Miroslaw Malek and M. Ray Mercer
University of Texas at Austin, Austin, TX 78712-1084
Phone: (512) 471-5704 or (512) 471-1804
E-mail: malek@emx.utexas.edu or mercer@cerc.utexas.edu
Contract Title: Testing and Fault-Tolerant Design Techniques
for Advanced Digital Architectures
Contract Number: N00014-86-K-0554
Reporting Period: 1 January 1989 - 31 December 1991

FINAL TECHNICAL REPORT

Our 1989-1991 research on Testing and Fault-Tolerant Design Techniques culminated in numerous results reflected in over 37 publications, a large number of presentations at national and international meetings and followup implementations in computer industry. Here is a brief overview of the most important projects.

1. Topological Testing.

We have continued working on our new concept, topological testing, and demonstrated several applications in the area of multiprocessor testing. Topological testing uses graph theoretic optimization methods, such as the Traveling Salesman Problem, the Chinese Postman Problem, coloring, path covering and partitioning to minimize the test time. The topological testing techniques can be applied to test a system's behavior and its organization at each level of the system's hierarchy; namely, circuit, logic, register transfer, instruction and processor-memory-switch levels. Specifically, the topological testing approach was demonstrated by developing tests for the multistage interconnection network and the hypercube network. We have also developed optimal test sets for mesh networks and further increased a test coverage by including solutions for priority circuits. Optimization is accomplished by transforming a priority circuit testing problem into a graph coloring problem. Time optimization for the testing of these networks gives remarkable results by taking advantage of inherent parallelism and removing test redundancy. One to three orders of magnitude improvement is achieved by applying topological testing techniques to the testing of a wide range of existing networks. The universality and power of the proposed approach may prove to be useful, not only to system testing, but also in system integration. This approach has been successfully applied to conformance testing of communication protocols. The applicability of the method can further be demonstrated in testing of microprocessors at register transfer level, long-haul communication networks and others.

This method yields improvement of over 3,000 times in testing time over pseudoexhaustive testing of multistage networks and has been used in network testing of GF-11 parallel computers at IBM Corporation.

References:
Malek, M., A. Mourad and M. Pandya, "Topological Testing," Proceedings of the 1989 International Test Conference, 103-110, Washington, DC, August 1989.

Özden, B., "Topological Testing of Parallel Computer Networks," Master's Thesis, The University of Texas at Austin, May 1990.

Malek, M., and B. Özden, "Optimized Testing of Meshes," Proceedings of the 1990 International Test Conference, 627-637, Washington, DC, September 1990.

Mourad, A., B. Özden and M. Malek, "Comprehensive Testing of Multistage Interconnection Networks," IEEE Transactions on Computers, 935-951, August 1991.

2. WSI Array Testing.

A novel method of testing components on the wafer has been proposed. This method may potentially surpass any existing method in terms of efficiency and fault coverage. The principle of testing numerous components is based on voting, and if a given element disagrees with the majority, then it is deemed faulty. The element continues to participate in voting, as it is considered a useful member of voting because it is usually given incorrect output for a small subset of tests. Additional study resulted in a method of testing an interconnection network of switches on a wafer. Both methods give a comprehensive approach to WSI Testing and will have an impact on WSI Testing methods in the future.

References:
Choi, Y-H., D. S. Fussell and M. Malek, "Token-Triggered Testing of Wafer-Scale Arrays," in the book titled *Wafer Scale Integration*, C. Jesshope and W. Moore (eds.), Adam Hilger, Bristol and Boston, 246-258, 1986 (presented at Wafer-Scale Integration Workshop in England, July 1985).

Fussell, D. S., M. Malek and S. Rangarajan, "Wafer-Scale Testing/Design for Testability," in the book titled *Wafer Scale Integration*, E. E. Swartzlander, Jr. (ed.), Chapter 9, 413-472, Kluwer, Boston, 1989.

Fussell, D. S., S. Rangarajan and M. Malek, "Fault Diagnosis of Linear Processor Arrays," in the book titled *Defect and Fault Tolerance in VLSI Systems*, I. Koren (ed.), Chapter 6, 149-160, Plenum Press, New York, 1989.

Rangarajan, S., D. S. Fussell and M. Malek, "Built-in Testing of Integrated Circuit Wafers," IEEE Transactions on Computers, C-39 (2), 195-205, February 1990.

Rangarajan, S., D. S. Fussell, and M. Malek, "Efficient Fault Diagnosis of Switches in Wafer Arrays," Proceedings of the 1992 International Conference on Wafer Scale Integration, 341-351, San Francisco, CA, January 22-24, 1992.

3. Fault-Tolerant Parallel Computer Networks.

Design, analysis and a measure of graceful degradation method have been proposed. The new network, cylindrical banyan, yields the best cost/distance properties among a plethora of parallel computer networks. Multistage networks such as SW-banyan, CC-banyan and their extra-stage variations have been analyzed with respect to performance and dependability properties. Our analysis methods allow designers to choose a network which fits their requirements best.

References:

Cherkassky, V., and M. Malek, "Partitioning and Permuting Properties of CC-Banyan Networks," IEEE Transactions on Computers, C-38 (2), 274-278, February 1989.

Cherkassky, V., and M. Malek, "A Measure of Graceful Degradation in Parallel-Computer Systems," IEEE Transactions on Reliability, 38 (1), 76-81, April 1989.

Malek, M., and E. Opper, "The Cylindrical Banyan Multicomputer: A Reconfigurable Systolic Architecture," Journal of Parallel Computing, 10, 319-327, 1989.

4. Reliability Tools.

The most comprehensive-to-date survey of reliability evaluation tools has been published in ACM Computing Surveys and is now one of the most widely used references in this area. A followup on this project is a company called Rainbow Systems which was incorporated by my student Allen Johnson. An extended Petri Net model for dependability evaluation was developed and my student is now incorporating this model to new dependability evaluation software. This software will allow more complete and accurate assessment of software/hardware dependability.

References:

Johnson, A. M., and M. Malek, "Survey of Software Tools for Evaluating Reliability, Availability and Serviceability," ACM Computing Surveys, 20 (4), 227-269, December 1988; translated and reprinted in Japanese, Kyoritsu Shuppan Co., Ltd., publisher, 1990.

Johnson, A. M., "Fault Modeling and Fault Contamination in Multiprocessor Systems," Ph.D. Dissertation, The University of Texas at Austin, December 1989.

5. Naturally Redundant Algorithms.

We have characterized a class of algorithms suitable for fault-tolerant execution in multiprocessor systems by exploiting the existing embedded redundancy in the problem variables. Because of this unique property, no extra computations need be superimposed to the algorithm in order to provide redundancy for fault recovery, as well as fault detection in some cases. A forward recovery scheme is thus employed with very low time overhead. The method is applied to the implementation of two iterative algorithms: solution of

Laplace equations by Jacobi's method and the calculation of the invariant distribution of a Markov chain. Experiments show less than 15% performance degradation for significant problem instances in fault-free situations, and as low as 2.43% in some cases. The extra computation time needed for locating and recovering from a detected fault does not exceed the time necessary to execute a single iteration. The fault detection procedures provide fault coverage close to 100% for faults which affect the correctness of the computations. It is anticipated that this approach, combined with its small ultrareliable kernel, will form a foundation for fault-tolerant distributed computing system design.

References:

Laranjeira, L., M. Malek and R. M. Jenevein, "On Tolerating Faults in Naturally Redundant Algorithms," Proceedings of the Tenth Symposium on Reliable Distributed Systems, 118-127, Pisa, Italy, September 30 - October 2, 1991.

Laranjeira, L., M. Malek and R. M. Jenevein, "NEST: A Nested-Predicate Scheme for Fault Tolerance," accepted for publication in IEEE Transactions on Computers May 1992.

Laranjeira, L., M. Malek and R. M. Jenevein, "Space/Time Overhead Analysis and Experiments with Fault-Tolerant Techniques," accepted for publication in Proceedings of the 3rd IFIP Working Conference on Dependable Computing for Critical Applications, Sicily, Italy, September 14-16, 1992.

6. Responsive Computer Systems Design Framework.

Memory hierarchy was introduced as a natural solution to keep memory speed in pace with CPU speed to increase average system performance at an affordable memory cost. On the other hand, in real-time systems, not the average performance but the performance of each individual task needs to be improved to meet deadlines, and the execution time needs to be predictable to achieve the correct scheduling.

Memory hierarchy, combined with multiprogramming, introduces variable delays to access memory and yields to stochastic execution time of tasks. We have shown how to align each level of memory hierarchy to meet the real-time computing requirements. Our scheme to cache and memory allocation reduces the cold starts with no extra hardware cost.

We conclude that there is no need for special cache designs for real-time systems, and memory allocation can be easily extended to increase the predictability of caches. The I/O scheduling needs to use a scheduling policy similar to the CPU scheduling to achieve high CPU utilization.

We are currently testing this approach in a distributed system and plan to extend it by making it fault tolerant. Reassigning tasks to more than one processor will ultimately create a responsive memory system.

References:
Malek, M., "Responsive Systems, A Challenge for the Nineties," Keynote Address, Proceedings of Euromicro '90, Sixteenth Symposium on Microprocessing and Microprogramming, Amsterdam, The Netherlands, North-Holland, Microprocessing and Microprogramming 30, 9-16, August 29, 1990.

Malek, M., "Responsive Systems: A Marriage Between Real Time and Fault Tolerance," Keynote Address, Proceedings of the Fifth International Conference on Fault-Tolerant Computing Systems, Nürnberg, Germany, Springer-Verlag, Informatik-Fachberichte 283, 1-17, September 25, 1991.

Barborak, M., A. Dahbura and M. Malek, "The Consensus Problem in Fault-Tolerant Computing," Department of Computer Sciences Technical Report #TR-91-40, The University of Texas at Austin, December 1991.

7. Formalization of Fault Tolerance.

We developed a comprehensive formal model for fault-tolerant parallel algorithms and a general methodology for designing reliable applications for multiprocessor environments. The model relies on the formalization of fault-tolerant concepts by means of three nested system predicates and on properties ruling their interrelationships. This rigorous framework facilitates the study of the specific properties that enable an algorithm to tolerate faults. The consequence of that is the outline of systematic design techniques that can be utilized to add fault-tolerant properties to algorithms while preserving their functional characteristics. The proposed model also allows for the quantification of the costs of fault tolerance in terms of space and time redundancy, clarifying the tradeoffs which are inherent to the fault-tolerant design process. The model and design methodology are validated by the uniform application of their principles in the study of several well-known fault-tolerant techniques. The analysis of the cost of fault tolerance in each of these techniques points out that the exploitation of natural redundancy, in applications where this property is present, will lead to the design of fault-tolerant parallel algorithms with very attractive cost/benefit ratio.

References:
Malek, M., and K. H. Yau, "The Resiliency Triple in Multiprocessor Systems," Proceedings of 1988 International Conference on Parallel Processing, 351-358, Chicago, IL, August 15-19, 1988.

Yau, K. H., "The Analysis and Optimization of Fault Tolerance in Multiprocessor Systems: A Graph Theoretic Approach," Ph.D. Dissertation, The University of Texas at Austin, May 1989.

Harary, F., and M. Malek, "Quantifying Fault Recovery in Multiprocessor Systems, to appear in Computers and Mathematics with Applications, Pergamon Press, New York.

8. Innovative Search Techniques.

Design, test and even verification of complex systems may require search. Thus, search is fundamental and progress in search is crucial to numerous disciplines of computer science and operations research. We have been particularly interested in a new search technique called tabu search, and we have also developed a hybrid algorithm technique (HAT). We continue our experiments with a hybrid algorithm where multiple algorithms execute the same problem and exchange information. This promising approach, implemented on shared-memory parallel processors, gave us superlinear speedup. We have achieved an order-of-magnitude speedup on a two-processor system. The only way we can explain it is that, in fact, we have created a new algorithm. We are pursuing analytical and experimental techniques to further investigate this promising approach.

References:

Malek, M., M. Guruswamy, H. Owens and M. Pandya, "Serial and Parallel Simulated Annealing and Tabu Search Algorithms for the Traveling Salesman Problem," Annals of Operations Research, 21, 59-84, 1989.

Malek, M., M. Pandya, M. Guruswamy and H. Owens, "A Hybrid Algorithm Technique," Department of Computer Sciences Technical Report #TR-89-6, The University of Texas at Austin, June 1989.

9. Statistical Models for Delay Faults in Logic Circuits.

The methods for generating tests for classical stuck-at faults are maturing. Unfortunately, technological advances in integrated circuit fabrication and new design styles result in many other manufacturing defects which are not guaranteed to be detected by tests for stuck-at faults. When integrated circuits containing these undetected defects are used in a large complex computing system -- such as those planned for SDI -- the overall reliability of the system can be seriously compromised. The goal of this research has been to identify the most serious non-classical defects and either provide economical methods to detect their existence or identify the point of technological progression where such testing will be required.

If all stuck-at faults have been detected using static (not at speed) testing methods, then delay defects are the most likely to cause system failures. Even when stuck-at testing is done at speed, an adequate measure of protection against delay defects may or may not have been achieved. It is very important to quantify that level of protection so that informed decisions on how much additional testing is required can be made by the manufacturer. Single stuck-at fault coverage has served as such a metric for traditional faults, but delay defects are considerably more complex to model, and their potential impact on system performance is similarly more difficult to analyze.

The work described in the references below was the first to combine a practical engineering approach to delay testing with adequate metrics of the delay testing quality. In particular, the delay fault model has a cardinality which grows linearly with circuit size (in contrast to the path delay model which can grow exponentially in circuit size). At the same time, the good circuit path delay along which each delay defect is detected is used to quantify the test quality. The resulting statistical delay fault coverage figure can also be used to predict defective integrated circuit levels due to delay defects. This work is the follow-on to a previous publication which was recognized via Honorable Mention for the Best Paper at the 1988 International Test Conference.

References:

Park, E. S., M. R. Mercer and T. W. Williams, "A Statistical Model for Delay-Fault Testing," IEEE Design and Test of Computers, 45-55, February 1989.

Park, E. S., M. R. Mercer and T. W. Williams, "Statistical Delay Fault Coverage and Defect Level for Delay Faults," IEEE Transactions on Computers, accepted for publication, September 1990.

10. Test Pattern Generation for Delay Faults in Logic Circuits.

Analysis using the delay defect modeling approach described above may result in a statistical delay fault coverage which is inadequate (the associated defective integrated circuit level is too great). In such situations, additional delay tests must be generated to achieve adequate product quality.

Two different types of delay tests are possible. First, tests for transition defects can be produced -- where the sensitized path along which the detection occurs is not considered. Second, tests for small delay defects can be produced along paths which are the longest (in terms of path delay) or along paths which are almost the longest. The first type of tests are much more economical to generate, and the first reference cited below documents methods for simple modification of stuck-at fault test generation methods to produce transition tests. Again, the metrics described above can be used to determine if adequate product quality has been achieved by transition testing. If not, then methods in the two later references cited below can be used to target small delays and rapidly guarantee delay test quality by targeting defects along paths with large delays.

References:

Glover, C.T., and M. R. Mercer, "A Deterministic Approach to Adjacency Testing for Delay Faults," Proceedings of the 26th ACM/IEEE Design Automation Conference, 351-356, Las Vegas, NV, June 25-29, 1989.

Park, E. S., and M. R. Mercer, "An Efficient Delay Test Generation System for Combinational Logic Circuits," Proceedings of the 27th ACM/IEEE Design Automation Conference, 522-528, Orlando, FL, June 24-28, 1990.

Park, E. S., and M. R. Mercer, "An Efficient Delay Test Generation System for Combinational Logic Circuits," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, accepted for publication July 1991.

11. The Impact of Synthesis Based Timing Optimization on Testing Delay Faults in Logic Circuits.

While the previous two sections deal with problems which exist today in delay testing, work in this section anticipates the effects which technological advances in design automation will have on delay testing. In particular, as design automation systems for logic synthesis become more sophisticated, they will be able to rapidly determine worst case delay paths in the synthesized circuits, and modifications can be done automatically to more equitably allocate resources to equally distribute delay along all paths in the circuit. The result -- in the limit -- will be that the distribution of good circuit path delays will be much more compact because all paths will have nearly the same delay.

There are significant implications to delay testing which result from these observations. First, in the long term, transition testing is more and more likely to be adequate for delay test quality so that more costly small defect delay testing will be less important. Second, in order to maintain a reasonable level of system integrity against delay defects, system clock rates must allow more slack as the good circuit path delay distributions become more compact. This work was recognized with a Best Paper Award at the 1991 Design Automation Conference.

References:

Williams, T. W., B. Underwood and M. R. Mercer, "The Interdependence Between Delay-Optimization of Synthesized Networks and Testing," Proceedings of the 28th ACM/IEEE Design Automation Conference, 87-92, San Francisco, California, June 17-19, 1991. (Best Paper Award at 1991 Design Automation Conference.)

Park, E. S., B. Underwood, T. W. Williams and M. R. Mercer, "Delay Testing Quality in Timing-Optimized Designs," Proceedings of the 1991 International Test Conference, 897-905, Nashville, TN, October 28 - November 1, 1991.

12. Methods for Efficiently Representing and Manipulating Logic Switching Functions.

In the last few years, a significant amount of interest has been generated in the manipulation of switching functions for design automation and testing applications. In particular, Ordered Binary Decision Diagrams (OBDDs) and several variations have been shown to be powerful tools for design verification and testing at the logic level. For example, design verification information can be produced at rates which are thousands of times faster than traditional logic simulation methods when functional methods using OBDD-like representations are employed. Such tools will be invaluable in the design

verification and testing of integrated circuits, modules, boards, and systems such as those required by SDI.

One significant area of improvement possible in functional representations and manipulations is that of switching variable ordering. The key current problem with functional approaches today is the potential size required to represent certain switching functions, and the time required to manipulate a switching function is also determined by the amount of memory space required in the representation. For certain switching functions, the representation size is very dependent upon the order in which switching variables appear in the representation. In the limit, the best ordering may result in a representation which is linear in the number of switching variables while the worst ordering may result in a representation which is exponential in the number of switching variables.

Previous research by others has attempted to analyze the topology of the circuit and generate switching variable orderings which are near optimal. In contrast, our approach is to make very accurate estimates of the quality of a given ordering with very reasonable computational effort. The result is that many potential orderings can be rapidly evaluated and compared until an acceptable ordering is found. Our results using this method have produced the smallest representations to date for a set of widely distributed benchmark circuits. Such advances significantly increase the utility and efficiency of functional approaches to design verification and testing.

References:

Ross, D. E., K. M. Butler, R. Kapur and M. R. Mercer, "Fast Functional Evaluation of Candidate OBDD Variable Orderings," Proceedings of The European Conference on Design Automation, 4-10, Amsterdam, The Netherlands, February 25-28, 1991.

Butler, K. M., D. E. Ross, R. Kapur and M. R. Mercer, "Heuristics to Compute Variable Orderings for Efficient Manipulation of Ordered Binary Decision Diagrams," Proceedings of the 28th ACM/IEEE Design Automation Conference, 417-420, San Francisco, California, June 17-19, 1991.

13. Analysis of Fault Models and Logic Circuit Characteristics.

In addition to delay defects, bridging defects are a source of concern in terms of adequate test quality and product defect levels. Historically, tests for stuck-at faults have also detected almost all bridging defects. It is many times not the case that stuck-at faults and bridging faults cause the circuit to fail in exactly the same way. Instead, the bridging faults are just fortuitously detected as a byproduct of stuck-at testing.

Because of the power of functional analysis tools which are described above, we have been able to analyze the interaction between target stuck-at fault detection and fortuitous bridging fault detection in more detail and with more statistical accuracy than any previous work. In particular, the ensemble sizes for some of these investigations exceeds

10^{4000} cases. Such a large number of cases could never be handled explicitly by traditional simulation approaches, but the use of OBDDs and switching function representations have made such analyses possible.

In contrast to delay faults, we have determined that fortuitous detection of bridging defects by tests for stuck-at faults will be adequate until integrated circuit part defect levels drop below one hundred defective ICs per million. Around that point, more sophisticated testing methods will be required to maintain acceptable quality levels. For example, functional methods for automatic test pattern generation may produce many alternative tests which can be used to detect a given stuck-at fault. However, among all those tests, one may be much superior in the detection of non-target bridging defects.

References:

Butler, K. M., and M. R. Mercer, *Assessing Fault Model and Test Quality*, Kluwer Academic Publishers, 1991, ISBN 0-7923-9222-1.

Kapur, R., K. M. Butler, D. E. Ross and M. R. Mercer, "On Bridging Fault Controllability and Observability and Their Correlations to Detectability," Proceedings of The European Test Conference, 333-339, Münich, Germany, April 10-12, 1991.

Kapur, R., and M. R. Mercer, "Bounding Signal Probabilities for Testability Measurement Using Conditional Syndromes," IEEE Transactions on Computers, accepted for publication December 1991.

14. Summary.

In summary, we have made a number of significant contributions to the testing and fault-tolerant design of complex logic systems and computer architectures. Such contributions will improve the efficiency and accuracy of the design of such systems; they will improve the quality of tested manufactured products used to build such systems; and they will provide fault-tolerance to assure dependable operation in the field. These contributions have been documented in the technical literature, and transfer of these technological innovations to industry has been initiated.