

DesignCon 2010

A New Method for Receiver Tolerance Testing Using Crest Factor Emulation

Ransom Stephens, Ransom's Notes
[ransom@ransomsnotes.com]

John Calvin, Tektronix Instruments
[john.c.calvin@tektronix.com]

Abstract

Emerging receiver tolerance tests require a calibrated mix of sinusoidal jitter/noise, intersymbol interference, and Gaussian random jitter/noise. We review the motivations for these requirements, survey the standard application methods and introduce a new one. Crest Factor Emulation addresses the difficulty of including random noise/jitter stress with both large crest factor and wide bandwidth in a way that mitigates inaccuracies of conventional techniques and decreases test time by orders of magnitude for specifications requiring Bit Error Ratios (BER) at 10^{-12} renders measurements to $\text{BER} < 10^{-18}$ possible in less than a minute, and provides designers a new diagnostic handle.

Author(s) Biographies

Ransom Stephens' company, Ransom's Notes, produces and presents content at every level of technical sophistication to help engineers advance to technology's cutting edge. He spent 13 years in basic-research laboratories and universities across the United States and Europe specializing in precise measurements of noisy signals. He is the author of over 200 articles in the electronics industry, science journals, and magazines, has introduced new measurement techniques for electrical and optical systems, invented methods for extracting signals from noise, led an engineering commando team, and served on high data-rate standards committees. Contact him at www.RansomsNotes.com.

John Calvin currently is the chairman of the Serial ATA International Organization's Interoperability working group. John is a principal engineer at Tektronix where he has worked for the last 15 years with a focus on high speed serial measurements solutions for industry standards. He has worked as a contributor to SATA testing since 2000. John holds a Bachelors Degree in Electrical Engineering from Washington State University and has been awarded 7 patents in measurement-related technology.

Introduction

The successful operation of links at multi GB/s data rates requires either an extraordinarily high quality transmission path or a receiver architecture capable of tolerating crosstalk, jitter, and amplitude noise. Over the last decade communications and computer standards such as PCI Express, Serial ATA and 10 GbE increasingly require that receivers include components that enable them to tolerate impairments. Clock data recovery and equalization circuits allow receivers to accommodate signals that may be so distorted that they are unrecognizable as digital signals.

A “receiver tolerance test” probes the ability of a receiver to work with a degraded input signal. The idea is to subject the receiver to a well defined worst case signal and require that it operate at a specified Bit Error Ratio (BER), usually 10^{-12} or lower.

We begin with a review of receiver tolerance testing by emphasizing how each stress – the compliant pattern, rise/fall time, Sinusoidal Jitter (SJ), Inter-Symbol Interference (ISI), Random Jitter (RJ) and noise, and Spread Spectrum Clocking (SSC) – plays a unique role in testing the clock recovery, equalizer and decision circuit elements of receivers. Then we examine difficulties in tolerance testing such as accuracy, expense and test time. Finally, we introduce the concept of Crest Factor Emulation and show how stressed receiver tolerance testing can be performed faster, more accurately and at lower expense than with traditional techniques.

Review of High Speed Serial Technology

Figure 1 shows the standard components of a serial data system.

The transmitter – serializer, reference clock and multiplier

Starting on the left end of Figure 1, eight parallel inputs are multiplexed into a serial data stream. At the transmitter, timing of the serial logic transitions is governed by a reference clock. The reference clock, which might be synthesized or based on a crystal oscillator usually at 100 MHz, is multiplied up to the data rate by a Phase Locked Loop (PLL).

Most of the emerging serial data standards include an option for Spread Spectrum Clocking (SSC). SSC is low frequency modulation of the clock which spreads the radiated energy of the system over a larger frequency band and makes it easier for the transmitter to pass government electromagnetic interference regulations. The most common form of SSC is 33 kHz triangle-wave frequency modulation with amplitude of a few thousand parts per million.

The primary sources of signal degradation at the transmitter are:

- (1) Phase noise of the reference clock, which increases as the square of the multiplication factor of the PLL [1] and is the primary RJ source [2].
- (2) Duty Cycle Distortion (DCD) caused by skew in the parallel input stream and/or asymmetry in the duration of adjacent clock cycles. The result is a difference in the duration of high and low logic levels. DCD doesn't combine linearly with other signal degradations [3].

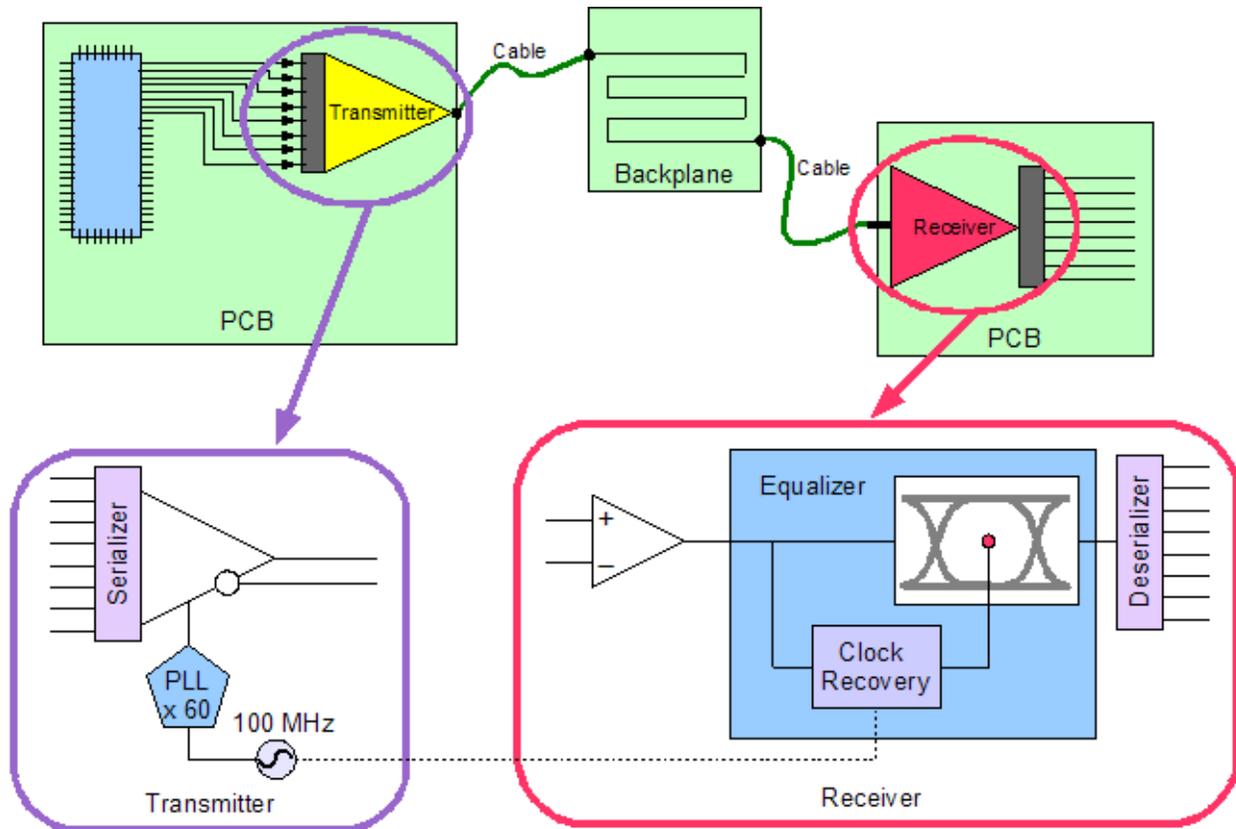


Figure 1: Serial data straw diagram.

The transmission path – backplanes and cables

The serialized data then propagates along a transmission path to the receiver. Differential signaling is almost always used to reduce stray fields and crosstalk – the net electromagnetic radiation of two neighboring transmission paths, each transmitting the opposite waveform, is very nearly zero, even when there is a common mode voltage.

The transmission path may include both Printed Circuit Board (PCB) and cables. Circuit boards and backplanes are usually made of Flame Retardant Type 4 fiberglass weave (FR-4). The cables might be high quality matched cables, but are more likely twisted pair. Figure 2 shows the progressive impairment of a signal as it traverses increasing lengths of transmission paths. The same thing happens as the data rate increases; ISI rapidly degrades a pristine digital signal into a nasty looking waveform, Figure 3.

The resistance of the conductor causes signal attenuation; the skin effect and dispersion cause non-uniform frequency response [4]. The dominant frequency component of a given bit is determined by the bit pattern that immediately surrounds it. In Figure 4 a simple Resistor-Capacitor (RC) time constant is used to illustrate ISI. In Figure 4a, where the data signal, 01010101, is a clock signal at half the data rate, the response of the circuit is sufficient for each bit to cross the logic-decision voltage threshold and be accurately identified. In Figure 4b, the data signal, 00001111, is a clock signal at one-eighth the data rate. Over the string of

Consecutive Identical Bits (CIB or CID) the time constant is sufficiently short for the signal to reach the voltage rail but too long for the signal to cross the voltage threshold during the first logic 1 following the string of 0s – resulting in an error: the 00001111 string would be identified as 10000111. A mixed example is shown in Figure 4c; here the 00001011 signal would be identified as 10000001.

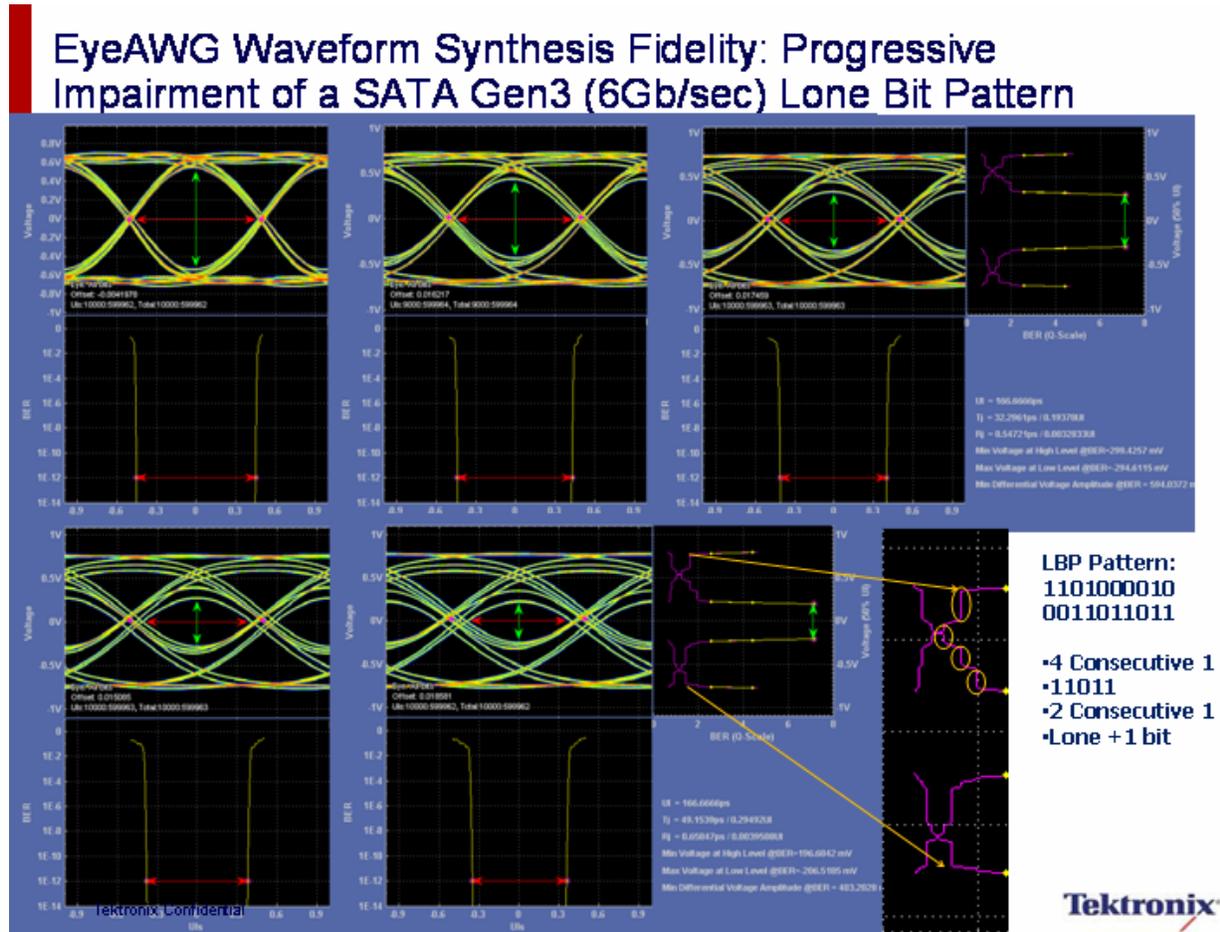


Figure 2: From left to right and top to bottom, the progressive impairment of a signal as it traverses successively longer transmission paths.

The ISI problem is compounded by impedance mismatches at connectors, which cause reflections, and unfortunate trace layout, which causes multipath interference. Further aggravating the situation, DCD and ISI do not combine in a linear way. This is one of the difficulties that face standards bodies as they attempt to define specifications that assure component interoperability.

The receiver – clock recovery, equalizer, decision circuit, and deserializer

At the receiver, a comparator first converts the differential signal which then enters a complicated circuit. In Figure 1, the equalizer is portrayed to encompass the clock recovery and decision circuits because in most designs they're interrelated.

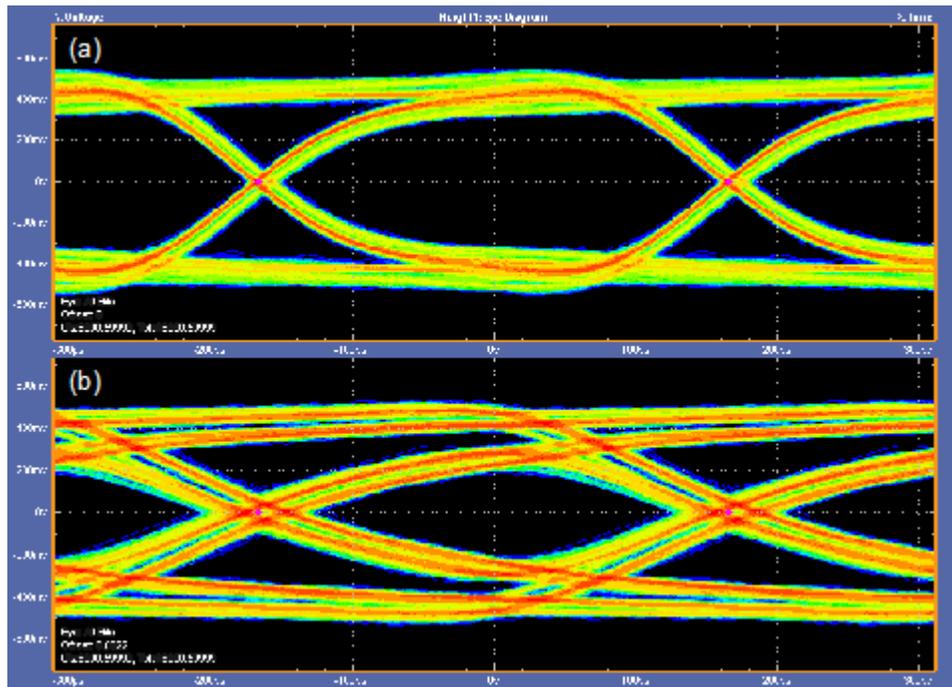


Figure 3: Two signals on the same length of PCB, (a) at 3 Gb/s and (b) at 6 Gb/s.

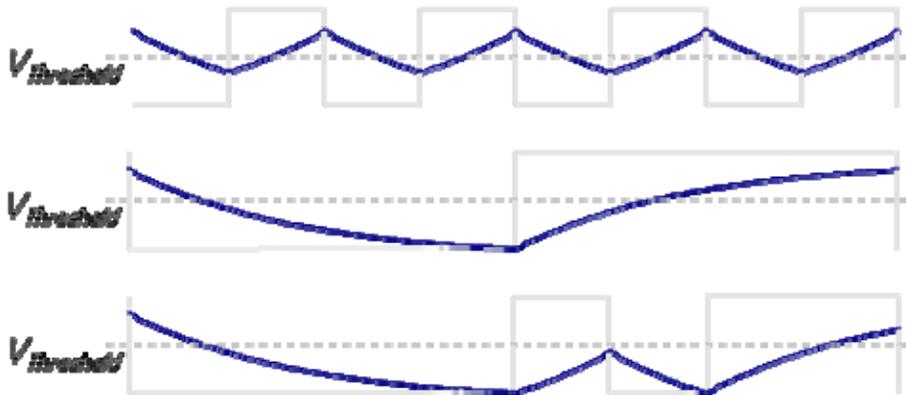


Figure 4: Simple examples of Inter-Symbol Interference (ISI). $V_{Threshold}$ is the logic-decision threshold, if the observed voltage is greater than $V_{Threshold}$ then the bit is identified as a 1, if less than $V_{Threshold}$, a 0.

Clock Recovery

Recovering the clock at the receiver, whether or not the reference clock is distributed (dashed line in Figure 1) provides the opportunity to effectively filter the signal's low frequency jitter. The clock recovery circuit, whether based on a PLL or a Phase Interpolator (PI) and whether or not the reference clock is distributed, fabricates a data-rate clock signal based on the timing of the incoming waveform logic transitions. The resulting clock includes the low frequency jitter that is on the data. This way, when the recovered clock sets the timing of the decision circuit, the

sampling point dances in rhythm with the lower frequency jitter. In other words: clock recovery serves as a high-pass jitter filter [2,5].

The clock recovery circuit must have sufficient bandwidth for the recovered clock to track SSC. The clock recovery bandwidth is typically $f_{data}/1667$ – which should be ample.

De-emphasis and equalization

At ever higher data rates, ISI closes the eye. Since ISI is caused by the media and geometry of the transmission path, we can quantify the effect and account for it. At the transmitter, signal transitions can be *de-emphasized*: increasing the voltage magnitude of bits prior to transitions increases their high frequency content mitigating frequency response properties of the transmission path sufficient to open the eye at the receiver, Figure 5.

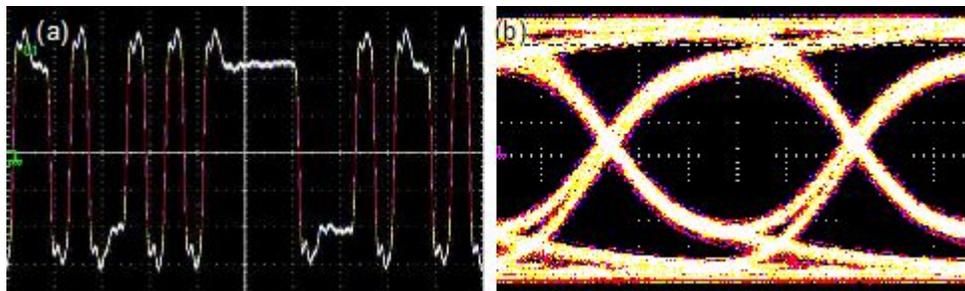


Figure 5: (a) de-emphasized signal at transmitter and (b) at receiver.

Receiver equalization involves a decision circuit that inverts the effects of the transmission path; we *equalize* the cause of eye-closure so that a degraded waveform can be interpreted. Many equalization techniques are being developed in addition to the standards: Feed-Forward Equalizer (FFE) and Decision Feedback Equalizer (DFE) [6].

Decision Circuit

The decision circuit decides whether a given bit is identified 1 or 0. A simple decision circuit uses a sampling point positioned at a voltage threshold and a time-delay defined with respect to the recovered clock. If the voltage is larger than the threshold, V_{th} , at the time-delay, x_{sp} , it must be a 1, else it's a 0. Of course the (x_{sp}, V_{th}) position of sampling point is not an ideal point, both setup and hold times and voltage-slice sensitivity consume jitter and noise margin.

Stressed Receiver Tolerance Testing in Serial Data Systems

The receivers in serial data technologies have to be specified to assure their interoperability with transmitters, clocks, backplanes, cables, et cetera from different vendors. To assure that a receiver is adequate, it is tested under the most stressful conditions consistent with the technology specification. If the receiver can perform under the worst-case conditions at or better than the specified Bit Error Ratio (BER) then it's good to go.

Errors occur when logic transitions fluctuate across the sampling point of the decision circuit. For example, if an ISI trajectory causes an error, the BER is at least the ratio of the number of occurrences of that trajectory to the length of the stress pattern – if ISI alone causes errors the BER is typically higher than 10^{-4} .

Most specifications require $BER < 10^{-12}$.

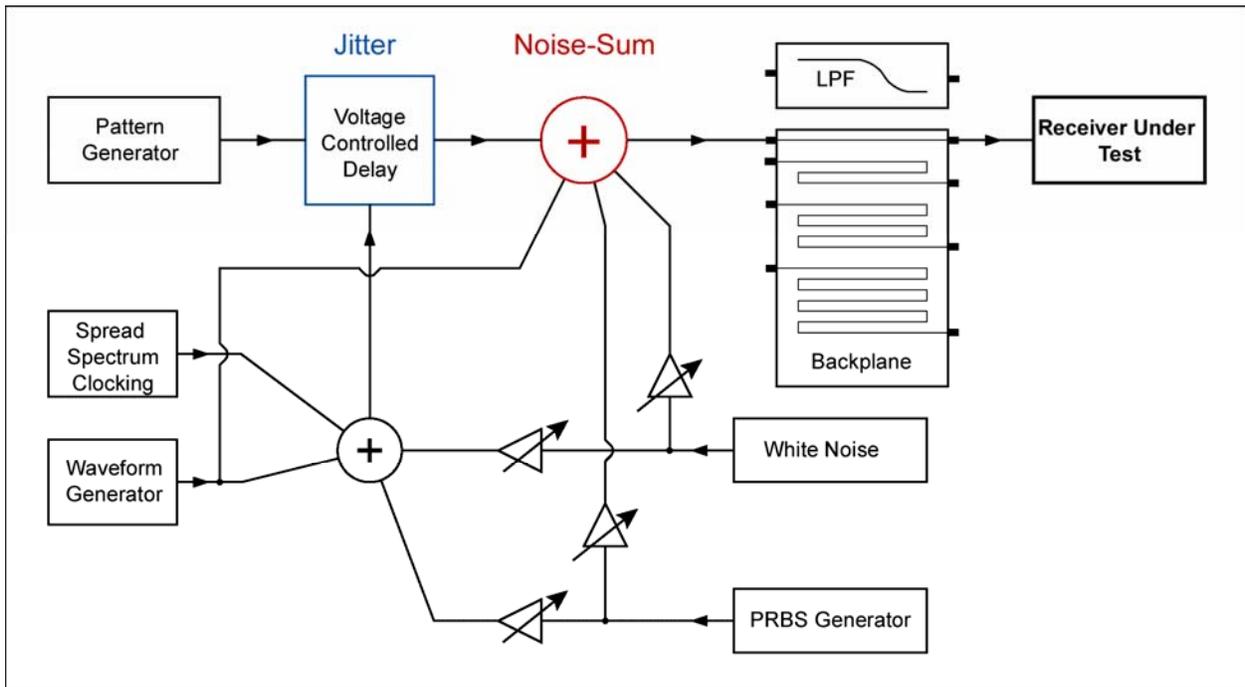


Figure 6: A common approach to implementation of receiver stress testing.

Typical stressed signal transmitters generate a stressful test pattern with adjustable levels of Sinusoidal Jitter (SJ), ISI, SSC, random white noise and random white jitter (RJ), and an impairment that emulates crosstalk. Figure 6 shows a common approach.

The test pattern

Most standards use specially designed compliance stress patterns like the Continuous Jitter Test Pattern (CJTPAT) [7].

Test pattern effect on baseline wander

Serial data receivers are typically AC coupled. If the average voltage of a signal varies significantly, the baseline voltage wanders. AC coupling is stressed by varying the mark density of the signal. The mark density is the average fraction of logic ones in the data: the ratio of the number of logic ones to the total number of bits. Most standards require a mark density of $\frac{1}{2}$; that is, equal numbers of ones and zeros averaged over the time scale of the system. The time scale is determined by the RC time constant of the receiver's AC coupling. Baseline wander is stressed by signals that vary from mark density of $\frac{1}{2}$ over times shorter than, but close to, the specified RC constant. For example, a signal like 0010 repeated many times, followed by 1101 repeated the same number of times provides back-to-back AC coupling stress while satisfying the overall mark density = $\frac{1}{2}$ requirement.

Test pattern effect on clock recovery

The clock recovery circuit functions by synthesizing a data-rate clock from the logic transitions of the input waveform. The more logic transitions, the easier it is for the receiver's clock recovery circuit. To this end, almost every standard requires that compliant signals have an

average transition density of $\frac{1}{2}$ – the transition density is the average number of logic transitions ($1 \rightarrow 0$ and $0 \rightarrow 1$) in a signal. Transition density of $\frac{1}{2}$ means that, on average, half of the bits in the signal precede a transition. The most difficult signal to derive a clock from is a long string of Consecutive Identical (CID) bits – the fewer transitions, the less clock content in the signal. To stress the receiver, then, we put CID strings in the test pattern that challenge the clock recovery bandwidth.

Many serial data technologies incorporate 8B/10B data-coding to guarantee a minimum number of logic transitions.

Test pattern effect on inter-symbol interference

Inter-Symbol Interference (ISI) determines the average waveform, or “trajectory,” of each bit in a signal. The trajectory is determined by the bit pattern immediately neighboring a given bit. The number of bits that affect the trajectory is given by the length of the pulse response. Thus, it is important to stress the receiver with every bit sequence that extends over the length of the pulse response – typically less than about eight bit periods.

Deterministic Jitter (DJ)

In some specifications, DJ has been re-coined Bounded High Probability Jitter (BHPJ) – a more cumbersome but also more descriptive term – we use DJ here.

DJ can be separated into two categories: that which is correlated to the signal, like ISI, and that which is not, like Sinusoidal Jitter (SJ) [3]. As the name implies, SJ is sinusoidally varying phase modulation. ISI causes both jitter and voltage noise – eye closure in both the vertical and horizontal directions – and occurs at rational fractions of the data rate and with amplitudes that depend on the loss character of the transmission path.

Most standards require a combination of SJ and ISI, though some just require DJ and let the user choose the type. Emerging standards and, it is safe to expect, standards yet to come will not be so forgiving.

Several standards (e.g., SATA, SAS, DisplayPort, RapidIO, USB 3.0) require that SJ be applied according to a template, Figure 7. To probe the clock recovery frequency response, multi Unit Interval (UI) amplitude SJ is applied below the clock recovery bandwidth and fractional UI amplitudes above. It’s worth keeping in mind that the actual frequency response includes phase information that is lost in this test.

SJ is commonly applied, Figure 6, by modulating the data-rate clock of the pattern generator. Alternatively, SJ can be applied by digitally generating the test signal with an Arbitrary Waveform Generator (AWG). Virtually any SJ amplitude and frequency can be combined with the pattern and other stresses to form the transmitted waveform.

Since ISI can be accommodated with equalization, new and emerging standards, like USB 3.0, have increasingly specific ISI requirements. The requirements are given by equivalent trace lengths on “typical” FR-4 PCB and cable. Since no two traces or cables are identical, the “equivalence” is defined through *S*-parameter masks. The idea is to permit as much backplane/cable design freedom as possible in a specification that also provides a receiver designer everything necessary to implement a sufficient equalization scheme.

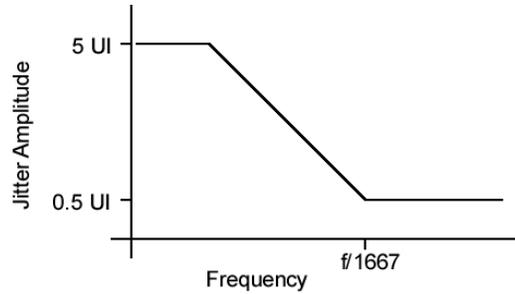


Figure 7: Applied SJ template for Serial ATA.

The traditional approach to applying ISI, Figure 6, uses pre-calibrated pieces of backplane. There are a couple of serious problems with this approach. The calibration depends on the condition of the connectors and quality of the connections and, to a lesser extent, by temperature and humidity. A more reliable and flexible approach is to synthesize the ISI directly from the S -parameters and transmit the calculated waveform with an AWG. In this approach the “equivalent length” can be varied by scaling the S -parameters, which also enables measurement of the dispersion penalty tolerance of the receiver (i.e., the maximum transmission path it can handle).

Random Jitter (RJ)

RJ is caused by the sum of many small effects like thermal oscillations in the clock, tiny variations in trace widths and conductor radii in cables and so forth. It is universally assumed to follow a Gaussian time domain distribution, Figure 8, characterized by the standard deviation, σ , Figure 8, and to have a white frequency spectrum. The Gaussian assumption is the obvious statistical choice [8]. The white frequency spectrum is a less obvious choice because the primary cause of random jitter, the oscillator driving the reference clock at the transmitter, tends to have a pink, flicker-dominated frequency spectrum [9].

Since RJ is unbounded, a peak-to-peak value for jitter can only be defined in reference to a Bit Error Ratio, hence the definition of Total Jitter at a Bit Error Ratio, TJ(BER) [10].

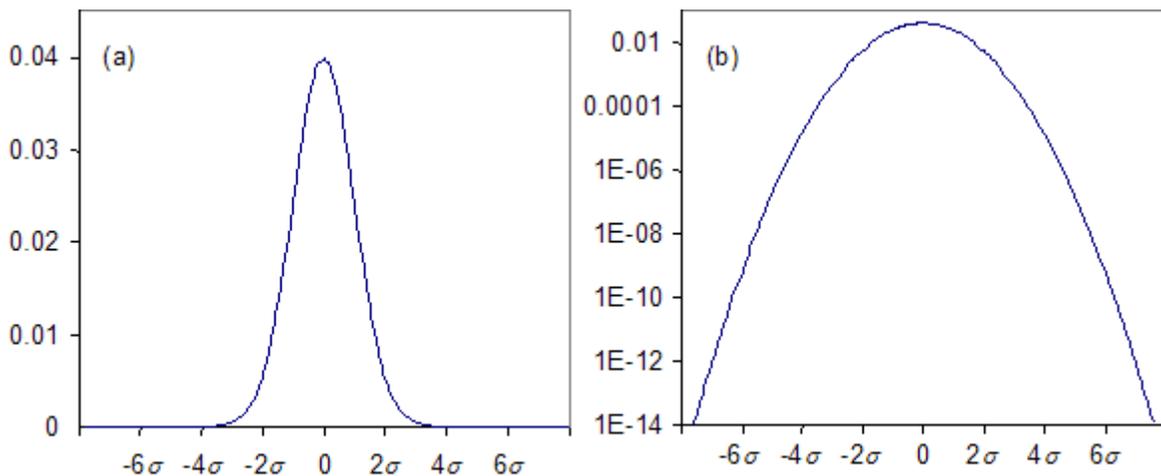


Figure 8: A Gaussian distribution on (a) linear and (b) logarithmic scales.

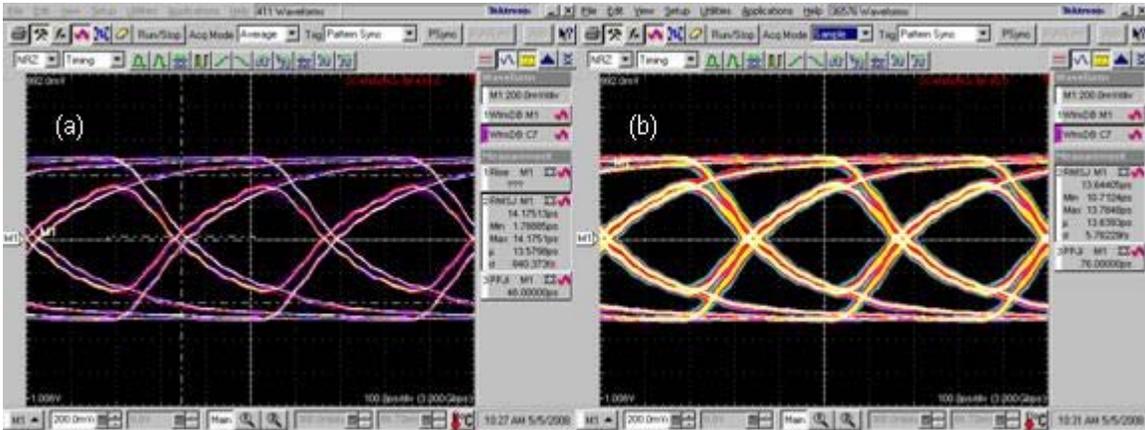


Figure 9: A signal with the same ISI but (a) very low and (b) significant RJ.

The common approach for applying RJ is to use a random noise source, Figure 6. Commercial noise sources [11] are based on a reverse-biased diode emitting amplified Johnson noise. They provide signals that are statistically consistent with a Gaussian distribution of ample crest factor and have white frequency spectra.

The effect of RJ is to smear signal trajectories. Consider a signal with only ISI, Figure 9a. The eye-diagram trajectories of every bit are clearly defined. Adding RJ smears the trajectories, Figure 9b.

The role of RJ in stressed receiver tolerance tests

Where DJ provides an immediate high probability challenge to a receiver's performance – hence the synonym, Bounded High Probability Jitter – RJ is more subtle. In smearing bit trajectories, RJ exacerbates receiver response to all types of DJ.

RJ presents a small but important additional stress to the clock recovery circuit. It causes the timing of every DJ-defined edge to randomly fluctuate a small amount. Small random fluctuations stress the ability of the clock recovery circuit to remain steady and locked.

We can derive from Figure 8 that the vast majority of RJ stresses are very small: over 99.7% of RJ occurrences are less than 10% of the peak-to-peak DJ. The vast majority of RJ effects are small amplitude fluctuations that occur with high probability.

Compliance requires testing to a BER of 10^{-12} . The fundamental difficulty in testing to such low probabilities is obtaining a stressed signal generator (sig-gen) with sufficiently large crest factor.

Crest factor and the probability of large fluctuations

The crest factor characterizes the maximum divergence of a distribution from its mean in terms of its standard deviation. For a voltage distribution, it is the peak voltage divided by the rms voltage,

$$\text{Crest Factor} \equiv \frac{|V_{\text{Peak}}|}{V_{\text{rms}}}$$

For a finite set of random variables that follow a Gaussian distribution, it is half the maximum spread of the observed distribution divided by the standard deviation,

$$\text{Crest Factor} = \frac{\frac{1}{2}(x_{\max} - x_{\min})}{\sigma}$$

An ideal random noise source generates a Gaussian distribution whose tails extend infinitely in both directions. Real, hardware random noise sources have limited but sufficient crest factors; the longer they run, the larger the observed crest factor.

Here's some terminology: an "outlier" is an "instance" of jitter at low probability. The probability of the occurrence of an outlier decreases the more outlying the instance.

Consider Figure 10, BER as a function of sampling-point time-delay position for a 6 Gb/s SerialATA Gen-3 stress signal (a bathtub plot). It is essentially the cumulative distribution function of the jitter probability density function – the integral of the convolution of the prescribed pattern and the RJ and DJ stresses [12]. The structure of the bathtub curve is caused by a displacement inward at high probability dictated by DJ followed by long smooth tails caused by RJ [13]. The eye-opening at different BERs is given by the distance between the two curves and TJ(BER) is given by the nominal bit period minus the eye opening at that BER.

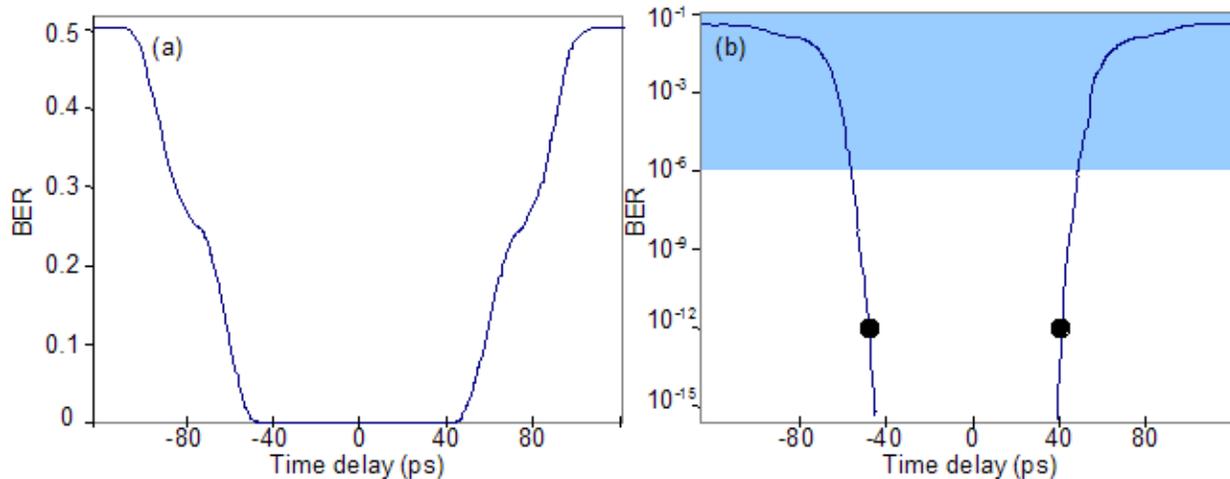


Figure 10: The bit error ratio as a function of the time-delay, BER(x) – a bathtub plot – for the SerialATA stress signal on (a) linear and (b) logarithmic scales.

The most important tests are in the high probability shaded region of Figure 10b where the high probability low amplitude random jitter and bounded high probability jitter (a.k.a. DJ) dominate. The effect of DJ at high probabilities, here for BER > 10⁻³ or so and, generally, for BER > 10⁻⁵, causes most of the structure in the curves. For example, at the time-delay just larger than the SJ amplitude, we expect BER to experience a rapid decline. We expect similar abrupt drops in BER at the time-delays corresponding to the largest ISI variations. All this structure occurs in the shaded region, but once we get to lower probabilities, BER < 10⁻⁶ or so, the unbounded low probability RJ fluctuations give the curves their predictable unwavering smooth decay.

Below the shaded high BER region, the prescribed stress signal exhibits no appreciable structure – just smooth tails decaying fast.

The distance between the two points at BER = 10⁻¹² in Figure 10b gives the maximum compliant receiver sensitivity. That is, a perfect receiver would have a jitter margin given by the distance

between those two points and operate with $BER \ll 10^{-12}$; conversely, if the decision circuit setup and hold time is exactly the distance between those two points, then the receiver has no margin, and would operate at $BER=10^{-12}$, similarly, if the setup and hold is larger, than it would operate with $BER > 10^{-12}$ and fail the test.

Receiver compliance testing requires a stress sig-gen capable of generating a signal with a crest factor sufficient to probe the BER requirement of the technology standard; BER of 10^{-12} requires a 14σ spread in the RJ distribution: a crest factor of 7 or about 8.5 dB.

Compliance requires that $BER < 10^{-12}$ be verified on at least two frequency-amplitude points of the SJ template with appropriate confidence, usually a 95% confidence level upper limit is satisfactory. If we ignore the fact that certain logic transitions in the compliance pattern exert a greater stress than others and no errors occur, we need samples of 3×10^{12} bits [14] for each point in the SJ template.

But remember, each transition has a specific trajectory determined by ISI and each of those is modified by SJ. What if a large instance of RJ occurs on the logic transition with the highest level of ISI and causes an error? What if that same large instance of RJ, had it occurred on a transition with the median level of ISI wouldn't cause an error? We might reject that receiver rather than wait another 20 minutes to get enough non-errored bits to satisfy us that the actual BER is less than 10^{-12} . Or what if a large RJ instance occurred on a transition with almost no ISI and we pass a receiver that should have failed?

Truly random fluctuations are uncontrollable and irreproducible – unless there is no other way to produce a compliant signal, they have no place in a test lab.

Random tests require substantial statistical samples of data to attain accurate confidence levels. If we wait long enough, we'll get a sufficient statistical sample to test the receiver all the way down the bathtub curve to 10^{-12} (many minutes), 10^{-15} (days), or even, a favorite for some, 10^{-18} (years).

Crest factor emulation

In a test lab, it is imperative that engineers understand the initial conditions of a test. A truly random signal cannot be reproduced or controlled.

A better way to apply the requisite RJ is to introduce pseudorandom rather than genuinely random noise and to do so in a calculated fashion. The essence of crest factor emulation is to synthesize RJ and introduce the large amplitude, low probability instances at $BER=10^{-12}$ where they are most useful.

That is, crest factor emulation enables the use of a deep memory AWG to synthesize the entire stressed waveform – including SJ, ISI, SSC, and RJ.

Figure 11 shows a distribution of a million random variables that follow a Gaussian distribution including a single 10^{-12} probability outlier indicated by the red square at 7σ . Introduction of a low probability, large amplitude instance of jitter has no effect on the frequency spectrum – no more or less than its random occurrence would.

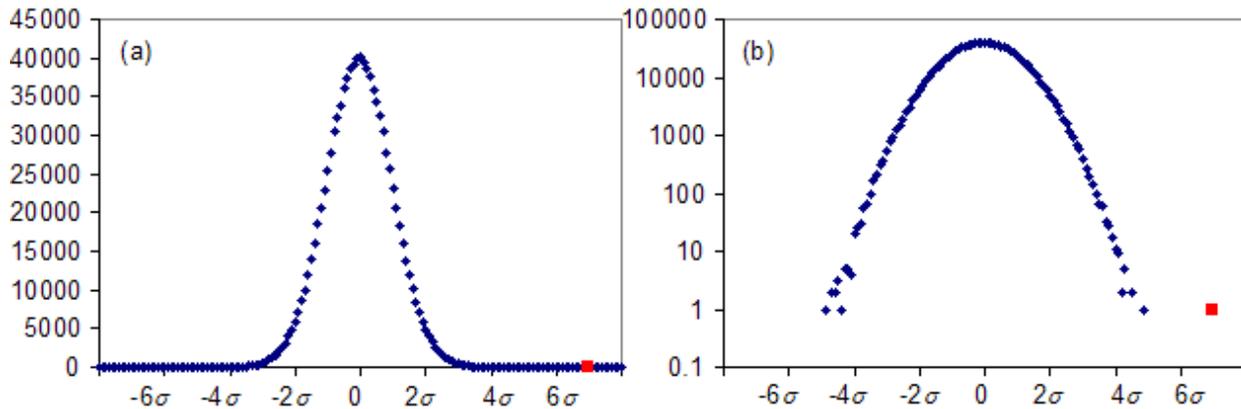


Figure 11: Gaussian distribution composed of 10^6 instances with a single low probability, 10^{-12} , outlier on (a) linear and (b) logarithmic scales.

The difference between the two waveforms in Figure 12 is indicated by the transition marked in red in Figure 12b – this is the transition with the 10^{-12} probability outlier – a displacement of 0.18 UI. Strictly speaking, the best specific transition to apply the outlier is that which experiences a level of DJ closest to the dual-Dirac model-dependent peak-to-peak DJ – see Ref [13] for a discussion of this rather annoying technicality. A reasonable and more intuitive alternative is to apply the outlier to the transition that experiences both the median level ISI and half the SJ amplitude. Median level ISI and half amplitude SJ is the most probable instance of DJ. We could put the outlier on the bit with the worst-case ISI, but then we'd be probing a BER smaller than 10^{-12} and might reject a compliant receiver. Obviously, you have the ability to decide for yourself where to put the outlying RJ instance. Anywhere in the pattern would be compliant to the letter of specifications that currently exist.

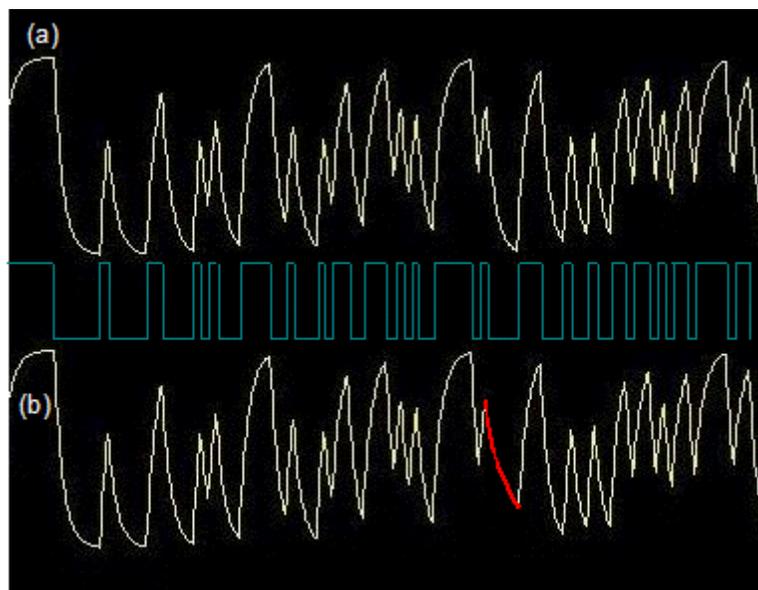


Figure 12: A segment of the SATA Gen 3 stress waveform, (a) without and (b) with an instance of 7σ RJ.

If the receiver can tolerate this carefully controlled worst case signal, then its BER is assured to be less than 10^{-12} – in a test that takes a few seconds. In fact, we can do the same thing to test tolerance down to 10^{-15} or 10^{-18} without any increase in test time.

If a receiver fails the compliance test – whether at $\text{BER} = 10^{-12}$ or 10^{-18} – the failure can be isolated, repeated, and diagnosed whether it results from an extrema in ISI, PJ, and/or RJ; whether it is in the part of the test pattern designed to probe baseline wander, clock recovery, or the decision circuit.

To use crest factor emulation, an AWG with sufficiently deep memory is required (plus the software necessary for synthesizing a waveform with all the necessary stress). Memory depth is ultimately constrained by the length of the stress pattern. DCD and ISI don't limit memory depth, but SJ does. The stress pattern must be repeated enough times that the sinusoidal swing is represented on every unique bit trajectory. Without crest factor emulation, the memory depth necessary to accommodate RJ would be in the tens of terabytes, with crest factor emulation, for most serial standards, it is in the tens of megabytes.

From the perspective of the receiver, there is no difference between a stress signal that includes the RJ shown in Figure 11 with outliers installed by hand, and a signal generated with a hardware noise source that includes a 7σ outlier – not in the time or frequency domains. Since the stress signal is prescribed by the standard, there are no surprises at low BERs. As long as the memory depth of the AWG is sufficient to amply accommodate the stress pattern and prescribed DJ, pseudorandom RJ with precisely placed outliers is indistinguishable from the conventional technique.

There are three differences between tests with the crest factor emulated in a synthesized waveform and tests with a hardware noise source: test time, access to low probabilities and repeatability.

Conclusion: The best way to perform a stressed receiver tolerance test

Test engineers need control of their signals. Without control, tests cannot be reproduced for verification and systematic uncertainty destroys precision.

The following recipe for compliant stressed receiver tolerance testing greatly simplifies the whole process. It also indicates the minimum memory depth required of the AWG:

1. Synthesize the appropriate binary stress pattern with the specified rise/fall time.
2. Apply ISI to the resulting waveform according to S -parameters that correspond to the backplane/cable DJ requirement.
3. Apply SJ with amplitude and frequency prescribed by a template like the one in Figure 7 to a couple thousand repetitions of the signal achieved in step 2.
4. Generate and apply pseudorandom Gaussian RJ to every transition in the signal except the one with median DJ – i.e., the edge with median ISI and $\frac{1}{2}$ amplitude SJ.

At this point we have a signal capable of testing every point in the shaded region of Figure 10b for the SJ amplitude/frequency we chose in step 3.

5. Apply a single 7σ displacement to the transition with median DJ deviation.

Now we have the signal at the $\text{BER}=10^{-12}$ point of Figure 11b with the edge in Figure 12b.

6. If the receiver tolerates the waveform, return to step 3 and apply SJ from another point on the template – repeat three times, once for SJ at a frequency within the clock recovery bandwidth/below the roll off, once for SJ frequency at the end of the roll off, and once at a frequency above the clock recovery bandwidth.

7. If the receiver fails one of the three tests, find the errored bit(s) in the synthesized waveform and start debugging.

Each test can be performed in seconds. This is not an approximation and it can be performed to arbitrarily low BERs with no increase in test time.

Crest factor emulation is compliant to any standard that does not explicitly specify the equipment required for testing.

References

- 1 - W. P. Robins, *Phase Noise in Signal Sources Theory and Applications*, (Peter Peregrinus Ltd., 1982).
- 2 - Ransom Stephens, *Jitter 360 Knowledge Series, Part 6: Reference Clock Jitter and Data Jitter*, published and distributed by Tektronix Corp., 2007, available at www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12844&lc=EN.
- 3 - Ransom Stephens, *Jitter 360 Knowledge Series, Part 3: All About the Acronyms: RJ, DJ, DDJ, ISI, DCD, PJ, SJ, ...*, published and distributed by Tektronix Corp., 2006, available at www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12839&lc=EN.
- 4 - Ransom Stephens, *Answering Next Gen Serial Challenges, Part 2: The Case of the Closing Eye*, published and distributed by Tektronix Corp., 2008 available at www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12854&lc=EN.
- 5 - Ransom Stephens, *Jitter 360 Knowledge Series, Part 5: Clock Recovery in Serial-Data Systems*, published and distributed by Tektronix Corp., 2007.
- 6 - Ransom Stephens, *Answering Next Gen Serial Challenges, Part 3: Acting on an Impulse, Equalization and Emphasis*, published and distributed by Tektronix Corp., 2008, available at <http://www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12855&lc=EN>.
- 7 - See the FibreChannel web site, www.t11.org.
- 8 - See any standard probability and statistics text, for example, Anthony J. Hayter, *Probability and Statistics For Engineers and Scientists*, 2nd ed. (Brooks/Cole Publishing, 2002).
- 9 - J.A. Barnes and D.W. Allan, "A Statistical Model of Flicker Noise," Proc. IEEE, vol. 43, pp. 176-178, February 1966; D. Halford, "A General Mechanical Model for Spectral Density Random Noise with Special Reference to Flicker Noise $1/|f|$," Proc. of the IEEE, vol. 56, no. 3, March 1968, pp. 215-258; Sullivan, D. B., et al. (eds.), *Characterization of Clocks and Oscillators*, National institute of Standards and Technology Technical Note 1337, National institute of Standards and Technology, Boulder, CO 80303-3328; R.B. Blackman and J.W. Tukey, *The Measurement of Power Spectra*, New York: Dover, 1958.
- 10 - Ransom Stephens, *Jitter 360 Knowledge Series, Part 1: The Meaning of Total Jitter*, published and distributed by Tektronix Corp., 2006, available at www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12836&lc=EN.

-
- 11 - Ransom Stephens and Bob Muro, *Characterization of Gaussian Noise Sources*, DesignCon 2008, Proceedings, International Engineering Consortium, 2008.
 - 12 - Ransom Stephens and Marcus Müller, Chapter 2 of *Digital Communications Test and Measurement: High-Speed Physical Layer Characterization*, edited by Dennis Derickson and Marcus Müller, (Prentice Hall, Boston, 2007).
 - 13 - Ransom Stephens, *Jitter 360 Knowledge Series, Part 2: What the Dual-Dirac Model Is and What it Is Not*, published and distributed by Tektronix Corp., 2006, available at www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=wpp&ci=12838&lc=EN.
 - 14 - Ransom Stephens, Marcus Müller and Russ McHugh, *Measurement of Total Jitter at Low Probability Levels Using the Optimized BERT Scan Method*, DesignCon2005 Proceedings, International Engineering Consortium, 2005.