

IEEE 1149.6 – A Practical Perspective

Bill Eklow, Cisco Systems Inc., San Jose CA, beklow@cisco.com

Carl Barnhart, Cadence Design Systems, Ojai, CA, cfb@cadence.com

Mike Ricchetti, Intellitech, Durham, NH, miker@intellitech.com

Terry Borroz, Teradyne, North Reading, MA, terry.borroz@teradyne.com

Abstract

The IEEE 1149.6 standard was approved in March of 2003. The standard extends the capability of the IEEE 1149.1 standard to include AC-coupled and/or differential nets. These nets are predominant in new, multi-gigabit serial technology.


This paper will present a short overview of the 1149.6 standard and the issues that it addresses. The paper will then discuss design, verification and test considerations that must be dealt with when implementing 1149.6 logic on a component or a board. Since 1149.6 addresses a problem that is in part analog in nature, there are several new issues that must be dealt with which are not covered in previous 1149 based standards. The paper discusses the nature of these issues and their impact on design, verification and test. The paper is intended to be tutorial, providing tips and techniques to allow the user to quickly understand and implement an 1149.6 based design.

1 Introduction

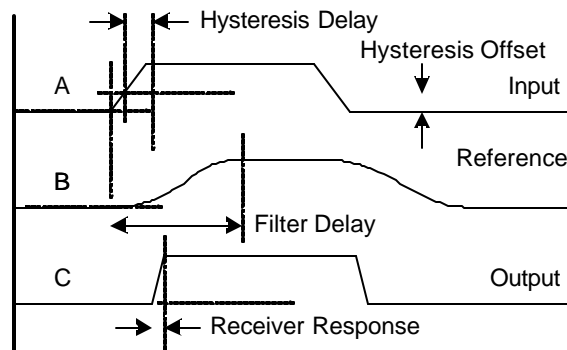
The IEEE 1149.6 standard was developed to address the requirements of boundary-scan testing of AC-coupled, differential nets. The use of AC-coupling creates an RC network which causes signals to decay over time. The data must be captured within a time window after the UPDATE-DR state. This is not supported by the current IEEE 1149.1 standard [1]. Differential signaling is also not adequately supported by 1149.1. The standard allows an “Analog” model in which the boundary scan cell is placed internal to the mission driver and receiver, or a digital model which requires a boundary cell on each pin. Both of these solutions are undesirable for high speed serial logic.

The 1149.6 standard proposes a solution which allows placing a boundary cell internal to the mission driver and

a “test receiver” on each pin of the receiving device. The test receiver captures transitions, not levels (as is done by 1149.1). This allows 1149.6 to capture data independent of the subsequent decay. To deal with potential unknown offsets 1149.6 proposes a “self-referenced” hysteretic comparator (Figure 1) to capture the signal transition [2]. The input to the receiver, a delayed version of the signal and the output of the hysteretic comparator are also shown in Figure 1. It is this implementation that is the core of the 1149.6 standard.



Self-referenced comparator



Response to AC coupled signal

Figure 1: Self referenced comparator with response

2 Design Issues

The IEEE 1149.6 Standard introduces a number of changes to the on-chip structures familiar from 1149.1. Most are quite straightforward, both to understand and implement, however, there is one requirement that is completely new.

For example, on the output side of the chip, there are only minor changes. An output boundary scan cell for an AC Test Pin requires a new input which is exclusive-ORed with the output of the Update cell. This new input is driven by a new signal output from the TAP, which will toggle on falling edge of TCK when an AC test instruction is active and the TAP is in the correct state. There are new mandatory instructions defined which will cause the signal to toggle. These are straightforward design changes to the TAP and to the boundary cell, and the new boundary cells themselves are described in detail in the Standard [2].

The typical use of a single boundary cell as input to a differential driver does not necessarily change. However, in some high performance SERDES designs, multiple boundary cells are placed on the inputs to the serializer to avoid placing the boundary cell in the high performance path.

On the input side of the chip is a completely new test receiver, and one of these test receivers is required on every AC Test Pin input pad (not one receiver for a differential pair, as in the past). In combination with a traditional observing boundary scan cell for each, the test receiver has two jobs: first, to detect transitions at the chip pad that meet minimum requirements in terms of voltage swing and transition time; and secondly, to detect one of the following three conditions: a rising transition, a falling transition, or no transition at all.

When an input to a chip may be AC coupled on a board, there may be no way of knowing what the time constant of that coupling will be at the time the chip is designed. Indeed a chip with AC pins may be used in a number of different environments, each one with a different time constant. Also, when designing a receiver, it is not always known exactly what the driver characteristics will be. It is not unheard of to mix drivers and receivers of different technologies, or even different protocols on a board. In this environment, the only signal characteristic that remains constant is the transition itself. Therefore, that is what we detect.

For the sake of understanding and defining the requirements, the standard assumes that the test receiver will be implemented as a hysteretic comparator with some special capabilities. The hysteresis levels must be set at a level that is appropriate for the channel transmission protocol, possible attenuation, and the test mode. Additionally, the hysteretic comparator should not respond unless the transition is sustained for a minimum amount of time that is again based on the transmission protocol.

In order to maximize the noise margin, the standard requires that the fastest time constant in an AC coupled path be allowed to decay away between transitions of the test signal when the new AC test instructions are active, allowing hysteresis to be set relative to the full amplitude of the transition in each direction.

In addition, the test receiver must be able to support both the new AC test instructions of this standard, and the static tests mandated in IEEE 1149.1. When an IEEE 1149.1 instruction is active, the test receiver must revert to level sensitive operation. At best, only half the transition amplitude is available for setting hysteresis levels, and that assumes that the transition level or common mode voltage at the input pin is known.

When looking at possible board defects, it becomes apparent that one of the syndromes that must be recognized is that there is no transition when there should be one. So this condition must be explicitly recognized in order to detect and diagnose the possible defects. In order to assist in reporting the lack of any transitions, the hysteretic comparator must have the means to maintain the value in the boundary-scan cell if no transition is observed during the test. One way to do this is to preload the hysteretic memory with the value in the boundary-scan cell. If no transition is detected, then that same value will be captured into the scan cell. This has the similar effect as if the value in the scan cell were being driven back onto the input pin by a weak driver, and for test purposes the value scanned into the AC pin receiver boundary-scan cells must be treated as a signal on the board.

Some of these characteristics of the hysteretic comparator, such as pre loading the hysteretic memory, are a matter of straightforward design. It is either present or it isn't. Understanding how to set the analog characteristics of the hysteretic comparator deserves a little more explanation.

Generally, the driver of a channel or signal has only a few characteristic specified, most notably the voltage swing, minimum and maximum, and either a transition time or maximum frequency, under certain loading conditions. These relatively sparse specifications must be understood and used to derive most of the analog characteristics of the test receiver. Complicating the calculations is the fact that we're really concerned about the voltage swing and transition time at the receiver pad. So the worst case conditions of use must also be understood (or assumed and documented) in order to estimate attenuation and other effects that might be observed at board test time.

Assuming that the driver and receiver are physically very close on the board, the voltage swings and transition times of the receiver would be virtually identical to those of the driver.

Assuming that the driver and receiver are physically far apart on the board, the voltage swings will be attenuated and the transition times increased. How much may be able to be determined with analog simulation of worst-case networks, or may have to be set simply by engineering judgment. The attenuation of the voltage swings is frequency dependent, but since the maximum frequency during board test for these signals is half the TCK frequency, and the maximum TCK frequency is specified in the BSDL, the chip and board designers may be able to force this attenuation to be dominantly resistive, and therefore calculable. High-frequency attenuation of the transition, however, is not dependent on the signal frequency during test. As stated in the Standard, the best estimate for transition time at the receiver is one-quarter of the cycle time at maximum frequency (one-half of the bit time). If the transition time were worse than this, then there would be severe voltage attenuation at the maximum frequency, possibly rendering the link inoperable.

Once these three values are known, minimum and maximum voltage swing and maximum transmission time at the receiver, the Standard provides straightforward rules for calculating hysteresis levels for both receiver test modes (AC and DC), the minimum duration of a signal above the hysteresis level for detection, and the minimum time constants for both the coupling and the edge detection low pass filters. The rules are described in Clause 6 of the standard [2], and several design examples with different protocols and different termination schemes are shown in Appendix C. While the rules create specific limits, there is still room for engineering judgment. An understanding of the board environment during test,

particularly the types of noise sources that may be present, and the variability that can be introduced to any one of the fundamental channel characteristics by chip and board process variations, is critical to successful design.

For instance, setting the hysteresis voltages too low will adversely affect the ability to reject noise. However, setting the hysteresis too high could prevent the test receiver from operating at all. The discussions in Clause 6 and Appendix C attempt to explore many of these ideas and possibilities, but cannot possibly be complete. It remains up to the design engineer to understand the behavior of the circuits.

The design of this test receiver is the major task required of technology suppliers for adoption of this Standard.

At the chip level, the designer must know which I/O pins require AC Test capability, and select and instantiate the drivers, receivers, test receivers, and boundary cells that support this Standard. The BSDL must be coded to contain the necessary descriptions of the AC pins, and to reflect the maximum TCK frequency for the chip. These steps may be automated for the designer as the Design For Test Synthesis tools start supporting the Standard.

The chip designer should include the assumptions about test receiver design, such as minimum voltage swings, tolerable threshold levels for DC tests, etc. in the chip specification sheet to guide the board and system designers.

At the board level, the designer must pay attention to the treatment of AC test pins, and to the design and layout of the AC coupling, if it is used. The chip designs may place constraints, documented in either the BSDL or the spec sheet, on the time constants of the coupling and the topology of the coupling and termination. Some pins may require AC coupling, others just allow it. In addition, it is often important to ensure that the coupling and termination, if on the board, is as close to the receiving chip as possible. As always, noise and noise coupling considerations are important. During board test, the chips may exceed their normal simultaneous switching limits at Update_DR, requiring very good power supply decoupling during test.

At the system level, if there are long signal paths from board to board, the attenuation of the signal should be compared to the assumed minimums for test if these connections are to be tested in the system.

3 Verification

Compliance checking involves verifying that the behavior of a IEEE Std. 1149.6 design, as observed from its I/O pins, conforms to the rules and permissions of the standard. Best practices for design verification requires that compliance checking be performed both as part of pre-silicon verification, and post-silicon verification, efforts. Pre-silicon compliance checking is generally done using simulation, both digital and analog in the case of 1149.6, and helps to assure that the IC complies with the standard prior to fabrication, and that the BSDL file is syntactically correct. In addition, some structural checks to verify IEEE Std. 1149.6 compliance may also be necessary during this phase of the design. Post-silicon compliance checking is recommended as a final certification step and to verify that the BSDL file that will be released to customers matches the silicon.

As 1149.6 utilizes and is compatible with the existing IEEE Std. 1149.1, it must first comply with 1149.1. As such, a major part of 1149.6 compliance checking is verification of the digital behavior of its 1149.1 infrastructure. This includes verification of behavior for the 1149.1 mandatory instructions (i.e., EXTEST, BYPASS, SAMPLE, PRELOAD), the length of the IR and DR registers, and association of Boundary Register cells with pins [1].

It should be noted that whereas pre-silicon compliance checking for IEEE Std. 1149.1 has required only digital simulation, the input test receivers for the 1149.6 standard mandate certain analog circuitry and characteristics that may require other methods for compliance verification. These may include analog simulations and/or structural checks of the design.

The following section provides an overview of some of the additional compliance checks that are necessary for IEEE Std. 1149.6. In depth aspects of compliance verification, conformance and documentation requirements for 1149.6 are beyond the scope of this paper; as such *the reader is encouraged to refer to Clauses 7 and 8 in the standard for more comprehensive details* [2].

3.1 Mandatory IEEE Std. 1149.6 Instructions

The following should be verified for the two new mandatory 1149.6 instructions, EXTEST_PULSE and EXTEST_TRAIN [2]

1. Verify that these instructions have been provided in the BSDL description. The opcodes for EXTEST_PULSE and EXTEST_TRAIN are defined by the designer and should be provided in the BSDL (or other device specification, if the BSDL is not yet available).
2. Verify that each instruction becomes effective at the falling edge of TCK in the *Update-IR* TAP Controller state. The instruction will drive values on the AC output pins at this time.
3. Verify that all output pins or bi-directional pins in output mode behave in accordance with the rules in 5.3.1 of the standard while the EXTEST_PULSE instruction is effective.
4. Verify that all output pins or bi-directional pins in output mode behave in accordance with the rules in 5.4.1 of the standard while the EXTEST_TRAIN instruction is selected.
5. For 3 and 4 above, AC pins should be tested with both AC behavior selected and DC behavior selected (if AC/DC selection cells are provided), as well as enabled and disabled (if CONTROL cells are provided). DC pins should be tested both enabled and disabled (if CONTROL cells are provided).
6. Check that system values are driven out onto AC pins in the *Test-Logic-Reset* state of the TAP Controller.

3.2 Analog Behavior of AC Pins

The analog pin behavior of input test receivers and output drivers must also be checked for compliance. It is expected that the generated test signals be driven out of the mission mode driver at the normal system drive levels and edge rates. Input test receivers must reconstruct the signal, as driven by an output driver, when an AC test instruction is in effect. However, when the IEEE Std. 1149.1 EXTEST instruction is in effect, the test receivers should behave as a level detector. AC/DC selection cells are optional and provide the ability to disable the AC test signal on individual output drivers, causing them to behave as DC EXTEST drivers.

3.2.1 Input Test Receivers

The following properties and behavior should be verified for input test receivers. Verify that all AC input and bi-directional pins have a test receiver monitoring the associated pin. Single-ended pins should have one test input receiver and differential pin pairs should have one input test receiver per leg. Each input test receiver should be verified to operate in both level detect and edge detect modes. Additionally, it should be verified that each input

test receiver meets the required designer specified parameters for ΔV_{Min} , ΔV_{Max} , T_{Trans} , V_{Hyst_Level} and V_{Hyst_Edge} , as defined by the standard [2]. Verification of level detection and edge detection operation is also described in further detail in the standard.

It should be noted that the voltage and time parameters above are not specified in the BSDL. They must be provided separately. It may be appropriate for detailed parametric verification of the I/O circuit to be performed by the I/O circuit designers prior to silicon, and incorporated into chip manufacturing test. The above parametric values will have to be communicated via means other than BSDL, such as in the device data sheet.

3.2.2 Output Drivers

In addition to the driver behavior verified during compliance checking of the EXTEST_PULSE and EXTEST_TRAIN instruction, the following additional properties and behavior should be verified for single-ended and differential drivers of an AC output pin or channel.

1. For standard IEEE Std. 1149.1 instructions that drive outputs (EXTEST, CLAMP and optionally RUNBIST or INTEST) [1], verify the following:
 - That these drivers drive data from a single associated Boundary Scan register cell.
 - When the content of the Update flip-flop of the associated Boundary-Scan Register cell changes, each enabled output driver produces the correct transitions on its output.
 - That a disabled output driver produces a quiescent (i.e., non-driven) state on its output(s).
 - Where implemented, verify that such drivers are disabled or enabled via the content of a control cell in an associated Boundary-Scan Register cell.
 - Where optionally implemented, verify that Boundary Scan register cells on output drivers of AC output pins capture a fixed logic one or zero value with the SAMPLE instruction.
2. For the standard EXTEST_PULSE and EXTEST_TRAIN AC test instructions, verify the following:
 - That such drivers drive data from a single associated Boundary Scan register cell and that the cell is the same register cell as the one associated with the pin in item 1 above.

- That when disabled, such output drivers produce a quiescent, i.e. non-driven state on its output(s).

3.2.3 AC/DC Selection Cells

Optional AC/DC selection cells allow the ability to disable the AC test signal generated at an output driver during the EXTEST_PULSE or EXTEST_TRAIN instruction. When the AC behavior is disabled the output driver provides a fixed logic level, as if the DC EXTEST instruction was operating. The following behavior should be verified for all AC/DC selection cells:

- Each AC/DC selection cell is associated with at least one AC pin driver and no output drivers are controlled by more than one AC/DC selection cell.
- The structure and operation via the TAP Controller of AC/DC selection cells should be compliant with IEEE Std. 1149.1 control cells.
- When each AC/DC selection cell is updated with a 0, any associated AC pin data cells controlled by the AC/DC selection cell should drive their fixed logic levels for both the EXTEST_PULSE and EXTEST_TRAIN instructions, verifying that AC pin output drivers behave as if EXTEST is in effect.
- AC/DC selection cells capture the values as specified by the design, i.e. either a fixed value or the contents of their Update stages.

3.3 Production Test

Following pre-silicon and post-silicon compliance verification, devices are understood to be properly designed and compliant with the IEEE Std. 1149.6. During volume manufacturing then, it is important that each device be tested for proper fabrication of its' testability circuitry.

While many of the compliance verification or characterization tests could be reused or leveraged here, these tests will likely be overly time-consuming and focused on verification problems (like the correspondence of BSDL to silicon), rather than detection of manufacturing defects. As such, production testing should focus on proving that the testability circuitry was fabricated properly, and this can be divided into two main tasks: verifying the digital 1149.1 infrastructure logic (e.g., TAP Controller and Boundary Scan cells) and the analog characteristics of the test receivers.

The digital logic in IEEE Std. 1149.6 may very well be subordinated into a higher level testability scheme such as internal scan, and thus tested as a portion of the general logic of the device for production test. If this is not the case, then some of the tests above can be adapted and streamlined for production test of the digital logic. Structured test methods are often preferred for the digital logic, as they can readily provide high defect coverage.

Tests for analog characteristics of the test receivers can be performed by applying waveforms with controlled edge rates and voltage swings to the AC inputs of the device. Referring to Figure 41 of the standard [2], this amounts to setting up input waveforms with known valid edges and known invalid edges, where these transitions should be detected and rejected, respectively.

4 Board Level Test

While test generation for 1149.6 is similar to 1149.1, there are several significant differences. This section describes the additional work that must be done within test tools to expand a regular 1149.1 interconnect test to also perform 1149.6 tests of AC interconnects. Much of this material is also described in section A.3 of Annex A in the 1149.6 standard [2].

At a high level the circuitry described in previous parts of this presentation allows the interconnect test generator to treat AC interconnects as if they were plain wires (assuming correct timing and sequencing). This concept should be applied to testing AC interconnects alongside the DC testing normally done with 1149.1.

4.1 Determining the interconnects to be tested

Traditional 1149.1 test generation engines will already provide a list of nets to be tested with traditional DC techniques. A new list must be developed which specifies those nets which are AC-coupled and include 1149.6 capable drivers and receivers, and also DC-coupled differential nets which with 1149.6 capable drivers and receivers. 1149.6 provides greater test coverage for differential nets and thus should be preferred over 1149.1 when the drivers and receivers are 1149.6 capable.

4.2 Assigning unique IDs for all interconnects

Traditional interconnect tests (Wagner [3]) drive a unique ID onto each DC net and verify that it is received properly. Uniqueness allows detection of shorts between interconnects, since a short will cause a failure during test

vectors where the ID's differ. Both AC and DC coupled nets should be assigned unique ID's in an 1149.6 environment. This would allow short detection between AC and DC coupled nets. Note, however, that for differential nets, the differential signals must be assigned complementary ID's.

4.3 Choosing timing and instructions

The previous section described how to assign ID's and generate test vectors for an AC and DC coupled design. However, for the 1149.6 circuitry to work properly, it is also necessary to use specific TAP state trajectories and timings when applying each bit in this sequence.

4.3.1 Calculating TTest

The "TTest" parameter defined in section 6.2 of the standard [2] is the minimum time between signal transitions that is acceptable to an 1149.6 test receiver. This value may vary based on the test receiver and coupling capacitor on the net. The test generation software must insure that the test's minimum transition time is greater than the largest TTest across all the AC interconnects on the board. So the test generation software must use the largest TTest across all of the interconnects on the board as a "whole board" value.

A receiver's TTest value can be calculated by going through the following steps in order and using the first one that applies:

Step 1: Check to see if there is a TLP time specified for the receiver in its BSDL. If so, then the receiver is a "self-referenced" receiver, as defined in rule 6.2.3.1.a)1). [2] In this case, set TTest to $3 * TLP$.

Step 2: Check to see if there is a THP time specified for the receiver in its BSDL. If so, and if the "On_chip" keyword is also present, then the receiver is a "fixed-reference" receiver, as defined in rule 6.2.3.1.a)2), [2] with an on-chip coupling capacitor. In this case, set TTest to $3 * THP$.

Step 3: If neither TLP time, nor an on-chip THP time is specified in the BSDL the receiver is a fixed-reference receiver with an off-chip coupling capacitor. The time constant of the board-level coupling circuit must be calculated and TTest set to three times this value.

The process described above should provide a reliable TTest value. It may be prudent, however, to also provide a mechanism for manually overriding the calculated value.

4.3.2 Choosing instructions, state trajectories, and timing

In an 1149.6 interconnect test, ICs with 1149.6 capability will use one of the two new instructions provided by the standard: EXTEST_PULSE or EXTEST_TRAIN. These instructions require the TAP to go through the Run-Test/Idle state between each update and capture. Depending on the board topology, additional sequencing may be required.

4.4 Using the EXTEST_PULSE Instruction

The simplest approach occurs in situations where all 1149.6 ICs can use EXTEST_PULSE. This is the preferred way to test 1149.6 capable logic, however, use of the EXTEST_PULSE instruction may be precluded in some instances (described below). In response to the EXTEST_PULSE instruction the driver will transition the net upon entry into the Run-Test/Idle state and transition the net again upon exit from the Run-Test/Idle state. The final transition will be to the level (value) which was loaded during UPDATE_DR. The final transition will be captured by the test receiver in the subsequent CAPTURE_DR state.

To ensure that the final transition occurs properly, all 1149.6 devices must be kept in the Run-Test/Idle state for a minimum time (TTest) as calculated in 4.3.1 above. The time spent in Run-Test/Idle could increase if a "wait duration" value is specified for any affected IC in the "AIO_EXTEST_Pulse_Execution" attribute in its BSDL. [2] In this case, the time spent in the Run-Test/Idle state must be the greater of the maximum TTest and the maximum wait duration. (Again – manual override may be prudent).

4.5 Using the EXTEST_TRAIN Instruction

It is possible that some circuits may not be able to operate properly with the single pulse provided by EXTEST_PULSE instruction. The EXTEST_TRAIN instruction provides a train of pulses which are synchronized to the falling edge of TCK. As with EXTEST_PULSE, these signals pass through the AC coupling network and the final transition will be sensed by the test receiver after exiting the Run-Test/Idle state.

The determination of which ICs (if any) require EXTEST_TRAIN will typically be made by the component and/or board designer based on their analysis of the circuits involved. The test software should provide a way for the test engineer to specify which ICs must use the EXTEST_TRAIN instruction.

If any component requires the EXTEST_TRAIN instruction, the simplest setup would be to use it in all 1149.6 ICs. In this case, the test engineer needs to be careful to make sure that the pulse width specified by the EXTEST_TRAIN instruction is longer than the TTest parameter which was described above. If the EXTEST_TRAIN pulses do not fully decay between transitions, it is possible that the data capture could be unreliable. If these two parameters cannot be reconciled, then it will be necessary to perform both EXTEST_TRAIN and EXTEST_PULSE instructions. Obviously, careful consideration of TCK frequency must be taken into account when making this analysis.

If both EXTEST_TRAIN and EXTEST_PULSE are required, it is important to keep in mind all of the considerations which must be taken into account for both instructions:

- There may be a minimum number of required transitions in the EXTEST_TRAIN instruction
- There may be a minimum frequency requirement in the EXTEST_TRAIN instruction
- Both instructions require the data signal to completely decay before the next transition

This can be controlled by the TCK frequency and the duration spent in the Run-Test/Idle state. However, the test software must be careful to make sure that these requirements are met for all devices. More information can be found in Annex A.3.4.1.4. [2]

4.6 Choosing hysteresis presets

In order to detect whether a valid transition has occurred on a given net, the memory element in the test receiver must be "pre-loaded" with the opposite value that the data is transitioning to. This way the transitioning data will overwrite the preloaded (hysteresis) value in the test receiver. A net which does not transition will cause the preloaded data to be retained. There is, however, value in considering cases where the preloaded data is the same as the level which the data is transitioning to. In addition to testing functionality in the test receiver itself, unexpected data transitions (transitions in the opposite direction) may also be detected by this method.

A set of 4 vectors will allow testing for potential board level failures and for correct operation of the test logic. The four vectors could be generated as follows:

- Drive high to low (0), expect a low (0)
- Drive low to high (1), expect a high (1)

- Preload low, no transition, expect a low
- Preload high, no transition, expect a high

There are several ways to derive these vectors. One way is to define a fixed pattern of preset values, based on the net ID's, that can be used for all receivers and is guaranteed to bear all four relationships at some point during the test. For some types of the unique short-detection IDs mentioned above, you can simply define one additional unique ID and use it as a universal preset sequence, and this will guarantee a 0 preset value when expecting a 1 and a 1 when expecting a 0 at some point during the test for each receiver. If this can be combined with presets guaranteed to match fixed capture values elsewhere in the pattern set, then this will achieve the desired result. In addition, the fixed pattern of preset values generated by this technique can also be used as expect values in unconnected receivers to test that they are not shorted.

Another way is to derive the preset values based on the algorithm used to generate the test vectors. Vectors could be generated which match alternating 0 and 1 preset values against successive zeros in the capture sequence and then independently match alternating 0 and 1 preset values against successive ones in the capture sequence.

A third possibility would be to add four additional scan vectors solely for the purpose of testing the four combinations of preset vs. expected value.

Note that the additional test capability provided by these "hysteresis presets" exists for IEEE 1149.1 testing as well.

4.7 Accounting for differential signaling

The test pattern generator must account for the fact that a value shifted in to a differential driver is received unaltered at the plus side test receiver but complemented at the minus side receiver.

4.8 Generating vectors assuming no AC coupling

An 1149.1 interconnect test can be used to improve shorts detection. The idea is to treat each AC interconnect as two separate nets (one on each side of the capacitor). The test should assume no signal propagation through the capacitors, and must allow enough time between update and capture for the effects of the coupling capacitors to die out.

This test will require a new set of vectors. The hysteresis preset capability will be used in this test in order to load

an "undriven" expected value. The sequence of presets for such a receiver should form an ID that differs from all other unique IDs in the shorts test. This will allow a short to one of these receivers to be detected. In addition, this will cause a failure if a capacitor is shorted. This is important because a shorted capacitor will not cause a failure using the EXTEST_PULSE and/or EXTEST_TRAIN instruction.

As mentioned above, this test must spend enough time in Run-Test/Idle to allow signals propagated through the coupling networks to die out before capture. The amount of time required to wait is calculated the same as TTest described above, except that step 1 of that process should be skipped. This in effect sets the waiting time to 3 times the time constant of the coupling network. Of course, the actual waiting time must be the maximum of all these values across the entire board. (And note that, once again, this is another case where a user override would also make sense.)

4.9 Miscellaneous test generation considerations

The following items were not described above but must be done in order to support 1149.6 test generation:

- Upgrading BSDL parsers in order to get at the new information added to the BSDL language to support this standard.
- Support for the use of AC/DC selection cells when steady states must be maintained during AC tests.
- 1149.6 tests of boards with multiple scan chains must guarantee simultaneous transitions through states that affect AC test circuitry (UPDATE_DR, UPDATE_IR, Run-Test/Idle and SELECT_DR_SCAN).
- Noise implications (SSO and crosstalk) should be considered when generating test vectors. A method for reducing noise by performing double scans is shown in section A.3.4.5 of the 1149.6 standard. [2]

4.10 Diagnostic issues

A basic strategy for diagnosing a failing DC interconnect test is to examine the behavior of the failing interconnects, looking for behavior that can be explained by likely faults. For example, in order to detect opens the diagnostic software typically looks for "stuck" receivers. To detect shorts, the diagnostic software can search for pairs of failing interconnects whose behavior can be explained by the short between them, with the assumption that a zero

will overpower a one in the event of a conflict.

It should be possible to detect board level defects in a similar manner for AC coupled interconnects, there are, however some differences. For example, opens will cause a test receiver to capture its preset value. In addition, a shorted capacitor in an AC interconnect will cause the test receiver to (incorrectly) capture the values being driven by its associated AC driver during the DC-only test mentioned in section 4.8 above.

In general, understanding the likely behavior of faulty interconnects is a subject that requires more study at this point, and is beyond the scope of this paper. (It probably deserves a paper in its own right.) Note that a good place to start in exploring this subject is the extensive library of simulation data that is available on the 1149.6 working group web site (<http://grouper.ieee.org/groups/1149/6/>).

One upgrade that will definitely have to be made to the diagnostic software is to somehow let it know the hysteresis preset values for all the test receivers throughout these interconnect tests.

4.11 Tester hardware

Note that AC interconnect testing will probably require special drive and detect circuitry on the tester for AC signals that go to the board edge.

4.12 Additional tests for special situations

There are special situations that might benefit from additional tests beyond those for which the 1149.6 standard was designed. These situations include:

- A plain 1149.1 driver AC coupled to an 1149.6 receiver.
- Any driver AC coupled to a plain 1149.1 receiver.

There are potentially useful testing methods for these situations that are deprecated but which might be useful in some situations. Describing these techniques in detail is beyond the scope of this paper. They are described in section A.3.4.4 of Annex A of the standard [2].

5 Conclusion

With the proliferation of high speed (multi gigabit) serial data communication protocols, it is expected that there

will be a significant number of devices on which IEEE 1149.6 is implemented. In fact, several 1149.6 enabled chips are already either in design or production. Tools are in development which will enable 1149.6 insertion and verification. Board test software is also being upgraded to enable testing of 1149.6 components and their interconnections to other boundary scannable components on the board.

While 1149.6 is an extension to IEEE 1149.1, there are several new and significant considerations which must be taken into account during design, verification and (chip and board) testing. The most significant aspect of the 1149.6 architecture is the “analog” test receiver. Designers will need to pay close attention to the selection of the analog parameters to ensure that functionality, performance, test coverage and noise immunity will not be compromised. Designers will also need to consider both board and chip level environments as each will impact the other. Some analog verification will be required to ensure correct performance of the test receiver. Changes in the digital verification will be primarily centered on the transition through the Run-Test/Idle state and should be relatively straightforward. Board level test software will need to differentiate 1149.6 enabled components from 1149.1 compliant components which are not 1149.6 capable. Some work will need to be done to develop vectors which can detect and diagnose defects in a “hybrid” 1149.6/1149.1 environment. 1149.6 requires preconditioning the “hysteresis” memory of the test receiver. Finally, test software must also consider the best way to transition through the Run-Test/Idle state in order to reliably generate and capture valid data transitions.

That being said, IEEE 1149.6 provides imperative test coverage for high speed serial components. Careful consideration of both board and chip parameters will provide significant test benefits for these high speed designs.

6 References

1. IEEE Std 1149.1-2001, IEEE Standard Test Access Port and Boundary-Scan Architecture.
2. IEEE Std 1149.6-2003, IEEE Standard for Boundary-Scan Testing of Advanced Digital Networks.
3. P.T. Wagner, “Interconnect Testing With Boundary Scan,” Proceedings, International Test Conference 1987.