PCI-X 2.0 White Paper

Revision 1.0

April 5, 2002

Authors:

Sujith Arramreddy
Chief Technology Officer
ServerWorks

Dwight Riley
Principal Member of Technical Staff
Compaq

Table of Contents

INTRODUCTION

The PCI architecture has proven to be successful beyond even the most optimistic expectations. Today nearly every new computer platform comes outfitted with multiple PCI slots. In addition to the unprecedented number of PCI slots being shipped, there are also hundreds of PCI adapter cards that are available to satisfy virtually every conceivable application. This enormous momentum is difficult to ignore.

Today there is also a need for a new, higher-performance I/O interface to support emerging, ultrahigh-bandwidth technologies such as 10 Gigabit Ethernet, 10 Gigabit FibreChannel, 4X and 12X InfiniBand, and others. A standard that can meet these performance objectives, while maintaining backward compatibility to previous generations of PCI would undoubtedly provide the ideal solution.

To meet these objectives, the PCI-X 2.0 standard has been developed. PCI-X 2.0 has the performance to feed the most bandwidth-hungry applications while at the same time maintaining complete hardware and software backward compatibility to previous generations of PCI and PCI-X.

The PCI-X 2.0 standard introduces two new speed grades: PCI-X 266 and PCI-X 533. These speed grades offer bandwidths that are two times and four times that of PCI-X 133 -- ultimately providing bandwidths that are more than 32 times faster than the original version of PCI that was introduced eight years ago. It achieves the additional performance via time-proven DDR (Double Data Rate) and QDR (Quad Data Rate) techniques that transmit data at either 2-times or 4-times the base clock frequency.

Because PCI-X 2.0 preserves so many elements from previous generations of PCI it is the beneficiary of a tremendous amount of prior development work. The operating systems, connector, device drivers, form factor, protocols, BIOS, electrical signaling, BFM (bus functional model), and other original PCI elements, are all heavily leveraged in the PCI-X 2.0 specification. In fact many of these elements remain identical in PCI-X 2.0. These similarities make implementation easy because these elements have already been designed and engineers are already familiar with them. As a result, the time-to-market is short, and risk is dramatically reduced.

The market migration to PCI-X 2.0 will also be easy because there are so many previous-generation PCI adapter cards already on the market. There are already hundreds of PCI adapter cards that are available today that can be utilized by every PCI-X 266 and PCI-X 533 slot. In addition, new PCI-X 266 and PCI-X 533 adapter cards have ready homes in any of the millions of PCI and PCI-X slots in existing systems.

Because of these factors, PCI-X 2.0 provides the ideal next-generation, local I/O solution for high-bandwidth applications. It offers the performance needed for today's and tomorrow's applications in an easy-to-adopt, backward-compatible standard.

WHY PCI-X 2.0 MAKES SENSE

Provides The Required Bandwidth
PCI-X 2.0 has more than enough bandwidth to support even the most cutting-edge technologies including 10 Gigabit Ethernet, 10 Gigabit FibreChannel, and 4X and 12X InfiniBand. In fact PCI-X 2.0 can support multiple high-bandwidth ports simultaneously, and at full bandwidth. For example, a PCI-X 533 adapter card could easily support two ports of either 10 Gigabit Ethernet, or 10 Gigabit FibreChannel.

Full Hardware Backward Compatibility
A key reason for the success of each new generation of PCI has been that they have been fully backward compatible with the previous versions. PCI-X 2.0 continues the PCI tradition with complete backward compatibility. Even legacy 32-bit/33MHz PCI adapter cards that are keyed for 3.3V or universal signaling can plug into the highest performance PCI-X 2.0 slots and operate seamlessly. And the converse is also true. The newest, highest-performance PCI-X 2.0 adapter cards can readily operate in 3.3V PCI slots from past-generation computers.

Immediate Support for PCI-X 2.0 Slots and Adapters
Not only is forward and backward compatibility convenient, but it also creates another significant advantage to the user: every slot is equally usable. Adapter cards for literally hundreds of applications are already available that can plug into all the new PCI-X 2.0 slots.

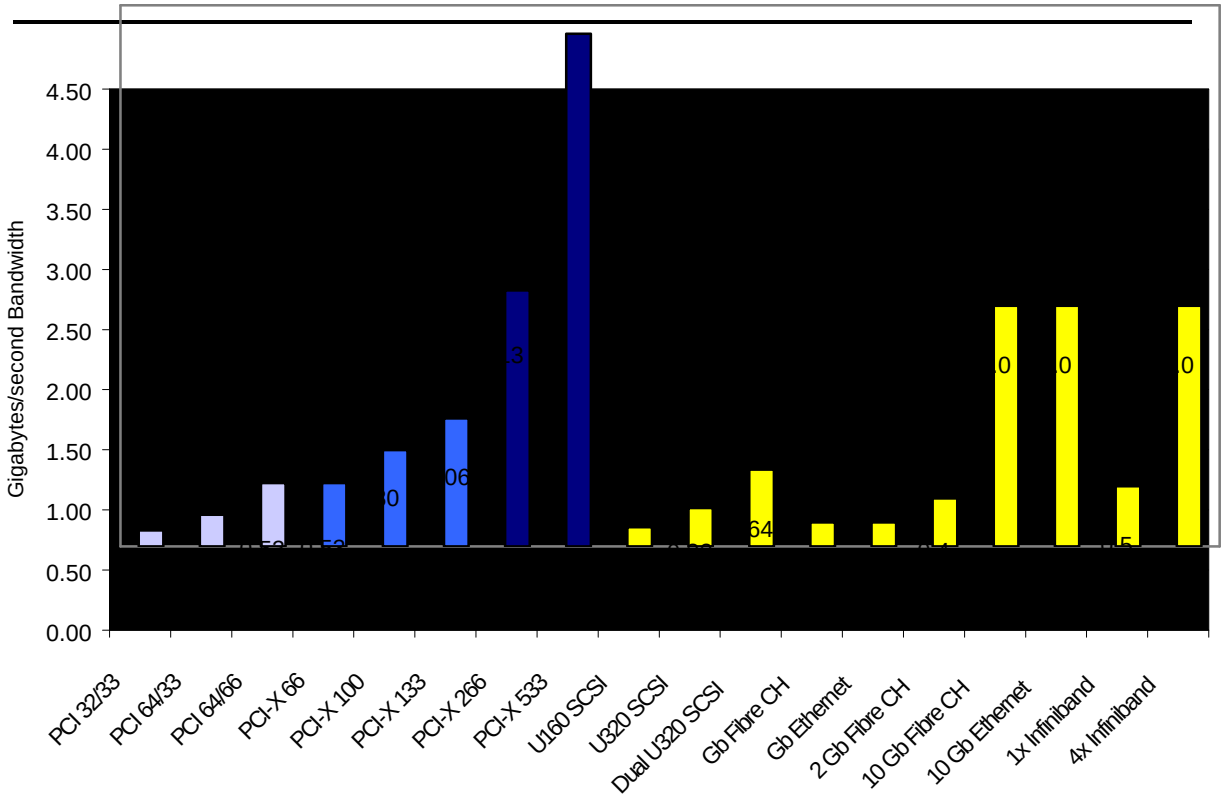

Bandwidth offered by different generations of PCI compared to the bandwidth demands of various target applications.

Immediate Support for PCI-X 2.0 Slots and Adapters (Continued)

From a system designer's perspective, this means there is no need to design a system with more than one type of slot – or to attempt to predict how many of each type of slot a user might possibly need. This can eliminate the need to design multiple systems supporting multiple types of users. From a user's perspective, there is no need to be concerned about which cards work in which slots, or how many of each type of slot are still free.

Full Software Backward Compatibility

To make the migration to PCI-X 2.0 easy, full software backward compatibility is also faithfully maintained. Operating systems that support current generations of PCI also support PCI-X 2.0: Window 3.0/95/98/2000/XP, Linux, OpenVMS, Novell, UNIX, OS/2, SCO, Nonstop Computing, etc. All the current addressing, registers, and protocols are accepted and interpreted correctly by PCI-X 2.0. Software compatibility is also true of the BIOS, and the device drivers. They don't require any change; by default they are compatible with all speeds of PCI-X 2.0 adapter cards.

Time to Market

The PCI-X 2.0 specification will be finalized and released in the spring of 2002. Yet companies that are involved in developing the specification are already designing PCI-X 2.0 systems and adapter cards. These systems and adapter cards will be ready for production in the first half of 2003. These cards and systems have the capability to come to market quickly because they leverage most of the features of the PCI-X 1.0 specification. In fact much of it was already designed and debugged when the PCI-X 1.0 generation was created. Consequently the time to market will be fast.

Speed Mismatch Compatibility

All speed generations of PCI are interoperable. For example, a 33MHz adapter card can plug into a 66MHz slot, a PCI-X 133 slot, or even a PCI-X 533 slot. If the card can't meet the performance of the slot, the slot frequency is throttled down to meet the maximum frequency of the card. In this example, since the card is capable of operating at 33MHz, the slot will change its frequency to 33MHz to match the speed of the card. The converse is also true; a higher performance card will also adjust its frequency to match the maximum performance of a slower slot, maintaining complete forward and backward compatibility.

These adjustments happen automatically and transparently, so there's no need for the user to reset jumpers, modify BIOS settings, upgrade software, or install new drivers. The user simply plugs in the card, and the PCI protocol takes care of the rest. This increases the utility of PCI-X 2.0 adapter cards by enabling the user to easily mix cards and systems from different generations.

High Bandwidth Systems Demand PCI-X 2.0

Although the PCI-X 2.0 technology is applicable to any computer platform, the natural fit for this high-performance I/O interface is in systems such as enterprise servers, professional workstations, high-end UNIX servers, mainframes, networking, and communication applications. These systems need the high bandwidth offered by PCI-X 266 and PCI-X 533.

ELEMENTS OF PCI-X 2.0 LEVERAGED FROM PCI-X 1.0 AND PCI

The PCI-X 2.0 specification leverages heavily from previous generations of conventional PCI and PCI-X 1.0. The fact that so many previously designed elements are in place significantly improves product development time and time-to-market, while lowering development risk.

Architecture

PCI-X 2.0 uses the same basic architecture that transcends all generations of PCI. All generations also use the same set of control signals, which maintain the same functions across all generations.

State Machine

Only minor changes need to be made to the state machine to implement PCI-X 266 and PCI-X 533. This was one of the key elements of the design that will ease the transition.

BFM (Bus Functional Model)

PCI-X 266 and PCI-X 533 use a Bus Functional Model (BFM) that is nearly identical to that of PCI-X 133. Only a few lines of code are changed to create the BFM for these new speed grades.

Device Drivers

PCI-X 266 and PCI-X 533 device drivers are identical to the device drivers of PCI-X 133 and PCI-X 66. They can be reused as they stand without any modification.

High Bandwidth per Pin: One Pin per Bit

PCI-X 2.0 makes more efficient use of its pins than some of the other new "serial" I/O technologies. In fact, PCI technologies, including PCI-X 2.0, require one fourth the number of data pins for each bit of data.

"Serial" I/O technologies require differential signaling, and they are also uni-directional. As a result, one pair of wires is required to transmit data, and a second pair is required to receive data. The result is that four physical wires are required for each bit of data. By contrast, PCI technologies, including PCI-X 2.0, require only one physical wire for each bit of data. This 4X advantage in wire efficiency compensates for the lower clock frequency of PCI-X 2.0 and brings its performance head-to-head with the newest and most esoteric "serial" I/O interfaces – all while retaining full backward compatibility with traditional PCI adapter cards.

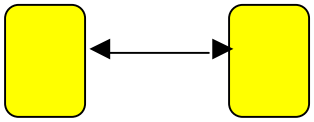

One bit of data in PCI technologies: 1 single bi-directional wire.

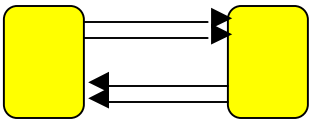

One bit of data in "serial" technologies: 4 wires -- 2 for transmit, 2 for receive

Component Pinouts
Components that interface to PCI-X 2.0 buses may be able to leverage the pinouts that have been used for their PCI-X-1.0-based components, reducing design time.

Designer Expertise
Because most designers and engineers are already familiar with PCI, there's no need for significant retraining or costly education. ASIC designers, signal integrity engineers, software engineers, and board designers can leverage much of their current expertise when they design with PCI-X 2.0.

Design Tools
Currently available design tools can be used for much of the design and testing of PCI-X 2.0 Although some tools require minor modifications, these modifications are already being made.

Connector
PCI-X 2.0 uses the same, proven connector that has served all generations of PCI. This standard connector is widely available and inexpensive because of significant economies of scale.

Power Management
For applications that are particularly power sensitive, power management solutions are already included in the PCI-X 2.0 specification. The power management protocols for PCI have been in use for years in notebook computers. These same protocols can be used with any PCI technology, and are included in the PCI-X 2.0 specification.

Hot-Plug Support
The ability to add or remove adapter cards from a powered-up server is a key capability for businesses that operate mission-critical server environments. This same Hot-Plug capability makes it possible for IT managers to upgrade and maintain systems without costly system downtime. The extensive Hot-Plug functionality in PCI-X 2.0 has already been defined, tested, and proven with previous generations of PCI technology. Hot-Plug capability is also supported in the PCI-X 2.0 specification.

NEW FEATURES

Strobes
To ensure that data is transferred accurately under all conditions, strobes have been added to the specification of PCI-X 266 and PCI-X 533. These strobes are used to trigger the clock input of the data latches and ensure that signals are latched at the precise time. Because strobes and signals are both subject to the same variations in process, temperature, and voltage, they will both drift in the same direction, and with the same magnitude. Therefore data is latched securely, even at the very high data rates of PCI-X 533.

1.5V Signaling
To achieve the very high frequencies of PCI-X 2.0, lower voltage signal swings were required. As a result, PCI-X 266 and PCI-X 533 require new 1.5V signaling. However, to maintain compatibility with previous-generations of 3.3V PCI technologies, the I/O buffers have been carefully designed to support both signal levels.

ECC Support

The PCI-X 2.0 specification includes ECC (Error Correcting Codes) to provide additional fault tolerance. Although the bit error rates for PCI-X 2.0 will be very low, this new feature will give PCI-X 2.0 additional robustness, making it ideal for the most error-sensitive applications. The ECC support in PCI-X 2.0 protects not only the data, but the header too. Because the protocol uses eight check bits it has the ability to correct single-bit errors and to detect dual-bit errors. Single-bit errors are automatically corrected, while dual-bit errors are tagged for retransmission. Also, the ECC in PCI-X 2.0 comes without incurring any performance penalty; it takes the same time to perform ECC functions as it does for the simple parity error detection of previous generations of PCI.

New PCI-X 2.0 Registers

To support the ECC functionality of PCI-X 2.0, new configuration registers specific to PCI-X 2.0 have been defined. Backward compatibility is still maintained because the new configuration registers automatically default to functional values for earlier generations of PCI.

Higher Bus Efficiency

Not only does PCI-X 2.0 technology operate at higher clock frequencies, and transfer more bits of data per clock cycle than previous generations of PCI, but the protocol is also more efficient. Analysis of previous generations of PCI highlighted the fact that overhead tasks were wasting significant bandwidth. This inefficiency has been addressed by enhancements to the PCI-X protocol to dramatically improve bus utilization. As a result, data is not only transferred at a higher clock rate, but it is also transferred more efficiently.

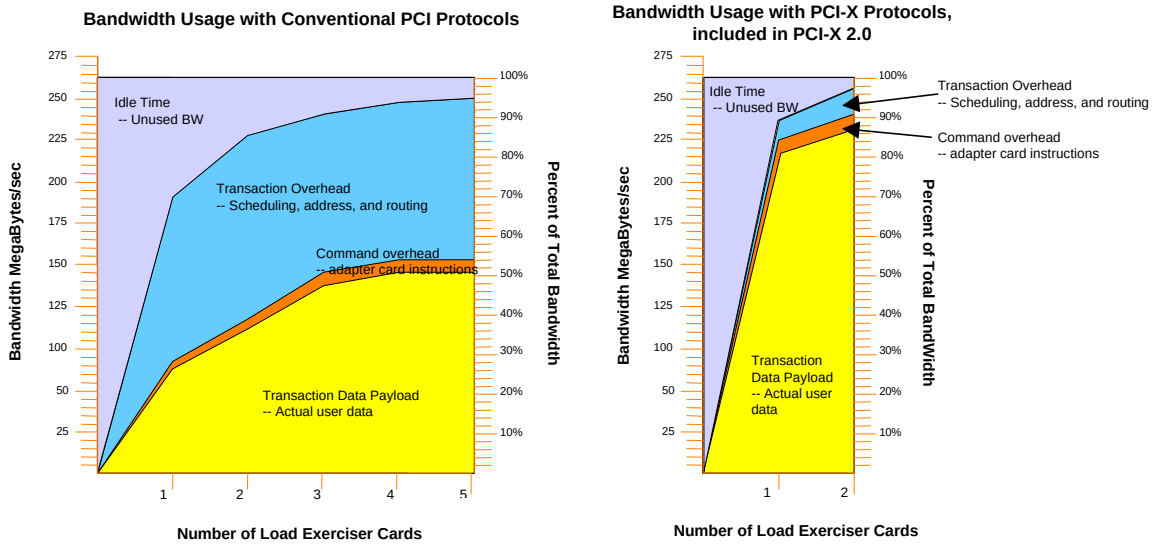

Conventional PCI technologies offered comparatively low usable bandwidth. These inefficiencies were corrected in the PCI-X protocol.

SUMMARY

The PCI standard has a long, successful history, beginning as a then-high-speed 33MHz standard and evolving over time to ultrahigh frequencies, multiple data rates, and highly efficient buses. Each new PCI standard has built upon the previous generation to meet the performance requirements of its time. A key to the success of PCI over many generations is the fact that software and hardware backward compatibility has been conscientiously maintained at every step. The design and implementation of each PCI generation has also been easy because designers and engineers can leverage their expertise acquired from previous generations. The PCI-X 2.0 specification is the next step in this highly successful legacy. PCI-X 2.0 offers the high performance needed for today's and tomorrow's applications in a standard that is easy to adopt and backward compatible with existing PCI technologies.