

Introduction to the Time Domain Characterization of Frequency Standards*

James Jespersen
Time & Frequency Division
National Institute of Standards & Technology
Boulder, Colorado 80303

Abstract

The simple mean, the standard deviation and its square, the variance are well known statistical measures of a set of data points. It seems natural, therefore, to apply these measures to signals generated by frequency standards. However, this approach quickly reveals some significant problems.

Normally, when we compute the mean and standard deviation of some process, we assume that including more data points in the computation brings us ever closer to the true mean and standard deviation of the process. Unfortunately this is not true, in general, for data obtained from signals generated by frequency standards. Here, the situation usually deteriorates as more data points are included.

This tutorial describes, largely in heuristic terms, modified statistical measures which do not “blow up” for the kinds of noise processes normally encountered in signals generated by frequency standards.

I. Introduction

This paper is about the question “How do we characterize the performance of frequency standards?” Historically two approaches have developed: one in the frequency domain and the other in the time domain. Our emphasis here will be on the time domain approach; however, we shall begin by considering the frequency domain approach because familiarity with this technique helps to understand why the two usual measures in the time domain, the standard deviation and its square, the variance, don’t work when applied to frequency standards.

Let’s begin by considering an ideal frequency standard which, of course, does not exist in nature. Such a standard generates a noise-free output with constant amplitude, frequency and phase. Real frequency standards fall short in all three categories. However, variations in amplitude with time are, for practical purposes, negligible, so we won’t consider that as a problem in the rest of this paper.

Mathematically, then, the output signal, $V(t)$, of our ideal standard is

$$V(t) = V_p \sin(2\pi\nu_0 t) \tag{1}$$

*Contribution of the U.S. Government, not subject to copyright.

where V_p , is the constant amplitude and ν_0 is some constant frequency.

What real frequency standards generate is

$$V = V_p \sin[2\pi\nu_0 t + \phi(t)] \quad (2)$$

where the phase, ϕ , is a noise perturbation of the desired frequency.

The phase noise, ϕ , falls into two categories: noise we can remove and those we can't. If the signal phase changes predictably with time, a constant frequency offset from the desired frequency for example, we can correct the output by appropriate compensating electronic circuitry or on paper. By definition, the unpredictable, or random part of the signal can't be corrected. Here our only avenue for improvement is to use statistical techniques. However, these statistical improvements are in reality based on assumptions—some model or models—which we believe, or at least suspect, and in a last ditch effort hope, underlie the random part of a signal.

All this adds up, often enough, to a kind of catch 22. Consider a clock whose resonator is a swinging pendulum. We suppose that the clock's misbehavior is related to changes in the weather. Specifically our model says that temperature changes alter the length of the pendulum and thus its natural frequency. We want to predict how much the pendulum clock's time departs, on a daily basis, from time kept by a perfect clock over, say, a period of a month. It's easy. All we have to do is foretell the weather day by day for the next month.

Here our model is probably correct but it doesn't help much in any detailed way. The best we can do, in the absence of actual temperature measurements, is depend on local weather forecasts which are notoriously unreliable.

The actual situation for frequency standards may be worse. A good deal of statistical information about frequency standards has accumulated from numerous measurements over the years, but nobody has any respectable models for what causes the most troublesome kinds of noise. Put differently, the models themselves are statistical models based on measurements, not statistical models based on the physics of the underlying behavior. It's as though we have reams of data giving the average temperature but we know nothing whatsoever about the laws governing atmospheric processes. With this state of affairs about the best we can do is say the temperature tomorrow will be about what it was today.

Nevertheless, much has been gained by inspecting frequency standards through the eyes of statisticians. As an example, in section VI, we consider how various statistical models of frequency standard performance can be applied to determining the performance of a clock over a specified time.

II. Measuring Performance in the Frequency Domain

Consider again the output of an ideal frequency standard as expressed by Equation 1. Figure 1a graphically represents Equation 1 in the frequency domain. This figure shows that the output is at frequency ν_0 with amplitude V_p . A variation on Figure 1a, Figure 1b, is to plot the power carried by the signal at frequency ν_0 , rather than the amplitude. This representation is called the power spectrum of the signal, and we shall use this representation from now on for reasons which will be apparent later. We can easily extend the power spectrum representation of a signal, as in Figure

1c, to represent, say, the superposition of three sinusoids at frequencies ν_1 , ν_2 , and ν_3 with powers P1, P2, and P3.

The power spectrum approach is a convenient way to represent a sinusoid which has been perturbed by noise. We can think of the noise as “blurring” the frequency of the sinusoid so that the power spectrum of a real oscillator is more like the one shown in Figure 2 than the idealization in Figure 1b. The frequency spread of the power spectrum curve is a measure of the amount of noise.

The power spectrum in Figure 3 includes both the signal at ν_0 and the noise components at other frequencies, but we are interested primarily in the power spectrum of the oscillator noise. We can isolate this noise as follows. The actual frequency output, $\nu(t)$, of our oscillator signal is just the rate of change in the total phase of the sinusoid given by Equation 2.

$$2\pi\nu(t) = \frac{d}{dt}(2\pi\nu_0 t + \phi(t)) = 2\pi\nu_0 + \frac{d}{dt}(\phi(t)) \quad (3)$$

That is $\nu(t) - \nu_0 = \frac{1}{2\pi} \frac{d}{dt}(\phi(t))$ is the amount by which the output frequency deviates from the desired frequency. A quantity, called the fractional frequency fluctuation, $y(t) = \frac{\nu(t) - \nu_0}{\nu_0}$ is a dimensionless quantity with greater utility than $(\nu(t) - \nu_0)$ as illustrated by the following example:

Suppose the quantity $(\nu(t) - \nu_0)$ equals 1 Hz for two different oscillators. Are the two oscillators equal in quality? Not if one oscillator operates at 10 Hz and the other at 10 MHz. In the first case, the average value of the fractional frequency fluctuation is 1/10, and in the second 1/10,000,000 or 1×10^{-7} . The 10 MHz oscillator is thus more precise—has a higher Q. If frequencies are multiplied or divided, using ideal electronics, the fractional stability is not changed. Since $y(t)$ represents the fractional frequency excursions of the signal frequency about the nominal frequency, ν_0 , the power spectrum of $y(t)$ characterizes the oscillator noise independent of the oscillator signal frequency.

III. Real Oscillator Noise Viewed from the Frequency Domain

Many workers have measured the noise associated with high quality oscillators. Typically the measurements yield the result shown in Figure 3. Here we show the power spectral density, $S_y(f)$ of the fractional frequency fluctuations $y(t)$. The power spectral density $S_y(f)$, is simply the power spectrum normalized so that the area under the $S_y(f)$ curve equals one. $S_y(f)$ is one of the two commonly used measures of frequency instability—the one in the frequency domain. The figure shows for frequencies ranging from 10 Hz to about 0.01 Hz the noise power spectral density is constant with frequency. This kind of noise, called “white noise FM”, can be reduced by averaging measurements taken over an extended period of time. The reason it can be “averaged out” is because these noise fluctuations, on the average, advance the phase of the signal as much as they delay it.

In the range of frequencies from about 0.01 Hz to about 10^{-4} Hz, the noise density increases as f^{-1} . This kind of noise is called “flicker noise FM” and its cause in high quality oscillators is not well understood. However, it is probably related to power supply voltage fluctuations, magnetic field fluctuations, and component changes, etc. The frequency range over which it dominates, varies from oscillator to oscillator.

At frequencies lower than 10^{-4} Hz, noise with an f^{-2} dependency is commonly found. This noise, "random walk FM noise", is probably related to environmental effects such as mechanical shock, temperature variations, vibrations, etc. which cause random shifts in the carrier oscillator signal frequency. It can be minimized by carefully isolating the oscillator from its environment.

IV. Effect of Record Length and Sampling Rate on Estimating $S_y(f)$

Suppose we have an analog record of oscillator noise, $y(t)$, 50 seconds long, Figure 4a. We sample this record once every second, producing 50 discrete data points y_1, y_2, y_3 , etc. This record contains a mixture of noise fluctuations at many different rates or Fourier frequencies. However, the sampling rate, once per second, and the length of the record (50 seconds), place bounds on the range of frequencies which can be recovered from the sampled data. First, the sampling rate dictates the highest frequency, the Nyquist frequency, which can be recovered or resolved from the data points, as illustrated in Figure 4b.

Samples are taken at the instants marked by the dots. At this sampling rate we can resolve the fluctuation labeled 1 Hz. Fluctuations at a higher rate, such as the one labeled 5 Hz will be confused with frequency fluctuations at lower frequencies. In this particular example there is no way to distinguish between the signal fluctuations at 1 Hz and at 5 Hz. To avoid this kind of confusion we must either sample at a rate which is high enough to resolve the highest rate of noise fluctuation we expect in the oscillator noise spectrum, or alternately, we must remove all frequencies above the highest frequency which can be resolved by our sampling rate. This could be done by low pass filtering the data before sampling. Later in this section we shall describe a digital filter which accomplishes this purpose; although, the data could be filtered with an analog filter to achieve the same result.

In general the highest frequency we can resolve by sampling every Δt seconds, is $\frac{1}{2\Delta t}$.

The length of the record dictates the lowest noise frequency we can detect. Obviously we cannot detect a fluctuation with a one year period by inspecting a record 50 seconds long. More precisely, the lowest frequency we can detect from a record T seconds long is $2/T$ Hz or .04 Hz when $T = 50$ seconds.

Most procedures developed to determine the power spectral density from a series of discrete data points assume that the sampling rate is several times faster than needed to resolve the highest fluctuation frequency expected in the data and the record several times longer than needed to resolve the lowest frequency expected.

V. Time Domain Characterization of Stability

In the preceding section we showed how oscillator noise can be characterized in the frequency domain. As a practical matter, the output signals are often sampled to obtain a discrete set of y_i and then mathematically manipulated to obtain an estimate of $S_y(f)$. This often requires considerable computation. However, for the types of noises which plague most oscillators, there is a considerably simpler approach (from a computational point of view) which is the second, commonly used, characterization of instability mentioned in the introduction.

In this method we again start with a set of y_i , but instead of converting these measurements to a frequency domain representation we develop a procedure for determining the dispersion or scatter of the y_i as a measure of oscillator noise. The bigger the scatter or dispersion of the y_i the greater the instability of the output signal of the oscillator.

The most common statistical measure of dispersion is the standard deviation (or equivalently the variance which is the square of the standard deviation). The variance, R^2 , is simply a measure of the numerical spread of a set of data points with respect to the average or mean value, \bar{y} , of that set of data points. The formula for computing the variance of a set of M data points taken at uniform intervals of time is:

$$R^2 = \frac{1}{M-1} \sum_{i=1}^M (y_i - \bar{y})^2$$

The reason for squaring the terms $(y_i - \bar{y})$ is that an excursion in y_i from the mean always gives a positive contribution in the calculation of scatter. That is an excursion in the positive direction is not canceled by one in the negative, as is the case for the mean.

A very simple measure of the instability of an oscillator is to simply compute the variance of the y_i obtained from a particular oscillator. We would expect our confidence in our measurement of the variance to increase as we increased the number of data points used in our computations. Unfortunately, for common types of oscillator noise, this is only true to a point. To understand this better refer back to Figure 3 which illustrates the spectral noise density typical of frequency standards. Suppose we have sampled the oscillator noise once per second. We decide to estimate the variance of $y(t)$ from 100 consecutive data points. As we have seen already, data taken at intervals of one second over a record length of 50 seconds effectively "isolates" those noise fluctuations with frequencies in the 0.5 to .04 Hz range. Or stated differently, only those noise fluctuations in the 0.5 to 0.04 frequency range contribute to our estimate of the variance. As we increase the data length, lower and lower noise frequencies contribute to the variance. As we see from Figure 3, however, the noise fluctuations have a constant power down to about 0.01 Hz. This means that as long as our data length does not exceed 200 seconds, the variance does not depend on record length. (It is true, however, that our confidence in the estimate of the variance depends on the number of data points, so that, for best confidence in the variance, we should use 200 data points.)

If we extend our calculations to include data taken over an interval which exceeds 200 seconds, the variance will be influenced by noise fluctuations from that part of the noise spectrum which is not flat with frequency. Now the variance grows in magnitude with data length. This is unfortunate. We would like our statistical measure of scatter to converge toward some value as we increase the number of values used in our estimation—not move off to new values with increasing data length.

We must modify our calculation of variance as applied to oscillator noise, so that the variance converges toward some value as we increase the number of data points employed in our calculations.

VI. Stationarity

In the previous section we saw that the variance does not yield a satisfactory result when applied to frequency standard noise. The problem we are encountering is similar to the problem that arises when we attempt to compute the average value for data which has a slow drift. This drift might be due to some instrumental fault so the drift is an artifact of the measurement procedure.

Consider the record shown in Figure 5-a. These data have random fluctuations superimposed on a slow drift. Obviously the mean value computed for the first half of the record does not equal the mean computed for the second half. In this example the mean depends on when the mean is computed.

With stationary data the mean and variance converge toward some value as the record length increased. Obviously data with a slow drift is non-stationary so the mean and variance never converge to any particular value.

We would like to have a measure of dispersion of frequency standard noise which converges toward some value as record length increases. If we apply the normal variance, R^2 , calculation to the data shown in Figure 5a, this will not be so; since, the mean, which enters into the variance calculation, changes with record length.

Suppose, however, that we calculate another quantity, d^2 , which is similar to R^2 except that instead of subtracting the mean from each data point before squaring their summation, we subtract the previous data point. That is:

$$d^2 = \frac{1}{M-1} \sum_{i=1}^M (y_i - y_{i-1})^2. \quad (4)$$

The terms $(y_i - y_{i-1}) = \Delta y_i$ take on nearly the same values that they would whether a trend is present or not i.e., the quantity d^2 is insensitive to the absence or presence of a trend, while R^2 is sensitive.

Let us examine d^2 from a different point of view. Figure 5b shows a plot of the Δy_i 's. The Δy_i plot is similar to Figure 5a except that the trend is missing. In fact, we can think of the process of taking successive differences (computing the Δy 's) as equivalent to running the data in Figure 5-a through a high pass filter which removes the trend. Equation 3, then, is akin to computing R^2 but contains the additional feature of removing the trend. With the trend removed, we would expect d^2 to converge with increasing data length where R^2 did not.

Let's consider now, in more detail, the nature of the high pass filter produced by taking successive differences. We could determine experimentally the frequency response of this filter by taking a large number of digitized sine waves with the same amplitude but with different frequencies. Suppose for example we took a 1 Hz sinusoid sampled ten times per second to produce a set of y_i 's, Figure 6a. We now take successive differences between the y_i to produce the Δy_i which are also shown in Figure 6b. As we see we obtain, again, a sinusoid of the same frequency but with reduced amplitude (and shifted phase which is not important here). Similarly we could take other sinusoids of unit amplitude and in a similar manner determine their amplitudes upon taking successive differences. The result of such an experiment is shown in Figure 7a. As we see the filter does not attenuate the frequency at the reciprocal of two times the sampling rate (0.05 Hz) but does attenuate all lower frequencies with the attenuation increasing as the frequency approaches zero.

In some situations we might want to leave the trend in and remove the fast fluctuations. For example, suppose we have data consisting of hourly temperature readings taken over a period of one year. If we average the hourly temperature readings, a day at a time, we remove the hourly fluctuations leaving just the day to day trend. The response of this low-pass averaging filter is shown in Figure 7b. It has unity gain at 0 Hz and zero gain at $1/2$ (averaging interval) = $1/[2 (24 \text{ hours})] = 5.8 \times 10^{-6}$ Hz.

Suppose, now, we have a set of measurements which we first average, corresponding to low-pass filtering, and then difference, corresponding to high-pass filtering. We can think of these two operations together as equivalent to band pass filtering the original data. The shape of this band pass filter is simply the product of the low and high pass filters, Figure 8a.

Figure 8a shows the response shape of the band pass filter corresponding to first averaging the data two seconds at a time and then taking successive differences. As we increase the time over which we average the data and then take differences, the frequency at which the band pass filter has maximum response moves toward a lower frequency and the width of the filter decreases. To demonstrate this we have plotted in Figure 8-b the band pass filter response corresponding to averaging and differencing times of 2, 4, and 10 seconds.

From this illustration we see that by adjusting the length of the averaging time of the original data points, we can move the center frequency of the filter. This suggests that we might use such a band pass filter to estimate the power spectral density of oscillator noise. Suppose we first average and difference our data, the y_i , with the filter centered at 1.0 Hz. The filtered data now contains only those noise components near 1 Hz. From this filtered data we compute a value for the variance. Next we center the filter at 0.5 Hz and again compute the variance from the filtered data. Now we have computed the variance for the noise fluctuations in the vicinity of 0.5 Hz. We can continue this process indefinitely letting the filter center frequency approach 0.0 Hz. If we now plot the variances obtained as a function of the associated center frequency of the filter, we obtain an estimate of the noise power as a function of frequency which is similar to the kind of information contained in a power spectral density, $S_y(f)$, plot. This is the reason we have chosen to work with the variance rather than the standard deviation since the standard deviation has the dimensions of amplitude.

It can be shown, on a rigorous mathematical basis, that if we pass noise with a $1/f$ or $1/f^2$ power spectral dependence through the type of digital filter we have been discussing, that the noise passing through this filter is stationary. If we now compute the variance of this filtered noise, we can be assured that the variance will converge as the record length increases. Further this variance is a measure of the noise fluctuations in a particular frequency band which is dependent upon the frequency at which the bandpass filter is centered.

VII. Pair Variance

With the background developed in the previous section, we can now define explicitly the second measure of frequency instability (the one in the time domain) which is commonly employed. This measure, often called the "pair variance," is defined by:

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{i=1}^M (y_i - y_{i-1})^2$$

where M is the number of data points and the y_i are obtained by averaging the data, in M segments, each τ seconds long.

We can consider this formula in the light of the bandpass filtering discussed in the previous section. First, we notice that the y_i are obtained by averaging the signal in chunks τ seconds long, which corresponds to low-pass filtering the data. Second, the differences are taken between successive pairs of data which are separated in time by τ seconds; because, the original data was averaged τ

seconds at a time to create the y_i . As we know the averaging and differencing procedures amount to band pass filtering the data to produce stationary noise—assuming the spectral noise density varies as $1/f$ or $1/f^2$. We notice that there is a factor $1/[2(M - 1)]$ in front of the summation sign which differs by a factor of $1/2$ from the formula for R^2 . This is because we want the two formulas for R^2 and σ_y^2 to give the same numerical result for white noise FM—which is stationary.

Let's apply the variance to oscillation noise of the type shown in Figure 3. Suppose we compute $\sigma_y(\tau)$, the square root of the pair variance, for $\tau = 10$ seconds. This corresponds to calculating the standard deviation for data which has passed through the filter labeled $\tau = 10$ seconds in Figure 8b. As we increase τ , our band pass filter moves toward lower frequencies and narrows. As long as the spectrum is constant with frequency, σ_y decreases with τ ; since, the filter narrows as it moves toward lower frequencies. As we move into the f^{-1} portion of the spectrum, the bandpass filter continues to narrow as the noise power increases. In this region the filter narrows at a rate which exactly compensates for the increase in noise power so that σ_y is flat in this range. This flat region is called the "flicker floor" of the oscillator, and from a practical point of view, once the flicker floor is reached, no further gains are made by averaging. As we move into the f^{-2} region, the noise power increases at a rate which is faster than is compensated for by the narrowing filter so that in this region σ_y grows with τ .

Obviously the variation of $\sigma_y(\tau)$ with τ is related to the spectral density noise $S_y(f)$ as indicated by the arrows connecting the two parts of Figure 9. In fact, it can be shown mathematically that the slopes of the σ_y vs τ curve are -1 , 0 , and $+1$ for f^0 , f^{-1} , and f^{-2} types of noise respectively. We see then that the σ_y vs (τ) curve is a fairly simple way, from a computational point of view, to infer the noise spectrum of the oscillator.

Although we have been particularly concerned with f^{-1} and f^{-2} noise, it can be shown as we said that the pair variance converges for noise which is divergent up to, but not including f^{-3} noise. At f^{-3} , and beyond, $\sigma_y(\tau)$ does not converge as the data length increases.

We have considered two measures of frequency stability: one in the frequency domain, $S_y(f)$ and the other in the time domain, $\sigma_y^2(\tau)$. From a computational point of view, σ_y^2 is a convenient measure of stability and can be used to characterize the spectrum of oscillation noise. On the other hand, the power spectral density, $S_y(f)$, is quite often more convenient for theoretical developments since it is a fundamental quality. Cutler and Searle (reference 1) have shown that σ_y^2 and $S_y(f)$ are related as follows:

$$\sigma_y(\tau) = \int_{-\infty}^{\infty} df S_y(f) \frac{\sin^4 f}{(\pi f \tau)^2}$$

This equation is the analog of what we have developed from a digital filter point of view. The term:

$$\frac{\sin^4 f}{(\pi f \tau)^2}$$

can be thought of as a band pass filter, through which the noise spectrum $S_y(f)$ is passed, during the integration over f , to determine σ_y for a particular τ .

VIII. Characteristics of Real Frequency Standards and Clocks

Now that we are armed with the primary time domain performance measure of frequency standards let's see how the various kinds of standards stack up. Figure 10 illustrates what we might expect from a number of different devices.

For a cesium frequency standard we have displayed three curves: one represents a standard portable commercial device, another a high performance commercial standard, and finally, a laboratory standard which would be found normally at a national standards laboratory. We include both a standard and a high performance rubidium standard. OCXO means an oven controlled crystal oscillator.

With these characterizations let's consider one example of how we might use these performance measures: the problem of predicting the performance of clocks whose resonators are based on the frequency standard characteristics discussed in the previous section.

By now we know that the kind of noise that dominates, for a particular frequency standard, depends on the sample time τ . For example we see, from Figure 3 that a standard rubidium device is dominated with flicker noise FM—the flat part of any of the curves in the figure—at about $\tau = 1000$ seconds while the transition for the passive hydrogen maser doesn't occur until about $\tau = 100,000$ seconds.

Since a clock is based on a frequency standard, the degree to which a clock departs from a perfect clock over time depends on the kind of noise dominating the frequency standard for the prediction time of interest. In our discussion we assume that constant frequency offsets and linear drifts in frequency have been removed from the frequency standard driving our clock, and that only the random noise is keeping our clock from perfection.

For the three regimes of noise illustrated in Figure 3 reference 2 shows that:

1. When white noise FM and random walk noise predominate, the optimum predictor for clock dispersion is simply $\tau\sigma_y(\tau)$
2. when flicker noise FM is dominant the optimum predictor is $\tau\sigma_y(\tau)\ln 2$.

Figure 11 shows, applying these rules, the time dispersion—the RMS time prediction error—we would expect for active and passive masers, a cesium standard, and a rubidium standard.

As a specific example, consider a large collection of clocks all of whose resonators are cesium frequency standards with the same noise properties. We want to know the RMS scatter of these clocks after, say, 100,000 seconds—a little over one day—after they have been synchronized. As Figure 11 shows at 100,000 seconds the RMS spread is about 5 nanoseconds. From a practical point of view this means that if we have a large collection of clocks, and want to keep them synchronized to about 5 nanosecond, we need to resynchronize them about once a day.

If these clocks were at nodes in a communications systems we would need to dedicate some of our communication channel to achieve this level of synchronization, or alternately recalibrate against some external time source, say a satellite time signal, about once a day.

Of course clocks based on other kinds of resonators would need resynchronization more or less often depending on the quality of the resonator.

IX. Concluding Remarks

The notion of averaging and differencing successive data points to remove both short term and long term variations is an old one. Meteorologists have used these techniques for at least 150 years in their studies of temperature and pressure variations. In more recent times a 1942 paper by von Neuman *et. al.* explicitly states the problem of applying the classical standard deviation to data containing a trend [reference 3]:

“There are cases, however, where the standard deviation may be held constant, but the mean varies from one observation to the next. If no correction is made for such variation of the mean, and the standard deviation is computed from the data in the conventional way, then the estimated standard deviation will tend to be larger than the true population value. When the variation in the mean is gradual, so that a trend (which need not be linear) is shifting the mean of the population, a rather simple method of minimizing the effect of the trend on dispersion is to estimate standard deviation from differences.”

Clock metrologists have developed these ideas into rather sophisticated but relatively simple tools for the study of the important noise processes in frequency standards. In the time and frequency literature you will find the pair variance also referred to as the “two sample variance” or the “Allan variance” after David Allan [reference 4] who has focused these techniques on frequency standards. But the problem of characterizing oscillator frequency stability is by no means closed. We have discussed the most commonly encountered measures in the frequency and time domains, but anyone who wishes to pursue current investigations will find reference 5 a good place to start.

References

1. Cutler, L. and Searle, C., “Some Aspects of the Theory and Measurement of Frequency Fluctuations in Frequency Standards.” Proc. IEEE, Vol. 54, pp. 136-154, Feb., 1966
2. Allan, D. and Hellwig, H., “Time Deviation and Time Prediction Error for Clock Specification, Characterization, and Application.” PLANS 1978 Conf. Proc., pp. 30-38
3. von Neuman, J., H.R. Kent, H.R. Bellinson, and B.I. Hart, “The Mean Square Successive Difference,” Ann. Math. Stat., Vol. 12, pp. 221-331, 1942
4. Allan, D., Marc Weiss and J. Jespersen, “A Frequency Domain View of Time-Domain Characterization of Clocks and Time and Frequency Distribution Systems.” 45th Annual Symposium on Frequency Control, May 1991
5. Rutman, J. and Walls, F., “Characterization of Frequency Stability in Precision Frequency Sources.” Proc. IEEE, Special Issue on Time and Frequency, Vol. 79, pp. 952-960, July 91




Figure 1a Signal Characterization of a Sinusoid in the Frequency Domain

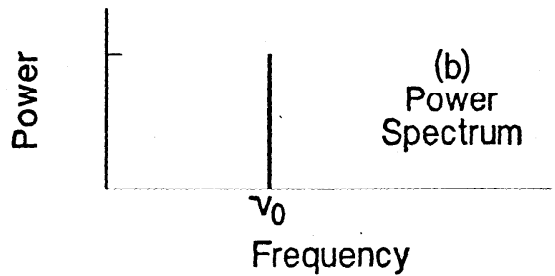


Figure 1b Power Characterization of a Sinusoid in the Frequency Domain




Figure 1c Power Characterization of Sinusoids at Three Different Frequencies in the Frequency Domain

Noise Power Spectrum




Figure 2 Power Spectrum of a Noisy Signal

Spectral Density vs. Fourier Frequency as Translated from Time Domain

Random Walk FM

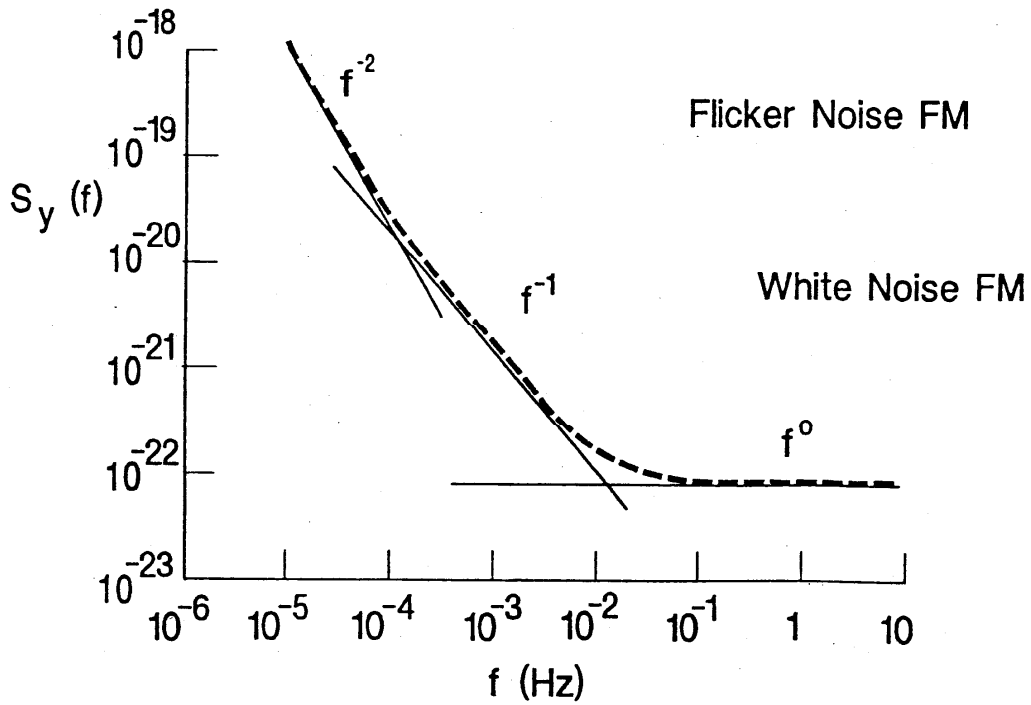


Figure 3 Power Spectral Density of a Typical Frequency Standard




Figure 4a A Sampled Data Record




Figure 4b Confusing a 1 Hz with a 5 Hz Signal

High Pass Filter

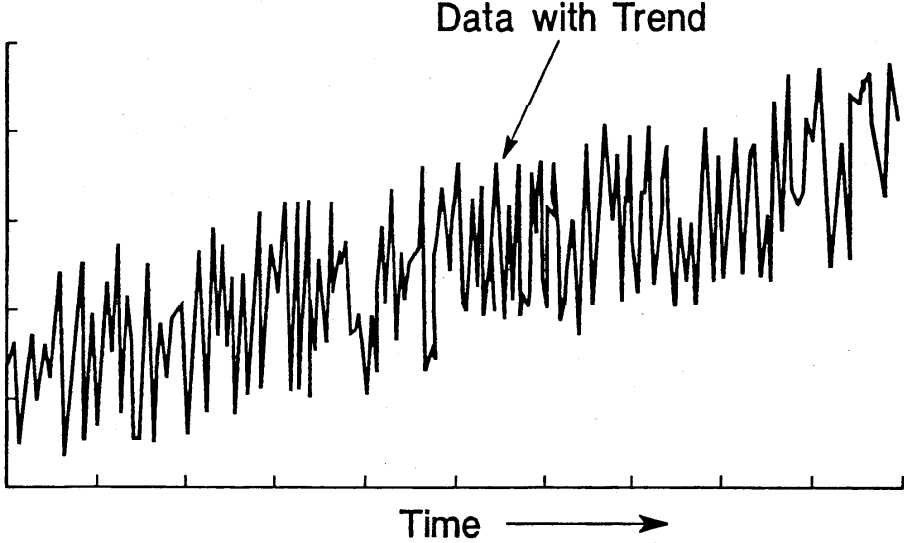


Figure 5a Data with a Trend




Figure 5b Data with Trend Removed




Figure 6a Original Signal




Figure 6b Signal after Differencing Data

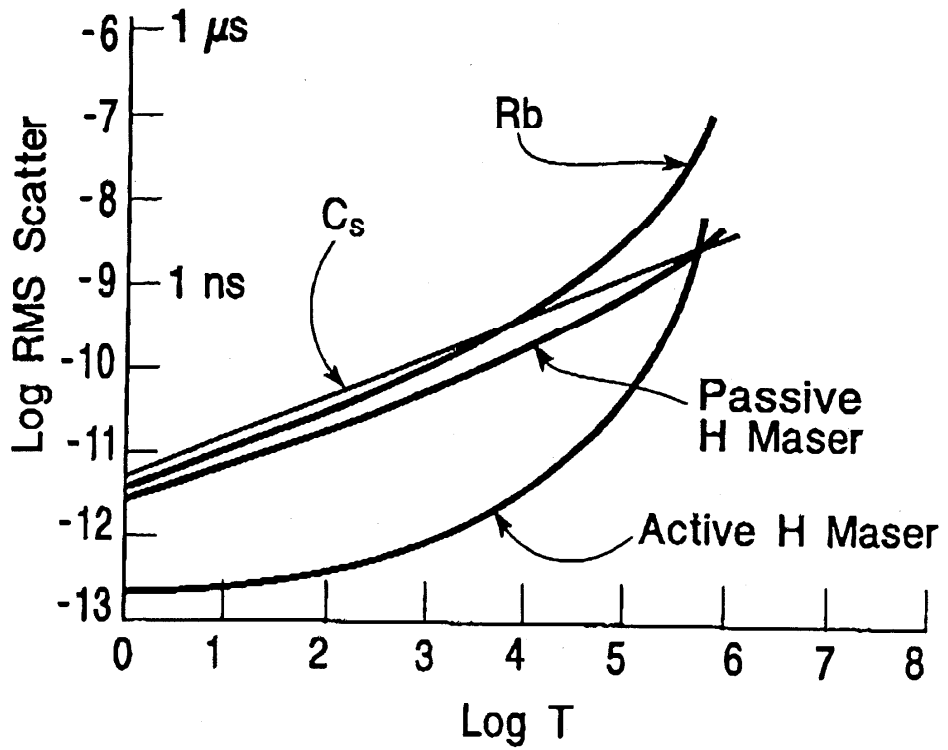


Figure 11 Predicting Clock Performance

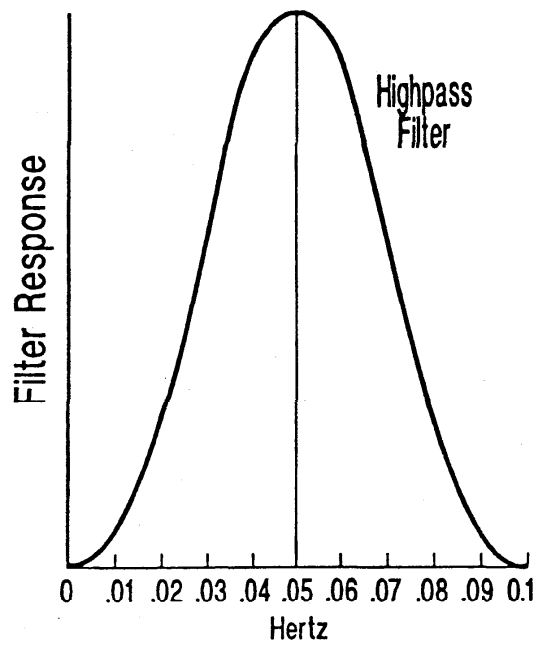


Figure 7a High-Pass Filter due to Differencing Data




Figure 7b Low-Pass Filter due to Averaging Data

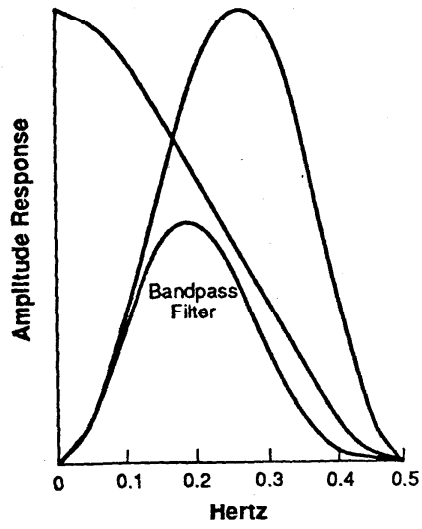


Figure 8a Band-Pass Filter due to Averaging and Differencing Data




Figure 8b Change in Filter Characteristics due to Changing Averaging Time




Figure 9 How Power Spectral Density and Pair Variance are Related




Figure 10 Characteristics of Typical Frequency Standards




Figure 11 Predicting Clock Performance