

Questions and Answers:

The Real 64/66 PCI from Xilinx

What does Xilinx mean by The Real 64/66 PCI?

PCI 64-bit, 66MHz is an extremely difficult interface to implement, not only in FPGA technology, but also in ASIC technology. Several companies – FPGA vendors and standard-chip vendors – have announced that they will have products in the future, or that they can meet the stringent 66MHz PCI timing, but all have yet to prove that they have a real solution. Xilinx is the first company to actually ship a complete, general-purpose solution for 64-bit, 66MHz PCI.

The Xilinx solution is "real" because it provides:

- **Real compliance** by providing a fully PCI v2.2 compliant core with guaranteed 66MHz PCI timing. No other company provides a soft core with guaranteed timing.
- **Real flexibility** by providing a standard, off-the-shelf FPGA that meets all timing requirements of a fully compliant, 64-bit, 66MHz PCI interface.
- **Real performance** by providing a PCI core that supports zero wait-state burst and scalable on-chip dual-port FIFOs. As a result, the designer can achieve up to the theoretical maximum throughput of 528 MBytes/second.
- **Real availability because** only Xilinx ships the 64-bit, 66MHz PCI solution today. Moreover, Xilinx has had an early access program with more than a dozen customers since September last year, all designing for 66MHz PCI.

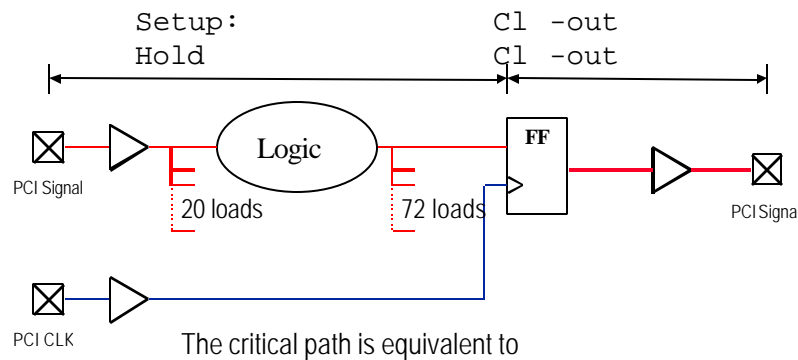
Why is 66 MHz PCI so difficult to do?

There are several reasons why PCI is such a challenge:

- PCI is an extremely complex protocol for which almost an infinite number of transaction combinations can occur. The PCI Special Interest Group (PCI-SIG) provides a checklist which all PCI vendors must comply with. However, this checklist only covers a small subset of all the possible combinations of transactions that can occur, hence, it is not good enough to only adhere to the checklist. As a result, many of the standard PCI interface chips are released with a long errata list of known issues and limitations. Xilinx has an internal testbench that runs more than six million test vectors to verify its PCI core designs.
- Meeting the stringent 66MHz PCI timing is a challenge in any technology. The 66MHz PCI v2.2 specification calls for a 3 ns (nanosecond) max setup-time, 0 ns hold-time, 6 ns max clock-to-out, and 2 ns min clock-to-out. For example, to achieve zero wait-state burst with a fully compliant behavior, a number of PCI control signals must be decoded live during the 3 ns setup-time. In a 64-bit 66MHz interface, those signals must go through an input buffer and fan out to more than 20 loads, go through one level of logic to 74 registers, in 3 ns or less. The following illustrations provides details.

Attaining 200MHz Plus:

The Critical Path for Zero Wait-state PCI64/66



Xilinx supports this path in its Virtex FPGA family, and it is implemented and guaranteed in Xilinx Real 64/66 PCI core design. Furthermore, Xilinx guarantees the 0 ns hold-time and both the 6 ns max clock-to-out and 2 ns min clock-to-out.

How can Xilinx guarantee the timing in its soft Real 64/66 PCI design?

Xilinx spends significant effort to fully verify its PCI interface cores through extensive simulation, and through characterization of both the core and the FPGA device. The verification includes both functional behavior and timing. When all timing parameters are met, placement constraints and routing constraints are applied to the core. The constraints ensure that the PCI core behaves correctly and meets timing every time a user implements it with a unique back-end design. The user's back-end design doesn't affect the PCI timing as long as the user follows the design rules described in the product documentation. To maintain the correct implementation constraints, the PCI core is parameterized through an easy-to-use graphical user interface available on the Xilinx web site. This technology is known as Smart-IP Technology and is only available from Xilinx. Because other FPGA vendors do not provide Smart-IP technology, their users must account for tuning the implementation and verifying the timing in order to claim full PCI compliance.

How does an FPGA with a PCI core compare to a standard PCI-bridge chip?

First, there are no general-purpose standard 64-bit, 66MHz PCI-bridge chips available yet. Second, a standard PCI-bridge chip is designed for a specific application, with a specific feature set, bridging from PCI to a given local bus. In an FPGA, a user can integrate a fully compliant PCI interface with customized FIFOs, DMA, and other functions in up to one

million gates of user logic. The result is higher integration, higher performance, and lower system cost.

How does a standard FPGA with a soft PCI core compare to a specialized FPGA with a PCI interface embedded in silicon?

FPGAs have become a very popular device for integrating glue logic – and lately entire logic systems – because it is a broad-based standard component. FPGAs come in a variety of sizes and packages, they are flexible to use, offer fast time-to-market, require short lead-times and simplify inventory logistics. Creating a specialized FPGA with a PCI interface violates the fundamental benefits of an FPGA – flexibility to change – that made the device so popular in the first place. Furthermore, since PCI is such a difficult application to get right the first time, suppliers of specialized PCI FPGAs will face the same quality issues and limitations as standard PCI chip vendors have for several years.

The benefits of using standard, rather than specialized FPGAs, for PCI applications are numerous:

- With a standard FPGA solution, users can choose from a range of device sizes and packages. For example, Xilinx offers 5,000 – one million system gates for user applications for the company's various PCI 32-bit and 64-bit solutions. Only one specialized PCI FPGA has announced to date.
- With a standard FPGA solution users can implement PCI features, DMA configurations and FIFO sizes that are required for their specific applications, and users do not have to pay for unused features. As with standard PCI-bridge chips, a specialized FPGA comes with a fixed set of features.
- With a standard FPGA and soft PCI core solution, users can adapt to future PCI revisions and add requirements that may be revealed when new PCI chip sets or PCs are introduced.

Isn't a specialized FPGA with embedded PCI a lower cost solution?

A PCI interface embedded in silicon may occupy less silicon area. However, this saving is offset by increased manufacturing costs resulting from lower manufacturing volume and more complex testing. Moreover, a standard FPGA with its regular architecture will be migrated more aggressively to new process technologies, which over time will result in both lower cost and higher performance.

So Xilinx PCI products are delivered over the Internet. How does that work?

A customer purchases the The Real 64/66 PCI products as regular software product. Xilinx will ship the "physical" product, not including any design files, with a unique serial number. The user then registers on Xilinx web site to get access to the PCI configuration tool, all the design files, reference designs, and online documentation. The unique serial number will be used to authenticate a new user. All customers will have instant access to new The Real 64/66 PCI releases, new reference designs and other updates.

What tools do I need to design with Xilinx PCI products?

The customer receives all design files required for a VHDL and Verilog design flow, including wrappers for instantiation, simulation models and a basic test bench. Xilinx tests and documents the Synopsys FPGA Compiler and FPGA Express tools for design entry, and Model Technology and Verilog XL for simulation.

Additionally, Xilinx has announced support for the Synplify tools from Synplicity, which has tested and documented the flow.

How many PCI designs has Xilinx supported to date?

Xilinx customers have completed more than 1,000 designs with the PCI32 cores, of which many have been tested at the PCI-SIG's Plug-fest for full compliance. Xilinx currently has more than a dozen customers using the new The Real 64/66 PCI products.

How can I order the The Real 64/66 PCI solution from Xilinx, and what is the price?

- Order part number DO-DI-PCI64
- Price \$14,995
- Available now from Xilinx local sales representatives. See the Xilinx web site for details (www.xilinx.com)

Xilinx PCI Customer Classes

What prerequisites are necessary to attend a Xilinx PCI course?

Attendees should have Xilinx design experience and some knowledge of PCI. The students need to have the following:

- Working experience with digital design
- Basic knowledge of Verilog or VHDL
- Some experience of Xilinx Foundation and FPGA Express including writing UCF files
- Some knowledge of PCI (Xilinx will provide PCI basics online before hand)

What topics will be covered?

The focus will be on design usage of Xilinx PCI products with hands-on participation during the lab session. In addition, the following topics will be covered:

- Basic PCI concepts
- Xilinx PCI solution
- Designing with Xilinx PCI - Configuration, Integration, and Verification
- Design debug

What will students have learned on completion of the course?

The goal of Xilinx Customer Education is to help customers successfully complete a PCI design. After completing this course, customers will be able to:

- Describe the basics of the PCI specification
- Select the appropriate PCI solution for a specific application
- Register for Xilinx PCI lounges on Xilinx web site, where all design files are configured and downloaded.
- Configure and download a The Real 64/66 PCI design
- Integrate the The Real 64/66 PCI design with a back-end design
- Verify and debug a PCI design with Xilinx tools

How can I order the PCI training course, and what is the price?

- Order part number TC-PCI-2DY
- The 2 day course is \$1,000
- For more information, see the web site <http://support.xilinx.com> – under Education, click Courses

###