JitterFest: Why Different Jitter Analyzers Give Different Results -- Which Is Right?

by Ransom Stephens

If you've been measuring the jitter alphabet soup -- R_J, D_J, DD_J, ISI, DCD, TJ (BER) -- then you're already frustrated by the variation of measurement results from different jitter analyzers. What puzzled me was the way the problems show up and, with all the fuss manufacturers make over which algorithm is best, I was surprised that the biggest effect turned out to be hardware quality. Certainly some algorithms are better than others, and you want the best hardware-software combination, but voltage noise is the culprit for the most obvious discrepancy.

Take your favorite jitter analyzer -- be it a real-time oscilloscope, an equivalent-time sampling oscilloscope, a time-interval error analyzer (TIA), or a bit error ratio tester (BERT) -- apply a data signal and write down the value it reports for random jitter (R_J). Now insert a length of transmission channel, a backplane or some PCB, into the signal path and compare the value for R_J with and without it. Two of the real-time oscilloscopes that we tested reported over 500% greater R_J after we introduced thirty inches of PCB trace on a 2.5-Gbit/s signal and only one reported the same R_J; the one with the lowest voltage noise, an equivalent-time sampling oscilloscope.

The PCB trace cannot affect the actual random noise of the signal, it can only introduce data-dependent jitter (DD_J). So how come the R_J numbers increased so much? How come each jitter analyzer except one reports a completely different R_J value with the PCB? And how come the jitter analyzer manufacturer didn't tell me this would happen?

But I am getting ahead of myself. A couple of years ago, a field engineer (that's a "salesman" to you and me) dragged me to a customer who was complaining that he was getting different results from different jitter analyzers. One of the analyzers was an Agilent BERT and the other was a real-time oscilloscope made by another company -- Agilent didn't have any jitter analysis capability on their real-time oscilloscopes at the time. The BERT was performing a simple bathtub-plot extrapolation[1] and the oscilloscope was doing I-don't-know-what. The only information I could find about the oscilloscope's jitter analysis capabilities described what was measured but not how the measurements were performed.

A couple of weeks later I was on a conference call with a group of Agilent engineers from the divisions that make BERTs, real-time oscilloscopes, and equivalent-time sampling oscilloscopes. There was a lot of discussion about competing ways of measuring jitter, references to the T11 MJSQ document[2], and a consensus that, yes, different jitter analyzers gave different answers, and, no, we don't know which is right. Fig. 1(a) shows a plot of fast estimates of total jitter defined at a bit error ratio of 10^{-12}, TJ (10^{-12}), the results among the different analyzers vary by about 25%. Fig. 1(b) shows rms R_J measurements that vary up to 800%. We agreed that the easiest way to tell which analyzers were accurate was simple: apply known levels of different combinations of jitter to the different jitter analyzers and see how they do.

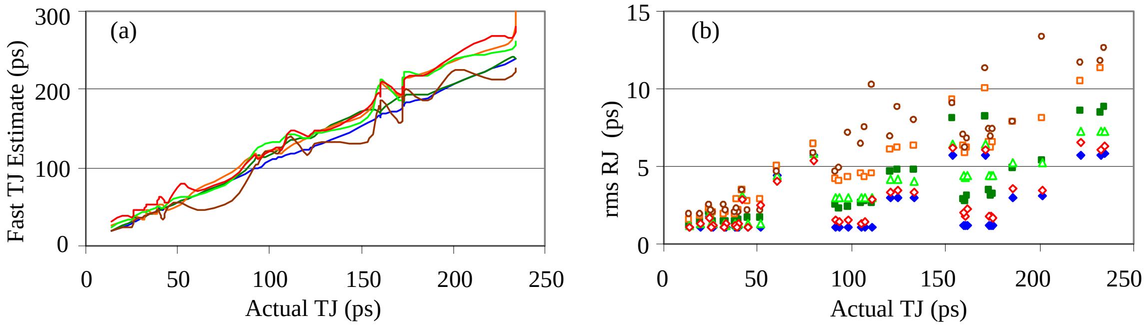

Fig. 1: (a) TJ(10^{-12}) Estimates By Different Test Sets Vs. True Value Of TJ(10^{-12}) (b) R_J Measured By Different Test Sets Vs. True Value Of TJ(10^{-12})

The next step was convincing management. But at Agilent this was easy, all we had to do was go to management and say "…Dave and Bill would let us do it, it's the HP way."[3] If that hadn't worked we could have played the trump card: around here if you mention Bill Hewlett's garage you can get almost anything (except a raise).

We assembled a team and designed a jitter transmitter, assembled jitter analyzers from all the major manufacturers, defined a set of jitter test conditions that we thought reasonably spanned the jitter space and scheduled a time and place to perform measurements. We called it JitterFest.

Our work with a simple jitter transmitter a couple of years ago taught us three things -- we needed a more accurately-calibrated jitter transmitter; TJ(BER) is important; and, yes, the jitter analyzers give different results. It took a long time to develop a truly precise jitter transmitter because we still had to do our usual jobs (management told us that "this too is the HP way").

By the time we were able to get all the equipment back together, the Agilent sampling oscilloscope team had built a new jitter analyzer, the 86100C DCA-J. With the precision jitter transmitter[4] we could tell which analyzers worked and which didn't and we were particularly interested in challenging the new guy, the DCA-J. At the time I worked mostly with BERTs -- I was the guy who originated the mantra: "you can only measure TJ on a BERT" -- and was a big skeptic of the DCA-J.

Along with the precision jitter transmitter, we assembled six different jitter analyzers, three 6-GHz bandwidth real-time oscilloscopes, a 4-GHz bandwidth time interval analyzer, an Agilent BERT (N4901B SerialBERT) and the DCA-J (Agilent 86100C) equipped with a 20-GHz electrical receiver. It was difficult to figure out how each jitter analyzer operates -- only the Agilent analyzers have complete descriptions of how they work.[5, 6] By supplementing the manufacturers' manuals and product notes with patent disclosures, we got an idea of their techniques. Two of the real-time oscilloscopes (labeled brand X and brand Z in the figures below) use algorithmic variations of the dual-Dirac model where the time interval error distribution is integrated into a partial bathtub plot and then fit by a complementary error function. The other real-time oscilloscope (brand Y) extracts both R_J and D_J from the time interval error spectrum. The time interval analyzer (TIA) uses a technique where the time interval error distribution is separated into correlated and uncorrelated subsets with R_J determined by applying a fitting algorithm to the uncorrelated set. D_J is determined from the time interval error spectrum and then combined with R_J to estimate TJ. The DCA-J separates the correlated and uncorrelated jitter, then measures R_J from the spectrum of the uncorrelated jitter, measures ISI, DCD, and DD_J from the correlated jitter distribution and extracts D_J from the whole jitter distribution.[5] The BERT uses the simplest application of the dual-Dirac model, R_J and D_J are determined by fitting the complementary error function to the slopes of the bathtub plot and the fits are extrapolated down to the BER of interest to estimate TJ.

Ultimately, all techniques for estimating TJ at a low BER are based on some form of the dual-Dirac model: they assume that the tails of the jitter distribution are determined by the R_J Gaussian.

Our goals were simple: determine the state of jitter analysis by resolving which techniques work and why. We configured the real-time oscilloscopes and TIA based on the manufacturer user manuals recommendation for greatest accuracy, regardless of test time. Since the jitter analyzers were only available to us for a limited time, we could only try one of the techniques available on each system.

We chose a set of test conditions with different combinations of R_J, sinusoidal and triangular PJ, ISI and DCD. R_J is caused by thermal effects and is assumed by the industry to follow a Gaussian distribution characterized by its width or standard deviation, σ. Deterministic jitter (D_J) includes data-dependent jitter (DD_J), and periodic jitter (PJ). DD_J can be further decomposed into the time component of inter-symbol interference (ISI) and duty-cycle-distortion (DCD). ISI is caused by network elements that modify the trajectories of logic transitions by changing the frequency and attenuation response of the system. The resulting bit transition time displacement is jitter induced by the time-component of ISI. DCD results from asymmetries in clock cycles. Since DCD occurs on clock signals, one could argue that it is not “data-dependent.” We file it under DD_J because DCD and ISI interfere -- changing one changes the other -- so DCD is data-dependent when coupled with ISI. PJ comes from electromagnetic pickup of periodic sources like power supply coupling. The distinguishing feature of D_J is that its peak-to-peak value is bounded. DCD and ISI are called *bounded correlated jitter* because they are correlated to the data signal; PJ and crosstalk are called *uncorrelated bounded jitter*; and R_J, uncorrelated unbounded jitter.

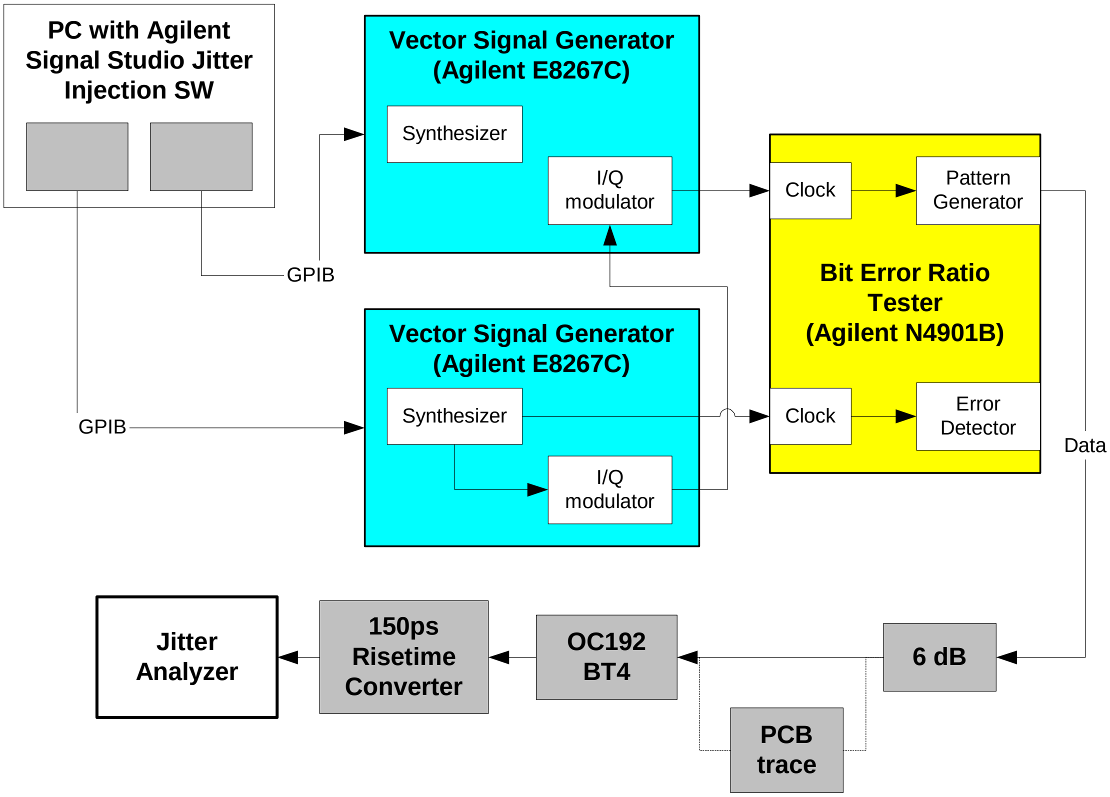

Fig. 2: Block Diagram Of Precision Jitter Transmitter

The precision jitter transmitter is described in detail in Reference[4], Fig. 2. The calibration of the applied jitter signals can be traced to accepted calibration standards. R_J was applied by I/Q modulation of the clock that drives a pattern generator and was calibrated with an uncertainty of 1.5% plus the baseline uncertainty of 0.27 ps. PJ is also applied by I/Q modulation and was calibrated with an uncertainty of 1%. DD_J was applied through combinations of DCD and ISI. DCD was applied by varying the crossing point setting of the pattern generator with an uncertainty less than 1 ps. ISI was applied by introducing different lengths of PCB traces and was calibrated with an accuracy of 0.5%. The calibration of the transmitter with no applied jitter signals, its baseline, was not traceable. The baseline was dominated by a combination of R_J, 0.685 ± 0.270 ps rms and ISI, 3.9 ± 1.0 ps peak-to-peak. TJ was calibrated by convolving the components of each jitter condition and integrating the results to obtain BER(x) and calculating TJ at a BER of 10^{-12}, TJ (10^{-12}). The uncertainty of TJ was calculated by propagating the uncertainties of each component through the calculation in the standard way.[7]

We chose the jitter conditions to satisfy three requirements:

1. We wanted to provide a set of conditions that sampled the typical space encountered in the lab
2. We set the levels of jitter to values near the margins of most specifications
3. We didn't apply any jitter conditions that the manufacturers said their equipment couldn't handle

So, it was a level field where everyone could play. The R_J ranged from less than 1 ps to about 5 ps, PJ was either 0, 7 or 27 ps, DD_J was either 4 ps or between 75 and 140 ps.

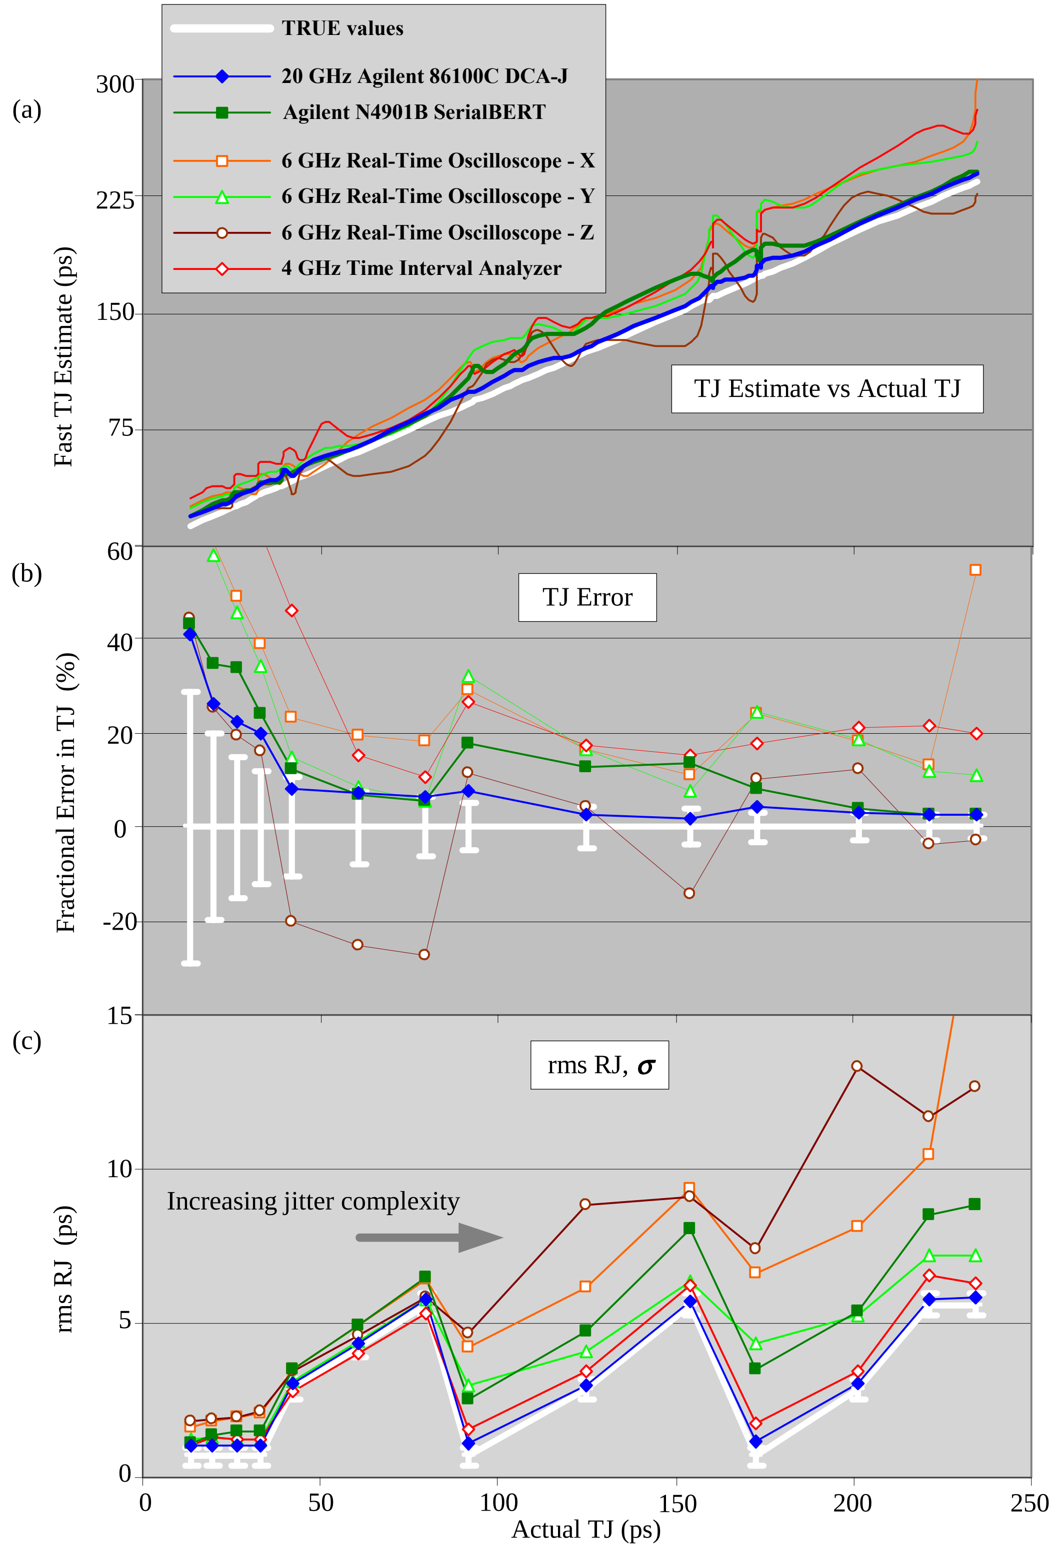

Fig. 3(a), (b) & (c): See Text And Analysis Below

Fig. 3(a) shows the fast estimates of TJ (10^{-12}) for each jitter analyzer plotted against the calibrated actual values of TJ (10^{-12}). Fig. 3(b) gives the percent error of the TJ (10^{-12}) estimates with error bars that indicate the calibration uncertainty -- any measurement within the vertical span should be considered consistent with the truth. Fig. 3(c) gives the R_J value for the points plotted in Fig. 3(b). As the jitter conditions become more complex, the accuracy of most of the jitter analyzers degrades.

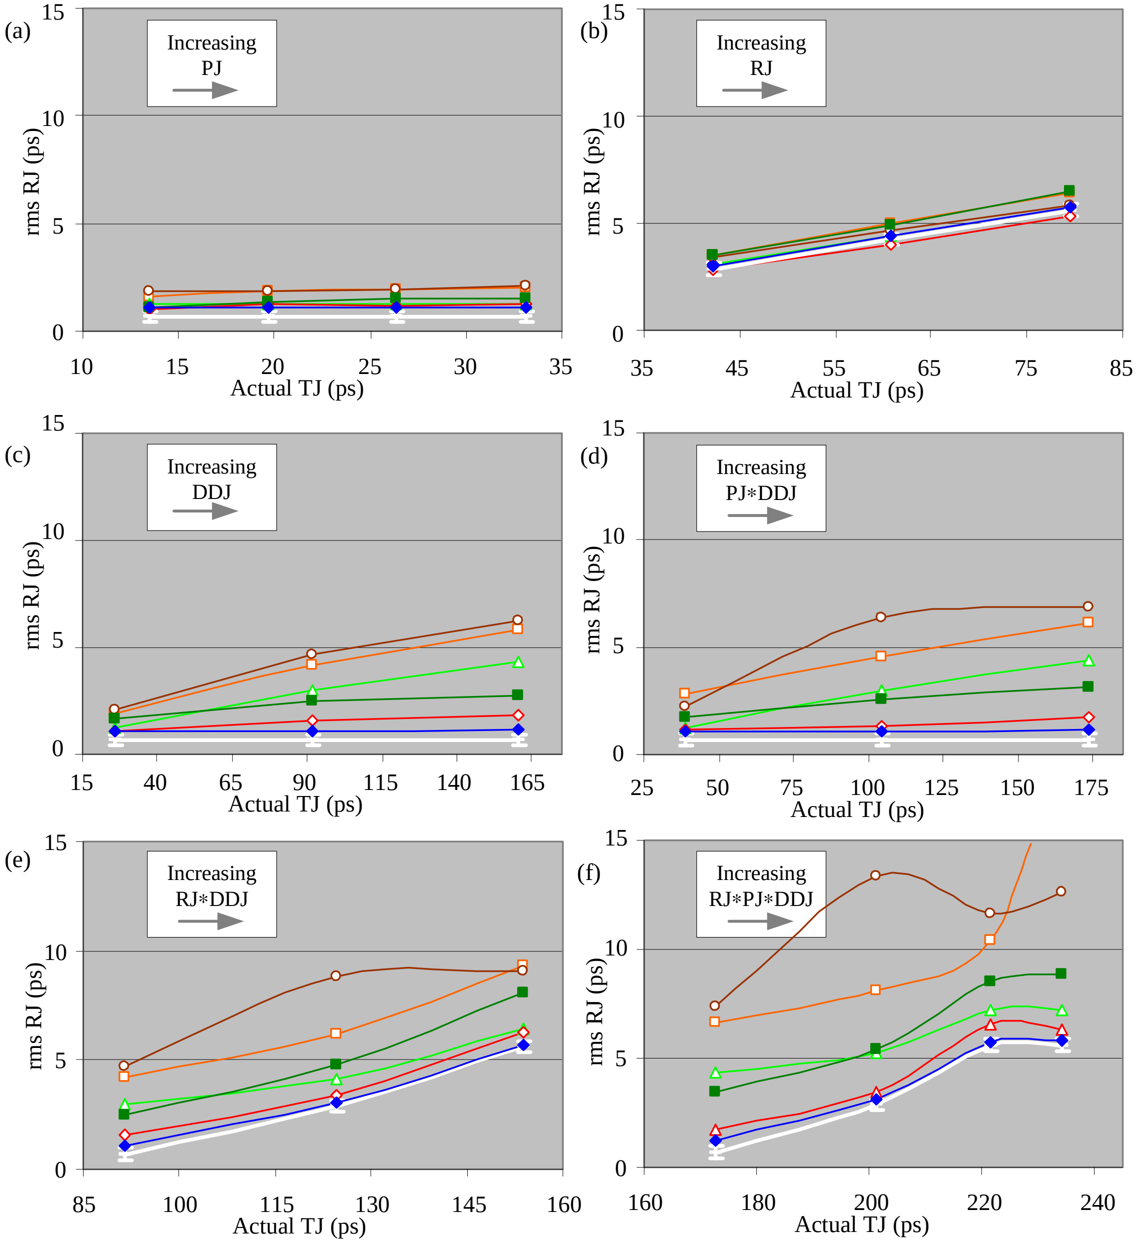

Fig. 4: R_J Measurements As A Function Of TJ (10^{-12}) With Conditions Dominated By PJ (a), R_J (b), DD_J (c), Mix Of PJ*DD_J (d), Mix Of R_J*DD_J (e), And Mix Of PJ*$R_J$$DD_J$ (f)

For values of TJ(10^{-12}) (<25 ps) the uncertainties are dominated by the transmitter baseline and the jitter analyzers are all near their TJ noise floors. Moving from left to right, in Figs. 3, the jitter conditions grow more complex, the fractional error in TJ, Fig. 3(b), is more or less constant for most analyzers, but the deviation of the measurements from the calibrated values (white line) in Figs. 3(a) and (c) show that as the jitter conditions become more complicated the accuracy of most jitter analyzers degrades.

The most aberrant behavior is demonstrated by real-time oscilloscope Z, the brown line with open circles. We'll figure out why oscilloscope Z underestimates TJ while nailing R_J later. But first, the largest cause of the discrepancies is demonstrated in Fig. 4 where R_J is separately plotted for different jitter conditions. Figs. 4(a) and (b) show that the analyzers all perform well in simple conditions with just R_J or PJ. Figs. 4(c) and (d) show that as increasing lengths of transmission channel are introduced -- increasing DD_J and holding R_J constant -- the reported R_J values of most analyzers increases; that is, DD_J is mistaken for R_J. There are at least two causes:

1. The analyzer voltage noise is converted to timing noise and interpreted as R_J
2. As the combination of D_J sources grows complicated, the tails of the D_J distribution become smoother and resemble a Gaussian distribution causing the techniques that fit the tails of either BER(x) or the jitter distribution to have trouble distinguishing R_J from D_J.

The voltage noise floor of the acquisition hardware affects jitter analysis by corrupting the time measurement of logic transitions. With three logic transitions (see Fig. 5) where each has the same amplitude noise and zero timing noise, moving from left to right the rise times of the signals increase and the amplitude noise has a larger affect on the transition time. When voltage noise is caused by the acquisition hardware, rather than the signal, the reported value of R_J -- the "R_J measurement floor" -- increases with increasing rise time. The precision jitter transmitter generates ISI through the filtering and attenuation effects of a PCB trace. In addition to increased DD_J from the time-component of ISI, the vertical component of ISI increases the rise/fall times of the signal and the analyzer noise is converted to something that jitter analysis algorithms report as R_J.

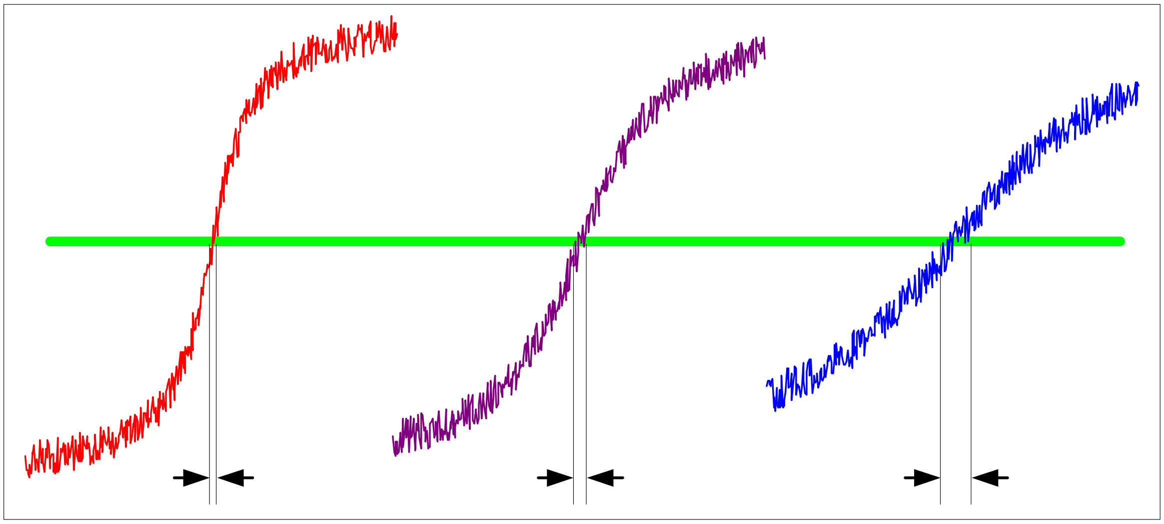

Fig. 5: Amplitude Noise Is Converted To Jitter

The voltage noise floor of real-time oscilloscopes is roughly proportional to the setting of their vertical sensitivity down to a minimum absolute noise floor. The real-time oscilloscopes that we studied have noise floors of 30 - 40 mV/div rms. With vertical sensitivity set to 100 mV/div, they had effective voltage noise floors of 3 - 4 mV rms. The voltage noise floor of the DCA-J is typically 0.25 mV.

On a BERT, voltage noise is described in terms of the error detector sensitivity, the minimum voltage difference required for the error detector to discriminate a logic '1' from a logic '0.' The error detector sensitivity determines the minimum observable TJ that a BERT can measure and causes a small bias in TJ measurements that is proportional to the signal's rise/fall time.

I was surprised that the effects of analyzer sampling clock jitter, trigger jitter, time-base linearity, and transition time accuracy were dwarfed by the effects of voltage noise. This is not to say that you shouldn't watch out for trouble from these sources, though.

The bandwidth of the analyzers wasn't a big issue for us because we chose a data rate that was easily accessible to the analyzers. Generally speaking more bandwidth is better.

On the algorithmic side, it is important to measure R_J on a data set that doesn't include DD_J. As more D_J terms are introduced, the D_J distribution appears increasingly Gaussian. The reason comes from the central limit theorem: the distribution of the combination of a large number of independent processes follows a Gaussian distribution -- regardless of the shapes of the component distributions. To prevent algorithms from reporting DD_J as R_J, one should measure R_J on the data set that is uncorrelated to the test data. We can see what's going on by looking at the green line in Fig. 3(a). Except for the region of TJ values between about 100 and 150 ps, the fast BERT estimate is spot-on. The region between 100 and 150 ps is dominated by jitter conditions with low R_J and high DD_J. The reason that R_J is increasingly overestimated as the jitter conditions become more complicated is that the edges of the bounded D_J distribution become smooth and resemble a Gaussian distribution as more and more sources of D_J are included. Remember the central limit theorem from statistics, "the convolution of an infinite number of independent distributions follows a Gaussian distribution." When triangular PJ is combined with DCD and ISI, the edges of the distribution are smeared. A smooth D_J distribution convolved with Gaussian R_J results in a distribution with tails that can be well parameterized by a Gaussian -- but not the R_J Gaussian, rather, a Gaussian with a larger width than the underlying R_J width, σ. Now look at the measurements from the sampling oscilloscope, solid blue diamonds and from the TIA, open red diamonds. They both use techniques that remove DD_J from the data prior to extracting values for R_J and give the most accurate R_J measurements.

Once the correlated jitter, ie DD_J, is removed from the distribution, equating R_J to the rms noise in the jitter-frequency spectrum gives the most accurate R_J results. The reason the spectral technique is most accurate is simple. Accurate fits to the tails of BER(x) or the jitter distribution require a statistical sample large enough to assure that the region included in the fit is dominated by R_J, not D_J. While removing the correlated jitter goes a long way, it's possible for uncorrelated jitter -- eg triangular PJ -- to distort the fit. It's also possible for crosstalk from neighboring channels to contaminate the uncorrelated distribution in a way that mimics noise in the spectrum. If a neighboring channel is not phase-locked to the signal, or if it is counter-propagating, then crosstalk energy can appear in the continuous jitter-frequency spectrum and be mistaken for noise. One way around this dilemma is to measure R_J with neighboring channels turned off, then turn on crosstalk channels and perform the jitter analysis with the R_J value fixed to the value measured without crosstalk.

That the brand Z real-time oscilloscope, indicated by the hollow brown circles, underestimates TJ in R_J-dominated environments, shows how algorithmic parameters can be tweaked too much. If we only consider jitter conditions with applied R_J, it's easy to see what's going on. Fig. 4(b) shows that oscilloscope Z gives accurate R_J measurements when there is no DD_J and Fig. 6 the TJ (10^{-12}) estimates for the conditions in Fig.4(b), show that oscilloscope Z underestimates TJ (10^{-12}) in those conditions. Oscilloscope Z uses a fitting technique to measure R_J and D_J similar in principle to the standard dual-Dirac fitting algorithm performed by the BERT. The simplest imaginable application of the dual-Dirac approximation under these conditions should give accurate TJ estimates. That R_J is accurate and TJ is not indicates that the designers of the jitter algorithm used by oscilloscope Z did something a little too fancy in tuning their parameters.

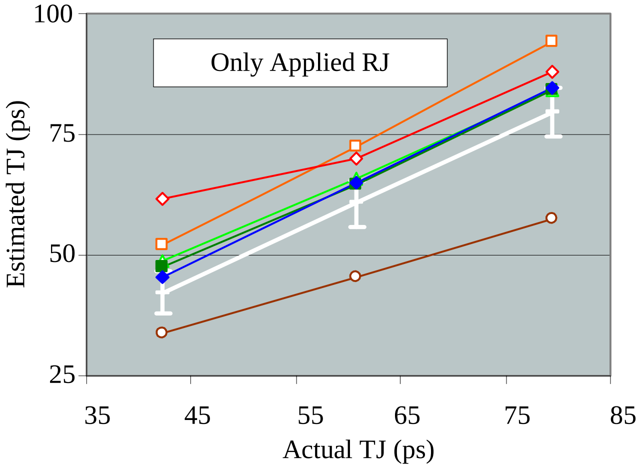

Fig. 6: TJ (10^{-12}) Estimates Vs Actual TJ For Conditions With Only Applied R_J

JitterFest taught us the not-very-shocking lesson that the combination of clean data acquisition hardware and well considered software provide the most accurate jitter measurements. The key algorithmic techniques[5] were to first separate the jitter distributions that were correlated to the data from that which was uncorrelated and then to use a spectral technique for measuring random noise in the uncorrelated sample. Having obtained rms R_J, σ, measurement of the dual-Dirac model-dependent D_J is simplified and accurate estimates of TJ at low BERs is straightforward.

That said, no thorough jitter analysis can be performed without a BERT for one simple reason: TJ at low BERs can only be measured on a BERT. While the fast TJ estimates performed on a BERT are not the most accurate, the BERT provides the only method for actually checking the TJ results.

The most poignant result of JitterFest was the failure of real-time oscilloscopes to perform accurate jitter analysis. I think that the fundamental root cause is simply that real-time oscilloscopes tend to have too much amplitude noise. As data rates increase, we're likely to see more and more engineers adopt equivalent-time sampling oscilloscopes as their default signal integrity tool. The combination of lower noise, wider bandwidth, new clock-recovery hardware that allow them to trigger on data signals, and lower price make them more desirable signal integrity analyzers than real-time oscilloscopes.

About The Author

Ransom Stephens (ransom_stephens@agilent.com) specializes in the analysis of electrodynamics in high-rate digital systems and the marketing of analysis tools developed by the Digital Verification Solutions division of Agilent Technologies. He has spent the last four years analyzing timing noise and dispersion and developing new techniques for extracting signals from noise. He received his PhD in physics at UCSB in 1990 and spent the succeeding ten years in basic research making precise measurements of rare processes in particle physics at laboratories across the United States and Europe.

Footnotes

[1] Ransom Stephens, "Jitter analysis: The dual-Dirac model, RJ/DJ, and *Q*-scale," Agilent Technologies Whitepaper, 2004. http://www.home.agilent.com/upload/cmc_upload/All/dualdirac1.pdf

[2] MJSQ stands for Methodologies for Jitter and Signal Quality and the document is available at http://www.t11.org/index.htm, T11 is a technical committee within the International Committee for Information Technology Standards. And, please, never use the term "methodology" when "method" works just as well, thank you.

[3] The references to the "HP way" come from the fact that Agilent Technologies was spun off from Hewlett-Packard Corporation several years ago. I think of Agilent as the high-tech outfit that Bill Hewlett and Dave Packard envisioned when they sold that first oscillator built in Bill's garage. (This is the opinion of the author, of course, but that doesn't change the fact that it's the truth.)

[4] Jim Stimple and Ransom Stephens, "Precision Jitter Transmitter," DesignCon 2005. http://www.home.agilent.com/upload/cmc_upload/All/precisionjittertransmitter.pdf

[5] Agilent Technologies Product Note, 86100C-1, "Precision jitter analysis using the Agilent 86100C DCA-J," Agilent Literature Number 5989-1146EN. http://cp.literature.agilent.com/litweb/pdf/5989-1146EN.pdf

[6] Agilent Technologies' Application Note, "Jitter Fundamentals: Agilent 81250 ParBERT Jitter Injection and Analysis Capabilities," Literature Number AN-5988-9756EN. http://cp.literature.agilent.com/litweb/pdf/5988-9756EN.pdf

[7] There is a terrific summary of probability and statistics applied to all sorts of measurements at the Particle Data Group's web site http://pdg.lbl.gov got to "Mathematical Tools" under "Reviews, Tables, and Plots."

as published in...
analogZONE