ASSET INTERTECH

Whitepaper:

IEEE P1687 Internal JTAG (IJTAG) taps into embedded instrumentation

By Al Crouch
Chief Technologist, Core Instrumentation
ASSET InterTech
Co-Chairman, IEEE P1687 IJTAG Standard Working Group

ASSET™
Driving Embedded Instrumentation

ASSET InterTech, Inc.
2201 N. Central Expressway, Suite 105
Richardson, TX 75080, USA

http://www.asset-intertech.com/

Table of Contents

Table of Figures

Executive Summary

For many years, chip makers have routinely embedded test and measurement functionality into their high-end, high-speed devices. This was done out of necessity by the chip makers. They had found that this was the most effective and cost-efficient way to characterize, validate and test their devices.

Now, these high-end chips have entered the mainstream. As a result, the industry is realizing that there is a great wealth of test and measurement, and design-for-test (DFT) intellectual property (IP) embedded on-chip that can be put to good use in a wide range of applications to more cost-effectively validate, test and debug chips, circuit boards and systems throughout their entire life-cycles. The new IEEE P1687 Internal JTAG (IJTAG) standard is a step in this direction.

This paper describes the objectives of the IEEE P1687 working group, how it achieved those objectives and the resulting IJTAG architecture that is embodied in the standard.

In the Beginning…

Like many standards bodies, the IEEE P1687 committee was intent on specifying a standard that met the real needs of the industry. In fact, before it was designated a standard committee, the IJTAG working group polled the industry extensively to determine the perceived deficiencies of existing standards that pertained to embedded instrumentation, such as the IEEE 1149.1 boundary-scan (JTAG) standard and the IEEE 1500 core test standard. In a very real sense, the five years of work that has been invested in developing the IEEE P1687 standard has been a long and involved response to the working group's original findings concerning the needs of the industry.

One of the primary objectives of the IJTAG working group was to streamline how embedded on-chip instruments functioned by defining a standards-based interface for these instruments. The committee believed that this would simplify the task of using these instruments and it would create an opportunity for the development of a third-party marketplace for tools that would complement this IP. In addition, committee members felt that eventually a market for portable third-party embedded instrumentation could develop.

To complete a first version of the IEEE P1687 IJTAG standard, the committee borrowed from the IEEE 1149.1 boundary-scan (JTAG) standard. As a result, IEEE P1687 initially reflects certain architectural features of the boundary-scan standard. For instance, IJTAG re-uses boundary scan's concepts of a Test Access Port (TAP) and controller. Moreover, the IJTAG access network for embedded instruments incorporates a set of registers that is similar to the Test Data Registers (TDR) found in the boundary-scan standard and which typically comprise every boundary-scan chain.

The On-Chip IJTAG Architecture

Although similar, boundary-scan and IJTAG scan chains have several very important differences. For example, boundary-scan chains are fixed in length and composition while IJTAG networks are dynamic and variable in their configuration. In fact, segments in an IJTAG scan chain or network can be added or subtracted as requirements change. Unlike IEEE 1149.1 boundary scan which is based on a concept of fixed instructions, the configuration of an IJTAG chain is controlled by the variable data passing through the chain's data registers.

Consequently, the IEEE P1687 IJTAG standard is flexible from an architectural standpoint. Designers can deploy various configurations of an IJTAG network in order to meet a wide range of engineering, cost, operational and other tradeoffs. These architectural configurations are documented by defining them in the IJTAG standard's Instrument Connectivity Language (ICL). ICL essentially describes where the IJTAG TDRs are, the scan paths that connect and access them, how and when these scan paths should vary, the connections between the IJTAG scan paths and the boundary-scan TAP controller on the device, and the parallel connections between the embedded IJTAG instruments and the IJTAG TDRs.

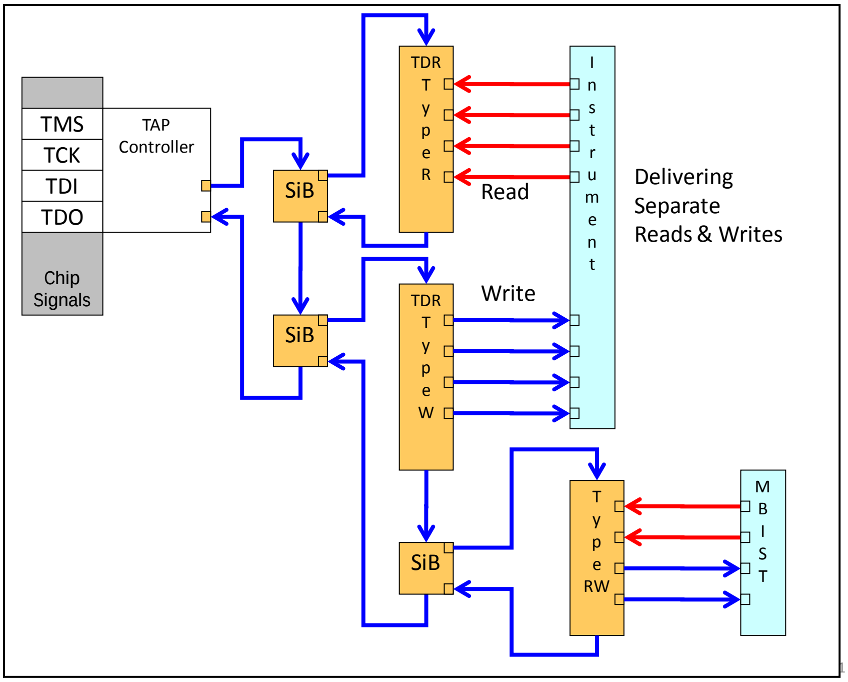

Figure 1: This example of an on-chip IEEE P1687 architecture includes a IEEE 1149.1 boundary-scan TAP controller with its four signals and three instances of IEEE P1687's Segment Insertion Bit (SIB) which allows on-demand access to instrument interface registers.

Segmenting the Network

In the IEEE P1687 standard the relative equivalent of IEEE 1149.1 boundary scan's Instruction Register (IR) would be the Segment Insertion Bit (SIB). Instructions to select any given boundary-scan TDR come from the chip's IEEE 1149.1 boundary scan TAP Test Controller IR or the Wrapper Instruction Register (WIR) which is defined in the IEEE 1500 core test standard. To create the active chip-level scan chain that includes the TDR associated with an instrument, an instruction encoding would be installed which would connect the selected TDR to the chip's Test Data In (TDI) and Test Data out (TDO) pins. This would allow serial scan data to pass through the chip and by the embedded instrument's internal signals.

Difficulties arise with a "per-instruction" arrangement when the number of embedded instruments exceeds approximately 10. At 50, the situation is absolutely unwieldy. Of course, many contemporary system-on-a-chip (SoC) devices can have more than 300 memories and anywhere from 30 to 300 built-in self test (BIST) instruments to test them.

The problem centers on the limitations of the typical IEEE 1149.1 boundary-scan architectures, of which there are two: 1.) a single, long daisy-chain configuration; or 2.) each instruction targets a specific TDR.

With a long daisy-chain scan path the test access time may increase since all instruments must be accessible on the daisy-chain and their access must be active whenever the daisy-chain instruction is selected. Additionally, instruments embedded in low-power partitions in a design will not functional well in this sort of architecture because the daisy-chain will be broken when these instruments are powered down or are put in a sleep mode for periods of time. Moreover, the daisy-chain architecture is at the highest risk of catastrophic breakdown. A problem on a single bit on a daisy-chain scan path, such as a blocked bit or a stuck-at, will render the entire architecture useless. When the daisy-chain is broken at a certain point, access will be denied to all instruments beyond that point.

Boundary scan's other common architecture, the one-instruction-per-TDR architecture, limits operations because only one instrument can be selected by the one active IEEE 1149.1 boundary-scan instruction. (IEEE 1149.1 allows only one instruction to be active at a time.) The problem here is twofold. First, fast test times based on the concurrent operation of multiple instruments is not possible since only one instrument can be selected at a time. This would be disconcerting to chip manufacturers who are under extreme pressure to reduce test time and thereby reduce test costs. Reducing test time means that a chip would spend less time on an expensive ATE test system where in-socket costs of eight cents a second can quickly add up to unacceptably high cost levels.

One way to reduce chip test time is to schedule multiple tests at the same time. Unfortunately, chip designers often don't know which test they can or want to run concurrently until the chip is on the ATE system for manufacturing test development. So, if there were 50 instruments on a chip and the developers knew they wanted to run at least six tests simultaneously, they would have to develop IEEE 1149.1 JTAG instructions for every possible combination of six instruments of the 50 that are on-chip. This would come to more than 15 million instructions! This instruction explosion is a problem that can be moderated to some degree with compound instructions that group specific instrument accesses together, but this leads to a loss of flexibility for test scheduling because simultaneous instrument operations are typically only determined at design time.

The IJTAG committee has foreseen these types of difficulties with both the instruction-based and the daisy-chain-based boundary-scan architectures and has attempted to solve them architecturally with SIBs as instruction-bits embedded in the data scan path. A SIB is a single TDR bit that can be placed anywhere in the P1687 architecture. There is no limit on the number of SIBs deployed. A SIB is something of a gatekeeper. It can open (add) or close (subtract) a scan path segment. In this way, the SIB acts as a single external bypass register capable of providing access to an embedded instrument. Scheduling tests through the various instruments on a chip is as simple as placing logic 'ones' in certain scan path configuration registers and executing an 1149.1 Data Register Update (commonly referred to as an UpdateDR). When this has completed, access to all of the needed instruments that have been identified will be available in an active scan chain so that multiple tests can be launched simultaneously. Instruments on scan path segments but with their SIBs set to 'zero' will remain off the active scan path and as a result they will not contribute any additional bit-length to the instrument access path.

Once instruments are started, the scan paths providing access to them can be closed without stopping the instruments – a capability that is unique to IJTAG networks. In fact, other IJTAG scan operations can be going on without disturbing the instruments on a closed scan path since the rules dictate that the TDR freezes to the last state before the SIB closes. Only an intentional access or an 1149.1 JTAG reset will change the state of a scan path that is not actively selected. Various algorithms may be applied to open SIBs, communicate (read and write) with instruments, and to close SIBs. These algorithms, known as SIB Handling Strategies can be applied to minimize test time, minimize access time, maximize concurrent instrument operations, minimize network serial data or other optimizations.

Test Portability for Life-Cycle Cost Savings

A major advantage of the IJTAG standard involves test vector re-use. Test vectors are written to the instrument interface and can be retargeted and re-used with any 1149.1 boundary-scan TAP. Since the IEEE P1687 ICL describes a particular device's IJTAG access network, retargeting already developed test vectors is simply a matter of reprocessing the ICL so that it reflects the targeted device's IJTAG architecture. In other words, test vectors can be re-used beginning with chip-level design verification and continuing through wafer probe, IC package test, circuit board test, system test, yield analysis, and even troubleshooting, debug and repair of user returns.

Another language that is defined in the IJTAG standard, the Procedural Description Language (PDL), represents the test vectors or operational procedures that are applied directly to instruments. PDL lets the engineer control and automate the operation and scheduling of embedded instruments independently of any IJTAG access network they may be connected to. For example, aside from the basic iReads and iWrites (that is, IJTAG reads and writes), PDL features a number of flow-control operatives such as if-then-else, for-next, do-while and others. The functionality of PDL surpasses that which is included in IEEE 1149.1 boundary scan's current *de facto* language, Serial Vector Format (SVF), which applies vectors to a device TAP and must be re-developed with each new architecture.

The limitations of SVF severely restrict the portability of test vectors. For example, SVF tests associate the necessary boundary-scan or IEEE 1500 instructions as well as the data and operating sequences with the embedded instrument in a particular device. Unfortunately, when that instrument is embedded in another chip, the SVF tests must be recreated all over again because SVF code only functions in relation to each device's Boundary-Scan Description Language (BSDL) file, a detailed description of the boundary-scan architecture on the device and, by necessity, device-specific. That is, when an embedded instrument is re-used in multiple chip designs, each design will likely have a different IEEE 1149.1 boundary-scan architecture. Consequently, a SVF test vector that was developed for a certain chip must be completely rewritten for every other device where the test vector is re-targeted.

IJTAG's PDL addresses these issues. PDL describes a process, test or operation in terms of the direct interface on the embedded instrument (PDL is meant to provide the documentation on how to use an instrument. This is missing from the current 1149.1 and 1500 documentation since this documentation spells out how to use the 1149.1/1500 architecture and instructions). PDL is complemented by ICL, which describes the chip's IEEE P1687 IJTAG architecture, although in far less detail than a Verilog or VHDL model would. Essentially, PDL creates a portable vector format that can accompany the instrument no matter where or in how many devices it is embedded. For example, the same instrument might be deployed several times in a chip, or at different levels within a chip's hierarchical architecture. Or it might be implemented in several different chips with distinct configurations and logic.

The 3 Portions of the P1687 Architecture

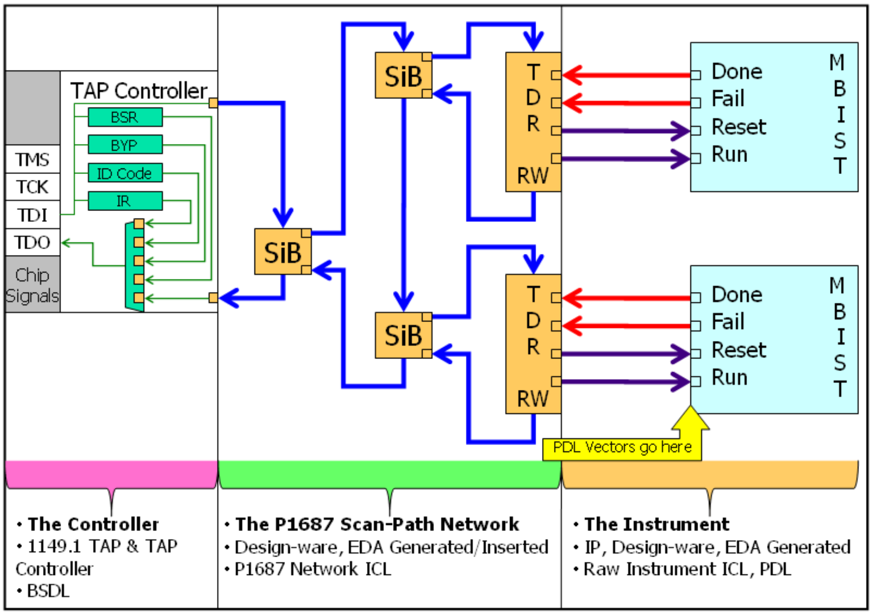

Figure 2: This illustration shows a chip's boundary-scan TAP controller, the device's internal IJTAG scan path network and how PDL provides a portable interface for the embedded instrument.

From the Inside Out…

Dating back to the advent of the electronics industry, the dominant approach to validation, test and debug has been from the outside inward (pins-in). Until recently, physical access points like test pads on circuit boards or pins on chips were rather abundant. It was a simple matter of placing a physical probe on an access point so that measurements could be taken or vectors inserted and observed.

But now, the industry is changing. In response to the current trajectory of technology, the test and measurement component of the electronics industry is becoming much more introspective. Out of necessity, test and measurement methods are taking on an inside-out approach. Device pins are hidden under silicon these days. And test pads are often forbidden on high-speed board designs because the capacitive effect of a pad will disrupt the signaling it was meant to validate or test.

The cumulative effects of these and many other forces at play in the industry are rapidly bringing non-intrusive and embedded validation, test and debug methodologies like IEEE P1687 IJTAG to the forefront.