Oscillator Phase Noise: A Tutorial

Thomas H. Lee, *Member, IEEE,* and Ali Hajimiri, *Member, IEEE*

Abstract—Linear time-invariant (LTI) phase noise theories provide important qualitative design insights but are limited in their quantitative predictive power. Part of the difficulty is that device noise undergoes multiple frequency translations to become oscillator phase noise. A quantitative understanding of this process requires abandoning the principle of time invariance assumed in most older theories of phase noise. Fortunately, the noise-to-phase transfer function of oscillators is still linear, despite the existence of the nonlinearities necessary for amplitude stabilization. In addition to providing a quantitative reconciliation between theory and measurement, the time-varying phase-noise model presented in this tutorial identifies the importance of symmetry in suppressing the upconversion of $1/f$ noise into close-in phase noise, and provides an explicit appreciation of cyclostationary effects and AM–PM conversion. These insights allow a reinterpretation of why the Colpitts oscillator exhibits good performance, and suggest new oscillator topologies. Tuned *LC* and ring oscillator circuit examples are presented to reinforce the theoretical considerations developed. Simulation issues and the accommodation of amplitude noise are considered in appendixes.

Index Terms—Jitter, low-noise oscillators, noise, noise measurement, noise simulation, oscillators, oscillator noise, oscillator stability, phase jitter, phase-locked loops, phase noise, phase-noise simulation, voltage-controlled oscillators.

I. INTRODUCTION

IN GENERAL, circuit and device noise can perturb both the amplitude and phase of an oscillator's output. Of necessity, however, all practical oscillators inherently possess an amplitude-limiting mechanism of some kind. Because amplitude fluctuations are usually greatly attenuated as a result, phase noise generally dominates. Thus, even though it is possible to design oscillators in which amplitude noise is significant (particularly at frequencies well removed from that of the carrier), we will focus primarily on phase noise in the body of this tutorial. The issue of amplitude noise is considered in detail in the Appendix, as are practical issues related to how one performs simulations of phase noise.

We begin by identifying some very general tradeoffs among key parameters, such as power dissipation, oscillation frequency, resonator Q, and circuit noise power. These tradeoffs are first studied qualitatively in a hypothetical ideal oscillator in which linearity of the noise-to-phase transfer function is assumed, allowing characterization by the impulse response. Although linearity is defensible, time invariance fails to hold even in this simple case. That is, oscillators are linear time-varying (LTV) systems. By studying the impulse response,
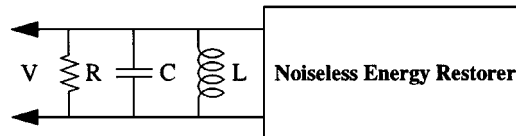
Fig. 1. "Perfectly efficient" *RLC* oscillator.

we discover that periodic time variation leads to frequency translation of device noise to produce the phase-noise spectra exhibited by real oscillators. In particular, the upconversion of $1/f$ noise into close-in phase noise is seen to depend on symmetry properties that are potentially controllable by the designer. Additionally, the same treatment easily subsumes the cyclostationarity of noise generators, and helps explain why class-C operation of active elements within an oscillator may be beneficial. Illustrative circuit examples reinforce key insights of the LTV model.

II. GENERAL CONSIDERATIONS

Perhaps the simplest abstraction of an oscillator that still retains some connection to the real world is a combination of a lossy resonator and an energy restoration element. The latter precisely compensates for the tank loss to enable a constant-amplitude oscillation. To simplify matters, assume that the energy restorer is noiseless (see Fig. 1). The tank resistance is therefore the only noisy element in this model.

To gain some useful design insight, first compute the signal energy stored in the tank

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

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

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

where we have assumed a sinusoidal waveform.

The total mean-square noise voltage is found by integrating the resistor's thermal noise density over the noise bandwidth of the *RLC* resonator

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

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

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

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 defined generally as proportional to the energy stored, divided by the energy dissipated:

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

Therefore

$$\frac{N}{S} = \frac{\omega kT}{Q P_{\text{diss}}}. \tag{6}$$

The power consumed by this model oscillator is simply equal to P_{diss}, the amount dissipated by the tank loss. The noise-to-carrier ratio is here 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.

III. 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. By postulate, the only source of noise is the white thermal noise of the tank conductance, which we represent as a current source across the tank with a mean-square spectral density of

$$\frac{\overline{i_n^2}}{\Delta f} = 4kTG. \tag{7}$$

This current noise becomes voltage noise when multiplied by the effective impedance facing the current source. In computing this impedance, however, it is important to recognize that the energy restoration element must contribute an average effective negative resistance that precisely 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 Δw (called the offset frequency) from the center frequency ω_0, the impedance of an LC tank may be approximated by

$$Z(\omega_0 + \Delta\omega_0) \approx j \cdot \frac{\omega_0 L}{2\frac{\Delta\omega}{\omega_0}}. \tag{8}$$

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

$$Q = \frac{R}{\omega_0 L} = \frac{1}{\omega_0 G L}. \tag{9}$$

Solving (9) for L and substituting into (8) yields

$$|Z(\omega_0 + \Delta\omega| = \frac{1}{G} \cdot \frac{\omega_0}{2Q\Delta\omega}. \tag{10}$$

Thus, we have traded an explicit dependence on inductance for a dependence on Q and G.

Next, 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

$$\frac{\overline{v_n^2}}{\Delta f} = \frac{\overline{i_n^2}}{\Delta f} \cdot |Z|^2 = 4kTR \left(\frac{\omega_0}{2Q\Delta\omega}\right)^2. \tag{11}$$

The power spectral density of the output noise is frequency dependent because of the filtering action of the tank, falling as the inverse-square of the offset frequency. This $1/f^2$ behavior simply reflects the fact that the voltage frequency response of an RLC tank rolls off as $1/f$ to either side of the center frequency, and power is proportional to the square of voltage. 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 (11) includes their combined effect. The equipartition theorem of thermodynamics tells us that, in equilibrium, amplitude and phase-noise power are equal. Therefore, the amplitude limiting mechanism present in any practical oscillator removes half the noise given by (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 the normalized single-sideband noise spectral density:

$$L\{\Delta\omega\} = 10\log\left[\frac{2kT}{P_{sig}} \cdot \left(\frac{\omega_0}{2Q\Delta\omega}\right)^2\right]. \tag{12}$$

These units are thus proportional to the log of a density. Specifically, they 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 ω_0. For example, one might speak of a 2-GHz oscillator' phase noise as "–110 dBc/Hz at a 100-kHz 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. As lacking in rigor as "dBc/Hz" is, it is common usage [1].

Equation (12) tells us that phase noise (at a given offset) improves as both the carrier power and Q increase, as predicted earlier. These dependencies make sense. Increasing the signal power improves the ratio simply because the thermal noise is fixed, while increasing Q improves the ratio quadratically because the tank"s impedance falls off as $1/Q\Delta\omega$.

Because many simplifying assumptions have led us to this point, it should not be surprising that there are some significant differences between the spectrum predicted by (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 (12), because there are additional important noise sources besides tank loss. For example, any physical implementation of an energy restorer will be noisy.
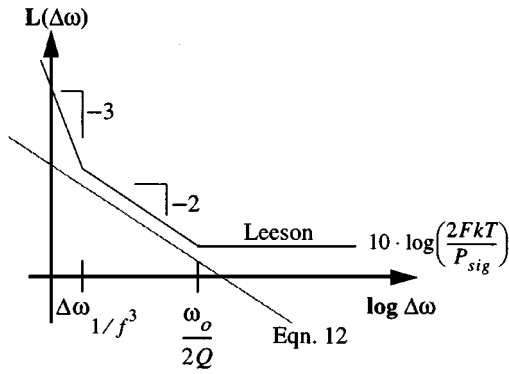
Fig. 2.   Phase noise: Leeson versus (12).

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. Last, there is almost always a $1/(\Delta\omega)^3$ region at small offsets (we will ignore here the eventual flattening of the spectrum at extremely small offsets [4], [15]).

A modification to (12) provides a means to account for these discrepancies

$$L\{\Delta\omega\}$$
$$= 10\log\left[\frac{2FkT}{P_{\text{sig}}}\left\{1 + \left(\frac{\omega_0}{2Q\Delta\omega}\right)^2\right\}\left(1 + \frac{\Delta\omega_{1/f^3}}{|\Delta\omega|}\right)\right].$$
(13)

These modifications, due to Leeson, 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 [2]. With these modifications, the phase-noise spectrum appears as in Fig. 2.

It is important to note that the factor F is an empirical fitting parameter and therefore must be determined from measurements, diminishing the predictive power of the phase-noise equation. Furthermore, the 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. Also, it is not clear what the corner frequency will be in the presence of more than one noise source with $1/f$ noise contribution. Last, 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 ways to reduce phase noise. The Leeson model additionally introduces the factor F, but without knowing precisely what it depends on, it is difficult to identify specific ways to reduce it. The same problem exists with $\Delta\omega_{1/f^3}$ as well. Last, blind application of these models has periodically led to earnest but misguided attempts to use active circuits to boost Q. Sadly, increases in Q through such means are necessarily accompanied by increases in F as well, preventing the anticipated improvements in phase noise. Again, the lack of analytical expressions for F can obscure this conclusion, and one continues to encounter various doomed oscillator designs based on the notion of active Q boosting.

That neither (12) nor (13) can make quantitative predictions about phase noise is an indication that at least some of the assumptions used in the derivations are invalid, despite their apparent reasonableness. To develop a theory that does not possess the enumerated deficiencies, we need to revisit, and perhaps revise, these assumptions.

IV. A Linear, Time-Varying Phase-Noise Theory

The foregoing derivations have all assumed linearity and time invariance. Let' us 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 observations as a consequence of nonlinear behavior. One observation is that a single-frequency sinusoidal disturbance injected into an oscillator gives rise to two equal-amplitude sidebands, symmetrically disposed about the carrier [7]. Since LTI systems cannot perform frequency translation and nonlinear systems can, nonlinear mixing has occasionally been proposed to explain phase noise. Unfortunately, the amplitude of the sidebands must then depend nonlinearly on the amplitude of the injected signal, and this dependency is not observed. One must conclude that memoryless nonlinearity cannot explain the discrepancies, despite initial attractiveness as the culprit.

As we shall see momentarily, amplitude-control nonlinearities certainly do affect phase noise, but only incidentally, by controlling the detailed shape of the output waveform. An important insight is that disturbances are just that: *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 phase disturbance, we ought to expect doubling the injected noise to produce double the disturbance. Linearity would therefore appear to be a reasonable assumption *as far as the noise-to-phase transfer function is concerned.* It is therefore particularly important to keep in mind that, when assessing linearity, it is essential to identify explicitly the input–output variables. Linear relationships may exist between certain variable pairs at the same time nonlinear ones exist between others. Linearization of some of these relationships need not imply linearization of the fundamentally nonlinear behavior of the active devices. Indeed, we will perform a linearization around the steady-state solution, which automatically takes the effect of device nonlinearity into account. There is therefore no contradiction here with the prior acknowledgment of nonlinear amplitude control.

We are left only with the assumption of time invariance to re-examine. In the previous derivations, we have extended time invariance to the noise sources themselves, meaning that the measures that characterize noise (e.g., spectral density) are time invariant (*stationary*). 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 more accurate 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 (Fig. 3). 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, as seen in Fig. 4.

If the impulse happens to coincide with a voltage maximum (as in the left 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 preexisting 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 of oscillation and the timing of the zero crossings, as in the right 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.

Because 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

$$h_\phi(t, \tau) = \frac{\Gamma(\omega_0 \tau)}{q_{\max}} u(t - \tau) \qquad (14)$$

where $u(t)$ is the unit step function. 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 2π. As its name suggests, it encodes information about the sensitivity of the oscillator to an impulse injected at phase $\omega_0 \tau$. In the *LC* oscillator example, $\Gamma(x)$ 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 (and most accurate) to determine $\Gamma(x)$ through simulation, but there are also analytical methods (some approximate) that apply in special cases [4], [8]. In any event, to develop a feel for typical shapes of ISF's, consider two representative examples, first for an *LC* and a ring oscillator in Fig. 5(a) and (b).

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

$$\phi(t) = \int_{-\infty}^{\infty} h_\phi(t, \tau) i(\tau)\, d\tau = \frac{1}{q_{\max}} \int_{-\infty}^{t} \Gamma(\omega_0 \tau) i(\tau)\, d\tau. \qquad (15)$$
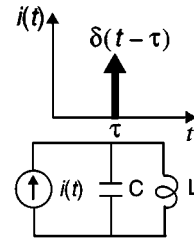

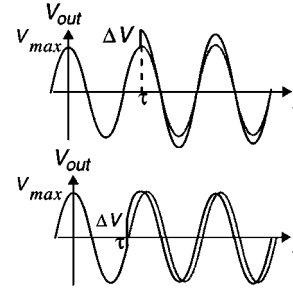Fig. 3.   *LC* oscillator excited by current pulse.

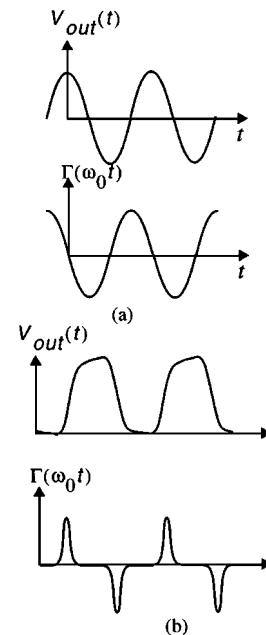Fig. 4.   Impulse responses of *LC* tank.

Fig. 5.   Example ISF for (a) *LC* oscillator and (b) ring oscillator.

This computation can be visualized with the help of the equivalent block diagram shown in Fig. 6.

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

$$\Gamma(\omega_0 \tau) = \frac{c_0}{2} + \sum_{n=1}^{\infty} c_n \cos(n\omega_0 \tau + \theta_n) \qquad (16)$$

where the coefficients c_n are real and θ_n is the phase of the nth harmonic of the ISF. We will ignore θ_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
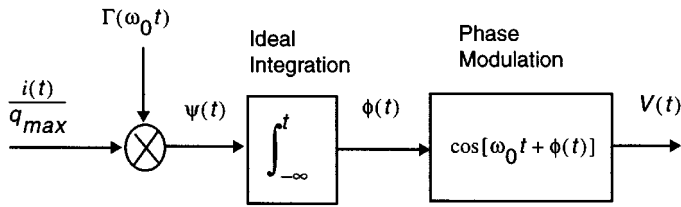
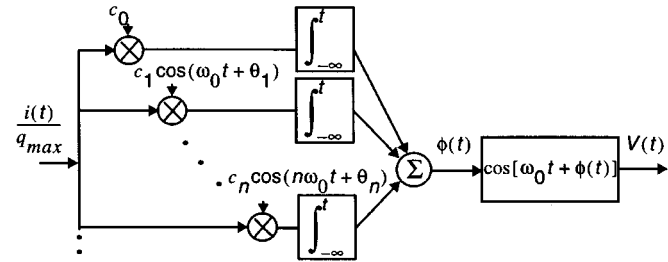Fig. 6.   The equivalent block diagram of the process.

Fig. 7.   The equivalent system for ISF decomposition.

converges rapidly, so that it is often well approximated by just the first few terms of the series.

Substituting the Fourier expansion into (15), and exchanging summation and integration, one obtains

$$\phi(t) = \frac{1}{q_{\max}} \left[\frac{c_0}{2} \int_{-\infty}^{t} i(\tau)\, d\tau + \sum_{n=1}^{\infty} c_n \int_{-\infty}^{t} \cdot\, i(\tau) \cos(n\omega_0 \tau)\, d\tau \right]. \tag{17}$$

The corresponding sequence of mathematical operations is shown graphically in the left half of Fig. 7. Note that the block diagram contains elements that are analogous to those of a superheterodyne receiver. The normalized noise current is a broad-band "RF" signal, whose Fourier components undergo simultaneous downconversions (multiplications) by "local oscillator" signals at all harmonics of the oscillation frequency. It is important to keep in mind that multiplication is a linear operation if one argument is held constant, as it is here. The relative contributions of these multiplications are determined by the Fourier coefficients of the ISF. Equation (17) thus 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, as implied by the superheterodyne imagery invoked in the preceding paragraph. To demonstrate this property explicitly, consider injecting a sinusoidal current whose frequency is near an integer multiple m of the oscillation frequency, so that

$$i(t) = I_m \cos[(m\omega_0 + \Delta\omega)t] \tag{18}$$

where $\Delta\omega \ll \omega_0$. Substituting (18) into (17) and noting that there is a negligible net contribution to the integral by terms

other than when $n = m$, one obtains the following approximation:

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

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 ω_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 are not quite done: (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 $v_{\text{out}}(t) = \cos[\omega_0 t + \phi(t)]$. 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 (19) results in two equal-power sidebands symmetrically disposed about the carrier

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

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). This relationship can be, and has been, verified experimentally.

The foregoing result may be extended to the general case of a white noise source [6]

$$P_{\text{SBC}}(\Delta\omega) \approx 10 \cdot \log \left(\frac{\dfrac{\overline{i_n^2}}{\Delta f} \displaystyle\sum_{m=0}^{\infty} c_m^2}{4 q_{\max}^2 \Delta\omega^2} \right). \tag{21}$$

Equation (20) implies both upward and downward frequency translations of noise into the noise near the carrier, as illustrated in Fig. 8. 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 ultimately becomes $1/f^3$ 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 *downconversion*, 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.
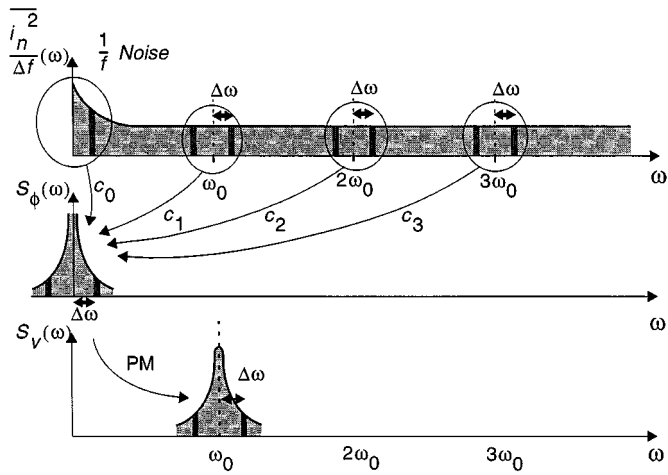
Fig. 8. Evolution of circuit noise into phase noise.

It is clear from Fig. 8 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

$$\sum_{n=0}^{\infty} c_m^2 = \frac{1}{\pi} \int_0^{2\pi} |\Gamma(x)|^2 \, dx = 2\Gamma_{\text{rms}}^2 \tag{22}$$

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

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

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

Among other attributes, (23) 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. Using (23), we obtain the following for the noise in the $1/f^3$ region:

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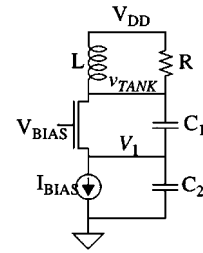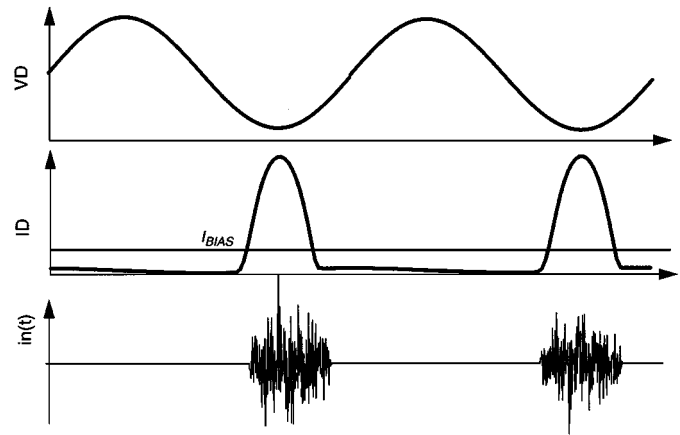

Fig. 9. Colpitts oscillator (simplified).

Fig. 10. Approximate incremental tank voltage and drain current for Colpitts oscillator.

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

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

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 Γ_{dc} is the dc value of the ISF, there is a possibility of reducing *by large factors* the $1/f^3$ phase-noise corner. The ISF is a function of the waveform, and hence potentially under the control of the designer, usually through adjustment of the rise- and fall-time symmetry. 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 MESFET's. 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. As alluded to earlier, the noise sources in many oscillators cannot be well modeled as stationary. A typical example is the nominally white drain or collector noise current in a MOSFET. Noise currents are a function of bias currents, and the latter 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 [14]

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

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 unitless function with a peak value of unity. Substituting this into (15) allows us to treat cyclostationary noise as a stationary noise source provided we define an effective ISF as follows:

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

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

Having identified the factors that influence oscillator noise, we are now in a position to articulate the requirements that must be satisfied to make a good oscillator. First, in common with the revelations of LTI models, both the signal power and resonator Q should be maximized, all other factors held constant. In addition, note that an active device is always necessary to compensate for tank loss, and that active devices always contribute noise. Note also that the ISF's tell us that there are sensitive and insensitive moments in an oscillation cycle. Of the infinitely many ways that an active element could return energy to the tank, this energy should be delivered *all at once*, where the ISF has its minimum value. In an ideal *LC* oscillator, therefore, the transistor would remain off almost all of the time, waking up periodically to deliver an impulse of current at the signal peak(s) of 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. Last, the best oscillators will possess the symmetry properties that lead to small Γ_{dc} for minimum upconversion of $1/f$ noise. In the following section, we consider several circuit examples of how to accomplish these ends in practice.

V. CIRCUIT EXAMPLES

A. LC Oscillators

Let us examine first a Colpitts oscillator and its relevant waveforms, now that we have developed these insights (see Figs. 9 and 10). Note that the drain current flows only during a short interval coincident with the most benign moments (the peaks of the tank voltage). 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.

Both LTI and 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 Clapp's modification to the Colpitts oscillator (Fig. 11). Differential implementations of oscillators with tapped resonators have recently made an appearance in the literature [5], [9], [10]. These types of oscillators become increasingly attractive as supply voltages scale downward, where conventional resonator connections lead to V_{DD}-constrained signal swings. Use of tapping allows signal energy to remain high even with low supply voltages.
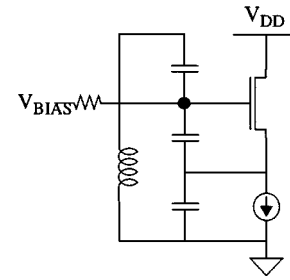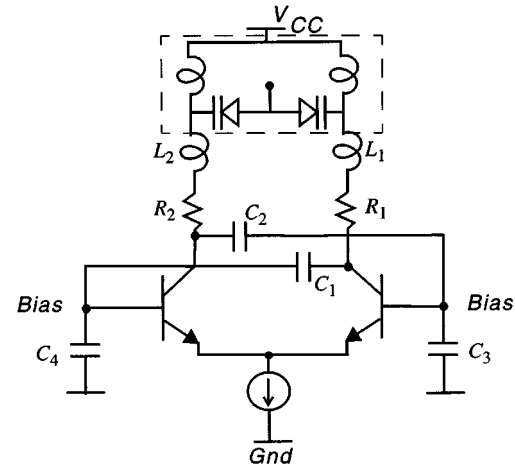
Fig. 11.   Clapp oscillator.

Fig. 12.   Simplified schematic of the VCO in [10].

Phase-noise predictions using the LTV model are frequently more accurate for bipolar oscillators due to availability of better device noise models. In [10], impulse response modeling was used to optimize the noise performance of a differential bipolar voltage-controlled oscillator (VCO) with an automatic amplitude control loop using an external resonator. A simplified schematic of this oscillator is shown in Fig. 12. A tapped resonator is used to increase the tank signal power P_{sig}. The optimum capacitive tapping ratio is calculated to be around 4.5 based on simulations that take the cyclostationarity of the noise into account. The effect of $1/f$ noise reduction is clearly seen, as a $1/f^3$ corner of 3 kHz is both predicted and measured, in comparison with a device $1/f$ noise corner of 200 kHz. The measured phase noise of -106 dBc/Hz at 100-kHz offset in the $1/f^2$ region is also in excellent agreement with the predicted value of -106.2 dBc/Hz. The automatic amplitude control loop allows for independent optimization of the steady-state and startup conditions in terms of phase noise.

As mentioned, a key insight of the LTV theory concerns the importance of symmetry. A configuration that exploits this knowledge is the symmetrical negative resistance oscillator shown in Fig. 13 [6]. This configuration is hardly 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 (Γ_{dc}) for each half-circuit, one may minimize the upconversion of $1/f$ noise. Through exploitation of symmetry in this manner, the $1/f^3$ corner can be dropped to exceptionally low values,

even though device $1/f$ noise corners may be high (as is typically the case for CMOS). Furthermore, the bridge-like arrangement of the transistor quad allows for greater signal swings, compounding the improvements in phase noise. As a result of all of these factors, a phase noise of -121 dBc/Hz at an offset of 600 kHz at 1.8 GHz has been obtained with low-Q on-chip spiral inductors, on 6 mW of power consumption in a 0.25-μm CMOS technology [6]. This result rivals what one may achieve with bipolar technologies. With a modest increase in power, the same oscillator"s phase noise becomes compliant with specifications for GSM1800.

B. Ring Oscillators

As an example of a circuit that does not well approximate ideal behavior, consider a ring oscillator. First, the "resonator" Q is poor since the energy stored in the node capacitances is reset (discharged) every cycle. Hence, if the resonator of a Colpitts oscillator may be likened to a fine crystal wine glass, the resonator of a ring oscillator is mud. Next, energy is restored to the resonator during the edges (the worst possible times), rather than at 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 noncritical applications, or inside wide-band phase-locked loops that clean up the spectrum.

However, there are certain aspects of ring oscillators that can be exploited to achieve better phase-noise performance in a mixed-mode integrated circuit. Noise sources on different nodes of an oscillator may be strongly correlated due to various reasons. Two examples of sources with strong correlation are substrate and supply noise, arising from current switching in other parts of the chip. The fluctuations on the supply and substrate will induce a similar perturbation on different stages of the ring oscillator.

To understand the effect of this correlation, consider the special case of having identical noise sources on all the nodes of the ring oscillator, as shown in Fig. 14. If all the inverters in the oscillator are the same, the ISF for different nodes will differ only in phase by multiples of $2\pi/N$, as shown in Fig. 15. Therefore, the total phase due to all the sources is given by (15) through superposition [13]

$$\phi(t) = \frac{1}{q_{max}} \int_{-\infty}^{t} i(\tau) \left[\sum_{n=0}^{N-1} \Gamma\left(\omega_0\tau + \frac{2\pi n}{N}\right) \right] d\tau. \quad (29)$$

Expanding the term in brackets in a Fourier series, it can be observed that it is zero except at dc and multiples of $N\omega_0$, i.e.,

$$\phi(t) = \frac{N}{q_{max}} \sum_{n=0}^{\infty} c_{(nN)} \int_{-\infty}^{t} i(\tau) \cos(nN\omega_0\tau) d\tau \quad (30)$$

which means that for fully correlated sources, only noise in the vicinity of integer multiples of $N\omega_0$ affects the phase. Therefore, every effort should be made to maximize the correlations of noise arising from substrate and supply perturbations. This can be done by making the inverter stages and the noise sources on each node as similar to each other as possible by proper layout
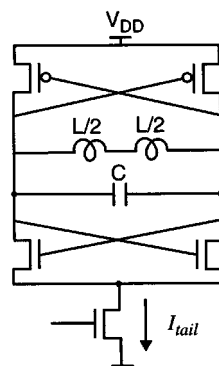

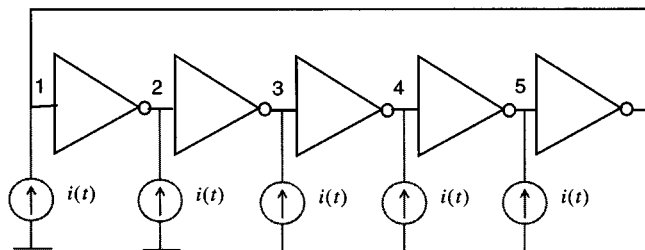Fig. 13. Simple symmetrical negative resistance oscillator.

Fig. 14. Five-stage ring oscillator with identical noise sources on all nodes.
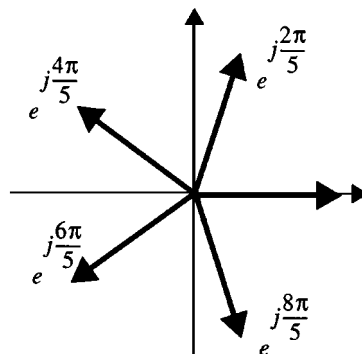

Fig. 15. Phasors for noise contributions from each source.

and circuit design. For example, the layout should be kept symmetrical, and the inverter stages should be laid out close to each other so that substrate noise appears as a common-mode source. This latter consideration is particularly important in the case of a lightly doped substrate, since such a substrate may not act as a single node [11]. It is also important that the orientation of all the stages be kept identical. The interconnecting wires between the stages must be identical in length and shape, and a common supply line should feed all the inverter stages. Furthermore, the loading on all stages should be kept identical, perhaps, for example, by using dummy buffer stages as necessary. Use of the largest number of stages consistent with oscillation at the desired frequency will also be helpful because, as a practical matter, fewer cn coefficients will then affect the phase noise. Last, as the low-frequency portion of the substrate and supply noise then dominates, one should exploit symmetry to minimize Γ_{dc}.

Another common conundrum concerns the preferred topology for MOS ring oscillators, i.e., whether a single-ended or differential topology results in better jitter and phase-noise

performance for a given center frequency f_0 and total power dissipation P. Using the noise equations for transistors and an approximate expression for the ISF, one may derive expressions for the phase noise of MOS differential and single-ended oscillators [13]. Based on these expressions, the phase noise of a single-ended (inverter chain) ring oscillator is found to be independent of the number of stages for a given power dissipation and frequency of operation. However, for a differential ring oscillator, the phase noise (jitter) grows with the number of stages. Therefore, even a properly designed differential CMOS ring oscillator underperforms its single-ended counterpart, with a disparity that increases with the number of stages. The difference in the behavior of these two types of oscillators with respect to the number of stages can be traced to the way they dissipate power. The dc current drawn from the supply is independent of the number and slope of the transitions in differential ring oscillators. In contrast, inverter-chain ring oscillators dissipate power mainly on a per transition basis and therefore have better phase noise for a given power dissipation. However, a differential topology may still be preferred in IC's with a large amount of digital circuitry because of the lower sensitivity to substrate and supply noise, as well as lower noise injection into other circuits on the same chip. The decision of which architecture to use should be based on both of these considerations.

Yet another commonly debated question concerns the optimum number of inverter stages in a ring oscillator to achieve the best jitter and phase noise for a given f_0 and P. For single-ended CMOS ring oscillators, the phase noise and jitter in the $1/f^2$ region are not strong functions of the number of stages [13]. However, if the symmetry criteria are not well satisfied, and/or the process has large $1/f$ noise, a larger N will reduce the jitter. In general, the choice of the number of stages must be made on the basis of several design criteria, such as $1/f$ noise effect, the desired maximum frequency of oscillation, and the influence of external noise sources, such as supply and substrate noise, that may not scale with N.

The jitter and phase-noise behavior is different for differential ring oscillators. Jitter and phase noise increase with an increasing number of stages. Hence, if the $1/f$ noise corner is not large, and/or proper symmetry measures have been taken, the minimum number of stages (three or four) should be used to give the best performance. This recommendation holds even if the power dissipation is not a primary issue. It is not fair to argue that burning more power in a larger number of stages allows the achievement of better phase noise, since dissipating the same total power in a smaller number of stages with larger devices results in better jitter and phase noise, as long as it is possible to maximize the total charge swing.

VI. 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 ISF minimum, 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. The theory also shows that the inferior broad-band noise performance of ring oscillators may be offset by their potentially superior ability to reject common-mode substrate and supply noise.

APPENDIX A
NOTES ON SIMULATION

Exact analytical derivations of the ISF are usually not obtainable for any but the simplest oscillators. Various approximate methods are outlined in [3] and [4], but the most generally accurate method is direct evaluation of the time-varying impulse response. In this direct method, an impulsive excitation perturbs the oscillator, and the steady-state phase perturbation measured. The timing of the impulse with respect to the unperturbed oscillator's zero crossing is then incremented and the simulation repeated until the impulse has been "walked" through an entire cycle.

The impulse must have a small enough value to ensure that the assumption of linearity holds. Just as an amplifier"s step response cannot be evaluated properly with steps of arbitrary size, one must judiciously select the area of the impulse rather than blindly employing some fixed value (e.g., one coulomb). If one is unsure if the impulse chosen is sized properly, linearity may always be tested explicitly by scaling the size of impulse by some amount and verifying that the response scales by the same factor.

Last, some confusion persists about whether the LTV theory properly accommodates the phenomenon of amplitude-to-phase conversion that some oscillators exhibit. As long as linearity holds, the LTV theory provides the correct answer, provided that an exact ISF has been obtained. This is due to the fact that changes in the phase of oscillator due to an amplitude change appear in the impulse response of the oscillator. As noted in the preceding paragraphs, the direct impulse response method is the most reliable one, as it makes no assumptions other than linearity. This reliability is in contrast with the approximate analytical approaches offered in [3, Appendix].

APPENDIX B
AMPLITUDE RESPONSE

While the close-in sidebands are dominated by phase noise, the far-out sidebands are greatly affected by amplitude noise. Unlike the induced excess phase, the excess amplitude $A(t)$, due to a current impulse, decays with time. This decay is the direct result of the amplitude restoring mechanisms always present in practical oscillators. The excess amplitude may decay very slowly (e.g., in a harmonic oscillator with a high-quality resonant circuit) or very quickly (e.g., a ring oscillator). Some circuits may even demonstrate an underdamped second order amplitude response.
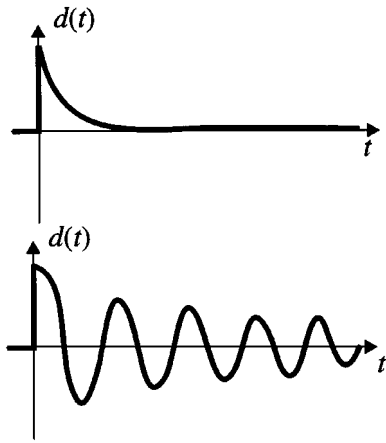
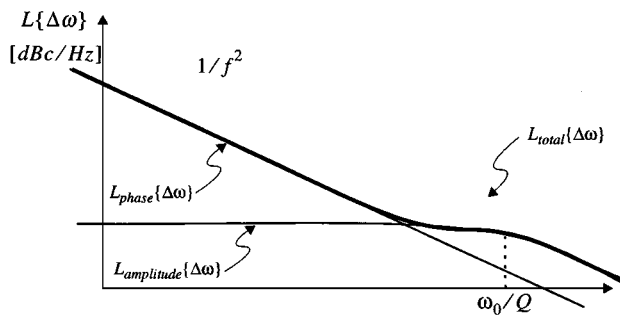Fig. 16.   Overdamped and underdamped amplitude responses.

Fig. 17.   Phase, amplitude, and total output sideband power for the overdamped exponentially decaying amplitude response.

In the context of the ideal *LC* oscillator of Fig. 1, a current impulse with an area Δq will induce an instantaneous change in the capacitor voltage, which, in turn, will result in a change in the oscillator amplitude that depends on the instant of injection, as shown in Fig. 4. The amplitude change is proportional to the instantaneous normalized voltage change[1] $\Delta V/V_{\max}$ for small injected charge, i.e.,

$$\Delta A = \Lambda(\omega_0 t)\frac{\Delta V}{V_{\max}} = \Lambda(\omega_0 t)\frac{\Delta q}{q_{\max}} \qquad \Delta q \ll q_{\text{swing}} \quad (31)$$

where $\Lambda(\omega_0 t)$ is a periodic function that determines the sensitivity of each point on the waveform to an impulse and is called the amplitude impulse sensitivity function. It is the amplitude counterpart of the phase impulse sensitivity function $\Gamma(\omega_0 t)$. From a development similar to that of Section IV, the amplitude impulse response can be written as

$$h_A(t, \tau) = \frac{\Lambda(\omega_0 t)}{q_{\max}} d(t - \tau) \tag{32}$$

where $d(t - \tau)$ is a function that defines how the excess amplitude decays. Fig. 16 shows two hypothetical examples of $d(t)$ for a low Q oscillator with overdamped response and a high Q oscillator with underdamped amplitude response. For most oscillators, the amplitude limiting system can be approximated as

[1]Note that the amplitude change is only equal to the normalized voltage change $\Delta V/V$ if the impulse is injected at the peak of the sinusoidal tank voltage.

first or second order [12]. The function $d(t - \tau)$ typically will thus be either a dying exponential or a damped sinusoid.

For a first-order system

$$d(t - \tau) = e^{-\omega_0(t-\tau)/Q} \cdot u(t - \tau). \tag{33}$$

Therefore the excess amplitude response to an arbitrary input current $i(t)$ is given by the superposition integral

$$A(t) = \int_{-\infty}^{t} \frac{i(\tau)}{q_{\max}} \Lambda(\omega_0\tau) e^{-\omega_0(t-\tau)/Q} \, d\tau. \tag{34}$$

If $i(t)$ is a white noise source with power spectral density $\overline{i_n^2}/\Delta f$, the output power spectrum of the amplitude noise $A(t)$ can be shown to be

$$L_{\text{amplitude}}\{\Delta\omega\} = \frac{\Lambda_{\text{rms}}^2}{q_{\max}^2} \cdot \frac{\overline{i_n^2}/\Delta f}{2 \cdot \left(\frac{\omega_0^2}{Q^2} + \Delta\omega^2\right)} \tag{35}$$

where Λ_{rms} is the rms value of $\Lambda(\omega_0 t)$. If L_{total} is measured, the sum of both $L_{\text{amplitude}}$ and L_{phase} will be observed, and hence there will be a pedestal in the phase-noise spectrum at ω_0/Q, as shown in Fig. 17. Also note that the significance of the amplitude response depends greatly on Λ_{rms}, which, in turn, depends on the topology.

ACKNOWLEDGMENT

REFERENCES

[1] T. H. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*.   Cambridge, U.K.: Cambridge Univ. Press, 1998.

[2] D. B. Leeson, "A simple model of feedback oscillator noise spectrum," in *Proc. IEEE*, vol. 54, Feb. 1966, pp. 329–330.

[3] A. Hajimiri and T. Lee, "A general theory of phase noise in electrical oscillators," *IEEE J. Solid-State Circuits*, vol. 33, pp. 179–194, Feb. 1998.

[4] ——, *The Design of Low-Noise Oscillators*.   Norwell, MA: Kluwer, 1999.

[5] J. Craninckx and M. Steyaert, "A 1.8 GHz CMOS low-phase-noise voltage-controlled oscillator with prescaler," *IEEE J. Solid-State Circuits*, vol. 30, pp. 1474–1482, Dec. 1995.

[6] A. Hajimiri and T. Lee , "Design issues in CMOS differential *LC* oscillators," *IEEE J. Solid-State Circuits*, vol. 34, pp. 716–724, May 1999.

[7] B. Razavi , "A study of phase noise in CMOS oscillators," *IEEE J. Solid-State Circuits*, vol. 31, pp. 331–343, Mar. 1996.

[8] F. X. Kaertner , "Determination of the correlation spectrum of oscillators with low noise," *IEEE Trans. Microwave Theory Tech.*, vol. 37, pp. 90–101, Jan. 1989.

[9] T. I. Ahrens and T. H. Lee, "A 1.4GHz, 3mW CMOS *LC* low phase noise VCO using tapped bond wire inductance," in *Proc. ISLPED*, Aug. 1998, pp. 16–19.

[10] M. A. Margarit, J. I. Tham, R. G. Meyer, and M. J. Deen, "A low-noise, low-power VCO with automatic amplitude control for wireless applications," *IEEE J. Solid-State Circuits*, vol. 34, pp. 761–771, June 1999.

[11] T. Blalack, J. Lau, F. J. R. Clement, and B. A. Wooley , "Experimental results and modeling of noise coupling in a lightly doped substrate," in *Tech. Dig. IEDM*, Dec. 1996, pp. 623–626.

[12] K. K. Clarke and D. T. Hess, *Communication Circuits: Analysis and Design*.   Reading, MA: Addison-Wesley, 1971.

[13] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, pp. 790–804, June 1999.
[14] W. A. Gardner, *Introduction to Random Processes*. New York: McGraw-Hill, 1990.
[15] W. A. Edson, "Noise in oscillators," *Proc. IRE*, pp. 1454–1466, Aug. 1960.

Thomas H. Lee (S'87–M'91) received the S.B., S.M., and Sc.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1983, 1985, and 1990, respectively.

He joined Analog Devices in 1990, where he was primarily engaged in the design of high-speed clock recovery devices. In 1992, he joined Rambus, Inc., Mountain View, CA, where he developed high-speed analog circuitry for 500 Mbyte/s CMOS DRAM's. He has also contributed to the development of PLL's in the StrongARM, Alpha, and K6/K7 microprocessors. Since 1994, he has been an Assistant Professor of Electrical Engineering at Stanford University, where his research focus has been on gigahertz-speed wireline and wireless integrated circuits built in conventional silicon technologies, particularly CMOS. He has received 12 U.S. patents. He is the author of *The Design of CMOS Radio-Frequency Integrated Circuits* (Cambridge, U.K.: Cambridge Univ. Press, 1998) and is a coauthor of two additional books on RF circuit design. He is also a Cofounder of Matrix Semiconductor.

He has twice received the "Best Paper" award at the International Solid-State Circuits Conference (ISSCC), was coauthor of a "Best Student Paper" at ISSCC, and was awarded a Packard Foundation Fellowship in 1997. He is a Distinguished Lecturer of the IEEE Solid-State Circuits Society and was recently named a Distinguished Microwave Lecturer.

Ali Hajimiri (S'95–M'99) received the B.S. degree in electronic engineering from the Sharif University of Technology, Tehran, Iran, in 1994 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1996 and 1998, respectively.

He was a Design Engineer with Philips, where he worked on a BiCMOS chipset for GSM cellular units from 1993 to 1994. During the summer of 1995, he was with Sun Microsystems, where he worked on the UltraSPARC microprocessor's cache RAM design methodology. During the summer of 1997, he was with Lucent Technologies (Bell Labs), where he investigated low-phase-noise integrated oscillators. In 1998, he joined the Faculty of the California Institute of Technology, Pasadena, as an Assistant Professor of electrical engineering, where his research interests are high-speed and RF integrated circuits. He is the author of *The Design of Low Noise Oscillators* (Boston, MA: Kluwer Academic, 1999) and has received several U.S. and European patents.

Dr. Hajimiri was the Bronze Medal winner of the First International Physics Olympiad, Groningen, the Netherlands. He was a corecipient of the International Solid-State Circuits Conference 1998 Jack Kilby Outstanding Paper Award.