PART THREE

DECADE 90

The future of system design

Design for testability creates better products at lower cost

Chips, boards, and systems of the 1990s will be far more sophisticated than those of today. Engineers who adopt a design-for-test (DFT) philosophy will create easily testable, more-reliable products that cost less to manufacture and operate. Without DFT methods, the cost of testers, fixtures, and test programs will soar.

Steven H Leibson, *Regional Editor*

Many design engineers' and managers' attitudes about product testability seem to have frozen during an earlier era in electronics, when a technician could troubleshoot almost any problem in 20 minutes with a scope and a little savvy. For products designed without a design-for-testability (DFT) philosophy, today's most advanced in-circuit ATE testers can do no more than automate the time-honored tradition of sticking a test probe into a failing test node to find the problem. But as electronic systems grow in complexity, this approach grows less and less effective and increasingly costly.

Today's electronic systems are already making test probing impractical because surface-mount technology, VLSI, and other advanced electronic-packaging schemes hide thousands of circuit nodes from the probe tip. The majority of systems built in the 1990s will certainly present even more difficulty in testing. As electronic-device technology and system design advance, therefore, the traditional attitude toward testing is becoming more and more unrealistic.

In the past, many design engineers ruled out design for testability because they thought it would cost too much, take up too much space, and delay their projects. Project managers pressed by schedule and cost constraints still encourage these attitudes because, for the most part, testing provides feedback about manufacturing quality—information that's important to the manufacturing department but not necessarily to the design lab.

Today, as they have in the past, many design engineers use every available transistor for speed, capacity, or other performance features, leaving no room for testability circuitry. The designs they create, therefore, are often difficult or impossible to test. In the 1990s, however, successful designs will allocate some circuitry for testability, which leads to better product quality and a shorter development cycle.

Testing provides quality feedback

Testing and DFT methodologies are initially useful to validate a design once it's actually built, but they serve best to provide valuable feedback on the manufacturing process once a product reaches production. "There's no sense in building a product that you can't test and can't build reliably," says Paul Gifford, manager of central systems engineering at Sequent Computers (Beaver-

ton, OR). Sequent plans to incorporate scan design in its third-generation, multiprocessor computer systems to ensure that the products will be manufactured correctly. Gifford believes that the benefits of the DFT approach far outweigh the extra effort required to design testability into the computers. Further, he says, the performance penalty (if any) for adding DFT is only 5 to 6%.

Ignorance of DFT methods and the perception that testability is of secondary importance are the last major obstacles that a company must overcome before it can adopt a DFT orientation. When you consider a product's entire life cycle—including manufacturing, testing, and field service—you find that DFT methods actually save time and money in comparison with the traditional approach of "tossing it over the wall and letting the test engineer handle it."

Today's unknown testing costs

Sadly, many companies don't know what product testing really costs them. Ask a test manager what it costs to test a product and you'll often get a figure derived from the total number of products tested divided by the cost of the entire test operation. That figure gives you an average test cost for all the products run through the testing department, but it doesn't paint a very accurate picture of the true cost for any particular product. In addition, such an estimate of test-department costs generally doesn't include field-maintenance and repair costs.

Logical Solutions Technology Inc (a Campbell, CA, testability consulting firm) estimates that the average electronics manufacturer spends between 35 and 45% of a product's total cost on testing parts, subsystems, and final assemblies. The company says that its customers have saved, on the average, about $1.5 million per year by following its DFT recommendations. Note that these figures represent savings for systems built at today's complexities, not for the more complex systems that will be built in the 1990s. Test experts predict that test problems will be much worse for very complex systems unless design engineers add testable-design and DFT methods to their lexicon.

At the 1987 Government Microcircuit Applications Conference (held in Orlando, FL), Mitre Corp (Bedford, MA) reported on a study it performed for the Electronic Systems Div of the US Air Force to determine the impact built-in test (BIT) circuitry would have on equipment maintenance. The study indicated that BIT would improve both instantaneous and steady-state equipment availability, and that it would improve mission reliability by identifying weak modules before a critical failure occurred. From a field-service viewpoint, BIT reduces the occurrences of "cannot duplicate" (CND) and "retest OK" (RTOK) situations, because BIT circuitry pinpoints failing components.

Intermittent failures cease to be difficult to find and repair in systems that have BIT, because the BIT circuitry can store the identity of the faulty module. In addition, the incorporation of BIT circuitry reduces the mean time to repair (MTTR) by eliminating fruitless CND and RTOK maintenance actions; the BIT circuits immediately indicate the problem source, eliminating the troubleshooting phase of repair.

Mitre's report indicates that a BIT design with a 90% chance of isolating a problem incurs 10 to 30% extra design time during a product's development cycle. However, such incremental development costs add very little to the product's overall life-cycle costs **(Fig 1a)**. Considering the time and money required to develop and debug tests and test fixtures, as well as to perform field maintenance and repair, DFT methods ultimately save your company both time and money. According to the Mitre report, although the system-design phase of product development represents only about 15% of the product's total life-cycle cost, it has a 70% impact on that product's operation and support costs **(Fig 1b)**.

Testability also translates into product quality during production, because it lets you ship fewer products with undiscovered faults. In the field, products that are designed to be testable can be repaired more quickly, resulting in less down time for the customer. Though Mitre's report specifically applies to military systems, which have longer development and life cycles, the report's conclusions have equal validity for commercial product development, even if the numbers aren't exactly the same.

Three keys to testability

Because of the overwhelming evidence that DFT is simply part of a good overall design strategy, many companies are actively developing DFT methods. Though these methods differ, they all focus on the three keys to testability: partitioning (to break complex systems into testable blocks), control (to allow a test to stimulate testable blocks), and visibility (to extract the system's response to the test stimuli).

DFT methods take several approaches, which include

divide-and-conquer, several types of serial scanning, and built-in self-test (BIST) or built-in test (BIT) (see box, "DFT methods focus on scan-path testing.") Each of these methods recognizes that you can no longer test increasingly complex systems simply by increasing the number of test probes on an ATE tester. Such an approach has grown prohibitively expensive as system complexities soar. Instead, the current DFT methods focus on adding test circuitry to the product. This extra circuitry allows less-complex test equipment to perform simpler tests with better fault coverage.

A divide-and-conquer test scheme works well in systems that can be divided into easily testable blocks or in blocks that have existing tests. For example, RAM and ROM blocks are relatively simple to test, yet they consume large portions of a system's transistor budget, so testing them is worthwhile. Today's test methods can verify the operation of these structures quite easily when they're isolated from the rest of a system.

In a paper presented at the 1987 Custom Integrated Circuits Conference, National Semiconductor (Sunnyvale, CA) discussed techniques for isolating blocks of circuitry embedded in an IC. If such a block is based on an existing standard part, such as the 82C50 asynchronous communications element in National's paper, you can use an existing test to verify that block's operation by employing data multiplexers to bring the block's input and output signals to the chip's leads.

National Semiconductor's paper also compares parallel and serial methods of accessing such isolated blocks. Parallel-access methods allow faster testing, but require more points of contact between the tester and the system. Serial-access methods are slower, but don't require as many test points. Because serial test methods require fewer test probes and less-expensive test equipment, engineers are adopting such techniques more and more frequently.

When engineers at NCR's Microelectronics Div (Fort Collins, CO) developed the PLM (Prolog machine) microprocessor in conjunction with the Computer Science Div of the University of California at Berkeley, they knew that the complexity of the chip would make testing difficult unless they included some on-chip test circuitry. Designed to act as a coprocessor in an engineering workstation, the PLM implements a tagged architecture and five hardware stacks to support the Prolog language environment. The resulting IC, representing a system with a complexity of about 45,000 gates, incorporates eleven 32-bit data buses, sixty-four

Text continued on pg 140

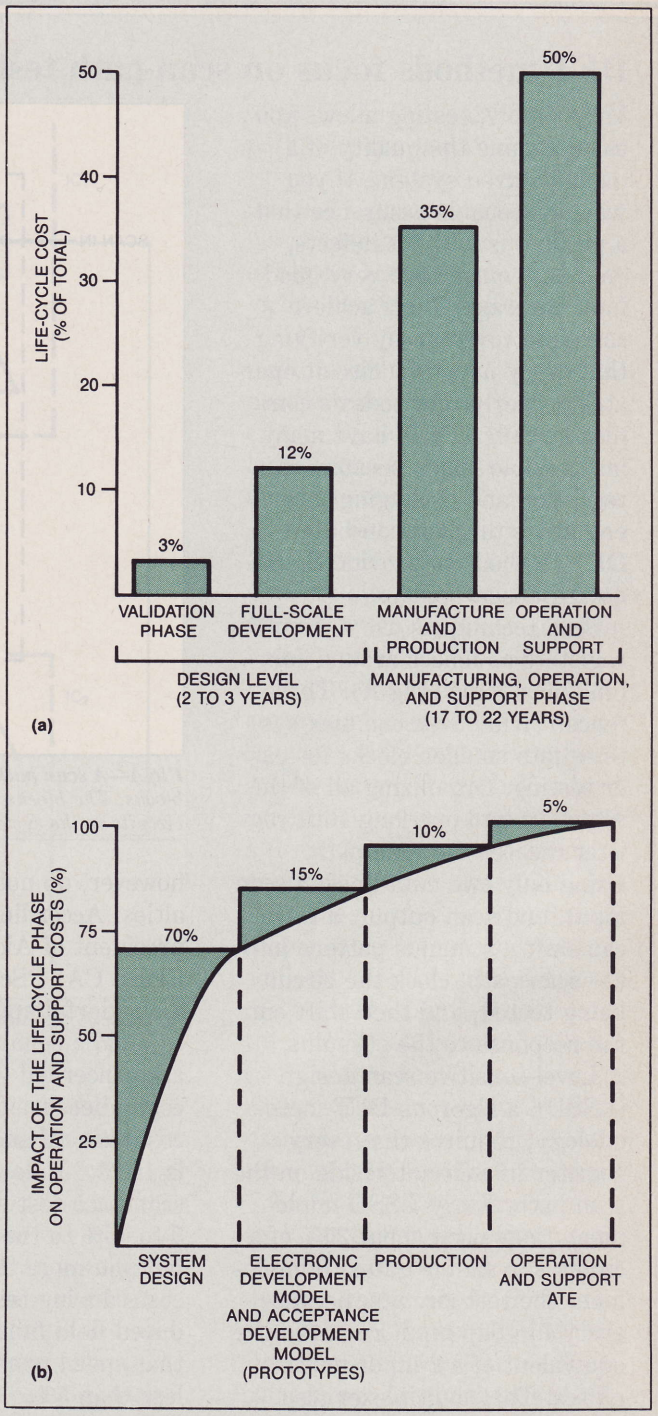

Fig 1—Although DFT methods add to the cost of a product's development phase, *the production, operation, and support phases represent the largest portion of the product's total life-cycle cost* ***(a).*** *And although the system-design phase of product development represents only about 15% of the total product-life-cycle cost, it has a 70% impact on that product's operation and support costs* ***(b).***

PART THREE
DECADE 90
The future of system design

DFT methods focus on scan-path testing

Very simply, testing allows you to determine the quality of a manufactured system. If you want reasonable assurance that a product is without defects, your test must have very good fault coverage. Tests achieve complete coverage by verifying that every node in a circuit operates properly. For today's complex circuits, which have many inaccessible nodes because of integration and packaging, most engineers turn to scan-based DFT methods to provide the required observability.

Scan techniques use a circuit's registers as fences around combinatorial logic (**Fig A**). These fences divide even complex systems into smaller blocks for easier testing. Organizing all of the registers into one long shift register creates the scan path. Using only two data lines, a scan input, and scan output, a tester can shift a stimulus pattern into the scan path, clock the circuit being tested, and then shift out the response to the stimulus.

Level-sensitive scan design (LSSD), a rigorous DFT methodology, requires that every register in a circuit reside on the scan path. Early LSSD implementations used about 25% of a circuit's available gates to implement the test circuits, because every flip-flop input required the equivalent of a 2-input multiplexer. That multiplexer also added an extra delay to the circuitry, slowing the system's maximum speed.

Recent test-circuit designs, however, do not inflict such penalties. According to Fred Bulow, president of Aida Corp (Santa Clara, CA), "Scan-path testing is a wonderful approach if you have to build a reliable product and are concerned with development costs, because test-development costs are nil and fault coverage is 100%." Bulow claims that scan-path test circuits add about 5 to 15% to the cost of a raw IC, but you more than recover those costs during testing and from reduced field failures. He adds that speed penalties amount to less than 5%.

As a concrete example of scan-path costs, consider the gate arrays from Integrated Logic Systems Inc's (ILSI, Colorado Springs, CO). The firm adds scan-path logic to its gate arrays by incorporating five extra transistors in each of the arrays' sequential cells. This test circuitry consumes a mere 0.4% of the total silicon die area and uses about 4.5% of the total interconnection on the circuit to link the sequential cells into a scan path. The scan path has a negligible effect on the gate array's maximum clock speed.

Though the costs of ILSI's test logic are low, the scheme provides tremendous benefits. For a 1438-gate design, ILSI's automatic test generator (ATG) created a set of test vectors in 13 minutes that provides 99.3% fault coverage. To produce a test

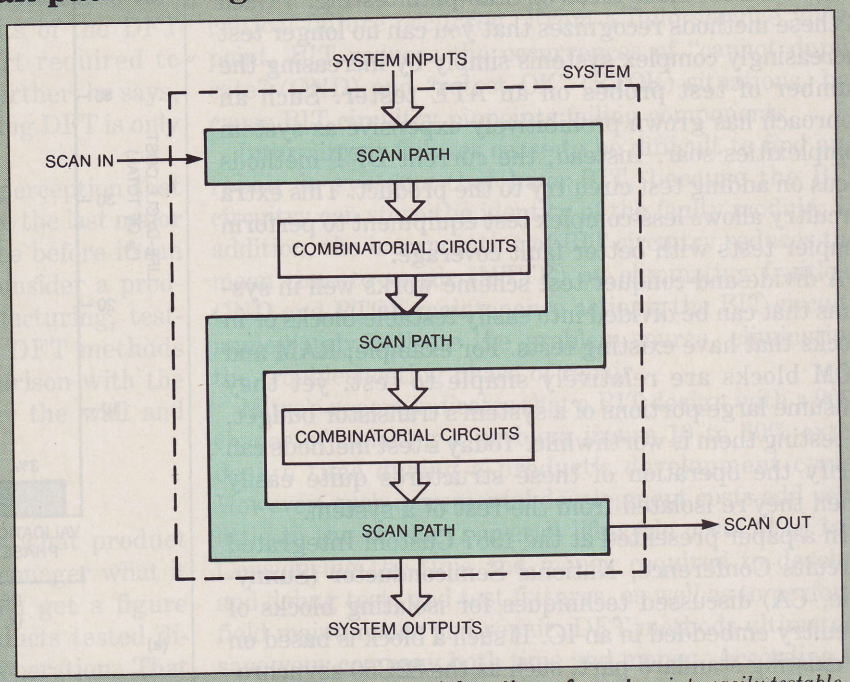

Fig A—A scan path divides the combinatorial portions of a system *into easily testable blocks. The blocks are surrounded by a scan path composed of all of the sequential circuits in the system, which are connected to form one or more long shift registers.*

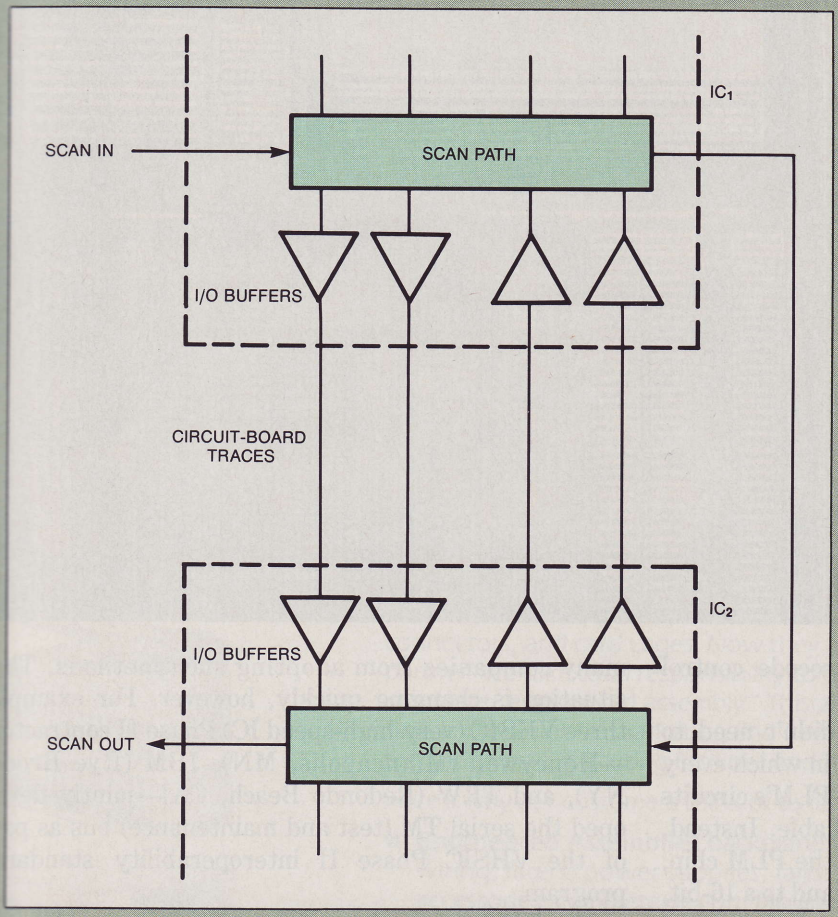

Fig B—The boundary-scan method lets you test all of the I/O buffers *and associated circuit-board traces by using serial test techniques.*

for the same design without testability circuits, a fault-grading program required almost 14 hours and provided only 85% fault coverage. Further, for a 2456-gate design, the ATG required 14 minutes to create a test with 99.88% fault coverage; for the same circuit without scan-path logic, the fault-grading program required a little more than 57 hours to create a test that yielded 59% fault coverage.

You can also use a form of scan-path testing called "boundary scan" to check interconnections between ICs. If each IC has a scannable register attached to its input and output buffers, the registers create a scan path surrounding the buffers and the pc-board traces (**Fig B**). An ATG can create a test with 100% fault coverage for this simple topology in a very short time. In addition, if the boundary scan registers are part of the IC's level-sensitive scan path, you can use those boundary-scan registers to test the IC's internal circuitry as well. JTAG's latest testability-bus proposal encompasses both boundary-scan and IC-testing capabilities.

Even if you merely add a scan path to your design, you'll still need a tester to check the circuitry. By adding a little more logic to the scan path, you can build the entire tester into your system. Such built-in self-test (BIST) and built-in test (BIT) circuits allow a system to verify its own operation on an ongoing basis.

Engineers designing BIST circuits usually employ a linear-feedback shift-register configuration (a procedure that's also called "signature analysis") and a pseudorandom test-pattern generator. These items both reduce the number of stimulus and response vectors stored in the self-test circuitry, and decrease the amount of time required for the test. Thus, to add BIST capability to a circuit that has scan logic, you require only the test-pattern generator and the signature-analysis feedback registers.

BIST also allows you to test an IC at full speed, a situation that is becoming less and less feasible on testers as clock speeds climb beyond 100 MHz. Because BIST circuits use the same types of transistors that the chip's other circuits use, the tests can easily run at the maximum possible clock rate.

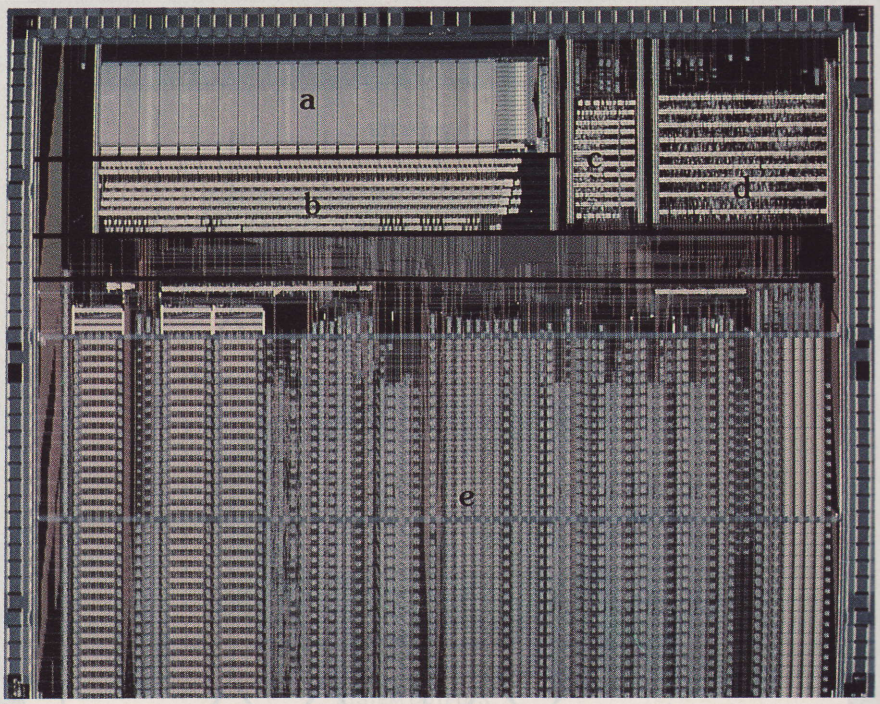

Fig 2—The PLM (Prolog Machine) μP developed jointly by NCR and the University of California at Berkeley incorporates two serial scan paths and extra microcode to aid in testing the chip. Engineers can completely test the microcode ROM ***(a)*** *with the first scan path* ***(b)****. The second scan path* ***(c)*** *reads a 16-bit status register that provides a gross indication of the operability of the chip's data path* ***(e)****. Microcode test instructions perform a more detailed check of the data-path's integrity and of the operation of the microsequence controller* ***(d)****.*

32-bit registers, and an 80,000-bit microcode control store **(Fig 2)**.

NCR's engineers decided that they didn't need to create a fully scannable design (a design in which every flip-flop is in the scan chain), because the PLM's circuits are already very observable and controllable. Instead, they designed two serial scan paths for the PLM chip, providing test access to the control store and to a 16-bit, data-path status register. Using these scan paths, test engineers can completely check the integrity of the microcode ROM and obtain a gross indication of the data path's operability. The scan paths add less than 5% to the chip's total area, but they allow a tester to check 70 to 80% of the μP's circuitry.

In addition, NCR's engineers incorporated extra microinstructions in the PLM's control store to facilitate detailed testing of the chip's data path. The company uses these microinstructions to check the IC's operation during manufacture. Self-test programs running on the PLM in a system can use them as well to monitor the chip's function while the system is operating.

So far, test standards have been lacking

Although companies such as NCR are already employing serial test methods to build complex systems, the lack of serial test-bus standards has prevented many companies from adopting such methods. That situation is changing quickly, however. For example, three VHSIC (very-high-speed IC) Phase II contractors —Honeywell (Minneapolis, MN), IBM (Rye Brook, NY), and TRW (Redondo Beach, CA)—jointly developed the serial TM (test and maintenance) bus as part of the VHSIC Phase II interoperability standards program.

The TM bus consists of four unidirectional lines, including a 6.25-MHz clock, a control line, and two data lines. A master test and maintenance controller uses the synchronous, backplane-level TM bus to check the status of as many as 32 modules in a system, sending data and control bits out on one unidirectional data line and receiving module status back on the second data line. In addition, Honeywell and IBM created an ETM (element test and maintenance) bus to allow an embedded test and maintenance processor to monitor as many as 32 individual devices within a module.

Other serial test standards are starting to appear as well. In 1985, Philips (Eindhoven, the Netherlands) started a test-bus study group that eventually became known as JTAG (the joint test action group), an ad hoc committee with representatives from European and US companies. JTAG hopes to create one serial test standard that IC vendors, board manufacturers, and systems

integrators can all use. To that end, JTAG directed its efforts towards developing a standard for a boundary-scan test system. JTAG's latest proposal, version 2.0, specifies both a boundary-scan test scheme and a standard test-access port that supports boundary-scan and other serial test methods.

The IEEE has also pursued a serial test standard through its P1149 working group. Jon Turino, co-chair of the working group and president of Logical Solutions Technology Corp, says JTAG's most recent specification has become the first element in the IEEE's serial test standard. That specification, P1149.1/JTAG, is a 4-wire subset of the full IEEE P1149 test-bus interface. However, the IEEE P1149 working documents incorporate extra levels of test capability. The P1149.2 subset consists of seven wires and supports additional serial interfaces besides the JTAG specification. P1149.3 and P1149.4 specify real-time digital and analog test-bus interfaces, respectively. A full implementation of the IEEE P1149 test bus requires 25 wires. P1149's proposal also includes fixed test protocols so that test-generation software can automatically create tests for boards and systems that incorporate scan circuitry.

Standard ICs lack test ports

Standards such as the JTAG and IEEE proposals promise to make DFT methods far more popular with engineers who develop systems based on ASICs, because engineers can include testability circuits in an ASIC definition without having to invent a DFT scheme. However, designers who design systems with standard ICs still face a major obstacle: Chip vendors have not taken a leadership position in offering testable parts.

One reason for this omission, of course, has been the lack of a standard test bus. But another reason, say IC vendors, is that customers have not requested testability features. Without market demand, the chip makers had little reason to add testability features to standard ICs. As standard semiconductor products grow in complexity, however, the same pressures that encourage engineers to use DFT methods for ASICs are forcing the IC vendors to add a variety of testability circuits to their standard parts. Such test circuits make the job of testing the individual ICs much faster and easier. System designers can then employ these on-chip test circuits for board- and system-level tests as well.

For example, Intel (Santa Clara, CA) added substantial testability circuitry to its 80386 μP. The circuitry included linear-feedback shift registers and pseudorandom counters for built-in self-testing, plus additional circuits that give the μP direct access to its translation look-aside paging buffer. These test circuits consume approximately 2% of the total silicon, but exercise 52% of the chip's 285,000 transistor sites. The company also took the unusual step of documenting the operation of those test circuits in the processor's data sheet so that any designer developing an 80386-based system could make use of the test circuits with a power-up, self-test software routine. Intel claims that several companies designing 80386-based systems are taking advantage of the μP's on-chip testability features.

Convincing management is tough

Pat Gelsinger, the Intel engineer who spearheaded the drive to add testability circuits to the 80386, says he had a hard time convincing the Intel's management to allocate silicon for those circuits. However, the DFT methodology made a great contribution to the 80386 project: It let Intel both obtain a fully functional device quickly and test production parts quickly. These facts produced a fundamental change in the company's attitude towards DFT methods. You can expect to see more testability circuits in future Intel products, says Gelsinger.

Another chip vendor, Texas Instruments (Dallas, TX), incorporated a complete serial scan path in its TMS320C30 DSP processor. A 4-wire serial test port emerges from the chip's package on four dedicated test pins. The DSP processor contains about 700,000 transistors. It will be available in the third quarter and will be one of the company's first ICs to incorporate testability circuits. Test circuitry uses about 10% of the chip. Although Texas Instruments perceived only a sma

Standard-test-bus information

If you would like more information about the JTAG testability bus, contact Rod Tulloss, the chair for JTAG's North-American working group: You can reach him at AT&T, (609) 639-6116. To obtain more information about the IEEE's P1149 testability-bus efforts, contact Jon Turino at Logical Systems Technology Inc, (408) 374-3650.

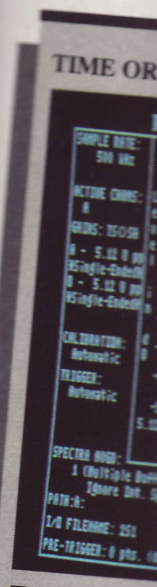

The R350 is a
time, FFT spe
ously. 32K data
disk. Two mode
inputs. Linear o
frequency curso
Print, store, retri
aliasing filters o

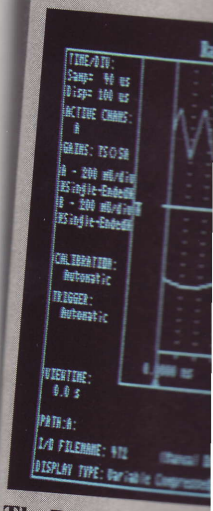

The R350 is a PC-b
Sample 2 channels si
per channel. EMI pro
Switchable differentia
Software-selectable g
500 volts. Full analog
adjust potentiometer.
Vertical waveform zoo
"C", BASIC, and Turb

N March 31, 1988

PART THREE
DECADE 90
The future of system design

amount of customer interest in testability, the company decided that increasingly complex designs such as the TMS320C30 required a DFT approach—just to help the company test the device during production. Now the company simply considers DFT methodology a part of good overall design practice.

Scan-path design will help the company's engineers fabricate the DSP processor and develop future versions of the part. "The TMS320C30 has a very modular architecture, and we can test each module independently to isolate fabrication problems," says Ray Simar, chief architect and program manager for the processor. In addition, when Texas Instruments broadens the TMS320C30 family, it will do so simply by adding new modules to the existing architecture. "That module will not be considered complete until it has the 4-wire test interface and test vectors," Simar says. The scan path provides an added benefit: It makes software development much easier by allowing an in-circuit emulator to read the μP's complete internal state through the scan path.

Companies attempting complex designs—such as Sequent, Intel, and Texas Instruments—already embrace DFT methods, because these methods allow them to build better-quality products at lower cost. More companies will follow their lead in the 1990s. Considering the benefits that DFT affords, and the fact that standards for DFT are imminent, the way is clear for you to adopt DFT methodologies now. **EDN**

References

1. Aadsen, Duane R and Sunil K Jain, "Automation of BIST for embedded RAM," Custom Integrated Circuits Conference, 1987, pg 66.
2. Breitenwischer, Thomas G, "Logic verification and production testing of nonstructured embedded VLSI blocks," Custom Integrated Circuits Conference, 1987, pg 62.
3. Gelsinger, Patrick P, "Built-in self test of the 80386," International Test Conference, 1986, pg 169.
4. Grasso, A and D J Clancy, "Built-in-test: An integrated maintenance concept," Government Microcircuit Applications Conference, 1987, pg 245.
5. Liu, Chun-Yeh, Kewal K Saluja, and J S Upadhyaya, "BIST-PLA: A built-in self-test design of large programmable logic arrays," 24th ACM/IEEE Design Automation Conference, 1987, pg 385.
6. *Tales of Testability*, Logical Solutions Technology Inc, Campbell, CA, 1986.

Article Interest Quotient (Circle One)
High 491 Medium 492 Low 493