips. Since the 5 μA test limit does not represent a "good" manufacturing limit [14], the

difference in I_{DDQ}-only failed chips detected by one method but not the other would be statistically significant to draw meaningful conclusions.

The overkill and defect level are computed as follows:

$$Overkill = \frac{Number\ of\ chips\ that\ pass\ after\ BI}{Total\ number\ of\ chips\ rejected} * 100$$

$$Defect\ Level = \frac{Number\ of\ chips\ that\ fail\ after\ BI}{Total\ number\ of\ chips\ accepted} * 100$$

These values are scaled appropriately considering the entire population that was not burned in. Table 4 shows overkill and defect level (DL) values for both methods and their combination. In the combined method a chip is rejected if it is considered faulty by either method. Since I_{DDQ}-only fail chips are not conclusively defective, overkill and DL are computed by considering all such chips fault-free and then by considering all such chips faulty. This is indicated in the second row. The columns headed "Good" ("Faulty") have overkill or defect levels values computed by considering all I_{DDQ}-only failed chips fault-free (faulty). The actual values would lie between these two extremes.

Table 3: Distribution of chips for different test methods

SEMATECH Wafer Probe Test Result	CR Accept		CR Reject		Post 6-hour BI SEMATECH Test Result
	NCR accept (1153)	NCR reject (149)	NCR accept (280)	NCR reject (359)	
All pass	949	0	100	3	All pass
	19	0	7	1	I_{DDQ} fail
	17	0	2	0	Other
I_{DDQ}-fail	69	26	76	54	All pass
	94	119	90	295	I_{DDQ} fail
	5	4	5	6	Other

Table 4: Overkill and DL for different test methods

Metric	Overkill %		Defect Level %		Effective Yield %
I_{DDQ} fail	Good	Faulty	Good	Faulty	
CR	97.97	36.46	2.00	19.82	67.08
NCR	98.03	16.34	2.02	16.68	77.83
CR+NCR	97.84	32.87	1.91	11.71	59.40

5. Discussion

Table 3 reveals many interesting findings. Since both NCR and CR thresholds are high, only 3 of the all pass chips at wafer probe are rejected by both methods. The NCR test rejects fewer chips that pass all wafer tests than the CR test. The SEMATECH data has many healer chips that have reduced I_{DDQ} after burn-in (and thus pass all tests). Since NCR threshold is very loose it accepts more healer chips than CR test. Since healers represent unstable or unreliable chips, they would typically get rejected up front in the test flow and will not be subjected to burn-in[4]. In practice, NCR threshold should be selected by observing wafer-level variation in I_{DDQ}. NCR test rejects more chips that fail wafer level and post burn-in I_{DDQ} test than the CR test. Considering NCR threshold of the 21, these chips are more likely to be defective.

Figure 8 shows the distribution of post-BI failures of chips according to their maximum NCR. The healer chips are shown separately. As expected, the bins with maximum NCR

---

[4] SEMATECH project S-121 was a research study where the effect of burn-in on I_{DDQ} was studied. Hence failed chips were also subjected to burn-in.

values less than 2.5 have a high percentage of fault-free chips. As expected the bins for high NCR values have higher failure rates. Since the BI sample was non-uniform, some higher NCR bins have very few chips resulting in low failure rates. But for all practical purposes, it is safe to assume that NCR values higher than 10 would have high failure rate. The bins with NCR values much less than 1 represent good dice in bad neighborhood (*spatial dips*). Previous work has shown that the probability of failure of such dice is high [15]. Several chips having maximum mean NCR less than 1 pass all post burn-in tests. Some of these chips are more likely to contain subtle defects and fail sooner.

The bins with NCR values greater than 10 essentially represent bad dice in good neighborhoods (*spatial peaks*). Such spatial outliers can be easily identified by the NCR test method. As Figure 8 indicates the probability that such chips will pass burn-in falls with higher NCR values. Many of these chips are healers and hence unreliable.

NCR test seems to rejects more I_{DDQ}-only failed chips than the CR test. This gives a yield penalty for a fast wafer region. Another metric like flush delay can be combined with NCR values to reduce the yield loss [16]. Even if a single die in the neighborhood yields high NCR for the center die, the center die should be regarded as defective. If this causes unacceptable yield loss, such dice could be selectively burned in.

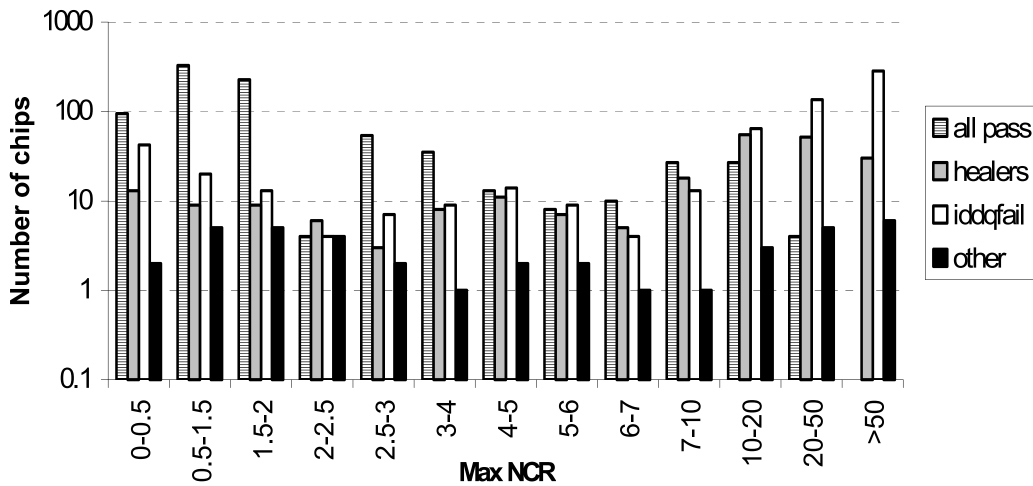

Figure 8: Distribution of post-BI results of chips for different maximum NCR values

Advantages of NCR

NCR is an intuitively simple metric. It can be used to screen high leakage current chips surrounded by good chips as well as detect good chips in a bad neighborhood (*spatial dips*). Only a single good neighboring die is necessary for NCR test. This is generally not a problem except in a poor yield zone or on the wafer edge. Since NCR is self-calibrating and self-scaling it provides an easy way to use spatial information for I_{DDQ} pass/fail decisions for any technology node.

Limitations of NCR

The basic assumption of NCR is that neighboring chips have I_{DDQ} values that are correlated. If the neighboring chips are missing, the NCRs cannot be determined. This is true for dice on the wafer edge and in a poor yield zone. However, this can be resolved by considering dice at longer distances [10]. To account for lot and wafer level variations, it is helpful to find the best predictors for each die position. Dice on the wafer edge could be correlated to dice on the edge either on the same wafer or on another wafer [17].

6. Conclusions and Future Work

This paper has shown that comparing the I_{DDQ} of a die with that of its neighboring chips and observing the variation trend can be useful in spotting local spatial outliers

caused by a process glitch, spot defects, etc. For the same defect level, the NCR-based test results in higher yield and lower overkill compared to current ratios. This test method assumes that wafer level variation in fault-free parameters is smooth, which is usually the case. If stepper field patterns are observed in the data, the NCR-based test method alone would cause more yield loss. In this case, it is necessary to identify the best predictors for each die position similar to study reported in [18].

It would be interesting to see if combination with other parameters or other test methods such as delta-I_{DDQ} improves the results. For dice having no adjacent neighbors, we will consider dice at longer distances. The results obtained in this direction will be reported in future. Also similar to the study reported in [17], correlating dice from other wafers for same XY locations could be useful.

Acknowledgements

This research was funded in part by the National Science Foundation under grant CCR-9971102 and Texas Advanced Research Program (ARP) 2001 under grant 512-186-2001. Thanks to Phil Nigh of IBM for providing the SEMATECH data and Peter Maxwell of Agilent Technologies for allowing reproduction of Figure 1.

References

[1] J. Figueras and A. Ferre, "Possibilities and Limitations of I_{DDQ} Testing in Submicron CMOS," *IEEE Trans. on Components, Packaging, and Manu. Technology*, part B, Vol. 21, No. 4, Nov. 1998, pp. 352-359.

[2] *International Technology Roadmap for Semiconductors (ITRS),* Semiconductor Industry Association, 2001, available online at http://public.itrs.net.

[3] A. Keshavarzi et al., "Effectiveness of Reverse Body Bias for Lowe Power CMOS Circuits," *8th NASA Symp. VLSI Design*, IEEE Press, Piscataway, NJ 1999, pp. 2.3.1-2.3.9.

[4] S. Kundu, "I_{DDQ} Defect Detection in Deep Submicron CMOS ICs," *Asian Test Symp.*, 1998, pp. 150-152.

[5] V. Szekely et al., "Cooling as a Possible Way to Extend the Usability of I_{DDQ} Testing," *Electronics Letters*, vol. 33, no. 6, Dec. 1997, pp. 2117-2118.

[6] C. Thibeault, "An histogram Based Procedure for Current Testing of Active Defects," *Intl. Test Conf.*, 1999, Atlantic City, NJ, pp. 714-723

[7] P. Maxwell et al., "Current Ratios: A Self-scaling Technique for Production I_{DDQ} Testing", *Intl. Test Conf.,* Atlantic City, NJ, 1999, pp. 738-746.

[8] S. Jandhyala et al., "Clustering Based Techniques for I_{DDQ} Testing", *Intl. Test Conf.,* 1999, Atlantic City, NJ, pp. 730-737.

[9] A. Gattiker and W. Maly, "Current Signatures: Application", *Intl. Test Conf.,* Washington D.C., October 1997, pp. 156-165.

[10] W. R. Daasch et al. "Variance Reduction Using Wafer Patterns in I_{DDQ} Data", *Intl. Test Conf.,* 2000, Atlantic City, NJ, pp. 189-198.

[11] S. Sabade and D. M. H. Walker, "Improved Wafer-level Spatial Analysis for I_{DDQ} Limit Setting," *Intl. Test Conf.,* 2001, Baltimore, MD, pp. 82-91.

[12] A. Keshavarzi et al., "Multiple-Parameter CMOS IC Testing with Increased Sensitivity for I_{DDQ}," *Intl. Test Conf.,* Atlantic City, NJ, 2000, pp. 1051-1059.

[13] C. Hawkins and J. Soden, "Deep Submicron CMOS Current IC Testing: Is There a Future?," *IEEE Design and Test of Computers*, Oct.-Dec. 1999, pp. 14-15.

[14] P. Nigh et al., "So What is an Optimal Test Mix? A Discussion of the SEMATECH Methods Experiment" *Intl. Test Conf.*, Washington D.C., October 1997, pp. 1037-1038.

[15] A. D. Singh et al., "Screening for Known Good Die Based on Defect Clustering: An Experimental Study", *Intl. Test Conf.,* Washington D.C., October 1997, pp. 362-369.

[16] S. Sabade and D. M. H. Walker, "Wafer-level Spatial and Flush Delay Correlation Analysis for I_{DDQ} Estimation," *IEEE Intl. Workshop on Defect Based Testing*, Monterey, 2002, pp. 47-52.

[17] R. B. Miller and W. C. Riordan, "Unit Level Predicted Yield: a Method of Identifying High Defect Density Die at Wafer Sort," *Intl. Test Conf.*, Baltimore, MD, 2001, pp. 1118-1127.

[18] R. Daasch et al., "Neighbor Selection for Variance Reduction in I_{DDQ} and Other Parametric Data," *Intl. Test Conference,* 2001, Baltimore, MD, pp. 92-100.