Testing Bridges to Nowhere - Combining Boundary Scan and Capacitive Sensing

Stephen Sunter
LogicVision
sunter *at* logicvision *dot* com

Kenneth P. Parker
Agilent Technologies
kenneth_parker *at* agilent *dot* com

Abstract

As printed circuit board dimensions continue to decrease, in-circuit tester (ICT) access using a bed-of-nails plus capacitive sensing is increasingly difficult. Stimulus injection using IEEE 1149.1 Boundary-Scan has been proposed as an alternative, but without modification it has significant limitations. An IEEE-supported Working Group is developing an extension entitled, "1149.8.1 - Draft Standard for Boundary-Scan-Based Stimulus of Interconnects to Passive and/or Active Components". It would add capabilities to 1149.1 that facilitate testing of connections to non-Boundary-Scan components, especially passive components and vacant connectors that are connected to devices equipped with 1149.8.1 facilities. This paper describes existing limitations, IC design changes that would address them, some experimental results, and a summary of how this proposed standard is evolving.

1 Introduction

Recent papers [2][8] studied the concept of using Boundary-Scan driver resources to substitute for ICT resources for testing components on boards using capacitive sensing technologies.[1] As bed-of-nails access becomes more difficult, it was noted that with some effort, Boundary-Scan [3] with its EXTEST capability could both "ground" (guard) and stimulate board nodes that were connected to a passive device-under-test (DUT), typically a vacant connector or empty socket, but still use a floating capacitive sense plate and associated tester resources for detecting test signals. The papers reported on experiments that demonstrated the concept working on real boards in production environments, and some of the limitations.

Several limitations were observed to be caused by the definition of EXTEST itself, and one paper [2] offered some ideas for a new instruction and stimulus capabilities that would make such testing more effective. These ideas became the seed for a proposed standard: IEEE P1149.8.1[2]. The "8.1" suffix anticipates that there could be other new standards, "8.2" etc., that could further extend Boundary Scan.

2 Capacitive Opens Test Background

Capacitive opens test was invented over 15 years ago and makes use of ICT bed-of-nail access to apply a small (400 mV peak-to-peak) AC signal to one board node at a time that supplies a signal to a component under test, as diagramed in Figure 1. The other nodes that are attached to the connector are connected to ground by the ICT. The injected AC signal is then detected by a floating sense plate (cut to fit the outline of the connector) gently mounted over the mouth of the open connector, and not *into* the connector. A photograph of a floating sensor plate assembly above a device is shown in Figure 2.

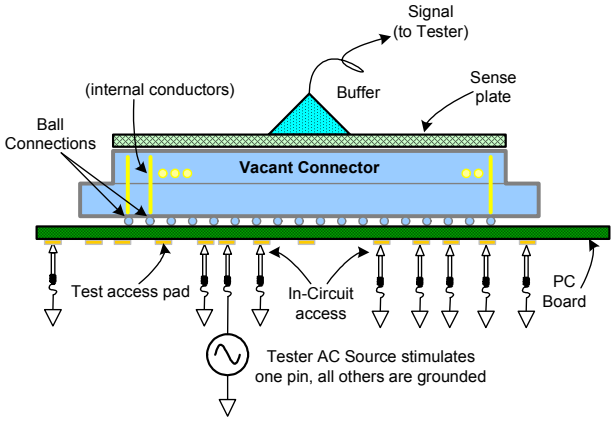

Figure 1: Capacitive opens test setup for a ball-grid array connector

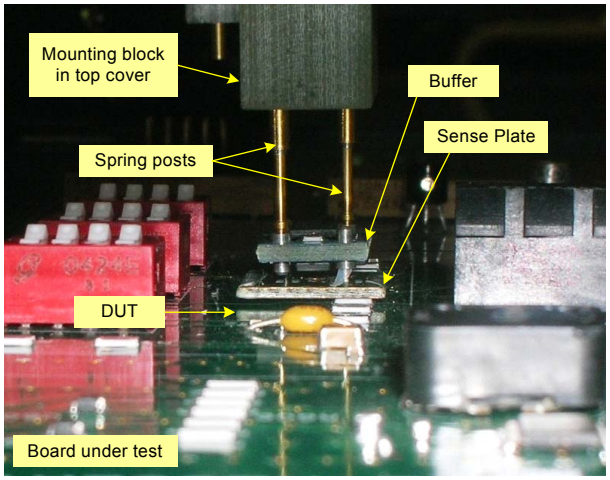

Figure 2: A DUT with a suspended sense plate and buffer.

This plate is capacitively coupled to the tip of each connector pin inside the connector, so the AC signal appears on the sense plate – the capacitance is typically less than 1 pF. A local buffer atop the plate amplifies this

[1] Capacitive sensing technology is over 15 years old and is familiar to many ICT users of "TestJet ®" technology.

[2] The authors are members of the "Selective Toggle" working group (P1149.8.1). This paper is our interpretation of what this working group is developing. As with any developing standard, it and the technology described herein are subject to change.

978-1-4244-4867-8/09/$25.00 ©2009 IEEE

signal and transmits it back to the ICT system where the measured capacitance is calculated. The capacitance is compared to the expected "good" value and if the measured value is within measurement tolerance, the path from the signal source to the sense plate is judged as complete. If the solder ball that should exist between the connector and board is missing (open) then the capacitance reading will be significantly smaller and an open to that pin can be diagnosed. Each pin in sequence is stimulated while the others are guarded and each expected capacitance can be verified. The guards help to channel the electric field from the tested pin to the sense plate and help to block unwanted signals from the PC board which might confuse the measurement. This technology is proven in production testing for finding open solder joints between board components and the board itself.

This paper will mostly discuss test of connections to empty connectors ("bridges to nowhere"), but this type of testing is widely used for checking connections to a variety of other components too, such as non-boundary-scan ICs, DIP switches, sockets, resistor packs, etc. for open defects. The technique depends on ICT nail access, although as access decreases, measurements are generally still possible on remaining pins – that is, the technique is relatively graceful in coverage degradation as access is eliminated. However, as access becomes severely limited, a new approach is needed to regain fault coverage.

It is important to note that with significant access, good coverage of shorts between the nodes attached to the tested device is achieved. As access is eliminated, conventional shorts test coverage degrades as well.

Another important feature of capacitive opens test is that the board is unpowered while the test is running. Thus, non-functional test conditions, like guarding many nodes to ground, can be applied which would be impractical with the power applied.

To detect opens and shorts at the site of these devices, 1149.1 Boundary-Scan [3] and 1149.6 [5] should be useful, but they are aimed at testing interconnections between ICs that are compliant with those standards. When a printed wire to a passive component such as a connector must be tested, Boundary-Scan is inadequate.

Nevertheless, Boundary-Scan appears to have the resources to implement a surrogate for the ICT stimulus and guards[3] shown in Figure 1. The referenced papers from ITC'08 [2][8] showed this is possible. However, the example shown in Figure 3 reveals one of the coverage holes: 1149.1 does not require all IC pins to have drive capability when EXTEST is in effect, so input-only pins present a problem.

Briefly stated, 1149.1 has three major shortfalls that affect capacitive measurements, and a shortfall that impairs shorts detection:

a. Not all Boundary-Scan pins can drive (for stimulus or guarding);
b. The frequency at which a stimulus pin can toggle depends on the TCK frequency divided by the boundary register length;
c. Differential outputs generate complementary waveforms that cancel at a capacitive sensing plate.
d. Not all Boundary-Scan driven pins can monitor their own pin states (called self-monitoring) to determine if shorts may be present.

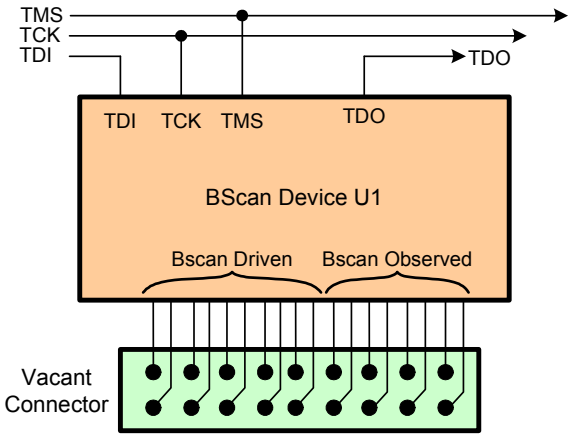

Figure 3: Boundary Scan device connected to connector.

The papers [2][8] noted these problems and the new P1149.8.1 working group set out to define a new instruction (or several) that would help solve these problems. The goal is to recover both opens and shorts coverage in the area of passive components that is lost due to increasingly dense and inaccessible circuit nodes. This would allow board testing with sustained efficacy in the face of reduced access, and would lead to lower fixturing costs as well.

3 Possible solutions, and challenges

To address the objective of generating a toggling signal on any selected pin of an IC, the first solutions considered were the existing 1149.1 [3], 1149.4 [4], and 1149.6 [5] boundary-scan standards (especially since the authors are still active members of the respective IEEE working groups). It was found that several fundamental rules in those standards prevented their addressing the toggling problem.

3.1 IEEE 1149.1 Boundary Scan

IEEE 1149.1 requires use of the whole boundary scan register (BSR) to load in another logic state for a pin, so producing an AC signal is very inefficient when scan chains are long. The standard permits new instructions to refer to portions of the BSR, so this potentially facilitates limiting BSR length during toggling. However, to maintain a constant toggle frequency, the TCK frequency would need to continually change depending on the BSR length in each IC. Mandating a constant BSR partition length was not considered.

[3] The "guard" concept is different. Rather than a DC guard to ground, Boundary-Scan gives us the ability to hold a signal at steady-state logic 0 or 1, as an "AC-guard". This is sufficient for capacitive testing.

1149.1 does not adequately address differential pins, considering them to be analog, and only requires observing and controlling the single-ended result, at which point it is difficult or impossible to detect faults in only one leg of a differential pair. In a low-noise situation, an open circuit in one leg of a differential pair will typically have no impact on the received signal since the unconnected input will be terminated to the common mode voltage, V_{COM}, and the other leg's signal will swing around V_{COM}.

IEEE 1149.1 cannot detect a fault on an input-only pin if that net has no output connected to it in another IC. Nor does 1149.1 require pins to be bidirectional, and if they are 3-state bidirectional, they must be declared as bidir and require at least two boundary scan cells.

IEEE 1149.1 uses only function-mode output drive and input threshold, so when two outputs having equal drive are shorted, they can potentially produce a voltage that is mid-way between logic 1 and 0. As discussed in [2], when the voltage is capacitively sensed, the two shorted nets will have similar 50% voltage swings, and if they have similar coupling to the capacitive sensor, the sum is 100% and becomes indistinguishable from the fault-free case. In theory, it is also a problem for detection with 1149.1 inputs because mid-way between logic 0 and 1 is the optimal and common choice for input threshold. As a result, the captured logic value for a bridging short is indeterminate since it will be sensitive to process variations.

Requiring input thresholds to be *not* exactly mid-way between logic 0 and 1, as is noted in 1149.4 ("input threshold should be a voltage different than that which two shorted equal drive outputs would deliver" (see [4], rule 6.2.1.1, Note 1), would improve detection by Boundary-Scan but would not affect detection by capacitively-coupled sensors.

3.2 1149.4 Mixed-Signal Test Bus

IEEE 1149.4, the mixed-signal standard bus for boundary scan, was also considered. This standard facilitates the diagnosis of opens in wires to nowhere: the capacitance of wire connected to a pin can be measured, for example, by applying an AC signal to the pin via the AB1 analog bus and measuring its amplitude via the AB2 bus [15]. The standard is also suitable for diagnosing differential pin faults, and almost any shorts, but requires 2 or 4 analog TAP pins, analog changes to all the pad I/O cells, and analog buses to all signal pins, which may be impractical for more than a few hundred pins due to leakage concerns [16]. The standard was not intended for use in purely digital ICs (though it could be useful for contactless test of I/O parameters [9]). Its use requires board-level analog buses, stimuli, and measurement, which are usually considered inappropriate for a digital board.

3.3 1149.6 Boundary Scan for Advanced I/O

IEEE 1149.6, the boundary scan standard for advanced I/Os (specifically, AC-coupled or differential signals) that is gaining widespread adoption, has the infrastructure to toggle any pins at one half the TCK frequency. The standard requires conveying the half-TCK-frequency "AC Test Signal" to all outputs to generate the AC stimulus needed to detect the presence of series capacitors. However, it requires an extra boundary scan bit for every group of pins to be toggled. To be able to toggle every pin, one at a time, would double the number of boundary scan bits, and would not address the other 1149.1 problems discussed earlier. Furthermore, toggling at half the TCK frequency can be too high a frequency when the TCK is above 1 MHz because the noise bandwidth becomes too large, yet a high TCK frequency is needed to minimize the time to serially shift boundary scan data in and out.

3.4 Make Input-Only Pins Bi-Directional

Clearly, input-only pins will require a stimulus source if those pins may be connected to board wires that only travel to empty sockets. Therefore, they must become bi-directional. But how much output drive should they have? If all "input" pins have equal output drive, then the problem arises as described earlier for mid-rail voltages from bridging shorts. Requiring too much drive will significantly increase input capacitance, and requiring too little drive may be insufficient to overcome pull-ups.

The challenge to adding a toggle capability to output or bidirectional pins is to do it without doubling the number of boundary scan cells (as 1149.6 would), and without significantly increasing the number of signals routed to all boundary scan cells. Several new signals appear necessary: one originating from the BSR bit that indicates which pin is to be toggled, another which indicates the logic level of each non-toggled output pin, a global signal to indicate that toggle mode is in effect, and another to convey the toggling clock (like the AC Test Signal used by 1149.6).

3.5 Differential pins

Differential outputs present a particularly challenging requirement because a capacitively-coupled sensor near a board can only detect the sum of nearby AC signals, and the sum of a differential pair of AC signals is intentionally designed to be as constant as possible. To sense a non-zero sum, the differential output drivers must deliver a non-differential signal. We investigated altering a differential driver's design so that the non-inverting output had a different rise time than the inverting output. Designers who were asked about this rejected it as too intrusive and performance-affecting (even though the signal requirements would only apply in test mode).

Differential drivers have very little capacitance in their high-speed signal path, to minimize power, and a symmetrical design to maintain complementary behavior. Any technique for unbalancing them must not add any capacitance to high-speed signal nodes, and must not

affect one side differently than the other.

Any new boundary-scan standard is more likely to be accepted by industry if pad I/O cells do not need to be modified, especially circuit nodes that are connected directly to pads, because the design of electro-static discharge (ESD) tolerant circuitry is exceptionally time-consuming and risky – after a manufacturer has qualified a library of pad cells, they are extremely reluctant to make any changes. With this in mind, an additional specification objective for differential output drivers (and any output drivers) is that the desired test-mode behavior be obtainable using pre-existing I/O pad cells. The 1149.6 standard has been relatively quickly adopted by many companies because it does not require any changes to differential output drivers. (Unfortunately, that standard requires significant analog changes to differential inputs but no other solution has been found that facilitates diagnosing which leg of a differential pair is faulty.)

3.6 Detecting Pull-up Resistors

Pull-up (and pull-down) resistors are common on many boards, and their resistance is easily tested by bed-of-nails ICT. With diminished access, consideration was given to detecting at least the existence of pull-up resistors, and possibly their resistance too.

Detecting the resistance value is not presently a priority because only the first and last placed resistors of each resistance value need to be tested to ensure that the correct reel of resistors was used (this is referred to as "coupon" testing). However, with further decreased physical access, other solutions may be needed.

Detecting the presence or value of a pull-up resistance is not presently addressed by capacitively-coupled sensing. So we considered using boundary scan.

One way to detect the existence of a pull-up, using boundary scan, is to first drive the relevant output pin to logic 0, then disable or tristate the pin during the UpdateDR state of the TAP controller, and then quickly capture the logic value at the pin, 2½ TCK clock periods later in CaptureDR state, as described in [10] for IC testing. For board testing, this approach requires another IC's input pin to be connected to the same wire or the output pin to be bidirectional. It also requires the TCK clock frequency to be fast enough that normal leakage current will not cause the signal to rise to the threshold voltage in those 2½ clock cycles.

For example, the specified maximum leakage current for most I/O pins of today's FPGA pins is 10 µA, and their capacitance is 5 pF, so for an I/O pin that is only connected to an empty socket with 5 pF capacitance, the rise time to a 1.65 V threshold is 10 pF × 1.65 V / 10 µA = 1.65 µs. The minimum tolerable TCK clock period is 1.65 µs / 2.5 cycles = 0.66 µs, corresponding to 1.5 MHz. This is a reasonable constraint since typical TCK frequencies for board testing are 0.5~5 MHz. However, to detect the resistance of a pull-up (without using 1149.4), requires a much higher TCK frequency. For example, to detect the difference in rise time (to an input's threshold voltage) for 1 kΩ versus 10 kΩ, would require a TCK frequency of 30 MHz, which is too fast for most boards and board testers.

3.7 Detecting Shorts on High-Drive Outputs

Detecting very low impedance, short circuits between high-drive outputs and power rails is easily achieved with unpowered ICT. For power-on ICT, the time interval between applying power and then detecting via boundary scan that a high-drive output is shorted, becomes crucial to preventing damage or "wounding".

This time interval is usually not under control of the test engineer because it is mostly a function of how a board must be powered up. For example, a board may have several supply voltages that need to be applied in a particular sequence with time between them for stabilization. Many board designs regulate the supplied voltages to obtain other voltages, in a particular sequence. These regulators must be tested first for accuracy and stability, since any defects in the power supplies will likely cause other functional fails. The time needed for power sequence, stabilization and verification can be a few seconds. After all this is done, Boundary-Scan test can begin.

Many ICs are designed to implement the 1149.1 HIGHZ instruction, which tri-states all outputs so that bed-of-nails ICT can over-drive outputs, but this may be seconds too late. Boards that are to be tested only with power-on ICT may need to be designed so that any high-drive outputs are immediately tri-state on power-up – some ICs are designed this way. Even with that precaution, the time between enabling these drivers via boundary scan, scanning out to check their actual logic level, and then disabling them requires a full circulation of the boundary scan register, an interval of at least a millisecond. If this is too long, one way to detect these shorts [13] uses positive feedback from the pad to overdrive the driver's input signal, which requires changes to both the boundary scan and pad cells. Another way [12] is to de-assert the global enable signal during the HIGHZ instruction for one TCK cycle in the CaptureDR state, which does not require changes to the boundary scan or pad cells but is not possible within a standard Boundary-Scan TAP controller.

4 Details of the Proposed Boundary Scan

This section briefly describes each proposed requirement or recommendation being considered for the P1149.8.1 boundary scan standard, and then explains how each of the problems identified in the previous section is solved.

In this proposed standard, the term "selective toggle" means using the Boundary Scan infrastructure to deliver an AC stimulus to a single pin or a small number of pins of an IC, while all other signal pins (except TAP pins) "hold" their signals at constant logic 0 or 1.

4.1 Not all signal pins

While it is preferable that *all* signal pins have selective toggle (ST) capability, since any pin may be connected to non-boundary scan devices or empty connectors, it may be impractical for some pins, such as sensitive analog and high-speed differential inputs (recall they are labeled as "analog" in 1149.1). Differential inputs that are connected to an empty socket are common and are a target of this standard, but ultimately IC buyers can choose ICs based on whether the pins that will connect to empty sockets in their application have ST capability.

Pins that have ST circuitry are designated as "ST pins" to distinguish them from other signal pins (which should be as few as possible).

4.2 Toggle during Run-Test/Idle

Since we wish 1149.8.1 to be compliant with 1149.1 and other standards based on 1149.1, especially so that all existing 1149.1 tests will work for boards containing a mixture of ICs with and without ST, the toggling must take place during the Run-Test-Idle state, just as it does for 1149.6. As the TAP controller exits from Run-Test/Idle state, any ST pin toggling ceases and reverts to the logic value it would have in 1149.1 EXTEST mode. Toggling only occurs when the SELECTIVE_TOGGLE instruction is active. Non-toggling ST pins are required to behave as though the EXTEST instruction were active.

4.3 All ST pins can drive out during ST mode

Since the primary objective of this new standard is to stimulate selected pins with an AC waveform, while all other pins are driven to logic 1 or logic 0, it becomes mandatory that all ST pins be able to drive out a signal, including pins whose function mode requires input-only behavior. Requiring input-only pins to be bi-directional is a significant departure from 1149.1, but is common practice in the industry. Bi-directional pins permit ICs to be tested in IC production without contacting the pins, which facilitates reduced pin-count (RPC) testing and multi-site testing, both of which significantly reduce the cost of IC testing.

The primary test capability that bi-directionality adds to an input-only pin, at the board-level, is the ability to test for shorts and opens in the board wire connected to that pin even if no other IC is connected to that wire. It also permits testing of inputs that are connected to tri-stated outputs of non-boundary-scan ICs. Of course, once bi-directionality has been added for 1149.8.1, it is available for 1149.1-style testing too.

Input-only pins could be toggled via on-chip pull-up or termination resistors, which might permit an existing input-only pad cell to be used for an ST pin, or it might permit toggling sensitive analog or differential input pins without adding parasitic capacitance or ESD-sensitive elements. See the later section on required output drive.

4.4 Update latch holds non-toggling pin value

To be consistent with 1149.1, the logic value at the output of each boundary scan cell's Update latch must be the logic value that a non-toggling output pin would drive in EXTEST mode. The logic value is shifted in using the normal scan-in and update sequence while a PRELOAD or EXTEST instruction is active.

4.5 Logic 1 in BSR selects toggling pins

To select which pins are to toggle and which are to hold a constant logic value, the following scheme has been chosen. While the SELECTIVE_TOGGLE instruction is active, logic 0 values are shifted into all boundary-scan register cells, except for cells that control internal logic, output enables, and ST pins that are to toggle. The scheme exploits (and requires) the existence of two latches in a typical boundary scan cell: the shift register, Q_1, and the UpdateDR-clocked latch, Q_2, as shown in the example boundary scan cell in Figure 4a.

To enable the UpdateDR latch and the shift register latch to hold different logic values, the UpdateDR clock pulse can be suppressed to prevent it occurring immediately after shifting in all logic values for the boundary scan cell design of Figure 4b, or the UpdateDR clock can be combined with the AC Signal for the boundary scan cell design of Figure 4c. After updating the "hold" logic values in PRELOAD or EXTEST, the logic values shifted in during SELECTIVE_TOGGLE will be mostly logic 0s, with pin toggling enabled by a logic 1.

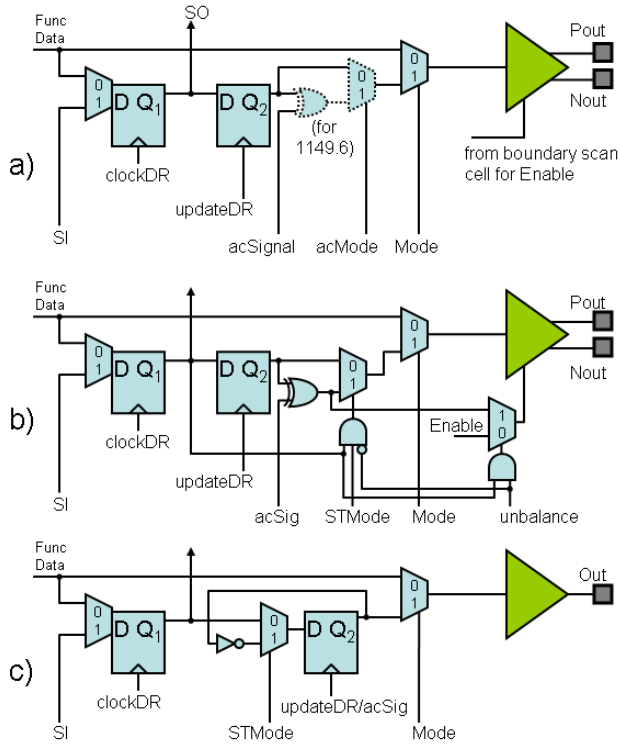

Figure 4: Example P1149.8.1 Boundary Scan outputs (and example 1149.6 output cell for comparison)

4.6 Global toggling signal

From 1149.6 experience, we know that it is practical to distribute a toggling signal, and if an IC already implements 1149.6 (or not), that same AC signal can be reused. Figures 4b and 4c show two different global signal schemes.

However, the frequency needs to be considerably different than for 1149.6 – that standard requires generation of a frequency that is always half that of TCK, or a single pulse whose width is any number of TCK cycles in duration (i.e. for the duration of the Run-Test/Idle state). Nevertheless, it is relatively easy to design a boundary scan cell that would be compliant with 1149.6 and P1149.8.1.

A toggling frequency as low as 8 kHz is needed to be consistent with present capacitive-sense board tests (to minimize noise bandwidth), yet we must accommodate TCK frequencies ranging from perhaps 50 kHz to at least 10 MHz. Higher TCK frequencies minimize the time to shift out the BSR data values. To be consistent with 1149.1 requirements, all toggling edges must be synchronous to TCK falling edges. Therefore, a programmable frequency divider is necessary, with TCK as its input clock.

4.7 Global ST-mode signal

To control whether pins respond to the toggling signal, it is necessary to have a new mode control signal, as is provided for 1149.1, 1149.4, and 1149.6 standards. It is possible to combine these mode signals at each boundary scan cell to permit additional modes.

4.8 Output drive

As stated earlier, capacitive-detection of shorts between equal strength drivers is ineffective, and Boundary-Scan based detection with a mid-rail threshold voltage has problems too. To detect these shorts, the Working Group considered a novel approach: ensure that non-toggling ST pins have weaker drive in test mode. For this approach, toggling pins would use their function-mode drive strength, but ST pins that are not toggling would deliver a logic 0 or logic 1 with "weak" drive, e.g., between 0.1 and 0.001 mA, or 10~1000 kΩ. Unfortunately, simulation analysis revealed that the capacitively-sensed signal from a toggling signal surrounded by weakly driven DC signals was insufficiently distinguishable from shorted signals because even 1 pF of AC-coupling between signals was sufficient to overdrive the weakly driven signals (and most existing on-chip pull-ups are 30~100 kΩ, so requiring lower resistance than this would require re-design and more power). Therefore, the approach suggested in 1149.4 – using a non-mid-rail threshold for detection – was adopted, though only as a recommendation since it does not guarantee detection: shorted drivers may naturally have unequal strength and thus produce a non-mid-rail voltage.

A capacitively-coupled sensor detects a signal proportional to the capacitance between the stimulus wire and the sensor plate, and to the derivative of the stimulus signal waveform. Therefore, it is proportional to both the slew rate and the amplitude of the signal. To be consistent with a 1 V amplitude, 8 kHz sine wave, a minimum slew rate of 1 V/ms is sufficient. However, since this is easily met using digital waveforms, and to permit more noise margin, the rule adopted is that the 10~90% rise time must be less than 5 microseconds, and the recommended minimal acceptable amplitude is 400 mV peak-to-peak, for each wire. This corresponds to a minimum acceptable slew rate of 64 V/ms.

Loading conditions affect output signal swing: resistive load can reduce the amplitude and capacitive load can reduce the slew rate. For most output drivers, the logic swing and slew rate will easily exceed the toggle requirements, but for input pins the output toggle drive must be documented. Specifically, the IC's datasheet must state the smallest resistance to either logic rail, and the largest capacitance, for which the IC can deliver the required toggle amplitude and slew rate. Furthermore, the pin's driver must be able to deliver a steady-state logic 0 or 1 into the minimum resistance with sufficient voltage margin relative to the pin's own input threshold voltage for the logic level to be detected.

For example, for a 3.3 volt IC's pin that is input-only in function mode, and whose datasheet recommends a minimum pull-up resistance of 10 kΩ, the output drive at mid-rail in toggle mode must be greater than 0.215 mA to achieve 0.5 V margin relative to a 1.65 V threshold. If its threshold is set instead at a non-mid-rail threshold, such as 1 V, then the output drive must exceed 0.28 mA to achieve the same margin (this drive could be delivered via a 3.5 kΩ on-chip switchable, pull-down resistance).

Specifying a high value for the minimum off-chip resistance, to ease on-chip drive requirements for an input-only pin, can reduce fault coverage. If the drive current is too weak, then any AC-coupling between wires driven by steady-state ST pins and a wire driven by a toggling pin (as mentioned earlier in this section) can prevent reliable detection of shorts between these wires.

4.9 Differential outputs

The requirement for a differential output is simply that the sum of the pair of signals should toggle with the recommended minimum amplitude specified earlier: 400 mV peak-to-peak. In other words, the differential signal must become quite unbalanced. Of course, it is useless to specify this if there are no practical ways to achieve it. The following paragraphs describe one simple technique for generating non-complementary output signal pairs, suitable for many existing, well-designed differential drivers, including multi-gigahertz differential drivers.

First consider a common-mode logic (CML) driver as shown in Figure 5. This type of driver design is used in SerDes outputs operating up to 10 Gb/s [1] and higher [6]. When enabled, both output signals swing between

two voltages: logic 1 corresponds to minimum current flowing through the non-inverting output's source resistance, and logic 0 corresponds to maximum current flowing through that resistance; the opposite is true for the inverting output current. When disabled, the current source becomes an open circuit and no current flows through either source resistance. If the input logic value is held at a constant logic 1, and only the Enable is toggled, the non-inverting output's source resistance current toggles between zero and minimum, whereas the inverting output's current toggles between zero and maximum. Hence, the output signals are no longer complementary, as can be seen in the simulation waveforms of Figure 5 for an unterminated output.

A similar effect can be achieved by toggling the common source current between maximum and minimum (instead of enabling and disabling it).

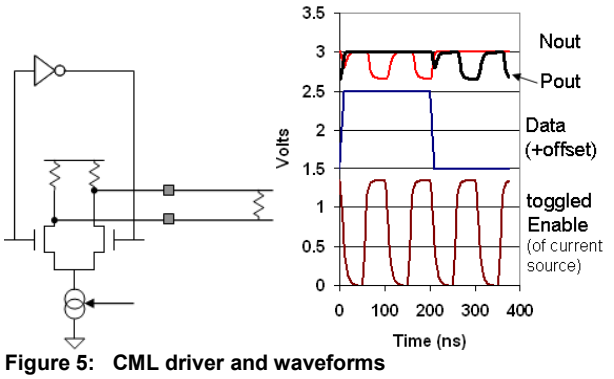

Figure 5: CML driver and waveforms

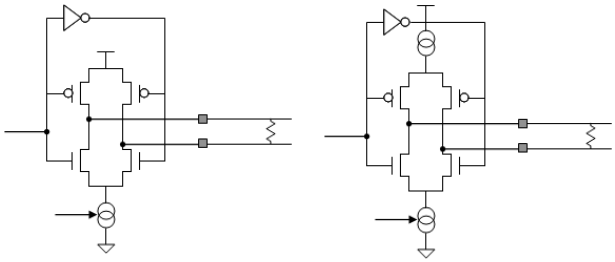

Figure 6: a) Simple LVDS driver b) Balanced LVDS driver

Next consider a low voltage differential signaling (LVDS) driver that uses a single current source, as shown in Figure 6a [7]. The results for this type of driver are very similar to those for CML, if the driver is disabled by turning off its current source – turning off all the output transistors will not achieve the desired effect since the output will become high impedance.

Lastly consider a balanced LVDS driver like the one shown in Figure 6b, and used in many ICs [14]. For this case, regardless of the constant-current value, the two outputs always swing symmetrically around their V_{COM}, and their sum is always constant – not what we want. In reality, for maximum current flowing, one of the output transistors will saturate more than the other, which will produce non-complementary behaviour. To ensure a more predictable result, it is necessary to make one of the two current sources less ideal – its DC bias voltage can be increased so that the current source transistor goes out of saturation and into its linear region of operation where it behaves more as a resistance than a current source, i.e., it is fully on.

For all cases, the data input value determines which side of the differential pair toggles with greater amplitude, which helps diagnosis of the open connection.

As will be seen later in this paper, it is useful to toggle a differential driver in its normal balanced complementary mode in addition to its unbalanced mode. To permit this, an additional control is needed and a simple way to provide this is with a balance/unbalance control bit in the BSR. The advantage relative to using an instruction bit is that the control bit may be placed closer to the differential outputs, and there could be multiple bits (all having the same value) to minimize global routing.

When a short exists between the outputs of a differential driver, while it is toggled in unbalanced mode, the resulting signal might not be detectable. For an LVDS driver whose termination resistor is shorted, the resulting sum is only reduced by 10~20%. For a CML driver, the amplitude is not reduced at all. To detect such shorts, a low-speed differential comparator that has hysteresis could be used as a self-monitor. For logic level changes in EXTEST mode, a short-circuit would prevent changes in the logic level of the hysteretic comparator. The amount of hysteresis must distinguish between a few ohms and a normal termination resistance. Since a self-monitor may have significant performance impact for high-speed differential outputs, the standard is likely to recommend this instead of requiring it.

It is interesting to observe that if the data and enable inputs of a differential driver are toggled simultaneously, the two output waveforms overlap each other in voltage, whereas they do not overlap if only the enable is toggled. This can be useful for detecting whether the path to the differential receiver is DC-coupled or AC-coupled, since two signals that do not overlap will produce a steady-state logic value at a DC-coupled differential receiver. This feature might eventually be useful for testing like 1149.6 but without requiring changes to the receiver.

5 Other capabilities of the infrastructure

Although not discussed by the Working Group, simulations show that when the enable to a single-ended output driver is toggled, the value of a pull-up resistance or the presence of a short circuit can be detected

5.1 Detecting pull-up resistance

For the case where a pull-up resistor exists on the output wire, and the driver's BSR data value is logic 0, toggling the enable will result in toggling the wire's signal and its amplitude is proportional to the pull-up resistance, as shown in Figure 7 for three different pull-up values and a 5 MHz TCK (recall that updates occur on falling edges only). The signal detected by a capacitively-coupled sensor (in effect, a high-pass filter) is signi-

ficantly different for the three pull-up resistances.

It is feasible to perform a capture outside of the CaptureDR state [11], though this is not prescribed by 1149.1. Specifically, capturing in the Run-Test/Idle state, a half TCK cycle after the enable signal tri-states the driver, as shown in Figure 8, gives better time resolution for detecting a pull-up's resistance by monitoring the captured logic value. Recall that 1149.1 limited us to 2½ TCK cycles minimum time between Update and Capture, so a capture during Run-Test/Idle reduces the minimum interval by a factor of five. To retain the captured logic value for shifting out, the normal capture in CaptureDR state must be suppressed, as described in [10].

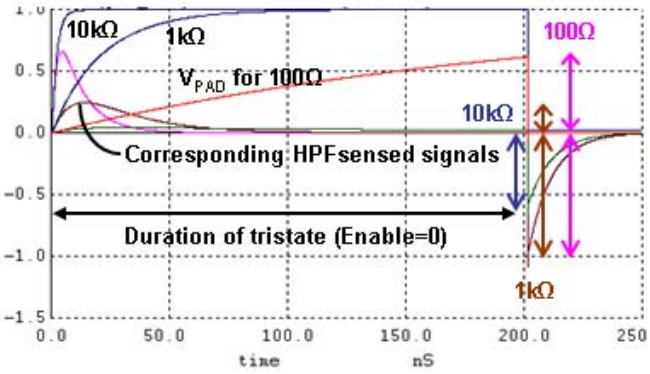

Figure 7: Waveforms for tri-stated output with pull-ups of 100, 1k and 10 kΩ, and the corresponding signals detected via 1.6 MHz high-pass filter (1 pF + 10 kΩ).

The number of TCK cycles in Run-Test/Idle can be increased to detect the rise times at 0.5, or 1.5, or 2.5, etc. cycles after the driver is tri-stated. The frequency of TCK can be increased (or varied) to permit much finer timing resolution.

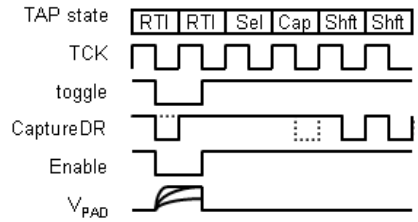

Figure 8: Waveforms for testing pull-up resistance.

5.2 Detecting shorts on high-drive outputs

If the Enable for a high-drive output is toggled for only one or two TCK cycles instead of being steady-state on for many cycles as discussed earlier, and the pad's logic level is simultaneously captured, then the impact of a short circuit may be reduced greatly and will be independent of the number of bits in the BSR. The Capture during Run-Test/Idle technique just described for pull-up detection can be used.

6 Results

Typical capacitances between each pin in an empty socket and a sensor plate are 20~200 fF. For open connections, this capacitance is typically less than 10 fF. If there are many pins, the signal is effectively divided by the ratio of one pin's capacitance to the sum of all other pin capacitances. Note, however, that the capacitance between a sense plate and any two pins that are closer together than their distance from the sense plate can be similar to the capacitance of only one of those pins, due to field fringing effects.

The ideal unbalanced differential signal to maximize the capacitive-sensor detected signal comprises a full-swing signal on one wire and a zero amplitude signal on the other wire, which is very similar to the existing TestJet approach described in the Introduction where a single-ended stimulus is applied to one wire and all others are guarded. During unbalanced mode, the two signals may be in-phase or complementary, but only if the sum of the signals toggles at least 400 mV.

A test set would comprise two tests for differential pairs. If sufficient signal is detected by the sense plate during *unbalanced* differential toggling (without caring how it is unbalanced or which signal swing is larger), it proves that the differential driver is connected to one or both of the two pins: *Signal detected = Pass*. If sufficient signal is detected during *balanced* differential toggling, it proves that some fault exists (and the detected phase may permit diagnosis): *Signal not detected = Pass, if balanced test passed too*.

6.1 Coverage increases and/or access reductions

The list of potential faults was considered with reference to the schematic shown in Figure 9, and is listed in Table 1 along with detection results based on Spice simulations.

For boards with many connectors, which is increasingly common for applications as diverse as computer motherboards and cell phones, adding P1149.8.1 to the main ICs is expected to improve coverage by more than 25% [2].

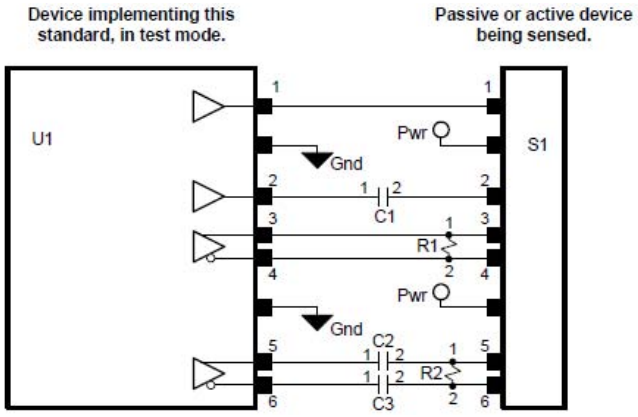

Figure 9: Board connections between a Boundary-Scan device U1 and capacitively sensed device S1.

ID	Defect site	Detection result
	Single-ended signals	
1	U1 pin 1 short to DC	Detected by CS, SM
2	U1 pin 1 open	Detected by CS
3	C1 pin 1 open	Detected by CS
4	U1 pin 1 short to pin 2	Detected by CS, SM
5	S1 pin 2 short to DC	Detected by CS, SM*
6	C1 pin 1 short to pin 2	*Not detected*
	Differential signals	
7	U1 pins 3 and 4 open	Detected by CS, unbalanced
8	U1 pin 3 short to pin 4	Detected by CS for LVDS *but not for CML*, SM**
9	R1 pin 1 open	Detected by CS, excessive amplitude when unbalanced
10	U1 pin 3 open	Detected by CS, balanced
11	U1 pin 5 open	Detected by CS, balanced
12	S1 pin 4 short to pin 5	Detected by CS, balanced
13	S1 pin 5 open	Detected by CS, balanced
14	R2 pin 1 short to pin 2	*Not detected*

SM= self monitor on outputs (bidir) of IC U1.
CS = capacitive sense at socket S1, and whether fault was detected while differential signals were balanced or unbalanced
* for small C1, and/or long TCK period.
** for large R1 and/or small SM hysteresis voltage.

Table 1: List of targeted board faults, and detection results

For one analysis of board defects detected during six months of production testing by Cisco, some interesting observations were presented to the P1149.8.1 Working Group by Steve Butkovic regarding their TestJet-detected failures:

- 54% of failures were open pins to leaded ICs (mostly due to coplanarity issues) – capacitive sensing for BGAs is inherently poor, so boundary scan is used more widely for BGAs;
- 28% were false (no defect found) due to poor signal-to-noise level, and half of all false fails were TestJet tests;
- 10% were missing components (mostly surface mount resistors and capacitors);
- 5% were open pins to leaded connectors;
- Shorts detected by previous tests (e.g., boundary scan).

From these observations, we can conclude: opens to ICs occur much more often than for connectors and the opens are not detected by boundary scan; a good signal-to-noise ratio is needed to prevent false fails; and missing resistors and capacitors are a significant portion of failures.

Another study showed two potential benefits of this technology. First, defect coverage can be improved without increasing bed-of-nail access. This is important when changes in technology decrease accessibility. Second, test coverage can be preserved when access is deliberately decreased to reduce cost. Eliminating nails is desirable since current industry quotes for fixture costs can exceed $15.00 per node nail, and because clustering many nails in a small region of a board can induce board flexure that can damage solder ball joints.

Figure 10a shows some coverage provided by TestJet, dependent on nail access given by test pads. This includes all three indicated solder connections along the path, plus the resistor presence is also tested (but not its resistance). The cost of this is the area needed to provide the test pad and the cost of fixture nail(s).

Figure 10b shows coverage that can be obtained using Boundary-Scan stimulus. No test pads or nails are required and an extra solder joint is tested under the Boundary-Scan device.

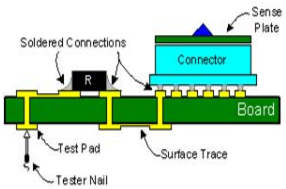

Figure 10a: TestJet coverage for solder opens along a path

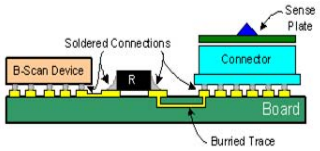

Figure 11b: Coverage using Boundary-Scan based stimulus

In some real production boards recently outfitted with a precursor to 1149.8.1 technology, the following was observed:

- On an ultra-mobile personal computer board (i.e., very dense layout), 298 previously inaccessible nodes could be tested for opens with capacitive sensing (and no other access), with two or more pins per node becoming testable;
- On a multiprocessor CPU board that once required 264 test pads with 20 mil diameters on 39 mil spacing (i.e., costly fixturing) in a very dense area under the processors (i.e., risk of cracked ball joints), 176 of these could be eliminated with no coverage loss.

This precursor technology was not able to take full advantage of all the 1149.8.1 concepts being developed, so these results, though nice, were not as good as we would expect with the full 1149.8.1 feature set. Further, the precursor capability used here is non-standard and slated for obsolescence.

7 Limitations

One limitation is undetected fault types. For example, series capacitors are much larger than the typical 0.5 pF capacitance to a sensor plate, so it is difficult to detect whether a series capacitance to an empty socket has a short circuit across it. One possible solution is to add a high resistance pull-up, on the vacant connector side of the capacitor, so that the pull-up could be detected as described in section 5.1. But differential signals that are AC-coupled and then terminated are presently undetectable, as shown in Table 1.

LVDS drivers must be toggled between low and high current, not between enabled and disabled. Also, a balanced LVDS driver must be re-designed so that in test mode, the source current and sink currents are not balanced – fortunately this only affects DC voltages in the driver.

It remains to be seen whether the pad I/O cell rules are acceptable to IC designers. The most significant new requirement is that differential inputs (that are declared as ST pins) must be bidirectional in test mode, however toggling via on-chip termination resistors might be an acceptable implementation that does not impact performance.

8 Discussion

In one analysis presented to the P1149.8.1 Working Group, Thai-Minh Nguyen of LSI Logic stated that adding 1149.6 capability to their high-speed differential receivers (excluding boundary scan) increased cell area by 30% and required 6~8 months design time (for all receivers); the bandwidth of the I/Os was not significantly affected (but jitter was). They expected 1149.8.1 to have a similar area, design time, and performance impact. This study shows the importance of having 1149.8.1 specifications that do not require I/O cell design changes.

9 Acknowledgements

The authors would like to recognize the contributions of the P1149.8.1 Working group led by Jeff Burgess, and specific work by JJ Grealish, Scott Bowdon, Sophocles Metsis, Thai-Minh Nguyen and Tony Suto.

The P1149.8.1 web site contains contributors, minutes, and presentations from all Working Group discussions: http://grouper.ieee.org/groups/1149/atoggle/

10 Conclusions

This paper presented reasons why extensions to IEEE 1149.1, 1149.4, and 1149.6 are needed to test boards that contain empty sockets and non-Boundary-Scan components: achieving high test coverage for these boards typically requires bed-of-nails ICT and capacitive sensing but this approach is being strained as boards become increasingly dense.

The specific 1149.x limitations identified and addressed are related to connections between sockets and input-only pins, delivering a specific-frequency stimulus to one or a few selected pins despite long boundary registers, capacitively sensing differential signals, and reliably detecting shorts between equal drive outputs.

The proposed P1149.8.1 standard is still evolving, but thus far has solved the targeted limitations by requiring:

- additional combinational logic in each boundary scan cell to permit toggling any selected pin at the TCK frequency divided by any integer regardless of BSR length;
- input-only pins to be bi-directional, but permitting pins to be toggled via pull-ups;
- EXTEST-equivalent output drive for non-toggling pins, and non-mid-rail threshold voltages for self-monitors, so that shorts between them and equal-drive toggling pin(s) would be detected;
- differential output signals to be toggled in a non-complementary way, for example by toggling their enable signals, which was shown to need no design changes in many high-speed differential drivers.

The Working Group is presently soliciting feedback from industry to determine the practicality and effectiveness of the proposed solutions and specification. So far, for a board with Boundary Scan devices compliant with the P1149.8.1, it seems it will be practical to test wire paths connected to them that travel to only non-Boundary-Scan devices, passive components, or empty sockets – from a Boundary Scan perspective these paths are bridges to nowhere.

References

[1] J. Bulzacchelli *et al*, "A 10-Gb/s 5-Tap DFE/4-Tap FFE Transceiver in 90-nm CMOS Technology", *IEEE J. Solid-State Circuits*, vol. 41, pp.2885-2900, Dec 2006

[2] D. F. Dubberke, J. J. Grealish, W. I. van Dick, "Solving In-Circuit Defect Coverage Holes with a Novel Boundary-Scan Application", *Proceedings Int. Test Conf.*, paper 11.2, Oct 2008

[3] IEEE Std. 1149.1-2001, "IEEE Standard Test Access Port and Boundary-Scan Architecture"

[4] IEEE Std. 1149.4-1999, "IEEE Standard for a Mixed Signal Test Bus", 2000

[5] IEEE Std. 1149.6-2003, "IEEE Standard for Boundary-Scan Testing of Advanced Digital Networks",

[6] D. Kehrer *et al*, "40-Gb/s 2:1 Multiplexer and 1:2 Demultiplexer in 120-nm Standard CMOS", *IEEE J. Solid-State Circuits* , vol. 38, November 2003

[7] *LVDS Owner's Manual*, National Semiconductor, 2000, page 1

[8] D. Norrgard, K.P. Parker, "Augmenting Boundary-Scan Tests for Enhanced Defect Coverage", *Proceedings Int. Test Conf.*, paper 11.3, Oct 2008

[9] S. Sunter, B. Nadeau-Dostie, "Complete, Contactless I/O Testing – Reaching the Boundary in Minimizing Digital IC Testing Cost", *Proceedings Int. Test Conf.*, pp. 446-455, Oct 2002

[10] S. Sunter, C. McDonald, G. Danialy, "Contactless Digital testing of IC Pin Leakage Currents" *Proceedings Int. Test Conf.*, pp 204-210, Oct 2001

[11] S. Sunter, "Method and circuit for testing DC parameters of circuit input and output nodes", *US patent 6586921*, 2003

[12] S. Sunter, P. Gauthier, B. Nadeau-Dostie, "Boundary scan with strobed pad driver enable", *US patent 7219282*, 2007

[13] L. Whetsel, "Bi-directional scan design with memory and latching circuitry", *US patent 5706296*, 1998

[14] A. Tajalli, Y. Leblebici, "A Slew Controlled LVDS Output Driver Circuit in 0.18 um CMOS Technology", *IEEE J. Solid-State Circuits*, vol.44, pp. 538-48, Feb. 2009

[15] S. Sunter, "Circuit and method for determining the location of defect in a circuit", *US patent 6717415*, 2004

[16] Minutes of P1149.4 Working Group, March 1997 http://grouper.ieee.org/groups/1149/4/min0397.html