Empirical Validation of Yield Recovery Using Idle-Cycle Insertion

Donghui Lee

Stanford University

Intaik Park

Stanford University

Erik Volkerink

Verigy

Jeff Rearick

Advanced Micro Devices

Altering the details of a delay test application protocol can significantly affect test results and thus yield. Specifically, inserting wait states between scan shifting and the launch clock can improve yield considerably. Although the exact mechanisms involved remain elusive, the authors eliminate several possibilities through a series of experiments.

The act of shifting the scan chains causes many scan flip-flops (and much of the combinational logic fed by them) to be toggled. This is due to the low correlation between consecutive vectors and the presence of large amounts of pseudorandomly filled don't-care

■ **A TRANSITION FAULT** model typically guides test pattern generation to cover defects that cause late signal transitions.¹ Such a delay test requires a two-pattern test: The first pattern initializes a node signal value. The second pattern causes this value to switch to the opposite value and sensitize the transition to an observable output. These two events should happen on two successive clock pulses at the rated speed of the devices. These two clock pulses are responsible for launching and capturing the transition, respectively.

Two scan-based delay test protocols are supported by most commercial ATPG tools: launch on shift (LOS) and launch on capture (LOC). In LOS test patterns, scan data is loaded through scan chains. Nodes in the logic under test are initialized at the last shift clock (the launch clock). For the flip-flops to capture the logic's response in the downstream flip-flops, the scan-enable signal must switch from 1 to 0 before the subsequent capture clock is applied. In LOC test patterns, the scan-enable signal does not need to operate at system speed. Once scan data is loaded through scan chains, the scan-enable signal transitions from 1 to 0 in whatever time frame this signal requires. Subsequently, the launch and capture clocks are applied at full system speed, after the scan-enable signal switches to 1, and then the logic's response is scanned out.

bits within a vector, resulting in several nodes that will be toggled when the vectors are scanned in and scanned out.² Therefore, scan-based test patterns' switching activity often leads to high power consumption and large currents during test mode. This current demand can cause supply voltage droop (also called *IR* drop) because of the supply network's resistance. As a result, the logic gates will operate below their nominal voltage and consequently below their rated speed. Thus, *IR* drop caused during scan test can cause a circuit to function more slowly than under normal operating conditions, leading to the mistaken conclusion that the circuit is defective. This mistake results in a yield loss known as a *false failure*.

In this article, we evaluate the impact of idle-cycle insertion on yield by quantifying the number of test escapes. We empirically quantify false failures resulting from *IR* drop by inserting idle cycles at appropriate points during the scan test application protocol. Launch delay (LD) test delays the launch clock by inserting a certain amount of time (called idle cycles or idle time) after the last shift cycle. During these idle cycles, the chip's power supply network is given time to recover from the *IR* drop induced by excessive switching activity during scan shifting. The amount of *IR* drop depends on the power supply network design;

different chip designs need different idle times. We present LD test results for two different chip designs: the ELF13 graphics processor and the ELF18 DSP processor.

Experimental results show that altering the details of the delay test application protocol can significantly impact test results and thus the yield of the product under test. Inserting wait states between scan shifting and the launch clock results in a measurable yield improvement. For example, in ELF13 experiments, we achieved a 20% reduction in false failures without increasing the number of test escapes or requiring additional ATPG time. Although the exact mechanisms involved remain elusive, we eliminated several possibilities through a series of experiments. It is clear from these experiments that yield recovery is a real phenomenon and that LD tests help to recover from *IR* drop. (The "Previous work" sidebar discusses other approaches that have attempted to minimize the effects of *IR* drop during scan test.)

Adding idle cycles

An *LOC LD test set* is an LOC test set with idle cycles inserted between the last shift and the launch clock in each test pattern. The *regular LOC test set* is the LOC test set before idle cycles are inserted, which is directly generated from a commercial ATPG tool. Figure 1 presents two timing-diagram variants for LOC LD.

As Figure 1 shows, there are two possible ways to implement an LOC LD transition delay test set: insert idle cycles before the scan-enable signal transition or insert idle cycles after the scan-enable signal transition. The key properties of the LD test set are as follows:

- The expected values of test patterns do not change. Therefore, no additional ATPG time is required.
- No clocks pulse during idle cycles.
- No outputs are compared during idle cycles.
- Primary input values are not changed during idle cycles.
- The LD test set has the same fault coverage as the regular test set.
- The LD test merely increases the wait time between the end of scan-shifting activity and the start of the launch clock pulse.
- The LD test increases overall test time.

Implementation of LD test sets takes advantage of commercial ATPG tools. After a commercial ATPG tool

Previous work

Several other approaches have attempted to minimize the effects of *IR* drop during scan test. One such method is to generate *quiet patterns*, which contain low switching activity.¹ In this method, the logic under test has the same switching activity as the original patterns, but the rest of the circuit has lower switching activity. Alternatively, test scheduling provides a way to generate patterns that satisfy a chip's power budget.² Other previous work on generic test power reduction includes a static compaction technique,³ test vector reordering,⁴ and gating of scan elements.⁵ Although all of this work demonstrates power reduction during scan-based testing, none of it attempts to correlate excess power consumption with false failures.

References

1. J. Saxena et al., "A Case Study of IR-Drop in Structured At-Speed Testing," *Proc. Int'l Test Conf. (ITC 03)*, IEEE CS Press, 2003, pp. 1098-1104.
2. R.M. Chou, K. Saluja, and V.D. Agrawal, "Scheduling Tests for VLSI Systems under Power Constraints," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 5, no. 2, June 1997, pp. 175-185.
3. R. Sankaralingam, R.R. Oruganti, and N. Touba, "Static Compaction Techniques to Control Scan Vector Power Dissipation," *Proc. 18th IEEE VLSI Test Symp. (VTS 00)*, IEEE CS Press, 2000, pp. 35-40.
4. V. Dabholkar et al., "Techniques for Minimizing Power Dissipation in Scan and Combinational Circuits during Test Application," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 17, no. 12, Dec. 1998, pp. 1325-1333.
5. S. Gerstendorfer and H.J. Wunderlich, "Minimized Power Consumption for Scan-Based BIST," *Proc. Int'l Test Conf. (ITC 99)*, IEEE CS Press, 1999, pp. 77-84.

generates a compact test set, a script can be used to easily insert idle cycles. Thus, implementing LD test sets incurs no ATPG overhead. However, test application time increases because idle cycles increase the number of tester cycles per test pattern. If the purpose of the LD test is to reduce false failures, it's not necessary to add idle cycles to passing test patterns, because idle-cycle insertion would not change the pass/fail results of passing patterns.

We provided a control group for the experiment by modifying another test set based on the SSF to include idle cycles. An *SSF LD test set* is a regular SSF test set with idle cycles inserted between the

Figure 1. Timing diagram of a launch-on-capture (LOC) launch delay (LD) test pattern: idle cycles before (a) and after (b) the scan-enable transition.

last shift and the single capture clock in each test pattern.

ELF13 experimental setup

ELF13 is a processor manufactured with 0.13-micron copper technology. It has about 7.2 million logic gates, with more than 10 different clock domains. We generated a variety of LD test sets for both single stuck-at faults and transition delay faults. We generated 20 test sets for ELF13, using two fault models; and we added five different numbers of idle cycles (100, 200, 300, 500, and 1,000), each 142.84 ns in duration. We applied these 20 LD test sets to chips that failed the regular LOC or 2-detect SSF test sets. (Lee and McCluskey describe the process of selecting ELF13 sample chips.³⁾)

ELF13 experimental results

Table 1 presents the experimental results for the LD test sets on the ELF13 chip. We used system-level test (SLT) and functional test to discriminate between false failures and defective chips for ELF13. That is, false failures passed both SLT and functional test; defective chips failed either SLT, functional test, or both. Although the fault grading of functional test was unavailable, Lee and McCluskey present the experimental results of functional-test coverage and failing behaviors by examining failing flip-flops and the

design blocks containing potential defects.⁴

Four important results are apparent from Table 1. First, the greater the number of idle cycles inserted, the greater the number of false failures that pass the LD test sets. This is direct evidence that waiting between scan shifting and capturing directly benefits yield. Second, no chips passed before the scan-enable transition. This points to the scan-enable transition event, rather than the scan-shifting activity, as the cause of the electrical disturbance that reduces yield. Third, the insertion

of idle cycles is seemingly less effective for SSF tests than for LOC transition tests, though the sample size is rather small. Finally, no test escapes were observed from the LD test sets.

Avoidance of *IR* drop by repeat fill

An alternative method for recovering yield loss due to *IR* drop is to avoid *IR* drop in the first place (instead of simply allowing recovery from it by adding idle cycles). It's possible to reduce the amount of switching power (and hence current and *IR* drop) by minimizing the switching activity during scan shifting. We can reduce the switching activity of test patterns by filling the don't-care bits with the last care bit (a *repeat-fill scheme*).⁵ We used the ELF13 chips tested here to collect empirical data on the ability of repeat-fill test patterns to minimize *IR*-drop-induced yield loss.

We generated two LOC test sets using different don't-care bit assignments to measure the effect of power reduction on decreasing the number of false failures. For the first LOC test set, an ATPG pseudo-randomly assigned don't-care bits (a *random-fill scheme*). For the second LOC test set, we used a repeat-fill scheme. For the ELF13 design, the size of the repeat-fill test set is prohibitively large if the entire fault list is targeted, so we generated both test sets by targeting only the faults detected by failing test patterns in the experimental chip samples. We identified these

Table 1. Experimental results of launch delay (LD) tests for the ELF13 chip.

Idle-cycle position	No. of idle cycles	Launch on capture (LOC)				Single stuck-at fault (SSF)			
		False failures		Defective chips		False failures		Defective chips	
		Pass	Fail	Pass	Fail	Pass	Fail	Pass	Fail
Before scan-enable transition	100	0	88	0	20	0	14	0	4
	200	0	88	0	20	0	14	0	4
	300	0	88	0	20	0	14	0	4
	500	0	88	0	20	0	14	0	4
	1,000	0	88	0	20	0	14	0	4
After scan-enable transition	100	3	85	0	20	0	14	0	4
	200	3	85	0	20	0	14	0	4
	300	5	83	0	20	0	14	0	4
	500	14	74	0	20	0	14	0	4
	1,000	17	71	0	20	1	13	0	4

faults by fault-grading the failing test patterns. The resulting test sets would be expected to fail when applied to these sample chips, unless there were indeed *IR*-drop issues related to the fill data.

Table 2 indicates the yield recovery made possible by using repeat fill instead of random fill. Specifically, 58% of false failures (46 of 78) and 50% of defective chips (6 of 12) passed the repeat-fill LOC test set but failed the random-fill LOC test set. Hence, these results confirm that using the repeat-fill scheme can bring about significant yield recovery by suppressing *IR* drop. Unfortunately, this scheme could also degrade the test quality because of the loss of nontarget defect detection by random patterns. (The “Impact on test time” sidebar shows the impact of test execution time on scan-based delay test.)

ELF18 experimental results

The results observed on the ELF13 chip prompted further study using the ELF18 chip.⁶ Each of the sample ELF18 chips contained six DSP cores, which we tested individually. The ELF18 chips were fabricated using 0.18-micron aluminum and copper technology. Each chip has about 53,000 logic gates. Unfortunately, functional test was unavailable for the ELF18; there-

fore, false failures were not identified. The experimental setup paralleled that used for the ELF13 chip, although we varied the number of idle cycles more (nine different numbers of idle cycles instead of five) and these idle cycles were smaller in duration (50 ns each). Table 3 presents the experimental results for the 36 LD test sets on the ELF18 chip.

Again, some important results are apparent from Table 3, though often different from those in Table 1. First, the greater the number of idle cycles inserted, the greater the number of DSP cores that pass the LD test sets, but only for the stuck-at-fault tests (unlike ELF13). That trend was not present in the ELF18 LOC LD tests, in which only two ELF18 cores benefited from the addition of a single idle cycle. Those same two cores did not pass the LOC LD test set with further insertion of additional idle cycles. Second, once again, no chips passed the LOC LD tests before the scan-enable transition, but now many pass the SSF-LD tests before the scan-enable transition. This points to the scan-shifting activity rather than the scan-enable transition event as the cause of the electrical disturbance that reduces yield—just the opposite of the ELF13 case. Third, and again in opposition to the ELF13 results, the insertion of idle cycles is now more effective for SSF tests than for LOC transition tests.

Possible explanations for yield recovery

Despite the somewhat inconsistent results between the two experiments, we can make the following generalization: Adding a sufficiently large number of idle cycles both before and after the scan-enable

Table 2. ELF13 test results for repeat-fill LOC on the chips that failed random-fill LOC.

ELF13 chips	Pass	Fail	Total
False failures	46	32	78
Defective chips	6	6	12

Impact on test time

Table A shows the impact of the various ELF13 LD test sets on test execution time.

Table A. Summary of test application time for the ELF13 chip.

Test set	No. of idle cycles	Test time (s)	Additional test time (s)
Reference LOC	0	11.3	0
LOC LD	100	11.6	0.3
	200	12.0	0.7
	300	12.4	1.1
	500	13.2	1.9
	1,000	15.0	3.7
Reference SSF	0	5.1	0
SSF LD	100	5.3	0.2
	200	5.5	0.4
	300	5.6	0.5
	500	6.0	0.9
	1,000	6.8	1.7

*LOC: launch on capture; SSF: single stuck-at fault; LD: launch delay.

Although the additional test time seems large (a 25% increase in test application time with 1,000 idle cycles), the additional test time can be less than 25% of the time for the entire test flow. LD test patterns do not need to be applied to all chips. If false failures must be reduced, a test flow can be created so that it applies LD test patterns only to chips that fail regular patterns. In this case, the additional test time due to idle cycles depends on the yield. For

example, if there is a 25% test time increase due to idle cycles, the additional test time in terms of total test application time is $(1 - Y)(1.25)$, where Y is the yield. If the yield is 90%, the additional test time increases by 12.5%. We can modify this equation according to the test suites used in the entire test flow. For example, the 12.5% test time increase just mentioned assumes that only the LOC test is applied in the test flow. If the SSF and LOC tests are applied, and false failures due to the LOC test must be reduced, the additional test time is $Y_{SSF}(1 - Y_{LOC})(1.25)$, where Y_{SSF} and Y_{LOC} refer to the yield of the SSF and LOC tests. In this case, the test time increase is smaller than in the previous case. We can further reduce total test application time by considering the failing pattern data. Note that idle cycles do not need to be added to passing patterns. When chips are tested using regular patterns, failing regular patterns can be identified by logging failing

cycles. Then, idle cycles can be added only to those failing regular patterns. In this case, the total test application time can be further reduced, depending on the number of failing patterns. If f represents the percentage of regular LOC patterns that fail, the total application time is $Y_{SSF}(1 - Y_{LOC})(1.25)(0.01)(f)$. If both yields are 90%, and 20% of regular LOC patterns fail, then the additional test time increases by 2.25%.

transition is likely to recover yield loss for either the LOC test sets, the SSF test sets, or both. Here, we consider different possible explanations for this yield recovery.

Clock stretching

One possible explanation for the observed data that more chips pass when extra idle cycles are inserted is that these chips are in fact defective but are being undertested in the idle-cycle environment. Earlier work has demonstrated that the clock network itself can run 15% more slowly during scan-based delay test than during functional test.⁷ Clock stretching may occur because the supply network's inductance opposes the sudden change in current that the launch cycle requires following a relatively quiet period after scan-shifting activity ceases. At first glance, idle-cycle insertion would seem to exacerbate this problem by prolonging that quiet period.

To resolve this issue, we performed Spice simulations. Previous work analyzes global clock distribution networks.^{8,9} We combined this information with clock generator models (which use a five-stage inverter chain with one inverter replaced by a NAND gate to enable the clock generator), a power grid, and a lossy RLC transmission line for the clock distribution tree to set up the simulation conditions.⁸ During the simulations, we measured the clock period from the first clock pulse to the second clock pulse after the idle time shown in Figure 2.

Table 4 presents simulation conditions for the lossy transmission line model of the clock network.⁸ In all four simulations, we set the R , L , and C values of the power grid model as follows: $R = 15$ ohms, $L = 0.1$ nH, and $C = 0.2$ pF. We estimated all values from reasonable assumptions consistent with the ELF13 design, although we did not extract these values from layout information.

Figure 3 presents the simulation results. In all four simulations, clock stretching saturated after less than 90 ns of idle time. In ELF13 LOC testing, the intrinsic dead time measured from the last shift to the launch clock (with no idle cycles inserted) was greater than 300 ns. This implies that clock stretching was already saturated before idle-cycle insertion, thus refuting the claim that the yield recovery observed might have been false due to undertesting exacerbated by idle-cycle insertion. In other words, although the Spice simulations indeed confirmed that clock stretching was occurring, it occurred in the same measure between the regular LOC tests and the LOC LD tests, so any yield difference between those two tests must be attributed to some other factor.

Another compelling argument against clock stretching being the possible explanation for the observed yield improvement is that the yield also improved for SSF tests when we added idle cycles. Because the SSF test application method uses only a single capture clock, stretching from one clock period to the next is irrelevant. Some other mechanism must be at work.

Power grid recovery after *IR* drop

The large current surge generated by simultaneously switching internal nodes can cause supply voltage droop. This supply voltage droop (the *IR* drop) occurs because of the supply network's resistance, along with its inductive components [the *L(di/dt)* drop].¹⁰ After a current surge (such as from scan shifting) ends, the power supply network recovers its intended value, governed by the network's time constant. A simple lumped RLC circuit can model the supply-ground network,¹¹ as Figure 4 shows. Each logic gate experiences different *V_{DD}* and ground, depending on the physical layout.

Many previous studies have presented analytical solutions to the supply network noise problem.¹⁰ In this experiment, we performed Spice simulations using a 0.13-micron technology file to explore the dynamic behavior of this network and to determine whether its recovery time was consistent with the observed results from LD testing.

We used inverters as logic gates in the simulation, which we performed with the following circuit configurations:

Table 3. Experimental results for LD tests on the ELF18 chip.					
Idle-cycle position	No. of idle cycles	LOC		SSF	
		Pass	Fail	Pass	Fail
Before scan-enable transition	1	0	141	5	128
	2	0	141	5	128
	3	0	141	7	126
	5	0	141	8	125
	10	0	141	9	124
	15	0	141	14	119
	30	0	141	15	118
	60	0	141	16	117
	140	0	141	17	116
After scan-enable transition	1	2	139	6	127
	2	0	141	7	126
	3	0	141	9	124
	5	0	141	10	123
	10	0	141	12	121
	15	0	141	16	117
	30	0	141	16	117
	60	0	141	17	116
	140	0	141	18	115

*Passing and failing entries refer to the number of DSP cores, not chips.

1. Each inverter switched independently and drove capacitive load *C_L*.
2. Each inverter chain consisted of four inverters and switched independently. The capacitive load of the last inverter in each chain was *C_L*.
3. This was like the second configuration, but each inverter chain consisted of five inverters.
4. This was also like the second configuration, but in this case each inverter chain consisted of six inverters.

The total number of logic gates (inverters) was the same in all four configurations. We wanted to match the size of the circuit used in the Spice simulations to the size of ELF13 chip. Therefore, we performed the Spice simulations on 12,000 logic gates. We measured

Figure 2. Timing diagram of clock-stretching simulation.

Table 4. Clock network simulation conditions.

Simulation	Total <i>R</i> transmission (ohms)	Total <i>L</i> transmission (nH)	Total <i>C</i> transmission (pF)
1	100	100	10
2	500	100	10
3	1,000	100	10
4	2,000	100	10

voltage droop recovery time from the onset of input transition of logic gates to the time when the supply voltage value recovered to the nominal value (100% recovery). This recovery time depends on the position of the logic gates. In this experiment, we measured the worst-case recovery time. We set the L and C values of the power grid model equal to those of the ground model ($L = 0.1$ nH and $C = 0.2$ pF), and we varied the R values.

Figure 5 shows the 100% recovery times for the four configurations, and the 95% recovery time (the amount of time required to recover 95% of the nominal value) for Configuration 4, the one with the largest recovery time of the four configurations. The x -axis gives the different R and C_L values that we used in the simulations, and the y -axis represents measured voltage droop recovery time. In the worst case, the recovery time exceeded 150 μ s. The 95%

recovery time for Configuration 4 was about 140 μ s, or 1,000 idle cycles.

The simulation results in Figure 5 suggest that the power supply network, disturbed by either scan-shifting activity or the transition of the scan-enable signal, can take more than 1,000 idle cycles to recover fully (even the 95% recovery time takes 1,000 idle cycles). Tests performed during that recovery period may be subject to overly strenuous conditions due to power supply droop, which artificially limits their speed performance. Thus, the longer the capture clock is delayed, the greater the amount of IR drop that can be eliminated; and the greater the number of idle cycles inserted, the faster the chip can operate. This is consistent with ELF13 test results shown in Table 1: The greater the number of idle cycles inserted, the greater the number chips that pass the LD test.

Minimum V_{DD} as an indicator of IR drop

Another data point related to IR drop is the minimum V_{DD} value, which is the minimum supply voltage at which a chip operates correctly. Earlier work has demonstrated that structural test patterns that cause less IR drop make chips operate at a lower supply voltage.¹² Thus, the minimum V_{DD} measurement can serve as an indirect method to measure IR drop, which is otherwise very difficult to view directly. In this experiment, we measured minimum V_{DD} for sample chips that failed the regular LOC test set at the

nominal supply voltage but passed at an elevated supply voltage. The measured minimum V_{DD} values are thus higher than the nominal V_{DD} . We measured the minimum V_{DD} values for two groups of ELF13 chips (all of which failed the regular LOC test): those that failed the LOC LD tests (LD failing chips) and those that passed the LOC LD tests (LD passing chips). Seventy-seven percent of the LD passing chips had minimum V_{DD} values of less than 1.06 of nominal V_{DD} (that is, less than 6% extra supply voltage). On the

Figure 3. Clock period versus idle time for the four clock network simulations.

Figure 4. Simplified equivalent circuit for supply-ground network: single logic gate (a) and multiple logic gates (b).

other hand, 92% of the LD failing chips had minimum V_{DD} values greater than 1.06 of nominal V_{DD} (that is, greater than 6% extra supply voltage). This shows that chips with small minimum V_{DD} values are more likely to pass LD tests, meaning the additional time allotted to recover from IR drop during LD tests improves yield. Conversely, the benefit of LD test is limited by the maximum IR drop for each chip: The higher the minimum V_{DD} (because of large IR drop), the less likely a chip is to pass an LD test with a given number of idle cycles.

Using F_{max} and excess delay as recovery indicators

An alternative method of assessing the effect of scan-test-induced IR drop on yield is to quantify the size (in units of time) of the delays that cause LOC tests

to falsely fail. The excess delay, D_E , is the additional delay that causes a chip to fail an LOC test. We can calculate it using the following equation:

$$D_E = (1/F_{max}) - (1/F_S)$$

where F_{max} is the maximum operating frequency (the maximum frequency of a chip at which it passes a regular LOC test) and F_S is the system frequency (the speed of a chip during functional operation). When we ran this experiment on the ELF13 chips, we measured the excess delay for each chip that failed the regular LOC test. As before, there were two groups (LD passing and LD failing). Almost all the LD passing chips (94%) had excess delays smaller than 2 ns, but almost half (46%) of the LD failing chips had excess delays greater than 2 ns. As with minimum V_{DD} , excess delay is a good

Figure 5. Voltage droop recovery time. (The units of R and C_L are ohms and pF.)

indicator of the likelihood of improving yield by inserting idle cycles so that the power supply network can rebound from the IR drop induced from applying a scan test.

Comparing ELF13 and ELF18 test results

As Tables 1 and 3 show, the LD test results for the ELF13 differ from those for the ELF18. In ELF13 tests, LOC LD tests passed many chips (17 of 88), and SSF LD tests passed only one of 14 chips. In ELF18 tests, LOC LD tests passed only a few DSP cores (two of 141), but SSF LD tests passed many DSP cores (18 of 133). These somewhat inconsistent test results can be due to different test conditions and chip selections. The necessary conditions for a chip to pass SSF LD tests are as follows: First, a chip should have correct values before the single capture clock is applied. Second, a chip should not have a stuck-at defect.

The first condition implies that a chip can fail an SSF test (without idle insertion) because the single capture clock is applied before internal nodes have settled on the basis of the final scanned-in values. This settling between the last shift and the single capture clock can be assisted by adding idle cycles, during which chips can recover from IR drop. The ELF18 has a faster shift speed (20 MHz) than the ELF13 (7 MHz), whereas the ELF18 has a slower operating speed than the ELF13. Thus, ELF18 chips are designed to operate more slowly than ELF13 chips, but shift speed is faster

in ELF18 chips. Therefore, it is more difficult for ELF18 chips to meet the first condition, which explains why we can see more ELF18 DSP cores passing SSF LD tests than ELF13 chips. This is consistent with the ELF18 test results, which indicate that the scan-shifting activity rather than the scan-enable transition event causes the electrical disturbance; that is, chips pass both before and after LD. However, in ELF13 experiments, only one chip passed the SSF LD test, indicating that the

scan-enable transition rather than the scan-shifting activity causes false failures; that is, chips pass after LD but not before.

The reason the SSF LD test has only one passing ELF13 chip, whereas the LOC LD test has many (17) passing ELF13 chips, is that in ELF13 chips the scan-enable transition event causes the electrical disturbance. IR drop caused by the scan-enable transition increases a signal's propagation delay. The increased propagation delay can affect the LOC test results because the propagation of a signal is captured at the subsequent capture clock applied at a rated speed. However, this increased delay does not affect the SSF test results, because there is no capture-clock event that captures the late signal transition in the SSF test set.

The second condition implies that a chip cannot pass the SSF LD test, regardless of how much idle time is inserted, if it has a stuck-at defect. LD tests help chips recover from IR drop by adding idle cycles. But if an internal node is stuck at V_{DD} or ground due to a defect, a chip cannot pass an LD test even if that chip has fully recovered from IR drop. This condition is related to the ELF18 results in which few ELF18 DSP cores (two of 141) passed LOC LD tests.

In addition to the two conditions mentioned, a chip requires a third condition to pass an LOC LD test: Delays of the chip must be small enough for the chip to pass the LOC LD test. A chip cannot pass LOC LD even if it has fully recovered from IR drop if the delay after recovery is still larger than the chip's operation speed.

The ELF13 chips used for the LOC LD experiments passed 2-detect and LOC tests applied at a slow speed (10 MHz).³ Therefore, we can assume that all 88 ELF13 chips meet the first two conditions. On the other hand, only nine of 141 ELF18 DSP cores used for LOC LD passed the SSF test. So, we can assume that only nine of 141 ELF18 DSP cores meet the first two conditions. In other words, 132 ELF18 DSP cores (141-9) that failed the SSF test cannot pass LOC LD, because they may have stuck-at defects. Our test results show that the two ELF18 DSP cores that passed the LOC LD test came from the nine DSP cores that passed the SSF test; that is, no ELF18 DSP cores passed LOC LD if they failed the SSF test. However, 18 ELF18 DSP cores passed the SSF LD test, yet none passed LOC LD; this discrepancy could be due to the third condition just discussed. If we exclude SSF failing chips in considering the passing rate of LOC LD, both chip designs show similar rates (19.3% for ELF13 and 22.2% for ELF18).

EXPERIMENTS ON THE finer granularity of idle cycles would be helpful to identify the optimum idle cycles for yield improvement. Future work should include Spice simulations of the supply-ground network based on physical layout and measurements of IR drop within the chip. ■

Acknowledgments

This work was supported by the Semiconductor Research Corp. under contract 1175.001. We thank Nvidia and Philips for making the chips available and for the test support they provided. We thank Bill Price, Stefan Eichenberger, and Fred Bowen of Philips Semiconductors, and Bruce Cory of Nvidia. The experiments described in this article were conducted on the Verigy V93000 SoC tester.

References

1. J.A. Waicukauski et al., "Transition Fault Simulation," *IEEE Design & Test*, vol. 4, no. 2, Apr. 1987, pp. 32-38.
2. S. Wang and S.K. Gupta, "ATPG for Heat Minimization during Test Application," *IEEE Trans. Computers*, vol. 47, no. 2, Feb. 1998, pp. 256-262.
3. D. Lee and E.J. McCluskey, *Experimental Data on the Extent of Yield Loss*, annual report, Semiconductor Research Corp., 2006.
4. D. Lee and E.J. McCluskey, *Techniques to Identify the Potential Cause of Overkill*, annual report, Semiconductor Research Corp., 2005.

5. E.J. McCluskey et al., "ELF-Murphy Data on Defects and Test Sets," *Proc. 22nd IEEE VLSI Test Symp. (VTS 04)*, IEEE CS Press, 2004, pp. 16-22.
6. K.A. Brand et al., "Speed Clustering of Integrated Circuits," *Proc. Int'l Test Conf. (ITC 04)*, IEEE CS Press, 2004, pp. 1128-1137.
7. J. Rearick and R. Rodgers, "Calibrating Clock Stretch during AC Scan Testing," *Proc. IEEE Int'l Test Conf. (ITC 05)*, IEEE CS Press, 2005, pp. 1-8.
8. Y.I. Ismail, E.G. Friedman, and J. Neves, "Exploiting the On-Chip Inductance in High-Speed Clock Distribution Networks," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 9, no. 6, Dec. 2001, pp. 963-973.
9. P.J. Restle et al., "A Clock Distribution Network for Microprocessors," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, May 2001, pp. 792-799.
10. N. Srivastava, X. Qi, and K. Banerjee, "Impact of On-Chip Inductance on Power Distribution Network Design for Nanometer Scale Integrated Circuits," *Proc. 6th Int'l Symp. Quality of Electronic Design (ISQED 05)*, IEEE CS Press, 2005, pp. 346-351.
11. K.T. Tang and E.G. Friedman, "Simultaneous Switching Noise in On-Chip CMOS Power Distribution Networks," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, vol. 10, no. 4, Aug. 2002, pp. 487-493.
12. J. Saxena et al., "A Case Study of IR-Drop in Structured At-Speed Testing," *Proc. Int'l Test Conf. (ITC 03)*, IEEE CS Press, 2003, pp. 1098-1104.

Donghwi Lee is a PhD candidate in the Department of Electrical Engineering at the Center for Reliable Computing at Stanford University. He is also a DFT engineer at Nvidia in San Jose, California. His research interests include delay testing and ATPG. Lee has a BS in electrical engineering and a BS in materials science, both from Yonsei University, Korea, and an MS in electrical engineering from Stanford University.

Erik Volkerink leads the Labs Network at Verigy, serves as chief scientist at Memory Test Solutions, and is assistant director of the Center for Reliable Computing and a consulting assistant professor at Stanford University. His research interests focus on semiconductor design and test. Volkerink has a BS and an MS from the University of Twente, the Netherlands, and a PhD

from Stanford University, all in electrical engineering. He is a member of the IEEE.

Intaik Park is a PhD candidate in the Department of Electrical Engineering at the Center for Reliable Computing at Stanford University. His research interests include ATPG and test for delay defects. Park has a BS from Worcester Polytechnic Institute and an MS from Stanford University, both in electrical engineering.

Jeff Rearick is an AMD Fellow in the Mile High Design Center at Advanced Micro Devices in Fort Collins, Colorado, where he works as a DFT architect. His technical interests in-

clude a wide range of DFT strategies for leading-edge microprocessors, from memory and logic BIST to small-delay-defect ATPG, to embedded instrumentation for high-speed serial I/O circuits. Rearick has a BS from Purdue University and an MS from the University of Illinois at Urbana-Champaign, both in electrical engineering.

■ Direct questions and comments to Donghwi Lee, Center for Reliable Computing, Stanford University, Room 234, William Gates Hall, 353 Serra Mall, Stanford, CA 94305; dhlee@crc.stanford.edu.

For further information on this or any other computing topic, visit our Digital Library at <http://www.computer.org/publications/dlib>.

Here now from the IEEE Computer Society IEEE ReadyNotes

Looking for accessible tutorials on software development, project management, and emerging technologies? Then have a look at ReadyNotes, another new product from the IEEE Computer Society.

ReadyNotes are guidebooks that serve as quick-start references for busy computing professionals.

Available as immediately downloadable PDFs (with a credit card purchase), ReadyNotes sell for \$19 or less.
www.computer.org/ReadyNotes

IEEE

IEEE
computer
society