

Jitter: Measurements and Instrument Solutions

Introduction

Jitter is one of the primary reasons for transmission errors in circuits and products. When choosing an instrument to address jitter, consider the types of tests you will be conducting, the characteristics of the devices you will be testing, and also the testing environment. For instance, some instruments are more appropriate for an R&D environment. Others deliver higher speed and lower-cost-per-test and are therefore better suited to a manufacturing environment. The data rate of the devices you will be testing and any governing specifications also play roles in your tool choices. This application note discusses jitter measurements at high data rates, issues impacting measurement accuracy, and looks at the instrument choices available for observing and analyzing jitter.

Jitter Measurements at High Data Rates

In telecommunications and enterprise arenas jitter specifications and measurements are governed and documented by standards bodies. For instance, as in the telecommunications world, random/deterministic jitter is not a governing measurement. In fact, in many high-speed I/O sectors, new bus standards are being introduced with little commonality in specifying and measuring jitter. This forces designers and test engineers to clearly understand the different jitter measurement vantage points—and which measurement techniques will be best for them.

What follows concentrates on the jitter measurements required by many of today's high-speed digital transmission system, as well as bus and interconnect standards. These include SONET/SDH/OTN, Ethernet, InfiniBand, PCI Express®, and Serial ATA. The various transmission systems, market sectors, standards, and the corresponding jitter measurement techniques are identified in Table 1.

Table 1. Jitter measurements organized by transmission system, technologies, and standards.

	Telecom	Enterprise	High-Speed I/O
Transmission system (applications)	Metro/long distance network traffic. (Voice/data backbone)	Up to 10-km point-to-point, VOIP/TCP-IP. (Building, site, campus)	Point-to-point traffic from centimeters to meters (memory, storage, peripherals)
Technologies	SONET, SDH, OTN, OC-48/192/768, STM-16/64/192, OTU-1/2/3	100-Mb, 1-GbE (copper), 10-GbE (optical), 3.125-G XAU electrical	InfiniBand, S-ATA, USB, PCI Express, FibreChannel, RapidIO, HyperTransport
Standards	ITU-T G.783, G.825, G.8251, Telcordia GR-253	IEEE-802.3ae (various)	Various NCITS-TR25
Jitter measurements	Unit Interval (UI) vs. frequency. Sinusoidal test with jitter analyzers	Stressed-eye, transmitter-dispersion penalty (TDP), worst-case statistical test with BERTs and DCA	Eye diagram, bathtub, histogram, RJ/DJ. Jitter specifications are evolving. Test with oscilloscopes, DCAs and BERTs

SONET/SDH/OTN jitter measurements

The international SONET/SDH/ OTN standards specify measuring jitter generation, jitter transfer, and jitter tolerance in efforts to confirm that jitter is within bounds. In the SONET/SDH/OTN sectors, the term 'jitter' is employed where the frequency of the unwanted phase modulation is greater than 10 Hz. When the frequencies are less than 10 Hz, the standards refer to unwanted modulation is referred to as 'wander'.

International SONET/SDH/OTN standards further band-limit the spectrum when verifying compliance. Components of the baseband are excluded from the measurement. This means that different filter values are used for different data rates. For example, for 10-Gb/s SONET/SDH signals, bandpass filters rejecting jitter components below 20 kHz and above 80 MHz are employed. Whereas an additional high-pass filter is used to measure frequencies above 4 MHz.

Jitter generation

Jitter generation, sometimes called intrinsic jitter, refers to the jitter present at the output of a single device (output jitter refers to that from the network). It is specified in unit intervals and the result is quantified as an RMS or a peak-to-peak value. Most measurement systems use high-pass and low-pass filters to band-limit the spectrum.

Output jitter results are strongly influenced by the data being carried. Test results can vary widely between those for a simple repetitive pattern such as 101010... and those for a complex PRBS (pseudo-random-bit-sequence) pattern. It is important that the data being transmitted be defined when measuring and specifying jitter generation. Likewise, it is crucial when measuring jitter generation that the result not be affected by the intrinsic jitter of the measurement system.

Jitter tolerance

To make sure that your devices will operate error free in the presence of the worst-case jitter from preceding sections in the network, you need to measure jitter tolerance. Jitter tolerance measurement requires a source of sinusoidal jitter and a method to assess BER, such as an instrument like that shown in Figure 1.



Figure 1. Testing instruments such as the Keysight Technologies, Inc. OmniBER OTN provide both a source of sinusoidal jitter and a method to assess BER in a single enclosure (depicted here on the instrument's screen).

A signal is generated with added sinusoidal jitter and applied to the DUT. At each jitter frequency, the amplitude of the jitter is increased until transmission errors are detected. Alternatively, a specified level of input jitter is generated and error-free operation is checked. In the real world, jitter is unlikely to be sinusoidal, but it is easy to generate and gives repeatable results.

Jitter transfer

Jitter transfer describes how a clock recovery module or repeater locks and tracks data as jitter is placed on it. Jitter transfer requirements for clock recovery circuits specify a minimum amount of jitter gain vs. frequency up to a given cut-off frequency, beyond which the jitter must be attenuated. The jitter transfer specification is intended to prevent the buildup of jitter in a network consisting of cascaded regenerators. When measuring jitter transfer, you need to make sure that your DUT does not transfer more than the SONET/SDH-mandated 0.1 dB of peaking.

Wander

Whereas jitter is normally measured with reference to a clock extracted from the data signal, wander is measured against an external reference clock. The fundamental measurement is time interval error (TIE)—the instantaneous time deviation of the clock signal under test relative to the reference source. TIE samples are used to calculate wander results such as maximum TIE (MTIE) and time deviation (TDev).

Selecting Instruments for Jitter Measurement

Instruments to use for jitter evaluation can be grouped into two categories: stimulus sources and instruments that perform jitter measurement and analysis. Pulse/pattern generators are the principal instrument in the stimulus category, while there are several instruments that perform jitter measurement analysis. This group provides the variety of jitter vantage points needed to obtain a comprehensive picture of the jitter affecting your device or system.




Figure 2. The Keysight Technologies, Inc. Infiniium Z-Series oscilloscopes provide the industry's lowest noise and jitter measurement floors.

Pulse/pattern generator

The pulse/pattern generator creates and delivers arbitrary patterns in both differential and single-ended structures with a minimum of phase noise. Instruments in this family enable the user to select pseudo-random structures (PSRSs) that emulate random data streams with lengths ranging from bits to megabits. Usually pulse/pattern generators are equipped with a jitter or delay control so that the user can set up a pattern with a precise amount of jitter. This, in turn, enables the user to characterize a circuit's response to it—a requisite for both jitter-tolerance and jitter-transfer tests.

Low-level, phase-noise/spectrum analyzer

To obtain an intrinsic jitter spectrum, using a phase-noise or low-level jitter analysis system is necessary. These instruments provide the highest level of accuracy in measuring the frequency content of the jitter in your design. They derive their accuracy from a high degree of over-sampling and a narrow measurement bandwidth. Such systems uncover deterministic jitter mechanisms not detected by oscilloscopes. These systems exhibit extremely low noise floors and are immune to amplitude noise.

Phase-noise systems are preferred for this purpose over spectrum analyzers for a variety of reasons. Low-level jitter analysis can be used as a design or troubleshooting aid when examining noise-floor mechanisms, phase-locked loops, VCOs, crystal oscillators, and other clocks and references—situations in which random jitter needs to be carefully monitored. Using a precision system based on frequency-domain phase measurement techniques is critical for this type of analysis. Spectral evaluation often leads to great insights into design and system idiosyncrasies not observable by other techniques. However, low-level analysis is limited to jitter components less than 200 MHz. This is substantially below the full bandwidth many standards stipulate, so other jitter analysis tools are also required for spectral analysis requiring greater bandwidth.

Real-time sampling oscilloscopes

High-speed, real-time digital storage oscilloscopes (DSOs) are versatile and flexible, and therefore a frequent choice for jitter analysis. A rule of thumb is that the bandwidth of analysis should be at least 1.8 times the maximum bit rate for a non-return-to-zero (NRZ) serial signal. So this confines a DSO with a 6 GHz bandwidth to a maximum data rate of 3.2 Gb/s. These oscilloscopes, generally exhibit a jitter measurement floor of 1.5 picoseconds of jitter, or less.

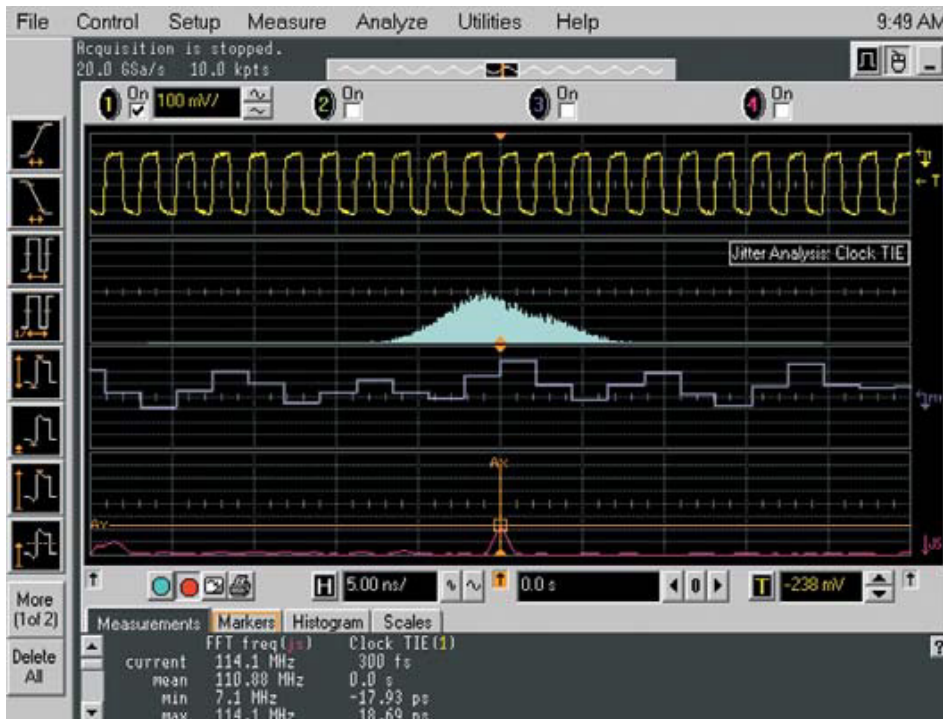


Figure 3. High-speed, real-time digital storage oscilloscopes can display a number of views of the same signal. Shown here are a number of views of the 456 MHz clock signal at (A): a histogram of a transition point (B), TIE vs. time (C), and an FFT of the TIE data (D).

To capture the entire signal in a single acquisition, DSOs oversample a signal at least two times and usually at more than 3.5 times the maximum bandwidth. They may interpolate between these samples to increase the effective time resolution of a waveform capture. Oversampling oscilloscopes acquire a signal in about one-tenth the time of under-sampling oscilloscopes, such as digital communications analyzers, but they are considered slow compared to a BERT with respect to appraising edge-to-edge metrics. You can route signals directly to an oscilloscope or tap into an active circuit using high-bandwidth probes.

Be careful to select and use probes correctly. Often they are the weakest link in a data acquisition setup. Once a waveform is captured, many measurement and display functions can be employed such as producing eye diagrams, recovering the transmitter clock, and determining time interval error (TIE), duty-cycle, rise-time, and fall-time parameters.

Oscilloscopes can display histograms and trend lines for any of these parameters and perform FFTs (fast Fourier transforms) to determine frequency content of the data. You can display a variety of jitter parameters such as cycle-to-cycle, n-cycle, period, and delay. Also you can display simultaneously the data waveform, time-trend data, and the FFT of these measurements, as shown in Figure 1, thereby providing you with powerful diagnostic capabilities. Oscilloscopes can also be used in conjunction with jitter analysis software to provide further capabilities such as random/deterministic jitter separation and to confirm compliance with specific standards.

Digital communications analyzers (DCAs)

DCAs are members of the oscilloscope family, but the undersampling method they employ results in significant differences from a DSO. They can reach bandwidths exceeding 80 GHz. Using the above-mentioned rules of thumb for NRZ-serial communications, these instruments are also the only oscilloscopes that can be used for data rates above 3.2 Gb/s. They accomplish this with an under-sampling method that requires a repetitive signal.

Based upon the arriving repetitive signal and a reference related to it, the DCA samples at times internally derived and offset in time from each successive trigger. It thus constructs the waveform in a piecemeal fashion. This requires a very precise time base. The under-sampling method requires more time to acquire a waveform than analyzing oversampling methods. DCAs require a low-jitter trigger event and cannot capture contiguous cycles on a single trigger.

Though DCAs are the only oscilloscope option for data rates exceeding 3.2 Gb/s, they can be used for eye diagram analysis at all data rates. They also have the capabilities to perform other roles such as time-domain reflectometry. Finally, DCAs are also less expensive than comparable-bandwidth DSOs.

Bit error ratio testers

Bit error ratio (BER) is a key performance metric in digital communications designs and BERTs play a significant role in testing these systems. They provide comprehensive stimulus/response testing, they can also perform stimulus-only and response-only functions. BERTs provide solutions up to 45 Gb/s and come in different configurations: serial-only, serial-to-parallel, parallel-to-serial, and parallel-to-parallel. These configurations enable them to test a wide variety of different systems. A good example is a serializer-to-deserializer (SERDES) module.

Because BERTs count every edge or transition on a waveform, they provide the fastest measurement on a per-transition basis. BERTs enable adjustment of the sample-time location and the threshold levels. These features are useful for generating point eye diagrams and iso-BER diagrams (with contour lines delimiting equal probability areas in the eye diagram). BERT tools enable bathtub plot creation, bathtub plot extrapolation—which speeds the creation of bathtub plots—and random/deterministic jitter separation.

Logic analyzers with EyeScan

Logic analyzers are seldom considered for obtaining parametric measurements or for studying jitter. However, when equipped with a feature called EyeScan, they can monitor every edge of up to 300 parallel lines simultaneously on systems up to 1.5 Gb/s (shown in Figure 4). They are also well suited for Skew on parallel buses. EyeScan can provide measurements with 10 ps resolution. If a trace exhibits excessive skew or eye closure, the user interface allows you to activate that single trace, and thereby to focus in on it for detailed analysis.




Figure 4. Logic analyzers with EyeScan can monitor every edge of up to 300 parallel lines simultaneously on signals up to 1.5 Gb/s, with 10-picosecond resolution.

Jitter Measurement Accuracy

The accuracy of jitter measurements depends not only on the characteristics of the measurement system, but also on the characteristics of the measured signal and the type of jitter being measured. This means that the same measurement system, measuring the same signal will have a different accuracy when measuring period jitter than it does when measuring cycle-to-cycle jitter. It also means that a given measurement system will exhibit a different accuracy measuring the same type of jitter on two different signals.

For example, voltage noise is typically the dominant source of error in jitter measurements. Since the voltage noise is converted to time jitter by the slope of the signal's transitions, jitter measurement accuracy depends on the slew rate of the signal being measured.

So, the trick to characterizing the accuracy of jitter measurements is to match the target measurement system and target signal characteristics, while measuring a known ideal source. Sine wave sources work well for this application because they are easily adjustable in amplitude and frequency. Also sources with arbitrarily low jitter are readily available.

The procedure described below measures the accuracy of a specific target jitter measurement on a specific target signal using a specific target measurement system. Therefore, it is often useful to perform the target measurement prior to measuring the measurement's accuracy. Performing the target measurement first helps identify the specific characteristics of the target measurement that are needed to perform the accuracy measurements.

Determining the jitter measurement floor

Although some jitter measurements underestimate the true jitter value, most measurements overestimate it. This comes about because the dominant jitter measurement errors tend to be uncorrelated additive random processes, like jitter is itself. Specifying jitter measurement floor (JMF) is one easy way to quantify the amount that a jitter measurement over estimates the true jitter value. Jitter measurement floor (JMF) refers to the lowest value a jitter measurement would produce if it were applied to a perfect source that had zero jitter.

Here are steps in determining the measurement floor in a period jitter measurement:

1. Configure the measurement system to match that of the target measurement system.
2. Determine the slew rate of the target signal. Note that an oscilloscope's built-in differentiate function can be useful for measuring slew rate.
3. Determine the nominal period of the target signal.
4. Connect a low-jitter, sinewave source to your oscilloscope. Duplicate the target jitter measurement setup as much as possible. If the target jitter measurement will be made through a probe, then connect the sinewave source through the same probe.
5. Set the initial input sinewave amplitude so that it fills approximately three vertical divisions on the oscilloscope display.
6. Set the sinewave frequency so that the sum of an integer number of its periods matches the period of the target signal, and the sinewave's slew rate equals the slew rate, approximately, of the target signal. This may take some trial and error. For real-time oscilloscopes, do not use sinewave frequencies that are larger than half the oscilloscope's sampling frequency.
7. Adjust the sinewave amplitude until its slew rate equals the slew rate of the target signal.
8. Measure the jitter of the sinewave at those threshold crossings that match the period of the target signal. Note that you may need to use a more general delta-time measurement feature to measure an integer number of sinewave cycles.

In Figure 5 is an example of a period-jitter, measurement floor measurement for a PCI Express application. The target signal to be measured is a $1.0\text{-V}_{\text{p-p}}$, 2.5-GHz clock signal with a nominal slew rate of 10 V/ns. In this example, two cycles of a $0.6\text{ V}_{\text{p-p}}$, 5.0 GHz sinewave are used to mimic the period and the slew rate of the target PCI Express signal. The oscilloscope is configured to measure the time interval between two clock transitions that are two cycles apart. Measurement statistics are then used to calculate the standard deviation of a large number of these two-cycle time interval measurements.

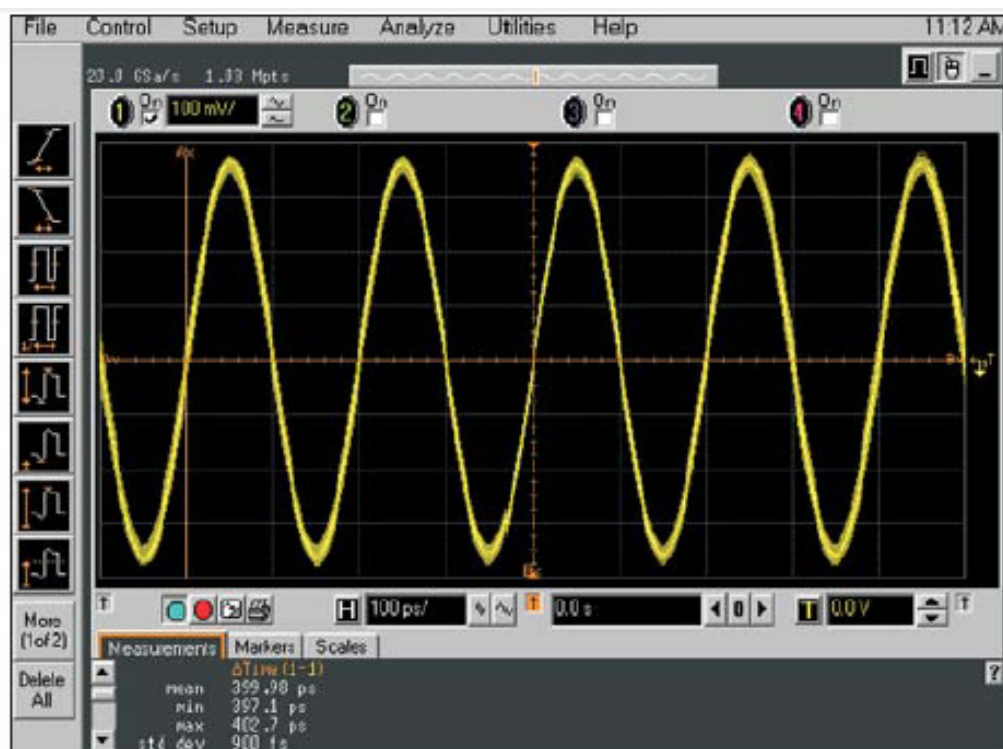


Figure 5. An example of a period jitter measurement floor measurement for a PCI Express application as described in the text. The measured jitter measurement floor is 900 fs RMS.

The procedure for measuring the JMF of a cycle-to-cycle jitter or n-cycle jitter measurement is very similar to that for period jitter. The value of n in the n-cycle jitter measurement must be multiplied by the number of sinewave periods used to represent one target signal period. Since two sinewave cycles are used to represent one target signal period, an n value of two would be used on the signal from the previous example.

The JMF results shown in Table 2 were all measured using the same input signal and measurement system as was used in the example of Figure 4. Notice that the jitter measurement floor is different for all three examples above, even though they are measuring the same signal over the same time interval. These measurement floors differ because they are all calculated from different functions of waveform transitions.

Table 2. Comparison of jitter measurement floor values for three different jitter types.

Jitter type	JMF
Period jitter	0.90 ps RMS
Cycle-to-cycle jitter	1.52 ps RMS
Time interval error	0.66 ps RMS

Time interval error is calculated from the time variation of individual waveform transitions. Period jitter, on the other hand, is calculated from the difference between the time variations of pairs of transitions. Depending on the correlation between these transition pairs, the variation in their time difference may be larger or smaller than the variation of either of the individual transitions.

Though this technique is simple and very effective in most applications, it does have limitations. It does not, for example, include the effects of aliasing. Nor does it include the frequency response flatness effects that appear when measuring non-periodic signals like data signals. In addition, some TIE applications will calculate the time interval error of the test sinewave from more transitions than would be used in the calculation of the target signal over the same total time range. In some cases, this difference can result in overestimation of the jitter measurement floor. Fortunately, this difference is rarely significant and can be avoided by calculating the TIE value using external post-processing.

Subtracting out the jitter measurement floor

Subtracting the jitter measurement floor from raw jitter measurements is always tempting, and often effective, but be careful. Subtracting the variance of the measurement floor from the variance of the measured value, as shown in equation (1), only applies if the measurement errors are completely uncorrelated to the true jitter.

$$J_{corrected} = \sqrt{(J_{meas})^2 - (JMF)^2} \quad (1)$$

Also, as the magnitude of the measured value becomes closer to the measurement floor, the relative uncertainty of the reported value grows. This technique may allow you to report a smaller jitter value, but the relative uncertainty of this value may be greatly increased.

Conclusion

The need to control jitter is influenced by the product's data rate and the standards for product application areas. Understanding these requirements, as well as factors that impact measurement accuracy, can help engineers select from the instrument options available for evaluating jitter.

Related Literature

- Jitter Fundamentals: Sources, Types, and Characteristics
Application Note, literature number 5992-1637EN
- Separating and Time-Correlating Deterministic Jitter Components
Application Note, literature number 5992-1639EN
- Jitter Solutions for Telecom, Enterprise, and Digital Designs
Brochure, literature number 5988-9592EN

Download your Next Insight

Keysight software is downloadable expertise. From first simulation through first customer shipment, we deliver the tools your team needs to accelerate from data to information to actionable insight.

- Electronic design automation (EDA) software
- Application software
- Programming environments
- Productivity software

Learn more at www.keysight.com/find/software

Start with a 30-day free trial.

www.keysight.com/find/free_trials

PCI-SIG®, PCIe® and the PCI Express® are US registered trademarks and/or service marks of PCI-SIG.

Learn more at: www.keysight.com

For more information on Keysight Technologies' products, applications, or services, please contact your local Keysight office. The complete list is available at:
www.keysight.com/find/contactus

