

Noise in Mixers, Oscillators, Samplers, and Logic

An Introduction to Cyclostationary Noise

Joel Phillips and Ken Kundert

Version 1a, May 2000

The proliferation of wireless and mobile products has dramatically increased the number and variety of low power, high performance electronic systems being designed. Noise is an important limiting factor in these systems. The noise generated is often cyclostationary. This type of noise cannot be predicted using SPICE, nor is it well handled by traditional test equipment such as spectrum analyzers or noise figure meters, but it is available from the new RF simulators.

The origins and characteristics of cyclostationary noise are described in a way that allows designers to understand the impact of cyclostationarity on their circuits. In particular, cyclostationary noise in time-varying systems (mixers), sampling systems (switched filters and sample/holds), thresholding systems (logic circuitry), and autonomous systems (oscillators) is discussed.

Search Terms

Oscillator phase noise, jitter, noise folding, modulated noise, time-varying noise, cyclostationary noise.

Published in the Proceedings of the IEEE Custom Integrated Circuits Conference, May 2000. A companion document that contains the slides and speaker notes for the associated presentation is also available from designers-guide.org/theory.

Last updated on March 10, 2019. You can find the most recent version at designers-guide.org. Contact the author via e-mail at ken@designers-guide.com.

Permission to make copies, either paper or electronic, of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage and that the copies are complete and unmodified. To distribute otherwise, to publish, to post on servers, or to distribute to lists, requires prior written permission.

1.0 Introduction

SPICE noise analysis is not able to compute valid noise results for many common classes of circuit for which noise is of interest. Circuits such as mixers, oscillators, samplers, and logic gates either produce noise at their output whose power varies significantly with time, or whose sensitivity to noise varies significantly with time, or both. New simulation algorithms have recently become available that can be used to predict the noise performance of these types of circuits [9,12,14]. However, noise of this type is unfamiliar to most designers. This paper introduces the ideas needed to understand and model noise in these types of circuits using terminology and concepts familiar to circuit designers.

1.1 Ensemble Averages

Noise free systems are deterministic, meaning that repeating the same experiment produces the same result. Noisy systems are stochastic — repeating the same experiment produces slightly different results each time. An experiment is referred to as a trial and a group of experiments is referred to as an ensemble of trials, or simply an ensemble. Noise can be characterized by using averages over the ensemble, called expectations, and denoted by the operator $E\{\cdot\}$. The expectation is the limit of the ensemble average as the number of trial approaches infinity.

Let v_n be a noisy signal. It can be separated into a purely noise free, or deterministic, signal v , and a stochastic signal that is pure noise, n , where

$$v_n(t) = v(t) + n(t). \quad (1)$$

The mean of the noisy signal is the noise free signal, $E\{v_n(t)\} = v(t)$, and the mean of the noise is zero, $E\{n(t)\} = 0$. The variance of $n(t)$, $\text{var}(n(t)) = E\{n(t)^2\}$, is a measure of the power in the noise at a specific time. A more general power-like quantity is the autocorrelation, $R_n(t, \tau) = E\{n(t) n(t-\tau)\}$, a measure of how points on the same signal separated by τ seconds are correlated. The autocorrelation is related to the variance by $\text{var}(n(t)) = R_n(t, 0)$. By performing the Fourier transform of the autocorrelation function with respect to the variable τ and then averaging over t , we obtain the time-averaged power spectral density, or PSD, that is measured by spectrum analyzers.

1.2 Colored or Time-Correlated Noise

Noise that is completely uncorrelated versus time is known as white noise. For white noise the PSD is a constant and the autocorrelation function is an impulse function centered at 0, $R_n(t, \tau) = R(t)\delta(\tau)$.

If the noise passes through a circuit that contains energy storage elements, such as capacitors and inductors, the PSD of the resulting signal will be shaped by the transfer function of the circuit. This shaping of the noise versus frequency is referred to as coloring the noise.

Energy storage elements also cause the noise to be correlated versus time. This occurs simply because noise produced at one point in time is stored in the energy storage element, and comes out some time later. This results in the autocorrelation function having nonzero width in τ .

The energy-storage elements cause the noise spectrum to be shaped and the noise to be time-correlated. This is a general property. If the noise has shape in the frequency domain then the noise is correlated in time, and vice versa.

1.3 Cyclostationary or Frequency-Correlated Noise

Circuits with time-varying operating points can cause the ensemble averages that describe noise to vary with time t . If they vary in a periodic fashion, the noise is said to have cyclostationary properties, and the ensemble averages referred to as being cyclostationary [3,4,5]. If they vary in a quasiperiodic fashion, they are referred to as polycyclostationary, though in this paper there will be no distinction made between cyclostationary and polycyclostationary processes.

Cyclostationarity occurs when the time-varying operating point modulates the noise generated by bias-dependent noise sources or when the time-varying circuit modulates the transfer function from the noise source to the output. As suggested by the name, modulated noise sources can be modeled by modulating the output of stationary noise sources.

FIGURE 1.

The resistor generates white thermal noise. The switch opens and closes periodically, so the noise at the output is cyclostationary.



Figure 1 shows a simple example of cyclostationary noise. A periodically operating switch between the noise source (the resistor generating white thermal noise) and the observer causes the output noise to have periodically varying statistics.

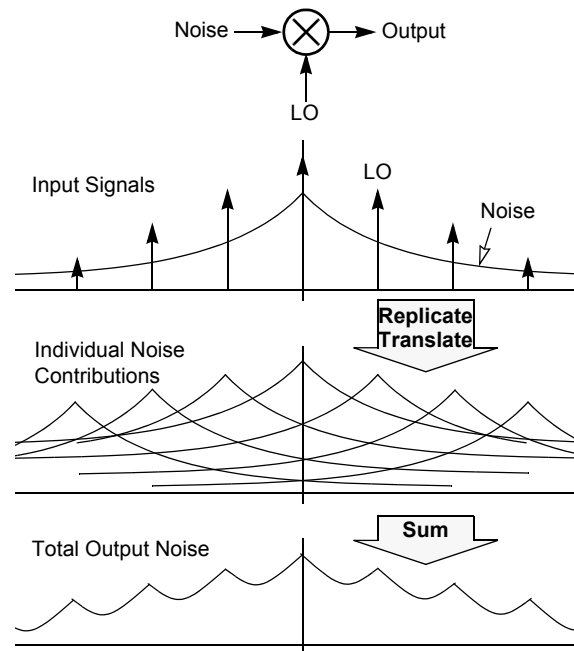
Noise is transmitted from the resistor to the observer only when the switch is closed. It can be said that cyclostationary noise is “shaped in time t ”. However, with no energy storage elements the noise is completely uncorrelated versus time τ (noise at a particular time is uncorrelated with the noise at any previous time) and therefore is white, even though it is cyclostationary. One cannot tell that noise is cyclostationary by just observing the time-average PSD.

In the example shown in Figure 2, stationary noise with an arbitrary PSD is modulated by a periodic signal. This is representative of both ways in which cyclostationary noise is generated (modulated noise sources and modulated signal paths). It is also representative of how noise is modulated in many types of circuits. In a mixer, the noise is modulated by the LO. In a sampler, it is modulated by the clock. In a digital logic, the noise is modulated by the logic signals. And in an oscillator, it is modulated by the oscillation signal itself.

Modulation can be interpreted as multiplication in the time domain or convolution in the frequency domain. Thus, the modulation by a periodic signal causes the noise to mix up and down in multiples of the modulation frequency in a process that is often referred to as noise folding.

FIGURE 2.


How noise is moved around by a mixer. The noise is replicated and translated by each harmonic of the LO, resulting in correlations at frequencies separated by kf_{LO} .



Noise from the source at a particular frequency f is replicated and copies appear at $f \pm kf_0$, where k is an integer and f_0 is the fundamental frequency of the periodic signal. Conversely, noise at the output at a particular frequency f has contributions from noise from the sources at frequencies $f \pm kf_0$.

FIGURE 3.

With a complex phasor representation of noise, noise at frequencies separated by $k\omega_0$ is correlated. When converted to real signals, the complex conjugate of the noise at negative frequencies is mapped to positive frequencies. As a result, the upper and lower sidebands contain correlated noise.



Because of the translation of replicated copies of the same noise source, noise separated by kf_0 is generally correlated. Remember that noise folds across DC, so noise in upper and lower sidebands will be correlated. Consider the top of Figure 3 where noise is

shown at both negative and positive frequencies. This implies a complex phasor representation is being used. When this complex signal is converted to a real signal, the complex conjugate of signals at negative frequencies is mapped to positive frequencies. In this way, the signal at frequencies $\Delta\omega$ above and below a harmonic are correlated. These frequencies are referred to as upper and lower sidebands of the harmonic.

Recall from the previous section that

$$\text{shape in frequency} \Leftrightarrow \text{correlation in time}$$

Now from this section also see that

$$\text{shape in time} \Leftrightarrow \text{correlation in frequency}$$

This is the duality of shape and correlation. If one is known, the other can be recovered. This is important because it allows us to choose either the time or frequency domain to describe noise in any particular system by simply noting whether the dominant statistical effects are more easily described by the shape or the correlation.

2.0 Calculating Noise

Noise is generally so small that it does not cause the circuit to behave nonlinearly (one exception is with oscillators, which is discussed later). Therefore noise is calculating using perturbation techniques, that is, by splitting the noisy signal into large and small components. The large signal component (the operating point) is periodic and the small component (the noise) is stochastic. First we set the small stochastic portion of the stimulus to zero by disabling all of the noise sources and solve for the large-signal periodic steady-state solution that determines the circuit operating point. We then linearize the circuit about the periodic large signal operating point and apply the small stochastic signal to this linearized system. The linearized system is time-varying and unlike linear time-invariant systems, can model frequency conversion effects that create cyclostationarity. The linear time-varying system is solved numerically. These linear time-varying systems generally are quite large and require special numerical techniques to be practical. The reader is referred to [12,14] for details of numerical implementations.

3.0 Characterizing Cyclostationary Noise

There are three common methods of characterizing cyclostationary noise.

The time-average power spectral density is similar to what would be measured with a conventional spectrum analyzer. Since the analyzer has a very small effective input bandwidth, it ignores correlations in the noise and so ignores the cyclostationary nature of the noise (assuming that the frequency of the cyclostationarity is much higher than the bandwidth of the analyzer).

The second method is to use the spectrum along with information about the correlations in the noise between sidebands. This is a complete description of the cyclostationarity in the noise. It is used when considering the impact of cyclostationary noise from one stage on a subsequent synchronous stage. Two stages would be synchronous if they were driven by the same LO or clock, or if the output of one stage caused the subsequent

stage to behave nonlinearly. From this form it is relatively easy to determine the amount of power in the AM or PM components of the noise.

The third method is to track the noise at a point in phase, or noise versus phase. The noise at a point in phase is defined as the noise in the sequence of values obtained if a noisy periodic signal¹ is repeatedly sampled at the same point in phase during each period. It is useful in determining the noise that results when converting a continuous-time signal to a discrete-time signal. It is also useful when determining the jitter associated with a noisy signal crossing a threshold.

3.1 Time-Average Power Spectral Density

If a stage that generates cyclostationary noise is followed by a filter whose passband is constrained to a single sideband (the passband does not contain a harmonic and has a bandwidth of less than $f_0/2$, where f_0 is the fundamental frequency of the cyclostationarity), then the output of the filter will be stationary. This is true because noise at any frequency f_1 is uncorrelated with noise at any other frequency f_2 as long as both f_1 and f_2 are within the passband.

Consider a stage that generates cyclostationary noise with modulation frequency f_1 that is followed by a stage whose transfer characteristics vary periodically at a frequency of f_2 (such as a mixer, sampler, etc.). Assume that f_1 and f_2 are non commensurate (there is no f_0 such that $f_1 = n f_0$ and $f_2 = m f_0$ with n and m both integers). Then there is no way to shift f_1 by a multiple of f_2 and have it fall on a correlated copy of itself. As a result, the cyclostationary nature of the noise at the output of the first stage can be ignored. With regard to its effect on the subsequent stage, the noise from the first stage can be treated as being stationary and we can characterize it using the time-average power spectral density [8,13].

If f_1 and f_2 are commensurate, but m and n are both large with no common factors, then many periods of f_1 and f_2 are averaged before the exact phasing between the two repeats. In this case, the cyclostationary nature of the noise at the output of the first stage can often be ignored.

The time-averaged power spectral density (PSD) can be used as the basis of a noise model when the subsequent stages eliminate or ignore the cyclostationary nature of the noise. Filtering eliminates the cyclostationary nature of noise, converting it to stationary noise, if the filter is a single-sideband filter with bandwidth less than $f_0/2$. The cyclostationary nature of the noise is ignored if the subsequent stage is not synchronous with the noise, or if it is synchronous but running at a sufficiently different frequency so that averaging serves to eliminate the cyclostationarity.

When a stage producing cyclostationary noise drives a subsequent stage that has a time-varying transfer function that is synchronous with the first, then ignoring the cyclostationary nature of the noise from the first stage (say by using the time-average PSD) generates incorrect results. One common situation where this occurs is when a switched-capacitor filter is followed by a sample-and-hold, and both are clocked at the same rate (or a multiple of the same rate). Another common situation is when the first stage produces a periodic signal that is large enough to drive the subsequent stage to behave non-

1. By *noisy periodic signal* we mean a signal of the form $v_n(t) = v(t) + n(t)$ where $v(t)$ is T -periodic and $n(t)$ is T -cyclostationary but is not periodic.

linearly. In this case, the large periodic output signal will modulate the gain of the subsequent stage synchronously with the cyclostationary noise produced by the first stage. This occurs when an oscillator drives the LO port of a mixer or sampler, when one logic gate drives another, or when a large interfering signal drives two successive stages into compression.

In these situations, the cyclostationary nature of the noise produced in the first stage must be considered when determining the overall noise performance of the stages together.

3.2 AM & PM Noise

One can separate noise near the carrier into AM and PM components [8,11]. Consider the noise at sidebands at frequencies $\Delta\omega$ from the carrier. Treat both these sidebands and the carrier as phasors. Individually add the sideband phasors to the carrier phasor. The sideband phasors are at a different frequency from the carrier, and so rotate relative to it. One sideband will rotate at $\Delta\omega$, and the other at $-\Delta\omega$. If the noise is not cyclostationary, then the two sidebands will be uncorrelated, meaning that their amplitude and phase will vary randomly relative to each other. Combined, the two sideband phasors will trace out an ellipse whose size, shape, and orientation will shift randomly. However, if the noise is cyclostationary, then the sidebands are correlated. This reduces the random shifting in the shape and orientation of the ellipse traced out by the phasors. If the noise is perfectly correlated, then the shape and orientation will remain unchanged, though its size still shifts randomly.

The shape and orientation of the ellipse is determined by the relative size of the AM and PM components in the noise. This is demonstrated in Figure 4. For example, oscillators almost exclusively generate PM noise near the carrier whereas noise on the control input to a variable gain amplifier results almost completely in AM noise at the output of the amplifier. Having one component of noise dominate over the other is a characteristic of cyclostationary noise. Stationary noise can also be decomposed into AM and PM components, but there will always be equal amounts of both.

It is a general rule that combining stationary noise with a large periodic or quasiperiodic signal and is passing it through a stage undergoing compression or saturation results primarily in phase noise at the output. Stationary noise contains equal amounts of amplitude and phase noise. Passing it through a stage undergoing compression causes the amplitude noise to be suppressed, leaving mainly the noise in phase.


4.0 Oscillator Phase Noise

It is the nature of all autonomous systems, such as oscillators that they produce relatively high levels of noise at frequencies close to the oscillation frequency. Because the noise is close to the oscillation frequency, it cannot be removed with filtering without also removing the oscillation signal. It is also the nature of nonlinear oscillators that the noise be predominantly in the phase of the oscillation. Thus, the noise cannot be removed by passing the signal through a limiter. This noise is referred to as oscillator phase noise.

In a receiver, the phase noise of the LO can mix with a large interfering signal from a neighboring channel and swamp out the signal from the desired channel even though

FIGURE 4.

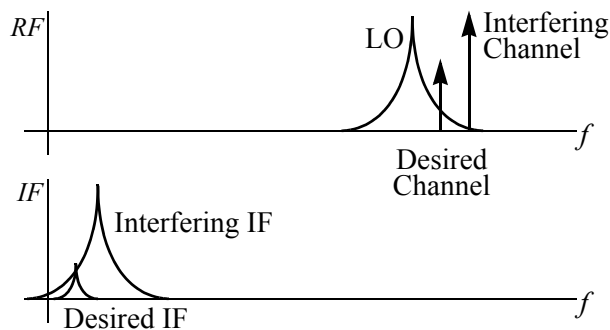
How the amplitude and phase relationship between sidebands cause AM and PM variations in a carrier. The phasors with the hollow tips represents the carrier; the phasors with the solid tips represent the sidebands. The upper sideband rotates in the clockwise direction and the lower in the counterclockwise direction. The composite trajectory is shown below the individual components. (a) Single-sideband modulation (only upper sideband). (b) Arbitrary double-sideband modulation where there is no special relationship between the sidebands. (c) Amplitude modulation (identical magnitudes and phase such that phasors point in same direction when parallel to carrier). (d) Phase modulation (identical magnitudes and phase such that phasors point in same direction when perpendicular to carrier).



most of the power in the interfering IF is removed by the IF filter. This is referred to as *reciprocal mixing* and is illustrated in Figure 5.

FIGURE 5.

In a receiver, the phase noise of the LO can mix with a large interfering signal from a neighboring channel and swamp out the signal from the desired channel even though most of the power in the interfering IF is removed by the IF filter. This process is referred to as *reciprocal mixing*.



Similarly, phase noise in the signal produced by a nearby transmitter can interfere with the reception of a desired signal at a different frequency produced by a distant transmitter.

4.1 Feedback Oscillators

Consider a feedback oscillator with a loop gain of $H(f)$. $X(f)$ is taken to represent some perturbation stimulus and $Y(f)$ is the response of the oscillator to X . The Barkhausen condition for oscillation states that the effective loop gain equals unity and the loop phase shift equals 360 degrees at the oscillation frequency f_0 . The gain from the perturbation stimulus to the output is $Y(f)/X(f) = H(f)/H(f) - 1$, which goes to infinity at the oscillation frequency f_0 .

The amplification near the oscillation frequency is quantified by assuming the loop gain varies smoothly as a function of frequency in this region [10]. If $f = f_0 + \Delta f$, then $H(f) \approx H(f_0) + dH/df \Delta f$ and the transfer function becomes

$$\frac{Y(f + \Delta f)}{X(f + \Delta f)} \approx \frac{H(f) + dH/df \Delta f}{H(f) + dH/df \Delta f - 1}. \quad (2)$$

Since $H(f_0) = 1$ and $dH/df \Delta f \ll 1$ in most practical situations, the transfer function reduces to

$$\frac{Y(f + \Delta f)}{X(f + \Delta f)} \approx \frac{1}{dH/df \Delta f}. \quad (3)$$


Thus, for circuits that contain only white noise sources, the noise voltage (or current) is inversely proportional to Δf , while the noise power spectral density is proportional $1/\Delta f^2$ near the oscillation frequency.

So far we have assumed that the oscillator is linear time-invariant (LTI). This has allowed us to see that the amplification of noise near the carrier frequency is created by an LTI phenomenon that is a natural consequence of the oscillator's complex pole pair on the imaginary axis of the s -plane at f_0 . However, the LTI model does not explain why the noise is predominantly in the phase of the oscillation. Nor is it a good foundation for further analysis. It is easy to be misled by this model because it does not include effects that are fundamentally important to the behavior of the oscillator. To include these effects would require modeling the periodically time-varying nature of the transfer functions [15], which is beyond the scope of this paper. Instead, this model will be 'fixed-up' to explain phase noise with qualitative arguments and the next section presents a more solid and general model.

The Barkhausen criterion for oscillation in a feedback oscillator states that the effective gain around the loop must be unity for stable oscillation (loop gain magnitude equals 1 and loop phase shift equals 360°). To assure the oscillator starts, the initial loop gain is designed to be greater than one, which causes the oscillation amplitude to grow until the amplifier goes into compression far enough so that the effective loop gain reduces to 1. If, for some reason the amplitude of the oscillation decreases, the amount of compression reduces, causing the loop gain to go above 1, which causes the oscillation amplitude to increase. Similarly, if the oscillation amplitude increases, the amplifier goes further into compression, causing the loop gain to go below 1, which causes the amplitude to decrease. Thus, the nonlinearity of the amplifier is fundamental to providing a stable oscillation amplitude, and also causes amplitude variations to be suppressed. As shown in Figure 6, any amplitude variations that result from noise are also suppressed, leaving only phase variations. As a result, the noise at the output of an oscillator is generally referred to as oscillator phase noise.

FIGURE 6.

A linear oscillator along with its response to noise (left) and a nonlinear oscillator with its response to noise (right). For the nonlinear oscillator to have a stable amplitude, the average conductance exhibited by the nonlinear resistor must be negative below, positive above, and zero at the desired amplitude. The open-tipped arrows are phasors that represents the unperturbed oscillator output, the carriers, and the circles represent the response to perturbations in the form of noise. With a linear oscillator the noise simply adds to the carrier. In a nonlinear oscillator, the nonlinearities act to control the amplitude of the oscillator and so to suppress variations in amplitude, thereby radially compressing the noise ball and converting it into predominantly a variation in phase.




4.2 Oscillator Limit Cycle

The above explanation only addresses feedback oscillators. In this section, an alternative approach is taken that only assumes that the oscillator has a stable limit cycle and so applies to oscillators of all kinds.

FIGURE 7.

The trajectory of an oscillator shown in state space with and without a perturbation Δv . By observing the time stamps (t_0, \dots, t_6) one can see that the deviation in amplitude dissipates while the deviation in phase does not.



Consider plotting two state variables for an oscillator against each other, as shown in Figure 7. In steady state, the trajectory is a stable limit cycle, v . Now consider perturbing the oscillator with an impulse and assume that the response to the perturbation is Δv . Separate Δv into amplitude and phase variations,

$$\Delta v(t) = (I + \alpha(t))v(t + \phi(t)/2\pi f_o) - v(t). \tag{4}$$

where $v(t)$ represents the unperturbed output voltage of the oscillator, $\alpha(t)$ represents the variation in amplitude, $\phi(t)$ is the variation in phase, and f_o is the oscillation frequency.

Since the oscillation is stable and the duration of the disturbance is finite, the deviation in amplitude eventually decays away and the oscillator returns to its stable orbit ($\alpha(t) \rightarrow 0$ as $t \rightarrow \infty$). In effect, there is a restoring force that tends to act against amplitude noise.

This restoring force is a natural consequence of the nonlinear nature of the oscillator and at least partially suppresses amplitude variations.

The oscillator is autonomous, and so any time-shifted version of the solution is also a solution. Once the phase has shifted due to a perturbation, the oscillator continues on as if never disturbed except for the shift in the phase of the oscillation. There is no restoring force on the phase and so phase deviations accumulate. A single perturbation causes the phase to permanently shift ($\phi(t) \rightarrow \Delta\phi$ as $t \rightarrow \infty$). If we neglect any short term time constants, it can be inferred that the impulse response of the phase deviation $\phi(t)$ can be approximated with a unit step $s(t)$. The phase shift over time for an arbitrary input disturbance u is

$$\phi(t) \sim \int_{-\infty}^{\infty} s(t-\tau)u(\tau)d\tau = \int_{-\infty}^t u(\tau)d\tau, \quad (5)$$

or the power spectral density (PSD) of the phase is

$$S_{\phi}(\Delta f) \sim \frac{S_u(\Delta f)}{(2\pi\Delta f)^2} \quad (6)$$

This shows that in all oscillators the response to any form of perturbation, including noise, is amplified and appears mainly in the phase. The amplification increases as the frequency of the perturbation approaches the frequency of oscillation. Various approaches are available to improve the relative noise performance of the oscillator, such as using a resonator with a higher Q , increasing the output signal level relative to the noise (increases power dissipation), or using cleaner devices. However the $1/\Delta f^2$ amplification of noise that occurs in oscillators can only be removed by constraining the phase of the oscillator. This is accomplished by entraining the oscillator to another, cleaner signal, either by injection locking it to that signal, or by embedding it in a phase-locked loop for which that signal is the reference.


4.3 Oscillator Voltage Noise and Phase Noise Spectra

There are two different ways commonly used to characterize noise in an oscillator. S_{ϕ} is the spectral density of the phase and S_v is the spectral density of the voltage. S_v contains both amplitude and phase noise components, but with oscillators the phase noise dominates except at frequencies far from the carrier and its harmonics. S_v is directly observable on a spectrum analyzer, whereas S_{ϕ} is only observable if the signal is first passed through a phase detector. Another measure of oscillator noise is L , which is simply S_v normalized to the power in the fundamental.

As $t \rightarrow \infty$ the phase of the oscillator drifts without bound, and so $S_{\phi}(\Delta f) \rightarrow \infty$ as $\Delta f \rightarrow 0$. However, even as the phase drifts without bound, the excursion in the voltage is limited by the diameter of the limit cycle of the oscillator. Therefore, as $\Delta f \rightarrow 0$ the PSD of v flattens out, as shown in Figure 8. The more phase noise, broader the linewidth (the higher the corner frequency), and the lower signal amplitude within the linewidth. This happens because the phase noise does not affect the total power in the signal, it only affects its distribution. Without noise, $S_v(f)$ is a series of impulse functions at the harmonics of the oscillation frequency. With noise, the impulse functions spread, becoming fatter and shorter but retaining the same total power.

FIGURE 8.

Two different ways of characterizing noise in the same oscillator. S_ϕ is the spectral density of the phase and S_v is the spectral density of the voltage. S_v contains both amplitude and phase noise components, but with oscillators the phase noise dominates except at frequencies far from the carrier and its harmonics. S_v is directly observable on a spectrum analyzer, whereas S_ϕ is only observable if the signal is first passed through a phase detector.



The voltage noise S_v is considered to be a small signal outside the linewidth and thus can be accurately predicted using small-signal analyses. Conversely, the voltage noise within the linewidth is a large signal (it is large enough to cause the circuit to behave nonlinearly) and cannot be predicted with small-signal analyses. Thus, small-signal noise analysis, such as is available from RF simulators, is valid only up to the corner frequency (it does not model the corner itself).

4.4 Oscillators and Frequency Correlation

With driven cyclostationary systems that have a stable time reference, the correlation in frequency is a series of impulse functions separated by $f_0 = 1/T$. Thus, noise at f_1 is correlated with f_2 if $f_2 = f_1 + kf_0$, where k is an integer, and not otherwise. However, the phase produced by oscillators that exhibit phase noise is not stable. And while the noise produced by oscillators is correlated across frequency, the correlation is not a set of equally spaced impulses as it is with driven systems [3]. Instead, the correlation is a set of smeared impulses. That is, noise at f_1 is correlated with f_2 if $f_2 = f_1 + kf_0$, where k is close to being integer.

Technically, the noise produced by oscillators is not cyclostationary [1]. This distinction only becomes significant when the output of an oscillator is compared to its own output from the distant past. This might occur, for example, in a radar system where the current output of an oscillator might be mixed with the previous output after it was delayed by traveling to and from a distant object. It occurs because the phase of the oscillator has drifted randomly during the time-of-flight. If the time-of-flight is long enough, the phase difference between the two becomes completely randomized and the two signals can be treated as if they are non-synchronous (see Section 3.1 on page 6). Thus, the noise in the return signal can be taken as being stationary because it is ‘non-synchronous’ with the LO, even though the return signal and the LO are derived from the same oscillator. If the time-of-flight is very short, then there is no time for the phase difference between the two to become randomized and the noise is treated as if it is simply cyclostationary. Finally, if the time-of-flight significant but less than the time it takes the oscillator’s phase to become completely randomized, then the phase is only partially randomized. In this case, one must be careful to take into account the smearing in the correlation spectrum that occurs with oscillators. Because of these difficulties in interpreting the oscillator frequency spectrum, it is wise to refer to the time-domain model implied in (4) when interpreting noise from autonomous oscillators.

4.5 Phase Noise Calculations

To see how oscillator phase noise can be calculated, consider the effect of a small phase perturbation on the oscillator signal. With observation times that are short (in other words, if we do not attempt to resolve frequencies to within the linewidth of the oscillator), we can linearize (4) to obtain

$$\Delta v(t) = \frac{dv(t)}{dt} \frac{\Delta\phi(t)}{2\pi f_0}. \quad (7)$$

This equation simply says that phase perturbations are those that align with tangential perturbations to the oscillator limit cycle. To analyze phase noise, we must determine how much each noise source contributes to perturbations in the oscillator state along the direction of the limit-cycle tangent. Because noise perturbations that contribute to tangential movements are *not*, in the general case, strictly tangential, accurate oscillator noise analysis requires some rather involved linear algebraic calculations [15] that are derived from Floquet theory.

5.0 Jitter

Jitter is an undesired fluctuation in the timing of events. One models jitter in a signal by starting with a noise-free signal v and displacing time with a stochastic process j . The noisy signal becomes

$$v_j(t) = v(t + j(t)). \quad (8)$$

Jitter is equivalent to phase noise in (4) where $j = \phi/2\pi f_0$. It is used in situations where it is more natural to think of the noise being in the timing of events rather than in the phase or in the signal level.

5.1 Sources of Jitter

In systems where signals are continuous valued, an event is usually defined as a signal crossing a threshold in a particular direction. The threshold crossings of a noiseless periodic signal, v , are precisely evenly spaced. However, when noise is added to the signal,

$$v_n(t) = v(t) + n(t), \quad (9)$$

each threshold crossing is displaced slightly. Thus, a threshold converts additive noise to jitter. This is the way jitter is created in nonlinear circuits such as logic circuitry.

The noise n and the jitter j can be related by expanding (8) into a Taylor series, setting $v_n(t) = v_j(t)$, and dropping the high order terms,

$$v(t) + n(t) = v(t + j(t)) = v(t) + \frac{dv(t)}{dt} j(t) + \dots, \quad (10)$$

$$n(t) \cong \frac{dv(t)}{dt} j(t). \quad (11)$$

Then, the variance in the time of the threshold crossing is

$$\text{var}(j(t_c)) \cong \frac{\text{var}(n(t_c))}{\left(\frac{dv(t_c)}{dt}\right)^2}, \quad (12)$$

where t_c is the expected time of the threshold crossing.

Another important source of jitter is oscillator phase noise. To predict the jitter in an oscillator, assume that u in (5) is a white stationary process and define a such that

$$S_\phi(f) = a \frac{f_0^2}{f^2}, \quad (13)$$

where $f_0 = 1/T$ is the oscillation or carrier frequency. Demir [1] shows that the variance of the length of a single period is aT . The variance of the length of each period is uncorrelated and so the variance in the length of k periods is simply k times the variance of one period. The jitter J_k is the standard deviation of the length of k periods, and so

$$J_k = \sqrt{kaT}. \quad (14)$$

In the case where u represents flicker noise, $S_u(f)$ is generally pink or proportional to $1/f$. Then $S_\phi(f)$ would be proportional to $1/f^3$ at low frequencies [6]. In this case, there are no explicit formulas for J_k .

5.2 Effect of Jitter

Jitter in the time at which a signal is sampled creates noise in the result if the signal is changing at the time when it is sampled. This is one way in which noise is generated when converting continuous-time signals to discrete-time signals. Using (11), the variance of the noise can be computed from the variance of the jitter at the time of the sampling and the slewrate (or time derivative) of the input signal at the time of the sampling.

$$\text{var}(n(t_s)) \cong \left(\frac{dv(t_s)}{dt}\right)^2 \text{var}(j(t_s)), \quad (15)$$

If one samples a constant valued signal, jitter in the time at which the sampling occurs does not create noise in the output. Thus, during flat portions of waveforms, an uncertainty in the sampling time creates no noise

6.0 Noise and Jitter in Logic Circuits

Logic circuits are thresholding circuits and so ignore noise at the input when the input signal is far from the threshold. As such, logic circuits are only sensitive to noise at an input when that input is undergoing a transition. Similarly, logic circuits produce their highest noise levels at the output when the output is transitioning. Because of the strong variability in both the level of noise produced at the output and the sensitivity to noise at the input, traditional approaches to describing noise, such as signal-to-noise ratio, are not very helpful when working with logic circuits. Instead, it is best to characterize the noise in terms of jitter. Once the jitter is known for the logic blocks that make up a system, it is generally relatively straight-forward to compute the jitter of the system (the


variance of the jitter for a cascade of uncorrelated jitter sources is simply the sum of the variance of the jitter of each source individually). The difficulty, of course, is determining the jitter of the individual blocks.

6.1 Cyclostationary Noise from Logic Circuits

The noise produced by a logic circuit, such as the inverter shown in Figure 9, comes from different places depending on the phase of the output. When the output is high, the output is insensitive to small changes on the input. The transistor M_P is on, however, and the noise at the output is predominantly due to the thermal noise from its channel. When the output is low, the situation is reversed and most of the output noise is due to the thermal noise from the channel of M_N . When the output is transitioning, thermal noise from both M_P and M_N contribute to the output. In addition, the output is sensitive to small changes in the input. In fact, any noise at the input is amplified before reaching the output. Thus, noise from the input tends to dominate over the thermal noise from the channels of M_P and M_N in this region. Noise at the input includes noise from the previous stage and thermal noise from the gate resistance. In addition, with significant current flowing in the transistors, flicker noise from the channel also contributes.

FIGURE 9.

Schematic of a inverter.



6.2 Characterizing the Jitter of a Logic Circuit

One can apply (12) to compute jitter of this circuit. To do so, one must drive the circuit with a representative periodic signal while accurately modeling the input source and output load, both of which are typically other logic circuits. Both the slewrate and the noise must be determined at the time of the threshold crossing. This last point is very important. The total output noise power of a logic circuit would be dominated by the thermal noise produced by the output devices if the circuit spends most of its time with an unchanging output. This noise is usually ignored by subsequent stages and does not contribute to jitter. Thus, using the time-averaged spectral density to characterize the noise in a logic circuit is misleading. Only the noise produced by a circuit at the point where its output crosses the threshold of the subsequent stage should be taken into account when characterizing the jitter of a logic circuit.

There are several different ways of determining the noise produced by a logic circuit at the time when its output crosses the threshold. All assume the availability of a circuit simulator that can perform a cyclostationary noise analysis. If the simulator can directly compute the noise level as a function of time, it is a simple matter to determine the time of the threshold crossing and use noise computed for that time. If the noise is output as a spectral density, it is necessary to integrate the noise over all frequencies to determine the total noise before applying (12). If the simulator can only compute the time-average noise, one can use a limiter or a sample-and-hold to isolate the noise at the threshold

crossing [7]. Each of these approaches make assumptions as to how sensitive a subsequent stage will be to noise produced away from the threshold. If the simulator is capable of producing a summary of noise contributions from each noise source, then an alternative approach would be to simulate both stages together and use the above techniques to measure the jitter at the output of the subsequent stage. When applying (12), only include the output noise contributed by noise sources within the stage being characterized. In this way both the loading and the noise sensitivity of the subsequent stage are accurately modeled. It is also possible and desirable to include a representative driver stage. Noise generated by the driver and load stages are ignored by this method.

7.0 If You Have Questions

If you have questions about what you have just read, feel free to post them on the *Forum* section of *The Designer's Guide Community* website. Do so by going to designers-guide.org/forum.

References

- [1] A. Demir, A. Mehrotra, and J. Roychowdhury. Phase noise in oscillators: a unifying theory and numerical methods for characterisation. *Proceedings of the 35th Design Automation Conference*, June 1998.
- [2] A. Demir and A. Sangiovanni-Vincentelli. *Analysis and Simulation of Noise in Nonlinear Electronic Circuits and Systems*. Kluwer Academic Publishers, 1997.
- [3] W. Gardner. *Introduction to Random Processes with Applications to Signals and Systems*. McGraw-Hill, 1989.
- [4] W. A. Gardner. *Statistical Spectral Analysis: A Nonprobabilistic Theory*. Prentice-Hall, 1987.
- [5] W. A. Gardner. Exploitation of spectral redundancy in cyclostationary signals. *IEEE Signal Processing Magazine*, Vol. 8, No. 2, pp. 14-36, 1991.
- [6] F. Käertner. Analysis of white and $f^{-\alpha}$ noise in oscillators. *International Journal of Circuit Theory and Applications*, vol. 18, pp. 485-519, 1990.
- [7] K. Kundert. Predicting the phase noise and jitter of PLL-based frequency synthesizers. Available from designers-guide.org/analysis.
- [8] Ken Kundert. Introduction to RF simulation and its application. *Journal of Solid-State Circuits*, vol. 34, no. 9, September 1999. Available from designers-guide.org/analysis.
- [9] M. Okumura, T. Sugawara and H. Tanimoto. An efficient small signal frequency analysis method for nonlinear circuits with two frequency excitations. *IEEE Transactions of Computer-Aided Design of Integrated Circuits and Systems*, vol. 9, no. 3, pp. 225-235, March 1990.
- [10] B. Razavi. A study of phase noise in CMOS oscillators. *IEEE Journal of Solid-State Circuits*, vol. 31, no. 3, pp. 331-343, March 1996.

- [11] W. Robins. *Phase Noise in Signal Sources (Theory and Application)*. IEE Telecommunications Series, 1996.
- [12] J. Roychowdhury, D. Long and P. Feldmann. Cyclostationary noise analysis of large RF circuits with multi-tone excitations. *IEEE Journal of Solid-State Circuits*, March 1998.
- [13] M. Terrovitis, R. Meyer. *Cyclostationary Noise in Communication Systems*. ERL memo M99/36, 1999. Available through Electronics Research Laboratory Publications, University of California, Berkeley.
- [14] R. Telichevesky, K. Kundert and J. White. Efficient AC and noise analysis of two-tone RF circuits. *Proceedings of the 33rd Design Automation Conference*, June 1996.
- [15] J. Phillips, D. Feng, K. Kundert and J. White. Numerical characterization of noise in oscillators. *International Symposium on the Mathematical Theory of Networks and Systems*, July, 1998.