PCI Express Basic Info

*This info also applies to Laptops

- PCI Express Laptops
- PCI Express Motherboards
- PCI Express Video Cards
- PCI Express CPU Motherboard Combo's
- PCI Express Barebone Systems
- PCI Express SLI
- PCI Express Video Cards Review

One of the most significant changes to come to PC systems in a decade, PCI Express is a new interconnect technology designed to provide universal connectivity for use as a chip-to-chip and chip to adapter card interconnect. PCI Express architecture provides for extremely high bandwidth at low cost. There are two main reasons to care about PCI-Express: 1) PCI is now an old standard dating back to the early 90's and no longer fits our needs in terms of speed/performance. 2) AGP also is in a similar position as PCI now, and chipset manufacturers are killing AGP motherboard support in favor of the much faster PCI Express interface. This means you are looking at a forced transition in the graphic sector, thus you really don't have a lot of choice in the coming years.

While we've spent plenty of time and energy improving the speed of processors, memory, and other parts of the PC we've gone virtually nothing with the main connection between many devices-PCI. As such we are stuck with a technology in our PCs and Servers that still runs at the speeds and bandwidth we were comfortable with in the 90's. PCI as we know it is holding us back - it is a bottle neck - a limitation to the maximum performance of our systems.

We all want the most from our PC. To get the most out of our PC we must remove all bottlenecks (obstacles to performance). To that end we must turn to the next best alternative: PCI Express.

PCI Express can offer up to 70 times the bandwidth of today's PCI architecture and is scaleable for the future. PCI Express will be featured across all Intel platforms including desktop, server, workstation and in the latter half of 2004 with mobile platforms as well. PCI Express will be the I/O architecture for everything from graphics adapters to Ethernet cards to TV tuners. This massive bandwidth will alleviate many current and future performance bottlenecks on the adapter bus.

PCI Express is based on a type of serial communications technology somewhat like that in USB or SATA hard drives. The mechanical (physical) board connectors come in one of four types: x1, x2, x4, and x16 (see illustration below) in order to meet different peak bandwidth requirements.

PCI Express Technical Specs:

- **Full duplex point-to-point topology**
- **Differential low voltage interconnect**
- **Embedded clocking**
- **Scalable frequency: Initial Bit Rate: 2.5Gb sec/lane/direction**
- **Scalable bandwidth - data layer is scalable to 1x, 2x, 4x, 8x, 12x, 16x, 32x lane widths**
- **Each PCI Express "lane" uses 4 wires - one differential pair for transmit and one pair for receive**

* Note: PCI Express is NOT the same as PCI-X slots, it is a totally new technology.

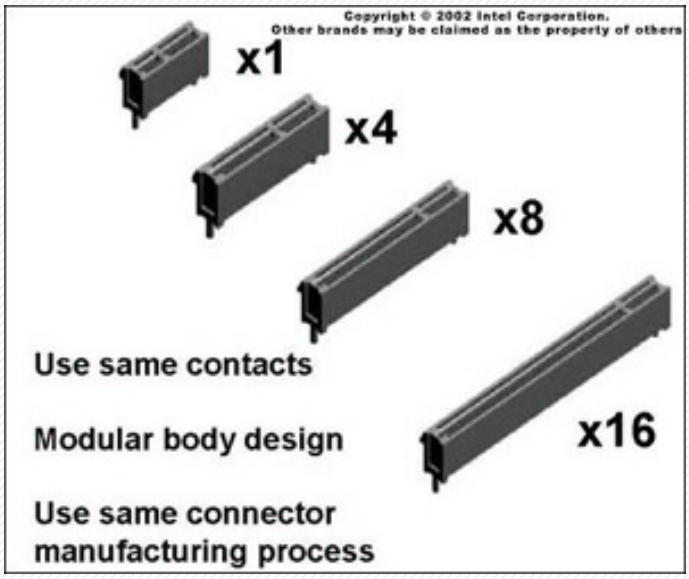

PCI Express Bandwidth

LANES	Peak Bandwidth (Duplex Mode)
x1	500MB/s
x4	2 GB/s
x8	4 GB/s
x16	8 GB/s

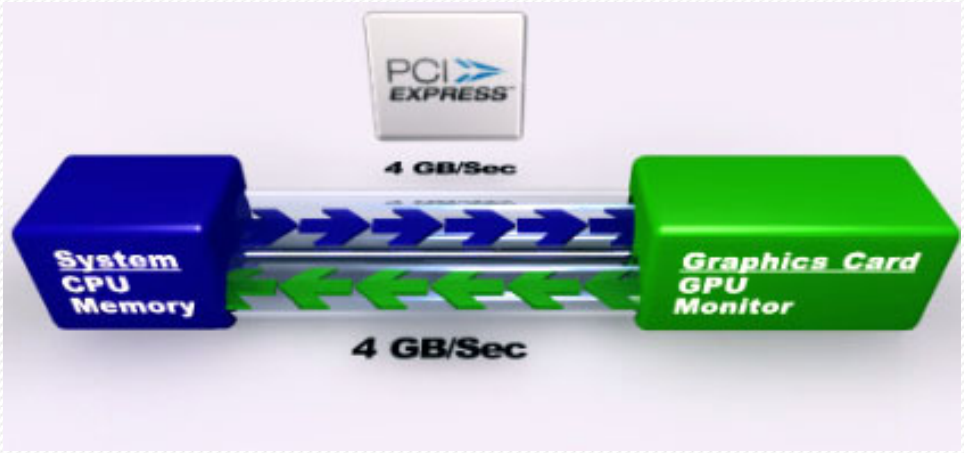

PCI express is a highly flexible, reliable, modular and scalable design that will eventually replace all PCI slots on the motherboard and AGP slots. It has better power management, native hot-plug support, backwards compatibility with PCI software, support for streaming media (such as video camera or TV), and truly scalable configurations. In addition:

- Compatible with existing PCI drivers and software and operating systems
- High bandwidth per pin. Low overhead. Low latency
- Ability to scale speeds by forming multiple lanes
- A point-to-point connection, allows each device to have a dedicated connection without sharing bandwidth
- Ability to comprehend different data structures
- Low power consumption and power management features

- Hot swap-ability and hot plug-ability for devices
- Supported by nearly 500 system hardware vendor

PCI Express and 3D Graphics

The x1 PCI Express slots will easily replace the standard 32-bit PCI slots and have four times the bandwidth.

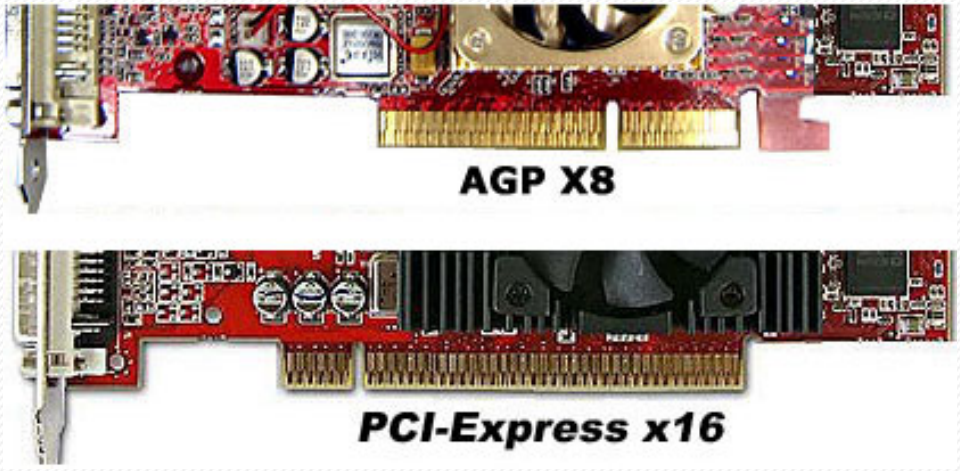

The high-performance x16 configuration will have up to 4GB/sec bandwidth to replace AGP technology and will also have four times the bandwidth of AGP 8x!

With the advent of PCI Express video cards whole new worlds of 3D gaming and superior graphics performance will be possible. ATI and nVidia have already developed video cards using PCI express architecture and in just a short couple years will be the dominant video card interface and only choice for 3D graphics power users. The new video processors have a native, or “true” PCI Express interface. They can communicate directly with the PCI Express bus at PCI Express speeds (do not need to use a bridge).

With double the bandwidth of the AGP 8X graphics bus, NVIDIA’s PCI Express solutions deliver a new level of PC performance for graphics and networking. NVIDIA’s desktop, mobile, and workstation PCI Express solutions provide something for every type of PC user; delivering faster graphics and system performance for your PC.

While the classic PCI Bus is based upon a parallel architecture, PCI Express is serial based, drastically reducing pin count. It is a point to point protocol much like AGP. Devices do not share bandwidth.

The usual PCI system is laid out in something like the following diagram:

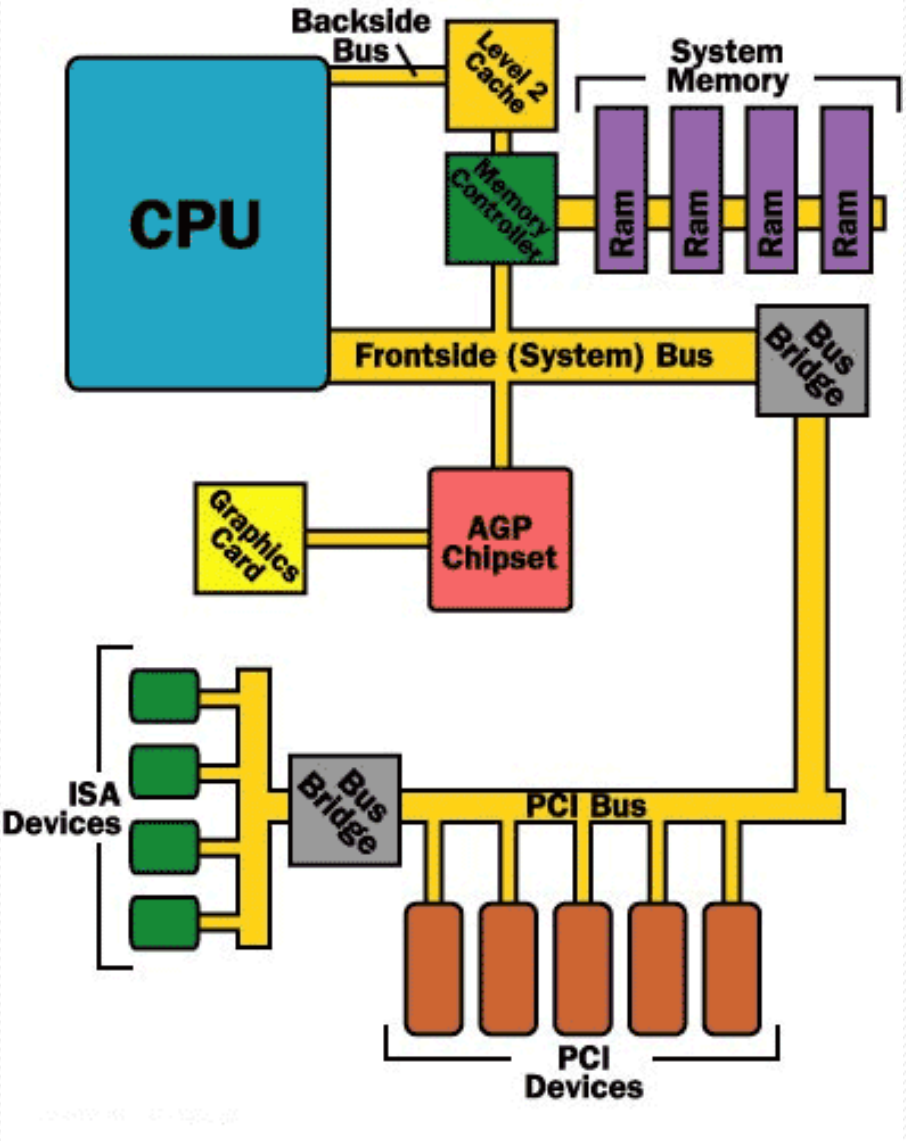

PCI system layout

Intel introduced a new bus standard for consideration, the **Peripheral Component Interconnect** (PCI) bus. PCI presents a hybrid of sorts between ISA and VL-Bus. It provides direct access to system memory for connected devices, but uses a bridge to connect to the frontside bus and therefore to the CPU. Basically, this means that it is capable of even higher performance than VL-Bus while eliminating the potential for interference with the CPU.

The frontside bus is a physical connection that actually connects the processor to most of the other components in the computer, including main memory (RAM), hard drives and the PCI slots. These days, the frontside bus usually operates at 400-MHz, with newer systems running at 800MHz.

As processor speeds steadily climb in the GHz range, many companies are working feverishly to develop a next-generation bus standard. Many feel that PCI, like ISA before it, is fast approaching the upper limit of what it can do.

PCI-Express is aimed at the home computer market, and could revolutionize not only the performance of computers, but also the very shape and form of home computer systems. This new bus isn't just faster and capable of handling more bandwidth than PCI. PCI-Express is a point-to-point system, which allows for better performance and might even make the manufacturing of motherboards cheaper. PCI-Express slots will also accept older PCI cards, which will help them become popular more quickly than they would if everyone's PCI

components were suddenly useless.

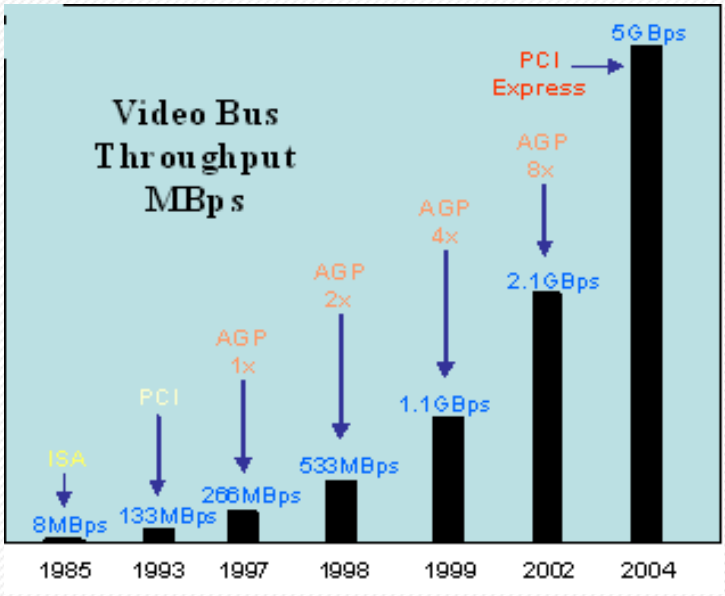

A basic PCI-Express slot will be a 1x connection. This will provide enough bandwidth for high-speed Internet connections and other peripherals. The 1x means that there is one lane to carry data. If a component requires more bandwidth, PCI-Express 2x, 4x, 8x, and 16x slots can be built into motherboards, adding more lanes and allowing the system to carry more data through the connection. In fact, PCI-Express 16x slots are already available in place of the AGP graphics card slot on some motherboards. PCI-Express 16x video cards are at the cutting edge right now, costing more than $500. As prices come down and motherboards built to handle the newer cards become more common, AGP could fade into history.

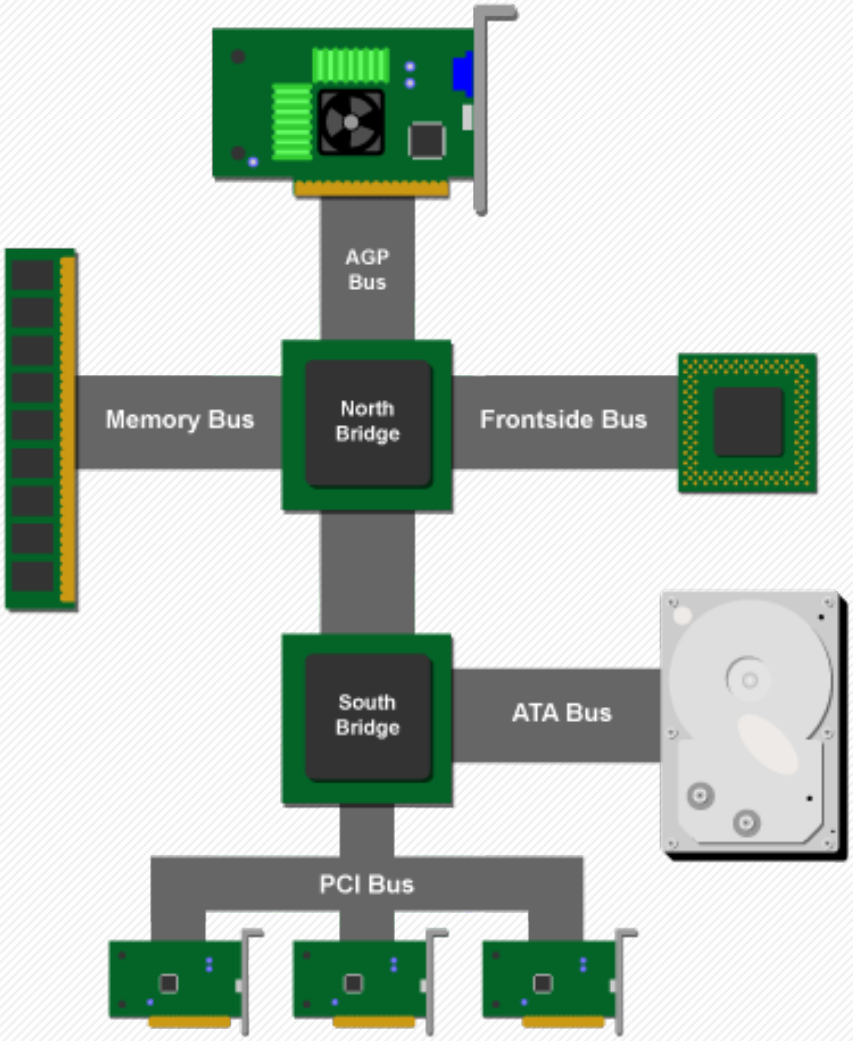

Northbridge and Southbridge

PCI uses a **shared bus topology** to allow for communication among the different devices on the bus; the different PCI devices (i.e., a network card, a sound card, a RAID card, etc.) are all attached to the same bus, which they use to communicate with the CPU.

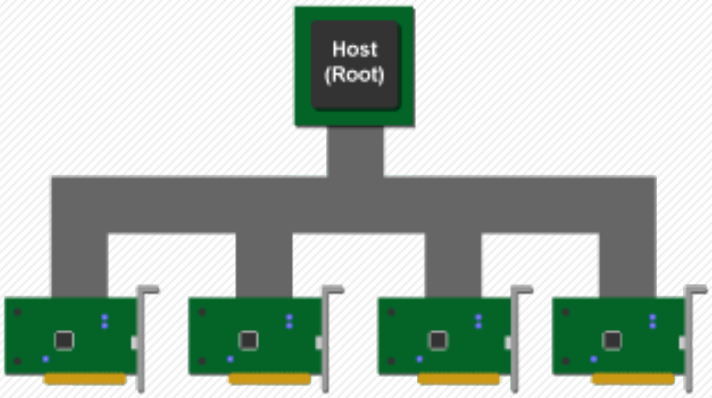

Shared bus

Because all of the devices attached to the bus must share it among themselves, there has to be some kind of **bus arbitration** scheme in place for deciding who gets access to the bus and when, especially in situations

where multiple devices need to use the bus at the same time. Once a device has control of the bus, it becomes the **bus master**, which means that it can use the PCI bus to talk to the CPU or memory via the chipset's southbridge.

The shared bus topology's main advantages are that it's simple, cheap, and easy to implement — or at least, that's the case as long as you're not trying to do anything too fancy with it. Once you start demanding more performance and functionality from a shared bus, then you run into its limitations. Let's take a look at some of those limitations, in order to motivate our discussion of PCI Express's improvements.

PCI Express is going to be an important step in bringing convergence to the PC industry by modernizing and standardizing on an I/O standard. The immediate benefactors will be manufacturers more so than the end user. PCI Express is supposed to be cheaper to implement than the traditional PCI Bus and the AGP slot. It should also alleviate the need for custom buses like the CSA which the Intel 875 chipset uses for Gigabit Ethernet.

PCI Express is focused towards the graphics segment and the big jump in bandwidth offered by the PCI Express x16 slot. The biggest problem is that current AGP cards will not work which means potentially costly upgrades on the graphics front particularly at the enthusiast segment which is more likely to be buying cutting edge parts more so than the mainstream crowd. The implication that the extra bandwidth offered by PCI Express will somehow translate into big performance gains will not likely hold up unfortunately. Just as we have seen little real world improvement transitioning from AGP 4X to AGP 8X, the extra bandwidth offered by PCI Express x16 is likely not required initially nor for some time as clearly, the AGP bus is not the limiting factor on current generation video cards. Other architectural changes like improved latency may help somewhat but users should not be surprised to see little difference between AGP and PCI Express cards.

More Images and a few FAQs

PCI-Express was a development effort, led by the likes of Intel, to advance the I/O (input output) functionality of today's computers. With the ever greater demand for fast processing and with the CPU being bottlenecked by aging system busses, this development was designed to regain the balance between raw CPU speed and system speed.

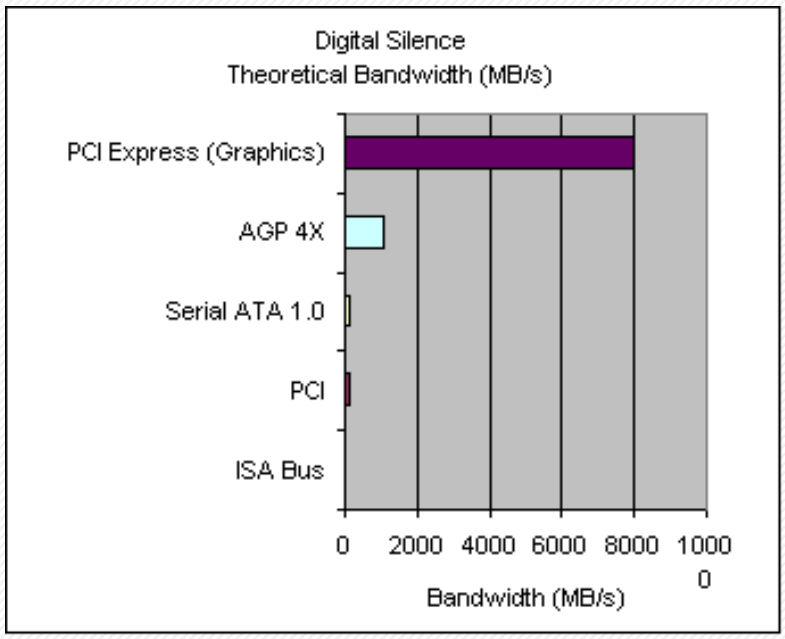

Image above referenced Here

PCI-Express is the industry's attempt to unify all of the current different types of I/O bus into a single "future proof" standard. Over the past ten years, PCI has handled the large and varying uses it has been given, most of which were never foreseen when the specification was made. Ports like the AGP, ATA and USB, were developed to cope with data transmissions that needed greater support for time dependant data. Things like video streaming and other real-time applications were not explicitly address in previous PCI specifications (PCI 2.2 and PCI-X included). Current PCI specifications are based on a multi-drop, parallel bus implementation that is coming very close to its performance limits.

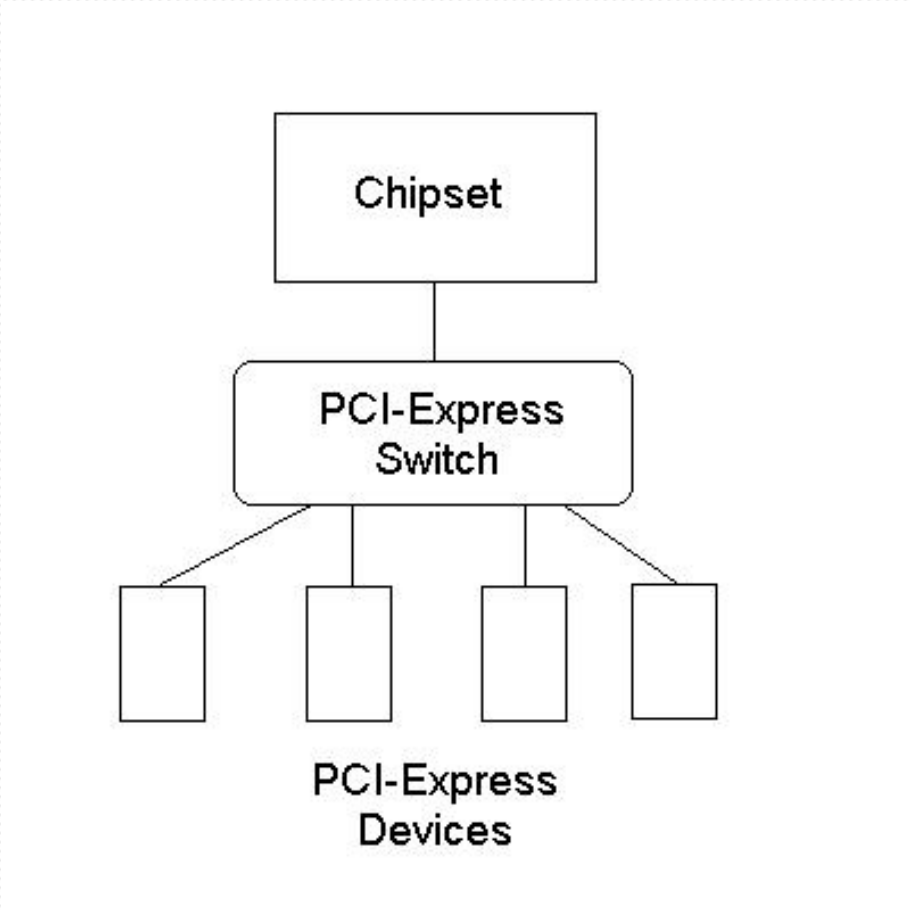

PCI-Express takes a leaf out of the book of the LAN and adds a switch to the system topology. The switch replaces the multi-drop bus and is used to distribute I/O messages on a peer-to-peer basis. This means that if one PCI-Express device wants to send data to another, it doesn't necessarily need to go through the chipset (even though the switch may be part of the chipset). This reduces the amount of messages that the chipset has to process itself. Next, from the switch comes PCI-Express links. Each link can contain many 'lanes' making each device link individually scaleable, in turn, adding more bandwidth with the addition of each lane. This is where the term 1x, 2x 4x, 16x etc comes in. Standard add-in cards may be 1x (low bandwidth), graphics cards may be 16x (very high bandwidth), depending on the needs of the cards.

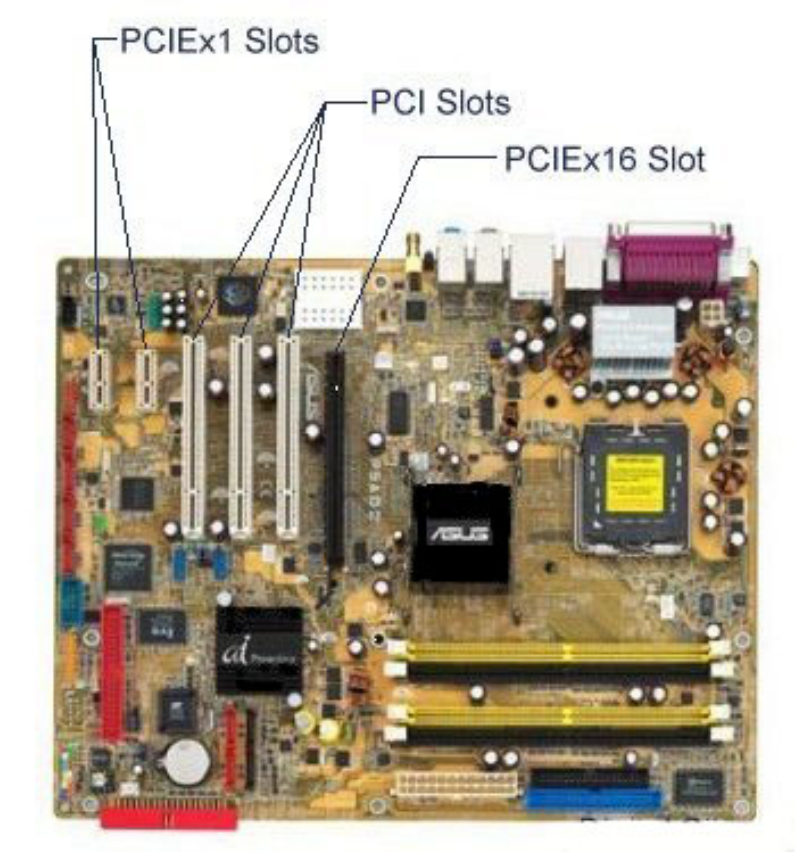

ASUS P5-AD2 Deluxe

So what do I get?

The biggest benefit is the speed of PCI-Express. The basic speed of 2.5 Giga transfers per second per channel direction gives each channel a 200MB per second in data capacity (That's 100MB per second per pin!!). This is almost twice that of current PCI.

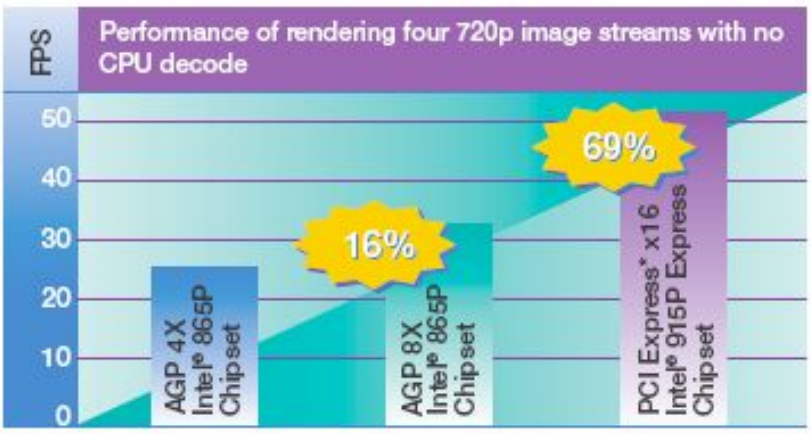

The point-to-point design also means that many wild and wacky system designs are possible. PCI-Express connections can be made with connectors and cables due to the low signal count.

Frequently Asked Questions

After we released the first versions of this info, we received a few questions.

Q. *How will the addition of PCIe to my system affect interrupt handling?*

A. Hopefully the only change will be the addition of the new slots for you to use. PCIe is fully backwardly compatible with existing PCI technology. Current operating systems like Windows XP will quite happily use the PCIe slots without issue. Although some versions of Windows NT/2000 will not be able to use the new PCIe slots without some tweaking. Longhorn, according to some random errata on microsoft.com, will have native support for PCIe, but until then, the operating system will emulate the slots like they were typical PCI slots. This does not mean that they will perform like current PCI, only that they may get better in the future when the operating system is written to support PCIe functions fully.

Q. *So after all this techno-babble what does it mean for the end user?*

A. From the system performance point of view, the advantages to the end user should be obvious. The increase in bandwidth between devices and the fact that devices can interact without using CPU time is a distinct performance advantage over the current PCI set-up. The first thing will be the requirement for a new motherboard with PCIe on it. Some manufacturers (like the ASUS P5-AD2 Deluxe shown above) are providing boards with both the current PCI we know and the new PCIe slots. Therefore, transitioning your hardware can take place in phases.

One other point that was raised by HardOCP was to do with PCIe graphics cards and PSU's. "24-pin PSUs will be a requirement for the GeForce 6800 on a motherboard that might not be that well engineered." So your old faithful PSU that has served you so well in the past may also need replacing. Then, as an after thought, your operating system may need updating. As mentioned above, some versions of windows do not support PCIe straight away, so if for some strange reason your still running Windows 95, be prepared.

Q. *Over the years, consumers complained that games were not taking full advantage of AGP4X, let alone AGP8X. What makes PCI Express any different?*

A. While developers had more than enough bandwidth available in AGP8X, they claimed that the AGP bus was too slow for their needs. The speed of the PCI Express will change that...albeit very slowly. When developers determine what features and technologies to support in their upcoming games, they project what the hardware base will be when the game is released. Since it will be quite a while before the vast majority of computers have PCI Express, don't look for too many games to take advantage of PCI Express for awhile.

Will a PCIe video card work on Windows XP?

Yes. A PCIe video card will work on current operating systems, because PCIe is software compatible with PCI. PCIe hardware will work on operating systems that support PCI. For example, on current versions of Windows an x16 video card will be able to transfer data at the higher x16 rate of 4 Gbps.

Will PCIe graphics coexist with AGP?

There will be chipsets that support both AGP and X16 PCIe. Some motherboards will have both AGP and X16 PCIe slots using such a chipset.

ECS 915P-A (1.0) PCI Express Motherboard

Will multimonitor configurations work on PCIe graphics?

Multimonitor configurations of PCIe are expected to work just like PCI. Whether they do will depend on the motherboard manufacturers. For example, x16, x8, and x8 triple monitor configuration will necessitate the existence of one x16 and two x8 slots on the motherboard.

What are the performance implications of using PCIe Graphics?

High-speed PCIe graphics will have better performance than AGP will. The PCIe graphics cards will use the x16 PCIe slot. This translates into a bandwidth of 4 Gbps (Gigabytes per second). This is already a twofold increase over AGP 8X. In this case, "x1" means that the slot has one PCIe lane, which will give it a bandwidth of 264 Mbps. This is equal to the bandwidth provided by AGP 1X and twice that of PCI (132 Mbps).

PCIe version	AGP	Bandwidth
PCIe x1	AGP 1X	264 Mbps
PCIe x4	AGP 4X	1 Gbps
PCIe x8	AGP 8X	2 Gbps
PCIe x16	2 x AGP 8X	4 Gbps

Additionally, the AGP specification does not support "snooping." It implies that memory used by devices needs to be mapped uncached or write combined by the processor in order to prevent the processor from caching that memory, or else an expensive cache flush needs to be done between handoff of a surface between CPU and GPU. Thus, processor read access to that memory will be very slow.

PCIe will support snooping. It will now be possible to map such shared memory as cacheable and still be able to maintain coherency between the CPU and the GPU. Snooped transactions are slower than nonsnooped transactions, but since the CPU can read the shared memory at full speed and we do not need to flush any caches, the tradeoff might mean better performance in some scenarios.

One of the main goals of the design of PCI-Express was to achieve a low cost solution capable of high volume application. In other words, the add-ins cards can become simpler with fewer pins, which means that unit costs

can be reduced. So today's modem will get smaller and cheaper but it also gives future devices room to grow.