System Bus Technology

PCI, PCI-X and Beyond
A Technical Brief

Summary

As input/output devices, host bus adapters rely on matching data throughput requirements between the host system bus and that of attached storage devices. While Ultra320 SCSI and 2-Gigabit Fibre Channel standards have increased available bandwidth, the PCI standard has remained relatively static since the introduction of the PCI 64-bit, 66 MHz bus. The PCI-X standard provides a more effective match for these emerging storage standards while maintaining full backward-compatibility with the current PCI standard.

This technical brief explains the benefits of PCI-X and differences between this standard and the current PCI implementation.

PCI Development

The PCI (Peripheral Component Interface) standard is designed for high-speed, efficient data movement. It achieves this by taking advantage of the low overhead associated with local bus architecture – vastly improving the functionality and scalability of desktop computers and servers. This local bus uses a data path that directly interfaces with and accesses the computer CPU. This facilitates high-speed data transfers and requires processing power independent of the CPU so that commands can be carried out autonomously on the PCI bus with minimal CPU intervention. This allows the CPU to handle important tasks without the additional overhead of responding to all PCI activity.

The PCI architecture is, by far, the most widely adopted local bus standard (S-Bus, VME and NuBus are other examples). PCI bus performance is dependent on two factors: clock speed and bus width. Using a highway as an analogy, consider clock speed as the speed limit, while bus width is the number of lanes.

Increasing either of these allows you to move more automobiles. A PCI bus acts in a similar manner. Increasing clock speed or bus width allows you to move more data per second.

Since its introduction in the early 1990s, the PCI architecture has seen a number of improvements allowing faster data transfers. The following chart shows the available PCI implementations. Like SCSI technology, the PCI standard allows for backward compatibility, meaning that users can protect their investment in legacy devices.

Clock Speed	Bus Width	Overall Data Throughput
33 MHz	32-bit	132 MB/sec.
66 MHz	32-bit	264 MB/sec.
33 MHz	64-bit	264 MB/sec.
66 MHz	64-bit	528 MB/sec.

However, as computers get faster and more powerful the need for faster communication and 'fatter' data pipes continues to increase. The latest improvement, known as PCI-X, is slated to raise the performance bar to the 1-Gigabyte threshold. It is designed to increase I/O requirements for high bandwidth applications such as Gigabit Ethernet, Fibre Channel and Ultra320 SCSI.

PCI-X is the latest implementation of PCI bus technology. Using the same 64-bit architecture as the current standard, PCI-X has doubled the clock speed to 133 MHz, allowing transfer speeds up to 1 GB/sec. It is expected that PCI-X will have its first implementation on servers and workstations with eventual migration to the desktop market.

The adoption rate of PCI-X will depend on how quickly users upgrade current technology to take advantage of PCI-X. Already available to the high-end marketplace in late-2001 and expected in mid-range servers by late 2002, PCI-X will facilitate the migration of high-end video editing and Enterprise-class data management to servers and workstations.

PCI-X hardware is available for testing purposes and has been implemented in Fibre Channel and SCSI host adapters. Companies such as Compaq have integrated PCI-X into their product lines.

Value Proposition of PCI Architectures

When deployed properly, the PCI bus architecture facilitates the use of bandwidth intensive devices such as attached SCSI storage. The foremost example of this is the creation and distribution of rich media content.

Creating and editing video using a computer requires 10-160 MB/sec. throughput for smooth playback. Considering that a 30-second commercial shot in high definition format requires up to three terabytes of storage, it is easy to imagine the complexity of being able to deliver each frame of this clip on time and in order. Working hand in hand with SCSI, not only is PCI able to deliver this performance, it is capable of delivering multiple streams of high definition video simultaneously on the same computer system. PCI-X and Ultra320 SCSI, both due for deployment in early 2002, take storage performance to a new frontier with bandwidth available to edit dual-stream high definition in real time.

PCI *and Beyond* - InfiniBand

The next generation of PCI technology, code named Arapaho, has the ability to deliver throughputs 10x faster than the current technology of 1.1 GBsec. Early chip products would enable products to begin showing up in late 2003 with software devices drivers being compatible with what PCI is today.

InfiniBand represents an industry wide effort to develop a replacement for the PCI bus and as thus represents a significant change in server architecture. It draws heavily from Intel's Virtual Interface architecture and was designed to support storage, networking and inter-processor communications.

Because InfiniBand involves such sweeping changes in server architecture the advantages InfiniBand brings over current I/O technology may not be readily available in server architectures for at least two years.

There is an ongoing debate over whether PCI and InfiniBand are competing or complementary technologies. InfiniBand advocates argue that it is a superior technology in all aspects and will eventually replace PCI in all applications. The reality is that the PCI will be around for a long, long time to come. Present in virtually 100% of desktop systems, there is a huge legacy presence that will demand support for years to come. Additionally, while InfiniBand is optimized for distributed I/O, PCI is superior for local I/O. So while InfiniBand has advantages in a SAN environments, network environments and inter-processor communications, it does not operate as effectively in a direct attached storage environment.

ATTO - PCI and InfiniBand

All ATTO host bus adapters are based on the PCI standard. Current Fibre Channel and SCSI products use the standard PCI interface (both 32-bit and 64-bit). Fibre Channel host adapters are scheduled to be upgraded to PCI-X later this year. The FC 3320 dual-channel adapter will be released as a PCI-X product. The UL4D/UL4S Ultra320 host adapter will be PCI-X compatible when released later this year.

ATTO will debut its first generation InfiniBand products in late 2002. ATTO's initial InfiniBand products will provide InfiniBand to Fibre Channel connectivity, enabling InfiniBand switches to connect to Fibre Channel SANs.

We will initially develop target channel adapters (TCA), which will enable customers to connect InfiniBand devices to SANs using the same systems. These TCAs will initially provide 2.5 Gb/sec. InfiniBand and 2Gb/sec. Fibre Channel connectivity.

In time, server architectures will migrate to InfiniBand. But until that, with it's large installed base PCI and PCI-X will continue to be the predominant bus architecture.

Terms

Bus – is analogous to a highway system for the computer. There are several busses inside each computer. The bus is responsible for moving data between system components.

Bus Width – like a highway, a bus can be described by the number of "lanes" it has. This refers to the number of data bits that can be moved simultaneously per clock cycle. Common PCI bus widths are 32 and 64 bits.

CPU – Central Processing Unit. The computer's processor (i.e. Intel Pentium, AMD Thunder K7).

Local Bus – a bus that directly accesses the CPU. This direct link allows faster data movement because the movement of data is less affected by other components. The PCI bus is a local bus.

Clock Speed – refers to the number of clock cycles a certain component generates per second. This is typically described in Megahertz (MHz). PCI architectures run as slow as 8 MHz and as fast as 133 MHz. Most commonly, it runs at 33 or 66 MHz.

Device Driver – specific code tailored to tell hardware how to communicate with the operating system. Drivers are required for SCSI host adapters to properly transfer data to and from the storage device and across the PCI bus.

CompactPCI - as part of the general PCI specifications, CompactPCI (cPCI) has found widespread acceptance in the telecommunications industry and in Europe. With similar performance as the general PCI implementations, cPCI has a more robust connector that lends itself to rugged field applications or embedded systems like the ATTO Diamond storage array. cPCI is slightly more expensive than standard PCI and is primarily a niche technology.