Balanced Excitation and its Effect on the Fortuitous Detection of Dynamic Defects

Jennifer Dworak, Brad Cobb, James Wingfield, and M. Ray Mercer

Department of Electrical Engineering
Texas A&M University
College Station, Texas 77843-3259
mercer@ee.tamu.edu

Abstract

Dynamic defects are less likely to be fortuitously detected than static defects because they have more stringent detection requirements. We show that (in addition to more site observations) balanced excitation is essential for detection of these defects, and we present a metric for estimating this degree of balance. We also show that excitation balance correlates with the parameter τ in the MPG-D defective part level model.

1. Introduction

Historically, great effort has been expended toward quickly generating short test pattern sets that obtain very high stuck-at fault coverages. Ideally, these test sets are composed of as few test patterns as possible. Unfortunately, such test pattern sets are not necessarily effective for the detection of all of the kinds of defects that may actually occur in integrated circuits [Mill 88], [Max 92], [Ma 95], [Butl 91]. This is because many of the defects actually present in manufactured IC's do not behave exactly like any of the faults targeted during test pattern generation. Instead, the detection requirements of many actual defects often include additional constraints. For example, in the case of an unintentional short, both ends of the bridge must be set to opposite values for detection to occur.

Because the number of potential defects is huge, and because the exact nature of the defects that are actually present in a given set of chips is unknown, generating and simulating test patterns to deterministically detect all possible defects with all possible detection requirements is extremely impractical. However, just as non-targeted faults may be fortuitously detected by a test that targets a different fault, unmodeled defects may be fortuitously detected by a test that targets a mismatched fault. Previous research has shown that detecting faults multiple times and maximizing the number of observations of the least-observed sites increases the probability of satisfying the detection requirements for unmodeled defects [Ma 95], [Grim 99]. This leads to the detection of more defects and reduces the overall defective part level (DPL).

In previous work, we have shown that the reduction of static defective part level as test patterns are applied can be estimated based upon the number of times circuit sites are observed. We have developed a model (named MPG-D) to estimate this reduction. In addition, we have developed a test pattern generation method known as DO-RE-ME (Deterministic Observation, Random Excitation, and MPG-D defective part level Estimation) to generate effective test pattern sets that try to optimize fault site observation. Test patterns generated with this method have been shown to be superior to those generated by a commercial ATPG tool [Dwor 01].

Our previous work has focused on generating better test pattern sets for the detection of *static* defects. However, as clock speeds increase, the detection of *dynamic* (delay) defects has become increasingly important. These dynamic defects are generally harder to detect than static defects. At a minimum, at least two test patterns must be applied at-speed so that an unacceptable delay in the transition of a site's logic value can be detected. Thus, at least one additional constraint—the need for a transition at a circuit site—is required for the detection of a dynamic defect at that site. More precise models of physical defects can only add to these constraints. This paper investigates the relationship between the detection of static and dynamic defects and will show that the additional difficulty inherent in detecting dynamic defects will influence defective part level reduction, test set size, and what constitutes effective test pattern generation methods. We will show that the importance of optimizing the "excitation balance" of test patterns that detect a given fault grows as the non-targeted defects become more difficult to detect, and we will show how this excitation balance relates to one of the parameters in our MPG-D defective part level model.

1530-1591/04 $20.00 (c) 2004 IEEE

2. Predicting static DPL using MPG-D

In the past, we have used the MPG-D defective part level model to analyze the effectiveness of *static* test pattern sets. The underlying basis for this model is the fact that two conditions must be satisfied for any defect to be detected: defect excitation and site observation. Defect excitation refers to setting the test pattern input values so that there is a difference between the logic value in the good circuit and the logic value in the defective circuit at the site where the defect occurs. In contrast, observation of the defective site allows an incorrect logic value to propagate to a primary output or scan chain element. When both excitation and observation occur simultaneously, the defect is detected.

Unfortunately, because the exact defects occurring in a set of ICs are unknown, the excitation requirements for those defects are also unknown. However, regardless of the type of defect present, the site where the defect occurs must be observed for detection to take place. Furthermore, every time the defect site is observed (with a different test pattern and different excitation conditions) there is some probability of *fortuitously* fulfilling the excitation requirements for the undetected defects that may be present there. In the past, we have shown that the probability of exciting an undetected defect at a site i (given that site i is observed) is a decaying exponential function of the number of previous observations (found via fault simulation) and τ [Dwor 01]:

$$P_{excite|obs_i} = e^{-\#obs_i / \tau}$$

The MPG-D model assumes that initially (before any test patterns have been applied) the total defective part level is equally distributed among all circuit sites. So, the initial defect level contribution of site i is:

$$DL_i(0) = \frac{1 - Yield}{\text{\# of sites}}$$

As each test pattern p is applied, the DL contribution of every observed site decreases according to the following formula:

$$\Delta site_i = DL_i(p-1) * A * P_{excite|obs_i} \quad \text{if site } i \text{ is}$$

observed by pattern p and zero otherwise.

Here, $DL_i(p-1)$ refers to the defect level contribution of site i after the first $p-1$ test patterns have been applied. In the simplest case, after pattern p is applied, the defect level contribution of site i can be calculated as:

$$DL_i(p) = DL_i(p-1) - \Delta site_i$$

Then, the overall defect level can be obtained by summing the defect level contributions over all sites:

$$DL(p) = \sum_{i=1}^{\text{\# of sites}} DL_i(p)$$

Thus, the MPG-D model describes how additional observations of a circuit site lead to additional decreases in the defective part level associated with that site. Values for A and τ are chosen to match historical or simulation data for each circuit design. In particular, the constant τ determines how quickly additional observations of a circuit site "run out of steam" and fail to contribute significant additional reductions to DL.

3. Fortuitous detection of dynamic defects

Extending defective part level prediction to dynamic defects and optimizing test pattern sets to detect those defects requires an understanding of how dynamic and static defects are different. Most importantly, the detection requirements of dynamic defects include additional constraints. Detection of those defects is thus inherently harder, and the amount of fortuitous detection that can occur will generally be lower.

Any non-redundant static defect will be capable of being detected by some number of input combinations, or test patterns. Generally, the more test patterns that will detect a defect, the more likely it is to be detected fortuitously when a different fault is being targeted.

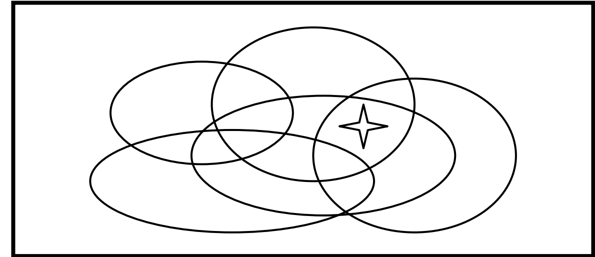

Figure 1

For example, consider the box in Figure 1 to be the set of all possible test patterns (2^n for n inputs) and each oval within the box to be the set of test patterns that will detect a particular static defect. All of the circles are relatively large, and there is significant overlap among them. Thus, if one of the defects is modeled as a fault and explicitly targeted during test generation, there is a high likelihood that several other defects (which may or may not be modeled as faults) will be fortuitously detected simultaneously. For example, if the test pattern corresponding to the location of the star is chosen, three defects will be detected.

When we consider dynamic defects instead of static defects, two things happen. First, the box now represents all test pattern pairs (2^{2n} for n inputs), and second, the relative sizes of the ovals—and thus their overlap areas become smaller, as shown in Figure 2.

In the limit, as more and more requirements must be fulfilled for the detection of difficult defects, the overlaps shrink to zero, and a separate test pattern will be needed for the detection of every defect.

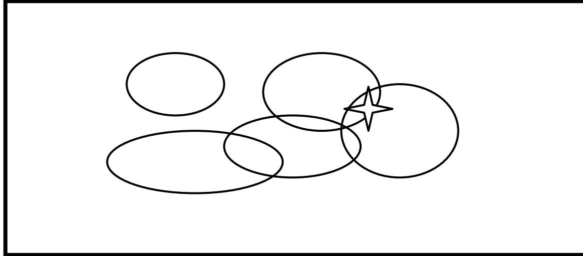

Figure 2

We can investigate this change in fortuitous detection as defects become harder to detect by comparing the expected fortuitous detection of non-targeted stuck-at faults to non-targeted transition faults. (Recall that the fortuitous detection of *defects* given the detection of a targeted fault is analogous to the detection of a *non-targeted fault* given the detection of a targeted fault). Thus, for our experiments, we can consider two defect spaces: a defect space in which all defects correspond to stuck-at faults, and a defect space in which all defects correspond to the more difficult-to-detect transition faults. For the first defect space, we can find how many stuck-at faults we expect to simultaneously detect (on average) given that a particular stuck-at fault has been targeted. We can compare this to how many transition faults (in the second defect space) we expect to simultaneously detect given that a particular transition fault has been targeted.

We can calculate this change in expected fortuitous detection experimentally by using OBDD's to find the set of all tests for every fault considered. Then, the expected number of fortuitous detections of other faults (given that a particular fault has been detected) can be calculated by ANDing the detection OBDD's and finding the size of the overlaps. We can then find the expected fortuitous detection (EFD) for a fault *f* with the following formula:

$$EFD(f) = \frac{\sum_{i=1}^{\# \text{faults} \neq \text{f}} \left| \text{test set f} \cap \text{test set i} \right|}{\left| \text{test set f} \right|}$$

We performed this experiment twice for c432—once for stuck-at faults and once for transition faults. The results of these two experiments can be found in Figure 3. In this graph, the faults are ordered such that each stuck-at fault is paired with its associated transition fault, and the faults are in order of increasing EFD for the stuck-at faults. It is apparent that (in every case) the expected number of fortuitous detections of stuck-at faults (given the detection of a stuck-at fault) is higher than the expected number of fortuitous detections of transition faults (given the detection of a transition fault). The ratio between the EFD of a stuck-at fault and its associated transition fault varies between 1.75 and 3.41, which matches the fact that the transition faults are about twice as hard to detect as the stuck-at faults.

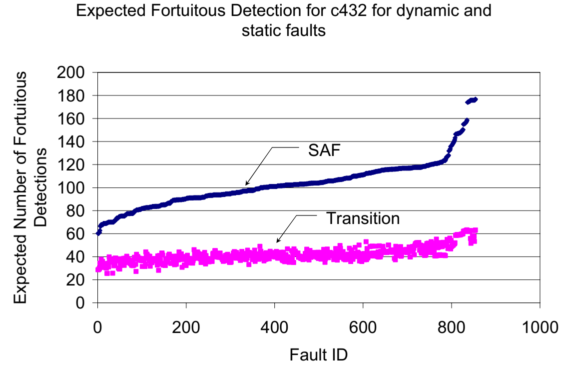

Figure 3

One important effect of the decrease in EFD as faults/defects become more difficult to detect is a corresponding increase in the number of test patterns that must be applied to achieve the same number of minimum detections. When the EFD of *targeted faults* decreases, the test set lengths will naturally become longer to achieve the same fault coverage—especially if aggressive compaction techniques are not employed. In contrast, when the EFD of *unmodeled defects* decreases, the number of observations of sites required to achieve an adequate reduction in defective part level will increase.

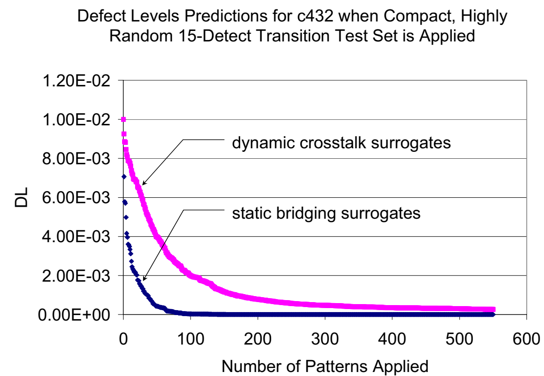

Figure 4

For example, consider the graph in Figure 4. The lower curve shows defective part levels predicted from static AND/OR bridging surrogate simulation for c432 when the observation (second) patterns of a compact 15-

detect transition test set are applied. (S*urrogates* represent unmodeled defects that, unlike faults, are not targeted during test pattern generation.) The test set was generated using the method described in [Lee 02], and almost 92% of the 550 patterns are required for 15 detections of every stuck-at fault. Significant reduction of bridging defective part level is rapidly achieved.

The upper curve shows defective part levels predicted from crosstalk surrogate simulation for c432 for the same test set. These dynamic surrogates are obviously harder to detect than the AND/OR bridges. They can only be detected at a single site and require opposite transitions on the two involved sites. The additional constraints for detection lead to a much slower defective part level reduction. Furthermore, the final defective part levels actually differ by several orders of magnitude.

Thus, although this test set was more than adequate for detection of the static bridging surrogates, a significant number of crosstalk surrogates remained undetected. This indicates that, as contrasted with the static case, a 15-detect test set may not be adequate for the detection of unmodeled, difficult dynamic defects.

4. The importance of balanced excitation

Another effect of the increasing difficulty of dynamic defect detection is the importance of truly random excitation. The DO-RE-ME test generation strategy specifies that each site (especially those with few observations) should be *deterministically observed* as many times as possible while *randomly exciting* whatever defects may be present there. (Note that random excitation refers to randomly exciting the unknown *defects* - not the fault being targeted.) Obviously, observing a site multiple times with identical test patterns will not further reduce the defective part level (unless some other operating conditions are changed). Thus, randomness is needed during test generation to ensure that patterns targeting a given fault are actually different.

Although randomness is required in the test pattern generation procedure, we have not found that a great deal of additional effort was necessary to ensure high randomness in the static case. However, we have found that the degree of randomness does affect the quality of test sets when more difficult dynamic defects must be detected. This can be seen by observing Figure 5.

In Figure 5, we see the results of the application of two different transition test pattern sets to c432 where the resulting defective part level is predicted from crosstalk surrogate simulation. The two test pattern sets were generated with different ATPG tools and have a different number of minimum detections. One ATPG tool used a randomized OBDD-based method described in [Dwor 02]. This tool creates detection OBDD's for every fault in the fault list and chooses a random test from all possible tests for that fault. Specifically, when the detection OBDD is traversed, the choice of which path to follow at a given node is a random choice weighted by the number of minterms that lead to detection along each of the two paths. A 30-detect test set was generated with this tool.

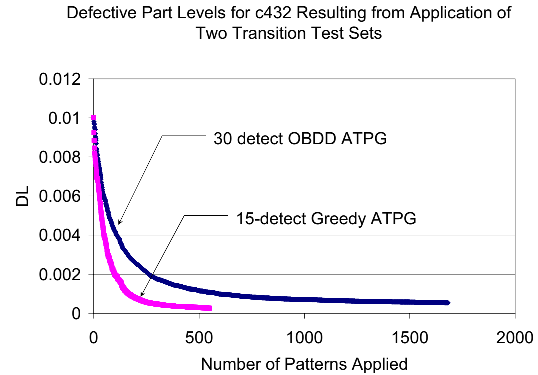

Figure 5

The second test pattern set was generated using a greedy algorithm to create compact transition test sets and is described in [Lee 02]. The greedy algorithm is a combination of the FAN algorithm for fault targeting and parallel pattern simulation. Significant compaction is achieved by choosing primary inputs in groups of five and simultaneously simulating all 32 partial patterns. The numbers of potential and guaranteed fault detections for each partial pattern are calculated and evaluated with a greedy metric so that the best input combination can be chosen. Randomness is achieved by selecting backtrace paths for target nodes randomly and by choosing the five primary inputs randomly. For this experiment, a 15-detect test set was generated with this tool.

Surprisingly, as can be seen in Figure 5, even though the 30-detect test set had many more observations of every circuit site and many more detections of the transition faults, the 15-detect set detected significantly more surrogates. The most reasonable explanation for these seemingly counterintuitive results is that the quality of the excitation of the unmodeled defects is unequal. The 15-detect set generated by the greedy algorithm must have more balanced excitation than the 30-detect test set.

To show the difference in excitation balance between these two test sets, we need to develop a way of quantifying this balance. In previous work on increasing the diversity among test patterns in an n-detect set, some have suggested that two different detections should be counted only if a pattern composed of the common primary inputs (where primary inputs that disagree are set to X) no longer detects the fault [Pom 01]. While this method would help ensure different patterns, it is focused on faults instead of defects. In addition, it does not truly

provide a convenient measure for evaluating excitation balance.

To find a good metric for estimating the balance of excitation, we need to consider how to maximize the probability of detection for an unmodeled defect with arbitrary detection requirements. Specifically, if we have no knowledge about what type of defect is present, then we cannot say with certainty that any particular site should be equal to either a one or a zero to excite that defect. The ideal would be for every site to have an equal chance of being a one or zero (independent of all other circuit values). Of course, this perfectly even division at a given site is not always possible, especially if a pattern has been generated to detect a particular fault. At the very least, the value at the site where the fault is located and some of the circuit values necessary to allow that site to be observed may be deterministically set. However, it is still possible to quantify how close a pattern set is to the ideal with respect to the patterns that detect a particular fault. We propose to estimate this excitation balance for a fault j in the following manner:

$$\text{Exc_Bal}_j = 0.25 - \frac{\sum_{i=1}^{\#\text{ of sites}} \left(0.5 - \text{ones_prob}_{i|j\,\text{detected}}\right)^2}{\#\text{ of sites}}$$

This formula estimates the excitation balance of a test pattern set for those patterns that detect fault j. In the ideal case, a circuit site i will be equal to one exactly half of the time when fault j is detected, and a zero will be added to the sum. In the worst case, a site will always be equal to one or always equal to zero, and a value of 0.25 will be added to the sum. The quantity within the parentheses is squared to give more importance to sites that have highly skewed probabilities. Finally, the average value from the sites is subtracted from 0.25 so that larger values will indicate more balanced excitation. We can use this definition to look at the quality of the excitation balance for the two test sets described earlier from Figure 5. The resulting graph is shown in Figure 6.

Figure 6 shows a dramatic difference in the excitation balance numbers of the two test pattern sets. The graph shows excitation balance values calculated from the observation (second) patterns of the pattern pairs from the two sets. An excitation balance value has been calculated for the two test pattern sets for each non-redundant transition fault. The resulting excitation balance numbers have been sorted from smallest to largest for each test pattern set. The lower curve corresponds to the excitation balance values of the 30-detect OBDD-generated test set. The upper curve corresponds to the excitation balance values for the 15-detect greedy test pattern set. Because higher excitation balance values according to our metric indicate more balanced excitation, the 15-detect greedy test pattern set is significantly better from an excitation balance point of view. This matches the results from Figure 5 and explains why a set with fewer detections of every fault could have significantly better surrogate detection capabilities. Thus, excitation balance does indeed seem to be an important consideration when generating test patterns to detect difficult dynamic defects.

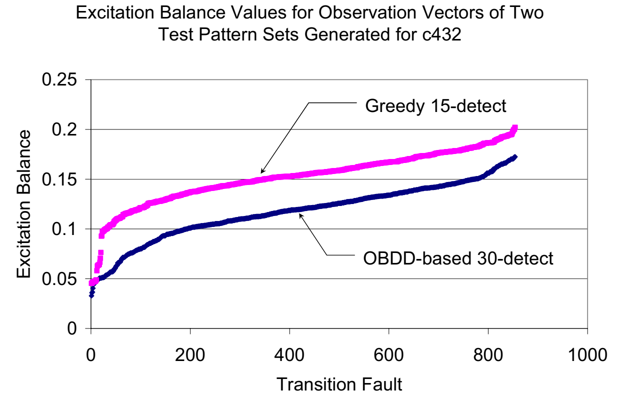

Figure 6

5. Correlation of excitation balance with τ

Although multiple detect test pattern sets are generated by explicitly targeting faults, if we improve the quality of the excitation balance of the test set (by increasing the randomness of defect excitation), we can make the set less biased towards any particular circuit configuration when that fault is detected. This improves defect detection by increasing the effectiveness of additional observations of a circuit site, and this effectiveness is represented by the constant τ in the MPG-D model. Thus, we performed several experiments in which test pattern sets were generated with different ATPG methods to investigate the correlation between our measure of excitation balance and the value of the τ in the MPG-D model. The results can be found in Figure 7.

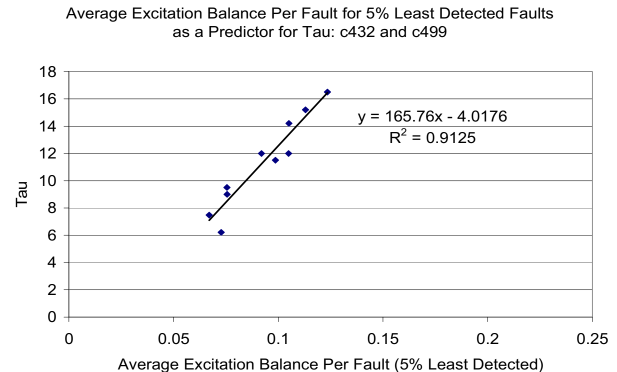

Figure 7

For this graph, we collected data for two ISCAS benchmark circuits: c432 and c499. We evaluated 5 test sets for each. Each test set for a circuit was generated with a different ATPG algorithm. These methods include the greedy algorithm described earlier, the OBDD-based method described earlier, another OBDD-based method in which the weighting of the path to take through the OBDD was altered, an OBDD based method where multiple faults were targeted with a single pattern, and the method employed by a commercial ATPG tool. For each test pattern set created, crosstalk surrogate simulation was performed, and the "ideal" value of τ was determined.

Because the least detected faults contribute the most to the overall defective part level, we concentrated on the 5% least detected faults when calculating excitation balance. For each of these faults, an excitation balance value was calculated for the observation patterns of a given test set. These values were then averaged. (Averaging is required because the circuits are of different sizes and have different numbers of faults.) These average values of excitation balance for the test sets are plotted vs. τ in Figure 7. Although the correlation is not perfect, it is still apparent that excitation balance has a significant relationship with the value of τ, even for two very different circuits. Thus, excitation balance of a test pattern set can be used to estimate the value of τ, at least for defects of the same character.

6. Conclusions

In this paper we have investigated the relationship between static and dynamic defects and explained why dynamic defects are likely to be considerably more difficult to detect. We have shown experimentally that as defects become more difficult to detect, the overlap of their test spaces decreases, and the expected number of fortuitous detections given the detection of a particular fault decreases as well. This leads to longer test pattern sets because each test pattern detects fewer defects.

We have also shown that the minimum number of observations of circuit sites to achieve reasonable defective part level reductions increases as defects become more difficult to detect, but that additional observations are not necessarily enough. Specifically, these additional observations must be balanced to effectively excite the unmodeled defects. We have developed a metric to estimate this excitation balance and have shown that it relates to the time constant τ in the MPG-D model. Our average excitation balance values were able to predict τ even for very different circuits and different ATPG methods. Thus, we can compare test pattern sets with respect to not only the number of site observations, but also with respect to the balance of the defect excitation. In addition, for defects of a given difficulty, we can use the excitation balance of the test set to guide the selection of the constant τ in the MPG-D model and compute more accurate defective part level predictions.

Acknowledgements

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship. Partial support was also provided by the Texas Advanced Technology Program Project No. 000512-0260-2001 and by the Semiconductor Research Corporation under contract SRC 2000-TJ-831. The authors would also like to express their sincere appreciation to Sooryong Lee and Li-C. Wang for their valuable comments and discussion.

References

[Mill 88] S.D. Millman and E.J. McClusky, "Detecting bridging faults with stuck-at test sets," *Proc. International Test Conference*, 1988, pp. 773-783.

[Max 92] P.C. Maxwell, R. C. Aiken, V. Johansen, and I. Chiang, "The effectiveness of I_{DDQ}, functional and scan tests: how many fault coverages do we need?" *Proc. International Test Conference*, 1992, pp. 168-177.

[Ma 95] S. C. Ma, P. Franco, and E. J. McCluskey, "An experimental chip to evaluate test techniques: experiment results," *Proc. International Test Conference,* 1995, pp. 663-672.

[Butl 91] K. M. Butler and M. R. Mercer, "Quantifying non-target defect detection by target fault test sets," *Proc. European Test Conference*, Munich Germany, April 10-12, 1991, pp. 91-100.

[Grim 99] M. R. Grimaila, S. Lee, J. Dworak, K. M. Butler, F. Stewart, H. Balachandran, B. Houchins, V. Mathur, J. Park, Li-C. Wang, and M. R. Mercer, " REDO – Random Excitation and Deterministic Observation – first commercial experiment," *Proc. VLSI Test Symposium*, 1999, pp. 268-274.

[Dwor 01] J. Dworak, J. Wicker, S. Lee, M. R. Grimaila, K. M. Butler, B. Stewart, L-C. Wang, and M. R. Mercer, "Defect-Oriented Testing and Defective-Part-Level Prediction," *IEEE Design and Test of Computers*, January-February, 2001, Vol. 18, No. 1, pp. 31 - 41.

[Lee 02] S. Lee, B. Cobb, J. Dworak, M.R. Grimaila, and M.R. Mercer, "A new ATPG algorithm to limit test set size and achieve multiple detections of all faults," *Proc. Design, Automation and Test in Europe Conference and* Exhibition, 2002, pp. 94-99.

[Dwor 02] J. Dworak, J. Wingfield, B. Cobb, S. Lee, Li-C. Wang, and M.R. Mercer, "Fortuitous Detection and Its Impact on Test Set Sizes Using Stuck-at and Transition Faults," *Proc. 17th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, 2002, pp. 177-185.

[Pom 01] Pomeranz, I. and Reddy, S., "Definitions of the Numbers of Detections of Target Faults and their Effectiveness in Guiding Test Generation for High Defect Coverage," *Design, Automation, and Test in Europe*, March 13-16, 2001, pp. 504-508.