


New 100-Gb/s Transmission Standards Put Testing to the Test

Though 100 Gb/s is significantly faster than more mature high-speed technologies, enough similarities exist to help ease navigation through the new test landscape.

Varistors, also called metal-oxide varistors (MOVs), are used to protect sensitive circuits from a variety of overvoltage conditions. Essentially, these voltage-dependent, nonlinear devices have electrical characteristics similar to back-to-back Zener diodes.

Many data-communications and telecommunications technologies, such as Infiniband and Fibre Channel, now use 100 Gigabit Ethernet (100 GbE) for transport. High-speed-serial (HSS) technology makes it possible (Fig. 1). It encompasses noise-resistant differential signaling and jitter-resistant embedded clocking, along with closed-eye equalization. HSS enables 25+ Gb/s on printed circuit boards (PCBs). Furthermore, 100-



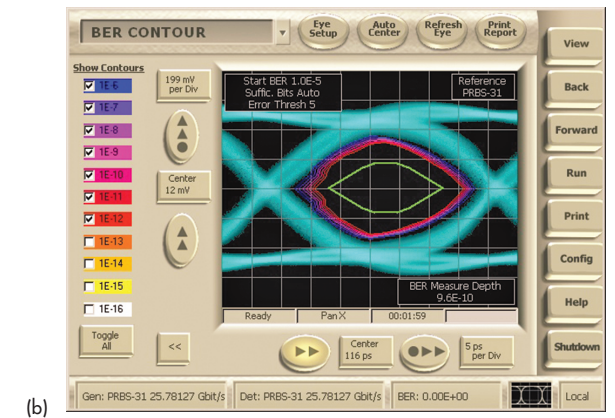
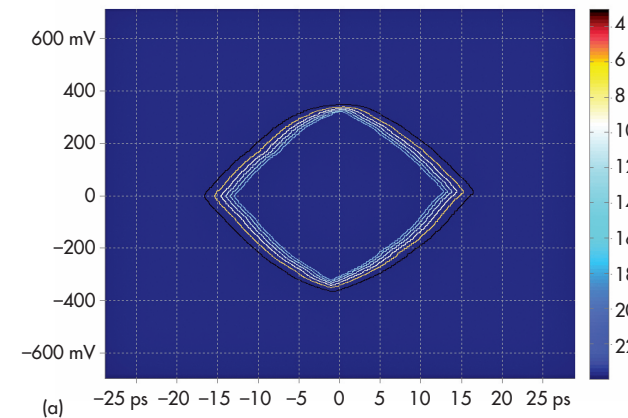
Gb signal transmission to optical transceivers can be achieved by paralleling HSS links, thus easing connectivity to the fiber-optic backbone.

Dealing with HSS requires an understanding of the way analog waveforms relate to digital-signal bit-error rates (BERs). Looking at signals on a scope, less than 3 ps of random jitter closes the eye altogether. The latest standards typically permit less than 700 fs of random jitter.

With that in mind, the question is: “How does one test a 100G system?” Fortunately, commonalities exist with 100 GbE, Fibre Channel’s 32GFC, and the Agreements of the Optical Inter-networking Forum’s Common Electrical Interface (OIF-CEI). Looking at those technologies, one can find familiar themes in the interplay of jitter, noise, and crosstalk (Table 1).

As standards evolve, the committees that generate them may use terms in different ways. For example, it’s important to distinguish between “data” rate and “payload” rate. Data rate is the propagation rate of raw data. Payload rate doesn’t

1. In a typical 100G system, a SerDes serializes a signal and transmits four over-25-Gb/s differential pairs. For each pair, the electrical signals from the SerDes go to an optical transceiver, which re-times the signals and transmits over either single-mode (SM) or multi-mode (MM) optic fibers. A second transceiver turns the optical signals back into electrical signals for a deserializing SerDes. A purely electrical version would follow the same scheme without the intermediate transceiver-driven optical signaling. At the top is a typical 4x25-Gb/s 100G SerDes-transceiver WDM optical system, and below that is a 4x25-Gb/s 100G SerDes-transceiver optical system. At the bottom is a 4x25-Gb/s 100G SerDes-to-SerDes electrical system. (The figure does not show the symmetric return paths.)



2. In eye-mask testing with BER contours, the BER=10⁻⁶ contour (the outer yellow-orange contour) corresponds to a 5x10⁻⁵ hit ratio (a). A BERT Scope using a 30-Gb/s signal displays a typical BER contour (b).

include overhead from error correction and coding, so payload rate is always less than or equal to data rate.

In addition, because signaling is always non-return to zero (NRZ), the proper unit for data rate is Gb/s rather than Gbaud. On the other hand, the terms “symbol” and “bit” may be used for the same fundamental communications element, depending on the standard being discussed.

ANALYZING THE STANDARDS

100 GbE – IEEE Std 802.3ba

To begin an examination of existing standards, consider the Long Reach (100GBase-LR4) and Extended Reach (100GBase-ER4) 100-GbE standards noted in the first column of Table 1. They’re part of 100-GbE optical transmission specifications, under IEEE Std 802.3ba.

One can see that the differences between them are primarily at the receive end. The ER4 receiver needs greater sensitivity and must pass a more difficult stress tolerance test than the LR4 receiver. Moving down the table, the Short Reach, 100GBase-SR4, 4x25 Gb/s low-cost, multimode (MM) standard, along with the 100GBase-CR4 and 100GBase-KR4 for electrical transport over cables and backplanes, are works in progress. When finished, the 100GBase specifications will provide a complete suite of optical interconnect systems.

Ethernet

Below the 100GBase specifications in the table’s first column are SONET/SDH standards. With Ethernet rapidly becoming the default for both datacom and telecom, it’s becoming obsolete, so skip discussion on these.

OIF CEI

This is where changes in testing start to get interesting. Implementation Agreements (IAs) from the OIF-CEI do not pre-

scribe compliance tests like that of the aforementioned specs. The emphasis is on “informative” and “normative” tests that are intended to assure interoperability. “Normative” tests are akin to compliance tests, while “informative” tests are intended to develop a more thorough understanding of performance and margin. For example, consider the two OIF-CEI IA tests summarized in Table 1.

Short-reach IA, OIF-28G-SR has the following features:

- Multiple lanes at 19.90-28.05 Gb/s
- Uses differential pairs
- Spans across 300 mm of PCB
- Up to one connection
- Operates at a BER of less than 10⁻¹⁵

As of this writing, IA, OIF-28G-VSR reach has yet to be officially published. However, taking guidance from the preliminary version of the standard, it will probably be multiple electrical lanes at 19.60 to 28.05 Gb/s for signaling between SerDes (called “hosts” in the IA) and transceivers (called “modules”).

The SerDes and transceiver may be separated by 100 mm of PCB to a connector, with an additional 50 mm or so of conducting trace; the system is required to operate at a BER below 10⁻¹⁵.

Fibre Channel 32GFC

The Fibre Channel standard, 32GFC, calls for a 28.05-Gb/s data rate. Why does “32GFC” refer to 28.05-Gb/s technology? The idea is that the name of each generation reflects a doubling of the payload rate in each generation. The payload rate for 32GFC is 25.6 Gb/s, twice that of 16GFC, even though the data rate didn’t double.

As the table indicates, the 32GFC document has not been published, and the preliminary version has few reference values.

TEST TYPES

Whether for transmitter or receiver testing, optical or electrical, test patterns that put every aspect of a component and every component of a system to the test are essential. One example is pseudo-random binary sequences (PRBSn), which are standardized patterns with every permutation of n bits.

The OIF CID jitter tolerance pattern is designed to have the most aggressive elements of the PRBS31, plus 72-bit sequences of consecutive identical (CID) bits. However, it's at a manageable length.

Transmitter Tests

All transmitter tests, both electrical and optical, should be performed with all system channels active in both directions in order to include all reasonable sources of crosstalk interference. To prevent unrealistic data-dependent interference, test patterns on the crosstalk channels should be different than the test signal pattern (or at least introduce sufficient delay between them so that the patterns aren't synchronized).

Crosstalk channels should also operate asynchronously. For one thing, except in special cases, each channel operates with a clock that's been independently recovered from its incoming data. Although each clock operates at the same nominal rate, they're neither frequency- nor phase-locked. In addition, synchronous-crosstalk degradation occurs in the same region of the test-signal eye diagram every time an aggressor makes a logic transition. However, synchronous crosstalk itself is random.

If all of the SerDes are integrated on one chip with multiple serialized outputs, there may be inter-chip crosstalk. Also in that case, if the outputs use a common clock, then they need to be active and synchronized with the test channel, and each channel should transmit a unique signal.

Do not overlook the frequency response of the printed-circuit board. Proper testing requires pre-emphasis at the transmitter and equalization at the receiver. Stressed receiver tolerance tests assure that receivers can operate at the specified BER even with the worst-case compliant input signal.

Eye-Mask Tests

Eye-mask tests (Fig. 2) on wide-bandwidth optical-to-electrical receivers and clock-recovery units are essential. Note that the clock recovery -3-dB bandwidth differs among the specs.

To perform the tests, the optical-to-electrical receiver should apply a 4th-order Bessel-Thompson filter with a reference frequency of three-fourths the data rate. It's needed so

TABLE 1: HSS STANDARDS					
Standard		Geometry	Reach	Data rate	Bif-error rate (BER)
100 GbE	100GBase-LR4	4 SM fibers	10 km	4 x 25.78125 Gb/s	≤10 ⁻¹²
	100GBase-ER4		40 km		
	100GBase-SR4*	4 MM fibers	≤10 m	iOS, Android	≤10 ⁻¹²
	100GBase-CR4*	4 cables, backplane	*	*	≤10 ⁻¹²
OIF-CEI	OIF-28G-SR	n traces on PCB	30 cm	19.90-	≤10 ^{-15*}
	OIF-28G-VSR*		15 cm*	28.05 Gb/s	
Fibre Channel	32GFC	n channels, optical and electrical	TBA*	28.05 Gb/s	≤10 ^{-12*}

*In recent discussions with standards committees

that different test platforms can operate under uniform measurement conditions.

Eye-mask testing is inherently random. This can be addressed by requiring a minimum "hit ratio," defined as the ratio of the number of mask violations to the total number of samples acquired per unit interval. The more hits, the greater the accuracy. A transmitter is compliant if it achieves a hit ratio of less than 5×10⁻⁵.

Optical Receiver Testing

This kind of testing resembles the optical receiver stress tests for the long and extended reach 4×25-Gb/s topologies. Dedicated test equipment obviously facilitates such testing. Table 2 summarizes some of the test patterns.

Essentially, what's happening during optical receiver testing is that the test equipment generates inter-symbol interference (ISI) with a 4th-order Bessel-Thompson filter that removes higher-order harmonics from the test-generator output. Then, random jitter is applied via a precision Gaussian noise generator. This can be imposed by adding the noise and then running the signal through a limiting amplifier.

TABLE 2: OPTICAL RECEIVER TEST PATTERNS	
0 x 0 off square wave	8 bits low, 8 bits high
PRBS 9	511 bits
PRBS 15	32,767 bits
PRBS 31	2.1 Gbits
Scrambled idle	
OIF CEI jitter tolerance pattern	(72 CID bits + ≥ 10328 from PRBS31 + seed) + complement

Electrical Receiver Tests

Stressed receiver-tolerance testing subjects the receiver to the worst-case signal while looking at BER. For 100 GbE and 32GFC, the receiver is compliant if BER ≤ 10⁻¹². For OIF-CEI, compliancy equates to BER ≤ 10⁻¹⁵. There are different levels and types of stress. Some only require sinusoidal jitter, while others need more.

ISOLATING FAILURES

At this point, an interesting question arises: "What if some element in the chain fails to meet compliance requirements—what's the process for isolating a failure source?"

Compliance tests tend to include too many elements to provide simple answers. To determine which elements or components of a system might be causing problems, product designers need to plan diagnostic tests that probe specific weaknesses. This means adding complexity, test upon test, to find problems and determine margins.

For example, if the transmitter fails, one should simplify the test conditions by removing any test-compliance boards. Subsequently, the transmitter output would be analyzed with as direct a connection as possible.

In addition, it's essential to perform jitter and noise analysis. Analyze the breakdown while applying more complex patterns, or introducing longer PCB copper runs. Apply pre-emphasis, and turn on crosstalk aggressors. For each set of conditions, analyze eye diagrams, BER-eye, BER-contours, and the jitter and noise breakdown.

This is where dedicated test equipment comes into play. For instance, jitter maps can distinguish different types of jitter. Non-periodic bounded uncorrelated jitter (BUJ) may imply that crosstalk isn't being sufficiently shielded. Duty-cycle distortion (DCD) suggests a look at the transmitter; ISI probably means trouble with the output path. And high levels of random jitter suggest problems with transmitter clocking, while sinusoidal or periodic jitter can be a symptom of electromagnetic interference from a nearby component, such as a switching power supply.

For more information on this topic, including test-equipment screen shots, download "Physical Layer Tests of 100 Gb/s Communications Systems" (<http://info.tek.com/www-physical-layer-tests-of-100gbs-comm-systems.html>).

INSTRUMENTATION

Essentially, eye diagrams are created by collecting and displaying voltage/time data samples, acquired at a sample rate that's orders of magnitude below the data rate. For example, the sampling rate for a 10-Gb/s (1010 bits/second) data rate might be 105 samples per second. These eye diagrams offer a way to quickly and intuitively assess the quality of a digital signal, providing information about rise times, fall times, jitter, overshoot,


and other aspects of the digital signal.

Test equipment for creating eye diagrams requires a bit-error-ratio tester (BERT) that generates pseudo-random repeating test patterns. BERT testers usually have a range of trigger choices.

For example, a "clock trigger" provides a classical eye diagram that implements all possible bit transitions in one display. Alternatively, a "divided clock trigger" may be handy when the instrument used to construct the eye has a narrower trigger input bandwidth than the signal data rate. And with a "pattern trigger," it's possible to scroll through the signal on the display device. In general, the test engineer should avoid triggering on the data itself, because long runs of identical characters don't supply enough transitions to trigger from.

Sometimes it might be necessary to trigger from a recovered clock. This would occur, for example, if a clock signal weren't available. Or, in the case of long-distance fiber-optic communications, corruption may occur to either the clocks at the transmit end, or the data at the receive end link. In either of those cases, the clock must be recovered.

Some recovery circuits that are usually characterized by narrow loop bandwidth can be accomplished by a filtering function. It removes some of the jitter from the clock signal and some of the jitter that was present on the data signal so that delay between jittered clock and jittered data adds destructively.

For a more extensive treatment of the above information, check out "Anatomy of an Eye Diagram" (www.tek.com/application-note/anatomy-eye-diagram). 

CHRIS LOBERG, senior technical marketing manager at Tektronix, is responsible for oscilloscopes in the Americas Region. During his more than 13 years with the company, he has held multiple positions, including Marketing Manager for Tektronix's Optical Business Unit. Loberg earned an MBA in marketing from San Jose State University, Calif..

IDEAS FOR DESIGN WANTED

Send us your Ideas For Design. We'll pay you \$150 for every Idea For Design that we publish. In addition, this year's top design as selected by our readers will earn an additional \$500, with two runners-up each receiving \$250. You can submit your Ideas For Design via:

• E-mail: roger.engelke@penton.com

OR BY

• Postal mail to:
Ideas For Design
Electronic Design
1166 Avenue of the Americas, 10th Floor
New York, NY 10036

Go to www.electronicdesign.com for our submission guidelines.