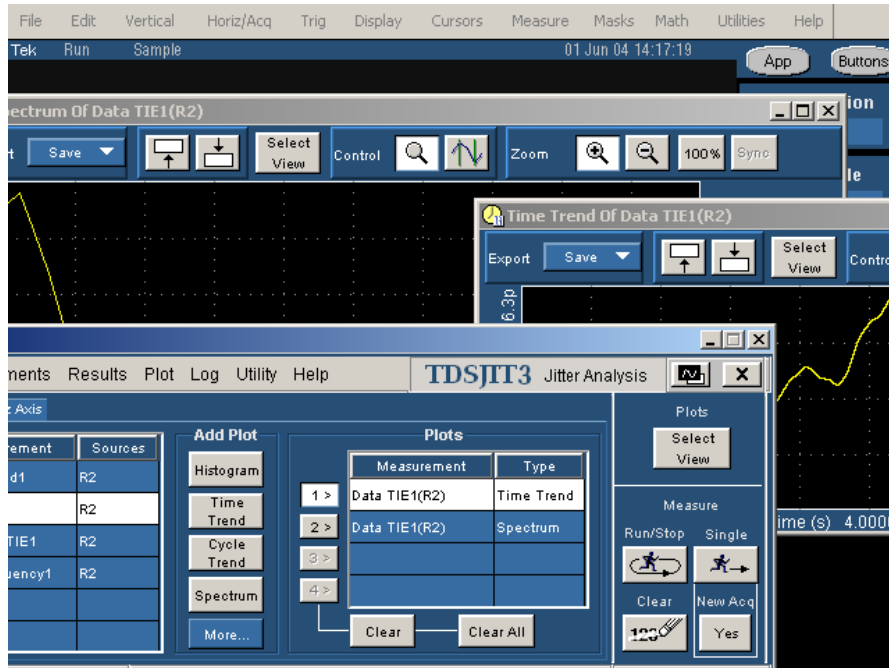# Jitter Analysis: A Brief Guide to Jitter

▶ What is Jitter?

Jitter is noise in the temporal or timing domain.

Yes, it really is that simple. Normally though we apply the term to electrical or optical signals we can measure with oscilloscopes or other time measurement equipment.

We can think of jitter two ways.

- As an instantaneous effect: This one edge isn't where I wanted it to be.
- As an accumulation of effects: in this series of edges, each edge is displaced an equal amount, and the last edge shows the sum of the displacement times.

▶ Who Needs to Measure Jitter?

Any one interested in high-speed digital or serial design, debug, or verification needs to measure jitter. Any time the data rates being implemented includes the terms Mb/s or Gb/s (Mega or Giga-bits per second) or MT/s or GT/s (Mega or Giga-transfers per second), jitter tools are needed.

Even engineers making copiers are concerned with jitter: even obscure things like drum rotation speed variation is critical (also very difficult to measure with any accuracy).

▶ Is There a Difference Between Short-term and Long-term Jitter?

In determining instantaneous effects, we can view jitter as a time variation of one edge relative to another edge. Using a clock PERIOD measurement as an example, we can measure many periods. We really can't tell by period alone whether there are instantaneous changes. Sure, we can see changes by looking

at the standard deviation of all the measurements. We can see a distribution limit by looking at the minimum and maximum periods in the measurement sample we collected. But we don't know anything about HOW the timing changed, or why.

We need to use a bit of detective work. Using real-time oscilloscopes we can measure sequences of adjacent periods. By measuring the differences from one period to the next period we can begin to see how the timing changes in any one instance of a period. A clock period that changes over one period has cycle-to-cycle jitter. Jitter over one cycle can have very high frequency content relative to the fundamental frequency of the circuit being tested.

For example, measuring a 100MHz clock that changes period every cycle – say a longer period one cycle and a shorter period the next cycle – will have jitter at 50MHz. That's one-half the fundamental or center frequency.

We can also have jitter that is totally random, and not periodic as in the previous example. Each cycle can change a random amount: longer for several cycles, shorter for one, longer for one, shorter for two, longer for one, and so on. Given a sufficient number of samples we can determine the changes are random and follow certain statistical profiles.

These types of jitter are instantaneous changes in clock period are best viewed with real-time instrumentation. Looking at cycle-to-cycle jitter has one negative shortcoming: only high frequency jitter is visible. The act of looking at adjacent cycles filters out any longer term effects that are happening in the signal – any effect that occurs in phase with the number of cycles measured will fall into a measurement null point – the window is equivalent to a notch filter.

In looking for accumulated effects, many methods are available. In prior years, even analog oscilloscopes were used to view jitter by triggering on one edge in a repeating signal, like a clock, and looking at a portion of the signal some time after trigger. The time between the trigger and the viewed area was the delay value. With the fastest instruments fitted with delay lines even the trigger point and next edge could be viewed, but more often the delay was great enough that many cycles of the clock would transpire before the instrument could display the signal. Because the full statistical view wasn't available in these analog instruments, it became common practice to look at edges away from the trigger point. Often, this delay value was chosen at random – with twenty microseconds (20µs) one common value in use. Sometimes the delay was a calculated number of cycles later.

This delayed technique has a few flaws. The most important flaw to consider is that if the count of cycles isn't precise, the signal may have far more noise than what is measured. With a timing noise source that has a fixed modulation, for example a switching power supply injecting noise into an oscillator, the noise may change the frequency abruptly between two common values: a rectangular modulation. Using the delayed technique to measure jitter, the phase change induced by the power supply noise may be greater than one clock period. The instrument will show this as a slightly wider jitter, with the fact the jitter is greater than one period hidden from view.
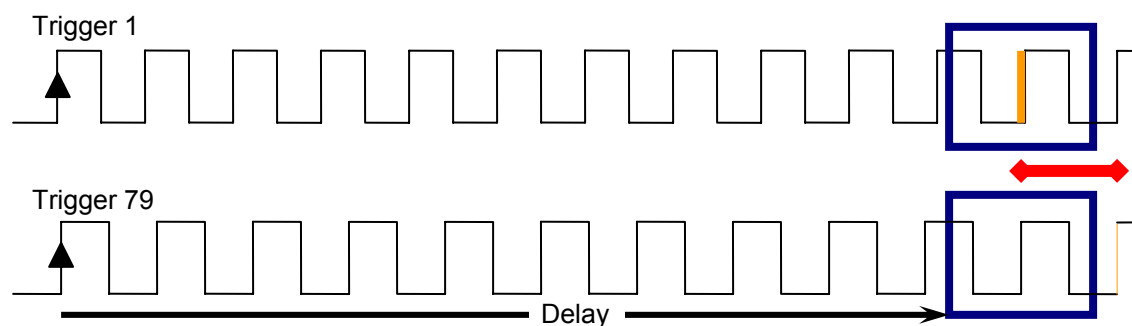

Figure 1 - Accumulated Jitter

In Figure 1 the apparent jitter is the width of the vertical orange line, yet the true jitter is the length of the red line.

When measuring accumulated jitter, it is absolutely necessary to ensure a constant number of edges so that jitter greater than 1UI is properly measured. An alterative method of measuring accumulated jitter is to measure the signal relative to reference signal. This allows a direct view of the relationships between the reference and the signal; both over a cycle-to-cycle basis or over the entire measurement period.

How Does One Determine if Jitter Affects Them?

There is no absolute answer to this question. Generally, jitter will affect any signal. The better question is "How will the signal be affected?"

The answer lies in how the signal is used. In most low speed systems, typical jitter won't have any effect. In high speed systems, jitter will generally degrade performance. Whether performance degrades (measured in error ratio – valid bits versus invalid bits transmitted), or an error simply stops the system, depends on system design. Most systems today can accommodate a small number of errors.

Jitter is usually viewed two ways: transmitted jitter and tolerance to jitter. We can measure jitter coming from a circuit; this is the transmitted jitter and what most jitter measurement is for. To measure circuit tolerance for jitter, a known amount of jitter is injected into a receiver, then increased until the receiver fails or hits a specified error ratio measured in bit error rate or frame error rate. To know how much jitter is injected at the failure threshold, the injection source is measured just like the transmitter jitter is measured.

Most standards dictate how much transmit jitter is allowed, and how much jitter a receiver must tolerate. Measurements are needed for both tests.

How Does One Measure any Jitter Found?

Jitter can be measured using a variety of available tools:

- Bit Error Rate Testers (BERT)
- Jitter Analyzers
- Counter Timers aka SIA/TIA
- Spectrum Analyzers
- Sampling Oscilloscopes
- Real-Time Oscilloscopes

▸ BERT

Bit error rate testers are in a class by themselves. They include both source and measurement hardware. The source and capture pieces are usually called a Pattern Generator and Pattern Detector. The pattern generator can source signals at many data rates, from old serial standards like RS-232 at kilo-baud rates, to hundreds of Gigabit per second PRBS 2^{51} (pseudo-random bit sequences with lengths in excess of 2,251,799,813,685,249 bits).

The BERT can send a packet of data, and compare it to what is received at the detector. Any errors are counted and the ratio of correct to incorrect bits is recorded. The value of this method is that it is a 100% test of the signal path and bits transmitted through the circuit being tested. The BERT pattern detector can be used alone to monitor signals coming from a transmitter under test – if the pattern is known in advance and a synchronization signal is available. In some BERTs, the synchronization can be via encoded header data.

BERTs are available in a variety of configurations: single-channel high speed models; and parallel use multi-channel models, usually with lower frequency ratings.

BERTs have very good timebases and accurate timing circuits with the pattern generators and detectors. This affords them the ability to measure very small amounts of jitter with reasonable accuracy. Jitter is measured by moving the sample point in fixed steps and detecting whether the signal is above or below a threshold. Measuring many thousands of edges this way builds a substantial sample set on which statistics are used to determine the standard deviation and peak to peak changes in edge timing.

BERTs tend to have large price tags, and test times are typically on the order of dozens of minutes for scans, and hours for true measurements. While they are very good at doing the specific task of testing BER, they do not offer the versatility of a more generalized test instrument, such as a real-time oscilloscope.

Useful between DC and 100Gb/s.

▶ Jitter Analyzer

Jitter analyzers are dedicated instruments or systems designed to measure timing changes in a circuit. Generally targeted at very high speed components for specific communications standards, jitter analyzers are very expensive tools and have limited functionality outside the frequency range and standards they were designed for.

There's really no typical design or measurement method for jitter analyzers. Some use phase analysis similar to spectrum analyzers, some use bit counters similar to BERTs. The design is usually dictated by which standard is being measured: some standards require measuring system wander (jitter below 10 Hertz) while monitoring 622Mb/s (mega-bit-per-second) or 40Gb/s traffic.

Useful between DC and 50Gb/s.

▶ Counter Timers

- **Also known as: Time Interval Analyzer, Signal Integrity Analyzer**

Counter timers measure jitter using a statistical approach. By measuring hundreds, thousands or even millions of signal periods – or the time between the signal and some reference – counter timers can measure jitter with a fair amount of accuracy, depending on the quality of the timebase and the resolution of the counters used.

Common instruments use timers and counters with resolutions down to 200fs (femto-seconds). This is equivalent to a counter running at 5THz (tera-Hertz). These timers don't use real counters at these frequencies. They use alternative methods. One uses a ramp generator based on a current source into a capacitor. A signal starts the current flowing into the capacitor; another signal stops the current flow. The instrument then measures the voltage across the capacitor with a voltage converter (ADC) and each bit of the ADC equals a 200fs time change.

Counter-timers under-sample the signal and can't measure cycle-to-cycle period or phase change over time. This is a serious limitation because the jitter frequency can't be determined directly and must be estimated using a mathematical autocorrelation technique. The derived frequency contains no phase information; harmonic relationships are unknown and jitter amplitudes are often misread.

Under-sampling also causes aliasing. Higher frequency jitter can manifest as lower frequency jitter and lead to incorrect assumptions about jitter sources. While aliasing can be a problem with any sampled system, having the problem start below 1MHz makes it more serious than when it starts at 20GHz (Nyquist for current 40GSa/s oscilloscopes).

Useful between DC and 4.5Gb/s.

▶ Spectrum Analyzers

A sub-class of spectrum analyzers is Phase Noise Analyzers.

Until the recent advent of real-time spectrum analyzers, spectrum analyzers used swept sinewaves and frequency mixing to deliver very good frequency range and sensitivity. This non-real-time spectrum analyzer is also known as a swept mode spectrum analyzer.

Combined with exceptional frequency sources used to mix down the measured signal, this allows them to be used for measuring minute changes in a signal's phase. After measuring the signal phase deviation, a direct conversion to time is possible: a signal that changes phase by 1 degree has a 1/360 change in period.

When jitter is measured with a spectrum analyzer, it is termed phase noise and is measured relative to the measured signal in dBc. The signal under test is the fundamental or carrier frequency, and phase noise results are normalized to a unit Hertz value based on several aspects of the spectrum analyzer measurement system: filter bandwidths, sweep rates, etc.

An advantage of spectrum analyzers is that they can have very good phase measurement resolution and themselves have very good noise floors. This provides a platform capable of measuring frequency and phase change well below what is possible in conventional time-domain instrumentation. For example, several common spectrum analyzers have noise floors below -105dB at 10kHz offset. This is compared to a current real-time oscilloscope which comes in with -50dB at 10kHz offset.

Spectrum analyzers can be quite effective at finding small sources of phase noise, and can be useful in understanding different types of intentional modulation. To their deficit, swept mode spectrum analyzers can't measure cycle-to-cycle changes in a signal. Swept mode spectrum analyzers can't make measurement of data signals. Real-time spectrum analyzers can measure cycle-to-cycle changes, and like their older siblings are quite useful for measuring small phase variations and providing detailed views of modulation.

Useful between DC and 500GHz.

▶ Real-time Oscilloscopes

This class of instrumentation provides a wonderful way to look at jitter.

A real-time oscilloscope captures a time-sequential array of voltage samples. These samples provide a coherent record of voltage variation from which edge timing information is derived.

From the edge timing information array, all sorts of data is available. The period of any cycle is measurable; the difference between adjacent periods is measurable; the mean of all the captured periods is measurable; etc. Using the mean period and the phase of the first and last cycles, the error of each individual edge is measurable – this result is time interval error (TIE). With suitable compute power, post processing this edge data results in many measurements that are impossible with the other instruments mentioned above.

A real-time oscilloscope provides edge information with a time-stamp of exactly when the edge occurred thus providing phase information that can be used to monitor harmonic effects buried deep within the sample set. This also allows real-time measurement of non-clocklike signals: real data transmitted by a circuit. This real-time approach allows detailed time domain and frequency domain analysis of clock and data signal jitter.

With time, voltage, and phase information intact, real-time instrumentation can determine jitter amplitude and frequency complete with harmonic relationships. This information can then be used to resolve down to individual jitter components, and even greatly assist determining jitter sources. This uniquely positions real-time instruments as the best jitter debugging tool available, and provides a platform for the best jitter analysis tools available.

The only shortcoming of real-time instrumentation is not having adequate sample rate and bandwidth to capture the highest data rates that exist today.

Useful between DC and 10Gb/s.

▸ What are Period, Cycle-to-Cycle, N-Cycle and TIE Jitter?

Period Jitter

Period jitter is the measurement of a signal, generally a repeating signal, from one edge to another similar repeating edge. Common period measurement tools measure the rising edge of a signal to the next rising edge. With some data transmission methods, like DDR memory, both the rising and falling edges are used to clock data bits. In this case, the measurement becomes a half period measurement. After taking a significant sample of period measurements, the standard deviation and peak values can be resolved: this statistical information is the period "jitter" in the signal.

Cycle-to-Cycle Jitter

Cycle-to-cycle jitter is the application of simple arithmetic to the period measurements just taken. If timing information for two adjacent periods is known, the difference between them is the cycle-to-cycle change: period one minus period two. Again, after taking a significant sample of periods, and measuring the difference between them, the standard deviation and peak values can be resolved. This statistical information is the cycle-to-cycle "jitter".

N-Cycle and N-Period Jitter

N-Cycle jitter extends cycle-to-cycle jitter. Instead of measuring the difference between adjacent periods, a multiple grouping of N periods is measured. Where N = 3, the width of three periods is measured. N can be any value from 2 to n where only the available sample population limits n. Once again, after taking a significant sample of periods and measuring the differences in the multiple groupings, the standard deviation and peak values can be resolved. This statistical information is the N-Cycle "jitter".

There are two permutations of N-Cycle jitter. The first reports statistics from the difference between adjacent groupings (best called N-Cycle jitter). The second reports statistics from the variation of the multiple periods (best called N-Period jitter).

Time Interval Error (TIE)

The time interval error measurement differs substantially from the period based measurements described above. TIE requires measuring a signal against a reference clock. In sampling oscilloscopes, and in equivalent time modes of real-time oscilloscopes, the reference clock is generally a clock recovered from the signal under test. A circuit monitors the signal and locks to the edges it finds. This can be a PLL, a phase interpolator, an over-sampling circuit, or even calculated from a real-time record of edge crossings.

The measurement tool measures the time difference between the reference edge and the signal edge. This method works with both clock and data signals. The measured difference is the time interval error. After measuring a significant sample of time intervals and their related timing errors, the standard deviation and peak values can be resolved. This statistical information is the TIE "jitter". In cases where the reference is a calculated linear value, the mean TIE should be zero and the standard deviation and peak values are of interest. In cases where the reference is calculated to emulate a PLL or is an electrically recovered clock, the mean value is rarely zero, but should be close to zero. Note: With a PLL, the mean result shows the relative tracking error and is only a gauge of the recovery method. Tracking error doesn't generally affect the jitter amplitude in a significant way, but on occasion it may. Thus a very large value relative to the TIE standard deviation may give cause for investigation – perhaps the PLL has had to track a large frequency shift and has a significant overshoot - a condition that may yield better results if manually reconfigured.

TIE is very useful, especially in real-time instruments, TIE maintains a record of error versus time. From this record, accumulated phase error measurements become possible. From this the frequency and magnitude of jitter components are easily found using simple frequency domain transforms like the FFT.

▸ What is Phase Noise?

Phase noise is just jitter by a different name and different terms. Consider a sinusoidal clock signal at 1GHz. The period of the signal is 1ns. If the signal has 0.1UI of peak jitter (10% of one unit interval), the phase error is calculated by determining the phase angle of subtended by 0.1UI. If one rotation of phase is 360 degrees or 2pi radians, then the phase error of the signal is 36 degrees, peak or 0.1 * 2pi r = 0.628r. That's the simple view.

In many cases, phase noise isn't measured as a simple broadband effect shown in the above example. Rather, phase noise is a measurement of jitter at a specified frequency away from carrier. In this case, the carrier is the signal under test, and the phase noise is the relative amplitude of noise integrated into a point some offset from carrier. This type of phase noise measurement results in Phase Noise in dBc√Hz units. This gets rather complicated as we begin to consider the resolution bandwidth, the offset, the final frequency resolution, the instrumentation noise floor, among other things that can compromise any jitter measurement. Some spectrum analyzers and jitter measurement instruments can convert phase noise from dBc√Hz into time (second). TDSJIT3 v2 offers a similar conversion.

▸ What are the Limiting Factors for Making Quality Jitter Measurements?

- Jitter noise floor (JNF)
- Jitter accuracy (DTA)
- Jitter correlation

Jitter noise floor and accuracy topics are covered in depth in a white paper on the tektronix website, www.tektronix.com/jitter . To summarize, JNF typifies the level of performance of the combined voltage and time measurement components within an oscilloscope.

Jitter accuracy is directly related to delta-time-accuracy (DTA). The product manuals have the actual specification – generally a formula based on sample rate, signal risetimes, and timebase stability.

Jitter correlation is a very complex topic, just like the fundamentals of jitter itself. There are several factors of correlation that are important. The first is how well does the result match any time standard. This is easily answered with the delta-time-accuracy specification which is based on calibration to NIST standards. Thus the acquisition system can measure timing and deliver results that are traceable to NIST. This is good for basic timing measurements like PERIOD, WIDTH, and even TIE.

But with measurements like Rj and Dj, along with the estimations for Tj @ BER, there are no NIST standards.

What can be used as a reference for Tj @BER is a BERT: used as a reference because it can actually measure Tj @ BER. Since most BERT's have noise floors much better than real-time oscilloscopes, this makes sense as a viable transfer standard. Several correlation studies have been performed, comparing the Tektronix TDSJIT3 real-time scope solution to a commercially produced BERT. These studies have shown extremely good correlation between the two measurement methods. Fig.1 below shows some of these test results. This shows true BERT measurements at several BER levels (red dots). The blue lines are TDSJIT3 Tj @ BER estimates – the actual BER curves or bathtub curves from TDSJIT3.

Correlation between the BERT and other test methods have shown substantial divergence in test results.
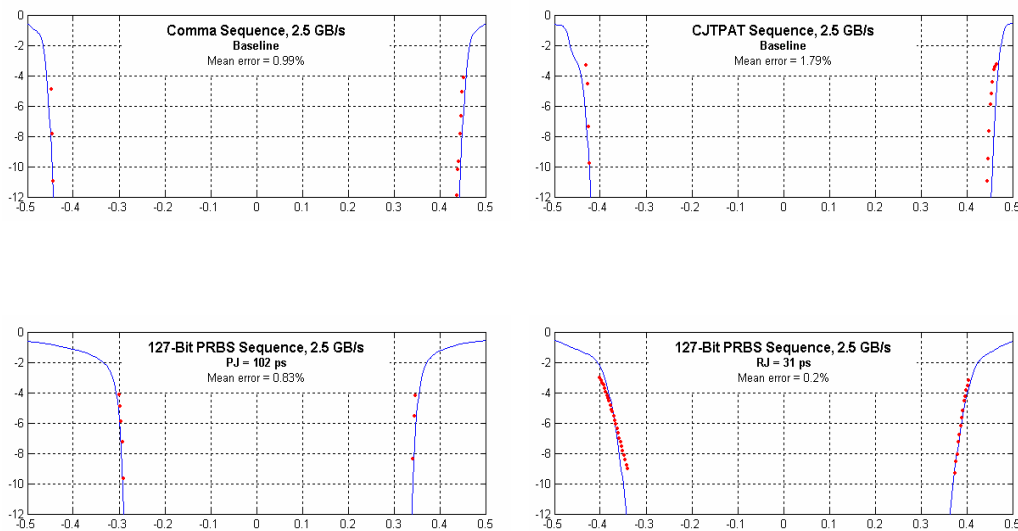
▶ TDSJIT3 – More Factory Correlation

Figure 2

‣ How are Scopes Used to Measure Jitter?

There are two basic oscilloscope methods available today to measure jitter: real-time methods and equivalent-time methods.

Equivalent time methods rely on placing multiple edges of a repeating pattern of data edges relative to a reference or clocking signal. This is essentially what we call an eye-diagram. The results of an eye-diagram correlate to the real-time TIE measurement as shown in Figure 3.

▶ Time Interval Error
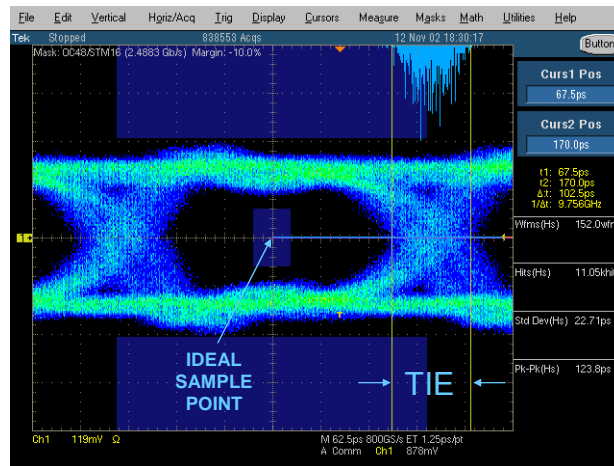
Figure 3

By overlaying multiple edges, an eye diagram or edge envelope can be established and measured using cursors, histograms or built-in measurements. The shortcoming of this method is that is masks any jitter that is lower in frequency than the clock recovery bandwidth. A PLL is the most common from of establishing a timing reference signal. This is true of BERT's, sampling oscilloscopes, real-time oscilloscopes and other TIA methods that draw eye-diagrams.

Real-time oscilloscopes can overcome this shortcoming two ways. The first method is to recover a linear time reference based on the captured waveform, which eliminates any masking of lower frequency noise. This method can also reproduce a model of a PLL and allows direct correlation to sampling methods. This first method is a clock to edge measurement. The second method directly measures an edge relative to the prior edge. From this data the signal period and cycle-to-cycle effects are easily derived. This method is known as a direct edge-to-edge measurement.

TDSJIT3 allows both clock-to-edge and and edge-to-edge measurement types: TIE and PERIOD are examples of these types of measurements. While TDSRTE allows both measurement types, it goes further and allows plots of an eye-diagram based on the recovered clock used for the TIE measurement. This eye plot allows for meaningful measurements of both time and voltage details relative to how a receiver will "see" the signal.

TDSJIT3 focuses on analysis and debugging applications for both parallel and serial data communications, though it is useful for some compliance tasks.

TDSRTE focuses on compliance applications for serial data communications, though it has some moderately useful analysis capabilities.

▶ What's new in TDSJIT3 v2.0

TDSJIT3 began life in 2002. It introduced Rj/Dj analysis to real-time oscilloscopes. This was a major introduction of capability and dramatically enhanced the ability to make jitter measurements using Tektronix real-time scopes. At the same time that Rj/Dj was introduced, the entire user interface was

revamped to make the software easier to use. Similar improvements in usability have been made in each subsequent TDSJIT3 release.

Here are some highlights from the latest release of TDSJIT3 (v2):

- v2 includes an application wizard: this allows novice users to quickly setup and make the most common jitter measurements. Note: the wizard is limited to setting up only one measurement.

- v2 includes reference level autoset algorithms: this eliminates the often frustrating need to always reset the reference levels prior to measurements. The user can override the autoset functions and manually set desired thresholds. An indication appears on the configuration panel to note when the reference levels will be updated: generally, after any application or oscilloscope settings have changed the reference levels will be updated.

- v2 includes an arbitrary Rj/Dj analysis module: some users are required to measure jitter on circuits and systems that do not support repeating traffic. Examples are serial scramblers, many live serial links, and cases of very long repeating pseudo random patterns. The arbitrary method allows measurement in these cases. While the repeating pattern method is faster and generally the most accurate method, the arbitrary method will measure in the fringe random data applications.

- v2 includes improved vertical and horizontal autoset algorithms: this allows repeatable application of the autoset function, with better optimization of horizontal and vertical scaling relative to the signal and expected measurements.

- v2 includes both high-pass and low-pass filters for most measurements: this important capability is essential for viewing small modulation effects like spread-spectrum clock profiles. It also is useful for eliminating unwanted noise in results: for example standards that only required measurements with particular frequency bands.

- v2 includes equivalent Rj/Dj: this is a dual-Dirac implementation of estimating Tj @ BER by using curve fitting to two or more higher but measured Tj @ BER values to estimate a lower Tj @ BER. For example, using measurements of Tj @ BER at 10^{-4} and 10^{-6} to estimate Tj @ 10^{-12} performance and to estimate Rj and Dj numbers that are normalized to the Tj BER curve. Note: Results of using this method do not correlate well with actual Rj and Dj measurements; to avoid confusion with the other more accurate methods used in TDSJIT3, this feature must be enabled via the TDSJIT3 command line interface. This feature is available for users who must correlate to measurements made for specific standards. For example, PCI-Sig is evaluating a similar technique for inclusion into the next generation of it's standards. Having this capability ready to go, provides an advantage for users.

- v2 includes a jitter transfer function plot: this eliminates the need to use external tools to determine jitter changes between two points in a system. Commonly used for testing PLL performance by measuring the jitter on the input or reference clock compared to the jitter on the PLL output. The plot essentialy shows jitter transfer in dB by frequency. With proper setup of the reference clock, this plot is useful to determine PLL loop bandwidth.

- v2 includes a phase noise plot: phase noise is a common method of quantifying jitter on clocks. Phase noise is a measurement that shows jitter amplitude in dB relative to the carrier, or dBc, and is usually normalized to a resolution bandwidth of 1 Hertz, dBc√Hz. The conversion from time to dBc√Hz is complex so having this plot available makes it easier to correlate measurements with different types of equipment (e.g. spectrum analyzer). The plot configuration panel has the added ability to measure phase noise in the more friendly time domain, with settable integration limits so that phase noise between two frequencies is measurable. This is accomplished without using filters and may ease some applications.

- v2 includes the ability to relocate the entire application to a second monitor or even run from a remote PC to take advantage of improvements in CPU performance over time and to allow a better viewing experience. With this comes the ability to resize plots and move the main window off the primary oscilloscope display (undock).

- v2 contains many other minor improvements: among the most mentionable are better measurement throughput, better plotting performance, and better GPIB controllability.

Contact Tektronix:

ASEAN / Australasia / Pakistan (65) 6356 3900
Austria +41 52 675 3777
Balkan, Israel, South Africa and other ISE Countries +41 52 675 3777
Belgium 07 81 60166
Brazil & South America 55 (11) 3741-8360
Canada 1 (800) 661-5625
Central Europe & Greece +41 52 675 3777
Central East Europe, Ukraine and Baltics +41 52 675 3777
Denmark +45 80 88 1401
Finland +41 52 675 3777
France & North Africa +33 (0) 1 69 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) 56666-333
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) 528-5299
Russia, & CIS 7 095 775 1064
South Africa +27 11 254 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
USA (Export Sales) 1 (503) 627-1916
For other areas contact Tektronix, Inc. at: 1 (503) 627-7111
Updated April 6, 2005

09/05   61W-18897-1   JH/WOW

Tektronix®

Enabling Innovation