

LI-BIST: A Low-Cost Self-Test Scheme for SoC Logic Cores and Interconnects ^{*}

Krishna Sekar

Sujit Dey

Department of Electrical and Computer Engineering, University of California, San Diego
{ksekar, dey}@ece.ucsd.edu

Abstract

For deep sub-micron system-on-chips (SoC), interconnects are critical determinants of performance, reliability and power. Buses and long interconnects being susceptible to crosstalk noise, may lead to functional and timing failures. Existing at-speed interconnect crosstalk test methods are based on either (i) inserting dedicated interconnect self-test structures (leading to significant area overhead), or (ii) using existing logic BIST structures (e.g., LFSRs), which often result in poor defect coverage. Additionally, it has been shown that power consumed during testing can potentially become a significant concern.

In this paper, we present Logic-Interconnect BIST (LI-BIST), a comprehensive self-test solution for both the logic of the cores and the SoC interconnects. LI-BIST reuses existing LFSR structures but generates high-quality tests for interconnect crosstalk, while minimizing area overhead and interconnect power consumption. On applying LI-BIST to a DSP chip, we achieved crosstalk defect coverage of 99.7% for the interconnects and single stuck-at-fault coverage of 91.36% for the logic cores, while incurring an area overhead of only 4% over conventional LBIST.

1. Introduction

Advances in device technology have led to an era where entire systems can be implemented on a single chip, referred to as System-on-Chip (SoC). As SoC complexity grows with increasing integration and reducing feature sizes, the on-chip interconnect architecture, which is responsible for inter-core communication, plays a much more critical role since it starts dominating system performance [11] and power consumption [8]. Reliability of SoCs depends increasingly on the error-free operation of such interconnects. Testing of SoCs, hence, implies testing not only the logic cores but also the interconnect architecture.

The use of deep sub-micron (DSM) technology in SoCs increases the capacitive coupling between adjacent wires leading to severe crosstalk noise, which causes the functionality or performance of the chip to deviate significantly from expected behavior. Several physical design [14, 3] and analysis [9, 10] techniques have been developed to allow design for margin and to minimize signal integrity problems. However, these may be prohibitive in terms of design cost. Also, it is impossible to take into account all the possible process variations and physical defects during design. Hence, we need to address the crosstalk issue by means of testing techniques.

Increased cross-coupling capacitance between a pair of interconnects can produce either glitches or delays depending upon the signal transitions at the interconnects as shown in Figure 1(a) and 1(b) respectively. Also, previous studies [9, 5] have shown that crosstalk noise is more pronounced for long interconnects. Current and future SoCs will be dominated by a large number of very long interconnects and buses needed for the integration and communication among the cores in the chip [11]. In this paper, therefore, we focus on developing a comprehensive test solution to address the test needs of both the logic cores as well as the buses and global interconnects of an SoC.




Figure 1. Types of crosstalk faults

Crosstalk effects most adversely affect high performance circuits operating at GHz clock frequencies. At-speed testing is essential for testing such chips adequately since many crosstalk effects are not manifested at lower speeds. However, the gap between ASIC speeds and external testers' accuracy for timing signal resolution at ASIC pins is constantly growing [11]. Furthermore, test equipment having high speed, large pin count, large memory, and good timing accuracy can be prohibitively expensive. We, therefore, employ self-testing techniques to address the problem of testing for crosstalk in SoCs.

In recent years, the power consumption of digital systems during testing has become a major concern as it may increase significantly as compared to normal operational mode [15]. Also, empirical studies have shown that the power dissipation associated with long interconnects accounts for a significant fraction of the overall system power [8]. This power consumption is dominated by the increasing inter-wire cross-coupling capacitances in DSM technology [12]. The energy dissipated due to cross-coupling capacitances can vary depending on the type of transitions on the interconnects [12]. We, therefore, also focus on making our self-test scheme extremely power-efficient.

Bai *et al* [2], have proposed inserting dedicated interconnect self-test structures in the SoC to generate vectors which have 100% crosstalk defect coverage. This scheme is based on the Maximal Aggressor Fault Model proposed by Cuvillo *et al* [5]. However, this method has a prohibitively high area overhead. To reduce this overhead, existing logic BIST struc-

^{*}This work is supported by MARCO/DARPA Gigascale Silicon Research Center (GSRC).

tures, like linear feedback shift registers (LFSRs), could be reused to generate interconnect tests. But, LFSR vectors have very poor crosstalk defect coverage. In this paper, we address the issue of how to generate high quality interconnect crosstalk tests, efficiently reusing existing test structures so as to minimize the area overhead. Our proposed methodology, called *Logic-Interconnect BIST* (LI-BIST), produces high crosstalk defect coverage with low area penalty and low interconnect power consumption.

The interconnect fault model used in this paper is the Maximal Aggressor Fault Model that was reported and validated in [5]. Next, we briefly review this fault model, and the corresponding Maximal Aggressor test vectors.

1.1. Interconnect Crosstalk Fault Model

If the crosstalk problem is addressed at the process level, the number of possible process variations and physical defects that need to be considered even for a pair of interconnects is very large. For wide buses, considering all such variations is clearly prohibitive. At the circuit level, the cumulative effect of process variations can be described behaviorally by a coarser mesh of lumped circuit elements, but the resulting fault universe is still too large. Hence, we need an abstract fault model that can represent all crosstalk defects with a small number of faults.

The Maximal Aggressor Fault Model (MAFM) [5] is a functional fault model representing all the process variations and physical defects that lead to any of the following four crosstalk errors on a wire designated as the victim wire among the set of interconnects under test: positive glitch (g_p), negative glitch (g_n), rising delay (d_r) and falling delay (d_f). All the other wires are designated as aggressors and act collectively to generate the glitch or delay error on the victim. Figure 2 shows the transitions needed on the aggressors and victim wires to produce the maximum error effect for all four error types on the victim wire Y_i . These transitions constitute the Maximal Aggressor (MA) tests; they are necessary and sufficient for detecting the corresponding four crosstalk faults for the victim wire Y_i .




Figure 2. Maximal Aggressor Fault Model

From Figure 2 we see that the MA vectors are such that all the aggressor wires have a same direction transition while the victim wire remains at 0/1 or has an opposite direction transition. Such vectors cause the maximum cross-coupling capacitive effect on the victim wire. Hence, vectors which cause a majority of the wires to transition in the same direction and relatively few wires to remain at 0/1 or transition in the opposite direction (and thus exciting high cross-coupling

capacitive effect on the victim wires) achieve high interconnect crosstalk defect coverage.

1.2. Paper Outline

In Section 2, we briefly describe the interconnect test scheme proposed by Bai *et al* in [2]. In Section 3, we discuss how existing BIST structures could be reused to deliver interconnect tests, and the shortcomings thereof. These shortcomings are addressed in Section 4, where we propose a new test generator design which can generate vectors for both logic as well as interconnects and also present the overall test architecture for LI-BIST. Experimental results are presented in Section 5. Section 6 concludes the paper.

2. Interconnect Test using Dedicated Self-Test Structures

A self-test methodology for testing interconnects based on MAFM has been proposed by Bai, Dey & Rajski [2]. Here, we briefly describe their scheme.

The scheme is based on the fact that since the required Maximal Aggressor tests are known *a-priori*, if suitable self-test structures can be inserted in the SoC to generate all the required vectors, then the self-test methodology will be able to achieve 100% crosstalk defect coverage. For each core to core test transaction, the methodology requires a test generator in the interconnect interface of the source core and an error detector in the interconnect interface of the destination core. For example, in the SoC shown in Figure 3, to test the transaction from core $C1$ (CPU) to core $C3$ (RAM), a test generator is inserted at the output of the CPU core, and an error detector is inserted at the inputs of the RAM core. The test vectors are launched on the interconnect under test by the test generator of the source core and measured for logical consistency at the other end of the interconnect by the error detector in the destination core. Since the drivers and loads of the core play a crucial role in crosstalk noise, the test generators/error detectors are located between the core outputs/inputs and the core's buffer connections to the bus. A global test controller, which selects and activates appropriate test generators/error detectors is also described in the paper.




Figure 3. SoC with embedded self-test structures

The total hardware overhead of this scheme applied to a Digital Signal Processing chip, CMUDSP, as reported in the paper, is about 22% which is clearly prohibitive. Hence, although this scheme achieves 100% crosstalk defect coverage, yet it is infeasible due to the high area overhead involved.

It should be noted that the proposed self-test structures are used only to test the SoC interconnects. They are exclusive of any other test structures that may be present on-chip to test the logic of the SoC. If existing logic BIST structures could be efficiently reused to deliver interconnect tests, then the need for separate interconnect test generators/error detectors can be done away with and the area overhead kept at a minimum. We describe such a scheme in the next section.

3. Reuse of Logic BIST for Crosstalk Testing

In this section, we describe how logic BIST structures could be reused to deliver interconnect tests, and the shortcomings thereof. In the next section, we propose our new test generator design which addresses these shortcomings.

Consider an SoC in which logic BIST structures, consisting of Linear Feedback Shift Registers (LFSRs) and Multiple Input Signature Registers (MISRs), are used to test the logic of the circuit. Since in an SoC, each core is itself a large and complex circuit, all cores are assumed to have their own dedicated LFSR/MISR. In our proposed methodology, the LFSRs will be used to generate test vectors not only for the logic cores but also for the interconnects. Similarly, the MISRs will be used for compacting the responses of both the logic cores as well as the interconnects. There are two test phases, the logic test phase and the interconnect test phase. During the logic test phase, the cores are tested using their own LFSRs/MISRs as in traditional logic BIST. During the interconnect test phase, for each core-to-core test transaction, the LFSR of the source core generates the test vectors, which are then delivered onto the bus and they are compressed by the MISR at the destination core for signature analysis.




Figure 4. Testing for crosstalk defects on interconnects reusing Logic BIST structures

This scheme is illustrated by Figure 4. The example SoC contains two cores $C1$ and $C2$ that communicate using interconnects $I1$ and $I2$. Each core has a dedicated LFSR/MISR for logic BIST ($L1/M1$ for core $C1$ and $L2/M2$ for core $C2$). The figure shows an example logic BIST configuration. The core flip-flops are scanned and the I/Os are boundary scanned. There is a multiplexer at each core ($MUX1$ and $MUX2$) to choose whether the values going on the interconnects are test vectors generated by the LFSR or normal oper-

ational values from the core. Similarly there is a multiplexer ($MUX3$ and $MUX4$) to choose whether the MISR compresses the output responses of the logic core or the values on the interconnects. Suppose we are testing for crosstalk defects on $I1$, for the core-to-core transaction $C2 \rightarrow C1$. The vectors generated by $L2$ (the source core's LFSR) are driven on $I1$, and are compacted by $M1$ (the destination core's MISR). At the end of the test, MISR $M1$'s signature is analyzed to see if there was any crosstalk error for this particular transaction. Similarly, for testing the interconnect $I2$, for the core-to-core transaction $C1 \rightarrow C2$, the vectors are generated by $L1$ and compacted by $M2$. Note that the interconnect test scheme does not depend upon the actual logic BIST configuration as long as an LFSR and MISR are present at each core.

This scheme looks attractive since the hardware overhead incurred is very minimal. It essentially reuses existing logic BIST structures to act as test generators and error detectors for testing crosstalk defects in interconnects. But, unfortunately the vectors generated by an LFSR, though good for logic, have poor crosstalk defect coverage [13]. This is because the kind of vectors which achieve high stuck-at-fault coverage for logic is very different from those which achieve high interconnect crosstalk defect coverage. As shown in Section 1.1, good vectors for interconnect crosstalk test are such that they cause a majority of the wires to transition in the same direction and a few wires to remain at 0/1 or transition in the opposite direction, thus exciting high cross-coupling capacitive effect on the victim wires. The vectors produced by an LFSR are pseudo-random and do not have such a characteristic. Hence, they have poor defect coverage.

We would ideally like to use this framework of testing the interconnects (as it has low area overhead) but generate test vectors having much higher crosstalk defect coverage. Thus, our goal is to develop a test generator, which while not compromising on logic fault coverage, produces high quality interconnect crosstalk tests with minimal area overhead. We call this integrated self-test scheme for both logic and interconnects as LI-BIST which we present in the next section.

4. New Test Generator Design and LI-BIST Test Architecture

In this section, we first present the design of the test generator for LI-BIST and discuss the motivation behind such a design. We then present the test architecture of LI-BIST.

4.1. New Test Generator Design

We first discuss the desired properties of LI-BIST and the corresponding requirements of the test generator:

Low Cost: The test generator should be designed so as to maximize the reuse of existing self-test structures (LBIST) so that additional test circuitry area is minimized.

High Logic and Interconnect defect coverage: The new test generator should achieve as high logic fault coverage as existing logic BIST schemes. Also, it should produce high quality interconnect crosstalk tests. Hence, the profile of the interconnect test vectors produced by the test generator

should be such that the majority of the wires transition in the same direction (to act as aggressors) and relatively few wires remain 0/1 or transition in the opposite direction (to act as victims) so as to cause a high cross-coupling capacitive effect on the victim wires, as described in Section 1.1.

Low Power: The vectors generated by the test generator should be interconnect power efficient. Hence, we should minimize opposite direction transitions on adjacent wires, so that the *total amount* of coupling capacitances excited is low.

Figure 5 shows the structure of a 4 bit LI-BIST test generator. It consists of an LFSR that produces vectors for testing the logic core (as in traditional logic BIST), and an extension circuit that modifies the LFSR vectors such that they have high interconnect crosstalk defect coverage. In the extension circuit, the adjacent bits of the LFSR are both *ANDed* and *ORed*. For example, LFSR bits 2 and 3 are input to *AND2* and *OR2*. The outputs of the AND gate and the OR gate are connected to a 2:1 multiplexor. The multiplexor selects which of these values drives the corresponding output line. For example, *MUX2* selects whether the output line *L2* is driven by the output of *AND2* or *OR2*. The select line of all the multiplexors is driven by a toggle flip-flop. Hence, the value of an output line alternates at every cycle between the output of the respective *AND* gate and *OR* gate. For example, at clock cycle *i* if the *AND* gates drive the lines *L0* to *L3*, then the *OR* gates drive them at cycle *i+1*, the *AND* gates drive them at cycle *i+2*, and so on.

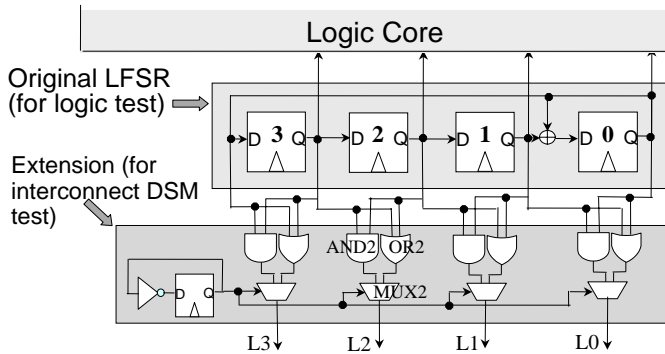


Figure 5. A 4 bit test generator for LI-BIST

With this small extension to a regular LFSR, the interconnect test vectors generated have the required profile as described earlier in this section. For an LFSR configured with a primitive polynomial, the probability of any bit *i* being 0, $P_i(0)$, is $1/2$; $P_i(1) = 1/2$. So the probability that the *AND* of two bits, *i* and *j*, is 0 is $3/4 (1 - P_i(1) * P_j(1))$. Similarly, the probability that the *OR* of two bits, *i* and *j*, is 0 is $1/4 (P_i(0) * P_j(0))$. Now, since the interconnect vectors generated are driven alternately by the *AND* gates and the *OR* gates, hence the probability that an output line is 0 varies alternately between $3/4$ and $1/4$. So, we are essentially generating weighted random patterns which are alternately weighted with a high probability of 0 and a high probability of 1. Hence, the profile of the interconnect test vectors generated will be such that many wires transition in the same

direction, and a few wires remain static or transition in the opposite direction; exactly the kind of vectors we require.

In order to minimize the correlation among the generated interconnect vectors due to the structural dependencies among the LFSR bits, ideally a separate LFSR should be used for each output line, all of them differently configured. However, such a scheme is expensive in terms of area. Our proposed scheme using only a single LFSR to generate all the bits, is more than adequate for our requirements since we do not really care what the exact probabilities are as long as alternate vectors have a high probability of 0 and 1. For the same reason, the seed value of the LFSR and the primitive polynomial used does not affect the crosstalk defect coverage much.

The size of the extension circuit which is added to the LFSR to appropriately weight the interconnect test vectors is relatively small. For testing an *n* bit bus, for example, we require *n* 2-input *AND* gates, *n* 2-input *OR* gates, *n* 2-bit *MUX* gates, 1 *DFF* and 1 *NOT* gate.

4.2. Test Architecture for LI-BIST

Figure 6 illustrates the test architecture of *LI-BIST* on an example SoC consisting of 2 cores. Each core is surrounded by a *Test Wrapper*, which consists of self-test structures like a Test Pattern Generator (TPG) (described in the previous section) and a MISR. The TPG generates test vectors for both the logic core as well as the interconnects. Multiplexors are used to select between the normal core outputs and the interconnect test vectors. The MISR compacts the output responses of both the core and the interconnects; this is selected by multiplexors. There is a centralized *LI-BIST Controller* which controls the whole test sequence. It gives control signals to all the test structures and is also responsible for seeding the TPG and unloading the MISR signature.

There are three operational modes of the SoC; normal mode, core test mode and interconnect test mode. In the normal mode, the SoC functions in its normal system operational mode. During the core test mode, the internal logic of the cores is tested via the TPG and the MISR, as in logic BIST. The core F/Fs are scanned and the I/Os are boundary-scanned for controllability and observability. In the interconnect test mode, for each core-to-core transaction, the vectors generated by the TPG of the source core are routed onto the interconnect and they are compacted by the MISR of the destination core. For example in Figure 6, there are three core-to-core transactions which are tested: $Core1 \xrightarrow{I_1} Core2$, $Core1 \xrightarrow{I_2} Core2$, $Core2 \xrightarrow{I_3} Core1$.

5. Validation of LI-BIST Methodology

To validate the LI-BIST methodology, we applied it to a Digital Signal Processing chip, CMUDSP [4], which corresponds to Motorola DSP56002. In this section, we first briefly describe the CMUDSP architecture and discuss how LI-BIST is applied to this chip. Next, we describe how the interconnect crosstalk defect coverage of LI-BIST is measured using a high-level interconnect DSM error model. Finally,

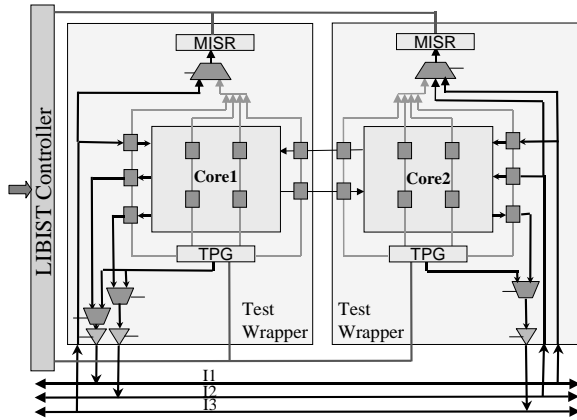


Figure 6. Test architecture of LI-BIST

we present our experimental results by comparing LI-BIST with the previous two schemes outlined in Sections 2 and 3.

5.1. Application of LI-BIST to CMUDSP

Figure 7 shows the CMUDSP chip, which consists of four components: Arithmetic and Data Logic Unit (ALU), Address Generation Unit (AGU), Bus Switch, and Program Control Unit (PCU). The ALU contains the X, Y, A and B registers along with the multiply-accumulate and adder units. The AGU generates the addresses for accessing the data memories. The PCU contains the program counter and flag bits for controlling the whole CMUDSP chip. The PCU also performs program address generation and instruction decoding. The Bus Switch is used to control data flow between buses. CMUDSP consists of four sets of separate data buses (XDB, YDB, PDB, GDB) and address buses (XAB, YAB, PAB). The X and Y buses are connected to the data memory, and the P buses are connected to the program memory.

In order to apply LI-BIST, each of the components is enclosed in a test wrapper as described in Section 4.2. The test wrapper contains test structures for testing both the cores as well as the interconnects. An LI-BIST controller is also inserted, which interfaces with the test wrappers by means of control and data signals, as shown in Figure 7. The LI-BIST controller initiates logic or interconnect tests when it receives corresponding external signals. At the end of the tests, it signals whether the chip is good or not through the error flag.

5.2. Fault Simulation Methodology

To calculate the interconnect crosstalk defect coverage of LI-BIST for the CMUDSP chip, we have to inject and simulate crosstalk defects in the interconnects, and examine how many of these defects are detected by the interconnect test vectors. The most accurate way of doing this is spice-level simulation. However, it is not feasible to simulate the entire chip at this level since it takes prohibitively long. Hence, we use a high-level crosstalk defect simulation method [1], which allows simulation of the crosstalk effects on the interconnects, together with the HDL models of the rest of the chip components including the test structures. We describe this fault-simulation methodology briefly here.

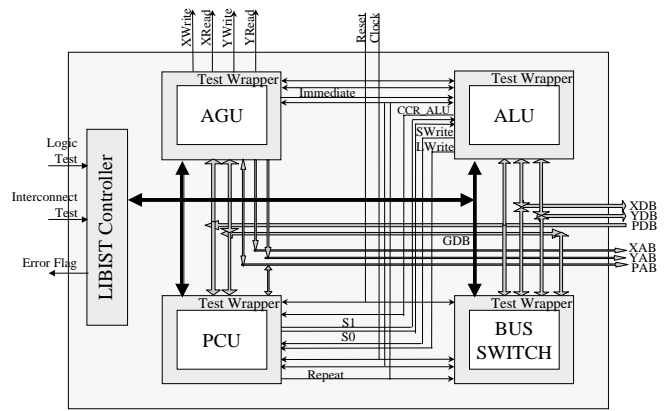


Figure 7. Application of LI-BIST to CMUDSP

The high-level crosstalk defect simulation environment for the CMUDSP chip is shown in Figure 8. Behavioral level interconnect DSM error models [1] are inserted corresponding to the interconnects under test so that we can inject crosstalk defects in interconnects and simulate the behavior of defect effects (like glitches and delays) at the high-level. Coded in HDL, the error model takes as input a parameter file containing the values of the coupling capacitances among the interconnects. Given an input transition on the driver end of the bus, the error model determines whether a crosstalk error occurs at the receiver end. A crosstalk defect can be injected by perturbing the values and distribution of cross-coupling parameters beyond a threshold value [5] to reflect process variation beyond design margins.




Figure 8. Crosstalk defect simulation environment

Fault-simulation is performed as follows: An HDL test generator of the source core generates the test vectors for a bus, which are input to the corresponding DSM error model. Depending upon the test vector transitions and the coupling parameters of the bus from the parameter file, the model generates output vectors, which may or may not contain digitally encoded glitch or delay errors. Then, the MISR at the destination core compresses these output vectors. At the end of the test, the MISR signature is analyzed to see whether the injected defect was detected or not. To estimate the defect

coverage, the same defect simulation process is repeated on all defects from a preconstructed defect library.

To generate the defect library, we randomly perturb the nominal values of coupling capacitances among the interconnects according to a given defect distribution. Given the resulting perturbation, we use the criteria in [5], to determine if the perturbation is large enough to be detected by *any* tests. If so, we record the perturbation as a defect. This process is repeated until a satisfactory number of defects are generated.

In our experiments, we used a Gaussian distribution to model the defect distribution in terms of the variation of capacitance values (in %). A standard variance of 50% was chosen. A total of 1000 defects were generated for each bus.

5.3. Experimental Results

In this sub-section, we compare the LI-BIST scheme with the MA Test scheme described in Section 2 and the LBIST scheme described in Section 3.

Table 1 compares the three schemes in terms of the additional area overhead over conventional LBIST for testing the interconnects, the interconnect crosstalk defect coverage, and the logic fault coverage for the CMUDSP chip. The LBIST scheme and the LI-BIST scheme were both implemented on the CMUDSP. The values for the MA Test scheme are reported from [2]. We used Synopsys' Design Compiler [6] synthesis tool to synthesize CMUDSP along with the inserted test structures. The area overhead is in terms of the additional number of gates (over conventional LBIST) required to test the interconnects of the chip. The defect coverage is measured using the high-level fault-simulation methodology described in the previous sub-section. The logic fault coverage is in terms of single stuck-at-fault coverage and is measured using Mentor Graphics' FastScan [7].

From the Table, we observe that the MA tests have 100% defect coverage, but with a high area penalty. The LBIST tests using LFSRs are not feasible due to their poor defect coverage. The LI-BIST scheme, however, achieves very high defect coverage with very low area overhead. Both LBIST and LI-BIST achieve the same logic fault coverage since the test methodology for logic is the same in both the schemes.

Table 1. Comparison of LI-BIST with MA Tests and LBIST

Test Scheme	Area Overhead over LBIST (%)	Defect Coverage (%)	Logic Fault Coverage (%)
MA Tests	22	100	-
LBIST	1.22	1	91.36
LI-BIST	4	99.7	91.36

In order to assess the power-efficiency of LI-BIST, we conducted another set of experiments to measure the average interconnect power consumption of all the three schemes. The results are presented in Table 2. The interconnect power consumption was measured using a DSM-accurate power estimation technique proposed in [12].

The LFSR test vectors consume the maximum amount of interconnect power since the vectors are pseudo-random and

Table 2. Interconnect Power Consumption

Test Methodology	Avg. Int. Power (mW)
MA Tests	0.43
LBIST	3.66
LI-BIST	0.77

they have many opposing direction transitions which excite a large *amount* of coupling capacitance. The MA vectors have the least power consumption because all the wires, except the victim, transition in the same direction. So, although there is a high coupling capacitive effect on the victim wire, the total amount of coupling capacitance excited is very low. The LI-BIST generated vectors have marginally higher interconnect power consumption as compared to the MA vectors since there may be more than one victim wires. So, the LI-BIST test methodology is also power-efficient.

6. Conclusion

In this paper, we presented LI-BIST, a comprehensive SoC test methodology for both logic cores and interconnects. It efficiently reuses existing on-chip test structures to generate high quality interconnect crosstalk tests. We have compared LI-BIST with existing solutions and validated it using a high-level crosstalk error model. Experiments on a DSP chip confirm that LI-BIST yields high crosstalk defect coverage, low area overhead and low interconnect power consumption.

References

- [1] X. Bai and S. Dey. High-level crosstalk defect simulation for system-on-chip interconnects. In *Proc. IEEE VLSI Test Symposium*, Apr. 2001.
- [2] X. Bai, S. Dey, and J. Rajsiki. Self-test methodology for at-speed test of crosstalk in chip interconnects. In *Proc. Design Automation Conf.*, pages 619–624, June 2000.
- [3] Z. Chen and I. Koren. Crosstalk minimization in three-layer HVH channel routing. In *Proc. IEEE Int. Symp. on Defect and Fault Tolerance in VLSI Systems*, pages 38–42, 1997.
- [4] The Carnegie Mellon Synthesizable Digital Signal Processor Core (<http://www.ece.cmu.edu/lowpower/benchmarks.html>).
- [5] M. CuvIELLO, S. Dey, X. Bai, and Y. Zhao. Fault modeling and simulation for crosstalk in system-on-chip interconnects. In *Proc. Int. Conf. Computer-Aided Design*, pages 297–303, Nov. 1999.
- [6] Design Compiler 2000.05, Synopsys Inc. (<http://www.synopsys.com>).
- [7] FastScan v8.6.4.5, Mentor Graphics Corp. (<http://www.mentor.com>).
- [8] R. Mehra, L. M. Guerra, and J. M. Rabaey. A partitioning scheme for optimizing interconnect power. *IEEE J. Solid-State Circuits*, 32(3):433–443, Mar. 1997.
- [9] P. Nordholz, D. Treytnar, J. Otterstedt, H. Grabinski, D. Niggemeyer, and T. W. Williams. Signal integrity problems in deep submicron arising from interconnects between cores. In *Proc. IEEE VLSI Test Symposium*, pages 28–33, April 1998.
- [10] K. Rahmat, J. Neves, and J. Lee. Methods for calculating coupling noise in early design: a comparative analysis. In *Proc. Int. Conf. Computer Design VLSI in Computers and Processors*, pages 76–81, 1998.
- [11] Semiconductor Industry Association. International technology roadmap for semiconductors, 1999.
- [12] C. N. Taylor, S. Dey, and Y. Zhao. Modeling and minimization of interconnect energy dissipation in nanometer technologies. In *Proc. Design Automation Conf.*, June 2001.
- [13] Y. Zhao and S. Dey. Analysis of Interconnect Crosstalk Defect Coverage of Test Sets. In *Proc. Int. Test Conf.*, pages 492–501, Oct. 2000.
- [14] H. Zhou and D. F. Wang. Global routing with crosstalk constraints. In *Proc. Design Automation Conf.*, pages 374–377, 1998.
- [15] Y. Zorian. A distributed BIST control scheme for complex VLSI devices. In *Proc. IEEE VLSI Test Symposium*, pages 4–9, Apr. 1993.