

SYSTEM APPLICATIONS AND DESIGN GUIDELINES WITH IDT'S ZERO-DELAY BUFFERS

Overview

This application report presents different methods of achieving a defined phase relationship between the input reference clock and output clock for IDT PLL-based zero-delay buffers. The report focuses on the effects of varying trace lengths in the PLL feedback loop and capacitive loading at the feedback clock input in generating an early (or late) clock. Early clock is defined as a case when the phase of the output clock is advanced in relation to the reference input clock.

Introduction

There are two major types of clock driver architectures: a buffer-type and a feedback-type with phase-locked loop (PLL). In a buffer-type (non-PLL) clock driver, the input wave propagates through the device and is “re-driven” by the output buffers. The output signal directly follows the input signal with some device propagation time delay (t_{PD}). For additional information on Fanout buffers please refer to IDT Clock Distribution Application Note.

The feedback-type clock drivers utilize an additional feedback input path from one of the device clock outputs. The feedback path can be internal or external to the device. In the case of an external feedback, a trace on the PCB is used to connect an output pin to the feedback input pin. This type of device is usually based upon one or more PLLs that are used to align the phase and frequency of the feedback and reference input. Board-level trace-length mismatches can be compensated in order to achieve defined skew relationship between the input and output clock. Outputs can be selectively divided, multiplied or inverted while still maintaining very low input to output skew.

What is a PLL?

A basic PLL is a feedback system that receives an incoming oscillating signal and generates an output waveform that oscillates at an integer or fractional multiple of the input signal. It is comprised of a phase or frequency detector, a filter and a voltage-controlled oscillator as shown in Figure 1. In order for the PLL to align the reference input (CLKIN) with an output, the output must be fed back to the feedback (FBIN) input of the PLL.

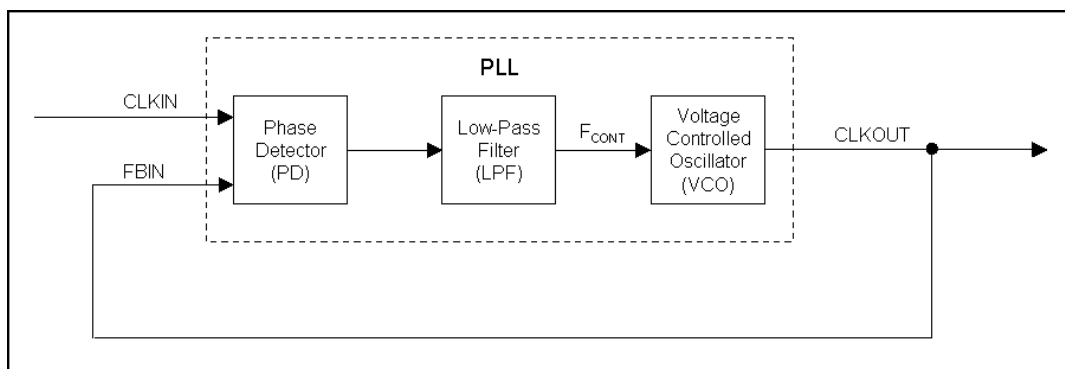


Figure 1. Phase-Locked Loop (PLL) Block Diagram

The Phase Detector (PD) typically evaluates the rising edge of the CLKIN input with respect to the FBIN input. If the CLKIN input leads the FBIN input, indicating that the voltage controlled oscillator (VCO) is running slower than ideal, the PD produces a 'Pump Up' signal that is triggered on the rising edge of the CLKIN input and lasts until the rising edge of the FBIN input. This 'Pump Up' pulse forces the VCO to run faster and reduces the time difference between the CLKIN input and the FBIN input.

If the FBIN input leads the CLKIN input, the PD produces a 'Pump Down' signal that is triggered on the rising edge of the FBIN input and lasts until the rising edge of the CLKIN input. The 'Pump Down' pulse forces the VCO to run slower and reduces the time difference between the FBIN input and the CLKIN input.

The PD forces the VCO to run faster or slower based on the relationship between the CLKIN and FBIN input edge. The output of the VCO is the internally generated oscillator waveform. The input voltage that controls the frequency of the VCO is a measure of the phase deviation between the reference and feedback inputs. The PLL is designed to operate within a limited band of frequencies (tracking range). If the input frequency is outside this band, the circuit will not lock-on to the input signal and the frequencies of CLKIN (F_{CLKIN}) and CLKOUT (F_{CLKOUT}) will be different. As long as F_{CLKIN} remains within the tracking range of the PLL, F_{CLKOUT} will equal F_{CLKIN} .

The low-pass filter converts the 'Pump Up' and 'Pump Down' signals into a DC control voltage (FCONT) and the magnitude of this voltage is proportional to the width of the 'Pump Up' or 'Pump Down' pulse. The range of the voltage produced by the filter is guaranteed to force the VCO to operate at any frequency within its allowable operating range.

The PLL may respond to the absence of an reference input (CLKIN) clock in different ways. The output path from the PLL may be disabled thereby disabling the device outputs or the VCO could be forced to operate at its minimum operating frequency. The specifics of the device need to be known if the design will be placed in this condition. There are buffers that support dual clock inputs if the loss of an input clock is expected.

PLLs have a number of desirable properties that include the ability to minimize input-to-output skew (zero-delay buffers), multiply and generate different clock frequencies (clock synthesizers), correct input clock duty cycle distortion, reduce timing jitter (jitter attenuators) and extract clock signals from data streams (clock data recovery). The focus of this application report is on the use of PLLs in zero-delay buffers (ZDB). A ZDB aids clock tree designs that require zero propagation delay from the input signal to the output. A completely integrated PLL allows alignment in both the phase and the frequency of the reference input clock with the output clock.

What is a Zero-Delay Buffer?

A zero-delay buffer is a device that can fan out one clock signal into multiple clock signals with near zero delay. The device is well suited for a variety of clock distribution applications requiring tight input-output and output-output skew relationship. A simplified diagram of a ZDB is shown in Figure 2.

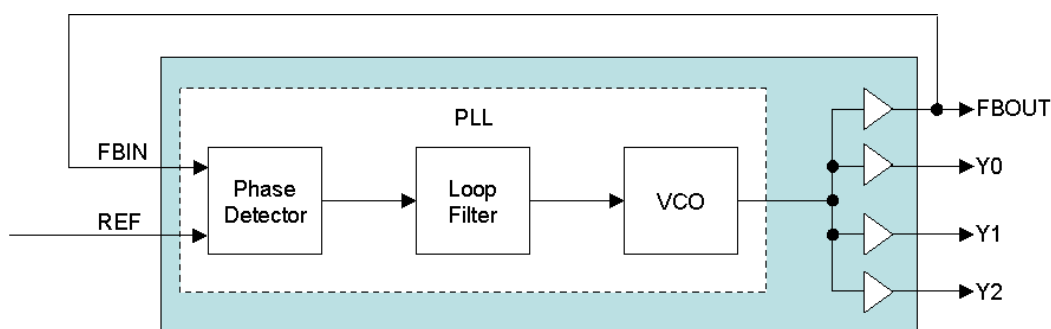


Figure 2. Simplified Zero-Delay Buffer (ZDB) Block Diagram

A ZDB is built with a PLL that uses a reference and feedback input. The feedback input is driven by one of the outputs. The phase detector adjusts the output frequency of the VCO so that its two inputs have no phase or frequency offset. Since the PLL feedback path includes one of the outputs and its associated load, it will dynamically compensate for the load placed on that output. This means that it will have near zero delay from the reference input to the output from which the feedback is derived, irrespective of the loading on that output. Note that this is only true for the output that provides the feedback. All other outputs have an input-to-output delay that is affected by the differences in the output loads. Please see the section “*Tuning a ZDB for Phase Adjustment*” for a detailed discussion on this topic.

Tuning a ZDB for Phase Adjustment

Depending on the application and delay requirement, the system designer can adjust the input-output delay for Zero Delay Buffers.

The output signal of the buffer is said to lead when it transitions earlier in time than the reference input signal. It can also be viewed as negative delay. On the other hand, the output signal is said to lag when it transitions later in time than the input signal (positive delay). To tune the device and adjust the lead or lag of the buffer outputs, the system designer must understand the relationship between REF and FBIN, and the relationship between the output driving FBIN and the other outputs.

The PLL compares the phase of the signal at the FBIN pin at a threshold of $VDD/2$ (where VDD is the device supply voltage) and compares it to the signal at the REF input pin at the same threshold. All the outputs (including the output driving FBIN) transition at the same time.

Changing the load on an output clock affects its rise time and therefore alters the time it takes for the output to reach the $VDD/2$ threshold. Using these properties to their advantage, the system designers can then adjust the time when the outputs reach the $VDD/2$ threshold relative to when the REF input reaches the $VDD/2$ threshold. The output driving FBIN, however, cannot be adjusted; it will always have near zero delay from the REF input at $VDD/2$.

In order to have a positive phase error (i.e. REF leads the outputs), the FBIN pin should be less loaded compared to the Y_n outputs. On the other hand, if negative phase error is desired (i.e. the Y_n outputs lead the REF input); the output used for the feedback can be loaded heavily to advance in time the other outputs.

When the difference in loading between the output used for FBIN and the other outputs is zero, the signals will be aligned in phase. If the loading on the output used for FBIN is different from the other outputs, the signals will not be aligned in phase. As a rule of thumb, 1pF will induce about 50ps to 60ps of skew. Figure 3 shows example waveforms of the input and output clock signals before and after mounting a 1pF capacitor between the FBIN input and GND. After mounting the 1pF capacitor, the CLKOUT signal is retarded by 60ps.

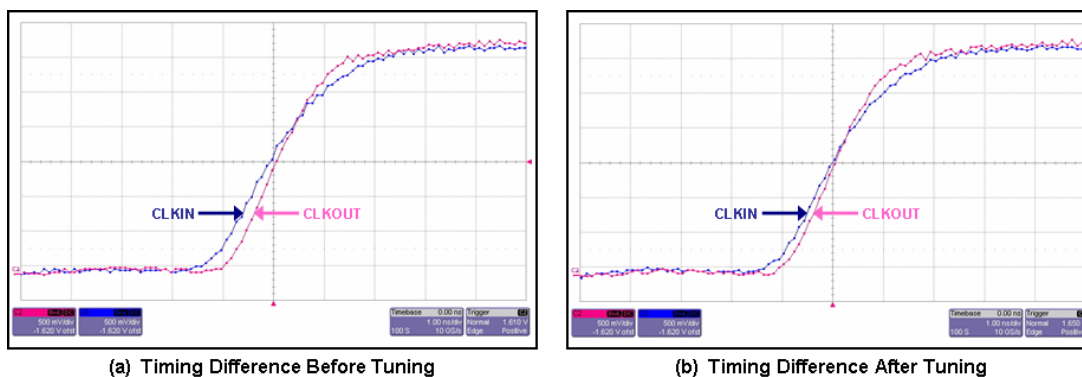


Figure 3. Timing Effect of Adding 1pF Capacitance Between FBIN and GND

Note that the zero-delay buffer will always adjust itself to keep the $VDD/2$ point of the output at near zero delay compared to the $VDD/2$ point of the reference.

A one inch trace adds roughly 3pF parasitic capacitance, which translates to approximately 150ps to 180ps delay. Therefore, the skew can also be altered by adjusting the feedback trace length relative to the trace length of the other outputs. The system designer can advance the Y_n outputs by increasing the feedback trace length, thereby generating an early clock. On the other hand, the system designer can shorten the feedback trace length and advance the REF input signal relative to the Y_n signal.

It is important to note that the capacitive loading on the FBIN pin is the best way to fine tune the phase error and this loading should be placed as close to the FBIN pin as physically possible. Adding additional capacitance beyond 30pF is not suggested due to the possibility of degrading the clock edges and adding more jitter to the outputs. Adjusting the trace length of the feedback loop is used more for the coarse tuning of phase error.

If the application requires the outputs of the ZDB to have near zero delay with respect to one another (output-output skew), the output clock of the device that is driving the zero-delay buffer must be loaded the same as the other outputs of the device.

Adjusting the lead or lag of the output skew with a capacitor has its benefits. However, it does have imperfections due to variation in the capacitor itself. As an alternative to the methods outlined above, a programmable skew device such as the IDT5T9890 or IDT5T9891 could also be considered.

Using External Feedback

Many ZDBs have an open external feedback path that is simply closed by driving one of the output clocks into the FBIN pin for zero delay buffer operation.

Alternatively, the feedback path can be used for other interesting applications. Using a discrete delay element in the feedback path will generate outputs that lead the input signal. Sometimes, designs require copies of a clock that are early compared to the input clock. Figure 4 shows a circuit implemented to generate such early clocks using a ZDB.

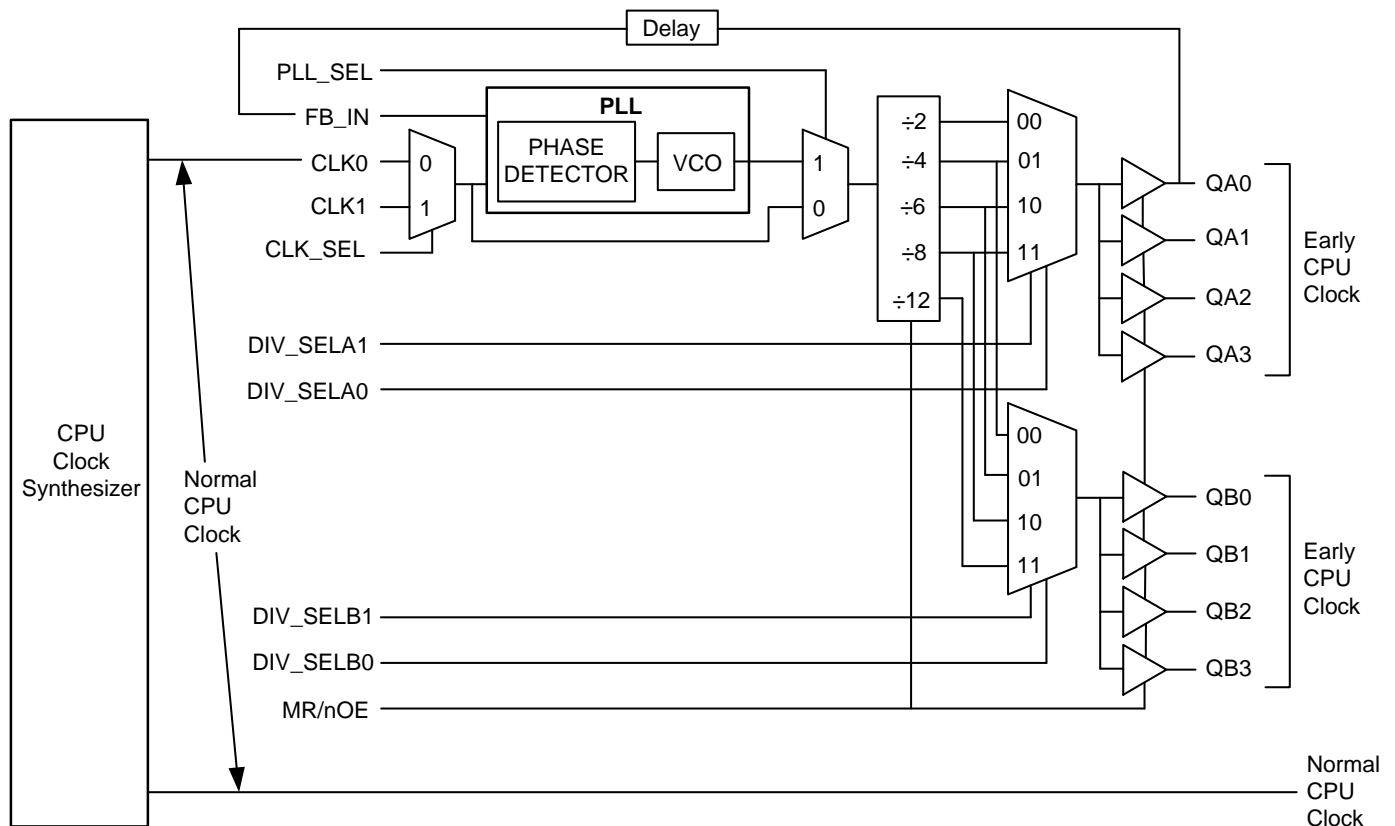


Figure 4. ICS8752i Used to Generate Early Clocks

Another simple approach to lead or lag output clocks is to insert a trace delay into the feedback path. The output of the buffer will lead the input by the amount of trace delay added in the feedback path. This approach provides a precise method for delay adjustment. Some designers will embed a very long trace into the board from an output pin to the feedback pin. At the ends of each trace segment, the designer places pads for zero ohm resistors. This allows for incremental addition of delay into the feedback path to align the output to the precise phase. Figure 5 shows an example where the feedback path is n-inches shorter than the other output traces. This initially allows the remaining outputs to be later in time than the input. By placing a capacitor ($3 \times n$ pF) on the feedback line, the outputs can be moved forward in time.




Figure 5. ICS8752i Used with Trace Skew Adjustments

Application Example

In the following section, three examples are presented. The first (Figure 6) shows zero delay while the second (Figure 7) and the third (Figure 8) show delay adjustment of CLKOUT output and CLKIN respectively.

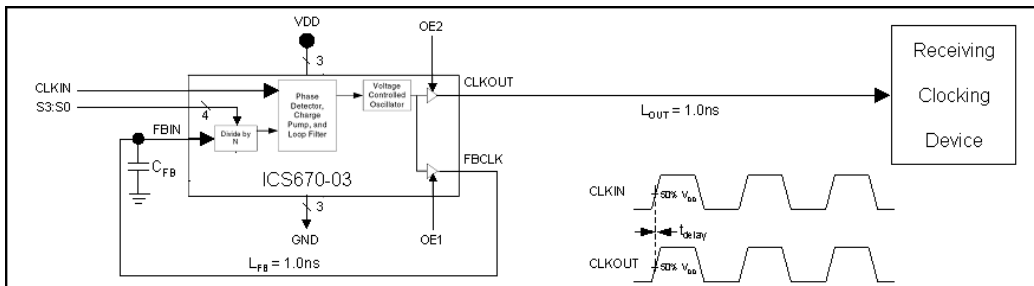


Figure 6. $L_{OUT} = L_{FB}$; Zero Phase Delay Between CLKIN and CLKOUT

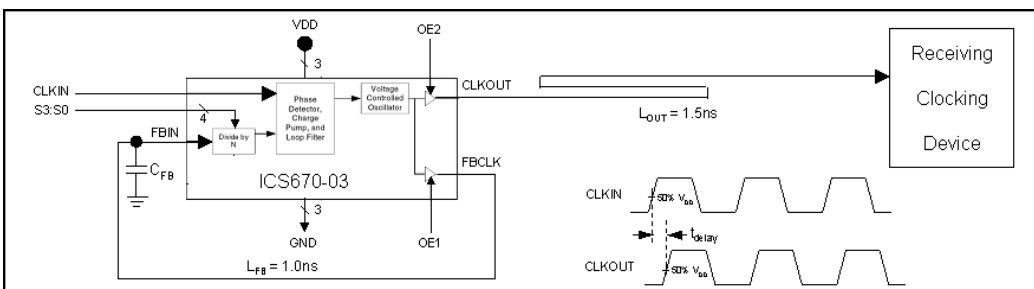


Figure 7. L_{OUT} Longer than L_{FB} by 0.5 ns

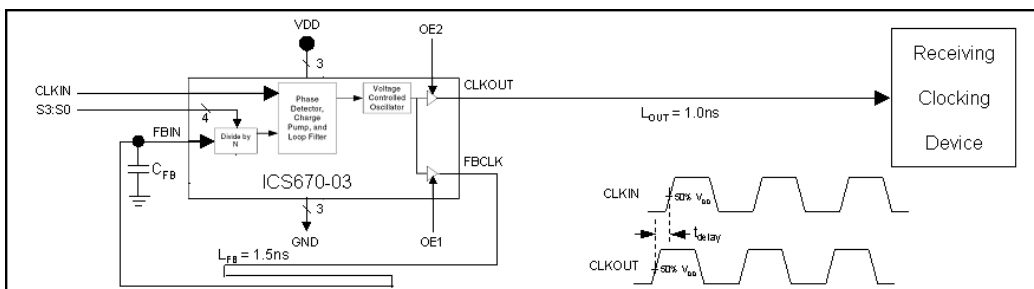


Figure 8. L_{FB} Longer than L_{OUT} by 0.5 ns

One of the systems designer's goals, when implementing a ZDB, is to match the edges of CLKIN and CLKOUT so that there is zero delay (no phase shift) between the input signal and the output signal. This is achieved by matching the feedback trace length with the output clock trace length as shown in Figure 6. However, in practical implementation, this may be difficult to achieve.

A possible solution is to design the feedback trace length to be slightly shorter than the output clock trace length as shown in Figure 7. In the lab, the output waveforms will be slightly delayed and skewed to the right of the input waveform. As an example, the ICS670-02, configured in the 1X mode ($F_{CLKOUT} = F_{CLKIN} = 50$ MHz), was tested in the lab with a feedback trace length which was shorter than the output trace length and the CLKOUT signal was delayed by 700ps. The measurements are shown in Figure 9.




Figure 9. CLKIN and CLKOUT with Shorter Feedback Trace

The user can then tune the ZDB and adjust the timing by adding (or increasing) the feedback loading capacitance accordingly. In the ICS670-02 example, a 10pF capacitance was added to the feedback capacitance and the CLKOUT signal better aligned with the CLKIN signal as shown in Figure 10.




Figure 10. CLKIN and CLKOUT After Tuning the ICS670-02

Care should be taken when designing with ZDB to ensure that the feedback trace length does not exceed the output trace length. If the feedback trace length is longer than the output trace length, the CLKOUT signal will be “early” and skewed to the left of the CLKIN signal as shown in Figure 8. In this case, adding a feedback capacitance will only increase the skew time and CLKOUT will drift more to the left of CLKIN, increasing the timing difference. Therefore, it is impossible to tune the ZDB and adjust the timing with a feedback capacitance if the feedback trace length is longer than the output trace length.

Layout Guidelines

As the device pin density and system frequency increases, the printed circuit board (PCB) layout becomes more complex and critical. A successful high-speed board must effectively integrate devices and other elements while avoiding signal transmission problems associated with high-speed I/O standards. Since IDT zero delay buffers include a variety of high-speed features, including fast I/O pins and high edge rates, it is imperative that an effective design successfully achieve the following:

- Reduces system noise by filtering and evenly distributing power to all devices
- Matches impedance and terminates the signal line to diminish signal reflection
- Minimizes crosstalk between parallel traces
- Reduces the effects of ground bounce.

These high-speed design goals are achieved by adhering to the following board design guidelines:

- Filter any noise that can interfere with device PLL components such as the VCO and phase detector and cause output clock jitter. Each VDD pin must be decoupled individually, by placing $0.1\mu\text{F}$ and $10\mu\text{F}$ capacitors physically close to the VDD pin, to prevent power supply noise generated by one device circuit block from interfering with another circuit block.
- Isolate the power planes of the clock driver from the power plane of the board by a ferrite bead. The ferrite blocks high frequency noise from reaching the main power supply. Use a non-resonant, surface-mount ferrite bead large enough to handle the current in series with the power supply and place a $10\mu\text{F}$ to $100\mu\text{F}$ bypass tantalum capacitor next to the ferrite bead (see Figure 11).




Figure 11. Filtering Noise with a Ferrite Bead

- Place power supply decoupling capacitors and filter components as close to VCC as possible. For decoupling, it is recommended to use low-inductance, low-ESR (Equivalent Series Resistance) capacitors as they provide best performance. The capacitors filter low-frequency noise from the power supply and supply extra current when many outputs switch simultaneously in a circuit.
- Keep power and ground planes adjacent. This reduces power supply noise.
- Create a localized ground plane, where possible, for the clock trace on the PCB top layer. These localized planes provide a return path for RF current.
- Configure unused I/O pins as output pins, and drive the outputs high or low. In this configuration, the pins may act as virtual power or virtual ground pins. The unused output pins that are driven high may be connected to VCCIO to prevent VCC sag and the unused output pins that are driven low may be connected to GND on the board, to reduce ground bounce.
- Keep trace impedance short and balanced to minimize reflections and ringing.

- Reduce crosstalk in the layout by widening the spacing between signal lines as much as routing restrictions will allow, trying not to bring traces closer than three times the dielectric height and minimize parallel run lengths between single-ended signals. Route with short parallel sections and minimize long, couple sections between nets. The trace conductor should be as close to the ground plane as possible to couple the transmission line tightly to the ground plane and help decouple it from adjacent signals.
- Use differential routing techniques where possible, especially for critical nets (i.e. match the lengths as well as the turns that each trace goes through).
- Minimize EMI by avoiding the use of vias to route clock signals. Vias add unwanted inductance to the trace and in general reduce the effectiveness of bypass capacitors. Further, vias can cause impedance change and reflections.
- Place clock signals far away from I/O area, but avoid routing traces near the edge of the PCB.
- Route clock traces to each load when driving multiple loads and terminate them individually. Load all outputs equally. Clock traces should not intersect with each other. The clock traces should be of equal length to minimize clock skew.
- Avoid using right angle bends as they increase capacitive loading of the transmission line and, in turn, causes unwanted reflection, ringing, overshoot and undershoot, and degraded signal rise and fall times. Instead, use smooth rounded or chamfered traces.

In Conclusion

There are four methods that are used to adjust the phase in PLL-based clock buffers:

1. By adjusting the feedback trace length.
2. By varying the capacitive load on the feedback pin.
3. By varying the capacitive load on the output clock.
4. By adjusting the trace on the output clock.

This report discussed all four methods.

References:

1. Johnson, H. W., and Gram, M. "*High-Speed Digital Design*". Prentice Hall, 1993.
2. ICS Application Report, "*Ferrite Bead Recommendations*", Rev. 101404.

Innovate with IDT and accelerate your future networks. Contact:

www.IDT.com

For Sales

800-345-7015
408-284-8200
Fax: 408-284-2775

For Tech Support

www.idt.com/go/clockhelp

Corporate Headquarters

Integrated Device Technology, Inc.
www.idt.com

