

Eyeing jitter

SHAKING
OUT WHY
SIGNALS
SHAKE

JITTER HAS BECOME A HOT TOPIC among system designers. Seemingly easy to understand, it provides a quantifiable and—thanks to the now-ubiquitous eye-diagram display—graphical indication of the severity of a host of phenomena that damage data integrity.

Jitter's importance is undeniable, but whether it deserves all the attention it has been getting is another matter. For example, manufacturers of BERTs (bit-error-rate testers) think that some of that attention should instead focus on BER, which in one number tells you the probability of errors in a data stream. However, a common belief is that jitter can provide guidance on how to fix data-integrity problems but BER cannot. To this idea, BERT manufacturers reply that modern BERTs, which produce more than simple numerical BER values, are probably the best and fastest tools for unearthing problems that underlie data errors. But these manufacturers, as well as manufacturers of other types of signal-integrity instruments, are finding ways to get their instruments to display eye diagrams, which until recently only oscilloscopes could produce. Sometimes, the reason is to comply with standards that require such displays. Often, though, the reason is marketers' realization that eye diagrams build confidence in jitter meas-

urements, and many engineers won't buy signal-integrity instruments that don't produce eye displays.

In fact, many knowledgeable engineers feel that, to ensure data integrity in multigigabit-per-second data streams traveling over copper paths, you need more than one of the many hardware and software tools (**Table 1**) that fight data corruption. *EDN* adapted **Table 1** from material supplied by Agilent Technologies, which offers more types of tools than does any other manufacturer for measuring and analyzing jitter and related phenomena in multigigabit data streams. Even so, the **table** is not comprehensive; despite the breadth of its product lines, Agilent does not supply *every* type of software and measurement instrument used in signal-integrity work. For example, although the company—when it was Hewlett-Packard Test and Measurement—was among the first to offer TIAs (time-interval analyzers), it has never manufactured TIAs suited to analyzing multigigabit streams.

This article discusses time jitter—the

At a glance.....44

For more information52

JITTER MAY BE THE ENEMY OF DATA INTEGRITY, BUT ATTACKING JITTER HEAD-ON IS ONLY ONE—AND NOT ALWAYS THE FASTEST—WAY TO END DATA CORRUPTION.




Illustration by Mike O'Leary

variability in the point of occurrence of an event in a clock period or UI (unit interval). The event is almost always a logic-state transition. Ideally, the transition point would be the same in all UIs, but jitter always occurs, even if it is too small to notice or worry about. For visualizing how jitter can cause errors in data, probably no mechanism surpasses the eye diagram (Figure 1). Eye diagrams overlay waveforms from multiple UIs using either the real clock or a reconstructed clock as the timing reference. The diagrams show how transition-time variations can sometimes cause systems to sample data when its value is a logic 1

rather than the expected 0 or vice versa. When UIs are long, finding a sampling point at which data is reliably in the proper state usually is no problem, but, as data rates go up and UIs shrink, finding a suitable sampling point becomes increasingly difficult.

JITTER AND BER

Jitter can contain deterministic components. To the extent that such components are present, jitter is predictable (Figure 2). Measurements of finite duration can capture the worst-case value of D_j (deterministic jitter). Also, the PDF (probability-density function) of D_j is

bounded (Figure 3). Jitter typically also contains random components, however. You usually assume that the distribution of time intervals that exhibit R_j (random jitter) follows a Gaussian PDF (Figure 4); as you continue to measure instantaneous R_j in successive UIs, you acquire larger and larger peak-to-peak R_j values. Were you to wait long enough, you would measure infinite peak-to-peak jitter. Therefore, the magnitude of R_j is usually expressed as an rms value, which is also the 1σ (1-standard deviation) value of the Gaussian PDF. If the jitter is truly Gaussian and your system is immune to peak-to-peak jitter equal to some multi-

ple of the 1σ value R_j , **Table 2** can help you to determine the system's BER.

In many designs, the error-rate target is no more than one error per trillion transmitted bits (a BER of less than 10^{-12}). To produce an error rate that low, a system must introduce no errors when subjected to random jitter of $\pm 7\sigma$; that is, the system's jitter tolerance must be at least 14σ peak-to-peak. In theory, if the jitter has a Gaussian distribution, an average of $1.3 \cdot 10^{-12}$ bit transitions fall outside the $\pm 7\sigma$ permissible-jitter band, so achieving a BER of $1.0 \cdot 10^{-12}$ requires that the system introduce no errors when the jitter covers a slightly larger, 14.069σ , band, but many system designers accept 14σ as close enough.

Whereas you can most directly observe time jitter on its home turf—in the time domain—the inextricable relationship between the time and the frequency domains allows you to also observe time jitter in the frequency domain. Spectrum analyzers—the quintessential frequency-domain-measuring instruments—can detect time jitter. R_j widens the spectral peaks of periodic phenomena. P_j (periodic jitter), which is one type of D_j , produces peaks at frequencies that are related to its repetition rate. When you work with multigigabit data streams, you are working at multigigahertz frequencies. If you are fortunate enough to own a spectrum analyzer that can resolve spectral peaks spaced mere tens or hundreds of hertz apart at those high frequencies, the peaks' relationship to one or more frequencies can lead straight to the jitter's cause.

Companies such as Agilent and Anritsu offer very sophisticated frequency-domain instruments for measuring and characterizing time jitter. Although these instruments work well with data streams in copper networks, their forte is measurements on the even higher speed streams in optical networks. Sometimes, the line of demarcation between copper and optical networks becomes a bit blurred. For example, despite a name that suggests purely optical communication, Fibre Channel specifications also cover high-speed copper networks. Some signal-integrity experts call FCIA (Fibre Channel Industry Association) specifications the most authoritative documents you can find on methods of measuring data integrity in multigigabit copper networks (**Reference 1**).

Unless you've given the situation some

AT A GLANCE

▷ At today's multigigabit data rates, controlling jitter—the uncertainty in the time of occurrence of repetitive phenomena—has become a key factor in ensuring data integrity.

▷ Many kinds of instruments measure jitter and related phenomena, such as BER (bit-error rate). More and more of these units now display eye diagrams, which provide insights into the nature and causes of jitter. Eye displays are no longer the exclusive province of oscilloscopes.

▷ Contrary to popular belief, BERTs (bit-error-rate testers) do more than measure BER. In many situations, they provide the fastest insights into how and why systems behave as they do.

▷ Even passive components can cause data-integrity errors. With the aid of modeling software, you can predict these errors with high accuracy—and thus eliminate them.

thought, you may not realize that passive circuit elements can introduce or exacerbate time-jitter in high-speed signals. Passive components can do this damage by causing reflections, by picking up crosstalk and noise, and by degrading signal amplitudes and rise times—that is, by reducing bandwidth. The longer a signal takes to cross an active device's threshold region, the more uncertain is the time at




Figure 1 The familiar eye diagram, obtained by overlaying multiple UIs (unit intervals), shows much about your signals. You can see the rise and fall times and how jitter displaces the state transitions. The clearances between the waveforms and the specified forbidden region in the middle of the eye give you some idea of how closely system performance comes to not meeting the spec. But when the spec defines acceptable performance as no more than one error in a trillion UIs, you need a more rigorous tool.

which the IC detects a state transition. This uncertainty constitutes jitter. Similarly, attenuation, which often degrades SNR, can also introduce jitter, particularly when the signals are none too big to begin with—as is the case in LVDS (low-voltage differential signaling) and other ultra-high-speed systems.

Using measurements made with high-performance TDR/TDT (time-domain-reflectometer/time-domain-transmission) systems, specialized EDA software packages, such as TDA Systems' IConnect, can model imperfections in transmission lines and connectors and predict their effect on jitter. The models yield simulated eye diagrams that are indistinguishable from the real eye diagrams that you obtain by using a scope to observe the components' responses to pulse patterns. (TDRs send fast-rise-time pulses into cables and components and observe the reflected signals at the port at which they apply the pulses. TDT systems apply pulses at one port and observe the transmitted signals at another.)

Of the many types of jitter-measurement instruments, specialized TIAs made by two companies, GuideTech and Wavecrest, have become the standard in production test. Wavecrest, in particular, holds key patents in areas such as separating R_j and D_j . The company has also published an excellent introduction to jitter (**Reference 2**). Wavecrest says that, despite its reputation for supplying instruments that make jitter measurements in production, more than half of the units it produces go to design and development laboratories.

Although neither the GuideTech nor the Wavecrest instruments require a trigger, and both are TIAs that identify and determine the timing of state transitions in the signals under test, the instruments' architectures differ fundamentally. Wavecrest TIAs directly measure UI durations and waveform duty cycles. GuideTech TIAs record with great precision the times at which state transitions occur and compute the desired quantities from the stored time stamps. When working with multigigabit data streams, neither type of instrument attempts to measure the time of every transition or the duration or duty cycle of every UI. Both types of instruments must, however, recognize and count the state transitions whose time of occurrence they don't measure. GuideTech claims that its



instruments can make more measurements per unit of time. By making series of measurements at varying threshold levels, both companies' instruments construct eye diagrams that look just like those you can see on oscilloscopes.

MEASURING BER INSTEAD OF JITTER

The idea behind a BERT is simple enough. In its most basic form, the instrument consists of a pulse-pattern generator that produces a serial data stream and applies it to the UUT (unit under test). The second key functional block is a decision comparator, which compares the UUT output with the value expected at a specified time within the UI. The final major elements are counters, which accumulate the total number of UIs in the data stream and the number of UIs in which the measured and expected data differ at the decision point.

Besides varying the decision point, high-performance BERTs can also vary the comparator's 1/0 threshold. You might think that the lowest (best) BER always occurs when the instrument samples the received data exactly at the middle of the UI and when the 1/0 threshold is exactly at the midpoint of the signal swing. However, you can often obtain useful insights into the causes of D_j when the optimum sample timing or threshold voltage differs from these ideals. Moreover, by varying the sampling point and

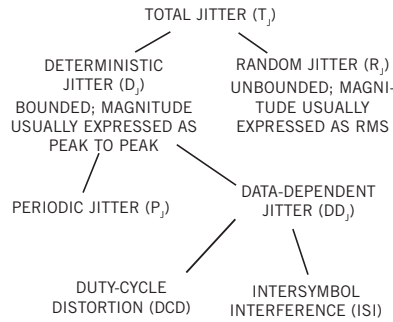


Figure 2 Jitter is either random or deterministic. Deterministic jitter is either periodic or data-dependent. You can characterize data-dependent jitter as either duty-cycle distortion or intersymbol interference (courtesy Wavercrest).

the threshold, a BERT can construct both eye diagrams and jitter "bathtub" curves (Figure 5).

The bathtub curve is, in effect, an eye diagram viewed in cross section. Eye diagrams often use color to indicate the frequency of occurrence of signals at the various pixel positions. However, a modern video adapter in a PC can construct a monochrome 3-D eye diagram in which the frequency of occurrence is represented not by color but by the simulated height of the pixels above the plane of the screen. If you "slice through" such a 3-D eye by holding the "knife" at a designated voltage level, you get a bathtub curve, although, to maximize the dynamic range of the presented data, the

bathtub curve's vertical, frequency-of-occurrence scale is usually logarithmic—not linear as it usually appears in the simulated 3-D representation.

Another way of obtaining a jitter bathtub curve is to plot BER (displayed on a logarithmic scale) versus the decision point within the UI. Such a bathtub curve is subtly different from and often more useful than the one described in the previous paragraph. Whereas the eye diagram shows a PDF, BER is a CDF (cumulative-distribution function).

SEQUENTIAL VERSUS RT SCOPES

Two types of digital oscilloscopes are useful for jitter measurements on multi-gigabit data streams. The more familiar type is the RT (real-time)-sampling scope, in which the current maximum bandwidth is 6 GHz. A more specialized type of instrument with much wider bandwidth—as much as 80 GHz—is the sequential-equivalent-time sampling scope. Agilent calls its sequential-sampling scopes DCAs (digital-communications analyzers). Tektronix refers to its sequential-sampling scopes as CSAs (communications-signal analyzers). Not all Tek scopes with CSA model numbers are sequential-sampling instruments, however. Some sample in real time and also allow users to operate them in the random equivalent-time-sampling mode. The third manufacturer of ultra-wideband digital scopes, LeCroy Corp,

TABLE 1—TOOLS FOR MEASURING JITTER AND SIGNAL INTEGRITY IN MULTIGIGABIT DATA STREAMS*

Where tested	Product types tested					
	Components		Subsystems		Systems	
	R and D	Manufacturing	R and D	Manufacturing	R and D	Manufacturing
Data rate						
Less than 3.2 Gbps	6-GHz-bandwidth real-time-sampling scope	20-GHz-or-greater-bandwidth sequential-sampling scope	6-GHz-bandwidth real-time-sampling scope	20-GHz-or-greater-bandwidth sequential-sampling scope	6-GHz-bandwidth real-time-sampling scope	20-GHz-or-greater-bandwidth sequential-sampling scope
	BERT capable of making input-output comparisons at user-selected voltages and times within the UI		BERT capable of making input-output comparisons at user-selected voltages and times within the UI			
	Low-jitter pulse generator		Low-jitter pulse generator			
	High-performance frequency-domain jitter-characterization system		High-performance frequency-domain jitter-characterization system			
Equal to or greater than 3.2 Gbps	20-GHz-or-greater-bandwidth sequential-sampling scope		20-GHz-or-greater-bandwidth sequential-sampling scope		20-GHz-or-greater-bandwidth sequential-sampling scope	
	BERT		BERT		BERT	
	Parallel BERT		Parallel BERT		Parallel BERT	
Media	Frequency-domain physical-layer test system	High-performance TDR scope	Frequency-domain physical-layer test system	High-performance TDR scope	Frequency-domain physical-layer test system	High-performance TDR scope

* Courtesy Agilent Technologies



does not make sequential-sampling instruments. LeCroy's highest performance instruments, the SDA series, are RT-sampling units, which also support random equivalent-time sampling. LeCroy calls this mode RIS (random-interleaved sampling).

Both types of equivalent-time-sampling scopes require repetitive waveforms that are preceded by a trigger event and a stable delay. The waveforms need not be periodic, however. Despite their enormous bandwidth, sequential-sampling scopes acquire waveforms more slowly than do RT and RIS instruments. By taking one sample after each trigger and then slightly advancing the sample point in preparation for the next trigger, sequential-sampling scopes translate high-frequency waveforms into much lower frequency signals. This mode of operation allows effective sampling rates of hundreds of billions—or even trillions—of samples per second, even though the actual sampling rate rarely exceeds 200k samples/sec.

Part of the reason that sequential-sampling scopes can achieve such high bandwidth is that the sampling circuit is a so-called zero-order-hold circuit or boxcar sampler. Once sampled, the analog waveforms that the instrument must handle are low-frequency replicas of the ultra-high-frequency input signals. Another advantage of the sequential-sampling architecture is that the relatively leisurely analog-to-digital-conversion rate permits the use of very-high-resolution ADCs; converters with 14-bit resolution are common in these scopes. Ultra-wideband RT scopes use 8-bit ADCs, though some such scopes can obtain resolution greater than 8 bits through averaging of readings on repetitive waveforms.

Among the requirements for faithful jitter measurements with any digital scope is low aperture uncertainty. If you are making eye-diagram measurements, the external delay from issuance of the trigger signal to the start of the measured waveform must be extremely stable. In a sequential-sampling scope, the internal delay from receipt of the trigger to capture of a waveform sample must also be very stable. Although jitter on these time intervals adds to the jitter you are trying to measure, it need not be troublesome, because the state of the art for the stability of each of these intervals is often measured in femtoseconds. Suppose that,

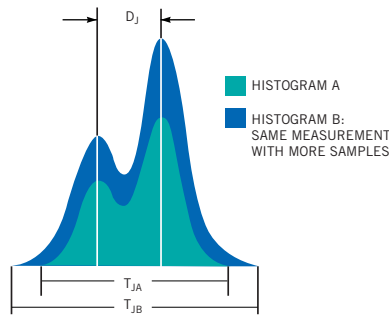


Figure 3 The observed peak-to-peak value of random jitter grows without limit the more measurements you make. Therefore, the upper curve is both wider and higher than the lower one. These histograms are for a combination of random and deterministic jitter. The random component is what makes the tails (sides) of the distribution not vertical. The double peak joined by a curve indicates that the jitter includes a sinusoidal component (courtesy Wavecrest).

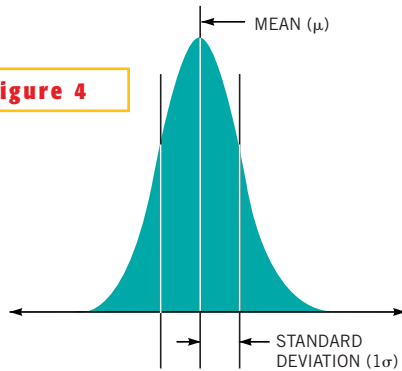


Figure 4 The most common assumption about the shape of the probability density function of random jitter is that it conforms to a Gaussian distribution, such as this one. What you can't see because of the scale is that the tails extend outward to plus and minus infinity.

in addition to the jitter you want to measure, your measurements must account for three 200-fsec-rms R_j sources, which you assume exhibit Gaussian PDFs. Because the distributions exhibit similar shapes and the values are expressed as rms, you combine them by taking the square root of the sum of the squares. Thus, these sources add a combined jitter of $\sqrt{120,000}$ fsec=346.4 fsec rms, which, for Gaussian distributions at the $\pm 7\sigma$ points, is equivalent to approximately 4.874 psec p-p. If you are working with a 3.125-Gbps data stream, the UI duration is 320 psec. Therefore the 4.874-psec p-p of additional jitter is only a little more than 1.5% of a UI. Even when you consider that a UI can contain two

state transitions, the associated peak-to-peak jitter still totals only about 3% of a UI, which is rarely significant.

Although scope manufacturers say that the 6-GHz, -3-dB bandwidth of today's top-of-the-line RT-sampling scopes is adequate for measurements on 3.125-GHz data streams, many engineers remain unconvinced. They prefer the additional bandwidth of sequential-sampling instruments, whose -3-dB frequencies can be more than an order of magnitude greater. This point of view may have some justification. Aliasing—the appearance in a signal of spurious frequency components that can mimic P_j —occurs if the scope's ADC receives a signal that contains significant energy beyond f_N (the Nyquist frequency—half the sampling rate). In RT instruments that take 20G samples/sec—the highest sampling rate in all of today's fastest-sampling RT scopes— $f_N = 10$ GHz. These scopes' 6-GHz, -3-dB frequencies are only 60% of f_N , and the scopes use brick-wall filters to attenuate frequency components near and above f_N , so aliasing should not be a problem. If, however, components with peak-to-peak amplitudes greater than 1 LSB (least significant bit) reach the ADC input, their presence could justify concerns about aliasing and its influence on measured jitter.

YOUR LIFE ISN'T LONG ENOUGH...

How do sequential-sampling scopes' notoriously low throughputs—numbers of waveforms per second the instruments can acquire and display—affect their usability for jitter measurements? The answer is: compared with RT scopes, not as much as you might think. The sequential-sampling units are *almost* always invariably slower than the RT scopes. Depending on their architecture, RT scopes can spend substantial time processing data—for example, performing sin x/x interpolation to “connect the dots” between samples. Sequential-sampling scopes do not usually have to perform such processing because their samples are much more closely spaced on the waveform.

The fastest RT scopes take 20G samples/sec—one sample every 50 psec. The UI of a 3.125-Gbps bit stream is 320 psec wide. In other words, the RT scope samples the bit stream only six or seven times in each UI. From those six or seven samples plus a few more just outside the UI,

the scope can reconstruct the analog waveform within the UI. If the input data is bandlimited to 10 GHz or less, such reconstruction is not a problem.

But the scope does have another problem: At 3.125 Gbps and 20G samples/sec, an RT scope with the deepest waveform memory currently available—96M samples per active channel—can capture just 15 million UIs. The data set might thus be long enough to determine whether the UUT achieves a BER of 10^{-7} , but when the scope has filled its memory once, it has insufficient data to say whether the BER is much lower. Moreover, in each of the 15 million UIs, the scope must test for a potential error by determining whether the waveform entered the eye mask's forbidden (open) zone.

This testing is bound to be time-consuming, and determining whether the BER meets a spec of 10^{-12} requires a minimum of 10^{12} tests. In other words, the scope must repeat the set of 15 million tests at least 66,667 times. Some people say that you should test long enough to accumulate 100 errors—assuming that the BER just equals the specification. Such a conservative approach would require that you test 10^{14} UIs. I couldn't determine how much time the scope needs to conduct a test, but if it can test one UI per microsecond, the testing would take 10^{12} μ sec (to capture one error) to 10^{14} μ sec (to capture 100 errors). That period, 10^6 to 10^8 sec, is equivalent to 11.57 to 1157 days, or as much as 3.17 years of round-the-clock testing.

A BERT is much faster; allowing time for a normal amount of processing, the BERT needs less than 10 minutes to directly determine whether an error has occurred in each of 10^{12} 320-psec-duration UIs. If you want to allow the BERT enough time to accumulate 100 errors at the specified 10^{-12} BER, you can run the test overnight; it takes less than 1000 minutes—less than 533 minutes if processing overhead is negligible. Because of the BERT's speed, even scope manufacturers that don't make BERTs recommend that you have both a BERT and a scope. Moreover, the BERT contains a serial-data generator, which you need for many tests in which you also use the scope. Now, by displaying eye diagrams,

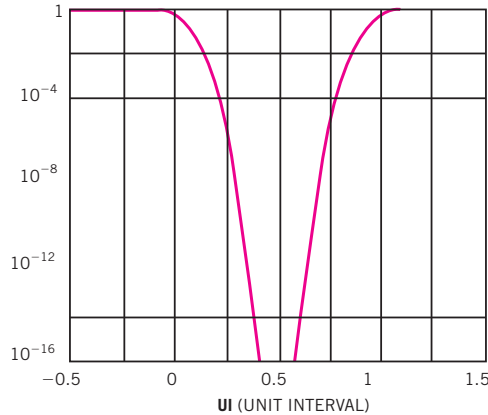


Figure 5 This bathtub curve for an ideal situation shows that, by carefully positioning the point within the UI (unit interval) at which your system samples the data stream, you can achieve an error rate of 10^{-16} , or one error in 10 quadrillion bits.

some BERTs, such as SyntheSys Research's BitAnalyzer 1500, allow you to eliminate the scope from certain test setups. Although SyntheSys doesn't suggest that this capability will enable many of its customers to get rid of their scopes or avoid buying new ones, it can save a lot of time. If you need both a scope and a BERT and the enhanced BERT lets you do without the scope in some tests, you have to move cables from the scope to the BERT and back again much less often.

TYING UP SOME LOOSE ENDS

The industry groups that set standards for various high-speed buses have developed many standards for jitter measurement. Among those is a data pattern called K28.5, which is widely used in Fibre Channel testing. Some workers in the field find this pattern especially useful because, even though it is only 20 bits long (00111110101100000101), it contains strings of five successive ones and five successive zeros, which can shift the data stream's average voltage level. When you repeat the pattern, you also get 0101001

and 1010110 strings. The pattern thus reveals DD_J (data-dependent jitter).

An instrument that is new to jitter measurement is the protocol analyzer. Protocol analyzers have heretofore dealt with signal-integrity issues only by inference. They track down protocol violations, which are often a consequence of signal-integrity problems. Generally, though, you are responsible for making the connection to the cause of the violation. Information from the protocol analyzer helps you to identify possible causes, but, to home in on the physical-layer problem, you may need an instrument, such as a scope, that makes analog-waveform measurements. Now, however, Data

Transit has introduced an eye-diagram-display pod for its BusDoctor Rx protocol analyzer. The company says that the EyePod is no replacement for a BERT, but it may save you from having to connect a scope to the UUT—and you can view the eye diagrams on the screen of the PC that hosts the protocol analyzer.

A term that you may have heard in connection with jitter measurement is "golden PLL." A golden PLL is usually a circuit that extracts a trigger signal from a data stream. A common use of this trigger is for making eye-diagram measurements. Sometimes, the golden PLL is implemented not in hardware but in software algorithms that run in the computer within the scope that displays the eye diagrams. Different scope manufacturers prefer the different approaches and make strong arguments in favor of both hardware and software golden PLLs.

Other frequently used jitter-measurement terms are TIE (time-interval error) and jitter track. If you have a stable trigger source—whether or not derived from a golden PLL—you can use it as the reference for eye diagrams. You can also use it to accumulate a jitter track. The jitter track is a record of a series of measurements of successive time intervals, all of which would have the same duration if there were no jitter. Often, the beginning point of each of the measurements in the series is the aforementioned trigger. If you calculate the mean (average) value of all of the measurements in the jitter track and subtract the mean from each measurement, you have a record of TIE versus time. You can now plot a histogram

TABLE 2—BIT-ERROR RATE VERSUS PEAK-TO-PEAK JITTER*

T _J (peak to peak)	BER
6 • σ	1.3 • 10^{-3}
8 • σ	3.17 • 10^{-5}
10 • σ	2.87 • 10^{-7}
12 • σ	9.87 • 10^{-9}
14 • σ	1.28 • 10^{-12}
14.069 • σ	1 • 10^{-12}

Note: T_J (peak-to-peak) is the maximum jitter the system accepts without error.

* Courtesy Wavecrest



of TIE and determine such values as rms and peak-to-peak jitter.

Although this treatment of jitter measurement is hardly definitive, one additional point deserves comment: the role of proprietary technology. Because jitter is so important to data integrity, many companies are working on jitter measurement. Several of them have positioned key technical people on standards-setting bodies. In most cases, these engineers' mission is to ensure that the standards are written based on techniques on which their employers hold or have applied for patents. When a patented technology becomes part of a standard, the company that holds the patent doesn't always license the technology to competitors or does so only for such high fees that the competitors cannot profitably sell products that embody the technology. This situation forces the companies that don't hold the patents to find alternative ways of making equivalent measurements. Sometimes, the alternatives are technically superior, but the company proposing them bears the burden of proving that the results are equivalent to those of the patented technique. This situation is wasteful, and it often results in widespread confusion in the marketplace. One consequence is a plethora of conflicting standards that inhibit adoption of the underlying technology. The companies responsible for this situation probably pride themselves on their cleverness. However, they may

be—as is said—too clever by half. □

REFERENCES

1. Nguyen, John, Pierre Raymond, and Hu Yoshida, *Multiprotocol Fibre Channel: the foundation for server and storage data networks*, Hitachi Data Systems, Santa Clara, CA, www.hds.com/pdf/wp94_multiprotocol_fc.pdf
2. Wavecrest Corp, Eden Prairie, MN, *Understanding Jitter: Getting Started*, www.wavecrest.com/technical/VISI_6_Getting_Started_Guides/6understanding.pdf
3. Rowe, Martin, "Measure jitter three ways," *Test and Measurement World*, March 2003, pg 17.

AUTHOR'S BIOGRAPHY

This article is yet another byproduct of EDN Senior Technical Editor Dan Strassberg's involvement with "The Scopes Trial" (EDN, Feb 6, 2003, pg 44). Strassberg became fascinated with the jitter-measurement and -characterization software embedded in the scopes under test and decided that readers might want to learn more about measuring jitter. Strassberg holds a BSEE from Rensselaer Polytechnic Institute and an MSEE from MIT. He has been writing about test and measurement for EDN since 1987. You can contact him at dstrassberg@edn.com.



FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to www.edn.com/info and enter the reader service number. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

Agilent Technologies

1-800-452-4844
www.agilent.com
Enter No. 301

Amherst Systems Associates

1-413-596-5354
www.amherst-systems.com
Enter No. 302

Anritsu Co

1-800-267-4878
1-972-644-1777
www.anritsu.com
Enter No. 303

Data Transit Corp

1-408-279-1555
www.datatransit.com
Enter No. 304

Digital Lightwave Corp

1-877-442-3445
1-727-442-6677
www.lightwave.com
Enter No. 305

GuideTech

1-408-733-6555
www.guidetech.com
www.jitter.com
Enter No. 306

LeCroy Corp

1-800-453-2769
www.lecroy.com
Enter No. 307

Picosecond Pulse Labs

1-303-443-1249
www.picosecond.com
Enter No. 308

SyntheSys Research

1-650-364-1853
www.synthesysresearch.com
Enter No. 309

TDA Systems

1-503-246-2272
www.tdasystems.com
Enter No. 310

Tektronix

1-800-426-2200
www.Tektronix.com
Enter No. 311

Wavecrest

1-800-733-7128
www.wavecrest.com
Enter No. 312

OTHER ORGANIZATIONS MENTIONED

FCIA (Fibre Channel Industry Association)
www.fibrechannel.org/

SUPER INFO NUMBER

For more information on the products available from all of the vendors listed in this box, enter no. 313 at www.edn.com/info.