How Stressing a Receiver Removes Stress from a Designer

Ransom Stephens, Ph.D.

December, 2007

Abstract:

The ideal way to determine if a receiver is fit for an application is to test it with the worst signal it is likely to ever see. Techniques for testing the tolerance of receivers to different types of stress have been around for decades and have proven their utility at rates up to 10 Gb/s. Now, stress tests applied in serial data standards are stretching the concept to its limit. We show how different types of stress arise in systems, how they can be applied to a receiver for diagnosing problems and then look at a typical serial standard to see how stress tests are used for interoperability compliance.

The idea of stress-testing receivers is simple: If a receiver can operate with a low enough Bit Error Ratio (BER) when it's being fed the worst-possible signal, then it is guaranteed to operate even better with a normal signal. One problem with stress-testing always seems to arise when working with interoperability standards: compliance tests don't provide the information needed to diagnose problems. A bigger problem though, is how can one define the "worst-possible signal"? One receiver might be robust to high amplitude, white random noise, and random jitter, but cough at the first sign of periodic noise in a particular frequency band.

In this paper our goal is to understand how to execute stress tests for both diagnostics and compliance in serial technologies. The question is: what types of stress should we apply for the new serial-data technologies? The answer lies in the challenges that serial-data technologies like SATA, 10 GbE, PCI Express, et cetera, present; the use of asynchronous architecture means that clock recovery circuits need to be robust to challenging data sequences and since they employ long stretches of FR-4 circuit board at multi-gigabit rates, receivers must be able to work with ever smaller eye openings. But I'm getting ahead of myself.

Signal Degradation

A catalog of signal degradation would make a fine reference book. We'll use a big diagram to sort it all out. One glance at Figure 1 and you'll immediately know what I mean by eye stress. I got a headache just

drawing it. Seriously though, with the setup in Figure 1 we can produce the necessary stresses for testing and debugging receivers.

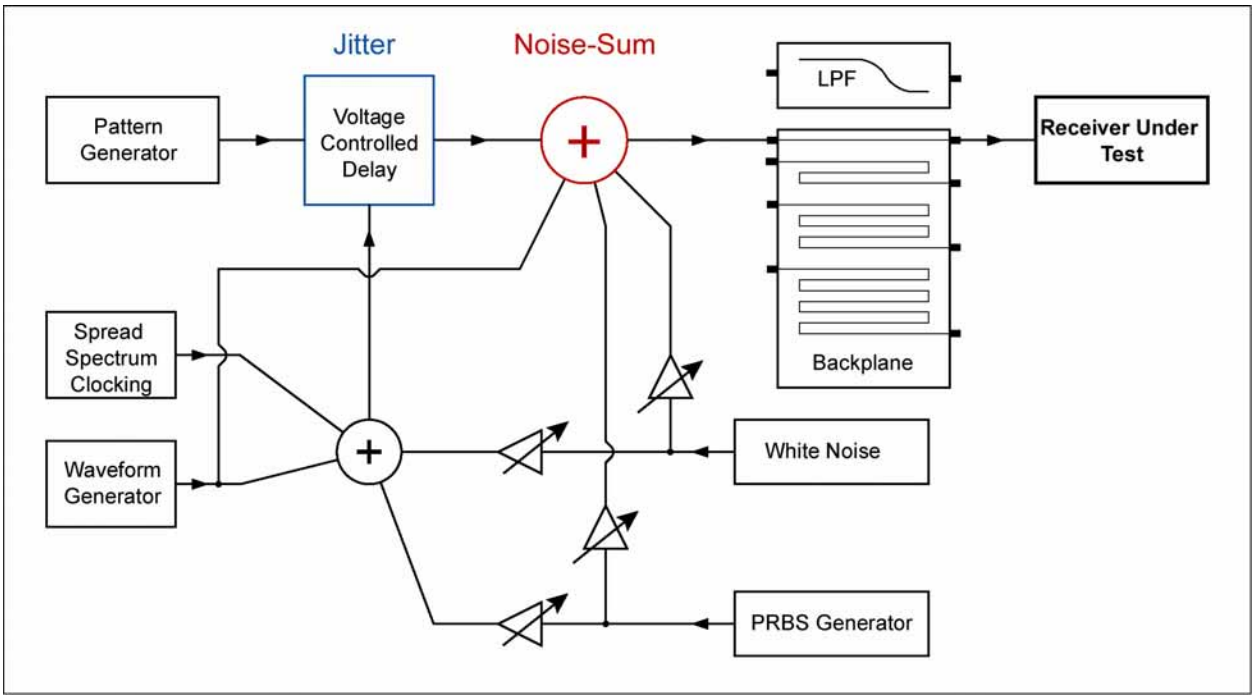

Figure 1: Big diagram of signal degradation – ultimate how to stress an eye reference.

We start in the top left with an instrument-grade pattern generator. In this model the pattern generator provides a perfect signal with whatever rise/fall times we desire and no consequent inter-symbol interference. The pattern generator is also the source of two key stresses. First, different patterns stress the receiver in different ways. Typical receivers in serial-data systems recover a clock signal from the data itself. The recovered clock is used to determine the times at which to sample the signal and identify bits. A pattern with a long string of Consecutive Identical (CID) bits, i.e., no logic transitions, can cause clock recovery failure. A pattern where the average number of ones or zeros differs from 50% can cause the receiver slice-threshold to drift. Interoperability standards define many different stress patterns. Alternating patterns (0101…) and Pseudo-Random Binary Sequence patterns of different lengths and combinations of the two are the most common. For Example, PCI-Express has a compliance pattern consisting of a combination of symbols that intersperse alternating patterns between patterns with CID bit strings: 0011111010 1010101010 followed by its inverse.

The second stress that the pattern generator causes, in our model anyway, is Duty Cycle Distortion (DCD) – a difference in the widths of logic ones and zeros, also called Pulse Width Distortion (PSD). DCD

is a special type of periodic jitter that is correlated to the data. We'll return to DCD in a few paragraphs when we discuss Inter-Symbol Interference (ISI).

The voltage-controlled delay just to the right of the pattern generator allows us to apply jitter to the signal. As its name implies, a voltage applied to the voltage-controlled delay causes a delay in the signal that is proportional to the applied voltage. For example, Figure 2, if we set the arbitrary waveform generator at the bottom left of the diagram to transmit a sine wave, then the output signal has Periodic Jitter (PJ) at the frequency of the sine wave and peak-to-peak jitter proportional to the voltage of the sine wave.

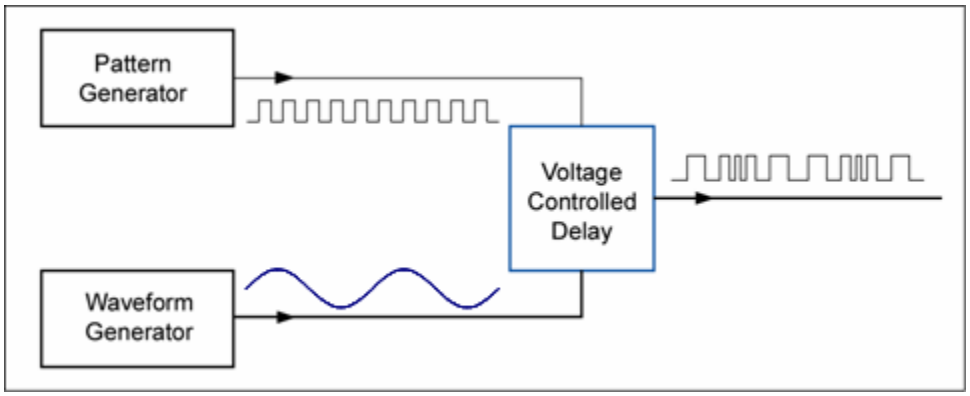

Figure 2: Voltage-controlled delay.

To the right of the voltage-controlled delay in Figure 1, the summing element that I labeled "noise-sum" is used to apply amplitude noise to the signal. The sine wave that caused the PJ when applied to the voltage controlled delay, would cause Periodic Noise (PN) when summed with the generated signal, Figure 3.

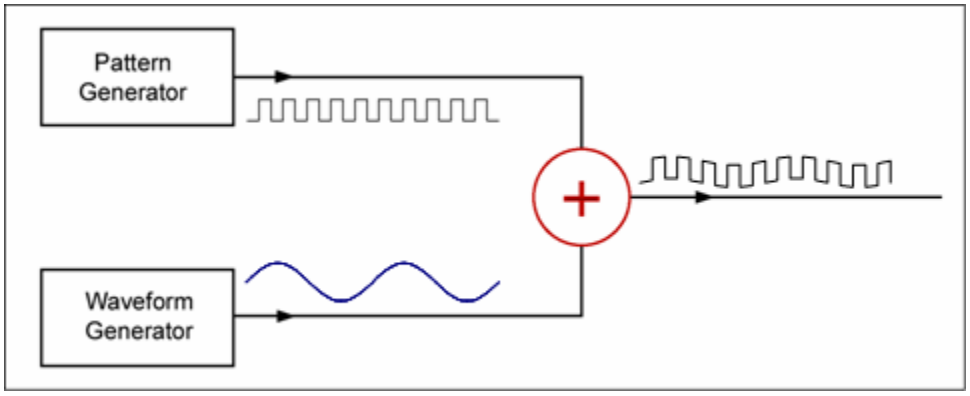

Figure 3: Waveform with and without Periodic Noise (PN).

PJ and PN are caused by electromagnetic interference, most notably power-supply feed through, but PJ and PN can be used to diagnose many problems in receivers. Some receiver designs, especially digitally-based designs like Phase Interpolators, can exhibit odd behavior in different frequency bands.

There are two random noise sources in Figure 1. The white noise generator applied to the voltage controlled delay puts Gaussian Random Jitter (RJ) and/or Random Noise (RN) on the generated signal. Remember, RJ follows an unbounded Gaussian distribution that is specified by its width or rms value, σ. Tthat it is *unbounded* means that a peak-to-peak measurement of RJ is ill-defined. When the white noise generator is applied to the noise-sum, the signal exhibits random voltage noise; due to the non-zero rise/fall times, the random voltage noise causes both Random Noise (RN) and Random Jitter (RJ).

The second random noise source in Figure 1 is more obscure. A Pseudo-Random Binary Sequence (PRBS) generator is used to apply random, but bounded noise. The timing of the PRBS signal shouldn't be correlated with the pattern generator and the rate of the PRBS signal needn't be near the rate of the signal. The PRBS waveform applied to the noise-sum puts stress on the signal similar to what we might expect from crosstalk in a real system. PRBS applied to the voltage-controlled delay can be configured in many different ways. If the PRBS generator is controlled by a clock that is neither frequency nor phase-locked to the pattern generator, then the consequent jitter doesn't correspond to any obvious physical stress in a real system, though it could be used to test a receiver in a nightmare situation with untraceable unexplainable Bounded Uncorrelated Jitter (BUJ). If the PRBS generator is phase-locked to the pattern generator, then its jitter would emulate a system that had a defective multiplexer/demultiplexer near the Receiver Under Test (RUT).

Perhaps the most interesting (and confounding) stress is caused by the backplane just to the left of the RUT in Figure 1. The backplane has several different trace-lengths that can be used to apply different levels of Inter-Symbol Interference (ISI). ISI is caused by the frequency and attenuation response of the transmission medium, e.g., cables, circuit board traces, et cetera. We can think of the channel as a chain of RC circuits. The farther the signal propagates, the greater is the signal distortion. Figure 4 shows the ISI resulting from increasing distances of propagation through FR-4 backplane.

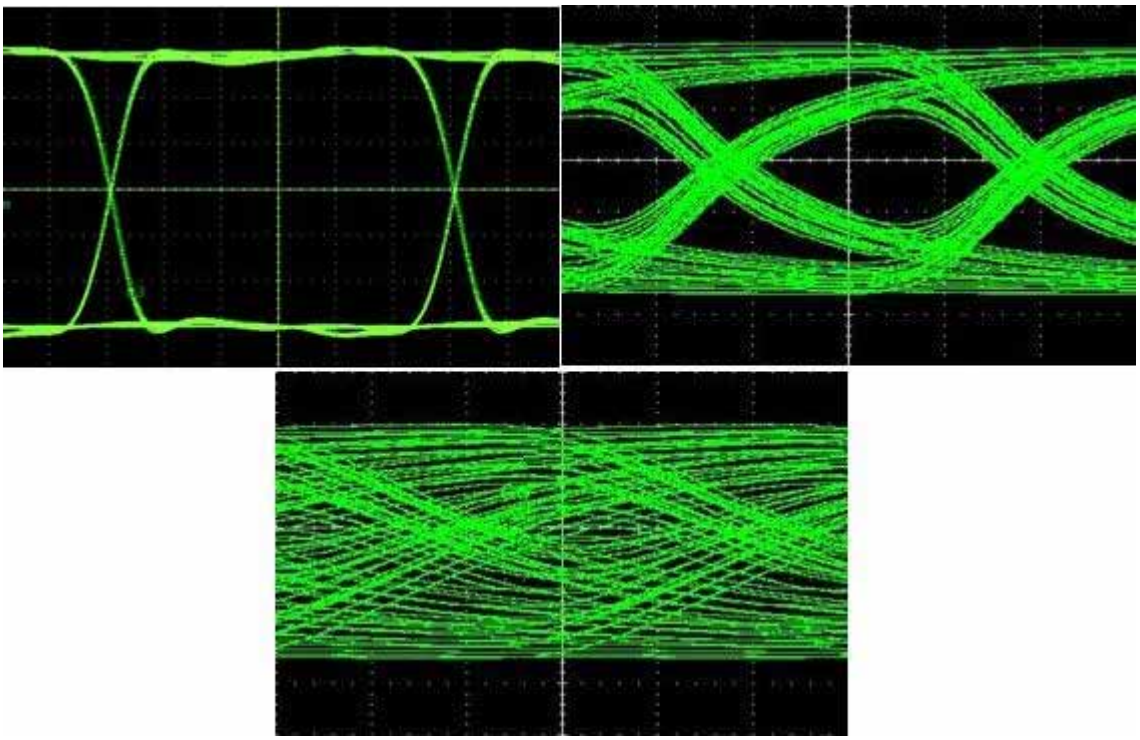

Figure 4: Screen shots of waveforms and eye diagrams with different amounts of ISI, (a) no ISI, perfect signal, (b) medium ISI, (c) high ISI.

There is an added complication with ISI. When ISI is applied to a signal with DCD, the two signals do not combine in a well-behaved fashion.

In addition to the imperfections described above, the receiver must also operate under the stress of the technology *features*. The Spread Spectrum Clock (SSC) generator in Figure 1 *spreads the spectrum* of the signal. That is, by applying low frequency (typically 33.3 kHz) low amplitude (0.5% of the data rate) periodic jitter to the signal, the transmitted power is smeared about its peak. While the net transmitted power is unchanged, the power transmitted at the nominal frequency is reduced, making it easier to pass government regulations.

Another common feature of serial-data technology is de-emphasis: Immediately following a logic transition the signal is driven at the full voltage swing, but when consecutive identical bits follow the transition, the transmitted voltage is reduced or *de-emphasized*, Figure 5. Well implemented de-emphasis reduces the impact of ISI at the receiver, so de-emphasis should mean lower stress on the receiver. Still, since a compliant transmitter could apply more de-emphasis than is ideal, the receiver should be tested with de-emphasis signals.

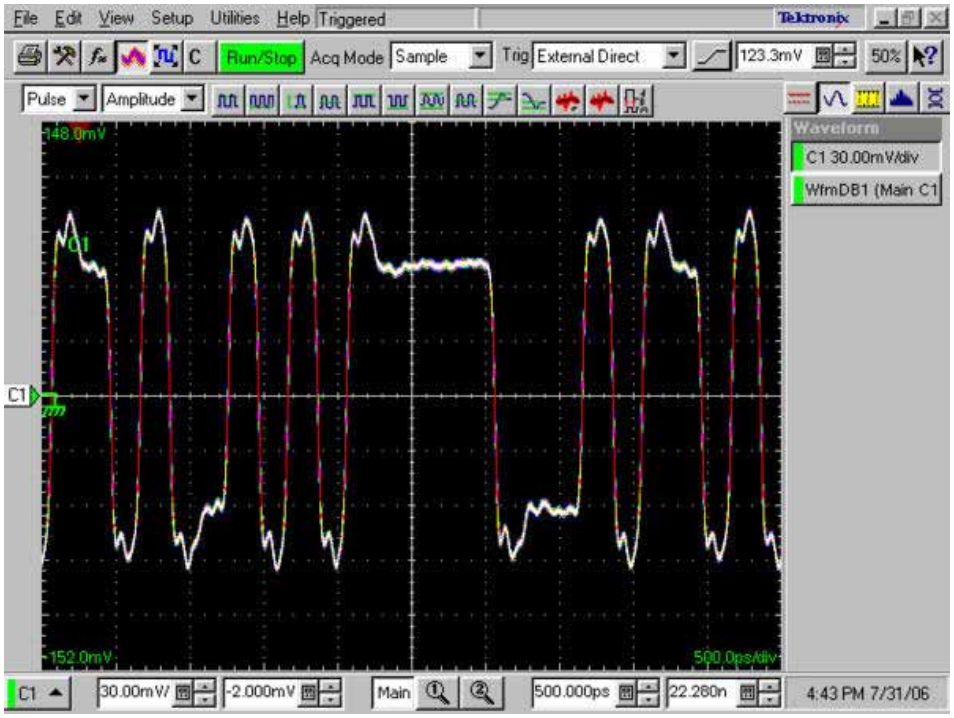

Figure 5: (a) Waveform of a de-emphasized signal.

To summarize, the system shown in Figure 1 is capable of applying any stress that you would need to test the tolerance of a receiver, but it begs several questions.

1. Do I really have to buy all the equipment in Figure 1?
2. How do I know how much jitter and noise I'm applying in any configuration?
3. Doesn't this thing have to be calibrated?
4. How do I know what the "worst-case compliant" signal is?

The answers are, "no", "you don't", "yes", and "see next section."

The solution to Figure 1 is that the whole mess can be emulated by a high-end arbitrary waveform generator. All the signals in Figure 1 can be reproduced in well controlled amounts and combinations by synthesizing the effects in a high-bandwidth, high-resolution, deep-memory, programmable arbitrary waveform generator. Tektronix AWG7000 series with direct synthesis software converts Figure 1 into Figure 6.

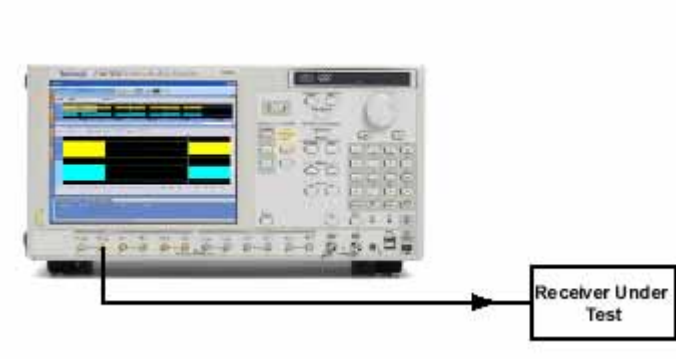

Figure 6: Implementation of Figure 1 with a high-bandwidth, high-resolution, deep-memory, programmable arbitrary waveform generator.

Stressed-Eye Compliance Tests

The idea of testing a receiver by subjecting it to the worst signal allowed by a particular technology is a good one. Except for one thing: different receivers respond to different stresses differently so there is no single worst-case-fits-all signal. Further, it isn't practical to perform an exhaustive test. Figure 1 includes at least five separate noise and jitter stresses. If we decided to test the receiver with all permutations of four different settings for each stress we'd end up making 4^5 = 1024 measurements.

Thus, the standards committees take a pragmatic approach. They require that receivers tolerate a few combinations of random and deterministic noise and jitter without being too picky about how the mixture is composed.

For example, the stressed eye of Serial Advanced Technology Attachment (SATA) Generation 2 (3 Gb/s) is composed of the SATA compliance pattern with

(1) Long rise/fall times – 67-136 ps.
(2) Narrow vertical eye-opening – 275 mV peak-to-peak.
(3) Sinusoidal noise of amplitude 100 mV peak-to-peak at frequencies from 2 to 200 MHz in 2 MHz steps.
(4) Deterministic Jitter (DJ) of 0.42 UI.
    This is where the pragmatism comes in. It can be any type of DJ: sinusoidal jitter or Bounded Uncorrelated Jitter (BUJ) from the PRBS generator.
(5) Random Jitter of 0.18 UI – where 0.18 UI is the "Total Jitter Defined at a Bit Error Ratio of 10^{-12}, TJ(10^{-12})," which corresponds to a standard deviation of the RJ Gaussian distribution of σ = 0.18 UI/ 14 = 13 mUI.

If the receiver operates at a BER < 10^{-12} across the sinusoidal noise frequency spectrum, then it is compliant!

There is a subtlety in determining whether or not the BER is less than 10^{-12}. At 3 Gb/s it would require 15 minutes to process the 2.6×10^{12} bits necessary to demonstrate BER < 10^{-12} at the 99% Confidence Level at each sinusoidal noise frequency. The result is a total test time of several hours. Most engineers can't afford to spend that much time so, of course, there's a faster method. The well-established dual-Dirac method extrapolates a several second measurement to determine whether or not BER < 10^{-12}.

One of the difficulties in configuring a stressed eye measurement is obtaining a truly Gaussian white noise generator for the RJ source. If you perform the full 15 minute measurement at each frequency setting, then the RJ source should faithfully follow a Guassian all the way out to 7σ. Most engineers use the faster technique which requires the RJ source to follow a Gaussian out to less than 4σ which reduces the price of the stress transmitter shown in both Figure 1 and Figure 6.

Conclusion

Stressed eye receiver testing is the natural evolution from SONET/SDH receiver tolerance tests. The new stress tests accommodate nearly every type of stress. Problems with receivers can be diagnosed by applying different stresses to isolate those signal degradations that most challenge a receiver. The brute-force method of applying stresses to signals was shown in Figure 1; it is a good graphic to help you keep track of what all the stresses are and where they come from in real systems, but it would be unpleasant to configure such a jitter-noise transmitter and borderline impossible to calibrate it. Fortunately there are high-rate programmable arbitrary waveform generators that can apply calibrated levels of the different stresses. In fact, these "fancy arb's" can capture signals from a system like Figure 1 for later retransmission.

Stressed-eye compliance tests are increasingly common in different standards. PCI Express, SATA, SAS, FibreChannel, Fully Buffered DIMM, OIF's CEI, 10GbE all employ stressed-eye compliance tests. The ideal is to test the receiver with a worst-case but compliant signal. If the receiver tolerates the signal at a low enough bit error ratio, then it passes. But it's impossible to come up with a small number of signals that are worst-case for every possible receiver so, instead, a few pragmatic stressed signals are required.