Designer's

to verifying signal integrity



▶ Primer

Table of Contents

Introduction and Overview	3 - 5)
Design Project Overview	3	
Measurement Steps Keep Pace with Design Steps	4	
Review of Signal Integrity Concepts	5	
Chapter 1 – Impedance Measurements	6 - 1	10
The Foundation of Signal Integrity	6	
Overview of Impedance Measurement Tools and Techniques	6	
An Impedance Measurement Procedure and Its Conclusions	7	
Setting Up and Making Connections	7	
True Differential TDR Measurements	8	
TDR Measurements and Results	9	
Chapter 2 – ASIC Verification	11 -	- 13
Ensuring Signal Integrity at the Source	11	
Digital Verification Steps Range from Basic to Advanced	12	
Digital volmoution otopo hango nom Daolo to havanood		
Chapter 3 – Basic Functional Verification	14 -	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design	14 – 14	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design Verification Tools	14 - 14 14	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design Verification Tools Checking the Pulse of a New Design	14 - 14 14 16	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design Verification Tools Checking the Pulse of a New Design Spread Spectrum Clocking and Serial ATA	14 - 14 14 16 17	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design Verification Tools Checking the Pulse of a New Design Spread Spectrum Clocking and Serial ATA SSC Example: The Serial ATA Clock	14 - 14 14 16 17 16	- 18
Chapter 3 – Basic Functional Verification Checking the "Heartbeat" of an Emerging Design Verification Tools Checking the Pulse of a New Design Spread Spectrum Clocking and Serial ATA SSC Example: The Serial ATA Clock Chapter 4 – Internal Bus Compliance	14 - 14 14 16 17 16 19 -	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 - 19	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 16 17 16 19 - 19 19	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 16 17 16 19 - 19 19 19	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 - 19 19 19 20	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 - 19 19 19 20 20	- 18 - 23
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 19 19 20 20 20	- 18
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 - 19 19 19 20 20 20 20 21 21	- 18
 Chapter 3 – Basic Functional Verification	14 - 14 14 16 17 16 19 - 19 19 19 20 20 20 21 21	- 18

- 31
ļ
ļ
,
5
,
)
2 – 35
;
;
ļ
5
;
;
5
,
,
, ,

Introduction and Overview

This primer is intended for engineering professionals with an interest in high-performance digital system design under real-world conditions: time constraints, cost limitations, quality requirements, and manufacturability concerns.

The primer focuses on signal integrity, a key measurement issue affecting digital designs. The term "signal integrity" encompasses the analog factors that affect both the performance and the reliability of digital designs. As system speeds increase, signal integrity becomes a greater challenge: it is ever more difficult to maintain clean pulse edges, low noise and aberrations, and nominal amplitude and timing characteristics. A rigorous regime of signal integrity measurements can trace these problems to their root causes.

Many designers are concerned with compliance measurements. Compliance with industry standards ensures interoperability among system elements, and the measurements usually entail a series of prescribed acquisition and analysis steps. But successful compliance testing often depends on finding and eliminating the signal integrity problems that are a frequent cause of errors. Consequently this primer will examine not only codified compliance measurements, but also the signal integrity issues that can affect them.

The design process is made up of multiple steps, each with its own particular signal integrity challenges and measurement needs. To illustrate some of the solutions that serve these needs, we will follow a new server motherboard design as it proceeds from raw printed circuit board (PCB) to finished product.



Figure 1-1. The server motherboard incorporates several industry -standard serial buses

Design Project Overview

The design is a server motherboard, illustrated in Figure 1-1*1. The performance of this motherboard, its processors, and chipsets will determine the overall data throughput of a flagship system-level product.

It is critical to maximize the data throughput, but not at the cost of reliability. Every functional element of the new motherboard must be validated along the way, from impedances on the raw circuit board to compliance tests and characterization of the completed design. Doing so will speed development (catching problems early and minimizing rework) and enable the new product to reach the market on time. In the highly competitive marketplace for which the product is being developed, the majority of the profits will go to the product that delivers the most performance, soonest.

^{*1} The motherboard and system described in this primer are constructs created to illustrate the steps of the design process. The motherboard may include combinations of features not commonly found within a single device. Nevertheless, each of the steps summarized in this document resolves a real-world measurement challenge.

Primer

Pre-Assembly Ve	erification	Functional Verification & Troubleshooting			System Verification & Characterization		
Impedance Measurements	ASIC Verification	Basic Functional Verification	Internal Bus Compliance	Operational Validation, Fault Detection, Debug	External Bus Compliance	Characterization	
Check impedance on PCB traces & connectors	Verify ICs will support system specs	Check clocks, buses, enables, etc.	Perform eye diagram analysis	Run tests; detect errors; identify and track problems	Confirm signal integrity and run eye diagram tests on outputs	Measure and document signal margins and limits	

Figure 1-2. The steps in the motherboard design process.

The motherboard will be part of an end product designed for customers who need a high-throughput server implementation. To deliver the requisite performance and reliability, the motherboard will incorporate advanced technologies such as PCI-Express, Serial ATA, XAUI, Gigabit Ethernet (GbE), and double data rate RAM (DDR RAM; in this case, the second-generation DDR 2).

PCI-Express replaces traditional chip-to-chip buses such as the graphics interface bus and the PCI system bus. Advanced PCI-Express implementations can provide a robust 10 Gb/s interface for the network fabric, supporting protocols as diverse as Gigabit Ethernet, USB 2.0, 10 Gb FibreChannel, and others. This motherboard design includes a 4X PCI Express bus for graphics and a 16X PCI Express bus for I/O.

Like many server designs, the motherboard relies on a pair of CPUs to execute instructions and move data speedily. A high-speed Memory Controller Hub (MCH or "North Bridge" coordinates transactions and carries data between the DDR 2 channel, the CPU pair, and other functional elements on the board. The I/O Control Hub (IOCH or "South Bridge") oversees transactions with external peripheral buses as well as connections to internally-mounted PCI devices.

Multi-gigabit data rates are present on buses throughout the board. Unit Intervals (UI) are brief; in the low hundreds of picoseconds (ps). Many signals have risetimes below 100 ps. Data is transported both internally and externally on differential transmission lines that must meet rigorously specified industry standards. Most importantly, the board's high-speed operating characteristics require careful attention to signal integrity issues.

Measurement Steps Keep Pace With Design Steps

The design project includes critical measurements at every step. This calls for solutions that include high-performance real-time and sampling oscilloscopes; Logic Analyzers; sophisticated probes and fixturing; and in some cases, specialized acquisition modules for TDR (time-domain reflectometry) and optical measurements. From a measurement perspective the steps of the design process proceed as shown in Figure 1-2.

After a quick review of signal integrity definitions, we will follow this sequence of steps to see how signal integrity measurements play an enabling role in the development of high-performance digital systems.

Review of Signal Integrity Concepts

The notion of signal integrity^{*2} pertains to noise, distortion and anomalies that can impair a signal in the analog domain. At frequencies in the gigahertz range, a host of variables can affect signal integrity: signal path design, impedances and loading, transmission line effects...even power distribution on the circuit board. It is the designer's responsibility to minimize such problems in the first place, and correct them when they appear.

There are two fundamental sources of signal degradation.

- Digital issues—typically timing-related. Bus contentions, setup and hold violations, metastability, and race conditions can cause erratic signal behavior on a bus or device output.
- Analog issues—low-amplitude signals, slow transition times, glitches, overshoot, crosstalk, and noise; these phenomena may have their origins in circuit board design or signal termination, but there are other causes as well.

Not surprisingly, there is a high degree of interaction and interdependence among digital and analog signal integrity issues. For example, a slow rise time on a gate input can cause the output pulse to be delayed, in turn causing a bus contention in the digital environment further downstream. A thorough solution for signal integrity measurement and troubleshooting involves both digital and analog tools.

*² The term "signal integrity" applies to the device under test. Another issue is "signal fidelity" which applies to the measurement equipment. Factors such as probe loading, measurement system bandwidth, and more determine signal fidelity. Meaningful signal integrity and compliance measurements require tools with good signal fidelity.

Primer



Impedance Measurements

The Foundation of Signal Integrity

Seasoned engineers know that signal integrity is the result of constant vigilance during the design process. It's all too easy for signal integrity problems to get compounded as a design evolves, and to become more difficult to track down. A tiny aberration that goes unnoticed in the first prototype board can bring the whole system to a crashing halt when the board is merged with others.

Given these realities, where does signal integrity begin? Designers working on the most critical high-speed technologies often start their signal integrity work at the very beginning—with the raw, unpopulated circuit board. Most high-speed protocols require a 50Ω impedance; for example PCI Express specifications call for a 50Ω transmission line with 10Ω tolerance. Analyzing transmission line principles is beyond the scope of this document, but suffice to say that the tolerances are critical to high-fidelity signal transmission. And less deviation is always better.

Modern layout tools implement the applicable impedance rules for high-speed protocols, but physics, circuit board materials, and human error can introduce unforeseen departures. As a result, many developers have learned that a rigorous process of verifying impedance characteristics can help them detect and correct problems early. Design choices can be reconsidered, if need be, before quantity orders are placed with a vendor.

Our server motherboard layout includes many connectors. Some are edge-connector types, such as those in the upper left quadrant of the board. Others are industrystandard cable connections for buses such as 10 Gb



Figure 2-1. Block diagram of a TDR acquisition setup using a TDR module

Ethernet CX4. A board's complement of connectors makes up a significant part of the materials costs involved in manufacturing the device. This simple fact prompts a constant search for lower-cost substitutes for components such as the PCI Express connectors. What impact might these proposed substitutes have on the high-speed signals that pass through them? Impedance measurements can help us decide whether the low-cost connectors will do the job.

The impedance measurement process will verify the traces and their vias, pads, and lands. At the same time, it will determine the effects of two brands of PCI Express connectors under consideration.

Overview of Impedance Measurement Tools and Techniques

The tool of choice for measuring impedances is a sampling oscilloscope such as the Tektronix CSA8000 or TDS8000 Series equipped with the 80E04 Time Domain Reflectometer (TDR) module. These are high-performance sampling oscilloscopes whose plug-in architecture enables them to implement both conventional sampling acquisition and time domain reflectometry applications. The TDR module permits the signal transmission environment be analyzed in the time domain, just as the signal integrity of live signals will be analyzed in the time domain.

Time domain reflectometry measures the reflections that result from a signal traveling through transmission environments such as circuit board traces, cables, or connectors. The TDR instrument sends a fast step pulse through the medium and displays the reflections from the observed transmission environment. Figure 2-1 is a simplified block diagram of this scheme. The TDR display is a voltage waveform that includes the incident step and the reflections from the transmission medium. The reflections increase or decrease the step amplitude depending on whether the nature of the discontinuity is more inductive or capacitive, respectively.

A reflection from an impedance discontinuity has rise time equal to or—more likely—slower than that of the incident step. The physical spacing of any two discontinuities in the circuit determines how closely their reflections will be positioned relative to one another on the TDR waveform. Two neighboring discontinuities may be indistinguishable to the measurement instrument if the distance between them amounts to less than half the system rise time.

The quality of the incident step pulse is critical, especially when measuring short traces. In addition to its fast risetime, the step must be accurate in terms of amplitude and free from aberrations. The 80E04 TDR module's incident risetime sets an industry benchmark in its class, and its amplitude accuracy and step response aberrations are state-of-the-art.

An Impedance Measurement Procedure and its Conclusions

The impedance measurement procedure that follows has three objectives:

- To detect any fundamental layout problems that are causing the impedance of a trace, via, or pad to deviate beyond its permissible tolerance; to find the physical location of such problem areas
- To compare and evaluate two brands of PCI Express connectors with the intent of qualifying the less expensive of the two for adoption into the design
- To ensure that the impedance environment is compliant with applicable industry standards for all of the serial buses on the motherboard, including PCI Express

Setting Up and Making Connections

Tektronix recommends calibrating the TDR before conducting PCB measurements. This is a process which the 8000 Series instruments can execute automatically when equipped with software available at no charge from Tektronix. The calibration procedure, developed on behalf of the IPC (Association Connecting Electronics Industries), is known as IPC-TM-650. Its intent is to ensure consistency and a known degree of traceability in PCB impedance measurement results. To summarize the procedure:

- Connect the TDR module to the same fixturing connectors, cables etc. – that will be used in the actual measurement procedures
- In place of the device under test, substitute a reference standard air line or PCB coupon, and perform a conventional TDR measurement
- Save the result as a Reference waveform

The steps above are simplified for the sake of clarity (certain calculations have been omitted), but the calibration software handles these additional tasks. Once a reference trace is stored, it becomes the template against which "return" signals are compared for amplitude differences over time.

The next step is to connect the 80E04 TDR module to the Device Under Test (DUT). Note again, the project objectives are threefold: verifying circuit traces, comparing PCI Express sockets, and compliance verification. The best way to introduce the incident step to a socket is through SMA connectors mounted on a fixture board that plugs into the socket. The same path(s) can also conduct the step to PCB traces that branch from the socket. Primer



Figure 2-2. True differential incident steps. The step and its complement are simultaneously launched.

True Differential TDR Measurements

Many of today's high-speed serial standards rely on differential transmission techniques using complementary signals. Two "wires" or PCB traces carry simultaneous mirror images of the signal. Though more complex than the single-ended approach, differential transmission is less vulnerable to external influences such as crosstalk and induced noise, and generates less of the same.

But differential paths require TDR measurements just as other transmission environments do. The incident pulse must be sent down both sides of the differential pair and the reflections measured. There are two ways this can be done:

The "virtual" or "calculated" differential TDR method. The TDR sends out two positive-going incident steps alternating in time, and corrects both the polarity and the alignment before the measurement waveform is displayed on the oscilloscope screen.

The true differential TDR method. The TDR sends out complementary signals that are accurately and correctly aligned in time. The DUT receives a differential stimulus signal more like those it will encounter in its end user application, potentially producing better insights into the device's real-world response. The TDR system does not need to manipulate the displayed step placements.



Figure 2-3. Serially-launched incident steps. The lower waveform is not a true "complement" of the step waveform but rather a mathematical inversion.

The calculated differential TDR method can provide valid results in many practical situations. However, the more modern true differential TDR method has gained broad acceptance among designers for several important reasons:

- The IPC-TM-650 impedance measurement standard (V2.5.5.7 or later) for printed circuit boards requires true differential measurements for appropriate protocols
- Some protocol standards stipulate that measurements be performed on powered links, which necessitates true differential signaling
- The calculations used to synthesize a differential display from non-differential signals will produce errors if the signal contains nonlinearities. Such errors are extremely difficult to detect.

The Tektronix 80E04 TDR Sampling Module is a true differential solution. Figure 2-2 illustrates the simultaneous (and complementary) incident pulses launched by the 80E04. Compare this with Figure 2-3, which depicts the serially-launched pulses of a calculation-based TDR (both illustrations represent an oscilloscope's view of the TDR signals). If the DUT requires powered measurements using differential transmission, then the serially-launched approach will not suffice.



Primer

Figure 2-4. TDR impedance measurement screen

Manual TDR probes and hardware positioners are also available but a dedicated SMA connection, when practicable, will always maintain the best signal environment. The threaded SMA connection provides excellent shielding and controlled impedance. This ensures that the fixturing itself will have minimal impact on the outgoing step or its returning reflections. Tektronix recommends tightening the SMA collar with a special wrench that applies the prescribed 5 inch-pounds (0.56 Nm) of torque. This provides the tightest "seal" on the signal path without risk of ruining the connector.

The PCI Express bus uses differential signaling. There are two methods of performing TDR measurements in a differential environment (see Sidebar). The 80E04 module is capable of true differential operation, and so will carry out the necessary TDR measurements without the need for special calculations.

TDR Measurements and Results

The TDR is connected to the specific PCI Express trace through an SMA fixture board plugged into Socket A. The incident step will see the effect of Sockets A and B, as well as the vias and the Ball Grid Array (BGA) pad at the end of the trace. With no termination present at the BGA pad, the TDR trace will drive toward an infinite impedance.

The TDR display tells a detailed story about the impedance variations in the signal path. Figure 2-4 is a TDR screen from the PCI Express measurement.



Figure 2-5. The TDR impedance display reveals how PCB features such as connectors and vias can cause deviations in the impedance environment.

Figure 2-5 is simplified sketch of the same result. The horizontal scale has been compressed to depict all of the events occurring over a fairly long trace. To see all these events in sufficient detail on the TDR itself, it would be necessary to expand horizontal (time) scale and scroll through the waveform record. The grey band across the top of the trace indicates the limits of the $\pm 5\%$ impedance tolerance set forth in the PCI Express standard.

The impedance display reads from left to right. The leftmost events are those physically closest to the step generator, the origin of the signal. Again, these are reflections produced by the elements along the signal path when they receive the energy from the incident step. A perfect path with perfect terminations would produce no reflections. But there is no such signal path in the real world, and it is the designer's job to understand and account for impedance variations.

Primer

At the top of the step transition in Figure 2-5, there is a momentary overshoot caused by the step generator itself. Shortly after the waveform settles, it sees the trace impedance, then it encounters the impedance effects of the first PCI Express connector—Event A in the illustration. This is the connector that is currently specified in the bill of materials. Note that its impedance falls well within the $\pm 5\%$ tolerance range.

But the next aberration is clearly out of bounds. Event B defines the character of the second PCI Express socket in the signal path. This is the low-cost connector being considered as a substitute for the one currently designed into the motherboard. Variations on this scale can affect signal integrity by attenuating the signal, distorting it, or even (if the impedance irregularity is mostly capacitive) giving it a path to adjacent traces.

Unfortunately the connector that produced Event B is not suitable for adoption in the high-performance motherboard. Its impedance deviation disqualifies it.

Continuing down the path, we see reflections from two vias, then the normal rise in impedance when the incident step hits the open circuit at the BGA pad. If there was an IC device installed on the BGA site, then it would be possible to characterize the impedance of the device pins and their internal bonding to the IC pads as well as the on-chip terminations. Alternatively, these measurements can be performed on the IC device separately from the PCB.

Repeating the TDR measurement process on other traces of interest will yield a detailed, high-confidence impedance profile of the raw PCB. "Up-front" work with the TDR can preempt potential signal integrity problems after assembly.

Conclusion

The TDR measurement has met all of its objectives. The low-cost connector has been evaluated and disqualified, potentially saving untold headaches in later design and production steps. And the PCI Express layout has been proven to meet the applicable standards for impedance.

Positive or Negative?

Even though the true differential TDR delivers an inverted (negative-going or complementary) pulse simultaneously with its positive counterpart, the resulting TDR impedance display shows a seemingly uninverted trace—all values on the trace are positive. Why?

Consider what the display is showing: impedance in ohms. It is comparable to measuring a potentiometer with digital multimeter. You expect a range of resistance values beginning at 0Ω and increasing positively toward some maximum. The same is true when using a TDR to observe a passive element such as a circuit board trace: the impedance is always a positive construct with the lowest possible value being near 0Ω . There are no negative ohms or Rho values.



ASIC Verification

Ensuring Signal Integrity at the Source

TDR measurements on the unpopulated motherboard are the first critical step in the hardware verification process. The active devices that make the raw PCB into a server motherboard are the next candidate for evaluation as the design moves ahead.

"Off-the-shelf" semiconductor devices are usually tested by their manufacturers, but many designers prefer to confirm these performance claims rather than risking an entire design on unproven components. More importantly, early samples of custom devices must be verified. Short-run ASIC prototypes (used to fabricate prototypes of the end user product itself) will be tested and ultimately the device tests might be released to production.

The verification process may include up to four types of tests:

- DC parametric—the most basic measurements to ensure that the device provides proper logic levels under loading and adequate current at each binary state and very low leakage in the high-impedance "compliant" state.
- Analog functional—in mixed-signal devices, this test verifies the analog features. Again, simple digital commands set up the device and execute the gain and range changes that are unique to analog components.
- Digital functional—a simple command set initializes the device in its various modes, then test vectors (binary data) are applied at a low data rate. The resulting outputs are compared with expected data.
- At-Speed functional—in effect, a digital functional test at the device's maximum clock and data rates.

The DC parametric and analog functional tests are beyond the scope of this signal integrity discussion. But the digital functional tests examine timing details that can impact signal quality.

Recommended solutions for ASIC evaluation include the CSA8200 or TDS8200 sampling oscilloscopes, or a realtime oscilloscope such as the TDS7404B or TDS6804B; TLA Series Logic Analyzers (for situations requiring a deep acquisition memory for digital data); and DTG5000 Data Timing Generators (signal sources).

The issue of signal access is a challenge when testing complex ASIC devices. Particularly in the case of the atspeed tests, fixturing must be designed to connect—with the least possible signal degradation—oscilloscopes and high-speed signal sources to a sequence of selected device pins. Most designers also characterize the fixture itself (using a TDR as explained in the previous section) to prevent impedance-related signal distortions.

It is common practice to design fixtures with relay matrices to distribute oscilloscope inputs and signal source outputs to respective pins on the DUT. Test programming environments such as National Instruments LabVIEW[®] provide an expedient means to automate measurements and control signal routing on the fixture. If properly designed, test programs developed for validation and evaluation can be ported to production test applications with minor modifications.

Primer



Figure 3-1. Setup and hold timing diagram

Digital Verification Steps Range from Basic to Advanced

"Setup" and "hold" timing values are essential specifications for any clocked digital device, and the heart of the ASIC functional verification measurements. Setup time is defined as the length of time the data must be present (and in a stable, valid state) before the clock edge occurs. Hold time defines the amount of time that data state must be maintained after the clock edge. In high-speed digital devices used for computing and communications, both values may be as low as a few hundred picoseconds. Setup and hold violations can cause signal integrity problems in the form of transients, edge aberrations, glitches and other intermittencies.

Figure 3-1 depicts a typical setup and hold timing diagram. In this example, the data envelope is narrower than the clock. This emphasizes the fact that when using today's high-speed logic, transition times and setup and hold values can be very brief even when the cyclical rate is relatively slow.

There are two common approaches to evaluating the setup and hold time performance of a device (other timing parameters also can be verified using these methods): Low-speed tests: Depending on the requirements foreseen in the end-user application, a fairly coarse functional verification procedure may be adequate. In some cases it is not necessary to take quantitative measurements of the actual setup and hold values. If the device will tolerate a broad clock placement range, the timing "test" may be as simple as running a low-speed functional data pattern, adjusting the clock edge's position relative to the data, and observing the results on an oscilloscope. The oscilloscope's trigger is set to a 50% level.

The device tends to become "metastable" as it exceeds its setup and hold timing limitations. Metastability is an unpredictable state in which the device output may switch to either "1" or "0" without regard to the logical input conditions. Similarly, excessive jitter may appear on the output when setup and hold tolerances are violated.

At-speed tests with "burst" data: The burst functional test exercises the device at rates approximating its intended operational frequency. A signal source such as the Tektronix DTG5000 Series Timing Generator can deliver a block of data to the DUT at rates much higher than the basic functional test just explained. The process is still one of empirically finding a range of setup and hold values and specifying the system's clock placement accordingly.

Using the data generator to drive all of the device's inputs, repetitive data patterns help isolate the recurring skew problems that cause repeated setup and hold violations.





▶ Figure 3-2. ASIC measurement system and fixture

Figure 3-3. The oscilloscope's Transition trigger can be set to detect edges are are taking too long to complete their transition. In this example the trigger has found an edge that takes more than 3.2 ns to switch from a high state to low state.

Figure 3-2 depicts an integrated ASIC evaluation setup. An oscilloscope works in conjunction with a Tektronix DTG5000 Series signal source that addresses the device via an automated fixture. A relay matrix, not visible in the photo, switches in DC and lower-speed signals for functional tests. Critical signal paths include impedance-controlled connection points for high-speed signals as well. Each of these is a potential test point to which the oscilloscope connects, depending on the needs of the specific device test.

The oscilloscope may be a sampling instrument if the device under test has data rates in the multi-gigabit range. Alternatively, a real-time Digital Storage Oscilloscope (DSO) or Digital Phosphor Oscilloscope (DPO) can offer powerful tools such as integrated jitter measurement software and a broad selection of specialized triggering functions. These capabilities are indispensable when it is necessary to scrutinize the signal's analog characteristics over/undershoot, edge details, pulse flatness, glitches, and more.

For example, suppose a setup time violation occurs occasionally on a device that has been proven to be comfortably within specifications. The error might even be pattern-dependent, that is, it occurs only when a particular data combination is present. The real-time oscilloscope can draw on a wealth of trigger types to isolate the problem. If a slow or distorted transition is suspected, for example, the instrument's Transition trigger can be set to detect the edge (either rising or falling) that is taking too long to make its excursion (Figure 3-3). This trigger condition can be combined with logic qualification or other timing variables. Similarly, jitter-related failures can be tracked down with the aid of jitter measurement software.

Conclusion

The ASIC evaluation step enables custom and programmable IC devices to be isolated from the system environment and qualified for installation in the emerging motherboard design. Moreover, it helps detect and avoid circumstances that cause metastability, jitter, and noise which can cause signal integrity problems that impact other system elements.

▶ Primer



Basic Functional Verification Checking the "Heartbeat" of an

Emerging Design

With PCB impedance measurements and ASIC verification completed, it is time to assemble a prototype motherboard and begin "live" measurements. The next few steps in the design process are where the majority of functional problems emerge and must be solved. Many of these problems stem from signal integrity issues that are impossible to detect until all of the board's components are in place and working together.

As explained earlier, escalating data rates are at the heart of many signal integrity problems in today's designs. In most modern bus protocols, the data rate is the basic currency of a device's throughput and efficiency. Raw bit rates extend from about 1 Gb/s to 4.25 Gb/s depending on the protocol, and multi-lane topology is often used to increase net throughput. Some standards use a single lane; for FibreChannel 4X and Serial ATA, the data throughput is equal to the signaling rate. Multi-lane configurations scale their throughput in direct proportion to the number of lanes used. InfiniBand, PCI Express, and XAUI all have 4X variants that provide good optimization for 10 Gb/s, while InfiniBand and PCI Express carry the concept even further. PCI Express offers up to 32 lanes (80 Gb/s) on the circuit board and 16 lanes (40 Gb/s) at the connector. Note that the raw data is not the same as data transfer rate, which is expressed in gigabytes per second (GB/s). Since data transfer rate depends so much on processes occurring in higher layers of the architecture, a physical-layer comparison is not useful.

Verification Tools

For the verification and troubleshooting steps in this part of the design process, the real-time DSO or the DPO are the tools of choice. Both types of instruments offer not only the necessary performance (bandwidth, sample rate, etc.), but also a wealth of triggering choices, probing options, application-specific software packages, and more. Most importantly, these real-time platforms make it easy to probe a series of test points and reliably acquire waveforms ranging from power supply noise to multi-GHz data streams. Automated setup routines find, scale, and display the signal, while cursors and built-in automated measurements simplify analysis.

The DSO is ideal for low- or high-repetition-rate signals with fast edges or narrow pulse widths. The DSO also excels at capturing single-shot events and transients, and it is the best solution for multi-channel acquisition at high bandwidths and sample rates. For the motherboard design project, Tektronix' TDS6000B Series DSOs are an appropriate solution.

The DPO is the right tool for digital troubleshooting, for finding intermittent signals, and for many types of eye diagram and mask testing. The DPO's extraordinary waveform capture rate overlays sweep after sweep of information more quickly than any other oscilloscope, presenting frequency-of-occurrence details—in color—with unmatched clarity. The Tektronix TDS7000 Series exemplifies a DPO solution that meets the exacting needs of digital system design applications.

In any high-speed measurement, the choice of probes is critical. The oscilloscope and the probe work together as a measurement system. If possible the probe should provide the same bandwidth as the oscilloscope itself. In addition to its absolute bandwidth rating, the probe should have the least possible loading impact on the signal. Many serial buses require differential signal paths which are best served by true differential probes.

Primer



 Figure 4-1. Differential Probe (foreground shows handheld tip assembly)

Ideally the measurement system bandwidth, including that of the probe, should be at least three times (3X) the frequency of the signals to be observed. That equates to about 7.5 GHz for the 2.5 Gb/s data rates used on the new motherboard. The probe most suitable for this application is the Tektronix P7380 probe (Figure 4-1), an 8 GHz true differential probe whose risetime, only 35 ps (20% to 80%) meets the needs of fast serial signals. Its ultra-low loading and diversity of attachment methods ensure fast, positive connection with minimal effect on signals. The tiny probe tip makes it easy to reach into narrowly-spaced traces and device pins.

Protocol/Bit Rate (per lane)	Oscilloscope	Probe
PCI Express 2.5 Gb/s	• TDS6000B with opt. RTE, PCE	• P7350SMA/P7380SMA
• Serial ATA II 3.0 Gb/s	• TDS7704B with TDSJIT3 (opt. JT3), Serial ATA Test SW	• P7380SMA
Serial RapidlO 3.125 Gb/s	TDS7704B with TDSJIT3 (opt. JT3)	 P7260, P7240, P7330, P7350
 Infiniband 2.5 Gb/s 	TDS6000B with opt. RTE, IBA	• P7350SMA/P7380SMA
• IEEE 1394 (Firewire) 3.2 Gb/s	TDS7704B with TDSJIT3 (opt. JT3)	 P7260, P7240, P7330, P7350
• 10 GbE XAUI 3.125 Gb/s	TDS6804B with opt RTE, TDSJIT3 (opt. JT3)	• P7380SMA, P7260
• FibreChannel 2.125 Gb/s (2X) 4.25 Gb/s (4X)	 TDS7704B with opt. RTE, SM, TDSJIT3 (opt. JT3) CSA/TDS8200, 80A03, 80E02/3/4, A0A05 	• P7380SMA, P7260
• USB 2.0 480 Mb/s	• TDS7404B with opt. USB	 P7260, P6248, P6247, P6245 (3), TCA-BNC, TCP202

Table 4-1. Summary of high-speed protocols and measurement solutions

Table 4-1 lists some representative high-speed buses and devices, showing the oscilloscope and probe combinations that match their technical requirements.

The oscilloscopes and probes described here will be used in all of the design verification and compliance steps to follow. Primer

Checking the Pulse of a New Design

The motherboard prototype is assembled. The board has passed its visual inspection for common hardware problems such as solder bridges and misplaced components. The system BIOS chip is installed. External power supplies are connected and powered up. Subsequent steps in the process will track down code-level operational problems and their sources, but a preliminary series of checks must confirm the presence, timing, and integrity of the basic signals that support the board's functionality.

In the past, the "heartbeat" checks on prototype devices were straightforward. But all that changed when high-speed transmission became the norm, and serial buses introduced challenges such as encoded signaling and embedded clocks.

With an embedded clock, there is no clock trace to probe. The clock signal is woven into the data stream, and special measurement tools are required to extract the clock information. Some oscilloscopes such as the Tektronix TDS6000B or TDS7000B Series have clock extraction features built in. These features can be supplemented with powerful software tools to automate measurements on specific serial protocols, or to measure jitter attributes.

The Tektronix TDSJIT3 jitter measurement toolset is a valuable asset when working with Spread Spectrum Clock (SSC) technology. SSC minimizes radiated electromagnetic interference (EMI), and is gaining favor as a solution that can forestall problems during EMI compliance testing. Because SSC is new to many digital system designers, it merits special attention during the prototype checkout.

Embedded SSC resists conventional measurement approaches. Yet it is essential to ensure that the embedded SSC clock is present and operating within specifications.

SSC Example: The Serial ATA Clock

Spread spectrum clocking intentionally varies the frequency of the clock signal to distribute its energy over many frequencies. In contrast, jitter can be defined as an unintended timing variation. But the same tool, the Tektronix TDS Series application package known as TDSJIT3, can distinguish either type. We can take advantage of this capability to observe the embedded spread spectrum clock and confirm that its characteristics and transitions (Off -On and On-Off) comply with industry standards.

TDSJIT3 software provides a comprehensive set of measurements for jitter parameters. These terms include cycle-to-cycle jitter, time interval error (TIE), and more. The TDSJIT3 feature best suited for our SSC measurement is the Time Trend plot. This plot displays the jitter value versus time to reveal repeating patterns such as modulation, various periodic frequency components, and... spread spectrum frequency variations.

The object of this example is to evaluate the SSC functionality on a Serial ATA II (SATA II) bus. The embedded clock runs at 1.5 GHz and varies by less than 3% over a 33.3us interval. To capture the signal, we will use the P7380 differential probe's variable-spacing tip to probe near the SATA connector.

Ten cycles of the 33 kHz modulation frequency on the 1.5 Gb/s data must be captured. The Tektronix TDS6804B oscilloscope will perform the acquisition. It is among the few available solutions that can capture the necessary 330 uS of signal data at its full 20 GS/s sample rate.

Once the signal data is acquired, TDSJIT3 establishes a reference clock from the data signal. The result is a jitterfree clock with the same mean frequency and phase as the data. Then TDSJIT3 automatically plots the changes in the recovered clock period.



 Figure 4-2. Modulation profile of a SATA spread spectrum clocking scheme

Spread Spectrum Clocking and Serial ATA

In most parts of the world, electronic products sold commercially must meet radiated emissions limits. High-frequency digital operations, particularly clock signals, tend to emit energy that can leave the "box" and act like stray RF signals in the surrounding environment. These signals can cause trouble in other equipment ranging from cellular phones to heart pacemakers.

Radiated emissions from the clock signal can be reduced by modulating the clock frequency using a technique known as Spread Spectrum Clocking (SSC). Just a tiny amount of modulation—as low as a few picoseconds—is all that is necessary to minimize emission problems. Many serial buses, notably PCI Express and Serial ATA, rely on SSC techniques to control radiation.

The characteristics of the SSC implementation are precisely spelled out in industry specifications. Ranges and tolerances are critical, since the clock affects every transaction on the bus. For Serial ATA (SATA) for example, standard V1.0a outlines a group of SSC requirements, among them:

- When SSC is employed, all device timings (including jitter, skew, clock period, output rise/fall times, etc.) must meet the existing non-SSC specifications
- The minimum Unit Interval (UI) time is 666.43 ps; the maximum is 670.12 ps
- The preferred modulation method is "downspreading;" adjusting the spread technique to preclude modulation above the nominal frequency.



- Figure 4-3. The SSC clock period increases in proportion to the modulation amount (dotted line).
- For triangular modulation, the clock frequency deviation must be down-spread and shifted no more than 0.5% from the nominal frequency; that is, +0%/-0.5%.
- The modulation frequency of the SSC must be in the range of 30-33 KHz

Figure 4-2 depicts the triangular SSC modulation profile defined in the SATA I V1.0a specification. Here, fnom is the nominal clock frequency (without SSC), fm is the modulation frequency, and δ is the modulation amount. The horizontal axis expresses time.

To understand how the SSC modulation operates on the clock frequency, see Figure 4-3. For the sake of clarity, this scenario shows the clock period unchanging for several cycles. Interval T1 is the same as T2, T3, etc. up to the point at which the modulation envelope begins to turn upward on the graph. Then the clock period begins to get longer with each successive cycle. Mid-way through the excursion, the period is Tn, that is, some value greater than T1, T2, and so on. Eventually the degree of modulation peaks and begins to turn down, following the modulation profile just explained.

The small amount of SSC modulation is virtually impossible to observe in real time with a conventional oscilloscope measurement. To analyze the behavior of an SSC-modulated signal, it is necessary to capture and store at least ten cycles' worth of modulation activity, recover the embedded clock, then plot the change in the clock period over time. The analysis features in the latest TDSJIT3 toolset are indispensable for this task. Primer



Figure 4-4. Modulation profile for a spread-spectrum clock. The vertical axis represents the amount of modulation being applied.

Figure 4-4 depicts a TDSJIT3 record showing the SSC clock modulation profile observed on the SATA bus. Note that this is not a waveform! It is a plot showing frequency deviation in the signal mapped over the allowable range of unit interval tolerances. The baseline value is the minimum UI: 666.4 ps. The peak excursion approaches—but does not exceed—the maximum UI, that is, 670.1 ps. The period of the 1.5 GHz clock is changing by less than 4ps. TDSJIT3 has extracted the needed information from a complex, fast-changing signal.

In Figure 4-4, the modulation profile accurately follows the guideline for a down-spread clock signal. This confirms that 1) the SSC implementation is applying modulation correctly and within specifications, and 2) the transition between states (Off-to-On in this case) is clean and free of aberrations. Note the similarity between this plot and Figure 4-2, the profile taken from the SATA V1.0a specification. A second acquisition is required to capture the On-to-Off transition.

Design or layout errors can have a radical impact on the performance of the SSC implementation. Fortunately, these too can be captured with TDSJIT3. Figure 4-5 illustrates a spread spectrum clock signal that has a large transient at the beginning of the Off-to-On transition. Moreover, the



Figure 4-5. This spread-spectrum modulation profile traces the effect of a transient.

modulation amount is well below the range allowed in the specification (this may be a deliberate choice). The transient alone is enough to disturb operations on the SATA bus. The aberration may have its origins in the SSC design itself, or it may be an artifact of other system behaviors such as ground bounce or crosstalk. The low amplitude of the modulation implies a flaw in the SSC design. In any event, a problem in the modulation profile has been identified by the TDSJIT3 application, and the problem must be resolved before proceeding with the compliance and functional evaluation steps to come.

Conclusion

The TDSJIT3 jitter analysis application has confirmed that the embedded SSC is functioning correctly. Other measurements, performed on key control signal lines using conventional point-to-point probing, prove that the motherboard is ready to proceed with further operational and compliance testing.

Until recently, heartbeat checks consisted of little more than verifying the presence of a clock signal. Today we must extricate the clock from high-speed data and examine its static and dynamic (modulation) behavior. Fortunately, this increased complexity has been met by tools that rely on built-in application expertise to deliver results quickly and accurately.

Functional Verification	& Troubleshooting
-------------------------	-------------------

Basic Functional Verification	Internal Bus Compliance	Operational Validation, Fault Detection, Debug
-------------------------------------	----------------------------	---

Internal Bus Compliance

Compliance Tests Pave the Way to Interoperability

With the server motherboard assembled and its power supplies and clocks verified, internal buses such as SATA II are the next order of business. These devices must be checked not only for functionality but also for their compliance with applicable industry standards. It is the only way to ensure interoperability with the other subsystems that will reside within the server. When a hard disk drive array, for example, is integrated with the motherboard, we want to be as confident as possible that the motherboard is delivering signals with the proper characteristics. Signal integrity plays a key role, affecting the eye diagram measurements that are the basis of most compliance tests. In addition, some signal integrity factors (for example, jitter) may have their origins on the motherboard and must be controlled at their source.

Like the other measurement steps in the motherboard development process, verifying compliance on the board's internal serial buses requires rigorous probing, acquisition, and analysis methods. The compliance tests in this section will focus on the Serial ATA II bus, which is the transmission channel to the externally-mounted hard drives serving the motherboard.

Probing Requirements for Compliance Measurements

Every measurement begins with a reliable contact to the device under test, and compliance measurements are no exception. High-performance oscilloscopes suitable for compliance work can be outfitted with a range of probe types. These include differential probes, various highimpedance FET probes, and various modular "browsing" formats with adapters for diverse contact situations. Not only electrical requirements, but also physical realities, guide the choice of a probe for the internal bus compliance tests.

Probing: Electrical Issues

Most serial buses transmit signals differentially, therefore compliance tests must capture signals differentially. One approach is to use a pair of single-ended probes feeding two oscilloscope inputs, then process the incoming signals with a math operation. Most oscilloscopes can do this. However this method consumes two inputs, meaning that a four-channel oscilloscope can monitor only two active signals (two differential pairs) at a time. In addition, the math operation precludes some types of measurements and eye diagrams.

In contrast, a true differential probe combines the two branches of the differential signal outside the oscilloscope, delivering one single-ended signal per probe to the acquisition system. Skew problems are minimized and the instrument's full channel count is available.

Other electrical considerations include bandwidth, of course; like the oscilloscope itself, the probe's bandwidth and edge response (risetime) must meet the needs of the DUT. As a rule of thumb, the measurement system bandwidth should be at least 1.8 times the data rate that of the device being observed, equivalent to 3.6 times the fundamental frequency of data in the non-return to zero format. The low-voltage differential signaling (LVDS), Current Mode Logic (CML), and Positive Emitter Coupled Logic (PECL) signals common in many serial buses require a low-capacitance, low loading probe. Both DC resistance and AC impedance must be maximized in order to provide a stable, very low current burden on the signal.

Primer



Figure 5-1. This P7380 differential probe includes adapters that aid access to hard-to-reach PCB features. Some types of leads can be soldered down for the most secure contact

Probing: Physical Considerations

The choice of probes is also affected by the physical accessibility the signals of interest.

In high-speed designs such as our motherboard, it is common practice to route critical signals to impedancecontrolled coaxial test points, usually SMA high-frequency connectors, for the most positive connection to instrumentation. This is particularly useful during prototype development. A probe designed to mate directly to such test points without intervening tips or jumpers is the preferred solution for this application. An example of this approach is the Tektronix P7380SMA probe, an 8 GHz tool aimed at high-speed serial measurements and other digital signal applications. The P7380SMA is commonly used in applications which can afford to dedicate PCB real estate to test connectors.

In contrast, when the signal is available only on the pins of a fine-pitch surface-mount device (SMD) or on a via in the midst of densely packed traces, a probe with very small, flexible contact points is best for the job. An example is the Tektronix P7380 probe (Figure 5-1), which uses a tiny passive probe tip element, separate from the probe's internal amplifier circuit, to extend the usable reach of the probe. Its distributed attenuator topology preserves highfrequency signal fidelity and edge response characteristics.



Figure 5-2. The P7380SMA probe connects to the SATA bus through an extender plugged into the device's output connector

Making the Connection

On our motherboard, the SATA II bus goes to the edge of the PCB, where a connector meets the cable that carries the signals to a hard disk drive. Following the process explained earlier in this document, both the PCB signal path and the connector have been pre-screened with a TDR. The connector provides an ideal access point for our compliance measurements.

A simple "extender" makes it possible to connect the oscilloscope to the signal path via the motherboard's SATA II connector. The extender is a small PCB mounted with a pair of SMA connectors. Each SMA connector brings out one side of the differential signal. The extender plugs directly into the SATA II connector using the pinout specified in the applicable industry standard. Figure 5-2 depicts such a fixture in a typical compliance measurement setting.

When the P7380SMA differential probe is connected to the fixture, the whole apparatus appears as a properly terminated load for the signal. This ensures an optimal, controlled environment for compliance tests. Moreover, the connection is positive and stable, a requirement when the test plan includes multiple measurements performed sequentially.

Some probes, notably the P7380SMA, provide the flexibility of an adjustable termination voltage. The termination voltage can be adjusted to match the input signal common mode voltage for minimum DC loading on the input signal. In addition the voltage can also to varied to stress-test the input signal driver during the evaluation process.

Compliance Tests Look at Signals from Many Perspectives

All serial standards include amplitude, timing, jitter, and eye diagram measurements within their compliance testing specifications. Of course, standards differ from one to the next, and not all measurements are required for compliance with every standard. Following is an overview of some key compliance measurements.

- Amplitude tests ensure that the signal has enough amplitude tolerance to do its job under worst case conditions
- Differential Voltage tests confirm that a specific minimum differential voltage will arrive at the receiver under worstcase media conditions (maximum loss). This ensures proper data transfer.
- Common Mode Voltage Measurements detect any common mode imbalances and noise, and help locate crosstalk and noise effects that may be coupling into one side of the differential pair and not the other.
- Waveform Eye Height tests characterize the data eye opening in the amplitude domain. It is measured at the .5 Unit Interval (UI) point, where the UI timing reference is defined by the recovered clock.
- Timing Measurements detect aberrations and signal degradation that arise from distributed capacitance, crosstalk, and other causes that affect signal integrity.
- Unit Interval and Bit Rate measurements look at variations in the embedded clock frequency over a large number of consecutive cycles.

- Rise/Fall Time measurements confirm that the transition times are within acceptable limits. Edges that are too fast can cause EMI issues, while those that are too slow rise can cause data errors.
- Waveform Eye Width measurements can verify the general "health" of the signal if accompanied by statistical details about the number of edges used in the measurement.
- Jitter measurements to determine factors such as random jitter, deterministic jitter, and total jitter (at a specific Bit Error Rate) are required in many serial bus specifications.

This is a partial list of compliance measurements. Full compliance tests may include 50 individual parameters or more, though many of these are extracted from the basic eye diagram acquisition.

Acquisition, Then Analysis

With today's tools, the actual measurements might just be the easiest part of compliance testing. A case in point is the TDS RT-Eye[™] compliance measurement package, which transforms a host TDS Series oscilloscope into a fullfeatured automated compliance tester. Once the hardware connections to the DUT are in place and the basic test requirements selected from a menu, the eye diagram portion of the compliance test procedure becomes a "onebutton" operation. This is the procedure we will follow for the compliance tests on the motherboard project.

The application's first step toward creating an eye diagram is recovering the clock from the serial bit stream. Methods range from Phase Lock Loop (PLL) recovery to oversampling; the Serial ATA II standard calls for the PLL method. TDS RT-Eye[™] package, a general-purpose toolset, can be equipped with a Serial ATA module that is preconfigured with this clock extraction scheme. The application's flexible software-based clock recovery approach enables it to meet the needs of a wide range of standards and their variants.

Primer

The measurement application automatically sets up the trigger points, waveform scaling, and performs the clock extraction. TDSRT-Eye™ "knows" about details of the specific standard: spread spectrum clocking, pre/de-emphasis, and other variables that affect the acquisition. These would be very complex to program into the oscilloscope manually.

Serial ATA II requires that compliance-related jitter measurements be performed on a specified number of consecutive (contiguous) bits. This requirement can only be satisfied by a real-time oscilloscope; a sampling instrument with its equivalent-time acquisition can not capture the necessary consecutive cycles except with special repeating test patterns designed for the purpose.

The TDSRT-Eye[™] application software deals with the setup and acquisition. The tool derives eye diagrams and various industry-specified jitter parameters from a single real -time waveform acquisition. In critical applications, additional confidence can be gained by accumulating statistics over multiple acquisitions.

To perform real time jitter measurements such as Random Jitter (RJ), Deterministic Jitter (DJ), and Total Jitter (TJ) at a 10-12 bit error rate (one part per trillion), a jitter test pattern length must be specified. This is a decision left to the user, who may have proprietary guidelines to meet in addition to industry compliance specifications. Popular jitter test patterns include TS1, CJTPAT, CSPAT, CRPAT, and more; some measurement packages offer these as an integral part of the toolset.

The TDSRT-Eye[™] application is known for its versatile software clock recovery algorithm. It is equally at home with a PLL methodology or phase interpolation, to name just two examples. It encompasses many serial standards, and importantly, it can adapt to evolutionary changes in those standards. As part of the TDSRT-Eye[™] analysis application, the software clock recovery algorithm is designed to support post-acquisition analysis and reporting.

But what if there is a need for troubleshooting the bus, probing and observing waveforms in real time? Software clock recovery is less suitable for this purpose. "Touchdown" probing applications are better served by a hardware PLL clock recovery architecture such as that built into the Option SM toolset for TDS Series oscilloscopes. With this feature, each time the probe contacts the signal, the clock



▶ Figure 5-3. Serial ATA II compliance measurement as displayed by the automated TDSRT-Eye measurement package.

recovery is less flexible than its software counterpart, it is the right solution for fast debug work.

Figure 5-3 presents the results of the SATA II compliance test. The eye diagram appears to be well within the mask tolerances, and the quantitative measurements listed in the Results Summary table confirm that. The measurement application has compared its actual acquired values with the limits drawn from the published SATA II standard These terms reside within the application. All of the relevant parameters have passed the compliance test.

The Serial ATA II bus is just one of several serial buses on the motherboard, but the compliance test procedures for all of them are similar. When TDSRT-Eye[™] has checked all the buses for compliance, the device is ready to move on to the functional testing and debug step.

Conclusion

Compliance testing is not optional—it is mandated under industry standards and expected by end-users. Eye diagrams are the most important tool for most compliance tests. Smart automated software tools integrated into highperformance oscilloscopes minimize the burden of eye diagram measurements. They provide easily-documented proof of compliance and can reveal hidden problems due to noise, jitter, aberrations, and other signal integrity issues.



Figure 5-4. Eye diagram. The blue areas are the mask violation zones.

Eye Diagrams: The Cornerstone of Compliance Measurements

The eye diagram has become the definitive tool for validation and compliance testing of digitallytransmitted signals. It is a display, typically viewed on an oscilloscope, that quickly reveals impairments in a high-speed digital signal.

The eye diagram is built up by overlaying the waveform traces from many successive unit intervals (UI). Eye diagrams display serial data with respect to a clock recovered from the data signal using either hardware or software tools. In the illustration shown here, the clock was recovered by a hardware-based "Golden PLL." The diagram displays all possible transitions (edges), positive-going and negative-going, and both data states in a single window. The result is an image that (using some imagination) resembles an eye, as shown in Figure 5-4.

In an ideal world, each new trace would line up perfectly on top of those that came before it. In the real world, signal integrity factors such as noise and jitter cause the composite trace to "blur" as it accumulates

The blue regions in Figure 5-4 have special significance. They are the violation zones used as mask boundaries during compliance testing. The blue polygon in the center defines the area in which the eye is widest. This encompasses the range of safe decision points for extracting the data content (the binary state) from the demodulated signal. The upper and lower blue bars define the signal's amplitude limits.

If a signal peak penetrates the upper bar, for instance, it is considered a "mask hit" that presumably will cause the compliance test to fail (some standards may tolerate a small number of mask hits). More commonly noise, distortion, transients, or jitter cause the traces to thicken. The eye opening "shrinks," touching the inner blue polygon. This too is a compliance failure, since it reveals an intrusion into the area reserved for evaluating the state of the data bit.

The eye diagram's compelling advantage is that it enables a quick visual assessment of the signal's quality. The information-rich display looks like it might be a challenge to set up and acquire. But modern digitizing oscilloscopes such as those in the Tektronix TDS family can be optioned with tools that expedite the complex clock recovery, triggering, and scaling automatically, then perform quantitative measurements on the data. With these software and hardware facilities, the eye diagram measurement has become a "one-button" operation.

▶ Primer



Operational Validation, Fault Detection, and Debug

Looking at Digital Data and Analog Interactions

Everything is in readiness: clocks checked, ASICs verified, compliance proven. Now begins the operational validation step, the first look at the integrated features of the motherboard. For the first time we will see whether the system will accept, transfer, process, and store information as intended. The proof of the design lies in its ability to handle instructions and deliver the expected response.

A design as complex as a server motherboard is almost certain to encounter some errors during development. These may include incorrectly placed components, logic design problems, improper terminations, and more. A bus channel or memory signal whose analog characteristics risetime, amplitude—are flawed can cause an error in a high-level system instruction. Conversely, fast-changing digital switching on a bus can impact the analog behavior of signals on nearby traces. The engineer is left with the challenge of detecting and correcting issues like these in the least possible time.

It is a matter of quickly identifying the real cause of the problem. Analog or digital? It means choosing tools and troubleshooting methods that address both domains. The favored solution for such challenges is a pairing of familiar instruments: a high-bandwidth DSO or DPO and a Logic Analyzer. The DSO is the best tool for observing individual events such as glitches, as well as distortions, transition times and critical setup and hold timing values. The logic analyzer captures logic signals in their elemental form—binary values with associated timing information—as they



Figure 6-1. The iLink[™] toolset includes the iCapture[™] acquisition feature, which allows the logic analyzer to acquire both digital and analog signals simultaneously through a single probe; and the iView[™] integrated analog digital display capability, which brings time-correlated analog waveforms to the logic analyzer screen.

move through the system. Capturing the interaction between the two domains is the key to efficient troubleshooting.

Some modern solutions, notably the Tektronix Series Logic Analyzers and the TDS Series DPOs include features to integrate the two instruments, sharing triggers and timecorrelated displays. Figure 6-1 diagrams this approach. In this chapter we will see how these instruments, working together, can drill down to low-level design problems.

Logic Analyzer is the Right Tool for Digital Data Analysis

The Logic Analyzer is indispensable for digital signal validation and troubleshooting. Equipped with accessories ranging from probes and fixturing to bus support packages and software disassemblers, the Logic Analyzer captures information on a bus and displays both timing diagrams and state information. It can trigger on a particular data word or an error such as a setup and hold violation. Alternatively it can receive an external trigger and record synchronized cycle-by-cycle state data.



Figure 6-2. The connectorless 32-channel probe conducts both digital and analog signals to the logic analyzer. The unused PCB test pads next to the probe accept a second probe.

Making the Connection

It is common practice to design in dedicated test instrument connectors on high-performance digital PCBs. Our motherboard includes pads designed for the Tektronix P6980 compression probe. The P6980 dispenses with costly on-board test connectors, mating instead to the circuit boards pads via cSpring (cLGA) connections (Figure 6-2). Equally important, it is designed to deliver selected analog signals to an oscilloscope connected to the logic analyzer via the TLA iCapture[™] multiplexer. At the same time, it drives digital signals to the Logic Analyzer. With the motherboard powered up and running, the Logic Analyzer acquires data from the test points. Typically this data takes the form of hexadecimal machine code, or it may be correlated to the high-level instructions that created the code. Alternatively, the Logic Analyzers "timing" mode displays binary waveform representations. In any event, data accumulates in the Logic Analyzers memory, which may be as deep as 256.

In the example application, assume that two separate buses—"A3," a four-bit bus, and "A2," an eight-bit bus must be verified. A logic analyzer probe is connected to the test points as shown in Figure 6-2. Initially the two buses seem to be transferring data just as the motherboard design models predicted.

But after a few cycles of operation, errors begin to appear on the A3 bus. A value that should be "8" shows up as a "0" not just once, but repeatedly. It is not a "stuck bit" or misrouted signal; that would cause the same error to occur continuously. The erratic nature of the problem implies some intermittent event is being mistaken for a legitimate data bit, altering the value of the hexadecimal results. The repeating nature of the problem points toward a glitch that is caused by an error in the layout or assembly of the prototype motherboard. And the hexadecimal value, 8 instead of 0, implies a problem in the most significant bit of the A3 bus (Note: the logic analyzer detects errors on other buses, but this discussion will focus on a particular issue affecting A3 for the sake of clarity).

What are the circumstances under which the A3 bus error appears?

▶ Primer



Figure 6-3. The LA has triggered on the glitch and flagged the individual glitch locations.

Tracing the Error

The true nature of the error begins to reveal itself after an acquisition with the logic analyzer's glitch capture trigger and display, which activates when it detects glitches, then flags their location in the timing display. The Logic Analyzer defines a "glitch" as an occurrence of more than one signal transition between sample points. Figure 6-3 depicts the resulting display on a TLA Series Logic Analyzer.

Note that there are two types of "waveforms" displayed in Figure 6-3. Each of the two busses, A2 and A3, is summarized with a bus waveform that reflects the word value on the respective bus. Bus waveforms provide an at-a-glance indication of the state of many individual signals, saving time when troubleshooting. In addition the display can be configured to break out each individual signal line and again flag the glitch locations. In Figure 6-3, the period between clocks is 4.00 ns. In TLA Series Logic Analyzers, the MagniVu[™] acquisition captures the signal values at 125 ps intervals, and this information can be displayed as a separate, high-resolution view of any signal. This feature acquires the high-resolution data at the same time as the main timing data, through the same probe.

Figure 6-4 illustrates the display with the MagniVu[™] acquisition traces added. Here, both the 125 ps clock ticks and the more detailed view of signal A3-3 are shown along with the bus waveform views. The signal shows a brief transition in the latter half of the cycle. Since we already know that the cycle is producing an incorrect bus value, this transition is likely the cause of the error. But what is causing the invalid transition?

Prime	r
-------	---

LA 1 AJ CI	0	C2 6	Deta		0.00				
LA 1: Sample	-12.000 m								20.000 m
LA 1: A3	•	X		X		X	2	X	
LA 1: A2	61	X				00			
LA 1: A3(3)				-					
A 1: Magnil/u Sample	State Manual	diameter and							
LA 1. HageWu: A3(3)						an of the part of the			

▶ Figure 6-4. The logic analyzer's MagniVu[™] acquisition display reveals a glitch in the A3(3) signal

Taking the Analog Perspective

Frequently, digital signal aberrations arise from analog signal integrity problems. The iLink™ toolset shown earlier (Figure 6-1) makes it easy to explore the analog characteristics of digital anomalies. The TLA iCapture™ multiplexer delivers any four signals from the P6980 probe to an attached TDS Series oscilloscope via an analog multiplexer inside the Logic Analyzer. Because it provides a path for both analog and digital signals, the probe eliminates double-probing and the associated double-loading of the device signals. Another feature of the iLink[™] toolset, iView[™] measurements, displays the resulting time-correlated digital and analog waveforms on the logic analyzer screen.

Figure 6-5 shows the TLA Series display that results when analog signal A3(3) is aligned onscreen with its digital equivalent. It is a picture that tells the whole story: at the exact moment of the digital glitch, the analog signal's amplitude is degraded in the area of the logic threshold. It apparently dips below the threshold voltage for an instant, creating a momentary "low," or logic 0 level. Then it increases just enough to cross above the threshold and return to the "high" or logic 1 level before switching to logic 0 again at the cycle boundary.

This analog behavior is the origin of the glitch, and for the error in the hexadecimal output on the bus. The instability is such that it does not affect every falling edge the same way; many pulses pass without errors. Of course, it will be necessary to review the design models determine when the valid edge should occur; before or after the unstable portion of the waveform in this bus cycle.

Primer



▶ Figure 6-5. iView™ measurements show the analog behavior underlying the digital glitch on A3(3)

The experienced engineer will recognize clues in this distorted waveform. A degraded logic level such as this is usually the result of a reflection coming back from an improperly terminated transmission line. In the case of our motherboard, the signal's fast edges encountered a missing termination resistor at the signal's destination. The result is an erratic but damaging erosion of the falling edge.

To summarize, troubleshooting with the logic analyzer/oscilloscope combination is a matter of proceeding from a high-level, global view to a zoomed-in closeup of individual signals using the four signal formats available on an instrument equipped with the iLink[™] toolset:

- The bus waveform gives an at-a-glance indication of problems occurring somewhere on the bus.
- The deep timing waveform reveals exactly which signal line is involved.
- The high-resolution MagniVu[™] timing waveform pinpoints the time placement of the error to a resolution of 125 ps.
- The analog waveform, provided by the DSO connected via the iCapture[™] multiplexer and the iLink[™] toolset, captures the specific analog characteristics of the signal, revealing potential causes.



► Figure 6-6. iVerify[™] analysis uses the eye diagram format to display multiply signals at once.

A Shortcut that Detects Signal Integrity Problems Quickly

There is an additional troubleshooting tool available to logic analyzers equipped with the iLink[™] toolset: iVerify[™] analysis, which bring analog eye diagram analysis to the logic analyzer screen.

The eye diagram is a visual tool to observe the data valid window and general signal integrity on clocked buses. It is a required compliance testing tool for many of today's buses, particularly serial types, but any signal line can be viewed as an eye diagram.

iVerify[™] analysis speeds troubleshooting by incorporating up to hundreds of eye diagrams into one view that encompasses the leading and trailing edges of both positive-going and negative-going pulses. Figure 6-6 depicts the result. Here, 12 (twelve) signals—the entire A3 and A2 buses—are superimposed. The benefit of "twelve at once" is clear—And is easy extended in a world of 32-bit and 64-bit buses, every shortcut helps. Any group of signals currently connected to the P6980 probe can be selected.

Because an eye diagram presents all possible logic transitions in a single view, it can also provide a fast assessment of a signal's health. It reveals analog problems such as slow risetimes, transients, attenuated levels, and more. Some engineers start their evaluation by looking at first at the eye diagrams, then track down any aberrations.

Primer



▶ Figure 6-7. The Nerify™ analysis tool brings the errant waveform to the front and highlights it for easy evaluation.

The eye in Figure 6-6 reveals an anomaly in the signal—a thin blue line whose blue color indicates a relatively infrequent transition. Yet it proves that at least one of the twelve signals has an edge that is outside the normal range. A mask feature helps locate the specific signal causing the problem. By drawing the mask in such a way that the offending edge penetrates the mask area, the relevant signal can be isolated, highlighted and brought to the front layer of the image. The result is shown in Figure 6-7, in which the flawed signal has been brought to the front and highlighted in white.

In this example the aberrant edge indicates a problem on the A3 (0) signal. The origin of the problem is crosstalk—the edge change is being induced by signals on an adjacent trace on the motherboard.

Conclusion

The functional verification and troubleshooting phase of the server motherboard project has detected and resolved several signal integrity-related problems. The logic analyzer is the first line of defense when testing digital functionality. However, digital problems can stem from analog signal issues, including edge degradation due to improper termination or crosstalk as demonstrated here. By teaming the logic analyzer with an oscilloscope and evaluating timecorrelated digital and analog signals on the same screen, problems affecting either domain are easy to solve.

With its full digital feature set verified, the server motherboard is ready to move on to its final compliance and interoperability tests.

iLink[™] Toolset: Two Powerful Measurement Tools Team Up

Although Logic Analyzers and oscilloscopes have long been the tools of choice for digital troubleshooting, not every designer has seen the dramatic benefits that come with integrating these two key instruments.

Logic Analyzers speed up debugging and verification by wading through the digital information stream to trigger on circuit faults and capture related events. Oscilloscopes peer behind digital timing diagrams and show the raw analog waveforms, quickly revealing signal integrity problems.

Several Tektronix Logic Analyzer models offer the iLink[™] toolset, a Logic Analyzer/oscilloscope integration package that is unique in the industry. The iLink[™] toolset joins the power of Tektronix TLA Series Logic Analyzers—memory depths to 256, MagniVu[™] acquisition with 125 ps resolution and advanced state machine-based triggering—to selected TDS Series oscilloscope models. A powerful set of iLink[™] toolset features brings timecorrelated digital and analog signals to the logic analyzer display. While the Logic Analyzer acquires and displays a signal in digital form, the attached TDS Series oscilloscope captures the same signal in its analog form and displays it on the Logic Analyzer screen. Seeing these two views simultaneously makes it easy to see, for example, how a timing problem in the digital domain was a result of a glitch in the analog realm.

The iLink[™] Toolset is a comprehensive package designed to speed problem detection and troubleshooting:

- iCapture[™] multiplexing provides simultaneous digital and analog acquisition through a single logic analyzer probe
- iView[™] display delivers time-correlated, integrated logic analyzer and oscilloscope measurements on the logic analyzer display
- iVerify[™] analysis offers multi-channel bus analysis and validation testing using oscilloscope-generated eye diagrams

▶ Primer



External Bus Compliance

Critical Measurements Support Final Pass/Fail Testing

We have gone to great lengths to prevent signal integrity problems on the motherboard where possible, and solve them where necessary. Now it is time to ensure that the motherboard's "deliverables"—valid signals for external subsystems—are compliant with standards. Not surprisingly, some of the recommended compliance measurements are in effect tests for signal integrity.

The motherboard's 10 Gigabit Ethernet (10 GbE) optical output is one of its most prominent high-speed serial features. The port is a XENPAK module that transfers data from a XAUI bus carrying four differential 3.125 Gb/s signals (one quad in each direction). Looking at the transmitter side, the XENPAK module converts the XAUI electrical signals received from the I/O Controller Hub (IOCH) to an optical 10GbE signal for transmission. The specific optical unit used in our design is a commercially-purchased module that mates to a standard receptacle on the motherboard. The module is designed to meet the IEEE 802.3ae standard for 10 GbE standard.



Figure 7-1. The loopback test uses its own transmitted data to analyze receiver performance. Note the attenuator on the Receive side. Some receivers cannot directly accept the levels sent by their own transmitters, therefore the attenuator steps down the level as needed.

Compliance tests range from relatively simple loopback exercises to thorough eye diagram analysis and specialized parametric measurements. The more demanding, oscilloscope-based eye measurement offer some challenges not seen earlier in the development process. Most importantly, data rates may exceed the bandwidth limits of real-time oscilloscopes and their automated compliance tools. The best solution for signal integrity and eye diagram analysis in this situation is a full-featured sampling oscilloscope such as the Tektronix CSA8200 or TDS8200.

Self-Test Methods Depend on Reliable Signals

Loopback tests are a form of self-test employed by many types of external buses, notably 10 GbE. The loopback method is simple: send a specific known data pattern out through the transmitter, physically connect the transmitter output back to the receiver input, and ensure that the receiver is capturing the same information that was sent. It may be necessary to include an optical attenuator in the loop to reduce the signal level coming into the receiver or to add "stress" to the test. The setup, including the attenuator, is depicted in Figure 7-1. The intelligence to generate and compare the pattern may be built into the bus module, the IOCH, or even the CPU itself. Many digital system manufacturers are choosing the latter approach, since it permits the most comprehensive patterns as well as the most "automated" loopback testing.

The loopback approach is widely accepted for receiver testing. It eliminates the need for costly bit error rate testers (BERTs) and associated tools, while providing results that correlate well with real-world operations.

Like most self-tests on this type of equipment, loopback tests start with the assumption that the transmitter is able to provide the correct signals. Unless the transmitter section has been thoroughly evaluated for functionality and compliance, it may have signal integrity problems of its own. Unfortunately these may be overlooked due to the short loopback connection. Yet they can cause failures in the system environment.

Consequently, loopback receiver testing may be meaningless unless it starts from the vantage point of a "known-good" transmitter. Oscilloscope-based measurements are a proven method to evaluate transmitter performance and signal integrity.

Transmitter Tests Lay the Foundation

The IEEE 802.3ae standard specifies a host of compliance measurements, including eye diagram tests and quantitative measurements on key power parameters. There are several steps in the compliance test but for the sake of clarity we will examine one crucial measurement known as Optical Modulation Amplitude to show how the procedure can be completed quickly and easily. Once the compliance step is done, the loopback tests can be performed confidently in the knowledge that the transmitter is sending valid signals.

Optical Modulation Amplitude is Typical of 10 GbE Transmitter Measurements

Most designers of high-bandwidth Ethernet equipment are accustomed to measuring two critical entities for compliance: Extinction Ratio and Average Optical Power (AOP). 10 GbE introduces a new measurement term, Optical Modulation Amplitude (OMA), that supplants AOP and supplements Extinction Ratio. Both AOP and OMA are expressed in dBm, that is decibels relative to 1mW. The absolute OMA value is simply the High (Binary 1) power level minus the Low (Binary 0) power level on a square wave signal. The formula for OMA expressed in dBm is as follows:

$$OMA_{dBm} = 10 \log\left(\frac{High-Low}{1mW}\right)$$

There are two methods for deriving the OMA parameter. The first is a "recommended" approach that is actually the binding requirement of the standard. The second is an "approximate" method that is useful for estimation but does not fully satisfy the standard. The approximate method is used when it is necessary to determine OMA in the presence of live traffic.

But live traffic is not a problem during the design process, and our motherboard must be proven compliant with the IEEE 802.3ae standard. Therefore we will pursue the recommended approach for the OMA measurement.

Primer



Figure 7-2. OMA measurement requirements, as shown in the IEEE 802.3ae specification.

The measurement solution for this application will be a CSA8200 sampling oscilloscope equipped with an 80C08C optical module. The module includes the built-in clock recovery capability required for the 10 GbE environment. The pairing of the sampling mainframe with the 80C08C module enables key 10 GbE compliance tests to be performed automatically (there are also comparable sampling modules designed for purely electrical standards, and for standards encompassing both optical and electrical functionalities. Connecting the optical signal to the DUT is very simple: a conventional patch cable connects the oscilloscope's optical input to the output of the XENPAK module.

Following the Standard

The IEEE 802.3ae standard puts forth some very specific guidelines for the OMA measurement. It requires a square wave signal of a frequency ranging from one-fourth to oneeleventh the bit rate of the 10 GbE signal. This is accomplished by concatenating from four to eleven consecutive ones followed by the same number of consecutive zeros. To ensure the most reliable clock



 Figure 7-3. Results of OMA measurement. This screen depicts a passing, compliant waveform.

recovery, the one-fourth frequency is recommended. As explained earlier, the capability to generate a signal with the appropriate duty cycle and frequency is often embedded in the output module itself, or within the firmware supporting either the I/O Control Hub (IOCH) or the CPU. If no such provisions are made, it will be necessary to manually program the IOCH to create the signal; or the "approximate method" (explained below) may be used. This alternative method works with the eye diagram of any data or random pattern.

Figure 7-2 is adapted from the IEEE 802.3ae specification. It illustrates the exact measurement points for the OMA parameter using the recommended measurement method. Just one cycle iteration is shown here, although the signal is continuous. Essentially, the power (amplitude) is captured during the middle 20% portion of the four consecutive high and the four consecutive low binary states. This should be the most stable portion of the power waveform.

With the CSA8200, the OMA test itself is a matter of selecting the OMA measurement from a menu and pressing one button. The oscilloscope takes care of the clock recovery, positions the measurement windows, acquires the instantaneous power values, and displays the result as a quantitative figure. To place the measurement windows, the oscilloscope first finds the midpoint of the positive-going and negative-going edges to determine the actual pulse and cycle width. Then, knowing that the signal is a square wave, the instrument centers the windows accordingly.

The outcome of the acquisition taken from our motherboard's XENPAK 10 GbE module is shown in Figure 7-3. Note that the screen layout is effectively the same as that depicted in the standard, though more cycles are shown. The display makes it easy to confirm that the measurement occurred during the middle 20% of the top and bottom of the pulse sequence for four consecutive 1s by four consecutive Os, as specified.

The OMA measurement is the first step toward proving that the 10 GbE transmitter is compliant with the standard and free of signal integrity problems. Subsequent mask tests will produce information about the signal's noise, jitter, aberrations, and rise/fall time. If these values are compliant, then the transmitter is for our purposes a "known-good" element. Now the loopback test can be performed with the assurance that any errors that may appear are arising from the receiver.



Figure 7-4. OMA measurement using the "approximate" method. This approach allows testing in the presence of live traffic but does not qualify as a full compliance test.

An "Approximate" Method for Estimating OMA

The IEEE 802.3ae specification offers an alternative means of measuring OMA. It is intended to provide a useable approximation when the square wave pattern is notavailable, and it allows the OMA measurement to be performed even when traffic is present. For engineers who feel the square-wave pattern is not representative of real-world operating conditions, this "approximate" approach can be used to confirm results derived using the recommended method.

The CSA8200 oscilloscope runs the eye diagram in Figure 7-4 automatically when invoked from the front panel. It determines the High and Low levels by deriving the mean value of the histograms at the top and bottom, respectively, of the eye diagram—directly above and below the crossing point rather than in the center 20%. From this information it calculates the appropriate measurement points and captures an OMA value. Using the eye diagram crossing point in this way is very unusual in the field of eye diagram analysis, but the results are very close to those achieved with the recommended (offline) measurement method.

▶ Primer



Characterization

Defining and Specifying the Final Product

The motherboard is complete, verified for compliance and functionality. Problems arising from signal integrity and other issues have been tracked down and resolved but more importantly, many potential problems have been forestalled by a thorough test plan that began with the raw PCB.

The final step is to characterize the product and determine its performance margins, limits, and tolerances. The information becomes the basis of data sheets and promotional materials.

In many cases these parameters must be tested over a range of temperature, humidity, altitude, vibration, and more. Power consumption is usually considered, and battery life is a factor in many types of products. The device's resistance to electromagnetic interference (EMI), as well as its own tendency to radiate EMI, must be evaluated and documented.

Much of the same equipment used in the earlier design process steps—real-time and sampling oscilloscopes, high-bandwidth probes, and specialized measurement and analysis software—is used to explore the limits of our finished product, the motherboard. By this point in the process, signal integrity problems should not be an issue unless external effects (loading, EMI, and so forth) cause deviations.

Characterization may include detailed stress testing. In this process, the device is subjected to variations in supply voltage, for example, to evaluate its ability to function normally under less-than-ideal circumstances. Another approach is to drive deliberately flawed signals into the input. Impairments may include marginal amplitudes, slow risetimes, and aberrations such as overshoot. An arbitrary waveform generator (AWG) such as the Tektronix AWG710B is the preferred tool to produce these symptoms.

Specifications for the Designer and for the End-User

Most high-performance electronic products have two specification levels:

- Published specifications guaranteed to the end-user.
 These appear in the product manuals, brochures, etc.
 A device such as the motherboard, whose functional connections to other equipment must comply with industry standards, will have a set of specifications affirming that compliance. There may be additional parameters such as clock frequencies that have concretely specified values and tolerances.
- Non-Published specifications intended to support the design process and provide guardbands above the published specifications. Guardbands are margins that ensure normal production units will fall within the published specifications.

Both sets of specifications are usually summarized in an Engineering Specification that acts as a guideline throughout the development project. A final round of measurements can confirm that the full range of specifications, both published and non-published, have been met. Depending on the type of device involved, these measurements might include parameters such as setup-and-hold and other timing tolerances; pulse amplitudes and risetimes; clock frequency stability; noise characteristics; jitter; certain impedance limits, and more. All of these measurements are well within the capability of the instruments already discussed in this primer. Primer

Summary and Conclusion

Over the course of the design process for the motherboard, we have relied on critical measurements to evaluate components, to detect and solve signal integrity-related problems, and to perform compliance tests. We have used an array of measurement tools and procedures:

Pre-Assembly Verification

- Impedance Measurements: Used the sampling oscilloscope with TDR module to check the impedances of PCB traces and connectors, eventually concluding that a proposed low-cost connector might cause signal integrity problems in the final design.
- ASIC Verification: Used the sampling oscilloscope and digital timing generator (signal source) to verify prototype ASICs and confirm that specifications such as Setup/Hold timing are within tolerances.

Functional Verification and Troubleshooting

- Basic Functional Verification: Used the real-time oscilloscope and high-speed SMA probing tools to check the board's clock signals for aberration-free performance; detected and analyzed a signal integrity issue (a glitch) in the spread-spectrum clock.
- Internal Bus Compliance: Used the real-time oscilloscope with automated eye diagram analysis software to perform internal compliance tests on Serial ATA and other on-board buses.
- Operational Validation, Fault Detection, Debug: Used a logic analyzer to monitor functional tests to verify that logic operations proceed correctly; detected a signal integrity problem using the real-time oscilloscope connected to the logic analyzer utilizing the iLink[™] toolset.

System Verification and Characterization

- External Bus Compliance: Used a sampling oscilloscope and optical acquisition module to verify signal integrity on 10 GbE transmitter signals prior to eye diagram compliance testing.
- Characterization: Described the use of real-time and sampling oscilloscopes to characterize performance limits, define margins, and generate information for publication.

In this primer, we have demonstrated the importance of being alert for signal integrity problems during the design process. Glitches, anomalies, and impairments can emerge at virtually any point along the way, and must be eliminated before proceeding to the next step. We have shown how signal integrity problems can be quickly solved with the aid of automated acquisition and analysis tools in conjunction with state-of-the-art Oscilloscopes, Logic Analyzers, and Signal Sources.

Contact Tektronix:

ASEAN / Australasia / Pakistan (65) 6356 3900 Austria +41 52 675 3777 Balkan, Israel, South Africa and other ISE Countries +41 52 675 3777 Belgium 07 81 60166 Brazil & South America 55 (11) 3741-8360 Canada 1 (800) 661-5625 Central East Europe, Ukraine and Baltics +41 52 675 3777 Central Europe & Greece +41 52 675 3777 Denmark 80 88 1401 Finland +41 52 675 3777 France & North Africa +33 (0) 1 69 81 81 Germany +49 (221) 94 77 400 Hong Kong (852) 2585-6688 India (91) 80-22275577 Italy +39 (02) 25086 1 Japan 81 (3) 6714-3010 Luxembourg +44 (0) 1344 392400 Mexico, Central America & Caribbean 52 (55) 56666-333 Middle East, Asia and North Africa +41 52 675 3777 The Netherlands 090 02 021797 Norway 800 16098 People's Republic of China 86 (10) 6235 1230 Poland +41 52 675 3777 Portugal 80 08 12370 Republic of Korea 82 (2) 528-5299 Russia, CIS & The Baltics 7 095 775 1064 South Africa +27 11 254 8360 Spain (+34) 901 988 054 Sweden 020 08 80371 Switzerland +41 52 675 3777 Taiwan 886 (2) 2722-9622 United Kingdom & Eire +44 (0) 1344 392400 USA 1 (800) 426-2200 USA (Export Sales) 1 (503) 627-1916 For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com

43

Copyright © 2004, Tektronix, Inc. All rights reserved, Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies. 11/04 DOR/WOW

55W-17671-0

