

On Modern and Historical Short-Term Frequency Stability Metrics for Frequency Sources

Michael S. McCorquodale
Mobius Microsystems, Inc.
Sunnyvale, CA USA 94086
mccorquodale@mobiusmicro.com

Richard B. Brown
College of Engineering, University of Utah
Salt Lake City, UT USA 84112
brown@utah.edu

Abstract—Modern frequency stability metrics employed in industry are discussed within the context of the historical development of the Allan variance. It is shown that phase noise and the first and second jitter differences are often reported in lieu of the Allan variance for frequency sources utilized in consumer electronics. This inconsistency between the historical and modern treatment of frequency stability is addressed through the presentation of analytical relationships and bounds for these common metrics. The presented expressions are compared with experimental results where good agreement is found while employing several different measurement approaches. The clear requirement for standardized metrology is motivated.

I. INTRODUCTION

The Allan variance was introduced over three decades ago to address the fact that the classical variance does not converge for common noise processes when represented by a power-law spectrum [1]. Outside of the frequency control community, much recent work (such as that in [2]) has treated phase noise and jitter in solid-state oscillators where expressions have been derived to relate phase noise to timing variance, though they have not been reconciled with measurement. In modern frequency control devices, phase noise and the first and second jitter differences are reported in lieu of the Allan variance.

This trend has emerged because the classical variance is simpler to measure, intuitive and broadly useful, though different measurement techniques exist, including real-time and sampling approaches. Additionally, the classical variance does, in fact, converge utilizing these modern measurement techniques even in presence of pink noise. In a large segment of the frequency control industry today, there is an inconsistency between the manner in which short-term frequency stability is treated and the historical practice.

Further, several analytical expressions have not been reconciled with measurement and there is a lack of consistent results when employing different measurement approaches for timing jitter. Building on related work in [3], in this work we present the relationships between the phase noise and the common jitter metrics employed in industry. Bounds are discussed along with their physical interpretations in the context of the underlying noise process. These analytical expressions are confirmed by measurement of crystal and solid-state oscillators utilizing industry-standard instrumentation. Measurement nuances and techniques are discussed and consistent results are demonstrated.




Figure 1. The time-domain voltage output of an oscillator, $v(t)$, with period T_o : (a) ideal (b) with noise. The discrete sequence $\{t_k\}$ represents the time-stamps of the positive zero-crossing transitions of the signal. The discrete sequence of periods is given by $\{T_k := t_{k+1} - t_k\}$.

II. HISTORICAL TREATMENT OF SHORT-TERM FREQUENCY STABILITY

A. Oscillator Modeling

The time-domain voltage output of an ideal oscillator is shown in Fig. 1(a) and can be represented mathematically by $v(t) = V_o \sin \omega_o t$ where V_o is the oscillation amplitude, ω_o is the radian frequency, $T_o = 2\pi/\omega_o$ is the ideal period and t is time. The same oscillator, under the influence of noise, is shown in Fig. 1(b) and can be modeled mathematically by $v_n(t) = (V_o + \varepsilon(t)) \sin(\omega_o t + \phi(t))$ where $\varepsilon(t)$ and $\phi(t)$ represent amplitude and phase noise respectively. In this work, and as in Fig. 1(b), only phase noise is considered.

As shown in Fig. 1, the discrete sequence, $\{t_k\}$, represents the time-stamps of the positive zero-crossing transitions of the signal. The corresponding discrete sequence of periods is defined by $\{T_k := t_{k+1} - t_k\}$. With this sequence, the first jitter difference is defined as,

$$\sigma_T = \sqrt{\text{var}(T_k)}. \quad (1)$$

σ_T is often referred to as the period or one-cycle jitter and it is the root of the classical variance. Similarly, the second jitter difference, or cycle-to-cycle jitter, is defined as,

$$\sigma_{\Delta T} = \sqrt{\text{var}(T_{k+1} - T_k)}. \quad (2)$$




Figure 2. Fourier-domain spectrum of an oscillator: (a) ideal (b) with noise

Historically, the instantaneous radian frequency, $\omega(t)$, is considered and given by,

$$\omega(t) = \frac{d}{dt}(\omega_o t + \phi(t)) = \omega_o + \frac{d}{dt}\phi(t). \quad (3)$$

The frequency deviation of the oscillator is defined by $\Delta\omega(t) = \omega(t) - \omega_o$ and the fractional frequency deviation is the normalized instantaneous deviation: $y(t) = \Delta\omega(t)/\omega_o$.

The oscillator may also be characterized in the Fourier-domain. As shown in Fig. 2(a), the power spectrum of the ideal oscillator is a Dirac delta function at the fundamental frequency, $f_o = \omega_o/2\pi$, while the noisy oscillator exhibits phase noise sidebands as shown in Fig. 2(b).

Historically, the single sideband (SSB) spectrum of either the phase noise process, $\phi(t)$, or the normalized frequency deviation, $y(t)$, is considered and given by $S_\phi(f_m)$ and $S_y(f_m)$ respectively. Both may be represented by a power-law model as shown in Fig. 3 and given by:

$$S_\phi(f_m) = \sum_{\alpha=-4}^0 b_\alpha f_m^\alpha = \frac{f_o^2}{f_m^2} \sum_{\alpha=-2}^2 h_\alpha f_m^\alpha \quad (4)$$

$S_\phi(f_m)$ and $S_y(f_m)$ are related by: $S_\phi(f_m) = (f_o^2/f_m^2)S_y(f_m)$. In this work, $S_\phi(f_m)$ will be used because it is easily related to the SSB phase noise power spectral density (PSD), which is commonly measured with standard instrumentation. The relationship between the two and the measurement is given by:

$$\left(\frac{N_o}{P_o}\right)_{f_m} = \frac{1}{2}S_\phi(f_m) = \frac{4S_{v_n}(f_o + f_m)}{V_o^2}, \quad (5)$$




Figure 3. The power-law phase noise model with five distinct regions.

where N_o and P_o are the noise and signal power respectively, f_m is the frequency offset from the fundamental and S_{v_n} is the double sideband spectral density of the oscillator waveform. In typical oscillators, the SSB phase noise PSD is dominated by white of frequency and flicker of frequency phase noise which exhibit f_m^{-2} and f_m^{-3} relationships in $S_\phi(f_m)$ respectively. These regions of $S_\phi(f_m)$ are considered next.

B. Application of the Model and a Mathematical Pathology

The spectrum of $S_\phi(f_m)$ can be related to the first and second jitter differences by the application of the Wiener-Khinchin theorem [2]-[4] assuming that $S_\phi(f_m)$ is stationary and integrable. The relationships are given by:

$$\sigma_T^2 = \frac{4}{\omega_o^2} \int_0^\infty S_\phi(f_m) \sin^2(\pi f_m T_o) df_m, \quad (6)$$

$$\sigma_{\Delta T}^2 = \frac{16}{\omega_o^2} \int_0^\infty S_\phi(f_m) \sin^4(\pi f_m T_o) df_m. \quad (7)$$

Consider the specific case of white of frequency noise defined by $S_\phi(f_m) = h_0(f_o^2/f_m^2)$ and substitute into (6) and (7):

$$\sigma_T^2 = \frac{4}{\omega_o^2} \int_0^\infty h_0 \frac{f_o^2}{f_m^2} \sin^2(\pi f_m T_o) df_m = \frac{h_0 T_o}{2}, \quad (8)$$

$$\sigma_{\Delta T}^2 = \frac{16}{\omega_o^2} \int_0^\infty h_0 \frac{f_o^2}{f_m^2} \sin^4(\pi f_m T_o) df_m = h_0 T_o. \quad (9)$$

As shown, both the first and second jitter differences have well-defined and tractable solutions.

Next consider the specific case of flicker of frequency noise defined by $S_\phi(f_m) = h_{-1}(f_o^2/f_m^3)$ and substitute again:

$$\sigma_T^2 = \frac{4}{\omega_o^2} \int_0^\infty h_{-1} \frac{f_o^2}{f_m^3} \sin^2(\pi f_m T_o) df_m = \infty, \quad (10)$$

$$\sigma_{\Delta T}^2 = \frac{16}{\omega_o^2} \int_0^\infty h_{-1} \frac{f_o^2}{f_m^3} \sin^4(\pi f_m T_o) df_m = 4h_{-1} \ln(2) T_o^2. \quad (11)$$

Here the first jitter difference, or the classical variance, does not converge. This mathematical pathology is illustrated in Fig. 4 where $\sin^2(f_m)/f_m^3$ diverges as f_m approaches zero. However, the second jitter difference is integrable because $\sin^4(f_m)/f_m^3$ is finite for all f_m . Interestingly, if a high-pass filter is applied to $S_\phi(f_m)$, then divergence of the classical variance will not be observed in the presence of pink noise. This concept will be revisited again later in this manuscript.

The Allan variance was introduced to address this pathology. It is computed by measuring the frequency of the signal under test repeatedly, and without dead-time between samples, over the interval τ . The fractional frequency error, y , is computed for each sample and the ensemble average of the square of the difference between samples yields the Allan




Figure 4. A mathematical pathology. The function converting the spectrum of S_ϕ to the first jitter difference (root of the classical variance) is unbounded as f_m approaches zero. The second jitter difference is well-defined.

variance. In practice, the sample count is finite; thus, for N samples, the Allan variance is estimated by:

$$\sigma_y^2(\tau) = \frac{1}{2} \langle (\bar{y}_2 - \bar{y}_1)^2 \rangle \approx \frac{1}{2(N-1)} \sum_{k=1}^{N-1} (\bar{y}_{k+1} - \bar{y}_k)^2. \quad (12)$$

The root of (12) is the Allan deviation. In [5], the relationship between the Fourier-domain representation of $S_y(f_m)$ and time-domain Allan variance was determined and is given by:

$$\sigma_y^2(\tau) = 2 \int_0^\infty S_y(f_m) \frac{\sin^4(\pi f_m \tau)}{(\pi f_m \tau)^2} df_m. \quad (13)$$




Figure 5. Visualizing the Allan deviation for varying sample intervals, τ . (a) A short sample interval dominated by white of frequency noise. (b) A long sample interval dominated by flicker of frequency noise. (c) The resulting Allan deviation against sample time where white of frequency noise decreases as τ increases until the flicker of frequency floor is reached.

Referring to (13) and Fig. 4, it is clear that the Allan variance is not pathological in the case of flicker of frequency noise because of the \sin^4 function.

The behavior of the Allan deviation is illustrated in Fig. 5 for the case of white and flicker of frequency noise only. In Fig. 5(a), τ is small and the deviation is larger as high frequency white noise is captured. However, in Fig. 5(b) τ is larger and the white noise has been averaged. Fig. 5(c) illustrates that the Allan deviation saturates to a floor determined by the flicker of frequency noise.

III. MODERN TREATMENT OF SHORT-TERM FREQUENCY STABILITY

Despite this elegant and tractable historical treatment of short-term frequency instability, the Allan variance is not utilized in large segments of the frequency control industry. Instead, and particularly in applications including wireline interfaces for consumer electronics, the first and second jitter differences and phase noise are reported despite the fact that the first jitter difference does not converge theoretically. Moreover, the first jitter difference is often extrapolated to estimate the bit error rate (BER) for wireline communication links. Fig. 6 illustrates a typical eye-diagram and the extrapolated bathtub curve. In such a measurement, the signal must not cross into the template keep-out region during the measurement interval. Then, once the first jitter difference has converged, it is extrapolated to estimate the BER across the unit interval (UI).

Similarly, it is often of interest to estimate the peak-to-peak jitter for a frequency source or wireline link. The distribution of the periods, $\{T_k\}$, is unbounded, thus the peak is not defined as illustrated in Fig. 7. However, an estimate of the peak jitter can be determined for a bounding cycle count, N , and by applying probability measure. Specifically, assume that $\{T_k\}$ exhibits a Gaussian distribution with a well-defined standard deviation, σ . The peak jitter is determined by solving (14) for α and scaling σ by α . The resulting computation indicates that no more than one in N cycles will exceed the peak jitter determined by $\alpha\sigma$.

$$\frac{1}{2} \operatorname{erfc}\left(\frac{\alpha}{2\sqrt{2}}\right) = \frac{1}{N}, \text{ where } \operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^\infty e^{-x^2} dx. \quad (14)$$




Figure 6. An illustration of a typical eye-diagram and bathtub curve measurement. The signal must not cross into the template keep-out region over the measurement interval. Once the first jitter difference has converged, the BER across the unit interval (UI) is estimated by extrapolation




Figure 7. The distribution of $\{T_k\}$ is unbounded. However, peak-to-peak jitter can be bounded for a finite cycle count and estimated by scaling the root of the classical variance, or the standard deviation.

For these, and many other related applications, the root of the classical variance is commonly employed in industry. The motivation for using this metric is that it is intuitive, and it is useful for the extrapolation of unbounded processes, which would otherwise require an inordinant amount of time to observe. Considering this, bounds and relationships for these commonly employed metrics are considered next.

IV. BOUNDS AND RELATIONSHIPS

In [3], the upper bound relating σ_T and $\sigma_{\Delta T}$ was derived. Those results were achieved by considering the expressions in (6) and (7) along with trigonometric identities. The resulting relationship is:

$$\sigma_{\Delta T}^2 = 4\sigma_T^2 - \sigma_{2T}^2 \quad (15)$$

Next consider that variances are positive with a lower bound of zero. Using (15), the following relationships are derived:

$$\sigma_{\Delta T}^2 \leq 4\sigma_T^2 - 0, \text{ thus } \frac{\sigma_{\Delta T}}{\sigma_T} \leq 2 \quad (16)$$

$$0 \leq 4\sigma_T^2 - \sigma_{2T}^2, \text{ thus } \frac{\sigma_{2T}}{\sigma_T} \leq 2 \quad (17)$$

$$\sigma_{\Delta T}^2 \geq 4\sigma_T^2 - 2\sigma_T^2, \text{ thus } \frac{\sigma_{\Delta T}}{\sigma_T} \geq \sqrt{2} \quad (18)$$

Therefore, the second jitter difference is related to the first jitter difference by both an upper and lower bound given by:

$$\sqrt{2}\sigma_T \leq \sigma_{\Delta T} \leq 2\sigma_T. \quad (19)$$

Identical results can be determined by definition. Expanding the expression in (2) yields:

$$\sigma_{\Delta T}^2 = \text{var}(T_{k+1}) + \text{var}(T_k) - 2\text{cov}(T_{k+1}, T_k). \quad (20)$$

If all T_k are independent and identically distributed, the covariance is zero and (20) reduces to the lower bound in (19). Interestingly, the upper bound in (19) is achieved when the covariance in (20) is -1.

Next consider the solutions to the Fourier-to-time domain transformation for the case of white noise in (8) and (9). The ratio of these expressions is identical to the lower bound in (19). Therefore, if the underlying phase noise process is white, the ratio of $\sigma_{\Delta T}$ to σ_T is $\sqrt{2}$. However, due to the pathology in

(10), the same treatment cannot be applied for the case of flicker of frequency noise. Nevertheless, and returning to the expanded definition in (20) and the bound in (19), the upper bound is approached as the covariance approaches -1. This indicates that the underlying noise process exhibits memory, which is well-known to be the case for flicker noise processes.

Lastly, the SSB phase noise PSD can be related to both the first and second jitter difference using (5)–(7). With these bounds and relationships, experimental results are considered next.

V. EXPERIMENTAL RESULTS

A. Frequency Control Devices Under Test

Free-running and self-referenced solid-state oscillators (SSO) have been introduced recently by the authors where representative examples include [4] and [6]. It has been shown that devices achieve performance comparable to crystal oscillators (XOs) [4]. For the purpose of experimentation, a free-running 24MHz SSO and XO were placed in a temperature-controlled chamber at 25°C to minimize the influences of the native temperature coefficient.

B. Integrated SSB Phase Noise PSD

The SSB phase noise PSD was measured using a signal source analyzer for both the XO and SSO. The upper limit of the offset frequency from the fundamental was 5MHz. The data were exported from the instrument and imported into a math CAD program. Here the far-from-carrier phase noise was extrapolated from the measurement at the last offset frequency (5MHz). Next, the \sin^2 mask in (6) was generated and projected against the measurement. These three curves for the two devices under test (DUTs) are shown in Fig. 8. The first jitter difference was computed by numerically integrating the result of the projection of the data onto the trigonometric mask for several different upper limits of the integral in (6). Results are summarized in Table I. Here, the limit $f_o/2$ represents the first peak of the \sin^2 mask while the limit f_o corresponds to the first null. Results indicate that the majority of the computed period jitter is captured over this first interval. The additional upper limits are even offset frequencies and correspond to the odd harmonics of the signal under test. These results merely accumulate the white noise floor. For the case of the SSO, it appears that the phase noise had not reached the floor by the maximum offset frequency (5MHz). Thus, the rate of accumulation is greater than for the XO for the higher integration limits.

C. Real-Time Approach

Both σ_T and $\sigma_{\Delta T}$ were measured using a 20GSa/s real-time oscilloscope (RTO) with a 6GHz analog bandwidth. In this approach, a high-speed A/D converter captures a real-time record of the waveform up to the full memory depth of the instrument. Additional data captures can be performed to add to the database for statistical analysis. The waveform is constructed by interpolating between samples and a variety of statistics can be computed from the resulting data set. However, even at a sample rate of 20GSa/s, the sample spacing is 50ps while the expected values of σ_T and $\sigma_{\Delta T}$ are below 5ps based on results from the integrated SSB phase noise PSD. Thus, the

technique is highly sensitive to interpolation between samples. Consequently, the instrument must be configured to utilize the entire dynamic range of the front-end A/D. Additionally, interpolation between samples must be performed over a wide and linear transition region of the signal where anomalies, such as cross-over distortion, do not exist. Results are shown

in Fig. 9 and summarized in Table I. Both σ_T and $\sigma_{\Delta T}$ converged for over 1MSa. The bounds presented in (19) were not violated for either the XO and SSO. Both devices exhibited performance closer to the upper bound, thus indicating cycle-to-cycle correlation and the presence of device flicker noise as expected based on the phase noise data captured in Fig. 8.




Figure 8. Measured SSB phase noise PSD, the \sin^2 mask in (6) and the resulting projection. The first jitter difference is estimated by numerically integrating the projection over frequency for: (a) 24MHz SSO (b) 24MHz XO.

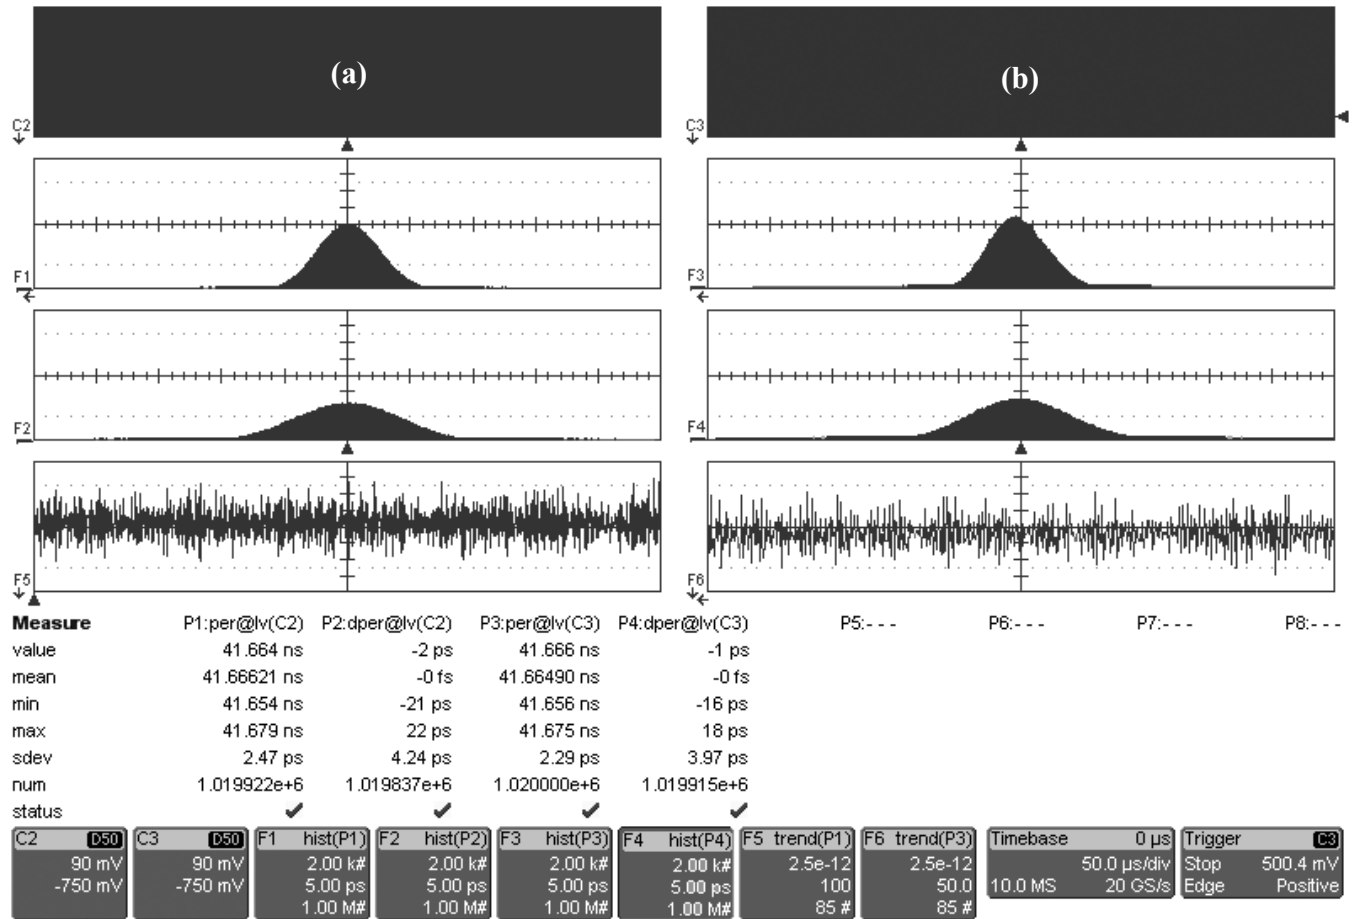


Figure 9. Real-time oscilloscope measurement of the first jitter differences (period or one-cycle jitter), the second jitter differences (cycle-to-cycle jitter) and the trends of frequency normalized to the mean. The first jitter differences are reported in the columns labeled “per@lv.” The second jitter differences are reported in the columns labeled “dper@lv.” Standard deviations are labeled “sdev.” Histograms are shown on the same scale. Data are reported for the following DUTs: (a) XO: $\sigma_T = 2.47\text{ps}$, $\sigma_{\Delta T} = 4.24\text{ps}$ (b) SSO: $\sigma_T = 2.29\text{ps}$, $\sigma_{\Delta T} = 3.97\text{ps}$.




Figure 10. Digital sampling oscilloscope measurement of the first jitter differences (period or one-cycle jitter). At least 100kSa are captured in each measurement. Measurements were captured on identical voltage and time scales. Data are reported for the following DUTs: (a) XO: $\sigma_T = 2.61$ ps (b) SSO: $\sigma_T = 2.27$ ps.

TABLE I. SUMMARY OF EXPERIMENTAL RESULTS

Metrology Approach	Variables or Conditions	Measured Metric (Units)	24MHz XO	24MHz SSO
Integrated SSB phase noise PSD from $f_m = 0$ to $f_m = f_h$	$f_h = 8f_o$	σ_T (ps)	1.95	3.00
	$f_h = 6f_o$		1.69	2.60
	$f_h = 4f_o$		1.38	2.12
	$f_h = 2f_o$		0.988	1.50
	$f_h = f_o$		0.713	1.06
	$f_h = f_o/2$		0.523	0.752
Digital sampling oscilloscope (DSO)	100kSa	σ_T (ps)	2.61	2.27
Real-time oscilloscope (RTO)	1MSa	σ_T (ps)	2.47	2.29
		$\sigma_{\Delta T}$ (ps)	4.24	3.97

D. Digital Sampling Approach

σ_T was measured using a 200kHz digital sampling oscilloscope (DSO) with a 3GHz analog bandwidth. In this approach, the waveform is reconstructed from low frequency samples of the signal across several different trigger events. Therefore, real-time events are not captured. The instrument was configured as shown in [6], where the signal from the DUT is split via a power splitter. One path triggers the instrument and the other is measured. σ_T is captured by accumulating samples of the first edge after the trigger event. Results are shown in Fig. 10. Over 100kSa were captured for both devices which required a net sampling time of approximately 20 minutes. The vertical (voltage) and horizontal (time) scales are identical for both measurements.

VI. DISCUSSION AND CONCLUSIONS

Good experimental agreement was found between measurements utilizing the RTO and DSO approaches for both DUTs. Integration of the phase noise PSD yielded much less consistent results for the SSO, which is due to the fact that the measured phase noise for the device was not at the floor at the maximum offset frequency. Better results were achieved for the XO. The presented bounds relating the first and second jitter differences were not violated and indicated the presence of device flicker noise as expected. The first jitter difference converged reliably and repeatably for both time-domain measurement approaches. However, both approaches accumulate the

results of short individual captures. Thus, as suggested previously, these approaches effectively apply a high-pass filter to the \sin^2 function in (6) and avoid the discussed pathology in the presence of pink noise. Although it has been shown that such measurements are common in industry and not mathematically tractable, the approach is acceptable because it does not cause measurement divergence. However, variable results across different instruments are certainly expected because the filtering effect is dependent on the capture dynamics. Clearly, standardized frequency stability metrology, similar to the Allan variance, is needed in the frequency control industry.

REFERENCES

- [1] D. Allan, "Time and Frequency (The-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 34, no. 6, Nov. 1987, pp. 647–654.
- [2] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. of Solid State Circuits*, vol. 34, no. 2, June 1999, pp. 790–804.
- [3] D. C. Lee, "Analysis of Jitter in Phase-Locked Loops," *IEEE Trans. Circuits Syst. II*, vol. 49, no. 11, Nov. 2002, pp. 704–711.
- [4] M. S. McCorquodale, "Self-Referenced, Trimmed and Compensated RF CMOS Harmonic Oscillators as Monolithic Frequency Generators," in *IEEE Int. Frequency Control Symposium*, 2008, pp. 408–413.
- [5] J. A. Barnes, *et al.*, "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, vol. IM-20, 1971, pp. 105–120.
- [6] M. S. McCorquodale, *et al.*, "A monolithic and self-referenced RF LC clock generator compliant with USB 2.0," *IEEE J. of Solid State Circuits*, vol. 42, no. 2, Feb. 2007, 385–399.