CONFIDENCE ON THE SECOND DIFFERENCE ESTIMATION OF FREQUENCY DRIFT
A Study based on Simulation

Marc A. Weiss, David W. Allan, and David A. Howe
Time and Frequency Division
National Institute of Standards and Technology
325 Broadway
Boulder, CO 80302

Abstract

We use simulation to compare the confidence in estimating frequency drift in the presence of stochastic noise between two different estimators: (1) a mean second difference using neighboring data, and (2) a single second difference over the entire data set. In each case we simulate 100 s data sampling and 10,000 samples (11 and 1/2 days of data). The two estimators show a similar confidence when the noise is purely random walk frequency modulation, but when there is also white frequency modulation, as is common in oscillators, method (2) is more efficient. An advantage of (1) is that it provides an internal confidence to an estimate on a single data set. As a practical example, when we simulate a typical rubidium gas cell frequency standard, the confidence of the drift estimate was ten times better for method (2) over method (1).

Theory

We studied the confidence on two second difference estimators of frequency drift in the presence of white noise frequency modulation (WHFM) and of random walk frequency modulation (RWFM) using simulation methods in Monte Carlo runs. We generated time series simulating clock data (10,000 points) taken every 100 s against a perfect reference, continuing for 10^6 s, about 11 and 1/2 days. In each case we repeated the simulation 100 times, that is, we generated 100 different data sets with the same stochastic and deterministic parameters but with a different starting seed number. We repeated our two different second difference estimators on each data set, taking the mean and standard deviation across all data for a given set of parameters. In this paper we refer to a "knee" in the Allan variance plot of $\sigma_y(\tau)$

versus τ. The knee is the τ value where we see a transition from a region where one power law dominates, such as $\sigma_y(\tau) \sim \tau^{1/2}$ for RWFM, to the region where another process dominates, such as $\sigma_y(\tau) \cong \tau^1$ for frequency drift.

We simulated clock data for two types of experiments. The first type may be typical of some quartz crystal oscillators. In this case, there was a frequency drift of 1 part in 10^{10} per day, and the stochastic noise was purely RWFM. This was done using four different levels of RWFM to make the knee in the Allan variance curve appear at 10^3 s, 10^4 s, 10^5 s, and equivalent of 10^6 s, if we had generated more data. This allowed us to study the ability of the second difference estimator to find the true drift when the random walk effects dominated increasingly more of the data set. In a second type of experiment we introduced a combination of WHFM and RWFM with frequency drifts of either 1 part in 10^{14} per day or 1 part in 10^{13} per day, modelling the performance of two types of rubidium gas-cell frequency standards.

The second difference operator, Δ^2, estimates drift by computing the change in average frequency from one interval to the next. The second difference was computed in two ways: (1) a mean second difference using neighboring data, and (2) a single second difference over the entire data set. For (1), the mean second difference of a data set, we take first differences over the minimum time interval, τ_0, the average frequencies over each τ_0, and then difference them again to obtain an estimate of drift, the change in frequency, for the interval τ_0. We then average these second differences over the entire data set. In this case the individual estimates of average frequency which are used to estimate drift are poor in the presence of the random walk frequency process, compared to those for a longer averaging time as in (2). But the second difference operator converts the RWFM process to a white noise process, hence allowing averaging of the individual drift estimates in

method (1), since in our simulation, the drift remains constant throughout the data length. This compares with method (2) where we take the average frequency over the first half of the data set and subtract it from the average frequency over the second half. In this case we have a better estimate of the average frequency, thus a better estimate of drift in the single second difference. It turns out from our simulation that methods (1) and (2) provide similar confidence on the estimates in the case of pure RWFM. However in real oscillators we usually find stochastic noise which is a mixture of RWFM as well as WHFM or what is called flicker frequency modulation. In these more realistic situations, method (2) is significantly more efficient as a drift estimator. We will demonstrate the increased efficiency if WHFM is present. An advantage of method (1) is that it provides an internal confidence of the estimate, the standard deviation of the mean of the series of second differences.

If WHFM noise is included with RWFM and drift, then the application of the second-difference operator Δ^2 yields a non-white process. In a WHFM process the frequency is a white noise process of deviations around a mean frequency. Since a WHFM process is equivalent to a random walk phase modulation (RWPM), then the first difference of phase of WHFM data yields a white noise process. The longer we average white noise, the better the estimate of the mean. This means, the larger the value of τ we use to compute first differences, as long as the dominant noise process is WHFM, the better the estimate of mean frequency. As the interval τ increases for the Δ^2 operator, the estimate of average frequency $\overline{y}$ improves because high-frequency energy contained in the WHFM is "filtered" by a greater extent. In the extreme case in which τ is 1/2 the length of the full data set and there is only a single Δ^2 to compute (i.e., three points as shown in figure (1), then the high-frequency energy is maximally filtered. Hence the knowledge of the frequency component of the data set is the best and is better than computing $\overline{\Delta^2}$ over shorter τ intervals. Though as τ increases, the effects of RWFM and then drift become more dominant, yet this can explain why method (2), using a single over-all second difference, is more practically efficient than method (1).

From a variance point of view, the variance of frequency is largest at τ_0, the minimum data spacing, and decreases as τ. This is because with white noise, data are random uncorrelated. From a frequency-domain viewpoint, the spectrum resulting from applying Δ^2 to WHFM phase data yields a non-white process -- having increased power at high Fourier frequencies compared to low Fourier frequencies.

We will consider a few subtleties of Δ^2 in the presence of RWFM. The Δ^2 operator applies a sequential "difference-of-slopes" operation to a data set. If the set is the time difference of two clocks (or phase error data) from, say, a time-interval counter, then the Δ^2 operator represents the "difference-of-frequencies" (since the slope of the phase is proportional to frequency) determined from three sequential time difference readings. Subsequent three-somes can be taken and the Δ^2 operator applied to each; an average over the full data set can then be computed as Barnes has discussed [1]. The Δ^2 operator approximates the familiar $\dfrac{d^2}{dt^2}$ operator used in a continuous function.

To say that noise on a signal is consistent with a model of RWFM has several implications. We assume that there is a continuous process which is the double integral of a white noise process. Any measurements on the signal, however, are discrete band-limited measurements. Thus, if we doubly difference the discrete measurements, we might not obtain the theoretical underlying white noise process. There are at least three reasons for this: (1) initial data points are lost for each difference operation, i.e., there are N-1 differences for N data samples; (2) the functions involved are sampled at τ_0 thus the data always include power at the sampling frequency of $1/\tau_0$ which has an associated window function, and hence has power introduced within the bandwidth of our operators; and (3) as shown by Greenhall [2], adjacent sets of data can be highly correlated in RWFM, yielding a bias when double differencing. These problems can be mollified. If the initial data points on differencing are maintained, then item (1) is not a concern. Regarding item (2), the sampling window function can be shaped to reduce sidelobe leakage [3]. The correlation effect of item (3) may still cause the second difference of the sampled data to yield a different white-noise process than the theoretical, continuous process. As an aside, the addition of the drift term to the RWFM can perturb the estimate of WHPM using the Δ^2 operator, since it is added and subsequently subtracted at different steps in the process. This problem can also be resolved. Thus, with proper care we can find that double differencing yields a white process, though we

301

lose some information about the theoretical underlying continuous process.

Experiment

Measurements are commonly taken by a time-interval counter, producing a sampled, usually equally spaced, time series with interval τ_0. The WHFM and RWFM models are invariably useful for studying the performance of Cs, Rb, Qz, H, and other frequency standards [4]. Optimal estimation involves converting the noise term to a white process so that a simple average and standard deviation can be computed. The $\overline{\Delta^2}$ operator will do this to pure RWFM. This paper shows that the result is essentially the same for pure RWFM plus drift if we apply Δ^2 to each data three-some (at interval τ_0) or if we compute Δ^2 by simply taking the first, middle, and last points as shown in figure 1.

converts white noise into a model of the phase for RWFM. These noises are then added to a quadratic term whose slope increase corresponds to a frequency drift ranging from 10^{-14} per day to 10^{-10} per day to generate the final data set. The analysis is then applied to these data to compute the drift and confidence. Since we know the drift a priori, we can test the accuracy of the drift and confidence estimation.

A family of plots of $\sigma_y(\tau)$ showing (theoretically) four levels of RWFM and drift of 10^{-10} per day are shown in figure 2. RWFM and drift go as $\tau^{+\frac{1}{2}}$ and τ^{+1} respectively. The "knee" (or change of slope in the $\sigma_y(\tau)$ plot) occurs at 10^3, 10^4, and 10^5 s. We note that in the highest level of RWFM used in our simulation, the knee occurs at a sample time equal to the length of the data set, 10^6 s; so in this case, the RWFM essentially masks the drift.
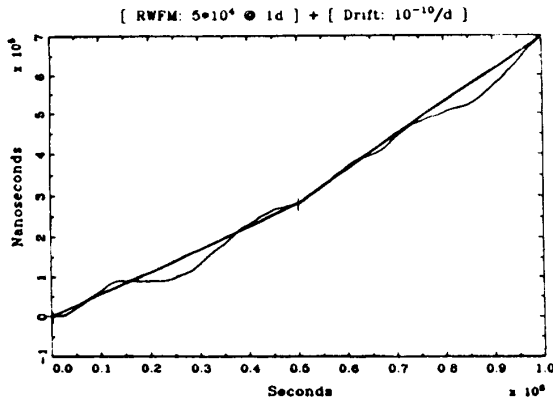

Figure 1. This figure illustrates the three points used in a simple second difference for estimating the frequency drift. The random line is simulated RWFM super-imposed on a linear frequency drift. The two straight lines indicate the slopes corresponding to the two frequencies from which the frequency drift is estimated.

A Monte Carlo method was used to test the concept on a set of simulated data. The simulated data consist of drift equal to 10^{-10} per day with RWFM added at four levels of amplitude. The RWFM data were generated by twice integrating band-limited white noise. The first integration converts white noise into random walk noise. This output can be used to model RWPM, which is the same as WHFM. Using a different "seed" and two integrations
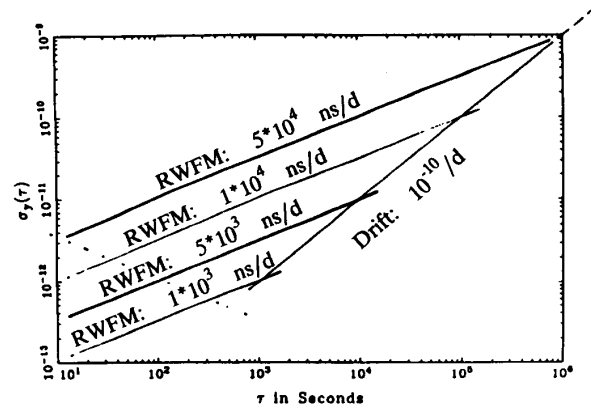

Figure 2. An illustration of a frequency stability plot, $\sigma_y(\tau)$, for different simulated levels of RWFM and a fixed frequency drift of 1 part in 10^{10} per day as is typical for some precision quartz-crystal oscillators. Notice the intersections "the knee" between the RWFM curves and the frequency drift curve occur at 1000 s, 10,000 s, 100,000 s, and at 1,000,000 s.

Also shown in figure 2 is white frequency modulation (WHFM) (the single dotted line with $\tau^{-\frac{1}{2}}$ behavior) at a level common to high-performance commercial cesium-beam frequency standards and to rubidium gas-cell frequency standards. Note, however, that this much frequency drift would be unusually poor performance for a cesium standard. Depending on the level of RWFM, we may or may not see its $\tau^{+\frac{1}{2}}$ behavior if masked by the WHFM plus drift. If there is no RWFM and we have pure WHFM and drift, then Δ^2 applied to the phase data will not convert the

302

WHFM to a white process. In this case, the residuals around a linear regression to the frequency are white, and this regression line is the appropriate estimator of drift. Note, WHFM is the same as RWPM, the first integral of a WHPM process.

Figure 3 is a model with a WHFM level of 10 ns and a RWFM of 30 ns/d, both at 1 d, and a drift of 10^{-13} per day. The last two numbers are typical for commercial rubidium gas-cell frequency standard and the first number for traditional cesium-beam frequency standard. Though not plotted, we also did the simulation for the anticipated performance of the rubidium frequency standards that are planned for the GPS Block 2-R satellites to be used in cross-link ranging for that system. The model elements in this second case were 3 ns/d for the WHFM, 4 ns/d at 1 d for the RWFM and 10^{-14} per day for the frequency drift.
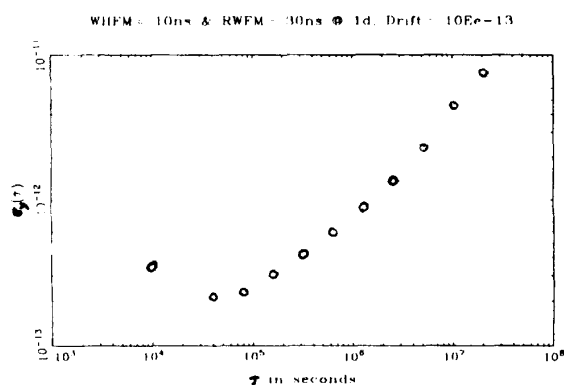
WHFM 10ns & RWFM 30ns @ 1d, Drift 10Ee-13

Figure 3. An illustration of a frequency stability plot, $\sigma_y(\tau)$, for simulated data as may be typical in some atomic frequency standards. The noise model includes both WHFM and RWFM. The frequency drift model is typical of commercial rubidium frequency standards.

Results

Results of using the simulated data and computing a single second difference (first, mid, and last points only) are shown in Table I. One can imagine the data representing four different experiments each involving 100 independent clocks and each clock having the same level of RWFM (and value of drift, of course). The third column is the average of each of the 100 cases; the results closely match the expected value of 10^{-10} per day with a +18% error being the worst in the last (highest level) case of

RWFM. The fourth column represents the standard deviation over each 100 cases.

TABLE I
100 Single 2nd Difference Estimates of
10^{-10}/d Drift + RWFM

RWFM @ 1d	"Knee" [s]	Mean of 2d's	S.D. of 2D's
1000 ns/d	10^3	$1.00*10^{-10}$	$.004*10^{-10}$
5000 ns/d	10^4	$1.00*10^{-10}$	$.190*10^{-10}$
10000 ns/d	10^5	$1.02*10^{-10}$	$.344*10^{-10}$
50000 ns/d	10^6	$1.18*10^{-10}$	$2.07*10^{-10}$

Table II shows the results of computing the mean second differences using the same original data. The average of the mean second differences over each of the 100 cases is shown in column three. Again the worst match is in the last (highest level) case of RWFM being in error by +12%. The fourth column shows the standard deviations of the mean second differences used in column three and is a measure of confidence on the result. The fifth column is the RMS of the standard deviations of the means and these are essentially identical to the standard deviations because the Δ^2 operator yields a pure white process from the pure RWFM data set.

TABLE II
100 Mean 2nd Difference Estimates of
10^{-10}/d Drift + RWFM

RWFM (1d)	"Knee" [s]	Mean of M2d	S.D. of M2D	RMS(SD Mean)
1000 ns/d	10^3	$1.00*10^{-10}$	$0.035*10^{-10}$	$0.034*10^{-10}$
5000 ns/d	10^4	$0.99*10^{-10}$	$0.173*10^{-10}$	$0.170*10^{-10}$
10000 ns/d	10^5	$1.00*10^{-10}$	$0.333*10^{-10}$	$0.340*10^{-10}$
50000 ns/d	10^6	$1.12*10^{-10}$	$1.670*10^{-10}$	$1.670*10^{-10}$

A typical plot of the simulated phase data with the "knee" at 10^4 s is shown in figure 4. We see the obvious quadratic trend and the added slight RWFM noise accompanying it. The corresponding $\sigma_y(\tau)$ plot is shown in figure 5. Here we see the $\tau^{+\frac{1}{2}}$ RWFM behavior over most of the plot and then the curvature in long-term toward τ^{+1} frequency drift. It is interesting to note that even with the "knee" being comparable to the data length, the frequency drift is still reasonably measurable.

303

For the typical rubidium frequency standard data (see figure 3) where the modeled drift was 10^{-13} per day, we obtained $(0.61 \pm 0.42) \times 10^{-13}$ per day for the

$$[\text{ RWFM: } 1\cdot10^4 \;\bullet\; 1d \;] \; + \; [\text{ Drift: } 10^{-10}/d \;]$$
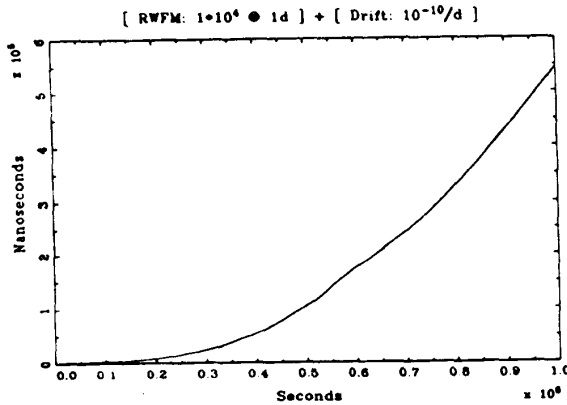

Figure 4. A typical plot of the simulated phase data with the "knee" at 10^4 s. We see the obvious quadratic trend and the added slight RWFM noise accompanying it.

$$[\text{ RWFM: } 1\cdot10^4 \;\bullet\; 1d \;] \; + \; [\text{ Drift: } 10^{-10}/d \;]$$
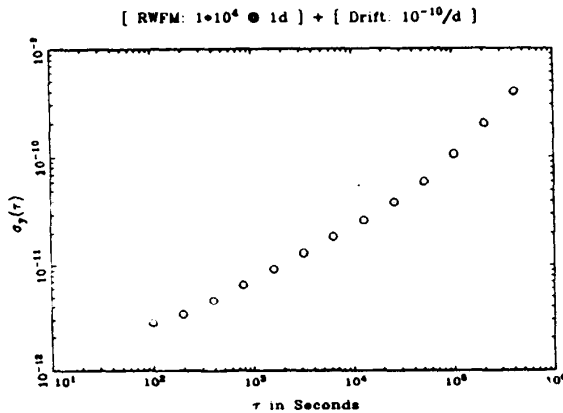

Figure 5. This frequency stability plot, $\sigma_y(\tau)$, corresponds to the data shown in figure 4. One sees the $\tau^{1/2}$ (RWFM) behavior transitioning into the τ^{+1} (frequency drift) behavior beyond 10^4 s.

external estimate of the confidence using method (1) and $(0.91 \pm 0.12) \times 10^{-13}$ per day for the estimate using method (2). For the improved performance rubidium planned for the GPS Block 2-R with a modeled drift of 10^{-14} per day we estimated $(0.3 \pm 1.4) \times 10^{-14}$ per day from the internal estimate of the confidence using method (1). The external estimate of the confidence gave a very large uncertainty. For

method (2) the frequency drift estimate was $(0.88 \pm 0.17) \times 10^{-14}$ per day.

Conclusions

We have compared two different estimators of drift in the presence of pure RWFM, and the more realistic combination of WHFM and RWFM. We find with pure RWFM that even though one is computationally simpler, they both produce estimates of drift with comparable confidence. Thus, the drift can be estimated with confidence using only three points from this data set: the first, the middle, and the last points. The real advantage of this simple second-difference frequency drift estimator appears when a significant level of WHFM is present. Then the confidence on this estimate becomes significantly better than computing second differences over τ_0, then computing the mean of these. We note that similar simulation would show this second difference estimator giving significantly better estimates of drift with WHFM and RWFM stochastic noise, than linear regression on the frequency or quadratic regression on the phase.

In general, taking the results of this paper, which are based on pure simulation and modeling, along with the results of Barnes and Allan (see ref. [1] and [4]) which are based on real data, we have the following conclusions concerning the estimation of frequency drift. If the predominate noise process is white noise phase modulation (WHPM), then the frequency drift can be best determined by calculating a quadratic-least-squares to the phase (or the time) data values. If the predominate noise process in white noise frequency modulation (WHFM), then the frequency drift can be best determined by a linear-least-squares to the frequency data values. Otherwise, the simple second difference, determined from the first, middle and last phase (or time) data values will be a more efficient estimator of frequency drift for typical frequency standards.

References

[1] J. A. Barnes, "The Measurement of Linear Frequency Drift in Oscillators," in NIST Technical Note 1337, pp. 264-295, 1990.

[2] C. Greenhall, "A Structure Function Representation Theorem with Application to Frequency Stability Estimation," IEEE Trans. Instrum. Meas., vol. IM-32, no. 2, pp.364-370, 1983.

[3] D.A. Howe, D.W. Allan, J.A. Barnes, "Properties of Signal Sources and Measurement Methods," in NIST Technical Note 1337, pp. 14-60, 1990.

[4] D.W. Allan, "Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators," in NIST Technical Note 1337, pp. 121-128, 1990.