Beyond the Byzantine Generals: Unexpected Behavior and Bridging Fault Diagnosis

David B. Lavo
Tracy Larrabee
Computer Engineering University of California, Santa Cruz CA 95064

Brian Chess
Hewlett-Packard Corporation, Palo Alto CA 94306

Abstract

Physical defects cause behaviors unmodeled by even the best fault simulators, which complicates predictive diagnosis. This paper reports on a diagnosis procedure that uses modified composite signatures constructed from single stuck-at information combined with a lexicographic matching and ranking algorithm. The diagnosis procedure is used to perform high-quality bridging fault diagnosis for more than 400,000 diagnostic experiments involving dropping or adding behaviors from the simulations of faulty circuits.

1 Introduction

Accurate fault diagnosis is an integral part of failure analysis. The purpose of fault diagnosis is to specify a part of the chip to examine in order to determine the cause of failure. Traditional fault diagnosis compares the observed faulty behavior to behaviors predicted by a chosen fault model and fault simulator. As identified by Aitken and Maxwell [4], the problem has been approached in two ways: via a simple fault model (almost always the single stuck-at model) and a sophisticated matching algorithm, or a more accurate fault model (such as a bridging fault model) and a comparatively simple matching algorithm. A good example of an approach using a simple model is that of De and Gunda [11], where a parameterized algorithm ranks all possible stuck-at faults in a circuit via a *merit* value. An example of a system using an accurate fault model is that of Chakravarty and Gong [6], where observed faulty behavior is matched with the expected behavior of bridging fault candidates.

Simple models have the advantage of being easy to construct, and they can be flexibly applied for more complex diagnoses, but they are not amenable to defect location; the most precise answer, the identification of a single node, cannot identify the failure mechanism. In addition, it may occur that a single implicated node spans a large physical area. Often the resulting diagnoses consist of a large number of nodes, and it is possible that none of them may be associated with the true failure mechanism.

Sophisticated fault models, on the other hand, usually provide smaller diagnoses that are rarely misleading, but the increased accuracy has its cost. Not only are sophisticated fault models difficult to create and expensive to apply, but by being extremely specific about the behavior of the candidate faults, diagnosis using sophisticated models may not recognize any behavior resulting from failure mechanisms that were unmodeled or imperfectly modeled.

Previously we presented a technique that offers the precision of an accurate model approach using only simple model information [10]. This paper presents an adaptation of the middle-ground method to overcome a primary weakness of accurate model approaches; the resulting technique is not only precise, it is robust. It provides high-quality diagnoses of bridging faults, even when the bridging fault's behavior is unpredictable. This paper presents the results of several experiments that validate the new approach—even when applied to faulty behaviors that contain a high percentage of noise.

Section 2 discusses the common foundations of diagnosis, and Section 3 briefly describes our technique, Section 4 describes the performance of the technique when noise is incorporated into the observed faulty behavior, and Section 5 analyzes the efficacy and robustness of the new technique and points out the importance of the sparse n-cube space of the candidate signatures and faulty behaviors.

2 Foundations of Diagnosis

Traditional fault diagnosis involves taking the observed behavior of an actual faulty circuit along with a list of candidate faults and choosing the candidate faults that most closely explain the observed behavior. In a simple model approach, the candidate faults are not expected to correspond directly to failure mechanisms—they are simply meant to act as signposts to guide the physical search for the defect. In contrast, a more accurate fault model will ideally point directly to the site of the defect—perhaps even suggesting the process step at which the defect was introduced [16].

Because, when using a simple model approach, the candidates do not directly correspond to the expected failure mechanisms, a sophisticated matching technique is required to produce a diagnosis [17]. Conversely, for sophisticated models, because observed behavior is expected to correspond very closely to modeled behavior, only simple matching algorithms are used to produce a diagnosis.

For either diagnosis approach, the list of faults that best explain the response of the faulty circuit is the *diagnosis*. Every diagnosis can be put in one of two major categories: If it contains the candidate fault that most closely corresponds to an actual failure mechanism, it is a *correct* diagnosis; if it does not contain that candidate, it is an *incorrect* diagnosis. If a correct diagnosis contains the correct candidate and no other, it is an *exact* diagnosis; if a correct diagnosis contains the correct candidate but some additional candidates, it is a *partial* diagnosis. If an incorrect diagnosis contains no candidates, it is a *failed* or *empty* diagnosis.

2.1 Previous Work

Millman, McCluskey, and Acken introduced a method for diagnosing bridging faults: Each potential bridging fault is represented by a *composite signature*, which is the union of the four single stuck-at behaviors associated with the two bridged nodes [20]. A composite signature containing errors that are a superset of the errors contained in the observed faulty behavior is said to be a *match*. The diagnosis of a bridging fault is the list of matches. Note that a diagnosis does not require explicit simulation of bridging faults.

We have previously published an improvement to the technique of Millman, McCluskey, and Acken that continues to use only single stuck-at information but improves on the original technique in three ways: considering only realistic bridges [12], incorporating *match restriction* (flagging some vectors as incapable of detecting a particular bridging fault), and incorporating *match requirement* (flagging some vectors as dependably detecting a particular bridging fault).

Figure 1 illustrates the composite signature of a fault candidate for node **X** bridged with node **Y**; it shows how the candidate is composed of the four single stuck-at behaviors for **X** and **Y**. When **X** stuck-at 0 and **Y** stuck-at 0 are both detected (or when both stuck-at 1s are detected), the bridging fault cannot be stimulated, and these *restricted* vectors are removed from the fault candidate. When a particular vector detects both **X** stuck-at 0 and **Y** stuck-at 1 (or vice-versa), that vector should detect the bridging fault, and it is flagged as a *required* vector.

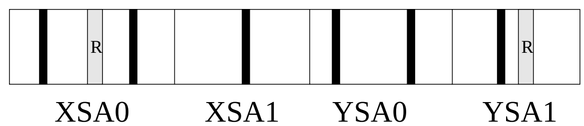

Figure 1: The composite signature of X bridged to Y with match restrictions (in black) and match requirements (labeled R). Corresponding vertical line positions indicate identical vectors.

The new hybrid technique uses elements of the simple model approach and the sophisticated model approach. Like a simple model approach, it uses only single stuck-at information; like a sophisticated model, it requires a candidate list corresponding to expected modes of failure. The matching algorithm used by Millman, McCluskey, and Acken was very simple (behavior containment). Our matching algorithm is more involved, and it provides more precise diagnoses, but neither algorithm will succeed when confronted with unmodeled behavior. If the observed behavior contains errors that fall outside the composite signature of the correct candidate, the candidate will no longer match and the diagnosis will fail.

Our technique was highly successful at diagnosing bridging behaviors. However, variable logic thresholds downstream from the fault site (also known as the Byzantine Generals Problem for bridging faults [1, 2]) caused the diagnosis procedure to provide no candidate faults (an empty diagnosis) roughly ten percent of the time. The Byzantine Generals Problem is only the first of numerous details that can affect the behavior of real bridges.

2.2 Bridging Fault Modeling

Predicting the behavior of bridging faults in CMOS ICs is a difficult problem. Acken and Millman demonstrated that wired-logic is a poor predictor of CMOS bridge behavior and suggested the voting model to take into account variable drive strengths in CMOS logic gates [1]. However, they pointed out that their method did not address the Byzantine Generals Problem for bridging faults [2]. Maxwell and Aitken suggested biased voting as a solution to the Byzantine Generals Problem [19], but biased voting fails to simulate all the repercussions of feedback bridging faults, and it fails to take into account input logic thresholds more than one level away from the fault site. Greenstein and Patel suggested using a mixed-mode simulator with SPICE-level accuracy in a region around the fault site [13], but their method fails to account for timing errors introduced by bridges and may not accurately predict when a feedback bridging fault creates latching behavior [22]. None of the above schemes address the fact that the resistance of a bridging fault is not a known quantity and may significantly change the behavior of the bridge [21]. Even minor modeling deficiencies may cause largely unrecognizable behavior if the state of the modeled machine and the chip being diagnosed begin to diverge.

Although many of these shortcomings are acceptable for fault-grading, the accuracy required by a sophisticated model-based diagnosis algorithm is much greater. Aitken [3] has demonstrated some of the problems with fault prediction, and presented a process of validation and refinement of fault models to improve their accuracy vis-a-vis actual behaviors. Even with such additional work, however, perfect fault modeling is not within reach; sophisticated model-based diagnosis must be able to tolerate some unexpected behavior in the form of both unexpected errors and the absence of expected errors. The original idea of diagnosis using sophisticated models was to model exact behavior and then do simple comparison—which makes the requirement for tolerance somewhat ironic. The emphasis of this research has been to avoid the difficult job of perfecting fault models, and instead to design a matching algorithm that allows for the vagaries of unexpected fault behavior.

3 Unexpected Behavior

A diagnostic algorithm has two functions. First, it selects a subset of the candidate faults to constitute the diagnosis. Second, it orders the candidates in the diagnosis to express the likelihood that each candidate is an explanation of the observed faulty behavior. (Some algorithms may appear to omit one of these two functions, but most of these algorithms are performing a trivial version of the omitted function.)

Our strategy for the diagnosis of bridging faults is to use an algorithm that can (1) relax its subset selection criteria in the face of unexpected behaviors and (2) order the candidates according to how well they meet the expectations. Consider the $\mathbf{n}$ candidate faults to be $\mathbf{C_1}$ through $\mathbf{C_n}$. Each candidate is a set comprised of ⟨test vector:output pin⟩ pairs representing the errors caused by the presence of the fault. The number of individual bit-errors for candidate fault $\mathbf{i}$ is represented by $|\mathbf{C_i}|$. Similarly, the observed faulty behavior is $\mathbf{B}$ (the set of ⟨test vector:output pin⟩ pairs where deviations from fault-free behavior were detected), and the number of individual bit-errors is $|\mathbf{B}|$.

Figure 2 shows diagrammatically the set comparison between the observed behavior ($\mathbf{B}$) and a candidate fault ($\mathbf{C}$). The observed output errors that are correctly predicted by the candidate are represented as set $\mathbf{I}$ (the *intersection*), the output errors that are predicted by the candidate but not observed are set $\mathbf{M}$ (the *mispredictions*—this is also known as the number of *nodetects* [11]), and the output errors that are observed but not predicted by the candidate are $\mathbf{N}$ (the *nonpredictions*).

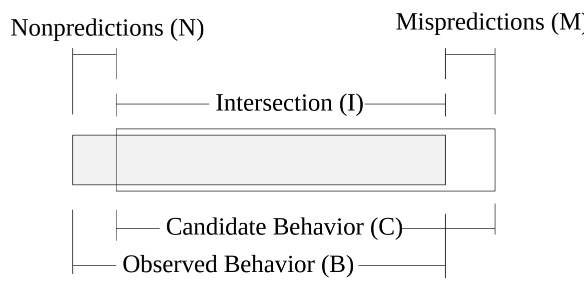

Figure 2: Overlapping of candidate behavior (unshaded) and observed behavior (shaded).

The parameters in the algorithm of De and Gunda [11] control the importance of misprediction and nonprediction with respect to each other. Because of the manner of construction of a composite signature, there are clear indications which should be more important in the diagnosis procedure. The composite signatures are deliberately inclusive—they predict many behaviors that are not expected to occur due to the presence of a bridging fault. If the fault's behavior has been deliberately over-predicted, then an unpredicted behavior is much more unexpected than a prediction that did not occur. Strictly applying this expectation gives $\mathbf{D}$, the diagnosis, to be

$$D = \{C_i \mid B \subseteq C_i\}. \tag{1}$$

Another way to state this is that $|\mathbf{N_i}|$ must be 0 in order for candidate $\mathbf{C_i}$ to be included in the diagnosis. (Note that restricted vectors have already been removed from the candidates before they are considered.)

There are additional expectations represented by the set $\mathbf{R_i}$, the required vectors for candidate $\mathbf{C_i}$: the observed behavior must contain all of the required vectors of the candidate, or the candidate is eliminated. The diagnosis is now

$$D = \{C_i \mid (B \subseteq C_i) \wedge (R_i \subseteq B)\}. \qquad (2)$$

The matching criteria in Equation (2) results in a strict matching algorithm, applicable when ideal bridging fault behavior is assumed. But any unexpected behavior may result in all candidates being excluded from the diagnosis. The resulting failed diagnoses provide the opportunity to recover likely candidates by relaxing the matching conditions, or lowering the expectations expressed by the matching criteria.

The standard technique used by simple model approaches is to first select an acceptable diagnosis size, $\mathbf{d}$, and then construct a diagnosis of the $\mathbf{d}$ best candidates. The ranking criteria are crucial: what is it that makes one candidate superior to another?

Unexpected faulty behavior is the presence of unanticipated errors or the absence of anticipated errors in the observed behavior. The effects of unexpected behavior on the matching process described above are relatively clear: If errors are missing, then the correct candidate may be eliminated because one or more of its required vectors did not appear; if unexpected errors appear, some of the added behaviors may violate the subset criterion (expressed in Equation (1)), and the correct candidate will no longer match.

Our ranking scheme is based on the expectation that the best candidates are the ones that contain the largest amount of the faulty behavior. If the first ranking does not provide enough information, there is another expectation to be formalized: the required vectors of the correct candidates will usually be contained in the observed behavior. $\mathcal{R}_i$, the percentage of predicted required vectors that is fulfilled by the observed behavior (if there are no required vectors, $\mathcal{R}_i$ is equal to 1) is needed to formalize this expectation.

$$\mathcal{R}_i = \begin{cases} 1 & \text{if } |R_i| = 0 \\ \frac{|B \cap R_i|}{|R_i|} & \text{otherwise} \end{cases}$$

Additionally, by removing restricted vectors, the expected size of $\mathbf{M}$ has been reduced. Given the simplifying assumption that all mispredictions are equally likely, if two candidates have the same size intersection with the observed behavior and contain the same percentage of matched required vectors, but one candidate has a much larger $\mathbf{M}$ than the other, the candidate with the larger $\mathbf{M}$ is a less-likely description for the faulty behavior.

Our ranking is a lexicographic ordering in which the nonprediction index has priority, followed by the prediction of the required vectors (used to break nonprediction index ties), and finally the misprediction value (used to break ties of the first two metrics). Formally,

$$\forall C_i \ (|I_i| \geq |I_{i+1}|), \qquad (3)$$

$$\forall C_i \ (|I_i| = |I_{i+1}|) \longrightarrow (\mathcal{R}_i \geq \mathcal{R}_{i+1}), \qquad (4)$$

$$\forall C_i \ (|I_i| = |I_{i+1}|) \wedge (\mathcal{R}_i = \mathcal{R}_{i+1}) \longrightarrow (|M_i| \leq |M_{i+1}|), \qquad (5)$$

and to select the best $\mathbf{d}$ candidates,

$$D = \bigcup_{i=1}^{d} \{C_i\}. \qquad (6)$$

Figure 3 shows four candidate faults compared with the observed behavior and their resulting ranking. $\mathbf{C_1}$, $\mathbf{C_2}$, and $\mathbf{C_3}$ are ranked higher than $\mathbf{C_4}$ because they contain more of the observed behavior, as specified in Equation (3); $\mathbf{C_1}$ and $\mathbf{C_2}$ are ranked higher than $\mathbf{C_3}$ because they contain a higher percentage of matched required vectors, as specified in Equation (4); and $\mathbf{C_1}$ is ranked higher than $\mathbf{C_2}$ because it contains less misprediction, as specified in Equation (5).

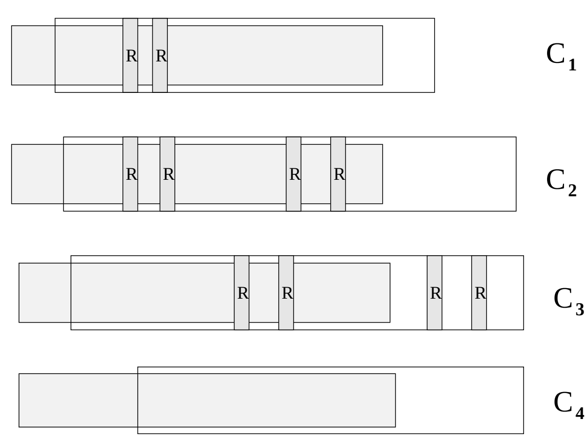

Figure 3: Four candidate faults compared with the observed behavior and their corresponding ranking.

It is important to note that the ranking criteria impose a very general set of expectations, avoiding any assumptions about the nature of unmodeled behaviors. Further experimentation, or greater knowledge about the likely sources of error, may suggest refinements or additional criteria that can be imposed on the candidates to increase accuracy and precision.

4 Diagnosis Experiments

As stated previously, no simulator is a perfect predictor of the behavior of real circuits. If a bridging fault causes timing-dependent behavior or latching behavior, or if it corresponds to a defect that has a non-negligible resistance, the number of errors seen at circuit outputs may increase or decrease from the simulation result. These problems are exacerbated when sequential circuits are considered: fault modeling errors compound themselves as faulty behavior begins to affect the state of the circuit. How robust is the diagnosis procedure outlined above when observed behaviors are unpredictable?

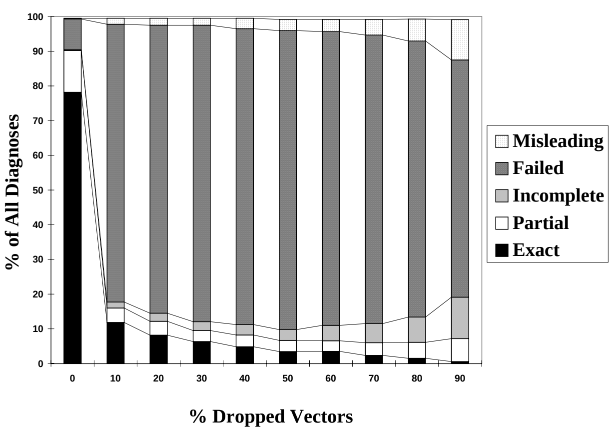

Figure 4: First pass results: Percentage of diagnoses that are misleading, failed, incomplete, partial, and exact as a percentage of dropped output errors.

In order to answer that question, we executed a number of diagnostic trials designed to evaluate the technique. To get a statistically valid sample of potential input for the trials, we first generated observed faulty behaviors for the 10% most likely bridging faults, as determined by Inductive Fault Analysis [15], for each of the MCNC layouts of the ISCAS-85 circuits [5]. The bridging faults chosen covered about 50% of the weighted critical area for the circuits [8].[1] The number of diagnostic trials ranged from 160 for the C432 to 5,379 for the C7552. The time taken for a single diagnosis, on a DEC Alphastation 250 4/266, varied from 1.4s for the C432 to 46.4s for the C7552.

The test set used was generated by a diagnostic test pattern generator to assure the best possible stuck-at diagnosis information [7, 14]. The Nemesis bridging fault simulator [9, 18] produced the observed faulty behaviors but was not used during the diagnostic procedure. Note that while we did not generate observed behaviors for all realistic faults, we did produce composite signatures for all realistic faults (the number of composite signatures ranged from 1,596 for the C432 to 53,790 for the C7552). We then took the bridging fault behaviors generated by Nemesis and modified them by including various types of random noise.

For the first experiment, we randomly removed from 10 to 90 percent of the errors from the observed behaviors. Figure 4 through Figure 6 show the results of this experiment. For each figure, there are ten bars. Each bar is labeled with a percentage of errors dropped from the observed behaviors. The bar values are averages over the ten benchmark circuits, where the values for each circuit are weighted equally.

Figure 4 shows that the first pass, in which strict matching criteria are applied, produces largely failed (empty) diagnoses even when dropping 10 percent of the output errors. Match requirement throws out many normally-matching signatures, so candidate ranking criteria must be used to obtain a meaningful diagnosis. With diagnosis size **d** equal to 10, Figure 5 shows that candidate ranking is very successful even up to the point that 80 percent of the original errors have been deleted. Not only does the correct match appear in the first 10 entries, but, as Figure 6 illustrates, its average position is in the top four entries with, as shown by the standard deviation bars, very little variation over the ten circuits.

Similarly, what if a simple two-line bridging fault were not the only problem with a circuit? If more than one fault occurred, if the bridge spanned three lines, if the bridge created a latch, or if there were any other sources of error in the circuit, the observed errors might be a superset of the predicted errors. How robust is the diagnosis procedure in this case?

In order to answer that question, the observed behaviors from the original experiments were modified by adding new errors to the observed behavior. The added errors were subsets of other simulated bridging

[1] The faults were sampled to reduce the number of trials to a feasible number; the high-probability faults were considered statistically most interesting (and they cover about half of the bridging fault potential for these circuits, as reflected by their weighted critical area). We performed experiments using other samples, and they show no significant difference in diagnostic results.

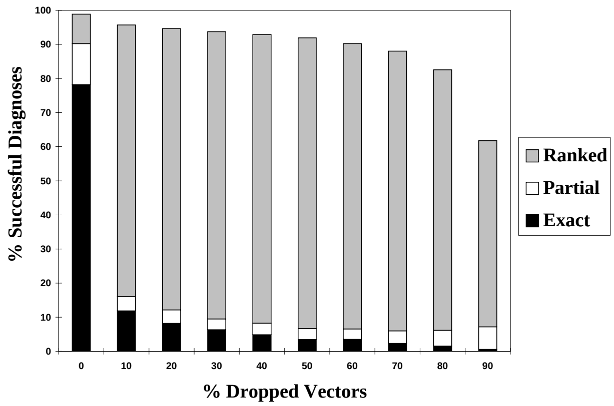

Figure 5: Percentage of correct diagnoses as a percentage of dropped output errors.

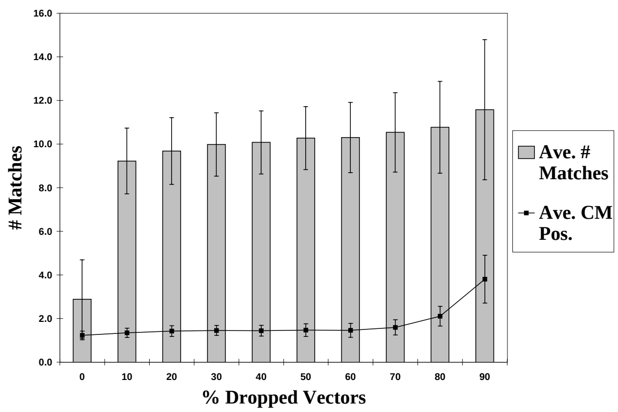

Figure 6: Evaluation of correct diagnoses: Size of diagnosis and position of correct match in candidate ranking as a percentage of dropped output errors.

fault behaviors (to approximate probable behavior). From 10 percent up to 90 percent new errors were added; that is, the total number of errors would never be more than twice as large as the original number of errors. As Figure 7 shows, failed diagnoses dominate. However, Figure 8 shows that candidate ranking produces a correct diagnosis more than 90 percent of the time up to the point where the added errors are 20 percent of the observed behavior. Even when the observed errors are twice as numerous as the original errors, over 60 percent of the diagnoses are successful.

A third set of experiments was conducted in which the same percentage of errors were added and dropped from the observed behaviors. Once again the errors that were added were subsets of other sim-

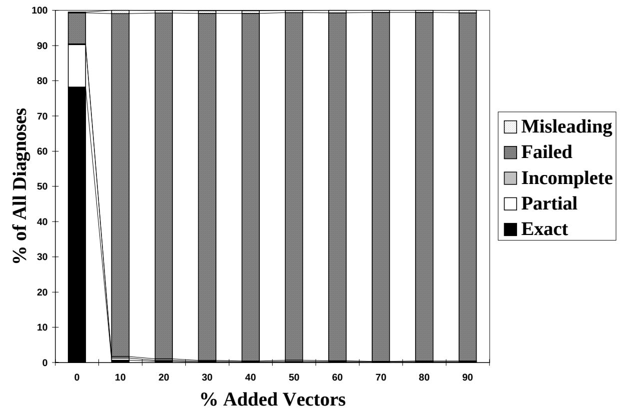

Figure 7: First pass results: Percentage of diagnoses that are misleading, failed, incomplete, partial and exact as a percentage of added output errors.

ulated bridging fault behaviors for the same circuit. Again failed diagnoses dominate, but as Figure 10 shows, candidate ranking produces a correct diagnosis more than 80 percent of the time up to the point where 30 percent of the observed behavior is unrelated to the original bridging fault.

5 Analysis

The diagnoses obtained are surprisingly accurate, even for severely altered behaviors. How can the algorithm still identify the correct candidate so often amidst such extreme noise levels? The question can be answered by rephrasing it: Why does an altered behavior not begin to resemble (alias to) other candidates in the list? Some insight can be found by examining the behaviors and candidates themselves.

Fault signatures can be considered strings of individual detection bits, where each bit represents a ⟨test vector:output pin⟩ pair. The Hamming distance between two fault signatures is the number of bit positions in which the two strings differ. We measured the Hamming distance between 10,000 randomly-selected observed behavior pairs for each of the ISCAS-85 circuits. Figure 12 shows the results for one circuit, the C880. The striking feature of this plot is the distance, or degree of mismatch, commonly found between fault signatures. The C880 (which has 26 output pins and 71 test vectors) has an average number of observed behavior signature bits of 49.2. The mean Hamming distance between sampled pairs is 94.2 bits, and the median Hamming distance

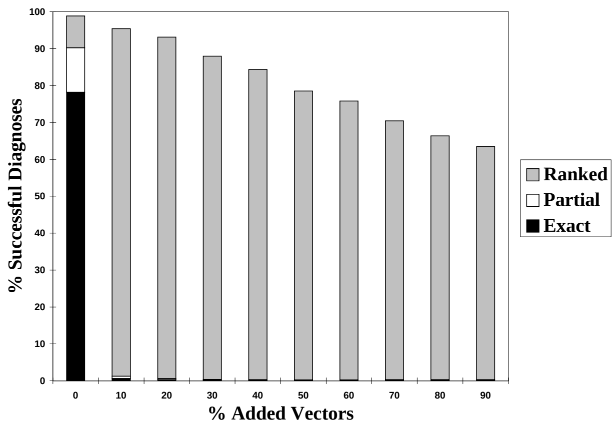

Figure 8: Percentage of correct diagnoses as a percentage of added output errors.

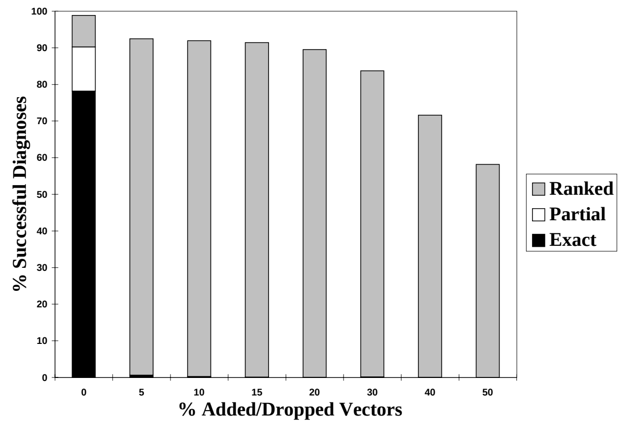

Figure 10: Percentage of correct diagnoses as a percentage of added/dropped output errors.

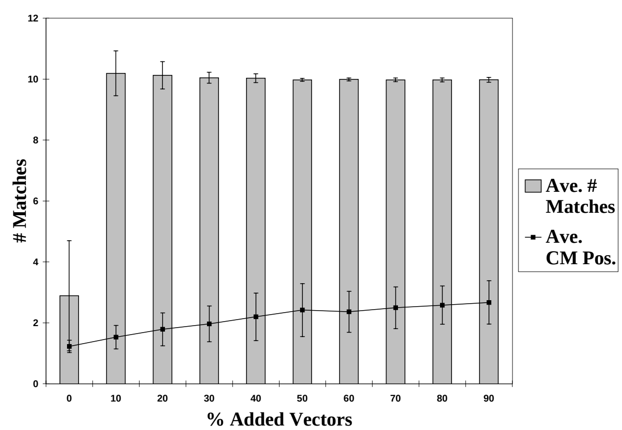

Figure 9: Evaluation of correct diagnoses: Size of diagnosis and position of correct match in candidate ranking as a percentage of added output errors.

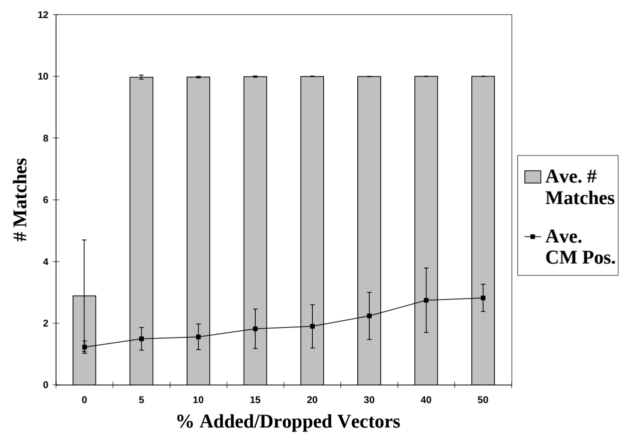

Figure 11: Evaluation of correct diagnoses: Size of diagnosis and position of correct match in candidate ranking as a percentage of added/dropped output errors.

is 91.0 bits. Therefore, the average pair of observed behavior signatures is almost completely disjoint (the average distance is twice the average number of bits). The C880 is shown as a representative circuit; all of the ISCAS-85 circuits display similar results: the average distance between signatures pairs ranges from 1.77 (C432) to 1.97 (C5315) times the average signature size.

The observation of large average distances between behaviors, while quite general, has important implications for diagnosis. If the fault signatures produced by a circuit are not sufficiently distinct, the problem of diagnosis become intractable for any procedure that relies on these signatures, regardless of the accuracy of the model or the presence of noise.

When using sophisticated models for diagnosis, the goal is to have candidates closely correspond to targeted behaviors. The candidates constructed by our diagnosis procedure, as stated, are a compromise between cost and precision: using only stuck-at information, we hope to build candidate composite signatures that approximate, with acknowledged imprecision, the signature of the bridging fault.

Figure 13 is the plot of the Hamming distance between 10,000 randomly-selected composite signatures pairs for the C880. The plot has curves for three sets of samples: original MMA composite signatures, our composite signatures using only stuck-at signature information to restrict (drop) non-detecting vec-

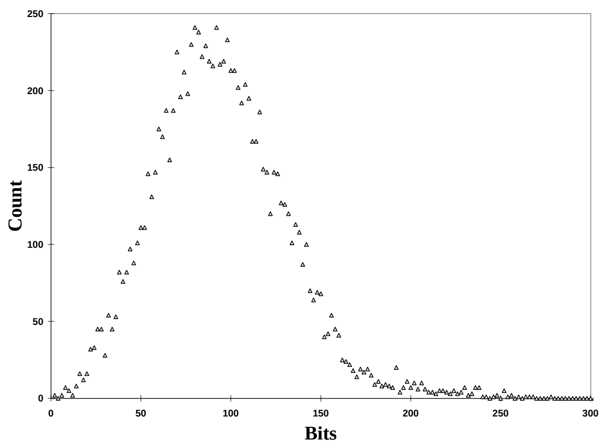

Figure 12: Hamming distance between 10,000 randomly selected observed behavior pairs for the C880.

tors, and our composite signatures when information from a logic simulator is used to restrict vectors [10]. As the curves indicate, improvements arising from additional information reduce the size of the signatures and decrease the variance. For example, for the leftmost curve, the average number of composite signature bits is 88.8; the mean and median distance between sampled composite signature pairs is 165. As expected, the average pair of composite signatures—like the observed behaviors—is almost completely disjoint.

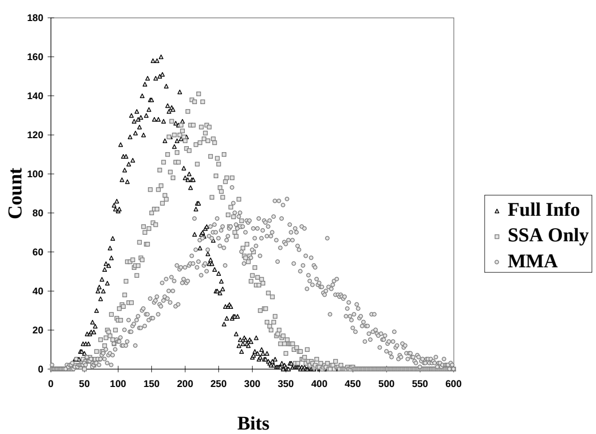

Figure 13: Hamming distance between 10,000 randomly selected composite signature pairs for the C880.

The rankings applied to the candidate matches are designed to reflect, with some additional sophistication, the distance of the candidate signatures to that of the observed behavior. Without the noise of unexpected behaviors, the composite signature of the correct fault for a typical diagnosis is closer to the observed behavior than any other candidate: This can be seen by the high percentage of exact diagnoses in the baseline results, and the high ranking of the correct match in all diagnoses. If, however, many other candidates are ranked almost as high as the correct match, the presence of unexpected behavior may cause these other candidates to be promoted above, and possibly eliminate from the diagnosis, the correct match.

The definition of a successful diagnosis is relatively strict: only 10 candidates are considered, and it only takes the promotion of 10 unwanted candidates to make a diagnosis unsuccessful. Not only must the signatures of the average candidate pair be separated; the rankings should reflect the superiority of the candidates and themselves be separated enough to resist noise. The lexicographic ordering described in Section 3 is converted to a numerical ranking by the diagnosis program.

To show that these requirements are being met by our ranking scheme, we revist the baseline results in Figure 14. Here we plot the top 20 average rankings for all circuits, relative to the ranking of the top candidate (where the error bars represent one standard deviation as measured over all circuits, and the top candidate has no error bar because all percentages are reported relative to it). As demonstrated, the rankings decay significantly. In general the top candidate (which is usually correct) is ranked 30% higher than the 10th or 11th ranked candidate, which makes the top candidate quite resitant to demotion, even in the face of severely altered behaviors.

These experiments support the proposition that unexpected behaviors need not invalidate a model-based diagnosis approach, as long as three conditions hold. First, the behaviors encountered are sufficiently distinct. Second, a candidate behavior closely corresponds to its targeted expected behavior. Third, the matching algorithm can appropriately rank candidates to reflect their correspondence to predicted behavior.

6 Conclusions

Highly accurate fault diagnosis cannot depend on fully predictable fault behavior. We have presented a method of bridging fault diagnosis via matching and ranking with composite signatures that is very successful even when diagnosing largely unpredictable behaviors. Our ranking criteria are based on the

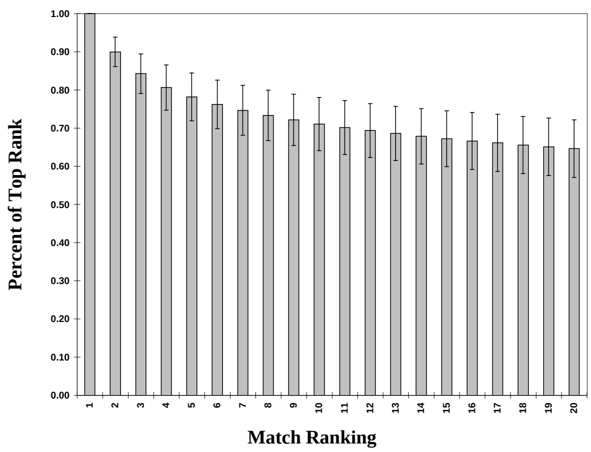

Figure 14: Average rankings, relative to top ranking, for top 20 candidates over all circuits (in the absence of noise).

construction of the composite candidates. Our analysis of signature n-cube space explains some of our success: the candidates and observed behaviors are sparsely distributed over the space, making it possible to alter observed behaviors without aliasing.

References

[1] J. M. Acken and S. D. Millman. Accurate modeling and simulation of bridging faults. *Proceedings of the Custom Integrated Circuits Conference*, pages 17.4.1–17.4.4, 1991.

[2] J. M. Acken and S. D. Millman. Fault model evolution for diagnosis: Accuracy vs precision. *Proceedings of the Custom Integrated Circuits Conference*, pages 13.4.1–13.4.4, 1992.

[3] R. C. Aitken. Finding defects with fault models. In *Proceedings of International Test Conference*, pages 498–505. IEEE, 1996.

[4] R. C. Aitken and P. C. Maxwell. Better models or better algorithms? on techniques to improve fault diagnosis. *Hewlett-Packard Journal*, pages 110–116, February 1995.

[5] F. Brglez and H. Fujiwara. A neutral netlist of 10 combinational benchmark circuits and a target translator in fortran. In *Proceedings of the IEEE International Symposium on Circuits and Systems*, 1985.

[6] S. Chakravarty and Y. Gong. An algorithm for diagnosing two-line bridging faults in combinational circuits. *Proceedings of Design Automation Conference*, pages 520–524, 1993.

[7] B. Chess. Diagnostic test pattern generation and the creation of small fault dictionaries. Master's thesis, University of California, Santa Cruz, June 1995.

[8] B. Chess, A. Freitas, F. J. Ferguson, and T. Larrabee. Testing CMOS logic gates for realistic shorts. In *Proceedings of International Test Conference*, pages 395–402. IEEE, 1994.

[9] B. Chess and T. Larrabee. Bridge fault simulation strategies for CMOS integrated circuits. In *Proceedings of Design Automation Conference*, pages 458–462, 1993.

[10] B. Chess, D. B. Lavo, F. J. Ferguson, and T. Larrabee. Diagnosis of realistic bridging faults with single stuck-at information. In *Proceedings of International Conference on Computer-Aided Design*, pages 185–192. IEEE, 1995.

[11] K. De and A. Gunda. Failure analysis for full-scan circuits. In *Proceedings of International Test Conference*, pages 636–645. IEEE, 1995.

[12] F. J. Ferguson and J. P. Shen. A CMOS fault extractor for inductive fault analysis. *IEEE Transactions on Computer-Aided Design*, 7(11):1181–1194, November 1988.

[13] G. Greenstein and J. Patel. EPROOFS: a CMOS bridging fault simulator. In *Proceedings of International Conference on Computer-Aided Design*, pages 268–271. IEEE, 1992.

[14] T. Gruning, H. Koopmeiners, and U. Mahlstedt. DIATEST: A fast diagnostic test pattern generator for combinational circuits. In *Proceedings of International Conference on Computer-Aided Design*, pages 194–197. IEEE, 1991.

[15] A. Jee and F. J. Ferguson. Carafe: An inductive fault analysis tool for CMOS VLSI circuits. In *Proceedings of the IEEE VLSI Test Symposium*, pages 92–98, 1993.

[16] J. Khare and W. Maly. Inductive contamination analysis (ICA) with SRAM application. In *Proceedings of International Test Conference*, pages 552–560. IEEE, 1995.

[17] R. P. Kunda. Fault location in full-scan designs. In *International Symposium for Testing and Failure Analysis*, pages 121–127. IEEE, 1993.

[18] T. Larrabee. Test pattern generation using Boolean satisfiability. *IEEE Transactions on Computer-Aided Design*, pages 4–15, January 1992.

[19] P. Maxwell and R. Aitken. Biased voting: a method for simulating CMOS bridging faults in the presence of variable gate logic thresholds. In *Proceedings of International Test Conference*, pages 63–72. IEEE, 1993.

[20] S. D. Millman, E. J. McCluskey, and J. M. Acken. Diagnosing CMOS bridging faults with stuck-at fault dictionaries. In *Proceedings of International Test Conference*, pages 860–870. IEEE, 1990.

[21] R. Rodriguez-Montanes, E. M. J. G. Bruls, and J. Figueras. Bridging defects resistance measurements in a CMOS process. In *Proceedings of International Test Conference*, pages 892–899. IEEE, 1992.

[22] C. D. Roth. Simulation and test pattern generation for bridge faults in CMOS ICs. Master's thesis, University of California, Santa Cruz, Department of Computer Engineering, June 1994.