Test in the Emerging Intellectual Property Business

In his keynote speech at the 1998 International Test Conference, held last October in Washington, D.C., Robin Saxby offered insight into his company's IP licensing business model and the importance of efficient and effective test to the success of products based on embedded CPU cores. Robin Saxby is chairman, president, and CEO of ARM and has led the company since its founding in 1990. Here, he and his co-author Peter Harrod, a manager in the Silicon Design group at ARM, summarize the presentation's main points.

ROBIN SAXBY
PETER HARROD
ARM Ltd

PRODUCT TEST, always a challenge even to traditional manufacturing companies, presents a special challenge to the "chipless chip" or intellectual property vendors like ARM. IP companies of course have emerged in recent years as part of the industrywide trend toward product specialization, as a result of the conflicting challenges of increasing complexity with reduced time to market. At ARM, for instance, we are primarily architects in that our CPU and System on Chip (SoC) designs are licensed to and sold by our global partners.

Before we can license a design, we must ensure that it is fully testable. In defining our test approach, we need to look first at potential design applications, at how we license the designs, and at the designs' intrinsic complexity. All are important considerations in defining an IP test strategy.

The advent of deep-submicron technology brought a move to system-on-chip design, with yesterday's printed circuit boards becoming today's chips. Such chips require that functional blocks designed by different engineering groups—often from different companies—be interconnected. Moreover, time to market is a major and familiar pressure on the system chip designer responsible for interconnecting these functional blocks. Of course, these chips must be fully functional and reliable, with a high degree of testability. Typical system chip designs today contain up to 10 million transistors, with 100-million-transistor devices on the horizon. Clearly, as design complexity mounts each year, so does the complexity of testing such designs.

ARM strategic thrust

With so many designs, so many potential applications, and a variety of licensee partners, how do we test to ensure that our partners encounter minimal production difficulties with our designs? After all, when combined, these applications and partners involve hundreds of test variables.

Our strategy is pragmatic and focuses on reuse and ways to optimize test vectors.

Reusable IP. Reuse is an implicit goal of all intellectual property vendors, and test is essential to achieving reuse. As a provider of high-quality IP, ARM supplies a complete set of deliverables for each of its IP blocks. An IP block is a block of electronic circuit

0740-7475/99/$10.00 © 1999 IEEE

function—it could be a complete microprocessor core or just a simple peripheral. An essential component of these deliverables is a set of manufacturing test vectors that the system integrator can use easily and effectively. Over the years, ARM has developed various kinds of test vectors using equally diverse techniques, each optimized for the particular chip design or application. Innovative approaches have often served as the key to efficient testing, just as they served our initial CPU core development.

Table 1. Vector count for the ARM7TDMI CPU.

Test technique	No. of vectors	Ratio
Parallel functional	23,519	1:1
Serialized functional (via 1149.1 port)	922,266	39:1
AMBA bus-based vectors (access via on-chip bus)	119,640	5:1
Full scan (ARM7TDMI-S)	3,173,865	135:1

In many cases, we've taken a hybrid approach to the testing of embedded cores and systems on chip, choosing the technique or combination of techniques that best fits the application. We naturally apply the traditional functional parallel test, as we did, for example, to our ARM7500FE SoC design, in which we tested each of the four major macrocells in turn, using multiplexed isolation. The ARM7500FE is a highly integrated, multimedia single-chip computer, comprising an ARM processor with cache, a floating-point unit, video and sound controller, memory controller, and I/O controller. Only a minimum number of external components are needed to create a complete computing system, such as an Internet appliance. But we have also, when the application required it,

- used serialized vectors applied via the IEEE 1149.1 port
- applied vectors to modules in turn using the Advanced Microcontroller Bus Architecture (AMBA) on-chip bus for test access
- used full-scan techniques (as, for example, with the recently announced, fully synthesizable ARM7TDMI-S).

The ARM7TDMI-S is suited to all traditional ARM7TDMI applications but has the benefits of flexibility and portability for those customers that want to exploit modern synthesis-based design flows. The hard macro approach in the ARM7TDMI, by comparison, gives lower power consumption.

A recent system-on-chip design combined full-scan for one block, block testing via an on-chip test bus for others, and a traditional functional test approach for the remainder. For each technique, our engineers trade off silicon area, number of test pins required, test time, time to generate the vectors, portability, and ease of debug.

The ARM7TDMI[1] CPU design is our best-known example of reusable intellectual property—it has been licensed to over 30 semiconductor companies and manufactured on more than 60 different manufacturing processes to date. The particular application in which ARM7TDMI is used can dictate the test technique that our engineers apply. For example, the high fault coverage achieved by the fully scannable ARM7TDMI-S may be a distinct requirement for automotive applications.

Table 1 shows a comparison of the number of vectors needed to test the ARM7TDMI and ARM7TDMI-S cores. Admittedly, the full-scan vector count for the ARM7TDMI-S looks high. Typically, test engineers would reduce this to match the vector depth of the rest of the system by splitting the scan chain into multiple parallel scan chains. It is also worth noting that another advantage of the full-scan vectors is that they achieve 99.9% fault coverage, whereas the other techniques achieve about 95%.

Functional test vectors. Test economics for reusable intellectual property blocks dictate that traditional functional testing still has a part to play. The high cost of developing high-coverage vectors without resorting to scan or automatic test pattern generation can be amortized over the many uses and reuses of the intellectual property. Our CPU cores—those that would be classified as hard virtual components (VCs) by the Virtual Socket Interface Alliance (VSIA)—have typically used full custom structures to minimize area and power. These structures don't lend themselves to automatic test insertion. A hard VC is one that is supplied as layout and cannot be altered by the user, as opposed to a soft VC, which could typically be synthesized to suit specific requirements (for example, speed or silicon area). These functional vectors can also be reused for at-speed testing, an important requirement for many customers and an essential tool for finding delay faults.

In 1995, we released AMBA with the goal of simplifying embedded CPU system design.[2] A key feature of this architecture is that it enables test access through reusing the on-chip bus infrastructure. Each separate AMBA-compliant module on a system on chip can thereby be tested independently. Engineers can use the on-chip bus for test isolation with a test harness as shown in Figure 1, next page.

Designers can add internal test registers so that users can apply block inputs independently from their connections to the rest of the chip. The only requirements for test re-

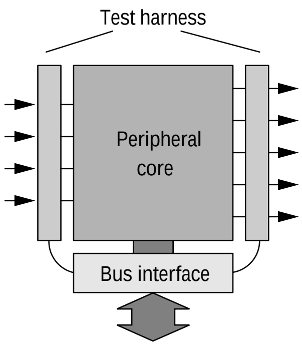

***Figure 1.** On-chip bus provides test access, and a test harness provides test isolation.*

configuration are to provide a 32-bit bidirectional port for test vector access and a method for controlling the system clock. In many systems, the 32-bit data bus provides a natural test bus port; on 8- or 16-bit systems, engineers can reconfigure a combination of data and address lines as a bidirectional test access port.

We are continually evaluating and exploring new and innovative test techniques. I_{DDQ} testing is ideally suited to fully static chip designs, such as those produced by ARM. Built-in self-test techniques also have a part to play, especially in system-on-chip designs where time to market is key. As the silicon area available to designs expands, so the opportunity for adding more on-chip test logic grows. This opportunity must, however, be balanced by the trade-offs of silicon area, power consumption, test cost, vector length, and time to market.

IP deliverables

When it comes to delivering a complete package of test deliverables to intellectual property licensees, manufacturing test vectors are just part of the story. Power and ac characterization vectors are equally important. We have developed an effective pre-fab characterization flow for our CPU cores. This means that we can characterize setup and hold times and propagation delays of our core I/Os without having to manufacture a test chip. As you can imagine, characterizing a tiny core inside a much larger test chip is not at all easy.

ARM's generic test vector format makes it easy to translate our vectors to the wide variety of formats in use by our partners. However, we are tracking the development of the P1450 STIL standard vector format and look forward to supporting an industry-standard format.

Verification and debug are other areas in which ARM is active. Verifying that a design will work in a system before committing it to silicon has been made easier for system integrators by means of codesign and coverification tools. Innovative products such as ARM's Multi-ICE ease the problems of multiple processor system debug. Also, provision of some on-chip hardware for postsilicon validation is now easier to justify with the increased silicon budget available to designers.

When validating CPU cores during simulation, our engineers often use test logic outside of the core (for testing interrupts, for example). Once the CPU core has been fabricated, it is useful to run the same validation sequences on the final silicon, especially each time the layout is ported to a new silicon process. If the validation test logic is instantiated on chip, this enables the validation sequences to be run at full speed on the final chip. The provision of a small bank of SRAM on-chip can also enable at-speed self-test.

Our partnership with the EDA community has resulted in a rich set of models that interface with most simulators. These models enable our customers to carry out codesign and coverification before a design is committed to silicon, leading to faster, more accurate IC solutions.

We have always recognized the importance of efficiently testing reusable intellectual property. In so doing, we fully support the initiatives of the VSIA (see http://www.vsi.org for more information). The VSIA Manufacturing Test Development Working Group works in two main areas to address the fact that efficient manufacturing test creation and application is key to the successful exploitation of reusable IP:

- Test data interchange between virtual component providers and system chip integrators
- Guidelines to virtual component providers regarding DFT and test isolation

The efforts of the P1500 Embedded Core Test standards committee will also play an important part in defining industry standards in this important field (see http://grouper.ieee.org/groups/1500 for more information).

THE FUTURE LOOKS ROSY for both IP developers and integrators. We are finding ways of more easily exploiting the extra silicon area that is now becoming available as process technology shrinks yet further. Initiatives such as the VSIA and the Virtual Component Exchange (VCX; http://www.vcx.co.uk) in Scotland will make IP use and reuse even more commonplace. Yet testing all this functionality will remain a major challenge, one that IP developers, IP integrators, tools providers, and the DFT community must continue to tackle with enthusiasm. D&T

References

1. S. Segars, K. Clarke, and L. Goudge, "Embedded Control Problems, Thumb and the ARM7TDMI," *IEEE Micro*, Vol. 15, No. 5, Oct. 1995, pp. 22–30.
2. D. Flynn, "AMBA: Enabling Reusable On-Chip Designs," *IEEE Micro*, Vol. 17, No. 4, July/Aug. 1997, pp. 20–27.

Send general questions and comments to Robin Saxby, ARM Ltd, Liberty House, Moorbridge Road, Maidenhead, Berkshire, SL6 8LT, UK; Robin.Saxby@arm.com. Specific questions about DFT should be sent to Peter Harrod, ARM Ltd, 90 Fulbourn Road, Cambridge, CB1 9JN, UK; Peter.Harrod@arm.com.