

Figure 1. A repeating sine function is the basis of an oscillating signal.

sine wave repeats itself. A cycle (2π radians of phase) of the oscillation is produced in one period T .

It is convenient for us to express angles in radian units rather than in units of degrees, and positive zero crossings of the voltage will occur every 2π radians. The *frequency* ν is the number of cycles in one second (Hz), which is the reciprocal of period (seconds per cycle). The expression describing the voltage V produced by a sine-wave signal generator is given by

$$V(t) = V_0[1 + a(t)] \sin[\Phi(t)]$$

where V_0 is the peak voltage amplitude, $a(t)$ is amplitude noise, and $\Phi(t)$ is the total accumulated phase. Equivalent expressions are

$$V(t) = V_0[1 + a(t)] \sin\left(2\pi \frac{t}{T}\right)$$

and

$$V(t) = V_0[1 + a(t)] \sin(2\pi\nu t)$$

For the following discussion, we will assume the amplitude noise $a(t)$ is zero. Consider Fig. 2. Let's assume that the maximum value of V equals 1, hence $V_0 = 1$. We say that the voltage $V(t)$ is normalized to unity.

If we are given the frequency of the sine wave, then no matter how big or small Δt may be, we can determine ΔV . Let us look at this from another point of view. Suppose that we can measure ΔV and Δt . From this, there is a sine wave at a unique minimum frequency corresponding to the given ΔV and Δt . For infinitesimally small Δt , this frequency is called the *instantaneous frequency* at this t . The smaller the interval Δt , the better the approximation of

FREQUENCY STABILITY

DAVID A. HOWE
National Institute of Standards
and Technology (NIST)
Boulder, Colorado

1. THE SINE WAVE AND STABILITY

A sine-wave signal generator produces a voltage that changes in time in a sinusoidal manner as shown in Fig. 1. The signal is an *oscillating signal* because the




Figure 2. For a given phase, ΔV vs. Δt of the sine-wave signal corresponds to a unique minimum frequency called the instantaneous frequency if Δt is diminishingly small.

instantaneous frequency at t . In practice, because of finite bandwidths, we cannot measure the instantaneous frequency.

When we speak of oscillators and the signals they produce, we recognize that an oscillator has some nominal frequency at which it operates. The “frequency stability” of an oscillator is a term used to characterize how small the frequency fluctuations of the oscillator signal are. We usually refer to frequency stability when comparing one oscillator with another. As we shall see later, we can define particular aspects of an oscillator’s output, then draw conclusions about its relative frequency stability. People often speak of “frequency stability” when they actually mean “frequency instability.” Frequency stability is the degree to which an oscillating signal produces the same value of frequency for any interval Δt throughout a specified period of time. An internationally recommended definition of “frequency instability” is: “The spontaneous and/or environmentally caused frequency change within a given time interval.”¹

Let’s examine the two waveforms shown in Fig. 3. Frequency stability depends on the amount of time involved in its measurement. Of the two oscillating signals, it is evident that “2” is more stable than “1” from time t_1 to t_3 assuming that the horizontal scales are linear in time. From time t_1 to time t_2 , there may be some question as to which of the two signals is more stable, but it’s clear that from time t_2 to time t_3 , signal “1” is at a frequency different from that in interval t_1 to t_2 .

If we want an oscillator to produce a particular frequency ν_0 , then we’re correct in stating that if the oscillator signal frequency deviates from ν_0 over any interval, this is a result of something that is undesirable. In the design of an oscillator, it is important to consider the sources of mechanisms that degrade the oscillator’s frequency stability. These undesirable mechanisms cause random (noise) or systematic processes to exist on top of the sine-wave signal of the oscillator. To account for the noise components at the output of a sine-wave signal generator, we can express the output as

$$V(t) = V_0[1 + a(t)] \sin[2\pi\nu_0 t + \phi(t)] \quad (1)$$

where

- V_0 = nominal peak voltage amplitude
- $a(t)$ = deviation of amplitude from nominal (i.e., $\delta V/V_0$)
- ν_0 = nominal fundamental frequency
- $\phi(t)$ = deviation of phase from nominal

Ideally a and ϕ should equal zero for all time. However, in the real world there are no perfect oscillators. To determine the extent of the noise components a and ϕ , we turn our attention to measurement techniques.

¹The present IEEE standard for the measure of frequency stability is the one-sided spectral density $S_y(f)$ in the frequency domain or the two-sample or Allan variance $\sigma_y(\tau)$ in the time domain. These are explained later.




Figure 3. Top (1): Instantaneous frequency is inconsistent and less stable from t_2 to t_3 . Bottom (2): Instantaneous frequency is consistent and more stable throughout.

The typical precision oscillator, of course, has a presumably stable sinusoidal voltage output with a frequency ν and a period of oscillation T (which is the reciprocal of the frequency: $\nu = 1/T$). One goal is to measure the frequency and/or the frequency stability of the sinusoid. Instability is actually what is measured, but with little confusion it is usually called “stability” in the literature. Naturally, fluctuations in frequency correspond to fluctuations in the period. Almost all frequency measurements, with very few exceptions, are measurements not of frequency but of the phase or of the period fluctuations in an oscillator, 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.

We must also realize that any frequency measurement involves two oscillators. In some instances, one oscillator is in the counter. It is impossible to purely measure only one oscillator. In some instances one oscillator may sufficiently outperform the other, and 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

$$y(t) = \frac{\nu_1 - \nu_0}{\nu_0} \quad (2)$$

as the fractional frequency difference or offset of oscillator one ν_1 with respect to a reference oscillator ν_0 divided by the nominal frequency ν_0 . Conceptually, we can also think of Eq. (2) as the free-running frequency of an individual oscillator ν_1 , differentiated with respect to its own nominal value ν_0 . Now, $y(t)$ is a dimensionless quantity and useful in describing oscillator and clock performance; that is, the time fluctuation or difference $x(t)$ of an oscillator over a period of time t is given simply by

$$x(t) = \int_0^t y(t') dt' = \frac{\phi(t)}{2\pi\nu_0} \quad (3)$$

We see that the time deviations and the phase deviations are related by a constant, $1/2\pi\nu_0$. Since it is impossible to

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

$$\bar{y}(t) = \frac{x(t + \tau) - x(t)}{\tau} \tag{4}$$

Tau (τ) may be called the sample time or averaging time; it may be determined, for example, by the gate time of an electronic counter.

What happens in many cases is that we sample 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 accumulated count of cycles so that it can be read out, printed, or stored in some other way. Then there is a delay time for such processing of the data before the counter arms or initializes and resumes on the next cycle of the oscillation. During the delay time (or process time), information is lost. This is called "dead-time", and in some instances it becomes a problem. Unfortunately for data processing in typical oscillators the effects of deadtime often hurt most when it is the hardest to avoid. In other words, for times that are short compared to a second, when it is very difficult to avoid deadtime, this is usually where deadtime can make a significant difference in the data analysis. Typically, for many oscillators, if the sample time is long compared to a second, the dead-time makes little difference in the data analysis, unless it is excessive [1]. New equipment or techniques are now available that contribute zero or negligible deadtime [2].

In reality, of course, the sinusoidal output of an oscillator is not pure; it contains noise (frequency) fluctuations as well. We will describe three different methods of measuring the frequency fluctuations in precision oscillators other than measuring the frequency directly with a frequency counter, listed as a fourth method. The direct frequency counter technique is often very limiting because the number of resolvable digits on the counter are often inadequate for precision oscillators, and counter input noise masks oscillator noise for short sample times. In all the methods one also needs to properly match the impedances of different connected electronic instruments, use short connecting cable lengths, and use high-quality, stable connectors.

1.1. Common Methods of Measuring Frequency Stability

1.1.1. Beat-Frequency Method. The first technique is called a heterodyne frequency-measuring method or beat-frequency method. The signals from two independent oscillators are fed into the two ports of a double balanced mixer, as illustrated in Fig. 4. The device labeled "Amp" is an amplifier.

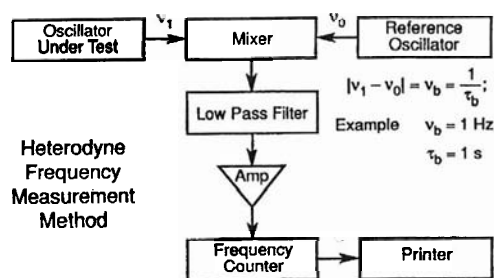


Figure 4. Measurement of the frequency difference ("beat note") between oscillators can increase measurement precision. State-of-the-art oscillators can readily be measured by this method.

The difference frequency, or the beat frequency v_b , is obtained as the output of a lowpass filter (to suppress carrier frequency harmonics) that follows the mixer. This beat frequency is then amplified and fed to a frequency counter and printer or other recording device. The fractional frequency is obtained by dividing v_b by the nominal carrier frequency v_0 . This system has excellent precision; one can measure essentially all state-of-the-art oscillators.

1.1.2. Dual-Mixer Time-Difference (DMTD) System. This technique uses two heterodyne measurements operating simultaneously. The time difference of the zero crossings of each beat frequency is measured and yields an excellent precision, 10^{-13} seconds. A block diagram is shown in Fig. 5. It should be mentioned that if time or time fluctuations can be measured directly, an advantage is obtained over just measuring frequency. The reason is that we can readily calculate the frequency from the time, only if there is no deadtime. In the past, frequency was not inferred from the time (for sample times of the order of several

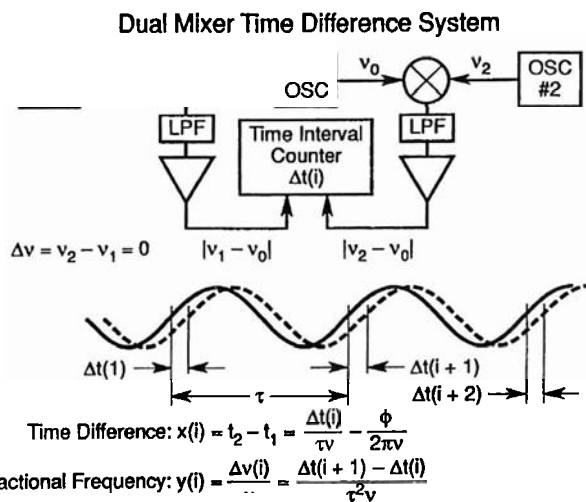


Figure 5. Measurement of the time difference between two beat notes from two oscillators with a common transfer oscillation can further increase measurement precision. Instability of transfer oscillator cancels to first order.

seconds and less) because the time difference between a pair of oscillators operating as clocks could not be measured with sufficient precision. However, now the precision of DMTD opens the door to measuring time fluctuations as well as frequency fluctuations for sample times as short as a few milliseconds, all without deadtime.

In Fig. 5, oscillator 1 could be considered to be under test and oscillator 2 could be considered to be the reference oscillator. Their outputs go to the ports of a pair of double-balanced mixers. Another oscillator with separate symmetric buffered outputs is fed to the other two ports of the pair of double-balanced mixers. This common oscillator's frequency is offset by a desired amount from those of the other two oscillators. Then two different beat frequencies are produced by the two mixers as shown. These two beat frequencies will be out of phase by an amount proportional to the time difference between oscillators 1 and 2—excluding the differential phase shift that may be inserted (component “4” is a phase shifter). Further, the beat frequencies differ in frequency by an amount equal to the frequency difference between oscillators 1 and 2.

This measurement technique is very useful where oscillators 1 and 2 outputs are at very nearly the same frequency. This is typical for atomic standards (cesium, rubidium, and hydrogen frequency standards).

Illustrated at the bottom of Fig. 5 is what might represent the beat frequencies from the two mixers. A phase shifter may be inserted as component “4” 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 (for certain types of noise) when the time difference is determined. After these beat signals are amplified, the start port of a time interval counter is triggered with the positive zero crossing of the other beat. Taking the time difference between the zero crossings of these beat frequencies, we measure the time difference between zero crossings of oscillators 1 and 2, but with a precision that 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

$$x(i) = \frac{\Delta t(i)}{\tau_b \nu_0} - \frac{\phi}{2\pi \nu_0} \frac{k}{\nu_0} \tag{5}$$

where $\Delta t(i)$ is the i th time difference as read on the counter, τ_b is the beat period, ν_0 is the nominal carrier frequency, ϕ is the phase delay in radians added to the signal of oscillator 1, and k is an integer number of cycles of ν_0 to be determined in order to remove the cycle ambiguity. It is important to know k only if the absolute time difference is desired; for measurements of frequency and of time fluctuations, k may be assumed zero unless we go through a cycle during a data run. The fractional frequency $y(i, \tau)$ between oscillators 1 and 2 can be derived in the normal

way from the time fluctuations:

$$y_{1,2}(i, \tau) = \left\{ \begin{array}{l} \frac{\nu_1(i, \tau) - \nu_2(i, \tau)}{\nu_0} \\ \frac{x(i+1) - x(i)}{\tau} \\ \frac{\Delta t(i+1) - \Delta t(i)}{\tau_b^2 \nu_0} \end{array} \right. \tag{6}$$

In Eqs (5) and (6), it is assumed that the transfer (or common) oscillator is set at a frequency lower than those of oscillators 1 and 2, and that the voltage zero crossing of the beat frequency $\nu_1 - \nu_c$ starts—and that $\nu_2 - \nu_c$ stops—the time interval counter. The fractional frequency difference may be averaged over any integer multiple of τ_b ,

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

where m is any positive integer. If needed, τ_b can be made to be very small by having very high beat frequencies. The transfer (or common) oscillator may be replaced with a low phase noise frequency synthesizer, which derives its basic reference frequency from oscillator 2. In this setup the nominal beat frequencies are given simply by the amount by which the output frequency of the synthesizer is offset from ν_2 . Sample times as short as a few milliseconds with subpicosecond (< 1 ps) resolution are obtained. Note that logging the data at such a rate usually requires special equipment. The National Institute of Standards and Technology (NIST) timescale measurement system is based on the DMTD.

1.1.3. Loose Phase-Locked Loop Method. This type of method is illustrated in Fig. 6. 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° out of phase, so that the average voltage out of the new mixer is nominally zero, and the instantaneous voltage fluctuations correspond to phase fluctuations rather than to 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

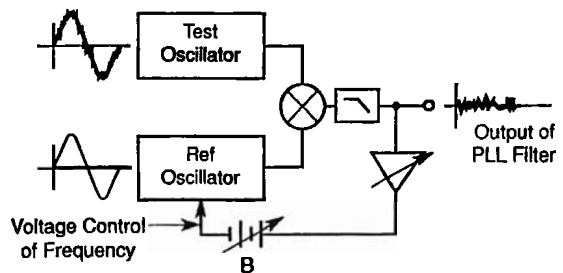


Figure 6. Direct measurement of the phase difference between two oscillators yields excellent precision. The technique requires electronic frequency control of a clean reference oscillator to maintain a loose phase lock, hence a zero beat.

making low-noise precision stability measurements. The output of this mixer is fed through a lowpass filter and then amplified in a feedback loop, causing the voltage-controlled oscillator (reference) to be phase-locked to the test oscillator. The response time of the loop is adjusted such that a very loose phase-lock (long-time-constant) condition exists.

The response (or attack) time is the time it takes the servo system to make 70% of its ultimate correction after being slightly disturbed. The response time is equal to $1/\pi w_h$, where w_h is the servo bandwidth. If the response time of the loop is about a second, then the voltage fluctuations will be proportional to the phase fluctuations for sample times shorter than one second. Depending on the coefficient of the tuning capacitor and the quality of the oscillators involved, the amplification used may vary significantly, but may typically range from 40 to 80 dB via a good low-noise amplifier. In turn this signal can be fed to a spectrum analyzer to measure the Fourier components of the phase fluctuations. It is of particular use for sample times shorter than one second (for Fourier frequencies greater than 1 Hz) in analyzing the characteristics of an oscillator. It is particularly useful if one has discrete sidebands such as 60 Hz, or detailed structure in the spectrum.

One may also take the output voltage from the above-mentioned amplifier and feed it to an analog-to-digital (A/D) converter. This digital output becomes an extremely sensitive measure of the short-term time or phase fluctuations between the two oscillators. Resolutions of the order of a picosecond (ps) are easily achievable.

1.1.4. Time-Difference Method Using a Counter. The last measurement method we will illustrate is very commonly used, but typically does not have the measurement precision that is more readily available in the first three methods illustrated above. This method, called the *time-difference method*, is shown in Fig. 7. Because of the wide bandwidth needed to measure fast-risetime pulses, this method is limited in signal-to-noise ratio. However, some commercially available counters allow us to do sig-

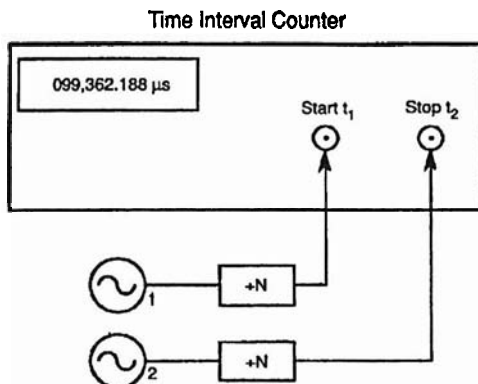


Figure 7. Measurement of the time difference between two oscillators, usually after division by N to obtain 1 pulse-per-second, yields only moderate measurement performance compared to previous methods. The technique is dependent on several properties of the counter and its trigger circuits.

nal averaging or precision risetime comparisons (precision of time-difference measurements in the range of 10 ns–10 ps are now available). Such a method yields a direct measurement of $x(t)$ without any translation, conversion, or multiplication factors. However, even if adequate measurement resolution is available, caution should be exercised in using this technique because it is not uncommon to have significant instabilities in the frequency dividers shown in Fig. 7—of the order of 100 ps. The technique is therefore suitable for long, not short, averaging times.

A trick to bypass divider problems is to feed the oscillator signals directly into the time interval counter and observe the zero-voltage crossing. The divided signal can be used to resolve cycle ambiguity of the carrier; otherwise the carrier phase at zero volts may be used as the time reference. The slope of the signal at zero volts is $2\pi V_0/\tau_1$, where $\tau_1 = 1/\nu_1$ (period of oscillation). For $V_0 = 1$ V and a 5-MHz signal, this slope is 3 mV/ns, which is a good sensitivity. (*Caution:* A correct impedance match of less than 1.5 VSWR is critical for this setup to be stable.)

2. CHARACTERIZATION

Given a set of data for 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 random (may be predicted only statistically) or nonrandom (i.e., systematic—environmentally induced or that have a causal effect that can be determined and in many cases be predicted).

2.1. Nonrandom Fluctuations

Nonrandom fluctuations are usually the main cause of departure from “true” time or “true” frequency in the long term.

If, for example, we have values of the frequency over a period of time, and a frequency offset from nominal is observed, one may calculate directly that the phase error will accumulate as a ramp. If, on the other hand, the frequency values drift linearly, then the time fluctuations will behave quadratically. In almost all oscillators, these “systematics,” as they are sometimes called, are the primary cause of time and/or frequency departure. A useful approach for determining 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 least-squares quadratic fit to the phase or to the time deviations is rarely as efficient an estimator of the frequency drift for most oscillators. Precision frequency standards are affected by their environment. These environmental perturbations often cause long-term departures of frequency and time, which in a data run can look like drift, but are not.

2.2. Random Fluctuations

After the systematic or nonrandom effects of a dataset have been calculated or estimated, they may be subtracted

from the data, leaving the residual random fluctuations. They can usually be best characterized statistically using $\sigma_y(\tau)$, the Allan deviation, for short τ (values) and “Theol” for long τ , the agreed-on standards (IEEE) in the time domain, to be explained in the next section. It is often the case for precision oscillators that these random fluctuations may be effectively modeled with power-law spectral densities. This topic and measurements of spectrum are discussed later. We have

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

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 coefficient indicating the level of that type of noise, and α is a number modeling the most appropriate type of power law for the data. If we observe from a $\log \sigma_y^2(\tau)/\log \tau$ diagram a particular slope (call it μ) over certain regions of sample time, τ , this slope has a correspondence to a power-law spectral density or a set of the same with some amplitude coefficient h_α . In particular, $\mu = -(\alpha + 1)$ for $-3 < \alpha < 1$ and $\mu \cong -2$ for $\alpha \geq 1$. Further, a correspondence exists between h_α and the coefficient for $\sigma_y(\tau)$ [1]. The transformations for some of the more common power-law spectral densities have been tabulated, making it quite easy to transform the frequency stability modeled in the time domain over to the frequency domain and vice versa. Examples of some power-law spectra and other types of noise that have been simulated by computer are shown in Fig. 8. The root Allan variance (an RMS or deviation called “Adev”) and Theol-deviation are constructed to extract frequency instability and not measurement system noise. Synchronization and measurement system noise is phase or time instability characterized by other statistics such as time deviation (TDEV) and maximum time interval error (MTIE) [3].

Once the noise characteristics have been determined, one is often able to deduce whether the oscillators are performing properly, and whether they are meeting either the design specifications or the manufacturer’s specifications. For example, a cesium beam frequency standard or a rubidium gas cell frequency standard, when working

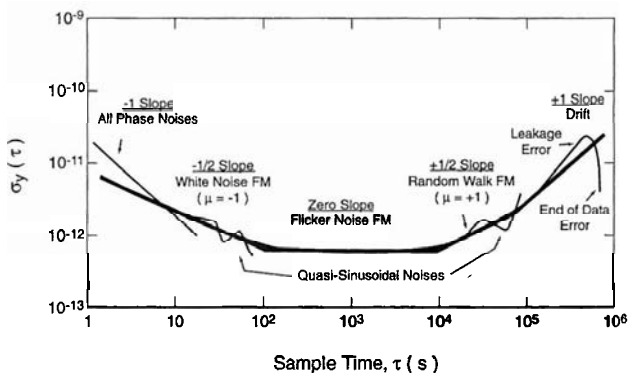


Figure 8. “Adev” (root Allan variance estimate) showing power-law noise as straight lines in addition to other errors. Our goal is to properly interpret this kind of plot of frequency stability.

properly, should exhibit white frequency noise (slope of $-\frac{1}{2}$) for values of τ of the order of a few seconds to several thousand seconds.

3. ANALYSIS OF TIME DOMAIN DATA

Suppose now that we are given the time or frequency fluctuations between a pair of precision oscillators measured, for example, by one of the techniques outlined in Section 1, and a stability analysis is desired. Let this comparison be depicted by Fig. 9. The minimum sample time is determined by the measurement system. If the time difference or time fluctuations are available, then the frequency or the fractional frequency fluctuations may be calculated from one period of sampling to the next. 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. 9, $\sigma_{\text{std.dev.}}(\tau)$, where y_i is the average fractional frequency over the dataset, 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. We have studied what happens to the standard deviation when the dataset may be characterized by power-law spectra that are more dispersive than classical white-noise frequency fluctuations. In other words, if the fluctuations are characterized by flicker noise or any other non-white-noise frequency deviations, what happens to the standard deviation for that dataset? We can show that the standard deviation is a function of the number of data points in the set (discussed next), and it is also a function of the dead-time and of the measurement system bandwidth. For example, using flicker-noise frequency modulation as a model, as the number of data points increases, the standard deviation increases monotonically without limit. Some statistical measures have been developed that do not depend on the datalength and that are readily usable for characterizing the random fluctuations in precision oscillators. The IEEE has adopted a standard measure

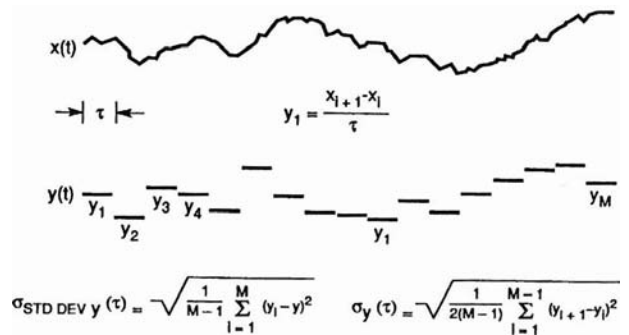


Figure 9. A simulated plot of the time fluctuations $x(t)$ between a pair of oscillators and of the corresponding fractional frequencies calculated from the time fluctuations each averaged over a sample time τ . At the bottom are the equations for the standard deviation (left) and for the time-domain measure of frequency stability as recommended by the IEEE (right).

known as the “Allan variance” taken from the set of useful variances developed, and an experimental estimation of the square root of the Allan variance is shown as the bottom right equation in Fig. 9 [2,4]. This equation is very easy to implement experimentally, as we need to simply add up the squares of the differences between adjacent values of y_i , divide by the number of them and by 2, and take the square root. We then have the quantity that the IEEE subcommittee has recommended for specification of stability in the time domain, denoted by $\sigma_y(\tau)$

$$\sigma_y(\tau) = \left\langle \frac{1}{2} (\bar{y}(t+\tau) - \bar{y}(t))^2 \right\rangle^{1/2} \tag{9}$$

where the brackets “ $\langle \rangle$ ” denote infinite time average. In practice this is easily estimated from a finite dataset as follows:

$$\sigma_y(\tau) = \left[\frac{1}{2(M-1)} \sum_{i=1}^{M-1} (y_{i+1} - y_i)^2 \right]^{1/2} \tag{10}$$

where the y_i are the discrete frequency averages as illustrated in Fig. 9.

We would like to know how $\sigma_y(\tau)$ varies with the sample time τ . A simple and very useful trick that we can use if there is no deadtime is to average the values for y_1 and y_2 and call that a new y_1 averaged over 2τ ; similarly average the values for y_3 and y_4 and call that a new y_2 averaged over 2τ , and so on, and finally apply the same equation as before to get $\sigma_y(2\tau)$. One can repeat this process for other desired integer multiples m of τ , and from the same dataset 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 we have deadtime in the measurements, adjacent pairs cannot be averaged in an unambiguous way to simply increase the sample time. We have to retake the data for each new sample time—often a very time-consuming task. This is another instance where deadtime can be a problem.

The classical variance (standard deviation squared) is the wrong statistic for measurements of frequency stability, because in most cases it depends on the number of data samples. Fig. 10 plots the ratio of the standard deviation squared for N samples to the standard deviation squared for two samples, $\langle \sigma^2(2, \tau) \rangle$, which is the same as the Allan variance, $\sigma_y^2(\tau)$. We can see the dependence of this standard deviation on the number of samples for various kinds of power-law spectral densities commonly encountered as reasonable models for many important precision oscillators. Note that $\sigma_y^2(\tau)$ has the same value as the classical variance for the classical noise case (white-noise FM). Figure 10 shows that with the increasing length of data the standard deviation of the common classical variance is not well behaved.

We may combine Eqs. (4) and (9) to obtain an equation for $\sigma_y(\tau)$ in terms of the time-difference or time-deviation measurements:

$$\sigma_y(\tau) = \left\langle \frac{1}{2\tau^2} (-x(t+2\tau) + 2x(t+\tau) - x(t))^2 \right\rangle^{1/2} \tag{11}$$




Figure 10. 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. This figure illustrates one reason why the standard deviation is not a suitable measure of frequency stability [4].

which for N discrete time readings, also called N_x , may be estimated as

$$\sigma_y(\tau) \cong \left[\frac{1}{2(N-2)\tau^2} \sum_{i=1}^{N-2} (-x_{i+2} + 2x_{i+1} - x_i)^2 \right]^{1/2} \tag{12}$$

where the i integer denotes the number of the reading in the set of N and the nominal spacing between readings is τ . If there is no deadtime in the data and the original data were taken with a sample time τ_0 , a set of x_i values can be obtained by integrating the y_i values:

$$x_{i+1} = x_i + \tau_0 \sum_{j=1}^i y_j \tag{13}$$

Once we have the x_i values, we can pick τ in Eq (13) to be any integer multiple m of τ_0 , specifically $\tau = m\tau_0$:

$$\sigma_y(m\tau_0) \cong \left[\frac{1}{2(N-2m)m^2\tau_0^2} \sum_{i=1}^{N-2m} (-x_{i+2m} + 2x_{i+m} - x_i)^2 \right]^{1/2} \tag{14}$$

Equation (14), called the “max-overlap estimator,” is regarded as the best estimator of $\sigma_y(m\tau_0)$.

Example 1. Find the two-sample (Allan) variance, $\sigma_y^2(\tau)$, of the following sequence of fractional frequency fluctuation

values y_k , each value averaged over one second:

$$\begin{aligned} y_1 &= 4.36 \times 10^{-5} & y_5 &= 4.47 \times 10^{-5} \\ y_2 &= 4.61 \times 10^{-5} & y_6 &= 3.96 \times 10^{-5} \\ y_3 &= 3.19 \times 10^{-5} & y_7 &= 4.10 \times 10^{-5} \\ y_4 &= 4.21 \times 10^{-5} & y_8 &= 3.08 \times 10^{-5} \end{aligned}$$

(assume no deadtime in measurement of averages).

Since each average of the fractional frequency fluctuation values is for one second, then the first variance calculation will be at $\tau=1$ s. We are given $M=8$ (eight values); therefore, the number of pairs in sequence is $M-1=7$. We have

Data Values y_k ($\times 10^{-5}$)	First Differences ($y_{k+1} - y_k$) ($\times 10^{-5}$)	First Difference Squared ($(y_{k+1} - y_k)^2$) ($\times 10^{-10}$)
4.36	—	—
4.61	0.25	0.06
3.19	-1.42	2.02
4.21	1.02	1.04
4.47	0.26	0.07
3.96	-0.51	0.26
4.10	0.14	0.02
3.08	-1.02	1.04
		Sum = 4.51

$$\sum_{k=1}^{M-1} (y_{k+1} - y_k)^2 = 4.51 \times 10^{-10}$$

Therefore the Allan variance is

$$\sigma_y^2(1\text{ s}) = \frac{4.51 \times 10^{-10}}{2(7)} = 3.2 \times 10^{-11}$$

and the Allan deviation is

$$\sigma_y(1\text{ s}) = [\sigma_y^2(1\text{ s})]^{1/2} = [3.2 \times 10^{-11}]^{1/2} = 5.6 \times 10^{-6}$$

Using the same data, we can calculate the Allan variance for $\tau=2$ s by averaging pairs of adjacent values and using these new averages as data values for the same procedure as above. For three second averages ($\tau=3$ s), take adjacent threesomes and find their averages and proceed in a similar manner. More data must be acquired for use of longer averaging times.

The confidence of the estimate on $\sigma_y(\tau)$ improves nominally as the square root of the number of data values used. In this example $M=8$, and the confidence can be expressed as being no better than $1/\sqrt{8} \times 100\% = 35\%$. This is a one-sigma (1σ) uncertainty (68% confidence interval) in the estimate for the $\tau=1$ s average. We can dramatically improve confidence using a combination of signal processing, as discussed next.

For the particularly difficult measurement problem of determining the frequency stability of frequency

standards and oscillators for long averaging times, we can use the special-purpose statistic, the estimator of a theoretical variance 1 ("Theo1"), given in native form by [5]

$$\widehat{\text{Theo1}}(m, \tau_0, N_x) = \frac{1}{0.75(N_x - m)(m\tau_0)^2} \times \sum_{i=1}^{N_x - m} \sum_{\delta=0}^{(m/2) - 1} \frac{1}{(m/2 - \delta)} [(x_{i-x_i-\delta+\frac{m}{2}}) + (x_{i+m-x_i+\delta+\frac{m}{2}})]^2 \tag{15}$$

for m even, $10 \leq m \leq N_x - 1$, and $\tau = 0.75 m \tau_0$. It has statistical properties like those of the Allan variance, with the significant enhancement that it can evaluate frequency stability at longer averaging times than by using the Allan definition. We can remove bias relative to "Avar" by a composite statistic given by

$$\begin{aligned} &\text{TheoH}(m, \tau_0, N_x) \\ &= \begin{cases} \text{Avar}(m, \tau_0, N_x) & \text{for } 1 \leq m < \frac{k}{\tau_0} \\ \text{TheoBR}(m, \tau_0, N_x) & \text{for } \frac{k}{0.75\tau_0} \leq m \leq N_x - 1, m \text{ even} \end{cases} \end{aligned} \tag{16}$$

where k is the largest $\tau \leq T/10$ where $\text{Avar}(m, \tau_0, N_x)$ has sufficient confidence. In this equation TheoBR is defined

$$\begin{aligned} &\text{TheoBR}(m, \tau_0, N_x) \\ &= \left[\frac{1}{n+1} \sum_{i=0}^n \frac{\text{Avar}(m=9+3i, \tau_0, N_x)}{\widehat{\text{Theo1}}(m=12+4i, \tau_0, N_x)} \right] \widehat{\text{Theo1}}(m, \tau_0, N_x), \end{aligned} \tag{17}$$

where

$$n = \left\lfloor \frac{0.1N_x}{3} - 3 \right\rfloor$$

(where $\lfloor \cdot \rfloor$ means the integer part). Equation (16) computes a function that is Avar in short term and Theo1 in long term.

4. SPECTRUM ANALYSIS

Another method of characterizing the noise in a signal source is by means of spectrum analysis [6-8]. To understand this approach, let's examine the waveform shown in Fig. 11.

Here we have a sine wave that for short instances is perturbed by noise. Some workers loosely refer to these types of noises as "glitches." The waveform has a nominal frequency over one cycle that we'll call ν_0 ($\nu_0 = 1/T_0$). At times, noise causes the instantaneous frequency to differ markedly from the nominal frequency. If a pure sine-wave signal of frequency ν_0 is subtracted from this waveform,




Figure 11. Sine wave that is perturbed by periodic glitches.

the remainder is the sum of the noise components. These components are of various frequencies and the sum of their amplitudes is nearly zero except for the intervals during each glitch, when their amplitudes momentarily reinforce each other. This is shown graphically in Fig. 12.

We can construct a graph plotting RMS power against frequency for a given signal into a given load. This kind of plot is called the *power spectrum*. For the waveform of Fig. 11, the power spectrum will have a high value at ν_0 and lower values for the signals produced by the glitches. Closer analysis reveals that there is a recognizable, somewhat constant, repetition rate associated with the glitches.

In fact, we can deduce that there is a significant amount of power in another signal whose period is the period of the glitches as shown in Fig. 12. Let's call the frequency of the glitches ν_s . Since this is the case, we will observe a noticeable amount of power in the spectrum at ν_s with an amplitude that is related to the characteristics of the glitches. The power spectrum shown in Fig. 13 has this feature. A predominant ν_s component has been depicted, but other harmonics also exist.

Some noise will cause the instantaneous frequency to "jitter" around ν_0 , with a distribution that is higher and lower than ν_0 . We thus usually find a "pedestal" associated with ν_0 as shown in Fig. 14.

The process of breaking a signal down into all of its various components of frequency is called *Fourier expansion*. In other words, the addition of all the frequency components, called *Fourier frequency components*, produces the original signal. The value of a Fourier frequency is the difference between the frequency component and the fundamental frequency. The power spectrum can be normalized to unity such that the total area under the curve equals one. The power spectrum normalized in this way is the *power spectral density*.

The power spectrum of $V(t)$, often called the *RF spectrum*, is very useful in many applications. Unfortunately, if we are given the RF spectrum, it is impossible to determine whether the power at different Fourier frequencies is a result of amplitude fluctuations " $a(t)$ " or phase fluctua-




Figure 12. Periodic glitches are undesirable and can be separated from the desired sine wave and characterized in the frequency domain.




Figure 13. A plot of power (into a load) vs. frequency is a *power spectrum*.

tions " $\phi(t)$." The RF spectrum can be separated into two independent spectra, one of which is the *spectral density of $\phi(t)$* .

For the purpose here, the phase fluctuation components are the ones of interest. The spectral density of phase fluctuations is denoted by $S_\phi(f)$, where f is Fourier frequency. For the frequently encountered case where the AM power spectral density is negligibly small and the total modulation of the phase fluctuations is small (mean-square value is much less than 1 rad^2), the RF spectrum has approximately the same shape as the phase spectral density.

However, a main difference in the representation is that the RF spectrum includes the fundamental signal (carrier), and the phase spectral density does not. Another major difference is that the RF spectrum is a *power spectral density* and is measured in units of watts/hertz. The phase spectral density involves no "power" measurement of the electrical signal. The units are radians²/hertz. It is tempting to think of $S_\phi(f)$ as a "power" spectral density because in practice it is measured by passing $V(t)$ through a phase detector and measuring the detector's output power spectrum. The measurement technique makes use of the relation that for small deviations ($\delta\phi \leq 1$ radian)

$$S_\phi(f) = \left[\frac{V_{\text{RMS}}(f)}{V_s(f)} \right]^2 \tag{18}$$

where $V_{\text{RMS}}(f)$ is the root-mean-square noise voltage in a 1 Hz bandwidth (i.e., per $\sqrt{\text{Hz}}$) at a Fourier frequency f , and $V_s(f)$ is the sensitivity (volts per radian) at the phase quadrature output of a phase detector that is comparing the two oscillators. In the next section, we will look at a scheme for directly measuring $S_\phi(f)$ by determining $V_s(f)$.

SPECTRAL DENSITY




Figure 14. The power spectrum of an oscillator includes its "noise pedestal."

One question we might ask is “How do frequency changes relate to phase fluctuations?” After all, it’s the frequency stability of an oscillator that is a major consideration in many applications. The frequency is equal to a rate of change in the phase of a sine wave. This tells us that fluctuations in an oscillator’s output frequency are related to phase fluctuations since we must change the rate of $\phi(t)$ to accomplish a shift in $\nu(t)$, the frequency at time t . A rate of change of total $\phi_T(t)$ is denoted by $\dot{\phi}_T(t)$. We then have

$$2\pi\nu(t) = \dot{\phi}_T(t) \tag{19}$$

The dot denotes the mathematical operation of *differentiation* on the function ϕ_T with respect to its independent variable t .² From Eqs. (19) and (1) we get

$$2\pi\nu(t) = \dot{\phi}_T(t) = 2\pi\nu_0 + \dot{\phi}(t)$$

Rearranging, we have

$$2\pi\nu(t) - 2\pi\nu_0 = \dot{\phi}(t)$$

$$\nu(t) - \nu_0 = \frac{\dot{\phi}(t)}{2\pi} \tag{20}$$

The quantity $\nu(t) - \nu_0$ can be more conveniently denoted as $\delta\nu(t)$, a change in frequency at time t . Equation (20) tells us that if we differentiate the phase fluctuations $\phi(t)$ and divide by 2π , we will have calculated the frequency fluctuation $\delta\nu(t)$. Rather than specifying a frequency fluctuation in terms of shift in frequency, it is useful to denote $\delta\nu(t)$ with respect to the nominal frequency ν_0 . The quantity $\delta\nu(t)/\nu_0$ is called the *fractional frequency fluctuation*³ at time t and is signified by the variable $y(t)$. We then have

$$y(t) = \frac{\delta\nu(t)}{\nu_0} = \frac{\dot{\phi}(t)}{2\pi\nu_0} \tag{21}$$

The fractional frequency fluctuation $y(t)$ is a dimensionless quantity. When talking about frequency stability, its appropriateness becomes clearer if we consider the following example. Suppose that in two oscillators $\delta\nu(t)$ is consistently equal to +1 Hz and we have sampled this value for many times t . Are the two oscillators equal in their ability to produce their desired output frequencies? Not if one oscillator is operating at 10 Hz and the other at 10 MHz. In one case, the average value of the fractional frequency fluctuation is 1 in 10, and in the second is 1 in 10,000,000 or 1×10^{-7} . The 10 MHz oscillator is then more accurate. If frequencies are multiplied or divided using ideal electronics, the fractional stability is not changed.

In the frequency domain, we can measure the spectrum of frequency fluctuations $y(t)$. The spectral density of fre-

quency fluctuations is denoted by $S_y(f)$ and is obtained by passing the signal from an oscillator through an ideal FM detector and performing spectral analysis on the resultant output voltage. $S_y(f)$ has dimensions of (fractional frequency)²/Hz or Hz^{-1} . Differentiation of $\phi(t)$ corresponds to multiplication by f/ν_0 in terms of spectral densities. With further calculation, one can deduce that

$$S_y(f) = \left(\frac{f}{\nu_0}\right)^2 S_\phi(f) \tag{22}$$

We will address primarily $S_\phi(f)$, that is, the spectral density of phase fluctuations. For the purpose of noise measurements, $S_\phi(f)$ can be measured with a straightforward, easily duplicated equipment setup. Whether one measures phase or frequency spectral densities is of minor importance since they bear a direct relationship. It is important, however, to make the distinction and to use Eq. (22) when necessary.

4.1. The Loose Phase-Locked Loop

In Section 1.1.3 we described a method of measuring phase fluctuations between two phase-locked oscillators. Now we will review a common procedure for measuring $S_\phi(f)$.

Suppose that we have a noisy oscillator. We wish to measure the oscillator’s phase fluctuations relative to nominal phase. One can do this by phase-locking another oscillator (called the *reference oscillator*) to the test oscillator, and mixing the two oscillator signals 90° out of phase (phase quadrature). This is shown schematically in Fig. 15. The two oscillators are at the same frequency in the long term, as guaranteed by the phase-locked loop (PLL). A lowpass filter (to filter the RF sum component) is used after the mixer since the difference (baseband) signal is the one of interest. By holding the two signals at a relative phase difference of 90°, short-term phase fluctuations between the test and reference oscillators will appear as voltage fluctuations from the mixer.

With a PLL, if we can make the servo time constant very long, then the PLL bandwidth as a filter will be small. This may be done by lowering the gain A_v of the loop amplifier. We want to translate the phase modulation

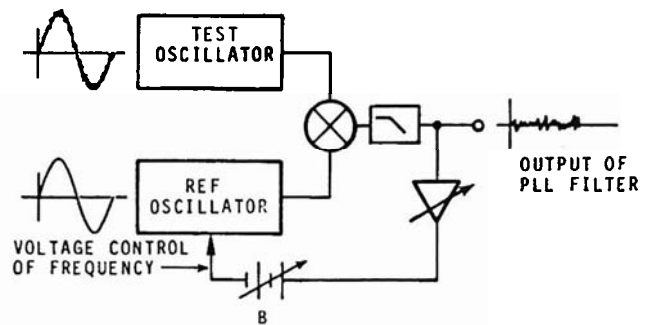


Figure 15. The phase noise of a test oscillator is usually measured by a loose phase-locked loop. The test and reference oscillators will naturally lock so that their signals have a phase difference of 90 deg. and the PLL output voltage fluctuations correspond to phase fluctuations between the oscillators.

²As an analogy, the same operation relates the velocity of an object to its acceleration.

³Some international recommendations replace “fractional” by “normalized.”

spectrum to baseband spectrum so that it is easily measured on a low-frequency spectrum analyzer. With a PLL filter, we must keep in mind that the reference oscillator should be as good as or better than the test oscillator. This is because the output of the PLL represents the noise from both oscillators, and if not properly chosen, the reference can have noise masking the noise from the test oscillator. Often, the reference and test oscillators are of the same type and have, therefore, approximately the same noise levels. We can acquire a meaningful measurement by noting that the noise we measure is from two oscillators. Many times a good approximation is to assume that the measured noise power is twice that associated with either single oscillator. $S_{\phi}(f)$ is general notation depicting spectral density on a reciprocal hertz (Hz^{-1}) basis. The output from PLL filter necessarily yields noise from two oscillators.

The output of the PLL filter at Fourier frequencies above the loop bandwidth is a voltage representing phase fluctuations between reference and test oscillator. It is necessary to make the time constant of the loop long compared to the inverse of the lowest Fourier frequency that we wish to measure, that is, $\tau_c > [1/2\pi f(\text{lowest})]$. This means that if we want to measure $S_{\phi}(f)$ down to 1 Hz, the loop time constant must be greater than $1/2\pi$ seconds. We can measure the time constant by perturbing the loop (momentarily disconnecting the battery is convenient) and noting the time it takes for the control voltage to reach 70% of its final value. The signal from the mixer can then be inserted into a spectrum analyzer. A preamplifier may be necessary in the signal path into the spectrum analyzer.

The analyzer determines the mean-square voltage that passes through the analyzer's bandwidth centered around a prechosen Fourier frequency f . It is desirable to normalize results to a 1 Hz bandwidth. Assuming white phase noise (white PM), this can be done by dividing the mean-square voltage by the analyzer bandwidth in hertz. We may have to approximate for other noise processes. [The phase noise sideband levels will usually be indicated in RMS volts per root hertz ($\text{V}/\sqrt{\text{Hz}}$) on most analyzers.]

4.2. Equipment for Frequency-Domain Stability Measurements

4.2.1. Low-Noise Mixer. This should be a high quality, double-balanced type as shown in Fig. 16 but single-ended types may be used. The oscillators should have well-buffered outputs to be able to isolate the coupling between the two input RF ports of the mixer. Results that are too good




Figure 16. A low noise mixer is a key component for precise phase-noise measurements.

may be obtained if the two oscillators couple tightly via signal injection through the input ports. We want the PLL to control locking. One should read the specifications in order to prevent exceeding the maximum allowable input power to the mixer. However, it is best to operate near the maximum for best signal-to-noise ratio out of the IF port of the mixer, and, in some cases, it is possible to drive the mixer into saturation without burning out the device.

4.2.2. Low-Noise DC Amplifier. The amount of gain A_v needed in the loop amplifier will depend on the amplitude of the mixer output and the degree of varactor control in the reference oscillator. We may need only a small amount of gain to acquire lock. On the other hand, it may be necessary to add as much as 80 dB of gain. Good low-noise DC amplifiers are available from a number of sources, and with cascading stages of amplification, each contributing noise, it will be the noise of the first stage that will add most significantly to the noise being measured. Amplifiers with very low equivalent input noise performance are available from many manufacturers. The response of the amplifier should be flat from DC to the highest Fourier frequency one wishes to measure. The loop time constant is inversely related to the gain A_v , and A_v is best determined experimentally by sweeping the system with known modulation applied at the output of one oscillator [9].

4.2.3. Voltage-Controlled Reference Quartz Oscillator. This oscillator should be a good one with specifications available on its frequency-domain stability. The reference must be no worse than the test oscillator in the frequency domain. The varactor control should be sufficient to maintain phase lock of the reference. In general, test oscillators of moderate quality may have varactor control of as much as 1×10^{-6} fractional frequency change per volt. Some provision should be available on the reference oscillator for tuning the mean frequency over a frequency range that will enable phase lock. Many factors enter into the choice of the reference oscillator, and often it is convenient to simply use two test oscillators phase-locked together. In this way, we can assume that the noise out of the PLL filter is no worse than 3 dB greater than the noise from each oscillator. If it is uncertain whether both oscillators contribute approximately equal noise, then we should perform measurements on three oscillators, taking two at a time. The noisier-than-average oscillator will reveal itself.

4.2.4. Spectrum Analyzer. The signal analyzer should typically be capable of measuring the noise in RMS volts in a narrow bandwidth from near 1 Hz to the highest Fourier frequency of interest. This may be 50 kHz for carrier frequencies of 10 MHz or lower and several megahertz for microwave carrier frequencies. For voltage measuring analyzers, it is typical to use units of "volts per $\sqrt{\text{Hz}}$ ". The spectrum analyzer and any associated input amplifier will exhibit high-frequency rolloff. The Fourier frequency at which the voltage has dropped by 3 dB is the measurement system bandwidth f_h , or $\omega_h = 2\pi f_h$. This can be measured directly with a variable signal generator.

A frequency-domain measurement setup is shown schematically in Fig. 17. The component values for the lowpass filter out of the mixer are suitable for oscillators operating at around 5 MHz.

The active gain element (A_v) of the loop is a DC amplifier, hopefully with flat frequency response, or, if not, a known frequency response. One may replace this element by an integrator to achieve high gain near DC and hence, maintain better lock of the reference oscillator in long term. Otherwise long-term drift between the reference and test oscillators might require manual re-adjustment of the frequency of either oscillator [1,6].

Rather than measure the spectral density of phase fluctuations between two oscillators, it is possible to measure the phase fluctuations introduced by a device such as an active filter or amplifier. Only a slight modification of the existing PLL filter equipment setup is needed. The required scheme is shown in Fig. 18.

Figure 18 is a *differential phase noise measurement* setup. The output of the reference oscillator is split so that part of the signal passes through the device under test. We want the two signals going to the mixer to be 90° out of phase; thus, phase fluctuations between the two input ports cause voltage fluctuations at the output. The voltage fluctuations then can be measured at various Fourier frequencies on a spectrum analyzer.

To estimate the noise inherent in the test setup, one can in principle bypass the device under test and compensate for any change in amplitude and phase at the mixer. In order to measure inherent test equipment noise, the PLL filter technique must be converted to a differential phase noise technique. It is good practice to measure the system noise before proceeding to measurement of device noise.

4.3. Procedure and Example

For all of these setups, at the input to the spectrum analyzer, the voltage varies as the phase fluctuations in short term. The conversion to spectral density is

$$S_{\phi}(f) = \left(\frac{V_{RMS}(f)}{V_s(f)} \right)^2,$$

and $\mathcal{L}(f) = \frac{1}{2} S_{\phi}(f)$




Figure 17. Typical hookup for a phase-noise measurement using a loose PLL.

where V_s is the phase sensitivity of the mixer in volts per radian at offset frequency f . Using the setup described previously, V_s can be measured by disconnecting the feedback loop to the varactor of the reference oscillator. The peak voltage swing is equal to V_s in units of V/rad (volts per radian) if the resultant beat note is a sine wave at frequency f . This may not be the case for state-of-the-art $S_{\phi}(f)$ measurements, where one must drive the mixer very hard to achieve low mixer noise levels. Hence the output will not be a sine wave, and the V/rad sensitivity must be estimated by the slew rate (through zero volts) of the resultant square wave from the mixer/amplifier.

The value for the measured $S_{\phi}(f)$ in decibels is given by

$$S_{\phi}(f) = 20 \log \frac{V_{RMS}(f)}{V_s \text{ full-scale } \phi - \text{detector voltage at } f}$$

Example 2. Given a PLL with two oscillators such that at the mixer output: $V_s = 1 \text{ V/rad}$ with a beat frequency $f = 45 \text{ Hz}$, $V_{RMS}(45 \text{ Hz}) = 100 \text{ nV per root hertz}$. Solve for $S_{\phi}(45 \text{ Hz})$:

$$S_{\phi}(45 \text{ Hz}) = \left(\frac{100 \text{ nV/Hz}^{-1/2}}{1 \text{ V/rad}} \right)^2 = \left(\frac{10^{-7}}{1} \right)^2 \text{ rad}^2/\text{Hz}$$

$$= 10^{-14} \frac{\text{rad}^2}{\text{Hz}},$$

$$\mathcal{L}(45 \text{ Hz}) = 5 \times 10^{-15} \frac{\text{rad}^2}{\text{Hz}}$$

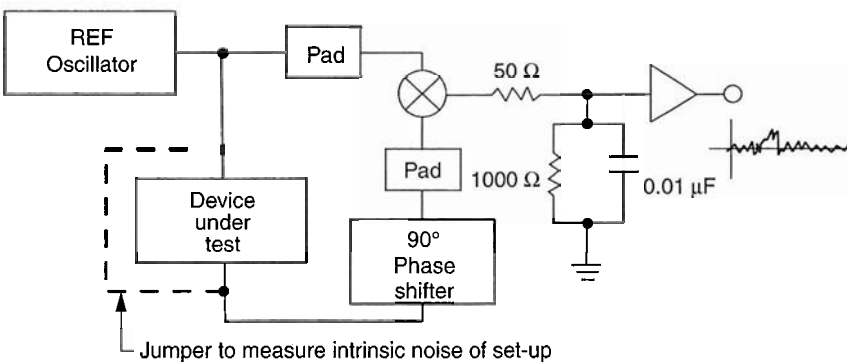


Figure 18. A differential measurement and measurement of the noise floor (by replacing the device under test with a through-cable) requires a 90 deg. phase shift at one of the mixer inputs.

In decibels

$$S_{\phi}(45 \text{ Hz}) = 20 \log \frac{100 \text{ nV}}{1 \text{ V}} = 20 \log \frac{\sim}{10^0}$$

$$= 20(-7) = -140 \text{ dB at } 45 \text{ Hz,}$$

$$\mathcal{L}(45 \text{ Hz}) = -143 \text{ dB at } 45 \text{ Hz}$$

In the example, note that the mean frequency of the oscillators in the PLL was not essential to computing $S_{\phi}(f)$. However, in the application of $S_{\phi}(f)$, the mean frequency ν_0 is necessary information. Along with $S_{\phi}(f)$, one should always refer to ν_0 . In the example above, where $\nu_0 = 5 \text{ MHz}$, we have

$$S_{\phi}(45 \text{ Hz}) = 10^{-14} \frac{\text{rad}^2}{\text{Hz}}, \nu_0 = 5 \text{ MHz}$$

From Eq. (22), $S_y(f)$ can be computed as

$$S_y(45 \text{ Hz}) = \left(\frac{45}{5 \times 10^6} \right)^2 10^{-14} \frac{\text{rad}^2}{\text{Hz}}$$

$$= 5.1 \times 10^{-25} \text{ Hz}^{-1}, \nu_0 = 5 \text{ MHz.}$$

5. POWER-LAW NOISE PROCESSES

Power-law noise processes are models of precision oscillator noise that produce a particular *slope* on a spectral density plot. We often classify these noise processes into one of five categories. For plots of $S_{\phi}(f)$, they are

1. Random-walk FM (random walk of frequency), S_{ϕ} plot goes down as $1/f^4$.
2. Flicker FM (flicker of frequency), S_{ϕ} plot goes down as $1/f^3$.
3. White FM (white of frequency), S_{ϕ} plot goes down as $1/f^2$.
4. Flicker PM (flicker of phase), S_{ϕ} plot goes down as $1/f$.
5. White PM (white of phase), S_{ϕ} plot is flat.

Power-law noise processes are characterized by their functional dependence on Fourier frequency. Equation (22) relates $S_{\phi}(f)$ to $S_y(f)$, the spectral density of frequency fluctuations.

The spectral density plot of a typical oscillator's output is usually a combination of different power-law noise processes. It is very useful and meaningful to categorize the noise processes. The first job in evaluating a spectral density plot is to determine which type of noise exists for a particular range of Fourier frequencies. It is possible to have all five noise processes generated from a single oscillator, but in general only two or three noise processes are dominant. Figure 19 is a graph of $S_{\phi}(f)$ showing the five noise processes on a log-log scale. Figure 20 shows the

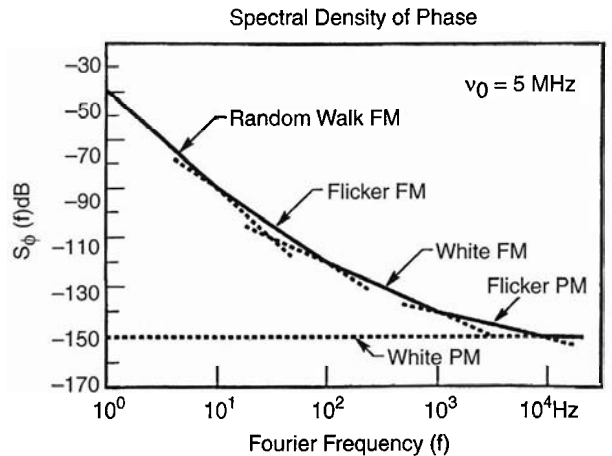


Figure 19. Power-law noise is indicated by a particular slope in the phase-noise measurement.

spectral density of phase fluctuations for a typical 5-MHz quartz oscillator.

6. CAUSES OF NOISE TYPES IN A SIGNAL SOURCE

6.1. Power-Law Noise Processes

Section 5 pointed out the five commonly used power-law models of noise. With respect to $S_{\phi}(f)$, one can estimate a straight-line slope (on a log-log scale) that corresponds to a particular noise type. This is shown in Fig. 19.

We can make the following general remarks about power-law noise processes:

1. Random-walk FM ($1/f^4$) noise is difficult to measure since it is usually very close to the carrier. "Random-walk FM" usually relates to the oscillator's physical environment. If random walk FM is a predominant feature of the spectral density plot, then mechanical shock, vibration, temperature, or other environmental effects may be causing "random" shifts in the carrier frequency.
2. Flicker FM ($1/f^3$) is a noise whose physical cause is seldom fully understood but may typically be related

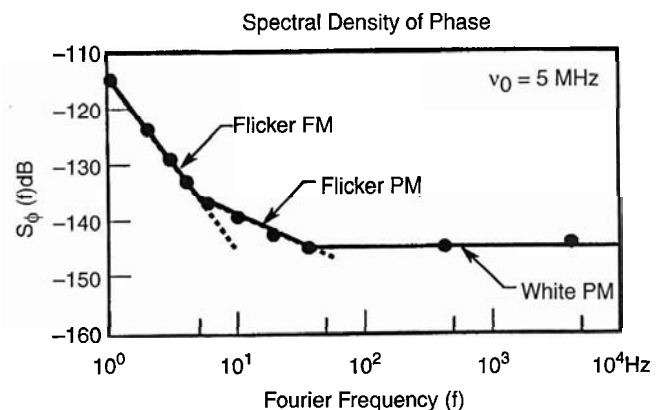


Figure 20. Different power-law noises have different causes in an oscillator's output signal.

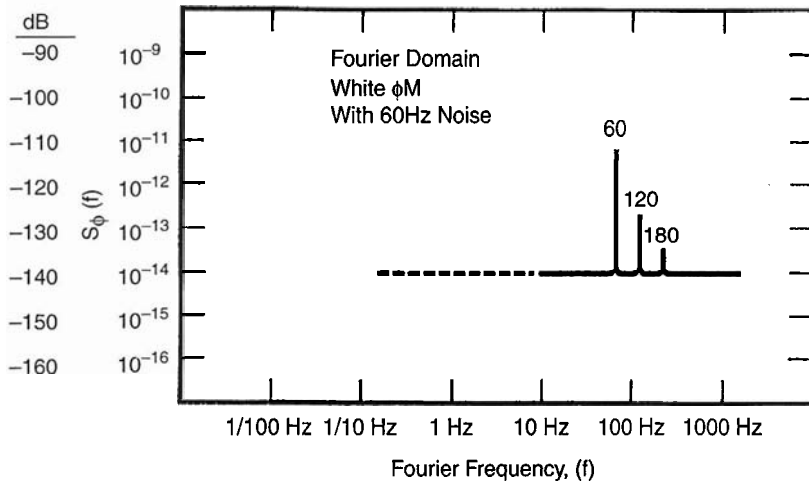


Figure 21. 60 Hz and harmonics are easily distinguished in a phase-noise measurement.

to the physical resonance mechanism of an active oscillator, or the design or choice of parts used for the electronics, or environmental properties. Flicker FM is common in high-quality oscillators, but may be masked by white FM ($1/f^2$) or flicker PM ($1/f$) in lower-quality oscillators.

3. White FM ($1/f^2$) noise is a common type found in passive-resonator frequency standards. These contain a slave oscillator, often quartz crystal, which is locked to a resonance feature of another device that behaves much like a high-Q filter. Cesium and rubidium standards have white FM noise characteristics.
4. Flicker PM ($1/f$) noise may relate to a physical resonance mechanism in an oscillator, but it usually is added by noisy electronics. This type of noise is common even in oscillators of the highest quality because in order to bring the signal amplitude up to a usable level, amplifiers are used after the signal source. Flicker PM noise may be introduced in these stages. It may also be introduced in a frequency multiplier. Flicker PM can be reduced with good low-noise amplifier design (i.e., using RF negative feedback) and hand-selecting transistors and other electronic components.
5. White PM (f^0) noise is broadband phase noise and has little to do with the resonance mechanism. It is probably produced by phenomena similar to that of flicker PM ($1/f$) noise. Stages of amplification are usually responsible for white PM noise. This noise can be kept at a very low value with good amplifier design, hand-selected components, the addition of narrowband filtering at the output, or, if feasible, increasing the power of the primary frequency source.

6.2. Other Types of Noise

A commonly encountered type of noise from a signal source or measurement apparatus is the presence of

60-Hz AC line noise. Shown in Fig. 21 is a constant white PM noise source with 60-, 120- and 180-Hz components added. This kind of noise is usually caused by AC power getting into the measurement system or the source under test. In the plot of $S_\phi(f)$, we observe discrete line spectra. Although $S_\phi(f)$ is a measure of spectral density, we can interpret the line spectra with no loss of generality, although one seldom refers to spectral densities when characterizing discrete lines. Figure 22 is the time-domain representation of the same white phase modulation level with 60-Hz noise. Note that the amplitude of $\sigma_y(\tau)$ varies up and down depending on sampling time. This is because in the time domain the sensitivity to a periodic wave varies directly as the sampling interval. This effect (which is an aliasing effect) can be used as a tool for filtering out a periodic wave imposed on a signal source. By sampling in the time domain at integer periods, we can be virtually insensitive to the periodic (discrete line) term, which is a useful strategy for removing the effect of the periodic wave.

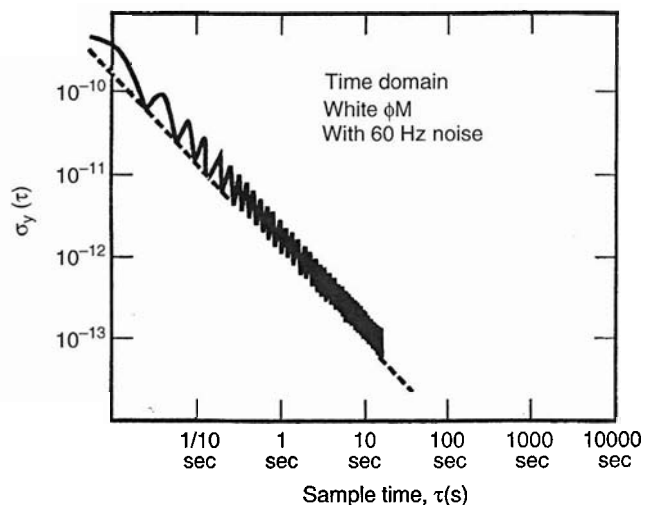


Figure 22. It is not easy to interpret an Allan deviation plot when 60 Hz noise is present on an oscillating signal.

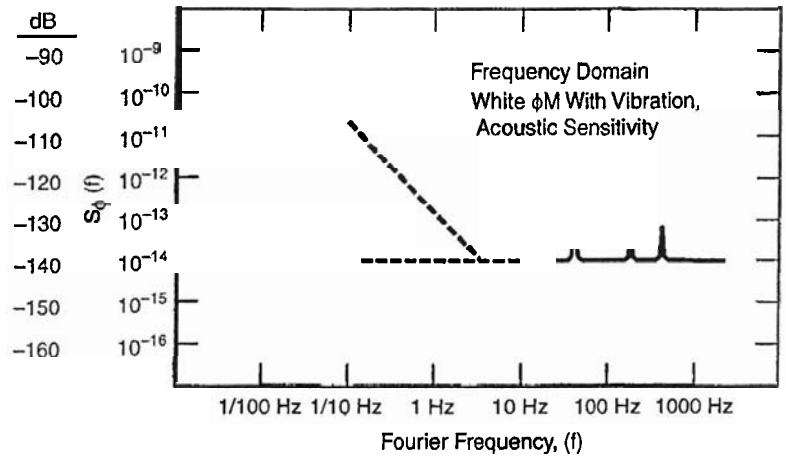


Figure 23. An oscillator under vibration causes sideband noise modulation that is apparent in a phase-noise measurement.

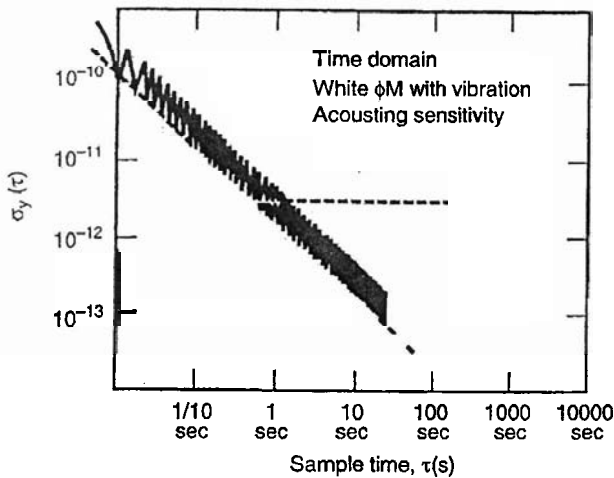


Figure 24. The Allan deviation of an oscillator under vibration causes a general increase in the level of frequency instability.

For example, diurnal variations in data due to day-to-day temperature, pressure, and other environmental effects can be eliminated by sampling the data once per day. This approach is useful for data with only one periodic term.

Figure 23 shows the kind of plot one might see of $S_{\phi}(f)$ with vibration and acoustic sensitivity in the signal source with the device under vibration. Figure 24 shows the translation of this effect to the time domain. Also noted in Fig. 23 is a (typical) flicker FM behavior in the low-frequency region. In the translation to time domain (Fig. 24), the flicker FM behavior masks the white PM (with the superimposed vibration characteristic) for long averaging times.

2. J. A. Barnes, A. R. Chi, L. S. Cutler, D. J. Healey, D. B. Leeson, T. E. McGunigal, J. A. Mullen, Jr., W. L. Smith, R. L. Sydnor, R. F. C. Vessot, and G. M. R. Winkler, *Characterization of Frequency Stability*, NBS Technical Note 394, 1970; *Proc. IEEE Trans. Instrum. Meas.* IM-20:105 (1971).
3. C. Audoin and B. Guinot, *The Measurement of Time: Time, Frequency and the Atomic Clock*, Cambridge Univ. Press, 2001.
4. D. W. Allan, Statistics of atomic frequency standards, *Proc. IEEE* 54:221 (1966).
5. D. A. Howe and T. F. Pepler, Very long-term frequency stability: Estimation using a special-purpose statistic, *Proc. 2003 Joint Meeting IEEE Int. Frequency Control Symp. and EFTF Conf.*, 2003.
6. D. A. Howe, NIST Technical Note 679, *Frequency Domain Stability Measurements: A Tutorial Introduction*, 1976.
7. J. Rutman, Characterization of phase and frequency instabilities in precision frequency sources: fifteen years of progress, *Proc. IEEE* 66(9):1048-1075 (1978).
8. IEEE Std 1139-1999, *Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE-SA Standards Board, J. R. Vig, chairperson (1999).
9. F. L. Walls, Secondary standard for PM and AM noise at 5, 10, and 100 MHz, *Proc. IEEE Trans. Instrum. Meas.* IM-42:136-143 (1993).

BIBLIOGRAPHY

1. D. B. Sullivan, D. W. Allan, D. A. Howe, and F. L. Walls, eds., *Characterization of Clocks and Oscillators*, NIST Technical Note 1337, 1990.