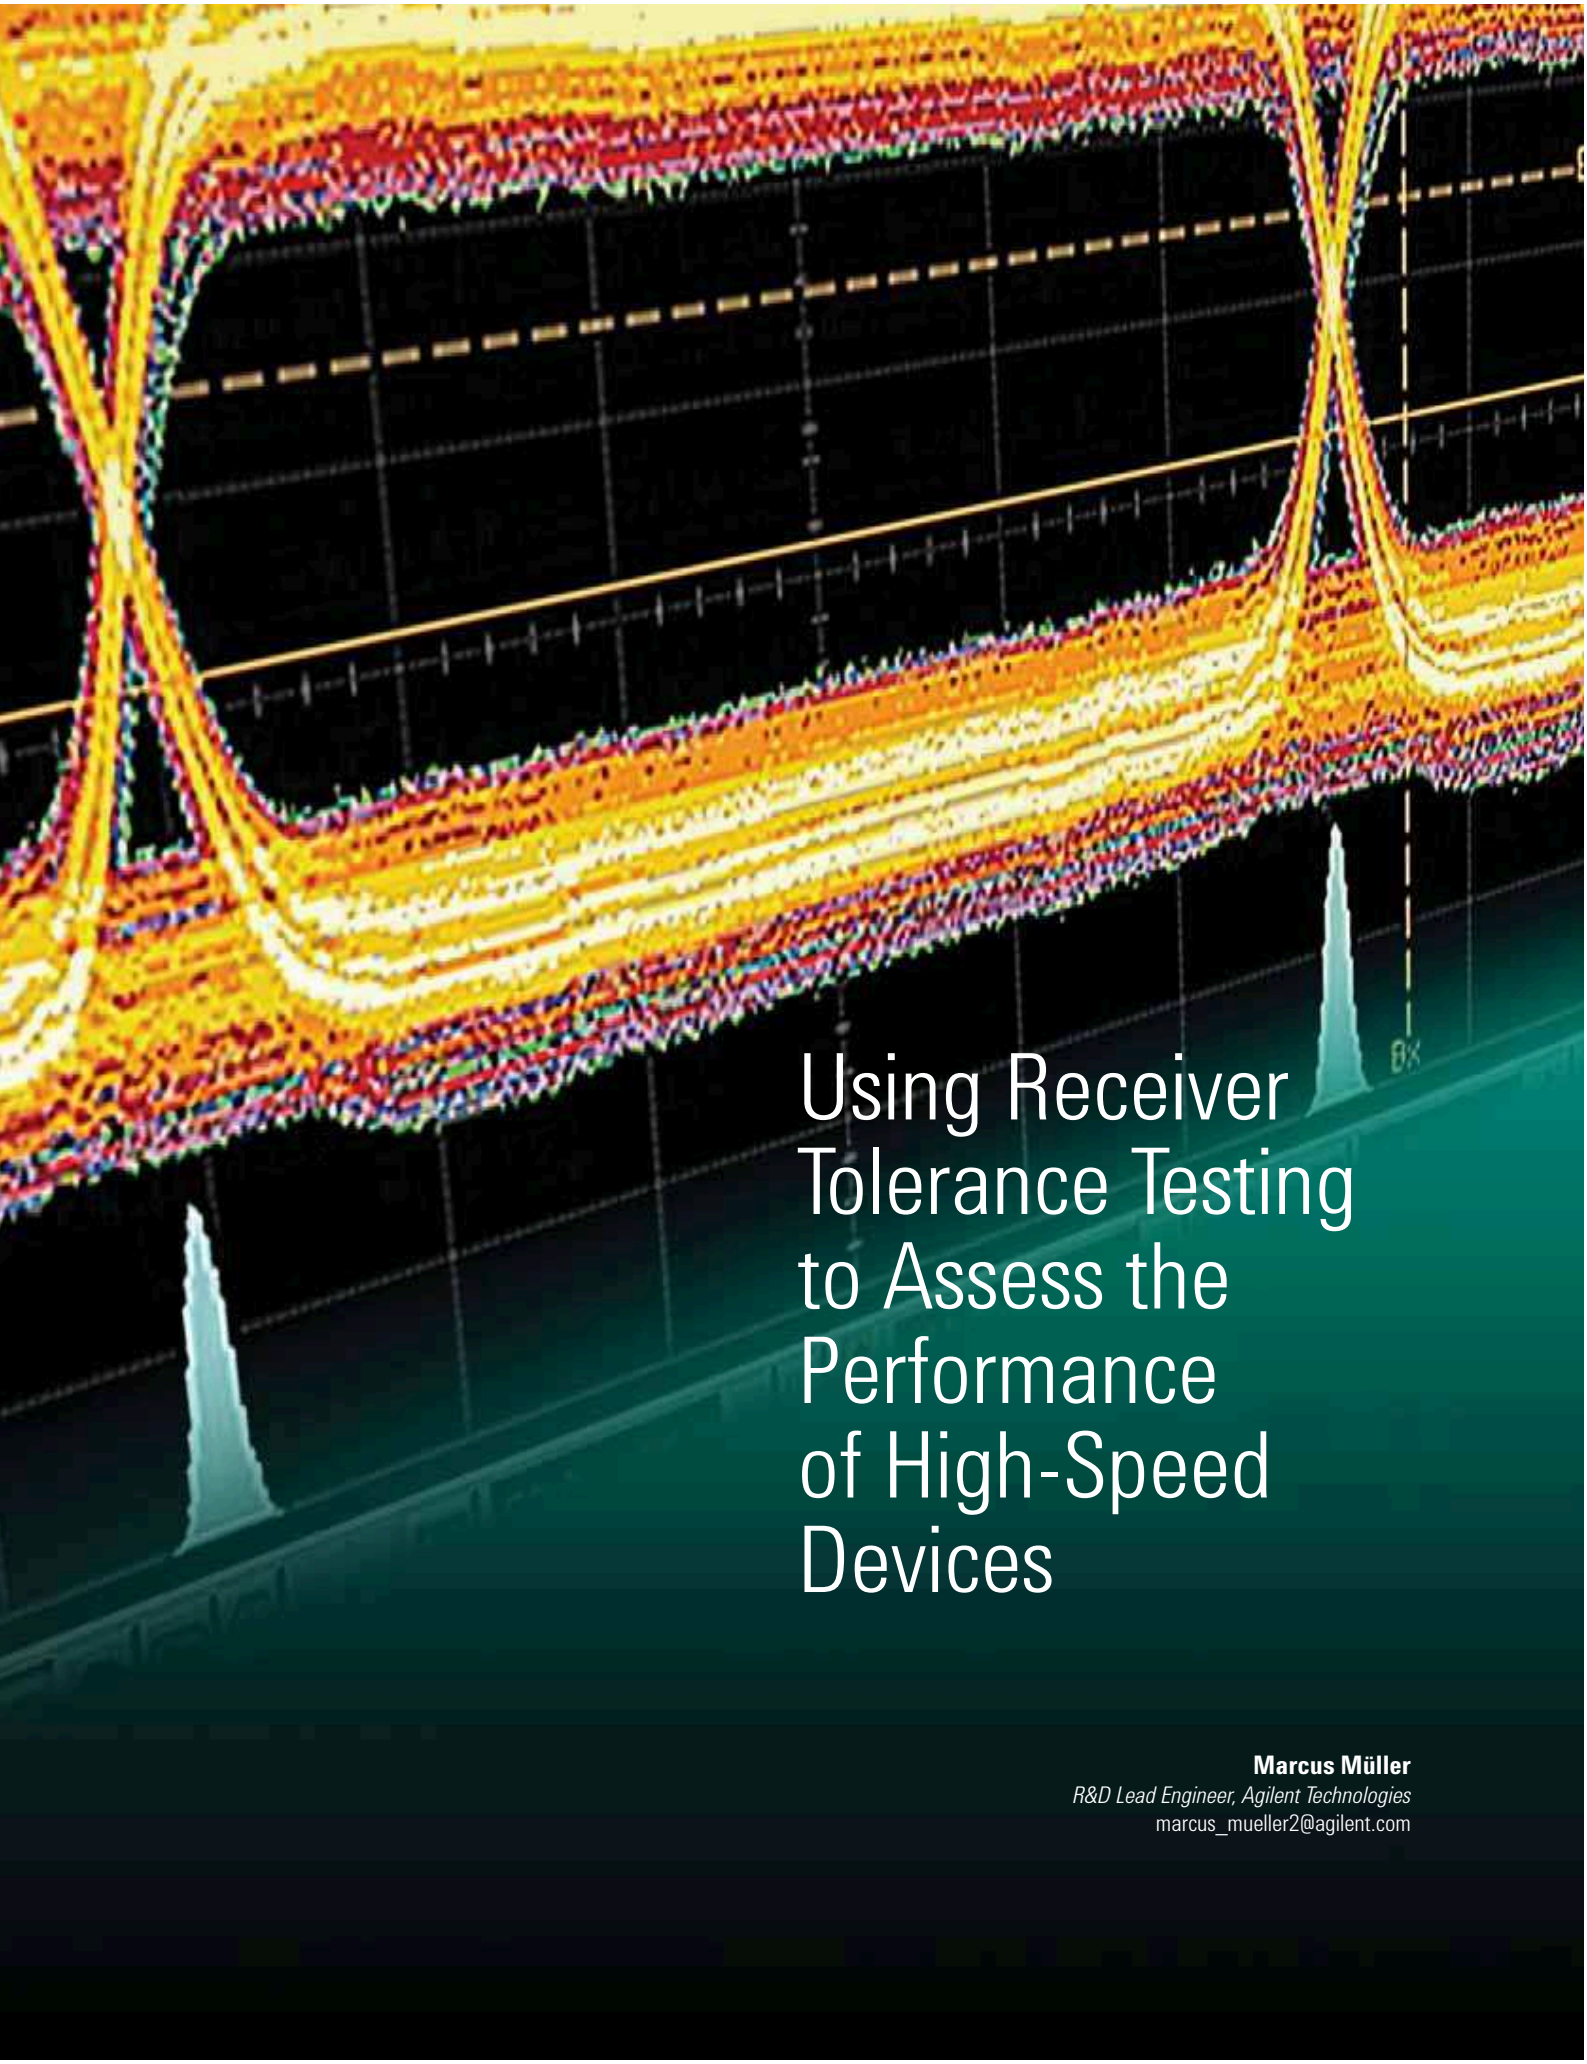

Using Receiver Tolerance Testing to Assess the Performance of High-Speed Devices

Marcus Müller
R&D Lead Engineer, Agilent Technologies
marcus_mueller2@agilent.com

The landscape of data communications and computing has undergone dramatic changes in the past five years. What used to be parallel-clocked or source-synchronous buses have been replaced by high-speed serial architectures with embedded clocks. Today, a typical PC contains a PCI Express graphics port (serial at 2.5 Gb/s, soon to be 5 Gb/s) instead of a parallel Advanced Graphics Port (AGP) and Serial ATA drives (at 1.5 or 3 Gb/s) instead of parallel IDE or SCSI — either of which were recently state-of-the-art technologies. The dominance of high-speed serial links, especially in the cost-sensitive PC market, was made possible by incredible advances in transmitter and (especially) receiver technologies: clock-data recovery, clever coding schemes, receiver equalization and transmitter pre-emphasis have been major factors.

As receivers become more complex, thorough testing becomes more difficult. Until recently, typical receiver-tolerance test systems included numerous pieces of equipment. These systems were expensive, bulky, difficult to set up, tough to control and extremely hard to calibrate. Today, the necessary functionality is available in standalone bench instruments such as the Agilent J-BERT N4903A high-performance serial bit error ratio tester (BERT). This instrument can perform full receiver characterization over a wide range of parameters, in addition to other critical measurements such as bit error ratio (BER) and jitter.

Basics of receiver tolerance testing

The basic point of receiver tolerance testing is to assess how well the device under test (DUT) can handle bad signals. There are two important aspects to this: the definition of a "bad" signal and the parameters that define expected operation of the DUT.

A receiver tolerance setup tests multiple features of the receiver: clock recovery from the data stream, jitter tolerance, noise immunity and the ability to compensate for the effects of band-limited transmission channels. Receiver tolerance testing is often termed "jitter tolerance" testing because measurements are jitter related; however, noise immunity in particular is typically included. Additional aspects of receiver testing include operation with low signal levels, at signal offsets and with stress-data patterns. The actual worst-case bad signal depends on the application, which is usually defined in the relevant standard.

It is often difficult to determine if the device works with a given signal. In transmitter measurements, we attach a test instrument such as an oscilloscope to its output and directly measure the parameter of interest. In contrast, receiver testing relies on indirect observations. The only method that comes close to a direct measurement of receiver performance is a BER measurement: The data stream at the output of the receiver is compared to expected data and the BER is calculated by dividing the total number of compared bits into the number of error bits. This depends on access to the device output, which cannot be taken as a given.

Test setup and device types

On a high level, the general setup for receiver-tolerance testing is fairly straightforward (Figure 1). A pattern generator creates a digital data signal with a certain amount of jitter, noise or other impairments and sends it to the DUT. However, determining the device's ability to tolerate the input signal requires access to the data after it is received by the DUT.

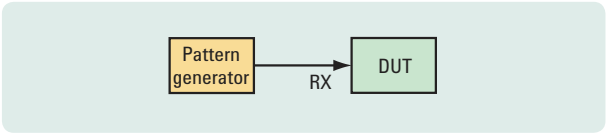

Figure 1. High-level setup for a receiver tolerance test

Typical high-speed serial integrated circuits are serializing/de-serializing transceivers, integrating both a transmitter and receiver in a single device. In this case, it is possible to use the device's transmit output to resend the received data and then measure the BER with an external error detector. In most cases, the loopback is done within the chip; however, the received data can also be looped externally on a test adapter or test board.

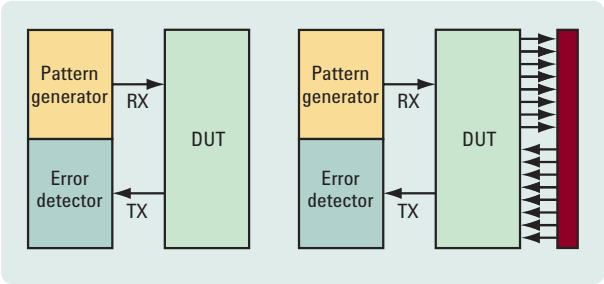

Figure 2. Internal (left) and external (right) loop-back setups for receiver tolerance test of serial high speed transceivers

Measurement procedure

Whether the receiver works or not is seldom a yes/no decision: A receiver may be able to handle the bad test signal most of the time, but may occasionally cause bit errors. If the device output is available, the BER can be used to establish a good/bad threshold. Most standards specify a limit for the acceptable error ratio, often at 10^{-12} but sometimes even lower. Unfortunately, proving that the device operates at 10^{-12} or lower requires at least $3x10^{12}$ error-free bit compares. Even at high data rates, this takes a long time (e.g., more than five minutes at 10 Gb/s). If the error ratio is high, the measurement can be terminated after detecting just a few errors, reducing the time spent on a bad device.

Verification of receiver compliance with a standard begins with knowledge of the required test signal components from the specification. Compliance testing requires application of the specified signal impairments at the receiver input, measurement of BER and a comparison of the results with the specified BER limit. In many cases, though, a single test is not enough and it is necessary to test different combinations of jitter and noise components, sometimes at different signal levels. In fact, almost every standard specifies the amount of sinusoidal jitter the device must tolerate as a function of the jitter frequency; the lower the frequency, the more jitter must be tracked. This procedure tests the clock data recovery circuit of the receiver. Figure 3 shows an example: The solid black line is the standard mask, the dots show each measurement point and the color of the dots indicates the BER at each point (green if the BER is lower than 10^{-12} and red otherwise). Since all points are green, the example device passed the test.

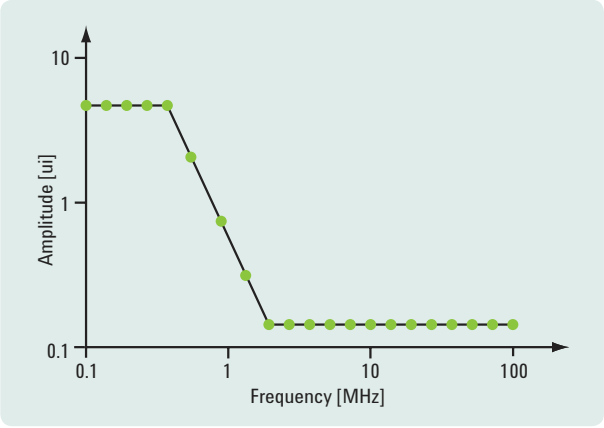

Figure 3. Compliance test with a sinusoidal jitter tolerance mask. Green dots show measured points that passed a BER test at 10^{-12}.

While a compliance test results in a pass/fail decision and is therefore used mostly in production environments, it is also useful to know how well the device works. Such a characterization typically sweeps a parameter such as sinusoidal jitter frequency as in the preceding example. However, instead of just checking the receiver's ability to tolerate a given level of sinusoidal jitter at the current frequency, the goal is to quantify exactly how much jitter it can tolerate. This is done by increasing the jitter level step by step until the device fails, as shown in Figure 4. Receiver characterization is not limited to sinusoidal jitter: BER can also be measured as a function of inter-symbol interference, random jitter or any other variable the test setup can control.

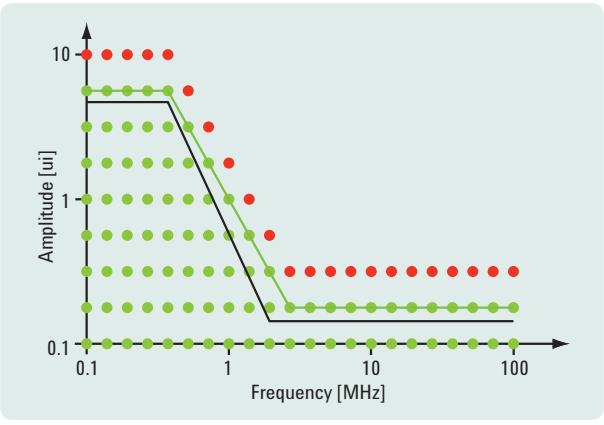

Figure 4. Characterization measurement with sinusoidal jitter. Green dots show measured points that passed; red dots measured points that failed a BER test at 10^{-12}.

Test signal generation

So far, our discussion has assumed the ability to create test signals with various kinds of jitter and other impairments. It is now time to examine how these signals are generated and thereby gain a better understanding of the test instrumentation.

Random and periodic jitter generation

The most important part of the signal generation for receiver tolerance testing is non-correlated jitter: random jitter and deterministic period jitter. The two main methods for jitter generation are phase/frequency and delay modulation. For historical reasons dating back the first SONET standard, typical tests often use periodic jitter with a sinusoidal waveform; such signals seldom occur with real devices, but they allow for easy testing of CDR/PLL performance versus jitter frequency.

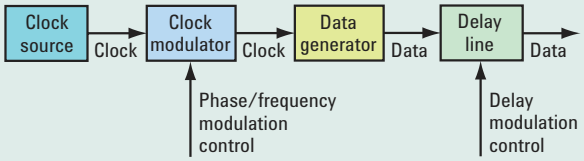

Figure 5. Clock and delay modulation in the pattern generator

Phase/frequency modulation, which is typically located in the clock synthesizer section of the pattern generator, consists of the clock modulator itself (e.g., an I/Q modulator) and control-signal generation (Figure 5). Due to bandwidth limitations in the modulator, the choices of modulation frequency and modulation amplitude are limited; the maximum jitter amplitude depends on the frequency, as illustrated by the blue line in Figure 6. At low frequencies, it is possible to achieve very high modulation of several hundred unit intervals (UI). Clock modulation is used mainly for lower speed sinusoidal jitter, utilized in clock data recovery testing.

The delay modulator is located in the data path of the instrument and consists of a control-signal generator and a voltage-controlled delay line (Figure 5). Typical delay lines have a maximum modulation amplitude of several hundred picoseconds, independent of the modulation frequency (black line in Figure 6). Because the delay time is constant, the maximum modulation in UI depends on the data rate. The main advantage of voltage-controlled delay lines is an extremely high modulation bandwidth, which makes them the method of choice for higher speed period and wideband random-jitter generation.

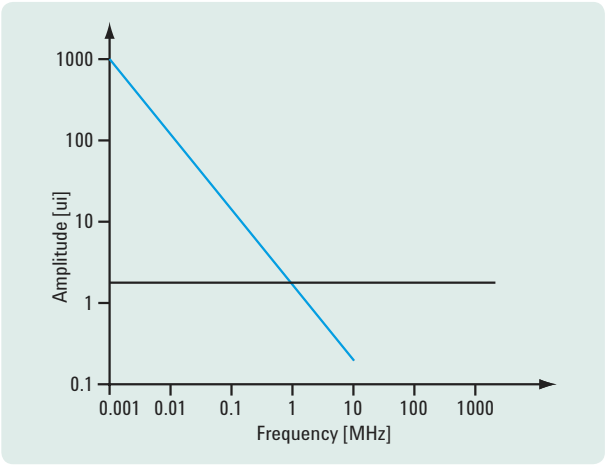

Figure 6. Maximum sinusoidal jitter amplitude, as a function of jitter frequency. Blue line: frequency/phase modulation; Black line: delay modulation.

Voltage control of the delay makes it very easy to use an arbitrary waveform generator to create jitter distributions with arbitrary time waveforms and probability distributions. This enabled the OIF/CEI standard to specify a very special jitter tolerance signal, bounded uncorrelated jitter (BUJ). Generated from a heavily filtered pseudo-random binary sequence (PRBS), it emulates crosstalk and other bounded Gaussian effects.

Inter-symbol interference

Today, most high-speed serial receivers feature some type of equalization to compensate for inter-symbol interference (ISI), the effect of a band-limited channel. Testing these equalizers requires a signal with a certain amount of ISI. The method of choice uses a real trace on a lossy board material, which has a more realistic loss characteristic than a simple RF low-pass filter. Several standards (e.g., XAUI) have specified compliance backplanes, which are commercially available (Figure 7). For custom designs, it may be necessary to develop an application-specific ISI test board.

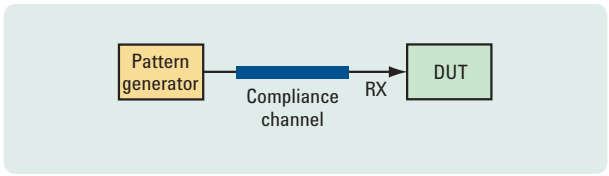

Figure 7. Tolerance test setup for inter-symbol interference using an external compliance backplane

The disadvantage of external board traces is that they need to be inserted manually into the signal path. If we want to make measurements over different trace lengths, or make comparisons with and without ISI, automated measurements are almost impossible, or require numerous expensive and bulky microwave-grade relays. One solution is to build a special test board that has several traces in binary increments (e.g., 1", 2", 4", 8", 16", 32"), connecting the traces via relays such that each one can be switched into the signal path, making it possible to create every trace length between 1" and 63" (Figure 8). Of course, the relays may cause the results to differ from a real backplane; however, with the fine trace-length granularity, it will almost always be possible to find a trace that is close enough.

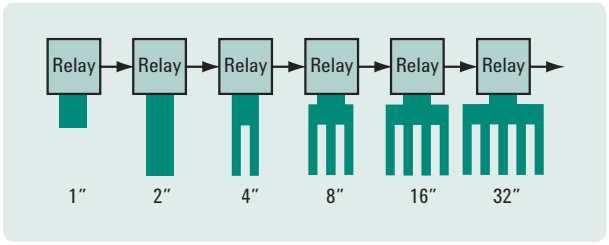

Figure 8. ISI generation using a backplane with switchable traces

Level interference

The last major signal impairment is level noise, which is added to the test-data signal through passive adders (Figure 9). The modulation control signal is usually a sinusoid at frequencies ranging from several hundred megahertz to several gigahertz. Level interference testing comes in two variants: common mode and differential mode. Common mode uses the same signal on both signal rails and tests the common mode rejection of the receiver. In differential mode, the two control signals have a 180-degree phase offset, which tests the receiver's ability to work with small differential amplitudes. Level interference is usually applied after the compliance channel or at the near-end, as seen from the device (refer back to Figure 7).

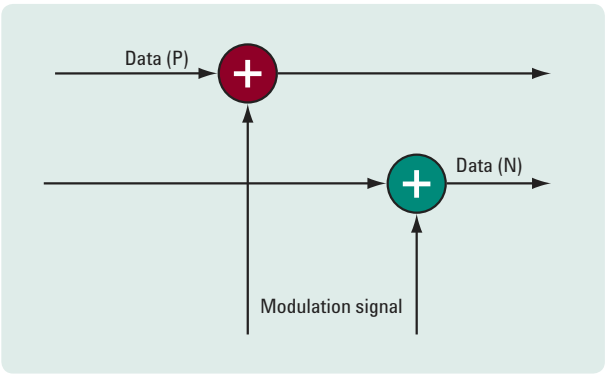

Figure 9. Level interference circuit using passive adders

Calibration

Calibration may be the most challenging aspect of receiver tolerance testing. While is it fairly easy to create a very bad signal, it is difficult to create a signal with precise amounts of jitter, inter-symbol interference and level noise. Without a good calibration strategy, measurements will not be reproducible, making the results essentially meaningless.

Calibration of ISI and level interference is fairly straightforward: Channel S-parameters can be measured directly with a network analyzer and level noise is easily measured on an oscilloscope. Unfortunately, many jitter measurement instruments have tremendous difficulty calibrating the jitter components of the receiver test signal. The first issue is the ability to measure large jitter amplitudes. Most jitter measurement methods are limited to about 0.5 UI total jitter, peak-to-peak, which is sufficient for real-world signals. As shown above, however, a jitter tolerance setup can create several hundred UI of total jitter. The second issue comes from the artificial nature of test-signal jitter. In a sense this violates the valid real-world assumptions measurement instruments must make in order to separate jitter into its components. For example, a common assumption is that spectrally wide components are random jitter, which is not necessarily the case for complicated period jitter waveforms.

To date, the best strategy for jitter calibration is to use special test equipment for each jitter component and measure them one by one. Low-frequency sinusoidal jitter can be measured with a special spectral technique, Bessel Nulls; this works even at extremely high modulation amplitudes. Random jitter and higher speed period jitter are measured with an oscilloscope within the range of the instrument; outside of the range, the delay control signal can be measured directly. The drawback to the "divide and conquer" approach to calibration is the assumption that the individual components are independent and add via convolution; however, in practice this assumption continues to remain valid.

Conclusion

It seems likely that the dramatic pace of evolution in data communications and computing will continue unabated. The ability to test ever-faster serial architectures — and their receiver devices — depends on measurement tools that provide meaningful measurements of critical parameters such as jitter and BER. In both R&D and manufacturing, receiver tolerance is perhaps the most challenging measurement task in high-speed serial device characterization. This task is made easier with standalone bench instruments such as the Agilent J-BERT N4903A high-performance serial bit error ratio tester. Introduced in 2006, the N4903A was the first commercially available instrument to include sources for testing of all jitter components (Figure 10).

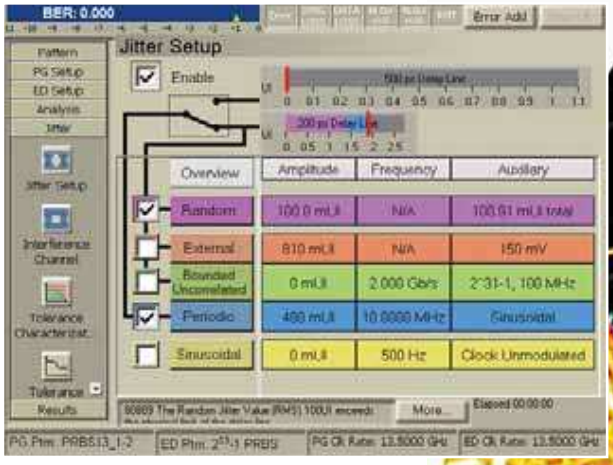

Figure 10. Example of the N4903A user interface showing the manual jitter set-up page

References

Agilent application note *Understanding Jitter and Wander Measurements and Standards* (second edition) available from www.agilent.com

IEEE: 10 GigaBit Ethernet (10GbE) Standard, IEEE 802.0ae

OIF: *Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps and 11G+ bps I/O,* IA #OIF-CEI-01.0, December 13, 2004, Optical Internetworking Forum

Channel Compliance Testing Utilizing Novel Statistical Eye Methodology, Anthony Sanders, Mike Resso, John D'Ambrosia, DesignCon 2004

Agilent J-BERT N4903A High-Performance Serial BERT with Complete Jitter Tolerance Testing, data sheet available from www.agilent.com/find/jbert