This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

Understanding Phase Error and Jitter: Definitions, Implications, Simulations, and Measurement

Ian Galton, *Fellow, IEEE*, and Colin Weltin-Wu, *Member, IEEE*

(Invited Paper)

***Abstract*— Precision oscillators are ubiquitous in modern electronic systems, and their accuracy often limits the performance of such systems. Hence, a deep understanding of how oscillator performance is quantified, simulated, and measured, and how it affects the system performance is essential for designers. Unfortunately, the necessary information is spread thinly across the published literature and textbooks with widely varying notations and some critical disconnects. This paper addresses this problem by presenting a comprehensive one-stop explanation of how oscillator error is quantified, simulated, and measured in practice, and the effects of oscillator error in typical oscillator applications.**

***Index Terms*— Oscillator phase error, phase noise, jitter, frequency stability, Allan variance, frequency synthesizer, crystal, phase-locked loop (PLL).**

I. Introduction

ELECTRONIC oscillators are in nearly all electronic devices. They make it possible to modulate and demodulate wireless signals, clock both digital and sampled-data analog circuits, and measure time intervals. Moreover, they are astonishingly precise. For example, present-day mobile telephone transceivers incorporate several highly-tunable multi-GHz oscillators whose output signals have integrated jitter of less than a few hundred femtoseconds. Nevertheless, oscillator precision is often a limiting factor in high-performance applications such as modern wireless and wireline communications, radar, and high-speed digital circuits. Accordingly, it is critical for engineers to understand how oscillator performance is quantified, simulated, and measured, and how oscillator error affects application performance.

Unfortunately, acquiring this knowledge is not easy. The information is spread thinly across the published literature, and there are fundamental disconnects between much of the published information and actual practice. One disconnect is that most textbooks and papers define oscillator error in terms of sinusoidal oscillator signals, but most practical circuits use squared-up versions of such oscillator signals that approximate square waves. A related disconnect is that most communication textbooks define a mixer as performing multiplication by a sinusoidal oscillator signal whereas most practical mixer circuits perform multiplication by squared-up oscillator signals. Fundamental questions typically are left unanswered such as: How does the squaring-up process change the oscillator error? How does multiplying by a squared-up oscillator signal instead of a sinusoidal signal change a mixer's behavior and response to oscillator error?

Another obstacle to learning the material is that there are three distinct oscillator error metrics in common use: phase error, jitter, and frequency stability. Each metric offers advantages in certain applications, so it is important to understand how they relate to each other, yet with the exception of [1], the authors are not aware of prior publications that provide this information directly.

The goal of this tutorial is to comprehensively and systematically present this information. There are many publications that describe and model the circuit-level mechanisms of oscillator error, e.g., [2]–[9], so this material is not covered here. Instead, the paper describes how to evaluate, simulate, and measure oscillator error, and the system-level effects of oscillator error in typical circuit applications.

II. Phase Error

A. Sinusoidal Oscillator Signal Model

Oscillator accuracy is critical in so many electronics applications that the recommended quantities with which to quantify it have been specified in an IEEE standard [10]. As with most other papers and textbooks that touch on the subject, the IEEE standard starts from the premise that an ideal oscillator is purely sinusoidal. It expresses the instantaneous output of a non-ideal oscillator as

$$v(t) = [V_0 + \varepsilon(t)] \sin(\omega_0 t + \phi(t)), \tag{1}$$

where V_0 is the nominal peak amplitude, $\varepsilon(t)$ is the *amplitude error*, ω_0 is the nominal frequency, and $\phi(t)$ is the *phase error*. The amplitude error represents the oscillator's deviation from the nominal amplitude, and the phase error represents the oscillator's deviation from the ideal phase, $\omega_0 t$. It is assumed that $|\varepsilon(t)| < V_0$ for all t and $\omega_0 t + \phi(t)$ monotonically increases with t, which are reasonable assumptions given that the purpose of the model is to characterize precision oscillators.

The oscillator signal's zero-crossing times, i.e., the values of t for which $v(t) = 0$, are of particular importance. Throughout this paper, the zero-crossing times are denoted as t_n and ordered such that $t_n > t_{n-1}$ for all integers n, so t_n is an increasing sequence that comprises all values of t at which

Manuscript received March 6, 2018; revised May 16, 2018 and June 14, 2018; accepted July 10, 2018. This work was supported by the National Science Foundation under Grant 1617545. This paper was recommended by Associate Editor E. Blokhina. *(Corresponding author: Ian Galton.)*

I. Galton is with the University of California at San Diego, La Jolla, CA 92093 USA (e-mail: galton@ucsd.edu).

C. Weltin-Wu is with Analog Devices, Inc., San Jose, CA 95134 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCSI.2018.2856247

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

$v(t) = 0$. It follows from (1) and the $|\varepsilon(t)| < V_0$ assumption that, for each integer n,

$$\phi(t_n) = \pi n - \omega_0 t_n. \tag{2}$$

Samples of the phase error at the zero-crossing times, i.e., $\phi(t_n)$ for every integer n, can be measured directly from $v(t)$ by substituting the nth time at which $v(t)$ crosses zero into (2). In this sense, the $\phi(t_n)$ values are uniquely determined by $v(t)$.

However, $\phi(t)$ is not uniquely determined by $v(t)$ for values of t that are not zero-crossing times, because at such times there are an infinite number of $\varepsilon(t)$ and $\phi(t)$ functions that yield the same value of $v(t)$. This can be verified by observing that whenever $t \neq t_n$, the right side of (1) can be written as $[V_0 + \varepsilon_1(t)]\,\sin(\omega_0 t + \phi_1(t))$, where

$$\varepsilon_1(t) = [V_0 + \varepsilon(t)]\,\frac{\sin(\omega_0 t + \phi(t))}{\sin(\omega_0 t + \phi_1(t))} - V_0 \tag{3}$$

and $\phi_1(t)$ is any function for which $\omega_0 t + \phi_1(t)$ is not an integer multiple of π.

Furthermore, $\varepsilon(t)$ is not uniquely determined by $v(t)$ for *any* value of t. This follows directly from the argument above for each t that is not a zero-crossing time. For each t that is a zero-crossing time, it follows because the sine function in (1) evaluates to zero, so $\varepsilon(t)$ can take on any value.

These ambiguities generally make it impossible to separate the effects of $\phi(t)$ from those of $\varepsilon(t)$ on the performance of a circuit driven by $v(t)$. Yet it is often the case in practice that the effects of $\varepsilon(t)$ on circuits driven by oscillators are negligible compared to those of $\phi(t)$. One reason is that the mean squared value of $\varepsilon(t)$ is often very small for practical oscillators. Another reason is that typical circuits driven by oscillators only change state when $v(t)$ is relatively close to zero where the magnitude of the sinusoid in (1) is small, so it attenuates the effect of $\varepsilon(t)$. In such cases, $\phi(t)$ represents the oscillator's only significant non-ideal behavior, so the above-mentioned ambiguities are avoided.

B. Squared-Up Oscillator Signal Model

Oscillator signals that switch as abruptly as possible between $\pm V_0$ at each zero-crossing are often used instead of the sinusoidal oscillator signal given by (1). Such oscillator signals are sometimes generated by amplifying and clipping sinusoidal oscillator signals, so they are often called *squared-up* oscillator signals. Squared-up oscillator signals are widely used in practice, because, as described shortly, they offer practical benefits when used to drive circuits that are only sensitive to the oscillator signals near their zero-crossings.

In analogy with (1), a squared-up oscillator signal can be modelled as

$$v(t) = [V_0 + \varepsilon(t)]\, r\,(\omega_0 t + \phi(t))\,, \tag{4}$$

where $r(\theta)$ is a 2π-periodic function which is as close as possible to a unit square wave given by

$$r_{ideal}(\theta) = \begin{cases} 1, & \text{if } \sin(\theta) > 0, \\ 0, & \text{if } \sin(\theta) = 0, \\ -1, & \text{if } \sin(\theta) < 0. \end{cases} \tag{5}$$

Thus, $\phi(t_n)$ is still given by (2). In cases where the squared-up oscillator signal is obtained by amplifying and clipping a sinusoidal oscillator signal, the $\phi(t)$ and $\varepsilon(t)$ associated with the squared-up oscillator signal are generally different from those associated with the sinusoidal oscillator signal. However, to the extent that the amplifying and clipping circuitry has negligible noise and input offset voltage, the zero-crossing times, t_n, and the phase error sampled at the zero-crossing times, $\phi(t_n)$, are not altered by the squaring up process.

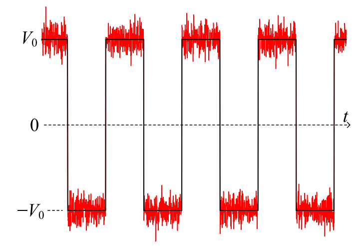

Fig. 1. Superimposed squared-up oscillator signals with and without amplitude noise.

It is convenient to rewrite the squared-up oscillator signal model as

$$v(t) = V_0 r\,(\omega_0 t + \phi(t)) + e(t), \tag{6}$$

which is equivalent to (4) when $e(t) = \varepsilon(t) r(\omega_0 t + \phi(t))$ but emphasizes that $\varepsilon(t)$ can be viewed as a type of additive error. Practical oscillator circuits typically introduce other types of additive error as well, and error introduced by the circuits they drive often can be modeled as input-referred error added to the oscillator signal. Therefore, $e(t)$ is more realistically modelled as $e(t) = \varepsilon(t) r(\omega_0 t + \phi(t)) + e_a(t)$, where $e_a(t)$ is all additive error other than $\varepsilon(t) r(\omega_0 t + \phi(t))$.

To the extent that $r(\theta)$ approximates (5), $v(t)$ is only near zero at its zero-crossing times, a consequence of which is that $e(t)$ only has a significant effect on the oscillator signal when the magnitude of $v(t)$ is large. This phenomenon is illustrated in Fig. 1, wherein the oscillator signal was generated by (6) with an $r(\theta)$ that well-approximates (5). Two versions of the waveform are superimposed: an ideal version in which $e(t)$ is zero, and a non-ideal version in which $e(t)$ is a noise waveform. The two waveforms are nearly identical except when their magnitudes are much greater than zero. If such an oscillator signal drives a circuit that is insensitive to the oscillator signal when its magnitude is far from zero, $e(t)$ has little effect on the circuit's behavior.

The common practice of using squared-up oscillator signals to drive such circuits renders even relatively strong amplitude and additive error negligible in terms of circuit performance. Nevertheless, as described in Sections VI and VII, care must be taken to prevent $e(t)$ from corrupting phase error simulations and laboratory measurements.

Unlike $e(t)$, phase error cannot be neglected in practical circuits. However, as explained shortly, it is not the phase error waveform itself, but the phase error sampled at times t_n, i.e., $\phi(t_n)$, that matters, and this sampling process leads to several practical issues because of aliasing.

C. Reconciliation of the Two Oscillator Signal Models

For the reasons described above, it is neither common, nor desirable in many cases, for oscillator signals to be sinusoidal in practical circuits. Yet the official IEEE definition and most communication system textbooks describe oscillator signals as sinusoidal. As shown in this subsection, these two viewpoints can be reconciled to an extent by observing that the sinusoidal

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

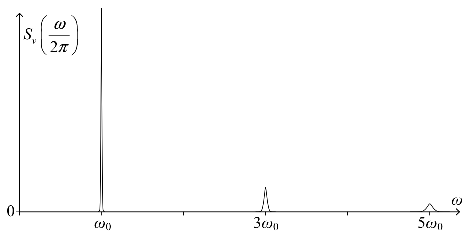

Fig. 2. Power spectrum of the first three terms of the squared-up oscillator signal given by (8).

and squared-up oscillator signal models often have nearly equivalent spectra for frequencies from 0 to well above ω_0.

Given that $r(\theta)$ in (6) is a 2π-periodic real function, it can be expressed as a Fourier series. Substituting the Fourier series representation of $r(\theta)$ into (6) yields

$$v(t) = V_0 \sum_{n=0}^{\infty} \rho_n \sin\left(n\omega_0 t + n\phi(t) + \theta_n\right) + e(t), \tag{7}$$

where ρ_n and θ_n are constants that depend on the Fourier series coefficients. For the case where $r(\theta) = r_{ideal}(\theta)$, (7) reduces to

$$v(t) = \frac{4}{\pi} V_0 \sum_{n=1,3,5,\ldots}^{\infty} \frac{1}{n} \sin\left(n\omega_0 t + n\phi(t)\right) + e(t). \tag{8}$$

Except for the $4/\pi$ constant scale factor, the $n = 1$ term of (8) is identical to the right side of (1) in the absence of amplitude error. The remaining terms are therefore the only significant differences between the sinusoidal and squared-up oscillator signal models in this case.

The nth term of (8) for $n > 1$ is a sinusoid with phase modulation $n\phi(t)$, so Carson's rule suggests that its bandwidth is approximately n times larger than that of the first term [11]. However, the power spectrum of the nth term is centered n times further away from zero than that of the first term. Furthermore, its power is $1/n^2$ times that of the first term. Consequently, it is reasonable to expect that the terms in (8) for $n > 1$ are negligible at frequencies from 0 to well above ω_0, as illustrated conceptually in Fig. 2.

This phenomenon is particularly relevant to frequency translation via mixers in communication systems. The classic communication system textbook definition of a mixer is a circuit that multiplies an input signal by a sinusoidal oscillator signal and filters the result. An example of such a mixer that translates the center frequency of a bandpass signal from a non-zero frequency, ω_1, to another non-zero frequency, $\omega_1 - \omega_0$, is shown in Fig. 3.

In most practical situations, the sinusoidal oscillator signal in the communication system textbook version of a mixer could be replaced by a squared-up oscillator signal without significantly changing the output of the mixer. For example, consider the mixer of Fig. 3 except with the sinusoidal oscillator signal replaced by a squared-up oscillator signal. It is straightforward to verify that each of the $n > 1$ terms in (8) contribute signal components that lie outside the passband of the filter provided the bandwidths of the mixer's input signal and bandpass filter are sufficiently narrow. In this case, only the $n = 1$ term in (8) significantly affects the output of the mixer, and this term is identical to the sinusoidal oscillator signal model in the absence of amplitude error except for a scale factor of $4/\pi$.†

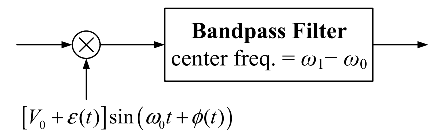

Fig. 3. Communication system textbook version of a mixer with the addition of oscillator phase and amplitude error.

This behavioral equivalence between sinusoidal and squared-up oscillator signals in mixers provides some rationale for the persistence of the sinusoidal oscillator signal model despite the widespread use of squared-up oscillator signals. Mixers are of great importance in communication systems so they are taught early at the undergraduate level, and it is easier to explain the basic operation of communication system textbook style mixers with sinusoidal oscillator signals than with squared-up oscillator signals. Mixers of one type or another have been in use for over a century [13], [14], and for the first several decades they were the primary application of oscillators. Moreover, early mixers did approximate multiplication by sinusoidal oscillator signals, and this remains so in very high-frequency (e.g., millimeter wave) applications wherein parasitic capacitances and other electronic device nonidealities make squared-up oscillator signals impractical.

Nevertheless, as described in more detail below, the rationale breaks down somewhat with typical present-day mixers. Furthermore, in other oscillator applications, such as sampling and clocking, the sinusoidal oscillator model is not particularly useful. These applications use the zero-crossing times of oscillator signals to initiate specific events, so there is no benefit to thinking of the oscillator signals as being sinusoidal in such cases.

D. Circuit Applications of Oscillators

This subsection describes the types of analog and mixed-signal circuits commonly driven by oscillators. A comprehensive description of this subject could easily fill a textbook, so it is far beyond the scope of this paper. Instead, the goal is to provide qualitative descriptions that explain the practical reasons which underlie the statements made above regarding squared-up oscillator signals and why the circuits driven by them usually are insensitive to $e(t)$.

1) Mixers: As described in Section II-C, communication system textbooks describe mixers as ideal multipliers followed by filters. In contrast, circuit design textbooks typically consider just the multiplier to be the mixer. Filtering is always performed following the multiplier, but it is not considered by most circuit designers to be part of the mixing operation.

Practical mixers do not implement true multiplication. Instead, as depicted in Fig. 4 a practical mixer usually can be modelled as a device that multiplies the input signal, $x(t)$, by a strongly nonlinearly distorted and limited version of the oscillator signal, i.e., $f(v(t))$. The nonlinearity, $f(v)$, is monotonic and is limited in the sense that $f(v) \cong 1$ when $v > V_{s+}$ and $f(v) \cong -1$ when $v < V_{s-}$, where V_{s+} and V_{s-}

†An exception occurs if the mixer input has unwanted signal components at $\omega_1 - k\omega_0$, for any $k = 3, 5, \ldots$, in which case the unwanted components corrupt the desired component in the mixer's output. A harmonic rejection mixer that consists of multiple conventional mixers with squared-up oscillator signals of different phases can be used to avoid this problem if necessary [12].

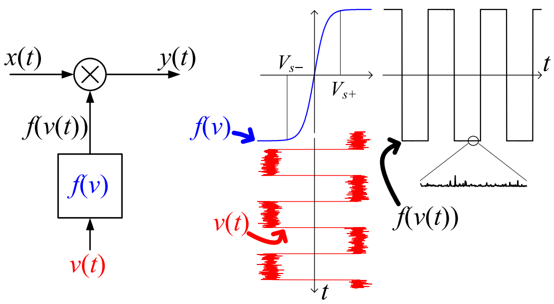

Fig. 4. Model of a practical mixer with an example that demonstrates how $f(v)$ suppresses the squared-up oscillator signal's $e(t)$ error.

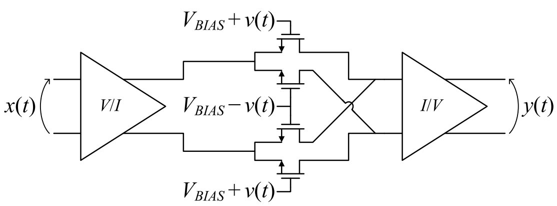

Fig. 5. Simplified diagram of a passive mixer circuit.

denote the approximate values of v where $f(v)$ saturates at 1 and -1, respectively. The limiting causes the mixer to be nearly insensitive to $v(t)$ when $v(t) < V_{s-}$ or $v(t) > V_{s+}$.

Suppose $v(t)$ is the squared-up oscillator signal given by (6) with an $r(\theta)$ that well-approximates $r_{ideal}(\theta)$, and V_0 is large enough that $V_0+e(t) > V_{s+}$ and $-V_0+e(t) < V_{s-}$. In this case $v(t)$ is nearly independent of $e(t)$ when $V_{s-} < v(t) < V_{s+}$ for the reasons described in Section II-B. While $v(t)$ does depend on $e(t)$ when $v(t) < V_{s-}$ or $v(t) > V_{s+}$, these are the regions over which the mixer is insensitive to $v(t)$. Consequently, the mixer output is insensitive to $e(t)$ for all values of $v(t)$. Given that the oscillator signal is only in the range $V_{s-} < v(t) < V_{s+}$ very near the oscillator signal's zero-crossing times, the mixer output is also insensitive to the nonlinearity imposed by $f(v)$. It follows that under these circumstances the behavior of the mixer is nearly identical to that of an ideal multiplier with the important and beneficial exception that it practically ignores $e(t)$. Specifically, the mixer's output under these circumstances well approximates

$$y(t) = \frac{4}{\pi} \sin\left(\omega_0 t + \phi(t)\right) x(t). \tag{9}$$

A practical example of such a mixer is shown in Fig. 5. The figure shows the general structure of a MOS transistor based mixer of the type currently used in mobile telephone receivers [15], [16]. The mixer consists of a differential voltage to current (V/I) converter, four MOS transistors controlled by the oscillator signal, and a differential current-to-voltage (I/V) converter.

Each of the four transistor gates is driven by $V_{BIAS}+v(t)$ or $V_{BIAS}-v(t)$, where V_{BIAS} is a constant bias voltage and $v(t)$ is the squared-up oscillator signal. The circuit is designed such that the four transistors approximate switches which connect the differential outputs of the V/I converter to the differential inputs of the I/V converter directly when $v(t)$ is sufficiently greater than zero and with swapped polarity when $v(t)$ is sufficiently less than zero.

Although the on-resistances of the transistors are never zero, the low differential input impedance of the I/V converter combined with the much higher differential output impedance of the V/I converter causes the mixer's behavior to be fairly insensitive to non-zero transistor on-resistance. In this context, V_{s-} and V_{s+} are voltage levels such that $v(t) < V_{s-}$ and $v(t) > V_{s+}$ are the ranges of $v(t)$ over which the on-resistance of the transistors is sufficiently low to negligibly affect mixer performance. By design, the value of V_{BIAS} is chosen to cause $V_{s-} \cong -V_{s+}$ and V_0 is chosen large enough to ensure $V_0+e(t) > V_{s+}$ and $-V_0+e(t) < V_{s-}$ for all t, so the mixer is insensitive to $e(t)$ for the reasons described above.

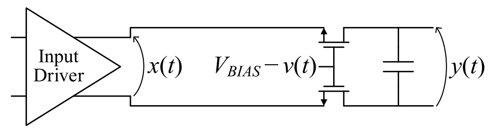

Fig. 6. Simplified diagram of an ADC's sampling circuit.

2) Oscillators as Clocks in Non-Mixer Applications: In most other applications, the oscillator signal is used as a clock to mark a set of times, τ_n for $n = \ldots, -1, 0, 1, 2, \ldots$, that are as close to being uniformly spaced as possible. The zero-crossing times of a squared-up oscillator are ideal for this purpose. Typically, only the positive-going zero-crossings or only the negative-going zero-crossings are used to mark τ_n in any given application, so the oscillator signal's duty cycle can deviate from 50% without causing timing errors. In the following, without loss of generality τ_n is taken to be the nth positive-going zero-crossing time of the oscillator signal.

Ideally, $\tau_n = nT_0$, where $T_0 = 2\pi/\omega_0$ is the oscillator signal's nominal period, but oscillator phase error causes deviations from this ideal. It follows from (2) that $\tau_n = t_{2n} = [2\pi n - \phi(t_{2n})]/\omega_0$. This can be rewritten as

$$\tau_n = nT_0 + \Delta\tau_n, \quad \text{where } \Delta\tau_n = -\frac{\phi(\tau_n)}{\omega_0}. \tag{10}$$

Thus, $\Delta\tau_n$ is the deviation of τ_n from its ideal value of nT_0, and it is proportional to the oscillator signal's phase error sampled at τ_n. It is often called *absolute jitter* or *aperture jitter* and it is described further in Section V.

3) Analog-to-Digital Converters: A typical ADC uses a squared-up oscillator signal, $v(t)$, as a clock to mark the times, τ_n, at which to sample its input signal, $x(t)$. It converts the resulting sequence of sampled values, $x(\tau_n)$, to a digital output sequence. The sampling circuitry is designed to be insensitive to the clock signal, $v(t)$, except near its positive-going zero-crossings, so ADCs, like mixers, are insensitive to $e(t)$.

There are several types of ADC sampling circuits, but the circuit shown in Fig. 6 is representative of the core operation performed by most of them. The gates of the MOS transistors are driven by $V_{BIAS} - v(t)$, where V_{BIAS} is a constant bias voltage and $v(t)$ is the squared-up oscillator signal given by (6) with an $r(\theta)$ that well-approximates $r_{ideal}(\theta)$. When $v(t)$ is close to $-V_0$ the transistors are turned on, and the voltage across the capacitor tracks $x(t)$. At the positive-going zero-crossing of the oscillator signal, $v(t)$ rapidly transitions toward V_0 which abruptly turns off the transistors, thereby freezing the voltage stored on the capacitor until the subsequent negative-going zero-crossing time.

The circuit is designed such that $V_{BIAS} - V_0 - e(t) < V_{off}$ and $V_{BIAS} + V_0 - e(t) > V_{on}$ for all t, where V_{off} and V_{on} are voltage levels for which the transistors have high-enough off-resistance when $V_{BIAS} - v(t) < V_{off}$ and low-enough on-resistance when $V_{BIAS} - v(t) > V_{on}$ so as not to degrade performance. Hence, the use of a squared-up oscillator signal for $v(t)$ ensures that $V_{BIAS} - v(t)$ is only between V_{off} and V_{on} very near the zero-crossing times of $v(t)$. One consequence is that $e(t)$ can be neglected much like in the case of the mixer described above. Another consequence is that each transistor's nonlinear transition between its on and off states has little effect on the sampling circuit's performance. Hence, the frozen voltage stored on the capacitor when $v(t)$ is close to V_0 well-approximates $x(\tau_n)$.

4) Other Clocking Applications: Oscillators are also used to clock a wide range of other analog, mixed-signal, and digital circuits including DACs, frequency dividers, phase-frequency detectors (PFDs), time-to-digital converters (TDCs), digital-to-time converters (DTCs), timers, and synchronous digital logic. In most cases, the components within these circuits that are clocked by the oscillator signals are edge-triggered latches or flip-flops.

The clock input of each such edge-triggered latch or flip-flop is driven by a clock signal of the form $V_{BIAS} + v(t)$, where, once again, V_{BIAS} is a constant bias voltage and $v(t)$ is a squared-up oscillator signal. By design, the edge-triggered circuit only changes its state when $v(t)$ is in a narrow range of values about its midpoint, and V_0 is large enough that that $-V_0 + e(t)$ and $V_0 + e(t)$ are outside this range for all t. Therefore, by exactly the same reasoning described above for mixers and sampling circuits, such circuits are insensitive to $e(t)$.

E. Phase Error Versus Phase Noise

An oscillator signal's phase error, $\phi(t)$, usually contains both a random component and a deterministic component. The random component is caused by device noise, such as thermal and flicker noise, introduced by the transistors and other components that make up the oscillator circuit. The deterministic component is the result of deterministic disturbances that are inadvertently generated within or parasitically coupled into the oscillator circuitry. For example, the deterministic component of $\phi(t)$ in a PLL-based oscillator inevitably contains spurious tones, also known as spurs, that are not harmonics of the oscillator frequency.

In this paper, $\phi(t)$ is called *phase error* and the random component of $\phi(t)$ is called *phase noise*. This choice was made to avoid confusion because the word *noise* is considered by many people to denote purely random phenomena. It is also consistent with much of the published literature, e.g., most papers that report the measured performance of PLL-based oscillators refer to phase noise and spurious tones as distinct types of phase error. Unfortunately, the literature is not consistent in this respect. For instance, in the IEEE standard $\phi(t)$ is called *phase fluctuations* and, somewhat confusingly, the term *phase noise* is reserved exclusively for a function $\mathcal{L}(f)$ (pronounced "script-ell of f") equal to half the one-sided time average power spectrum of $\phi(t)$ [10].

F. Time Average Power Spectra

Laboratory measurements and circuit simulations of power spectra inevitably estimate time average power spectra as opposed to statistical power spectra [17]. The two-sided time average power spectrum of a real-valued continuous-time signal, in this case $\phi(t)$, is defined as

$$S_{\phi\phi}(f) = \lim_{T\to\infty} \frac{1}{T} \left|FT\{\phi_T(t)\}\right|^2 \text{ for } -\infty < f < \infty \quad (11)$$

provided the limit exists, where f has units of Hz, and $FT\{\phi_T(t)\}$ is the Fourier transform of $\phi(t)$ restricted to the interval $0 \le t \le T$. Specifically,

$$FT\{\phi_T(t)\} = \int_{-\infty}^{\infty} \phi_T(t) e^{-j2\pi f t} dt \quad (12)$$

where

$$\phi_T(t) = \begin{cases} \phi(t), & \text{if } 0 \le t \le T, \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

As is well known and can be verified from the definition above, $S_{\phi\phi}(f)$ for any value of f can be interpreted as $1/\Delta f$ times the time average power of $\phi(t)$ in the frequency band between f and $f + \Delta f$ in the limit as $\Delta f \to 0$. Hence, each value of $S_{\phi\phi}(f)$ represents a *power density* per Hz. In this case, $\phi(t)$ has units of radians, so $S_{\phi\phi}(f)$ has units of radians^2 per Hz.

Replacing $|FT\{\phi_T(t)\}|^2$ in (11) by the product of the right side of (12) and the complex conjugate of the right side of (12) and rearranging the result yields

$$S_{\phi\phi}(f) = \int_{-\infty}^{\infty} R_{\phi\phi}(\tau) e^{-j2\pi\tau f} d\tau, \quad (14)$$

where $R_{\phi\phi}(\tau)$ is given by

$$R_{\phi\phi}(\tau) = \lim_{T\to\infty} \frac{1}{T} \int_0^T \phi(t)\phi(t+\tau) dt \quad (15)$$

and is called the *time average autocorrelation* of $\phi(t)$. Equation (14) implies that $S_{\phi\phi}(f)$ is the Fourier transform of the time average autocorrelation. Applying the inverse Fourier transform to (14) gives

$$R_{\phi\phi}(\tau) = \int_{-\infty}^{\infty} S_{\phi\phi}(f) e^{j2\pi\tau f} df. \quad (16)$$

It follows from (15) that the time average power of $\phi(t)$ is $R_{\phi\phi}(0)$, so (16) implies that the time average power of $\phi(t)$ is the integral of $S_{\phi}(f)$ over all frequencies.

The Fourier transform of any real-valued function is conjugate symmetric, and $|c| = |c^*|$ for any complex number c, so (11) implies that $S_{\phi\phi}(f) = S_{\phi\phi}(-f)$. Consequently, the one-sided time average power spectrum of $\phi(t)$, defined as

$$S_{\phi}(f) = 2S_{\phi\phi}(f) \text{ for } f \ge 0, \quad (17)$$

is often used. The factor of 2 ensures that the time average power of $\phi(t)$ is the integral of $S_{\phi}(f)$ over all positive frequencies.‡

As described above,

$$\mathcal{L}(f) = \frac{1}{2} S_{\phi}(f) \quad (18)$$

by definition. This definition is important because $\mathcal{L}(f)$ is what laboratory phase noise measurement instruments estimate. For reasons explained in Section III-A, when $\mathcal{L}(f)$ is expressed in decibels (dB), i.e., $10\log_{10}(\mathcal{L}(f))$, its units are *defined* to be dBc/Hz.

‡For the remainder of the paper, all *time average power spectra* are referred to as just *power spectra* for brevity.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

G. Spurious Tones

A signal or its power spectrum is said to contain a spur at a frequency f_m if the signal's power spectrum contains a Dirac delta function at f_m. The power of a phase error spur at frequency f_m is therefore given by

$$P_{spur} = \lim_{\varepsilon \to 0} \int_{f_m - |\varepsilon|}^{f_m + |\varepsilon|} S_{\phi\phi}(f) df \tag{19}$$

with units of radians^2.

Due to the symmetry of $S_{\phi\phi}(f)$ about $f = 0$, its spurs appear in pairs, e.g., a sinusoidal component $B\sin(2\pi f_m t)$ in $\phi(t)$ gives rise to two $B^2/4$-powered spurs, at frequencies f_m and $-f_m$. Nevertheless, the power of each spur is quantified separately by (19), a justification for which is explained in the next section.

As shown in Section IV, spurs are especially detrimental to the performance of wireless systems, so spur mitigation techniques in PLL-based oscillators are an active research topic [18]–[21].

III. Oscillator Signal Power Spectrum

A. Relationship to Phase Error Spectrum

As described in Section II, it is usually the case in practice that only the fundamental term of an oscillator signal (e.g., the $n = 1$ term in (8)) contributes to the oscillator signal's power spectrum in a wide bandwidth around $f = \omega_0/(2\pi)$, and in many applications $e(t)$ can be neglected. In these cases, the oscillator signal of interest has the form

$$v(t) = A\sin(\omega_0 t + \phi(t)) \tag{20}$$

where A is a constant amplitude.

Applying the angle sum trigonometric identity to (20) gives

$$v(t) = A\sin(\omega_0 t)\cos(\phi(t)) + A\cos(\omega_0 t)\sin(\phi(t)). \tag{21}$$

If $|\phi(t)| << 1$ for all t as is often the case for oscillators, applying the sine and cosine small angle approximations to (21) results in

$$v(t) \cong A\sin(\omega_0 t) + A\cos(\omega_0 t)\phi(t). \tag{22}$$

The first and second terms on the right side of (22) are the ideal oscillator signal and an additive noise term caused by oscillator phase error, respectively.

The power spectrum of $v(t)$ and its properties are given by (11)-(17) except with ϕ replaced by v. Given that the ideal oscillator signal term in (22) has zero bandwidth, the one-sided power spectrum of $v(t)$ for $f \neq \omega_0/(2\pi)$ only depends on the noise term. It follows from well-known Fourier transform properties and the power spectrum definition that the one-sided power spectrum of $v(t)$ can be written as

$$S_v(f) \cong \frac{A^2}{4} S_\phi\left(\left|f - \frac{\omega_0}{2\pi}\right|\right) \quad \text{for } f \neq \frac{\omega_0}{2\pi}. \tag{23}$$

This can be rearranged and combined with (18) to give

$$\mathcal{L}(f) \cong \frac{S_v\left(\frac{\omega_0}{2\pi} + f\right)}{A^2/2} \quad \text{for } f > 0. \tag{24}$$

The right side of (24) used to be the *definition* of $\mathcal{L}(f)$ [22]. When this was the case, (18) followed from this definition as an approximation. The old definition had the advantage that it was a proxy for $S_\phi(f)/2$ that could easily be measured directly from an oscillator signal's power spectrum. Unfortunately, the simple relationship between the old definition of $\mathcal{L}(f)$ and $S_\phi(f)$ breaks down unless $|\phi(t)| << 1$. Given that practical cases of interest exist for which $|\phi(t)| << 1$ does not hold, the definition was eventually changed to (18) [10].

The error term in (22) caused by oscillator phase error can be viewed as an amplitude modulated (AM) waveform with *carrier signal* $A\cos(\omega_0 t)$ and modulation signal $\phi(t)$. The time average power of the carrier signal is $A^2/2$, and for $f \neq \omega_0/(2\pi)$ the one-sided power spectrum of the AM waveform is $S_v(f)$. The portion of an AM waveform's power spectrum on either side of the carrier frequency is traditionally called a *sideband* of the signal. Therefore (24) implies that $\mathcal{L}(f)$ is approximately equal to the power density per Hz in a sideband of $v(t)$ at an offset of f from the carrier divided by the power of the carrier.[§]

It follows that 10 times the logarithm of the right side of (24) can be interpreted as having units of dBc/Hz, i.e., the power density per Hz in a sideband relative to the power of the carrier in dB. Therefore, with the old definition of $\mathcal{L}(f)$, the units of $10\log_{10}(\mathcal{L}(f))$ are dBc/Hz. This is not strictly true for the current definition of $\mathcal{L}(f)$, but for consistency, and because (24) holds approximately with the current definition, the units of $10\log_{10}(\mathcal{L}(f))$ are *defined* to be dBc/Hz.

If $S_\phi(f)$ contains a spur at f_m, the argument above can be repeated while integrating both sides of (23) over an ε-interval around $f_0 + f_m$ and applying (19). This shows that P_{spur} is equal to the power of the spur in $S_v(f)$ at $f_0 + f_m$, divided by the power of the carrier. Hence, the units of $10\log_{10}(P_{spur})$ are *defined* to be dBc.

By the symmetry of $S_{\phi\phi}(f)$, if $S_v(f)$ has a spur at $f_0 + f_m$ caused by oscillator phase error, it also has a spur of the same power at $f_0 - f_m$. Therefore, it might seem redundant to evaluate the power of each spur individually, e.g., as in (19). However, for reasons explained in Section VII, measurements of oscillator signal power spectra are seldom perfectly symmetric for frequency offsets above and below f_0, and therefore evaluating the power of each spur separately is not redundant in practice.

B. Continuous-Time Versus Discrete-Time Paradox

It is argued in Section II that only the samples of a squared-up oscillator signal's phase error at the oscillator signal's zero-crossing times, i.e., $\phi(t_n)$, affect the oscillator output waveform in the absence of $e(t)$ error. Yet it is argued in Section III-A that in many cases of interest the power spectrum of the fundamental term of such an oscillator signal is a function of the power spectrum of the continuous-time phase error waveform, $\phi(t)$. This apparent contradiction can be resolved as follows.

As described in Section II-C, a squared-up oscillator signal can be used in place of a sinusoidal oscillator signal provided the power spectra of the terms in the squared-up signal centered at integer multiples of $f_0 = \omega_0/(2\pi)$ do not overlap the power spectrum of the fundamental term centered at f_0. This requires that $S_{\phi\phi}(f)$ be bandlimited to a bandwidth that is lower than $f_0/2$. In such cases the definition of $S_{\phi\phi}(f)$ given by (11)-(13) implies that $FT\{\phi_T(t)\}$ is also bandlimited with this bandwidth in the limit as $T \to \infty$. It follows from the sampling theorem that $FT\{\phi_T(t)\}$ in (11) can be replaced

[§] Older literature often shows phase noise as $\mathcal{L}(\Delta f)$ to underscore this fact.

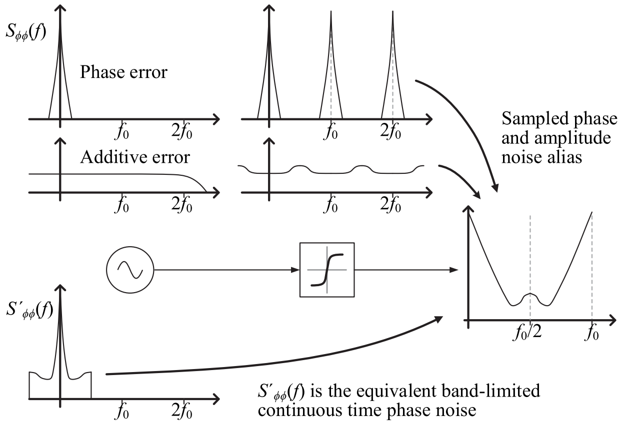

Fig. 7. Sampling behavior of the limiting amplifier on the oscillators phase and additive error.

by T_0 times the discrete-time Fourier transform of $\phi(nT_0)$ for $|f| < f_0/2$. Hence, in such cases

$$S_{\phi\phi}(f) = \begin{cases} \lim_{T\to\infty} \frac{T_0^2}{T} \left| \sum_{n=-\infty}^{\infty} \phi_T(nT_0)\, e^{-j2\pi f n T_0} \right|^2, & \text{if } |f| < \frac{f_0}{2}, \\ 0, & \text{otherwise.} \end{cases} \quad (25)$$

Provided the bandwidth of $S_{\phi\phi}(f)$ is sufficiently low, (10) implies that $\phi(\tau_n) \cong \phi(nT_0)$ to a high degree of accuracy. In such cases, it follows from (25) that $S_{\phi\phi}(f)$ can be expressed as a function of $\phi(t)$ sampled at the oscillator signal's positive-going zero-crossing times. A nearly identical argument shows that it can also be expressed as a function of $\phi(t)$ sampled at the oscillator signal's negative-going zero-crossing times. Thus, provided $|\phi(t)| << 1$, it follows from (17), (23), and (25) that $S_v(f)$ can be expressed as a function of $\phi(t_n)$ instead of $\phi(t)$.

C. Sampled Phase Error

The assumption that $S_{\phi\phi}(f) = 0$ for $|f| \geq f_0/2$, which led to (25), is usually acceptable in wireless transceivers, because the various signal paths in wireless transceivers typically have bandwidths much less than $f_0/2$ by design. In contrast, the assumption is rarely valid in clocked systems which generally lack equivalent band-limiting circuits. In such cases, (25) does not hold because sampling the phase error at a rate of f_0 causes aliasing.

A common approach to avoid this problem is to replace the actual phase error power spectrum, $S_{\phi\phi}(f)$, with a conceptual modified phase error power spectrum, $S'_{\phi\phi}(f)$, equal to the right side of (25). Even though $S'_{\phi\phi}(f) \neq S_{\phi\phi}(f)$ when $S_{\phi\phi}(f)$ is not bandlimited to $|f| < f_0/2$, $S'_{\phi\phi}(f)$ is bandlimited to $|f| < f_0/2$ and already contains all content that would have aliased into the frequency band $|f| < f_0/2$ had $S_{\phi\phi}(f)$ been sampled at a rate of f_0. Accordingly, $S'_{\phi\phi}(f)$ can be used in place of $S_{\phi\phi}(f)$ in all of this paper's results that relate the phase error power spectrum to f_0-rate samples of the phase error or, by extension, to absolute jitter.

Fig. 7 shows how this approach can be extended to handle broadband additive error that gets converted to phase error when a sinusoidal oscillator signal is amplified and clipped by a limiting amplifier to generate a squared-up oscillator signal. Both the broadband phase and the broadband additive error are sampled at the zero crossings of the oscillator signal. The limiting amplifier's wide bandwidth typically causes aliasing of many bands of broadband noise into the $|f| < f_0/2$ frequency band.

The conversion of additive error to phase error also affects the clock buffers which follow the limiting amplifier. The purpose of a clock buffer is to maintain the squared-up edges of the clock signal over long transmission distances, and therefore is susceptible to the same error-sampling effect discussed for limiting amplifiers. While each clock buffer may only contribute a small amount of phase error, in ICs such as network routers or FPGAs where GHz clocks are distributed several centimeters, their cumulative error usually dominates the oscillator's intrinsic phase error [23], [24].

The simulation methodology described in Section VI-A ensures that the sampling and conversion effects described above are properly captured, so that any amplitude error on the resulting squared-up clock signal can usually be ignored, as explained in Section II-B. The phase error power spectrum of this squared-up clock signal can then be "input-referred" as shown in Fig. 7 to generate the equivalent band-limited $S'_{\phi\phi}(f)$ which contains all the sampled noise.

Throughout the remainder of the paper, $S_{\phi\phi}(f)$ is used in equations that relate the phase error power spectrum to samples of the phase error or, equivalently, to jitter, as a proxy for either the actual phase error power spectrum or $S'_{\phi\phi}(f)$. In cases where the actual phase error is bandlimited to $|f| < f_0/2$, the reader should view $S_{\phi\phi}(f)$ as denoting the power spectrum of the actual phase error. Otherwise, the reader should view $S_{\phi\phi}(f)$ as denoting $S'_{\phi\phi}(f)$. In this way, the quantity denoted by $S_{\phi\phi}(f)$ in such equations is bandlimited to $|f| < f_0/2$ by definition.

D. Why $S_{\phi\phi}(f)$ Increases With Oscillator Frequency

As can be seen from (1) and (6), oscillator frequency, ω_0, and phase error, $\phi(t)$, are separate variables in both the sinusoidal and squared-up oscillator signal models. This separation obscures an implicit dependency of phase error on oscillator frequency that arises because phase error represents the oscillator's timing error as a fraction of its period but its period is inversely proportional to its frequency. For example, (10) implies that changing an oscillator's frequency by X dB without changing the mean squared value of its absolute jitter would increase its phase error power by $2X$ dB.

Indeed, it has been shown for many different types of oscillators that the effect of circuit noise on absolute jitter is relatively independent of the frequency to which a given oscillator is tuned [25]–[28]. In such cases, changing the oscillator's frequency does not significantly change the mean squared value of its absolute jitter.

This effect can cause confusion when comparing the phase noise performance of oscillators tuned to different frequencies. The confusion can be avoided by *normalizing* the phase noise of the oscillators to a common frequency prior to comparing their phase noise spectra. The phase noise of an oscillator normalized to a frequency ω_n is $(\omega_n/\omega_0)^2 S_{\phi\phi}(f)$ where ω_0 is the nominal frequency and $S_{\phi\phi}(f)$ is the phase noise spectrum of the oscillator. In general, the phase noise of two oscillators normalized to the same frequency will be approximately equal if they have equivalent jitter performance. For example, two oscillators, one with frequency ω_0 and phase error power spectrum $S_{\phi\phi}(f)$ and one with frequency $2\omega_0$ and phase error power spectrum $4S_{\phi\phi}(f)$ would have the same normalized phase noise power spectrum. Normalizing the phase noise spectra of different oscillators to a common frequency in this fashion allows for a fair phase noise performance comparison.

A commonly-used oscillator figure of merit (FOM) that implicitly performs this normalization is

$$FOM = 10\log\left[\left(\frac{f_0}{\Delta f}\right)^2 \frac{1}{\mathcal{L}(\Delta f)\,P}\right] \tag{26}$$

where $f_0 = \omega_0/(2\pi)$, Δf is any frequency offset, and P is the oscillator's power consumption in mW [25]. For instance, if an oscillator's frequency is doubled (increased by 3 dB) without changing its phase noise or power consumption, its FOM increases by 6 dB, which is consistent with the above discussion.

IV. Effects of Phase Error on Mixing

In wireless receivers, a mixer's input signal, $x(t)$, usually consists of a *desired signal* plus multiple unwanted signals at other frequencies called *interferers* or *blockers*. The interferers not only arise from unwanted signals received through the antenna, but they can also result from non-ideal circuit behavior. For example, in frequency division duplex (FDD) transceivers it is not possible to fully isolate the receiver from the transmitted signal, so a partially attenuated version of the transmitted signal appears to the receiver as an interferer.

Ideally, the mixer frequency-translates $x(t)$ such that the desired signal resides within the passband of the filter following the mixer, and the interferers are either suppressed by the filter or subsequently suppressed by digital filtering after analog-to-digital conversion. Unfortunately, oscillator phase error mixes with the interferers which can result in error that ends up in the same frequency range as the desired signal, thereby corrupting the desired signal and reducing the sensitivity of the receiver. This phenomenon is known as *reciprocal mixing*.

It follows from (9), (20), and (22) that the output of a typical mixer in the band of interest can be written as

$$y(t) = \frac{4}{\pi}\sin(\omega_0 t)\,x(t) + \frac{4}{\pi}\cos(\omega_0 t)\,\phi(t)x(t). \tag{27}$$

The first term on the right side of (27) represents the ideal behavior of the mixer, and the second term represents the effect of oscillator phase error. The sine and cosine factors in the two terms both perform the same frequency translations; the sine factor frequency-translates $x(t)$ by both $\omega_0/(2\pi)$ and $-\omega_0/(2\pi)$ Hz, and the cosine factor frequency-translates $\phi(t)x(t)$ by both $\omega_0/(2\pi)$ and $-\omega_0/(2\pi)$ Hz. Therefore, any portion of the power spectrum of $\phi(t)x(t)$ that overlaps the power spectrum of the desired signal portion of $x(t)$ causes reciprocal mixing error.

For example, suppose $x(t)$ contains a desired signal component centered at frequency f_d and an interferer centered at frequency f_i, and suppose $\phi(t)$ contains a spurious tone given by $B\sin[2\pi(f_d - f_i)t]$ where B is a constant. The spurious tone causes $\phi(t)x(t)$ in the second term of (27) to contain a copy of the interferer scaled by $B/2$ and centered at frequency f_d, which corrupts the desired signal component unless the power spectrum of the interferer times $B^2/4$ is sufficiently small that it can be neglected.

More generally, multiplication in the time domain is equivalent to convolution in the frequency domain, so the power

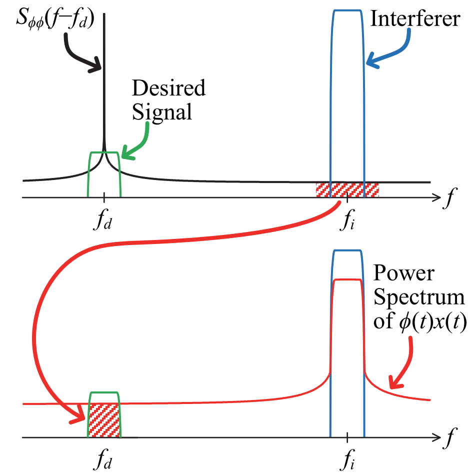

Fig. 8. Power spectra on a dBc/Hz scale of: the desired signal and an interferer in $x(t)$, $\phi(t)$ (frequency-shifted by f_d), and $\phi(t)x(t)$. Hash marks show a portion of $\phi(t)$ that causes reciprocal mixing error (top plot) and the resulting reciprocal mixing error (bottom plot).

spectrum of $\phi(t)x(t)$ has a non-zero component in the frequency range $f_d - \Delta f/2$ to $f_d + \Delta f/2$ if an interferer with bandwidth Δf is centered at frequency f_i and $S_{\phi\phi}(f - f_d)$ is non-zero for $f_i - \Delta f \le |f| \le f_i + \Delta f$. This component is centered on and therefore corrupts the desired signal component of $x(t)$. As illustrated in Fig. 8, if the interferer has a high-enough power relative to the desired signal, the reciprocal mixing error can be significant even when $S_{\phi\phi}(f - f_d)$ is small for $f_i - \Delta f \le |f| \le f_i + \Delta f$.

Thus, reciprocal mixing error is caused by phase error mixing with interferers via the $\phi(t)x(t)$ term in (27), and the corresponding signal-to-noise ratio (SNR) over the desired signal band decreases with increasing interferer power. This often places stringent restrictions on the phase error, because interferers can be extremely powerful relative to the desired signal in typical wireless receivers.** For example, in LTE cellular handset receivers, interferers with $|f_i - f_d| \ge$ 10 MHz can be up to 87 dB more powerful than the desired signal [29].

Of course, the phase error also mixes with the desired signal component via the $\phi(t)x(t)$ term, and this causes error in the desired signal band too. In applications such as GSM handset receivers wherein the required SNR over the signal band is modest, this error is less problematic than reciprocal mixing error because the SNR does not decrease with increasing signal power. However, in receivers for high-order modulation formats, such as 256 QAM in LTE handset receivers, the required SNR over the signal band is high enough that close-in oscillator phase error, i.e., phase error below several MHz, becomes one of the most challenging oscillator specifications to meet [21].

Oscillator phase error also causes mixing error in wireless transmitters. Mixers are used in a wireless transmitter to frequency-translate the desired signal to the RF transmit frequency. As in receivers, each mixer behaves according to (27); it performs the same frequency translations on both its input signal, $x(t)$, and the phase error term, $\phi(t)x(t)$. After frequency translation, these signals have spectral shapes comparable to the interferer and $\phi(t)x(t)$ power spectra,

**Typically, large interferers occur relatively far from f_d, so they place stringent requirements on the *far-out* phase error, i.e., on the phase error power spectrum above at least several MHz.

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

respectively, shown in the bottom plot of Fig. 8, except in this case f_i corresponds to the transmit frequency.

The combined signals are passed through a power amplifier and then bandpass filtered. Typically, the bandpass filter is a SAW filter or duplexer which attenuates the out-of-band portion of $\phi(t)x(t)$ by about 50 dB, but the amplification prior to filtering is often so high that the residual out-of-band emission from the term corresponding to $\phi(t)x(t)$ can still be significant. For example, in FDD transceivers, its overlap with the receive band can desensitize the receiver. This out-of-band emission phenomenon generally dictates the maximum far-out phase error that can be tolerated in a given application.

The close-in phase error results in the portion of the error power spectrum arising from $\phi(t)x(t)$ that overlaps the desired transmitted signal. As in receivers, this term can be problematic in applications that use high-order modulation formats such as the 256 QAM option in LTE handset receivers unless the close-in phase error is kept very low.

V. Jitter

In mixers it is convenient for the reasons described above to represent non-ideal oscillator behavior in terms of continuous-time phase error, $\phi(t)$. In most other applications, it is more useful to represent the non-ideal behavior in terms of the absolute jitter, $\Delta\tau_n$, as defined in (10), or quantities derived from the absolute jitter.

Like phase error, jitter typically consists of random and deterministic components. Accordingly, the modifiers *random* or *deterministic* are applied to jitter metrics in cases where the metrics represent only the effects of the random or deterministic components, respectively. The modifier *total* is also often used to indicate when a jitter metric represents the effects of both components, e.g., total absolute jitter is comprised of random absolute jitter and deterministic absolute jitter. Without a modifier, the total quantity is usually implied.

A. Jitter Density

In many clocked systems, the clock signal undergoes several frequency multiplications and divisions between its source and each circuit that it drives. In such cases, phase error is an inconvenient metric, because, as discussed in Section III-D, it increases with clock frequency. To avoid the tedium of scaling phase error after every frequency translation, a normalized power spectrum called jitter density and defined as

$$J(f) = \frac{1}{\omega_0^2} S_\phi(f) \tag{28}$$

is often used in place of the phase error power spectrum when analyzing such systems. Jitter density is so-named because for $f < f_0/2$ it is proportional to the single-sided discrete-time power spectrum of the absolute jitter sequence, $\Delta\tau_n$. This can be seen by dividing both sides of (25) by ω_0^2, and using (10) and the symmetry property (17). The units of $\Delta\tau_n$ are seconds, so the units of $J(f)$ are seconds^2 per Hz.

The jitter densities of oscillator signals at different nominal frequencies do not need to be normalized to a common frequency for comparison as would be necessary for the corresponding phase error power spectra, because the normalization is built into the definition of jitter density. This property simplifies many practical computations as demonstrated in the following sections.

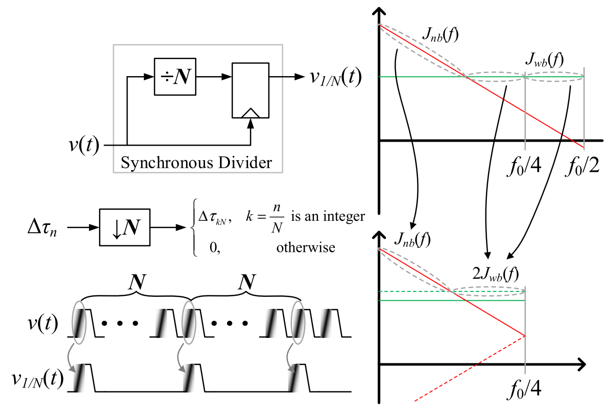

Fig. 9. Divide-by-N circuit block diagram and effect on input clock jitter.

B. Jitter of Divided Clocks

A clock divider that reduces the frequency of its input clock by an integer factor, N, is a fundamental building block in many clocked systems. As shown in Fig. 9, such a clock divider typically consists of edge-selection logic, i.e., an N-fold digital divider driven by the oscillator signal, $v(t)$, followed by a retiming flip-flop. The retiming operation renders the absolute jitter of the output clock signal insensitive to noise and other non-ideal circuit behavior within the edge-selection logic, thereby isolating the jitter-contributing circuitry to just the retiming flip-flop. To the extent that the absolute jitter added by the retiming flip-flop is negligible, the output clock signal's absolute jitter is approximately equal to a sub-sampled version of the input clock signal's absolute jitter.

The jitter density of a typical clock signal has regions with different slopes [30]. The top right plot in Fig. 9 shows the jitter density of the divider's input clock, where the sloped regions have been simplified into a single low-frequency (close-in) sloped region with jitter density $J_{nb}(f)$, and a broadband (far-out) flat region with jitter density $J_{wb}(f)$. The bottom right plot in Fig. 9 shows the jitter density of the divider's output clock signal for the case of division by 2 (i.e., $N = 2$). In this example, the power of the input clock signal's close-in jitter density falls far below the power of the far-out jitter density in the frequency band $[f_0/4, f_0/2)$, so the aliased contribution of the close-in jitter density to the jitter density of the output clock signal is negligible. The far-out jitter density of the output clock signal is, however, doubled in power because of aliasing introduced by the sub-sampling.

Generalizing this example to N-fold division, the output clock's close-in jitter density remains that of the input clock, whereas its far-out jitter density is N times that of the input clock. In theory, it is possible for the division ratio, N, to be large enough that aliasing of the close-in jitter density is non-negligible, but this rarely happens in practical low-jitter circuits.

C. Integrated Jitter

In some applications, the mean squared value of the absolute jitter, i.e., the time average of $(\Delta\tau_n)^2$, is of interest. This can usually be related to the phase error spectrum as shown below.

Equation (25) can be rearranged as

$$S_{\phi\phi}(f) = T_0 \sum_{k=-\infty}^{\infty} R_{\phi\phi}[k] e^{-j2\pi f k T_0}, \text{ for } |f| < \frac{f_0}{2} \tag{29}$$

where

$$R_{\phi\phi}[k] = \lim_{N\to\infty} \frac{1}{N} \sum_{n=0}^{N-1} \phi\,(nT_0)\,\phi\,((n+k)\,T_0). \tag{30}$$

The summation on the right side of (29) has the form of a discrete-time Fourier transform. Applying the inverse discrete-time Fourier transform equation gives

$$R_{\phi\phi}[k] = \int_{-\infty}^{\infty} S_{\phi\phi}\,(f)\,e^{j2\pi f k T_0} df. \tag{31}$$

It follows from (30) that the mean squared value of $\phi(nT_0)$ is $R_{\phi\phi}[0]$, so (31) implies that the mean squared value of $\phi(nT_0)$ is equal the integral of $S_{\phi\phi}(f)$ over all frequencies. To the extent that $\phi(\tau_n) \cong \phi(nT_0)$, this with (10) implies that the mean-squared value of the absolute jitter is given by

$$\sigma_{abs}^2 = \frac{1}{\omega_0^2}\int_{-\infty}^{\infty} S_{\phi\phi}\,(f)df = \frac{1}{\omega_0^2}\int_{0}^{\infty} S_{\phi}\,(f)df. \tag{32}$$

This and other mean-squared quantities derived from the absolute jitter that are defined in the next section are often stated as their square-root equivalents, i.e., the root-mean-squared (rms) absolute jitter σ_{abs}.

Section III-C defined the use of $S_\phi(f)$ in this paper as a proxy for either the one-sided power spectrum of the continuous-time phase error, or $S'_\phi(f)$, the conceptual modified power spectrum. As explained in that section, clocked circuits tend to contain sources of broadband noise which result in broadband jitter, and the applications which they serve also tend to have wide bandwidths. For these reasons, it is almost always $S'_\phi(f)$ that is used in applications that specify jitter. Irrespective of whether $S_\phi(f)$ or $S'_\phi(f)$ is appropriate, both are bandlimited to the frequency range $[0,\ f_0/2)$. Therefore, (32) can be rewritten as

$$\sigma_{abs}^2 = \frac{1}{\omega_0^2}\int_{0}^{f_0/2} S_{\phi}(f)df = \int_{0}^{f_0/2} J(f)df. \tag{33}$$

Actually integrating the jitter density down to $f = 0$ would yield an unbounded result in practice, because oscillator phase error power spectra (and correspondingly, jitter density) tend to be unbounded at low frequencies [30]. Usually this is not a practical concern, though, as many clocked systems are designed to be insensitive to low-frequency jitter density. For example, transmitter-receiver pairs in serial communication links use PLL-based oscillators and clock and data recovery (CDR) circuits configured in a manner that effectively bandpass-filters the transmitter clock jitter density before it reaches timing-sensitive circuits in the receiver [31].

Hence, in practice the integration limits of 0 and $f_0/2$ in (33) are usually replaced by limits that are greater than 0 and less than $f_0/2$, respectively. In particular, integration limits of 12 kHz to 20 MHz are widely used. These limits were first seen in the specifications for the PLL-based oscillators designed for the synchronous optical network (SONET) [32]. They have since proliferated to applications where they no longer have any architectural significance, but because most PLL-based oscillators have qualitatively similar phase error power spectra, they provide a useful means by which to quickly compare the performance of different PLL-based oscillators.

In some applications, it is of interest to evaluate the contribution of a phase error spur at a particular frequency to σ_{abs}^2. It follows from (33) that this contribution can be written as

$$\sigma_{abs-spur}^2 = \frac{1}{\omega_0^2}\lim_{\varepsilon\to 0}\int_{f_m-|\varepsilon|}^{f_m+|\varepsilon|} S_\phi(f)df \tag{34}$$

where f_m is the frequency of the spur. Hence, (17) and (19) imply that

$$\sigma_{abs-spur}^2 = \frac{2}{\omega_0^2}P_{spur} = \frac{2}{\omega_0^2}10^{P_{dBc}/10} \tag{35}$$

where P_{dBc} is the spur power in units of dBc, obtained through measurements as described in Section VII.

D. Other Jitter-Based Metrics

In typical clocked digital circuits, multiple flip-flops whose inputs and outputs are interconnected through combinational logic are driven by the same clock signal. Each timing path in such a digital circuit is defined as a path that begins at the output of a flip-flop, passes through combinational logic, and ends at the input of a flip-flop. Correct functionality of the overall digital circuit requires that the maximum delay among all such timing paths, which is called the *critical path delay*, is less than the minimum clock period. The minimum clock period depends on the jitter: the higher the jitter the lower the minimum clock period and, therefore, the lower the maximum possible clock frequency of the digital circuit.

In particular the variation of a clock signal's period from its nominal value, i.e., $\Delta\tau_n - \Delta\tau_{n-1}$, is the jitter metric of interest. It is called *period jitter*, and is frequently used in place of absolute jitter to specify the performance of clock signals for digital circuits.

The concept of period jitter can be generalized to *accumulated jitter*, which is also known as *N-cycle*, or *long-term jitter*. Accumulated jitter is defined as the variation of the time interval spanned by a clock edge and the Nth edge preceding it minus N times the nominal period, so it can be expressed as $\Delta\tau_n - \Delta\tau_{n-N}$. Thus, accumulated jitter for $N = 1$ is just period jitter.

By definition, accumulated jitter is equivalent to absolute jitter passed through a discrete-time filter with transfer function $1 - z^{-N}$, so its mean squared value can be written as

$$\sigma_{acc}^2(N) = \int_{0}^{f_0/2} \left|1 - e^{-j2\pi N f/f_0}\right|^2 J(f)df. \tag{36}$$

E. Effects of Jitter on Analog to Digital Conversion

A common use of precision oscillators is the clocking of ADCs. The operation of many ADCs can be abstracted into two sequential operations. The first operation is sampling the continuous-time, continuous-valued input signal, $x(t)$, at a discrete set of times defined by the oscillator, τ_n. The second operation is quantizing the sequence of continuous-valued samples, $x(\tau_n)$, into a sequence of digital codes. In many ADC architectures, the second operation is insensitive to clock jitter, so only the initial sampling operation is analyzed here [33].

Fig. 6 shows a simplified input sampling circuit. Ignoring all non-idealities except clock jitter, the nth sample of the input voltage is given by $y[n] = x(nT_0 + \Delta\tau_n)$. Equation (10) implies that the $|\phi(t)| << 1$ assumption made in previous sections is equivalent to $|\Delta\tau_n| << T_0$. Assuming $x(t)$ is bandlimited to prevent aliasing, the sampling error caused by

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

jitter can be approximated as $y[n] \cong x(nT_0) + x'(nT_0) \cdot \Delta\tau_n$, where $x'(t)$ is the first derivative of $x(t)$.

Even with these assumptions and approximation, the jitter-induced error, $x'(nT_0) \cdot \Delta\tau_n$, is hard to analyze without simulation. Textbooks typically consider a worst-case scenario in which $x(t)$ is a maximum-frequency, maximum-amplitude sinusoid, $x(t) = A_{max} \sin(\omega_x t)$ [33]. In this case, the sampled voltage error is approximately $A_{max}\omega_x \Delta\tau_n \cos(\omega_x nT_0)$. Thus, prior to quantization, this jitter-induced error already sets an upper bound on the maximum achievable SNR given by

$$SNR_{max} = 10\log_{10}\left(\frac{\frac{1}{2}A_{max}^2}{\frac{1}{2}A_{max}^2\omega_x^2\sigma_{abs}^2}\right) = 10\log_{10}\left(\frac{1}{\omega_x^2\sigma_{abs}^2}\right). \tag{37}$$

Restricting $x(t)$ to be a single sinusoid makes the analysis tractable, but it has been shown empirically that the SNRs of ADCs with broadband input signals achieve much higher performance than the single-tone result would suggest [34]. Yet (37) remains useful as pessimistic performance bound.

Moreover, σ_{abs}, the clock quality metric on which SNR_{max} depends, is often used instead of $S_\phi(f)$ to stipulate the required ADC clock quality in applications. When an ADC's input signal, $x(t)$, is broadband, the error caused by jitter, i.e., $x(nT_0 + \Delta\tau_n) - x(nT_0)$, tends to be spread across the Nyquist band in a way that does not depend strongly on the shape of $S_\phi(f)$, so σ_{abs} is a more convenient metric. Notable exceptions are software-defined and reconfigurable radios in which ADCs capture broad spectrum signals, but the constituent desired signals and interferers each occupy relatively narrow bandwidths [35]. In such cases, $S_\phi(f)$ is required to characterize the full radio performance [36].

F. Jitter Histograms

Equation (33) provides a means with which to calculate σ_{abs} from $J(f)$, but it is possible to approximate σ_{abs} without knowing $J(f)$ by plotting a histogram of the $\Delta\tau_n$ sequence. The histogram provides an estimate of the probability distribution of $\Delta\tau_n$ from which σ_{abs} can be calculated directly. This is useful in applications like ADC clocking wherein $J(f)$ is not particularly useful. Furthermore, with reasonable assumptions on the sources of the random and deterministic absolute jitter contributions to the total absolute jitter, it is possible to estimate the relative magnitudes of their contributions [37].

For many clocked systems, jitter histograms are also a useful visualization of clock signal phase error. For example, in a serial communication receiver there is a circuit which samples incoming data at time intervals controlled by a sampling clock. As described in Section VII-C, it is usually possible to observe the waveforms of the data and sampling clock in the time domain. Knowing the histogram of the sampling clock jitter shows the distribution of when the data is sampled in time, which is usually informative because realistic absolute jitter histograms are non-Gaussian. In contrast, simply calculating σ_{abs} by integrating $J(f)$ would give no information on the distribution of $\Delta\tau_n$.

VI. Phase Error Simulation

In principle, an oscillator's phase error could be estimated via transistor-level transient simulations. However, this is rarely practical, even with present-day high-performance computers. The difficulty arises from the large timescales necessary to capture phase error quantities of interest. For example, to measure the phase error of a 1 GHz oscillator down to 1 kHz would require running a transient simulation for one million oscillator periods just to capture a single period of a 1-kHz phase error component.

For the past few decades, simulation methods have been available that greatly accelerate the analysis of relatively simple circuits that generate and process periodic signals, such as oscillators and clock buffers. While such methods are much faster than direct transient simulation, they have limitations, as described below, which make them generally inapplicable to the simulation of larger systems such as PLL-based oscillators.

For larger systems, it is more common to use the data generated in block transistor-level simulations to develop behavioral models, which are then used to simulate the complete system in an event-driven simulator. As explained below, event-driven simulators are ideally suited for simulating large systems that process discrete events, such as an oscillator's zero crossings.

A. Transistor-Level Simulation

To perform a phase noise or random jitter simulation, the simulator must first compute the circuit's periodic steady state, which is comprised of the periodic waveforms at the various circuit nodes after all transient settling effects have died out. For example, when an oscillator with a high-Q resonator is started, it usually takes many cycles for the oscillation envelope to stabilize. In contrast, a clock buffer driven by an external clock signal may only need a few cycles to stabilize [38]. There are two main techniques used to compute this solution: *shooting* and *harmonic balance* [39]–[41].

The shooting method is a time-domain approach wherein the circuit is iteratively simulated over a time interval equal to the expected period. After each iteration, the differences between the circuit's initial and final states are measured, and the initial state for the next iteration is modified. This process continues until the circuit's final state is equal to its initial state, indicating the periodic solution is achieved [39]. The majority of the computation time is spent on transient simulations, so this method is especially effective for circuits with short periodicities, e.g., multi-GHz RF oscillators. Moreover, as the waveforms are computed in the time-domain, sharp transitions resulting in broadband frequency-domain content, such as occur in ring oscillators and CMOS clock buffers, are not problematic [42].

In contrast, harmonic balance is considered a frequency-domain method. The circuit is first divided into two partitions, one containing all the linear elements, and another containing the nonlinear elements. The simulator iteratively solves for the voltages at, and current through, the nodes at the interface between these two partitions. Given an initial guess at the node voltages, the currents into the linear partition are solved in the frequency domain, which are computed using simple matrix multiplications. This is much faster than a transient simulation which requires iteratively solving differential equations. The currents into the nonlinear partition are solved by first applying the Inverse-Fast-Fourier Transform (IFFT) to the node voltages, computing the currents in the time domain, then returning them to the frequency domain via the FFT. By Kirchoff's current law, the current into the linear partition must equal the current leaving the nonlinear partition, so the simulator iterates until this condition is satisfied.

The computation of the nonlinear currents is cumbersome, but the huge speed advantage gained in the frequency-domain computation of the linear currents gives the harmonic balance method an overall advantage for circuits which are weakly nonlinear, only have a few nonlinear elements, or have extremely long periodicities. For weakly nonlinear circuits, the overhead of the FFT and IFFT operations is modest because the circuit solution can be represented by only a few harmonics in the frequency domain, and by having few nonlinear elements, the quantity of time-domain computations is reduced. For circuits with long periodicities, to achieve a given accuracy the harmonic balance method is usually more efficient than the shooting method because it does not suffer from accumulated numerical noise which plagues long transient simulations. Ideal candidates for harmonic balance analysis are low-noise, extremely high-Q systems such as crystal oscillators [42].

Regardless of how the circuit's periodic steady state solution is calculated, the second step in a phase noise or random jitter simulation is to use this solution to compute how the different circuit noise sources contribute to voltage noise on the periodic signal at a chosen node. This method takes advantage of the periodicity of the signals within the circuit so that the effect of very low frequency circuit noise can be analyzed using only a single period of the circuit's behavior [43], [44].

The distinctions made in the above discussion regarding the types of circuits best suited to either the shooting or harmonic balance methods were more critical when these methods were first developed. However, the rapid improvement in simulation and computer technology over the intervening decades has blurred this separation. While there do still exist specialized problems which are better served by one or the other method, there are now general-purpose industry-standard simulation tools based on both methods, as well as hybrid methods [45], [46].

Moreover, while early incarnations of these simulators only estimated the raw signal power spectrum, present-day simulators are able to estimate both phase and amplitude error power spectra, and decompose the spectra into upper- and lower-sideband contributions [47]. Furthermore, provided the necessary information, they can accurately compute the modified phase error power spectrum resulting from sampling and/or noise folding, as discussed in Section III.C [48].

Simulation technology has reached a level of maturity where the simulator tool is typically no longer the design bottleneck. Rather, assuming the device models are accurate, most analysis errors are the result of improper or incomplete simulation methodology. For example, a complete clock distribution system comprised of an oscillator, limiting amplifier, and clock buffers is usually still too large to be simulated efficiently in one circuit. Rather, it is partitioned into sections which are simulated independently. Unfortunately, oscillators and buffers tend to be extremely sensitive to the circuits they drive, and buffers additionally tend to be sensitive to the circuits that drive them [49], [50]. Therefore, realistic driving and loading conditions must be included when simulating each partition.

In particular, the oscillator and its limiting amplifier is a critical combination, as the interface between them is especially susceptible to broadband additive noise to phase noise conversion. When simulating the jitter of an oscillator and its limiting amplifier, the clock buffer following the limiting amplifier should be included. This only serves as a realistic load, for reasons described above, as the simulation analysis is performed on the clock signal at the output of the limiting amplifier.

In spite of the advancements in simulator technology, there are still situations where simulating the combined oscillator and limiting amplifier can cause the simulator to have convergence issues. For example, this often occurs when simulating a high-Q crystal oscillator driving a CMOS limiting amplifier. For the reasons discussed above, the oscillator and its limiting amplifier are optimally simulated using the harmonic balance and shooting methods, respectively, so their combination is challenging for either simulator. In this situation, what is often done is to perform separate simulations on each block, while loading the oscillator with an approximation of the CMOS buffer load, and driving the CMOS buffer with an approximation of the oscillator's sinusoidal output. The two results are then added in power to represent the total phase noise at the CMOS buffer's output.

When simulating clock buffer chains, their layout strongly influences the clock signal waveforms, which in turn influences their jitter performance. It is therefore important to electromagnetically model the interconnect between the clock buffers, including both the signal and ground return path, to account for potential transmission-line effects that are not accurately captured in RC parasitic extractors [51]. The results from electromagnetic simulators are invariably high-order frequency-domain models in the form of s- or y-parameters. Such models are most efficiently simulated in their native form using the harmonic balance method. Alternatively, their dominant effects can usually be captured in lumped element models to be used with transient-based simulators, and some simulators can automatically perform this conversion [46], [47].

Clock buffer chains are also highly sensitive to their power supplies and power distribution networks. This includes the interconnect between the clock buffers' voltage regulators, the voltage regulators themselves, and even the package models. Clock buffer supply current waveforms usually have broadband content at frequencies well beyond a voltage regulator's bandwidth. At these high frequencies, there are often parasitic resonances. Therefore, accurate clock buffer simulations require not only the circuits and their drivers and loads, but electromagnetic models of their complete on-chip environment, and their supply circuitry as well [52]–[54].

B. Behavioral Simulation

There are two major obstacles to using the simulation methods described above for more complicated systems such as PLL-based oscillators. The first obstacle is that a periodic steady-state solution must exist and be computed [42]. This precludes the simulation of fractional-N type PLL-based oscillators, as well as systems which use digitally-generated pseudorandom dither, because such circuits tend to have extremely long periods (e.g., several days) [48]. The second obstacle is that even in cases where the period is relatively small, the complexity of a complete PLL-based oscillator imposes impractical processing and memory requirements, and even when these resources are available, computing the periodic steady-state solution can still take days. This simulation methodology would provide little utility during the design phase, where many iterations usually are required.

A more practical method for simulating PLL-based oscillators is to use behavioral event-driven simulations [55]–[57]. These simulations offer a huge speed advantage over transient

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

transistor-level simulations, because behavioral event-driven simulations only capture the relevant functionality of circuits rather than all of the voltages and currents of every node with very small time steps. For example, a transistor-level oscillator simulation iteratively solves differential equations to compute many points per period of the oscillator waveform. Yet, as has been argued throughout this paper, many circuits driven by an oscillator signal are insensitive to the oscillator signal at times other than its zero crossings. Therefore, PLL-based oscillators are almost always behaviorally-modeled by generating the sequence of their rising or falling zero-crossing times, τ_n. This sequence can be processed using (10) and (25) to generate the conceptual band-limited phase error power spectrum, $S'_{\phi\phi}(f)$. Even though the shape of the oscillator signal waveform is lost in a behavioral simulation, many non-idealities such as noise and nonlinearity can still be accurately modeled [55]–[59].

The main drawback of behavioral event-driven simulations is that their results are only as accurate as their behavioral models of the circuit's sub-blocks. Care is therefore required to continually cross-verify the behavioral models against transistor-level simulations of the corresponding circuit sub-blocks.

Mixed-signal simulators, or *co-simulators*, bridge the gap between the speed of behavioral event-driven simulators and the functional accuracy of transistor-level simulators [60]. A co-simulator can efficiently simulate a system comprised of transistor circuits, behaviorally-modeled circuits, and fully-digital circuits. The behavioral models are written in a mixed-signal language such as Verilog-AMS, while the digital circuits can be described in Verilog. The co-simulator partitions the system into its different types, and runs simulators optimized for each type in lock-step. Data between the simulated partitions is continually communicated, so that the whole system is simulated in unison.

The power of a co-simulation is that a design can start with all behavioral models, and then transistor-level circuit blocks can be used in place of the corresponding behavioral models as they are designed. As different combinations of transistor-level circuits and behavioral models can be run, this implicitly improves the behavioral model verification. Unfortunately, because the analog portions of the system are simulated with a transistor-level simulator, co-simulations still take orders of magnitude more simulation time than fully behavioral simulations.

VII. Phase Error Measurement

A. Spectrum Analyzers

A spectrum analyzer performs the equivalent of passing its input signal through a bandpass filter whose center frequency is swept across a user-defined range of frequencies, called the *span*.[††] This frequency-swept bandpass filter is called the *equivalent bandpass filter*. The power of the equivalent bandpass filter's output signal is measured, and this power is plotted on the spectrum analyzer display's vertical axis on a dB scale against the equivalent bandpass filter's center frequency on the horizontal axis. The bandwidth of the equivalent bandpass filter is called the *resolution bandwidth* (RBW), and the time it takes for the equivalent bandpass filter's center frequency to traverse the span is called the *sweep time*.

If the spectrum analyzer input is driven by an oscillator signal, then the displayed power on the spectrum analyzer at a frequency f is equal to the power of the oscillator signal in the frequency band $(f - \text{RBW}/2, f + \text{RBW}/2)$. This is proportional to an estimate of the power spectrum of the oscillator signal, $S_v(f)$, as described in Section II-F. The accuracy of this estimate improves as the RBW decreases.

If the assumptions that $\varepsilon(t) \cong 0$ and $|\phi(t)| << 1$ made in Section III-A are valid, then (23) can be used to estimate $S_\phi(f)$ from $S_v(f)$. As oscillator phase error power spectra tend to rise at low frequencies, this approximation becomes invalid at very low offsets from the carrier [61]. Extremely low frequency oscillator phase error is indistinguishable from frequency drift, which can also corrupt spur measurements, as explained shortly.

Most spectrum analyzers have modes specifically for phase *noise* measurement which perform the carrier power normalization and apply any instrument-specific correction factors automatically. Examples of correction factors include accounting for the exact transfer function of the equivalent bandpass filter, and accounting for non-idealities in the power detector that follows the equivalent bandpass filter [62], [63].

If amplitude error is not negligible, then replacing A with $[A+\varepsilon(t)]$ in (20) results in an extra additive term $\varepsilon(t)\sin(\omega_0 t)$ in (22). Furthermore, any additive error on the oscillator signal will also appear in its power spectrum. Therefore, the oscillator signal power spectrum contains both amplitude error as well as additive error, and these errors are indistinguishable from phase error in $S_v(f)$. This results in a pessimistic estimate of $S_\phi(f)$ from $S_v(f)$.

The full form of the right side of (22) including amplitude modulation is the carrier signal plus $\varepsilon(t)\sin(\omega_0 t) + A\cos(\omega_0 t)\phi(t)$, where $\varepsilon(t)$ and $A\phi(t)$ can be viewed as complex modulation of the carrier. Therefore, when the oscillator signal has both phase and amplitude error, the sidebands of its power spectrum are no longer symmetric, in contrast to the case of pure phase error described in Section III-A. Parasitic coupling in both the circuit as well as the measurement setup can result in amplitude modulation of the oscillator signal, leading to such asymmetries. This effect is often noticeable in the spurious tones of PLL-based oscillators.

As explained in Section III-A, the power of a spur in the oscillator phase error power spectrum at frequency f_m is equal to the power of the spur in the oscillator signal power spectrum at frequency $f_0 + f_m$, divided by the power of the carrier. Thus, the phase error spur power in dBc can be read from a dB-scale oscillator signal power spectrum as the difference between the values of the power spectrum at $f_0 + f_m$ and f_0.

However, there is a caveat to this spur measurement technique. As described above, the value of the spectrum analyzer measurement at frequency $f_0 + f_m$ is the power at the output of the conceptual RBW-bandwidth bandpass filter centered at frequency $f_0 + f_m$. Therefore, the measured power is approximately the sum of the spur power plus the oscillator signal noise power density integrated over the band $(f_0 + f_m - \text{RBW}/2,\ f_0 + f_m + \text{RBW}/2)$. As the power of the spur of interest decreases, the RBW must be decreased so that the measured power at that frequency is dominated by the spur power, not the integrated noise power.

The left plot in Fig. 10 shows a PLL-based oscillator's signal measured power spectrum over a 100 kHz span centered at the 2.4 GHz carrier frequency, measured by a Rohde & Schwarz FSW13 spectrum analyzer. Symmetric spurs are visible at

[††]This is a behavioral description, not an explanation of how spectrum analyzers are implemented (a good explanation of which can be found in [17]).

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

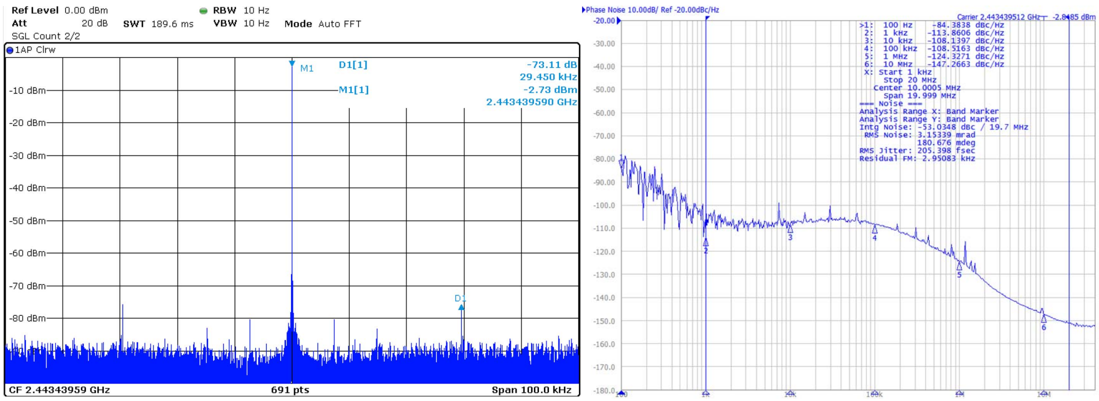

Fig. 10. Measured power spectra of a PLL-based oscillator signal (left) and its phase noise (right).

approximately 7.4, 14.7, and 29.5 kHz offsets from the carrier, with the largest shown by the delta marker as −73 dBc. The phase noise of this PLL-based oscillator in the 10 kHz to 100 kHz region is known to be −108 dBc/Hz. However, because this measurement was taken with a 10 Hz RBW and without averaging, any spur below about −80 dBc would become indistinguishable from the noise, so to resolve a lower-power spur would require a narrower RBW.

Unfortunately, narrow RBWs can complicate spur measurements and lead to erroneously optimistic results, because the time it takes for the spectrum analyzer to measure the power at the output of the equivalent bandpass filter increases as the RBW decreases. Oscillator frequency drift may cause the spur frequency to drift in the measurement. Often, this has the effect of spreading the spur power across a range of frequencies that can extend beyond the bandwidth of the spectrum analyzer's conceptual bandpass filter, thereby underestimating the spur power. Moreover, because frequency drift varies randomly over time, its effect on a spur may differ from its effect on the carrier, as they are measured at different points during the sweep time. These effects can cause spur power measurement to be optimistic, often by several dB.

The above discussion also explains why it is necessary to disable spectrum analyzer averaging when measuring spurs. Spectrum analyzer averaging causes the instrument to average the results of multiple sequentially captured measurements. This mode is useful for increasing the precision of noise measurements. Unfortunately, oscillator frequency drift causes the spur to appear at different frequencies across different measurement sweeps. When these results are averaged, the resulting spur power is reduced as it is averaged with noise, and just as explained above, the averaging may affect the spur power and carrier power differently, due to any randomness in the frequency drift.

Most spectrum analyzers implement the conceptual bandpass filtering operation described above by either downconverting the input signal to a fixed intermediate frequency and passing the result through a fixed-frequency bandpass filter, or digitizing the input signal and performing fast Fourier transform (FFT) based spectrum estimation [17]. Either way, the accuracy of the center frequency of the conceptual bandpass filter depends on the accuracy of the spectrum analyzer's local oscillator frequency, which, in turn, depends on the accuracy of the spectrum analyzer's reference oscillator frequency. Most spectrum analyzers offer the option of using an external reference oscillator signal. When measuring PLL-based oscillator phase error spurs, this feature makes it possible to lock the spectrum analyzer's local oscillator to the PLL-based oscillator's reference oscillator [62]. In this case, the spectrum analyzer's local oscillator has the same frequency drift as the PLL-based oscillator, which avoids the spur-measurement problem described above.

B. Phase Noise Analyzers

The lowest-noise phase noise measurements are made with phase noise analyzers, which are also called signal source analyzers depending on the manufacturer. Phase noise analyzers are essentially extremely sensitive receivers, and are capable of differentiating between the amplitude and phase noise components of an oscillator signal over a limited bandwidth [64].

Two different measurement techniques are commonly used, homodyne phase discrimination and heterodyne frequency discrimination [65]. The phase discrimination method is typically called narrowband-optimized mode, as it is optimized for measuring the close-in phase noise of oscillators with very little frequency drift, such as crystal oscillators and PLL-based oscillators locked to crystal oscillators. The frequency discriminator method is often called wideband-optimized mode, and is better-suited to measure the phase noise of oscillators with relatively high low-frequency phase noise or frequency drift, such as free-running voltage-controlled oscillators. Both of these measurement methods are further improved by correlation techniques, which essentially average out the uncorrelated noise between multiple identical measurement channels inside the instrument [66]. The total added noise from these instruments can be less than 10 femtoseconds of integrated jitter.

The right side of Fig. 10 shows the phase noise plot of the same PLL-based oscillator discussed in the previous subsection, as measured by a Keysight E5052B signal source analyzer. For the reasons given in the previous section, the averaging techniques used by this instrument to improve the noise measurement precision adversely affects its ability to resolve spurs. For example, the spurs at 7.4, 14.7, and 29.5 kHz in the left-side power spectrum are also visible in the phase noise plot, but their powers are not directly readable as they are spread over a range of frequencies.

C. Time-Domain Methods

For serial communication applications, phase error is usually measured in the time domain with high-speed digitizing oscilloscopes, as they can make measurements that would

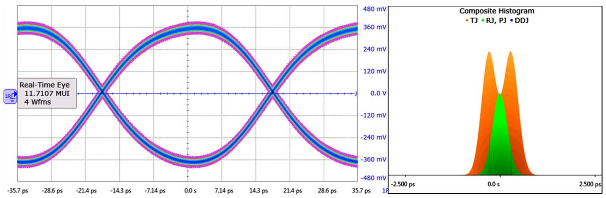

Fig. 11. Measured 14 GHz clock eye diagram and jitter histogram.

be impossible with other instruments [67]. Digitizing oscilloscopes fall into two categories, *realtime* and *sampling* (also called equivalent-time). A realtime oscilloscope operates like a digital version of a traditional analog oscilloscope, in that continuous segments of the input signal are digitized and displayed onscreen. The maximum bandwidth of the input signal is limited by the sample-rate of the ADC, which is currently up to well over 100 GS/s [68].

Realtime oscilloscopes often have internal configurable zero-crossing detectors, so that finite-duration segments of the oscillator signal absolute jitter can be extracted. An absolute jitter sequence can be processed via an FFT, which, by (25), gives an estimate of the oscillator phase error's modified conceptual power spectrum. An absolute jitter sequence can also be viewed as a histogram, and by analyzing the histogram as described in Section V-F, various aspects of the type and magnitude of the jitter can be inferred.

Fig. 11 shows measurement results for a 14 GHz clock signal via a Keysight DSAZ634A realtime oscilloscope. The clock signal was generated by a 28 Gb/s serial transmitter, transmitting a repetitive "01" sequence. The left plot in Fig. 11 is comprised of over 11 million segments of the measured clock waveform shifted and overlaid atop one another, in a plot known as an eye diagram [69]. The right plot in Fig. 11 shows the histogram of the clock signal's zero crossings. It shows the total jitter, random jitter, and two additional terms: periodic jitter (PJ) and data-dependent jitter (DDJ). The complete categorization of jitter used by this instrument is described in [70].

A realtime oscilloscope's unique ability to debug circuit problems in the time domain is particularly valuable in that it makes it possible to compare, edge by edge, the absolute jitter on a clock signal against other signals in the system. For example, in a serial transmitter circuit, coupling may cause transitions in the data being transmitted to disturb the edges of its PLL-based oscillator signal [31]. By viewing the clock and data signals simultaneously on a realtime oscilloscope, it is possible to verify this effect, whereas if the clock signal were viewed on a spectrum analyzer, a spur would be visible with no indication of its phase.

In contrast to the Nyquist-rate sampling in a realtime oscilloscope, a sampling oscilloscope sub-samples its input signal at a much lower rate, usually below 1 MHz [67]. On the other hand, its maximum input frequency is no longer limited by the sampling rate, but by the bandwidth of the analog frontend and sampling circuitry. Furthermore, as the sampling oscilloscope's ADC has more time to process each sample, it can achieve higher voltage resolution. Currently, realtime oscilloscopes have 8-10 bits of resolution, whereas sampling oscilloscopes typically have over 12 bits of resolution [67].

Due to its sub-sampling behavior, a sampling oscilloscope cannot be used to generate phase error power spectra, but it can generate jitter histograms, and its ability to debug time-domain issues is only slightly limited compared to a realtime oscilloscope. For example, a serial transmitter can be configured to transmit a repeating data sequence. Provided the duration of each data sequence transmission is slower than the maximum sub-sampling rate of the sampling oscilloscope, it is possible to sub-sample a very high-speed clock signal relative to the start time of each repeated data sequence. The full clock signal can then be slowly reconstructed in the time domain, hence the "equivalent-time" name for this instrument. The limitation of this method stems from the instrument's sub-sampling behavior, in that any periodic disturbances in the oscillator signal not harmonically-related to the sub-sampling rate can corrupt the measurement.

A drawback of time-domain measurement instruments is caused by their wide analog bandwidths. The analog circuitry present in both types of oscilloscopes prior to their ADCs adds broadband thermal noise to the oscillator signal, and this additive noise corrupts the oscillator signal via the mechanism described in Section III-C. Therefore, even the highest-performance realtime and sampling oscilloscopes have integrated jitter floors well above those of phase noise analyzers.

VIII. Frequency Stability

In the preceding sections, oscillator error was analyzed in the frequency domain via phase error and jitter power spectra, $S_\phi(f)$ and $J(f)$. There exist, however, a broad class of problems in fields such as astronomy [71], satellite navigation [72], and precision metrology [73], [74] in which oscillator error is usually characterized in the time-domain with a quantity called the Allan variance, $\sigma_y^2(\tau)$, which is described in this section. The main utility of the Allan variance and its variants is that they allow for simple characterization of very long-term oscillator behavior (days to years). For example, they are useful in analyzing the behavior of a system where the frequency of a low-precision oscillator is periodically re-calibrated with a high-precision oscillator, a situation which frequently occurs in mobile devices [75]. The development of the time-domain approach to oscillator error was primarily driven by the invention of the atomic clock, whereas the phase error frequency-domain approach was motivated by the development of Doppler radar [76], [77].

This section considers only the phase noise component of phase error. Spurious tones and deterministic modulation in the oscillator phase error have been extensively analyzed within this framework [78], [79], but they require special consideration which is beyond the scope of this tutorial.

A. Frequency Fluctuations

The instantaneous frequency, $\omega(t)$, of the oscillator signal given by (1) is defined as the derivative of the argument of the sine term, i.e., $\omega(t) = \omega_0 + \phi'(t)$ [80]. The *frequency fluctuations* function is defined as $\omega(t) - \omega_0$, i.e., the deviation of the instantaneous frequency from its nominal value. The *fractional frequency fluctuations* function, $y(t)$, is the frequency fluctuations function normalized by ω_0, i.e., $y(t) = \phi'(t)/\omega_0$.

The use of frequency rather than phase as the variable of interest stems from early research characterizing oscillators by counting its cycles [81]. A digital counter can be used to approximately measure sequences of time averages of $y(t)$ of the form

$$\bar{y}_k = \int_{\lambda_k}^{\lambda_k+\tau} \frac{y(t)}{\tau} dt = \frac{\phi(\lambda_k + \tau) - \phi(\lambda_k)}{\omega_0 \tau}, \tag{38}$$

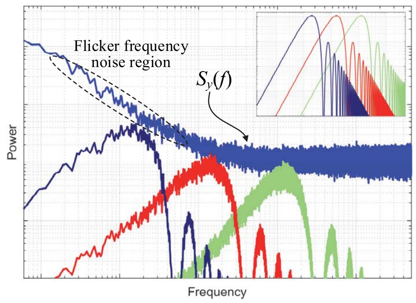

Fig. 12. Frequency-domain signal processing of the Allan variance. Inset shows the equivalent bandpass filter transfer function.

where λ_k denotes the time of the kth measurement, and τ is the duration of the average, also known as the counter gating interval, or stride [82]. It follows from (38) that the $\bar{y}_k$ sequence is generated by sampling the result of a τ-width moving average filter applied to $y(t)$, which is a type of lowpass filter. Hence, $\bar{y}_k$ can be interpreted as samples of the continuous time sequence whose power spectrum is a lowpass-filtered version of $S_y(f)$.

Counters only measure the zero crossings of a clock, so the measurement of (38) incurs quantization noise. Laboratory frequency counters use various methods to improve the measurement accuracy of $\bar{y}_k$, especially for short gating intervals relative to the clock period [83], [84]. However, for the types of problems where Allan variance is a useful metric, the values of τ of interest are typically many orders of magnitude larger than the clock period, so this quantization error can often be ignored. However, there are other non-idealities within the measurement instrument such as drift and noise in its internal time base. These cannot be ignored, and must be considered as part of the uncertainty in the measurement [85].

B. Allan Variance

A focus of early research on precision oscillators was characterizing flicker frequency noise, which manifests as the low-frequency region of $S_y(f)$ that decays at -10 dB/decade [86], shown in Fig. 12. Unlike other types of noise, such as white frequency noise (flat in $S_y(f)$) or white phase noise ($+20$ dB/decade in $S_y(f)$), measurement uncertainty due to flicker frequency noise cannot be removed by averaging, which leads to an accuracy limitation for measurements that depend on averaging over time intervals. Its measurement also defies characterization by the variance, because the integral over positive frequencies including 0 of a lowpass-filtered $S_y(f)$ is unbounded, due to the divergence of the integral of $1/f$ at $f = 0$ [87].

The divergence of the flicker noise model at low frequencies is shown to be accurate over extremely long observation intervals, well beyond practical measurement durations [88], [89]. One method to bound the integral of $S_y(f)$ is to use arguments similar to those in Section V-C, where a restricted jitter density integration interval is used to bound the calculation of absolute jitter. However, there is no unique interval of frequency noise integration for the characterization of oscillators that would be applicable in all scenarios [86]. Therefore, a different metric is needed that is both generally applicable and practically measurable.

The problem with computing the variance of $\bar{y}_k$ occurs because the integral of $S_y(f)$ diverges at $f = 0$, so the most straightforward solution is to take the variance of the first difference of $\bar{y}_k$. The Allan variance, or two-sample variance, is thus defined as

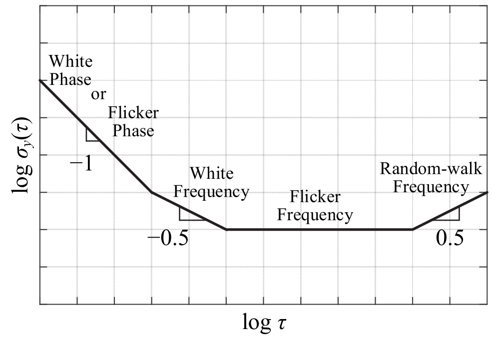

Fig. 13. The σ-τ plot, annotated with the slopes of different regions indicating the different noise type contributions.

$$\sigma_y^2(\tau) = \frac{1}{2} E\left[(\bar{y}_2 - \bar{y}_1)^2\right]. \tag{39}$$

Hence, the value $\sigma_y^2(\tau)$ is equal to the variance of the difference between two adjacent τ-duration average frequency measurements, normalized to the nominal frequency. This may seem contrived, but it has an intuitive interpretation in relation to timekeeping, which is explained in the next sub-section.

Taking the first difference in the time domain is equivalent to a highpass filter in the frequency domain, so the combined response of the first difference and the τ-width moving average filter is equivalent to a bandpass filter applied to $S_y(f)$ [90]. The behavior of this filter is shown in Fig. 12, for three different values of τ. The bandwidth and center frequency of this filter is proportional to $1/\tau$, so it has a constant logarithmic bandwidth. As flicker noise has constant integrated power over logarithmic bandwidths, $\sigma_y^2(\tau)$ is constant when the value of τ places the filter's center frequency in regions of $S_y(f)$ dominated by flicker frequency noise.

The square root of the Allan variance, $\sigma_y(\tau)$, is called the Allan deviation. This is typically plotted on a log-log scale versus τ, in what is called a σ-τ plot, shown in Fig. 13. Like a phase noise plot, different regions of the σ-τ plot are dominated by different types of noise. A typical σ-τ plot shows, for increasing τ, a downward sloping region, a flat region, and an upward sloping region. As indicated, all but the upward-sloping region (random-walk frequency) can be attributed to the intrinsic noise processes encountered in IC design. The random-walk frequency region tends to be dominated by external effects, such as random temperature fluctuations.

The value of τ required to measure the upward sloping region depends on the type of oscillator. For example, precision temperature-compensated crystal oscillators (TCXO) can show this region after τ's of minutes or hours, whereas the Rubidium oscillators in GPS satellite clocks require a τ of several days [91]! It is important to note that, as τ (and hence measurement duration) becomes large, it becomes increasingly difficult to isolate the performance of the oscillator under test from the measurement environment. In addition to temperature fluctuations, issues such as physical vibration due to office foot traffic, power surges causing equipment failure, or failure of the oscillator itself can all be considered random events that tend to manifest during long-term measurements.

C. Interpretation and Comparison With Phase Noise

Mobile devices such as cellular phones and GPS receivers contain circuits called realtime clocks (RTC). An RTC contains an internal oscillator tuned to a known nominal frequency. By counting the oscillator's cycles, a sense of time is kept. The frequency stability of the mobile device RTC oscillator is typically poor, so the RTC is periodically recalibrated using the base station time as described below.

Every τ seconds, the base station transmits its time to the mobile device. The mobile device updates its internal time to the received time, but also computes the error between the RTC's measurement of the τ-second interval and the actual interval as defined by the base station. The mobile device then uses this measured error to correct the RTC's prediction of the subsequent τ-second interval.

In this manner, the mobile device's time error from the base station time is reset to 0 every τ seconds, but its time error grows as time elapses from the most recent reset up to the next reset. In this scenario, the standard deviation of the mobile device's time error immediately before each recalibration is closely approximated by $\tau\sigma_y(\tau)$. While this idealized example ignores practical challenges faced by such systems, e.g., temperature fluctuations between recalibrations, it illustrates how the Allen variance is a useful tool to characterize the precision of time measurements based on the performance of the underlying oscillator used to measure the time interval.

Using a σ-τ plot in the above example, if the chosen τ falls in the -1 sloped region, then $\tau\sigma_y(\tau)$ is a constant with respect to τ. If τ falls in the flat region, then $\tau\sigma_y(\tau)$ grows linearly with τ, and if τ falls in the rising region, then $\tau\sigma_y(\tau)$ grows at a much higher rate. From the σ-τ plot, it is also easy to calculate the maximum allowable time between recalibrations given a maximum time error standard deviation requirement. In contrast, such information is not evident from a plot of either $S_\phi(f)$ or $J(f)$.

The time between the negative-slope region and the flat region in a σ-τ plot is of particular interest. This time, called τ_0, corresponds to the time after which averaging a frequency measurement longer than τ_0 seconds does not increase the measurement precision. Thus, the *flicker noise floor* of frequency stability is a critical specification of precision oscillators for metrology applications. Recent publications have reported reaching a frequency stability of below 1.5×10^{-18} with just 30 minutes of averaging [92]. For a point of scale, this is equivalent to measuring the diameter of the earth to a precision of less than the width of a hydrogen atom [93]!

The Allen variance is a convenient metric not only because it is simple to estimate through measurement, but also because its estimate converges for all the commonly encountered power-law noise types. There exist several other types of variances that offer advantages over the Allen variance. For example, the modified Allan variance (MVAR) introduces a τ-dependent scale factor to the Allan variance that gives the white phase and flicker phase noise regions different slopes on the σ-τ plot [94]. The total variance (TVAR) is another Allan variance like metric that has better convergence of the estimator sequences for long τ in the presence of effects such as linear frequency drift [95]. Another metric is the Hadamard variance which has a sharp frequency selectivity and can estimate very high order frequency noise, namely flicker walk frequency (-50 dB/dec phase noise) and flicker run frequency (-60 dB/dec phase noise) noise [96]. This was needed to analyze the Rubidium oscillators in GPS satellites that exhibit complex frequency drift behavior [91].

The approach to selecting a type of variance is similar to the choice of window type for power spectrum estimation. The different types emphasize (or obscure) different aspects of the signal, and depending on the features of interest one may be more appropriate than another. The references [90], [97], [98] provide more details on many commonly-used types of variance.

IX. Conclusion

Oscillators and the clock signals they generate are ubiquitous in modern electronic systems. Due to their wide range of applications, different disciplines have each developed their own metrics and methods by which to characterize oscillator non-idealities, each uniquely suited to their particular application. Wireless systems are predominantly concerned with the oscillator's phase error power spectrum. Digital clocked circuits, and wireline communication systems are more frequently analyzed using time-domain jitter calculations. Metrology and physics applications frequently employ the Allan variance. Although these metrics, and the instruments used to measure them, appear very different, as explained in this tutorial they are all just different perspectives on oscillator phase error.

Acknowledgements

The authors wish to thank Walid Ali-Ahmad, Enrique Alvarez-Fontecilla, Ajay Balankutty, Brian Buell, Matthew Felmlee, Raghavendra Haresamudram, Sassan Tabatabaei, Yiwu Tang and the anonymous reviewers for their many excellent comments and suggestions which helped us improve this paper.

References

[1] N. Da Dalt and A. Sheikholeslami, *Understanding Jitter and Phase Noise a Circuits and Systems Perspective*. Cambridge, U.K.: Cambridge Univ. Press, 2018.

[2] E. Rubiola, *Phase Noise and Frequency Stability in Oscillators*. Cambridge, U.K.: Cambridge Univ. Press, 2009.

[3] A. Hajimiri and T. H. Lee, "A general theory of phase noise in electrical oscillators," *IEEE J. Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, Feb. 1998.

[4] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790–804, Jun. 1999.

[5] T. H. Lee and A. Hajimiri, "Oscillator phase noise: A tutorial," *IEEE J. Solid-State Circuits*, vol. 35, no. 3, pp. 326–336, Mar. 2000.

[6] A. Demir, A. Mehrotra, and J. Roychowdhury, "Phase noise in oscillators: A unifying theory and numerical methods for characterization," *IEEE Trans. Circuits Syst. I, Fundam. Theory Appl.*, vol. 47, no. 5, pp. 655–674, May 2000.

[7] K. A. Kouznetsov and R. G. Meyer, "Phase noise in LC oscillators," *IEEE J. Solid-State Circuits*, vol. 35, no. 8, pp. 1244–1248, Aug. 2000.

[8] A. Mazzanti and P. Andreani, "Class-C harmonic CMOS VCOs, with a general result on phase noise," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2716–2729, Dec. 2008.

[9] D. Murphy, J. J. Rael, and A. A. Abidi, "Phase noise in LC oscillators: A phasor-based analysis of a general result and of loaded Q," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 57, no. 6, pp. 1187–1203, Jun. 2010.

[10] *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*, IEEE Standard 1139-2008, Feb. 2009.

[11] R. S. Carson, *Radio Communications Concepts: Analog*. New York, NY, USA: Wiley, 1990.

[12] J. A. Weldon *et al.*, "A 1.75-GHz highly integrated narrow-band CMOS transmitter with harmonic-rejection mixers," *IEEE J. Solid-State Circuits*, vol. 26, no. 12, pp. 2003–2015, Dec. 2001.

[13] R. A. Fessenden, "Method of signalling," U.S. Patent 1 050 728, Jan. 14, 1913.

[14] E. W. Herold, "The operation of frequency converters and mixers for superheterodyne reception," *Proc. Inst. Radio Eng.*, vol. 20, no. 2, pp. 84–103, Feb. 1942.

[15] P. Sivonen, J. Tervaluoto, N. Mikkola, and A. Parssinen, "A 1.2-V RF front-end with on-chip VCO for PCS 1900 direct conversion receiver in 0.13-μm CMOS," *IEEE J. Solid-State Circuits*, vol. 41, no. 2, pp. 384–394, Feb. 2006.

[16] H. Darabi, *Radio Frequency Integrated Circuits and Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2015.

[17] "Spectrum analysis basics," Keysight Technol., Santa Rosa, CA, USA, Appl. Note 150, 2016. [Online]. Available: http://literature.cdn.keysight.com/litweb/pdf/5952-0292.pdf

[18] K. J. Wang, A. Swaminathan, and I. Galton, "Spurious tone suppression techniques applied to a wide-bandwidth 2.4 GHz fractional-N PLL," *IEEE J. Solid-State Circuits*, vol. 43, no. 12, pp. 2787–2797, Dec. 2008.

[19] M. Zanuso *et al.*, "A wideband 3.6 GHz digital $\Delta\Sigma$ fractional-N PLL with phase interpolation divider and digital spur cancellation," *IEEE J. Solid-State Circuits*, vol. 46, no. 3, pp. 627–638, Mar. 2011.

[20] X. Gao, E. A. M. Klumperink, G. Socci, M. Bohsali, and B. Nauta, "Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector," *IEEE J. Solid-State Circuits*, vol. 45, no. 9, pp. 1809–1821, Sep. 2010.

[21] C.-W. Yao *et al.*, "A 14-nm 0.14-ps$_{rms}$ fractional-N digital PLL with a 0.2-ps resolution ADC-assisted coarse/fine-conversion chopping TDC and TDC nonlinearity calibration," *IEEE J. Solid-State Circuits*, vol. 52, no. 12, pp. 3446–3457, Dec. 2017.

[22] P. Kartaschoff, "Frequency and time," in *Monographs in Physical Measurement*. New York, NY, USA: Academic, 1978.

[23] Y. Frans *et al.*, "A 56-Gb/s PAM4 wireline transceiver using a 32-way time-interleaved SAR ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits*, vol. 52, no. 4, pp. 1101–1110, Apr. 2017.

[24] *Ultrascale+ FPGA Selection Guide*, Xilinx Corp., San Jose, CA, USA, 2017. [Online]. Available: https://www.xilinx.com/support/documentation/selection-guides/ultrascale-plus-fpga-product-selection-guide.pdf

[25] P. Kinget, "Integrated GHz voltage controlled oscillators," in *Analog Circuit Design*. Boston, MA, USA: Springer, 1999, pp. 353–381.

[26] D. Ham and A. Hajimiri, "Concepts and methods in optimization of integrated LC VCOs," *IEEE J. Solid-State Circuits*, vol. 36, no. 6, pp. 896–909, Jun. 2001.

[27] J. A. McNeill, "Jitter in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 32, no. 6, pp. 870–879, Jun. 1997.

[28] P. F. J. Geraedts, E. van Tuijl, E. A. M. Klumperink, G. J. M. Wienk, and B. Nauta, "A 90 μW 12 MHz relaxation oscillator with a -162 dB FOM," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2008, pp. 348–618.

[29] *Evolved Universal Terrestrial Radio Access (E-UTRA); User Equipment (UE) Radio Transmission and Reception*, document TR 36.101, 3GPP, Jun. 2018. [Online]. Available: http://www.3gpp.org

[30] D. B. Leeson, "A simple model of feedback oscillator noise spectrum," *Proc. IEEE*, vol. 54, no. 2, pp. 329–330, Feb. 1966.

[31] B. Casper and F. O'Mahony, "Clocking analysis, implementation and measurement techniques for high-speed data links—A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 1, pp. 17–39, Jan. 2009.

[32] *GR-253-CORE*, Telcordia Technol., Piscataway, NJ, USA, Sep. 2000, no. 3.

[33] R. J. Van de Plassche, *CMOS Integrated Analog-to-Digital and Digital-to-Analog Converters*, vol. 742. New York, NY, USA: Springer, 2013.

[34] W. Kester, "Converting oscillator phase noise to jitter, tutorial MT-008," Analog Devices, Norwood, MA, USA, Oct. 2008. [Online]. Available: http://www.analog.com/media/en/training-seminars/tutorials/MT-008.pdf

[35] Q. Wang, H. Shibata, A. Liscidini, and A. C. Carusone, "A digital filtering ADC with programmable blocker cancellation for wireless receivers," *IEEE J. Solid-State Circuits*, vol. 53, no. 3, pp. 681–691, Mar. 2018.

[36] J. Wu *et al.*, "A 2.7 mW/channel 48–1000 MHz direct sampling full-band cable receiver," *IEEE J. Solid-State Circuits*, vol. 51, no. 4, pp. 845–859, Apr. 2016.

[37] R. Stephens, "Jitter analysis: The dual-dirac model, RJ/DJ, and Q-scale," Agilent Technol., Santa Clara, CA, USA, Appl. Note, 2004. [Online]. Available: https://www.keysight.com/upload/cmc_upload/All/dualdirac1.pdf

[38] N. M. Nguyen and R. G. Meyer, "Start-up and frequency stability in high-frequency oscillators," *IEEE J. Solid-State Circuits*, vol. 27, no. 5, pp. 810–820, May 1992.

[39] T. J. Aprille and T. N. Trick, "Steady-state analysis of nonlinear circuits with periodic inputs," *Proc. IEEE*, vol. 60, no. 1, pp. 108–114, Jan. 1972.

[40] K. S. Kundert and A. Sangiovanni-Vincentelli, "Simulation of nonlinear circuits in the frequency domain," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. CAD-5, no. 4, pp. 521–535, Oct. 1986.

[41] M. Nakhla and J. Vlach, "A piecewise harmonic balance technique for determination of periodic response of nonlinear systems," *IEEE Trans. Circuits Syst.*, vol. CAS-23, no. 2, pp. 85–91, Feb. 1976.

[42] K. S. Kundert, "Introduction to RF simulation and its application," *IEEE J. Solid-State Circuits*, vol. 34, no. 9, pp. 1298–1319, Sep. 1999.

[43] M. Okumura, H. Tanimoto, T. Itakura, and T. Sugawara, "Numerical noise analysis for nonlinear circuits with a periodic large signal excitation including cyclostationary noise sources," *IEEE Trans. Circuits Syst. I. Fundam. Theory Appl.*, vol. 40, no. 9, pp. 581–590, Sep. 1993.

[44] J. Roychowdhury, D. Long, and P. Feldmann, "Cyclostationary noise analysis of large RF circuits with multitone excitations," *IEEE J. Solid-State Circuits*, vol. 33, no. 3, pp. 324–336, Mar. 1998.

[45] K. Mayaram, D. C. Lee, S. Moinian, D. A. Rich, and J. Roychowdhury, "Computer-aided circuit analysis tools for RFIC simulation: Algorithms, features, and limitations," *IEEE Trans. Circuits Syst. II. Analog Digit. Signal Process.*, vol. 47, no. 4, pp. 274–286, Apr. 2000.

[46] B. Wan and X. Wang, "Overview of commercially-available analog/RF simulation engines and design environment," in *Proc. 12th IEEE Int. Conf. Solid-State Integr. Circuit Technol.*, Oct. 2014, pp. 1–4.

[47] *Spectre Circuit Simulator and Accelerated Parallel Simulator RF Analysis User Guide, Ver. 16.1*, Cadence Des. Syst., San Jose, CA, USA, 2017.

[48] K. Kundert, "Modeling and simulation of jitter in phase-locked loops," in *Analog Circuit Design: RF Analog-to-Digital Converters; Sensor and Actuator Interfaces; Low-Noise Oscillators, PLLs and Synthesizers*. New York, NY, USA: Springer, 1997.

[49] R. Achar and M. S. Nakhla, "Simulation of high-speed interconnects," *Proc. IEEE*, vol. 89, no. 5, pp. 693–728, May 2001.

[50] E.-P. Li, E.-X. Liu, L.-W. Li, and M.-S. Leong, "A coupled efficient and systematic full-wave time-domain macromodeling and circuit simulation method for signal integrity analysis of high-speed interconnects," *IEEE Trans. Adv. Packag.*, vol. 27, no. 1, pp. 213–223, Feb. 2004.

[51] S. Kapur and D. E. Long, "EMX user's manual, ver. 4.6," Integrand Softw., Berkeley Heights, NJ, USA, Tech. Rep., 2014.

[52] F. O'Mahony, C. P. Yue, M. A. Horowitz, and S. S. Wong, "Mahony, "A 10-GHz global clock distribution using coupled standing-wave oscillators," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1813–1820, Nov. 2003.

[53] S. C. Chan *et al.*, "A resonant global clock distribution for the cell broadband engine processor," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 64–72, Jan. 2009.

[54] Z. Bai, X. Zhou, R. D. Mason, and G. Allan, "Low-phase noise clock distribution network using rotary traveling-wave oscillators and built-in self-test phase tuning technique," *IEEE Trans. Circuits Syst., II, Exp. Briefs*, vol. 62, no. 1, pp. 41–45, Jan. 2015.

[55] M. H. Perrott, M. D. Trott, and C. G. Sodini, "A modeling approach for $\Sigma\Delta$ fractional-N frequency synthesizers allowing straightforward noise analysis," *IEEE J. Solid-State Circuits*, vol. 37, no. 8, pp. 1028–1038, Aug. 2002.

[56] J. Zhuang, Q. Du, and T. Kwasniewski, "Event-driven modeling and simulation of an digital PLL," in *Proc. IEEE Int. Behav. Modeling Simulation Workshop*, Sep. 2006, pp. 67–72.

[57] I. L. Syllaios, R. B. Staszewski, and P. T. Balsara, "Time-domain modeling of an RF all-digital PLL," *IEEE Trans. Circuits Syst., II, Exp. Briefs*, vol. 55, no. 6, pp. 601–605, Jun. 2008.

[58] R. B. Staszewski, C. Fernando, and P. T. Balsara, "Event-driven simulation and modeling of phase noise of an RF oscillator," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 52, no. 4, pp. 723–733, Apr. 2005.

[59] C.-C. Kuo, M.-J. Lee, C.-N. Liu, and C.-J. Huang, "Fast statistical analysis of process variation effects using accurate PLL behavioral models," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 56, no. 6, pp. 1160–1172, Jun. 2009.

[60] K. Kundert and O. Zinke, *The Designer's Guide to Verilog-AMS*. New York, NY, USA: Springer, 2006.

[61] T. Decker and R. Temple, "Choosing a phase noise measurement technique," in *Proc. RF Microw. Meas. Symp. Exhib.*, 1999, pp. 1–64.

[62] "Spectrum and signal analyzer measurements and noise," Keysight Technol., Santa Rosa, CA, USA, Appl. Note 1303, 2017. [Online]. Available: https://literature.cdn.keysight.com/litweb/pdf/5966-4008E.pdf

[63] J. Wolf, "Phase noise measurements with spectrum analyzers of the FSE family," Rohde & Schwarz, Munich, Germany, Appl. Note 1EPAN 16E, 1995. [Online]. Available: https://cdn.rohde-schwarz.com/pws/dl_downloads/dl_application/application_notes/1epan16/1epan16e.pdf

[64] G. Feldhaus and A. Roth, "A 1 MHz to 50 GHz direct down-conversion phase noise analyzer with cross-correlation," in *Proc. Eur. Freq. Time Forum*, Apr. 2016, pp. 1–4.

[65] "E5052B signal source analyzer—Advanced phase noise and transient measurement techniques," Keysight Technol., Santa Rosa, CA, USA, Appl. Note 5989-7273EN, 2017. [Online]. Available: http://literature.cdn.keysight.com/litweb/pdf/5989-7273EN.pdf

[66] W. F. Walls, "Cross-correlation phase noise measurements," in *Proc. 46th IEEE Freq. Control Symp.*, May 1992, pp. 257–261.

[67] B. Fetz and G. LeCheminant, "Measuring Mult-Gb/s signals…'what type of oscilloscope should i use?"' Agilent Technologies, Santa Clara, CA, USA, Appl. Note, 2006. [Online]. Available: https://www.keysight.com/upload/cmc_upload/All/Which_scope_to_use06.pdf

[68] *Infiniium Z-Series Oscilloscopes, Data Sheet*, Keysight Technol., Santa Rosa, CA, USA, 2017. [Online]. Available: https://literature.cdn.keysight.com/litweb/pdf/5991-3868EN.pdf

[69] G. Breed, "Analyzing signals using the eye diagram," *High Freq. Electron.*, vol. 4, no. 11, pp. 50–53, 2005.

[70] *Clock Jitter Analysis With Femto-Second Resolution, Tutorial*, Keysight Technol., Santa Rosa, CA, USA, 2008. [Online]. Available: https://www.keysight.com/upload/cmc_upload/All/Clock_Jitter_Analysis_2008.pdf

[71] R. A. Hulse and J. H. Taylor, "Discovery of a pulsar in a binary system," *Astrophys. J.*, vol. 195, pp. 51–53, Jan. 1975.

[72] X. Li *et al.*, "Accuracy and reliability of multi-GNSS real-time precise positioning: GPS, GLONASS, BeiDou, and Galileo," *J. Geodesy*, vol. 89, no. 6, pp. 607–635, 2015.

[73] P. Kartaschoff and J. A. Barnes, "Standard time and frequency generation," *Proc. IEEE*, vol. 60, no. 5, pp. 493–501, May 1972.

[74] S. P. Benz and C. A. Hamilton, "Application of the Josephson effect to voltage metrology," *Proc. IEEE*, vol. 92, no. 10, pp. 1617–1629, Oct. 2004.

[75] S. A. Gronemeyer, "Calibrated real time clock for acquisition of GPS signals during low power operation," U.S. Patent 6 662 107, Dec. 9, 2003.

[76] M. A. Lombardi, T. P. Heavner, and S. R. Jefferts, "NIST primary frequency standards and the realization of the SI second," *NCSLI Meas.*, vol. 2, no. 4, pp. 74–89, 2007.

[77] *Proceeding of the IEEE-NASA Symposium on the Definition and Measurement of Short-Term Frequency Stability*, Goddard Space Flight Center, Greenbelt, MD, USA, 1964.

[78] C. A. Greenhall, "Spectral ambiguity of Allan variance," *IEEE Trans. Instrum. Meas.*, vol. 47, no. 3, pp. 623–627, Jun. 1998.

[79] K. L. Senior, J. R. Ray, and R. L. Beard, "Characterization of periodic variations in the GPS satellite clocks," *GPS Solutions*, vol. 12, no. 3, pp. 211–225, 2008.

[80] L. S. Cutler and C. L. Searle, "Some aspects of the theory and measurement of frequency fluctuations in frequency standards," *Proc. IEEE*, vol. 54, no. 2, pp. 136–154, Feb. 1966.

[81] W. A. Marrison, "The crystal clock," *Proc. Nat. Acad. Sci. USA*, vol. 16, no. 7, pp. 496–507, 1930.

[82] J. A. Barnes *et al.*, "Characterization of frequency stability," *IEEE Trans. Instrum. Meas.*, vol. IM-20, no. 2, pp. 105–120, May 1971.

[83] J. Kalisz, "Review of methods for time interval measurements with picosecond resolution," *Metrologia*, vol. 41, no. 1, pp. 17–32, 2004.

[84] S. T. Dawkins, J. J. McFerran, and A. N. Luiten, "Considerations on the measurement of the stability of oscillators with frequency counters," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 54, no. 5, pp. 918–925, May 2007.

[85] J. E. Gray and D. W. Allan, "A method for estimating the frequency stability of an individual oscillator," in *Proc. 28th Annu. Symp. Freq. Control*, May 1974, pp. 243–246.

[86] D. W. Allan, "Statistics of atomic frequency standards," *Proc. IEEE*, vol. 54, no. 2, pp. 221–230, Feb. 1966.

[87] J. A. Barnes, "Atomic timekeeping and the statistics of precision signal generators," *Proc. IEEE*, vol. 54, no. 2, pp. 207–220, Feb. 1966.

[88] W. H. Press, "Flicker noises in astronomy and elsewhere," *Comments Astrophys.*, vol. 7, no. 4, pp. 103–119, 1978.

[89] E. Milotti. (2002). "1/f noise: A pedagogical review." [Online]. Available: https://arxiv.org/abs/physics/0204033

[90] D. W. Allan, M. A. Weiss, and J. L. Jespersen, "A frequency-domain view of time-domain characterization of clocks and time and frequency distribution systems," in *Proc. 45th Annu. Symp. Freq. Control*, May 1991, pp. 667–678.

[91] S. T. Hutsell *et al.*, "Operational use of the Hadamard variance in GPS," Naval Observatory Alternative Master Clock, Schreiver Air Force Base, Colorado Springs, CO, USA, Tech. Rep., 1996.

[92] M. Schioppo *et al.*, "Ultrastable optical clock with two cold-atom ensembles," *Nature Photon.*, vol. 11, no. 1, pp. 48–52, 2017.

[93] N. Hinkley *et al.*, "An atomic clock with 10^{-18} instability," *Science*, vol. 341, no. 6151, pp. 1215–1218, 2013.

[94] D. W. Allan and J. A. Barnes, "A modified 'allan variance' with increased oscillator characterization ability," in *Proc. 35th Annu. Freq. Control Symp.*, May 1981, pp. 470–475.

[95] D. A. Howe, "The total deviation approach to long-term characterization of frequency stability," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 47, no. 5, pp. 1102–1110, Sep. 2000.

[96] R. A. Baugh, "Frequency modulation analysis with the Hadamard variance," in *Proc. IEEE 25th Annu. Symp. Freq. Control*, Apr. 1971, pp. 222–225.

[97] W. Riley, "Handbook of frequency stability analysis," NIST, Gaithersburg, MD, USA, Tech. Rep. 1065, 2008.

[98] J. Rutman, "Characterization of phase and frequency instabilities in precision frequency sources: Fifteen years of progress," *Proc. IEEE*, vol. 66, no. 9, pp. 1048–1075, Sep. 1978.

Ian Galton (F'11) received the Sc.B. degree from Brown University in 1984, and the M.S. and Ph.D. degrees from the California Institute of Technology in 1989 and 1992, respectively, all in electrical engineering. Since 1996, he has been a Professor of electrical engineering with the University of California, San Diego, where he teaches and conducts research in the field of mixed-signal integrated circuits and systems for communications. His research involves the invention, analysis, and integrated circuit implementation of critical communication system blocks such as data converters and phase-locked loops.

Colin Weltin-Wu (M'15) received the B.S. degree in math and the M.E. degree in electrical engineering from the Massachusetts Institute of Technology and the Ph.D. degree from Columbia University. He spent 15 months, starting in 2014, collaborating with I. Galton at the University of California at San Diego. He is currently involved in high-performance phase-locked loops at Analog Devices, Inc.