# **OSCILLATOR PHASE NOISE: A TUTORIAL (Invited)**

Thomas H. Lee Stanford University http://www-smirc.stanford.edu

#### Abstract

An understanding of how device noise becomes phase noise in practical oscillators is complicated by the presence of nonlinearities (for amplitude stabilization), and by the failure of time invariance. This tutorial reviews the qualitative insights of older (linear, time-invariant) models, and supplements those with powerful additional insights provided by a recently developed time-varying model. Among the most significant are the importance of symmetry in suppressing the upconversion of 1/f noise into noise near the carrier, and an appreciation of cyclostationary effects.

#### Introduction

Circuit and device noise can perturb both the amplitude and phase of an oscillator's output. Because all practical oscillators possess some amplitude-limiting mechanism, amplitude fluctuations are thus usually greatly attenuated. We therefore focus exclusively on *phase noise* in this tutorial.

We begin by identifying some fundamental trade-offs among key parameters, such as power dissipation, oscillation frequency, resonator Q and noise. After studying these trade-offs qualitatively in a hypothetical ideal oscillator, we consider quantitatively how various noise processes corrupt the output spectrum of real oscillators.

## **General Considerations**

A lossless resonator oscillates on its own, given some initial energy. In practice, of course, all real resonators are lossy to some extent (and we also presumably want to drive a load eventually), so we need to arrange for a way to replace the lost energy in order to make an oscillator. Suppose for the moment that the resonator is connected to an energy restoring element that has the remarkable (and unrealizable) property that it is noiseless:

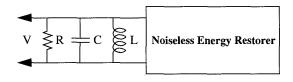


FIGURE 1. "Perfectly efficient" RLC oscillator

The noiseless energy restorer here supplies just enough energy to the tank to compensate for the dissipation by the tank's resistance, thereby leading to a constant amplitude oscillation. The tank resistance is the only noisy element in this model. To gain some useful design insight, first compute the signal energy stored in the tank:

$$E_{stored} = \frac{1}{2}CV_{pk}^2 \tag{1}$$

so that the mean-square signal (carrier) voltage is:

$$\overline{V_{sig}^2} = \frac{E_{stored}}{C} \tag{2}$$

where we have assumed a sinusoidal waveform.

The only source of noise in our idealized model is the tank resistance. Its total mean-square noise voltage is found by integrating its thermal noise density over the noise bandwidth of the *RLC* filter:

$$\overline{V_n^2} = 4kTR \int_0^\infty \left| \frac{Z(f)}{R} \right|^2 df = 4kTR \cdot \frac{1}{4RC} = \frac{kT}{C}$$
 (3)

Combining Eqn. 2 and Eqn. 3, we obtain a noise-to-signal ratio (the reason for this "upside-down" ratio is one of convention, as will be seen shortly):

$$\frac{N}{S} = \frac{\overline{V_n^2}}{\overline{V_{sig}^2}} = \frac{kT}{E_{stored}} \tag{4}$$

Sensibly enough, one therefore needs to maximize the signal levels to minimize the noise-to-carrier ratio.

We may bring power consumption and resonator Q explicitly into consideration by noting that Q can be generally defined as proportional to the energy stored, divided by the energy dissipated:

$$Q = \frac{\omega E_{stored}}{P_{diss}} \tag{5}$$

Therefore,

$$\frac{N}{S} = \frac{\omega kT}{QP_{diss}} \tag{6}$$

The power consumed by our "perfectly efficient" oscillator is simply equal to  $P_{diss}$ , the amount dissipated by the tank loss. For such an oscillator, the noise-to-carrier ratio is inversely proportional to the product of resonator Q and the power consumed, and directly proportional to the oscillation frequency. This set of relationships still holds approximately for real oscillators, and explains the near

obsession of engineers with maximizing resonator Q, for example.

# **Detailed Considerations: Phase Noise**

To augment the qualitative insights of the foregoing analysis, let us now determine the actual output spectrum of the ideal oscillator.

# A. Phase Noise of an Ideal Oscillator

Assume that the output in Fig. 1 is the voltage across the tank, as shown. The only source of noise is the white thermal noise of the tank resistance, which we represent in Norton form as a current source across the tank with a mean-square spectral density of

$$\frac{\vec{i}_n^2}{\Delta f} = 4kTG \tag{7}$$

where G is the reciprocal of the tank resistance.

This current noise becomes voltage noise when multiplied by the impedance facing the current source. In computing this impedance, however, it is important to recognize the contribution of the energy restoration element. Since, by postulate, the circuit oscillates with a constant amplitude, the energy restoration element must contribute an average effective negative resistance that just cancels the positive resistance of the tank. Hence, the net result is that the effective impedance seen by the noise current source is simply that of a perfectly lossless LC network.

For relatively small displacements  $\Delta\omega$  from the center frequency  $\omega_0$ , the impedance of an LC tank may be approximated by

$$Z(\omega_o + \Delta\omega) \approx j \cdot \frac{\omega_o L}{2\frac{\Delta\omega}{\omega}} \quad . \tag{8}$$

We may write the impedance in a more useful form by incorporating an expression for the unloaded tank Q:

$$Q = \frac{R}{\omega_o L} = \frac{1}{\omega_o GL} \tag{9}$$

Solving Eqn. 9 for L and substituting into Eqn. 8 yields:

$$\left| Z(\omega_o + \Delta \omega) \right| \approx \frac{1}{G} \cdot \frac{\omega_o}{2Q\Delta\omega}$$
 (10)

All we've done is exchanged an explicit dependence on inductance with a dependence on Q and G.

Next, we multiply the spectral density of the mean-square noise current by the squared magnitude of the tank impedance to obtain the spectral density of the mean-square noise voltage:

$$\frac{\overline{v_n^2}}{\Delta f} = \frac{\overline{i_n^2}}{\Delta f} \cdot |Z|^2 = 4kTR \left(\frac{\omega_o}{2Q\Delta\omega}\right)^2$$
 (11)

The spectral density of the noise is now frequency dependent because of the filtering action of the tank, and in fact appears to increase without bound as the frequency approaches  $\omega_0$ . Note also that an increase in tank Q reduces the noise density, when all other parameters are held constant, underscoring once again the value of increasing resonator Q.

In our idealized *LC* model, thermal noise affects both amplitude and phase, and Eqn. 11 accounts for both of these effects. However, all practical oscillators employ some form of amplitude limiting, as noted previously. Consequently, phase fluctuations dominate in all well-designed oscillators. The *equipartition theorem* of thermodynamics tells us that, in equilibrium, noise energy splits evenly between amplitude and phase domains. Amplitude limiting thus removes half the noise given by Eqn. 11.

It is traditional to normalize the mean-square noise voltage density to the mean-square carrier voltage, and report the ratio in decibels, thereby explaining the "upside down" ratios presented previously. Performing this normalization yields the following equation for phase noise:

$$L\{\Delta\omega\} = 10\log \left[ \frac{2kT}{P_{sio}} \cdot \left( \frac{\omega_o}{2Q\Delta\omega} \right)^2 \right]$$
 (12)

We see that Eqn. 12 is thus proportional to the log of a density. Its units are commonly expressed as "decibels below the carrier per hertz", or dBc/Hz, specified at a particular offset frequency  $\Delta\omega$  from the carrier frequency  $\omega_o$ . For example, one might speak of a 1GHz oscillator's phase noise as "–110dBc/Hz at a 100kHz offset." It is important to note that the "per Hz" actually applies to the argument of the log, not to the log itself; doubling the measurement bandwidth does not double the decibel quantity. Misleading as "dBc/Hz" is, it is common usage (1).

Eqn. 12 tells us that phase noise (at a given offset) improves as both the carrier power and Q increase, as predicted earlier. It also shows that the noise varies as the inverse square of the frequency offset. These dependencies make sense. Increasing the signal power improves the ratio simply because the thermal noise is fixed, while increasing Q (or the offset) improves the ratio quadratically because the tank's impedance falls off as  $1/Q\Delta\omega$ , and the square of the noise voltage is proportional to the square of the impedance.

Because of the many simplifying assumptions that have led us to this point, real oscillators conform only approximately to these expectations. It should not be surprising that there are some significant differences between the spectrum predicted by Eqn. 12 and what one typically measures in practice. For example, although real spectra do possess a region where the observed density is proportional to  $1/(\Delta\omega)^2$ , the magnitudes are typically quite a bit larger than predicted by Eqn. 12, mainly because there are additional important noise sources in practical oscillators besides than tank loss. For example, any physical implementation of an energy restoration device will contain noisy elements. Furthermore, measured spectra eventually flatten out for large frequency offsets, rather than continuing to drop quadratically. Such a floor may be due to the noise associated with any active elements (such as buffers) placed between the tank and the outside world, or it can even reflect limitations in the measurement instrumentation itself. Even if the output were taken directly from the tank, any resistance in series with either the inductor or capacitor would impose a bound on the amount of filtering provided by the tank at large frequency offsets and thus ultimately produce a noise floor.

Finally, there is almost always a  $1/(\Delta\omega)^3$  region at small offsets. In an effort to account for all three discrepancies, Leeson proposed a widely used, but ultimately *ad hoc*, modification to Eqn. 12 (2):

$$L\{\Delta\omega\} = 10\log\left[\frac{2FkT}{P_{sig}}\left\{1 + \left(\frac{\omega_o}{2Q\Delta\omega}\right)^2\right\} \left(1 + \frac{\Delta\omega_{1/f^3}}{|\Delta\omega|}\right)\right]$$
(13)

His modifications to Eqn. 12 consist of a factor F to account for the increased noise in the  $1/(\Delta\omega)^2$  region, an additive factor of unity (inside the braces) to account for the noise floor, and a multiplicative factor (the term in the second set of parentheses) to provide a  $1/|\Delta\omega|^3$  behavior at sufficiently small offset frequencies. With these modifications, the phase noise spectrum appears as follows:

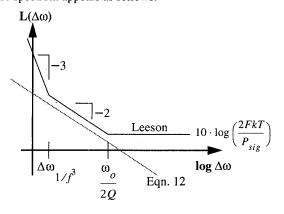


FIGURE 2. Phase noise: Leeson vs. Eqn. 12

Here, the factor F is empirical and therefore must be determined from measurements, diminishing greatly the predictive power of the phase noise equation. Furthermore, the Leeson model asserts that  $\Delta\omega_{1/f}$ 3, the boundary between the  $1/(\Delta\omega)^2$  and  $1/|\Delta\omega|^3$  regions, is precisely equal to the 1/f corner of device noise. However, measurements frequently show no such equality, and thus one must generally treat  $\Delta\omega_{1/f}$ 3 as an empirical fitting parameter as well. Finally, the frequency at which the noise flattens out is not always equal to half the resonator bandwidth,  $\omega_0/2Q$ .

Both the ideal oscillator model and the Leeson model suggest that increasing resonator Q and signal amplitude are the only ways to reduce phase noise. Unfortunately, nothing in Leeson's model guides us in the computation or reduction of F, and we have already noted that  $\Delta\omega_{1/f}3$  is an empirical factor as well.

That neither Eqn. 12 nor Eqn. 12 can make quantitative predictions about phase noise tells us that at least some of the assumptions used in the derivations are invalid, despite their apparent reasonableness. To develop a correct theory, we need to revisit and revise these assumptions.

## A Linear, Time-Varying (LTV) Phase Noise Theory

The foregoing derivations have all assumed linearity and time invariance. Let's reconsider each of these assumptions in turn.

Nonlinearity is clearly a fundamental property of all real oscillators, as it is necessary for amplitude limiting. Several phase noise theories have consequently attempted to explain certain features of phase noise as a consequence of nonlinear behavior. One of these features is that a singlefrequency sinusoidal disturbance injected into an oscillator gives rise to two distinct sidebands, symmetrically disposed about the carrier, and with equal amplitudes. Since a linear, time-invariant (LTI) system can only produce responses at the same frequencies as those of the input, and nonlinearities are well known to exhibit behavior qualitatively similar to that observed, nonlinear mixing has been proposed to explain the sidebands and, by extension, phase noise. Unfortunately, this facile chain of reasoning implies that the amplitude of the sidebands must then depend nonlinearly on the amplitude of the injected signal, and this dependency is simply not observed (except at absurdly large injection amplitudes). We must conclude, therefore, that nonlinear phenomena cannot explain the discrepancies observed between the Leeson model and real oscillator behavior, despite their initial attractiveness as the culprit.

As we shall see momentarily, nonlinearities affect phase noise only *incidentally*, through controlling the detailed shape of the output waveform. An important insight is that both amplitude and phase disturbances are *perturbations* 

superimposed on the main oscillation. They will always be much smaller in magnitude than the carrier in any oscillator worth designing or analyzing. Thus, if a certain amount of injected noise produces a certain amount of amplitude or phase disturbance, we ought to expect that doubling the injected noise would produce double the disturbance. Thus linearity would appear to be a reasonable assumption as far as the noise-to-phase transfer function is concerned, so we will continue to invoke it. As with all assumptions, this one must be tested (and we will do so).

We are left only with the assumption of time invariance to reexamine. Here we consider time invariance to extend to the noise sources themselves, where previously we have implicitly assumed *stationarity*, meaning that the measures that characterize noise (e.g., spectral density) are time-invariant. In contrast with linearity, the assumption of time invariance is less obviously defensible. In fact, it is surprisingly simple to demonstrate that oscillators are fundamentally time-varying systems. Recognizing this truth is the main key to developing a correct theory of phase noise (3).

To show that time-invariance fails to hold, consider explicitly how an impulse of current affects the waveform of the simplest resonant system, a lossless *LC* tank:

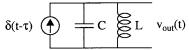


FIGURE 3. LC oscillator excited by current pulse

Assume that the system is oscillating with some constant amplitude until the impulse occurs, then consider how the system responds to an impulse injected at two different times:

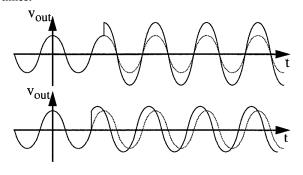


FIGURE 4. Impulse responses of LC tank

If the impulse happens to coincide with a voltage maximum (as in the upper plot), the amplitude increases abruptly by an amount  $\Delta V = \Delta Q/C$ , but because the response to the impulse superposes exactly in phase with the pre-existing oscillation, the timing of the zero crossings does not change. On the

other hand, an impulse injected at some other time generally affects both the amplitude and the timing of the zero crossings, as in the lower plot. Interpreting the zero-crossing timings as a measure of phase, we see that the amount of phase disturbance for a given injected impulse depends on when the injection occurs; time-invariance thus fails to hold. An oscillator is therefore a linear, but (periodically) time varying (LTV) system.

To the extent that linearity remains a good assumption, the impulse response still completely characterizes the system, even with time variation thrown in. Noting that an impulsive input produces a step change in phase, the impulse response may be written as:

$$h_{\phi}(t,\tau) = \frac{\Gamma(\omega_o \tau)}{q_{max}} u(t-\tau)$$
 (14)

where u(t) is the unit step. Dividing by  $q_{max}$ , the maximum charge displacement across the capacitor, makes the function  $\Gamma(x)$  independent of signal amplitude.  $\Gamma(x)$  is called the *impulse sensitivity function* (ISF), and is a dimensionless, frequency- and amplitude-independent function periodic in  $2\pi$ . As its name suggests, it encodes information about the sensitivity of the system to an impulse injected at phase  $\omega_0 \tau$ . In our example of the LC oscillator,  $\Gamma(x)$  has its maximum value near the zero crossings of the oscillation, and a zero value at maxima of the oscillation waveform. In general, it is most practical to determine  $\Gamma(x)$  through simulation, but there are also analytical methods that apply in special cases (4). In any event, to develop a feel for typical shapes of ISFs, consider two representative examples, first for an LC oscillator:

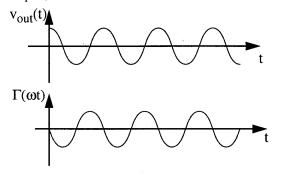


FIGURE 5. Example ISF for LC oscillator

Note that, in this case, the ISF is *approximately* proportional to the derivative of the oscillation waveform itself, a

relationship that holds crudely even for other types of oscillators:

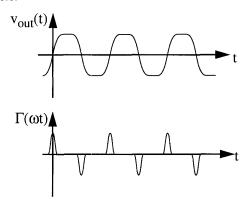


FIGURE 6. Example ISF for ring oscillator

Once the ISF has been determined (by whatever means), we may compute the excess phase through use of the superposition integral. This computation is valid here since superposition is linked to linearity, not time invariance:

$$\phi(t) = \int_{-\infty}^{\infty} h_{\phi}(t, \tau) i(\tau) d\tau = \frac{1}{q_{max}} \int_{-\infty}^{t} \Gamma(\omega_{o} \tau) i(\tau) d\tau \quad (15)$$

To cast this equation in a more practically useful form, note that the ISF is periodic and therefore expressible as a Fourier series:

$$\Gamma(\omega_o \tau) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n \cos(n\omega_o \tau + \theta_n)$$
 (16)

where the coefficients  $c_n$  are real, and  $\theta_n$  is the phase of the *n*th harmonic of the ISF. We will ignore  $\theta_n$  in all that follows because we will be assuming that noise components are uncorrelated, so that their relative phase is irrelevant.

The value of this decomposition is that, like many functions associated with physical phenomena, the series typically converges rapidly, so that it is often well approximated by just the first few (e.g., two) terms of the series.

Substituting the Fourier expansion into Eqn. 15, and exchanging summation and integration, one obtains:

$$\phi(t) = \frac{1}{q_{max}} \left[ \frac{c_0}{2} \int_{-\infty}^{t} i(\tau) d\tau + \sum_{n=1}^{\infty} c_n \int_{-\infty}^{t} i(\tau) \cos(n\omega_o \tau) d\tau \right]$$
(17)

This equation allows us to compute the excess phase caused by an arbitrary noise current injected into the system, once the Fourier coefficients of the ISF have been determined. Earlier, we noted that signals (noise) injected into a nonlinear system at some frequency may produce spectral components at a different frequency. We now show that a linear, but time-varying system can exhibit qualitatively similar behavior. To demonstrate this property explicitly, consider injecting a sinusoidal current whose frequency is near an integer multiple, m, of the oscillation frequency, so that

$$i(t) = I_m \cos \left[ \left( m\omega_o + \Delta\omega \right) t \right] \tag{18}$$

where  $\Delta \omega \ll \omega$ . Substituting Eqn. 18 into Eqn. 15 and noting there is a negligible net contribution to the integral by terms other than when n = m, one obtains the following approximation:

$$\phi(t) \approx \frac{I_m c_m \sin(\Delta \omega t)}{2q_{max} \Delta \omega}$$
 (19)

The spectrum of  $\phi(t)$  therefore consists of two equal sidebands at  $\pm\Delta\omega$ , even though the injection occurs near some integer *multiple* of  $\omega_0$ . We see that we do not need to invoke nonlinearity to explain this frequency conversion (or "folding"). This observation is fundamental to understanding the evolution of noise in an oscillator.

Unfortunately, we're not quite done: Eqn. 19 allows us to figure out the spectrum of  $\phi(t)$ , but we ultimately want to find the spectrum of the output voltage of the oscillator, which is not quite the same thing. The two quantities are linked through the actual output waveform, however. To illustrate what we mean by this linkage, consider a specific case where the output may be approximated as a sinusoid, so that  $v_{out}(t) = \cos{[\omega_0 t + \phi(t)]}$ . This equation may be considered a phase-to-voltage converter; it takes phase as an input, and produces from it the output voltage. This conversion is fundamentally nonlinear because it involves the phase modulation of a sinusoid.

Performing this phase-to-voltage conversion, and assuming "small" amplitude disturbances, we find that the single-tone injection leading to Eqn. 19 results in two equal-power sidebands symmetrically disposed about the carrier:

$$P_{SBC}(\Delta\omega) \approx 10 \cdot \log \left( \frac{I_m c_m}{4q_{max} \Delta\omega} \right)^2$$
 (20)

To distinguish this result from nonlinear mixing phenomena, note that the amplitude dependence is linear (the squaring operation simply reflects the fact that we are dealing with a power quantity here).

The foregoing result may be extended to the general case of a white noise source:

$$P_{SBC}(\Delta\omega) \approx 10 \cdot \log \left( \frac{\overline{i_n^2}}{\Delta f} \sum_{m=0}^{\infty} c_m^2 \right)$$
 (21)

Eqn. 21 implies both upward and downward frequency translations of noise into the noise near the carrier, as illustrated in the following figure:

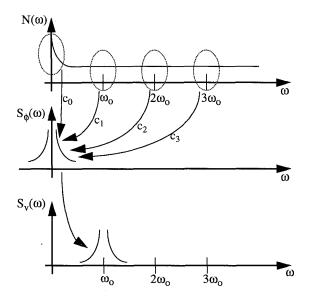


FIGURE 7. Evolution of circuit noise into phase noise

This figure summarizes what the foregoing equations tell us: Components of noise near *integer multiples* of the carrier frequency all fold into noise near the carrier itself.

Noise near DC gets *upconverted*, weighted by coefficient  $c_0$ , so 1/f device noise becomes  $1/f^2$  noise near the carrier; noise near the carrier stays there, weighted by  $c_1$ ; and white noise near higher integer multiples of the carrier undergoes *down-conversion*, turning into noise in the  $1/f^2$  region. Note that the  $1/f^2$  shape results from the integration implied by the step change in phase caused by an impulsive noise input. Since an integration (even a time-varying one) gives a white voltage or current spectrum a 1/f character, the power spectral density will have a  $1/f^2$  shape.

It is clear from Fig. 7 that minimizing the various coefficients  $c_n$  (by minimizing the ISF) will minimize the phase noise. To underscore this point quantitatively, we may use Parseval's theorem to write:

$$\sum_{n=0}^{\infty} c_m^2 = \frac{1}{\pi} \int_{0}^{2\pi} |\Gamma(x)|^2 dx = 2\Gamma_{rms}^2$$
 (22)

so that the spectrum in the  $1/f^2$  region may be expressed as:

$$L(\Delta \omega) = 10 \cdot \log \left( \frac{\frac{\overline{i_n^2}}{\Delta f} \Gamma_{rms}^2}{2q_{max}^2 \Delta \omega^2} \right)$$
 (23)

where  $\Gamma_{rms}$  is the rms value of the ISF. All other factors held equal, reducing  $\Gamma_{rms}$  will reduce the phase noise at all frequencies. Eqn. 21 (or Eqn. 23) is the rigorous equation for the  $1/f^2$  region, and is one key result of the LTV model. Note that, unlike the Leeson model, no empirical curve-fitting parameters are present in Eqn. 21.

Among other attributes, Eqn. 21 allows us to study quantitatively the upconversion of 1/f noise into close-in phase noise. Noise near the carrier is particularly important in communication systems with narrow channel spacings. In fact, the allowable channel spacings are frequently constrained by the achievable phase noise. Unfortunately, it is not possible to predict close-in phase noise correctly with LTI models.

This problem disappears if the LTV model is used. Specifically, assume that the current noise behaves as follows in the 1/f region:

$$\overline{i_{n,1/f}^2} = \overline{i_n^2} \cdot \frac{\omega_{1/f}}{\Delta \omega} \tag{24}$$

where  $\omega_{1/f}$  is the 1/f corner frequency. Substitution into Eqn. 21 gives us

$$L(\Delta\omega) = 10 \cdot \log \left( \frac{\frac{\overline{l_n^2}}{\sqrt{l_n^2}} c_0^2}{8q_{max}^2 \Delta\omega^2} \cdot \frac{\omega_{1/f}}{\Delta\omega} \right)$$
 (25)

which describes the phase noise in the  $1/f^3$  region. The  $1/f^3$  corner frequency is then

$$\Delta \omega_{1/f^3} = \omega_{1/f} \cdot \frac{c_0^2}{4\Gamma_{pmr}^2} \approx \omega_{1/f} \cdot \left(\frac{c_0}{c_1}\right)^2 \tag{26}$$

from which we see that the  $1/f^3$  phase noise corner is not necessarily the same as the 1/f device/circuit noise corner; it will generally be lower. In fact, since  $c_0$  is the DC value of the ISF, there is a possibility of reducing by large fac-

tors the  $1/f^3$  phase noise corner. The ISF is a function of the waveform, and hence potentially under the control of the designer. This result is not anticipated by LTI approaches, and is one of the most powerful insights conferred by this LTV model. This result has particular significance for technologies with notoriously poor 1/f noise performance, such as CMOS and GaAs MESFETs. A specific circuit example of how one may exploit this observation follows shortly.

One more extremely powerful insight concerns the influence of *cyclostationary* noise sources. In most oscillators, the noise sources cannot be well modeled as stationary. A typical example is the nominally white collector shot noise (or MOSFET drain current noise), which vary because device currents vary periodically with the oscillating waveform. The LTV model is able to accommodate a cyclostationary white noise source with ease, since such a source may be treated as the product of a stationary white noise source and a periodic function:

$$i_n(t) = i_{n0}(t) \cdot \alpha(\omega_0 t) \tag{27}$$

Here,  $i_{n0}$  is a stationary white noise source whose peak value is equal to that of the cyclostationary source, and  $\alpha(x)$  is a periodic function with a peak value of unity. Substituting this into Eqn. 15 allows us to treat cyclostationary noise as a stationary noise source provided we define an effective ISF as follows:

$$\Gamma_{eff}(x) = \Gamma(x) \cdot \alpha(x)$$
 (28)

Thus, none of the foregoing conclusions changes as long as  $\Gamma_{eff}$  is used in all of the equations.

Having identified the factors that influence oscillator noise, we're now in a position to articulate the requirements that must be satisfied to make a good oscillator. First, note that an active device is always necessary to compensate for tank loss, and that active devices always contribute noise. Note also that the ISFs tell us that there are sensitive and insensitive moments in an oscillation cycle. Of the many possible ways that an active element could return energy to the tank, then, this energy should be delivered all at once, at the peak of the tank voltage, where the ISF has its minimum value. In an ideal oscillator, therefore, the transistor would deliver an impulse of current at the peak(s), then go into a coma each cycle. The extent to which real oscillators approximate this behavior determines the quality of their phase noise properties. Since an LTI theory treats all instants as equally important, such theories are unable to anticipate this important result.

Let us examine a typical oscillator, now that we have developed these insights:

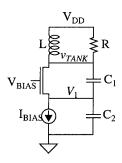


FIGURE 8. Colpitts oscillator (simplified)

The relevant waveforms for this oscillator appear approximately as follows:

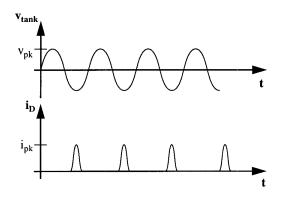


FIGURE 9. Approximate incremental tank voltage and drain current for Colpitts oscillator

Note that the drain current flows only during a short interval coincident with the peaks of the tank voltage. Hence, a Colpitts oscillator approximates ideal behavior. Its corresponding excellent phase noise properties account for the popularity of this configuration. It has long been known that the best phase noise occurs for a certain narrow range of tapping ratios (e.g., a 4:1 capacitance ratio), but before the LTV theory, no theoretical basis existed to explain a particular optimum.

As an example of a circuit that does *not* well approximate ideal behavior, consider a ring oscillator. First, the "resonator" Q is poor; in fact, it is unity, since the energy stored in the node capacitances is reset (discharged) every cycle. Hence, if the resonator of a Colpitts oscillator is a fine wine glass, the resonator of a ring oscillator is a lump of clay. Next, energy is restored to the resonator during the edges (the worst possible times), rather than the voltage maxima. These factors account for the well-known terrible

phase noise performance of ring oscillators. As a consequence, ring oscillators are found only in the most non-critical applications, or inside wideband PLLs that clean up the spectrum.

## **Circuit Examples**

We close with two brief examples that underscore important design insights. First, both LTI and the LTV models point out the value of maximizing signal amplitude. To evade supply voltage or breakdown constraints, one may employ a tapped resonator to decouple resonator swings from device voltage limitations. A common configuration that does so is the Clapp modification to the Colpitts oscillator:

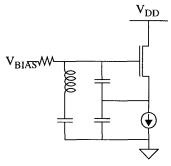


FIGURE 10. Clapp oscillator

Differential versions of this oscillator have recently made an appearance in the literature, without explicit reference to Clapp, or an explanation that its advantage derives from the tapped resonator configuration to allow an increase in signal energy (5).

It was stated earlier that a key insight of the LTV theory concerns the importance of symmetry. A configuration that exploits this knowledge is the symmetrical negative resistance oscillator (6):

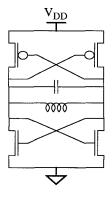


FIGURE 11. Simple symmetrical negative resistance oscillator

This configuration is not new, but an appreciation of its symmetry properties is. Here, it is the half-circuit symmetry that is important, because noise in the two half circuits is only partially correlated at best. By selecting the relative widths of the PMOS and NMOS devices appropriately to minimize the DC value of the ISF for each half-circuit, one may minimize the upconversion of 1/f noise. Through exploitation of symmetry in this manner, the 1/f³corner can be dropped to exceptionally low values (below 100Hz), even though device 1/f noise corners may be well above 10-100kHz. As a result, a phase noise of -121dBc/Hz at an offset of 600kHz has been obtained with on-chip spiral inductors at 1.8GHz, on 6mW of power consumption in a 0.25µm CMOS technology (6). This result rivals what one may achieve with bipolar technologies.

# **Summary**

The insights gained from LTI phase noise models are simple and intuitively satisfying: One should maximize signal amplitude and resonator Q. An additional, implicit insight is that the phase shifts around the loop generally must be arranged so that oscillation occurs at or very near the center frequency of the resonator. This way, there is a maximum attenuation by the resonator of off-center spectral components.

Deeper insights provided by the LTV model are that the resonator energy should be restored impulsively at the voltage maximum, instead of evenly throughout a cycle, and that the DC value of the effective ISF should be made as close to zero as possible to suppress the upconversion of 1/f noise into close-in phase noise.

## Acknowledgments

It is a great pleasure to acknowledge many enlightening discussions about phase noise with Prof. Ali Hajimiri of Caltech during his development of the LTV theory at Stanford.

# References

- (1) T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, Cambridge University Press, 1998.
- (2) D. B. Leeson, "A Simple Model of Feedback Oscillator Noise Spectrum," *Proceedings of the IEEE*, vol. 54, pp. 329-330, Feb. 1966.
- (3) A. Hajimiri and T. Lee, "A General Theory of Phase Noise in Electrical Oscillators," *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179-194, Feb. 1998.
- (4) A. Hajimiri and T. Lee, *The Design of Low-Noise Oscillators*, Kluwer, 1999.
- (5) J. Craninckx and M. Steyaert, "A 1.8GHz CMOS Low-Phase-Noise Voltage-Controlled Oscillator with Prescaler," *IEEE Journal of Solid-State Circuits*, vol. 30, no. 12, pp. 1474-1482, Dec. 1995.
- (6) A. Hajimiri and T. Lee, "Design Issues in CMOS Differential LC Oscillators," *IEEE Journal of Solid-State Circuits*, May. 1999.