PCI Express

PCI Express

Year created	2004
Created by	Intel · Dell · IBM · HP
Supersedes	AGP · PCI · PCI-X
Width in bits	1–32
Number of devices	One device each on each endpoint of each connection. PCI Express switches can be used to create multiple endpoints out of one endpoint to allow sharing of one endpoint with multiple devices.
Capacity	Per lane (each direction): • **v1.x**: 250 MB/s (2 Gbit/s) • **v2.x**: 500 MB/s (4 Gbit/s) • **v3.0**: 1 GB/s (8 Gbit/s) 16 lane slot (each direction): • **v1.x**: 4 GB/s (32 Gbit/s) • **v2.x**: 8 GB/s (64 Gbit/s) • **v3.0**: 16 GB/s (128 Gbit/s)
Style	Serial
Hotplugging interface	Yes, if ExpressCard or PCI Express ExpressModule
External interface	Yes, with PCI Express External Cabling such as Intel Thunderbolt.

PCI Express (**Peripheral Component Interconnect Express**), officially abbreviated as **PCIe**, is a computer expansion card standard designed to replace the older PCI, PCI-X, and AGP bus standards. PCIe has numerous improvements over the aforementioned bus standards, including higher maximum system bus throughput, lower I/O pin count and smaller physical footprint, better performance-scaling for bus devices, a more detailed error detection and reporting mechanism, and native hot plug functionality. More recent revisions of the PCIe standard support hardware I/O virtualization.

The PCIe electrical interface is also used in a variety of other standards, most notably ExpressCard, a laptop expansion card interface.

Format specifications are maintained and developed by the PCI-SIG (PCI Special Interest Group), a group of more than 900 companies that also maintain the Conventional PCI specifications. PCIe 3.0 is the latest standard for expansion cards that is available on mainstream personal computers.[1] [2]

Applications

PCI Express is used in consumer, server, and industrial applications, as a motherboard-level interconnect (to link motherboard-mounted peripherals), a passive backplane interconnect and as an expansion card interface for add-in boards.

In virtually all modern PCs, from consumer laptops and desktops to enterprise data servers, the PCIe bus serves as the primary motherboard-level interconnect, connecting the host system processor with both integrated-peripherals (surface mounted ICs) and add-on peripherals (expansion cards.) In most of these systems, the PCIe bus co-exists with 1 or more legacy PCI busses, for backward compatibility with the large body of legacy PCI peripherals.

Architecture

Conceptually, the PCIe bus can be thought of as a high-speed serial replacement of the older PCI/PCI-X bus,[3] an interconnect bus using shared address/data lines.

A key difference between PCIe bus and the older PCI is the bus topology. PCI uses a shared parallel bus architecture, where the PCI host and all devices share a common set of address/data/control lines. In contrast, PCIe is based on point-to-point topology, with separate serial links connecting every device to the root complex (host). Due to its shared bus topology, access to the older PCI bus is arbitrated (in the case of multiple masters), and limited to 1 master at a time, in a single direction. Furthermore, the older PCI's clocking scheme limits the bus clock to the slowest peripheral on the bus (regardless of the devices involved in the bus transaction). In contrast, a PCIe bus link supports full-duplex communication between any two endpoints, with no inherent limitation on concurrent access across multiple endpoints.

In terms of bus protocol, PCIe communication is encapsulated in packets. The work of packetizing and depacketizing data and status-message traffic is handled by the transaction layer of the PCIe port (described later). Radical differences in electrical signalling and bus protocol require the use of a different mechanical form factor and expansion connectors (and thus, new motherboards and new adaptor boards); PCI slots and PCIe slots are not interchangeable. At the software level, PCIe preserves backward compatibility with PCI; legacy PCI system software can detect and configure newer PCIe devices without explicit support for the PCIe standard, though PCIe's new features will not be accessible.

The PCIe link between 2 devices can consist of anywhere from 1 to 32 lanes. In a multi-lane link, the packet data is striped across lanes, and peak data-throughput scales with the overall link width. The lanecount is automatically negotiated during device initialization, and can be restricted by either endpoint. For example, a single-lane PCIe (x1) card can be inserted into a multilane slot (x4, x8, etc.), and the initialization cycle will autonegotiate the highest mutually supported lanecount. The link can dynamically downconfigure the link to use fewer lanes, thus providing some measure of failure tolerance in the presence of bad/unreliable lanes. The PCIe standard defines slots and connectors for multiple widths: x1, x4, x8, x16, x32. This allows PCIe bus to serve both cost-sensitive applications where high throughput is not needed, as well as performance-critical applications such as 3D graphics, network (10 Gigabit Ethernet, multiport Gigabit Ethernet), and enterprise storage (SAS, Fibre Channel.)

As a point of reference, a PCI-X (133 MHz 64 bit) device and PCIe device at 4-lanes (x4), Gen1 speed have roughly the same peak transfer rate in a single-direction: 1064MB/sec. The PCIe bus has the potential to perform better than the PCI-X bus in cases where multiple devices are transferring data communicating simultaneously, or if communication with the PCIe peripheral is bidirectional.

Interconnect

PCIe devices communicate via a logical connection called an interconnect[4] or *link*. A link is a point-to-point communication channel between 2 PCIe ports, allowing both to send/receive ordinary PCI-requests (configuration read/write, I/O read/write, memory read/write) and interrupts (INTx, MSI, MSI-X). At the physical level, a link is composed of 1 or more *lanes*.[4] Low-speed peripherals (such as an 802.11 Wi-Fi card) use a single-lane (×1) link, while a graphics adapter typically uses a much wider (and thus, faster) 16-lane link.

Lane

A lane is composed of a transmit and receive pair of differential lines. Each lane is composed of 4 wires or signal paths, meaning conceptually, each lane is a full-duplex byte stream, transporting data packets in 8 bit 'byte' format, between endpoints of a link, in both directions simultaneously.[5] Physical PCIe slots may contain from one to thirty-two lanes, in powers of two (1, 2, 4, 8, 16 and 32).[4] Lane counts are written with an *x* prefix (e.g., *x16* represents a sixteen-lane card or slot), with x16 being the largest size in common use.[6]

Serial bus

The bonded serial format was chosen over a traditional parallel format due to the phenomenon of timing skew. Timing skew is a direct result of the limitations imposed by the speed of an electrical signal travelling down a wire, which it does at the finite speed of electricity. Because signal paths across an interface have different finite lengths, parallel signals transmitted simultaneously arrive at their destinations at slightly different times. When the interface clock rate increases to the point where the wavelength of a single bit is less than the smallest difference between path lengths, the bits of a single word do not arrive at their destination simultaneously, making parallel recovery of the word difficult. Thus, the speed of the electrical signal, combined with the difference in length between the longest and shortest path in a parallel interconnect, leads to a naturally imposed maximum bandwidth. Serial channel bonding avoids this issue by not requiring the bits to arrive simultaneously. PCIe is just one example of a general trend away from parallel buses to serial interconnects. Other examples include Serial ATA, USB, SAS, FireWire and RapidIO. Multichannel serial design increases flexibility by allowing slow devices to be allocated fewer lanes than fast devices.

Form factors

PCI Express (standard)

A PCIe card will fit into a slot of its physical size or larger (maximum x16), but may not fit into a smaller PCIe slot (x16 in a x8 slot). Some slots use open-ended sockets to permit physically longer cards and will negotiate the best available electrical connection. The number of lanes actually connected to a slot may also be less than the number supported by the physical slot size. An example is a ×8 slot that actually only runs at ×1; these slots will allow any ×1, ×2, ×4 or ×8 card to be used, though only running at ×1 speed. This type of socket is described as a *×8 (×1 mode)* slot, meaning it physically accepts up to ×8 cards but only runs at ×1 speed. The advantage gained is that a larger range of

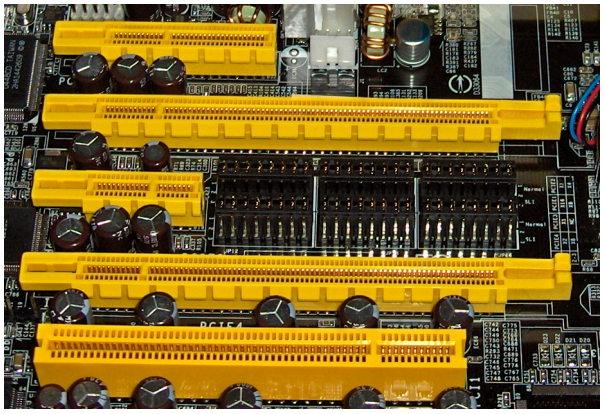

PCI Express slots (from top to bottom: ×4, ×16, ×1 and ×16), compared to a traditional 32-bit PCI slot (bottom), as seen on DFI's LanParty nF4 SLI-DR

PCIe cards can still be used without requiring the motherboard hardware to support the full transfer rate, which keeps design and implementation costs down.

Pinout

PCI express ×4 connector pinout

Pin	Side B	Side A	Comments
1	+12V	PRSNT1#	Pulled low to indicate card inserted
2	+12V	+12V	
3	Reserved	+12V	
4	Ground	Ground	
5	SMCLK	TCK	SMBus and JTAG port pins
6	SMDAT	TDI	
7	Ground	TDO	
8	+3.3V	TMS	
9	TRST#	+3.3V	
10	+3.3Vaux	+3.3V	Standby power
11	WAKE#	PWRGD	Link reactivation, power good.
Key notch			
12	Reserved	Ground	
13	Ground	REFCLK+	Reference clock differential pair
14	HSOp(0)	REFCLK-	Lane 0 transmit data, + and –
15	HSOn(0)	Ground	
16	Ground	HSIp(0)	Lane 0 receive data, + and –
17	PRSNT2#	HSIn(0)	
18	Ground	Ground	
19	HSOp(1)	Reserved	Lane 1 transmit data, + and –
20	HSOn(1)	Ground	
21	Ground	HSIp(1)	Lane 1 receive data, + and –
22	Ground	HSIn(1)	
23	HSOp(2)	Ground	Lane 2 transmit data, + and –
24	HSOn(2)	Ground	
25	Ground	HSIp(2)	Lane 2 receive data, + and –
26	Ground	HSIn(2)	
27	HSOp(3)	Ground	Lane 3 transmit data, + and –
28	HSOn(3)	Ground	
29	Ground	HSIp(3)	Lane 3 receive data, + and –
30	Reserved	HSIn(3)	

31	PRSNT2#	Ground	
32	Ground	Reserved	

An ×1 slot is a shorter version of this, ending after pin 18. ×8 and ×16 slots extend the pattern.

Legend

Ground pin	Zero volt reference
Power pin	Supplies power to the PCIe card
Output pin	Signal from the card to the motherboard
Input pin	Signal from the motherboard to the card
Open drain	May be pulled low and/or sensed by multiple cards
Sense pin	Tied together on card
Reserved	Not presently used, do not connect

Power

PCI Express cards are allowed a maximum power consumption of 25W (×1: 10W for power-up). Low profile cards are limited to 10W (×16 to 25W). PCI Express Graphics (PEG) cards may increase power (from slot) to 75W after configuration (3.3V/3A + 12V/5.5A).[7] Optional connectors add 75W (6-pin) or 150W (8-pin) power for up to 300W total.

PCI Express Mini Card

PCI Express Mini Card (also known as Mini PCI Express, Mini PCIe, and Mini PCI-E) is a replacement for the Mini PCI form factor based on PCI Express. It is developed by the PCI-SIG. The host device supports both PCI Express and USB 2.0 connectivity, and each card uses whichever the designer feels most appropriate to the task. Most laptop computers built after 2005 are based on PCI Express and can have several Mini Card slots.

A WLAN PCI Express Mini Card and its connector.

Physical dimensions

PCI Express Mini Cards are 30×50.95 mm. There is a 52 pin edge connector, consisting of two staggered rows on a 0.8 mm pitch. Each row has 8 contacts, a gap equivalent to 4 contacts, then a further 18 contacts. A half-length card is also specified 30×26.8 mm. Cards have a thickness of 1.0 mm (excluding components)..

Electrical interface

PCI Express Mini Card edge connector provide multiple connections and buses:

- PCIe x1
- USB 2.0
- SMBus
- Wires to diagnostics LEDs for wireless network (i.e., WiFi) status on computer's chassis

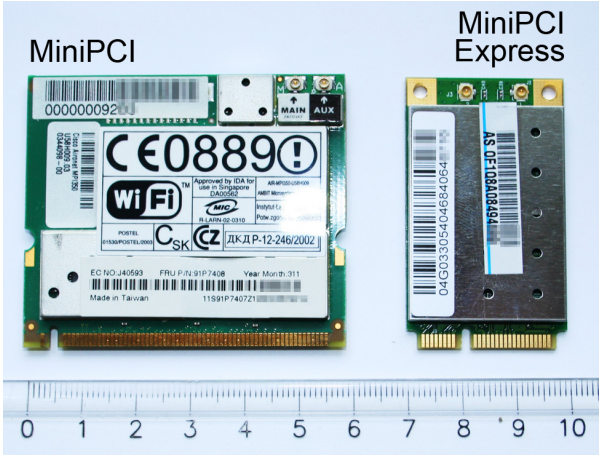

MiniPCI and MiniPCI Express cards in comparison

- SIM card for GSM and WCDMA applications. (UIM signals on spec)
- Future extension for another PCIe lane
- 1.5 and 3.3 volt power

Mini PCIe SSD & mSATA SSD

Some notebooks (notably the Asus Eee PC, the MacBook Air, and the Dell mini9 and mini10) use a variant of the PCI Express Mini Card as an SSD. This variant uses the reserved and several non-reserved pins to implement SATA and IDE interface passthrough, keeping only USB, ground lines, and sometimes the core PCIe 1x bus intact.[8] This makes the 'miniPCIe' flash and solid state drives sold for netbooks largely incompatible with true PCI Express Mini implementations.

Also, the typical Asus miniPCIe SSD is 71mm long, causing the Dell 51mm model to often be (incorrectly) referred to as half length. A true 51mm Mini PCIe SSD was announced in 2009, with two stacked PCB layers, which allows for higher storage capacity. The announced design preserves the PCIe interface, making it compatible with the standard mini PCIe slot. No working product has yet been developed, likely as a result of the popularity of the alternative variant.

To confuse matters even more, there is also an industry standard called mSATA, which uses the same principle of SATA passthrough on a PCI Express Mini connector, but which is electrically incompatible with normal PCIe and ASUS/DELL's connectors. So for example Intel's 310 series may sometimes be advertised as mini PCIe SSD cards, but needs an mSATA enabled motherboard.

PCI Express External Cabling

PCI Express External Cabling (also known as *External PCI Express*, *Cabled PCI Express*, or *ePCIe*) specifications were released by the PCI-SIG in February 2007.[9] [10]

Standard cables and connectors have been defined for x1, x4, x8, and x16 link widths, with a transfer rate of 250 MB/s per lane. The PCI-SIG also expects the norm will evolve to reach the 500 MB/s as found in PCI Express 2.0. The maximum cable length hasn't been determined yet. An example of the uses of Cabled PCI Express is a metal enclosure, containing a number of PCI slots and PCI-to-ePCIe adapter circuitry. This device would not be possible had it not been for the ePCIe spec.

Derivative forms

There are several other expansion card types derived from PCIe. These include:

- Low height card
- ExpressCard: successor to the PC card form factor (with ×1 PCIe and USB 2.0; hot-pluggable)
- PCI Express ExpressModule: a hot-pluggable modular form factor defined for servers and workstations
- XMC: similar to the CMC/PMC form factor (with ×4 PCIe or Serial RapidI/O)
- AdvancedTCA: a complement to CompactPCI for larger applications; supports serial based backplane topologies
- AMC: a complement to the AdvancedTCA specification; supports processor and I/O modules on ATCA boards (×1, ×2, ×4 or ×8 PCIe).
- FeaturePak: a tiny expansion card format (43 x 65 mm) for embedded and small form factor applications; it implements two x1 PCIe links on a high-density connector along with USB, I2C, and up to 100 points of I/O.
- Universal IO: A variant from Super Micro Computer Inc designed for use in low profile rack mounted chassis. It has the connector bracket reversed so it cannot fit in a normal PCI Express socket, but is pin compatible and may be inserted if the bracket is removed.
- Thunderbolt: A variant from Intel that combines DisplayPort and PCIe protocols in a form factor compatible with Mini DisplayPort.

History

While in early development, PCIe was initially referred to as *HSI* (for *High Speed Interconnect*), and underwent a name change to *3GIO* (for *3rd Generation I/O*) before finally settling on its PCI-SIG name *PCI Express*. It was first drawn up by a technical working group named the *Arapaho Work Group* (AWG) which, for initial drafts, consisted of an Intel only team of architects. Subsequently the AWG was expanded to include industry partners.

PCIe is a technology under constant development and improvement. The current PCI Express implementation is version 3.0.

PCI Express 1.0a

In 2003, PCI-SIG [11] introduced PCIe 1.0a, with a data rate of 250 MB/s and a transfer rate of 2.5 GT/s.

PCI Express 1.1

In 2005, PCI-SIG [11] introduced PCIe 1.1. This updated specification includes clarifications and several improvements, but is fully compatible with PCI Express 1.0a. No changes were made to the data rate.

PCI Express 2.0

PCI-SIG announced the availability of the PCI Express Base 2.0 specification on 15 January 2007.[12] The PCIe 2.0 standard doubles the per-lane throughput from the PCIe 1.0 standard's 250 MB/s to 500 MB/s. This means a 32-lane PCI connector (x32) can support throughput up to 16 GB/s aggregate. The PCIe 2.0 standard uses a base clock speed of 5 GHz, while the first version operates at 2.5 GHz.

PCIe 2.0 motherboard slots are fully backward compatible with PCIe v1.x cards. PCIe 2.0 cards are also generally backward compatible with PCIe 1.x motherboards, using the available bandwidth of PCI Express 1.1. Overall, graphic cards or motherboards designed for v2.0 will be able to work with the other being v1.1 or v1.0.

The PCI-SIG also said that PCIe 2.0 features improvements to the point-to-point data transfer protocol and its software architecture.[13]

Intel 's first PCIe 2.0 capable chipset was the X38 and boards began to ship from various vendors (Abit, Asus, Gigabyte) as of October 21, 2007.[14] AMD started supporting PCIe 2.0 with its AMD 700 chipset series and nVidia started with the MCP72.[15] All of Intel's prior chipsets, including the Intel P35 chipset, supported PCIe 1.1 or

1.0a.[16]

PCI Express 2.1

PCI Express 2.1 supports a large proportion of the management, support, and troubleshooting systems planned to be fully implemented in PCI Express 3.0. However, the speed is the same as PCI Express 2.0. Most motherboards sold currently come with PCI Express 2.1 connectors.

PCI Express 3.0

PCI Express 3.0 Base specification revision 3.0 was made available in November 2010, after multiple delays. In August 2007, PCI-SIG announced that PCI Express 3.0 would carry a bit rate of 8 gigatransfers per second, and that it would be backwards compatible with existing PCIe implementations. At that time, it was also announced that the final specification for PCI Express 3.0 would be delayed until 2011.[17] New features for the PCIe 3.0 specification include a number of optimizations for enhanced signaling and data integrity, including transmitter and receiver equalization, PLL improvements, clock data recovery, and channel enhancements for currently supported topologies.[18]

Following a six-month technical analysis of the feasibility of scaling the PCIe interconnect bandwidth, PCI-SIG's analysis found out that 8 gigatransfers per second can be manufactured in mainstream silicon process technology, and can be deployed with existing low-cost materials and infrastructure, while maintaining full compatibility (with negligible impact) to the PCIe protocol stack.

PCIe 2.0 delivers 5 GT/s, but employs an 8b/10b encoding scheme which results in a 20 percent ((10-8)/10) overhead on the raw bit rate. PCIe 3.0 removes the requirement for 8b/10b encoding and instead uses a technique called "scrambling" in which "a known binary polynomial is applied to a data stream in a feedback topology. Because the scrambling polynomial is known, the data can be recovered by running it through a feedback topology using the inverse polynomial"[19] and also uses a 128b/130b encoding scheme, reducing the overhead to approximately 1.5% ((130-128)/130), as opposed to the 20% overhead of 8b/10b encoding used by PCIe 2.0. PCIe 3.0's 8 GT/s bit rate effectively delivers double PCIe 2.0 bandwidth. According to an official press release by PCI-SIG on 8 August 2007:

> The final PCIe 3.0 specifications, including form factor specification updates, may be available by late 2009, and could be seen in products starting in 2010 and beyond.
> —[20]

As of January 2010, the release of the final specifications had been delayed until Q2 2010.[21] PCI-SIG expects the PCIe 3.0 specifications to undergo rigorous technical vetting and validation before being released to the industry. This process, which was followed in the development of prior generations of the PCIe Base and various form factor specifications, includes the corroboration of the final electrical parameters with data derived from test silicon and other simulations conducted by multiple members of the PCI-SIG.

On May 31, 2010, it was announced that the 3.0 specification would be coming in 2010, but not until the second half of the year.[22] Then, on June 23, 2010, the PCI Special Interest Group released a timetable showing the final 3.0 specification due in the fourth quarter of 2010.[23]

Finally, on November 18, 2010, the PCI Special Interest Group officially published the finalized PCI Express 3.0 specification to its members to build devices based on this new version of PCI Express.[24]

Current status

PCI Express has replaced AGP as the default interface for graphics cards on new systems. With a few exceptions, all graphics cards being released as of 2009 and 2010 from AMD (ATI) and NVIDIA use PCI Express. NVIDIA uses the high bandwidth data transfer of PCIe for its Scalable Link Interface (SLI) technology, which allows multiple graphics cards of the same chipset and model number to be run in tandem, allowing increased performance. ATI has also developed a multi-GPU system based on PCIe called CrossFire. AMD and NVIDIA have released motherboard chipsets which support up to four PCIe ×16 slots, allowing tri-GPU and quad-GPU card configurations.

PCI Express has displaced a major portion of the add-in card market. PCI Express was originally only common in disk array controllers, onboard gigabit Ethernet, Wi-Fi and graphics cards. Most sound cards, TV/capture-cards, modems, serial port/USB/Firewire cards, network/WiFi cards that would have used the conventional PCI in the past have moved to PCI Express x8, x4, or x1. While some motherboards have conventional PCI slots, these are primarily for legacy cards and are being phased out as operating systems no longer support the older hardware.

Hardware protocol summary

The PCIe link is built around dedicated unidirectional couples of serial (1-bit), point-to-point connections known as *lanes*. This is in sharp contrast to the earlier PCI connection, which is a bus-based system where all the devices share the same bidirectional, 32-bit or 64-bit parallel bus.

PCI Express is a layered protocol, consisting of a *transaction layer*, a *data link layer*, and a *physical layer*. The Data Link Layer is subdivided to include a media access control (MAC) sublayer. The Physical Layer is subdivided into logical and electrical sublayers. The Physical logical-sublayer contains a physical coding sublayer (PCS). The terms are borrowed from the IEEE 802 networking protocol model.

Physical layer

The PCIe Physical Layer (**PHY**, PCIEPHY, PCI Express PHY, or PCIe PHY) specification is divided into two sub-layers, corresponding to electrical and logical specifications. The logical sublayer is sometimes further divided into a MAC sublayer and a PCS, although this division is not formally part of the PCIe specification. A specification published by Intel, the PHY Interface for PCI Express (PIPE),[25] defines the MAC/PCS functional partitioning and the interface between these two sub-layers. The PIPE specification also identifies the *physical media attachment* (PMA) layer, which includes the serializer/deserializer (SerDes) and other analog circuitry; however, since SerDes implementations vary greatly among ASIC vendors, PIPE does not specify an interface between the PCS and PMA.

At the electrical level, each lane consists of two unidirectional LVDS or PCML pairs at 2.525 Gbit/s. Transmit and receive are separate differential pairs, for a total of 4 data wires per lane.

A connection between any two PCIe devices is known as a *link*, and is built up from a collection of 1 or more *lanes*. All devices must minimally support single-lane (x1) link. Devices may optionally support wider links composed of 2, 4, 8, 12, 16, or 32 lanes. This allows for very good compatibility in two ways:

- a PCIe card will physically fit (and work correctly) in any slot that is at least as large as it is (e.g., an ×1 sized card will work in any sized slot);
- a slot of a large physical size (e.g., ×16) can be wired electrically with fewer lanes (e.g., ×1, ×4, ×8, or ×12) as long as it provides the ground connections required by the larger physical slot size.

In both cases, PCIe will negotiate the highest mutually supported number of lanes. Many graphics cards, motherboards and bios versions are verified to support ×1, ×4, ×8 and ×16 connectivity on the same connection.

Even though the two would be signal-compatible, it is not usually possible to place a physically larger PCIe card (e.g., a ×16 sized card) into a smaller slot —though if the PCIe slots are open-ended, by design or by hack, some motherboards will allow this.

The width of a PCIe connector is 8.8 mm, while the height is 11.25 mm, and the length is variable. The fixed section of the connector is 11.65 mm in length and contains 2 rows of 11 (22 pins total), while the length of the other section is variable depending on the number of lanes. The pins are spaced at 1 mm intervals, and the thickness of the card going into the connector is 1.8 mm.[26] [27]

Lanes	Pins		Length	
	Total	Variable	Total	Variable
×1	2×18 = 36[28]	2×7 = 14	25 mm	7.65 mm
×4	2×32 = 64	2×21 = 42	39 mm	21.65 mm
×8	2×49 = 98	2×38 = 76	56 mm	38.65 mm
×16	2×82 = 164	2×71 = 142	89 mm	71.65 mm

Data transmission

PCIe sends all control messages, including interrupts, over the same links used for data. The serial protocol can never be blocked, so latency is still comparable to conventional PCI, which has dedicated interrupt lines.

Data transmitted on multiple-lane links is interleaved, meaning that each successive byte is sent down successive lanes. The PCIe specification refers to this interleaving as "data striping". While requiring significant hardware complexity to synchronize (or deskew) the incoming striped data, striping can significantly reduce the latency of the n^{th} byte on a link. Due to padding requirements, striping may not necessarily reduce the latency of small data packets on a link.

As with other high data rate serial transmission protocols, clocking information is embedded in the signal. At the physical level, PCI Express 2.0 utilizes the 8b/10b encoding scheme[19] to ensure that strings of consecutive ones or consecutive zeros are limited in length. This was used to prevent the receiver from losing track of where the bit edges are. In this coding scheme every 8 (uncoded) payload bits of data are replaced with 10 (encoded) bits of transmit data, causing a 20% overhead in the electrical bandwidth. To improve the available bandwidth, PCI Express version 3.0 employs 128b/130b encoding instead: similar but with much lower overhead.

Many other protocols (such as SONET) use a different form of encoding known as *scrambling* to embed clock information into data streams. The PCIe specification also defines a scrambling algorithm, but it is used to reduce electromagnetic interference (EMI) by preventing repeating data patterns in the transmitted data stream.

Data link layer

The Data Link Layer performs two vital services for the PCIe express link: (1) sequence the transaction layer packets (TLPs) that are generated by the transaction layer, (2) ensure reliable delivery of TLPs between two endpoints via an acknowledgement protocol (ACK and NAK signaling) that explicitly requires replay of unacknowledged/bad TLPs (3) initialize and manage flow control credits

On the transmit side, the data link layer generates an incrementing sequence number for each outgoing TLP. It serves as a unique identification tag for each transmitted TLP, and is inserted into the header of the outgoing TLP A 32-bit cyclic redundancy check code (known in this context as Link CRC or LCRC) is also appended to the end of each outgoing TLP.

On the receive side, the received TLP's LCRC and sequence number are both validated in the link layer. If either the LCRC check fails (indicating a data error), or the sequence-number is out of range (non-consecutive from the last valid received TLP), then the bad TLP, as well as any TLPs received after the bad TLP, are considered invalid and discarded. The receiver sends a negative acknowledgement message (NAK) with the sequence-number of the invalid TLP, requesting re-transmission of all TLPs forward of that sequence-number. If the received TLP passes the LCRC check and has the correct sequence number, it is treated as valid. The link receiver increments the sequence-number

(which tracks the last received good TLP), and forwards the valid TLP to the receiver's transaction layer. An ACK message is sent to remote transmitter, indicating the TLP was successfully received (and by extension, all TLPs with past sequence-numbers.)

If the transmitter receives a NAK message, or no acknowledgement (NAK or ACK) is received until a timeout period expires, then the transmitter must retransmit all TLPs for which no positive acknowledgement (ACK) was received. Barring a persistent malfunction of the device or transmission medium, the link-layer presents a reliable connection to the transaction layer, since the transmission protocol ensures delivery of TLPs over an unreliable medium.

In addition to sending and receiving TLPs generated by the transaction layer, the data-link layer also generates and consumes DLLPs, data link layer packets. ACK and NAK signals are communicated via (DLLP), as are flow control credit information, some power management messages and flow control credit information (on behalf of the transaction layer.)

In practice, the number of in-flight, unacknowledged TLPs on the link is limited by two factors: the size of the transmitter's replay buffer (which must store a copy of all transmitted TLPs until they the remote receiver ACKs them), and the flow control credits issued by the receiver to a transmitter. PCI Express requires all receivers to issue a minimum number of credits, to guarantee a link will allow PCIConfig TLPs and message TLPs to be sent.

Transaction layer

PCI Express implements split transactions (transactions with request and response separated by time), allowing the link to carry other traffic while the target device gathers data for the response.

PCI Express utilizes credit-based flow control. In this scheme, a device advertises an initial amount of credit for each of the receive buffers in its transaction layer. The device at the opposite end of the link, when sending transactions to this device, will count the number of credits consumed by each TLP from its account. The sending device may only transmit a TLP when doing so does not result in its consumed credit count exceeding its credit limit. When the receiving device finishes processing the TLP from its buffer, it signals a return of credits to the sending device, which then increases the credit limit by the restored amount. The credit counters are modular counters, and the comparison of consumed credits to credit limit requires modular arithmetic. The advantage of this scheme (compared to other methods such as wait states or handshake-based transfer protocols) is that the latency of credit return does not affect performance, provided that the credit limit is not encountered. This assumption is generally met if each device is designed with adequate buffer sizes.

PCIe 1.x is often quoted to support a data rate of 250 MB/s in each direction, per lane. This figure is a calculation from the physical signaling rate (2.5 Gbaud) divided by the encoding overhead (10 bits per byte.) This means a sixteen lane (x16) PCIe card would then be theoretically capable of 16×250 MB/s = 4 GB/s in each direction. While this is correct in terms of data bytes, more meaningful calculations will be based on the usable data payload rate, which depends on the profile of the traffic, which is a function of the high-level (software) application and intermediate protocol levels.

Like other high data rate serial interconnect systems, PCIe has a protocol and processing overhead due to the additional transfer robustness (CRC and acknowledgements). Long continuous unidirectional transfers (such as those typical in high-performance storage controllers) can approach >95% of PCIe's raw (lane) data rate. These transfers also benefit the most from increased number of lanes (×2, ×4, etc.) But in more typical applications (such as a USB or Ethernet controller), the traffic profile is characterized as short data packets with frequent enforced acknowledgements.[29] This type of traffic reduces the efficiency of the link, due to overhead from packet parsing and forced interrupts (either in the device's host interface or the PC's CPU.) Being a protocol for devices connected to the same printed circuit board, it does not require the same tolerance for transmission errors as a protocol for communication over longer distances, and thus, this loss of efficiency is not particular to PCIe.

Uses

External PCIe cards

Theoretically, external PCIe could give a notebook the graphics power of a desktop, by connecting a notebook with any PCIe desktop video card (enclosed in its own external housing, with strong power supply and cooling); This is possible with an ExpressCard interface, which provides single lane v1.1 performance.
[30] [31] [32] [33] [34]

IBM/Lenovo has also included a PCI-Express slot in their Advanced Docking Station 250310U. It provides a half sized slot with an x16 length slot, but only x1 connectivity.[35] However, docking stations with expansion slots are becoming less common as the laptops are getting more advanced video cards and either DVI-D interfaces, or DVI-D pass through for port replicators and docking stations.

Additionally, Nvidia has developed Quadro Plex external PCIe Video Cards that can be used for advanced graphic applications. These video cards require a PCI Express x8 or x16 slot for the interconnection cable.[36] AMD has recently announced the ATI XGP technology based on proprietary cabling solution which is compatible with PCIe x8 signal transmissions.[37] This connector is available on the Fujitsu Amilo and the Acer Ferrari One notebooks. Only Fujitsu has an actual external box available, which also works on the Ferrari One. Recently Acer launched the Dynavivid graphics dock for XGP. Shuttle introduced their own external graphics solutions, GXT.

There are now card hubs in development that one can connect to a laptop through an ExpressCard slot, though they are currently rare, obscure, or unavailable on the open market. These hubs can have full-sized cards placed in them.

Magma and ViDock also makes use of ExpressCard and implements the usage of External graphic solutions .ViDock are expansion chassis tailored specifically for adapting PCI Express graphics cards for use with ExpressCard equipped laptop PCs. This enables user to make use of connecting PCIe cards externally. Although, the developments in these technologies are still ongoing. Other examples that underwent are - MSI GUS, Asus XG Station.

Recently, Intel and Apple introduced Thunderbolt, which allows for external PCI(e) devices.

Competing protocols

Several communications standards have emerged based on high bandwidth serial architectures. These include InfiniBand, RapidIO, HyperTransport, QPI and StarFabric. The differences are based on the tradeoffs between flexibility and extensibility vs latency and overhead. An example of such a tradeoff is adding complex header information to a transmitted packet to allow for complex routing (PCI Express is not capable of this). The additional overhead reduces the effective bandwidth of the interface and complicates bus discovery and initialization software. Also making the system hot-pluggable requires that software track network topology changes. Examples of buses suited for this purpose are InfiniBand and StarFabric.

Another example is making the packets shorter to decrease latency (as is required if a bus is to be operated as a memory interface). Smaller packets mean that the packet headers consume a higher percentage of the packet, thus decreasing the effective bandwidth. Examples of bus protocols designed for this purpose are RapidIO and HyperTransport.

PCI Express falls somewhere in the middle, targeted by design as a system interconnect (local bus) rather than a device interconnect or routed network protocol. Additionally, its design goal of software transparency constrains the protocol and raises its latency somewhat.

Development tools

When developing and/or troubleshooting the PCI Express bus, examination of hardware signals can be very important to find the problems. Logic analyzers and bus analyzers are tools which collect, analyze, decode, store signals so people can view the high-speed waveforms at their leisure.

References

[1] "PCI Express 2.0 (Training)" (http://www.mindshare.com/learn/?section=132B11E3). MindShare. . Retrieved 2009-12-07.

[2] "PCI Express Base specification" (http://www.pcisig.com/specifications/pciexpress/base3/). PCI_SIG. . Retrieved 2010-10-18.

[3] "HowStuffWorks "How PCI Express Works"" (http://computer.howstuffworks.com/pci-express.htm). Computer.howstuffworks.com. . Retrieved 2009-12-07.

[4] "PCI Express Architecture Frequently Asked Questions" (http://www.pcisig.com/news_room/faqs/faq_express/). PCI-SIG. . Retrieved 23 November 2008.

[5] "PCI Express Bus" (http://www.interfacebus.com/Design_Connector_PCI_Express.html). . Retrieved 2010-06-12.

[6] "PCI Express – An Overview of the PCI Express Standard - Developer Zone - National Instruments" (http://zone.ni.com/devzone/cda/tut/p/id/3767). Zone.ni.com. 2009-08-13. . Retrieved 2009-12-07.

[7] PCI-SIG: Board Design Guidelines for PCI Express Architecture 2004 p. 19

[8] "Eee PC Research" (http://beta.ivancover.com/wiki/index.php/Eee_PC_Research). . Retrieved 26 October 2009.

[9] "PCI Express External Cabling 1.0 Specification" (http://www.pcisig.com/specifications/pciexpress/pcie_cabling1.0/). . Retrieved 9 February 2007.

[10] "February 7, 2007" (http://www.pcisig.com/news_room/news/press_release/02_07_07). Pci-Sig. 2007-02-07. . Retrieved 2010-09-11.

[11] http://www.pcisig.com

[12] PCI-SIG (15 January 2007). "PCI Express Base 2.0 specification announced" (http://www.pcisig.com/news_room/PCIe2_0_Spec_Release_FINAL2.pdf) (PDF). Press release. . Retrieved 9 February 2007. — note that in this press release the term *aggregate bandwidth* refers to the sum of incoming and outgoing bandwidth; using this terminology the aggregate bandwidth of full duplex 100BASE-TX is 200 Mbit/s

[13] Tony Smith (11 October 2006). "PCI Express 2.0 final draft spec published" (http://www.reghardware.co.uk/2006/10/11/pic-sig_posts_pcie_2_final_draft/). The Register. . Retrieved 9 February 2007.

[14] Gary Key & Wesley Fink (21 May 2007). "Intel P35: Intel's Mainstream Chipset Grows Up" (http://www.anandtech.com/cpuchipsets/showdoc.aspx?i=2993). AnandTech. . Retrieved 21 May 2007.

[15] Anh Huynh (8 February 2007). "NVIDIA "MCP72" Details Unveiled" (http://www.dailytech.com/article.aspx?newsid=6021). AnandTech. . Retrieved 9 February 2007.

[16] "Intel P35 Express Chipset Product Brief" (http://download.intel.com/products/chipsets/P35/317304.pdf) (PDF). Intel. . Retrieved 5 September 2007.

[17] Hachman, Mark. "PC Magazine" (http://www.pcmag.com/article2/0,2817,2351266,00.asp). Pcmag.com. . Retrieved 2010-09-11.

[18] "PCI Express 3.0 Bandwidth: 8.0 Gigatransfers/s" (http://www.extremetech.com/article2/0,1697,2169018,00.asp). ExtremeTech. 9 August 2007. . Retrieved 5 September 2007.

[19] "PCI Express 3.0 Frequently Asked Questions" (http://www.pcisig.com/news_room/faqs/pcie3.0_faq/). PCI-SIG. . Retrieved 23 November 2010.

[20] "PCI-SIG Announces PCI Express 3.0 Bit Rate For Products In 2011 And Beyond" (http://www.pcisig.com/news_room/08_08_07/). 8 August 2007. . Retrieved 24 March 2008.

[21] "PCI Express 3.0 Spec Pushed Out to 2010" (http://www.pcmag.com/article2/0,2817,2351266,00.asp). 5 August 2009. . Retrieved 30 August 2009.

[22] "PCI-Express 3.0 specs to finally be released in 2H" (http://www.tcmagazine.com/tcm/news/misc/28387/pci-express-30-specs-finally-be-released-2h). 31 May 2010. . Retrieved 7 June 2010.

[23] "PCI Express 3.0 Nears Completion" (http://www.extremetech.com/article2/0,2845,2365574,00.asp). 23 Jun 2010. . Retrieved 16 August 2010. "The PCI Special Interest Group released a detailed timetable for the next-generation PCI Express 3.0 specification on Wednesday, with a final specification due by the fourth quarter and products possibly in early 2011."

[24] "PCI Special Interest Group Publishes PCI Express 3.0 Standard." (http://www.xbitlabs.com/news/other/display/20101118151837_PCI_Special_Interest_Group_Publishes_PCI_Express_3_0_Standard.html). 18 November 2010. . Retrieved 18 November 2010.

[25] "PHY Interface for the PCI Express Architecture, version 2.00" (http://download.intel.com/technology/pciexpress/devnet/docs/pipe2_00.pdf) (PDF). . Retrieved 21 May 2008.

[26] "Mechanical Drawing for PCI Express Connector" (http://www.interfacebus.com/Design_Connector_PCI_Express.html#d). . Retrieved 7 December 2007.

[27] "FCi schematic for PCIe connectors" (http://portal.fciconnect.com/Comergent/en/US/fci/drawing/10018783.pdf). . Retrieved 7 December 2007.

[28] "PCI Express 1x, 4x, 8x, 16x bus pinout and wiring @" (http://pinouts.ru/Slots/pci_express_pinout.shtml). Pinouts.ru. . Retrieved 2009-12-07.
[29] "Computer Peripherals And Interfaces" (http://www.google.com/books?id=Xp7-NKsJ8_sC&pg=PA35&dq=frequent+enforced+acknowledgements/). Technical Publications Pune. . Retrieved 23 July 2009.
[30] "Magma ExpressBox1: Cabled PCI Express for Desktops and Laptops" (http://www.magma.com/products/pciexpress/expressbox1/index.html). Magma.com. . Retrieved 2010-09-11.
[31] "TheInquirer" (http://www.theinquirer.net/default.aspx?article=41876). TheInquirer. . Retrieved 2010-09-11.
[32] "Custompcmag.co.uk" (http://www.custompcmag.co.uk/news/601237/msi-demos-external-PCIe-graphics-box.html). Custompcmag.co.uk. . Retrieved 2010-09-11.
[33] ASUSTeK Computer (http://www.asus.com/news_show.aspx?id=5369)
[34] "Technology Beats. News and Reviews" (http://www.vr-zone.com/index.php?i=4728). VR-Zone. 1995-09-09. . Retrieved 2010-09-11.
[35] "IBM/Lenovo Thinkpad Advanced Dock Overview" (http://www-307.ibm.com/pc/support/site.wss/document.do?sitestyle=lenovo&lndocid=MIGR-61232). 307.ibm.com. 2010-03-07. . Retrieved 2010-09-11.
[36] "NVIDIA Quadro Plex VCS – Advanced visualization and remote graphics" (http://www.nvidia.com/page/quadroplex.html). Nvidia.com. . Retrieved 2010-09-11.
[37] "Advanced Micro Devices, AMD – Global Provider of Innovative Microprocessor, Graphics and Media Solutions" (http://ati.amd.com/technology/xgp/index.html). Ati.amd.com. . Retrieved 2010-09-11.

Further reading

- *PCI Express System Architecture*; 1st Ed; Ravi Budruk / Don Anderson / Tom Shanley; 1120 pages; 2003; ISBN 9780321156303.
- *Introduction to PCI Express : A Hardware and Software Developer's Guide*; 1st Ed; 325 pages; 2003; ISBN 9780970284693.
- *Complete PCI Express Reference : Design Implications for Hardware and Software Developers*; 1st Ed; 1056 pages; 2003; ISBN 9780971786196.

External links

- Introduction to PCI Protocol (http://electrofriends.com/articles/computer-science/protocol/introduction-to-pci-protocol/)
- *PCI Express Base Specification Revision 1.0* (http://www.pcisig.com/specifications/pciexpress/base). PCI-SIG. 29 April 2002. (Requires PCI-SIG membership)
- PCI-SIG, the industry organization that maintains and develops the various PCI standards (http://www.pcisig.com/)
- An introduction to how PCIe works at the TLP level (http://billauer.co.il/blog/2011/03/pci-express-tlp-pcie-primer-tutorial-guide-1/)
- Intel Developer Network for PCI Express Architecture (http://www.intel.com/technology/pciexpress/devnet/)
- PCI-E Graphics Cards and Specs (http://www.gpubench.com/)
- pcie specs @TPB (http://thepiratebay.org/torrent/6094173/PCI.Express.Specifications)

Article Sources and Contributors

PCI Express *Source*: http://en.wikipedia.org/w/index.php?oldid=444329609 *Contributors*: 10nitro, 1ForTheMoney, 5lithy, Aaron mcd, Acerperi, Ad88110, Adoniscik, Aethandor, Agoode, Ahoerstemeier, Aidoor, AirDjurdan, Aknorals, Aldie, Alereon, Aleshik, Alison, AlistairMcMillan, Alobodig, Aluvus, Alvin-cs, Amfar018, Amor001, AndreKR, Andrew sh, Andromeda451, Andy Dingley, Arch dude, Arkrishna, Armando, Arnero, ArnoldReinhold, Ascidian, Asestar, Aspects, Atracht, Audriusa, Avb, Avoided, Azuris, Balabiot, Bart.vanassche, Ben Ben, Bender2k14, Berkut, Bidgee, Bilepie66, Billauer, Blakegripling ph, Blazingluke, BlueDevil, Bmicomp, Bobblewik, Bogey97, Borgx, Bovineone, Brian Patrie, Brianski, Bthetford, Btilm, BubblesFromBelow, Bytre, Captain-tucker, Cbreaker, Ceros, CesarB, Chovain, Chowbok, Ciphergoth, Ckmac97, ClockworkSoul, Closedmouth, Cncxbox, Coffee4binky, Cometstyles, CommonsDelinker, Cootiequits, Cristan, Cryogenius, Cwolfsheep, Cybercobra, Cybermaus0, Cynix, DARTH SIDIOUS 2, DMahalko, DOSGuy, DaProx, Darin-0, Darkhorse, Darxus, Davewho2, David Latapie, DavidDouthitt, DeathHamster, Defunctzombie, Deineka, Dennis Schmitz, Denniss, Dgies, DhurdPLX, Diablo-D3, Discospinster, DmitryKo, Don't give an Ameriflag, DoriSmith, Dpark, Drjt87, Drrll, Dsalt, Dsf7183, Dsm, Dwvisser, E rulez, Eatcacti, Ebikeofacookbook, Editore99, Efeinberg, Electron9, EoGuy, Epbr123, Epiteo, Eptin, ErisDiscord, Etchesketch, Evice, Evil Monkey, ExpatEgghead, FIGATT, Falcon8765, Fdrake, Ferret141, Fethers, Firefox, Foobaz, FracJackMac, Francs2000, Frankenpuppy, Frap, Frediloc, Fredrik, Freewol, Fsiler, Furrykef, Gdo01, Geek94, Geekosaurus, Gerixau, GermanX, Gimmetrow, Glome83, Goodone121, GraemeL, Graham1428, GrandDrake, Gt4kill3r, Gurch, Halframed, Haris Akkool, HarisM, Hcberkowitz, Hellbus, Heron, Hideyuki, Hopp, Hotdogger125, Hpcanswers, Imroy, J.delanoy, Jack74661, Jacob Poon, Jcolbyk, Jdwinx, Jebinece, Jeffq, Jel, Jesse Viviano, Jflins, Jgui, Jk050902, Jlundell, Joebishpie, JogyB, JohnnySolc, Johnnydc, Jonverve, Joshuakh, Jsmethers, Juux, K20, Karam.Anthony.K, Kate, KatherinePankratz, Kbh3rd, Kbolino, Keelec, KelleyCook, Kendelarosa5357, Kenyon, Kestral fire, Kiko3d, Kinema, Kingvashy, Ksiddique, KuSh, La Parka Your Car, Leandrod, Lee Pavelich, Leo soundhar, LilHelpa, Ling.Nut, Lkesteloot, Lotje, Lpolyak, Luna Santin, M7, MER-C, MSTCrow, Madmaxx, Mahumphrey, Materialscientist, Matneh, Mattgirling, Mattventura, Maury Markowitz, MaxVT, Maxis ftw, Mdoc7, Meco, Metageek, Mgdunn, Mgolden, MightyWarrior, Mikepelley, Mikiemike, Mild Bill Hiccup, Mjaxon, Modster, Moreati, MorganHaggis, Moxfyre, Mr. Mentor, Mr. Wheely Guy, Mrand, Mulad, Music Sorter, NAHID, Nageh, NeOak, Neckername, Neelix, Nextop, Nicknomo, Nikpapag, Nisselua, Nixdorf, Nosforit, Nz17, OSP Editor, Omicronpersei8, Oneiros, Orphan Wiki, OwenX, PCI-SIG Administration, Pabouk, Pale blue dot, Paranoidmage, Pauli133, Pedantic of Purley, Peyre, Phatom87, Philip Trueman, Placi1982, Plasticup, Plonk420, Plugwash, Poompt, Psrgchx, Qdotdot, Qwe, R!SC, RBBrittain, Rayisthechosenone, Rbo159, Reedy, RevRagnarok, RexNL, Riana, Rich Farmbrough, Riluve, Rjwilmsi, Rl, Rlcantwell, Rmallins, Robertwharvey, Rockstone35, Ronark, RoyBoy, Rsduhamel, Rsmoore, Ruggedsys, RuineR, SNEST2, Sahrin, Salam32, Samrawlins, Sarahstirland, Sarenne, SatsukiMikata, Sbmeirow, Schodge, Scohoust, Scott 110, Sct72, Selket, Shaddack, Shafran, Shandris, Shawnc, Shnout, Silveroblivion, Sleigh, Sloverlord, Smathw, Smoove Z, Snickerdo, Spartan, SpeedyGonsales, Spinningspark, Spyd3r0us, Ssri1983, Stasdm, SteveSims, Strandist, Strom, StuffOfInterest, Supaplex, Sweetness46, SystemBuilder, Szlevi, Tabledhote, Tansm, Tarquin, The Belgain, The Thing That Should Not Be, The Wild Falcon, The1physicist, TheDJ, TheJosh, TheNightFly, Thingg, Thinkingatoms, ThreeBlindMice, Thue, Thunderbird2, Tide rolls, Timc, Timharwoodx, Todd Vierling, Tom.Reding, Towel401, Tpbradbury, Turionaltec, Tverbeek, Undeference, Vmaldia, Voidxor, Vsood 007, Wadsworth, Warpozio, Warren, Wasted Sapience, WegianWarrior, Welguisz, Whiner01, WhosAsking, Widefox, Winterst, Wkafig, Wolrahnaes, Wuhwuzdat, Xaje, Zac67, ZeroOne, Милан Јелисавчић, 941 anonymous edits

Image Sources, Licenses and Contributors

file:PCI Express logo.svg *Source*: http://en.wikipedia.org/w/index.php?title=File:PCI_Express_logo.svg *License*: Fair Use *Contributors*: Cristan, Jesse Viviano, Phatom87, Sfan00 IMG, TheDJ

Image:PCIExpress.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:PCIExpress.jpg *License*: GNU Free Documentation License *Contributors*: w:user:snickerdo

Image:A PCI Express Mini Card and its connector.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:A_PCI_Express_Mini_Card_and_its_connector.jpg *License*: Creative Commons Attribution 2.0 *Contributors*: Bastiaan van den Berg

Image:MiniPCI and MiniPCI Express cards.jpg *Source*: http://en.wikipedia.org/w/index.php?title=File:MiniPCI_and_MiniPCI_Express_cards.jpg *License*: Public Domain *Contributors*: Cvdr

License

Creative Commons Attribution-Share Alike 3.0 Unported
http://creativecommons.org/licenses/by-sa/3.0/