

A Low-Cost Programmable PCI Express Solution

Our solution comprises a Philips PX1011A-EL1 device and a Spartan-3/E FPGA containing the PCI Express Endpoint core.

by David "Andrew" Brierley-Green
Senior Principal Engineer
Philips Semiconductors
david.brierley-green@philips.com

Ho Wai Wong-Lam
Program Manager
Philips Semiconductors
ho.wai.wong-lam@philips.com

PCI Express is a new interconnect standard that provides a serial replacement for the PCI, AGP, and PCI-X buses, which are commonly used in computer and embedded systems. The market has embraced the adoption of PCI Express in computers, and a wide variety of PCI Express plug-in cards and ExpressCards are now available.

Many plug-in cards and embedded systems that use FPGAs to implement PCI and PCI-X are expected to migrate to PCI Express in the coming years. Higher throughput, lower printed circuit board complexity, and the unification of various interconnectivity standards into one PCI Express standard are all factors that entice more developers to adopt PCI Express technology.

A programmable solution offers flexibility, short time to market, and low upfront costs, and is ideal for emerging and low- to mid-volume applications. Philips Semiconductors offers the PX1011A-EL1 x1 PCI Express PHY to form a low-cost programmable PCI Express solution with

Xilinx® Spartan-™3/E FPGAs containing the Xilinx PCI Express Endpoint core. This combination achieves a price point that is a fraction of previously available programmable solutions for PCI Express. It enables designers of high-volume applications to take advantage of a programmable and compliant PCI Express solution.

PCI Express and Its Physical Layer

Like all modern networking and connectivity standards, PCI Express uses a layered protocol model. Data is transferred at 2.5 Gbps over a single PCI Express lane; this configuration is referred to as an x1 link. PCI Express is scalable in that you can bundle multiple lanes together. For example,

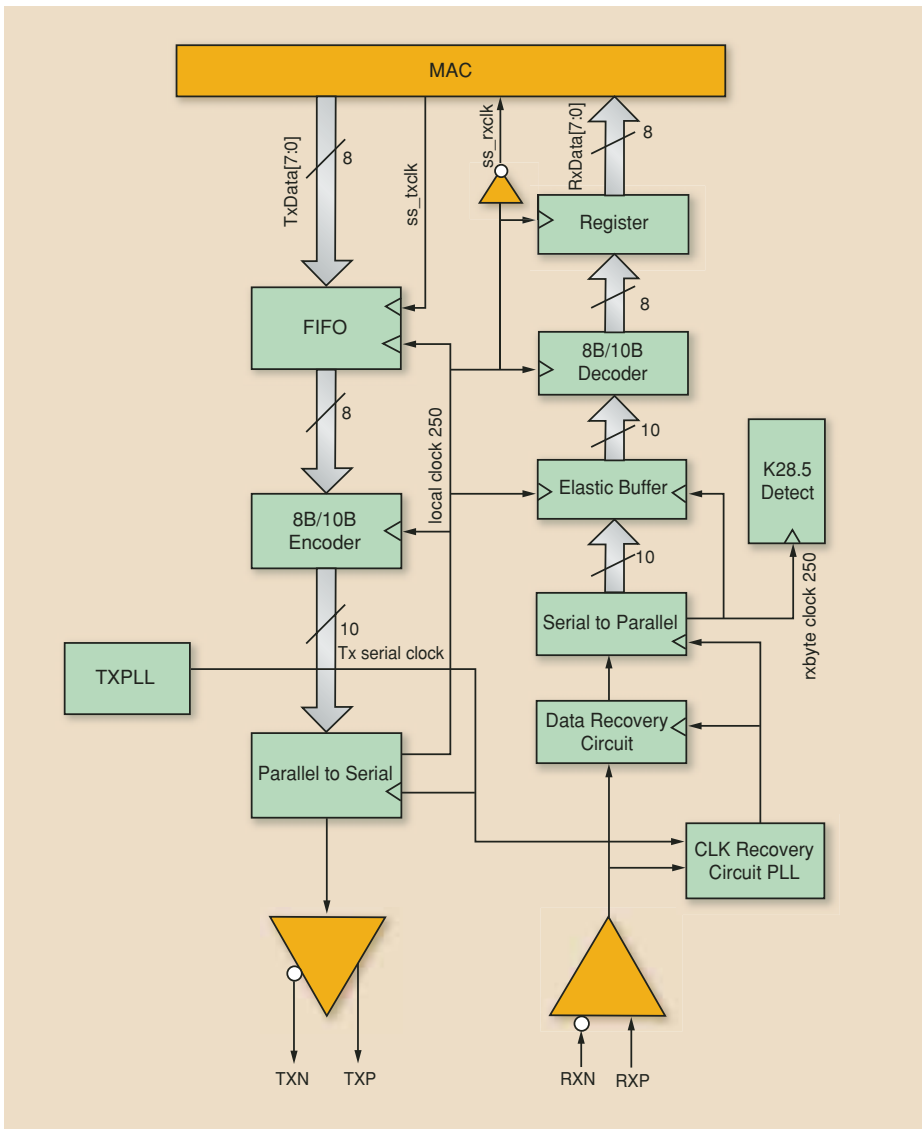


Figure 1 – Block diagram of PX1011A-EL1

an x4 link consists of four x1 lanes. A maximum of 32 lanes is allowed in a link, resulting in an aggregate bandwidth of 80 Gbps in each direction.

The physical layer's main function is to get packets of data across the link with a 10-12 or lower bit error rate. Factors such as signal voltage levels, equalization, receiver performance, transmit bit order, coding, and link initialization and training are dealt with in the physical layer. The main purpose of link initialization and training is to configure the link width (number of lanes), lane ordering, and correct polarity reversal within a differential conductor pair. The logical physical sub-layer handles link training.

Intel defines a specification for the PIPE

– Physical Interface for PCI Express – between the media access layer (MAC) and PHY. Part of the logical sub-layer of the physical layer resides in the MAC portion of the PIPE interface. The physical coding sub-layer (PCS) of the physical layer and the electrical physical layer reside in the PHY portion of the PIPE interface. Therefore, PCI Express PHY devices that are based on the PIPE interface do not contain all of the functions described in the physical layer of the PCI Express specification.

The data link layer provides packets to and receives packets from the physical layer. The data link layer makes the unreliable link appear perfect to higher layers by retransmitting packets with errors. The transaction layer translates PCI transaction

requests from the software layer into packets called transaction layer packets (TLPs) and relies on the underlying data link and physical layers to deliver the TLPs to the corresponding transaction layer at the destination node. An important function of the transaction layer is flow control.

PX1011A-EL1

A block diagram of the PX1011A-EL1 is shown in Figure 1. The interface between the PX1011A-EL1 and the higher layer logic is a variant of the PIPE interface.

On the transmit side, the PX1011A-EL1 receives 8-bit words of data from the MAC, along with a control bit that indicates whether the 8-bit word is data or a control character. The data is transferred at a rate of one word per cycle of a 250 MHz clock. The data is first buffered in a FIFO, which allows for phase differences between the PIPE clock and the internal 250 MHz transmit clock generated by the transmit phase-locked loop (TXPLL). The data is then 8B/10B encoded, with the 10-bit data serialized and differentially transmitted onto the transmission line.

On the receive side, the PX1011A-EL1 receives serial differential data from the transmission line. A clock is recovered from the data; this clock is used to sample the serial data in the center of the data eye. The retimed serial data is then passed through a serial-parallel converter.

The next step is to find the 10-bit symbol boundaries with a special 10-bit character called a “comma,” or K28.5 character. The K28.5 character cannot occur by chance because of the concatenation of any other legal 10-bit characters. Once symbol synchronization has been achieved, the realigned 10-bit characters are passed through an elastic buffer that compensates for the frequency difference (if any) between the recovered and locally generated transmit clocks. This frequency compensation happens when you add or remove special characters (called “skip” characters) that are provided in the PCI Express protocol for this purpose. The retimed 10-bit characters are then decoded and the resulting 8-bit data words are registered and output from the PX1011A-EL1 to the MAC device.

The PX1011A-EL1 (shown in Figure 2) uses an 81-pin LFBGA package with 0.8 mm pitch. The package is approximately 9 x 9 mm and has a height of 1.05 mm. The power dissipation is less than 300 mW total during active operation. A leaded version is available, with a lead-free version to follow.

PXPIPE Interface

The PIPE specification was defined for an on-chip interface, and it does not address very well the difficulties that arise with a chip-to-chip connection. The PIPE interface defines a single 250 MHz clock, called PCLK, to which both the transmit and receive sides of the interface are synchronized. PCLK is an output of the PHY and an input to the MAC. This means that the MAC responds to the rising edge of PCLK by shifting out a word for transmission. Taking into account of the time of flight for signals across the channel, which includes PCB traces and possibly a connector, we see that the two instances of one-way propagation delay must be included in the timing budget. The total of the round-trip propagation delay plus the clock-to-out delay of the MAC plus the setup time of the PHY must be less than one PCLK period (4 ns).

The classic solution to this problem is to use source-synchronous clocking. With source-synchronous clocking, the clocks propagate in the same direction as the data rather than in the opposite direction. Because both clock and data incur the same propagation delay, they cancel out the timing budget. You could, in theory, have a propagation delay longer than the clock period.

By comparison, the Intel PIPE specification solves this timing budget problem by introducing the option of a 16-bit data interface, thus doubling the timing budget

from 4 ns to 8 ns. This solution comes at a cost of higher pin count (thus a larger package and higher cost) and introduces an extra layer of logic to convert from 16 to 8 bits, thus increasing latency.

Philips Semiconductors has defined an alternative version of the PIPE interface named PXPIPE. Rather than a single byte-rate clock, we provide two clocks: a transmit byte clock and a receive byte clock. The PXPIPE interface also specifies the use of SSTL2 signaling. Being an on-chip interface, the original PIPE specification does not specify any signaling levels.




Figure 2 – PX1011A-EL1 device

Applications

The small size and low power consumption make the PX1011A-EL1 ideally suited for ExpressCard applications. ExpressCards are the new generation of PCMCIA PC cards used to add functionality to portable computers. ExpressCards use either a USB2.0 or PCI Express x1 host interface. They have restrictive component height requirements and stringent power consumption limits.

The PX1011A-EL1/Spartan-3/E PCI Express solution has many potential applications. This list shows some of the potential areas, and by no means represents an exhaustive enumeration. The Philips/Xilinx solution opens up a new horizon of potential applications that can make use of a low-cost and programmable PCI Express solution. The scope of applications is limited only by your imagination.

- Storage/RAID
- PC controlled and embedded test equipment
- Digital TV tuners
- Home printers
- Disk recorders
- Professional graphics boards
- Professional cameras
- Image capture and processing
- Digital media creation
- Digital music mixers
- Network security systems
- Voice over IP
- DSL modems
- Medical imaging (ultrasound, X-rays)

“PX Surfboard” Demo Design

A demo design called the “PX Surfboard” showcases the Philips/Xilinx PCI Express solution. The demo can be used for interoperability and compliance testing. In addition, it shows how you can design a system using the PX1011A-EL1/Spartan-3 solution. Schematics and Gerber files are also available.

Conclusion

The Philips PX1011A-EL1 device, used in conjunction with a Xilinx Spartan-3/E FPGA containing the PCI Express Endpoint LogiCORE™ solution, provides a low-cost, low-power, fully standards-compliant PCI Express solution with proven interoperability with other vendors' PCI Express solutions. For more information, please visit the following websites:

- www.semiconductors.philips.com/markets/connectivity/wired/pciexpress/solution/index.html
- www.xilinx.com/xlnx/xebiz/designResources/ip_product_details.jsp?key=DO-DI-PCIE-PIPE