PennWell®

LIGHTWAVE®

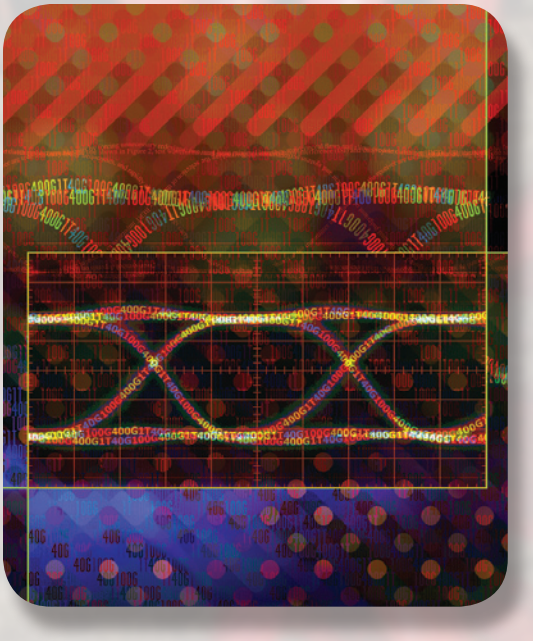

EDITORIAL GUIDE

Advances in optical design testing

Demands for support of 100 Gbps data rates in smaller packages have challenged technology developers' creativity. Optical communications engineers have responded with a variety of approaches, some of which are extremely complex. Fortunately, advances in test and measurement instruments have kept pace, as these articles illustrate.

SPONSORED BY:

Anritsu Discover What's Possible™ EXFO EigenLight PRECISE OPTICAL MONITORING JDSU

Originally published January 1, 2013

The challenges of high-speed technology test and measurement

Is your problem rooted in your high-speed technology design or in the equipment you're using to test it?

By **HIROSHI GOTO**

THE IMPORTANCE OF test-signal fidelity grows as data rates increase. That's because waveform inadequacies cause one of two business problems: One, they make component flaws seem worse than they really are and increase costs by forcing unnecessary redesigns; two, and more troublesome, they may conceal flaws and allow products with performance deficiencies to be shipped. Bit-error-rate testers (BERTs) such as the one shown in Figure 1 cast the final judgment of component and system performance — BER — making the selection of the proper BERT imperative to successful design and production of high-speed devices and systems.

Evaluating channels and receivers begins with the test waveform. For accurate compliance and useful diagnostic testing, pattern generators must have sufficient bandwidth to excite at least three harmonics. Figure 2 shows an instrument-grade eye diagram that has smooth, continuous edges with no harmonic distortion, bandwidth limiting, overshoot, or ringing, even at 25 Gbps. BERTs that produce such eye diagrams can support all three harmonics.

Test waveforms free of ISI and DCD

To achieve eye diagrams like the one shown in Figure 2, test waveforms must have minimal inter-symbol interference (ISI) and duty-cycle distortion (DCD). ISI, the dominant signal impairment at high data rates, is caused by the frequency response of the channel. It gets its name because the frequency content

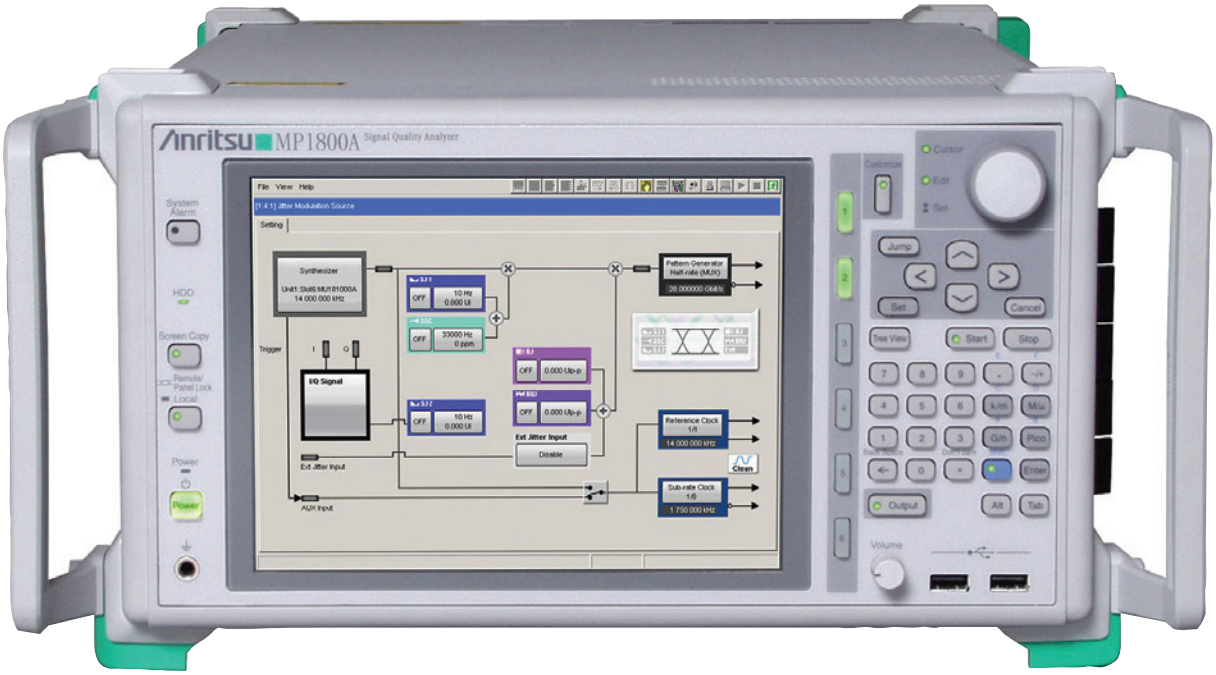

FIGURE 1. *BERTs are essential tools for characterizing the performance of communications components, subsystems, and systems.*

associated with a given symbol depends on the values of neighboring symbols. Long strings of identical symbols carry lower frequencies, and rapidly alternating strings carry higher frequencies. Pre-/de-emphasis at the transmitter and equalization at the receiver are techniques for solving ISI problems.

Simulations can predict the effects of a channel on a waveform at the earliest design stages. Of course, simulation accuracy depends on channel characterization accuracy. Simulations have trouble with real world edge effects caused by connectors, vias, dielectric anisotropies, etc., and have to be verified by measurement.

Extracting ISI from a simulation enables starting with an ideal waveform and examining it at any point in its propagation. The benefit of a measurement — provided the engineer starts with an instrument-quality waveform and takes care with any cables and connections — is that the engineer sees reality. Internally generated ISI is a big problem for engineers trying to make such accurate measurements.

Any intrinsic ISI in the test waveform corrupts measurements of channel ISI. It's possible to unfold the ISI of the initial waveform from the ISI of the outgoing waveform. But extensive, detailed analysis is required.

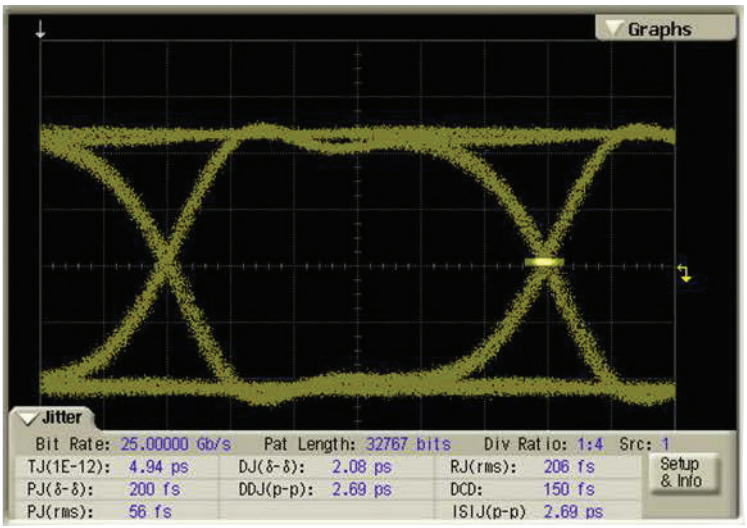

FIGURE 2. *Today, BERTs must provide eye diagrams with enough resolution to capture all three harmonics at data rates of 25 Gbps or more.*

DCD is the variation in the widths of positive and negative pulses, i.e., differences in the widths of 1s and 0s. It can be identified by vertical asymmetries in eye-diagram crossing points of alternating clock-like data patterns. DCD introduces extra frequency content that alters measurements of channel ISI and is difficult to untangle.

Quality pattern generators provide adjustable crossing points for a variety of reasons, including the ability to remove DCD. When selecting a pattern generator, examine the transmitted waveform and make certain that neither memory-based patterns nor pseudo-random binary sequence (PRBS) patterns exhibit ISI or DCD.

De-emphasis opens high-data-rate eyes

Since ISI is caused by a channel's frequency response, its effects can be mitigated by transmitting a signal that exhibits the inverse of that response. That's the idea behind pre- and de-emphasis.

For example, the first approximation of channel-frequency response is a low-pass filter. To mitigate this low-pass nature, high-frequency signal components are amplified or low-frequency components are attenuated. Pre-emphasis amplifies high frequencies and increases the total signal power, while de-emphasis attenuates the lower frequencies and signal power either decreases or remains constant. Since the effects on the signal are equivalent up to an overall constant, both can be referred to as "emphasis."

Taps and cursors

High-frequency components of a digital signal occur at logic transitions. The first order correction applies a larger voltage swing to the bit that follows a data

transition and is called a "two-tap" correction. In principle, this process can continue with finer corrections until the complete inverse frequency response of the channel is encoded in the transmitted signal. In practice, we're limited to applying correction factors to individual bits or else face the increased cost of requiring a transmitter with much greater bandwidth than the signal itself requires.

Each tap is a corrective boost or reduction applied to the amplitude of a bit neighboring the bit of interest. The bits to which these taps are applied are called "cursors." The terminology for bits that precede those of interest is a "pre-cursor" and for bits that follow, "post-cursor."

Since tap values reflect the inverse frequency response of the channel, their optimum values can be calculated from that frequency response. The most accurate way to characterize a channel's response is to measure its S-parameters with a vector network analyzer. To determine the optimal taps, transmission analysis software can be used with the BERT and a "four-tap" emphasis module.

To assure BER measurements accurately reflect the performance of the devices being tested, it's imperative that the pattern generator have minimal random jitter (RJ). A 1-psec RMS error in RJ becomes a 14-psec (peak-to-peak) mistake in total jitter (TJ) for a 10-12 error ratio. At 25 Gbps, 3-psec RMS of RJ is sufficient to cause a >10-12 BER. At extreme data rates, the intrinsic RJ of an instrument-quality waveform must be <350-fsec RMS to prevent the test equipment from contributing more than 10% of uncertainty to BER measurements.

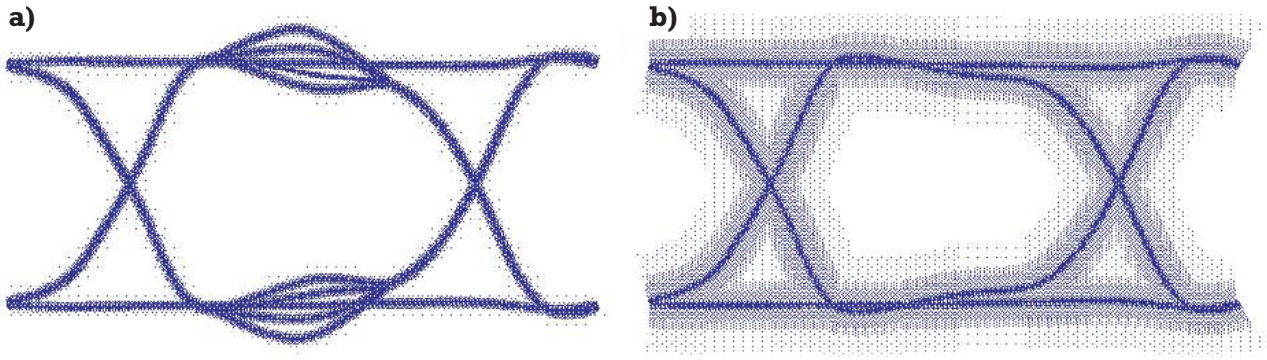

FIGURE 3. *There are two kinds of crosstalk: synchronized (a) and unsynchronized (b).*

De-embedding — the process of subtracting transmission-line effects from measurements — can help limit jitter. Effective de-embedding has two strict requirements. First, the S-parameters of any channel (including cables, PCB traces, and the connections between them) must be scrupulously characterized to bandwidths of at least 3X the data rate. Second, no matter how accurate the characterization of the network elements being de-embedded, the accuracy of the waveform extracted at the position of interest is still limited by the quality of the test waveform. Any ISI or DCD intrinsic to the test waveform must also be de-embedded, an arduous task that cannot be performed with S-parameters and simulation alone.

Effects of crosstalk

To create systems at 100 Gbps and higher, the advantages of high-speed serial technology — differential signaling, embedded clocking, transmitter emphasis, and receiver equalization — are combined with the data-rate scaling properties of parallel architectures. For example, 100 Gigabit Ethernet achieves 100 Gbps by combining either 10 separate differential parallel lanes, each at 10 Gbps, or four differential lanes, each at 25 Gbps. The problem of skew is solved by adding complexity to the data-link layer of the protocol stack. Crosstalk — the other major problem of parallel architectures — can't be corrected so easily.

There are two categories of crosstalk. Synchronized crosstalk occurs in systems where the separate lanes operate with the same clock. Since the victim and aggressors are frequency locked, crosstalk noise on the victim has fixed phase relationships (see Figure 3a). Unsynchronized crosstalk occurs when the victim and aggressors operate with distinct clocks. Here, the timing of aggressor noise varies over the victim eye diagram (see Figure 3b).

Designers can make headway against crosstalk through simulation, but just as is the case for ISI, the system ultimately has to be tested with parallel streams of data. Measuring the BER effects of crosstalk requires that the pattern generator have special features. It must be able to transmit several simultaneous signals. At 10—20 Gbps, three parallel signals are usually sufficient. At higher rates, where the unexpected must be anticipated, it may be necessary to excite every lane in the system.

The first step toward evaluating channel crosstalk is to assure that the multichannel pattern generator itself doesn't already have crosstalk. An easy way to check is to perform a BERT loopback test comparing bathtub plot measurements between single-channel and multichannel operation.

Crosstalk BER measurements can't be performed on an oscilloscope. While oscilloscopes can separate different types of jitter and estimate BER performance by extrapolation, they come up short with crosstalk. In general, TJ can only be measured on a BERT, and in the case of crosstalk, TJ can't even be reliably estimated on an oscilloscope. Simulation can provide insight into what level of crosstalk to expect; if the channels are synchronized, simulation also can help engineers identify bathtub plot asymmetries. But the magnitude of the crosstalk impact on BER can only be measured.

Beyond the cutting edge

Good test engineers eliminate systematic errors caused by worn-out cables, dirty connections, and poorly trained technicians. But systematic uncertainty caused by test equipment can be invisible right up until a system fails.

Engineers working at the cutting edge of technology must use test tools just beyond cutting edge. Accurate channel and receiver testing begins with the quality, calibration, and reliability of the pattern generator.

Tight budgets may make it tempting to use inexpensive equipment in a lab. Yet cutting corners by using flawed pattern generators with poorly integrated error counters will increase the net cost of both development and production. Inaccurate product characterization costs more in market share than the difference between a high-end BERT that produces high-quality waveforms and a less expensive alternative that generates inadequate waveforms.

HIROSHI GOTO is Anritsu Co.'s business development manager for general purpose test equipment. He has more than 25 years of experience in communications. He began his career in design engineering of optical products and has spent the past 10 years in product marketing for telecom and datacom products.

Originally published January 12, 2012

New test techniques required for CFP 100G optics

By **PAUL BROOKS**

AS WITH EVERY new technology, 100G is part evolution and part revolution, with the new CFP (C form-factor pluggable) optics a revolutionary step. Via new electronic, photonic, and software technology, the CFP offers the enabling step for cost-effective and successful 100G deployment. In particular, the 100G LR4 CFP module will economically address the challenge of aggregating traffic from routers to transport equipment at distances up to 10 km.

Most operators recognize the need to perform deep testing to ensure they make the right equipment choices and deliver maximum quality of service (QoS) to end users. During the testing process, it's important to identify signal integrity concerns such as noise, crosstalk, and impedance, and diagnose PCB and connector issues.

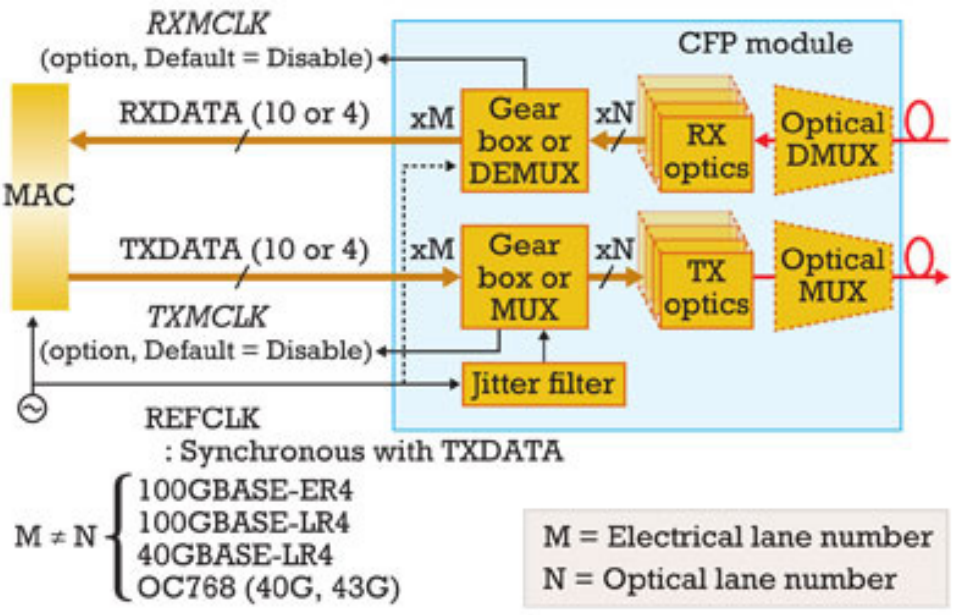

FIGURE 1. *Block diagram of a CFP gearbox IC.*

Major testing challenge

The IEEE 802.3ba standard, ratified in June 2010, details specifications for support of 40- and 100-Gbps Ethernet (40-GbE and 100-GbE) data rates, including 100GBase-LR4. Meanwhile, the CFP Multi-Source Agreement (MSA) defines a hot-pluggable optical transceiver form factor to enable 40- and 100-Gbps applications, including 40 GbE and 100 GbE. Client interfaces based on pluggable CFP optics are now appearing on routers, switches, and transport equipment.

The integration of so many complex, high-speed photonics and electronics inside a CFP module presents major challenges both for network equipment manufacturers and operators. The gearbox in particular is a very challenging technology to get right (see Figure 1), and there are also multiple issues with both the photonic and electrical interfaces that play critical roles in interoperability.

Operators, who have on occasion seen early technology adopters struggle with novel photonics, want to make sure that the new CFPs work reliably without errors and interoperate with other standards-compliant modules. The fact that the first-generation CFP modules are expensive, above $30,000 in many cases, only heightens this desire. For these reasons, the CFP testing process must go deeper than simply unframed bit error rate testing (BERT), which has traditionally been used to validate optical modules.

The testing source should mimic real-world signals to verify the transparency of the gearbox and electrical/optical and optical/electrical interfaces. The use of parallel data lanes makes it essential to measure inter-lane skews (the difference in propagation delay among the different lanes). With four optical wavelengths, all of the current optical parameters such as power, stability, stressed eye sensitivity (SRS), and eye pattern have to be measured simultaneously to validate the time delay among signals and skew on the different wavelengths. With data carried over 20 virtual lanes, 10 electrical lanes, and four optical lanes, identifying the domain in which errors have originated is critical.

Basic testing

The first test should normally be a simple optical loopback with attenuation. The CFP is inserted into the tester and the tester is set up to run a simple traffic test - either unframed, PCS, Ethernet, or Optical Transport Network (OTN) as

FIGURE 2. *CFP validation requires a series of tests. Here is a typical parameter setup for a stress test.*

required. The optical output from the CFP is looped via an optical attenuator back to the receive (Rx) port of the CFP. The tester first validates the CFP with an MDIO computability check. The next step is running a traffic test and looking for errors.

An unframed pseudorandom bit sequence (PRBS) or digital word that can fully stress the electrical and optical layers provides the basis for first-pass testing for CFP transponders. Normally a 10x10-Gbps PRBS/digital word can be used to validate continuity, but a 10-Gbps per lane pattern is not transparent to a 10:4 gearbox. The mux/demux process can scramble the bit order, leading to an invalid pattern. So it is necessary to use a gearbox-transparent pattern generator that offers 20 logical lanes, each with a 5-Gbps PRBS or digital word bit-sequence muxed into 10 physical lanes.

It is important that test systems report both erroneous zeros and erroneous ones as well as the conventional error count. This capability can help determine error bias, which can be an indication of which circuit function is failing. Longer patterns and digital words can also be used to stress clock recovery and DC balance aspects of the circuitry. The test system also should allow for offsetting

the PRBS patterns in time to remove any bit-wise correlation between lanes that could lead to crosstalk. The latest generation of test systems also provides an algorithm that can predetermine which 25-Gbps wavelength each 5-Gbps PCS virtual lane is mapped to, enabling the user to easily determine if errors are correlated with a 10-Gbps lane (electrical issue) or a 25-Gbps wavelength (photonic issue).

An optical eye test is performed by connecting a high-bandwidth scope with a suitable optical/electrical interface to the test system. An eye diagram is commonly used to review the transmitter waveform and evaluate transmitter performance. The data patterns are superimposed upon each other using a common time axis, typically a little under two bit periods in width.

The CFP transmitter should avoid gray areas in the center of the diagram that are midway between logic 0 and logic 1. Masks are often placed in the center of the eye in areas where the waveform should not exist to aid in verifying that the eye is open. To perform this test with accuracy, it is important that the test equipment driving the transponder has high bandwidth and a low jitter trigger.

Static and dynamic skew

At 10 Gbps per lane, only 100 psec exist per unit interval (UI), so even a small change in timing per lane may represent a significant part of the whole UI. The latest test systems enable control of two types of skew variation. The first controls variation on the whole UI on a per-lane basis, which is the standard static skew approach. This provides a useful indication at the PCS layer of the performance of the per-lane first-in/first-out (FIFO) buffer depth.

The second measurement, dynamic skew variation, is established as a mandatory test by IEEE 803.3ba. Dynamic skew variation ensures that the individual lanes continue to track correctly despite skew changes across a range of ±2 UI. Maintaining the dynamic skew tolerance ensures that any propagation delay variation in individual 10-Gbps output buffers will not cause bit errors in the line receivers.

The MDIO serial control also requires performance validation over a wide range of operating conditions. A debug application can ensure that the MDIO correctly

identifies parameters such as optical power, temperature, voltage, and vendor ID. The user can exercise peek/pole, block read, etc., and check the operation of control pins.

Stress testing parameters such as dynamic skew and clocking pulling range can be used to provide complete confidence in module performance (see Figure 2). The user can select options for skew, clock offset, PRBS, etc., or pick a fully automatic canned transponder validation test. Impairments in the optical lane can be augmented by electrical lane skew to fully test the margins of the electro-optics. Monitoring the performance of the receiver makes it possible to see the precise point at which the line card starts to produce errors. PCS layer alarms and errors can be fully validated against standards.

CFPs are the key to cost-effective and reliable 100G deployment, and so they require extensive testing and characterization. The testing challenges posed by 100G are formidable but the latest generation of test and measurement equipment is up to the challenge. New testers deliver complete coverage from the physical layer through the photonics layer while analyzing PCS/Ethernet traffic and providing deep OTN analysis with dynamic skew.

Breadth and depth of test coverage for CFPs coupled with applications that quickly highlight and diagnose issues is the key to quick and effective deployments.

PAUL BROOKS is product manager for Communications Test and Measurement at JDSU.

Originally published December 17, 2012

Real-time oscilloscope analysis for 28/32-Gbps SerDes measurements

By **BRIG ASAY**

WITH THE RECENT bandwidth breakthroughs in real-time oscilloscopes, measurements that previously could only be achieved using a sampling oscilloscope now can be handled by real-time oscilloscopes. Real-time oscilloscopes have not only broken the 60-GHz bandwidth barrier, they continue to improve in their measurement performance. While their noise and jitter levels still aren't as low as those of a sampling oscilloscope with a precision time base, they have narrowed what was once a large signal integrity gap.

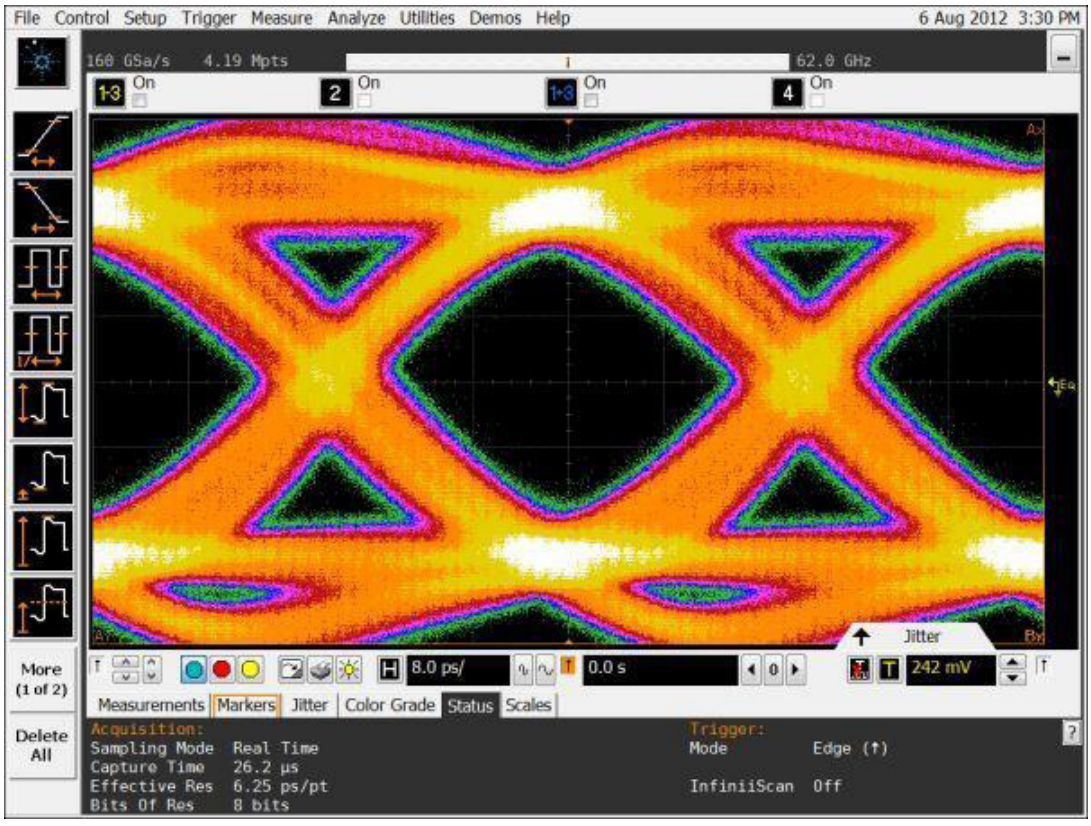

FIGURE 1. *A 28-Gbps eye captured by with one of the latest generation of high-bandwidth real-time oscilloscopes.*

As a result, oscilloscope users that previously would only use sampling oscilloscopes now should also consider a real-time oscilloscope. Real-time oscilloscopes offer the advantage of making measurements in a single acquisition and can capture milliseconds of data in that single acquisition, providing greater flexibility for debugging. The disadvantage tends to be cost (real-time oscilloscopes can be more than twice as expensive as sampling oscilloscopes) and the slightly worse signal integrity (i.e., higher noise and jitter).

With these tradeoffs in mind, many designers feel that the advantages of a real-time oscilloscope are very appealing, especially when designing with 28- and 32-Gbps signals (Figure 1). Choosing a real-time oscilloscope can be the right decision for these applications; however, it is important to know how to properly use these scopes to maximize design margins in 28- and 32-Gbps designs.

Bandwidth

Probably the most important characteristic of an oscilloscope is its bandwidth. To properly depict an edge, you must have the right amount of bandwidth. One rule of thumb for many years has been that you need to capture the fifth harmonic to have enough bandwidth on an oscilloscope. Of course, for 28- and 32-Gbps signals this would mean that you would need at least 70 GHz of bandwidth.

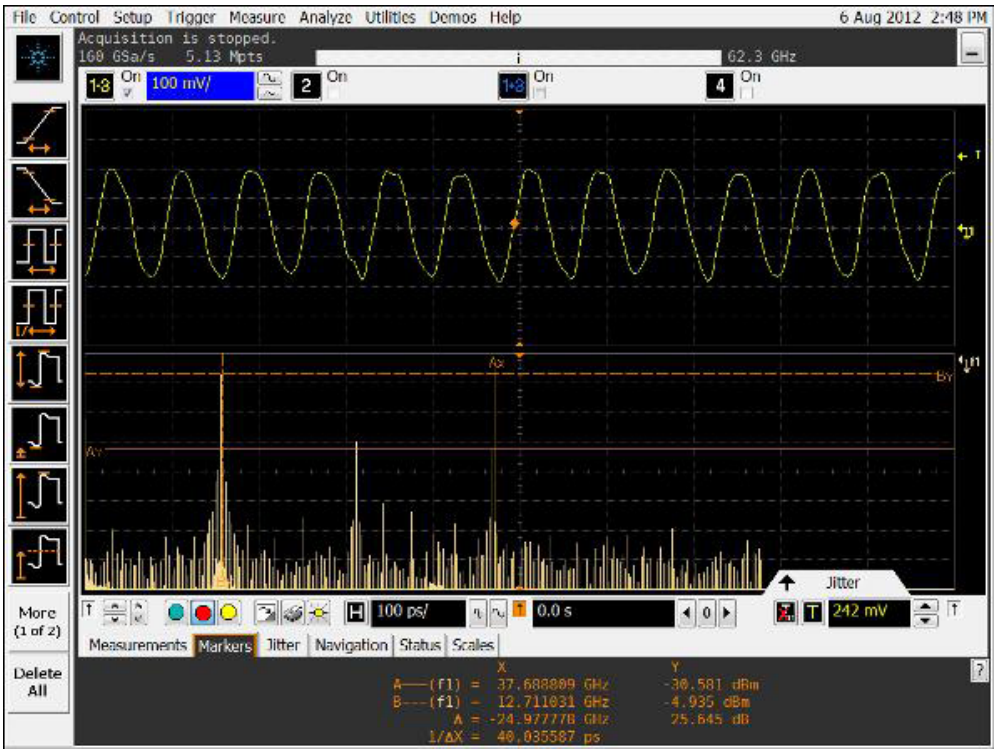

FIGURE 2. *FFT of a 28-Gbps SerDes. Notice that the third and fifth harmonics have significantly lower power than the first harmonic.*

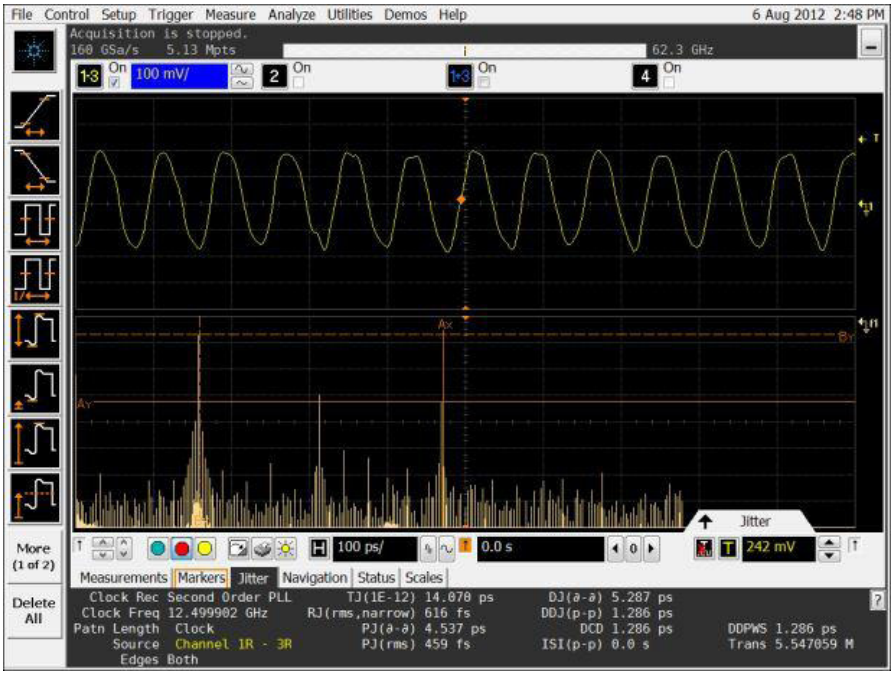

FIGURE 3. *In this 28- Gbps signal with 15-ps rise time, notice that the fifth harmonic still has energy.*

The reality is that the fifth harmonic rule of thumb does not apply for 28- and 32-Gbps signals. What is interesting about these high-speed SerDes signals is that they tend to have slower rise times. Slower rise times mean that there is less high-frequency content. Again, enough oscilloscope bandwidth is needed to properly digitize the rise time of the signal. If a signal has 40-ps (10/90) rise time then only 16 GHz of bandwidth is needed. In the case of a 40-ps rise time, the first harmonic is all the frequency content that exists. Figure 2 shows the high-frequency content of a typical 25-Gbps signal.

Of course, not all 28- and 32-Gbps signals have rise times as slow as 40 ps, which makes having more bandwidth important. For instance, a rise time of 20 ps could have harmonic content to the third harmonic. This would require a bandwidth of 50 GHz for the real-time oscilloscope. As rise times increase in speed to sub-10 ps, the fifth harmonic becomes a valid rule of thumb once again, as harmonic content would then be available to greater than 70 GHz. Figure 3 depicts a 15-ps rise time on a 28 Gbps device; notice that the frequency content of the signal reaches higher than 40 GHz.

It should be noted that one of the demands for high-speed devices is to minimize power; unfortunately, to drive a device to sub-10 ps requires very high power. For 28- and 32-Gbps signals, low power overrides faster rise times, which means 50

GHz of bandwidth is more than enough bandwidth for testing most high-speed SerDes devices.

Noise floor

As mentioned in the introduction, sampling oscilloscopes will have a lower noise floor than real-time oscilloscopes; depending on the full scale setting of the device, the difference can be significant. As a result, the noise floor is extremely important when using a real-time oscilloscope. In many ways, noise floor is nearly as important as the bandwidth of the oscilloscope. Oscilloscope noise erodes margins; the more noise that the oscilloscope has, the more the oscilloscope could influence the measurement.

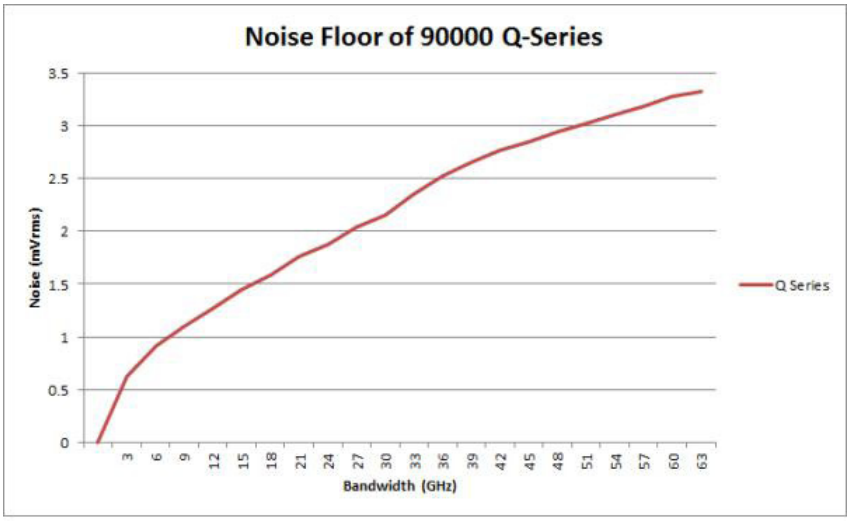

FIGURE 4. *As oscilloscope bandwidth increases, so does the noise floor.*

As data rates continue to shrink, the unit interval of the signal shrinks as well, which means that there are smaller margins. Oscilloscope users can ill afford to have these margins eroded due to instrument noise. In fact, if the oscilloscope noise exceeds the noise of the target system, it will become the principal source of the jitter and eye closing. Designers who spend much of their time ensuring their chips are fully functional do not want their measurements and decisions to be based on the oscilloscope noise rather than the design itself. Clearly, choosing a real-time oscilloscope with the right mix of bandwidth and low noise must be a key part of an oscilloscope purchase decision.

Once making 28G measurements, it is possible to improve the noise floor of the oscilloscope by using the correct amount of bandwidth. Due to high-frequency noise, as oscilloscope bandwidths increase, so does the noise (see Figure 4). As discussed earlier, typically 28- and 32-Gbps signals have slow rise times and very

little high-frequency energy above 30 GHz. Oscilloscopes also have a feature which is known as bandwidth limiting, where you can choose the bandwidth of the oscilloscope. By correctly choosing the bandwidth on your oscilloscope, excess noise can be eliminated, increasing design margins.

Jitter measurement floor

Typically noise will affect the height of the eye. But for 28- and 32-Gbps signals, many key specifications revolve around jitter. In most cases the jitter must be measured on very long patterns such as PRBS23 and PRBS31. Measuring jitter specifications such as total jitter, random jitter, and deterministic jitter with long patterns is difficult for high-performance oscilloscopes unless they are properly designed.

High-performance oscilloscopes have two key hardware specifications that contribute to 28/32-Gbps measurements. The first is known as the sample clock jitter and the second is known as the jitter measurement floor. Sample clock jitter is often known as the intrinsic jitter; it shows how good an oscilloscope clocking system is at aligning samples. Each high-performance oscilloscope has an internal clock that aligns the scope's data points. The internal clock ties the ADC to the sampler to the memory controller.

The sample clock jitter contributes to the jitter measurement floor. The internal clock of the oscilloscope is stressed with long patterns to properly separate random from deterministic jitter. The oscilloscope requires very deep memory and deep memory stresses the internal clock. The stressed internal clock can drift and thus misplace important samples. This misplacement of samples causes increased random jitter, not from the designer's device, but rather from the oscilloscope.

Oscilloscope designers must use extremely accurate clocks that maintain low sample clock jitter, even when stressed by deep memory.

Deep memory

As just implied, to measure jitter on long patterns accurately requires more memory. The deeper the oscilloscope memory and more accurate the oscilloscope sample clock, the more accurate the jitter decomposition will be for PRBS23 and PRBS31 patterns.

A PRBS23 pattern has 8.4 million bits in its pattern; at 28 Gbps and 160 GS/s, an oscilloscope must have at least 50 MPts of data to capture a single pattern. However, it is better to acquire more than one pattern in a single acquisition, which is where deeper memory means higher accuracy. Fortunately, real-time oscilloscopes capable of handling 2 GPts of data are now on the market.

Analysis tools

In recent years, real-time oscilloscopes have begun including more advanced analysis tools as part of their user interface or as optional additions. These tools include the ability to separate jitter, de-embed, equalize a signal, and look at a Fast Fourier Transform (FFT). The tools make measurements for 28-Gbps signals possible. For example, some of the measurements needed for the 28G Common Electrical Interface (CEI) include baud rate, rise time, fall time, output differential voltage, output common mode voltage, uncorrelated unbounded Gaussian jitter, total jitter, and duty cycle distortion. All of these measurements can be performed by a real-time oscilloscope with the right tools.

Just making the measurements may not be enough, however. As mentioned earlier, one of the dilemmas of a real-time oscilloscope is that it has higher noise than a sampling oscilloscope. One tool that is becoming increasingly popular among the real-time oscilloscope vendors is the ability to calibrate out oscilloscope noise and jitter. Another key tool is the ability to fix the random jitter in jitter measurements. Fixed random jitter means that regardless of the rise-time speed, which influences the amount of scope-induced jitter, the increased oscilloscope jitter can be ignored. Fixed random jitter also makes debugging cross-talk situations possible.

Conclusion

Real-time oscilloscopes can be effective tools for measuring 28- and 32-Gbps devices. It is important to understand the tradeoffs of a real-time oscilloscope that can make the most accurate measurements. When used properly a real-time oscilloscope will accurately depict the device and offer the flexibility needed in a single acquisition system.

BRIG ASAY is product manager, high performance oscilloscopes at Agilent Technologies.