NetSeminar Q&A for Signal Integrity Series - Jitter Measurements for High-Speed Digital Transmission (April 29, 2003)

The following questions pertain to general jitter measurements:

Q: How do mask shapes and dimensions correlate to bit-error-ratio measurements?

A: My understanding is there is not a direct relationship between a mask shape and the bit-error-ratio that you'll achieve. Perhaps the most fundamental reason is that we're only characterizing the transmit side with a mask shape and a bit-error-ratio is a system-level measurement. So you wouldn't be able to directly predict how a system is going to behave because you're only looking at the transmit side. The design for the mask shape is determined when a standard is designed, and the standard includes the entire communications system. So the receiver specs are also known and it should have some influence on the mask shape. In my experience with standards development, I've yet to see where there was a direct correlation between bit-error-ratio and a mask shape. So I think it's just not there.

Q: Which is better for recovering a clock for a trigger, a wide loop bandwidth or narrow loop bandwidth?

A: It depends on how well you'd like to know what the jitter performance is. The wide loop bandwidth effectively acts as a jitter filter, removing the lower frequency jitter from the waveform measurement, essentially the jitter that's within the loop bandwidth. Sometimes that's desired and some standards even require it. Other times if you want to know what's the complete, whole spectrum jitter on that waveform, you want to have a narrow loop bandwidth. It turns out that in the clock recovery schemes that we've developed for the 86100 DCA we have two loop bandwidths. So that if you want to go to a narrow loop bandwidth, you can see the full jitter spectrum. If you'd like to remove the low frequency jitter, you use a different setting.

Another way to look at the method you want to choose is to consider what the hardware in the system uses. What type of PLL clock does the receiver have? Does it have a specified loop bandwidth? What you want to do is emulate your clock recovery and your measurement system, making it emulate what is actually in the system. So if the system receiver is able to track out lower frequency jitter, you may want your waveform measurement system to do exactly the same thing.

Q: What process is used to remove the noise from the clock below the 4 MHz range of the clock recovery circuit?

A: When we say noise, I'll assume that it's jitter noise. The phase-lock loop that the clock recovery system is built around has a Voltage Controlled Oscillator, a mixer, and a filter. The mixer produces a difference signal proportional to the frequency difference between the incoming data and the VCO. That difference signal or error signal is proportional to how far off the VCO is from the incoming data rate. The error signal drives the VCO to lock to the incoming data. There's a low-pass filter between the mixer output and the VCO to have a stable system. Essentially the bandwidth of that low-pass filter sets whether it's wide loop or narrow loop bandwidth. This is what dictates how the jitter noise is limited. If the jitter is very fast, it will be outside of the filter bandwidth and the VCO will not be able to follow the changing rate. Since the jitter is not on the VCO, and the VCO is triggering the oscilloscope, the jitter will not be common-moded out of the data being displayed on the oscilloscope.

Q: If I am using a recovered clock to trigger the oscilloscope, is the waveform display wrong?

A: It's not to say whether it's wrong or right, we just have to say it's different. We have to realize that some of the jitter spectrum has been removed. Sometimes that's desirable, sometimes it's not.

Q: Is the software recovered clock similar in principle to a hardware CDR?

A: Yes. The software recovery clock actually could be any mathematical algorithm that we choose to use to recover an ideal clock. We typically choose a software algorithm that emulates the characteristics of a hardware realizable clock data recovery circuit with a single pole roll off of a specified 3 dB bandwidth.

Q: Is there a limit to the magnitude of jitter that can be seen with an oscilloscope?
A: Yes and no. When we're looking at an eye-diagram with an 86100 DCA, as soon as you have one full bit period of jitter, the eye-diagram closes and beyond that you can't tell whether the jitter is beyond a bit period a little bit or a lot. So in the eye-diagram, you're limited to one unit interval or bit period of jitter. With the real time oscilloscope, we can see beyond one unit interval when we measure a very long record length.

Q: What is wander and how is it defined?
A: Wander and jitter are both descriptions of timing instability. The typical rule of thumb is that wander describes jitter that's very slow. In the telecommunications world they say variations that are less than 10 Hz in rate are wander. Anything above 10 Hz would be jitter.

Q: Can you specify jitter in a unit interval when it's over a frequency range?
A: Jjitter can be expressed as a magnitude when it's measured in the frequency domain. It's not too hard if the jitter is a sinusoid. The modulation index can be used to convert to unit intervals, but I can't remember the formula, but it's something like beta over 2 pi.

Q: To reduce the jitter contribution of the oscilloscope, can I just increase the vertical sensitivity setting in order to increase the effective slew rate of the captured signal relative to the peak-to-peak deflection of the data signal?
A: Yes. It's true that the faster the slew rate is the less contribution that vertical noise will have to the jitter measurement. In the case of sampling oscilloscopes like the 86100 DCA, you can increase the slew rate as much as you want as long as your signal does not violate the maximum input range that's specified for the oscilloscope. You can do that even if your signal is clipped vertically on the screen, if it exceeds the screen by even a couple of screen diameters, that's usually okay in a sampling oscilloscope. However, in a real-time oscilloscope, you cannot do that. You need to be careful to make sure that your signal is as large as possible but it should not clip the A/D converter that is used to digitize the signal, or exceed the A/D converter's range if the signal does not exceed what's displayed on the oscilloscope.

Q: How does vertical noise impact timing jitter?
A: This is an interesting thing. If you take it from the perspective of a measurement that as we look at a rising edge with some slope to it on an oscilloscope waveform, there's going to be some thickness to it. Is that thickness due to jitter or is it due to noise? It turns that there's really no adequate way looking at just the edge, to know for sure whether the trace thickness is due to noise or jitter. As you get to a flat region of the waveform where the amplitude is generally constant, any trace thickness is likely due to noise. However, if the flat region is quite thin, then you know on the edge the trace thickness was probably jitter. So to answer the original question, vertical noise can degrade a jitter measurement. Averaging can be used to eliminate noise, but unfortunately it will also eliminate random jitter too.

Q: The oscilloscope's time base stability must be very important for jitter measurements. How good is the time base stability of most digital oscilloscopes today, and how good are your time bases?
A: At first thought, the time base stability of the oscilloscope seems very important. It turns out that almost all the oscilloscopes on the market today are very accurate and that the stability of their time base typically does not enter into jitter measurements. It can in some cases, but predominantly, the limitations to jitter have more to do with trigger jitter and the vertical signal-to-noise ratio. As for how good are they? For fast slew rate signals, signals that are close to the bandwidth of the oscilloscope, or the rise time of the oscilloscope, most of the high bandwidth real-time oscilloscopes can measure jitter down to the less than one picosecond RMS range.

Q: How fast does the oscilloscope sampling need to be relative to data and edge rate of the signal in order to give accurate measurements of jitter?
A: For real time acquisition, sample rate is very important to accurately capture high-speed signals and jitter. And the bandwidth of the oscilloscope and its maximum sample rate are directly related. Many of the signals we showed in this seminar are examples for 2.5 Gb/s data signals. The highest bandwidth real-time oscilloscope on the market today is 6 GHz with a sample rate of 20 GSa/s. With Agilent oscilloscopes you can capture signals on all four channels with no tradeoff in sampling. With 20 GSa/s, you can reliably capture signals with data rates up to about 3.2 Gb/s. So that's pretty much the limitation today for real-time oscilloscopes.

Q: I have a 500 Mb/s system. Can I get accurate jitter measurements with a 5 GSa/s sampling rate oscilloscope in either real or equivalent time?

A: The question is one of accuracy -- how accurate do you want to be? In the equivalent time/sampling oscilloscope the amount of data you get is probably not adequate to measure to 10^{-12} BERT level. In the real-time oscilloscope it would require a lot more data and we would potentially use extrapolation for those measurements. Extrapolations usually work, but there are times when the BERT lets you know for sure where you're at.

If the question was is there enough sample rate for the bandwidth involved to accurately digitize the bit stream? In a real-time oscilloscope, sample rate is important because you cannot violate Nyquist. For that particular case, 5 GSa/s would be enough sample rate for a 500 Mb/s system. In terms of a equivalent time oscilloscope, sample rate really does not apply. Because it's equivalent time, the bandwidth of the oscilloscope can be independent of the sample rate, so you have to go primarily by its bandwidth.

Q: What is the sampling rate requirement of real-time oscilloscopes versus the data rate of the system under test?

A: There is a direct relationship between bandwidth, sample rate, and data rate. To give you an example, at 2.5 Gb/s, assume you have a 1, 0, 1, 0… repeating pattern. Since we typically work with NRZ-type signals today, the highest fundamental frequency of this signal would be half of this data rate, or 1.25 GHz. We recommend that the sample rate be at least 3 to 4 times the highest fundamental frequency of this signal. So for a 2.5Gb/s data rate, which translates into a maximum fundamental signal frequency of 1.25 GHz, we would recommend at least 5 GHz bandwidth of your oscilloscope, and even higher bandwidth would produce better measurement results. Nyquist says that the sample rate must be twice the highest signal frequency. In practicality, you need more than twice to eliminate aliasing caused by signal frequency components beyond the Nyquist rate, which would be present on signals with very fast rise and fall times. Agilent's newest real-time oscilloscopes sample at a maximum rate of 20 GSa/s on all four channels. At 20 GSa/s on all four channels you get a 6 GHz system bandwidth for each channel, which is sufficient for data rates up to about 3.3 Gb/s. There is archived Agilent seminar available on Netseminar.com that talks about bandwidth, sample rate, and exactly what to choose for different applications.

Q: You talk about oscilloscopes and bit-error-ratio tests, but what about timing analyzers? When should you use those?

A: A timing analyzer is another technique to verify jitter. It's not to say that different techniques are not valid for making jitter measurements. We've tended to look at the oscilloscope and the bit-error-ratio tester. An oscilloscope is a very effective tool to see the waveform and find the root cause of jitter. The BERT allows you to work with precision down to very low levels. Timing interval analysis is another valid way to characterize jitter.

Q: Why is there so much disagreement in extrapolating random and deterministic jitter in a waveform?

A: Partly it's because people have different models and algorithms in use. There is not a standard for how these things are extracted, so how the algorithm's used, on what hardware, etc. is not consistent. It's something that I think that the community as a whole would like to get some resolution on.

Q: If you're looking at a single edge on say a 1010 data pattern, how good is the RJ measurement you get by taking the histogram, 1 sigma x 14? I know it can include parametric jitter effects.

A: I don't know what you mean by "measuring RJ on a single edge". What are you measuring the jitter on that edge with respect to; the oscilloscope trigger, another edge? If your idea is to use an alternating 1, 0, 1, 0 pattern to remove the inter-symbol interference component of jitter and so that the rest of the jitter is all random, seems to be very risky. You still have contributions from asynchronous periodic jitter, which can be pretty large, and plus contributions from duty cycle distortion, which could also be very large. If you just take a standard deviation measurement of that jitter or the RMS value, multiply by 14, assuming that it's Gaussian, and that gives you a 10^{-12} error rate. That seems like a pretty risky thing to do.

Q: When you say to apply averaging to the signal to remove random jitter, are you referring to the data waveform or the TIE waveform?

A: I was referring to averaging the TIE waveform. However, you can also use signal averaging on the data waveform to eliminate some of the random jitter components, but this will probably produce less accurate results. When you average the data waveform, you are actually averaging out the vertical jitter in the signal, which will ultimately translate into timing jitter. But what I was referring to during this presentation is averaging the TIE jitter trend waveform using one of the waveform math functions in the oscilloscope. This will average out the random jitter components, which is a vertical component in the TIE trend waveform.

Q: Since the TIE is not a continuous signal, but rather a set of points representing time deviation data taken at the sample times, why is it displayed as points with the only important line showing the TIE equal to a zero value rather than a continuous line?

A: It's a user preference as to how you would like to see the data plotted or displayed. We investigated several different algorithms when we developed this for the real-time oscilloscope and the display that we chose takes each real discrete value and extends it on to the next value so that it's easier to see. As you zoom in and out of waveforms it's easy to see the trend in the jitter or the time error with this plotting technique as it moves through the waveform.

Q: Is there software available for 86100 DCA to do TIE analysis using a frame trigger?

A: No. We don't have any software for the 86100 DCA at this time. It would be interesting to know your application if you'd willing to contact us.

Q: What is the recommended method to use for less than 1 ps RMS in jitter?

A: I know that in the equivalent time sampling oscilloscopes that we can actually get the oscilloscope jitter down into the 100-150 femtoseconds range, which allows you to get an accurate measurement of signals in the 200-300 femtoseconds range. You can get a very accurate measurement because the oscilloscope jitter is so low.

As for using a real-time oscilloscope, the smallest amount of jitter that you can measure depends on what jitter you're trying to measure. Are you measuring period jitter or time interval error or cycle-to-cycle? It also depends on the slew rate of your input signal and some other things. In the case of our 6 GHz bandwidth real-time oscilloscopes, the jitter measurement floor is on the order of 0.7 picosecond.

Q: For analyzing jitter with real-time oscilloscope techniques, isn't more acquisition memory always better?

A: Yes and no. If you're trying to capture very low frequency phenomenon such as uncorrelated periodic jitter, the deeper memory you have, the better. However with clock recovery, often some of this low frequency uncorrelated jitter that may be coupling in is going to be filtered out by the PLL clock recovery approach in the system. So, you don't always have to capture more and more data in order to get something you really don't care about anyway. But there is a minimum that you do need. For instance, the example we showed of capturing and measuring the spread spectrum clock had a frequency of 33 kHz. When you're sampling at the maximum rate of 20 GSa/s, you need enough memory to capture at least one period of this modulation. For a 33 kHz spread spectrum clock, you will need at least 700 kB of high-speed acquisition memory. In addition, if you're trying to capture low-frequency periodic jitter such as interference from a switching power supply operating at 100 kHz, then the oscilloscope will need sufficient memory to capture this signal at the full sample rate.

Q: How deep are your oscilloscopes' memory depth and what would be the largest PRBS pattern you could effectively evaluate?

A: Our newest 6 GHz real-time oscilloscopes have 1 MB of memory when sampling at the maximum sample rate, which is 20 GSa/s. Memory can be extended up to 32 MB of acquisition memory if the oscilloscope is sampling at a lower sample rate of 2 GSa/s and less. But if you are trying to capture and analyze high-speed data signals in the multi-gigahertz range, you're oscilloscope will probably be operating at the higher end, or 20 GSa/s. So in this case, you're pretty much limited to 1 MB of captured data. This translates into 50 microseconds of single-shot acquisition time with a 50 ps sample resolution. So if you know what your data rate is and the length of your PRBS pattern, you can figure out how much time range you need.

The following questions relate to bathtub curve measurements:

Q: You didn't mention separation of types of jitter, RJ, DJ, or PJ. Isn't this important for characterizing these types of data links?

A: It is very important and we'll talk a little bit about some of the jitter separation techniques. Effectively what happens when you construct the bathtub curve there are algorithms that will determine which elements of the jitter are deterministic and which are random. When you analyze the deterministic jitter, you'd like to know what are the different mechanisms, whether it is periodic jitter, correlated jitter and so on. So it is important and some standards actually specifically call out the separation. It is part of what a bathtub analysis can do with the right sort of software algorithms to derive those parameters.

Q: How long would it take to make a bathtub jitter measurement to 10^{-12} without any extrapolation?

A: There's not a direct answer because there are several parameters that you have some flexibility in defining and setting up the bathtub measurement. One is, how many errors are you going to have to log before you know the bit-error-ratio? Another is, how much time do you want to spend at any given point as you walk your way through the eye-diagram?

One example is the 2.5 Gb/s test where we can test down to the 10^{-12} level in about 20 minutes per side, so it is in this case 40 minutes. You could reduce this a little bit by reducing the number of errors you need to log, or you could increase it and get a higher precision measurement by increasing the number of errors you need to log. So you can do it in under an hour at 2.5 Gb/s. If you go to 10 Gb/s, it would be four times faster, or one-fourth the time. If you were to go to a lower rate, it's proportionally slower.

Q: Doesn't the amplitude sensitivity of the bit-error-ratio test affect bathtub measurements since you're measuring at points in the eye with very low amplitude?

A: Yes. The bit-error-ratio test method is just looking for errors and it won't know if the errors were due to amplitude or due to jitter. It turns out that when we're characterizing transmitters, we usually have quite a bit of signal strength such that the error performance that is measured is dominated by the jitter as long as we stay in the center of the eye. If we were looking at a very low signal level, then your accuracy would be affected by the amplitude.

Q: Which BERT systems are capable of producing a bathtub curve?

A: Any bit-error-ratio system that has the ability to move the sampling point is capable of producing a bit-error-ratio bathtub curve. The Agilent BERTs, the new 3G/7G/13G Serial BERTs (N4906, N4902, N4901) and the ParBERT, have it embedded within the function set, making it extremely easy to use. We also offer a free PC software package to create bathtub curves for the 86130A and 71612C BERTs.

Q: How can I get a copy of the bathtub curve software for the 86130 and 71612?

A: The quickest way is to go the http://www.agilent.com, then search for "bathtub". Under the search results for "Software, Firmware & Drivers", you'll find "Bathtub Jitter Software". Click on the link to download the software for free.

Q: Can you summarize the procedure of how to make a bathtub curve measurement? Are any reference papers available?

A: If you perform the search above for the free bathtub curve software, you'll also find Application Note 1550-12, "Using Bathtub Jitter Software", lit. no. 5988-8434EN. This note describes bathtub curve measurements in general, in addition to using the free software on the BERTs.

Q: If you mathematically integrate the area under the Gaussian curve and present the result with a logarithmic vertical axis, do you then get a bathtub curve?

A: The bathtub curve and the histogram are both looking at the same type of data. I'm not sure if you can reconstruct a bathtub curve through an analysis of that histogram. Typically with the oscilloscope histogram, your precision is going to be extremely low. You won't come close to trying to be able to see it down to the rate 10^{-8} and clearly not down to the 10^{-12} area of measurement level. It turns out that the information provided by the histogram and the information provided by the bathtub curve are similar presentations of the same type of information.

Q: Should I use a recovered clock to drive the bit-error-ratio test in a bathtub measurement?

A: This is a similar issue to what happened with the oscilloscope measurement where a recovered clock is valid to use, but you have to be aware of the issues that come with recovered clocks. Recovering a clock will basically remove the low frequency jitter, the jitter within the loop bandwidth of the clock recovery circuit and that will not be part of the bathtub measurement. It turns out that some standards call out a clock recovery system and specify the loop bandwidth, so it can be a valid, even a required technique. One thing about the Agilent clock recovery systems is that they have dual loop bandwidths. A narrow loop bandwidth if you want to see all the jitter or a wide loop bandwidth if you want to screen out the jitter that is perhaps called out by a standard.

Q: How can one convert the waveform database results of a sampling oscilloscope to a bathtub curve to allow a direct comparison to BERT bathtub curves?

A: The information for the bathtub curve is there. It's not something I think I could describe easily. When we try to do a correlation, we end up getting a bit fuzzy in those Gaussian tails. These are hard to come by with just the waveform data from an eye-diagram.

Q: Can I get a copy of the presentation?

A: Yes, go to eSeminar archive. When the video starts up a row of 10 bars will appear beneath the video. Click on the last bar to jump to the last portion of the presentation. After a little while the Agilent Resource page will be displayed, where you can get a copy of the presentation plus find many valuable links for more information on the subject. There is also a Digital discussion forum (http://www.agilent.com/find/forums) if you would like to pose/review other questions.