

## Measuring the Allan variance by sinusoidal fitting

Ralph G. DeVoe

Citation: [Review of Scientific Instruments](#) **89**, 024702 (2018); doi: 10.1063/1.5010140

View online: <https://doi.org/10.1063/1.5010140>

View Table of Contents: <http://aip.scitation.org/toc/rsi/89/2>

Published by the [American Institute of Physics](#)

---

### Articles you may be interested in

[A practical double-sided frequency selective surface for millimeter-wave applications](#)

[Review of Scientific Instruments](#) **89**, 024703 (2018); 10.1063/1.5023406

[Note: An approach to 1000 T using the electro-magnetic flux compression](#)

[Review of Scientific Instruments](#) **89**, 016106 (2018); 10.1063/1.5003411

[Simple, low-noise piezo driver with feed-forward for broad tuning of external cavity diode lasers](#)

[Review of Scientific Instruments](#) **89**, 023102 (2018); 10.1063/1.5009643

[Invited Article: Concepts and tools for the evaluation of measurement uncertainty](#)

[Review of Scientific Instruments](#) **88**, 011301 (2017); 10.1063/1.4974274

[Cryogenic positioning and alignment with micrometer precision in a magnetic resonance force microscope](#)

[Review of Scientific Instruments](#) **89**, 013707 (2018); 10.1063/1.5008505

[The performance and limitations of FPGA-based digital servos for atomic, molecular, and optical physics experiments](#)

[Review of Scientific Instruments](#) **89**, 025107 (2018); 10.1063/1.5001312

---



**PFEIFFER VACUUM**

**VACUUM SOLUTIONS FROM A SINGLE SOURCE**

Pfeiffer Vacuum stands for innovative and custom vacuum solutions worldwide, technological perfection, competent advice and reliable service.

[Learn more!](#)

## Measuring the Allan variance by sinusoidal fitting

Ralph G. DeVoe

Physics Department, Stanford University, Stanford, California 94305, USA

(Received 23 October 2017; accepted 19 January 2018; published online 6 February 2018)

The Allan variance of signal and reference frequencies is measured by a least-square fit of the output of two analog-to-digital converters to ideal sine waves. The difference in the fit phase of the two channels generates the timing data needed for the Allan variance. The fits are performed at the signal frequency ( $\approx 10$  MHz) without the use of heterodyning. Experimental data from a modified digital oscilloscope yield a residual Allan deviation of  $3 \times 10^{-13}/\tau$ , where  $\tau$  is the observation time in s. This corresponds to a standard deviation in time of  $<300$  fs or  $20 \mu\text{rad}$  in phase. The experimental results are supported by statistical theory and Monte Carlo simulations which suggest that optimized devices may have one or two orders of magnitude better performance. *Published by AIP Publishing.* <https://doi.org/10.1063/1.5010140>

### I. INTRODUCTION

The Allan variance is a widely used statistical tool<sup>1–3</sup> used to characterize the noise and instability of frequency and time standards<sup>4</sup> such as quartz oscillators, rubidium and cesium standards, and trapped and cooled atomic frequency references. Two methods of measuring the Allan variance are common today, both of which rely on heterodyning for their resolution.<sup>5</sup> The first is the dual mixer time difference (DMTD) technique.<sup>6–8</sup> In this method, the unknown and reference signals are heterodyned by a local oscillator down to a low frequency (usually from 1 to 1000 Hz) where measurements of the zero-crossing yield timing information. The resolution is enhanced by the ratio of oscillator to beat frequency, usually a factor of  $10^4$ – $10^7$ . Although the DMTD method can be highly precise,<sup>8</sup> with a residual Allan deviation (ADEV)  $\approx 5 \times 10^{-14}/\tau$ , where  $\tau$  is the observation time in s, its resolution depends on the analog design of the zero-crossing detector, filters, and amplifiers. The second method is the direct-digital method<sup>5,9–12</sup> in which a high speed analog-to-digital converter (ADC) digitizes both unknown and reference signals, and the results are digitally down-converted into in-phase (I) and quadrature (Q) signals by synthesized local oscillators. The arctangent function is then used to compute the phase, from which timing information is generated. We call this the DD method to differentiate it from other methods using ADC's. The DD method has minimal dependence on analog circuit issues but requires substantial hardware and firmware to perform the I/Q demodulation, synthesize the local oscillators, and process the data through the digital logic. The limiting ADEV of that method is comparable to that of DMTD systems, with commercial devices specifying ADEV of  $10^{-13}$ – $10^{-15}$  for  $\tau = 1$  s.<sup>13</sup>

In this article, we demonstrate a digital Allan variance technique which differs from both the DMTD and DD methods but which has competitive resolution. It does not use heterodyning in either analog or digital form. Instead the timing or phase information is extracted by a least-square fit of the outputs of two ADC's to an ideal sine wave. The resolution derives from the averaging inherent in a least-square

fit: a fit of  $\mathcal{O}(10^4)$  points will enhance the resolution for stationary random noise processes by a factor  $\mathcal{O}(10^2)$ . Initial experiments using a modified general purpose digital oscilloscope yield a residual ADEV of  $<3 \times 10^{-13}/\tau$ , corresponding to a time deviation  $\sigma = 220$  fs or a phase deviation of  $20 \mu\text{rad}$ .

The technique resembles a conventional frequency counter in some respects. Where a counter initiates a measurement with a single zero-crossing, the fitting method determines the phase of the signal and reference by least-square fits of  $\mathcal{O}(10^4)$  points. For example, a 100 MHz ADC would sample a 10 MHz sine wave every 10 ns and  $10^4$  points would occupy  $100 \mu\text{s}$  at the beginning of the interval  $\Delta t$ . The fit yields the most probable phase over this  $100 \mu\text{s}$  interval. Since the fit phase depends on the values of the signal over many cycles, it is insensitive to artifacts near the zero-crossing and to dc offsets. The resolution enhancement due to fitting is dependent on prior knowledge of waveform; the method only works with sinusoidal signals, unlike a counter which can measure intervals between any two pulses.

This technique can be easier to implement than either the DMTD or DD methods. The ADC's would typically be part of a digital storage oscilloscope (DSO), which would write blocks of signal and reference data to files using its standard external trigger, once per  $\Delta t$ . The least-square fit is performed by an independent Python routine which reads each file as it is created and outputs the phase difference for each file. This eliminates the need for special-purpose hardware or firmware.

Several different implementations of the DD approach have been developed recently. One system<sup>10</sup> uses a commercial SDR (software defined radio) module to perform the phase measurement via I/Q demodulation. Another<sup>11</sup> uses a programmable gate array driven by two ADC's to mathematically execute the I/Q demodulation. A third<sup>12</sup> downconverts the signal from 5 MHz to 4 Hz by subsampling and then extracts the phase by a least-square fit. Sinusoidal fitting differs from previous work in that it does not employ downconversion or I/Q demodulation; comparable resolution is achieved by directly fitting the ADC output at the signal frequency. Sinusoidal

fitting should also be distinguished from other methods of processing timing data using least-square fits, such as the Omega counter<sup>14</sup> and the parabolic variance,<sup>15</sup> which fit timing data after they are generated by existing methods. Finally, high-resolution phase measurements have also been achieved with a “tracking DDS”<sup>16</sup> (direct digital synthesizer) servo which has a different design than Ref. 9.

## II. DESCRIPTION OF THE EXPERIMENT

A general purpose DSO has been used to measure the Allan deviation of several quartz and rubidium frequency standards, as well the residual deviation of the device itself. The DSO was a Digilent Analog Discovery<sup>17</sup> which contains an AD9648 dual 14-bit ADC, running at a clock rate near 100 MHz. The standard external trigger of the DSO was driven at intervals  $\Delta t$ , where  $0.1 < \Delta t < 10$  s, by an auxiliary timer, and 8000 points of signal and reference data were recorded at  $\approx 10$  ns per point. The DSO wrote a data file once per trigger containing 80  $\mu$ s of data, covering  $\approx 800$  cycles of the signal. See Fig. 1. A run typically consisted of a set of  $10^3$ – $10^4$  separate files.

Data analysis was performed by a Python routine which read each file and performed a least-square fit of the unknown and reference signals to a function of the form

$$S(t) = A \sin(2\pi f_0 t + \varphi) + \epsilon, \quad (1)$$

where  $A$  is the amplitude,  $\varphi$  is the phase,  $f_0$  is the frequency, and  $\epsilon$  is the DC offset. The fitting routine used was the `curve_fit` method of the `scipy.optimize` library. The fit finds the values of these parameters which minimize the residual  $R$  defined by

$$R^2 = \frac{1}{N} \sum_{i=0}^N [D(i) - S(i \times t_c)]^2, \quad (2)$$

where  $D(i)$  is the ADC output,  $t_c$  is the period of the clock, and  $N$  is the number of points in a file. Only  $\varphi$  is used for the timing data; the fit frequency  $f_0$  is poorly determined during

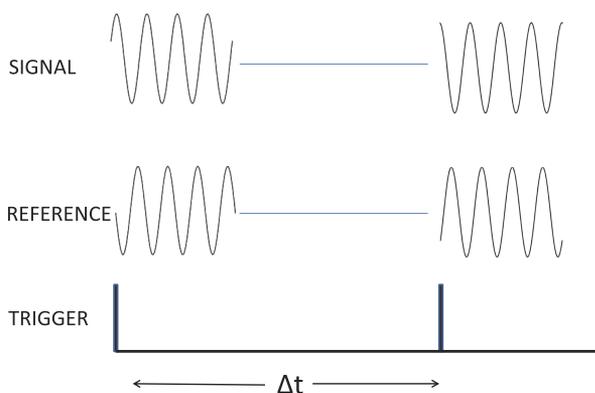


FIG. 1. Timing diagram (not to scale). The digital oscilloscope is triggered every interval  $\Delta t$ , where  $0.1 < \Delta t < 10$  s. Each channel records 8000 points of signal and reference at a  $\approx 100$  MHz rate, which occupy the first 80  $\mu$ s of  $\tau$ . For clarity, only 4 cycles of data are shown in an expanded scale, instead of  $\approx 800$  cycles of the 10 MHz signal. An independent Python routine reads each file and performs a least-square fit of the signal and reference, yielding values of the signal and reference phase  $\varphi_S$  and  $\varphi_R$ . The phase difference  $\varphi_D$  is used to compute the Allan variance. See text.

the short 80  $\mu$ s digitizing interval. The clock period  $t_c$  must be known in advance to moderate accuracy ( $10^{-8}$ ), but its value is not critical for the same reason. For convenience, the fit used only 4096 points from the 8000 point file. The fitting routine requires approximate initial values of these parameters for convergence. The starting phase was estimated by a routine which computes which quadrant  $\varphi$  lies in by comparing adjacent points. Good timing data were produced when the sine wave fits had residuals  $R < 1.5 \times 10^{-3}$  of the amplitude  $A$ , where  $A \approx 2$  V. Typical values of  $\epsilon$  were 20 mV or about  $10^{-2}$  of  $A$  and were stable over the course of the experiment. Data from the signal  $S$  and reference  $R$  channels were independently fitted and the phase difference  $\varphi_D = \varphi_S - \varphi_R$  was written to a file together with  $\varphi_S$  and  $\varphi_R$  and the residuals. The relative frequency of the  $S$  and  $R$  channels for the  $j$ -th trigger was computed from  $(\varphi_D(j+1) - \varphi_D(j))/2\pi\tau$ .

Two modifications were made to the DSO for these tests. First, the internal 100 MHz clock is an integral multiple of the standard frequencies of 5 and 10 MHz. This could limit the resolution by repeated sampling of a few points on the waveform. An external clock was therefore coupled into the clock generator<sup>17</sup> through a ground-isolated transformer so that the clock frequency was 97.2 MHz. Other clock frequencies were tried but did not change the experimental results. Second, this DSO has a frequency response which is limited by the preamplifiers to  $< 30$  MHz. The preamps generate second-harmonic distortion at 10 MHz which increases the residuals  $R$ . Therefore the preamplifiers were bypassed and the ADC inputs were driven directly by ground-isolated transformers. No other changes were made to the hardware or software of the scope.

Figure 2 shows the results of a residual noise measurement, in which the  $S$  and  $R$  channels were driven by the same 10 MHz signal. Low pass filters were used to reduce second and higher order harmonics to 75 db below the carrier. The signal was run through a resistive power splitter and then through 6 db attenuators to isolate the inputs from each other. The average delay is about 57.25 ps, corresponding to  $\approx 11$  mm of cable

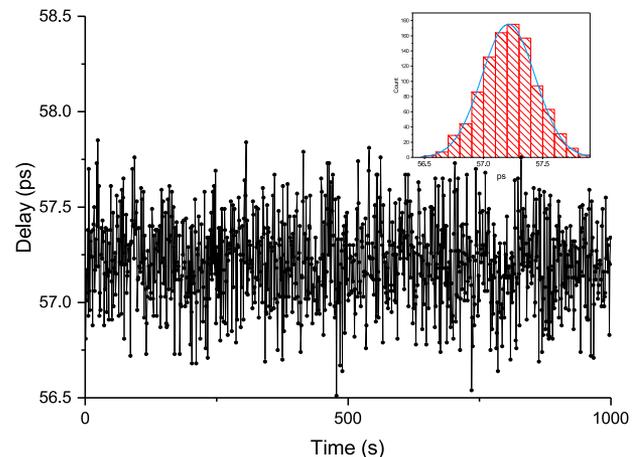


FIG. 2. Experimental results showing the instrumental noise with both channels connected to the same source. The phase difference  $\varphi_D$  is expressed in terms of time delay by  $\varphi_D/2\pi f_0$ . The data were taken once per s for 1000 s. The mean delay between channels A and B is 57.25 ps which corresponds to 11 mm of cable length. Inset: histogram of the data. The data fit a Gaussian with a standard deviation of  $\sigma = 220$  fs.

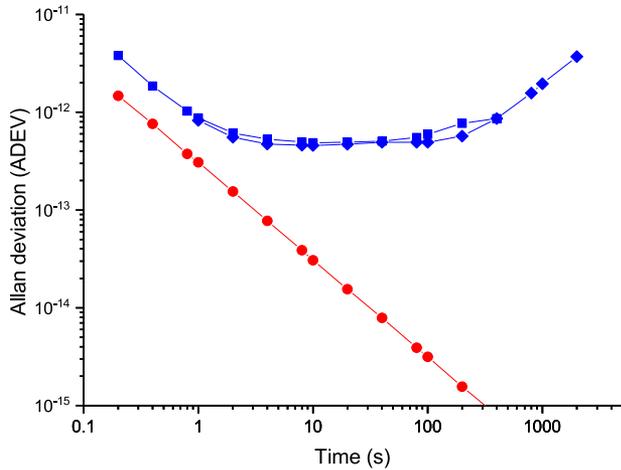


FIG. 3. Allan deviation (ADEV) of a residual measurement (red dots) and for a pair of identical quartz oscillators (blue squares and diamonds). For the residual measurement, both channels were driven by the same source where 10 000 points were taken at 0.2 s intervals. The ADEV of  $3 \times 10^{-13}$  at  $\tau = 1$  s is consistent with  $\sigma \approx 300$  fs in Fig. 2. For the quartz oscillators, the data were taken in two overlapping runs of 10 000 points each: one at  $\Delta t = 0.2$  s (squares) and one at  $\Delta t = 1$  s (diamonds). See text.

mismatch between the two channels. The inset shows a histogram of the data which follows a Gaussian distribution with a standard deviation  $\sigma$  of 220 fs. This is somewhat greater than the timing jitter specified for this ADC which is about  $\approx 140$  fs. Note that during this run, which lasted for 1000 s, the phase of each channel varied over the full 100 nS period of the inputs while their *difference* had  $\sigma = 220$  fs, a ratio of  $3 \times 10^{-6}$ .

Figure 3 shows the Allan deviation of the residual noise as well as a measurement of 2 identical FTS-1050a quartz crystal standards. The Allan deviation was computed from the Python data by two different routines: one based on the AllanTools Python library and the other using TimeLab. As expected, 300 fs  $\sigma$  corresponds to  $\approx 3 \times 10^{-13}$  ADEV at  $\tau = 1$  s and drops linearly with  $\tau$ . The upper trace of the two quartz standards shows an ADEV below  $10^{-12}$  in the 1-100 s range and the characteristic rise in ADEV for  $\tau < 1$  s and for  $\tau > 100$  s.

It is important to note that least-square fitting creates an implicit low-pass filter for the phase noise sidebands of the signal. Sidebands with a period less than the fitting time of 40  $\mu$ s will be filtered out since the fit assumes that the phase is constant during the interval. Conversely sidebands with a period longer than 40  $\mu$ s will vary little during the fit and pass through. This corresponds to a low pass filter for the fit phase with a cutoff  $f_H \approx 10$  kHz. The cutoff frequency  $f_H$  plays an important role in the relationship between the Allan variance and the phase noise spectral density,<sup>2</sup> as discussed in detail in Ref. 11.

### III. ESTIMATING RESOLUTION

Least-square fitting of data to a sine wave is a standard topic in signal processing theory.<sup>18</sup> It has been used for many years to measure the effective number of bits (ENOB) of an ADC<sup>19</sup> and is embodied in an IEEE standard.<sup>20</sup> However, it has not previously been applied to Allan variance measurements. We shall show that our measured resolution  $\sigma$  is consistent

with theory by three different approaches: first, signal processing theory yields a simple analytic expression for  $\sigma$ ; second, a graphical intuitive picture agrees giving a similar result without the statistical formalism; and third, a Monte Carlo routine is used to test these approximations by generating simulated data which are processed by the same analysis program used in the experiments.

One of the simplest methods of estimating the power of a fitting algorithm is to compute the Cramers-Rao lower bound.<sup>18</sup> In this case, we want to estimate the standard deviation of the phase  $\sigma_\varphi$  resulting from a least-square fit to a noisy sinusoidal signal. Assume a sinusoidal signal of amplitude  $A$  with the white Gaussian additive noise with a standard deviation  $\sigma_A$ . Then the Cramers-Rao lower bound is given by

$$\sigma_\varphi = \frac{\sigma_A C}{A \sqrt{M}}, \quad (3)$$

where  $\sigma_\varphi$  is in radians,  $C$  is a constant of order unity, and  $M$  is the number of samples in the fit. This result assumes a 3-parameter fit; the offset  $\epsilon$  is ignored. The samples  $M \gg 1$  are assumed to be distributed over many cycles of the sine wave. The constant  $C$  is of order unity and depends on how many parameters are free and how many fixed.<sup>18</sup> For a 1-parameter fit, where only  $\varphi$  is free,  $C = \sqrt{2}$ , while for a 3-parameter fit, where neither  $\varphi$ ,  $f_0$ , or  $A$  are known,  $C = \sqrt{8}$ .

Next we relate the white noise  $\sigma_A$  to the quantization noise of the ADC. Quantization of a sinusoid has been studied in detail in Ref. 21. A general rule is that  $\sigma_A \approx 1/\sqrt{12}$  of the least significant bit (lsb). Then the effective additive white noise due to the discreteness of the ADC is

$$\frac{\sigma_A}{A} = \frac{1}{2^N \sqrt{12}}, \quad (4)$$

where  $N$  is the ENOB. The ENOB represents the dynamic accuracy of an ADC in following a rapidly changing signal. AD9648 ADC in the current device is nominally 14 bits, but the ENOB is specified<sup>22</sup> at 12.0 bits for a 10 MHz input. Substituting in Eq. (2) yields

$$\sigma_\varphi = \frac{1}{2^N} \frac{C}{\sqrt{M} \sqrt{12}} \quad (5)$$

with  $\sigma_\varphi$  in radians. Choosing  $C = \sqrt{8}$  from above,  $C/\sqrt{12} = 0.816 \approx 1$  and may be ignored. Converting from radians to seconds with  $1/2\pi f_0$ , we get

$$\sigma \approx \frac{1}{2\pi f_0 2^N \sqrt{M}}. \quad (6)$$

It is important to emphasize that this result is a statistical lower bound assuming only quantization noise. For the conditions of the above experiment, where  $N \approx 12$ ,  $M = 4096$ , and  $f_0 = 10$  MHz, Eq. (6) yields  $\sigma = 60$  fs for a single fit. The analysis program computes the difference of two uncorrelated fits for the signal and reference channels so that the errors add in quadrature, giving a lower bound of 85 fs. This should be compared to the 220 fs measured value. Note that the differential timing jitter of the ADC<sup>17</sup> is  $\approx 140$  fs, which may account for the difference.

One application of this theory is to answer the question: which is more important, ADC resolution or ADC speed? That

is, is it better to have a 12 bit ADC at 100 MHz or an 8 bit ADC at 1 GHz? Eq. (6) suggests the former since the ADC clock rate does not enter into the equation. See, however, the Monte Carlo discussion below.

There is a simple graphical construction which provides a more intuitive explanation of Eq. (6). Consider first how a single ADC measurement can be used to predict the zero-crossing time of a sinusoid, which is equivalent to a phase measurement. Assume that the measurement is made during a linear part of the sine function, e.g., where the sine is between  $-0.5$  and  $+0.5$  so that  $\sin(x) \approx x$ . Further assume that the ADC measurement yields the integer value  $n$ , that is, an amplitude  $n/2^{N-1}$ . The exponent  $N-1$  arises since the  $N$  bits must cover both polarities. Then an extrapolation to the zero crossing at time  $t_0$  gives

$$t_0 = t_n - \frac{1}{2\pi f_0} \frac{n}{2^{N-1}}, \tag{7}$$

where  $t_n$  is the time of the ADC measurement. A spread of ADC values between  $-0.5$  lsb and  $+0.5$  lsb will correspond to a variation in the zero-crossing time  $\delta t$  of

$$\delta t = \frac{1}{2\pi f_0 2^{N-1}}, \tag{8}$$

as shown in Fig. 4. The r.m.s. value of this is given by  $\delta t/\sqrt{12}$ , using a derivation similar to Eq. (4). A least-square fit of  $M$  measurements in effect combines  $M$  measurements in a statistically optimum way so that the timing error is reduced by a factor  $\sqrt{M}$ . However, many measurements do not contribute to the zero-crossing since they lie near a maxima or minima of  $\sin(x)$ . Assume for simplicity that  $1/2$  of these contribute. Then the standard deviation of the zero-crossing is given by

$$\sigma_{\delta t} \approx \frac{\delta t}{\sqrt{12}\sqrt{M/2}} = \frac{\sqrt{2/3}}{2\pi f_0 2^N \sqrt{M}} \approx \frac{1}{2\pi f_0 2^N \sqrt{M}}, \tag{9}$$

which is the same as Eq. (6). This derivation separates the resolution into two factors: the single-shot timing resolution  $\propto 1/2^N$  and an averaging factor of  $1/\sqrt{M}$ . Since the predicted resolution of  $\approx 100$  fs is about  $10^{-6}$  of the 100 ns period of  $f_0$ ,

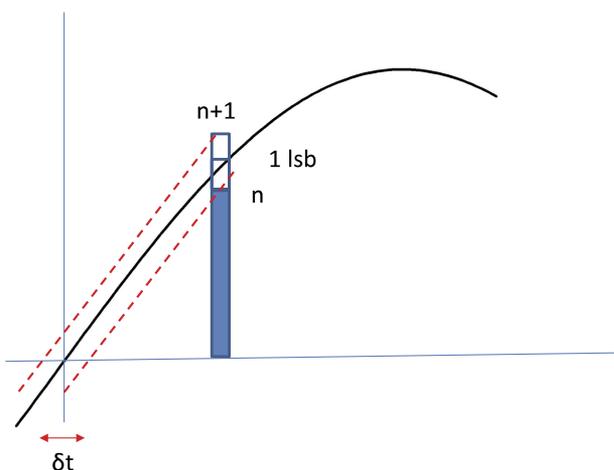


FIG. 4. Showing how the ADC amplitude resolution  $1/2^N$  translates into the timing resolution of the zero-crossing. The slope of the sine function is  $1/2\pi f_0$ . The size of 1 lsb is exaggerated for clarity.

it is important to realize that  $\approx 10^4$  of this (10 ps) is due to the ADC timing resolution and only  $\approx 10^2$  is due to averaging.

Equation (6) is a convenient approximation, but it gives only general guidance to the performance of a specific device. For example, it does not contain the sampling frequency. This is clearly incorrect in general and makes Eq. (6) insensitive to repeating samples, where  $f_0$  is an integral divisor of the sampling rate (e.g., 10 MHz and 100 MHz). Mathematically the sampling frequency disappears from Eq. (6) because terms such as

$$\frac{1}{N} \sum_{n=0}^{N-1} \sin[2\pi n f_0 + \varphi] \tag{10}$$

in the derivation<sup>18</sup> are assumed to approach 0, while in practice they stop converging once the samples repeat.

A Monte Carlo (MC) program was written to check the above approximations and to uncover other details of the design. Two sine waves were generated and digitized by truncation to  $2^N$  different levels. The data were then written to files which were read by the same program that analyzed the ADC output. The sinusoids were initially generated with a fixed phase difference, for example,  $\pi/4$  or 12.5 ns. The common phase or “start” phase of the two was then randomized over  $2\pi$  by a random number generator. This is required because the trigger of the DSO is not coherent with the signal and reference oscillators, certainly not at the ps level required. The analysis program computed the phase difference between the two fits which was subtracted from the known Monte Carlo value to give the fitting error, as shown in Figs. 5 and 6 below. The amplitude and phase noise of the sources were set to zero in the data below.

Figures 5 and 6 show the MC results as functions of the number of bits  $N$  and the number of points in the fit  $M$ . These test the  $2^N$  and  $\sqrt{M}$  factors in Eq. (6). Figure 5 shows an exponential dependence similar to Eq. (6) but with a slightly smaller slope  $\approx 2^{0.9N}$ . The MC data are about a factor of 2 larger at  $N = 16$ , which is consistent with Eq. (6) being a lower bound. This can be considered good agreement, given the approximations involved. Figure 6 shows the fitting error

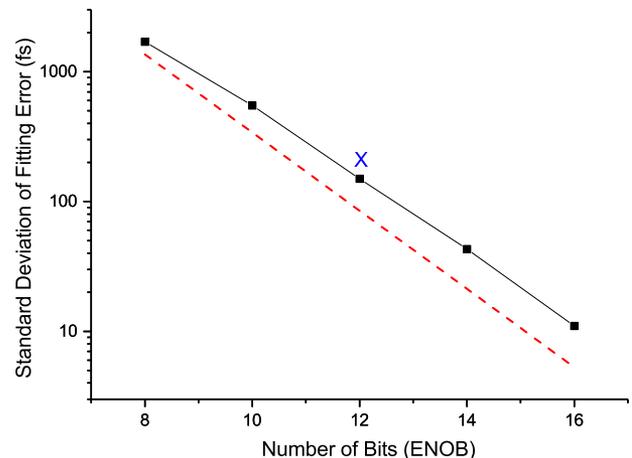


FIG. 5. Monte Carlo result (black squares) for the standard deviation of the fitting error as the effective number of bits (ENOB) is varied. This tests the  $2^N$  factor in Eq. (6).  $M$  is fixed at 4096. The analytic lower bound of Eq. (6) is shown (dotted red line). The cross marks the experimental data of Fig. 2.

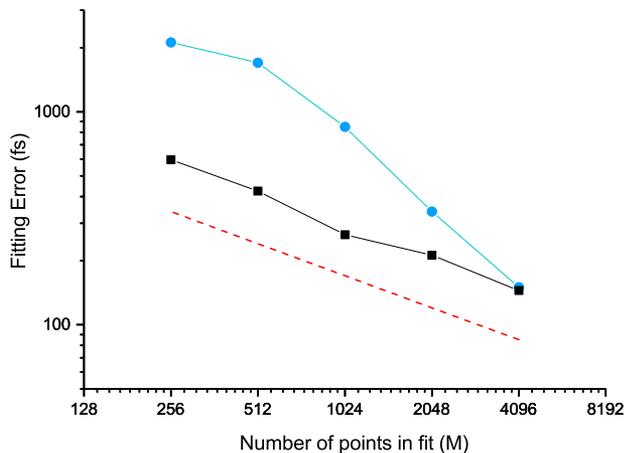


FIG. 6. Monte Carlo result for the standard deviation of the fitting error as the number of points  $M$  in the fit is varied. This tests the  $\sqrt{M}$  factor in Eq. (6).  $N$  is fixed at 12 bits. Black squares show the effect of symmetrically removing points from the fit while distributing them over the same time interval. Blue dots show results of using a shorter dataset but with a constant time between points. See text. Dotted red line is Eq. (6).

as a function of  $M$ , with  $N$  fixed at 12 bits. The two MC curves were computed by selecting two different subsets of an  $M = 4096$  point dataset. The black squares result from datasets that were reduced symmetrically without changing the duration of the fit. For example, the  $M = 256$  dataset was produced by selecting every 16-th point so that the points extended over the same  $41 \mu\text{s}$  interval as for the  $M = 4096$  case. In contrast, the blue dots result from using a shorter dataset with a constant interval between points ( $\approx 10 \text{ ns}$ ), which for the  $M = 256$  case would occupy only  $2.6 \mu\text{s}$ . At the level of approximation of Eq. (6), both results should be the same. However the symmetric case (black) follows the  $\sqrt{M}$  rule of Eq. (6) with a multiplier of  $\approx 2$ , while the short dataset case violates the  $\sqrt{M}$  behavior in this region, and is up to  $\approx 8$  times larger. An analytic explanation for this has not yet been developed. This behavior shows the limitations of the approximate theory and the necessity of careful modeling of specific devices.

Overall these three theoretical approaches confirm that the excellent time resolution observed in these experiments is a predictable result of the ADC resolution and the averaging properties of a least-square fit. Figures 5 and 6 show how these results may be applied to other ADC's and DSO's with different values of  $N$  and  $M$ .

#### IV. CONCLUSIONS

These initial experiments have been intended as a proof-of-principle of the sinusoidal fitting method. The DSO was chosen for simplicity and accessibility rather than for performance. Nevertheless the residual ADEV of  $3 \times 10^{-13}$  at  $\tau = 1 \text{ s}$  is better than many high-performance frequency standards, e.g., commercial cesium beam standards.<sup>13</sup> Figures 5 and 6 imply that substantial improvement is possible with longer datasets and higher resolution ADC's.

These results demonstrate that heterodyning is not necessary for high resolution Allan variance measurements. It is apparent that heterodyning plays a different role in DMTD

systems than it does in DD systems. In the former, it increases the resolution by orders of magnitude, while in the latter it is used as an algorithm to generate phase data from a stream of ADC outputs. In a DD system, heterodyning is done after the ADC so that the ADC time jitter and amplitude resolution at the signal frequency are preserved in the digital data. Fitting at the signal frequency is simpler and yields comparable results.

The counter-like arrangement of the data taking in Fig. 1 was dictated by the relatively short file size of this DSO (8K). Other DSO's can create much bigger files and may be optimized with different algorithms. This paper has not addressed the problem of computing the phase noise spectral density nor has it considered the use of cross-correlation techniques.<sup>5,9</sup> Algorithms using both are being considered.

The fitting method is relatively easy to implement. It does not require dedicated hardware or firmware since the DSO operates in the standard mode, writing output files when triggered by an external timer. The least-square fit is contained in an independent Python routine. Synchronizing the DSO's clock to an external source is facilitated by the clock generator chip<sup>17</sup> used in many DSO's which typically have a capture bandwidth sufficient for the several percent offset required to avoid repeated sampling.

#### ACKNOWLEDGMENTS

Thanks to Leo Hollberg for a critical reading of the manuscript.

- <sup>1</sup>D. W. Allan, *Proc. IEEE* **54**, 221 (1966).
- <sup>2</sup>E. Rubiola, *Phase Noise and Frequency Stability in Oscillators* (Cambridge University Press, Cambridge, 2009).
- <sup>3</sup>*NIST Technical Note 1337*, edited by D. B. Sullivan, D. W. Allan, D. A. Howe, and F. W. Walls (U.S. Department of Commerce, National Institutes of Standards and Technology, 1990).
- <sup>4</sup>J. Levine, *Rev. Sci. Instrum.* **70**, 2567 (1999).
- <sup>5</sup>S. Stein, *IEEE Trans. Sonics Ultrason.* **57**, 540 (2010).
- <sup>6</sup>L. S. Cutler and C. L. Searle, *Proc. IEEE* **54**, 136 (1966).
- <sup>7</sup>D. W. Allan and H. Daams, in *Proceedings of the 29th Annual Symposium on Frequency Control, 1975* (IEEE, 1975), p. 404.
- <sup>8</sup>G. Brida, *Rev. Sci. Instrum.* **73**, 2171 (2002).
- <sup>9</sup>J. Grove, J. Hein, J. Retta, P. Schweiger, W. Solbrig, and S. Stein, in *Frequency Control Symposium and Exposition, 2004, Proceeding of the 2004 IEEE International* (IEEE, 2004), p. 287.
- <sup>10</sup>J. A. Sherman and R. Joerdens, *Rev. Sci. Instrum.* **87**, 054711 (2016).
- <sup>11</sup>K. Mochizuki, M. Uchino, and T. Morikawa, *IEEE Trans. Instrum. Meas.* **56**, 1887 (2007).
- <sup>12</sup>S. Roemisch, S. R. Jefferts, and T. E. Parker, in *General Assembly and Scientific Symposium, 2011 XXXth URSI* (IEEE, 2011), p. 4.
- <sup>13</sup>Microsemi Corp., Aliso Viejo, CA 92656.
- <sup>14</sup>E. Rubiola, M. Lenczner, P. Y. Bourgeois, and F. Vernotte, *IEEE Trans. Sonics Ultrason.* **63**, 961 (2016).
- <sup>15</sup>F. Vernotte, M. Lenczner, P. Y. Bourgeois, and E. Rubiola, *IEEE Trans. Sonics Ultrason.* **63**, 611 (2016).
- <sup>16</sup>C. E. Calosso, in *Proceedings of the 2013 IEEE IFCS* (IEEE, 2013), p. 747.
- <sup>17</sup>Digilent Corp., Analog Discovery Technical Reference Manual, 2013, Pullman, WA 99163.
- <sup>18</sup>S. M. Kay, *Fundamentals of Statistical Signal Processing* (Prentice Hall, 1993).
- <sup>19</sup>I. Kollar and J. J. Blair, *IEEE Trans. Instrum. Meas.* **54**, 1978 (2005).
- <sup>20</sup>IEEE Standard 1241-2010, IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters.
- <sup>21</sup>B. Widrow and I. Kollar, *Quantization Noise* (Cambridge, 2008).
- <sup>22</sup>Analog Devices, AD 9648 Datasheet, Norwood, MA 02062.