

Zero-Aliasing Space Compaction of Test Responses Using Multiple Parity Signatures

Krishnendu Chakrabarty, *Member, IEEE*, and John P. Hayes, *Fellow, IEEE*

Abstract—We present a parity-based space compaction technique that eliminates aliasing for any given fault model. The test responses from a circuit under test with a large number of primary outputs are merged into a narrow signature stream using a multiple-output parity tree. The functions realized by the different outputs of the compactor are determined by a procedure that targets the desired fault model. Experimental results for the ISCAS-85 benchmarks show that zero aliasing of single stuck-line faults can be achieved with a two-output parity tree compactor. Our findings corroborate recent results on the fundamental limits of space compaction.

Index Terms—Aliasing, built-in self-testing (BIST), parity trees, test response compaction.

I. INTRODUCTION

SPACE compaction is necessary for built-in self-testing (BIST) to compress a k -bit-wide data stream to a q -bit-wide signature stream, where $q \ll k$. A basic problem associated with compaction is aliasing, which occurs when a faulty response maps to the fault-free signature. Aliasing affects BIST testing quality by reducing fault coverage, therefore design techniques that eliminate aliasing are of considerable interest.

A desirable property of a space compaction circuit is that it propagate all errors from its inputs to its outputs. A parity element such as an exclusive-or gate has very good signal propagation properties, therefore, parity tree circuits are often considered ideal for space compaction [2], [9]. A parity tree propagates all errors that appear on an odd number of its inputs; however, it masks errors that appear on an even number of inputs. While experimental results indicate that a high percentage of single stuck-line (SSL) faults in typical logic circuits are detected with a parity tree space compactor [2], [5], zero-aliasing compaction is rarely achieved.

A multiple-output parity tree has better error propagation properties than a single-output tree. This is exploited by Tarnick [10] to achieve high coverage of the errors induced by a set of SSL faults. He uses a greedy covering procedure to find the number of compactor outputs q required to ensure zero aliasing for the given set of errors. A drawback of this method is that the value of q is high and zero aliasing is not achieved for all SSL faults.

Manuscript received October 15, 1995; revised October 15, 1997. This work was supported by the National Science Foundation under Grant MIP-9503463.

K. Chakrabarty is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA.

J. P. Hayes is with the Advanced Computer Architecture Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA.

Publisher Item Identifier S 1063-8210(98)02946-1.

Zero-aliasing space compaction for a specific fault model has been studied recently using a graph model [3]. For a given circuit under test, a fault model, and a test set, maximal compaction (smallest value of q) is achieved by relating the compactor design to the graph coloring problem. It is shown in [3] that the number of test observation points can be reduced to one or two for all the ISCAS-85 benchmarks [1]. However, with this approach it is difficult to determine the hardware overhead, in part because maximal compaction does not always yield the lowest cost compactor.

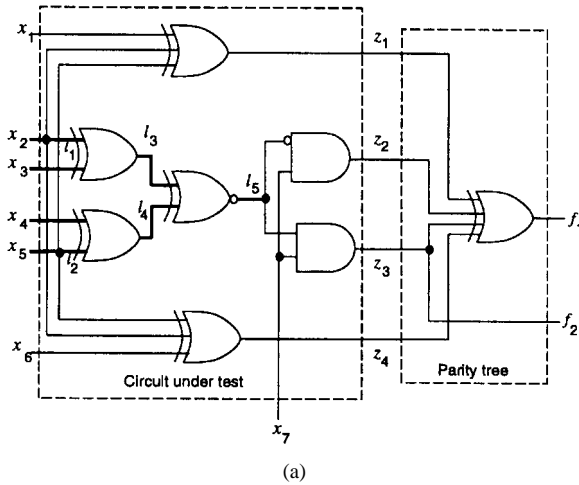
We present a new technique for designing zero-aliasing space compactors tailored to a specific fault model. We show that zero aliasing for SSL faults can be achieved for all the ISCAS-85 benchmark circuits using a parity tree with at most two outputs. This is consistent with the results on maximum compaction presented in [3]. The advantages of using parity trees for space compaction are that the hardware overhead is determined a priori, and close to 100% coverage can be easily achieved for faults in the compactor itself [4].

II. MULTIPLE PARITY SIGNATURES

We first review some terminology [2] related to fault sensitization. A fault f is *odd-sensitized* by a test set \mathcal{T} if there exists at least one test pattern $t \in \mathcal{T}$ that sensitizes f to an odd number of primary outputs. A fault that is not odd-sensitized is called *even-sensitized*. To achieve zero aliasing, we have to eliminate even sensitization. This can be done either by altering the test set or by modifying the parity tree [2]. However, both these methods lead to a longer response sequence with a corresponding increase in the test application time. Multiple parity functions (signatures) allow us to achieve zero aliasing without requiring any changes to the test set or any increase in the test application time.


Let the circuit under test have k outputs z_1, z_2, \dots, z_k , and let $f_1 = z_1 \oplus z_2 \oplus \dots \oplus z_k$ be the function realized by a single-output parity tree. Our goal is to determine $q - 1$ additional parity functions f_2, f_3, \dots, f_q such that an error in the circuit under test is propagated by at least one of the f_i 's, $1 \leq i \leq q$. These functions are of the form: $f_i = x_{i,1}z_1 \oplus x_{i,2}z_2 \oplus \dots \oplus x_{i,k}z_k$, $2 \leq i \leq q$, and where each $x_{i,j} \in \{0, 1\}$. First, we assume that $x_{i,1} + x_{i,2} + \dots + x_{i,k} = 1$, $1 \leq i \leq q$. This implies that the additional parity functions are some selected output functions of the circuit under test.

Parity functions can be selected from a *sensitization cover* [2], which is a minimal set of primary outputs to which every even-sensitized fault is sensitized by the test patterns in \mathcal{T} . It can be obtained from the sensitization table for \mathcal{T} , whose



Even-sensitized faults	Primary outputs			
	z ₁	z ₂	z ₃	z ₄
x ₂ s-a-0, s-a-1	1	1	1	1
x ₄ s-a-0, s-a-1	0	1	1	0
x ₅ s-a-0, s-a-1	0	1	1	0
x ₅ s-a-0, s-a-1	1	1	1	1
l ₁ s-a-0, s-a-1	0	1	1	0
l ₂ s-a-0, s-a-1	0	1	1	0
l ₃ s-a-0, s-a-1	0	1	1	0
l ₄ s-a-0, s-a-1	0	1	1	0
l ₅ s-a-0, s-a-1	0	1	1	0

(b)



Even-sensitized faults	Primary outputs				
	z ₁	z ₂	z ₃	z ₄	z ₅
x ₂ s-a-0, s-a-1	1	1	1	1	0
x ₃ s-a-0, s-a-1	0	1	1	0	0
x ₅ s-a-0, s-a-1	1	1	1	1	0
x ₆ s-a-0, s-a-1	0	0	0	1	1
l ₁ s-a-0, s-a-1	0	1	1	0	0
l ₂ s-a-0, s-a-1	0	1	1	0	0
l ₃ s-a-0, s-a-1	0	1	1	0	0
l ₄ s-a-0, s-a-1	0	1	1	0	0
l ₅ s-a-0, s-a-1	0	1	1	0	0

(b)

Fig. 1. Zero-aliasing space compaction via multiple parity signatures: (a) circuit with even-sensitized faults and zero-aliasing parity functions f_1 and f_2 and (b) its sensitization table.

columns denote primary outputs z_1, z_2, \dots, z_k , and whose rows denote the even-sensitized faults. The entry in the i th row and j th column is 1 if and only the fault corresponding to row i is sensitized to output z_j by some test pattern in the test set. Consider the example circuit of Fig. 1(a), whose sensitization table for a complete (exhaustive) test set is shown in Fig. 1(b). This implies that the minimal sensitization cover is either $\{z_2\}$ or $\{z_3\}$. Therefore, zero aliasing can be ensured if the space compactor generates two functions f_1 and f_2 , as shown in Fig. 1(b).

The above method for ensuring zero aliasing is complete, since it is always possible to find a sensitization cover for a set of (detectable) even-sensitized faults. However, q can be large, thus reducing the effectiveness of space compaction. For example, consider the circuit of Fig. 2(a). A sensitization cover for this circuit using a complete test set is always of size two, e.g., $\{z_2, z_5\}$ and $\{z_2, z_4\}$ are sensitization covers. Therefore, in order to ensure zero aliasing using only the sensitization cover, say $\{z_2, z_4\}$, we have to use a compactor with three outputs, $f_1 = z_1 \oplus z_2 \oplus \dots \oplus z_5$, $f_2 = z_2$, and $f_3 = z_4$. To alleviate this problem, we propose an alternative approach that examines fault sensitization to groups of primary outputs.

III. PAIRWISE SENSITIZATION

We turn next to the case where each parity function f_i is the exclusive-or of two primary outputs. This implies that

Even-sensitized faults	Primary output pairs									
	1 2	3 4	5 6	7 8	9 10					
x ₂ s-a-0, s-a-1	0 0	0 1	0 0	1 0	1 1					
x ₃ s-a-0, s-a-1	0 0	0 1	0 0	1 0	1 0					
x ₅ s-a-0, s-a-1	0 0	1 1	0 1	1 1	1 1					
x ₆ s-a-0, s-a-1	0 0	1 1	1 1	1 1	1 0					
l ₁ s-a-0, s-a-1	1 1	0 0	0 1	1 1	1 0					
l ₂ s-a-0, s-a-1	1 1	0 0	0 1	1 1	1 0					
l ₃ s-a-0, s-a-1	1 1	0 0	0 1	1 1	1 0					
l ₄ s-a-0, s-a-1	1 1	0 0	0 1	1 1	1 0					
l ₅ s-a-0, s-a-1	1 1	0 0	0 1	1 1	1 0					

(c)

Fig. 2. Obtaining multiple parity signatures using a pairwise sensitization cover: (a) circuit with even-sensitized faults, (b) its sensitization table, and (c) its pairwise sensitization table.

$x_{i,1} + x_{i,2} + \dots + x_{i,k} = 2, 2 \leq i \leq q$. In order to determine the value of q and assign values to the $x_{i,j}$'s, we construct a *pairwise sensitization table* (PST) by considering all possible pairs of primary outputs. The even-sensitized faults make up the rows of this table, while the $k(k-1)/2$ output pairs constitute the columns. We map the pair (z_i, z_j) to the column $s = (i-1)k + j - i(i+1)/2$ of the PST. This ensures that the PST columns $1, 2, \dots, k(k-1)/2$ denote the output pairs $(1, 2), (1, 3), \dots, (1, k), (2, 3), \dots, (2, k), \dots, (k-1, k)$, respectively. The entries in the PST are assigned values (0, 1) according to the following rule: $\text{PST}[r, s] = 1$ if and only if the even-sensitized fault corresponding to row r of the PST is sensitized by at least one test pattern to $z_i(z_j)$, but not to $z_j(z_i)$; otherwise, we set $\text{PST}[r, s] = 0$.

The next step is to find a set of columns of the PST that cover all the rows. The number of parity functions (apart from f_1) is equal to the number of columns in this cover. The values of the $x_{i,j}$'s are determined by the (z_i, z_j) pairs corresponding to this cover. Now, from the pairwise sensitization cover, we find that columns 7 and 9, i.e., the pairs, $\{z_2, z_5\}$ and $\{z_3, z_5\}$, are both pairwise sensitization covers. Therefore, $q = 2$, and we can ensure zero aliasing with only two signatures by setting f_2 to either $z_2 \oplus z_5$ or $z_3 \oplus z_5$.

We next pose the following question: Is it always possible to find a pairwise sensitization cover? The following theorem shows that such a cover always exists.

Theorem 1: For a circuit under test C with a test set \mathcal{T} and a set of even-sensitized faults F , it is always possible to find a pairwise sensitization cover for C .

Proof: A pairwise sensitization cover cannot be found only if there exists an even-sensitized fault $y_i \in F$ such that $\text{PST}[i][j] = 0$ for all j . This in turn is possible only if (a) y_i is sensitized to every primary output, and (b) the number of primary outputs is even. However, in such cases, we can detect y_i by letting $f_2 = z_i$ for some $i, 1 \leq i \leq k$. \square

It might be expected that the entries in the PST can be obtained by simply exclusive-oring the corresponding entries in the sensitization table. Let $\text{ST}[i][j]$ denote an entry in the sensitization table and let ϕ be the function that maps an (i, j) pair to a column of the PST. Then, we ask, "Is $\text{PST}[k][\phi(i, j)] = \text{ST}[k][i] \oplus \text{ST}[k][j]$?" We show below that this is not true. Suppose test pattern t_1 (and only t_1) sensitizes the fault f_k to primary output z_i . Also suppose test pattern t_2 (and only t_2) sensitizes f_k to primary output z_j ($i \neq j$). Then, $\text{ST}[k][i] \oplus \text{ST}[k][j] = 0$ but $\text{PST}[k][\phi(i, j)] = 1$.

Since the number of outputs of the space compactor is one more than the size of the PST cover, the effectiveness of the pairwise sensitization approach is measured by the size of the PST cover. Let the compaction efficiency η be defined as $\eta = (k - q) / (k - 1)$ where k is the number of primary outputs of the circuit under test, and q is the number of outputs of the compactor. Note that $0 \leq \eta \leq 1$, and $\eta = 1$ (0) if $q = 1(k)$. For intermediate values of q , η lies between zero and one. The higher the value of η , the more effective is the space compaction. We use this measure in the following section.

IV. EXPERIMENTAL RESULTS AND ANALYSIS

We performed a series of simulation experiments for the ISCAS-85 circuits to determine the size of their PST covers. The problem of finding sensitization and pairwise sensitization covers is equivalent to the NP-hard set-covering problem [6]. Therefore, we used a fast, heuristic covering procedure from [2]. The simulation results are listed in Table I. (We know from [2] that every SSL fault in c880 is odd-sensitized by the test patterns generated by COMPACTEST or ATALANTA.) The CPU times for these experiments on a SUN Sparc 10 workstation ranged from 1 to 421 s. For comparison purposes, the sizes of the sensitization covers are also listed. The results show that the use of the PST cover improves the compaction efficiency significantly.

These results also raise the following interesting question: Why is the compaction efficiency higher using PST covers

than using sensitization covers? We explain this below in probabilistic terms.

The sensitization table and the PST table are used to "cover" all the even-sensitized faults. These covers are called the sensitization covers and the PST covers, respectively. One way to compare the two approaches is to determine, for each case, the probability that an entry in the corresponding table is "1." Recall that $\text{ST}[i][j] = 1$ if there exists a test pattern in the test set that sensitizes fault f_i to z_j . Similarly, $\text{PST}[i][jk] = 1$ if there exists a test pattern that sensitizes f_i to $z_j(z_k)$ but not to $z_k(z_j)$.

Let \mathcal{E}_j^i be the event that a given fault f_u is sensitized by test pattern t_i to primary output z_j . Let the number of test patterns be m . We assume that the propagation of a fault effect to primary output z_j by test pattern t_{i_1} is independent of its propagation to z_j by t_{i_2} . This assumption is justified because the output of a combinational circuit depends only on the values applied at its inputs. This implies that for all $1 \leq i \leq m$, the events \mathcal{E}_j^i 's are independent. Let $P[\mathcal{E}_j^i] = p_j$, and let \mathcal{E}_j be the event that the column j entry of the sensitization table for the row corresponding to f_u is one. (We let p_j be the fraction of test patterns that sensitize f_u to z_j .) The expected number of 1's in this row of the sensitization table is $kP[\mathcal{E}_j]$, where k is the number of primary outputs, i.e., the number of columns in the sensitization table.

Now, $P[\mathcal{E}_j] = P[\cup_i \mathcal{E}_j^i]$, and since the events \mathcal{E}_j^i are not mutually exclusive, the probability of the union of these events is not equal to the sum of the probabilities of the individual events. Therefore

$$\begin{aligned} P[\mathcal{E}_j] &= mp_j - \binom{m}{2} p_j^2 + \binom{m}{3} p_j^3 + \dots + (-1)^{m-1} p_j^m \\ &= (-1)^{m-1} (p_j - 1)^m + 1. \end{aligned} \quad (1)$$

Next, let \mathcal{E}_{jk}^i be the event that for the given fault f_u , the column (jk) entry of the PST for the row corresponding to f_u is one. The probability associated with this event is given by $P[\mathcal{E}_{jk}^i] = p_j(1 - p_k) + (1 - p_j)p_k = p_j + p_k - 2p_j p_k$. Let \mathcal{E}_{jk} be defined in a similar manner as \mathcal{E}_j as above, and therefore

$$P[\mathcal{E}_{jk}] = (-1)^{m-1} (p_j + p_k - 2p_j p_k - 1)^m + 1. \quad (2)$$

We can easily see that $P[\mathcal{E}_{jk}^i] > P[\mathcal{E}_j^i]$ if $p_j < 0.5$. Similarly, $P[\mathcal{E}_{jk}] > P[\mathcal{E}_j]$ if $p_j < 0.5$. Thus if there are k primary outputs, and $p_j < 0.5$ for $1 \leq j \leq k$, then the probability of a "1"-entry in the PST is higher than the probability of a "1"-entry in the sensitization table.

If a fault were typically sensitized by a small fraction of test patterns to any given primary output, we would have $p \ll 0.5$, and therefore the compaction efficiency would be higher for multiple parity signatures using pairwise sensitization. To investigate this issue further, we carried out experiments for some of the ISCAS benchmark circuits. In each case, we considered all the even-sensitized faults for the reduced test sets generated by COMPACTEST and ATALANTA. We found that in almost all cases, an even-sensitized fault is sensitized to any given primary output by only a small fraction of test patterns, i.e., $p \ll 0.5$. For example, for the c432 circuit, p is at most 0.143 for both test sets. For c499, the maximum

TABLE I
SIMULATION RESULTS FOR THE ISCAS BENCHMARK CIRCUITS USING REDUCED TEST SETS GENERATED (a) BY COMPACTEST AND (b) BY ATALANTA

ISCAS benchmark circuit	Number of even-sensitized faults	Number of primary outputs	Size of sensitization cover	Size of PST cover	Compaction efficiency η		CPU time (sec)
					Sensitization cover	PST cover	
c432	6	7	2	1	0.67	0.83	1.02
c499	63	32	15	1	0.52	0.97	10.43
c1355	8	32	3	1	0.90	0.97	5.15
c1908	14	25	3	1	0.87	0.96	8.85
c2670	79	140	8	1	0.94	0.99	137.35
c3540	77	22	9	3	0.56	0.86	43.73
c5315	87	123	10	1	0.92	0.99	148.78
c6288	105	32	18	8	0.42	0.74	152.61
c7552	58	108	10	1	0.91	0.99	421.33

(a)

ISCAS benchmark circuit	Number of even-sensitized faults	Number of primary outputs	Size of sensitization cover	Size of PST cover	Compaction efficiency η		CPU time (sec)
					Sensitization cover	PST cover	
c432	7	7	2	1	0.67	0.83	1.08
c499	9	32	1	1	0.97	0.97	1.75
c1355	7	32	5	1	0.84	0.97	5.42
c1908	7	25	3	2	0.88	0.960	9.57
c2670	5	140	1	1	0.99	0.99	246.72
c3540	17	22	8	2	0.62	0.95	53.70
c5315	45	123	1	1	0.99	0.99	68.13
c6288	9	32	4	1	0.87	0.97	157.88
c7552	59	108	10	1	0.91	0.99	420.77

(b)

TABLE II
SIZES OF THE THIRD-ORDER SENSITIZATION COVERS FOR CIRCUITS FOR WHICH THE PST COVERS ARE OF SIZE GREATER THAN ONE

ATPG program	Circuit	Size of the third-order sensitization cover	Compaction efficiency η	CPU time (sec)
COMPACTEST	c3540	1	0.95	177.65
	c6288	1	0.97	111.05
ATALANTA	c1908	1	0.96	36.98
	c3540	1	0.95	132.77

value of p for tests generated by ATALANTA is only 0.031. There are 63 even-sensitized faults in c499 for the test set generated by COMPACTEST, of which there are ten faults that are sensitized to two outputs with $p > 0.5$. Note that there are 32 primary outputs, and thus for *most* (fault, output) pairs, $p < 0.5$.

We next compare our results to those obtained by Tarnick [10] and Table III. The method in [10] used pseudorandom test patterns and a sample of only 100 faults for each circuit. It is clear that our method, which directly addresses SSL faults, provides considerably higher compaction efficiency and is also computationally feasible.

The notion of sensitization covers can be extended to include larger groups of primary outputs. If j primary outputs are combined to form the parity signatures f_2, f_3, \dots, f_q , the sensitization cover thus obtained is the j th-order sensitization cover. Hence, the sensitization and PST covers are the first-order and second-order, respectively.

Returning to Table I, we see that c3540 and c6288 in (a), and c1908 and c3540 in (b) have second-order sensitization

TABLE III
COMPARISON BETWEEN THE NUMBER OF OUTPUTS q AND THE COMPACTION EFFICIENCY η FOR THE COMPACTOR OBTAINED USING PST COVERS TO THE ONE OBTAINED USING TARNICK'S GREEDY COVERING METHOD [10]

ISCAS benchmark circuit	Number of primary outputs	PST covers using ATALANTA test sets		Tarnick's greedy covering method [10]	
		q	η	q	η
c432	7	2	0.83	7	0
c499	32	2	0.97	6	0.84
c880	26	1	1	8	0.72
c1355	32	2	0.97	5	0.87
c1908	25	3	0.96	4	0.88
c2670	140	2	0.99	5	0.97
c3540	32	3	0.95	14	0.38
c5315	123	2	0.99	6	0.96
c6288	32	2	0.97	9	0.74
c7552	108	2	0.99	5	0.96

covers of size > 1 . Now, it might be expected that for these examples, there would exist third-order sensitization covers of size one. This is indeed the case, as illustrated by Table II. Every column of the third-order sensitization table corresponds to a group of three primary outputs (z_i, z_j, z_k) . If the test set contains a pattern that propagates an even-sensitized fault f_u to either exactly one or all three of these outputs, then the entry in the table corresponding to row f_u and column (z_i, z_j, z_k) is one.

REFERENCES

- [1] F. Brglez and H. Fujiwara, "A neutral netlist of 10 combinational benchmark circuits and a target simulator in Fortran," in *Proc. 1985 Int. Symp. Circuits Syst. (ISCAS)*, 1985, pp. 695-698.

- [2] K. Chakrabarty and J. P. Hayes, "Efficient test response compression for multiple-output circuits," in *Proc. 1994 Int. Test Conf.*, 1994, pp. 501–510.
- [3] K. Chakrabarty, B. T. Murray, and J. P. Hayes, "Optimal space compaction of test responses," in *Proc. 1995 Int. Test Conf.*, 1995, pp. 834–843.
- [4] K. Chakrabarty, "Test response compaction for built-in self testing," Ph.D. dissertation, Comput. Sci. Eng. Program, Univ. Michigan, Ann Arbor, 1995.
- [5] H. Fujiwara and A. Yamamoto, "Parity-scan design to reduce the cost of test application" *IEEE Trans. Computer-Aided Design*, vol. 12, pp. 1604–1611, Oct. 1993.
- [6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman, 1979.
- [7] H. K. Lee and D. S. Ha, "On the generation of test patterns for combinational circuits," Dep. Elect. Eng., Virginia Polytech. Inst. State Univ., Tech. Rep. 12_93, Dec. 1993.
- [8] I. Pomeranz, L. N. Reddy, and S. M. Reddy, "COMPACTEST: A method to generate compact test sets for combinational circuits," in *Proc. 1991 Int. Test Conf.*, 1991, pp. 194–203.
- [9] S. M. Reddy, K. K. Saluja, and M. G. Karpovsky, "A data compression technique for built-in self-test," *IEEE Trans. Comput.*, vol. 37, pp. 1151–1156, Sept. 1988.
- [10] S. Tarnick, "Bounding error masking in linear output space compression schemes," in *Proc. 1994 Asian Test Symp.*, 1994, pp. 27–32.

Krishnendu Chakrabarty (S'92–M'96) received the B.Tech. degree from the Indian Institute of Technology, Kharagpur, and the M.S.E. and Ph.D. degrees from the University of Michigan, Ann Arbor, all in computer science and engineering, in 1990, 1992, and 1995, respectively.

While at the University of Michigan, he was a Research Assistant at the Advanced Computer Architecture Laboratory in the Department of Electrical Engineering and Computer Science. Since 1995 he has been Assistant Professor of Electrical and Computer Engineering at Boston University, Boston, MA. He will join Duke University, Durham, NC, as Assistant Professor of Electrical and Computer Engineering in August 1998. His research interests are in computer-aided design of VLSI circuits and systems, testing, design verification, and fault-tolerant computing.

Dr. Chakrabarty is a member of Sigma Xi.

John P. Hayes (S'67–M'70–SM'81–F'85) received the B.E. degree from the National University of Ireland, Dublin, in 1965, and the M.S. and Ph.D. degrees from the University of Illinois, Urbana-Champaign, in 1967 and 1970, respectively, all in electrical engineering.

While at the University of Illinois, he participated in the design of the ILLIAC III computer. In 1970, he joined the Operations Research Group at the Shell Benelux Computing Center, The Hague, The Netherlands, where he worked on mathematical programming. From 1972 to 1982, he was a faculty member of the Departments of Electrical Engineering—Systems and Computer Science of the University of Southern California, Los Angeles. Since 1982, he has been a Professor in the Electrical Engineering and Computer Science Department of the University of Michigan, Ann Arbor. He was the Founding Director of the University of Michigan's Advanced Computer Architecture Laboratory. He is the author of over 150 technical papers, two patents, and five books, including *Computer Architecture and Organization*, 3rd ed. (New York: McGraw-Hill, 1998), *Layout Minimization for CMOS Cells* (New York: Kluwer, 1992; co-authored with R. L. Maziasz), and *Introduction to Digital Logic Design* (Reading, MA: Addison-Wesley, 1993). His current research interests are in the areas of computer architecture, computer-aided design, verification and testing, VLSI design, and fault-tolerant embedded systems.

Dr. Hayes is a member of the Association for Computing Machinery (ACM) and Sigma Xi. He was Technical Program Chairman of the 1977 International Conference on Fault-Tolerant Computing, Los Angeles, CA, and the 1991 International Computer Architecture Symposium, Toronto, Ont., Canada. He has served as Editor of various technical journals, including the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS and the *Journal of Electronic Testing*.