I_{DDQ} Test: Will It Survive the DSM Challenge?

Sagar S. Sabade and D.M.H. Walker
Texas A&M University

Deep-submicron technologies pose difficult challenges for I_{DDQ} testing in the future. The low threshold voltage used by DSM devices decreases the defect resolution of I_{DDQ}. However, because I_{DDQ} is a valuable test method, researchers are working to augment I_{DDQ} with other test parameters to prolong its effectiveness.

THE COMPLEX TASK of ensuring the correct operation of ICs becomes more challenging as design complexity increases. Methods of testing ICs fall into two broad categories, operational and defect-based. These tests complement each other. Operational tests verify the chip's functionality—for example, functional test and at-speed test. Defect-based tests target the physical defects using their abstract representation (fault). The absence of a defect passes the test. Such tests include

- structural or stuck-at (scan) test,
- structural delay test (ac scan),
- leakage current (I_{DDQ}) test, and
- very-low-voltage test.

Functional test applies predetermined patterns called test vectors at the IC's inputs and compares the output with the expected pattern. A scan-based testing method is a structural test that checks combinational logic, flip-flops or latches, and connectivity, by placing the device in a logic state and changing it by shifting patterns through the flip-flops when they are configured into shift registers (scan chains). Detection of a mismatch indicates the presence of a defect in the device. The at-speed test applies patterns at the device's rated frequency to verify it can operate at the required speed. The ac scan delay test uses scan chains to deliver vector pairs that verify timing behavior on the specific paths (path delay), or on identified gates or connections that are represented as faults (transition delay).

Leakage current—or I_{DDQ} (direct drain quiescent current), as it is popularly known—test is a defect-based test that measures device supply current under steady-state conditions. Fully static CMOS circuits consume little power when their inputs are stable, because there is no direct path from the V_{DD} supply rail to ground. Hence, an IC that draws a large amount of current when inputs are stable is likely to be defective. This is the basic philosophy behind I_{DDQ} testing. Thus, I_{DDQ} test can detect shorts (bridges) between two signals or between signal and power supply lines (both categories called *active* or *pattern-dependent defects*), or between V_{DD} and ground (called *pattern-independent* or *passive defects*).

The inverter circuit in Figure 1a illustrates the basic philosophy of I_{DDQ} testing. A stable input and the absence of defects results in the low quiescent current flowing from V_{DD} to ground (because there is no direct path from V_{DD} (to ground), as shown in Figure 1b. Defects like the source-drain short shown in Figure 1a, however, would have significant current flow

0740-7475/02/$17.00 © 2002 IEEE

through the transistors. Thus, by measuring the elevated leakage current, we can identify a defective chip. Several shorts in the circuit (V_{DD} to ground, gate to source, and so forth) would also cause elevated I_{DDQ}.

The I_{DDQ} test differs from the functional test because no Boolean pass or fail decision occurs. A chip consuming very high current can still function correctly. Therefore, semiconductor manufacturers face the dilemma of whether to reject a chip failing only the I_{DDQ} test or to ship it to market with some defect that was not detected by other tests. Such a chip may operate correctly or may result in a customer return.

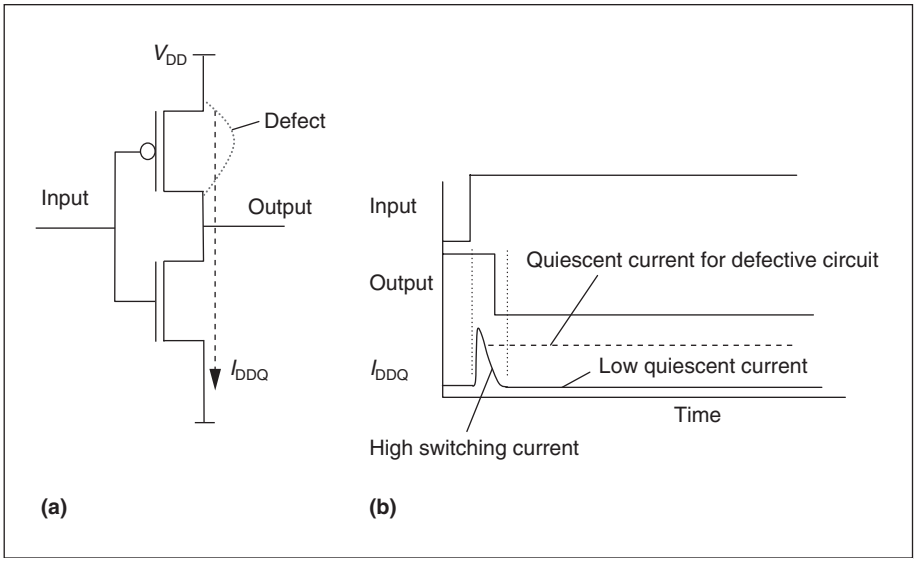

Figure 1. Static leakage current through a fault-free inverter is low (a); the presence of a defect increases it considerably (b).

I_{DDQ} testing offers several advantages. Defect detection requires an activated defect and a noticeable effect of its presence. Unlike stuck-at test, I_{DDQ} uses power supply lines for observation so there are no propagation requirements. Thus, it offers 100% observability. Generally, achieving reasonably high fault coverage takes only a few vectors. I_{DDQ} detects weak ICs or those with latent defects. Thus, I_{DDQ} test can screen chips for high-reliability applications and could possibly replace or reduce burn-in—a capability that will become more important as burn-in loses its effectiveness in accelerating latent defects for deep-submicron technologies.

In this article, we describe the challenges to I_{DDQ} testing posed by DSM technologies and discuss several I_{DDQ} test methods that aim at reducing variance in fault-free leakage current using various correlation and data analysis methods.

Basics of I_{DDQ} testing

Traditional I_{DDQ} testing follows a simple approach called single or static threshold method. Any chip with I_{DDQ} test measurements exceeding a certain threshold value—determined either by circuit simulation or empirically—is considered defective. For reasons explained later, this method will not work for new and emerging technologies.

Moore's law predicts that chip complexity will double every 18 months, and test complexity typically increases at least as fast. Thus, testing DSM chips containing millions of transistors becomes extremely difficult.

As designers reduce transistor geometries, they must decrease the supply voltage to avoid electrical breakdown. To retain or improve performance, however, it is necessary to reduce the threshold voltage (V_{TH}) as well. The subthreshold leakage current (current flowing through the transistor when gate-to-source voltage is less than the threshold voltage) is given by

$$I_{sub} = \mu C_{ox} \frac{W}{L} V_t^2 e^{(V_{GS}-V_{TH})/\eta V_t}\left(1 - e^{-V_{DS}/V_t}\right)$$

where μ is the carrier mobility, C_{OX} is the gate capacitance per unit area, W is the channel width, L is the channel length, V_{GS} is the gate-to-source voltage, V_t is the thermal voltage, V_{TH} is the threshold voltage, and η is the technology-dependent parameter. Thus, reducing threshold voltage V_{TH} causes an exponential increase in the subthreshold leakage current.

Because of an increasing number of transistors, mixed-signal designs, and reduced threshold voltages, leakage current levels are rising with each technology node. Table 1 (next page) shows the projections of the 2001 *International Technology Roadmap for Semiconductors* (*ITRS*)(http://public.itrs.net) for fault-free leakage currents at 25 °C for future

Table 1. *ITRS* projections for leakage currents of high-performance ICs.

Year	Maximum I_{DDQ} (mA)
2003	30 to 70
2005	70 to 150
2008	150 to 400
2011	400 to 1,600
2014	8,000 to 20,000

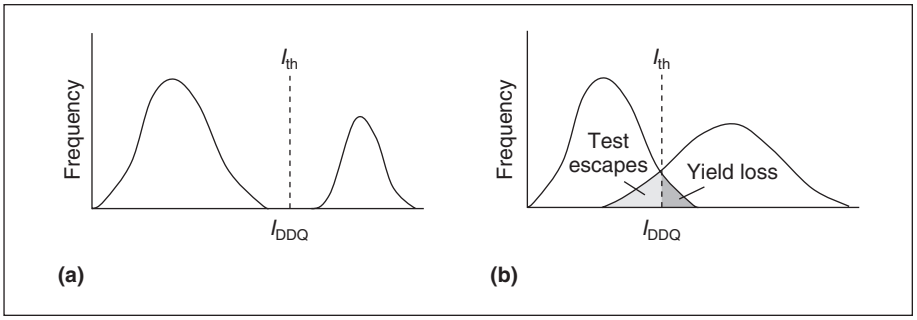

Figure 2. A single threshold I_{DDQ} test could distinguish between faulty and fault-free die distributions for earlier technologies (a), but causes yield loss and/or test escapes for DSM technologies due to overlapping distributions (b).

high-performance ICs. A large background current makes distinguishing defect current extremely difficult, thus reducing the defect resolution of I_{DDQ} test.

Moreover, smaller transistors make it difficult to precisely control transistor geometries, causing large variations in fault-free I_{DDQ} along with high values. With overlapping fault-free and faulty I_{DDQ} distributions, it is impossible to identify defective chips. As Figure 2 shows, any single I_{DDQ} threshold invariably results in false rejects (yield loss) and/or false accepts (test escapes). Industry analysts expect that the I_{DDQ} levels and variations will continue to increase as transistor geometries shrink further,[1] and the *ITRS* considers this a difficult challenge. It seems the advent of DSM technologies has heralded the end of single-threshold I_{DDQ} testing.

Measurement issues

I_{DDQ} can be measured only after inputs are stabilized and internal toggling is settled. This makes I_{DDQ} testing comparatively slow. Although some researchers report speeds of up to 100 MHz, these speeds are mostly design specific. Measuring leakage over the large background current presents a great challenge due to extremely small signal-to-noise ratios. External sensor approaches that measure the voltage drop across a resistance in the power network have proven slow, fueling research into built-in current sensors. BICS are acceptable only if they do not require any special manufacturing, cause no performance penalty, and occupy only a tiny area of valuable silicon real estate. BICS should also be able to measure I_{DDQ} values (for limit-setting techniques) as well as direction (for defect diagnosis).

Solutions

Fundamentally, two ways exist to approach the problem: either reduce the leakage current itself, or find a way to distinguish faulty I_{DDQ} from fault-free leakage. We can divide the proposed solutions into three categories:

- technology solutions,
- design solutions, and
- data analysis solutions.

The first two categories attempt to reduce background leakage; the last attempts to remove the effects of process variation. Technology solutions include reverse body bias (aimed at reducing stand-by leakage),[2] and silicon on insulator—IBM's approach to reduce power demands on high-performance CPUs. Performing I_{DDQ} testing at low temperatures, which reduces background leakage, is often expensive or impractical. Design solutions include multiple-threshold transistors,[3] model-based estimation of I_{DDQ},[4] built-in current sensors,[5] external partitioning of the power supply network, and several others not discussed in this article. Data analysis solutions rely on different (statistical) methods for finding faulty chips. Some methods use within-chip (intradie) vector-to-vector I_{DDQ} correlation. These include current signatures,[6] and ΔI_{DDQ}.[7] Some use information from other chips (interdie) for variance reduction.[8] Other methods combine both inter- and intradie variations. These include current ratios,[9] statistical clustering,[10] and correlation with other process parameters.[11,12] We will limit our discussion to data analysis solutions, since technology and design solutions are discussed elsewhere in detail.[2,3,5]

Current signatures

In its simplest form, a current signature is a graphical display of I_{DDQ} readings sorted in ascending order. It relies on the premise that I_{DDQ} for an active defect is higher (for vectors that excite it) than normal leakage. Thus, the presence of *steps* or *jumps* in a signature means at least two distinct leakage paths or an active defect. Testers screen the dies in case any step size exceeds the predefined limit. Figure 3 shows current signatures for three chips. Chip A, the fault-free chip in Figure 3, exhibits a smooth signature and small variation in I_{DDQ}. Chip B has a passive defect; this elevates all of its I_{DDQ} values, but the signature is smooth. Chip C has an active defect, as the steps in its signature indicate. Chip C also shows several orders of magnitude of variation in I_{DDQ}.

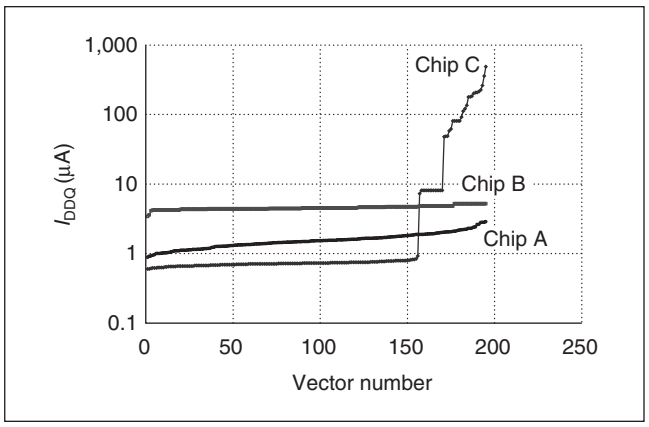

Figure 3. Current signatures reveal a fault-free chip (chip A) and a faulty chip having an active defect (chip C), as indicated by the presence of jumps, but does not reject a chip with a passive defect (chip B).

To obtain a step in the signature, we need at least two vectors: one exciting the defect, and one not. The absence of large steps in the signature can mean either a fault-free chip or a chip with a passive defect (for example, a V_{DD}-to-ground resistive short). Thus, the current-signatures approach cannot detect passive defects. For reasonable defect resolution, several I_{DDQ} readings are necessary. Because I_{DDQ} test is slow, collecting data for current signatures consumes more tester time and costs more. Moreover, deciding the step size threshold for determining a faulty device is not trivial. Step size threshold cannot be judicially set, so unacceptable yield loss/test escapes can occur.

Delta I_{DDQ}

ΔI_{DDQ} is the difference between two consecutive I_{DDQ} readings. Other definitions for the difference between the I_{DDQ} of two chips—I_{DDQ} at different temperatures or voltages, before and after burn-in—are also sometimes called ΔI_{DDQ}. This approach relies on the assumption that fault-free I_{DDQ} variation between vectors is smaller than that introduced by a defect. The random variation in I_{DDQ} causes deltas to be positive and negative with equal probability. For a fault-free chip, the mean ΔI_{DDQ} is close to 0, and standard deviation due to intrinsic variation is very small. Figure 4 (next page) shows histograms of ΔI_{DDQ} for three chips: a fault-free chip, a chip with an active defect, and a chip with a passive defect. It is easy to spot the large standard deviation of the faulty chip in Figure 4b, but the chip in Figure 4c has a passive defect yet is not rejected, because ΔI_{DDQ} values are small and have small variation. As with current signatures, deciding the maximum fault-free delta is challenging.

Current ratios

Despite the increased magnitude and variation in I_{DDQ} in new technologies, the ratio of maximum-to-minimum I_{DDQ} for fault-free chips remains relatively constant. A leaky but fault-free chip consumes proportionately more current for all vectors, and thus its maximum-to-minimum ratio is comparable to fault-free ratios of other chips. This is the basic idea behind the current-ratios test. In this method, the input vectors that cause minimum and maximum I_{DDQ} are determined by characterization, and the current ratio is obtained. The addition of a guard band accounts for process variation. In production, testers measure the I_{DDQ} for the vector causing minimum I_{DDQ} and dynamically adjust the upper limit for other vectors.

As with other approaches, deciding the appropriate fault-free current ratio threshold can be challenging. Figure 5 (page 13) shows current ratios sorted in ascending order for several fault-free chips that passed all Sematech tests, and chips that failed only the 5-μA threshold I_{DDQ} test. Clearly, the faulty chips exhibit more spread in current ratios than the fault-free

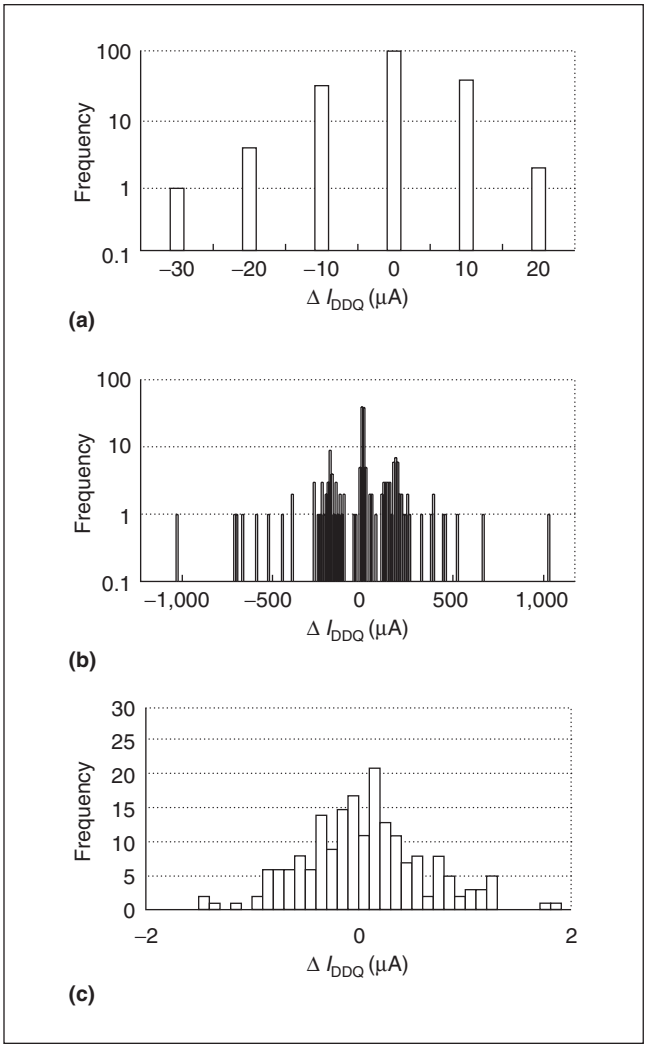

Figure 4. In a Δ I_{DDQ} test, smaller standard deviation distinguishes a fault-free chip, with mean = 0 and standard deviation = 7.4 (a), from a faulty chip, with mean = 0.002 and standard deviation = 216 (b), but does not reject a chip with a passive defect, with mean = 0.004 and standard deviation = 0.57 (c).

chips. Even for the fault-free chips, however, current ratios vary by an order of magnitude. Several defective chips have current ratios comparable to fault-free chip current ratios. Because the minimum I_{DDQ} vector varies from chip to chip, it may be necessary to measure all I_{DDQ} vectors to identify the minimum vector, to avoid unacceptable yield loss or test escapes.

Statistical clustering

Statistical clustering sorts data into groups having a high degree of *natural association* among members of the same group, and a low degree between members of different groups. It uses correlation or other such measures of association to classify the groups. You may loosely consider it a multidimensional regression. Researchers have applied clustering techniques to I_{DDQ} testing.[10] They use the I_{DDQ} data to divide the chips into multiple clusters. Chips with similar I_{DDQ} are sometimes clustered into different groups because of different vector-to-vector correlation. In principle, testers can use any parameter for clustering. For example, clustering I_{DDQ} data combined with chip speed helps find outliers.

Because of its nature of grouping data, clustering inherently accounts for process variations. However, producing meaningful results requires several readings. Additionally, after cluster division, it is up to the user to determine the groups that represent faulty and fault-free chips. Research suggests clustering may be useful to decide appropriate pass/fail limits for each lot/wafer.[10]

Spatial correlation

Neighboring dies on a wafer have highly correlated fault-free parameters because they undergo similar processing. Spatial correlation techniques exploit this feature to estimate fault-free I_{DDQ}. For example, a simple scheme takes the I_{DDQ} average of neighboring dies to estimate the fault-free I_{DDQ} of the center die. Alternatively, we can perform weighted linear regression between the I_{DDQ} of the center die and its neighbors. The best method for estimating the I_{DDQ} is a function of many process parameters.

Figure 6a shows wafer-level variation in I_{DDQ}. Some chips, called gross spatial outliers, have I_{DDQ} values far higher than in neighboring chips. Their absolute I_{DDQ} level, however, falls well within the global distribution, illustrating the difficulty in separating out faulty chips using a single threshold method. Figure 6b shows the local spatial variation in I_{DDQ} without the gross outliers. Some spatial outliers are still visible. Spatial outliers show large variation in residuals (estimated minus actual I_{DDQ}) and have high burn-in fallout rates (chips that fail during or after burn-in).[8] The magnitude and variation of residuals reveal the anomalous behavior of

these chips, which may contain subtle defects not detected by other tests. Even if a chip passes all other tests, it is still a potential customer return. Spatial-correlation-based schemes can reduce burn-in by selectively burning in these chips. For example, in burn-in reduction, a manufacturer might decide to burn in only spatial outliers. Alternatively, in burn-in avoidance, a manufacturer might reject all spatial outliers and ship the remaining chips without burn-in. Such approaches require mature high-yield processes.

The challenges in using such methods include choosing the appropriate outlier rejection method, setting the proper rejection threshold/criterion, and converting I_{DDQ} data to proper standard statistical distributions. Defect clustering causes defects to cluster on the wafer. Neighboring defective estimators can lead to overestimation of fault-free I_{DDQ} and an increase in the defect level. Moreover, the estimate is less accurate when data is available for only a few adjacent chips or when the process environment changes rapidly with distance. This includes dies on the wafer edge or in a poor yield zone. One solution uses dies at longer distances.[13] Selected neighboring dies must exhibit high correlation. Determining highly correlated wafer regions requires careful study of wafer-level spatial patterns. A reliability prediction study found that for wafer edge dies in a microprocessor product, the best predictors were dies on the other edge from the same wafer, or at the same location but on different wafers within the lot.[14] Moreover, unless global outliers are rejected, I_{DDQ} is always overestimated and several faulty chips are accepted, thereby increasing the defect level.

Multiparameter correlation

The necessity to rely on multiple parameters to better identify outliers and faulty chips seems clear. One such parameter is the maximum frequency of operation (F_{max}). The transistors with smaller effective channel length (L_{eff}) switch more quickly and have higher I_{DDQ}. Researchers have found the correlation between F_{max} and I_{DDQ} useful to identify outliers.[11] An alternative to F_{max}, easily measured for latch-based designs, is the flush delay. We obtain the flush delay by turning on all scan clocks, thus converting the

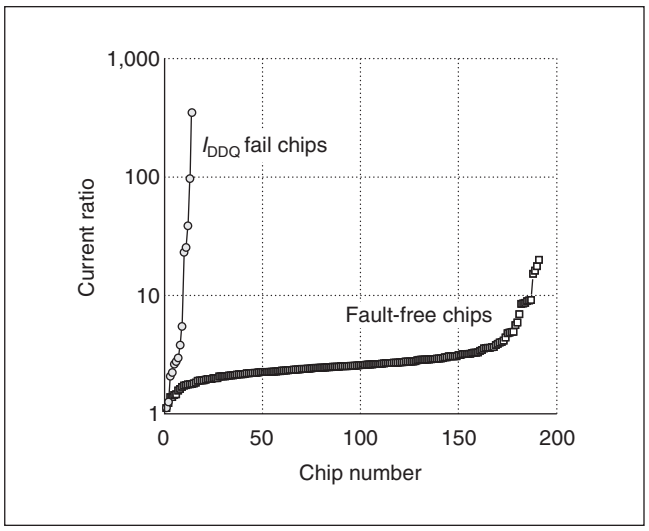

Figure 5. Current ratios of fault-free chips show small variation compared to those of faulty chips.

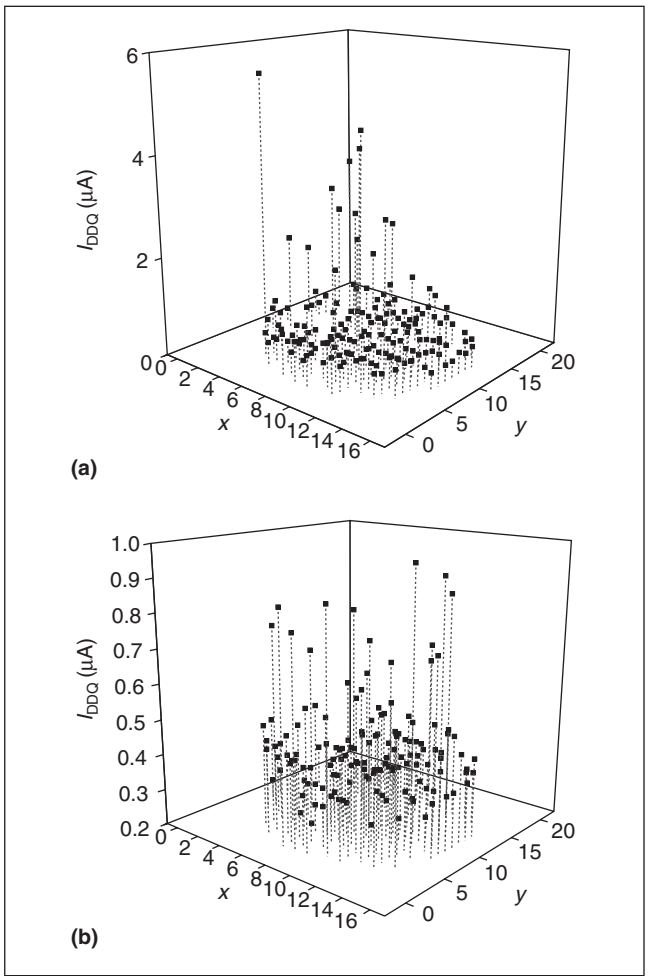

Figure 6. Wafer-level variation in I_{DDQ} with gross spatial outliers (a), and after gross spatial outliers are rejected (b).

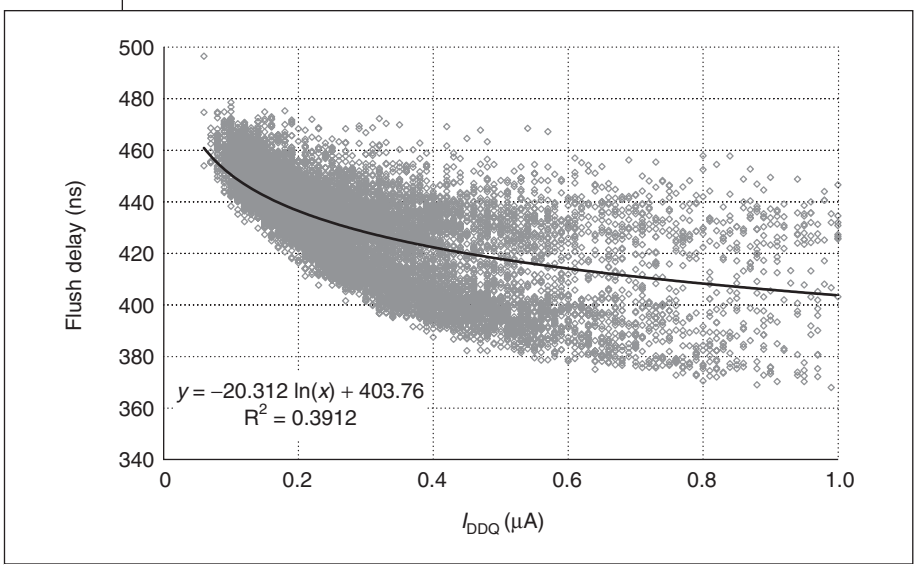

Figure 7. Correlation between I_{DDQ} and flush delay can help estimate maximum fault-free leakage.

scan chain into a chain of buffers. The flush delay equals the time it takes for a signal transition to traverse across the entire chain. As transistor geometries deviate from their nominal values across the chip, flush delay provides a good estimate of chip performance. Figure 7 illustrates the correlation between I_{DDQ} and flush delay for chips that passed stuck-at, ac scan, and functional test. The correlation is higher for intrinsic leakage (smaller I_{DDQ} values) and is useful to improve the estimate of fault-free I_{DDQ}.[12] The correlation varies from wafer to wafer and from lot to lot.

Statistical methods

Researchers have suggested several different statistical approaches for I_{DDQ} data analysis. No method, however, is a panacea. The method used must be suited to the available data. For example, spatial correlation is useful for smooth wafer-level variations, but data with many spatial outliers can result in many test escapes because of overestimation of I_{DDQ}. The limitation of all statistical methods is that only post-test analysis is possible. Some tests need secondary information, such as the die coordinates on the wafer or values from neighboring dies. We can do this at the wafer level using inkless wafer mapping, but application of such techniques at package-level test is difficult. The handler must bin the part (good/bad) after testing is complete. At this point, the tester needs all the necessary information to make a pass/fail decision. This limits the test to using only measurements from the part under test. The tester could also use information from previously tested parts in the lot if access to die-level tracing is available. Statistical means provide greater confidence in defect-level prediction and can provide valuable feedback to improve yield (such as what the most occurring defect was, which process step caused it, and whether there is an equipment malfunction or process parameter drift). Many statistical methods assume data follows a certain standard distribution (for example, normal). It is necessary to convert the data to the appropriate distribution. None of the data analysis approaches completely resolves the threshold setting issue or eliminates yield loss or test escapes. However, such a threshold setting is not arbitrary, but a way of trading off quality (defect level) for cost.

AS DESIGNERS SCALE transistor geometries further, I_{DDQ} values and variation will increase. Researchers project leakage currents of 8 A to 20 A for performance-optimized chips by the year 2014. Understanding the components of the variation in I_{DDQ} is essential to developing the most suitable screening method. Moreover, the defective component of current may fall as the supply voltage is scaled down. This would increase the overlap between faulty and fault-free I_{DDQ} distributions. As each I_{DDQ} technique loses its resolution, test engineers will have to combine them for accuracy. I_{DDQ} dependence on temperature, voltage, input pattern, and correlation with parameters such as flush delay or die position must be exploited to define a multidimensional outlier identification method. Without this, there would be an unacceptable amount of yield loss and test escapes.

I_{DDQ} test poses different challenges for low-

power and high-performance chips. We will be able to use statistical, analysis-based techniques for low-power devices. High-performance chips, however, require a multidirectional approach to the problem. Designers must try to make circuits I_{DDQ} testable by minimizing background leakage. As the external current measurement becomes impractical, using BICS is the only option. To be effective, the performance penalty for BICS must be negligible. Researchers must produce fast, accurate, and sensitive BICS. I_{DDQ} measurement requires rigorous statistical data analysis to reduce yield loss. Manufacturers must be able to define their own statistical procedures to optimally tune pass/fail criteria. It may not be possible to bin the chips until the data from the lot/wafer is collected. Testers must monitor trends in lot-to-lot or wafer-to-wafer variation in I_{DDQ} and use this data in their analyses. The use of inkless flows and electronic databases can support the post-process analysis. In some cases, testers will need the capability to do on-the-fly calculations to make a pass/fail decision. All this is far more complicated than a simple comparison with a threshold.

Despite all these efforts, I_{DDQ} test will continue to lose its resolution to detect defects. Nevertheless, I_{DDQ} will remain a valuable component of the test suite. More research is necessary to understand new defect mechanisms to accurately predict defect levels for bulk CMOS technology with new materials and for emerging technologies. ■

Acknowledgments

This work was funded in part by the National Science Foundation under grant CCR-9971102. We gratefully acknowledge helpful comments of anonymous reviewers that improved the quality of this manuscript. We also thank Phil Nigh of IBM for providing Sematech data. Several figures in this article use data from Sematech experiment S-121.

■ References

1. J. Soden and C. Hawkins, "I_{DDQ} Testing: Issues Present and Future," *IEEE Design and Test of Computers*, vol. 13, no. 4, Winter 1996, pp. 61-65.
2. Z. Chen et al., "I_{DDQ} Testing for Deep-Submicron ICs: Challenges and Solutions," *IEEE Design & Test of Computers*, vol. 19, no. 2, Mar.-Apr. 2002, pp. 24-33.
3. M.C. Johnson et al., "Leakage Control with Efficient Use of Transistor Stacks in Single Threshold CMOS," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 1, Feb. 2002, pp. 1-5.
4. P. Maxwell and J. Rearick, "Estimation of Defect-Free I_{DDQ} in Sub-micron Circuits Using Switch Level Simulation," *Proc. Int'l Test Conf.* (ITC 98), IEEE Press, Piscataway, N.J., 1998, pp. 882-889.
5. W. Maly and P. Nigh, "Built-In Current Testing-Feasibility Study," *Proc. Int'l Conf. Computer-Aided Design* (ICCAD 88), IEEE CS Press, Los Alamitos, Calif., 1988, pp. 340-343.
6. A. Gattiker and W. Maly, "Current Signatures: Application," *Proc. Int'l Test Conf.* (ITC 97), IEEE Press, Piscataway, N.J., 1997, pp. 156-165.
7. C. Thibeault, "An Histogram-Based Procedure for Current Testing of Active Defects," *Proc. Int'l Test Conf.* (ITC 99), IEEE Press, Piscataway, N.J., 1999, pp. 714-723.
8. S. Sabade and D.M.H. Walker, "Improved Wafer-Level Spatial Analysis for I_{DDQ} Limit Setting," *Proc. Int'l Test Conf.* (ITC 01), IEEE Press, Piscataway, N.J., 2001, pp. 82-91.
9. P. Maxwell et al., "Current Ratios: A Self-Scaling Technique for Production I_{DDQ} Testing," *Proc. Int'l Test Conf.* (ITC 99), IEEE Press, Piscataway, N.J., 1999, pp. 738-746.
10. S. Jandhyala et al., "Clustering Based Techniques for I_{DDQ} Testing," *Proc. Int'l Test Conf.* (ITC 99), IEEE Press, Piscataway, N.J., 1999, pp. 730-737.
11. A. Keshavarzi et al., "Multiple-Parameter CMOS IC Testing with Increased Sensitivity for I_{DDQ}," *Proc. Int'l Test Conf.* (ITC 00), IEEE Press, Piscataway, N.J., 2000, pp. 1051-1059.
12. S. Sabade and D.M.H. Walker, "Wafer-Level Spatial and Flush Delay Correlation Analysis for I_{DDQ} Estimation," *Proc. IEEE Int'l Workshop Defect Based Testing,* IEEE CS Press, Los Alamitos, Calif., 2002, pp. 47-52.
13. W.R. Daasch et al., "Variance Reduction Using Wafer Patterns in I_{DDQ} Data," *Proc. Int'l Test Conf.* (ITC 00), IEEE Press, Piscataway, N.J., pp. 189-198.
14. R.B. Miller and W.C. Riordan, "Unit Level Predicted Yield: A Method of Identifying High Defect Density Die at Wafer Sort," *Proc. Intl. Test Conf.* (ITC 01), IEEE Press, Piscataway, N.J., pp. 1118-1127.

Sagar S. Sabade is a PhD student in the Department of Computer Science at Texas A&M University. His research interests include I_{DDQ} test-based outlier identification, burn-in reduction, and defect diagnosis. He has a BE in instrumentation from Pune University, India, and an MTech in electronic design and technology from the Indian Institute of Science, Bangalore. He is a member of the IEEE and the Community of Science.

D.M.H. Walker is an associate professor and associate head in the Department of Computer Science, Texas A&M University. His research interests include defect-based test, realistic fault modeling, I_{DDQ} test, and yield modeling. He has a BS in engineering from the California Institute of Technology, and an MS and PhD in computer science from Carnegie Mellon University. He is a member of the IEEE, the ACM, and Sigma Xi. He is general chair of the 2003 Defect Based Testing Workshop.

■ Direct questions and comments about this article to Hank Walker, Dept. of Computer Science, MS 3112, Texas A&M University, College Station, TX 77843-3112; walker@cs.tamu.edu.