12 Frequency and Time—Their Measurement and Characterization

Samuel R. Stein

Time and Frequency Division
National Bureau of Standards
Boulder, Colorado

PRECISION FREQUENCY CONTROL
Volume 2
Oscillators and Standards
ISBN 0-12-280602-6

LIST OF SYMBOLS

df	Number of degrees of freedom
f	Fourier (cycle) frequency
$G_{eq}(j\omega)$	Open-loop transfer function of a phase-locked loop
F(s)	Transfer function of the loop filter of a phase-locked loop
$H(f-f_0)$	Transfer function of a filter
K_d	Sensitivity of a linear phase detector in volts per radian
P	Number of counts accumulated by a time-interval counter
Q	Quality factor of a resonance, equal to the ratio of the stored energy to the energy lost in one cycle
s	Independent variable of the Laplace transform
s_y^2	Estimate of the two-sample or Allan variance
$S_y(f)$	One-sided power spectral density of the fractional-frequency deviations
$S_\phi(f)$	One-sided power spectral density of the phase deviations
$S_\phi^{TS}(f)$	Two-sided power spectral density of the phase deviations
$S_{\phi,c}^{TS}(f)$	Two-sided power spectral density of the phase deviations of the carrier
$S_{\phi,p}^{TS}(f)$	Two-sided power spectral density of the phase deviations of the pedestal
$S_V^{TS}(f)$	Two-sided power spectral density of the instantaneous voltage
V_0	Peak voltage of a signal generator
V_d	Output voltage of a phase detector
V_t	Voltage applied to the tuning element of the voltage-controlled oscillator of a phase-locked loop
$V(t)$	Instantaneous voltage of a signal generator or other device
$x(t)$	Time deviation required by a signal generator operating at nominal frequency ν_0 to accumulate phase equal to $\phi(t)$.
$X(f)$	Fourier transform of $x(t)$
$y(t)$	Instantaneous fractional-frequency offset from nominal
$\bar{y}_k$	Mean fractional-frequency offset over the kth interval
α	Exponent of f for a power-law spectral density
$\Delta\nu_Q$	Frequency uncertainty due to the quantization of measurements
ζ	Damping constant of a phase-locked loop
μ	Exponent of τ for a power-law Allan variance
$\nu(t)$	Instantaneous frequency
ν_0	Nominal frequency of a signal generator
$\bar{\nu}(t_1; t_2)$	Mean frequency over the interval $t_1 \le t \le t_2$
ν_H	Heterodyne frequency or difference frequency between two oscillators
$\rho(\chi^2)$	Probability density of the chi-squared distribution
$\sigma_y^2(N, T, \tau)$	Sample variance of N fractional-frequency deviations, each averaged over a period τ spaced at intervals T
$\sigma_y^2(\tau)$	The two-sample or Allan variance of the fractional-frequency deviations
mod $\sigma_y^2(\tau)$	Modified Allan variance
τ	Averaging time
τ_0	Period of the time base of a counter
$\phi(t)$	Instantaneous phase deviation
χ^2	Chi-squared distribution
ω	Fourier (angular) frequency
ω_n	Natural frequency of a phase-locked loop

2.1 CONCEPTS, DEFINITIONS, AND MEASURES OF STABILITY

This chapter deals with the measurement of the frequency or time stability precision oscillators. It is assumed that the average output frequency is termined by a narrow-band circuit so that the signal is very nearly a sine ave. To be specific, it is also assumed that the output is a voltage, which is nventionally (Barnes *et al.*, 1971) represented by the expression

$$V(t) = [V_0 + \varepsilon(t)] \sin[2\pi\nu_0 t + \phi(t)], \tag{12-1}$$

here V_0 is the nominal peak voltage amplitude, $\varepsilon(t)$ the deviation of nplitude from nominal, ν_0 the nominal fundamental frequency, and $\phi(t)$ the viation of phase from nominal.

When either specifying or measuring the noise in an oscillator, one must nsider the nature of the reference. This may be either a passive circuit such a narrow-band filter, another similar oscillator, or a set of oscillators, nthesizers, and other signal-generating equipment. A reference with lower ise than the device under test may be available, and in this case the pressions developed in this chapter describe the noise in the oscillator one. However, a state-of-the-art device will have lower noise than any ailable reference. In this case all the expressions below refer to the sum of vice and reference noise. The most common approach to solving this roblem is to compare two or more nearly identical devices. Under most rcumstances it is then reasonable to assume that each oscillator contributes alf of the measured noise.

The most direct and intuitive method of characterizing the properties of signal is to determine the two-sided spectrum of $V(t)$, which is denoted $S_V^{rs}(f)$ (Rutman, 1978). The variable f is called a Fourier frequency and is ry closely related to the concept of a modulation frequency. A positive indicates a frequency above the carrier frequency ν_0, while a negative f dicates a frequency lower than the carrier. Since the noise can in theory odulate the carrier at all possible frequencies, a continuous function is quired to describe the modulation of $V(t)$. S is called a spectral density and $S_V^{rs}(f)$ is the mean-square voltage $\langle V^2(t)\rangle$ in a unit bandwidth centered at f. It proportional to the rf power per unit bandwidth delivered by the oscillator a matched load. The total signal power is proportional to the mean square oltage, which is also called the variance of the signal since the mean value of (t) is zero. The variance is therefore equal to the two-sided spectral density ntegrated over all frequencies.

The two-sided spectrum is usually measured by an rf spectrum analyzer, a evice that functions like a bandpass filter followed by a bolometer, as shown

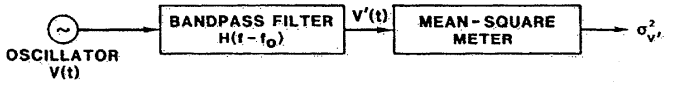

FIG. 12-1 An rf spectrum analyzer. The device produces an output proportional to the mean-square value of the signal passing through a tunable narrow-band filter centered at frequency f_0.

in Fig. 12-1. The spectrum of the filtered voltage $V'(t)$ is equal to the square of the magnitude of the filter transfer function $H(f-f_0)$ multiplied by the spectrum of the input signal (Cutler and Searle, 1966). The variance of the filtered voltage is obtained from Parseval's theorem:

$$\sigma_{V'}^2(f_0) = \int_{-\infty}^{\infty} |H(f - f_0)|^2 S_V^{TS}(f)\, df. \tag{12-2}$$

If the bandpass filter is sufficiently narrow, so that $S_V^{TS}(f)$ changes negligibly over its bandwidth, then Eq. (12-2) may be inverted. With this assumption, the power spectrum is estimated from the measurement using Eq. (12-3):

$$S_V^{TS}(f_0) = \sigma_{V'}^2(f_0)/B, \tag{12-3}$$

where $B = \int_{-\infty}^{\infty} |H(f' - f_0)|^2\, df'$ is the noise bandwidth of the filter and f_0 its center frequency. Figure 12-2 shows a typical two-sided rf spectrum. For many oscillators the spectrum has a Lorentzian shape, that is,

$$S_V^{TS}(f) = \frac{2\langle V^2\rangle/\pi\, \Delta f_{3dB}}{1 + (f/(\Delta f_{3dB}/2))^2}. \tag{12-4}$$

The Lorentzian lineshape is completely described by the mean square voltage $\langle V^2\rangle$ and the full width at half maximum Δf_{3dB}.

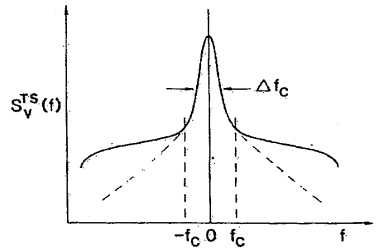

FIG. 12-2 The rf spectrum of a signal. It is often useful to divide the spectrum into the carrier and the noise pedestal. The spectral density of the carrier exceeds that of the noise pedestal for Fourier frequencies smaller in magnitude than f_c.

12.1.1 Relationship between the Power Spectrum and the Phase Spectrum

The power spectrum differs from a delta function $\delta(f)$ due to the presence of the amplitude- and phase-noise terms, $\varepsilon(t)$ and ϕ(t), respectively, included in Eq. (12-1). Usually the noise modulation separates into two distinct components, so that one observes a very narrow feature called the carrier above the level of a relatively broad pedestal. The frequency that separates the carrier and pedestal is denoted f_c. Below this frequency the spectral density of the carrier exceeds that of the pedestal. Assuming that the amplitude noise is negligible compared to the phase noise and that the phase modulation is small, the relationship between the power and phase spectra is given by (Walls and DeMarchi, 1975)

$$S_V^{TS}(f) \simeq \frac{V_0^2}{2} e^{-I(f_c)} \times \begin{cases} \delta(f), & \text{if } -f_c < f < f_c, \\ S_\phi^{TS}(f), & \text{otherwise.} \end{cases} \tag{12-5}$$

$S_\phi^{TS}(f)$ is the two-sided spectrum of the phase fluctuations, which divides into a carrier component $S_{\phi,c}^{TS}(f)$ and a pedestal component $S_{\phi,p}^{TS}(f)$; $I(f_c)$ is given by

$$I(f_c) = \int_{f_c}^{\infty} S_{\phi,p}(f)\, df \simeq \int_{f_c}^{\infty} S_\phi(f)\, df \tag{12-6}$$

since above f_c the pedestal dominates the noise spectrum. The variance of the carrier is equal to $(V_0^2/2)e^{-I(f_c)}$, with the remaining variance in the pedestal. If $\Delta\nu_p$ is the width of the pedestal and $\Delta\nu_c$ the width of the carrier, then the power density in the carrier is equal to that in the pedestal when $I(f_c) = \ln(\Delta\nu_p/\Delta\nu_c)$. For the pedestal, one may use the 3-dB linewidth for $\Delta\nu_p$ provided that $\int_{f_c}^{\infty} S_\phi(f)\, df < \ln 2$. Otherwise the pedestal width is estimated from $\int_{\Delta\nu/2}^{\infty} S_\phi(f)\, df = \ln 2$. For the carrier, the linewidth is estimated by calculating $\int_{\Delta\nu_c/2}^{\infty} S_{\phi,c}(f)\, df = \ln 2$.

The foregoing analysis makes it possible to draw certain conclusions concerning detection of the carrier. We use $\int_p df$ to denote the integral over the phase noise pedestal. If $\int_p S_\phi(f)\, df < \ln 2$, then the carrier may be resolved irrespective of detector bandwidth. When $\ln(\Delta\nu_p/\Delta\nu_c) > \int_p S_\phi(f)\, df > \ln 2$, the carrier may be resolved by restricting the detection bandwidth. But when $\int_p S_\phi(f)\, df > \ln(\Delta\nu_p/\Delta\nu_c)$, the carrier can no longer be distinguished from the pedestal since its spectral density is smaller.

12.1.2 The IEEE Recommended Measures of Frequency Stability

By the mid-1960s the problem of the specification of precision oscillators had become extremely important, but there was very little uniformity among manufacturers, metrologists, and applications engineers in the methods of

performing measurements or the description of measurement results. This situation was complicated by the difficulty of comparing the various descriptions. A measure of stability is often used to summarize some important feature of the performance of the standard. It may therefore not be possible to translate from one measure to another even though the respective measurement processes are fully described and all relevant parameters are given. This situation resulted in a strong pressure to achieve a higher degree of uniformity.

In order to reduce the difficulty of comparing devices measured in separate laboratories, the IEEE convened a committee to recommend uniform measures of frequency stability. The recommendations made by the committee are based on the rigorous statistical treatment of ideal oscillators that obey a certain model (Barnes *et al.*, 1971). Most importantly, these oscillators are assumed to be elements of a stationary ensemble. A random process is stationary if no translation of the time coordinate changes the probability distribution of the process. That is, if one looks at the ensemble at one instant of time, then the distribution in values for a process within the ensemble is exactly the same as the distribution at any other instant of time. The elements of the ensemble are not constant in time, but as one element changes value other elements of the ensemble assume previous values. Thus, it is not possible to determine the particular time when the measurement was made.

The stationary noise model has been adopted because many theoretical results, particularly those related to spectral densities, are valid only for this case. It is important for the statistician to exercise considerable care since experimentally one may measure quantities approximately equal to either the instantaneous frequency of the oscillator or the instantaneous phase. But the ideal quantities approximated by these measurements may not both be stationary. The instantaneous angular frequency is conventionally defined as the time derivative of the total oscillator phase. Thus,

$$\omega(t) = \frac{d}{dt}[2\pi\nu_0 t + \phi(t)], \tag{12-7}$$

and the instantaneous frequency is written

$$\nu(t) = \nu_0 + \frac{1}{2\pi}\frac{d\phi}{dt}. \tag{12-8}$$

For precision oscillators, the second term on the right-hand side is quite small, and it is useful to define the fractional frequency

$$y(t) = \frac{\nu(t) - \nu_0}{\nu_0} = \frac{1}{2\pi\nu_0}\frac{d\phi}{dt} = \frac{dx}{dt}, \tag{12-9}$$

where

$$x(t) = \phi(t)/2\pi\nu_0 \tag{12-10}$$

is the phase expressed in units of time. Alternatively, the phase could be written as the integral of the frequency of the oscillator:

$$\phi(t) = \phi_0 + \int_0^t 2\pi[\nu(\theta) - \nu_0]\, d\theta. \tag{12-11}$$

However, the integral of a stationary process is generally not stationary. Thus, indiscriminate use of Eqs. (12-7) and (12-11) may violate the assumptions of the statistical model. This contradiction is avoided when one accounts for the finite bandwidth of the measurement process. Although a more detailed consideration of the statistics goes beyond the scope of this treatment, it is very important to keep in mind the assumption that lie behind the statistical analysis of oscillators. In order to analyze the behavior of real oscillators, it is necessary to adopt a model of their performance. The model must be consistent with observations of the device being simulated. To make it easier to estimate the device parameters, the models usually include certain predictable features of the oscillator performance, such as a linear frequency drift. A statistical analysis is useful in estimating such parameters to remove their effect from the data. It is just these procedures for estimating the deterministic model parameters that have proved to be the most intractable. A substantial fraction of the total noise power often occurs at Fourier frequencies whose periods are of the same order as the data length or longer. Thus, the process of estimating parameters may bias the noise residuals by reducing the noise power at low Fourier frequencies. A general technique for minimizing this problem in the case of oscillators actually observed in the laboratory is discussed below.

It has been suggested that measurement techniques for frequency and time constitute a hierarchy (Allan and Daams, 1975), with the measurement of the total phase of the oscillator at the peak. Although more difficult to measure with high precision than other quantities, the total phase has this status owing to the fact that all other quantities can be derived from it. Furthermore, missing measurements produce the least deleterious effect on a time series consisting of samples of the total phase. Gaps in the data affect the computation of various time-dependent quantities for times equal to or shorter than the gap length, but have a negligible effect for times much longer than the gap length. The lower levels of the hierarchy consist of the time interval, frequency, and frequency fluctuation. When one measures a quantity somewhere in this hierarchy and wishes to obtain a higher quantity, it is necessary to integrate one or more times. In this case the problem of missing data is quite serious. For example, if frequency is measured and one wants to

know the time of a clock, one needs to perform the integration in Eq. (12-11). The missing frequency measurements must be bridged by estimating the average frequency over the gap, resulting in a time error that is propagated forever. Thus, it is preferable to always make measurements at a level of the measurement hierarchy equal to or above the level corresponding to the quantity of principle interest. In the past this was rather difficult to do. Measurement systems constructed from simple commercial equipment suffered from dead time, that is, they were inactive for a period after performing a measurement. To make matters worse, methods for measuring time or phase had considerably worse noise performance than methods for measuring frequency. As a result, many powerful statistical techniques were developed to cope with these problems (Barnes, 1969; Allan, 1966). The effect of dead time on the statistical analysis has been determined (Lesage and Audoin, 1979b). Other techniques have been developed to combine short data sets so that the parameters of clocks over long periods of time could be estimated despite missing data (Lesage, 1983). The rationale for these approaches is considerably diminished today. Low-noise techniques for the measurement of oscillator phase have been developed. Now, commercial equipment is capable of measuring the time or the total phase of an oscillator with very high precision. Other equipment exists for measuring the time interval. These devices use the same techniques that were previously employed for the measurement of frequency and are very competitive in performance.

The proliferation of microcomputers and microprocessors has had an equally profound effect on the field of time and frequency measurement. There has been a dramatic increase in the ability of the metrologist to acquire and process digital data. Many instruments are available with suitable standard interfaces such as IEEE-583 or CAMAC (IEEE, 1975) and IEEE-488 (IEEE, 1978). As a result, there has been a dramatic change in direction away from analog signal processing toward the digital, and this process is accelerating daily. Techniques once used only by national standards laboratories and other major centers of clock development and analysis are now widespread. Consequently, this chapter will focus first on the peak of the measurement hierarchy and the use of digital signal processing. But the analysis is directed toward estimating the traditional measures of frequency stability. Considerable attention will be paid to problems associated with estimating the confidence of these stability measures and obtaining the maximum information from available data.

The IEEE has recommended as its first measure of frequency stability the one-sided spectral density $S_y(f)$ of the instantaneous fractional-frequency fluctuations $y(t)$. It is simply related to the spectral density of phase fluctuations since differentiation of the time-dependent functions is equivalent to

multiplication of their Fourier transforms by $j\omega$:

$$S_y(f) = (f/\nu_0)^2 S_\phi(f) = (2\pi f)^2 S_x(f) \tag{12-12}$$

Section 12.1.1 on the relationship between the power spectrum and the phase spectrum described the analog method for the measurement of a spectral density. If a voltage V is the output of the oscillator, then the result of the measurement is proportional to the rf power spectral density. But if the voltage were proportional to the frequency or phase of the oscillator, then the result of the measurement would be proportional to the spectral density of the frequency or phase. The most common units of $S_\phi(f)$ are radians squared per hertz.

Alternatively, the spectral density can be obtained by digital analysis of the signal. For example, the quantity $S_x(f)$ can be calculated from the Fourier transform of $x(t)$. The relevant continuous Fourier-transform pair is defined as follows:

$$X(f) = \int_{-\infty}^{\infty} x(t) e^{-j2\pi ft}\, dt \tag{12-13}$$

and

$$x(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} X(f) e^{j2\pi ft}\, df. \tag{12-14}$$

However, one does not generally have continuous knowledge of the phase of the oscillator. Since it is relatively easy to measure $x(t)$ at equally spaced time intervals, we assume the existence of the series x_l, where $x_l = x(l\tau)$ for integer values of l. The discrete Fourier transform is defined by analogy to the continuous transform (Cochran *et al.*, 1967):

$$X(f) = \sum_{l=-\infty}^{\infty} x(l) e^{-j2\pi fl}. \tag{12-15}$$

In practice the time series has finite length T consisting of N intervals of length τ, and it is not possible to compute the infinite sum. Nevertheless, it remains possible to compute a spectrum that is not continuous in f but rather has resolution Δf, where

$$\Delta f = 1/T = 1/N\tau. \tag{12-16}$$

The need to sum over all values of the index l is removed by assuming that the function $x(t)$ repeats itself with period T. The resulting spectrum contains no information on the spectrum at Fourier frequencies less than $1/T$. Truncation of the time series also introduces spurious effects due to the turn-on and turn-off transients. These problems can be minimized through the use of a window function. The computed spectrum is actually the square of the magnitude of

the window function multiplied by the desired spectrum. The use of a window function reduces the variance of the spectrum estimate at the expense of smearing out the spectrum to a small degree. With these changes but no window function, we arrive at the discrete finite transform

$$X(n\,\Delta f) = \frac{1}{N}\sum_{k=0}^{N-1} x(k\tau)e^{-j2\pi n\,\Delta f\,k\tau}. \tag{12-17}$$

The spectral density of $x(t)$ is computed from Eq. (12-17) by squaring the real and imaginary components, adding the two together, and dividing by the total time T:

$$S_x(m\,\Delta f) = \frac{\{R[X(m\,\Delta f)]\}^2 + \{I[X(m\,\Delta f)]\}^2}{T}. \tag{12-18}$$

The digital method of estimating spectral densities has many advantages over analog signal processing. Most important is the fact that it may be computed from any set of equally spaced samples of a time series. As a result, the technique is compatible with other methods of characterizing the signal, that is, the sampled data can be stored and processed using a variety of algorithms. In addition, each record of length T produces a single estimate of the spectrum for each of the N frequencies $\Delta f, 2\,\Delta f, \ldots, N\,\Delta f$. It is therefore possible to estimate the entire spectrum much more quickly using the digital technique than it would be using analog methods. The fast Fourier transform, a very efficient algorithm for the computation of the discrete finite transform, has opened the way to versatile self-contained, commercial spectrum analysis. It is also very straightforward to compute the spectrum from data acquired by computerized digital data acquisition systems.

A result of the finite sampling rate is that the upper frequency limit of the digital spectrum analysis is $1/2\tau$, called the Nyquist frequency (Jenkins and Watts, 1968). Power in the signal being analyzed that is at frequencies higher than the Nyquist frequency affects the spectrum estimate for lower frequencies. This problem is called aliasing. The out-of-band signal is rejected by only approximately 6 dB per octave above the Nyquist frequency. Thus, when significant out-of-band signals exist, they must be reduced by analog filtering. One or more low-pass filters are usually sufficient for this purpose.

As its second measure of frequency stability, the IEEE recommended the sample variance $\sigma_y^2(\tau)$ of the fractional-frequency fluctuations. It is a measure of the variability of the average frequency of an oscillator between two adjacent measurement intervals. The average fractional-frequency deviation $\bar{y}_k$ over the time interval from t_k to $t_k + \tau$ is defined as

$$\bar{y}_k = \frac{1}{\tau}\int_{t_k}^{t_k+\tau} y(t)\,dt, \tag{12-19}$$

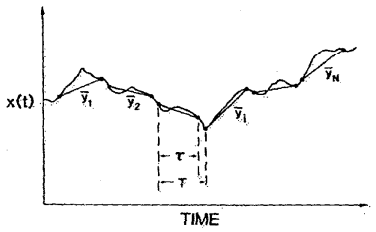

FIG. 12-3 Measurement process for the computation of the sample variance. The phase difference between two oscillators is plotted on the ordinate. The measurement yields a set of frequencies averaged over equal intervals τ separated by dead time $T - \tau$.

from which it follows that

$$\bar{y}_k = \frac{x(t_k + \tau) - x(t_k)}{\tau}. \tag{12-20}$$

The quality τ is often referred to as the sampling time or the averaging time. Equations (12-19) and (12-20) are not the only way to define mean frequency, but they are the simplest. Other definitions lead to alternative measures of stability that may have desirable properties.

Suppose that one has measured the time or frequency fluctuations between a pair of precision oscillators and a stability analysis is desired. The process is illustrated in Fig. 12-3. These are N values of the fractional frequency $\bar{y}_i$. Each one is measured over a time τ, and measurements are repeated after intervals of time T. If the measurement repetition time exceeds the averaging time, then there is a dead time equal to $T - \tau$ between each frequency measurement, during which there is no information available.

There are many ways to analyze these data. A fairly general approach is the N-sample variance defined by the relation

$$\langle \sigma_y^2(N, T, \tau) \rangle = \left\langle \frac{1}{N-1} \sum_{n=1}^{N} \left(\bar{y}_n - \frac{1}{N} \sum_{k=1}^{N} \bar{y}_k \right)^2 \right\rangle, \tag{12-21}$$

where the angle brackets denote the infinite time average. Frequently, Eq. (12-21) does not converge as $N \to \infty$, since some noise processes in oscillators diverge rapidly at low Fourier frequencies. This implies that the precision with which one estimates the variance does not improve simply as the sample size is increased. For this reason, the two-sample variance with no dead time is preferred. Also called the Allan variance, it converges for all the major noise types observed in precision oscillators. It may be written as

$$\sigma_y^2(\tau) = \langle \tfrac{1}{2}(\bar{y}_{k+1} - \bar{y}_k)^2 \rangle. \tag{12-22}$$

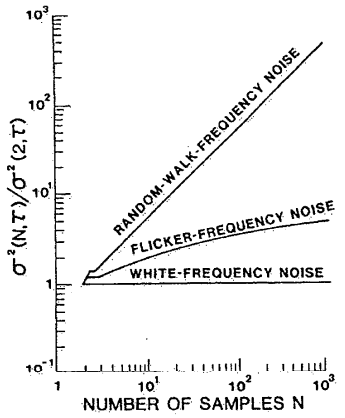

FIG. 12-4 N-sample variance versus Allan variance. The two-sample variance converges for the important types of noise observed in frequency standards but the ratio of the traditional variance to the two-sample variance is an increasing function of sample size for flicker frequency noise and random-walk frequency noise.

The dependence of the classical variance on the number of samples is shown in Fig. 12-4 for the case of no dead time. The quantity plotted is the ratio of the N-sample variance to the Allan variance. Note that $\sigma_y^2(\tau)$ has the same value as the classical variance for the white-noise frequency modulation. However, the classical variance grows without bound for flicker-frequency and random-walk-frequency noises.

One may combine Eqs. (12-20) and (12-22) to obtain an equation for $\sigma_y(\tau)$ in terms of the time-difference or time-deviation measurements:

$$\sigma_y^2(\tau) = \langle \tfrac{1}{2}\tau^{-2}[x(t+2\tau) - 2x(t+\tau) + x(t)]^2 \rangle. \tag{12-23}$$

N discrete time readings may be used to estimate the variance

$$\sigma_y^2(\tau) \cong \frac{1}{2(N-2)\tau^2} \sum_{i=1}^{N-2} (x_{i+2} - 2x_{i+1} + x_i)^2, \tag{12-24}$$

where i denotes the number of the measurement in the set of N and the nominal spacing between measurements is τ. Since it has been assumed that there is no dead time between measurements, one can write τ in Eq. (12-24) as an integer multiple of τ_0, that is, $\tau = m\tau_0$, where τ_0 is the smallest spacing of the data. In this case

$$\sigma_y^2(m\tau_0) \cong \frac{1}{2(N-2m)m^2\tau_0^2} \sum_{i=1}^{N-2m} (x_{i+2m} - 2x_{i+m} + x_i)^2. \tag{12-25}$$

12.1.3 The Concepts of the Frequency Domain and the Time Domain

Spectral densities are measures of frequency stability in what is called the equency domain since they are functions of Fourier frequency. The Allan ariance, on the other hand, is an example of a time-domain measure. In a trict mathematical sense, these two descriptions are connected by Fourier ransform relationships (Cutler and Searle, 1966). However, for many years he inadequacy of measurement equipment created artificial barriers between hese two characterizations of the same noise process. As a result, many pecialized techniques have been developed to translate between the various neasures of stability (Allan, 1966; Burgoon and Fischer, 1978). The preceding ections have demonstrated how easily both types of stability measures can e computed from the same data provided that the measurement process rovides complete information. For example, both $\sigma_y^2(m\tau_0)$ and $S_x(m\,\Delta f)$ can e computed from evenly spaced samples of $x(t)$. However, incomplete nformation can result from either measurement dead time or interruptions in he data acquisition process. In these cases translation techniques remain aluable.

Both the spectral density and the Allan Variance are second-moment measures of the time series $x(t)$. However, it is only possible to translate unambiguously from the spectral density to the Allan variance, not the reverse. To calculate the spectral density it is necessary to use the autocorrelation function of the phase. The following discussion on power-law noise processes further demonstrates this dichotomy. As we shall see, the Allan variance for a fixed measurement bandwidth does not distinguish between all of the noise processes that are commonly observed in precision oscillators.

12.1.4 Translation between the Spectral Density of Frequency and the Allan Variance

The power-law model is most frequently used for describing oscillator phase noise. It assumes that the spectral density of frequency fluctuations is equal to the sum of terms, each of which varies as an integer power of frequency. Thus, there are two quantities that completely specify $S_y(f)$ for a particular power-law noise process: the slope on a log–log plot for a given range of f and the amplitude. The slope is denoted by α and therefore f^α is the straight line on a log–log plot that relates $S_y(f)$ to f. The amplitude is denoted h_α. When we examine a plot of the spectral density of frequency fluctuations, we represent it by the addition of all the power-law processes (Allan, 1966; Vessot *et al.*, 1966) with the appropriate coefficients:

$$S_y(f) = \sum_{\alpha=-\infty}^{\infty} h_\alpha f^\alpha. \tag{12-26}$$

TABLE 12-1

Correspondence between Common Power-Law Spectral Densities and the Allan Variance[a]

Noise type	$S_y(f)$	$\sigma_y^2(\tau)$
White phase	$h_2 f^2$	$3f_h h_2/(2\pi)^2\tau^2$
Flicker phase	$h_1 f$	$\dfrac{[1.038 + 3\ln(2\pi f_h\tau)]}{(2\pi)^2} h_1 \dfrac{1}{\tau^2}$
White frequency	h_0	$\frac{1}{2}h_0(1/\tau)$
Flicker frequency	$h_{-1}f^{-1}$	$2\ln(2)h_{-1}$
Random-walk frequency	$h_{-2}f^{-2}$	$\frac{1}{6}(2\pi)^2 h_{-2}\tau$

[a] Where necessary for convergence the spectral density has been assumed to be zero for frequencies greater than the cutoff frequency f_h.

This technique is most valuable when only a few terms in Eq. (12-26) are required to describe the observed noise and each term dominates over several decades of frequency. This situation often prevails. Five power-law noise processes (Allan, 1966; Vessot *et al.*, 1966) are common with precision oscillators:

(1) random-walk frequency modulation $\alpha = -2$
(2) flicker frequency modulation $\alpha = -1$
(3) white frequency modulation $\alpha = 0$
(4) flicker phase modulation $\alpha = 1$
(5) white phase modulation $\alpha = 2$

The spectral density of frequency is an unambiguous description of the oscillator noise. Thus, the spectrum can be used to compute the Allan variance (Barnes *et al.*, 1971):

$$\sigma_y^2(\tau) = \frac{2}{(\pi\nu_0\tau)^2}\int_0^\infty S_\phi(f)\sin^4(\pi f\tau)\,df. \tag{12-27}$$

However, Eq. (12-27) shows that the Allan variance is very sensitive to the high frequency dependence of the spectral density of phase, thereby necessitating a detailed knowledge of the bandwidth-limiting elements in the measurement setup. The integral has been computed for each of the power-law noise processes, and the results are summarized in Table 12-1 (Barnes *et al.*, 1971). For α in the range $-2 \le \alpha \le 0$, the Allan variance is proportional to τ^μ, where $\mu = -\alpha - 1$. When the log of the Allan variance is plotted as a function of the log of the averaging time, the graph also consists of straight-line segments with integer slopes. However, Table 12-1 also shows that even if

e oscillator is reasonably modeled by power-law spectra, it is not practical distinguish between white phase noise and flicker phase noise from the pendence of the Allan variance on τ. In both cases $\sigma_y^2 \simeq 1/\tau^2$.

1.5 The Modified Allan Variance

Table 12-1 also shows that the Allan variance has very different bandwidth pendence for white phase noise and flicker phase noise. Therefore, these ise types have been distinguished by varying the bandwidth of the easurement system. If $x(t)$ were measured, the noise type could be identified computing the spectrum. However, both the approach of making easurements as a function of bandwidth and the computation of the ectrum can be avoided by calculating a modified version of the Allan riance. The algorithm for this variance has the effect of changing the ndwidth inversely in proportion to the averaging time (Snyder, 1981; Allan d Barnes, 1981).

Each reading of the time deviation x_i has associated with it a measurement-stem bandwidth f_h. Similarly, we can define a software bandwidth $f_s = f_h/n$, hich is $1/n$ times narrower than the hardware bandwidth. It can be realized y averaging n adjacent x_i's. Based on this idea it is possible to define a odified Allan variance that allows the reciprocal software bandwidth to ange linearly with the sample time τ:

$$\operatorname{mod}\sigma_y^2(\tau) = \frac{1}{2\tau^2}\left\langle\left[\frac{1}{n}\sum_{i=1}^{n}(x_{i+2n} - 2x_{i+n} + x_i)\right]^2\right\rangle, \tag{12-28}$$

here $\tau = n\tau_0$. Equation (12-28) reduces to Eq. (12-23) for $n = 1$. One can see at $\operatorname{mod}\sigma_y^2(\tau)$ is the second difference of three time values, each of which is a onoverlapping average of n of the x_i's. As n increases the software andwidth decreases as f_h/n.

For a finite data set of N readings of x_i ($i = 1$ to N), $\operatorname{mod}\sigma_y^2(\tau)$ can be stimated from the expression

$$\operatorname{mod}\sigma_y^2(\tau) \cong \frac{1}{2\tau^2 n^2(N - 3n + 1)}\sum_{j=1}^{N-3n+1}\sum_{i=j}^{n+j-1}(x_{i+2n} - 2x_{i+n} + x_i)^2, \tag{12-29}$$

hich is easy to program but takes more time to compute than the orresponding equation (12-24) for $\sigma_y^2(\tau)$.

Table 12-2 gives the relationship between the time-domain measure $\operatorname{mod}\sigma_y^2(\tau)$ and its power-law spectral counterpart. In the right-hand column re the asymptotic values of the ratio of the modified Allan variance to the llan variance. It is clear from the table that $\operatorname{mod}\sigma_y^2(\tau)$ is very useful for white

TABLE 12-2

Correspondence between Common Power-Law Spectral Densities and the Modified Allan Variance[a]

Noise type	$S_y(f)$	mod $\sigma_y^2(\tau)$	mod σ_y^2/σ_y^2
White phase	$h_2 f^2$	$h_2 \frac{3f_h}{n(2\pi)^2} \frac{1}{\tau^2}$	n
Flicker phase	$h_1 f$	$h_1 \frac{[1.038 + 3\ln(2\pi f_h \tau)]}{(2\pi)^2} \frac{1}{\tau^2}$	1
White frequency	h_0	$h_0/4\tau$	0.5
Flicker frequency	$h_{-1} f^{-1}$	$h_{-1}(0.936)$	0.674
Random-walk frequency	$h_{-2} f^{-2}$	$h_{-2}(5.42)\tau$	0.824

[a] Where necessary the spectral density has been assumed to be zero for frequencies greater than the cutoff frequency f_h. The constant n is the number of adjacent phase values that are averaged to produce the bandwidth reduction. The values in the last two columns are for the asymptotic limit $n \to \infty$. In practice, n only needs to be 10 or larger before the asymptotic limit is approached within a few percent. When $n = 1$ the ratio in the last column is 1 in all cases.

phase modulation and flicker phase modulation, but for $\alpha \leq 1$ the conventional Allan variance gives both an easier-to-interpret and an easier-to-calculate measure of stability.

It is interesting to make a graph of α versus μ for both the ordinary Allan variance and the modified Allan variance, such as the one shown in Fig. 12-5.

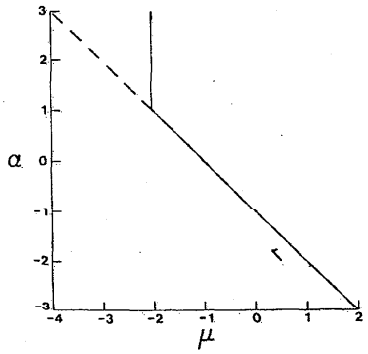

FIG. 12-5 Relationship between a power-law spectral density whose slope on a log-log plot is α and the corresponding sample variance whose slope on a log-log plot is μ. The solid line describes the behavior of the Allan variance, while the dashed line shows the advantage of the modified Allan variance for white phase noise and flicker phase noise.

This graph allows one to determine power-law spectra for noninteger as well as integer values of α. In the asymtotic limit the equation relating μ and α for the modified Allan variance is

$$\alpha = -\mu - 1 \qquad \text{for} \quad -3 < \alpha < 3. \tag{12-30}$$

12.1.6 Determination of the Mean Frequency and Frequency Drift of an Oscillator

Before the techniques of the previous four sections can be meaningfully applied to practical measurements, it is necessary to separate the deterministic and random components of the time deviation $x(t)$. Suppose, for example, that an oscillator has significant drift, such as might be the case for a quartz crystal oscillator. With no additional signal processing, the Allan variance would be proportional to τ^2. The variance of the Allan variance would be very small, further demonstrating that deterministic behavior has been improperly described in statistical terms and the oscillator's predictability is much better than the Allan variance indicated. Unfortunately, it is difficult to estimate the oscillator's deterministic behavior without introducing a bias in the noise at Fourier frequencies comparable to the inverse of the record length. In practice, it has been sufficient to consider two deterministic terms in $x(t)$:

$$x(t) = x_0 + (\Delta\nu/\nu_0)t + \tfrac{1}{2}Dt^2 + x_1(t) \tag{12-31}$$

The first term on the right-hand side is the synchronization error. The second term is due to imperfect knowledge of the mean frequency and is sometimes called syntonization error. The quadratic term, which results from frequency drift, is the most difficult problem for the statistical analysis because the Allan variance is insensitive to both synchronization and syntonization errors.

For white noise, the optimum estimate of the process is the mean. Therefore, a general statistical procedure that can be followed is to filter the data until the residuals are white (Allan *et al.*, 1974; Barnes and Allan, 1966). For example, at short times the frequency fluctuations of atomic clocks are usually white. Taking the first difference of Eq. (12-31), we find that

$$\bar{y}(t) = \frac{\Delta\nu}{\nu_0} + Dt + \frac{x_1(t+\tau) - x_1(t)}{\tau}, \tag{12-32}$$

and a linear least square fit to the frequency data yields the optimum estimate of $\Delta\nu$. However, the drift in atomic clocks is generally so small that the value obtained for D will not be statistically significant when τ is small enough to

satisfy the assumption of white frequency noise. Thus, we are led to consider the first difference of the frequency,

$$\frac{\bar{y}(t+\tau)-\bar{y}(t)}{\tau} = D + \frac{x_1(t+2\tau)-2x_1(t+\tau)+x_1(t)}{\tau^2}. \quad (12\text{-}33)$$

Many atomic clocks are dominated by random walk of frequency noise for long averaging times. Thus, the first difference of the frequency data (the second difference of the phase data) is white, and the optimum estimate of the drift is just the simple mean. If instead, a linear least square fit were removed from the frequency data in this region of τ, then the random-walk residuals would be biased, and it is likely that an optimistic estimate of $\sigma_y(\tau)$ would be obtained.

The optimum procedure would be different if the dominant noise type were flicker of frequency, rather than random walk. But there is no simple prescription that can be followed to estimate the drift in that case. Fortunately, a maximum likelihood estimate of the prameters for some typical cases has shown that the mean second difference of phase is still a good estimator of frequency drift in the sense that it introduces negligible bias in the Allan variance. Thus, in practice there is a simple prescription for computing the Allan variance in the presence of significant drift. Starting with the phase data, one forms the second difference and uses the simple average to estimate the mean. The value of τ chosen for creating these second differences must be long enough so that the predominant noise process is random-walk frequency modulation. After subtracting this estimate, the second-difference data is integrated twice to recover phase data with drift removed, and further analysis, including the computation of the Allan variance, may proceed. Figures 12-6 through 12-10 illustrate the estimation of drift. The quadratic dependence of the phase data in Fig. 12-6 nearly obscures the noise. The first difference of this data produces the nearly linear frequency dependence shown in Fig. 12-7, and the second difference produces the residuals shown in Fig. 12-8, which appear to be nearly white. Rigorous statistical analysis of this data indicates that the first difference of the frequency is indeed white with 90% confidence. Next, the mean frequency difference is subtracted. Then the residuals of Fig. 12-8 are integrated twice, and the result is the estimate of the phase deviation with drift removed shown in Fig. 12-9. Fig. 12-10 illustrates the Allan variance of this data calculated by three techniques. The squares were computed from the data of Fig. 12-6, while the open circles were computed following the recommended procedure for estimating the drift. The validity of the approach is illustrated by the black dots, which are the result of a statistically optimum parameter estimation procedure.

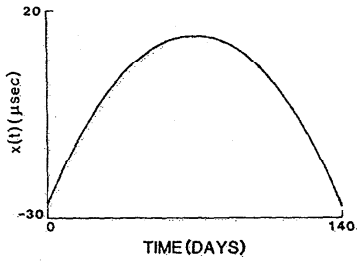

FIG. 12-6 Measured phase difference between a frequency standard and a reference during a 140-day experiment. The nearly quadratic form of the data effectively obscures the noise.

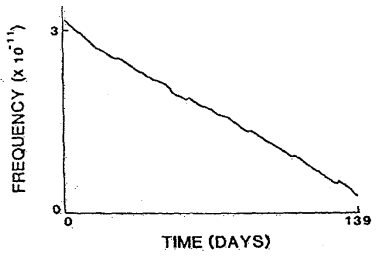

FIG. 12-7 One-day frequency averages obtained by taking the first differences of the data in Fig. 12-6. The ordinate is the fractional difference of the daily frequency from a nominal value. The nearly linear change in frequency with time is apparent, although the random deviations are visible.

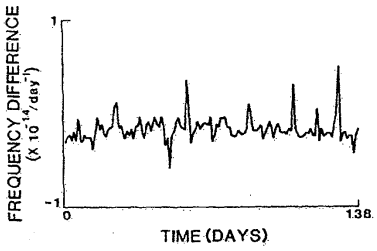

FIG. 12-8 Second difference of the data in Fig. 12-6. The second difference operation has removed the nonrandom behavior and the residuals appear to be nearly white.

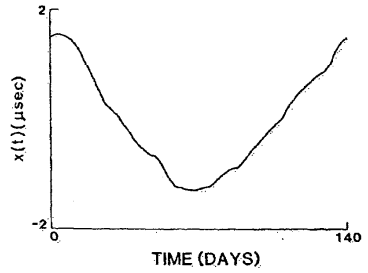

FIG. 12-9 Phase variations of the frequency standard due to the residuals, obtained by performing two integrations on the data of Fig. 12-8. The ordinate scale is expanded approximately 10 times compared to Fig. 12-6.

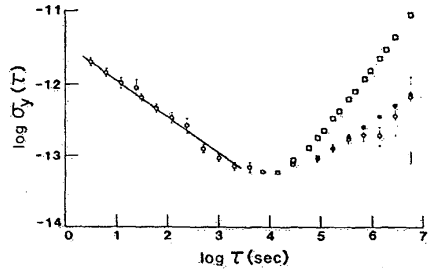

FIG. 12-10 Logarithm of the square root of the Allan variance as a function of the logarithm of the averaging time for three different computation methods. The squares were computed from the data of Fig. 12-6 and show the effect of the drift. The open circles were computed from the data of Fig. 12-9. The closed circles were computed using an optimum-parameter estimation procedure.

12.1.7 Confidence of the Estimate and Overlapping Samples

Consider three phase or time measurements of one oscillator relative to another at equally spaced intervals of time. From this phase data one can obtain two adjacent values of average frequency and one can calculate a single sample Allan variance (see Fig. 12-11). Of course, this estimate does not have high precision or confidence, since it is based on only one frequency difference.

For most commonly encountered oscillators, the first difference of the frequency is a normally distributed variable with zero mean. However, the square of a normally distributed variable is not normally distributed. This is so because the square is always positive and the normal distribution is

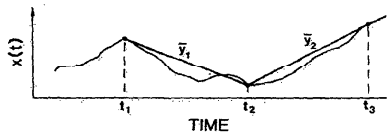

FIG. 12-11 Calculation of two average frequencies $\bar{y}_1$ and $\bar{y}_2$ by measuring the phase of an oscillator $x(t)$ at times t_1, t_2, and t_3.

completely symmetric, with negative values being as likely as positive ones. The resulting distribution is called a chi-squared distribution, and it has one "degree of freedom" since the distribution was obtained by considering the squares of individual (i.e., one independent sample), normally distributed variables (Jenkins and Watts, 1968).

In contrast, from five phase values four consecutive frequency values can be calculated, as shown in Fig. 12-12. It is possible to take the first pair and calculate a sample Allan variance. A second sample Allan variance can be calculated from the second pair (i.e., the third and fourth frequency measurements). The average of these two sample Allan variances provides an improved estimate of the true Allan variance, and one would expect it to have a tighter confidence interval than in the previous example. This could be expressed with the aid of the chi-squared distribution with two degrees of freedom.

However, there is another option. One could also consider the sample Allan variance obtained from the second and third frequency measurements, that is, the middle sample variance. This last sample Allan variance is not independent of the other two, since it is made up of parts of each of the others. But this does not mean that it cannot be used to improve the estimate of the true Allan variance. It does mean that the new average of three sample Allan variances is not distributed as chi squared with three degrees of freedom. The

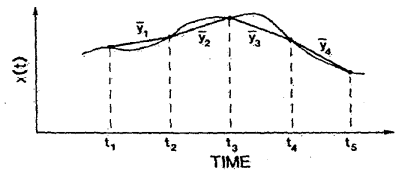

FIG. 12-12 Calculation of four frequency values $\bar{y}_1$, $\bar{y}_2$, $\bar{y}_3$, and $\bar{y}_4$ from five phase measurements at times t_1, t_2, t_3, t_4, and t_5. The sample variance formed from $\bar{y}_1$ and $\bar{y}_2$ and the one formed from $\bar{y}_3$ and $\bar{y}_4$ are independent. The sample variance formed from $\bar{y}_2$ and $\bar{y}_3$ is not independent of the other two but does contain some additional information useful in estimating the true sample variance.

number of degrees of freedom depends on the underlying noise type, that is, white frequency, flicker frequency, etc., and may have a fractional value.

Sample Allan variances are distributed as chi square according to the equation

$$\chi^2 = (\mathrm{df})s_y^2/\sigma_y^2, \tag{12-34}$$

where s_y^2 is the sample Allan variance, df the number of degrees of freedom (possibly not an integer), and σ_y^2 the true Allan variance, which we are interested in knowing but can only estimate imperfectly.

The probability density for the chi-squared distribution is given by the relation (Jenkins and Watts, 1968)

$$\rho(\chi^2) = \frac{1}{2^{\mathrm{df}/2}\Gamma(\mathrm{df}/2)}(\chi^2)^{\mathrm{df}/2-1}e^{-\chi^2/2}, \tag{12-35}$$

where $\Gamma(\mathrm{df}/2)$ is the gamma function, defined by the integral

$$\Gamma(t) = \int_0^\infty x^{t-1}e^{-x}\,dx. \tag{12-36}$$

A typical distribution is shown in Fig. (12-13).

Chi-squared distributions are useful in determining confidence intervals for variances and standard deviations, as shown in the following example. Suppose one has a sample variance $s^2 = 3.0$ and it is known that this variance has 10 degrees of freedom. The object is to calculate a range around the sample value of $s_y^2 = 3.0$ that *probably* contains the true value σ_y^2. The desired confidence is, say, 90%. That is, 10% of the time the true value will actually fall outside of the stated bounds. The usual way to proceed is to allocate 5% to the low end and 5% to the high end for errors, leaving 90% in the middle. This is arbitrary and a specific problem might dictate a different

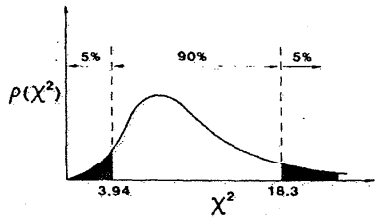

FIG. 12-13 Approximate form of a typical chi-squared distribution. For 10 degrees of freedom, 5% of the area under the curve corresponds to values of χ^2 less than 3.94, and an additional 5% corresponds to values of χ^2 greater than 18.3.

allocation. By referring to tables of the chi-squared distribution, one finds that for 10 degrees of freedom (df = 10) the 5% and 95% points correspond to

$$\chi^2(0.05) = 3.94, \qquad \chi^2(0.95) = 18.3. \tag{12-37}$$

Thus, with 90% probability the calculated sample variance $s_y^2 = 3$ satisfies the inequality

$$3.94 < (\mathrm{df})s_y^2/\sigma_y^2 < 18.3, \tag{12-38}$$

and this inequality can be rearranged in the form

$$1.64 < \sigma_y^2 < 7.61. \tag{12-39}$$

The estimate $s_y^2 = 3$ is a *point* estimate. The estimate $1.64 < \sigma_y^2 < 7.61$ is an *interval* estimate and should be interpreted to mean that 90% of the time the interval calculated in this manner will contain the true σ_y^2.

12.1.8 Efficient Use of the Data and Determination of the Degrees of Freedom

Typically, the sample variance is calculated from a data set using the relation

$$s^2 \cong \frac{1}{N-1} \sum_{n=1}^{N} (z_n - \bar{z})^2, \tag{12-40}$$

where it is implicitly assumed that the z_n's are random and uncorrelated (i.e., white) and where $\bar{z}$ is the sample mean calculated from the same data set. If all of this is true, then s^2 is chi-squared distributed and has $N - 1$ degrees of freedom.

Consider the case of two oscillators being compared in phase with N values of the phase difference obtained at equally spaced intervals τ_0. From these N phase values one obtains $N - 1$ consecutive values of average frequency, and from these one can compute $N - 2$ individual sample Allan variances (not all independent) for $\tau = \tau_0$. These $N - 2$ values can be averaged to obtain an estimate of the Allan variance at $\tau = \tau_0$.

The variance of this Allan variance has been calculated (Lesage and Audoin, 1973; Yoshimura, 1978). This approach is less versatile than the method of the previous section since it yields only symmetric error limits. However, it is simple and easy to use. Let $\Delta(N)$ be the relative difference between the sample Allan variance and the true value. Thus,

$$s_y^2 = [1 + \Delta(N)]\sigma_y^2(\tau). \tag{12-41}$$

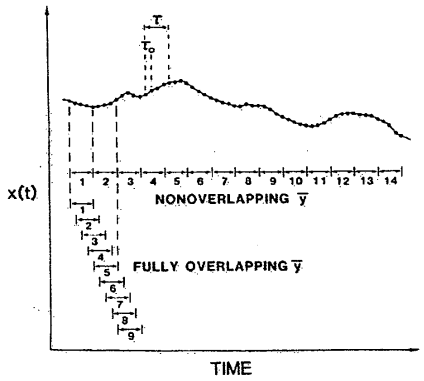

FIG. 12-14 Illustration of the case of $\tau = 4\tau_0$, for which the ratio of the number of fully overlapping to nonoverlapping estimates of the variance is more than 8 for the 57 phase points shown. When the averaging time for the computation of mean frequencies τ exceeds the sampling time τ_0, the number of fully overlapping mean frequencies is far larger than the number of nonoverlapping frequencies. In general, for large N approximately $2m$ times as many estimates of the sample variances can be computed using the fully overlapping technique.

TABLE 12-4

Number of Degrees of Freedom for Calculation of the Confidence of the Estimate of a Sample Allan Variance[a]

Noise type	df
White phase	$\dfrac{(N+1)(N-2m)}{2(N-m)}$
Flicker phase	$\exp\left[\ln\left(\dfrac{N-1}{2n}\right)\ln\left(\dfrac{(2m+1)(N-1)}{4}\right)\right]$
White frequency	$\left[\dfrac{3(N-1)}{2m} - \dfrac{2(N-2)}{N}\right]\dfrac{4m^2}{4m^2+5}$
Flicker frequency	$\dfrac{2(N-2)}{2.3N-4.9}$ for $m=1$
	$\dfrac{5N^2}{4m(N+3m)}$ for $m \geq 2$
Random-walk frequency	$\dfrac{N-2}{m}\dfrac{(N-1)^2-3m(N-1)+4m^2}{(N-3)^2}$

[a] For $\tau = m\tau_0$ from N phase points spaced τ_0 apart.

TABLE 12-5

Number of Degrees of Freedom for Calculation of the Confidence of the Estimate of a Sample Allan Variance for the Major Noise Types[a]

N	m	White phase	Flicker phase	White frequency	Flicker frequency	Random-walk frequency
9	1	3.665	4.835	4.900	6.202	7.000
	2	3.237	3.537	3.448	3.375	2.866
	4	1.000	1.000	1.000	1.000	0.999
129	1	65.579	79.015	84.889	110.548	127.000
	2	64.819	66.284	71.642	77.041	62.524
	4	63.304	52.586	42.695	36.881	29.822
	8	60.310	37.306	21.608	16.994	13.567
	16	54.509	22.347	9.982	7.345	5.631
	32	44.761	9.986	4.026	2.889	2.047
	64	1.000	1.000	1.000	1.000	1.000
1025	1	526.373	625.071	682.222	889.675	1023.000
	2	525.615	543.863	583.622	636.896	510.502
	4	524.088	459.041	354.322	316.605	253.755
	8	521.038	366.113	186.363	156.492	125.398
	16	514.952	269.849	93.547	76.495	61.241
	32	502.839	179.680	45.947	36.610	29.210
	64	478.886	104.743	21.997	16.861	13.288
	128	432.509	50.487	10.003	7.281	5.516
	256	354.914	17.429	4.003	2.861	2.005
	512	1.000	1.000	1.000	1.000	1.000

[a] N is the number of equally spaced phase points that are taken m at a time to form the averaging time.

12.1.9 Separating the Variances of the Oscillator and the Reference

A measured variance contains noise contributions from both the oscillator under test and the reference. The individual contributions are easily separated if it is known a priori that the reference is much less noisy than the device under test or equal to it in performance. Otherwise, the individual contributions can be estimated by comparing three devices (Barnes, 1966). The three possible joint variances are denoted by σ_{ij}^2, σ_{jk}^2, and σ_{ik}^2, while the individual device variances are σ_i^2, σ_j^2, and σ_k^2. The joint variances are composed of the sum of the individual contributions under the assumption that the oscillators are independent:

$$\begin{aligned} \sigma_{ij}^2 &= \sigma_i^2 + \sigma_j^2, \\ \sigma_{jk}^2 &= \sigma_j^2 + \sigma_k^2, \\ \sigma_{ik}^2 &= \sigma_i^2 + \sigma_k^2. \end{aligned} \tag{12-43}$$

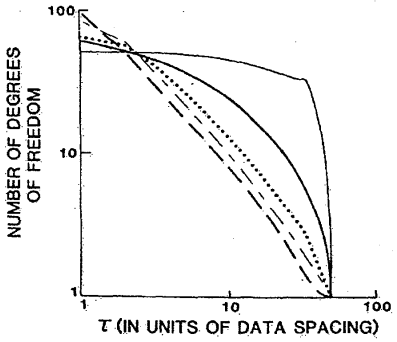

FIG. 12-15 Number of degrees of freedom as a function of averaging time for the case of 101 phase measurements: The heavy broken line is for random-walk frequency noise, the light broken line is for flicker frequency noise, the dotted line is for white frequency noise, the heavy solid line is for flicker phase noise, and the light solid line is for white phase noise.

An expression for each individual variance is obtained by adding two joint variances and subtracting the third:

$$\begin{aligned} \sigma_i^2 &= \tfrac{1}{2}(\sigma_{ij}^2 + \sigma_{ik}^2 - \sigma_{jk}^2), \\ \sigma_j^2 &= \tfrac{1}{2}(\sigma_{jk}^2 + \sigma_{ij}^2 - \sigma_{ik}^2), \\ \sigma_k^2 &= \tfrac{1}{2}(\sigma_{jk}^2 + \sigma_{ik}^2 - \sigma_{ij}^2). \end{aligned} \tag{12-44}$$

This method works best if the three devices are comparable in performance. Caution must be exercised since Eqs. (12-44) may give a negative sample Allan variance despite the fact that the true Allan variance is positive definite. This is possible because the confidence interval of the estimate is sufficiently large to include negative variances. Such a result is an indication that the confidence intervals of the sample Allan variances are too large and that more data is required.

12.2 DIRECT DIGITAL MEASUREMENT

12.2.1 Time-Interval Measurements

A common technique for measuring the phase difference between oscillators having nearly equal nominal frequencies is the use of direct time-interval measurements. In this section and those that follow, the symbols ν_{10} and ν_{20} are used to indicate the nominal values of ν_1 and ν_2, respectively. In the simplest form of this technique, a time-interval counter is started on some

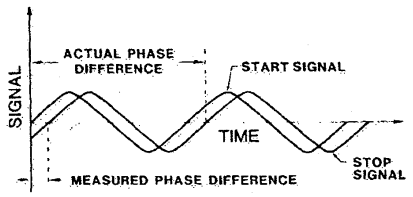

FIG. 12-16 The phase difference measured by a time-interval counter is the phase difference between the start signal and the stop signal modulo the period of the stop signal.

arbitrarily selected positive-going zero crossing of the signal from one oscillator (started on ν_{10} at time t_1) and stopped in the next positive-going zero crossing of the second oscillator (stopped on ν_{20} at time t_2). The measured time difference is

$$x_2(t_1) - x_1(t_1) \cong -P\tau_c[1 + (\nu_{20} - \nu_{10})/\nu_{10}], \tag{12-45}$$

where P is the reading of the time-interval counter and τ_c the period of its time base (Allan *et al.*, 1974). The units of the time difference is seconds of oscillator number 1. Equation (12-45) demonstrates an important characteristic of both time- and phase-difference measurements. Because of distortion the phase of an oscillator is generally not well known except at zero crossings. Thus, the quantity usually measured is $x_2(t_2) - x_1(t_1)$. However, all analysis techniques require the phase difference at the same time, and the translation requires a correction that takes into account the difference in frequency between the two oscillators. This correction is the reason for the second term in the brackets on the right-hand side of Eq. (12-45).

The simple scheme described above measures a maximum accumulated phase difference of one cycle of the signal. When the phase difference exceeds one cycle the counter reading is periodic, as shown in Fig. 12-16. This ambiguity can be reduced by dividing the signals from each oscillator before the time-interval measurement. The complete system is shown in Fig. 12-17. The effect of the divider is to increase the time interval before an ambiguity

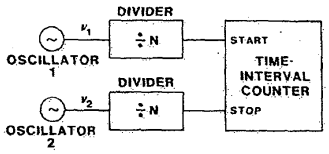

FIG. 12-17 Schematic diagram of the dividers used in conjunction with a time-interval counter to increase the maximum measurable phase difference to N cycles of the stop signal.

occurs to N/ν_{10}, where N is the divisor. Such measurement systems are used at many standards laboratories for the long-term measurement of atomic clocks, whose output is usually divided down to 1 pulse/sec. Since time-interval counters with resolution better than 0.1 nsec are available, this measurement scheme is suitable for long-term performance monitoring, yielding frequency-measurement precision of 10^{-14} for 1-day averages.

12.2.2 Frequency Measurements

Average frequency is measured most directly using a frequency counter. Used this way, the counter determines the number of whole cycles M occurring during a time interval τ given by the counter's time base. Thus

$$\bar{\nu}(0;\tau) = (M + \Delta M)/\tau \simeq M/\tau, \tag{12-46}$$

where $\bar{\nu}(t_1;t_2)$ denotes the average frequency over the interval from t_1 to t_2 and ΔM, the fractional cycle, is not measured by the counter. The starting time is arbitrarily called $t = 0$. Thus, the quantization error is given by

$$\Delta\nu_Q/\langle\nu\rangle < 1/M. \tag{12-47}$$

12.2.3 Period Measurements

For low frequencies, the number of cycles counted may be small and the quantization error can be very large. By measuring the period instead of the frequency, it is possible to decrease the error without increasing the duration of the measurement. A period counter measures the duration of M whole cycles of the signal as N cycles of the time base τ_c. The fraction of a cycle ΔN is not measured. Thus, we have

$$M = \bar{\nu}(0; M/\nu_0)(N + \Delta N)\tau_c, \tag{12-48}$$

and therefore

$$\bar{\nu}(0; M/\nu_0) \cong M/N\tau_c \tag{12-49}$$

and the quantization error is

$$\Delta\nu_Q/\langle\nu\rangle < 1/N. \tag{12-50}$$

Frequency measurements are almost never used to characterize precision oscillators, but period measurements are very common. A straightforward extension of this method eliminates the bias potentially introduced by the quantization error and permits the measurement of accumulated phase. The counter must be capable of being read without halting the counting process.

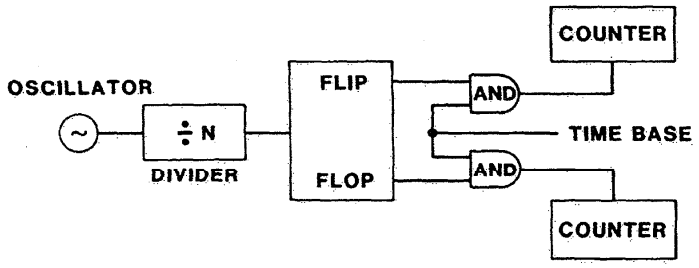

FIG. 12-18 Two-counter system to eliminate dead time in period measurements. The two counters alternately count the number of cycles of the time base in N periods of the oscillator under test.

Alternatively, a second counter may be used to begin counting the same time base when the first counter stops. The second approach is illustrated in Fig. 12-18. This type of measurement system is sometimes called a chronograph.

12.3 SENSITIVITY-ENHANCEMENT METHODS

12.3.1 Heterodyne Techniques

It is possible for oscillators to be very stable, and values of $\sigma_y(\tau)$ can be as small as 10^{-16} in some state-of-the-art standards. Thus, one often needs measuring techniques capable of resolving very small fluctuations in $y(t)$. One of the most common techniques is the heterodune or beat-frequency technique. In this method the signal from the oscillator under test is mixed with a reference signal of almost the same frequency so that one is left with a lower average frequency for analysis without reducing the frequency (or phase) fluctuations themselves.

In principle, it is possible to analyze the most general measurement case, where no restrictions are placed on the average frequency or phase difference between the two oscillators under test. Equation 12-1 can be inverted as

$$2\pi\nu_0 t + \phi(t) = \arcsin[V(t)/V_0] \tag{12-51}$$

and used to obtain the series $\phi(m\tau)$ by sampling the voltage at regular time intervals. This direct technique is not used, because it requires unobtainable mixer performance characteristics. The high-level rf signals that are required for low-noise phase measurements produce significant harmonic distortion, so that the output of the phase detector deviates significantly from a sine wave. Furthermore, the distortions are generally sensitive to level and environmental perturbations. However, the phase relationships among the various harmonics are very stable, so it is possible to use the repetition of one

point on the waveform in order to count cycles. The positive-going zero crossings are normally chosen in order to provide immunity from changes in both the amplitude and symmetry of the waveform.

Consider two signals whose frequency difference is much less than the frequency of either oscillator:

$$V_1(t) = V_{10} \sin[2\pi\nu_{10}t + \phi_1(t) + \phi_{10}]$$

and (12-52)

$$V_2(t) = V_{20} \sin[2\pi\nu_{20}t + \phi_2(t) + \phi_{20}],$$

where $|\nu_{10} - \nu_{20}| \ll \nu_{10}$ and the constants ϕ_{10} and ϕ_{20} represent the nominal phases of the two signals.

Suppose that the two signals are mixed in a linear product detector and filtered so that the signal at the sum frequency $\nu_{10} + \nu_{20}$ is highly attenuated. The result is

$$V(t) \cong V_0 \cos[2\pi(\nu_{10} - \nu_{20})t + \phi_{10} - \phi_{20} + \phi_1(t) - \phi_2(t)], \quad (12\text{-}53)$$

which may be characterized by any of the measurement techniques discussed in Section 12.2. The amplitude V_0 of the mixer output is a function of the mixer design, the input amplitudes, and the output termination (Walls *et al.*, 1976). Using the definition (12-10), we find that for the heterodyned signal

$$x_H(t) = (1/2\pi\nu_H)\,\Delta\phi(t), \quad (12\text{-}54)$$

where

$$\nu_H = |\nu_{10} - \nu_{20}| \quad (12\text{-}55)$$

and

$$\Delta\phi(t) = \phi_1(t) - \phi_2(t). \quad (12\text{-}56)$$

Equation (12-54) may be rewritten as

$$x_H(t) = (\nu_0/\nu_H)x(t), \quad (12\text{-}57)$$

from which we conclude that a given phase change corresponds to a larger time deviation for the heterodyne signal than for the original signal. As a result, the quantization error for the period measurement technique is reduced by the factor ν_H/ν_0.

12.3.2 Homodyne Techniques

The limit of the heterodyne method, called homodyne, occurs when $\nu_{10} = \nu_{20}$. In this case the output of the phase detector is given by

$$V(t) \cong V_0 \cos[\phi_{10} - \phi_{20} + \phi_1(t) - \phi_2(t)]. \quad (12\text{-}58)$$

The analysis of phase noise is accomplished by arranging that $\phi_{10} - \phi_{20} = \pi/2$, which can be achieved with a phase shifter. Then,

$$V(t) \cong -V_0 \sin[\phi_1(t) - \phi_2(t)] \cong V_0[\phi_2(t) - \phi_1(t)]. \tag{12-59}$$

There are various methods by which one can control the signal $V_2(t)$ so that $\nu_{10} = \nu_{20}$ without producing significant correlation between $\phi_2(t)$ and $\phi_1(t)$. When any one of these methods is used, it is possible to use $V(t)$ as a measure of $\phi(t)$. Two methods, delay lines and phase-locked loops (Gardner, 1966), are described below.

12.3.2.1 DISCRIMINATOR AND DELAY LINE

The circuit of a discriminator or delay-line system for measuring phase noise is illustrated in Fig. 12-19. The delayed signal is given by

$$V_2(t) = V_1(t - t_d) = V_{20} \sin[2\pi\nu_{10}(t - t_d) + \phi_1(t - t_d) + \phi_{10} + \phi_s]. \tag{12-60}$$

When the phase shifter is set for quadrature, $\phi_s - 2\pi\nu_{10}t_d = \pi/2$ and

$$V_2(t) = V_{20} \sin[2\pi\nu_{10}t + \phi_1(t - t_d) + \phi_{10} + \pi/2]. \tag{12-61}$$

The output of the phase detector is given by

$$V(t) = V_0[\phi_1(t - t_d) - \phi_1(t)]. \tag{12-62}$$

Substituting Eq. (12-62) into Eq. (12-20), we obtain

$$\bar{y}(t - t_d; t) = -V(t)/2\pi\nu_0 V_0 t_d \tag{12-63}$$

and we see that the delay-line method can be used to produce samples of $\bar{y}(m\tau_0)$ by varying the delay time. However, the technique is used more frequently with a fixed delay by restricting its application to the region of τ much greater than the delay time, so that $\bar{y}(t - t_d; t)$ is a good approximation for the instantaneous frequency. Under this assumption spectrum analysis of

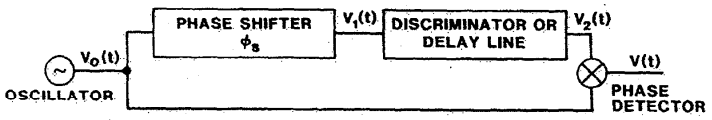

FIG. 12-19 A delay-line phase-noise-measurement system. When the phase shifter is adjusted so that $V_2(t)$ is in phase quadrature with $V_0(t)$, the output of the phase detector is approximately equal to the instantaneous frequency deviation of the oscillator. The spectral density of the source may be estimated for Fourier frequencies small compared to the inverse of the delay time.

: signal from the mixer can be used to estimate the spectral density of the quency fluctuations:

$$S_y(f) \cong \frac{1}{(2\pi\nu_0 t_d)^2} S_{V/V_0}(f) \qquad \text{for} \quad f \ll 1/\pi t_d. \tag{12-64}$$

equency discriminators are applied in an analogous fashion. A resonant cuit is often used to provide discrimination since it produces a phase shift oportional to the frequency deviation from the resonant frequency. For ample, the phase shift on reflection from a resonance with loaded quality tor Q is

$$\phi = \arctan(2Qy) \cong 2Qy, \tag{12-65}$$

ovided that the frequency deviation is small compared to the bandwidth of e resonance and the applied signal is nearly at the center frequency of the scriminator. This can be accomplished either manually or with a frequency- cked loop. The design of such a loop is similar to the phase-locked loop of e next section. Once again, one can spectrum analyze the signal from the ixer to obtain

$$S_y(f) \simeq \frac{1}{(2Q)^2} S_{V/V_0}(f) \qquad \text{for} \quad f \ll \nu_0/Q. \tag{12-66}$$

ıe noise floor for measurements made with either a delay line or discrimi- ıtor normally results from white voltage noise in the analysis circuitry and independent of the Fourier frequency. We denote the noise floor S_{V/V_0} inimum) and find the noise floor for frequency or phase measurements by

$$S_\phi(\text{noise limit}) = \frac{\nu_0^2}{f^2} S_y(\text{noise limit}) = \frac{\nu_0^2}{f^2(2Q)^2} S_{V/V_0}(\text{minimum}). \tag{12-67}$$

onsequently, the discriminator or delay-line technique is limited in sensi- ity since the output voltage is proportional to the frequency deviations. reater sensitivity is possible using two oscillators in a phase-locked loop. ıe noise in the reference is an important consideration, even though the ference is passive in the case of a discriminator or a delay line. If the cillator has sufficiently low noise, then the circuits described measure the ıriations of the discriminator center frequency or the delay variations in the lay line.

2.3.2.2 PHASE-LOCKED LOOP

The block diagram for the most general phase-locked loop that will be nsidered here is shown in Fig. 12-20. The noise voltage summed into the op is a schematic way of representing $\phi_n(t)$, the open-loop phase noise of the

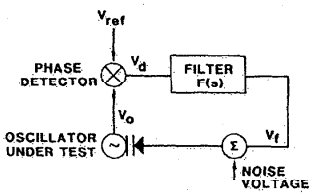

FIG. 12-20 Block diagram of a phase-locked loop. The order of the loop is determined by the filter transfer function. For convenience, noise in the oscillator under test is introduced at the summing junction.

oscillator under test. Phase noise in the reference oscillator is denoted by $\phi_{ref}(t)$.

The purpose of using a phase-locked loop is simply to guarantee that the two oscillators are, on the average, in phase quadrature. When the oscillators are near quadrature, the voltage output of the phase detector is proportional to the difference in phase between the two output signals.

Analysis of the phase-locked loop yields the result

$$\phi_o(s) = \phi_n(s)\left[\frac{1}{1 + G_{eq}(s)}\right] + \phi_{ref}(s)\left[\frac{G_{eq}(s)}{1 + G_{eq}(s)}\right], \tag{12-68}$$

where $G_{eq}(s)$ is the open-loop transfer function defined by

$$G_{eq}(s) = \frac{K_o K_d F(s)}{s} \tag{12-69}$$

and $\phi_n(s)$ and $\phi_{ref}(s)$ are the Laplace transforms of the corresponding time-varying quantities. We can also calculate the voltage output of the phase detector,

$$V_d(s) = \frac{K_d[\phi_{ref}(s) - \phi_n(s)]}{1 + G_{eq}(s)}, \tag{12-70}$$

as well as the feedback voltage to the varactor,

$$V_f(s) = F(s)V_d(s) = \frac{sG_{eq}(s)}{G_{eq}(s)}[\phi_{ref}(s) - \phi_n(s)]. \tag{12-71}$$

Assuming that the phase noise of the two oscillators is not correlated,

$$S_{V_d}(\omega) = \frac{K_d^2}{|1 + G_{eq}(j\omega)|^2}[S_{\phi_{ref}}(\omega) + S_{\phi_n}(\omega)], \tag{12-72}$$

$$S_{V_f}(\omega) = \frac{\omega^2 G_{eq}(j\omega)^2}{|1 + G_{eq}(j\omega)|^2}[S_{\phi_{ref}}(\omega) + S_{\phi_n}(\omega)]. \tag{12-73}$$

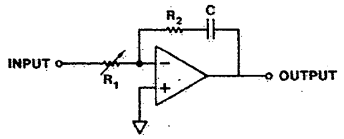

FIG. 12-21 Circuit diagram of the most common loop filter for a second-order phase-ked loop. Resistor R_2 is required for stable operation. Capacitor C provides the low-quency gain needed to reduce the phase errors of the first-order loop.

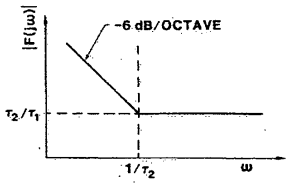

FIG. 12-22 Bode plot for the loop filter of Fig. 12-21.

hus, if we know the behavior of $G_{eq}(j\omega)$, then we can relate the measured ectrum of the voltage at the output of the phase detector or at the varactor ner to the sum of the spectral densities of the phase noise of the two scillators.

The loop filter is often chosen to be a pure gain. The resulting first-order op has a significant drawback: the two oscillators are offset from quadra-ire by a phase shift proportional to their open-loop frequency difference. In rder to maintain system calibration, the operator must remove the fre-uency offset from time to time. This problem can be eliminated by using a cond-order loop. Figure 12-21 illustrates one loop filter that can be used to chieve the desired frequency response. The transfer function of this filter is

$$F(s) = (1 + s\tau_2)/s\tau_1, \tag{12-74}$$

here $\tau_2 = R_2C$ and $\tau_1 = R_1C$. Figure 12-22 shows the Bode plot of the equency-response function of this filter. Substitution of Eq. (12-74) into Eq. 2-69) yields the open-loop frequency-response function

$$G_{eq}(j\omega) = -\frac{\omega_n^2 + 2j\zeta\omega_n\omega}{\omega^2}, \tag{12-75}$$

here

$$\omega_n = [K_0 K_a \quad]^{1/2} \tag{12-76}$$

nd

$$\zeta = \tfrac{1}{2}\tau_2\omega_n. \tag{12-77}$$

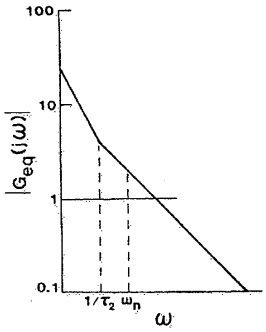

FIG. 12-23 Bode plot of the open-loop frequency-response function for a phase-locked loop having the loop filter of Fig. 12-21. Parameters were chosen to illustrate a stable condition.

The first requirement to be satisfied by the loop parameters is that the closed loop be stable. Since the transfer function $G_{eq}(s)$ has no poles or zeros for $s > 0$, a sufficient requirement for the phase-locked loop to be stable is that the slope of the Bode plot of $|G_{eq}(j\omega)|$ be less steep than -12 dB/octave at the point where $|G_{eq}(j\omega)| = 1$. The Bode plot of $|G_{eq}(j\omega)|$ is shown in Fig. 12-23 for a case where the loop operation is stable.

It is desirable for the loop to be nearly critically damped, that is, $\zeta = 1$. At critical damping the natural frequency of the loop is related to τ_2 by

$$\omega_{n,\zeta=1} = 2/\tau_2. \tag{12-78}$$

Under the same conditions the unity gain frequency is

$$\omega_{1,\zeta=1} = 4.12/\tau_2. \tag{12-79}$$

The second requirement to be satisfied by the phase-locked loop is related to the accuracy with which spectral-density measurements can be made. Substitution of Eq. (12-75) into Eq. (12-72) yields

$$S_{V_d}(\omega) = \frac{K_d^2\omega^4}{(\omega^2 - \omega_n^2)^2 + 4\zeta^2\omega^2\omega_n^2}[S_{\phi_{ref}}(\omega) + S_{\phi_n}(\omega)]. \tag{12-80}$$

Since the proportionality factor has a high pass response, it is possible to use an essentially constant calibration to relate $S_{V_d}(\omega)$ and $S_{\phi}(\omega)$. For example, if we require that

$$S_{V_d}(\omega) \cong K_d^2[S_{\phi_{ref}}(\omega) + S_{\phi_n}(\omega)] \tag{12-81}$$

with no more than 10% error for all Fourier frequencies greater than 2π rad/sec, then for the critically damped loop the requirement on τ_2 is $\tau_2 > 1.4$ sec.

The third requirement on loop performance is that the frequency offset tween the two oscillators produce negligible phase shift of the oscillators)m quadrature. In the ideal loop the phase error for a frequency error $\Delta\nu$ troduced at time $t = 0$ is

$$\phi_{\text{error}} = 2\pi\,\Delta\nu t\, e^{-\omega_n t}.$$

owever, in the actual circuit there is a finite phase error due to the limited op gain of the amplifier of Fig. 12-21. Nevertheless, the phase error is duced by 10^5 compared to its value for a first-order loop. Typically, the ror is less than the residual phase error due to the voltage offset at the mixer ıtput and should be much less than 1°.

The feedback loop reduces the sensitivity of the system for measurements the phase spectral density for Fourier frequencies less than the unity-gain)quency of the phase-locked loop. One way to avoid this problem is to ilize the feedback voltage V_f. Substituting Eq. (12-75) into Eq. (12-73), we ıd that

$$S_{V_f}(\omega) = \frac{2\pi\nu_0^2(\omega_n^4 - 4\zeta^2\omega_n^2\omega^2)}{(\omega^2 - \omega_n^2)^2 + 4\zeta^2\omega_n^2\omega^2}\,[S_{y_{\text{ref}}}(\omega) + S_{y_n}(\omega)]. \qquad (12\text{-}82)$$

)r this case, the proportionality factor has a low pass response and a ınstant calibration factor may be used to relate $S_{V_f}(\omega)$ to $S_y(\omega)$.

3.3 Multiple Conversion Methods

Quite often the beat frequency between the signal under test and the)oratory reference is unsuitable or inconvenient for frequency-stability :asurements. The frequency may be too high for the available counters or e heterodyne factor may be too small to yield the required noise enhance- :nt. Under these circumstances a second mixing stage in series with the first n be used to produce the desired beat frequency. On the other hand, the rect beat frequency between two oscillators may be too small. For example, e frequencies of commercial cesium-beam frequency standards are usually close together that the beat frequency between two devices would be near cycle/day, making it impossible to observe the stability at shorter times. ıis limitation can be overcome by the use of two parallel mixing stages.

3.3.1 FREQUENCY SYNTHESIS

A commercial frequency synthesizer is usually the most convenient way to oduce arbitrary reference frequencies for stability measurements. A mixing ıge preceding the synthesizer can be used both to bring the signal into the

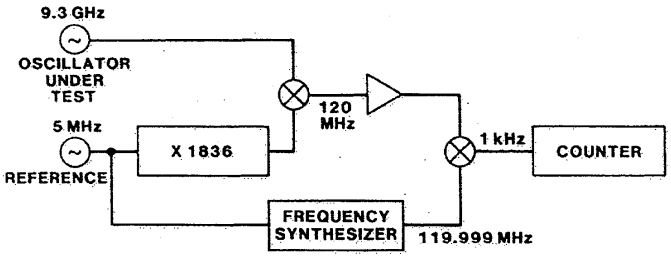

FIG. 12-24 Use of frequency synthesis to measure oscillators whose frequency differs significantly from the available low-noise reference. It may be necessary to use a frequency multiplier to bring the signal into the range of the available synthesizer or to overcome the synthesizer's phase noise.

appropriate range and to enhance the oscillator noise compared to the short-term phase noise of the synthesizer. Figure 12-24 demonstrates both aspects of the technique.

The initial mixing stage from the microwave frequency to the rf results in a substantial heterodyne factor, 77.5 for the example chosen. The output of the first conversion stage lies within the range of low-noise commercial frequency synthesizers, which makes it possible to obtain a fixed, low beat frequency over a wide range of input frequencies. The initial mixing stage also reduces the frequency synthesizer's contribution to the measurement-system noise. Figure 12-25 shows the typical phase excursions of a high-quality commercial synthesizer operated near 5 MHz.

Under some circumstances a frequency divider may be used to provide the signal for the second mixing stage, as shown in Fig. 12-26. This technique has the disadvantage of requiring a custom divider but results in much lower measurement noise than the direct use of a synthesizer with a single heterodyne stage.

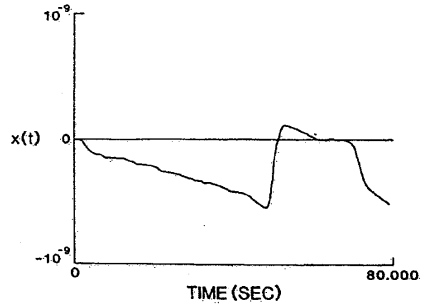

FIG. 12-25 Typical phase excursions of a commercial frequency synthesizer.

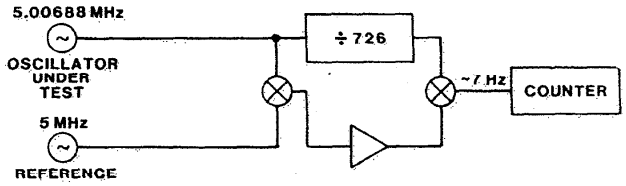

FIG. 12-26 Use of a simple divider as a substitute for a commercial frequency synthesizer in a heterodyne measurement system. Better noise performance can result from the initial mixing stage.

12.3.3.2 THE DUAL-MIXER TIME-DIFFERENCE TECHNIQUE

There is no best answer to the question of how to make frequency-stability measurements. However, by combining versatility with low-noise performance, the dual-mixer time-difference technique (Cutler and Searle, 1966; Allan and Daams, 1975) shown in Fig. 12-27 comes close to the ideal. The original motivation for this method was to use a transfer oscillator and two mixers in parallel to permit short-term frequency-stability measurements between oscillators that have an inconveniently small frequency difference. The transfer oscillator is most easily realized with a frequency synthesizer locked to one of the oscillators, designated oscillator 1 in Fig. 12-27. By convention the frequency of the synthesizer is set low compared to the oscillator under test, so we write the frequency of the synthesizer as

$$\nu_s = \nu_1(1 - 1/R). \tag{12-83}$$

The constant R is equal to the heterodyne factor, which can be seen by calculating the beat frequency between oscillator 1 and the synthesizer:

$$\nu_{B1} = \nu_1 - \nu_s = \nu_1/R. \tag{12-84}$$

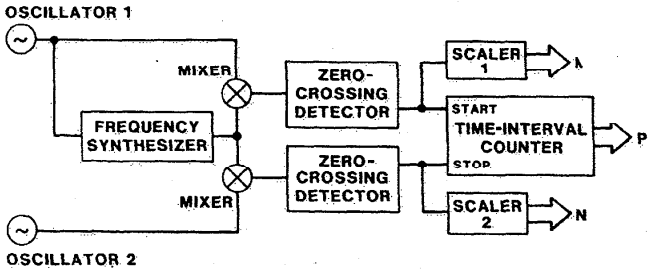

FIG. 12-27 A dual-mixer measurement system. The scalars measure the number of whole cycles of elapsed phase, while the time-interval counter measures the fractional cycle.

The combination of oscillator, frequency synthesizer, and mixer functions as a divider and scaler 1 functions as the system clock, recording elapsed time in units of cycles of oscillator 1.

The signals from oscillators 1 and 2 are represented according to Eq. (12-52) with $\phi_{10} = \phi_{20} = 0$, and the signal from the synthesizer is written

$$V_s(t) = V_{s0} \cos[2\pi\nu_{s0}t + \phi_s(t)]. \tag{12-85}$$

The phase of the synthesizer retards nearly linearly in time compared to the phase of oscillators 1 and 2. At time t_M the synthesizer reaches phase quadrature with oscillator 1 and the beat signal crosses zero (in the positive direction), producing a pulse from the zero-crossing detector and starting the time-interval counter. At a later time t_N the continued sweep of the synthesizer has brought it into quadrature with oscillator 2, and a pulse is produced that stops the time-interval counter. The phase difference between the oscillators can be written in terms of the three counter readings:

$$\phi_2(t_M) - \phi_1(t_M) = 2(N - M)\pi - 2\pi[\bar{\nu}_{B2}(t_M; t_N)]\tau_c P, \tag{12-86}$$

where N is the reading of scaler 2, M the reading of scaler 1, P the reading of the time-interval counter, and τ_c the period of its time base (Stein *et al.*, 1983). Comparison with Eq. (12-45) for direct time-interval measurements reveals that the role of the scalers is to accumulate the coarse phase difference between the oscillators, while the time-interval counter provides fine-grain resolution of the fractional cycle. This process is illustrated in Fig. 12-28. The advantage of the technique over direct time-interval measurements is that the noise performance is improved by the large heterodyne factor, allowing time resolution of 0.1 psec to be obtained. The synthesizer degrades the noise performance very little since it contributes to the noise only over the interval $\tau_c P$.

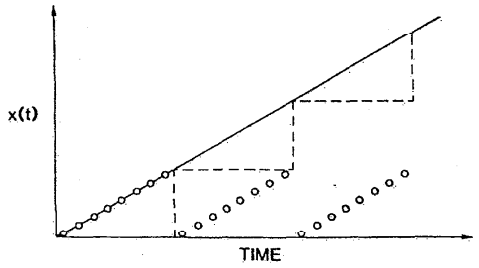

FIG. 12-28 Total elapsed phase measured by the dual-mixer system of Fig. 12-27 (solid line). This phase measurement consists of two components: the number of full cycles that have elapsed is the step function plotted as a dashed line; the fractional cycle is the saw-tooth function plotted as open circles.

The average beat frequency $\bar{\nu}_{B2}(t_M; t_N)$ cannot be known exactly, but it may ɔe estimated with sufficient precision if it changes slowly compared to the nterval between measurements. If the primed and unprimed variables epresent two independent measurements, then

$$\bar{\nu}_{B2}(t_M; t_N) \cong (N' - N)/[R(M' - M)/\nu_{10} + \tau_c(P' - P)]. \quad (12\text{-}87)$$

2.3.3.3 FREQUENCY MULTIPLICATION

A frequency multiplier produces n full cycles of the output signal for each ycle of the input signal, where n is an integer determined by the design of the levice. Such a device is also a phase multiplier, that is, the total phase ıccumulation of the output signal is n times as great as the phase accumuation of the input signal:

$$\Phi_{out} = 2\pi\nu_{out}t + \phi_{out}(t) = 2\pi(n\nu_{in})t + n\phi_{in}(t). \quad (12\text{-}88)$$

t follows that the spectral density of the output signal is enhanced by a factor ɔf n^2 compared to the input signal,

$$S_{\phi_{out}}(f) = n^2 S_{\phi_{in}}(f).$$

naking it easier to perform the necessary noise measurements. Similarly, it is ılso easier to make Allan-variance measurements. If the oscillator under test ınd the reference are both multiplied by the same factor, the beat frequency vill be n times larger than with no multiplication but the heterodyne factor vill be the same. The zero crossings that must be detected by the counter ıave n times higher slope and more easily overcome the voltage noise in the ɔounter trigger circuits. The ability to measure frequency stability is only ɔnhanced if the multipliers have extremely low phase noise themselves. This is he case for many modern multipliers that are triggered by the zero crossings ɔf the input signal. As a result, the use of multipliers can reduce the ɔerformance requirements on the phase detector and the following low-noise ımplifiers.

12.4 CONCLUSION

The IEEE recommendations have achieved the goal of introducing ıubstantial uniformity in the specification of oscillator performance. The Allan variance and the one-sided power spectral density of phase have proved ıufficient to evaluate oscillators for all common applications. In a few cases nore specialized measures are helpful in relating performance to the specific ıpplication. For example, the rms time-prediction error is helpful in judging a ɔlock's ability to keep time over long intervals (Allan and Hellwig, 1978).

However, the specialized performance measures are generally calculable in terms of the IEEE recommended measures.

Significant progress has been made during the last 15 years in measurement techniques and data processing. These advances have obscured the dividing line between the frequency domain and the time domain. Today the spectral density and the variance are most often computed from the identical input data set, the equally spaced time series of the phase deviations. The choice of a specific measurement setup can be made mostly on a cost versus performance basis. Perhaps the biggest advance in commercially available equipment is the introduction of heterodyne measurement techniques for time-domain (counter-based) measurements. As a result, the noise performance of these systems has improved dramatically.

One recommendation that should be made is to perform measurements as high up in the measurement hierarchy as possible. Direct measurement of the phase deviation is most desirable. This approach places the largest share of the burden on the measurement equipment, minimizes long-term errors, and maximizes data processing flexibility.