NetSeminar Q&A for Signal Integrity Series - Jitter Measurements in Digital Circuits (January 24, 2003)

Q: Isn't dispersion predictable and bounded?

A: Yes, you are right it is bounded. Dispersion can be modal or chromatic. Modal refers to the spreading of a pulse as light waves travel down and reflect about the fiber which results in slightly different times of arrival. This is clearly a bounded phenomena dependent on the length of the fiber, fiber material, and wavelength of the light. Chromatic dispersion refers to multi-wavelength light that travels at different speeds. Again, this would be bounded and reasonably predictable.

Q: Is the determination of N in the random jitter evaluation N times sigma, purely dependent on what your bit error rate budget is?

A: It's dependent on the bit error rate that you desire from your system. The budget is a result of that. So, most of these standards want a 10^{-12} bit error rate. You can look up the appropriate sigma to get you that error rate; the table for this is readily available in various sources. So the N that you use for that is strictly dependent on the bit error rate you want and then from there you can build your spec budget or your jitter budget.

Q: What are PRBS mechanisms?

A: PRBS are pseudo-random bit sequences -- sequences of 1s and 0s that are specifically designed with linear shift registers with XOR gates and they have special characteristics, namely equal probability of 1s and 0s and equal probability of groups of 1s and 0s. They are typically characterized by a 2 to the N minus 1 type length. These repeat in time and therefore will have periodic spectra that get more dense as the repeating interval lengthens. This allows for the possibility to confuse Deterministic Jitter as Random Jitter. In the limit, an infinite length repeating interval the spectrum would be continuous and would seemingly represent random jitter.

Q: Does your depiction of the eye diagram include consideration of NRZ encoding?

A: Yes, non-return to zero (NRZ) encoding is assumed on these.

Q: What defines the unit interval on the eye diagrams? Is there an ideal clock reference?

A: Usually the ideal unit interval is provided given the standard that you are trying to measure. The other thing I would add is that the sensor, usually the clock is placed in the center of the eye diagram and the data is placed relative to that either recovered or given clock. For real-time oscilloscopes, the unit interval can be calculated from the data waveform.

Q: How can we estimate jitter by using a spectrum analyzer?

A: The phase noise based system, which essentially is a spectrum analyzer, looks at the close-in spectrum to the, given clock and does an analysis of the jitter that way. If instead you are using a conventional spectrum analyzer, it's a little tricky to try to estimate the noise spectrum, which is what your random jitter would be. However, deterministic elements tend to standout prominently in the graphs (as spurs) but you have to remember that this is phase modulation, so you have to do a Bessel term analysis to convert those spectral heights to an equivalent jitter magnitude.

Q: How do you estimate jitter by using the noise floor measurement of a spectrum analyzer?

A: It's a tricky thing to do with the noise floor. Because there are deterministic elements and you have to make sure that you don't integrate the deterministic elements and you have to also be assured your amplitude noise isn't present. You also need to know that the spectrum analyzer performance is low enough noise.

Q: You mentioned at slide 18 the maximum value of random distribution depends on the time of observation. However, I was under the impression that it did not. I thought the probability of random processes does not depend on the time.

A: I thought the same sort of thing, the word is “depends”. It's not a function of time. As time goes on, the peak-to-peak value will just increase because lengthening your observation time allows for greater likelihood of seeing rare events. You will ultimately see the highest peaks if you wait long enough. Another way to look at it is if you trying to estimate what the Gaussian process is, you are basically trying to fit a curve to it. The more samples you have the more accurate your estimation will be. So from a measurement sense, to get more samples takes more time.

Q: Is jitter an important concern in source synchronous signaling systems? Is there a relationship between jitter and skew?

A: Two different questions here. In synchronous systems, similar to SONET and SDH, the difficulty is that the timing propagates through the system even as the system is received and retransmitted. Part of that process is actually taking the timing and recovering a clock and sending it back out on that clock. So the issues there are one that we didn't talk about, things like jitter transfer, where jitter can actually build up through the repeater process. So in a synchronous system you do have to treat jitter a little bit differently. As for skew, skew is kind of DC jitter and wander is jitter that's less than 10 Hz to DC. The official term jitter refers to elements that are greater than 10 Hz.

Q: What software exists for bit error rate testing for RJ, DJ and bathtub curves?

A: For the 86130 3.6G BERT and the 71612 12.5G BERT there is a free software package that you can run from a PC to do that extraction. The 81250 ParBERT has a built-in feature to do the bathtub jitter and the RJ/DJ extraction. For the 54850 Oscilloscopes, the Amherst Systems M1 software package can create a bathtub curve and separate RJ/DJ.

Q: How does the bathtub plot relate to the histogram of an eye diagram? It seems as if they are directly related to one another.

A: They are directly related. We talked about the tails of the Gaussian functions. The bathtub curve deals with the complementary error function or the integral of those Gaussian functions, so they are related. The histogram, if you have a lot of samples and could pick the actual curve out, would be related. As a matter of fact, on a bathtub curve you can estimate the histograms, so you are absolutely right, they are related and they are related by the Gaussian curve function.

Q: What does Agilent offer in the area of jitter tolerance testing for sinusoidal jitter above 10 MHz?

A: We have several products that do that. There is the JS-1000 Performance Jitter System, which has the ability to do sinusoidal jitter testing into the 80 MHz range for SONET/SDH compliance testing. The 71501D Jitter Analyzer also does sinusoidal jitter analysis to 80 MHz, essentially jitter tolerance and jitter transfer. And then there are the OmniBER products, which have a built-in ability to do testing according to SONET/SDH standards.

Q: Does the ParBERT system allow variable bit rates? For example, can it run between 1 and 10 Gb/s and every frequency in between?

A: The 81250 ParBERT system provides parallel BER testing up to 45 Gb/s for high-speed digital components. The modular design lets you mix and match analyzer channels, generator channels and speed classes (333 kHz – 675 MHz, 333 Mb/s - 1.65 Gb/s or 2.7 Gb/s, 21 Mb/s -3.35 Gb/s, 9.5 - 10.8 Gb/s and 38 - 45 Gb/s) to create a system that meets your needs. Many of these speed ranges are optimized for SONET/SDH or Ethernet testing. If you need to test from 1 to 10 Gb/s, you will need to use the 71612C 12.5G Serial BERT.

Q: Is the bit error rate floor below 10^{-12} depending on random or deterministic jitter or wander mechanisms?

A: It is due to both deterministic and random jitter. If you recall the bathtub curve had a flat portion, that's where the deterministic jitter dominates. Then you transition into the steep wall of the bathtub curve, where the curve precipitously declines is dominated by the random jitter component. That's supposed to be on the skirts of the Gaussian function. Where the Gaussian curve starts however is a function of the deterministic jitter. However, occasionally and it's rare, but you will get very low probability deterministic events and there have been cases where you are basically running down the edges of the bathtub curve and then people reach what's called an error floor. That's where something, again very low probability, is deterministic. So that's one of the reasons sometimes we are just estimating, that curve fit just doesn't do it. You want to know for sure that you have that bit error ratio, then you need to look that far down because occasionally there will be error floors and those are deterministic.

Q: On slide 29 bullet 4, shouldn't sigma L and sigma R add power-wise if they are different processes?

A: I thought about that myself, why are we adding 2 sigmas and then dividing by 2? Why aren't we taking the average when in fact when we normally have processes that multiply each other? We are not adding distributions at this point. The assumption is that there is one composite random process that is convolving with the deterministic processes and we are making two different measurements of that random process. Due to many reasons, amongst which is the measurement algorithm, its sensitivities, and the actual data set, the two values will always be different. If these truly reflect one process and are relatively close together it is appropriate to average them for the best estimate of the sigma of the composite random process. This is commonly stipulated in the digital bus standards (i.e. Fibre Channel).

Q: Why is the distribution of random jitter characterized as Gaussian and not uniform?
A: A uniform distribution does not describe what's really going on here. A truly random will have the characteristic of having infinite tails, although very low probability. The frequency content would be uniform but the distribution magnitude would be Gaussian.

The assumption that random jitter (RJ) follows a Gaussian distribution simply reflects the Central limit theorem in statistics theory. Basically, it states that a sum of small and uncorrelated effects follows a Gaussian distribution. The trouble with this is not that a Gaussian isn't the best assumption among possible choices (e.g., your suggestions as well as other possibilities like a Breit-Wigner or what have you) it's that the underlying effects are individually non-Gaussian. So only in the limit of an infinite data set can we really expect the RJ to follow a Gaussian. On the other extreme, with a very limited data set, we have to expect large fluctuations from the Gaussian shape and, to an even greater extreme, if the course of RJ could be narrowed down to individual processes (instead of a sum of uncorrelated small effects) then we should use those distributions. For example, shot-noise should use a Poisson distribution, and in the dramatic extreme, for thermal noise at low temperatures we could write down the wave function of the conductor and propagate individual thermal excitations into the jitter function. I advise against this because it'd probably take months and you'd end up with the same answer that you got assuming a Gaussian.

I think the point is not that some other function makes more sense than the Gaussian assumption, rather, that the Gaussian makes the most sense in the usual case where we really can't narrow down the RJ sources to things that are easier to parameterize. That said, it's possible that for a given system you might observe a particular empirical distribution for RJ that is both reproducible and not particularly Gaussian. In this (rather weird case) I'd recommend trying to parameterize a smooth function (an exponentiated polynomial can mimic almost any analytic function) to fit the empirical distribution and then use that instead of a Gaussian.

Q: Would you say that histograms are a means to quantify what is seen by eye diagram? Perhaps someone with enough experience really doesn't need histograms because they can discern this information with the eye.
A: It is interesting that with the eye diagram display, particularly when you are in a gray scale or color graded type display, you can kind of pick out what the distribution is. If there is definite deterministic jitter, you see the bright stripes in the eye diagram. So people with experience learn to pick those out. Once things get a little more subtle, and don't jump out at you, then the histogram can give you a little better insight as you look to that shape and see is it truly Gaussian. Then you have a good idea it's dominated by random, whereas if you see some little bumps and wiggles in the histogram, it may not have been visible with the naked eye but you get a little better more precise view when you actually do use the histogram.

Q: Is any software compatible with the 86100 DCA that separates RJ and DJ?
A: No. It turns out there is not anything at this time that does that. We know it's an important feature and something that we hope to come up within the near future.

Q: Does the 86100 digital communication analyzer require a periodic signal to achieve high bandwidth?
A: There is a general rule of thumb when using these wide-bandwidth sampling scopes. You have to have a repetitive signal and then a trigger that is synchronous to that. That's true if you want to see a pulse train or something like that, but for an eye diagram the data does not need to be repetitive or essentially periodic, but you do have to have a trigger that is synchronous to the data to be able to display that eye diagram. That trigger typically being a clock at the line rate.

Q: When using sampling scope with an external trigger signal, are there any requirements for the trigger signal regarding the rise time in order to reduce the jitter?

A: You wouldn't want to go slower than the edge of the sine wave. Normally people are using a digital type clock and that would then typically then have an edge speed faster than what you get with a sine wave, but the rule of thumb is that the sine wave is likely to be adequate.

Q: With regards to the histogram window, is the height set to one pixel? It sounds like we need to know the display's pixel width.

A: What happens is if you have too high of a histogram window, that you tend to have the edge speed start to influence the histogram. So you would really like to keep that slice just as narrow as possible and obviously the narrowest slice you can take would be one pixel, to not be influenced by the edge speed.

Q: Can the 86100B be used to get high bandwidth transients, random errors?

A: Because it is a sampling scope where it does not really see everything that goes by, the only way you would be able to capture something that was transient is if it happened over and over again. That sort of violates the description of what a transient is. So typically for that you would need to go to a real-time scope architecture.

Q: What method is used to perform DJ and RJ separation on your real-time oscilloscopes?

A: The real-time oscilloscope does not have the DJ and RJ built in but there are some optional add-on packages that you can employ that will do the separation of deterministic and random jitter such as Amherst Systems M1 software.

Q: Is there a better way to measure jitter or just a different way specified by the different standards?

A: It seems like the different standards all want to approach it in their own unique way. One Gigabit Ethernet and FibreChannel tend to lean on the bathtub approach. Serial ATA looks at specific durations between edges to characterize jitter. And in the SONET/SDH world, they tend to have a sinusoidal analysis. They have all taken their own approach to it because you have different people who were on the standards committees and they just felt that that was the way they should go after that. So they are unique architectures.

Q: Do you have tools for SDRAM interface jitter measurements?

A: No. Nothing is specifically targeted for that with the 54850 series Infiniium oscilloscope.

Q: How high a frequency can the ADS SI tool go to?

A: How high do you want to go? There is no fundamental limitation in ADS with respect to frequency. There are limitations to the models so frequency limitations have to be determined on a model-by-model basis. The Advanced Design System is really designed for RF, microwave, and millimeter-wave designs so engineers are finding that ADS really excels where the other tools become inaccurate. With the higher bandwidth needs for today's digital designs, ADS is an ideal tool for interconnect analysis.

Q: Is ADS a Spice based simulation tool?

A: No, it's not. ADS is a suite of simulation tools and built-in models. Compared to Spice, ADS offers a linear circuit simulation, non-linear circuit simulation, modulation envelope simulation, transient simulation, and EM simulations. Its models are optimized for frequency dependent effects like dielectric constant, skin effect, etc. While Spice is a non-linear circuit simulator, its' models do not include frequency dependent effects. You can create a Spice model to have some frequency dependence, but they will have limited frequency range. ADS can create models that can be used in Spice, but it is its own environment and it can be a circuit simulator as well as a modeler. ADS also can be integrated into IC design tools from Cadence for mixed signal analysis.

Q: Does slide 33 show an improperly terminated signal in the 30 inch picture, or is this just the result of XAUI 30 inches?

A: It's a XAUI measurement at 30 inches of the backplane that we had. So, is it improperly terminated? It's not perfectly matched, and it's very hard to get a perfectly matched signal, especially when you are on a backplane with a lot of interconnects, et cetera. So that is going to be implementation specific.

Q: Can I get a copy of the presentation?

A: Yes, go to slide 50 of the eSeminar archive. From there 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.