



PCI Express

- origins, advantages and likely evolution.

The History

Back in 1984, IBM shipped its PC AT with the venerable ISA (Industry Standard Architecture) bus. ISA was a 16-bit interface, which meant that data could only be transferred two bytes at a time. Initially this wasn't a problem for inherently slow devices available at the time - COM ports, printer ports, sound cards or CD-ROMs. However, the ISA bus proved too slow for later high performance disk access and display adapters. When the ISA bus started running out of steam, other architectures were put forward to open the bottleneck but it was the PCI bus that successfully brought the much needed bandwidth to the system

The Bottleneck

Since the 1990's a lot of progress has been made in improving the speed of processors, memory and other PC peripherals. However, the technology connecting these devices to each other - the PCI bus, has remained largely unchanged, creating a system bottleneck.

Figure 1 shows many of the contributors to the PCI bus bottleneck. This figure shows the bandwidth required by various communications, video, and external devices that are serviced by the PCI

bus. It can be seen that the multidrop, shared PCI bus is hard pressed to keep up with today's devices. This situation worsens with upcoming peripheral devices with even higher data rates. For example, Gigabit Ethernet requires a bandwidth of 125 MB/sec, and the IEEE 1394b bus has a maximum bandwidth of 100 MB/sec, more or less saturating a standard PCI bus.

As computers became increasingly graphically-oriented, the bandwidth limitations of PCI became especially apparent in the plug-in

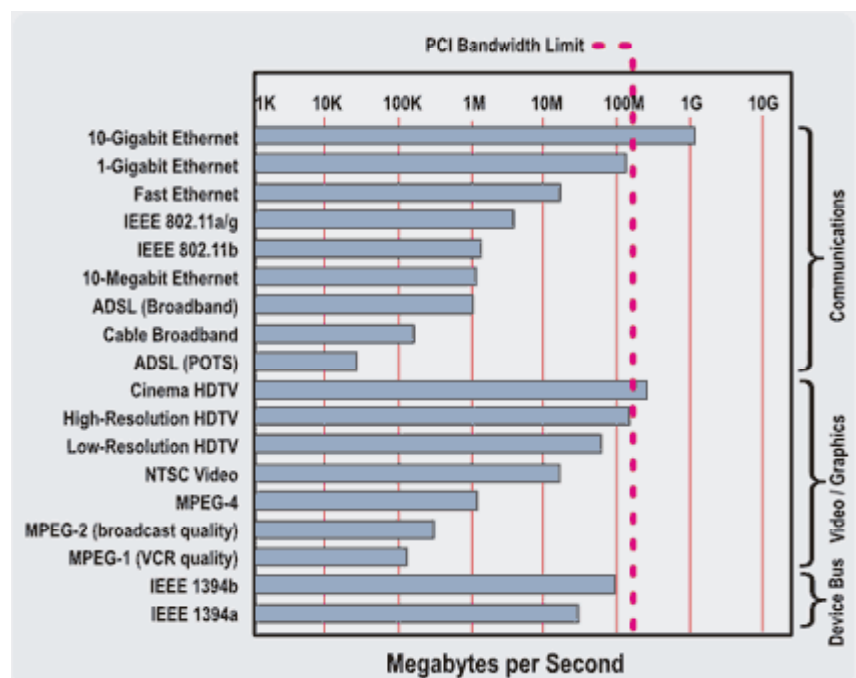


Figure 1: Bandwidth of Devices serviced by the PCI Bus

graphics card. The development of the AGP slot went some way in opening up the bandwidth but never displaced PCI. Although AGP is still incorporated into many motherboards today, it is gradually being phased out in favour of PCI Express. Many newer graphics cards are available exclusively in PCI Express, with AGP versions of the same cards becoming correspondingly rare. At this rate, it is expected that AGP will be almost totally replaced within a few years.


The Emergence of PCI Express

PCI-Express (formerly 3GIO) was a development effort, led by the likes of Intel. This development was designed to regain the balance between raw CPU speed and system speed. Early 2002, ownership of 3GIO was transferred to the PCI-SIG where it was re-named the PCI Express Architecture – signalling industry consensus on PCI and AGP’s successor. Often abbreviated to PCIe or PCIx, it has nothing to do with the PCI-X bus found on many workstation and server class boards.

The PCI Express standard was designed to be completely transparent to software developers - an operating system and device drivers designed for PCI can boot in a PCI Express system without any code modification.

Architecture

The PCI Express Architecture defines a flexible, scalable, high-speed, point-to-point, hot pluggable/hot swappable interconnect that is software-compatible with PCI. The underlying design of the PCI Express architecture differs completely from that of the “classical” PCI bus. For one thing, PCI Express uses a serial scheme, which allows for higher clockspeeds.



This, coupled with the fact that several PCI Express channels (or “lanes”) can be bundled

together in an expansion slot leaves room for stages of expansion, such as x8 and x12. Standard add-in cards may be 1x (low bandwidth), graphics cards may be x16 (very high bandwidth), depending on the needs of the cards. The same data rate is available in both directions simultaneously.

PCI Express cards will physically fit into slots designed for their lane configuration or higher (up-plugging) but not into slots designed for lower lane configurations (down-plugging). So, for example, a x1 card will fit into x1, x4, x8, and x16 slots but a x16 card will only fit into a x16 slot. A x1 card in any compliant PCI Express slot will always run in x1 mode

It’s anticipated that the standard expansion slot on future motherboards will be the PCI Express x1 slot. In this case, “x1” means the slot has one PCI Express lane, which will give it a bandwidth of 250MB/s (500MB/s if you count the bandwidth in both directions, i.e. full-duplex), making it almost twice as fast as PCI. On top of that, devices will no longer find themselves competing for bandwidth, as each slot will have the full 250MB/s for itself.

Performance

Increased bandwidth can be equated into increased system performance. It has long been known that to better utilise a processor, data must be transferred to it as quickly as possible. Chipset designers have consistently addressed this by increasing Front Side Bus speeds.

Common Buses and their Theoretical Max Bandwidth	
FC	~ 32 MB/s
AGP EX	2,100 MB/s
PCI Express x1	250 [500]* MB/s
PCI Express x2	500 [1000]* MB/s
PCI Express x4	1,000 [2000]* MB/s
PCI Express x8	2,000 [4000]* MB/s
PCI Express x16	4,000 [8000]* MB/s
PCI Express x32	8,000 [16000]* MB/s
IDE (ATA-70)	~ 100 MB/s
IDE (ATA-133)	~ 133 MB/s
SATA	~ 150 MB/s
Serial ATA	~ 25 MB/s
Serial ATA (FireWire)	~ 100 MB/s

The problem with this is that Front Side Bus speed increases the speed of transfer between the memory and CPU but does not address data from other sources that require CPU or memory access, such as drives, network traffic, video etc. PCI Express addresses this problem head on by increasing bandwidth between devices allowing them to interact without using CPU time.

However, it’s worth remembering that bus bandwidths are theoretical measurements.

These will only be fully realised when the market designs in PCI Express, optimising silicon and software to take full advantage of the additional bandwidth offered by this architecture. For example - although a x16 PCI Express connection is at least 190% faster than AGP 8x, many current chipsets are designed around AGP, so a huge performance increase may not be observed.

Summary

It's a familiar model in the IT industry to see the market driving the consumer in new technology adoption. As of 2005, PCI Express appears to be well on its way to becoming the new backplane standard in personal computers. The latest generation of desktops are based around Intel's 9XX chip sets (aka Grantsdale and Alderwood) with PCI Express bus slots on the motherboard. Dell began to transition its consumer desktop computers from AGP to PCI Express from 2004.

Most of the new graphics cards from both ATI Technologies and NVIDIA use PCI Express.

NVIDIA also uses the high-speed data transfer of PCI Express for its newly developed Scalable Link Interface (SLI) technology, which allows two graphics cards of the same chipset and model number to be run at the same time, allowing increased performance. Most new Gigabit Ethernet chips and some 802.11 wireless chips also use PCI Express.

Other hardware such as RAID controllers and network cards are also starting to make the switch. In 2005, Apple updated both the consumer iMac and workstation PowerMac to use PCI Express exclusively, hence supplanting the AGP and PCI-X connectivity that they had formerly utilized. The main point here is that a system upgrade is not necessary just to go to PCI Express, but when buying an industrial computer, a unit combining PCI and PCI Express options will future proof your investment and support the next generation of network, RAID and IO cards.