

Technology Note


PCI Express

The PCI Express logo is centered within a large, stylized circular graphic. The graphic consists of several overlapping, semi-transparent blue and white rings that create a sense of depth and motion. The text "PCI EXPRESS" is written in a bold, black, sans-serif font. The word "PCI" is on the top line, and "EXPRESS" is on the bottom line. A blue, stylized arrowhead or "wing" symbol is positioned between "PCI" and "EXPRESS". A small "TM" trademark symbol is located to the right of "EXPRESS".

PCI EXPRESS™

www.euresys.com
info@euresys.com

© Copyright 2006 Euresys s.a. Belgium. Euresys® is registered trademarks of Euresys s.a. Belgium.
Other product and company names listed are trademarks or trade names of their respective manufacturers.
Euresys reserves the right to modify product specifications and price without previous notice.

The Euresys logo features a stylized, multi-colored 'E' shape composed of several rectangular blocks in shades of blue, green, and red. Below this graphic, the word "EURESYS" is written in a bold, black, sans-serif font. Underneath "EURESYS", the tagline "Excellence in vision" is written in a smaller, italicized, black, sans-serif font.

EURESYS™
Excellence in vision

PCI Express

Contents

Introduction	3
PCI Bus Variation	3
PCI Background	3
Structure	4
Lanes	4
Architecture	5
Layer Stack	6
Cost	6
Mechanical	6
Power	6
Conventional PCI Comparison	7
Performance	8
Number of Lanes	8
Clock Speed	9
Interoperability	9
General Rule	9
Remark	10
References	10
PCI and PCI Express	10
Euresys	10

WARNING

EURESYS S.A. shall retain all property rights, title and interest in the documentation and trademarks of EURESYS S.A.

The licensing, use, leasing, loaning, translation, reproduction, copying or modification of the marks or documentation of EURESYS S.A. contained in this book, is not allowed without prior notice.

EURESYS S.A. may modify or change the information given in this documentation at any time, in its discretion, and without prior notice. EURESYS S.A. shall not be liable for any loss of or damage to revenues, profits, goodwill, data, information systems or other special, incidental, indirect, consequential or punitive damages of any kind resulting of omissions or errors in this book.

www.euresys.com

© Copyright 2006 Euresys s.a. Belgium. Euresys® is registered trademark of Euresys s.a. Belgium.
Other product and company names listed are trademarks or trade names of their respective manufacturers.

Introduction

Today, PCI Express is mainstream technology for connecting a number of devices in a PC. Most motherboards have PCI Express architecture and additional slots, in consumer as well as in high-end PCs.

The Conventional PCI bus has served the industry well for the past 10-12 years and it will continue to play an important role in the coming years.

Following successful Conventional PCI tracks, the new PCI Express standard was developed. PCI Express performs as a general purpose high data rate I/O interconnect in multiple market segments and applications.

The key Conventional PCI attributes, such as its usage model and software interfaces, are maintained. At the same time, its bandwidth-limiting, parallel bus implementation is replaced by a long-life, fully-serial interface.

PCI Express provides industry leading performance and cost-effectiveness for the coming decade.

PCI Bus Variation

Euresys has published an application note dedicated to Conventional PCI and PCI-X: *PCI Bus Variation* (available from Euresys website). In addition to compatibility issues, this note exposes important performance figures.

Please refer to this application note for comparison material with PCI Express performance and characteristics.

PCI Background

PCI Express inherits from ten years of Conventional PCI usage. The current status of PCI busses and the important milestones of the Conventional PCI and PCI-X history are summarized hereafter.

Conventional PCI is characterized by a bus width of 32 or 64 bits, and a clock speed of 33 or 66 MHz. Any combination of width and speed can be found. The majority of existing desktop PCs are, at least partly, equipped with Conventional PCI technology.

The original PCI revision 1.0 local bus specification was introduced in the early 1990s. The industrial inception of this technology started with PCI revision 2.0 in April 1993, which included expansion slots and 66 MHz clock rate support.

In the following years, three revisions (2.1 in 1995, 2.2 in 1998 and 2.3 in 2002) introduced significant improvements. The final specification -revision 3.0 published in February 2004- definitely precludes the use of 5 Volt add-in cards.

PCI-X is a high performance evolution of Conventional PCI. It uses the same hardware structure, which makes it possible to operate a PCI-X add-in card in a Conventional PCI slot, and vice-versa. The PCI-X technology is dedicated to server applications, and is therefore found on many high-end PC motherboards.

PCI-X allows for a considerably higher bandwidth than Conventional PCI. It is characterized by a higher clock speed (up to 133 MHz), and is usually found in the 64-bit bus width, although the 32-bit width can exist.


Revision 1.0, initially issued in July 2000, specifies PCI-X as an addendum to Conventional PCI. Revision 2.0, an improvement to PCI-X specification, published in November 2002, extends the clocking capability to 266 and 533 MHz.



Structure

Lanes

The PCI Express physical layer interfaces fast and robust serial point-to-point bus. A *lane* -the basic PCI Express communication channel- combines two serial busses, one transmitting data and one receiving data. From the application side, a PCI Express lane behaves as full-duplex channel.



The delivery bandwidth from and to a device is increased by using multiple lanes in parallel. The PCI Express architecture permits the parallelization of 1, 2, 4, 8, 12, 16 and 32 lanes. This parallelization only involves the hardware layer and is application transparent.

PCI Express lanes are clocked at 2.5 GHz, providing theoretical communication channels at 200 MByte/s in each direction. This bandwidth is close to twice the base Conventional PCI data rate.

In the draft PCI Express 2.0 specification, the clock will be increased to 5 GHz -and expectedly more in a future specification- increasing significantly the available bandwidth while keeping backward compatibility with lower speed components.

Technology Note. PCI Express


Architecture

In PCI Express based computers, the CPU and the memory exchange data through a new device, called the *root complex*. This root complex exposes multiple PCI Express lanes.

Usually, sixteen lanes are dedicated to the graphical interface. The remaining lanes are connected to PCI Express devices, or *end points*, and to PCI Express connectors. According to the required data rate of the end points and connectors, one or more lanes will be wired in parallel.

The last component introduced in PCI Express architecture is the *switch* that easily adds connection capabilities to a system, as shown on the following figure.

Implementing PCI Express to PCI bridges offers compatibility with Conventional PCI devices.




Technology Note, PCI Express

Layer Stack

The PCI Express structure is specified in layers. Compatibility with the PCI addressing model is maintained to ensure that all existing applications and drivers operate unchanged.

PCI Express configuration uses standard mechanisms as defined in the PCI specification. The software layers generate read and write requests that are transported by the transaction layer to the I/O devices using a packet-based, split-transaction protocol. The data link layer adds sequence numbers and CRC error correction to these packets to insure the data integrity. The basic physical layer implements a transmit pair and a receive pair.



PCI Express software layers are backward compatible with Conventional PCI. Applications moving from Conventional PCI to PCI Express architecture do not require software upgrading.

Future operating systems will take advantage of new PCI Express capabilities such as increased memory space management, virtual channels, and packet prioritization.


Cost

In many applications, the simplified connection scheme between components will reduce the overall cost of PCI Express systems compared to Conventional PCI.

Mechanical

The PCI Express form factor complies with Conventional PCI. Boards, equivalent in size and using the same brackets, integrate straight in current PC enclosures. Only the PCI Express connectors are modified and often reduced.

Four edge connectors are specified for PCI Express add-on boards; the x1 or 1-lane smaller connector, the x4, the x8 and the x16 connectors. The identical four connector sizes are specified for motherboards slots.



Power

For simple integration, the PCI Express connector provides a +12 V and a +3.3 V voltage to power the add-on boards.

The effective power consumption of a PCI Express board should be marked on the board, according to a standard addendum currently in review.




Technology Note, PCI Express

Conventional PCI Comparison

PCI Express implements serial point-to-point communication channels where Conventional PCI uses parallel multipoint busses. This fundamental move in the physical layer leads to a renewed architecture of computers.

For reminder, in Conventional PCI based computers, the CPU, the memory and the graphical interface exchange data through the host bridge. One or several PCI busses connect the host bridge with additional lower data rate devices, including expansion PCI connectors.

Implementing PCI to PCI bridges adds connection capabilities to a system.




Performance

Number of Lanes

The number of lanes is, for now, the main parameter influencing the performance of a PCI Express interface.

Parallelizing lanes increases the available data rate accordingly. The physical PCI Express interface automatically splits the data across the available lanes, as shown in the following figure.

No firmware is required to split the data packets.



The theoretical usable data rate of one PCI Express lane reaches about 200 MByte/s in each direction (for a total of 2 x 200 MBytes/s).

Theoretical Data Rate at 2.5 GHz

Number of Lanes	Bus Data Rate	
	MByte/s	GByte/s
1	2 x 200	0.4
2	2 x 400	0.8
4	2 x 800	1.6
8	2 x 1600	3.2
12	2 x 2400	4.8
16	2 x 3200	6.4
32	2 x 6400	12.8



Technology Note, PCI Express

Euresys development team reports the available data rate as slightly higher than 180 MByte/s in both directions. These figures have been measured on marketed standard PCs.

As of today, in typical conditions, a data rate of 2 x 180 MByte/s is a realistic performance for a 1-lane PCI Express board. An equivalent performance of 2 x 720 MByte/s is expected for a 4-lane PCI Express board.

Available Data Rate at 2.5 GHz

Number of Lanes	Bus Data Rate	
	MByte/s	GByte/s
1	2 x 180	0.36
4	2 x 720	1.44

Clock Speed

The PCI Express revision 1.1 specification, released in March 2005, specifies a 2.5 GHz clock. Expectedly future revision will specify higher clock rates. At that time, the interface performance will directly depends on the clock speed.

Interoperability


General Rule

A PCI Express board may be plugged in any PCI Express connector with as many lanes or more. As a result, a 1-lane board is universal as it is compatible with all connectors.

A PCI Express board cannot be plugged in a connector with fewer lanes. As a safeguard, the connector mechanical design prevents any forbidden configuration.

The PCI Express Card Electromechanical Specification revision 1.1, released in March 2005, specifies the x1, x4, x8 and x16 connectors.

The following figure shows examples of configurations.



Technology Note. PCI Express

Remark

Please note that according to the PCI Express specification, boards and connectors are only required to work in the basic x1 and in their nominal configurations. They are not required to work in intermediate configurations.

As shown in the following table, the effective number of lanes used for data transfer is unknown in advance in three combinations of board and slot sizes, and must be negotiated by both devices. In these cases of up-plugging -plugging a smaller board in a larger connector- the guaranteed effective number of lanes is one, as per specification. A successful negotiation will increase the effective number of lanes up to the intermediate state.

Effective Number of Lanes used for Data Transfer		Motherboard Slot Size			
		x1	x4	x8	x16
Add-on Board Size	x1	1	1	1	1
	x4	-	4	Negotiated (1 or 4)	Negotiated (1 or 4)
	x8	-	-	8	Negotiated (1, 4 or 8)
	x16	-	-	-	16

For example, a 4-lane board is inserted in an 8-lane connector. If the slot is only compatible with x1 and x8 configurations, the communication channel will establish on a single lane, with a data rate four times lower than expected. If the slot is 4-lane compatible, the communication channel will use these four lanes up to the expected data rate of 2 x 800 MByte/s. An addendum to the standard is in review and will specify a marking of both boards and connectors in order to show to which configurations a device is compliant with and help the user to determine the available data rate of its system.

References

PCI and PCI Express

The PCI bus specification is handled by the PCI-SIG (PCI Special Interest Group), which is an association of numerous industrial actors in the field.

PCI Express Base Specification
Revision 1.1
March 28, 2005

PCI Express Card Electromechanical Specification
Revision 1.1
March 28, 2005

The web site address of PCI-SIG is www.pcisig.com

Euresys

Euresys is a member of the PCI-SIG (vendor ID h1805).

America, Euresys Inc.
500 Park Boulevard, suite 525, Itasca, Illinois 60143
Toll free: 1-866-EURESYS - Phone: 630-250-2300 - Fax: 630-250-2301

Asia, Euresys Pte. Ltd.
627A Aljunied Road, #08-09 BizTech Centre, Singapore 389842
Phone: +65 6748 0085 - Fax: +65 6841 2137

Japan, Euresys s.a. Japan Representative Office
AIOS Hiroo Building 8F, Hiroo 1-11-2, Shibuya-ku, Tokyo 150-0012
Phone: +81 3 5447-1256 - Fax: +81 3 5447-0529

Europe, Euresys s.a. Corporate Headquarters
14, Avenue du Pré-Ailly, B-4031 Angleur, Belgium
Phone: +32 4 367 72 88 - Fax: +32 4 367 74 66

