

Introduction to Precision Analysis of High-Speed Serial Systems and Components

Ransom Stephens, Ph.D.

INTRODUCTION

Advances in multi-gigabit data transfer have been dominated by serial data technologies like USB (Universal Serial Bus) and PCI Ex (Peripheral Component Interface Express), as well as technologies that have converted from parallel to serial, like SAS (Serial Attached SCSI).

Now, as we leap another order of magnitude in data rate from the 3rd generations of serial technologies like USB3 at 5 Gb/s and PCI Ex Gen 3 at 8 Gb/s to 40 and 100 Gb/s Ethernet (40 GbE and 100 GbE), parallel architectures are coming back. The most ambitious is 100 GbE's four lanes at 25 Gb/s each.

In this paper we start with a quick, high level view of the emerging multi-Gb/s architectures and then delve into the tricks that make them work and how to analyze them.

High-speed serial systems are analyzed with at least one of three goals: diagnostics, compliance, or functional test. Compliance testing is an exhaustive

checklist of performance benchmarks designed to assure the interoperability of components made by different vendors. Diagnostics, or hardware debug, involves providing well-understood conditions so that problems can be traced to their causes. Functional test, usually associated with manufacturing, employs a limited number of fast tests to determine if a product works.

HIGH-SPEED SERIAL TECHNOLOGY

Figure 1 shows the essential components of HSS (High Speed Serial) technology. The transmitter serializes a parallel data stream and transmits it through a channel that typically includes conducting cables and backplanes, and/or optic fibers. The most challenging technology is at the receiver because ones and zeros can't be distinguished at these data rates with a simple slicer. Eye diagrams of signals at several Gb/s are more often than not closed. Tricks at the transmitter, like pre-emphasis and de-emphasis, and equalization at the receiver effectively reopen the eye so that symbols can be accurately decoded.

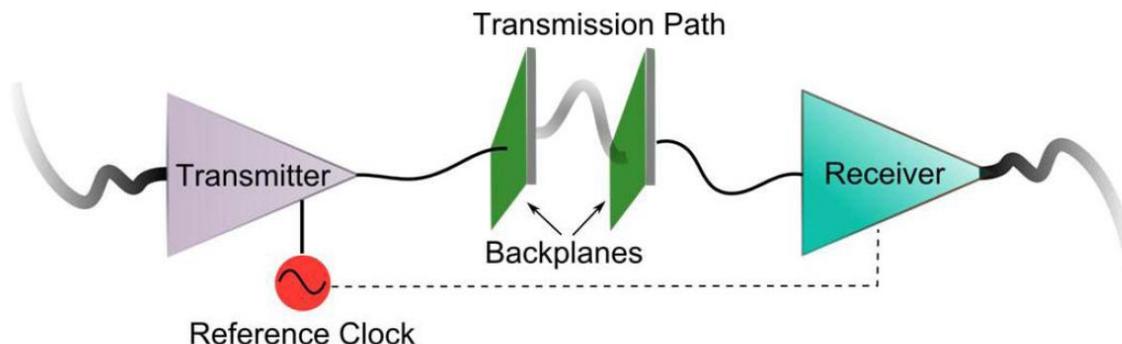


Figure 1. Principal components of a HSS (High Speed Serial) system.

THE TRANSMITTER

The transmitter, Figure 2, serializes a parallel bus into a single serial data stream. It then encodes and scrambles the data to aid clock recovery, for reasons discussed below in The Receiver subsection.

The reference clock is the ultimate timing source for logic transitions. Reference clocks are almost always oscillators at much lower rates than the data rate, typically 100 MHz. To get up to the data-rate, the reference clock is multiplied by a Phase Locked Loop (PLL). A pitfall is that the PLL amplifies the oscillator phase noise by the square of the multiplication factor. The resulting phase noise appears as random jitter on the signal.

One of the enabling features of HSS technology is differential signaling. In an ideal system, two complementary signals are situated arbitrarily close together. The electromagnetic radiation produced by one line interferes destructively with the other, canceling out crosstalk and electromagnetic interference. Of course, in real systems the two signals are separated and can never be exactly the same length; still, the result is a much cleaner system than the single-ended equivalent.

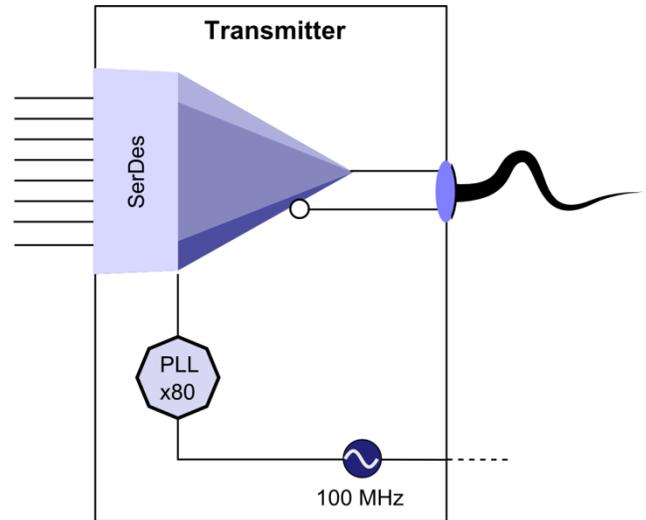


Figure 2. Typical transmitter architecture.

TRANSMITTER TEST

Figure 3 shows a typical transmitter test configuration; this example is from the USB3 specification [1]. Compliance test specifications set maximum allowed levels of Random Jitter (RJ), Deterministic Jitter (DJ) and Total Jitter (TJ) defined at a Bit Error Ratio (BER). The BER requirement is usually $BER < 10^{-12}$. Values of RJ, DJ, and TJ are all calculated *after* receiver equalization.

Thus, the test tool, usually a high bandwidth oscilloscope, must be capable of emulating equalization and clock recovery schemes. Tests are performed with compliance test patterns designed to replicate the most challenging symbol sequences a transmitter can face. In compliance tests, the transmitter operates with all the features turned on, including de-emphasis and Spread Spectrum Clocking (SSC).

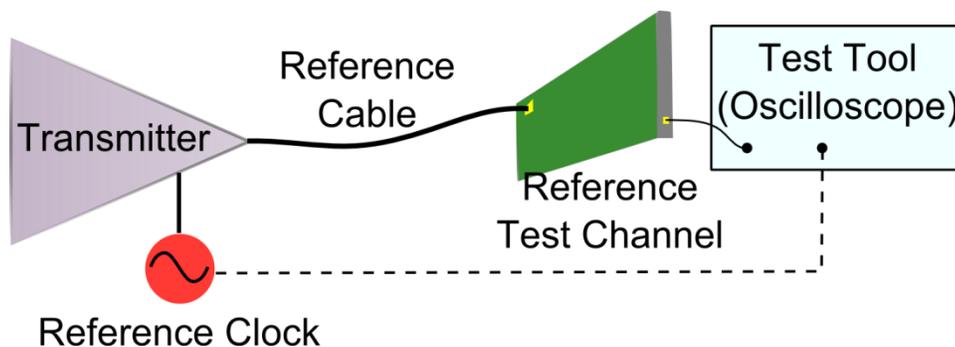


Figure 3. Typical transmitter test setup.

SPREAD SPECTRUM CLOCKING (SSC)

SSC is a technique for reducing electromagnetic interference by smearing the generated power in the frequency domain. The standard technique is to apply triangular wave modulation at 30 kHz with an amplitude of 4000 ppm to the signal.

TRANSMITTER DE-EMPHASIS

Transmitter de-emphasis is a trick used to compensate the distortions introduced by the channel.

Since distortions are worse at high frequencies, and since the highest frequencies occur at logic transitions, transmitter de-emphasis increases the voltage swing of the symbols that immediately follow transitions, as shown in Figure 4. The bits before and after the transition have smaller voltage swings and are said to be *de-emphasized*. By de-emphasizing the lower frequency components of the signal, the channel response is partially compensated. Figure 4c shows the eye diagram of a standard signal, Figure 4a, after propagation through a channel, and Figure 4d shows the result for a de-emphasized signal.

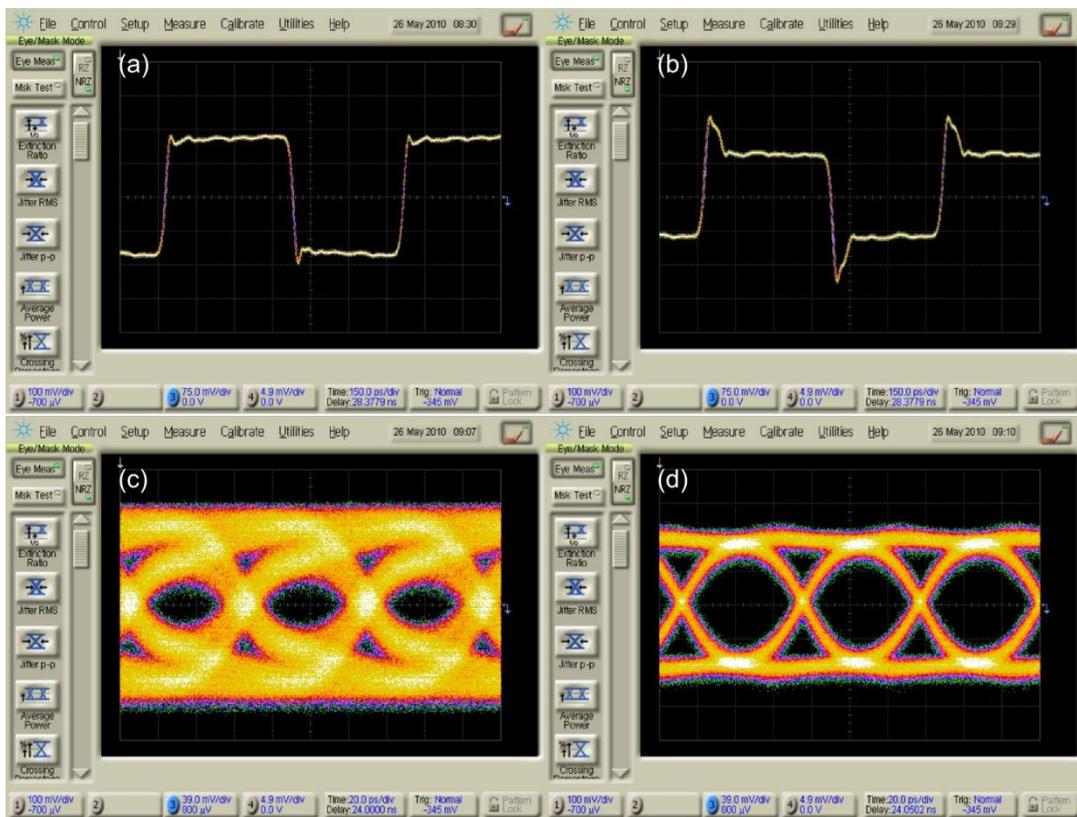


Figure 4. A 16 Gb/s signal (a) without de-emphasis and (b) with de-emphasis (produced by the [PatternPro Serial Data Generator](#) with a [programmable pattern amplifier](#)) and the resulting eye diagrams, (c) and (d) respectively, after propagating through a backplane.

Transmitter diagnostic testing should start as close to the transmitter as possible and should begin with the simplest possible data patterns. A clock-like signal, 101010..., has no symbol-sequence dependent effects like Inter-Symbol Interference (ISI), which make it a good signal for investigating Duty Cycle Distortion (DCD) and RJ. If the transmitter looks okay, then advance to increasingly complicated patterns to see if the transmitter introduces ISI.

You can see ISI by averaging the waveform on the oscilloscope. Averaging removes RJ so that ISI should appear as bold lines in the transition regions of the eye diagram. These lines are called transition trajectories. Only after the transmitter has earned your confidence right out of the spigot should you introduce cables and then backplanes. With each addition, watch how the eye diagram closes.

THE CHANNEL

The transmission path is a combination of cables and backplanes and, for 40 GbE and 100 GbE, optic fibers. The first thing that digital engineers must do when working at Gb/s data rates is to embrace the analog nature of real signals. Circuit boards and backplanes are complicated waveguides. Rather than thinking of the signal as a succession of logic highs and lows, think of it as an electromagnetic field propagating wildly through a dielectric barely hanging onto the conducting trace.

CHANNEL TEST

To limit the distortion introduced, compliant cables and backplanes have to meet S-parameter requirements. S-parameters can be measured on either a Vector Network Analyzer (VNA) or an ultra-wide bandwidth oscilloscope.

A VNA measures the response of passive devices to carefully calibrated standing waves. Frequencies are stepped from very low to quite high. The differential scattering matrix is calculated from the response of all frequencies.

The oscilloscope technique, Figure 5, is faster, easier, and less expensive, but it is also less accurate. Where a VNA makes the measurement in the frequency domain, an oscilloscope can make it in the time domain. The transmitted and reflected impulse response is the time-domain equivalent of the S-parameters. Since the first derivative of a step function is an impulse, the response of the channel to any known pattern includes sufficient information to calculate the impulse response and, from that, derive S-parameters. It works like this: An instrument-quality signal is transmitted through the channel. The oscilloscope identifies the pattern and then calculates the S-parameters by resolving the impulse response.

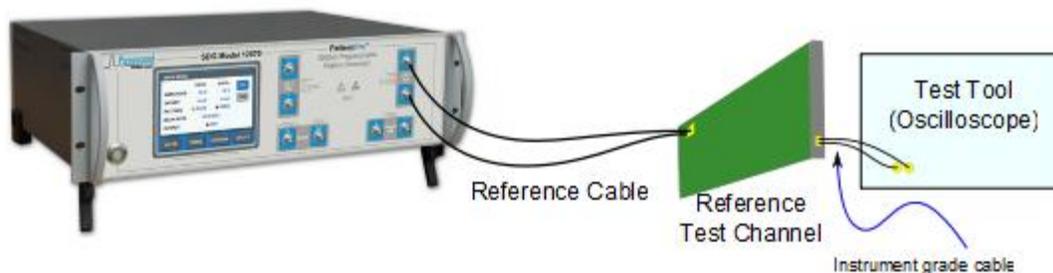


Figure 5. Measuring S-parameters with a [PatternPro Serial Data Generator](#) and an oscilloscope.

Since the oscilloscope technique loses accuracy at bandwidths above about 20 GHz, it should be used only with caution for data rates above 8 Gb/s or so (NB: this benchmark is advancing all the time so check with your scope vendor). At extreme data rates, like the 4x25 Gb/s of 100 GbE or the single 40 Gb/s of 40 GbE that must survive one meter of backplane, the VNA technique should be used.

THE RECEIVER

Receiver technology is the heart of serial data technology's success, Figure 6. The trick to identifying bits at over a few Gb/s is to provide the symbol decoder a clock recovered from the data. If a local clock is used, jitter increases the bit error ratio above the specified maximum.

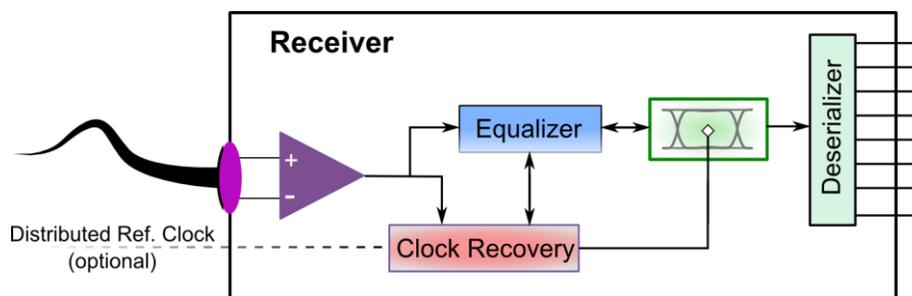


Figure 6. Typical receiver architecture.

There are two primary architectures. First, pure clock-forwarding (also called embedded clocking) is used in USB3 and 100 GbE. The receiver recovers a clock purely from the incoming signal. PLLs and Phase Interpolators (PIs) synchronize a local oscillator to data logic transitions. This technique mitigates the effect of any jitter below the PLL or PI bandwidth. Here's how it works: The clock determines the sampling point used by the symbol decoder to distinguish logic levels. By reconstructing the clock from symbol transitions much of the jitter on the data is also on the clock. The result is that the sampling point jitters the same way as the data, dancing in rhythm, so that the jitter they share doesn't cause errors.

There is a complication. To reconstruct a clock, logic transitions must occur fairly often. That is, the signal must have sufficient *transition density*. If the transition density is too low, the clock recovery circuit can't lock. USB2 and 3, sATA, SAS, PCI Ex Gen 2, DisplayPort, etc. all use 8B/10B data encoding to assure runs of no more than five identical bits. The drawback of 8B/10B coding is that two of every ten bits are used by the coding scheme, a 20% bandwidth overhead. The overhead is reduced in 40 GbE and 100 GbE by combining 64B/66B coding with both data scrambling and Forward Error Correction (FEC).

To reduce the encoding overhead without substantially increasing system cost, PCI Ex Gen 3 uses a combination of distributed and embedded clocks. By connecting the dotted line in Figure 1 the same oscillator is used at the receiver as at the transmitter.

Introducing identical phase noise at both ends of the system provides the freedom to use more liberal clock recovery schemes.

In particular, Delay Locked Loops (DLLs) compare the timing of logic transitions to a set of quantized phases of a fixed reference. At each transition, the DLL determines whether the clock phase is ahead or behind the symbol transition. The early and late intervals are accumulated and, after a predetermined time, effectively the DLL bandwidth, the recovered clock phase is either advanced or reduced. Long runs of identical symbols don't affect the clock phase. Once the recovered clock is aligned with the data, the distributed reference clock effectively fixes the timing. With this immunity from the problem of long runs of identical bits, PCI Ex Gen 3 can use a combination of 128B/130B encoding and scrambling. The bandwidth overhead is reduced from the 20% of PCI Ex Gen 2 to 1.5 % for Gen 3.

CLOSED EYE ANALYSIS

The intimidating aspect of HSS technology is that the received signal is usually a closed eye, Figure 7. The eye is closed by signal distortion caused by the frequency response of the channel. The frequency response depends on everything from impedance matching at connectors to the many paths that electromagnetic fields can take through these highly nonlinear waveguides. The gross effect is that cables and backplanes behave like complex low pass filters.

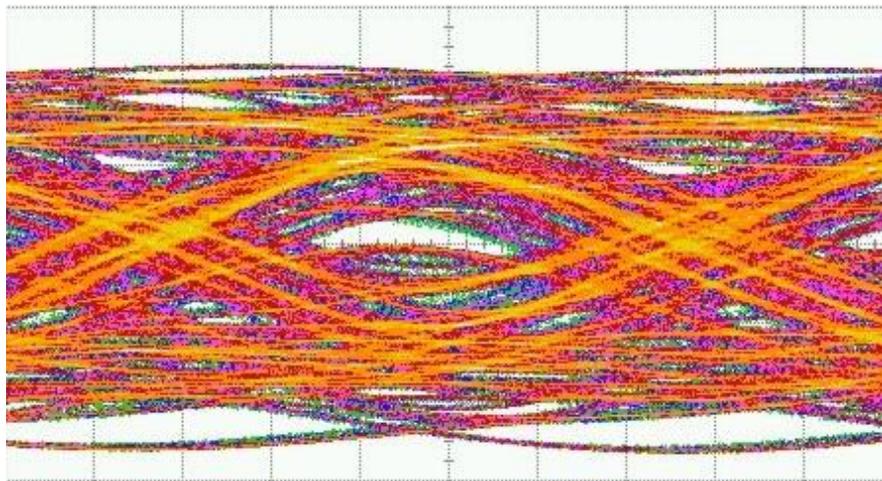


Figure 7. 20 Gb/s signal after propagating through many feet of backplane.

Consider Figure 8. There are two sources of frequency content. First, Figure 8a, the Fourier components required to make a square wave come in half-integer multiples of the data rate. At multigigabit rates, transmitters rarely include more than the third harmonic, $5/2 f_d$.

Second, in Figure 8b, lower frequency components are introduced by data sequences at half-integer fractions of the data rate — fractions limited by the longest run of consecutive identical bits allowed by the specific data encoding/scrambling requirement.

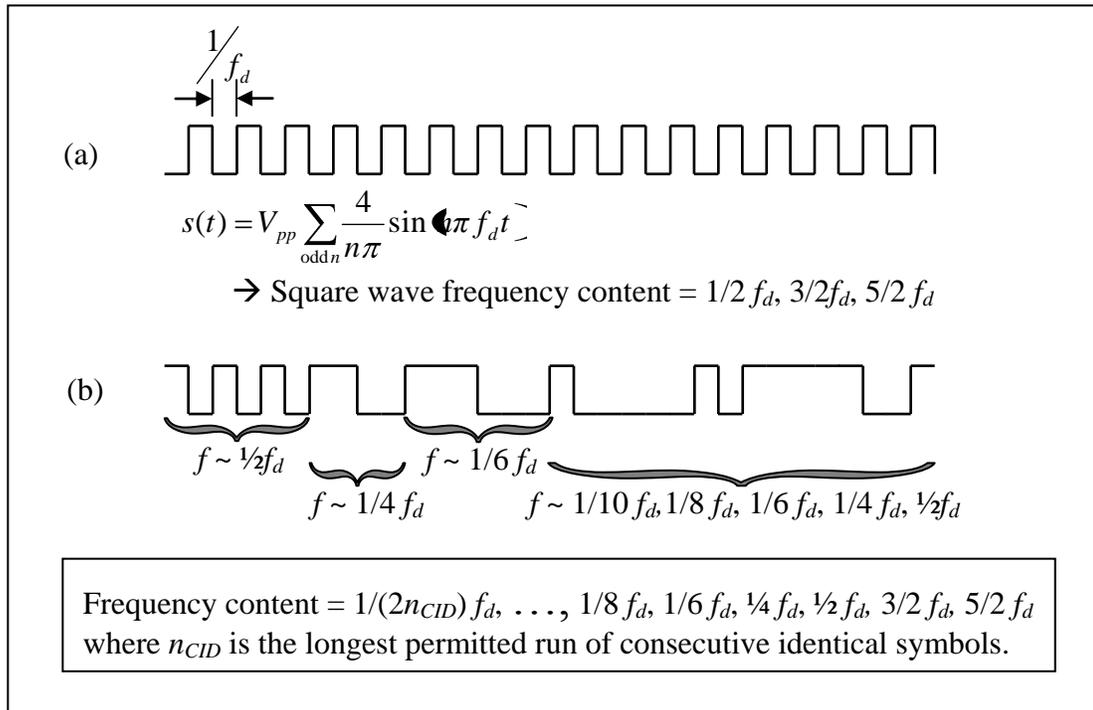


Figure 8. Frequency content of a digital signal:
(a) Fourier components of the fundamental square wave and
(b) frequencies introduced by the data.

Since the response of the channel is essentially static, and since the frequency content of the signal is reasonably well known, the channel response can be corrected or *equalized*. This is the essence of equalization.

The idea can be extended to include manipulation of the voltages of any number of bits, or *cursors*, surrounding the bit being identified. In current HSS technologies, no more than three tap emphasis is employed.

**EQUALIZATION AT THE TRANSMITTER:
DE-EMPHASIS AND PRE-EMPHASIS**

De-emphasis is a type of “one-tap” equalization. A *tap* is the amount that the voltage of a given bit is adjusted. De-emphasis is one-tap equalization where the voltage of the bit is adjusted only at a transition. If the voltage of the bit that precedes the transition is modified, it’s called “two-tap” transmitter equalization.

EQUALIZATION AT THE RECEIVER

The goal of equalization is to remove channel distortion from the received signal so that the symbol decoder can operate at a sufficiently low Bit Error Ratio, usually BER < 10⁻¹². There are three categories of equalization: linear, nonlinear and adaptive.

The simplest equalizer is a filter that inverts the gross effects of the channel by suppressing low frequencies, amplifying high frequencies and blocking noise at very high frequencies. These are called Continuous Time

Linear Equalizers (CTLE) and are designed by assigning poles and zeros and their frequencies to develop an appropriate transfer function, like the one in Figure 9.

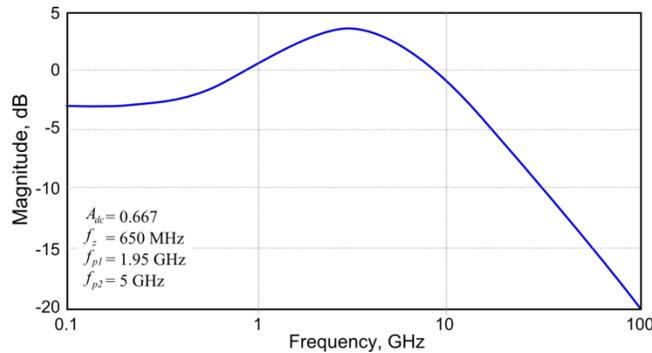


Figure 9. The CTLE transfer function suppresses low frequencies, amplifies the first few data-rate harmonics, and blocks high frequency noise.

To move up in complexity, think of the received signal as a long waveform. Let the voltage at the center of the bit we're concerned with, the n^{th} bit, be $c(n)$, and let's call the voltages of neighboring bits *ursors*. The first *pre-cursor*, $c(n-1)$, is the voltage of the bit immediately preceding the bit of interest; similarly, the voltage of the first *post-cursor* is $c(n+1)$.

The pulse response in Figure 10 shows how an isolated bit is smeared by a channel. The fact that bits

smear into adjacent bits provides a technique for correcting them. A discrete, linear equalizer, like a Feed Forward Equalizer (FFE), uses the voltages of surrounding bits to correct the voltage of the one being identified. The corrections are called *taps*. For example, a "5-tap Feed Forward Equalizer" is given by:

$$eq(n) = t_{n-4}c(n-4) + t_{n-3}c(n-3) + t_{n-2}c(n-2) + t_{n-1}c(n-1) + t_n c(n)$$

where t_n is the value of the tap multiplying the voltage of cursor n , $c(n)$.

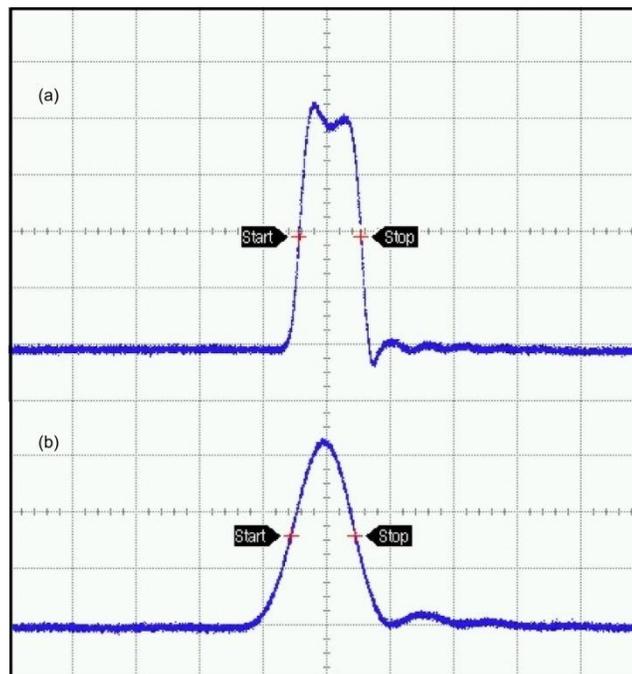


Figure 10. (a) a single logic one among a long string of zeros and (b) that pulse after propagating through a cable.

The most common equalizer is the nonlinear Decision Feedback Equalizer (DFE), Figure 11. In a DFE, another layer of taps is imposed over an FFE. The output of the FFE goes to the symbol decoder. The decoded symbol is delayed and then looped back, and a separate set of taps are applied to this digital signal. The output of the loop is combined with the output of the linear equalizer to give the decision feedback equalized signal. By applying taps to the decoded digital symbols, a DFE is manifestly nonlinear.

There are two primary reasons why adding this nonlinear layer improves equalizer performance. First, the linear equalizer can only correct distortion that extends over the cursers it taps; the nonlinear layer extends the reach. Second, and more importantly, the role of the linear equalizer is different when it's a

component of a DFE. Alone, the optimal FFE taps for opening the eye also amplify high frequency noise and jitter. By including the decision feedback loop, the linear equalizer can be optimized for minimum noise gain while still reducing ISI. The feedback loop then finishes the job, providing a more optimal solution. The only major drawback is that, since decision feedback is predicated on the assumption that the decisions are accurate, a few consecutive mistakes can lead to an avalanche of errors.

Adaptive equalizers modify the tap values for different conditions. Most adaptive equalizers are proprietary, all of them benefit from having a known training sequence included in the protocol, as is done in the USB3 protocol.

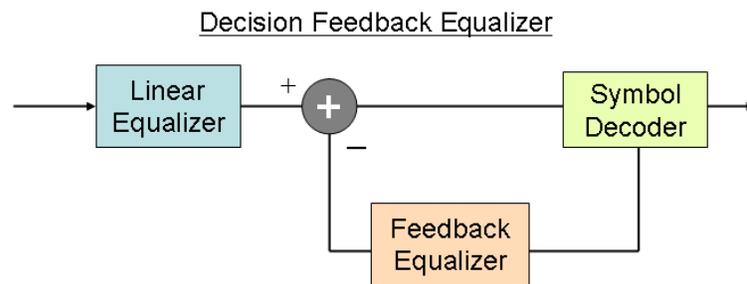


Figure 11. Decision Feedback Equalizer (DFE) block diagram.

**RECEIVER TEST:
STRESSED EYE TOLERANCE TESTING**

The idea of stressed eye tolerance testing, Figure 12, is to subject the receiver to the worst-case compliant signal.

If the receiver can operate at or below the maximum specified BER (again, usually $BER < 10^{-12}$), then it should operate with the signal of any compliant transmitters and channels. To this end, a stress signal is composed of a challenging pattern with the maximum allowed RJ and ISI, plus sinusoidal jitter.

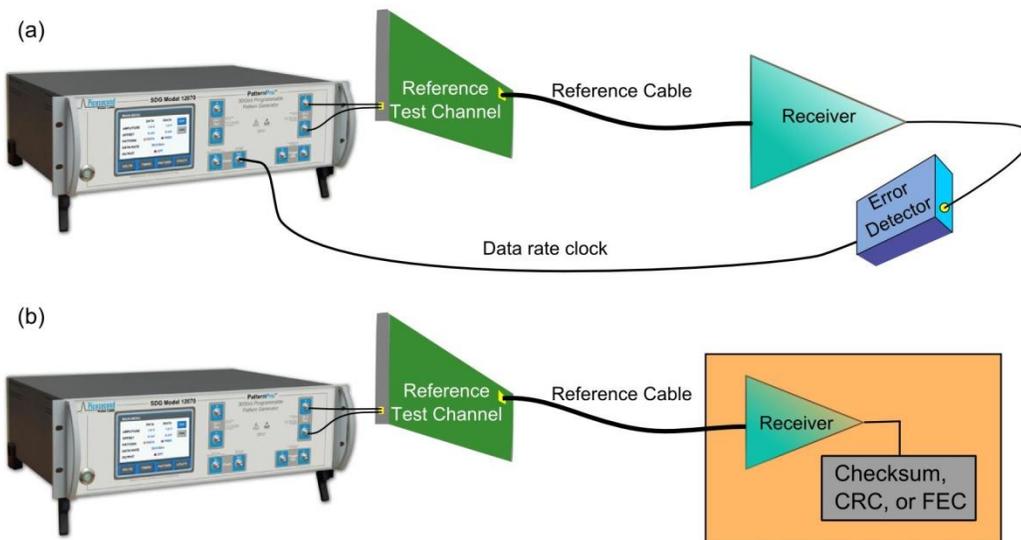


Figure 12. Stressed eye tolerance test setup (a) with an error detector and (b) without an error detector.

Compliance test patterns, such as CJTPAT [2], are designed to challenge the internal components of the receiver. Sequences of low and high fractions of logic 1s, that is, low and high *mark densities*, challenge receiver stability. If the RC time constant of AC-coupled receivers isn't large enough, varying the mark density causes the baseline to drift.

To challenge the clock recovery circuit's ability to synchronize and lock, compliance patterns include long runs of consecutive identical bits, that is, long runs of low transition density. The mark density and transition density over the entire scrambled and encoded pattern must both be ½, but can vary over pattern segments.

The signal then has other stressors applied to it. The specifications differ on which types of stress must be applied. The vast majority require random jitter and sinusoidal jitter. The sinusoidal jitter is applied according to a template, Figure 13, in order to probe the clock recovery bandwidth. Each frequency is tested separately, though typically, measurements just below and above the roll-off are sufficient. Some standards require that SSC also be applied. Applying SJ across the template together with SSC is redundant because SSC is very low amplitude periodic jitter, and the applied SJ is very high amplitude periodic jitter.

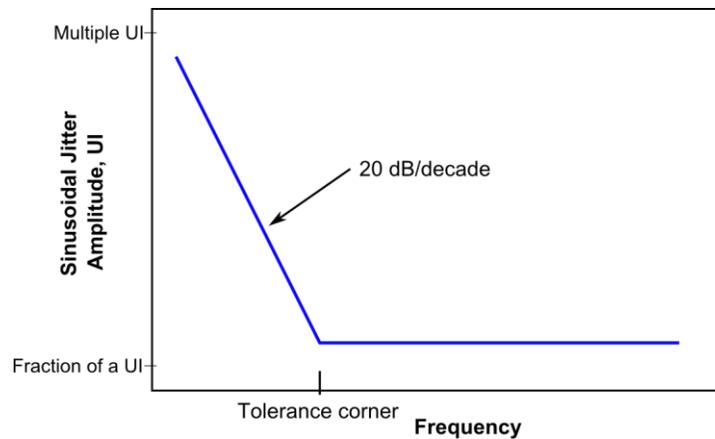


Figure 13. Typical sinusoidal jitter template for HSS compliance tests.

Most HSS technologies above 5 Gb/s, including USB3, 100 GbE, and PCI Ex Gen 3, proscribe reference channels with the worst-case but compliant response. To encourage the reference channel to impose maximum ISI, the compliance pattern includes many symbol sequence permutations. The idea is that each sequence of 8 or so symbols produces a unique logic transition trajectory.

Figure 12 begins with a Serial Data Generator. The [SDG 12070](#) applies the requisite RJ and SJ to the compliance pattern. The stress signal is transmitted through the reference channel to the receiver being tested. In Figure 12a, the performance of the receiver is measured by an error detector (which requires a data-rate clock signal). The error detector compares the receiver output to the transmitted pattern and calculates the BER (Bit Error Ratio). If no errors are observed over 3×10^{12} bits, then $BER < 10^{-12}$ is assured at the 95% confidence level — which is sufficient.

While there is no substitute for a [high-quality programmable pattern generator](#), the added expense of an [error detector is usually not necessary](#). If the

protocol includes calculation of a checksum, Cyclic Redundancy Check (CRC), frame error detection, or if it incorporates Forward Error Checking (FEC), the errors are reported by the receiver itself, as in Figure 12b.

RECEIVER DIAGNOSTIC TESTING

If a design fails the receiver tolerance test, use the fact that each imposed stress, including elements of the pattern, is designed to stress different receiver components.

First, isolate the problem by removing stressors until the BER is okay. The removed stressors should bound the problem. Next, start with a simple pattern with no stress and add stressors individually. For example, a simple clock signal shouldn't cause the receiver baseline to drift and is the easiest signal for clock recovery. Insert longer consecutive identical bits in the pattern, vary the mark and transition densities, introduce DCD (Duty Cycle Distortion) by varying the crossing point voltage and apply sinusoidal jitter of

different frequencies to increasingly complicated patterns to determine if it is a clock recovery problem.

If the system passes the compliance test with all stresses included except for the ISI-introducing effect of the reference channel, check the connectors. Any impedance mismatch can aggravate the ISI of the system beyond the compliant stress.

If you conclude that the equalizer is the problem, and all you have at the channel output is a closed eye, then implement your equalizer design on the scope. Most scopes include options for equalization emulators and many accommodate MATLAB scripts. If the emulated equalizer opens the eye, then the problem is in the implementation of the equalizer on the receiver. If it doesn't, then the equalizer design is inadequate.

The mixed serial/parallel nature of 40 GbE and 100 GbE, whether 4x10 or 10x10 Gb/s lanes or the even more challenging 4x25 Gb/s architecture, introduces the potential for skew and crosstalk. Skew is the variation in propagation time between different channels. The 40 and 100 GbE specification implements a Physical Coding Sublayer (PCS) that permits a generous 180 ns of skew.

Crosstalk is a different story. Spikes of electromagnetic noise occur at data transitions where electric fields vary drastically. The higher the data rate, the faster the rise/fall times, the sharper the spike. Differential signaling cancels a lot of the radiation, but it's not perfect. With 10 adjacent 10 Gb/s channels, crosstalk is likely to be a problem.

FUNCTIONAL TEST

Functional testing needs to be fast, efficient and effective. The idea is not to shake down every feature of every part, but to zero in on a few tests that are problem indicators. Figure 14 shows a receiver functional test setup. Multiple channels are driven by a single programmable [PatternPro](#) generator, in this case the [SDG 12070](#), accompanied by a [PSPL 8020B](#). The PSPL 8020B is a multiple output pattern amplifier that splits the signal into as many as 4 differential outputs. Each output is driven by a separate programmable limiting amplifier. The result is a source that can drive four separate test stations.

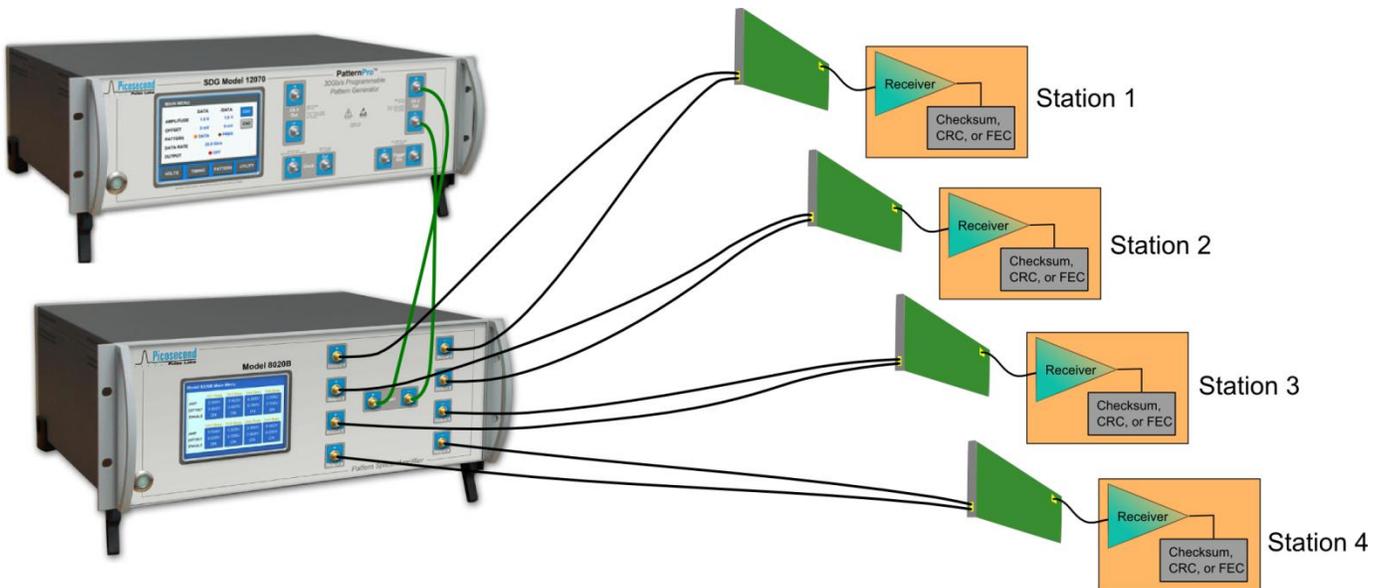


Figure 14. Multiple-station functional receiver and channel test setup featuring [PatternPro Serial Data Generators](#).

YOUR SOURCE FOR HIGH-SPEED SERIAL ANALYSIS

High-speed serial data technology combines the natural noise-canceling properties of differential signaling and the jitter tolerant features of embedded (or forward) clocking with channel-response correcting de-emphasis at the transmitter and/or equalization at the receiver. The results are robust, multi-gigabit, data transfer systems ranging from a few Gb/s to over 10 Gb/s, and even with closed eyes, BERs can be held below 10^{-12} .

With the huge bandwidth demands of video-on-demand, cloud computing and data storage/retrieval, 40 GbE and 100 GbE are coming on strong. To get to 100 Gb/s, GbE reintroduces parallel architecture while retaining the strengths of serial signaling.

Development and manufacture of HSS components and systems requires cutting edge test equipment. Picosecond Pulse Lab is dedicated to helping you assemble just the equipment that you need in a way that is easily scalable to meet your future needs, but and with no excess features or cost. Our **PatternPro** serial data instruments separate BERT functionality such as pattern generation and error checking so that you can buy what you need as you need it.

The **PatternPro** line of products covers the latest serial data test needs for all generations of the emerging high rate standards like PCI-Express, USB, SAS/SATA, et cetera as well as 100G Ethernet including complicated optical modulation like DPQPSK (dual-polarization quadrature phase shift keying).

Our **Serial Data Generator (SDG)** products are the cornerstone of the **PatternPro** product line, key features include:

- Maximum data rates range from 1.6 Gb/s to 30 Gb/s – generously covering PCI-Ex Gen 1 to 100 GbE
- Programmable data rate, amplitude, offset and crossing point – necessary for diagnostic and compliance testing
- Programmable data patterns – with PRBS and common test patterns pre-programmed
- Multi-channel generators with independent signal generation – key for crosstalk analysis
- Data rate clock – crucial for use of BERT error detectors and some high bandwidth oscilloscopes
- Option for built-in [jitter insertion](#) RJ, SJ, and external delay modulation – necessary for stressed eye testing

- Multi-channel pattern amplifiers – to scale-up your pattern generator investment
- Up to 3-tap programmable [pre/de-emphasis](#)
- Touch screen GUI and USB remote control – to ease automated functional testing

The **PatternPro** SDG products provide multi-channel, instrument-grade sources that generate crisp, clean, test signals with all the features necessary for compliance, diagnostic, and system testing right now.

At [Picosecond Pulse Labs](#) we're building on our legacy of precision, NIST-quality components to become your go-to source for high-speed serial data analysis.

REFERENCES

- [1] http://www.usb.org/developers/docs/usb_30_spec_060910.zip
- [2] See the FibreChannel web site, www.t11.org