

Testing NASA's 3D-Stack MCM Space Flight Computer

KOPPOL SASIDHAR

Georgia Institute of
Technology

LEON ALKALAI

Jet Propulsion Laboratory

ABHIJIT CHATTERJEE

Georgia Institute of
Technology

Advanced packaging technologies pose major testing challenges for complex designs. The authors present test strategies used in the space flight computer designed for NASA's Deep Space-1 mission.

ADVANCED PACKAGING TECHNOLOGIES such as 3D chip stacking, multichip modules (MCMs), and 3D stacks of MCMs provide opportunities for significant reductions in system mass, volume, and power. NASA's New Millenium Program (NMP) is using these technologies to develop a highly integrated, 3D microavionics architecture for future space science applications. The flight computer under development for the NMP Deep Space-1 mission consists of a 3D MCM stack that includes a CPU, large local memory, nonvolatile memory, and a high-bandwidth I/O interface.^{1,2} This set of modules will weigh about 1.5 kg, representing more than an order-of-magnitude reduction in mass relative to current spacecrafts. The reduction in mass will help make frequent, low-cost, miniature scientific missions to the outer planets possible.

Figure 1 shows a block diagram of the space flight computer. The processor module (one slice of the stack), designed by Lockheed Martin Federal Systems, consists of a RAD6000-5L processor, a local input/output (LIO) chip, oscillators, and start-up ROM.² The second module contains 400 Mbytes of local memory implemented with five 3D memory cubes of 16-Mbit DRAM chips stacked 40 high. The nonvolatile memory module, a solid-state recorder (SSR) designed by TRW, contains 1 Gbit of flash memory, using 32-Mbit plastic-package parts stacked four high. Finally, the I/O module,

designed by Boeing, includes the following: a PCI (peripheral component interface) bus for intermodule communication, a fiber optic I/O (AS1773) bus, and a PCI-bus-to-VMEbus bridge for communication with the rest of the spacecraft.

The four modules are packaged together in a low-cost, flexible MCM stacking technology developed by Space Computer Corporation and integrated at NASA Jet Propulsion Laboratory. The MCMs are mounted on printed circuit boards stacked in three dimensions with elastomeric interconnections. Figure 2 depicts the 3D MCM stack approach.

The 3D MCM space flight computer must undergo exhaustive testing and diagnosis before integration into the spacecraft. JPL is using methodologies based on the IEEE 1149.1 boundary-scan architecture for electrical test of the 3D stack. Special test chips perform reliable test and diagnosis of the computer architecture. The test hardware has been fabricated and is currently being tested. Before we describe the space flight computer's test methodology in detail, let's look at the general problem of MCM testing and available test strategies.

MCM testing

Computer designers widely use the IEEE 1149.1 boundary-scan architecture³ to facilitate MCM testing. This architecture associates a memory cell with each I/O pin of a

chip. These memory cells, or boundary-scan (BS) cells, connect in series to form the BS shift register. The test standard contains a four-port connection, the test access port. The TAP consists of four lines: test clock (TCK), test mode select (TMS), test data in (TDI), and test data out (TDO). The TDI line carries test instructions and data. The TDO line transfers test results to the external tester. The TCK signal carries the test clock. State transitions on the TMS line, which are decoded by the TAP controller, define the state of the chip's test circuitry. Figure 3 shows the general test architecture of a chip supporting the boundary-scan standard.

Researchers have proposed several techniques for testing MCMs. Zorian's structured-testability approach⁴ is well accepted in industry. It is based on die-level built-in self-test (BIST) and module-level boundary scan (Figure 4, next page). BIST is a design-for-testability technique that allows a circuit to test itself, typically to detect stuck-at faults in the die's application logic.

For testing, the BS cells of all dies in an MCM are connected into one scan path. The TDO of one die is tied to the TDI of another die, except for the initial TDI and last TDO ports, which are tied to the MCM's distinct terminals (Figure 4). Using this configuration, engineers can carry out interconnection testing, snapshot observation of normal system data, and chip testing. To enable chip testing, they send special test instructions to a die through the boundary-scan path to activate its BIST hardware. They can then observe the results of the BIST process through the boundary-scan signals to determine whether the die is "good" or "bad."

In this article, we assume that the ICs in each stack slice are compatible with the IEEE 1149.1 boundary-scan standard (which we refer to as JTAG for convenience). We also assume that these ICs are equipped with BIST mecha-




Figure 1. Block diagram of the 3D microavionics stack.




Figure 2. A 3D MCM stack.




Figure 3. Chip supporting IEEE 1149.1 boundary-scan standard.

nisms that can be activated under JTAG control. Our goals are to test all IC interconnections, within each slice as well




Figure 4. Boundary scan in an MCM.

as across slices, and to activate the BIST mechanisms of each IC under JTAG control. We can thus validate each IC's functionality as well as all the IC interconnections, providing comprehensive test of the stack electronics.

Motivating the test hardware design is the need to test the stack electronics reliably in the presence of faults in the boundary-scan circuitry, such as an open TDI line. Of particular concern are faults that can render the entire 3D stack untestable. Hence, we must evaluate the proposed test methodologies in the context of test reliability—that is, the ability to test the 3D stack in the presence of point failures in the test hardware itself. We also want the test architecture to

be scalable—that is, to enable the addition of stack slices without a change in test methodology. Further, the time necessary for testing an embedded IC must depend only on the test sequence length for that IC, not on the number of ICs in the stack or the test complexities of other embedded ICs. Because the manufacture of NASA's 3D MCM stack involves different industrial partners, an open and pluggable architecture that allows seamless testing across different slice technologies is also necessary. Finally, a key design goal is that we use the same approach to test the 3D stack in flight as on the ground.

Specialized BS bus configurations

In this section, we present three simple strategies for testing a 3D stack of MCMs. A test bus configuration and a number of test lines characterize each strategy. We assume that no special test chips are included in the MCM slices.

Serial-module scan test. This strategy is a direct extension of the IEEE Std 1149.1 approach for testing multiple chips on a board. As shown in Figure 5, we connect the BS registers in the ICs of the MCM slices in a single serial-scan chain running through the 3D stack. We distribute the test data to all the MCM slices through module 1's TDI input, which is connected to the external tester (each module in the figure corresponds to a slice). We scan the test responses from all the slices into the external tester through module 4's TDO line. The TMS and TCK lines feed in parallel to all the slices.

This test configuration has several drawbacks:

- The test is not very reliable because a fault in the BS chain between two ICs can prevent the entire stack from being tested.
- The time needed to test an MCM slice individually depends on the complexity of other slices in the 3D stack.




Figure 5. Serial-module scan test configuration. TRST: test reset.




Figure 6. Multiple-module control test configuration.

During the test of a single slice, the BS hardware in other slices is in bypass mode. A chip is in bypass mode when a single flip-flop is connected between the TDI and TDO lines of its BS hardware. Placing all its ICs in bypass mode places an entire slice of the 3D stack in bypass mode. In this mode, the number of flip-flops connected between the slice's TDI and TDO lines equals the number of ICs connected in the slice's BS chain. In Figure 5, for example, when module 1 (slice 1) is tested, seven flip-flops are connected in the serial chain between module 1's TDO line and module 4's TDO line. This adds a delay of seven clocks for every test vector applied to slice 1. The delay is substantial when the number of test vectors applied and the number of ICs in other slices are large.

- One cannot test a slice without knowing the test architecture of the other slices (the number of chips to be placed in bypass mode and the delay incurred in scanning test data must be known).
- During test of a single MCM slice, other slices in the stack must be powered up to enable all of them to be in bypass mode.

The serial-module scan technique has two advantages. It requires only four test lines to test the complete 3D stack. Further, it requires no extra test hardware at the MCM and 3D-stack levels, other than the hardware provided by the JTAG standard. Overall, however, the inability of the test process to withstand point failures in the BS bus and test hardware is of serious concern. In contrast, the following BS test bus configurations allow reliable testing in the presence of point failures.

Multiple-module control test. In this approach, each

slice has a separate TMS line associated with it, as shown in Figure 6. The TDI and TDO lines from each slice connect directly to the external tester. Overall, seven test lines are required: TCK, TDI, TDO, and four TMS lines (one connected to each slice).

Multiple-module control test offers these advantages:

- It restricts the continuous BS register chains to lie within a single slice, unlike serial-module test, in which the chain spans multiple slices. Hence, test reliability improves since most faults in a slice do not prevent other slices from being tested.
- It tests each MCM slice independently. A slice's test data is not transmitted through other slices. We scan test data into a slice's TDI line by placing the data on the stack's TDI port. Further, we scan out results from the TDO line to the external tester directly. Hence, this technique ensures that at any instant, test data is scanned out from only one MCM slice into the external tester.
- An MCM slice can be tested in isolation without powering up other slices in the 3D stack. We do this by activating the TMS line connected to that slice while holding other TMS lines idle.
- The structured test of an MCM involves module-level interconnect test and IC-level BIST. Using multiple TMS lines enables BIST and interconnect test in different slices simultaneously. For example, while subjecting slice 1 to interconnect test, we can initiate BIST in the other slices. By suitably configuring operations, we can minimize test time.

Thus, this technique is more robust from a test perspective than the serial-chain scan approach. However, faults on the




Figure 7. Parallel-module test configuration.




Figure 8. Test features included in an MTC.

TCK line will prevent the entire stack from being tested. Also, specific stuck-at faults on the TDI and TDO lines could be critical. Hence, we next present a strategy that provides separate test buses to each MCM slice in the 3D stack.

Parallel-module test. In this approach, each slice has a separate TMS, TDI, TCK, and TDO line associated with it (Figure 7). All the slices can be tested independently in parallel. This approach is more reliable than the previous approaches because faults on the TCK line of a slice do not

prevent other slices from being tested. Also, slices can be tested at different clock frequencies, unless the test procedure relies on synchronization of circuitry in different slices or the entire stack assembly is being tested for reliability. This approach provides maximum test application parallelism and allows test results for each slice to be analyzed independently. However, it requires the maximum number (12) of test lines, four for each slice.

All the test strategies just presented involve varying the number of test lines to increase reliability, speed, and parallelism. The number of test lines varies from a minimum of four to a maximum of 16. No extra hardware is

used at the MCM and 3D-stack levels. The choice of one of these approaches for a particular test application depends on three factors: the number of test buses allowed by the external tester, the test parallelism allowed by the external tester's software, and the number of pins provided for test in each slice.

MTC-based strategies

Next we present test techniques based on inclusion of a test controller die in each slice of the 3D stack. An MCM test

controller (MTC) die incorporated in each slice controls its scan path during test. The MTC includes a TAP controller and an optional BS register through which test vectors are scanned in and applied to the non-BS dies.^{5,6} It has a single-bit bypass register that enables an MCM slice to be treated as a single JTAG device and placed in bypass mode when necessary. Figure 8 depicts the MTC test features.

Further, we can reconfigure the scan path itself with a scan path linker^{7,8} to control an MCM's scan path configuration (Figure 9). The scan path linker partitions the scan path into multiple secondary scan paths. By scanning in a control word, the scan path linker selects one or more of the secondary scan paths and includes it in the primary scan path. This increases the scan chain's reliability. For example, when the control word is 1010, chips 1 and 3 are connected in a chain. The TDI line is connected to TDO1, TDI1 is connected to TDO3, and TDI3 is connected to TDO. Faults in a particular scan chain section are easily isolated without affecting test operations in fault-free parts of the MCM.

The scan path linker's test hardware consists of approximately 75 gates.⁷ The proposed MTC architecture consists of 240 gates⁹ and includes a TAP controller and BS register as well as a scan path linker. Next we briefly describe two




Figure 9. Scan path linker in an MCM.

strategies that include an MTC in each slice and use the test bus configurations described in the previous section.

Serial-MTC test. Hilla⁸ developed this technique (Figure 10) from the serial-module scan test approach. This technique has an advantage over its parent scheme: When an MCM slice is placed in bypass mode, a 1-bit register is placed between the MCM's external TDI and TDO lines, irrespective of the number of chips connected in the scan path. A fault in the scan path between two slices will prevent test of the 3D stack. However, a fault in the scan path in a slice will not affect test of the other slices, since the MTC in the faulty slice is in bypass mode. Hence, this technique is more reliable than serial-module scan test.




Figure 10. Serial-MTC test configuration.




Figure 11. Parallel-MTC test configuration.




Figure 12. Master MTC test configuration.

Parallel-MTC test. Similarly, we can define techniques for incorporating MTCs into each slice for multiple-module control test and parallel-module test. The parallel-MTC test approach corresponds to the other end of the test spectrum, with the maximum number of test lines and maximum additional test circuitry (Figure 11). It provides very reliable testing.

Autonomous BIST

So far, we have identified techniques for offline testing of a 3D stack after assembly, using an external tester. For space applications, it is equally important to consider autonomous testing with BIST mechanisms that function in the absence of an external tester. To enable BIST, we must use a master test controller that manages the test process. The following two test strategies rely on the use of a master test controller for BIST of a 3D MCM stack.

Master MTC test. This test strategy, diagrammed in Figure 12, identifies the MTC in slice 1 of the 3D stack as the master MCM test controller (MMTC). The MMTC controls the

flow of test data to the 3D-stack slices. During assembly test, the MMTC connects to an external tester through a single JTAG bus and to the MTCs in other slices through separate JTAG buses. By sending a control word to the MMTC from the tester, we connect the JTAG bus from the external tester to one of the three buses that transfer test data to the slices. The test hardware in an MMTC consists of 310 gates.⁹ For autonomous test of the 3D stack, the MMTC contains an optional control and data buffer (CDB) that stores test data. This test data is loaded into the buffer after assembly of the flight computer module.

The assembly test process for the 3D MCM stack proceeds as follows. First, the external tester enables self-test of the MMTC. Next, the MMTC tests the JTAG buses connected to other slices and the other MTCs. Then, it tests the 3D-stack slices. The external tester initially sends a control word to the MMTC to choose a slice to test. This connects the JTAG bus from the external tester to the appropriate slice through the MMTC. Subsequently, the MMTC sends control words to the MTC associated with the slice to configure the scan chains in the slice. Test data is then scanned into the MCM




Figure 13. Reconfigurable MTC test configuration.

slice from the external tester.

This technique permits BIST of a 3D stack by the MMTC. For autonomous BIST, a test program running on the CPU performs the actions of the external tester. However, the 3D stack cannot be tested if the MMTC is faulty or if the JTAG bus between the external tester and MMTC fails. Also, if the JTAG bus connected between a slice and the MMTC fails, the slice cannot be tested. The next test technique addresses some of these problems.

Reconfigurable MTC test. The reconfigurable test technique enhances BIST reliability. It provides two test features in a 3D stack: redundant test lines and a reconfigurable test architecture. The test ICs in all slices, called reconfigurable MCM test controllers (RMTCs), are identical. We select one RMTC as the master test controller of the 3D stack. Redundant test lines enable testing in the event that some lines are faulty. All test lines are bidirectional and can be configured as input or output lines by the RMTCs.

The boundary-scan architecture requires a minimum of four test lines: TDI, TDO, TCK, and TMS. Our approach provides additional test lines connected to the RMTCs, as shown in Figure 13. The test bus is common to all slices in the stack. In the example shown, we have provided two extra redundant lines.

Each RMTC contains the basic circuitry included in an MMTC. It also contains circuitry to detect faulty test lines and configure the test bus appropriately. The RMTC test hardware consists of 460 gates.⁹ Initially, all the RMTCs perform self-test in parallel. Subsequently, they mutually verify the self-test response. Then, they select one of the fault-free

Table 1. Comparison of test strategies for a 3D stack of MCMs.

Strategy	Test lines	Hardware (gates)	Test time	Reliability
Serial-module scan	4	—	T	1
Multiple-module control	$4 + \log_2(N)$	—	$< T$	4
Parallel-module	$4N$	—	T/N	5
Serial-MTC	4	$N \times 240$	T	2
Parallel-MTC	$4N$	$N \times 240$	T/N	6
Master-MTC	4	$310 + (N-1) \times 240$	T	3
Reconfigurable-MTC	$4 + X$	$N \times 460$	T	7

RMTCs as the master. The master RMTC tests the JTAG bus and configures the test lines. Subsequently, the 3D-stack slices are tested by their respective RMTCs.

This test approach fails only when all RMTCs or all test lines in a 3D stack are faulty, which is highly improbable. Hence, this test approach is very reliable. Further, the technique is easily scalable to larger-size 3D stacks. Additional slices can be added and connected to the test bus without changes to the test architecture of the original slices.

Comparison of test strategies

We now compare the approaches with respect to test hardware, test time, test lines, and test reliability. We assume the number of MCM slices in a 3D stack to be N . We assume the test scan chains in all the slices are the same length. We denote the time taken to test a 3D stack using the serial-module approach as T and use it as a standard of comparison. X refers to the number of extra test lines required by the technique. Table 1 presents the results.

In the table, “test lines” refers to the number of lines that must be connected from an external tester to the associated 3D MCM stack for test application. “Test hardware” refers to all the additional circuitry (in test ICs) incorporated in every slice of the stack for enhanced test. Since the addi-




Figure 14. Test scheme for 3D flight computer module.




Figure 15. Redundant JTAG buses in the 3D MCM stack.

tional logic required for test in each slice is not very large, we can incorporate it in any IC in the MCM slice if we have the flexibility to do so. The table does not include the logic complexity of the TAP controller circuitry for the boundary-scan standard. The table ranks the strategies from 1 to 7 in test reliability, with 7 being the most reliable.

Serial-module scan is the least reliable approach. The MTC-based serial-module technique is more reliable but has some of the deficiencies of serial-module scan test. The master MTC test technique ranks 3 because faults in the MMTC are catastrophic and prevent test of the entire module. The multiple-module control and parallel-module strategies follow. The parallel-MTC approach is the second most reliable technique and requires the maximum number of test lines. Finally, the reconfigurable test approach is the most reliable and requires

the most complicated test hardware: the RMTCs.

The number of test lines needed by the reconfigurable approach is variable and depends on the availability of test pins on the 3D stack. The test time is minimal for parallel test strategies because the MCM slices in the 3D stack are tested in parallel. However, these approaches require the most test lines, four lines from the external tester to each MCM slice.

Microavionics stack test architecture

The test strategy implemented for NASA's 3D microavionics stack incorporates features of the multiple-module control test technique discussed earlier, with additional redundant test lines. NASA chose this approach over MTC-based test strategies due to practical difficulties in procuring

radiation-hardened RMTC IC implementations, a requirement for deep-space-mission electronics. We expect such radiation-hardened ICs to become available in the future.

Each MCM slice in the 3D stack has a separate TMS line associated with it. However, only one TMS line is connected to the microavionics stack from the external tester. A 4:1 demultiplexer connects the external TMS line to the TMS lines corresponding to the different slices. Figure 14 illustrates this test scheme.

To allow testing of the 3D stack in the presence of specific failures in the test lines, such as line opens, a pair of redundant buses runs vertically along opposite sides of the stack, feeding each MCM slice. The two buses are connected in wired-OR configuration at their destination terminals in each slice of the stack, as shown in Figure 15. For example, the TMS line connected to an IC in the I/O MCM slice is the wired-OR of the TMS lines of test buses JTAG1 and JTAG2 in Figure 15. Thus, if one TMS line is faulty, test control data can be sent over the other TMS line.

Figure 16 shows the test architecture of the processor MCM slice. The LIO chip contains a JTAG diagnostic interface with provisions for a JTAG master controller and all JTAG bus signals. The LIO JTAG bus is controlled by the external diagnostic interface (COP) and the RAD6000 CPU through JTAG interface registers. Test data is placed on the LIO JTAG bus by the RAD6000 CPU directly or by an external source through the COP interface. The slice contains an additional JTAG bus (external JTAG) that can be controlled directly from an external source. This enables easy test access to the modules even when the JTAG diagnostic interface in the LIO is faulty.

The external JTAG bus also enables direct observability of the test data being scanned by the RAD6000 CPU or the COP into the modules under test. We test the JTAG diagnostic interface in the LIO, the RAD6000 JTAG interface, and the COP interface by scanning in the test data from the COP or the processor into the external tester through the external JTAG bus. This allows diagnosis of faults in the processor's test hardware. The internal JTAG bus is configured as the wired-OR of the external JTAG bus and the LIO JTAG bus as shown in Figure 16. Hence, while one JTAG bus (LIO or external)




Figure 16. Test structures in the processor slice.

is scanning the test data into the slices, the other must be tristated. The LIO chip controls the demultiplexer in the processor slice.

Figures 17 through 19 (next page) show the test structures of the other three modules. The DRAM controller in the local memory slice is JTAG-compatible. Hence, we scan test data into the controller and apply it to the data and address buses of the memory stacks. A similar setup permits extensive test of the SSR module's memory banks and FPGA controllers. The I/O module has a boundary-scan chain through the PCI and VME controller ASICs. However, the DR 1773 ASIC is not JTAG-compatible. Figure 19 shows the scan chain. All the ASICs have BIST capability. The boundary-scan RunBIST instruction enables BIST of the ASICs.

Microavionics stack test procedure

The test procedure for NASA's 3D microavionics stack involves detailed test and diagnosis of the test architecture and MCM slices as well as all related interconnections.

Initially, we test the JTAG test architecture of all the MCM slices. This involves verifying the JTAG controllers and buses in all slices. Because the JTAG buses to all the slices are scanned out through the processor slice, it is essential that we test the JTAG diagnostic interface in the LIO first. This involves testing the external JTAG bus, the LIO JTAG bus, the COP interface, the RAD6000 CPU test interface, and the demultiplexer control. After the test architecture of the processor slice is validated, we check the test setup on the other slices by scanning in serial data through the associated JTAG buses. This validates the correct functioning of the JTAG test hardware.




Figure 17. Test structures in the local memory slice.




Figure 18. Test structures in the SSR slice.




Figure 19. Test structures in the I/O slice.

Subsequently, we test the interconnections between slices. We test the PCI bus passing through all the modules, and the memory data and address bus from the processor slice to the local memory slice.

After testing all the stack slice interconnections, we test the slices individually. This process starts with exhaustive test




Figure 20. Test flow diagram for the 3D microavionics stack.

of the processor MCM slice. The interconnections in the processor slice are tested first. Next, we test the RAD6000 CPU and the LIO IC by enabling BS and BIST test. Next, the local memory slice is tested, followed by the SSR and I/O slices.

To generalize the test procedure, we can describe it by the following recursive algorithm:

```

Test_Algorithm()
{
    1. Test the JTAG architecture of the 3D stack.
    2. Test(3D MCM stack).
}

Test(Component)
{
    Divide the component into subcomponents.
    Test interconnections between subcomponents.
    For all subcomponents do
        Test(subcomponent)
}

```

In the first iteration of the subroutine Test(), a component corresponds to the entire 3D MCM stack. First, the in-

terconnections between the MCM slices are tested. Then each slice is tested individually. When Test() is invoked again, a component corresponds to an individual MCM slice. The subcomponents in this case are ICs in the MCM slice. Figure 20 shows the test flow diagram. If a test fails at any step, the external tester diagnoses the fault and repair begins. Subsequently, the complete test routine repeats.

NASA has had the test structures fabricated and has used them to test the 3D microavionics stack after assembly. While the spacecraft is in orbit, the RAD6000 CPU can initiate autonomous test of the stack. Personnel on earth can observe the test data and results communicated through the external JTAG bus.

THE METHODS we have described greatly enhance the ability to test stacks in the presence of failures that can inhibit the operation of the test hardware itself and render the entire stack untestable. Such failures are of great concern in the fabrication of 3D MCM stacks (as opposed to MCMs themselves), due to their fabrication complexities, longer line lengths, and mechanical stresses. Achieving acceptable assembly yields and meeting product quality requirements through appropriate test methodologies are critical problems. We can alleviate both problems by adopting test approaches that guarantee the integrity and performance of the underlying packaging technologies. 

Acknowledgments

The research described here was performed in part at the Jet Propulsion Laboratory, California Institute of Technology. The work was sponsored by NASA's New Millennium Program.

References

1. L. Alkalai, J. Klein, and M. Underwood, "The New Millenium Program Microelectronics Systems Advanced Technology Development," *Proc. New Millenium Microspacecraft Technology AIAA 34th Aerospace Sciences Meeting & Exhibit*, AIAA Tech. Library, 1996.
2. L. Alkalai and B. Jarvis, "The Design and Implementation of NASA's Advanced Flight Computing Module," *Proc. IEEE Multi-Chip Module Conf.*, IEEE, Piscataway, N.J., 1995, pp. 40-44.
3. *IEEE Std 1149.1-1990 Test Access Port and Boundary-Scan Architecture*, IEEE, Piscataway, N.J., 1993.
4. Y. Zorian, "A Structured Testability Approach for Multi-Chip Modules Based on BIST and Boundary Scan," *IEEE Trans. Component, Packaging, and Manufacturing Technology, Part B*, Vol. 17, No. 3, Aug. 1994, pp. 283-290.
5. N. Jarwala, "Designing Dual Personality IEEE 1149.1-Compliant Multi-Chip Modules," *Proc. IEEE Int'l Test Conf.*, IEEE Computer Society Press, Los Alamitos, Calif., 1994, pp. 446-455.
6. W.D. Ballew and L.M. Streb, "Board Level Boundary Scan Regaining Observability with an Additional Chip," *Proc. IEEE Int'l Test Conf.*, IEEE CS Press, 1989, pp. 182-189.
7. "Scan Path Linkers with 4-Bit Identification Buses, SN54ACT8997 & SN74ACT8997," Texas Instruments, Inc., TI0286-D3597, 1990.
8. S.C. Hilla, "Boundary Scan Testing for Multi-Chip Modules," *Proc. IEEE Int'l Test Conf.*, 1992, pp. 224-230.
9. K. Sasidhar, *Parallel Test Techniques for Multi-Chip Modules*, PhD dissertation, Dept. of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, Ga., 1997.



Koppol Sasidhar received his BTech degree in computer science and engineering from the Indian Institute of Technology, Kharagpur. He received his MS and PhD degrees in electrical engineering from the Georgia Institute of Technology. His research interests are MCM testing, graph theory, and parallel algorithms.



Leon Alkalai is the center lead of the Center for Integrated Space Microsystems, a NASA Center of Excellence at the Jet Propulsion Laboratory. The center focuses on developing advanced microelectronics, microavionics, and advanced computing technologies for future deep-space missions. Alkalai joined JPL after receiving his PhD in computer science from UCLA.



Abhijit Chatterjee is an assistant professor in the School of Electrical and Computer Engineering, Georgia Institute of Technology. Previously, he was a research staff member at the General Electric Research and Development Center in Schenectady, N.Y. His research interests are VLSI and mixed-signal design and test, low-power circuit design, computer algorithms, and design automation. Chatterjee received the BTech from the Indian Institute of Technology, Kanpur, the MS in electrical engineering and computer science from the University of Illinois at Chicago, and the PhD in electrical and computer engineering from the University of Illinois at Urbana-Champaign. He is a collaborating partner in NASA's New Millennium Project. Chatterjee is also the author of a US patent and over 90 journal and conference publications. He is a senior member of the IEEE.

Send correspondence to Abhijit Chatterjee, Georgia Institute of Technology, School of Electrical and Computer Eng., 777 Atlantic Dr., Atlanta, GA 30332-0250; abhijit.chatterjee@ee.gatech.edu.