Considerations on the Measurement of the Stability of Oscillators with Frequency Counters

Samuel T. Dawkins, John J. McFerran and André N. Luiten

***Abstract*— The most common time-domain measure of frequency stability, the Allan variance, is typically estimated using a frequency counter. Close examination of the operation of modern high-resolution frequency counters shows that they do not make measurements in the way commonly assumed. The consequence is that the results typically reported by many laboratories using these counters are not, in fact, the Allan variance, but a distorted representation (as pointed out in a previous publication [1]). We elucidate the action of these counters by consideration of their operation in the Fourier domain, and demonstrate that the difference between the actual Allan variance and that delivered by these counters can be very significant for some types of oscillators. We also discuss ways to avoid, or account for, a distorted estimation of Allan variance.**

***Index Terms*— frequency stability, frequency counter, Allan variance, frequency standards.**

I. Introduction

THIS paper considers the issue of characterizing the frequency fluctuations of an oscillator or clock. Let us take a clock with an output signal of the form

$$V(t) = A\sin(2\pi\, y(t) f_0\, t), \tag{1}$$

where f_0 is the average frequency of the oscillator over the entire measurement period and $y(t)$ is the instantaneous fractional frequency of the oscillator. The frequency stability of this oscillator can be characterized with the commonly used Allan variance (or its square root, the Allan frequency deviation), which is defined as [2]:

$$\sigma_{\mathrm{A}}^2(\tau) \equiv \langle \frac{1}{2}(\bar{y}_{k+1} - \bar{y}_k)^2 \rangle, \tag{2}$$

where we define the kth sample of the normalized frequency, averaged over some measurement time (sometimes called the integration time), τ, as

$$\bar{y}_k = \frac{1}{\tau}\int_{t_k}^{t_k+\tau} y(t)\, dt. \tag{3}$$

In practice, the ensemble average in (2) is usually replaced by a summation of m consecutive measurements of the kernel shown inside the angular brackets:

$$\sigma_{\mathrm{A}}^2(\tau) \approx \frac{1}{m}\sum_{k=1}^{m} \frac{(\bar{y}_{k+1} - \bar{y}_k)^2}{2}. \tag{4}$$

The integral in (3) corresponds to a single (normalized) measurement of a traditional frequency counter for a selected measurement time, τ, usually referred to as the gate time in counter nomenclature. Following the approach of Rubiola [1], a single counter measurement over a gate time of τ can be written as a weighted integral:

$$f_{\mathrm{meas,k}} = \int_{-\infty}^{\infty} y(t) f_0\; w_{\Pi}(t - t_k)\, dt = f_0\, \bar{y}_k, \tag{5}$$

where $w_{\Pi}(t) = 1/\tau$ for $0 < t < \tau$ and 0 elsewhere, referred to in this paper as a Π-estimator. Similarly, we can rewrite (2) as:

$$\sigma_{\mathrm{A}}^2(\tau) = \langle \left[\int_{-\infty}^{\infty} y(t) w_A(t - t_k)\, dt\right]^2 \rangle, \tag{6}$$

where we have combined the consecutive integrations and scaling factors into a single temporal windowing function

$$w_A(t) = \begin{cases} -\frac{1}{\sqrt{2}\tau} & 0 < t \le \tau \\ \frac{1}{\sqrt{2}\tau} & \tau < t \le 2\tau, \\ 0 & \text{elsewhere} \end{cases} \tag{7}$$

which we will refer to as the Allan variance weighting function (shown in Fig. 1a).

However, some frequency counters (including those that are commonly found in many laboratories) do not implement the frequency measurement as described above; in other words their operation cannot be modeled by (5). These counters implement an internal averaging algorithm, which changes the shape of the temporal windowing function. This results in a strong suppression of any frequency fluctuations that have a Fourier frequency significantly higher than the reciprocal

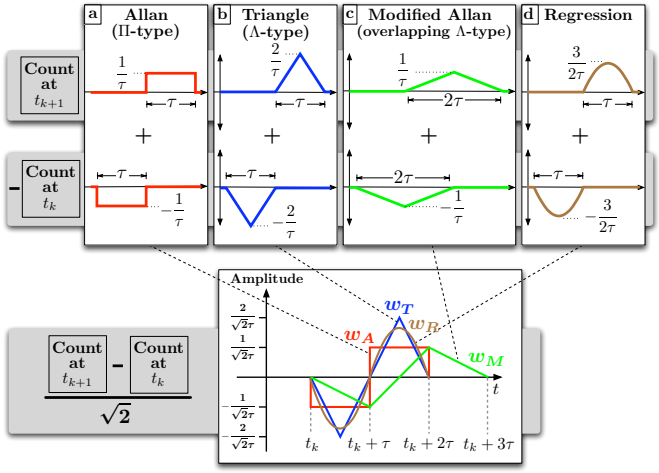

Fig. 1. Each variance weighting function arises from the subtraction of two consecutive frequency measurements: (a) using a traditional/reciprocal counter (w_A), (b) using a high-resolution counter (w_T), (c) using high-resolution counters with overlap to get the modified Allan variance (w_M), (d) using the regression style of counter described in [4] (w_R).

Manuscript received July 15, 2006; revised December 28, 2006.

The authors are with the School of Physics, University of Western Australia, Crawley, WA 6009, Australia (email: sam@physics.uwa.edu.au).

of the gate time. The intention is to reject noise at these frequencies to enhance the resolution of the frequency measurement for a given gate time, but it also has very important consequences when the output of the counter is subsequently used in the calculation of the Allan variance. In this paper we consider the effect of this change in counter operation, as well as how to overcome this limitation. In addition, we consider a proper treatment for the response of the counter to frequency fluctuations at frequencies much greater than the reciprocal of the counter gate time. We note that this paper builds on Rubiola's previous work [1] and corrects some misconceptions presented in that paper.

II. Counter Function

A. High-Resolution Counters

A frequency counter operates by tracking the phase of a signal by detection of its zero-crossings. The rate that zero-crossings (and therefore half-cycles) occur with time is used to estimate the time derivative of the signal phase, $\phi(t)$, and thus deliver the frequency of the signal, $f(t)$, through the relationship

$$f(t) \equiv f_0\, y(t) = \frac{1}{2\pi}\frac{d\phi(t)}{dt}. \tag{8}$$

A traditional frequency counter estimates the derivative by counting the (integer) number of half-cycles in a nominated measurement time τ. A reciprocal counter avoids the error associated with partial cycles by altering the measurement time to coincide with an exact integer number of half-cycles of the input signal [3], [4]. We note that both of these approaches effectively measure the change in absolute phase between the beginning and end of the measurement period, ignoring the time spacing of the zero-crossings in between. This is functionally equivalent to the Π-estimator in frequency space, which is revealed by substituting (8) into (5):

$$f_{\mathrm{meas,k}} = \int_{-\infty}^{\infty} f(t)\, w_{\Pi}(t - t_k)\, dt = \frac{1}{2\pi\tau}[\phi(t_k + \tau) - \phi(t_k)]. \tag{9}$$

Some modern counters improve the resolution of the measurement by using the information contained in the spacing of the zero-crossings within the measurement time [4]. In particular, this paper concerns the Agilent 53131A/53132A in 'time arming mode' or 'digits arming mode' [5] (or the B+K Precision 1856D in internal arming mode), which subtracts the average absolute phase over the first half of the measurement time (i.e. for $t - t_k \in [0, \tau/2]$) from the average absolute phase of the second half (for $t - t_k \in [\tau/2, \tau]$). In frequency terms, this equates to a weighted average of a series of traditional counter integrations, within the single user-selected gate time. The duration of each integration component is half the user-selected gate time (in contrast to the full gate time as suggested in [1]), each delayed from the previous by a small fraction of τ. These individual frequency measurements are then summed equally to give rise to a windowing function that is reasonably well approximated by a triangular weighting function (see Fig. 2). We will refer to this weighted averaging process as a Λ-estimator. We note that while both the Π-estimator (e.g. used by the Agilent 53131A/2A in external arming mode or the Stanford Research Systems SR-620) and the Λ-estimator deliver the average frequency, their response to fluctuations with Fourier frequencies higher than the reciprocal of the gate time is entirely different. Indeed, the advantage of the Λ-estimator is that it is less sensitive to noise at those Fourier frequencies.

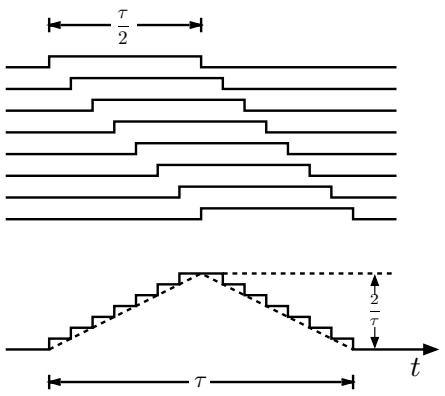

Fig. 2. The operation of modern high-resolution counters involves multiple averaging within a single gate time, τ, well approximated by a Λ-estimator.

We emphasize here that the definition of the Allan variance stipulates the use of the Π-estimator for the estimation of $\bar{y}_k$ in (6) [2]. Therefore, the use of the Λ-estimator in its place does not yield the correct estimation of the Allan variance. If we naively implement the Allan variance algorithm by subtracting adjacent frequency measurements from one of these modern frequency counters, then we obtain what we will refer to as the triangle variance weighting function (see Fig. 1b):

$$w_T(t) = \begin{cases} -\frac{2\sqrt{2}}{\tau^2}t & 0 < t \le \frac{\tau}{2} \\ \frac{2\sqrt{2}}{\tau^2}(t-\tau) & \frac{\tau}{2} < t \le \frac{3\tau}{2} \\ -\frac{2\sqrt{2}}{\tau^2}(t-2\tau) & \frac{3\tau}{2} < t \le 2\tau \\ 0 & \text{elsewhere.} \end{cases} \tag{10}$$

We further note (and show in Fig. 1c) that the triangle variance is not equivalent to the modified Allan variance, as suggested in [1]:

$$w_M(t) = \begin{cases} -\frac{1}{\sqrt{2}\tau^2}t & 0 < t \le \tau \\ \frac{1}{\sqrt{2}\tau^2}(2t-3\tau) & \tau < t \le 2\tau \\ -\frac{1}{\sqrt{2}\tau^2}(t-2\tau) & 2\tau < t \le 3\tau \\ 0 & \text{elsewhere.} \end{cases} \tag{11}$$

For interest, we have also derived (see Fig. 1d) the weighting function for a counter using a linear regression technique [4]:

$$w_R(t) = \begin{cases} \frac{3\sqrt{2}}{\tau^3}t(t-\tau) & 0 < t \le \tau \\ -\frac{3\sqrt{2}}{\tau^3}(t-\tau)(t-2\tau) & \tau < t \le 2\tau \\ 0 & \text{elsewhere.} \end{cases} \tag{12}$$

For reasons of brevity we have not considered this type of counter further, but note that it can be analyzed in the same way as the other counter algorithms.

B. Fourier Analysis

The action of the different temporal windowing functions may be more easily comprehended in the Fourier domain. Using the Power theorem [6], we can equivalently express the

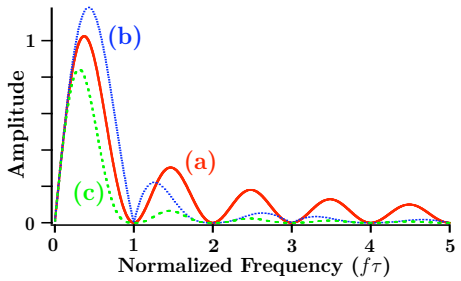

Fig. 3. Fourier transform of the weighting function for (a) the Allan variance measured with a traditional reciprocal counter, ($|W_A(f)|$: solid line), (b) the triangle variance measured with a high-resolution counter with internal averaging, ($|W_T(f)|$: dotted line), and (c) the modified Allan variance, ($|W_M(f)|$: dashed line).

action of the temporal weighting functions as an integration in the Fourier domain:

$$\sigma^2(\tau) = \langle \left[\int_{-\infty}^{\infty} w(t)y(t)dt \right]^2 \rangle = \langle \left[\int_{-\infty}^{\infty} W(f)Y^*(f)df \right]^2 \rangle, \quad (13)$$

where $w(t)$ is an arbitrary temporal windowing function and $W(f) = \int_{-\infty}^{\infty} w(t)e^{-\imath 2\pi ft}dt$ is its corresponding Fourier transform. Here, $Y(f)$ is the Fourier transform of the instantaneous fractional frequency $y(t)$, which is related to the more familiar power spectral density of the fractional frequency, $S_y(f)$ [2]. In fact, we demonstrate in the appendix that (13) can be rewritten in terms of the one-sided $S_y(f)$ as

$$\sigma^2(\tau) = \int_0^{\infty} S_y(f)|W(f)|^2 df. \quad (14)$$

It is thus clear that the squared magnitude of the Fourier transform of the temporal weighting function reveals the spectral sensitivity of the particular variance. We now consider in detail the Fourier transform applied to the three variances of interest in this paper; the Allan variance weighting function $w_A(t)$

$$|W_A(f)| = \frac{\sqrt{2}\,\sin^2(\pi f\tau)}{\pi f\tau}, \quad (15)$$

the triangle variance weighting function $w_T(t)$

$$|W_T(f)| = \frac{\sqrt{32}\,\sin^2(\frac{\pi f\tau}{2})\,|\sin(\pi f\tau)|}{(\pi f\tau)^2}, \quad (16)$$

and finally, for completeness, the modified Allan variance [7] weighting function $w_M(t)$

$$|W_M(f)| = \frac{\sqrt{2}\,\sin^2(\pi f\tau)|\sin(\pi f\tau)|}{(\pi f\tau)^2}. \quad (17)$$

Fig. 3 shows $|W_A(f)|$, $|W_T(f)|$ and $|W_M(f)|$, which reveals the rapid decrease of sensitivity at high frequencies of the triangle variance and the modified Allan variance, as compared with the Allan variance.

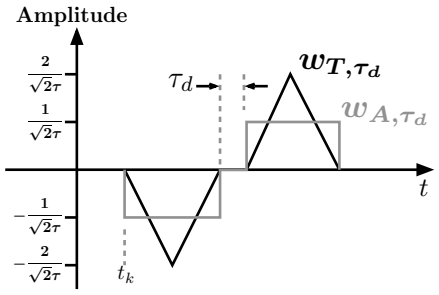

Fig. 4. The triangle variance (black) and Allan variance (gray) weighting functions with dead-time τ_d.

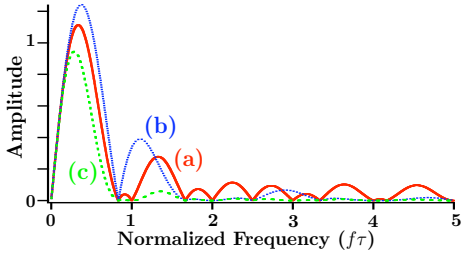

Fig. 5. Fourier transform of the weighting functions for (a) the Allan variance, $|W_{A,\tau_d}(f)|$ (solid line), (b) the triangle variance, $|W_{T,\tau_d}(f)|$ (dotted line), and (c) the modified Allan variance, $|W_{M,\tau_d}(f)|$ (dashed line), all with a dead-time of $0.2\,\tau$.

C. Consideration of Dead-time

In practice when a single frequency counter is used to make a sequence of frequency measurements there is some time between consecutive frequency samples needed for processing. It is possible to avoid significant measurement dead-time through the use of two coordinated counters, but many measurements will involve some amount of dead-time. The effect of this dead-time on the variance calculation is easily modeled by inserting a temporal gap, τ_d, between subtracted counts when deriving the weighting functions (see Fig. 4). The weighting functions in the Fourier domain then become:

$$|W_{A,\tau_d}(f)| = \frac{\sqrt{2}\,|\sin(\pi f\tau)|\,|\sin(\pi f(\tau+\tau_d))|}{\pi f\tau}, \quad (18)$$

$$|W_{T,\tau_d}(f)| = \frac{\sqrt{32}\,\sin^2(\frac{\pi f\tau}{2})\,|\sin(\pi f(\tau+\tau_d))|}{(\pi f\tau)^2}, \quad (19)$$

and

$$|W_{M,\tau_d}(f)| = \frac{\sqrt{2}\,\sin^2(\pi f\tau)\,|\sin(\pi f(\tau+\tau_d))|}{(\pi f\tau)^2}. \quad (20)$$

These functions are plotted in Fig. 5 for a dead-time of $0.2\,\tau$. The most obvious effect of dead-time is the down-shift in frequency of the nodes of zero sensitivity in the spectrum.

D. Implications

From the preceding sections we see that the naive application of the Allan variance algorithm to a sequence of measurements generated by a Λ-type frequency counter does not yield the intended result. The magnitude of the error depends strongly on the nature of the noise in the signal

TABLE I

COMPARISONS OF ALLAN VARIANCE (TRADITIONAL COUNTER), TRIANGLE VARIANCE (HIGH-RESOLUTION COUNTER) AND MODIFIED ALLAN VARIANCE RESULTING FROM CHARACTERISTIC NOISE. HERE, PM STANDS FOR PHASE MODULATED AND FM STANDS FOR FREQUENCY MODULATED. (NOTE: A CUTOFF FREQUENCY, f_H, IS INTRODUCED FOR THE ALLAN VARIANCE OF WHITE PHASE NOISE AND FLICKER PHASE NOISE TO AVOID AN INFINITE RESULT. WE ALSO IGNORE THE SMALL SINUSOIDAL TERM.)

Noise Type	$S_y(f)$	Allan ($\sigma_{\rm A}^2$)	Modified Allan	Triangle
White PM	$h_2 f^2$	$\frac{3\,f_H}{4\,\pi^2}\,h_2\tau^{-2}$ $= \sigma_{\rm A}^2(\tau)$	$\frac{3}{8\,\pi^2}\,h_2\tau^{-3}$ $= \frac{1}{2\,f_H\,\tau}\,\sigma_{\rm A}^2(\tau)$	$\frac{2}{\pi^2}\,h_2\tau^{-3}$ $= \frac{8}{3\,f_H\,\tau}\,\sigma_{\rm A}^2(\tau)$
Flicker PM	$h_1 f$	$\frac{1.038+3\,ln(2\,\pi f_H \tau)}{4\,\pi^2}\,h_1\tau^{-2}$ $= \sigma_{\rm A}^2(\tau)$	$\frac{3\,ln(\frac{256}{27})}{8\,\pi^2}\,h_1\tau^{-2}$ $= \frac{3.37}{3.12+3\ln \pi f_H \tau}\,\sigma_{\rm A}^2(\tau)$	$\frac{6\,ln(\frac{27}{16})}{\pi^2}\,h_1\tau^{-2}$ $= \frac{12.56}{3.12+3\ln \pi f_H \tau}\,\sigma_{\rm A}^2(\tau)$
White FM	h_0	$\frac{1}{2}h_0\tau^{-1}$ $= \sigma_{\rm A}^2(\tau)$	$\frac{1}{4}h_0\tau^{-1}$ $= 0.50\,\sigma_{\rm A}^2(\tau)$	$\frac{2}{3}h_0\tau^{-1}$ $= 1.33\,\sigma_{\rm A}^2(\tau)$
Flicker FM	$h_{-1}f^{-1}$	$2\,ln(2)\,h_{-1}$ $= \sigma_{\rm A}^2(\tau)$	$2\,ln(\frac{3\,3^{11/16}}{4})\,h_{-1}$ $= 0.67\,\sigma_{\rm A}^2(\tau)$	$(24\,ln(2) - \frac{27}{2}\,ln(3))\,h_{-1}$ $= 1.30\,\sigma_{\rm A}^2(\tau)$
Random Walk FM	$h_{-2}f^{-2}$	$\frac{2}{3}\,\pi^2\,h_{-2}\,\tau$ $= \sigma_{\rm A}^2(\tau)$	$\frac{11}{20}\,\pi^2\,h_{-2}\,\tau$ $= 0.82\,\sigma_{\rm A}^2(\tau)$	$\frac{23}{30}\,\pi^2\,h_{-2}\,\tau$ $= 1.15\,\sigma_{\rm A}^2(\tau)$
Frequency Drift ($\dot{y} = D_y$)	-	$\frac{1}{2}D_y^2\tau^2$	$\frac{1}{2}D_y^2\tau^2$	$\frac{1}{2}D_y^2\tau^2$

TABLE II

FIRST ORDER ERROR, δ, IN THE ALLAN, TRIANGLE AND MODIFIED ALLAN VARIANCES CAUSED BY THE INCLUSION OF DEAD-TIME. THE VARIANCE WITH DEAD-TIME, $\sigma_{\tau_d}^2$, IS THE ORIGINAL VARIANCE, σ^2, AUGMENTED BY δ: $\sigma_{\tau_d}^2 \simeq (1+\delta)\,\sigma^2$. (*FOR THE MOST DIVERGENT NOISE, THE ALLAN VARIANCE HAS A COMPLICATED DEPENDENCE ON THE CUTOFF FREQUENCY, f_H; HOWEVER, FOR SIMPLICITY WE GIVE THE MAXIMUM VALUE.)

Noise Type	Allan	Modified Allan	Triangle
White PM	$2\,\frac{\tau_d}{\tau}$*	$-0.33\,\frac{\tau_d}{\tau}$	0
Flicker PM	$2\,\frac{\tau_d}{\tau}$*	$0.67\,\frac{\tau_d}{\tau}$	$0.43\,\frac{\tau_d}{\tau}$
White FM	0	0	0
Flicker FM	$\frac{\tau_d}{\tau}$	$1.33\,\frac{\tau_d}{\tau}$	$0.62\,\frac{\tau_d}{\tau}$
Random Walk FM	$1.50\,\frac{\tau_d}{\tau}$	$1.67\,\frac{\tau_d}{\tau}$	$1.30\,\frac{\tau_d}{\tau}$
Frequency Drift	$2\,\frac{\tau_d}{\tau}$	$2\,\frac{\tau_d}{\tau}$	$2\,\frac{\tau_d}{\tau}$

being measured. Table I compares the Allan variance, triangle variance and modified Allan variance resulting from several common characteristic noise types. For an arbitrary signal, however, Greenhall [8] has shown that it is not possible to derive the power spectral density from the Allan variance; this is equally true for the triangle variance calculation. Therefore, it is not possible to manipulate data taken from a Λ-type counter to yield the Allan variance that would have been measured by a Π-type frequency counter, except for the special case where the shape of $S_y(f)$ is already known. The corrections for listed $S_y(f)$ are also presented in Table I. Where the measurement involves a significant dead-time, a further correction may be necessary, and we include Table II for convenience. We note that the Allan variance values are consistent with the bias functions tabulated by Barnes and Allan [9].

The easiest way to be sure of obtaining the true Allan variance is to ensure that a Π-type counter is used. Unlike in [1], where the averaging of a series of Λ-type counter

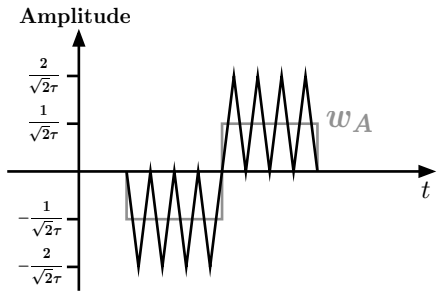

Fig. 6. Averaging Λ-counts before calculating the Allan variance produces a triangularly modulated Π-count. Here we show four averages ($N = 4$).

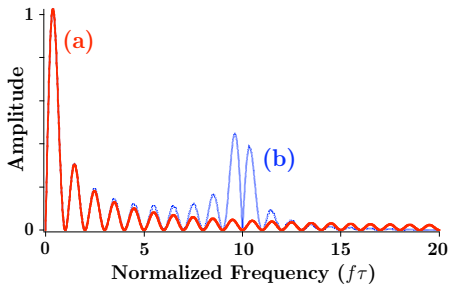

Fig. 7. The variance calculated by averaging consecutive Λ-counts (b) approximates the Allan variance (a) at low frequencies but increases sensitivity around $\frac{N}{\tau}$ (here $N = 10$).

measurements is assumed to approximate the Π-estimator, we find that a series of Λ-estimators produces a Π-estimator modulated by a triangle wave (see Fig. 6). For a number of samples, N, the variance produced by this method approaches the Allan variance at frequencies comparable to the reciprocal of the gate time, but introduces sensitivity near N times the reciprocal of the gate time (see Fig. 7).

A second alternative is to use a spectrum analyzer to generate the frequency noise, $S_y(f)$, and then calculate the

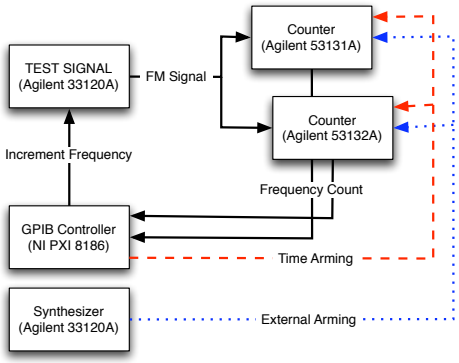

Fig. 8. Experimental setup to measure the windowing function with negligible dead-time. The counters take measurements alternately, armed either by computer instruction (time arming) or a rising edge (external arming).

Allan variance by combining (14) and (18):

$$\sigma_{\rm A}^2(\tau)=\int_0^\infty S_y(f)\frac{2\,\sin(\pi f\tau)^2\,\sin(\pi f(\tau+\tau_d))^2}{(\pi f\tau)^2}df. \quad (21)$$

To derive the triangle variance, we combine (14) with (19) to get:

$$\sigma_{\rm T}^2(\tau)=\int_0^\infty S_y(f)\frac{32\,\sin(\frac{\pi f\tau}{2})^4\,\sin(\pi f(\tau+\tau_d))^2}{(\pi f\tau)^4}df. \quad (22)$$

These integral forms are easily calculated because the $S_y(f)$ produced by the spectrum analyzer is already discretized into frequency bins. Accordingly, the variance in (14) becomes a summation over N bins:

$$\sigma^2(\tau)=\sum_{n=1}^{N}[\overline{S_{y,n}}\int_{f_n}^{f_{n+1}}|W(f)|^2df], \quad (23)$$

where $\overline{S_{y,n}}$ is the power spectral density calculated by the spectrum analyzer for the nth bin. We also note that the ratio of the Allan and triangle Fourier transforms takes the particularly simple form, which is independent of dead-time:

$$\frac{W_{T,\tau_d}(f)}{W_{A,\tau_d}(f)}=\frac{\tan(\frac{\pi f\tau}{2})}{(\frac{\pi f\tau}{2})}=\mathrm{tanc}(\frac{\pi f\tau}{2}). \quad (24)$$

III. Results

A. Model Testing

We tested our model with two types of commonly used counters that we had available (Agilent 53131A/53132A). They were supplied by a synthesizer (shown schematically in Fig. 8) with a carrier signal that was modulated at a single well-defined frequency:

$$V_{mod}=V_0\sin(2\pi f_c t+M\sin(2\pi f_m t)). \quad (25)$$

The resulting set of contiguous frequency measurements (100s of data at each modulation frequency) were then applied to (2) to provide a measure of the response of the variance to that modulation frequency, f_m. Repetition of the procedure for a range of f_m allows us to generate a plot of the spectral sensitivity of the windowing function, $w(t)$, as a function of frequency. In fact, this measurement gives the Fourier sine transform, so we scale the result by $\frac{\sqrt{2}}{M}$ to derive the general Fourier transform, normalized by the modulation index, M. The result obtained when using external gating (i.e. the Allan variance spectral window $|W_A(f)|$) is compared with the theoretical function from (15) in Fig. 9. Similarly, Fig. 10 shows that the measured spectral window of the counter while in time arming mode (i.e. the triangle variance spectral window $|W_T(f)|$) agrees with (16). Finally Fig. 11 shows that the expressions given in (18) and (19) including dead-time also hold, and that the calculated sensitivity functions are in excellent agreement with the measurements. The results verify that when modern enhanced resolution counters are used to gather frequency data, they generate the triangle variance rather than the intended Allan variance. Experimental verification of the cusp on Fig. 10 at a frequency of $\frac{1}{\tau}$, which is not present on either the Allan or modified Allan variance spectral window, is a prominent visual indication that the triangle variance is different from both.

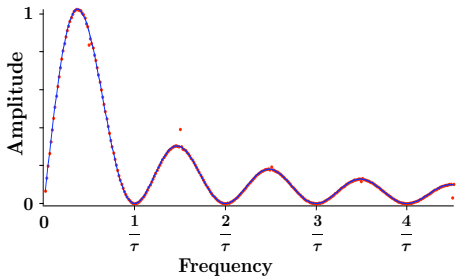

Fig. 9. The Allan variance spectral window, $|W_A(f)|$, as measured with the counters in external arming mode (dots) compared with that calculated with (15) (line).

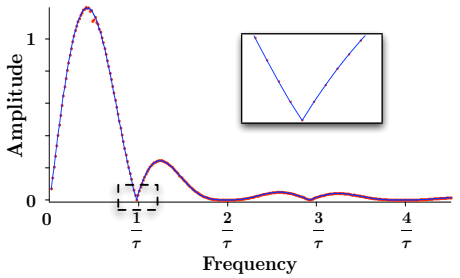

Fig. 10. The triangle variance spectral window, $|W_T(f)|$, as measured with the counters in time arming mode (dots) compare with that calculated with (16) (line). Inset: Magnification of the cusp at $f=\frac{1}{\tau}$.

We note that measurements deviate from the model prediction where the modulation frequencies are close to a multiple of $\frac{1}{2\tau}$. At such frequencies, the modulation and the measurement rate maintain relative phase, which compromises the estimation of the ensemble average in (2).

B. Counter Response to Real Oscillators

We have tested the effect of the triangle variance algorithm on three different types of oscillator. We first synthesized signals that were dominated by white phase noise and white frequency noise to approximate the characteristic noise of an active hydrogen maser and a primary frequency standard

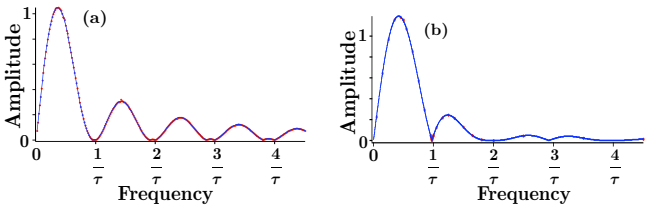

Fig. 11. Experimental verification of (a) the Allan variance spectral window, $|W_{T,\tau_d}(f)|$, with dead-time of 5% of τ and (b) the triangle variance spectral window, $|W_{T,\tau_d}(f)|$, with dead-time of 2.6% of τ.

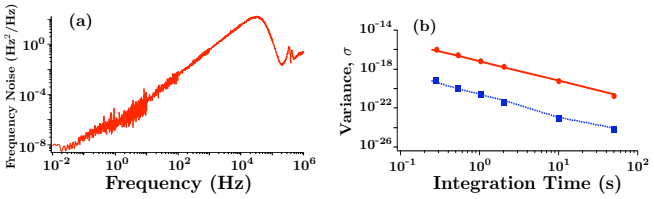

Fig. 12. The left panel (a) shows synthesized noise, dominated by white phase noise, which was used in (23) to generate the Allan variance (solid lines) and the triangle variance (dotted lines) in the right panel (b). Also shown in (b) are measurements made on the same source with an externally armed counter (dots) and a time-armed counter (squares).

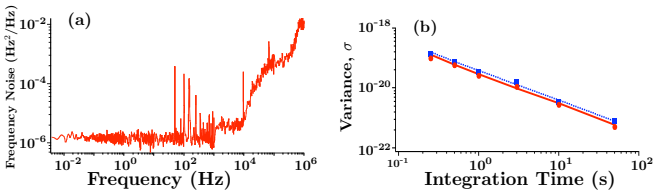

Fig. 13. The left panel (a) shows synthesized white frequency noise, which was used in (23) to generate the Allan variance (solid lines) and the triangle variance (dotted lines) in the right panel (b). Also shown in (b) are measurements made on the same source with an externally armed counter (dots) and a time-armed counter (squares).

respectively. In all cases we applied a low pass filter of 50 kHz to provide a well-defined cut-off frequency, f_H, as is required for the interpretation of Tables I and II. These signals were used, in turn, to simultaneously drive two counters - one in external arming mode and another in time arming mode. Fig. 12 shows the white phase noise dominated signal in the Fourier domain, as measured by an FFT spectrum analyzer, and the corresponding Allan and triangle variances measured by the respective counters (shown as error bars). The lines on Fig. 12 are calculated directly from the frequency noise power spectral density as per (23), using an upper integration limit of 50 kHz. Fig. 13 shows the equivalent results for the white frequency noise dominated signal. Together these figures demonstrate that for conventional oscillators that are dominated by white frequency noise, the difference between the two types of measurement is small, but for an oscillator with more rapidly diverging frequency noise, the difference can be very significant.

We also tested the effect of this measurement on an optical frequency synthesizer we have built at UWA [10]. Indeed, we first noticed the importance of the counter gating mode when making Allan variance measurements of the synthesizer's repetition rate signal. In this case, due to the limited bandwidth of the frequency control transducer (∼30 kHz), the signal exhibits a strong increase in the frequency fluctuations with Fourier frequency (see Fig. 14), which is similar to flicker phase noise. The Allan variance shows a marked difference in its behavior in comparison with the triangle variance, which is due to a stronger rejection of the noise at high frequencies. We predict that some frequency standards based on macroscopic resonators will produce a similar discrepancy. The frequency fluctuations of such oscillators can be dominated by environmental sensitivity at acoustic frequencies [11], [12], which will contribute more strongly to the triangle variance than the Allan variance.

IV. Discussion

Despite the existence of several alternative techniques for characterizing frequency stability, such as the modified Allan variance and the Total variance [13], the Allan variance has remained the most common measure of frequency stability. This is most likely due to its simplicity and ease of implementation.

We suggest that, instead of viewing the effect of counter averaging as a hindrance, the resulting triangle variance be put forward as a possible new measure for characterizing frequency stability. It shares two of the advantages of the modified Allan variance; in particular, being convergent in response to very divergent noise (up to $S_y(f) \propto f^3$), and secondly the ability to distinguish between white phase noise and flicker phase noise because of their different dependence on the integration time ($\frac{1}{\tau^3}$ and $\frac{1}{\tau^2}$ respectively). Furthermore, it offers the same ease of implementation as the Allan variance, since it is directly delivered by some counters in their default high resolution mode. This is in contrast to the modified Allan variance which requires additional computational effort. Perhaps most importantly though, when the counters are operating in their most sensitive mode, there is no choice but to accept frequency measurements that have been filtered in this way. Thus we suggest that the triangle variance provides a useful measure of frequency stability that is closely related to the traditional Allan variance, whilst also exploiting the highest sensitivity mode of modern counters.

V. Conclusion

We have demonstrated that the usual method of estimating the Allan variance of an oscillator is incorrect with many

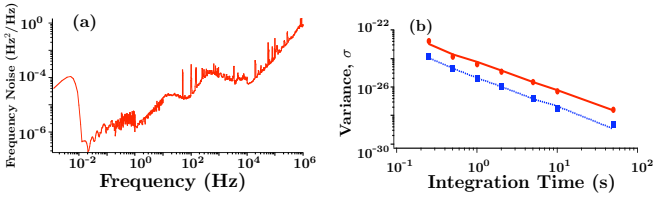

Fig. 14. The left panel (a) shows the frequency fluctuations of the repetition rate of the optical frequency synthesizer, which was used in (23) to generate the Allan variance (solid lines) and the triangle variance (dotted lines) in the right panel (b). Also shown in (b) are measurements made on the same source with an externally armed counter (dots) and a time-armed counter (squares).

modern high-resolution frequency counters. The internal averaging processes used by such counters delivers reduced sensitivity to frequency fluctuations at high Fourier frequencies (as compared with the reciprocal of the gate time), resulting in an incorrect estimate of the Allan variance as it has been defined. The error is moderate (33 % of the true value) for oscillators dominated by white frequency noise, but it can be of great significance (larger than an order of magnitude) for some sources, especially those dominated by white phase noise. Either one must take care to avoid the error in the measurement process or, alternatively, consider the difference meaningful by incorporating this triangle variance into a new definition of frequency stability.

APPENDIX

In this appendix, we demonstrate the analysis to go between (13) and (14). We follow the form of the proof presented in Appendix I in Barnes [2], except we include any variations in the definition of the variance, such as the consideration of multiple samples or dead-time, within an arbitrary weighting function $w(t)$. Let us start with the first half of (13) and re-write the squared term as the product of two independent integrals and manipulate:

$$\begin{aligned}\sigma^2(\tau) &= \langle [\int_{-\infty}^{\infty} w(t)y(t)dt]^2 \rangle \\ &= \langle \int_{-\infty}^{\infty} w(t')y(t')dt' \int_{-\infty}^{\infty} w(t'')y(t'')dt'' \rangle \\ &= \langle \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} w(t')w(t'')y(t')y(t'')dt'dt'' \rangle \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} w(t')w(t'')\langle y(t')y(t'')\rangle dt'dt'' \end{aligned} \quad (26)$$

We then use the definition of the autocorrelation function and the Wiener-Khinchin theorem for the one-sided $S_y(f)$ to get the relationship:

$$\begin{aligned}\langle y(t')y(t'')\rangle &= R_y(t'-t'') \\ &= \int_0^{\infty} S_y(f)\exp^{\imath 2\pi f(t'-t'')}\, df. \end{aligned} \quad (27)$$

Note that this requires that the integral $\int_0^{\infty} S_y(f)df$ exists which can be forced by introducing a high frequency cutoff f_H (also required for the calculation of Allan variance for white phase noise and flicker phase noise as per Table I). Substituting (27) into our variance expression and re-arranging yields:

$$\begin{aligned}\sigma^2(\tau) &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} w(t')w(t'')\int_0^{\infty} S_y(f)e^{\imath 2\pi f(t'-t'')}df\,dt'dt'' \\ &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_0^{\infty} w(t')w(t'')S_y(f)e^{\imath 2\pi f(t'-t'')}dt'dt''df \\ &= \int_0^{\infty} S_y(f)\int_{-\infty}^{\infty} w(t')e^{\imath 2\pi f t'}dt' \int_{-\infty}^{\infty} w(t'')e^{-\imath 2\pi f t''}dt''df. \end{aligned}$$

The final line above reveals the familiar form of the Fourier transform and we can attain the result by knowing that $W(f)$ is Hermitian because $w(t)$ is a real function:

$$\begin{aligned}\sigma^2(\tau) &= \int_0^{\infty} S_y(f)W(-f)W(f)df \\ &= \int_0^{\infty} S_y(f)W^*(f)W(f)df \\ &= \int_0^{\infty} S_y(f)|W(f)|^2 df. \end{aligned} \quad (28)$$

ACKNOWLEDGMENT

We would like to thank Frank van Kann, Paul Abbott, James Anstie and Milan Marić for their useful input. This work is funded by the Australian Research Council.

REFERENCES

[1] E. Rubiola, "On the measurement of frequency and of its sample variance with high-resolution counters," *Review of Scientific Instruments*, vol. 76, p. 054703, 2005.

[2] J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen Jr., W. L. Smith, R. L. Sydnor, R. F. C. Vessot, and G. M. R. Winkler, "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, vol. IM-20, pp. 105–120, 1971.

[3] J. Kalisz, "Review of methods for time interval measurements with picosecond resolution," *Metrologia*, vol. 41, pp. 17–32, 2004.

[4] S. Johansson, "New frequency counting principle improves resolution," in *Proceedings of the 2005 IEEE International Frequency Control Symposium and Exposition*, 2005, p. 8.

[5] *Agilent 53131/132A Universal Counter Operating Guide*, Agilent Technologies, Inc., 815 14th Street S.W., Loveland, Colorado 80537 U.S.A., 2003.

[6] R. N. Bracewell, *The Fourier transform and its applications*, 3rd ed. McGraw-Hill, 2000.

[7] D. W. Allan and J. A. Barnes, "A modified "allan variance" with increased oscillator characterization ability," in *Proc. 35th Ann. Freq. Control Symposium*, 1981, pp. 470–475.

[8] C. A. Greenhall, "Spectral ambiguity of allan variance," *IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT*, vol. 47, no. 3, pp. 623–627, Jun 1998.

[9] D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, "Characterization of clocks and oscillators," *NIST Technical Note 1337*, 1990.

[10] J. J. McFerran, S. T. Dawkins, P. L. Stanwix, M. E. Tobar, and A. N. Luiten, "Optical frequency synthesis from a cryogenic microwave sapphire oscillator," *OPTICS EXPRESS*, vol. 14, no. 10, pp. 4316–4327, May 2006.

[11] F. Bondu, P. Fritschel, C. N. Man, and A. Brillet, "Ultrahigh-spectral-purity laser for the virgo experiment," *Optics Letters*, vol. 21, no. 8, pp. 582–584, 1996.

[12] M. Notcutt, L.-S. Ma, A. D. Ludlow, S. M. Foreman, J. Ye, and J. L. Hall, "Contribution of thermal noise to frequency stability of rigid optical cavity via Hertz-linewidth lasers," *Phys. Rev. A*, vol. 73, no. 3, p. 031804, March 2006.

[13] C. A. Greenhall, D. A. Howe, and D. B. Percival, "Total variance, an estimator of long-term frequency stability," *IEEE TRANSACTIONS ON ULTRASONICS FERROELECTRICS AND FREQUENCY CONTROL*, vol. 46, no. 5, pp. 1183–1191, Sep 1999.

Samuel T. Dawkins was born in Perth, Western Australia in 1976. He received the B.Sc. (Physics) and B.E. (Mech. Eng.) degrees from the University of Western Australia in 1999. After two years working as a mechanical engineer with a thermoelectric refrigeration company, he returned to the University of Western Australia as a Ph.D. candidate. As part of this, he has built an optical frequency reference based on all-sapphire Fabry-Perot resonators at room temperature with the Frequency Standards and Metrology Group in the School of Physics.

John J. McFerran received a B.Sc.Eng. degree from the University of Western Australia in 1993 with majors in mechanical engineering and physics. After a three year absence from academia he returned to the School of Physics, U.W.A. and was awarded a Ph.D. in 2002. Since 2002 he has held a post-doctoral research associate position in the Frequency Standards and Metrology Group at U.W.A. In 2004 and 2005 he was a Guest Researcher at the National Institute of Standards and Technology, Boulder, CO (in the Time and Frequency Division and Optoelectronics Division, respectively). His research interests include optical frequency synthesis with both fibre-laser and solid-state-laser based frequency combs, as well as optical atomic frequency standards.

André N. Luiten is the Head of the Optical Division of the Frequency Standards and Metrology Group in the School of Physics at the University of Western Australia. He received his Ph.D. entitled ”Sapphire Secondary Frequency Standards” from the University of Western Australia in 1996. A/Prof. Luiten’s research interests are characterised by a desire to develop highly precise instrumentation and then using these devices to make measurements of high importance. For example, the microwave sapphire oscillator developed out of his Ph.D. demonstrates a frequency stability in the 10^{-16} range and has been used in precision tests of relativity. More recently André has been interested in developing optical frequency metrology technology including phase-coherent optical synthesizers and frequency stable optical oscillators and clocks. André has attracted $4.5M of research funding and published over 50 papers in journals during his academic career.