Hidden Complexities of PLLs Are Revealed

By John G. Maneatis
True Circuits Inc.
E-mail: maneatis@truecircuits.com

To achieve the timing goals needed in their ASIC designs, many engineers include phase-locked loops. PLLs have a number of desirable properties that include the ability to multiply clock frequencies, correct clock duty cycles and cancel out clock distribution delays. These properties allow designers to use inexpensive, low-frequency crystals as their off-chip clock source and subsequently multiply the frequencies on-chip to produce any number of higher-frequency internal clock signals. They also let designers control setup-and-hold time windows and clock-to-output delays at the chip interface by aligning the windows to the edges of the chipís clock source.

While seemingly simple in structure and function, phase-locked loops are filled with hidden complexity that can cause trouble for even the best circuit designers. The design of PLLs in modern ASIC processes is becoming increasingly difficult due to the limited supply-voltage headroom over the thresholds of core thin-oxide devices. Such devices are often required to meet the operating-frequency targets and to maintain supply-voltage flexibility. However, the reduced supply-voltage headroom will adversely affect the PLLís noise performance.

ASIC designers must be aware of potential performance pitfalls for PLL designs and know how to properly characterize the performance of PLLs and detect issues in performance results that can affect chip timing budgets. With this understanding, they will be better able to successfully decide which PLLs to use and how to best integrate them into their chip designs.

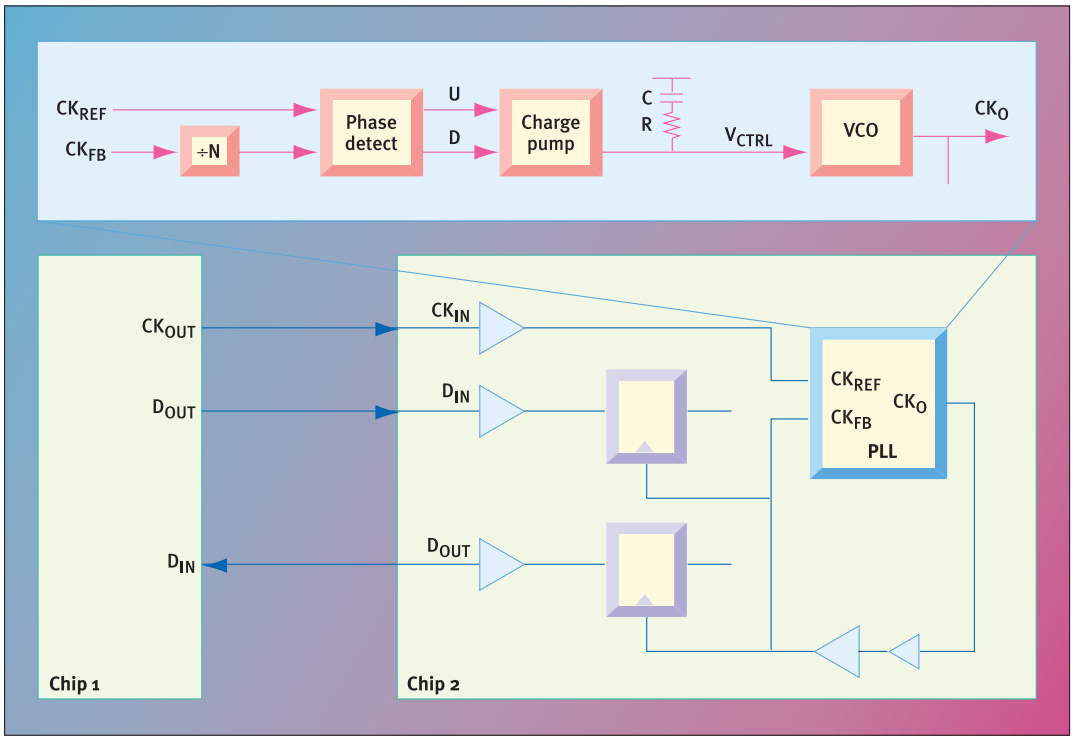

Figure 1: The high-level structure of a PLL is composed of a phase detector, a charge pump, a loop filter and a voltage-controlled oscillator (VCO).

Structure and operation

To truly appreciate the performance issues within a PLL, one must first understand its structure and how it works. The high-level structure of a PLL is seemingly straightforward. It is composed of a phase detector, a charge pump, a loop filter and a voltage-controlled oscillator. Immediately after the PLL circuit is activated, it is in an ìunlockedî state, since the divided VCO output frequency is unrelated to that of the reference.

However, the negative feedback in the loop adjusts the VCO output frequency by integrating the phase error between the rising clock edges of the periodic reference input and the divided VCO output. The integrated phase error causes the divided VCO output frequency to approach that of the reference. As the PLL reaches ìlock,î the phase error detected by the phase detector approaches zero, because the divided VCO output frequency and phase align with that of the reference. Since the phase detector only compares against the divided VCO output, the PLL output will be N times higher in frequency than the reference and feedback inputs, thus allowing the PLL to perform frequency multiplication.

In addition, if the clock distribution is added to the feedback path, the PLL will align the distributed

clock signal to the reference, effectively eliminating the clock-distribution delay.

The blocks within a PLL can consist of varying amounts of analog and digital circuitry, even to the extreme of being completely digital. However, whether composed of digital or analog circuitry, PLLs perform the analog functions of generating and aligning the phase of clock signals. Like analog blocks, they are susceptible to analog issues such as noise, which is commonplace and unavoidable in the harsh mixed-signal environment of todayís ASICs. If a PLL does not respond well to noise, it can introduce time-varying offsets in the phase of the output clock from its ideal value.

These time-varying offsets in the output-clock phase are commonly referred to as jitter. Jitter can have disastrous effects on internal timing paths by causing setup-time violations, and can affect off-chip interfaces by causing setup-and-hold-time violations that lead to data transmission errors. While other performance issues, such as instability, inadequate frequency range, locking problems and static-phase offset, can also affect PLL designs, output jitter is one of the most significant issues and one of the most difficult to adequately address in the design of a PLL.

The supply and substrate noise generated by the on-chip and off-chip sources is highly data-dependent and can have a wide range of frequency components that include low frequencies. Substrate noise tends not to have as large low-frequency components as supply noise, since no significant dc drops develop between the substrate and the supply voltages. Under worst-case conditions, PLLs may experience supply and substrate noise levels as large as 10 percent and 5 percent of the nominal supply voltage, respectively.

The actual level of substrate noise depends on the nature of the substrate used in the IC process. To reduce the risk of latchup, many IC processes use lightly doped epitaxy on the same-type heavily doped substrate. These substrates tend to transmit substrate noise across large distances on the chip, which makes the noise difficult to eliminate through guard rings and additional substrate taps.

Supply and substrate noise affect a PLL by causing frequency shifts in the VCO output, which lead to phase shifts that accumulate for many cycles until the noise pulses subside or the PLL can correct the frequency error at a rate limited by its loop bandwidth. Because the phase error can accumulate over many cycles, the worst-case output jitter will usually be caused by a low-frequency square-wave noise signal. If a PLL is underdamped, noise frequencies near the loop bandwidth can be even more significant. In addition, a PLL can amplify reference input jitter at frequencies near the loop bandwidth, especially if it is underdamped.

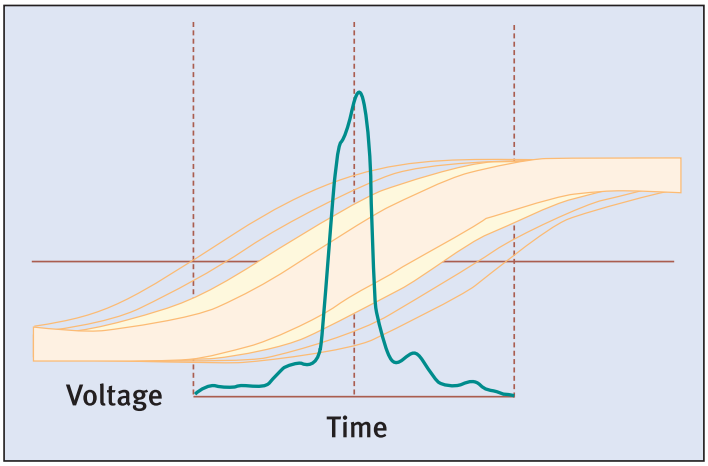

Figure 2: The measured output jitter histogram shows outlying edges that are low-frequency noise events with fast rise times.

Output-jitter types

Output jitter can be measured in a number of ways, relative to absolute time, another signal or the output clock itself. The first is commonly referred to as absolute jitter or long-term jitter, the second as tracking jitter or input-to-output jitter, when the other signal is the reference signal. If the reference signal is perfectly periodic such that it has no jitter, absolute jitter and tracking jitter for the output signal are equivalent. The third measurement, relative to the output clock, is often called periodic, or cycle-to-cycle, jitter. Cycle-to-cycle jitter can be measured as the time-varying deviations in the period of single clock cycles, or in the width of several clock cycles (referred to as cycle-to-Nth-cycle jitter).

Output jitter can also be reported as either rms or peak-to-peak jitter. Rms jitter is interesting only to applications that degrade gracefully when presented with a small number of edges with large time displacements well beyond the rms specification. Such applications can include video and audio signal generation. Peak-to-peak jitter specification is typically the only useful one for jitter in synchronous digital systems, since most setup- or hold-time failures are catastrophic to the operation of a chip.

The significance of a particular measurement of jitter also depends on the application for the PLL. Cycle-to-cycle jitter is usually important in all PLL applications. Tracking jitter is important in applications in which the PLL output clock is used to drive or sample data to or from another clock domain, typically in interface applications. Long-term jitter is sometimes important in applications involving clock multiplication.

Because the phase error in PLLs accumulates over many cycles, the tracking jitter for PLLs that results from supply and substrate noise can be several times larger than the cycle-to-cycle jitter. However, due to the added jitter from on-chip clock-distribution networks, which typically have poor supply- and substrate-noise rejection, the observable difference may be less than a factor of two for well-designed PLLs.

Cycle-to-cycle jitter can also be increased in frequency-multiplying PLLs by periodic disturbances in the period of the first one or two output cycles at the beginning of each reference cycle. This disturbance is caused by systematic residual error from the phase detector.

Measuring jitter correctly can be challenging. Given that PLLs must operate in a noisy mixed-signal environment, it is important to measure them within an equivalently noisy environment. Measuring PLLs in a quiet, low-noise environment can yield optimistic and misleading jitter results. Also, when artificial noise is applied to the PLLís analog supplies, care must be taken to capture the worst-case noise-frequency content. For long-term jitter and tracking jitter, this worst-case noise signal is a square wave at or below the loop-bandwidth frequency, which is usually about a factor of 20 below the minimum PLL operating frequency. For cycle-to-cycle jitter, the worst-case noise signal is a square wave with edge transition times less than a PLL output-clock period and a frequency less than that of the reference. The frequency of this noise signal can be above the

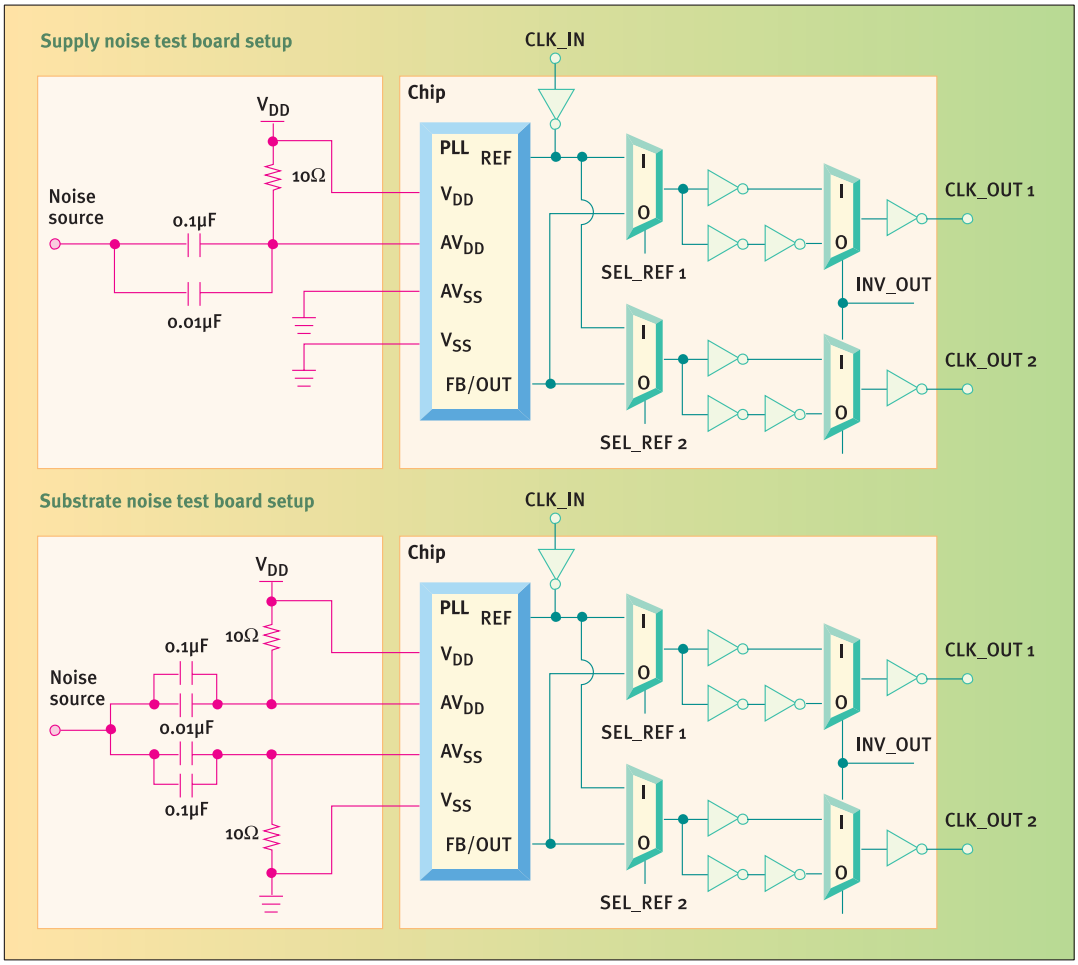

Figure 3: Simple charges to the power supply network on a test board make it possible to characterize the response of a PLL to supply and substrate noise. Optional changes on chip isolate the tracking jitter and duty-cycle of the PLL from distortions in the output channel.

loop bandwidth.

The accompanying figure shows an example of a board setup and optional chip setup for characterizing PLLs with added noise. An external pulse generator couples low-frequency square-wave noise into either just AVDD (the positive analog PLL supply) for supply-noise tests or both AVDD and AVSS (the negative analog PLL supply) for substrate-noise tests. Applying noise common to both AVDD and AVSS with respect to VSS, which dominates the chip substrate potential, is equivalent to applying noise to just the substrate. These board features can be added through rework to any board, including production boards, as long as the PLL supplies are accessible.

First, characterize noise

Only surface-mount components should be used in the supply-noise coupling network. Before performing any jitter measurements, the noise on the supplies should be characterized. While the PLL will impose additional high-frequency noise on the supplies, this extra noise should be ignored, since it is correlated to the PLL output.

Cycle-to-cycle jitter can be measured by triggering an oscilloscope from the PLL output and observing the movement in time of the next edge of the same type one cycle later. Tracking jitter and long-term jitter can be measured by triggering an oscilloscope from the PLL reference input and observing the movement in time of the first PLL output edge. When the reference input and PLL output signals are both driven off-chip through a similar path to the oscilloscope, undesired jitter in the clock-output path, which is unrelated to the PLL, can be canceled out.

Both measurements should be performed with a relatively noise-free reference clock.

[Integrated Systems Design]