PROPERTIES OF SIGNAL SOURCES AND
MEASUREMENT METHODS

D. A. Howe, D. W. Allan, and J. A. Barnes

Time and Frequency Division
National Bureau of Standards
Boulder, Colorado 80303

Summary

This paper is a review of frequency stability measurement techniques and of noise properties of frequency sources.

First, a historical development of the usefulness of spectrum analysis and time domain measurements will be presented. Then the rationale will be stated for the use of the two-sample (Allan) variance rather than the classical variance. Next, a range of measurement procedures will be outlined with the trade-offs given for the various techniques employed. Methods of interpreting the measurement results will be given. In particular, the five commonly used noise models (white PM, flicker PM, white FM, flicker FM, and random walk FM) and their causes will be discussed. Methods of characterizing systematics will also be given. Confidence intervals on the various measures will be discussed. In addition, we will point out methods of improving this confidence interval for a fixed number of data points.

Topics will be treated in conceptual detail. Only light (fundamental) mathematical treatment will be given.

Although traditional concepts. will be detailed, two new topics will be introduced in this paper: (1) accuracy limitations of digital and computer-based analysis and (2) optimizing the results from a fixed set of input data.

The final section will be devoted to fundamental (physical) causes of noise in commonly used frequency standards. Also transforms from time to frequency domain and vice-versa will be given.

Key Words. Frequency stability; Oscillator noise modeling; Power law spectrum; Time-domain stability; Frequency-domain stability; White noise; Flicker noise.

Introduction

Precision oscillators play an important role in high speed communications, navigation, space tracking, deep space probes and in numerous other important applications. In this paper, we will review some precision methods of measuring the frequency and frequency stability of precision oscillators. Development of topics does not rely heavily on mathematics. The equipment and set-up for stability measurements are outlined. Examples and typical results are presented. Physical interpretations of common noise processes are discussed. A table is provided by which typical frequency domain stability characteristics may be translated to time domain stability characteristics and vice-versa.

I. THE SINE WAVE AND STABILITY

A sine wave signal generator produces a voltage that changes in time in a sinusoidal way as shown in figure 1.1. The signal is an oscillating signal because the sine wave repeats itself. A cycle of the oscillation is produced in one period "T". The phase is the angle "Φ" within a cycle corresonding to a particular time "t".
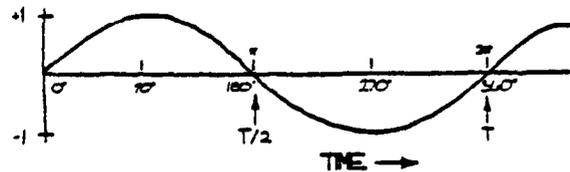

FIGURE 1.1

It is convenient for us to express angles in radians rather than in units of degrees, and positive zero-crossings will occur at even multiples of π-radians. The frequency "ν" is the number of cycles in one second, which is the reciprocal of period (seconds per cycle). The expression describing the voltage "V" out of a sine wave signal generator is given by $V(t) = V_p \sin [\Phi(t)]$ where V_p is the peak voltage amplitude. Equivalent expressions are

$$V(t) = V_p \sin \left(2\pi \frac{t}{T} \right)$$

1

and

$$V(t) = V_p \sin(2\pi\nu t).$$

Consider figure 1.2. Let's assume that the maximum value of "V" equals 1, hence "V_p" = 1. We say that the voltage "V(t)" is normalized to unity. If we know the <u>frequency</u> of a signal and if the signal is a <u>sine wave</u>, then we can determine the incremental change in the period "T" (denoted by Δt) at a particular angle of phase.
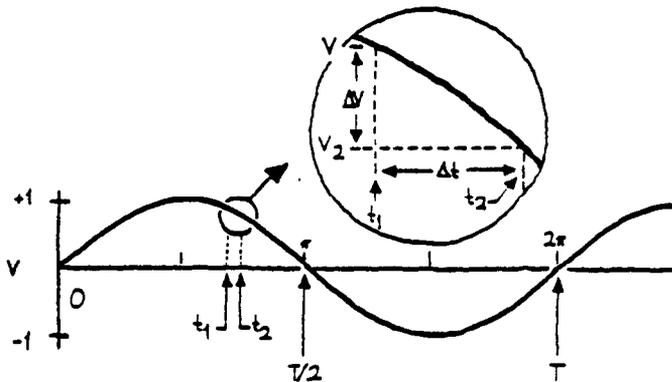

FIGURE 1.2

Note that no matter how big or small Δt may be, we can determine ΔV. Let us look at this from another point of view. Suppose 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 <u>instantaneous frequency</u> at this t. The smaller the interval Δt, the better the approximation of instantaneous frequency at t.

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 the frequency fluctuations of the oscillator signal. There is no formal definition for "frequency stability". However, one usually refers 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. In general terms,

"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".

Let's examine the two waveforms shown in figure 1.3. Frequency stability depends on the amount of time involved in a measurement. Of the two oscillating signals, it is evident that "2" is more stable than "1" from time t_1 to t_3 assuming the horizontal scales are linear in time.
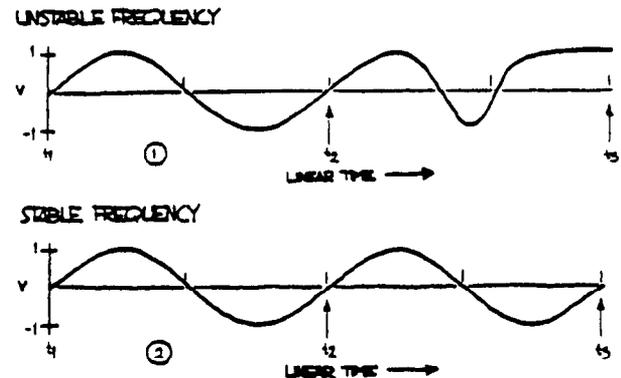

FIGURE 1.3

From time t_1 to 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 t_3, signal "1" is at a different frequency 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 which is undesirable. In the design of an oscillator, it is important to consider the sources of mechanisms which degrade the oscillator's frequency stability. All undesirable mechanisms cause random (noise) or systematic processes to exist along with 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

$$V(t) = [V_0 + \varepsilon(t)] \sin [2\pi\nu_0 t + \phi(t)]. \quad (1.1)$$

2

where V_0 ≡ nominal peak voltage amplitude,

$\varepsilon(t)$ ≡ deviation of amplitude from nominal,

ν_0 ≡ nominal fundamental frequency,

$\phi(t)$ ≡ deviation of phase from nominal.

Ideally "ε" 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 "ε" and "ϕ", we shall turn our attention to measurement techniques.

The typical precision oscillator, of course, has a very 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 measured, but with little confusion it is often 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 of phase or of the period fluctuations in an oscillator, not of frequency, even though the frequency may be the readout. As an example, most frequency counters sense the zero (or near zero) crossing of the sinusoidal voltage, which is the point at which the voltage is the most sensitive to phase fluctuations.

One must also realize that any frequency measurement 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 be enough better than the other that the fluctuations measured may be considered essentially those of the latter. However, in general because frequency measurements are always dual, it is useful to define:

$$y(t) = \frac{\nu_1 - \nu_0}{\nu_0} \qquad (1.2)$$

as the fractional frequency difference or deviation of oscillator one, ν_1, with respect to a reference oscillator ν_0 divided by the nominal frequency ν_0. Now, $y(t)$ is a dimensionless quantity and useful

in describing oscillator and clock performance; e.g., the time deviation, $x(t)$, of an oscillator over a period of time t, is simply given by:

$$x(t) = \int_0^t y(t')dt' \qquad (1.3)$$

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 is always some sample time. So when determining a fractional frequency, $y(t)$, in fact what is happening is that the time deviation is being measured say starting at some time t and again at a later time, $t + \tau$. The difference in these two time deviations, divided by τ gives the average fractional frequency over that period τ:

$$\bar{y}(t) = \frac{x(t + \tau) - x(t)}{\tau} \qquad (1.4)$$

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

What happens in many cases is that one samples a number of cycles of an oscillation during the preset gate time of a counter; after the gate time has elapsed, the counter latches the value of the number of cycles so that it can be read out, printed, or stored in some other way. Then there is a delay time for such processing of the data before the counter arms and starts again on the next cycle of the oscillation. During the delay time (or process time), information is lost. We have chosen to call it dead time and in some instances it becomes a problem. Unfortunately for data processing in typical oscillators the effects of dead time 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 dead time, that is usually where dead time can make a significant difference in the data analysis. Typically for many oscillators, if

3

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 which contribute zero or negligible dead time.[2]

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

We will describe five different methods of measuring the frequency fluctuations in precision oscillators.

1.1 Common Methods of Measuring Frequency Stability

A. Beat frequency method

The first system is called a heterodyne frequency measuring method or beat frequency method. The signal from two independent oscillators are fed into the two ports of a double balanced mixer as illustrated in figure 1.4.
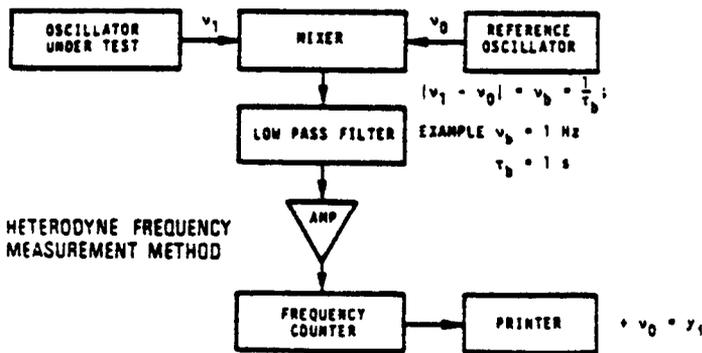

HETERODYNE FREQUENCY
MEASUREMENT METHOD

$$|v_1 - v_0| = v_b = \frac{1}{\tau_b}$$

EXAMPLE $v_b = 1$ Hz

$\tau_b = 1$ s

$v_0 = v_1$

FIGURE 1.4

The difference frequency or the beat frequency, v_b, is obtained as the output of a low pass filter which follows the mixer. This beat frequency is then amplified and fed to a frequency counter and printer or to some 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.

B. Dual mixer time difference (DTMD) system

This system shows some significant promise and has just begun to be exploited. A block diagram is shown is figure 1.5.

DUAL MIXER TIME DIFFERENCE SYSTEM

TIME DIFFERENCE: $x(1) = t_2 - t_1 = \frac{\Delta\epsilon(1)}{\tau v} - \frac{\tau}{2\pi v}$

FRACTIONAL FREQUENCY: $y(1) = \frac{\Delta v(1)}{v} = \frac{\Delta\epsilon(1+1) - \Delta\epsilon(1)}{\tau^2 v}$
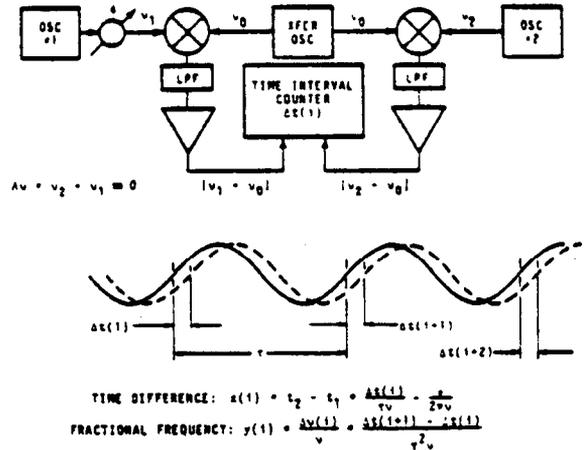
FIGURE 1.5

To preface the remarks on the DMTD, it should be mentioned that if the time or the time fluctuations can be measured directly, an advantage is obtained over just measuring the frequency. The reason is that one can calculate the frequency from the time without dead time as well as know the time behavior. The reason, in the past, that frequency has not been inferred from the time (for sample times of the order of several seconds and shorter) is that the time difference between a pair of oscillators operating as clocks could not be measured with sufficient precision (commercially the best that is available is 10^{-11} seconds). The system described in this section demonstrates a precision of 10^{-13} seconds. Such precision opens the door to making time measurements as well as frequency and frequency stability measuements for sample times as short as a few milliseconds and longer, all without dead time.

In figure 1.5, oscillator 1 could be considered under test and oscillator 2 could be considered the reference oscillator. These signals go to the ports of a pair of double balanced mixers. Another oscillator with separate symmetric buffered outputs is fed to the remaining other two

4

ports of the pair of double balanced mixers. This common oscillator's frequency is offset by a desired amount from the other two oscillators. Then two different beat frequencies come out of the two mixers as shown. These two beat frequencies will be out of phase by an amount proportional to the time difference between oscillator 1 and 2—excluding the differential phase shift that may be inserted. 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 one has oscillator 1 and oscillator 2 on the same frequency. This is typical for atomic standards (cesium, rubidium, and hydrogen frequency standards).

Illustrated at the bottom of figure 1.5 is what might represent the beat frequencies out of the two mixers. A phase shifter may be inserted as illustrated to adjust the phase so that the two beat rates are nominally in phase; this adjustment sets up the nice condition that the noise of the common oscillator tends to cancel (for certain types of noise) when the time difference is determined. After amplifying these beat signals, the start port of a time interval counter is triggered with the positive zero crossing of one beat and the stop port with the positive zero crossing of the other beat. Taking the time difference between the zero crossings of these beat frequencies, one measures the time difference between oscillator 1 and oscillator 2, but with a precision which has been amplified by the ratio of the carrier frequency to the beat frequency (over that normally achievable with this same time interval counter). The time difference $x(i)$ for the i^{th} measurement between oscillators 1 and 2 is given by eq (1.5).

$$x(i) = \frac{\Delta t(i)}{\tau_b \nu_o} - \frac{\phi}{2\pi \nu_o} + \frac{k}{\nu_o} \qquad (1.5)$$

where $\Delta t(i)$ is the i^{th} time difference as read on the counter, τ_b is the beat period, ν_o is the nominal carrier frequency, ϕ is the phase delay in radians added to the signal of oscillator 1, and k is an integer to be determined in order to remove

the cycle ambiguity. It is only important to know k if the absolute time difference is desired; for frequency and frequency stability measurements and for time fluctuation measurements, k may be assumed zero unless one goes through a cycle during a set of measurements. The fractional frequency can be derived in the normal way from the time fluctuations.

$$y_{1,2}(i, \tau) = \begin{cases} \dfrac{\nu_1(i, \tau) - \nu_2(i, \tau)}{\nu_o} \\[2ex] \dfrac{x(i + 1) - x(i)}{\tau} \\[2ex] \dfrac{\Delta t(i + 1) - \Delta t(i)}{\tau_b^2 \nu_o} \end{cases} \qquad (1.6)$$

In eqs (1.5) and (1.6), assumptions are made that the transfer (or common) oscillator is set at a lower frequency than oscillators 1 and 2, and that the voltage zero crossing of the beat $\nu_1 - \nu_c$ 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:

$$y_{1,2}(i, m\tau_b) = \frac{x(i + m) - x(i)}{m\tau_b} \qquad (1.7)$$

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 set-up the nominal beat frequencies are simply given by the amount that the output frequency of the synthesizer is offset from ν_2. Sample times as short as a few milliseconds are easiliy obtained. Logging the data at such a rate can be a problem without special equipment. The latest NBS time scale measurement system is based on the DMTD and is yielding an excellent cost benefit ratio.

C. Loose phase lock loop method

This first type of phase lock loop method is illustrated in figure 1.6. The signal from an oscillator under test is fed into one port of a

TN-18

mixer. The signal from a reference oscillator is fed into the other port of this mixer. The signals are in quadrature, that is, they are 90 degrees out
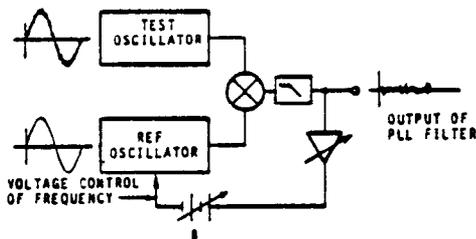

FIGURE 1.6

of phase so that the average voltage out of the 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 making low noise precision stability measurements. The output of this mixer is fed through a low pass filter and then amplified in a feedback loop, causing the voltage controlled oscillator (reference) to be phase locked to the test oscillator. The attack time of the loop is adjusted such that a very loose phase lock (long time constant) condition exists. This is discussed later in section VIII.

The attack time is the time it takes the servo system to make 70% of its ultimate correction after being slightly disturbed. The attack time is equal to $1/\pi w_h$, where w_h is the servo bandwidth. If the attack time of the loop is about a second then the voltage fluctuations will be proportional to the phase fluctuations for sample times shorter than the attack time. 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 specturm analyzer to measure the Fourier components of the phase fluctuations. This system of frequency-domain analysis is discussed in sections VIII to X.

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 specifically very useful if one has discrete side bands such as 60Hz or detailed structure in the spectrum. How to characterize precision oscillators using this technique will be treated in detail later in section IX and XI.

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

D. Tight phase lock loop method

The second type of phase lock loop method (shown in figure 1.7) is essentially the same as the first in figure 1.6 except that in this case the loop is in a tight phase lock condition; i.e., the attack time of the loop should be of the order of a few milliseconds. In such a case, the phase fluctuations are being integrated so that the voltage output is proportional to the frequency
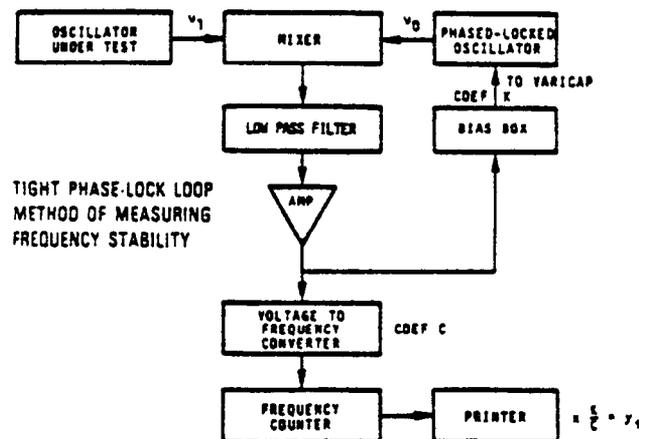

FIGURE 1.7

fluctuations between the two oscillators and is no longer proportional to the phase fluctuations for sample times longer than the attack time of the loop. A bias box is used to adjust the voltage on the varicap to a tuning point that is fairly linear and of a reasonable value. The voltage

TN-19

fluctuations prior to the bias box (biased slightly away from zero) may be fed to a voltage to frequency converter which in turn is fed to a frequency counter where one may read out the frequency fluctuations with great amplification of the instabilities between this pair of oscillators. The frequency counter data are logged with a data logging device. The coefficient of the varicap and the coefficient of the voltage to frequency converter are used to determine the fractional frequency fluctuations, y_i, between the oscillators, where i denotes the i^{th} measurement as shown in figure 1.7. It is not difficult to achieve a sensitivity of a part in 10^{14} per Hz resolution of the frequency counter, so one has excellent precision capabilities with this system.

E. Time difference method

The last measurement method we will illustrate is very commonly used, but typically does not have the measurement precision more readily available in the first four methods illustrated above. This method is called the time difference method, and is shown in figure 1.8. Because of the wide
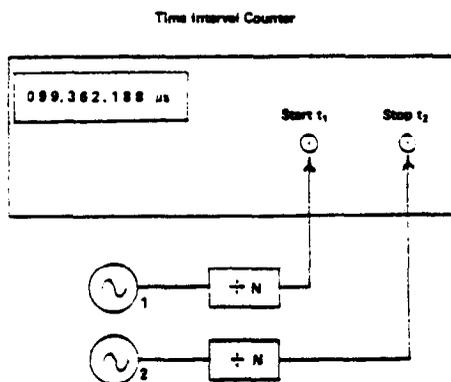

FIGURE 1.8

bandwidth needed to measure fast rise-time pulses, this method is limited in signal-to-noise ratio. However, some counters are commercially available allowing one to do signal averaging or to do precision rise-time comparison (precision of time difference measurements in the range of 10 ns to 10 ps are now available). Such a method yields a direct measurement of x(t) without any translation,

conversion, or multiplication factors. Caution should be exercised in using this technique even if adequate measurement precision is available because it is not uncommon to have significant instabilities in the frequency dividers shown in figure 1.8--of the order of several nanoseconds. The technology exists to build better frequency dividers than are commonly available, but manufacturers have not yet availed themselves of state-of-the-art techniques in a cost beneficial manner. A trick to by-pass divider problems is to feed the oscillator signals directly into the time interval counter and observe the zero voltage crossing into a well matched impedance. (In fact, in all of the above methods one needs to pay attention to impedance matching, cable lengths and types, and connectors). 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 $2\pi V_p/\tau_1$, where $\tau_1 = 1/\nu_1$ (the period of oscillation). For $V_p = 1$ volt and a 5 MHz signal, this slope is 3m volts/ns, which is a very good sensitivity.

II. MEASUREMENT METHODS COMPARISON

When making measurements between a pair of frequency standards or clocks, it is desirable to have less noise in the measurement system than the composite noise in the pair of standards being measured. This places stringent requirements on measurement systems as the state-of-the-art of precision frequency and time standards has advanced to its current level. As will be shown, perhaps one of the greatest areas of disparity between measurement system noise and the noise in current standards is in the area of time difference measurements. Commercial equipment can measure time differences to at best 10^{-11}s, but the time fluctuations second to second of state-of-the-art standards is as good as 10^{-13}s.

The disparity is unfortunate because if time differences between two standards could be measured with adequate precision then one may also know the time fluctuations, the frequency differences, and the frequency fluctuations. In fact, one can set

7

Hierachy Status	Example of Measurement Method	Data deducible from the measurement			
		Time	δTime	Freq.	δFreq.
1	Dual Mixer Time Diff. or Time Interval Counter	$x(t) = \frac{\Delta t}{\tau\nu}$	$\delta x(t,\tau) = x(t+\tau)-x(t)$	$y(t,\tau) = \frac{\delta x(t,\tau)}{\tau}$	$\delta y(t,\tau) = y(t+\tau,\tau)-y(t,\tau)$
2	Loose Phase-locked reference oscillator	can't measure	$\delta x(t,\tau) = \frac{\delta \phi(t,\tau)}{2\pi\nu}$	"	"
3	Heterodyne or beat frequency	can't measure	can't measure	$\|y(t,\tau)\| = \frac{\nu_{beat}}{\nu_o}$	"
4	Tight phase-locked reference oscillator	can't measure	can't measure	can't measure	$\delta y(t,\tau) = c\cdot\delta V$

TABLE 2.1

Measurement Method	Time accuracy	Time stability	Frequency accuracy (1 day)	Frequency stability	Advantages	Disadvantages	Approximate Cost (kS)
Dual Mixer Time Difference	~100ps	0.1ps	~10^{-16}	~1.10^{-13} x τ^{-1}	No dead time; many choose sample time (1ms to as large as desired); oscillators may be at zero beat or different; measurement bandwidth easily changed; measurement noise typically below oscillators noise for tau~1s and longer; measures times, time stability, frequency, and frequency stability.	More complex than other methods, and hence is more suceptable to extranious signal pickup, e.g., ground loops; the time difference is module the beat period, e.g., 200ns @ 5MHz.	~5 to 30
Time Interval Counter	~100ps	10ps	~10^{-14}	~1.10^{-11} x τ^{-1}	Wide bandwidth input allows a variety of signals; simple to use; cycle ambiguity almost never a problem; measures time, time stability, frequency, and frequency stability.	Measurement noise typicaly in excess of oscillator instabilities for tau values up to and of the order of several thousand seconds; hence, typically not useful for short term stability analysis.	1 to 7
Loose Phase-lock reference[a] oscillator	---	~0.1ps	depends on calibration of varicap curve	~1.10^{-11} x τ^{-1}	Useful for short term time stability analysis as well as spectrum analysis; excellent phase sensitivity—typically better than most oscillators.	Calibration depends on appropriate reference oscillator varicap curve characteristics.	<1
Heterodyne or Beat Frequency	---	---	~10^{-16}	~1.10^{-13} x τ^{-1}	Measurement noise can typically be made less than oscillator instabilities for tau =1s and longer.	Minimum tau determined by period of beat frequency—typically not adjustable; can not compare oscillators near zero beat; cannot tell which oscillator is high or low in frequency; dead time often associated with these measurements.	0.2 to 1
Tight phase-locked reference[a] oscillator	---	---	---	~1.10^{-13} x τ^{-1}	Measurement noise typically less than oscillator instabilities for tau=1s and longer; good measurement system band-width control; dead time can be made small or negligible.	Need voltage controlled reference oscillator; frequency sensitivity is a function of varicap tuning curve, hence not conducive to measuring absolute frequency differences.	<1

[a] Assumes reference oscillator is available to the metrologist.

TABLE 2.2

TN-21

up an interesting hierarchy of kinds of measurement systems: 1) those that can measure time, $x(t)$; 2) those that can measure changes in time or time fluctuations $\delta x(t)$; 3) those that can measure frequency, $\nu(y \equiv (\nu-\nu_0)/\nu_0)$; and 4) those that can measure changes in frequency or frequency fluctuations, $\delta\nu$ ($\delta y \equiv \delta\nu/\nu_0$). As depicted in table 2.1, if a measurement system is of status 1 in this hierarchy, i.e., it can measure time, then time fluctuations, frequency and frequency fluctuations can be deduced. However, if a measurement system is only capable of measuring time fluctuations (status 2 - table 2.1), then time cannot be deduced, but frequency and frequency fluctuations can. If frequency is being measured (status 3 - table 2.1), then neither time nor time fluctuations may be deduced with fidelity because essentially all commercial frequency measuring devices have "dead time" (technology is at a point where that is changing with fast data processing speeds that are now available). Dead time in a frequency measurement destroys the opportunity of integrating the fractional frequency to get to "true" time fluctuations. Of course, if frequency can be measured, then trivially one may deduce the frequency fluctuations. Finally, if a system can only measure frequency fluctuations (status 4 - table 2.1), then neither time, nor time fluctuations, nor frequency can be deduced from the data. If the frequency stability is the primary concern then one may be perfectly happy to employ such a measurement system, and similarly for the other statuses in this measurement hierarchy. Obviously, if a measurement method of Status 1 could be employed with state-of-the-art precision, this would provide the greatest flexibility in data processing. From section 1, the dual mixer time difference system is purported to be such a method.

Table 2.2 is a comparison of these different measurement methods. The values entered are nominal; there may be unique situations where significant departures are observed. The time and frequency stabilities listed are the nominal second to second rms values. The accuracies listed are taken in an absolute sense. The costs listed are nominal estimates in 1981 dollars.

Figure 2.1 is a diagram indicating the sample time regions over which the various methods
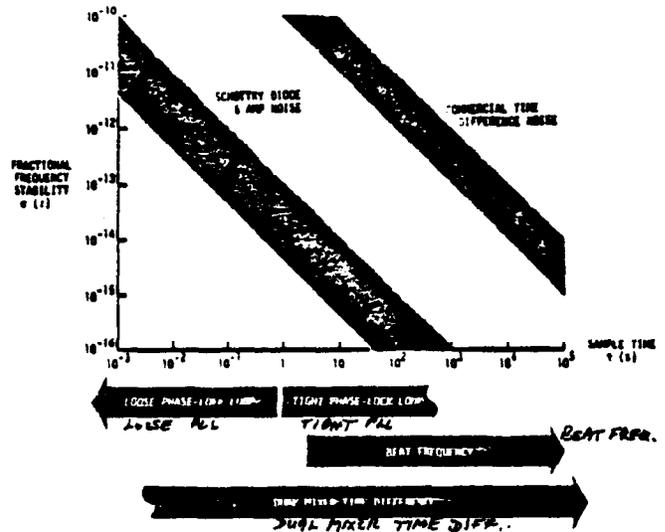

FIGURE 2.1

are most appropriately applied. The large diagonally oriented area indicates the typical noise limits of the measurement technique (at particular values of sample time indicated on the horizontal scale).

III. CHARACTERIZATION

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

3.1 Non-random Fluctuations

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

9

If, for example, one has the values of the frequency over a period of time and a frequency offset from nominal is observed, one may calculate directly that the phase error will accumulate as a ramp. If the frequency values show some linear drift then the time fluctuations will depart as a quadratic. In almost all oscillators, the above systematics, as they are sometimes called, are the primary cause of time and/or frequency departure. A useful approach to determine the value of the frequency offset is to calculate the simple mean of the set, or for determining the value of the frequency drift by calculating a linear least squares fit to the frequency. A least squares quadratic fit to the phase or to the time derivative is typically not as efficient an estimator of the frequency drift for most oscillators.

3.2 Random Fluctuations

After calculating or estimating the systematic or non-random effects of a data set, these may be subtracted from the data leaving the residual random fluctuations. These can usually be best characterized statistically. It is often the case for precision oscillators that these random fluctuations may be well modeled with power law spectral densities. This topic is discussed later in sections VIII to X. We have

$$S_y(f) = h_\alpha f^\alpha, \tag{3.1}$$

where $S_y(f)$ is the one-sided spectral density of the fractional frequency fluctuations, f is the Fourier frequency at which the density is taken, h_α is the intensity coefficient, and α is a number modeling the most appropriate power law for the data. It has been shown[1,3] that in the time domain one can nicely represent a power law spectral density process using a well defined time-domain stability measure, $\sigma_y(\tau)$, to be explained in the next section. For example, if one observes from a log $\sigma_y^2(\tau)$ versus τ 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 $1 \leq \alpha$. Further a correspondence exists between h_α and the coefficient for $\sigma_y(\tau)$. These coefficients have been calculated and appear in section XI. 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 that have been simulated by computer are shown in figure 3.1. In descending order these
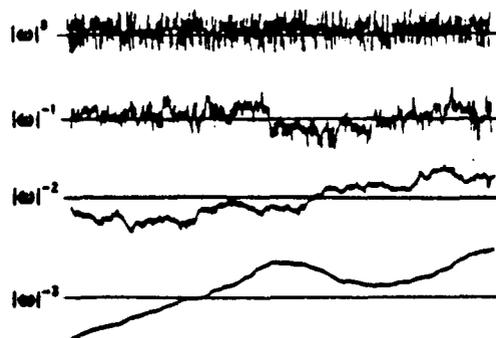
POWER LAW SPECTRA

FIGURE 3.1

have been named white noise, flicker noise, random walk, and flicker walk (the ω in fig. 3.1 is angular Fourier frequency, $\omega = 2\pi f$).

Once the noise characteristics have been determined, one is often able to deduce whether the oscillators are performing properly or not and whether they are meeting either the design specifications or the manufacturers specifications. For example a cesium beam frequency standard or a rubidium gas cell frequency standard when working properly should exhibit white frequency noise, which is the same as random walk phase (or time) for tau values of the order of a few seconds to several thousand seconds (see also sec. XI).

IV. ANALYSIS OF TIME DOMAIN DATA

Suppose now that one is given the time or frequency fluctuations between a pair of precision oscillators measured, for example, by one of the techniques outlined in section I, and a stability analysis is desired. Let this comparison be depicted by figure 4.1. The minimum sample time

is determined by the measurement system. If the time difference or the time fluctuations are available then the frequency or the fractional frequency fluctuations may be calculated from one period of sampling to the next over the data length as indiciated in figure 4.1. 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 figure 4.1, $\sigma_{std. dev.}(\tau)$, where $\bar{y}$ is the average fractional frequency over the data set and is subtracted from each value of y_i before squaring, summing and dividing by the number of values minus one, (M-1), and taking the square root to get the standard deviation. At NBS, we have studied what happens to the standard deviation when the data set may be characterized by power law spectra which are more dispersive than classical white noise frequency fluctuations. In other words, if the fluctuations are characterized by flicker noise or any other non-white-noise frequency deviations, what happens to the standard deviation for that data set? One can show that the standard deviation is a function of the number of data points in the set; 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 monotonically increases without limit. Some statistical measures have been developed which do not depend upon the data length and which are readily usable for characterizing the random fluctuations in precision oscillators. An IEEE subcommittee on frequency stability has recommended what has come to be 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 figure 4.1. This equation is very easy to implement experimentally as one simply need add up the squares of the differences between adjacent values of y_i, divide by the number of them and by two, and take the square root. One then has the quantity which the IEEE subcommittee has recommended for

specification of stability in the time domain—denoted by $\sigma_y(\tau)$.

$$\sigma_y(\tau) = \left\langle \frac{1}{2} \ (\bar{y}(t+\tau) - \bar{y}(t))^2 \right\rangle^{\frac{1}{2}}, \qquad (4.1)$$

where the brackets "<>" denote infinite time average. In practice this is easily estimated from a finite data set as follows:

$$\sigma_y(\tau) \cong \left[\frac{1}{2(M-1)} \ \sum_{i=1}^{M-1} \ \left(y_{i+1} - y_i \right)^2 \right]^{\frac{1}{2}}, \qquad (4.2)$$

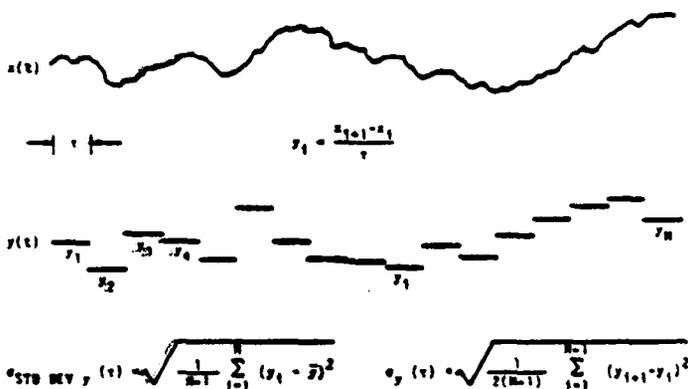where the y_i are the discrete frequency averages as illustrated in figure 4.1.

FIGURE 4.1

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

One would like to know how $\sigma_y(\tau)$ varies with the sample time, τ. A simple trick that one can use that is very useful if there is no dead time, is to average the previous values for y_1 and y_2 and call that a new y_1 averaged over 2τ, similarly average the previous values for y_3 and y_4 and call that a new y_2 averaged over 2τ etc., and finally apply the same equation as before to get $\sigma_y(2\tau)$. One can repeat this process for other desired integer multiples of τ and from the same data set

11

be able to generate values for $\sigma_y(m\tau)$ as a function of $m\tau$ from which one may be able to infer a model for the process that is characteristic of this pair of oscillators. If one has dead time in the measurements adjacent pairs cannot be averaged in an unambiguous way to simply increase the sample time. One has to retake the data for each new sample time--often a very time consuming task. This is another instance where dead time can be a problem.

How the classical variance (standard deviation squared) depends on the number of samples is shown in figure 4.2. Plotted is the ratio of the standard deviation squared for N samples to the standard deviation squared for 2 samples; $<\sigma^2(2,\tau)>$ is the same as the Allan variance, $\sigma_y^2(\tau)$. One can see the dependence of the standard deviation upon the number of samples for various kinds of power
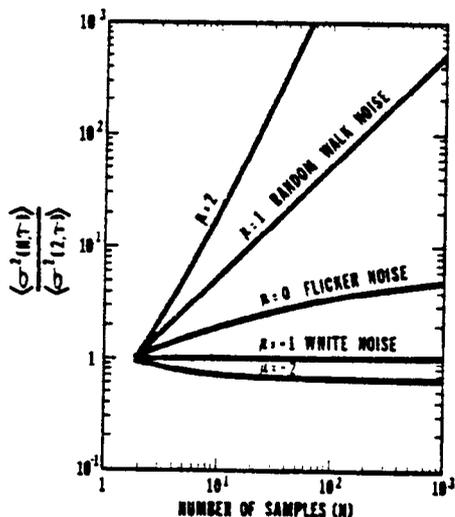
FIGURE 4.2

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

law spectral densities commonly encountered as reasonable models for many important precision oscillators. Note, $\sigma_y^2(\tau)$ has the same value as

the classical variance for the classical noise case (white noise FM). One main point of figure 4.2 is simply to show that with the increasing data length the standard deviation of the common classical variance is not well behaved for the kinds of noise processes that are very often encountered in most of the precision oscillators of interest.

One may combine eq (1.4) and eq (4.1), which yields an equation for $\sigma_y(\tau)$ in terms of the time difference or time deviation measurements.

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

(4.3)

which for N discrete time readings may be estimated as,

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

(4.4)

where the i denotes the number of the reading in the set of N and the nominal spacing between readings is τ. If there is no dead time in the data and the original data were taken with the x's spaced by τ_0, we can pick τ in eq (4.4) to be any integer multiple of τ_0, i.e., $\tau = m\tau_0$:

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

(4.5)

Equation (4.5) has some interesting consequences because of the efficient data usage in terms of the confidence of the estimate as will be explained in the next section.

EXAMPLE: Find the Allan variance, $\sigma_y^2(\tau)$, of the following sequence of fractional frequency fluctuation values y_k, each value averaged over one second.

$y_1 = 4.36 \times 10^{-5}$ $y_5 = 4.47 \times 10^{-5}$

$y_2 = 4.61 \times 10^{-5}$ $y_6 = 3.96 \times 10^{-5}$

$y_3 = 3.19 \times 10^{-5}$ $y_7 = 4.10 \times 10^{-5}$

$y_4 = 4.21 \times 10^{-5}$ $y_8 = 3.08 \times 10^{-5}$

(assume no dead-time in measurement of averages)

TN-25

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

Data values y_k (x 10^{-4})	First differences $(y_{k+1} - y_k)$ (x 10^{-5})	First difference squared $(y_{k+1} - y_k)^2$ (x 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
		4.51

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

Therefore,

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

and

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

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

One sees that with large numbers of data values, it is helpful to use a computer or programmable calculator. 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 the $1/\sqrt{8} \times 100\% \approx 35\%$. This then is the allowable error in our estimate for the $\tau = 1s$ average. The next section shows methods of computing and improving the confidence interval.

V. CONFIDENCE OF THE ESTIMATE AND OVERLAPPING SAMPLES[4]

One can imagine taking 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. From these two frequency measurements, one can calculate a single sample Allan (or two-sample) variance (see fig. 5.1). Of course this variance does not have high precision or confidence since it is based on only one frequency difference.
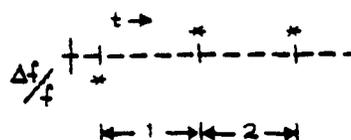

FIGURE 5.1

Statisticians have considered this problem of quantifying the variability of quantities like the Allan Variance. Conceptually, one could imagine repeating the above experiment (of taking the three phase points and calculating the Allan Variance), many times and even calculating the distribution of the values.

For the above cited experiment we know that the results are distributed like the statistician's chi-square distribution with one degree of freedom. That is, we know that for most common oscillators the first difference of the frequency is a normally distributed variable with the typical bell-shaped curve and zero mean. However, the square of a normally distributed variable is NOT normally distributed. That is easy to see since the square is always positive and the normal curve is completely symmetric and negative values are as likely as positive. The resulting distribution is called a chi-square 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.

In contrast, if we took five phase values, then we could calculate four consecutive frequency values, as in figure 5.2. We could then take the
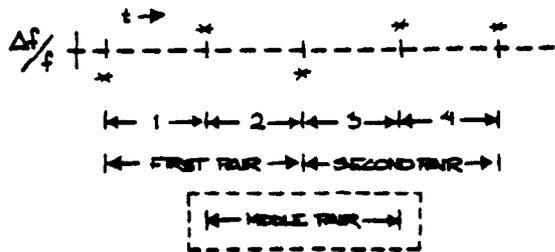
13

FIGURE 5.2

first pair and calculate a sample Allan Variance, and we could calculate a second sample Allan Variance 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 we would expect it to have a tighter confidence interval than in the previous example. This could be expressed with the aid of the chi-square distribution with TWO degrees of freedom.

However, there is another option. We could also consider the sample Allan Variance obtained from the second and third frequency measurements. That is the middle sample variance. Now, however, we're in trouble because clearly this last sample Allan Variance is NOT independent of the other two. Indeed, it is made up of parts of each of the other two. This does NOT mean that we can't use it for improving our estimate of the "true" Allan Variance, but it does mean that we can't just assume that the new average of three sample Allan Variances is distributed as chi-square with three degrees of freedom. Indeed, we will encounter chi-square distributions with fractional degrees of freedom. And as one might expect, the number of degrees of freedom will depend upon the underlying noise type, that is, white FM, flicker FM, or whatever.

Before going on with this, it is of value to review some concepts of the chi-square distribution. Sample variances (like sample Allan Variances) are distributed according to the equation:

$$\chi^2 = \frac{(d.f.) \cdot s^2}{\sigma^2} \qquad (5.1)$$

where S^2 is the sample Allan Variance, χ^2 is

chi-square, d.f. is the number of degrees of freedom (possibly not an integer), and σ^2 is the "true" Allan Variance we're all interested in knowing--but can only estimate imperfectly. Chi-square is a random variable and its distribution has been studied extensively. For some reason, chi-square is defined so that d.f., the number of degrees of freedom, appears explicitly in eq (5.1). Still, χ^2 is a (implicit) function of d.f., also.

The probability density for the chi-square distribution is given by the relation

$$p(\chi^2) = \frac{1}{2^{d.f.} \; \Gamma\left(\frac{d.f.}{2}\right)} (\chi^2)^{\frac{d.f.}{2} - 1} e^{\frac{-\chi^2}{2}} \qquad (5.2)$$

where $\Gamma\left(\frac{d.f.}{2}\right)$ is the gamma function, defined by the integral

$$\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} dx \qquad (5.3)$$

Chi-square distributions are useful in determining specified confidence intervals for variances and standard deviations. Here is an example. Suppose we have a sample variance $s^2 = 3.0$ and we know that this variance has 10 degrees of freedom. (Just how we can know the degrees of freedom will be discussed shortly.) Suppose also that we want to know a range around our sample value of $s^2 = 3.0$ which "probably" contains the true value, σ^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 our 90% in the middle. This is arbitrary and a specific problem might dictate a different allocation. We now resort to tables of the chi-square distribution and find that for 10 degrees of freedom the 5% and 95% points correspond to:

$$\chi^2(.05) = 3.94$$

$$\text{for d.f.} = 10 \qquad (5.4)$$

$$\chi^2(.95) = 18.3$$

14

Thus, with 90% probability the calculated sample variance, s^2, satisfies the inequality:

$$3.94 \leq \frac{(d.f.) \cdot s^2}{\sigma^2} \leq 18.3 \qquad (5.5)$$

and this inequality can be rearranged in the form

$$1.64 \leq \sigma^2 \leq 7.61 \qquad (5.6)$$

or, taking square roots:

$$1.28 \leq \sigma \leq 2.76 \qquad (5.7)$$

Now someone might object to the form of eq (5.7) since it seems to be saying that the true sigma falls within two limits with 90% probability. Of course, this is either true or not and is not subject to a probabilistic interpretation. Actually eq (5.7) is based on the idea that the true sigma is not known and we estimate it with the square root of a sample variance, s^2. This sample variance is a random variable and is properly the subject of probability, and its value (which happened to be 3.0 in the example) will conform to eq (5.7) nine times out of ten.

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

$$s^2 = \frac{1}{N-1} \sum_{n=1}^{N} (x_n - \bar{x})^2 \qquad (5.8)$$

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

Thus, for the case of white x_n, and a conventional sample variance (i.e., eq (5.8)), the number of degrees of freedom are given by the equation:

$$d.f. = N-1 \qquad (5.9)$$

The problem of interest here is to obtain the corresponding equations for Allan Variances using overlapping estimates on various types of noise (i.e., white FM, flicker FM, etc.).

Other authors (Lesage and Audoin, and Yoshimura) have considered the question of the variance of the Allan Variances without regard to the distributions. This is, of course, a closely related problem and use will be made of their results. These authors considered a more restrictive set of overlapping estimates than will be considered here, however.

VI. MAXIMAL USE OF THE DATA AND DETERMINATION OF THE DEGREES OF FREEDOM.

6.1 Use of Data

Consider the case of two oscillators being compared in phase and exactly N values of the phase difference are obtained. Assume that the data are taken at equally spaced intervals, τ_0. From these N phase values, one can obtain 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 variance has been calculated by the above cited authors.

Using the same set of data, it is also possible to estimate the Allan Variances for integer multiples of the base sampling interval, $\tau = n\tau_0$. Now the possibilities for overlapping sample Allan Variances are even greater. For a phase data set of N points one can obtain exactly N-2n sample Allan Variances for $\tau = n\tau_0$. Of course only a fraction of these are generally independent. Still the use of ALL of the data is well justified (see fig. 6.1).

Consider the case of an experiment extending for several weeks in duration with the aim of getting estimates of the Allan Variance for tau values equal to a week or more. As always the purpose is to estimate reliably the "true" Allan Variance as well as possible--that is, with as tight an uncertainty as possible. Thus one wants
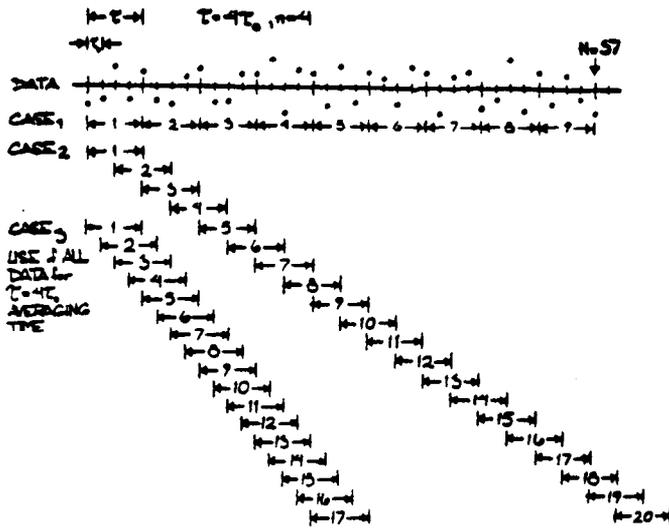
TN-28

FIGURE 6.1

to use the data as efficiently as possible since obtaining more data can be very expensive. The most efficient use is to average all possible sample Allan Variances of a given tau value that one can compute from the data.

The problem comes in estimating how tight the confidence intervals really are—that is, in estimating the number of degrees of freedom. Clearly, if one estimates the confidence intervals pessimistically, then more data is needed to reach a specified tolerance, and that can be expensive. The other error of over-confidence in a questionable value can be even more expensive. Ideally, one has realistic confidence estimates for the most efficient use of the data, which is the intent of this writing.

6.2 Determining the Degrees of Freedom

In principle, it should be possible to determine analytically the equations corresponding to eq (5.9) for all cases of interest. Unfortunately the analysis becomes quite complicated. Exact computer algorithms were devised for the cases of white phase noise, white frequency modulation and random walk FM. For the two flicker cases (i.e., flicker FM and PM) a completely empirical approach was used. Due to the complexity of the computer programs, empirical fits were devised for all five noise types.

The approach used is based on three equations relating to the chi-square distribution:

$$\frac{\chi^2}{(d.f.)} = \frac{s^2}{\sigma^2} \qquad (6.1)$$

$$E[\chi^2] = d.f. \qquad (6.2)$$

$$Var[\chi^2] = 2(d.f.) \qquad (6.3)$$

where the expression $E[\chi^2]$ means the "expectation," or average value of χ, $Var[\chi^2]$ is the variance of χ^2, and d.f. is the number of degrees of freedom.

A computer was used to simulate phase data sets of some length, N, and then Allan Variances with $\tau = n\tau_0$ were calculated for all possible samples. This "experiment" was repeated at least 1000 times using new simulated data sets of the same spectral type, and always of the same length, N. Since the data were simulated on a computer, the "true" Allan Variance, σ^2, was known for many of the noise models and could be substituted into eq (6.1). From the 1000 values of s^2/σ^2, distributions and sample variances were obtained. The "experimental" distributions were compared with theoretical distributions to verify that the observed distributions truely conformed to the chi-square distribution.

The actual calculation of the degrees of freedom were made using the relation:

$$d.f. = \frac{2(\sigma^2)^2}{Var(s^2)} \qquad (6.4)$$

which can be deduced from eqs (6.1), (6.2), and (6.3). The $Var(s^2)$ was estimated by the sample variance of the 1000 values of the average Allan Variances, each obtained from a phase data set of length N.

Of course this had to be repeated for various values of N and n, as well as for each of the five common noise types: white PM, flicker PM, white FM, flicker FM, and random walk FM. Fortunately, certain limiting values are known and these can be used as checks on the method. For example, when $(N-1)/2 = n$, only one Allan Variance is obtained from each data set and one should get about one degree of freedom for eq (6.4), which was observed

16

in fact. Also for n=1 the "experimental" conditions correspond to those used by Lesage and Audoin, and by Yoshimura. Indeed, the method also was tested by verifying that it gave results consistent with eq (5.9) when applied to the conventional sample variance. Thus, combining eq (6.4) with the equations for the variance of the Allan Variances from Lesage and Audoin and Yoshimura, one obtains:

White PM $\qquad$ d.f. $= \dfrac{18(N-2)^2}{35N-88}$, for $N \geq 4$

Flicker PM $\qquad$ d.f. $= ?$

White FM $\qquad$ d.f. $= \dfrac{2(N-2)^2}{3N-7}$ $\qquad$ (6.5)

Flicker FM $\qquad$ d.f. $= \dfrac{2(N-2)^2}{2.3N-4.9}$

Random Walk FM $\qquad$ d.f. $= N-2$

for n=1. Unfortunately, their results are not totally consistent with each other. Where inconsistency arose the value in best agreement with the "experimental" results was chosen.

The empirical equations which were fit to the "experimental" data and the known values are summarized below:

White PM $\qquad$ d.f. $\cong \dfrac{(N+1)(N-2n)}{2(N-n)}$

Flicker PM $\qquad$ d.f. $\cong \exp\left(\ln \dfrac{N-1}{2n} \; \ln \dfrac{(2n+1)(N-1)}{4}\right)$

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ (6.6)

White FM $\qquad$ d.f. $\cong \left[\dfrac{3(N-1)}{2n} - \dfrac{2(N-2)}{N}\right] \cdot \dfrac{4n^2}{4n^2 + 5}$

Flicker FM $\qquad$ d.f. $\cong \begin{cases} \dfrac{2(N-2)}{2.3N - 4.9}, & \text{for } n=1 \\[2mm] \dfrac{5N^2}{4n(N + 3n)}, & \text{for } n \geq 2 \end{cases}$

Random Walk FM $\quad$ d.f. $\cong \dfrac{N-2}{n} \; \dfrac{(N-1)^2-3n(N-1)+4n^2}{(N-3)^2}$

The figures in Appendix I demonstrate the fit to the "experimental" data.

It is appropriate to give some estimate of just how well these empirical equations approach the "true" values. The equations have approximately (a few percent) the correct assymptotic behavior at n=1 and n=(N-1)/2. In between, the values were tested (using the simulation results) over the range of N=5 to N=1025 for n=1 to n=(N-1)/2 changing by octaves. In general, the fit was good to within a few percent. We must acknowledge that distributional problems with the random number generators can cause problems, although there were several known values which should have revealed these problems if they are present. Also for three of the noise types the exact number of degrees of freedom were calculated for many values of N and n and compared with the "Monte Carlo" calculations. The results were all very good.

Appendix I presents the data in graphical form. All values are thought to be accurate to within one percent or better for the cases of white PM, white FM, and random walk FM. A larger tolerence should be allowed for the flicker cases.

VII. EXAMPLE OF TIME-DOMAIN SIGNAL PROCESSING AND ANALYSIS

We will analyze in some detail a commercial portable clock, Serial No. 102. This cesium was measured against another commercial cesium whose stability was well documented and verified to be better than the one under test. Plotted in figure 7.1 are the residual time deviations after removing
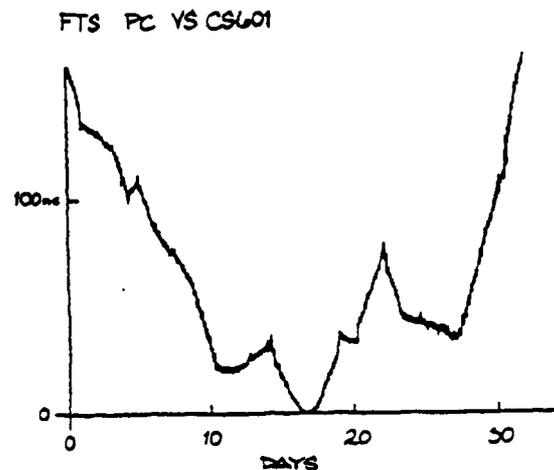
FTS PC VS CS601

FIGURE 7.1

17

a mean frequency of 4.01 parts in 10^{13}. Applying the methods described in section IV and section V, we generated the $\sigma_y(\tau)$ diagram shown in figure 7.2.
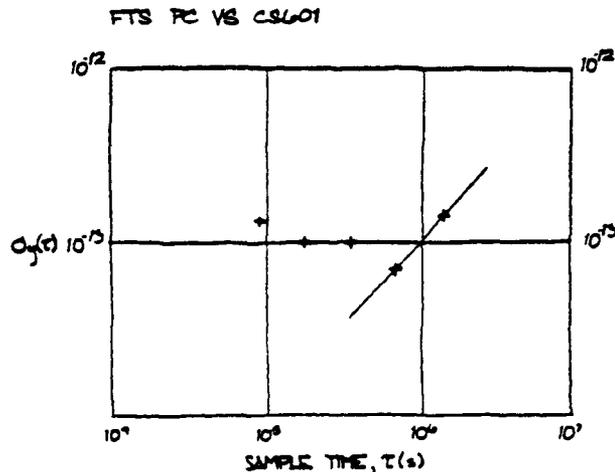
FTS PC VS CS601

FIGURE 7.3

FIGURE 7.2

One observes that the last two points are proportional to τ^{+1} and one is suspicious of a significant frequency drift.

If one calculates the drift knowing that $\sigma_y(\tau)$ is equal to the drift times $\frac{\tau}{\sqrt{2}}$ a drift of 1.22×10^{-14} per day is obtained. A linear least squares to the frequency was removed and sections IV and V were applied again. The linear least squares fit showed a drift of 1.23×10^{-14} per day, which is in excellent agreement with the previous calculated value obtained from $\sigma_y(\tau)$. Typically, the linear least squares will give a much better estimate of the linear frequency drift than will the estimate from $\sigma_y(\tau)$ being proportional to τ^{+1}.

Figure 7.3 gives the plot of the time residuals after removing the linear least squares and figure 7.4 is the corresponding $\sigma_y(\tau)$ vs. τ diagram. From the 33 days of data, we have used the 90% confidence interval to bracket the stability estimates and one sees a reasonable fit corresponding to white noise frequency modulation at a level of $4.4 \times 10^{-11} \tau^{-\frac{1}{2}}$. This seemed excessive
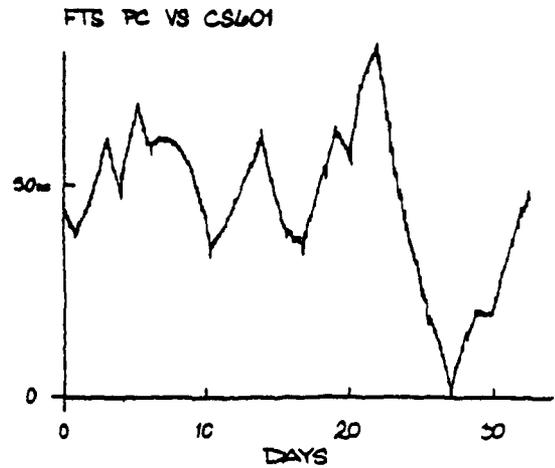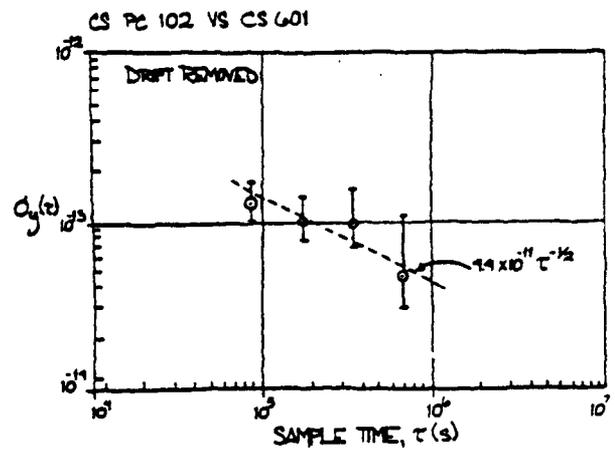
FIGURE 7.4

FIGURE 7.5

18
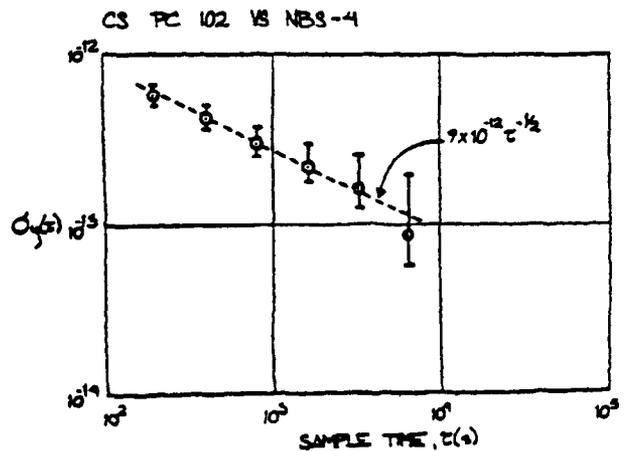
in terms of the typical performance of this particular cesium and in as much as we were doing some other testing within the environment, such as working on power supplies and charging and discharging batteries, we did some later tests. Figure 7.5 is a plot of $\sigma_y(\tau)$ after the standard had been left alone in a quiet environment and had been allowed to age for about a week. One observes that the white noise frequency modulation level is more than a factor of 4 improved over the previous data. This led us to do some studies on the effects of the power supply on the cesium frequency as one is charging and discharging batteries, which proved to be significant. One notices in figure 7.4 that the $\sigma_y(\tau)$ values plotted are consistent within the error bars with flicker noise frequency modulation. This is more typical of the kind of noise one would expect due to such environmental perturbations as discussed above.

Careful time- and/or frequency-domain analyses can lead to significant insights into problems and their solutions and is highly recommended by the authors. The frequency-domain techniques will be next approached.

VIII. SPECTRUM ANALYSIS

Another method of characterizing the noise in a signal source is by means of spectrum analysis. To understand this approach, let's examine the waveform shown in figure 8.1.
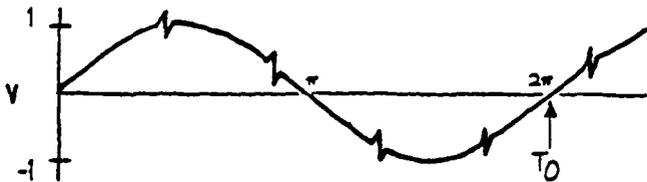
FIGURE 8.1

Here we have a sine wave which is perturbed for short instances by noise. Some loosely refer to these types of noises as "glitches". The waveform has a nominal frequency over one cycle which we'll call "ν_0" ($\nu_0 = \frac{1}{T}$). 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, the remainder is the sum of the noise components. These components are of a variety of 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 figure 8.2.
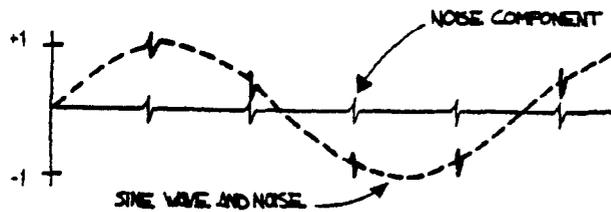

FIGURE 8.2

One can plot a graph showing rms power vs. frequency for a given signal. This kind of plot is called the power spectrum. For the waveform of figure 8.1 the power spectrum will have a high value at ν_0 and will have 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 figure 8.2. 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 which is related to the characteristics of the glitches. The power spectrum shown in figure 8.3 has this feature. A predominant ν_s component has been depicted, but other harmonics also exist.
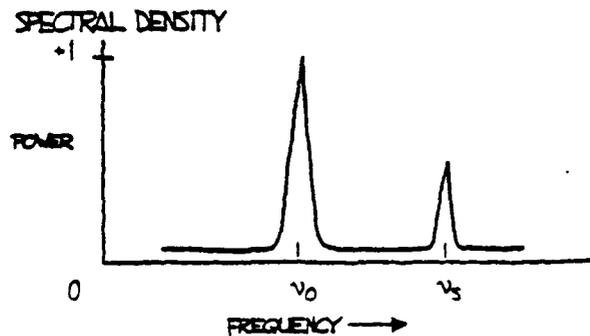

FIGURE 8.3

19

TN-32

Some noise will cause the instantaneous fre-
quency to "jitter" around ν_0, with probability of
being higher or lower than ν_0. We thus usually
find a "pedestal" associated with ν_0 as shown in
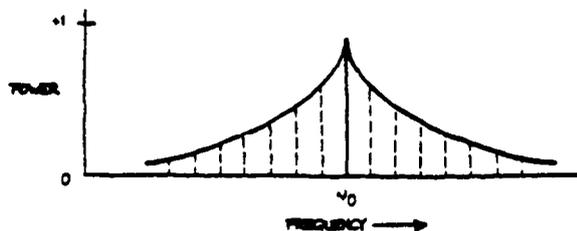figure 8.4.

SPECTRAL DENSITY

FIGURE 8.4

The process of breaking down a signal into
all of its various components of frequency is
called Fourier expansion (see sec. X). 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, often called the RF
spectrum, of V(t) is very useful in many appli-
cations. Unfortunately, if one is given the RF
spectrum, it is impossible to determine whether
the power at different Fourier frequencies is a
result of amplitude fluctuations "$\varepsilon(t)$" or phase
fluctuations "$\phi(t)$". The RF spectrum can be separ-
ated into two independent spectra, one being the
spectral density of "$\varepsilon(t)$" often called the AM
power spectral density and the other being the
spectral density of "$\phi(t)$".

For the purposes 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 fre-
quently encountered case where the AM power spec-
tral density is negligibly small and the total
modulation of the phase fluctuations is small
(mean-square value is much less than one 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 (car-
rier), 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 radians^2/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 measure-
ment technique makes use of the relation that for
small deviations ($\delta\phi \ll 1$ radian),

$$S_\phi(f) = \left(\frac{V_{rms}(f)}{V_s} \right)^2 \qquad (8.3)$$

where $V_{rms}(f)$ is the root-mean-square noise voltage
per $\sqrt{Hz}$ at a Fourier frequency "f", and V_s is the
sensitivity (volts per radian) at the phase quadra-
ture output of a phase detector which is comparing
the two oscillators. In the next section, we will
look at a scheme for directly measuring $S_\phi(f)$.

One question we might ask is, "How do fre-
quency 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 fre-
quency 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 have then

$$2\pi\nu(t) = \dot\phi_T(t) \qquad (8.4)$$

The dot denotes the mathematical operation of
differentiation on the function ϕ_T with respect to
its independent variable t.* From eq (8.4)

---

* As an analogy, the same operation relates the
  position of an object with its velocity.

20

and eq (1.1) we get

$$2\pi v(t) = \dot{\phi}_T(t) = 2\pi v_0 + \dot{\phi}(t)$$

Rearranging, we have

$$2\pi v(t) - 2\pi v_0 = \dot{\phi}(t)$$

or

$$v(t) - v_0 = \frac{\dot{\phi}(t)}{2\pi} \qquad (8.5)$$

The quantity $v(t) - v_0$ can be more conveniently denoted as $\delta v(t)$, a change in frequency at time t. Equation (8.5) tells us that if we differentiate the phase fluctuations $\phi(t)$ and divide by 2π, we will have calculated the frequency fluctuation $\delta v(t)$. Rather than specifying a frequency fluctuation in terms of shift in frequency, it is useful to denote $\delta v(t)$ with respect to the nominal frequency v_0. The quantity $\frac{\delta v(t)}{v_0}$ is called the fractional frequency fluctuation** at time t and is signified by the variable y(t). We have

$$y(t) = \frac{\delta v(t)}{v_0} = \frac{\dot{\phi}(t)}{2\pi v_0} \qquad (8.6)$$

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 in two oscillators $\delta v(t)$ 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 1/10, and in the second case is 1/10,000,000 or 1×10^{-7}. The 10 MHz oscillator is then more precise. 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

---

** Some international recommendations replace "fractional" by "normalized".

spectral density of frequency 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 $\frac{f}{v_0}$ in terms of spectral densities. With further calculation, one can derive that

$$S_y(f) = \left(\frac{f}{v_0}\right)^2 S_\phi(f) \qquad (8.7)$$

We will address ourselves primarily to $S_\phi(f)$, that is, the spectral density of phase fluctuations. For noise-measurement purposes, $S_\phi(f)$ can be measured with a straightforward, easily duplicated equipment set-up. 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 (8.7) if necessary.

8.1 The Loose Phase-Locked Loop

Section I, 1.1, C described a method of measuring phase fluctuations between two phase-locked oscillators. Now we will detail the procedure for measuring $S_\phi(f)$.

Suppose we have a noisy oscillator. We wish to measure the oscillator's phase fluctuations relative to nominal phase. One can do this by phaselocking 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 figure 8.9. The two oscillators are at the same frequency in long term as guaranteed by the phase-lock loop (PLL). A low-pass filter (to filter the R.F. 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 out of the mixer.
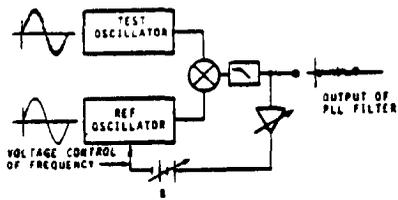
FIGURE 8.9

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 spectrum to base-band 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 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. 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 noise power is twice that which is associated with one oscillator. $S_\phi(f)$ is general notation depicting spectral density on a per hertz basis. A PLL filter output 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 with the inverse of the lowest Fourier frequency we wish to measure. That is, $\tau_c > \frac{1}{2\pi\ f(lowest)}$. This means that if we want to measure $S_\phi(f)$ down to 1 Hz, the loop time-constant must be greater than $\frac{1}{2\pi}$ seconds. One 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 preamp may be necessary before the spectrum analyzer.

The analyzer determines the mean square volts that pass through the analyzer's bandwidth centered around a pre-chosen Fourier frequency f. It is desireable 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 Hz. One may have to approximate for other noise processes. (The phase noise sideband levels will usually be indicated in rms volts-per-root-Hertz on most analyzers.)

8.2 Equipment for Frequency Domain Stability Measurements

(1) Low-noise mixer

This should be a high quality, double-balanced type, 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 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. It is best to operate near the maximum for best signal-to-noise 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.
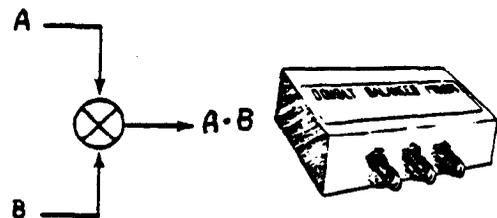

FIGURE 8.10

(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

---

* See Appendix Note # 3

** See Appendix Note # 4

TN-35

varactor control in the reference oscil-
lator. 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
which will add most significantly to the
noise being measured. If a suitable
low-noise first-stage amplifier is not
readily available, a schematic of an
amplifier with 40 dB of gain is shown in
figure 8.11 which will serve nicely for
the first stage. Amplifiers with very
low equivalent input noise performance
are also available from many manufac-
turers. 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 the determi-
nation of A_v is best made by experimen-
tation knowing that $\tau_c < \dfrac{1}{2\pi f \text{ (lowest)}}$.

LOW NOISE AMPLIFIER
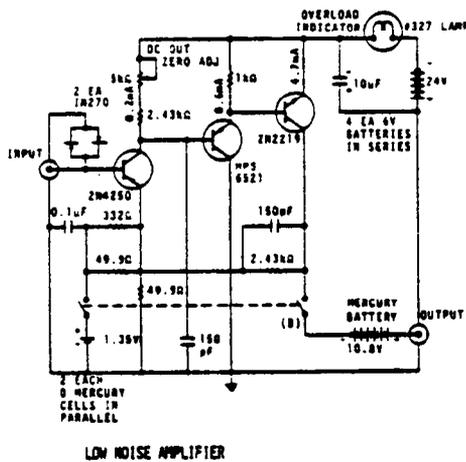
FIGURE 8.11

(3) Voltage-controlled reference quartz
    oscillator

This oscillator should be a good one
with specifications available on its
frequency domain stability. The refer-
ence should be no worse than the test
oscillator. The varactor control should

be sufficient to maintain phase-lock of
the reference. In general, low quality
test oscillators may have varactor
control of as much as 1×10^{-6} fractional
frequency change per volt. Some provi-
sion 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, one 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
that both oscillators are contributing
approximately equal noise, then one
should perform measurements on three
oscillators taking two at a time. The
noisier-than-average oscillator will
reveal itself.

(4) Spectrum analyzer

The signal analyzer typically should 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.
For voltage measuring analyzers, it is
typical to use units of "volts per $\sqrt{Hz}$".
The spectrum analyzer and any associated
input amplifier will exhibit high-fre-
quency 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 gener-
ator.

Section X describes how analysis can be
performed using a discrete fourier transform
analyser. Expanding digital technology has made
the use of fast-fourier transform analysis affor-
dable and compact.

* See Appendix Note # 5

23

TN-36

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 set up is needed. The scheme is shown in figure 8.12.
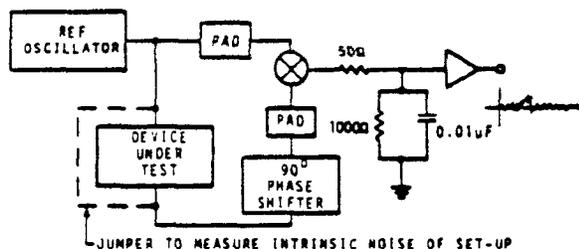
FIGURE 8.12

Figure 8.12 is a _differential_ _phase_ _noise_ _measurement_ set-up. 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 set-up, one can in principle bypass the device under test and compensate for any change in amplitude and phase at the mixer. The PLL filter technique must be converted to a differential phase noise technique in order to measure inherent test equipment noise. _It is a good practice to_ _measure the system noise before proceeding to mea-_ _surement of device noise._

A frequency domain measurement set-up is shown schematically in figure 8.13. The component
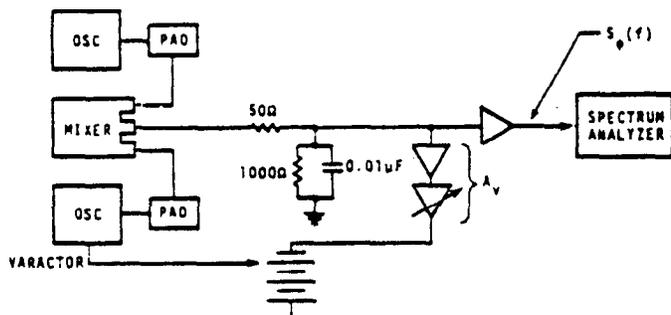
FIGURE 8.13

values for the low-pass 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 with flat 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 one or the other oscillator.[5]

8.3 Procedure and Example

At the input to the spectrum analyzer, the voltage varies as the phase fluctuations in short-term

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

V_s is the phase sensitivity of the mixer in volts per radian. Using the previously described equipment set-up, 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 volts/rad if the resultant beat note is a sine wave. 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 volts/rad sensitivity must be estimated by the slew-rate (through zero volts) of the resultant square-wave out of the mixer/amplifier.

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

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

EXAMPLE: Given a PLL with two oscillators such that, at the mixer output:

$$V_s = 1 \text{ volt/rad}$$

$V_{rms}(45 \text{ Hz}) = 100 \text{ nV per root hertz}$
solve for $S_\phi(45 \text{ Hz})$.

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

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

In decibels,

$$S_\phi(45 \text{ Hz}) = 20 \log \frac{100 \text{ nv}}{1 \text{ V}} = 20 \log \frac{10^{-7}}{10^0}$$

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

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 an $S_\phi(f)$, one should always attach ν_0. In the example above $\nu_0 = 5$ MHz, so we have

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

From eq (8.7), $S_y(f)$ can be computed as

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

$$S_y \ (45 \text{ Hz}) = 8.1 \times 10^{-25} \text{ Hz}^{-1}, \ \nu_0 = 5 \text{ MHz}.$$

IX. POWER-LAW NOISE PROCESSES

Power-law noise processes are models of precision oscillator noise that produce a particular <u>slope</u> 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 8.7 relates $S_\phi(f)$ to $S_y(f)$, the spectral density of frequency fluctuations. Translation of $S_y(f)$ to time-domain data $\sigma_y(\tau)$ for the five model noise processes is covered later in section XI.

The spectral density plot of a typical oscillator's output usually is 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 being generated from a single oscillator, but, in general, only two or three noise processes are dominant. Figure 9.1 is a graph of $S_\phi(f)$ showing the five noise processes on a log-log scale. Figure 9.2 shows the spectral density of phase fluctuations for a typical high-quality oscillator.
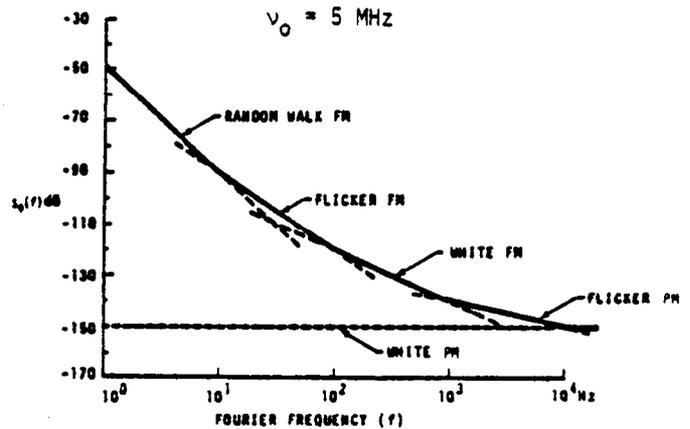
SPECTRAL DENSITY OF PHASE
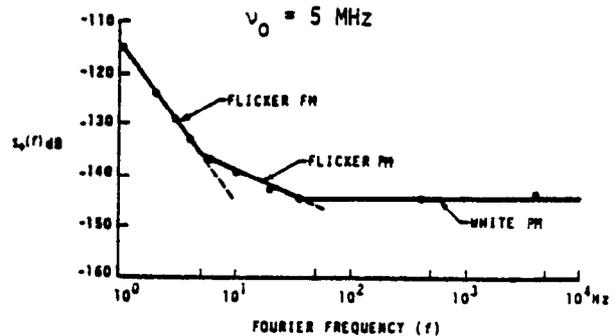

FIGURE 9.1

SPECTRAL DENSITY OF PHASE

FIGURE 9.2

25

X. PITFALLS IN DIGITIZING THE DATA

The advent and prolific use of digital computers has changed the manner in which processing of analog signals takes place if a computer is used. This section addresses the most common problems in such analyses.

10.1 Discrete-Continuous Processes

Digital processing implies that data must be presented to a computer or other processor as an array of numbers whether in a batch or in a time series. If the data are not already in this form (it usually is not when considering frequency stability measurements), then it is necessary to transform to this format by digitizing. Usually, the signal available for analysis is a voltage which varies with frequency or phase difference between two oscillators.

10.2 Digitizing the Data

Digitizing the data is the process of converting a continuous waveform into discrete numbers. The process is completed in real time using an analog-to-digital converter (ADC). Three considerations in the ADC are of importance here:

1. Conversion time
2. Resolution (quantization uncertainty)
3. Linearity

An ADC "looks at" an incoming waveform at equi-spaced intervals of time T. Ideally, the output of the ADC is the waveform (denoted by y(t)) multiplied by a series of infinitely narrow sampling intervals of unit height as in figure 10.1. We have at t = T

$$y_1(t) = y(t)\delta(t-T) = y(T)\delta(t-T) \qquad (10.1)$$

where $\delta(t-T)$ is a delta function. If y(t) is continuous at t = nT and n = 0, ±1, ±2,..., then

$$y_\ell(t) = \sum_{n=-\infty}^{\infty} y(nT)\ \delta\ (t-nT)$$

$$\qquad (10.2)$$

ℓ = integer

The delta function respresentation of a sampled waveform eq (10.2) is useful when a subsequent continuous integration is performed using it.[6]
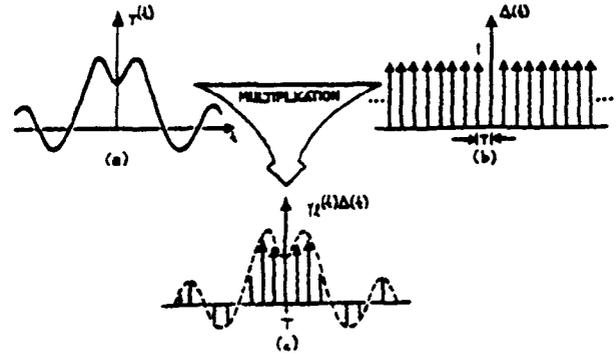

FIGURE 10.1

In ADC's, the input signal is sampled during an aperture time and held for conversion to a digital number, usually binary. Sampling and processing takes time which is specified as the conversion time. This is the total time required for a complete measurement at one sample to achieve a given level of accuracy. If $y_\ell(t)$ is the ideal discrete-time representation of continuous process y(t), then the ADC output denoted by $y'_\ell(t)$ is:

$$y'_\ell = y_{\ell-d} \pm \varepsilon(\frac{dy}{dt}) \qquad (10.3)$$

where "d" is the conversion time and $\varepsilon \cdot \frac{dy}{dt}$ the accuracy tolerance at "d" as a function of rate-of-change in y(t). In general a trade-off exists between d and ε. For example, for a commonly available, high-quality 10-bit ADC, a conversion time of d = 10 μs yields a maximum error of 3%. Whereas given a 30 μs conversion time, we can obtain 0.1% maximum error.

The error due to conversion time "d" is many times negligible since processing in digital filters and spectrum analysis takes place after the converter. Conversion time delay can be of critical concern, however, where real-time processing at speeds of the order of "d" become important

26

such as in digital servo loops where corrections are needed for fast changing errors.

A portion of the conversion-time error is a function of the rate of change $\frac{dy}{dt}$ of the process if the sample-and-hold portion of the ADC relies on the charging of a capacitor during an aperture time. This is true because the charge cycle will have a finite time-constant and because of aperture time uncertainty. For example, if the time-constant is 0.1 ns (given by say a 0.1 Ω source resistance charging a 0.001 µfd capacitor), then a 0.1% nominal error will exist for slope $\frac{\Delta y}{\Delta t}$ = 1V/µs due to charging. With good design, this error can be reduced. The sampling circuit (before charge) is usually the dominant source of error and logic gate-delay jitter creates an aperture time uncertainty. The jitter typically is between 2-5 ns which means an applied signal slewing at, say, 1 V/µs produces an uncertainty of 2-5 mV. Since $\varepsilon \cdot \frac{dy}{dt}$ is directly proportional to signal slewing rate, it can be anticipated that high-level, high-frequency components of y(t) will have the greatest error in conversion. For typical ADC's, less than 0.1% error can be achieved by holding $\frac{\Delta y}{\Delta t}$ to less than 0.2 V/µs.

The continuous process y(t) is partitioned into 2^n discrete ranges for n-bit conversion. All analog values within a given range are represented by the same digital code, usually assigned to the nominal midrange value. There is, therefore, an inherent quantization uncertainty of ± ½ least-significant bit (LSB), in addition to other conversion errors. For example, a 10-bit ADC has a total of 1024 discrete ranges with a lowest order bit then representing about 0.1% of full scale and quantization uncertainty of ± 0.05%.

We define the dynamic range of a digital system as the ratio between the maximum allowable value of the process (prior to any overflow condition) and the minimum discernable value. The dynamic range when digitizing the data is set by the quantizing uncertainty, or resolution, and is the ratio of 2^n to ½ LSB. (If additive noise makes coding ambiguous to the ½ LSB level, then the dynamic range is the ratio of 2^n to the noise uncertainty, but this is usually not the case.)

For example, the dynamic range of a 10-bit system is 2^n = 1024 to ½, or 2048 to 1. Expressed in dB's, this is

$$20 \log 2048 = 66.2 \text{ dB}$$

if referring to a voltage-to-code converter.

The converter linearity specifies the degree to which the voltage-to-code transfer approximates a straight line. The nonlinearity is the deviation from a straight line drawn between the end points (all zeros to all ones code). It is usually not acceptable to have nonlinearity greater than ½ LSB which means that the sum of the positive errors or the sum of the negative errors of the individual bits must not exceed ½ LSB (or ± ½ LSB). The linearity specification used in this context includes all effects such as temperature errors under expected operating temperature extremes and power supply sensitivity errors under expected operating supply variations.

10.3 Aliasing

Figure 10.1 illustrates equispaced sampling of continuous process y(t). It is important to have a sufficient number of samples/second to properly describe information in the high frequencies. On the other hand, sampling at too high a rate may unnecessarily increase the processing labor. As we reduce the rate, we see that sample values could represent low or high frequencies in y(t). This property is called aliasing and constitutes a source of error similiar to "imaging" which occurs in analog frequency mixing schemes (i.e., in the multiplication of two different signals).

If the time between samples (k) is T seconds, then the sampling rate is $\frac{1}{T}$ samples per second. Then useful data in y(t) will be from 0 to $\frac{1}{2T}$ Hz and frequencies higher than $\frac{1}{2T}$ Hz will be folded into the lower range from 0 to $\frac{1}{2T}$ Hz and confused with data in this lower range. The cutoff frequency is then given by

$$f_s = \frac{1}{2T} \qquad (10.4)$$

and is sometimes called the "Nyquist frequency."

27

We can use the convolution theorem to simply illustrate the existence of aliases. This theorem states that multiplication in the time domain corresponds to convolution in the frequency domain, and the time domain and frequency domain representations are Fourier transform pairs.[7] The Fourier transform of y(t) in figure 10.1(a) is denoted by Y(f); thus:

$$Y(f) = \int_{-\infty}^{\infty} y(t)e^{-j2\pi ft}dt \qquad (10.5)$$

and

$$y(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} Y(f)e^{j2\pi ft}df \qquad (10.6)$$

The function Y(f) is depicted in figure 10.2(a). The Fourier transform of $\Delta(t)$ is shown in figure 10.2(b) and is given by $\Delta(f)$ where applying the discrete transform yields:

$$\Delta(f) = \frac{1}{T} \sum_{n=-\infty}^{\infty} \delta(f - \frac{n}{T}), \qquad (10.7)$$

recalling that

$$\Delta(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT), \qquad (10.8)$$

from eq (10.2).
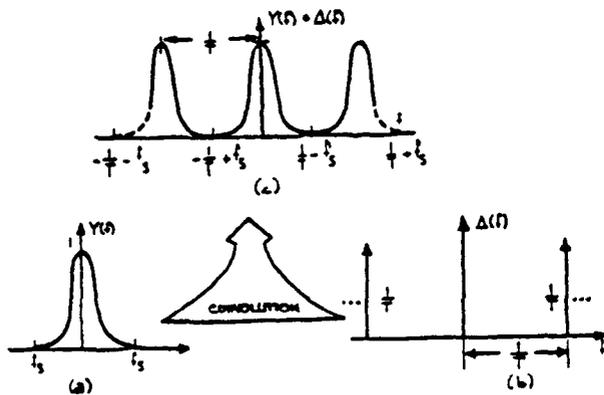
FIGURE 10.2

Y(f) convolved with Δf is denoted by Y(f)* $\Delta(f)$ and is shown in figure 10.2(c). We see that the transform Y(f) is repeated with origins at $f = \frac{n}{T}$. Conversely, high frequency data with information around $f = \frac{n}{T}$ will fold into the data around the origin between $-f_s$ and $+f_s$. In the computation of power spectra, we encounter errors as shown in figure (10.3).
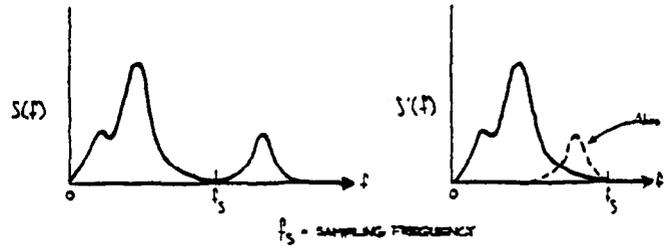
FIGURE 10.3

Aliased power spectra due to folding. (a) True Spectra, (b) Aliased Spectra.

Two pioneers in information theory, Harold Nyquist and Claude Shannon, developed design criteria for discrete-continuous processing systems. Given a specified accuracy, we can convey time-domain process y(t) through a finite bandwidth whose upper limit f_N is the highest significant spectral component of y(t). For discrete-continuous process $y_k(t)$, ideally the input signal spectrum should not extend beyond f_s, or

$$f_N \leq f_s \qquad (10.9)$$

where f_s is given by eq (10.4). Equation (10.9) is refered to as the "Shannon limit."

In practice, there is never a case in which there is absolutely no signal or noise component above f_N. Filters are used before the ADC in order to suppress components above f_N which fold into the lower bandwidth of interest. This so-called anti-aliasing filter usually must be quite sophisticated in order to have low ripple in the passband, constant phase delay in the passband, and steep rolloff characteristics. In examining the rolloff requirements of the anti-aliasing filter, we can apply a fundamental filter property that the output spectrum is equal to the input

28

TN-41

spectrum multiplied by the square of the frequency response function; that is,

$$S(f) [H(f)]^2 = S(f)_{out} \qquad (10.10)$$

The filter response must be flat to f_N and attenuate aliased noise components at $\frac{n}{T} \pm f = 2nf_s \pm f$. In digitizing the data, the observed spectra will be the sum of the baseband spectrum (to f_N) and all spectra which are folded into the baseband spectrum

$$S(f)_{observed} = S_0(f) + S_{-1} (2f_s - f) + S_{+1} (2f_s + f)$$
$$+ S_{-2} (4f_s - f).... + S_i (2 i (f_s + \frac{1}{i} f))$$
$$\qquad (10.11)$$
$$= S_0(f) + \sum_{i=-M}^{M} S_i (2 i (f_s + f))$$

where M is an appropriate finite limit.

For a given rejection at an upper frequency, clearly the cutoff frequency f_c for the anti-aliasing filter should be as low as possible to relax the rolloff requirements. Recall that an nth order low-pass filter has frequency response function

$$H(f) = \frac{1}{1 + j\left(\frac{f}{f_c}\right)^n} \qquad (10.12)$$

and output spectrum

$$S(f)_{out} = \frac{S(f)}{1 + \left(\frac{f}{f_c}\right)^{2n}} \qquad (10.13)$$

and after sampling, we have (applying eq (10.11))

$$S(f)_{observed} = \frac{S_0(f)}{1 + \left(\frac{f}{f_c}\right)^{2n}} + \sum_{i=-M}^{M} \frac{S_i(2 i (f_s + \frac{1}{i}f))}{1 + \frac{2 i (f_s + \frac{1}{i}f)}{f_c}^{2n}}$$
$$\qquad (10.14)$$

If f_c is chosen to be higher than f_N, then the first term (baseband spectrum) is negligibly affected by the filter, which is our hope. It is

the second term (the sum of the folded in spectra) which causes an error.

As an example of the rolloff requirement, consider the measurement of noise process n(t) at f = 400 Hz in a 1 Hz bandwidth on a digital spectrum analyzer. Suppose n(t) is white; that is,

$$S_n(f) = k_0 \qquad (10.15)$$

$$k_0 = constant$$

Suppose further that we wish to only measure the noise from 10 Hz to 1 kHz; thus f_N = 1 kHz. Let us assume a sampling frequency of f_s = $2f_N$ or 2 kHz. If we impose a 1 dB error limit in $S_{observed}$ and have 60 dB of dynamic range, then we can tolerate an error limit of 10^{-6} due to aliasing effects in this measurement, and the second term in eq (10.14) must be reduced to this level. We can choose f_c = 1.5kHz and obtain

$$S(f)_{observed} = k_0 + \sum_{i=-M}^{M} \frac{k_0}{1 + \frac{2 i (f_s + \frac{1}{i}f)}{f_c}^{2n}}$$
$$\qquad (10.16)$$

The term in the series which contributes most is at i = -1, the nearest fold-in. The denominator must be 10^6 or more to realize the allowable error limit and at n $\geq$ 8 this condition is met. The next most contributing term is i = +1 at which the error is < 10^{-7} for n = 8, a negligible contribution. The error increases as f increases for a fixed n because the nearest fold-in (i = -1) is coming down in frequency (note fig. 10.2(c)) and power there is filtered less by the anti-aliasing filter. Let us look at the worst case (f = 1kHz) to determine a design criteria for this example. At f = 1 kHz, we must have n $\geq$ 10.

Thus the requirement in this example is for a 10-pole low-pass filter (60 dB/octave rolloff).

10.4 Some History of Spectrum Analysis Leading to the Fast Fourier Transform

Newton in his Principia (1687) documented the first mathematical treatment of wave motion al-

though the concept of harmonics in nature was pointed out by Pythagoras, Kepler, and Galileo. However, it was the work of Joseph Fourier in 1807 which showed that almost any function of a real variable could be represented as the sum of sines and cosines. The theory was rigorously treated in a document in 1822.

In using Fourier's technique, the periodic nature of a process or signal is analyzed. Fourier analysis assumes we can apply fixed amplitudes, frequencies, and phases to the signal.

In the early 1900's two relatively independent developments took place: (1) radio electronics and electric power hardware were fast growing technologies; and (2) statistical analysis of events or processes which were not periodic became increasingly understood. The radio engineer explored signal and noise properties of a voltage or current into a load by means of the spectrum analyzer and measurement of the power spectrum. On the other hand, statisticians explored deterministic and stochastic properties of a process by means of the variance and self-correlation properties of the process at different times. Wiener (1930) showed that the variance spectrum (i.e., the breakdown of the variance with Fourier frequency) was the Fourier transform of the autocorrelation function of the process. He also theorized that the variance spectrum was the same as the power spectrum normalized to unit area. Tukey (1949) advocated the use of the variance spectrum in the statistical treatment of all processes because (1) it is more easily interpreted than correlation-type functions; and (2) it fortuitously is readily measureable by the radio engineer.

The 1950's saw rigorous application of statistics to communication theory. Parallel to this was the rapid advancement of digital computer hardware. Blackman and Tukey (1959) and Welch (1961) elaborated on other useful methods of deriving an estimate for the variance spectrum by taking the ensemble time-average sampled, discrete line spectra. The approach assumes the random process is ergodic. Some digital approaches estimate the variance spectrum using Wiener's theorem if correlation-type functions are useful

in the analysis, but in general the time-averaged, sample spectrum is the approach taken since its implementation is direct and straightforward. Most always, ergodicity can be assumed.[8,9]

The variance of process $y(t)$ is related to the total power spectrum by

$$\sigma^2[y(t)] = \int_{-\infty}^{\infty} S_y(f) \, df. \qquad (10.17)$$

Since

$$\sigma^2[y(t)] = \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} y^2(t) \, dt \qquad (10.18)$$

we see that if $y(t)$ is a voltage or current into a 1-ohm load, then the mean power of $y(t)$ is the integral of $S_y(f)$ with respect to frequency over the entire range of frequencies $(-\infty, \infty)$. $S_y(f)$ is, therefore, the power spectrum of process $y(t)$. The power spectrum curve shows how the variance is distributed with frequency and should be expressed in units of watts per unit of frequency, or volts squared per unit of frequency when the load is not considered.

Direct estimation of power spectra has been carried out for many years through the use of analog instruments. These have variously been referred to as sweep spectrum analyzers, harmonic analyzers, filter banks, and wave analyzers. These devices make use of the fact that the spectrum of the output of a linear system (analog filter) is the spectrum of the input multiplied by the square of the system's frequency response function (real part of the transfer characteristic). Note eq (10.10). If $y(t)$ has spectrum $S_y(f)$ feeding a filter with frequency response function $H(f)$, then its output is

$$S(f)_{filtered} = [H(f)]^2 \, S_y(f) \qquad (10.19)$$

If $H(f)$ is rectangular in shape with width Δf, then we can measure the contribution to the total power spectrum due to $S_y(f \pm \frac{\Delta f}{2})$.

The development of the fast Fourier transform (FFT) in 1965 made digital methods of spectrum

TN-43

estimation increasingly attractive. Today the choice between digital or analog methods depends more on the objectives of the analysis rather than on technical limitations. However, many aspects of digital spectrum analysis are not well known by the casual user in the laboratory while the analog analysis methods and their limitations are understood to a greater extent.

Digital spectrum analysis is realized using the discrete Fourier transform (DFT), a modified version of the continuous transform depicted in eqs (10.5) and (10.6). By sampling the input waveform $y(t)$ at discrete intervals of time $t_\ell = \Delta t$ representing the sampled waveform by eq (10.2) and integrating eq (10.5) yields

$$Y(f) = \sum_{\ell=-\infty}^{\infty} y(\ell t) e^{-j2\pi f \ell T} \qquad (10.20)$$

Equation (10.20) is a Fourier series expansion. Because $f(t)$ is specified as being bandlimited, the Fourier transform as calculated by eq (10.20) is as accurate as eq (10.5); however, it cannot extend beyond the Nyquist frequency, eq (10.4).

In practice we cannot compute the Fourier transform to an infinite extent, and we are restricted to some observation time T consisting of $n\Delta t$ intervals. This produces a spectrum which is not continuous in f but rather is computed with resolution Δf where

$$\Delta f = \frac{1}{n\Delta t} = \frac{1}{T} \qquad (10.21)$$

With this change, we get the discrete finite transform

$$Y(m\Delta f) = \sum_{n=0}^{N-1} y_\ell(t) e^{-j2\pi m\Delta f n t} \qquad (10.22)$$

The DFT computes a sampled Fourier series, and eq (10.22) assumes that the function $y(t)$ repeats itself with period T. $Y(m\Delta f)$ is called the "line spectrum." A comparison of the DFT with the continuous Fourier transform is shown later in part 10.7.

The fast Fourier transform (FFT) is an algorithm which efficiently computes the line spectrum by reducing the number of adds and multiplies involved in eq (10.22). If we choose $T/\Delta t$ to equal a rational power of 2, then a symmetric matrix can be derived through which $y_\ell(t)$ passes and quickly yields $Y(m\Delta f)$. An N-point transformation by the direct method requires a processing time proportional to N^2 whereas the FFT requires a time proportional to $N \log_2 N$. The approximate ratio of FFT to direct computing time is given by

$$\frac{N \log_2 N}{N^2} = \frac{\log_2 N}{N} = \frac{Y}{N} \qquad (10.23)$$

where $N = 2^Y$. For example, if $N = 2^{10}$, the FFT requires less than 1/100 of the normal processing time.

We must calculate both the magnitude and phase of a frequency in the line spectrum, i.e., the real and imaginary part at the given frequency. N points in the time domain allow N/2 complex quantities in the frequency domain.

The power spectrum of $y(t)$ is computed by squaring the real and imaginary components, adding the two together and dividing by the total time T. We have

$$S_y(m\Delta f) = \frac{R[Y(m\Delta f)]^2 + I[Y(m\Delta f)]^2}{T} \qquad (10.24)$$

This quantity is the sampled power spectrum and again assumes periodicity in process $y(t)$ with total period T.[10]

10.5 Leakage

Sampled digital spectrum analysis always involves transforming a finite block of data. Continuous process $y(t)$ is "looked at" for T time through a data window which can functionally be described by

$$y'(t) = w(t) \cdot y(t) \qquad (10.25)$$

where $w(t)$ is the time domain window. The time-discrete counterpart to eq (10.25) is

$$y_\ell'(t) = w_\ell(t) \cdot y_\ell(t) \qquad (10.26)$$

31

and $w_{\ell}(t)$ is now the sampled version of $w(t)$ derived similarly to eq (10.2). Equation (10.26) is equivalent to convolution in the frequency domain, or

$$Y'(m\Delta f) = W(m\Delta f)*Y(m\Delta f) \qquad (10.27)$$

$Y'(m\Delta f)$ is called the "modified" line spectrum due to convolution of the original line spectrum with the Fourier transform of the time-domain window function.

Suppose the window function is rectangular, and

$$w_{\ell}(t) = 1, \qquad \frac{-T}{2} \le t < \frac{T}{2}$$
$$= 0, \qquad t > \left|\frac{T}{2}\right| \qquad (10.28)$$

This window is shown in figure 10.4(a). The Fourier transform of this window is

$$W(m\Delta f) = T\frac{\sin\pi\ m\Delta f\ NT}{\pi m\Delta f NT} \qquad (10.29)$$

and is shown in figure 10.4(b). If $y(t)$ is a sine wave, we convolve the spectrum of the sinusoid, a delta function, with $W(m\Delta f)$.

(a)

(b)
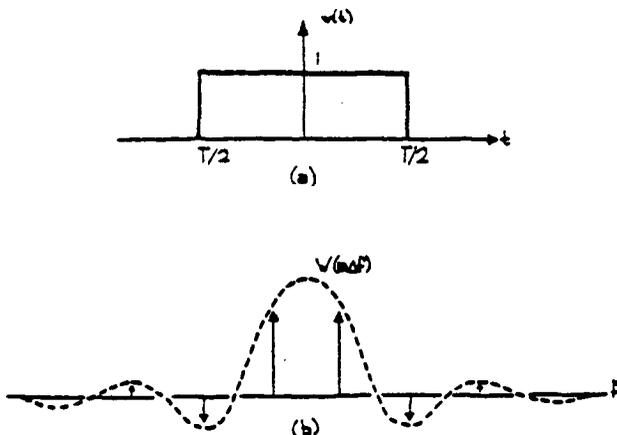
FIGURE 10.4

The transform process (eq 10.22) treats the sample signal as if it were periodically extended. Discontinuities usually occur at the ends of the window function in the extended version of the sampled waveform as in figure 10.5(c). Sample spectra thus represent a periodically extended sampled waveform, complete with discontinuites at its ends, rather than the original waveform.
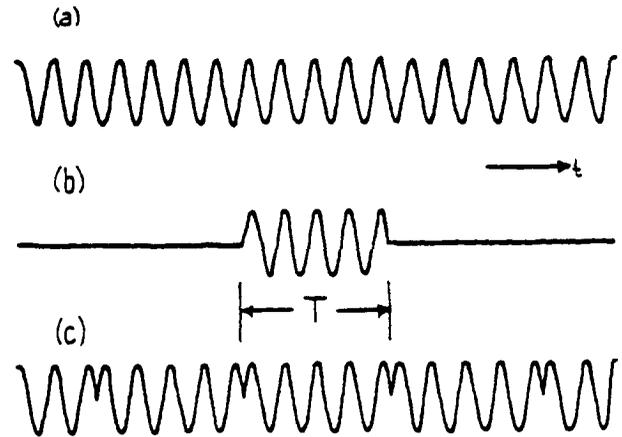
FIGURE 10.5

Spurious components appear near the sinusoid spectrum and this is referred to as "leakage." Leakage results from discontinuites in the periodically extended sample waveform.

Leakage cannot be eliminated entirely, but one can choose an appropriate window function $w(t)$ in order to minimize its effect. This is usually done at the expense of resolution in the frequency domain. An optimum window for most cases is the Hanning window given by:

$$w(t) = [\tfrac{1}{2} - \tfrac{1}{2} \cos (\tfrac{2\pi t}{T})]^a \qquad (10.30)$$

for $0 \le t \le T$ and "a" designates the number of times the window is implemented. Figure 10.6(a) shows the window function and 10.6(b) shows the Hanning line shape in the frequency domain for various numbers of "Hanns." Note that this window

32

$$-w(t) \cdot \left[\frac{1}{2} - \frac{1}{2} \cos \frac{2\pi t}{T} \right]$$

(a)

HANNING LINE SHAPES

3 IS FREQUENCY NORMALIZED BY
1/T WHERE NT IS WINDOW WIDTH
CURVE A SIN πSΔt (RECT WINDOW)
B HANNING ONCE
C HANNING TWICE
D HANNING 3 TIMES
E HANNING 4 TIMES

(b)
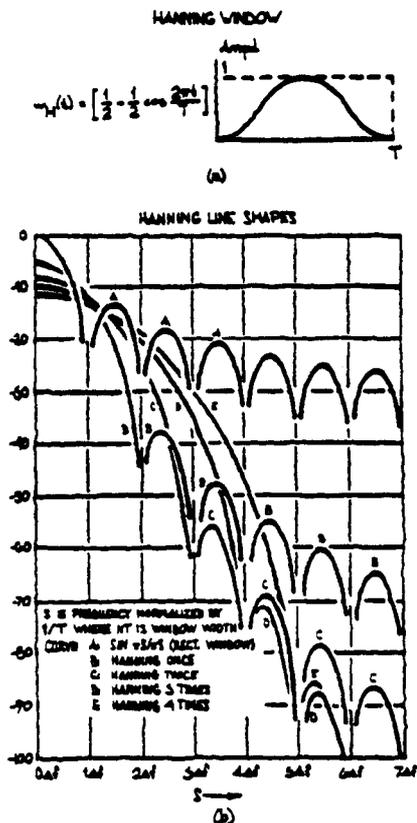
FIGURE 10.6

eliminates discontinuities due to the ends of sample length T.

Each time the Hanning window is applied, the sidelobes in the transform are attenuated by 12 dB/octave, and the main lobe is widened by $2\Delta f$. The amplitude uncertainty of an arbitrary sine wave input is reduced as we increase the number of Hanns; however, we trade off resolution in frequency.

The effective noise bandwidth indicates the departure of the filter response away from a true rectangularly shaped filtered response (frequency domain). Table 10-I lists equivalent noise bandwidth corrections for up to three applications of the Hanning window.[11]

Number of Hanns	Equivalent Noise Bandwidth
1	1.5 Δf
2	1.92 Δf
3	2.31 Δf

TABLE 10.I

10.6 Picket-Fence Effect

The effect of leakage discussed in the previous section gives rise to a sidelobe type response that can be tailored according to the time-window function through which the analyzed signal passes as a block to be transformed to the frequency domain. Using the Hanning window diminishes the amplitudes of the sidelobes, however, it increases the effective bandwidth of the passband around the center frequency. This is because the effective time-domain window length is shorter than a perfect rectangular window. Directly related to the leakage (or sidelobe) effect is one called the "picket-fence" effect. This is because the sidelobes themselves resemble a frequency response which has geometry much like a picket fence.

The existence of both sidelobe leakage and the resultant picket-fence effect are an artifact of the way in which the FFT analysis is performed. Frequency-domain analysis using analog filters involves a continuous signal in and a continuous signal out. On the other hand, FFT analysis involves a continuous signal in, but the transform to the frequency domain is performed on blocks of data. In order to get discrete frequency information from a block, the assumption is made that the block represents one period of a periodic signal. The picket-fence effect is a direct consequence of this assumption. For example, consider a sinewave signal which is transformed from a time-varying voltage to a frequency-domain representation through an FFT. The block of data to be transformed will be length, T, in time. Let's say that the block, T, represents only $4\frac{1}{2}$ cycles of the input sinewave as in figure 10.5. Artificial sidebands will be created in the transform to the frequency domain, whose frequency spacing equals $\frac{1}{T}$, or the reciprocal of the block length. This represents a worst-case condition for sidelobe generation and creates a large number of spurious discrete frequency components as shown in figure 10.7(b).

If, on the other hand, one changes the block time, T, so that the representation is an integral number of cycles of the input sinewave, then the transform will not contain sidelobe leakage compo-
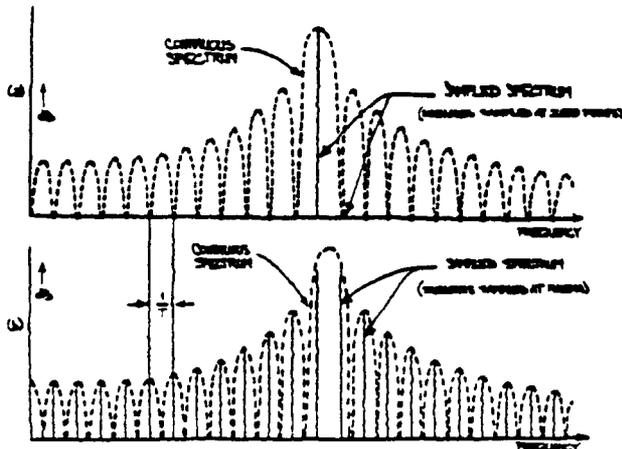
33

FIGURE 10.7

nents and the artificial sideband frequency compo-
nents disappear. In practice, when looking at
complex input signals, the block time, T, is not
synchronous with any component of the transformed
part of the signal. As a result, discrete fre-
quency components in the frequency domain have
associated with them sidebands which come and go
depending on the phase of the time window, T,
relative to the sine components of the incoming
signal. The effect is much like looking through a
picket fence at the sidebands.[12]

An analogy to the sidelobe leakage and picket-
fence effect is to record the incoming time-varying
signal on a tape loop, which has a length of time,
T. The loop of tape then repeats itself with a
period of T. This repeating signal is then coupled
to a scanning or filter-type spectrum analyzer.
The phase discontinuity between the end of one
passage of the loop and the beginning of the loop
on itself again represents a phase-modulation
component, which gives rise to artificial sidebands
in the spectrum analysis. A word of caution —
this is not what actually happens in a FFT analyzer
(i.e., there is no recirculating memory). However,
the Fourier transform treats the incoming block as
if this were happening.

10.7 Time Domain-Frequency Domain Transforms

A. Integral transform

Figure 10.8 shows the well-known integral
transform, which transforms a continuous time-
domain signal extending over all time into a
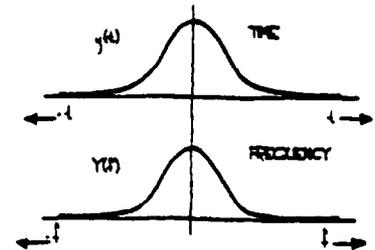
FIGURE 10.8

continuous frequency-domain signal extending over
all frequency. This is the ideal transform. In
practice, however, one deals with finite times and
bandwidths. The integral transform then, at best,
is an estimate of the transform and is so for only
short, well-behaved signals. That is, the signal
goes to zero at infinite time and at infinitely
high frequency.

B. Fourier series

The Fourier-series transform assumes perio-
dicity in the time-domain signal for all time.
Only one period of the signal (for time T) is
required for this kind of transform. The Fourier
series treats the incoming signal as periodic with
period, T, and continuous. The transformed spec-
trum is then discrete with infinite harmonic
components with frequency spacing of $\frac{1}{T}$. This is
shown in figure 10.9.

FIGURE 10.9

34

SAMPLED FUNCTIONS

$$y(n) = \sum_{n=-\infty}^{\infty} y(k)_n e^{j2\pi k}\, dk$$

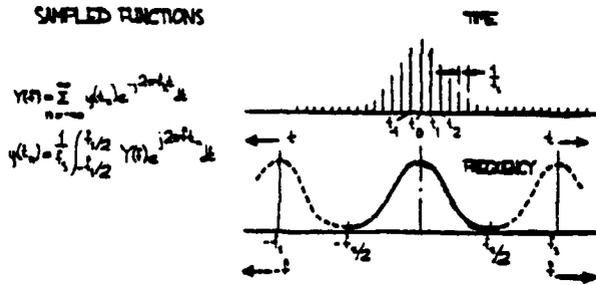$$y(k)_n = \int_{-1/2}^{1/2} y(n)_k e^{j2\pi k}\, dk$$

FIGURE 10.10

Figure 10.10 shows the transform from a sampled time-domain signal to the frequency domain. Note that in the frequency domain, the result is repetitive in frequency. This effect, commonly called aliasing, is discussed earlier in section 10.3. Figure 10.9 and figure 10.10 show the symmetry between the time- and frequency-domain transforms of discrete lines.

C. Discrete fourier transform

Finally, we have the sampled, periodic (assumed) time-domain signal which is transformed to a discrete and repetitive (aliased) frequency-domain representation. This is shown in figure 10.11

DISCRETE FOURIER TRANSFORM

$$Y(t_k) = \frac{1}{N} \sum_{n=0}^{N-1} y(t_n)_k e^{-j\frac{2\pi nk}{N}}$$

$$y(t_n) = \sum_{k=0}^{N-1} Y(t_k)_n e^{j\frac{2\pi nk}{N}}$$

FIGURE 10.11

XI. TRANSLATION FROM FREQUENCY DOMAIN STABILITY MEASUREMENT TO TIME DOMAIN STABILITY MEASUREMENT AND VICE-VERSA.
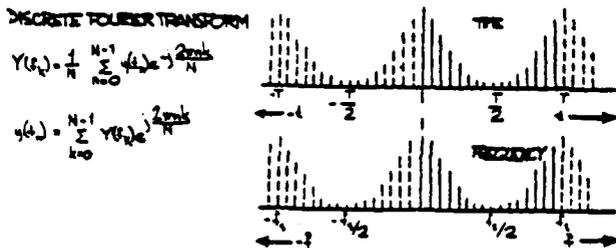
11.1 Procedure

Knowing how to measure $S_\phi(f)$ or $S_y(f)$ for a pair of oscillators, let us see how to translate the power-law noise process to a plot of $\sigma_y^2(\tau)$. First, convert the spectrum data to $S_y(f)$, the spectral density of frequency fluctuations (see sections III and VIII). There are two quantities which completely specify $S_y(f)$ for a particular power-law noise process: (1) the slope on a log-log plot for a given range of f and (2) the amplitude. The slope we shall denote by "α"; therefore f^α is the straight line (on log-log scale) which relates $S_y(f)$ to f. The amplitude will be denoted "h_α"; it is simply the coefficient of f for a range of f. When we examine a plot of spectral density of frequency fluctuations, we are looking at a representation of the addition of all the power-law processes (see sec. IX). We have

$$S_y(f) = \sum_{\alpha=-\infty}^{\infty} h_\alpha\, f^\alpha \qquad (11.1)$$

In section IX, five power-law noise processes were outlined with respect to $S_\phi(f)$. These five are the common ones encountered with precision oscillators. Equation (8.7) relates these noise processes to $S_y(f)$. One obtains

		with respect to $S_y(f)$
1.	Random Walk FM	(f^{-2}) . . . $\alpha = -2$
2.	Flicker FM	(f^{-1}) . . . $\alpha = -1$
3.	White FM	(1) . . . $\alpha = 0$
4.	Flicker ϕM	(f) . . . $\alpha = 1$
5.	White ϕM	(f^2) . . . $\alpha = 2$

slope on log-log paper

Table 11.1 is a list of coefficients for translation from $\sigma_y^2(\tau)$ to $S_y(f)$ and from $S_\phi(f)$ to

35

$\sigma_y{}^2(\tau)$. In the table, the left column is the designator for the power-law process. Using the middle column, we can solve for the value of $S_y(f)$ by computing the coefficient "a" and using the measured time domain data $\sigma_y{}^2(\tau)$. The rightmost column yields a solution for $\sigma_y{}^2(\tau)$ given frequency domain data $S_\phi(f)$ and a calculation of the appropriate "b" coefficient.

In the phase spectral density plot of figure 11.1, there are two power-law noise processes for oscillators being compared at 1 MHz. For region 1, we see that when f increases by one decade (that is, from 10 Hz to 100 Hz), $S_\phi(f)$ goes down by three decades (that is, from 10^{-11} to 10^{-14}). Thus, $S_\phi(f)$ goes as $1/f^3 = f^{-3}$. For region 1, we

$S_y(f) = H_\alpha f^\alpha$ $\alpha =$	$S_y(f) = a\,\sigma_y^2(\tau)$ $a =$	$\sigma_y^2(\tau) = b\,S_\phi(f)$ $b =$
2 (white phase)	$\dfrac{(2\pi)^2\,\tau^2\,f^2}{3\,f_h}$	$\dfrac{3\,f_h}{(2\pi)^2\,\tau^2\,\nu_0^2}$
1 (flicker noise)	$\dfrac{(2\pi)^2\,\tau^2\,f}{1.038 + 3\,\ln(\omega_h\tau)}$	$\dfrac{[1.038 + 3\,\ln(\omega_h\tau)]f}{(2\pi)^2\,\tau^2\,\nu_0^2}$
0 (white frequency)	$2\,\tau$	$\dfrac{f^2}{2\,\tau\,\nu_0^2}$
-1 (flicker frequency)	$\dfrac{1}{2\,\ln(2)\cdot f}$	$\dfrac{2\,\ln(2)\cdot f^3}{\nu_0^2}$
-2 (random walk frequency)	$\dfrac{6}{(2\pi)^2\,\tau\,f^2}$	$\dfrac{(2\pi)^2\,\tau\,f^4}{6\,\nu_0^2}$

TABLE 11.1
Conversion table from time domain to frequency domain and from frequency domain to time domain for common kinds of interger power law spectral densities; $f_h (= \omega_h/2\pi)$ is the measurement system bandwidth. Measurement reponse should be within 3 dB from D.C. to f_h (3 dB down high-frequency cutoff is at f_h).

$$S_\phi(f) = \frac{\nu_0^2}{f^2}\,S_y(f)$$

SPECTRAL DENSITY OF PHASE

FIGURE 11.1

identify this noise process as flicker FM. The rightmost column of table 11.1 relates $\sigma_y^2(\tau)$ to $S_\phi(f)$. The row designating flicker frequency noise yields:

$$\sigma_y^2(\tau) = \frac{2\ln(2)\cdot f^3}{\nu_0^2}\,S_\phi(f)$$

One can pick (arbitrarily) a convenient Fourier frequency f and determine the corresponding values of $S_\phi(f)$ given by the plot of figure 11.1. Say, $f = 10$, thus $S_\phi(10) = 10^{-11}$. Solving for $\sigma_y^2(\tau)$, given $\nu_0 = 1$ MHz, we obtain:

$$\sigma_y^2(\tau) = 1.39 \times 10^{-20}$$

therefore, $\sigma_y(\tau) = 1.18 \times 10^{-10}$. For region 2, we have white PM. The relationship between $\sigma_y^2(\tau)$ and

$S_\phi(f)$ for white PM is:

$$\sigma_y^2(\tau) = \frac{3f_h}{(2\pi)^2 \, \tau^2 \, \nu_0^2} S_\phi(f)$$

Again, we choose a Fourier frequency, say $f = 100$, and see that $S_\phi(100) = 10^{-14}$. Assuming $f_h = 10^4$ Hz, we thus obtain:

$$\sigma_y^2(\tau) = 7.59 \times 10^{-24} \cdot \frac{1}{\tau^2}$$

therefore,

$$\sigma_y(\tau) = 2.76 \times 10^{-12} \cdot \frac{1}{\tau}.$$

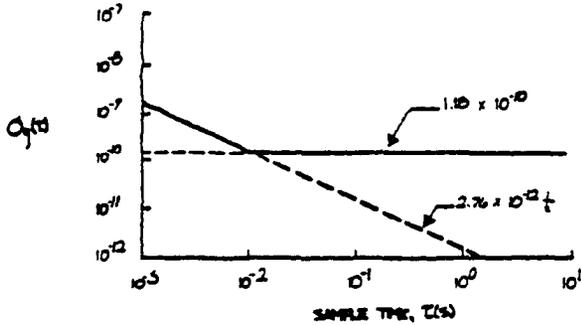The resultant time domain characterization is shown in figure 11.2.

FIGURE 11.4

Figures 11.3 and 11.4 show plots of time-domain stability and a translation to frequency domain. Since table 11.1 has the coefficients which connect both the frequency and time domains, it may be used for translation to and from either domain.

XII. CAUSES OF NOISE PROPERTIES IN A SIGNAL SOURCE

12.1 Power-law Noise Processes

Section IX pointed out the five commonly used power-law models of noise. With respect to $S_\phi(f)$, one can estimate a staight line slope (on a log-log scale) which corresponds to a particular noise type. This is shown in figure 12.1 (also fig.9.1).

FIGURE 11.2

The translation of $S_\phi(f)$ of figure 11.1 yields this $\sigma_y(\tau)$ plot.

FIGURE 11.3

FIGURE 12.1

37
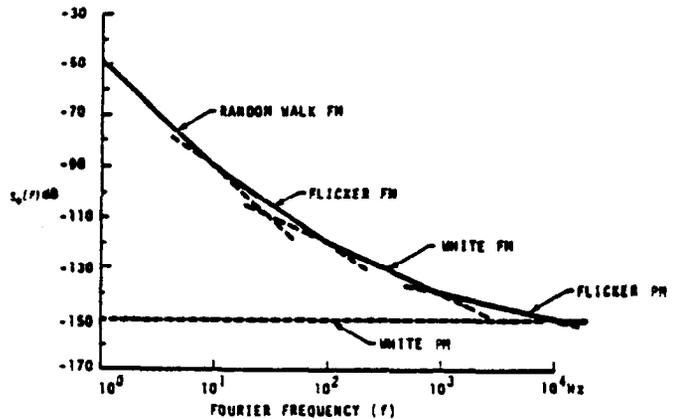
We can make the following general remarks about power-law noise processes:

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 usually not fully understood but may typically be related 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, which is locked to a resonance feature of another device which 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 the highest quality oscillators, 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 (e.g., 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 similar phenomena as 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 increasing, if feasible, the power of the primary frequency source.[13]

12.2 Other types of noise

A commonly encountered type of noise from a signal source or measurement apparatus is the presence of 60 Hz A.C. line noise. Shown in figure 12.2 is a constant white PM noise source with 60 Hz, 120 Hz 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 $S_\phi(f)$, one observes discrete line spectra. Although $S_\phi(f)$ is a measure of spectral density, one can interpret the line spectra with no loss of generality, although one usually does not refer to spectral densities when characterizing discrete lines. Figure 12.3 is the time domain representation of the same white phase modulation level with 60 Hz noise. Note that the amplitude of $\sigma_y(t)$ 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 alias effect) is a very powerful tool for filtering out a periodic wave imposed on a signal source. By sampling in the time domain at integer periods, one is virtually insensitive to the periodic (discrete line) term.

* See Appendix Note # 7

38

FM behavior masks the white PM (with the super-
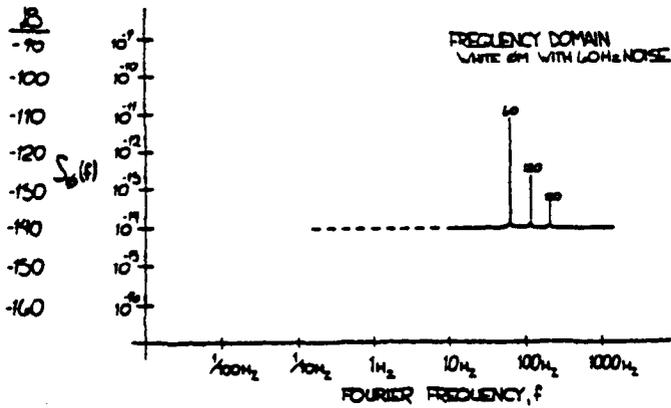imposed vibration characteristic) for long aver-
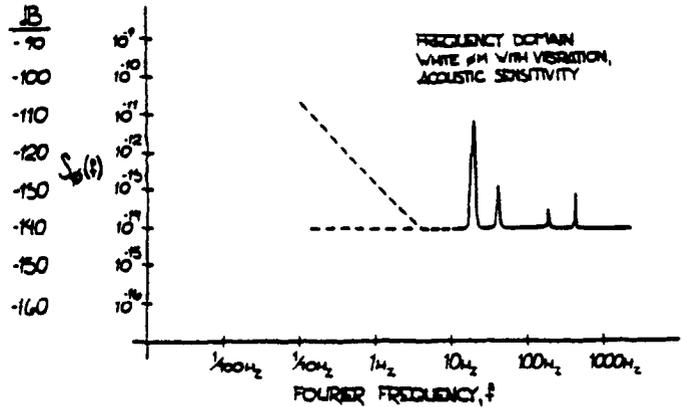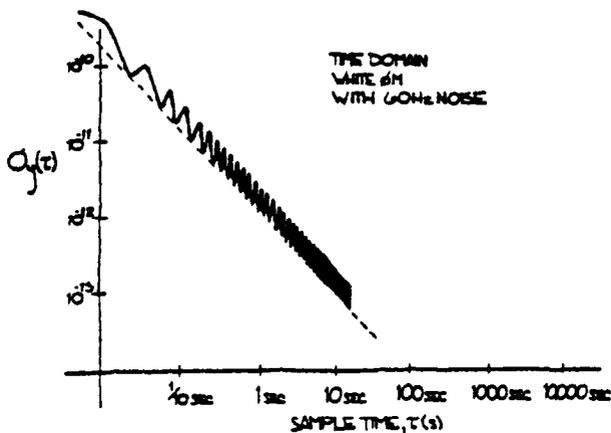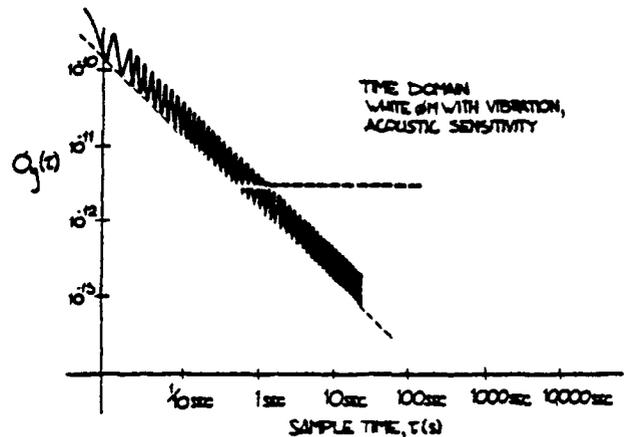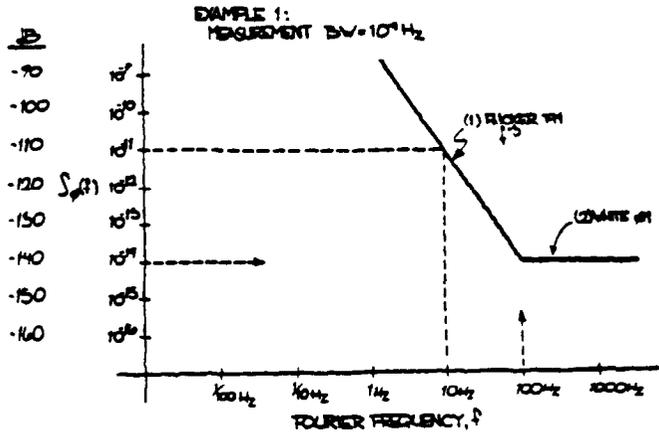aging times.

FIGURE 12.2

FIGURE 12.4

FIGURE 12.3

FIGURE 12.5

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

Figure 12.4 shows the kind of plot one might
see of $S_\phi(f)$ with vibration and acoustic sen-
sitivity in the signal source with the device
under vibration. Figure 12.5 shows the transla-
tion to the time domain of this effect. Also
noted in figure 12.4 is a (typical) flicker FM
behavior in the low frequency region. In the
translation to time domain (fig. 12.5), the flicker

Figure 12.6 shows examples of plots of two
power law processes $(S_\phi(f))$ with a change in the
flicker FM level. (Example 1 is identical to the
example given in sec. XI.) Figure 12.7 indi-
cates the effect of a lower flicker FM level as
translated to the time domain. Note again the
existence of both power law noise processes.
However for a given averaging time (or Fourier
frequency) one noise process may dominate over the
other.

39

TN-52

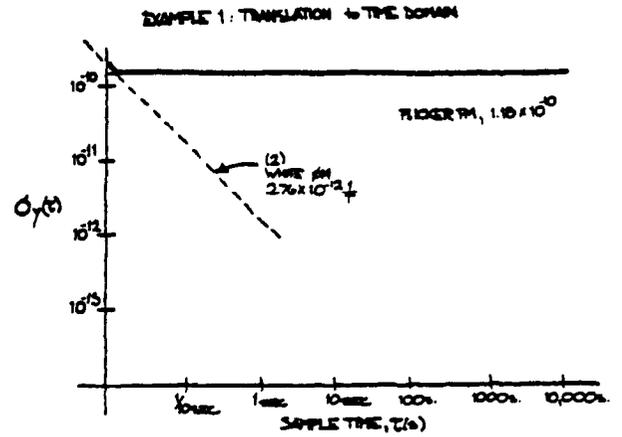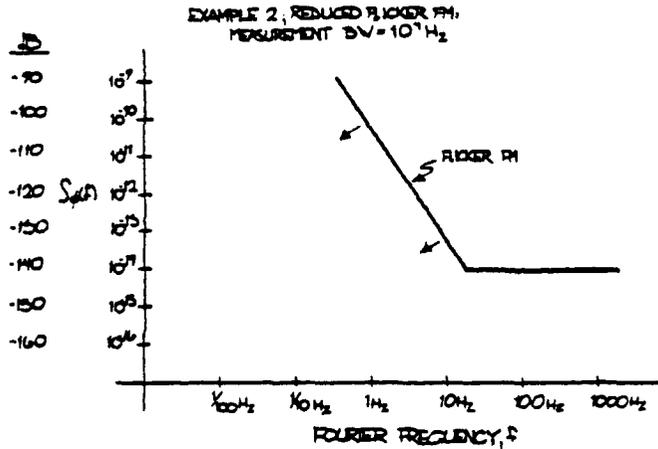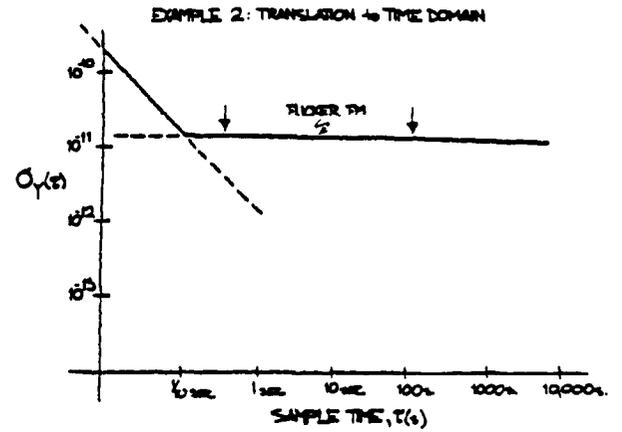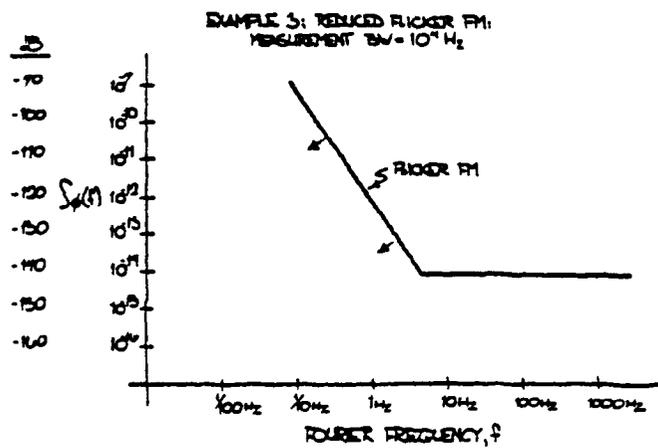EXAMPLE 1:
MEASUREMENT BW = 10⁴ Hz

FIGURE 12.6a

EXAMPLE 1: TRANSLATION to TIME DOMAIN

FIGURE 12.7a

EXAMPLE 2: REDUCED FLICKER FM:
MEASUREMENT BW = 10⁴ Hz

FIGURE 12.6b

EXAMPLE 2: TRANSLATION to TIME DOMAIN

FIGURE 12.7b

EXAMPLE 3: REDUCED FLICKER FM:
MEASUREMENT BW = 10⁴ Hz

FIGURE 12.6c

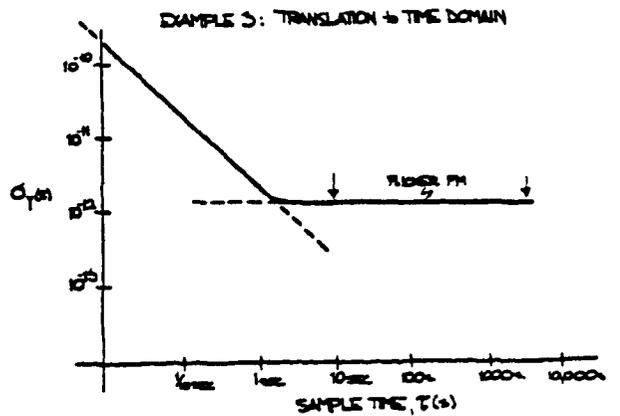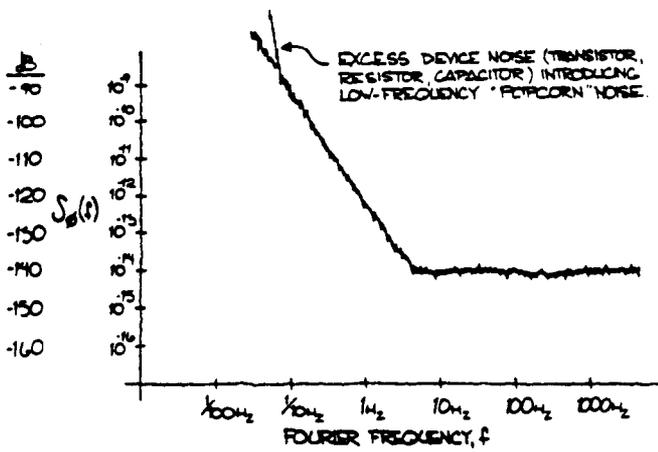EXAMPLE 3: TRANSLATION to TIME DOMAIN
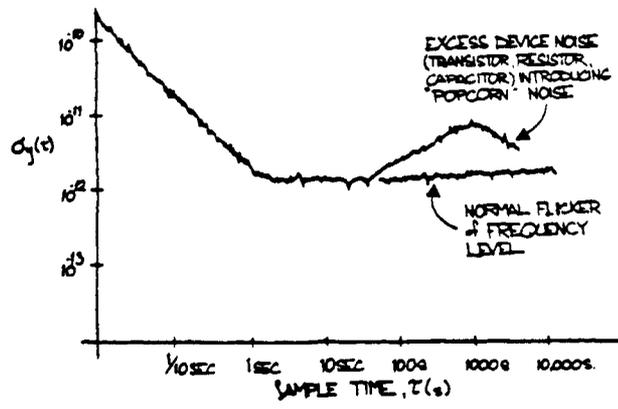
FIGURE 12.7c

40
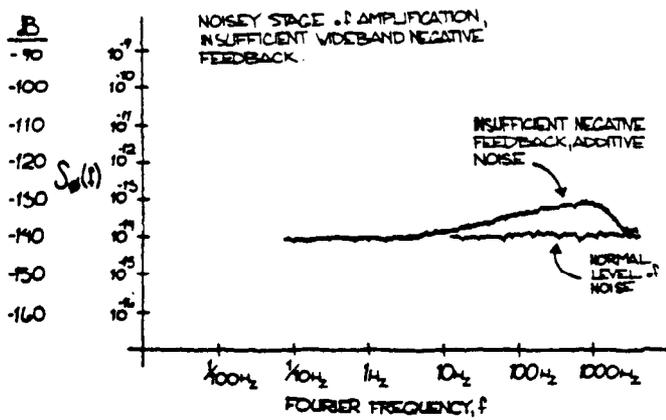
FIGURE 12.8

FIGURE 12.9

FIGURE 12.10

Excess device noise from transistors, capacitors, resistors, and the like can introduce a low frequency noise which has been referred to as "popcorn" noise because of its sonic qualities. Figure 12.8 shows a plot of $S_\phi(f)$ from a signal source having such excess low frequency noise. Figure 12.9 is the time domain representation. The rise in amplitude of σ_y for long averaging times is particularly aggravating. The solution to this kind of problem if it is introduced by devices is to carefully grade the devices in the assembly and testing process.

Stages of amplification following a signal source many times rely on local degenerate or overall negative feedback schemes in order to minimize the excess noise from active gain elements (such as transistors). This is the recommended design approach. However, phase shift in the negative feedback circuit or poor bandwidth in the gain elements can result in poor high frequency noise behavior. Figure 12.10 shows a kind of result one might see as a gradual rise in $S_\phi(f)$ because of insufficient negative feed back at high Fourier frequencies.

Section X discussed aliasing in the frequency domain. Figure 12.11 shows the resultant measurement anomaly due to digital sampling of a poorly bandlimited (anti-aliased) white noise source. Noise voltage above the sampling frequency f_s is folded into the analysis region of interest. Note also that the stopband ripple characteristics are folded into the high-frequency portion of the passband. For a given sampling frequency, a compromise exists between increasing the high-
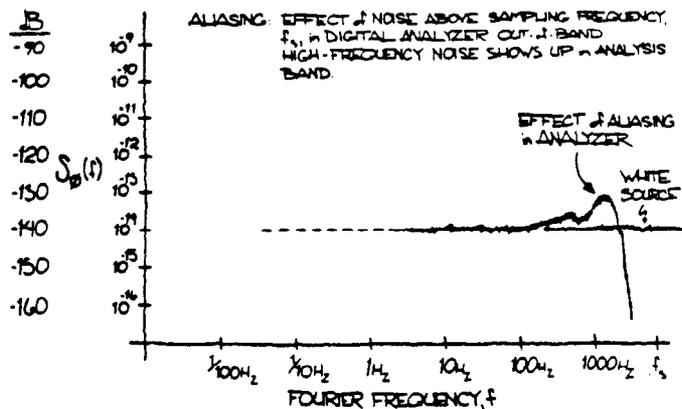

FIGURE 12.11

41

frequency extent of the analysis band and improving the anti-aliasing filter's stopband rejection. Section X has an example of the filter requirements for a particular case.

XIII. CONCLUSION

This writing highlights major aspects of time-domain and frequency-domain oscillator signal measurements. The contents are patterned after lectures presented by the authors. The authors have tried to be general in the treatment of topics, and bibliography is attached for readers who would like details about specific items.

References

1. J. A. Barnes, Andrew R. Chi, Leonard S. Cutler, Daniel J. Healey, David B. Leeson, Thomas E. McGunigal, James A. Mullen, Jr., Warren L. Smith, Richard L. Sydnor, Robert F. C. Vessot, and Gernot M. R. Winkler, "Characterization of Frequency Stability," NBS Tech. Note 394 (1970); Proc. IEEE Trans. on Instrum. and Meas. IM-20, 105 (1971).

2. D. W. Allan, "The Measurement of Frequency and Frequency Stability of Precision Oscillators," NBS Tech. Note 669 (1975); Proc. 6th PTTI Planning Meeting, 109 (1975).

3. D. W. Allan, "Statistics of Atomic Frequency Standards," Proc. IEEE 54, 221 (1966).

4. Sections V and VI are from notes to be published by J. A. Barnes.

5. S. R. Stein, "An NBS Phase Noise Measurement System Built for Frequency Domain Measurements Associated with the Global Positioning System, NBSIR 76-846 (1976). See also F. L. Walls and S. R. Stein, "Servo Technique in Oscillators and Measurement Systems," NBS Tech. Note 692 (1976).

6. C. D. McGillen and G. R. Cooper, Continous and Discrete Signal and System Analysis, Holt, Rinehart, and Winston, Inc. (1974).

7. E. O. Brigham, The Fast Fourier Transform, Prentice-Hall, Inc., New Jersey (1974).

8. Private communication with Norris Nahman, June 1981.

9. G. M. Jenkins and D. G. Watts, Spectral Analysis and Its Application, Holden-Day, San Francisco, 1968.

10. P. F. Panter, Modulation, Noise, and Spectral Analysis, McGraw-Hill, New York, 1965, pp. 137-141.

11. D. Babitch and J. Oliverio, "Phase Noise of Various Oscillators at Very Low Frequencies, Proc. 28th Annual Symp. on Frequency Control, May 1974. U. S. Army Electronics Command, Ft. Monmouth, NJ.

12. N. Thrane, "The Discrete Fourier Transform and FFT Analyzers," Bruel and Kjaer Technical Review, No. 1 (1979).

13. Much of the material presented here is contained in NBS Tech. Note 679, "Frequency Domain Stability Measurements: A Tutorial Introduction," D. A. Howe (1976).

Bibliography

General References

1. November or December of even-numbered years, IEEE Transactions on Instrumentation and Measurement (Conference on Precision Electromagnetic Measurements, held every two years).

2. Proceedings of the IEEE. Special issue on frequency stability, vol. 54, February 1966.

3. Proceedings of the IEEE. Special issue on time and frequency, vol. 60, no. 5, May 1972.

4. Proceedings of the IEEE-NASA Symposium on the Definition and Measurement of Short-Term Frequency Stability at Goddard Space Flight Center, Greenbelt, MD, November 23-24, 1964. Prepared by Goddard Space Flight Center (Scientific and Technical Information Division, National Aeronautics and Space Administration, Washington, DC, 1965). Copies available for $1.75 from the US Government Printing Office, Washington, DC 20402.

5. Proceedings of the Annual Symposium on Frequency Control. Prepared by the US Army Electronics Command, Fort Monmouth, NJ. Copies from Electronics Industry Association, 2001 Eye Street, N. W., Washington, DC 20006, starting with 32nd Annual or call (201) 544-4878. Price per copy, $20.00.

6. J. A. Barnes, A. R. Chi, L. S. Cutler, et al., "Characterization of Frequency Stability," IEEE Trans. on Intsrum. and Meas. IM-20, no. 2, pp. 105-120 (May 1971). Prepared by the Subcommittee on Frequency Stability of the Institute of Electrical and Electronics Engineers.

7. Time and Frequency: Theory and Fundamentals, Byron E. Blair, Editor, NBS Monograph 140, May 1974.

8. P. Kartaschoff, Frequency and Time, London, England: Academic Press, 1978.

9. G. M. R. Winkler, Timekeeping and Its Application, Advances in Electonics and Electron Physics, Vol. 44, New York: Academic Press, 1977.

10. F. M. Gardner, Phaselock Techniques, New York: Wiley & Sons, 1966.

11. E. A. Gerber, "Precision Frequency Control and Selection, A Bibliography," Proc. 33rd Annual Symposium on Frequency Control, pp. 569-728, 1979. Available from Electronics Industries Association, 2001 Eye St. NW, Washington, DC 20006.

12. J. Rutman, "Characterization of Phase and Frequency Instabilities in Precision Frequency Sources: Fifteen Years of Progress," Proc. IEEE, vol. 66, no. 9, September 1978, pp. 1048-1075.

Additional Specific References (selected)

1. D. W. Allan, "Statistics of Atomic Frequency Standards," Proc. IEEE, vol. 54, pp. 221-230, February 1966.

2. D. W. Allan, "The Measurement of Frequency and Frequency Stability of Precision Oscillators," NBS Tech. Note 669, July 1975.

3. J. R. Ashley, C. B. Searles, and F. M. Palka, "The Measurement of Oscillator Noise at Microwave Frequencies," IEEE Trans. on Microwave Theory and Techniques MTT-16, pp. 753-760, September 1968.

4. E. J. Baghdady, R. D. Lincoln, and B. D. Nelin, "Short-Term Frequency Stability: Theory, Measurement, and Status," Proc. IEEE, vol. 53, pp. 704-722, 2110-2111, 1965.

5. J. A. Barnes, "Models for the Interpretation of Frequency Stability Measurements," NBS Technical Note 683, August 1976.

6. J. A. Barnes, " A Review of Methods of Analyzing Frequency Stability," Proc. 9th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, Nov. 1977, pp. 61-84.

7. J. A. Barnes, "Atomic Timekeeping and the Statistics of Precision Signal Generators," Proc. IEEE, vol. 54, pp. 207-220, February 1966.

8. J. A. Barnes, "Tables of Bias Functions, B_1 and B_2, for Variances Based on Finite Samples of Processes with Power Law Spectral Densities," NBS Tech. Note 375, January 1969.

9. J. A. Barnes and D. W. Allan, "A Statistical Model of Flicker Noise," Proc. IEEE, vol. 43, pp. 176-178, February 1966.

10. J. A. Barnes and R. C. Mockler, "The Power Spectrum and its Importance in Precise Frequency Measurements," IRE Trans. on Instrumentation I-9, pp. 149-155, September 1960.

11. W. R. Bennett, "Methods of Solving Noise Problems," Proc. IRE, May 1956, pp. 609-637.

12. W. R. Bennett, Electrical Noise, New York: McGraw-Hill, 1960.

13. C. Bingham, M. D. Godfrey, and J. W. Tukey, "Modern Techniques of Power Spectrum Estimation," IEEE Trans. AU, vol. 15, pp. 56-66, June 1967.

14. R. B. Blackman and J. W. Tukey, The Measurement of Power Spectra, New York: Dover, 1958.

15. E. O. Brigham and R. E. Morrow, "The Fast Fourier Transform," IEEE Spectrum, vol. 4, no. 12, pp. 63-70, December 1967.

16. J. C. Burgess, "On Digital Spectrum Analysis of Periodic Signals," J. Acoust. Soc. Am., vol. 58, no. 3, September 1975, pp. 556-557.

17. R. A. Campbell, "Stability Measurement Techniques in the Frequency Domain," Proc. IEEE-NASA Symposium on the Definition and Measurement of Short-Term Frequency Stability, NASA SP-80, pp. 231-235, 1965.

18. L. S. Cutler and C. L. Searle, "Some Aspects of the Theory and Measurement of Frequency Fluctuations in Frequency Standards," Proc. IEEE, vol. 54, pp. 136-154, February 1966.

19. W. A. Edson, "Noise in Oscillators," Proc. IRE, vol. 48, pp. 1454-1466, August 1960.

20. C. H. Grauling, Jr. and D. J. Healey, III, "Instrumentation for Measurement of the Short-Term Frequency Stability of Microwave Sources," Proc. IEEE, vol. 54, pp. 249-257, February 1956.

21. D. Halford, "A General Mechanical Model for $|f|^\alpha$ Spectral Density Random Noise with Special Reference to Flicker Noise $1/|f|$," Proc. of the IEEE, vol. 56, no. 3, March 1968, pp. 251-258.

22. H. Hellwig, "Atomic Frequency Standards: A Survey," Proc. IEEE, vol. 63, pp. 212-229, February 1975.

23. H. Hellwig, "A Review of Precision Oscillators," NBS Tech. Note 662, February 1975.

24. H. Hellwig, "Frequency Standards and Clocks: A Tutorial Introduction," NBS Tech. Note 616-R, March 1974.

25. Hewlett-Packard, "Precision Frequency Measurements," Application Note 116, July 1969.

26. G. M. Jenkins and D. G. Watts, Spectral Analysis and its Applications, San Francisco: Holden-Day, 1968.

27. S. L. Johnson, B. H. Smith, and D. A. Calder, "Noise Spectrum Characteristics of Low-Noise Microwave Tubes and Solid-State Devices," Proc. IEEE, vol. 54, pp. 258-265, February 1966.

43

28. P. Kartaschoff and J. Barnes, "Standard Time and Frequency Generation," Proc. IEEE, vol. 60, pp. 493-501, May 1972.

29. A. L. Lance, W. D. Seal, F. G. Mendoza, and N. W. Hudson, "Automating Phase Noise Measurements in the Frequency Domain," Proc. 31st Annual Symposium on Frequency Control, 1977, pp. 347-358.

30. D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," Proc. IEEE L., vol. 54, pp. 329-330, February 1966.

31. P. Lesage and C. Audoin, "A Time Domain Method for Measurement of the Spectral Density of Frequency Fluctuations at Low Fourier Frequencies," Proc. 29th Annual Symposium on Frequency Control, 1975, pp. 394-403.

32. P. Lesage and C. Audoin, "Estimation of the Two-Sample Variance with Limited Number of Data," Proc. 31st Annual Symposium on Frequency Control, 1977, pp. 311-318.

33. P. Lesage and C. Audoin, "Characterization of Frequency Stability: Uncertainty due to Finite Number of Measurements," IEEE Trans on Instrum. and Meas. IM-22, pp. 157-161, June 1973.

34. W. C. Lindsey and C. M. Chie, "Specification and Measurement of Oscillator Phase Noise Instability," Proc. 31st Annual Symposium on Frequency Control, 1977, pp. 302-310.

35. D. G. Meyer, "A Test Set for Measurement of Phase Noise on High Quality Signal Sources," Proc. IEEE Trans on Instrum. and Meas. IM-19, No. 4, pp. 215-227, November 1970.

36. C. D. Motchenbacher and F. C. Fitchen, Low-Noise Electronic Design, New York: John Wiley & Sons, 1973.

37. D. Percival, "Prediction Error Analysis of Atomic Frequency Standards," Proc. 31st Annual Symposium on Frequency Control, 1977, pp. 319-326.

38. D. Percival, "A Heuristic Model of Long-Term Clock Behavior," Proc. 30th Annual Symposium on Frequency Control, 1976, pp. 414-419.

39. J. Shoaf, "Specification and Measurement of Frequency Stability," NBS Internal Report 74-396, November 1974.

40. J. L. Stewart, "Frequency Modulation Noise in Oscillators," Proc. IRE, vol. 44, pp. 372-376, March 1956.

41. R. F. C. Vessot, L. Mueller, and J. Vanier, "The Specification of Oscillator Characteristics from Measurements in the Frequency Domain," Proc. IEEE, pp. 199-206, February 1966.

42. F. L. Walls, S. R. Stein, J. E. Gray, and D. J. Glaze, "Design Considerations in State-of-the-Art Signal Processing and Phase Noise Measurement Systems," Proc. 30th Annual Symposium on Frequency Control, 1976, pp. 269-274.

43. F. Walls and A. Wainwright, "Measurement of the Short-Term Stability of Quartz Crystal Resonators and the Implications for Quartz Oscillator Design and Applications," IEEE Trans. on Instrum. and Meas., pp. 15-20, March 1975.

44. G. M. R. Winkler, "A Brief Review of Frequency Stability Measures," Proc. 8th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, Nov. 1976, pp. 489-527.
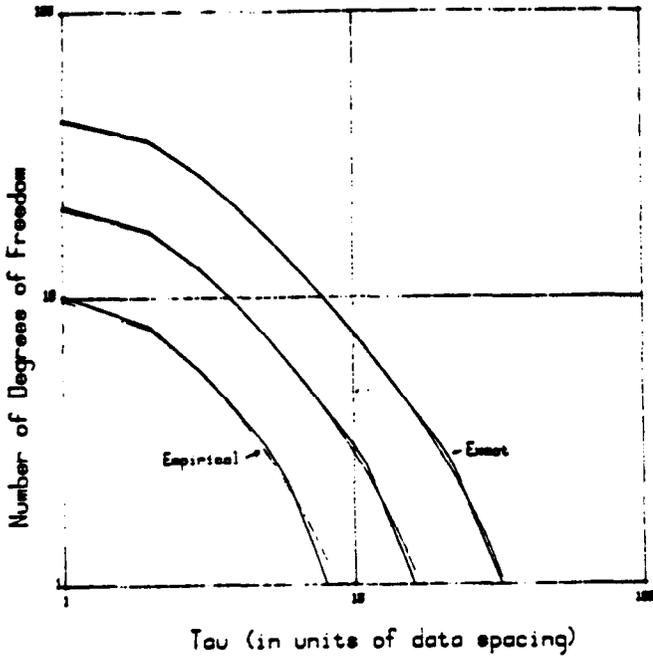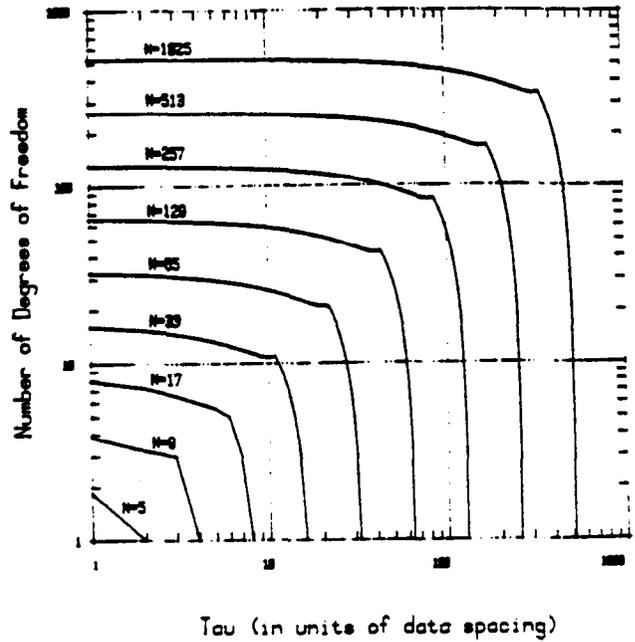
Appendix

DEGREES OF FREEDOM FOR ALLAN VARIANCE
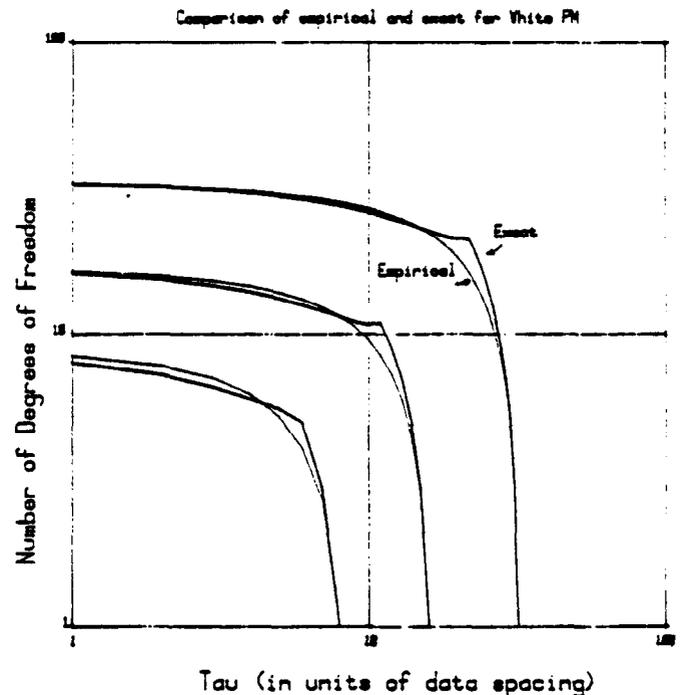
Comparison of empirical and exact for White PM

Tau (in units of data spacing)

TN-57

DEGREES OF FREEDOM FOR ALLAN VARIANCE

Comparison of empirical and exact for White FM

Number of Degrees of Freedom

Empirical

Exact

Tau (in units of data spacing)

DEGREES OF FREEDOM FOR ALLAN VARIANCES

White Phase Noise

Number of Degrees of Freedom

N=1025
N=513
N=257
N=129
N=65
N=29
N=17
N=9
N=5

Tau (in units of data spacing)

DEGREES OF FREEDOM FOR ALLAN VARIANCE

Comparison of empirical and exact for Random Walk FM

Number of Degrees of Freedom

Exact

Empirical

Tau (in units of data spacing)
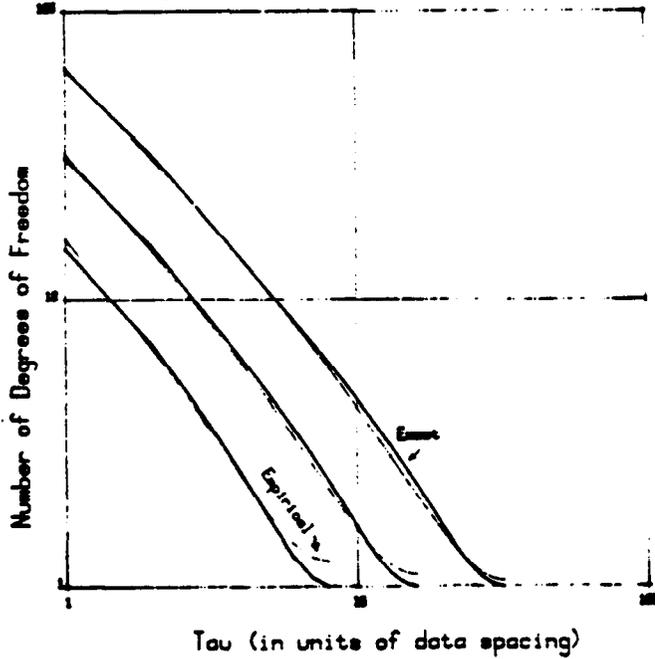
DEGREES OF FREEDOM FOR ALLAN VARIANCES

White Frequency Noise

Number of Degrees of Freedom

N=1025
N=513
N=257
N=129
N=65
N=29
N=17
N=9
N=5

Tau (in units of data spacing)

(Empirical formula used)

45
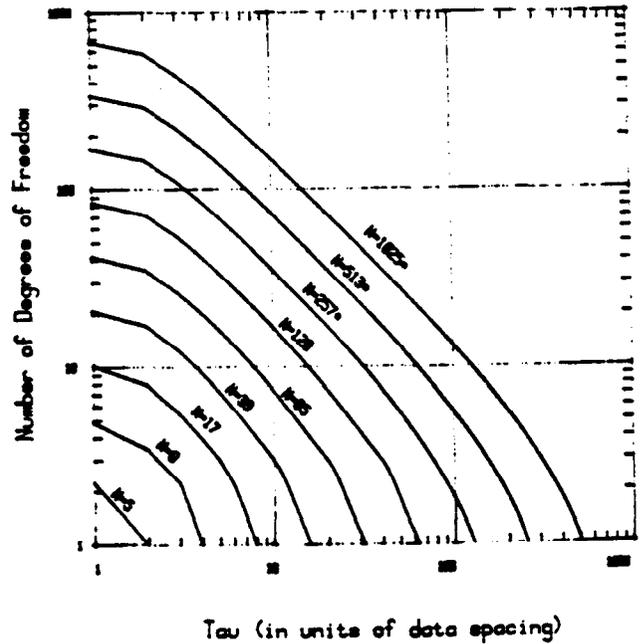
DEGREES OF FREEDOM FOR ALLAN VARIANCES

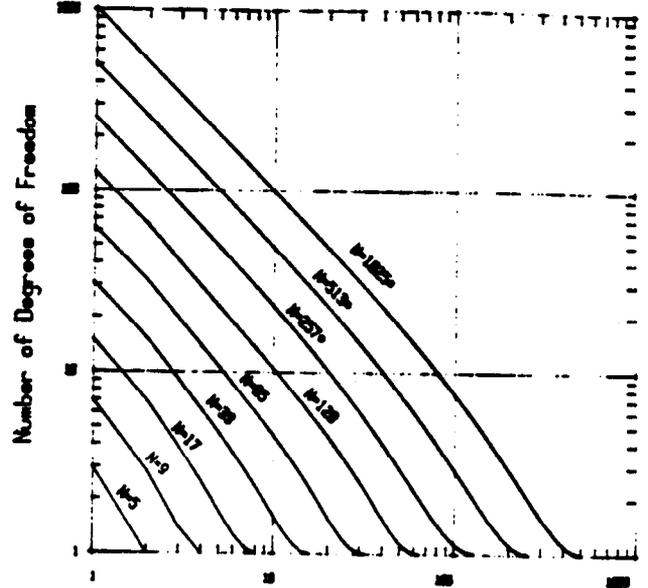### Flicker Phase Modulation
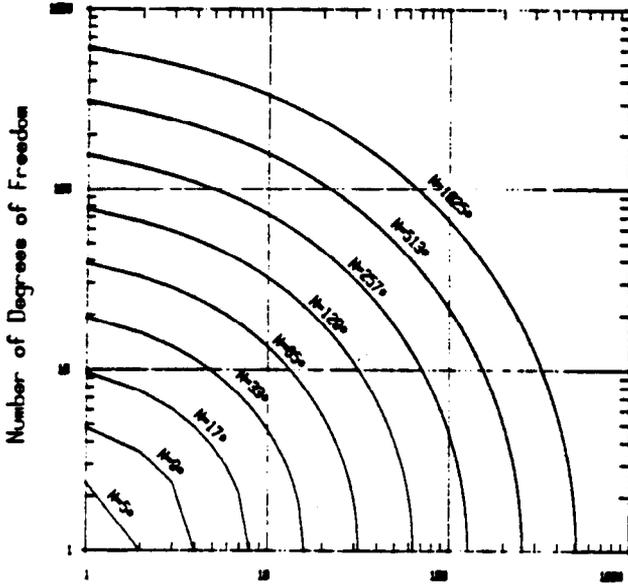
Tau (in units of data spacing)

(• -Empirical Formula used)

DEGREES OF FREEDOM FOR ALLAN VARIANCES
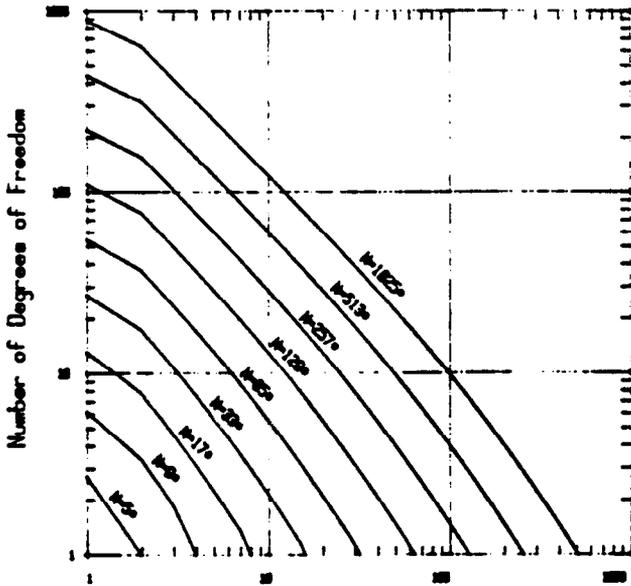
Random Walk FM

Tau (in units of data spacing)

(• -Empirical Formula used)

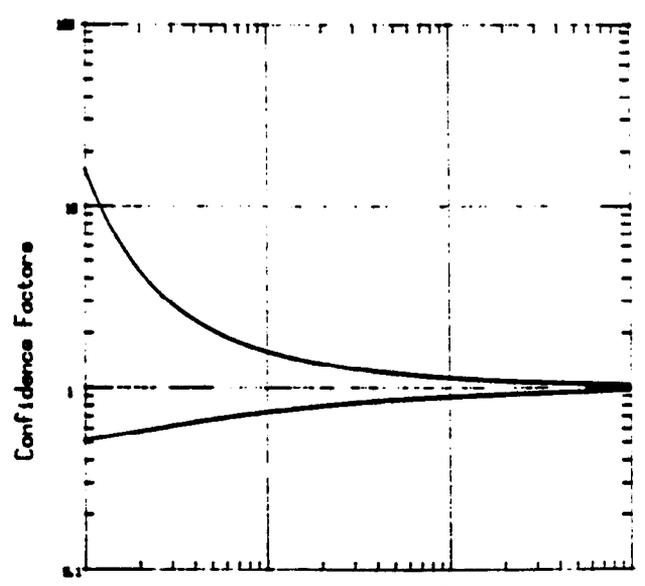## DEGREES OF FREEDOM FOR ALLAN VARIANCES

Flicker Frequency Modulation

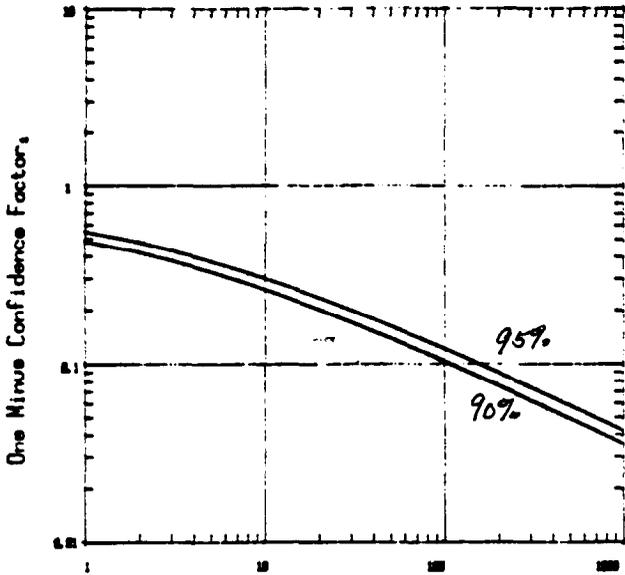Tau (in units of data spacing)

(• -Empirical Formula used)

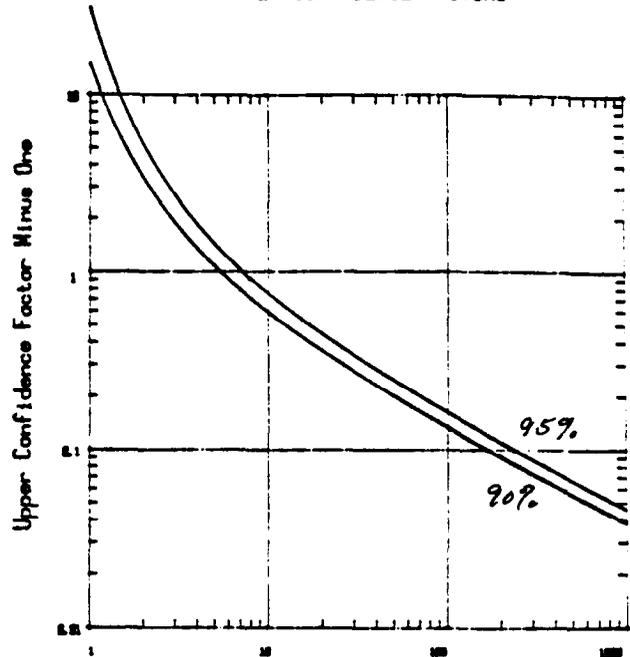## CONFIDENCE FACTORS FROM CHI-SQUARE

90% Intervals

Number of Degrees of Freedom

46

TN-59

LOWER CONFIDENCE FACTORS

UPPER CONFIDENCE FACTORS

Number of Degrees of Freedom

Number of Degrees of Freedom

Distribution of Ratio of Observed Variance to Actual

Chi Square/Degree of Freedom

47

TN-60