

Online BIST and BIST-Based Diagnosis of FPGA Logic Blocks

Miron Abramovici, *Fellow, IEEE*, Charles E. Stroud, *Senior Member, IEEE*, and John M. Emmert, *Senior Member, IEEE*

Abstract—We present the first online built-in self-test (BIST) and BIST-based diagnosis of programmable logic resources in field-programmable gate arrays (FPGAs). These techniques were implemented and used in a roving self-testing areas (STARs) approach to testing and reconfiguration of FPGAs for fault-tolerant applications. The BIST approach provides complete testing of the programmable logic blocks (PLBs) in the FPGA during normal system operation. The BIST-based diagnosis can identify any group of faulty PLBs, then applies additional diagnostic configurations to identify the faulty look-up table or flip-flop within a faulty PLB. The ability to locate defective modules inside a PLB enables a new form of fault-tolerance that reuses partially defective PLBs in their fault-free modes of operation.

Index Terms—Adaptive computing, built-in self-test (BIST), fault diagnosis, fault tolerance, field-programmable gate arrays (FPGAs), online testing.

I. INTRODUCTION

FIELD-PROGRAMMABLE gate arrays (FPGAs) that can be reprogrammed in a system and allow a system using reconfigurable hardware to adapt to changes in its external environment, are used to extend its initial capabilities by implementing new functions on the same hardware platform. In-system reprogrammability results in increased functional density and reduced power consumption—features very important in many domains, such as space missions or mobile devices. The reprogrammability and the regular structure of FPGAs are ideal to implement low-cost fault-tolerant hardware, which makes them very useful in systems subject to strict high-reliability and high-availability requirements, such as systems deployed in harsh and/or hostile environments or in remote unmanned long-life missions or systems whose operation is critical and must continue uninterrupted. These types of applications rely on online testing to protect against both transient failures and permanent faults that appear during the lifetime of the system. Online testing must also provide high-resolution

diagnosis, because to bypass faulty resources, they must first be accurately located.

However, most previous methods for FPGA testing address only offline testing using either externally applied vectors [1]–[9] or built-in self-test (BIST) [10]–[17]. Online FPGA testing methods based on coding techniques [18], [19] do not diagnose the detected faults. Online testing based on triple (or multiple) modular redundancy have a very high overhead. The “fault-scanning” method [20] achieves low-overhead online testing, but it is applicable only to bus-based FPGA architectures, it assumes that certain parts of the FPGA are fault free, and it requires a modified design of the FPGA blocks.

One problem with conventional online testing is that it detects only faults that affect the current operation performed in the system. New faults, however, are equally likely to occur in spare resources or in the currently unused portion of the operational part of the system. These dormant faults may accumulate and have a very detrimental effect on the system reliability [21]. Thus, to guarantee a reliable operation, the online test must completely check all the system resources, including spares.

In this paper, we present new techniques for online testing and diagnosis of nontransient faults in programmable logic blocks (PLBs), developed as part of the *roving self-testing area (STARs)* project [22], [23]. A STAR, is a temporarily offline section of the FPGA where self-testing goes on without disturbing the normal system activity in the rest of the chip. The roving of the STARs periodically brings every section of the device under test. This approach guarantees complete testing of the FPGA (PLBs, programmable interconnect, and configuration memory), including all spare resources, and does not require any part of the chip to be fault free. Recently, the roving STARs approach has been extended for testing delay faults [24]. Our techniques are applicable to any FPGA supporting incremental run-time reconfiguration, which is the ability to reconfigure only a part of the FPGA, while normal operation continues uninterrupted in the rest of the device. FPGAs supporting this feature include ORCA (from Lattice) and Virtex and Spartan (from Xilinx).

Our online BIST approach detects any combination of multiple faulty PLBs, with the exception of a few pathological cases unlikely to occur in practice. Our diagnosis approach is based on analyzing the BIST results, followed (when needed) by an adaptive divide-and-conquer method to obtain maximum diagnostic resolution. These techniques uniquely identify any combination of multiple faulty PLBs, again with the exception of a few pathological cases. Additional BIST-based diagnostic techniques identify the defective parts inside a faulty PLB, such as a faulty look-up table (LUT) or flip-flop. We used ORCA

Manuscript received November 11, 2002; revised December 17, 2003. This work was supported by the Defense Advanced Research Projects Agency (DARPA) Adaptive Computing Systems program under Contract F33615-98-C-1318.

M. Abramovici is with Design Automation for Flexible Chip Architectures (DAFCA), Framingham, MA 01701 USA (e-mail: miron@dafca.com).

C. E. Stroud is with the Department of Electrical and Computer Engineering, Auburn University, Auburn, AL 36849 USA (e-mail: cestroud@eng.auburn.edu).

J. M. Emmert is with the Department of Electrical and Computer Engineering, University of North Carolina at Charlotte, Charlotte, NC 28223 USA (e-mail: jm Emmert@uncc.edu).

Digital Object Identifier 10.1109/TVLSI.2004.837989

TABLE I
ACRONYMS AND ABBREVIATIONS

Acronym	Meaning
BIST	Built-In Self-Test
BISTER	group of PLBs forming a BIST unit
BUT	Block Under Test
FPGA	Field-Programmable Gate-Array
H-STAR	Horizontal STAR
LUT	Look-Up Table
ORA	Output Response Analyzer
ORCA	Optimized Reconfigurable Cell Array
PLB	Programmable Logic Block
PUB	Partially Usable Block
STAR	Self-Testing ARea
TPG	Test Pattern Generator
TREC	Test and REconfiguration Controller
V-STAR	Vertical STAR

[25] for our implementation, but we emphasize that our technique can be applied to any FPGA that features in-system incremental run-time reconfiguration, such as Virtex and Spartan series FPGAs.

The remainder of this paper is organized as follows. Section II outlines the roving STARs approach and describes the online BIST method. The fault-detection capabilities of the approach is described in Section III, and Section IV presents diagnosis based on analysis of the BIST results. Section V discusses diagnosis of faults internal to PLBs, and Section VI presents experimental results obtained on known defective ORCA FPGAs. Finally, Section VII presents conclusions and suggestions for future research directions. The acronyms and abbreviations used in the paper are summarized in Table I.

II. THE ROVING STARS

We assume that an embedded processor controls the test, diagnosis, and fault-tolerance functions associated to all the FPGAs in the system. This test and reconfiguration controller (TREC), which has separate processor-specific fault-tolerance mechanisms [26], also maintains all the configurations that may be used in the future. The roving STARs approach targets nontransient faults that appear during the lifetime of the system. Here a fault is one faulty PLB that may have several arbitrary internal faults. Once faulty resources have been detected and identified, our fault-tolerant techniques, presented in [27], reuse faulty logic resources whenever possible, and bypass faults using fault-free spare resources otherwise. As a result, any combination of faulty logic resources can be tolerated given sufficient spare resources in the FPGA. For online testing of transient faults, the application logic implemented in FPGAs employs a concurrent error-detection technique, such as the one described in [19]. Fault tolerance for transient faults is achieved by periodically saving the state of the system (checkpointing) and restoring the last saved state when a transient is detected.

Fig. 1(a) illustrates the initial floor plan of an FPGA, having a contiguous working area and a spare area consisting of a vertical STAR (V-STAR) with two columns and an horizontal STAR (H-STAR) with two rows. A STAR tests both the PLBs and the programmable interconnect in its area, but in this paper, we




Fig. 1. Roving STARs: (a) initial position and (b) during roving.




Fig. 2. BISTER tile for PLBs.




Fig. 3. Integrated ORA/scan cell.

will focus only on PLB testing. A roving step involves relocating a slice of the system logic adjacent to the STAR in the current STAR position and recreating the STAR logic in the area vacated by the system logic. This is illustrated in Fig. 1(b) after one roving step across of the V-STAR and one roving step down of the H-STAR. Testing the entire FPGA takes one full horizontal sweep of V-STAR and one full vertical sweep of H-STAR. While only one STAR would have been sufficient to test all the PLBs in the FPGA, having two roving STARs makes possible complete testing of interconnect [28] and also provides essential features for diagnosis and fault tolerance [27]. For an $N \times N$ FPGA, a full sweep of one STAR requires $N/2$ positions; thus, the total number of roving positions is $N^2/4$. The roving process is based on run-time reconfiguration and it is controlled by TREC [22].

A STAR consists of several disjoint tiles, with the PLBs in each tile configured as shown in Fig. 2, in a structure called a BISTER. All BISTERS in the same STAR work concurrently. A BISTER inherits some of the concepts used in our offline FPGA BIST methods for PLBs [11], [14], [29]. Specifically, the BISTER contains PLBs configured as a test pattern generator (TPG) that applies pseudoexhaustive test patterns to two identically configured blocks under test (BUTs), whose outputs are compared by an output response analyzer (ORA). The ORA latches and reports mismatches as test failures. Direct comparison avoids possible aliasing errors introduced by compressing responses into signatures [30]. The additional testing typically required in comparison-based testing to ensure that the

comparator is fault free is not needed in our approach, since every PLB in turn is tested as a BUT. The TPG and ORA are reset by TREC just prior to initiating the BIST sequence. Fig. 3 illustrates an ORA comparing four pairs of outputs (denoted O1 through O4) from two BUTs (denoted B1 and B2). The flip-flop stores the result of the comparison, and the feedback loop latches the first mismatch in the flip-flop. In the *Pass/Fail* result, flip-flops of all ORAs are connected to form the scan chain [11]. A *test phase* is defined as a test configuration that tests a BUT in one of its modes of operation, while a *test session* is defined as the set of test phases that test a BUT in all of its modes of operation [29]. In every test phase, the scan chain is also tested, to assure the integrity of the test results.

The TREC accesses the FPGA using its boundary-scan mechanism [31], so that this access is transparent to the normal function of the chip (most FPGAs support boundary scan). The TREC employs the boundary-scan interface to reconfigure the STARs for different test and diagnosis operations, to initiate the BISTERS, and to scan out the test results. BISTERS work with the boundary-scan test clock, *TCK*, and the configuration clock is also derived from *TCK*. The first configuration of a BISTER checks the proper operation of the scan register, inducing mismatches by comparing BUTs with different configurations. This also protects against the case of all ORA flip-flops being stuck at the “Pass” value. The TREC also controls the system clock(s) and has the capability to stop the system operation for short intervals to allow for safe relocation of a slice of the system logic in the last tested area. If faults are detected, the TREC starts the diagnosis process.

III. DETECTING FAULTY PLBs

The BUTs are repeatedly configured to be tested in all their modes of operation. In order to minimize fault latency, it is important to minimize the number of BIST configurations required to completely test the PLBs; this is particularly important since the configuration download time is much greater than the BIST application time. The modes of operation of a PLB may be determined only from the information available in the FPGA data book, without having a detailed knowledge of its implementation. Since the BISTER provides complete testing only for the BUTs, we have to reconfigure every BISTER several times so that each PLB will be a BUT in at least one configuration. The number of PLBs for an ORA and for a TPG depend on the target PLB architecture. In our implementation in the ORCA FPGA, a TPG needs three PLBs and an ORA only one. The same numbers of PLBs are needed for Virtex and Spartan II FPGAs.

Fig. 4(a) illustrates six floor plans of a 3×2 BISTER tile, where T, B, and O denote, respectively, a TPG cell, a BUT, and an ORA. The goal of the six configurations is to systematically rotate the functions of the PLBs, so that eventually every PLB in the tile is completely tested twice, each time being compared with a different BUT. This rotating strategy assures that every single faulty PLB and almost all combinations of multiple faulty PLBs are guaranteed to be detected; these results are proven next.

Claim 1: Any single faulty PLB is guaranteed to be detected in at least two BISTER configurations.




Fig. 4. Rotations for a 3×2 BISTER tile: (a) 6 rotations; (b) with two faulty BUTs in one BISTER configuration; and (c) one faulty BUT in one BISTER configuration.

Proof: The faulty PLB is a BUT in two BISTER configurations, where exhaustive input patterns are produced by a fault-free TPG, and its outputs are compared with a fault-free BUT by a fault-free ORA. Hence, no fault (single or multiple) detected in the BUT can escape detection in these two configurations. ■

Claim 2: Except for a few pathological cases, any pair of faulty PLBs is guaranteed to be detected in at least one BISTER configuration.

Proof: Since any single faulty PLB is detected, a pair of faulty BUTs will escape detection only if they have a circular masking relation [32], where each faulty PLB masks the detection the other one. This masking should occur in any BISTER configuration where a single faulty PLB would be detected. Because a TPG or an ORA containing a faulty PLB may still work correctly, we analyze only configurations where at least one of the faulty PLBs is a BUT.

Case 1) Both faulty PLBs are BUTs in one BISTER configuration. Then the TPG and the ORA are fault-free, and the only way the faulty pair can escape detection is to have functionally equivalent faults, since then the compared outputs will not mismatch. Assume that the faulty PLBs are BUTs in the configuration 1 in Fig. 4(b). Then in configurations 3 and 5, one of them is a BUT (compared with a fault-free BUT by a faultfree ORA), and the other is a TPG cell. The only way the faulty pair can escape detection is for the faulty TPG to skip exactly the patterns that detect the faulty BUT. We say that this is a pathological case, because it has an extremely low probability of occurrence: we need two faulty PLBs and they must have functionally equivalent faults, and when a faulty PLB is part of the TPG, the TPG must skip those patterns that detect the faulty BUT.

Case 2) At most one faulty PLB is a BUT in any BISTER configuration. Assume, without loss of generality, that the faulty cells are in the first row of the BISTER, as shown in Fig. 4(c). Let X be the faulty BUT and Y the faulty TPG cell in configuration 1. To have

1	2	3	4	5	6	7	8
T S	T S	S B	B O	O B	B T	S T	S T
T S	S B	T O	T B	B T	O T	B T	S T
T B	T O	T B	T S	S T	B S	O T	B T
B O	T B	T S	T S	S T	S T	B S	O B

 Fig. 5. 4×2 BISTER tile rotations.

circular masking between X and Y , all of the following conditions must be true: 1) in configuration 1, the TPG must skip those patterns that detect X ; 2) in configuration 4, the TPG must skip those patterns that detect Y ; 3) in configuration 2, the ORA X must not record the mismatch between Y and the fault-free BUT; and 4) in configuration 3, the ORA Y must not record the mismatch between X and the faultfree BUT. Again, this is a pathological situation because we need the AND of four conditions, which are very unlikely by themselves. ■

Clearly, similar arguments can be made for more than two faulty PLBs. Hence we can conclude that:

Claim 3: In practice, any combination of faulty PLBs is detected in at least one BISTER configuration.


To overcome routing congestion problems, spare PLBs may be incorporated into the BISTER tile to provide additional routing resources. Here, “spare” means “not involved in testing,” as all the PLBs in a STAR are “spare” with respect to normal operation. Fig. 5 illustrates the eight rotations of the 4×2 BISTER tile with two spare cells (denoted by S) used in our implementation. While these are not “pure” rotations, since they maintain the three TPG cells in the same column, they do achieve the same property as the rotations of the 3×2 tile. The presence of the spare cells makes circular masking relations even more unlikely, since any masking relation between two faulty PLBs disappears whenever one of them becomes a spare.

IV. LOCATING FAULTY PLBs

In every BIST configuration, the TREC records the set of failures obtained at the ORA outputs. Here we present a diagnosis method that attempts to locate faulty PLBs based on analyzing the test results, and we will show that in many cases, accurate resolution can be obtained without additional reconfigurations. Most importantly, the frequent case of a single faulty PLB falls in this category.

In any fault location procedure, maximum diagnostic resolution is achieved when faults are isolated within an equivalence class containing all the faults that produce the observed response. If every equivalence class has only one fault, we say that the fault is *uniquely diagnosed*, that is, there is no other fault that can produce the same response. In our case, a fault is one faulty PLB that may have several internal faults, and the response is the set of failing test phases detected at the outputs of the ORAs. We begin by assuming a single faulty PLB in the FPGA, then we analyze the case of multiple faulty PLBs; in Section V we will discuss locating faults inside a defective PLB.

We can observe that the PLBs in a BISTER tile can be partitioned into two disjoint sets, so that every set contains the PLBs


 Fig. 6. Combined test sessions for (a) 3×2 tile and (b) 4×2 tile.

that are pairwise compared when configured as BUTs. For example, for the 3×2 tile, the BUTs in configurations 1, 3, and 5, are compared only among themselves and they are never compared to the BUTs in phases 2, 4, and 6 [see Fig. 4(c)]. A similar partition exists for the 4×2 tile as well. We can represent the relations between BUTs that are pairwise compared and the ORAs that observe them by the graphs shown in Fig. 6, where a BUT is denoted by a node B_i , and the ORA that observes BUTs B_i and B_j , is denoted by O_{ij} . We can view such a graph as representing a combined test session in which half of the BUTs in the tile are concurrently tested. Each tile is completely tested in two combined test sessions, where the BUTs in one session are ORAs in the other one, and *vice-versa*.

Theorem 1: Any single faulty PLB is guaranteed to be uniquely diagnosed.

Proof: When the faulty PLB is configured as a BUT, it is detected at its two adjacent ORAs, and no other BUT is detected at the same two ORAs. Note that when the faulty PLB is configured as a TPG cell, no error is generated even when the fault modifies the TPG patterns, because the two BUTs still receive the same patterns. When configured as an ORA, the faulty PLB may generate an error, but this will be in addition to the failures observed in the combined session when it is a BUT, and these results are sufficient to distinguish it from any other single faulty PLB (in Section VI we will describe this situation as encountered in an actual faulty FPGA). ■

As a consistency check, we can verify that the sets of failures obtained at the two ORAs are identical.

The following results deal with the location of a group G of faulty PLBs that is detected in at least one test session. To uniquely diagnose G means to identify all its faulty blocks such that no other group (including subsets of G) can produce the same result. Although unique diagnosis is not always possible, there are many situations when it can be guaranteed.

We will analyze the interactions among faulty PLBs under the following assumption:

Assumption A1: A TPG with faulty PLBs does not skip the patterns that detect faults in a BUT. (Later in this section we will analyze the situation when this assumption is not true.)

Based on this assumption, the set of failing phases obtained at ORA O_{ij} is given by

$$FO_{ij} = (FB_i \cup FB_j) - Feq_{ij} \quad (1)$$

where FB_k is the set of failing phases of BUT B_k , and Feq_{ij} is the set of failing phases of both B_i and B_j that have identical responses (and thus, do not cause mismatches at O_{ij}). In Fig. 7,




Fig. 7. Graphical representation of (1).

the area of FO_{ij} is marked by diagonal lines. Note that FO_{ij} is empty (\emptyset) when both B_i and B_j are fault free, or when the faults in B_i and B_j are equivalent (since then $FB_i = FB_j = Feq_{ij}$). Also note that there exists one situation when FO_{ij} is the same, no matter if only one or both of the BUTs observed at O_{ij} are faulty; this occurs if the two BUTs have the same sets of failing phases ($FB_i = FB_j$), but their faulty responses are never the same ($Feq_{ij} = \emptyset$).

Knowing the set of failing phases observed at O_{ij} and the complete set of failing phases of one of the two faulty BUTs, we can determine the set of failing phases where the two BUTs have identical responses by

$$Feq_{ij} = FB_i - FO_{ij} = FB_j - FO_{ij}. \quad (2)$$

Based on FO_{ij} and one of the two sets, we can also compute a lower bound on the other set by

$$FB_i \supseteq (FO_{ij} - FB_j) \cup Feq_{ij} \quad (3)$$

or by

$$FB_j \supseteq (FO_{ij} - FB_i) \cup Feq_{ij}. \quad (4)$$

Note that (3) becomes an equality when FO_{ij} and FB_j are disjoint

$$FB_i = FO_{ij} \cup Feq_{ij}. \quad (5)$$

This occurs when $FB_i \subseteq FB_j$ and $Feq_{ij} = FB_i$.

For the diagnosis procedure, we make one more assumption:

Assumption A2: No more than two faulty BUTs have identical responses in the same failing phase.

To justify this assumption, recall that in every configuration the TPG applies exhaustive tests for that mode of operation. If assumption A2 is not satisfied, it means that we would have three or more BUTs with equivalent faults for that mode of operation, which is quite an unlikely event.

The following results will be used by our diagnosis procedure.

Lemma 2: If none of the two ORAs observing BUT B fails in phase p , then B does not fail phase p .

Proof: Assume, by contradiction, that B fails phase p . Based on Assumption A1, the vectors detecting its faults are generated by the TPG. Then the other two BUTs observed by the two ORAs must have equivalent faults to the fault of B in phase p , because p is not reported by either of the two ORAs. But then we would have three BUTs with equivalent faults in the same phase, which contradicts Assumption A2. Therefore, B does not fail phase p . ■




Fig. 8. Graphical representations for \blacklozenge for: (a) Theorem 2, (b) Lemma 7, and (c) Lemma 8.

Lemma 3: If the two ORAs observing the same BUT report no failures, then their common BUT is fault-free.

Proof: Based on Lemma 2, the common BUT of the two ORAs does not fail any phase. ■

Lemma 4: Any BUT failure appears at least at two ORAs.

Proof: Assume, by contradiction, that phase p is reported only at one ORA. Let B be the BUT failing p and O be the other ORA observing B . Since O does not report p as a failure, the other BUT (say C) observed by O must have identical response in p . But the other ORA observing C does not report p as a failure either, so we have more than two BUTs with equivalent faults in p , which contradicts assumption A2. Therefore, p must appear at the output of at least two ORAs. ■

Lemma 5: If we have failures only at two ORAs, they must have identical failures.

Proof: Based on Lemma 4, any failure must appear at least twice, and in this case we have only two failing ORAs. Hence, their failures must be identical. ■

Theorem 2: (For the 4×2 tile) If only two ORAs observing the same BUT report failure in phase p , their common BUT fails in phase p .

Proof: In Fig. 8(a), without loss of generality, assume that O_{12} and O_{14} are the two failing ORAs reporting p as a failure. Assume, by contradiction, that their common BUT B_1 does not fail phase p . Then both B_2 and B_4 must fail p . But since no other ORA reports p , then B_3 must also fail in phase p , and its response must be identical to that of B_2 and B_4 . Hence, we would have three BUTs with equivalent faults in phase p , which contradicts assumption A2. Therefore, B_1 must fail phase p . ■

Lemma 6: If two ORAs observing the same BUT report no failures, then all three BUTs observed by them are fault free.

Proof: Based on Lemma 3, the common BUT of the two ORAs is fault-free. So each one of the other BUTs is compared with a fault-free one by an assumed fault-free ORA, and no failure is reported. Hence, the other two BUTs are also fault free. ■

Lemma 7: (For the 4×2 tile) If only two ORAs without a common BUT fail phase p , then at least one pair of BUTs between the two ORAs are faulty and have identical response in phase p .

Proof: In Fig. 8(b), without loss of generality, assume that O_{12} and O_{34} are the two ORAs without a common BUT failing phase p . At O_{12} , p may be caused by B_1 or B_2 . Let us assume that B_1 fails p . Because p is not observed at O_{14} , B_4 must also fail p with an identical response. Similarly, had we assumed that B_2 fails p , we would have concluded that B_2 and B_3 have equivalent faults in phase p . Note that all four BUTs may be

faulty as well, but in this case the equivalent faults in the two pairs will be different. ■

Note that under the condition of Lemma 7, we cannot reach a unique diagnosis. However, we can move to a second phase of diagnosis where we divide the set of suspected PLBs and test each subset in separate BISTER tiles. The BIST results of each separate BISTER is then analyzed to achieve a unique diagnosis. If unique diagnosis cannot be obtained for a given subset of suspected PLBs, it can be further subdivided and retested. A subsequent example will show how this technique can uniquely diagnose the faulty PLBs.

Lemma 8: Assume that two out of three ORAs fail, such that the middle one has no failures, and the other two have disjoint failures. Then the two BUTs observed by the nonfailing ORA are fault-free and the other two BUTs are faulty.

Proof: In Fig. 8(c), without loss of generality, assume that O_{34} has no failures, while O_{14} and O_{23} have disjoint failures. Then we want to prove that B_3 and B_4 are fault-free and B_1 and B_2 are faulty.

Assume, by contradiction, that one of the two BUTs observed by the nonfailing ORA (say, B_3) fails phase p . Since O_{34} reports no failures, B_4 must have an equivalent fault in p . At most one of O_{14} and O_{23} may fail p , because their failures are disjoint. Then we have to analyze two cases.

- Case 1) One of the failing ORAs (say, O_{23}) shows p as a failure. Since p is not observed at O_{14} , while B_4 fails p , B_1 must also have an equivalent fault in p . In this case we would have three BUTs with equivalent faults in p , which would contradict assumption A2.
- Case 2) None of the failing ORAs reports p as a failure. But this would mean that both B_1 and B_2 must also have equivalent faults in p , since p does not appear at O_{14} or O_{23} . In this case we would have four BUTs with equivalent faults in p , which would contradict assumption A2.

Therefore both B_3 and B_4 are fault-free. Then B_2 is the source of the failures recorded at O_{23} , while those observed at O_{14} originate at B_1 . ■

As a consistency check, we can verify that the failures recorded at the ORA between the faulty BUTs are the union of the disjoint sets of failures obtained at the other two failing ORAs (for the example above, $FO_{12} = FO_{14} \cup FO_{23}$).

The following examples deal with the results of only one combined test session. Under the assumption that a faulty PLB is detected only when configured as a BUT, the two sessions can be independently analyzed, but the their results must be cross-checked for compatibility.

Example 1: Fig. 9(a) shows the set of failing phases obtained at the four ORAs in one combined test session for a 4×2 BISTER tile. Because failures in test phase 1 are reported only at O_{23} and O_{34} , then B_3 is faulty and fails in phase 1 (by Theorem 2). In the same manner, B_4 is identified as failing phase 5, B_1 is identified as failing phase 9, and B_2 is identified as failing phase 7. Since failing phase 8 is reported at 3 ORAs, Theorem 2 does not apply and we cannot determine which BUTs fail phase 8 (should be at least two). Note that we have identified all four BUTs as faulty.




Fig. 9. Graphical representations for (a) example 1, (b) example 3, and (c) example 4.

If our goal is also to identify the exact set of failing phases for every faulty PLB (to be able to reuse its fault-free modes of operation as a PUB), we can divide the set of suspected PLBs and retest each subset to obtain additional information. Let us assume that we test B_3 in a separate BISTER with known fault-free PLBs and we find out that $FB_3 = \{1, 8\}$. Then from (2) we obtain $Feq_{34} = \{1, 8\} - \{1, 5\} = \{8\}$, and from (4) we determine that B_4 must also fail phase 8. At this point we still do not know whether B_1 or B_2 is causing O_{12} to report failure in phase 8, so we again divide and re-test B_1 in a separate BISTER. Let us assume that we get $FB_1 = \{9\}$. Then from (2) we obtain $Feq_{12} = \emptyset$, and from (4) we determine that B_4 must also fail phase 8. Now we know the exact sets of failing phases for every BUT. ■

Example 2: Assume that we have $FO_{23} = FO_{34} = \{3, 5\}$ and no failures at O_{12} and O_{14} . Based on Lemma 6, we conclude that B_1 , B_2 , and B_4 are fault free, and from Theorem 2 we determine that B_3 is failing phases 3 and 5. ■

This example has illustrated the common pattern of a single faulty PLB.

Example 3: Assume that we have $FO_{23} = \{2, 5\}$, $FO_{14} = \{3, 4\}$, and no failures at O_{34} [see Fig. 9(b)]. Because FO_{23} and FO_{14} are disjoint while O_{34} reports no failures, the conditions of Lemma 8 are satisfied, and we can conclude that B_3 and B_4 are fault free, and B_1 and B_2 are faulty. Moreover, from (1) we find that $FB_1 = \{2, 5\}$ and $FB_2 = \{3, 4\}$. As a consistency check, we can verify that $FO_{12} = \{2, 3, 4, 5\}$. ■

Example 4: Assume that we have $FO_{23} = \{4\}$, $FO_{12} = \{3\}$, $FO_{34} = \{3, 4\}$, and no failures at O_{14} [see Fig. 9(c)]. Since only two ORAs with a common BUT fail phase 4, from Theorem 2 we determine that B_3 is faulty and fails phase 4. Because phase 3 is a common failure for two ORAs without a common BUT, from Lemma 7 we know that at least one of the pairs $\{B_1, B_4\}$ or $\{B_2, B_3\}$ have identical failures in phase 3. This is not a unique diagnosis, so we select one of the four suspects, say B_4 , to be tested in a separate BISTER, and let us assume that we find that B_4 is fault free. From Lemma 7 it follows that B_1 is also fault free, and then B_2 and B_3 must be the faulty PLBs with identical responses in phase 3. ■

Note how, by dividing suspect faulty PLBs into separate BISTER tiles, retesting, and analyzing BIST results, we can uniquely diagnose three faulty PLBs that cannot be diagnosed by analyzing the failing BIST results alone.

In practice, we have extended the *multiple faulty cell locator* (MULTICELLO) algorithm, originally developed for offline testing [29], to be used for identification of faulty PLBs in a 4×2 BISTER. Fig. 10 illustrates the algorithm analyzing the

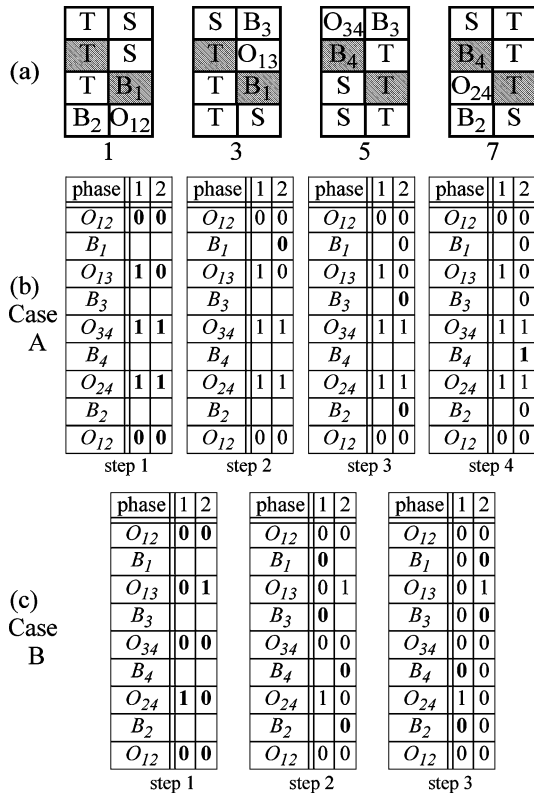


Fig. 11. A 4 × 2 BISTER tile with two faulty PLBs and MULTICELLO diagnosis for two cases where Assumption A1 does not hold including: (a) rotations of interest; (b) case A; and (c) case B.




Fig. 12. Basic PLB architecture.

MULTICELLO cannot make any entry in the table beyond step 1. For phase 2, MULTICELLO correctly identifies B₄ as faulty, but the fault in B₄ does not explain the failures at O₁₃ in phase 1. Therefore, in step 6 we divide the suspect PLBs into subsets and retest each subset. In at most two such steps, the two faulty PLBs will be separated and we will achieve unique diagnosis.

Case B: When either of the two faulty PLBs is part of the TPG, the TPG skips the patterns that detect the other faulty BUT. However, the faulty BUTs will be detected in different phases. As shown in Case B in Fig. 11, MULTICELLO completely fills the diagnosis table in step 3, but in step 5 it encounters an inconsistency, because all BUTs are identified as fault

TABLE II
BIST PHASES FOR ORCA 2C AND 2CA SERIES LOGIC BLOCKS

Phase	Flip-Flop/Latch Modes and Options					LUT Mode	Number Outputs
	FF/Latch	Set/Reset	Clock	Clock Enable	FF Data Input		
1	-	-	-	-	-	async. RAM	4
2	-	-	-	-	-	adder/subtractor	5
3	-	-	-	-	-	5-variable MUX	4
4	-	-	-	-	-	5-variable XOR	4
5	FF	async. reset	falling edge	active low	LUT output	count up	5
6	FF	async. set	falling edge	enabled	PLB input	count up/down	5
7	Latch	sync. set	active low	active high	LUT output	count down	5
8	FF	sync. reset	rising edge	active low	PLB input	4-variable	4
9	Latch	-	active high	active low	dynamic select	4-variable	4
10	-	-	-	-	-	multiplier	5
11	-	-	-	-	-	greater/equal comparator	5
12	-	-	-	-	-	not equal comparator	5
13	-	-	-	-	-	sync. RAM	4
14	-	-	-	-	-	dual-port RAM	4

free. Therefore we apply the divide and retest procedure, which will eventually will separate the two faulty PLBs.

V. DIAGNOSIS WITHIN A FAULTY PLB

During each test session, TREC records all ORA failures. This allows us to identify the failing mode(s) of operation of the faulty PLB and its faulty internal module(s). For example, ORCA 2C, Virtex, and Spartan II series PLBs contain four 4-input LUTs and four flip-flops with an output multiplexer network that can select any LUT or flip-flop output. The LUTs can also function as RAMs or fast adders. The flip-flops can also function as level-sensitive latches and have programmable set/reset and clock-enable features. The ORCA 2CA series FPGA PLB has the same architecture as the 2C series but has additional modes of operation including comparator, multiplier, and dual-port RAM [25]. Lines 1–9 in Table II summarize the test phases developed for the ORCA 2C series FPGA, while lines 10–14 describe the additional modes of operation tested for ORCA 2CA series. Phases 1–4 and 10–15 test the operation of the LUT/RAM module of the PLB, while phases 5–9 test the flip-flops and their modes of operation. Let us assume that a PLB fails only phases 5, 6, and 7. Since this faulty PLB passed the exhaustive tests for all its other modes of operation, it may be safely used for any function that does not involve counting (since phases 8 and 9 pass, the flip-flops work correctly as registers). This is the concept of a *partially usable block* (PUB) [22], [27], where such a defective block is reused whenever possible in its fault-free modes of operation, to increase the effective spare capacity of the FPGA in fault-tolerant applications. If phases 5–9 are the only ones that fail, the LUT module can still be used as a RAM or to perform any combinational logic function.

When the BIST results do not clearly indicate which component(s) of a faulty PLB is at fault, additional test phases are

TABLE III
ADDITIONAL DIAGNOSTIC PHASES FOR ORCA 2C AND 2CA

Phase	Flip-Flop/Latch Modes & Options					Logic Tested			
	FF/ Latch	Set/ Reset	Clock	Clock Enable	FF Data Input	LUT Mode	LUT	FF	MUX
1-4	Latch	-	active high	enabled	PLB input	-			√
5-8	-	-	-	-	-	4-vari- able	√		√
9	FF	async. set	falling edge	active low	PLB input	-		√	
10	FF	async. reset	rising edge	active high	PLB input	-		√	
11	FF	sync. set	rising edge	active low	PLB input	-		√	
12	FF	sync. reset	falling edge	active high	PLB input	-		√	
13	Latch	async. set	active low	active low	PLB input	-		√	
14	Latch	async. reset	high	high	PLB input	-		√	
15	Latch	sync. set	active high	active low	PLB input	-		√	
16	Latch	sync. reset	active low	active high	PLB input	-		√	

added to further isolate and identify the fault(s) and obtain even higher diagnostic resolution. Table III presents a list of additional diagnostic phases for the ORCA 2CA FPGA. Phases 1 through 8 test the output multiplexer, ensuring that every LUT or flip-flop output can reach each output of the multiplexer. Phases 9–12 isolate and test the flip-flop modes of operation, while phases 13–16 exercise the LATCH operation.

Diagnosis within a PLB to determine which LUT(s) or flip-flop(s) are faulty requires additional BIST configurations. Once the normal BIST and diagnostic configurations for PLBs have been executed and the faulty PLBs are identified, special PUB diagnostic configurations are applied for the two or more rotations that will put the PLBs identified as faulty under test. A set of eight PUB diagnostic configurations are used to test the LUTs, while another set of four configurations tests the flip-flops. Four of the eight LUT diagnostic configurations program the LUT with a four-input exclusive-OR function while the other four program the LUT with a four-input exclusive-NOR function in order to test all LUT bits for stuck-at-0 and stuck-at-1. During each of the four LUT configurations for PUB diagnosis, the outputs of the LUTs are rotated through the four outputs of the multiplexer. In this way, if the error follows the rotation, the fault is known to be in a LUT and the faulty LUT is identified as illustrated in the example of a faulty LUT in Fig. 13. On the other hand, if the error does not rotate, then the output multiplexer is faulty and the faulty output is identified. However, a combination of faulty LUT and multiplexer output requires additional diagnostic configurations. The same rotation strategy is used for the flip-flop PUB diagnostic configurations as can be seen in Fig. 13 by substituting flip-flops for LUTs.

Any single faulty LUT or flip-flop can be identified and any combination of a single faulty LUT and single faulty flip-flop can be identified with the current ORA design in the ORCA 2C FPGA. In Fig. 3, the ORA flip-flop stores the combined result of comparing four pairs of BUT outputs which limits the PUB diagnostic resolution for this ORA design multiple faulty LUTs




Fig. 13. PUB diagnosis.




Fig. 14. ORA with greater diagnostic resolution.

and flip-flops are encountered. Additional diagnostic resolution can be obtained by comparing each pair of outputs separately. In this way we can determine any combination of faulty LUT's and flip-flops, but at the expense of 24 PUB diagnostic configurations instead of 12 configurations. With the PLB architecture of more recent ORCA PLBs (like the 4C series FPGA), any combination of multiple faulty LUTs and flip-flops in the PLB can be identified. Fig. 14 illustrates this ORA implementation, which is also feasible in the Xilinx Spartan and Virtex series FPGAs [33].

Once the PUB diagnosis has completed for a given BISTER tile, the diagnostic procedure moves on to the next BISTER tile identified as having faulty PLBs until all PLBs that have been identified as faulty within a STAR have undergone PUB diagnosis. While this is the procedure in the current implementation, PUB diagnosis of multiple faulty BISTER tiles can be performed in parallel at the expense of slightly more complicated diagnostic software in TREC. When PUB diagnosis has been completed within a STAR, the identified faulty logic resources are processed by the fault-tolerance software to determine whether a faulty PLB can be utilized as a PUB by the next system function to occupy that position.

VI. EXPERIMENTAL RESULTS

The diagnosis approach presented here was applied to three known-faulty ORCA 2C15A FPGAs [25] (originally from Lucent Technologies) that had been previously tested with the off-line approaches presented in [15] and [16]. All three FPGAs failed the online tests introduced in this paper. Applying the logic diagnosis techniques described in the previous sections produced inconsistent results for two FPGAs (Chip 1 and Chip 2). This was the expected result, because the faults in these devices are interconnect faults [16]. The logic fault located in the third faulty FPGA (Chip 3) matched the faulty PLB identified in [15].

Next, we present the details of the online diagnosis of Chip 3. Fig. 15 shows the results of the test phases for both combined test sessions, where the failing test phases refer to Table II. In the first session, the same failures appear at two ORAs observing a common BUT, hence, this BUT is faulty. In the second




Fig. 15. Diagnosis of an actual faulty PLB for: (a) test session 1 and (b) test session 2.

session, all test phases fail with failure indications at only one ORA which, at face value, results in an inconsistency in step 5 of the MULTICELLO diagnostic procedure since it contradicts Lemma 4. However, this ORA is being implemented by the faulty BUT identified in the first session, where the failing test phases indicate one or more faulty flip-flops in the BUT; then the failures in the second session are consistent with a single faulty PLB with faulty flip-flops. Application of the additional diagnosis configurations discussed in Section V and shown in Table III confirmed that the fault only affects the four PLB flip-flops, leaving the LUTs unaffected. This defective PLB was successfully reused as a PUB to implement combinational logic functions in a fault-tolerant application [27].

Total roving and testing time in a fault-free ORCA 2C15 FPGA is approximately 1.34 s when the ORCA 2C Boundary Scan interface is operated at its maximum specified clock rate of 10 MHz. As a result, in the worst case a fault could escape detection for almost 1.34 s if it were to occur in a STAR position that has just been tested. Since diagnosis of faulty PLBs within a BISTER tile can be performed based on the failing BIST results, the 1.34 s roving and testing time also includes diagnosis to a faulty PLB. On the other hand, PUB diagnosis requires a maximum of 7.2 ms in the ORCA 2C15 to be able to identify a faulty LUT and/or flip-flop within the faulty PLB.

VII. CONCLUSION

In this paper, we presented BIST and BIST-based diagnostic methods for the roving STARS approach for online FPGA testing and diagnosis. The first step is to use the BIST approach to test all PLBs in the BISTER tile. Next, an analysis of the BIST results facilitates locating most common faults—one or two faulty PLBs—without any reconfiguration. If a unique diagnosis cannot be obtained, the suspected faulty PLBs are divided into subsets and retested. This combined technique achieves unique diagnosis with reasonably fast diagnosis time. The main result is the development of an online BIST and diagnostic method able to accurately diagnose any single faulty PLB and most combinations of multiple faulty PLBs within a BISTER tile. The combinations of multiple faulty PLBs that cannot be diagnosed appear to be very restrictive situations unlikely to occur in practice. We have applied our BIST-based diagnosis to actual faulty ORCA FPGAs and we correctly identified their known faulty PLBs. Our diagnostic technique provides the foundation necessary for implementing low cost and efficient fault-tolerant approaches in FPGAs [23], [27].

ACKNOWLEDGMENT

The authors wish to acknowledge S. Wijesuriya, C. Hamilton, and B. Skaggs from the Department of Electrical Engineering, University of Kentucky, Lexington, as well as T. Slaughter from the Department of Electrical and Computer Engineering, University of North Carolina at Charlotte for their contributions to BIST and BIST-based diagnosis of FPGA logic blocks during the Roving STARS project.

REFERENCES

- [1] W. Huang and F. Lombardi, "An approach to testing programmable/configurable field programmable gate arrays," in *Proc. IEEE VLSI Test Symp.*, 1996, pp. 450–455.
- [2] W. Huang, F. Meyer, X. Chen, and F. Lombardi, "Testing configurable LUT-based FPGAs," *IEEE Trans. VLSI Syst.*, vol. 6, pp. 276–283, Apr. 1998.
- [3] T. Inoue, S. Miyazaki, and H. Fujiwara, "Universal fault diagnosis for lookup table FPGAs," *IEEE Des. Test Comput.*, vol. 15, pp. 39–44, Jan. 1998.
- [4] C. Jordan and W. P. Marnane, "Incoming inspection of FPGAs," in *Proc. European Test Conf.*, 1993, pp. 371–377.
- [5] F. Lombardi, D. Ashen, X. Chen, and W. Huang, "Diagnosing programmable interconnect systems for FPGAs," in *Proc. ACM Int. Symp. FPGAs*, 1996, pp. 100–106.
- [6] M. Renovell, J. Figueras, and Y. Zorian, "Test of RAM-based FPGA: methodology and application to interconnects," in *Proc. IEEE VLSI Test Symp.*, 1997, pp. 230–237.
- [7] M. Renovell, J. Portal, J. Figueras, and Y. Zorian, "Testing the interconnect of RAM-based FPGAs," *IEEE Des. Test Comput.*, vol. 15, pp. 45–50, Jan. 1998.
- [8] —, "SRAM-based FPGA: Testing the LUT/RAM modules," in *Proc. IEEE Int. Test Conf.*, 1998, pp. 1102–1111.
- [9] —, "SRAM-based FPGA: testing the embedded RAM modules," *J. Electron Testing: Theory Applicat.*, vol. 14, no. 1/2, pp. 159–167, 1999.
- [10] B. Culbertson, R. Amerson, R. Carter, P. Kuekes, and G. Snider, "Defect tolerance on the Teramac custom computer," in *Proc. IEEE Symp. Field-Programmable Custom Computing Machines*, 1997, pp. 116–123.
- [11] C. Hamilton, G. Gibson, S. Wijesuriya, and C. Stroud, "Enhanced BIST-based diagnosis of FPGAs via boundary scan access," in *Proc. IEEE VLSI Test Symp.*, 1999, pp. 413–418.
- [12] C. Stroud, P. Chen, S. Konala, and M. Abramovici, "Evaluation of FPGA resources for built-in self-test of programmable logic blocks," in *Proc. ACM Int. Symp. FPGAs*, 1996, pp. 107–113.
- [13] C. Stroud, S. Konala, P. Chen, and M. Abramovici, "Built-in self-test for programmable logic blocks in FPGAs," in *Proc. IEEE VLSI Test Symp.*, 1996, pp. 387–392.
- [14] C. Stroud, E. Lee, S. Konala, and M. Abramovici, "Using ILA testing for BIST in FPGAs," in *Proc. IEEE Int. Test Conf.*, 1996, pp. 68–75.
- [15] C. Stroud, E. Lee, and M. Abramovici, "BIST-based diagnostics for FPGA logic blocks," in *Proc. IEEE Int. Test Conf.*, 1997, pp. 539–547.
- [16] C. Stroud, S. Wijesuriya, C. Hamilton, and M. Abramovici, "Built-in self-test of FPGA interconnect," in *Proc. IEEE Int. Test Conf.*, 1998, pp. 404–411.
- [17] S.-J. Wang and T.-M. Tsai, "Test and diagnosis of faulty logic blocks in FPGAs," in *Proc. IEEE Int. Conf. Computer-Aided Design*, 1997, pp. 722–727.
- [18] A. Burress and P. Lala, "Online testable logic design for FPGA implementation," in *Proc. IEEE Int. Test Conf.*, 1997, pp. 471–478.
- [19] C. Zeng, N. Saxena, and E. McCluskey, "Finite state machine synthesis with concurrent error detection," in *Proc. IEEE Int. Test Conf.*, 1999, pp. 672–679.
- [20] N. Shnidman, W. Mangione-Smith, and M. Potkonjak, "Online fault detection for bus-based field programmable gate arrays," *IEEE Trans. VLSI Syst.*, vol. 6, pp. 656–666, Dec. 1998.
- [21] A. Steininger and C. Scherrer, "On the necessity of online BIST in safety critical applications," in *Proc. Fault-Tolerant Computing Symp.*, 1999, pp. 208–215.
- [22] M. Abramovici, C. Stroud, S. Wijesuriya, C. Hamilton, and V. Verma, "Using roving STARS for online testing and diagnosis of FPGAs in fault-tolerant applications," in *Proc. Int. Test Conf.*, 1999, pp. 973–982.
- [23] M. Abramovici, J. Emmert, and C. Stroud, "Roving STARS: An integrated approach to online testing, diagnosis, and fault tolerance for FPGAs in adaptive computing systems," in *Proc. NASA/DoD Workshop on Evolvable Hardware*, 2001, pp. 73–92.

- [24] M. Abramovici and C. Stroud, "BIST-based delay-fault testing in FPGAs," in *Proc. IEEE Int. OnLine Test Workshop*, 2002, pp. 131–134.
- [25] Lattice Semiconductor Available. [Online] <http://www.latticesemi.com/products>
- [26] D. Siewiorek and R. Swarz, *The Theory and Practice of Reliable System Design*. Bedford, MA: Digital Press, 1982.
- [27] J. Emmert, C. Stroud, B. Skaggs, and M. Abramovici, "Dynamic fault tolerance in FPGAs via partial reconfiguration," in *Proc. IEEE Symp. Field-Programmable Custom Computing Machines*, 2000, pp. 165–174.
- [28] C. Stroud, M. Lashinsky, J. Nall, J. Emmert, and M. Abramovici, "On-line BIST and diagnosis of FPGA interconnect using roving STARS," in *Proc. IEEE Int. OnLine Test Workshop*, 2001, pp. 31–39.
- [29] M. Abramovici and C. Stroud, "BIST-based test and diagnosis of FPGA logic blocks," *IEEE Trans. VLSI Syst.*, vol. 9, pp. 159–172, Apr. 2001.
- [30] C. Stroud, *A Designer's Guide to Built-In Self-Test*. Norwell, MA: Kluwer, 2002.
- [31] *Standard Test Access Port and Boundary-Scan Architecture*, IEEE Standard P1149.1-1990, 1990.
- [32] M. Abramovici, M. A. Breuer, and A. D. Friedman, *Digital Systems Testing and Testable Design*. Piscataway, NJ: IEEE Press, 1994.
- [33] Xilinx, Inc. [Online]. Available: <http://www.xilinx.com/products>

Miron Abramovici (S'76–M'80–SM'86–F'93) received the Ph.D. degree from the University of Southern California, Los Angeles, in 1980.

He is co-founder and CTO of Design Automation for Flexible Chip Architectures (DAFCA), an electronic design automation (EDA) startup company providing a reconfigurable infrastructure platform for SoCs. He was Adjunct Professor of Computer Engineering at the Illinois Institute of Technology, Chicago. He was Principal Investigator and Project Leader of a three year Defense Advanced Research Projects Agency (DARPA) sponsored project on adaptive computing systems. He was also a distinguished member of the Technical Staff at Bell Labs, Murray Hill, NJ. He co-authored *Digital Systems Testing and Testable Design* (Piscataway, NJ: IEEE Press, 1994), and he has over 70 publications and 19 issued U.S. patents.

Dr. Abramovici has been an editorial board member for *IEEE Design and Test of Computers* and *Journal of Electronic Testing: Theory and Applications*.

Charles E. Stroud (S'74–M'88–SM'90) received the Ph.D. degree from the University of Illinois at Chicago in 1991.

Currently, he is a Professor in the Department of Electrical and Computer Engineering, Auburn University, Auburn, AL. Previously, he was a Distinguished Member of Technical Staff at Bell Labs, Naperville, IL, where he worked for 15 years as a VLSI and printed circuit board designer, with additional work in CAD tool development and BIST for digital and mixed-signal VLSI. He is author of *A Designer's Guide to Built-In Self-Test* (Norwell, MA: Kluwer, 2002). He has over 100 publications and 13 issued U.S. patents for various BIST techniques for VLSI and field-programmable gate arrays (FPGAs).

Dr. Stroud has been an editorial board member of the *IEEE Design and Test of Computers* and IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS.

John M. Emmert (S'92–M'93–SM'04) received the Ph.D. degree from University of Cincinnati, Cincinnati, OH, in 1999.

Currently, he is Director of Computer Engineering, and an Associate Professor in the Department of Electrical and Computer Engineering at the University of North Carolina, (UNC) Charlotte. He is also an officer in the United States Air Force. While on Active Duty in the Air Force he flew Teledyne Ryan's AQM-34L Unmanned Air Vehicles (UAVs); developed, implemented, and successfully tested a digital control system for flying multiple UAVs on a single channel; and developed computer-aided design tools for mapping circuits onto ASIC and field-programmable gate array (FPGA) platforms. Currently, he performs research in the areas of CAD tool development and mixed signal BIST for high speed, electronic warfare applications. He is also an active member of the U.S. Air Force Reserve.