XILINX®

XAPP225 (v1.3) February 18, 2009

Data to Clock Phase Alignment

Author: Nick Sawyer

Summary

When designing digital systems, there is often a requirement to synchronize incoming data and clock signals with an internal system clock, i.e., the internal and external clock are at exactly the same frequency, but due to variable backplane, board, or application-specific standard product (ASSP) delays, the phase relationship is not known. The circuit described in this application note addresses this issue for both single traces and data buses up to 210 MHz in a Virtex®-II -5 device. The speed limitation is imposed by the maximum frequency that can be accepted by the Digital Clock Manager (DCM), in a mode where it is capable of providing both a new clock and a new clock shifted by 90 degrees. The maximum speed would be 210 MHz in a Virtex-II -5 device or a Virtex-II Pro -6 device, and around 160 MHz in a Spartan®-3 generation device. A typical application is shown in Figure 1.

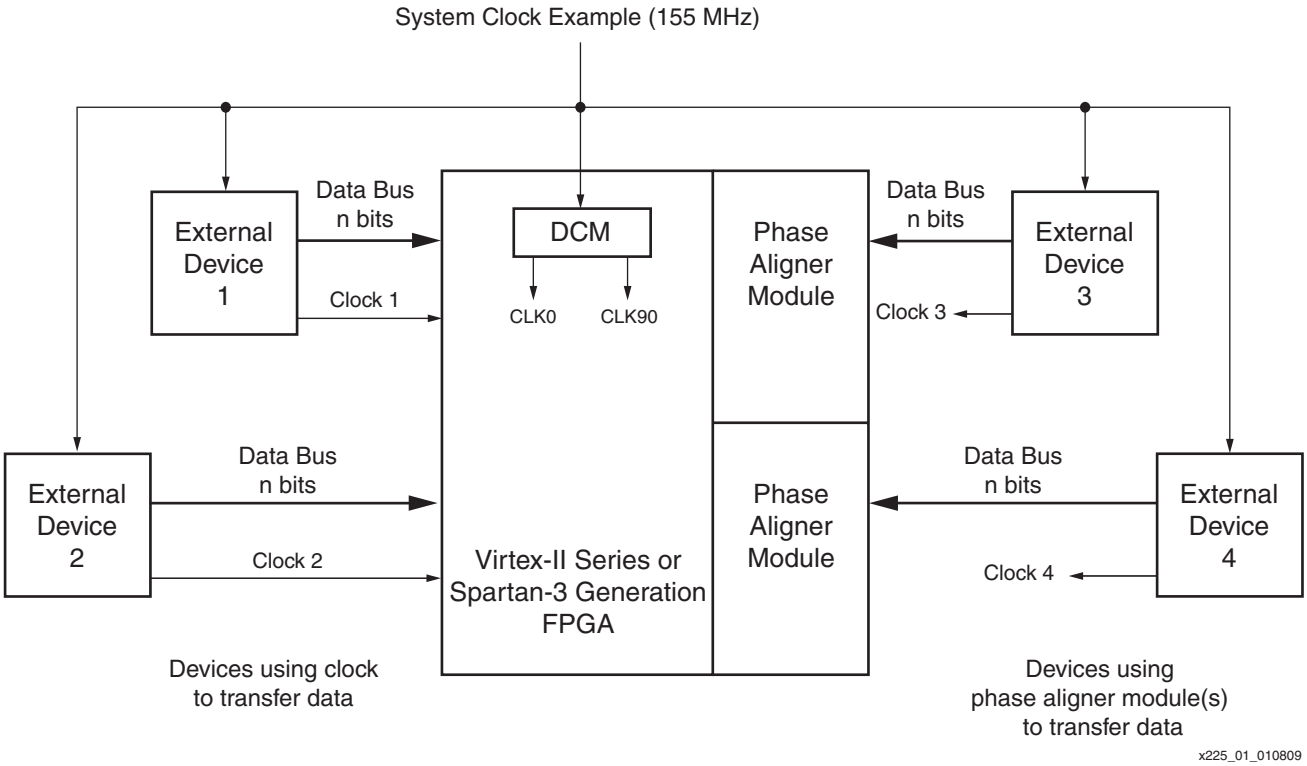

Figure 1: **Typical Application of Data to Clock Phase Alignment**

© 2002–2009 Xilinx, Inc. XILINX, the Xilinx logo, Virtex, Spartan, ISE, and other designated brands included herein are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.

Introduction

There are two classical methods of achieving synchronization. One is to sample the data with four different clocks and decide which of the four possible clocks is most valid. A multiplexer is used to provide the most valid clock to the rest of the required logic running at the communications link speed. However, a large percentage of the global buffers available in some of the Spartan-3 generation devices would be used up immediately with this method. In addition, inserting a multiplexer in the clock path effectively disables the DCM clock zeroing function. The Virtex-II series is improved with sixteen global clock buffers. However, using a global clock buffer for each incoming clock can be an inefficient use of system resources at low clock speeds.

The other method (described in detail here) is to sample the data on four phases of the same clock, decide which phase is the most "valid", and then resynchronize the data to the system clock. This method is more resource efficient, requiring only two global buffers and one DCM.

The incoming system clock is fed to a DCM component, and the DCM CLK0 is used to provide a clock (CLK) for the synchronizer circuit, as well as feedback for the DCM. There is also a version of the input clock delayed by 90 degrees (CLK90), and synchronized with the original clock. These waveforms are shown in the Figure 2 timing diagram along with the four possible data arrival cases used in the next section.

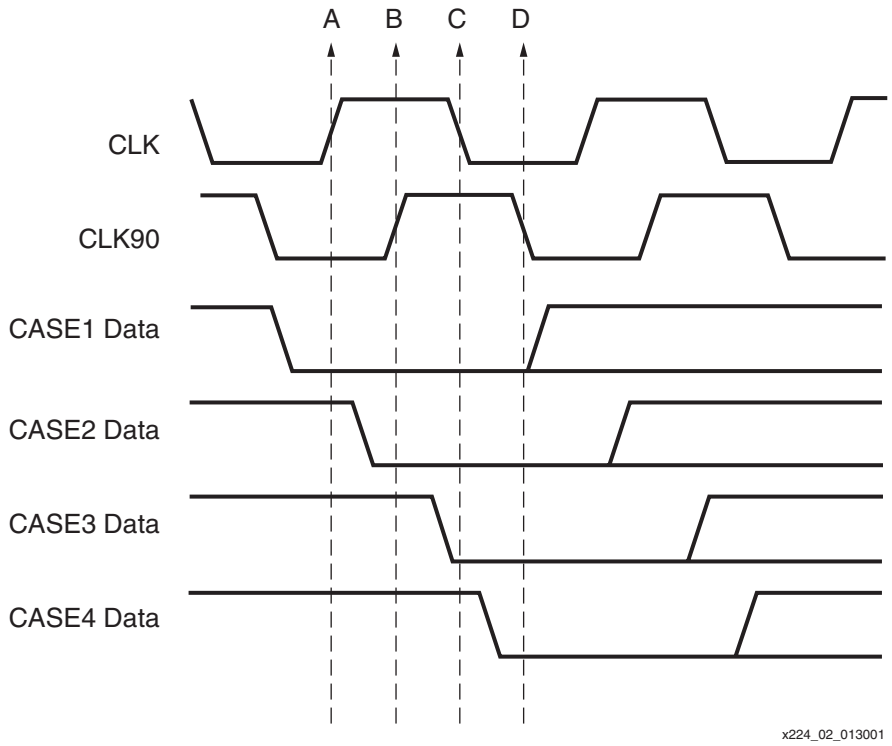

Figure 2: **Timing Diagram**

As shown in Figure 3, the incoming data is applied to four flip-flops, two clocked by CLK (one rising edge and one falling edge), and two by CLK90 (rising and falling edges). It is important that the delay from the input pin to these four flip-flops be almost equal. This is easily achieved by giving the software a MAXSKEW parameter for this net, of 500 ps, for example. The absolute delay is irrelevant; only the skew is important. However, when dealing with data buses, it is also prudent to apply a MAXDELAY parameter of approximately 1.0 ns to each net in the bus.

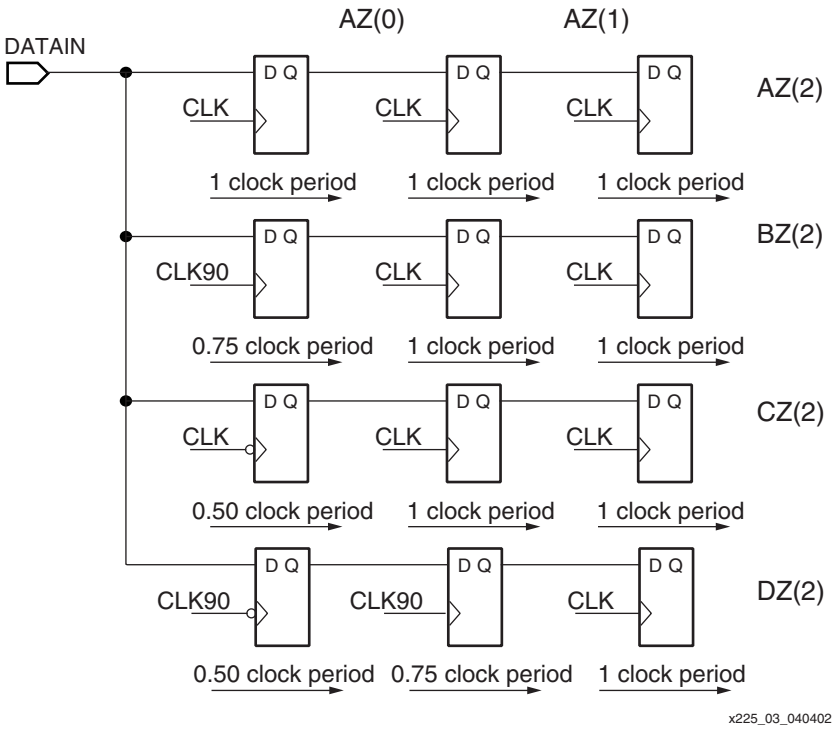

Figure 3: **Input Stage**

The first flip-flop is clocked by the rising edge of the clock described as time domain A. The second flip-flop is clocked by the rising edge CLK90 (time domain B); the third flip-flop is clocked by the falling edge of CLK, (time domain C); and the fourth is clocked by the falling edge CLK90, (time domain D). As shown in the timing diagram (Figure 2), this gives four data sample points, each separated by 90 degrees of the original clock frequency. In the case of a 160 MHz system clock, this logic is effectively running at 640 MHz.

These four sample points are then clocked once more, to remove any metastability issues and to move them into the same time domain. This actually takes place in two stages (again to avoid any four times clock frequency logic paths).

In the decision stage, shown in Figure 4, the circuit detects transitions on the data lines. The signals AAP to DDP recognize positive transitions, and the signals AAN to DDN recognize negative transitions. Eight signals are now available for the decision process. Four mutually exclusive signals can now be decoded, where only one transitions High whenever there is a data transition. These four conditions are as follows:

1. AAP = BBP = CCP = DDP = 1, or AAN = BBN = CCN = DDN = 1. Time domain A was the first to see the data. Therefore, the data from C is used for forwarding. This condition is indicated by the signal USEC in Figure 5 being High.
2. AAP = 1 and BBP = CCP = DDP = 0 or AAN = 1 and BBN = CCN = DDN = 0. Time domain B was the first to see the data. Therefore, the data from D is used for forwarding. This condition is indicated by the signal USED in Figure 5 being High.
3. AAP = BBP = 1 and CCP = DDP = 0, or AAN = BBN = 1 and CCN = DDN = 0. Time domain C recognized the transition first. Use the data clocked in during time domain A for forwarding into the system. This is the data that has been sampled midway through its period, i.e., the best noise margin. This condition is indicated by the signal USEA in Figure 5 being High.
4. AAP = BBP = CCP = 1 and DDP = 0, or AAN = BBN = CCN = 1 and DDN = 0. Time domain D was the first to see the data. Therefore, the data from B is used for forwarding. This condition is indicated by the signal USEB in Figure 5 being High.

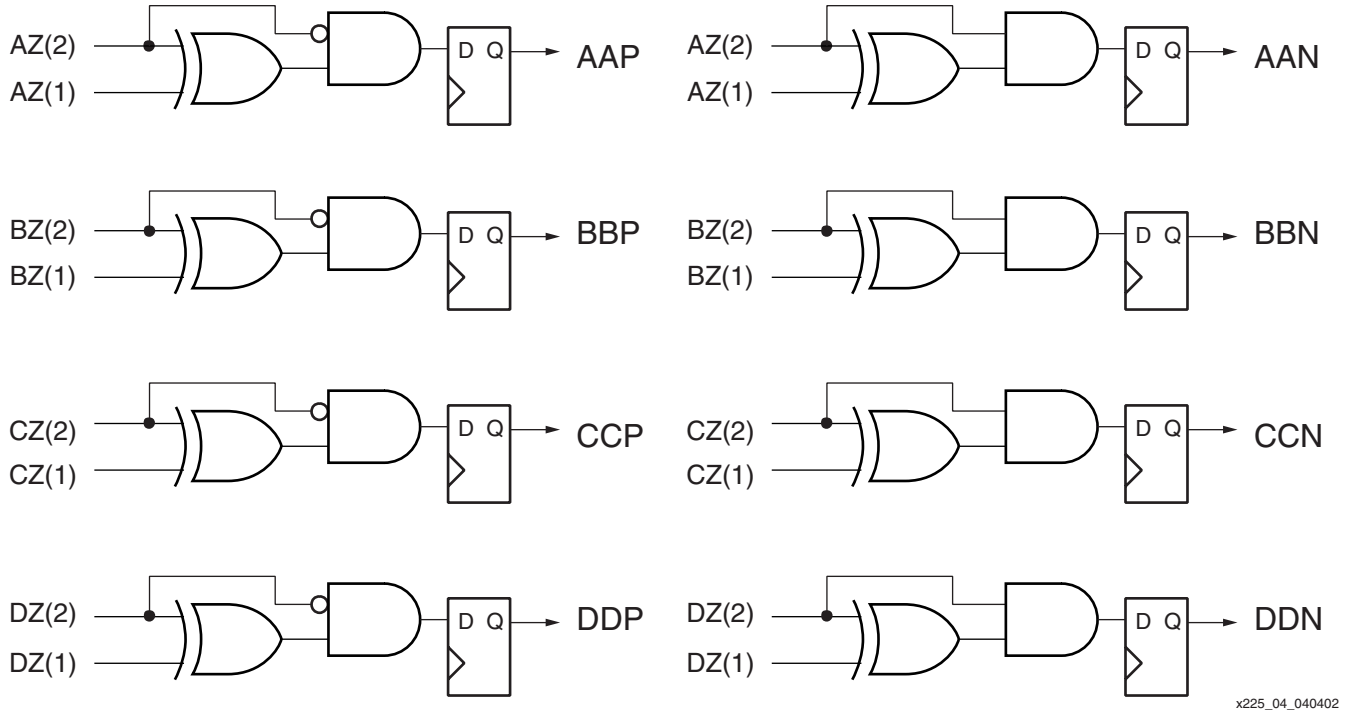

Figure 4: **Decision Stage**

The selection of data is done with a very simple multiplexer used to select data from the appropriate time domain (Figure 5). The complexity increases, in the real world, since the data is not fixed in time, but moves slightly with temperature and voltage changes. That is, it is quite possible that the circuit will synchronize with one time domain, but will resynchronize later to another. At 160 MHz, this implies that the data could have “moved” at least 1.6 ns. This also illustrates another use for the circuit, in bus systems, where clock delays will change as boards are added and removed. The circuit will automatically compensate for these delays.

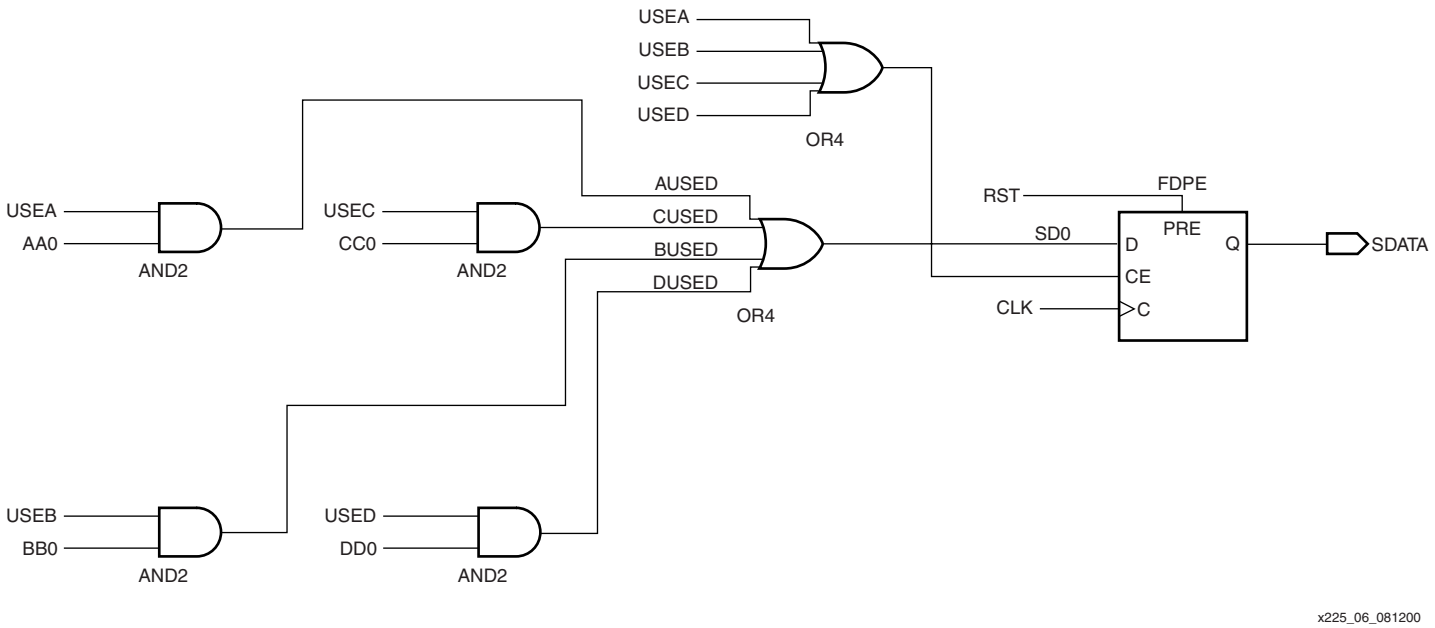

Figure 5: **Data Multiplexer**

As the data “moves”, the multiplexer is used to select the data from one of the four available time domains, with one exception. If the data moves from domain A to domain D, then effectively one stage of delay to the circuit is added, i.e., a “wraparound” occurs. This is done by detecting this transition (and also the opposite, i.e., going from D to A) and using a variable length shift register. This is easily achieved by using the SRL16 primitives in the Virtex series. When the circuit starts up, the SRL16s are programmed to be a 3-bit delay, if the transition mentioned above (A to D) is detected, then the SRL16s are reconfigured to be a 4-bit delay, or a 2-bit delay for D to A (Figure 6).

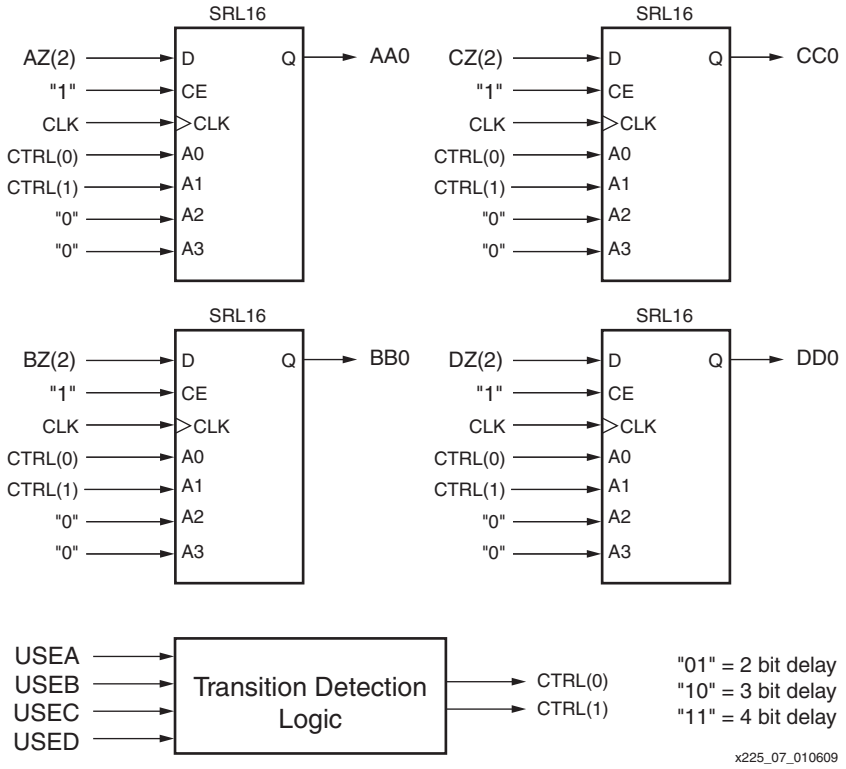

Figure 6: **Variable Length Shift Register**

If no transition has occurred, the circuit is locked, i.e., the selection logic does not change state, and if the transition wanders into another phase, then the selection logic is changed appropriately, as described above.

Metastability

Since it is possible that the data will not properly change state between clock sample points, there is a possibility of metastability. As the data transition approaches a sample point (for example, clock sample A in Figure 2), it will eventually end up inside the setup time to the clock CLK, causing one of several outcomes to occur.

If the flip-flop is fast enough to still see the transition, then AAx = BBx = CCx = DDx = 1, (case 1) and all will work as previously described. If the flip-flop does not see the transition, then AAx = 1 and BBx = CCx = DDx = 0, (case 2) and again all will work properly.

Finally, the flip-flop could briefly enter a metastable state. If this occurs, then the second synchronizing flip-flop will still “see” a 1 or 0 and will register that state, leading to the same arguments as above. There is no problem as long as the load on the potentially metastable flip flop is one, except when the metastable period is *exactly* equal to the input clock period. However, this event is extremely unlikely with today’s very fast silicon, and even then, there are three further registers in the data path. Therefore, the chance of a metastable event upsetting the operation of the circuit is vanishingly small.

Conclusion

Virtex-II series and Spartan-3 generation devices can be used to interface to external components running off the same clock but with an unknown phase relationship. This phase relationship can also be allowed to vary up to plus-or-minus one clock cycle in operation without affecting data integrity. When using the phase aligner described in this application note, data transfers can therefore occur at speeds up to 210 MHz without need of a clocked mechanism, which would waste clock resources.

Reference Design

The reference design circuit is implemented in HDL. It is fully synthesizable and comes in two variations. One variation is a master unit (`sync_master.vhd`), which will deskew one or two data lines and provide the control signals for further slave units. The slave modules (`sync_slave.vhd`) also deskew two data lines and because all the clocks are on global buffers the slave modules can be added to a design ad infinitum.

The reference design files include a `top.ucf` file, containing all the timing constraint information. It is important to use this file because some paths are very fast. Of course, in the case of buses, the data input delays have to be very closely matched. The best way to achieve this is to locate the bus input pins closely together.

The reference design files can be downloaded from:
https://secure.xilinx.com/webreg/clickthrough.do?cid=115124

The reference design matrix is shown in Table 1.

Table 1: **Reference Design Matrix**

Parameter	Description
General	
Developer Name	Xilinx
Target Devices (Stepping Level, ES, Production, Speed Grades)	Virtex-II, Virtex-II Pro, Spartan-3 Generation FPGAs
Source Code Provided	Yes
Source Code Format	VHDL, Verilog
Design Uses Code/IP from Existing Application Note, Reference Design, 3rd party, or CORE Generator™ software	No
Simulation	
Functional Simulation Performed	Yes
Timing Simulation Performed	Yes
Testbench Used for Functional and Timing Simulations	Yes
Testbench Format	VHDL and Verilog
Simulator Software/Version Used	ModelSim Xilinx Edition-III, version 6.2c
SPICE/IBIS Simulations	No
Implementation	
Synthesis Software Tools/Version Used	Xilinx Synthesis Technology (XST), version 10.1
Implementation Software Tools/Versions Used	ISE® Design Suite 10.1
Static Timing Analysis Performed	Yes

Table 1: **Reference Design Matrix *(Continued)***

Parameter	Description
Hardware Verification	
Hardware Verified	Yes
Hardware Platform Used for Verification	Virtex-II XLVDS demonstration board

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
09/18/00	1.0	Initial Xilinx release.
04/04/02	1.1	Revised for Virtex-II Series devices.
04/19/07	1.2	• Added Spartan-3 generation devices. • Figure 6 and accompanying text: Corrected bit delays and codes.
02/18/09	1.3	• Removed Virtex series FPGAs. • Replaced DLL with DCM. • Updated **Summary** and **Introduction**. • Updated Figure 1 and Figure 6. • Added reference design matrix (Table 1).

Notice of Disclaimer

Xilinx is disclosing this Application Note to you “AS-IS” with no warranty of any kind. This Application Note is one possible implementation of this feature, application, or standard, and is subject to change without further notice from Xilinx. You are responsible for obtaining any rights you may require in connection with your use or implementation of this Application Note. XILINX MAKES NO REPRESENTATIONS OR WARRANTIES, WHETHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, IMPLIED WARRANTIES OF MERCHANTABILITY, NONINFRINGEMENT, OR FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL XILINX BE LIABLE FOR ANY LOSS OF DATA, LOST PROFITS, OR FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR INDIRECT DAMAGES ARISING FROM YOUR USE OF THIS APPLICATION NOTE.