CONFIDENCE ON THE THREE-POINT ESTIMATOR OF FREQUENCY DRIFT*

Marc A. Weiss and Christine Hackman
Time and Frequency Division
National Institute of Standards and Technology
325 Broadway
Boulder, CO 80303

Abstract

It has been shown that a three–point second difference estimator is nearly optimal for estimating frequency drift in many common atomic oscillators. We derive a formula for the uncertainty of this estimate as a function of the integration time and of the Allan variance associated with this integration time.

Theory

The three–point drift estimator is a useful tool for estimating the frequency drift in many atomic oscillators [1]. In this paper we derive a formula for the uncertainty of the three–point drift estimate; as we shall demonstrate, there is a simple relationship between the uncertainty of the drift estimate and the Allan variance of the residuals which remain after the estimated drift is removed. We explain how to apply the uncertainty formula and then we use it to assess the uncertainty of the drift estimate in several examples.

Let us begin by discussing the three–point drift estimator. To define it, let x(t) be a time series of time difference measurements between two oscillators drifting in frequency relative to each other. An optimal estimator, $\hat{D}$, of drift uses the first, middle, and last time–difference points. We estimate the average frequency over the first and second halves of the data, subtract the first frequency from the second, and then divide by τ, the time elapsed between the first and middle or middle and last data points. This yields:

$$\begin{aligned} \hat{D} &= \frac{1}{\tau}\left(\frac{x(2\tau)-x(\tau)}{\tau} - \frac{x(\tau)-x(0)}{\tau}\right) \\ &= \frac{1}{\tau^2}\left((x(2\tau)-2x(\tau)+x(0)\right) \end{aligned} \tag{1}$$

That is, we estimate drift as $1/\tau^2$ times the second difference of the time series x, where we take the second difference over as large an interval as possible.

Let us separate the time offset $x(t)$ into the part due to the frequency drift D and the part due to everything else (initial offsets, stochastic noise, systematics):

$$x(t) = x'(t) \frac{D}{2} t^2 \ . \tag{2}$$

We now will show that the uncertainty of the drift estimate, $\hat{D}$, is functionally related to the Allan variance of the $x'(t)$ time series.

If we substitute (2) into (1) we obtain

$$\hat{D} = \frac{1}{\tau^2} \left(x'(2\tau) - 2x'(t) + x'(0) + D\tau^2 \right) \ . \tag{3}$$

Rearrangement yields:

$$\hat{D} - D = \frac{1}{\tau^2} \left(x'(2\tau) - 2x'(\tau) + x'(0) \right) \ . \tag{4}$$

The expected variance of our drift estimate, $\hat{D}$, around the true drift D will thus be

$$\langle \left(\hat{D} - D \right)^2 \rangle = \frac{1}{\tau^4} \langle \left(x'(2\tau) - 2x'(\tau) + x'(0) \right)^2 \rangle \ , \tag{5}$$

where $\langle \ \rangle$ is the expectation operator. The square root of this quantity is the expected deviation of $\hat{D}$ around the true value.

If we compute an Allan variance of x' for the integration time τ we obtain [2,3]

$$\sigma_{y'}(\tau)^2 = \frac{1}{2\tau^2} \langle \left(x'(2\tau) - 2x'(\tau) + x'(0) \right)^2 \rangle \ . \tag{6}$$

Substitution of (6) into (5) yields our result, the relationship between the expected deviation in the drift estimator, $\hat{D}$, and the Allan variance of x', the drift–removed data:

$$\boxed{\langle \left(\hat{D} - D \right)^2 \rangle = \frac{2}{\tau^2} \sigma_{y'}(\tau) \ ,} \tag{7}$$

where $\sigma_{y'}(\tau)$ is the Allan variance of the x' data, and y' refers to the frequency data derived from x'. Thus we see that the uncertainty in our drift estimate is a function of the Allan variance of the drift–removed data.

Application of Equation 7

The application of (7) requires a bit of finesse. First of all, the alert reader has probably noticed that, since we don't know the value of D, the true drift, we cannot obtain the time series $x'(t)$. To circumvent this problem, we obtain an approximation of $x'(t)$ by removing the estimated drift from the $x(t)$ series. We then compute the Allan variances for the approximate $x'(t)$ series. However, it is at

this point that we encounter another problem: It is generally true that if you 1) use a second–difference estimator (such as the three–point estimator) to estimate drift, 2) remove this estimated drift from the time series, and then 3) compute the Allan variances for the residual time series, the Allan variances obtained for large integration times (such as $\tau = 1/2$ the data length) will be biased low, i.e. the Allan variance will not be an accurate measure of the frequency variability at large integration times. In fact, if we were to take a data set with constant drift, compute $\hat{D}$ using (1), remove $\hat{D}t^2/2$ from each data point, and then compute $\sigma_{y'}(\tau)$ for this same τ, we would obtain exactly zero.

We need to have $\sigma_{y'}(\tau)$ for $= 1/2$ the data length in order to use (7). Yet we know that after removing $\hat{D}$ from $x(t)$, we are going to get the incorrect value of 0 for $\sigma_{y'}2(\tau)$ for $\tau = 1/2$ the data length. However, while $\sigma^2_{y'}(\tau)$ is incorrectly low for large τ, it does accurately represent the frequency variability for smaller τ. Furthermore, the noise processes of atomic oscillators are such that, for a given range of integration times τ, it is usually the case that $\sigma^2_{y'}(\tau) = k\tau^n$, where k is a constant and n is an integer ranging from 1 to -2. The result of this power–law behavior of $\sigma_{y'}(\tau)$ is that log–log plots of $\sigma_{y'}(\tau)$ versus τ exhibit linear behavior. This can be seen in Figure 1. Therefore, in order to obtain $\sigma^2_{y'}(\tau)$ for $\tau = 1/2$ the data length, we look at the log–log plot of $\sigma_{y'}(\tau))$ versus τ and discard the incorrectly–low values of $\sigma_{y'}(\tau)$ which occur at large τ (For example, in Figure 1, we would discard the point for which log τ (seconds) ≈ 7. In Figure 3 we would discard the point for which log τ (seconds) ≈ 6.75). Then, we use the $\sigma_{y'})\tau)$ points which correspond to the largest remaining τ values to determine k and n (i.e., we determine the equation of the line on the log–log plot formed by the remaining valid data points). Then, knowing k and n, we use the equation $\sigma^2_{y'}(\tau) = k\tau n$ to determine the value of $\sigma^2_{y'}(\tau)$ at $\tau = 1/2$ the data length. This value is what we need to apply (7).

For cesium beam and rubidium gas–cell oscillators, the dominant noise types at large integration times are flicker frequency modulation and random walk frequency modulation (FLFM and RWFM, respectively). FLFM corresponds to an n value of 0 and RWFM corresponds to an n value of $+1$. For very large τ, RWFM generally dominates. Therefore, if the last (i.e. largest τ) valid linear trend that we see on the log–log plot is consistent with a model of RWFM, we may use this slope with a measure of confidence to estimate the value of $\sigma^2_{y'}(\tau)$ at $\tau = 1/2$ the data length. If, however, the last linear trend corresponds to FLFM, we need to ask ourselves whether the FLFM noise type continues out to $\tau = 1/2$ the data length, or whether RWFM is the correct noise type for $\tau = 1/2$ the data length. The assumption of RWFM as the noise type always leads to a larger computed value of $\sigma^2_{y'}(\tau)$ than the assumption of FLFM. Thus, simply assuming that RWFM dominates at $\tau = 1/2$ the data length yields a conservative estimate. The uncertainty in the Allan variance estimate will limit the accuracy of our uncertainty estimate. Nevertheless, we can make conservative estimates of uncertainty and obtain meaningful results.

In summary, to use (7) to estimate the uncertainty of we take the following steps:

1. Compute using the second difference estimator (3), where in that equation, $\tau = \tau_{max}$, the time interval for one–half the data length. Remove $\hat{D}t^2/2$ from each of the time–difference data points $x(t)$.

2. Compute the Allan deviations $\sigma_{y'}(\tau)$, for $\tau = n\tau_0$, where n is an integer multiple of the sampling interval τ_0. Make a log–log plot of $\sigma_{y'}(\tau)$ versus τ.

3. Look for abnormally low values of $\sigma_{y'}(\tau)$ at large values of τ. Discard them.

4. Determine the parameters k and n in the equation $\sigma^2_{y'}(\tau) = k\tau n$ for the last valid linear trend

on the log–log plot. Then use this equation to compute $\sigma_y^2(\tau)$ for τ_{max}. Remember to consider the possibility that the noise type might change past the last valid $\sigma_y(\tau)$ value on the log–log plot (i.e., the noise type might change from FLFM to RWFM).

5. Substitute this value of $\sigma_y^2(\tau)$ into (7). Solve (7) for the variance of $\hat{D}$. The square root is the expected deviation.

Examples

As examples we use atomic standards aboard GPS satellites studied from July 1, 1991, to September 15, 1992, a period of 443 d. Satellites are referred to by their pseudo–random code number (PRN), the number by which users identify satellites , or by their satellite vehicle number (SVN), the number used by the GPS control segment. Clocks on the GPS satellites are measured at NIST against the AT1 time scale. For clocks which ran for this entire period, drift could be estimated using a second difference with τ=221.5 d. Not all clocks analyzed were on line for this entire period, in which case shorter τ values were found. We found an assortment of dominant noise types at various integration times, with FLFM and RWFM dominating at times equal to one–half the data length. Table I gives our example results and indicates associated figure numbers.

PRN#2 and figure 1 illustrate the difficulty in determining noise type. Looking at figure 1, we see that, while FLFM, τ^0, is the probable slope for the last valid $\sigma_y(\tau)$ values, the uncertainty allows for the possibility of a $\tau^{1/2}$ slope, indicating RWFM. Furthermore, RWFM is usually the dominant noise process for cesium frequency standards at integration times such as 221.5 d [6]. We compute a more conservative value in the second line of the table. Similarly for PRN#25 we have assumed FLFM in its first line. If we assume RWFM we see we find only a small change.

Another consideration is that equation (7) applies to the Allan variance, not the modified Allan variance. In figures 1 and 5 we used the modified Allan variance. We can account for this as follows. Asymptotically, if we define

$$R_\infty = \lim_{\tau \to \infty} \frac{\text{mod } \sigma_y^2(\tau)}{\sigma_y^2(\tau)}, \tag{8}$$

then $R_\infty = 0.91$ for RWFM and $R_\infty = 0.82$ for FLFM [5]. These corrections have been included in the table.

Conclusions

We have derived a relationship that allows us to estimate the uncertainty of the three–point estimator of frequency drift. It does not give a lot of precision but it is adequate for determining a confidence level. With the procedure outlined, we can determine an upper bound on the uncertainty of the estimate of frequency drift.

References

[1] M.A. Weiss, D.W. Allan, and D.A. Howe, *"Confidence on the Second Difference Estimation of Frequency Drift, A Study Based on Simulation,"* Proceedings of the 1992 IEEE Frequency Control Symposium, pp. 300–305, 1992.

[2] D.A. Howe, D.W. Allan, and J.A. Barnes, *"Properties of Signal Sources and Measurement Methods,"* 35th Annual Symposium on Frequency Control, 1981.

[3] Eds. D.B. Sullivan, D.W. Allan, D.A. Howe, and F.L. Walls, NIST Tech Note 1337: Characterization of Clocks and Oscillators, 1990.

[4] S.R. Stein, *"Frequency and Time–Their Measurement and Characterization,"* Precision Frequency Control, vol. 2, chap. 12, pp. 191–232, 1985.

[5] D. B. Sullivan, D.W.Allan, D.A. Howe, and F.L. Walls *"Characterization of Clocks and Oscillators,"* in [3], pp.1–13.

[6] D.W. Allan, *"Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators,"* in [3], pp. 121–128.

Table 1

PRN# (Type)	Dates	τ_{max} days	Dominant Noise Type	$\sigma_y(\tau)$ @ τ(s)	Estimated Drift $\pm$ σ parts in 10^{-15}/d	Figure Nos.
2 (Cs)	1Jul91 - 15Sep92	221.5	FLFM	$0.4 \cdot 10^{-13}$	-2.7 $\pm$0.3	1
2 (Cs)	1Jul91 - 15Sep92	221.5	RWFM	$0.2 \cdot 10^{-13}$ @ 10^6	-2.7 $\pm$0.6	1
3 (Rb)	1Jul91 - 15Sep92	221.5	RWFM	$2.0 \cdot 10^{-13}$ @ 10^6	-98 $\pm$6	2
12 (Rb)	8Apr - 15Sep92	80.5	RWFM	$2.5 \cdot 10^{-13}$ @ 10^6	-130 $\pm$10	3
19 (Cs)	1Jul - 18Dec91	85.5	RWFM	$1.2 \cdot 10^{-13}$ @ 10^6	35 $\pm$5	4
25 (Rb)	30Jun - 15Sep92	39	FLFM	$0.7 \cdot 10^{-13}$	-183 $\pm$3	5
25 (Rb)	30Jun - 15Sep92	39	RWFM	$0.6 \cdot 10^{-13}$ @ 10^6	-183 $\pm$4	5

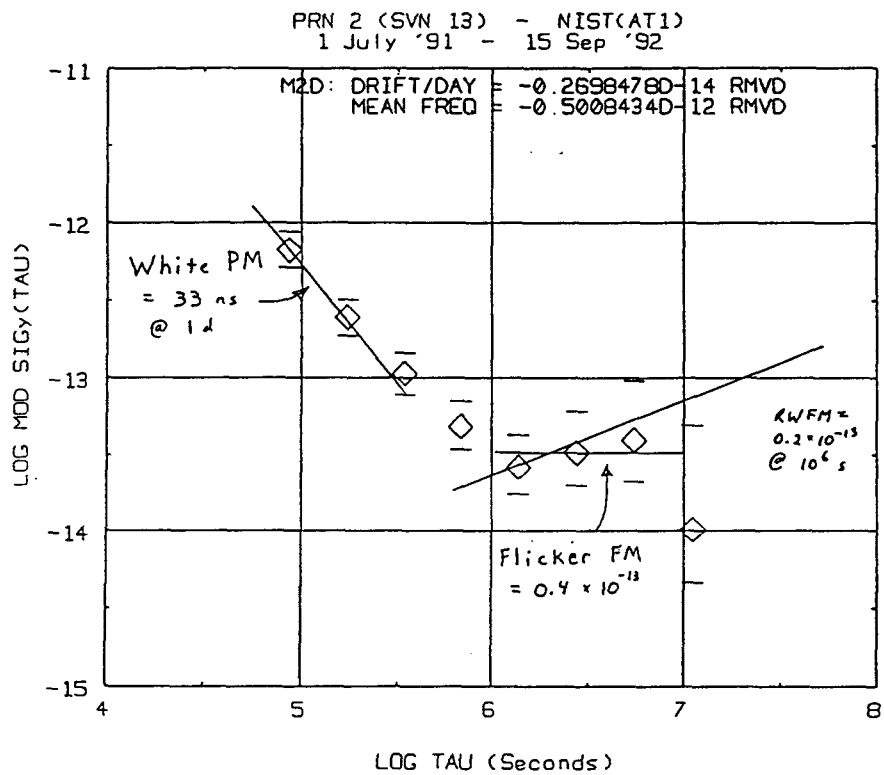

Figure 1: The modified Allan variance of the Cs clock on PRN#2 as measured at NIST against the AT1 time scale from July 1, 1991 to September 15, 1992.

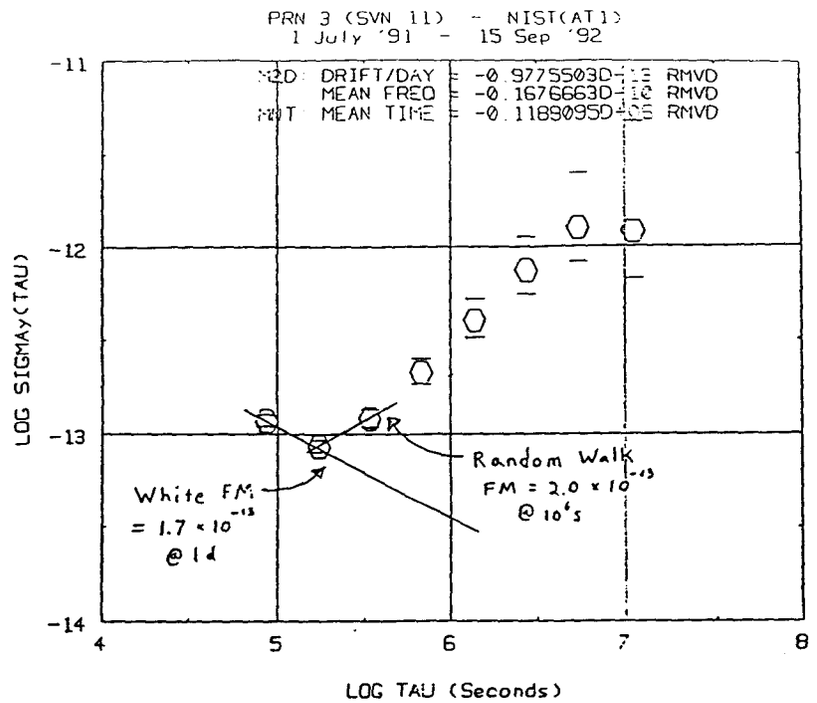

Figure 2: The Allan variance of the Cs clock on PRN#3 as measured at NIST against the AT1 time scale from July 1, 1991 to September 15, 1992.

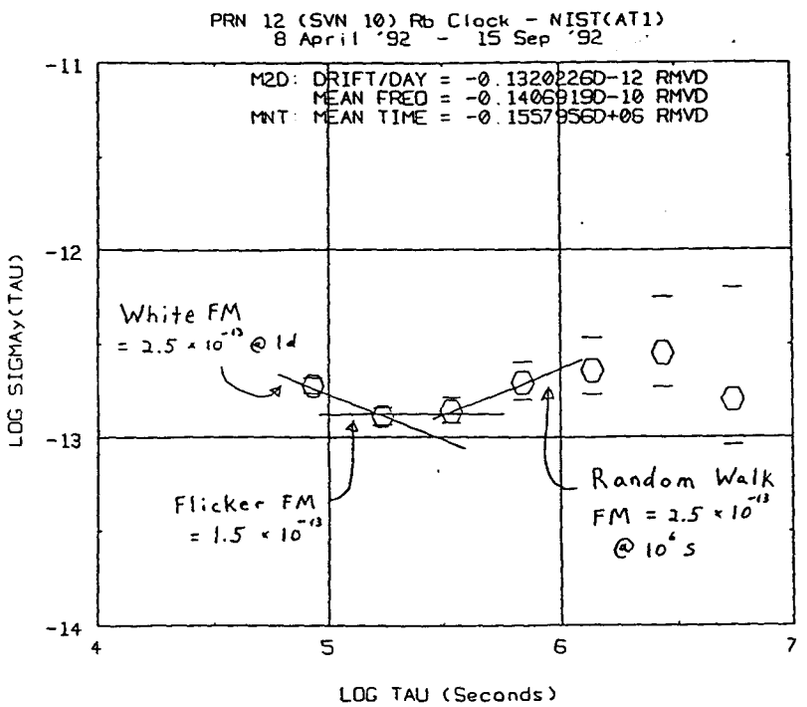

Figure 3: The Allan variance of the Rb clock on PRN#12 as measured at NIST against the AT1 time scale from April 8, 1991 to September 15, 1992.

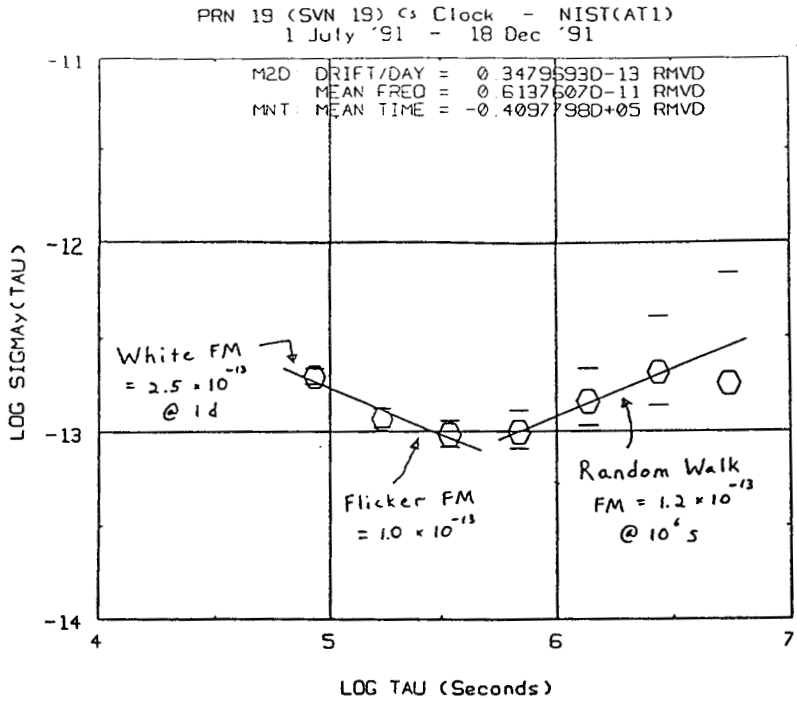

Figure 4: The Allan variance of the Cs clock on PRN#19 as measured at NIST against the AT1 time scale from April 8, 1991 to September 15, 1992.

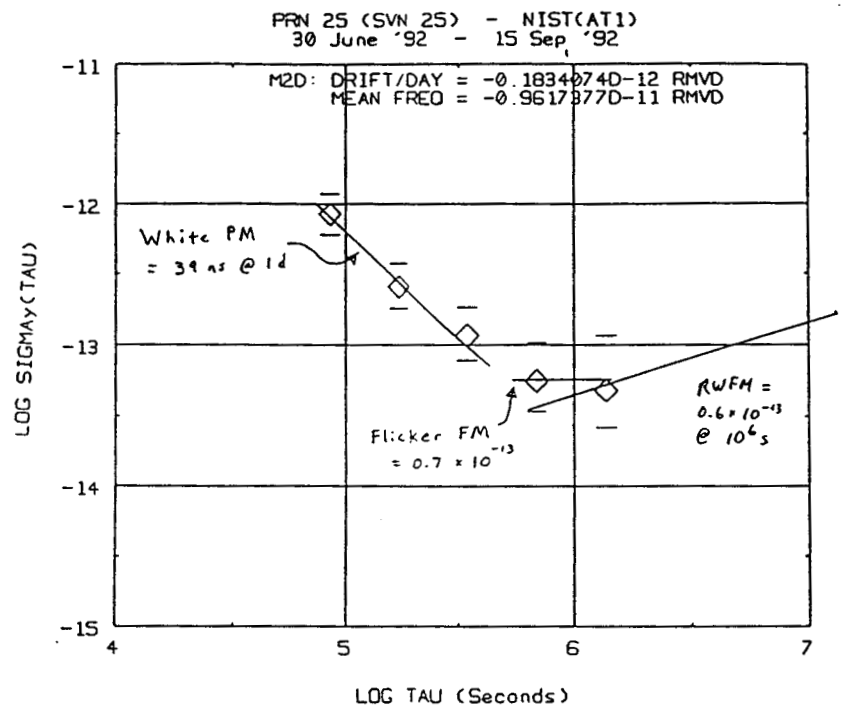

Figure 5: The Allan variance of the Rb clock on PRN#25 as measured at NIST against the AT1 time scale from April 8, 1991 to September 15, 1992.