Removing Oscilloscope Noise from RMS Jitter Measurements

Contents

1 Introduction

We present here a new method [1] to accurately measure jitter using a real-time oscilloscope when the level of jitter added to a signal from the measurement environment approaches or exceeds the signal's intrinsic jitter. This method builds on previous work [2] that combined measurement and modeling data to eliminate false spurs in peak-to-peak jitter data. We focus here to eliminate random amplitude noise introduced by the test environment in RMS jitter data.

A test environment can add phase and amplitude noise to a signal under test (SUT). Phase noise modulates a signal's edges directly, whereas amplitude noise converts to phase error during the oscilloscope sampling process. Both cases increase the measured SUT jitter above its true value. Perhaps the dominant source of amplitude noise in a test environment is vertical (quantization) noise in an oscilloscope's sampling system [3]. This can be optimized when setting up the oscilloscope [4], but is always present to some extent. Imperfections in the oscilloscope's interleaving architecture also add amplitude noise, which distorts the measured waveform [5]. Additional sources of amplitude noise from the test environment may include baluns [6], EMI, crosstalk, power-line noise, etc., which can inject noise into the SUT at the PCB or connector level, external to the oscilloscope. Although many of these random-amplitude noise sources can be eliminated using the method below, it cannot remove phase noise introduced by the test environment, such as from an oscilloscope's internal oscillator.

Within the industry, other methods exist to remove an oscilloscope's contribution to its reported RMS jitter values. One method [7] calibrates an oscilloscope's jitter contribution using phase noise data from a reference clock source. Another method [8] uses a SUT's slew rate and an oscilloscope's vertical noise (with the SUT disconnected) to calibrate the oscilloscope's jitter contribution using performance characteristics equations published in its data sheet. After the oscilloscope's jitter contribution determined, it can be removed from the measured SUT jitter using quadrature subtraction.

The below method bears some resemblance to the latter method above, but relies on empirical modeling, rather than data sheet equations, to calibrate jitter contributed by the test environment. We'll introduce this method by way of example, using it to evaluate a PCIe® v4.0 [9] reference clock. This is a practical application since (1) the PCI-SIG® association requires [10] clock jitter analysis using a real-time oscilloscope, (2) each new generation of specifications has lower clock-jitter requirements, (3) the noise floor of oscilloscopes approach or exceed that of today's precision oscillators, (4) the PCI Express® marketplace is cost sensitive, which handicaps over-paying for clock performance, and (5) accurately reporting clock jitter (e.g., without environment noise) can increase the number of solutions offered by clock vendors, which provides more options and flexibility for their customers.

Although the discussion below evaluates PCI-SIG® reference-clocks for time-interval error (TIE) jitter, the method can generally be used to remove (random amplitude) environment noise from any jitter or voltage measurement on clock or data signals.

2 Methodology

The method consists of several key steps, discussed below and illustrated in Figure 1 and Figure 2. Note that certain steps may be performed in one or more ways (not shown) and achieve similar results.

Step 1: Acquire Signal. Power up a device-under-test (DUT) and configure it to output a SUT. For PCIe® v4.0 applications, a compliance load board is connected between the DUT and oscilloscope (with 15 dB loss at 4 GHz, 2 pF termination, and probe connections to the oscilloscope). Setup a real-time oscilloscope to measure jitter [4], and perform one continuous acquisition of the SUT. Apply a noise-reduction filter (typically 2-5 GHz, depending on the edge rate of the signal) to the acquired voltage waveform to remove broadband oscilloscope noise.

Step 2: Acquire Noise. To remove as much jitter from the test environment as possible, power off the DUT and leave it connected to the oscilloscope. To remove oscilloscope-only jitter, power off the DUT, disconnect it from the oscilloscope, and terminate the oscilloscope inputs. Acquire a waveform using the same oscilloscope settings as Step 1. This waveform represents noise from the test environment to remove. Apply the Step 1 filter to this noise waveform.

Step 3: Model of Average Rising Edge. Detect the location of rising edges in the signal waveform. Using the midpoint-crossing voltage (e.g., 0V for a differential signal) of each edge as its origin, create a continuous-time model for a time segment representing the average rising edge. The voltage range of this model only needs to extend far enough such that adding the noise waveform to this location in the model no longer changes the data points near the midpoint of the waveform (which are the only data points involved in computing jitter).

Step 4: Model of Average Falling Edge. Perform Step 3 for falling edges.

Step 5: Model of Ideal Signal. Model a discrete square (or other) wave sampled at the same time interval, and with roughly the same amplitude, as the signal waveform. Use the rising- and falling-edge models to replace points in corresponding edges of the square wave model, such that the average period of the model and signal waveform agree within machine precision (e.g., 16+ digits of floating point precision). This results in an ideal, jitter-free model of the signal (within machine precision), whose 0V transition regions follow the average shape of their respective edges. Only the midpoint (e.g., 0V) crossing regions of this model will be used.

Step 6: Model of Ideal Signal + Noise. Add the noise and ideal model waveforms to create a model-plus-noise waveform.

Steps 7 and 8: Compute Edge Jitter; Jitter Time Trend. Compute time-interval error (TIE) jitter for each edge in the signal waveform as traditionally done, to obtain a jitter time trend for the signal waveform (blue waveform in Figure 1, Step 8). Repeat for the model-plus-noise waveform, using the ideal model waveform as a reference (as illustrated for one rising edge in Figure 1, Step 7; and gray waveform in Figure 1, Step 8).

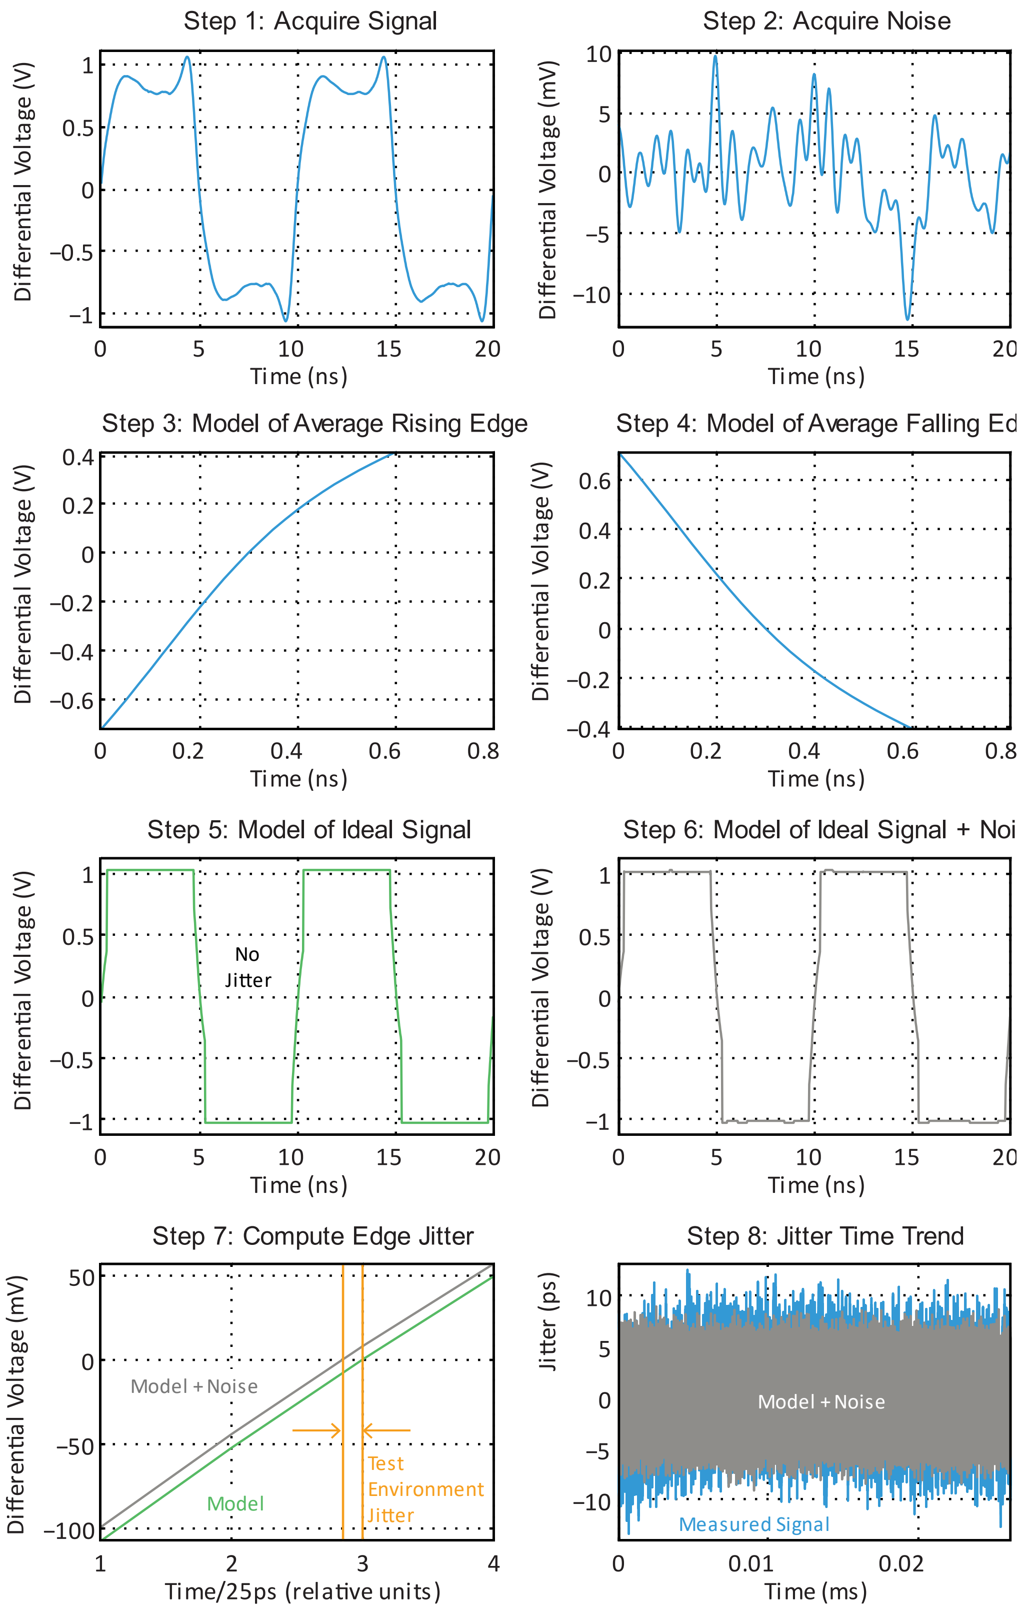

Figure 1: Illustration of key steps to remove environment noise from oscilloscope jitter data

Step 9: Filter Jitter Spectrum. Transform the jitter time trends into the frequency domain and apply the same required jitter filter to each. An example spectrum is shown in Figure 2.

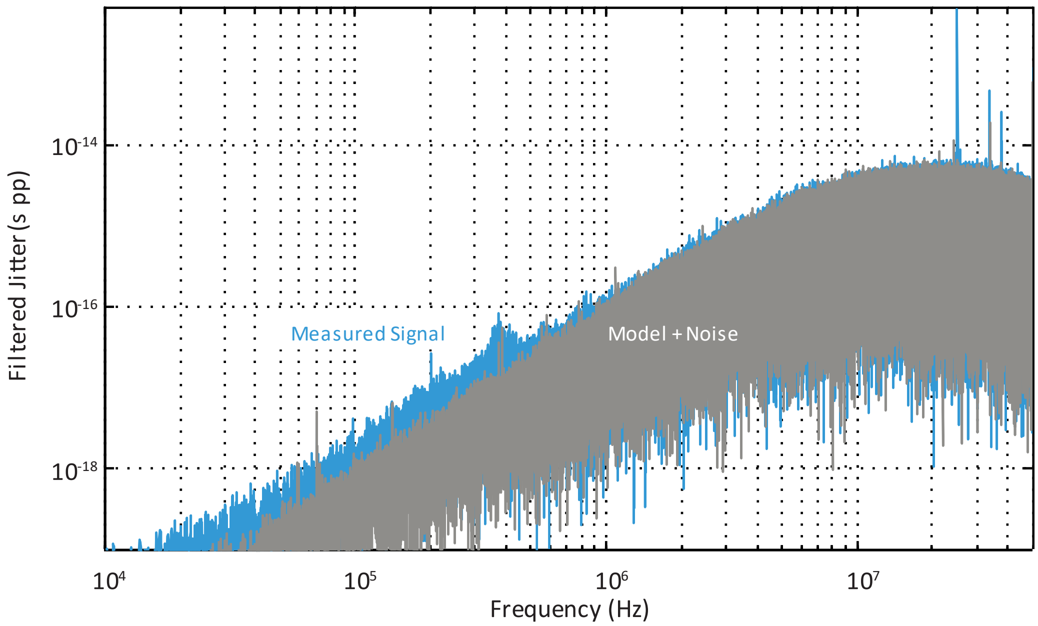

Figure 2: Filtered jitter spectrums for (1) the signal as measured (blue), and (2) an empirical model of an ideal (jitter free) signal with noise added from the test environment (gray)

Step 10: Compute RMS Jitter. Transform the filtered jitter spectrums back into the time domain. Compute the RMS jitter for the filtered signal (J_S), and the filtered model-plus-noise (J_N). Estimate the true intrinsic RMS jitter for the DUT as,

$$J_{DUT} = \sqrt{J_S^2 - J_N^2}$$

Step 11: Compute Error Bars. Compute error bars for the computed DUT jitter. These bars represent the uncertainty in estimating the DUT jitter (JDUT) from the noise removal process.

3 Results

Figure 3 compares jitter results before (JS) and after (JDUT) removing noise from the test environment, for all 64 PCIe® GEN-4 jitter filter combinations. Error bars are drawn on the DUT jitter data to indicate 98% confidence intervals. The worst-case filter combination in Figure 3 is shown in Figure 2 above.

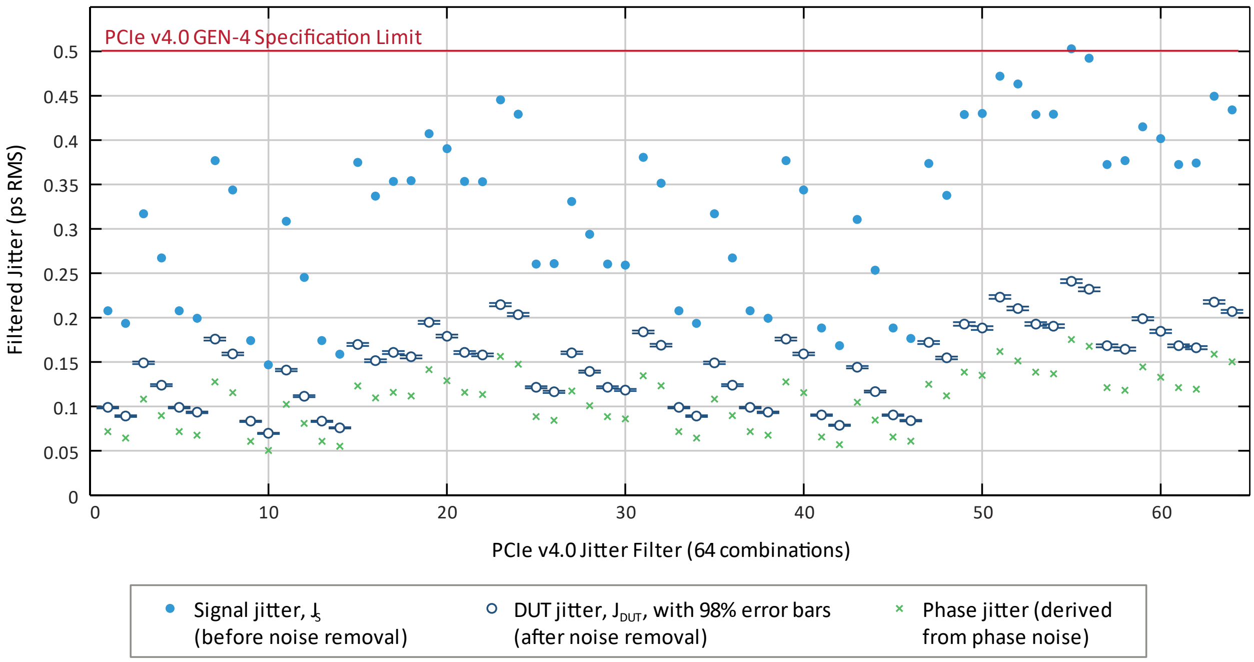

Figure 3: Computed PCIe v4.0 GEN-4 jitter results, comparing before (•) versus after (○) noise removal, for all required filter combinations. For reference, phase jitter (×) is also shown

Using a traditional analysis (e.g., JS), the clock device analyzed above fails the PCIe® v4.0 GEN-4 compliance jitter limit of 0.5 ps RMS. However, after removing jitter introduced by the test environment, the device (e.g., JDUT) easily passes.

The PCI-SIG® association assumes the reference-clock jitter is almost all random [11], and computes RMS jitter without removing spurs (other than GEN-4 SSC spurs). In practice, some amount of deterministic jitter can be present without impacting results much since the RMS calculation is fairly tolerant to outliers. To follow current PCIe® v4.0 practices, the spectrums used to create Figure 3 did not specifically remove these spurs. However, in general, all significant spurs should be removed when computing an RMS value intending to represent random jitter. Alternatively, spurs should be retained when intending to include all statistical components of RMS jitter. The latter scenario is common in telecom markets (e.g., SONET), whereas the former is common in non-telecom markets (e.g., Ethernet).

For reference, the device's phase noise is also measured, filtered, integrated, and plotted as phase jitter in Figure 3. As a sanity check, observe that the DUT jitter (JDUT) is larger than its corresponding phase jitter for each filter combination. In practice, phase jitter serves as a lower-bound for DUT jitter. Compared with phase jitter, the DUT jitter in Figure 3 can be larger due to phase noise from the oscilloscope's sampling clock, and spurious noise (to the extent that spurs are present) as mentioned above.

4 Error Bars

A statistical model and formal analysis of the noise removal process is provided in the Appendix. It assumes the device jitter (JDUT) and environment noise (JN) are independent, uncorrelated, and normally distributed. Some of its interesting results are discussed below.

Figure 4 illustrates how the upper error bar of DUT jitter changes with increasing jitter (JN) from the test environment (note that the lower error bar is nice to plot but plays no role in determining compliance, and so is largely ignored in our analysis). Figure 4 may be viewed, for example, as one of the filter combinations shown in Figure 3.

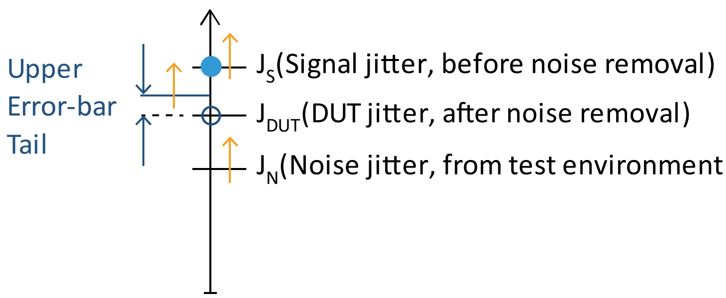

Figure 4: With J_{DUT} fixed, increasing J_N causes the upper error bar (and J_S) to increase

Figure 5 shows that the 98% upper error bar increases with JN, but decreases with larger populations of jitter measurements. To generalize Figure 5, the axes are normalized to the computed DUT jitter (e.g., after noise removal). An example interpretation of chart data is also included in Figure 5. Note that a 98% upper error bar bounds 99% of the total population below it.

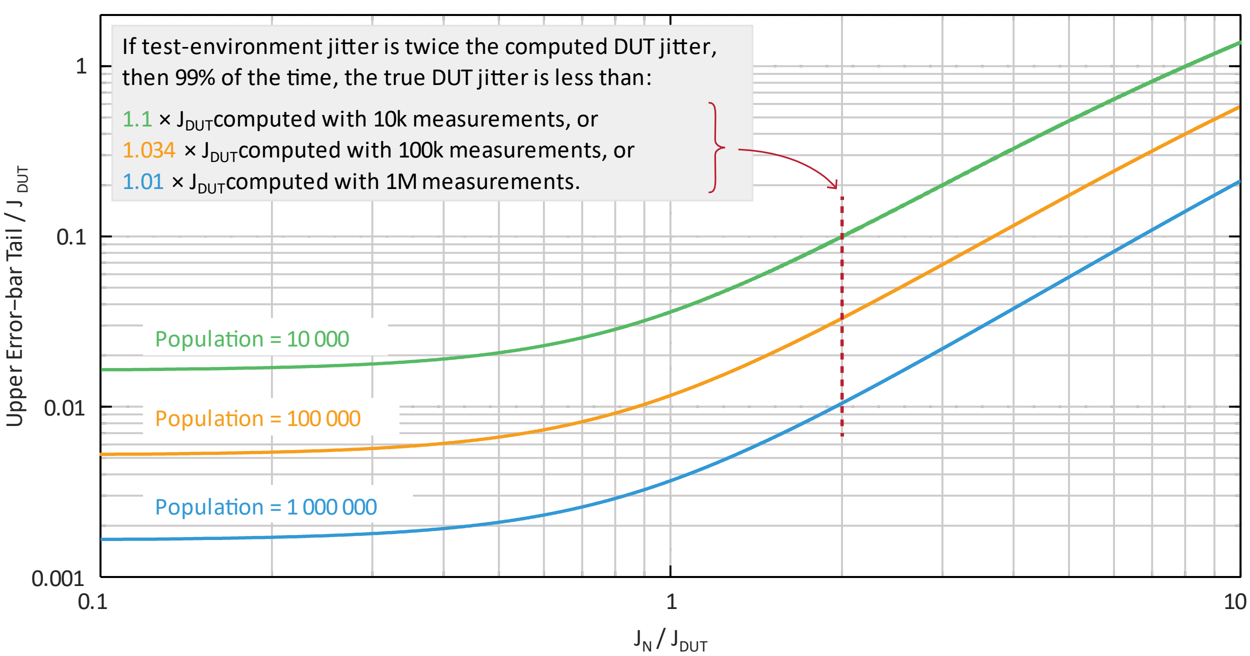

Figure 5: 98% error bars for JDUT can be made insignificant with proper choice of population

In the extreme case when an RMS signal jitter (JS), as traditionally measured using an oscilloscope, is exactly equal to the RMS jitter (JN) added by the test environment (e.g., imagine the gray and blue curves in Figure 2 have the same magnitude), then the DUT jitter computed from quadrature subtraction is 0. Still, the error bars on the computed DUT jitter can remain relatively small if enough samples are analyzed. For example, if JS = JN = 7.3 ps RMS with 1M samples, the 98% upper error bar on the JDUT = 0 data point is positioned at 497 fs RMS, just meeting the PCIe® v4.0 GEN-4 requirement of 0.5 ps RMS. In practice, the jitter introduced by the test environment (e.g., oscilloscope vertical noise) is generally much lower. For example, the worst-case filtered JN in Figure 3 was below 0.5 ps RMS. Thus, the methodology can realistically be applied to noisy environments with high confidence.

5 Conclusion

A method was presented to remove random jitter added by the test environment from a device's measured RMS jitter. The dominant source of environment jitter is often vertical noise in a real-time oscilloscope's sampling system. As such, this method can effectively lower an oscilloscope's random-jitter noise floor. The method itself requires no additional hardware, is spread-spectrum agnostic, is fast and accurate, and requires only a few lines of code to implement. The error bars on the computed RMS jitter values are predictable and can be made arbitrarily small by analyzing larger data sets. Although the above analysis applies to clocks, the same methodology can be applied to data, and other types of jitter (e.g., period jitter, cycle-to-cycle jitter, etc.) and voltage (e.g., eye diagram voltage-margin) measurements.

6 References

[1] Patent pending, "Characterizing a signal in the presence of noise," JitterLabs.

[2] US Patent 8,891,602 B1, "Analyzing jitter with noise from the measurement environment," JitterLabs (Nov. 18, 2014).

[3] "Evaluating Oscilloscope Vertical Noise Characteristics," Keysight Technologies, Application Note 5989-3020EN (2015).

[4] Application Note "AN10073 How to Setup a Real-time Oscilloscope to Measure Jitter," SiTime, , available at AN10073 How to Setup a Real-time Oscilloscope to Measure Jitter | SiTime.

[5] "Evaluating Oscilloscope Sample Rates vs. Sampling Fidelity," Keysight Technologies, Application Note 5989-5732EN (2011).

[6] "Measuring Phase Noise with Baluns," G. Giust, D. Jorgesen, Microwave Journal, v. 59, ed. 10 (October 2016), pp. 84-100.

[7] "Refclk Fanout Best Practices for 8GT/s and 16GT/s Systems," G. Richmond, SiLabs, presented at PCI-SIG Developers Conference (June 7, 2017) in Santa Clara, CA.

[8] N5400, "EZJIT+ Jitter Analysis Software Option," Keysight Technologies, Santa Rosa, CA.

[9] "PCI Express® Base Specification Revision 4.0 Version 0.9," (June 1, 2017), PCI-SIG® association, available at http://www.pcisig.com.

[10] Reference [7], section 8.6.7 states "Jitter measurements shall be made with a ... real time oscilloscope."

[11] Reference [4], section 8.6.7 states for GEN-2 and above, that "these signaling speeds utilize a lower PLL BW and a higher CDR BW, and the effect is to suppress SSC harmonics such that almost all the jitter appears as Rj."

7 Appendix

We have two devices: (1) a signaling device, and (2) a measuring device. Both devices contribute noise to a measured signal. Let t index each replication of the following experiment: the signaling device generates a signal that is measured with the measuring device. Let X_t denote the observed value of the signal (e.g., JS) on replication t, where $t = 1, \ldots, n$. Then

$$X_t = \mu + u_t + v_t \qquad \text{Eq. 1}$$

where

μ = the true value of the signal (e.g., 0, the mean value of jitter),
u_t = a stochastic error term associated with the signaling device (e.g., JDUT),
v_t = a stochastic error term associated with the measurement environment (e.g., JN).

We assume that

$$u_t \sim N(0, \sigma_u^2)$$

$$v_t \sim N(0, \sigma_v^2)$$

and that u_t and v_t are stochastically independent. Also, note that we assume that the signaling device is prepared identically over all n replications, so μ does not depend on t. If, instead, the signal μ depends on t, no inference is possible.

For the signaling device to be compliant to a standard, we must have

$$\sigma_u \leq M \qquad \Leftrightarrow \qquad \sigma_u^2 \leq M^2 \qquad \text{Eq. 2}$$

where M (e.g., the specification limit) is a given positive number.

The n values of X_t allow us to make some inferences about

$$\sigma^2 \equiv Var(u_t + v_t) = \sigma_u^2 + \sigma_v^2 \qquad \text{Eq. 3}$$

but not about σ_u^2 and σ_v^2 separately. However, we may perform m *calibration* experiments with the signaling device powered off. That is, these observations are modeled as

$$Y_t = v_t$$

for $t = 1, \ldots, m$. These calibration experiments allow us to form an estimate of σ_v^2. Let S^2 and S_v^2 denote the estimates of σ^2 and σ_v^2, respectively. Then we may estimate σ_u^2 by

$$S_u^2 \equiv S^2 - S_v^2 \qquad \text{Eq. 4}$$

For statistical inference, we need to find the "standard error" of this estimator.

Typically, $n = m = 10^5$ or higher. We therefore assume that both n and m are "large," in the sense that various asymptotic statistical approximations are acceptable.

From the random sample $(X_1, \ldots, X_n)$ we compute the following statistics.

$$\bar{X} \equiv \frac{1}{n}\sum_{t=1}^{n} X_t$$

$$S^2 \equiv \frac{1}{n-1}\sum_{t=1}^{n}(X_t - \bar{X})^2.$$

Then $\bar{X}$ is an unbiased estimator of μ, and S^2 is an unbiased estimator of σ^2. We know under these assumptions that

$$\frac{(n-1)S^2}{\sigma^2} \sim \chi^2(n-1), \qquad \text{Eq. 5}$$

and hence that

$$Var(S^2) = \frac{2\sigma^4}{n-1}\ . \qquad \text{Eq. 6}$$

Similarly, from $(Y_1, \ldots, Y_m)$ we compute

$$S_v^2 \equiv \frac{1}{m}\sum_{t=1}^{m} Y_t^2, \qquad \text{Eq. 7}$$

which is an unbiased estimator of σ_v^2. Under these assumptions

$$\frac{mS_v^2}{\sigma_v^2} \sim \chi^2(m) \qquad \text{Eq. 8}$$

and hence

$$Var(S_v^2) = \frac{2\sigma_v^4}{m}. \qquad \text{Eq. 9}$$

Combining these facts, we see that S_v^2 given by Eq. 4 is an unbiased estimator of σ_u^2 and that

$$Var(S_u^2) = \frac{2\sigma^4}{n-1} + \frac{2\sigma_v^4}{m}. \qquad \text{Eq. 10}$$

Substituting estimates of σ^2 and σ_v^2 into this equation gives us a formula for an estimate of the variance of S_u^2:

$$Est.Var(S_u^2) = \frac{2(S^2)^2}{n-1} + \frac{2(S_v^2)^2}{m}. \qquad \text{Eq. 11}$$

Taking the square root of this expression gives us an estimate of the "standard error" of S_u^2:

$$r \equiv \sqrt{Est.Var(S_u^2)} = \sqrt{\frac{2(S^2)^2}{n-1} + \frac{2(S_v^2)^2}{m}}. \qquad \text{Eq. 12}$$

Under the assumptions above, the Central Limit Theorem implies that S_u^2 is distributed approximately as $N(\sigma_u^2, r^2)$, and hence that the standardized random variable

$$\frac{S_u^2 - \sigma_u^2}{r}$$

is distributed approximately as a standard normal.

These results allow us to test the compliance of the signaling device to the standard specified by Eq. 2. Under the given assumptions, $\sigma_u \leq M$ holds if and only if $\sigma_u^2 \leq M^2$, so we may test for compliance by testing the null hypothesis

$$H_0: \sigma_u^2 > M^2. \qquad \text{Eq. 13}$$

Let

$$p \equiv Pr(H_0), \qquad \text{Eq. 14}$$

so the probability that the device is compliant equals $1 - p$. If p is very small, then we may reject the null hypothesis that the device is not compliant at the p level of significance.

If S_u^2, the estimate of σ_u^2, is greater than or equal to M^2, then we have no reason to reject H_0. If S_u^2 is less than M^2, we may evaluate the statistical significance of this fact by the test statistic

$$T \equiv \frac{M^2 - S_u^2}{r} \qquad \text{Eq. 15}$$

which measures the distance between M^2 and S_u^2 in units of standard error r. In these terms,

$$p \approx Pr(Z \geq T) = 1 - \Phi(T), \qquad \text{Eq. 16}$$

where Z denotes a standard normal random variable and $\Phi()$ denotes the cumulative distribution function of a standard normal random variable.

Example 1. Suppose that

$$n = 10^6, \qquad m = 10^6,$$

$$S^2 = 1.9, \qquad S_v^2 = 1.8.$$

Hence, from Eq. 4 and Eq. 11,

$$S_u^2 = 0.1,$$

$$Est.Var(S_u^2) \approx 1.37 \times 10^{-5},$$

and the later equation implies that

$$r \approx 0.003701.$$

Now suppose that $M^2 \equiv 0.11$, i.e., just slightly larger than S_u^2. The question is: is M^2 *significantly* larger than S_u^2? The test statistic T with this data is

$$T = \frac{0.11 - 0.10}{r} \approx 2.7017,$$

which implies that

$$p \approx Pr(\sigma_u^2 > M^2) \approx 1 - \Phi(T) \approx 0.003449,$$

so we may reject the null hypothesis of non-compliance at the 0.003449 level of significance. In other words, the probability that the signaling device is compliant with $M^2 = 0.11$ is

$$Pr(\sigma_u^2 \leq M^2) \approx \Phi(T) \approx 0.996551.$$

Given the logical equivalence

$$\sigma_u \leq M \qquad \Leftrightarrow \qquad \sigma_u^2 \leq M^2,$$

we see that these results may be rewritten in terms of σ_u:

$$Pr(\sigma_u > M) \approx 0.003449 \text{ and } Pr(\sigma_u \leq M) \approx 0.996551.$$

Instead of approaching statistical inference from a hypothesis testing viewpoint, we might choose to construct appropriate confidence intervals for σ_u^2 and for σ_u. We begin with the construction of a "two-sided" confidence interval for σ_u^2 that is centered at S_u^2. As noted above, the random variable

$$\frac{S_u^2 - \sigma_u^2}{r} \qquad \text{Eq. 17}$$

is distributed approximately as a standard normal. Let α be a (small) number in the range $(0,1)$. Typical values are 0.05 or 0.01. Given α, define z_α as the solution to

$$1 - \alpha = \Phi(z_\alpha). \qquad \text{Eq. 18}$$

For example, $z_{0.05} \approx 1.64485$, $z_{0.025} \approx 1.96$, and $z_{0.01} \approx 2.32635$. Combining these facts, we see that

$$Pr\left(-z_\alpha \leq \frac{S_u^2 - \sigma_u^2}{r} \leq z_\alpha\right) \approx 1 - 2\alpha. \qquad \text{Eq. 19}$$

As the double inequality

$$-z_\alpha \leq \frac{S_u^2 - \sigma_u^2}{r} \leq z_\alpha$$

may be rewritten as

$$S_u^2 - rz_\alpha \leq \sigma_u^2 \leq S_u^2 + rz_\alpha,$$

it follows that

$$Pr(L^2 \leq \sigma_u^2 \leq U^2) \approx 1 - 2\alpha$$

where

$$L^2 \equiv S_u^2 - rz_\alpha \text{ and } U^2 \equiv S_u^2 + rz_\alpha. \qquad \text{Eq. 20}$$

Hence, $[L^2, U^2]$ is a $100(1 - 2\alpha)$% confidence interval for σ_u^2.

Example 2. Let $\alpha = 0.05$. Given the same data as in Example 1, we find $U^2 \approx 0.106088$ and $L^2 \approx 0.093912$. Hence, $[0.093912, 0.106088]$ is a 90% confidence interval for σ_u^2.

Let's call the confidence interval specified by Eq. 20 the **primary** interval. From the primary interval, we may derive several other confidence intervals.

(1) Let L^2 and U^2 be given by Eq. 20. As

$$L^2 \leq \sigma_u^2 \leq U^2 \quad \text{if and only if} \quad L \leq \sigma_u \leq U,$$

it follows that $[L, U]$ is a $100(1 - 2\alpha)$% confidence interval for σ_u. For example, using the same data as in Example 2, we find that $[L, U] = [0.30645, 0.32571]$ is a 90% confidence interval for σ_u. Note. While S_u^2 is midway between L^2 and U^2 in the primary interval, it is *not* true that S_u is midway between L and U. In fact, it may be shown that

$$S_u > \frac{1}{2}(L + U) \quad \Leftrightarrow \quad S_u - L > U - S_u.$$

For example, in the interval $[L, U] = [0.30645, 0.32571]$ computed above, we find $S_u \approx 0.31623$, so

$$S_u - L \approx 9.777 \times 10^{-3},$$

$$U - S_u \approx 9.484 \times 10^{-3}.$$

(2) There are two ways the double inequality

$$L^2 \leq \sigma_u^2 \leq U^2$$

may fail to hold: we could have either

$$\sigma_u^2 < L^2 \quad \text{or} \quad \sigma_u^2 > U^2.$$

Because of the symmetry of a normal probability density function around its mean, these two possibilities are equally likely; in fact, by the way L^2 and U^2 are defined,

$$Pr(\sigma_u^2 < L^2) = Pr(\sigma_u^2 > U^2) = \alpha.$$

Hence, we may create a one-sided $100(1-\alpha)\%$ confidence interval for σ_u^2 by eliminating one of the two tails. In particular, $(0, U^2]$ is such an interval:

$$Pr(\sigma_u^2 \leq U^2) = 1 - \alpha.$$

For example, using the same data as in Examples 1 and 2, we find that $(0, 0.106088]$ is a 95% confidence interval for σ_u^2.

(3) Combining the ideas behind the two derived intervals given above, we find that $(0, U]$ is a $100(1-\alpha)\%$ confidence interval for σ_u:

$$Pr(\sigma_u \leq U) = 1 - \alpha.$$

For example, with the same data as used above, $(0, 0.32571]$ is a 95% confidence interval for σ_u.

Table 1: Revision History

Version	Release Date	Change Summary
1.0	26-Jul-2016	Initial Release JitterLabs LLC
1.0	30-Mar-2021	Initial Release SiTime with various edits and updates

SiTime Corporation, 5451 Patrick Henry Drive, Santa Clara, CA 95054, USA | **Phone:** +1-408-328-4400 | **Fax:** +1-408-328-4439