

PCI Express Advances Machine Vision

PCI Express is the peripheral bus now being adopted by next-generation PCs, servers, and industrial computers. It provides a scaleable, high-bandwidth, point-to-point pathway between peripheral cards and the computing core while retaining application software compatibility with previous generations. For machine-vision systems, the architecture and higher bandwidth of PCI Express yield major increases in achievable frame rate and image size as well as simplifying the implementation of multi-channel capability.

The PCI Special Interest Group (PCISIG) created PCI Express as a solution to the increasing mismatch between the PC's peripheral card bus and the I/O demands of high-performance graphics, communications, and storage. The traditional PCI bus had run into hard limits to its clock speed caused by skew across the parallel bus. The highest performance PCI bus achievable runs at 66 MHz with a 64-bit data pathway, for a data rate of 528 Mbytes/second.

The PCISIG developed an intermediate solution called PCI-X that aimed to provide higher bandwidth while still utilizing a parallel bus structure by running as fast as 266 MHz. The design of PCI-X peripherals was both complex and costly, however, and developers typically limited their systems to 133 MHz for a data rate of 1 Gbyte/second. Even at this lower speed, skew and loading considerations limited the fan-out to only a few peripheral cards. PCI-X thus did not prove satisfactory; an entirely different solution was needed.

PCI Express (PCIe) eliminates the skew and fan-out limitations of high-speed parallel buses by adopting a serial bus structure, instead. A PCIe connection is made point-to-point through a switch matrix, providing a direct link between the two communicating entities (see Figure 1). This direct link ensures that a data transfer is able to utilize the full bus bandwidth; there is no sharing of the bus with other connections during the transfer.

The physical arrangement of a PCIe link is through a set of serial lanes, with groupings of 1, 2, 4, 8, 12, 16, or 32 lanes allowed. Each lane operates at 2.5 Gbits/second using 8b/10b coding to provide a self-clocking 250 Mbyte/second data rate in each direction. Thus, a four-lane link provides 1 Gbyte/second and a 16-lane link provides 4 Gbytes/second per direction, significantly higher than the PCI bus can achieve. Further, the PCIe bus has considerable expansion potential. The recently-introduced PCIe 2.0 specification allows the links to run at 5.0 Gbits/second for double the data rate, and higher link rates are in development.

Software Compatibility Retained

Even though PCIe uses a serial connection at the physical layer, however, it retains backward compatibility with PCI peripherals at the driver, operating system, and application levels. The hardware for a PCIe link handles the transition from parallel, memory-mapped data transfers of PCI to the switched serial transactions of PCIe. Even the interrupt and other out-of-band signals of PCI map into the serial PCIe, so that the software is unaware of the bus change.

PCIe hardware handles this translation in three stages. At the physical layer, hardware automatically converts parallel data transfers into serial blocks and stripes the data across available lanes to fully utilize the link bandwidth. Each serial block contains addressing and identification bits that allow the physical layer hardware at the receiving end to deserialize and reassemble the data in the correct order, eliminating any skew present between lanes. The physical layer also handles lane width negotiations during system initialization, matching its output to the channel available.

Link-layer hardware provides error detection and correction for data transfers, operating at the block level. If a data transfer error occurs, the link-layer hardware automatically resends the block without the intervention of the upper layers. From the standpoint of the transaction and higher layers, an error event simply looks like a longer than usual response time for acknowledging the transaction.

The transaction layer of PCIe provides all of the conversion between memory-mapped PCI data transfer, command, and out-of-band signaling and the address-driven serial PCIe signals. The command and data interfaces between the transaction layer and the rest of the computer system are identical to those used by the PCI transaction layer. Thus, the hardware drivers, operating system, and applications software of a computer system remain unaltered when interacting with a PCIe peripheral.

Advantages of PCIe

This move to PCI Express provides PC-based system developers with a number of advantages. One is an increased ability to balance cost and performance. Designers only need to use as many lanes in their design as their performance requirements dictate. Given that the bus width is only four wires per lane, connector costs drop, as well.

Another advantage of PCIe is that the bus's dedicated connection helps eliminate system bottlenecks. The bandwidth of the PCI bus had to be shared among all the peripherals and other system functions, so that the available bandwidth decreased with each additional peripheral. PCIe is a switched point-to-point connection and allows peer-to-peer communications. This means that a data transfer has a dedicated link available to it and does not need to share bandwidth during the transfer. In addition, the peer-to-peer capability allows multiple transactions to occur simultaneously if a non-blocking switch is used. Further, peripheral devices can stream data direct to system memory without CPU intervention while the CPU is performing other tasks.

PCIe also helps simplify software design when upgrading systems. Its compatibility with older PCI designs means that applications, drivers, and operating system software for PCI devices run unmodified when using comparable PCIe devices. Further, PCIe allows additional software capabilities, such as an ability to offer quality-of-service controls on data transfers, that simplify the addition of new functions if the operating system supports them. Peripherals with a native PCIe interface are also positioned to support the expanded memory space of newer 64-bit operating system software such as Windows XP and Vista. Most PCI system cards were designed with older versions of Windows that only



supported 32-bit addressing, so they restrict their DMA addressing to 4 Gbytes, of which half is reserved for system use. This leaves only a 2 Gbyte block that can be directly accessed, forcing the use of virtual memory (see Figure 2) to handle larger data sets such as the output of a line-scan camera. PCIe peripherals designed to work with 64-bit operating system are able to stream data into an address range that runs to thousands of terabytes.

PCIe Implementation Challenges

Implementing a PCIe system is not without its challenges, however. One of the problems that system developers face is that most commercially-available PCIe motherboards were designed for desktop computer use and have limited PCIe functionality. A typical desktop motherboard offers one or two 16-lane PCIe slots for high performance, such as graphics functions, and a few 1-lane slots.

This arrangement causes problems for developers needing more performance than the 256 Mbyte/second bandwidth that the 1-lane slots provide or wishing to use cards developed for 4- or 8-lane slots. While PCIe hardware will negotiate a link to use fewer lanes when there is a mismatch, the slots themselves will not accept cards designed for more lanes (see Figure 3). So in a typical commercial motherboard a 4-lane machine vision card could only be installed in the 16-lane graphics card slot.

Because the intended use for the 16-lane slot is graphics, however, the use of another card type in that slot has caused problems with some motherboards. The BIOS expects a graphics card and may not properly handle system resource allocation for another card type. As a result, the system may not achieve the performance levels that the card can support.

The solution to this dilemma is for developers to use motherboards that target the server market. While not typically available from consumer sources, they can be obtained from industrial sources. These server motherboards usually incorporate 4-lane or 8-lane PCIe slots allowing system developers to readily achieve the performance levels they require.

The Right Design Approach

The design of a peripheral card for a PCIe system can follow one of two approaches. Because of software compatibility, it is possible to take an existing PCI or PCI-X design and simply connect it to a PCIe slot through a bridge chip. This approach has the advantage that it offers a low-cost upgrade path to the new bus structure, but results in a design that is limited to the performance that the original bus interface provided. This includes being limited to the 32-bit addressing of the original PCI DMA controller. Experience has also shown that simply using a bridge chip to upgrade a design runs into system initialization issues. As with the slot problems of commercial motherboards, the BIOS may create challenges. During allocation of system resources, some BIOS software has difficulty reaching through a bridge to determine card resource requirements. Some bridge chip vendors have augmented

their designs to address these issues, but this leaves system developers with a need to check the compatibility of motherboard BIOS and card bridge chip to avoid problems.

The second design approach is to make PCIe the peripheral card's native interface. Taking this approach eliminates the need to check for compatibility as there is no bridge device to confuse the BIOS. A native PCIe design also allows the card to offer maximum performance. Without the restriction of the PCI interface, the card's hardware can fully utilize the bandwidth of the PCIe link.

Many vendors have followed the bridge approach to enter the PCIe market quickly at low cost. DALSA, however, has chosen to use a native PCIe design approach in its next generation Camera Link frame grabber and vision processor product lines. This choice provides DALSA customers with the maximum degree of motherboard compatibility attainable, freeing their system design from artificial limitations. It also ensures that developers have a single point of contact for system design questions; there is no need to coordinate with a BIOS provider to resolve issues.

DALSA's PCIe products are available with 1-lane and 4-lane bus interfaces. The 1-lane designs offer the greatest flexibility to users because they will work in any PCIe slot in any motherboard. The 4-lane designs offer higher performance and are suitable for use in conjunction with workstation and server-grade motherboards.

PCIe Applications in Machine Vision

The high speed and dedicated bandwidth available to DALSA's frame-grabber boards because of their native PCIe interface enables users to address a wider variety of machine vision applications than earlier PCI devices. One application that benefits is the use of high-speed line-scan imaging for web inspection. Under the 32-bit addressing limit of PCI, line-scan systems needed to artificially break the image into frames of a size that would fit into the 2-GigaByte addressable memory space. This framing activity interrupts the data flow and complicates the image processing used in the inspection process by introducing arbitrary image boundaries. The 64-bit compatibility designed into DALSA's native PCIe interface allows virtually unlimited frame sizes, permitting the web inspection to proceed on a more continuous basis.

Another application that benefits from PCIe's attributes is multi-angle inspection. Such systems use multiple cameras to examine an object from several directions simultaneously. This multiple view allows a system to inspect all surfaces of an object without requiring the object to be manipulated. A conventional PCI-based system that offers multiple image-capture channels, however, must share the system bandwidth among the channels. This sharing quickly becomes a bottleneck, lowering the inspection system's throughput. A PCIe-based system can employ multiple frame-grabber boards and provide each with a link to system memory. Because the bandwidth of a PCIe link is dedicated to the data transfer it carries, each frame grabber in the system can operate at its full speed without affecting the others. This, in turn, allows the inspection system to provide maximum throughput.

DALSA aims to exploit additional features of PCIe, such as its independent, bi-directional bandwidth. This bi-directional capability is poised to help simplify co-processing in machine vision systems. As Gigabit Ethernet cameras began making inroads into general purpose machine vision applications, the need to offload from the host CPU the task of converting the GigE Vision packets into usable images became increasingly evident. A Bayer color GigE camera sending data at 100Mbytes/sec., for example, requires the host computer to spend tremendous amounts of valuable CPU clock cycles simply converting and decoding images into usable formats.

The CPU can be off-loaded, however, by utilizing a general-purpose co-processing board and a network interface card (NIC) on the PCIe bus. The camera sends its data to the network, and the NIC forwards the data to the coprocessor across PCIe. The coprocessor performs both packet conversion and Bayer decoding in real-time, then streams the data to system memory using the transmit half of its PCIe link. This high-speed data transmission, processing, and storage can all occur without CPU intervention, freeing the CPU to perform more critical analysis. This approach also allows the creation and use of standardized PCIe coprocessor cards rather than the proprietary mezzanine cards being used in PCI systems, simplifying system design and lowering cost.

This kind of growth potential along with the current benefits of PCI Express will ensure its adoption in many coming generations of PCs. Its attributes applied to machine vision systems will, in turn, allow these systems to provide greater throughput and multi-channel capabilities unattainable utilizing the PCI bus. DALSA is embracing the technology, providing its customers with design options that ensure maximum performance and widest compatibility with present and future generations of industrial computers.

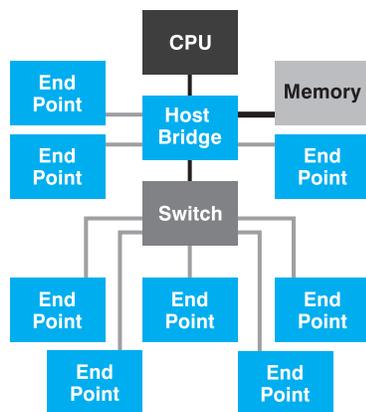


Figure 1: Rather than sharing a set of connections among many peripheral cards as in the traditional PCI bus, PCI Express uses switched serial links to provide a direct, full-bandwidth connection between any two nodes on the bus.

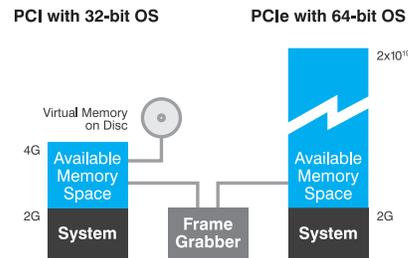


Figure 2: By supporting 64-bit addressing, PCIe can bypass the memory addressing limitations of older 32-bit PCI systems, providing direct access to a virtually unlimited address space with no need for virtual memory.

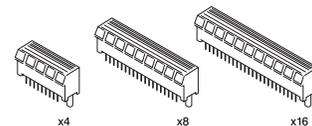


Figure 3: PCIe connections are downward compatible, so a 4-lane card can plug into an 8-lane slot and the PCIe protocol will negotiate lane usage to re-configure the slot for 4-lane operation.

Inder Kohli, Product Manager

About the Author

Inder Kohli is a Product Manager at DALSA, Montreal. He holds a Bachelor of Commerce degree with honors from the University of Delhi in India, a Certificate in Management Information Systems from McGill University, Montreal, as well as a Diploma in Electrical Engineering. Before joining DALSA in 1994, Mr. Kohli worked as an application support engineer, a software product manager, and an automation technologist.

Inder may be reached at: inder.kohli@dalsa.com