

Hidden Complexities of PLLs Are Revealed

by John G. Maneatis

In order to achieve the timing goals needed in their ASIC designs, many engineers decide to include phase-locked loops (PLLs). Phase-locked loops 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 to multiply the frequencies on-chip to produce any number of higher-frequency internal clock signals. They also allow designers to control setup and hold-time windows and clock-to-output delays at the chip interface by aligning the windows to the edges of the chip's clock source.

While seemingly simple in structure and function, PLLs are filled with hidden complexity that can give even the best circuit designers trouble. 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. The use of such devices is often required to meet the operating frequency targets and to maintain supply voltage flexibility. However, the reduced supply-voltage headroom will adversely affect the PLL noise performance. ASIC designers must be aware of potential performance pitfalls for PLL designs, how to properly characterize the performance of PLLs, and how to detect these issues in performance results that can affect chip timing budgets. With this understanding, they will be better able to successfully decide upon which PLL to employ and how to best integrate them into their chip designs.

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, shown in Fig. 1, is seemingly straightforward. It is composed of a phase detector, charge pump, loop filter, and voltage-controlled oscillator (VCO). Immediately after the PLL circuit is activated, it is in an "unlocked" state as 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 "lock", 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.

In order for the PLL to work properly as a closed loop control system, the behavior of each block must be carefully established. The overall behavior of the PLL can be described in terms of its bandwidth and damping factor. The bandwidth reflects the rate at which the PLL responds to changes in the reference

frequency. The damping factor characterizes the stability of the PLL. Low damping factors (1) cause the PLL output frequency to ring when the reference frequency changes, (2) cause the PLL to amplify phase noise on the reference signal, and (3) may cause the PLL to be unstable. High damping factors cause the PLL to initially respond rapidly but later slowly to a change in reference frequency or phase, which can be undesirable.

Noise Issues

The blocks within a PLL can be comprised 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 today's 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 on off-chip interfaces by causing setup and hold-time violations which lead to data transmission errors. While other performance issues like 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.

Consider, for example, the measured output jitter histogram in Fig. 2. It shows the traces of many PLL output clock edges triggered from edges on the reference input and a

histogram with the number of output edges as a function of their center voltage crossing time. Most of the edge samples occur very close to the reference, while a few outlying edges occur far to either side of the center peak. These outlying edges must be within the jitter tolerance of the interface. These few edges are typically caused by infrequent data-dependent noise events with fast rise times.

In general, output jitter for PLLs can be caused by both jitter in the reference signal and by a number of noise sources. The noise sources include thermal noise, flicker noise, and supply and substrate noise. Thermal noise is generated by electron scattering in the devices within the PLL and can be significant at low bias currents. Flicker noise is generated by mobile charge in the gate oxides of devices within the PLL and can be significant for low loop bandwidths. Supply and substrate noise is generated by on-chip sources external to the PLL, including chip output drivers and functional blocks such as adders and multipliers, and by off-chip sources. This noise can be very significant in digital ICs.

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 is possible for 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 by the IC process. To reduce the risk of latch-up, 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 make 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.

Output Jitter Types

Output jitter can be measured in a number of ways. It can be measured relative to (1) absolute time, (2) another signal, or (3) the output clock itself. The first is commonly referred to as absolute jitter or long-term jitter. The second is commonly referred to 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, long-term jitter and tracking jitter for the output signal are equivalent. The third is commonly referred to as period jitter 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 is interesting to applications that cannot tolerate any edges with time displacements beyond some absolute level. The peak-to-peak jitter specification is typically the only useful specification 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 where 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

Measuring jitter correctly can be difficult and challenging. Given that PLLs must operate in a noisy mixed-signal environment, it is important to measure them within an

equivalent noise environment. Measuring PLLs in a quiet, low-noise environment can yield optimistic and misleading jitter results. Also, when artificial noise is applied to the PLL's 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 that is less than the reference frequency. The frequency of this noise signal can be above the loop bandwidth.

Fig. 3 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. The application of 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 re-work to any board, including production boards, as long as the PLL supplies are accessible. 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 additional noise should be ignored since it is correlated to the PLL output.

If the chip is being designed as a test chip, additional multiplexers can be added around the PLL as shown in the figure to facilitate

the measurement of tracking jitter and duty cycle. The two parallel paths make any added jitter outside of the PLL common to both signals, allowing only the jitter from the PLL to be characterized. The multiplexers allow the delay through each output path to be independently characterized by allowing the reference input to be bypassed directly to the output. The selectively inverting multiplexers allow the duty cycle of the PLL output to be measured without being distorted by the off-chip drivers.

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.

Careful review of jitter measurements performed under real or simulated noise conditions is critical for the proper evaluation of proposed PLL designs. Jitter results are often measured improperly or obtained under noise-free conditions which can be very misleading. If a PLL has inadequate noise rejection, chip timing budgets will be exceeded, which will cause the chip to fail in unusual and seemingly random ways that are difficult to diagnose and rectify. Such failures may not surface until late in system bring-up or even in manufacturing and may lead to additional costly silicon fabrication runs and schedule delays. With a good understanding of the noise sensitivity issues

that can compromise PLL jitter performance, designers can make informed decisions on the best PLL design to use in their ASIC designs.

Figure Captions:
(Figures not available at this time)

Fig. 1 A PLL, composed of a phase detector, charge pump, loop filter, and voltage-controlled oscillator (VCO), can be used to cancel-out clock distribution delays.
(Title: Chip Clock Distribution and Generation with a PLL)

Fig. 2 This measured output jitter histogram shows outlying edges that are typically caused by infrequent data-dependent noise events with fast rise times.
(Title: Sample Measured Output Jitter Histogram)

Fig. 3 Simple changes 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.
(Title: Board and Chip Setup for Jitter Characterization)

Biography:

John G. Maneatis is president of True Circuits, Inc. (Los Altos, CA). Maneatis has been developing PLLs and other mixed-signal circuits for 13 years. He holds a Ph.D. degree in Electrical Engineering from Stanford University (Stanford, CA). Maneatis is the author of many publications on PLL circuits and is an associate editor for the IEEE Journal of Solid-State Circuits.