THE MEASUREMENT OF FREQUENCY AND FREQUENCY STABILITY OF PRECISION OSCILLATORS

David W. Allan

Time and Frequency Division
National Bureau of Standards

ABSTRACT

The specification and performance of precision oscillators is a very important topic to the owners and users of these oscillators. This paper presents at the tutorial level some convenient methods of measuring the frequencies of precision oscillators -- giving advantages and disadvantages of these methods.

Conducting such measurements, of course, gives additional understanding into the performance of the given pair of oscillators involved. Further it is shown that by processing the data from the frequency measurements in certain ways, one may be able to state more general characteristics of the oscillators being measured. The goal in this regard is to allow the comparisons of different manufacturers' specifications and more importantly to help assess whether these oscillators will meet the standard of performance the user may have in a particular application.

The methods employed for measuring frequency are designed for state-of-the-art oscillators, and an effort has been made to allow for fairly simple, inexpensive, and/or commonly available componentry to be used in the measurement systems. The method for measuring frequency stability is basically that recommended by the IEEE subcommittee which wrote the paper "Characterization of Frequency Stability," IEEE Transactions on Instrumentation and Measurement, IM-20, No. 2, pp. 105-120, (May 1971).

INTRODUCTION

Precision oscillators play an important role, in high speed communications, navigation, space tracking, deep space probes and in numerous other important applications. In this paper I will review some precision methods to measure the frequency and frequency stability of precision oscillators. The paper will be tutorial in nature and will concentrate on fairly well established methods; however, it will present one apparently unexploited and useful method. I will first define some terms and some basic concepts that will be useful later on and then discuss four different ways of measuring frequency and frequency stability. Finally, I will discuss briefly some useful methods of analyzing the results--to more nearly maximize on the information that may be deduced from the data.

The typical precision oscillator, of course, has a very stable sinusoidal voltage output with a frequency ν and a period of oscillation τ , which is the reciprocal of the frequency, $\nu = 1/\tau$, as illustrated in Fig. 1. The goal is to measure the frequency and/or the frequency stability (instability is actually measured, but is often called, with little confusion, stability in the literature) of the fluctuations of the sinusoid. The voltage out of the oscillator may be modeled by equation 1:

$$V_1 = V_p \sin (2\pi\nu_1 t) . \tag{1}$$

Of course, one sees that the period of this oscillation is the number of seconds per cycle or the inverse of the frequency in cycles per second. Naturally, fluctuations in frequency correspond to fluctuations in the period. Almost all frequency measurements, with very few exceptions, are measurements of phase or of the period fluctuations in an oscillator, not of frequency, even though the frequency may be the readout. As an example, most frequency counters sense the zero (or near zero) crossing of the sinusoidal voltage, which is the point at which the voltage is the most sensitive to phase fluctuations.

One must also realize that any frequency measurement always involves two oscillators. In some instances the oscillator is in the counter. One can never measure purely only one oscillator. In some instances one oscillator may be enough better than the other that the fluctuations measured may be considered essentially those of the latter. However, in general because frequency measurements are always dual, it is useful to define:

$$y\ (t) = \frac{\nu_1 - \nu_o}{\nu_o} \tag{2}$$

as the fractional frequency deviation of say oscillator one, ν_1, with respect to a reference oscillator ν_0 divided by the nominal frequency ν_0. Now, y(t) is a dimensionless quantity and useful in describing oscillator and clock performance; e.g. the time fluctuations, x(t), of an oscillator over a period of time t are simply given by:

$$x(t) = \int_0^t y(t)\, dt \qquad (3)$$

Since it is impossible to measure instantaneous frequency, any frequency or fractional frequency measurement always involves some sample time, τ --some time window through which the oscillators are observed; whether it's a picosecond, a second, or a day, there is always some sample time. So when determining a fractional frequency, y(t), in fact what happens in the device is that the time fluctuation is being measured say starting at some time t and again at a later time, $t + \tau$. The difference in these two time fluctuations, divided by τ gives the average fractional frequency over that period τ :

$$y(t, \tau) = \frac{x(t + \tau) - x(t)}{\tau} . \qquad (4)$$

Tau, τ, may be called the sample time or averaging time; e.g. it may be determined by the gate time of a counter.

What happens in many cases is that one samples a number of cycles of an oscillation during the preset gate time of a counter; after the gate time has elapsed the counter latches the value of the number of cycles so that it can be read out, printed or stored in some other way, and then there is a delay time for such processing of the data before the counter arms and starts again on the next cycle of the oscillation. During the delay time or process time information is lost. We have chosen to call it dead time and in some instances it becomes a problem. Unfortunately it seems that in typical oscillators the effects of dead time hurt the most where it is the hardest to avoid. In other words, for times that are short compared to a second, where it is very difficult to avoid dead time, that is usually where whether you do or do not have dead time makes a difference in the data. Typically for common oscillators, if the sample time is long compared to a second, the dead time makes little difference except in data analysis unless it is excessive [3].

SOME METHODS OF MEASUREMENT

In reality of course, the sinusoidal output of an oscillator is not pure; but it contains noise fluctuations as well. This section deals with the measurement of these fluctuations to determine the quality of a precision signal source.

I will describe four different methods of measuring the frequency fluctuations in precision oscillators.

A. The first is illustrated in Fig. 2. The signal from an oscillator under test is fed into one port of a mixer. The signal from a reference oscillator is fed into the other port of this mixer. The signals are in quadrature, that is, they are 90 degrees out of phase so that the average voltage out of the mixer is nominally zero, and the instantaneous voltage corresponds to phase fluctuation rather than to the amplitude fluctuations between the two signals. The mixer is a key element in the system. The advent of the Schottky barrier diode was a significant breakthrough in making low noise precision stability measurements and in all four measurement methods described below the double balanced Schottky barrier diode mixer is employed. The output of this mixer is fed through a low pass filter and then amplified in a feedback loop, causing the voltage controlled oscillator (reference) to be phase locked to the test oscillator. The time constant and gain are adjusted such that a very loose phase lock condition exists. Caution: the attack time is not the time constant of the RC network shown.

The attack time is the time it takes the servo system to make 70% of its ultimate correction after being slightly disturbed. The attack time is equal to the inverse of π times the servo bandwidth. If the attack of the loop is about a second then the voltage fluctuation will be proportional to the phase fluctuation for sample times shorter than the attack time or for Fourier frequencies greater than about 1 Hz. Depending on the quality of the oscillators involved, the amplification used may be from 40 to 80 dB via a good low noise amplifier, and in turn this signal can be fed to a spectrum analyzer, for example, to measure the Fourier components of the phase fluctuation. This system of frequency-domain analysis has been well documented in the literature [1,2,3] and has proven very useful at NBS; specifically, it is of use for sample times shorter than one second or for Fourier frequencies greater than 1 Hz in analyzing the characteristics of an oscillator. It is also specifically very useful if you have discrete side bands such as 60 Hz or detailed structure in the spectrum.

B. The second system (shown in Fig. 3) is essentially the same as in Fig. 2 except that in this case the loop is in a tight phase lock condition; i.e. the attack time of the loop should be of the order of a few milliseconds. In such a case, the phase fluctuations are being integrated so that the voltage output is proportional to the frequency fluctuations between the two oscillators and is no longer proportional to the phase fluctuations for sample times longer than the attack time of the loop. The bias box is used to simply adjust the voltage on the varicap so that you are at a tuning point that is fairly linear and of a reasonable value. Typically, the oscillators we have used at NBS are about 1 part in 10^9 per volt. The voltage fluctuations prior to the bias box (biased slightly away from zero) are fed to a voltage to frequency converter which in turn is fed to a frequency counter where one may read out the frequency fluctuations with great amplification of the instabilities between this pair of oscillators. The frequency counter data are logged with a printer or some other data logging device. The coefficient of the varicap and the coefficient of the voltage to frequency converter are used to determine the fractional frequency fluctuations, y_i, between the oscillators, where i denotes the i^{th} measurement as shown in Fig. 3. The sensitivity of the system that we have set up at NBS is about a part in 10^{14} per Hz resolution of the frequency counter, so one has excellent precision capabilities with this system.

The advantages and disadvantages of this type of tight phase lock system are as follows:

ADVANTAGES: The component cost is not too expensive unless one does not have a voltage controllable oscillator. Voltage to frequency converters can now be purchased for about \$150.00. Most people involved with time and frequency measurements already have counters and oscillators and so I have not entered these as expenses. In addition, good bandwidth control is obtainable with this system and the precision is adequate to measure essentially any of the state-of-the-art oscillators. The sample time can be of the order of a second or longer; it is difficult to go shorter than one second or an interaction will occur with the attack time of the tight phase lock loop. The dead time can be small; in fact, if you have a very fast counter, that is a counter which can scan the data more quickly than the attack time of the loop, the dead time will be negligible.

DISADVANTAGES: An oscillator that is controllable is necessary. For the price of increased precision, one has increased complexity over simply measuring with a direct frequency counter. The varicap tuning curve is nonlinear and so that curve must be calibrated and doing so is sometimes a bit of a nuisance. For that reason and some other reasons it is not useful in measuring the absolute frequency difference between the pair of oscillators involved in the measurement. The system is basically conducive to measuring frequency stability.

C. Beat Frequency Method. The next system I would like to describe is what is called a hetrodyne frequency measuring method or beat frequency method. The signal from two independent oscillators are fed into the two ports of a double balanced mixer as illustrated in Fig. 4. The difference frequency or the beat frequency out, ν_b, is obtained as the output of a low pass filter which follows the mixer. This beat frequency is then amplified and fed to a frequency counter and printer or some recording device. The fractional frequency can simply be obtained by dividing ν_b, by the nominal carrier frequency ν_o.

ADVANTAGES: This system has excellent precision; one can measure essentially all state-of-the-art oscillators. The component cost is quite inexpensive.

DISADVANTAGES: The sample time must be equal to or greater than the beat period, and for good tunable quartz oscillators this will be of the order of a few seconds; i.e. typically, it is difficult to have a sample time shorter than a few seconds. The dead time can be a problem for this measurement system because it will be equal to or greater than the beat period unless, for example, one uses a second counter which starts when the first one stops. Observing the beat frequency only is insufficient information to tell whether one oscillator is high or low in frequency with respect to the other one--a significant disadvantage for making absolute frequency measurements. However, it is often not difficult to gain this additional information to determine the sign (+ or -) of the beat frequency. The frequencies of the two oscillators must be different.

D. Dual Mixer Time Difference System. The last system is one that has just recently been developed at NBS* that

*Dr. Costain informed me that Herman Daams has developed a similar system at NRC.

shows some significant promise. A block diagram is shown in Fig. 5. In preface it should be mentioned that if the time or the time fluctuations can be measured directly an advantage is obtained over just measuring the frequency. The reason being that you can calculate the frequency from the time without dead time as well as know the time behavior. The reason in the past that frequency has not been inferred from the time for sample times of the order of several seconds and shorter is that the time difference between a pair of oscillators operating as clocks could not be measured with sufficient precision (commercially the best that is available is 10^{-10} seconds). The system described in this section demonstrated a precision of 10^{-13} seconds with the potential of doing about 10^{-14} seconds. Such a precision opens the door to making time measurements as well as frequency and frequency stability measurements for sample times as short as a few milliseconds as well as for longer sample times and all without dead time. In Fig. 5, oscillator 1 could be considered under test and oscillator 2 could be considered the reference oscillator. These signals go to the ports of a pair of double balanced mixers. Another oscillator with separate symmetric buffered outputs is fed to the remaining other two ports of the pair of double balanced mixers. This common oscillator's frequency is offset by a desired amount from the other two oscillators. In which case two different beat frequencies come out of the two mixers as shown. These two beat frequencies will be out of phase by an amount proportional to the time difference between oscillator 1 and 2--excluding the differential phase shift that may be inserted; and will differ in frequency by an amount equal to the frequency difference between oscillators 1 and 2. Now this system is also very useful in the situation where you have oscillator 1 and oscillator 2 on the same frequency. The hetrodyne or beat frequency method, in contrast, cannot be used if both oscillators are on the same frequency. Quite often it is the case with Atomic standards (Cesium, Rubidium and Hydrogen frequency standards) that oscillator 1 and 2 will nominally be on the same frequency.

Illustrated at the bottom of Fig. 5 is what might be represented as the beat frequencies out of the two mixers. A phase shifter may be inserted as illustrated to adjust the phase so that the two beat rates are nominally in phase; this adjustment sets up the nice condition that the noise of the common oscillator tends to cancel when the time difference is determined in the next step--depending on the level and the type of noise as well as the sample time involved. After amplifying these beat signals, the start port of a time interval counter is triggered with the zero crossing of one beat and the stop port with the zero crossing of the other beat. If the phase fluctuations of the common oscillator are small during this interval as compared to the phase fluctuations between oscillators 1 and 2 over a full period of the beat frequency the noise of the common oscillator is insignificant in the measurement noise error budget, which means the noise of the common oscillator can in general be worse than that of either oscillator 1 or 2 and still not contribute significantly. By taking the time difference between the zero crossings of these beat frequencies, what effectively is being measured is the time difference between oscillator 1 and oscillator 2, but with a precision which has been amplified by the ratio of the carrier frequency to the beat frequency over that normally achievable with this same time interval counter. The time difference x(i), for the i^{th} measurement between oscillators 1 and 2 is given by equation 5:

$$x(i) = \frac{\Delta t(i)}{\tau\nu} - \frac{\phi}{2\pi\nu} + \frac{n}{\nu} \tag{5}$$

where $\Delta t(i)$ is the i^{th} time difference as read on the counter, τ is the beat period, ν is the nominal carrier frequency, ϕ is the phase delay in radians added to the signal of oscillator 1, and n is an integer to be determined in order to remove the cycle ambiguity. It is only important to know n if the absolute time difference is desired; for frequency and frequency stability measurements and for time fluctuation measurements, n may be assumed zero unless one goes through a cycle during a set of measurements. The fractional frequency can be derived in the normal way from the time fluctuations.

$$y_{1,2}(i,\tau) = \begin{cases} \dfrac{\nu_1(i,\tau) - \nu_2(i,\tau)}{\nu} \\[2ex] \dfrac{x(i+1) - x(i)}{\tau} \\[2ex] \dfrac{\Delta t(i+1) - \Delta t(i)}{\tau^2\,\nu} \end{cases} \tag{6}$$

In equations (5) and (6), the assumptions are made that the transfer or common oscillator is set at a lower frequency than oscillators 1 and 2, and that the beat $\nu_1 - \nu_2$ starts and $\nu_2 - \nu_o$ stops the time interval counter. The sample time by appropriate calculation can be any integer multiple of τ:

$$y_{1,2}(i, m\tau) = \frac{x(i+m) - x(i)}{m\tau}, \qquad (7)$$

where m is any positive integer. If needed, τ can be made to be very small by having very large beat frequencies. In the system set up at NBS the common or transfer oscillator was replaced with a low phase noise synthesizer, which derived its basic reference frequency from oscillator 2. In this set up the nominal beat frequencies are simply given by the amount the output frequency of the synthesizer is offset from ν_2. Sample times as short as a few milliseconds were easily obtained. Logging the data at such a rate can be a problem without special equipment, e.g. magnetic tape. In the NBS set up, a computing counter was used with a processing time of about 1.5 ms, and sample time stabilities were observed for 2 ms and longer (see appendix for some computing counter program possibilities).

ADVANTAGES: If the oscillators, including the transfer oscillator, and a time interval counter are available, the component cost is fairly inexpensive ($500, most of which is the cost of the phase shifter). The measurement system bandwidth is easily controlled (note, that this should be done in tandum with both low pass filters being symmetrical). The measurement precision is such that one can measure essentially all state-of-the-art oscillators. For example, if the oscillators are at 5 MHz, the beat frequencies are 0.5 Hz, and the time interval counter employed has a precision of 0.1 μs, then the potential measurement precision is 10^{-14}s (10 femto seconds) for τ = 2s; other things will limit the precision such as noise in the amplifiers. As has been stated above, there is no dead time which is quite convenient for very short sample times (of the order of milliseconds). Dead time problems are difficult to avoid in this region. One obtains as long a sample time as is desired. This is determined essentially by the beat period or multiple of the same. If one replaces the common oscillator by a synthesizer then the beat period may be selected very conveniently. The synthesizer should have fairly low phase noise to obtain the maximum precision from the system. The system measures time difference rather than frequency and hence has that advantage. One may calculate from the data both the magnitude and the sign of the frequency difference. This system, therefore, allows the measurement of time fluctuations as well as time difference, and the calculation of frequency fluctuations as well as absolute frequency differences between the two oscillators in question. The system may be calibrated and the system noise be measured by simply feeding a signal from one oscillator symmetrically split two ways to replace oscillators 1 and 2.

DISADVANTAGES: The system is somewhat more complex than the others. Because of the low frequency beats involved, precautions must be taken to avoid ground loop problems; there are some straight forward solutions; e.g. in the NBS system a saturated amplifier followed by a differentiator and isolation transformer worked very well in avoiding ground loops. Buffer amplifiers are needed because the mixers present a dynamic load to the oscillator--allowing the possibility of cross-talk. The time difference reading is modulo the beat period. For example, at 5 MHz there is a 200 nanosecond per cycle ambiguity that must be resolved if the absolute time difference is desired; this ambiguity is usually a minor problem to resolve for precision oscillators.

As an example of the system's use, Fig. 6 illustrates a plot of a strip chart recording of a digital to analog output of the significant digits from the time interval counter between a quartz oscillator and a high performance commercial cesium oscillator. In other words this is a plot of the time fluctuations between these two oscillators as a function of time. The high frequency fluctuations (over fractions of a second) would most probably be those between the quartz oscillator and the quartz oscillator in the cesium servo system. The low frequency fluctuations (over seconds) would most probably be those induced by the cesium servo in its effort to move the frequency of its quartz oscillator to the natural resonance of the cesium atom--causing a random walk of the time fluctuations for sample times longer than the servo attack time.

SOME METHODS OF DATA ANALYSIS

Given a set of data of the fractional frequency or time fluctuations between a pair of oscillators, it is useful to characterize these fluctuations with reasonable and tractable models of performance. In so doing for many kinds of oscillators it is useful to consider the fluctuations as those that are random (may only be predicted statistically) and those that are non-random (e.g. systematics--those that are environmentally induced or those that have a causal effect that can be determined and in many cases can be predicted).

A. Non-random Fluctuations

Non-random fluctuations are usually the main cause of departure from "true" time or "true" frequency.

If for example one has the values of the frequency over a period of time and a frequency offset from nominal is observed, one may calculate directly that the time fluctuations will depart as a ramp (see Fig. 7). If the frequency values show some linear drift then the time fluctuations will depart as a quadratic. I mention this because in almost all oscillators the systematics, as they are sometimes called, are the primary cause of time and/or frequency departure. A useful approach to determine the value of the frequency offset is to calculate the simple mean of the set, or for determining the value of the frequency drift by calculating a linear least squares fit to the frequency. A precaution is to not calculate a least squares quadratic fit to the phase or time departure--such is not as efficient an estimator of the frequency drift for most oscillators.

B. Random Fluctuations:

After calculating or estimating the systematic or non-random effects of a data set, these may be subtracted from the data leaving the residual random fluctuations. These can usually be best characterized statistically. It is often the case for precision oscillators that these random fluctuations may be well modeled with power law spectral densities, [4,5,6,7]:

$$S_y(f) = h_\alpha f^\alpha, \tag{8}$$

where $S_y(f)$ is the one-sided spectral density of the fractional frequency fluctuations, f is the Fourier frequency at which the density is taken, h_α is the intensity coefficient, and α is a number modeling the most appropriate power law for the data. It has been shown [3,4,5,8], that in the time domain one can nicely represent a power law spectral density process using a well defined time-domain stability measure, $\sigma_y(\tau)$, which I will explain later. For example, if you have a log $\sigma_y(\tau)$ versus log τ diagram and you observe a particular slope--call it μ --over certain regions of sample time, τ ; this slope has a correspondance to a power law spectral density or a set of the same with some amplitude coefficient h_α , i.e. $\mu = -\alpha-1$ for $-3 < \alpha < 1$ and $\mu \simeq -2$ for $1 \leq \alpha$. Further, a correspondance exists between h_α and the coefficient for $\sigma_y(\tau)$. These coefficients and relationships have been calculated and appear in the literature [2,3,4]. The transformation for some of the more common power law spectral densities has been tabulated, [2,3,4]--making it quite easy to transform the frequency stability as may have been modeled in the time-domain over to the frequency domain and vise-versa. Some examples of some processes modeled by power law spectra that have been simulated by computer are shown in Fig. 8. In descending order these have been named, white noise, flicker noise, random walk, and flicker walk (the ω in Fig. 8 is angular Fourier frequency, $\omega = 2\pi f$). In Fig. 9 are plotted the actual data of the Atomic Time Scale of the National Bureau of Standards versus International Atomic Time (TAI) over a four year period. A least squares fit to the frequency drift has been subtracted from these data. The plot then is just the time fluctuations of the AT(NBS) scale with respect to TAI. There is a peak-to-peak deviation of about 6 microseconds. Figure 10 shows a plot of the same thing for the United States Naval Observatory atomic time scale versus TAI over the same four year period, and again a least squares fit to the frequency drift has been subtracted from the data. The peak-to-peak fluctuations are again about 6 microseconds. Figure 11 is a plot of the residual time fluctuations between a high performance cesium standard and our primary frequency standard, NBS-5, over about one-half day. The peak-to-peak fluctuations in this case are less than a nanosecond. Just by visual comparison of Figures 9, 10 and 11 with the simulated noises shown in Figure 8 indicates that these random processes are not white noise--hence the need for better frequency stability characterization.

Suppose now that you are given the time or frequency fluctuations between a pair of precision oscillators measured, for example, by one of the techniques outlined above, and you wish to perform a stability analysis. Let this comparison be depicted by Fig. 12. The minimum sample time is determined by the measurement system. If the time difference or the time fluctuations are available then the frequency or the fractional frequency fluctuations may be calculated from one period of sampling to the next over the data length as indicated in Fig. 12. Suppose further there are M values of the fractional frequency, y_i. Now there are many ways to analyze these data. Historically, people have typically used the standard deviation equation shown in Fig. 12, $\sigma_{\text{std. dev}}(\tau)$, where $\bar{y}$ is the average fractional frequency over the data set and is subtracted from each value of y_i before squaring, summing and dividing by the number of values minus one, (M-1), and taking the square root to get the standard deviation. At NBS, we have studied what happens to the standard deviation when the data set may be characterized by power law spectra which are more dispersive than classical white noise frequency fluctuations. In other words, if you have flicker noise or any other non white noise frequency deviations, one may ask what happens to the standard deviation for that data set. In fact, one can show that the standard deviation is a function of the number of data points in the set, of the dead time, and of the measurement system bandwidth (5,9). For example, using as a model flicker noise frequency modulation, as the number of data points increase, the standard deviation monotonically increases without limit. Some statistical measures have been developed which do not depend upon the data length and which are readily usable for characterizing the random fluctuations in precision oscillators (2-5,9). An IEEE subcommittee on Frequency Stability has recommended a particular variance taken from the set of useful variances developed, and an experimental estimation of the square root of this particular variance is shown as the bottom right equation in Fig. 12. This equation is very easy to implement experimentally as you simply add up the squares of the differences between adjacent values of y_i, divide by the number of them and by two, take the square root and you then have the quantity which the IEEE subcommittee has recommended for specification of stability in the time domain.

One would like to know how $\sigma_y(\tau)$ varies with the sample time, τ. A simple trick that one can use, that is very useful if there is no dead time, is to average y_1 and y_2 and call that $\bar{y}_1$ averaged over 2τ, then average y_3 and y_4 and call that $\bar{y}_2$ as averaged over 2τ, etc., and finally apply the same equation as before to get $\sigma_y(2\tau)$. One can repeat this process for other desired integer multiples of τ and from the same data set be able to generate values for $\sigma_y(m\tau)$ as a function of $m\tau$ from which one may be able to infer a model for the process that is characteristic of this pair of oscillators. If you have dead time in the measurements you cannot average adjacent pairs in an unambiguous way to simply increase the sample time. You have to retake the data for each new sample time--often a very time consuming task. This is another instance where dead time can be a problem.

How the classical variance (standard deviation squared) depends on the number of samples is shown in Fig. 13. Plotted is the ratio of the standard deviation squared for N samples to the standard deviation squared for 2 samples ($\sigma^2(2, \tau)$ is the same as $\sigma_y^2(\tau)$). One can see the dependence of the standard deviation with the number of samples for various kinds of power law spectral densities commonly encountered as reasonable models for many important precision oscillators. Note, $\sigma^2_y(\tau)$ has the same value as the classical variance for the classical noise case (white noise FM). One main point of Fig. 13 is simply to show that with the increasing data length the standard deviation of the common classical variance is not well behaved for the kinds of noise processes that are very often encountered in most of the precision oscillators of interest.

Figure 14 is an actual $\sigma_y(\tau)$ versus τ plot for a rubidium standard that was analyzed at the Bureau. One observes apparent white noise FM with the slope of $\tau^{-1/2}$ and then flicker noise frequency modulation, τ^0; and some random walk FM for sample times of the order of a tenth of a day and longer. Having this time-domain analysis, one can use the equations and the tables mentioned before to transform to the frequency domain, $S_y(f)$ versus Fourier frequency f, and this transformation is plotted in Fig. 15. An equation which shows directly the mapping for a model that is often used for cesium decives for sample times longer than 1 second is given by the following pair of equations:

$$S_y(f) = h_o + h_{-1}f^{-1} \tag{9}$$

$$\sigma_y^2(\tau) = \frac{h_o}{2} + 2\ \ln 2\ h_{-1} \tag{10}$$

The h_o term in each case is due to the white noise FM fundamentally induced by the shot noise in the cesium beam. The second term is flicker noise FM (flicker floor) that seems to always appear as a reasonable model for cesium as well as other standards. It does not have a well understood cause. As an example of equation (9) and (10), suppose from a $\sigma_y(\tau)$ versus τ plot we determined that

$\sqrt{h_{o/2}} = 2 \times 10^{-12}[s]^{1/2}$ and $\sqrt{2\ \ln\ 2h_{-1}} = 1 \times 10^{-14}$ as

for one comparison made between the NBS primary frequency standards, NBS-4 and NBS-5, then $h_o = 4 \times 10^{-24}$ and $h_{-1} = 7.2 \times 10^{-29}$.

If the frequency drift is not subtracted from the data then the $\sigma_y(\tau)$ versus τ plot as shown in Fig. 15, takes on a τ^{+1} behavior. Often such is the case with quartz crystal oscillators. The equation relating $\sigma_y(\tau)$ and the drift, D, is as follows:

$$\sigma_y(\tau) = \frac{D\tau}{\sqrt{2}} \qquad (11)$$

where D has the dimensions of fractional frequency per unit of τ , i.e. if τ is in days then D could be, for example, 10^{-10} per day. Suppose also that the data contain discrete side bands such as 60 Hz then the $\sigma_y(\tau)$ versus τ diagram may appear as shown in Fig. 17. The model for this figure, calculated by Sam Stein, was for the situation where the white phase noise power in a 1 KHz bandwidth was equal to the power in the 60 Hz side bands.

In Fig. 18 I have used the dual mixer time difference measuring system in order to observe the $\sigma_y(\tau)$ versus τ behavior for a high performance cesium standard versus a quartz crystal oscillator. The plot contains a lot of information. The measurement noise of the dual mixer system is indicated. One can see the short term stability performance of the quartz oscillator in the cesium versus the comparison quartz oscillator (Diana). One can see a little bit of 60 Hz present as indicated by the humps at 1/2 and 3/2 of τ = 1/60 Hz. One observes the attack time of the servo in the cesium electronics perturbing the short term stability of the quartz oscillator and degrading it to the level of the shot noise of the cesium resonance. The white noise frequency modulation characteristic then becomes the predominant power law causing $\sigma_y(\tau)$ to improve as $\tau^{-1/2}$ until the flicker floor of the quartz crystal oscillator (Diane), in this case 6 parts in 10^{13}, becomes the predominant noise source. Thus using this particular measurement system I was able to well characterize for sample times of a few milliseconds all the way out to 1000 seconds, the stability characteristics of this particular pair of oscillators. Longer sample times are of course easily achievable.

CONCLUSIONS

Some inexpensive (less than $1000) methods of precisely measuring the time difference, time fluctuations, frequency difference, and frequency fluctuations between a pair of state-of-the-art time and/or frequency standards have been reviewed or introduced. One novel method introduced demonstrated the capability of measuring all four of the above, plus being able to cover an impressive segment of sample times ($\tau \geq$ few milliseconds) with a time difference precision of better than 1 picosecond. Fraction frequency instabilities due to the noise in this novel measurement method were demonstrated to be less than one part in 10^{16} for $\sigma_y(\tau \geq 2 \times 10^4$ s).

Also reviewed were some efficient methods of data analysis--which allow one to gain insight into models that would characterize both the random and non-random deviation between a pair of frequency standards. A specific example was shown demonstrating the time domain fractional frequency stability, $\sigma_y(\tau)$, between two state-of-the-art commercial standards, i.e. a quartz oscillator and a high performance cesium standard for 2 ms $\leq \tau \leq 10^3$ s.

ACKNOWLEDGEMENTS

Many stimulating discussions preceeded the writing and experimental results, as well as influenced the content of the material in the text. To all those so involved I express gratitude. In particular I wish to thank Dr. Fred L. Walls for his contributions. In addition I wish to thank Jorge Valega and Howard Machlan for instrumentation assistance and data processing, respectively.

APPENDIX:
COMPUTING COUNTER PROGRAM

A. The following program may be used in a computing counter, is useful in determining the fractional frequency stability, $\sigma_y(\tau)$, and is unique as compared with other similar types of programs to determine stability in that it does so with no dead time. The following program actually determines the root-mean-square second difference, $(\Delta^2(\Delta t))_{rms}$, of the time difference readings between a pair of clocks or oscillators, and therefore complements very nicely the dual mixer time difference measurement system described in the text. The fractional frequency stability may be calculated from computer program results as follows:

$$\sigma_y(\tau) = \frac{1}{\sqrt{2}\,\tau^2\nu}\,(\Delta^2(\Delta t))_{rms} \tag{A1}$$

If additional programing steps were available, of course one could program the computing counter to calculate an estimate of $\sigma_y(\tau)$ directly. Following is the program procedure to generate $(\Delta^2(\Delta t))_{rms}$:

1. clear x
2. $c \rightleftarrows x$
3. Plug-in
4. $a \rightleftarrows x$
5. Plug-in
6. $a \rightleftarrows x$
7. $\overrightarrow{a\ x\ y}$
8. -(subtract)
9. $b \rightleftarrows x$
10. Xfer Program
11. Plug-in
12. $a \rightleftarrows x$
13. $\overrightarrow{a\ x\ y}$
14. -(subtract)
15. $b \rightleftarrows x$
16. $\overrightarrow{b\ x\ y}$
17. - (subtract)
18. $\overrightarrow{xy}$
19. x (multiply)
20. $\overrightarrow{c\ x\ y}$
21. + (add)
22. $c \rightleftarrows x$
23. Repeat
24. Xfer Program
25. $c \rightleftarrows x$
26. $\overrightarrow{N\ x\ y}$
27. $\div$ (divide)
28. $\sqrt{x}$
29. Display x
30. Pause

The confidence of the estimate will improve approximately as the square root of the number of times (N) the sub loop is repeated as preset by the programmer [10].

REFERENCES

[1] D. G. Meyer, "A test set for measurement of phase noise on high quality signal sources," IEEE Trans. on Instrumentation and Measurement, IM-19, No. 4, pp. 215-227, Nov. 1970.

[2] J. H. Shoaf, D. Halford, A. S. Risley, "Frequency stability specification and measurement: High frequency and microwave signals," NBS Tech. Note 632, Jan. 1973.

[3] David W. Allan, John H. Shoaf, and Donald Halford, "Statistics of time and frequency data analysis," Chapter 8 of NBS Monograph 140, pp. 151-204, May 1974.

[4] J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen, W. L. Smith, R. Sydnor, R. F. C. Vessot, G.M.R. Winkler, "Characterization of frequency stability," NBS Technical Note 394, Oct. 1970; or IEEE Trans. on Instr. and Meas., Vol. IM-20, No. 2, pp. 105-120, May 1971.

[5] David W. Allan, "Statistics of atomic frequency standards," Proc. IEEE, Vol. 54, No. 2, pp. 221-230, Feb. 1966.

[6] Helmut Hellwig, "A review of precision oscillators," Proc. 6th PTTI, Washington DC. To be published.

[7] F. L. Walls and A. E. Wainwright, "Measurements of the short-term stability of quartz crystal resonators - a window on future developments in crystal oscillators," Proc. 6th PTTI, Washington DC. To be published.

[8] R. F. C. Vessot, L. Mueller, and J. Vanier, "The specification of oscillator characteristics from measurements made in the frequency domain," Proc. IEEE, Vol. 54, No. 2, pp. 199-206, Feb. 1966.

[9] James A. Barnes, "Atomic time keeping and the statistics of precision signal generators," Proc. IEEE, Vol. 54, No. 2, pp. 207-220, Feb. 1966.

[10] Patrick LeSage and Claude Audoin, "Characterization of frequency stability: Uncertainty due to the finite number of measurements," IEEE Trans. on Instr. and Meas., Vol. IM-22, No. 2, pp. 157-161, June 1973.

PERIOD OF AN OSCILLATOR

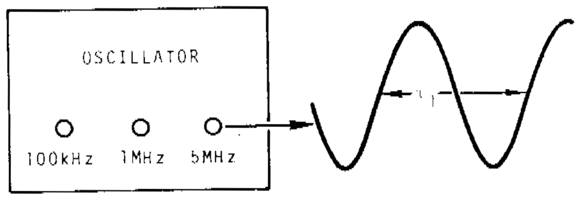

FIG. 1 Depiction of the sinusoidal voltage output from a precision oscillator, e.g. using quartz, rubidium, cesium or hydrogen as the frequency determining element. If the frequency out is ν_1, the the period of oscillation is $\tau_1 = 1/\nu_1$

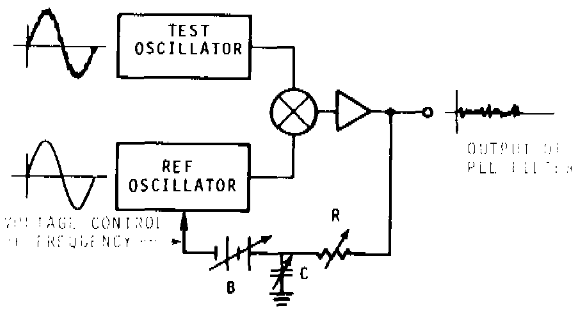

FIG. 2 A phase (or time) fluctuation measurement system. The reference oscillator is loosely phase-locked to the test oscillator--attack time is about 1 second. The reference and test oscillators are fed into the two ports of a Schottky barrier diode double balanced mixer whose output is fed through a low pass filter and low noise amplifier, hence to an RC network, a battery bias box and to the varicap of the reference oscillator. The instantaneous output voltage of the phase locked loop (PLL) following the low noise amplifier will be proportional to the phase or time fluctuations between the two oscillators.

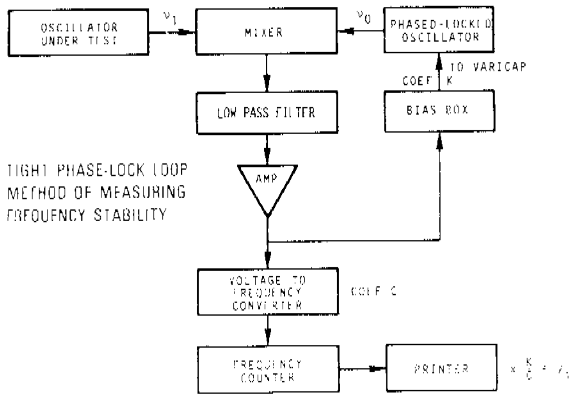

FIG. 3 A frequency fluctuation measurement system. The attack time of the phase lock loop in this case is much less than a second. The amplifier (AMP) output voltage fluctuations for sample times significantly larger than the servo loop attack time will be proportional to the frequency fluctuations between the oscillators.

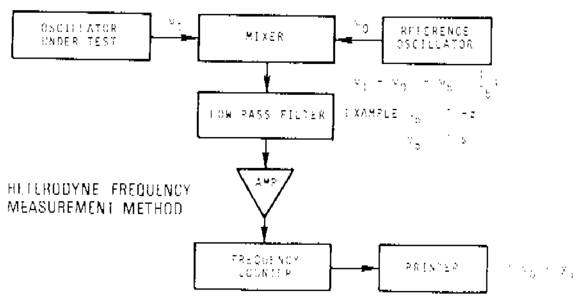

FIG. 4 A frequency and frequency fluctuation measurement system. The difference frequency, $|\nu_1 - \nu_0|$ is measured with a frequency counter. A counter measuring the period (or multiple period) of the beat (difference) frequency could equivalently be used.

DUAL MIXER TIME DIFFERENCE SYSTEM

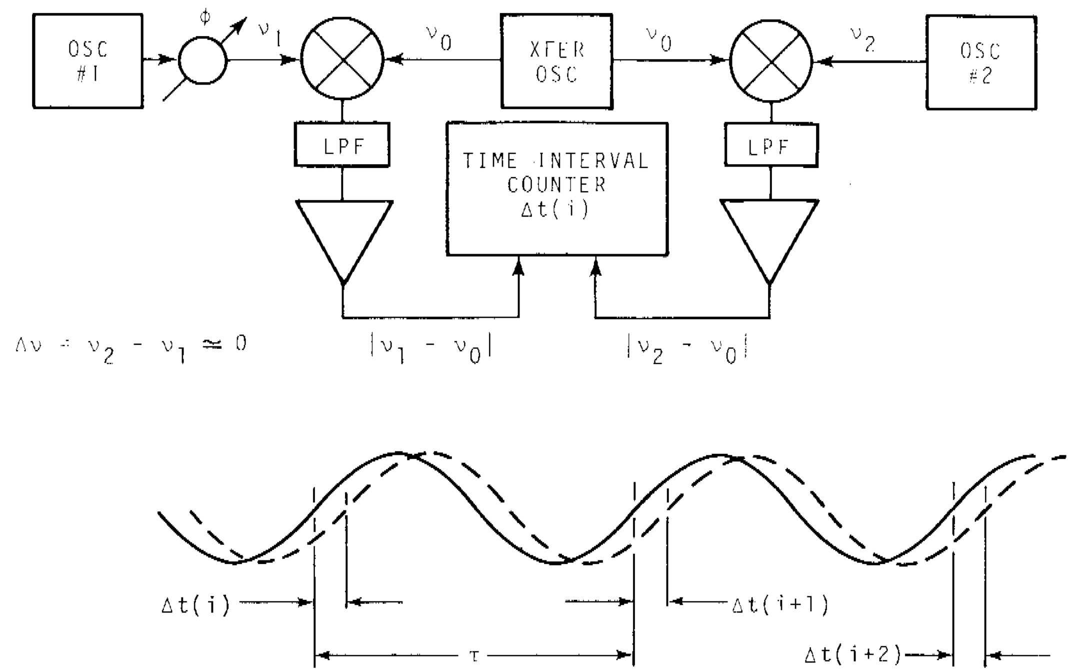

FIG. 5 A time difference and time fluctuation measurement system. The low pass filters (LPF) determine the measurement system bandwidth and must pass the difference frequencies which are depicted by the solid-line and dashed-line sinusoids at the bottom of the figure. The positive going zero volts crossing of these difference (beat) frequencies are used to start and stop a time interval counter after suitable low noise amplification. The i^{th} time difference between oscillator 1 and 2 is the $\Delta t(i)$ reading of the counter divided by $\tau\nu$ and plus any phase shift added, ϕ, where $\nu \simeq \nu_1 \simeq \nu_2$ is the nominal carrier frequency. The frequency difference is straight forwardly calculated from the time difference values.

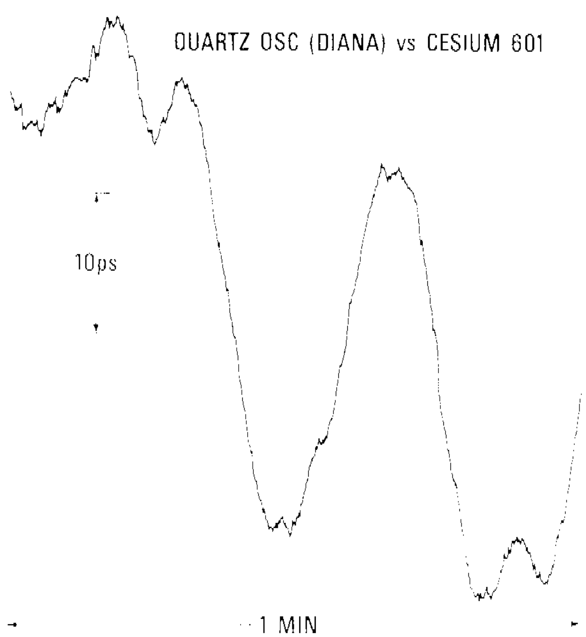

FIG. 6 A copy of a strip chart recording of the time fluctuations versus running time using the dual mixer time difference measurement system. The oscillators involved were a high performance commercial cesium standard and a high quality quartz crystal oscillator. The common oscillator employed was a low noise synthesizer. The measurement system noise was about 0.1 ps.

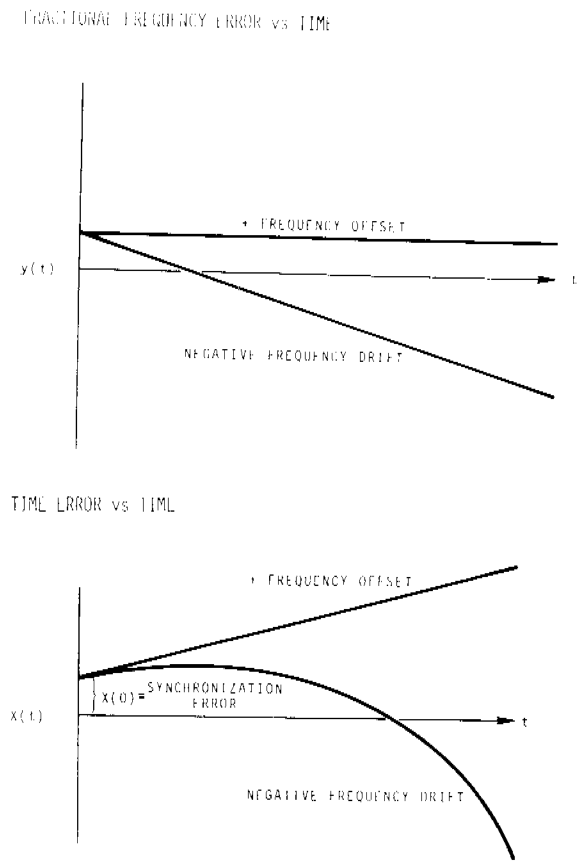

FIG. 7 Depiction of some commonly encountered non-random frequency and time deviations; i.e. a frequency offset error which maps into a linear time drift, and a linear frequency drift which maps into a quadratic time deviation.

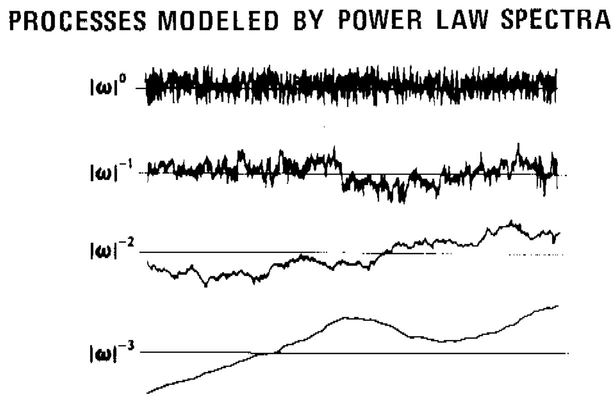

FIG. 8 Some sample plots of processes which may be modeled by power law spectral densities, $S(f) = h_\alpha f^\alpha$ ($\omega = 2\pi f$), as simulated with a computer. The white noise, $|\omega|^0$, is bandwidth limited. The subscript is left off of $S(f)$ as these plots may represent anything, e.g. frequency or time fluctuations.

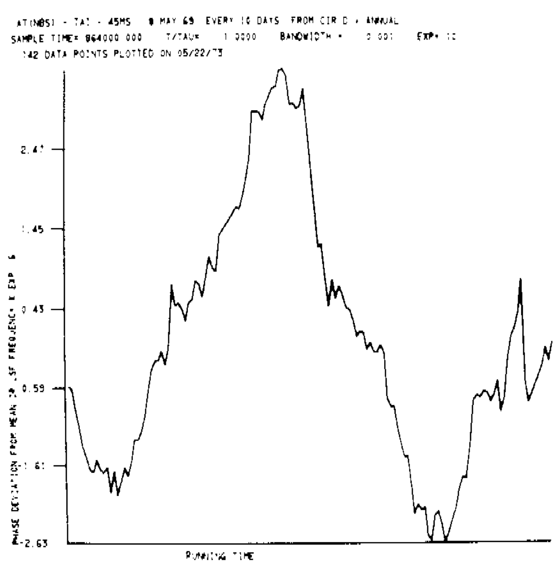

FIG. 9 The residual time fluctuations between the National Bureau of Standards Atomic Time Scale, AT(NBS), and the International Atomic Time Scale, TAI, after subtracting a least squares fit to the frequency. The vertical scale is in microseconds and the abscissa shows 1420 days following 8 May 1969. The peak-to-peak deviation is about 6μs.

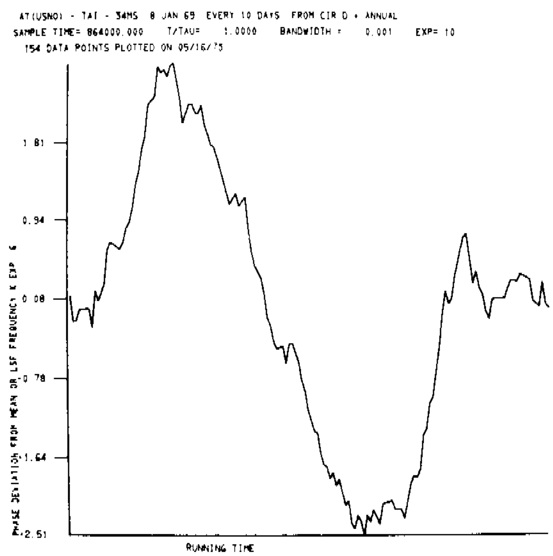

FIG 10 The residual time fluctuations between the United States Naval Observatory's Atomic Time Scale, AT (USNO), and the International Atomic Time Scale, TAI, after subtracting a least squares fit to the frequency. The vertical scale is in microseconds and the abscissa shows 1540 days following 8 Jan. 1969. The peak-to-peak deviation is less than 6μs.

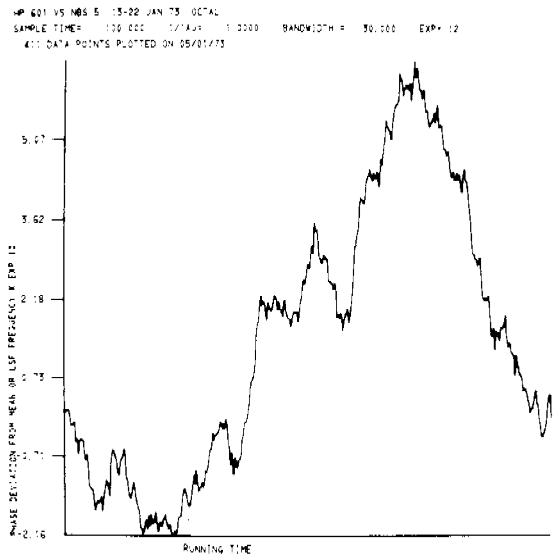

FIG 11 The residual time fluctuations between a high performance commercial cesium standard and one of the NBS primary frequency standards, NBS-5, after subtracting a mean frequency difference. The vertical scale is in units of 0.1 μs, and the abscissa shows 41100 s duration(~1/2 day). The peak-to-peak deviation is about 0.9 μs.

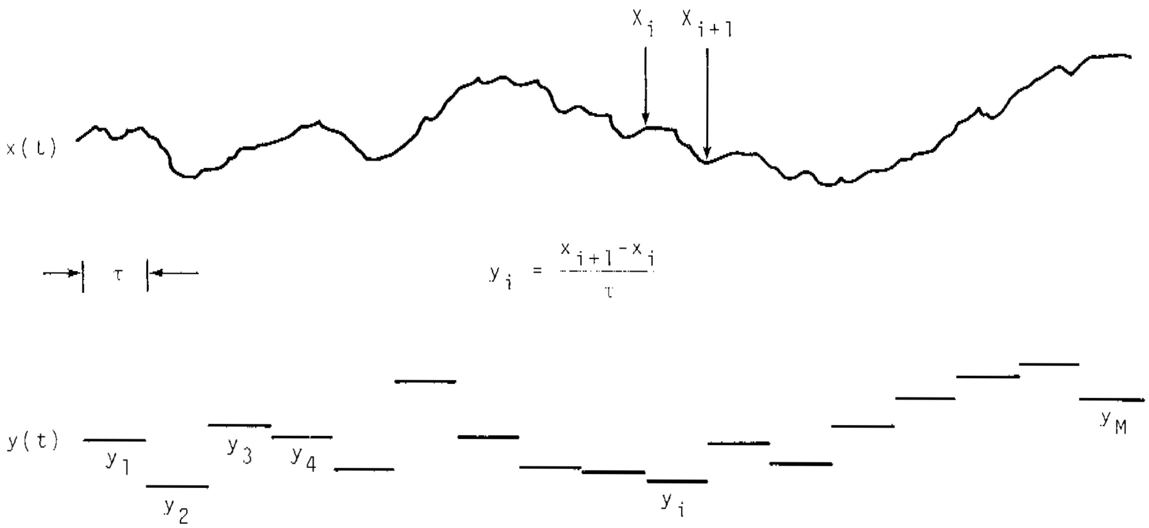

$$\sigma_{STD\ DEV\ y}(\tau) = \sqrt{\frac{1}{M-1} \sum_{i=1}^{M} (y_i - \bar{y})^2} \qquad \sigma_y(\tau) = \sqrt{\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (y_{i+1} - y_i)^2}$$

FIG 12 A simulated plot of the time fluctuations, x(t) between a pair of oscillators and of the corresponding fractional frequencies calculated from the time fluctuations each averaged over a sample time τ. At the bottom are the equations for the standard deviation (left) and for the time-domain measure of frequency stability as recommended by the IEEE subcommittee on Frequency Stability (right).

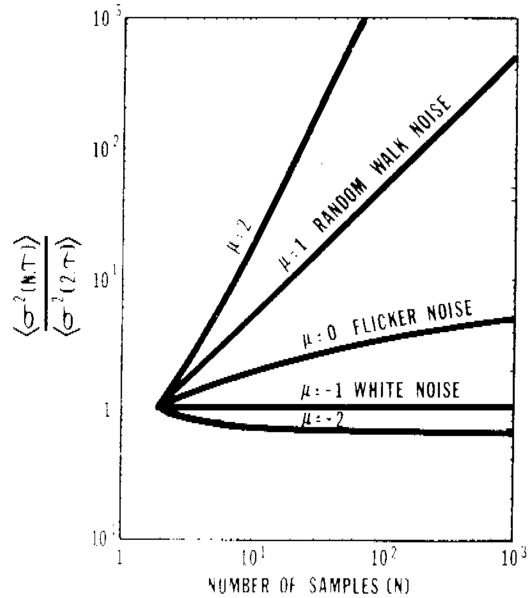

FIG 13 The ratio of the time average of the standard deviation squared for N samples over the time average of a two sample standard deviation squared as a function of the number of samples, N. The ratio is plotted for various power law spectral densities that commonly occur in precision oscillators. The figure illustrates one reason why the standard deviation is not a convenient measure of frequency stability; i.e. it may be very important to specify how many data points are in a data set if you use the standard deviation.

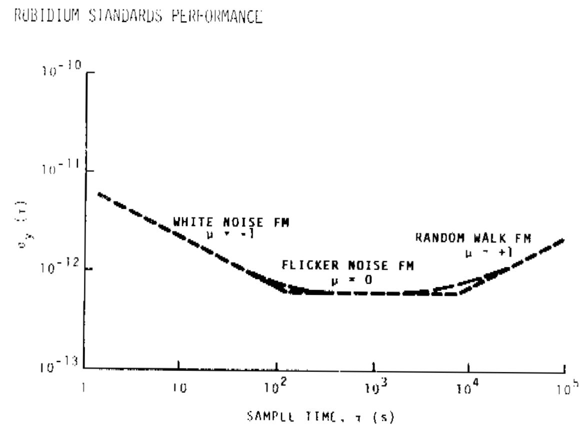

FIG 14 A $\sigma_y(\tau)$ versus τ plot modeling some actual data taken at NBS on some commercial rubidium standards. Notice that if $\sigma_y^2(\tau) \sim \tau^{\mu}$, then $\sigma_y(\tau) \sim \tau^{\mu/2}$ hence the $\tau^{-1/2}$ slope for $\mu = -1$. etc.

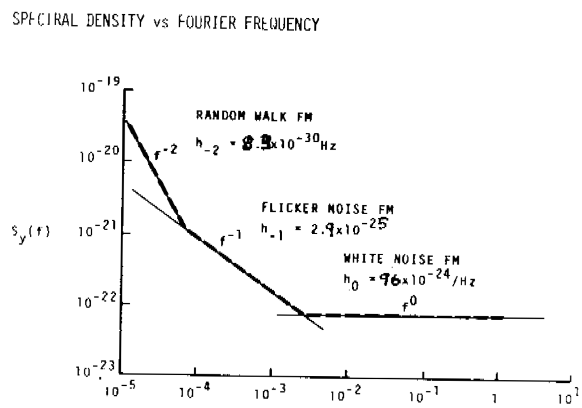

FIG 15 A plot of $S_y(f)$ versus f as transformed from the time-domain data plotted in Fig. 14.

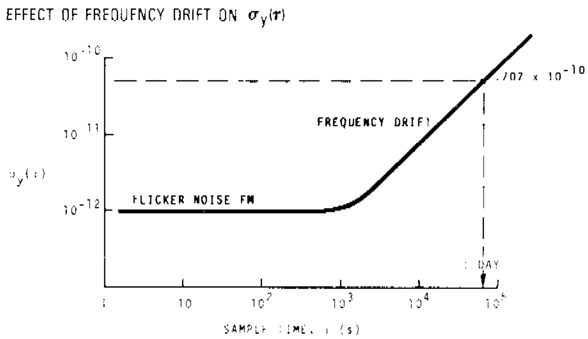

FIG 16 An example $\sigma_y(\tau)$ versus τ plot of an oscillator with both random fluctuations of flicker noise FM and a non-random linear fractional frequency drift of 10^{-10} per day. A plot appearing similar to this would be common for quartz crystal oscillators, for example.

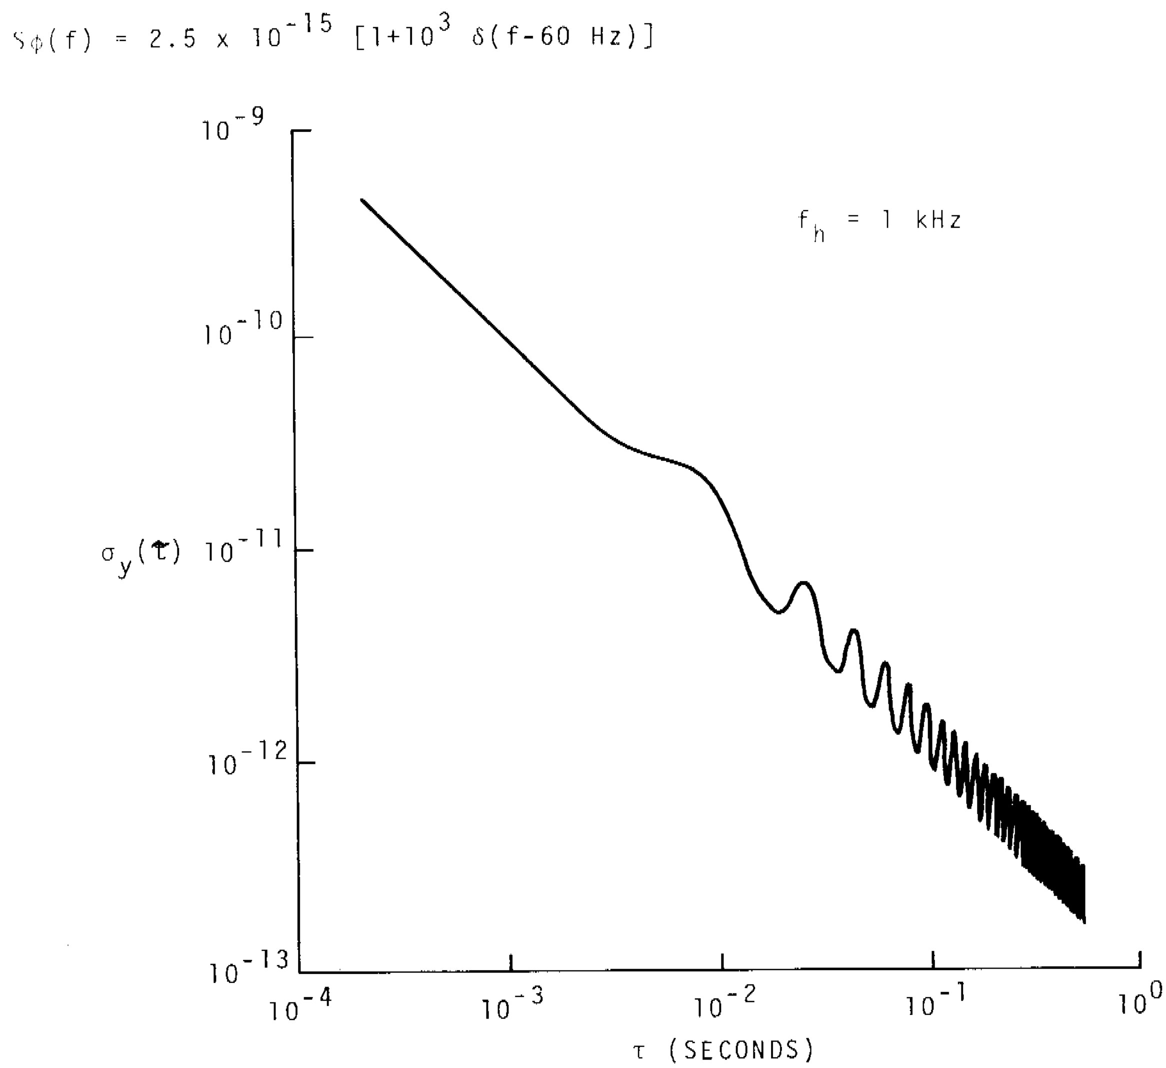

FIG 17 A calculated $\sigma_y(\tau)$ versus τ plot of white phase noise ($\alpha = +2$) with some 60 Hz FM superimposed. The power in the 60 Hz sidebands has been set equal to the power of the white phase noise in a 1 kHz bandwidth, f_h. Note: $S_\phi(f) = (\nu^2/f^2)\, S_y(f)$ and $\phi(t) = (2\pi)^2\, \nu^2 \cdot x(t)$.

QUARTZ OSCILLATOR (DIANA) VS COMMERCIAL CESIUM (#601)

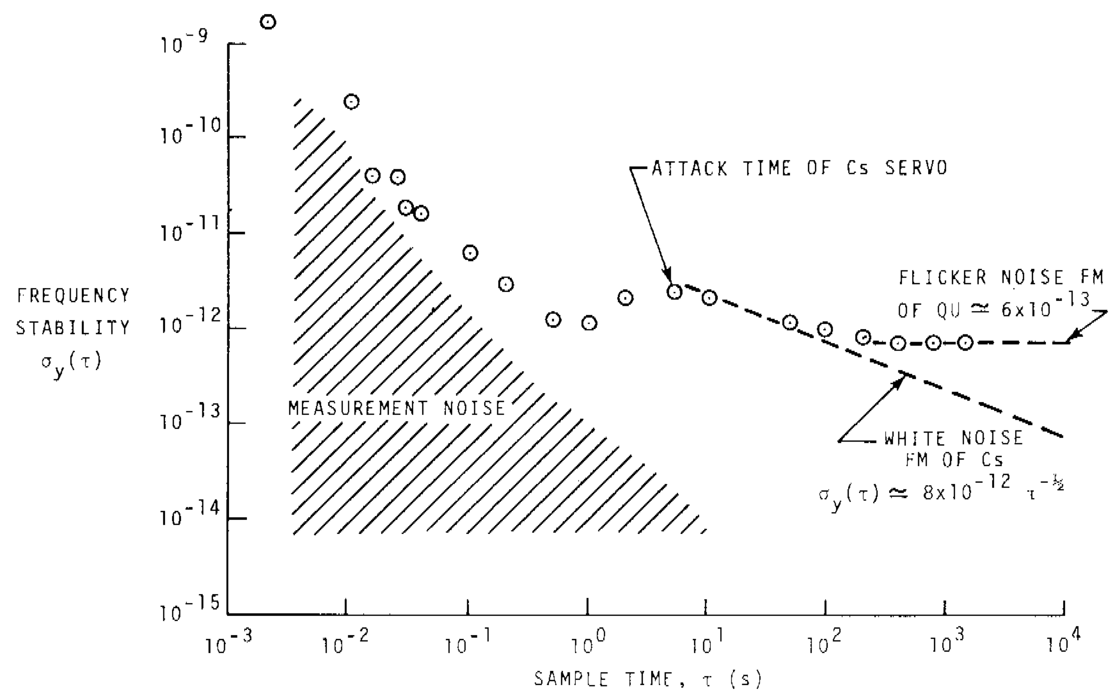

FIG.18 A $\sigma_y(\tau)$ versus τ plot of the fractional frequency fluctuations, $y(t)$ between a high performance commercial cesium beam frequency standard and a commercial quartz crystal oscillator.

QUESTION AND ANSWER PERIOD

DR. COSTAIN:

National Research Council.

Herman Damms had been doing a very similar dual mixer system and it does give a fantastic advantage when you have the two oscillators on the same frequency. As a matter of fact you can only test the system by putting the same signals into the two channels.

MR. ALLAN:

That's a good point. It's very conducive to testing the measurement noise, whereas with other systems it's sometimes quite difficult. This one is very amenable to looking at the measurement noise.