Comparing Functional and Structural Tests

Peter Maxwell, Ismed Hartanto and Lee Bentz

Imaging Electronics Division
Agilent Technologies

ABSTRACT

This paper describes an experimental study to understand issues and requirements for structural-based testing using low cost testers, compared to functional-based testing using expensive testers. Several studies have been directed at the effectiveness of various test methods, but none explicitly addressed issues involved in attempting to replace functional vectors with scan vectors and none carried the experiment further by placing defective chips into systems and running system tests. This paper describes the results of such an experiment and offers insight into necessary requirements for reduction or elimination of functional tests.

1. INTRODUCTION

There is increasing concern today about the escalating cost of test and the ability of ATE to meet stringent timing/accuracy requirements for those test strategies which use at-speed functional tests. Such tests require not only high speed clocks but also accurate edge placement for a large number of pins. The SIA roadmap (SIA, 1999) predicts that if present trends continue, by the year 2014 it may cost more to test a transistor than it costs to manufacture it.

Whereas it might be possible for ATE manufacturers to meet technical the requirements of functional test (e.g. West, 1999) the cost of such equipment remains an issue. Another, less obvious but critically important issue is the impact of a functional test approach on time-to-market. Typically, functional vectors are extracted from system simulations in the form of dump files of values at chip boundaries. These are then converted so that they can run on a tester. The translation process is not always easy since signal changes which occur at arbitrary times in a simulation must be adjusted to fit with timing constraints imposed by test equipment. A limited number of available edges means that the relative timing of some signals has to change, resulting in stimuli which do not match those in the original simulation. Frequently this results in tests which do not work on silicon and have to be modified. The debugging process of such tests is difficult and time-consuming and adds to time to market.

Complex system-on-chip designs, involving several embedded cores, pose substantial problems for a functional test approach. Even if a functional test of satisfactorily high fault coverage were available for each core, it can be very difficult to apply these tests using only the chip pins. Attempting to reverse map a functional test for an embedded core through other on-chip logic and other cores while preserving timing may prove an insurmountable task.

Due to the above there is increasing interest in structural test approaches, involving some degree of DFT. These invariably involve some form of scan-based design, both to allow satisfactory fault coverage of individual cores and to provide a delivery mechanism to get tests to the cores from the chip boundary (Marinissen et al., 1999). The use of built-in self test is also a way of reducing the reliance on test equipment to provide suitable stimuli at required clock rates. Arguments have been put forward that BIST is essential for complex SOC designs (Zorian, 1999).

A test strategy which is based on a structural approach, involving scan and BIST, clearly has reduced requirements on ATE. Scan pins require a large amount of memory behind them, but do not have high speed requirements. Also, scan-based tests typically do not require accurate edge placement on all pins, which means the number of high performance tester channels is reduced.

Traditional test approaches use either functional tests alone, or use them in conjunction with

Faults Detected	Traditional Method	Structural Method
stuck-at	functional/scan	scan/BIST
delay	at-speed functional	AC scan/BIST
IDDQ	sampled functional	scan
memory	functional/BIST	BIST

Table 1. Comparison of traditional and structural test methods

scan tests. Table 1 compares the traditional approach to test with a structured one. There is a heavy reliance on functional tests for all fault types. For static faults they may be augmented by scan tests, and for memory faults they may be augmented with BIST. However, for delay (timing) faults they constitute the major method of detection. In the structured approach, functional vectors are de-emphasized for test cost reasons, but it is important that overall quality does not suffer. Replacing functional vectors by scan vectors requires confirmation that the scan vectors can satisfactorily take over the role of functional vectors or if they cannot, under what conditions are functional vectors required.

There have been several studies in which the effectiveness of various test methods have been evaluated (Maxwell et al., 1991, 1992, 1996; Nigh et al., 1997; Ma et al., 1995; Chang et al., 1998). All these studies showed that each test method rejects parts that no other test method rejects. On the surface this implies one cannot eliminate any of the methods, but there are tradeoffs which are possible. Introducing a more rigorous scan test can reduce reliance on other methods such as functional tests, with no change in part quality. The above studies did not explicitly address issues involved in attempting to replace functional vectors with scan vectors and none carried the experiment further by placing defective chips into systems and running system tests. This paper describes the results of such an experiment and offers insight into necessary requirements for reduction or elimination of functional tests.

2. AC SCAN

The term "AC scan" refers to using a scan chain to deliver transitions to a circuit and capture the responses to those transitions within a given period of time, usually the period of the system clock. Transitions require two vectors two be applied to the circuit, and we will denote the first, or setup, vector by V1 and the second, or stimulus, vector by V2. Different implementations of AC scan vary in the way V1 and V2 are generated and applied.

2. 1 Stuck-at At-Speed

In this method the transition is achieved by a shift of one place of the scan chain. That is, V2 is a shifted version of V1. There are different constraints depending on the type of scan element:

- A muxed-scan design requires that the scan chain can shift at speed since V2 is accomplished by a complete shift of one bit in the scan chain. In addition, the chip has to be able to go from shift phase to capture phase within one system clock cycle.
- A design using separate scan clocks allows a low frequency shift into the master of each scan element, with the application of V2 occurring in normal mode using the system clock.
- Some combinations of signals are impossible due to V2 being V1 shifted. Taking a 3-bit scan chain as an example, V1 = {010}, V2 = {000} cannot be achieved.

In this method V1 is not considered for ATPG since the stuck-at vector applied is V2. The test is similar to a regular stuck-at test except the time allowed for the circuit to respond is limited.

2. 2 Double Clock

In this method ATPG is carried out to determine what both V1 and V2 have to be to obtain a desired transition. V2 is the circuit response to V1 so that the circuit is reverse-mapped to obtain the required V1 from a given V2. After scanning in V1, as in a normal stuck-at test, two system clocks are issued. The first captures the circuit reponse to V1, which launches transitions into the circuit. The circuit response to these transitions is captured on the second system clock pulse. Unlike the first method, this does not require either an at-speed shift capability or a high speed mode change capability.

2. 3 Shift-Update Elements

Special scan elements can be used which use more than the normal two latches (e.g. DasGupta et al, 1981; Dervisoglu & Stong, 1991). This allows V1 to be scanned in, then V2 scanned in without disturbing the value of V1. The system clock is then used to transfer V2 into the output latch, launching transitions, and capturing the response one clock cycle later. Alternatively, the structure can be used for either stuck-at at-speed or double clock tests. For stuck-at at-speed there is no restriction on possible combinations as is the case for two-latch elements. Such a structure allows completely arbitrary V1/V2 pairs, but comes at the expense of a larger scan element which may not be acceptable for area reasons.

3. DETAILS OF THE STUDY

3. 1 Overview

A relatively high volume production chip was chosen and the test program changed to allow full datalogs on all tests. The normal stop-on-first-fail was

inhibited so that data could be collected on each test that was applied to the part.

After a preliminary analysis of the wafer data, some "interesting" parts were identified. These were parts which uniquely failed either functional or ac scan tests, and parts which failed a combination of ac scan and functional, but passed all other tests. The parts were identified and tracked through packaging, then retested using the production package test program, also modified to allow more complete data logging. On the basis of the package test, some interesting parts were selected and sent to the manufacturing site where the parts were being assembled into systems. The manufacturing tests as used at that site were then applied to the parts to determine how the selected packages actually performed in a system.

3. 2 Chip Details

The vehicle was a standard cell ASIC with the following characteristics:

- 268K gates.
- Fully static, 3.3V operation.
- 0.35 µm (Leff), 3 metal layers.
- 7676 flip-flops, full scan design.
- 11 SRAMs.
- Three clock domains, 55 MHz, 33 MHz and 14.67 MHz.

3. 3 Tests Applied

- 3521 stuck-at vectors with 90.4% single stuck-at fault coverage.
- Functional vectors to bring combined coverage to 96%.
- 1327 IDDQ vectors with 81% pseudo-stuck-at coverage. The test method used a fixed 50 µA threshold.
- AC scan vectors (stuck-at at-speed)
- BIST and special tests for RAMs

To make the scan-versus-functional comparison valid, failures due to functional tests which target portions of the chip which are not covered by scan tests need to be treated separately. There are three categories of functional tests directed at areas not covered by scan:

RAM structural tests
PLL tests
"special" functional tests which were added to catch RAM timing paths on the dual-port RAMs which were not detected by the BIST tests.

A large sample of parts from 13 wafer lots were tested. The analysis excluded any parts which failed continuity, DC or pad leakage tests.

3. 4 Breakdown of Chip Area

The proportion of chip area occupied by different categories of circuits is:

standard cell	52%
custom	13%
RAM	18%
pads/drivers	17%

To get an idea of how these map to proportions of failing devices, they need to be multiplied by appropriate weights, to produce an "effective area". This is because different types of logic have different densities, and are more or less susceptible to processing defects. A set of weights which we used are:

standard cell	1.0
custom	1.0
RAM	2.25
pads/drivers	0.4

Multiplying the above areas by the weights and renormalizing, the proportions of effective areas are:

standard cell	46%
custom	12%
RAM	36%
pads/drivers	6%

3. 5 Issues in Application of AC Scan Tests

The basic source for the AC scan vectors were the original stuck-at vectors. However, due to design aspects of the chip and tester loading, not all patterns could be clocked at full speed. The chip had some multi-cycle paths which required two clock cycles in order to settle. The maximum frequency that tests involving these paths could be run was one half of the normal clock frequency.

The second limiting factor in some tests was the fact that the tester imposed a much larger load than is normally present on the pads. This slowed down some pads so that paths that involved reflection from an I/O pad back to internal flip-flops could not be run at full speed. The maximum operating speed for these paths was ¾ of the normal clock frequency.

As a result of these contraints, there were three sets of AC scan vectors:

H - full set (all the stuck-at vectors) run at half speed

M - set with multi-cycle paths removed run at ¾ speed

F - set with slow pad reflections additionally removed run at full speed

3. 6 Functional Vector Coverage

By themselves, the functional vectors had a single stuck-at coverage of 46.5%. Although this figure is relatively low, it does not accurately represent "true" coverage. The chip had a relatively large amount of logic that was not used in the shipping system. Consequently, functional vectors were not written to test this logic so the coverage of used logic is actually higher. Identifying which undetected faults are associated with unused logic is difficult and so far has not been done. The "true" fault coverage is therefore unknown.

4. RESULTS FROM WAFER TESTS

There are 5 major components of the test suite to compare, namely, IDDQ, stuck-at scan, ac scan, RAM and functional tests. Since some of these involve logic not covered by scan tests, separate Venn diagrams are shown in Fig. 1. Fig. 1a compares IDDQ, stuck-at scan, ac scan and logic functional tests (functional tests targetting the scanned logic). The diagram includes IDDQ fails, since although the experiment was not intended to compare other tests with IDDQ, it is still part of the test strategy and the unique contribution of any test type must take IDDQ test into account. Fig. 1b represents only those parts not accounted for in Fig. 1a, i.e. parts failing RAM, PLL or special tests.

A significant number of defective parts were due to RAM failures, but this was predictable from area considerations. The total number of defective chips from all causes was 6297 (note this is larger than the sum of defective parts in Fig. 1 since it includes the excluded parts discussed previously). From the weighted area calculations one would expect anywhere from 20-36% (depending on how accurate the weights are) of these to correspond to RAM failures. In terms of absolute numbers, this corresponds to the number of expected failures falling between 1259 and 2267 parts. From Fig. 1b the total number of guaranteed RAM failures is 918 + 369 + 1 + 503 + 389 + 2 = 2182, which is in good agreement with predictions.

A number of observations may be made from examining Fig. 1a:

1. 87 parts failed functional tests but did not fail any scan test. Of these, 73 (84%) failed IDDQ, which demonstrates the effectiveness of this test type. Without an IDDQ test there would be a much higher reliance on the functional tests.
2. The remaining 14 parts represent unique functional fails which would escape the tests if the functional vectors were removed. Since the functional vectors were not run at slow speed it is not possible to know whether these failures were due to static or speed related defects.
3. The number of unique AC scan fails (16) is almost the same number as unique functional fails. Detecting an equivalent number of defective parts as the number of unique functional fails implies one can trade off one test with the other with little

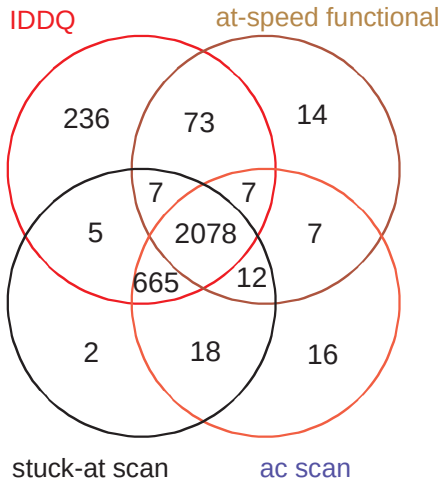

Figure 1a. Breakdown of failures in logic covered by scan

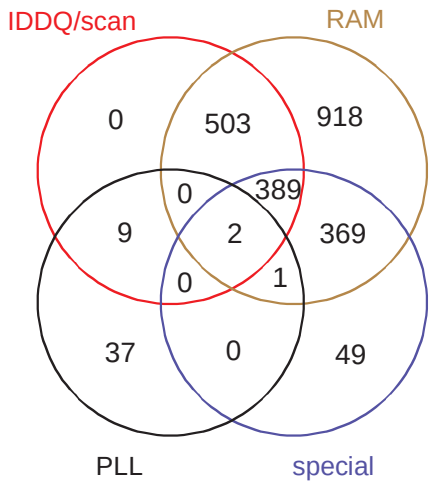

Figure 1b. Breakdown of failures in logic not covered by scan

change in outgoing quality. The assumption is that the AC scan fails are "real". This will be discussed further later.

Fig. 1b shows a number of unique failures from the "special" tests. The fact that these tests were necessary highlights potential limitations in BIST tests. Whereas BIST did an excellent job of testing the memory arrays, they did not satisfactorily exercise timing paths through the RAMs, despite the fact that the BIST ran "at-speed".

Fig. 2 illustrates the situation in simplified form. Due to timing requirements, some of the RAMs were embedded without wrappers and the inputs and outputs were not scannable. BIST tested the RAM only from the input register to the output register. Critical timings paths existed from scannable registers A to B, but an AC scan test could not be generated since it involved propagation through the RAM, which was not

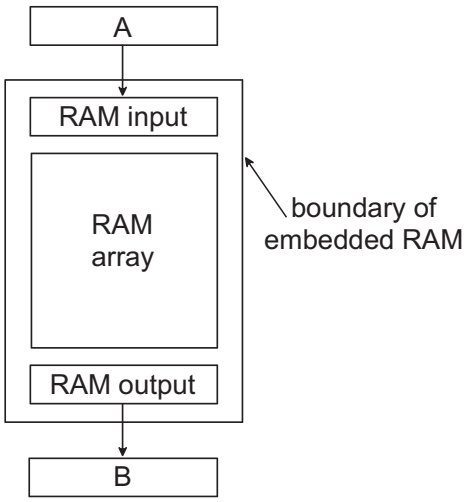

Figure 2. Critical path through RAM.

modelled in the ATPG tools. Functional tests were therefore the only method of testing these paths.

All the tests in the "RAM" portion of Fig. 1b are structural tests in the form of march and/or checkerboard tests. Most of them were applied using BIST, with the exception of some very small RAM blocks for which BIST would have represented too high an area overhead. The only functional tests involving RAMs appear in the "special tests" portion. It is therefore possible to combine the results of Figs 1a and

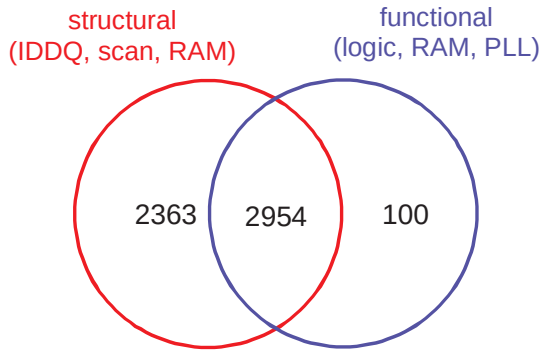

Figure 3. Breakdown of failures detected by structural and functional tests.

1b to produce a breakdown of failures detected by structural and functional tests. Such a breakdown is shown in Fig. 3.

5. RESULTS FROM SYSTEM TESTS

5. 1 Package Screening

As was discussed in Section 3.1, "interesting" die were sent to be packaged. The tracking system could not identify individual die, which were mixed in with normal packages, so the packages of interest had to be screened out in the normal production package test.

Category	Description of IC Test Failure	Failed in System	Passed in System
1	known bad (scan + functional)	20	0
2	unique functional (multiple fails)	19	6
3	fail single functional test only	3	6
4	fail single functional + AC scan	5	5
5	functional (multiple fails) + AC scan	7	0
6	unique AC scan (all M/F, some H)	0	12
7	unique AC scan (only half speed)	2	8
8	unique scan (AC + static)	0	1

Table 2. Breakdown of system test results.

This test was run at 100°C, compared to the wafer test at room temperature, so complete correlation with wafer fallout was not expected. No detailed data logging was done at package test because the majority of defective die were not packaged. The main purpose was to screen out packages which fell into the "interesting" categories.

5. 2 Overview of Tests and Expectations

Packages were placed in a special system which had been used to develop the manufacturing tests for the production line. Although the wafer screens had been done at both high and low voltage the system tests were run at nominal voltage. The source of tests, which were functional in nature, was the same as the functional tests which were run on the IC tester. Naturally there were differences in the way the tests were applied due to the different environments, but the expectation was a high correlation between functional failures during the IC test and system failures.

Whereas functional tests target circuit function, scan tests are structural, targetting specific faults. Unless the fault coverage of the functional vectors was very high, many faults detected by scan vectors (e.g. in many large counters which were infeasible to completely test with functional vectors because of the large number of clock cycles required) would not be detected by the functional vectors and would therefore not fail any system test. In particular, parts which fail ONLY a scan test are unlikely to fail a system test (at least the ones that were run). This is not to say the parts are not defective - when the system is used in a customers site it is quite possible the faults detected by scan tests would cause a failure.

No functional tests were written to cover the unused logic but the scan tests target all the logic. Some unique scan fails and IDDQ fails could therefore be expected to be in this unused logic.

5. 3 Breakdown of Passes and Fails

Table 2 shows the different categories of parts that were tested, and the number that passed/failed the system tests. All categories are mutually exclusive. Note that the numbers should not be correlated with those in Fig. 1a since the defective parts were selected from the package tests, as explained earlier.

Category 1 consitituted known bad parts and were included as a sanity check of the system tests. The single part in category 8 represented one which passed static stuck-at at wafer but failed at package. It was included for interest since there were no more like this (no static scan fails at wafer test were packaged).

Some observations which can be made about these results are:

1. There are some parts in both category 2 and 3 that did not fail system tests. The reason for this is probably due to differences in the way the tests are applied on the IC tester and in the system, resulting in the IC tester failing normally operational parts. The implication is that the impact of removing the functional vectors is not as severe as would otherwise be indicated by the Venn diagrams showing unique functional fails.
2. Perhaps the most surprising result was that none of the parts which uniquely failed the medium and/or full speed tests (category 6) failed in the system. Some parts in category 7 failed, which represent parts tested at half speed. Coverage is probably a major reason for this. Many tests had to be removed from the half-speed tests to get the remainder to operate at higher speeds. Further, the coverage loss is more than just tests involving pad reflections or multi-cycle paths. Because slow pads also affect their output value, strobing outputs could also not be done at full speed. Consequently, no output assertions were made for any test other than the half speed test. If critical paths in the chip involved output pads they may not have been well tested at only half speed.
3. Comparing categories 3 and 4, a higher proportion failed the system tests when the single functional test failure was coupled with an ac scan failure. This indicates the additional defect detection capability of the ac scan tests.

6. DISCUSSION AND CONCLUSIONS

The main objective of the experiment was to compare functional and scan vectors to better understand requirements to enable a test strategy that is mainly based on scan tests. The unique functional failures represent the risk involved in not using functional tests. For this chip the additional 6% coverage of the functional tests over the scan tests means that it is not possible to separate timing from static failures. However, regardless of the cause, such chips escaping the IC test would have a high probability of failing in their intended system. Clearly, the way to minimize such problems is to minimize the amount of logic which is not covered by scan tests, and provide means for testing such logic which does not rely on functional tests.

The 2286 parts in Fig. 1b are a well-known but graphic example of this principle. The RAMs are tested using BIST, and are not covered by any scan tests. Without the provision of the BIST tests (which are

structural, not functional) a scan-based test would have substantial numbers of defective parts escaping. The requirement for the functional "special" tests was a result of some RAM cores being effectively merged cores, with no scannable boundaries.

In some cases functional tests may represent the most practical way of testing some small portions of the logic which are blocked in test mode while scan tests are being applied. Such tests are likely to be much easier to generate than general functional tests and require much less in the way of tester resources.

Further work is needed to understand the unique AC scan fails. It needs to be demonstrated that these failures are "real" failures, as opposed to failing by invoking paths that are never sensitized in normal operation. The system tests showed that parts which failed the half speed test were more likely to fail in the system, which is to be contrasted with the fact that at wafer test more parts failed the full speed test than the half speed. Whereas the system tests had relatively low coverage and are therefore not necessarily a good final arbiter on quality, the result suggests that some of the AC scan failures are exercising functionally unsensitizable paths. This is a potential problem in any scan-based test which places a chip in test mode, where operation is different to that in normal mode. Work is in progress on diagnosis and failure analysis to address this issue.

Conversely to the above, it is necessary to attempt to understand the unique functional failures, in particular, if they are detecting faults in the additional 6% coverage. Diagnosis of functional failures is very difficult, however and it is unclear how successful such a task is likely to be. One of the major problems is trying to determine atwhich clock cycle a fault was activated, since the fault may not cause an incorrect value on a primary output until after many more clock cycles. The sequential nature of the circuit means that traditional fault dictionaries cannot be used. Also, since most defects do not behave as true stuck-at faults, fault simulation becomes less and less reliable the more clock cycles there are after a fault is activated, due to divergence of the faulty circuit from that being fault simulated.

Finally, work is also in progress to determine if any of the unique functional failures can be detected by augmented scan tests. Studies have shown that test sets which target stuck-at faults multiple times can be more effective that those in which multiple detection is fortuitous (Ma et al., 1995; Grimaila et al, 1999; Pomeranz & Reddy, 1999). By applying multiple-detection test sets at-speed a larger number of paths is likely to be exercised with potential detection of some of the unique functional failures.

Improved methods of obtaining speed coverage of I/O pins would help the effects of tester loading discussed earlier. Good results have recently been reported in using a wrap DFT structure to avoid the tester having to contact pads (Gillis et al, 1998).

The experiment has verified that AC scan tests are effective in detecting defects that would otherwise be detected only by at-speed functional tests. It has also demonstrated that in a structurally based test, careful DFT is required to avoid the necessity of some functional vectors. Sometimes these may be needed in areas which were considered well covered by at-speed BIST tests. Fig. 3 shows that by far the majority of defective parts were detectable using structural tests. Most of the remaining parts could also be detected by structural tests by implementing more extensive DFT on the RAMs and by pursuing recent results in BIST of PLLs (Kim et al., 2000).

Further work, as outlined above, is needed to better understand the effects of applying scan vectors at-speed, and what is required to enable them to achieve sufficient timing coverage that the expensive at-speed functional tests are not only reduced to an absolute minimum, but also do not require a large number of high accuracy pins.

The results support a different paradigm for production test, even if functional testing is not completely eliminated. From Fig. 3, of the 5417 defective parts, 98% were detected by purely structural tests. Since these tests could be done using a low cost tester at wafer sort, functional tests using a more expensive tester would need to be done only at package sort. The functional tests require only a fraction of the time required to run the structural tests, so that package throughput would also be increased. Overall test cost would decrease, more than offsetting the slightly higher cost of packaging a small number of defective chips missed at wafer sort.

7. ACKNOWLEDGEMENTS

The implementation of an experiment as described in this paper is not possible without considerable effort from many people. The authors would like to particularly acknowledge the work of Richard Rodgers and Shannon Sawyer in Product Manufacturing, and of Tim Henning in Hewlett-Packard who organized the system tests.

8. REFERENCES

1. J.T.-Y. Chang, C-W. Tseng, Y-C. Chu, S. Wattal, M. Purtell, E.J. McCluskey (1998) "Experimental Results for IDDQ and VLV Testing", *Proc. VLSI Test Symp.*, pp. 118-123.

2. S. DasGupta, R.G. Walther & T.W. Williams (1981) "An Enhancement to LSSD and Some Applications of LSSD in Reliability, Availability and Serviceability", *Proc. FTCS-11*, pp. 32-34.

3. B.I. Dervisoglu and G.E. Stong (1991) "Design for Testability: Using Scanpath Techniques for Path-Delay Test and Measurement", *Proc. Int. Test Conf.*, pp. 365-374.

4. P. Gillis, F. Woytowich, K. McCauley & U. Baur (1998) "Delay Test of Chip I/Os Using LSSD Boundary Scan", *Proc. Int. Test Conf.*, pp. 83-90.

5. M. Grimaila, S. Lee, J. Dworak, K. Butler, B. Stewart, H. Balachandran, B. Houchins, V. Mathur, J. Park, Li-C. Wang & M. Mercer (1999) "REDO - Random Excitation and Deterministic Observation - First Commercial Experiment", *Proc. VLSI Test Symp.*, pp. 268-274.

6. S. Kim, M. Soma & D. Risbud (2000) "An Effective Defect-Oriented BIST Architecture for High-Speed Phase-Locked Loops", *Proc. VLSI Test Symp.*, pp. 231-236.

7. S.C. Ma, P. Franco and E.J. McCluskey (1995) "An Experimental Chip to Evaluate Test Techniques Experimental Results", *Proc. Int. Test Conf.*, pp. 663-672.

8. E.J. Marinissen, Y. Zorian, R. Kapur, T. Tyler & L. Whetsel (1999) "Towards a Standard for Embedded Core Test: An Example", *Proc. Int. Test Conf.*, pp. 616-627.

9. P.C. Maxwell, R.C. Aitken, V. Johansen, and I. Chiang (1991) "The Effect of Different Test Sets on Quality Level Prediction: When is 80% Better Than 90%?", *Proc. Int. Test Conf.*, pp. 358-364.

10. P.C. Maxwell and R.C. Aitken (1992) "The Effectiveness of I_{DDQ}, Functional and Scan Tests: How Many Fault Coverages do we Need?", *Proc. Int. Test Conf.*, pp. 168-177.

11. P.C. Maxwell, R.C. Aitken, K.R. Kollitz & A.C. Brown (1996), "IDDQ and AC Scan: The War Against Unmodeled Defects", *Proc. Int. Test Conf.*, pp. 250-258.

12. P. Nigh, W. Needham, K. Butler, P. Maxwell, R.Aitken & W. Maly (1997) "So What IS an Optimal Test Mix? A Discussion of the SEMATECH Methods Experiment", *Proc. Int. Test Conf.*, pp. 1037-1038.

13. I. Pomeranz and S. M. Reddy (1999) "On n-Detection Test Sets and Variable n-Detection Test Sets for Transition Faults", *Proc. VLSI Test Symposium*, pp. 173-179.

14. SIA (1999) "International Technology Roadmap for Semiconductors 1999 Edition", Semiconductor Industry Association, San Jose, California.

15. B.G. West (1999) "Accuracy Requirements in At-Speed Functional Test", *Proc. Int. Test Conf.*, pp. 780-787.

16. Y. Zorian (1999) "Testing the Monster Chip", *IEEE Spectrum*, July, pp. 54-60.