Noise in Mixers, Oscillators, Samplers, and Logic

An Introduction to Cyclostationary Noise

Joel Phillips and Ken Kundert

Version 2, 9 September 2001

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.

This document contains the slides and speaker notes for a presentation on cyclostationary noise given at the Custom Integrated Circuits Conference in May of 2000. The companion paper 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.

1.0 Motivation

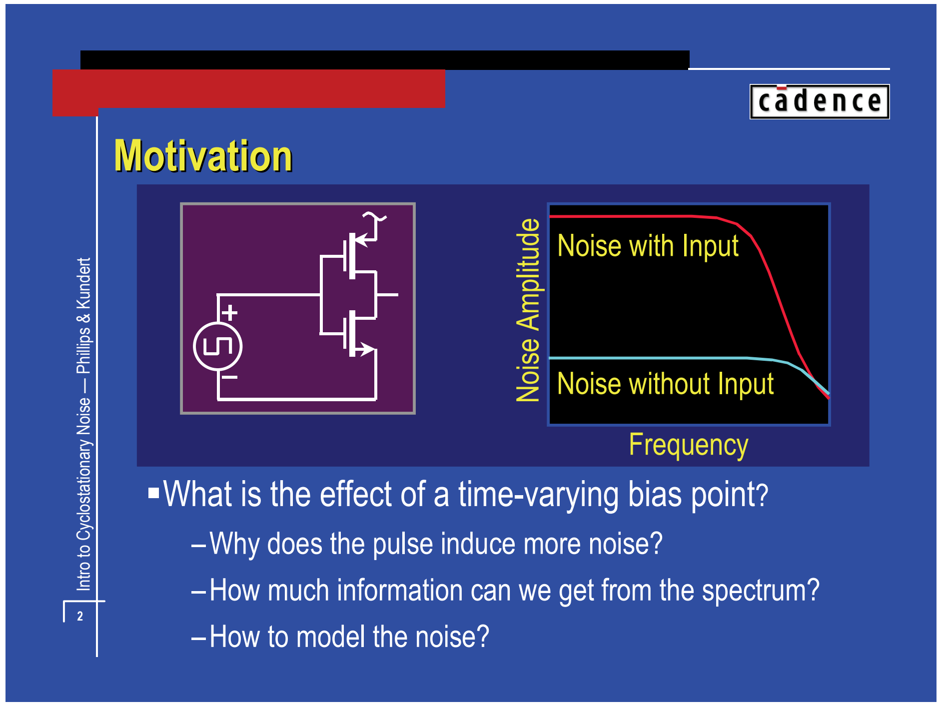

SPICE noise analysis is not able to compute valid noise results for many common classes of circuit for which noise is of interest. Noise in these circuits is of different character than that can be computed by SPICE. We show why and how in this presentation. We explain how this noise differs from the noise discussed in introductory text books on noise, and how this difference affect your circuits.

2.0 What is Noise?

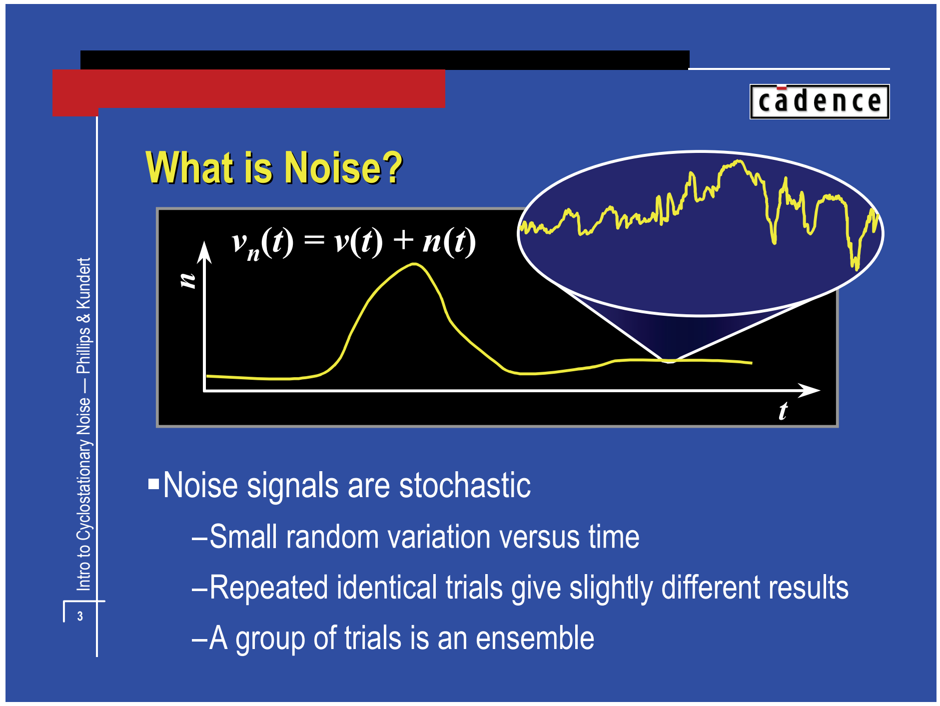

Noise free systems are deterministic, meaning that repeating the same experiment produces the same result. Noisy systems are stochastic — repeating the same experiment produces slightly different results each time. An experiment is referred to as a trial. A group of experiments is referred to as an ensemble of trials, or simply an ensemble.

3.0 Ensemble Averages

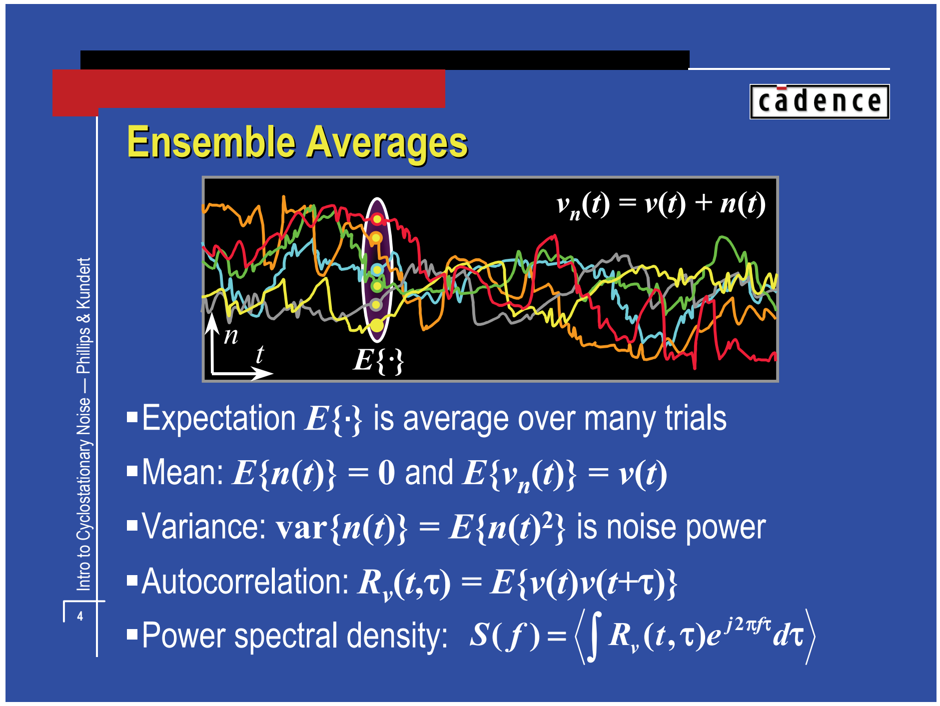

Assume that $v_n(t) = v(t) + n(t)$, where $v(t)$ is the deterministic signal, $n(t)$ is the noise, and $v_n(t)$ is the combined signal.

Noise is characterized using ensemble averages. An average over many trials is referred to as an expectation, and is denoted $E\{\cdot\}$. One ensemble average is the mean. The mean of the combined signal is an estimate of the noise free signal, $E\{v_n(t)\} = v(t)$. The mean of the noise alone is generally 0, $E\{n(t)\} = 0$. The variance, $\mathrm{var}(n(t)) = E\{n(t)^2\} - E\{n(t)\}^2$, is a measure of the power in the noise. The autocorrelation, $R_n(t, \tau) = E\{n(t)\, n(t-\tau)\}$, is a measure of how points on the same signal separated by τ seconds are correlated. The autocorrelation is related to the variance by $\mathrm{var}(n(t)) = R_n(t, 0)$. The Fourier transform of the autocorrelation function averaged over t is the time-averaged power spectral density, or PSD. It is the power in the noise as a function of frequency.

4.0 Cyclostationary Noise

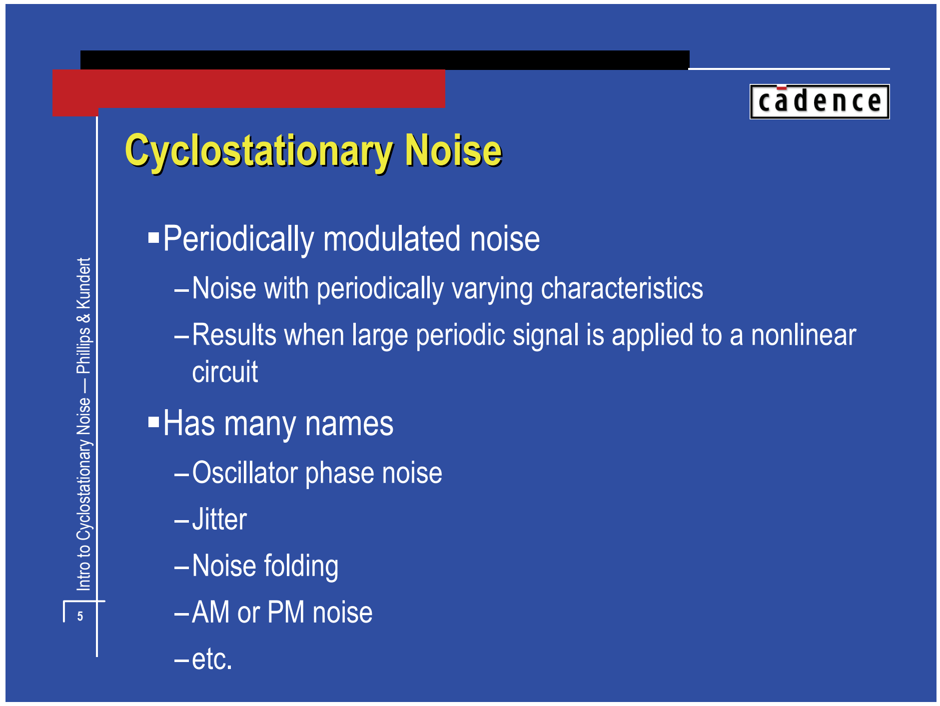

Though the term *cyclostationary noise* may be unfamiliar, it is critically important to most designers. They typically know it by other names: oscillator phase noise, jitter, noise folding, AM & PM noise, etc. By looking at noise using the common theoretical foundation of cyclostationary noise allows designers to better understand and predict all of these manifestations of noise.

5.0 White Noise

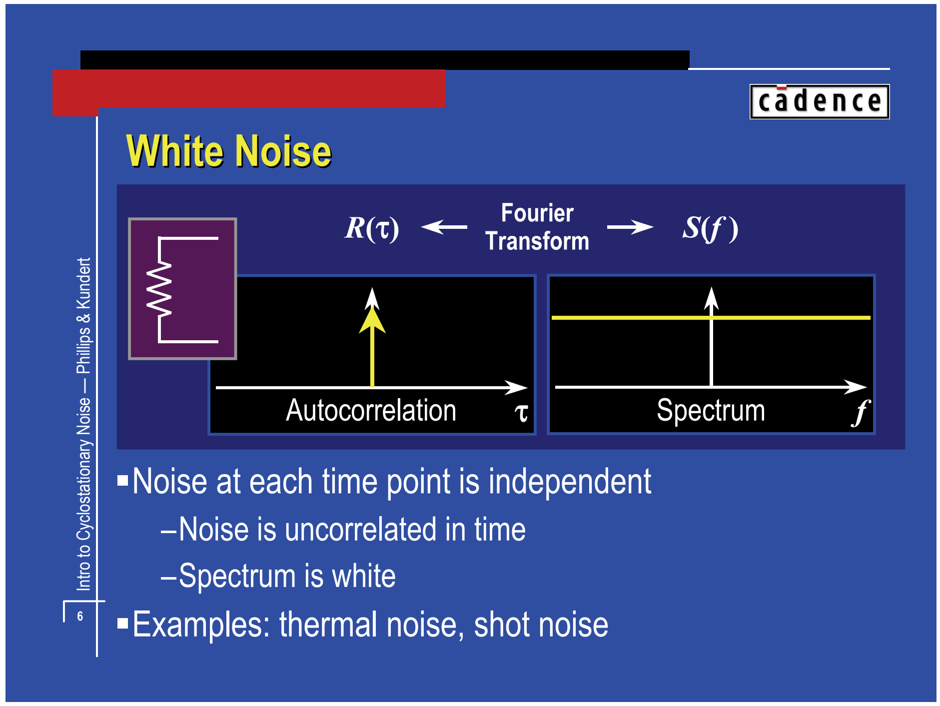

There are two ways of characterizing noise, the power spectral density or PSD, $S(f)$, and the autocorrelation function, $R(t,\tau)$. These two representations are related by the Fourier transform.

Completely uncorrelated noise is known as white noise. For white noise the PSD is a constant and the autocorrelation function is an impulse function centered at 0.

6.0 Colored Noise

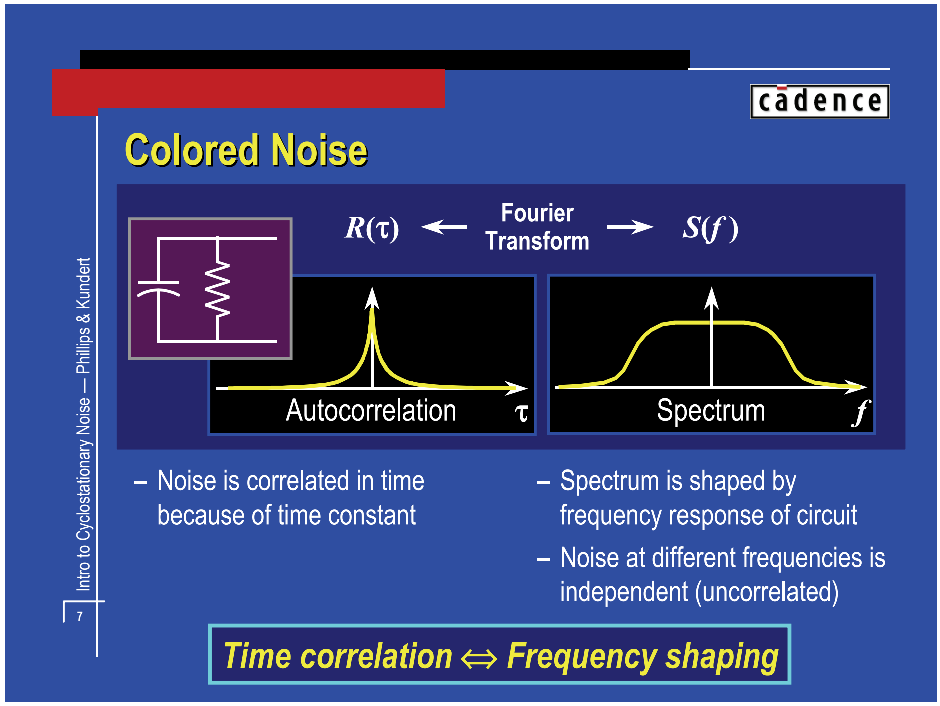

Energy storage elements cause the circuit to exhibit a frequency response that is shaped; it is not flat with frequency. As such, it causes the PSD to be shaped. This is referred to as coloring the noise. Colored noise has a PSD that varies with frequency.

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

Notice that in this case, both the spectrum is shaped and the noise is correlated over time. This is a general property, shaping the noise in the frequency domain implies that the noise is correlated in time, and visa versa.

7.0 Cyclostationary Noise

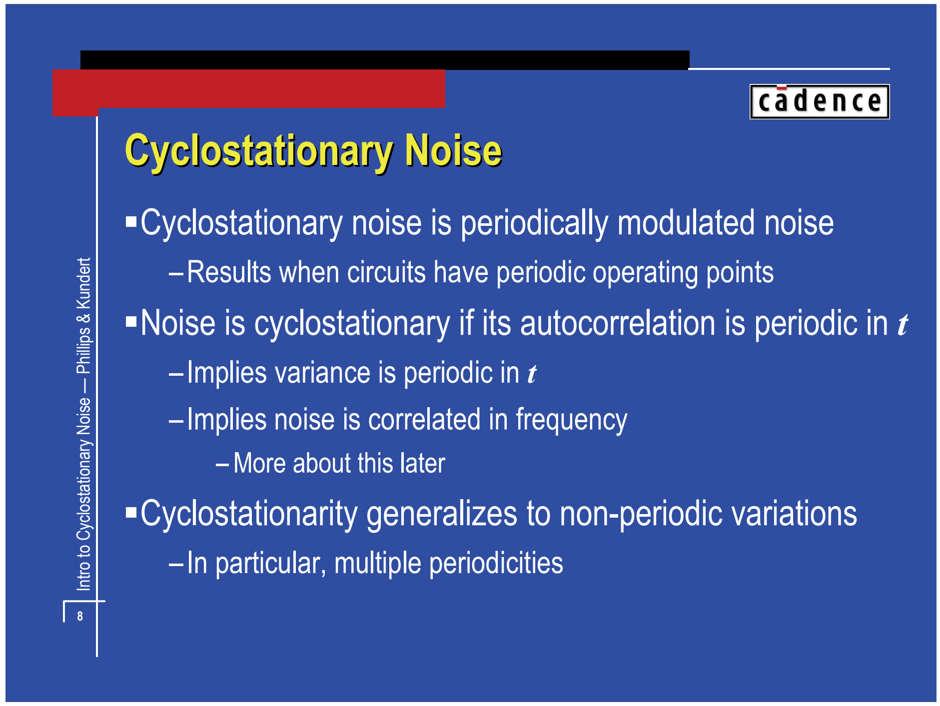

Cyclostationary noise results when noise is periodically modulated. This occurs in any nonlinear circuit that is driven by, or exhibits, large periodic signals.

Noise is cyclostationary if, and only if, its autocorrelation function $R(t,\tau)$ is periodic in t. Since the variance is $R(t,0)$, the instantaneous noise power also periodic in t. Also shown later in this presentation is that noise being cyclostationary also implies that the noise is correlated versus frequency.

Cyclostationarity naturally generalizes to non-periodic variation, in particular, when there are multiple periodicities (also known as quasiperiodicity).

8.0 Origins of Cyclostationary Noise

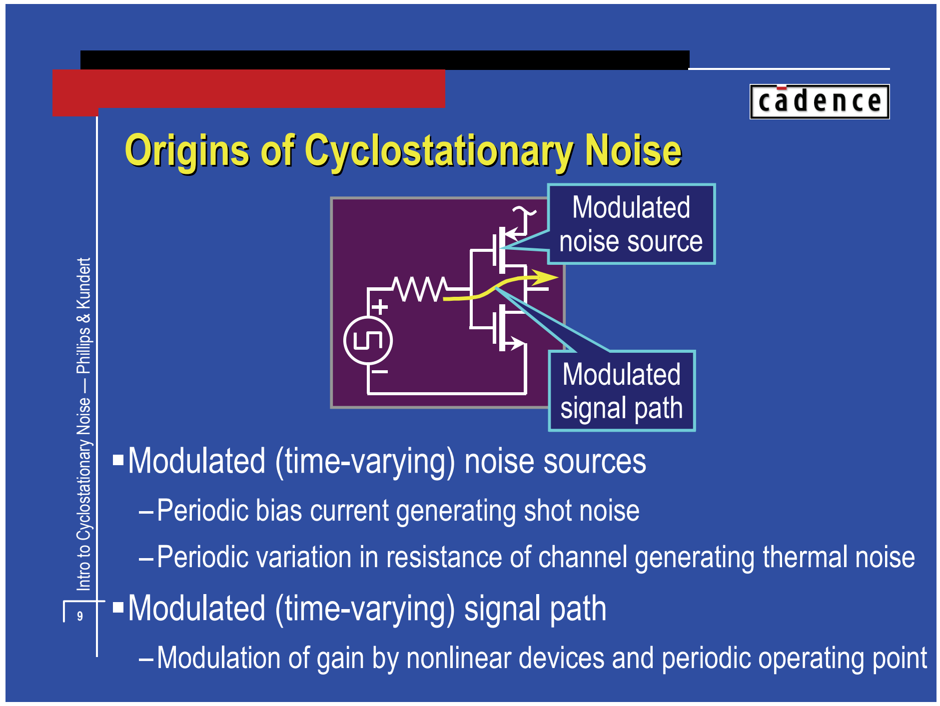

Cyclostationary noise is generated by circuits with periodic operating points. The time-varying operating point modulates the noise generated by bias-dependent noise sources, and modulates the transfer function from the noise source to the output. Both result in cyclostationary noise at the output.

As suggested by the name, modulated noise sources can be modeled by modulating the output of stationary noise sources.

9.0 Cyclostationary Noise vs. Time

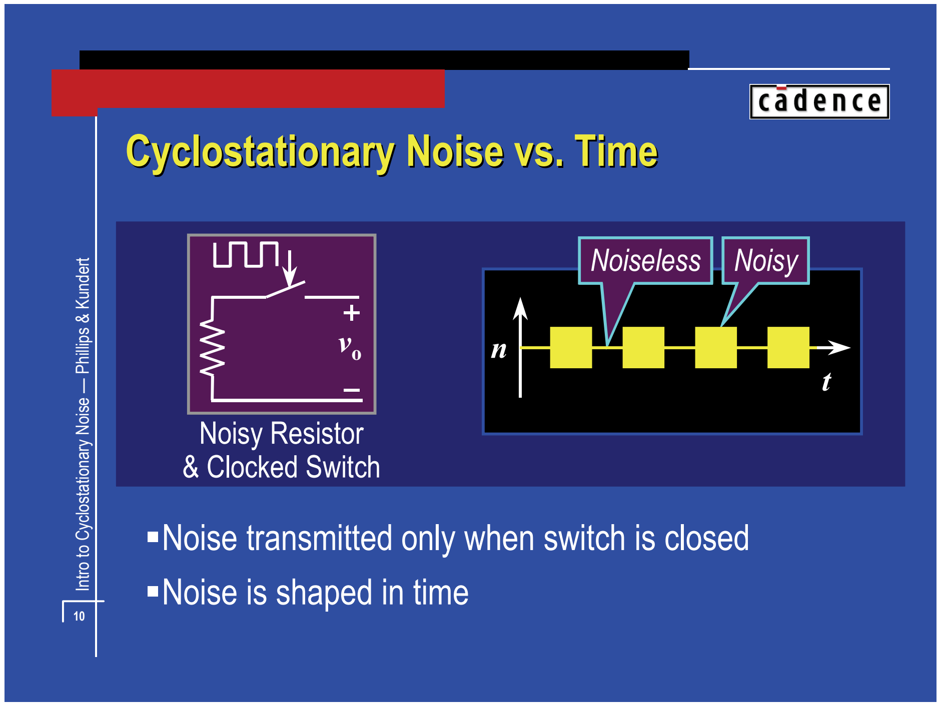

This is a simple example of cyclostationary noise. A periodically operating switch between the noise source (the resistor produces white thermal noise) and the observer causes the output noise to vary periodically.

It can be said that cyclostationary noise is “shaped in time”.

10.0 Cyclostationary Noise vs. Frequency

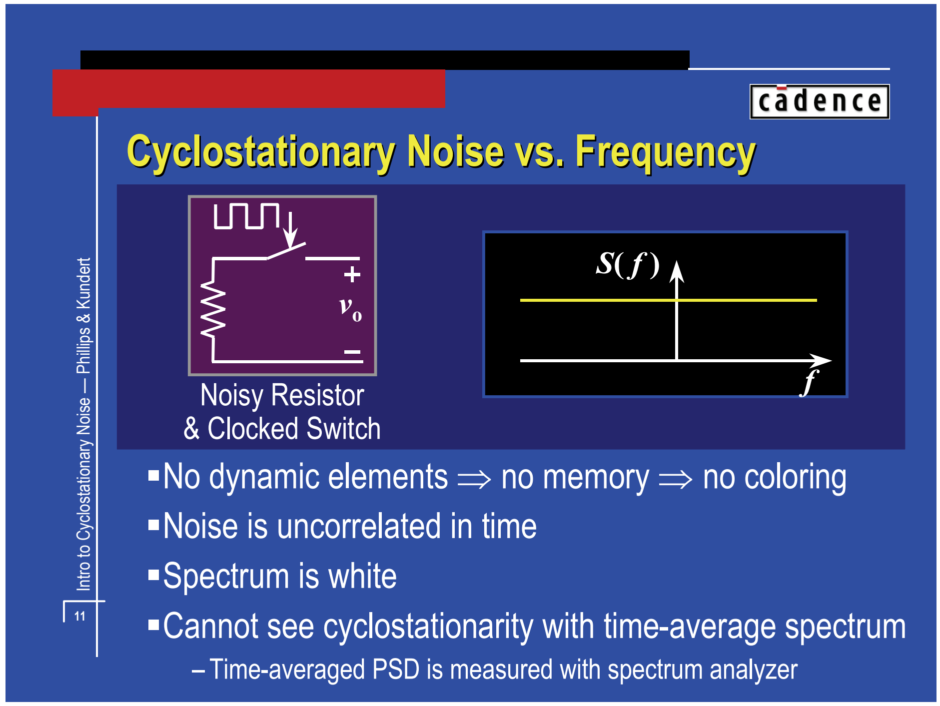

With no energy storage elements the noise is completely uncorrelated (noise at a particular time is uncorrelated with the noise at any previous time) and therefore is white, even though it is cyclostationary. One cannot tell that noise is cyclostationary by just observing a simple PSD.

The traditional form of a PSD is referred to as the *time-averaged PSD*. It is what is measured by a spectrum analyzer when the modulation frequency of the cyclostationary noise is well beyond the resolution bandwidth of the analyzer, meaning that it cannot track the periodic variations of the noise. As a result, it presents the time-average spectrum of the noise.

11.0 Cyclostationary Noise vs. Time & Frequency

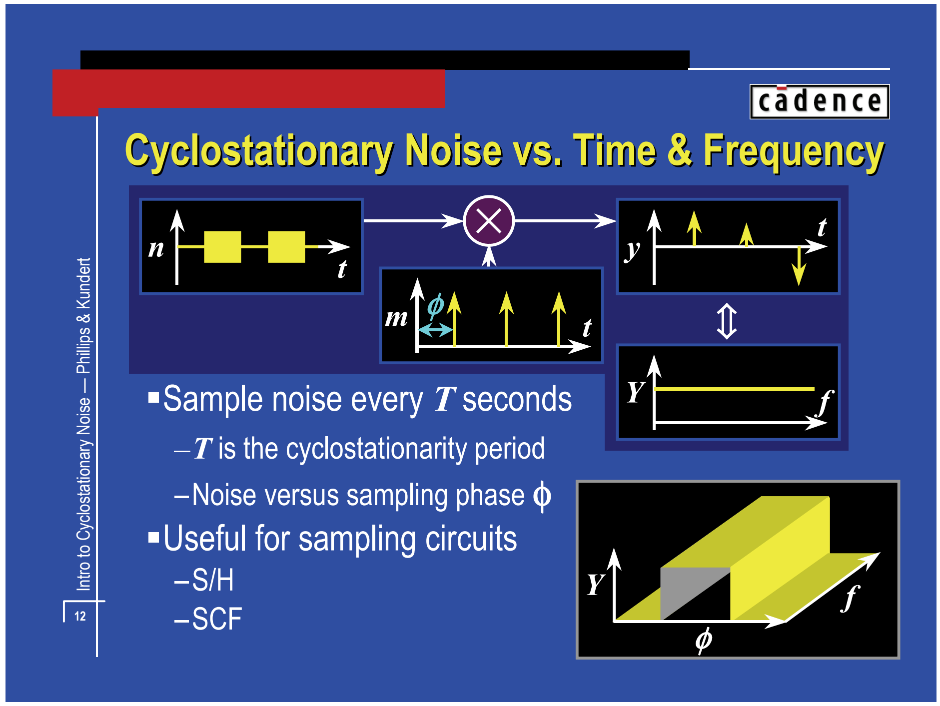

An alternative to the time-average PSD is the instantaneous PSD. It captures the time-varying nature of cyclostationary noise. One conceptually measures a instantaneous PSD by sampling the cyclostationary noise at the noise modulation frequency. The sampling process must be in exact synchronism with the noise modulation process and be offset in phase by ϕ. The result will be a random discrete-time noise process for which we will compute the PSD. As a side note, since it is a discrete-time random process, its spectrum is periodic in f_0, the sampling frequency. The phase offset f is then swept over an entire cycle and the process repeated. The result is a family of PSDs for the sampled noise parameterized over f — the instantaneous PSD.

Because of the sampling nature of the sample & holds and switched-capacitor filters, the instantaneous PSD is often the best way of characterizing noise in these circuits.

12.0 Cyclostationary Noise is Modulated Noise

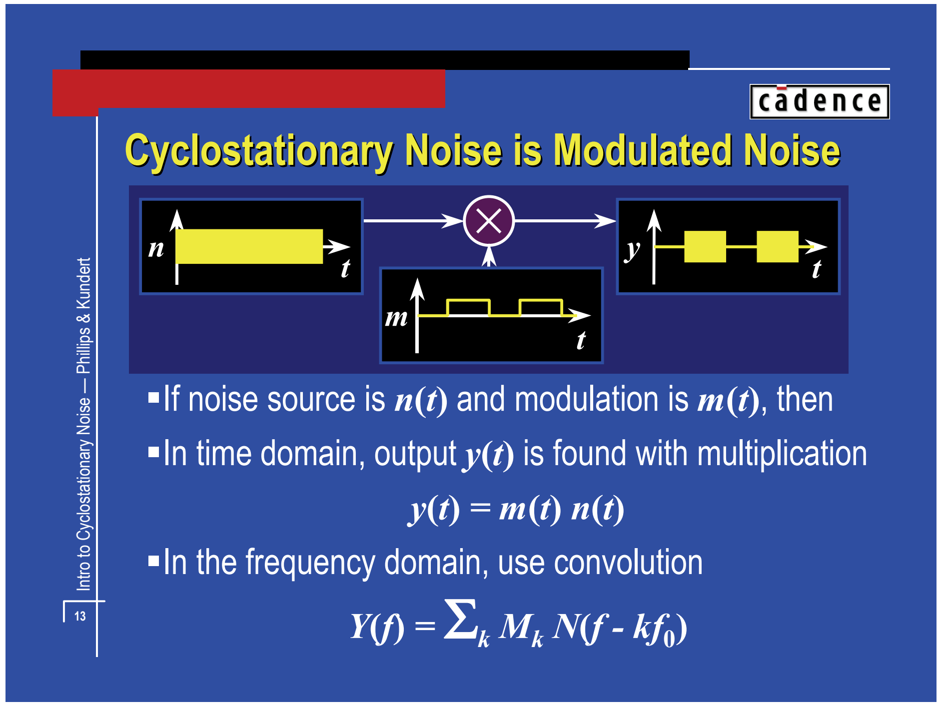

As mentioned before, cyclostationary noise is another name for modulated noise. Cyclostationary noise can be created by periodically modulating stationary noise. Similarly, cyclostationary noise can always be decomposed into, or modeled as, a stationary noise and a periodic modulation.

In the example shown, the stationary noise is modulation by a periodic pulse train. Recall that in the time domain modulation is performed by simply multiplying the stationary noise by the modulation function. In the frequency domain, the two signals are convolved.

13.0 Modulated Noise Spectrum

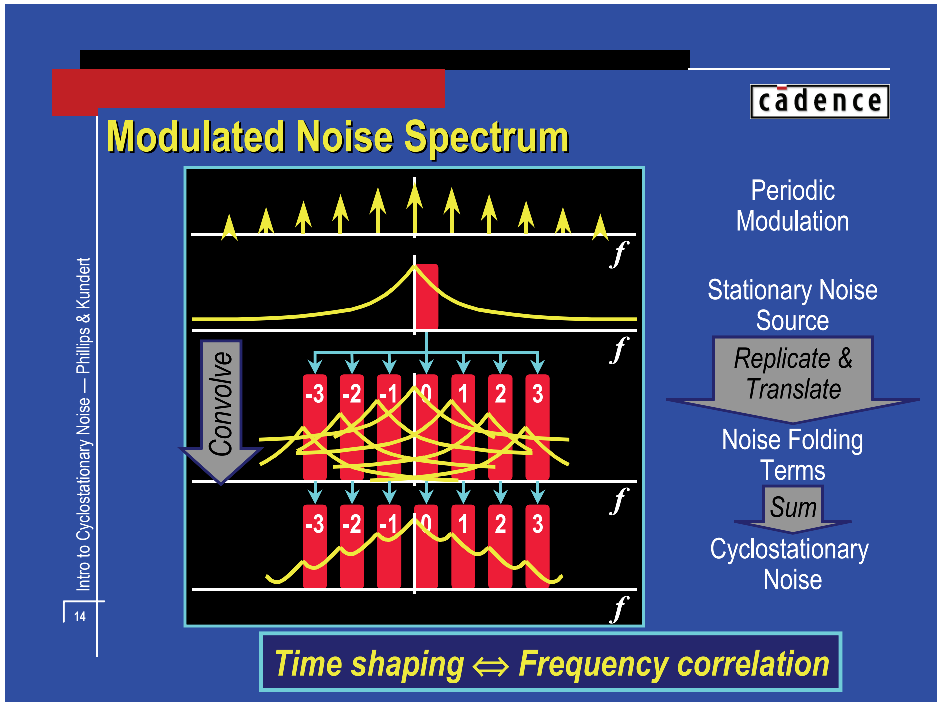

In the example, stationary noise with an arbitrary spectrum is modulated by some periodic signal. This is representative of both ways in which cyclostationary noise is generated (modulated noise sources and modulated signal paths).

Modulation is a 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.

Noise from the source at a particular frequency f is replicated and copies appear at $f \pm kf_0$ where f_0 is the modulation frequency. Conversely, noise at the output at a particular frequency f has contributions from noise from the sources at frequencies $f \pm kf_0$.

Modulation acts to shape the noise in the time-domain and correlate the noise in the frequency-domain.

14.0 Duality of Shape and Correlation

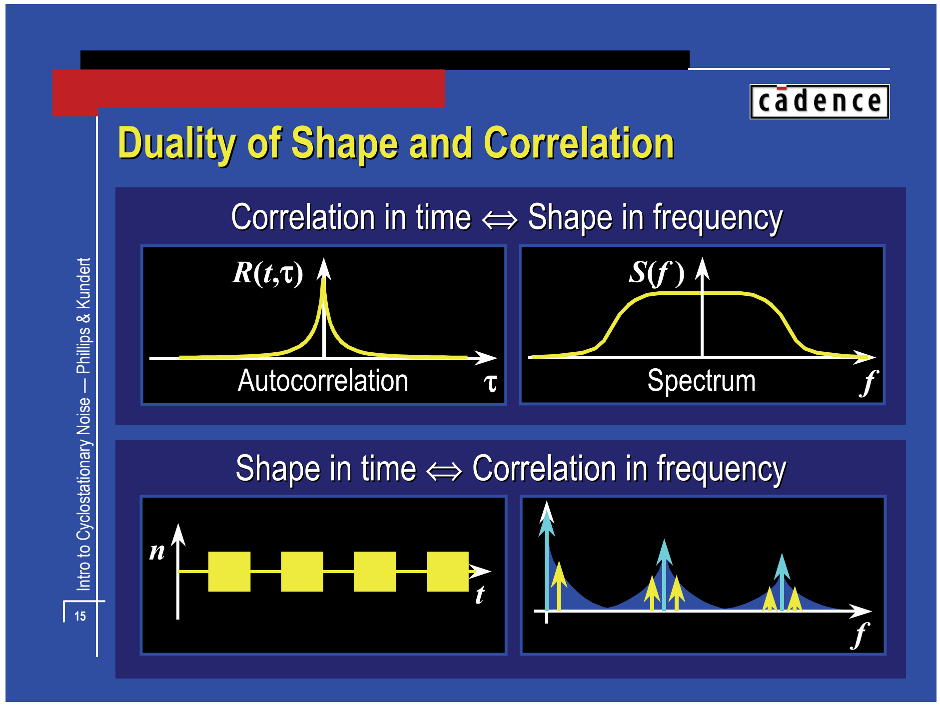

Let's review …

Recall that

shape in frequency ↔ *correlation in time*

Now also see that

shape in time ↔ *correlation in frequency*

This is the duality of shape and correlation. These relationships occur because in both cases the quantities on the left and right are related by the Fourier transform.

15.0 Correlations in Cyclostationary Noise

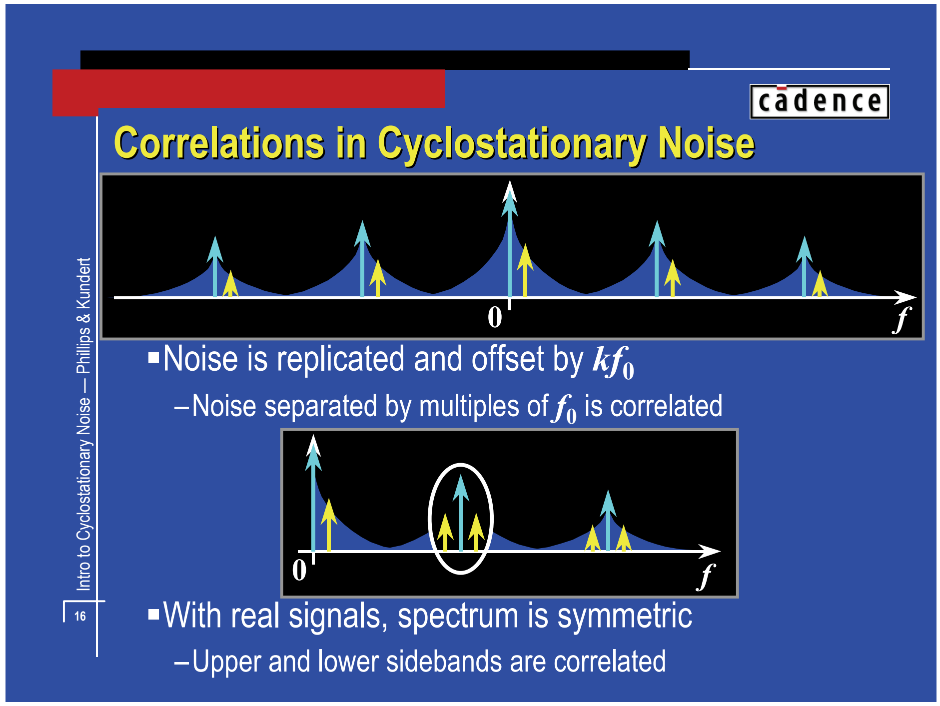

As just shown, modulation causes noise to be replicated and translated in frequency. Thus, noise separated by kf_0 is correlated where f_0 is the modulation frequency. Remember, noise folds across DC in a conjugate symmetric manner, so noise in upper and lower sidebands are correlated. In other words, in the top diagram 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 get 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.

16.0 Sidebands and Modulation

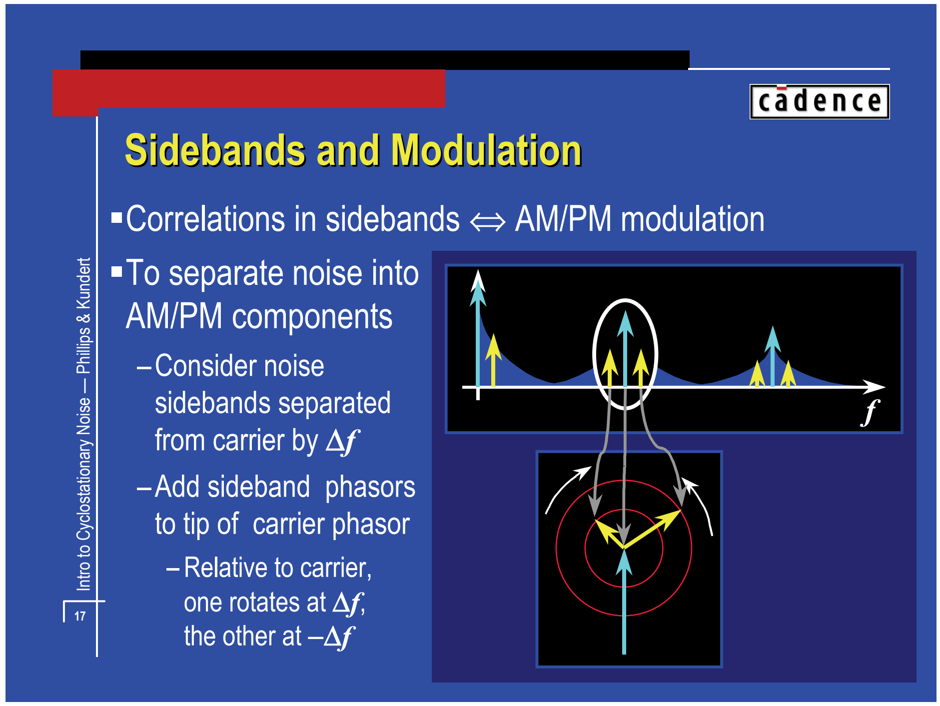

One can separate noise near the carrier into AM and PM components. Consider the noise sidebands at frequencies Δf from the carrier. Treat both these sidebands and the carrier as phasors. Individually add the sideband phasors to the carrier phasor. The sideband phasors are at a different frequency from the carrier, and so will rotate relative to it. One sideband will rotate at Δf, and the other at $-\Delta f$. 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, it the noise is cyclostationary, then the sidebands are correlated. This reduces the 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 remain unchanged, though its size still shifts randomly.

One can then decompose the ellipse into orthogonal components. The component parallel to the carrier phasor represents the amplitude modulation (AM noise) and the component perpendicular to the carrier represents the phase modulation (PM noise).

17.0 AM/PM Noise

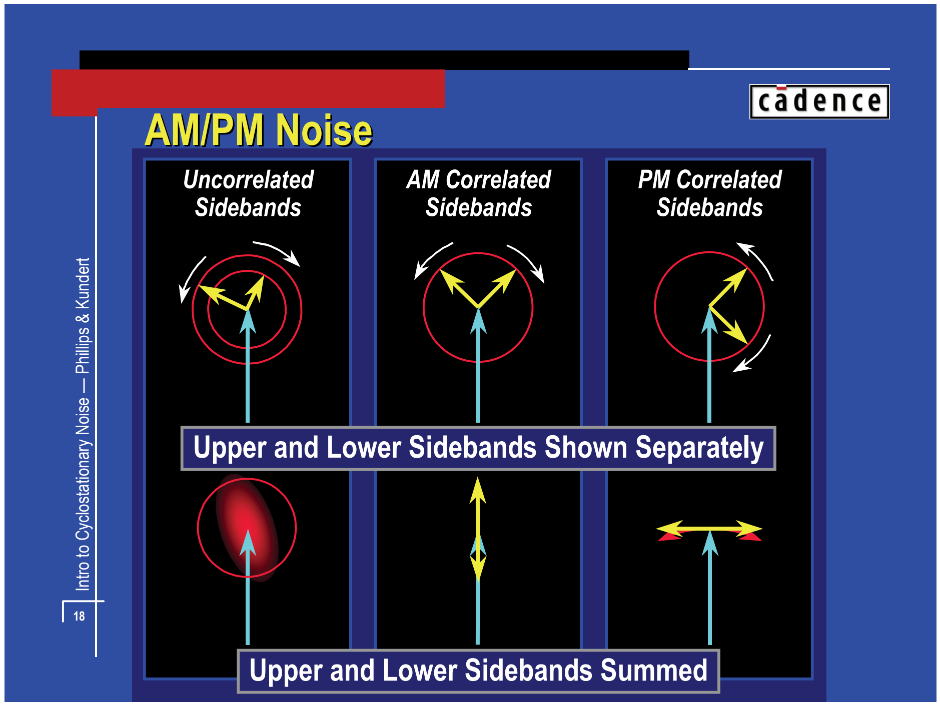

With cyclostationary noise, the noise in various sidebands is correlated. Depending on the magnitude and phase of the correlation, the noise at the output of the circuit can be AM noise, PM noise, or some combination of the two. For example, oscillators almost exclusively generate PM noise near the carrier whereas noise on the control input to a variable gain amplifier results almost completely in AM noise at the output of the amplifier.

When considering the noise about a carrier frequency, the noise can decomposed into AM and PM components. 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.

There are three cases illustrated. The first represents the case where stationary noise is added to the carrier. In this case the sidebands will have the same variance (average length) by are completely uncorrelated. The length and phase randomly vary, and in the short term trace out an ellipse that wobbles and changes it shape. But over the long term the noise envelope is a circle, which implies that there is an equal amount of AM and PM in the noise. This is always true with stationary noise.

When the noise is cyclostationary the sidebands will be correlated, which results in the AM component dominating over the PM, or visa versa.

18.0 Noise + Compression = Phase Noise

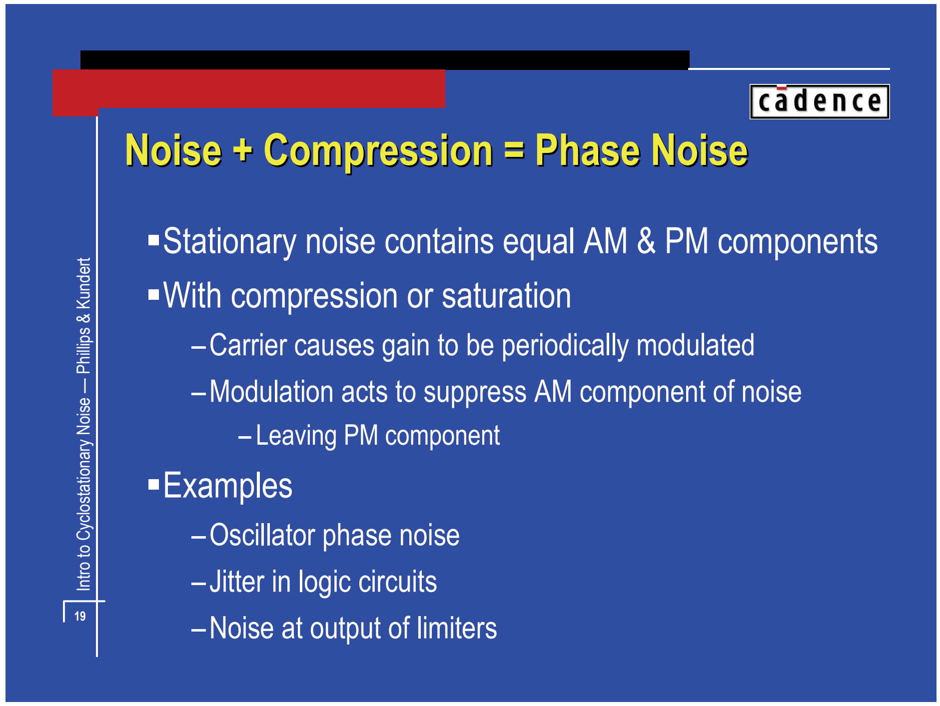

It is a general rule that when stationary noise is passed through a stage undergoing compression or saturation, the noise at the output is predominantly phase noise. Stationary noise contains equal amounts of amplitude and phase noise. Passing it through a stage undergoing compression cause the AM noise to be suppressed, leaving mainly the PM noise.

19.0 Ways of Characterizing Cyclostationary Noise

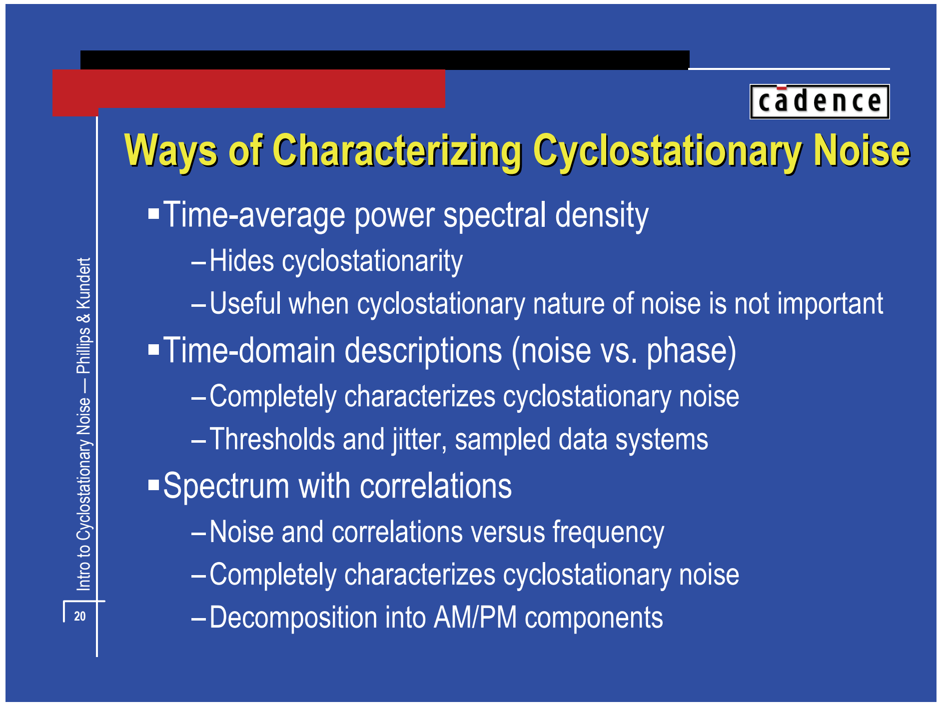

More review …

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 it 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 variations in the noise is much higher than the bandwidth of the analyzer). This is the primary output from SpectreRF's PNoise analysis.

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. This would be the case if two stages 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. SpectreRF outputs the correlations between sidebands if *noisetype=correlations*.

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 periodic signal is repeatedly sampled at the same point in phase during each period. It is useful in determining the noise that will result 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. SpectreRF outputs the noise versus phase if *noisetype=timedomain*.

20.0 When Can Cyclostationarity be Ignored?

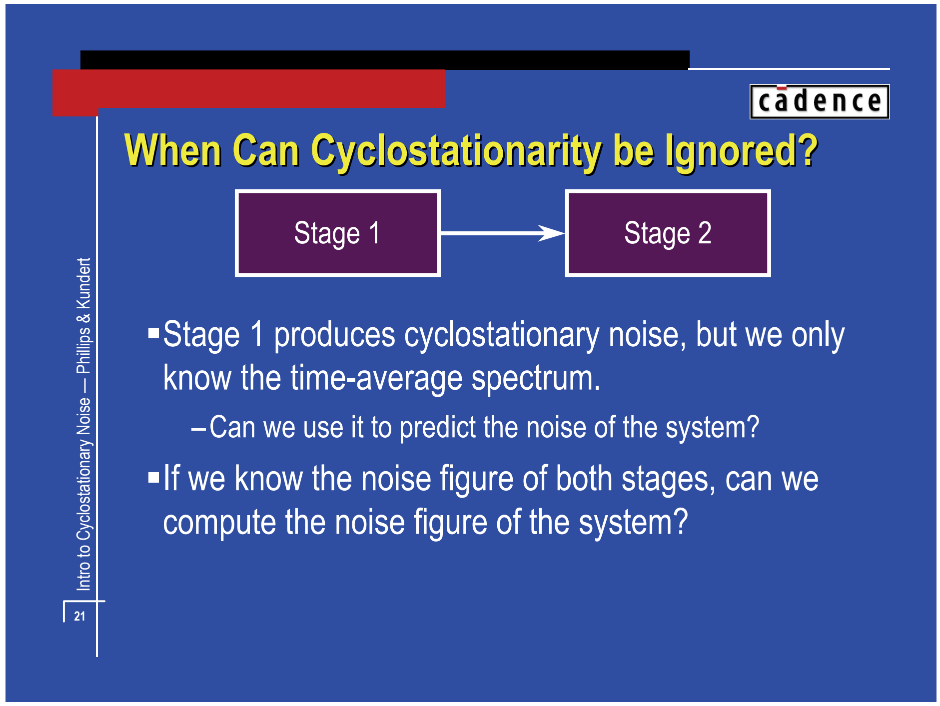

Sometimes it is only possible to determine the time-average PSD. So an important question is, when is it good enough? In other words, if you have a stage that is producing cyclostationary noise that drives another stage, and you know the time-average spectrum of the noise produced by the first stage, when is that enough to allow you to predict the effect of that noise on the subsequent stage.

It is enough to know the time-average spectrum of the noise produced by the first stage if the subsequent stage effectively averages the noise at its input. There are two cases where that occurs.

1. When the subsequent stage is a narrowband filter.

2. When the subsequent stage is non-synchronous with the one producing the cyclostationary noise.

Consider the case where one stage drives another. The first stage generates cyclostationary noise, but we only know the time-average PSD. In what situations do we know enough to characterize the noise of the whole system and when do we not. More specifically, assume that we know the noise figures of both stages, can we combine the noise figures using the standard formula to compute the noise figure of the entire system? These formulas assume that the noise produced by each stage is stationary, which is not true in this case.

21.0 Removing Cyclostationarity

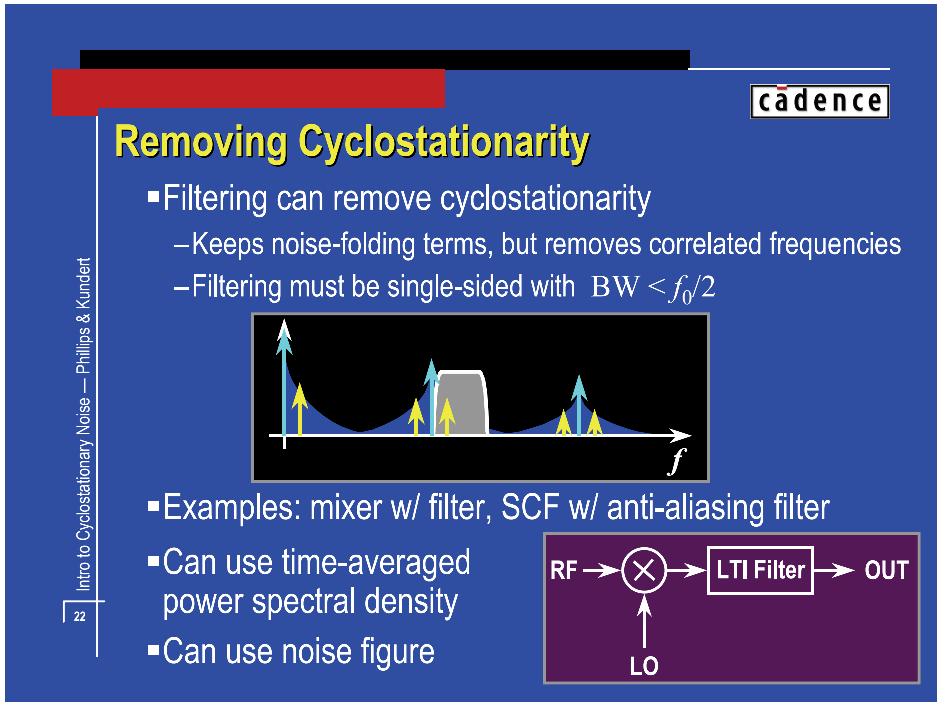

If a stage that generates cyclostationary noise is followed by a filter whose passband is constrained to a single sideband (the passband does not contain a harmonic and has a bandwidth of less than $f_0/2$, where f_0 is the fundamental frequency of the cyclostationarity), then the output of the filter 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.

In other words, the filter eliminates noise at any frequency that might be possibly be correlated with noise within the passband of the filter and so the resulting noise is not cyclostationary.

22.0 Ignoring Cyclostationarity

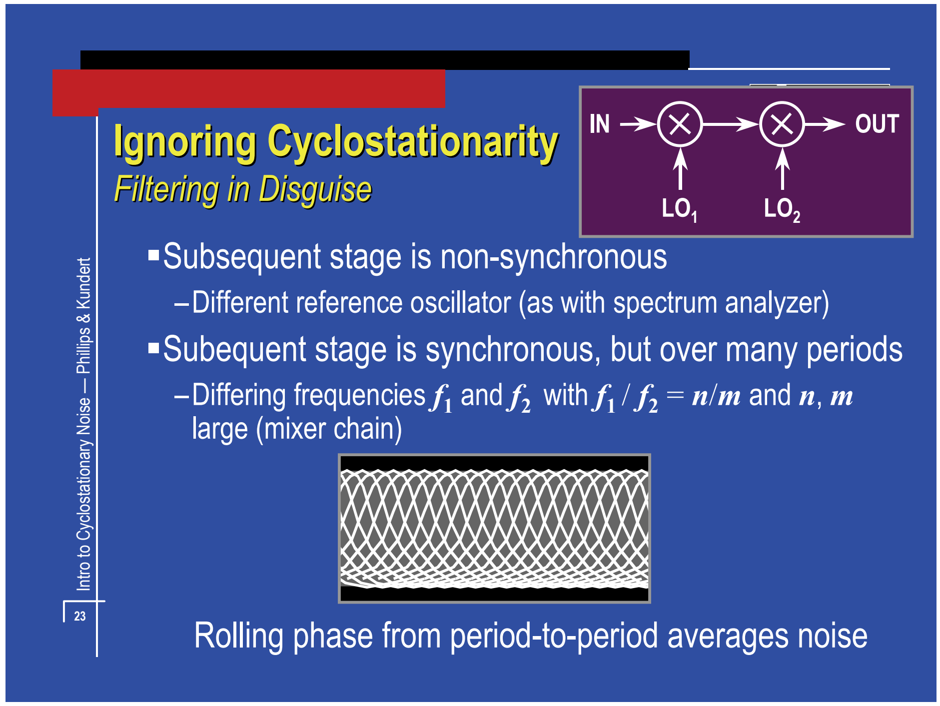

Consider a stage that generates cyclostationary noise with modulation frequency f_1 that is followed by a stage whose transfer characteristics vary periodically at a frequency of f_2 (such as a mixer, sampler, etc.). Assume that f_1 and f_2 are non commensurate (there is no f_0 such that $f_1 = n\, f_0$ 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 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 it 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).

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

23.0 When to Use the Time-Averaged PSD

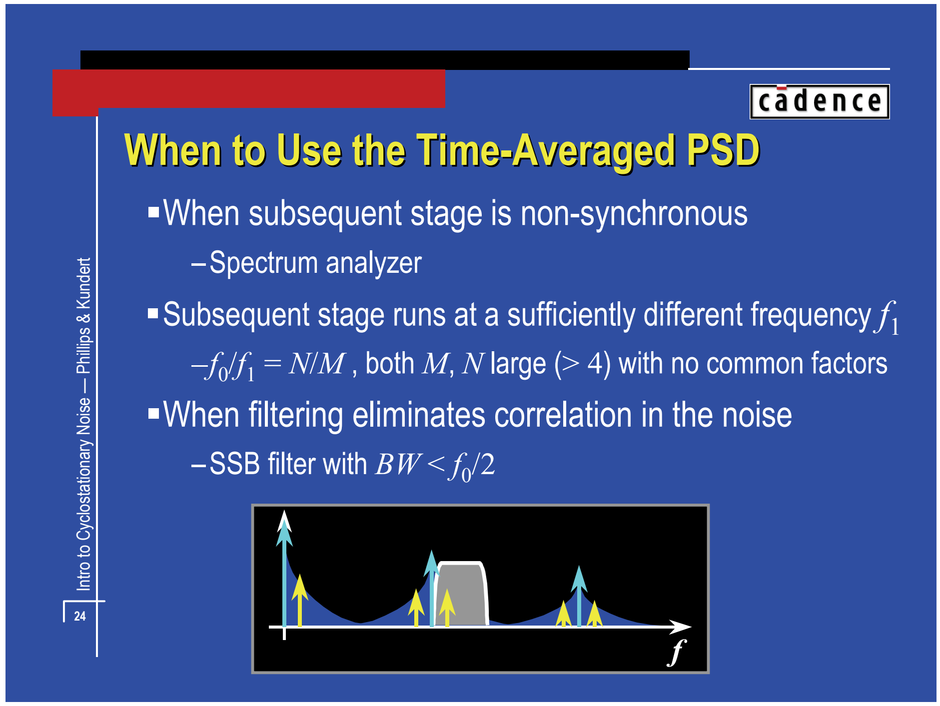

You are free to use the time-averaged power spectral density (PSD) when the cyclostationary nature of the noise will be eliminated or ignored by subsequent stages. Filtering eliminates the cyclostationary nature of noise, converting it to stationary noise, if the filter is a single-sideband filter with bandwidth less that $f_0/2$. 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.

24.0 When Knowing Time-Average PSD of a Stage is not Enough

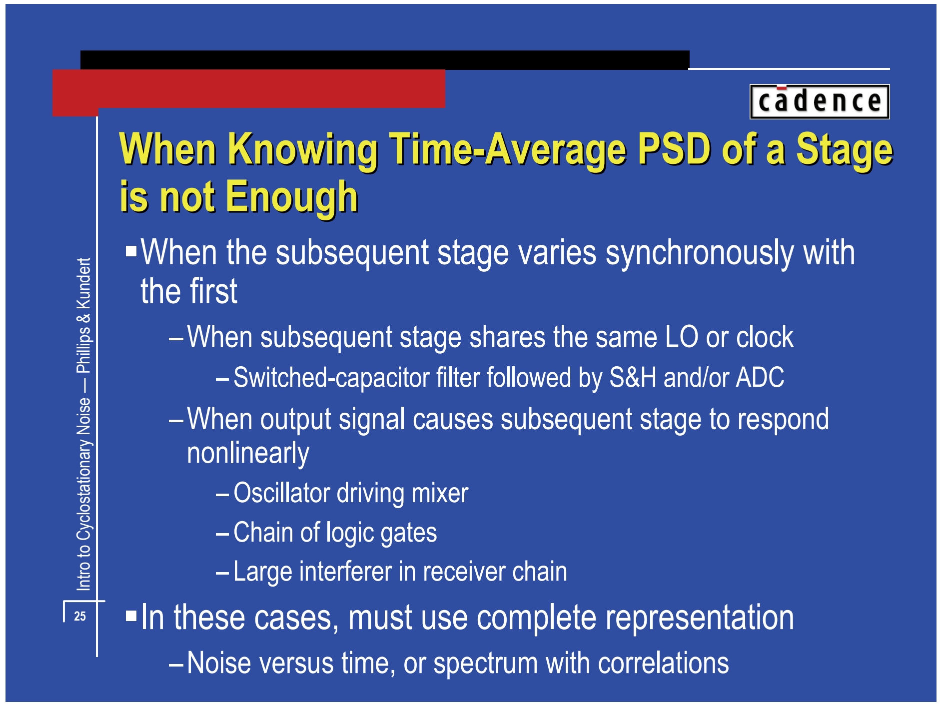

When a stage producing cyclostationary noise drives a subsequent stage that has a time-varying transfer function that is synchronous with the first, then ignoring the cyclostationary nature of the noise from the first (using the time-average PSD) generates incorrect results. One common situation where 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 nonlinearly. In this case, large periodic output signal that modulates the gain of the subsequent stage in synchronism with the cyclostationary noise produced by the first stage. This occurs with an oscillator drives the LO port of a mixer, when one logic gate drives another, or when a large interfering signal drives both stages into compression.

In these situations, one must consider the cyclostationary nature of the noise produced in the first stage when determining the overall noise performance of the stages together. Usually this must be done by either simulating or measuring the two stages together rather than apart.

25.0 Oscillator Phase Noise

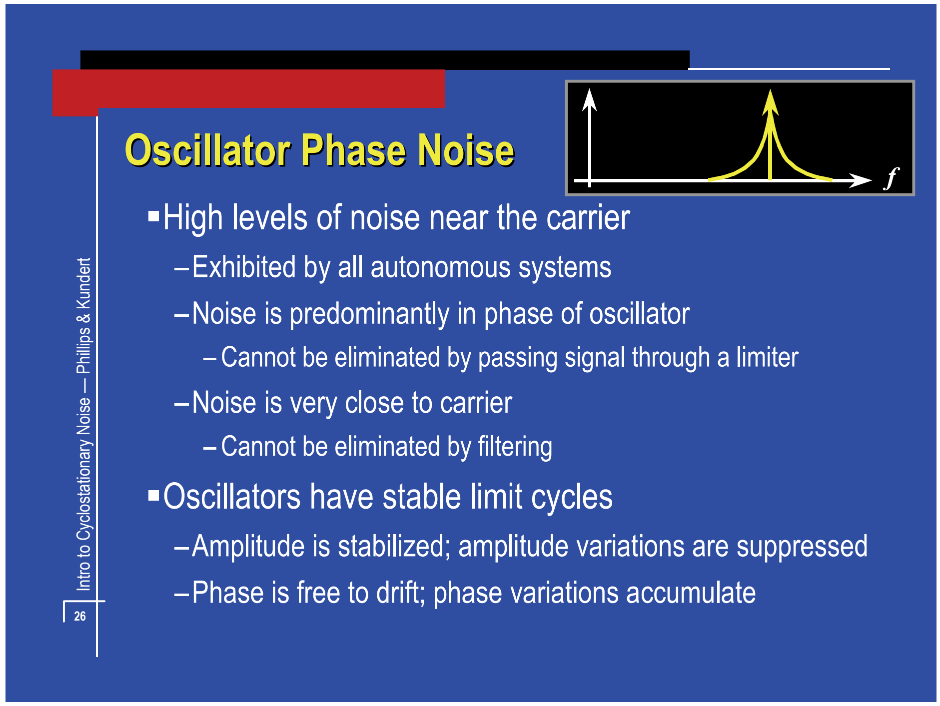

Another important class of cyclostationary noise is oscillator phase noise.

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 the nature of nonlinear oscillators that this noise be predominantly in the phase of the oscillation. Thus, the noise cannot be removed by passing the signal through a limiter.

26.0 Importance of Oscillator Phase Noise

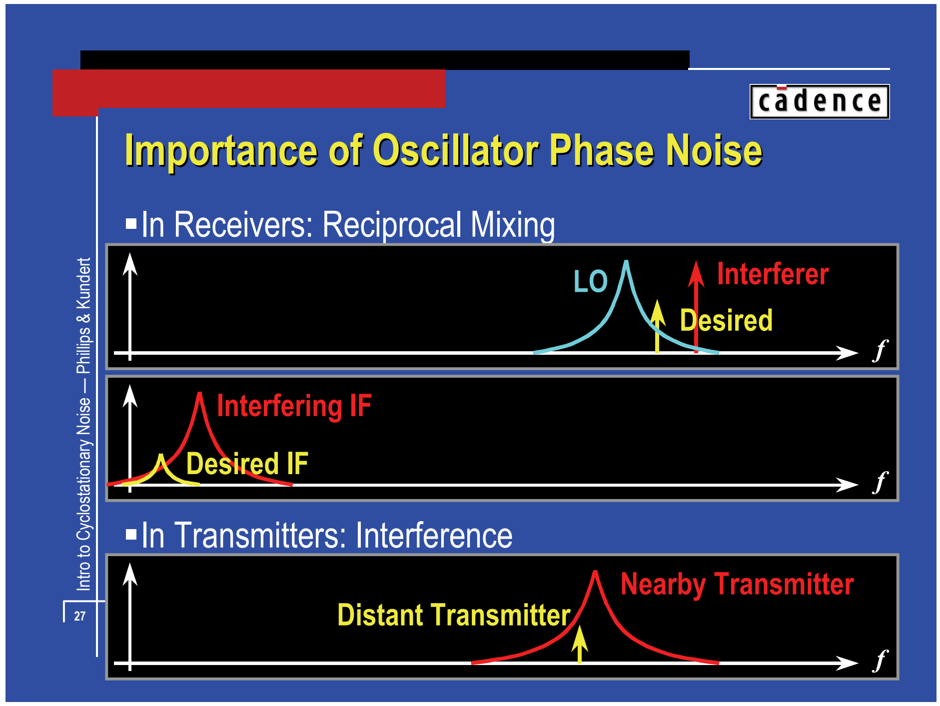

In a receiver, the phase noise of the LO can mix with a large interfering signal from a neighboring channel and swamp out the signal from the desired channel even though most of the power in the interfering IF is removed by the IF filter. This process is referred to as *reciprocal mixing*.

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.

27.0 The Oscillator Limit Cycle

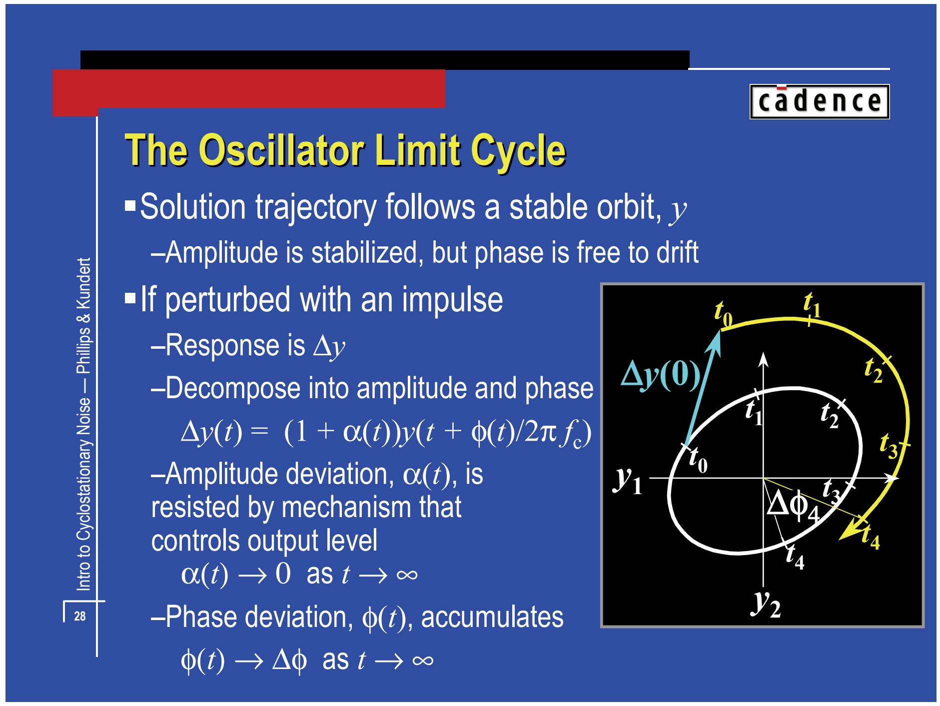

Consider plotting the capacitor voltage against the inductor current in a resonant oscillator. In steady state, the trajectory is a stable limit cycle. Now consider perturbing the oscillator with an impulse $x = \delta(t)$ and assume that the response to the perturbation is Δy. Separate Δy into amplitude and phase variations,

$$\Delta y(t) = (1 + \alpha(t))y(t + \phi(t)/2\pi f_c) - y(t).$$

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

Since the oscillator 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 ($\alpha(t) \to 0$ as $t \to \infty$). 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.

Since the oscillator is autonomous, 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 ($\phi(t) \to \Delta\phi$ as $t \to \infty$).

28.0 The Oscillator Limit Cycle (cont.)

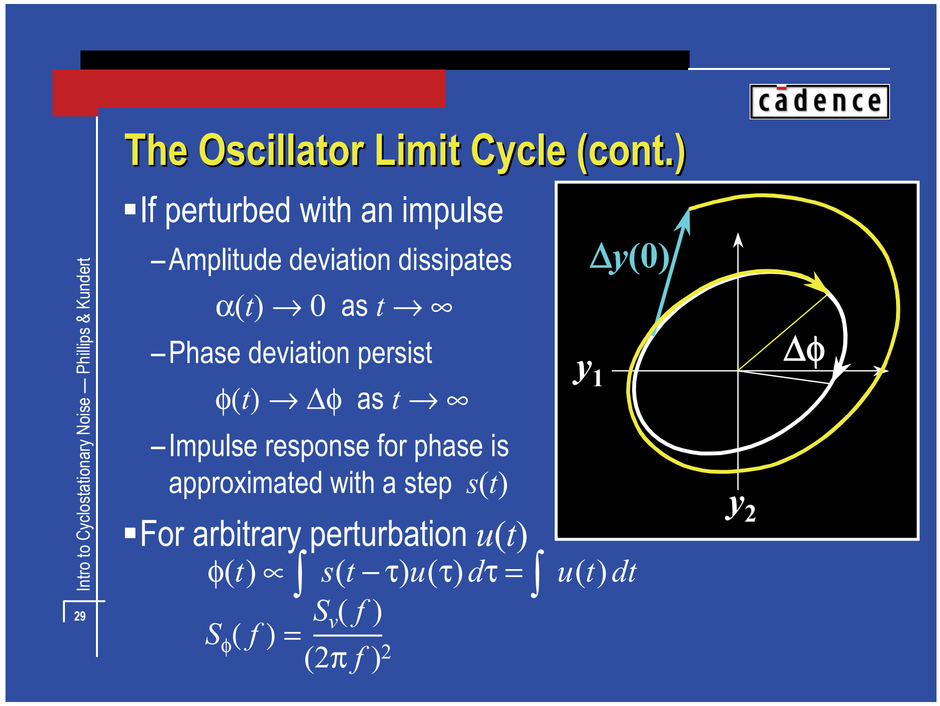

On the previous slide it was pointed out that $\alpha(t) \to 0$ as $t \to \infty$ and $\phi(t) \to \Delta\,\phi$ as $t \to \infty$. If we neglect any short term time constants, it can be inferred that the impulse response of the phase deviation $\phi(t)$ can be approximated with a unit step $s(t)$. The phase shift over time for an arbitrary input disturbance u is

$$\phi(t) \sim \int s(t-\tau)u(\tau)d\tau = \int u(t)\, dt,$$

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

$$S_\phi(f) \sim S_u(f)/(2\pi f)^2$$

This represents another way of explaining why oscillator noise is primarily phase noise and why the noise grows at frequencies close to the carrier frequency.

29.0 $1/\Delta f^2$ Amplification of Noise in Oscillator

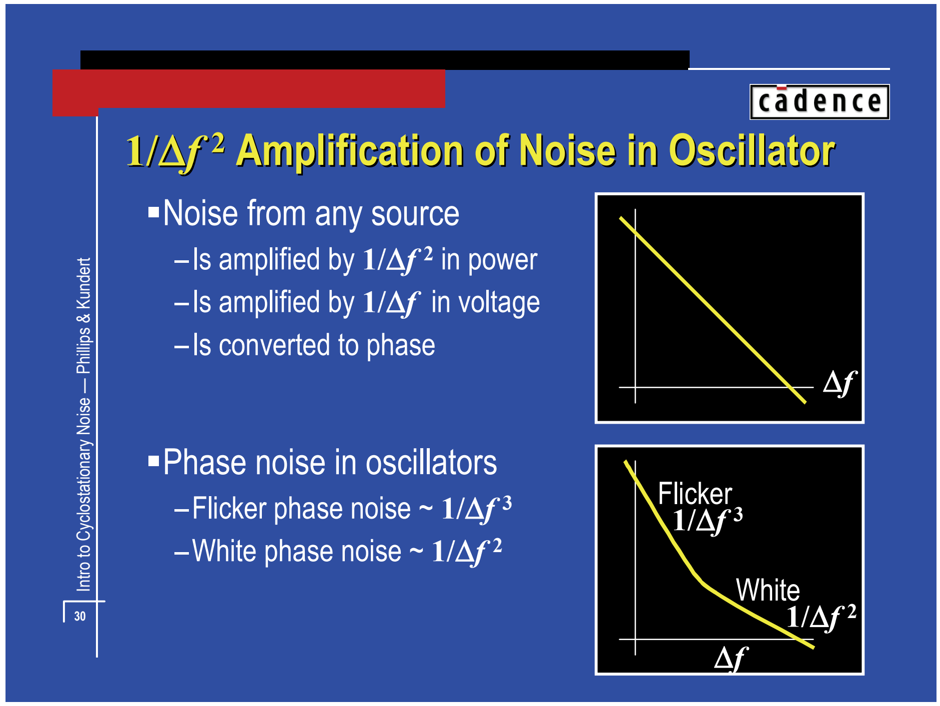

As shown previously, the gain from a noise source to the output phase is $1/\Delta f^2$ in power or $1/\Delta f$ in voltage. As such, flicker noise, whose power spectral density is proportional to $1/\Delta f$ in power, ends up producing a $1/\Delta f^3$ region in power in the power spectral density of the phase at the output of the oscillator. Similarly, white noise ends up producing a $1/\Delta f^2$ region in power in the power spectral density of the phase at the output of the oscillator.

30.0 Difference Between S_ϕ and S_v Noise

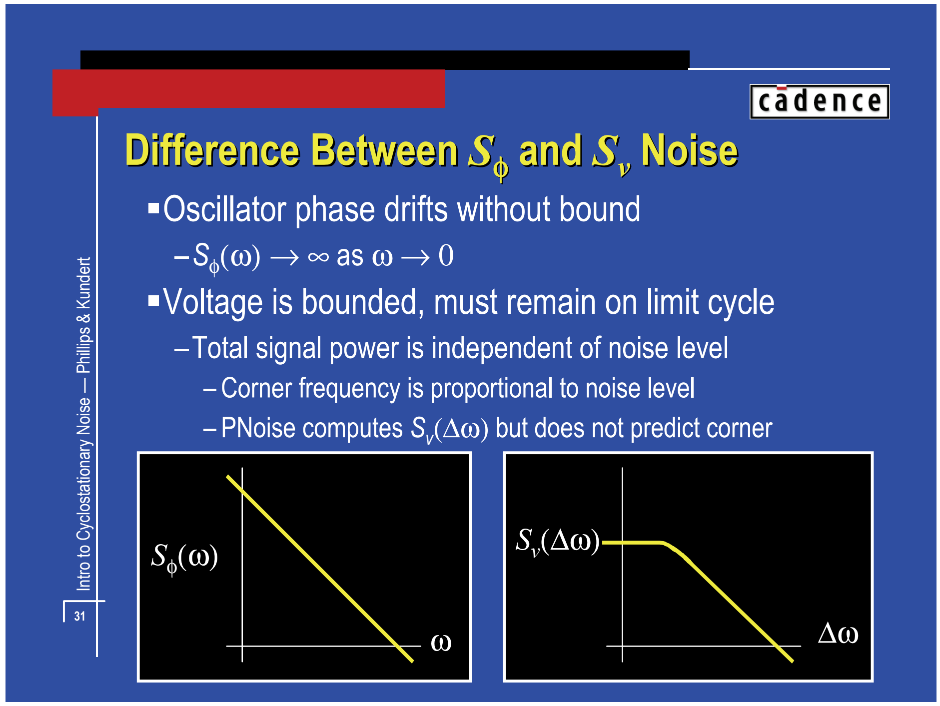

There are two different ways of characterizing noise in the same oscillator. S_ϕ is the spectral density of the phase and S_v is the spectral density of the voltage. S_v contains both amplitude and phase noise components, but with oscillators the phase noise dominates except at frequencies far from the 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. Another measure of oscillator noise is L, which is simply S_v normalized to the power in the fundamental.

As $t \to \infty$ the phase of the oscillator drifts without bound, and so $S_\phi(\omega) \to \infty$ as $\omega \to 0$. 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 $\Delta\omega \to 0$ the PSD of v flattens out. 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 phase noise, $S_v(\omega)$ is a series of impulse functions at the harmonics of the oscillation frequency. With phase noise, the impulse functions spread, becoming fatter and shorter but retaining the same total power.

The voltage noise S_v is considered small outside the linewidth and thus can be accurately predicted using small signal analyses. Conversely, the voltage noise within the linewidth is large 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).

31.0 Oscillator Phase Noise

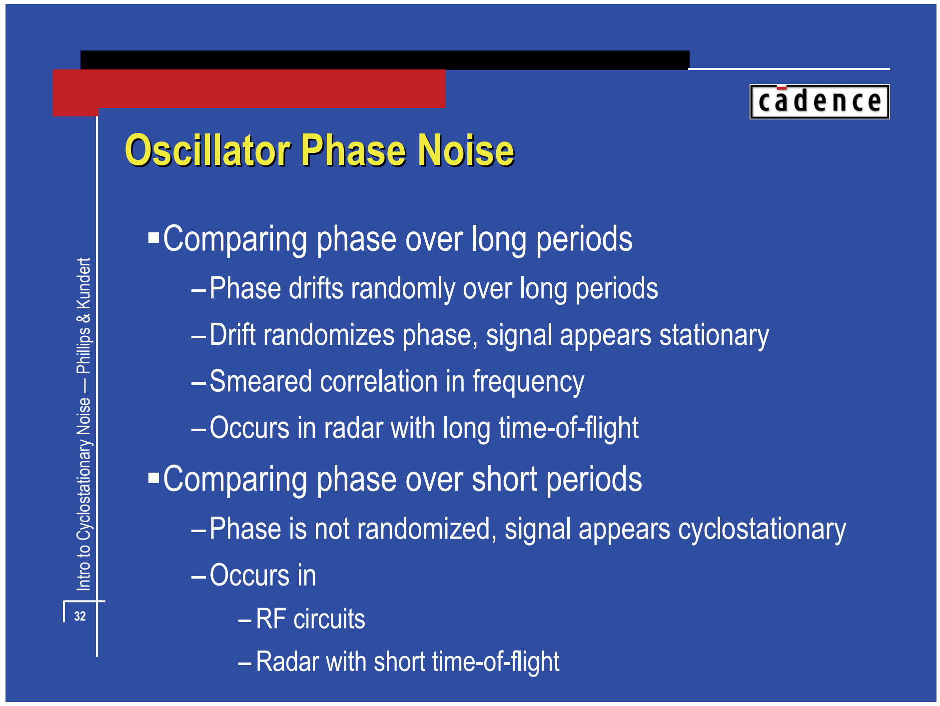

The phase of an oscillator is a random walk process. Thus, if you know the phase of an oscillator at some point, that will give you information about the phase of the oscillator in the near future but not the distant future. After the phase has been measured, each cycle brings a bit of uncertainty in the phase. That uncertainty accumulated over 1000's of cycles until eventually it builds up to be more than one complete cycle. At this point the phase is completely randomized and you no longer have any information about the phase.

Consider an oscillator used in a radar that is measuring the distance to the moon. Assume that the oscillator generates the signal that is transmitted to the moon, and that during the trip the phase of the oscillator becomes completely randomized. Once the signal bounces off the moon and returns, it is mixed with the oscillator signal again to convert it to baseband. Since the phase has been completely randomized, we can model this system with two different free running oscillators operating at the same frequency, one that generates the transmission signal, and the other that is used for the LO. Thus, the mixer is operating completely asynchronously with the transmit oscillator and so the noise produced by the transmit oscillator looks stationary to the mixer. Thus, over the long term oscillators produce stationary noise.

However, over the short term they produce cyclostationary noise. In a typical RF circuit, such as a LO driving a mixer, the mixer is operating in complete synchronism with the oscillator and so the noise produce by the oscillator appears cyclostationary to the mixer.

32.0 Jitter

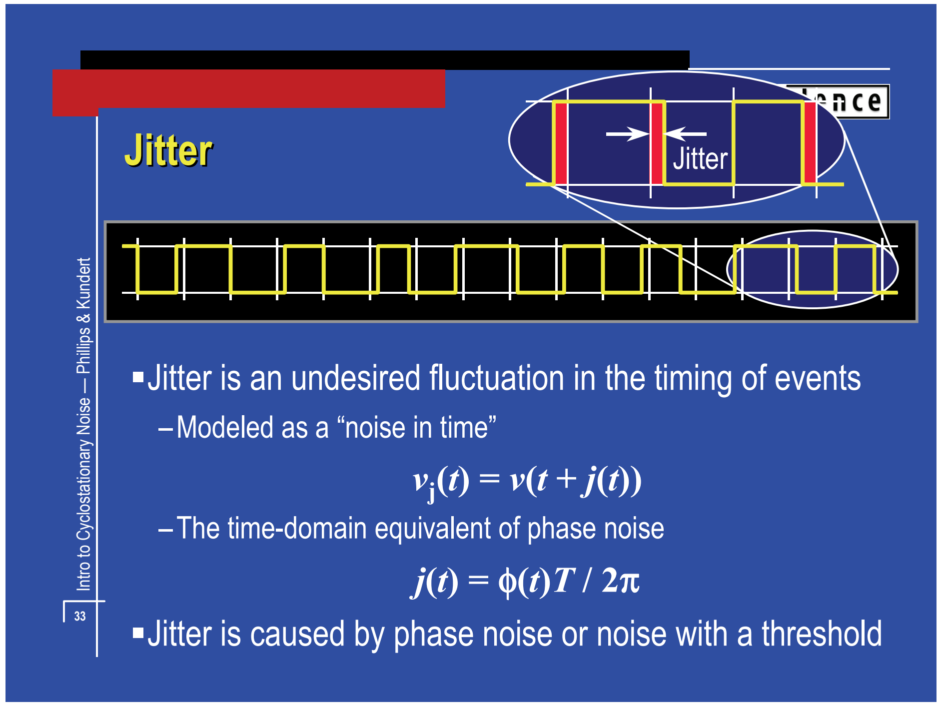

Jitter is an undesired fluctuation in the timing of events. It can be modeled as a noise in time. It is equivalent to phase noise and is used in situations where it is more natural think of the noise being in time rather than in phase or in signal level. Like phase noise, it serves to degrade system performance.

33.0 Noise + Threshold = Jitter

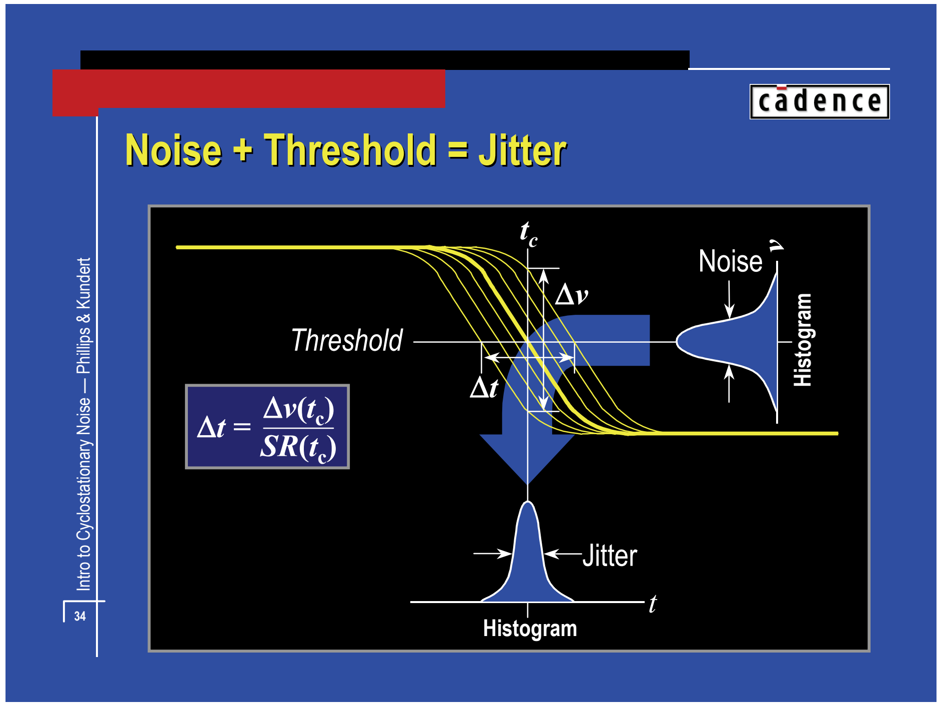

Noise added to a signal transitioning through a threshold adds jitter to the time at which the signal crosses the threshold. This is the way jitter is created in nonlinear circuits such as digital circuitry. The statistics of the jitter can be computed from the statistics of the noise at the time of the threshold crossing and the slewrate (or time derivative) of the large signal at the threshold crossing.

34.0 Jitter + Sampler = Noise

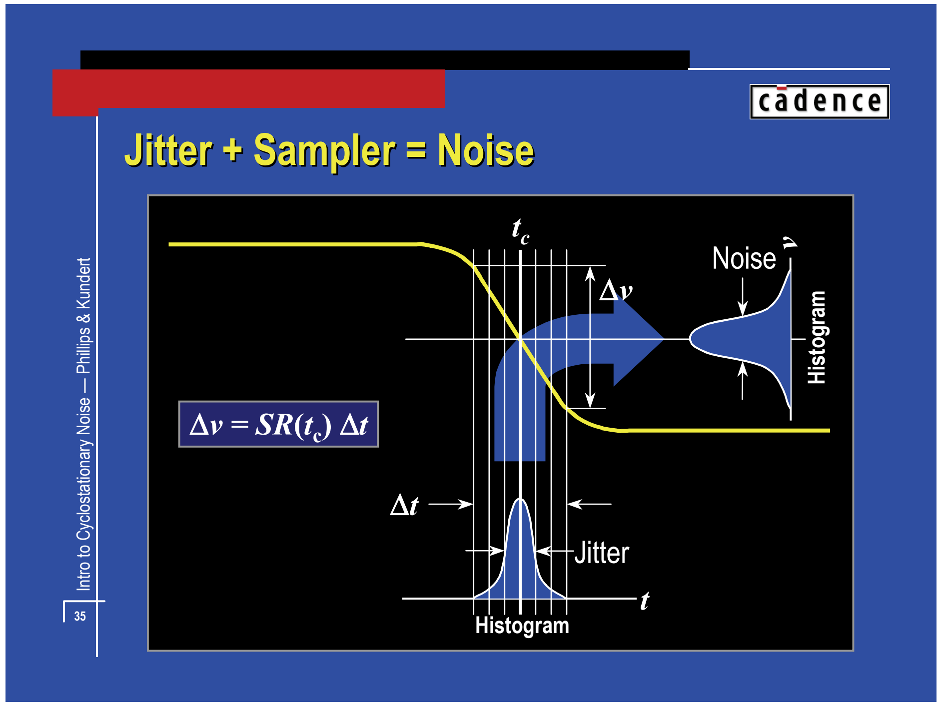

Jitter in the time at which a signal is sampled creates noise in the sampled result if the signal is changing at the time it is sampled. This is one way in which noise is generated when converting continuous-time signals are converted to discrete-time signals. The statistics of the noise can be computed from the statistics 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.

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

35.0 Cyclostationary Noise from Logic Gates

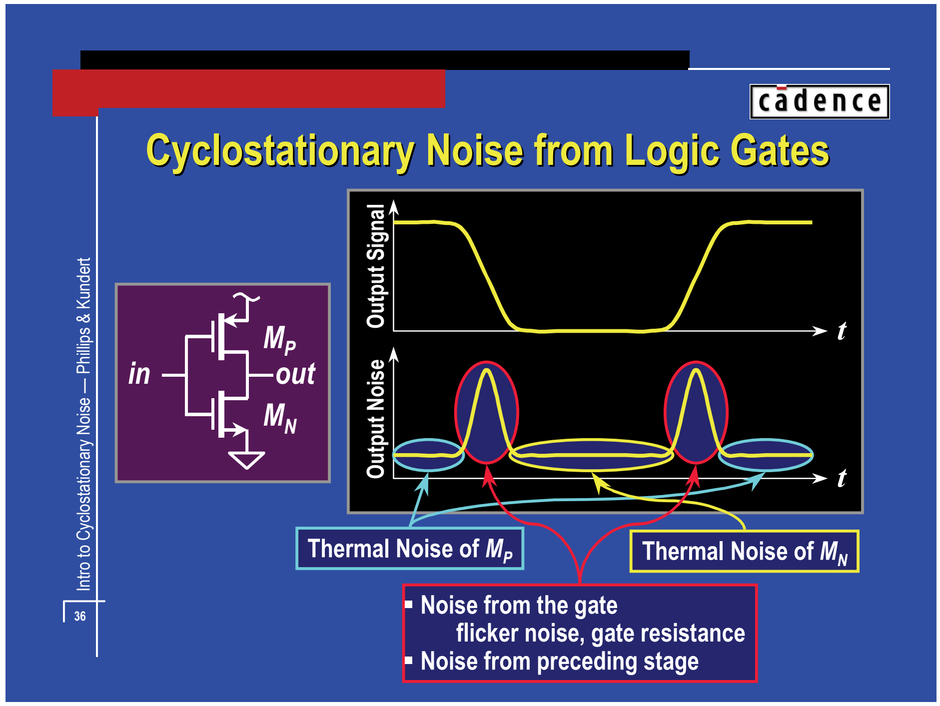

The noise produced by a logic gate comes from different places depending on the state of the output. When the output is high and M_P is on, the output is insensitive to small changes on the input and the noise at the output is predominantly due to the thermal noise from the channel of M_P. When the output 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 contributes to the output. In addition, the output is sensitive to small changes in the input, in fact, noise from 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.

36.0 Noise in a Chain of Logic Gates

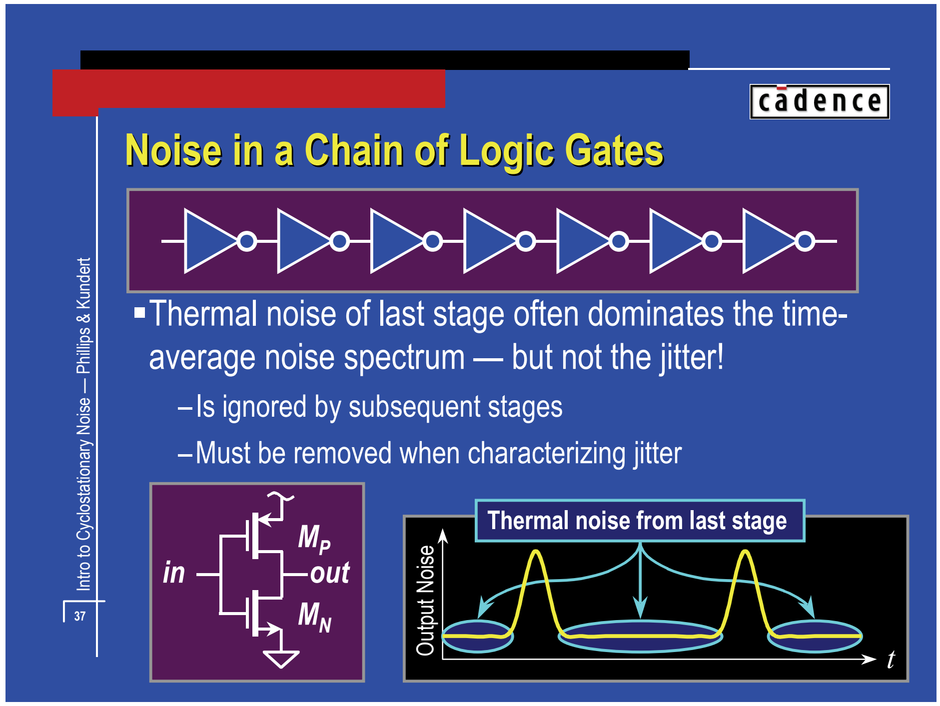

The total output noise power of a gate is dominated by the thermal noise produced by the output devices if the gate spends most of its time with an unchanging output. This noise is usually ignored by subsequent stages. Thus, using the time-averaged spectral density to characterize the noise in a gate is misleading. Only the noise produced by a gate at the point where its output crosses the threshold of the subsequent stage should be taken into account when characterizing the jitter of the gate.

37.0 Characterizing Jitter in Logic Gate

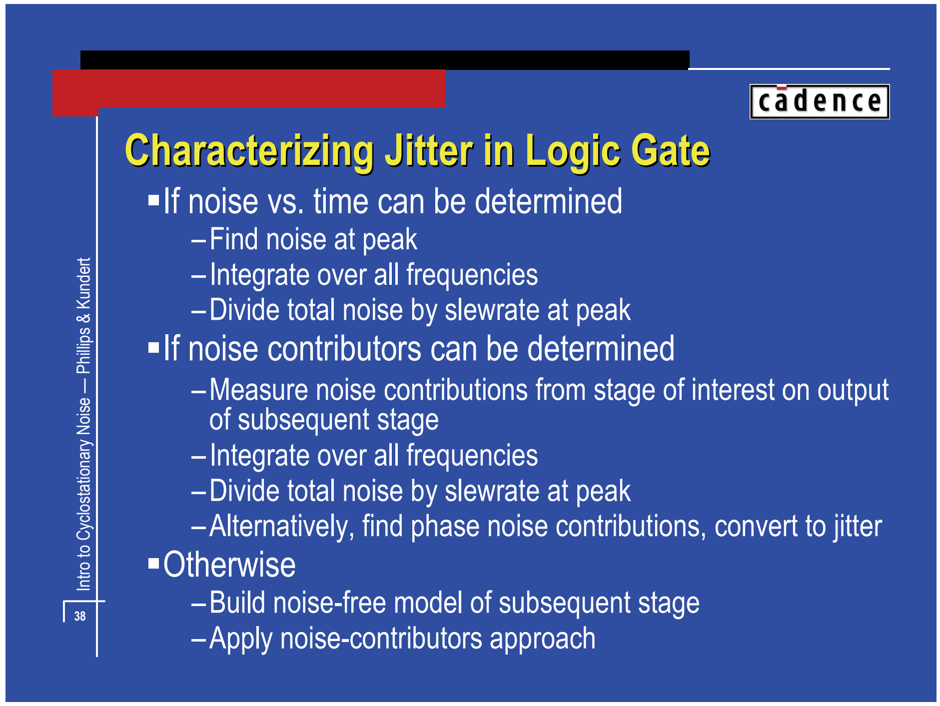

There are several different methods that can be used to determine the jitter of a logic block using a simulator such as SpectreRF. They all assume that the circuit is driven, so as to produce a periodic response, and that a cyclostationary noise analysis is performed about that periodic operating point.

The first approach can be used if the simulator is capable of determining the instantaneous noise power at a particular point in time (in this case the time of the threshold crossing) (SpectreRF can do this). Simply determine the time where the output signal crosses the threshold. Then find both the instantaneous noise power and the slewrate of the output signal at that time. Then convert the noise power to a noise voltage and divide through by the slewrate to determine the jitter.

The second approach can be used if the contributions of individual noise sources can be identified. In this case, the circuit is simulated along with the subsequent stage, or the stage it is expected to drive. The subsequent stage must be a thresholding circuit, and so the noise at the output will only contain contributions from the stage of interest if they contribute as jitter. Any noise from the stage of interest that does not manifest as jitter at the output of the subsequent stage will be blocked by the thresholding nature of the subsequent stage. To determine the jitter at the output due to the stage of interest, divide the average total noise at the output contributed by the stage of interest and divide by the RMS value of the derivative of the output signal.

Finally, the third approach is similar to the second, except that rather than following the stage of interest with the actual subsequent stage, you replace it with a model that produces no noise of its own. One can create such a model in the form of a limiter using Verilog-A (for proper operation it must have finite gain in its active region). In this way, only the stage of interest is producing noise and so one can use the total output noise rather than the individual noise contributions. Again, one divides the average total output noise by the RMS value of the time-derivative of the output signal to compute the jitter.

38.0 Characterizing Jitter

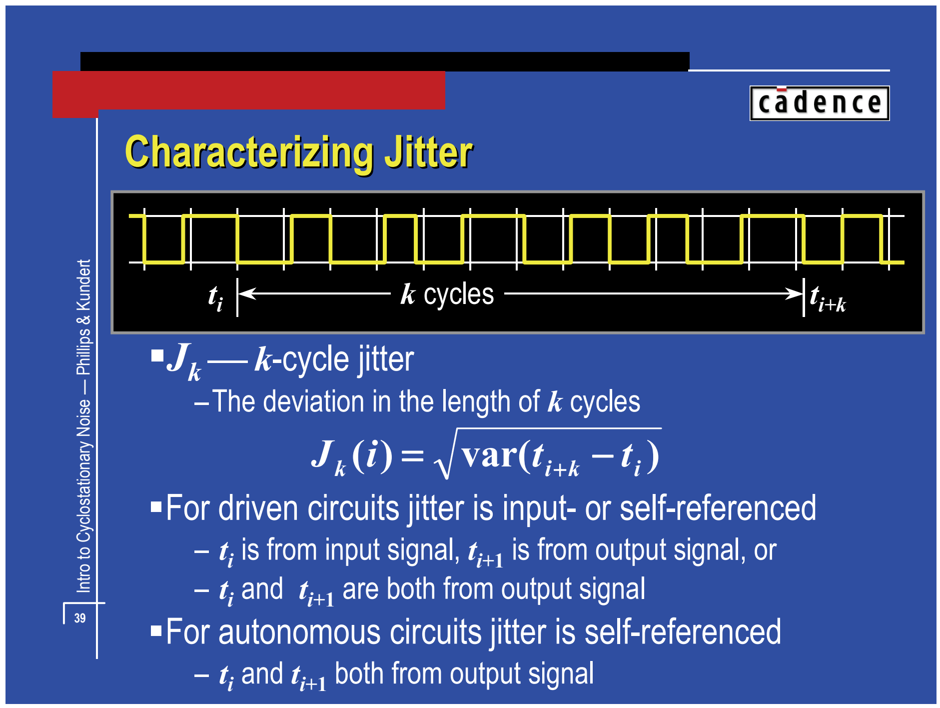

A common way to characterize jitter is k-cycle jitter, J_k. It represents the variation in the length of k cycles. If there is no jitter, then J_k will be 0, if the jitter on each cycle is completely uncorrelated with the jitter on previous cycles, then J_k will be constant with k, and if the jitter is correlated from cycle-to-cycle, then J_k will vary as a function of k.

For driven circuits, there are two different ways to measure k-cycle jitter. J_k is input-referred if a transition on the input signal starts the measurement interval. J_k is self-referred if a transition on the output signals starts the measurement interval. The results achieved using these two different approaches varies slightly because the output-referred jitter includes an extra noisy transition when compared to the input-referred jitter. Input-referred jitter is generally used for driven circuits.

With autonomous circuits there is no input transition that can be used to start the measurement interval, so self-referred jitter is used on oscillators.

39.0 Jitter in Simple Driven Circuits (Logic)

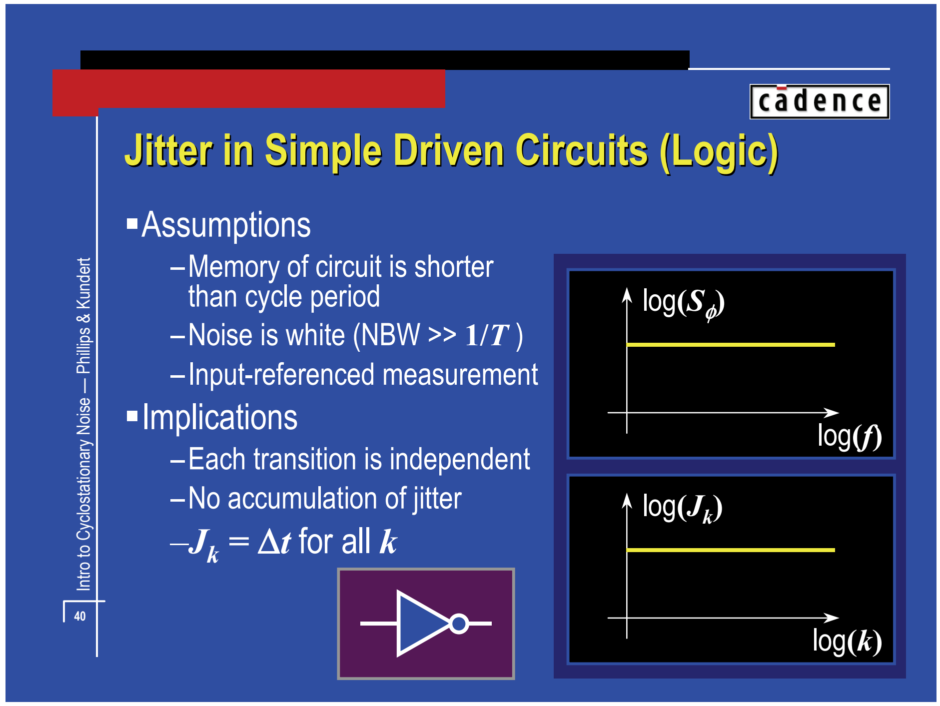

The next few slides attempt to give an feeling for what to expect with J_k on several important classes of circuits, First, consider a simple driven circuit such as an inverter. Assume that the bandwidth of the circuit is well beyond the input frequency, that the noise sources in the circuit are all white, and that J_k is measured using a input-referred measurement. In this case, the jitter on each transition is independent and uncorrelated to that on any other transition and $J_k = \Delta t$ for all k.

40.0 Jitter in Autonomous Circuits (Ring Osc, ...)

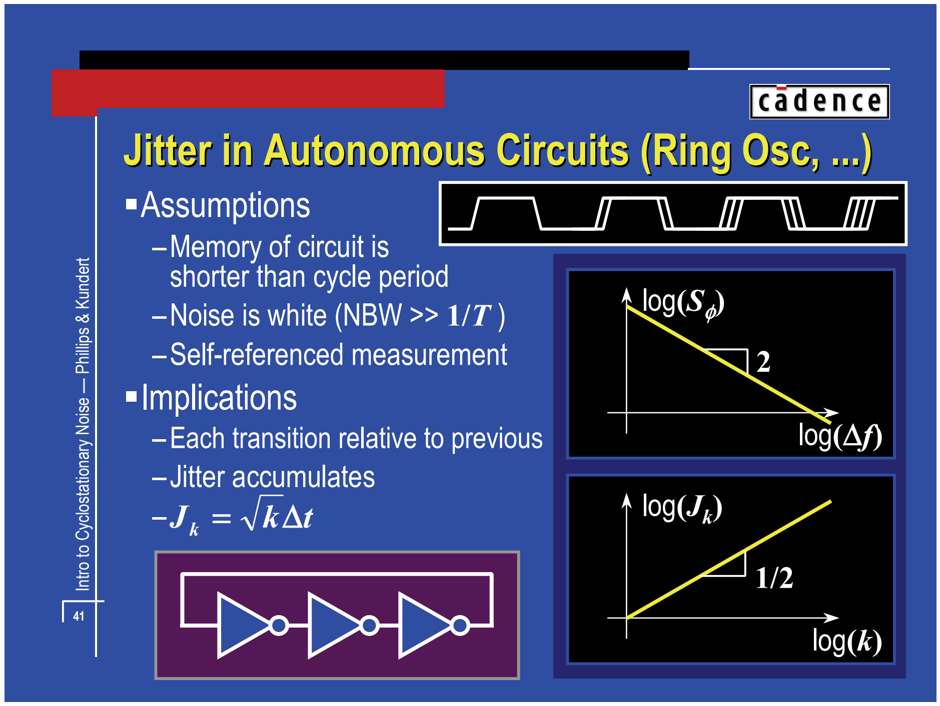

Next, consider a simple autonomous circuit such as an ring oscillator. Assume that the noise sources in the circuit are all white and that J_k is measured using a self-referred measurement. In this case, the jitter variance accumulates from cycle-to-cycle, and so $J_k = k^{1/2}\Delta t$.

41.0 Jitter in PLLs

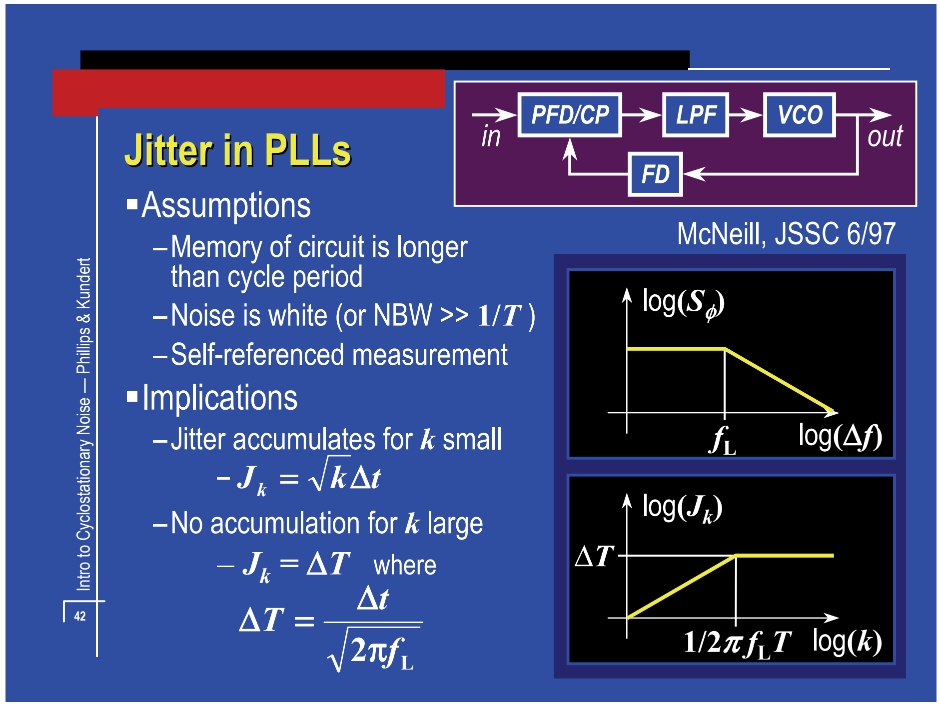

Finally, consider a circuit that combines the characteristic of the previous two circuits such as a phase locked loop. At low frequencies, the gain of the loop acts to suppress the jitter produced by the voltage-controlled oscillator (VCO). In this case, if we assume the reference is free of jitter, then the jitter is due to the phase-frequency detector & charge pump (PFD/CP) and the frequency divider (FD), which are wideband driven circuits. Thus, at low frequencies the noise is approximately flat with frequency. For the same reasons, J_k is independent of k for large k. At high frequencies, the low-pass filter (LPF) essentially opens the loop. Thus the noise and jitter characteristics of the output are basically those of the VCO alone. Thus, for large Δf the noise is proportional to $J_k f^{-2}$ and for small k J_k is proportional $k^{1/2}$.

42.0 Summary

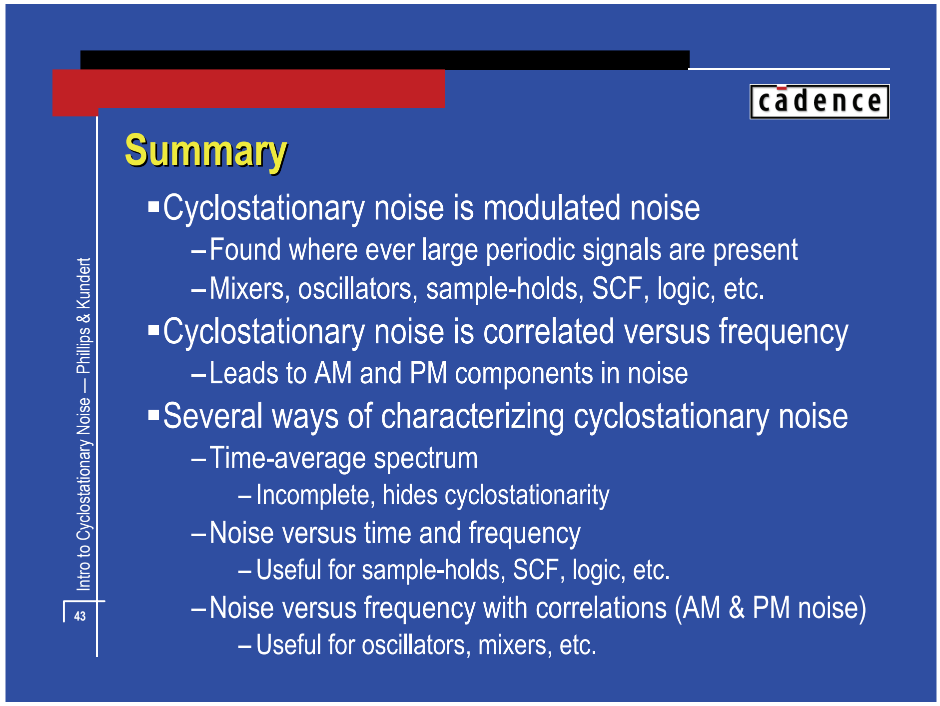

42.1 If You Have Questions

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*.