DesignCon 2007

A Comparison of Methods for Estimating Total Jitter Concerning Precision, Accuracy and Robustness

Martin Miller Ph.D., Chief Scientist, LeCroy Corporation
marty.miller@lecroy.com

Michael Schnecker, Product Specialist, LeCroy Corporation
michael.schnecker@lecroy.com

Abstract

Oscilloscopes have been used extensively to analyze the jitter performance of serial data links providing estimates (measurements) of total jitter as well as its "random" and "deterministic" parts. Higher speed serial data signals containing jitter from sources such as crosstalk and multi-Gaussian random noise ***can*** cause traditional oscilloscope-based methods to greatly over or under estimate the total jitter. This paper presents a comparison of several common methods for estimating total jitter using spectral and statistical techniques, as well as a new method (Normalized Q-Scale). Results are shown for both real and simulated (Monte-Carlo) measurement scenarios, giving insight into weaknesses and strengths of these methods. Specifically, issues of accuracy, precision, convergence or robustness are treated.

Authors Biographies

The authors are colleagues and collaborators on jitter related issues at LeCroy Corporation, based in Chestnut Ridge, New York, USA.

Martin T Miller, PhD, Chief Scientist at LeCroy Corp, has been a hands-on engineer and designer at the company for 29 years. His doctorate is from the University of Rochester (Rochester, NY) in high-energy physics. Miller has contributed analog, digital, and software designs, and during the last 16 years he has focused on measurement-and-display software capabilities for LeCroy scopes. A native of Baltimore, he holds several US patents and participates in task groups for JEDEC concerning jitter measurements.

Michael Schnecker, Product Specialist at LeCroy Corp, holds an MSEE from the Georgia Institute of Technology and has 18 years of experience in the test and measurement industry. Mr. Schnecker specializes in signal integrity measurements, including jitter, and has been key in specifying and deploying the LeCroy SDA series instruments.

Summary

Jitter measurements performed using an oscilloscope sometimes utilize the jitter spectrum to separate the random and deterministic jitter components. The resulting estimates of the ***Rj***, and ***Dj*** values are used to estimate the jitter CDF and bathtub curve from which the total jitter estimate is derived for a specified bit error rate. Digital oscilloscopes capture the digital signal waveform and use software to derive a phase reference (extracted clock or virtual CDR) from which the time interval error (TIE) is calculated. The spectrum of the TIE over time includes discrete peaks and a noise floor. The noise floor is used to estimate the standard deviation of the random part of the jitter which is assigned to a normal Gaussian distribution. This type of jitter estimation has been implemented in both real time and equivalent time sampling oscilloscopes.

Serial data bit rates have steadily increased to 5 Gb/s and beyond thus reducing the overall jitter margins while at the same time increasing the importance of effects of crosstalk and other jitter aggressors. These jitter sources have more subtle effect on the jitter spectrum than simple periodic jitter and can easily be mistaken for noise in the jitter separation procedures. Crosstalk and similar jitter sources are typically uncorrelated but bounded and they can be mistaken for random jitter by spectral jitter estimation methods, the random jitter as well as the total jitter can be severely overestimated in these cases. In other cases, the random jitter is composed of a number of Gaussian distributions only two of which have a significant impact on the total jitter. These dominant contributors are often at a much lower statistical weight and failure to account for this fact can lead to under estimation of the total jitter.

Definition of Total Jitter

Not so long ago engineers spoke of “worst case jitter”. There was a basic problem with this notion, in that as the observed number of measurements increased, so did the answer to the question. Without further qualification, this question had a non-convergent answer. No one seemed to notice, since frankly budgeting for timing was a pretty simple matter compared to today’s situation.
Today clock and bit rates are ever moving higher and so both the impact of error per time increases, while the total timing budget tightens. It would not be an exaggeration to say most of the electronics industry recognizes that jitter and its influence has become a matter of strong interest and importance. Thus the term “total jitter” was invented to circumvent the dilemma (of non-convergent pk-pk timing variations) and clarify this situation. ***Tj*** is the “expected value” of Peak-to-Peak observed *Time Interval Error* (***TIE***) jitter for ***a specified number of observations***.

The total jitter definition requires then, a further qualifying datum: at what number of measurements do we anticipate the value given for the pk-pk jitter variation, or at what bit error ratio, or alternately to what confidence level is this figure to be associated. The number of observations and the BER are related to the confidence level

- In this context, a 95% confidence level means that 95% of the measurements of timing error fall between some range of error. This range of error is the confidence interval and corresponds to ***Tj***. For this case, the rate of errors exceeding the range of ***Tj*** would be 5%
- Likewise for a BER of 10^{-9} (or one in a billion) the corresponding confidence level would be ($1.0- 10^{-9}$), and the confidence interval would indicate the range which is exceeded at a rate of 10^{-9} to 1.

It is worth noting that for some common specifications, the confidence level is not very demanding, and the question can be resolved directly by observing the ranges of values for a sample size (or number of observed measurements).

For example, if we wish to estimate the ***Tj*** for a BER of 10^{-5} (one part in 100,000) then we can either:

- Repeatedly measure 10^5 values of TIE, and report the ***average*** pk-pk variation
- Measure some larger number of values, and find the range that contains (the fraction) $1\text{-}10^{-5}$ of the total number of values

I classify measurements of this kind as “above the waterline” because the measurements can be made directly and do not require inference or extrapolation. Of course most of the more interesting cases in telecommunications and data communications require very high levels of confidence to attain low error rates, and for the most part measurements using oscilloscopes must resort to extrapolation to estimate ***Tj*** in such cases.

It is however worth mention that the sample sizes or number of timing variations that can be measured by an oscilloscope is rising rapidly as are sample rates, recording durations and data processing horsepower.

- Question for audience: If we accept the definition of Tj proposed by this author, what is the importance of ***Tj*** for Period, Cycle-Cycle, or Half-Cycle Jitter distributions (as opposed to TIE)? For clock sources in general, which "total" jitter measurements are significant?

Review

For some years now, we have been living with the (already) classic equation[1] to represent total Jitter for a given Bit Error Ratio (BER) and for a "dual-Dirac" jitter distribution.

$$Tj(BER) = Dj + \alpha(BER) \cdot Rj \qquad \textbf{(0.1)}$$

Where $\alpha(BER)$ is described as the confidence interval for a single "normal" Gaussian for any confidence level: $CL = 1 - BER$ (or the total jitter for a single Gaussian jitter distribution).

$$\alpha(BER) = 2 \cdot \left| CDF^{-1}_{Gausiian}\left(\frac{BER}{2\rho_{tx}}\right) \right| \cdot Rj \qquad \textbf{(0.2)}$$

We are encouraged to believe that ***Dj*** represents some kind of constant systematic error, and that ***Rj*** is the correct Gaussian sigma to apply to the estimation of total jitter over statistics.

There are a couple of problems with this formulation, which should be recognized before going much further, in particular if the reader wishes a more accurate representation of how a dual-Dirac actually manifests statistically, and moreover if the reader is interested in ***not*** very low Bit Error Rates or Ratios (e.g. 1 part in a million)

---

[1] Fibre Channel – Methods for Jitter and Signal Quality Specification – MJSQ, T11.2 project 1316-DT Rev. 14, June 9, 2004

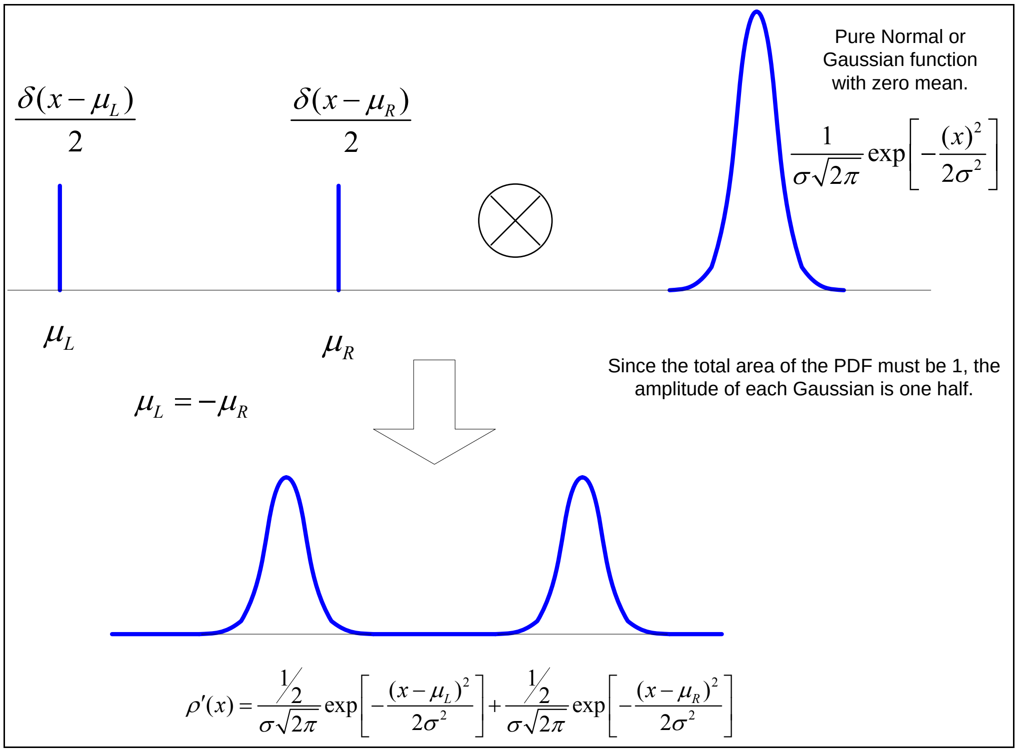

Figure 1: *convolution of a dual-Dirac distribution with a single Gaussian*

While this is a very simplified case, the real cases we encounter are more often two Gaussians with unequal sigma (Rj) but each with area ½.

The most general way to write the expression for total jitter dominated by two Gaussian contributions on each side would be:

$$Tj(BER) = Dj + \frac{\alpha}{2}\left(\frac{BER}{W_{left}}\right)\cdot Rj_{left} + \frac{\alpha}{2}\left(\frac{BER}{W_{right}}\right)\cdot Rj_{right} \tag{0.3}$$

So for the case of the dual-Dirac with equal weights (left and right) and a total weight of 1.0, a simplified expression would be:

$$Tj(BER) = Dj + \frac{\alpha}{2}\left(\frac{BER}{0.5}\right)\cdot Rj_{left} + \frac{\alpha}{2}\left(\frac{BER}{0.5}\right)\cdot Rj_{right} \tag{0.4}$$

If further the right and left-hand sigma's are equal, then:

$$Tj(BER) = Dj + \alpha\left(\frac{BER}{0.5}\right)\cdot Rj \tag{0.5}$$

Or in terms of the inverse *CDF*: (for convenience of calculation)

$$Tj(BER) = Dj + 2 \cdot \left| CDF^{-1}_{Gausiian}(BER) \right| \cdot Rj \qquad \textbf{(0.6)}$$

This expression is noticeably different from equation (0.1). What's surprising is the traditional equation is correct only for ***Dj*** much smaller than ***Rj***, which frankly is rarely the case. Once ***Dj*** is as large as half of ***Rj***, the expression **(0.5)** is more accurate.

The difference is most accentuated for small bit error rates (as are often used for oscillator manufacturer's specifications). This is significant because clock jitter is the one case where dual-Dirac distributions are common and where in many cases "above the waterline" measurements are made since the BER specification is "loose". The two graphs below show the error for each expression as a function of ***Dj*** (normalized to ***Rj*** of 1.0). It is interesting to notice that for a given BER, the difference between the two predictions seems to be a constant.

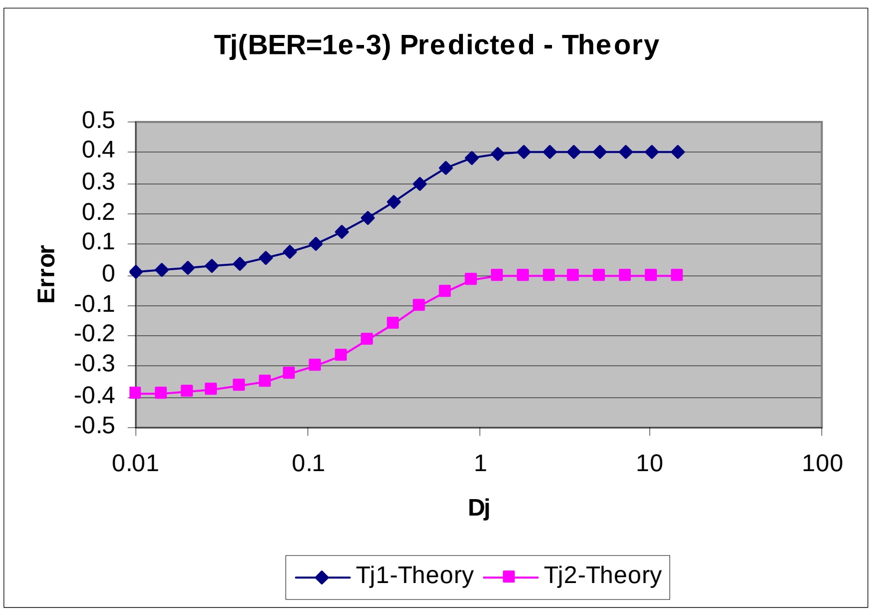

Figure 2*: Error in common* ***Tj*** *equation from real convolution, and error in modified equation for dual-Dirac* ***Tj****. Note for one part in a thousand and a separation of only 1 times* ***Rj****, the error is about 0.4* ***Rj****.*

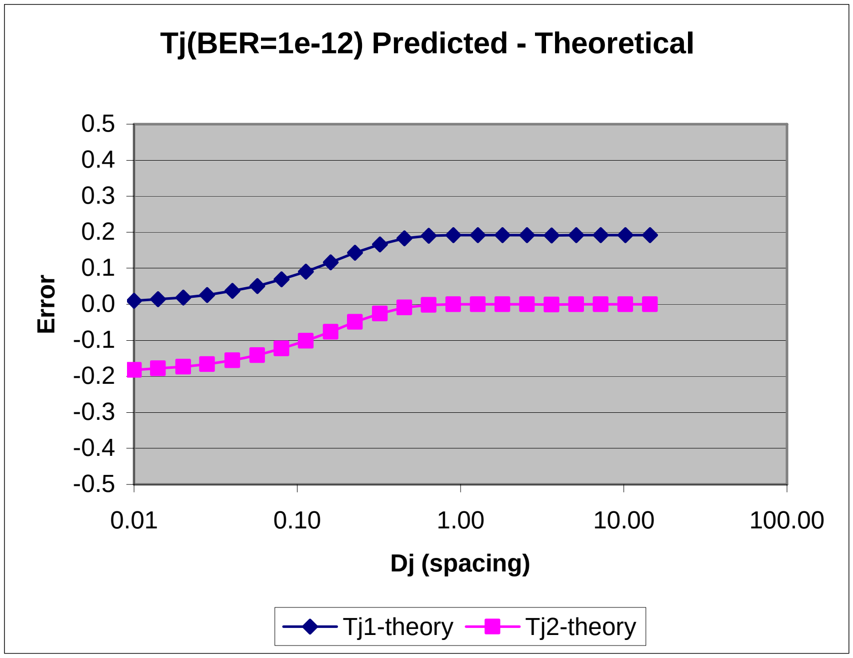

Figure 3**: Error in common **Tj** equation from real convolution, and error in modified equation for dual-Dirac **Tj**. Note for one part in 1e12, the difference between the two models only amounts to 0.2 sigma for **Dj** spacing ≥ 1 **Rj

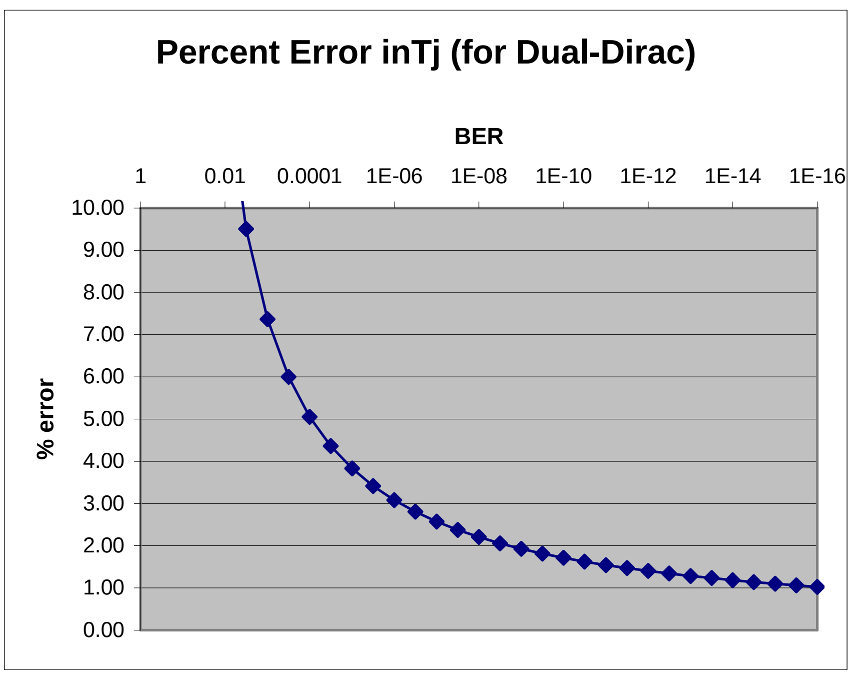

Figure 4:** This graph shows how the traditional equation overestimates **Tj** by nominally 1% at low bit error rates, for **Dj** = **Rj

So What?

Why would this be an interesting line of discourse? The difference between the two equations and results they predict are small (only 1/6 of 0.4 for BER = 1e-3, and 1/14 of 0.2 for BER=1e-12, both of which are small percent errors for ***Tj***). There are basically two reasons:

1. To be more mathematically rigorous in the understanding of the problem, and
2. To raise consciousness that the ***weight*** of the distributions which dominate the right and left sides of the distribution *do* play a role in the ***Tj*** observed for any measurement scenario. This is actually a foreshadowing of an important aspect of the normalized Q-Scale method

Methods for Estimation of Total Jitter (*Tj*):

BERT Scan:

One way to get a ***Tj*** estimate is to use a Bit Error Rate Tester or BERT to perform a BERT scan. The instrument generates a known pattern of data and artificially slides the sampling moment or time across an entire unit interval and looks for (counts) errors from the expected bit pattern.

Such a test is able to observe errors (not timing errors, but decoding errors) where the assumption is that the error is due to a timing or voltage error. The BERT can thus learn the boundaries of ***Tj*** (the confidence limits). It is not the purpose of this paper to discuss in detail the relative benefits of such a test method. It is mentioned because, while the BERT scan is a powerful verification of ***Tj*** estimates since it ***can*** be "above waterline" for many controlled measurement scenarios.

Equivalent-Time and Real-Time Oscilloscopes:

Oscilloscopes of both the "sampling" or equivalent time class as well as the Real-Time Oscilloscopes can be used to estimate ***Tj***. Similar methods of analysis can be applied to both classes of instrument.

It is not the purpose of this paper to enter into the various advantages of these classes of instruments; rather the essential matter is that oscilloscopes can provide measurements of TIE which then can be submitted to measurement methods (or calculations) which yield estimates of ***Tj*** … and how do those methods compare.

The first method to be discussed is the numerical fit of the extremes of the observed histogram distribution. This method assumes the tail is dominated by a Gaussian probability distribution ***PDF.*** This approach was the one this author first took when embarking on the mission of estimating total jitter. There are several pitfalls to the method, but it works well for well-behaved systems. The assumption is one often in the analysis of High Energy Physics data, is that a Gaussian "tail" has a quadratic behavior in the logarithmic domain.

Normal (Gaussian) Statistics: The probability distribution function is expressed as

$$\rho_n = \frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right]$$

where μ is the mean of the distribution, and sigma is the Gaussian standard-deviation,

So the relationship translated to a logarithmic scale would be:

$$\log(\rho_n) = -\log(\sqrt{2\pi}\sigma) - \left[\frac{(x-\mu)^2}{2\sigma^2}\right]$$

This form is *quadratic* but with the linear (first order) term exactly zero.

Observing this expected behavior for the extremes (or tails) of the distribution, we can obtain an estimate for the functional form, by performing an analytic weighted regression (or “weighted fit”). There are two approaches, and the final choice can be made to suit the occasion. One is to perform the weighted analytic procedure with three parameters “free” or to confine the functional for to the more expected “pure quadratic” (as indicated in the box above). This choice is less critical than the choice of how much of the distribution extreme to include in the fit, and how to assign weights and values to unpopulated distribution bins.

While this method is not without merit, it has been abandoned (by this author) in preference for the normalized Q-Scale described later in this paper. While the former method is easily adapted to be very sensitive to rare events, it is less robust and less rapidly converging than the superseding method. Although the author has put such a solution in public domain, it is believed that no-one is currently using such an algorithm in popular instrumentation.

Another approach is more mathematically rigorous, and is attributable to Levenberg-Marquardt[2], and is very general. Again, this method works well for well behaved systems and sports the major disadvantage that finding good first estimates of the means, weights and sigma’s of the distributions is critical for robust behavior. Since many real-world

---

[2] Levenberg, K. "A Method for the Solution of Certain Problems in Least Squares." Quart. Appl. Math. 2, 164-168, 1944. and Marquardt, D. "An Algorithm for Least-Squares Estimation of Nonlinear Parameters." SIAM J. Appl. Math. 11, 431-441, 1963.

measurements are NOT well behaved, this method has also been abandoned by this author (or rather displaced[3]) in favor of the normalized Q-Scale.

Spectral Method on the FFT of *TIE(t)*

This method[4] starts with a sequence of jitter measurements assumed to be uniformly spaced in time (i.e. nominally one jitter value at each nominal clock transition). This representation of jitter vs. time is transformed[5] into the frequency domain, and analyzed as a spectrum. The method assumes:

- That any spectral "peaks" are manifestations of deterministic jitter, and thus provide a measure of ***Dj***
- The remaining spectrum, or "noise floor" accounts for all of ***Rj***
- Also, that all of the noise floor is due to Gaussian processes, such that their RMS contributions may be combined (in quadrature) to yield and overall value for ***Rj***
- Further that the values obtained for ***Rj*** and ***Dj*** can be combined in the fashion described above to obtain an estimate for ***Tj***.

There are many cases where these assumptions are justified. In those cases this method has at least one thing strongly in its favor:

- It's rapid to obtain a converged result … (i.e. you don't need to acquire thousands of acquisitions to get a stable and repeatable result)

However, there are a few notable cases which do not meet these requirements and for which erroneous estimates will result. Such cases include:

- Non-Gaussian incoherent aggressors such as cross-talk from a nearby data channel (such an aggressor can be simulated by modulating jitter with a PRBS pattern). In such a case, the incoherency implies that the spectral influence is only to the noise floor and does not manifest as peaks, even though the contributions are not Gaussian in nature).
- Non-Stationary Periodic aggressors (like Spread Spectrum) manifest broadly in the jitter spectrum and cannot always be identified as "peaks" and so can pollute the noise floor with non-Gaussian contributions.

In effect, any wide-band aggressor which contributes bounded timing fluctuations will be indistinguishable from the noise floor, and consequently "counted" as ***Rj***. Many such

---

[3] By this, the author implies this method is useful as a secondary method, once the normalized Q-Scale method has been used to obtain useful and robust estimates for the parameterization of the jitter distribution.
[4] This method is described as "prior art" in Tektronix US patent 6,853,933 B2 apart from the methodology described as "invention" which concerns identifying the frequency domain "impulses" ascribed to *Dj* components of jitter.
[5] Obtaining the sequence of timing errors as a function of time generally requires that the oscilloscope be "real-time" or that it captures the entire evolution of timing errors over a single continuous interval of time.

cases can be shown to result in enormous overestimations of ***Tj*** when (either) of the equations for obtaining ***Tj*** as a function of ***Dj*** and ***Rj*** are used (since the multiplier on ***Rj*** is large for most interesting bit error ratios or in-confidence levels.

In addition, there is a critical difficulty with this method. Peak-detection is an error prone and not well standardized procedure. It is as easy to fail to identify peaks as it is to incorrectly identify them. As with most algorithms of this kind, a well-behaved system submits itself to the method, but anything not perfectly Gaussian is easily miss-analyzed.

Analyzing Statistics of the Empirical Distribution Function (EDF) and Q-Scale Renormalization:

This method is a variant on a method described nicely in the whitepaper [1], and does not presume to "invent" the Q-Scale view, which is actually much older than the reference. Like the reference, this method analyzes the Empirical Distribution Function (or integral of the histogram) rather than the histogram itself. This eliminates some subtle assumptions which must otherwise be made to deal with zero populations, and has same advantages as described in the Agilent document. However, it takes the model and method at least one step further.

This method is different in that it recognizes: ***in the presence of deterministic contributions, the tails of the overall distribution are governed by distributions which are not "whole" or unity normalized Gaussian distributions***.

Like the review of the ***Tj*** equation above, this method explicitly treats the weight of the dominating distributions at the boundaries of the distribution. The notable fact being that in most real-world cases (even for dual-Dirac) the extremes of the distribution are governed by a fraction (if not a *small* fraction) of the overall possible events in the distribution.

This method is based on statistical analysis, and first collects a histogram of the observed measurements of timing[6] errors or "jitter". If H_k is the population of the k^{th} histogram bin, and h_i is the left edge of the i^{th} bin, the empirical distribution function is then calculated from the histogram as:

$$EDF(x = h_i) = \frac{\sum_{k=0}^{k=i-1} H_k}{\sum_{k=0}^{k=N-1} H_k} = \frac{1}{P_{total}} \sum_{k=0}^{k=i-1} H_k \qquad \textbf{(0.7)}$$

[6] It is notable that the method can also be applied to vertical noise or to any statistical observation to determine extremal behavior and even to determine if Gaussian processes are "at work'.

The histogram is then summed from the left and right to the median coordinate and a symmetric *EDF* calculated. Each half of the distribution is calculated as follows:

$$EDF_{left}(x = h_i) = \frac{1}{P_{total}} \sum_{k=0}^{k=i-1} H_k \qquad \textbf{(0.8)}$$

$$EDF_{right}(x = h_i) = \frac{1}{P_{total}} \sum_{k=N-1}^{k=i} H_k \qquad \textbf{(0.9)}$$

This EDF is then plotted (internally if not displayed) on a variable normalization Q-Scale. The usual transformation (of the vertical axis in probability) is to a new variable *Q* obtained from *BER* as follows:

$$Q(BER) = CDF^{-1}_{Gaussian}(BER/2) \qquad \textbf{(0.10)}$$

For the sake of grasping the procedure, we show such a representation for a pure Gaussian Monte-Carlo generated histogram. The marvelous thing about this representation of the EDF, is that it's nice straight lines. In fact, the slope of the Q-Scale plot (for a single normal Gaussian) below gives directly the sigma or ***Rj*** of the distribution.

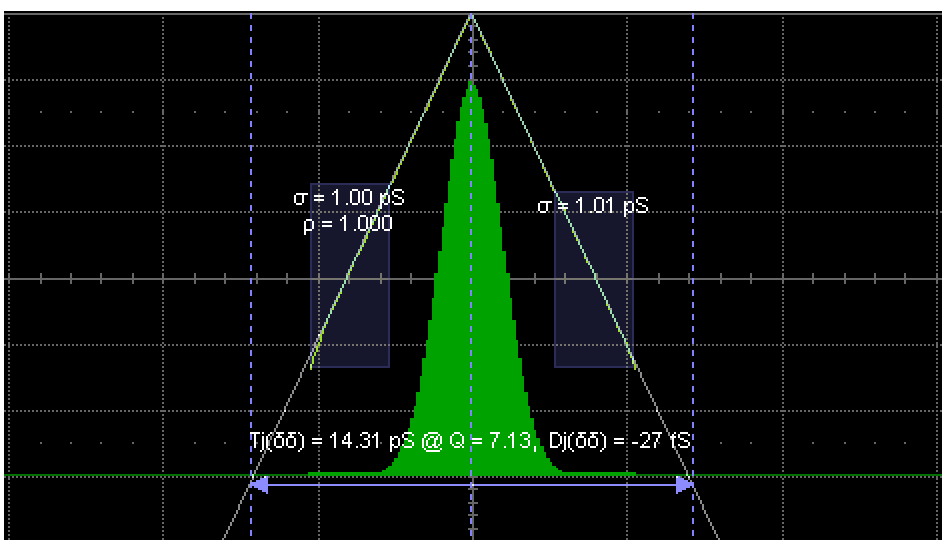

Figure 5*: A pure Gaussian distribution viewed on the standard Q-Scale looks like a perfectly linear triangle. The scale per division for this and all other Q-Scale plots is 1Q per division.*

However, in keeping with the notion of a normalization associated with the behavior, a new Q value is postulated for which there is a normalization factor, ρ_{norm}

$$Q_{norm}(BER) = CDF^{-1}(\frac{BER}{2\rho_{norm}}) \qquad \textbf{(0.11)}$$

If we were to remain with prior standard Q-Scale … and use it to observe an archetypical dual-Dirac distribution, we would see this (the green line), and obtain these results for the right-hand and left hand slopes (reciprocal of ***Rj***).

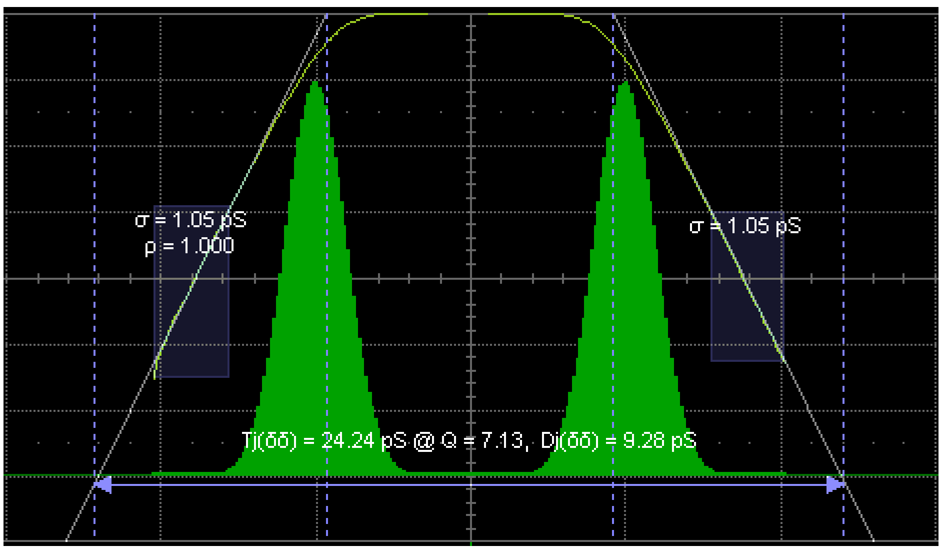

Figure 6*: This is probably the next most important case and represents a Dual-Dirac Normal Distribution. The total area of the distributions PDF is 1.0*

Now it is possible to "fit" to this EDF, and the resulting (right and left) ***Rj*** values are overestimated by about 5%. Note also that the ***Dj*** figure given from the intercept of the lines with Q = 0, is significantly underestimated. The lines associated with the observed EDF are curved at the top.

If however we use the proposed normalized Q-Scale instead, and set the normalization to ½ (since we know each distribution represents only half of the overall distribution, we obtain very linear behavior of the EDF, *we get the* **correct** values for ***Rj*** and ***Dj***

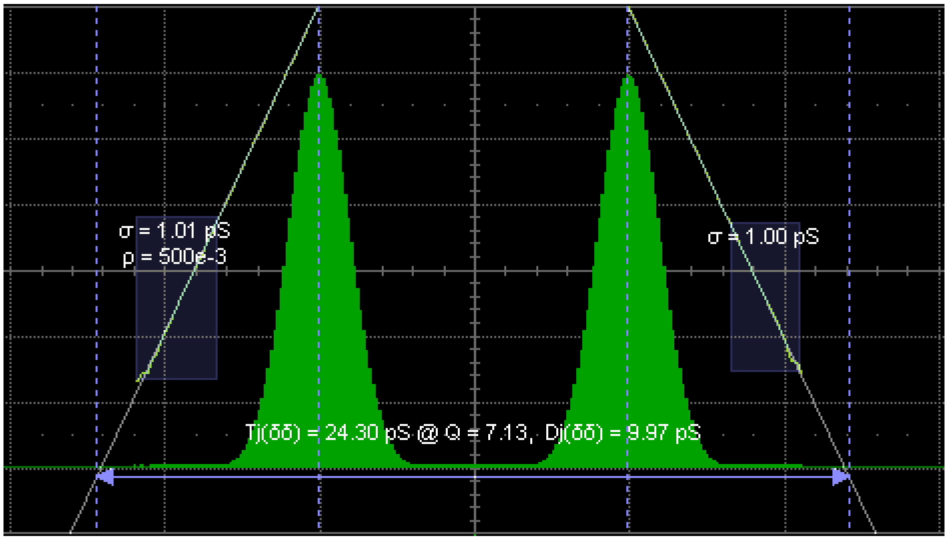

***Figure 7**: The same dual-Dirac distribution represented on the normalized Q-Scale is very linear, gives the correct positions of the peaks (so correct **Dj**) and the correct **Rj**'s for the right and left hand sides.*

The Normalized Q-Scale can be used with an arbitrary normalization factor, and indeed there are many real-world cases where the normalization becomes quite small. If the weight or ρ is known in advance (as it is for the dual-Dirac model) then choosing an appropriate value is simple. But what if we don't know in advance the weight of the dominant distributions?

Finding the "likely" weight for the extreme behavior of distributions is the subject of intellectual property applications by LeCroy Corp and this author.

The method, in a nutshell, is to vary the normalization to find one that yields the most linear[7] behavior for each extreme, and to thereby obtain estimates for the mean, sigma and weight of the dominant contributors.

Perhaps as important as the quantitative results is the confirmation of any linear behavior on the normalized Q-Scale. In some real-world cases we have observed, no believable linear behavior for any normalization exists. This is a very important qualitative observation or determination. The conclusion in such a case is that the extreme behavior of the distribution is ***NOT*** Gaussian.

Coming back to the cases which *do* show Gaussian behavior, look at the case of periodic aggressors. Viewed on a standard normalization, the sinusoidal aggressor (what is usually called ***Pj***) looks like this:

---

[7] Using numerical methods, a "best" compliance to linear behavior can be obtained. One example is to use the best "chi-square" result for linear regression to determine which normalization is "best". The constraints and assumptions may be manifold, but for even complex distributions, the effective dominant weights for right and left extreme behavior seem to be revealed by this procedure.

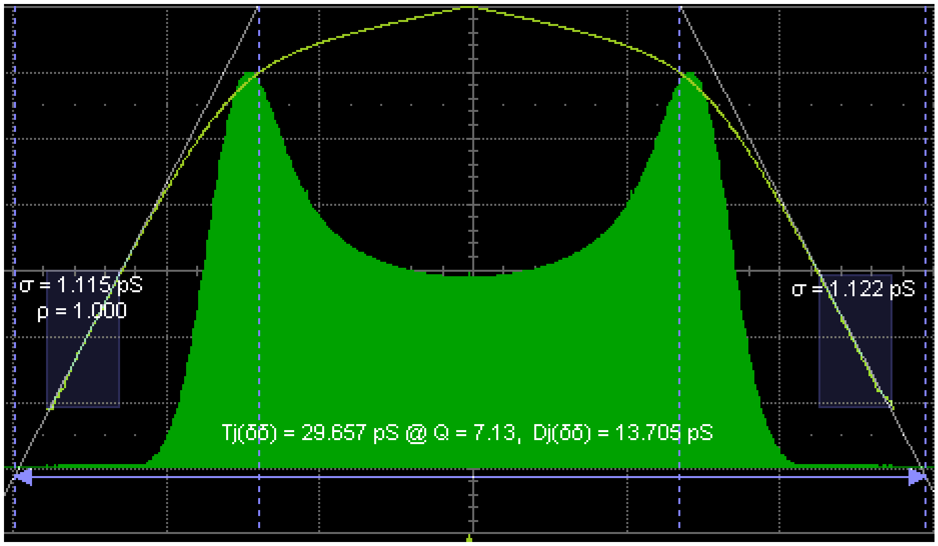

Figure 8*: Un-normalized Q-Scale plot for the PDF of a sinusoid with p-p amplitude 16ps convolved (Monte-Carlo) with a 1ps Gaussian, displays extremely non-linear extremal behavior although asymptotically linear (large Q)*

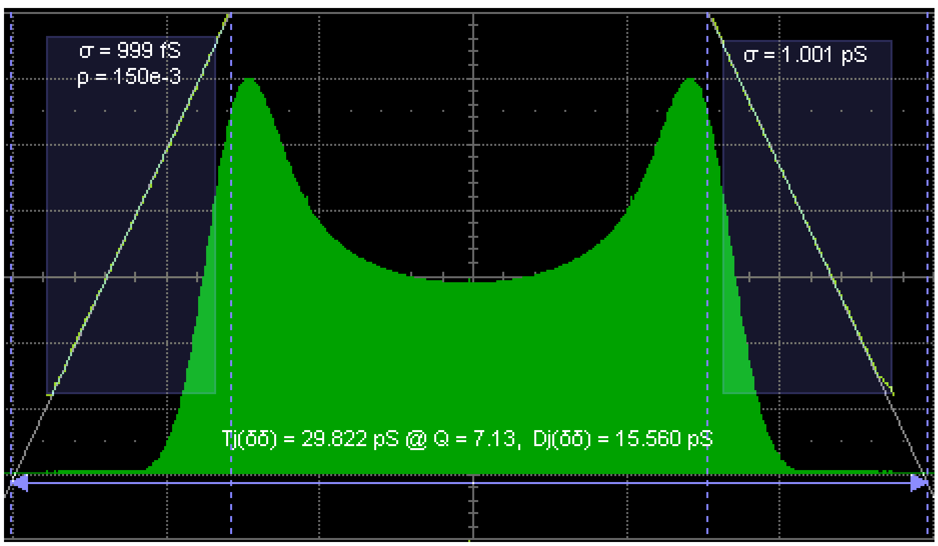

Figure 9*: Using an automatically normalized Q-Scale the correct details of the underlying probabilities are more closely estimated. The* ***Rj*** *values are more accurate(w.r.t. simulation parameters) , and the effective* ***Dj*** *is very plausible as well as the weight of the dominating distributions (about 0.15 each)*

Examples of Measurements and Monte-Carlo Simulations:

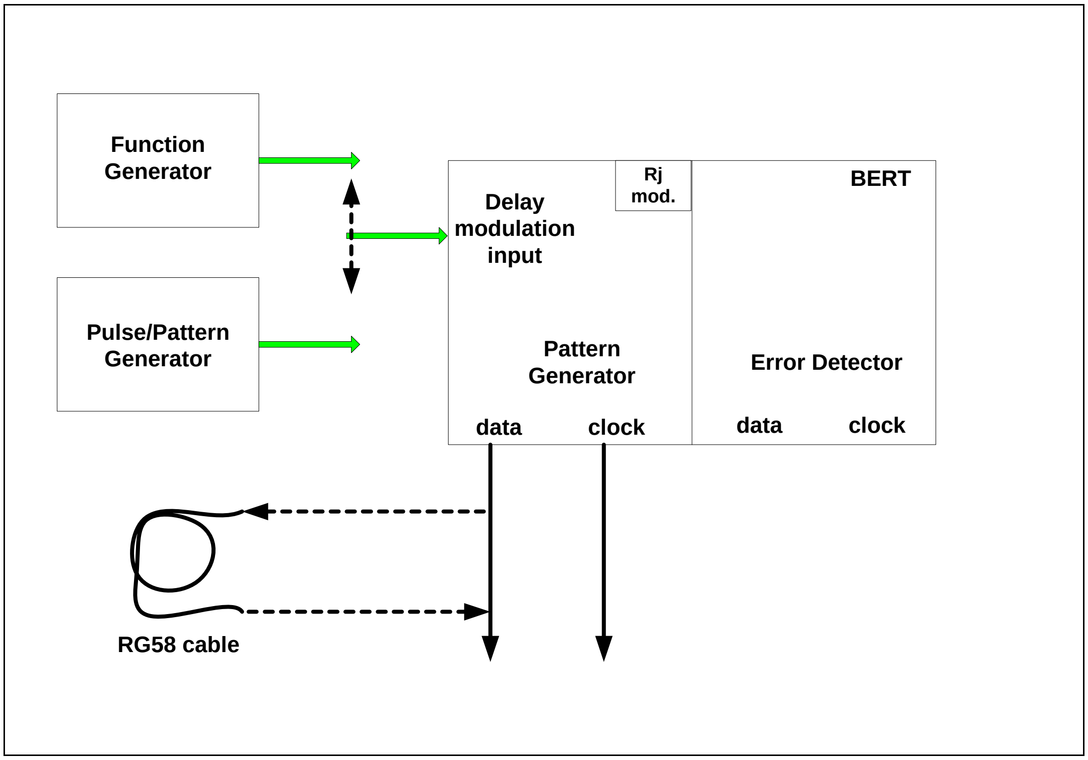

Figure 10: *Schematic of jitter generation system*

The test Instruments (and methods)

- A >20GHz Sampling oscilloscope using normalized Q-scale (NQ-Scale) jitter method.
- The **BERT** is an error detector and data generator pair. The data generator has features for stressed eye testing which include the addition of random, sinusoidal and bounded jitter types. This instrument measures total jitter directly (and estimates Rj and Dj … for those who are interested)
- A >10GHz GHz real-time oscilloscope with jitter measurement software employing the spectral decomposition method.

Test conditions

A few test cases were used to evaluate each instrument and its method. The BERT included a delay modulation input which was used to generate sinusoidal and bounded uncorrelated jitter by applying a sine wave or a PRBS10 data signal. Random jitter was added using the internal ***Rj*** injection feature of the BERT. The measurements from all 3 instruments were compared to the expected jitter level for each case and the differences among them were analyzed. A total of 4 scenarios were tested as follows:

Scenario	Description	Expected Tj	Expected Rj	Expected Dj
1	Baseline (low jitter source)	10.7 ps	0.37 ps	5 ps
2	8 ps pk-pk (sinusoidal) Sj @ 5 MHz	18.17 ps	0.37 ps	13 ps
3	8 ps pk-pk *BUj* @ 100 Mb/s PRBS 10	18.17 ps	0.37 ps	13 ps
4	Rj 10% @ 1E-12 BER	45.37 ps	2.88 ps	5 ps

Table 1: *Description of measurement scenarios: Note the* **Dj** *expected values of 5 ps are due to setup and cabling and possibly the signal source. These values came from extensive calibration of the test setup.*

Results

The results of the 4 cases run on the three instruments are summarized in figures 11 and 12. The total jitter shown in these plots are the primary measurements used for characterizing the performance of serial data links independent of which jitter model is used to predict the total jitter.

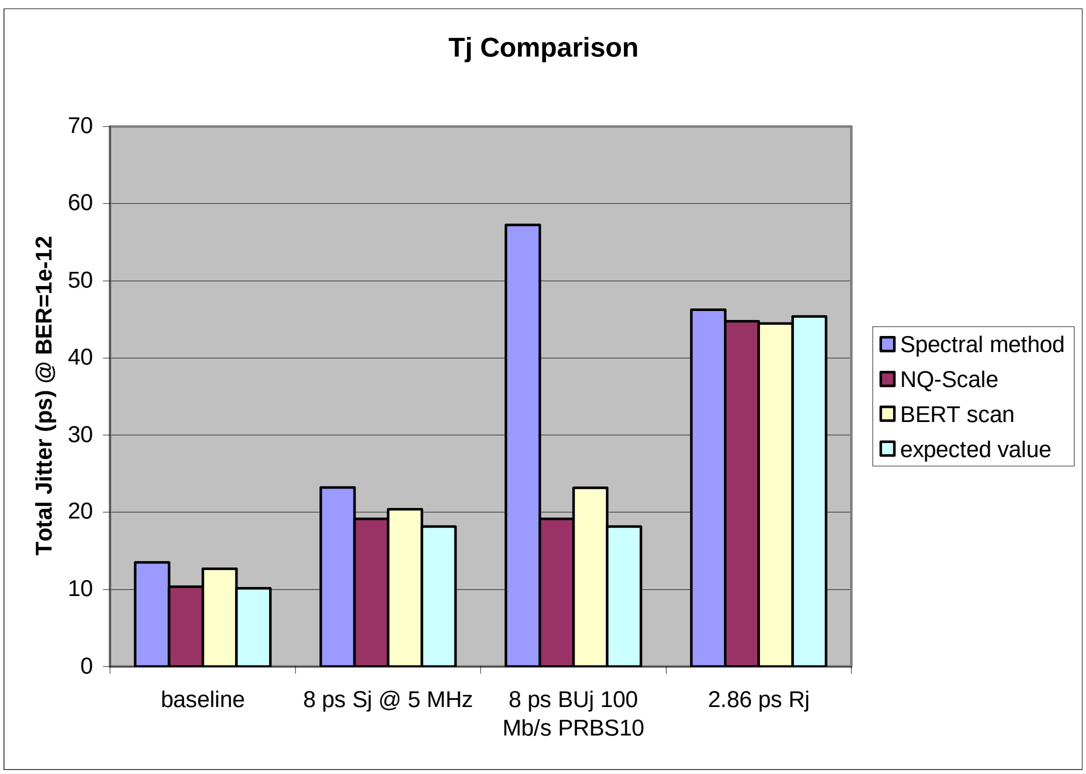

Figure 11: *Measurement results for 4 scenarios and 3 instruments (plus expectations)*

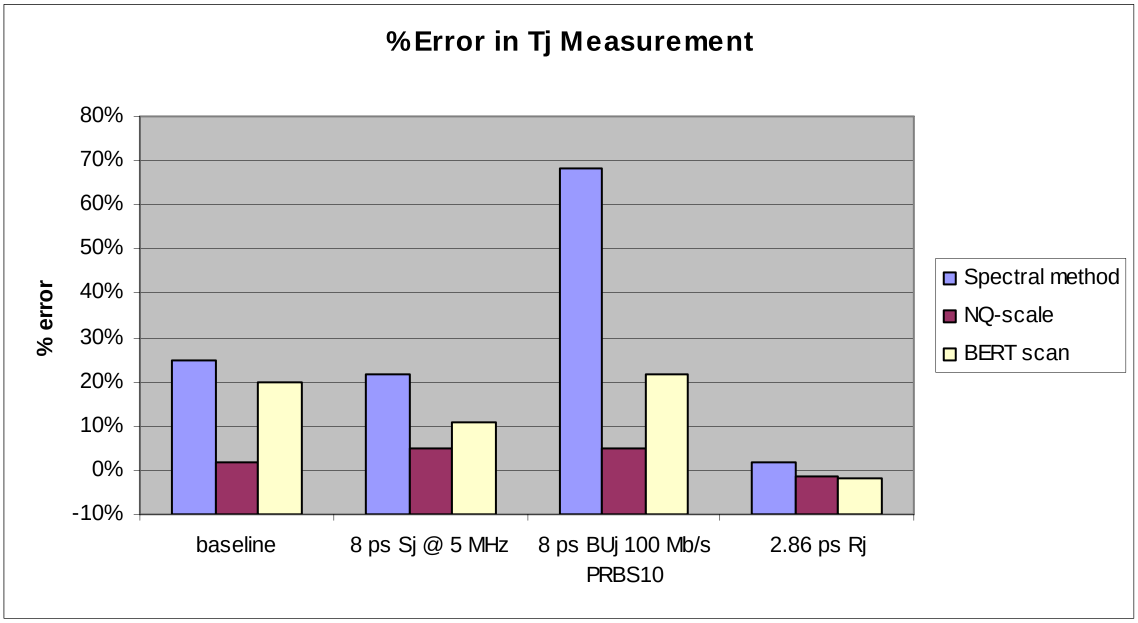

Figure 12: *Measurement results for 4 scenarios and 3 instruments expressed as a percentage of expected value*

As anticipated, the methods yield very similar results, except for the case of wideband (pseudo-random bit stream) aggressors. In this case, the spectral method cannot distinguish between the wideband contribution to the spectrum and any Gaussian noise-floor, and so severely overestimates the ***Rj*** and consequently the estimated ***Tj***. This error when expressed as a percent accentuates error.

Using Monte-Carlo simulations for the two oscilloscope methods (the authors have no simulation software for the BERT), the two methods can be compared for simulated jitter. These results are provided as complement to the measurement data, and confirm a large error for the spectral method when wide-band jitter is present.

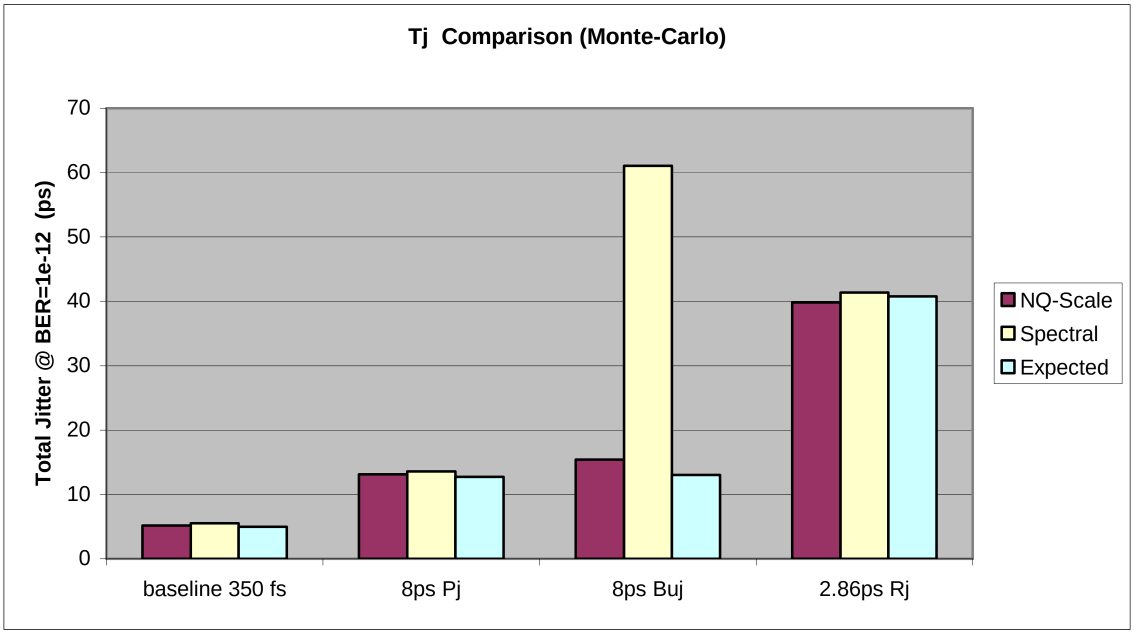

Figure 13: *Monte Carlo results for 4 scenarios and 3 methods (plus expectations)*

The simulations cover the same basic 4 scenarios and confirm the similarity in the methods for cases which contain only ***Rj*** or sinusoidal jitter. The same large overestimation of ***Tj*** is demonstrated for the simulation for injected Bounded Uncorrelated jitter (***BUj***) which is equivalent to wide-band jitter.

Conclusion

The matter of estimating Total Jitter (***Tj***) is not simple and is in general ***assumption laden.*** Direct and predictive methods exist, each with their own package of assumptions and limitations.

Models used to date, are lacking in detail and can produce significant errors in real-world situations. Since no model can predict jitter for every scenario, the model must be chosen to minimize the impact of different jitter types and levels on the overall jitter estimate. The dual-Dirac model is a good compromise but estimates of the random and deterministic parts of this model must be accurate for the models predictions to be valid.

A significant limitation of the spectral methods used in many oscilloscopes to estimate the total jitter is that signals containing crosstalk or other forms of bounded uncorrelated jitter tend to be counted as random (Gaussian) jitter resulting in over estimation of the total jitter. As these tests show, jitter estimates based on the measured statistics of signal timing can be more accurate than the spectral method.

The “normalized Q-Scale” (NQ-Scale) method is proposed as an improvement and extension to the traditional dual-Dirac model. This method yields more accurate estimates of the ***Rj*** and ***Dj*** values , at the same time recognizing that there is an effective “weight” associated with ***Rj*** in the jitter distribution. In addition it also has a useful qualitative result: it confirms or not if Gaussian behavior is observed in the measurement. When this occurs you have a warning that the predicted results are suspect.

The accuracy of jitter estimates is a result of the measurement instrument and the model used to predict the jitter statistics.

While it is generally accepted that BERTs are the most accurate method for measuring total jitter, they are not without their limitations. The vertical noise present in the front end amplifiers of BERTs contributes to the jitter measurement in cases where the eye is stressed by ISI. In these cases the BERT will over estimate the total jitter. Sampling scopes due to their lower vertical noise, on the other hand, give accurate estimates of random jitter even in these stressed eye cases. Real time oscilloscopes have higher noise levels than sampling scopes but the instruments can measure the slew rate of the signal and so the noise can be compensated.

References

[1] Jitter Analysis: The Dual-Dirac Model, RJ/DJ, and Q-scale, Ransom Stephens, December 31, 2004

[2] Anticipating, Characterizing, and Avoiding Problems with Crosstalk, Ransom Stephens and Al Neves, DesignCon 2006

[3] Fibre Channel - Methodologies for Jitter and Signal Quality Specification - MJSQ, T11.2 / Project 1316-DT/ Rev 1 June 9, 2004

[4] A new signal integrity test and measurement architecture for serial 10 Gb/s interfaces. Tom Lindsay, Euro DesignCon 2004

[5] Precision jitter analysis using the Agilent 86100C DCA-J, Agilent Product Note 86100C-1 5989-1146EN June 16, 2004

Appendix A:

Calibrating the jitter generator

In order to better understand the accuracy of each measurement method, it was necessary to understand the expected jitter level being generated for each test case. Three measurements were made manually using a sampling oscilloscope; baseline jitter, the ***Rj*** injection calibration, and the delay vs. peak to peak voltage for the modulation input. The ISI added by the cable was verified by measuring the peak to peak jitter in the eye pattern measured by overlaying all of the bits in the signal waveform using averaging. This method has been referred to as "eye line" and "frame scan".

Rj calibration

The ***Rj*** generator is set in terms of per cent of a unit interval at a 10e-12 bit error rate. The per cent was converted to time and then divided by 14 (the number of peak tp peak standard deviations at 10e-12 for a Gaussian) to arrive at the equivalent RMS ***Rj*** setting. The jitter was measured using the SDA100G for each setting using the ***Rj*** parameter in the jitter menu. The jitter noise floor of the SDA100G was determined by measuring an RF synthesizer whose jitter is known to be well below the noise floor of the oscilloscope. The measured jitter noise floor was then subtracted from the random jitter measured on the BERT pattern generator with all jitter injection features disabled to arrive at the expected baseline random jitter. The random jitter was then measured for ***Rj*** injection levels from 2.6% to 16%. The resulting calibration curve is shown in figure 2. The trend line was used to compute the expected random jitter for each ***Rj*** setting.

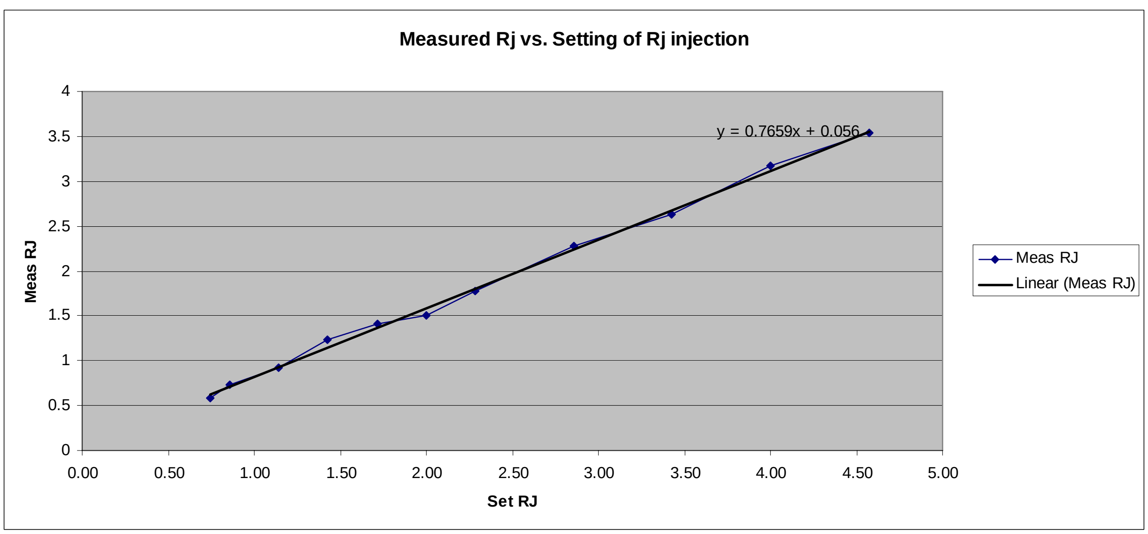

Figure A*1: calibration of random jitter injection*

Sj and *BUj* calibration

Both the sinusoidal and bounded, uncorrelated jitter were added using the delay modulation input of the BERT pattern generator. The input was calibrated by adjusting the DC voltage into the delay control while observing the absolute time shift of the data signal output of the BERT on an oscilloscope. The modulation rates in the test cases were chosen to be low enough so that they remain within the linear range of the modulator calibration curve. The calibration curve for the delay modulation control is shown in figure A2. The trend line gives the expected peak to peak deterministic jitter for each peak to peak voltage setting.

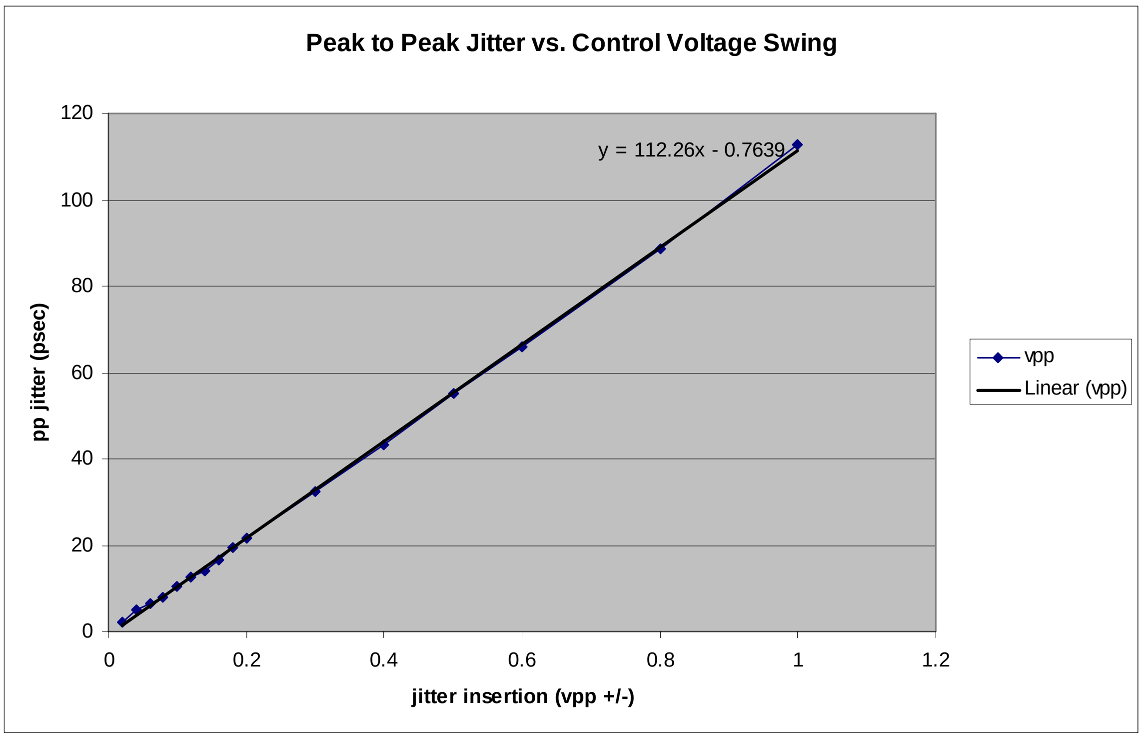

Figure A2: *Calibration curve for Sj and BUj injection*

Appendix B:

Pertinent Related Subjects (with brief comments)

- What effect does instrument and receiver "vertical sensitivity" (input noise) introduce into jitter measurements.

  This is the most important factor in "failure to correlate" issues between instruments and is not a matter of methodology, but rather a problem of AM to PM conversion. If the scope or receiver is noisy, then this imprecision in voltage translates to jitter. The jitter is real if the noise level is that of the real receiver, and is false if its noise is different than the receiver's noise. The manifestation of this sensitivity is almost always Gaussian (so affects ***Rj***) and is proportional to the reciprocal of the rise time multiplied by the RMS vertical noise.

- How Can Induced ISI/DDj (e.g. a cable) Modify ***Rj***?:

  ISI is usually introduced by the "medium" over which the data channel is transmitted and is generally due to a severe reduction in bandwidth. As such in a highly stressed data channel the rise times of the signal

- Crosstalk and other Non-Gaussian Contributors to Jitter Statistics

  Other data channel crosstalk can be either coherent (same clock rate and data pattern) or incoherent, phase unlocked or uncorrelated aggressor data channels). If the crosstalk is incoherent, then it will manifest in the spectrum of jitter as wide-band … and its power will be spread throughout the spectrum. As such the contribution of such aggressors cannot be distinguished from other baseline spectral contributors.

- Periodic Jitter and non-stationary Periodic Jitter (e.g. Spread Spectrum Clocking)

  If periodic or sinusoidal aggressors are not at constant frequency (e.g. FM) they also will have a spread-spectrum and peaks in the resulting analysis of the jitter spectrum will be difficult if not impossible to distinguish from the baseline.