

Multi-Chip Module Test and Diagnostic Methodology

The development of a manufacturing test and diagnostic methodology for multi-chip modules as used in the IBM 4300 processor models involves determining the most attractive compromise among a number of conflicting factors: a) high test coverage, b) high diagnostic resolution, c) test generation, d) test equipment, and e) test application and diagnosis. This paper describes a set of solutions which were developed to create a high-volume, low-cost manufacturing test operation for the product in question. This paper examines the role of the testing methodology in productivity and product quality, details the diagnostic approach chosen, and provides an example of the overall manufacturing system performance achieved by analyzing a large module production sample.

Introduction

The multi-chip module (MCM) is a multi-layer ceramic substrate [1] used in the IBM 4300 processor models. The MCM carries signal and power distribution in printed circuitry form for some three to nine LSI chips, each containing about 1200 logic blocks or up to 5000 memory cells [2]. The MCM test facility is one of a number of test and diagnostic operations which see the product from its formative stages, at the end of a semiconductor processing line, through its service life in a customer field installation. It is one link in a chain of test systems which eliminate defective components and cooperate interactively to improve production throughput.

A production test model is shown in Fig. 1. Note that the earlier defects are discovered, the fewer are the processes and tests that parts go through, with a resulting higher throughput for a given amount of equipment. The least test equipment usage is obtained for a good part testing *good* at each stage and not ever requiring a retest. Note also that this model permits changes to be made prior to testing if it is discovered that certain types of defects consistently escape detection. This is done through process parameter adjustments, manufacturing procedure changes, and module repair, chip replacement, and retesting. By these means, both present and future yield improvements are realized while, at the same time, product costs are reduced.

Since the MCM test area services the entire manufacturing site, it is required to handle several hundred MCM part numbers. Logic block densities approach 10 000 on a single MCM part number, because both logic and storage array (RAM) chips can be mounted on a substrate. These circuits may be accessed for testing by up to 320 signal input/output (I/O) pins. In addition, there is a requirement to repair all failing modules. This can be accomplished through diagnostic analysis and repair of each defect.

The software operations structure of the test area is shown in Fig. 2. The test approach was geared to an already existing wafer test operation. This operation had been designed to handle large volumes of test data and individual test result data, as well as individual test and rework history data for each part passing through the manufacturing line.

Two principal test methods were considered for the multi-chip modules: chip-in-place testing (CIP) and through-the-pins testing (TTP). The first involves testing each chip as an isolated entity while it is mounted on the multi-layer ceramic (MLC) package (i.e., surrounded by other chips). Electrical isolation is achieved by means of built-in inhibit circuitry which disables the drivers of all chips not under test. The implementation of this test method with minimum impact on the chip logic required the addition of five

© Copyright 1983 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.




Figure 1 Production test model. At each stage of manufacturing, a test is performed before product is released to the next stage. If the test is not 100% effective, product with defects escapes, only to be caught by a later stage in the manufacture, thereby incurring additional rework and retesting.

additional printed circuit layers in the MLC package, together with the attendant process steps. These, in turn, had an impact on the ultimate manufacturing capacity, yield, and throughput.

The main advantage of the CIP test method was the method's inherent ability to reuse the tests produced for chip manufacturing so that there was no need to generate a new set of tests. These tests had a greater than 98% test coverage, and the diagnosis for use in both repair and process modification was excellent. On the other hand, it took about six to eight times longer to apply the chip manufacturing tests to all nine chips individually than to test them with a new set of




Figure 2 MCM product and test information interaction and flow. The release interface tape (RIT) contains data for building and testing the MCM. Test data are extracted, verified, and processed for use by the testers connected to the host through minicomputers. Test results from tested MCM product are forwarded to MCMDIAG for analysis and diagnosis. MCMDIAG supplies dispositioning instructions and reports for use by test personnel and diagnostic engineers.

tests specifically generated for use by the TTP method. Also, substrate fault coverage is poor in CIP testing because the inter-chip logic connections are not tested. Finally, there was the incentive to simplify the MLC package by removing the additional MLC layers required to make CIP testing possible.

TTP testing is performed by accessing the same electrical connections which are used by the components to communicate in the final system assembly, namely the module I/O pins. This method of testing differs from card testing in that it is usual in the latter case to augment I/O pins as test points with internal-to-the-card test points. Internal test points were not needed because design engineering planned to adhere to LSSD (level-sensitive scan design) rules [3], thus




Figure 3 MCM through-the-pins (TTP) test methodology evaluation. All expected defect types and detection methods were studied to determine their effects on costs and product quality level (PQL). Note: TOST = terminals, opens, and shorts tests; PPT = partial power-up tests.

avoiding potential testing costs traditionally associated with the use of "internal" test points.

The TTP method promised a lower test running time and the saving of five additional printed-circuit layers in the substrate. However, these savings had to be balanced against the requirement for an MCM test generation step which involved significant additional computation time. New tests had to be generated to provide the required test coverage to ensure that a specific product quality level (PQL) [4] is met, both at the component-manufacturing and card-assembly levels. Diagnostic program enhancements were also required to support the need for MCM repair and recovery by providing sufficient resolution and discrimination to minimize repair action calls.

The advantages proved greater than the disadvantages. As a result, the MCM manufacturing test approach was to adopt the TTP method for all multi-chip module packages. Additional efforts were made to minimize test generation and to create a test and diagnostic software/hardware system which could execute computer-controlled testing and diagnostics and generate repair procedures in an automated line (described by Mescia and Woods [5] and by Burgess, Koens, and Pignetti [6]).

The rest of the paper is devoted to 1) formulating the test approach, 2) providing details of the off-line diagnostic technique, and 3) providing an example of the overall manufacturing system performance achieved.

Formulation of test methodology

The principal consideration in formulating a multi-chip module test strategy is the ability to maintain the PQLs. This

means that defect escapes from one package level to the next must be closely controlled to meet the production objectives of removal of defects at the chip and module levels. Actual quality levels are monitored continuously by the quality engineering group, which analyzes the total life cycle of machine parts from the wafer stage through to a part number in a system product, over the supported lifetime of a system. Any time the PQL levels are exceeded, an immediate and extensive corrective action is generated. For this reason, great emphasis was placed on evaluating the test coverage and diagnostic implications of the TTP test method alternatives as they applied to the complete menu of MCM defect distributions (i.e., known and theorized faults). The alternatives are shown in Fig. 3.

MCM defect mechanisms involve defective chip and substrate components, as well as defects introduced in the bonding and assembly operations. Recall that the multi-chip substrate is already unit-tested during the manufacturing process; so are the logic and storage array chips. The possible new defects could be chip-bonding shorts and opens (also called controlled collapse chip connection (C4) shorts and opens [7]), input/output (I/O) logic defects, wrong chip part numbers, placement defects, wrong substrate part numbers, chip misalignments or misorientations, defective chips, missing chips, improper or missing wire bonds, improper or missing engineering changes, etc.

These defect mechanisms result in a wide variety of electrical failure indications. Some defects, such as an open C4, result in a single stuck-fault indication that is detectable and traceable using standard stuck-fault test coverage and diagnostic schemes. Other defects, such as a wrong chip part number or a chip misalignment, result in catastrophic failure




Figure 4 Probability of MCM-shorts escapes although TOST were applied for a signal I/O to total I/O ratio, f_s , of 0.64.

mechanisms that are incomprehensible from a pattern failure standpoint. Still others, such as C4-to-C4 shorts, can result in intermittent failures or disturb-type failure indications under normal stuck-fault pattern exercises.

In order to detect successfully and to discriminate such a wide variety of failure mechanisms, a two-tiered test structure was proposed; it involved both prefunctional tests and functional pattern (stuck-fault) tests. The implications of a number of alternative test methods were examined (Fig. 3) against the entire MCM part number set for the IBM 4300 processor models. Major bonding and assembly test defect coverages were evaluated by creating a probabilistic model and examining the coverage obtained by prefunctional and also functional test schemes. A typical test coverage analysis is illustrated next for the case of MCM net-to-net shorts.

The prefunctional test efficiency is illustrated in Fig. 4. In this case, the probability of bonding defects (e.g., C4-to-C4 shorts) and substrate defects (e.g., net-to-net shorts) escaping a terminator opens/shorts test (TOST) is examined. The probability of escape p is plotted as a function of the number of internal pins and the net loading tolerances:

$$p = f_s^2 [(1 - 2f) X^2 + 2fX], \quad (1)$$

where f_s is the ratio of the signal I/O to the total I/O (a constant in each chip configuration), X is the percentage of C4 joints not accessible to module pins (i.e., the percentage of internal pins), and f is the probability that the combined net load of the two shorted nets does not exceed the allowed loading tolerances of either net. The projected internal-pin distribution points (mean, median, and worst-case) for the part number set were also plotted on these curves to help arrive at the TTP test decision.

A similar analysis was performed for the probability of MCM shorted nets escaping both the prefunctional and functional tests, as well as a *reduced* set of functional tests (Fig. 3). The proposed reduced set of functional tests fault-simulate only the chip I/Os visible to the module I/Os; all internal failures are ignored. Thus the effects of, as well as the number of, chip I/Os internal to the modules have to be accounted for. In this case, the probability of escape is a function of two additional factors, P_m and N . P_m is the probability of functional failures due to an inadvertent measurement of a stuck-fault failure in a section of logic theoretically independent of the net being tested but, by virtue of the short, actually being electrically dependent. N is the number of stuck-fault pattern sets executed through these inadvertently dependent shorted nets. The detailed analysis is beyond the scope of this paper.

As a result of such evaluations, a test methodology was evolved. We would only use prefunctional tests, called terminators, opens, and shorts tests (TOST), single stuck-fault functional patterns consisting of standard LSSD shift register tests, and the *reduced* set of stuck-fault tests. By allowing test pattern generation to select the reduced set of faults (i.e., chip and module I/O faults only instead of the full complement of stuck-faults), test preparation costs were cut to one-third and the test application time was cut in half. Additional savings were realized because individual tests covered more than a single chip. Overall, the implemented TTP method permitted us to increase the MCM throughput by between six and eight times over the CIP method, using the same test equipment.

MCM diagnostics

In the MCM manufacturing environment, two distinct types of diagnostics are required. The first type, repair diagnostics, is designed to determine the cause of a given failing module with the goal of repairing the defects. The second type, failure mode diagnostics, is designed to identify repetitive failure mechanisms which result in major short- or long-term yield losses, with the goal of yield improvement. To satisfy these requirements, programs have been developed to perform these functions efficiently in a high-volume production environment.

• Repair diagnostics

Unlike the scrapping practice in chip manufacturing, a module (which contains up to nine LSI chips and a multi-layer ceramic substrate) cannot simply be discarded if it fails final testing. The cause of the failure must be identified so that the module can be repaired and data collected for future yield improvements. Because of the large volumes, the large number of different modules, and the circuit complexity, it is desirable to perform the diagnosis automatically. For this purpose, a program called MCMDIAG was developed to isolate the causes of defective modules, issue rework instructions for repair, and store failure information in a tracking system data base. The plan was to perform the diagnosis independently of the tester, thereby causing no impact on tester throughput (see Fig. 2). For every failing module, the failure data were to be collected from the TOST and from the functional tests for later use in the diagnosis. Unlike the strategy used on the larger TCM modules [8], no additional tester probe information or on-line retesting was to be required for the MCM modules.

A unique feature of this form of diagnosis is that it does not rely on a precalculated fault dictionary. Instead, a post-test fault simulation is performed to identify stuck-faults that can explain the observed failures. This fault simulation is performed by a fault simulator that was designed specifically for MCM fault diagnosis. It differs from classical stuck-fault simulation in that it concentrates exclusively on trying to identify the circuits that are capable of providing the failure indications observed by the tester.

The fault simulation is performed on the logic structure of the entire module. Since the chips and the substrate are extensively tested before they are bonded during module assembly, the bonding between the chips and the substrate is the most likely source of defects. Logic blocks representing the bonding operation are added to the logic structure being simulated so that the effects of bonding failures can be simulated and the observable resultant failure patterns can be produced. Figure 5 illustrates how the addition of chip I/O logic blocks can provide faults which would otherwise not be simulated. For example, in Fig. 5(a), block A is an off-chip driver leaving the chip at pin X and going to pins Y and Z. Pin Y fans out to blocks B and C, while pin Z only goes to block D. In the original logic simulation structure, no fault represents an open on pin Y. The addition of chip input block Y to the simulation, as shown in Fig. 5(b), does provide such a fault. Both input and output blocks are simulated as single-input AND circuits.

• Failure mode diagnostics

A wide variety of potential defects are associated with chips bonded on MLC substrates. Given failing patterns from these defects, the program must not only be capable of




Figure 5 Method of simulating pin and electrical net failures that cross chip boundaries indicated by dashed lines. (a) Original logic structure showing chip I/O pins tied to the same electrical net; (b) modified simulated logic structure to isolate bonding problems associated with the illustrated interchip net.

selecting the chip(s) whose replacement would be most likely to correct the problem, but also of classifying the type of defect. Furthermore, consideration must be given to maximum resolution, multiple defects, and possible tester problems. To achieve this aggressive goal, MCMDIAG performs the following six analyses.

TOST failure analysis TOST failure analysis is performed on each module pin that fails TOST. The change in parallel resistance is calculated and the most probable cause is determined by considering potential open and shorted pins on the net. For modules with only one TOST failure, the rework instruction is to replace the chip containing the pin most likely to be faulty. When more than one TOST failure occurs, an evaluation is performed to identify the minimum set of chips that can explain all detected TOST failures. The chip which can account for the most TOST failures is selected. This method is repeated until all TOST failures are accounted for.




Figure 6 A shift register across multiple chip boundaries showing module pin connections required to isolate shift-register problems to a single chip.

TOST analysis is particularly effective in diagnosing chips that are misaligned or misoriented, or that have the wrong part number. However, for single-pin defects, it is limited to those which are connected to module pins in which the change of resistance due to the defect exceeds the product tolerance.

Precalculated fault dictionary analysis This analysis is used to determine single stuck-faults that can explain the first failing pattern. There are severe limitations to this analysis because only faults for the chip and module input-outputs are simulated. Thus, not all potential stuck-faults have been faulted, and fault information is only available for the first failing pattern. When results are available, the resolution is usually poor, and it is of little value for multiple defects. However, the data are available and may be used as another option whenever other analysis methods fail to identify a defect adequately.

Traceback analysis This is used in attempts to identify potential defective chips that can explain failures by tracing through the logic connections for all failing primary outputs (POs) and shift register latches (SRLs). This method has little diagnostic value because of its poor resolution and inability to distinguish defects associated with pins from those associated with chips. However, it is extremely effective in selecting a minimum number of faults to be simulated. Faults are selected that are logically connected to failing POs and SRLs of the first failing pattern. Consideration is given to inversion levels, and only those stuck-faults which are consistent with the failing states are selected. Occasionally, however, reconvergent fan-outs result in selecting both stuck-at-0 and stuck-at-1 faults for a given logic circuit.

SRLs are used as boundaries for terminating a path during traceback analysis. If the failure occurs after the test has exercised an LSSD system clock, a traceback path is allowed to go through at most two SRLs. A system clock in LSSD is a clock which brings data from the combinational part of the structure into the SRL [9]. The fault may therefore be on the other side of the SRL. Two SRLs are allowed in the path because data going into the SRL may be from another SRL which was correctly clocked. If no system clock was exercised prior to the failure, all traceback paths terminate at SRL boundaries.

Full fault dictionary analysis The full fault dictionary is produced by simulating a list of pin and internal stuck-faults of the logic that has been minimized by traceback analysis. A list of stuck-faults that can explain the first failing pattern is produced and all other faults are eliminated. These faults are then simulated over all failing patterns.

If simulation of a new failing pattern eliminates all potential faults, results are given up to that pattern, and analysis is repeated with the unexplained failing pattern becoming the new first failing pattern. The additional passes are necessary when more than one defect is causing the failures. The simulation over all failures improves the resolution and provides the capability for diagnosing multiple defects. In some cases, resolution may be lost because only the failing patterns are simulated. However, this is an advantage when diagnosing shorts and intermittent faults, where additional failures would be expected from the corresponding stuck faults. A different analysis is performed when it is determined that the defect is in a shift register.

Shift-register-latch fault analysis This is invoked when a shift-register output pin fails the shift-register test. In this case, it becomes impractical to perform a full fault dictionary analysis. The normal simulation method of directly filling the shift register is no longer valid in this situation, and traceback cannot stop at latch boundaries, creating an unacceptably large list of faults for simulation. Instead of a stuck-fault approach, an attempt is made to isolate the defective chip by observing points at which the shift register crosses chip boundaries. This requires that these points be connected to module pins which are measured during the shift register test. The defective chip will be the chip whose scan-out line is connected to the module pin that fails farthest from the shift-register scan-out line. Figure 6 shows a shift register which goes through six chips with each chip scan-out connected to a module pin. If only pins 5 and 6 fail during the shift-register test, the analysis concludes that chip 5 is the defective one.

Array pattern analysis This is invoked when a full fault dictionary analysis fails to find the defect and storage arrays are imbedded in the logic. Because memory elements are represented in the simulated logic as individual logic blocks, the fault dictionary approach does not provide good data on the causes of a storage array failure. However, information is available to correlate patterns with the storage array chip being exercised. Thus, a failure occurring while such a pattern is used points to the array as the most likely failure.

Automated diagnosis

The results of the six individual analyses are correlated to determine the defect classification and rework instructions with the best possible resolution. Analyses with the highest probability of success and defects with the highest probab-

ity of occurrence are given the greater weight in making this decision. The order of preference of selecting which analysis results are used is as follows:

1. Multi-TOST failure analysis.
2. Single TOST failure analysis with full fault dictionary correlation.
3. SRL failure analysis.
4. Full fault dictionary analysis (potential pin faults).
5. Full fault dictionary analysis (only internal chip faults).
6. Precalculated fault dictionary analysis.
7. Array pattern analysis.

Analyses are only considered when their rework instructions call for replacement of three or fewer chips.

The analyses were designed such that at least one analysis would be successful in identifying a typical defect. Occasionally, no analysis is successful or the resolution of the diagnosis is not adequate. In that case, the failing module, together with the failure analysis reports, is sent to the diagnostic engineering group for diagnosis (bottom right-hand side of Fig. 2).

Before the diagnosis is completed, the failure data are analyzed to determine if the failure could have been caused by a problem in the tester itself. All TOST measurements are checked to ensure that they indicate reasonable values on the pins. This identifies gross test problems and guarantees that the proper test program was applied. A failure here would call for a retest of the module. Whenever a diagnosis points to a module pin, a warning is given which indicates a potential module pin-to-tester connection problem.

A tracking system [5] has been installed to monitor the status of all modules in the test area and to keep a history of all modules tested. When MCMDIAG diagnoses a failing module, the suspected defective chip and the defect classification are automatically stored. To this entry are added visual verification results. A module may undergo several rework cycles, and for each test pass, information is accumulated. The tracking system can determine total yield or yield loss due to any given defect type for any given module for any given time period. Upon completion of diagnosis, a routing instruction is issued along with the rework instruction. The tracking system is automatically updated with the results of the diagnosis. If a module has failed identically in a prior test, it is routed to diagnostic engineers for analysis.

While a computer program may not always make decisions as intelligently as test engineers, the program is an encoding of the best engineering decisions to date, and is clearly cost-effective. It has the advantage of giving consistent results for equivalent failures. This uniformity of results makes the archival records more reliable. Despite extensive

Table 1 Classification of failing modules (%) encountered during the study using MCMDIAG.

Automatically diagnosed correctly	76
Not automatically diagnosed correctly	7
Required retesting before correct diagnosis could be made	16
Other (mechanical, etc.)	1

analyses and logic structures of up to 10 000 logic blocks, the average System/370 Model 168 time per diagnosis is only fifteen seconds.

Evaluation of the MCM test approach

In the first quarter of 1980, a study was conducted on a large sampling of MCM production modules in order to identify the causes of failures for all modules that failed the final test on the first pass. The objectives of the study were to determine the types and occurrences of yield detractors and to evaluate the effectiveness of the MCM test and diagnostic system. Table 1 shows how the failing modules were classified by the study using MCMDIAG.

There were four levels of investigation in determining the causes of failures. If a particular investigation was successful in identifying the defect, the defect type was recorded along with the MCMDIAG results. The unresolved modules were then sent to the next level of investigation. The first level of investigation involved a visual analysis of suspected bad chips. This could be used to identify chip placement problems, bonding problems, and physical damage to chips. The second level of investigation, electrical measurements on suspected faulty circuit nets, could be used to identify substrate problems, parametric problems on chips, and tester contact problems. The third level of investigation involved retesting of suspected bad chips in the chip test area. This allowed for identifying defects within the chips. The final level involved individual analysis of suspected bad chips. Modules that remained at this level were thought to be due either to misdiagnosis or to chip-test escapes. Extensive analysis revealed that in every case the suspected bad chip was defective.

All failing modules whose components were available for investigation were analyzed and the causes of failure were identified. Table 2 gives a breakdown of the defect categories that were observed, together with their percentage of occurrence. One of the results of the study was that no significant difference was observed between the number of chip rework cycles or the number of chips being pulled using MCMDIAG and those of earlier manual techniques. It was found that automatically issued rework instructions were successful 90% of the time, and that of the 10% of the cases

Table 2 MCM defect categories.

Defect in	Type of defect
Fabrication	Chip placement Bonding Mechanical
Device	Any
Substrate	Opens/shorts Dirty pins
Components	Alignments Missing
Tester	Mistest Contacts

which were not automatically diagnosed, MCMDIAG referred 4% of the cases to engineering for action, 5% were SRL analysis problems, and 1% were miscellaneous problems, possibly never automatable. The SRL analysis problem has been solved by implementing a rule requiring chip scan-out pins to be attached to substrate module pins. Of the 90% of cases which were automatically diagnosed, it was found that in 13% of the cases the failure could not be repaired with a chip replacement alone. While the failing net was identified, the repair sometimes also required rework on the substrate. Finally, in no case did the rework instruction fail to identify the proper chip if a failure really resided on a chip.

Summary

One major goal in formulating the manufacturing test approach for the MCM package was to maximize fault coverage over the entire defect fault menu while minimizing test generation and test application. Another major goal was to effect present and future yield improvements by means of component repair recovery and process parameter adjustments. This required a cost-effective diagnostic method with sufficient power and acuity to diagnose accurately the cause of module failures with sufficient resolution in most cases. This diagnostic method was automated and performed with little or no impact on tester capacity and throughput.

These approaches were implemented in a hardware/software MCM manufacturing test system over a period of about two and a half years. The productivity and efficiency of this system was tested and evaluated by means of a comprehensive engineering study. This study was performed on a large production run sample of multi-chip modules. Tester throughput, test times, diagnostic efficiency, turnaround times, product failure mechanisms, defect escapes, and PQL implications were all taken into account.

Ninety percent of the defects encountered during the study proved to be automatically diagnosable. This number could be increased by some five percent once all inter-chip shift register chip input-output pins were connected to module pins (as required by the new design rules). We believe it is possible to approach an automatic diagnostic effectiveness of 100% once we are able to automatically distinguish between chip-only defects and chip-with-substrate defects.

References and note

1. A. J. Blodgett, Jr., "A Multilayer Ceramic Multichip Module," *IEEE Trans. Components, Hybrids, Manuf. Technol.* CHMT-3, 634-637 (1980).
2. The term logic block is used instead of gate or circuit because the simulator considers a 2- to n -input logic gate as one logic block.
3. H. C. Godoy, G. B. Franklin, and P. S. Bottorff, "Automatic Checking of Logic Design Structures for Compliance with Testability Ground Rules," *Proceedings of the 14th Design Automation Conference*, New Orleans, LA, June 1977, pp. 469-478.
4. D. S. Cleverley, "Product Quality Level Monitoring and Control for Logic Chips and Modules," *IBM J. Res. Develop.* 27, 4-10 (1983, this issue).
5. N. C. Mescia and C. D. Woods, "Plant Automation in a Structured Distributed System Environment," *IBM J. Res. Develop.* 26, 497-505 (1982).
6. Robert M. Burgess, Kathleen B. Koens, and Emil M. Pignetti, Jr., "Semiconductor Final Test Logistics and Product Dispositioning System," *IBM J. Res. Develop.* 26, 605-612 (1982).
7. L. F. Miller, "Controlled Collapse Reflow Chip Joining," *IBM J. Res. Develop.* 13, 239-250 (1969).
8. P. L. Barry, "Failure Diagnosis on the LT1280," *IBM J. Res. Develop.* 27, 41-49 (1983, this issue).
9. E. B. Eichelberger and T. W. Williams, "A Logic Design Structure for LSI Testability," *Proceedings of the 14th Design Automation Conference*, New Orleans, LA, June 1977, pp. 462-468.

Received October 26, 1981; revised August 27, 1982

James J. Curtin IBM General Technology Division, East Fishkill facility, Hopewell Junction, New York 12533. Mr. Curtin is an advisory engineer working on bipolar array test and diagnostics. Since joining IBM in 1974 in East Fishkill, he has worked in LSI testing at the device, single-chip, and multi-chip module levels. His work has included test implementation, testing philosophy and strategy, and work in areas of prediction of quality levels. He received his B.S. in 1972 and his M.S. in 1974, both in electrical engineering from Rensselaer Polytechnic Institute, Troy, New York. Mr. Curtin is a member of Eta Kappa Nu and Tau Beta Pi.

John Waicukausk IBM General Technology Division, East Fishkill facility, Hopewell Junction, New York 12533. Mr. Waicukausk is located at East Fishkill, where he is an advisory engineer currently working in design and development of multi-chip module diagnostic systems using software techniques. He joined IBM in 1974 at East Fishkill, where he designed and developed the chip diagnostic system. Mr. Waicukausk received a B.S. in chemistry from Southern Illinois University, Carbondale, in 1970 and an M.S. in physical chemistry from the University of Illinois, Urbana, in 1971.