A FREQUENCY-DOMAIN VIEW OF TIME-DOMAIN CHARACTERIZATION OF CLOCKS AND TIME AND FREQUENCY DISTRIBUTION SYSTEMS

David W. Allan, Marc A. Weiss, and James L. Jespersen

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

Abstract

An IEEE standard (No. 1139-1988) now exists for "Standard Terminology for Fundamental Frequency and Time Metrology. As defined in this standard, the time-domain stability measure, $\sigma_y(\tau)$ has evolved into a useful means of characterizing a clock's frequency stability. There exists an ambiguity problem with $\sigma_y(\tau)$ for power - law spectral densities, $S_y(f)$, proportional to f^α, where $\alpha \geq +1$. For example, white noise phase modulation (PM) and flicker noise PM appear the same on a $\sigma_y(\tau)$ plot. Because of this ambiguity, $\text{Mod}\sigma_y(\tau)$ was developed.

More recently, it has become apparent there is no accepted measure for the performance of time and frequency distribution systems. At the current time, there is an important need for a good method for characterizing time and frequency transfer links in telecommunication networks.

Last year at this symposium suggestions were given for ways to characterize time and frequency distribution systems. Because of the above ambiguity problem, $\sigma_y(\tau)$ was shown to be a less useful measure than $\text{Mod}\sigma_y(\tau)$ for such systems. It was shown that $\sigma_x(\tau) = \tau \cdot \text{Mod}\sigma_y(\tau)/\sqrt{3}$ is a useful measure of time stability for distribution systems. For the case of white noise PM, $\sigma_x(\tau)$ is simply equal to the standard deviation for τ equal to the data spacing, τ_o and is equal to the standard deviation of the mean for τ equal to the data length (T).

In this paper, we recast these measures into the frequency-domain. We treat each of these measures as a digital filter and study their transfer functions. This type of measure is easily related to the passband characteristics of a given system, a particularly useful engineering approach.

Introduction

This paper concerns the characterization of frequency standards, clocks and associated systems. These associated systems may include: time and frequency measurement systems, time and frequency transmissions systems, time and frequency comparison systems, and telecommunication networks. As we shall see no single characterization is suitable. However, in this paper we discuss three statistical characterizations which cover most of the situations encountered in actual practice.

The characterizations that we require can be approached from two points of view: the time domain and the frequency domain. In this paper we describe in general terms the time and frequency domain approaches and then explain in some detail how the time domain approaches can be interpreted from the frequency domain point of view. We feel it is important to make this interpretation because of the significance of frequency domain approaches--particularly in engineering environments.

We will next explain why the mean and the standard deviation don't work for frequency standards. The simple mean, the standard deviation, and its square the classical variance, are well known statistical measures of a set of data points. It seems natural therefore that we should apply such measures to characterize various kinds of clock associated processes. However, application of these quickly reveals some significant problems.

Let's begin by considering the computation of the mean frequency output of a frequency standard. Normally, when we compute the mean of some process we suppose that including more data points in the computation brings us ever closer to the true mean of the process and that an infinite number of data points yields the true mean. Of course in the real world we must be content with a finite number of points always leaving some uncertainty in our attempt to find the mean. Nevertheless we assume that we can approach the true mean as nearly as we like if we are willing to collect enough data points.

Is this true of frequency standards? The answer surprisingly enough is "No!". How can this be? The answer is simple enough. Frequency standards do not generate a constant frequency output contaminated only by white noise. If they did we could average the output to get rid of the noise. That is, white noise is the kind of noise that can be averaged away. This is because for every phase advance it produces in the output of the frequency standard, it eventually produces a compensating phase retardation, so that the two cancel in the averaging process.

Much research has shown that a number of noise processes, in addition to white noise, afflict frequency standards. we can qualitatively say that these other kinds of noise represent trends--not necessarily linear trends--in the output frequency of a standard. In later sections of this paper, we will quantitatively identify these "trends" but for now we will stay with the generic "trend" since this notion is enough to demonstrate why standard statistical methods don't work for frequency standards.

As an example, consider a frequency standard whose frequency output is contaminated with white noise and also increases linearly with time--a very simple kind of trend. What can we say about the mean output frequency of such a standard? Not much, because there is no average; we have defined the standard as producing a monitonically increasing frequency. The frequency we find by averaging is a function of when we start the measurement and the length of time over which the average is made.

What can we do when confronted with such a situation? An obvious answer is to remove the trend, by whatever means, and then compute the mean in the normal way from the modified data. In this way we would expect to converge on the true mean as we average away the white noise by using ever more data points.

The process we have just described is the essence of a number of approaches that have been developed to produce useful statistical measures for the output signals of frequency standards and related devices.[1] We can think of it as a two step process: First we remove the offending trends, and then we compute the statistical measures in the standard way. This process is not always evident when we look at the statistical measures in common use. More often than not, the two steps are combined into one obscuring the underlying process.

We will next describe the frequency- and time-domain measures of frequency and time stability. The standard deviation and the mean are examples of what are called time domain measures. That is, we collect a number of data points, one after the other, and then use these points to construct some useful statistical measure. There are also frequency domain measures. The power spectral density of a set of data points is an example. It provides us with a picture of the deviations in the data having a particular (Fourier frequency) spectral component.

As you might suspect, frequency- and time-domain measures are related. In later sections of this paper we explore in some detail the relationships between time- and frequency-domain measures. However, as a simple example consider the following. Suppose we want to remove a long term trend from the output signal of a frequency standard. There are a number of ways we might proceed. We might, for example, fit a polynomial to the data and subtract this polynomial from the data to generate a new set of data which we could then treat in standard statistical fashion. If the trend were strictly linear, then the polynomial curve would simply be a straight line whose slope revealed the magnitude of the frequency drift.

Another approach is to pass the data through a high pass digital filter to remove the low frequency components--which is what a trend looks like to such a filter.[2]

A particularly simple high pass filter can be constructed by taking what are called "first differences" of the data. That is we subtract the kth data point from the kth + 1 data point, for all data points. This process effectively removes long term trends from the data.

As we shall see later the "first differencing" process corresponds to a digital filter in the frequency domain whose characteristics can be defined precisely. We should also add that the polynomial fitting procedure also corresponds to a particular digital filter. However the emphasis in this paper is on "difference" type procedures since they are easy to implement and are commonly used by the time and frequency community.

We now describe three different variances. One is particularly useful for characterizing the frequency stability of clocks and oscillators. The next is most useful for characterizing the frequency stability of time and frequency measurement systems, distribution and comparison systems as well as for distinguishing between white noise PM and flicker noise PM. The last is most useful for characterizing the time stability of any of the above as well as for network synchronization; e.g. telecommunications network. All three of these variances are built upon taking finite differences of the data.

The notion of taking differences to remove trends in data is an old one. We quote von Neuman et. al. from a 1942 paper [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 the standard deviation from differences."

Perhaps the most important part of this quote is the parenthetical "which need not be linear." As it turns out taking first differences and second differences--repeating the differencing process twice--is sufficient to remove most of the kinds of noise that one encounters in clocks. Much of the work in the last few decades in the statistical characterization of frequency standards has been directed toward understanding in detail the implications of using the differencing approach. Today this approach is the mainstay of time domain approaches to characterizing time related processes. In this paper we will focus on three such characterizations with their associated interpretations in the frequency domain. We briefly introduce them here with details following in later sections.

The first statistical measure we want to introduce, also historically the oldest, is called the "two-sample variance," the "pair variance" or the "Allan Variance" [1, 4-6] It is denoted $\sigma_y^2(\tau)$ and referred to herein as AVAR. It is defined as follows:

$$\sigma_y^2(\tau) = \frac{1}{2} <(\Delta y)^2>, \tag{1}$$

where the brackets "< >" denote expectation value, Δ is the first finite difference operator and y is the relative frequency offset as defined below.

AVAR was developed to address the problem of finding a suitable measure of the variability of the output frequency of a frequency standard. As we know, the computation of the standard variance will not work when applied to frequency standards because they contain noise processes which cannot be averaged out. The core idea of the Allan variance has already been introduced--the differencing method. Here the data to be differenced consists of a number of samples of the frequency of the standard taken over some period of time. The differencing procedure filters the noise processes that make the normal variance computation unsuitable. The actual formula, discussed later, accomplishes both the filtering and the variance computation in the same step.

Before we leave AVAR in this introductory section we should point out one thing. The output of a frequency standard is actually a signal whose phase advances in time with respect to some reference. We use phase or time difference almost interchangeably. This is so because they are directly proportional: $x(t) = \phi(t)/2\pi\nu_o$, where $\phi(t)$ is the phase difference reading in radians between two standards. The dimensions of x(t) are time. In practice, the frequency is derived by measuring the time or phase difference x(t) of the signal between the standard in question and the reference at two different times say t and $t+\tau$ giving us phases x(t) and $x(t+\tau)$.

Let $\nu(t)$ be the output frequency of the standard in question, and let ν_o be the frequency of the reference. We will assume, without loss of generality, that ν_o is perfect. The average relative frequency offset, $y(t) = (\nu(t)-\nu_o)/\nu_o$ of the standard in question over the time interval t to $t+\tau$ is then

$$y(t) = \frac{(x(t+\tau) - x(t))}{\tau} \tag{2}$$

If we think of AVAR from the point of view of phase measurements x(t) instead of frequency measurements y(t), then AVAR is constructed in terms of the second differences in phase but first differences in frequency since frequency by definition is obtained from first differences in phase. An alternative and very useful definition of AVAR is as follows:

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2} <(\Delta^2 x)^2>, \tag{3}$$

where "Δ^2" is the second finite-difference operator.

The second statistical measure we want to introduce is called the modified Allan variance or from here on "MVAR." [5, 7-9] It is defined as follows:

$$\sigma_y^2(\tau) = \frac{1}{2\tau^2} <(\Delta^2 \bar{x})^2>, \tag{4}$$

where $\bar{x}$ denotes phase averages being used in the second difference. We note that equations (3) and (4) are identical except for the phase averages. The three sequential phase averages are each taken over an interval τ. As τ changes, this changes the bandwidth in the software in just the right way to remove the ambiguity problem in AVAR. In other words, MVAR can distinguish between white noise PM and flicker noise PM, whereas AVAR cannot. In a later section this distinction will be more evident when we compare MVAR and AVAR from the frequency domain point of view.

Both AVAR and MVAR are particularly suited to characterizing the frequency instabilities of frequency standards. However there are situations where the emphasis is not on frequency but on time measurements. This brings up the final measurement we want to introduce, TVAR, where the "T" emphasizes the fact that we are focusing on time rather than frequency measurements. It is defined as follows:

$$\sigma_x^2(\tau) = \frac{1}{6} <(\Delta^2 \bar{x})^2>. \tag{5}$$

We see that $\sigma_x(\tau)$ is just $\tau \mathrm{Mod}\sigma_y(\tau)/\sqrt{3}$ and has many of the advantages of $\mathrm{Mod}\sigma_y(\tau)$, but is now a time stability measure.

How does the change in emphasis come about? [10-11] The notion of studying frequency instabilities has a local flavor to it in the sense that frequency is defined by a certain resonance frequency of the Cesium atom or quartz resonator while epoch time is an arbitrary manmade concept requiring coordination over time and space. Thus if we want to compare the frequencies of two remotely located standards we need to introduce some communication link which allows us to compare the phase or time difference between our two clocks. By measuring the change in the time or phase difference between these two standards over time we can determine the frequency offset between the two clocks.

Furthermore we might also want to determine the actual time offset between the clocks which again leads us to making time or phase difference measurements. Both of these examples point up the need for some statistical characterization where the focus is on time or phase rather than frequency, hence TVAR. So there is no confusion, we should point out that the time or phase difference between two clocks is a measurement. Whereas, the finite-difference operators, as in the above variance definitions, operates on a time series of measurement data. The advantages and disadvantages of these three variances will be more apparent later. [11]

We shall also look at the TVAR transfer function from the frequency domain point view This view and some other considerations reveal why TVAR is a more suitable measure for time related measurements than AVAR and MVAR.

Transfer Function Approach to Variances

A variance can be viewed from either the time domain or the frequency domain. [12-14] We intend to look at variances from both perspectives to aid in understanding what certain variances measure. We begin by describing the relationship between the variances in the two domains.

A. Variances in the Time Domain: Convolution

If, in the time domain, we have a time series of observables, $x(t)$, we may define a particular variance with the help of a convolving function, $h(t)$, where $h(t)$ is the impulse response function. We use the convolution, g, of x and h:

$$g(\lambda) = \int x(t)\cdot h(t-\lambda)dt. \qquad (6)$$

The variance corresponding to h, of the time series $x(t)$ is the infinite time average of g squared,

$$\sigma^2 = \lim_{T\to\infty} \frac{1}{T}\cdot \int_{\lambda=-T/2}^{T/2} g(\lambda)^2 d\lambda. \qquad (7)$$

The convolving function h, here, is the important definition. It determines how the variance selects data in the time domain, which is then squared and averaged. For example, the convolving function, h, for the classical variance is the step pulse: [2]

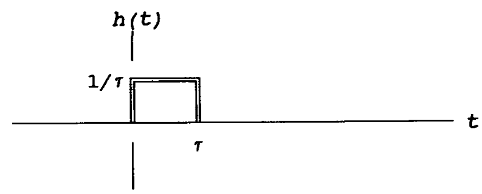

Figure 1. Impulse response function for classical variance. Each sample is taken over with an averaging time τ. Each sample is differenced with the mean, squared and averaged to obtain the classical variance.

This makes

$$g(t) = \overline{x(t)}_\tau\ , \qquad (8)$$

the average value of x from t to $t+\tau$. Thus the variance here is simply the second moment of $\overline{x}$.

For the Allan Variance (AVAR), h is the double pulse:

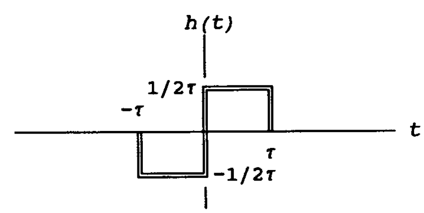

Figure 2. Impulse response function for the Allan or two-sample variance. Adjacent measurements--each averaged over an interval τ--are differenced. This change from one interval τ to the next is squared and averaged across the data, then divided by 2 for an AVAR estimate.

This makes the AVAR the expected value of the first difference squared. The normal use of AVAR is to characterize frequency stability. Thus, if $x(t)$ is a time series of clock time differences, the first difference of these divided by τ, $y(t)$, is the corresponding time series of frequency differences, averaged over the interval τ. Then convolution for the usual AVAR is

$$g(t) = \int y(\lambda)\cdot h_\tau(\lambda-t)d\lambda, \qquad (9)$$

giving a first difference of frequencies averaged over a time interval τ, or a second difference of time values. The integral of the square of this g results in the Allan, or two-sample variance

$$\sigma_y^2(\tau) = \lim_{T\to\infty} \frac{1}{T}\cdot \int_{\lambda=-T/2}^{T/2} g(\lambda)^2 d\lambda = \frac{\langle(\bar{y}_t - \overline{y_{t-\tau}})^2\rangle}{2}. \qquad (10)$$

For the Modified Allan Variance (MVAR), we first note that in the Allan variance the time interval for averaging frequency, τ, is a multiple of the basic sampling interval τ_0. Thus, it is possible to make n shifts of the pulses in figure 2 by τ_0, where $\tau = n \cdot \tau_0$. The modified Allan variance averages several first differences of frequency in this way, thus adjusting the software bandwidth to exploit the bandwidth dependence of white phase noise. As an example we show the convolving function for MVAR as the sum of two functions, as in figure 2, displaced by $\tau/2$.

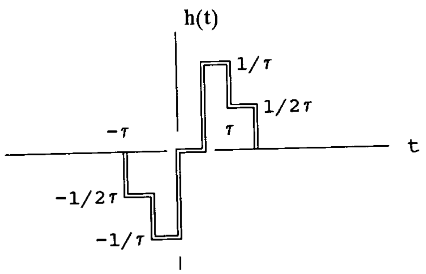

Figure 3. Impulse response function for the modified Allan or two-sample variance. This results from forming a second finite-difference from three contiguous intervals. Each interval contains the phase or time averaged over an interval τ. These second differences are squared and averaged across the data, then divided by $2\tau^2$ for an MVAR estimate.

The new variance, $\sigma_x^2(\tau)$ (TVAR), is simply $\tau^2 \cdot \mathrm{Mod}\sigma_y^2(\tau)/3$. Thus, the convolving function has the same shape as in figure 3, but the vertical scaling needs to be multiplied by $\tau^2/3$.

B. Variances in the Frequency Domain: Transfer Function

We now examine the impulse response functions of these variances as transformed into the frequency domain. We will see that the convolving function h again is the important definition. In this domain, the Fourier transform, H of h, becomes a kind of transfer function for defining the variance.

There are two steps to understanding the passage from the functions we've discussed in the time domain to the frequency domain. First we use the fact that an infinite time average of a function squared equals the integral of the spectrum of that function:

$$\lim_{T \to \infty} \int_{\lambda=-T/2}^{T/2} g(\lambda)^2 d\lambda = \int_0^{\infty} S_g(f)\, df \qquad (11)$$

Since the variance is an infinite time average of the square of the function g, it also equals the area under the power spectral density of g, the square of the Fourier transform of g.

Second, we use the mathematical relation that the Fourier transform of a convolution is the product of the Fourier transforms. Let us put these two facts together. For the general variance σ^2, defined as the integral of the square of the time series $x(t)$ convolved with $h(t)$, we have

$$\sigma^2 = \int_0^{\infty} S_x(f) \cdot |H(f)|^2 df, \qquad (12)$$

where $S_x(f)$ is the spectrum of x, and $H(f)$ is the Fourier transform of h.

This $H(f)$ is the transfer function of the variance. This differs from the usual use of the term "transfer function" in that instead of producing a signal sculpted by the shape of H, we produce a variance which is sensitive to frequencies according to the shape of H. [2]

This last equation, then, gives the relationship between the definitions of variance in the time and frequency domains. We see that the convolving function h that defines how the variance selects data in the time domain, also, via its Fourier transform, defines which frequencies the variance is sensitive to.

C. Transfer Functions of AVAR, MVAR, and TVAR

The transfer function for AVAR is shown in figure 4, a linear plot, for two different values of n, where $\tau = n \cdot \tau_0$. We see that the variance selects a band of frequencies for a given τ, and that the width of this band decreases as τ increases. Also note that the function goes to 0 at the origin. Indeed it goes to 0 fast enough that it remains integrable when multiplied by an f^{α} spectrum with α greater than -3. [6]

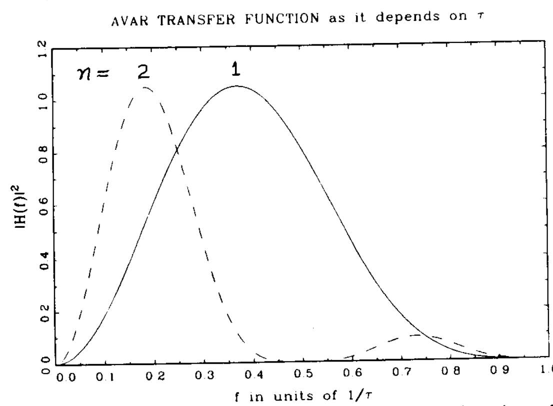

Figure 4. This is a plot of the squared transfer function of the impulse response function shown in figure 2. It is the function that multiplies the spectrum of the frequency deviations to obtain $\sigma_y^2(\tau)$. It is plotted hear for two values of τ (τ_0 and $2\tau_0$ where $1/\tau_0 = 1$). Note, the abscissa is linear and that the bandwidth of H(f) decreases as τ increases.

In figure 5, we see the transfer functions plotted on a logarithmic horizontal axis for n taking on the first eight powers of 2. We see that logarithmically, the bandwidth remains constant, and that the power-of-two transfer functions scan more-or-less independent frequency bands. Figure 6 shows the sum of these transfer functions. We see here that this sum yields a flat band-pass filter. The interpretation here is that this band pass represents the sum of information presented in $\sigma_y(\tau)$ versus τ plot. That is, the Allan variance plot of points chosen with n equal to a range of powers-of-two shows the stability of the data due to a certain band of frequencies. The sensitivity of AVAR to the different frequencies in this band is nearly constant. The band extends from $1/(2n\tau_o)$ to $1/(2\tau_o)$, where n is here the highest power of 2 chosen for the $\sigma_y(\tau)$ plot.

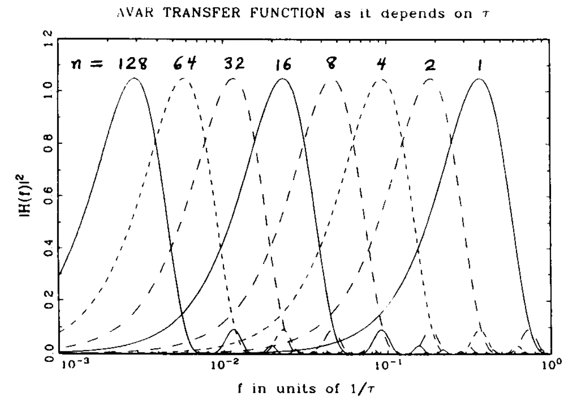

Figure 5. This is a plot of squared transfer functions of the impulse response function shown in figure 2 for 8 values of τ (1, 2, 4, 8....128 x τ_o). Note, the abscissa is logarithmic, and the apparent width of each transfer function is the same. They also appear distributed uniformly across a certain span of Fourier frequencies.

Next we look at the transfer function for MVAR. Analogous to AVAR, we see in figure 7 the MVAR transfer function in a linear plot for two different τ values. Here also, we see that the bandwidth decreases as τ increases, but in addition, the amplitude decreases also. This comes from the additional software filter in MVAR, the phase averaging, allowing MVAR to distinguish white phase noise from flicker phase noise. In figure 8 the MVAR transfer functions for powers-of-two τ values are summed as in figure 6; again, we see a flat band-pass. Thus, an MVAR power-of-two plot also presents the stability information due to a range of frequencies, with nominally equal sensitivity to frequencies within the range. For the same range of τ values, the MVAR cumulative transfer function is a little wider than that of AVAR and the high frequency end is slightly steeper.

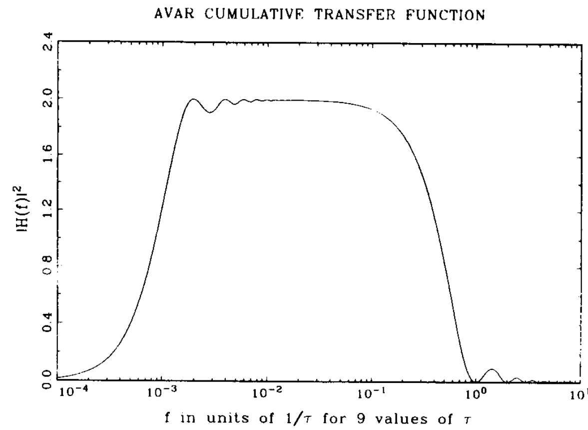

Figure 6. This figure shows the sum of the squared transfer functions for 9 values of τ (1, 2, 4, 8....256 x τ_o) for $\sigma_y^2(\tau)$. We conclude that a $\sigma_y(\tau)$ plot for such a set of τ values gives a nearly constant response to Fourier frequency over about two decades.

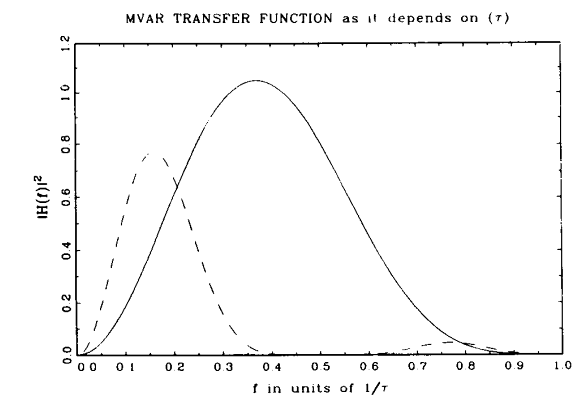

Figure 7. This is a plot of the squared transfer function of the impulse response function shown in figure 3. It is the function that multiplies the spectrum of the frequency deviations to obtain $\mathrm{Mod}\sigma_y^2(\tau)$. It is plotted hear for two values of τ (τ_o and $2\tau_o$ where $1/\tau_o = 1$). Note, the abscissa is linear and that the bandwidth of H(f) decreases as τ increases. Notice also, that the amplitude decreases with increasing τ. This is due to the software band-width change brought about by phase averaging.

In figure 9, we see the TVAR transfer function. Note here that the function "rings" forever. That is, neighboring sinusoidal lobes do not die out. Of course with finite data sampling, there is always a high-frequency

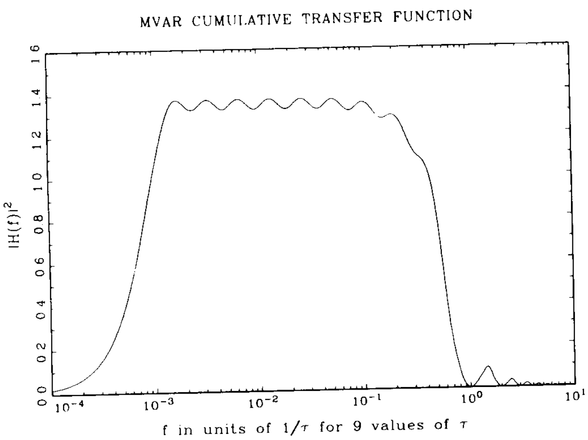

Figure 8. This figure shows the sum of the squared transfer functions for 9 values of τ (1, 2, 4, 8....256 x τ_o) for Mod$\sigma_y^2(\tau)$. We conclude that a Mod$\sigma_y(\tau)$ plot for such a set of τ values gives a nearly constant response to Fourier frequency over slightly more than two decades.

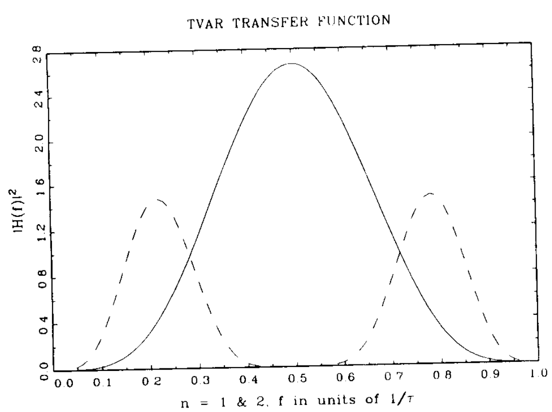

Figure 9. This is a plot of the squared transfer function of the impulse response function for TVAR. It is the function that multiplies the spectrum of the time deviations to obtain $\sigma_x^2(\tau)$. It is plotted hear for two values of τ (τ_o and $2\tau_o$ where $1/\tau_o = 1$). Note, the abscissa is linear and that the bandwidth and the amplitude of H(f) decrease as τ increases--similar to MVAR. As with MVAR, this is due to the software band-width change brought about by phase averaging. This transfer function, for a given τ, has repeat lobes into the higher Fourier frequencies indefinitely.

cut-off given by the Nyquist frequency $1/(2\tau_o)$. The sum of powers-of-two transfer functions, figure 10, shows a fairly flat band-pass up to the high-frequency end where there is greater frequency sensitivity--peaked at the Nyquist frequency, f_{Nyq}. Frequencies higher than the Nyquist frequency can be aliased into a TVAR computation up to the cut-off frequency, f_h. This points out the value of the general rule to have the sampling period equal to or less than $1/(2f_{Nyq})$. Then aliasing will not be a problem.

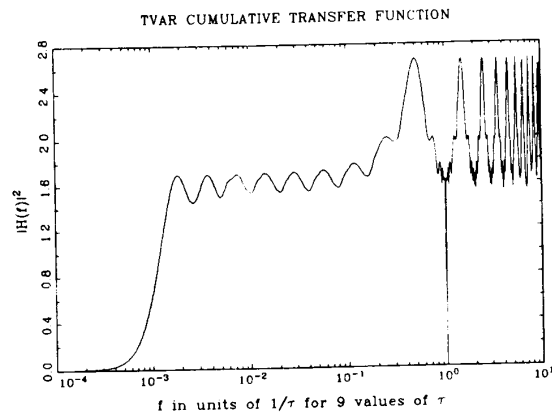

Figure 10. This figure shows the sum of the squared transfer functions for 9 values of τ (1, 2, 4, 8....256 x τ_o) for TVAR. We conclude that a $\sigma_x^2(\tau)$ plot for such a set of τ values gives a nearly constant response to Fourier frequency over slightly two decades but with increased sensitivity to Fourier frequencies at the Nyquist frequency, f_{Nyq}, and at 1/2 f_{Nyq}. Frequencies higher than $2f_{Nyq}$ can be aliased into the computation of TVAR. Frequencies will be aliased up to the measurement system cut-off frequency f_h. Maximum aliasing occurs at 3/2, 5/2, 7/2...x f_{Nyq}. A null occurs at twice the Nyquist frequency. Sensitivities at the non-aliased frequencies above $2f_{Nyq}$ are about the same as they are below the Nyquist frequency.

If the data sampling rate, τ_o, is greater than $1/(2f_{Nyq})$, then a TVAR plot will include Fourier energy from $1/(2n\tau_o)$ up to f_h with about equal sensitivity, except at the aliased values (3/2, 5/2, 7/2... X f_{Nyq} up to f_h). If f_h is equal to f_{Nyq}, then the TVAR cumulative transfer function looks very much like the MVAR cumulative transfer function.

Applications and Discussion

In the previous sections we have learned of three time domain statistical measures that are particularly appropriate for applications involving frequency standards, clocks and their associated measurement, comparison and distribution systems. We have also seen in some detail how these three time domain measures can be interpreted in the frequency domain. In this section we consider where each time domain measure is most appropriately applied. As we shall see, the selection of the appropriate time domain measure is a function of the types of noise which are characteristic of the process we are investigating as well as whether we want to study time stability or the frequency stability. We want to choose that measure which most clearly reveals the types and levels of noise involved in a particular application.

As we stated in the introduction, AVAR and MVAR were developed first, while TVAR is the newest member of our triad of statistical measures. In general terms, AVAR and MVAR are the measures to use when we are primarily interested in systems and devices where frequency is the quantity of interest, while TVAR is more appropriate where the quantity of interest is primarily time or phase.

Before we begin our discussion of specific applications, let's briefly review the five kinds of noise processes we are likely to encounter for the systems discussed in this paper. Although there are many ways to inventory these noise processes it is common in the time and frequency literature to list them as follows:

1. white noise PM (phase modulation)
2. flicker noise PM
3. white noise FM (frequency modulation)
4. flicker noise FM
5. random walk FM

Mathematically these noise processes have the power-law spectral density relationships shown in the Table 1. Table 2 shows the appropriate mathematical expression for each of the three time domain measures. Table 3 gives the coefficients needed to translate from the time domain to the frequency domain.

Figure 11 displays illustrative examples of the time variation, x(t), of these noise processes. As we proceed from type 1 through type 5 noise we notice that the amplitude variation with time grows increasingly more slowly. Generally speaking, for our applications, the physical explanation for this trend is as follows. The time variations with a f^{-3} and f^{-4} spectrum, for example, are often related to environmental factors such as temperature variations, mechanical shock, and path delay variations while the faster variations represented by f^{0} and f^{-1} processes are more likely related to internal characteristics of the device itself. Here, for example, we think of the noisy electronic components that make up the amplifying stages in a frequency standard.

As we have learned in previous sections AVAR, MVAR and TVAR have different characteristics in the frequency domain so it is not surprising that one time domain measure is better suited for one kind of noise process than another.

AVAR and MVAR are frequency-stability measures, and AVAR is particularly suited for measuring the intermediate to long-term stability of clocks and oscillators. MVAR is generally more suited for electronically

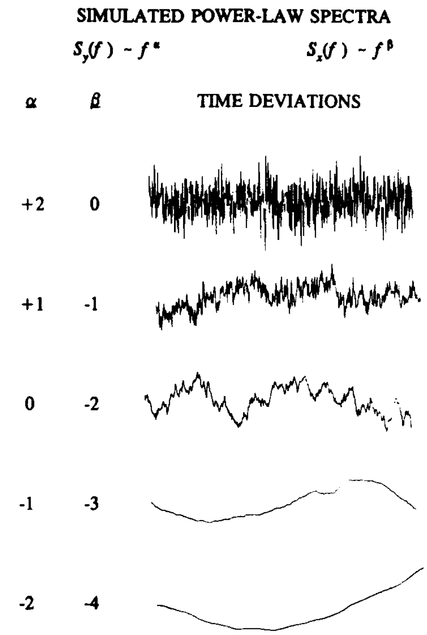

Figure 11. This is a display of the commonly occurring five power-law spectral density processes. These are often used as models for the time and/or frequency deviations in precision frequency sources.

generated noise processes and short-term frequency stability measurements. TVAR is a time-stability measure and is also suited for electronically generated noise processes. Both MVAR and TVAR are also sensitive to low frequency components often determined by environmental factors. TVAR is particularly suited for measuring the stability of time dissemination, comparison or measurement systems. It is also well suited as a measure of synchronization stability in telecommunication's networks.

We can see this from a different perspective by considering how AVAR, MVAR and TVAR vary with τ for our five dominant noise processes. Figures 12, 13 and 14 display the τ dependance for our three time domain measures. If we look at figure 12, we see that AVAR does not discriminate between white PM and flicker PM. This, as we said earlier, was one of the primary reasons for introducing MVAR, which as figure 13 shows, does discriminate between white PM and flicker PM. If we look at figure 14, we see that TVAR displays unambiguously the five noise types, as does MVAR, but that it also more clearly reveals the presence of white PM and flicker PM than does MVAR. This is, of course, the reason that TVAR was introduced since it "focuses" on the noise processes that are of most interest when we are making phase or time measurements.

Table 1. The power-law spectral density relationships for the five kinds of noise processes we are likely to encounter for the systems discussed in this paper.

NOISE TYPE	α	β	μ	μ'	η
White PM	2	0	-2	-3	-1
Flicker PM	1	-1	-2	-2	0
White FM	0	-2	-1	-1	1
Flicker FM	-1	-3	0	0	2
Random Walk FM	-2	-4	1	1	3

Where:

$\sigma_y^2(\tau) = a_\mu \tau^\mu$ $\quad S_y(f) = h_\alpha f^\alpha$

$\text{Mod}\ \sigma_y^2(\tau) = b_{\mu'} \tau^{\mu'}$ $\quad S_y(f) = h_\alpha f^\alpha$

$\sigma_x^2(\tau) = c_\eta \tau^\eta$ $\quad S_x(f) = h_\beta f^\beta$

$\sigma_x^2(\tau) = \frac{\tau^2}{3} \text{Mod}\ \sigma_y^2(\tau)$ $\quad S_y(f) = (2\pi f)^2 S_x(f)$

Table 2. The appropriate mathematical expression for each of the three time domain measures.

ABBREVIATION	NAME	EXPRESSION
AVAR	ALLAN VARIANCE	$\sigma_y^2(\tau) = \frac{1}{2}\langle(\Delta y)^2\rangle = \frac{1}{2\tau^2}\langle(\Delta^2 x)^2\rangle$
MVAR	MODIFIED ALLAN VARIANCE	$Mod\sigma_y^2(\tau) = \frac{1}{2\tau^2}\langle(\Delta^2 \bar{x})^2\rangle$
TVAR	TIME VARIANCE	$\sigma_x^2(\tau) = \frac{1}{6}\langle(\Delta^2 \bar{x})^2\rangle$

Table 3. The coefficients needed to translate from the time domain to the frequency domain.

NOISE TYPE	$S_y(f)$	$S_x(f)$
White PM	$\frac{(2\pi)^2}{3fh^2}\,[\tau^2\sigma_0^2(\tau)]\,f^2$	$\frac{1}{\tau_0 f_h}\,[\tau\,\sigma_x^2(\tau)]\,f^0$
Flicker PM	$\frac{(2\pi)^2}{A}\,[\tau^2\sigma_y^2(\tau)]\,f^1$	$\frac{3}{3.37}\,[\tau^0\sigma_x^2(\tau)]\,f^{-1}$
White FM	$2\,[\tau^1\sigma_y^2(\tau)]\,f^0$	$\frac{12}{(2\pi)^2}\,[\tau^{-1}\sigma_x^2(\tau)]\,f^{-2}$
Flicker FM	$\frac{1}{2\ell n2}\,[\tau^0\sigma_y^2(\tau)]\,f^{-1}$	$\frac{20}{(2\pi)^2 9\ell n2}\,[\tau^{-2}\sigma_x^2(\tau)]\,f^{-3}$
Random Walk FM	$\frac{6}{(2\pi)^2}\,[\tau^{-1}\sigma_y^2(\tau)]\,f^{-2}$	$\frac{240}{(2\pi)411}\,[\tau^{-3}\sigma_x^2(\tau)]\,f^{-4}$

$A = 1.038 + 3\ell n(2\pi f_h \tau)$

Table 4 shows in some detail the kinds of noise processes that are associated with various areas of application. The table also shows which of the three time domain measures is most appropriate for each particular application. We show the five different types of power-law-spectra and the ranges of applicability. These ranges include those for precision oscillators, for time distribution, network and comparison systems. From this table it is easy to see why MVAR and TVAR are better measures than AVAR for time distribution, network and comparison systems. On the other hand, AVAR estimator of frequency changes for white noise FM processes. This is important for commercial rubidium, cesium and for passive hydrogen masers. AVAR is also simpler to compute and is typically more intuitive than the other two variances. It nicely covers the range of applicability for precision oscillators except for the ambiguity problem in differentiating between white PM and flicker PM. This is only a problem for short-term stability in the case of active hydrogen masers and quartz crystal oscillators. In addition, $k\tau\sigma_y(\tau)$ is unbiased and useful measure of time error of prediction over the interval τ. The constant k depends on the power-law noise type, but is nominally equal to 1. [15]

The five power-law processes for the most part provide adequate modeling for time and frequency metrology. The higher values of alpha typically are used as models for the short-term stability of clocks and oscillators. The lower ends of the ranges are often appropriate models for the long-term stability of clocks and oscillators as well as for the time distribution, comparison, network and measurement systems. These lower values of alpha are often contaminated with diurnal and annual variations in these systems causing them to appear low-frequency dispersive. Some time comparison systems, such as GPS used in the common-view mode, are well modeled by white-noise PM in the day-to-day deviations. The bottom end of the variance ranges are those points where these variances are no longer convergent. If models were needed with lower values of alpha than those shown, then variances with higher order differences could be used, such as the Kolmogorov structure functions. [16] These three measures are convergent for the upper ranges of α and β ($\alpha > +2$ or $\beta > 0$), but only some of the Fourier transform relationships have not been worked out. [2] This is only because these models are not usual.

Sigma Tau Diagram

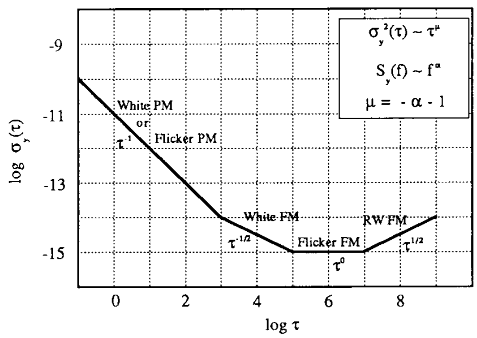

Modified Allan Variance
distinguishes White PM.

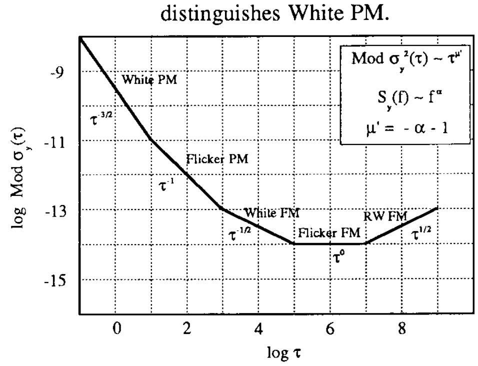

TVAR optimally estimates
time instability with White PM
and distinguishes other noise types.

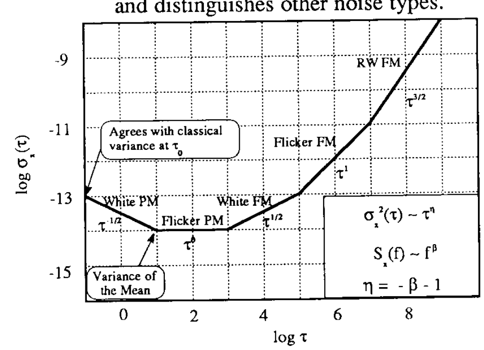

Figure 12, 13 and 14. These three figures are example plots for the square root of AVAR, MVAR, and TVAR, respectively. The sensitivity of these variances to the five power-law spectral density processes illustrated in figure 11 are depicted. Notice, that the slope on the $\sigma_y(\tau)$ plot is the same for white-noise PM as for flicker-noise PM (-1; ζ = -2). Notice also, that it is easier--visually--to distinguish between white phase, flicker phase and random walk phase on a $\sigma_x(\tau)$ plot than from a Mod$\sigma_y(\tau)$ plot. The slope changes are more dramatic to the eye. These three noise processes are particularly useful models for systems where time measurements are important.

Table 4. The kinds of noise processes that are associated with various areas of application. We also see which of the three time domain measures is most appropriate for each particular application.

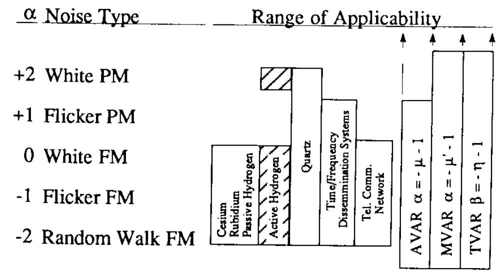

Conclusion and Summary

Over the last few years, the need for a measure of time stability has become apparent. A search of the literature reveals that the classical measures (standard deviation, mean and variance) have lead to confusing and often misinterpreted conclusions. A measure, TVAR, is shown to have the attributes needed for characterizing the random processes in systems where time stability or phase stability is important. TVAR was compared with and contrasted to the other two previously developed time-domain statistical measures. The need for the three measures is also explained.

In particular, we have discussed how these three measures are appropriate for frequency standards, clocks and their associated measurement and distribution systems. We have shown how these measures can be used to determine the five noise processes that dominate most of the systems of interest in this paper. However, we have also examined these measures from a frequency-domain point of view and have shown how each measure corresponds to a particular transfer function. This procedure reveals the way time domain measures treat the various noise processes from a frequency domain point of view. It also provides a link, for those more accustomed to working in the frequency domain, to the time domain measures which are frequently employed by the time and frequency community.

Acknowledgements

The authors are particularly grateful to the readers, David A. Howe and James A. Barnes, of this paper. We are also very appreciative to the telecommunication industries for their support of the research associated with this paper. We are also grateful for the continued support and encouragement in this project by Donald B. Sullivan.

References

[1] J. Rutman and F. L. Walls, "Characterization of Frequency Stability in Precision Frequency Sources," submitted to Proc. of the IEEE, 1991.

[2] L. S. Cutler and C. L. Searle, "Some Aspects of the Theory and Measurement of Frequency Fluctuations in Frequency Strandards," Proceedings of the IEEE, Vol. 54, No. 2, pp. 136-154, Feb. 1966.

[3] J. von Neuman, R. H. Kent, H. R. Bellinson, and B. I. Hart, "The Mean Square Successive Difference," Ann. Math. Stat. 12, 153-162 (1942).

[4] D. W. Allan, "Statistics of Atomic Frequency Standards," Proc. of the IEEE, 54, 221-231 (1966).

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

[6] J.A. Barnes, A.R. Chi, L.S. Cutler, D.J. Healey, D.B. Leeson, T.E. McGunigal, J.A. Mullen, Jr., W.L. Smith, R.L. Sydnor, R.F.C. Vessot, and G.M.R. Winkler, Characterization of Frequency Stability, IEEE Transactions on Instrumentation and Measurement, IM-20, No. 2, 105-120, 1971.

[7] D.W. Allan and J.A. Barnes, A Modified "Allan Variance" with Increased Oscillator Characterization Ability, Proceedings of the 35th Annual Frequency Control Symposium,, 470-475, 1981.

[8] P. Lesage and T. Ayi, "Characterization of Frequency Stability: Analysis of the Modified Allan Variance and Properties of Its Estimate," IEEE Trans. on Instrum. & Meas., IM-33, 332-336 (1984).

[9] L. G. Bernier, "Theoretical Analysis of the Modified Allan Variance," Proc. of the 41st Annual Frequency Control Symposium, 1987, pp. 116-121.

[10] D. W. Allan, D. D. Davis, J. Levine, M. A. Weiss, N. Hironaka, and D. Okayama, "New Inexpensive Frequency Calibration Service from NIST," Proc. of the 44th Annual Frequency Control Symposium, 1990, pp. 107-116.

[11] D. W. Allan, "Time and Frequency Metrology: Current Status and Future Considerations," submitted to Proc. of 5th European Frequency and Time Forum, March 13-15, Bescanson, France, 1991.

[12] D. B. Percival, "Reappraisal of Frequency Domain Techniques for Assessing Stability Measures," Proc. of the PTTI, 1987, pp. 69-80.

[13] J. Rutman, "Characterization of Phase and Frequency instabilities in Precision Frequency Sources: Fifteen Years of Progress," Proc. of the IEEE, Vol. 66, No. 9, pp. 1048-1075, September 1978.

[14] J. Rutman, "Characterization of Frequency Stability: A Transfer Function Approach and Its Application to Measurements via Filtering of Phase Noise," IEEE Trans. on I & M, Vol. IM-23, No. 1, March 1974, pp. 40-48.

[15] D.W. Allan, Time and Frequency (Time-Domain) Characterization, Estimation, and Prediction of Precision Clocks and Oscillators, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, UFFC-34, 647-654, 1987.

[16] W. C. Lindsey and C. M. Chie, "Theory of Oscillator Instability Based Upon Structure Functions," IEEE 64, 1652-1666 (1976).