

Interpreting Oscillatory Frequency Stability Plots*

D. A. Howe †

National Institute of Standards and Technology
325 Broadway
Boulder, CO 80303

Abstract - This writing discusses the appearance of peaks and valleys in Allan deviation plots (also known as sigma-tau, root Avar, or Adev plots). This distinctly oscillatory pattern, especially at long-term τ -values, usually means the presence of quasi-sinusoidal frequency modulation of an oscillator's signal. However, quasi-sinusoidal oscillatory behavior in sigma-tau plots at long τ may be due to statistical sampling and not to actual oscillator or clock data. Periodic variations in sigma-tau plots are often used as an indicator of periodic environmental perturbations such as a diurnal or other external influence, and it is important to know whether these variations are an analytical artifact or not.

Removal of drift can make the oscillatory pattern worse. Samples of clock data for a dispersive noise process look like a portion of a sinusoid because the sample duration is less than the inverse of the data's inherent low frequency extent. Drift removal also removes low frequency components of noise, which causes negative bias of root Avar at long tau. The root Avar and drift-removal transfer functions have peaks and nulls that interfere with each other and can cause an oscillatory pattern in the resulting sigma-tau plot.

The best way to determine whether sigma-tau periodic variations at long term are real or not is by substituting statistics such as root Totvar or the newer Theo1, which, in particular, shows no anomalous oscillatory behavior.

I. INTRODUCTION

One of the biggest problems of interpreting a plot of frequency stability, sigma *vs.* tau, or Allan deviation, is trying to determine whether a given feature is real and has physical cause, is an aberrant behavior of the statistic used to estimate frequency stability, or is an interaction of the data and statistical processing. For example, processing that removes one or more systematics, such as frequency drift when using the Allan deviation, causes

*Contribution of the U. S. Government, not subject to copyright.

†The author is with the Time and Frequency Metrology Group, Time and Frequency Division; E-mail: dhowe@boulder.nist.gov.

known significant variability in sigma at long-term values of τ . This is not to say that this sigma *vs.* tau plot is "biased" but rather that a particular realization of it contains a pattern or kind of error which looks like a feature of the data. Various errors occur at long term due to a limit on the length of the data run with combinations of noise when in reality the feature disappears if one were able to take more data. In a similar context, patterns and errors occur at low frequencies when applying a Fourier transform operator to a finite block of data.

This writing focuses on the appearance of peaks and valleys, or a distinctly oscillatory pattern, especially for large τ , whose interpretation usually means the presence of quasi-sinusoidal data. The appearance of an oscillatory pattern at long-term τ may not mean the presence of quasi-sinusoidal data but may be a property similar to what is called "side lobe leakage," an artifact of sampling in frequency-domain spectrum analysis. I show how particular data sampling can cause an oscillatory pattern in, for example, the Allan deviation, and how removal of drift can accentuate this pattern. Unfortunately, separating this statistical artifact from that due to actual quasi-sinusoidal data with period T is not easy, and I illustrate how samples of noise often *look like* one sinusoidal cycle over T . This appearance is real and is itself random whose distribution properties in terms of Fourier frequency is chi-square distributed with 1 or 2 degrees of freedom. Thus the level of the lowest Fourier frequency component when considering the frequency domain has a significantly asymmetric probability distribution that creates a familiar oscillatory pattern. "Smoothing" is a process by which an estimate of some sequence of parameters includes the neighboring values to reduce the variability in the result. I discuss how smoothing at the largest τ -values is accomplished in Total frequency stability estimators (Totdev, Total MDEV, Total TDEV, Total Hadamard) and substantially reduces oscillatory patterns that are anomalous in a plot of sigma *vs.* tau, while the oscillatory pattern is retained if quasi-sinusoidal data actually is present. This property of all Total estimators yields a way to determine whether or not the data are actually quasi-sinusoidal near the last τ values. Long-

term frequency stability can also be calculated using a new statistic called “Theo1” that shows virtually none of the oscillatory behavior inherent in the Allan deviation.

II. THE ALLAN VARIANCE

Where narrow band processing is not required for determining the characteristic noise level of broadband processes such as power-law processes, the Allan variance is used instead of the inherently narrow band, windowed discrete Fourier Transform, or DFT.¹ This variance can be used to form a broadband spectral estimate using well known conversion schemes discussed elsewhere [1–3]. The presence of nearly sinusoidal (or quasi-sinusoidal) data simply means that the data have at least one periodic component added to an otherwise random-noise process having no particular periodicity. This mixture makes narrowband spectral analysis more suitable than the Allan variance for characterizing levels of quasi-sinusoidal data components in the noise, but the Allan variance is still preferred for quantifying long-term frequency stability in the τ -domain where FLFM and RWFM may dominate. So we are often faced with having to interpret quasi-sinusoidal components using the Allan variance.




Figure 1: Sample root Allan variance $\sigma_y(\tau)$ (referred to as Avar the best, or maximum-overlap estimator) of a commercial Cs frequency standard vs. NIST’s AT1 time scale. In the long-term τ -region, we see variations which may be interpreted as having a quasi-sinusoidal modulation in the data at long term (note the existence of peaks and dips near the longest term values).

¹The DFT is also highly variable across frequencies, making interpretation of conventional DFT-based spectral estimates problematic.

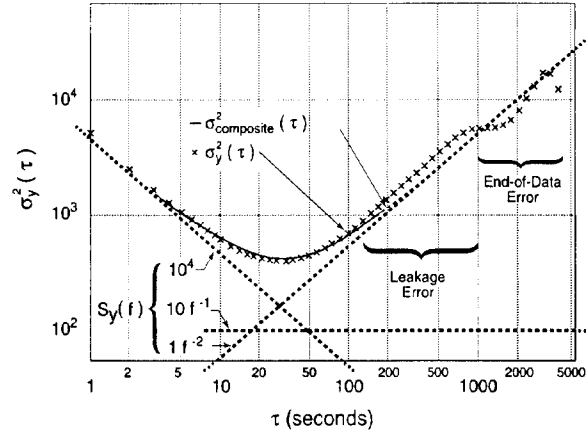


Figure 2: Note the apparent presence of quasi-sinusoidal data indicated by the oscillatory response at long-term of the sample Allan variance $\sigma_y^2(\tau)$ due to sampling errors (the noise simulation represents the linear addition of three power-law noise types and should result in the “composite” solid curve) [6].

A. Quasi-sinusoidal Allan Variance

As the averaging interval approaches the total interval T , even the best sample Allan variance (Avar) can show considerable variability and apparent negative bias that is dependent on known pathologies [4]. Total variance (Totvar), to be discussed later, has indicated more accurate estimates of characteristic long-term noise level [5]. This writing attempts to isolate errors associated with a less-known undesirable behavior that shows up as an oscillatory pattern at long term in the τ -domain using Avar. For example, Figure 1 is root Avar of data from a commercial Cs frequency standard vs. NIST’s AT1 time scale. In a τ -region of pure WHFM with no quasi-sinusoidal modulation, a quasi-sinusoidal modulation appears in Avar. At long term, note the existence of alternating peaks and dips in the last four τ -values. Figure 2 shows the response at all τ -values of Avar to pure FM noise simulation involving the linear addition of frequency-fluctuation spectral density ($S_y(f)$) slopes f^0 , f^{-1} , and f^{-2} , whose expected Avar result should be the line indicated as $\sigma_{\text{composite}}^2(\tau)$ [6]. While quasi-sinusoidal modulation appears, this interpretation at long-term τ (the end of data error) would be incorrect. An oscillatory pattern occurs randomly as shown in Figure 3. In 100 simulations of pure WHFM, we see similar patterns in many of those trials. In the case of the simulations, figures 2 and 3, we know by design that there is no quasi-sinusoidal modulation at the level indicated by Avar. Since there are no correlation artifacts known in Avar itself that might account for this response, what is happening? Furthermore,




Figure 3: Oscillatory response at long term shows up frequently in the classical sample root Allan variance $\sigma_y(\tau)$ with 100 simulation trials of WHFM noise.

root Totvar is recommended over root Allan variance as a measure for accurate long-term frequency stability characterization (see figure 4 simulation results). We find that the circularization of the data defined in Totvar and described later makes it less subject to these errors. But why?

The key for developing interpretations of sigma *vs.* tau at long-term is found in the relation of the Allan variance expressed in the frequency-domain model. The Allan variance *f*-to- τ transform can be written as:

$$\sigma_y^2(\tau) = \frac{1}{2} \langle (\bar{y}_{k+1} - \bar{y}_k)^2 \rangle \quad (1)$$

$$= \lim_{T \rightarrow \infty} \frac{1}{T} \left[2 \int_{\frac{1}{T}}^{f_h} df S_y(f) \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2} \right], \quad (2)$$

where τ is the averaging time, lag or correlation time $T = 2\tau$ by definition, and $S_y(f)$ is the Power Spectral Density (PSD) of $y(t)$, the fractional frequency deviation function. Note that $\frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2}$ is a filter response (squared) whose first and largest maximum is at $f\tau = \frac{1}{2}$. A plot of the overall filter response is shown as the solid curve in figure 5. Data runs have finite duration, T , so the integral in (1) has an operational lower limit frequency of $\frac{1}{T}$. At $\tau = T/2$, $f_{lower} = \frac{1}{T}$ has a high-pass characteristic for $f < \frac{1}{T}$ and a low-pass stage followed by a (and this is important) periodic filter response on the right side of the plot in figure 5 (of form $\sin x/x$). More specifically, we can separate Avar's filter response into two pieces, a high-pass (*hp*) to the left and a low-pass (*lp*) to the right, as:




Figure 4: Not only is root Totvar less variable, but its response is smoother with fewer oscillations at long-term with the same 100 simulation trials of WHFM noise as in figure 3.

$$|H_{hp}(f)|^2 \doteq (\pi f \tau)^2,$$

$$|H_{lp}(f)|^2 = \frac{\sin^4(\pi f \tau)}{(\pi f \tau)^2}.$$

The Allan variance can be interpreted as the standard variance of an FM noise process having passed through a constant- Q bandpass filter [7]. A close approximation is cascaded high-pass and low-pass single-pole filters with RC time-constant $\tau/2$, also shown in figure 5 [8].

B. Problems of Sampling and Processing

At the end of a data run the last possible τ -value using Avar is $T/2$ computed as half the squared subtraction of average frequency \bar{y}_1 for the first half of T and \bar{y}_2 for the second half. The maximum-lag response is at $\tau = T/2$, and furthermore only one sample value exists at $T/2$, which means only one value, hence a degree of freedom of 1, and it is used in Avar's estimate of the "true" Allan variance.

Before proceeding, it is somewhat enlightening to keep in mind that noisy phenomena can tend to oscillate by randomness alone. For example, let $X1, X2, X3$ be independent and have the same probability density. What is the probability that they are monotonic: $X1 < X2 < X3$ or $X1 > X2 > X3$? All 3! possible orders of $X1, X2, X3$ have probability 1/6, so the answer is 1/3. Thus the probability that they "oscillate", i.e., are not monotonic, is 2/3. In contrast, neighboring values of Avar

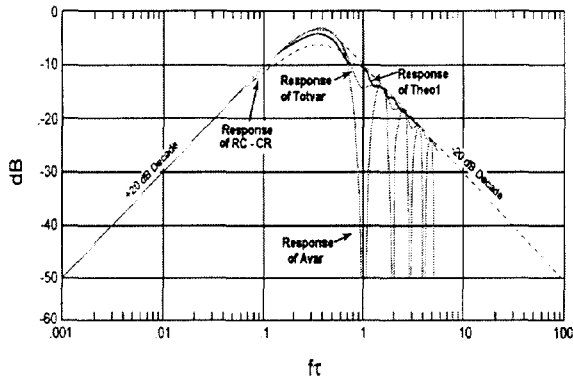


Figure 5: A comparison of frequency responses acting on $y(t)$ of Avar, Totvar, and a passband variance considered to be ideal consisting of a simple cascade of a single-pole high-pass followed by a low-pass filter with identical break points at $RC = \tau/2$ (dashed curve). Totvar side lobes are flatter and closer to ideal without Avar's deep nulls at large τ . The dark solid line is the frequency response of a new statistic called "Theol" and which exhibits the lowest passband ripple.

are not independent, so this example isn't entirely appropriate to a root Avar plot.

Regarding the issue of finite sampling and processing, there are three basic concepts involving sampling of data. First, a finite-observation period given by T always acts as a high-pass filter because we cannot be sensitive to any change taking longer than T . Low-frequency catastrophic divergence, a property of non-stationary models of processes with "red" spectra such as FLFM and RWFM, is never realized because of finite time limits [9]. For this reason, $\frac{1}{2T}$ is widely used as the usual sample low frequency f_{sl} and as the lower limit in frequency-domain analysis even though the real characteristic low frequency f_{rl} of the process being investigated may be far lower.

Note that one sample of data for a dispersive process (having a "red" spectrum) such as random walk does not look like noise about a nearly flat line in the time domain but will more apparently look like a complete, or a portion of, a sinusoidal cycle. This is because the sampling duration T is shorter than the inverse of the real minimum frequency of the flicker and random walk FM processes. Figure 6 shows 100 segments of $\{x_t\}$ as simulations of random walk FM. The Allan variance responds to this sample data as we would visually, namely as having sizeable power concentrated in one low frequency component in the vicinity $\frac{1}{T} < f_0 < \frac{1}{2T}$. The appearance of a sinusoidal cycle is a random occurrence whose distribution properties in terms of Fourier frequency is a skewed chi-square distribution if there are only 1 or 2 degrees of freedom. Thus the values at lower Fourier frequency

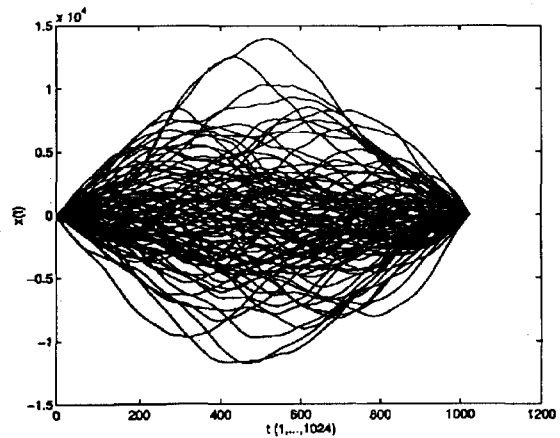


Figure 6: 100 segments of $\{x_t\}$ as simulations of random walk FM in which a slope is removed so that the ends are matched. Note that the random shape often appears as about one-half or one sinusoidal cycle over the entire sample interval.

components have an abnormally high probability of being indicated in the frequency-domain. This is consistent with τ -domain uncertainties in which lower values of $\sigma_y(\tau)$ have a higher probability of occurrence at the last few points.

Second, part of the statistical process or oscillator system (such as in a Cs frequency servo or disciplined oscillator) involves the routine matter of removing estimates of trends in frequency, usually modeled as functions of low order, most notably linear frequency drift. Avar characterizes residual noise measurements very well but not in the presence of systematics such as drift that mask the underlying stochastic process. But removing low-frequency drift also removes some of the stochastic noise. Drift removal suppresses, sometimes dramatically, the Fourier frequency components near $f = \frac{1}{2T}$. This is because an underlying quadratic function contained in $x(t)$ over interval T (that is, the usual model of linear drift in $y(t)$) is hard to distinguish from an underlying real half-cosine function over the same interval with other imposed noise. Estimates that remove one invariably remove the other and cause negative bias (depressed response) in Avar at long-term values of τ in the remaining residuals. It can also cause an oscillatory pattern in long-term Avar in the following way. Drift removers can be modeled as classes of Fourier filters with approximate lag T . At $\tau = T/2$, we are thus left with noise residuals having been processed by a notch filter whose center frequency is approximately the maximum center peak in Avar's frequency response (shown as the solid curve in figure 5) and its odd harmonics. Indeed the response of Avar using a common estimator of drift is exactly zero at $\tau = T/2$ [10]. Odd Fourier components at $\tau = T/4$ do not coincide with

Avar's nulls; however, at $\tau = T/8$ odd higher f components again end up in many of the nulls in Avar's response. Multiplying the transfer function of the drift remover with Avar's transfer function reveals a picket-fence pattern between the two. This pattern results from an interaction between Avar's frequency response, whose center frequency depends on the value of τ , and a drift remover, whose notch frequency does not. The result is that Avar exhibits oscillations *vs.* τ .

C. Leakage

It is clear that a large contributor to long-term behavior of frequency stability identifies with the periodic notches between the lobes in the bandpass characteristics of Avar. These side lobes have always been regarded as an undesirable property of narrow-band spectrum analysis. Spectrum analyzers typically compute a power spectrum for a time series $\{y_t\}$ by windowing, that is, tapering the series at both the beginning and end by means of tapering sequence $\{h_t\}$ and then taking the squared modulus of the DFT of the windowed series $\{h_t y_t\}$. The purpose of windowing is to reduce a potential bias known as leakage, in which power "leaks" from high power portions into low power portions of the spectrum. This author speculates that leakage from high-level short-term noise accounts for the error approximately over the region shown in figure 2.

Avar does not include a window series $\{h_t\}$, which tapers the original data because side lobe leakage is arguably a feature of the definition (if we regard a constant- Q bandpass as an idealized variance response as in the dashed curve of figure 5). Unfortunately Avar's response at long term depends more critically on the nulls between the side lobes than its short-term response because Avar has only 1 or 2 terms in its summand, meaning virtually no smoothing. Tapering would modify the high-frequency (low-pass) response, in particular its slope, and additionally would shorten the maximum averaging time to achieve the same confidence as the presently-used un-tapered series. Also a windowed version of the Allan variance naturally calls for wavelet-based variances as a more reasonable approach to analysis considering the vast literature on the subject of wavelet analysis. One rationale for the wavelet variance is that higher-order wavelets address the leakage problem [11]. Lastly, windowing of Avar is partially done using the modified Allan variance (Mvar) [12]; more windowing could overly complicate otherwise simple formulas. Nevertheless Allan variance leakage is a concern in common situations in which the short-term FM and PM noise level is higher than the long-term (usually FM only) noise level is low, but it is a contributing error in any situation in which high-power components can leak to low-power regions in the τ -domain.

D. How Leakage Creates an End-of-Data Error

Allan variance leakage originates from side lobe frequency response due to the window of observation of data in process, a consequence of starting a measurement at t_0 and stopping it at T . The solid curve of figure 5 with the downshoots (deep nulls) is the frequency response of Avar; the high-frequency response to the right of the main peak are the side lobe responses. Avar leakage error often shows up as a slight positive change in Avar's slope and is usually attributed to power from the noise at high frequency (short τ -values). This is because the Allan variance's *bandpass filtering on frequency* noise is equivalent to *high-pass filtering on phase* noise, so side lobe leakage is more pronounced in the presence of wide band (short-term) PM noise, a likely occurrence. For example, in figure 2, leakage error from high-frequency noise augments the oscillatory response in Avar which, in short- to mid-term, averages to a smooth change of positive slope in the response as possibly attributable to the indicated "leakage error" region in figure 2 except at the longest-term τ -values, where the number of averageable frequency deviates goes from a large number to only 1. At the longest τ -values, an oscillatory leakage error becomes indistinguishable from, and combines randomly with, other sampling errors attributable to an end-of-data.

The start and stop of a measurement from its non-measured state (assumed to be zero everywhere), sometimes introduces an artifact of sampling of a particular random shape in the data run which artifact appears as a non-negligible power localized at frequency $f_{sl} = \frac{1}{2T}$. Regarding $\frac{1}{2T}$ as itself a functional "frequency" declares that it is periodic (though it may not be) with harmonics across a wide spectrum; thus we have an end-of-data error attributable to an assumption about the state of the data outside the actual measurement run. (The same behavior occurs starting from zero at t_0 .)

For "red" processes, does the power below f_{sl} "leak" into the sample Allan variance in the same way as high-frequency noise? In a sense it does, but not because of its highly variable property alone. If $f_{rt} < f_{sl}$, then stopping the measurement (and likewise its abrupt turn-on at commencement) is likely to introduce a non-zero constant into Avar, which constant is indistinguishable from a pulse of excitation that produces a "ringing" damped oscillatory response out of Avar and whose period is roughly the reciprocal bandwidth of Avar's main one-octave, constant- Q frequency-domain filtering action indicated in (2) and shown as the solid curve of figure 5. Leakage described in this way causes the appearance of what looks like quasi-sinusoidal noise in the data itself. To make matters worse, the oscillatory Avar response is further accentuated by the deep nulls that, from a sig-

nal processing point of view, represent a highly dispersive filter-response at harmonics, one characterized by large phase-shifts *vs.* $f\tau$. The power in the harmonics of the pulse pass through the dispersive side lobe leakage response and are phase-shifted by a full π radians (or lag of $T/2$) at the nulls, which enhance Avar's oscillations. As mentioned, tapering the time-series data controls the power in these harmonics (or conversely, the Avar frequency response, the argument in the integral of (2)) but it compromises our basic objective, namely estimating long-term noise level.

In short, if an oscillatory pattern at long-term τ -values is observed in root Avar, is there real underlying quasi-sinusoidal noise in the processed data, or is the cause associated with sampling errors due to an end-of-data sampling artifact, only one of which might be leakage? *In toto*, it is the fact of the abrupt halt of the measurement of data that causes rather sizeable errors and particular patterns in the last few points (longest averaging times), which patterns combine with possible underlying quasi-sinusoidal noise in the data and confound a clean interpretation of root Allan variance.

III. THE TOTAL VARIANCE

Considerable recent work has been done to improve estimates of oscillator noise level even at the full extent of the data, and for intervals approaching the total data length, a new variance called "Total Variance" or Totvar has demonstrated less variability with only modest bias relative to the Allan variance. Totvar is a statistic whose strategy eliminates the abrupt halt of the measurement but without the use of tapering. It thereby reduces the deep nulls in the frequency response associated with Avar, while still characterizing long-term noise out to $\tau = T/2$ (see figure 5). Totvar incorporates a simple, but worthwhile, procedure before the application of the max-overlap sample Allan variance. In terms of phase (or time) deviates x_t , Totvar is Avar that is processing a larger, virtual set of phase or time data, which set is an odd, or reflected, extension at the beginning and end (left and right) of the original real set. Figure 7 illustrates the $x(t)$ double-sided mirror-reflection extension and, hence, the resulting circular or repeating representation of the original time series. The Totvar estimator is based upon the hypothesis that reasonable surrogates for unobserved frequency deviates y_t , $t < t_0$ or $t > T$, can be formed by tacking on reversed versions of $\{y_t\}$ at the beginning and end of the original series. The Totvar estimator makes use of certain of these surrogate values in order to extend the usual smoothing of the best Allan variance estimator (Avar) at the largest sampling times




Figure 7: Circular extension of the original $x(t)$ data set for computation of Total variance.

approaching $\tau = T/2$ [5].

A. Smoothing

"Smoothing" is a process by which an estimate of some sequence of parameters incorporates in some manner, usually by a weighted average, the neighboring values to reduce the variability in a result. In the case of the power spectrum, estimates are less variable (smoother) when computing overlapped ensemble averages of DFT's. Unfortunately these estimators lack the ability to be smoothed when we are concerned with the very lowest Fourier frequencies. The smoothing procedure is occasionally biased using a periodic extension, an inherent property of using the DFT. Nevertheless, smoothing is an accepted way of improving the confidence of spectral estimators. Often the bias is known or can be calculated under a certain set of conditions.

The root Allan variance *vs.* τ can incorporate neighboring τ -value estimates to reduce the variability in the statistical result. This was first pointed out by Howe, Allan, and Barnes [3] in what was called the "max-overlap estimator" and which has become the standard estimator. The procedure simply computes every possible estimator of the root Allan variance and overlaps them in an equally weighted average to the maximum extent possible. Root Totvar, as a sigma-tau plot, extends this smoothing to the longest τ -values by its periodic extension of the original dataset.

Since the frequency-error function $y(t)$ is always sampled for data duration T to realize series $\{y_t\}$ used in estimating (2), Totvar applies the usual max-overlap sample Allan variance so that it treats $\{y_t\}$ as periodic much in the same way that the DFT, for example, would interpret $\{y_t\}$ as periodic (by the convolution theorem), as contrasted with the continuous FT of continuous $y(t)$.

Why not use evaluations of Totvar at $\tau > \frac{T}{2}$? Sampling theory dictates that there is no advantage to starting the integral in (2) with a lower limit that is less than the Nyquist frequency $f_s = \frac{1}{2T}$ since there are an infinite number of reproductions of the integral's argument which extend to $f = 0$. In other words, there is no new information below f_s . The lower limit to the integral in (2) is usually written as 0 but because the first sampled frequency can't be 0 (implying infinite T), it ends up being $\frac{1}{T}$ as a practical matter since we must resort to a sample (finite T) Allan variance estimate and find the available range of τ to be $\tau_0 \leq \tau \leq T/2$ [13]. This is also verified in simulation. Evaluations of frequency stability at $\tau > \frac{T}{2}$ have been demonstrated by a statistic called "Theo1" and are discussed in the last section [14].

There is some evidence that Totvar responds more smoothly when estimates of systematics such as frequency drift have been removed from the data. Common estimators, which cause severely depressed values of Avar at long-term τ , do not adversely depress the characteristic noise at long term reported by Totvar [15]. This property may manifest as reduced oscillatory behavior using Totvar, but this has not been investigated beyond the discussion here.

With regard to side lobe frequency response, figure 5 shows frequency responses of Totvar and Avar as the shaded curves. The frequency response of Total variance as a function of τ is calculated by averaging the squares of the Fourier transforms of the many (Total) sampling functions and find that it resembles the frequency response of the Allan variance but without its deep nulls. Total variance has a smoother approximation to a constant- Q filter response. This indicates that Total variance has an interpretation like that of the Allan variance but has lower uncertainty at and near the total time interval of a measured data set [5, 13]. For the discussion at hand, Totvar more accurately represents a constant- Q bandpass filter with flatter side lobes in the τ -region of interest as shown in figure 5. In particular, root Totvar responds with less oscillatory ripple in the last τ -values, without the pronounced, alternating peaks and dips occurring as often. This is evidenced by the same 100 WHFM simulation trials of figure 3 but using root Totvar as shown in figure 4.

In light of the smoother long-term τ response of Totvar to random noise, an obvious question is whether it will respond properly, like Avar, with an expected oscillatory pattern in the presence of actual quasi-sinusoidal data. As mentioned, a finite-observation period given by T always acts as a high-pass filter because we become increasingly insensitive to gradual changes slower than T . The periodic data extension introduced into Totvar does not alter this fact, but does assume that any overall change is periodic with minimum frequency of $f_l = \frac{1}{2T}$. Given data that actually contain a periodic function, the

methodology in Totvar will extend this periodicity, although there will likely be a phase shift at the extension points t_0 and T , which phase shift doesn't seriously alter the result. Even though Totvar, like Avar, analyzes broadband spectra such as power-law noise processes, it naturally approaches narrow-band spectrum analysis in long-term τ -values because of its constant- Q frequency response. Totvar's periodic extension is likely to be a valid assumption for characterizing real quasi-sinusoidal data of any Fourier frequency including components at very low Fourier frequency. This feature is added incentive to use Totvar for more accurate interpretations of long-term frequency stability, particularly if oscillatory patterns show up with the use of Avar.

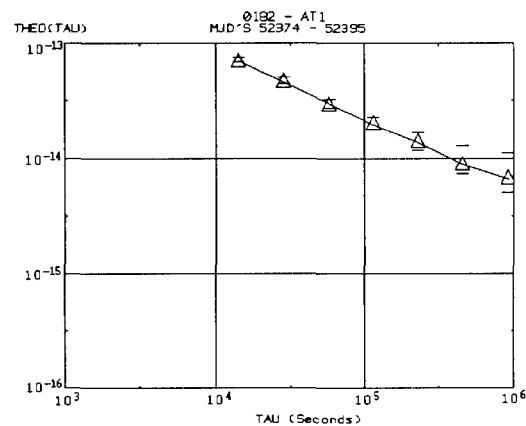


Figure 8: Long-term frequency stability can also be calculated using a new statistic called "Theo1," which shows virtually none of the oscillatory behavior inherent in the Allan deviation. Shown above is a plot of Theo1 (interpreted exactly as the usual sigma-tau, or Allan deviation, plot) that uses the same data used to compute the Allan deviation, or root Avar, in figure 1.

B. THEO1, Evaluations at Long-term

Avar and Totvar have useable properties as an estimator of frequency stability only to a longest-term τ -value of $T/2$. A new statistic has been tested that extracts frequency stability beyond $\tau = T/2$. This statistic, dubbed "Theo1," computes selected linear combinations of frequencies at every possible averaging time τ , but in a fixed interval T [14]. T in the sense of Theo1 has an equivalence to τ in the sense of Avar and Totvar. While Theo1 is a species unlike Avar and Totvar, it retains all of their desirable properties, plus, by its definition, can report frequency stability up to the full duration of a data run. Central to this discussion, Theo1 has very low ripple in its transfer function, as shown by the solid dark line in

figure 5, and has shown virtually no anomalous oscillatory behavior in simulation trials.

Figure 8 is a plot of Theol (interpreted exactly as the usual sigma-tau, or Allan deviation, plot) that uses the same data used to compute the Allan deviation, or root Avar, in figure 1. The confidence intervals in the indicated τ -range are significantly smaller than those in figure 1. Also, one notes the absence of an oscillatory pattern.

IV. CONCLUSION

This writing has shown that the problem of interpreting oscillatory patterns in root Avar (an Allan deviation plot) is not straightforward. The best way to determine whether periodic variations at long-term using root Avar are real or not is to use new statistics such as root Totvar (known as Totdev) or Theol, which, in particular, shows no oscillatory behavior.

V. ACKNOWLEDGEMENT

The author thanks Bill Riley of Datum, Corporation, Francois Vernotte of the Observatoire de Besançon, and Charles Greenhall of the Jet Propulsion Laboratory (Caltech) for very helpful comments and discussions.

VI. REFERENCES

- [1] J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healy, D. B. Leeson, T. E. McGunigal, J. A. Mullen, Jr., W. L. Smith, R. L. Sydnor, R. F. C. Vessot, G. M. R. Winkler, "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, **IM-20**, 105-120 (1971).
- [2] D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls (Editors), *Characterization of Clocks and Oscillators*, Natl. Inst. Stand. Technol. Technical Note 1337 (1990; available from NIST, 325 Broadway, Boulder, CO 80303-3328).
- [3] D. A. Howe, D. W. Allan, and J. A. Barnes, "Properties of Signal Sources and Measurement Methods," *Proc. 35th Freq. Cont. Symp.*, 1-47 (1981; reprinted in [2]).
- [4] D. A. Howe, "An Extension of the Allan Variance with Increased Confidence at Long Term," *Proc. IEEE International Frequency Control Symposium*, pp. 321-329, 1995.
- [5] D. A. Howe and C. A. Greenhall, "Total variance: a progress report on a new frequency stability characterization," *29th Ann. PTTI Systems and Applications Meeting*, pp. 39-48, 1997.
- [6] F. Vernotte, E. Lantz, J. Gros Lambert, J. J. Gagnepain, "Oscillator Noise Analysis: Multivariate Measurement," *IEEE Trans. Instrum. Meas.*, **IM-42**, no. 2, pp. 342-350, April 1993. Also "A New Multivariate Method for the Oscillator Noise Analysis," *Proc. IEEE International Frequency Control Symposium*, pp. 284-289, 1992.
- [7] J. Rutman, "Characterization of frequency stability: A transfer function approach and its application to measurements via filtering of phase noise," *IEEE Trans. Instrum. Meas.*, vol. **IM-23**, pp. 40-48, Mar. 1974. Also J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress," *Proc. IEEE* **66**, 1048-1075 (1978).
- [8] R. G. Wiley, "A Direct Time-Domain Measure of Frequency Stability: The Modified Allan Variance," *IEEE Trans. Instrum. Meas.*, vol. **IM-26**, pp. 38-41, Mar. 1977.
- [9] J. A. Barnes, "Models for the interpretation of frequency stability measurements," NBS TN-683, August, 1976.
- [10] M. A. Weiss and C. Hackman, "Confidence on the Three-Point Estimator of Frequency Drift," *Proc. 24th Ann. PTTI Meeting*, 451-460 (1992).
- [11] D. A. Howe and D. B. Percival, "Wavelet Variance, Allan Variance, and Leakage," *IEEE Trans. Instrum. Meas.*, **IM-44**, 94-97 (1995).
- [12] D. W. Allan, M. A. Weiss, and J. L. Jespersen, "A Frequency-domain View of Time-domain Characterizations of Clocks and Time and Frequency Distribution Systems," *Proc. 45th Freq. Cont. Symp.*, 667-678 (1991).
- [13] C. A. Greenhall, D. A. Howe and D. B. Percival, "Total Variance, an Estimator of Long-Term Frequency Stability," *IEEE Trans. Ultrasonics, Ferroelectrics, and Freq. Control*, **UFFC-46**, no. 5, 1183-1191, Sept, 1999.
- [14] D. A. Howe and T. K. Peppler, "Estimating Extremely Long-term Frequency Stability: A New Statistic," to be published.
- [15] D. A. Howe and K. J. Lainson, "Effect of Drift on TOTALDEV," *Proc. 1996 IEEE Intl. Freq. Control Symp.* 883-889.