Dynamic Test Set Selection Using Implication-Based On-Chip Diagnosis

N. Alves*, Y. Shi*, N. Imbriglia*, J. Dworak†, K. Nepal‡, R. I. Bahar*
*School of Engineering, Brown University, Providence, RI 02906
†Department of Computer Science and Engineering, Southern Methodist University, Dallas, TX 75205
‡Electrical Engineering Department, Bucknell University, Lewisburg, PA 17837

***Abstract*—We propose using *logic implications* as a source of online diagnostic data for on-chip test set selection by taking advantage of their ability to automatically identify a restricted set of faults as the potential cause of an observed error. This information will be used to dynamically choose a test set to detect systematic latent defects or wearout in a multicore system.**

I. Overview

Logic implications have been previously used for runtime error detection [1] where the error signals from each implication checker were OR'ed together to generate a single error signal. This error signal indicates a failure but does not provide information about which implication was violated. If each implication's individual error signal were saved in a distinct flip-flop, diagnostic information could be obtained at the expense of additional overhead.

This diagnostic information could be used to identify a test set that can be run whenever an implication flags a failure during normal operation. Thus, when applied in a multicore system, this procedure could be used to dynamically select and apply highly effective short test sets that are targeted to sources of previous failures to each core. This will enhance our ability to identify systematic issues and wearout in additional cores before they also have a chance to fail.

Furthermore, to reduce the overheard while retaining as much diagnostic information as possible, we group implications using a greedy algorithm that iteratively compares the faults covered by each implication/implication group, and combines the two that report the fewest fault coverage differences. The error signals from all the implications in a particular group are then OR'ed together and fed into a single flip-flop. As a result, the tradeoff between overhead and diagnostic resolution can be easily calibrated by changing the size of the groups.

Once the implication groups are determined, our test set selection procedure can be easily implemented on-chip. First, a *test superset* is stored in memory (e.g. flash memory, hard drive) alongside an array containing information regarding which patterns are associated with each implication. This array is called the *implication assignment table*. On an implication failure, the implication assignment table is read, and all patterns corresponding to a the failing implication are selected for the test set.

This work was supported in part by NSF under Grants CCF-110290 CCF-0915884, and CCF-1061164.

The selected test set must (a) focus on areas of the circuit that could have caused the error, (b) provide multiple detections of the faults, and (c) be short. To accomplish this, we start by creating a test superset. Next, we select an implication/implication group i in the list and determine a list of suspect faults it can potentially detect. These faults must be well-tested by the patterns we subsequently select.

We assign a weight $W_{i,j} = 1$ if implication i detects fault j, else 0. We then remove all faults with $W_{i,j} = 0$ from the fault list of i. We assign each pattern p in the superset a score S_p based on the faults it detects: $S_p = \sum_{j=1}^{\#\text{faults}} W_{i,j} \times Det_j$, where, Det_j is 0 if pattern p does not detect fault j and 1 otherwise. The pattern with the highest score is added to the pattern set for implication i. Note that while multiple detections of a fault are valuable, the first detection of a fault is more valuable than the second for fortuitous defect detection; the second is more valuable than the third, etc. This decrease in value follows a decreasing exponential function with a time constant τ [2]. Thus, once a pattern is chosen, we update the weights, $W_{i,j}$ using $W_{i,j}(t) = e^{\frac{-Ndet_j}{\tau}}$, where $W_{i,j}(t)$ represents the weight to be used for fault j when the t^{th} pattern of the test subset is to be chosen and $Ndet_j$ represents the number of times that fault j has been detected by the $t-1$ patterns that have already been added to the test subset. Once the new weights are obtained, we recompute the score for each pattern remaining in the superset. The unselected pattern with the highest score is then added to this implication's test pattern assignment list. We repeat this process until the desired number of patterns are included in the list. The same procedure is followed for all of the implications/groups.

Our results show that high diagnostic resolution may be obtained with the proposed method. When implications are not grouped, we can achieve highly targeted test sets that detect the suspect faults multiple times. Even when we group the implications to constrain the hardware overhead, the number of detections per suspect fault is still high, and thus these test sets are appropriate for fortuitous detection.

References

[1] N. Alves, A. Buben, K. Nepal, J. Dworak, and R. I. Bahar, "A cost effective approach for online error detection using invariant relationships," *IEEE Trans. CAD*, vol. 29, no. 5, pp. 788–801, May 2010.

[2] J. Dworak, M. Grimaila, S. Lee, L. Wang, and M. Mercer, "Enhanced DO-RE-ME based defect level prediction using defect site aggregation-MPG-D," in *ITC*, 2000, pp. 930–939.