# Cyclostationary Noise in RF Circuits

Joel Phillips & Ken Kundert
Cadence Design Systems

## 1. Introduction

The proliferation of wireless and mobile products have 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 strongly 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, such as SpectreRF.

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) and autonomous systems (oscillators) is discussed.

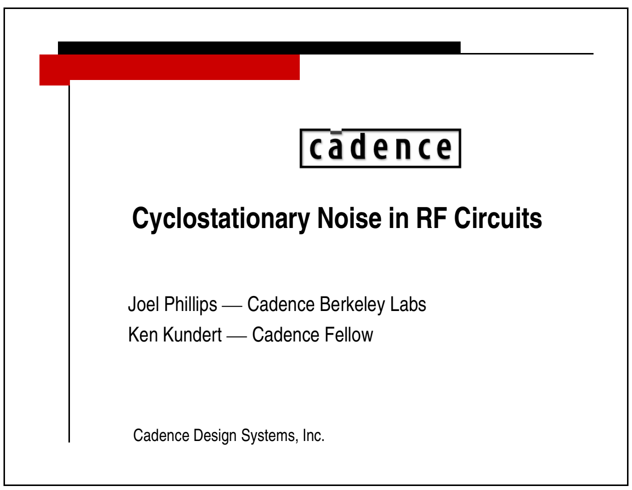


Slide 1

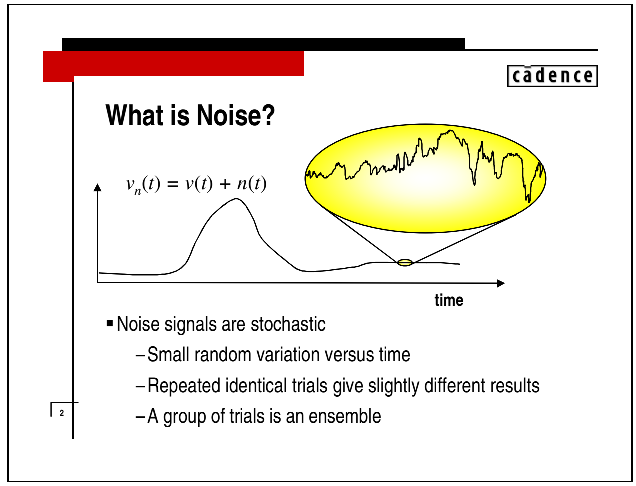


Slide 2

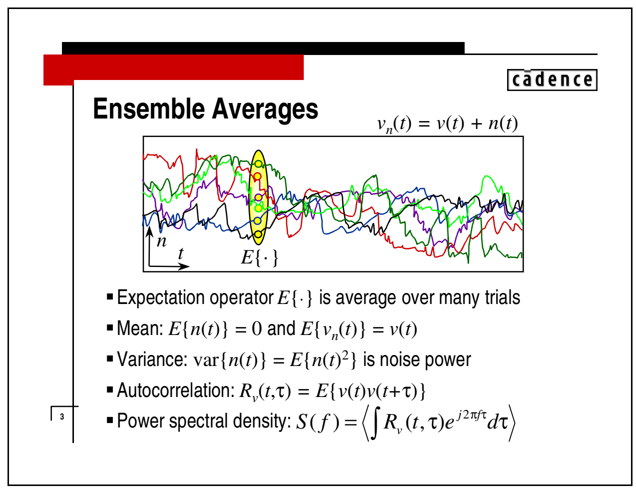


Slide 3

## 2. 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. 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 (as opposed to time averages). An averages over many trials is referred to as an expectation, and 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 $\tau$ 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. White Noise

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.

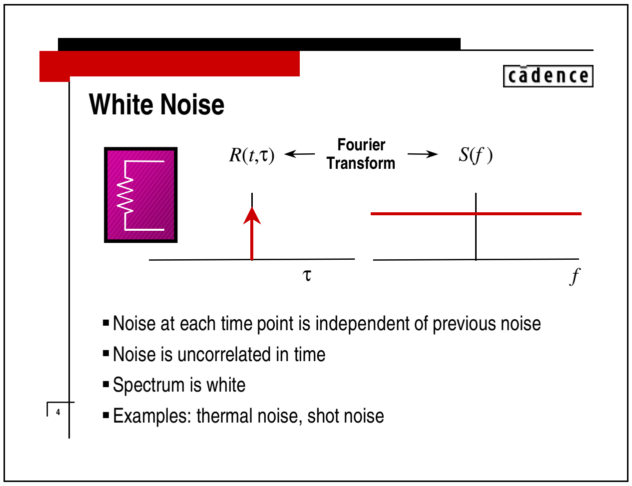


Slide 4

### 5. 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. 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.

### 6. Cyclostationary Noise

The ensemble averages for noise can vary with time. If they vary in a periodic fashion, they are referred to as being cyclostationary.

### 7. Origins of Cyclostationarity

Cyclostationary noise is generated by circuits with periodic or quasiperiodic 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.

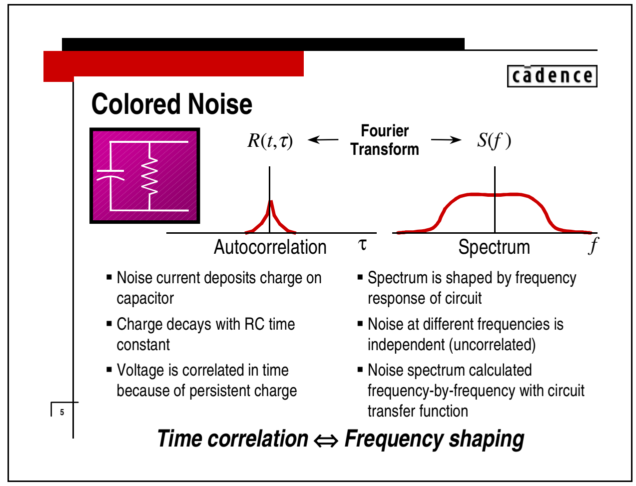


Slide 5

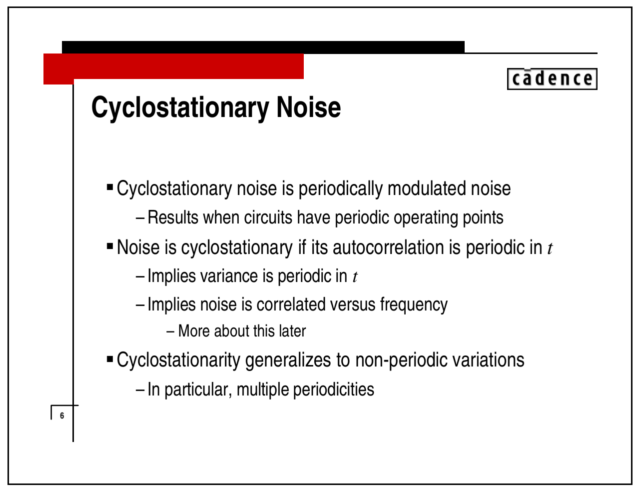


Slide 6

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

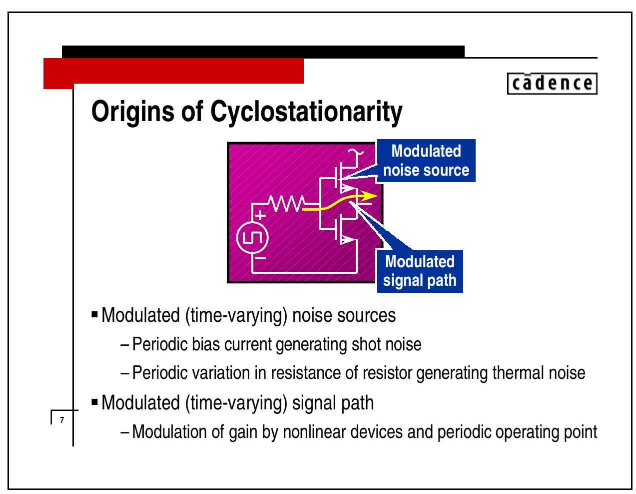


Slide 7

## 8. 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".

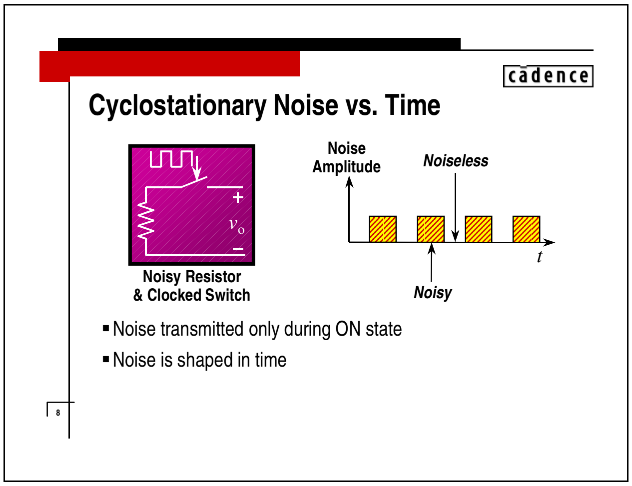


Slide 8

## 9. 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 the time-average PSD.

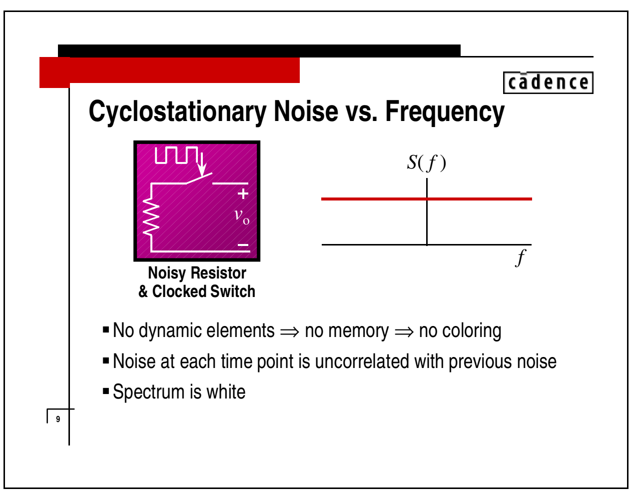


Slide 9

## 10. 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 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.

Noise from the source at a particular frequency $f$ is replicated and copies appear at $f \pm kf_0$. 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.

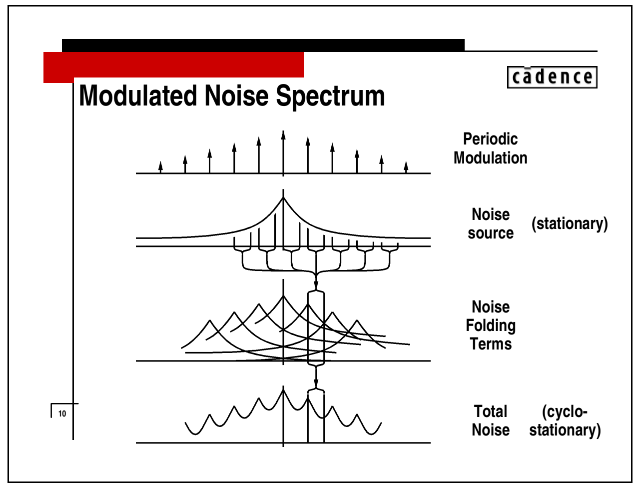


Slide 10

## 11. Frequency Correlations in Spectrum

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, 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.

## 12. Duality of Shape and Correlation

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.

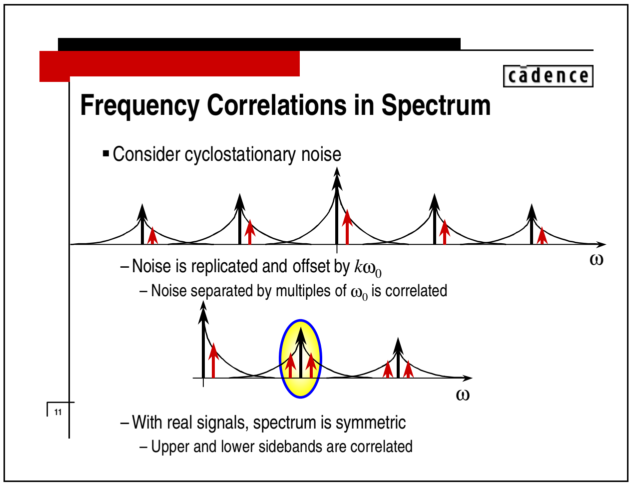


Slide 11

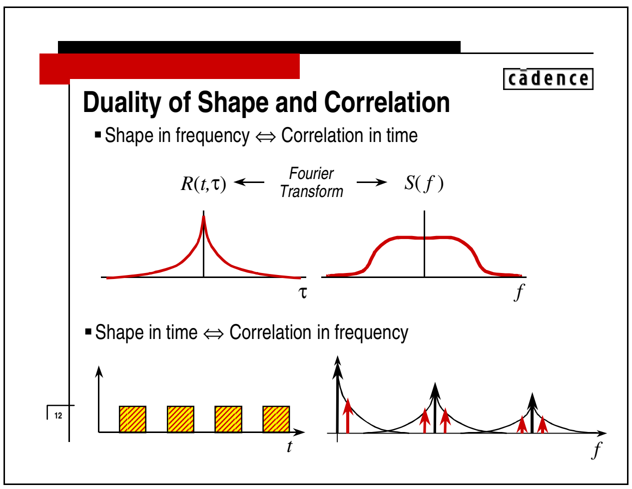


Slide 12

### 13. Ways of 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 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 noisy but otherwise 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*.

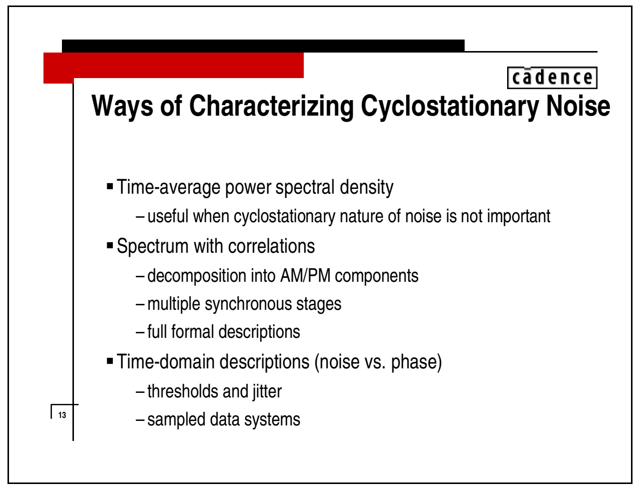


Slide 13

### 14. 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 any two frequencies within the passband is uncorrelated.

### 15. 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

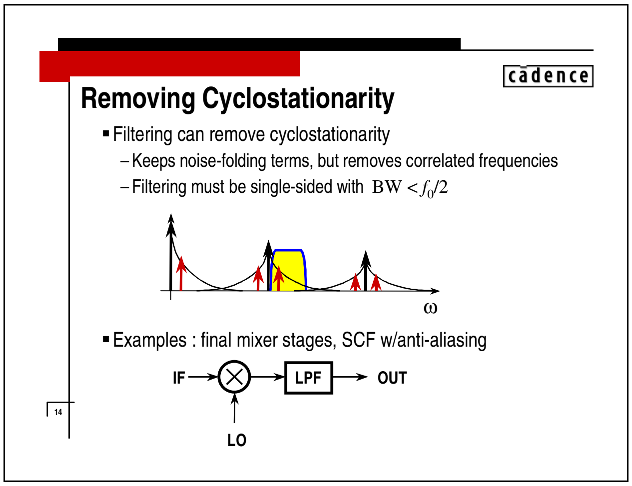


Slide 14

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.

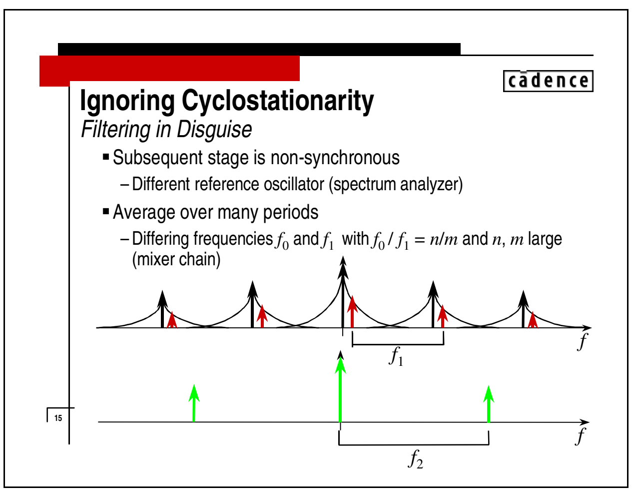


Slide 15

## 16. 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.

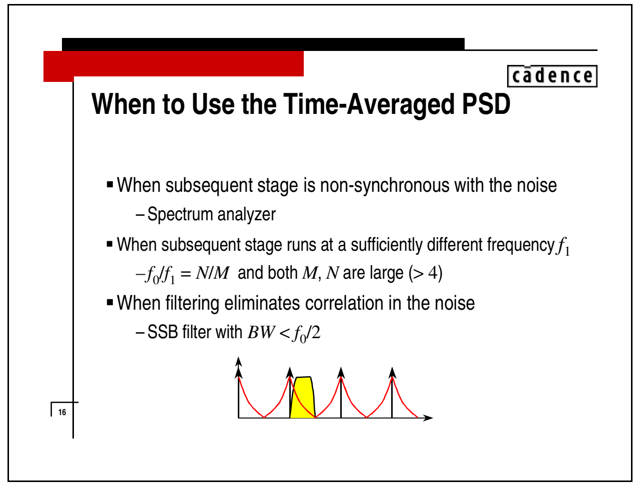


Slide 16

## 17. When Not to Use the Time-Averaged PSD

When a stage producing cyclostationary noise drives a

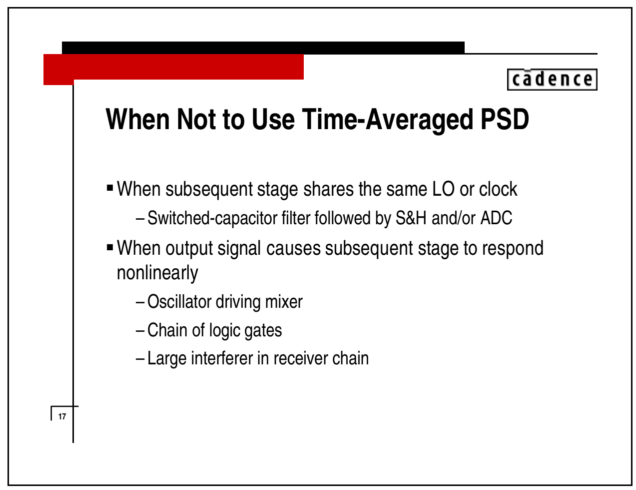


Slide 17

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, the large periodic output signal modulates the gain of the subsequent stage in synchronism with the cyclostationary noise produced by the first stage. This occurs when 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, consider the cyclostationary nature of the noise produced in the first stage when determining the overall noise performance of the stages together.

### 18. Sidebands and Phasors

One can separate noise near the carrier into AM and PM components. 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 will 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, 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.

The shape and orientation of the ellipse is determined by the relative size of the AM and PM components in the noise.

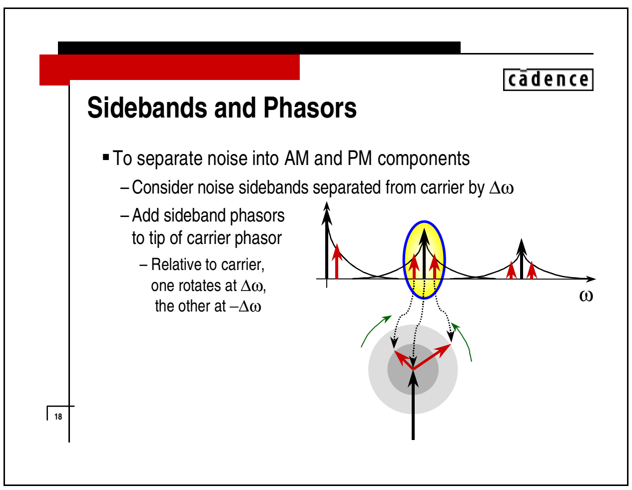


Slide 18

### 19. 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.

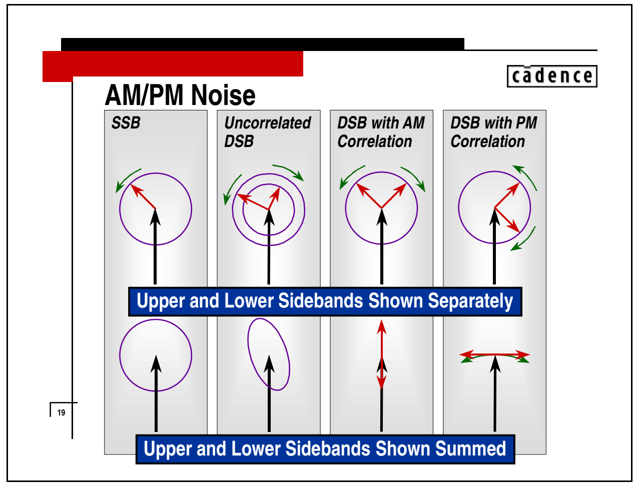


Slide 19

### 20. 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 causes the AM noise to be suppressed, leaving mainly the PM noise.

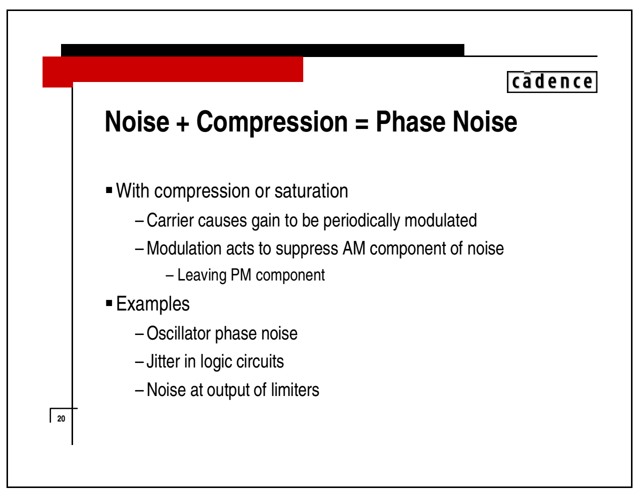


Slide 20

### 21. 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. This noise is referred to as oscillator phase noise.

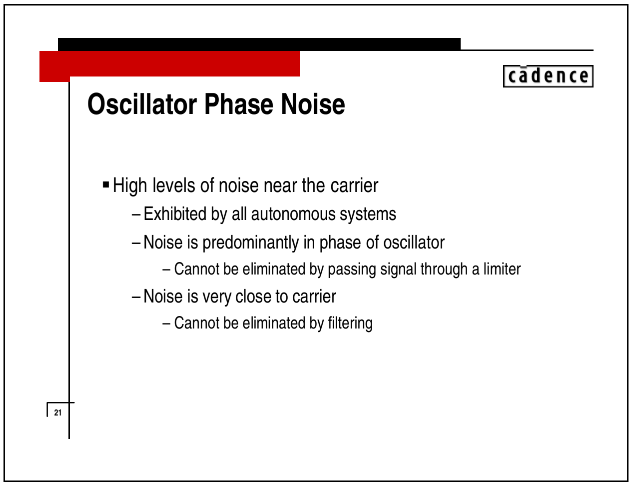


Slide 21

### 22. 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.

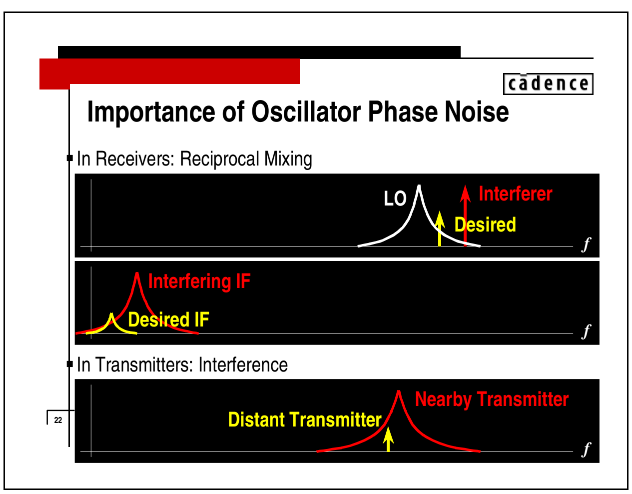


Slide 22

### 23. Noise in a Linear Oscillator

Consider a feedback oscillator with a loop gain of $H(j\omega)$. $X(j\omega)$ is taken to represent some perturbation stimulus and $Y(j\omega)$ 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 $\omega_0$. The gain from the perturbation stimulus to the output is $Y(j\omega)/X(j\omega) = H(j\omega)/H(j\omega)-1$, which goes to infinity at the oscillation frequency $\omega_0$.

The amplification near the oscillation frequency is quantified by assuming the loop gain varies smoothly as a function of frequency in this region. If $\omega = \omega_0 + \Delta\omega$, then $H(j\omega) \approx H(j\omega_0) + dH/d\omega\ \Delta\omega$ and the transfer function becomes $Y(j\omega{+}\Delta\omega)/X(j\omega{+}\Delta\omega) \approx (H(j\omega) + dH/d\omega\ \Delta\omega)\ /\ (H(j\omega) + dH/d\omega\ \Delta\omega - 1)$. Since $H(j\omega_0) = 1$ and $dH/d\omega\ \Delta\omega << 1$ in most practical situations, the transfer function reduces to $Y(j\omega{+}\Delta\omega)/X(j\omega{+}\Delta\omega) \approx 1/(dH/d\omega\ \Delta\omega)$.

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

The amplification of noise near the carrier frequency is created by a linear phenomenon that is an natural consequence of the oscillator's complex pole pair on the $j\omega$ axis at $\omega_0$.

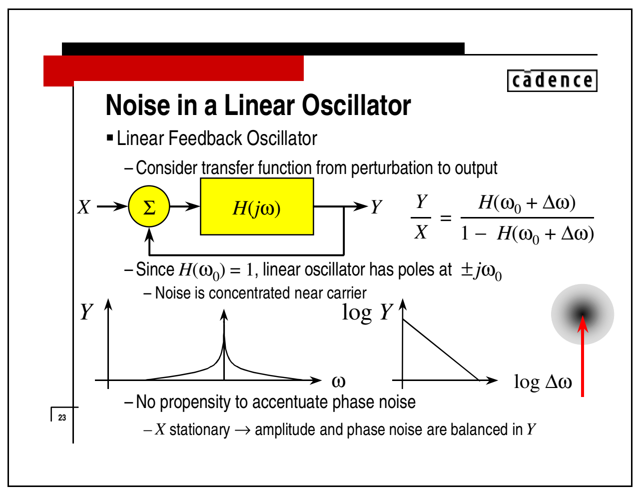


Slide 23

### 24. Barkhausen Criteria for Oscillation — The Origins of Phase Noise

The Barkhausen criterion for oscillation in a feedback oscillator is 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 greater than one, which causes the oscillation amplitude to grow until the amplifier goes into compression far enough so that the effective loop gain equals 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. 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*.

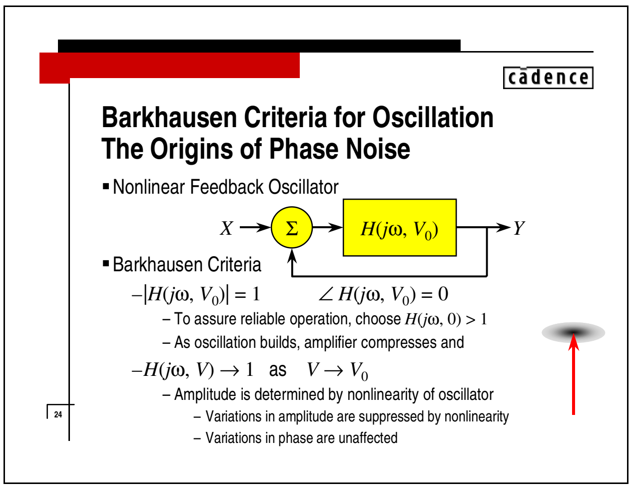


*Slide 24*

### 25. 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 $x$ and assume that the response to the perturbation is $\Delta y$. Separate $\Delta y$ 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 perturbed 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 $t \to \infty$).

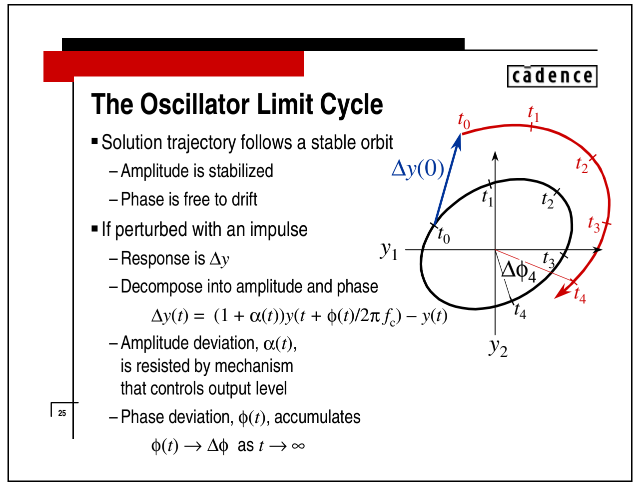


*Slide 25*

### 26. The Oscillator Limit Cycle (cont.)

On the previous slide it was pointed out that after an oscillator has been perturbed by an impulse, $\alpha(t) \to 0$ 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 with $1/f^2$ at frequencies close to the carrier frequency.

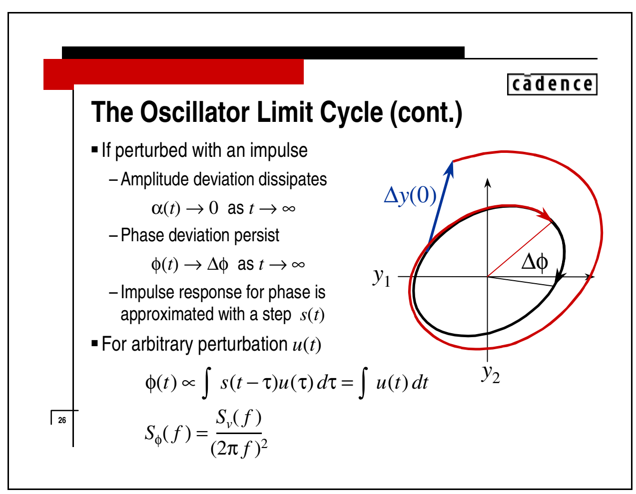


*Slide 26*

## 27. Difference Between $S_\phi$ and $S_v$ Noise

There are two different ways of characterizing noise in the same oscillator. $S_\phi$ 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 $S_\phi$ is only observable if the signal is first passed through a phase detector. Another measure of oscillator noise is $\mathcal{L}$, 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. $S_v$ has a corner frequency that defines its linewidth. The more phase noise, the broader the linewidth (the higher the corner frequency), and the lower the 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).

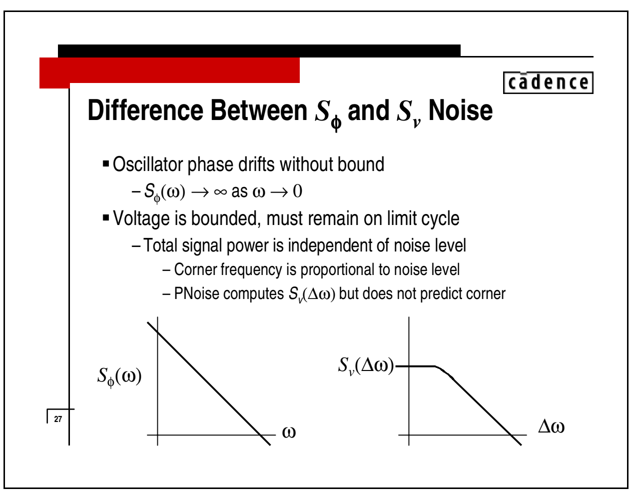


*Slide 27*

## 28. 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 $f = 1/T$. Thus, noise at $f_1$ is correlated with noise at $f_2$ if $f_2 = f_2 + kf$, where $k$ is an integer, and not otherwise. However, the frequency 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. Instead, the correlation is a set of smeared impulses. Thus, noise at $f_1$ is correlated with $f_2$ if $f_2 = f_2 + kf$, where $k$ is close to being integer. The correlation impulses have a finite linewidth just like the phase noise itself. Thus, technically, the noise produced by oscillators is not cyclostationary. Though in practice, the distinction does not appear to be very significant.

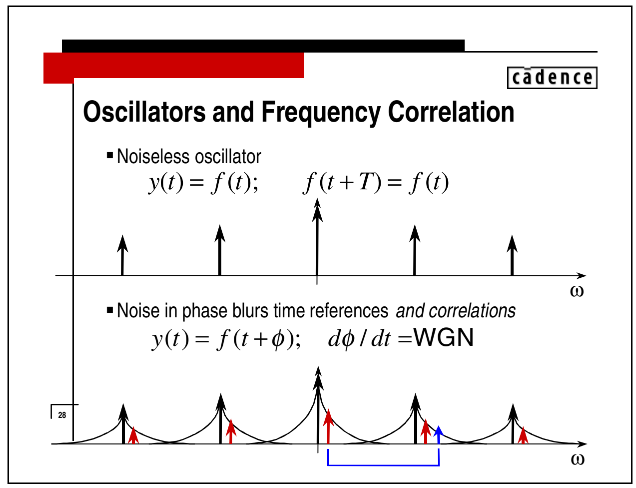


*Slide 28*

## 29. Summary

This presentation reviewed some basic noise concepts and then introduced how a periodically-varying operating point results in the generation of cyclostationary noise. Cyclostationary noise is modulated noise, and it was shown that in the frequency domain this implies that the noise is correlated for frequency offsets of exactly $kf_0$. The duality of shape and correlation in colored and cyclostationary noise was pointed out. There are different ways of characterizing cyclostationary noise, and two were discussed further: the time-averaged PSD, which ignores the cyclostationary nature of the noise, and the PSD along with frequency correlations. It was shown when it is possible to use the simpler time-averaged PSD and when it is not. Finally, cyclostationary noise can be decomposed into AM and PM components. How this happens was illustrated and it was pointed out that circuits that exhibit compression naturally convert noise into phase noise. Compression is a inherent part of any oscillator, and this is why they create phase noise. Oscillator phase noise was described.

Missing from this presentation is a third way in which cyclostationary noise is characterized, as a noise versus time. This is useful when considering sampled-data circuits (such as switched-capacitor filters and sample-and-holds) and thresholding circuits (such as comparators and digital logic). It is hoped that this presentation will be expanded to include these topics and presented at CICC 2000.

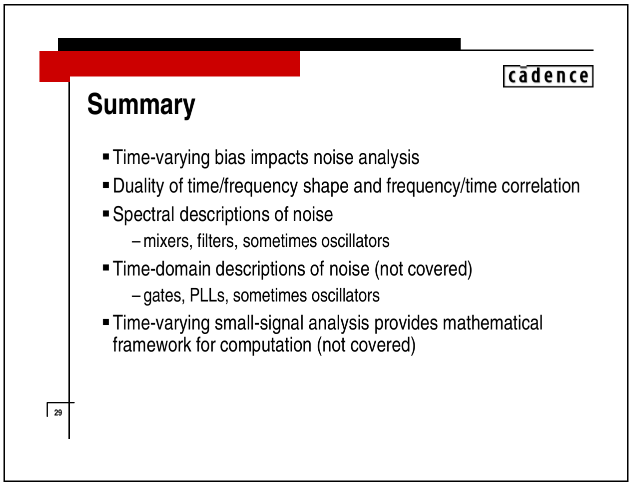


*Slide 29*