

Phase-Lock Loop-Based (PLL) Clock Driver: A Critical Look at Benefits Versus Costs

***Stevan Plote
Advanced System Logic – Semiconductor Group***

SCAA033



IMPORTANT NOTICE

Texas Instruments (TI) reserves the right to make changes to its products or to discontinue any semiconductor product or service without notice, and advises its customers to obtain the latest version of relevant information to verify, before placing orders, that the information being relied on is current.

TI warrants performance of its semiconductor products and related software to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are utilized to the extent TI deems necessary to support this warranty. Specific testing of all parameters of each device is not necessarily performed, except those mandated by government requirements.

Certain applications using semiconductor products may involve potential risks of death, personal injury, or severe property or environmental damage ("Critical Applications").

TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, INTENDED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT APPLICATIONS, DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS.

Inclusion of TI products in such applications is understood to be fully at the risk of the customer. Use of TI products in such applications requires the written approval of an appropriate TI officer. Questions concerning potential risk applications should be directed to TI through a local SC sales office.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards should be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance, customer product design, software performance, or infringement of patents or services described herein. Nor does TI warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used.

Contents

	<i>Title</i>	<i>Page</i>
Introduction	1
Acknowledgements	4
References	4

Introduction

Today, system clock frequencies continue to increase; we are now approaching the 50–100-MHz range. The clock period with which systems designers must work is shrinking, as is the tolerance for high propagation delays (t_{pd}) and high output skew ($t_{sk(o)}$, $t_{sk(p)}$, and $t_{sk(pr)}$) in clock-distribution systems:

$t_{sk(o)}$ is output-to-output skew in the same device.

$t_{sk(p)}$ is the difference between the low-to-high and high-to-low transition for a given output terminal:

$$t_{sk(p)} = |t_{PHL} - t_{PLH}| \quad (1)$$

$t_{sk(pr)}/t_{sk(pv)}$ is process skew or part-to-part skew and is a measure of the difference between the minimum low-to-high or high-to-low transition and the maximum high-to-low or low-to-high transition on the same terminal of two different clock drivers under the same operating conditions:

$$t_{sk(pv)} = |\min t_{pd} \text{ LH (device 1)} - \max t_{pd} \text{ HL (device 2)}| \quad (2)$$

Some manufacturers specify this parameter as:

$$t_{sk(pv)} = |\min t_{pd} \text{ LH (device 1)} - \max t_{pd} \text{ LH (device 2)}| \quad (3)$$

This is a less-stringent specification.

Table 1. Clock-Driver Timing Requirements

	SYSTEM CLOCK		
	50 MHz	66 MHz	100 MHz
Clock cycle time, t_c^\dagger	20 ns	15 ns	10 ns
Clock pulse duration, t_w^\ddagger	10 ns	7.5 ns	5 ns

† Clock cycle time $\geq 1/\text{system clock frequency}$

‡ Assumes a 50% duty cycle

These timing requirements imply that the t_{PLH} and t_{PHL} through the clock buffer and the maximum allowable output skews $t_{sk(o)}$ and $t_{sk(pr)}$ in a given system need to be less than or equal to the clock-pulse duration to not violate system timing specifications:

$$|t_{PLH} + t_{PHL} + t_{sk(o)} + t_{sk(pr)}| \quad (4)$$

Maximum allowable $t_{sk(o)}$ for a 50% duty-cycle clock system is 10% of the clock cycle time. This is an estimate.

None of the very fastest gate- and buffer-based clock-distribution devices can meet this timing. A system designer must therefore turn to a PLL-based clock driver as shown in Figure 1.

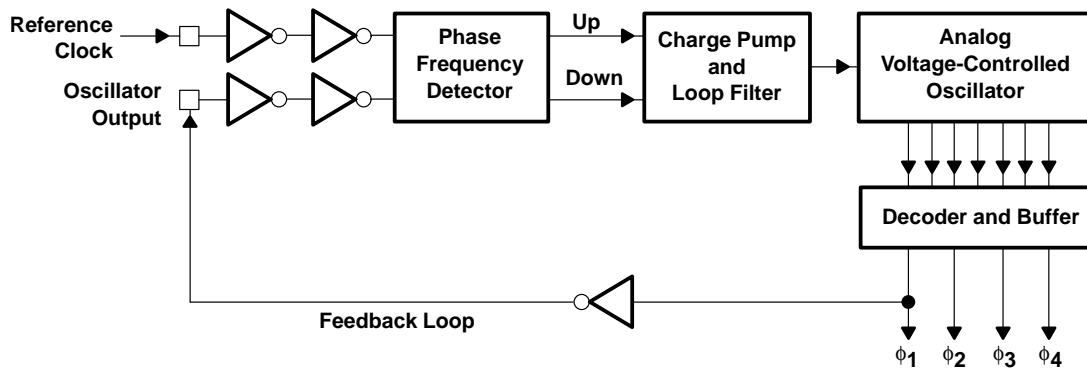


Figure 1. Block Diagram of a Clock Generator Based on a Charge-Pump PLL

The beauty of phase-lock loops is that they receive an input signal, compare this to the feedback of their internal clock generated by a voltage-controlled oscillator (VCO), and adjust the VCO by way of the charge pump to match the new input frequency and synchronize the internal and external clocks.

In Figure 1, the analog VCO is either a ring-oscillator type or a multivibrator type. The VCO can also be designed using a multistage tapped delay line that is calibrated to a precise delay per stage (a digital approach). The charge-pump design has various approaches using inverters, switches, and a passive RC low-pass filter. The phase detector is the second most important element (see Figure 2). The input from the external clock enters the phase detector, which is a set of balancing buffers and highly balanced D-type flip-flops. The phase detector must always be active. This clock input is compared to the feedback input D from the VCO.

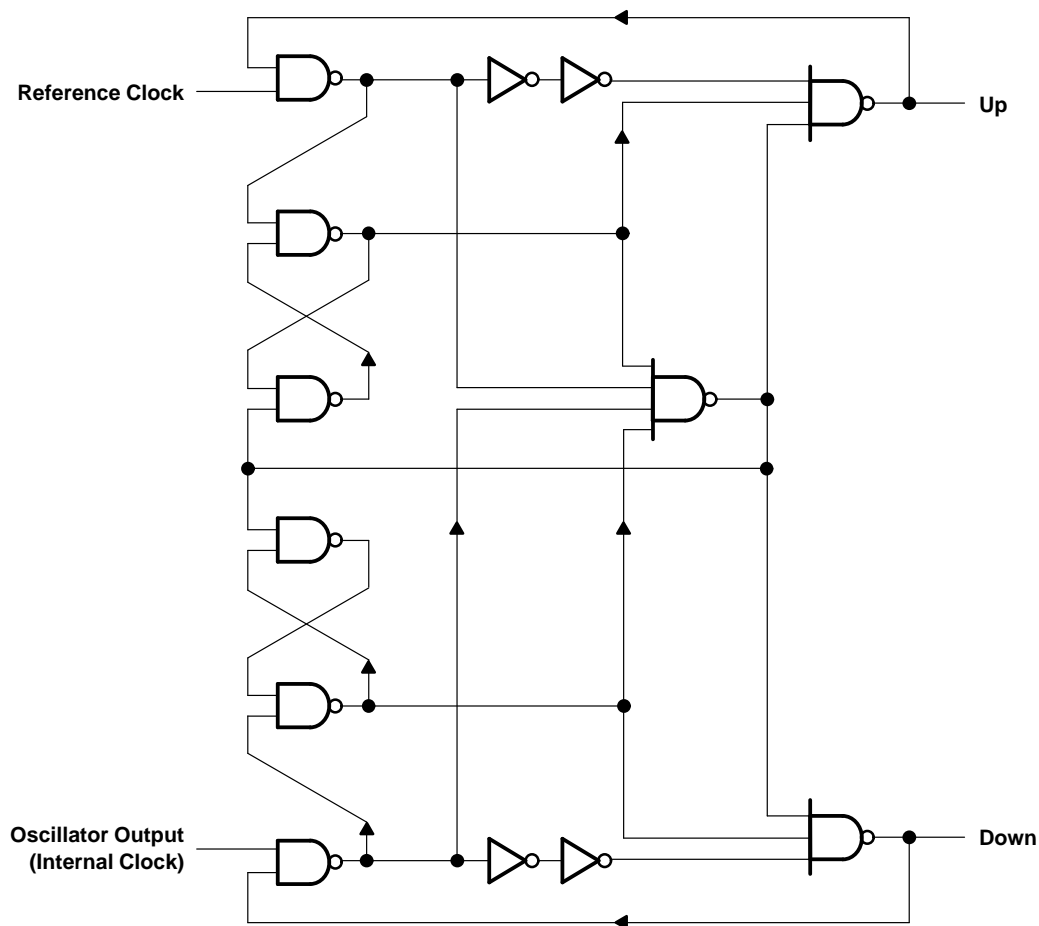


Figure 2. Phase-Detector Block

If the D input to the flip-flop is high before the rising edge of the clock, then D is phase advanced, the voltage to the VCO is reduced via the charge pump, and the internal clock is slowed.

If the D input to the flip-flop is low before the rising edge of the clock, then D is phase retarded and the voltage to the VCO is increased through the charge pump to speed up the VCO.

This process is repeated with every external input-clock pulse so that the feedback clock and the external clock are synchronized; the circuit then locks onto itself within a narrow frequency band. If the input clock slightly varies (within that frequency band), the PLL frequency does not vary. This narrow frequency band is known as the *dead zone* of the phase detector. Within this zone, the feedback clock and the external clock are so close in phase that there are no correction pulses out of the phase-detector circuit. Once the phase drifts out of this frequency band, the phase detector starts correcting again.

Circuit designers have tried to design nondead zone phase detectors such that the phase detector is always active and correcting. This, however, is very difficult to implement due to t_{jitter} , an important specification that results from the phase-detector circuit and noise in the VCO:

$$t_{\text{jitter}} \text{ of the phase detector} = \text{dead zone} + \text{correction pulse}$$

The phase detector must be very accurate and balanced to reduce the dead zone and keep the correction pulses small. The design result is a more analog than digital implementation.

The phase-lock loop locks in a very short time (less than 50 ms). In fact, after lock, many PLL-based clock drivers are termed *zero delay*. The reality is that most have a t_{pd} of ± 1 ns from the input frequencies. Comparing this to the 3- to 12-ns t_{pd} of other gate-based or divider clock drivers, there is a great advantage to a high-performance system designer in using a PLL-based clock driver.

Another feature of a PLL is the skew control of the device. Since the input signal is locked onto and regenerated at the output of the device, the variation of the signal from output to output is no longer a function of the chip layout and process as it is in gate- and flip-flop-based devices. PLL designs achieve maximum output skew [$t_{\text{sk(o)}}$] of 500 ps or less and process skew [$t_{\text{sk(pr)}}$] of < 1 ns; very important features from the designer's point of view.

Two other aspects of a PLL design allow for additional functionality. One of these is external feedback (not all PLL devices have this). This allows a designer to use an external gate- or flip-flop-based clock driver to drive multiple loads off of one PLL output. The PLL output drives the input of the external clock driver and, by feeding one output of that external buffer back into the PLL-based clock driver through the feedback terminal, one can synchronize the remaining PLL outputs with the remaining external buffer outputs and the input clock to the PLL-based clock driver. The other aspect is output jitter (t_{jitter}). Jitter occurs when a signal deviates in phase or frequency from that of an ideal clock. This shows up as noise on the outputs of the device. When a PLL locks into the input frequency, there is a limited variation of that signal at the outputs. This provides for good, low t_{jitter} specifications both on individual outputs and the entire device.

Multiple-board systems derive a huge benefit from this feature. A designer can use a master PLL based on a motherboard and synchronize the other boards in the system by driving the oscillator signal through the master PLL and out through the backplane, then recovering it on each of the boards via a PLL-based clock driver on each of the system boards. This works best because the master PLL has the lowest signal loss/output skew and can have high-drive outputs. Each system board is then synchronized to the master clock, and each board individually drives the clock via its own clock driver. This application also works best when the oscillator signal is divided down, driven across the backplane, and multiplied back up through the clock recovery PLLs to the system operating frequency desired across all of the boards. By reducing the clock frequency driven across the backplane, the level of extraneous noise in the signal is also reduced.

The main drawback to the PLL-based clock driver is cost. Due to the complexity of the circuitry and speeds at which the VCO needs to run, PLLs are expensive. In general, a PLL-based clock driver costs two to five times the price of a gate-based clock driver. This price is based upon the value of the product. If the accuracy and speed of the PLL is not needed, other solutions can be used.

Other disadvantages of PLLs are:

1. They are inherently very noise sensitive.
2. They are untestable by the end user, but can have the VCO bypassed in a system to allow examination of other parts of the device circuitry (mainly inputs) during debug stages of board-level design.
3. Some PLLs can require expensive, high-quality external components to implement the loop-filter design.
4. The external loop filter may have to be modified from part to part due to processing variations in a vendor's PLL silicon.
5. t_{jitter} can also occur due to substrate conditions. Isolation of key components such as the VCO and charge pump from the outputs and on-chip digital circuitry can reduce t_{jitter} . t_{jitter} is also increased if an external feedback terminal is implemented in the design along with the reference clock terminal. If the feedback terminal is tied to a separate V_{CC} and ground from the analog circuitry, it can increase jitter via shifting rail and ground levels between the feedback and the analog VCO and reference input. This can be combatted by tying the feedback terminals/clock input/ VCO/charge pump to the same V_{CC} and ground.

Texas Instruments has developed a group of three low-voltage, high-performance, PLL-based clock drivers. These devices are targeted at the high-performance (50 MHz–100 MHz), 3.3-V power-supply markets for RISC processors, Intel Pentium™ microprocessors, and synchronous DRAMs.

Each of these markets require multiple outputs, twelve on each device, and a way to configure the device outputs to be one half the input frequency, two times the input frequency, and the same frequency at the input frequency. Some of the devices also incorporate dampening resistors on the outputs to reduce reflections caused by transmission-line effects and increase the integrity of the signal at the load.

The TI devices incorporate a five-stage ring oscillator VCO, and internal loop filter, and an external feedback terminal (which allows for doubling the input-clock frequency at the outputs). The VCO, charge pump, reference-clock input, and feedback terminal are all tied to the same V_{CC} and GND and fully isolated from the remaining on-chip logic and outputs of the device.

This design is inherently stable and exhibits low t_{jitter} on each of its outputs. Measured t_{jitter} in spice simulation is 48 ps or less with a typical number of 23 ps.

The CDC2582, CDC2586, and CDC586 are very competitive solutions for the telecommunications, workstation, and PC-equipment clock distribution requirements.

Acknowledgements

Tim Ten Eyck for suggestions and design work.

Brett Clark for his dedicated lab work.

Craig Spurlin for his design work.

References

Jeong, Deog-Kyoon, Gaetano Borriello, David Hodges, and Randy Katz, "Design of PLL Based Clock Generation Circuits," IEEE Journal of Solid-State Circuits, Vol. 22, No. 2, April 1987.

Manasce, Victor, and Steve Sadler, "A Fully Digital Phase Locked Loop," Bell Northern Research.

Young, Ian, Jeffrey Greason, and Ken Wong, "A PLL Clock Generator With 5 to 110 MHz of Lock Range for Microprocessors," IEEE Journal of Solid-State Circuits, Vol. 27, No. 11, November 1992.