Measure specific parameters of an IEEE 1394 interface with Time Domain Reflectometry.

Michael J. Resso, Hewlett-Packard and Michael Lee, Zayante

Evaluating Signal Integrity of IEEE 1394 Physical Layer with Time Domain Reflectometry, Part 2

The data rates of today's high-speed digital systems are rapidly increasing. New serial data communication standards such as IEEE 1394 (FireWire) are being developed that push the gigabit per second limit and beyond. Cables, interconnects, PCBs, and silicon behave much differently as rise times in the sub-nanosecond regime become commonplace. Extreme challenges are presented to the digital designer as various physical layer components create signal integrity problems under these extreme conditions. Controlling the impedance environment of the FireWire physical layer is now paramount to assuring that a design meets functional requirements. By using Time Domain Reflectometry (TDR), the maximum amplitude of reflection can be measured directly to certify compliance with the 1394 standard.

The basics of TDR were presented in the first part of this article (*Insight*, Volume 4, Issue 2, 1999, page 22). This part describes methods of using TDR to measure the input impedance, input capacitance, and jitter of an IEEE 1394 interface.

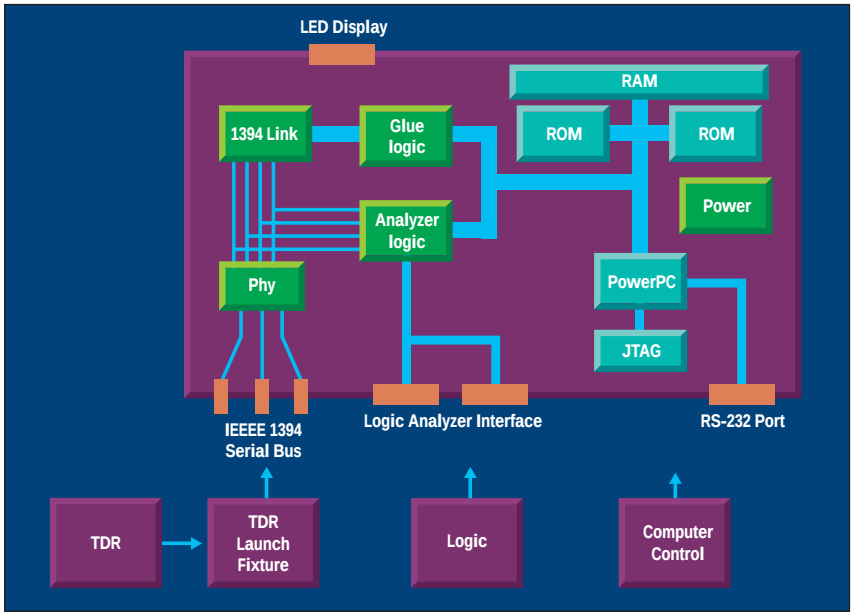

Figure 1. The basic components of the IEEE 1394 TDR test setup.

Focus on the Physical Layer

The IEEE 1394 compliance measurements described here focus on the physical layer silicon chip, referred to as the PHY chip, and the transceiver board into which it is mounted. The 1394 transceiver board has two differential signal pairs called TPA and TPB. To comply with the standard, both TPA and TPB must conform to the specification while in transmit mode and receive mode. They must also comply with an input capacitance specification. This ensures that a particular device does not load down the bus excessively.

The IEEE 1394-1995 standard is written from the perspective of the cable looking into the transceiver. Two different input impedance measurements are made: one when the transceiver is in transmit mode and another when it is in receive mode. The transceiver is in transmit mode when it is in the process of driving a 1b or 0b. It is in receive mode when it is not driving and is in an idle state. Signal

pairs TPA and TPB are very different and therefore both need to be tested. The TPA pair on the transceiver drives TPBias and has a small current source called Icd. The TPB pair has only the termination resistors and a 270-pF capacitor in parallel to ground. All of the external termination components on both TPA and TPB exist common mode only, an important design feature that enables the balanced differential circuit topology. This also allows the method of measuring a differential circuit using the single-ended TDR test techniques discussed later.

Standard Dictates the Input Impedance

The 1394 standard dictates an input impedance specification for both TPA and TPB. It states that for all devices the differential input impedance should be between 109 and 111 Ω when in receive mode and between 105 and 111 Ω when in transmit mode. Two methods are proposed by the standard to verify the input impedance compliance of TPA and TPB. Both methods use Differential TDR to accomplish this measurement. With Differential TDR, the TDR step generators inside the oscilloscope launch two fast edges (one into each leg of the twisted pair). The two edges are equal in magnitude and opposite in polarity, thereby creating a differential stimulus as normally seen by the transceiver data-transport pairs. The oscilloscope is set up to receive the differential response to the differential stimulus. It is also possible to configure the stimulus and/or response as common mode instead of differential. This equipment flexibility results in a test matrix that can lead to valuable insights into the Device Under Test (DUT).

For the first method of measuring input impedance, the oscilloscope is in a mode in which it measures differential impedance directly by using its impedance extraction algorithm firmware. The vertical units are changed from millivolts to ohms, and the waveform markers are placed at the appropriate position on the DUT. The differential impedance is then read directly from the marker table on the oscilloscope display. Two single-ended lines with a characteristic impedance of 50 Ω should yield a differential impedance of 100 Ω, assuming a perfectly balanced differential system. However, if the data transport lines of the FireWire transceiver are not perfectly balanced, possibly due to impedance discontinuties on one line and not the other, the differential impedance would not be the expected 100 Ω. Two ways imbalance can occur in the FireWire cable media are if the individual legs are not symmetrically twisted or if one line has a larger gauge wire. Three ways imbalance can occur on a printed circuit board are if there is a ground-plane discontinuity under one leg and not the other, if one microstrip leg is thicker than the other, or if the dielectric material in the board is not homogeneous.

We chose the second method described in the standard to measure the input impedance of the FireWire transceiver. In this case, the requirement is measured in terms of the maximum amplitude of reflection from the DUT. When the vertical units are changed from millivolts to percent reflection on the oscilloscope, the amplitude of the reflection coefficient (Rho), is measured. If the waveform markers are put at the proper position, the measurement can be read numerically from the oscilloscope display. The 1394 standard states that for a transceiver capable of S400 speeds, the amplitude of reflection must be less than 21.4 percent; for an S200-capable transceiver, it must be less than 37.4 percent; and for an S100-capable transceiver, it must be less than 48.0 percent.

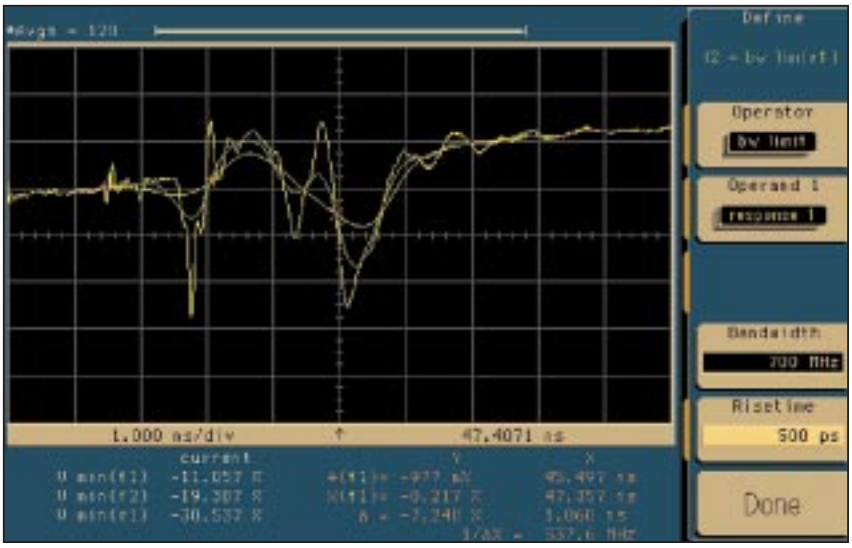

Figure 2. The measurements of input impedance of PHY #1 TPB Port were obtained using the amplitude of reflection method.

Rise Time Measurement

The 1394-1995 standard dictates the worst-case rise time for a data pulse to be 1 nanosecond. In order to measure the 1394 transceiver's response to this worst-case rise time, a novel variable-rise-time test methodology was implemented. A bandwidth-limiting math function allows the transceiver's response to a faster or slower edge to be simulated in software by the oscilloscope. This method has proven very useful because it indicates how much a specific discontinuity will contribute to overall signal integrity at real-world speeds. For example, if an S100 FireWire board is running at 100 Mb/s, it is much more tolerant of large impedance discontinuities than an S400 board running at 400 Mb/s. Increasing the TDR step will emphasize the areas of inherently lower signal integrity within the board. This is how high-speed connectors are characterized in the time domain. A gigabit FireWire copper interconnect would need to exhibit a low reflection coefficient when a very fast rise-time edge is launched into it. Even if the whole package passes the 1394 standard, its signal integrity can be improved by pinpointing discontinuities using the fastest edge speed.

The measurement of input impedance (of PHY #1 TPB Port in this case) is shown in **Figure 2**. The measurements were obtained using the amplitude of reflection method. There are two simulated waveforms and one real waveform. The two simulated waveforms represent signals into the transceiver with two different rise times. These rise times are 1 ns and 500 ps, representing the 1394-1995 standard and 1394b standard, respectively. The third waveform is the actual 35-ps rise time step from the TDR system. Naturally, the fastest edge results in the largest reflection, so the waveform with the highest-amplitude reflections is the one with the 35-ps rise time. The 500-ps and 1-ns waveforms yield reflections with correspondingly smaller amplitudes.

Measuring Input Capacitance

A second compliance test is transceiver input capacitance. As with the input impedance specification, different-speed transceivers have different performance requirements. The maximum input capacitance for S400 is 4 pF; for S200, it is 7 pF; and for S100, it is 9 pF. For this measurement, we used a TDR step with a 200-mV amplitude and a 1-ns rise time, the fastest allowed by the 1394-1995 standard. Using this fastest allowable rise time again provides the worst-case scenario for reflections from impedance discontinuities: The faster the rise time, the larger the amplitude of the reflection. This test shows what structure in the physical layer causes the most problems for signal integrity. The input capacitance specification is written to address the complete FireWire termination network: not only the physical silicon chip, but also the cable, connector, printed circuit board, and termination devices.

Normalization is a TDR measurement technique that allows any test fixture to be electrically removed from the measurement by calibration. The test fixture could be a cable assembly, a TDR probe, an SMA launch board, or a simple connector. There is normally some bandwidth lost as a 35-ps edge

passes through a test fixture, so it is useful to remove the test fixture from the measurement. For this TDR measurement the test fixture was a differential launch adapter board with two SMA 3.5-mm female connectors on one side and a FireWire male connector on the other side. A 5-inch custom FireWire female-to-female cable was then connected to the FireWire transceiver. Special FireWire male terminations were constructed by mating an SMA 3.5-mm female launch to a FireWire male connector. This allowed the 50-Ω termination and short to be easily connected and removed to complete the required reference plane calibration. With this configuration, the reference plane was set right at the input to the transceiver, enabling repeatable, accurate TDR measurements.

Why should the reference plane be so close to the transceiver input? The nature of TDR measurements dictates that a given step voltage is launched into the DUT. When the reflection returns, it is displayed as a percentage of the original voltage amplitude. However, after the first discontinuity reflects some portion of the incident wave, there is now less then 100 percent of the incident wave travelling to the second impedance discontinuity. This means that the second, third, and fourth discontinuity measurements are less accurate than the first. In fact, if the initial impedance discontinuity is too large, the resulting TDR waveforms can only be used qualitatively, to determine if the discontinuity is inductive or capacitive. So setting the reference plane right in front of the input to the transceiver produces the most accurate results.

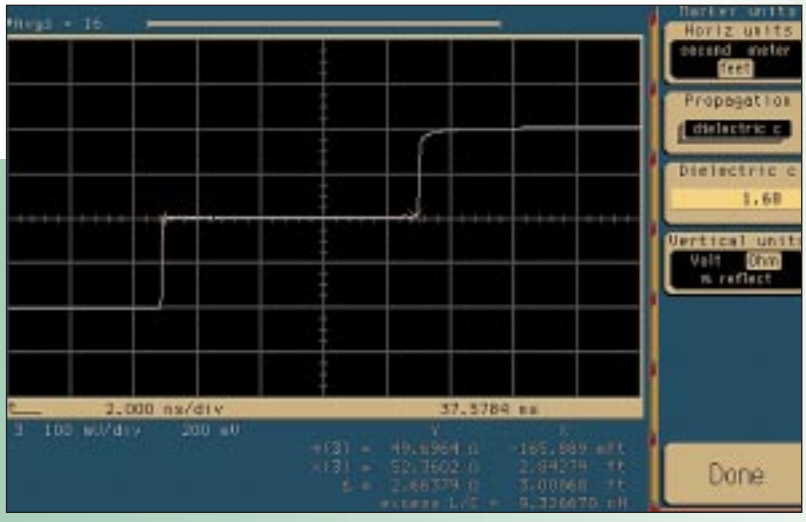

Scope Can Display Cable Length Directly

The 54750A Digitizing Oscilloscope with the 54754A Differential Time Domain Reflectometer plug-in module can easily measure the length of FireWire, or any other cable or trace. When the dielectric constant of the cable or circuit trace is set, the instrument can directly display the length to the end of the cable or to interesting TDR reflections.

Each leg of the differential pair of a FireWire cable can be measured separately using single-ended TDR to ensure that they are of equal length. To measure cable length, first calibrate the Time Domain Reflectometer and set the reference plane at the beginning of the cable. Special markers can then be placed on the TDR waveform at the reference plane. Next, scroll one of the two markers to the unterminated cable end, and input the known dielectric constant for the cable. You can read the cable length directly on the scope display.

The scope can also be used to measure the dielectric constant of a cable with a known length. Once the dielectric constant of a cable is known, then it can be used to measure the length of any of the same type of cable. To measure the dielectric constant, select the dielectric softkey and vary the dielectric-constant numeric value. When the delta marker function displays the length corresponding to the known cable length, the number being displayed in the dielectric-constant softkey menu is the correct dielectric constant.

For more information, check 1 on the reply card, or visit http://www.hp.com/info/insight3.

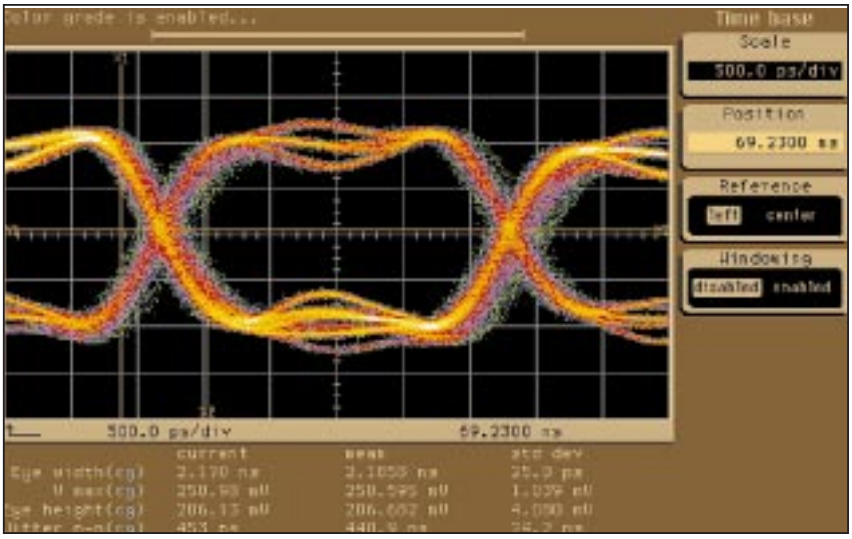

Figure 3. The digitizing oscilloscope performs automatic measurements and displays the jitter in the left-hand corner of the screen. This is the jitter of a failing S400 PHY.

Evaluating Jitter

The same TDR system used to measure input impedance and input capacitance can also be used to measure jitter, another performance parameter of high-speed digital circuits. Jitter is the deviation of the significant occurrences of a signal from their ideal positions in time. It is usually measured as a peak-to-peak or RMS value in units of nanoseconds or unit intervals. With the TDR step generators disabled, the oscilloscope is used as standard electrical input channels and measures jitter on an eye diagram. The 1394 standard states that the jitter transmitted between the two signal pairs, TPA and TPB, shall be no greater than 1.6-ns peak to peak for S100, 0.5-ns peak to peak for S200, and 0.3-ns peak to peak for S400. The jitter is measured from the rising and/or falling edge of TPA to the next (or as close as possible) rising and/or falling edge of TPB. The reason for looking at the next edge is that different clocks drive each of the 1394 devices. Each time data is transmitted from one node to the next, the data is synchronized to the local clocks. The longest data packet cannot exceed 83.3 µs, so there is no reason to measure jitter for longer than that. For 1394 this is called short-term jitter.

There are many reasons that jitter appears in the 1394 cable. The most common reason is that the Phase Locked Loop (PLL) design in the PHY chip is not designed to meet the 1394 transmit jitter requirements. The 1394 device runs off of a 49-MHz system clock that is based on a crystal input of 24.576 MHz. A PLL is commonly used to achieve this synchronization. If the PLL of one transceiver is bad, it could allow transmission of corrupt data and degrade the performance of the entire 1394 bus.

The 1394 system clock is embedded in the data-transport signal lines, TPA and TPB. By using exclusive-OR logic on these two differential signal pairs, a reliable clock can be extracted. Either TPA or TPB will change in one bit cell, but not both. This clock recovery scheme, called data-strobe encoding, gives better jitter budget than conventional clock/data recovery methods.

To measure 1394 transmit short-term jitter, the transceiver must be transmitting a changing data packet. This is accomplished by having the transceiver send a Write Quadlet Request broadcast packet addressed to a node that does not exist on the 1394 bus. (A quadlet of data is four bytes.) For example, in a two-node 1394 system, node 1 is the Root Node and node 0 is the Child Node. A command for a Write Quadlet Request transmitted to node 63 will never be acknowledged as received by either of the two nodes on the 1394 bus because it was not addressed to them. The absence of Acknowledge packet transmission is important when trying to measure transmit jitter. This protocol sets up the communication

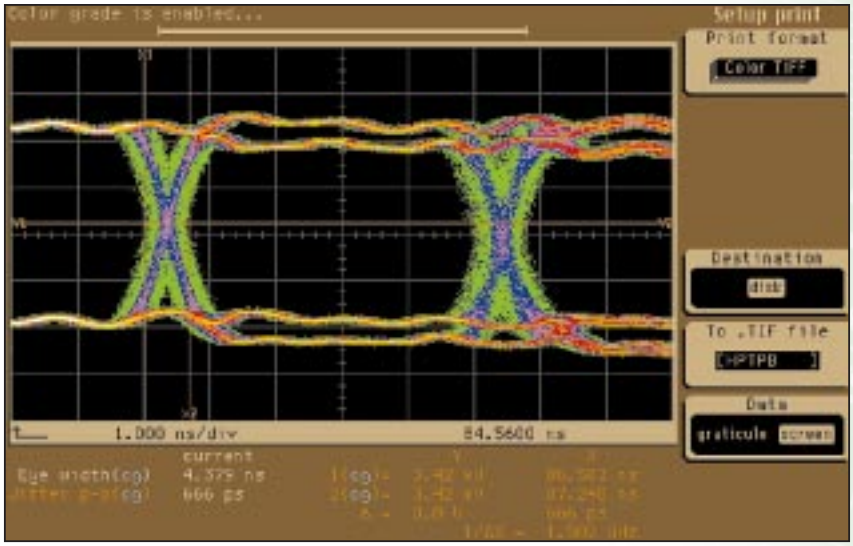

Figure 4. This is the jitter of a passing S400 PHY.

link between two nodes in such a way that a pseudo random binary sequence (PRBS) is transmitted. The result is a fully developed eye diagram. It is important to issue the packet from the Root Node, allowing it to transmit data with no request for the bus, thus making the data easier to trigger with less arbitration.

Along with this repeating Write Quadlet Request broadcast packet there are five quadlets of data. Within this data packet is a field called a transaction label field that increments 00h through 3Fh and back. It is used to generate an eye diagram on the digitizing oscilloscope that can then be analyzed. It is important to have an incrementing field because all possible logic-level transition combinations are needed to build a good eye diagram. The eye diagram is measured to get the peak-to-peak jitter measurement. Other critical measurements of the transmitted data can be made such as differential amplitude, bit cell time, overshoot, and rise and fall time for signal pairs TPA and TPB.

Once an eye diagram is obtained (**Figure 3 and Figure 4**), the measurement is a straightforward process. The digitizing oscilloscope performs automatic measurements and displays the jitter in the left-hand corner of the screen. Care must be taken to make sure that the oscilloscope is triggering on the right transition and the appropriate voltage level. Proper triggering is crucial because there is a 200-ns data prefix (which is not clocked data), and an improper trigger would erroneously include the 200 ns in the jitter measurement. Because the first logic transition of the first bit of data prefix is always of opposite polarity from the logic transition of the first bit of real data, the data prefix can be identified and excluded. When triggering on TPA and looking at TPB for jitter, triggering must be done on a rising edge. Conversely, when triggering on TPB and looking at TPA for jitter, triggering must be done on a falling edge. The best jitter measurement is obtained when triggering on one edge of clocked data and looking at the next edge of clocked data. A delay line was used to compensate for the 22-ns delay associated with the TDR oscilloscope used in the testing.

Signal integrity engineering is becoming a very important aspect of high-speed digital design. Impedance mismatch and the corresponding signal reflections within the physical layer can be identified using Time Domain Reflectometry. Only after these signal reflections are located can the signal integrity be evaluated, thus improving the reliability of high-speed data transmission.

For more information, check 1 on the reply card, or visit http://www.hp.com/info/insight3.