



January 2006 

I/O magazine

CONNECTIVITY SOLUTIONS FOR PROGRAMMABLE LOGIC PROFESSIONALS



INSIDE

A Paradigm Shift in Signal Integrity and Timing Analysis

Debugging and Validating PCI Express I/O


Understanding the PCI-SIG Compliance Program

How to Detect Potential Memory Problems Early in FPGA Designs

A New PCI Express Solution Simplifies Video Security Applications

 **XILINX®**

Support Across The Board.™



Design Kits Fuel Feature-Rich Applications

Build your own system by mixing and matching:

- Processors
- FPGAs
- Memory
- Networking
- Audio
- Video
- Mass storage
- Bus interface
- High-speed serial interface

Available add-ons:

- Software
- Firmware
- Drivers
- Third-party development tools

Avnet Electronics Marketing designs, manufactures, sells and supports a wide variety of hardware evaluation, development and reference design kits for developers looking to get a quick start on a new project.

With a focus on embedded processing, communications and networking applications, this growing set of modular hardware kits allows users to evaluate, experiment, benchmark, prototype, test and even deploy complete designs for field trial.

Gain hands-on experience with these design kits and other development tools by participating in a SpeedWay Design Workshop™ this spring.

For a complete listing of available boards, visit

www.avnetavenue.com

For more information about upcoming SpeedWay workshops, visit

www.em.avnet.com/speedway



Enabling success from the center of technology™



1 800 332 8638
em.avnet.com

I/O magazine

EDITOR IN CHIEF
Carlis Collins
carlis.collins@xilinx.com
408-879-4519

EXECUTIVE EDITOR
Forrest Couch
forrest.couch@xilinx.com
408-879-5270

MANAGING EDITOR
Charmaine Cooper Hussain

ONLINE EDITOR
Tom Pyles
tom.pyles@xilinx.com
720-652-3883

ART DIRECTOR
Scott Blair

ADVERTISING SALES
Dan Taie
1-800-493-5551



Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124-3400
Phone: 408-559-7778
FAX: 408-879-4780

© 2006 Xilinx, Inc. All rights reserved. XILINX, the Xilinx Logo, and other designated brands included herein are trademarks of Xilinx, Inc. PowerPC is a trademark of IBM, Inc. All other trademarks are the property of their respective owners.

The articles, information, and other materials included in this issue are provided solely for the convenience of our readers. Xilinx makes no warranties, express, implied, statutory, or otherwise, and accepts no liability with respect to any such articles, information, or other materials or their use, and any use thereof is solely at the risk of the user. Any person or entity using such information in any way releases and waives any claim it might have against Xilinx for any loss, damage, or expense caused thereby.

Making Sense of the Complex

Welcome to the second edition of *I/O Magazine*, the premier educational journal of I/O technology from Xilinx. This magazine was created for practicing engineers in the semiconductor and electronic design communities, with an emphasis on design challenges and solutions.

Gone are the days when FPGAs were used only for glue logic functions. Today's FPGAs perform central functions in a majority of systems in the communications, computing, storage, consumer, and automotive industries. Following Moore's law, advanced devices such as Xilinx® Virtex™-4 FPGAs are shipped with integrated 10 Gigabit transceivers, Ethernet MACs, and thousands of I/Os, able to morph from LVDS to HSTL to LVCMOS with the flip of a bit and making these advanced technologies available at a cost point previously unthinkable. If the past is any indication, next-generation FPGAs will bring even more capabilities to the design community.

Designing with such advanced technologies is incredibly exciting and always challenging. Rather than completing only a digital design, most designers now must deal with PC board and connector design and signal and power integrity issues. To successfully complete your projects, you must constantly update your knowledge – and what better way to do that than to learn from the people who designed these technologies? Xilinx and its partners are committed to helping you learn – and *I/O Magazine* is an excellent way to achieve that goal.

In this issue, you will find articles on relevant design issues such as PCI Express, memory interfaces, signal integrity, and PC board design. You will also find useful information about tools, IP, and training classes that can help you complete your design on time.

Thank you and happy reading!



Abhijit Athavale
Sr. Marketing Manager,
Connectivity Solutions
Xilinx, Inc.



Are Your Tools Up to the Next Challenge?

Reach New Heights with Quantum-SI™

**SiSoft's
High-Speed Design Closure™
Delivers First-Pass Success.**

Quantum-SI™ is the only system-level signal integrity tool that can deliver true High-Speed Design Closure™ by bringing together **signal integrity, timing, crosstalk and rules-driven design...**

...all in a single solution.

SiSoft can provide your organization a growth path to the future because our software incorporates the needs of our own signal integrity consultants who are solving next generation problems today. When you invest in SiSoft products you can be certain that you are investing in your future designs as well.

SiSoft provides software, design analysis kits and second-to-none consulting services. **Quantum-SI's Core-to-Core™** methodology enables our software to more accurately predict system-level noise and timing margins. **Quantum-SI** incorporates signal integrity, timing and crosstalk analysis with unparalleled accuracy, simulation capacity, and functionality. Only **Quantum-SI** integrates the capabilities necessary for **High-Speed Design Closure**, the key to achieving first-pass success.

To learn more about SiSoft's products and services or to request a product demo, visit us on the web at www.sisoft.com, or send email to info@sisoft.com.



CONTENTS

ARTICLES



A Paradigm Shift in Signal Integrity and Timing Analysis	6
Capturing Data from Gigasample Analog-to-Digital Converters.....	9
Xilinx/Micron Partner to Provide High-Speed Memory Interfaces	14
Implementing High-Performance Memory Interfaces With Virtex-4 FPGAs.....	16
Debugging and Validating PCI Express I/O	20
Using Complex Triggers in the Identify Debugger.....	24
Understanding the PCI-SIG Compliance Program	28
Successful DDR2 Design.....	31
Board Design Panacea	36
Deliver Efficient SPI-4.2 Solutions with Virtex-4 FPGAs	39
A Low-Cost PCI Express Solution	42
How to Detect Potential Memory Problems Early in FPGA Designs	44
Taking Rugged I/O Cabling and Connectors to Higher Speeds	48
A New PCI Express Solution Simplifies Video Security Applications	52
Designing a Spartan-3 FPGA DDR Memory Interface.....	56

PRODUCT REFERENCE

10-Gigabit Ethernet MAC.....	58
Tri-Mode Ethernet MAC.....	59
Virtex-4 Embedded Tri-Mode Ethernet MAC Wrapper.....	60
XAUI.....	61
Memory Interfaces Reference Design	62
Interfacing QDR II SRAM with Virtex-4 FPGAs.....	65
Xilinx PCI Express Solution.....	66
Spartan-3 Generation IP.....	68

EDUCATION

Signal Integrity for High-Speed Memory and Processor I/O	71
PCI Express Design Flow	72
Designing with Multi-Gigabit Serial I/O.....	73



A Paradigm Shift in Signal Integrity and Timing Analysis

Emerging high-speed interfaces are breaking traditional analysis approaches, forcing a paradigm shift in analysis tools and methodology.

by Barry Katz
President and CTO
SiSoft
barry.katz@sisoft.com

Simplistic rule-of-thumb approaches to interface analysis are proving to be woefully inadequate for analyzing modern high-speed interfaces like DDR2, PCI Express, and SATA-II. This situation will only worsen when emerging standards like DDR3 and 5-10 Gbps serial interfaces become commonplace.

Signal integrity analysis performed on only the shortest and longest nets in a design may not identify the worst-case inter-symbol interference, crosstalk, or pin timing scenarios caused by variations in stub length, number of vias, routing layers, AC specifications, package parasitics, and power delivery. An integrated, interface-centric approach that incorporates comprehensive signal integrity, timing, crosstalk,

and power integrity analysis is required to more accurately predict system-level noise and timing margins.

Figure 1 offers the results of a simplistic versus comprehensive analysis approach to illustrate the shortcomings associated with some analysis tools, which are built on outdated rule-of-thumb methodologies and assumptions. The first waveform in Figure 1 represents a high-speed differential network using Xilinx® Virtex™-II ProX RocketIO™ IBIS models, lossless transmission lines, and ideal grounds with no crosstalk or power noise.

It is quite apparent from viewing the results that the simplistic analysis approach fails to provide the accuracy of the more comprehensive approach. The second waveform represents the progressive effect on the eye as a longer stimulus pattern is used, along with more accurate modeling of interconnect structures. The analysis also used detailed SPICE I/O models,

accounting for power delivery, crosstalk, non-ideal grounds, and variations in process, voltage, and temperature.

When designers are fighting for tens of picoseconds and tens of millivolts, an approach that considers all of the factors affecting margin (see Figure 2) is essential to ensure that a design will meet its cost and performance goals.

Model Interconnect Topologies and Termination Schemes

Accurate modeling of interconnect structures and termination – including the component packaging, PCBs, connectors, and cabling – is critical for accurate simulations of high-speed networks. As edge rates have increased and interconnect structures have remained relatively long, the importance of modeling frequency-dependent loss has become much more crucial, which requires the use of two- and three-dimensional field solvers. Given the potential for wide varia-




Figure 1 – Xilinx Virtex-II RocketIO transceiver simplistic versus comprehensive analysis

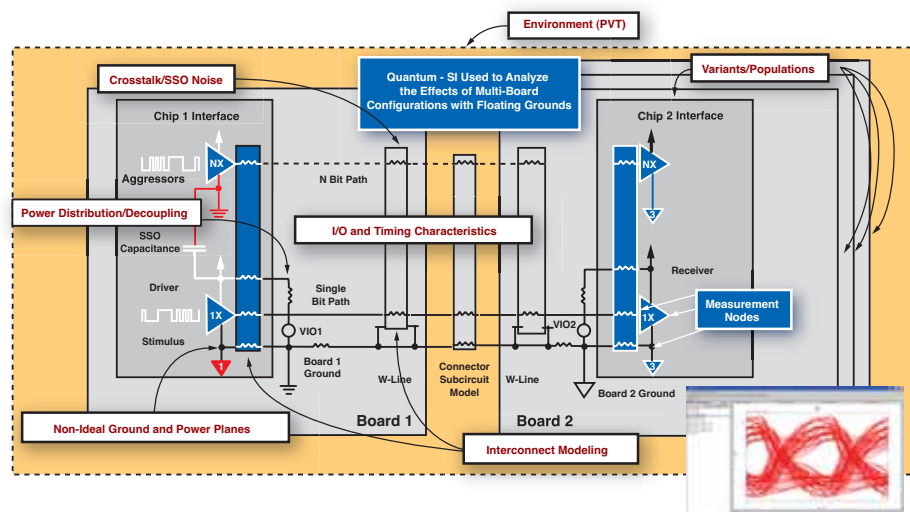


Figure 2 – Factors affecting system-level noise and timing margins

tion in the physical routing through packaging, PCBs, connectors, and cabling of many bus implementations, it is virtually impossible to identify the worst-case net without performing a comprehensive analysis on the entire interface.

Common analysis considerations that affect the analysis results include:

- Lossy versus lossless transmission lines
- Modeling vias as single- or multi-port structures
- Sensitivity to the number of vias in a net
- The use of two-dimensional distributed or three-dimensional lumped models for packages and connectors
- Modeling with S-parameters

Account for Inter-Symbol Interference

Traditional simulation approaches assume that signals are quiescent before another transition occurs. As the operating frequencies increase, the likelihood that a line has

not settled to its quiescent state increases. The effect on one transition from the residual ringing on the line from one or more previous transitions results in delay variations. These delay variations, called intersymbol interference, or ISI, require complex stimulus patterns that excite the different resonances of the network to create the worst-case scenarios. For some networks, these patterns may have a handful of transitions, but for multi-gigabit serial links, it is common to use long pseudo-random bit sequence (PRBS) patterns. Because the resonant frequency of a network is a function of the electrical length, the worst-case ISI effects may or may not occur on the shortest or longest net. In addition, interconnect process variations must be accurately accounted for, as this variation will cause changes in the resonant frequency (reflections) of the network.

Multi-gigabit serial link interfaces contain embedded clocks in the serial stream and use clock recovery techniques to

extract the serial data, which must meet stringent eye mask requirements. I/O buffer model accuracy that reflects pre-emphasis/de-emphasis and equalization is crucial for analyzing the effects of ISI.

Don't Forget the Effects of Crosstalk

Crosstalk is noise generated on a net from transitions on nearby interconnects in the circuit board, packages, connectors, and cables. Crosstalk can change the level of the signal on a net and therefore cause variations in the interconnect delays and reduce noise margins. Synchronous and asynchronous crosstalk are noise sources that must be fully analyzed to determine their effects on signal integrity and timing margins.

Model I/O Buffer Characteristics and Component Timing

I/O buffer electrical and timing characteristics play a key role in defining the maximum frequency of operation. A flexible methodology and automated analysis approach is required to support the wide variations in I/O technology models, including mixed IBIS and SPICE simulation. SPICE models are more accurate and very useful when simulating silicon-to-silicon. SiSoft implements this through its Core-to-Core Methodology, as shown in Figure 3. However, you should recognize that the improvement in accuracy comes at a price – a 5x to 100x simulation speed decrease.

Output buffers and input receivers are commonly characterized by numerous electrical/timing characteristics and reliability thresholds. These cells may include on-die termination, controlled impedances/slew rates, pre-emphasis, and equalization.

For high-speed parallel buses, data input timing is defined as a setup/hold time requirement with respect to a clock or strobe. Data output timing is defined by the minimum and maximum delay when driving a reference load with respect to a clock or strobe. With the advent of SSTL signaling, AC and DC levels were introduced for V_{il}/V_{ih} to more accurately characterize receiver timing with respect to an input signal. Further refinements have been made through slew rate derating (required for DDR2 and DDR3), which

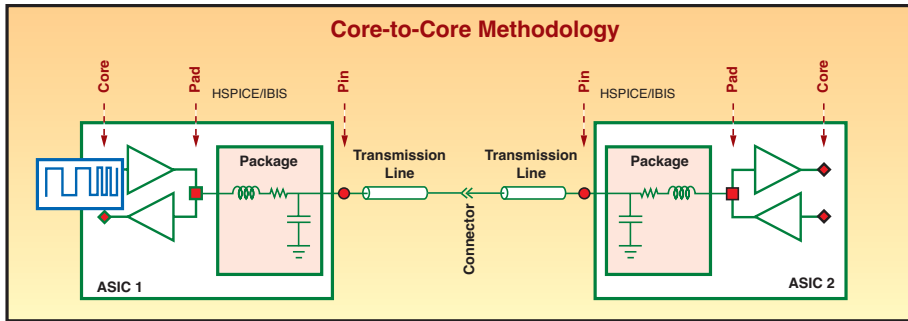


Figure 3 – SiSoft's Core-to-Core Methodology

uses tables to model the internal delay of a receiver at the core based on the slew rate at the pad. These refinements are not taken into account by simplified analysis approaches. This is why they cannot be used to accurately model the more complex behavior of many high-speed interfaces, where tens of picoseconds and tens of millivolts matter.

Don't Neglect PVT Variations

Many analysis tools and simplified methodologies neglect the effects of process, voltage, and temperature (PVT) variations, which can have disastrous results at high Mbps/Gbps signaling rates. It is especially important to consider IC process variations when modeling interconnect structures. Manufacturers typically supply data describing the AC specs and I/O buffer characteristics for fast, typical, and slow process parts, which bound the expected operating region. You should always analyze high-speed designs at the minimum/maximum operating extremes to avoid finding unpleasant surprises after the hardware is built.

Maintain Power Integrity

Maintaining the integrity of the power subsystems for both I/O and core power is critical. This requires analyzing stackups; PCB, package, and IC decoupling; routing layers and associated signal return paths. At a high level, the goal is to maintain a low impedance connection between associated voltage references across the operational frequency of interest. Simultaneous switching output (SSO) noise is commonly analyzed as part of power delivery to the I/O structures and also includes the effects of

package crosstalk. SSO is often quantified in terms of a timing uncertainty penalty applied to the AC timing specs of the chip.

Accurately Determine Setup and Hold Margins

Faster interfaces require maintaining very tight timing margins. Interfaces are typically classified as either synchronous (common-clock), source-synchronous, clock recovery, or a hybrid of these types. It is important that the clock distribution is accurately simulated and used in carefully correlated ways with data nets to accurately predict timing margins and optimal clock distribution. The integration of accurate signal integrity, timing, crosstalk, and rules-driven design is the basis of a new paradigm, which we call "High-Speed Design Closure."

Required Tools and Methodology Paradigms

To overcome the shortcomings of traditional analysis methodologies and inaccuracies associated with oversimplified rules-of-thumb, today's high-speed interface designers need to adopt a more comprehensive interface-centric system-level analysis approach that addresses many (if not all) of the issues discussed in this article.

High-quality I/O buffer models, interconnect models, and accurate component AC timing/electrical specifications are fundamental to any analysis approach. The process of capturing and managing multiple interface designs; performing comprehensive simulations over process, voltage, and temperature for a large solution space of variables; and analyzing the simulation results for waveform quality, timing, crosstalk, SSO, and ISI effects is a

daunting task without proper tools, which automate and integrate many manual steps and processes.

A highly automated analysis approach is also required to understand the loading effects associated with multi-board designs that include different board populations and part variants, and manage the complex set of variables within a multi-dimensional solution space. In pre-layout analysis, it is crucial to be able to mine the simulation results from different solution/space scenarios to pick an optimal solution for component placement and board routing.

Once the boards have been routed, it is equally important to verify the routed designs in the final system configuration, including different board populations and part variants to "close the loop" on signal integrity and timing. Accurate signal integrity analysis and crosstalk prediction in post-layout is essential to predicting system-level noise and timing margins.

With "High-Speed Design Closure," SiSoft is committed to providing tools for signal integrity, timing, crosstalk, and rules-driven design that meet rapidly changing signal integrity and timing requirements.

Conclusion

High-speed interface design and analysis complexity is only going to increase as edge rates and data rates get faster and voltage rails decrease. Engineering managers should recognize that setting up a high-speed interface analysis process requires an investment in simulation libraries, analysis products, and people.

When you invest in tools, do your homework first. Check to see if prospective tools can really address some of the tough issues presented in this article and that they provide you the growth path you need for the future. Perform thorough (and possibly lengthy) comparative evaluations of potential products to see if they address your current signal integrity, timing, power delivery, and crosstalk analysis needs, but also keep an eye to the future – it will arrive sooner than you think.

To learn more about SiSoft's products and services, visit www.sissoft.com or e-mail info@sissoft.com.

Capturing Data from Gigasample Analog-to-Digital Converters


Interfacing National Semiconductor's ADC08D1500 to the Virtex-4 FPGA allows quick-start customer application development.

by Ian King
Application Engineer
National Semiconductor
ian.king@nsc.com

Data conversion within the test and measurement domain and communications industry is moving into the gigasamples per second (GSPS) range. Developing a system capable of processing data at these speeds requires diverse engineering disciplines from the initial system concept through to board design, FPGA logic design, signal processing, and application software.

National Semiconductor has developed a leading-edge analog-to-digital (A/D) converter that can deliver as many as three billion samples per second to an 8-bit resolution. One of the main system design questions from customers regarding this product is how data can be reliably captured and processed at this speed. Therefore, National's applications team designed a development platform to provide a solution to this query and demonstrate a reliable data-capture method. This allows the design focus to shift away from the high-speed front end so that developers can focus on their intended application.

The platform also demonstrates that high clock speeds can be reached while maintaining low power dissipation sufficient for the entire system to be housed in a small enclosure, as would be required for a commercial or industrial system. In this article, I'll explain the techniques and analysis involved in achieving this goal.



Power Considerations

When selecting an FPGA for data capture that can achieve low power levels and performance, a 90 nm device is the first choice. In applications where data is captured in bursts (such as oscilloscopes and radar), the static power of the FPGA device becomes an important factor. This is because the high-speed data transfer between devices takes place over a very short time period, so the capture logic will be static while the application consumes the data.

Figure 1 shows a comparison of Xilinx® Virtex™-4 FPGA static power figures over device density. This indicates that the static power is significantly less than the power consumed by the National Semiconductor ADC08D1500 A/D converter, which is typically 1.8W when running from a 1.5 GHz sample clock. Therefore, for systems processing the captured data in bursts, the ADC can be the main source of heat and power dissipation. Having an ADC with low power figures is a key parameter in the design of products, especially those that are required to be small and portable. The design of this development platform confirms that these qualities are achieved by interfacing the ADC08D1500 to the Virtex-4 device.

Data Transmission

The next consideration for systems using the ADC08D1500 and Virtex-4 FPGA is the signaling between these devices. There are two key issues when handling two channels (each providing data at a rate of 1.5 billion (1.5×10^9) conversions per second):

- Signal integrity between the ADC and FPGA
- The rate of data transfer for each clock cycle

The ADC08D1500 uses low voltage differential signaling (LVDS) for each of its data outputs and clock signal. The main advantage of the LVDS signaling method is that you can achieve high data rates with a very low power budget. Two wires are used for each discrete signal that is to be carried across the circuit board, which should be designed to have a characteristic impedance

of 100 Ohms (defined by the LVDS standard). These traces are differentially terminated at the receiver with a 100 Ohm resistor to match the transmission line (see Figure 2).

A signal voltage is generated across the terminating resistor by a 3.5 mA current source within the driving output buffer, which provides a 350 mV signal swing for the receiving circuit to detect. The ADC08D1500 has a total of four 8-bit

data buses, plus a clock and over-range signal that require an LVDS type connection to the FPGA (Figure 3). This adds up to a total of 34 differential pairs, all of which require 100 Ohm termination.

The Virtex-4 device offers active digitally controlled impedance (DCI) and a simple passive 100 Ohm termination on-chip within the I/O buffers of the device. These on-chip termination methods eliminate the need to place passive resistors on




Figure 1 – Comparing the Virtex-4 static power over device density with the operating power of the ADC08D1500




Figure 2 – A typical LVDS circuit

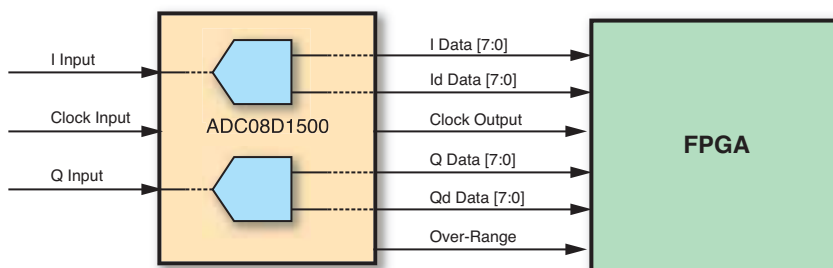


Figure 3 – ADC08D1500 connections to the FPGA

The ADC08D1500 provides a de-multiplexed data output for each of its two channels. Instead of providing a single 8-bit bus running at a data rate equal to the sampling speed, the ADC outputs two consecutive samples simultaneously on two 8-bit data buses (1:2 de-mux).

the circuit board and simplify the routing on the PCB. The DCI option consumes significantly more power than the passive option in this case, simply because of the number of discrete signal lines (68 total) that require termination. Therefore, I would advise turning on the DIFF_TERM feature within each of the IOBs (I/O buffers) to which the ADC signals are connected.

Data Capture

After transmitting data at high speeds using a robust signaling method, it is necessary to store this data into a memory array for post processing. The ADC08D1500 provides a de-multiplexed data output for each of its two channels. Instead of providing a single 8-bit bus running at a data rate equal to the sampling speed, the ADC outputs two consecutive samples simultaneously on two 8-bit data buses (1:2 de-mux).

If the ADC is configured as a single-channel device and put into DES (dual-edge sampling mode), then the sampling speed can be doubled (from 1.5 GSPS to 3.0 GSPS); thus, four consecutive samples are available simultaneously on each of the four buses (1:4 de-mux). This method of de-multiplexing the digital output reduces the data rate to at least half the sampling speed (1:2 de-mux), but increases the number of output data bits from 8 to 16.

For a 1.5 GHz sample rate, the conversion data will be output synchronous to a 750 MHz clock. Even at this reduced speed, FPGA memories and latches would not be able to accept this data directly. It is therefore beneficial to make use of a DDR method, where data is presented to the outputs on the both the rising and falling edges of the clock (Figure 4).

Although the data rate remains the same for DDR signaling, the clock frequency is halved again to a more manage-

able 375 MHz. This frequency is now in the realms of the FPGA IOB data latches. Before this data can be stored away to memory, a small pipeline constructed from a series of data latches is required. Starting with the inputs, for each data line connected to an IOB pair on the FPGA, two latches will be used to capture the incoming data. One latch is clocked on the rising edge of a phase-locked data clock, while the second latch is clocked using a signal that is 180 degrees out of phase.

The relative position of these clocks should be adjusted so that the edges are aligned with the center of the data eye, taking into account the propagation delay of the signal as it enters the FPGA (Figure 5). To simplify this clocking scheme, the Virtex-4 device is equipped with DCMs that allow these clock signals to be generated internally and can be phase-locked to the incoming data clock.

After latching the incoming data using a DCM, the clock domain must be shifted




Figure 4 – Oscilloscope plot of clock (top trace) and data from the ADC in DDR mode

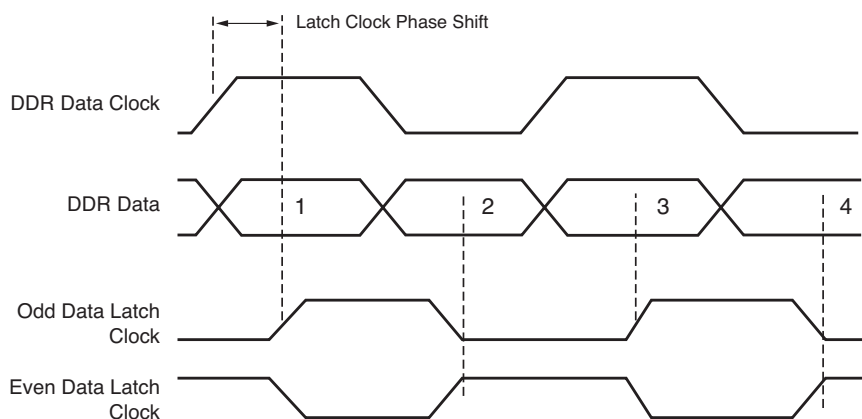


Figure 5 – DDR signaling with DCM-generated data-capture clocks

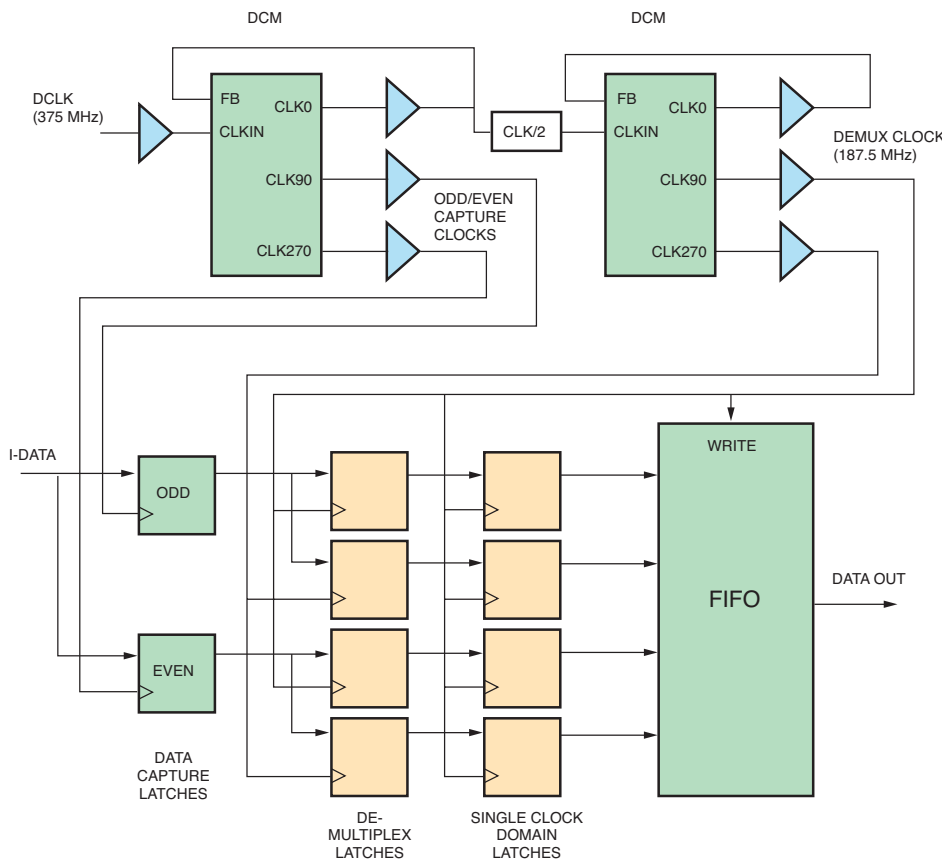


Figure 6 – Data-capture block diagram using two DCMs, latches, and a FIFO memory




Figure 7 – 128 bit input, 16 bit output, 4 KB deep FIFO

using an intermediate set of latches so that all of the data can be clocked into a memory array on the same clock edge. Because of the speed of the clock, there is not sufficient setup and hold time to re-clock the data; therefore the data must be de-multiplexed again to lower the data rate to 187.5 MHz. Once lowered, the data captured on the out-of-phase clock (even) can be re-captured using the in-phase clock (odd) running at the de-multiplexed rate (see Figure 6).

A second DCM is used to produce the de-mux clock. The clock input frequency is internally divided by two, which produces the 187.5 MHz clock signal. This DCM will provide an output that is phase-locked to the synchronous data clock (DCLK).

Data Storage

As shown in Figure 6, a single 8-bit data bus from the FPGA has been de-multiplexed by four. When all four data buses from the ADC are considered, this method produces a data word 128 bits wide running eight times slower than the sample speed for two-channel operation. The data can now be stored into a FIFO memory buffer.

Creating the custom FIFO for this application is made easy using the Xilinx LogiCORE™ FIFO Generator. Using this software wizard, you can create a FIFO with an input bus width as wide as 256 bits, having an aspect ratio (input-to-output bus width ratio) of 8 to 1. As this design has a 128 bit input bus, the minimum output bus width is 16 bits. This works out well, allowing one 8 bit output bus to be used for I Channel data and the other for the Q channel.

Because the aspect ratio is not 1:1, the FIFO generator will create the memory design using block RAM within the FPGA. A single block RAM can be configured as 36 bits wide by 512 locations deep, so to capture the 128-bit conversion word, the design will use four block RAMs. This gives each channel a 4 KB storage depth without having to cascade FIFO blocks (Figure 7). Having 4K bytes of storage is more than sufficient data for

The low power consumption of the two devices enables systems to operate without forced cooling in small enclosures and does not contribute to a large change in ambient temperature.

a Fast Fourier Transform (see Figure 8) to be applied to the digital conversion of the input signal and represents around 2.7 μ S of time-domain information at the 1.5 GHz conversion rate.

Conclusion

When used for the data capture application described, about 85% of the logic fabric inside the Virtex-4 (LX15) device

low switching noise and to be placed in very close proximity to a high-bandwidth, high-speed data converter without significantly downgrading the measured performance solved my FPGA design challenge.

The two-channel ADC development board discussed in this article is available to order from National Semiconductor in three speed grades: 500 MHz, 1 GHz,

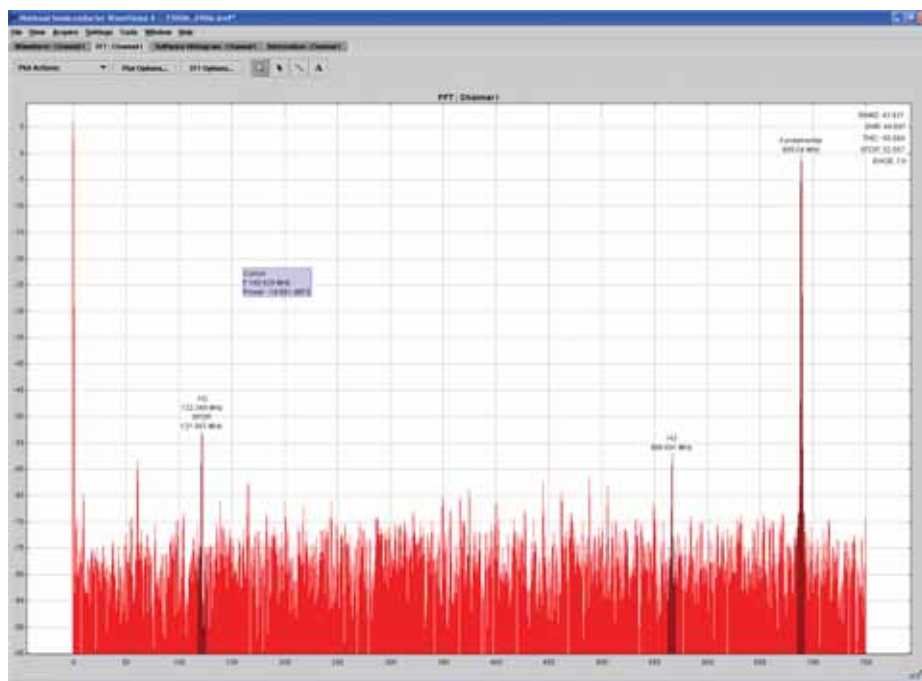


Figure 8 – FFT analysis of 689 MHz input captured by ADC08D1500 and Virtex-4 FPGA

remains available for proprietary firmware development. This leaves space for additional signal processing and data analysis to be performed in hardware, reducing the burden on the software application.

The low power consumption of the two devices enables systems to operate without forced cooling in small enclosures and does not contribute to a large change in ambient temperature. The ability of the Virtex-4 FPGA to operate with

and 1.5 GHz. On-board clocking is provided, so all that is required to get started is to provide an analog signal for sampling, plug in the power supply (included), and connect the USB interface to the host PC.

Single-channel device platforms are also available at 1 GHz and 1.5 GHz sample rates. For more information, visit www.national.com/xilinx and www.national.com/appinfo/adc/ghz_adc.html.

GET ON TARGET



IS YOUR MARKETING MESSAGE REACHING THE RIGHT PEOPLE?

Hit your target audience by advertising your product or service in *I/O Magazine*. You'll reach more than 30,000 engineers, designers, and engineering managers worldwide.

We offer very attractive advertising rates to meet any budget!

Call today:
(800) 493-5551
or e-mail us at
xcelladsales@aol.com



Xilinx/Micron Partner to Provide High-Speed Memory Interfaces

Micron's RLD RAM II and DDR/DDR2 memory combines performance-critical features to provide both flexibility and simplicity for Virtex-4-supported applications.

by Mike Black
Strategic Marketing Manager
Micron Technology, Inc.
mblack@micron.com

With network line rates steadily increasing, memory density and performance are becoming extremely important in enabling network system optimization. Micron Technology's RLD RAM™ and DDR2 memories, combined with Xilinx® Virtex-4™ FPGAs, provide a platform designed for performance.

This combination provides the critical features networking and storage applications need: high density and high bandwidth. The ML461 Advanced Memory Development System (Figure 1) demonstrates high-speed memory interfaces with Virtex-4 devices and helps reduce time to market for your design.

Micron Memory

With a DRAM portfolio that's among the most comprehensive, flexible, and reliable in the industry, Micron has the ideal solution to enable the latest memory platforms. Innovative new RLD RAM and DDR2 architectures are advancing system designs farther than ever, and Micron is at the forefront, enabling customers to take advantage of the new features and functionality of Virtex-4 devices.

RLD RAM II Memory

An advanced DRAM, RLD RAM II memory uses an eight-bank architecture optimized for high-speed operation and a double-data-rate I/O for increased bandwidth. The eight-bank architecture enables

RLD RAM II devices to achieve peak bandwidth by decreasing the probability of random access conflicts.

In addition, incorporating eight banks results in a reduced bank size compared to typical DRAM devices, which use four. The smaller bank size enables shorter address and data lines, effectively reducing the parasitics and access time.

Although bank management remains important with RLD RAM II architecture, even at its worst case (burst of two at 400 MHz operation), one bank is always available for use. Increasing the burst length of the device increases the number of banks available.

I/O Options

RLD RAM II architecture offers separate I/O (SIO) and common I/O (CIO) options. SIO devices have separate read and write ports to eliminate bus turnaround cycles and contention. Optimized for near-term read and write balance, RLD RAM II SIO devices are able to achieve full bus utilization.

In the alternative, CIO devices have a shared read/write port that requires one additional cycle to turn the bus around. RLD RAM II CIO architecture is optimized for data streaming, where the near-term bus operation is either 100 percent read or 100 percent write, independent of the long-term balance. You can choose an I/O version that provides an optimal compromise between performance and utilization.

The RLD RAM II I/O interface provides other features and options, including support for both 1.5V and 1.8V I/O lev-

els, as well as programmable output impedance that enables compatibility with both HSTL and SSTL I/O schemes. Micron's RLD RAM II devices are also equipped with on-die termination (ODT) to enable more stable operation at high speeds in multipoint systems. These features provide simplicity and flexibility for high-speed designs by bringing both end termination and source termination resistors into the memory device. You can take advantage of these features as needed to reach the RLD RAM II operating speed of 400 MHz DDR (800 MHz data transfer).

At high-frequency operation, however, it is important that you analyze the signal driver, receiver, printed circuit board network, and terminations to obtain good signal integrity and the best possible voltage and timing margins. Without proper terminations, the system may suffer from excessive reflections and ringing, leading to reduced voltage and timing margins. This, in turn, can lead to marginal designs and cause random soft errors that are very difficult to debug. Micron's RLD RAM II devices provide simple, effective, and flexible termination options for high-speed memory designs.

On-Die Source Termination Resistor

The RLD RAM II DQ pins also have on-die source termination. The DQ output driver impedance can be set in the range of 25 to 60 ohms. The driver impedance is selected by means of a single external resistor to ground that establishes the driver impedance for all of the device DQ drivers.

As was the case with the on-die end termination resistor, using the RLD RAM II

on-die source termination resistor eliminates the need to place termination resistors on the board – saving design time, board space, material costs, and assembly costs, while increasing product reliability. It also eliminates the cost and complexity of end termination for the controller at that end of the bus. With flexible source termination, you can build a single printed circuit board with various configurations that differ only by load options, and adjust the Micron RDRAM II memory driver impedance with a single resistor change.

DDR/DDR2 SDRAM

DRAM architecture changes enable twice the bandwidth without increasing the demand on the DRAM core, and keep the power low. These evolutionary changes enable DDR2 to operate between 400 MHz and 533 MHz, with the potential of extending to 667 MHz and 800 MHz. A summary of the functionality changes is shown in Table 1.

Modifications to the DRAM architecture include shortened row lengths for reduced activation power, burst lengths of four and eight for improved data bandwidth capability, and the addition of eight banks in 1 Gb densities and above.

New signaling features include on-die termination (ODT) and on-chip driver (OCD). ODT provides improved signal quality, with better system termination on the data signals. OCD calibration provides the option of tightening the variance of the pull-up and pull-down output driver at 18 ohms nominal.

Modifications were also made to the mode register and extended mode register, including column address strobe CAS latency, additive latency, and programmable data strobes.

Conclusion

The built-in silicon features of Virtex-4 devices – including ChipSync™ I/O technology, SmartRAM, and Xesium differential clocking – have helped simplify interfacing FPGAs to very-high-speed memory devices. A 64-tap 80 ps absolute delay element as well as input and output DDR registers are available in each I/O element, providing for the first time a run-time center alignment of data and clock that guarantees reliable data capture at high speeds.



Figure 1 – ML461 Advanced Memory Development System

Xilinx engineered the ML461 Advanced Memory Development System to demonstrate high-speed memory interfaces with Virtex-4 FPGAs. These include interfaces with Micron's PC3200 and PC2-5300 DIMM modules, DDR400 and DDR2533 components, and RDRAM II devices.

In addition to these interfaces, the ML461 also demonstrates high speed QDR-II and FCRAM-II interfaces to

Virtex-4 devices. The ML461 system, which also includes the whole suite of reference designs to the various memory devices and the memory interface generator, will help you implement flexible, high-bandwidth memory solutions with Virtex-4 devices.

Please refer to the RDRAM information pages at www.micron.com/products/dram/rldram/ for more information and technical details. 🌈

FEATURE/OPTION	DDR	DDR2
Data Transfer Rate	266, 333, 400 MHz	400, 533, 667, 800 MHz
Package	TSOP and FBGA	FBGA only
Operating Voltage	2.5V	1.8V
I/O Voltage	2.5V	1.8V
I/O Type	SSTL_2	SSTL_18
Densities	64 Mb-1 Gb	256 Mb-4 Gb
Internal Banks	4	4 and 8
Prefetch (MIN Write Burst)	2	4
CAS Latency (CL)	2, 2.5, 3 Clocks	3, 4, 5 Clocks
Additive Latency (AL)	No	0, 1, 2, 3, 4 Clocks
READ Latency	CL	AL + CL
WRITE Latency	Fixed	READ Latency - 1 Clock
I/O Width	x4/ x8/ x16	x4/ x8/ x16
Output Calibration	None	OCD
Data Strobes	Bidirectional Strobe (Single-Ended)	Bidirectional Strobe (Single-Ended or Differential) with RDQS
On-Die Termination	None	Selectable
Burst Lengths	2, 4, 8	4, 8

Table 1 – DDR/DDR2 feature overview

Implementing High-Performance Memory Interfaces with Virtex-4 FPGAs

You can center-align clock-to-read data at “run time” with ChipSync technology.

by Adrian Cosoroaba
Marketing Manager
Xilinx, Inc.
adrian.cosoroaba@xilinx.com

As designers of high-performance systems labor to achieve higher bandwidth while meeting critical timing margins, one consistently vexing performance bottleneck is the memory interface. Whether you are designing for an ASIC, ASSP, or FPGA, capturing source-synchronous read data at transfer rates exceeding 500 Mbps may well be the toughest challenge.

Source-Synchronous Memory Interfaces

Double-data rate (DDR) SDRAM and quad-data-rate (QDR) SRAM memories utilize source-synchronous interfaces through which the data and clock (or strobe) are sent from the transmitter to the receiver. The clock is used within the receiver interface to latch the data. This eliminates interface control issues such as the time of signal flight between the memory and the FPGA, but raises new challenges that you must address.

One of these issues is how to meet the various read data capture requirements to implement a high-speed source-synchronous interface. For instance, the receiver must ensure that the clock or strobe is routed to all data loads while meeting the required input setup and hold timing. But source-synchronous devices often limit the loading of the forwarded clock. Also, as the data-valid window becomes smaller at higher frequencies, it becomes more important (and simultaneously more challenging) to align the received clock with the center of the data.

Traditional Read Data Capture Method

Source-synchronous clocking requirements are typically more difficult to meet when reading from memory compared with writing to memory. This is because the DDR and DDR2 SDRAM devices send the data edge aligned with a non-continuous strobe signal instead of a continuous clock. For low-frequency interfaces up to 100 MHz, DCM phase-shifted outputs can be used to capture read data.

Capturing read data becomes more challenging at higher frequencies. Read data can be captured into configurable logic blocks (CLBs) using the memory read strobe, but the strobe must first be delayed so that its edge coincides with the center of the data valid window. Finding the correct phase-shift value is further complicated by process, voltage, and temperature (PVT) variations. The delayed strobe must also be routed onto low-skew FPGA clock resources to maintain the accuracy of the delay.

The traditional method used by FPGA, ASIC, and ASSP controller-based designs employs a phase-locked loop (PLL) or delay-locked loop (DLL) circuit that guarantees a fixed phase shift or delay between the source clock and the clock used for capturing data (Figure 1). You can insert this phase shift to accommodate estimated process, voltage, and temperature variations. The obvious drawback with this method is that it fixes the delay to a single value predetermined during the design phase. Thus, hard-to-predict variations within the system itself – caused by different routing to different memory devices, variations between FPGA or ASIC devices, and ambient system condi-

tions (voltage, temperature) – can easily create skew whereby the predetermined phase shift is ineffectual.

These techniques have allowed FPGA designers to implement DDR SDRAM memory interfaces. But very high-speed 267

also cause data and address timing problems at the input to the RAM and the FPGA's I/O blocks (IOB) flip-flop. Furthermore, as a bidirectional and non-free-running signal, the data strobe has an increased jitter component, unlike the clock signal.




Figure 1 – Traditional fixed-delay read data capture method

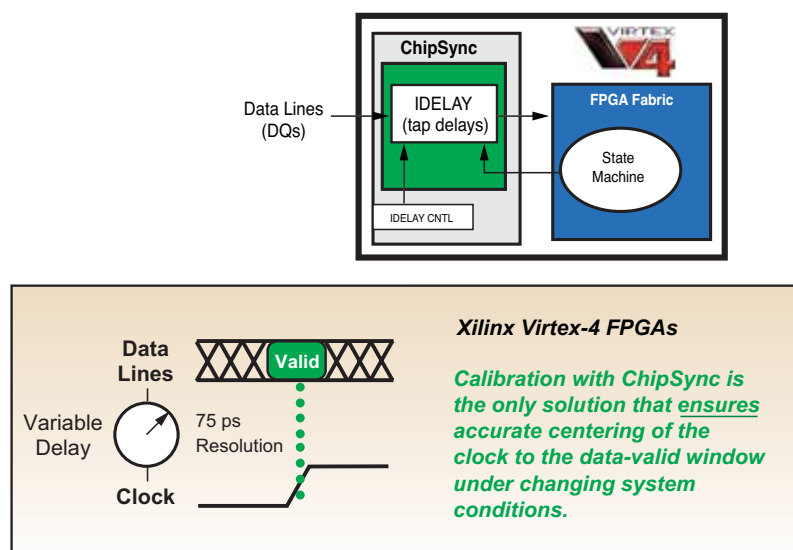


Figure 2 – Clock-to-data centering using ChipSync tap delays

MHz DDR2 SDRAM and 300 MHz QDR II SRAM interfaces demand much tighter control over the clock or strobe delay.

System timing issues associated with setup (leading edge) and hold (trailing edge) uncertainties further minimize the valid window available for reliable read data capture. For example, 267 MHz (533 Mbps) DDR2 read interface timings require FPGA clock alignment within a .33 ns window.

Other issues also demand your attention, including chip-to-chip signal integrity, simultaneous switching constraints, and board layout constraints. Pulse-width distortion and jitter on clock or data strobe signals

Clock-to-Data Centering Built into Every I/O

Xilinx® Virtex™-4 FPGAs with dedicated delay and clocking resources in the I/O blocks – called ChipSync™ technology – answer these challenges. These devices make memory interface design significantly easier and free up the FPGA fabric for other purposes. Moreover, Xilinx offers a reference design for memory interface solutions that center-aligns the clock to the read data at “run time” upon system initialization. This proven methodology ensures optimum performance, reduces engineering costs, and increases design reliability.

ChipSync features are built into every I/O. This capability provides additional flexibility if you are looking to alleviate board layout constraints and improve signal integrity.

ChipSync technology enables clock-to-data centering without consuming CLB resources. Designers can use the memory read strobe purely to determine the phase relationship between the FPGA's own DCM clock output and the read data. The read data is then delayed to center-align the

determine the phase relationship between the FPGA clock and the read data received at the FPGA. This is done using the memory read strobe. Based on this phase relationship, the next step is to delay read data to center it with respect to the FPGA clock. The delayed read data is then captured

transition (second-edge taps) in the FPGA clock domain.

Having determined the values for first-edge taps and second-edge taps, the state machine logic can compute the required data delay. The pulse center is computed with these recorded values as $(\text{second-edge taps} - \text{first-edge taps})/2$. The required data delay is the sum of the first-edge taps and the pulse center. Using this delay value, the data-valid window is centered with respect to the FPGA clock.

ChipSync features are built into every I/O. This capability provides additional flexibility if you are looking to alleviate board layout constraints and improve signal integrity.

Each I/O also has input DDR flip-flops required for read data capture either in the delayed memory read strobe domain or in the system (FPGA) clock domain. With these modes you can achieve higher design performance by avoiding half-clock-cycle data paths in the FPGA fabric.

Instead of capturing the data into a CLB-configured FIFO, the architecture provides dedicated 500 MHz block RAM with built-in FIFO functionality. These enable a reduction in design size, while leaving the CLB resources free for other functions.

Clock-to-Data Phase Alignment for Writes

Although the read operations are the most challenging part of memory interface design, the same level of precision is required in write interface implementation. During a write to the external memory device, the clock/strobe must be transmitted center-aligned with respect to data. In the Virtex-4 FPGA I/O, the clock/strobe is generated using the output DDR registers clocked by a DCM clock output (CLK0) on the global clock network. The write data is transmitted using the output DDR registers clocked by a




Figure 3 – Clock-to-data centering at “run time”

FPGA clock in the read data window for data capture. In the Virtex-4 FPGA architecture, the ChipSync I/O block includes a precision delay block known as IDELAY that can be used to generate the tap delays necessary to align the FPGA clock to the center of the read data (Figure 2).

Memory read strobe edge-detection logic uses this precision delay to detect the edges of the memory read strobe from which the pulse center can be calculated in terms of the number of delay taps counted between the first and second edges. Delaying the data by this number of taps aligns the center of the data window with the edge of the FPGA DCM output. The tap delays generated by this precision delay block allow alignment of the data and clock to within 75 ps resolution.

The first step in this technique is to

directly in input DDR flip-flops in the FPGA clock domain.

The phase detection is performed at run time by issuing dummy read commands after memory initialization. This is done to receive an uninterrupted strobe from the memory (Figure 3).

The goal is to detect two edges or transitions of the memory read strobe in the FPGA clock domain. To do this, you must input the strobe to the 64-tap IDELAY block that has a resolution of 75 ps. Then, starting at the 0-tap setting, IDELAY is incremented one tap at a time until it detects the first transition in the FPGA clock domain. After recording the number of taps it took to detect the first edge (first-edge taps), the state machine logic continues incrementing the taps one tap at a time until it detects the second

DCM clock output that is phase-offset 90 degrees (CLK270) with respect to the clock used to generate clock/strobe. This phase shift meets the memory vendor specification of centering the clock/strobe in the data window.

Another innovative feature of the output DDR registers is the SAME_EDGE mode of operation. In this mode, a third register clocked by a rising edge is placed on the input of the falling-edge register. Using this mode, both rising-edge and falling-edge data can be presented to the output DDR registers on the same clock edge (CLK270), thereby allowing higher DDR performance with minimal register-to-register delay.

Signal Integrity Challenge

One challenge that all chip-to-chip, high-speed interfaces need to overcome is signal integrity. Having control of cross-talk, ground bounce, ringing, noise margins, impedance matching, and decoupling is now critical to any successful design.

The Xilinx column-based ASMBL architecture enables I/O, clock, and power and ground pins to be located anywhere on the silicon chip, not just along the periphery. This architecture alleviates the problems associated with I/O and array dependency, power and ground distribution, and hard-IP scaling. Special FPGA packaging technology known as SparseChevron enables distribution of power and ground pins evenly across the package. The benefit to board designers is improved signal integrity.

The pin-out diagram in Figure 4 shows how Virtex-4 FPGAs compare with a competing Altera Stratix-II device that has many regions devoid of returns.

The SparseChevron layout is a major reason why Virtex-4 FPGAs exhibit unmatched simultaneous switching output (SSO) performance. As demonstrated by signal integrity expert Howard Johnson, Ph.D., these domain-optimized FPGA devices have seven times less SSO noise and crosstalk when compared to alternative FPGA devices (Figure 5).

Meeting I/O placement requirements and enabling better routing on a board requires unrestricted I/O placements for

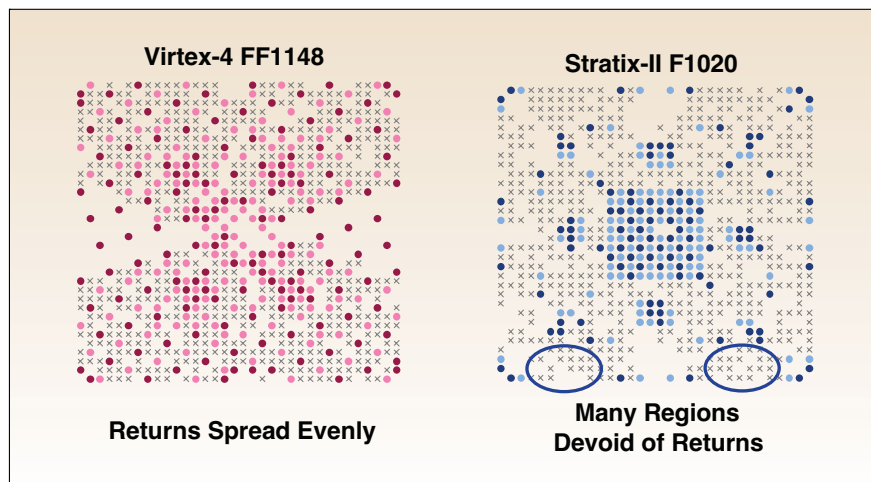


Figure 4 – Pin-out comparison between Virtex-4 and Stratix-II FPGAs




Figure 5 – Signal integrity comparison using the accumulated test pattern

an FPGA design. Unlike competing solutions that restrict I/O placements to the top and bottom banks of the FPGA and functionally designate I/Os with respect to address, data and clock, Virtex-4 FPGAs provide unrestricted I/O bank placements.

Finally, Virtex-4 devices offer a differential DCM clock output that delivers the extremely low jitter performance necessary for very small data-valid windows and diminishing timing margins, ensuring a robust memory interface design.

These built-in silicon features enable high-performance synchronous interfaces for both memory and data communications in single or differential mode. The ChipSync technology enables data rates

greater than 1 Gbps for differential I/O and more than 600 Mbps for single-ended I/O.

Conclusion

As with most FPGA designs, having the right silicon features solves only part of the challenge. Xilinx also provides complete memory interface reference designs that are hardware-verified and highly customizable. The Memory Interface Generator, a free tool offered by Xilinx, can generate all of the FPGA design files (.rtl, .ucf) required for a memory interface through an interactive GUI and a library of hardware-verified designs.

For more information, visit www.xilinx.com/memory.

Debugging and Validating PCI Express I/O

With these tips and tricks for using a logic analyzer, you can speed time to market and increase confidence in your design.

by Richard Markley
Logic Analysis Product Planning Manager
Agilent Technologies
richard_markley@agilent.com

Marco Davila
R&D Hardware Designer
Agilent Technologies
marco_davila@agilent.com

As PCI Express continues to replace PCI in many designs, engineers are finding themselves in uncharted territory. High-speed serial links running at 2.5 Gbps introduce new challenges that were not seen with traditional wider and slower parallel buses like PCI. Vias look like stubs. Data is 8b/10b encoded such that clocks are embedded. Signal swings are minimal. The list goes on and on. With these new challenges, you will need to rely more on test equipment than you have in the past.

One of these key pieces of test equipment is the logic analyzer. Although at first glance a logic analyzer may not appear to be suited for debugging a serial bus, recent advances have made the logic analyzer a powerful tool for system bring up and validation of serial buses like PCI Express (PCIe).

New technologies allow the logic analyzer interface (also known as an analysis probe) to use its hardware resources (instead of the logic analyzer's triggering resources) to look for packets.

Probing Advancements

Successfully probing a PCIe link is not a trivial task. Because of the gigabit speeds, test and measurement vendors need probing that is non-intrusive and easy to use.

The simplest method to probe a PCIe link is to use a slot interposer. Slot interposers require no forethought when it comes to probing – you simply plug the interposer into an available PCIe slot and plug your add-in card on top. Although they are simple to use, some interposers

specified a common footprint for all test vendors. This footprint is a “connector-less” design that uses landing pads for probing. Although very different from a slot interposer, the same potential concerns exist – electrical and mechanical non-intrusiveness.

In addition to these potential concerns, many designers should also consider how easy the probes are to use. Do they require special cleaning to get a reliable connection? Are they compatible with multiple board finishes such as hot air solder leveling

probe at the full link speed (2.5 Gbps) while keeping probe head volume to a minimum. An example of a flying lead set is shown in Figure 3.

Triggering Advancements

Because of the parallel nature of the logic analyzer, triggering on a packetized bus requires you to use many of the logic analyzer's triggering resources to define just the start of a packet. This is especially true in PCI Express, which has the option of multi-



Figure 1 – PCI Express slot interposer



Figure 2 – PCI Express midbus probe



Figure 3 – PCI Express flying lead set

are less intrusive than others. Obviously, an interposer cannot be so electrically intrusive that it breaks the link (that is, it doesn't allow the device under test to work). However, it is also important to pay attention to the mechanical intrusiveness of a slot interposer. Interposers that are shorter, with vertical egress (see Figure 1), provide more testing options to system designers.

Although interposers are simple to use, they are not helpful for chip-to-chip designs. Probing these designs (often called “midbus probing”) typically requires a designed in footprint. The PCI-SIG has

process (HASL) or gold plating? Do they require external cooling fans? An example of a midbus probe is shown in Figure 2.

Although a midbus probe is typically the preferred method for probing chip-to-chip designs, it does require a footprint to be designed in. Sometimes engineers do not have the room for a design in footprint, or they may have not considered debugging and validation early enough to design in the footprint. In these cases, a flying lead set can be very beneficial. As with all probing systems, the flying lead set must be electrically and mechanically non-intrusive. It should allow designers to

ple lane widths. The serial nature of the bus makes triggering significantly different from triggering on a parallel bus, where you would normally specify a value for a specific label.

New technologies allow the logic analyzer interface (also known as an analysis probe) to use its hardware resources (instead of the logic analyzer's triggering resources) to look for packets. These packet analysis probes contain “packet recognizers” specifically designed to help trigger on serial links. These allow you to define as many as four packets in each direction for the logic analyzer to trigger on. In addition, each packet recognizer allows you to define the entire

packet header, and as many as 8 bytes of the data payload (for a 3 double word [3DW]). These packet recognizers also provide the means for specifying “don’t cares” within the header/data fields. This stands in stark contrast to traditional logic analyzer resources that only allow you to define the packet type (transaction layer packet [TLP] or data link layer packet [DLLP]).

At first, the packet recognizer must determine the start of the packet. The packet may start in one of four lanes for a x16 link (lane 0, 4, 8, or 12), so the packet recognizer must look in each of these lanes. It does this automatically – you do not have to worry about defining the trigger steps to recognize this. Traditional logic analyzer triggering ends up using a large portion of its resources to determine only this event.

After resolving the start of packet and deskewing the lanes (just as the actual receiver does), the packet recognizers then look for matches to fields within the packet header and the data payload. The packet analysis probe will then send a signal back to the logic analyzer, which it can use in a trigger. These signals can be used with the full triggering resources of the analyzer (including counters, timers, sequencers, storing, and multi-way branching) to provide very robust, powerful triggering.

Common Debug Triggers

Using packet recognizers allows you to define an almost limitless amount of triggers. They are often used in debug techniques such as:

- Prestore and qualified capturing of packets
- Cross-bus triggering
- Triggering using an exerciser

During initial bring up of a PCIe device, you may want to capture a specific event and a large period of time before that event. Because you need to capture a long period in time, it is often beneficial to only store events that are of interest in the logic analyzer’s memory. However, this requires additional triggering and storage resources. If these resources are completely used in defining the type of packet, this may not be possible.

...test equipment like logic analyzers can help you as you move from the parallel world to the serial world.

A packet recognizer helps alleviate this problem. For example, you can define a specific packet header along with several bytes of data. We will call this “3DW with Data.” You can then define another packet that includes all of the types of events you want to store. In this case we only want to store other TLPs – all other fields in the recognizer are left as “don’t cares.” We call this “TLP only.” The logic analyzer will then use a simple pattern trigger to find the “3DW with Data” event, and you now have all of the analyzer’s resources left to qualify what is stored.

Often you will only want to see information before the trigger. In this case, you can set the logic analyzer to do what is called “prestore.” A 100% prestore will only store information before the trigger, so you can capture a larger period of time before your trigger event. When used in conjunction with the default storing, this allows you to capture the maximum amount of time before the trigger. In most logic analyzers, you can easily define the percent of “pre” or “post” store.

In a serial architecture like PCI Express, a disagreement between the perceived traffic viewed by the transmitter and receiver doesn’t always point to the root cause of a problem. Using a cross-bus triggering technique allows you to not only trigger on this disagreement, but also locate the source of problem. This problem might be caused by another bus in the system such as the processor system bus, DDR memory bus, SATA/SAS bus, or another I/O bus.

This is a very easy trigger to setup, but very powerful in the information that it provides. You can trigger from any one bus and capture time-correlated events on the other buses in their system. For example, a common trigger involves looking for a bus hang on the processor system bus.

This will then trigger and capture data on all of the additional buses you are looking at. Should the processor bus hang be caused by an event on the PCIe link, this is a quick way to see the events time-correlated together for maximum debug.

Another common cross-bus triggering technique involves looking at the PCIe link from the south bridge to a switch with multiple PCI slots. For example, it is often beneficial to trace a specific event as it occurs on the PCI bus and travels through the bridge to the PCIe link. Once again, packet recognizers can be very beneficial in this case, because they allow you to look for a very specific packet header with data. Traditional triggering using the logic analyzer’s resources would have a difficult time defining the packet with enough detail to capture this event easily.

Another common debug technique involves using an exerciser to generate traffic on the PCIe link while using the logic analyzer to capture the response to this stimulus. This is often known as “stimulus and response capture” and is a very powerful technique that is normally employed later in a designer’s program to test the compliance of their devices.


Conclusion

PCI Express is taking off as a common I/O interconnect for many designers. Although it has many benefits (scalable, backwards compatibility to PCI, fewer signals), it does present some significant design challenges. Because of this, test equipment like logic analyzers can help you as you move from the parallel world to the serial world.

To learn more about the equipment discussed in this article, please visit www.agilent.com/find/pciexpress or contact your local Agilent field engineer. 🌈



READY! SET! GO WITH PLANAHEAD PERFORMANCE!



Based on benchmark data from a suite of 15 real-world customer designs targeting Xilinx and competing FPGA Solutions.



Two speed grades faster with PlanAhead software and Virtex-4

With our unique PlanAhead software tool, and our industry-leading Virtex-4 FPGAs, designers can now achieve a new level of performance. For complex, high-utilization, multi-clock designs, no other competing FPGA comes close to the Virtex-4 PlanAhead advantage:

- 30% better logic performance on average = 2 speed grade advantage
- Over 50% better logic performance for complex multi-clock designs

MEET YOUR TIMING BUDGETS ... BEAT YOUR COMPETITION TO MARKET

Meeting timing budgets is the most critical issue facing FPGA designers*. Inferior tools can hit a performance barrier, impacting your timing goals, while costing you project delays and expensive higher speed grades. To maximize the Virtex-4 performance advantage, the new PlanAhead software tool allows you to quickly analyze, floorplan, and improve placement and timing of even the most complex designs. Now, with ISE and PlanAhead you can meet your timing budgets and reduce design iterations, all within an easy-to-use design environment.

Download a free eval today at www.xilinx.com/planahead, view the TechOnline web seminar, and prevent your next FPGA design from stalling.

* CMP: June 2005 FPGA EDA Survey



Highest customer satisfaction:

- ✓ Virtex & Spartan FPGAs
- ✓ ISE design tools
- ✓ Support



The Programmable Logic Company™

www.xilinx.com/planahead



View The
TechOnline
Seminar Today

BREAKTHROUGH PERFORMANCE AT THE LOWEST COST

©2006 Xilinx, Inc. All rights reserved. XILINX, the Xilinx logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

Using Complex Triggers in the Identify Debugger

You can obtain huge productivity gains with Synplicity's powerful and comprehensive FPGA debug tool.

by Dennis McCarty
Technical Marketing Manager
Synplicity, Inc.
dmccarty@synplicity.com

Hardware debuggers represent the ultimate system verification tool. Unlike simulators, debuggers show what the logic is actually doing inside the device while running in the system at full speed. When using a hardware debugger, it is crucial that you capture the precise data you need to discover bugs and verify system behavior. Not only must you locate the logic transitions around a certain event, you must also track bugs that may be rare events and trap them for closer examination.

The Identify RTL debugger from Synplicity offers you a view of logic behavior inside an FPGA operating within the system. It also offers a highly sophisticated set of trigger mechanisms and other features that you can use to isolate events germane to a particular problem.

In this article, I'll describe some of the features of Identify.




Figure 1 – Cross-trigger example

Triggering Across Clock Domains

Today's FPGA designers frequently use multiple clocks, as these devices come with numerous dedicated clock buffers. In multi-clock systems it is common to encounter timing problems related to clocking data between domains. Such problems include metastability, failure to meet setup or hold times, and dropped data. Detecting these often subtle problems is usually difficult. The problem may not appear in logic simulation at all, and may only be detected while debugging by over-sampling within a domain or by triggering from one domain and sampling in another.

Cross-triggering is a technique that enables you to trigger on an event in one domain and sample an event in another. As shown in Figure 1, the Identify product allows the trigger logic of one domain to drive and enable the trigger in another. You can use cross-triggering to view the timing of events that cross domains. You can also use it to see events occurring within a clock period by over-sampling the period with a faster clock.

Sampling Modes

Sampling modes control the way data is added to the buffer when a trigger condition is reached. These modes allow you to sort data inflows by mode and increase buffer efficiency by storing only relevant data.

Identify software offers four sampling modes:

- The normal mode fills the buffer completely in a single trigger event. Subsequent triggers are ignored unless you run the debugger again.
- In the always armed sampling mode, the buffer fills on every trigger until the debug is stopped using the stop icon.
- The qualified fill mode stores a single sample on each trigger. The buffer will contain only events that caused a trigger and will continue until the buffer is full or when sampling stops.
- The qualified interrupt sampling is like qualified fill, except that sampling will continue until it is interrupted. If sampling continues after the buffer is full, old data will be overwritten.

The qualified and always armed sampling modes must be enabled separately for each intelligent in-circuit emulator (IICE) module during instrumentation. You can enable these modes by clicking on the IICE configuration button in the Instrumentor and checking the boxes in the IICE sampler menu, as shown in Figure 2.



Figure 2 – Sampling modes

The sample mode is set during debugging using the pull-down sample mode icon menu, as shown in Figure 3.



Figure 3 – Sample mode pull-down menu

Trigger Modes

Trigger modes control the way data is added to the buffer upon reaching a trigger condition. There are four operating modes:

- The cycles mode triggers on the number entered in the value field representing the number of clock cycles after the condition.
- The events mode triggers on the n th instance of a trigger condition. In this mode the value field specifies the instance.
- The pulsewidth mode triggers after the trigger condition has remained active for n clock cycles.
- The watchdog mode triggers when the condition has not been active for n clock cycles since the last trigger event.

The default mode is cycles. To use the other modes, you must enable them by selecting the IICE configure button and clicking on the “complex counter triggering” box under the IICE controller menu. Use the arrow selectors to set the counter width to the maximum binary value you might need (Figure 4).



Figure 4 – Enabling trigger mode

To select trigger modes, use the down arrow, as shown in Figure 5.



Figure 5 – Specifying trigger mode (pulsewidth mode selected)

Bus Trigger Expressions

The Watchpoint setup display is used for single-bit data (see Figure 6).




Figure 6 – Watchpoint setup

Setting the trigger for a bus or a portion of a bus is more complicated, but offers a more powerful form of triggering. A right-click on a bus brings forth the menu shown in Figure 7. Several values or ranges of values are available. Entering a value in the left column but not the right causes a trigger on the exact value. Entering data in both columns will cause a trigger on the transition from the left value to the right value. To enable the trigger, check the box(es) next to each one.




Figure 7 – The four values 0-3 indicate that the currently selected IICE was configured for state machine triggering and that the four values correspond to C0-C3 in the state editor.

Partial Bus Trigger Values

Partial bus instrumentation is the definition of one or more bits of a bus such that it can be instrumented separately. Partial bus segments are defined using the menu, which you can invoke by right-clicking on the bus and selecting “add partial instrumentation.”

Each partial bus segment can be instrumented using the bus trigger menu displayed in Figure 8.



Figure 8 – Instrumenting partial bus segments

Trigger State Machine Editor

The most precise and powerful way to detect a unique condition is to use a state machine as a trigger. A state machine can traverse between states on any condition and trigger, or not, in any state. By using a state machine, you can create a sequence of steps and conditions that must be completed to arrive at a trigger condition. The Identify tool includes a state machine editor that allows you to graphically tailor the steps necessary to create the exact trigger condition you desire.

Although it is certainly possible to create a state machine directly in the source code for the purpose of triggering on an event, the Identify editor automates this process by providing a menu-based method. Moreover, a manual solution would require that you manually adjust the logic and specify new trigger nodes during instrumentation for each trigger adjustment and re-synthesis.

Adjustments such as whether to trigger on a state, under what conditions, and how the counter will be used to trigger are made in the debugger. You can dynamically make these adjustments during debugging without tampering directly with the design, making it easier and more efficient to use the Identify product’s integrated graphical state machine solution.

Configuring the IICE for State Machine Triggering

Configuring the IICE in advance is required for state machine debugging. The state machine trigger submenu is located in the IICE configuration menu, as shown in Figure 9. After specifying state machine triggering, you use the

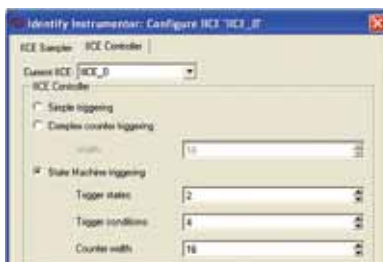


Figure 9 – State machine triggering through IICE menus

wheel switches to dial the number of states, number of trigger conditions, and the width of the counter. You do not have to use all of the resources specified at this stage during debugging.

Saving the IICE selection allows you to specify the behavior and triggering conditions when you are ready to debug. It is in the debugger where you define the state machine states and conditions. For any IICE that has been set to allow state machine triggering, an icon appears, as shown in Figure 10.



Figure 10 – Example of IICE module not enabled for state machine triggering

Those IICE modules not enabled for state machine triggering are shown with a gray box icon.

Defining the State Machine

Selecting the state machine icon invokes the state editor, as shown in Figure 11. The editor initializes to display a space for each of the states specified in the IICE configuration.



Figure 11 – Invoking the state machine editor

The editor has a pull-down insert macro selector from which you can select one of eight macros. The macros apply either one of the four trigger modes described above, one of two conditional modes, or one of two sample modes similar to those in the state machine.

Selecting a macro from the menu invokes the macro editor, which is used to define the macro function. The macro editor contains fields that determine which condition will be used for the state and the number of events or samples that will be counted. Select the condition(s) from among the numbered C values.

The Identify product brings uniquely powerful and comprehensive capabilities to FPGA debugging. The multiple clock triggering feature allows you to see events that are likely to remain undetected in a simulation environment.

Watchdog Timer Mode

The `st_watchdog` editor is shown in Figure 12 as an example. The editor defines the macro function and definition fields. Enter the transition condition in the A field. The transition is one of the state names among the number of states defined during instrumentation. The value for N is the number of clocks the timer counts before the trigger.

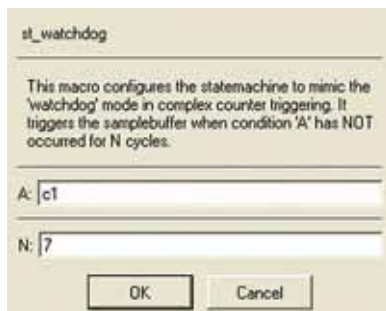


Figure 12 – `st_watchdog` editor

Conditional Modes

Two other macro examples are shown in Figure 13. On the left is the `st_B_after_A` macro. Here you enter two conditions (A and B) with the trigger based on the *n* number of times that B occurs after A has occurred. Condition A is then the qualifier to check for B one or more times for the trigger.




Figure 13 – Using conditional modes

State Editor

Each state has conditions under which it will transition to another state. The transition editor is used to describe the conditions of

one or more transitions from a state. You can invoke the editor by clicking on the pencil-and-paper icon. The editor includes fields and options for each state (Figure 14).



Figure 14 – The transition editor describes conditions of transitions from a state.

State Transitions

The first selection is the state number, from which the current state will transition. Use the thumbwheel to select the state. When you click OK to leave the editor, leave the “from” set to this state. If you select a “from” state other than the state where the editor was invoked, it will apply your changes to the other state and eliminate the transition altogether from the state you are editing. Remember, you can have any number of transitions to other states or remain in the current state.

Describing Conditions

In “on condition,” you specify the state condition under which the trigger will fire. The choices include any of the conditions (notated by a C) defined during the IICE configuration. These conditions are defined during instrumentation. Editing the value for any Watchpoint will display a value for each condition. Defining multiple Watchpoints as conditions will logically AND the conditions.

The default condition is “true,” meaning that the trigger will fire simply by entering the state. You can enter any of the C numbered values or “cntnull” by

typing in the value and negate the preceding value with an exclamation point.

State Machine Actions

The “actions” section works with the previous selections to allow another level of trigger control. The red T trigger box enables the trigger to fire when checked and when the previously described conditions exist. The remaining boxes control the counter and only affect triggering when the condition is selected as “cntnull.” That is when the counter reaches a value of zero.

The counter always decrements as represented by a counterclockwise arrow. The counter can be loaded to any value, as indicated by the down arrow. In any state the counter may be loaded, or enabled, to count down. If the counter reaches zero, it must be reloaded before its next use.

Checking the initialize counter box and entering a value starts the counter from that initial value. The trigger will, if enabled, fire when the counter rolls over.

You can add any number of additional state transition conditions to each state.

Transition values are cleared using the blank sheet icon. Transitions themselves are deleted using the X icon.

Conclusion

The Identify product brings uniquely powerful and comprehensive capabilities to FPGA debugging. The multiple clock triggering feature allows you to see events that are likely to remain undetected in a simulation environment. The sampling modes maximize buffer efficiency. The advanced triggering capabilities are a means for highly sophisticated refinement of data search methods.

The Identify product is a dynamic, in-system debugging environment that offers huge productivity gains, allowing you to debug in RTL code.

For more information, visit www.synplicity.com/products/identify/index.html.

Understanding the PCI-SIG Compliance Program

This program is the key to the successful launch of any product that incorporates PCI-SIG technologies such as PCI, PCI-X, or PCI Express.

by Eric Crabill
Staff Design Engineer
Xilinx, Inc.
eric.crabill@xilinx.com

The PCI-SIG Compliance Program, which is open to all members of the PCI-SIG, seeks to encourage and achieve the highest degree of voluntary compliance with PCI-SIG specifications where PCI-SIG technologies are used. The ultimate goal is to foster the development of high-quality products that offer reliable and hassle-free operation.

For most, the ultimate goal of participation is inclusion on the PCI-SIG Integrators List, which is a "quality pedigree" for a product. As a participant, you may elect to follow through to completion, or stop at any point along the way. The three parts of the program are:

- The Compliance Checklist
- The Compliance Workshop
- The Integrators List

In this article, I will present the utility of each of these steps to help you understand why the PCI-SIG Compliance Program should be an integral part of your product development.

Compliance Workshop Test Results Report				
(Press hard, you are making 2 copies. Turn white copy in to SIG)				
Add-in Card PCI Express		System PCI Express		
Vendor		Vendor		
Product Name		Product Name		
Product Rev.		Product Rev.		
Driver Rev.		BIOS Rev.		
Product Type		Motherboard		
<input type="checkbox"/> Graphics <input type="checkbox"/> SCSI <input type="checkbox"/> IDE <input type="checkbox"/> LAN <input type="checkbox"/> Other _____		CPU _____ Speed _____ # of Slots _____ MB Chipset _____ (Check all of the below that apply) <input type="checkbox"/> Planar <input type="checkbox"/> Graphics <input type="checkbox"/> SCSI <input type="checkbox"/> IDE <input type="checkbox"/> LAN <input type="checkbox"/> Other _____		
<input type="checkbox"/> Uses PCI to PCI Bridge				
	Test	Pass	Fail	Comments
System	Electrical Tests			
	BIOS Tests			
Add-in Card	Electrical Tests			
	Configuration Tests			
	Link Tests			
	Transaction Tests			
Overall Evaluation		Pass	Fail	
Notes				
Add-in Card Vendor		System Vendor		
Signature		Signature		
Print Name		Print Name		
Date		Date		

Figure 1 – PCI-SIG-focused testing results report (published with permission from PCI-SIG)

The Compliance Checklist

In addition to providing detailed and complete specifications, the PCI-SIG publishes a Compliance Checklist for each of its technologies. Although not a substitute for the original specification, Compliance Checklists provide an excellent design-time reference for product design and verification teams. Compliance Checklists are freely available on the PCI-SIG website.

Typically, a Compliance Checklist includes system, functional, electrical, timing, and mechanical assertions covering specification requirements that are deemed of paramount importance. If you are designing your product from scratch, the Compliance Checklist serves as a valuable guide for performing a critical review of

your product during the design phase.

Keep in mind that an interface IP core is not a complete application; some portions of the Compliance Checklist cover requirements that are beyond the scope of an IP core. Obvious examples of this are mechanical requirements; less obvious ones might be electrical and timing characteristics of an IP core delivered as source code.

If you are using a PCI, PCI-X, or PCI Express interface from an IP core provider, you should request Compliance Checklist information from the vendor. You will need this information to submit your own Compliance Checklist to the PCI-SIG for your finished product to be included on the Integrators List. The PCI-SIG suggests completing it after passing the Compliance

Workshop, but if you start reviewing the Compliance Checklist much earlier in the design cycle, you will have done yourself a great favor.

The Compliance Workshop

Several times a year, the PCI-SIG organizes free Compliance Workshops for members of the PCI-SIG. The Compliance Workshops provide three distinct opportunities:

- Focused compliance testing done directly by the PCI-SIG
- Interoperability testing done with other attendees
- A free lunch

As a participant, you fall in one of four categories: stationary PCI-SIG tester, traveling PCI-SIG tester, motherboard/system vendor, or add-in card vendor. Typically, the event is held in a hotel, with stationary PCI-SIG testers and motherboard/system vendors located in individual hotel suites. During the event check-in, participants are given a test schedule, where traveling PCI-SIG testers and add-in card vendors are given scheduled time slots in appropriate test suites. Participants have the option to decline testing with each other for any reason, and test results are confidential.

The details of the focused compliance testing done directly by the PCI-SIG depend on the type of interface involved. For example, PCI Express add-in cards are tested for electrical compliance, subjected to link and transaction protocol tests, and checked for a proper configuration space implementation. (Figure 1 shows the “report card” on which results are recorded.) To help participants pass the tests on their first visit to the Compliance Workshop, the PCI-SIG provides complete information about the tests on their website.

It is possible to run all of the tests in your own lab before attending the Compliance Workshop; this is a great strategy if you want to pass with flying colors on your first attempt. For PCI Express, the configuration tests do not require specialized test equipment. The electrical tests require a high-speed oscilloscope and a compliance base board, which is a hard-

ware test platform available from PCI-SIG. The link and transaction protocol tests require a specific Agilent protocol test card.

A complete lab setup might run close to \$150,000. Some of us are fortunate to have employers with this kind of capital equipment. If you do not have access to suitable test equipment, consider designing with transceivers and IP cores that have already passed these tests; you can participate in the focused compliance testing with confidence, even if you do not have the ability to perform it in advance.

The interoperability test sessions are less exacting than the focused tests. However, they are no less important, as they provide advance warning of problems

ity “report card” that is used for reporting results.) In the event that problems arise, I have observed that participants are highly motivated to resolve interoperability issues – often, someone with test or analysis equipment at the event is willing to help debug the issue and isolate the root cause.

The PCI-SIG recognizes that participants may bring designs that are not fully compliant, or have unknown or undisclosed bugs. For this reason, to pass the interoperability tests, you must only demonstrate a success rate of 80%. If you have also passed the PCI-SIG focused tests, you have met the additional requirements to have your device included on the

cations, the annual PCI-SIG Developer’s Conference, and frequent technical training events. Compared to many other standards organizations, membership in the PCI-SIG is very affordable.

The Integrators List

After you have successfully completed a Compliance Workshop and submitted a Compliance Checklist for your device, the PCI-SIG reviews the material and adds your device to the Integrators List under the appropriate category. Categories include components (silicon and IP cores), BIOS firmware, add-in cards, and PC-AT motherboards and systems. The Integrators List is your proof that your product passed the rigorous PCI-SIG tests and demonstrated interoperability with others.

This list is a valuable tool. As a developer, you might find yourself in the role of a customer, searching for silicon and IP cores that have been rigorously tested. Xilinx, as a vendor of silicon and IP cores, is proud to have a number of entries on the Integrators List. The low-cost Xilinx® LogiCORE™ PCI Express x1 Endpoint with PIPE Interface for Spartan™-3 devices is on the Integrators List. As of this writing, the Xilinx LogiCORE PCI Express x8 Endpoint for Virtex™-4 FX devices has passed the Compliance Workshop and Xilinx has submitted a Compliance Checklist for this product. By the time you read this, it should be on the Integrators List as well.

Similarly, if you are developing products that implement PCI-SIG technologies, you should make an effort to add your products to the Integrators List. Then, refer your customers to the list. Most customers welcome additional information to make intelligent purchases. Some discerning customers might even refuse to buy products that are not on the list.

If you are planning a product that integrates PCI, PCI-X, or PCI Express interfaces, join the PCI-SIG, participate in the Compliance Program, and get your product on the Integrators List. The success of your product may depend on it. For more information, visit the PCI-SIG website at www.pcisig.com.

Compliance Workshop Test Results Report				
Add-in Card		System		
Vendor:		Vendor:		
Product Name:		Product Name:		
Product Rev:		Product Rev:		
Driver Rev:		BIOS Rev:		
Product Type <input type="checkbox"/> Graphics <input type="checkbox"/> SCSI <input type="checkbox"/> IDE <input type="checkbox"/> LAN <input type="checkbox"/> Other _____ <input type="checkbox"/> Uses PCI to PCI Bridge		Motherboard CPU _____ Speed _____ # of Slots _____ MB Chipset _____ (Check all of the below that apply) Planar <input type="checkbox"/> Graphics <input type="checkbox"/> SCSI <input type="checkbox"/> IDE <input type="checkbox"/> LAN <input type="checkbox"/> Other _____		
	Test	Pass	Fail	Comments
	Detected/Configured			
	Demonstrated Operation			
Overall Evaluation		Pass	Fail	
Notes:				
Add-in Card Vendor		System Vendor		
Signature:		Signature:		
Print Name:		Print Name:		
Date:		Date:		

Figure 2 – PCI-SIG interoperability testing results report (published with permission from PCI-SIG)

your customers might encounter with your product. During these sessions, the participants set their own test procedure and must agree on what constitutes a pass or a fail. Generally, it is expected that you demonstrate some degree of functionality to substantiate that your interface is functional. (Figure 2 shows the interoperabil-

Integrators List. Should you fail, you can repeat the Compliance Workshop as many times as necessary.

Now, about this free lunch ... technically, it is not free, because you must be a PCI-SIG member, which currently costs \$3,000 per year per company. Membership also includes access to all the PCI-SIG specifi-

Successful DDR2 Design

Mentor Graphics highlights design issues and solutions for DDR2, the latest trend in memory design.

by Steve McKinney
HyperLynx Technical Marketing Engineer
Mentor Graphics
steven_mckinney@mentor.com

The introduction of the first SDRAM interface, in 1997, marked the dawn of the high-speed memory interface age. Since then, designs have migrated through SDR (single data rate), DDR (double data rate), and now DDR2 memory interfaces to sustain increasing bandwidth needs in products such as graphics accelerators and high-speed routers. As a result of its high-bandwidth capabilities, DDR and DDR2 technology is used in nearly every sector of the electronics design industry – from computers and networking to consumer electronics and military applications.

DDR technology introduced the concept of “clocking” data in on both a rising and falling edge of a strobe signal in a memory interface. This provided a 2x bandwidth improvement over an SDR interface with the same clock speed. This, in addition to faster clock frequencies, allowed a single-channel DDR400 interface with a 200 MHz clock to support up to 3.2 GB/s, a 3x improvement over the fastest SDR interface. DDR2 also provided an additional 2x improvement in bandwidth over its DDR predecessor by doubling the maximum clock frequency to 400 MHz. Table 1 shows how the progression from SDR to DDR and DDR2 has allowed today’s systems to maintain their upward growth path.

SDR		DDR				DDR2			
PC100	PC133	DDR - 200	DDR - 266	DDR - 333	DDR - 400	DDR2 - 400	DDR2 - 533	DDR2 - 667	DDR2 - 800
0.8	1.1	1.6	2.1	2.7	3.2	3.2	4.266	5.33	6.4
Single Channel Bandwidth (GB/s)									

Table 1 – The progression from SDR to DDR and DDR2 has allowed today’s systems to maintain their upward growth path. Speed grades and bit rates are shown for each memory interface.

With any high-speed interface, as supported operating frequencies increase it becomes progressively more difficult to meet signal integrity and timing requirements at the receivers. Clock periods become shorter, reducing timing budgets to a point where you are designing systems with only picoseconds of setup or hold margins. In addition to these tighter timing budgets, signals tend to deteriorate because faster edge rates are needed to meet these tight timing parameters. As edge rates get faster, effects like overshoot, reflections, and crosstalk become more significant problems on the interface, which results in a negative impact on your timing budget. DDR2 is no exception, though the JEDEC standards committee has created several new features to aid in dealing with the adverse effects that reduce system reliability.

Some of the most significant changes incorporated into DDR2 include on-die termination for data nets, differential strobe signals, and signal slew rate derating for both data and address/command signals. Taking full advantage of these new features will help enable you to design a robust memory interface that will meet both your signal integrity and timing goals.

On-Die Termination

The addition of on-die termination (ODT) has provided an extra knob with which to dial in and improve signal integrity on the DDR2 interface. ODT is a dynamic termination built into the SDRAM chip and memory controller. It can be enabled or disabled depending on addressing conditions and whether a read or write operation is being performed, as shown in Figure 1. In addition to being able to turn termination off or on, ODT also offers the flexibility of different termi-

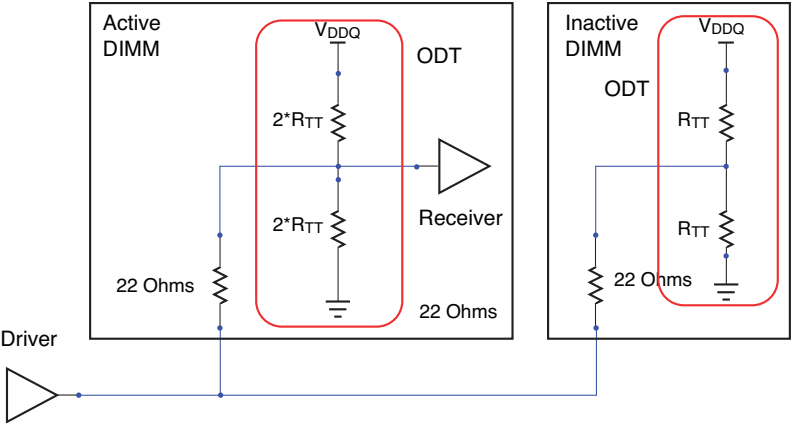


Figure 1 – An example of ODT settings for a write operation in a 2 DIMM module system where $R_{TT} = 150\text{ Ohms}$.

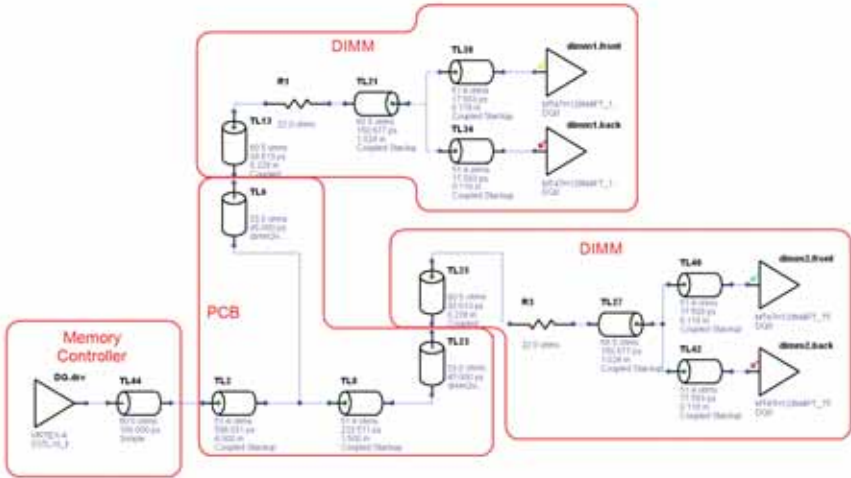


Figure 2 – The HyperLynx free-form schematic editor shows a pre-layout topology of an unbuffered 2 DIMM module system. Transmission line lengths on the DIMM are from the JEDEC DDR2 unbuffered DIMM specification.

nation values, allowing you to choose an optimal solution for your specific design. It is important to investigate the effects of ODT on your received signals, and you can easily do this by using a signal integrity software tool like Mentor Graphics’ HyperLynx product. Consider the example design shown in Figure 2, which shows a DDR2-533 interface (266 MHz) with two

unbuffered DIMM modules and ODT settings of 150 Ohms at each DIMM. You can simulate the effects of using different ODT settings and determine which settings would work best for this DDR2 design before committing to a specific board layout or creating a prototype. With the 150 Ohm ODT settings, Figure 3 shows significant signal degrada-




Figure 3 – The results of a received signal at the first DIMM in eye diagram form. Here, ODT settings of 150 Ohms are being used at both DIMM modules during a write operation. The results show there is an eye opening of approximately 450 ps outside of the VinAC switching thresholds.




Figure 4 – This waveform shows a significant improvement in the eye aperture with a new ODT setting. Here, the ODT setting is 150 Ohms at the first DIMM and 75 Ohms at the second DIMM. The signal is valid for 1.064 ns with the new settings, which is an increase of 614 ps from the previous ODT settings.

tion at the receiver, resulting in eye closure. The eye shows what the signal looks like for all bit transitions of a pseudo-random (PRBS) bitstream, which resembles the data that you might see in a DDR2 write transaction. Making some simple measurements of the eye where it is valid outside the VinhAC and VinIAC thresholds, you can see that there is roughly a 450 ps window of valid signal at the first DIMM module.

It is appropriate to try to improve this eye aperture (opening) at the first DIMM if possible, and changing the ODT setting is one of the options available for this. To improve the signal quality at the first

DIMM, you must change the ODT value at the second DIMM. Setting the ODT at the second DIMM to 75 Ohms and re-running the simulation, Figure 4 shows more than a 100 percent increase in the eye aperture at the first DIMM, resulting in a 1.06 ns eye opening. As you can see, being able to dynamically change ODT is a powerful capability to improve signal quality on the DDR2 interface.

With respect to a DDR interface, ODT allows you to remove the source termination, normally placed at the memory controller, from the board. In addition, the pull-up termination to VTT at the end of the data bus is no longer necessary. This reduces component cost and significantly improves the layout of the board. By removing these terminations, you may be able to reduce layer count and remove unwanted vias on the signals used for layer transitions at the terminations.

Signal Slew Rate Derating

A challenging aspect of any DDR2 design is meeting the setup and hold time requirements of the receivers. This is especially true for the address bus, which tends to have significantly heavier loading conditions than the data bus, resulting in fairly slow edge rates. These slower edge rates can consume a fairly large portion of your timing budget, preventing you from meeting your setup and hold time requirements.

To enable you to meet the setup and hold requirements on address and data

buses, DDR2's developers implemented a fairly advanced and relatively new timing concept to improve timing on the interface: "signal slew rate derating." Slew rate derating provides you with a more accurate picture of system-level timing on the DDR2 interface by taking into account the basic physics of the transistors at the receiver.

For DDR2, when any memory vendor defines the setup and hold times for their component, they use an input signal that has a 1.0V/ns input slew rate. What if the signals in your design have faster or slower slew rates than 1.0V/ns? Does it make sense to still meet that same setup and hold requirement defined at 1.0V/ns? Not really. This disparity drove the need for slew rate derating on the signals specific to your design.

To clearly understand slew rate derating, let's consider how a transistor works. It takes a certain amount of charge to build up at the gate of the transistor before it switches high or low. Consider the 1.0V/ns slew rate input waveform between the switching region, Vref to Vin(h/l)AC, used to define the setup and hold times. You can define a charge area under this 1.0V/ns curve that would be equivalent to the charge it takes to cause the transistor to switch. If you have a signal that has a slew rate faster than 1.0V/ns, say 2.0V/ns, it transitions through the switching region much faster and effectively improves your timing margin. You've added some amount of timing margin into your system, but that was with the assumption of using the stan-




Figure 5 – A 1V/ns signal has a defined charge area under the signal between Vref and VinhAC. A 2V/ns signal would require a + Δt change in time to achieve the same charge area as the 1V/ns signal. A 0.5V/ns signal would require a - Δt change in time to achieve the same charge area as the 1V/ns signal. This change in time provides a clearer picture of the timing requirements needed for the receiver to switch.

standard setup and hold times defined at 1.0V/ns. In reality, you haven't allowed enough time for the transistor to reach the charge potential necessary to switch, so there is some uncertainty that is not being accounted for in your system timing budget. To guarantee that your receiver has enough charge built up to switch, you have to allow more time to pass so that sufficient charge can accumulate at the gate.

Once the signal has reached a charge area equivalent to the 1.0V/ns curve between the switching regions, you can safely say that you have a valid received signal. You must now look at the time difference between reaching the V_{inAC} switching threshold and the amount of time it took for the 2.0V/ns to reach an equivalent charge area, and then add that time difference into your timing budget, as shown in Figure 5.

Conversely, if you consider a much slower slew rate, such as 0.1V/ns, it would take a very long time to reach the switching threshold. You may never meet the setup and hold requirements in your timing budget with that slow of a slew rate through the transition region. This could cause you to overly constrain the design of your system, or potentially limit the con-

figuration and operating speed that you can reliably support. But again, if you consider the charge potential at the gate with this slow slew rate, you would be able to subtract some time out of your budget (as much as 1.42 ns under certain conditions) because the signal reached an equivalent charge area earlier than when it crossed the V_{inAC} threshold.

To assist you in meeting these timing goals, the memory vendors took this slew rate information into account and have constructed a derating table included in the DDR2 JEDEC specification (JESD79-2B on www.jedec.com). By using signal derating, you are now considering how the transistors at the receiver respond to charge building at their gates in your timing budgets. Although this adds a level of complexity to your analysis, it gives you more flexibility in meeting your timing goals, while also providing you with higher visibility into the actual timing of your system.

Determining Slew Rate

To properly use the derating tables, it is important to know how to measure the slew rate on a signal. Let's look at an example of a slew rate measurement for the rising edge of a signal under a setup condition.

The first step in performing signal derating is to find a nominal slew rate of the signal in the transition region between the V_{ref} and $V_{in(h/l)AC}$ threshold. That nominal slew rate line is defined in the JEDEC specification as the points of the received waveform and V_{ref} and V_{inhAC} for a rising edge, as shown in Figure 6.

It would be a daunting task to manually measure each one of your signal edges to determine a nominal slew rate for use in the derating tables toward derating each signal. To assist with this process, HyperLynx simulation software includes built-in measurement capabilities designed specifically for DDR2 slew rate measurements. This can reduce your development cycle and take the guesswork out of trying to perform signal derating. The HyperLynx oscilloscope will automatically measure each of the edge transitions on the received waveform, reporting back the minimum and maximum slew rate values, which can then be used in the JEDEC derating tables. The scope also displays the nominal slew rate for each edge transition, providing confidence that the correct measurements are being made (see Figure 7).

The nominal slew rate is acceptable for use in the derating tables as long as the




Figure 6 – The waveform illustrates how a nominal slew rate is defined for a signal when performing a derating in a setup condition. The waveform is taken from the DDR2 JEDEC specification (JESD79-2B).

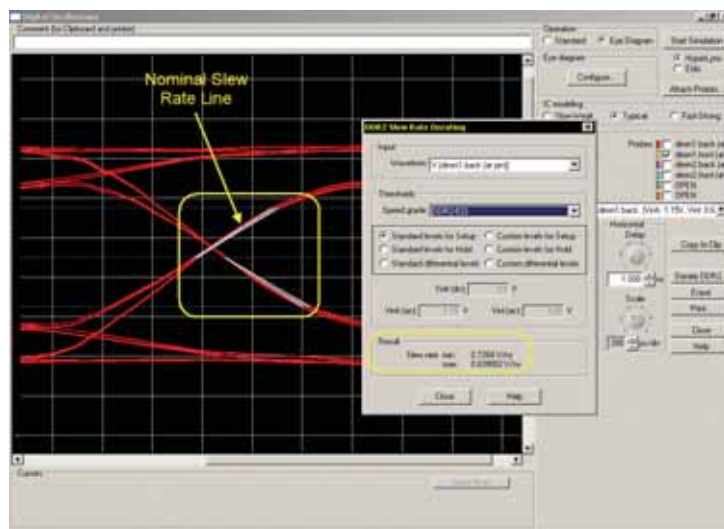


Figure 7 – The HyperLynx oscilloscope shows an automated measurement of the nominal slew rate for every edge in an eye diagram with the DDR2 slew rate derating feature. The measurement provides the minimum and maximum slew rates that can then be used in the DDR2 derating tables in the JEDEC specification.




Figure 8 – This waveform, taken from the DDR2 JEDEC specification, shows how a tangent line must be found if any of the signal crosses the nominal slew rate line. The slew rate of this tangent line would then be used in the DDR2 derating tables.

received signal meets the condition of always being above (for the rising edge) or below (for the falling edge) the nominal slew rate line for a setup condition. If the signal does not have clean edges – possibly having some non-monotonicity or “shelf”-type effect that crosses the nominal slew rate line – you must define a new slew rate. This new slew rate is a tangent line on the received waveform that intersects with V_{inhAC} and the received waveform, as shown in Figure 8. The slew rate

of this new tangent line now becomes your slew rate for signal derating.

You can see in the example that if there is an aberration on the signal edge that would require you to find this new tangent line slew rate, HyperLynx automatically performs this check for you. If necessary, the oscilloscope creates the tangent line, which becomes part of the minimum and maximum slew rate results. As Figure 9 shows, the HyperLynx oscilloscope also displays all of the tangent lines,




Figure 9 – The HyperLynx oscilloscope shows how the tangent line is automatically determined for you in the DDR2 slew rate derating feature. The slew rate lines in the display indicate that they are tangent lines because they no longer intersect with the received signal and V_{ref} intersection. The oscilloscope determines the slew rate of these new tangent lines for you and reports the minimum and maximum slew rates to be used in the derating tables.

making it easier to identify whether this condition is occurring.

For a hold condition, you perform a slightly different measurement for the slow rate. Instead of measuring from V_{ref} to the V_{inAC} threshold, you measure from V_{inDC} to V_{ref} to determine the nominal slew rate (shown in Figure 10). The same conditions regarding the nominal slew rate line and the inspection of the signal to determine the necessity for a tangent line for a new slew rate hold true here as well.

Conclusion

With the new addition of ODT, you’ve seen how dynamic on-chip termination can vastly improve signal quality. Performing signal derating per the DDR2 SDRAM specification has also shown that you can add as much as 1.42 ns back into your timing budget, giving you more flexibility in your PCB design and providing you with a better understanding of system timing.

Equipped with the right tools and an understanding of underlying technology, you will be able to move your designs from DDR to DDR2 in a reasonably pain-free process – realizing the added performance benefits and component-count reductions promised by DDR2. ●●

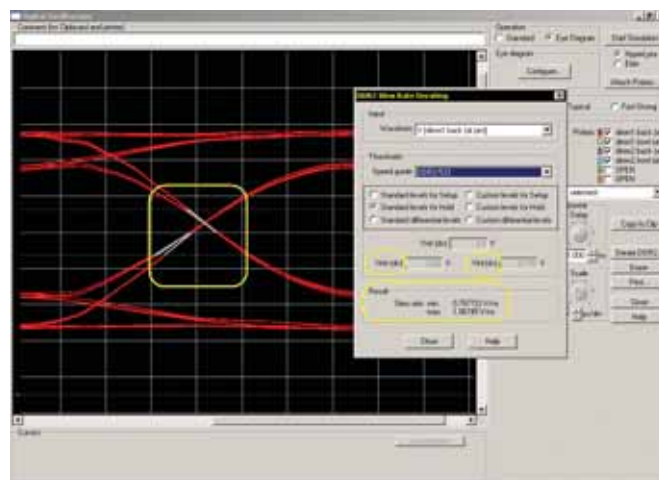


Figure 10 – The oscilloscope shows how a derating for a hold condition is being performed on the received signal. The DC thresholds are used in place of the AC switching thresholds, which are noted in the DDR2 derating dialog.

Board Design Panacea

The 7Circuits tool algorithmically solves FPGA pinout problems and synthesizes PC board schematics.

by Nagesh Gupta
Founder/CEO
Taray, Inc.
nagesh@tarayinc.com

PC board design is a cumbersome and time-consuming task. Although some of the steps require knowledge and intelligence to complete, most of the process is mundane and routine. Add FPGAs to the mix, and the complexity of the board grows significantly. FPGAs have a myriad of complex I/O rules that are multi-dimensional and can present difficult problems:

1. In most cases with large and complex designs, FPGA pinouts are hardly optimal, and non-optimal pinouts result in lower design performance. The cost of the PC board also increases because of the higher number of layers.
2. Today, pins for FPGAs are mostly selected manually. The pin selection is aided by large spreadsheets with signal names, I/O standards, clocking types, interface, and so on.
3. Drawing schematics is a fully manual process. The FPGA symbol has to be created, and then the FPGA pins have to be connected up to the interface pins. To avoid expensive mistakes, all of the pins have to be correctly connected. The configuration and power supply pins have to be connected as well.

Taray, which brought you the Xilinx® Memory Interface Generator, has developed a new tool called 7Circuits. 7Circuits solves these problems in an innovative way.

7Circuits

7Circuits is a highly intuitive tool that not only selects all of the FPGA pins but also generates PC board schematics for the FPGA and its interfaces.

7Circuits solves FPGA pin allocation problems algorithmically after considering the different constraints. At a higher level, the constraints that the tool considers are:

- Physical constraints. An example of a physical constraint is the physical placement of the FPGA and the interfaces on the PC board.

- Electrical constraints. I/O voltage levels, use of DCI termination, and I/O signaling standards form the electrical constraints.
- Logical constraints. The logical constraints are derived from the interface protocol. For example, if the FPGA is interfacing to a DDR2 memory, the DDR2 protocol will dictate the logical constraints of the interface.
- User preferences. You can tune the performance features of 7Circuits to achieve optimal results.
- FPGA. The location, type, and number of I/Os are among some of the parameters considered.

7Circuits comes with a board view on startup. You begin by placing the FPGA on the board. Next, you place the different components with which the FPGA interfaces. The FPGA and all of the components are shown to scale. The components should be located correctly with respect to the FPGA and the placement should be identical to the actual board placement. An example of the component and FPGA placement is shown in Figure 1.

7Circuits supports a large blend of standard components that you can select and place on the board. If a particular component is not already supported, 7Circuits provides a simple user interface to create the custom interface (alternately, Taray can help you create the interface). Defining the interface component correctly is key to the generation of correct outputs.

7Circuits can block off the pins selected outside the tool. Reading a UCF file with the pin location constraints supports this functionality. 7Circuits can also generate interfaces incrementally. In other words, you can open a saved project and add more interfaces to it without disturbing the existing connections.

If you want to use specific banks for certain interfaces, you can make 7Circuits do

it. You can also specify the percentage of pins to be used within each bank. This enables 7Circuits to be customized for any requirement.



Figure 1 – Placement of the FPGA and interface components on the board

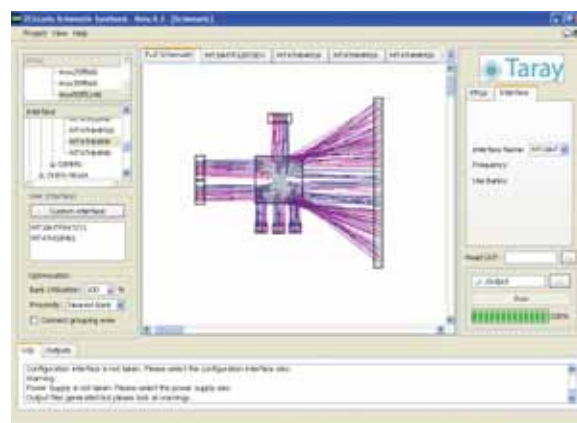


Figure 2 – A ratsnest view of the connections determined by 7Circuits

7Circuits goes through multiple optimization phases to select the pins optimally. After running through different optimization phases, 7Circuits displays the ratsnest connections to enable you to view any bowtie effects. Such interactive output at this stage is a key enabler to optimal results. You can try out different placements or different optimization options within 7Circuits to improve the bowtie effects. An example of the ratsnest is shown in Figure 2.

7Circuits produces a UCF file for pin locations; an EDIF schematics file for the FPGA, interface symbols, and schematics; and a top-level RTL file with all interface port declarations.

Key Advantages

7Circuits produces results with a holistic understanding of the problem space. This makes 7Circuits the first tool to bring system-level understanding into the FPGA solution. By doing so, 7Circuits comes up with the most optimal solution for pinout.

7Circuits reduces the time it takes to create an FPGA-based board from weeks to hours. The pinouts are very dependant on placement. In the current mode of operation, you do not have the luxury of trying out different placements to optimize results. Each placement and generation of the corresponding pinouts is at least a three-man-week task. This makes it impossible for you to try out various placements. With 7Circuits, you can try out four to five different placements and decide on the best placement within a few hours.

7Circuits offers you the added benefit of generating schematics for all of the mundane connections automatically. This task not only saves time, but also ensures correctness.

Here are some of the key advantages of using 7Circuits:

- 7Circuits connects all of the interface pins correctly. In addition, it connects up the power supplies to the right voltage levels.
- It connects Vref pins to the correct voltage levels depending on the I/O standard used.
- It reserves Vrp/Vrn pins when DCI is used. If DCI is used, the Vrp/Vrn pins are connected to the appropriate voltage levels.
- All configuration modes such as JTAG, slave serial, and master serial are supported. The connections are made automatically.

Because most of the mistakes are made in the unexciting and routine connections, the schematics are of a great benefit. They save greater than three man weeks of time and, more importantly, ensure correctness.




Figure 3 – Bowtie effects are significantly reduced, thus simplifying layout and reducing PCB layers.

Technology

The key to producing effective results is in the algorithms and the technology behind the tool. 7Circuits uses patent-pending technology to solve the issues identified in this article. Here are some of the key innovations in 7Circuits:

- Identifying and representing information. 7Circuits requires physical as well as architectural information on every interface and protocol. All of this information has been precisely identified for the components already supported. For new components, the tool provides a simple and intuitive GUI for you to give this information.
- Special signals are correctly identified and represented so that these signals can be associated to special pins. One example is the Xilinx RocketIO™ pins.
- 7Circuits also considers the logical and architectural aspects. Pins that are logically related will be placed together. This ensures quicker design convergence through the synthesis and PAR phases.
- 7Circuits constantly monitors the number of wire crossings and minimizes them, minimizing the number of board layers. This is key to reducing manufacturing costs.

- Length matching. Various heuristic algorithms are applied to reduce the delta length of signals that are to be length-matched. Applying these algorithms early on avoids long traces on the board. This improves signal quality and enables the PC board router to converge faster.

Results

7Circuits has been going through beta trials since Fourth Quarter 2005. Some of our customers have successfully laid out the board using our outputs.

Additionally, we have tested our results with many Xilinx reference designs. Our test process is as follows:

1. Generate a design for the same interfaces as the standard Xilinx reference board using 7Circuits.
2. Compare the ratsnest of the reference design against the ratsnest from the tool. In all cases, we found that 7Circuits produced a lower bowtie than the reference design.
3. Use the UCF generated by the tool and go through synthesis, build, map, PAR, and bitgen. Ensure that timing results from 7Circuits' UCF meet the reference design requirements.

Figure 3 shows an analytical comparison of the results for a memory reference board. The board has a Xilinx FPGA and interfaces with two DDR2 SDRAM DIMMs. This makes a 144-bit-wide interface. It also interfaces with DDR2 components to make a 24-bit-wide interface. The figure charts the frequency of line crossings against the number of line crossings. These comparisons clearly show the efficiency of the tool:

1. The original number of line crossings was 5,337. The line crossings with 7Circuits were reduced to 2,339 – a reduction of more than 50%.
2. There are 4,600 lines that cross each other manually. With 7Circuits, only 2,050 lines cross each other (1 point crossing each other).

Conclusion

Taray is committed to ensuring your success through the use of 7Circuits. Having created the Memory Interface Generator for Xilinx FPGAs, Taray's engineers have the depth of experience required to understand the issues facing you.

We are planning rich feature sets for future releases of 7Circuits, including:

- Schematics. 7Circuits will generate Orcad and DxDesigner schematics natively.
- Symbols. 7Circuits will be able to use symbols from your symbol library. Additionally, 7Circuits will also be able to use fractured (split) symbols to ensure that the schematics are consistent with your company standards.
- Parts. 7Circuits will support other Xilinx FPGA families and support more interface components.
- 7Circuits will offer a verification mode. This will be a great feature for you to check that your files are consistent and that your choices are optimal. You will be able to make incremental changes to improve your results.

A demo version of the 7Circuits tool is available at www.tarayinc.com. Revision 1.0 will be released in Second Quarter 2006.

Deliver Efficient SPI-4.2 Solutions with Virtex-4 FPGAs

Virtex-4 devices offer an ideal platform for source-synchronous designs like the widely adopted SPI-4.2 interface.

by Chris Ebeling
Principal Engineer
Xilinx, Inc.
chris.ebeling@xilinx.com

Krista Marks
Sr. Manager, IP Solutions Division
Xilinx, Inc.
krista.marks@xilinx.com

SPI-4.2 (System Packet Interface Level 4 Phase 2) is the Optical Internetworking Forum's recommended interface for the interconnection of devices for aggregate bandwidths of OC-192 (ATM and POS) and 10 Gbps (Ethernet), as illustrated in Figure 1.

In the last few years, this interface has become the de-facto standard on all leading 10 Gbps framer ASSPs and has been implemented directly on many next-generation network processors. SPI-4.2 has been broadly adopted because of its efficient interface, which offers high bandwidth with a low pin count and seamless handling of typical system requirements such as flow control, error insertion/detection, synchronization, and bus re-alignment.

The Xilinx® Virtex-4™ architecture provides an ideal platform for implementing SPI-4.2. The Xilinx SPI-4.2 LogiCORE™ IP targeting Virtex-4 devices provides a solution with one-third less resources, dramatic power savings, 1+ Gbps LVDS double-data-rate (DDR) I/O, and complete pin assignment flexibility.

SPI-4.2 LogiCORE IP

Xilinx has improved on its Virtex-II™ and Virtex-II Pro™ SPI-4.2 solution, already one of the smallest in the industry, and made it 30% smaller by leveraging new ChipSync™ technology in the Virtex-4 FPGA. ChipSync technology is supported on every pin of the Virtex-4 device family; thus the new SPI-4.2 LogiCORE IP can be targeted to any device pin-out. This allows you to select I/O pins that best fit your system and PCB requirements.

In addition, for those applications requiring multiple SPI-4.2 interfaces, the Virtex-4 FPGA's logic density, high pin count, and extensive clocking resources will support four or more full-duplex cores in a single device. Regardless of the performance your application requires,

Virtex-4 devices fully support the entire SPI-4.2 operating range, with high-speed LVDS support of data rates greater than 1 Gbps per pin.

ChipSync Technology

Xilinx introduced ChipSync technology in Virtex-4 FPGAs to enhance I/O capability when used for source-synchronous applications like SPI-4.2. ChipSync features are supported in every Virtex-4 I/O pin and include:

- New serial and de-serial (OSERDES and ISERDES) features. This enables logic built in the fabric to interface to the I/O at a fraction of the source-synchronous clock rate. The ISERDES also includes a Bitflip function. Bitflip allows you to shift the starting bit of deserialized data to achieve proper word alignment when linking multiple pins together (bus deskew).
- A new input delay (IDELAY) feature. This allows you to precisely adjust the input delay of each bit of a bus independently, in 78 ps increments. This provides a mechanism for tuning the interface timing to the system environment.

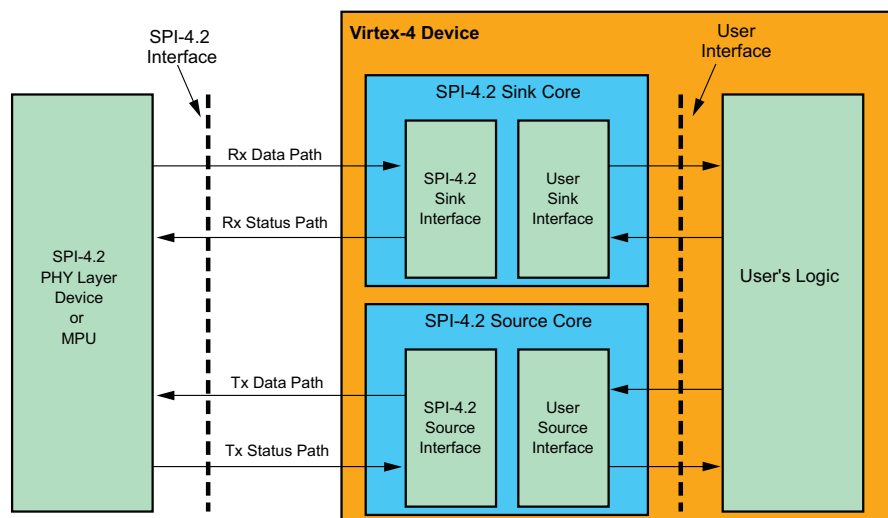


Figure 1 – Typical SPI-4.2 application

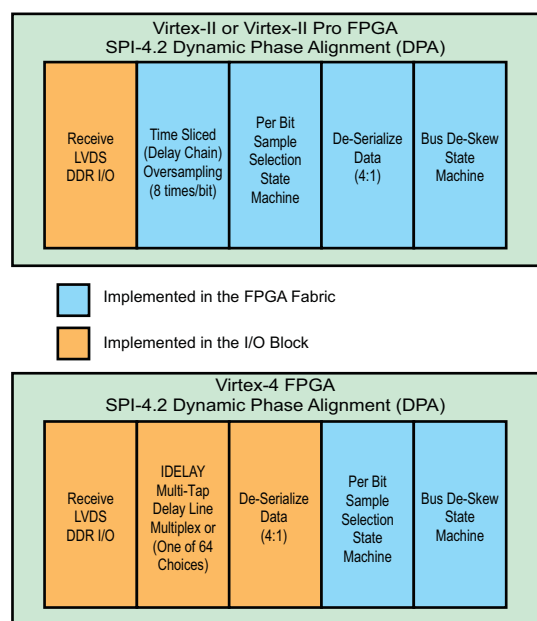


Figure 2 – DPA implementation in I/O logic for Virtex-II devices versus Virtex-4 devices

Additional DDR registers are now fully integrated into the input (ILOGIC) and output (OLOGIC) pins, simplifying the interface between the FPGA fabric and I/O blocks and supporting data transfer to and from the I/O logic on a single clock edge.

SPI-4.2 and ChipSync Technology

The SPI-4.2 interface has a DDR source-synchronous data bus that comprises 18 LVDS pairs (16 data bits, 1 control bit, and 1 clock). The SPI-4.2 source-synchronous clock varies from 311 MHz to 500 MHz.

For example, a typical OC-192 framer will require an aggregate bandwidth of 10 Gbps, which for a 16-bit dual data rate bus would require a data clock of at least 311 MHz, with 350 MHz a typical clock rate. The Xilinx SPI-4.2 LogiCORE IP easily meets your application requirements, regardless of performance, and with Virtex-4 ChipSync technology delivers a solution that is smaller and more flexible than prior FPGA implementations.

The SPI-4.2 core uses ChipSync technology to serialize egress data and de-serialize ingress data to a four-word (bus cycle) SPI-4.2 data stream at a lower clock rate. Operation of the core logic at a lower internal clock rate

allows you to implement high-frequency SPI-4.2 interfaces in the slowest speed grade Virtex-4 device.

The ISERDES and OSERDES functions allow the core logic to time multiplex and de-multiplex these four words to and from the I/O logic without using any CLB logic resources. The core logic need only operate at half the source-synchronous DDR clock rate. For example, a SPI-4.2 interface with a 500 MHz DDR reference clock would only require an FPGA fabric clock of 250 MHz – easily achievable in the Virtex-4 architecture.

As the frequency of the source-synchronous clock increases, data recovery at the receiving (sink) device becomes more challenging. The SPI-4.2 protocol provides a calibration data, or training pattern, that permits a receiving device to adjust its data sampling to the system interface timing. The process of tuning the interface to its particular timing is referred to as dynamic phase alignment (DPA).

Before Virtex-4 devices, Xilinx DPA solutions worked by over-sampling the input data and choosing the best sample from the group. This required valuable FPGA resources and careful control of the input data path in the FPGA fabric, restricting the SPI-4.2 interface pin placement. In Virtex-4 FPGAs, the IDELAY feature present in every I/O is ideally suited to perform this function, as shown in Figure 2. (See “Dynamic Phase Alignment with ChipSync Technology in Virtex-4 FPGAs,” also in this issue of the *Xcell Journal*).

The IDELAY features have two primary benefits for the SPI-4.2 core in Virtex-4 FPGAs:

- Integrating the IDELAY feature into the input pin (ILOGIC) reduces the FPGA resources required for DPA to less than 350 slices.
- The IDELAY function’s ability to adjust the data sampling point enables DPA to be implemented in the I/O – except for a small control state machine, which is implemented in the fabric. The state machine portion is fully synchronous and does not require a complex macro. Thus, there are no restrictions on SPI-4.2 pin assignments.

Clocking Resources

Virtex-4 FPGAs provide an unprecedented number of clock resources for implementing multiple SPI-4.2 interfaces in a single device. With the Virtex-II and Virtex-II Pro architectures, implementing more than two SPI-4.2 interfaces posed a clock management challenge. The abundance and flexibility of clock distribution in the Virtex-4 family solves this challenge, supporting as many SPI-4.2 interfaces as the device logic and I/O will allow.

In Virtex-4 devices, all devices have 32 global clock resources. No restrictions exist on global clock distribution other than a maximum of eight global clocks per clock region. All clock regions have access to any 8 of the 32 total global buffers, regardless of the requirements of other clock regions.

In addition to the eight global clocks, each region in the device has two regional clock buffers. The regional clock resources are ideal for interface clocking, like the source-synchronous clock scheme used by SPI-4.2. Note that even the smallest Virtex-4 device has a total of 48 available clock resources, each designed for low-skew clock distribution and clock power management. The SPI-4.2 LogiCORE IP can be configured to use either global or regional clock resources.

In Virtex-4 FPGAs, the global clock trees and associated buffers are implemented differentially, for best duty-cycle fidelity and greater common-mode noise rejection. With Virtex-II and Virtex-II Pro devices, if SPI-4.2 interface operates above 350 MHz, you must route the high-speed reference clock using two clock buffers to minimize duty-cycle distortion at the DDR registers.




Figure 3 – Illustration of four SPI-4.2 LogiCORE IP implemented on a Virtex-4 XC4VLX60 device

interfaces in the larger devices (Figure 3). The Virtex-4 clocking capability opens up a whole new class of SPI-4.2 applications, and provides an ideal platform for applications such as multiplexing and de-multiplexing, bridges, and switches.

	VIRTEX-II	VIRTEX-II PRO	VIRTEX-4
Power: Static Alignment @ 700 Mbps per LVDS Pair	1.9W	1.75W	1.55W
Power: Dynamic Alignment Performance per LVDS Pair	2.6W @800 Mbps	2.8W @944 Mbps	2.0W @1 Gbps
Speed Grades Supporting 800 Mbps per LVDS Pair	-6	-6, -7	-10, -11, -12

Table 1 – SPI-4.2 power estimates for Virtex-II, Virtex-II Pro, and Virtex-4 FPGAs

Because each global clock tree in Virtex-4 FPGAs is implemented differentially, only one clock buffer is required.

Not only does the Virtex-4 architecture have considerably more clock resources, but because they are distributed differentially, the SPI-4.2 LogiCORE IP requires fewer of them. These high-performance clock resources support as many as four SPI-4.2 interfaces in a mid-range device (LX40/LX60) and more than four SPI-4.2

Higher Performance at Lower Power

Virtex-4 silicon is manufactured with a triple-oxide process that reduces static power consumption by 40%. This will have a positive impact for all designs, including the SPI-4.2 interface, where the power savings are dramatic, as readily illustrated and summarized in Table 1.

With Virtex-4 devices, SPI-4.2 uses significantly less power than its Virtex-II and Virtex-II Pro predecessors, both because of

the enhanced 90 nm semiconductor process and because the LogiCORE IP uses 30% less fabric resources. At the same time, Virtex-4 FPGAs support 30% higher internal performance for SPI-4.2, with a maximum frequency of 250 MHz in the lowest speed grade (compared to 175 MHz in the lowest speed grade of Virtex-II and Virtex-II Pro devices). In addition, Virtex-4 FPGAs support 1+ Gbps LVDS for every I/O on the device.


This means that not only can you place multiple SPI-4.2 interfaces anywhere on the device, but for each implemented interface you get an aggregate bandwidth as high as 16+ Gbps. Designs that do not require this level of performance (such as more typical framer interfaces running at 10-12 Gbps) automatically get additional performance overhead that ensures ease of design integration and timing closure.

Conclusion

The Xilinx SPI-4.2 LogiCORE IP, coupled with Virtex-4 features, provides a highly efficient SPI-4.2 solution. We developed ChipSync technology that supports every I/O pin specifically for source-synchronous interfaces like SPI-4.2.

This technology enables you to design the most efficient SPI-4.2 solution, which uses significantly less resources (35% less), allows fully flexible device pin assignments (you choose the pinout), and supports extremely high interface speeds (1+ Gbps LVDS DDR I/O).

The higher performance is even more compelling because Virtex-4 FPGAs deliver it with lower power and significantly higher internal operating rates. The wealth of Virtex-4 clocking resources, combined with full pin assignment flexibility, opens up the possibility for new applications with multiple SPI-4.2 interfaces.

For more information about SPI-4.2 LogiCORE IP targeting Virtex-4 devices, please refer to this site at the Xilinx IP Center: www.xilinx.com/xlnx/xebiz/designResources/ip_product_details.jsp?key=DO-DI-POS4MC. A hardware demonstration is also available; for more information, contact your Xilinx representative. 

A Low-Cost PCI Express Solution

Spartan FPGAs are ideal for next-generation PCI applications and systems.

by Abhijit Athavale
Product Marketing Manager
Xilinx, Inc.
abhijit.athavale@xilinx.com

PCI has been the most widely used bus standard in the PC, server, and embedded markets for the past decade. Because PCI is limited by its shared, central arbitration-based architecture and system-synchronous clocking scheme, current and next-generation processors are outstripping its ability to keep up.

PCI's emerging replacement is PCI Express, a new connectivity standard that preserves the flexibility and familiarity of PCI while dramatically increasing bandwidth and performance. The controlling body for the PCI specification, the PCI SIG, has ratified PCI Express as the next-generation PCI. PCI Express-based products are now becoming available; shipments are expected to achieve high volume as early as 2006. Figure 1 shows the adoption forecast for PCI Express.

PCI Express uses serial I/O technology to create point-to-point connections and is reverse-compatible to PCI, preserving many original PCI advantages. It scales from a single lane (1x) to a 32 lane (32x) architecture, offering a bandwidth of 2.5 Gbps per lane. PCI 32/33 has a bandwidth of 1 Gbps, while PCI 64/66 has a bandwidth of 4 Gbps.

The 1x PCI Express implementation matches up very well with PCI 32/33, the most commonly used PCI interface across all markets. A two-lane implementation (5 Gbps) is an incremental improvement over

PCI 64/66. At the high end, a 32-lane PCI Express implementation supports a total of 80 Gbps, providing more than enough bandwidth to support the vast majority of next-generation applications.

Implementation Details

PCI Express is a three-layer specification: physical (PHY), logical, and transport, all defining separate functionalities. Also included in the specification are advanced features for hardware error recovery and system power management. (For more information about PCI Express, visit www.pcisig.com.)

Since 2000, Xilinx® has offered a line of PCI 32- and 64-bit solutions for Spartan™ series FPGAs. The most logical successor is a PCI Express solution using an external PHY chip paired with a Spartan-3 or Spartan-3E device. The PCI Express specification defines an interface to hook a PHY chip up to a separate device that houses the logical and transport layers

(called a PIPE interface – a white paper about this is available from Intel).

In the two-chip solution, the transport layer resides in a dedicated PHY chip, and the logic and transport layers reside in a Spartan FPGA. A broad range of PHY devices are available from manufacturers such as Genesys Logic, Philips Semiconductor, and Texas Instruments. PHY pricing will be less than \$10 for high volumes (250,000 units per year). (See the sidebar, "PHY Vendors," for contact information.) Xilinx has collaborated with Phillips Semiconductor and delivered this solution to our customers.

To implement the interface, Xilinx and several of our IP partners (including Eureka, GDA, and Northwest Logic) provide PIPE IP cores for Spartan-3 and Spartan-3E devices. A single-lane PCI Express controller requires approximately 500,000 gates (50% of a Spartan XC3S1000) for the logical and transport layer core, leaving the rest of the FPGA available for the user application (see

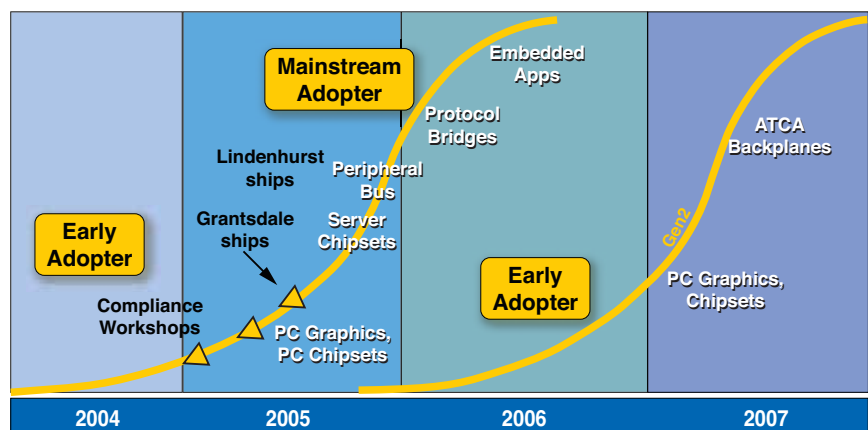


Figure 1 – PCI Express adoption forecast

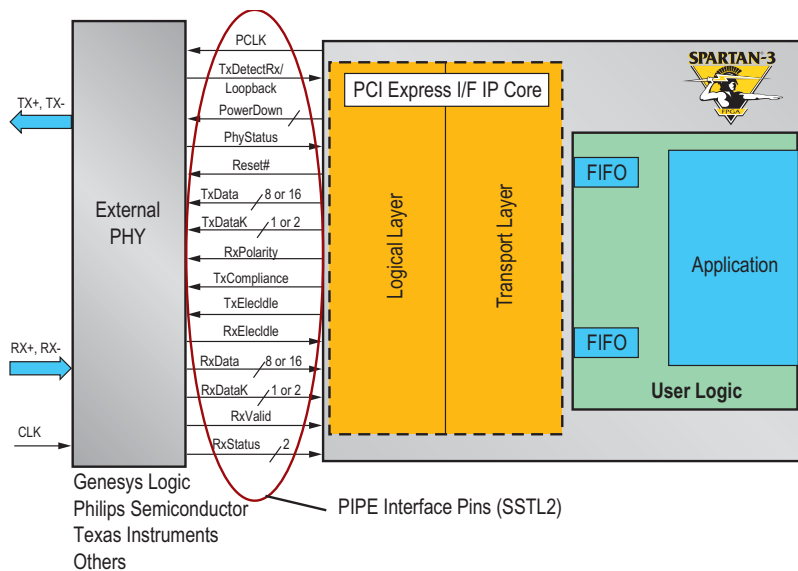


Figure 2 – PIPE interface between a Spartan FPGA and an external PHY

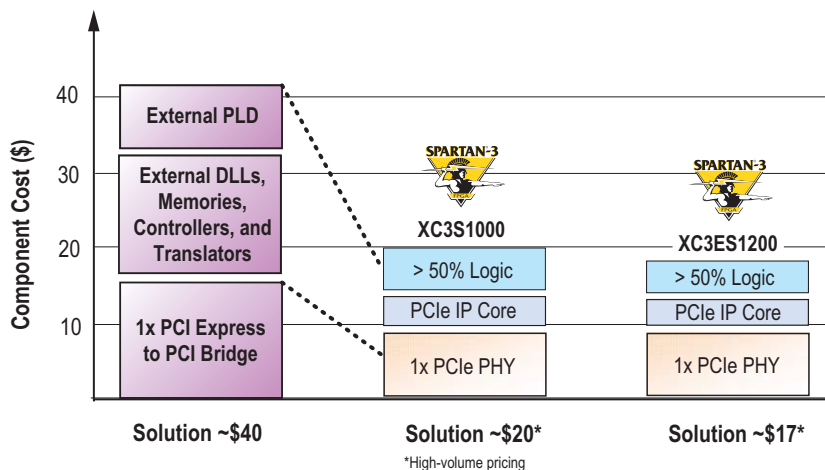


Figure 3 – Single-lane PCI Express implementation options

the “PCI Express Core IP” sidebar for details on Northwest Logic’s product and www.xilinx.com/pciexpress/ for details on PCI Express IP from our other IP partners.) Figure 2 shows the implementation of a PIPE interface using a Spartan FPGA and external PHY.

Figure 3 illustrates a range of options to implement a single-lane PCI Express interface. The cost of a standard-product option is fairly high (>\$40), making it tenuous for high-volume/low-cost applications. The Spartan options drop that cost substantially, and add the flexibility of programmable logic to integrate and implement other system capabilities. In 250K quantities (reasonable for typical consumer applications), the Spartan-3E version will cost approximately \$17.

Conclusion

In addition to reducing total costs, the Spartan FPGA + PHY option gives you substantial flexibility to build “PCI Express-to-anything” bridges and integrate other circuit elements. As most systems have a range of bandwidth requirements, preserving flexibility is important so that you can add lanes without dramatically changing the layout.

Spartan-3 and Spartan-3E FPGAs are available in a wide range of densities, and preserve migration up and down in overall bandwidth. And because FPGAs are fully reprogrammable post-deployment, they eliminate the risks associated with first-generation ASSPs and ASICs.

If you are currently using PCI for your interconnect standard and are architect-

ing your next-generation designs, you should consider the PCI Express option from Xilinx. We encourage you to find out how Spartan-3 and Spartan-3E FPGAs will help you meet your current and future design requirements. More information about Spartan-3 and Spartan-3E FPGAs, PCI Express IP, and compatible PHY devices is available at www.xilinx.com/pciexpress/.

PCI Express IP

PCI Express IP cores are available from multiple vendors including Xilinx and our partners. One such core from Northwest Logic is featured below.

Northwest Logic’s PCI Express Core is specifically designed for low-cost Spartan-3 FPGAs. A Spartan-3-based PCI Express design uses the Spartan-3 device with a low-cost physical interface for a PCI Express (PIPE)-compatible PHY chip. The PHY chip implements the low-level PCI Express physical layer, while the device takes care of the upper-level data link and transaction layers.

Another version of the PCI Express Core uses the internal MGTs in Virtex-II Pro and Virtex-4 FX FPGAs to provide a fully integrated PCI Express solution.

Northwest Logic’s PCI Express Core is one of the smallest PCI Express cores available, enabling you to target the smallest and consequently lowest cost FPGA. The core is provided with a comprehensive verification suite and expert support to ensure rapidly developed and validated designs.

Also available is a PCI Express Development Board for quickly prototyping a complete PCI Express System. A demo GUI, drivers, and PCI Express FPGA reference design are also included.

For more information (including pricing and core size for a particular FPGA family), visit the Northwest Logic website at www.nwlogic.com.

PHY Vendors

Genesys Logic

www.genesysamerica.com

Philips Semiconductor

www.semiconductors.philips.com

Texas Instruments

www.ti.com/pciexpress/

How to Detect Potential Memory Problems Early in FPGA Designs

System compatibility testing for FPGA memory requires methods other than traditional signal integrity analysis.

by Larry French
FAE Manager
Micron Semiconductor Products, Inc.
lfrench@micron.com

As a designer, you probably spend a significant amount of time simulating boards and building and testing prototypes. It is critical that the kinds of tests performed on these prototypes are effective in detecting problems that can occur in production or in the field.

DRAM or other memory combined in an FPGA system may require different test methodologies than an FPGA alone. Proper selection of memory design, test, and verification tools reduces engineering time and increases the probability of detecting potential problems. In this article, we'll discuss the best practices for thoroughly debugging a Xilinx® FPGA design that uses memory.

Memory Design, Testing, and Verification Tools

You can use many tools to simulate or debug a design. Table 1 lists the five essential tools for memory design. Note that this is not a complete list as it does not include thermal simulation tools; instead, it focuses only on those tools that you can use to validate the functionality and robustness of a design. Table 2 shows when these tools can be used most effectively.

This article focuses on the five phases of product development, as shown in Table 2:

- **Phase 1** – Design (no hardware, only simulation)
- **Phase 2** – Alpha (or Early) Prototype (design and hardware changes likely to occur before production)
- **Phase 3** – Beta Prototype (nearly “production-ready” system)

- **Phase 4** – Production

- **Phase 5** – Post-Production (in the form of memory upgrades or field replacements)

The Value of SI Testing

SI is not a panacea and should be used judiciously. SI should not be overused, although it frequently is. For very early or alpha prototypes, SI is a key tool for ensuring that your system is free of a number of memory problems, including:

- Ringing and overshoot/undershoot
- Timing violations, such as:
 - Setup and hold time
 - Slew rate (weakly driven or strongly driven signals)
 - Setup/hold time (data, clock, and controls)

Tool	Example
Electrical Simulations	SPICE or IBIS
Behavioral Simulations	Verilog or VHDL
Signal Integrity	Oscilloscope and probes; possibly mixed-mode to allow for more accurate signal capture
Margin Testing	Guardband testing and four-corner testing by variation of voltage and temperature
Compatibility Testing	Functional software testing or system reboot test

Table 1 – Memory design, test, and verification tools

- Clock duty cycle and differential clock crossing (CK/CK#)
- Bus contention

By contrast, SI is not useful in the beta prototype phase unless there are changes to the board signals. (After all, each signal net is validated in the alpha prototype.) However, if a signal does change, you can use SI to ensure that no SI problems exist with the changed net(s). Rarely – if ever – is there a need for SI testing in production.

SI is commonly overused for testing because electrical engineers are comfortable looking at an oscilloscope and using the captures or photographs as documentation to show that a system was tested (Figure 1). Yet extensive experience at Micron Technology shows that much more effective tools exist for catching failures. In fact, our experience shows that SI cannot detect all types of system failures.

Limitations of SI Testing

SI testing has a number of fundamental limitations. First and foremost is the memory industry migration to fine-pitch ball-grid array (FBGA) packages. Without taking up valuable board real estate for probe pins, SI is difficult or impossible because there is no way to probe under the package.

Micron has taken several hundred

Tool	Design	Alpha Proto	Beta Proto	Production	Post-Prod
Simulation – Electrical	Essential	Very Valuable	Limited Value	Rarely Used	No Value
Simulation – Behavioral	Essential	Very Valuable	Limited Value	Rarely Used	No Value
Signal Integrity	Unavailable	Critical	Limited Value	Rarely Used	No Value
Margin Testing	Unavailable	Essential	Essential	Essential	Essential
Compatibility	Unavailable	Valuable	Essential	Essential	Essential

Table 2 – Tools for verifying memory functionality versus design phase




Figure 1 – Typical signal integrity shot from an oscilloscope

thousand scope shots in our SI lab during memory qualification testing. Based on this extensive data, we concluded that system problems are most easily found with margin and compatibility testing. Although SI is useful in the alpha prototype phase, it should be replaced by these other tests during beta prototype and production.

Here are some other results of our SI testing:

- SI did not find a single issue that was not identified by memory or system-level diagnostics. In other words, SI found the same failures as the other tests, thus duplicating the capabilities of margin testing and software testing.

- SI is time-consuming. Probing 64-bit or 72-bit data buses and taking scope shots requires a great deal of time.
- SI uses costly equipment. To gather accurate scope shots, you need high-cost oscilloscopes and probes.
- SI takes up valuable engineering resources. High-level engineering analysis is required to evaluate scope shots.
- SI does not find all errors. Margin and compatibility testing find errors that are not detectable by SI.

The best tests for finding FPGA/memory issues are margin and compatibility testing.

Margin Testing

Margin testing is used to evaluate how systems work under extreme temperatures and voltages. Many system parameters change with temperature/voltage, including slew rate, drive strength, and access time. Validation of a system at room temperature is not enough. Micron found that another benefit of margin testing is that it detects system problems that SI will not.

Four-corner testing is a best industry practice for margin testing. If a failure is

How Does the Logic Analyzer (or Mixed-Mode Analysis) Fit In?

You may have noticed that Table 1 does not include logic analyzers. Although it is rare to find a debug lab that does not include this tool as an integral part of its design and debug process, we will not discuss logic analyzers in this article. Because of the cost and time involved, they are rarely the first tool used to detect a failure or problem in a system. Logic analyzers are, however, invaluable in linking a problem, after it has been identified, to its root cause. Like signal integrity (SI), logic analyzers should be used after a problem has been detected.

...margin and compatibility testing will identify more marginalities or problems within a system than traditional methods such as SI.

going to occur during margin testing, it will likely occur at one of these points:

- Corner #1: high voltage, high temperature
- Corner #2: high voltage, low temperature
- Corner #3: low voltage, high temperature
- Corner #4: low voltage, low temperature

There is one caveat to this rule. During the alpha prototype, margin testing may not be of value because the design is still changing and the margin will be improved in the beta prototype. Once the system is nearly production-ready, you should perform extensive margin testing.

Compatibility Testing

Compatibility testing refers simply to the software tests that are run on a system. These can include BIOS, system operating software, end-user software, embedded software, and test programs. PCs are extremely programmable; therefore, you should run many different types of software tests.

In embedded systems where the FPGA acts like a processor, compatibility testing can also comprise a large number of tests. In other embedded applications where the DRAM has a dedicated purpose such as a FIFO or buffer, software testing by definition is limited to the final application. Thorough compatibility testing (along with margin testing) is one of the best ways to detect system-level issues or failures in all of these types of systems.

Given the programmable nature of Xilinx FPGAs, you might even consider a special FPGA memory test program. This program would only be used to run numerous test vectors (checkerboard, inversions) to and from the memory to validate the DRAM interface. It could eas-

ily be written to identify a bit error, address, or row – in contrast to the standard embedded program that might not identify any memory failures. This program could be run during margin testing. It would be especially interesting for embedded applications where the memory interface runs a very limited set of operations. Likely, this type of test would have more value than extensive SI testing of the final product.

Tests Not To Ignore

The following tests, if ignored, can lead to production and field problems that are subtle, hard to detect, and intermittent.

Power-Up Cycling

A good memory test plan should include several tests that are sometimes skipped and can lead to production or field problems. The first of these is power-up cycling. During power-up, a number of unique events occur, including the ramp-up of voltages and the JEDEC-standard DRAM initialization sequence. Best industry practices for testing PCs include power-up cycling tests to ensure that you catch intermittent power-up issues.

Two types of power-up cycling exist: cold- and warm-boot cycling. A cold boot occurs when a system has not been running and is at room temperature. A warm boot occurs after a system has been running for awhile and the internal temperature is stabilized. You should consider both tests to identify temperature-dependent problems.

Self-Refresh Testing

DRAM cells leak charge and must be refreshed often to ensure proper operation. Self-refresh is a key way to save system power when the memory is not used for long periods of time. It is critical that the memory controller provide the prop-

er in-spec commands when entering and exiting self-refresh; otherwise, you could lose data.

Like power-up cycling, self-refresh cycling is a useful compatibility test. If an intermittent self-refresh enter or exit problem is present, repeated cycling can help detect it. Applications that do not use self-refresh should completely skip this test.


Sustaining Qualifications

One last area to consider is the test methodology for sustaining qualifications. That is, what tests should you perform to qualify a memory device once a system is in production? This type of testing is frequently performed to ensure that an adequate supply of components will be available for uninterrupted production.

During production a system is stable and unchanging. Our experience has shown that margin and compatibility testing are the key tests for sustaining qualifications. Because a system is stable, SI has little or no value.

Conclusion

In this article, our intent has been to encourage designers to rethink the way they test and validate FPGA and memory interfaces. Using smart test practices can result in an immediate reduction in engineering hours during memory qualifications. In addition, proper use of margin and compatibility testing will identify more marginalities or problems within a system than traditional methods such as SI. No “one-size-fits-all” test methodology exists, so you should identify the test methodology that is most effective for your designs.

For more detailed information on testing memory, see Micron’s latest DesignLine article, “Understanding the Value of Signal Integrity,” on our website, www.micron.com. 

Cabling Solutions 4X, 12X and Beyond

Contact Meritec
For 16X and More!

LATEST Signal Integrity Achievements!

- Near and crosstalk <2% at 100psec (20–80%) risetime (multiple aggressors)
- Meets InfiniBand Quad Data Rates
- Direct attach for unequalled signal integrity
- Equalized circuitry for long length signal integrity
- Innovative latch mechanism or thumbscrew alternative
- Angled egress for tight packaging

Solving interconnect problems for 39 years in MIL/AERO, Medical, Telecom/Datacom, Computers, Industrial, and other markets.

Specializing in High
Performance Connectors
& Cable Assemblies



888-MERITEC (637-4832) / www.meritec.com

Serving the
Automatic Test
Equipment Market



530-891-3551 / www.joysignal.com

Taking Rugged I/O Cabling and Connectors to Higher Speeds

You now have the option to take copper cabling and connectors to 12.8 Gbps and beyond.

by Tom Wirsing
Applications Engineer
Meritec
twirsing@meritec.com

John Sawdy
Manager of Signal Integrity Engineering
Meritec
jsawdy@meritec.com

Today's transceivers and chip sets are demanding better performance at higher data rates from the cables and connectors being used to carry serial data. Future systems promise to operate at even higher data rates. This performance is measured in terms of attenuation, crosstalk, and impedance control. The need for greater port density leads to that eternal conundrum: How can I package more signals in less space and at higher speeds without degrading performance?

Connector design is reaching a point where signal density requirements are severely limiting the ability to use interstitial ground planes to isolate single-ended or differential pair signals from one another. This same density requirement also makes the extensive use of ground pins problematic. Higher data speeds – and their correspondingly shorter signal wave lengths – also contribute to design problems by making which termination techniques you use and the number of terminations in a signal path critical.

Engineers are also encountering another problem caused by the need for increased density and the constraints imposed by signal integrity requirements. Increasing density and reducing the physical size of terminations leads to connectors and individual connections that are too fragile for many applications.

In this article, we'll describe some cutting-edge approaches that hold promise for taking copper to the speeds and densities that tomorrow's designs will require, and discuss how your designs can accomplish these goals while still providing a robust and reliable connection.

Faster and Farther

The speeds used today – and those proposed for the near future – were almost unthinkable just a few years ago. We can attribute this progress to significant developments in the understanding of (and ability to simulate) the conditions that high-speed signals encounter.

Advanced connector designs and manufacturing techniques allow connectors to approach transparency, enabling you to take advantage of the signal conditioning now embedded in many transceivers to design serial links between boards, racks, and cabinet bays. This enables copper cabling to be a feasible option for data center distances that exceed 15 meters at 6 Gbps – for example, T10 SAS-2 cabling.

Using software to simulate 2D, 2.5D, and 3D fields has allowed engineers to design for signal integrity before a part is molded or stamped. This has led to the use of different combinations of signal and ground pin size and placement. You can now optimize the placement of signal and ground pins to match the particular requirements of a given application, or use pins of different widths and thickness to increase the shielding available in an interconnect to control crosstalk.

Additionally, you must consider impedance control – and its direct impact on insertion loss and return loss – in the design phase. Signal and ground pins can have their size and shape contoured within the connector to minimize any changes in impedance. The physical shape of mating contacts at the

point(s) of contact also plays a role in defining the quality of the transmission line.

Using automated welded contact to wire terminations creates predictable and repeatable signal paths. You should take great care in establishing these contact weld programs to ensure that the termination zone size is kept to an absolute minimum.

Impedance control is also aided by bringing the shielding right up to the termination zone whenever possible. Tight manufacturing tolerances ensure a consistent physical geometry, which leads to consistent eye diagrams.

You can address all of these concerns in the design of a connector from its inception with simulation software. Figure 1 compares a next-generation differential contact with Meritec's current 4X cable.




Figure 1 – Near-end crosstalk (NEXT) at a 40 psec (20-80%) risetime: green = Meritec's current 4X with two nearest neighbors added together, crosstalk = 2%; white = Meritec's high-speed differential contact with six nearest neighbors added together, crosstalk = 1.4%.

I/O cabling offers unit-distance losses within the cable that are significantly less than those within the printed circuit board. For example, according to test report #335, conducted by co-author John Sawdy, the losses in a 3 meter, 26 AWG (American Wire Gauge) cable are roughly equivalent to a 12 inch, 4.5 mil trace in a low-loss substrate.

The silicon you choose can also help. Using signal conditioning techniques such as pre-emphasis, post-emphasis, and adaptive equalization can allow copper to meet the needs of the multi-gigabit data transmission community. The semiconductor industry continues to explore other, more advanced signaling techniques for the future, addressed in "Beyond 10 Gbps," presented at DesignCon 2005 by Tom Palkert of Xilinx.

Ten Pounds in a Five-Pound Bag

Two approaches spring immediately to mind when addressing the need for increased data density: increased signal density and increased data rate. The industry has chosen to attack the problem on both fronts simultaneously.

Increasing signal density is not as simple as putting more pins in a tighter grid. As the pins get closer, their fields encroach on each other and interact. This interaction can lead to substantial reductions in signal integrity.

Meritec has successfully designed and tested a co-planar board-to-board connector that allows for densities on the order of 44 differential pairs or 66 single-ended signals per square inch. We have also simulated designs for mezzanine connectors and I/O cables, which show great promise. This contact design should also be suitable for many zero insertion force (ZIF) applications.

Bolstered by this testing and simulation, Meritec is developing new higher density mezzanine connectors, co-planar board-to-board connectors, and 16X-48X rugged I/O cable assemblies that will meet current and future needs for data transmission at speeds from 2.5 Gbps per lane through 12.8 Gbps per lane and beyond. Figure 2 represents one proposed, stackable cable assembly configuration.




Figure 2 – High-speed differential cable assembly and mating board connector

Increasing the data rate requires you to use many of the techniques we have described. These techniques allow the connectors and cable assemblies to accommodate the data rate increases called for in such current and proposed standards and applications as PCIe, RapidIO, HyperTransport, custom Xilinx® RocketIO™ transceivers,

CX4, Qnet, NUMA, Myrinet, OIF CEI, Infiniband, SAS, Fibre Channel, or SONET extenders.

Rugged Construction is a Must

Engineers designing for high density and faster data rates will quickly encounter another serious obstacle. As pins become smaller and more tightly spaced, the possibilities of physical damage increase. Contact and connector designs that look great in a CAD program or signal integrity simulation can begin to show serious

As pin counts increase, so do mating forces. These forces can reach the point where they prohibit hand mating and require special mechanisms to engage. Standard pin and socket contacts typically have very high insertion forces; you should take steps to reduce these to manageable levels for high-pin-count connectors. The hermaphroditic contact design allows for extremely low insertion forces, despite high pin counts, while maintaining high normal force because of dual contact.

You must also offer protection from



Figure 3 – Edge view – individual row of mated contact pairs

flaws in the field. Fragile construction becomes an issue when connectors are subjected to use in uncontrolled environments and expected to perform.

Smaller contacts are a clear requirement of higher densities, so you must utilize another means of physical protection. One approach is to use a hermaphroditic contact. In a standard male/female interconnect, the female end occupies far more space than the male. Meritec's hermaphroditic contact eliminates the need for the female. Two identical contacts are used to mate with each other.

This contact design has been used in numerous applications, and we have developed a new version for higher data rate applications. The contact is extremely robust. Each mating pair provides two points of contact; these contacts are also more tolerant of minor contact misalignments that become more consequential as pitch decreases. To our knowledge, there has never been a single failure of our hermaphroditic contact. Figure 3 shows how one row of these contacts is mated.

stray tools, corners, and edges. One way to accomplish this is to recess the contacts in a unique honeycomb-like housing that surrounds the contacts to protect them in the unmated state. Aside from intentional damage, the contacts are well protected from physical damage during handling, mating, and un-mating.

Conclusion

It is clear that the demand for higher data rates and signal density will continue to grow. It is also clear that a fundamental understanding of signal integrity; the ability to accurately simulate electrical and magnetic fields; advances in semiconductor technology; innovative contact design; and progressive manufacturing techniques will allow copper to remain a viable signal conductor for the foreseeable future.

For more information, contact Meritec Customer Service at (440) 354-3148, e-mail info@meritec.com, or visit www.meritec.com.

GET PUBLISHED



WOULD YOU LIKE TO BE PUBLISHED IN I/O MAGAZINE?

It's easier than you think!

Submit an article draft for our Web-based or printed I/O Magazine and we will assign an editor and a graphic artist to work with you to make your work look as good as possible.

For more information on this exciting and highly rewarding program, please contact:

Forrest Couch
Executive Editor, Xcell Publications
xcell@xilinx.com



Developing a PCI Express Design?

JUMP START YOUR PROJECT, SAVE TIME, SAVE DEVELOPMENT COSTS



PX WAVE PCI EXPRESS DESIGN KIT Only \$1800 – next day shipping

FEATURES

- Passed June 2005 PlugFest #45 Compliance using Xilinx LogicCORE™ IP & EurekaTech Core.
- PCI-SIG® Integrators PCI Express® x1 add-in card
- Philips PX1011A PCI Express PHY
- 4 Video A-D inputs
- 1MB SRAM
- I2C for video and external controls
- 2 Logic Analyzer Mictor Connectors also usable for daughter boards
- Upto 4 XCF04VO20 Flash
- Standard: Spartan™-3 XC3S1000FG676-5C.
- Upgrade Options: XC3S1500/2000/4000
- All needed power supplies

WHAT IS INCLUDED

- Board User Manual, Printed Circuit Board, Full Schematics (PDF) & Gerbers, Full FPGA UCF, Example Bitfile

APPLICATIONS

- Prototyping of PCI Express x1 MAC Cores
- Prototyping of RTL applications to work with PCI Express Cores and PHY
- Analysis and Evaluation of PCI Express
- Software and Driver Development

PIPE-C BOARD AVAILABLE

Prototype any PCIe PHY with any core. Get the first commercially available PIPE-C board. Has a Virtex 4 FX with an optional MGT x8 interface. This board has lead times. Contact us ASAP for details.

Tentmaker Systems Consulting Group

Tentmaker Systems is a San Jose consulting group available for your one-stop fast turn-around of:

- Architecture
- High Speed Board Design, Layout, Fab and Assembly
- System Specification, Architecture and Design
- FPGA Architecture and RTL Design – at your location or off-site
- ASIC Verification in FPGAs (includes Board level architecture, design and debug)
- ASIC & FPGA Simulation

Experienced PCI Express & Networking Designers

- All our consultants have 13-25 years of experience.
- On PCI-SIG® Integrators List
- Over 8 PCI Express Boards Delivered; 3 more currently in design for customers
- Co-Contributor and participant for the PCI Express PIPE C Connector Specification – 3 boards designed
- Successful Architecture and Implementation of multiple FPGA-based PCIe® cores for Xilinx and Philips and others with Board validation and compliance at PlugFest
- ASIC validation in multiple FPGAs of complex PCI Express Chips – ASICs in production.
- Customer List Includes: Xilinx, Philips Semiconductors, NetLogic Microsystems, Luminous Networks, Xalted Info Systems, Reliancy, Yvent and others.



Tentmaker

Contact Information:

Neil Mammen
Tentmaker Systems
neil1@tentmakersystems.com
www.tentmakersystems.com

New PCI Express Solution Simplifies Video Security Applications

How to implement a video security system using PCI Express.

by Neil Mammen
Technologist
Tentmaker Systems Consulting Group
neil1@tentmakersystems.com

An ideal video security device would be able to collect live compressed or uncompressed video, monitor each stream for motion, record all of the streams, and save the video to a hard disk (or write it out to a shuttle DVD system). However, these systems would end up costing more than what an average security consumer would be willing to pay, especially when the per-stream cost expands after four to eight streams.

This ideal security device would also be able to monitor motion and only record those streams that had motion, saving both disk space and bandwidth. One way to do this is to capture the video streams on a low-cost PCI card and store the input streams to a hard drive after performing some CPU processing. This reduces the cost of a built-from-scratch dedicated system by using off-the-shelf hardware.

However, there's a problem with the densities and processing power. If you want to store uncompressed streams, you quickly run out of bus bandwidth on a shared PCI bus, as well as processing power on its CPU. For instance, the maximum bus bandwidth in both directions on an off-the-shelf, low-cost (\$800) PC would be between 2-4 Gbps. Spreading the video-capture capabilities among multiple PCI cards does not help the bus bandwidth issue because 4 Gbps is the top limit of the entire PCI bus – not just of any single slot. And once shared with other cards on the PCI bus, you would be lucky to get about 1 Gbps of throughput in both directions.

You could use compression chips to reduce the bandwidth on the bus, but this would increase your cost and limit you to existing MPEG chipsets, without an easy way to perform additional processing or special motion-detection functions that are key for the security market.

Uncompressed video (once stripped of blanking) is around 165 Mbps of data. Thus, with 1 Gbps of total bandwidth, you are limited to at most a mix of six capture or playback devices of uncompressed video on one PCI bus PC.

PCI Express to the Rescue

PCI Express (PCIe) technology provides a significant jump in throughput to PC users. PCI Express is broken down into lanes. Each lane comprises a differential pair in each direction. Each differential pair provides a 2.5 Gbps stream with an 8b/10b encoding scheme, with 2 Gbps of data throughput per pair in that direction. But even more impressive, each PCIe slot on a motherboard has its own lanes that are not shared with any other slot. Each slot comes in configurations of 16 lanes (also called a x16 or "by 16"), 8 lanes (x8), 4 lanes (x4), or 1 lane (x1).

Today you can purchase an off-the-shelf low-cost PC motherboard with one x16 PCIe graphics slot and two x1 PCIe card slots, as well as two or more regular PCI slots. Server models come with x4 or x8 PCIe slots. You can even use the x16 graphics slot for another function if you do not need a graphics function, or if it is already integrated into the motherboard.

Thus, PCIe allows each card to provide data from 2 Gbps in a x1 lane to as much as 32 Gbps in a x16 lane.

You can immediately see the advantages. Most low-cost motherboards are now capable of supporting more than 36 Gbps of video data in both directions (this is very dependent on the speed of the peripherals). Bandwidth-wise, this means that each PC motherboard could technically support more than 200 uncompressed video captures or playbacks in each direction (although you will run into limitations on the peripherals before you get to this point).

Using low-cost Xilinx® FPGAs, you can go one step further and provide motion

detection as well as some hardware assist in the FPGA. A high-speed DDR DRAM will allow the CPU to perform the easier portions of the compression and store data only when there is motion, thus reducing the storage requirements. Of course, you will have to make some compromises, depending on if the streams will be played back on standard DVD players.

The Tentmaker PCIe Prototyping Solution

The Tentmaker PX Wave PCIe Design Kit shown in Figure 1 (the block diagram is shown in Figure 2) is one possible video security solution, comprising four video capture devices from Philips, a Xilinx Spartan™-3 FPGA, and a Philips PCI Express x1 PHY. It is designed as a low-cost (\$1800) evaluation system for companies



Figure 1 – PX Wave PCIe Design Kit board

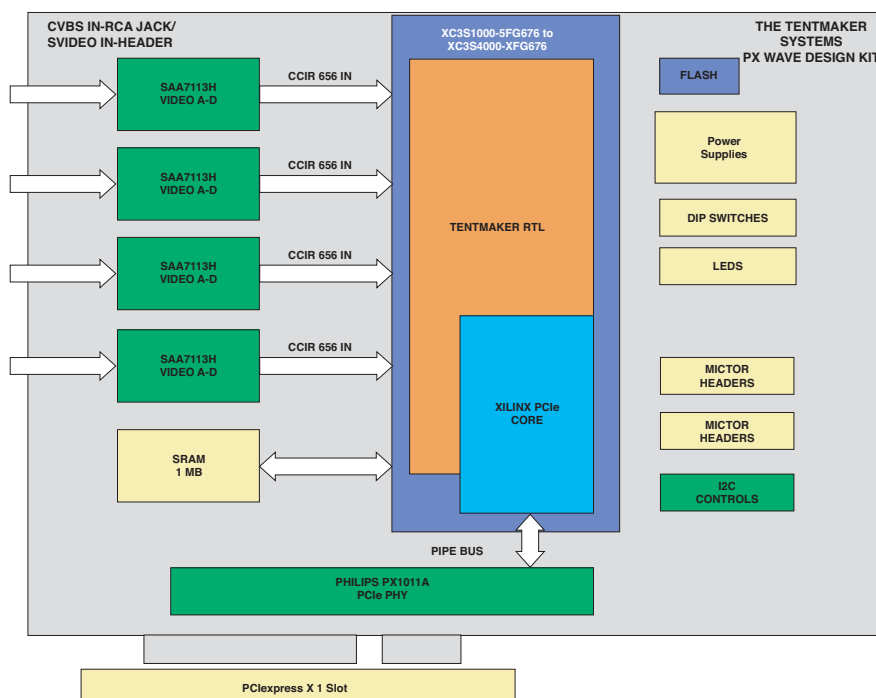


Figure 2 – PX Wave PCIe Design Kit block diagram

PCI Express is becoming more pervasive. As more applications like video continue to grow and require more bandwidth, PCI Express is well-suited to meet the related demands.

that want to get a jumpstart on designing boards to address this market. It contains all of the components except the high-speed DDR DRAM (an SRAM is used in this version). The PX Wave Design Kit allows companies to eliminate much of the learning curve associated with PCIe designs.

It would be easy to expand this design to use 16 video captures, the Xilinx PCIe x4 core, and an associated x4 PHY (Figure 3). Naturally, a cable harness would need to be

could include complex motion estimation, VLC (variable length code) generation, and other such preprocessing.

Other applications could use a high-speed, high-resolution camera that requires the extra bandwidth of PCIe for a single stream. You could also add hardware processing by using a preprocessing FPGA, as I've described.

For storage, it is also useful to be able to automatically add a graphic overlay showing the capture time and camera number.

ma) streams. The video analog-to-digital converters produce four independent digital CCIR656 streams, which are then fed into a low-cost Spartan-3 device for preprocessing. In the FPGA, the video data is stripped of blanks and syncs, packetized appropriately for PCIe, and fed to the Xilinx PCIe core. Software can then take the input video and display it, process it, or store it to disk.

PCI Express is straightforward if you follow some simple design principles. The high speed 2.5 Gbps lines are differential and thus simple to lay out, as long as the traces are length-matched and you adhere to some standard layout methodologies. More complicated is the PIPE bus that goes between the FPGA and the PHY. This bus must support signals at 250 MHz and each direction must be length-matched.

Conclusion

PCI Express is becoming more pervasive. As more applications like video continue to grow and require more bandwidth, PCI Express is well-suited to meet the related demands.

With connectors that allow you to add daughter boards and easily debug, the PX Wave PCIe Design Kit provides an easy way for companies to prototype generic PCI Express cards for security, video, and any generic application. In fact, the Xilinx PCIe core and the Philips PCIe PHY were prototyped and passed PCI-SIG PlugFest in the Summer of 2005 using the PX Wave Design Kit.

For more information, visit www.tentmakersystems.com.

Tentmaker Systems Consulting Group is part of a group of companies working on a PIPE-C specification. This is a connector specification that provides a standard connection between PHYs and PCIe cores, allowing various cores to be easily tested and used with various PHYs. The PIPE-C connector is available on various Tentmaker Systems boards.

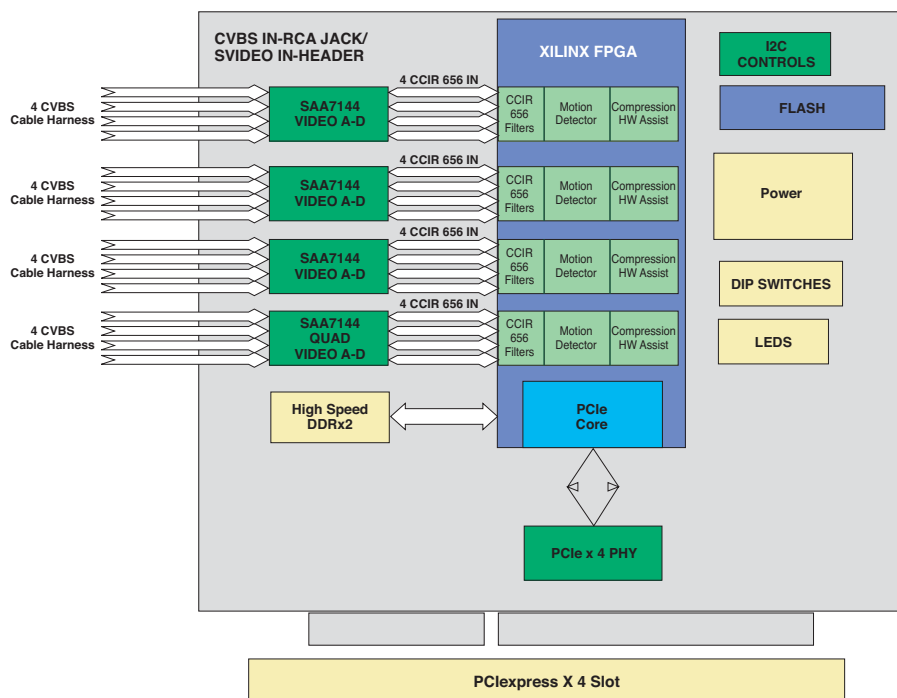


Figure 3 – Block diagram of a 16-input security solution with hardware assist

created for the 16 video inputs because of rear-plate surface-area limitations. The hardware assist could include a simple FPGA motion detector that provides an alarm and directs the PC to only record those streams that have motion, or you could dynamically allocate bandwidth so that cameras with the most motion get more bandwidth. Other hardware assists

You can do all of this very easily in software or hardware, assuming that the system has access to the compression algorithm. You can also insert text into the closed-caption fields at this point.

In the PX Wave PCIe Design Kit, four Philips SAA7113 chips are used to capture four input analog CVBS (composite video, blank, and sync) or Y/C (luma and chro-

FREE on-line training *with* Demos On Demand



A series of compelling, highly technical product demonstrations, presented by Xilinx experts, is now available on-line. These comprehensive videos provide excellent, step-by-step tutorials and quick refreshers on a wide array of key topics. The videos are segmented into short chapters to respect your time and make for easy viewing.

Ready for viewing, anytime you are

Offering live demonstrations of powerful tools, the videos enable you to achieve complex design requirements *and* save time. A complete on-line archive is easily accessible at your fingertips. Also, a *free* DVD containing all the video demos is available at www.xilinx.com/dod. Order yours today!



www.xilinx.com/dod



Pb-free devices
available now

Designing a Spartan-3 FPGA DDR Memory Interface

Xilinx provides many tools to implement customized DDR memory interfaces.

by Rufino Olay
Marketing Manager, Spartan Solutions
Xilinx, Inc.
rufino.olay@xilinx.com

Karthikeyan Palanisamy
Staff Engineer, Memory Applications Group
Xilinx, Inc.
karthi.palanisamy@xilinx.com

Memory speed is a crucial component of system performance. Currently, the most common form of memory used is synchronous dynamic random access memory (SDRAM).

The late 1990s saw major jumps in SDRAM memory speeds and technology because systems required faster performance and larger data storage capabilities. By 2002, double-data-rate (DDR) SDRAM became the standard to meet this ever-growing demand, with DDR266 (initially), DDR333, and recently DDR400 speeds.

DDR SDRAM is an evolutionary extension of “single-data-rate” SDRAM and provides the benefits of higher speed, reduced power, and higher density components. Data is clocked into or out of the device on both the rising and falling edges of the clock. Control signals, however, still change only on the rising clock edge.

DDR memory is used in a wide range of systems and platforms and is the computing memory of choice. You can use Xilinx® Spartan™-3 devices to implement a custom DDR memory controller on your board.

Interfacing Spartan-3 Devices with DDR SDRAMs

Spartan-3 platform FPGAs offer an ideal connectivity solution for low-cost systems, providing the system-level building blocks necessary to successfully interface to the latest generation of DDR memories.

Included in all Spartan-3 FPGA input/output blocks (IOB) are three pairs

of storage elements. The storage-element pair on either the output path or the three-state path can be used together with a special multiplexer to produce DDR transmission. This is accomplished by taking data synchronized to the clock signal's rising edge and converting it to bits synchronized on both the rising and falling edge. The combination of two registers and a multiplexer is referred to as double-data-rate D-type flip-flop (FDDR).

Memory Controllers Made Fast and Easy

Xilinx has created many tools to get designers quickly through the process of building and testing memory controllers for Spartan devices. These tools include reference designs and application notes, the Memory Interface Generator (MIG), and more recently, a hardware test platform.

Xilinx application note XAPP454, “DDR2 SDRAM Memory Interface for Spartan-3 FPGAs,” describes the use of a Spartan-3 FPGA as a memory controller,




Figure 1 – Read operation timing diagram

with particular focus on interfacing to a Micron MT46v32M16TG-6T DDR SDRAM. This and other application notes illustrate the theory of operations, key challenges, and implementations of a Spartan-3 FPGA-based memory controller.

DDR memories use non-free-running strobes and edge-aligned read data (Figure 1). For 333 Mbps data speeds, the memory strobe must be used for higher margins. Using local clocking resources, a

delayed strobe can be centered in the data window for data capture.

To maximize resources within the FPGA, you can explore design techniques such as using the LUTs as RAMs for data capture – while at the same time minimizing the use of global clock buffers (BUFGs) and digital clock managers (DCMs) – as explained in the Xilinx application notes. Results are given with respect to the maximum data width per FPGA side for either right and left

or top and bottom implementations. Implementation challenges such as these are mitigated with the new Memory Interface Generator.

Xilinx created the Memory Interface Generator (MIG 007) to take the guesswork out of designing your own controller. To create the interface, the tool requires you to input data including FPGA device, frequency, data width, and banks to use. The interactive GUI (Figure 2) generates the RTL, EDIF, SDC, UCF, and related document files.

As an example, we created a DDR 64-bit interface for a Spartan XC3S1500-5FG676

using MIG. The results in Table 1 show that the implementation would use 17% of the slices, leaving more than 80% of the device free for data-processing functions.

Testing Out Your Designs

The last sequence in a design is the verification and debug in actual hardware. After using MIG 007 to create your customized memory controller, you can implement your design on the Spartan-3 Memory Development Kit, HW-S3-SL361, as shown in Figure 3. The \$995 kit is based on a Spartan-3 1.5M-gate FPGA (the XC3S1500) and includes additional features such as:

- 64 MB of DDR SDRAM Micron MT5VDDT1672HG-335, with an additional 128 MB DDR SDRAM DIMM for future expansion
- Two-line LCD
- 166 MHz oscillator
- Rotary switches
- Universal power supply 85V-240V, 50-60 MHz



Figure 3 – Spartan-3 memory development board (HW-S3-SL361)

Conclusion

With the popularity of DDR memory increasing in system designs, it is only natural that designers use Spartan-3 FPGAs as memory controllers. Implementing the controller need not be difficult.

For more information about the application notes, GUI, and development board, please visit www.xilinx.com/products/design_resources/mem_corner/index.htm.

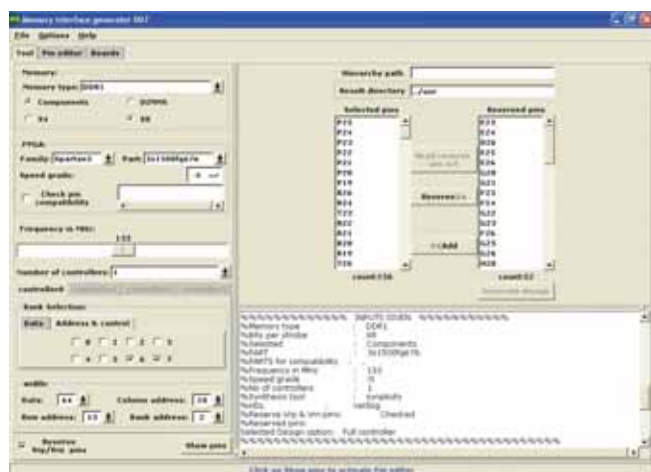


Figure 2 – Using the MIG 007 to automatically create a DDR memory controller

Feature	Utilization	Percent Used
Number of Slices	2,277 out of 13,312	17%
Number of DCMs	1 out of 4	25%
Number of External IOBs	147 out of 487	30%

Table 1 – Device utilization for a DDR 64-bit interface in an XC3S1500 FPGA

Introduction

The LogiCORE™ 10-Gigabit Ethernet MAC core is a single-speed full-duplex 10 Gbps Ethernet Media Access Controller (MAC) solution that enables the design of high-speed Ethernet systems and subsystems.

Features

- Designed to 10-Gigabit Ethernet specification IEEE 802.3ae-2002
- Choice of external XGMII or internal FPGA interface to PHY layer
- Cut-through operation with minimum buffering for maximum flexibility in client-side interfacing
- Supports Deficit Idle Count for maximum data throughput; maintains minimum IFG under all conditions and provides line rate performance
- Configured and monitored through a microprocessor-neutral management interface
- Comprehensive statistics gathering with statistic vector outputs
- Supports flow-control in both directions
- MDIO STA master interface to manage PHY layers
- Extremely customizable; trade off resource usage against functionality
- Available under SignOnce license program
- Supports VLAN, jumbo frames, and WAN mode

LogiCORE Facts				
Core Specifics				
Device Family	Virtex™-II, Virtex-II Pro, Virtex-4			
Speed Grades	-5 for Virtex-II -5 for Virtex-II Pro -10 for Virtex-4			
Resources Used ⁽¹⁾	Slices	LUTs	FFs	Block RAM
	3777	3703	4211	0
Special Features	Delivered through Xilinx CORE Generator™			
Provided with Core				
Documentation	Product Specification User Guide Getting Started Guide			
Design File Formats	EDIF and NGC netlist			
Constraints File	UCF			
Verification	VHDL test bench Verilog test fixture			
Example Design	VHDL and Verilog			
Design Tool Requirements				
Xilinx Implementation Tools	ISE™ 8.1i			
Simulation	Mentor ModelSim® Cadence® IUS			
Synthesis	XST			
Support				
Provided by Xilinx, Inc. @ www.xilinx.com/support				

1. Numbers are approximate for default configuration. See "Device Utilization" on page 19 for a complete description device utilization by configuration.

Introduction

The LogiCORE™ Tri-Mode Ethernet Media Access Controller (TEMAC) core supports half-duplex and full-duplex operation at 10 Megabits-per-second (Mbps), 100 Mbps, and 1 Gbps.

Features

- Designed to IEEE 802.3-2002 specification
- Reconciliation sublayer with GMII/MII or RGMII Interface
- Configurable half-duplex and full-duplex operation
- Configured and monitored through an optional independent microprocessor-neutral interface
- Configurable flow control through MAC Control pause frames; symmetrically or asymmetrically enabled
- Optional MDIO interface to managed objects in PHY layers (MII Management)
- Optional Address Filter with a selectable number of address table entries
- Optional clock enables to reduce clock resource usage
- Support of VLAN frames to specification IEEE 802.3-2002
- Configurable support of jumbo frames of any length
- Configurable inter-frame gap adjustment
- Configurable in-band FCS field passing on both transmit and receive paths
- Available under the terms of the [SignOnce IP Site License](#) agreement

LogiCORE Facts	
Core Specifics	
Supported Device Family	Virtex™-4, Virtex-II, Virtex-II Pro, Spartan™-3, Spartan-3E
Speed Grade	<ul style="list-style-type: none"> • Virtex-4 -10 • Virtex-II, Spartan-3, Spartan-3E¹ -4 • Virtex-II Pro -5
Performance	10 Mbps, 100 Mbps, 1 Gbps
Core Resources	
Slices	1019-1801 slices ²
LUTs	1273-2160
FFs	1030-1809
DCM	1 (Spartan-3, Spartan-3E, Virtex-4, or RGMII only)
BUFG	2-6
PPC	0
IOB-FF	79
Core Highlights	
Designed to IEEE 802.3-2002	Hardware Verified
Provided with Core	
Documentation	Product Specification User Guide Getting Started Guide
Design File Formats	NGC Netlist, HDL example Design, Demonstration test bench, scripts
Constraints File	User Constraints File (.ucf)
Example Design	Tri-Mode Ethernet MAC with GMII/MII or RGMII interface
	Demo test environment
Design Tool Requirements	
Supported HDL	VHDL and/or Verilog®
Synthesis	XST 8.1i
Xilinx Tools	ISE™ 8.1i
Simulation tools	Mentor ModelSim®, Cadence IUS ³
Support	
Provided by Xilinx @ www.xilinx.com/support	

1. Spartan-3E devices support only the GMII protocol.
2. Precise number of slices depends on user configuration.
3. Scripts provided for Mentor ModelSim and Cadence IUS only.

Introduction

The LogiCORE™ Virtex-4 Embedded Tri-mode Ethernet Media Access Controller (MAC) Wrapper automates the generation of HDL wrapper files for the Embedded Tri-Mode Ethernet MAC in Virtex-4™ FX devices using the CORE Generator™ tool.

VHDL and Verilog instantiation templates are available in the Libraries Guide for the Virtex-4 Ethernet MAC primitive; however, due to the complexity and the large number of ports, the CORE Generator simplifies integration of the Ethernet MAC by providing HDL examples based on user-selectable configurations.

Features

- Allows selection of one or both of the two Ethernet MACs (EMAC0/EMAC1) from the Embedded Ethernet MAC primitive
- Connects the EMAC0/EMAC1 Tie-off Pins based on user options
- Provides user-configurable Ethernet MAC physical interfaces
 - Supports MII, GMII, RGMII v1.3, RGMII v2.0, SGMII, and 1000BASE-X PCS/PMA interfaces
 - Instantiates clock buffers, DCMs, RocketIOs, and logic as required for the selected physical interfaces
- Provides a simple FIFO-loopback example design, which is connected to the MAC client interfaces
- Provides a simple demonstration test bench based on the selected configuration
- Includes an example of a low-level driver for DCR accesses
- Generates VHDL or Verilog

LogiCORE Facts	
Supported Family	Virtex™-4 FX ¹
Performance	10 Mbps, 100 Mbps, 1 Gbps
Example Design Resources	
Slices	422-1354 ²
LUTs	464-1706 ²
FFs	519-1530 ²
BRAMs	4-8 ²
DCM	0-2 ²
BUFG	2-8 ²
Wrapper Highlights	
Optimized Clocking Logic	HDL Example Design
Hardware Verified	Demonstration Test Bench
Provided with Wrapper	
Documentation	Product Specification Getting Started Guide User Guide ³
Design File Formats	HDL Example Design, Demonstration Test Bench, Scripts
Constraints File	User Constraints File (.ucf)
Example Designs	Example FIFO connected to client I/F Demonstration Test Environment
Design Tool Requirements	
Supported HDL	VHDL and/or Verilog
Synthesis	XST 8.1i
Xilinx Tools	ISE™ 8.1i
Simulation tools	Mentor ModelSim® 6.1b Cadence IUS ⁴

1. Virtex-4 FX solutions require the latest silicon stepping and are pending hardware validation.
2. The precise number depends on user configuration; see [Device Utilization](#).
3. The Virtex-4 Embedded Tri-Mode Ethernet MAC User Guide is available from www.xilinx.com/bvdocs/userguides/ug074.pdf.
4. Scripts provided for Mentor ModelSim and Cadence IUS only.

Introduction

The Xilinx LogiCORE™ XAUI core is a high-performance, low pin count 10 Gbps interface intended to allow physical separation between data-link layer and physical layer devices in a 10-Gigabit Ethernet system.

The XAUI core implements a single-speed full-duplex 10 Gbps Ethernet eXtended Attachment Unit Interface (XAUI) solution for the Xilinx Virtex™-II Pro and Virtex-4 families of FPGAs.

The Virtex™-II Pro and Virtex-4 FPGA families, in combination with the XAUI core, enable the design of XAUI-based interconnects whether chip-to-chip, over backplanes, or connected to 10-Gigabit optical modules.

Features

- Designed to 10-Gigabit Ethernet specification IEEE 802.3ae-2002
- Uses 4 RocketIO™ transceivers at 3.125 Gbps line rate to achieve 10 Gbps data rate
- Implements DTE XGXS, PHY XGXS, and 10GBASE-X PCS in a single netlist
- Uses Virtex-II Pro or Virtex-4 Digital Clock Management to implement optional XGMII interface clocking
- Uses Virtex-II Pro or Virtex-4 DDR I/O primitives for the optional XGMII interface
- Elastic buffering of inbound XGMII data (optional)
- Uses RocketIO transceivers for the XAUI interface
- 802.3ae-2002 Clause 45 MDIO interface (optional)
- 802.3ae-2002 Clause 48 State Machines (optional for Virtex-II Pro)
- Supports 10-Gigabit Fibre Channel (10-GFC) XAUI data rates and traffic
- Available under the [SignOnce IP Site License](#) program

LogiCORE Facts				
Core Specifics				
Supported Device Family	Virtex-II Pro -6, -7 (2VP4 or larger) Virtex-4 ¹ (4VFX60)			
Resources Used ²	Slices	LUTs	FFs	Block RAMs
	917	1327	700	0
Special Features	Delivered through the CORE Generator™			
Provided with Core				
Documentation	Product Specification Getting Started Guide User Guide			
Design File Formats	NGC netlist			
Constraints File	UCF			
Verification	VHDL test bench Verilog test fixture			
Example design	VHDL and Verilog			
Additional Items	UniSim-based simulation models			
Design Tool Requirements				
Xilinx Implementation Tools	ISE™ 8.1i			
Simulation	Mentor ModelSim® Cadence IUS			
Support				
Provided by Xilinx, Inc. a www.xilinx.com/support/				

- Virtex-4 FX solutions require the latest silicon stepping and are pending hardware validation.
- Figures quoted are approximate for Virtex-II Pro default configuration. See "Device Utilization" on page 13 for details on device utilization by configuration.

Memory Interfaces Reference Designs

Give your designs the Virtex-4 FPGA advantage.

by Adrian Cosoroaba
Marketing Manager
Xilinx, Inc.
adrian.cosoroaba@xilinx.com

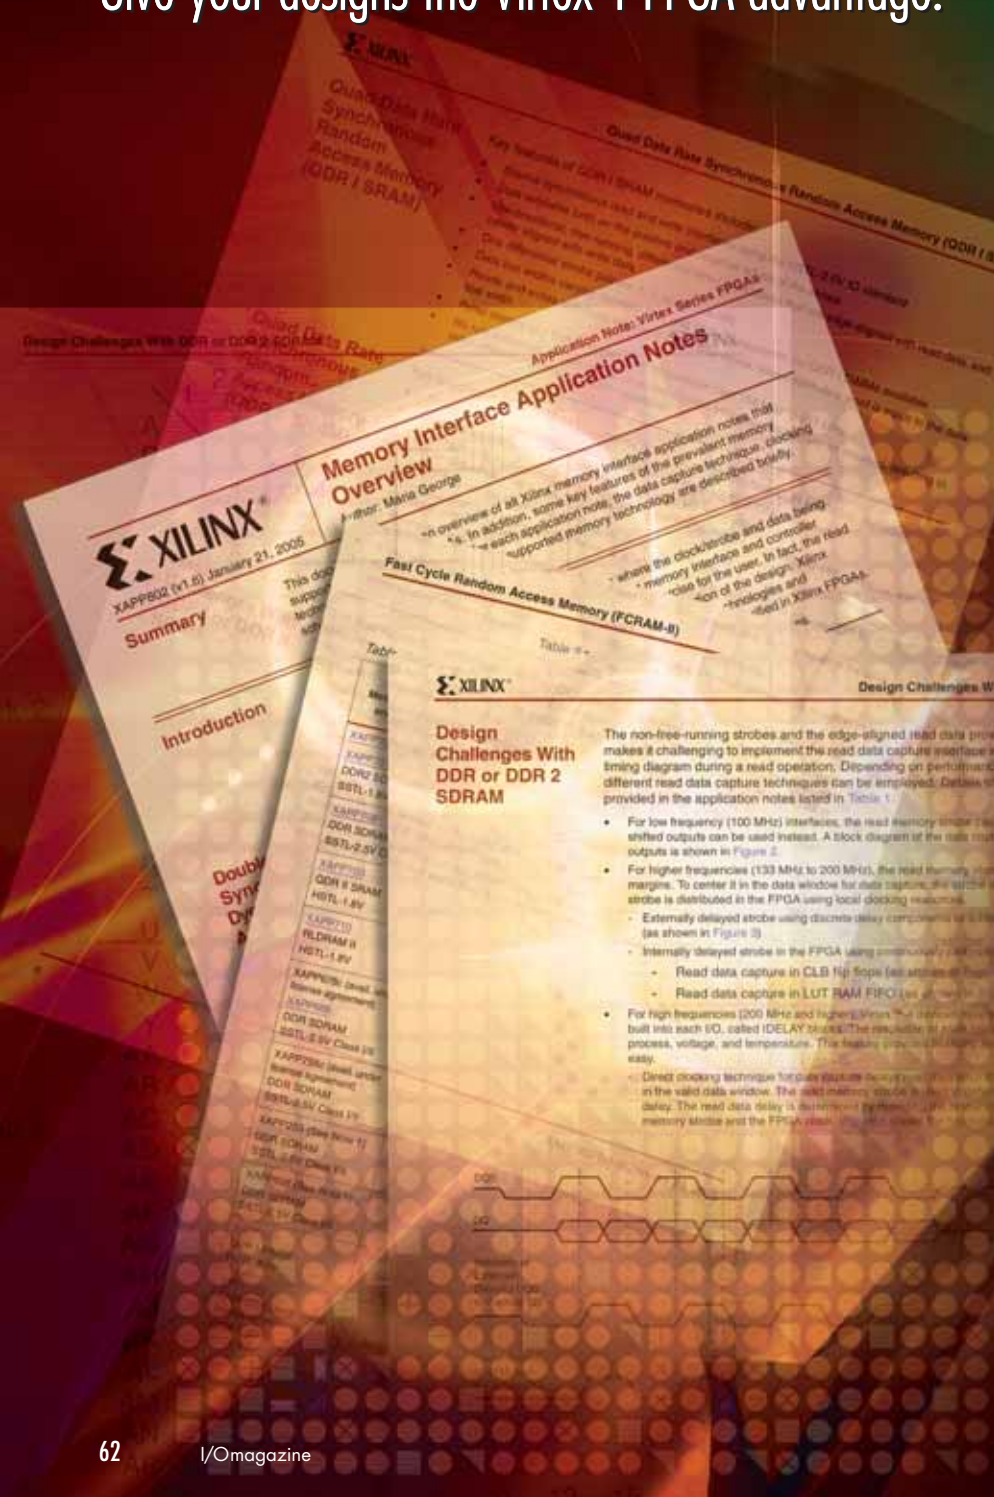
Memory interfaces are source-synchronous interfaces in which the clock/strobe and data being transmitted from a memory device are edge-aligned. Most memory interface and controller vendors leave the read data capture implementation as an exercise for the user. In fact, the read data capture implementation in FPGAs is the most challenging portion of the design. Xilinx provides multiple read data capture techniques for different memory technologies and performance requirements. All of these techniques are implemented and verified in Xilinx® FPGAs.

The following sections provide a brief overview of prevalent memory technologies.

Double Data Rate Synchronous Dynamic Random Access Memory (DDR SDRAM)

Key features of DDR SDRAM memories include:

- Source-synchronous read and write interfaces using the SSTL-2.5V Class I/II I/O standard
- Data available both on the positive and negative edges of the strobe
- Bi-directional, non-free-running, single-ended strobes that are output edge-aligned with read data and must be input center-aligned with write data
- One strobe per 4 or 8 data bits
- Data bus widths varying between 8, 16, and 32 for components and 32, 64, and 72 for DIMMs
- Supports reads and writes with burst lengths of two, four, or eight data words, where each data word is equal to the data bus width
- Read latency of 2, 2.5, or 3 clock cycles, with frequencies of 100 MHz, 133 MHz, 166 MHz, and 200 MHz
- Row activation required before accessing column addresses in an inactive row
- Refresh cycles required every 15.6 μ s
- Initialization sequence required after power on and before normal operation



Double Data Rate Synchronous Dynamic Random Access Memory (DDR 2 SDRAM)

Key features of DDR 2 SDRAM memories, the second-generation DDR SDRAMs, include:

- Source-synchronous read and write interfaces using the SSTL-1.8V Class I/II I/O standard
- Data available both on the positive and negative edges of the strobe
- Bi-directional, non-free-running, differential strobes that are output edge-aligned with read data and must be input center-aligned with write data
- One differential strobe pair per 4 or 8 data bits
- Data bus widths varying between 4, 8, and 16 for components and 64 and 72 for DIMMs
- Supports reads and writes with burst lengths of four or eight data words, where each data word is equal to the data bus width
- Read latency is a minimum of three clock cycles, with frequencies ranging from 200 MHz to 400 MHz
- Row activation required before accessing column addresses in an inactive row
- Refresh cycles required every 7.8 μ s
- Initialization sequence required after power on and before normal operation

Quad Data Rate Synchronous Random Access Memory (QDR II SRAM)

Key features of QDR II SRAM memories, the second-generation QDR I SRAMs, include:

- Source-synchronous read and write interfaces using the HSTL-1.8V I/O standard
- Data available both on the positive and negative edges of the strobe
- Uni-directional, free-running, differential data/echo clocks that are edge-aligned with read data and center-aligned with write data
- One differential strobe pair per 8, 9, 18, 36, or 72 data bits
- Data bus widths varying between 8, 9, 18, 36, and 72 for components (no QDR II SDRAM DIMMs available)

Memory Technology and I/O Standard	Supported FPGAs	Maximum Performance	Maximum Data Width	XAPP Number	XAPP Title	Data Capture Scheme
DDR 2 SDRAM SSTL-1.8V Class II	Virtex-4	333 MHz	8 bits (Components)	XAPP721 XAPP723	High Performance DDR 2 SDRAM Interface Data Capture Using ISERDES and OSERDES DDR2 Controller (267 MHz and Above) Using Virtex-4 Devices	Read data is captured in the delayed DQS domain and transferred to the FPGA clock domain within the ISERDES.
DDR 2 SDRAM SSTL-1.8V Class II	Virtex-4	267 MHz	16 bits (Components) 144-bit Registered DIMM	XAPP702 XAPP701	DDR 2 SDRAM Controller Using Virtex-4 Devices Memory Interfaces Data Capture Using Direct Clocking Technique	Read data delayed such that FPGA clock is centered in data window. Memory read strobe used to determine amount of read data delay.
DDR SDRAM SSTL-2.5V Class I/II	Virtex-4	200 MHz	16 bits (Components) 144-bit Registered DIMM	XAPP709	DDR SDRAM Controller Using Virtex-4 Devices	Read data delayed such that FPGA clock is centered in data window. Memory read strobe used to determine amount of read data delay.
QDR II SRAM HSTL-1.8V	Virtex-4	300 MHz	72 bits (Components)	XAPP703	QDR II SRAM Interface	Read data delayed such that FPGA clock is centered in data window. Memory read strobe used to determine amount of read data delay.
RLDRAM II HSTL-1.8V	Virtex-4	300 MHz	36 bits (Components)	XAPP710	Synthesizable CIO DDR RLDRAM II Controller for Virtex-4 FPGAs	Read data delayed such that FPGA clock is centered in data window. Memory read strobe used to determine amount of read data delay.

Table 1 – Virtex-4 memory interface application notes (XAPPs) currently available, with a brief description of the read data capture technique

XAPP Number Memory Technology and I/O Standard	Performance	Number of DCMs/DLLs	Number of BUFGs	Number of Interfaces with Listed DCMs and BUFGs	Device(s) Used for Hardware Verification	Requirements
XAPP721 XAPP723 DDR2 SDRAM SSTL-1.8V Class II	333 MHz	1 DCM 2 PMCDs	6	Multiple at Same Frequency	XC4VLX25 -11 FF668	All Banks Supported
XAPP702 XAPP701 DDR2 SDRAM SSTL-1.8V Class II	267 MHz	1	6	Multiple at Same Frequency	XC4VLX25 -11 FF668	All Banks Supported
XAPP709 DDR SDRAM SSTL-2.5V Class I/II	200 MHz	1	6	Multiple at Same Frequency	XC4VLX25 -11 FF668	All Banks Supported
XAPP703 QDR II SRAM HSTL-1.8V	300 MHz	1	3	Multiple at Same Frequency	XC4VLX25 -11 FF668	All Banks Supported
XAPP710 RLDRAM II HSTL-1.8V	300 MHz	1	5	Multiple at Same Frequency	XC4VLX25 -11 FF668	All Banks Supported

Table 2 – Resource utilization for all Virtex-4 memory interface application notes currently available

- Reads and writes with burst lengths of two or four data words, where each data word is equal to the data bus width
- Read latency is 1.5 clock cycles, with frequencies from 154 MHz to 300 MHz
- No row activation, refresh cycles, or initialization sequence after power on required, resulting in more efficient memory bandwidth utilization

Reduced Latency Dynamic Random Access Memory (RLDRAM II)


Key features of RLDRAM II memories include:

- Source-synchronous read and write interfaces using the HSTL-1.8V I/O standard
- Data available both on the positive and negative edges of the strobe
- Uni-directional, free-running, differential memory clocks that are edge-aligned with read data and center-aligned with write data
- One strobe per 9 or 18 data bits
- Data bus widths varying between 9, 18, and 36 for components and no DIMMs
- Supports reads and writes with burst lengths of two, four, or eight data words, where each data word is equal to the data bus width
- Read latency of five or six clock cycles, with frequencies of 200 MHz, 300 MHz, and 400 MHz
- Data-valid signal provided by memory device
- No row activation required; row and column can be addressed together
- Refresh cycles required every 3.9 μ s
- Initialization sequence required after power on and before normal operation

Conclusion

For application notes on various memory technologies and performance requirements, visit www.xilinx.com/memory. The summaries in Table 1 and Table 2 can help you determine which application note is relevant for a particular design. ●●

WHAT'S NEW



To complement our flagship publication *Xcell Journal*, we've recently launched three new technology magazines:

- ***Embedded Magazine***, focusing on the use of embedded processors in Xilinx® programmable logic devices.
- ***DSP Magazine***, focusing on the high-performance capabilities of our FPGA-based reconfigurable DSPs.
- ***I/O Magazine***, focusing on the wide range of serial and parallel connectivity options available in Xilinx devices.

In addition to these new magazines, we've created a family of Solution Guides, designed to provide useful information on a wide range of hot topics such as *Broadcast Engineering*, *Power Management*, and *Signal Integrity*.

Others are planned throughout the year.



See all the new publications on our website.

www.xilinx.com/xcell

Interfacing QDR II SRAM with Virtex-4 FPGAs

QDR II SRAM devices provide a suitable solution for memory requirements when partnered with Virtex-4 FPGAs.

by Veena Kondapalli
Applications Engineer Staff
Cypress Semiconductor Corp.
vee@cypress.com

The growing demand for higher performance communications, networking, and DSP necessitates higher performance memory devices to support such applications. Memory manufacturers like Cypress have developed specialized memory products such as quad data rate II (QDR II) SRAM devices to optimize memory bandwidth for a specific system architecture. In this article, I'll provide a general outline of a QDR II SRAM interface implemented in a Xilinx® Virtex™-4 XC4VP25 FF6688-11 device.

Figure 1 shows a block diagram of the QDR II SRAM design interface, with the physical interface to the actual memory device on the controller.

QDR II SRAM

QDR II can perform two data write and two data reads per clock cycle. It uses one port for writing data and one port for reading data. These unidirectional ports support simultaneous reads and writes and allow back-to-back transactions without the bus contention that may occur with a single bidirectional data bus.

Clocking Scheme

The FPGA generates all of the clock and control signals for reads and writes to memory. The memory clocks are typically generated using a double-data-rate (DDR) register. A digital clock manager (DCM) generates the clock and its inverted version. This has two advantages. First, the data, control, and clock signals all go through similar delay elements while exiting the FPGA. Second, the clock-duty cycle distortion is minimal when global clock nets are used for the clock and the 180° phase-shifted clock.

The reference design uses the phase-shifted outputs of the DCM to clock the interface on the transmit side. This configuration gives the best jitter and skew characteristics.

QDR II devices include the following features:

- Maximum frequency of operations - 250 MHz - tested up to 278 MHz
- Available in QDR II architecture with burst of 2 or 4
- Supports simultaneous reads/writes and back-to-back transactions without bus contention issues
- Supports multiple QDR II SRAM devices on the same bus to:
 - Increase the density of the memory resource

– Divide the speed of the interface by using multiple devices to achieve a given bandwidth

- Read: valid window worst-case 440 ps
- Write: valid window worst-case 460 ps
- Address and control signal timing analysis: command window worst-case 2360 ps

Conclusion

For more information about QDR II and Virtex-4 devices, see Xilinx application note XAPP703, "QDR II SRAM Interface for Virtex-4 Devices," at www.xilinx.com/bvdocs/appnotes/xapp703.pdf, as well as Cypress application note "Interfacing QDR-II SRAM with Virtex-4 Devices" at www.cypress.com.

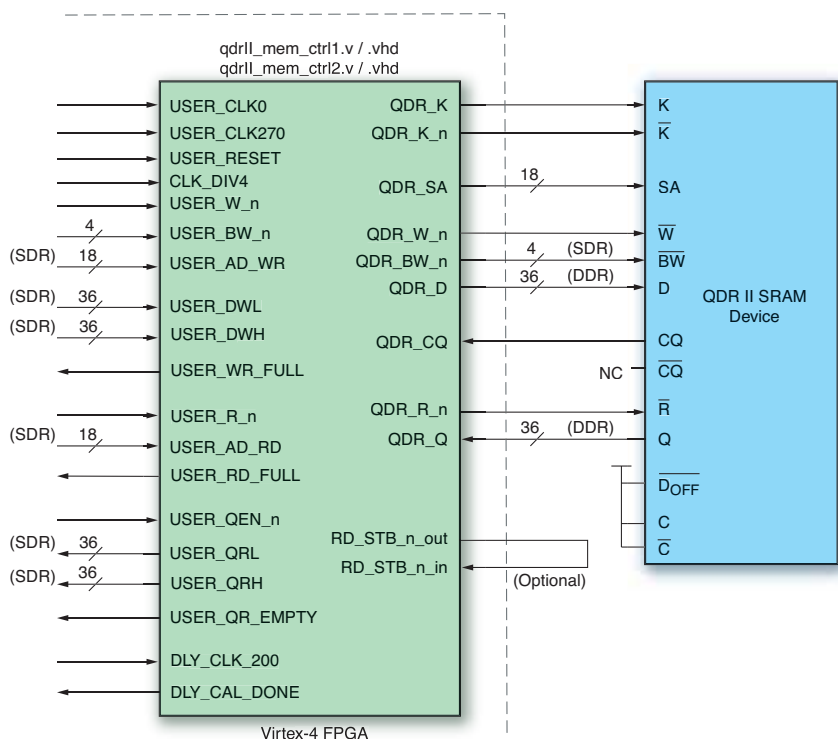


Figure 1 - Top-level architecture block diagram



Xilinx PCI Express Solution

PCI Express has emerged as the next generation technology replacing PCI. It provides higher performance and increased bandwidth while maintaining the flexibility and familiarity of PCI. Despite the advantages of PCI Express, design challenges associated with this new and complex protocol will directly affect time-to-market.

Xilinx provides a range of FPGA solutions to meet the needs of a variety of PCI Express applications. The breakthrough Virtex-4 and Virtex-II Pro FPGAs offer a fully integrated solution for applications with limited board real estate, utilizing built-in transceivers to implement the entire PCI Express interface in a single device. Alternatively, the low-cost Spartan-3 and Spartan-3E FPGA families can be used along with an external PHY device via the PHY Interface for PCI Express (PIPE).



The Xilinx PCI Express Advantage

The Xilinx PCI Express solution includes the PCI Express 1-lane, 4-lane and 8-lane endpoint IP cores for use with the Virtex-4 and Virtex-II Pro FPGA devices and the PCI Express PIPE 1-lane endpoint IP core for use with the Spartan-3 and Spartan-3E FPGA devices.

High Performance – The RocketIO™ Multi-Gigabit Transceivers (MGTS) on the Virtex-4 and Virtex-II Pro FPGAs give this core a 2.5 Gbps line speed in 1-lane configuration, 10 Gbps line speed in 4-lane configuration and 20 Gbps live speed in 8-lane configuration.

Low-Cost – The Xilinx PCI Express PIPE endpoint core is a high-bandwidth scalable and reliable IP building block for use with the Spartan-3 and Spartan-3E FPGAs. It is ideally suited for a broad range of high volume computing and communications applications requiring a low cost and 100% compliance with the PCI Express Base Specification v1.1a.

Flexibility – The inherently programmable nature of the FPGA allows you to continually modify your design as your performance and interoperability requirements evolve, reducing your risk in adopting the new PCI standard.



General Features

- High-performance, highly flexible, scalable, reliable, and general purpose I/O core
 - Compliant to the PCI Express Base Specification v1.1a
 - Compatible with current PCI software model
- Fully compliant with PCI Express transaction ordering rules
- Supports removal of corrupted packets for error detection and recovery
- Design verified by Xilinx proprietary test bench

PCI Express 1-Lane, 4-Lane and 8-Lane Endpoint Cores

- Incorporates Xilinx Smart-IP™ technology to guarantee critical timing
- Uses the RocketIO Multi-Gigabit Transceivers on the Virtex-4 and Virtex-II Pro FPGA devices to achieve high transceiver capability
 - 2.5 Gbps per lane line speed
 - Supports 1-lane, 4-lane and 8-lane operation (8-lane on Virtex-4 only)
 - Elastic buffers and clock compensation
 - Automatic clock data recovery
- 8b/10b encode and decode
- Offers standardized user interface
 - Easy-to-use packet-based protocol
 - Full-duplex communication
 - Back-to-back transactions enable greater link bandwidth utilization
 - Supports flow control of data and discontinuation of an in-process transaction in the transmit direction
 - Supports flow control of data in the receive direction
 - Transaction traffic class selection enabled
 - Support for automatically handling of error forwarded packets
 - Automatically decodes and removes error forwarding packet indicator from received data
 - Forward compatible with future link widths

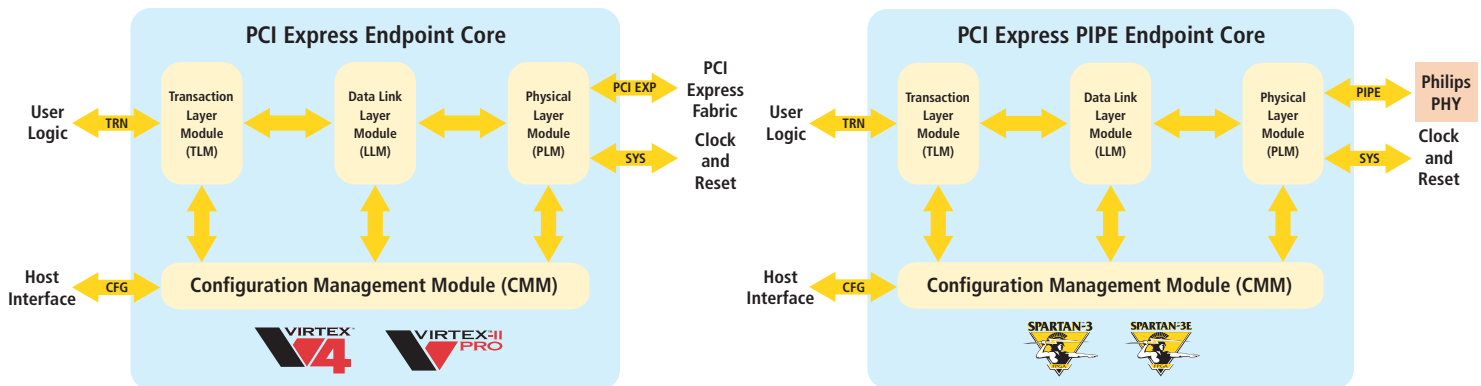
- Supports a maximum transaction payload of up to 4096 bytes
- Bandwidth scalability with frequency and/or interconnect width

PCI Express PIPE 1-Lane Endpoint Core

- Six individually programmable/configurable BARs and expansion ROM BAR
- Supports MSI and INTX emulation
- 32-Bit internal datapath
- Compatible with PCI/PCI-Express power management functions
 - Active state power management (ASPM)
 - Programmed power management (PPM)
- Used in conjunction with Philips PX1011A PCI Express standalone PHY to achieve high transceiver capability
 - 2.5 Gbps line speed
 - Elastic buffers and clock compensation
 - Automatic clock and data recovery
 - 8b/10b encode and decode
- Offers Xilinx-standardized easy-to-use LocalLink interface
 - Packet-based full-duplex communication
 - Back-to-back transactions enable greater link bandwidth utilization
 - Enables flow control of data and discontinuation of an in-process transaction in the transmit direction
 - Enables flow control of data in the receive direction
 - Automatically decodes and removes error forwarding packet indicator from received data
- Supports a maximum transaction payload of up to 512 bytes

Get Your PCI Express Solution Today

To learn more about the Xilinx PCI Express solution or to download the core, visit www.xilinx.com/pciexpress.



Corporate Headquarters

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124
Tel: (408) 559-7778
Fax: (408) 559-7114
Web: www.xilinx.com

European Headquarters

Xilinx
Citywest Business Campus
Saggart,
Co. Dublin
Ireland
Tel: +353-1-464-0311
Fax: +353-1-464-0324
Web: www.xilinx.com

Japan

Xilinx, K.K.
Shinjuku Square Tower 18F
6-22-1 Nishi-Shinjuku
Shinjuku-ku, Tokyo
163-1118, Japan
Tel: 81-3-5321-7711
Fax: 81-3-5321-7765
Web: www.xilinx.co.jp

Asia Pacific

Xilinx Asia Pacific Pte. Ltd.
No. 3 Changi Business Park Vista, #04-01
Singapore 486051
Tel: (65) 6544-8999
Fax: (65) 6789-8886
RCB no: 20-0312557-M
Web: www.xilinx.com

Distributed By:

FORTUNE 2005
100 BEST COMPANIES TO WORK FOR

XILINX
The Programmable Logic Company™

© 2005 Xilinx, Inc. All rights reserved. XILINX, the Xilinx logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

Spartan-3 Generation IP: Optimized for the World's Lowest-Cost FPGAs

The most comprehensive, optimized set of IP Cores:

Audio, Video, Broadcasting	JPEG Codecs, MPEG-2, MPEG-4, AES, H.264, Color Space Converters
Automotive	CAN Bus Controller 2.0A/B, 8051 Compatible Microcontroller, LIN Controller, MicroBlaze™, Video Compression Encoder, PCI
Communications & Networking	PCI, PCI Express, 10/100 MAC, GigE MAC, Mappers/Demappers and Framers/Deframers (STM0/OC1 – STM4/OC12), SDRAM Controller, DDR, DUC/DDC, HDLC Single Channel
Consumer	PCI, PCI Express, MPEG-2, MPEG-4, RSDS, USB 2.0, I2C, MicroBlaze
Industrial, Scientific & Medical	10/100 Ethernet, 1GB Ethernet, CAN Bus Controller 2.0A/B, Filters, Correlators, PicoBlaze™, MicroBlaze
Storage Area Networking	Serial Communication Controller, ATA and Serial ATA I/O Host Controller

Lowest Cost Connectivity for Chips, Boards and Backplanes

Xilinx offers IP cores for implementing your lowest cost system interconnectivity standards such as PCI Express, System Packet Interface (SPI-4.2/SPI-3)

PCI bus interfaces

- Best value for PCI 32/33 with effective cost below 75 cents*
- Programmable and Flexible PCI Express solution. The PCIe PIPE Endpoint LogiCORE™ combined with a discrete PCIe PHY offers a complete sub-\$12 PCIe Endpoint solution.

Gigabit Ethernet, 10/100 Ethernet, 10/100 Ethernet Lite

- Low cost Programmable and flexible 1 Gig (and under) Ethernet solution for less than \$10.00*
- A 10/100 Ethernet MAC core with OPB or PLB interface for embedded MicroBlaze and PowerPC™ solutions, a standalone Tri-Mode Ethernet MAC core.

CAN: Low cost automotive bus interface

- Effective cost of only \$1.27*

* Note: Pricing is based on typical implementation in the slowest speed grade, cheapest package with end of 2006 high volume pricing

Lowest Cost and Maximum Performance DSP Solution

Today, FPGAs and DSP processors often work together to meet the signal processing challenges in various high-performance digital communication systems, video/imaging, multimedia, and Aerospace and Defense systems. FPGAs complement DSPs in system logic consolidation, bus interfacing/bridging and signal processing acceleration. Xilinx and our partners offer a wide range of flexible DSP IP to help you get to market faster.

Error correction blocks

- Turbo Product Code Encoder, Turbo Product Code Decoder, Viterbi Decoder, Reed-Solomon Encoder, Reed-Solomon Decoder, Turbo Convolutional Code Encoder, Interleaver/De-interleaver

Modulation Demodulation

- Direct Digital Synthesizer, J.83 Universal Modulator Annex B, J.83 Universal Modulator Annex A/C, Digital Up Converter, Digital Down Converter

Transforms

- 2-D Discrete Cosine Transform, 1-D Discrete Cosine Transform, Fast Fourier Transform, 32 point Complex FFT

Filters

- Distributed Arithmetic FIR Filter, MAC filters, Cascaded integrator Comb (CIC)

Math Functions

- Floating Point Operator, Direct Digital Synthesizer, CORDIC

Lowest Cost Embedded Processing Solution

Xilinx offers complete range of processing IP solutions – ranging from the PicoBlaze 8-bit microcontroller to the high-end MicroBlaze 32-bit processor. This range of processing solutions lets you create high-performance, low-cost embedded systems for a wide range of applications in Spartan™-3 FPGAs.

To support processor-centric designs, Xilinx also offers a complete range of peripheral IP cores such as GPIO, Timer/Counter, UART 16450/16550, EMAC 10/100 and IIC. These allow you to customize your processor-based systems.

Low Cost Memory Controller Reference Designs

Xilinx provides free reference designs to help you interface to most popular DDR SDRAM memory from Micron, Samsung, and other companies. Xilinx provides a tool called the Memory Interface Generator (MIG) that allows users to quickly generate a HDL description of the kind of memory controller required for their application. Application notes now available on Memory Corner at: www.xilinx.com/products/design_resources/mem_corner/index.htm, describe the controller implementation in the silicon fabric.

Optimized For the World's Lowest Cost FPGA Family

With over 100 million units shipped, Spartan is the world's most popular low-cost FPGA architecture. With every generation of the Spartan architecture, Xilinx has delivered more logic and I/O at a lower price. The reason why Spartan-3 IP is so effective at reducing cost is the availability of embedded features such as:

Spartan-3 Feature For Lower Costs	Benefits
Shift Register Logic Functionality (SRL16)	Efficient pipelining and FIFO implementation. Reduces area used by multi-channel DSP functions.
Embedded Multipliers	Optimization of DSP IP cores such as FIR filters, Up and Down Converters
Distributed RAM	Efficient implementation of simple state machines and microcontrollers
18KB Block RAM	Ideal for memory intensive designs

IP Cores for Spartan-3 Speed Your System Design

Easy to Use IP Tools

Most IP is available in the ISE™ tools and accessible through the CORE Generator™ tool. The CORE Generator tool delivers a library of parameterizable and fixed netlist LogiCORE IP cores with the corresponding data sheets all designed and supported by Xilinx. For the latest updates visit Xilinx IP locator at www.xilinx.com/ipcenter today.

Simple Licensing Process

Xilinx and IP providers from around the world have combined efforts to form the Common License Consortium. The outcome is the simplification of the FPGA IP licensing process. Together, each company has agreed to license their IP cores to FPGA customers under a common set of terms known as the SignOnce IP License.



Evaluate before you buy

Before licensing an IP core, designers can download and evaluate them, free of charge, to ensure the cores meet their functionality requirements. To evaluate your IP, visit Xilinx IP locator at www.xilinx.com/ipcenter today.

Support

Xilinx provides world-class support for all Xilinx products, including IP cores. Visit www.xilinx.com/support for Documentation, Software Updates, Answers Database, and information on how to contact Xilinx Technical Support.

Quick Search with IP Locator

The Xilinx IP locator is the most comprehensive resource for intellectual property (IP) cores and development boards available from Xilinx and our third party partners. The advanced search feature allows quick and easy search based on functions (e.g., Bus Interface), sub function (e.g., PCI, Ethernet MAC), Xilinx devices and vendors. Visit www.xilinx.com/ipcenter for the most comprehensive set of IP available from Xilinx and partners.



Corporate Headquarters

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124
Tel: (408) 559-7778
Fax: (408) 559-7114
Web: www.xilinx.com

European Headquarters

Xilinx
Citywest Business Campus
Saggart,
Co. Dublin
Ireland
Tel: +353-1-464-0311
Fax: +353-1-464-0324
Web: www.xilinx.com

Japan

Xilinx, K.K.
Shinjuku Square Tower 18F
6-22-1 Nishi-Shinjuku
Shinjuku-ku, Tokyo
163-1118, Japan
Tel: 81-3-5321-7711
Fax: 81-3-5321-7765
Web: www.xilinx.co.jp

Asia Pacific

Xilinx Asia Pacific Pte. Ltd.
No. 3 Changi Business Park Vista, #04-01
Singapore 486051
Tel: (65) 6544-8999
Fax: (65) 6789-8886
RCB no: 20-0312557-M
Web: www.xilinx.com

Distributed By:

FORTUNE
100 BEST COMPANIES TO WORK FOR

2005



The Programmable Logic Company™

©2005 Xilinx, Inc. All rights reserved. XILINX, the Xilinx Logo, and other designated brands included herein are trademarks of Xilinx, Inc. PowerPC is a trademark of IBM, Inc. All other trademarks are the property of their respective owners.

Priced to go.



The industry's first 100K gate FPGA for under \$2.00*

Spartan-3E Platform FPGAs offer an amazing feature set for just \$2.00! You get 100K gates, embedded multipliers for high-performance/low-cost DSP, plenty of RAM, digital clock managers, and all the I/O support you need. All this in production now with a density range up to 1.6 million gates.

Perfect for digital consumer apps and much more!

With the Spartan-3E series, we've reduced the previous unit cost benchmark by over 30%. Optimized for gate-centric designs, and offering the lowest cost per logic cell in the industry, Spartan-3E FPGAs make it easy to replace your ASIC with a more flexible, faster-to-market solution. Compare the value for yourself . . . and get going on your latest design!



Easiest to use
Software

MAKE IT YOUR ASIC

Over **10** million
Spartan-3 devices
shipped!



The Programmable Logic Company™

For more information visit
www.xilinx.com/spartan3e

* Pricing for 500K units, second half of 2006



Pb-free devices
available now



Signal Integrity for High-Speed Memory and Processor I/O

SI20000-6-ILT (v1.0)

Course Specification

Course Description

Learn how signal integrity techniques are applicable to high-speed interfaces between Xilinx FPGAs and semiconductor memories. This course teaches you about high-speed bus and clock design, including transmission line termination, loading, and jitter. You will work with IBIS models and complete simulations using CAD packages. Other topics include managing PCB effects and on-chip termination. This course balances lecture modules and practical hands-on labs.

Level – Intermediate

Course Duration – 2 days

Price – \$1000 USD or 10 training credits

Course Part Number – SI20000-6-ILT

Who Should Attend? – Digital designers, board layout designers, or scientists, engineers, and technologists seeking to implement Xilinx solutions. Also end users of Xilinx products who want to understand how to implement high-speed interfaces without incurring the signal integrity problems related to timing, crosstalk, and overshoot or undershoot infractions.

Prerequisites

- Xilinx FPGA design experience preferred (equivalent of *Fundamentals of FPGA Design* course)

Software Tools

- Mentor Graphics HyperLynx®
- Cadence SPECCTRAQuest®

After completing this comprehensive training, you will have the necessary skills to:

- Identify when signal integrity is important and relevant
- Interpret an IBIS model and correct common errors
- Apply appropriate transmission line termination
- Understand the effect loading has on signal propagation
- Mitigate the impact of jitter
- Manage a memory data bus
- Understand the impact of selecting a PCB stackup
- Differentiate between on-chip termination and discrete termination

Course Outline

Day 1

- Introduction
- Transmission Lines
- **Mentor or Cadence Lab 1**
- IBIS Models
- **Mentor or Cadence Lab 2**
- **Mentor or Cadence Lab 3**
- High-Speed Clock Design
- **Mentor or Cadence Lab 4**
- SRAM Requirements
- **Mentor or Cadence Lab 5**

Day 2

- Physical PCB Structure
- On-Chip Termination
- SDRAM Design
- **Mentor Lab 6**
- Managing an Entire Design

Lab Descriptions

Note: Labs feature the Mentor Graphics or Cadence flow. For private training, please specify your flow to your registrar or sales contact. For public classes, flow will be determined by the instructor based upon class feedback.

- **Mentor Lab 1:** Opening the appropriate Mentor simulator
- **Mentor Lab 2:** Hands-on signal integrity observation of reflection and propagation effects
- **Mentor Lab 3:** Using an IBIS simulator to study basic transmission line effects
- **Mentor Lab 4:** Using saved simulation information to perform power calculation. Also, additional clock simulations
- **Mentor Lab 5:** Observing the effects of coupling on transmission lines
- **Mentor Lab 6:** Demonstrating how an SDRAM module can be handled with an EBD model
- **Cadence Lab 1:** Opening the appropriate Cadence simulator
- **Cadence Lab 2:** Analysis of a simple clock net
- **Cadence Lab 3:** Signal integrity effects caused by multidrop clock networks
- **Cadence Lab 4:** Crosstalk analysis
- **Cadence Lab 5:** Address and data analysis

Register Today

Xilinx delivers public and private courses in locations throughout the world. Please contact Xilinx Education Services for more information, to view schedules, or to register online.

Visit www.xilinx.com/education, and click on the region where you want to attend a course.

North America, send your inquiries to registrar@xilinx.com, or contact the registrar at 877-XLX-CLAS (877-959-2527). To register online, search by **Keyword** "High-Speed" in the Training Catalog at <https://xilinx.onsaba.net/xilinx>.

Europe, send your inquiries to eurotraining@xilinx.com, call +44-870-7350-548, or send a fax to +44-870-7350-620.

Asia Pacific, contact our training providers at: www.xilinx.com/support/training/asia-learning-catalog.htm, send your inquiries to education_ap@xilinx.com, or call: +852-2424-5200.

Japan, see the Japanese training schedule at: www.xilinx.co.jp/support/training/japan-learning-catalog.htm, send your inquiries to education_kk@xilinx.com, or call: +81-3-5321-7772.

You must have your tuition payment information available when you enroll. We accept credit cards (Visa, MasterCard, or American Express) as well as purchase orders and training credits.

Course Description

By learning PCI Express core protocol fundamentals, designers will gain a working knowledge of how PCI Express can be used in their systems. This course focuses on PCI Express protocol subjects that designers using the Xilinx PCI Express should understand in order to complete their designs faster and easier. Customers will also be introduced to each Xilinx PCI Express core product and will gain intimate knowledge of how the PCI Express core operates.

Level – Intermediate

Course Duration – 2 days

Price – \$1000 USD or 10 Training Credits

Course Part Number – PClxxxx-BETA

Who Should Attend? – Engineers who seek training in developing the necessary skills for designing PCI Express systems using Xilinx PCI Express cores

Prerequisites

- Basic PCI and/or PCI-X protocol knowledge
- Basic knowledge of Verilog or VHDL
- Basic experience with commonly used simulation tools like ModelSim
- Basic knowledge of Xilinx ISE™ software

Software Tools

- Xilinx ISE 8.1i
- ModelSim 6.0c PE

After completing this comprehensive training, you will have the necessary skills to:

- Effectively use the Xilinx PCI Express cores in your own design environments
- Select the appropriate PCI solution for a specific application
- Understand how PCI Express specification requirements apply to using Xilinx PCI Express cores.

Course Outline

Day 1

- PCI Express Overview
- Layers and Channels
- TLP Packet Fields and Packet Routing
- Local Link Interface
- **Lab 1:** Using the PCI Express Core Local Link Interface
- PCI Express Configuration Space
- **Lab 2:** Exploring the PCI Express Core Configuration Space

Day 2

- TLP Request and Completion Packets
- **Lab 3:** Designing with the PCI Express Core
- Physical Layer Electrical Subblock
- Physical Layer Logical Subblock
- Xilinx PCI Express Solutions
- **Lab 4:** Generating and Implementing a Xilinx PCI Express Core

Lab Descriptions

- **Lab 1:** Using the PCI Express Core Local Link Interface – This lab introduces the PCI Express core design that will also be used in Labs 2 and 3. It allows the user to become familiar with the cores user application interface (Local Link) and to modify the design to change the packets being sent
- **Lab 2:** Exploring the PCI Express Core Configuration Space – This lab reinforces lessons learned in the Configuration Space section by having users decode configuration packets to understand the requirements in configuring the core. In addition, users will be able to implement the user configuration space by modifying the Programmable I/O design receiver and transmit state machines
- **Lab 3:** Designing with the PCI Express Core – This lab takes an in-depth look at designing with the core. The user will become familiar with packet ordering, credits available, and allocating completion space for inbound completions
- **Lab 4:** Generating and Implementing a Xilinx PCI Express Core – This lab illustrates using the CORE Generator™ to generate a core. The core is then implemented and users can verify the implementation by studying the various reports created by Xilinx tools

Register Today

Xilinx delivers public and private courses in locations throughout the world. Please contact Xilinx Education Services for more information, to view schedules, or to register online.

Visit www.xilinx.com/education, and click on the region where you want to attend a course.

North America, send your inquiries to registrar@xilinx.com, or contact the registrar at 877-XLX-CLAS (877-959-2527). To register online, search by **Keyword "PCI"** in the Training Catalog at <https://xilinx.onsaba.net/xilinx>.

Europe, send your inquiries to eurotraining@xilinx.com, call +44-870-7350-548, or send a fax to +44-870-7350-620.

Asia Pacific, contact our training providers at: www.xilinx.com/support/training/asia-learning-catalog.htm, send your inquiries to education_ap@xilinx.com, or call: +852-2424-5200.

Japan, see the Japanese training schedule at: www.xilinx.co.jp/support/training/japan-learning-catalog.htm, send your inquiries to education_kk@xilinx.com, or call: +81-3-5321-7772.

You must have your tuition payment information available when you enroll. We accept credit cards (Visa, MasterCard, or American Express) as well as purchase orders and training credits.

Course Description

Learn how to employ RocketIO[™] MGT serial transceivers in your Virtex-II Pro[™] design! Understand and utilize the features of the RocketIO transceiver blocks, such as CRC, 8b/10b encoding, channel bonding, clock correction, and comma detection. Additional highlighted topics include debugging techniques, use of the Architecture Wizard, synthesis and implementation considerations, and standards compliance. This course balances lecture modules and practical hands-on labs.

Level – Intermediate

Course Duration – 2 days

Price – \$1000 USD or 10 Training Credits

Course Part Number – RIO22000-8-ILT

Who Should Attend? – FPGA designers and logic designers

Prerequisites

- Verilog or VHDL experience (or the *Introduction to Verilog* or the *Introduction to VHDL* course)
- Synthesis and simulation experience
- FPGA design experience or the *Fundamentals of FPGA Design* course
- Knowledge of high-speed serial I/O protocols and standards (SONET, Gigabit Ethernet, InfiniBand) is a plus.

Software Tools

- ISE 8.1
- ModelSim PE 6.0

After completing this comprehensive training, you will have the necessary skills to:

- Effectively use all of the advanced RocketIO features, such as CRC, channel bonding, clock correction, comma detection, 8b/10b encoding / decoding, programmable termination, and pre-emphasis
- Utilize the ports and attributes of RocketIO transceivers that control the RocketIO features
- Use the Architecture Wizard to instantiate RocketIO primitives in your design
- Achieve compatibility with high-speed I/O standards by using RocketIO transceivers

Course Outline

Day 1

- Introduction
- Clocking and Resets
- 8b/10b Encoder and Decoder Details
- **Lab 1:** 8b/10b Disparity and Bypass Lab
- Commas and Deserializer Alignment Details
- **Lab 2:** Commas and K-Characters Lab
- Cyclical Redundancy Check Details
- **Lab 3:** Cyclical Redundancy Check Lab
- Clock Correction Details
- **Lab 4:** Clock Correction Lab

Day 2

- Channel Bonding Details
- **Lab 5:** Channel Bonding Lab
- Architecture Wizard Overview
- Implementing a RocketIO Design
- **Lab 6:** Synthesis and Implementation Lab
- IP Overview: Aurora Reference Design
- **Lab 7:** Aurora Protocol Engine Lab
- Common Serial I/O Standards Compliance
- Physical Media Attachment Overview

Lab Descriptions

- **Lab 1:** 8B/10B Disparity/Bypass Lab – Learn how to use 8b/10b encoder / decoder and manipulate running disparity. Learn how to bypass the 8b/10b encoder / decoder.
- **Lab 2:** Comma and K-character Lab – Learn how to use programmable comma detection to align a serial data stream.
- **Lab 3:** CRC Lab – Modify a design to use the CRC feature for both the user mode and the Fiber Channel mode of CRC.
- **Lab 4:** Clock Correction Lab – Learn to use the clock correction logic to compensate for frequency differences on the TX and RX side of a link.
- **Lab 5:** Channel Bonding Lab – Modify a design to use two transceivers bonded together to form one virtual channel.
- **Lab 6:** Synthesis and Implementation Lab – Learn to use the Architecture Wizard to instantiate RocketIO primitives, synthesize a design, and implement the design.
- **Lab 7:** Aurora Protocol Engine Lab – Learn how to use the Aurora reference design to send and receive data.

Register Today

Xilinx delivers public and private courses in locations throughout the world. Please contact Xilinx Education Services for more information, to view schedules, or to register online.

Visit www.xilinx.com/education, and click on the region where you want to attend a course.

North America, send your inquiries to registrar@xilinx.com, or contact the registrar at 877-XLX-CLAS (877-959-2527). To register online, search by **Keyword** "High-Speed" in the Training Catalog at <https://xilinx.onsaba.net/xilinx>.


Europe, send your inquiries to eurotraining@xilinx.com, call +44-870-7350-548 or send a fax to +44-870-7350-620.

Asia Pacific, contact our training providers at: www.xilinx.com/support/training/asia-learning-catalog.htm, send your inquiries to education_ap@xilinx.com, or call: +852-2424-5200.

Japan, see the Japanese training schedule at: www.xilinx.co.jp/support/training/japan-learning-catalog.htm, send your inquiries to education_kk@xilinx.com, or call: +81-3-5321-7772.

You must have your tuition payment information available when you enroll. We accept credit cards (Visa, MasterCard, or American Express) as well as purchase orders and training credits.

HIGH-PERFORMANCE SIGNAL INTEGRITY SIMULATION SOFTWARE



NEXXIM®

State-of-the-art circuit simulation for high-capacity, multi-gigabit design

HFSS™/Q3D EXTRACTOR®

3D parasitic extraction for the design of on-chip passives and board/package interconnects

SIWAVE™

Full-board and full-package signal- and power-integrity analysis

DESIGNER SI™

System-level signal-integrity (SI) solution for IC/package/board co-design and verification



ANSOFT.COM

High VELOCITY LEARNING



Nu Horizons Electronics Corp. is proud to present our newest education and training program - **XpressTrack** - which offers engineers the opportunity to participate in technical seminars conducted around the country by experts focused on the latest technologies from Xilinx. This program provides higher velocity learning to help minimize start-up time to quickly begin your design process utilizing the latest development tools, software and products from both Nu Horizons and Xilinx.

Don't see a seminar in a city near you? Visit our website and let us know where you reside and what you are interested in learning about and we'll develop a curriculum just for you.

For a complete list of course offerings, or to register for a seminar near you, please visit:

www.nuhorizons.com/xpresstrack



Topics Covered

- What's New in ISE 8.1i
- Faster Performance - ISE Fmax and PlanAhead
- ISE 8.1i - Focus on New Features
- Design Optimization in ISE 8.1i using XST
- PlanAhead Design Analysis Tool
- Xilinx Platform Studio Update
- Providing Power for FPGAs
- Spartan Family Update
- Virtex Family Update
- CPLD Family Update

February 2006

- February 16 - Chicago, IL
Irvine, CA
San Jose, CA
Toronto, ON
- February 21 - Dallas, TX
Ottawa, ON
- February 22 - Los Angeles, CA
Montreal, ON
- February 23 - Boston, MA
Minnesota, MN
Shenzhen, China

March 2006


- March 1 - Orlando, FL
- March 7 - Beijing, China
Dallas, TX
- March 8 - Atlanta, GA
- March 13 - Portland, OR
- March 14 - Shanghai, China
- March 16 - Irvine, CA
Philadelphia, PA
- March 20 - Baltimore, MD
- March 21 - San Jose, CA
Toronto, ON
- March 22 - Austin, TX
Los Angeles, CA
- March 23 - Boston, MA
Chicago, IL
Minnesota, MN

Order Your **FREE DVD**
www.xilinx.com/sidvd



Dr. Howard Johnson
The world's foremost authority
on signal integrity

"HIGH SIGNAL INTEGRITY DEMANDS A LOW NOISE CHIP".



Design Example: 1.5 volt LVCMOS 4mA, I/O, 100 aggressors shown.



Dr. Howard Johnson, author of *High-Speed Digital Design*, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit www.sigcon.com to register.



Best Signal Integrity: 7x Less SSO Noise

Virtex-4 FPGAs deliver the industry's best signal integrity, allowing you to pre-empt board issues at the chip level, for high-speed designs such as memory interfaces. Featuring a unique SparseChevron™ pin out pattern, the Virtex-4 family provides the highest ratio of VCCO/GND pin pairs to user I/O pins available in any FPGA. By strategically positioning one hard power pin and one hard ground pin adjacent to every user I/O on the device, we've reduced signal path inductance and SSO noise to levels far below what you can attain with a virtual ground or soft ground architecture.

THE INDUSTRY'S HIGHEST SIGNAL INTEGRITY, PROVEN BY INDUSTRY EXPERTS

Incorporating continuous power and ground planes, plus integrated bypass capacitors, we're eliminating power-supply noise at its source. In addition, we provide on-chip termination resistors to control signal ringing. The lab tests speak for themselves. As measured by signal integrity expert Dr. Howard Johnson, no competing FPGA comes close to achieving the low-noise benchmarks of Virtex-4 devices.

Visit www.xilinx.com/virtex4/si today, and choose the right high-performance FPGA before things get noisy.



Highest customer satisfaction:
• Virtex & Spartan FPGAs
• ISE design tools
• Support



The Programmable Logic CompanySM

www.xilinx.com/virtex4/si



View The
TechOnLine
Seminar Today

BREAKTHROUGH PERFORMANCE AT THE LOWEST COST

©2006 Xilinx, Inc. All rights reserved. XILINX, the Xilinx logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

PN 0010917