

Characterizing Frequency Stability: A Continuous Power-Law Model with Discrete Sampling

Todd Walter

Abstract—This paper examines several aspects of the Allan variance and the modified Allan variance. New expressions for these variances are derived for noise processes that produce power spectral densities with both integer and noninteger powers (α) in their functional dependence on f . A single expression, continuous over α , is presented for each of these variances. Also investigated are the effects of discrete sampling and finite data length. Discrete equations are developed and compared with more familiar continuous expressions. In addition, the uncertainty of the estimates for the Allan variance and the modified Allan variance for fully overlapping data usage is presented. The uncertainties can be calculated for arbitrary α . The results presented are compared with computer simulations and found to be in excellent agreement.

I. INTRODUCTION

THE correct characterization of the effect of noise is an important issue to both users and manufacturers of high-performance oscillators because noise ultimately limits the precision of these devices. There are many possible sources of noise inherent both to the oscillators and to the frequency measurement systems. These sources are conventionally grouped into one of five noise categories. The categories are distinguished by the slopes of the power spectral densities (PSDs) of frequency fluctuations as functions of Fourier frequency in a log-log plot. In the standard power-law model, the PSD is assumed to be proportional to f^α . The five different noise categories considered to affect oscillators are white phase, flicker phase, white frequency, flicker frequency and random walk frequency. These correspond to the integer values of α ranging from 2 to -2 respectively. Although only integer values of α are commonly considered, there is no real reason to exclude noninteger values, as the model is largely empirical. The quality of the observed PSDs does not justify excluding the possibility of noise sources that would yield noninteger values of α . If one is interested in characterizing the limits of precision for a particular device, then the standard power-law model, coupled with deterministic effects, is more than adequate [21]. However, if one is interested in understanding the noise source itself, then correctly determining the true functional dependence of its PSD can be important. This paper examines the effects of noise which has a PSD proportional to f^α , but for a continuous range of α .

Recently a discrete noise generation routine capable of creating noise with arbitrary α was introduced [1]-[2]. In

Manuscript received September 22, 1992; revised June 15, 1993.
The author is with the Gravity Probe B Program, Hansen Experimental Physics Laboratory, Stanford University, Stanford CA 94305-4085.
IEEE Log Number 9216495.

addition to verifying the results presented in this paper via computer simulation, the noise generator addresses the issues of discretization and finite sample length. This paper will investigate these effects on both time domain and frequency-domain measures of noise.

We must first understand how noise affects the oscillator output. The output of an oscillator is often represented by [3]

$$V(t) = [V_0 + \varepsilon(t)] \sin[2\pi\nu_0 t + \varphi(t)] + V_1(t),$$

where V_0 and ν_0 are the respective nominal amplitude and frequency of the output, $\varepsilon(t)$ and $\varphi(t)$ are amplitude and phase fluctuations respectively, and $V_1(t)$ is additive noise. Provided ε and V_1 are much smaller than V_0 , the instantaneous frequency of the oscillator output can be written as

$$\nu(t) = \nu_0 + \frac{1}{2\pi} \frac{d\varphi(t)}{dt}.$$

From this, the instantaneous fractional frequency deviation from nominal may be defined:

$$y(t) \equiv \frac{\nu(t) - \nu_0}{\nu_0} = \frac{1}{2\pi\nu_0} \frac{d\varphi(t)}{dt}.$$

The frequency stability (or instability) of the oscillator can be specified through the characterization of $y(t)$. It is not possible to measure $y(t)$ directly. The measurement of $y(t)$ takes place over a finite time interval τ . Thus, it is usually more convenient to define the instantaneous phase deviation in units of time

$$x(t) \equiv \frac{\varphi(t)}{2\pi\nu_0}.$$

We can see that $y(t)$ is the time derivative of $x(t)$. The average value of $y(t)$ over the time interval τ can be obtained through integration:

$$\bar{y}_k = \frac{1}{\tau} \int_{t_k}^{t_k+\tau} y(t) dt = \frac{x(t_k + \tau) - x(t_k)}{\tau}.$$

The quantity $x(t)$ may contain both deterministic and non-deterministic effects. Examples of deterministic effects would include linear frequency drift and diurnal trends. In principle, it is possible to account for these effects. Nondeterministic effects or stochastic processes are more difficult to resolve. These processes lead to random fluctuations in $x(t)$. In the following sections, this paper will examine the effects of stochastic processes on stability measures.

II. FREQUENCY DOMAIN

The richest sources of information about noise and other disturbances to an oscillator are often the PSDs $S_x(f)$ and $S_y(f)$ ¹. The PSD graphically separates the various noise sources affecting the oscillator frequency. In fact, the most common noise types affecting oscillators were originally distinguished by the slopes of their PSDs. The standard power-law model is written as [3]

$$S_y(f) = \sum h_\alpha f^\alpha, \quad (1)$$

where h_α is the intensity coefficient for the noise category. As mentioned before, only integer values are traditionally considered. Some noise sources, such as white noise and those following a random walk, which produce PSDs having even integer values of α , are reasonably well understood. However, other noise sources yielding slopes in the region of the odd integers have been observed empirically. Analytically, these noise sources are not well understood. In some analyses [1], [4], fractional calculus is invoked to derive their dependencies. Because they depart from classical processes, their existence can lead one to widen the power-law model, and consider not only odd-integer values, but also noninteger values. In this new continuous model the sum in (1) could be replaced with an integral and the coefficients replaced with a noise intensity density function. This would permit the analysis of noise with a broader signature in the frequency domain. However, without any loss of generality, this paper will consider processes $x(t)$ which are affected by only a single noise process.

Because of their relationship in the time domain, the PSD of $x(t)$ is related to the PSD of $y(t)$:

$$S_x(f) = \frac{S_y(f)}{(2\pi f)^2}. \quad (2)$$

As already mentioned, $x(t)$ is the more convenient process, since it is a quantity we can directly measure. We can therefore translate the power-law model to a form that works with $x(t)$. I will use new parameters β and g_β defined such that

$$S_x(f) = \frac{h_\alpha}{(2\pi)^2} f^{\alpha-2} \equiv g_\beta f^\beta. \quad (3)$$

The five conventional noise categories now correspond to the integer values of β ranging from 0 to -4 . Most of the subsequent work will be done with these parameters and will later be converted back to traditional nomenclature (α, h_α).

One of the reasons that the PSD is such a powerful tool is that in addition to noise, any periodic modulation of the signal can be easily observed in the frequency domain. Examples of possible modulations include 60 Hz pick-up and mechanical resonances. These modulations appear as sharp peaks in the PSD. However, rather than being able to obtain the true PSD for a particular noise process, we can only form estimates from the discretely sampled data. The estimates are distorted by the effects of discrete sampling and finite data length. Because there is only a finite amount of data from which to form the

¹The spectral density is a function of Fourier frequency (f). The frequencies f are completely unrelated to the nominal oscillator frequency ν_0 . It is important to distinguish between these two.

estimate, we only have information about a finite frequency range. Any spectral information outside this range will fold down and distort the estimates in the region where we do have information. This effect is called aliasing (see [5] section X), and it will be mentioned again in sections VI and VII. Improper windowing may also create distortions and biases in the spectral estimates [6], [7]. One must therefore exercise great care when using the PSD estimate.

The timing community is primarily interested in clock accuracy. The PSD is not able to provide a convenient measure of the precision of time intervals. It does permit a measure of the types of noise affecting clock performance, but not of the uncertainty over a specific time interval. For this information, we will turn to measures in the time domain.

III. TIME DOMAIN

The time-domain analog of the PSD is the autocorrelation function. It is the inverse Fourier transform of the two-sided autospectral density. It should be noted that all of the spectral densities discussed in this paper are one-sided. The symmetric time/lag autocorrelation function for a continuous real zero-mean process $x(t)$ is given by [1], [20]

$$R_x(t, \tau) \equiv \langle x(t - \tau/2)x(t + \tau/2) \rangle$$

where the brackets $\langle \rangle$ denote an expected value or an infinite ensemble average. The explicit dependence on time is because many of the noise processes considered to affect oscillators are nonstationary. A (wide sense) stationary process is one, such that in the limit that time goes to infinity, its autocorrelation function converges to a finite value

$$\lim_{t \rightarrow \infty} R_x(t, \tau) \rightarrow R_x(\tau).$$

This is not the case for noise following the power-law model with $\beta \leq -1$.

For a stationary real process, $x(t)$, the autocorrelation function is related to the spectral density by the inverse Fourier integral

$$R_x(\tau) = \int_0^\infty S_x(f) \cos(2\pi f\tau) df. \quad (4)$$

Unfortunately, this function is not directly useful for characterizing oscillator stability. Unless the noise process is ergodic, it is impossible to obtain a reasonable estimate of the autocorrelation function. Because the process we are trying to characterize is a noise process, each individual measurement may vary widely. However, the average of many measurements should converge to a definite value. The difficulty is that we often have only one oscillator. Even if we had an ensemble of oscillators, we cannot guarantee that they all would be affected to the same degree by exactly the same noise processes. However, if the noise process is ergodic, the time average approaches the ensemble average for long averaging times. This would enable us to form an estimate from a single oscillator. It is only an estimate because any time series actually measured will be finite in length. There are two assumptions made with this approach: First, that the noise process will not change over time and second, that the process

is ergodic. A necessary, but not sufficient, requirement for ergodicity is that the process be stationary. It is obvious then why we cannot form estimates of the autocorrelation function for $\beta \leq -1$. We must, therefore, seek out other time-domain measures.

IV. ALLAN VARIANCE

The most common time-domain measure is the Allan (or two-sample) variance [9]. It has the advantage that experimentally it is convergent for $\alpha > -3$ (although a system bandwidth must be imposed for $\alpha \geq 1$). In terms of the phase deviation, the Allan variance is defined by

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2} \langle (x_{k+2} - 2x_{k+1} + x_k)^2 \rangle. \quad (5)$$

A new process can be defined by

$$z(t, \tau) \equiv \frac{x(t+2\tau) - 2x(t+\tau) + x(t)}{\tau}.$$

$z(t, \tau)$ is a process which must be stationary in order for the Allan variance to be convergent. Equation (5) is similar in appearance to the autocorrelation of $x(t)$. It can be expressed as a linear combination of autocorrelation functions

$$\begin{aligned} \sigma_y^2(\tau) = & \frac{1}{2\tau^2} [R_x(t+2\tau, 0) + 4R_x(t+\tau, 0) + R_x(t, 0) \\ & + 2R_x(t+\tau, 2\tau) - 4R_x(t+3\tau/2, \tau) \\ & - 4R_x(t+\tau/2, \tau)]. \end{aligned} \quad (6)$$

The Allan variance can also be related to the PSD, although the two are not a simple Fourier transform pair [10]:

$$\sigma_y^2(\tau) = \int_0^\infty 2S_x(f) \sin^4(\pi\tau f) df. \quad (7)$$

This relationship was derived for the case of stationary noise ($\alpha > 1$). However, the relationship is extended down to $\alpha > -3$ because the integral relationship converges and agrees with experimental results. Due to this convergence, many authors treat the autocorrelation function as though it represented a stationary process over this range. Instead, it is the linear combination of autocorrelation functions in (6) that represents the convergent measure.

Note that when α is greater than or equal to one, the integral relationship will not converge. The PSDs in these cases contain an infinite amount of energy when integrated out to higher frequencies. In reality the integrand must be multiplied by some filter function to reflect whatever filtering takes place in the measurement system. The finite bandwidth of the system ensures that the autocorrelation function will remain finite at zero time lag. For convenience, a rectangular window is often used. In many systems this can be a reasonable approximation. However, caution should be exercised as filter shapes can have some influence on the values obtained [11]. We will see such effects in section VII.

We now take two approaches towards deriving the Allan variance for continuous values of α . The first allows us to conveniently find $\sigma_y^2(\tau)$ as a function of α in the presence of a bandwidth limit. The second approach rigorously demonstrates the stationarity of $z(t, \tau)$ (and hence the convergence of the

Allan variance). The first method starts with noise in the frequency domain and uses special conditions to create a convergent autocorrelation function. To force stationarity and to eliminate the high-frequency problem, I impose upper and lower frequency limits on the noise. The PSD of the noise is given by (3) for $f_l < f < f_h$ and is zero outside this frequency range. This PSD can be substituted into (4). Defining a dimensionless argument $u = 2\pi f\tau$, the autocorrelation becomes

$$R_x(\tau) = g_\beta (2\pi\tau)^{-\beta-1} \int_{u_l}^{u_h} u^\beta \cos(u) du \quad (8)$$

for $\tau \neq 0$, and

$$R_x(0) = g_\beta \int_{f_l}^{f_h} f^\beta df \quad (9)$$

for $\tau = 0$.

These integrals must be solved separately for each integer value of β and for each region between adjacent integers. For purposes of demonstration, consider the region of $-2 < \beta < -1$. In this region $R_x(\tau)$ is given by

$$\begin{aligned} \frac{g_\beta (2\pi\tau)^{-\beta-1}}{\beta+1} & \left[u_h^{\beta+1} \cos(u_h) - u_l^{\beta+1} \cos(u_l) \right. \\ & \left. + \int_{u_l}^{u_h} u^{\beta+1} \sin(u) du \right]. \end{aligned}$$

The last term can be separated into two integrals:

$$\begin{aligned} & \int_{u_l}^{u_h} u^{\beta+1} \sin(u) du \\ & = \int_{u_l}^\infty u^{\beta+1} \sin(u) du - \int_{u_h}^\infty u^{\beta+1} \sin(u) du. \end{aligned}$$

In the limit that u_l is small, the first integral has the value

$$\int_{u_l}^\infty u^{\beta+1} \sin(u) du = -\Gamma(\beta+2) \sin(\pi\beta/2) + O(u_l^2).$$

In addition, in the limit that u_h is large, the second integral is

$$\int_{u_h}^\infty u^{\beta+1} \sin(u) du = -u_h^{\beta+1} \cos(u_h) + O(u_h^\beta).$$

$R_x(0)$ can be found with (9):

$$R_x(0) = \frac{g_\beta}{\beta+1} (f_h^{\beta+1} - f_l^{\beta+1}).$$

By substituting these values into (6) and taking the limit as f_l goes to zero and the limit as f_h becomes very large, we obtain an expression for the Allan variance:

$$\sigma_y^2(\tau) = \frac{(4 - 2^{-\beta-1})\Gamma(\beta+1) \sin(\pi\beta/2)}{(2\pi)^{\beta+1}\tau^{\beta+3}} g_\beta + \frac{3g_\beta f_h^{\beta+1}}{\tau^2(\beta+1)}. \quad (10)$$

Finally, we return to the more common parameter α to find

$$\sigma_y^2(\tau) = \frac{(2^{-\alpha+1} - 4)\Gamma(\alpha-1) \sin(\pi\alpha/2)}{(2\pi\tau)^{\alpha+1}} h_\alpha + \frac{3h_\alpha f_h^{\alpha-1}}{\tau^2(\alpha-1)(2\pi)^2}. \quad (11)$$

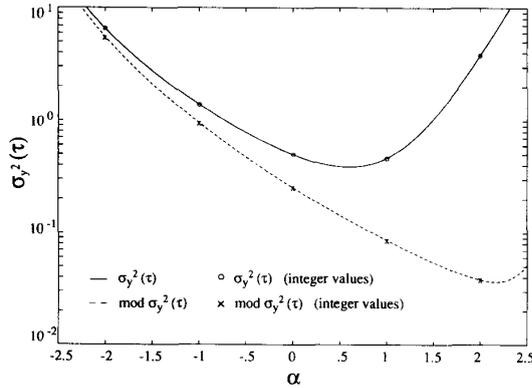


Fig. 1. This is a semi-log plot of the Allan variance (11) and the modified Allan variance (14) as functions of α . The solid line is the Allan variance for the case of $h_\alpha = 1$, $\tau = 1$ (second), and $f_h = 50$ (Hz). The open circles are plots of the five standard equations (Table I) for the same case. The dashed line is a plot of the modified Allan variance for the case of $h_\alpha = 1$, $\tau_0 = 0.01$ (seconds), $m = 100$ and $f_h = 50$ (Hz). The x 's represent the five-integer equations for the same case.

TABLE I
CONTINUOUS POWER-LAW MODEL WITH DISCRETE SAMPLING

α	$S_y(f)$	$S_x(f)$	$\sigma_y^2(\tau)^\dagger$	mod $\sigma_y^2(m\tau_0)^\dagger, \ddagger$
2	$h_2 f^2$	$\frac{h_2}{(2\pi)^2}$	$\frac{3f_h h_2}{(2\pi)^2 \tau^2}$	$\frac{3f_h h_2}{(2\pi)^2 m (m\tau_0)^2}$
1	$h_1 f$	$\frac{h_1}{(2\pi)^2 f}$	$\frac{[1.038 + 3 \ell n(2\pi f_h \tau)] h_1}{(2\pi)^2 \tau^2}$	$\frac{3.37 h_1}{(2\pi)^2 (m\tau_0)^2}$
0	h_0	$\frac{h_0}{(2\pi)^2 f^2}$	$\frac{h_0}{2\tau}$	$\frac{h_0}{4m\tau_0}$
-1	$\frac{h_{-1}}{f}$	$\frac{h_{-1}}{(2\pi)^2 f^3}$	$2 \ell n(2) h_{-1}$	$1.35 \ell n(2) h_{-1}$
-2	$\frac{h_{-2}}{f^2}$	$\frac{h_{-2}}{(2\pi)^2 f^4}$	$\frac{2\pi^2 \tau h_{-2}}{3}$	$\frac{1.65 \pi^2 m \tau_0 h_{-2}}{3}$

\dagger The assumption that $f_h \tau \gg 1$ has been made

\ddagger These are the limiting values for the case of large m

Although these equations were derived for a specific range of $\beta(\alpha)$, they are valid for all β greater than -5 ($\alpha > -3$). While these equations are undefined at most integer values, the limits as they approach the integer values exactly match the standard values given in literature [11] (Fig. 1). Thus, the five separate integer equations (Table I) can be replaced by the single equation (11). It is interesting to note that provided that f_i is small, its actual value does not enter into these equations.

The first term in each equation is negligible for values of α much above 1, while the second term is small for values of α below 1. Both terms are large and offsetting in the region of $\alpha = 1$. This equation demonstrates the necessity of imposing a bandwidth limit in the region of a $\alpha \geq 1$. The fact that the first term becomes negligible when α gets to be larger than 1 is one of the greatest drawbacks of the Allan variance. All noise processes in this region will yield the same dependence on τ , causing the Allan variance to become an ambiguous measure for distinguishing between different noise types. Other characterization tools must therefore be used (PSD or the modified Allan variance).

To prove that $z(t, \tau)$ is truly stationary, we must use (6) and the autocorrelation function for nonstationary $x(t)$. This

requires starting from the time domain and staying within it. Because of this, it is difficult to impose a rectangular frequency cutoff. Instead, I will consider the case of infinite bandwidth. In this case the Allan variance is only convergent for $-3 < \alpha < 1$. We can use the time-dependent autocorrelation given (in the limit $t \gg \tau$) by [1]:

$$R_x(t, \tau) = \frac{Q\Gamma(1+\beta)\tau^{-\beta-1}}{\Gamma(1+\beta/2)\Gamma(-\beta/2)} + \frac{Q\Gamma(-\beta-1)t^{-\beta-1}}{\Gamma(-\beta)\Gamma^2(-\beta/2)}$$

for $\beta \neq 1$ and

$$R_x(t, \tau) = Q\{\ell n(4t) - \ell n(|\tau|)\}$$

for $\beta = 1$. When these are substituted into (6) one finds that the time-dependent terms demonstrate remarkable cancellation to yield

$$\sigma_x^2(\tau) = \frac{Q\Gamma(1+\beta)(2^{-\beta-1} - 4)}{\Gamma(1+\beta/2)\Gamma(-\beta/2)\tau^{\beta+3}} + \mathbf{O}(t^{-\beta-5}).$$

The last term demonstrates that the Allan variance is in fact convergent for $\beta > -5$ which corresponds to $\alpha > -3$. This can be rewritten in terms of α , and we find that it agrees exactly with (11) for $\alpha < 1$ provided that the identification is made:

$$Q = \frac{h_\alpha}{2(2\pi)^\alpha}.$$

V. MODIFIED ALLAN VARIANCE

The modified Allan variance [8] was developed to resolve the difficulty that the Allan variance has distinguishing between noise types that have spectral densities with a $\alpha \geq 1$. It takes advantage of the fact that noise in this region has different dependencies on bandwidth. This is evident from (11). The modified Allan variance effectively changes the bandwidth with sampling time τ , by averaging adjacent measurements. The assumption made is that some number of equally spaced time measurements exist which have a basic time step of τ_0 . The modified Allan variance for time interval $\tau = m\tau_0$ (where m is an integer) is given by

$$\text{mod } \sigma_y^2(m\tau_0) \equiv \frac{1}{2(m\tau_0)^2} \left\langle \left\{ \frac{1}{m} \sum_{i=1}^m (x_{i+2m} - 2x_{i+m} + x_i) \right\}^2 \right\rangle \quad (12)$$

where x_k denotes $x(t)$ at $t = t_0 + k\tau_0$. This too is a linear combination of autocorrelation functions. We can see that the quantity inside the sum is $z(t_i, \tau_m)$, which we have already demonstrated to be stationary. The square can be expressed as a double sum, and the summations can be pulled outside the angle brackets to yield

$$\begin{aligned} \text{mod } \sigma_y^2(m\tau_0) &= \frac{1}{2(m^2\tau_0)^2} \\ &\times \sum_{i=1}^m \sum_{j=1}^m \{ \langle x_{i+2m} x_{j+2m} \rangle + 4 \langle x_{i+m} x_{j+m} \rangle \\ &+ \langle x_i x_j \rangle - 4 \langle x_{i+2m} x_{j+m} \rangle \\ &+ 2 \langle x_{i+2m} x_j \rangle - 4 \langle x_{i+m} x_j \rangle \}. \end{aligned}$$

The terms inside the curly brackets are identical to those for the Allan variance. Since this particular combination is independent of time, the number of summations can be reduced to one. The summation indices are replaced by $k = (i + j)/2$ and $\ell = |i - j|$. The modified Allan variance in terms of the discrete autocorrelation function is

$$\begin{aligned} \text{mod } \sigma_y^2(m\tau_0) &= \frac{1}{2(m^2\tau_0)^2} \\ &\times \{ m[R_x^d(k+2m, 0) + 4R_x^d(k+m, 0) \\ &+ R_x^d(k, 0) - 4R_x^d(k+3m/2, m) \\ &- 4R_x^d(k+m/2, m) + 2R_x^d(k+m, 2m)] \\ &+ \sum_{\ell=1}^{m-1} (m-\ell)[2R_x^d(k+2m, \ell) + 8R_x^d(k+m, \ell) \\ &+ 2R_x^d(k, \ell) - 4R_x^d(k+3m/2, m+\ell) \\ &- 4R_x^d(k+3m/2, m-\ell) \\ &- 4R_x^d(k+m/2, m+\ell) - 4R_x^d(k+m/2, m-\ell) \\ &+ 2R_x^d(k+m, 2m+\ell) + 2R_x^d(k+m, 2m-\ell)] \}, \end{aligned} \quad (13)$$

where the discrete autocorrelation function is defined by

$$R_x^d(k, m) \equiv R_x(k\tau_0 + t_0, m\tau_0).$$

The value of k in (13) is not important, because of the stationarity of $z(t, \tau)$. Using expressions for the autocorrelation from section IV it is possible to obtain a closed-form solution to the modified Allan variance similar to (11):

$$\begin{aligned} \text{mod } \sigma_y^2(m\tau_0) &= \frac{h_\alpha}{(2\pi m^2\tau_0)^2} \left\{ m \left[\frac{3f_h^{\alpha-1}}{(\alpha-1)} \right. \right. \\ &+ \left. \frac{(2^{-\alpha+1} - 4)\Gamma(\alpha-1)\sin(\pi\alpha/2)}{(2\pi m\tau_0)^{\alpha-1}} \right] \\ &\times \sum_{\ell=1}^{m-1} (m-\ell) \frac{\Gamma(\alpha-1)\sin(\pi\alpha/2)}{(2\pi\tau_0)^{\alpha-1}} \\ &\times [6(\ell)^{1-\alpha} - 4(m+\ell)^{1-\alpha} - 4(m-\ell)^{1-\alpha} \\ &+ (2m+\ell)^{1-\alpha} + (2m-\ell)^{1-\alpha}] \left. \right\}. \end{aligned} \quad (14)$$

While this equation is more complex than (11), it can be easily calculated with modern computers. This equation is similar to (11) in that it replaces the five integers for the modified Allan variance (Table I) with a single continuous function of α . Fig. 1 displays these results.

The modified Allan variance also has a relation to the PSD [12], [13]:

$$\text{mod } \sigma_y^2(m\tau_0) = \int_0^\infty 8S_x(f) \frac{\sin^6(\pi f m\tau_0)}{(m^2\tau_0)^2 \sin^2(\pi f\tau_0)} df. \quad (15)$$

The same comments that applied to (7) can be applied to (15). These relationships sometimes lead to the view that the Allan variance (or modified Allan variance) for a certain time interval τ represents the PSD sampled through a particular filter function. Each value of τ would then have a corresponding bandpass filter associated with it. The center of each bandpass

region would be inversely proportional to τ . It is possible to define other generalized variances by means of some filter function $|H(f)|^2$ [14]-[15]:

$$\text{General Variance} = \int_0^\infty S_x(f)|H(f)|^2 df. \quad (16)$$

where $H(f)$ may be the Fourier transform of the time-domain process or just a general filter. The two relationships (7) and (15) can be nonrigorously derived from this method.

VI. DISCRETE SAMPLING

Whether taking data or generating simulated data, one must face the reality of having only a finite number of points. This implies that it is only possible to obtain information about a finite number of frequencies (or time intervals). Usually, a set of N points, equally spaced in time, is sampled. It is hoped that these points represent instantaneous values of the continuous function being sampled. From these data points we form estimates of our various measures. To find the PSD estimate, a fast Fourier transform routine is usually invoked. For the Allan variance one assumes that the noise process is ergodic and replaces the ensemble average with a time average to form the estimate

$$\hat{\sigma}_y^2(m\tau_0, N) = \frac{1}{2(N-2m)(m\tau_0)^2} \sum_{k=1}^{N-2m} (x_{k+2m} - 2x_{k+m} + x_k)^2. \quad (17)$$

The $\hat{}$ denotes the fact that this is only an estimate. Because $z(t, \tau)$ is a Gaussian stochastic process, the estimate will be chi-square distributed, with a certain number of degrees of freedom. This estimate is designed to make maximal use of the data by using overlapping differences. This implies that the number of degrees of freedom of the estimate will be less than the $N - 2m$ points averaged. How much less will depend on the noise type. The number of degrees of freedom will be discussed in section IX. It is evident that the estimate becomes less accurate as m becomes larger.

For the modified Allan variance the estimate is

$$\begin{aligned} \text{mod } \hat{\sigma}_y^2(m\tau_0, N) &= \frac{1}{2(N-3m+1)(m^2\tau_0)^2} \sum_{k=0}^{N-3m} \\ &\times \left[\sum_{j=1}^m (x_{k+j+2m} - 2x_{k+j+m} + x_{k+j}) \right]^2. \end{aligned} \quad (18)$$

This estimate also makes use of overlapping data, and it becomes even worse at large m .

In all of the previous sections the equations were derived for the case of continuous functions. An important question is: What are the correct equations for discretely sampled data? Conventional noise simulation routines seek to generate noise which has a discrete PSD that samples from the continuous PSD. However, discretely sampling in the time domain for a finite period will cause some distortion of the PSD. It can be argued [1] that instead it should be the discrete autocorrelation function which directly samples the continuous function.

Because the discrete Allan and modified Allan variances would be formed from linear combinations of the discrete autocorrelation function, they too should directly sample the continuous functions given above. This holds true, except for the issue of bandwidth. The autocorrelation function will depend on the bandwidth of the system. The next question is: What is the correct way to implement the filtering? The simple response is to put the filter cut-off out at some reasonably high frequency, where it has the least effect. However, when examining the PSD, it is better to put the filter cutoff before the Nyquist frequency to prevent higher frequency components from wrapping themselves down into the lower frequency region. The problem this creates is that the upper bound on the integrals in (8) and (9) becomes $u_h = m\pi$, which for the first few values of m can hardly be considered to be approaching infinity. What then is the correct approach to filtering?

Recent papers [1], [2] demonstrate that for f^β noise discretely sampled in the time domain, the power spectral density will have the form

$$S_x^d(f) = g_\beta \left[\frac{\sin(\pi f \tau_0)}{\pi \tau_0} \right]^\beta, \quad (19)$$

where the superscript d specifically denotes the discrete sampling. This is consistent with other work [7] which states that for the case of discrete sampling the relationship between the spectral density of $x(t)$ and the spectral density of $y(t)$ becomes

$$S_x^d(f) = \frac{\tau_0^2 S_y^d(f)}{4 \sin^2(\pi f \tau_0)}. \quad (20)$$

The corresponding discrete autocorrelation for $\beta > -1$ is then given by [1]

$$R_x^d(m) = \frac{g_\beta \Gamma(m - \beta/2) \Gamma(1 + \beta)}{2(2\pi)^\beta \tau_0^{\beta+1} \Gamma(m + 1 + \beta/2) \Gamma(1 + \beta/2) \Gamma(-\beta/2)}. \quad (21)$$

Some form of filtering has been implicitly assumed as this expression is finite for zero lag. Following the same steps as in section IV, one can find the Allan variance

$$[\sigma_x^d(m\tau_0)]^2 = \frac{R_x^d(0)}{(m\tau_0)^2} \left[3 - 4 \frac{\Gamma(m - \beta/2) \Gamma(1 + \beta/2)}{\Gamma(m + 1 + \beta/2) \Gamma(-\beta/2)} + \frac{\Gamma(2m - \beta/2) \Gamma(1 + \beta/2)}{\Gamma(2m + 1 + \beta/2) \Gamma(-\beta/2)} \right]. \quad (22)$$

If one takes the limit as m becomes large, one finds

$$\lim_{m \rightarrow \infty} [\sigma_y^d(m\tau_0)]^2 \cong \frac{(4 - 2^{-\beta-1}) \Gamma(\beta + 1) \sin(\pi\beta/2)}{(2\pi)^{\beta+1} (m\tau_0)^{\beta+3}} g_\beta + \frac{3g_\beta \Gamma(\beta + 1)}{2(2\pi)^\beta \tau_0^{\beta+1} (m\tau_0)^2 \Gamma^2(1 + \beta/2)}.$$

The first term is in exact agreement with (10). The second term has the correct time dependence, but the value for f_h is dependent on β , and for some values is undefined. This is because the implicit filter in (21) is obviously not rectangular. Despite this, (22) is valid for all values of $\beta > -5$ and can be

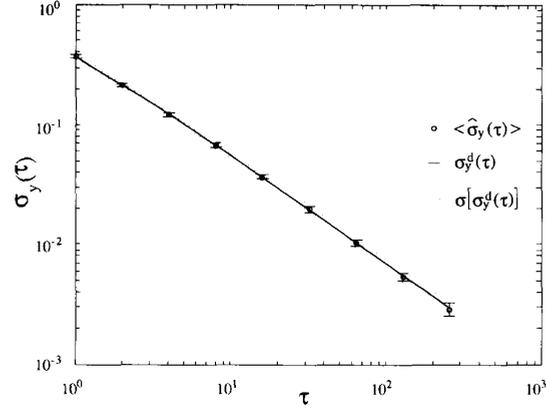


Fig. 2. This is a plot of the discrete Allan variance and mean Allan variance estimates for $\alpha = 1$. The solid line depicts (23) for the case of $h_\alpha = 1$ and $\tau_0 = 1$. The open circles represent the mean Allan variance estimates for 100 generated noise sequences. The variances of the 100 generated Allan variance estimates are represented by the error bars. The dotted lines show the expected variance (29).

rewritten in terms of the more familiar parameter α :

$$[\sigma_y^d(m\tau_0)]^2 = \frac{h_\alpha \pi \Gamma(\alpha - 1)}{m^2 (2\pi\tau_0)^{\alpha+1} \Gamma^2(\alpha/2)} \times \left[3 - 4 \frac{\Gamma(m + 1 - \alpha/2) \Gamma(\alpha/2)}{\Gamma(m + \alpha/2) \Gamma(1 - \alpha/2)} + \frac{\Gamma(2m + 1 - \alpha/2) \Gamma(\alpha/2)}{\Gamma(2m + \alpha/2) \Gamma(1 - \alpha/2)} \right]. \quad (23)$$

This equation is convergent for $\alpha > -3$. It is plotted for two different values of α in Figs. 2 and 3. Similarly we can find the modified Allan variance

$$\text{mod} [\sigma_y^d(m\tau_0)]^2 = \frac{h_\alpha Q \pi \Gamma(\alpha - 1)}{m^4 (2\pi\tau_0)^{\alpha+1} \Gamma^2(\alpha/2)} \times \left\{ m \left[3 - 4 \frac{\Gamma(m + 1 - \alpha/2) \Gamma(\alpha/2)}{\Gamma(m + \alpha/2) \Gamma(1 - \alpha/2)} + \frac{\Gamma(2m + 1 - \alpha/2) \Gamma(\alpha/2)}{\Gamma(2m + \alpha/2) \Gamma(1 - \alpha/2)} \right] + \sum_{\ell=1}^{m-1} (m - \ell) \frac{\Gamma(\alpha/2)}{\Gamma(1 - \alpha/2)} \left[6 \frac{\Gamma(\ell + 1 - \alpha/2)}{\Gamma(\ell + \alpha/2)} - 4 \frac{\Gamma(m + \ell + 1 - \alpha/2)}{\Gamma(2m + \ell + \alpha/2)} - 4 \frac{\Gamma(m - \ell + 1 - \alpha/2)}{\Gamma(m - \ell + \alpha/2)} + \frac{\Gamma(2m + \ell + 1 - \alpha/2)}{\Gamma(2m + \ell + \alpha/2)} + \frac{\Gamma(2m - \ell + 1 - \alpha/2)}{\Gamma(2m - \ell + \alpha/2)} \right] \right\}. \quad (24)$$

This equation also approaches the continuous expression in the limit that m becomes large.

VII. DISCRETE TRANSFER FUNCTIONS

The discrete variances in Section VI can also be related back to the discrete PSD. Rather than providing a rigorous

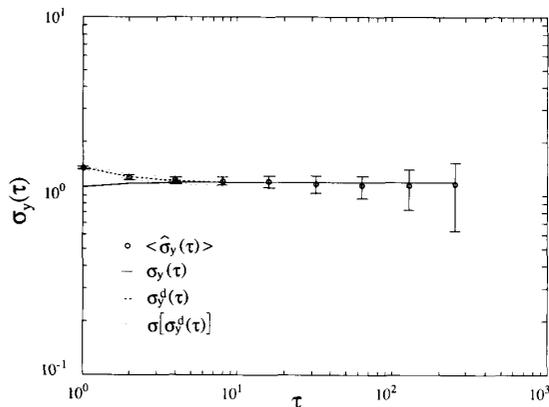


Fig. 3. This is a plot of the Allan variance and mean Allan variance estimates for $\alpha = -1$. The dashed line depicts the discrete Allan variance (23) for the case of $h_\alpha = 1$ and $\tau_0 = 1$. The solid line depicts the continuous Allan variance (11) for the same case with $f_h = 0.5$. The open circles represent the mean Allan variance estimates for 100 generated noise sequences. The variances of the 100 generated Allan variance estimates are represented by the error bars. The dotted lines show the expected variance (29). As expected, the estimates follow the discrete model rather than the continuous model.

derivation, we will adopt the generalized variance approach of (16) and demonstrate equivalence. First we must choose the expression for $|H(f)|^2$. For this, we will start with the discrete time-domain process

$$h_k = h(t_k) \equiv \frac{1}{\sqrt{2m\tau_0}} (\delta_{k,2m} - 2\delta_{k,m} + \delta_{k,0}),$$

where $\delta_{k,m}$ is the Kronecker-delta function (= 1 when the indices are identical, = 0 otherwise). This choice was made because of its resemblance to (5). Taking the absolute square of the discrete Fourier transform we find

$$|H(f)|^2 = \frac{8}{(m\tau_0)^2} \sin^4(m\pi f\tau_0),$$

which agrees with the continuous transfer function for the Allan variance. In the case of discrete sampling, we have no information about frequencies higher than the Nyquist frequency. All of the power in that region has been folded down into lower frequencies. Therefore, integration in the frequency-domain only takes place up to the Nyquist frequency. When all of this is incorporated into the discrete version of (16) with appropriate frequency limits we find

$$\begin{aligned} [\sigma_y^d(m\tau_0)]^2 &= \int_0^{\frac{1}{2\tau_0}} g\beta \left[\frac{\sin(\pi f\tau_0)}{\pi\tau_0} \right]^\beta \\ &\quad \times \frac{8}{(m\tau_0)^2} \sin^4(m\pi f\tau_0) df. \end{aligned}$$

When this is evaluated (with help from [16] 3.892) we find that it is exactly equivalent to (22).

This procedure can be repeated for the modified Allan variance, this time with some interesting results. The time-domain process in this case is

$$h_k \equiv \frac{1}{\sqrt{2m^2\tau_0}} \sum_{i=0}^{m-1} (\delta_{k,i+2m} - 2\delta_{k,i+m} + \delta_{k,i}).$$

which leads to

$$|H(f)|^2 = \frac{8 \sin^6(m\pi f\tau_0)}{(m^2\tau_0)^2 \sin^2(\pi f\tau_0)},$$

which again agrees with the continuous expression. Our expression for the modified Allan variance then becomes

$$\begin{aligned} \text{mod} [\sigma_y^d(m\tau_0)]^2 &= \int_0^{\frac{1}{2\tau_0}} g\beta \left[\frac{\sin(\pi f\tau_0)}{\pi\tau_0} \right]^\beta \\ &\quad \times \frac{8 \sin^6(m\pi f\tau_0)}{(m^2\tau_0)^2 \sin^2(\pi f\tau_0)} df. \end{aligned}$$

When this integral is evaluated it leads to

$$\begin{aligned} \text{mod} [\sigma_y^d(m\tau_0)]^2 &= \frac{g\beta \Gamma(\beta-1) \sin(\pi\beta/2)}{(m^2\tau_0)^2 (2\pi\tau_0)^{\beta+1}} \\ &\quad \times \left[10 \frac{\Gamma(1-\beta/2)}{\Gamma(\beta/2)} - 15 \frac{\Gamma(m+1-\beta/2)}{\Gamma(m+\beta/2)} \right. \\ &\quad \left. + 6 \frac{\Gamma(2m+1-\beta/2)}{\Gamma(2m+\beta/2)} - \frac{\Gamma(3m+1-\beta/2)}{\Gamma(3m+\beta/2)} \right], \end{aligned} \quad (25)$$

which has a considerably simpler form than earlier expressions for the modified Allan variance. We can express this as a function of α :

$$\begin{aligned} \text{mod} [\sigma_y^d(m\tau_0)]^2 &= \frac{h_\alpha \Gamma(\alpha-3)}{2(2\pi)^\alpha m^4 \tau_0^{\alpha+1} \Gamma^2(\alpha/2-1)} \\ &\quad \times \left[10 - 15 \frac{\Gamma(m+2-\alpha/2)\Gamma(\alpha/2-1)}{\Gamma(m-1+\alpha/2)\Gamma(2-\alpha/2)} \right. \\ &\quad \left. + 6 \frac{\Gamma(2m+2-\alpha/2)\Gamma(\alpha/2-1)}{\Gamma(2m-1+\alpha/2)\Gamma(2-\alpha/2)} \right. \\ &\quad \left. - \frac{\Gamma(3m+2-\alpha/2)\Gamma(\alpha/2-1)}{\Gamma(3m-1+\alpha/2)\Gamma(2-\alpha/2)} \right]. \end{aligned} \quad (26)$$

While we have been unable to prove that this expression is equal to (24) in the general case, the two equations reduce to the same form for all specific values of m examined. This then is a much simpler and much more easily computable equation for the modified Allan variance than any previous expression. Note that it contains only four terms as opposed to the $5m$ terms in (24). This equation is plotted for $\alpha = 2$ in Fig. 4.

The above relationships provide a clue as to the implicit filtering of (21). One can see that the filter is nearly rectangular and is in fact of the form

$$|H(f)|^2 = \left[\frac{\pi f\tau_0}{\sin(\pi f\tau_0)} \right]^\beta$$

for the region $f \leq 1/(2\tau_0)$ and has zero contribution beyond the Nyquist frequency. For white phase noise ($\beta = 0$) it is exactly a rectangular filter with cutoff at the Nyquist frequency. As β becomes smaller than 0, this window becomes distorted near the Nyquist frequency. However, this effect is unimportant as most of the noise power will be at the lower frequency end of the spectrum, away from the distortion.

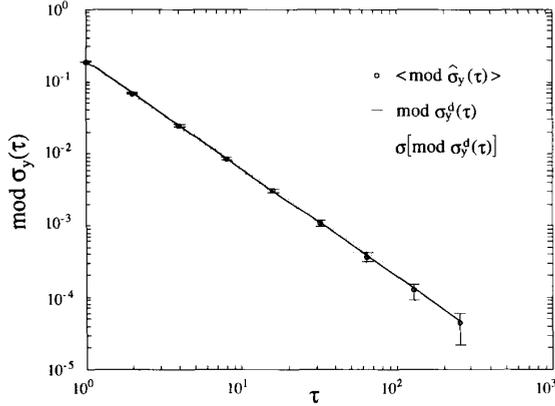


Fig. 4. This is a plot of the discrete modified Allan variance and mean modified Allan variance estimates for $\alpha = 2$. The solid line depicts (26) for the case of $h_\alpha = 1$ and $\tau_0 = 1$. The open circles represent the mean modified Allan variance estimates for 100 generated noise sequences. The variances of the 100 generated modified Allan variance estimates are represented by the error bars. The dotted lines show the expected variance (32).

For $\beta > -1$ there is little difference between the discrete and continuous equations. For β below this region there is a noticeable difference between the equations for the first few values of m . However, the discrete equations rapidly converge (by $m \approx 8$, see Fig. 3) to the continuous equations. In a real system there is often more than one noise process affecting the oscillator. Often the dominant contribution to the Allan and modified Allan variances for small values of τ comes from a noise in the vicinity of white phase noise. Other processes (flicker frequency, random walk frequency) dominate at large values of τ . Because the discrete variances approach the continuous variances for all values of τ as β approaches 0, the continuous and discrete equations cannot be distinguished for real systems. The practical implications of this are that as long as noise in the vicinity of $\beta = 0$ dominates for small values of τ , (23) is as good as (11), and (26) is better than (14) due to its relative simplicity.

VIII. FINITE DATA LENGTH

The Allan variance and modified Allan variance estimates defined in (17) and (18) will have some uncertainty due to the fact that they are formed with only a finite number of data points. Provided the noise processes are ergodic, the expected means of these estimates approach their respective variances (i.e. are unbiased). We are also interested in the variances of these random variables. This will enable us to find the uncertainty of our estimates and place confidence limits on them. This effect has been discussed elsewhere [17], [18], but it is worthwhile to examine it again for the case of arbitrary α . I shall follow the derivation used by Lesage and Adouin, except that I will use fully overlapping estimates. First, the variance of the Allan variance estimate is defined:

$$\sigma^2[\hat{\sigma}_y^2(m\tau_0, N)] \equiv \left\langle \left[\hat{\sigma}_y^2(m\tau_0, N) - \langle \hat{\sigma}_y^2(m\tau_0, N) \rangle \right]^2 \right\rangle. \quad (27)$$

Then, provided that $z(t, \tau)$ is normally distributed, we can use

$$\left\langle [\hat{\sigma}_y^2(m\tau_0, N)]^2 \right\rangle = 3 \left\langle \hat{\sigma}_y^2(m\tau_0, N) \right\rangle^2.$$

By substituting this equation and (5) into (27), we can find the variance of the Allan variance estimate in terms of the discrete autocorrelation function:

$$\begin{aligned} \sigma^2[\hat{\sigma}_y^2(m\tau_0, N)] &= \frac{1}{2(N-2m)^2(m\tau_0)^4} \\ &\times \sum_{\ell=-N+2m+1}^{N-2m-1} (N-2m-|\ell|) \\ &\times \{ R_x^d(k+2m, |\ell|) + 4R_x^d(k+m, |\ell|) \\ &+ R_x^d(k, |\ell|) - 2R_x^d(k+3m/2, |m+\ell|) \\ &- 2R_x^d(k+3m/2, |m-\ell|) \\ &- 2R_x^d(k+m/2, |m+\ell|) \\ &- 2R_x^d(k+m/2, |m-\ell|) \\ &+ R_x^d(k+m, |2m+\ell|) \\ &+ R_x^d(k+m, |2m-\ell|) \}^2. \quad (28) \end{aligned}$$

This is similar in appearance to the modified Allan variance. We can substitute in the specific case of (21) to find

$$\begin{aligned} \sigma^2[\hat{\sigma}_y^2(m\tau_0, N)] &= \frac{h_\alpha^2 \Gamma^2(\alpha-1) \sin^2(\alpha\pi/2)}{2(2\pi\tau_0)^{2\alpha+2}(N-2m)^2 m^4} \\ &\times \sum_{\ell=-N+2m+1}^{N-2m-1} (N-2m-|\ell|) \\ &\times \left\{ 6 \frac{\Gamma(|\ell|+1-\alpha/2)}{\Gamma(|\ell|+\alpha/2)} \right. \\ &- 4 \frac{\Gamma(|m+\ell|+1-\alpha/2)}{\Gamma(|m+\ell|+\alpha/2)} \\ &- 4 \frac{\Gamma(|m-\ell|+1-\alpha/2)}{\Gamma(|m-\ell|+\alpha/2)} \\ &+ \frac{\Gamma(|2m+\ell|+1-\alpha/2)}{\Gamma(|2m+\ell|+\alpha/2)} \\ &\left. + \frac{\Gamma(|2m-\ell|+1-\alpha/2)}{\Gamma(|2m-\ell|+\alpha/2)} \right\}^2. \quad (29) \end{aligned}$$

While this equation does not provide an intuitive feel for the variance of the Allan variance estimate, it can be readily programmed on a computer. This allows confidence intervals to be placed on the Allan variance estimate for noise processes with $\alpha > -3$. The measured variances are compared to these calculated values in Figs. 2 and 3.

The variance of the modified Allan variance estimate may be found in a similar manner. This variance is defined by

$$\sigma^2[\text{mod}\hat{\sigma}_y^2(m\tau_0, N)] \equiv \left\langle [\text{mod}\hat{\sigma}_y^2(m\tau_0, N) - \langle \text{mod}\hat{\sigma}_y^2(m\tau_0, N) \rangle]^2 \right\rangle. \quad (30)$$

Following the same procedure as above we can find

$$\begin{aligned} \sigma^2[\text{mod } \hat{\sigma}_y^2(m\tau_0, N)] &= \frac{1}{2(N-3m+1)^2(m^2\tau_0)^4} \\ &\times \sum_{\ell=-N+3m}^{N-3m} (N-3m+1-|\ell|) \\ &\times \left\{ \sum_{q=-m+1}^{m-1} (m-|q|) \right. \\ &\times [R_x^d(k+2m, |\ell+q|) \\ &+ 4R_x^d(k+m, |\ell+q|) + R_x^d(k, |\ell+q|) \\ &- 2R_x^d(k+3m/2, |m+\ell+q|) \\ &- 2R_x^d(k+3m/2, |m-\ell-q|) \\ &- 2R_x^d(k+m/2, |m+\ell+q|) \\ &- 2R_x^d(k+m/2, |m-\ell-q|) \\ &+ R_x^d(k+m, |2m+\ell+q|) \\ &\left. + R_x^d(k+m, |2m-\ell-q|) \right\}^2. \quad (31) \end{aligned}$$

The discrete autocorrelation function (21) may also be substituted into this equation to yield

$$\begin{aligned} \sigma^2[\text{mod } \hat{\sigma}_y^2(m\tau_0, N)] &= \frac{h_\alpha^2 \Gamma^2(\alpha-1) \sin^2(\alpha\pi/2)}{2(2\pi\tau_0)^{2\alpha+2} (N-3m+1)^2 m^8} \\ &\times \sum_{\ell=-N+3m}^{N-3m} (N-3m+1-|\ell|) \\ &\times \left\{ \sum_{q=-m+1}^{m-1} (m-|q|) \right. \\ &\times \left[6 \frac{\Gamma(|\ell+q|+1-\alpha/2)}{\Gamma(|\ell+q|+\alpha/2)} \right. \\ &- 4 \frac{\Gamma(|m+\ell+q|+1-\alpha/2)}{\Gamma(|m+\ell+q|+\alpha/2)} \\ &- 4 \frac{\Gamma(|m-\ell-q|+1-\alpha/2)}{\Gamma(|m-\ell-q|+\alpha/2)} \\ &+ \frac{\Gamma(|2m+\ell+q|+1-\alpha/2)}{\Gamma(|2m+\ell+q|+\alpha/2)} \\ &\left. \left. + \frac{\Gamma(|2m-\ell-q|+1-\alpha/2)}{\Gamma(|2m-\ell-q|+\alpha/2)} \right] \right\}^2. \quad (32) \end{aligned}$$

This also can be readily programmed on the computer. Unfortunately, because of the double sum, calculation time may be quite long for large values of N and m . In Fig. 4, the measured variances are again compared to the calculated values.

IX. RESULTS

In order to verify many of the equations presented here, a comparison was made to the results of computer simulations. Noise was generated corresponding to both integer and

TABLE II
NUMBER OF DEGREES OF FREEDOM FOR THE ALLAN VARIANCE ESTIMATE FOR $N = 1025$ AND VARIOUS VALUES OF m

m	White phase	Flicker phase	White frequency	Flicker frequency	Random-walk frequency
1	526.6	590.2	682.6	829.4	1024
2	525.9	554.3	584.3	606.2	526.0
4	524.4	453.2	354.7	306.8	244.0
8	521.4	336.1	186.5	150.0	118.4
16	515.3	232.0	93.53	73.51	58.10
32	503.2	150.8	45.83	35.76	28.25
64	479.2	92.31	21.84	16.97	13.35
128	432.8	52.12	9.852	7.616	5.922
256	355.2	26.19	4.016	3.012	2.246

noninteger values of α . The PSD estimates of the generated processes x_k were found to match the values predicted by (19). The Allan variance and modified Allan variance estimates were formed and averaged for 100 generated time series for many values of α . The means found in each case were in excellent agreement with (23) and (26). The variances of the estimates were also in agreement with (29) and (31). Figs. 2–4 display some of these results for the Allan and modified Allan variances for some different values of α .

The variance of the Allan variance estimate or modified Allan variance estimate is not always the best measure to use for placing confidence limits on the estimates. Because the estimates follow a chi-square distribution rather than a normal distribution, the upper and lower uncertainties should be determined separately and according to the proper distribution. We then need to find the number of degrees of freedom for each estimate. A much more complete treatment of this subject may be found elsewhere [5], [19]. Fortunately we can find the number of degrees of freedom from the variance of the estimates and the true variance. They are obtained from the relationship

$$d.f. = \frac{2(\sigma^2)^2}{\sigma^2[\hat{\sigma}^2]}. \quad (33)$$

This allows us to correctly calculate the upper and lower confidence limits for the estimates of our time-domain measures. Table II lists the degrees of freedom obtained for the Allan variance from (23) and (29). When these are compared to the values listed elsewhere [19], we find they are nearly identical for the cases of white phase, white frequency and random-walk frequency. The discrepancies are due to the fact that the algorithms used here were not made to accommodate integer values of α . Instead, an approximation of α very near the integer value is used. Thus the values listed in the subsequent tables are only accurate to about 0.1%. The flicker values differ from the values stated in literature [19]. This is most likely because the flicker values listed in the reference were obtained by empirically observing the variance of the estimates for computer-simulated noise. The values presented here are a numerical evaluation of an analytical model.

The number of degrees of freedom can be found for the modified Allan variance in a similar manner by substituting

TABLE III
NUMBER OF DEGREES OF FREEDOM FOR THE MODIFIED ALLAN VARIANCE
ESTIMATE FOR $N = 1025$ AND VARIOUS VALUES OF m

m	White phase	Flicker phase	White frequency	Flicker frequency	Random-walk frequency
1	526.6	590.2	682.6	829.4	1024
2	447.7	497.3	516.0	524.5	442.4
4	299.0	262.2	252.7	245.6	200.8
8	158.3	128.2	122.9	119.9	97.33
16	79.08	62.37	59.94	58.60	47.43
32	38.22	29.90	28.77	28.11	22.7
64	17.65	13.76	13.24	12.90	10.36
128	7.413	5.753	5.511	5.331	4.208
256	2.861	2.079	1.812	1.568	1.292

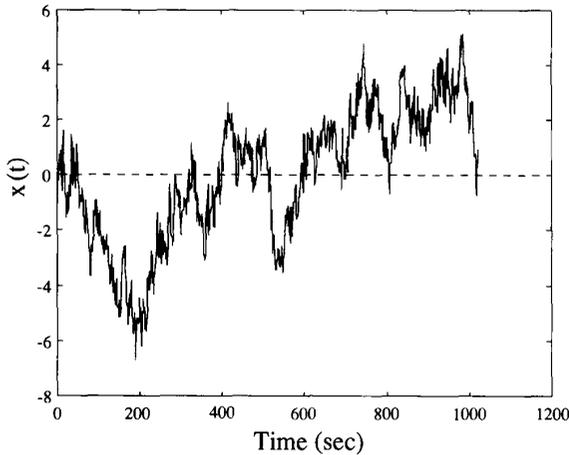


Fig. 5. This displays a single generated noise sequence $x(t_k)$ for the case of $\alpha = 0.4$, $h_\alpha = 1$ and $\tau_0 = 1$. The total number of points generated is 1024.

(26) and (31) into (33). The results are listed in Table III. It is evident that the numbers decrease much more rapidly for increasing m than they do for the Allan variance estimate.

Figs. 5 through 8 depict the results of many of the equations derived in this paper. These figures serve to summarize and illustrate the capabilities of some of the methods presented. A time series was generated with a new method [1]-[2] corresponding to $\alpha = 0.4$. Each point is separated by $\tau_0 = 1$ second, and a total of 1024 points was generated (Fig. 5). Visual inspection of the time series fails to show any obvious systematic trend. It is then appropriate to form the PSD estimate. The time series was multiplied by a Hanning window and discretely Fourier transformed. Fig. 6 demonstrates that the estimate follows the general trend predicted by (19). No sharp peaks are evident in the spectrum, so we return to the time domain. Figs. 7 and 8 depict the Allan variance estimates and modified Allan variance estimates respectively. The error bars indicate the 90% confidence intervals using the chi-square distribution with the number of degrees of freedom determined from (33). The figures demonstrate that the confidence limits include the expected theoretical results.

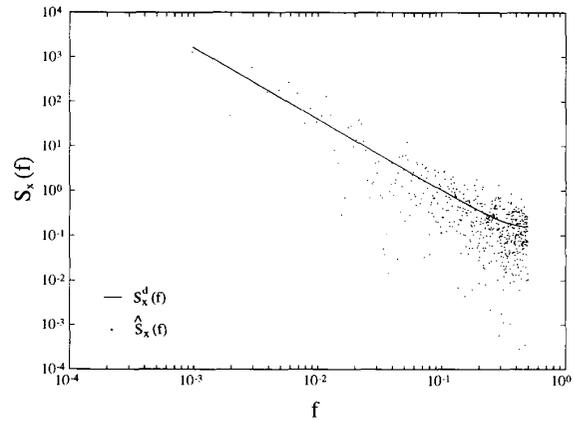


Fig. 6. This is a plot of the PSD estimate obtained via a discrete Fourier transform. The process was first multiplied by a Hanning window. The solid line is the expected PSD (19) for noise shown in Fig. 5.

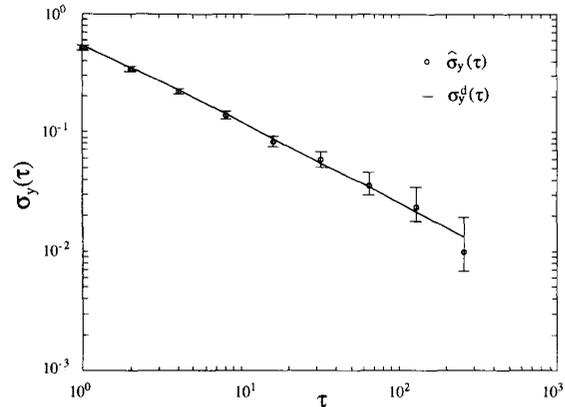


Fig. 7. This shows the Allan variance estimate formed for the sequence in Fig. 5. The solid line is the expected discrete Allan variance (23). The error bars correspond to the 90% confidence intervals using a chi-square distribution. The number of degrees of freedom for each point was calculated using (23), (29), and (33).

X. CONCLUSION

In this paper, new equations for the Allan variance and the modified Allan variance have been derived. In each case a single equation, continuous in α , has been presented. These equations can be used to augment or to replace the five standard integer equations for each variance. The recent introduction of a discrete form of the autocorrelation [1] has allowed for these methods to be extended to discrete processes. It has been shown that the discrete equations converge to their continuous counterparts. The differences between the discrete and continuous equations are in practice indistinguishable. However, the discrete equations offer some definite advantages. One such advantage is (26) has a much simpler form than its continuous analog (14).

This paper also re-examines the uncertainty of the variance estimates. The variances of the variance estimates have been derived for both the Allan variance and the modified Allan variance. Equations (29) and (32) were formed for estimates

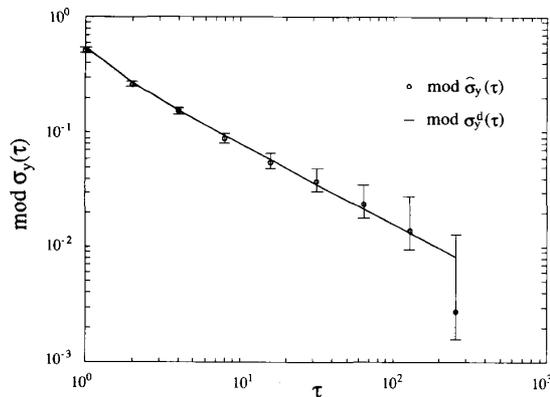


Fig. 8. This shows the modified Allan variance estimate formed for the sequence in Fig. 5. The solid line is the expected discrete modified Allan variance (26). The error bars correspond to the 90% confidence intervals using a chi-square distribution. The number of degrees of freedom for each point was calculated using (26), (32), and (33).

with overlapping data samples and are continuous functions of α . These equations allow for the calculation of the number of degrees of freedom for arbitrary α . From these results firm confidence limits can be placed on the estimates using the chi-square distribution.

The results presented in this paper can be used to more accurately determine the functional dependence of stability measures on α . This ability can be extremely useful for comparing experimental data to theoretical models. These comparisons should lead to a better understanding of noise sources which affect high-performance oscillators.

ACKNOWLEDGMENT

The author wishes to thank N. J. Kasdin for his numerous helpful discussions on this work and for his noise generation algorithm. The author is also grateful to Professor J. Turneaure and D. W. Allan for their valuable comments regarding this manuscript.

REFERENCES

- [1] N. J. Kasdin, "Discrete simulation of colored noise and stochastic processes and $1/f^\alpha$ power law generation," in *preparation for Proc. of the IEEE*.
- [2] N. J. Kasdin and T. Walter, "Discrete simulation of power law noise," in *Proc. 1992 IEEE Freq. Contr. Symp.*, pp. 274–283, May 1992.
- [3] J. A. Barnes, et al., "Characterization of Frequency Stability," *IEEE Trans. Instrum. Meas.*, vol. IM-20, pp. 105–120, May 1971.
- [4] J. A. Barnes and D. W. Allan, "A statistical model of flicker noise," in *Proc. IEEE*, vol. 54, pp. 176–178, Feb. 1966.
- [5] D. A. Howe, D. W. Allan and J. A. Barnes, "Properties of signal sources and measurement methods," in *Proc. 35th Ann. Freq. Contr. Symp.*, pp. 1–47, May 1981.
- [6] F. J. Harris, "On the use of windows for harmonic analysis with the discrete fourier transform," in *Proc. IEEE*, vol. 66, pp. 51–83, Jan. 1978.
- [7] D. B. Percival, "Characterization of frequency stability: frequency-domain estimation of stability measures," in *Proc. IEEE*, vol. 79, pp. 961–972, July 1991.
- [8] D. W. Allan and J. A. Barnes, "A modified allan variance with increased oscillator characterization ability," in *Proc. 35th Ann. Freq. Contr. Symp.*, pp. 470–474, May 1981.
- [9] D. W. Allan, "Statistics of atomic frequency standards," in *Proc. IEEE*, vol. 54, pp. 221–230, Feb. 1966.
- [10] L. S. Cutler and C. L. Searle, "Some aspects of the theory and measurement of frequency fluctuations in frequency standards," in *Proc. IEEE*, vol. 54, pp. 135–154, Feb. 1966.
- [11] D. B. Sullivan, D. W. Allan, D. A. Howe and F. L. Walls, "Characterization of clocks and oscillators," *NIST Technical Note*, 1337, 1990.
- [12] P. Lesage and T. Ayi, "Characterization of frequency stability: analysis of the modified Allan variance and properties its estimate," *IEEE Trans. Instrum. Meas.*, vol. IM-33, pp. 332–336, Dec. 1984.
- [13] J. Rutman and F. L. Walls, "Characterization of frequency stability in precision frequency sources," in *Proc. IEEE*, vol. 79, pp. 952–960, July 1991.
- [14] 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.
- [15] J. Rutman, "Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress," in *Proc. IEEE*, vol. 66, pp. 1048–1075, Sept. 1978.
- [16] I. S. Gradshteyn and I. M. Ryzhik, *Tables of Integrals Series and Products*, New York: Academic Press, 1965, 3.892, p. 476.
- [17] P. Lesage and C. Audoin, "Characterization of frequency stability: uncertainty due to the finite number of measurements," *IEEE Trans. Instrum. Meas.*, vol. IM-22, pp. 157–161, June 1973.
- [18] K. Yoshimura, "Characterization of frequency stability: uncertainty due to the autocorrelations of the frequency fluctuations," *IEEE Trans. Instrum. Meas.*, vol. IM-27, pp. 1–7, Mar. 1978.
- [19] S. R. Stein, *Precision Frequency Control*, vol. 2, edited by E. A. Gerber and A. Ballato, New York: Academic Press, 1985, pp. 191–416.
- [20] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, New York: McGraw-Hill, 1991.
- [21] T. Walter, "A multi-variance analysis in the time domain," in *Proc. of the 1992 PTTI Plan. App. Meeting*, Dec. 1992.

Todd Walter was born in New York in 1965. He received his B.S. in physics from Rensselaer Polytechnic Institute in 1987 and his M.S. in applied physics from Stanford University in 1989.

Since 1988 he has been with Stanford's Relativity Gyroscope program (Gravity Probe B) where he earned his Ph.D. degree in 1993.

His research concentrates on using gyroscopes as very precise mechanical clocks to perform a test of the Einstein Equivalence Principle, ground-based testing and analysis of the GPB gyroscopes, and analysis of noise affecting frequency sources.