

The Allan Variance – Challenges and Opportunities

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

Abstract—The Allan variance has historically been estimated 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 estimated 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 accurately and 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. Investigation of these difficulties has led to several opportunities to improve Allan variance estimation including the ability to estimate the instrumentation noise floor during a measurement and the existence of an optimum measurement bandwidth. In the end, this has led to 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 one input phase changes by one cycle compared to the other, the beat frequency signal shows a one cycle change in phase. Thus the carrier phase-difference information is preserved but the noise induced phase differences become a larger fraction of the period and much easier to measure. One limitation of heterodyne measurements is that atomic frequency standards rarely have the frequency offsets needed to produce the beat frequency.

The dual-mixer time difference measurement system is a form of heterodyne measurement system commonly used to measure the phase difference between two oscillators that have nearly equal frequencies, *i.e.* frequencies within a few Hz of one another [3]. This approach uses a transfer oscillator to measure the phase differences between each oscillator under test and the transfer oscillator. The phase difference between the inputs is estimated by subtracting the two measurements, once they have been resampled so they are time aligned. The technique has the benefit of cancelling most of the transfer oscillator noise. The advantage 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 frequency in dual-mixer time difference measurement systems or to extend the dual-mixer technique further to measure unequal frequency oscillators, *i.e.* when the difference frequency is a significant fraction of the RF frequency. Careful design of such measurement systems results in some of the synthesizer noise cancelling in the measurements [4].

Heterodyne measurement systems are not used to estimate the spectral density of phase noise for several reasons. Such measurement systems obtain phase information from the times of the zero crossings of the IF signal. There is a trade off between the low beat frequencies needed to achieve fine time resolution and the high beat frequencies needed to measure the spectrum far from the carrier. The heterodyne techniques are limited by the presence of the beat frequency signal to a linear gain of approximately 10, after which a combination of amplitude limiting and bandwidth increases are required to produce the fast rise times necessary for low-jitter zero-crossing detection. The high slew rate amplifiers used in zero crossing detection introduce unnecessary noise by using bandwidth in excess of the bandwidth required for the downstream analysis of AVAR [5]. These larger bandwidths are appropriate for spectrum analysis, but are useless without higher beat frequencies, which would then limit the measurement resolution. The study in this paper has also revealed that heterodyne measurement systems suffer from frequency domain aliasing that may corrupt the AVAR estimates and makes this approach unusable for correctly estimating the spectral density of phase noise.

The spectral density of phase noise has traditionally been measured using double-balanced mixers as phase detectors [6]. The usual method is to phase lock the device under test and the reference to one another. 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 suppression of the carrier at the output of the mixer increases the broadband




Figure 1: Cross correlation of phase-measurement subsystems

linear amplification that may be used before the mixer output is sampled. The amplification reduces the contribution of quantization noise of the sampler to an insignificant level. The noise bandwidth and sampling frequency may be chosen to satisfy the Nyquist-Shannon sampling theorem by employing adequate analog filtering in front of the sampler to reduce the signal level above one-half the sampling frequency to negligible levels [7]. This ability makes the phase-lock technique very well suited to estimate the spectral density of phase noise. Advanced techniques suppress the carrier of one of the inputs to the mixer reducing mixer flicker noise [8]. Outside the loop bandwidth, the time series of phase can be used to directly estimate the phase spectrum. Inside the loop bandwidth, the feedback voltage can be used to estimate the frequency spectrum, which can be converted to phase spectrum. However, the phase-lock loop used to maintain quadrature suppresses the carrier phase information inside the loop bandwidth and small calibration errors in the estimate of the frequency-voltage tuning sensitivity of the oscillator make it impossible to precisely recover the phase information. The small frequency errors integrate and the phase error grows linearly with time. Thus, the time series of phase can't be used to estimate AVAR with known accuracy for sample times much longer than the phase-lock loop time constant.

Direct-digital measurement systems resolve the incompatibility between the heterodyne and phase-lock measurement systems. The direct-digital measurement approach digitizes the input signals before any other signal processing and then 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 does not require full conversion to baseband (i.e. a phase lock loop). The total phase is “unwrapped” by adding π radians each time the baseband signal passes the branch cut. The principal limitation of the technique is the noise contribution of the analog-to-digital converters and the poor resolution of the arctangent near the branch cut. These noise contributions require the use of cross-correlation in order to make it possible to measure phase noise of high performance oscillators as shown in Figure 1. The use of the cross spectrum and the cross AVAR creates new challenges for achieving unbiased estimates of AVAR. Because of the widespread use of AVAR in specifying oscillator and clock performance, it is very advantageous to find approaches that use the cross statistics to produce unbiased estimates of AVAR or logical extensions of AVAR when appropriate. Many of the plots shown in the following report the Allan deviation – the square root of AVAR – which is abbreviated ADEV.

II. OVERVIEW

AVAR is a measure of how much the frequency of an oscillator changes from one sample interval of duration τ to the next interval with no intervening dead time. AVAR replaced the variance as a measure of frequency stability because the value of the estimate of the variance does not converge as the number of estimates increases for many of the noise processes that are commonly used to model precision clocks and oscillators [9]. Using the IEEE recommended definition of frequency based on the end point phases, AVAR can be written in terms of the phase series, ϕ_k [10],

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

where the angle brackets signify expectation value. AVAR can also be calculated from the phase noise spectrum, $S_\phi(f)$, as shown in (2)

$$\sigma_y^2(\tau) = 2 \int_0^{f_b} S_\phi(f) \frac{\sin^4(\pi f \tau)}{(\pi \nu_0 \tau)^2} df \quad (2)$$

III. CHALLENGES ESTIMATING AVAR

A. Using Heterodyne Phase Measurements

It is useful to distinguish between analog phase measurements that rely on a transducer to convert phase to a measurable quantity and digital phase measurements that only utilize techniques such as counting or taking ratios that don't require user calibration. The heterodyne method of estimating the phase difference between oscillators is a form of digital measurement since only a counter is required to determine the time of occurrence of each zero crossing. The down-conversion process preserves a cycle of phase [11]. 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 maximum and minimum. Only one sign of zero crossing can be used because the zero voltage level of the waveform moves relative to the zero phase point due to threshold noise in the comparator circuit. This noise appears as duty cycle variations – the positive and negative zero crossings move relative to one another – making it difficult, perhaps impossible, to use the second set to complement the phase information of the first set and thereby improve the performance level. The same phenomenon has been observed to limit the phase noise performance of digital dividers with analog input signals [12].

As a result, heterodyne phase-difference measurements for low-noise applications comprise at most one sample per full IF cycle. The IF is low-pass filtered to avoid the large phase




Figure 2: Frequency domain aliasing plotted on linear axes

shifts of band pass filters. The bandwidth is often chosen to be at least three times the IF frequency in order to pass the 3rd or higher harmonic and improve the slew rate of the mixer output near the zero crossing, but it must at the very minimum equal the IF frequency. A band pass filter can't be used without the filter's temperature coefficient dominating the low frequency noise of the measurement. Thus, the heterodyne measurement 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. The low-pass filter that passes the IF cannot control signal content between one-half the sampling frequency and the sampling frequency. All signal components above one-half the sampling frequency appear as aliases in the frequency band between 0 and one-half the sampling frequency as shown in Figure 2, where (a) shows the true spectrum and (b) shows the result of computing the spectrum from an under-sampled time series without anti-alias filtering. If there is significant noise power in the region from one-half to one times the IF frequency, then the aliased signals are present all the way to baseband and no further additional digital signal processing can eliminate their effect on ADEV and spectrum calculations. Other papers have discussed how some counters implement averaging algorithms that differ from the way the AVAR definition weights data [13]. Here we have seen how resolution enhancement by period expansion redistributes the noise in the frequency domain and may have significant impact on the AVAR calculation. There is no way to know whether this has happened without a priori knowledge of the phase noise spectrum.

This inevitable aliasing affects AVAR in some cases but not in others. If the dominant noise process is white phase noise, such as may occur in quartz crystal oscillators and hydrogen masers at short averaging times, AVAR is insensitive to aliasing. Suppose the noise density is σ_0^2 and the measurement bandwidth is f_h , which is assumed to be rectangular for simplicity. Substitution into (2) yields

$$\sigma_y^2(\tau) = \sigma_0^2 f_h \quad (3)$$

After aliasing, the noise density is $2f_h\sigma_0^2/f_s$ and the bandwidth is $f_s/2$. The product and AVAR are unchanged. However,




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

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. When the spectral density is a narrow line, AVAR is proportional to the window function of (2) evaluated at the frequency of the narrowband signal. The component of AVAR due to this signal has maxima at odd multiples of the inverse of twice the signal frequency and zeros at multiples of the inverse of the signal frequency. Thus, the true AVAR has a peak at 8.83 ms (half the period) and every odd multiple and a zero at every multiple of 16.67 ms (the period of the AC power). However, the signal appears as an alias at 40 Hz and the estimated AVAR has peak at every odd multiple of 12.5 ms and a zero at every multiple of 25 ms. Figure 3 shows the variation in the ADEV estimates between heterodyne phase-difference measurements with aliased spurs and simultaneous direct-digital measurements with spurs at the correct frequency for the case when ADEV has noticeable power-line spurious signals. Two trend lines are plotted that correspond to the peaks of the power line phase modulation and the oscillator phase noise respectively.

Although a bit outside of the focus of this paper, it is notable that aliasing can make heterodyne phase-difference measurements unreliable for spectrum estimation. This may occur if the down conversion folds noise, especially white phase noise, onto the heterodyne signal, causing spectral estimation biases. Figure 4 shows the result of estimating the single-sideband phase noise with a 100 Hz beat frequency heterodyne measurement system having 450-Hz noise bandwidth. A calibrated source with -110 dBc/Hz single sideband phase noise (SSB) was incorrectly estimated to have -100 dBc/Hz SSB. The excess noise in dB is calculated as $10 \cdot \log(450/50)$ or 9.5 dB which compares well with the 10 dB measured. The same measurement system estimates the correct ADEV because ADEV is insensitive to whether the white noise is spread over 450 Hz or folded into the 50-Hz half sampling rate. When used with care, 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 decreases rapidly with




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

frequency. The rapid roll-off of the noise with Fourier frequency prevents the noise originating above one-half the sampling frequency from significantly affecting the value measured much closer to the carrier.

B. Using Phase Noise Measurements

The method of estimating AVAR by integrating the spectral density of phase using (2) is generally 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 AVAR corresponding to the noise and the spurs can be computed separately. It is even more useful when the spurious signals can be shown to originate in the measurement process and may be deleted before they corrupt the AVAR estimate as discussed in Section III.D

The first peak in the integration kernel in (2) occurs at a frequency of $1/2\tau$. 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. The bandwidth dependence of AVAR shown in (2) and (3) makes it imperative to report the measurement bandwidth. Figure 5 shows the traditional plot of the measurement dependence. Both the calculated and measured ADEV of a calibrated white phase noise source are plotted as the measurement bandwidth is varied from 500 Hz to 0.5 Hz. For each sample rate f_s , the measurement bandwidth is $f_s/2$. The measured values are shown as individual plotted points, while the straight lines were calculated from (3) using the phase noise density of the source and the measurement bandwidth, which was varied over the range of 500Hz to 0.5 Hz. Section IV.B will show a more complete picture of the bandwidth dependence that has emerged since the direct digital phase-difference measurements have extended the lower limit to the range of available measurement bandwidths, which had previously been equal to f_s for heterodyne measurement systems.

C. Using Cross Correlation

Cross correlation is often used to reduce the instrument noise contribution to the measurement [14]. 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.

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

If heterodyne or direct-digital phase-difference measurements are used, then AVAR can be calculated even when the long-term oscillator frequency variations are dominated by temperature or aging effects. However, these deterministic frequency variations result in cross statistics that are quite noticeably different from the usual estimator. If a single reference oscillator is used as in Fig. 1, then the aging induced cross AVAR is given by

$$\sigma_y^{2\otimes}(\tau) = \frac{1}{2} (D_{IN} - D_{REF})^2 \tau^2, \quad (5)$$

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

The cross AVAR also performs a separation of variances if the independent reference oscillators are used as references for the two phase-difference measurement sub-systems as shown in Figure 6. Under these conditions, the cross AVAR may be a biased estimator of AVAR and doesn't necessarily represent it well. For example, in the region where the differential aging between the input and reference oscillators is the dominant source of frequency variation, the cross ADEV is given by

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




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

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 ADEV doesn't exist. The most obvious fix is to use the absolute value of the cross AVAR as an estimator of AVAR. However, this approach results in a notch in AVAR as shown in Fig. 7. The notch can be smoothed by integrating the cross spectrum to obtain AVAR.

$$\sigma_y^{2\otimes}(\tau) = 2 \int_0^{f_h} \text{Re} \left[S_\phi^\otimes(f) \right] \frac{\sin^4(\pi f \tau)}{(\pi v_0 \tau)^2} df \quad (7)$$

An example of fixing the AVAR notch using this technique is shown in Fig. 8.

D. Using Direct-Digital Phase-Difference Measurements

The input frequency v_{IN} is sampled with frequency v_s . Because this process is non-linear, spurs are created at all frequencies $M v_{IN} - N v_s$. When the frequency of the spur lies within the AVAR bandwidth and M and N are less than approximately 300, the AVAR estimate is noticeably affected. Most of these spurs can be identified and removed. Spurious modulation of the input signal must always produce a line in the spectrum whose phase is 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 AVAR.

Thus, when cross correlation is used to reduce the instrumentation contribution to the measurement noise, integration of the cross spectrum using (7) is a superior method of generating an unbiased estimate of AVAR than the use of (4) directly from very small sample times up to a few thousand seconds. From the practical point of view, it is very difficult to compute AVAR from the cross spectrum for times longer than about two-thousand seconds. The kernel of the integral in (7) peaks at $f = 1 / 2\tau$ and AVAR depends significantly on the first few bins of the spectrum estimate, which are known to be biased [15]. The impact of the lowest frequency portion of the spectrum on AVAR also depends on the method of numerical integration. As a result, it is not unusual to obtain a result that looks like Figure 9. Thus, for times longer than 1 s, a better AVAR estimate is obtained by selecting the largest of (7) and the absolute value of (4). Selection of the larger estimate eliminates both the notch when the cross variance is negative as well as the low values at long τ due to possible underestimation of first bin of the spectrum.




Figure 6: Cross correlation of phase-measurement subsystems and reference oscillators

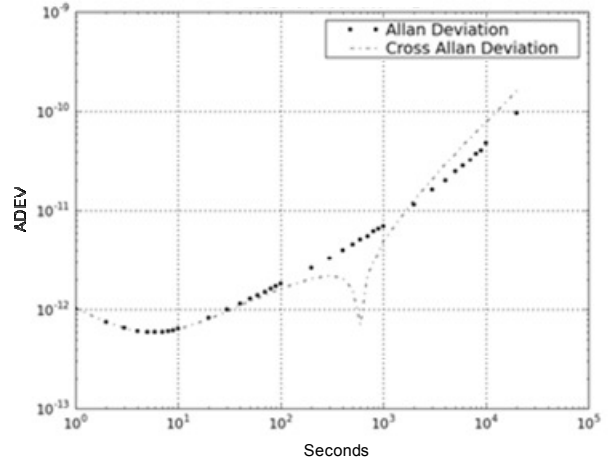


Figure 7: The absolute cross ADEV is a biased estimate of ADEV

IV. OPPORTUNITIES FOR AVAR ESTIMATION USING DIRECT-DIGITAL MEASUREMENT SYSTEMS

A. Flexible Filtering

Direct-digital measurement systems make it possible to satisfy the requirements of the sampling theorem. The RF signal can be sampled at a rate larger than twice the analog anti-alias filter bandwidth out to a frequency where there is sufficient rejection of undesirable signals. Digital anti-alias filtering and further sub-sampling allow total flexibility in the choice of measurement bandwidth. The maximum measurement bandwidth set by the sampling theorem is $(2\tau_{\min})^{-1}$, where τ_{\min} is sample period as well as the minimum sample interval for calculating AVAR.

Figure 10 and Figure 11 show the design of a typical anti-alias filter and its effect on ADEV respectively. The unfiltered data have $\tau_{\min} = 100$ ms and are filtered with a bandwidth of $\frac{1}{2}$ Hz. Fig. 11 shows that for $\tau < 1$ s, ADEV is highly filter dependent and should not be reported as demonstrative of device performance. Thus τ_{\min} is 1 s for the

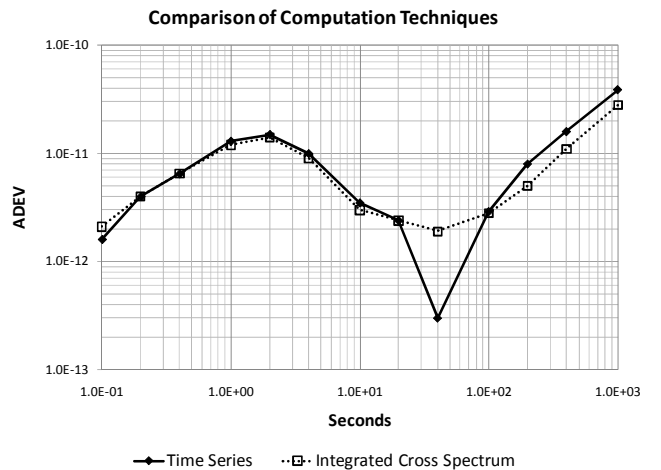


Figure 8: The integrated spectrum improves the ADEV estimate

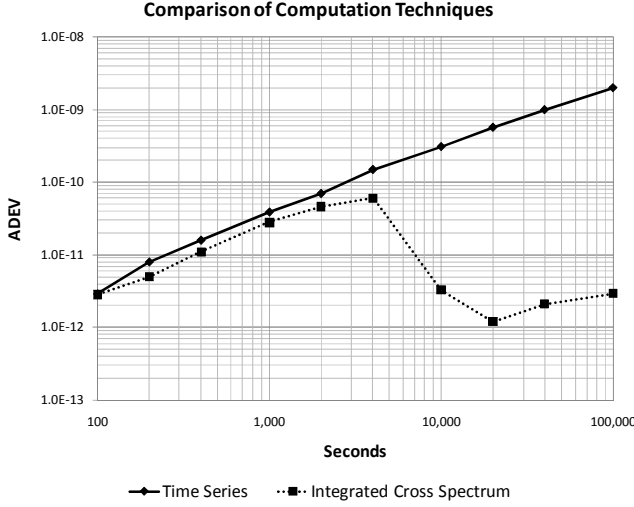


Figure 9: The integral of the spectrum produces poor ADEV estimates for $\tau > 2000$ s

anti-alias filtered data. The combination of the requirements of the sampling theorem and the effect of the anti-alias filter on AVAR leads to the conclusion that there is an optimum measurement bandwidth, f_h , that should be used to compute AVAR. The effect of the measurement bandwidth on AVAR sets the smallest bandwidth that is representative of the device. The sampling theorem sets the largest measurement bandwidth that is representative of the device and the two limits are equal. Thus, the optimum measurement bandwidth is $(2\tau_{\min})^{-1}$.

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 [16]. 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 affected the AVAR estimate. When cross-correlation is used

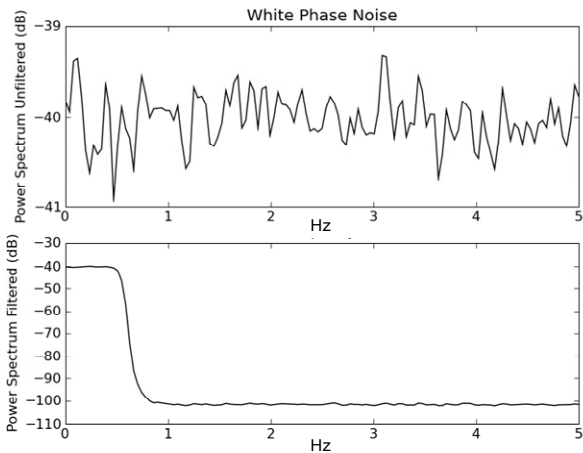


Figure 10: Anti alias filter design with $\frac{1}{2}$ Hz noise bandwidth




Figure 11: ADEV has an optimum bandwidth for any sample rate

to estimate AVAR of a pair of oscillators whose noise is less than the noise floor of a single measurement 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 is more difficult than the case of the standard AVAR.

However, more information is available from the cross statistics to help resolve this problem. The real part of the cross spectrum is an unbiased estimator of the spectrum [17]. Thus (7) is an unbiased estimator of 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 instrument's AVAR noise floor.

$$\sigma_y^{2\otimes}(\tau) = 2 \int_0^{f_h} \text{Im} \left[S_{\phi}^{\otimes}(f) \right] \frac{\sin^4(\pi f \tau)}{(\pi v_0 \tau)^2} df \quad (8)$$

Fig. 12 illustrates the use of the imaginary part of the spectrum to perform an in-situ estimate of the noise floor of the measurement. The upper black line is the measured spectrum between two commercial high-performance cesium beam frequency standards. Two estimates of the measurement noise are plotted below it. The gray curve shows the imaginary part of the spectrum, measured simultaneously with the devices under test. 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 noise. 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. A traditional noise floor measurement was performed subsequent to measuring the devices under test by splitting the signal from one source and applying it to both inputs of the measurement system. The results are over-plotted as grey dots. The dark grey area showing the overlap of the two noise floor estimates illustrates the quality of the in-situ noise floor measurement.




Figure 12: The cross spectrum contains noise floor information

Figure 13 illustrates AVAR and its noise floor calculated from (8) during the same set of measurements.

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 internally generated spurious signals that result from synthesis of arbitrary input frequencies. It is also used to estimate 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 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

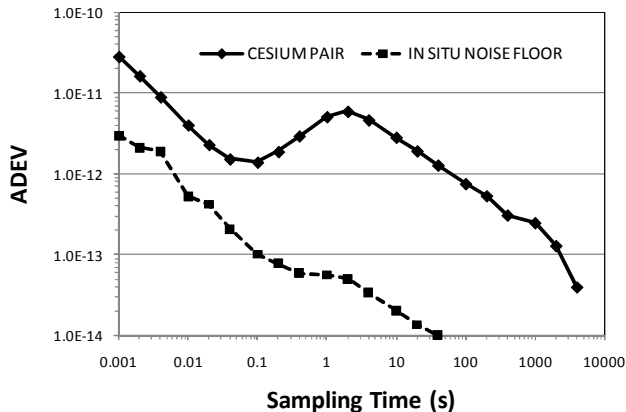


Figure 13: ADEV noise floor in gray calculated from the cross spectrum

complexity, but that complexity and the use of proper filtering result in improved measurement accuracy and reliability.

REFERENCES

- [1] S. R. Stein et al. "Direct-digital phase-noise measurement," Proc. IEEE Int'l Freq. Contr. Symp. pp. 287-291, (2004).
- [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] S.R. Stein et al. "Comparison of heterodyne and direct-sampling techniques for phase-difference measurements" Proc. of the 2005 NCSL International Workshop and Symposium (2005).
- [5] G.J. Dick, P.F. Kuhnle, R.L. Sydnor, "Zero-crossing detector with sub-microsecond jitter and crosstalk," Proc. Of the 22nd Annual Precise Time and Time Interval Applications and Planning Meeting, pp. 269-282, (1990).
- [6] L.S.Cutler and C.L. Searle, "Some aspects of the theory and measurement of frequency fluctuations in frequency standards", Proc. IEEE, vol. 54 no. 2, pp. 136-154, Feb. 1966.
- [7] C.E. Shannon, "Communication in the presence of noise", Proc. Institute of Radio Engineers, vol. 37, no. 1, pp. 10-21, Jan. 1949.
- [8] E. N. Ivanov et al. "Advanced phase detection technique for the real time measurement and reduction of noise in components and oscillators," Proc. IEEE Int'l Freq. Contr. Symp. pp. 509-514, (1997).
- [9] D.W. Allan, "Statistics of atomic frequency standards", Proceedings of the IEEE, Vol. 54, No. 2, 221-231, 1966.
- [10] IEEE Std 1139-1999, IEEE standard definitions of physical quantities for fundamental frequency and time metrology, July 1999.
- [11] 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.
- [12] J.D. Crockett and M. Siccardi, "Power supply noise conversion to phase noise in CMOS frequency digital divider", Proc. 22nd Ann. Symp. Freq. Contr. pp. 699-702, (2002).
- [13] S.T. Dawkins, J.J. McFerran, and A.N. Lutien, "Considerations on the measurement of the stability of oscillators with frequency counters," IEEE Trans. On UFFC, vol. 54 pp 918-925 (2007).
- [14] Warren F. Walls, "Cross correlation phase noise measurements," Proc. 46th Ann. Symp. Frequency. Contr., pp. 257-261, (1992).
- [15] F.L. Walls, D.B Percival, and W.R.Ireland, "Biases and variances of several FFT spectral estimators as a function of noise type and number of samples," Proc. 43rd Ann. Symp. Frequency. Contr., pp. 336-341, (1989).
- [16] D.A. Howe, D.W. Allan, and J.W. Barnes, "Properties of signal sources and measurement methods," Proc. IEEE Int'l Freq. Contr. Symp. pp. A1-A47, (1981).
- [17] E. Rubiola, "The magic of cross correlation in measurements from DC to optics," Proc. 22nd European Frequency and Time Forum, April 2008.