



PLL Design Techniques and Usage in FPGA Design

XBRF 006 August 28,1996 (Version 1.1)

Application Brief by Steve Sharp

Summary

This paper examines some general concepts concerning Phase Locked Loop (PLL) usage and their application in programmable logic devices. A critique of a newly-announced PLL implementation for FPGAs also is included.

Xilinx Family

Any

General PLL Usage

PLLs are primarily used for high performance designs. In large, fast gate arrays, the performance limiting element often is clock delay due to large clock networks. PLLs help reduce this clock delay and, thereby improve performance.

Figure 1 shows a simplified view of a chip with two internal clock tree branches. The goal is to minimize clock skew by having a near equal delay in each branch. As chip size grows, the number of branches is increased, and buffer sizes are tuned to equalize the delays between branches. This is often known as an "H tree" clock distribution scheme.

Figure 2 shows the relationship between total clock delay and skew, or the difference in delay between the branches of the clock tree. Both total clock delay and clock skew are factors in high performance systems.

The basic concept of a PLL is simple. By anticipating the edges of the input, the PLL can generate new clocks with edges slightly earlier than the input clock (The input clock must be of a stable frequency for this to work correctly). By tuning the amount of time these new clock edges precede the input clock edge to the delays of the various clock tree branches, all registers will see the clock edge at about the same time. This reduces the overall clock network delay as well as minimizing clock skew.




Figure 1: Simplified Clock "H Tree"

Figure 3 shows the internal clocks (dotted lines) being generated ahead of the input clock, so the actual internal clocks arrive at their destinations much closer to the input clock edge. There will still be some clock skew due to the delay in each branch (the signal takes a finite amount of time to propagate from the buffer to the end of the branch), but overall clock performance is improved.




Figure 2: Clock Network Without PLL




Figure 3: Clock Network With PLL

PLLs in FPGA Devices

As stated previously, a PLL locks onto an external clock and generates an internal clock with edges ahead of the edge of the external clock. This has the net effect of a reduction in clock delay, and therefore faster tco and tsu delays.

Altera announced Phase Locked Loop (PLL) support for the FLEX 10K FPGA family on May 27, 1996. This support is scheduled to be included in the 10K100 device in the September time frame. Altera refers to these new PLL features as "ClockLock" and "ClockBoost". ClockLock is the name for Altera's basic PLL function, while ClockBoost refers to their frequency multiplication capability. These two functions have been available in stand-alone clock chips for some time.

Apparently, Altera believes that the size of the device, and not performance, is the main reason to use PLLs. The FLEX 10K100 is a large device where the performance is limited primarily by logic and routing delays. Adding a PLL to speed up clock delays will have little effect on overall performance.

Figure 4 shows that, for gate arrays, the use of a PLL to reduce clock network delay will improve performance in large devices. In large gate arrays, the individual logic and routing delays are still small enough that reducing the clock network delay has an overall positive effect. For the Altera FLEX 10K family, even routing from one logic array block (LAB) to an adjacent one in the same row costs 4.7ns of routing delay for a -3 device. Routing to a logic element in another row costs 11.9ns of routing delay. This "step function" in routing delay makes efforts to prune a few nanoseconds out of the clock network delay of minimal value.

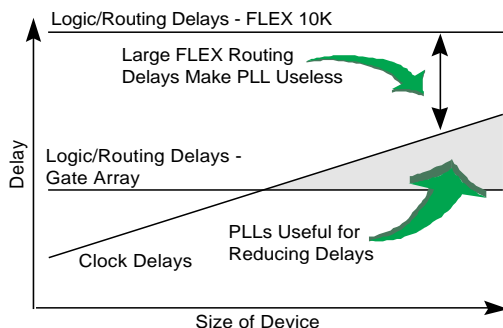


Figure 4: PLL Usage vs. Logic/Routing Delays

Altera's press release claims a reduction in clock-to-out delay from 9ns to 6ns and a reduction in setup time of "nearly in half" to 3.6ns, but does not state what speed

grade is associated with these claims. Xilinx has done tests that seem to indicate that these claims are reasonable, although the exact numbers cannot be verified due to the lack of information in the Altera press release. It appears that it would be possible to design a small amount of simple logic closely coupled to a few I/O pins that would realize an improvement in performance when using their PLL feature.

In the FLEX 10K family, however, the internal logic and routing delays typically are the limiting factor in system performance and not the clock delays. To examine this further, several test designs implemented in the FLEX 10K50 device were analyzed using both -4 and -3 speed data (advance data for both grades). Examining these more complex designs gives a more realistic picture of FLEX 10K performance.

Internal Routing Delays Limit FLEX 10K Performance

A simple 16-bit down counter with synchronous load was examined. This counter runs at 104 MHz in a -3 speed grade when using the carry chain, but at only 48 MHz without the carry chain. Routing the feedback paths without the carry chain represents a more accurate reflection of the performance of real-world designs, where logic paths span several logic levels and use more of the local routing.

Test circuits containing 32 16:1 multiplexers, a 16-bit parity generator, and 4 16-bit adders were also examined. These designs, which performed as shown in Table 1 for a -3 speed grade, also show that designs with real-world data path routing consistently perform in the sub-50 MHz range, rather than the 100 MHz range as Altera claims.

Table 1: Data Path Test Results for FLEX 10K10-3

Design	Delay (ns)	fmax (MHz)
32 copies, 16:1 MUX	26ns	38.5 MHz
16-bit parity generator	24ns	41.7 MHz
4 copies, 16 bit adder	23.1ns	43.3 MHz

The Altera FLEX 10K data sheet includes a parameter called tDRR. This parameter is stated to be "Register-to-register delay via 4 LEs, 3 row interconnects, and 4 local interconnects". The data sheet goes on to note that "A representative subset of signal paths is tested to approximate typical device applications." For the -4 speed grade this tDRR parameter is 23.8ns. This corresponds to a system frequency of 42 MHz. This further verifies that for designs that use multiple look up tables and local signal routing between logic array blocks, the performance of the FLEX 10K family rapidly falls into the sub-50 MHz region.

The clear conclusion from this data is that any performance improvement due to PLL use is likely to be insignificant in this frequency range (below 50 MHz).

Frequency Multiplication

Another use of a PLL is to perform frequency multiplication. This allows the use of a slower external clock while still achieving a faster internal clock. One benefit of slower external clocks is better noise immunity on the board. Altera offers X2, X3, and X4 options for their "ClockBoost" feature.

For most board designs, the clock noise effects are not a significant issue until frequencies exceed 50 MHz. At the minimum "X2" multiplication factor, this would translate to a 100 MHz internal clock rate. Even for the fastest (and as yet unavailable) speed grades, the 10K100 is limited by its slow logic/routing delays to 30-50 MHz clock rates; it has little use for a 100 MHz internal clock. Similarly the X3 and X4 options are of little use as well.

Summary

- PLLs are useful and are used frequently for high performance gate arrays to improve clock network delays and skew.

- Altera's claims in their PLL press release for the FLEX 10K100 FPGA seem to be substantiatable for setup/hold and clock-to-out times, but appear to be useful only for signals interfacing to the I/O pads and closely coupled to small blocks of logic.
- Internal logic and routing delays in the FLEX 10K100 FPGA are so slow that improvement in the clock delay is meaningless for most real applications with significant data flow.

With the restrictions Altera has imposed on PLL usage, and with the limitations in their logic and routing performance, the benefits to a user of their PLL become hard to find. Offering a performance feature first on a slow FLEX 10K100 FPGA device appears either misguided or a complete misunderstanding of PLL usage.

Xilinx is the industry leader in high density and high performance FPGAs, and will offer a PLL in a high-performance product in the near future. It will be the right implementation for a high-performance system, and one that will provide clear benefits to the user.



Headquarters

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124
U.S.A.

Tel: 1 (800) 255-7778
or 1 (408) 559-7778
Fax: 1 (800) 559-7114

Net: hotline@xilinx.com
Web: <http://www.xilinx.com>

North America

Irvine, California
(714) 727-0780

Englewood, Colorado
(303) 220-7541

Sunnyvale, California
(408) 245-9850

Schaumburg, Illinois
(847) 605-1972

Nashua, New Hampshire
(603) 891-1098

Raleigh, North Carolina
(919) 846-3922

West Chester, Pennsylvania
(610) 430-3300

Dallas, Texas
(214) 960-1043

Europe

Xilinx Sarl
Jouy en Josas, France
Tel: (33) 1-34-63-01-01
Net: frhelp@xilinx.com

Xilinx GmbH
Aschheim, Germany
Tel: (49) 89-99-1549-01
Net: dlhelp@xilinx.com

Xilinx, Ltd.
Byfleet, United Kingdom
Tel: (44) 1-932-349401
Net: ukhelp@xilinx.com

Japan

Xilinx, K.K.
Tokyo, Japan
Tel: (03) 3297-9191

Asia Pacific

Xilinx Asia Pacific
Hong Kong
Tel: (852) 2424-5200
Net: hongkong@xilinx.com

© 1996 Xilinx, Inc. All rights reserved. The Xilinx name and the Xilinx logo are registered trademarks, all XC-designated products are trademarks, and the Programmable Logic Company is a service mark of Xilinx, Inc. All other trademarks and registered trademarks are the property of their respective owners.

Xilinx, Inc. does not assume any liability arising out of the application or use of any product described herein; nor does it convey any license under its patent, copyright or maskwork rights or any rights of others. Xilinx, Inc. reserves the right to make changes, at any time, in order to improve reliability, function or design and to supply the best product possible. Xilinx, Inc. cannot assume responsibility for the use of any circuitry described other than circuitry entirely embodied in its products. Products are manufactured under one or more of the following U.S. Patents: (4,847,612; 5,012,135; 4,967,107; 5,023,606; 4,940,909; 5,028,821; 4,870,302; 4,706,216; 4,758,985; 4,642,487; 4,695,740; 4,713,557; 4,750,155; 4,821,233; 4,746,822; 4,820,937; 4,783,607; 4,855,669; 5,047,710; 5,068,603; 4,855,619; 4,835,418; and 4,902,910. Xilinx, Inc. cannot assume responsibility for any circuits shown nor represent that they are free from patent infringement or of any other third party right. Xilinx, Inc. assumes no obligation to correct any errors contained herein or to advise any user of this text of any correction if such be made.