

8 Hints for Debugging and Validating High-Speed Buses

Application Note 1382-10

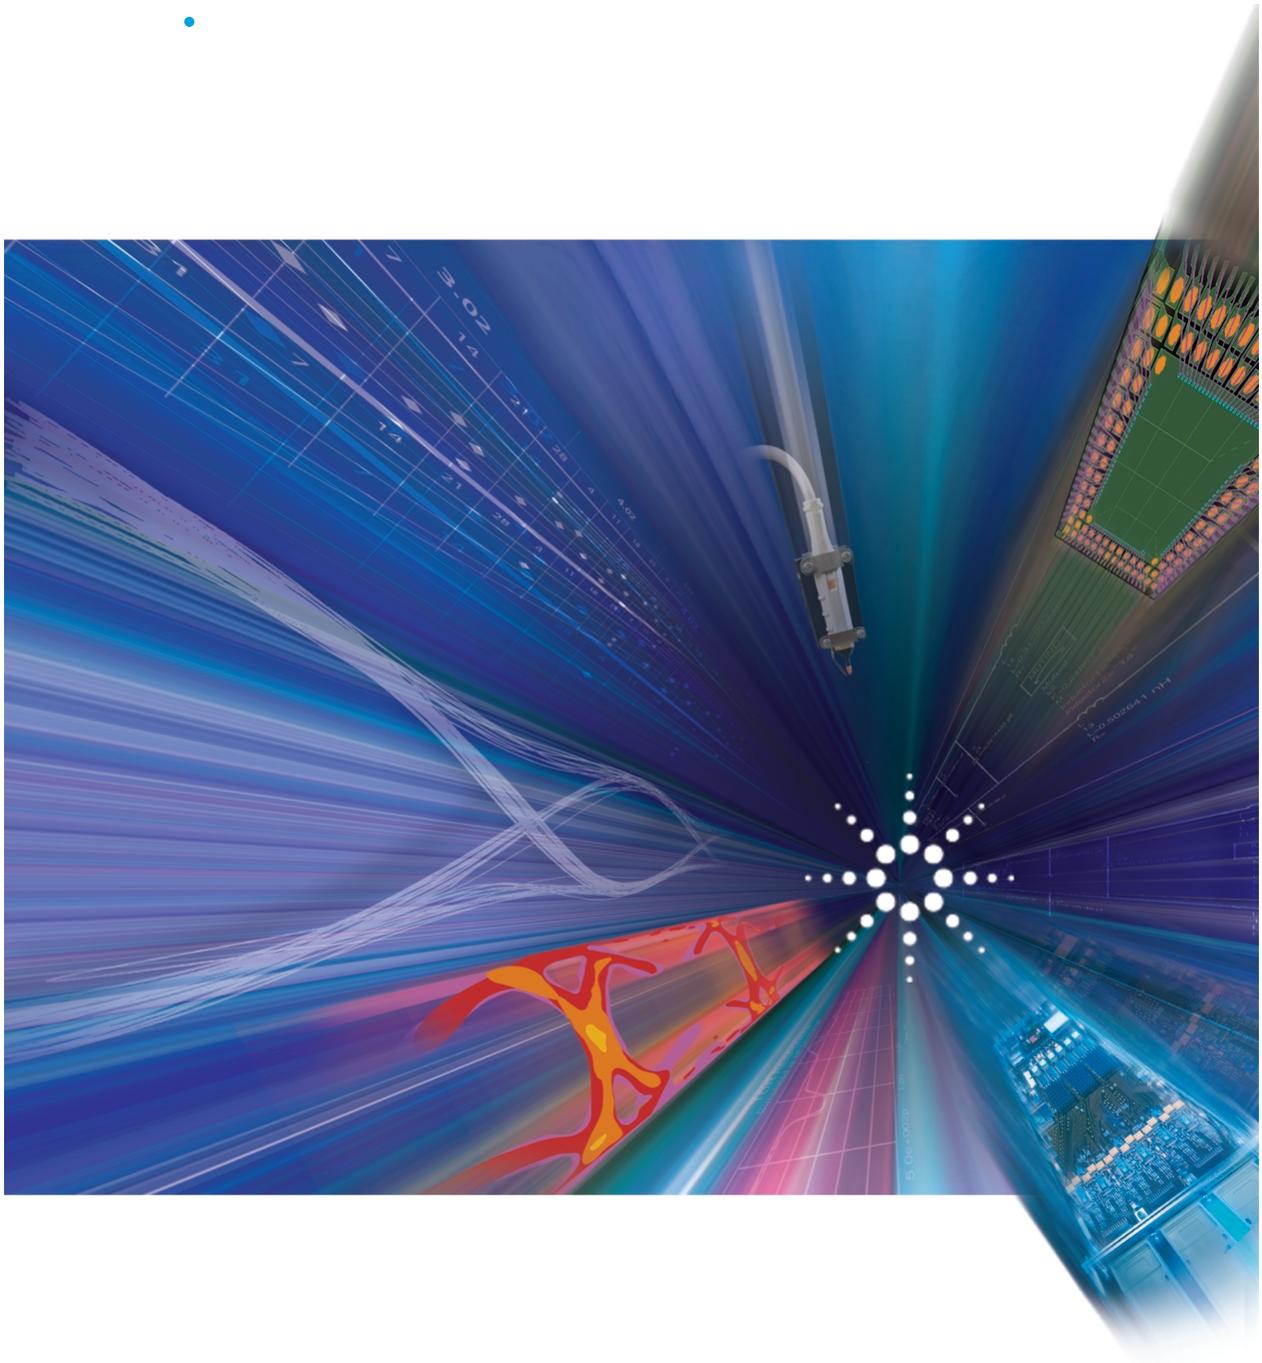


Table of Contents	
Introduction	2
Hint 1: Coordinate Your Tools Effectively	3
Hint 2: Use Frequency-Domain Instruments for Physical-Layer Testing of High-Speed Components	4
Hint 3: Understand Sources of Error in High-Speed Systems	6
Hint 4: Characterize Your Differential-Impedance Circuit Board	8
Hint 5: Validate Complex Systems with Test Cards	10
Hint 6: Integrate Your Tools	12
Hint 7: Save Time Making Signal Integrity Measurements with Eye Scan	13
Hint 8: Use LVDS for High-Speed Interconnects	14
Getting Aboard a High-Speed Bus Quickly	15
Support, Services, and Assistance	16

Who Should Read This Application Note?

Digital R & D engineers designing boards and subsystems with increasingly faster digital signals.

Introduction

The pace of innovation today is staggering as new technologies emerge onto the electronics scene and high speeds become even higher. It is becoming imperative to keep your design techniques current in order to remain competitive and meet your customers’ needs.

With edge rates operating at sub-nanosecond levels, digital design means effectively dealing with analog characteristics of digital signals, such as crosstalk, ground bounce, impedance, and EMI (electromagnetic interference), to maximize system performance and ensure signal quality. At these edge rates, a signal’s rise time, pulse width, timing, jitter, and noise content must be carefully measured and controlled.

While developing yesterday’s projects, you could move quickly through validation and debug using “rules of thumb” accumulated from design experience. Today you may find you need to spend more time in validation and debug. Design margins are much slimmer and it is a real challenge to understand and resolve problems. But this extra time is not available.

The following hints for debugging high-speed buses can get you started on updating those rules of thumb to meet the design and time-to-market challenges of today. The hints are abstracts of eight application notes available on our signal integrity web site at www.agilent.com/find/si. You can download these application notes and find a variety of other information to help you solve your signal integrity problems.

Hint 1: Coordinate Your Tools Effectively

InfiniBand System Level Debugging Case Study

InfiniBand interconnect technology uses a switched-fabric, point-to-point architecture to deliver new levels of scalability, availability, and performance. Beyond the high-speed data transfer rates, system performance is improved by distributing computing intelligence into the I/O system. This complex architecture can be a challenge to debug, but it can be accomplished with a system-level approach and the proper tools.

Three debugging concepts to follow are:

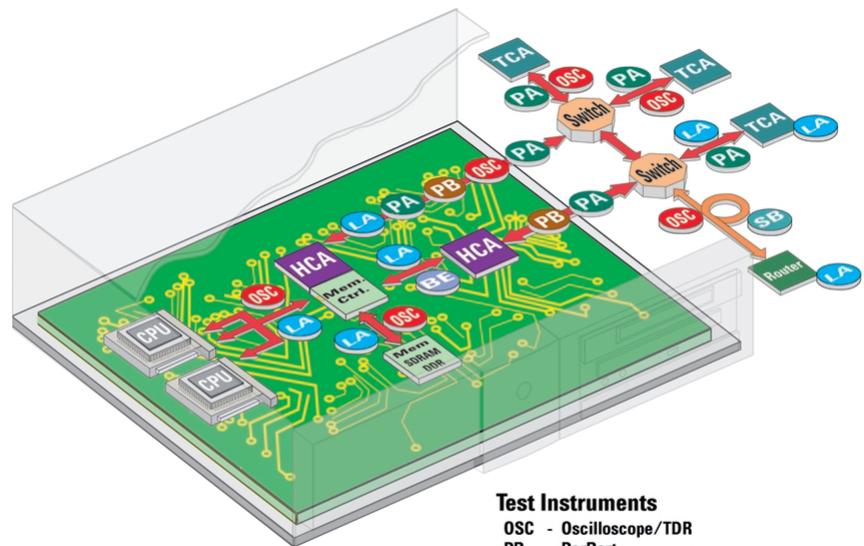
- Obtain broad visibility into all parts of your system.
- Cross-correlate activities in different parts of your system, including InfiniBand and other I/O buses such as PCI-X.
- Use a variety of stimuli, both artificially generated and from a real-world environment, to test your system.

As you prepare for system-level validation, you must select the appropriate tools for your test bench. Three basic types of test tools are most commonly needed for effective debugging and validation of InfiniBand systems at the higher protocol layers:— 1) a protocol analyzer, 2) a logic analyzer, and 3) a traffic generator. Figure 1 shows a typical InfiniBand system with the appropriate test equipment connected to the various links.

A protocol analyzer is optimized for protocol measurements, and focuses on providing a comprehensive view of information and data transfer on the InfiniBand link. A logic analyzer is most useful for looking at levels of the protocol up to the transport layer, and is optimized for providing cross-bus or multi-bus correlation of information. Traffic generators create controlled InfiniBand traffic for system validation, providing a controlled rather than a real-world stimulus.

It is important to understand the strengths of your tools, to know which one to select for a particular task, and to be able to use them efficiently. If you can coordinate your tools effectively, tracing your really tough system debugging problems to their root cause can be fairly straightforward.

For the complete application note, *InfiniBand System-Level Debugging*, and more signal integrity information, log onto: www.agilent.com/find/si.



InfiniBand Components
HCA - Host Channel Adapter
TCA - Target Channel Adapter

Test Instruments
OSC - Oscilloscope/TDR
PB - ParBert
SB - Serial BERT
LA - Logic Analyzer
PA - Protocol Analyzer/Traffic Generator
BE - Bus Exerciser/Analyzer (PCI-X)

Figure 1. A typical InfiniBand system showing interconnects to various components and appropriate test equipment.

Hint 2: Use Frequency-Domain Instruments for Physical-Layer Testing of High-Speed Components

Traditionally, time domain reflectometry (TDR) instruments have been chosen to analyze high-speed differential digital circuits and components. However, with increasing data rates, conventional validation tools are often inadequate for analyzing devices.

While there are many advantages to using frequency-domain measurements for physical-layer characterization of high-speed devices, two stand out among the most important. One is improved accuracy. In a vector network analyzer (VNA)-based system, a systematic error model is applied to raw measured data. The error terms for the model are derived from the calibration procedure (which involves connecting known standards to the instrument). The result is a measurement technique with bounded measurement uncertainties and traceability to a recognized standards laboratory.

The other significant advantage is exceptional dynamic range. For example, for a device having an effective 10/90 rise time of 35 ps, a VNA-based system can provide better-than-100-dB dynamic range in the time domain. This is important not only for crosstalk measurements, but also for mode-conversion and masking considerations.

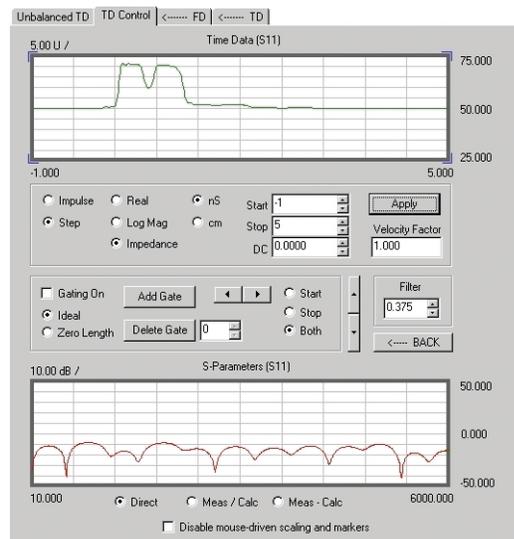


Figure 2. Multi-format displays improve data visualization.

Real-world differential devices often have undesirable common-mode responses, or exhibit unwanted mode conversions. Differential-to-common mode conversion can mean your device under test (DUT) will radiate interference signals or create unwanted ground currents that can couple to other circuitry. Common-to-differential mode conversion can mean your system will be susceptible to bit errors caused by external electromagnetic fields. The mode-conversion terms are ideally zero, and typically are very small in a well-designed device. Therefore, high dynamic range is necessary to pinpoint a problem with EMI (electromagnetic interference) generation or susceptibility.

Masking is another phenomenon addressed by high dynamic range. Time-domain responses are most accurate closest to the location of the source. A discontinuity in the DUT will reflect some power back to the source, so less power is transmitted to the rest of the DUT, reducing the measurement accuracy of discontinuities farther from the source. Greater dynamic range extends the ability of the instrument to accurately characterize devices that have several discontinuities or high loss.

Hint 2: Use Frequency-Domain Instruments for Physical-Layer Testing of High-Speed Components (continued)

Traditional two-port VNA systems are designed for characterizing unbalanced (single-ended) devices, and provide the test data in a format that may be of limited use in signal integrity applications. New four-port (two-channel) Agilent physical layer test systems have been developed specifically for differential devices, bringing the precision and accuracy of the VNA to the world of high-speed, balanced devices. These high-dynamic-range test systems let you see elusive EMI problems that previously might have gone undetected.

With a single set of connections, you can measure the single-ended, differential-mode, common-mode, and mode-conversion behavior of your DUT. When characterizing devices such as a pair of printed-circuit traces, you can

analyze each trace by itself to measure total delay and skew, or analyze them together as a balanced pair.

The format of the data can be changed depending on what is most meaningful for a given type of device, or for the type of information that is needed. These formats include time domain (TDR/TDT), frequency domain (S-matrix), eye diagram, and RLCG extraction (transmission line parameters).

Several additional capabilities are increasingly important as device speeds increase. These include de-embedding the effects of test fixtures and probes, simulating the effects of compensation networks on a DUT, translating the performance to an alternate reference impedance, and

examining the effects of phase skew. These types of analysis are most conveniently done in the frequency domain, so working with frequency parameters is more natural.

Frequency-domain measurements provide the type of information needed to perform physical-layer characterization of high-speed differential devices, and this methodology is becoming increasingly necessary. The VNA-based system's combined characteristics of accuracy, dynamic range, comprehensiveness, and flexibility make it ideal for such applications.

For the complete application note, *Physical-Layer Testing of High-Speed Components*, and more signal integrity information, log onto: www.agilent.com/find/si.

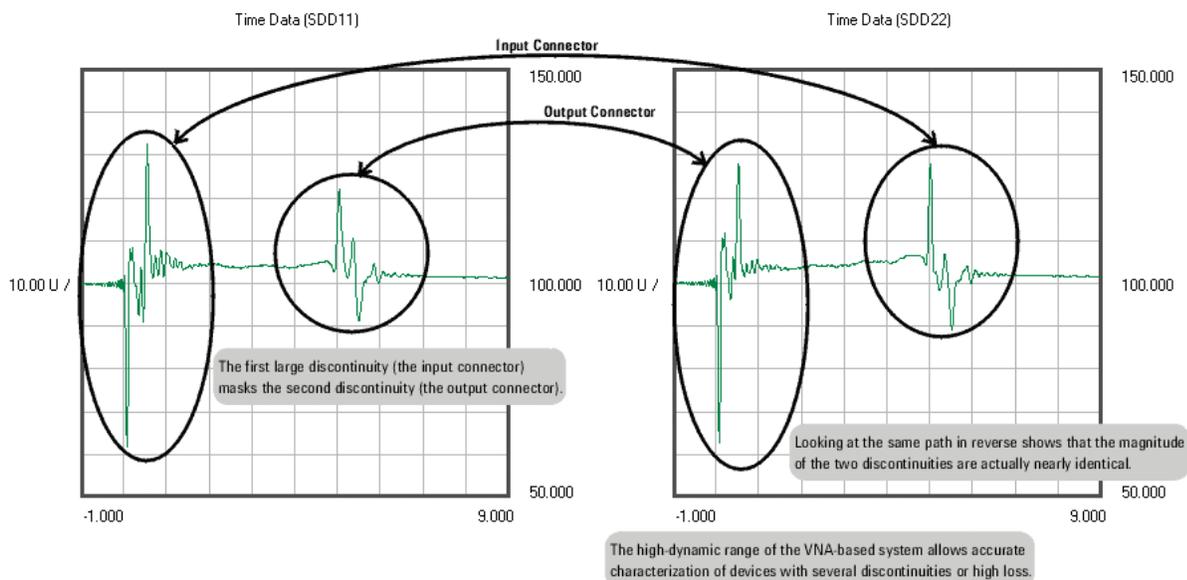


Figure 3. High-dynamic range defeats the masking effects of large discontinuities.

Hint 3: Understand Sources of Error in High-Speed Systems

Higher clock speeds are the obvious technology trend, but the many related changes have an equal or greater impact on designs. Faster clock speeds require smaller voltage swings and shorter setup and hold times. Data-valid windows become orders of magnitude smaller. Jitter caused by noise, crosstalk, and intersymbol interference further reduces their size, creating errors. Because the noise margins are so small, noise and timing budgets can no longer tolerate phenomena that were previously ignored. Figure 4 is an eye diagram of a single data channel. The first data bit has a data-valid window with an amplitude of 700 mV and a data-valid time of 1.5 ns. But a few bits later, a small amount of jitter has reduced the data-valid window to 1 ns.

Today's high-speed bus designer needs to carefully consider proper termination of signals and correct impedance matching. Higher frequencies result in shorter wavelengths, so board layout becomes critical. Traces need to be treated as transmission lines and even package leads need to be considered as potential antennas. Designers have to watch for breaks in the ground plane, match impedances to reduce reflections, and worry about trace separation, trace length, and even discontinuities in the FR4 material.

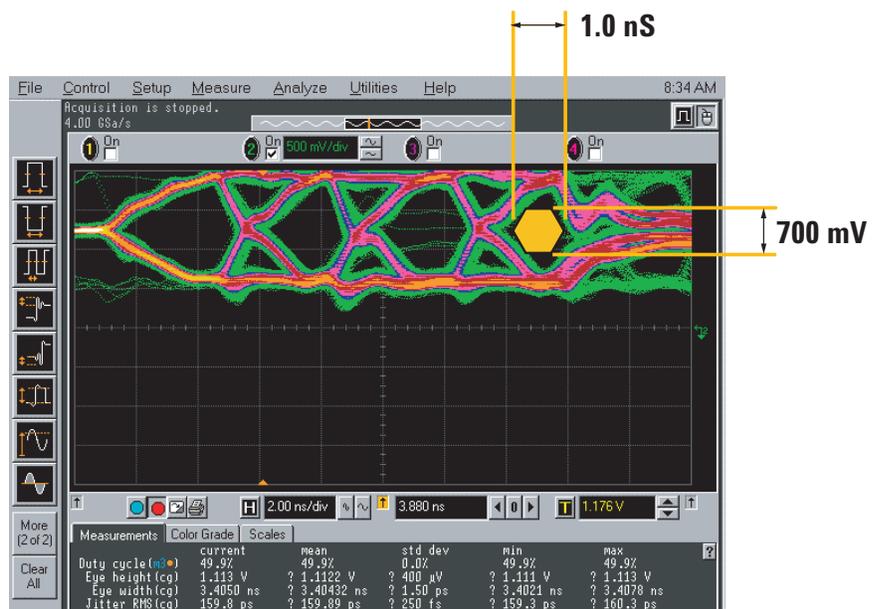


Figure 4. This eye diagram of a single data channel shows how the data-valid time can shrink from 1.5 ns to 1.0 ns with just a small amount of jitter.

Bus architectures are no longer clocked with one or two signals; source-synchronous bus designs can have a dozen or even more clocks or strobes moving the data. DDR (double data rate) is a source-synchronous bus. Instead of one clock, it has 19—one for the control signals and 18 strobes (or clocks) for the data groups. Each data group has its own strobe that is most likely skewed from the other data groups.

So an increase in bus speeds requires many changes to design methods. These same changes affect the way buses are tested and measured. One serious issue affecting testing is probing. Whenever a measurement is made, the connection of the test equipment affects the measurement to some extent. The effect of probing on a circuit becomes more pronounced as the frequency increases. Probing is no longer just clipping a probe to a test point; the probing system must be specifically designed for the bus under test.

Hint 3: Understand Sources of Error in High-Speed Systems (continued)

Source-synchronous data buses create an entirely new testing challenge. Each strobe has 4 data lines assigned. However, a logic analyzer cannot sample using 19 clocks—only one can be chosen. Furthermore, skews between sections may be large, and the data-valid windows are tiny. Realigning the data bus and clocking the data in at the right time are challenges for both the DDR bus designer and for the test equipment measuring the bus (figure 5a). It is possible to do this manually but it would take many hours, possibly days, to accomplish.

Figure 5b shows how Agilent Technologies' automated "eye finder" technology can solve the channel-to-channel skew problem created by a source-synchronous bus in a matter of minutes with 100-ps resolution. The sample and hold setting for each channel is aligned with DQS0.

Higher-performance buses are driving changes in design and testing methodologies. Signal integrity is critical because of the shrinking data-valid window. You need to seek out the most advanced probing and logic analysis tools to solve today's and tomorrow's complex design challenges.

For the complete application note, *Designing and Validating High-Speed Memory Buses*, and more signal integrity information, log onto: www.agilent.com/find/si.

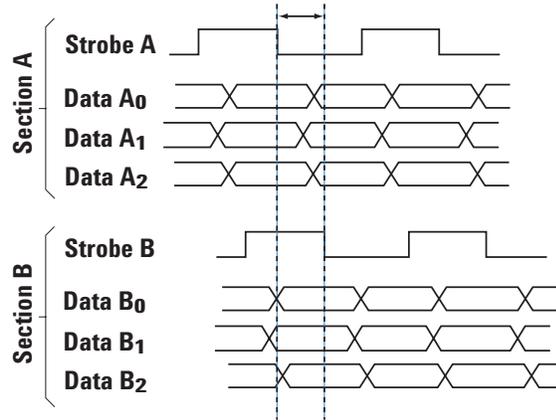


Figure 5a. It is not possible to sample with all 19 clocks of the DDR bus to compensate for the skew between sections. Using a single clock to sample all the skewed bus data can be a challenge.

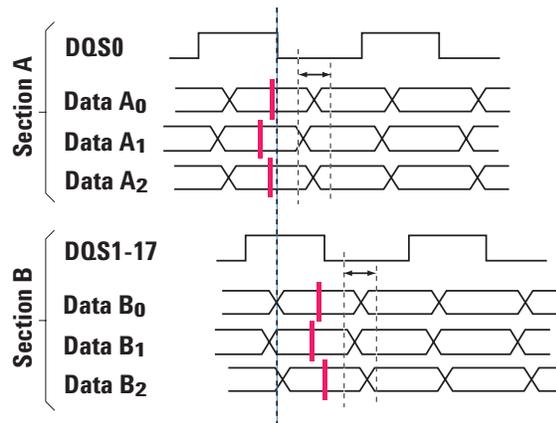


Figure 5b. Agilent's eye finder technology automatically solves the channel-to-channel skew problem in a matter of minutes.

Hint 4: Characterize Your Differential-Impedance Circuit Board

Differential-impedance circuit boards are becoming more common as low-voltage differential signaling (LVDS) devices proliferate. Yet there is much confusion in the industry about what differential impedance means, how to design for it, and how to leverage its benefits for noise rejection. Knowing the general features of differential-impedance transmission lines and how they can be characterized with traditional time domain reflectometry (TDR) instruments can help in the design and testing of high-speed digital systems.

Dual-channel TDR can be used to analyze features and real-world effects of high-speed differential transmission lines under a variety of conditions such as a gap in the return path or skew between the two channels. It can be hooked up in a variety of ways to maximize the amount

of information obtained. One channel of the TDR plug-in can be used to perform conventional TDR analysis. Using a second TDR channel allows analysis of the properties of differential transmission line pairs.

With time domain transmission (TDT), the first channel generates the exciting source into one end of a transmission line and the second TDR channel is the receiver at the other end. In this way, the TDR and TDT response of the device under test (DUT) can be measured simultaneously. The TDR response gives information about the impedance of the DUT, and the TDT gives information about the signal propagation time, signal quality, and rise-time degradation. In this mode, the TDT is emulating what a receiver might see at the far end.

A dual-channel TDR module allows the measurement of the

impedance characteristics of any coupled differential pair. With the addition of a 2-channel, 50-GHz plug-in module, the signals propagated to the end of the differential pair can be measured, emulating what actual far-end receivers might see.

In the example shown in figure 6, the signal at the far end when the pairs are driven differentially is measured. In the upper left screen shot, the TDR response without the DUT connected is shown. This highlights that one channel is driving a signal of 0 to 400 mV, while the other channel is driving a signal of 0 to -400 mV. What gets launched into a 50 Ω load is 0 to 200 mV in channel 1 and 0 to -200 mV in channel 2.

At the far end of the roughly 50 Ω differential pair, the two channels of the 50-GHz plug-in module measure the received voltage into a 50 Ω load. This shows the roughly 100 ps rise time from propagating down 8 inches of FR4. The individual channels are displayed as directly measured. In addition, the common signal (being the average of these two) and the differential signal can be automatically displayed. All received signals are displayed on the same scale. When the pair is driven with a well-balanced differential signal, the common signal created by the transmission down the path is virtually nonexistent.

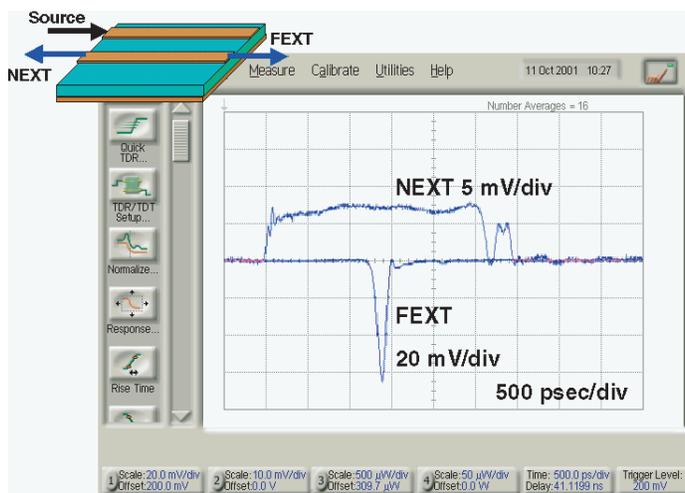


Figure 6. Measuring the impedance characteristics of a coupled differential pair.

Hint 4: Characterize Your Differential-Impedance Circuit Board (continued)

A frequent problem with differential drivers for differential-pair lines is skew between the two channels. This arises due to driver mismatch, different rise and fall times, different interconnect delays due to routing differences, or different loads on the two lines of the differential pair. Any signal imbalance at the receivers will create a common signal.

A variable skew can be introduced between the two driven TDR step generators, emulating what would happen if there were a skew in the drivers. In the example shown in figure 7, the common signal increases steadily as the skew increases from 0 to 100 ps, comparable to the rise time. If the skew is longer than 100 ps, the common signal at the receiver is basically constant. This suggests that to minimize the common signal, the skew should be kept to just a small fraction of the rise time.

Dual-channel TDR and a dual-channel amplifier module can provide full characterization of differential pairs of transmission lines, enabling you to better determine how your designs are working at the electrical level and suggesting ideas for how to improve them.

For the complete application note, *Differential Impedance Measurement with Time Domain Reflectometry*, and more signal integrity information, log onto: www.agilent.com/find/si.

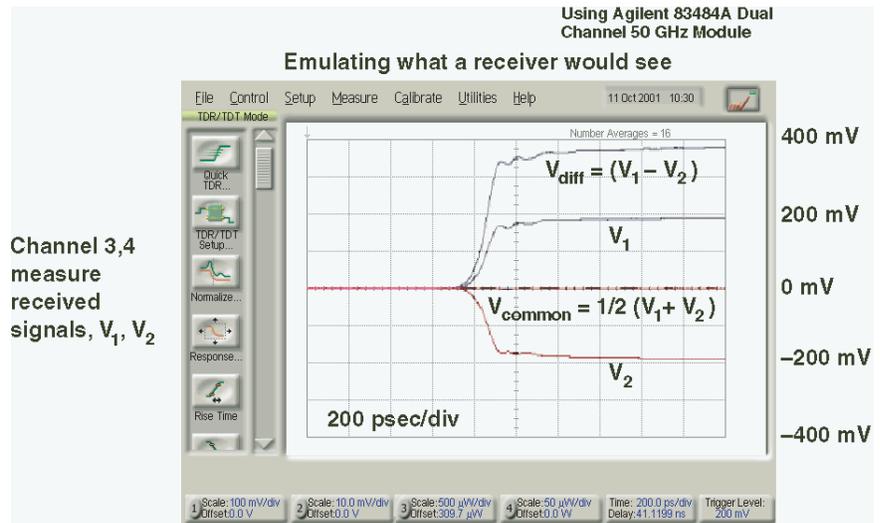


Figure 7. Emulation of a skewed differential pair by introducing a skew between the two TDR step generators .

Hint 5: Validate Complex Systems with Test Cards

PCI/PCI-X Validation Case Study

Validation of computer systems and subsystems is becoming increasingly complex as I/O systems and peripherals become more intelligent. The role of data transfer initiator is being delegated to the I/O systems instead of the CPU, causing data traffic to move in several directions simultaneously and freeing the CPU for data processing tasks. Increasing bandwidth needs require performance optimization in the subsystems, further increasing the complexity of the whole system.

Validation of these complex systems is becoming a very difficult task. It is necessary to

test systems and subsystems under real-life conditions while ensuring that corner-cases are covered, to confront the system under test (SUT) with all possible combinations of traffic, and to generate peak load conditions. Furthermore, tests must be reproducible to enable debugging.

How can you ensure that your products are well tested, corner-cases are covered, and the time spent is within reasonable limits? The solution is a combination of test cards and software running on the SUT. A test card is a device that is used specifically for testing. It is designed for a specific slot-based or cable-based I/O system—PCI, PCI-X, or others—and operates like any other device designed for

that system. It can allocate system resources, generate any type of traffic that is allowed, and also react to a transaction initiated somewhere else.

The test card has an external interface, so it can be controlled with an external controlling host, making it independent of the type of system and the operating system used. When used with an “internal” connection (directly from the SUT), it is possible to coordinate tests using the test card with other test tools or programs.

Using these test-card features, you can generate data traffic along any data path that is accessible by a connector. A bridge can be tested and stressed by placing two cards on either side of the bridge. Programming a test card to communicate directly with a device is a way to test that device. Because the traffic generated by the test card can be programmed directly, it is also completely repeatable, so generating reproducible results is much easier. Figure 8 shows a sample system with multiple I/O buses and test cards connected at appropriate points.

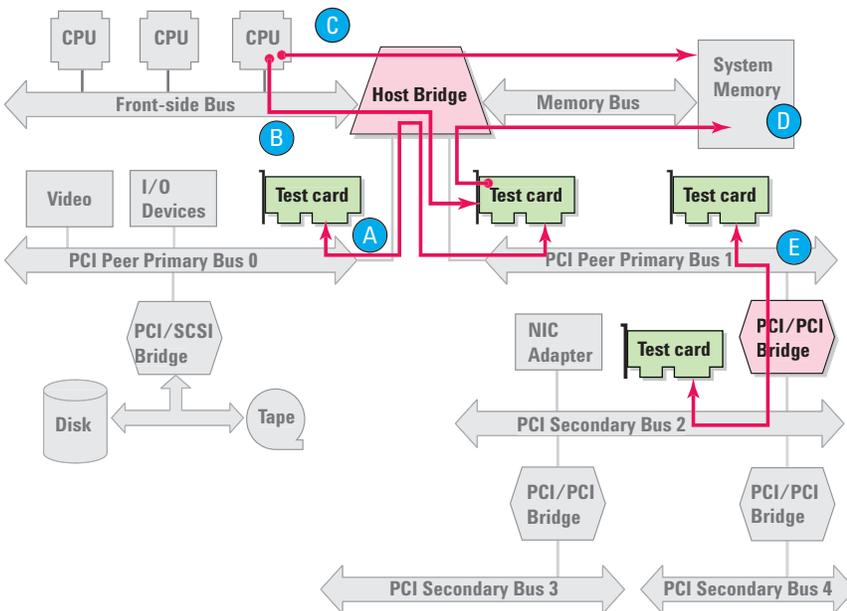


Figure 8. Using test cards, you can generate data traffic along any data path that is accessible by a connector. A bridge can be tested and stressed by placing two cards on either side of the bridge.

Hint 5: Validate Complex Systems with Test Cards (continued)

It is extremely useful if the test card has analyzing as well as exercising capabilities. These include the ability to monitor the transaction protocol, gather performance metrics, and provide traces of the ongoing transactions.

A test card placed at a central position within the system also can serve as a window into the system. It can read or write registers in system memory or the system's I/O space. It can read or write the configuration data of other devices including bridge devices, and can dump the contents of memory areas. This is possible even when the system itself is no longer operating.

Some system validation procedures cannot be accomplished by using test cards alone. In these cases, you can use a combined approach, or "validation framework". The validation framework handles all aspects of test-card testing—setup, running, and analysis—along with tests written specifically for devices that require CPU interaction. All testing is handled by a single piece of software, which has an application programming interface (API) that allows you to add new tests and to configure existing ones for different needs (figure 9).

Validating complex server and workstation systems has become as complex as the systems themselves. Using a combination of test cards and specific tests combined within a validation framework is an approach that can significantly reduce testing time and improve testing reproducibility.

For the complete application note, *I/O System and Chip Verification in PCI and PCI-X Systems*, and more signal integrity information, log onto: www.agilent.com/find/si.

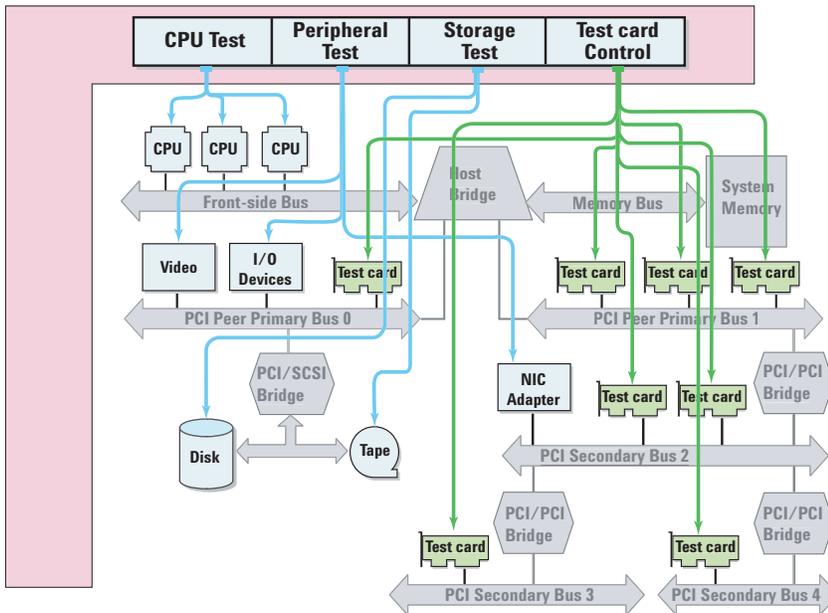


Figure 9. By using appropriate test cards, a single piece of software enables you to easily test very complex systems.

Hint 6: Integrate Your Tools

USB 2.0 Debugging Case Study

USB 2.0's high-speed data rate permits bandwidth-hungry peripherals to use a USB system, but it does create debugging problems. Signal integrity becomes critical at its data rate (480 Mb/s) and RF and even microwave design techniques must be used. Yet the digital world (data domain) also creates issues that affect the designer: interaction of devices, multi-speed traffic, and large amounts of data. Because the analog and digital behaviors are interrelated, designers of USB 2.0 devices need to take a holistic approach to debugging, validating, and characterizing their designs.

Many types of tools are needed to make accurate analog scope measurements, current measurements, and to capture and analyze data. The complexities of oscilloscopes, MATLAB®, software, logic analyzers, and breakout boards turn debugging, performance characterization, and specification compliance testing into daunting tasks. Fortunately, tools are available that simplify the process by taking advantage of the integration of PCs and test equipment. For example, MATLAB software is integrated into an oscilloscope; it eliminates the need to make a measurement and then transfer the data to a PC to do voltage or current analysis

(figure 10). Today's logic analyzers and oscilloscopes are integrated so that cross-triggering and even sharing of data on a single display simplify the debug process.

Another key test and debug technique is cross-bus analysis. It may be desirable to look at multiple USB hubs concurrently, to study a hub and the PCI bus coming out of a USB/PCI adapter card, or to observe the interaction between the USB, PCI, and PC memory system. Logic analyzers are capable of looking at multiple buses simultaneously, having one bus trigger other buses, and having all of the time-correlated data presented on a single display.

Using tools designed for the high-speed analog/digital world enables you to develop reliable USB 2.0 products and quickly bring them to the marketplace.

For the complete application note, *Debugging USB 2.0 for Compliance: It's Not Just a Digital World*, and more signal integrity information, log onto: www.agilent.com/find/si.

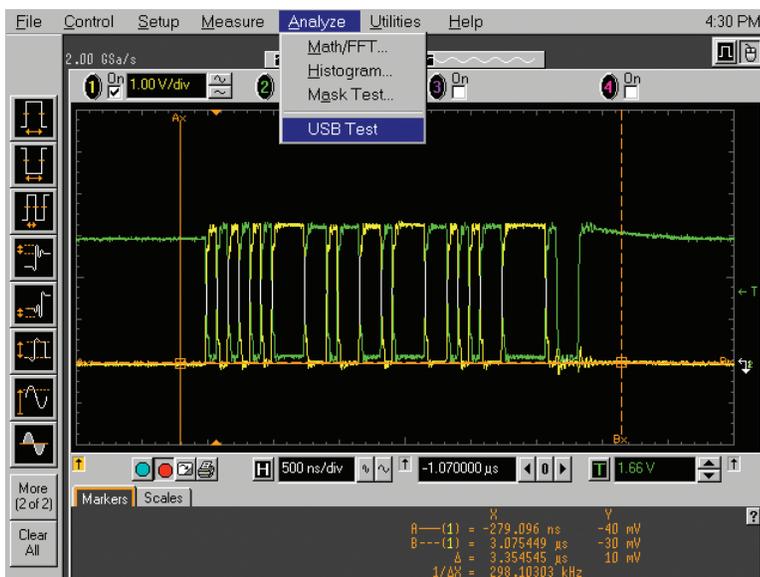


Figure 10. Including MATLAB software within an oscilloscope simplifies USB measurements.

Hint 7: Save Time Making Signal Integrity Measurements with Eye Scan

Oscilloscopes have traditionally been used to characterize and validate the signal integrity of prototype circuits. However, making multiple eye-diagram measurements with an oscilloscope can be a slow process due to the time required to connect and move probes. Measuring the hundreds of nodes on a complex system requires moving the probes hundreds of times. With eye scan, the connection is instantaneous if a connector has been designed into the board, so you can measure the signal integrity behavior on tens or hundreds of signal nodes much more quickly than with an oscilloscope. This enables you to acquire comprehensive signal integrity information on all buses in a design under a wide variety of operating conditions in a reasonable amount of time. Eye scan is a measurement tool that can reduce the time required to verify signal integrity in complex high-speed designs. When running eye scan, the logic analyzer scans all incoming signals for activity in a time range centered on the clock and over the entire voltage range of the signal. The results are displayed in a graph similar to an oscilloscope eye diagram (figure 11). Display colors correspond to the amount of signal activity detected.

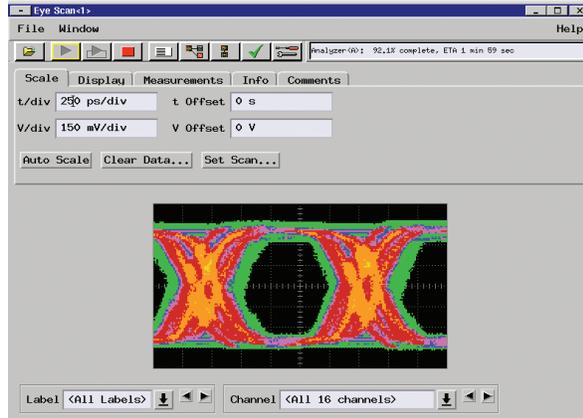


Figure 11. Measurements made with a logic analyzer are displayed in a format similar to an oscilloscope eye diagram.

Eye scan examines regions of time and voltage for signal transitions. The time regions are defined relative to active clock transitions in the user's system. The scan proceeds first along the time axis. When the specified range of time has been scanned, the voltage threshold is incremented and the time range is scanned again at a new threshold. This is repeated until all time and voltage regions have been scanned. The user can adjust the scan range and resolution in both time and voltage.

For the complete application note, *Saving Time with Multiple-Channel Signal Integrity Measurements*, and more signal integrity information, log onto: www.agilent.com/find/si.

Hint 8: Use LVDS for High-Speed Interconnects

Low-voltage differential signaling (LVDS) offers high-speed data transfer for innumerable types of interconnects. It uses high-speed analog circuit techniques to provide multi-gigabit data transfers on copper cables or printed circuit board traces.

LVDS, a generic interface standard, moves information on a board, between boards, modules, shelves, and racks, or box to box. The equivalent circuit structure of the LVDS physical layer is shown in figure 12.

Low-voltage signals have many advantages, including fast bit rates, lower power, and better noise performance. Because its voltage change between logic states is only 300 mV, LVDS can change states very fast. Low voltage swing reduces power consumption because it lowers the voltage across the termination resistors and lowers the overall power dissipation. To improve noise immunity, LVDS uses differential data transmission. Differential signals have the advantage of tolerating interference from outside sources such as inductive radiation from electric motors or crosstalk from neighboring transmission lines. A differential receiver responds only to the difference between the two inputs, so when noise appears commonly to both inputs, the input differential signal amplitude is undisturbed.

LVDS can also tolerate minor impedance mismatches in transmission paths. As long as the differential signal passes through balanced discontinuities in closely-coupled transmission paths, the signal can maintain integrity. The effect of non-impedance-controlled

connectors, printed circuit board vias, and chip packaging is not as detrimental to differential signals as it is to single-ended signals. The final LVDS system benefit is its integration capability.

LVDS is now spawning follow-on technologies that expand its applications. Bus LVDS (BLVDS) allows the low-voltage differential signals to work in bidirectional and multidrop configurations. Another LVDS derivative, ground-referenced LVDS (GLVDS), moves the differential signal's common-mode voltage close to ground, enabling chips operating from very low supply voltages to communicate over a high-speed standard interface.

When using LVDS for high-speed interconnects, you must verify the integrity of the complete signal path, which can be a difficult and time-consuming process for a complex device. A parallel bit-error-ratio tester, can speed up this process, allowing you to quickly verify the signal integrity of the physical layer.

For the full application note, *Low-Voltage Differential Signaling (LVDS)*, and more signal integrity information, log onto: www.agilent.com/find/si.

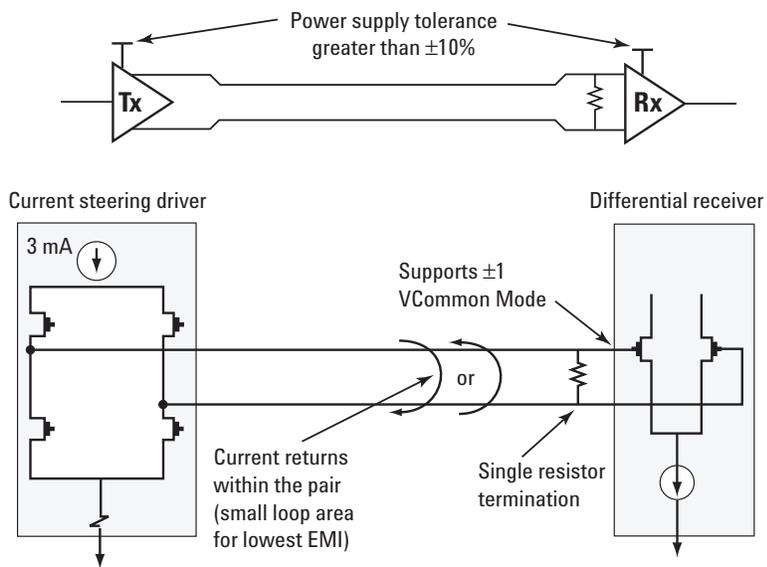


Figure 12. The equivalent circuit structure of the LVDS physical layer.

Getting Aboard a High-Speed Bus Quickly

The speed and complexity of today's leading-edge bus architecture requires much more effective test equipment than that used for past buses. High data rates, new differential logic families, complex protocols, and the RF nature of the latest buses make much of the equipment you may have been using up until now inadequate. At each stage of development, high-speed systems will test your existing knowledge and instrumentation.

The physical layer of any high-speed system must be designed carefully. It must be laid out properly to ensure that the digital signals reliably move along the bus. Great care must be taken to ensure that the effects of reflections, crosstalk, and other noise sources are reduced to an acceptable level. Time domain reflectometers (TDRs) based on the Agilent 86100 Infiniium DCA with TDR, and vector network analyzers (VNAs) such as Agilent's new four-port physical layer test systems are ideal for measuring the parameters of signal paths. You can use them to make sure that your signals will not be distorted and will result in an acceptably low error rate.

Once the physical layer of a high-performance bus is fully characterized and its properties meet the expected level of performance, it is necessary to verify that actual digital data will be properly transferred on the bus. Your goal is to verify that the design has an acceptable bit error

ratio (BER). Capable instruments such as the Agilent ParBERT 81250 enable you to quickly verify the data path.

After you are certain that the bus will accurately transfer digital data, you must then verify the protocol it uses. Most high-performance buses have some kind of protocol used to control data transfer on the bus. As these buses get more complex, the protocols can become unwieldy without appropriate test instruments that decode the protocol and enable you to verify the operation at a high level. For example, with InfiniBand, the Agilent E2950 Series of protocol analyzers and traffic generators/exercisers present bus traffic in a readily understood manner, making it possible for you to easily validate an InfiniBand system and make it compliant with InfiniBand Trade Association design specifications.

High-speed buses always operate within high-speed systems. These systems use the bus to transfer data between complex subsystems. Once a bus is known to work properly, the entire system must be debugged and validated. Often, the various subsystems operate somewhat independently and require very capable test instruments to monitor operations and then cross-correlate the information so you can understand what is going on. The Agilent 16700 logic analysis family enables you to combine triggering and data gathering with events throughout the system.

Agilent Technologies' Test and Measurement Support, Services, and Assistance

Agilent Technologies aims to maximize the value you receive, while minimizing your risk and problems. We strive to ensure that you get the test and measurement capabilities you paid for and obtain the support you need. Our extensive support resources and services can help you choose the right Agilent products for your applications and apply them successfully. Every instrument and system we sell has a global warranty. Support is available for at least five years beyond the production life of the product. Two concepts underlie Agilent's overall support policy: "Our Promise" and "Your Advantage."

Our Promise

Our Promise means your Agilent test and measurement equipment will meet its advertised performance and functionality. When you are choosing new equipment, we will help you with product information, including realistic performance specifications and practical recommendations from experienced test engineers. When you use Agilent equipment, we can verify that it works properly, help with product operation, and provide basic measurement assistance for the use of specified capabilities, at no extra cost upon request. Many self-help tools are available.

Your Advantage

Your Advantage means that Agilent offers a wide range of additional expert test and measurement services, which you can purchase according to your unique technical and business needs. Solve problems efficiently and gain a competitive edge by contracting with us for calibration, extra-cost upgrades, out-of-warranty repairs, and on-site education and training, as well as design, system integration, project management, and other professional engineering services. Experienced Agilent engineers and technicians worldwide can help you maximize your productivity, optimize the return on investment of your Agilent instruments and systems, and obtain dependable measurement accuracy for the life of those products.



www.agilent.com/find/emailupdates

Get the latest information on the products and applications you select.

Related Literature

Publication Title	Publication Type	Publication Number
<i>InfiniBand System-Level Debugging</i>	Application note	5988-4225EN
<i>Physical-Layer Testing of High-Speed Components</i>	Application note	5988-5075EN
<i>Designing and Validating High-Speed Memory Buses</i>	Application note	5988-4497EN
<i>Differential Impedance Measurement with Time Domain Reflectometry</i>	Application note	5988-4796EN
<i>I/O System and Chip Verification in PCI and PCI-X Systems</i>	Application note	5988-4795EN
<i>Debugging USB 2.0 for Compliance: It's Not Just a Digital World</i>	Application note	5988-4794EN
<i>Saving Time with Multiple-Channel Signal Integrity Measurements</i>	Application note	5988-5409EN
<i>Low-Voltage Differential Signaling (LVDS)</i>	Application note	5988-4797EN

For these application notes, signal integrity web casts and more information, log onto www.agilent.com/find/si.

By internet, phone, or fax, get assistance with all your test & measurement needs

Online assistance:
www.agilent.com/find/assist

Phone or Fax United States:
(tel) 800 452 4844

Canada:
(tel) 877 894 4414
(fax) 905 282 6495

China:
(tel) 800 810 0189
(fax) 800 820 2816

Europe:
(tel) (31 20) 547 2323
(fax) (31 20) 547 2390

Japan:
(tel) (81) 426 56 7832
(fax) (81) 426 56 7840

Korea:
(tel) (82 2) 2004 5004
(fax) (82 2) 2004 5115

Latin America:
(tel) (305) 269 7500
(fax) (305) 269 7599

Taiwan:
(tel) 0800 047 866
(fax) 0800 286 331

Other Asia Pacific Countries:
(tel) (65) 6375 8100
(fax) (65) 6836 0252
Email: tm_asia@agilent.com

Product specifications and descriptions in this document subject to change without notice.

©2002 Agilent Technologies, Inc. 2002
Printed in USA May 7, 2002
5988-5623EN