Selecting PLLs for ASIC Applications Requires Tradeoffs

John G. Maneatis, Ph.D., President, True Circuits, Inc.
Los Altos, California
October 7, 2004

Phase-Locked Loops (PLLs) are commonly used to perform a variety of clock processing tasks, such as clock frequency multiplication and clock deskewing. PLLs, like many other analog IP macros, come with many features and specifications. Selecting the correct PLL early in the design can help the design team make tradeoffs when they are less costly, improve the integration quality of the PLL, and avoid surprises close to tapeout.

A typical PLL application is shown in Figure 1. In this application, a PLL is used to align the setup and hold time window for chip input latches to the input clock edge which is at half the frequency. More specifically, the PLL is used to multiply the clock frequency by two and then align the edges of the distributed output clock to those of the received clock. The PLL accomplishes this task by adjusting the frequency of a voltage-controlled oscillator (VCO) which drives the output clock so that the distributed clock, once divided in half and fed back to the PLL, matches frequency and lines up with the received clock. This simple application requires certain features in the PLL, like the ability to accept a feedback clock and the ability to multiply the input frequency.

Figure 1: PLL application for deskewing a clock distribution network.

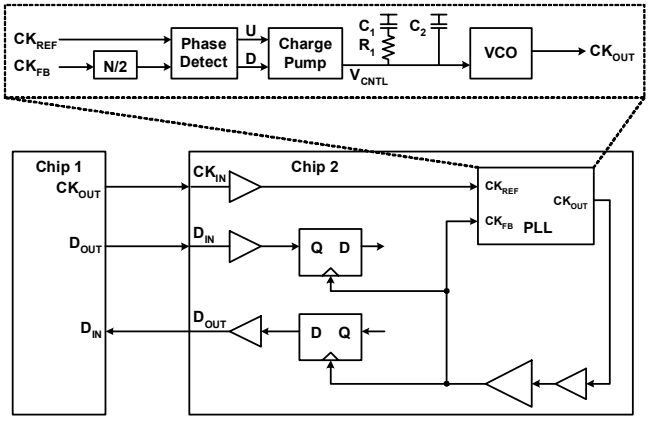

Given the large application base for PLLs, there are many possible features that might be needed in general. Like all features, they typically are accompanied by tradeoffs of one sort or another. Table 1 lists typical PLL features for digital applications and the advantages and disadvantages of each. They are categorized by features affecting the loop dynamics, circuit structures, loop configuration, output configuration, and counter configuration. Many of the tradeoffs relate the PLL loop dynamics and they affect the output jitter.

Table 1: Tradeoffs between PLL features.

Loop Dynamics	Advantages	Disadvantages
High bandwidth	• Low tracking jitter • Low long-term jitter • Needed for deskew PLLs	• No long-term or period jitter filtering • More difficult to avoid pattern jitter • Difficult to implement frequency spreading • Higher frequency overshoot
Low bandwidth	• Low period jitter • Period jitter filtering • Low frequency overshoot	• Poor long-term jitter (without high-Q VCO) • No long-term jitter filtering (without high-Q VCO)

Loop Dynamics	Advantages	Disadvantages
Tracking bandwidth (i.e. self-biased PLLs)	• Wide operating frequency range • Wide multiplication range • Constant loop dynamics • Maximizes tracking bandwidth • Minimizes PVT sensitivity	• Constrained analog circuits

Circuit Structure	Advantages	Disadvantages
Low-Q VCO (ring/relaxation oscillators)	• Wide VCO frequency range • Can have low period jitter • Most common approach for digital chips	• High phase noise • Not suitable for SONET transmit clocks
High-Q VCO (LC/Xtal oscillator)	• Low long-term and period jitter • Long-term and period jitter filtering • Low phase noise	• Narrow VCO frequency range • Requires a lot of characterization • Large area
High supply/substrate noise rejection	• Required for good jitter performance	• Chip may be doomed without it
Low power	• Reduces power requirements	• Increases jitter levels

Loop Configuration	Advantages	Disadvantages
Clock multiplication	• Input frequency flexibility	• Can exhibit pattern jitter
Deskewing	• Control chip input setup/hold windows	• Limited practical multiplication (tracking jitter)
Frequency spreading	• Can help reduce EMI emissions	• Typically requires low bandwidth
Fast locking	• Reduced lock time (important for low power, low bandwidth, and clock recovery applications)	• Possible frequency overshoot or instability
Clock recovery	• Used for data communication applications	• Typically have false lock problems • May need frequency acquisition aids

Output Configuration	Advantages	Disadvantages
Divided phase aligned outputs (i.e. /1, /2, /4 outputs)	• Facilitates applications like DDR/QDR	• Require matched clock distributions
Multiple output phases and adjustable output phases	• Useful for clock recovery applications • Best place to generate them is in PLL • Can achieve highest accuracy without calibration	• Increases PLL power dissipation • Potential pattern jitter issues

Counter Configuration	Advantages	Disadvantages
Built-in counters	• Avoids frequency headroom problems • Can minimize insertion delay for lower jitter	• Limits flexibility
Counter resetting	• Facilitates testing • Reduces reset time	• Slows down counters
Dynamic counter value changes	• Can implement software frequency adjustments	• More complicated synchronous counter interface
Lock detector	• Signals when PLL is locked • Can reduce lock time	• Can be unreliable (error threshold is double constrained, bandwidth constrained) • Can increase lock time
Cycle-slip detector	• Signals if the PLL is unlocked	• Cannot directly use as lock detector
Testing circuitry	• Increases fault coverage in digital circuitry • Enables vector testing without locking PLL • Enables trouble-shoot analog circuitry	• Must take care not to increase jitter (excessive MUXing, exposing analog voltages)

Tradeoffs between PLL Bandwidth and Output Jitter

Output jitter refers to the time-varying offset in the output clock edges. 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. Output jitter is commonly caused by jitter on the reference clock or by sensitivity in the VCO and other circuits to noise sources both inside the PLL block, such as thermal noise and flicker noise, and outside the PLL block, such as supply and substrate noise. While all of these noise sources are commonplace and unavoidable, supply and substrate noise will typically dominate in the harsh mixed-signal environment of today's ASICs for many applications.

There are several ways of measuring jitter. Jitter can be measured as variations in the clock period, commonly referred to as period jitter. Period jitter is important as it subtracts from the available setup time in logic paths. Jitter can also be measured as variations in the accumulated width of many clock periods, called long-term jitter. Long-term jitter is important in communication applications where other PLLs might have to directly or indirectly (through a serial data stream) track the output clock of a given PLL. In addition, jitter can be measured as the variation in time relationship between the edges of the reference clock and the output clock, called tracking jitter. Tracking jitter is important in clock deskewing applications like memory interfaces.

In order to understand the tradeoffs between the PLL loop dynamics and output jitter, it is important to understand how jitter is created and filtered inside PLLs. The sensitivity of the VCO to noise sources will lead to time variations in the VCO output periods. These variations will accumulate as time offsets in the generated clock edges over time. The negative feedback of the PLL will track out accumulated jitter over periods longer than the response time of the PLL, measured as its loop bandwidth. Under worst-case noise conditions, the time variations in each VCO output period can be correlated with successive output periods so that the jitter accumulates at a constant rate over time. With the accumulation time limited by the loop bandwidth, long-term jitter measured over periods longer than the response rate of the PLL is generally inversely proportional to the loop bandwidth.

The amount of long-term jitter that will result depends on the sensitivity of the VCO to noise. VCOs based on LC oscillators typically have high quality factor (Q), which gives them inherent stability that can substantially reduces their sensitivity to noise sources. Alternatively, VCOs based on RC oscillators, such as relaxation or ring oscillators, have low-Q, and thus are very sensitive to noise. The best such low-Q VCO circuits can do is employ isolation techniques to minimize their sensitivity to noise sources, but the sensitivities will be fundamentally higher than for high-Q oscillators. Thus, while high-Q VCOs can achieve low long-term jitter even with a low loop bandwidth, low-Q VCOs can only obtain low long-term jitter by maximizing the loop bandwidth and tracking the input frequency as close as possible.

Since PLLs only track the reference clock at a rate limited by their bandwidth, they can be used to filter out shorter-term jitter, such as variations in the reference clock period (period jitter), but will let longer-term jitter (jitter measured over many reference cycles) pass through. Since high-Q VCOs have low jitter accumulation, a low bandwidth can be used to filter out long-term jitter over large numbers of reference cycles. With low-Q VCOs, using a low bandwidth to filter out long-term jitter is a losing proposition, unless most of the jitter accumulates over a small number of reference cycles.

The aforementioned issues suggest that high-Q VCOs should be used in all PLLs. However, high-Q VCOs have limited frequency range and typically require laboratory characterization of silicon to determine the center frequency. They also tend to require significantly more chip area. Because of these limitations, PLLs with low-Q VCOs are the predominant design type used on ASICs because they have much larger tuning ranges, can be implemented with less area, and are generally more flexible. Clearly, there are independent advantages and disadvantages to

each type of VCO. However, it is important to note that it is not possible to achieve stringent phase noise requirements of many communications systems with RC oscillators.

Thus in summary, applications that require very low long-term jitter should be based on high-Q VCOs. Applications that only require very low period jitter can simply use medium to low bandwidth PLLs based on a low-Q VCOs. Applications that require low tracking jitter should use a high bandwidth PLL based on low-Q VCOs.

Other tradeoffs between loop bandwidth and jitter relate to the interconnection of PLLs. It is very common to have one PLL drive another. Ideally the second PLL should have a higher or lower bandwidth than the first. Since PLLs tend to slightly amplify input jitter at the loop bandwidth, cascading multiple PLLs at the same bandwidth will lead to more significant jitter amplification. In some SERDES interfaces, a high-frequency PLL might be used to generate the transmit clocks that require very low long-term jitter. Unless the PLL contains a high-Q VCO, it will simply pass the long-term jitter problem to its input reference clock. If another PLL is used to generate the mid-frequency reference for the output PLL, the long-term jitter of the final output clock will be fundamentally limited to some fraction of the reference clock period of the first PLL, independent of how many PLLs are used to multiply up that reference frequency.

Tradeoffs with Other PLL Features

Aside from loop bandwidth settings, other PLL features such as clock multiplication, clock deskewing, and frequency spreading can have significant tradeoffs inside the PLL, all of which, once again, are linked to output jitter.

Clock multiplication is a commonly desired feature in order to minimize the input clock frequency requirements and provide flexibility for on-chip clock frequencies. PLLs that perform clock multiplication are susceptible to a form of period jitter called pattern jitter. Pattern jitter results from the noise generated by the phase comparisons every PLL reference cycle and appears as pattern of output period jitter that repeats every reference cycle. Pattern jitter can be difficult to filter in a PLL that operates at a high bandwidth without causing instability.

PLLs that perform deskew function must be able to tolerate some amount of delay in the feedback path, typically composed of clock distribution networks or, in extreme cases, off-chip paths. In order to tolerate more feedback delay, a PLL must operate at a lower bandwidth, which will in turn increase the resultant tracking jitter. Also, as a PLL approaches a locked state, it will typically overshoot the target output frequency by some amount that increases with its loop bandwidth. The feedback path must be able to pass the maximum overshoot frequency without losing any edges or the VCO will become stuck at the maximum frequency and the PLL will fail to lock. This issue is most significant when the feedback path includes an off-chip path. Once again, reducing the amount of overshoot requires a reduction in the PLL bandwidth, leading to more tracking jitter.

PLLs performing frequency spreading functions can operate at either a low bandwidth or a high bandwidth, depending on the rate of frequency spreading. When operating at a high bandwidth, the PLL attempts to track the input clocks through a continuously adjusted feedback divider, leading to the spreading frequency characteristics. Operating at a high bandwidth is only possible if the adjustments to the feedback divider lead to smooth frequency changes rather than stair-step changes which are typically undesirable. When operating at a low bandwidth, the frequency spreading is obtained by modulating the VCO control directly, but at a rate much higher than the loop bandwidth so the PLL does not attempt to track out the modulation. If a low bandwidth is used, the PLL will not be able to achieve low long-term jitter as previously discussed.

PLL Applications

Given the tradeoffs between different PLL features, the features ultimately incorporated into the PLL will be set by the PLL application. Table 2 lists different digital applications for PLLs and their associated requirements and desired features. The applications can be categorized into a few classes: clock generation applications requiring low period jitter, clock deskew applications for synchronous interfaces, generating transmit and receive clocks for data communication, and generating frequencies that require high fractional precision.

Table 2: Feature requirements of various PLL applications.

Application	Requirements	Features
• Logic core clocks • Double-data rate (DDR) interfaces • Source-synchronous interfaces	• Minimum period jitter • Clock multiplication • May need frequency spreading	• Medium to low bandwidth • May need phase aligned divided clocks
• Synchronous interfaces (chips receiving data aligned with clock)	• Clock distribution deskewing • Clock multiplication • Low tracking jitter	• High bandwidth • External feedback path
• SERDES (SPI4, etc.) transmit path	• Low long-term jitter	• High bandwidth • Medium bandwidth with high-Q VCO
• High-speed SERDES (SONET OC192, 10G Ethernet) transmit path	• Very low long-term jitter	• Medium bandwidth with high-Q VCO
• SERDES receive path (clock and data recovery)	• Reject inter-symbol interference • Ignore missing transitions • Track data stream	• Medium bandwidth • May need multiple output phases
• Driving PLLs in SERDES cores	• May need low long-term jitter (if core does not use an high-Q VCO) • Clock multiplication	• High bandwidth
• Disk drive clock recovery	• Reject inter-symbol interference • Ignore missing transitions • Lock quickly to data stream	• Medium bandwidth • Fast locking
• Video clock generation	• Low long-term jitter • Low period jitter • Wide multiplication range for high fractional precision	• High tracking bandwidth • Pattern jitter rejection

Most data communication applications have tight jitter specifications on the transmitted and received data. These specifications make it possible to design the transmitter independent of the receiver and to address various physical effects on the channel. These jitter specifications must be considered carefully in choosing the best PLL for the application.

When selecting PLLs for ASIC applications it is important to understand the tradeoffs between different PLL features. Many of these features tradeoff directly with the PLL's jitter performance. The PLL loop bandwidth and the type of VCO are the biggest factors impacting output jitter. The best PLLs for a given application are PLLs that most flexibly address the design tradeoffs and avoid introducing tradeoffs not fundamental to the phase locking problem.