The Allan Variance – Challenges and Opportunities

Samuel R Stein
Symmetricom, Inc.
Boulder, Colorado, USA

***Abstract*—The Allan variance has historically been measured using heterodyne measurement systems, which have low noise and preserve the carrier phase information needed for long-term stability. The single-sideband phase noise has traditionally been measured using phase detectors that suppress the carrier in order to achieve even lower noise. The recent development of the direct-digital phase noise measurement technique makes it possible to estimate both statistics simultaneously from the same time series of the phase [1]. Our comparison of the three techniques has revealed several challenges to the accurate estimation of the Allan variance including undesired aliasing, biased estimators, and spurious signal generation. Successfully overcoming these challenges has resulted in faster, easier, more reliable, and more accurate measurement methods**

I. Introduction

The Allan variance (AVAR) is almost always estimated using some form of heterodyne measurement system [2]. The simple beat frequency measurement system is used to measure the phase difference between two oscillators that have a convenient frequency offset in the range of a few Hz. When the first input phase increases by one cycle compared to the second input, the beat frequency signal shows a one cycle increase in phase. Thus the carrier phase-differene information is preserved

The dual-mixer time difference measurement system is commonly used to measure the phase difference between two oscillators that have nearly equal frequency, which is generally true of standard frequency oscillators [3]. This approach uses a transfer oscillator to measure the phase differences between both oscillators under test and the transfer oscillator. The phase difference between the inputs is estimated by subtracting the two measurements, which has the benefit of cancelling most of the transfer oscillator noise. The benefit of the dual-mixer measurement system is the ability to use the heterodyne technique when the oscillators under test don't have a convenient frequency difference.

The utility of the heterodyne technique has been extended through the use of frequency synthesizers. They can be used as the reference in simple heterodyne systems, to produce the transfer oscillator offset in dual-mixer time difference measurement systems or to extend the dual mixer technique further to measure unequal frequency oscillators. Careful design of such measurement systems results in some or most of the synthesizer noise cancelling in the measurements.

Heterodyne measurement systems are not used to estimate the spectral density of phase noise for several reasons. The phase information is obtained from the times of the zero crossings of the IF signal. However, the wide voltage swing of this signal exceeds the dynamic range that would be needed to measure low broadband phase noise far from the carrier. In addition, the presence of the carrier signal in the mixer increases the flicker phase noise of that device. The study in this paper has also revealed that heterodyne measurement systems experience frequency domain aliasing that may corrupt the AVAR estimates and makes this approach unusable for estimating the spectral density of phase noise.

The spectral density of phase noise has traditionally been measured using double-balanced mixers as phase detectors [4]. The usual method is to phase lock the device under test to the reference. The control loop is used to maintain the two signals near quadrature so the output of the mixer is near zero and approximately proportional to the phase difference between the two oscillators. This approach eliminates the dynamic range limitation of the heterodyne technique. In addition, the use of the mixer as a phase detector reduces its flicker phase noise contribution. These two changes make the phase-lock technique suitable to estimate the spectral density of phase noise. However, the phase-lock loop used to maintain quadrature suppresses the carrier phase information inside the loop bandwidth and the time series of phase can't be used to estimate the AVAR for sample times longer than the loop time constants.

Direct-digital measurement systems resolve the incompatibility between the heterodyne and phase-lock measurement systems. The direct digital measurement approach digitizes the input signals and performs frequency conversion and phase detection numerically. The fundamental advance is the ability to measure the phase throughout the cycle using a phase detector based on the arctangent function, which requires no calibration and allows anti-alias filtering. The noise of analog-to-digital converters requires the use of cross-correlation as shown in Fig. 1. The use of the cross spectrum and the cross AVAR creates new challenges for achieving unbiased estimates of the AVAR. Many of the plots

This work was supported in part by NIST contract SB1341-04-C-0022

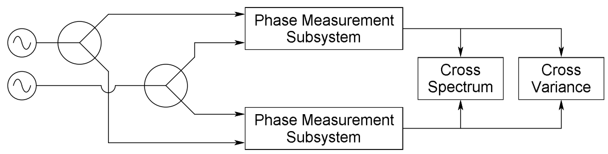

Figure 1: Cross correlation of phase-measurement subsystems

shown in the following report the Allan deviation – the square root of the AVAR – which is abbreviated ADEV.

II. OVERVIEW

AVAR is a measure of how much the frequency of and oscillator changes from one sample interval, τ, to the next interval with no intervening dead time. Using the IEEE recommended definition of frequency based on the end point phases, AVAR can be written in terms of the phase-time series, ϕ_k [5].

$$\sigma_y^2(\tau) = \left\langle \frac{1}{2}\left(\frac{\phi_{k+2} - 2\phi_{k+1} + \phi_k}{2\pi\nu_0\tau} \right)^2 \right\rangle \tag{1}$$

III. CHALLENGES ESTIMATING AVAR

A. Using Heterodyne Phase Measurements

The heterodyne method of estimating the phase difference between oscillators is a form of digital measurement in that no analog transducer is involved. The down-conversion process preserves a cycle of phase [6]. That is, one cycle at the input frequency is equal to one cycle at the intermediate frequency (IF). Heterodyning makes it easier to measure a fractional cycle because the period has increased compared to the period of the counter time base. The measurements are the times of the "zero crossings" at which time the signal is half way between the signal peaks and valleys. Only one sign of zero crossing can be used because signal sources have substantial duty cycle noise – the positive and negative zero crossings move relative to one another.

As a result, heterodyne phase-difference measurements comprise at most one sample per full IF cycle. The IF is low-pass filtered with a bandwidth that must be no smaller than the IF since bandpass filtering excessively perturbs its phase. This procedure violates the sampling theorem, which states that a signal is properly represented by digital samples only if it contains no power at Fourier frequencies above one-half the sampling frequency. All the signal components above one-half the sampling frequency appear as aliases in the frequency band between 0 and half the sampling frequency. Fig. 2 ilustrates frequency aliasing. The problem is exacerbated by the fact that the bandwidth is usually at least as great as the 3rd harmonic of the IF to achieve the best zero-crossing detection.

This inevitable aliasing affects the AVAR in some cases but not in others. If the dominant noise process is white phase noise such as may occur in a quartz crystal oscillator at short

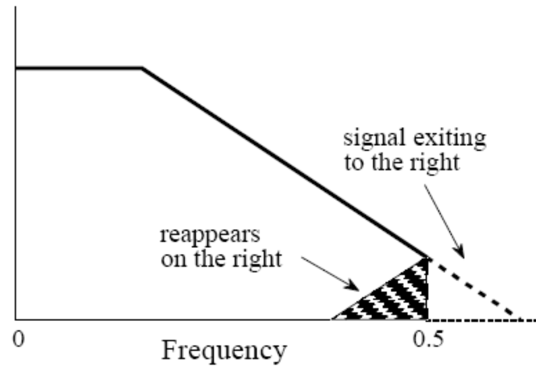

Figure 2: Frequency domain aliasing

averaging times, the AVAR is insensitive to aliasing because it depends only on the noise density, σ_0^2, multiplied by the measurement bandwidth, f_h, as shown in (2).

$$\sigma_y^2(\tau) = \sigma_0^2 f_h \tag{2}$$

Without aliasing the noise density is S_ϕ and the bandwidth is B. After aliasing the noise density is NS_ϕ and the bandwidth is $B/10$. However, when the noise is colored or spurious signals are significant, aliasing can have a dramatic affect on the AVAR estimate. Consider the example of a heterodyne measurement with 100 Hz sample rate and a dominant 60 Hz spurious signal from the power mains. The AVAR has a peak at 8.83 ms and every odd multiple and a zero at every even multiple. However, the signal appears as an alias at 40 Hz and the estimated AVAR has peak at 12.5 ms and every odd multiple and a zero at every even multiple. Fig. 3 shows the variation in ADEV estimates between a heterodyne phase-difference measurement with aliased spurs and a simultaneous direct-digital measurement with spurs at the correct frequency when the ADEV has noticable power-line spurious signals.

Although a little outside of the focus of this paper, it is notable that aliasing makes heterodyne phase-difference measurements unreliable for spectrum estimation. In this case, the aliased white noise causes inaccurate spectral density estimates. Fig. 4 shows the result of estimating the single-sideband phase noise with 100 heterodyne phase-

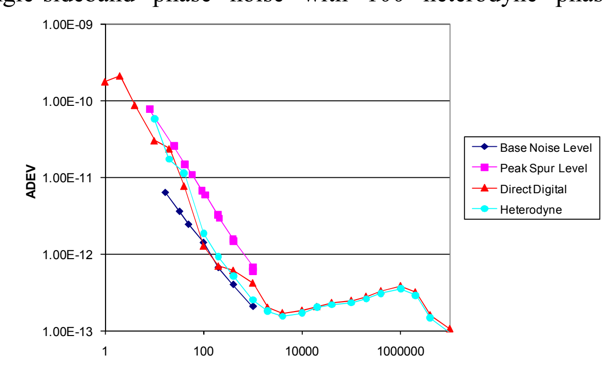

Figure 3: ADEV variations due to aliasing of power line spurs

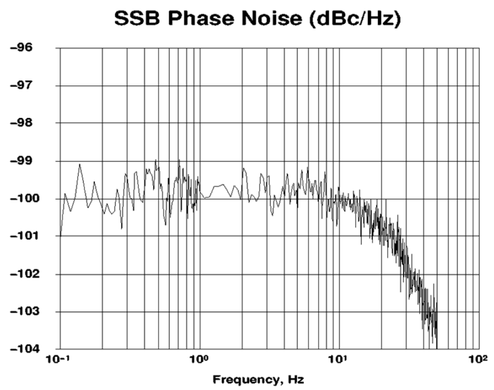

Figure 4: Phase noise of -110 dBc/Hz estimated from 100 Hz heterodyne measurements appears at -100 dBc/Hz

difference samples per second. A calibrated source with −110 dBc/Hz single sideband phase noise (SSB) was estimated to have −100 dBc/Hz SSB. The 10 dB discrepancy is accounted for within 0.5 dB by the 450 Hz calculated noise bandwidth of the zero-crossing detector. When used with care by a knowledgeable person, heterodyne measurements can be employed to estimate the spectrum of oscillators close to the carrier where the highly colored (SSB $\propto$ 1/f^3) flicker frequency noise negates the effect of aliasing.

B. Using Phase Noise Measurements

AVAR can be estimated by integrating the spectral density of phase as shown in (3). Thi s approach is very useful because it enables frequency domain filtering to separate the effects of various components of the oscillator signal. For example, power line spurs can be removed and the AVAR corresponding to the noise and the spurs computed separately.

$$\sigma_y^2(\tau) = 2\int_0^{f_h} S_\phi(f)\,\frac{\sin^4(\pi f \tau)}{(\pi \nu_0 \tau)^2}\,df \tag{3}$$

The first peak in the kernel occurs at a frequency of $1/\pi\tau$. Highly accurate phase noise measurements are required at Fourier frequencies well below the first peak frequency. Thus, when the time series of phase is produced by an analog phase detector with a PLL to maintain phase quadrature, extremely long loop time constants are required.

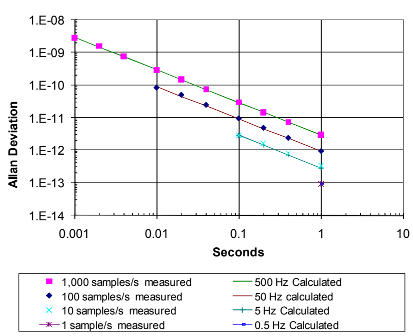

Figure 5: Calculated and measured ADEV for bandwidths from 500Hz to 0.5 Hz (top to bottom)

Equations (2) and (3) show that the AVAR of a white phase-noise process is linearly dependent on measurement bandwidth. Fig. 5 plots the calculated and measured ADEV as the measurement bandwidth is varied from 500 Hz to 0.5 Hz. This plot demonstrates the importance of reporting the measurement bandwidth.

C. Using Cross Correlation

Cross correlation is often used to reduce the noise contribution of the instrument to the measurement [7]. Independent measurements of the time series of the phase difference are made and (1) is rewritten as (4) using the $\otimes$ symbol to designate the cross version of the statistic.

$$\sigma_y^{2\otimes}(\tau) = \left\langle \frac{1}{2}\left(\frac{\phi_{k+2}^A - 2\phi_{k+1}^A + \phi_k^A}{2\pi\nu_0\tau}\right)\left(\frac{\phi_{k+2}^B - 2\phi_{k+1}^B + \phi_k^B}{2\pi\nu_0\tau}\right)\right\rangle \tag{4}$$

If heterodyne or direct-digital phase difference measurements are used, then the AVAR can be calculated even when the long-term oscillator frequency variations are dominated by temperature or aging effects. If a single reference oscillator is used the cross AVAR is given by

$$\sigma_y^{2\otimes}(\tau) = \tfrac{1}{2}\left(D_{IN} - D_{REF}\right)^2 \tau^2 \tag{5}$$

Where D_{IN}-D_{REF} is the relative aging of the input compared to the reference. This is the same result as the AVAR.

The cross AVAR also performs a separation of variances if the reference oscillators for the *A* and *B* measurements are independent. Under these conditions the cross AVAR may be a biased estimator of the AVAR and doesn't represent it well under some circumstances. For example, in the region where the differential aging between the input and reference oscillators is the dominant source of frequency variation

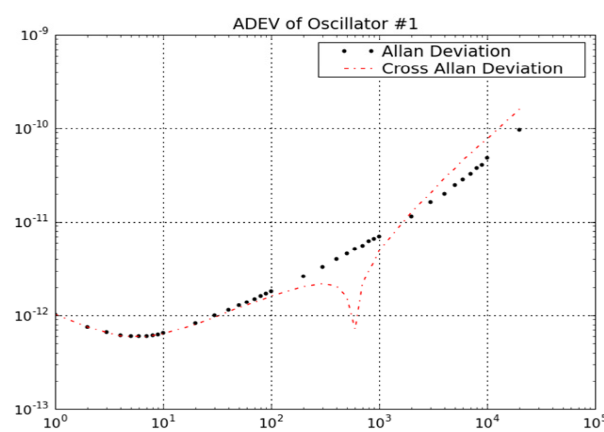

Figure 6: The absolute cross ADEV is a biased estimate of ADEV

the cross ADEV is given by

$$\sigma_y^{2\otimes}(\tau) = \tfrac{1}{2}(D_{IN} - D_{REF1})(D_{IN} - D_{REF2})\tau^2, \qquad (6)$$

where D_{IN}-D_{REF1} and D_{IN}-D_{REF2} are the relative aging of the input oscillator relative to the two independent references. Under these circumstances, the cross AVAR may be negative and the ADEV doesn't exist. The most obvious fix is to use the absolute value of the cross AVAR as an estimator of the AVAR. However, this approach results in a notch in the AVAR as shown in Fig. 6. The notch can be smoothed by integrating the cross spectrum to obtain the AVAR.

$$\sigma_y^{2\otimes}(\tau) = 2\int_0^{f_h} \mathrm{Re}\left[S_\phi^{\otimes}(f)\right]\frac{\sin^4(\pi f \tau)}{(\pi \nu_0 \tau)^2}\,df \qquad (7)$$

An example of fixing the AVAR notch using this technique is shown in Fig. 7.

D. *Using Direct-Digital Phase-Difference Measurements*

The input frequency ν_{IN} is sampled with frequency ν_s. Because this process is non-linear, spurs are created at all frequencies $M\,\nu_{IN} - N\,\nu_s$. When the frequency of the spur lies within the AVAR bandwidth and M and N are less than approximately 300, the AVAR estimate is affected. Most of these spurs can be identified and removed. Spurious modulation on the input must always produce a line in the spectum near the real axis. However, the spectra of internally generated sampling spurs are found empirically to have uniformly distributed phase angles. Those that don't lie along the real axis are known to not originate as modulation of the input signal and may be removed from the spectrum before integration to obtain the AVAR.

Thus, when cross correlation is used to reduce the instrumentation contribution to the measurement, integration of the cross spectrum using (6) is a superior method of generating an unbiased estimate of AVAR than use of (4) directly from very small sample times up to a few thousand seconds. For times longer than two-thousand seconds the best AVAR estimate is still given by the absolute value of the expression in (6).

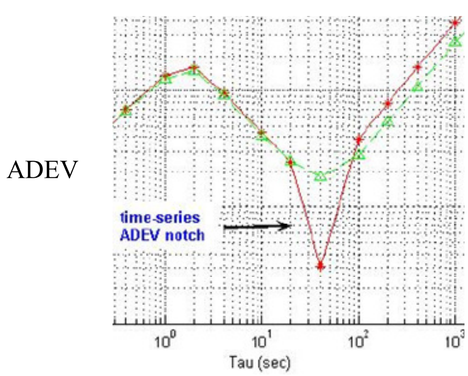

Figure 7: The integrated spectrum (green) improves the ADEV estimate

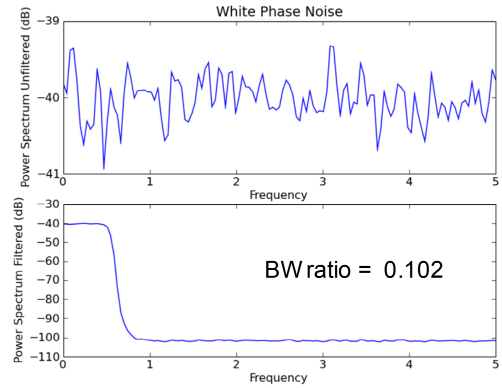

Figure 8: Anti alias filter design with ½ Hz noise bandwidth

IV. Opportunities for AVAR Estimation Using Direct-Digital Measurement Systems

A. *Flexible Filtering*

Direct-digital measurement systems satisfy the requirements of the sampling theorem. Anti-alias filtering and further sub-sampling allow total flexibility in the choice of measurement bandwidth. Exploring this flexibility has led to the conclusion that there is an optimum measurement bandwidth, f_h, that should be used when either analog or direct-digital phase measurements are used to compute AVAR. The maximum measurement bandwidth is set by the sampling theorem and is $(2\tau_{min})^{-1}$, where τ_{min} is sample period and the minimum sample interval for calculating AVAR. However, if a smaller measurement bandwidth were used, the AVAR estimates for $\tau < (2f_h)^{-1}$ would be more representative of the filter than the oscillator under test. Thus, the optimum measurement bandwidth is $(2\tau_{min})^{-1}$.

Figs. 8 and 9 show the design of a filter and its effect on ADEV respectively. The unfiltered data have τ_{min} = 100 ms and are filtered with a bandwidth of ½ Hz. Fig. 9 shows that the ADEV for τ < 1 s is highly filter dependent and should not be displayed.

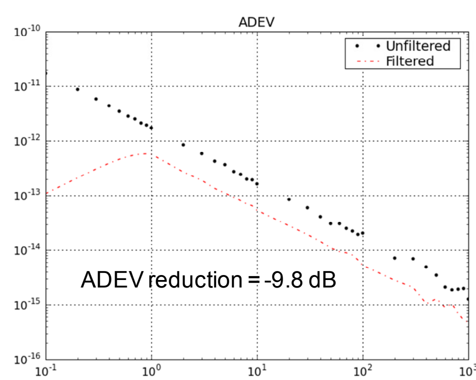

Figure 9: ADEV has an optimum bandwidth for any sample rate

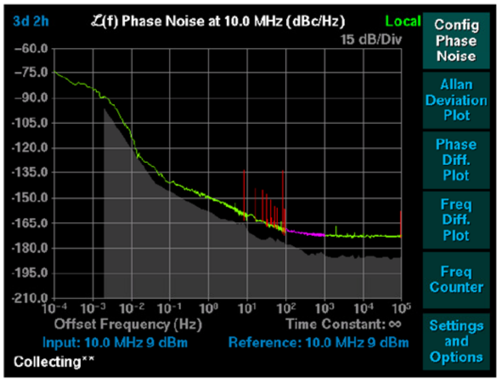

Figure 10: The cross spectrum contains noise floor information

B. *Cross Statistics*

The variance of the AVAR estimate calculated according to (1) is approximately inversely proportional to the square-root of the number of measurements. However, the AVAR estimate includes the mean-square instrument noise. A separate measurement of the instrumentation noise floor is usually performed to determine the degree to which it has biased the AVAR estimate. When cross-correlation is used to make estimate the AVAR of a pair of oscillators whose noise is less than the noise floor of a single measurment subsystem, the measurement is not representative of the oscillator AVAR until sufficient estimates have been averaged to reduce the instrument noise below the device under test noise. Since the device noise is unknown, the problem of estimating the quality of the measurement appears intractable.

However, there is more information available from the cross estimators. The real part of the cross spectrum is an unbiased estimator of the spectrum [8]. Thus (7) is an unbiased estimator of the AVAR. Similarly, the imaginary part of the spectrum can be used to calculate an AVAR estimate as shown in (8), which is an estimator of the intrument's AVAR noise floor.

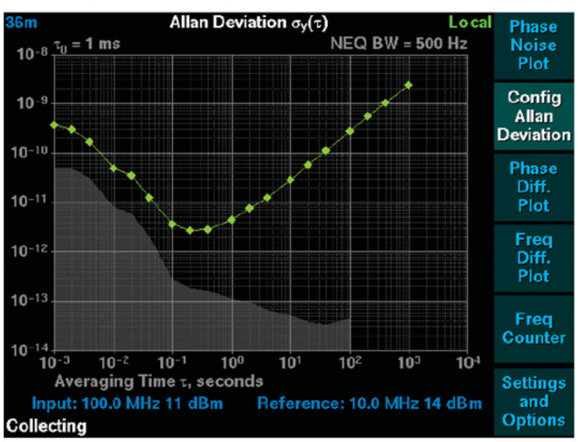

Figure 11: ADEV noise floor in gray calculated from the cross spectrum

$$\sigma_y^{2\otimes}(\tau) = 2\int_0^{f_h} \mathrm{Im}\left[S_\phi^{\otimes}(f)\right]\frac{\sin^4(\pi f\tau)}{(\pi\nu_0\tau)^2}\,df\,, \tag{8}$$

Fig. 10 illustrates the simultaneous measurement and display of the real part of the spectrum (green line) and the imaginary part of the spectrum (top of the grey area). Since the real part of the instrument noise is likely to have the same magnitude as the imaginary part, the grey area is a good estimate of instrumental bias. When it is more than 6 dB below the real part of the spectrum, the instrument noise bias is less than 1 dB and enough averages have been performed for a valid measurement. Fig. 11 illustrates the AVAR noise floor.

V. Conclusions

Direct-digital phase-difference measurement technology has made it possible to simultaneously estimate phase noise and AVAR without the use of phase-lock loops. The real and imaginary components of the cross spectrum have been used to supplement the AVAR estimation process. The real part of the cross spectrum is used for frequency domain filtering of spurious signals that result from synthesis of arbitrary input frequencies. It is also used to estimate the AVAR in the region where the cross AVAR is a biased estimator. Finally, the imaginary part of the cross spectrum is used to estimate the instrumental noise contribution to the AVAR.

Direct-digital measurements also enable more flexibility to set the measurement bandwidth and have led to the conclusion that there is an optimum noise bandwidth for AVAR estimation. Appropriate filtering eliminates aliasing common to heterodyne measurement systems. The increased information available through the use of cross statistics adds complexity but that complexity and the use of proper filtering result in improved measurement accuracy and reliability.

References

[1] S. R. Stein et al. "Comparison of heterodyne and direct-sampling techniques for phase-difference measurements," 2005 NCSL International Workshop and Symposium, unpublished.

[2] James A. Barnes et al. "Characterization of frequency stability," IEEE Trans. Instrum. Meas. vol. IM-20 No. 2, pp. 105–120, May 1971.

[3] David W. Allan and Howard Daams, "Picosecond time difference measurement system", Proc. 29th Ann. Symp. Freq. Contr. pp. 404-411, (1975)

[4] Algie L. Lance, Wendell.D.Seale, Frederick Labaar, "Phase noise and AM noise measurements", in Infrared and Millimetre. Waves, Volume 11, Chapter 7, Academic Press Inc., 1984 pp. 239–289

[5] IEEE Std 1139-1999 ,IEEE standard definitions of physical quantities for fundamental frequency and time metrology, Ju;y 1999.

[6] Samuel R. Stein, "Frequency and time – their characterization and measurement," in Precision Frequency Control, vol. 2, Eduard A Gerber and Arthur Ballato, Eds. New York: Academic Press, 1985, pp. 46–98.

[7] Warren F. Walls, "Cross correlation phase noise measurements," Proc. 46th Annu. Symp. Frequency. Contr., pp. 257-261, (1992)

[8] E. Rubiola, "The magic of cross correlation in measurements from DC to optics," Proc. 22nd European Frequency and Time Forum, April 2008.