

Choose the Right Platform for Your Jitter Measurements

Technical Brief

Signal Integrity

Jitter Measurements—Why Are They So Important?

Edge jitter in serial devices is a growing concern for designers and developers. Jitter matters. It affects system performance. It is difficult to trace and characterize. And jitter limits must be met in order to comply with industry standards such as PCI Express and Fibre Channel.

At the heart of the jitter measurement discipline, there is an understanding that jitter is the “canary in the coal mine”¹—an advance indicator that warns of unacceptable BER performance to come. Jitter causes eye diagram openings to shrink, in turn causing receiving elements to make incorrect decisions about the state of incoming data bits and packets.

¹ Before the advent of electronic gas-detection equipment, coal miners took caged canaries into the mine shafts to alert the workers to any accumulation of dangerous carbon monoxide in the shaft. Similarly, excessive jitter is a precursor of high BER.

And jitter issues aren’t restricted to the serial data path alone. Jitter affects timekeeping elements—clocks—as well. This too can spawn incorrect data decisions, since a misplaced clock edge can cause a data bit’s state to be read too early or too late.

These are the realities of jitter. And the whole problem is compounded with each successive increase in data and clock rates. Serial bus data rates are advancing well into the multi-gigahertz range, narrowing the tolerance for any kind of timing errors. Today, measurement vendors as well as serial bus designers are investing substantial resources into research about jitter behavior. The result is a broader understanding of jitter’s constituent parts and its impact on system performance. New approaches to jitter measurement are giving engineers more insight into their designs and better tools to confront the causes and effects of jitter.

This technical brief is aimed at designers, researchers, engineers, and technicians working with serial data communication architectures of all kinds. For these individuals, understanding and measuring jitter is the surest path toward solving BER problems at the system level. This document will explain some essential jitter terms, and then go on to discuss jitter measurements and the tools best suited for evaluating and quantifying jitter.

“Industry-Standard” Definition Doesn’t Tell The Whole Story

Most serial communication standards specify tolerances for jitter. In addition to PCI Express, other standards that define jitter limits include SATA, SAS, USB, Infiniband, HDMI, and more. Clearly, standards committees and developers are sensitive to the problem of jitter.

Even so, some standards are surprisingly vague in their specifications. The term “jitter” actually denotes several different jitter types, some more applicable to clock signals, others more pertinent to data lines. Standards documents tend to outline quantifiable jitter limits but may not offer much guidance toward determining which type is most critical in a given application. All forms of jitter have the potential to disrupt system BER, and different tools have different strengths in detecting them.

ANSI/INCITS² defines “timing jitter” as follows:

“The short-term variations of the significant instants of a digital signal from their ideal positions in time. Here short term implies phase oscillations of frequency greater than or equal to 10 Hz. Timing jitter may lead to crosstalk and/or distortion of the original analog signal and is a potential source of slips at the input ports of digital switches. It may also cause slips and resultant errors in asynchronous digital multiplexers.”
 [ANSI/INCITS T1.101-1999]

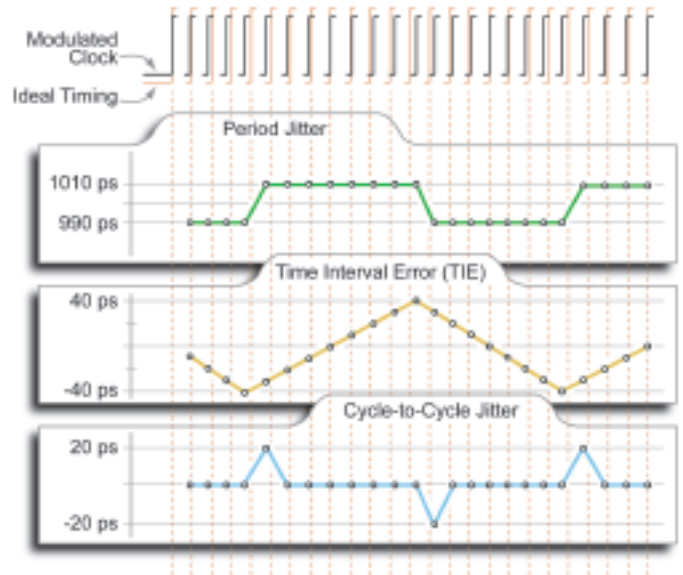


Figure 1. Three ways of analyzing jitter.

This is a good working definition of the term. But for designers attempting to make high-speed serial systems work flawlessly, it doesn’t go far enough. The jitter overview in Figure 1 is not all-inclusive but it expands the definition considerably.

The uppermost trace in Figure 1 is a rendering of a signal waveform. The pulses shown in red are perfectly-timed “ideal” clock cycles exactly one nanosecond (1000 ps) in duration. The black traces are actual edges that exhibit jitter. Here, the trailing edges of both waveforms have been omitted to more clearly show the time-varying nature of the jitter. Note that the black rising edge shifts in time, sometimes occurring later than the red ideal edge and sometimes earlier. The left-most edge (red/black pairing) is perfectly in time. For the purposes of discussion, the measurement begins with 2nd edge, in which the black edge (the actual observed signal) occurs 10 ps early.

² ANSI: American National Standards Institute;
 INCITS: International Committee for Information Technology Standards

Each dot on the three tabbed jitter plots represents one measurement, denoting the actual placement of one edge.

The **Ideal Timing** waveform shows a period of exactly 1000 ps. Dashed lines have been extended through all three plots to highlight the shifts between the ideal timing and the timing of the actual measured edges.

The uppermost plot, **Period Jitter**, is a simple view showing period measurements of the modulated clock waveform over time. Here the period varies between 990 ps and 1010 ps. Some of the periods are 10 ps too short, followed by others that are 10 ps too long. The signal repeats this pattern.

Developers working with clock circuitry, including spread-spectrum and frequency-hopping implementations, will be most interested in period jitter. They strive for clock designs that will deliver the most accurate, predictable frequency and overall stability. Short clock periods amount to instantaneous increases in frequency, and can lead to data errors due to violation of key parameters like setup and hold timing. Obviously period jitter runs counter to this and must be measured and minimized.

Time Interval Error (TIE) Jitter is the summation of Period Jitter in a discrete fashion, period by period. The TIE plot documents the accumulating error between the ideal edge and the actual edge placement. Looking at the first four dots, or measurement points, the TIE plot tracks steadily downward (goes more negative). Each TIE point is 10 ps further away from the zero line, since the successive dots on the Period Jitter plot are each -10 ps in error. Finally when the period jitter changes from every period being short to every period being long the TIE plot “turns the corner” and heads in the positive direction after accumulating -40 ps worth of error. The TIE values are obtained by measuring the time interval from the actual edge to a reference clock.

Note that the TIE jitter value can exceed the Unit Interval (UI), a characteristic that can easily cause confusion. But the problem is readily understood if you consider that TIE can grow with every cycle. For example, if the TIE accumulates 0.1 UI in each cycle, then its value will equal 5 unit intervals after 50 cycles.

TIE measurements are useful when examining the behavior of transmitted data streams, where the reference clock is often recovered from the data signal using a phase-lock-loop (PLL). If the PLL is slow to respond to the data stream’s changing bit rate, then TIE jitter will accumulate during the number of cycles it takes the PLL to catch up.

Moreover, a physical channel impairment phenomenon known as Inter-symbol Interference (ISI) can cause distorted signals. The result is “temporal spreading and consequent overlap of individual pulses to the degree that the receiver cannot reliably distinguish between changes of state” in the data stream. ISI manifests itself in the time domain as Data-Dependent Jitter (DDJ), one of the various forms of jitter included in TIE and other jitter measurements.

The **Cycle-to-Cycle Jitter** plot traces the dynamics of change in the clock signal. The plot depicts each cycle’s change from the previous cycle. Compare the Cycle-to-Cycle with the Period Jitter plot. If there is no change from one cycle to the next, then the Cycle-to-Cycle plot line is flat. When the edge timing switches from steadily accruing error in either direction (later or earlier) the Cycle-to-Cycle plot shows the inflection point with a peak, then returns to the steady-state value. The peaks, positive or negative, occur only when there is a change in this direction.

Cycle-to-cycle jitter is useful for monitoring a clock to ensure that it does not inadvertently produce short cycles that might exceed requirements such as setup and hold timing.

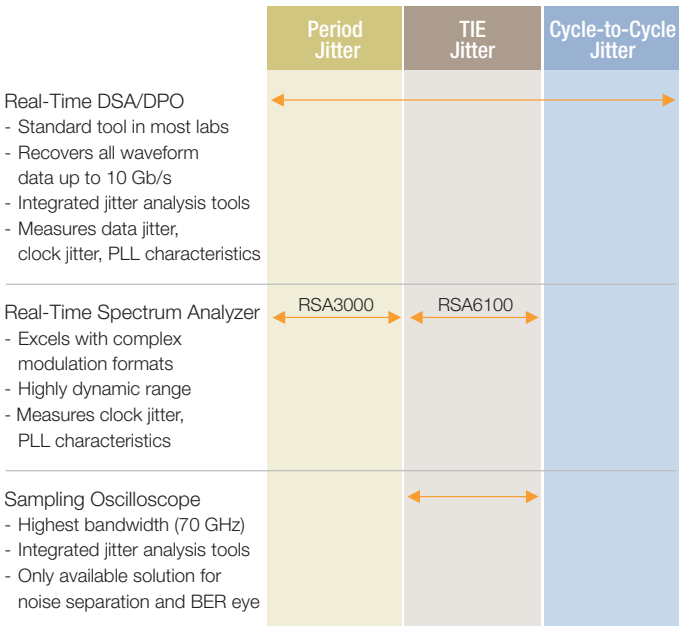


Figure 2. Each class of instrument has its particular strengths for jitter measurement.

Solutions Pursue Challenging Jitter Measurements

The most common solution for acquiring and analyzing jitter is the real-time oscilloscope. Modern digitizing instruments, digital phosphor oscilloscopes (DSA/DPO), can handle today’s multi-gigabit data rates. Equally important, they can be equipped with integrated software applications that perform detailed analysis of jitter and its components.

But the range of choices is not limited to DSA’s and DPO’s. Other entirely different platforms have their strengths, and some of their measurement capabilities overlap. Figure 2 depicts some leading jitter measurement and analysis platforms, summarizing the strengths of the various solutions in their respective application areas.

The Real-Time Oscilloscope

Because it is one of the most common measurement tools used in electronic research, development, and engineering, the real-time oscilloscope is likely to be the first line of defense when jitter issues need to be investigated. But the DSA/DPO is more than just an expedient choice; it is the ideal solution for many jitter applications and can address almost every kind of jitter measurement within reach of its bandwidth and resolution.

The DSA/DPO solution owes its jitter measurement versatility to the fact that it can capture the entire signal waveform, every detail, for many operational cycles of the DUT. Because a long history of waveform activity is preserved in the oscilloscope’s sample memory, we can study attributes as varied as rise time, pulse width, and... jitter of all kinds.

High-performance DSAs can acquire uncompromised signals from many of today’s leading serial buses, including Serial ATA Gen II, PCI-Express 2.0, Fibre Channel, and more.

Applicable features of industry-leading DSAs include:

- 20 GHz bandwidth
- Exceptionally low jitter noise floor: nominally 300 femtoseconds (300 x 10⁻¹⁵). Minimizes the oscilloscope’s contribution to the DUT jitter measurement.
- 8-bit acquisition provides dynamic range sufficient for most current serial standards; suitable for 16-level modulation schemes

An equally important part of the solution is a toolset that automates jitter measurements and analysis. Jitter measurement is a rarified discipline, but one that lends itself to application-specific software solutions (assuming the oscilloscope platform supports such functions).

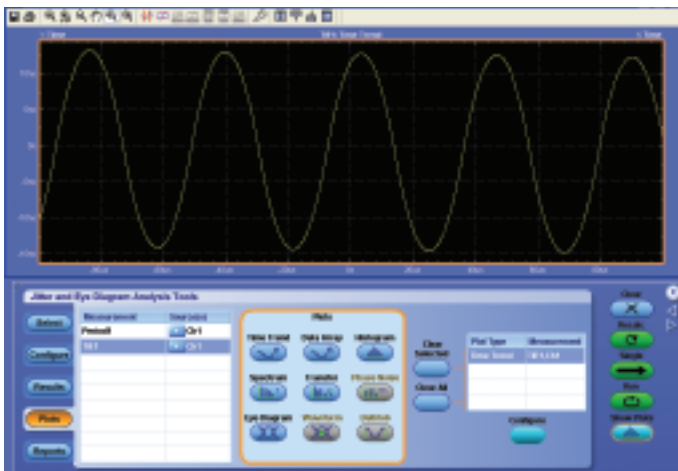


Figure 3. TIE Trend plot from DPOJET showing triangular SSC modulation. TIE integrates timing errors from successive periods, so the TIE trend of the triangular modulation is similar to a sinusoidal curve.

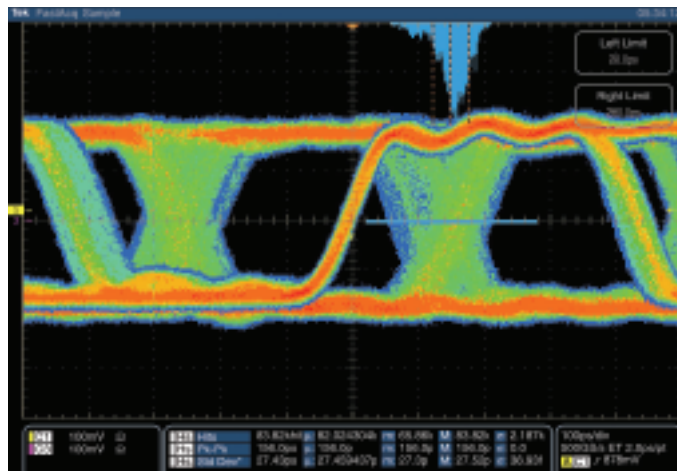


Figure 5. An equivalent-time eye diagram rendering of a serial data signal.




Figure 4. Period Trend plots can be useful for understanding modulation parameters. Here, the plot from DPOJET is showing triangular SSC Modulation. The plot shows SSC applied to a 3.0 Gb/s disk data signal.

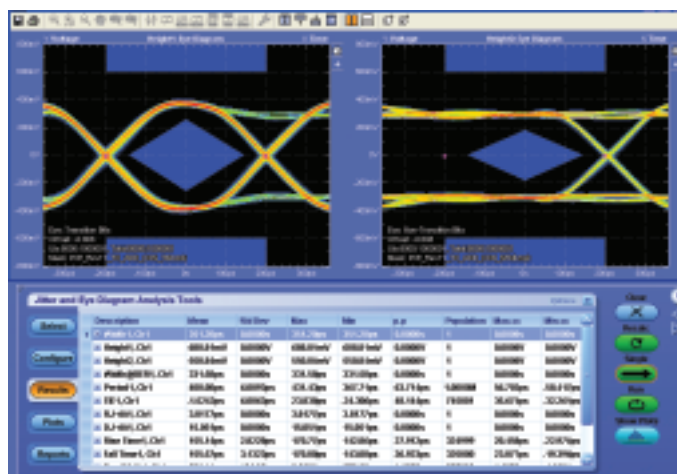


Figure 6. The DPOJET tool renders eye diagrams from single-shot acquisitions taken in real time. The advantage to this approach is its ability to distinguish the bit type and to measure and display transition and non-transition bits independently, as shown here.

An example of this type of product is Tektronix' DPOJET Jitter and Eye measurement package. Figures 3 through 6 illustrate results obtained with this tool.

There are some applications whose demands exceed the capabilities of the real-time DSA/DPO. The instrument's real-time bandwidth and resolution must be weighed against the DUT's data rate and its harmonics. And some esoteric forms of multi-level modulation may tax the instrument's

ability to distinguish between levels. In such instances, one of the other solutions will be more appropriate.

In general, though, a DSA or DPO—especially when equipped with integrated analysis tools—is an unsurpassed platform for jitter measurements. Because it embraces the entire continuum of jitter types, it will remain the preferred solution for most jitter analysis challenges.

Real-Time Spectrum Analyzer

Tektronix Real-Time Spectrum Analyzers (RTSA) are earning wide acceptance as general purpose frequency-domain tools and are gaining a foothold in the realm of jitter measurements. Like earlier swept spectrum analyzers, the RTSA has greater dynamic range than is available in a real-time oscilloscope.

Because the RTSA is by definition band-limited to its capture bandwidth, which in turn is processed by a resolution bandwidth (RBW) filter, the noise floor of an RTSA measurement is much lower than that of an oscilloscope. This makes the RTSA much more sensitive to low level spurious signals that are buried in the noise. If measurements such as harmonic content and low-amplitude spurious levels are required, the RTSA may be the best tool for the job.

Additionally, the phase noise of the RTSA can be several orders of magnitude better than an oscilloscope, yielding much better phase noise dynamic range for random and periodic jitter measurements. Sensitivity and phase noise dynamic range can be essential when measuring PLL and clock performance.

Current-generation RTSAs such as the Tektronix RSA6100 Series and RSA3000 Series offer excellent phase measurement resolution. They can measure frequency and phase changes with greater accuracy than the same information acquired with conventional time-domain methods. This applies to jitter measurements because tiny changes in a signal's phase equate to time. When a spectrum analyzer measures jitter, it is actually measuring phase noise.

To understand that concept, consider that a signal that changes phase by 1 degree has a 1/360 change in its period. Changes in the phase of a carrier create modulation sidebands in the signal's spectrum. These sidebands can be directly measured with a spectrum analyzer. Random jitter is equivalent to the spectral phase noise floor, while periodic jitter is equivalent to all the spurs in the spectrum.

Depending on the instrument used, RTSAs can measure TIE, or Random and Periodic Jitter. In the RSA6100 Series, phase noise measurement option 10 integrates all power




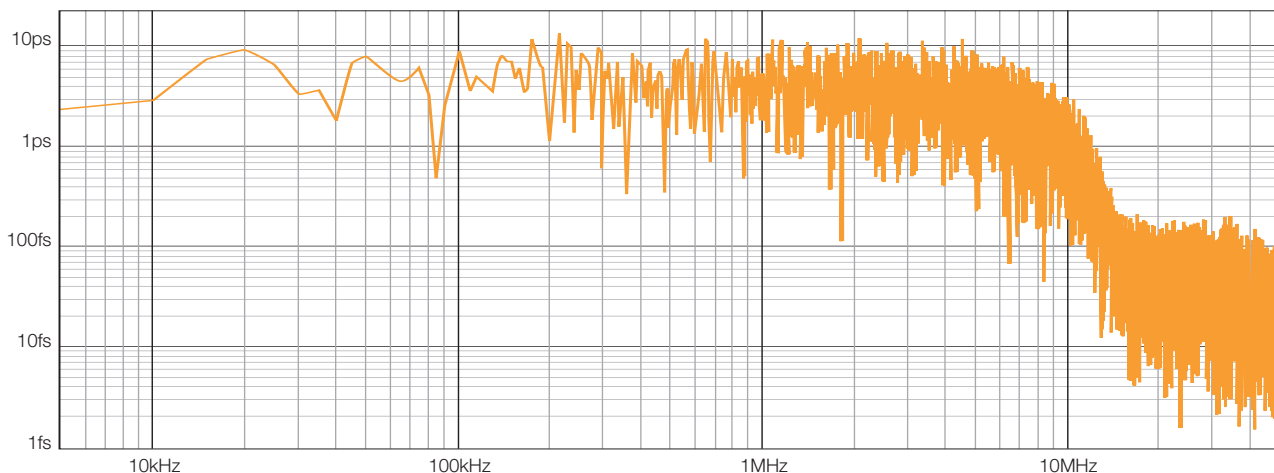
Figure 7a. RSA6100 Series Option 10 calculates Time Interval Error for a user-selectable integration bandwidth.

across the user-defined measurement bandwidth (up to 1 GHz) and calculates the Time Interval Error (see Figure 7a) in a manner similar to a swept spectrum analyzer. While RTSA do provide high precision, low residual measurements of jitter, they do not provide the rich time-domain data set available from real time oscilloscopes. Compare the single value of TIE shown in Figure 7a from the RTSA, to the statistical information available from the RT oscilloscope in Figure 7b.

While the RTSA Jitter and oscilloscope TIE Std Dev. Are in good agreement, the oscilloscope provides additional detail about the statistical nature of the signal including peak to peak, minimum and maximum TIE values. Conversely, the RTSA provides information with respect to the carrier frequency and power, RMS phase noise and residual FM present in the signal, information not available in the oscilloscope measurement.

In the RSA3000 Series with the Signal Source Analysis Suite, a part of Option 21, Random and Periodic Jitter measurements (Figure 8) are performed by analyzing contiguous time domain samples transformed to the frequency domain. Minimum Random Jitter measured on a 1 GHz Carrier (1 KHz to 10 MHz Bandwidth) can be as low as 0.2 picoseconds RMS. Minimum Periodic Jitter measured on the same signal is typically as low as 2 picoseconds peak—excellent precision by any standard.

TIE1: Spectrum



Description	Mean	Standard Dev	Maximum	Minimum	p-p
TIE1, Ch2	0.0000s	147.28ps	617.13ps	-595.41ps	1.2125ns
PN1, Ch2	152.93ps	0.0000s	152.93s	152.93s	NA

Figure 7b. A Real-Time oscilloscope calculates TIE for the same signal as in Figure 7a, including statistical detail of the signal.

The RTSA makes it possible to measure the jitter on a small signal in the presence of larger ones. The unique Frequency Mask Trigger function homes in on the particular instants in time when the signal's spectrum deviates from the normal, allowing you to troubleshoot sources of jitter. Note also that an RTSA is ideally equipped to handle complex modulation schemes and other aspects of mobile voice and data transmissions. And because the RTSA has a large memory coupled with a variable effective sample rate, it can store several seconds' worth of signal activity. The RTSA is ideal for measuring random and periodic jitter in clocks and local oscillators, and for examining clock recovery systems containing PLLs and DLLs. It is less suitable for data signals and signals containing data-dependent jitter.

The RTSA has among other features the ability to capture and store a seamless, continuous record of spectral "snapshots" and process the results in many different combinations.

Real-Time Spectrum Analyzers can measure transient events and cycle-to-cycle changes, and are very useful

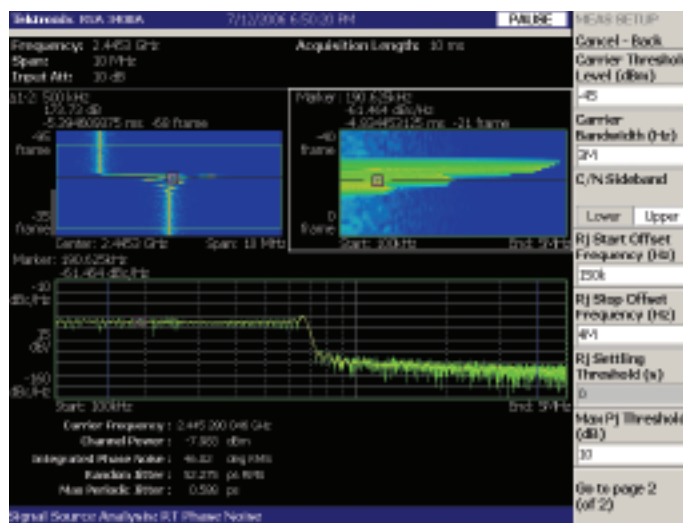


Figure 8. The RTSA calculates random and periodic jitter values from its phase noise readings.

for observing small phase variations and providing detailed views of modulation.

The Sampling Oscilloscope

The sampling oscilloscope brings prodigious bandwidth to the task of jitter measurement. At this writing, 70 GHz sampling oscilloscopes deliver unmatched signal fidelity for high-frequency time-domain measurements. A sampling oscilloscope may be the only effective solution when observing signals having data rates up to 60 Gb/s. Moreover, the sampling oscilloscope is the tool of choice when it is necessary to capture the harmonics of relatively “slow” signals. For example, a 10 Gb/s NRZ signal has a 5 GHz fundamental, but its 5th harmonic is 25 GHz--significantly beyond the reach of real-time DSA and DPO instruments today

Sampling oscilloscopes rely on repetitive input patterns to build up a waveform acquisition made up of samples taken over numerous cycles. Many types of serial devices offer diagnostic loops that can produce these repeating waveform streams, or an external data generator can be used as a driving source.

The sampling oscilloscope has a powerful ally in the field of jitter measurements: application-specific software packages for jitter/noise analysis. The first comprehensive tool to integrate features beyond purely temporal jitter is the Tektronix 80SJNB Jitter, Noise, and BER Analysis package. Its noise analysis features set the package apart from other jitter tools on the market. Tektronix sampling oscilloscopes extend the concept of jitter analysis to encompass jitter separation, noise separation, and BER eye estimation—all at bandwidths compatible with the highest serial data rates in use today.

Noise Analysis Joins Jitter Analysis To Confront BER

As explained earlier this document, designers combat jitter in order to minimize its effects on BER. But jitter is only one part of the BER story—it is the time, or horizontal component, of the total. A vertical component exists and plays a role in BER as well. This component is noise. A thorough jitter analysis

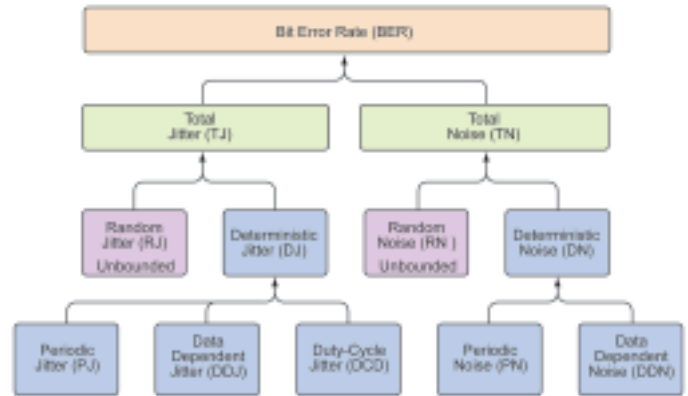


Figure 9. Both Jitter and Noise contribute to the total BER figure.

can be done equally well with a real-time oscilloscope, or an RTSA, or a sampling oscilloscope. But at this time the only solution for analyzing noise and its BER contribution is a sampling instrument.

Figure 9 depicts the many factors that roll up to the total BER figure. While a detailed description of all of these factors is beyond the scope of this discussion, it is important to note that most jitter elements have counterparts on the noise side. All are made up of components that arise from both periodic and random causes within a system.

Noise is an amplitude effect that reveals itself in eye measurements. It vertically constricts the eye, closing it from above and below. Less obviously, noise also affects jitter measurements based on samples taken at a threshold crossing on an edge. The impact of these effects depends on the magnitude of the noise, the slope of the signal, and more broadly, on the transient response of the system.

In some instances, noise may be the dominant effect that degrades BER. But most often, both horizontal and vertical effects play a role in BER. Understanding the difference between noise contributions and pure jitter contributions speeds up root-cause analysis of system problems.

Conclusion

Jitter measurement is a challenge that grows in importance with each advancement in system performance and data rates. There are several approaches to jitter measurement, each with its own strengths and tools. Many basic jitter tests can be handled by real-time oscilloscopes in conjunction with analysis applications that deliver both statistical and waveform information. When more dynamic range is needed, a real-time spectrum analyzer can approach jitter from a phase noise perspective, and additionally can measure jitter in the presence of interfering signals.

And when the signals in question are running at speeds beyond 20 GHz, the sampling oscilloscope is the only tool for jitter measurements. With its unmatched bandwidth and its ability to parse the contribution that noise makes to BER, the sampling oscilloscope delivers unique BER insights along with its jitter measurements.

Jitter Measurement References

1. 55W-16146-1 Understanding and Characterizing Timing Jitter

Contact Tektronix:

ASEAN / Australasia (65) 6356 3900
Austria +41 52 675 3777
Balkans, Israel, South Africa and other ISE Countries +41 52 675 3777
Belgium 07 81 60166
Brazil & South America (11) 40669400
Canada 1 (800) 661-5625
Central East Europe, Ukraine and the Baltics +41 52 675 3777
Central Europe & Greece +41 52 675 3777
Denmark +45 80 88 1401
Finland +41 52 675 3777
France +33 (0) 1 69 86 81 81
Germany +49 (221) 94 77 400
Hong Kong (852) 2585-6688
India (91) 80-22275577
Italy +39 (02) 25086 1
Japan 81 (3) 6714-3010
Luxembourg +44 (0) 1344 392400
Mexico, Central America & Caribbean 52 (55) 5424700
Middle East, Asia and North Africa +41 52 675 3777
The Netherlands 090 02 021797
Norway 800 16098
People's Republic of China 86 (10) 6235 1230
Poland +41 52 675 3777
Portugal 80 08 12370
Republic of Korea 82 (2) 6917-5000
Russia & CIS +7 (495) 7484900
South Africa +27 11 206 8360
Spain (+34) 901 988 054
Sweden 020 08 80371
Switzerland +41 52 675 3777
Taiwan 886 (2) 2722-9622
United Kingdom & Eire +44 (0) 1344 392400
USA 1 (800) 426-2200

For other areas contact Tektronix, Inc. at: 1 (503) 627-7111

Updated 12 November 2007

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2008, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

08/08 DM

37W-19858-1

Tektronix[®]
Enabling Innovation

