Gross delay defect evaluation for a CMOS logic design system product

by O. Bula
J. Moser
J. Trinko
M. Weissman
F. Woytowich

Randomly occurring gross delay defects allow chips to pass full stuck-fault testing at both wafer and module levels, but cause them to fail when operated at system speeds. This paper describes the results of an experiment designed to determine the actual delay defect component of shipped product quality level (SPQL) for a CMOS combination standard cell/gate array design system. More than 60 000 modules, representing chips from the same IBM computer system, have been delay-tested using the technique presented in this paper. The test technique uses the stuck-fault patterns for level-sensitive scan design (LSSD) product. The stuck-fault patterns are modified or "twisted" according to specific algorithms to propagate transitions through paths just prior to the output measure. The patterns are applied at system speed timings provided by the chip designers. Any gross delay defect present in a tested path causes a fail. The failing modules were characterized to determine the size of the delay defects. Failure diagnostics were performed on the defective modules by using existing stuck-fault diagnostic tools and a development version of a transition fault simulator. These were sent to physical failure analysis for delayering, visual verification, and electrical characterization. A summary of physical defects which produced gross delay defects is presented.

Introduction to shipped product quality level (SPQL)

Shipped product quality level can be defined as the upper limit on the number of modules shipped to the customer which fail to function properly in the card or system for reasons other than a chip design error or mishandling and misassembly after shipment. There are three reasons why a defective module could reach the customer in the card or system assembly plant:

1. The correct test procedures were not followed, or the module was not tested at all; this is known as *imperfect test.*
2. The dc[1] stuck-fault test did not achieve test coverage of 100% of the possible faults on the chip, resulting in a *dc untested* defect.

[1] The static or stuck-fault test is frequently referred to as dc test.

Table 1 Product descriptions for SPQL test vehicles.

Parameter	*Product part name*					
	SAdapt	*IUnit*	*SCU1*	*CAdapt*	*SCU2*	*Proc*
Chip size (mm)	7.5	7.5	7.5	9.4	9.4	9.4
Number of I/Os	166	166	166	183	182	229
Logic cell density (%)	83.2	71.0	67.8	49.9	55.8	87.5
RAM cell density (%)	0	0	0	0	27.4	4.5

3. A defect exists on the chip which allows it to pass dc stuck-fault testing but causes it to fail when operated at system speeds; this is known as *ac*[2] *escape.*

These are the three components of SPQL. A high SPQL imposes a higher rework cost on the box and system assembly plants. Generally, the "good" modules shipped to the customer are soldered directly to a card and are subjected to a card- or system-equivalent test. When the card fails, time and money are spent determining which module failed, replacing it, and retesting. Escapes can pass the card- or system-equivalent test, which is not exhaustive, only to cause a system to fail when the customer is using it.

Both the imperfect-test and dc-untested categories have a high probability of being caught by a card- or system-equivalent test at the next higher level of assembly, because both imply that the fault is manifested as a net stuck-at either a 1 or a 0. These categories can both be driven to low levels with some basic techniques. The minimum required dc test coverage for IBM Burlington logic design products is 97.5%. By driving the dc test coverage as close to 100% as possible, the number of dc untested fails can be greatly reduced. The imperfect test level can be evaluated with retest techniques which generally involve some type of guard band to account for tester tolerance, tester drift, and tester-to-tester differences.

The ac escape category has the greatest chance of reaching the customer; this type of defect is the most challenging to detect. The experiment described in this paper was designed to measure the ac SPQL of a CMOS logic design system and to determine the dominant physical failure mechanism contributing to the measured level of ac escapes.

Product/technology overview

All product tested in the ac experiment described was designed using the CMOS standard cell/gate array (CSC/CGA), a design system which offers combined gate array and standard cell logic capability with embedded RAMs and ROSs [1, 2]. The Engineering Design System (EDS) allows users to assemble a chip from a library of over 1000 unique logic books and macros, latches, embedded growable RAMs, and embedded personalizable ROSs. Offering multiple chip sizes, CSC/CGA ranges from 13 000 to 60 000 equivalent two-way NAND circuits, packaged as either single-chip modules (SCM) or multi-chip modules (MCM). The design system supports physical and electrical design as well as design verification, timing analysis, test generation, and release.

Overview of the ac SPQL experiment

The ac tests used in this experiment were created by modifying the dc stuck-fault patterns generated for LSSD-designed [3] product. The patterns were modified or "twisted" to propagate transitions through paths just prior to the output measure. The system speed timings were provided by the designers. For delay defects to be detected, the increase in delay for the logic blocks or global wiring nets to propagate transitions must be large enough to exceed the cumulative path delay beyond what can be tolerated within the system. These types of defects, referred to as gross delay defects, are caused by random process defects present during a chip's manufacture.

More than 60 000 dc good modules, representing chips from the same IBM computer system, have been ac delay-tested using the "twist" technique. **Table 1** shows the characteristics of the six primary test vehicles.

Test and reliability screen sequence

Each module in the experiment sample was packaged using devices fabricated in the IBM Burlington manufacturing process line. The chips were subjected to a manufacturing wafer test and reliability voltage screen on automatic test equipment (ATE) systems. The devices were then packaged into modules followed by a burn-in screen. A certified module final test was applied to each module after the burn-in process.

To this normal manufacturing test sequence, a one-corner (V_{dd} = 4.5 V) gross delay test was appended. The unique delay fails were identified, and the failing test patterns were saved for diagnostics. These modules were taken to the card manufacturing site for card-equivalent and/or system-equivalent test. Upon return to IBM Burlington, the modules were characterized to see what increase and changes in timings caused them to begin passing the gross delay test. This provided valuable information on the size of the delay defects. The overall test sequence is summarized as follows:

1. Subset of wafer parametric test.
2. Wafer-level voltage screen.
3. Wafer-level full parametric testing and dc stuck-fault test.

[2] Although delay defects may only be a subset of ac faults, in this paper the terms ac and delay defect are used interchangeably.

4. Module-level burn-in.
5. Module-level full parametric testing and dc stuck-fault test.
6. ac delay test.
7. ac timing characterization testing.

Each of the manufacturing tests and screens is examined in more detail in the following subsections, with emphasis placed on the types of manufacturing defects they were designed to find.

Parametric test performed at wafer and module levels
The wafers were parametrically tested to ensure input/output (I/O) integrity. These tests use the ATE's analog parametric measurement unit to either force current/measure voltage or force voltage/measure current to reject opens and shorts, excessive current leakages caused by gate-oxide defects or metal shorts, disconnection of pull-up or pull-down resistors, and high quiescent drain currents. Performance-screen ring-oscillator periods were measured as an indicator of chip speed. A goal of this test is to screen out overall chip process problems that degrade circuit performance over the chip in a systematic way.

Voltage screen performed at wafer level
At the wafer level, a reliability screen at elevated voltage is applied using a subset of the dc EDS-generated logic LSSD patterns and the embedded-array EDS array pattern generator (APG) tests. This screen has proven to remove typical CMOS reliability failure mechanisms such as shorts between various levels of metal, including M2 shorted to M1, M1 shorted to M1, and M2 shorted to M2, as well as M1 shorted to polysilicon and gate-oxide defects [4].

Stuck-fault test performed at wafer and module levels
After voltage screen, the chips receive the parametric tests followed by a dc application of the EDS-generated LSSD logic stuck-fault and embedded-array APG patterns. This approach ensures that most of the interconnections of the devices are not catastrophically shorted to ground, shorted to V_{dd}, or completely open.

Burn-in performed at module level
After module build, a burn-in stress at elevated temperature and voltage is done to reduce module early-life failure rates. Burn-in has been shown to accelerate latent defects culminating in the form of metal shorts, M1-to-polysilicon shorts, metal opens, gate-oxide defects, and cracked chips. A certified module final test is then executed, exercising and measuring the device under test parametrically and with the EDS LSSD logic and embedded-array APG patterns. This test culls out defects due to module build problems and burn-in-accelerated mechanisms.

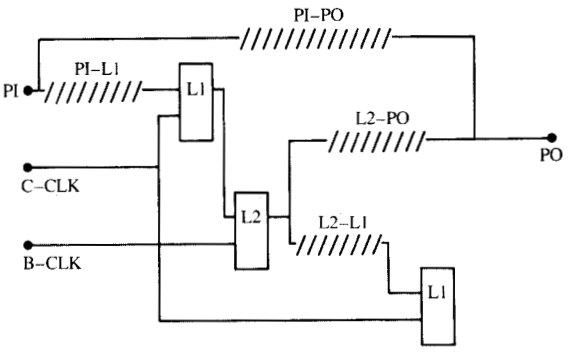

Figure 1
Paths within LSSD logic; slashes represent combinational logic.

Gross delay tests for LSSD designs

Using LSSD creates a more manageable combinatorial logic testing challenge from complicated sequential logic designs by introducing a scan path of latches to act as internal chip I/O for the control and observation of the logic blocks comprising the design. A dc benefit of following LSSD guidelines is that product test is independent of part-specific delay information or other transient effects. However, LSSD delay-test timings cannot always be run at system speeds, because LSSD scan paths may be slower than the system paths.

Many transition faults [5], defects which delay a rising or falling transition on a net but which do not create an outright stuck-at-logic state, are undetectable by the dc stuck-fault LSSD test. This is due to the test-pattern sequences generated by the current LSSD test-generation schemes and the generic, lengthy tester cycle times typical of a design system test methodology. The predominant system path on most LSSD designs is the L2 latch–combinatorial logic–L1 latch path. **Figure 1** shows the typical paths found within LSSD logic designs. The B-clock controls the L2, and the C-clock does the same for the L1. Transition faults, such as almost open metal line defects, may lie in such paths. In the typical operation of the LSSD double-latch design, the B-clock launches data from the L2 latches, followed by an allowance of time for data propagating from the L2 through logic to the data port of an L1, and is completed by a capturing C-clock latching data into L1. If valid data does not arrive at the time of the C-clock, erroneous data has been introduced into the system.

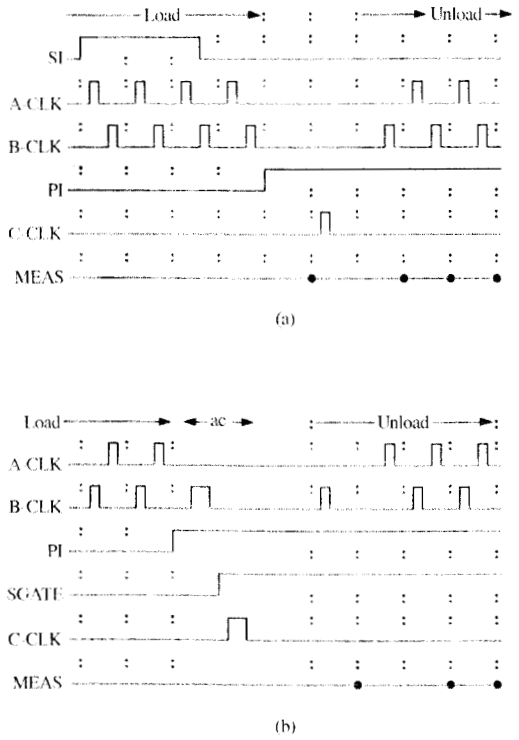

Figure 2

Comparison of a typical LSSD test (a) and an ac LSSD test (b).

A technique has been used by several semiconductor manufacturers to help sensitize transition faults along the L2–L1 path [6]. Basically, the last shift B-clock of the scan-in or load is delayed or merged in the same tester cycle with the C-clock. This is done via postprocessing of the EDS test-generation output, called an *all-events trace* (AET). This postprocessing has taken the form of modifying or inserting AET patterns into the original EDS AET to create a gross delay test AET. **Figure 2** shows a timing diagram comparison of a typical LSSD test [part (a)] and one which has been modified to provide an ac test [part (b)]. IBM developed an expert system approach to this transformation process. Within a test cycle, the B-clock is first applied, and is followed by the C-clock. The time between the rise of the B-clock and the fall of the C-clock must allow for the longest delay produced by a combinatorial logic path, taking into account both system and LSSD maintenance paths. Maintenance paths are logic paths on the chips which are exercised in manufacturing during test or in the field during maintenance analysis, but not exercised during the normal function of the machine. Also, the B- and C-clocks used in these tests must be of sufficient width not to violate set-up-and-hold timings of the latches.

The reversal of the shift A-clock and B-clock timings within the tester cycle, the addition of the B-clock to the same cycle as the C-clock, and the extra A-clock just before the scan-out or unload operation are the key "twists" to create the L2–L1 delay test for this experiment.

Gross delay test algorithms

Tests were created for other design paths that can show the same susceptibility to delay defects as the typical L2–L1. Each part's stuck-fault AET was transformed by an expert system to create a gross delay test AET for each of the following path types:

1. L2 latch–combinatorial logic–L1 latch.
2. Primary input (PI)–combinatorial logic–L1 latch.
3. L2 latch–combinatorial logic–primary output (PO).
4. Primary input–combinatorial logic–primary output (SCU1, SAdapt, and Proc only).

The timings used to test each path were generated by the designers. Special adjustments were made to account for the test environment. The timings, the transformed AET for each path, and the normal manufacturing test sequence were integrated into one test program per chip part.

The following algorithms were used to create the other delay tests via AET expert system manipulation.

Twist method for PI–PO paths

Only AET patterns where primary outputs are measured with no applied clocks are considered. A preceding pattern is inserted that complements all noncritical primary inputs not called out for stimulus by the AET pattern. Critical pins are inputs such as scan gates and test inhibits. The complemented inputs are then returned to the correct AET logic state in the cycle with the primary output measures. This delay test essentially checks combinatorial logic paths unbounded by latches.

Twist method for L2–PO paths

This test begins like the L2–L1 technique for launching data from the L2 latches. The shift register latch (SRL) load is modified to reverse A-clock and B-clock. The last SRL load B-clock is delayed, with all intervening primary output measures dropped. When the delayed B-clock occurs, the primary output measures are restored. Data are launched from L2, propagate through combinatorial logic, and are captured and measured at the primary outputs.

Twist method for PI–L1 paths
A pattern is inserted before any AET pattern in which a C-clock is pulsed. The inserted pattern complements all nontest function inputs. The time from the primary input stimuli to the C-clock application at the latch is critically controlled. This checks for transition faults on primary-input-to-L1 latch paths.

Card- and system-equivalent test

A prototype version of the computer system was used to do the card- and system-equivalent testing on the experiment fails. The power supply was set to 4.5 V to provide worst-case timing conditions. Repluggable sockets were installed on the processor card to allow easy changeability of modules. Experiment fails were plugged into a known-good low-end processor card. Only one failing module was inserted at a time. The card test executed the same microcode run during the initial program load. Scan chains, card functions, registers, and arrays were exercised. Although the card test had the diagnostic capability of identifying the failing module and card data bus, internal module diagnostics were beyond its capability. If a module passed the card test, it was further subjected to a system-equivalent test. In this test, the control store is loaded with run-time code, and major subsystems are tested. These functional tests check the workstation, DASD, diskette, main memory, and processor subsystems. The system-level diagnostics were subsystem-specific, and were unable to pinpoint the fail to a module.

ac SPQL rate calculation

The measured ac SPQL fallout rate is calculated in **Table 2** as a normalized function of the target ac SPQL, represented by the variable T.

This ac fallout rate must be adjusted for less than 100% ac test coverage in order to determine the ac SPQL. The average test coverage number has been determined to be 75%, based on transition-fault simulation of the test patterns for the six parts. Another factor to be considered is whether any false ac fails occurred in the sample. When the SAdapt modules were taken to the card manufacturing site for the system-equivalent test, it was determined that ten of the 31 modules exhibited a "glitch" due to a marginal design problem. These ten modules were not counted in the ac SPQL rate calculation. Taking these factors into account, we have the following equation:

$$ac\ SPQL = \left[\frac{(97 - 10)}{Total}\right]\left(\frac{1\,000\,000}{0.75}\right) = 8.5T \text{ ppm.}$$

The test results indicated that we were 8.5 times over our target rate for ac SPQL. It became imperative to identify the dominant physical defect mechanisms and implement activities to bring the rate back to our specified target.

Table 2 Raw and normalized ac SPQL fallout in parts per million (ppm).

Part	*ac fails*	*ac fails* (ppm)
CAdapt	13	$2.7T$
Proc	6	$3.8T$
IUnit	23	$7.0T$
SAdapt	31	$8.25T$
SCU1	18	$10.25T$
SCU2	6	$12.6T$
Total (cumulative)	97	$6.2T$

System test summary

System test consisted of replugging modules which had failed the gross delay test into a system environment. The results need to be examined in two groups, mainly because of the extreme differences seen between SAdapt and the other five chips. SAdapt had 5/31, or 16%, of the modules fail the card- or system-equivalent test, while three other modules failed because of the design error. On the other hand, 60/66, or 91%, of modules from the other five part numbers failed either the card- or system-equivalent test at the card manufacturing site. The fallout rate in the system is expected to be less than 100%, because the delay test detects defects whether they are on maintenance paths or system paths. The high correlation between modules failing the delay test and modules failing the card or system test indicates that our ac test methodology is detecting true gross delay defects. In other words, the failing modules represent the actual ac SPQL.

The drastic difference in fallout rates between SAdapt and the other five parts was explained by the designers as due to error detection/fault isolation which was used in the SAdapt design. Apparently, SAdapt was originally three chips in older technologies which were mapped over to CMOS. The three chips had built-in "handshaking" functions that allowed one chip to check for valid data in a cycle; if no valid data were present in that cycle, the chip would simply wait for the next cycle before proceeding with business. This type of design is less sensitive to gross delay defects.

Analysis of relative twist effectiveness

This section summarizes the number of modules failing each of the four twist test path groups and various combinations thereof. This information can be summarized across the six parts with a histogram (see **Figure 3**), which is useful in determining the effectiveness

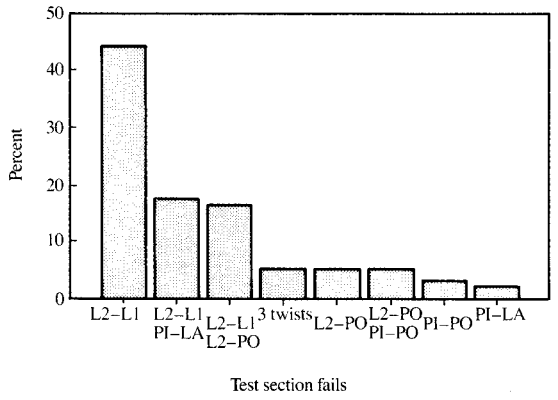

Figure 3

Twist fallout percentages for all six parts.

of the various twists. From this chart, we can see that 44.3% of the modules which failed the delay test failed only the L2–L1 path test, while 17.5% of the modules failed both L2–L1 and PI–L1 twists. All the bars added together equal 100% of the twist fail combinations seen in our ac test experiment. If we add (44.3% + 17.5% + 16.5% + 5.2%), we get a total of 83.5% of the ac failing modules in our experiment that could have been captured by administering only the L2–L1 twist. The next most effective twist was L2–PO, with 30.9% of the ac fail modules failing the L2–PO twist or in conjunction with one of the other twists. The PI–L1 test caught 23.7% of the modules. The PI–PO twist was only used for Proc, SCU1, and SAdapt. However, only nine of the 55 ac fails, or 16.4% of those three chips, failed this twist, making it the least effective.

This type of result is consistent with the test coverage numbers and the chip designs themselves. Since CSC/CGA is an LSSD-based design system, a high percentage of paths on our chips tend to be bounded by latches, making defects in those paths most likely to be caught by the L2–L1 twist. L2–PO may be the next most effective, because the higher fanout capacitance on nets driving to primary outputs may create larger delay defects, which are less likely to be canceled out by slack in the chip design and are therefore more likely to be caught by the delay test.

Summary of timing characterization results

The histogram shown in **Figure 4** summarizes the results of characterization testing done on 91 of the 97 ac fails to see how the delay test timings must be perturbed or lengthened to cause the modules to begin passing the gross delay test. This type of timing margin testing provides valuable information about whether the delay test catches transition faults or small delay defects. It is also extremely useful in conjunction with the diagnostics, as discussed in the next section. The timing margin test results show whether it is the C-clock pulse width (CW), B-clock pulse width (BW), B-to-C-clock delay (B–C Dly), or the Strobe delay which must be lengthened to allow

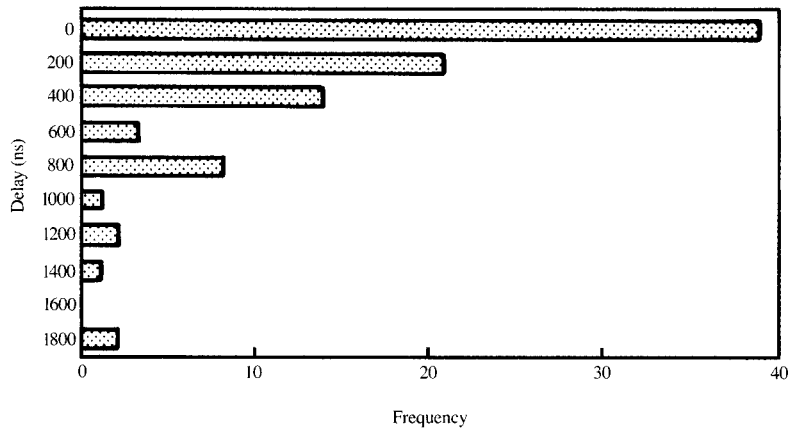

Frequency	Cumulative frequency	Percent	Cumulative percent
39	39	42.86	42.86
21	60	23.08	65.93
14	74	15.38	81.32
3	77	3.30	84.62
8	85	8.79	93.41
1	86	1.10	94.51
2	88	2.20	96.70
1	89	1.10	97.80
0	89	0.00	97.80
2	91	2.20	100.00

Figure 4

Timing distribution for all six parts.

the module to pass the delay test. The results of the timing margin test can be compared with the diagnostics to ensure that they are consistent. If the timing margin test shows that the delay defect measures hundreds of nanoseconds (ns), it can be assumed in the diagnostics that all the books fanning out from the defect could be affected by the additional delay. If the defect is only a few nanoseconds long, this may not be the case.

The timing characterization results produced values ranging in magnitude from 1–1900 ns. These delay values can be compared to the system cycle time of 120 ns. From the histogram and the column labeled *Cumulative percent*, it can be seen that 42.86% of the gross delay defects measure between 1 and 100 ns, while 65.93% of the delays measure between 1 and 200 ns. The number identifies the midpoint of the bar. The histogram could represent half of a normal distribution. There may be an additional group of small delay defects which are hidden by the chip slack and which could provide the other half of the normal distribution. It is recognized that the effectiveness of the delay test is reduced as the magnitude of the delay defect is reduced. For this reason, there are defects missing from the distribution which would be there if the delay test were 100% effective.

Diagnostic methods used to identify physical fails

The system used to analyze modules failing the delay test was the logic product diagnostic system (LPDS), which was originally developed for stuck-fault diagnosis and has been modified to support ac fault diagnosis.

The first failing pattern was simulated to determine what faults would explain the failure. This fault simulation used transition faults rather than stuck-faults, which meant that only "switching" logic could fail, and then only for one pattern in time. Additional failing patterns were then simulated to reduce the list of faults. In some cases this procedure would result in a small list of logic faults which explained all subsequent fails, and the diagnosis was considered complete.

Often there was no logic fault which would explain all of the failing tests and patterns, and the diagnostic procedure became more interactive. In this case, the LPDS TRACE routine was used to construct a logic diagram of the cone of logic feeding the failing net. TRACE is a graphical display of the chip logic design allowing the user to simulate the failing test pattern and display the logic states on each block and net. This allows the identification of critical logic within the cone. This logic diagram was reviewed and expanded for additional patterns where subsequent fails occurred. The cone trace-back was often truncated when a net with high fanout was encountered, unless failures were detected on various branches of the fanout.

Table 3 Failure analysis summary.

Defect type	*Product part name*						
	SAdapt	*IUnit*	*SCU1*	*CAdapt*	*SCU2*	*Proc*	*Total*
Total to FA	6	5	5	3	4	3	26
Resistive M1 open	4	0	4	1	1	1	11
Resistive M2 open	0	0	1	0	0	0	1
Diffusion defect	0	2	0	0	1	1	4
None found	2	3	0	2	2	1	10

The results of the timing characterization done on each module to determine what timing was critical to the failing module were used to refine the list of faults. Clock net problems would result in only the clock pulse width being critical. The amount of excessive delay also helped in truncating the logic cone at points of fanout, because large defects would be expected to affect all fanout points.

The logic trace, fault simulation, and results of timing margin tests usually resulted in a small area of the chip being labeled as suspect. Occasionally defects were detected at multiple primary outputs or shift register latches. This usually aided the diagnosis and in some cases was of primary importance. Global net defects appearing on nets with high fanout were often isolated to the physical section of the net, which was common to the failing logic sections. The physical net distribution was determined using LPDS NET DISPLAY. These defects were not precisely identified by the fault simulator, which considers block faults but not net faults.

We also found ac defects which caused increases in both rising and falling delays. These defects typically caused a large increase of resistance on a global net, and also could not be precisely diagnosed by the fault simulator, which only assumes unidirectional defects. In such cases, the fault simulator would diagnose two mutually exclusive sets of faults, with the defective block appearing in both but failing at opposite states.

Summary of physical failure analysis

A sample of ac failing modules was chosen for physical failure analysis (FA). Generally, modules were chosen that had excellent diagnostics based on first fail analysis and simulation of the identified faults over subsequent failing patterns with continued positive diagnosis. All of the part numbers were represented as equally as possible in the sample chosen for physical failure analysis. A summary of the dominant defect mechanisms can be seen in **Table 3**.

The predominant cause of delay defects is resistive first-layer-metal (M1) "opens." The most common reason for the M1 to be totally open, yet still able to conduct, is that in the CMOS process there is a layer of titanium

Table 4 CAdapt module 1 characterization information.

Failing sections	*Defect delay* (ns)	*Critical edge*	*Symptoms*	*Box fails*
L2–L1	130	B–C dly, CW, BW	100+ patts, 1 SRL stuck-at-0	C
PI–L1	120	CW, C dly		

Table 5 SCU1 module 1 characterization information.

Failing sections	*Defect delay* (ns)	*Critical edge*	*Symptoms*	*Box fails*
L2–PO	60	Strobe delay	50+ patts, 2 OCDs (0, 1, H)	S
PI–PO	100	Strobe delay		

Table 6 IUnit module 1 characterization information.

Failing sections	*Defect delay* (ns)	*Critical edge*	*Symptoms*	*Box fails*
L2–L1	500	B–C dly, CW, BW	100+ patts, 1 SRL (0, 1)	C
PI–L1	440	CW C dly		

deposited below the M1. This layer prevents aluminum from forming an alloy with silicon and creating "spikes" below the surface of the silicon into the junction. However, this titanium "skirt" has a sheet resistance in the range of 20 kΩ/sq. Instead of a dc-detectable M1 open, we have a gross delay defect. The mechanism for creating the M1 open is as follows. The titanium layer is evaporated onto an exposed resist pattern. Foreign material in the evaporator may be deposited onto the titanium layer prior to the evaporation of M1. Then, when the lift-off process is performed next, not only are the unwanted M1 and titanium removed with the resist, but also the foreign material and the M1 on top of it, creating a resistive M1 open.

Four modules were diagnosed as having manufacturing defects causing shallow diffusions for several adjacent transistors. The shallow field-effect transistor (FET) channels show early threshold voltage breakdown and cause the FETs to be slow. The shallow diffusions were determined to be caused by photoresist flaking off the outer areas of the wafers where there are few circuit patterns and being redeposited on the wafer to partially block the channel diffusions. The absence of circuit patterns causes poor adhesion and promotes flaking.

One module was found to have a delay defect due to a resistive M2 line caused by M2 notching. M2 notching is created when the light exposing the M2 resist-patterning layer is reflected on the underlying gate array topography. This causes ridges in the resist which result in a notched M2 line.

The physical failure analysis results for five representative ac fails are presented next in detail.

- *Analysis of CAdapt module 1*

The characterization results for one of the modules, of the product type CAdapt, can be found in **Table 4**. The table summarizes which delay test path groups the module failed, the timing characterization results, the critical clock and/or strobe edges which sensitized the delay fault, the failing test pattern symptoms, and the card-equivalent (C) or system-equivalent (S) test results.

Diagnostic summary A total of 128 patterns failed, with the same latch always failing with a 0 value. The fault was clearly in a latch-to-latch path, since it was sensitive to both the clock pulse widths and separations. The transition-fault simulator identified the data input to the SRL or the two-input AND gate feeding the SRL. The timing margin results are also consistent with a potential defect on the data input to the latch.

Failure analysis summary The failure analysis pinpointed a large M1 defect in the two-way AND book identified by diagnostics. The scanning electron microscope (SEM) photographs for CAdapt 1 are shown in **Figure 5**. The first photograph shows the chip with its metallization intact; the defect is centered in the lower half of the photo. In the second photograph, M2 has been removed and the defect is shown at a higher magnification. The defect has removed most of the V_{dd} contact to two p-channel FETs. This creates series resistance between V_{dd} and the sources of these FETs, making them slow to rise. This explains why the SRL failed only with a 0 value.

- *Analysis of SCU1 module 1*

Table 5 shows the characterization results for an SCU1 module.

Diagnostic summary One primary output was the dominant failure for SCU1 module 1. This TTL three-state off-chip driver (OCD) was failing at all three states: 0, 1, and high impedance (H). A failure at the high-impedance state signifies that the output of the OCD was at a voltage level somewhere between a valid "0" and "1" voltage level. No single transition fault explained all the fails during simulation; however, the OCD and the block feeding the data to the OCD were the only logical places to find defects. This was consistent with the timing characterization, which showed that the module would

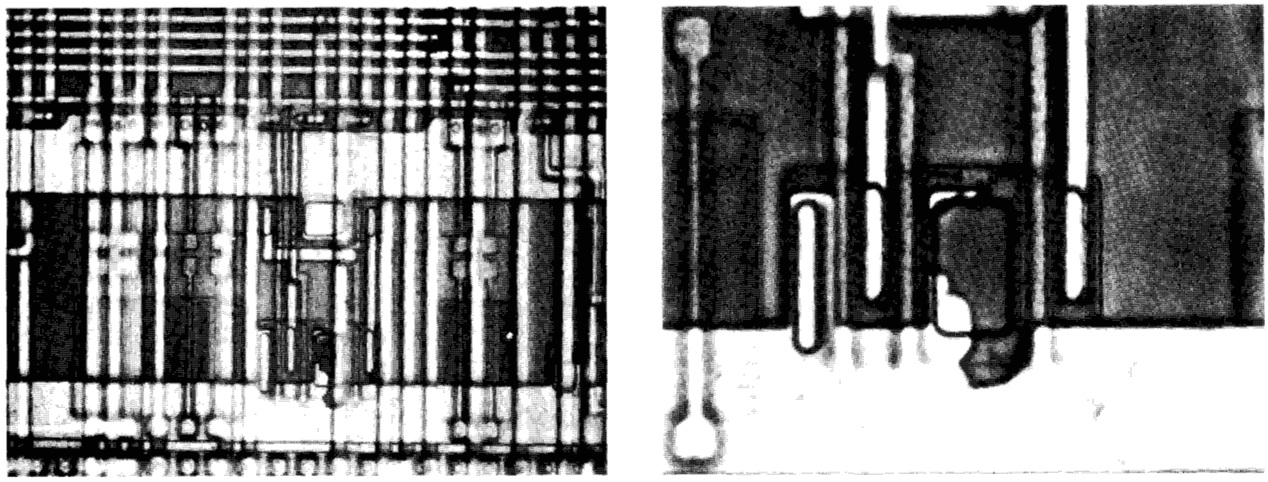

Figure 5

SEM photographs for CAdapt.

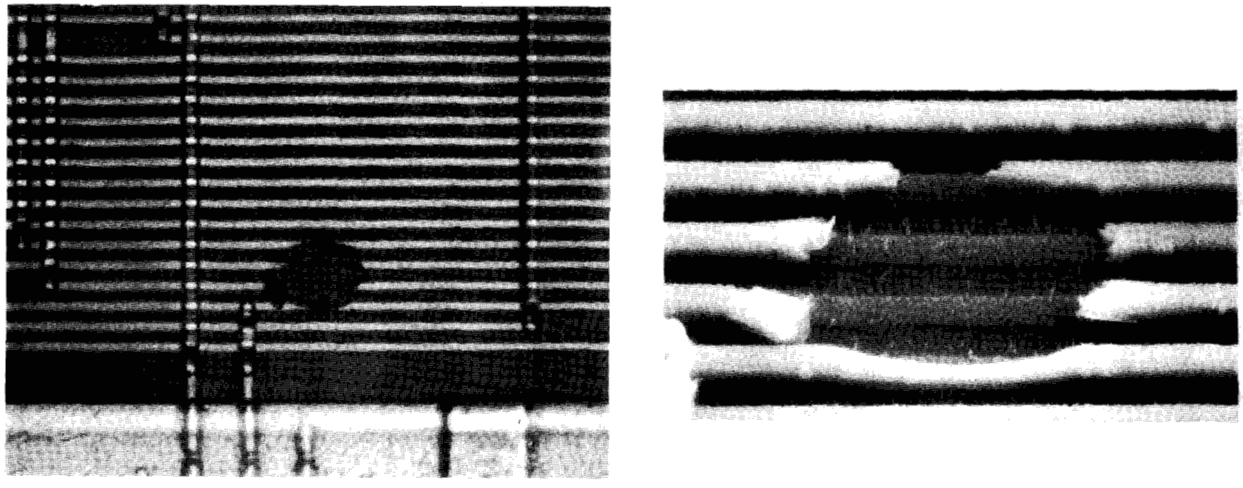

Figure 6

SEM photographs for SCU1 module 1.

pass if we lengthened the strobe time for the output measure by 100 ns.

Failure analysis summary The large M1 open shown in **Figure 6** was located during failure analysis on the global net feeding the data input of the OCD. All three M1 lines, which are incomplete due to the defect, were measured to be resistive prior to removal of the nitride. The top line measured 100 kΩ, and the middle line measured 400 kΩ. The lower line measured 333 kΩ, and is the one that feeds the failing OCD. Apparently, the other two lines are nets with a low enough capacitance so that the resulting RC time constant adder is not enough for them to fail the delay test.

- *Analysis of IUnit module 1*

The characterization results for an IUnit module are shown in **Table 6**.

Diagnostic summary Only one SRL was failing in both the 0 state and the 1 state. The diagnostic simulation results identified the data input of the SRL, the two-way

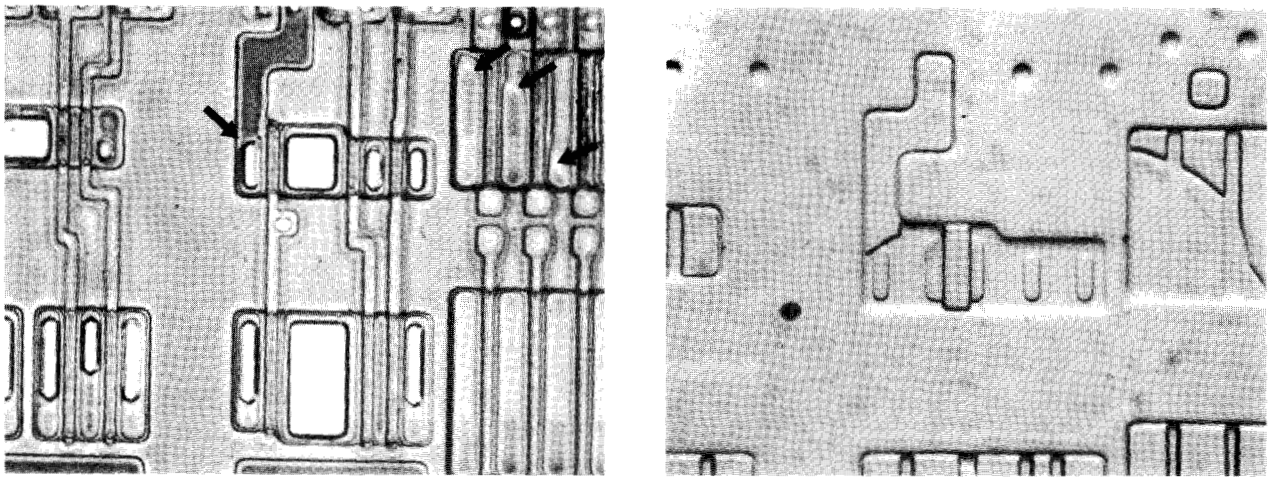

Figure 7

SEM photographs for IUnit module 1.

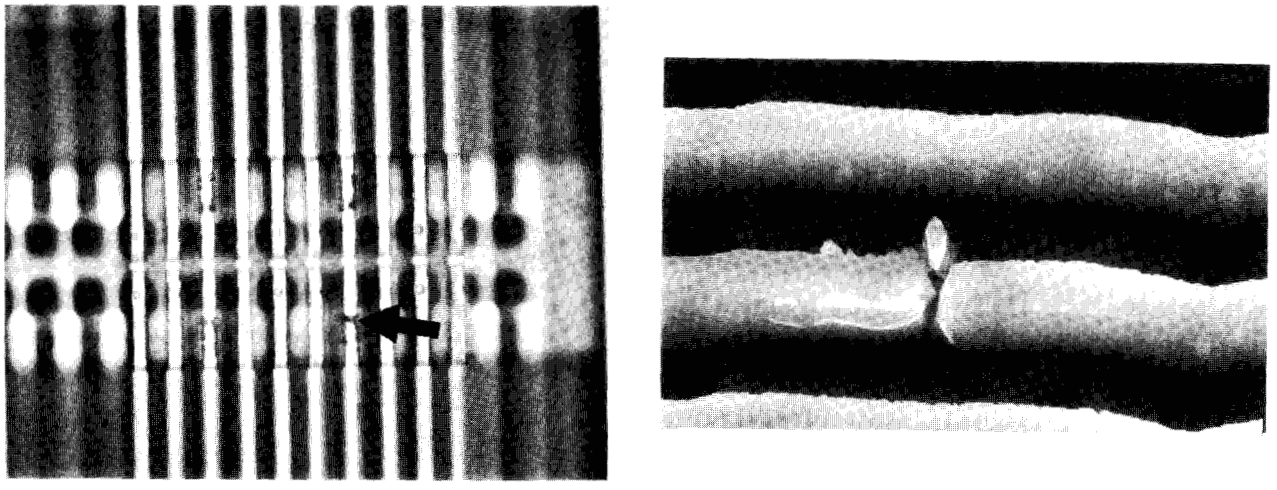

Figure 8

SEM photographs for SCUI module 2.

OR logic block feeding the data input, and the 2 × 3 AND*OR feeding the OR.

Failure analysis summary Physical failure analysis identified a diffusion defect in a two-way OR. In this defect, foreign material (photoresist) caused three shallow n-channel FET diffusions. These NFETs exhibited early threshold voltage breakdown when electrically characterized. The SEM photographs for IUnit module 1 are shown in **Figure 7**. The photograph on the left shows anomalies in the silicon, indicated with arrows. In the photograph on the right, the chip has been stripped to the surface, and the diffusions have been dyed. The darker, well-defined lines indicate a normal depth diffusion, and the shallow diffusions in the center and at the right can be differentiated.

- *Analysis of SCU1 module 2*

The characterization results for a second SCU1 module can be found in **Table 7**.

Diagnostic summary The simulation identified an input, output, and internal nodes of a six-way NAND logic block.

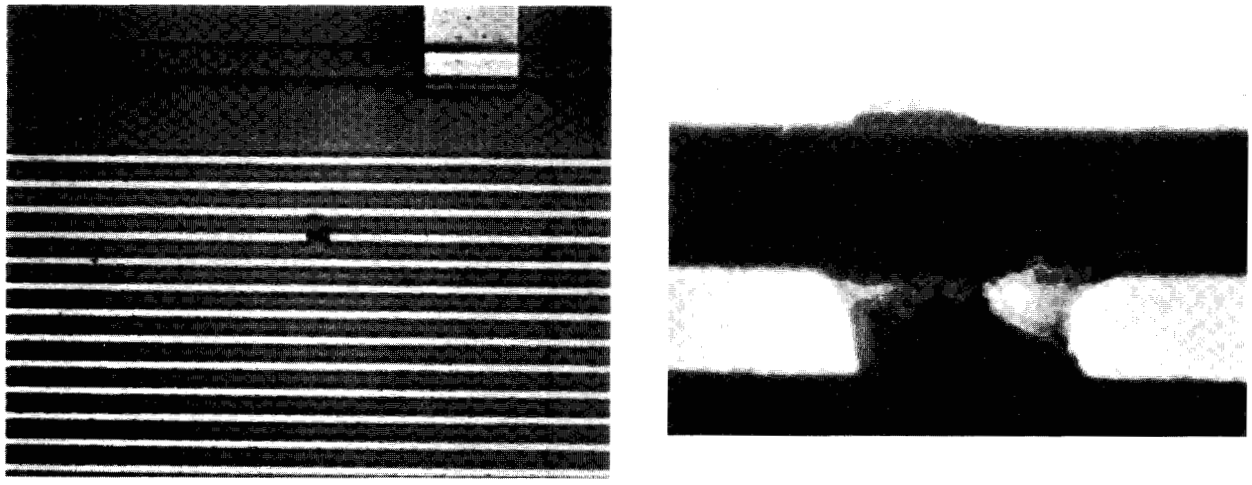

Figure 9

SEM photographs for SAdapt module 2.

Failure analysis summary A resistive M2 open was located on the global output of the six-way NAND. This was attributed to light reflecting off the gate array topography and causing uneven hardening of photoresist. The resultant M2 layer exhibits ridging and notching. The SEM photographs are shown in **Figure 8**. Four very regular notches can be seen in the left-hand SEM photograph, and the resistive M2 open is indicated with an arrow. The right-hand photograph is a greater magnification of the defect.

• *Analysis of SAdapt module 2*

Table 8 shows the characterization results for an SAdapt module.

Diagnostic summary Two three-state OCDs dominated the failing patterns for SAdapt module 2. The logic driving the enable/inhibit input of the OCDs was shared by the failing drivers plus 14 other OCDs, which were passing all the tests. When the enable net was displayed using NET DISPLAY RC analysis for SAdapt module 2 (described next), it was seen that the two failing OCDs were alone at the end of a long branch. This net was identified for failure analysis.

Failure analysis summary Once again, the failure was verified to be due to an M1 open defect. The resistance was measured to be about 80 kΩ and was found to be due to titanium being intact below the FM prior to M1 deposition. The defect was located on the enable net beyond the place where it branched to the last OCD which passed all tests. The SEM photographs are shown in **Figure 9**.

Table 7 SCU1 module 2 characterization information.

Failing sections	*Defect delay* (ns)	*Critical edge*	*Symptoms*	*Box fails*
L2–PO	50	Strobe delay	1 patt, 5 OCDs stuck-at-0	S

Table 8 SAdapt module 2 characterization information.

Failing sections	*Defect delay* (ns)	*Critical edge*	*Symptoms*	*Box fails*
L2–L1	56	B–C dly, CW, BW	1 patt/1 bit, many patts 2 PO, OCDs (H)	S
L2–PO PI–PO	72	Strobe delay		

RC analysis of SAdapt module 2

SAdapt module 2 was chosen as a good candidate for ASTAP circuit simulation of the defective net to verify the results of the timing characterization testing. **Figure 10** shows the NET DISPLAY plot of the defective net. The driving buffer book is located in the lower center of the figure. The net starts out on M1 and fans out to three OCDs before it branches in one direction to M2 and traverses the length of the chip. Its other branch continues to the right, where it fans out to ten additional OCDs, all of which never failed. The M2 branch reverts to M1, where it fans out immediately to one passing OCD and then reaches the two failing OCDs in the upper

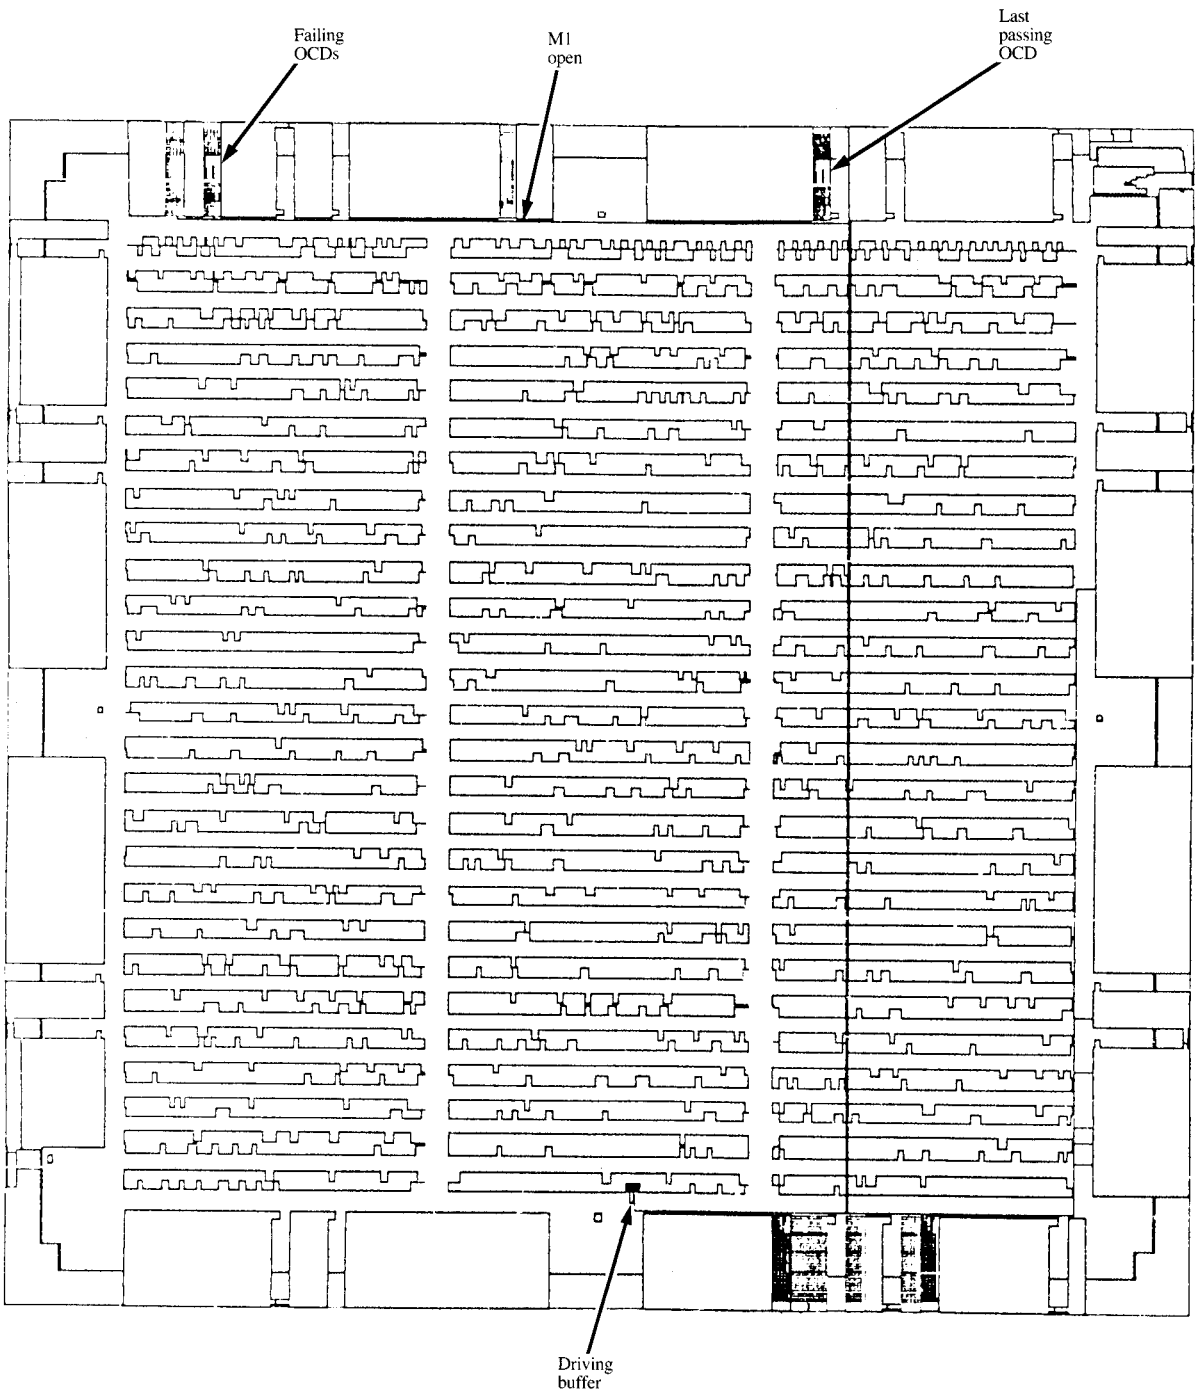

Figure 10

NET DISPLAY for SAdapt module 2.

left-hand corner of the chip. We can also see where the missing M1 occurred relative to the passing and failing OCDs.

Using the postphysical design wiring and fanout capacitance values from the chip-level delay calculator (PHIDO), an RC network was constructed in ASTAP. Physical failure analysis results showed that the defect added about 80 kΩ of resistance to the global net. From the ASTAP analysis, it was determined that the defect added about 75 ns to the net delay, which agrees very well with the timing margin results measuring 72 ns.

This same ASTAP deck was used to determine whether the defect would have been dc-detectable if the position of the defect were changed on the net. The timings used for the ASTAP run corresponded to the specified design system timing for a 7.5-mm chip. The actual defect position is very close to the sink end of the net. In the ASTAP circuit simulation deck, the defect position was moved to the source of the net, or directly after the driving buffer book. This provides the maximum possible RC delay adder to this net. The net was simulated and still switched in 388 ns. However, this was still fast enough to pass the dc test! Furthermore, this net, with a fanout of 16 OCDs, has a wiring + fanout capacitance of 10.3 pF. This is contrasted with a typical mean wiring + fanout capacitance of 0.5 pF. This net is 3-sigma worst case for capacitance, yet the presence of a worst-case-placed resistive M1 open will still allow the chip to pass the dc stuck-fault test. It can be concluded that a resistive M1 open on virtually any net on a chip will still result in a dc-good chip.

Conclusions

This experiment has yielded a large amount of data regarding the ac SPQL for CSC/CGA, but has also generated questions to be addressed by future analysis. The data showed that the level of ac SPQL for CSC/CGA was over target for the period of time in which the experiment sample was manufactured. Failure analysis has identified resistive M1 opens as the dominant failure mechanism causing delay defects. As a result of the data presented in this paper, activities have been implemented to reduce the delay defect component of shipped product quality level to the specified target level. A summary of the activities implemented follows.

All of the ac failure mechanisms identified by this experiment have been addressed by manufacturing. Additional dc-good modules which have been manufactured since the process improvements have been ac-tested. The results of these tests provide a monitor of the defect learning and its effect on ac SPQL. The total SPQL is also monitored on the basis of the manufacturing card test removal rate. In addition, the ac SPQL evaluation experiment has been expanded to other locations which manufacture CSC/CGA logic product. As a result of these activities, the ac SPQL has been improved by a factor of 10.

The ac test methodology has been effective at detecting gross delay defects. Further analysis is being performed to better understand the characteristics of delay defects, particularly small ones. In conclusion, this ac test experiment has improved the quality of CSC/CGA product and has also demonstrated the benefit of delay testing.

Acknowledgments

The authors wish to thank the following individuals for their contributions: P. Czahor and R. Walker, for performing the physical failure analysis; T. Cayea and D. Sprague, for test characterization contributions; T. Mechler, for work on logic diagnostics; M. Andraschko and G. Dancker, for providing delay test timings; E. Lakkis, for card and system testing; and W. Conyers, for tracking product fallout at the card and system manufacturing plant.

References

1. A. W. Aldridge, R. F. Keil, J. H. Panner, G. D. Pittman, and D. R. Thomas, "A 40K Equivalent Gate CMOS Standard Cell Chip," *Proceedings of the IEEE 1987 Custom Integrated Circuits Conference*, pp. 248–252.
2. Robert Hornung, Martine Bonneau, Bernard Waymel, James Fiore, Elliot Gould, Ronald Piro, John Martin, Lance McAllister, and Sze Tom, "A Versatile VLSI Design System for Combining Gate Array and Standard Cell Circuits on the Same Chip," *Proceedings of the IEEE 1987 Custom Integrated Circuits Conference*, pp. 245–247.
3. E. B. Eichelberger and T. W. Williams, "A Logic Design Structure for LSI Testability," *Proceedings of the ACM/IEEE 14th Design Automation Conference*, June 1977, pp. 462–468.
4. Jerry M. Soden and Charles F. Hawkins, "Electrical Properties and Detection Methods for CMOS IC Defects," *Proceedings of the IEEE 1st European Test Conference*, April 1989, pp. 76–85.
5. John Waicukauski, Eric Lindbloom, Barry Rosen, and Vijay Iyengar, "Transition Fault Simulation," *IEEE Design & Test of Computers* **4**, 32–38 (April 1987).
6. K. Kishida, F. Shirotori, Y. Ikemoto, S. Ishiyama, and T. Hayashi, "A Delay Test System for High Speed Logic LSIs," *Proceedings of the IEEE 23rd Design Automation Conference*, June 1986, pp. 786–790.

Received May 20, 1989; accepted for publication September 29, 1989

Orest Bula *IBM General Technology Division, Essex Junction, Vermont 05452.* Mr. Bula is currently a Staff Engineer/Scientist in the Technology Simulation Department in the Technology Development area. He received his B.S. in physics from the City

College of New York in 1969, and his M.A. in physics from the City University of New York in 1971. Mr. Bula joined IBM at Burlington in 1978. His work has focused primarily on transformation bridges for conversion of test data to formats acceptable to multiple logic testers and burn-in tools. Mr. Bula has published six IBM technical disclosures and has two patents pending. He is a member of the Institute of Electrical and Electronics Engineers.

John Moser *IBM General Technology Division, Essex Junction, Vermont 05452.* Mr. Moser is an Advisory Engineer in the High End ASIC Product Development area. He received his B.S. in electrical engineering in 1964 from Washington University, St. Louis, Missouri, and his M.S. in electrical engineering in 1977 from the University of Vermont. He joined IBM in 1964 and has held technical and management positions in Test and Product Development. Mr. Moser is an inventor, with three patents issued for test equipment design and product design techniques to improve testability. Recent assignments have allowed him to focus on the quality of shipped product. Mr. Moser received an IBM Outstanding Technical Achievement Award for development of improved methods of measuring quality and understanding mechanism-limited quality, especially as related to delay defects.

Jeanne Trinko *IBM General Technology Division, Essex Junction, Vermont 05452.* Ms. Trinko is a Senior Associate Engineer responsible for Test, Reliability, and SPQL in the Advanced BiCMOS Product Development area. She joined IBM in 1985 after receiving a B.S. in electrical engineering from the University of Vermont. Through the IBM Graduate Work/Study program, she received an M.S. in electrical engineering from the University of Vermont in 1989. Her past work involved the design and characterization of CMOS ASIC logic products, including on-chip and card-level timing and chip/package noise analysis. Ms. Trinko is the President of the Vermont Chapter of the Society of Women Engineers.

Mark Weissman *IBM General Technology Division, Essex Junction, Vermont 05452.* Mr. Weissman is a Senior Associate Programmer in the Logic Diagnostic Development Department of the Logic Test and Software Systems function. He joined IBM in 1984, after graduation from the University of Pennsylvania with a B.S. in computer science. Mr. Weissman has participated in the development and support of diagnostic models, tools, and simulators.

Frank Woytowich *IBM General Technology Division, Essex Junction, Vermont 05452.* Mr. Woytowich is a Senior Associate Engineer in Logic Design Systems Test Engineering, responsible for chip and module test methodology and its applications using automatic test equipment. He received his B.S. in electrical engineering in 1982, joining IBM that same year. Through the IBM Graduate Work/Study program, Mr. Woytowich received an M.S. in electrical engineering from the University of Vermont in 1987. His work has encompassed the test of custom microprocessors and microcontrollers as well as ASIC design systems logic/array components. His contributions have recently focused on ac test methodology and improvement of the ac quality of shipped components.